

# Microarchitecture Vulnerabilities

Past, Present and Future

Daniel Gruss (Graz University of Technology)  
Anders Fogh (Intel Corporation)

# Introduction

**Daniel Gruss**

Graz University of Technology

**Anders Fogh**

Intel

Daniel and Anders  
do not always agree!!



A stylized, monochromatic illustration of a vintage computer workstation. The central focus is a CRT monitor displaying the word "Past" in a bold, black, sans-serif font. The monitor sits on a large, rectangular base that houses a floppy disk drive. In front of the base is a keyboard and a mouse. To the left and right of the central workstation are smaller, similar computer setups, including monitors and keyboards. The background is filled with a dense, repeating pattern of various vintage computer components, such as circuit boards, monitors, and keyboards, creating a sense of a vast digital archive or a collection of old technology. The overall aesthetic is clean and minimalist, with a focus on the shapes and textures of the hardware.

**Past**

# Past – earliest days

Side Channels always existed

# Past – earliest days

Side Channels always existed

First scientific observations in 1943

~~SECRET~~

(b) (3) - P.L. 86-36



Approved for Release by NSA on  
09-27-2007, FOIA Case # 51633

## TEMPEST: A Signal Problem

The story of the discovery  
of various compromising radiations  
from communications and Comsec equipment.

impractical. Hydraulic techniques—to replace the electrical—were tried and abandoned, and experiments were made with different types of batteries and motor generators, in attempts to lick the power-line problem. None was very successful.

During this period, the business of discovering new TEMPEST threats, or refining techniques and instrumentation for detecting, recording, and analyzing these signals, progressed more swiftly than the art of suppressing them. Perhaps the attack is more exciting than the defense—something more glamorous about finding a way to read one of these signals than going through the drudgery necessary to suppress that whacking great spike first seen in 1943. At any rate, when they turned over the next rock, they found the acoustic problem under it. Phenomenon No. 5.

### *Acoustics*

We found that most acoustic emanations are difficult to exploit if the microphonic device is outside of the room containing the source equipment; even a piece of paper inserted between, say, an offending keyboard and a pick-up

~~SECRET~~

UNCLASSIFIED//FOR OFFICIAL USE ONLY

## Past – earliest days

Side Channels always existed

First scientific observations in 1943

Concept of “covert channels” in 1973

Operating  
Systems

C. Weissman  
Editor

---

# A Note on the Confinement Problem

Butler W. Lampson  
Xerox Palo Alto Research Center

**This note explores the problem of confining a program during its execution so that it cannot transmit information to any other program except its caller. A set of examples attempts to stake out the boundaries of the problem. Necessary conditions for a solution are stated and informally justified.**

Communications  
of  
the ACM

October 1973  
Volume 16  
Number 10

# Past – earliest days

Side Channels always existed

First scientific observations in 1943

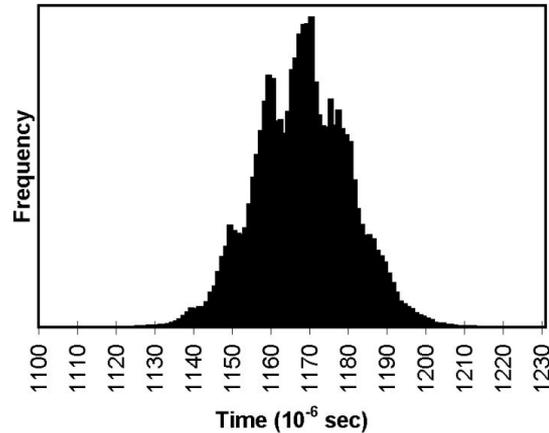
Concept of “covert channels” in 1973

1974-1980: Provable secure operating systems with exceptions for side channels

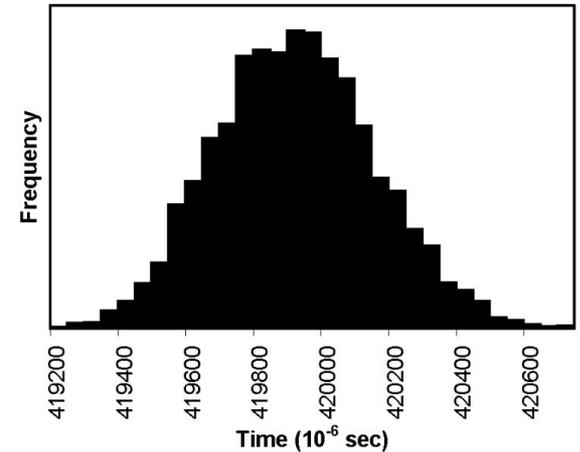
1985: Orange book. Covert channels with low bandwidth not a problem

1996: Paul Kocher’s seminal work on timing attacks

**FIGURE 1:** RSAREF Modular Multiplication Times

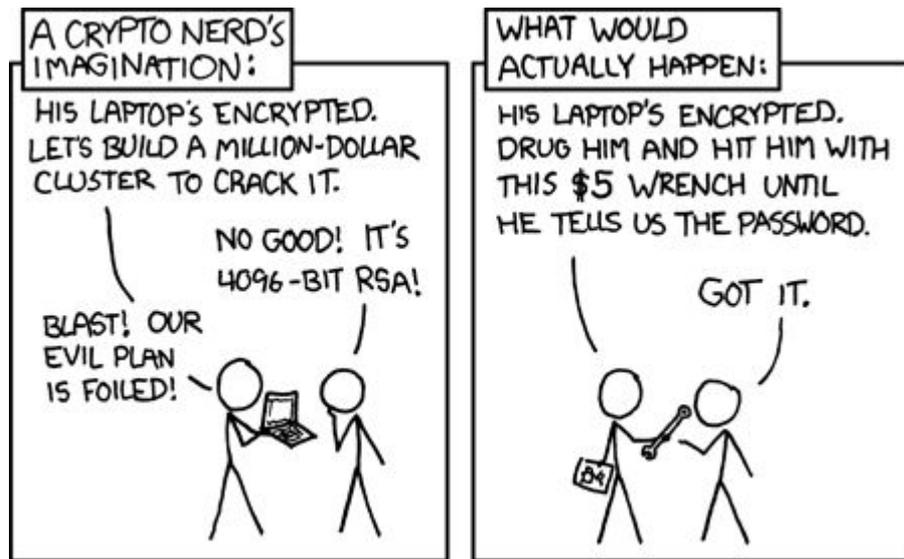


**FIGURE 2:** RSAREF Modular Exponentiation Times



# Past: cryptographic attacks

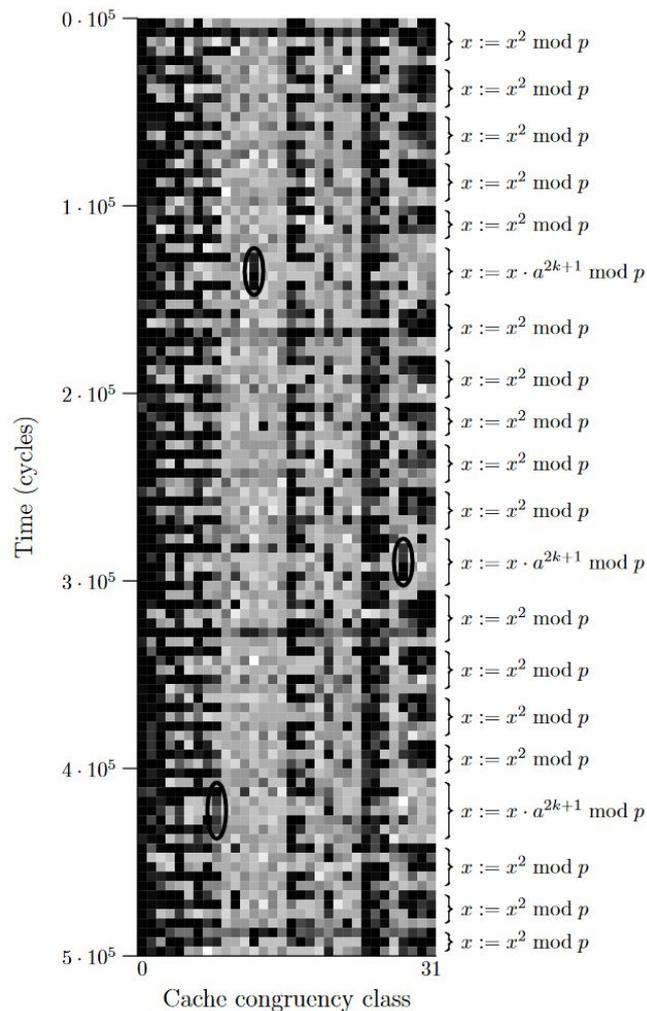
1996-2015 Mainly side channels on  
cryptography (threat model!)



# Past: cryptographic attacks

1996-2015 Mainly side channels on  
cryptography (threat model!)

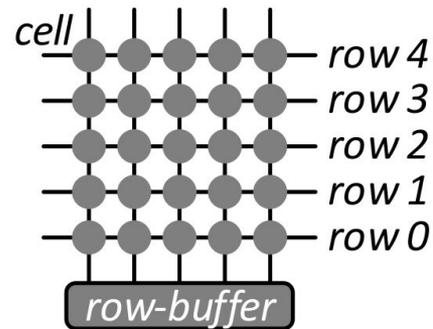
Colin Percival (2005): “Cache Missing  
for fun and profit”



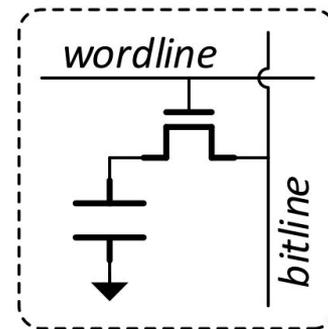
# Past: Moving beyond crypto

ISCA 2014 + BlackHat US 2015:

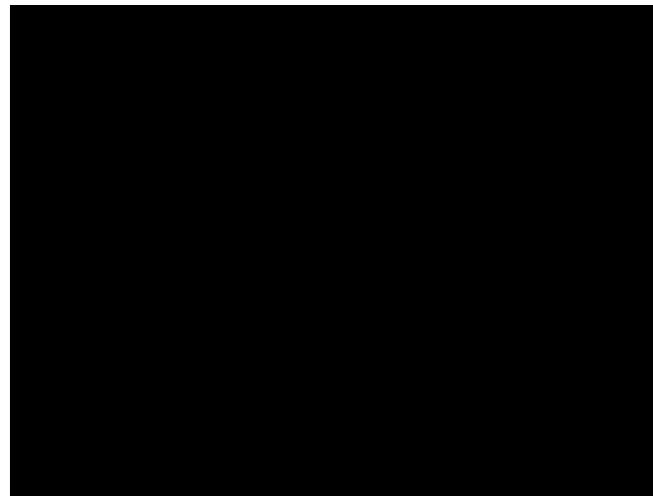
**Rowhammer**



**a.** Rows of cells



**b.** A single cell



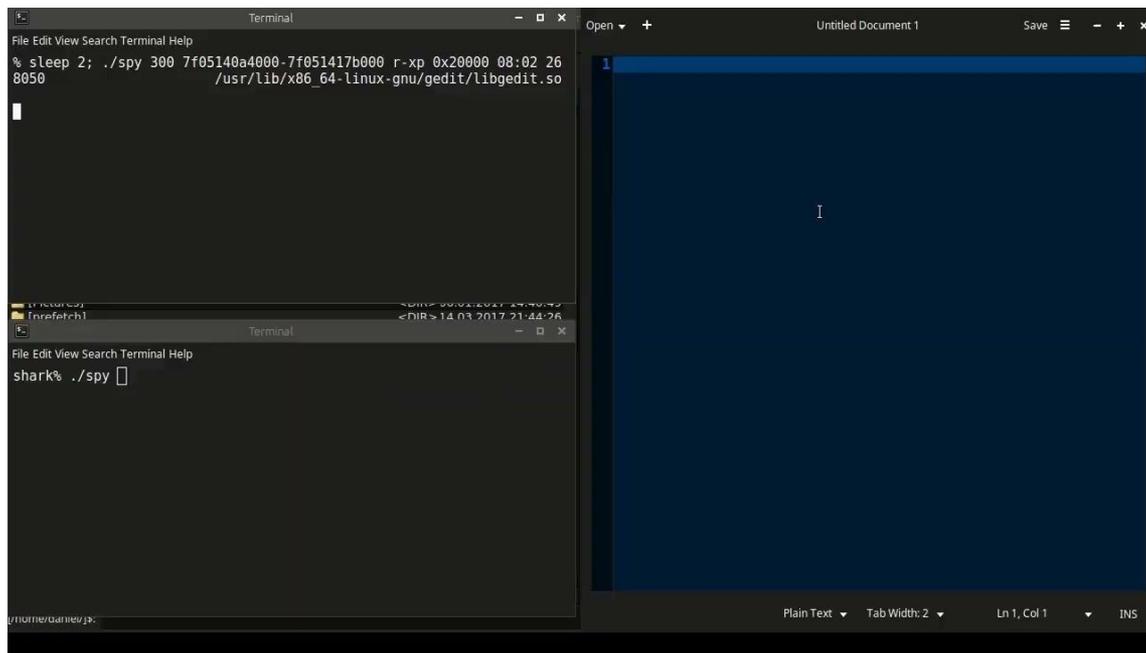
# Past: Moving beyond crypto

ISCA 2014 + BlackHat US 2015:

