



CPU Fuzzing: Automated Discovery of Microarchitectural Attacks

15 - 17 NOVEMBER 2022
RIYADH FRONT EXHIBITION CENTRE
SAUDI ARABIA

Daniel Weber, Michael Schwarz, Moritz Lipp

CO-ORGANISED BY:



الاتحاد السعودي للأمن
السيبراني والتكنولوجيا والذكاء
SAUDI FEDERATION FOR CYBERSECURITY
PROGRAMMING & DRONES

IN PARTNERSHIP WITH:



الهيئة العامة للرياضة
Saudi Federation for Cybersecurity



DEVELOPING STORY

COMPUTER CHIP FLAWS IMPACT BILLIONS OF DEVICES

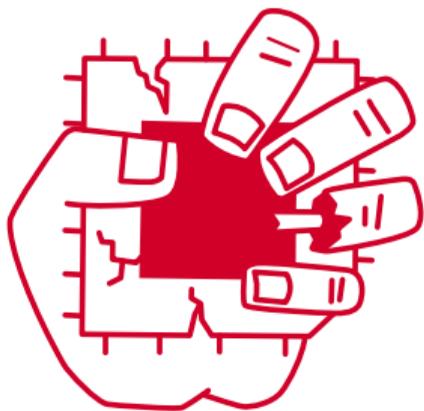
LIVE



DAX ▲ 164.69

NEWS STREAM







Microarchitectural attacks



Microarchitectural attacks



How do they **work**?



Microarchitectural attacks



How do they **work**?



How can we
find them **efficiently**?



Daniel Weber

PhD Student @ CISPA Helmholtz Center for Information Security

Research on CPU Security and Side Channels

 @weber_daniel

 daniel.weber@cispa.de



Michael Schwarz

Faculty @ CISA Helmholtz Center for Information Security

Research on CPU Security and Side Channels

🐦 @misc0110

✉ michael.schwarz@cispa.de



Moritz Lipp

Security Researcher @ Amazon Web Services

Research on CPU Security and Side Channels

 @mlqxyz

 mlipp@amazon.com

Osiris Research Team

Daniel Weber, *Ahmad Ibrahim, Hamed Nemati*, **Michael Schwarz**,
Christian Rossow

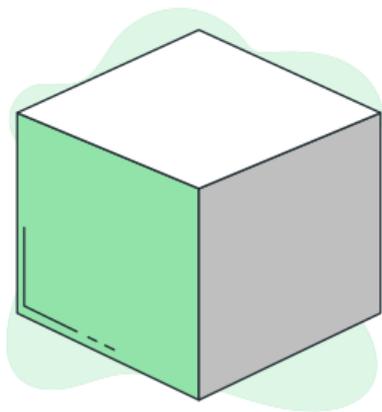


Transynther Research Team

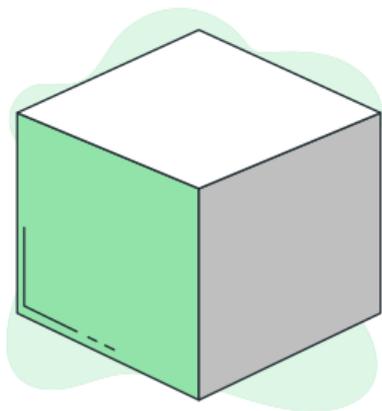
Daniel Moghimi, **Moritz Lipp**, *Berk Sunar*, **Michael Schwarz**

MSRevelio Research Team

Andreas Kogler, **Daniel Weber**, *Martin Haubenwallner*, **Moritz Lipp**,
Daniel Gruss, **Michael Schwarz**



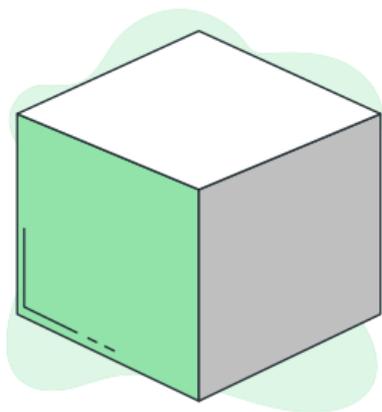
CPU



CPU



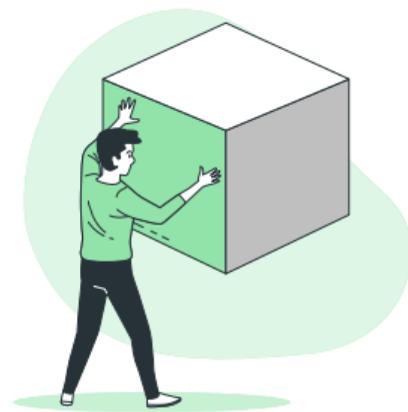
Specification (ISA)



CPU



Specification (ISA)



Interaction



Everything **works**
as expected



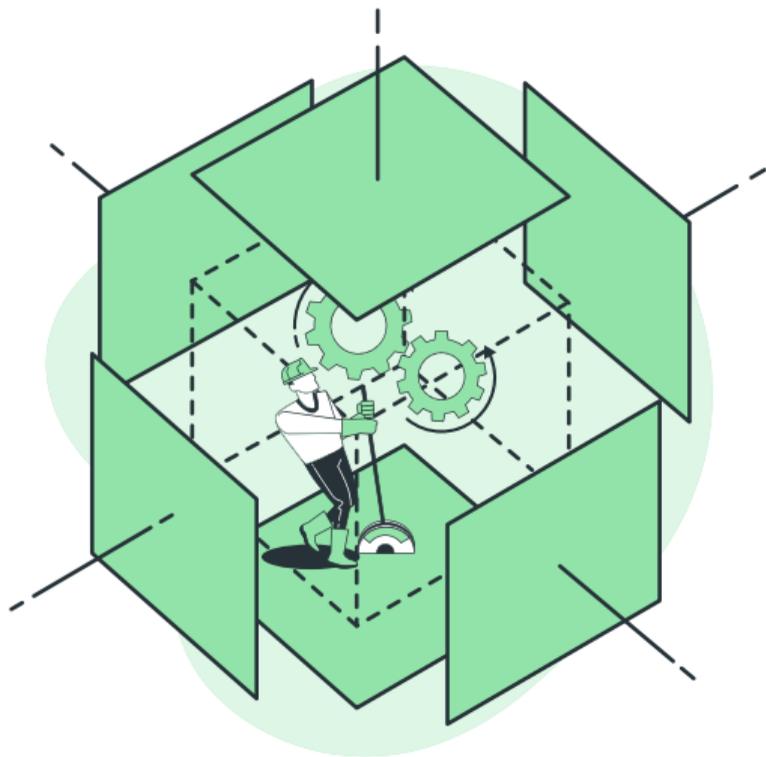
Everything **works**
as expected



No bugs



How can you attack this?



Information **leaks** through **side effects**



Power Consumption



Power Consumption



Temperature



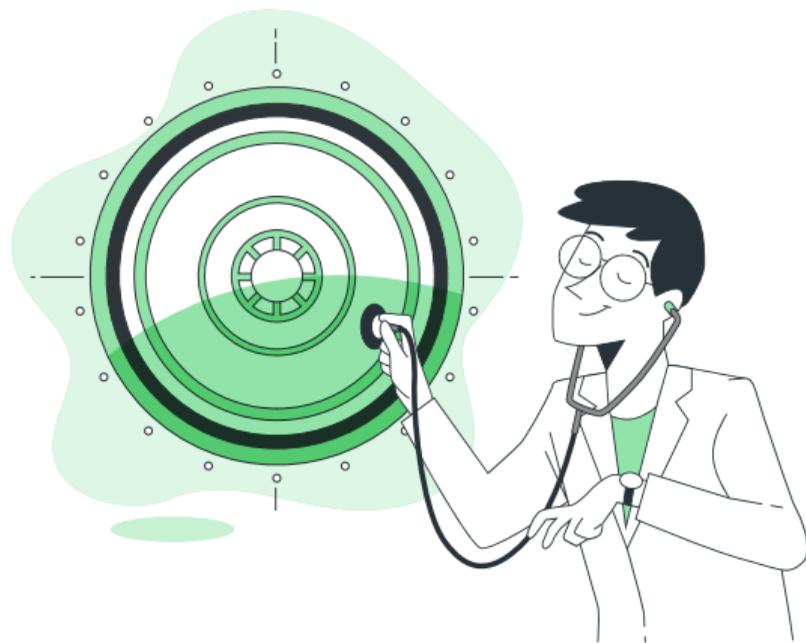
Power Consumption



Temperature



Execution Time



Observe Device

$\{x_n\} + \{y_n\} \stackrel{\text{df}}{=} \{x_n + y_n\}; \quad \|\{x_n\}\|_{CR} \downarrow n \rightarrow \infty$
 $\downarrow n \rightarrow \infty; \quad y_n \quad |g| < 1; \quad x: p \quad \sqrt[4]{4} \cdot \sqrt[3]{13} \cdot n^2$

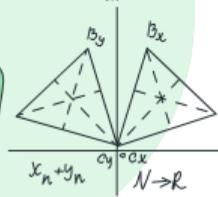
$x: p \quad \lim_{n \rightarrow \infty} \sqrt[n]{A} = 1$

$N \rightarrow \mathbb{R} \quad n \geq n_0: (x_n - g) < \varepsilon$

$\sqrt[4]{4^n - \cos 2n} \left(\frac{n^2 + n - 1}{n^2 - 2n + 3} \right)^5$
 $n \geq n_0: (x_n)$

$N \rightarrow \mathbb{R} \quad n \geq n_0: (x_n - g) < \varepsilon$

$\{x_n\} + \{y_n\} \stackrel{\text{df}}{=} \{x_n + y_n\}$







Evaluate data



Get the secrets. **Easy as that!**



Reality is **not** as easy



Software-only

No Physical Access required



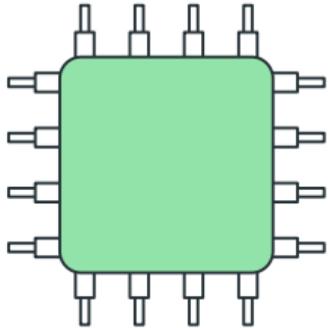
Software-only
No Physical Access required



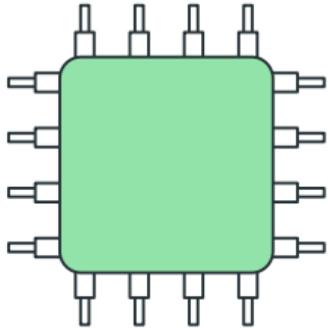
Understand Inner Workings
of the CPU



What can we observe? How does a CPU work internally?

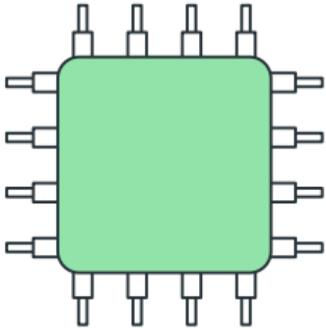


- Instruction Set Architecture (ISA) is an **abstract model** of a computer (x86, ARMv8, SPARC, ...)



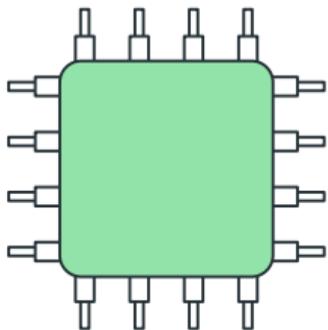
- Instruction Set Architecture (ISA) is an **abstract model** of a computer (x86, ARMv8, SPARC, ...)
- **Interface** between hardware and software

Architecture vs Microarchitecture



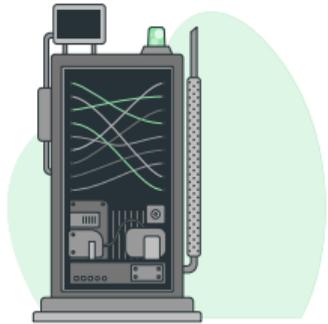
- Instruction Set Architecture (ISA) is an **abstract model** of a computer (x86, ARMv8, SPARC, ...)
- **Interface** between hardware and software
- Microarchitecture is an ISA **implementation**

Architecture vs Microarchitecture

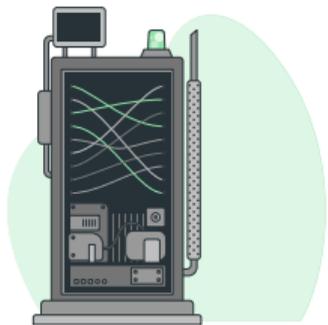


- Instruction Set Architecture (ISA) is an **abstract model** of a computer (x86, ARMv8, SPARC, ...)
- **Interface** between hardware and software
- Microarchitecture is an ISA **implementation**

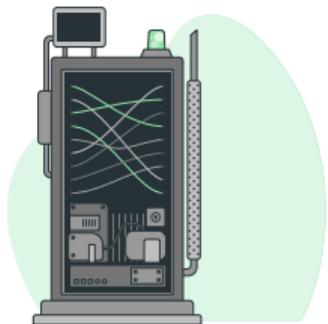




- Modern CPUs contain multiple **microarchitectural elements**



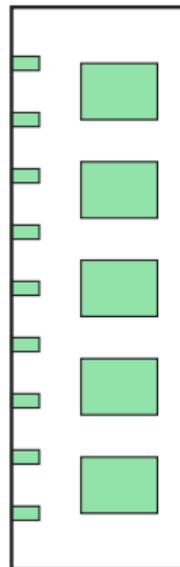
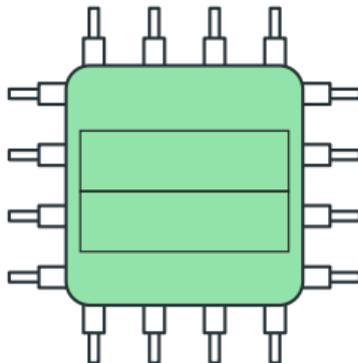
- Modern CPUs contain multiple **microarchitectural elements**
- **Transparent** for the programmer



- Modern CPUs contain multiple **microarchitectural elements**
- **Transparent** for the programmer
- **Optimize** for performance, power consumption, ...

CPU Optimization: Cache

```
printf("%d", i);  
printf("%d", i);
```

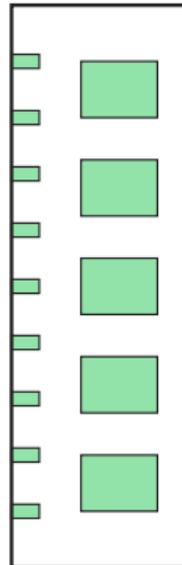
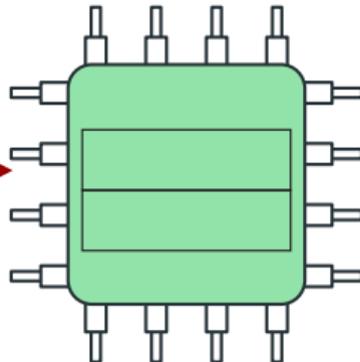


CPU Optimization: Cache

```
printf("%d", i);
```

```
printf("%d", i);
```

Cache miss

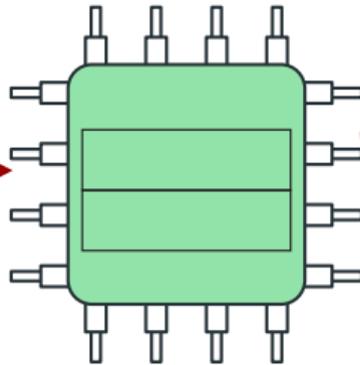


CPU Optimization: Cache

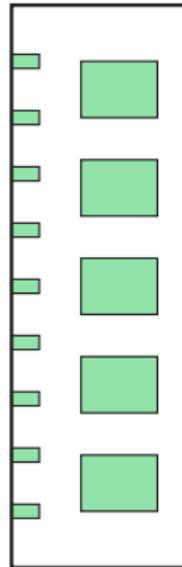
```
printf("%d", i);
```

```
printf("%d", i);
```

Cache miss



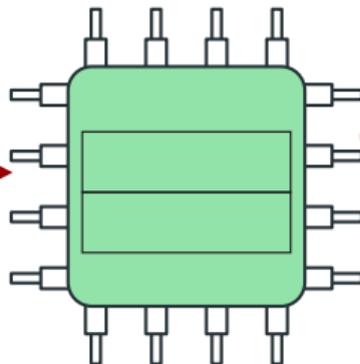
Request



CPU Optimization: Cache

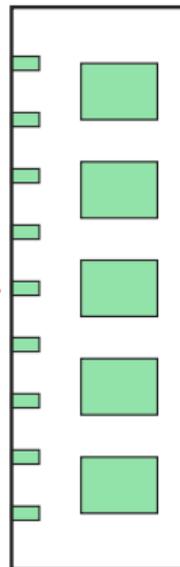
```
printf("%d", i);  
printf("%d", i);
```

