


Transient Execution Attacks

Exploiting the CPU's Microarchitecture

Michael Schwarz (@misc0110), Moritz Lipp (@mlqxyz)

October 7, 2019

Graz University of Technology



Michael Schwarz

PhD candidate @ Graz University of Technology

🐦 @misc0110

✉ michael.schwarz@iaik.tugraz.at



Moritz Lipp

PhD candidate @ Graz University of Technology

🐦 @mlqxyz

✉ mail@mlq.me

Let's Read Kernel Memory from User Space!



- Find something human readable, e.g., the Linux version

```
# sudo grep linux_banner /proc/kallsyms  
ffffffff81a000e0 R linux_banner
```



```
char data = *(char*) 0xffffffff81a000e0;  
printf("%c\n", data);
```

- Compile and run



```
segfault at ffffffff81a000e0
ip 000000000400535
sp 00007ffce4a80610
error 5 in reader
```



- Compile and run

```
segfault at ffffffff81a000e0
ip 0000000000400535
sp 00007ffce4a80610
error 5 in reader
```

- Kernel addresses are of course **not accessible**





- Compile and run

```
segfault at ffffffff81a000e0
ip 0000000000400535
sp 00007ffce4a80610
error 5 in reader
```


- Kernel addresses are of course **not accessible**
- Any invalid access throws an exception → **segmentation fault**

- Kernel is isolated from user space
- **Isolation**: combination of hardware and software
- Applications cannot access kernel
- Well-defined interface → **system calls**






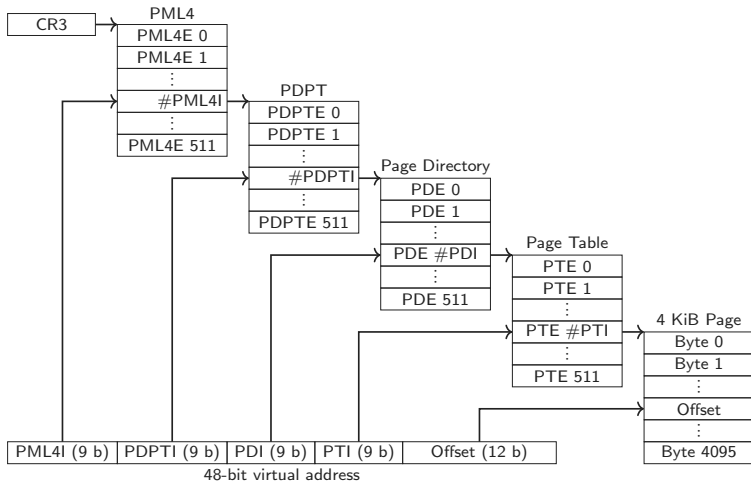
- **Virtual address spaces** isolate processes



- **Virtual address spaces** isolate processes
- Physical memory organized in **page frames**



- **Virtual address spaces** isolate processes
- Physical memory organized in **page frames**
- Virtual memory pages are **mapped** to page frames **using page tables**




P	RW	US	WT	UC	R	D	S	G	Ignored	
Physical Page Number										
									Ignored	X

- User/Supervisor bit defines in which **privilege level** the page can be accessed
















- **Catch** the segmentation fault!




- **Catch** the segmentation fault!
- Install a signal handler




- **Catch** the segmentation fault!
- Install a signal handler
- On exception → jump back and continue



- Still no kernel memory



- Still no kernel memory
- Privilege checks seem to work



- Still no kernel memory
- Privilege checks seem to work
- **Back to the drawing board**



- We cannot see architectural changes





- We cannot see **architectural** changes
- What about the CPU internals?





- We cannot see **architectural** changes
 - What about the CPU internals?
- The **microarchitectural** state


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




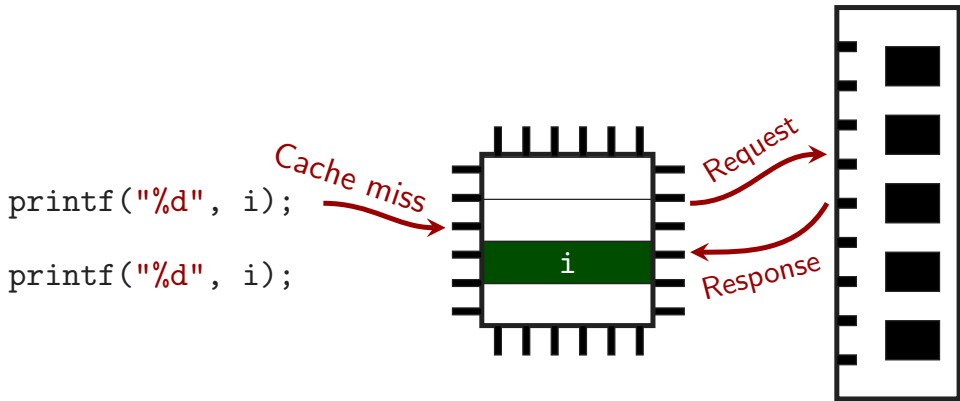
```
printf("%d", i);  
printf("%d", i);
```


Cache miss

















































- Can we see **illegal loads** in the cache?



- Can we see **illegal loads** in the cache?
- What happens if the load cannot continue?

Out-of-order Execution

```
int width = 10, height = 5;

float diagonal = sqrt(width * width
                      + height * height);
int area = width * height;

printf("Area %d x %d = %d\n", width, height, area);
```


Dependency



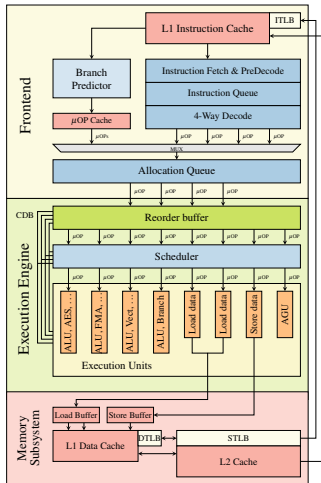
```
int width = 10, height = 5;  
  
float diagonal = sqrt(width * width  
                      + height * height);  
  
int area = width * height;  
  
printf("Area %d x %d = %d\n", width, height, area);
```

Parallelize





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



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

- Adapted code

```
*(volatile char*)0;  
array[84 * 4096] = 0;
```





- Adapted code

```
*(volatile char*)0;  
array[84 * 4096] = 0;
```

- volatile because compiler was not happy

```
1 warning: statement with no effect [-Wunused-value]  
2     *(char*)0;
```



- Adapted code

```
*(volatile char*)0;  
array[84 * 4096] = 0;
```

- volatile because compiler was not happy


```
1 warning: statement with no effect [-Wunused-value]  
2     *(char*)0;
```

- Static code analyzer is still not happy

```
1 warning: Dereference of null pointer  
2     *(volatile char*)0;
```




- Flush+Reload over all pages of the array



- “Unreachable” code line was actually executed



- Flush+Reload over all pages of the array



- “Unreachable” code line was actually executed
- Exception was only thrown afterwards



- Out-of-order instructions leave microarchitectural traces



- Out-of-order instructions leave microarchitectural traces
- We can see them for example in the cache



- Out-of-order instructions leave microarchitectural traces
- We can see them for example in the cache
- Give such instructions a name: **transient instructions**



- Out-of-order instructions leave microarchitectural traces
- We can see them for example in the cache
- Give such instructions a name: **transient instructions**
- We can indirectly observe the execution of transient instructions



- Maybe there is no permission check in transient instructions...



- Maybe there is no permission check in transient instructions...
- ...or it is only done when committing them



- Maybe there is no permission check in transient instructions...
- ...or it is only done when committing them
- Add another layer of indirection to test

```
char data = *(char*) 0xffffffff81a000e0;  
array[data * 4096] = 0;
```




- Maybe there is no permission check in transient instructions...
- ...or it is only done when committing them
- Add another layer of indirection to test

```
char data = *(char*) 0xffffffff81a000e0;  
array[data * 4096] = 0;
```

- Then check whether any part of array is cached




- Flush+Reload over all pages of the array



- Index of cache hit reveals data



- Flush+Reload over all pages of the array



- Index of cache hit reveals data
- Permission check is in some cases not fast enough



- Using out-of-order execution, we can read data at any address



- Using out-of-order execution, we can read data at any address
- Privilege checks are sometimes too slow




- Using out-of-order execution, we can read data at any address
- Privilege checks are sometimes too slow
- Allows to leak kernel memory




- Using out-of-order execution, we can read data at any address
- Privilege checks are sometimes too slow
- Allows to leak kernel memory
- Entire physical memory is typically also accessible in kernel address space




- Using out-of-order execution, we can read data at any address
- Privilege checks are sometimes too slow
- Allows to leak kernel memory
- Entire physical memory is typically also accessible in kernel address space
- Works on Intel CPUs and ARM Cortex-A75




- Assumed Meltdown can one only read data **from the L1**




- Assumed Meltdown can one only read data **from the L1**
- Leakage from L3 or memory is **possible**, just slower




- Assumed Meltdown can one only read data **from the L1**
- Leakage from L3 or memory is **possible**, just slower
- Even leakage of **UC (uncachable)** memory regions...




- Assumed Meltdown can one only read data **from the L1**
- Leakage from L3 or memory is **possible**, just slower
- Even leakage of **UC (uncachable)** memory regions...
 - ...if other hyperthread (legally) accesses the data




Kernel View



User View



context switch




- **Linux:** Kernel Page-table Isolation (KPTI)




- **Linux:** Kernel Page-table Isolation (KPTI)
- **Apple:** Released updates




- **Linux:** Kernel Page-table Isolation (KPTI)
- **Apple:** Released updates
- **Windows:** Kernel Virtual Address (KVA) Shadow




- Meltdown **fully mitigated** in software



- Meltdown **fully mitigated** in software
- Problem **seemed** to be solved



- Meltdown **fully mitigated** in software
- Problem **seemed** to be solved
- No attack surface left




- Meltdown **fully mitigated** in software
- Problem **seemed** to be solved
- No attack surface left
- That is what everyone thought




**There are no bugs,
just happy little accidents**






- Meltdown is a whole **category of vulnerabilities**



- Meltdown is a whole **category of vulnerabilities**
- Not only the user-accessible check



- Meltdown is a whole **category of vulnerabilities**
- Not only the user-accessible check
- There are more bits...

