







From Random Observations to Automated Leakage Discovery

Michael Schwarz

December 2022

CISPA Helmholtz Center for Information Security







- Mental model of CPU is simple



- Mental model of CPU is simple
- Instructions are executed **in program order**



- Mental model of CPU is simple
- Instructions are executed **in program order**
- Pipeline **stalls** when stages are not ready




- Mental model of CPU is simple
- Instructions are executed **in program order**
- Pipeline **stalls** when stages are not ready
- If data is **not cached**, we need to wait

INSTRUCTION


●	LOAD	5
●	MPY	6
●	MPY	6
●	PRINT	7

INFORMATION


●	33416
●	63




In-Order Execution




- Instructions are...




- Instructions are...
 - fetched (IF) from the L1 Instruction Cache




- Instructions are...
 - fetched (IF) from the L1 Instruction Cache
 - decoded (ID)



- Instructions are...
 - fetched (IF) from the L1 Instruction Cache
 - decoded (ID)
 - executed (EX) by execution units



- Instructions are...
 - fetched (IF) from the L1 Instruction Cache
 - decoded (ID)
 - executed (EX) by execution units
- Memory access is performed (MEM)




- Instructions are...
 - fetched (IF) from the L1 Instruction Cache
 - decoded (ID)
 - executed (EX) by execution units
- Memory access is performed (MEM)
- Architectural register file is updated (WB)


$$x = y + 1$$

Measuring Time

```
start = 🕒  
  x = y + 1  
end = 🕒
```


start = 

x = y + 1

end = 

$\Delta = \text{end} - \text{start}$

```
start = 🕒
```

```
  x = y + 1
```

```
end = 🕒
```

```
 $\Delta$  = end - start
```

1. run: Δ = 302

```
start = 🕒
```

```
  x = y + 1
```

```
end = 🕒
```

```
 $\Delta = \text{end} - \text{start}$ 
```

1. run: $\Delta = 302$

2. run: $\Delta = 54$

```
start = 🕒
```

```
x = y + 1
```

```
end = 🕒
```

```
 $\Delta = \text{end} - \text{start}$ 
```


1. run: $\Delta = 302$

2. run: $\Delta = 54$

Determinism?

Same code with **different** execution time **without** changes

Insanity: doing the same thing over and over again and expecting different results.



Albert Einstein

Measuring Time

start = 🕒

end = 🕒

start = 🕒

end = 🕒

$\Delta = \text{end} - \text{start}$

start = 🕒

end = 🕒

$\Delta = \text{end} - \text{start}$

1. run: $\Delta = 12$

start = 🕒

end = 🕒

$\Delta = \text{end} - \text{start}$

1. run: $\Delta = 12$

2. run: $\Delta = 12$


start = 🕒

end = 🕒

$\Delta = \text{end} - \text{start}$

1. run: $\Delta = 12$

2. run: $\Delta = 12$



**I HAVE NO
IDEA WHAT
I'M DOING**



Interrupts!

App



OS



Interrupts!

App



OS

Δ 



Interrupts!

App



OS

Δ 



Interrupts!

App




OS


Δ 




Interrupts!




Interrupts!




Interrupts!





Interrupts!



Interrupts!







```
int now = rdtsc();
while (true) {
    int last = now;
    now = rdtsc();
    if ((now - last) > threshold) {
        reportEvent(now, now - last);
    }
}
```



```
int now = rdtsc();  
while (true) {  
    int last = now;  
    now = rdtsc();  
    if ((now - last) > threshold) {  
        reportEvent(now, now - last);  
    }  
}
```


- Continuously acquire [high-resolution timestamp](#)




```
int now = rdtsc();  
while (true) {  
    int last = now;  
    now = rdtsc();  
    if ((now - last) > threshold) {  
        reportEvent(now, now - last);  
    }  
}
```


- Continuously acquire **high-resolution timestamp**
- Interrupt → large **difference** between timestamps

Interrupt-timing Attacks





Is that everything?



- Explains last experiment...
- ...but what about the simple calculation?
- Noise?



- **Memory** operations have different **runtimes**



- **Memory** operations have different **runtimes**
- Depends where the data is **located**



- **Memory** operations have different **runtimes**
- Depends where the data is **located**
 - Loading from **cache**: fast



- **Memory** operations have different **runtimes**
- Depends where the data is **located**
 - Loading from **cache**: fast
 - Loading from **memory**: slow




- **Memory** operations have different **runtimes**
- Depends where the data is **located**
 - Loading from **cache**: fast
 - Loading from **memory**: slow
 - Loading from **disk**: extremely slow




- Memory operations have different runtimes
- Depends where the data is located
 - Loading from cache: fast
 - Loading from memory: slow
 - Loading from disk: extremely slow
- Can also be measured