Cache miss



Request

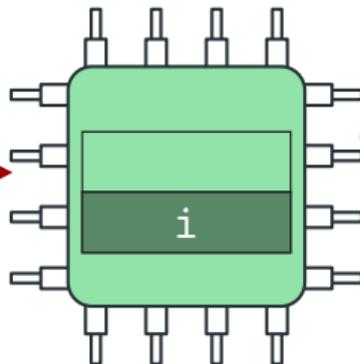
Response



CPU Optimization: Cache

```
printf("%d", i);  
printf("%d", i);
```

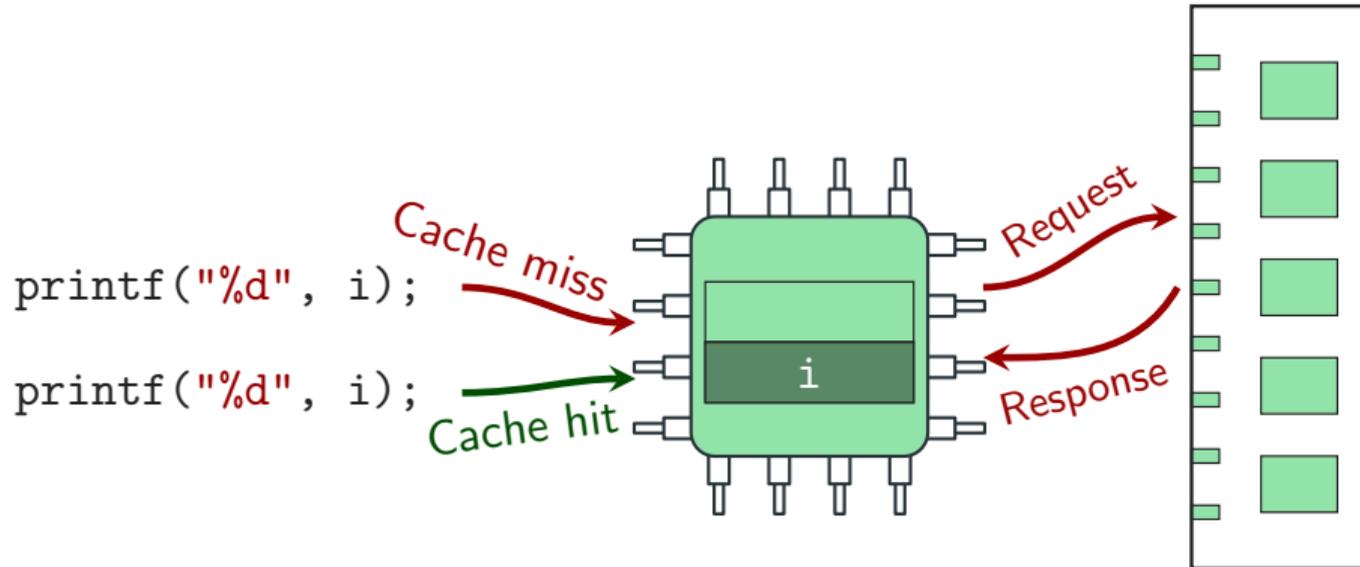
Cache miss



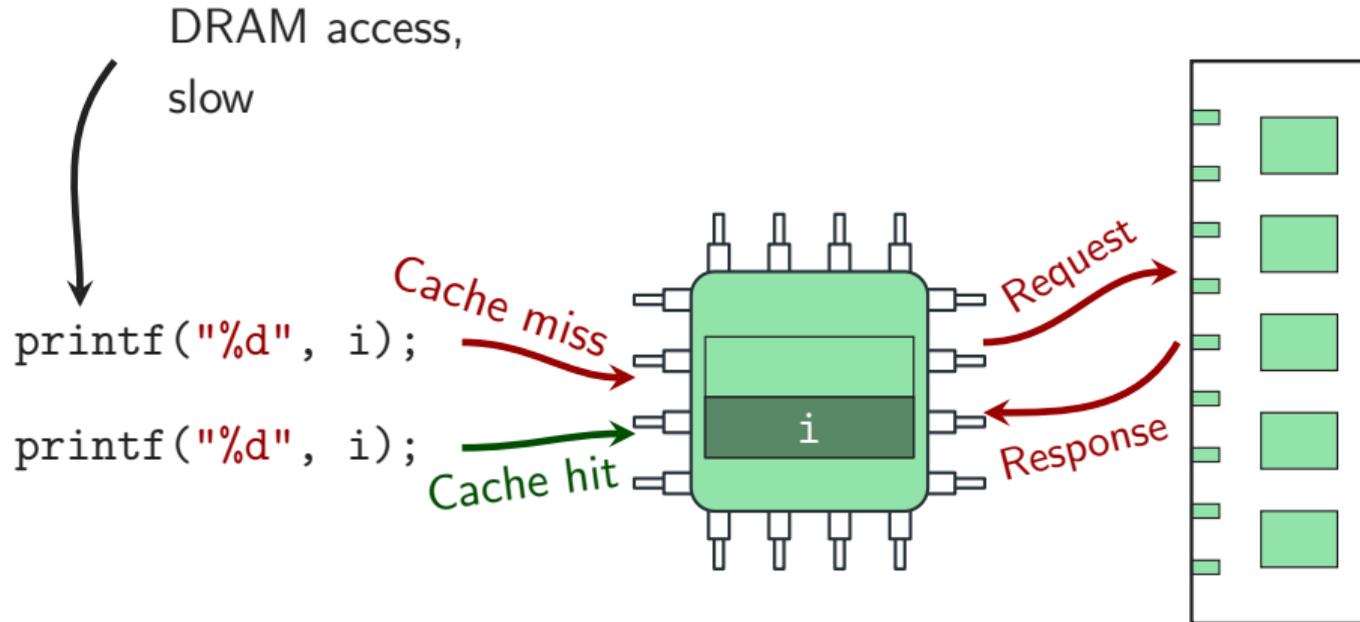
Request

Response

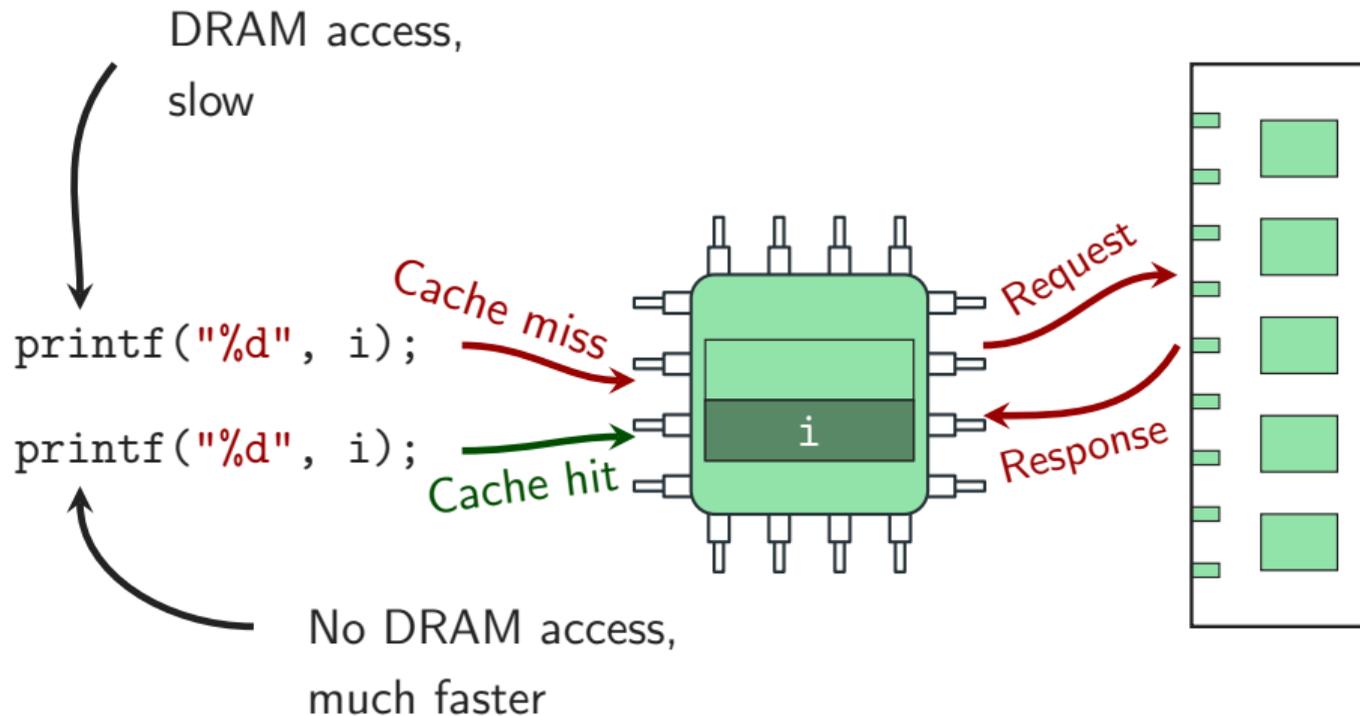
CPU Optimization: Cache



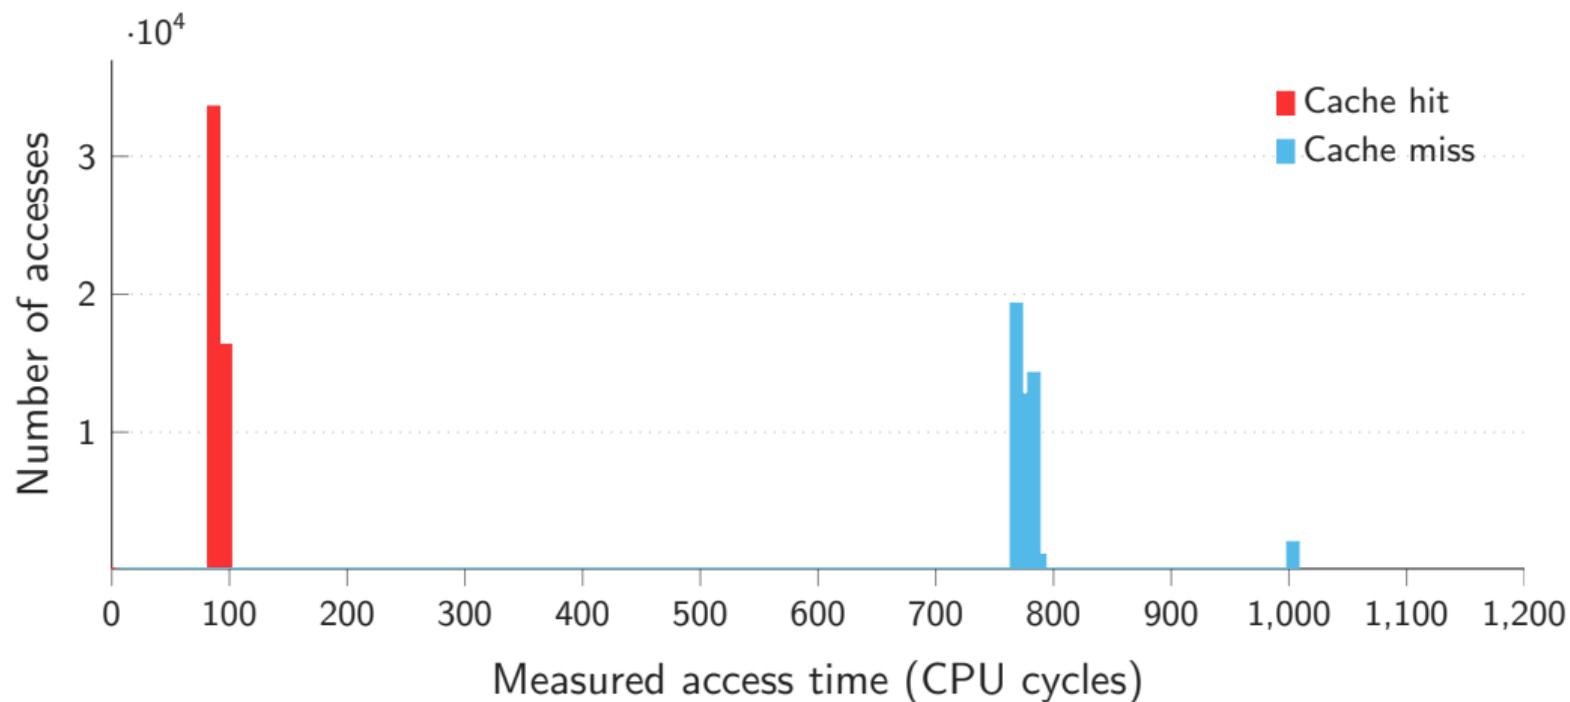
CPU Optimization: Cache



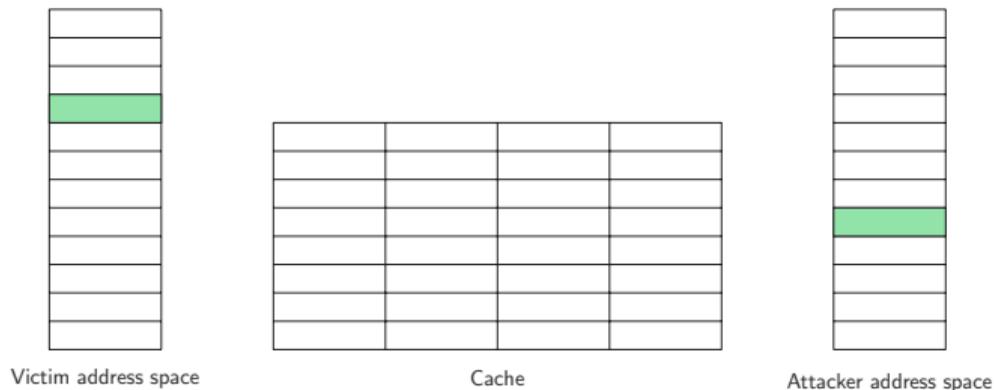
CPU Optimization: Cache



Caching speeds up Memory Accesses

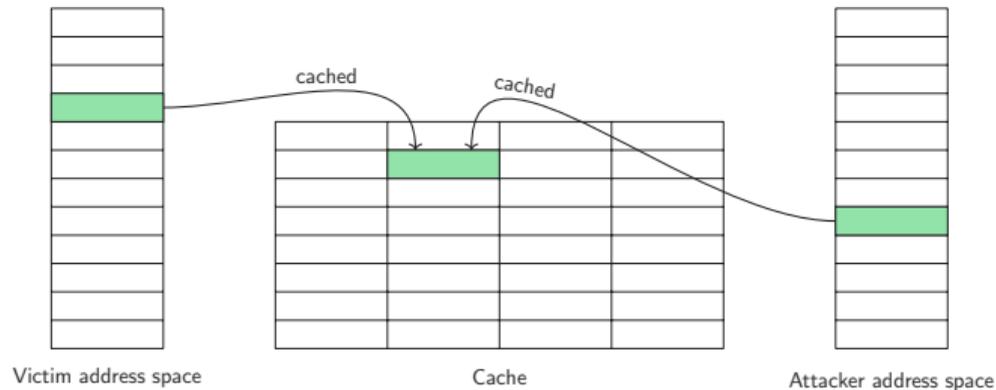


Flush+Reload



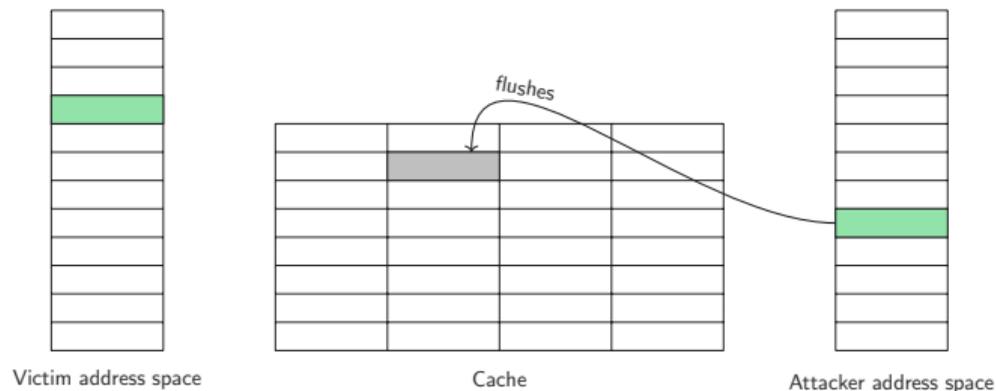
Step 1: Attacker maps shared library (shared memory, in cache)