P	RW	US	WT	UC	R	D	S	G	Ignored	
Physical Page Number										
									Ignored	X

P	RW	US	WT	UC	R	D	S	G	Ignored	
Physical Page Number										
									Ignored	X

- **Present** bit is the next obvious bit



- An even **worse** bug → Foreshadow-NG/L1TF



- An even **worse** bug → Foreshadow-NG/L1TF
- Exploitable from **VMs**



- An even **worse** bug → Foreshadow-NG/L1TF
- Exploitable from **VMs**
- Allows **leaking** data from the **L1** cache



- An even **worse** bug → Foreshadow-NG/L1TF
- Exploitable from **VMs**
- Allows **leaking** data from the **L1** cache
- Same mechanism as Meltdown



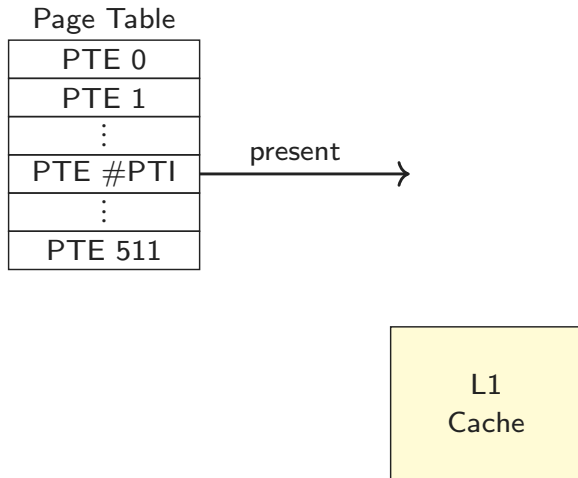
- An even **worse** bug → Foreshadow-NG/L1TF
- Exploitable from **VMs**
- Allows **leaking** data from the **L1** cache
- Same mechanism as Meltdown
- Just a **different bit** in the PTE

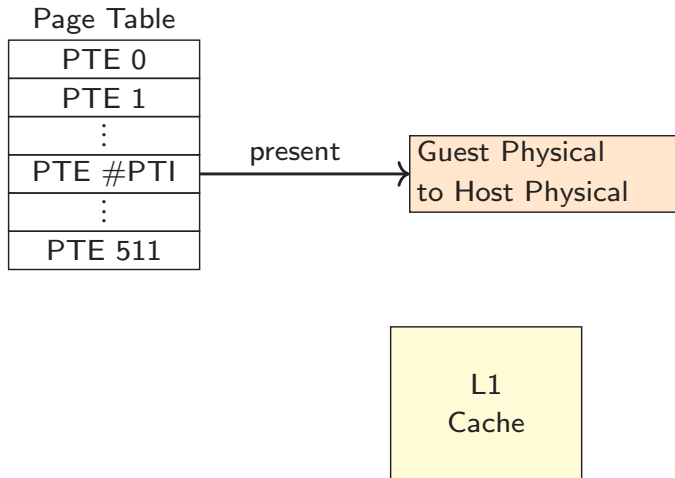
Page Table

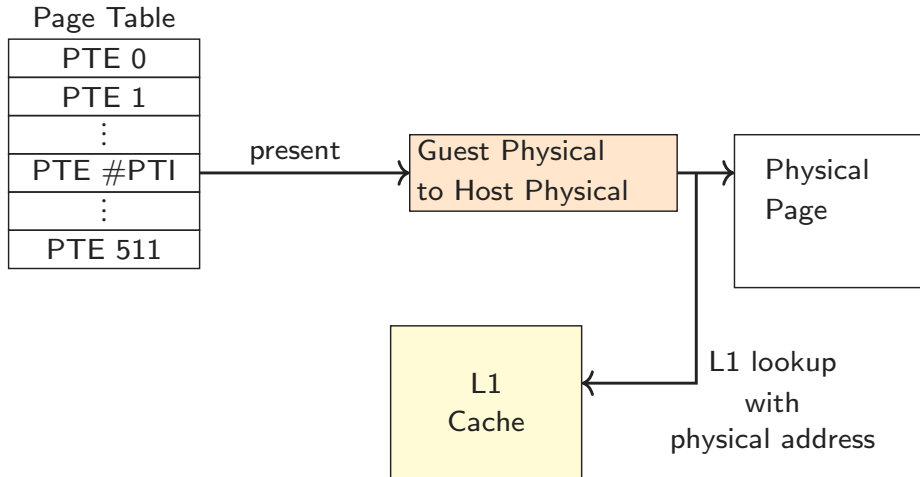
PTE 0
PTE 1
⋮
PTE #PTI
⋮
PTE 511





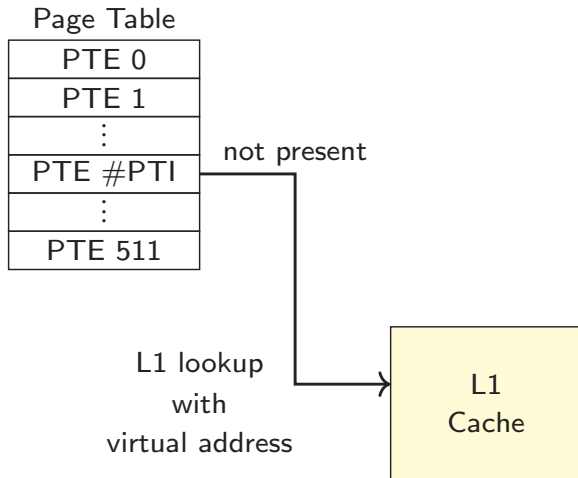
L1
Cache

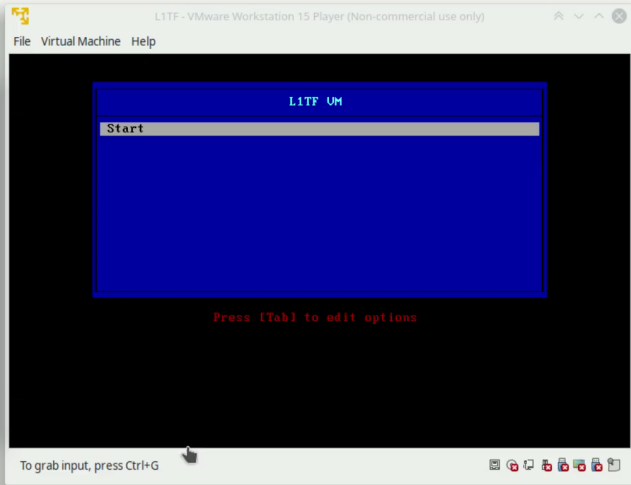














- KAISER/KPTI/KVA does not help



- KAISER/KPTI/KVA does not help
- Only **software workarounds**



- KAISER/KPTI/KVA does not help
- Only **software workarounds**
 - **Flush L1** on VM entry



- KAISER/KPTI/KVA does not help
- Only **software workarounds**
 - **Flush L1** on VM entry
 - Disable **HyperThreading**



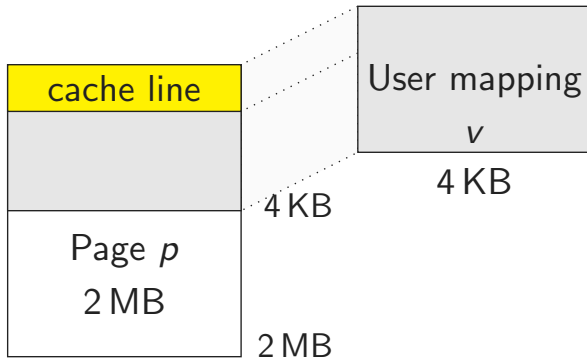
- KAISER/KPTI/KVA does not help
- Only **software workarounds**
 - **Flush L1** on VM entry
 - Disable **HyperThreading**
- Workarounds might not be complete

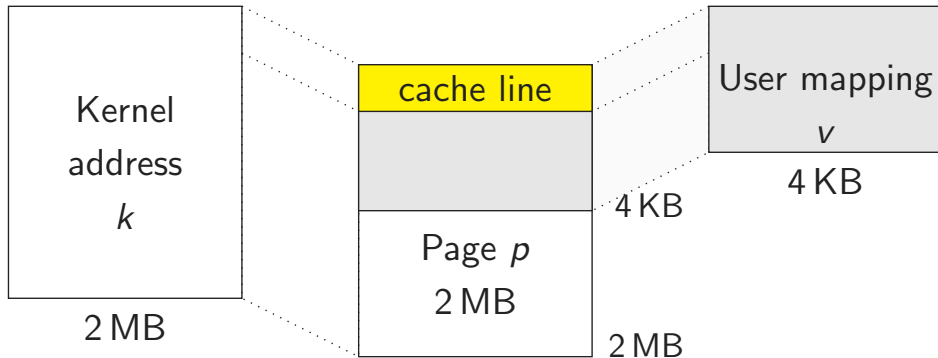



- What if the memory is **not cached**?




- What if the memory is **not cached**?
- No data → no leakage?