**Rowhammer**

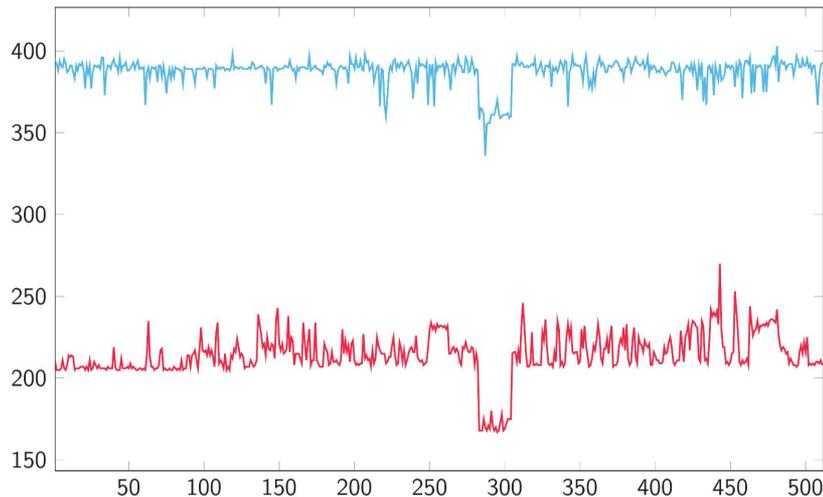
USENIX Security 2015:

**Cache Template Attacks**



## Breaking Kernel Address Space Layout Randomization with Intel TSX

Yeongjin Jang, Sangho Lee, and Taesoo Kim  
*Georgia Institute of Technology*



## Prefetch Side-Channel Attacks: Bypassing SMAP and Kernel ASLR

Daniel Gruss\*

Clémentine Maurice\*

Anders Fogh†

Moritz Lipp\*

Stefan Mangard\*

\* Graz University of Technology † G DATA Advanced Analytics

## Past: Moving beyond crypto

ISCA 2014 + BlackHat US 2015:

**Rowhammer**

USENIX Security 2015:

**Cache Template Attacks**

CCS + BlackHat US 2016:

**Breaking KASLR**

**CacheQuote: Efficiently Recovering Long-term Secrets of SGX EPID via Cache Attacks**

Fergus Dall<sup>1</sup>, Gabrielle De Micheli<sup>2</sup>, Thomas Eisenbarth<sup>3,4</sup>, Daniel Genkin<sup>2,5</sup>,  
Nadia Heninger<sup>2</sup>, Ahmad Moghimi<sup>4</sup> and Yuval Yarom<sup>1,6</sup>

**Application**

**Past:  
Moving beyond crypto**

ISCA 2014 + BlackHat US 2015:  
**Rowhammer**



USENIX Security 2015:  
**Cache Template Attacks**

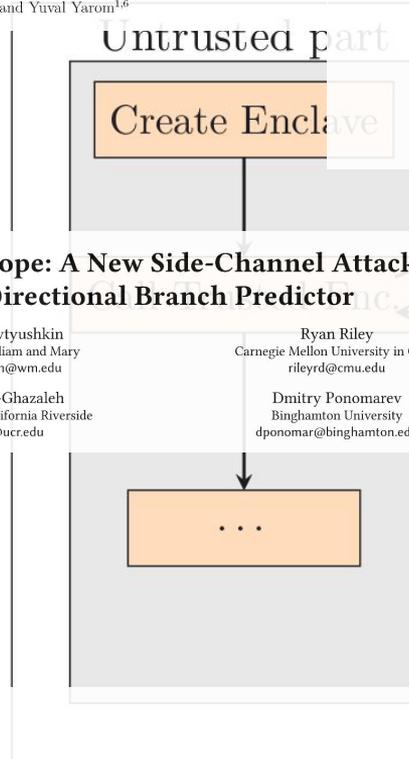
CCS + BlackHat US 2016:  
**Breaking KASLR**

2017: Many academic works on **attacking TEEs with side channels**

**BranchScope: A New Side-Channel Attack on Directional Branch Predictor**

Dmitry Evtushkin  
College of William and Mary  
devtyushkin@wm.edu  
  
Nael Abu-Ghazaleh  
University of California Riverside  
naelag@ucr.edu

Ryan Riley  
Carnegie Mellon University in Qatar  
rileyrd@cmu.edu  
  
Dmitry Ponomarev  
Binghamton University  
dponomar@binghamton.edu



**Controlled-Channel Attacks: Deterministic Side Channels for Untrusted Operating Systems**

Yuanzhong Xu  
The University of Texas at Austin  
yxu@cs.utexas.edu

Weidong Cui  
Microsoft Research  
wdcui@microsoft.com

Marcus Peinado  
Microsoft Research  
marcuspe@microsoft.com

Trusted Enc.

**COPYCAT: Controlled Instruction-Level Attacks on Enclaves**

Daniel Moghimi<sup>1</sup>, Jo Van Bulck<sup>2</sup>, Nadia Heninger<sup>3</sup>, Frank Piessens<sup>2</sup>, and Ber...

- <sup>1</sup>Worcester Polytechnic Institute, Worcester, MA, USA
- <sup>2</sup>imec-DistriNet, KU Leuven, Leuven, Belgium
- <sup>3</sup>University of California, San Diego, CA, USA

**How Trusted Executing Environments Fuel Power on Microarchitectural**

**CacheZoom: How SGX Amplifies The Power of Cache Attacks**

Ahmad Moghimi  
Worcester Polytechnic Institute  
amoghimi@wpi.edu

Gorka Irazoqui  
Worcester Polytechnic Institute  
girazoki@wpi.edu

**Leaky Cauldron on the Dark Land: Understanding Memory Side-Channel Hazards in SGX**

Wenhao Wang<sup>1</sup>, Guoxing Chen<sup>3</sup>, Xiaorui Pan<sup>2</sup>, Yinqian Zhang<sup>3</sup>, XiaoFeng Wang<sup>2</sup>,  
Vincent Bindschaedler<sup>4</sup>, Haixu Tang<sup>2</sup>, Carl A. Gunter<sup>4\*</sup>  
<sup>1</sup>SKLOIS, Institute of Information Engineering, Chinese Academy of Sciences & Indiana University Bloomington

# Past: Moving beyond crypto

ISCA 2014 + BlackHat US 2015:

**Rowhammer**

USENIX Security 2015:

**Cache Template Attacks**

CCS + BlackHat US 2016:

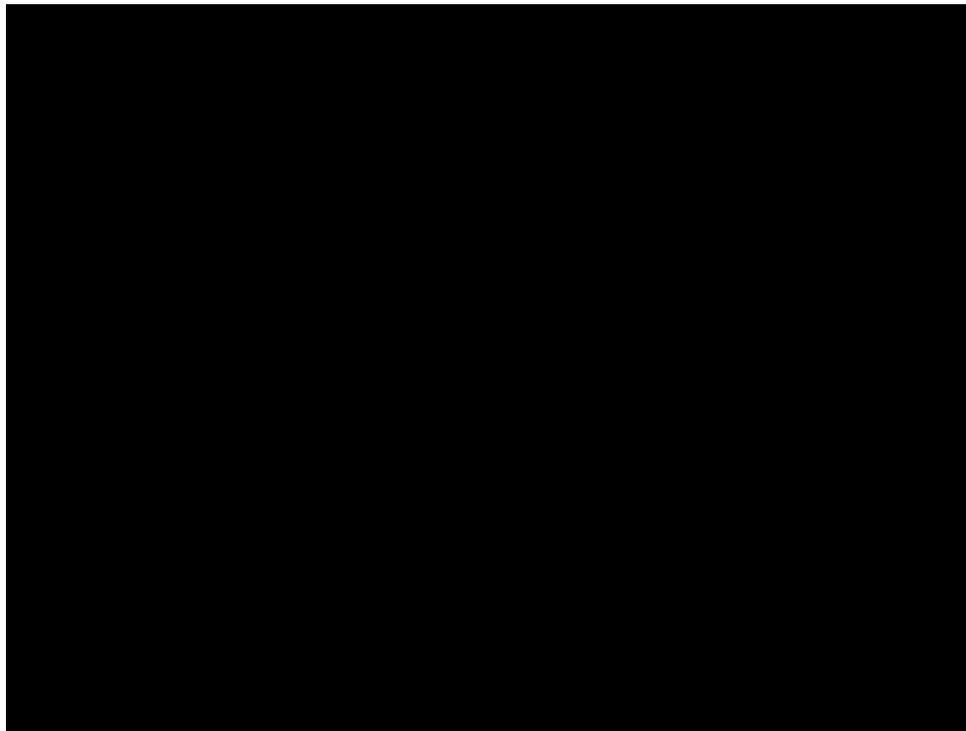
**Breaking KASLR**

2017: Many academic works on **attacking**

**TEEs with side channels**

USENIX + BlackHat US 2018, S&P 2019:

**Spectre & Meltdown**



# ① preface



*architectural*

**time**

① preface

② trigger instruction 



*architectural*

*transient execution*

**time** 

① preface



② trigger instruction 

③ transient access to secret

*architectural*

*transient execution*

**time** 

① preface



② trigger instruction 

③ transient access to secret

④ transmission of secret 

*architectural*

*transient execution*

**time** 

① preface



② trigger instruction 

⑤ fixup 

③ transient access to secret

④ transmission of secret 

*architectural*

*transient execution*

*architectural*

**time** 

① preface



*architectural*

② trigger instruction 

③ transient access to secret

④ transmission of secret 



*transient execution*

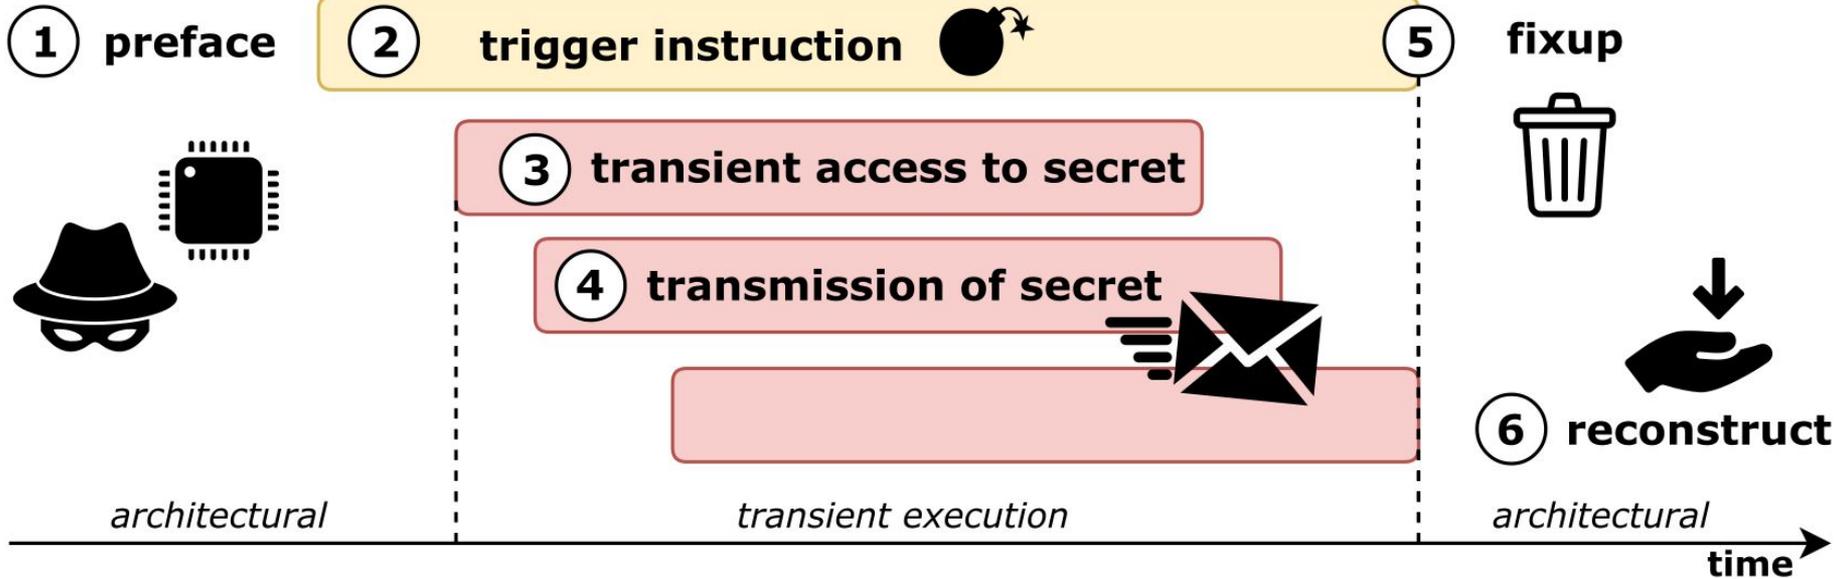
⑤ fixup 

⑥ reconstruct 

*architectural*

**time** 

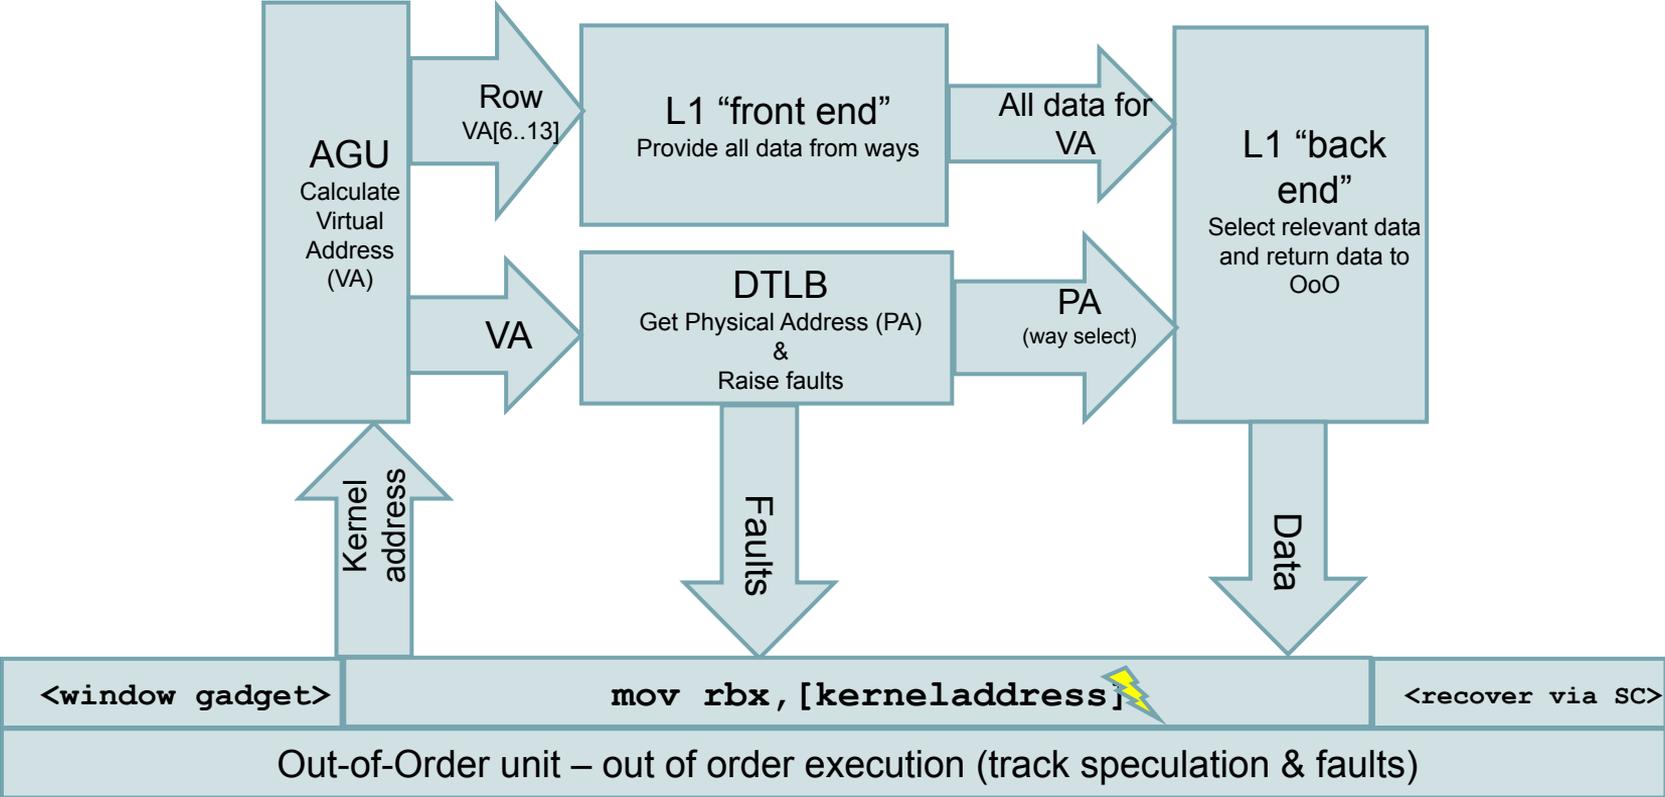
# Past: Meltdown



```
<window gadget> | mov rbx, [kerneladdress] | <recover via SC>
```

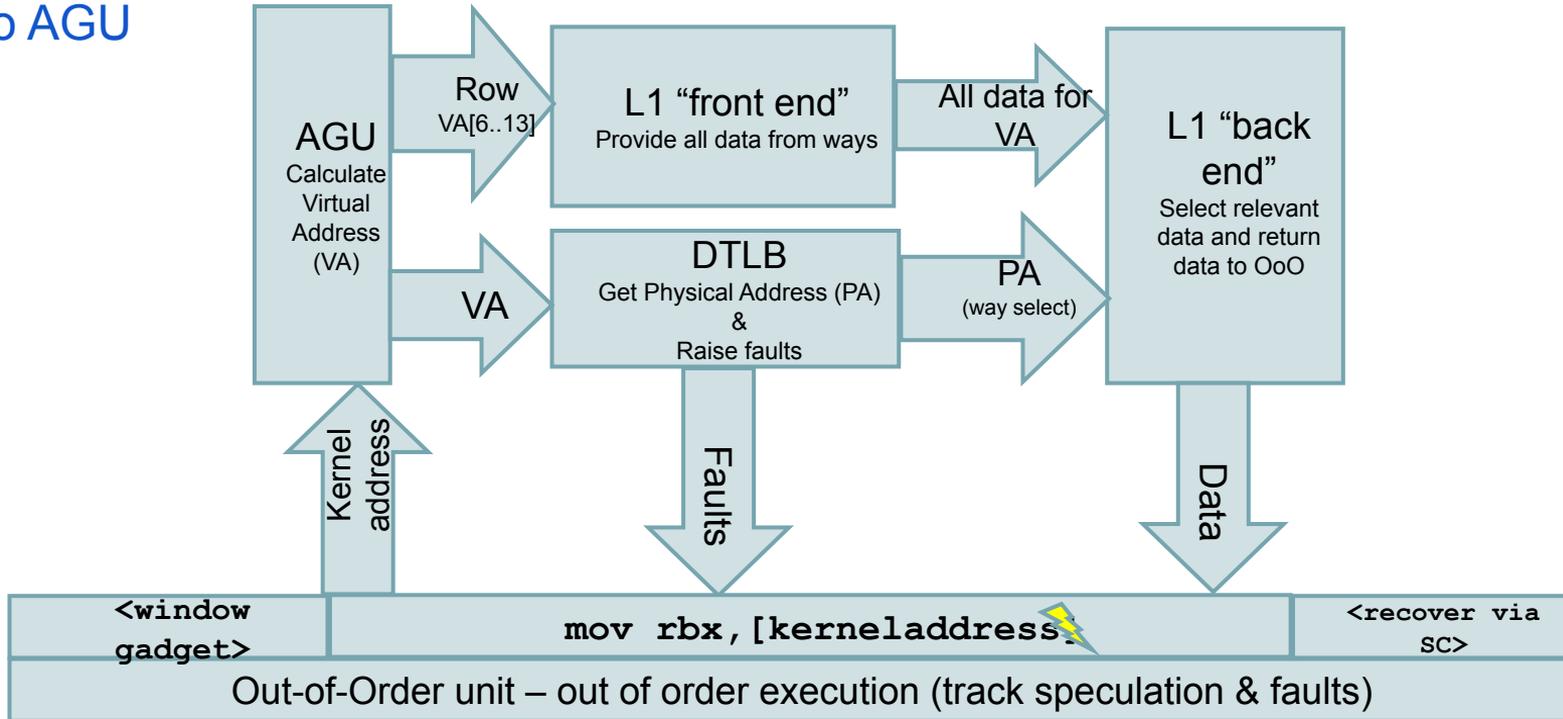
Out-of-Order unit – out of order execution (track speculation & faults)