Flush+Reload



Step 1: Attacker maps shared library (shared memory, in cache)

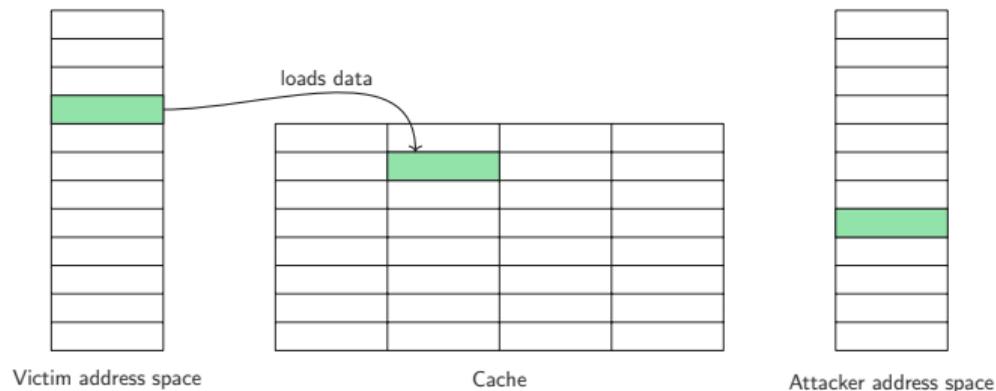
Flush+Reload



Step 1: Attacker maps shared library (shared memory, in cache)

Step 2: Attacker **flushes** the shared cache line

Flush+Reload

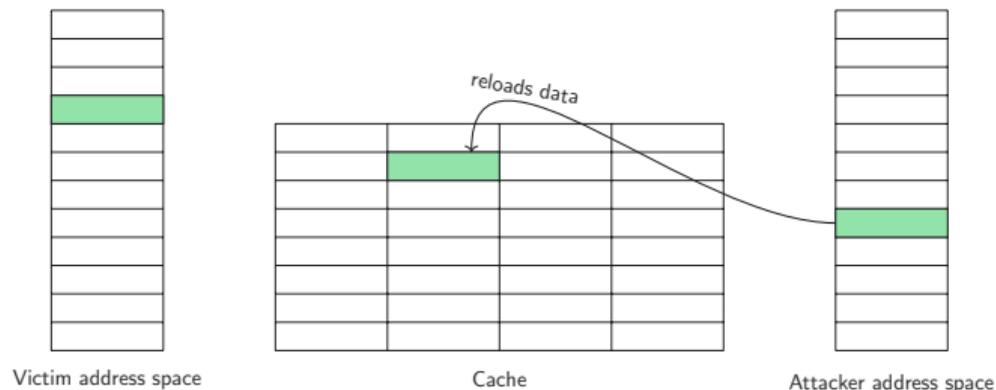


Step 1: Attacker maps shared library (shared memory, in cache)

Step 2: Attacker **flushes** the shared cache line

Step 3: Victim loads the data

Flush+Reload



Step 1: Attacker maps shared library (shared memory, in cache)

Step 2: Attacker **flushes** the shared cache line

Step 3: Victim loads the data

Step 4: Attacker measures the access time to **reload** the data



- Leak **cryptographic keys**



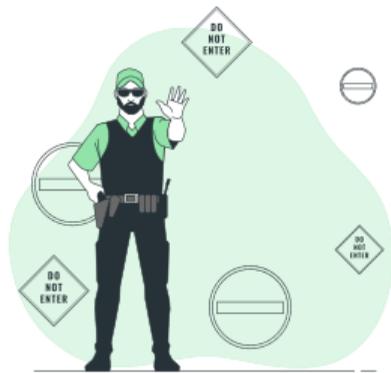
- Leak **cryptographic keys**
- Leak information on **co-located virtual machines**



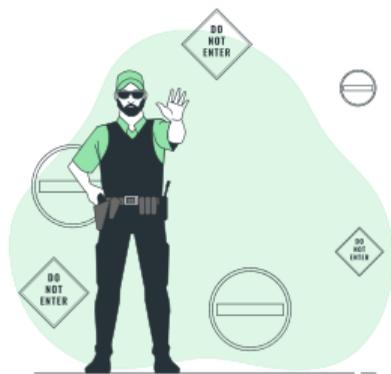
- Leak **cryptographic keys**
- Leak information on **co-located virtual machines**
- **Monitor** function calls of **other applications**



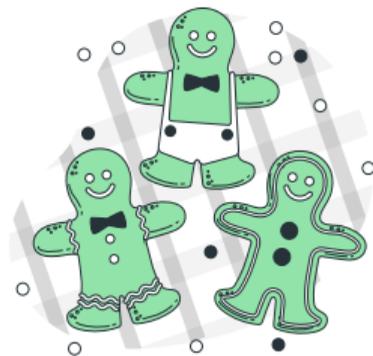
- Leak **cryptographic keys**
- Leak information on **co-located virtual machines**
- **Monitor** function calls of **other applications**
- Build **covert communication channels**
- ...



Leakage of **1** Instruction



Leakage of **1** Instruction



There are **>3000** instructions



How do you find **these** attacks?



Read documentation and
patents



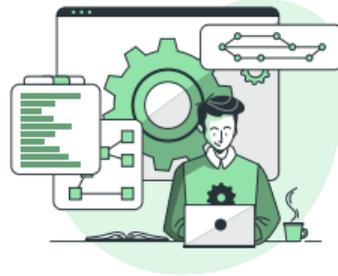
Read documentation and patents



Develop ideas



Read documentation and patents



Develop ideas



Test ideas and start over



New CPU, New Features



New CPU, New Features



Repeat from the beginning



Finding side channels and vulnerabilities is a **complex** and **time-consuming** process



Can we do **better**?



- How can you find bugs in **software** in an automated way?



- How can you find bugs in **software** in an automated way?
- **Fuzzing** is an automated software testing method
 - Inject invalid, malformed, unexpected input to a program



- How can you find bugs in **software** in an automated way?
- **Fuzzing** is an automated software testing method
 - Inject invalid, malformed, unexpected input to a program
- Can we **apply a similar approach** do find microarchitectural side-channels?

Challenges when Fuzzing CPUs



Limited Feedback

Challenges when Fuzzing CPUs



Limited Feedback



Dirty States

Challenges when Fuzzing CPUs



Limited Feedback



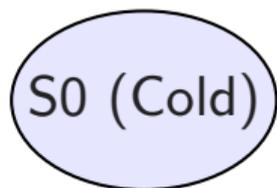
Dirty States

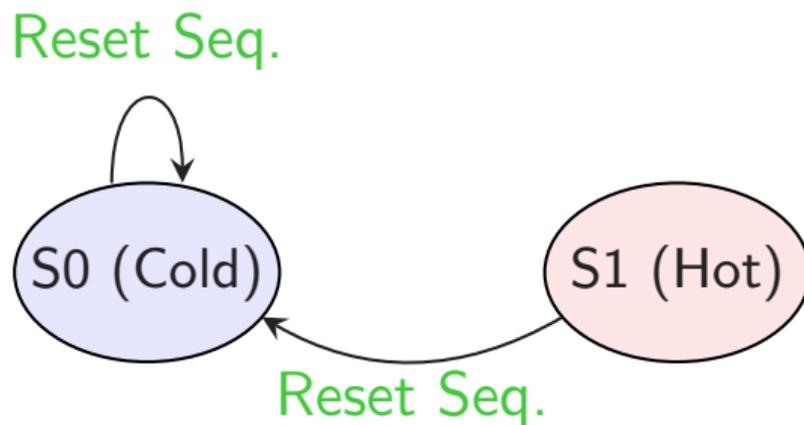


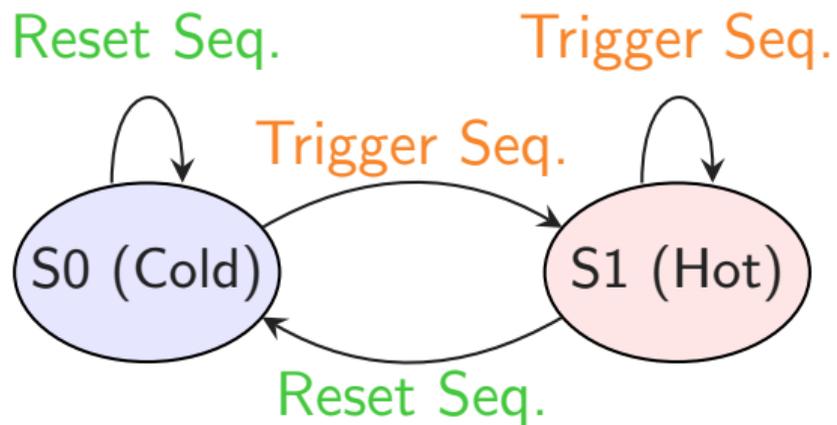
System Noise

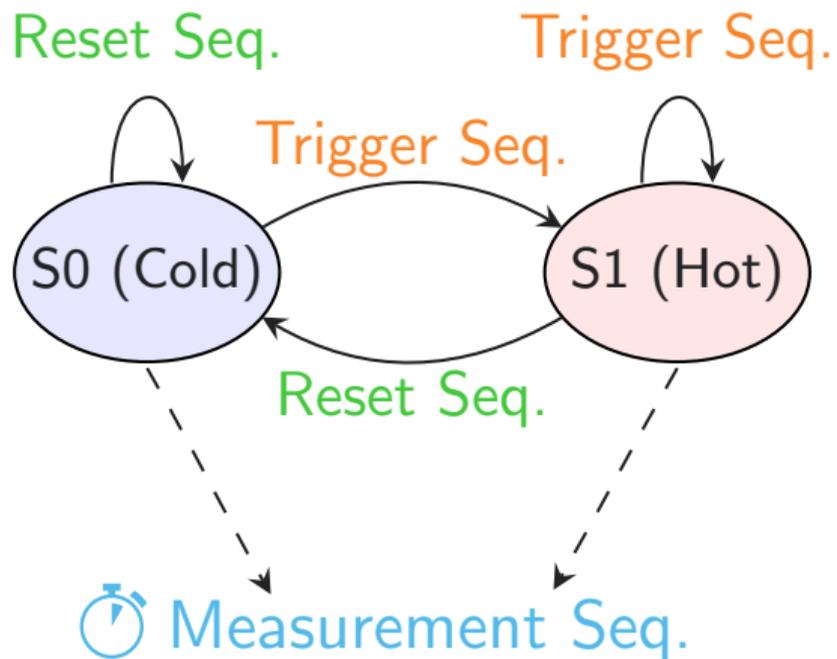


Can we **automatically decide** whether we **found a side channel**?