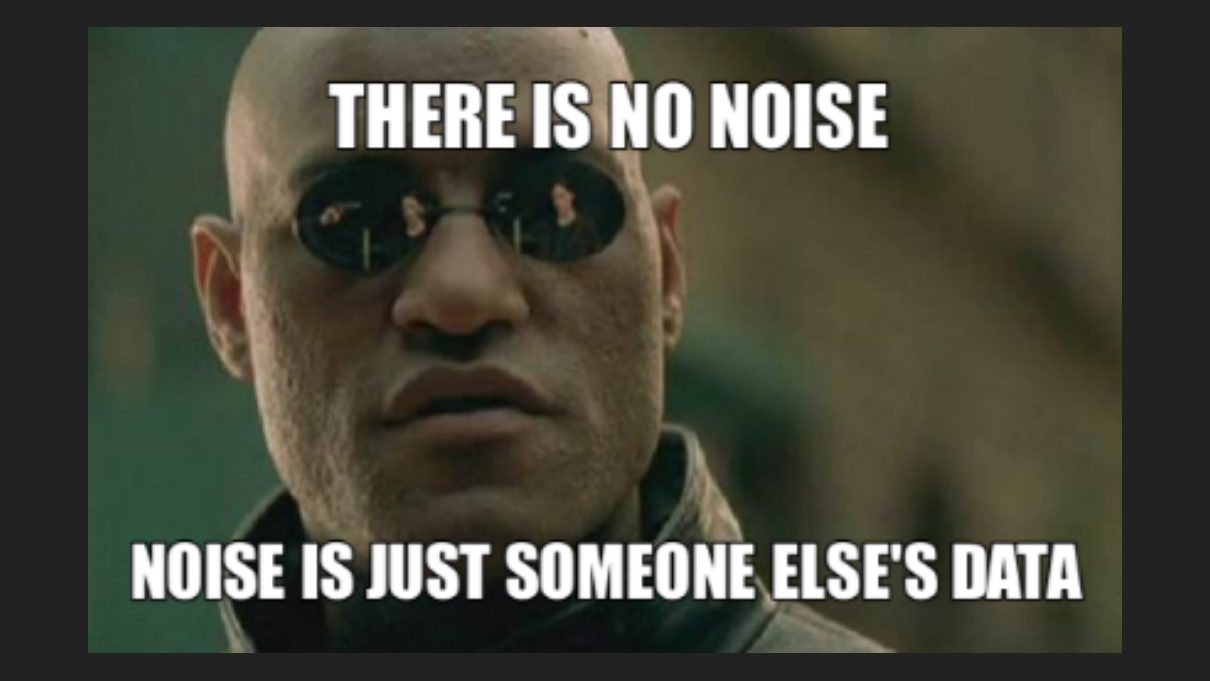






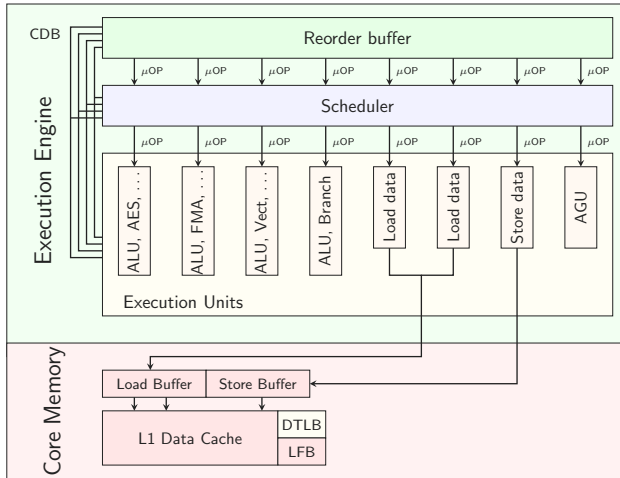


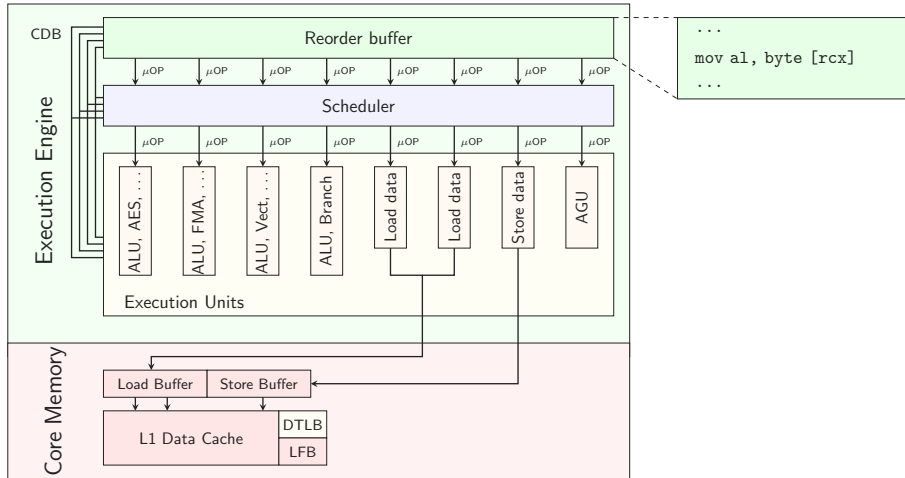


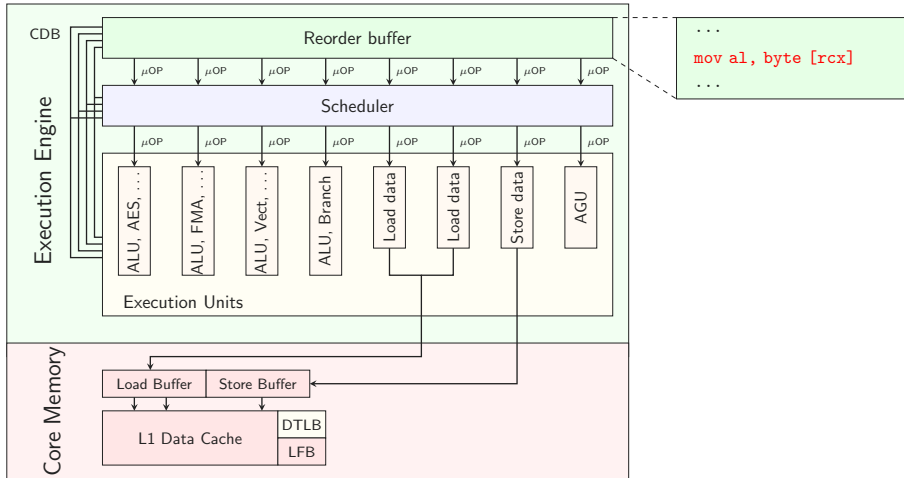
A close-up shot of Morpheus from the movie The Matrix. He is bald, wearing dark sunglasses, and has a serious expression. The background is blurred. The text is overlaid on the image in a white, bold, sans-serif font.