# Meltdown: Details



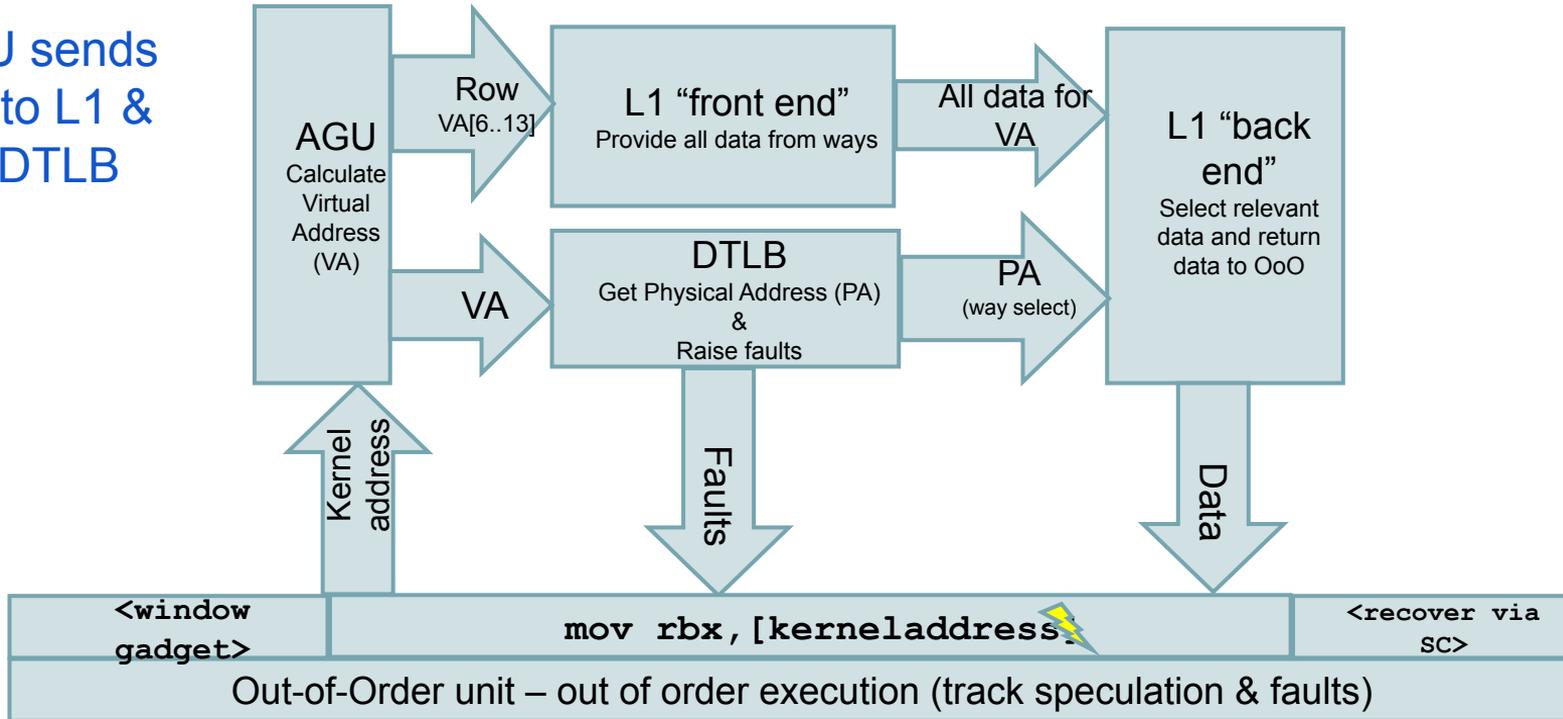
# Meltdown: Details

## 1. OoO Trigger load to AGU



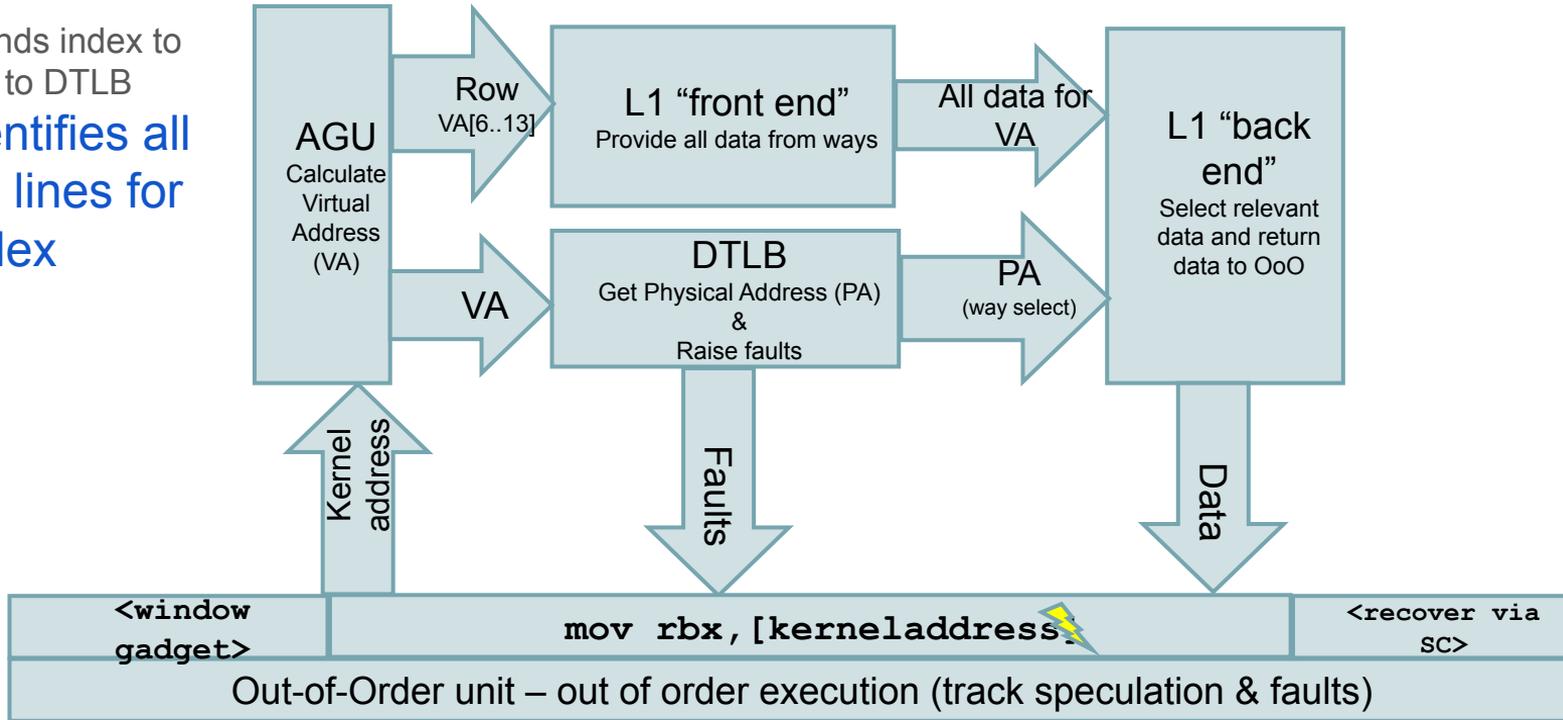
# Meltdown: Details

1. OoO Trigger load to AGU
2. AGU sends index to L1 & VA to DTLB



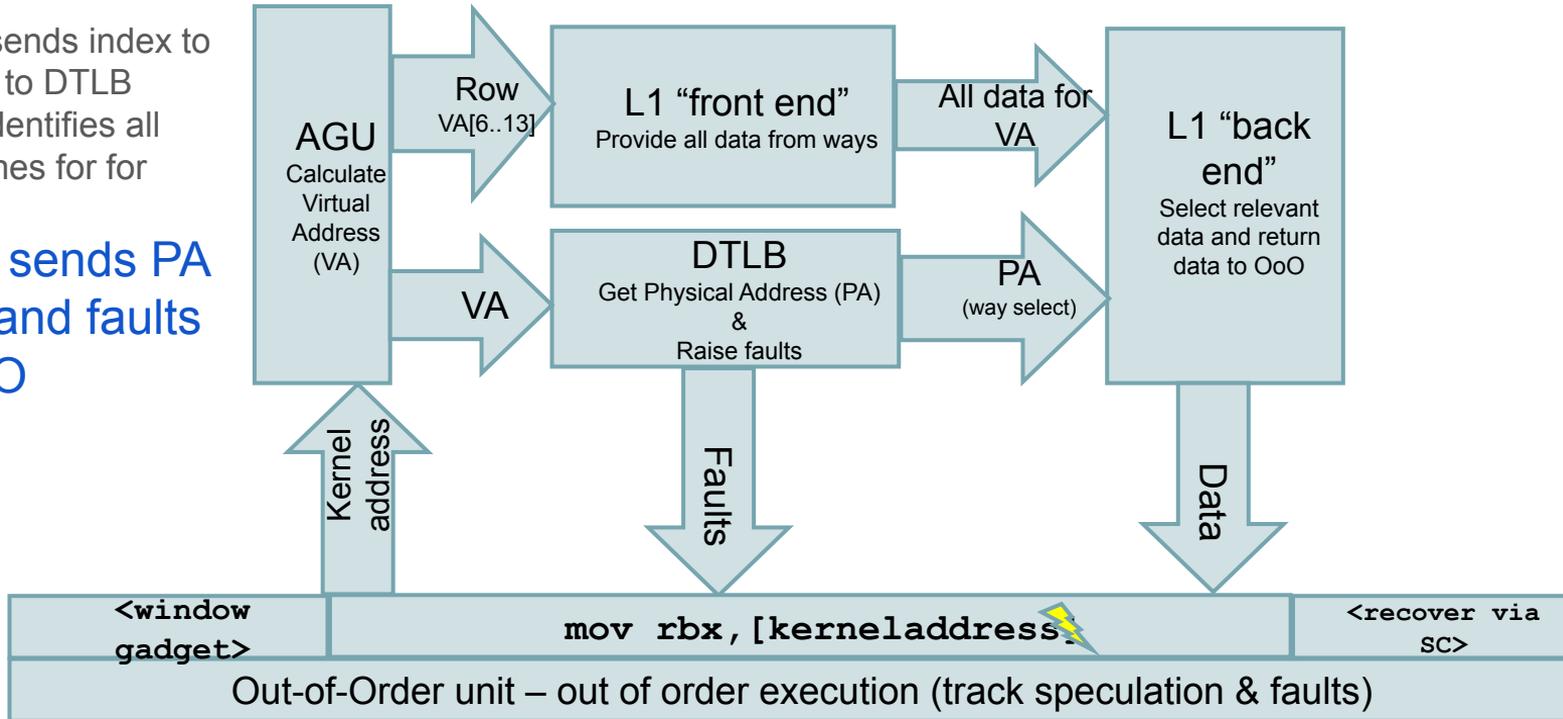
# Meltdown: Details

1. OoO Trigger load to AGU
2. AGU sends index to L1 & VA to DTLB
3. L1 identifies all cache lines for for index



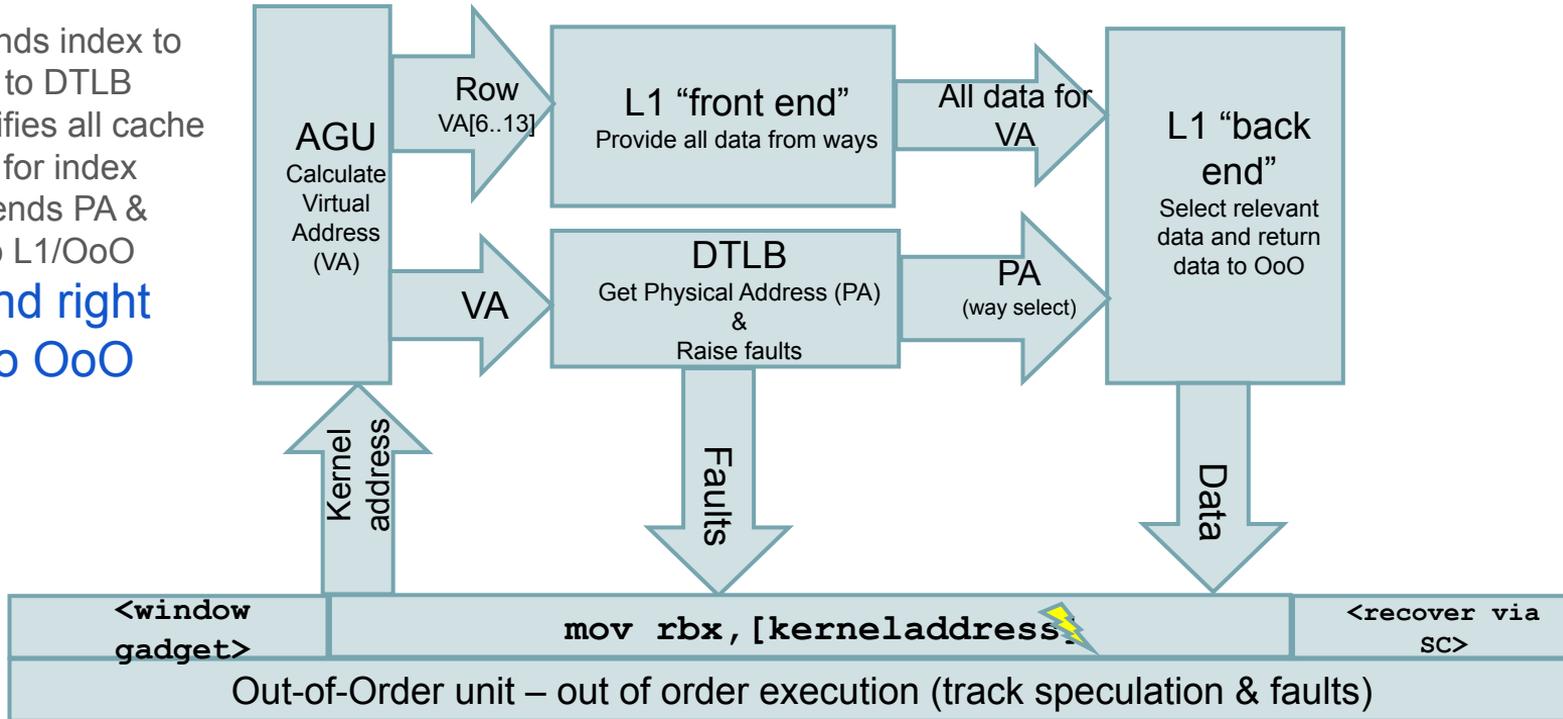
# Meltdown: Details

1. OoO Trigger load to AGU
2. AGU sends index to L1 & VA to DTLB
3. a L1 identifies all cache lines for for index
4. **DTLB sends PA to L1 and faults to OoO**



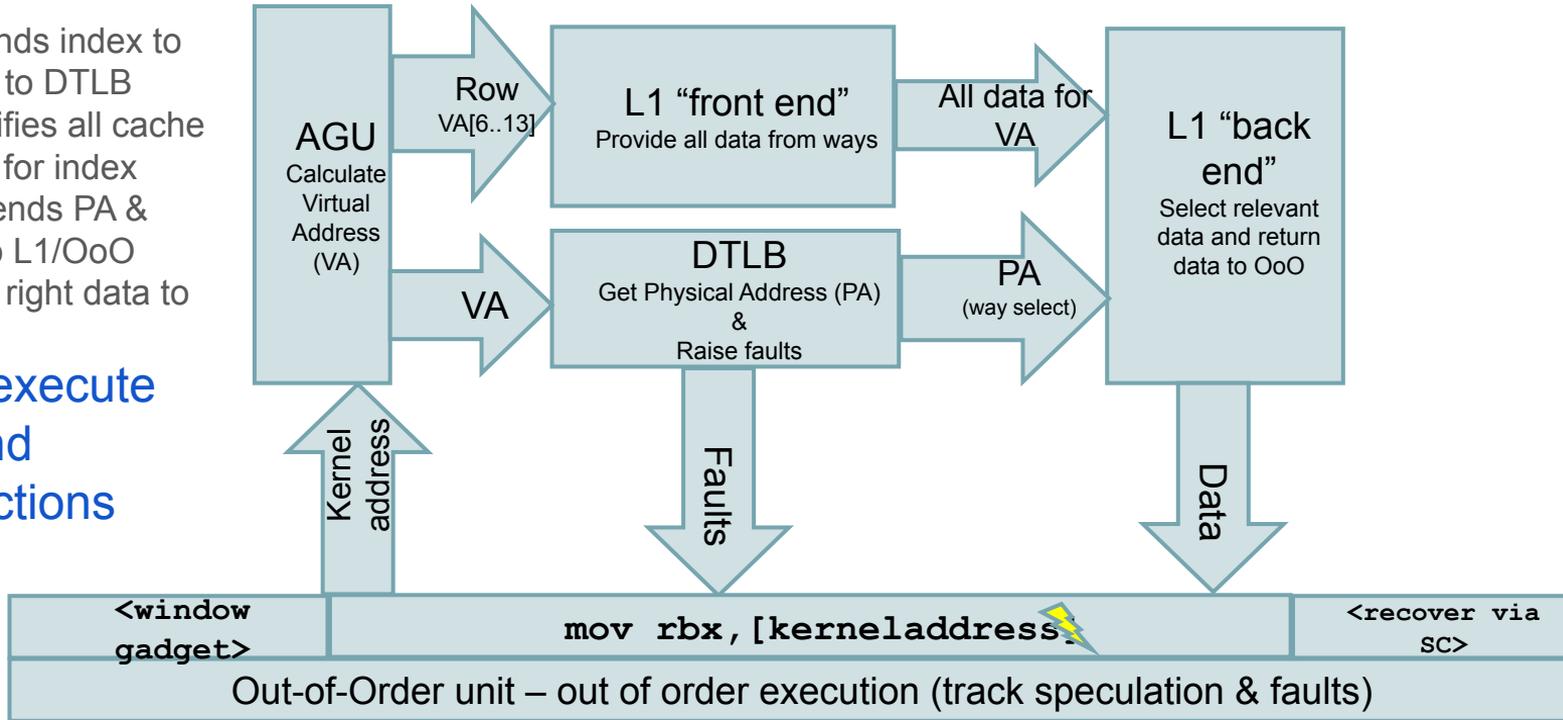
# Meltdown: Details

1. OoO Trigger load to AGU
2. AGU sends index to L1 & VA to DTLB
3. L1 identifies all cache lines for for index
4. DTLB sends PA & faults to L1/OoO
5. L1 send right data to OoO