Caching Speeds-up Memory Accesses




Caching Speeds-up Memory Accesses




**IS MY MENTAL
MODEL OF THE CPU WRONG?**



**NO, IT MUST BE THE
MEASUREMENTS THAT ARE WRONG**




Reality: (Simplified) Modern CPU





Reality: (Simplified) Modern CPU






Reality: (Simplified) Modern CPU




Reality: (Simplified) Modern CPU



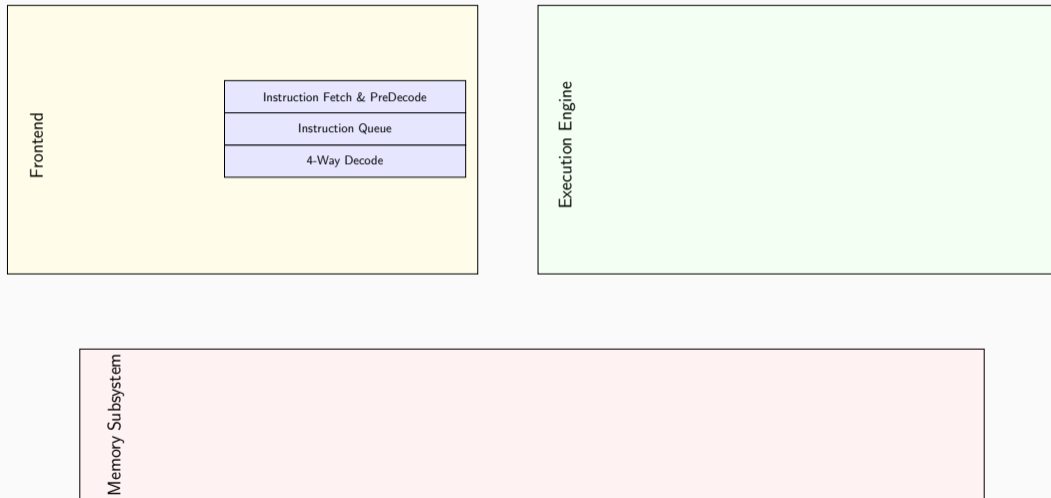
Reality: (Simplified) Modern CPU



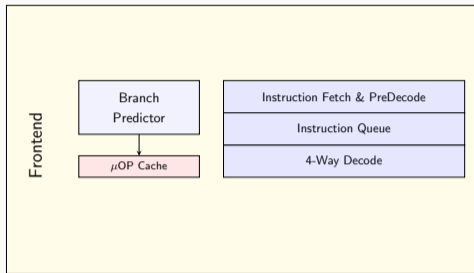
Reality: (Simplified) Modern CPU




Reality: (Simplified) Modern CPU



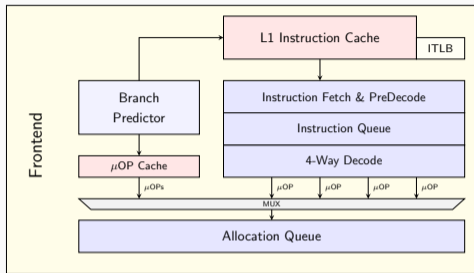
Reality: (Simplified) Modern CPU



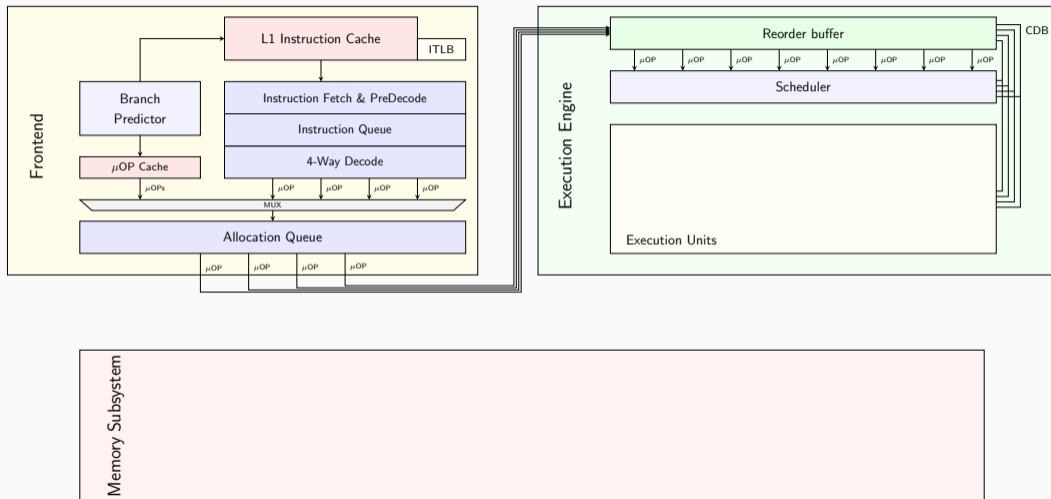
Reality: (Simplified) Modern CPU




Reality: (Simplified) Modern CPU




Reality: (Simplified) Modern CPU



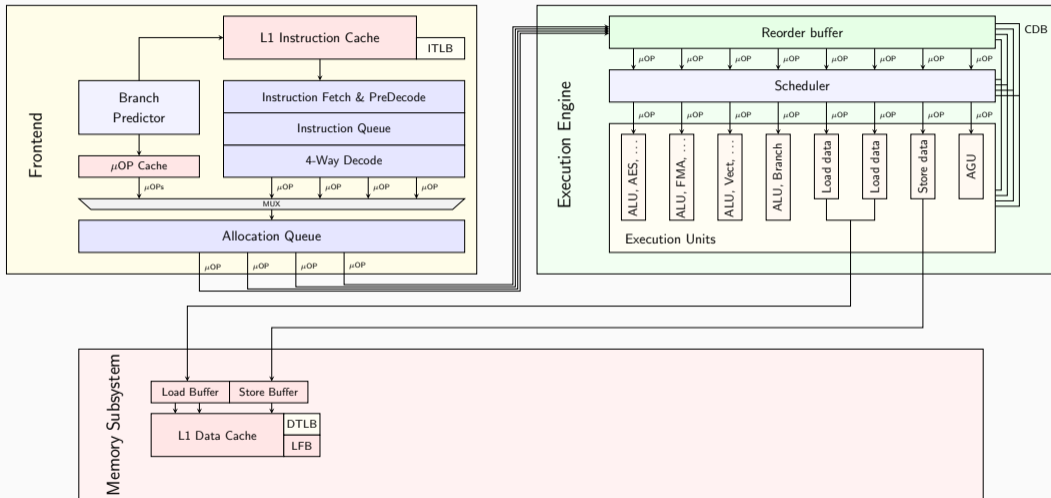
Reality: (Simplified) Modern CPU




Reality: (Simplified) Modern CPU



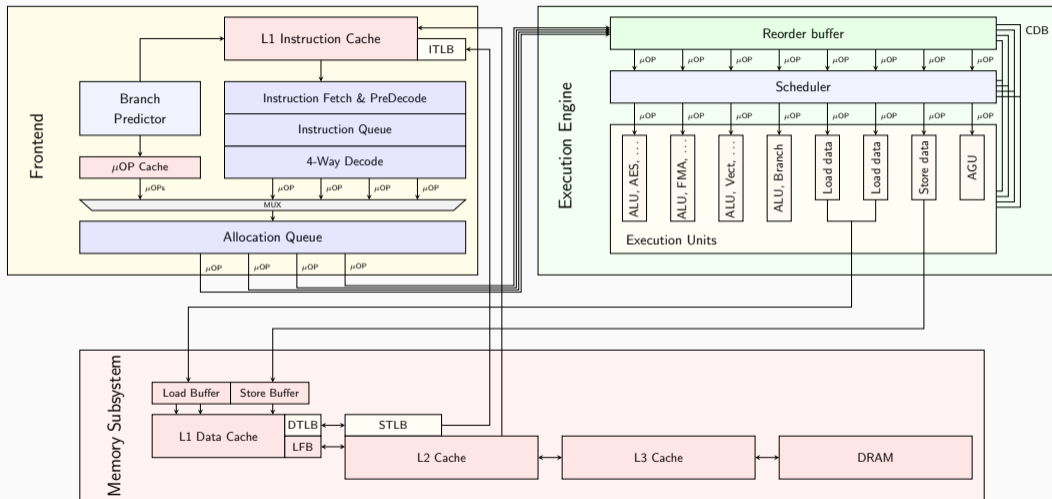
Reality: (Simplified) Modern CPU



Reality: (Simplified) Modern CPU




Reality: (Simplified) Modern CPU






- Cars all have the same **interface** (= architecture)
 - Steering wheel, pedals, gear stick, ...
- Some have special **extensions**
 - Air conditioning, cruise control, ...
- Driving skills are “compatible” with all cars




- Cars are **implemented** differently (= microarchitecture)
→ engine, fuel, motor control, ...
- Same car (“architecture”) with different engines
→ stronger or more efficient “microarchitecture”
- Drivers don't need to know anything about internals

No thorough Description



Intel manual
(full architecture)



Intel optimization manual
(microarchitecture parts)



4778 pages



868 pages

The prefetch Instruction




Use of software prefetch should be limited to memory addresses that are managed or owned within the application context. Prefetching to addresses that are not mapped to physical pages can experience **non-deterministic** performance **penalty**.


LETS DO IT!




Prefetch Timings on the Operating System



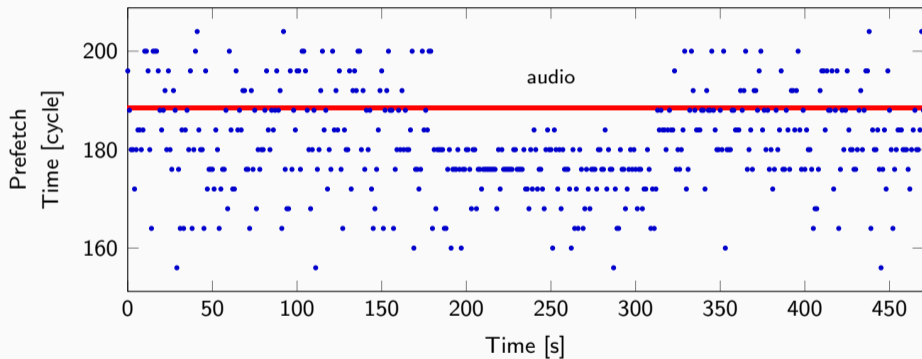
Prefetch Timings on the Operating System