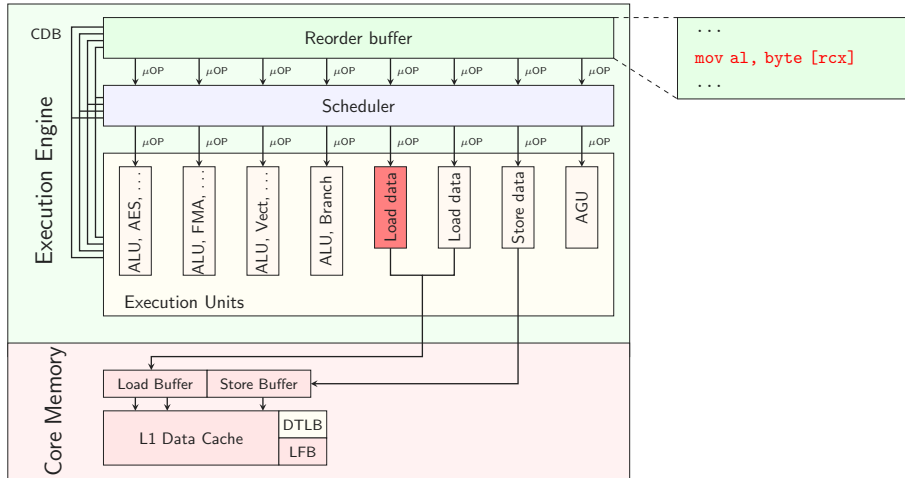
THERE IS NO NOISE

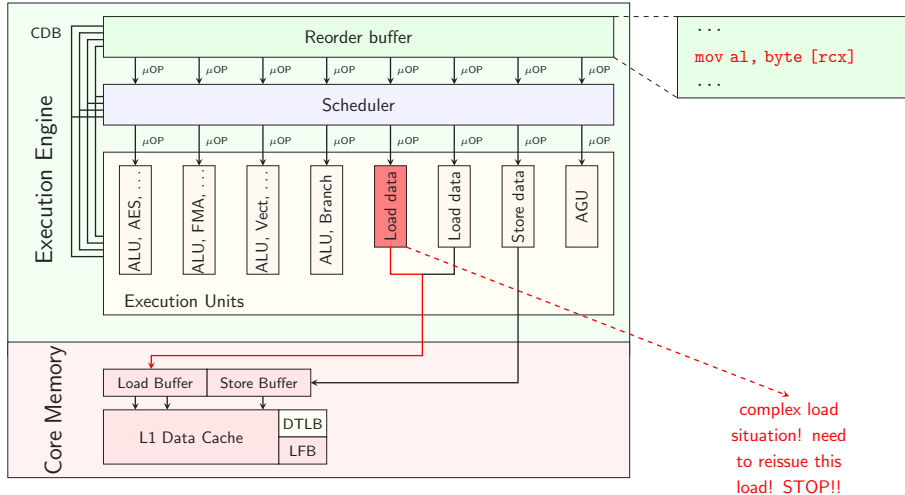
NOISE IS JUST SOMEONE ELSE'S DATA

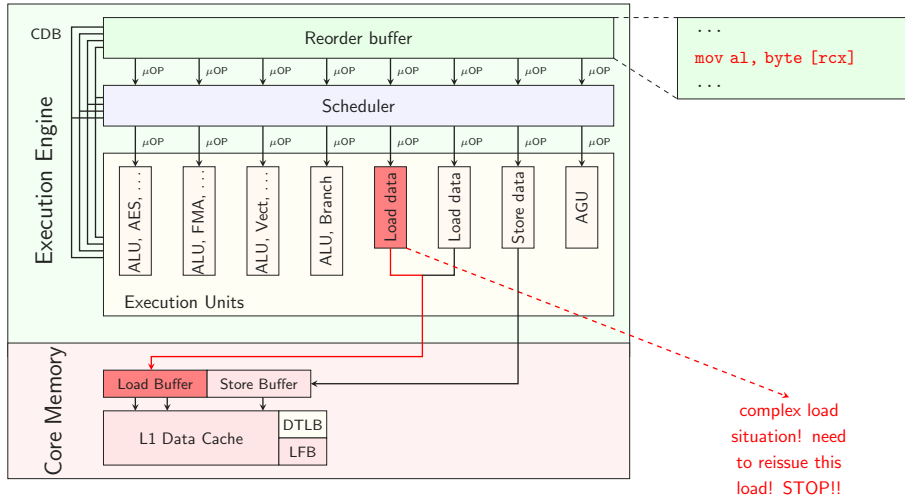


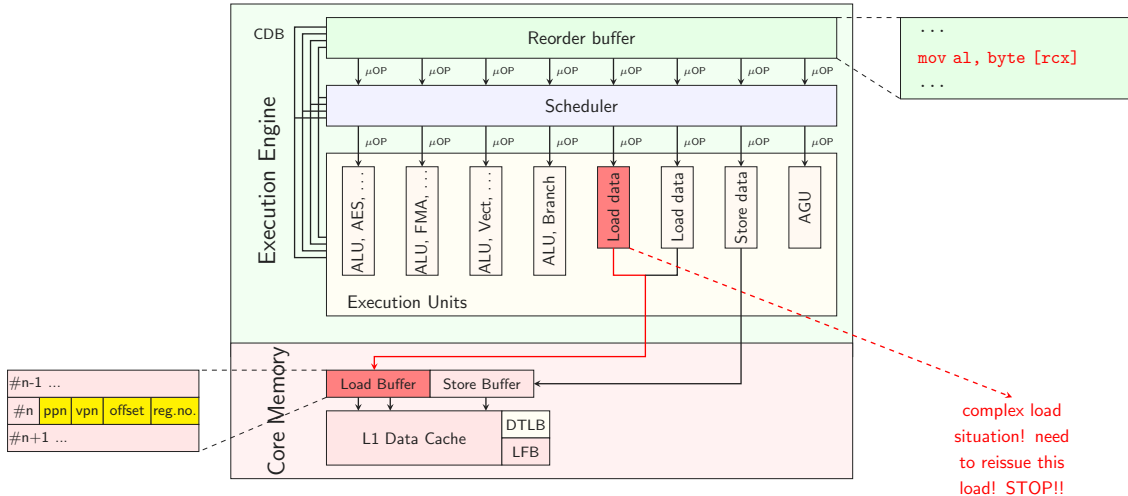


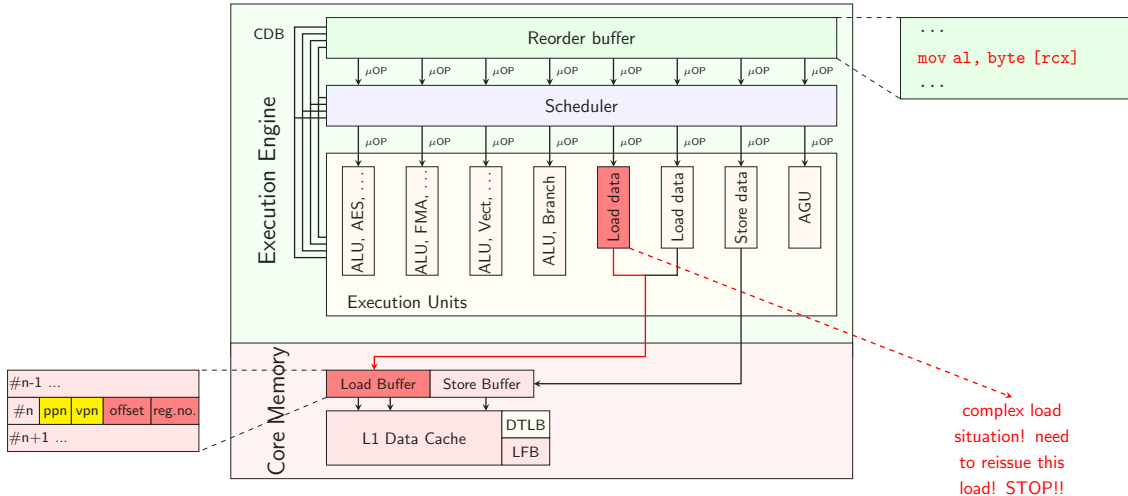


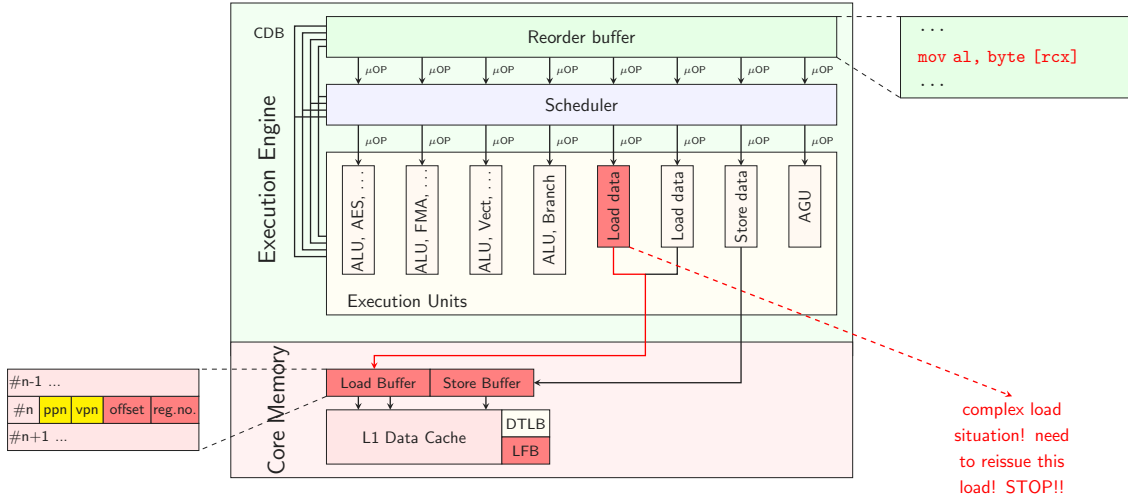















	Page Number			Page Offset			
Meltdown	51	Physical	12	11 0			
	47	Virtual	12				
Foreshadow	51	Physical	12	11 0			
	47	Virtual	12				
Fallout	51	Physical	12	11 0			
	47	Virtual	12				
ZombieLoad/ RIDL	51	Physical	12	11	6	5 0	
	47	Virtual	12				




- Leaks from the **fill buffer**



- Leaks from the **fill buffer**
- Crosses all privilege boundaries (Kernel, VM, SGX)



- Leaks from the **fill buffer**
- Crosses all privilege boundaries (Kernel, VM, SGX)
- Explored microcode assists as new type of faults



- Leaks from the **fill buffer**
- Crosses all privilege boundaries (Kernel, VM, SGX)
- Explored microcode assists as new type of faults
- Disadvantage: **minimal control** over leaked data

```
michael@hp /tmp/zombieload %
```




- Disable HyperThreading



- Disable HyperThreading
- **Microcode** Updates (to an extent)

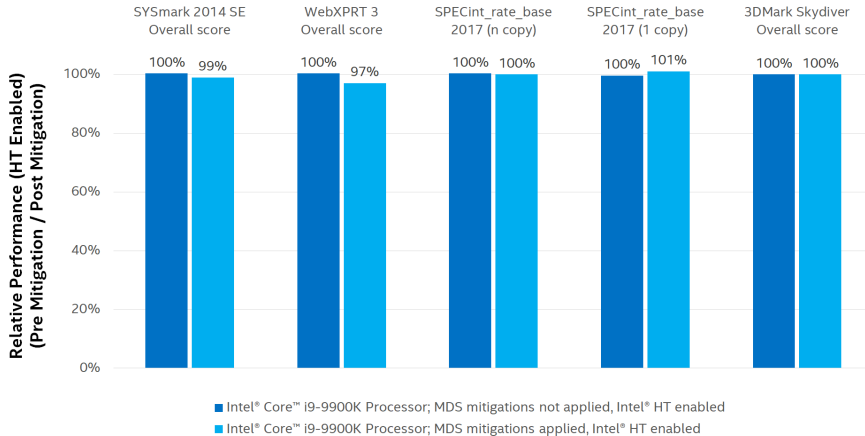


- Disable HyperThreading
- **Microcode** Updates (to an extent)
- VERW instruction → **clear buffers** on context switch




Performance results are based on testing as of May 8, 2018 and may not reflect all publicly available security updates. The configuration disclosure for details. No product or component can be absolutely secure. Software and workloads used in performance tests may have been customized for performance and workload requirements. Performance tests such as SPECint® and SPECint®_rate_base are intended to measure system components, software applications and functions. Any change in any of these factors may affect the results being. You should consult other information and performance.