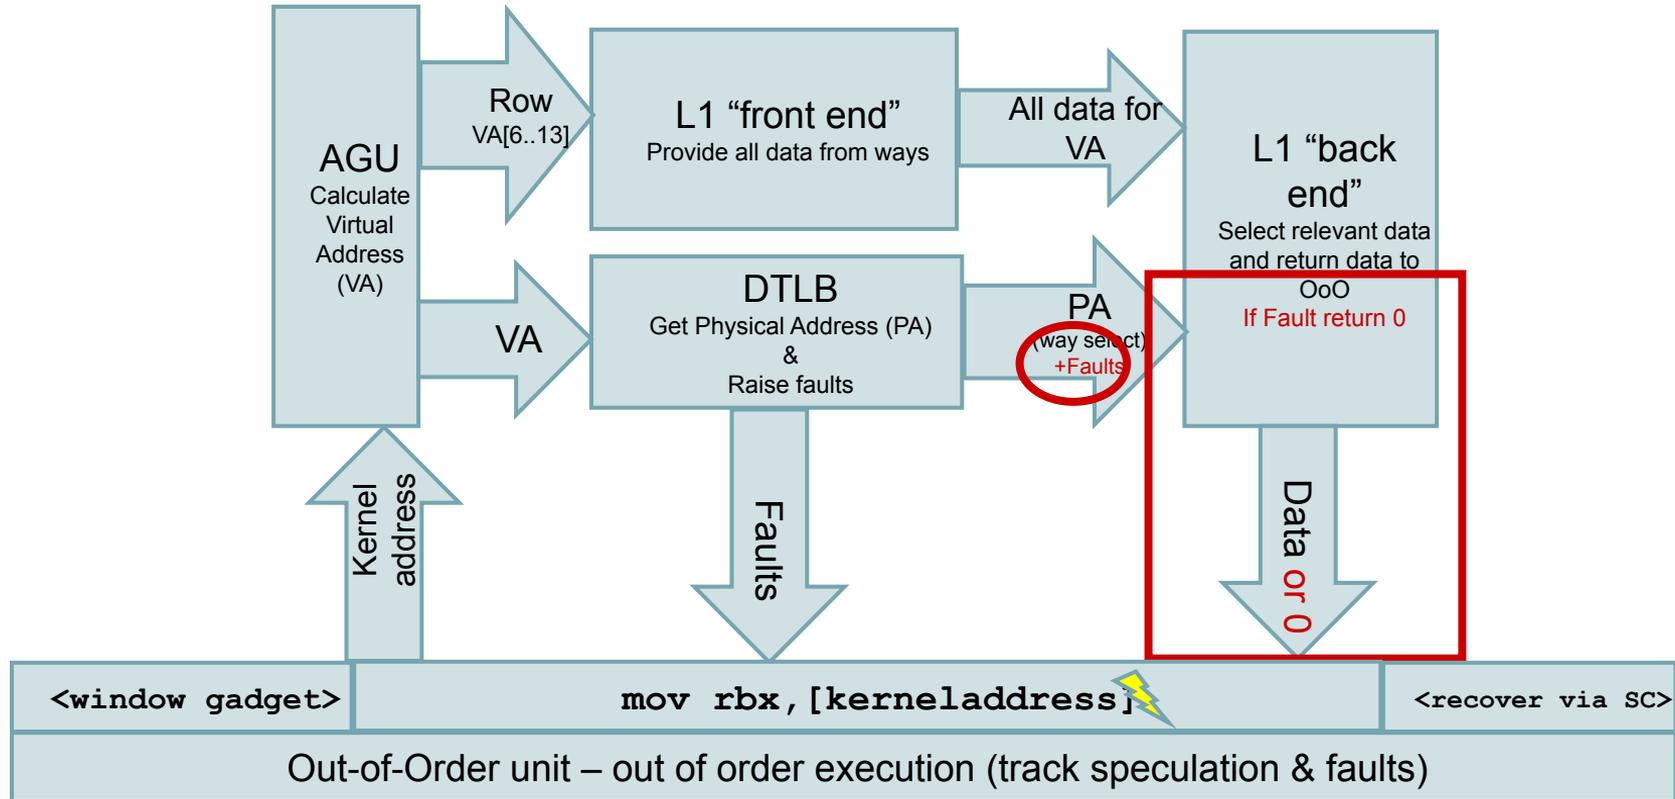


# Meltdown: Details

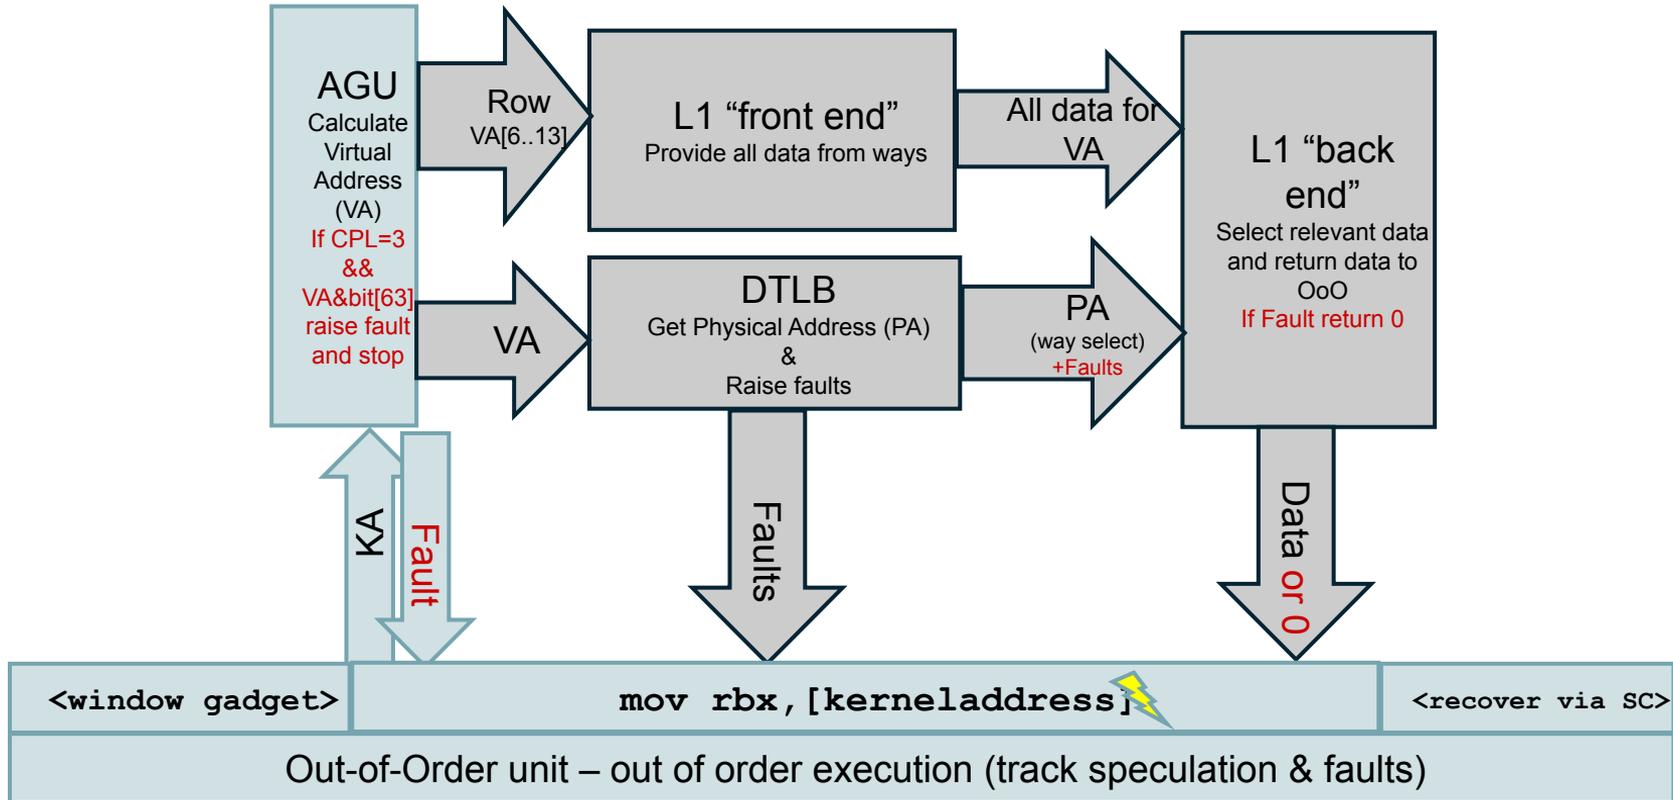
1. OoO Trigger load to AGU
2. AGU sends index to L1 & VA to DTLB
3. L1 identifies all cache lines for for index
4. DTLB sends PA & faults to L1/OoO
5. L1 send right data to OoO
6. OoO execute depend instructions



# The First Meltdown Mitigations



# Meltdown defense in depth (LASS)



# Spectre and LVI

<b>Methodology</b>		<b>Leakage</b> 	<b>Injection</b> 
<i><math>\mu</math>-Arch Buffer</i>			
<b>Prediction history</b>	PHT	BranchScope [79], Bluethunder [131]	Spectre-PHT [174]
	BTB	SBPA [8], BranchShadow [182]	Spectre-BTB [174]
	RSB	Hyper-Channel [46]	Spectre-RSB [177, 200]
	STL	—	Spectre-STL [128]
<b>Program data</b>	NULL	EchoLoad [49]	LVI-NULL [311]
	L1D	Meltdown [193], Foreshadow [310]	LVI-L1D [311]
	FPU	LazyFP [291]	LVI-FPU [311]
	SB	Store-to-Leak [270], Fallout [48]	LVI-SB [311]
	LFB/LP	ZombieLoad [276], RIDL [267]	LVI-LFB/LP [311]

An illustration of a server room. In the background, there are several tall server racks with glass doors. In the foreground, two laptops are open on a desk. To the left of the left laptop is a smartphone. To the right of the right laptop is another smartphone and a small blue device. The word "Present" is written in large, bold, black letters in the center of the image.

**Present**

# Present: Trends

Attack type	Activity level	(Point) Mitigation	Notable
Crypto side channels		Guidance & DOIT	Data dependent features for example data dependent prefetchers
Transient execution vulnerabilities		Hardware + Software +on/off switches Workarounds	Predictive store forwarding
Stale data vulnerabilities		Microcode Patches or SW Mitigation (if possible)	Not any recent attacks
Logical bugs		Microcode Patches (if possible)	Reptar, CacheWarp
Physical properties			Hertzbleed, Collide+Power
Exploitation methods			Spectre & Power

# Logic Issues

# Reptar - What's supposed to happen

REPZ is a prefix that will repeat an operation until the Z-flag becomes zero.

MOVSB will copy a single byte from DS:[RSI] to ES:[RDI] and increment both registers and decrement RCX & update flags.

REPZ MOVSB is thus a simple memcpy.

The REX-prefix (REX.PF) changes the meaning of how explicit operands of an instruction are interpreted. MOVSB doesn't have any explicit operands.

If you use the REX-prefix with REPZ MOVSB the CPU should ignore the prefix entirely

Opcode input

Rex.pf repnz movsb

Parsing input  
Dropping the  
REX-PF

ReX pf repnz movsb => 0xABCD

Issue uOps from  
Parsed input

0xABCD => uOps for repnz  
0xABCD

# Reptar - The bug

When the REX-prefix is parsed instead of ignored a single bit is overwritten.

This cause an invalid input to be used to generate uOps.

Under certain conditions this leads to a machine check. Careful analysis found that a condition could potentially lead to privilege escalation.

A microcode change that mitigates the issue has been made public.

Opcode input

Rex.pf repnz movsb

Parsing input  
Parses the  
REX-PF and  
Overwrite a bit

Rex.pf repnz movsb => 0xBBCD

Issue uOps from  
Parsed input

INVALID INPUT for uOps issue



# CacheWarp: Software-based Fault Injection using Selective State Reset

## Cachewarp

Confidential VM (encrypted but basically no data integrity)

**invd** instruction can invalidate a single cache line

Attack in three steps:

1. let confidential VM modify a target cache line
2. use **invd** to drop the modification
3. confidential VM continues with an outdated value

Ruiyi Zhang  
*CISPA Helmholtz Center  
for Information Security*

Lorenz Hetterich  
*CISPA Helmholtz Center  
for Information Security*

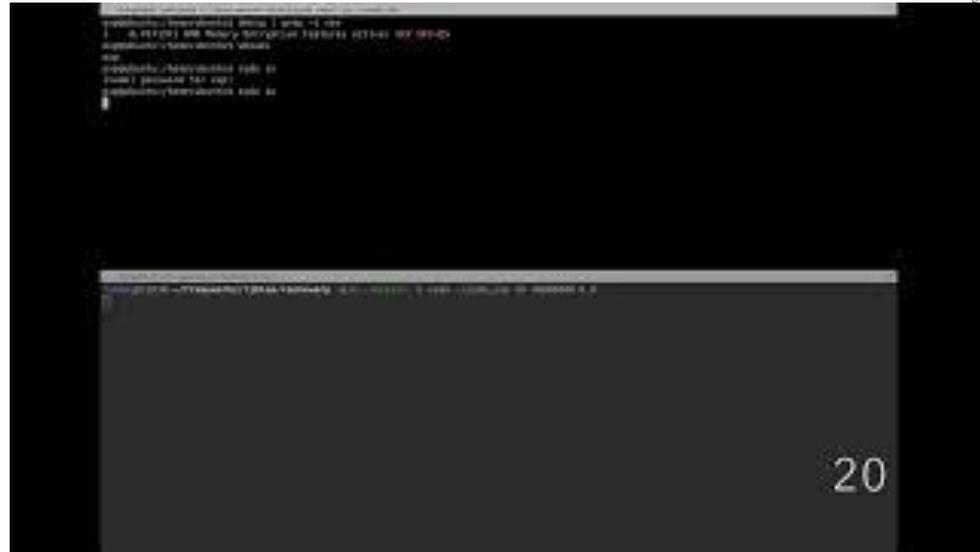
Lukas Gerlach  
*CISPA Helmholtz Center  
for Information Security*

Youheng Lü  
*Independent*

Michael Schwarz  
*CISPA Helmholtz Center for Information Security*

Daniel Weber  
*CISPA Helmholtz Center  
for Information Security*

Andreas Kogler  
*Graz University of Technology*



# Zenbleed

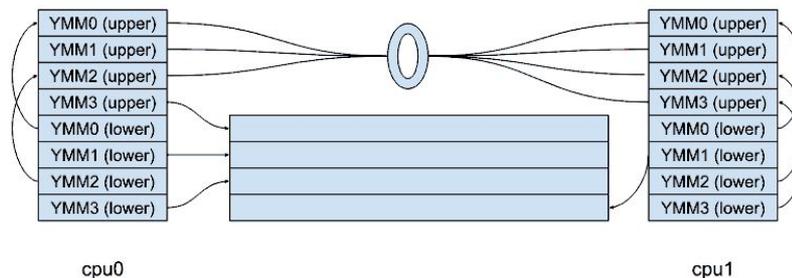
Register names are just for the user, CPU uses register file

XMM Register Merge Optimization: merge registers (e.g. zero registers)

also: for zero just set a zero-bit

Zenbleed:

1. misspeculation
2. **vzeroupper** → set zero-bit
3. merge → storage in register file released
4. victim stores data in this register
5. unroll misspeculation
6. architectural access to a victim data



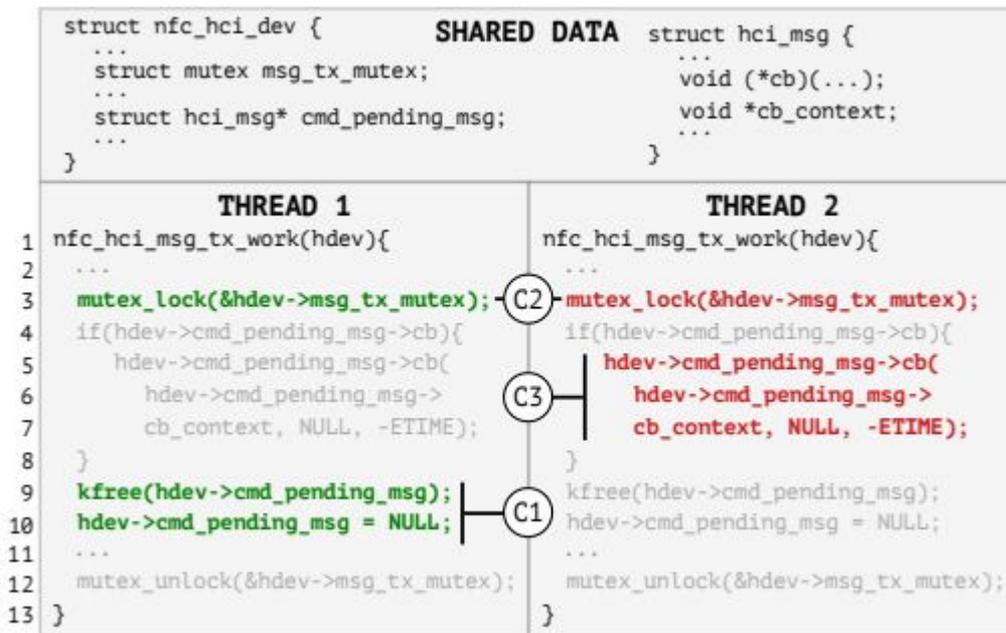
# Exploitation Techniques

# Exploitation techniques - example

GhostRace: Exploiting and Mitigating Speculative Race Conditions - Hany Ragab et. al.