Prefetch Timings on the Operating System



Prefetch Timings on the Operating System



I STILL HAVE NO IDEA

WHAT I'M DOING



Reset instruction



Instruction to measure



Instruction with possible
side effect

Testing A Sequence Triple



Testing A Sequence Triple



Testing A Sequence Triple



Cold path S0

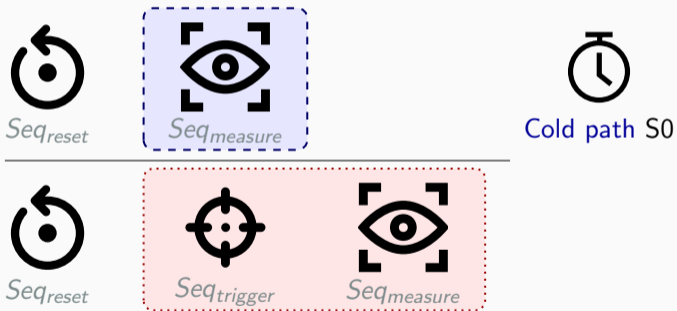
Testing A Sequence Triple



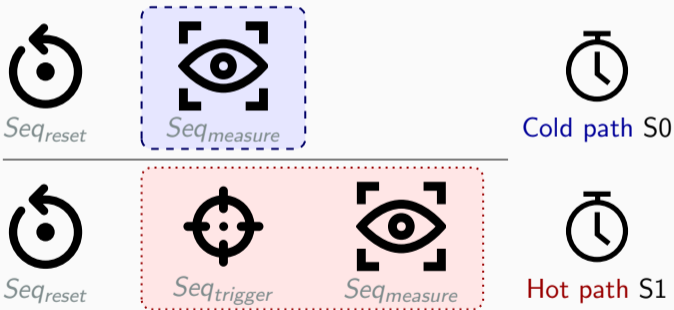
Cold path S0



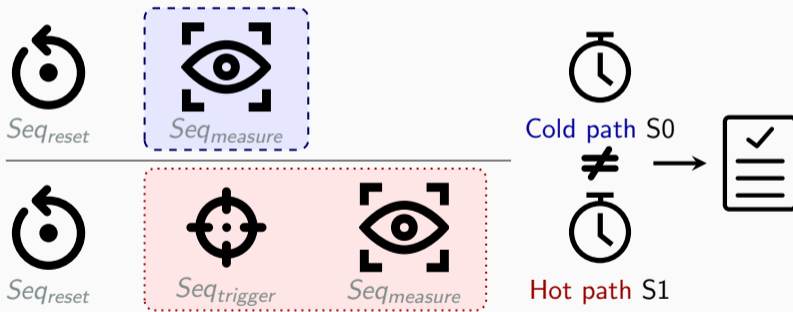
Testing A Sequence Triple



Testing A Sequence Triple

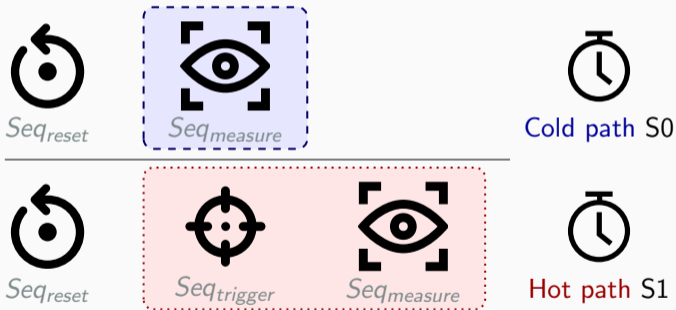


Testing A Sequence Triple



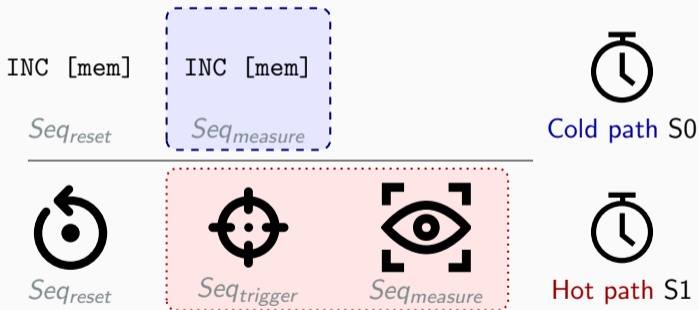
Testing A Sequence Triple

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




Testing A Sequence Triple

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



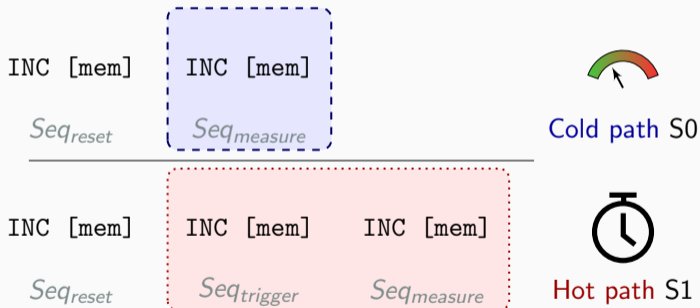
Testing A Sequence Triple

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



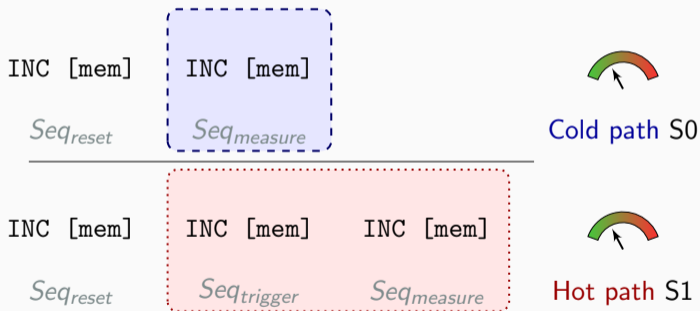
Testing A Sequence Triple

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




Testing A Sequence Triple

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




Testing A Sequence Triple

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




Testing A Sequence Triple

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




Testing A Sequence Triple

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



Testing A Sequence Triple

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



Recap: Measuring Time

```
start = ⌚
```

```
  x = y + 1
```

```
end = ⌚
```

```
 $\Delta$  = end - start
```



```
start = ⌚
```

```
  x = y + 1
```

```
end = ⌚
```

```
 $\Delta = \text{end} - \text{start}$ 
```

1. run: $\Delta = 302 \rightarrow$ cache miss
2. run: $\Delta = 54 \rightarrow$ cache hit

```
clflush [y]
```

```
start = ⌚
```

```
  x = y + 1
```

```
end = ⌚
```

```
 $\Delta = \text{end} - \text{start}$ 
```

1. run: $\Delta = 302 \rightarrow$ cache miss
2. run: $\Delta = 302 \rightarrow$ cache miss

```
clflush [y]
```

```
start = ⌚
```

```
  x = y + 1
```


```
end = ⌚
```

```
 $\Delta = \text{end} - \text{start}$ 
```


1. run: $\Delta = 302 \rightarrow$ cache miss
2. run: $\Delta = 302 \rightarrow$ cache miss

Determinism!


No randomness or non-determinism – just behavior we did not understand




Flush+Reload




Flush+Reload




Flush+Reload




Flush+Reload




Flush+Reload




Flush+Reload





$$M = C^d \bmod n$$

Flush+Reload on Square-and-Multiply




Flush+Reload on Square-and-Multiply

$$M = C^d \bmod n$$


1 1 0 0 1 1 0 ...

$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}} \times \underbrace{C}_{\text{multiply}}$$


Flush+Reload on Square-and-Multiply



Flush+Reload on Square-and-Multiply

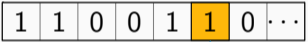


Flush+Reload on Square-and-Multiply

$$M = C^d \pmod n$$



$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}} \times \underbrace{C}_{\text{multiply}}$$

Flush+Reload on Square-and-Multiply

$$M = C^d \pmod n$$


$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}} \times \underbrace{C}_{\text{multiply}}$$

Flush+Reload on Square-and-Multiply



Problem



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

MEASURE



ALL THE THINGS


Offline





ISA





Instructions















- Fuzzed on 5 different CPUs
- AMD and Intel



~4 days per CPU




2 side channels rediscovered



4 new side channels




2 new attacks



- Random-number generator RDRAND
 - VM in the cloud (e.g., AWS) sees usage of other VMs
- Breaks VM isolation

RDRAND Covert Channel - Properties



AMD and Intel



VM and native



1000 bit/s




No memory

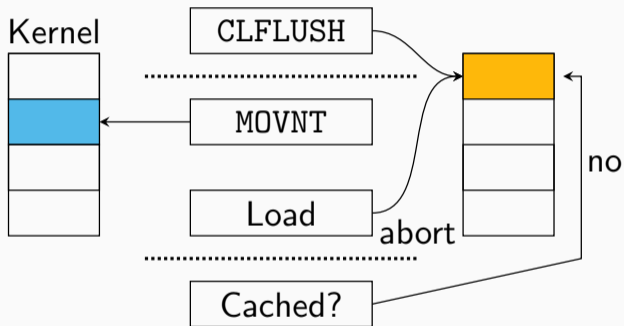



No detection

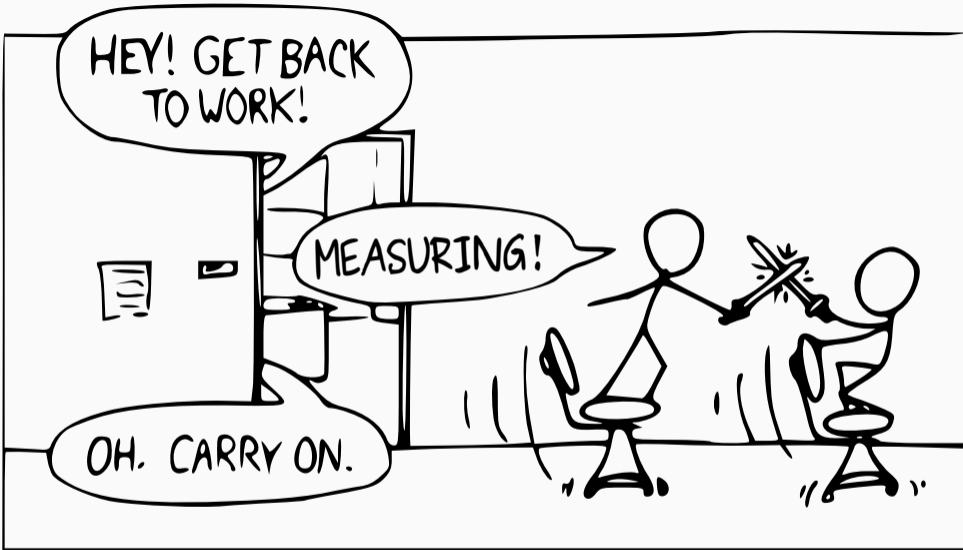


No mitigation









```
x = y + 1
```

```
start = ⌚
```

```
  x = y + 1
```

```
end = ⌚
```

```
 $\Delta = \text{end} - \text{start}$ 
```

```
x = y + 1
```

```
start = ⌚
```

```
  x = y + 1
```

```
end = ⌚
```

```
 $\Delta = \text{end} - \text{start}$ 
```

1. run: $\Delta = 54 \rightarrow$ cache hit
2. run: $\Delta = 54 \rightarrow$ cache hit

<crash>

$x = y + 1$ (never executed architecturally)

<restart>

start = ⌚

$x = y + 1$

end = ⌚

$\Delta = \text{end} - \text{start}$

1. run:

2. run:

<crash>

$x = y + 1$ (never executed architecturally)

<restart>

start = ⌚

$x = y + 1$

end = ⌚

$\Delta = \text{end} - \text{start}$

1. run: $\Delta = 54 \rightarrow$ cache hit

2. run: $\Delta = 54 \rightarrow$ cache hit


(Im)possible

Microarchitecture can do things impossible for the architecture

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 = kernel[0]
```




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 = kernel[0]
```

⚡ Page fault (Exception)