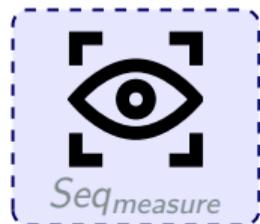




Testing Potential Side Channels



Testing Potential Side Channels



Testing Potential Side Channels



Seqreset



Seqmeasure



Cold path S0

Testing Potential Side Channels



Seq_reset



Seq_measure

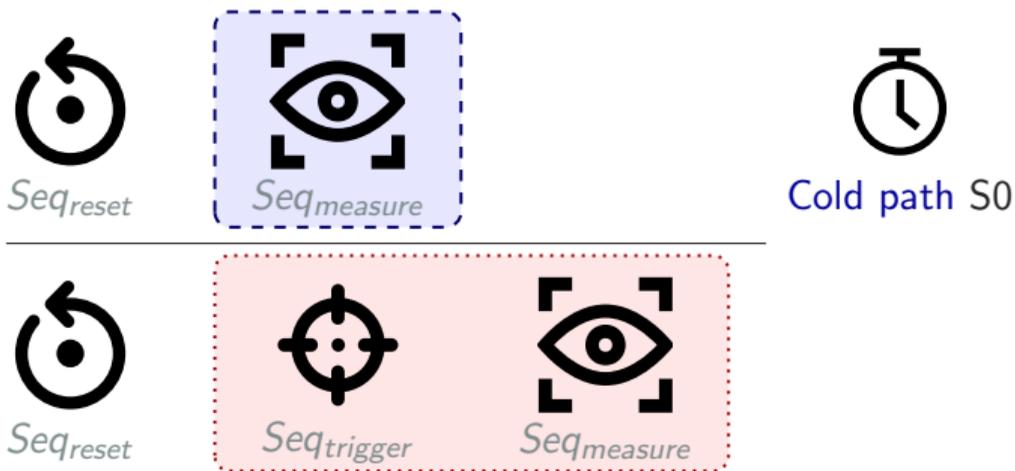


Cold path S0

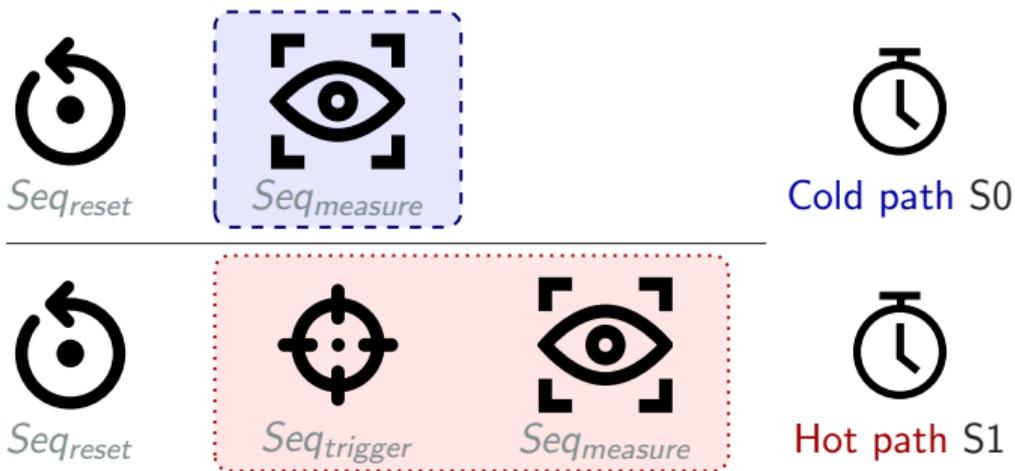


Seq_reset

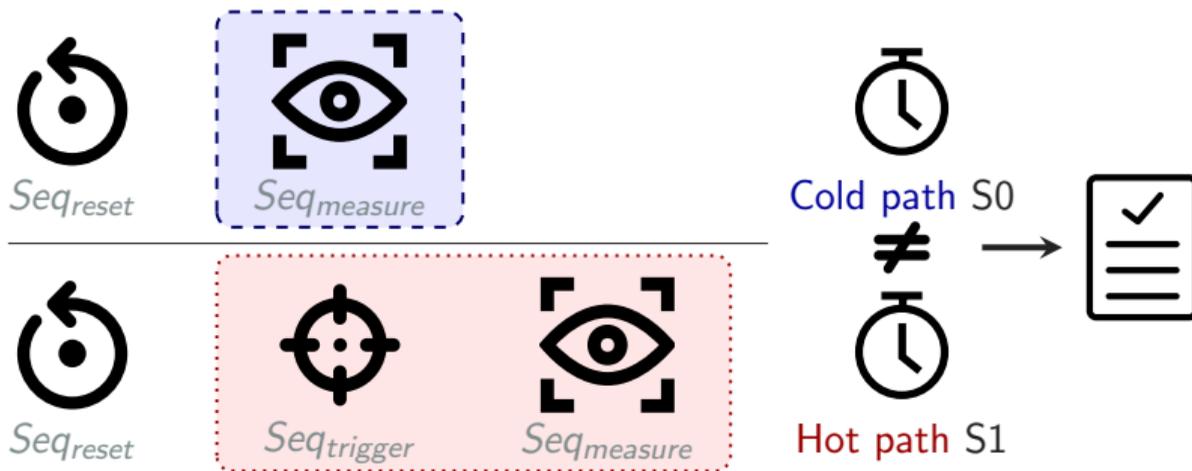
Testing Potential Side Channels



Testing Potential Side Channels

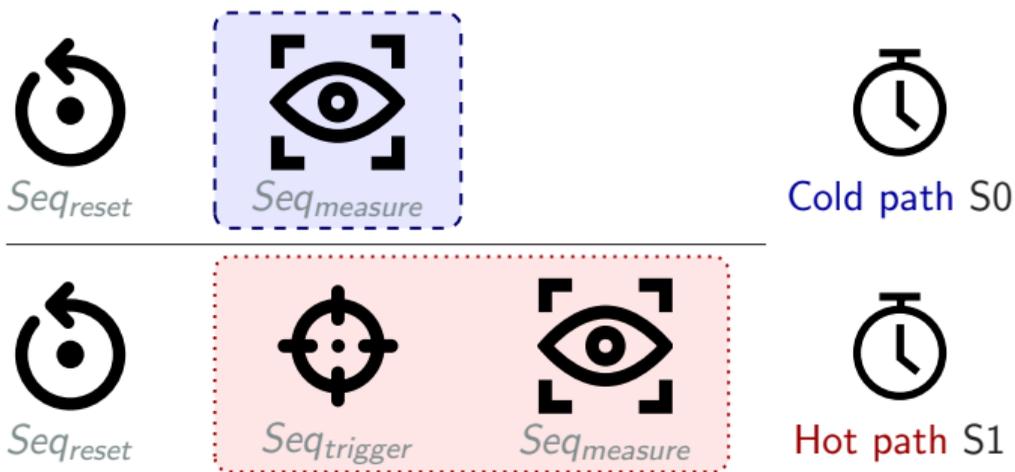


Testing Potential Side Channels



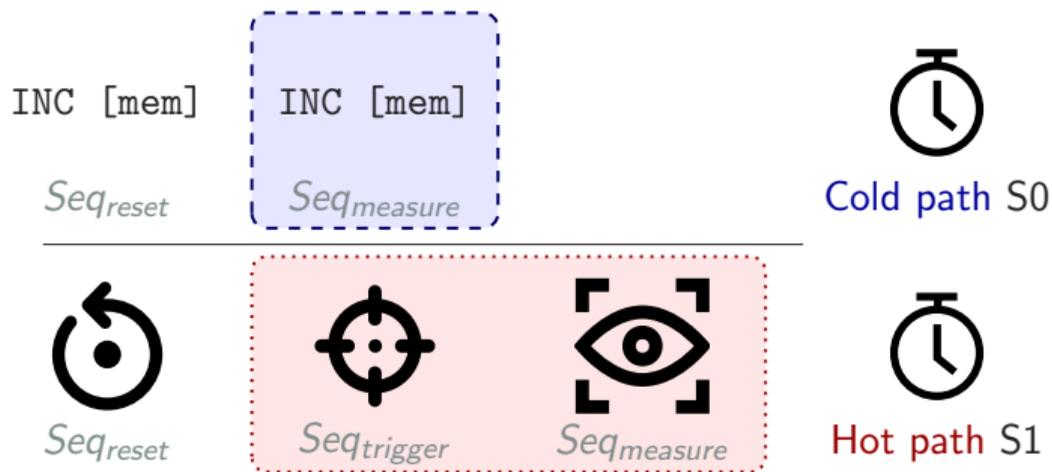
Testing Potential Side Channels

Example 1: $Seq_{measure} = Seq_{trigger} = Seq_{reset} = INC$ [mem]



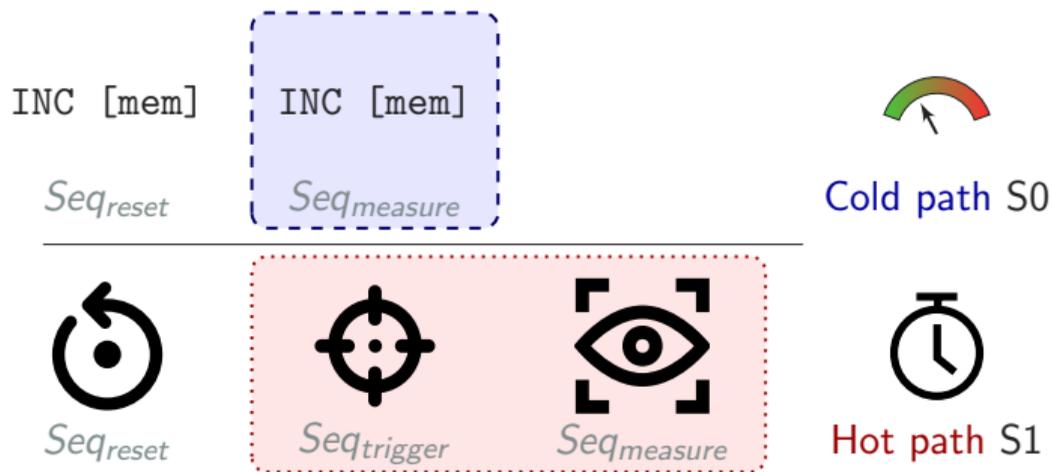
Testing Potential Side Channels

Example 1: $Seq_{measure} = Seq_{trigger} = Seq_{reset} = \text{INC} [\text{mem}]$



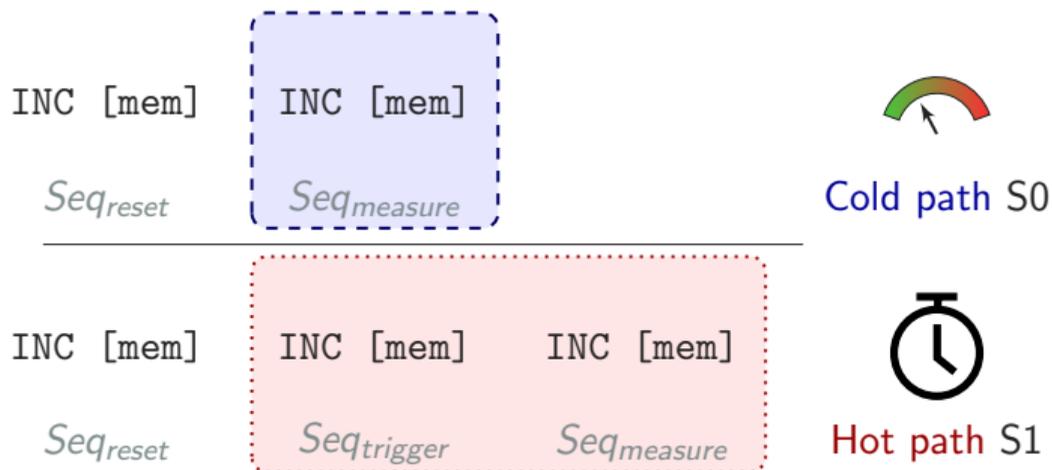
Testing Potential Side Channels

Example 1: $Seq_{measure} = Seq_{trigger} = Seq_{reset} = \text{INC} [\text{mem}]$



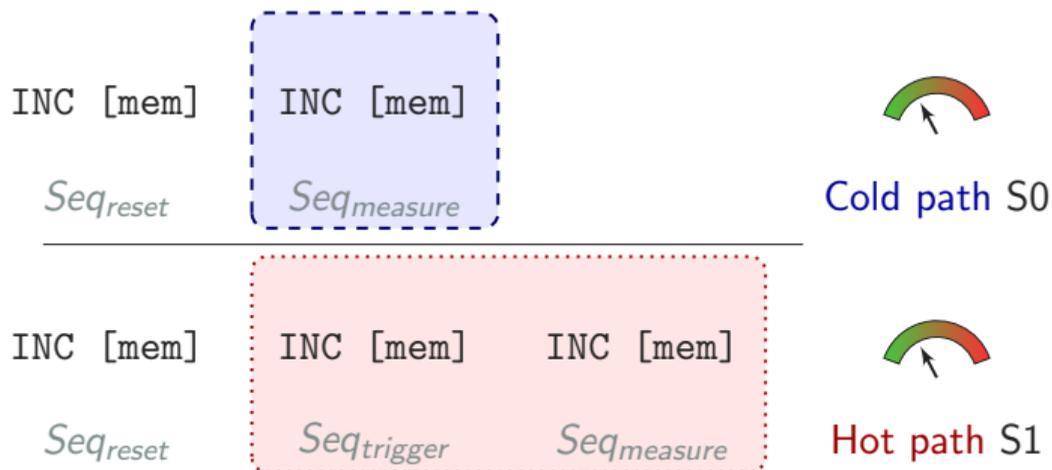
Testing Potential Side Channels

Example 1: $Seq_{measure} = Seq_{trigger} = Seq_{reset} = \text{INC} [\text{mem}]$



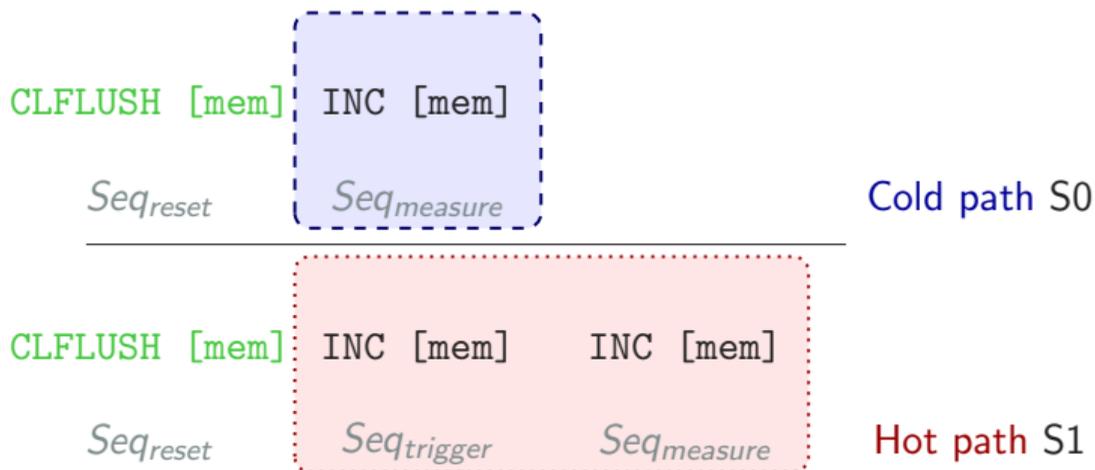
Testing Potential Side Channels

Example 1: $Seq_{measure} = Seq_{trigger} = Seq_{reset} = INC \ [mem]$



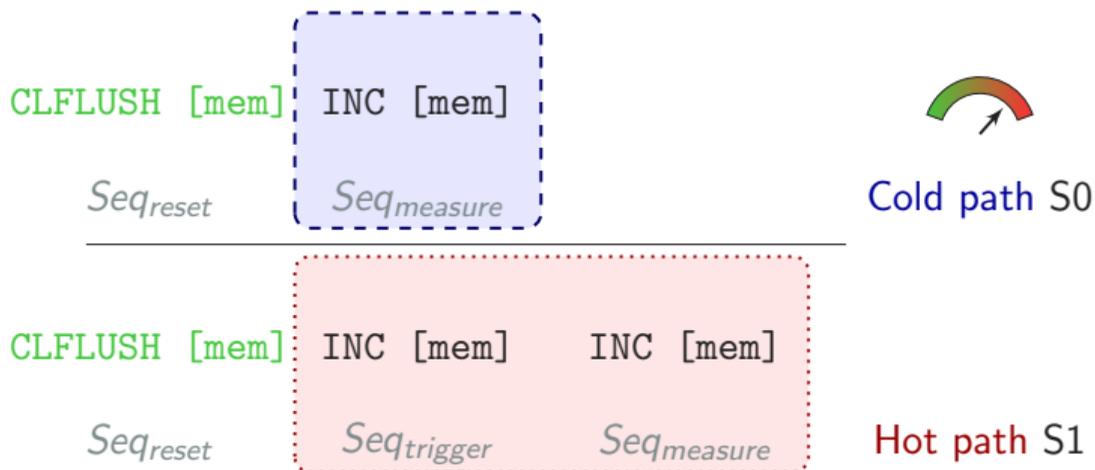
Testing Potential Side Channels

Example 2: $Seq_{measure} = Seq_{trigger} = \text{INC } [\text{mem}]$;
 $Seq_{reset} = \text{CLFLUSH } [\text{mem}]$



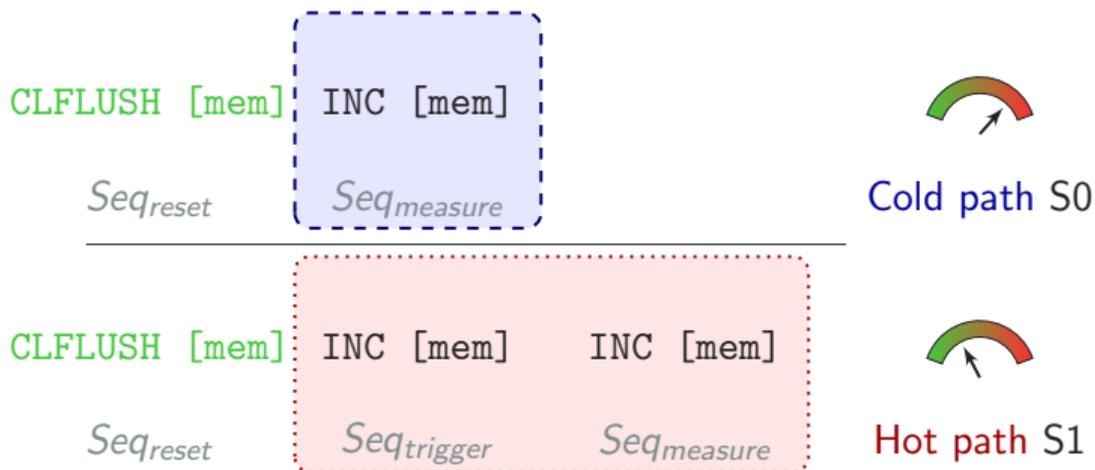
Testing Potential Side Channels

Example 2: $Seq_{measure} = Seq_{trigger} = \text{INC } [\text{mem}]$;
 $Seq_{reset} = \text{CLFLUSH } [\text{mem}]$



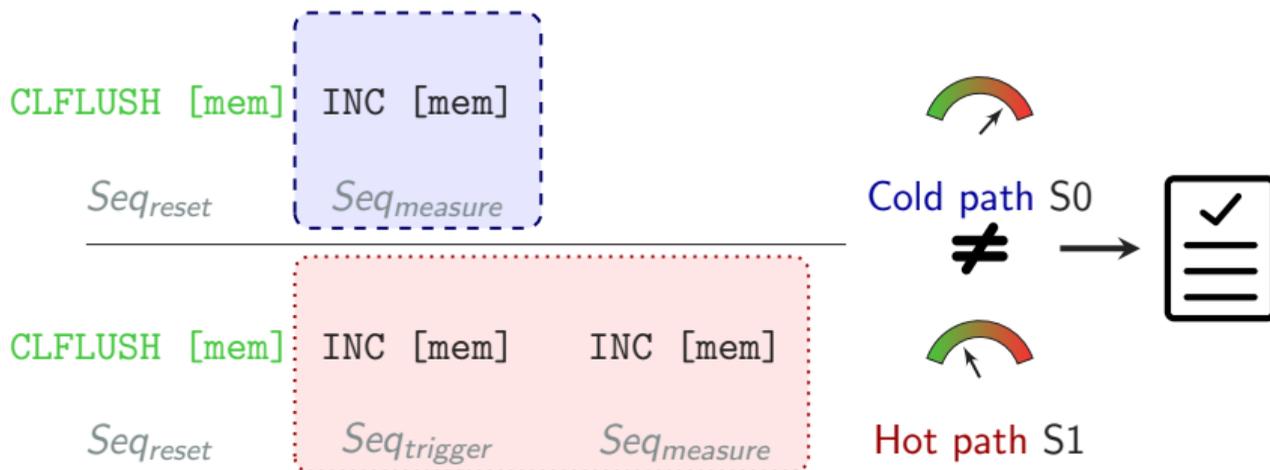
Testing Potential Side Channels

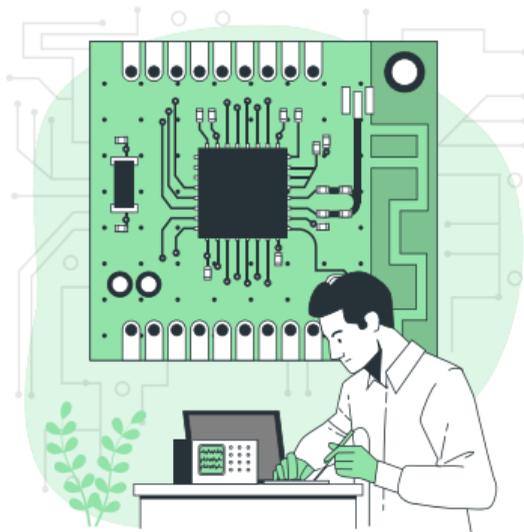
Example 2: $Seq_{measure} = Seq_{trigger} = \text{INC} [\text{mem}]$;
 $Seq_{reset} = \text{CLFLUSH} [\text{mem}]$