Spectre v1. variant that speculatively bypasses synchronization primitives.

Existing methods of mitigating Spectre v1 remain effective.



Quote from the papers abstract:

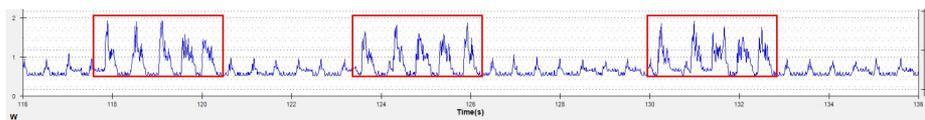
*“There’s is security, and then there’s just being ridiculous”* - Linus Torvalds, on Speculative Race Conditions

# Physical Domain in Software

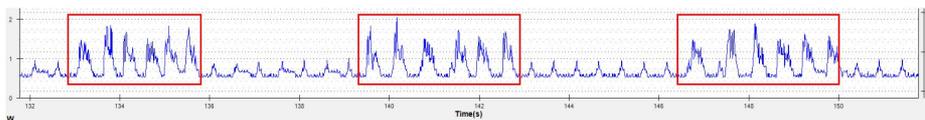
# Software-based Power Analysis

before 2020: mainly fingerprinting

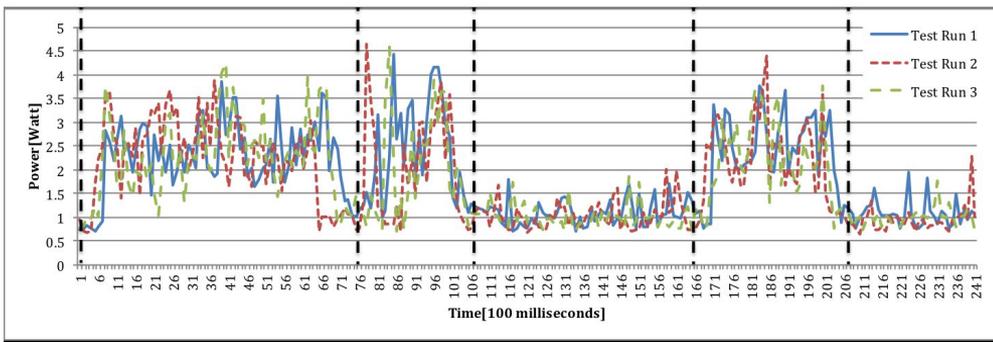
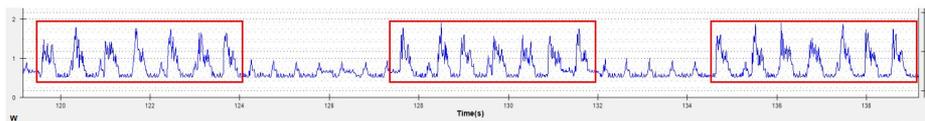
5 letters



6 letters



7 letters



# Software-based Power Analysis

before 2020: mainly fingerprinting

2020: Platypus

full recovery of cryptographic keys

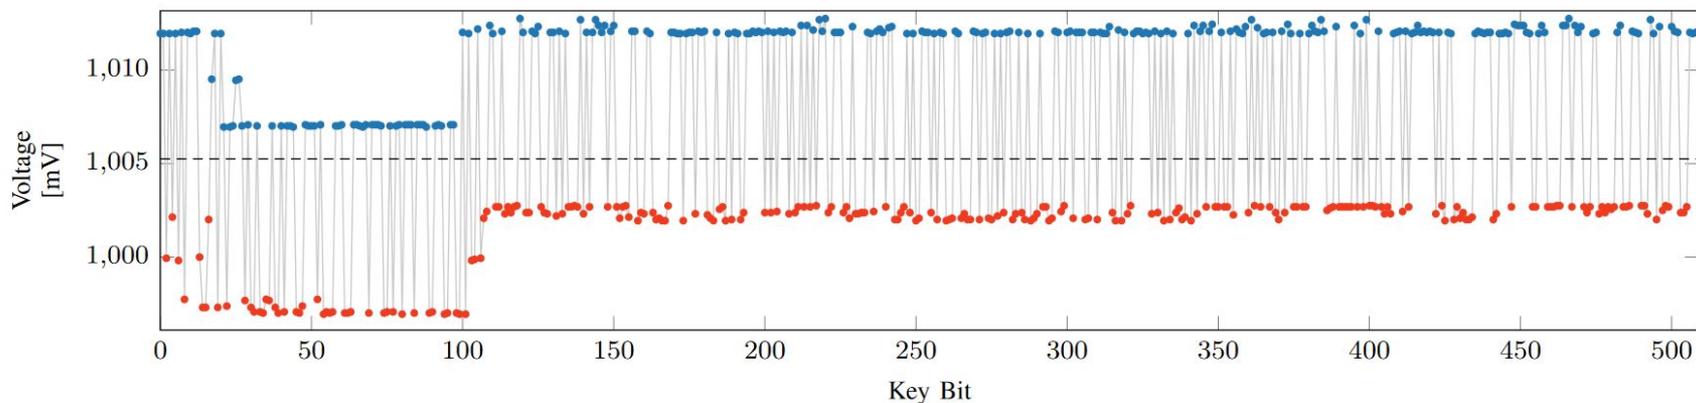


Fig. 13: Core voltage per measured instruction for each key bit offset in the fixed window length implementation of mbed TLS inside an SGX enclave on the Xeon E3-1275 v5. The blue marks represent 1 bits, while the red marks represent 0 bits. Using a threshold (dashed line), they can easily be distinguished.

# Software-based Power Analysis

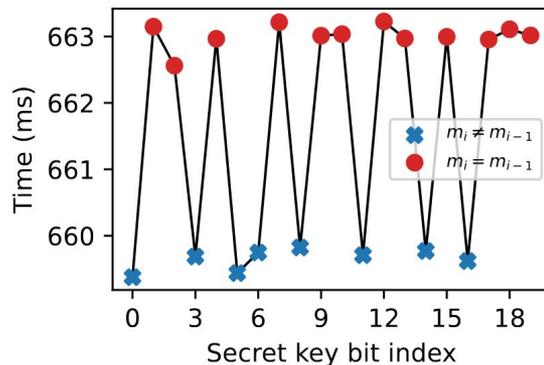
before 2020: mainly fingerprinting

2020: Platypus

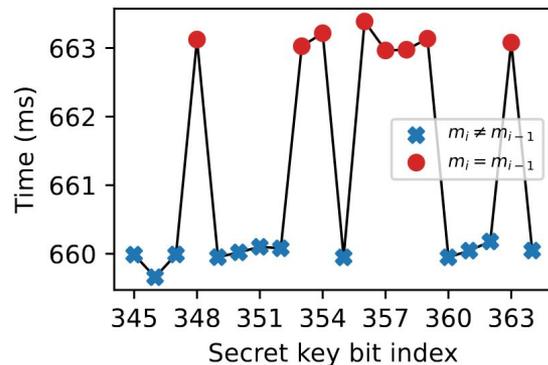
full recovery of cryptographic keys

2023: Hertzbleed

DVFS makes timing a proxy for energy consumption → remote attacks



(a) CIRCL first 20 bits



(b) CIRCL last 20 bits

# Software-based Power Analysis

before 2020: mainly fingerprinting

2020: Platypus

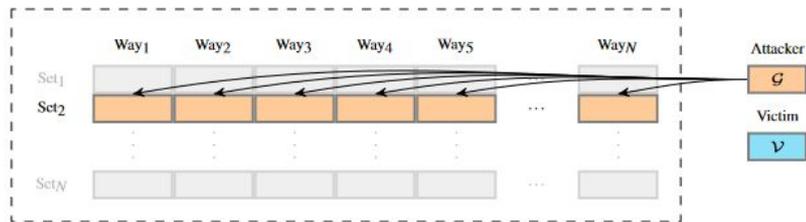
full recovery of cryptographic keys

2023: Hertzbleed

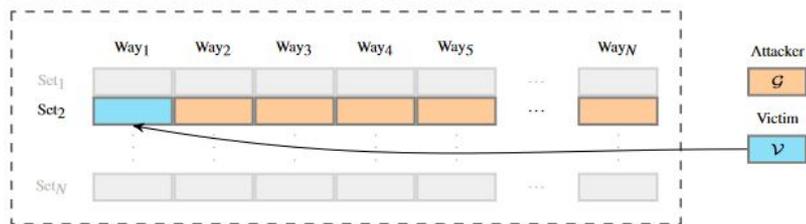
DVFS makes timing a proxy for energy consumption → remote attacks

2023: Collide+Power

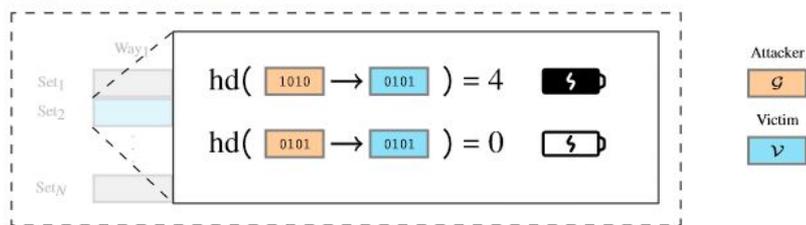
Generic Attacks (not just crypto)



(a) **Step 1:** The attacker primes each cache line of the target cache set with the attacker-controlled guess  $\mathcal{G}$ .



(b) **Step 2:** The victim accesses the secret  $\mathcal{V}$  and forces a cache line to change from  $\mathcal{G}$  to  $\mathcal{V}$ .



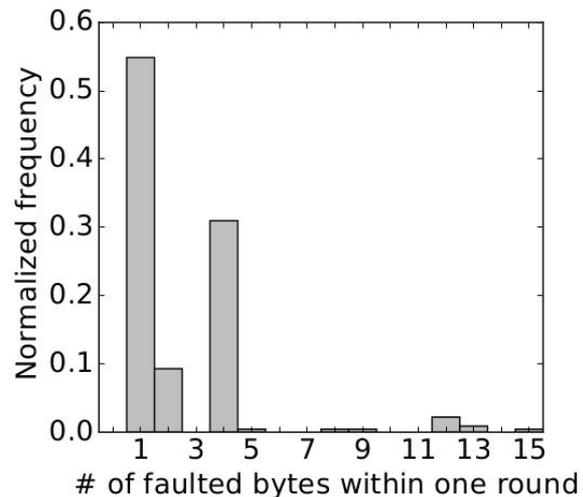
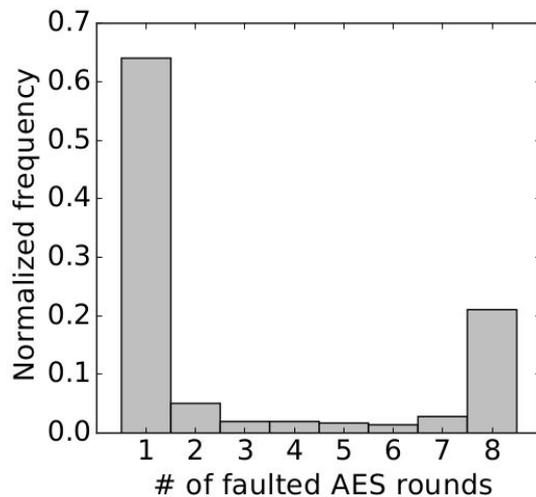
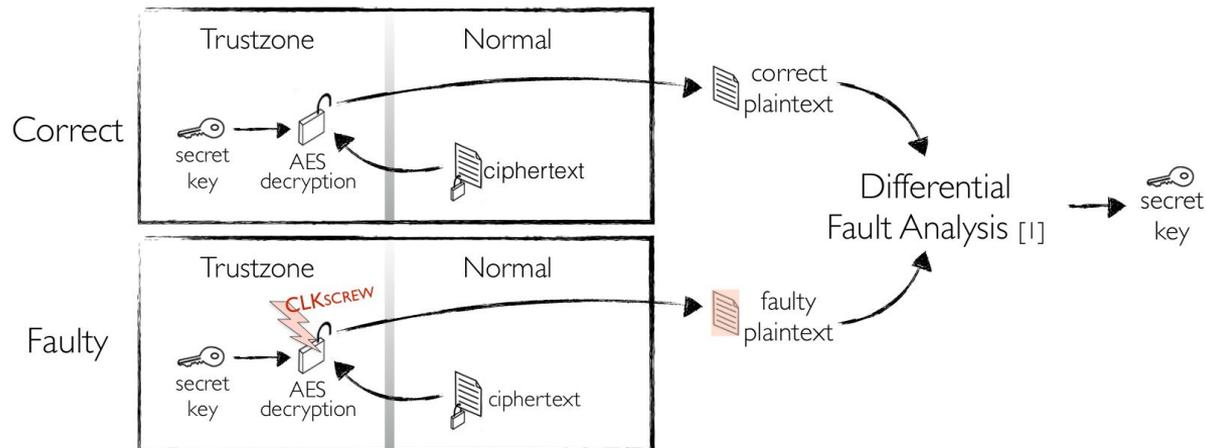
(c) **Step 3:** The energy consumption during this change is proportional to the number of bit changes between  $\mathcal{G}$  and  $\mathcal{V}$ .