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 = kernel[0]
```




```
mem[value]
```

K



Page fault (Exception)

Out of order



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 = kernel[0]`

`mem[value]`


K

Page fault (Exception)

Out of order




Meltdown Experiment




Kernel Memory

Meltdown Experiment



Meltdown Experiment



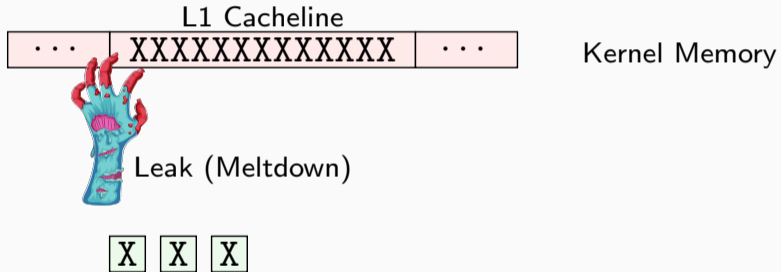
Kernel Memory




Leak (Meltdown)



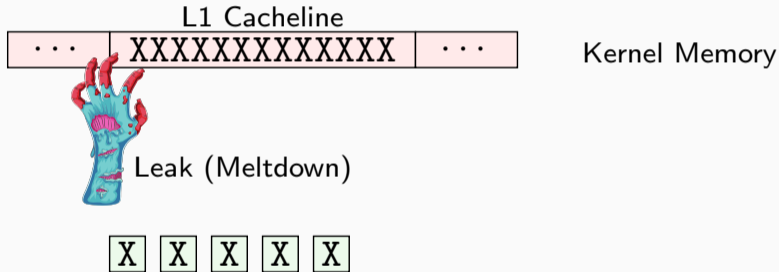
Meltdown Experiment



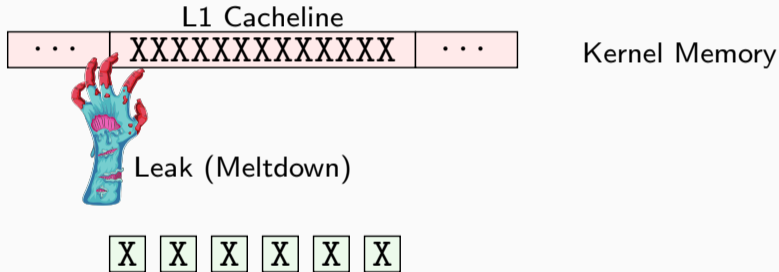
Meltdown Experiment




Meltdown Experiment




Meltdown Experiment




Meltdown Experiment




Meltdown Experiment




Meltdown Experiment




Meltdown Experiment




Leak (Meltdown)




Meltdown Experiment




Leak (Meltdown)




Meltdown Experiment



Leak (Meltdown)





Meltdown Experiment




X	X	X	X	X	X	X	P	X	X	X	X	X
---	---	---	---	---	---	---	---	---	---	---	---	---

Meltdown Experiment




Meltdown Experiment



Leak (Meltdown)




Meltdown Experiment




Leak (Meltdown)

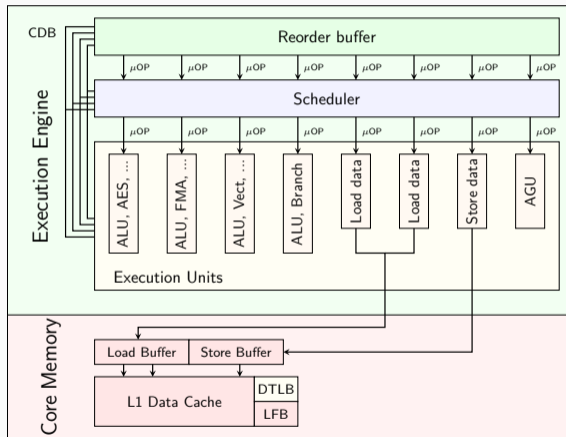




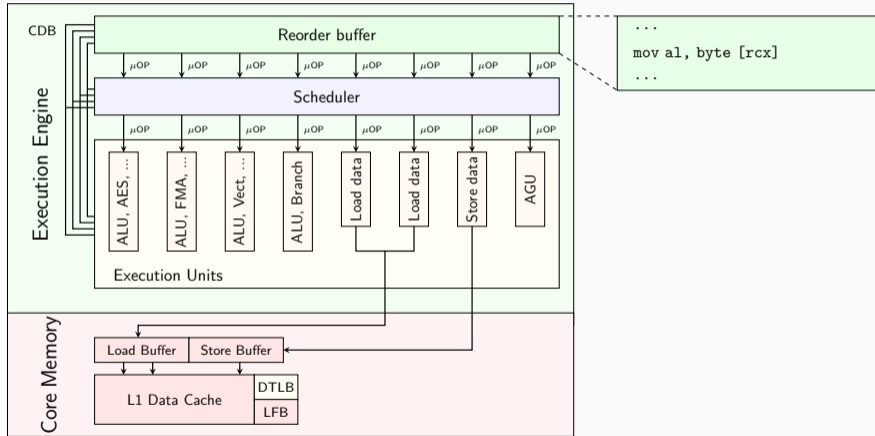
Leak (Meltdown)




Complex Load Situations




Complex Load Situations




Complex Load Situations



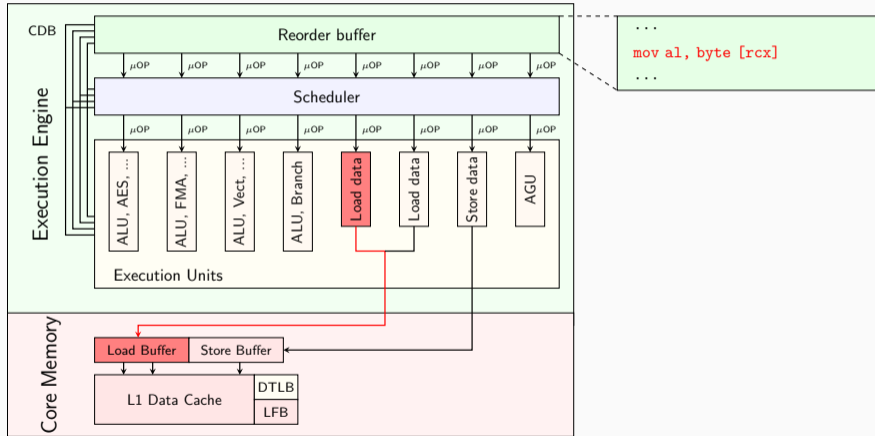
Complex Load Situations



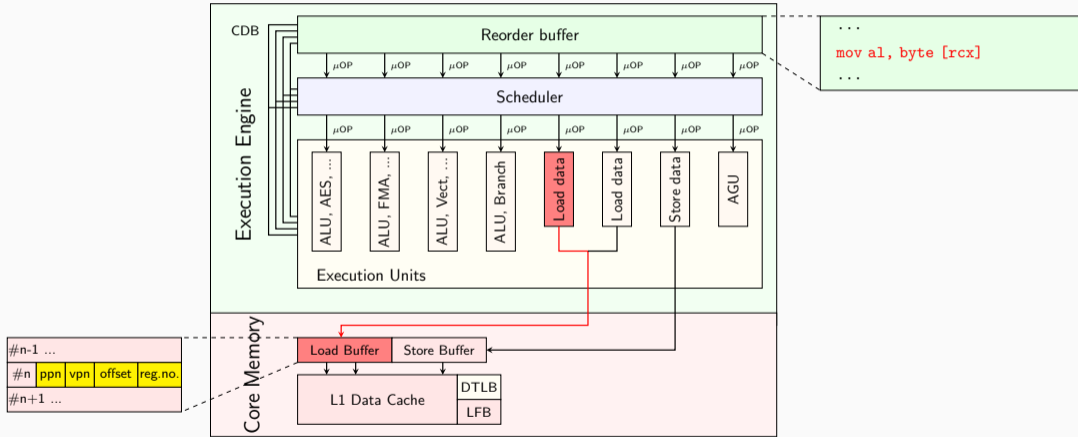
Complex Load Situations



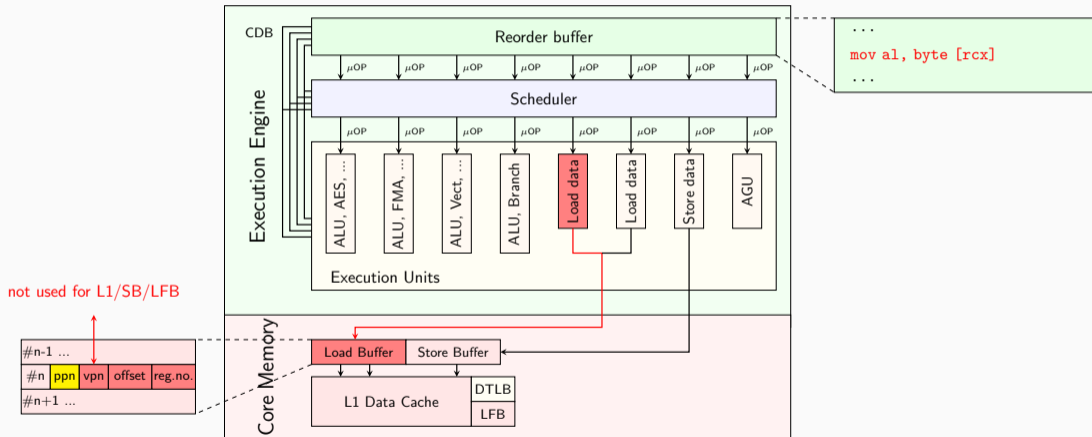
Complex Load Situations




Complex Load Situations



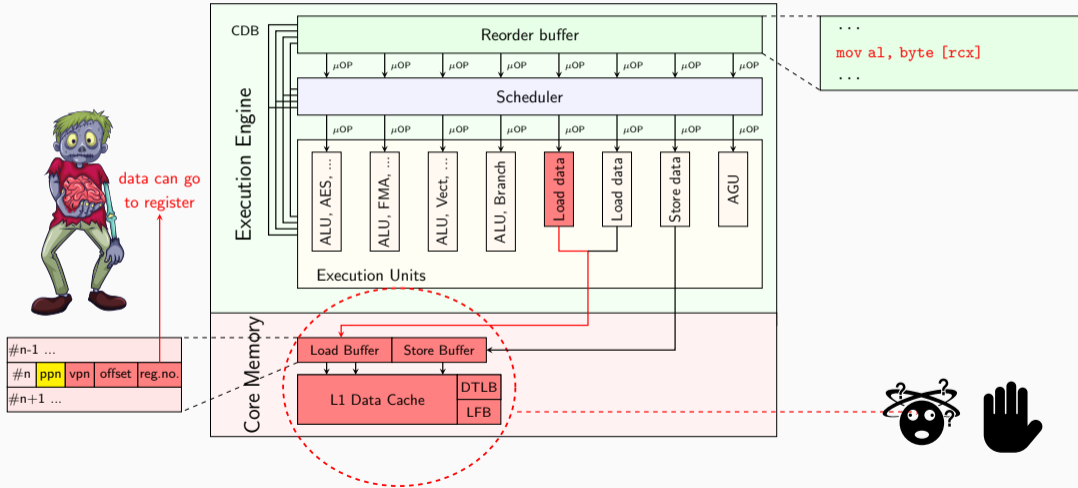
Complex Load Situations



Complex Load Situations



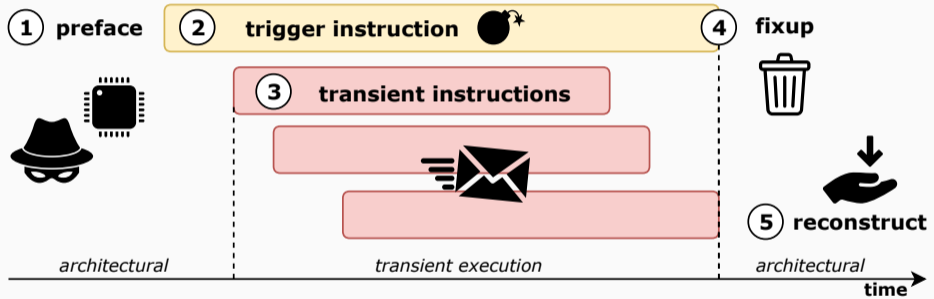
Complex Load Situations





There is no noise.

Noise is just
someone else's data



BBC

Intel Zombieload bug fix to slow data centre computers

THE VERGE

ZombieLoad attack lets hackers steal data from Intel chips

FORTUNE

'Zombieload' Flaw Lets Hackers Crack Almost Every Intel Chip Back to 2011. Why's It Being Downplayed?

How-To Geek

Only New CPUs Can Truly Fix ZombieLoad and Spectre


How it started



How it started



How it's going




I have some sort

**Of idea what I am
doing**



WITH GREAT INSIGHT
COMES GREAT AUTOMATION






- Many **Microarchitectural Data Sampling** (MDS) attacks
→ ZombieLoad, RIDL, Fallout, Meltdown-UC
- Different **variants** and leakage targets
- **Complex** to reproduce and test all variations
- Common: require a fault or microcode assist

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]
```