Testing Potential Side Channels

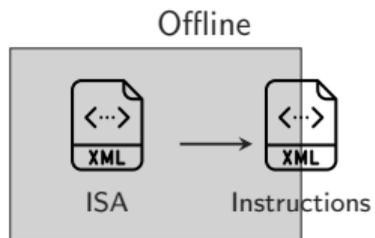
Example 2: $Seq_{measure} = Seq_{trigger} = \text{INC } [\text{mem}]$;
 $Seq_{reset} = \text{CLFLUSH } [\text{mem}]$



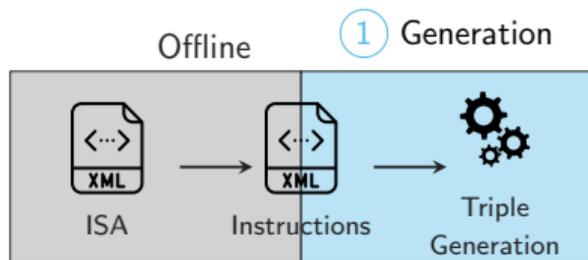


Let's fuzz for Side Channels

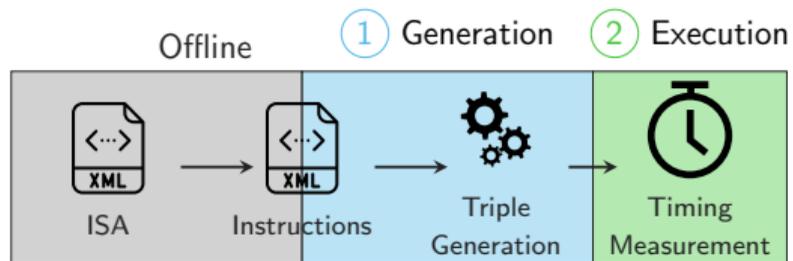
Osiris – Fuzzing x86 CPUs for Side Channels



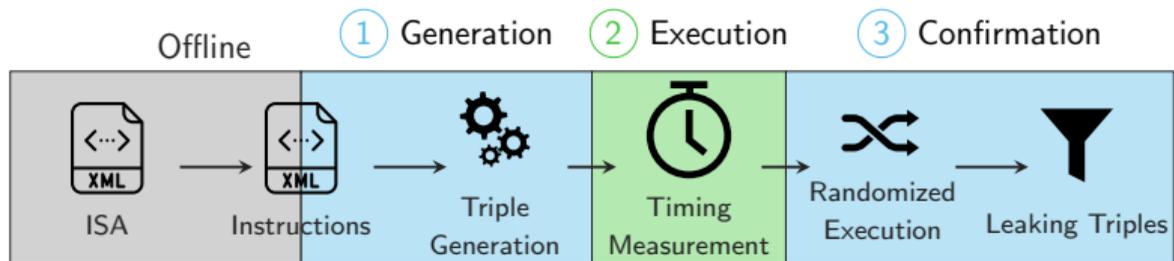
Osiris – Fuzzing x86 CPUs for Side Channels



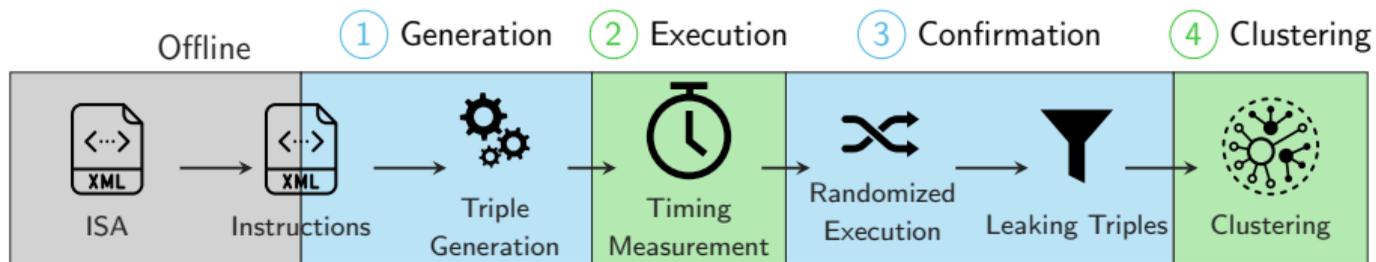
Osiris – Fuzzing x86 CPUs for Side Channels



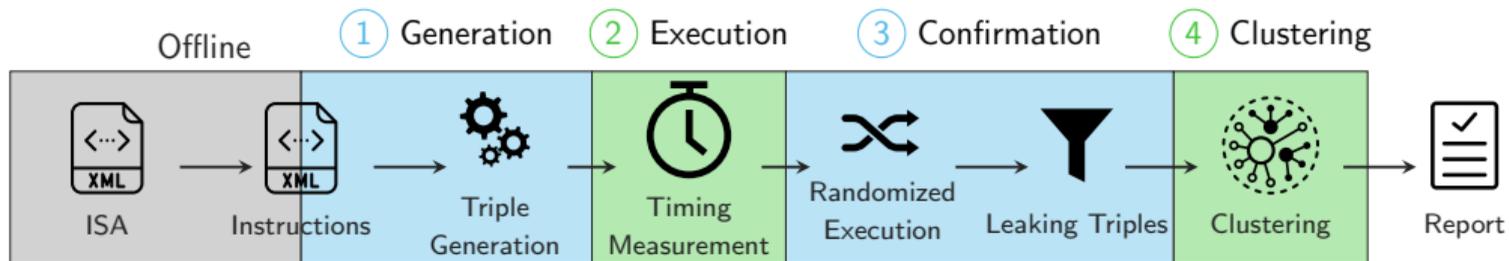
Osiris – Fuzzing x86 CPUs for Side Channels



Osiris – Fuzzing x86 CPUs for Side Channels



Osiris – Fuzzing x86 CPUs for Side Channels





4 days

Osiris – Results Overview



4 days



2 side channels reproduced

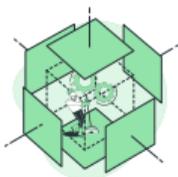
Osiris – Results Overview



4 days



2 side channels reproduced



4 new side channels

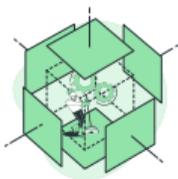
Osiris – Results Overview



4 days



2 side channels reproduced

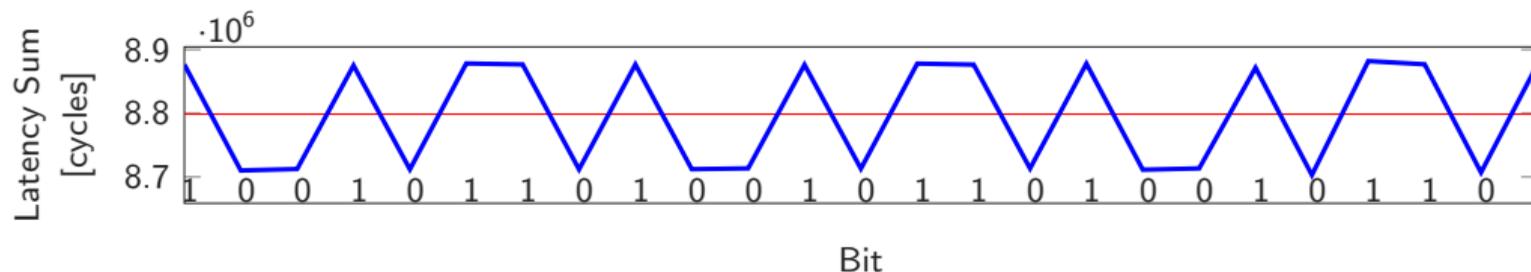


4 new side channels



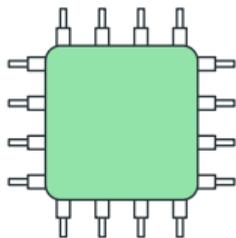
1 new attack

Findings – RDRAND Covert Channel



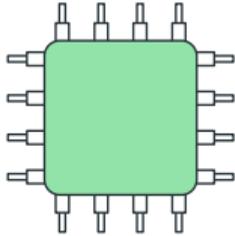
- RDRAND **cross-core** interference
- **Cross-core cross-VM** covert channel

Findings – RDRAND Covert Channel



AMD and Intel

Findings – RDRAND Covert Channel

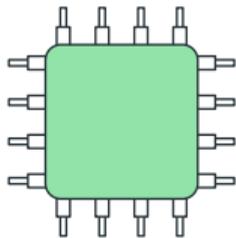


AMD and Intel



VM and native

Findings – RDRAND Covert Channel



AMD and Intel

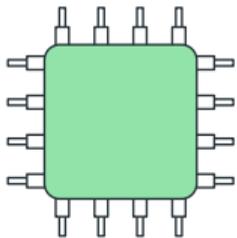


VM and native



1000 bit/s

Findings – RDRAND Covert Channel



AMD and Intel



VM and native

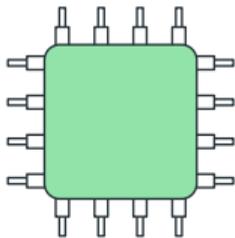


1000 bit/s



No memory

Findings – RDRAND Covert Channel



AMD and Intel



VM and native



1000 bit/s

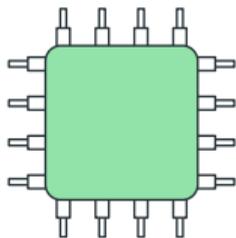


No memory



No detection

Findings – RDRAND Covert Channel



AMD and Intel



VM and native



1000 bit/s



No memory

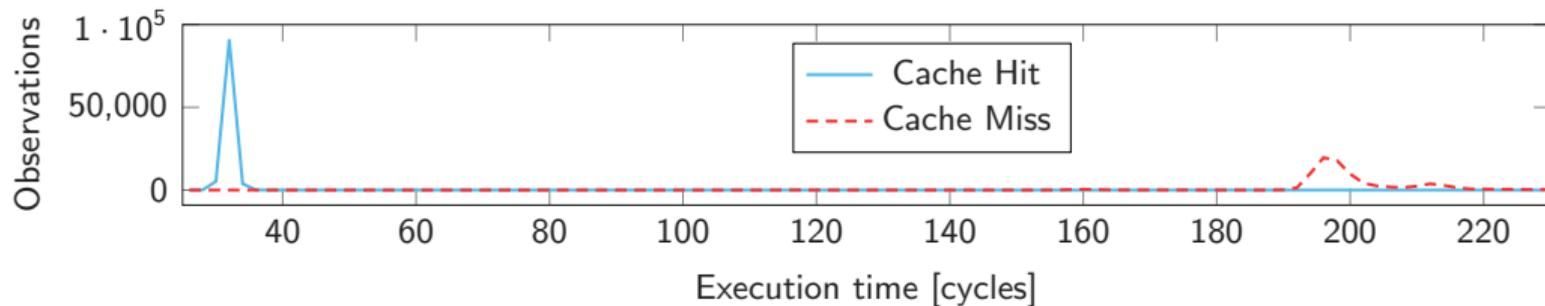


No detection



No mitigation

Findings – MOVNT Side Channel



- MOVNT can **replace** CLFLUSH
 - Flushes data from all cache levels



Faster reload

Findings – MOVNT Side Channel



Faster reload



Stealthy



Faster reload



Stealthy



Not prevented by any cache design



- Many exploits rely on the **knowledge of the memory location** of a certain function



- Many exploits rely on the **knowledge of the memory location** of a certain function
- OS Kernels are **protected by KASLR**

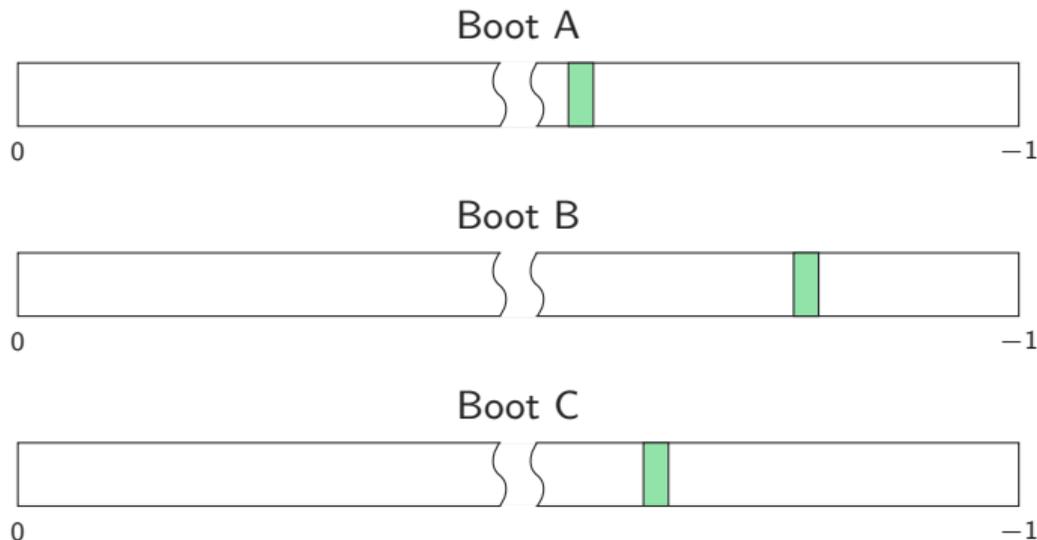


- Many exploits rely on the **knowledge of the memory location** of a certain function
- OS Kernels are **protected by KASLR**
- Memory location of the kernel is **randomized**



- Many exploits rely on the **knowledge of the memory location** of a certain function
- OS Kernels are **protected by KASLR**
- Memory location of the kernel is **randomized**
- Attacker **do not know where** to attack

KASLR: Kernel Address Space Layout Randomization



- Kernel is loaded to a different offset on every boot



- **KASLR break** based on MOVNT and transient execution

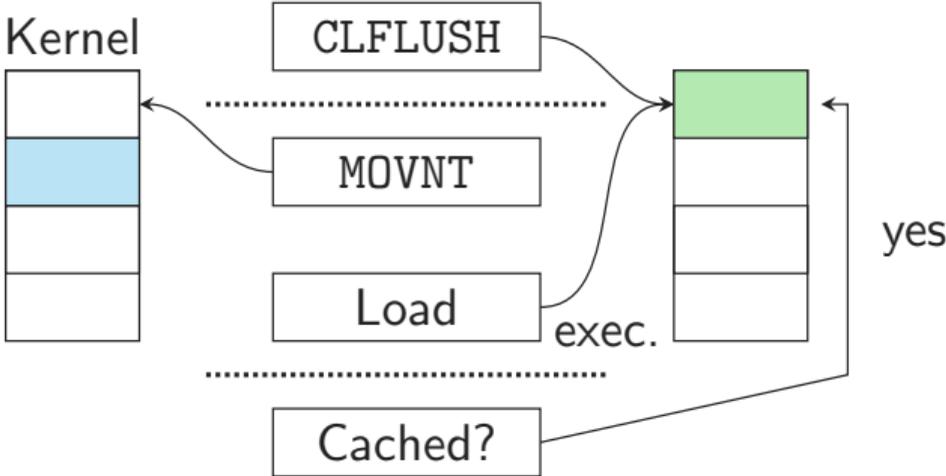


- **KASLR break** based on MOVNT and transient execution
- Works on **new Intel CPUs** (tested up to Tiger Lake)

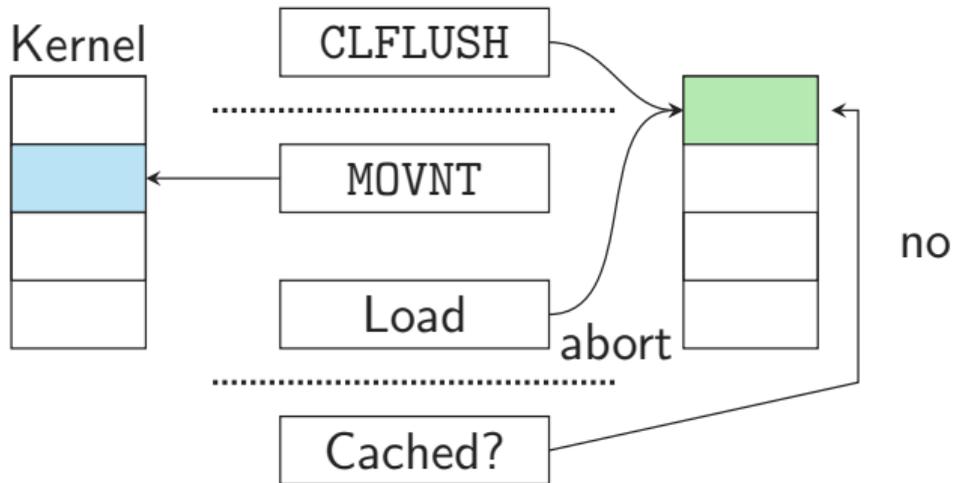


- **KASLR break** based on MOVNT and transient execution
- Works **on new Intel CPUs** (tested up to Tiger Lake)
- Breaks KASLR reliably in **136ms (average)**