[View Full Image](#)

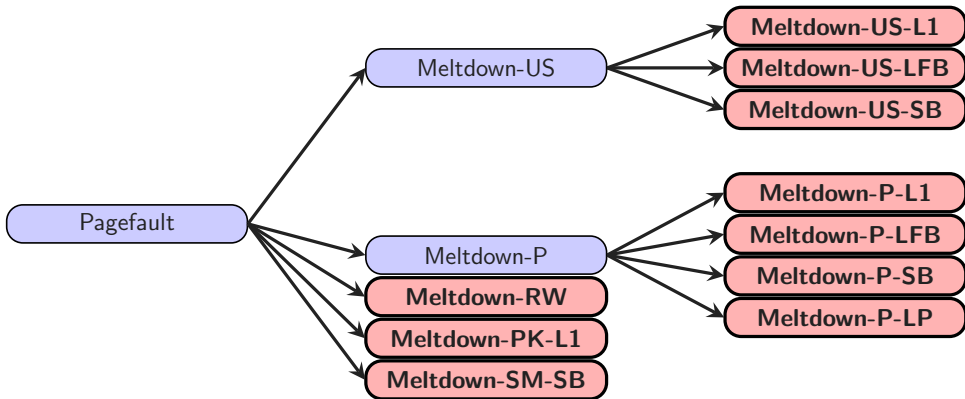


Performance results are based on testing as of May 8, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure. Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark® and MobileMark®, are measured using specific computer systems, components, software, operations and functions. Any change to any of these factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchase, including the performance of that product when combined with other products. For more complete information about performance and benchmark results, visit <http://www.intel.com/benchmarks>.

Software and workloads used in performance tests may have been optimized for the performance only on Intel microprocessors. [. . .] Any chance of any of those factors **cause the results to vary**.




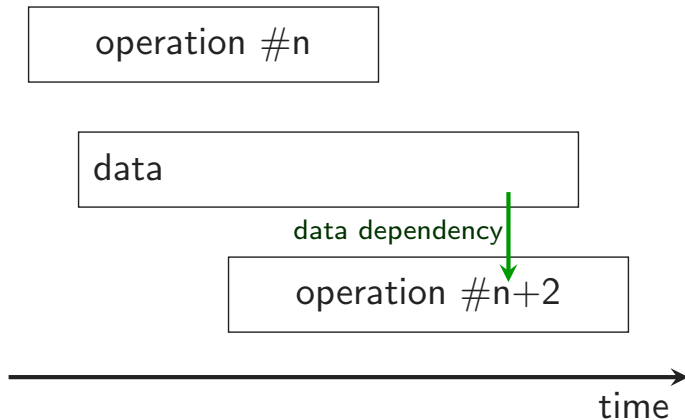
About 16% lower performance on average (phoronix.com)

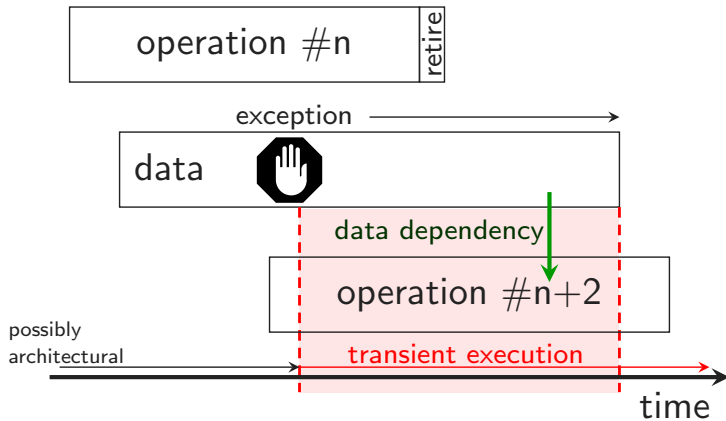


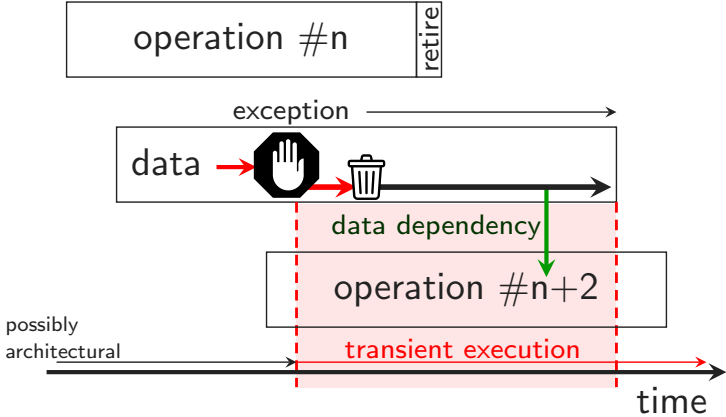
operation #n

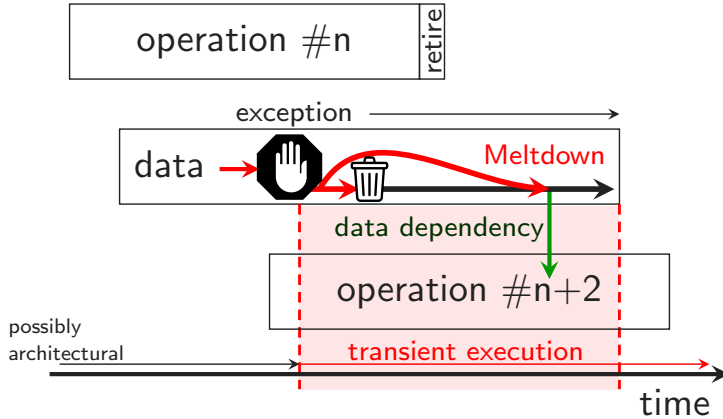


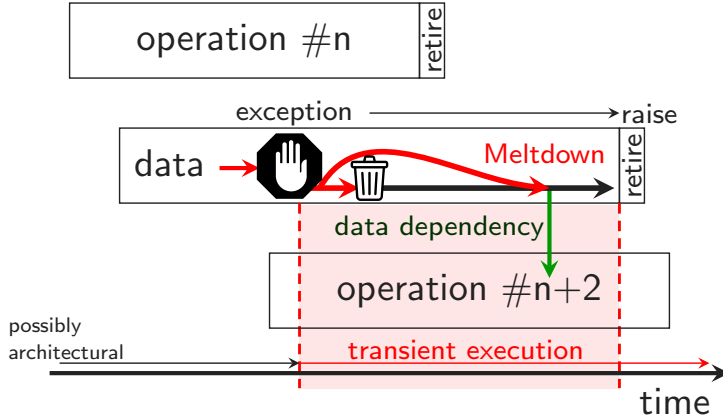


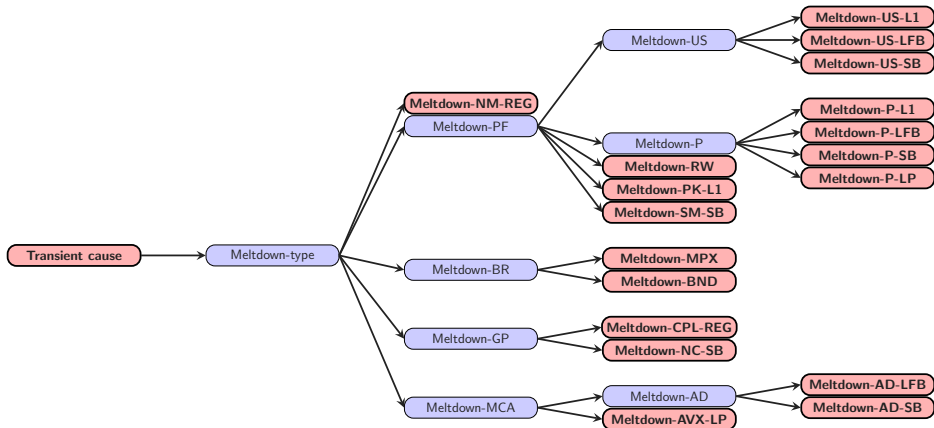














- Meltdown is **not** a fully **solved** issue



- Meltdown is **not** a fully **solved** issue
- The tree is extensible



- Meltdown is **not** a fully **solved** issue
- The tree is extensible
- **More** Meltdown-type **issues** to come



- Meltdown is **not** a fully **solved** issue
- The tree is extensible
- **More** Meltdown-type **issues** to come
- Silicon fixes might not be complete



- Meltdown not the only **transient execution attacks**



- Meltdown not the only **transient execution attacks**
- **Spectre** is a second class of transient execution attacks



- Meltdown not the only **transient execution attacks**
- **Spectre** is a second class of transient execution attacks
- Instead of faults, exploit control (or data) **flow predictions**



- CPU tries to predict the future (branch predictor), ...



- CPU tries to predict the future (branch predictor), ...
 - ...based on events learned in the past



- CPU tries to predict the future (branch predictor), ...
 - ...based on events learned in the past
- **Speculative execution** of instructions



- CPU tries to predict the future (branch predictor), ...
 - ...based on events learned in the past
- **Speculative execution** of instructions
- If the prediction was correct, ...