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]
```

⚡ Fault



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]
```




```
mem[value]
```



Fault

Out of order




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]
```


```
mem[value]
```

K



Fault

Out of order



DUMB WAYS to DIE™

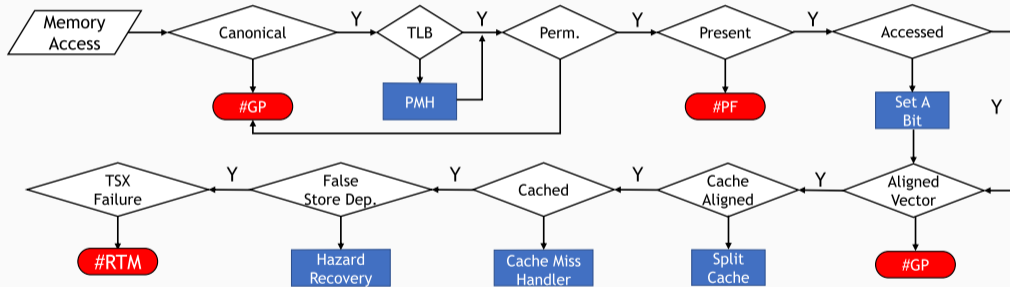


Memory Access Checks (simplified)


- Many possibilities for faults

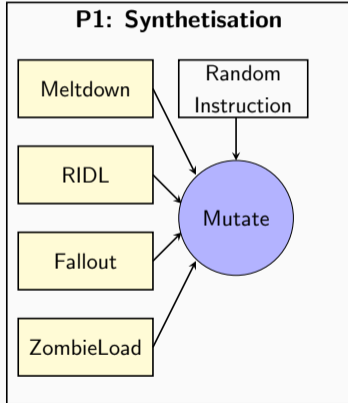
Memory Access Checks (simplified)

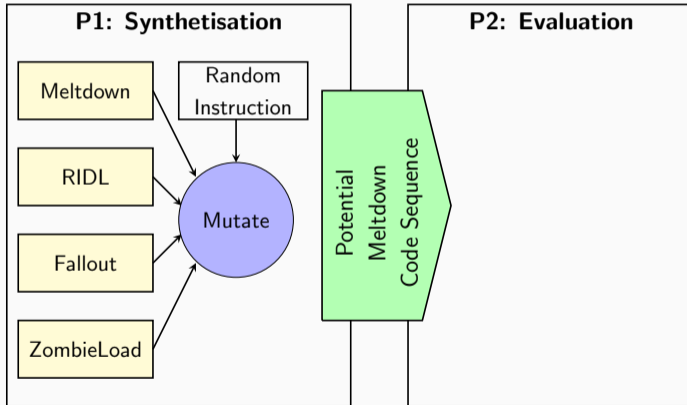
- Many possibilities for faults




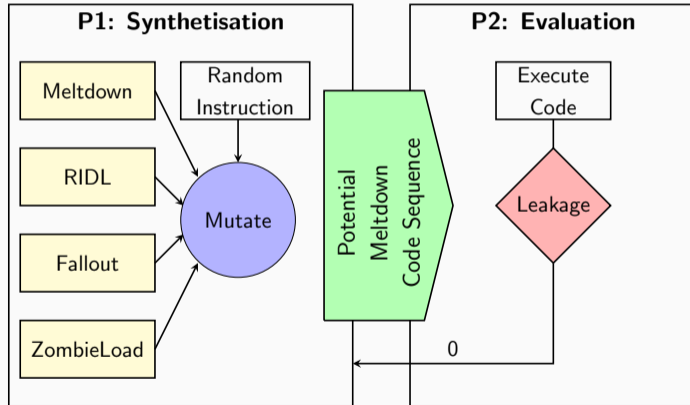
- Idea: mutation fuzzing for new variants




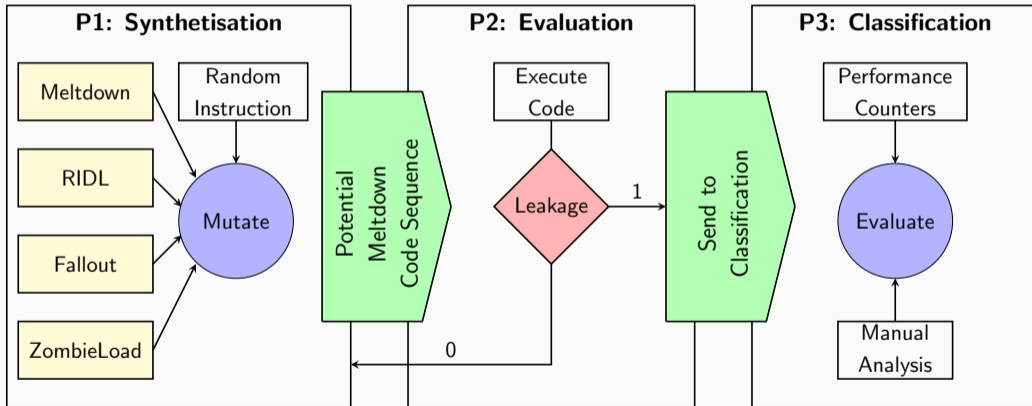













Transynther Results




26 hours runtime



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
 - Used for fast **memory copy**, e.g., in OpenSSL or kernel
- Leaked RSA key while decoding in OpenSSL



- Ice Lake microarchitecture reported **no vulnerabilities**
 - Transynther found a **regression** via a small mutation
- **Re-enabled** a “mitigated” variant
- Fixed via **microcode** update



ARCHITECTURE

STRANGER
-THINGS-

MICROARCHITECTURE

CONTENTION FAULTS

INTERRUPT CONFLICT



After all... why not?



**WHY SHOULDN'T I LOOK
AT ARCHITECTURAL INTERFACES?**

- Goal: disable a prefetcher for an experiment



- Goal: disable a prefetcher for an experiment

```
% sudo wrmsr -a 320 0x2
```





- Goal: disable a prefetcher for an experiment

```
% sudo wrmsr -a 320 0x2
```

- Typo: should be 420, not 320



- Goal: disable a prefetcher for an experiment

```
% sudo wrmsr -a 320 0x2
```

- Typo: should be 420, not 320

```
% sudo wrmsr -a 420 0x2  
Segmentation fault (core dumped)  
% top  
Segmentation fault (core dumped)  
% sudo reboot  
Segmentation fault (core dumped)
```




- Goal: disable a prefetcher for an experiment

```
% sudo wrmsr -a 320 0x2
```

- Typo: should be 420, not 320

```
% sudo wrmsr -a 420 0x2  
Segmentation fault (core dumped)  
% top  
Segmentation fault (core dumped)  
% sudo reboot  
Segmentation fault (core dumped)
```

→ Every application crashed on startup




- Goal: disable a prefetcher for an experiment