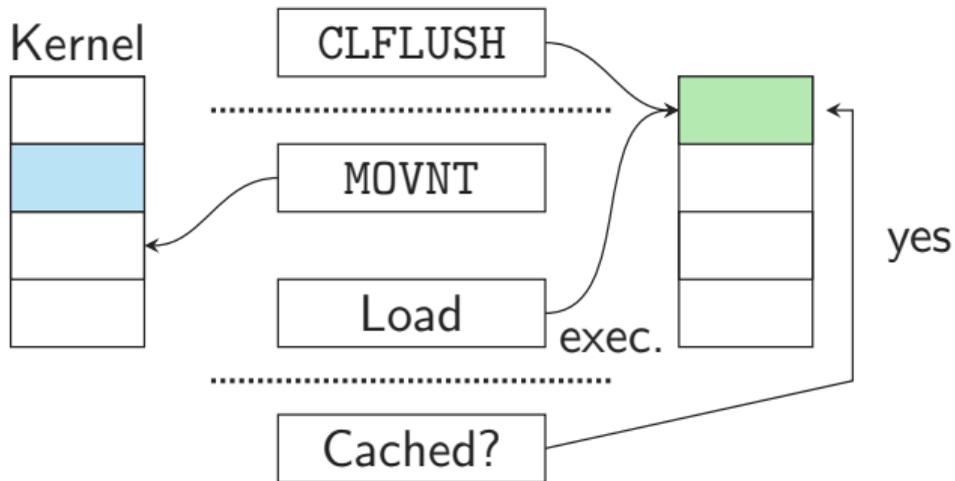
Findings – FlushConflict



Findings – FlushConflict



Findings – FlushConflict





FlushConflict Live Demo



- We can **automatically search** for side channels



- We can **automatically search** for side channels
- **Scaling** such tools is challenging



- We can **automatically search** for side channels
 - **Scaling** such tools is challenging
- Osiris prototype is currently limited to 1 instruction per sequence

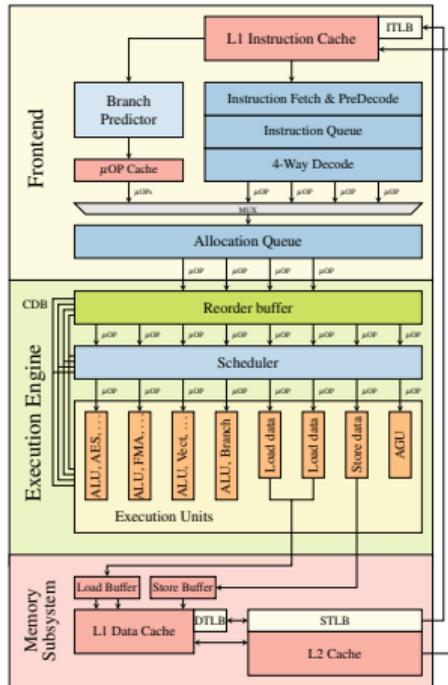


Can we find **more powerful attacks?**



Can we find **more powerful attacks**?
Can we **leak actual data** instead of meta data?

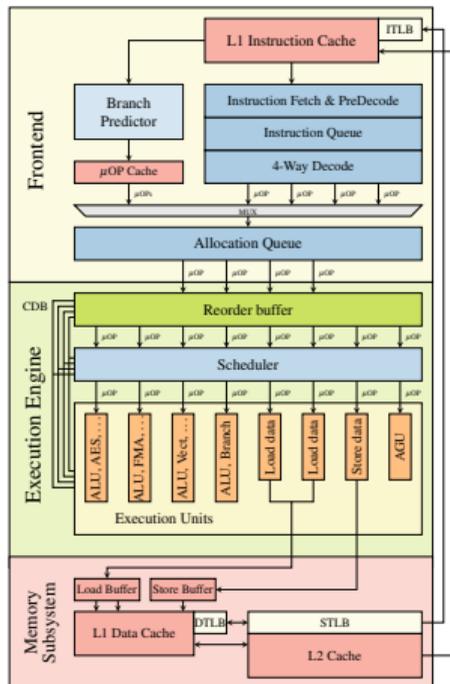
Out-of-Order Execution



Instructions are

- fetched and decoded in the **front-end**

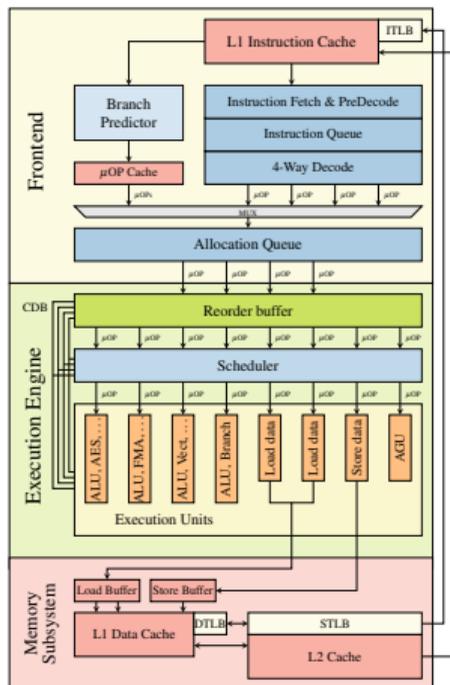
Out-of-Order Execution



Instructions are

- fetched and decoded in the **front-end**
- dispatched to the **backend**

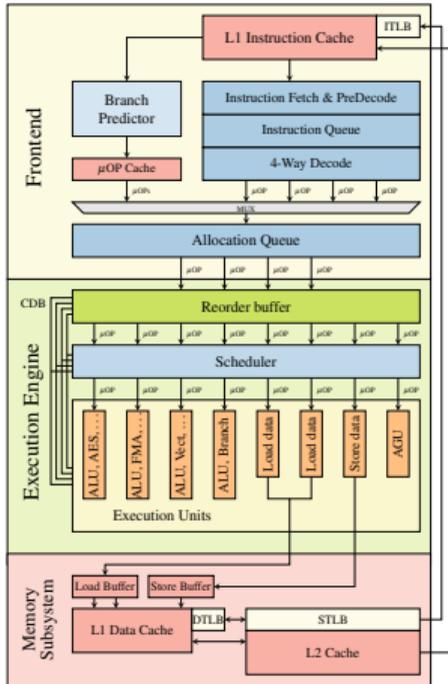
Out-of-Order Execution



Instructions are

- fetched and decoded in the **front-end**
- dispatched to the **backend**
- processed by **individual execution units**

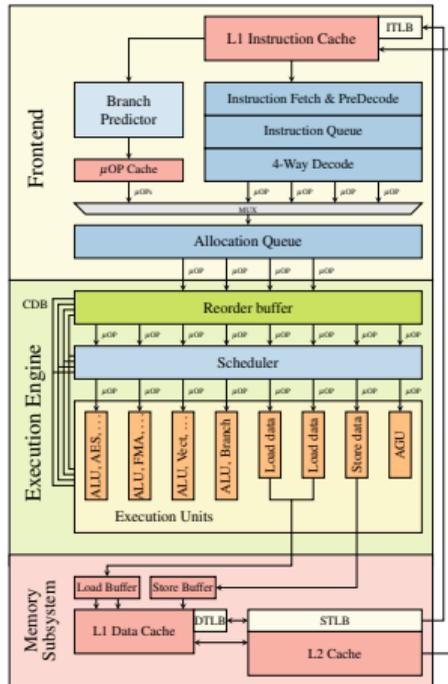
Out-of-Order Execution



Instructions

- are executed **out-of-order**

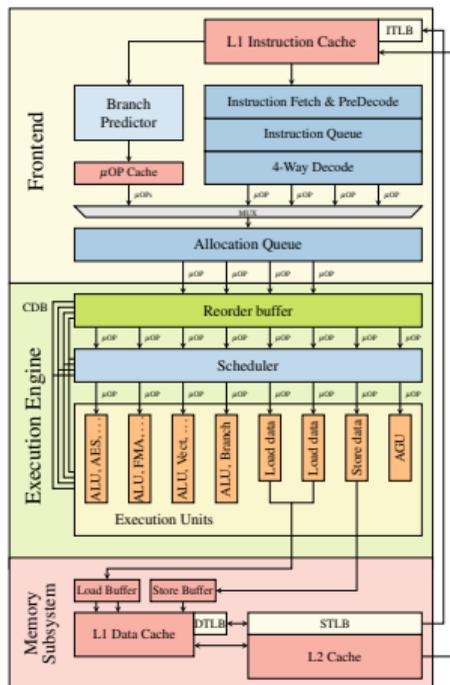
Out-of-Order Execution



Instructions

- are executed **out-of-order**
- wait until their **dependencies are ready**

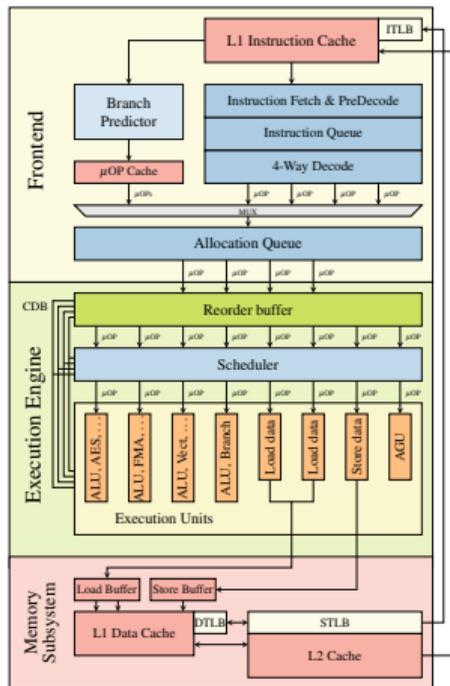
Out-of-Order Execution



Instructions

- are executed **out-of-order**
- wait until their **dependencies are ready**
 - Later instructions might execute prior earlier instructions

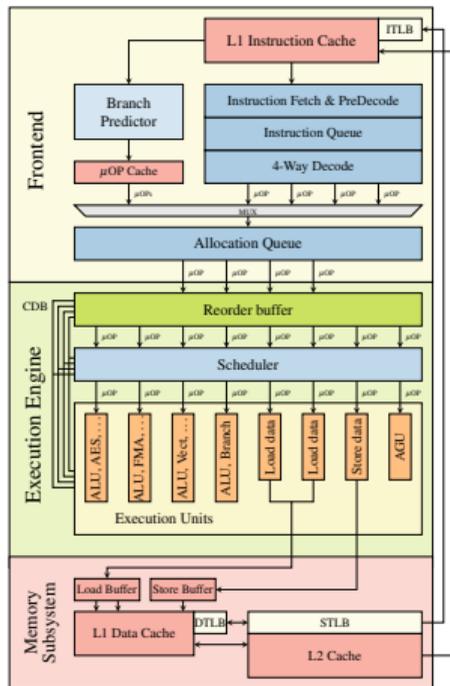
Out-of-Order Execution



Instructions

- are executed **out-of-order**
- wait until their **dependencies are ready**
 - Later instructions might execute prior earlier instructions
- **retire in-order**

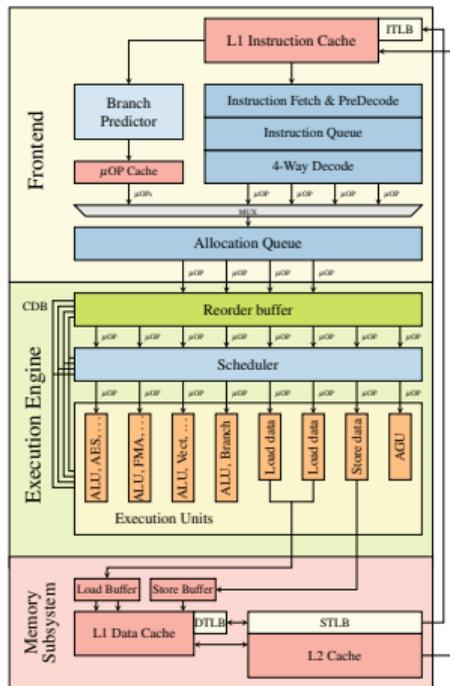
Out-of-Order Execution



Instructions

- are executed **out-of-order**
- wait until their **dependencies are ready**
 - Later instructions might execute prior earlier instructions
- **retire in-order**
 - State becomes architecturally visible

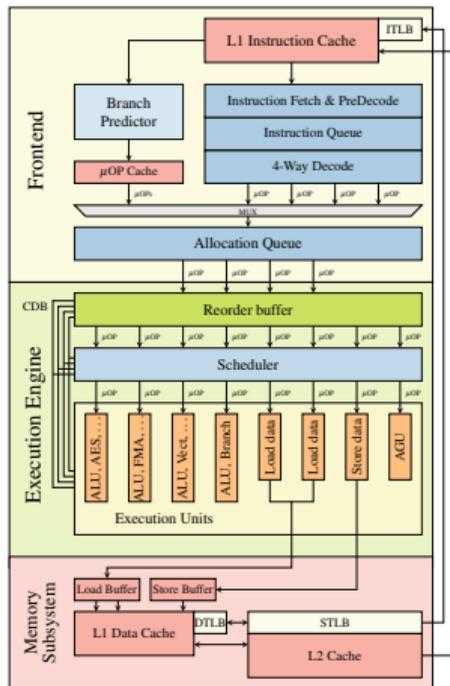
Out-of-Order Execution



Instructions

- are executed **out-of-order**
- wait until their **dependencies are ready**
 - Later instructions might execute prior earlier instructions
- **retire in-order**
 - State becomes architecturally visible
- **Exceptions** are checked during retirement

Out-of-Order Execution



Instructions

- are executed **out-of-order**
- wait until their **dependencies are ready**
 - Later instructions might execute prior earlier instructions
- **retire in-order**
 - State becomes architecturally visible
- **Exceptions** are checked during retirement
 - Flush pipeline and *recover state*

CPU Optimization: Lazy Error Handling (Out-Of-Order Execution)



- With **Meltdown**, we exploit out-of-order execution to leak the content of **inaccessible** kernel addresses



- With **Meltdown**, we exploit out-of-order execution to leak the content of **inaccessible** kernel addresses
 - Faulty state is **never architecturally visible**



- With **Meltdown**, we exploit out-of-order execution to leak the content of **inaccessible** kernel addresses
 - Faulty state is **never architecturally visible**
 - Use a side-channel (Flush+Reload) to make it visible



- Many different *attack variants*
 - ZombieLoad
 - TAA
 - MDS
 - Fallout
 - ...