- CPU tries to predict the future (branch predictor), ...
 - ...based on events learned in the past
- **Speculative execution** of instructions
- If the prediction was correct, ...
 - ...very fast



- CPU tries to predict the future (branch predictor), ...
 - ...based on events learned in the past
- **Speculative execution** of instructions
- If the prediction was correct, ...
 - ...very fast
 - otherwise: Discard results

`index = 0`

Shared 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


```
if (index < 4)
```

then

else

```
glyph[data[index]]
```

```
{
```



Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

index = 0

Shared 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

```
if (index < 4)
```

then

```
glyph[data[index]]
```


else

Speculate

```
{ }
```

Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 0


Shared 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




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	







`index = 1`

Shared 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


```
if (index < 4)
```

then

else

```
glyph[data[index]]
```

```
{
```




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

index = 1


Shared 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



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 1


Shared Memory

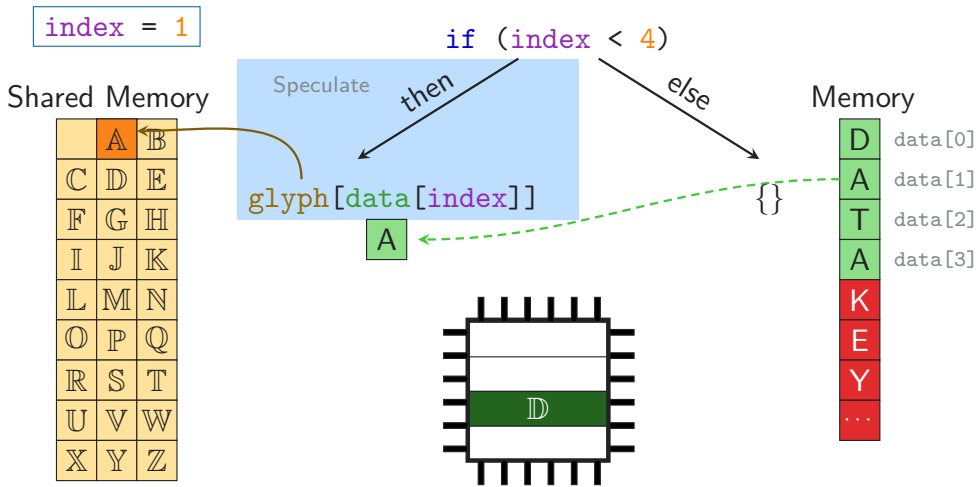
	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	







index = 1

Shared 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

```
if (index < 4)
```

Execute

then


```
glyph[data[index]]
```

else

```
{
```

Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



`index = 2`


Shared 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



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 2


Shared 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




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	








index = 2


Shared 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



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



`index = 3`


Shared 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



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 3

Shared 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

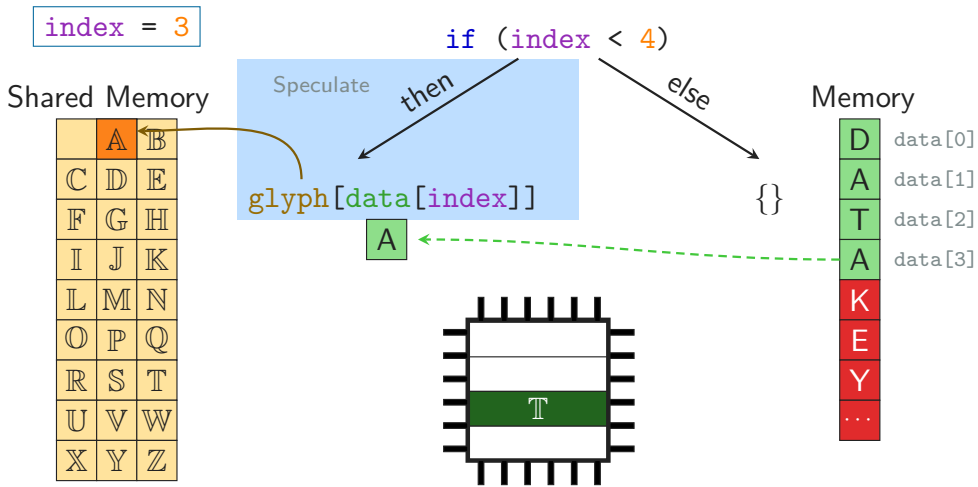



A

Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	









index = 3

Shared 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




Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

`index = 4`


Shared 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



Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	



index = 4

Shared 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

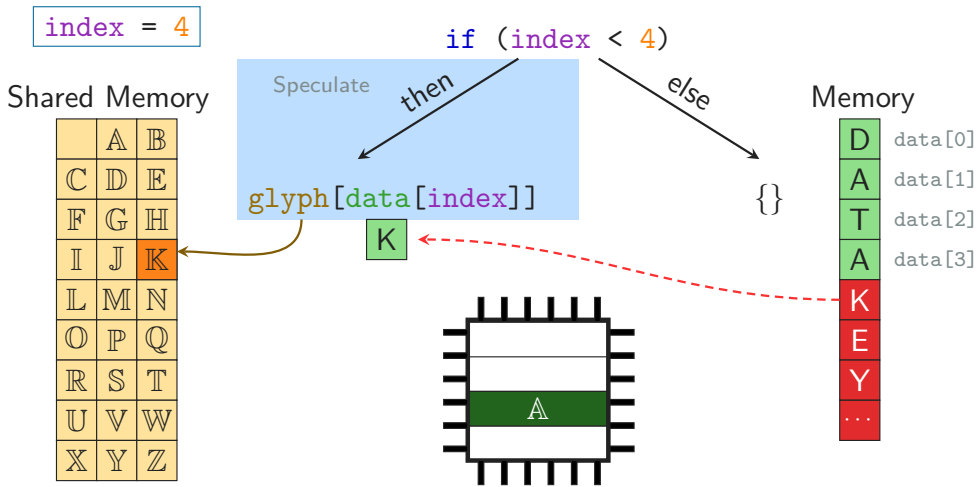



K

Memory

D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	








index = 4


Shared 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




Memory


D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

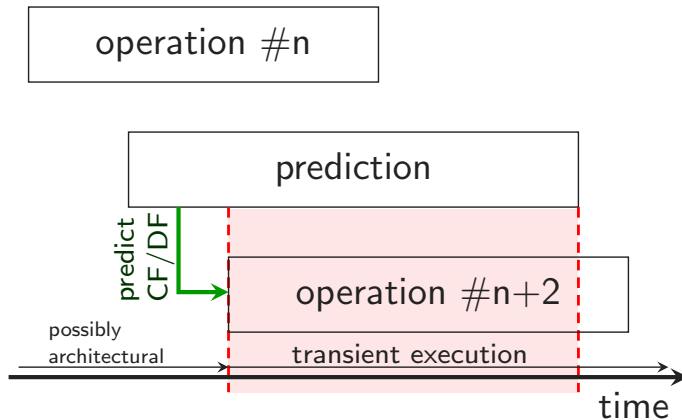


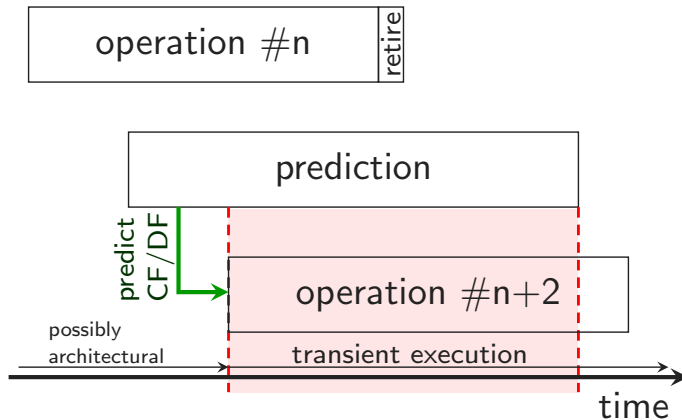
operation #n

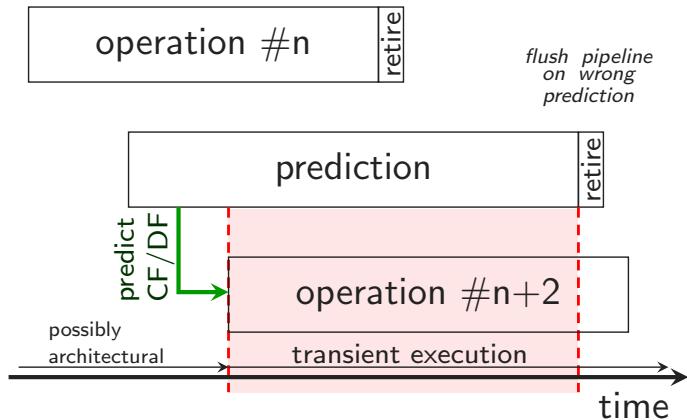


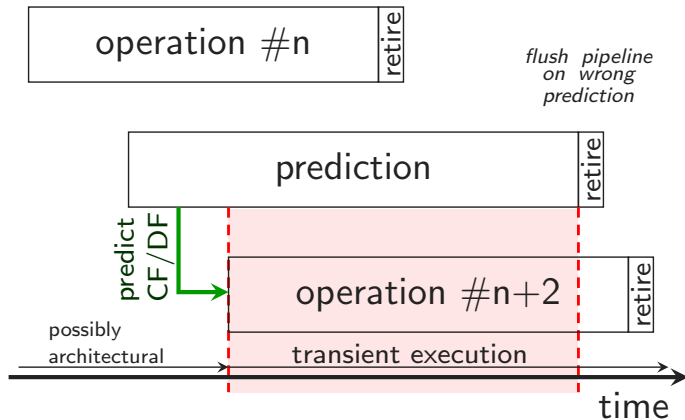


















- Many predictors in modern CPUs




- Many predictors in modern CPUs
 - Branch taken/not taken (PHT)




- Many predictors in modern CPUs
 - Branch taken/not taken (PHT)
 - Call/Jump destination (BTB)



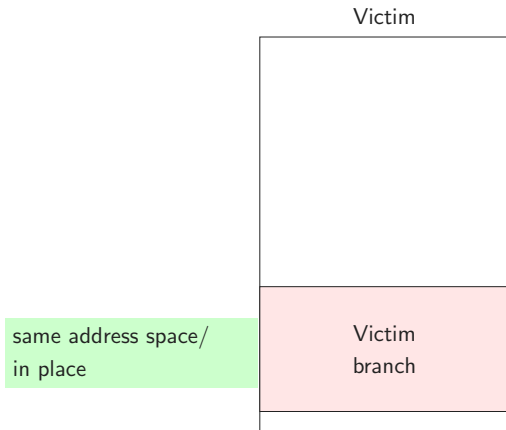
- **Many predictors** in modern CPUs
 - **Branch** taken/not taken (PHT)
 - **Call/Jump** destination (BTB)
 - Function **return** destination (RSB)