```
% sudo wrmsr -a 320 0x2
```


- Typo: should be 420, not 320

```
% sudo wrmsr -a 420 0x2  
Segmentation fault (core dumped)  
% top  
Segmentation fault (core dumped)  
% sudo reboot  
Segmentation fault (core dumped)
```


- Every application crashed on startup
- No information about this MSR in the Intel manual




- Model-Specific Registers (MSRs): CPU control registers





- Model-Specific Registers (MSRs): CPU control registers
- AMD K8: Undocumented debug mode via MSR

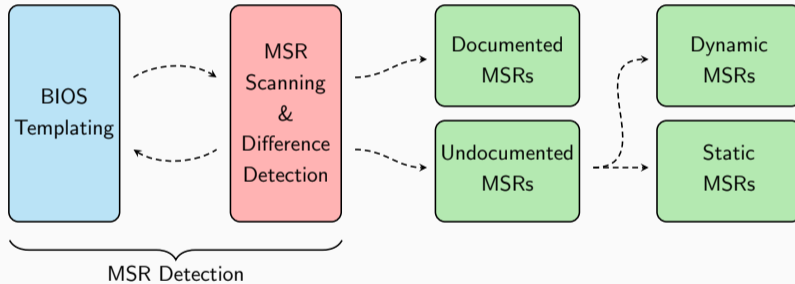


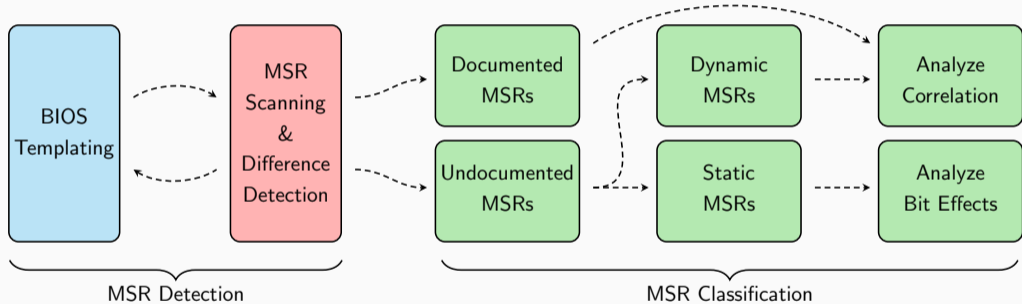
- Model-Specific Registers (MSRs): CPU control registers
- AMD K8: Undocumented `debug` mode via MSR
- VIA C3 CPU: `backdoor` access via undocumented MSR bit

















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 MSR



3 MSRs as attack mitigation



1 MSR allows TOCTOU vulnerability




New MSRs hint towards vulnerabilities




- MSRs often introduce **vulnerability fixes**




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



- Mitigate prefetch side-channel attacks



- 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

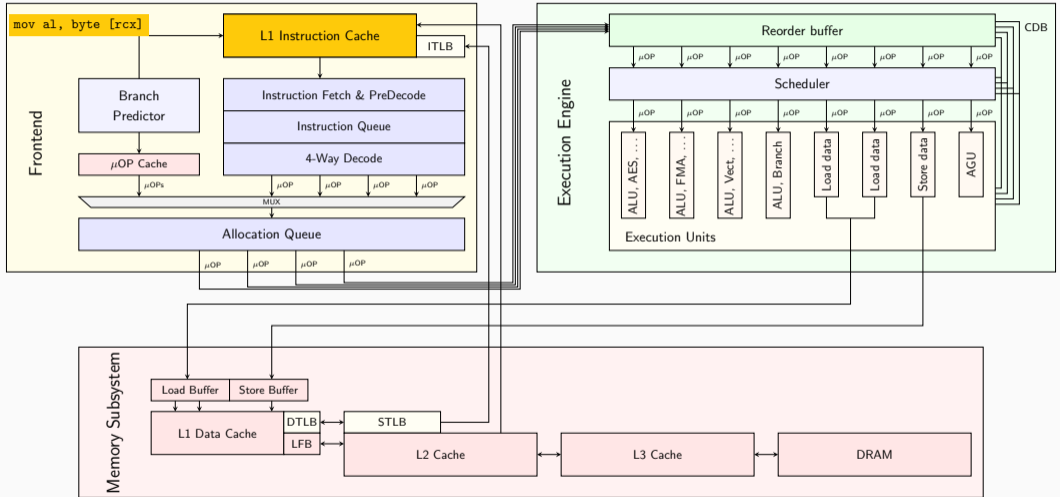


- Searching for **unknown behavior** is hard
- We can **automate the search** for undocumented MSR behavior
- Automation allows **tracing changes between releases**

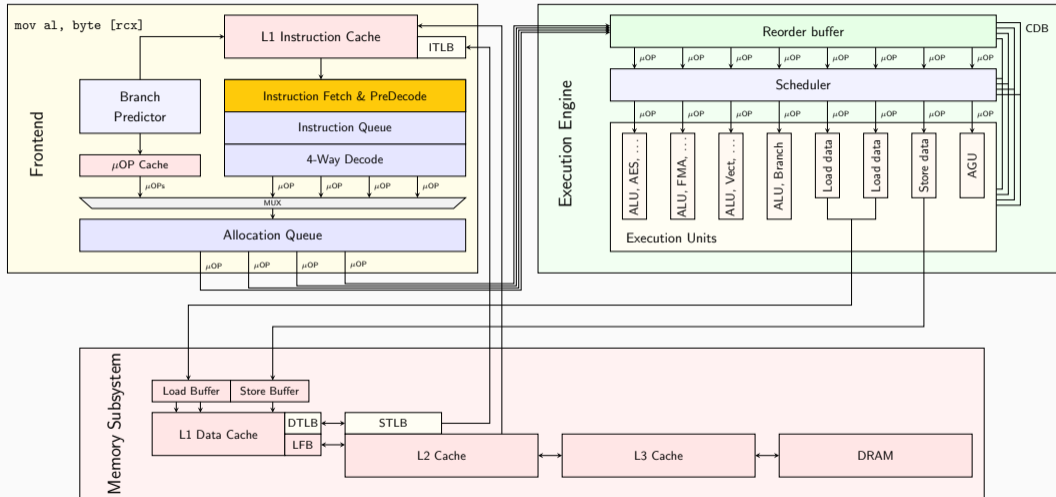


- Searching for **unknown behavior** is hard
- We can **automate the search** for undocumented MSR behavior
- Automation allows **tracing changes between releases**
- **Architectural interfaces** not fully explored

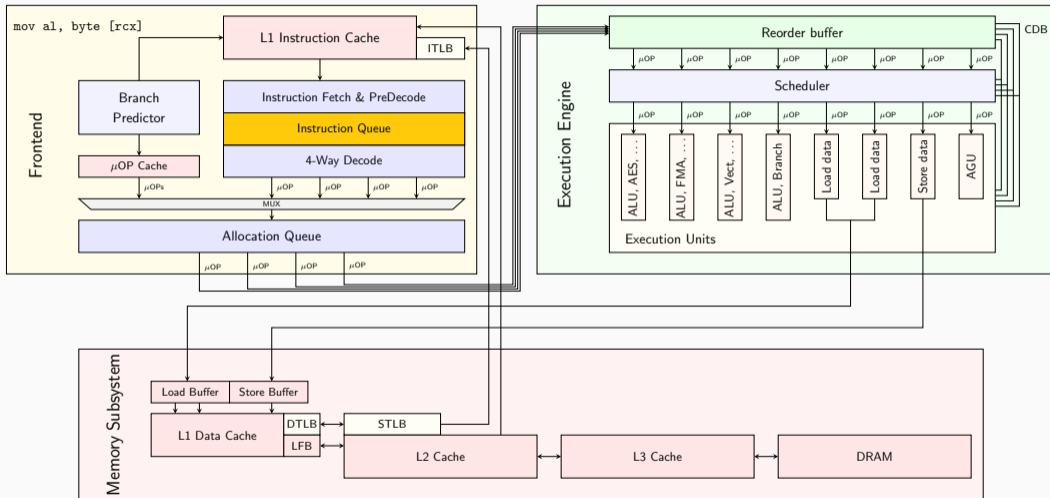
Loads are Complex



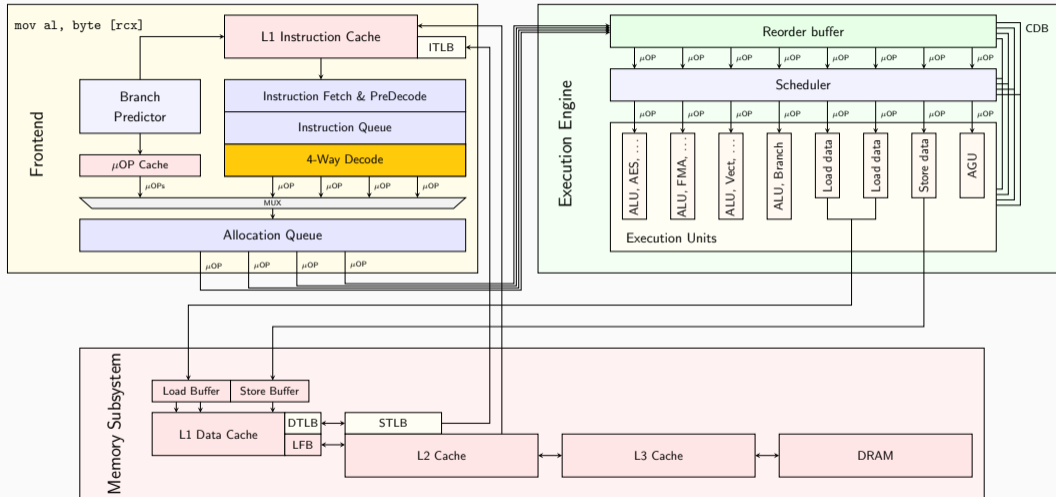
Loads are Complex



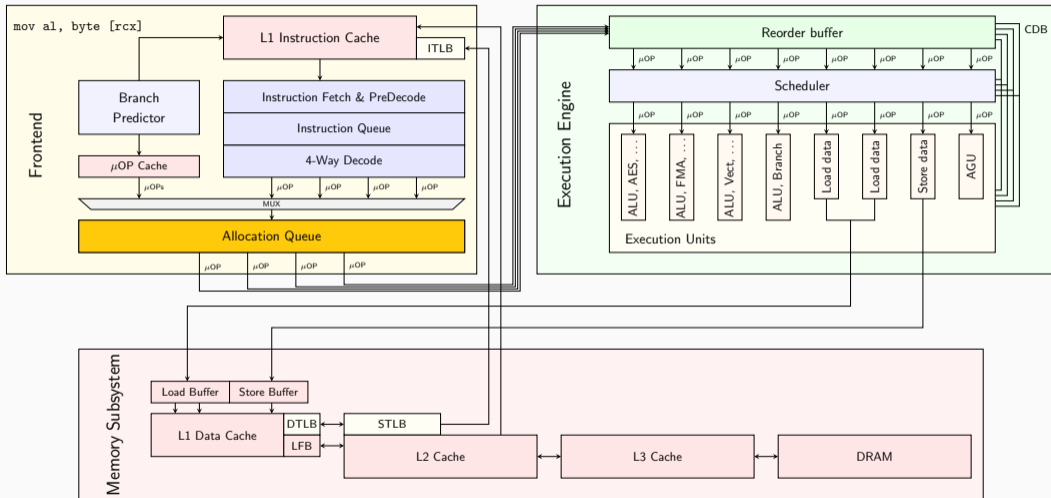
Loads are Complex



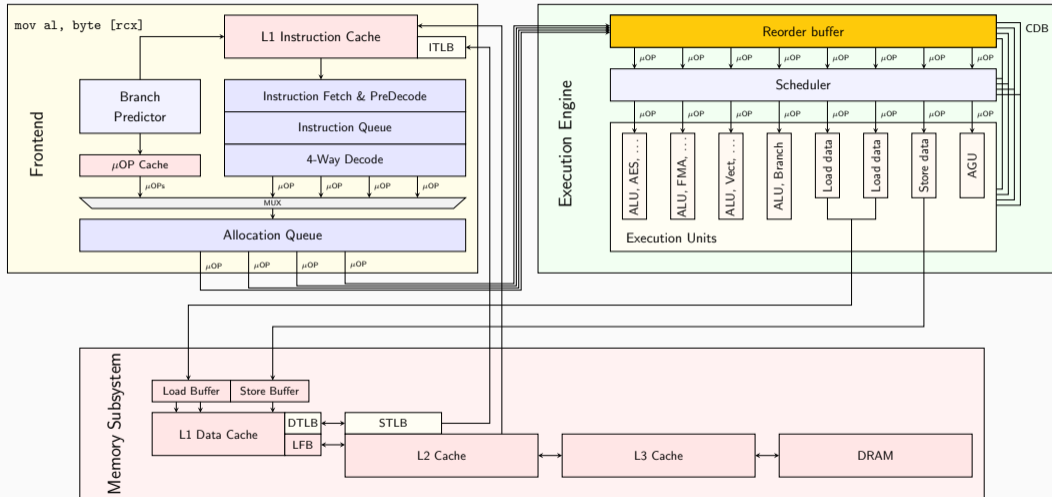
Loads are Complex



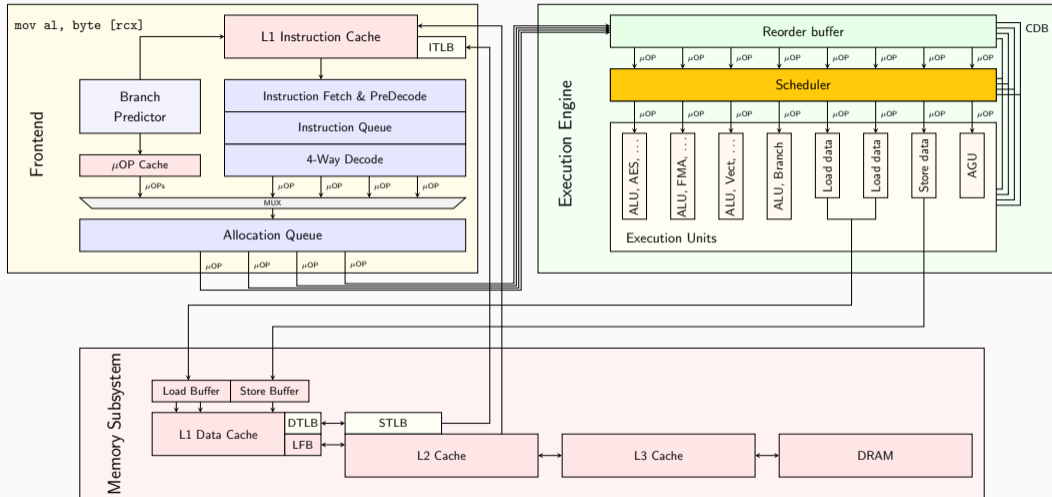
Loads are Complex



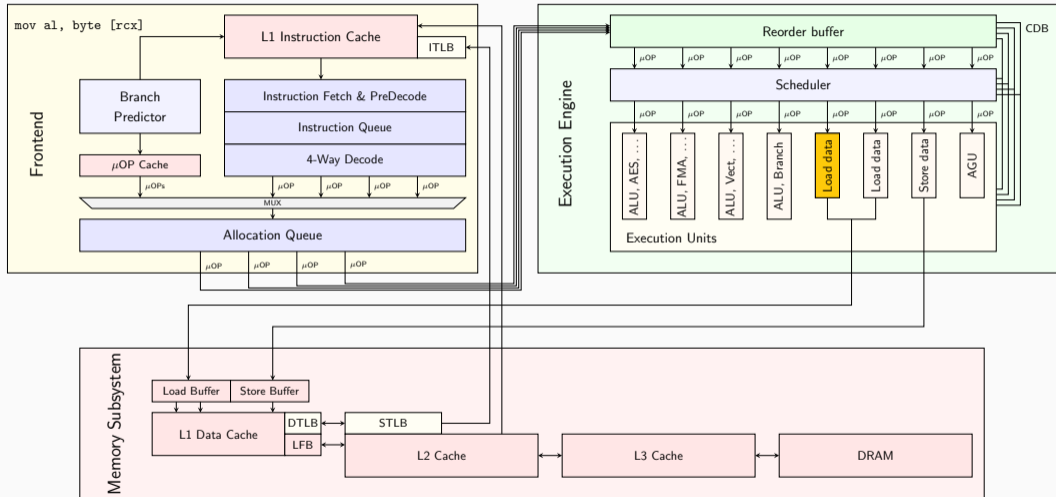
Loads are Complex



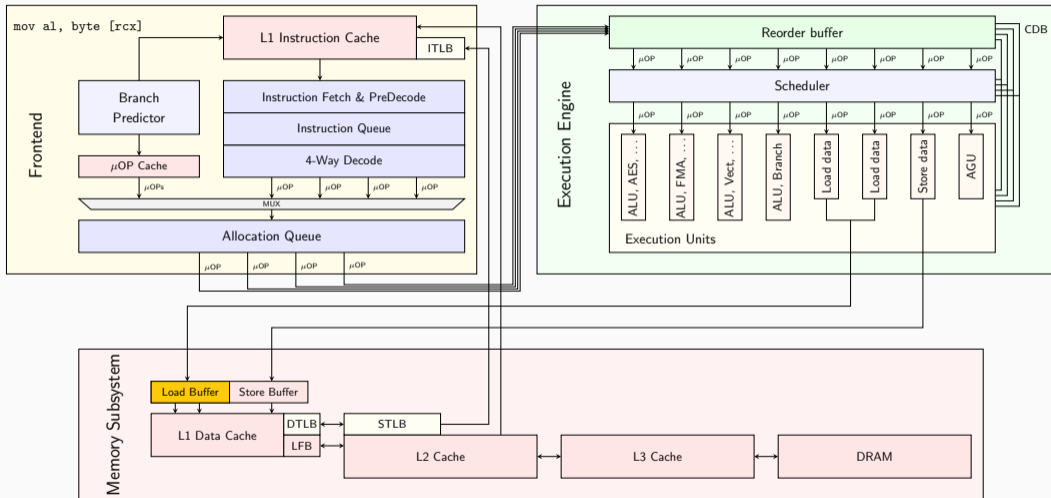
Loads are Complex



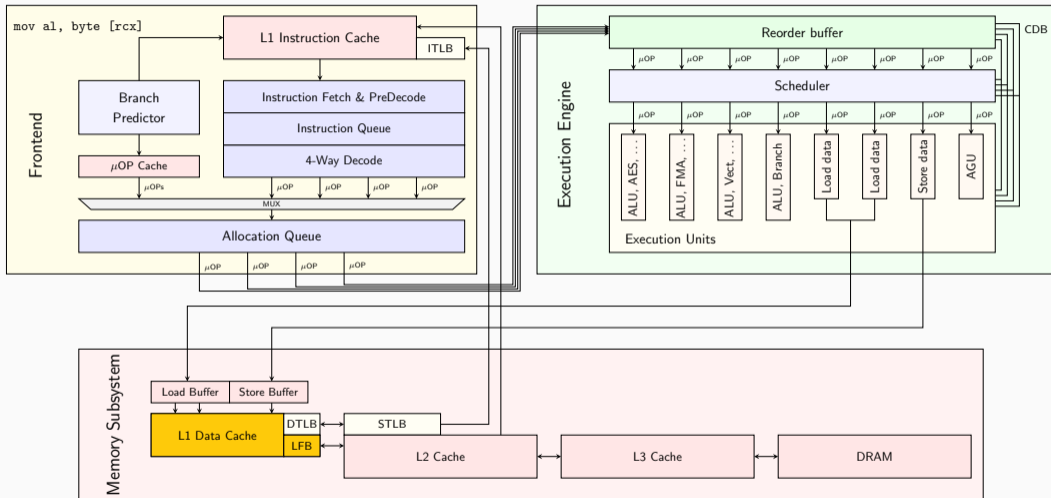
Loads are Complex



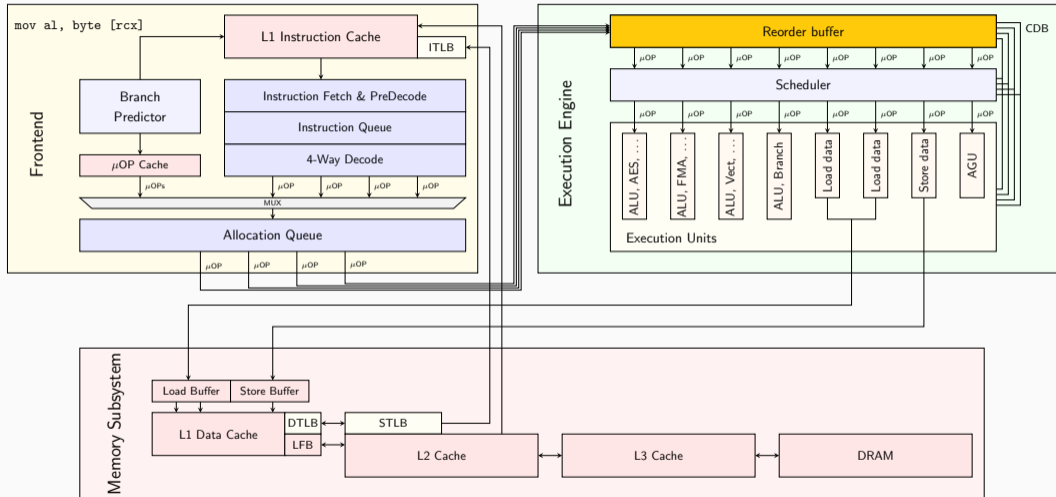
Loads are Complex



Loads are Complex



Loads are Complex



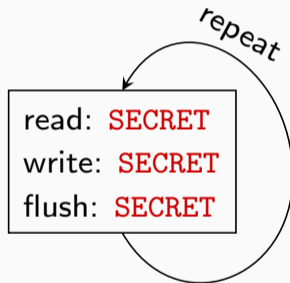
Bunch of Stuff Mapped