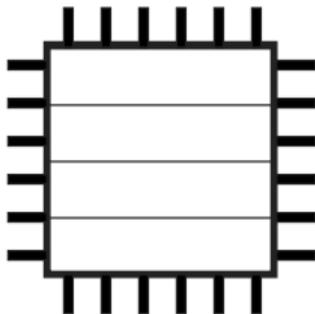
- Many different *attack variants*
 - ZombieLoad
 - TAA
 - MDS
 - Fallout
 - ...
- Leak from **different** microarchitectural buffers by triggering **different** microarchitectural conflicts

Meltdown-type Attacks

User Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z

```
char value = faulting[0]
```

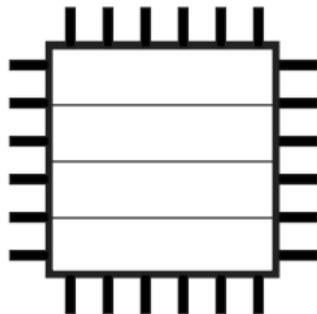


Meltdown-type Attacks

User Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z

```
char value = faulting[0]
```



Meltdown-type Attacks

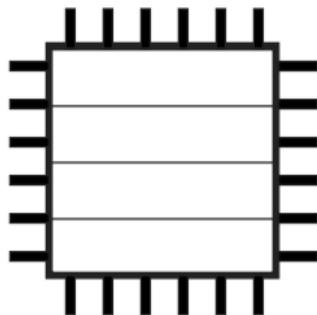
User Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z

```
char value = faulting[0]
```



K



Meltdown-type Attacks

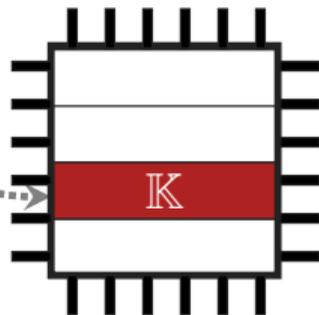
User Memory

	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z

```
char value = faulting[0]
```



K

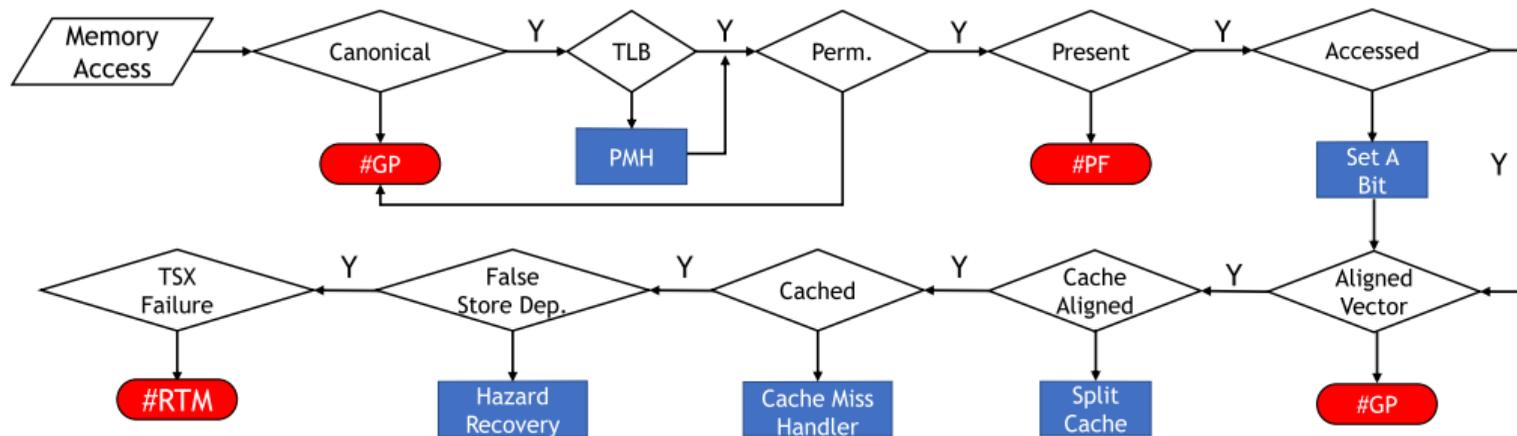


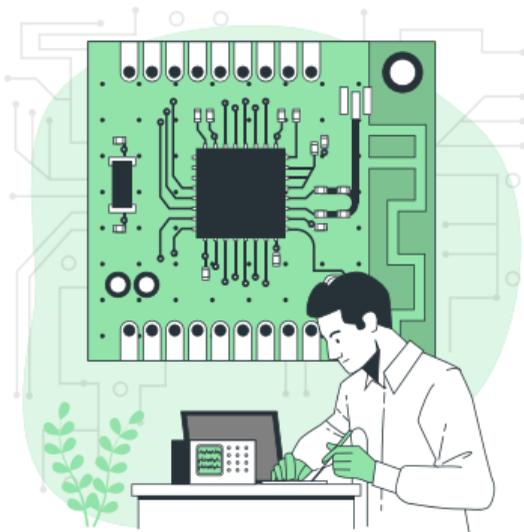
Memory Access Checks (simplified)

- Many possibilities for **faults** and **microcode assists**

Memory Access Checks (simplified)

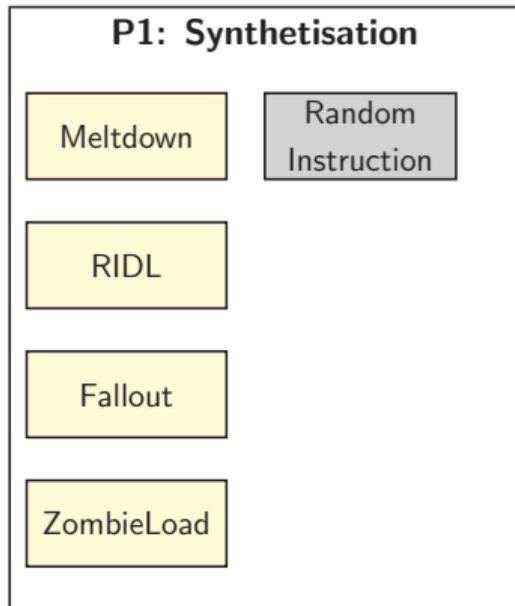
- Many possibilities for **faults** and **microcode assists**



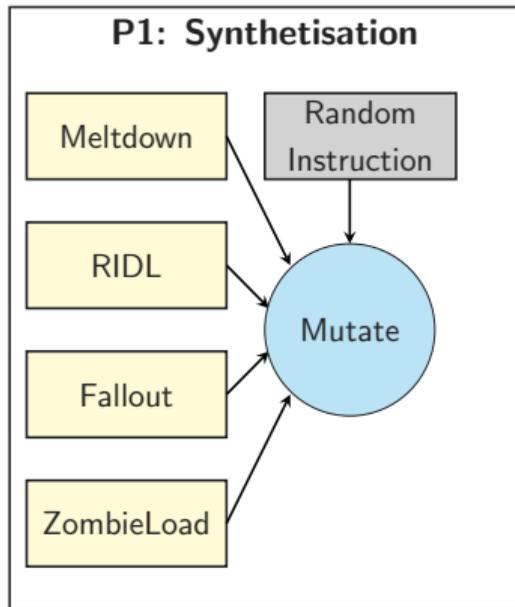


Fuzzing for Meltdown-type attacks

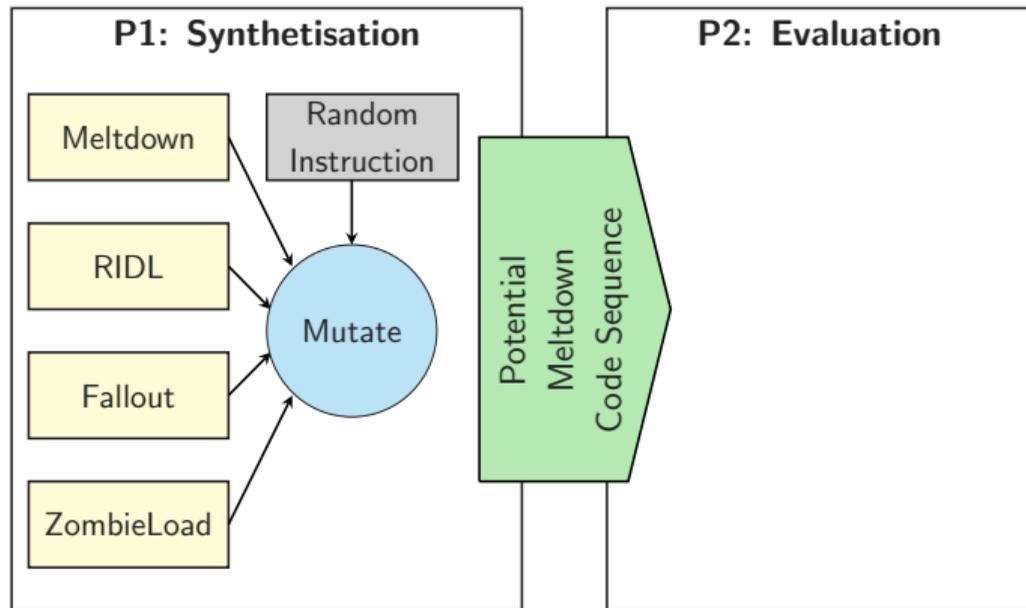
Transynther – Fuzzing for Meltdown-type Attacks



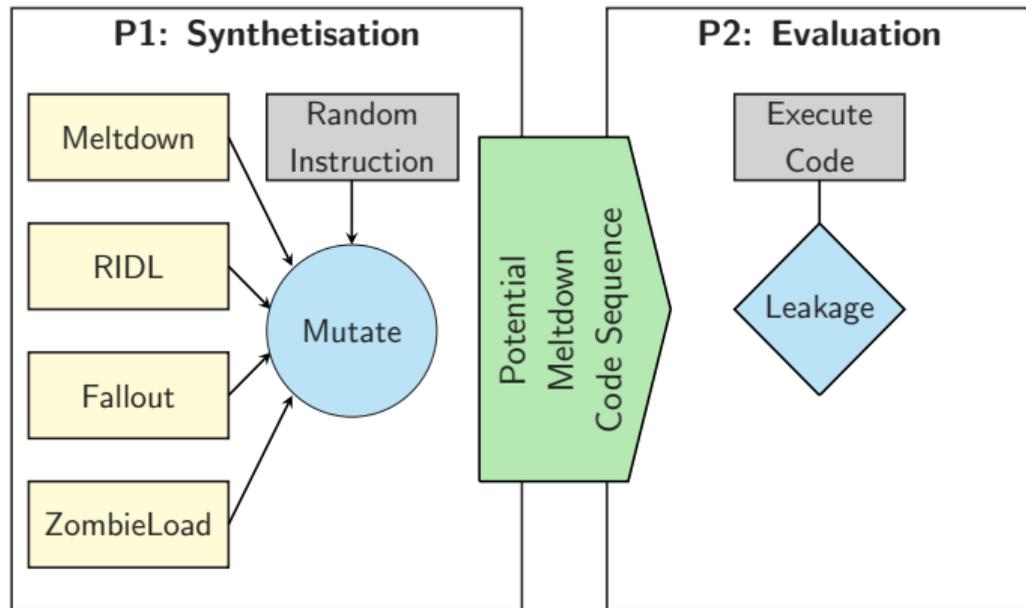
Transynther – Fuzzing for Meltdown-type Attacks



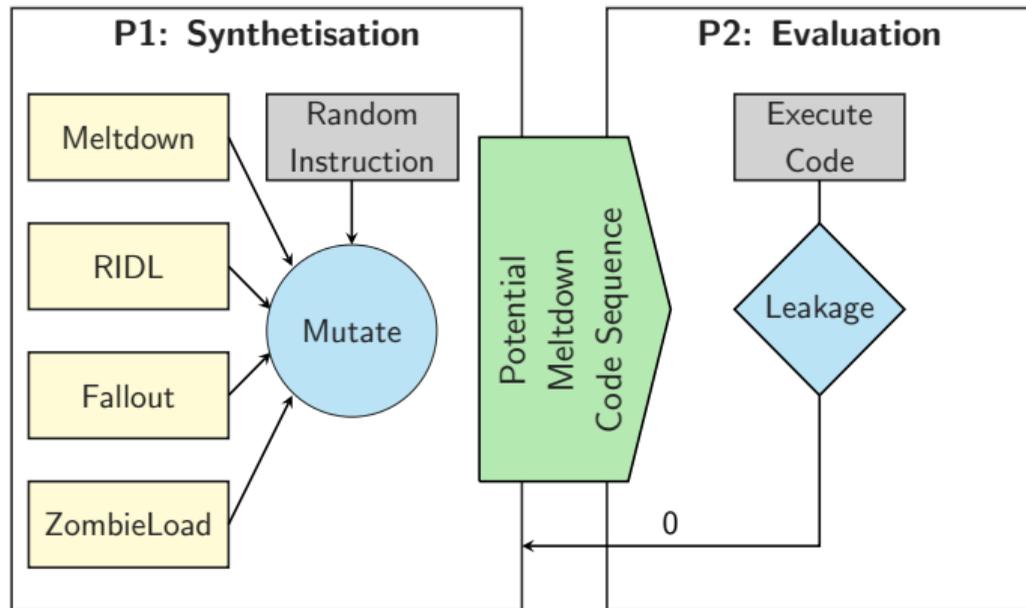
Transynther – Fuzzing for Meltdown-type Attacks



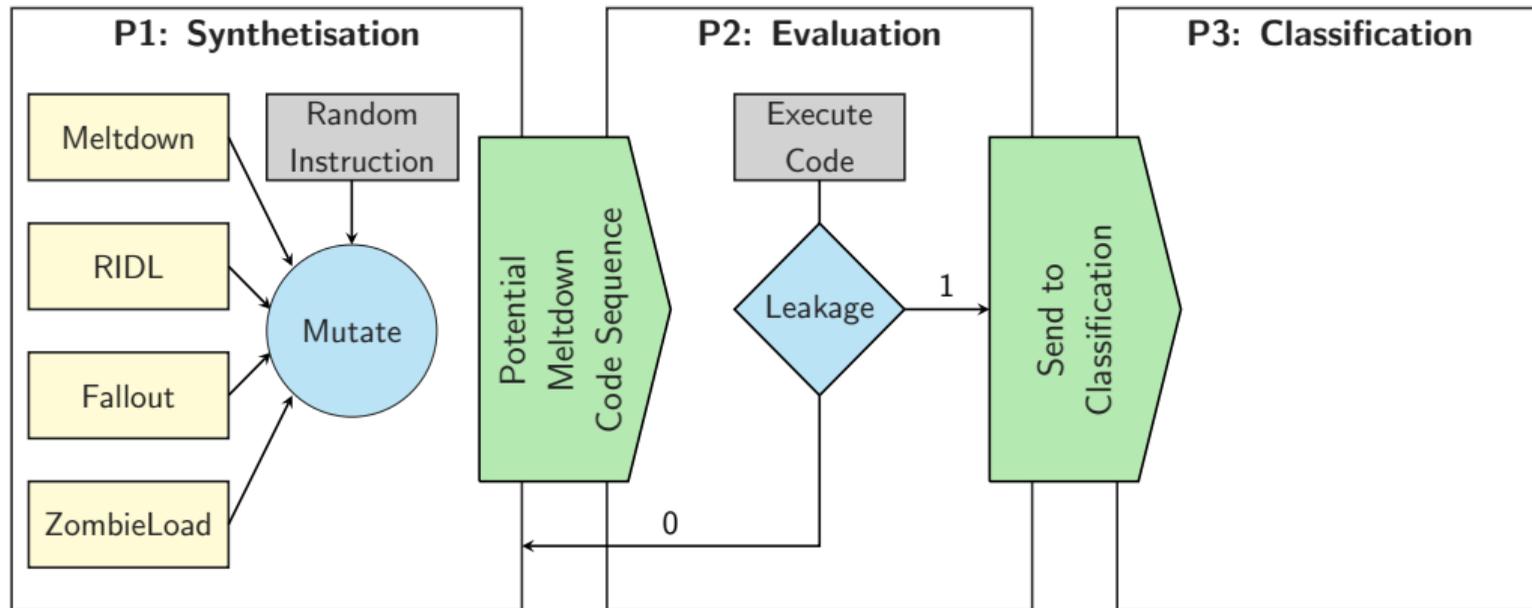
Transynther – Fuzzing for Meltdown-type Attacks



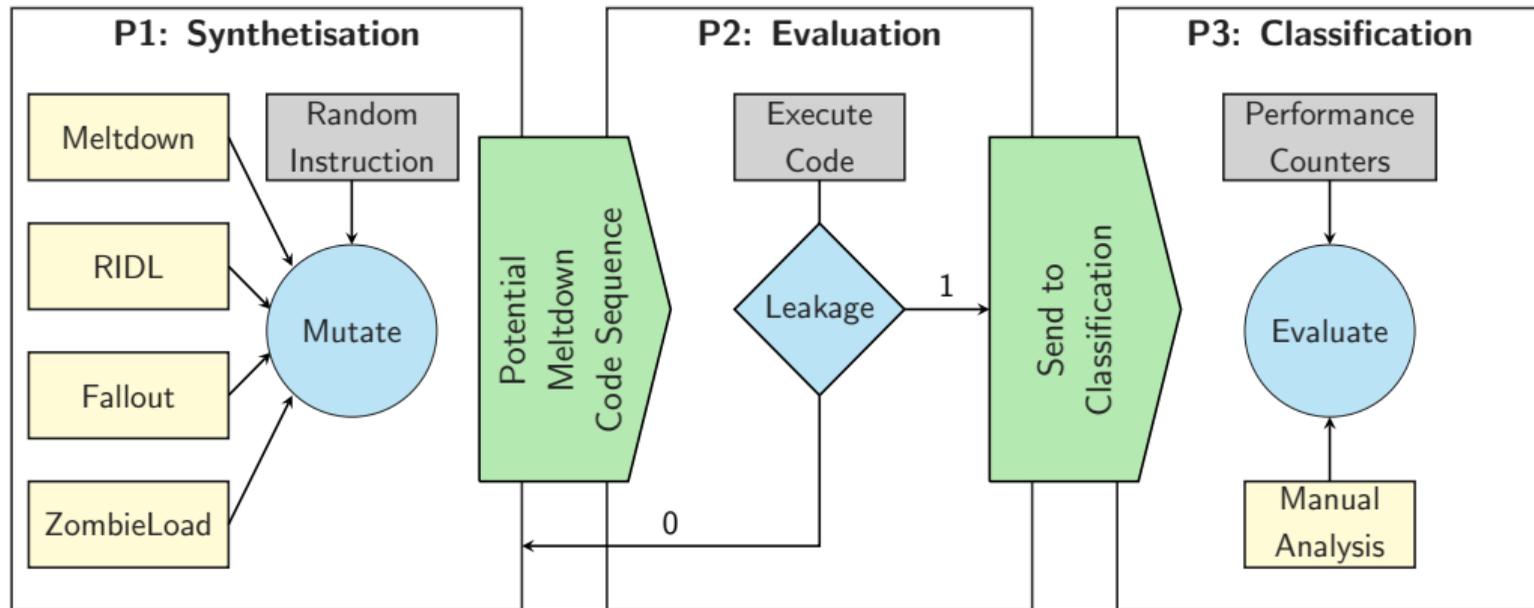
Transynther – Fuzzing for Meltdown-type Attacks