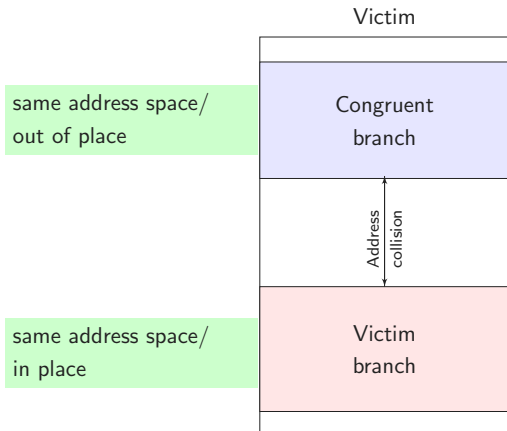


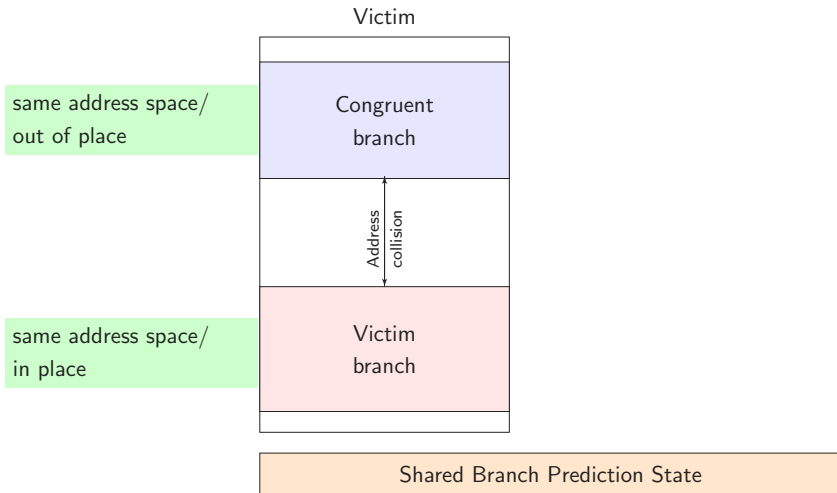
- **Many predictors** in modern CPUs
 - **Branch** taken/not taken (PHT)
 - **Call/Jump** destination (BTB)
 - Function **return** destination (RSB)
 - **Load** matches previous store (STL)

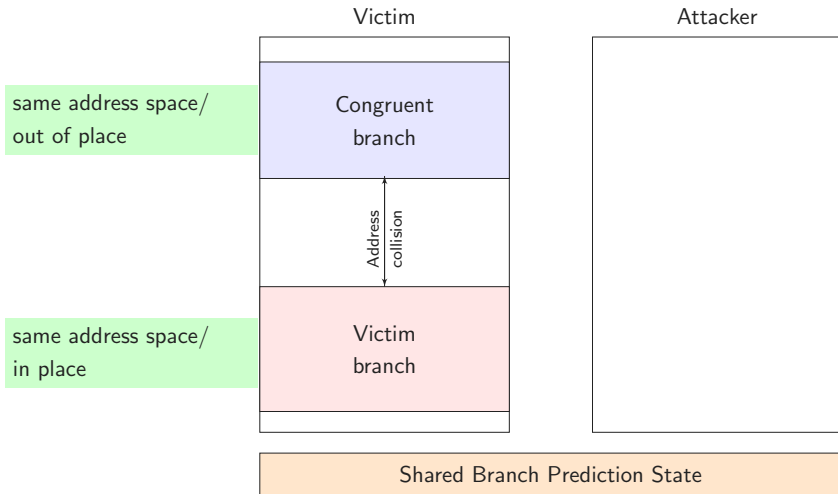


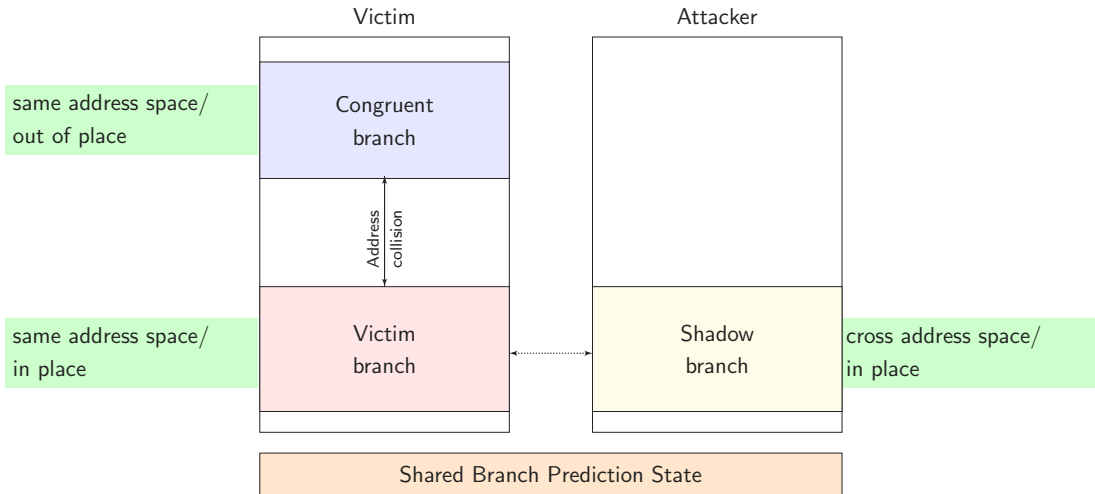
- **Many predictors** in modern CPUs
 - **Branch** taken/not taken (PHT)
 - **Call/Jump** destination (BTB)
 - Function **return** destination (RSB)
 - **Load** matches previous store (STL)
- Most are even **shared** among processes

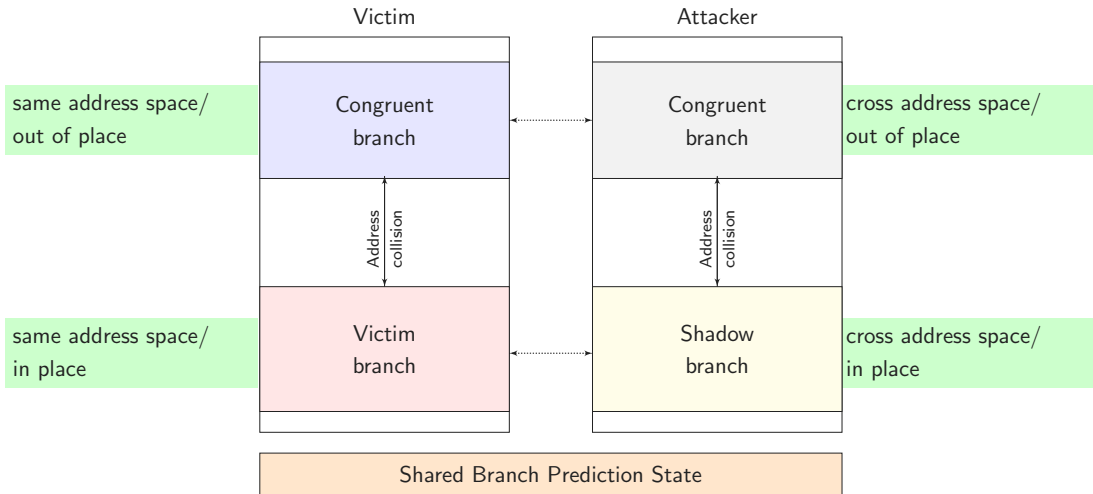










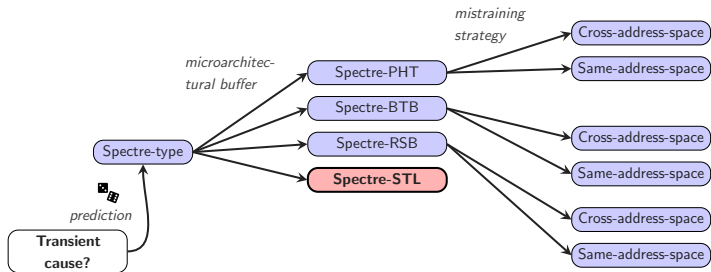


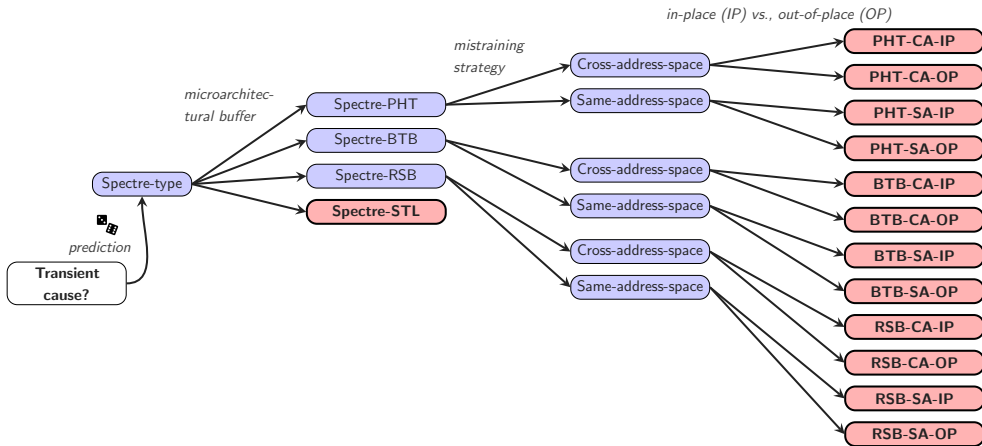



Transient
cause?














- Spectre is **not a bug**



- Spectre is **not a bug**
- It is an useful **optimization**




- Spectre is **not a bug**
 - It is an useful **optimization**
- Cannot simply fix it (as with Meltdown)



- Spectre is **not a bug**
 - It is an useful **optimization**
- Cannot simply fix it (as with Meltdown)
- **Workarounds** for critical code parts

Spectre defenses in 3 categories:



C1 Mitigating or reducing the accuracy of covert channels




C2 Mitigating or aborting speculation



C3 Ensuring secret data cannot be reached

		Attack	Defense	InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB	
Intel	Spectre-PHT																					
	Spectre-BTB																					
	Spectre-RSB																					
	Spectre-STL																					

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (▣), not theoretically impeded (□), or out of scope (◇).



Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (▣), not theoretically impeded (□), or out of scope (◇).

Attack \ Defense		InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB		
		Intel	Spectre-PHT					●	◐	◐	●						◐				
	Spectre-BTB				●			◐				●	◐	◐						◐	
	Spectre-RSB				◐			◐												◐	
	Spectre-STL							◐												◐	●

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (▣), not theoretically impeded (□), or out of scope (◇).

Attack \ Defense		InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB		
		Intel	Spectre-PHT					●	◐	◐	●	○					◐				
	Spectre-BTB				●			◐				●	◐	◐						◐	
	Spectre-RSB				◐			◐												◐	
	Spectre-STL							◐												◐	●

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (▣), not theoretically impeded (□), or out of scope (◇).

Attack \ Defense		InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB
		Intel	Spectre-PHT					●	◐	◐	●	○					◐	■	◐
	Spectre-BTB				●			◐				●	◐	◐		■	◐		
	Spectre-RSB			◐				◐								■	◐		
	Spectre-STL							◐								■	◐	■	●

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (◻), not theoretically impeded (◻), or out of scope (◇).

Attack \ Defense		InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB
		Intel	Spectre-PHT					●	◐	◐	●	○					◐	■	◐
	Spectre-BTB				●			◐				●	◐	◐		■	◐		
	Spectre-RSB			◐				◐								■	◐		
	Spectre-STL							◐								■	◐	■	●

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (◐), not theoretically impeded (◐), or out of scope (◇).

Attack \ Defense		InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB
		Intel	Spectre-PHT	□	□	□		●	◐	◐	●	○					◐	■	◐
	Spectre-BTB	□	□	□	●			◐				●	◐	◐		■	◐		
	Spectre-RSB	□	□	□	◐			◐								■	◐		
	Spectre-STL	□	□	□				◐								■	◐	■	●

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (◐), not theoretically impeded (□), or out of scope (◇).

Attack \ Defense		InvisiSpec	SafeSpec	DAWG	RSB Stuffing	Retpoline	Poison Value	Index Masking	Site Isolation	SLH	YSNB	IBRS	STIPB	IBPB	Serialization	Taint Tracking	Timer Reduction	Sloth	SSBD/SSBB
		Intel	Spectre-PHT	□	□	□	◇	◇	●	◐	◐	●	○	◇	◇	◇	◐	■	◐
	Spectre-BTB	□	□	□	◇	●	◇	◇	◐	◇	◇	●	◐	◐	◇	■	◐	◇	◇
	Spectre-RSB	□	□	□	◐	◇	◇	◇	◐	◇	◇	◇	◇	◇	◇	■	◐	◇	◇
	Spectre-STL	□	□	□	◇	◇	◇	◇	◐	◇	◇	◇	◇	◇	◇	■	◐	■	●

Attack is mitigated (●), partially mitigated (◐), not mitigated (○), theoretically mitigated (■), theoretically impeded (◐), not theoretically impeded (□), or out of scope (◇).



- Many countermeasures **only consider** the **cache** to get data...



- Many countermeasures **only consider** the **cache** to get data...
- ...but there are other possibilities, e.g.,



- Many countermeasures **only consider** the **cache** to get data...
- ...but there are other possibilities, e.g.,
 - Port contention (SMoTherSpectre)



- Many countermeasures **only consider** the **cache** to get data...
- ...but there are other possibilities, e.g.,
 - Port contention (SMoTherSpectre)
 - AVX (NetSpectre)



- Many countermeasures **only consider** the **cache** to get data...
- ...but there are other possibilities, e.g.,
 - Port contention (SMoTherSpectre)
 - AVX (NetSpectre)
- Cache is just the **easiest**

Linux 4.19.4 & 4.14.83 Released With STIBP Code Dropped

Written by [Michael Larabel](#) in [Linux Kernel](#) on 24 November 2018 at 09:00 AM EST. [6 Comments](#)



On Friday marked the release of the Linux 4.19.4 kernel as well as 4.14.83 and 4.9.139.

Greg Kroah-Hartman issued this latest round of stable point releases as basic maintenance updates. While these point releases don't tend to be too notable and generally go unmentioned on Phoronix, this round is worth pointing out since 4.19.4 and 4.14.83 are the releases that end up [reverting the STIBP behavior](#) that applied Single Thread Indirect Branch Predictors to all processes on supported systems. That is what was introduced in Linux 4.20 and then back-ported to the 4.19/4.14 LTS branches, which in turn hurt the performance a lot. So for now the code is removed.

As covered yesterday, [there is improved STIBP code on the way](#) for Linux 4.20 that by default just apply STIBP to SECCOMP threads and processes requesting it via prctl() but otherwise is off by default (that behavior can also be changed via kernel parameters).

Linux 4.19.4 & 4.14.83 Released With STIBP Code Dropped

Written by [Michael Larabel](#) in [Linux Kernel](#) on 24 November 2018 at 09:00 AM EST. [6 Comments](#)



On Friday marked the release of the Linux 4.19.4 kernel as well as 4.14.83 and 4.9.139.

Greg Kroah-Hartman issued this latest round of stable point releases as basic maintenance updates. While these point releases don't tend to be too notable and generally go unmentioned on Phoronix, this round is worth pointing out since 4.19.4 and 4.14.83 are the releases that end up [reverting the STIBP behavior](#) that applied Single Thread Indirect Branch Predictors to all processes on supported systems. That is what was introduced in Linux 4.20 and then back-ported to the 4.19/4.14 LTS branches, which in turn hurt the performance a lot. So for now the code is removed.

As covered yesterday, [there is improved STIBP code on the way](#) for Linux 4.20 that by default just apply STIBP to SECCOMP threads and processes requesting it via `prctl()` but otherwise is off by default (that behavior can also be changed via kernel parameters).

Linux 4.19.4 & 4.14.83 Released With STIBP Code Dropped

Written by Michael Larabel in Linux Kernel on 24 November 2018 at 09:00 AM EST. 6 Comments



On Friday marked the release of the Linux 4.19.4 kernel as well as 4.14.83 and 4.9.139.

Greg Kroah-Hartman issued this latest round of stable point releases as basic maintenance updates. While these point releases don't tend to be too notable and generally go unmentioned on Phoronix, this round is worth pointing out since 4.19.4 and 4.14.83 are the releases that end up reverting the STIBP behavior that applied Single Thread Indirect Branch Predictors to all processes on supported systems. That is what was introduced in Linux 4.20 and then back-ported to the 4.19/4.14 LTS branches, which in turn hurt the performance a lot. So for now the code is removed.

As covered yesterday, there is improved STIBP code on the way for Linux 4.20 that by default just apply STIBP to SECCOMP threads and processes requesting it via prctl() but otherwise is off by default (that behavior can also be changed via kernel parameters).

Linux 4.19.4 & 4.14.83 Released With STIBP Code Dropped


Written by Michael Larabel in Linux Kernel on 24 November 2018 at 09:00 AM EST. 6 Comments




On Friday marked the release of the Linux 4.19.4 kernel as well as 4.14.83 and 4.9.139.

Greg Kroah-Hartman issued this latest round of stable point releases as basic maintenance updates. While these point releases don't tend to be too notable and generally go unmentioned on Phoronix, this round is worth pointing out since 4.19.4 and 4.14.83 are the releases that end up reverting the STIBP behavior that applied Single Thread Indirect Branch Predictors to all processes on supported systems. That is what was introduced in Linux 4.20 and then back-ported to the 4.19/4.14 LTS branches, which in turn hurt the performance a lot. So for now the code is removed.


As covered yesterday, there is improved STIBP code on the way for Linux 4.20 that by default just apply STIBP to SECCOMP threads and processes requesting it via prctl() but otherwise is off by default (that behavior can also be changed via kernel parameters).




- Current mitigations are either **incomplete or cost performance**



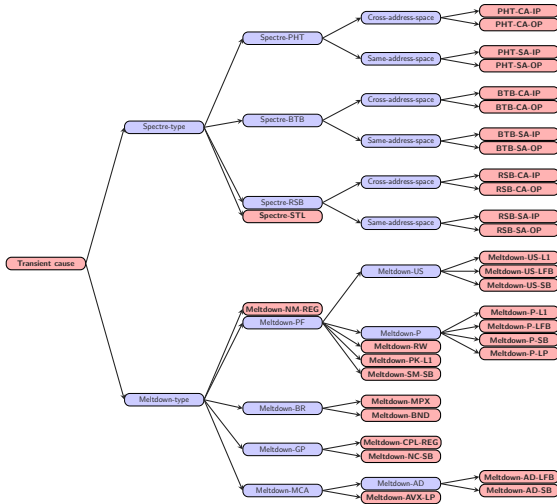
- Current mitigations are either **incomplete or cost performance**
- More **research** required




- Current mitigations are either **incomplete or cost performance**
- More **research** required
- Both on **attacks and defenses**




- Current mitigations are either **incomplete or cost performance**
- More **research** required
- Both on **attacks and defenses**
- Efficient defenses only possible when attacks are known







- **Transient Execution Attacks** are...




- **Transient Execution Attacks** are...
 - ...a **novel class** of attacks



- **Transient Execution Attacks** are...
 - ...a **novel class** of attacks
 - ...extremely **powerful**



- **Transient Execution Attacks** are...
 - ...a **novel class** of attacks
 - ...extremely **powerful**
 - ...only at the **beginning**



- **Transient Execution Attacks** are...
 - ...a **novel class** of attacks
 - ...extremely **powerful**
 - ...only at the **beginning**
- Many optimizations introduce side channels → now exploitable

BRACE YOURSELVES

MORE BUGS ARE COMING

Transient Execution Attacks

Exploiting the CPU's Microarchitecture

Michael Schwarz (@misc0110), Moritz Lipp (@mlqxyz)

October 7, 2019

Graz University of Technology