```
00001000-0009efff : System RAM
0009f000-000fffff : Reserved
    000a0000-000bffff : PCI Bus 0000:00
00100000-412f6017 : System RAM
45cba000-45cbafff : ACPI Non-volatile Storage
47f00000-64bfffff : Reserved
    61000000-64bfffff : Graphics Stolen Memory
64c00000-bfffffff : PCI Bus 0000:00
    66000000-721fffff : PCI Bus 0000:01
fee00000-fee00fff : Local APIC
    fee00000-fee00fff : Reserved
ff000000-ffffffff : Reserved
    ff000000-ffffffff : pnp 00:04
10000000-49b3ffff : System RAM
    1ad60000-1ae400df0 : Kernel code
    1ae400df1-1af25687f : Kernel data
    1af52a000-1af9fffff : Kernel bss
49b400000-49bfffff : RAM buffer
[...]
```

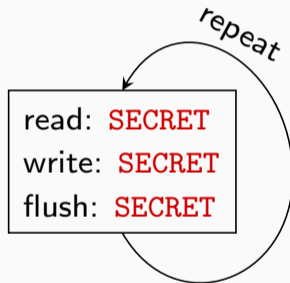

Scanning Memory

Physical Memory

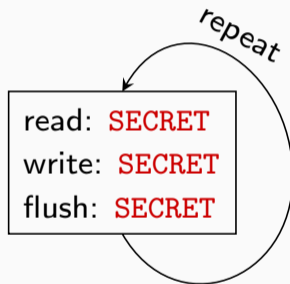



Scanning Memory

Physical Memory

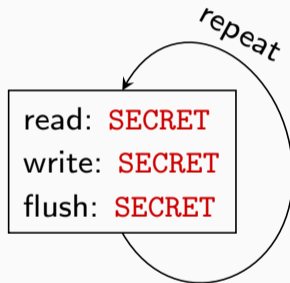



Physical Memory



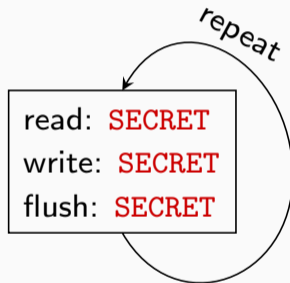

Scanning Memory

Physical Memory



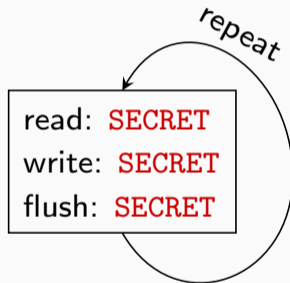

Scanning Memory

Physical Memory



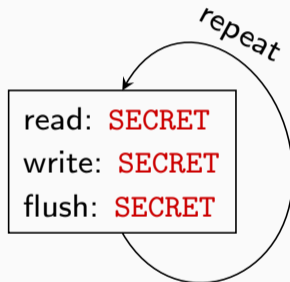
Scanning Memory

Physical Memory



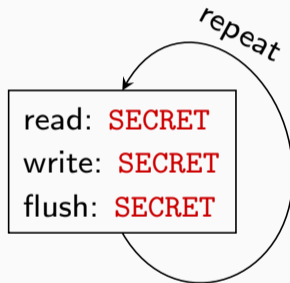
Scanning Memory


Physical Memory



Scanning Memory


Physical Memory