# Software-based Fault Attacks

since 2015: Rowhammer still not solved!

2017: CLKScrew overclock and attack Arm TrustZone

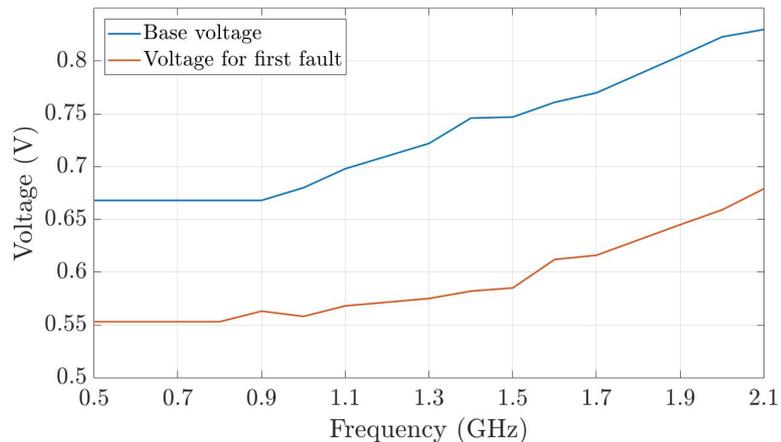


# Software-based Fault Attacks

since 2015: Rowhammer  
still not solved!

2017: CLKSkrew  
overclock and attack Arm TrustZone

2020: Plundervolt (VoltJockey,  
V0ltpwn, VoltPillager)  
undervolt and attack Intel SGX



```
Summary
-----
Iterations:      10000000
Start Voltage:   -252
End Voltage:     0
Stop after x drops: 10
Voltage steps:   1
Threads:         4
Operand1:        0x00000000deadbeef
Operand2:        0x1122334455667788
Operand1 is:     fixed value
Operand2 is:     fixed value
```

Mitigation efforts

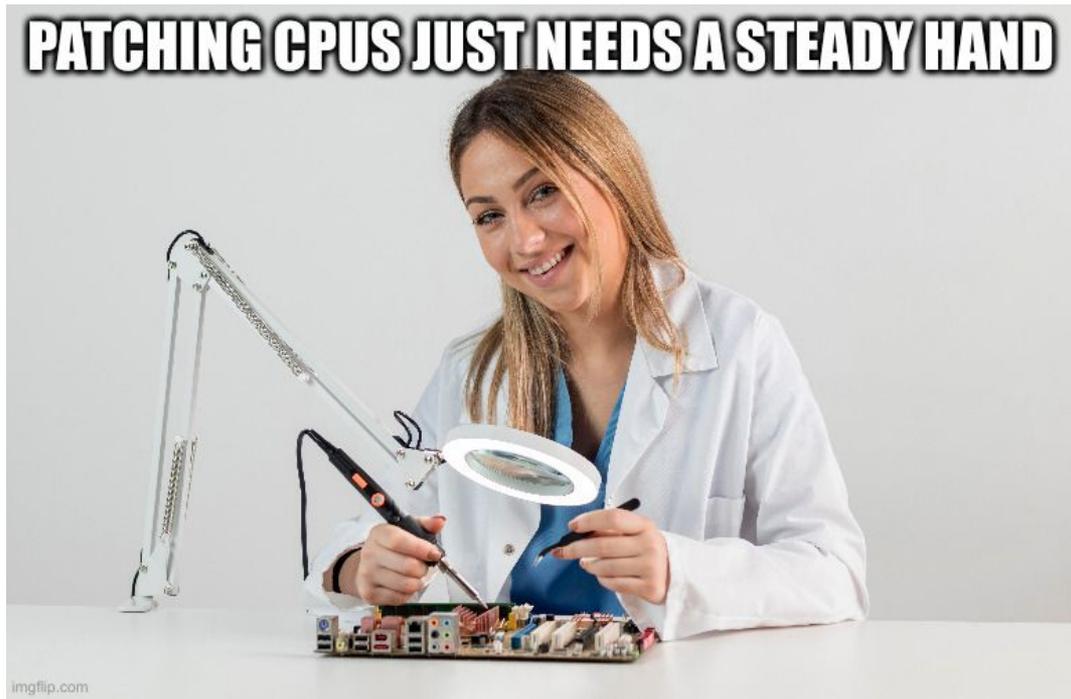
# Limitations of mitigations

Physical hardware cannot be changed in the field



## Limitations of mitigations

Physical hardware cannot be changed in the field



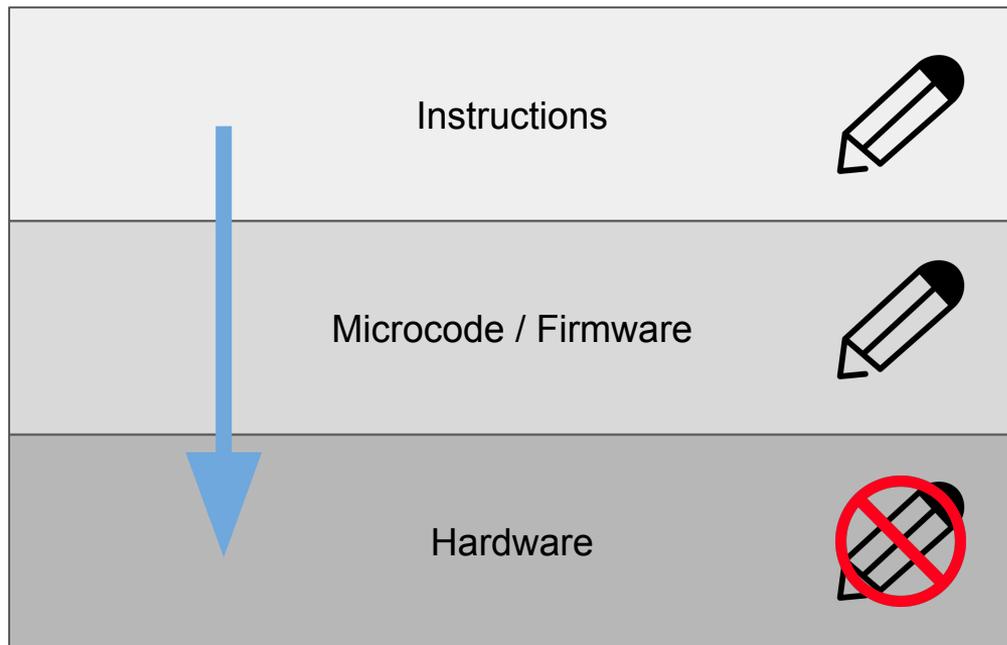
# Limitations of mitigations

Physical hardware cannot be changed in the field

Vendors build in “Survivability features”

**Microcode** is the most common used tool for mitigations.

**Other firmware** is also used



# Limitations of mitigations

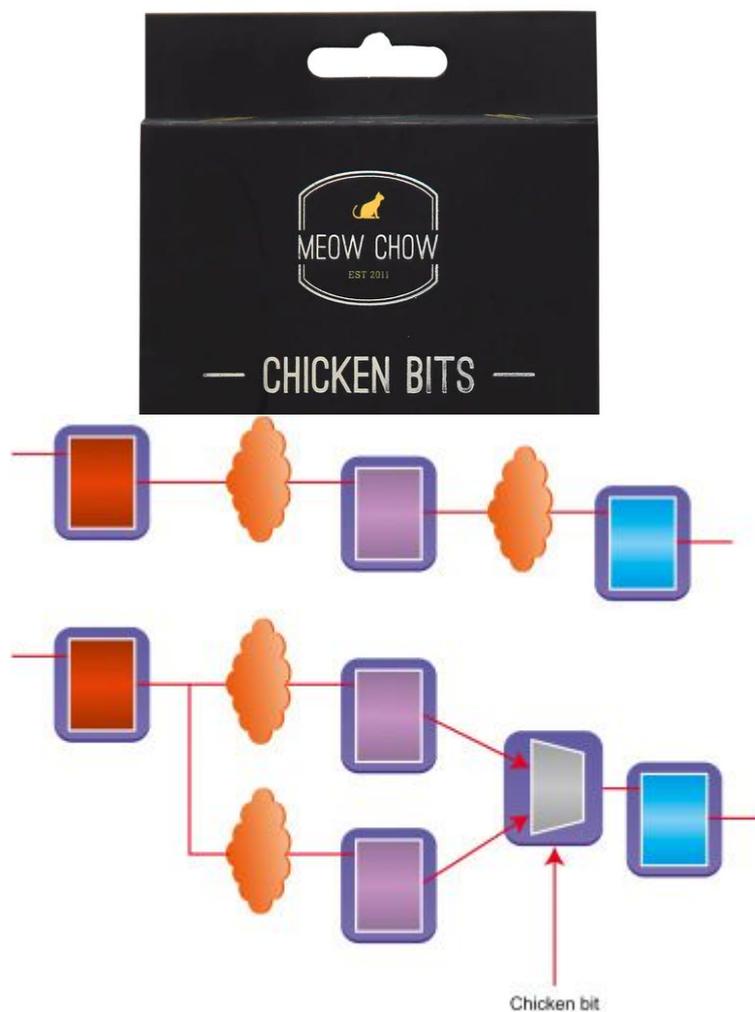
Physical hardware cannot be changed in the field

Vendors build in “Survivability features”

Microcode is the most common used tool for mitigations.

Other firmware is also used

“Chicken bits” to disable / change behavior



# Limitations of mitigations

Physical hardware cannot be changed in the field

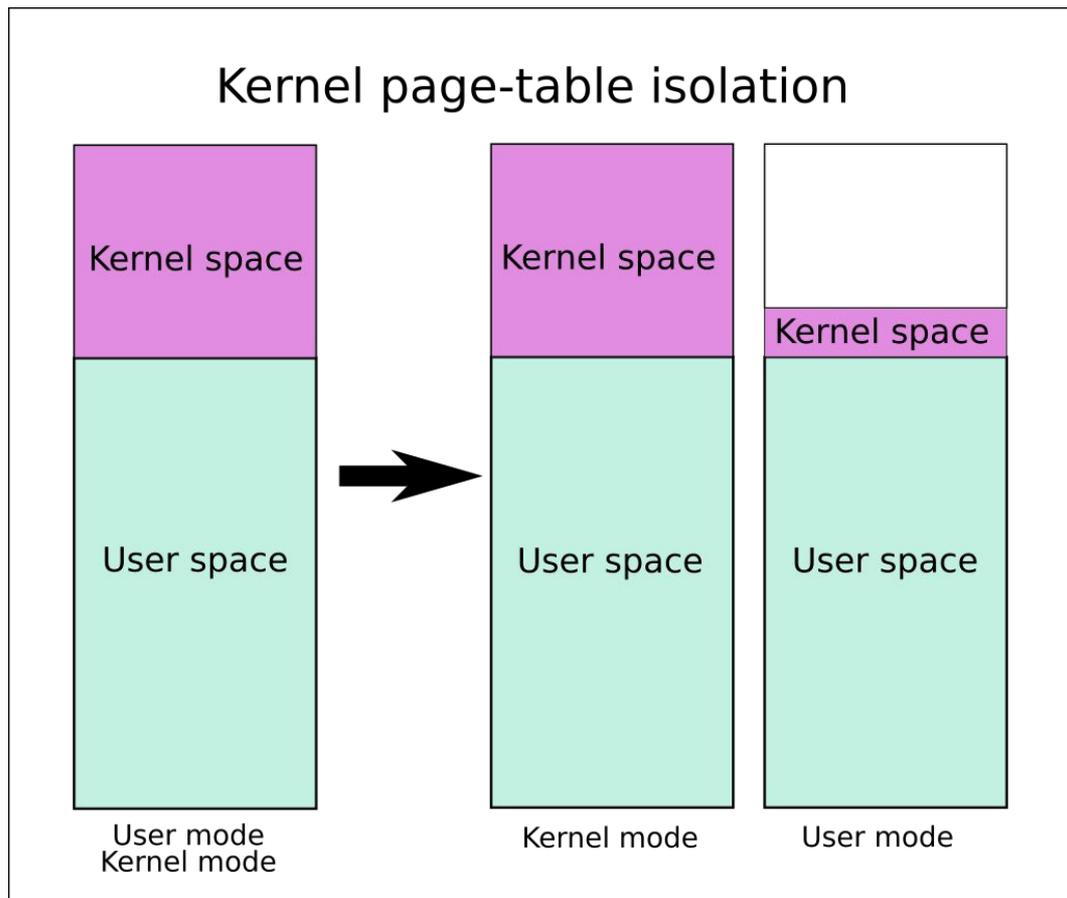
Vendors build in “Survivability features”

Microcode is the most common used tool for mitigations.

Other firmware is also used

“Chicken bits” to disable / change behavior

Some issues are best mitigated in software



# Limitations of mitigations

Physical hardware cannot be changed in the field

Vendors build in “Survivability features”

**Microcode** is the most common used tool for mitigations.

**Other firmware** is also used

“Chicken bits” to disable / change behavior

Some issues are best mitigated in software

Mitigations are **not always possible/reasonable** and almost always **difficult** and **time-consuming** to engineer

# Prevention Pre-silicon

Prevention starts before the product exist: pre-silicon

Pre-silicon is slow and cumbersome as the chips are emulated or simulated.