Transynther – Fuzzing for Meltdown-type Attacks



Transynther – Fuzzing for Meltdown-type Attacks





26 hours

Transynther – Results Overview



26 hours



100 unique leakage patterns

Transynther – Results Overview



26 hours



100 unique leakage patterns



7 attacks reproduced

Transynther – Results Overview



26 hours



100 unique leakage patterns



7 attacks reproduced



1 new vulnerability

Transynther – Results Overview



26 hours



100 unique leakage patterns



7 attacks reproduced



1 new vulnerability



1 regression



- **Medusa**: new variant of **ZombieLoad**



- **Medusa**: new variant of **ZombieLoad**
- Leaks from **write-combining buffer**, i.e., REP MOV

Findings – Reenabling MDS Attack



- **Medusa**: new variant of **ZombieLoad**
- Leaks from **write-combining buffer**, i.e., REP MOV
- Used for fast **memory copy**, e.g., in OpenSSL or kernel

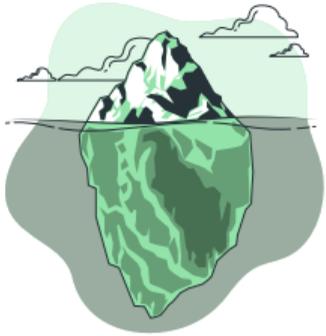


- **Medusa**: new variant of **ZombieLoad**
 - Leaks from **write-combining buffer**, i.e., REP MOV
 - Used for fast **memory copy**, e.g., in OpenSSL or kernel
- Leaked RSA key while decoding in OpenSSL



- Intel's Ice Lake CPUs were labeled **MDS-resistant**

Findings – MDS Attack on Ice Lake



- Intel's Ice Lake CPUs were labeled **MDS-resistant**
- Transsynther found a **working MDS attack**
 - Variant of a Fallout attack



- We can **automatically search** for transient-execution attacks



- We can **automatically search** for transient-execution attacks
- **Post-analysis** can also be automated



- We can **automatically search** for transient-execution attacks
- **Post-analysis** can also be automated
- **Regression bugs can be found** using automated tools



What about **unknown instructions**?



- **Sandsifter^a** searches for undocumented x86 instructions

Finding Unknown Instructions



- **Sandsifter**^a searches for undocumented x86 instructions
 - Trick to check for valid instruction bytes using page faults
- **Haruspex**^b uses speculative execution and performance counters

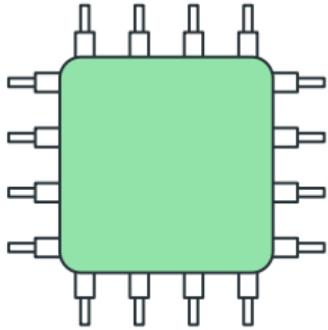
^aSandsifter (Christopher Domas): <https://github.com/xoreaxeaxeax/sandsifter>

^bHaruspex (Can Bölük): <https://github.com/can1357/haruspex>



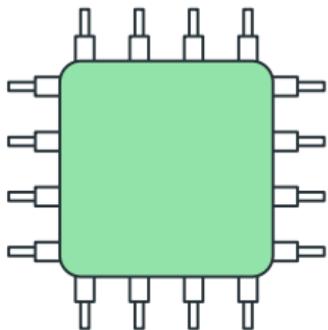
What else is **left**?

Model-Specific Registers and their Hidden Secrets



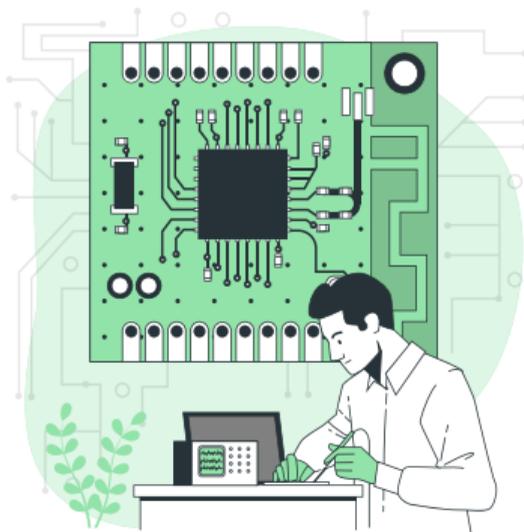
- Model-Specific Registers (MSRs) are **special CPU control registers**

Model-Specific Registers and their Hidden Secrets



- Model-Specific Registers (MSRs) are **special CPU control registers**
- VIA C3 CPU had **backdoor access**^a hidden behind an **undocumented MSR bit**

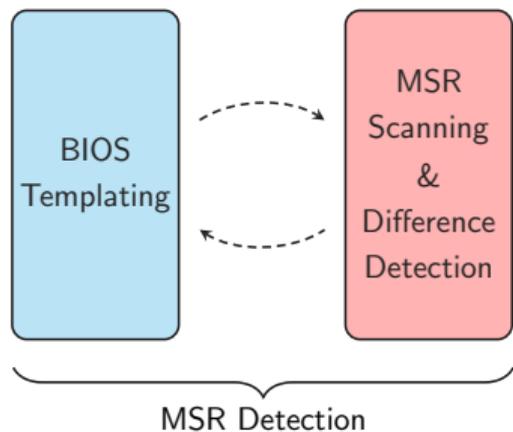
^aGOD MODE unlocked: Hardware backdoors in x86 CPUs (Christopher Domas – BlackHat USA '18)



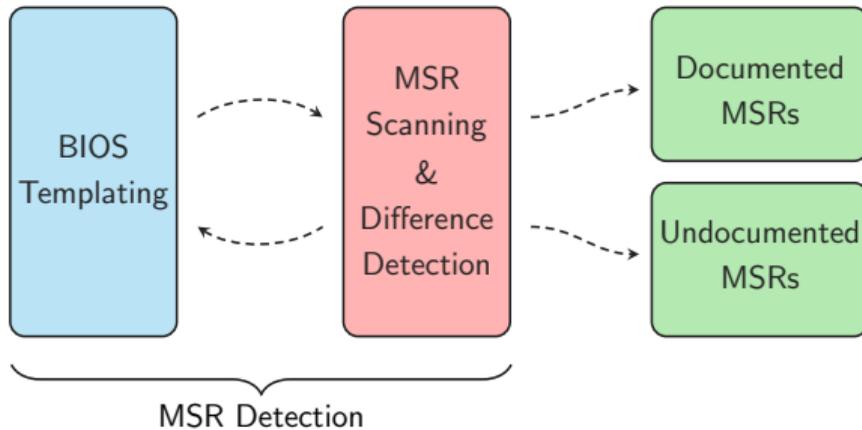
Searching for undocumented MSRs



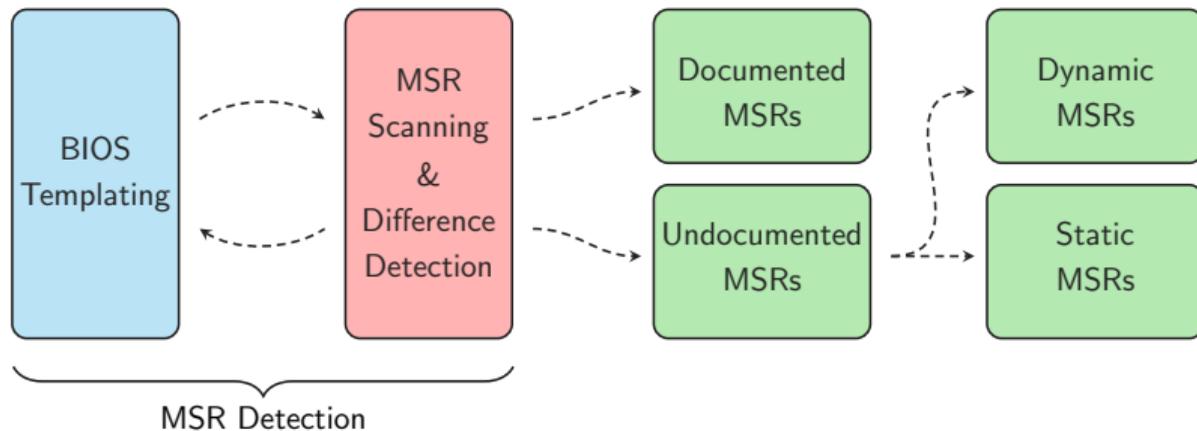
MSRevelio – Overview



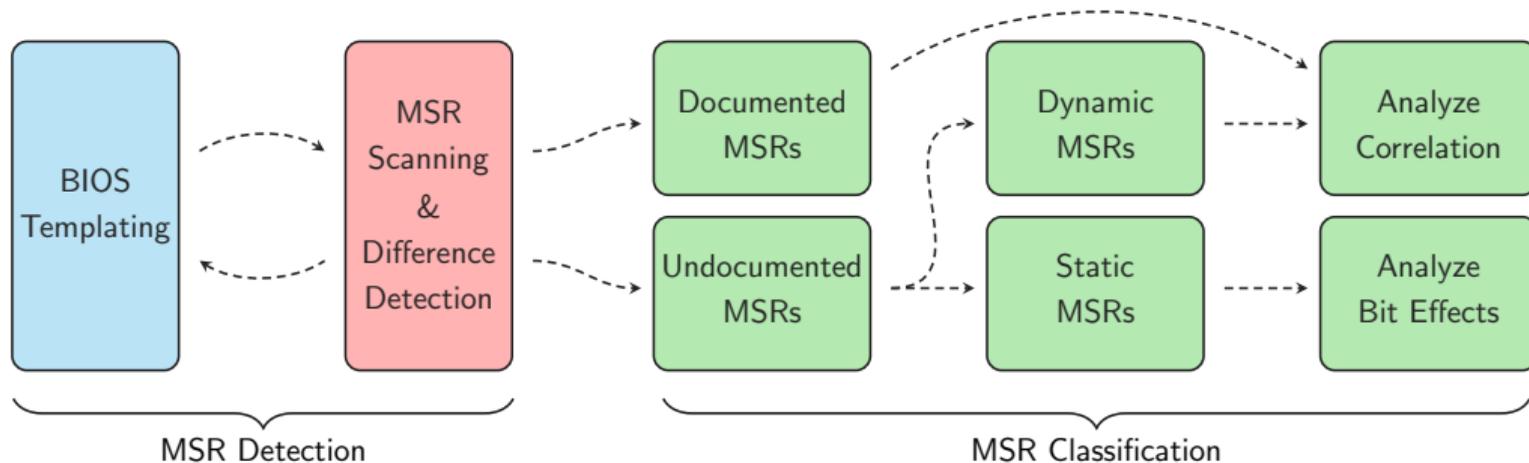
MSRevelio – Overview



MSRevelio – Overview



MSRevelio – Overview





5890 undocumented MSRs



5890 undocumented MSRs



3 MSRs as attack **mitigation**



5890 undocumented MSRs



3 MSRs as attack **mitigation**



1 MSR allows **TOCTOU** vulnerability



5890 undocumented MSRs



3 MSRs as attack **mitigation**



1 MSR allows **TOCTOU** vulnerability



New MSRs hint towards vulnerabilities

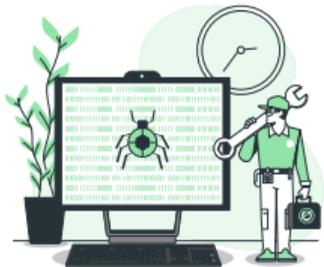


- MSRs often used to introduce **vulnerability fixes**

Findings – Undisclosed Attacks

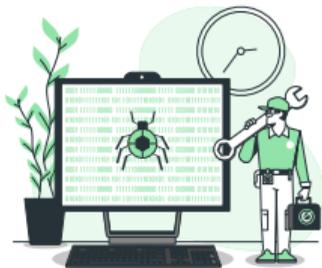


- MSRs often used to introduce **vulnerability fixes**
 - MSRs exist **before public disclosure**
- Useful for 1-Day Exploits

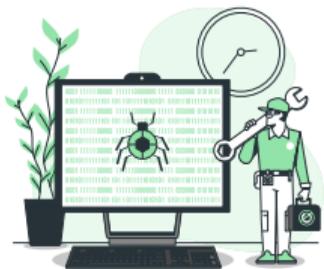


- Mitigate **prefetch side-channel attacks**

Findings – Attack Mitigation



- Mitigate **prefetch side-channel attacks**
- Reduce **Crosstalk leakage**



- Mitigate **prefetch side-channel attacks**
- Reduce **Crosstalk leakage**
- Reduce **Medusa leakage**



- Searching for **unknown behavior** is hard



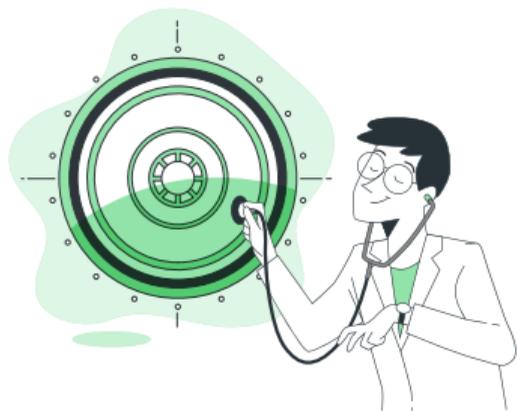
- Searching for **unknown behavior is hard**
- We can **automate the search** for undocumented MSR behavior



- Searching for **unknown behavior is hard**
- We can **automate the search** for undocumented MSR behavior
- Automation allows **tracing changes between releases**
(cf. Transynther)



Conclusion



Automated **side-channel**
discovery



Automated **side-channel**
discovery



Automated
transient-execution attack
discovery



Automated **side-channel**
discovery



Automated
transient-execution attack
discovery



Automated detection of
undocumented MSRs



- **Scalability:** Automation is easier than manual efforts (even if limited)



- **Scalability:** Automation is easier than manual efforts (even if limited)
- **Extensible:** Able to scan way **more variants** (often minor changes)



- **Scalability:** Automation is easier than manual efforts (even if limited)
- **Extensible:** Able to scan way **more variants** (often minor changes)
- **Versatile:** Execute on **various microarchitectures**



All our tools are **accessible** and **open source**:



All our tools are **accessible** and **open source**:

→ <https://github.com/D4nielMoghimi/medusa>

→ <https://github.com/CISPA/osiris>

→ <https://github.com/IAIK/msrevelio>



- Automated approaches are a **fruitful** way in



- Automated approaches are a **fruitful** way in
 - finding **new and unknown** attack variants



- Automated approaches are a **fruitful** way in
 - finding **new and unknown** attack variants
 - **regression** testing



- Automated approaches are a **fruitful** way in
 - finding **new and unknown** attack variants
 - **regression** testing
 - **extensible** for new building blocks and architectures



- Automated approaches are a **fruitful** way in
 - finding **new and unknown** attack variants
 - **regression** testing
 - **extensible** for new building blocks and architectures
- Despite **limitations**, great toolkit and support for manual efforts



- Automated approaches are a **fruitful** way in
 - finding **new and unknown** attack variants
 - **regression** testing
 - **extensible** for new building blocks and architectures
- Despite **limitations**, great toolkit and support for manual efforts
- Play an **essential role** in CPU research

CPU Fuzzing: Automated Discovery of Microarchitectural Attacks

15 - 17 NOVEMBER 2022
RIYADH FRONT EXHIBITION CENTRE
SAUDI ARABIA

Daniel Weber, Michael Schwarz, Moritz Lipp

CO-ORGANISED BY:



الاتحاد السعودي للأمن
السيبراني والبرمجة والذو
السيبراني
SAUDI FEDERATION FOR CYBERSECURITY
PROGRAMMING & DRONES

IN PARTNERSHIP WITH:



الهيئة العامة
للأمن السيبراني
Cyber Security Authority



- Icons and Images from [storyset.com](https://www.storyset.com)
- Images from CNN