```
00001000-0009efff : System RAM
0009f000-000fffff : Reserved
    000a0000-000bffff : PCI Bus 0000:00
00100000-412f6017 : System RAM
45cba000-45cbafff : ACPI Non-volatile Storage
47f00000-64bfffff : Reserved
    61000000-64bfffff : Graphics Stolen Memory
64c00000-bfffffff : PCI Bus 0000:00
    66000000-721fffff : PCI Bus 0000:01
fee00000-fee00fff : Local APIC
    fee00000-fee00fff : Reserved
ff000000-ffffffff : Reserved
    ff000000-ffffffff : pnp 00:04
10000000-49b3ffff : System RAM
    1ad60000-1ae400df0 : Kernel code
    1ae400df1-1af25687f : Kernel data
    1af52a000-1af9fffff : Kernel bss
49b400000-49bfffff : RAM buffer
[...]
```



Generate, receive and forward interrupts in modern CPUs.

- Local APIC for each CPU
- I/O APIC towards external devices
- Exposes registers

- **Memory-mapped APIC registers**

Timer					0x00
Thermal					0x10
ICR bits 0-31					0x20
ICR bits 32-63					0x30
	0	4	8	12	

- **Memory-mapped APIC registers**
 - Controlled by MSR IA32_APIC_BASE (default 0xFEE00000)

0xFEE00000:


Timer					0x00
Thermal					0x10
ICR bits 0-31					0x20
ICR bits 32-63					0x30
	0	4	8	12	

- **Memory-mapped APIC registers**
 - Controlled by MSR IA32_APIC_BASE (default 0xFEE00000)
 - Mapped as **32bit** values, **aligned to 16 bytes**

0xFEE00000:

Timer	0	4	8	12	0x00
Thermal	0	4	8	12	0x10
ICR bits 0-31	0	4	8	12	0x20
ICR bits 32-63	0	4	8	12	0x30

- **Memory-mapped APIC registers**
 - Controlled by MSR IA32_APIC_BASE (default 0xFEE00000)
 - Mapped as **32bit** values, **aligned to 16 bytes**
 - **Should not** be accessed at bytes 4 through 15.




Any access that touches bytes 4 through 15 of an APIC register may cause undefined behavior and must not be executed. This undefined behavior could include hangs, incorrect results, or unexpected exceptions.

Any access that touches bytes 4 through 15 of an APIC register may cause undefined behavior and must not be executed. This undefined behavior could include hangs, incorrect results, or unexpected exceptions.







Ruling out Microarchitectural Elements




Ruling out Microarchitectural Elements




Ruling out Microarchitectural Elements




Ruling out Microarchitectural Elements




Ruling out Microarchitectural Elements






- It's the **decoupling buffer** between L2 and LLC



- It's the **decoupling buffer** between L2 and LLC
- Contains **data** passed between L2 and LLC



- It's the **decoupling buffer** between L2 and LLC
- Contains **data** passed between L2 and LLC
- Like **Line Fill Buffers** for L1 and L2


Leaking from the Superqueue

APIC

EOI	???
ISR	???
IRR	???

Attacker

Victim (SGX)




Leaking from the Superqueue



APIC

EOI	???
ISR	???
IRR	???


Attacker



Leaking from the Superqueue



Leaking from the Superqueue





- First architectural CPU vulnerability
- Deterministically leak data from SGX
- Does not require hyperthreading



- All **low-hanging fruit**
- Approximately as sophisticated as software fuzzing in 1990
- Majority of fuzzers does **not** use **any guidance**
- More research on **feedback** necessary



- Simple models are sufficient to find leakage
- Dumb fuzzers find leakage within hours
 - New vulnerability variants
 - New side channels
 - Regression in new CPUs
 - New architectural vulnerabilities
- Prediction: smarter fuzzers → more vulnerabilities

<https://github.com/CISPA/Osiris>

 **USENIX'21**

D. Weber, A. Ibrahim, H. Nemati, M. Schwarz, C. Rossow.
Osiris: Automated Discovery of Microarchitectural Side Channels.



<https://github.com/vernamlab/Medusa>

 **USENIX'20**

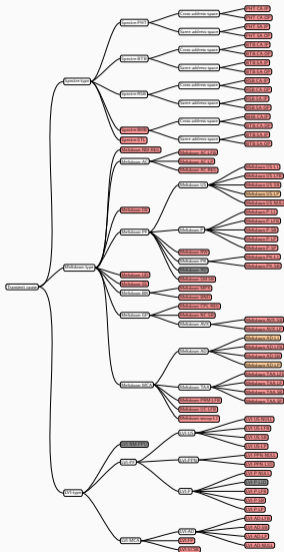
D. Moghimi, M. Lipp, B. Sunar, M. Schwarz.
Medusa: Microarchitectural Data Leakage via Automated Attack Synthesis.

<https://github.com/IAIK/MSRevelio>

 **USENIX'22**

A. Kogler, D. Weber, M. Haubenwallner, M. Lipp, D. Gruss, M. Schwarz.
Finding and Exploiting CPU Features using MSR Templating.

Spectre, Meltdown, and LVI Variants





Twittern



Josh Walden @jmw1123 · 19. Nov.

Case of beer on it's way/there later this week thanks Daniel! Thanks again for the partnership!



Daniel Gruss @lavados · 13. Nov.

Antwort an @Desertrold und @jmw1123

I'm in favor!





Daniel Gruss
@lavados

Antwort an @jmw1123

Thanks again Josh!

We already received the case a month ago but only found time this weekend to sit together and enjoy some!

We wish you a merry Christmas and look forward to continue working with Intel next year.

cc @cc0x1f @mlqxyz @misc0110 @tugraz_csbme #tugraz

[Tweet übersetzen](#)



Du und Claudio Canella

5:45 nachm. · 24. Dez. 2019 · [Twitter Web App](#)

23 „Gefällt mir“-Angaben



FUZZ



ALL THE THINGS



From Random Observations to Automated Leakage Discovery

Michael Schwarz

December 2022

CISPA Helmholtz Center for Information Security