This makes security validation & research significantly different from software validation

01	Architecture reviews	<ul style="list-style-type: none"><li>• Gives great ROI</li><li>• There is formal and informal reviews on arch</li></ul>
02	Taint tracking	<ul style="list-style-type: none"><li>• Taint tracking has proven useful for some issues</li><li>• Techniques such as CellFT used in production</li></ul>
03	Validation	<ul style="list-style-type: none"><li>• Security properties to standard validation</li><li>• Finds bugs during development</li></ul>
04	Formal validation	<ul style="list-style-type: none"><li>• Formal works well with hardware IP</li><li>• Formal definition of security properties can be done, but not easy</li></ul>
05	Defense in depth & hardening	<ul style="list-style-type: none"><li>• Bug analysis should lead to lessons learned</li></ul>

# Post-silicon

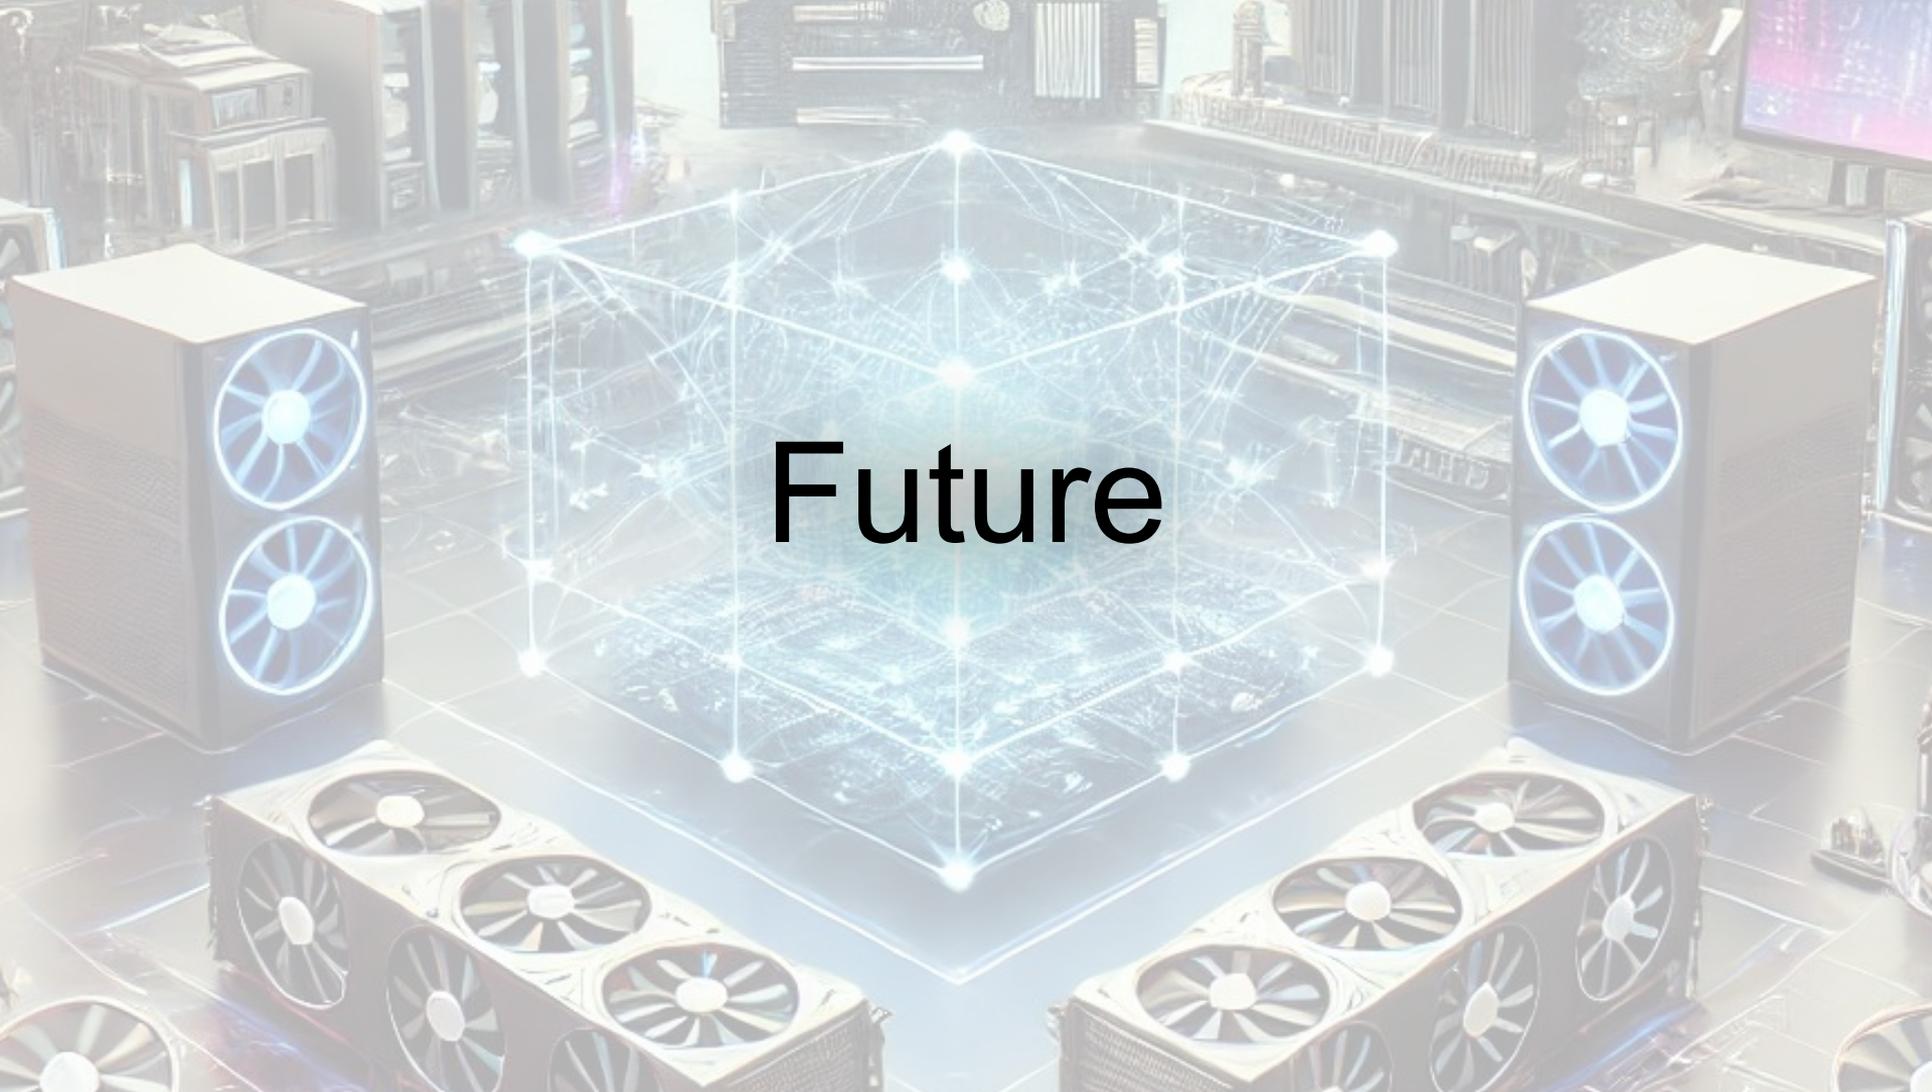
Prevention in silicon happens before product ship from A0 to shipping systems.

Some issues are best found in post-silicon.

Post-silicon issues are particularly difficult.

**Learning** from issues on last generation hardware is critically important.



A futuristic server room with a glowing blue cube of data points in the center, surrounded by server racks and cooling fans. The scene is dimly lit with a blue and purple color palette. The word "Future" is centered in the image.

Future

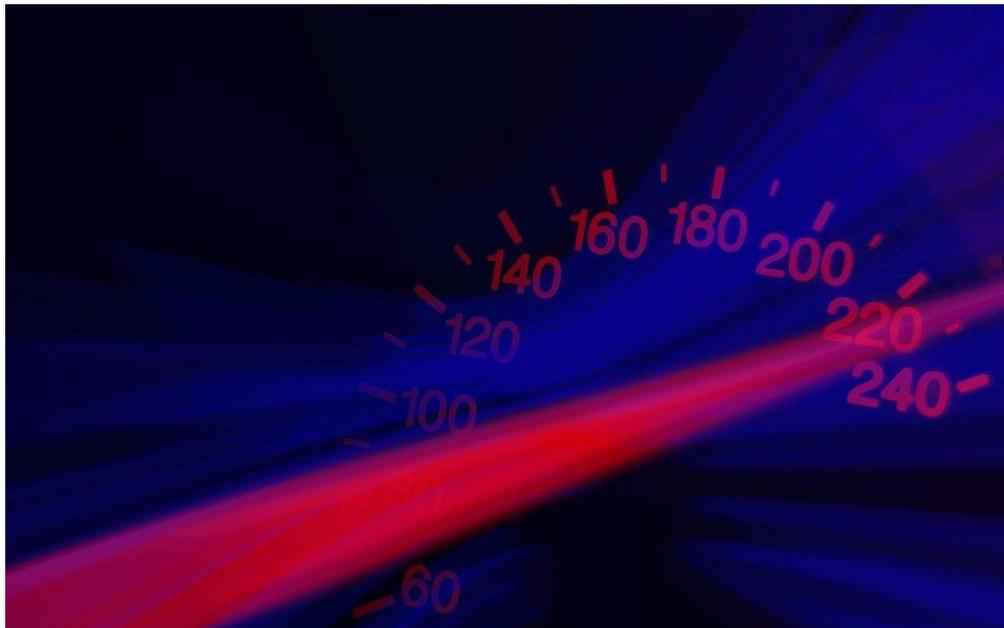
# Future of uArch security is future of uArch

Silicon performance is the main  
underlying driver for growth in compute  
ecosystem

## Performance comes from 3 sources

- New process technology
- uArch improvements
- Adaptation to changed  
workloads

uArch improvements & Changed  
workloads will lead to new security  
challenges



# uArch security future

## Offense

New kinds of prediction & data dependent behaviors (memory latency!). Memory is order of magnitude slower than compute. Some examples:

- New kinds of caches and bigger caches
- Work load specific prefetchers
- Different kinds of value prediction
- Cache & memory compression
- Growth in reorder buffer sizes
- New exploitation techniques

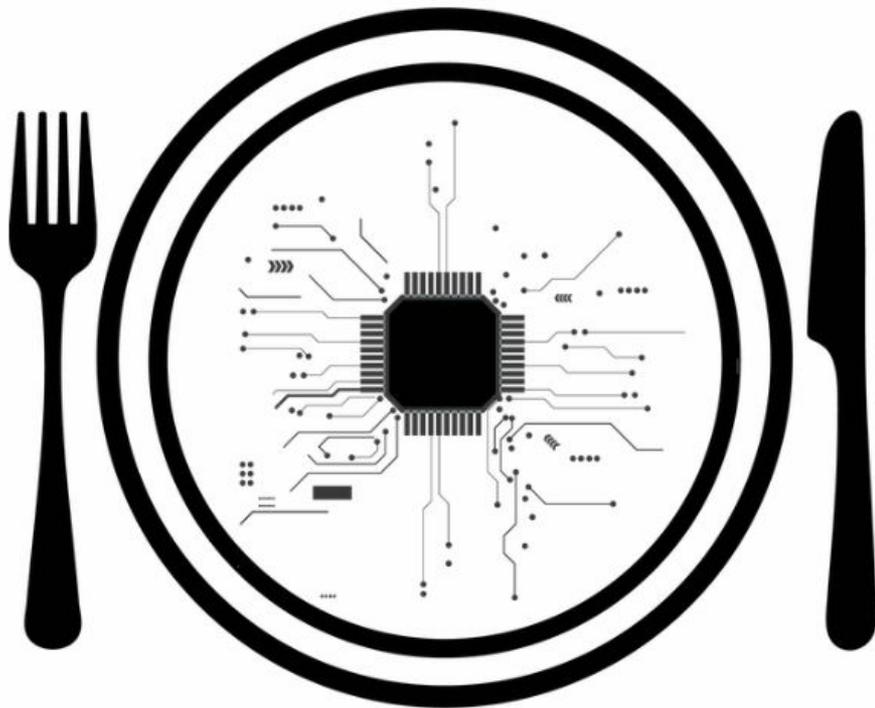
## Defense

- Increased maturity
  - Better tooling
  - More defense in depth
- New microarchitecture security features
- More configurability of security
  - Ex.PSF switch on AMD
- Improved support for software influence
  - Ex. Local configuration switches

# New kinds of compute

more heterogeneous - but all have uArch:

- **GPU (new use cases)**
  - Remote accessible
  - Increased complexity and new work loads
  - Example: “LeftoverLocals” by Trails of Bits
- **Neural Processing Units**
  - New model of compute
  - New threats: Integrity of models
  - Attack vector against system
- **AI training accelerators in the cloud**
  - Soon: shared resources + multi tenant
- **More generally:** More kinds of compute, more accelerators



# Defensive side of things

Huge gap between academia and industry:

## Academia

- provable Rowhammer mitigations available
- provable secure cache available

## Industry

- probabilistic Rowhammer mitigations
- secure caches not adopted (but non-inclusive LLCs)

# SNOWMEN



# uArch in uArch

Embedded processors everywhere --  
already with speculation:

Speculation vs confidentiality?

- Threat models rarely contain arbitrary execution  
→ constrains attackers
- Embedded processors often provide low-level access → new and different kinds of assets



# Take Aways

Side channels are **here to stay**

- Side channels **can be managed**

more aspects of microarchitecture and different kinds of issues

- Hard work for both offensive research and defense
- Defense is maturing

Microarchitecture is a **growth area**, so is microarchitecture security

Microarchitecture matters, so does microarchitecture security

# Microarchitecture Vulnerabilities

Past, Present and Future

Daniel Gruss (Graz University of Technology)  
Anders Fogh (Intel Corporation)