


# Exploiting the Microarchitecture: Transient Execution Attacks

**Michael Schwarz (@misc0110)**

April 11, 2019

Graz University of Technology



## Michael Schwarz

PhD candidate @ Graz University of Technology

🐦 @misc0110

✉ michael.schwarz@iaik.tugraz.at

**FOX**  
**BUSINESS**  
WASHINGTON, D.C.

WASHINGTON, D.C.

**NEWS**  
**ALERT**

**INTEL REVEALS DESIGN FLAW THAT  
COULD ALLOW HACKERS TO ACCESS DATA**

**WINTER STORM**



**FOX**  
**BUSINESS**  
NETWORK



**@FOXBUSINESS**



**DEVELOPING STORY**

# COMPUTER CHIP FLAWS IMPACT BILLIONS OF DEVICES

LIVE



DAX ▲ 164.69

NEWS STREAM



GLOBAL

## COMPUTER CHIP SCARE

The bugs are known as 'Spectre' and 'Meltdown'


**BBC** WORLD NEWS |

• £:HK\$ 10.58 •


EURO:£ 0.891 •


E

- Bug-free software does not mean safe execution



- Bug-free software does not mean safe execution
- Information leaks due to **underlying hardware**






- Bug-free software does not mean safe execution
- Information leaks due to **underlying hardware**
- **Exploit** leakage through **side-effects**



- Bug-free software does not mean safe execution
- Information leaks due to **underlying hardware**
- **Exploit** leakage through **side-effects**



Power consumption




Execution time




CPU caches







- 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




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




- 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**




Caches and buffers



Predictors





- Modern CPUs contain multiple **microarchitectural elements**



Caches and buffers




Predictors



- **Transparent** for the programmer





- Modern CPUs contain multiple **microarchitectural elements**



Caches and buffers




Predictors





- **Transparent** for the programmer
- Timing optimizations → side-channel leakage


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





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


*Cache miss*






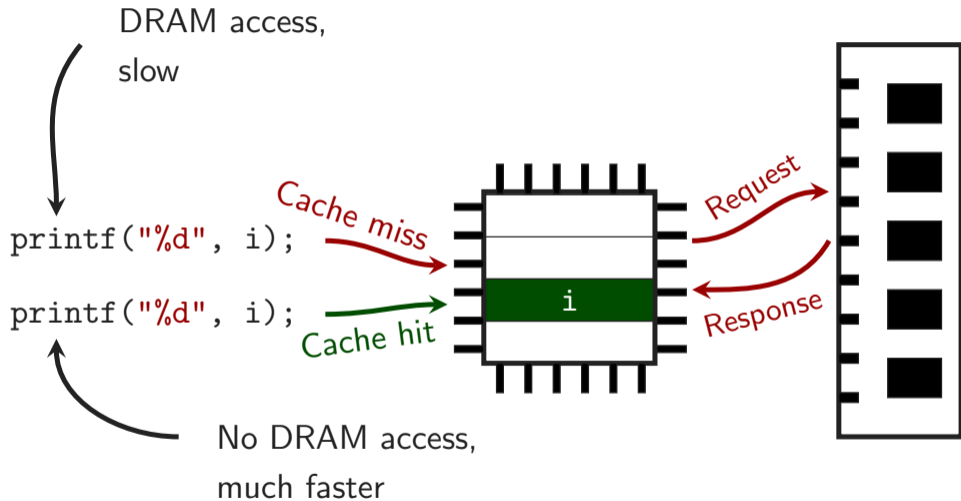


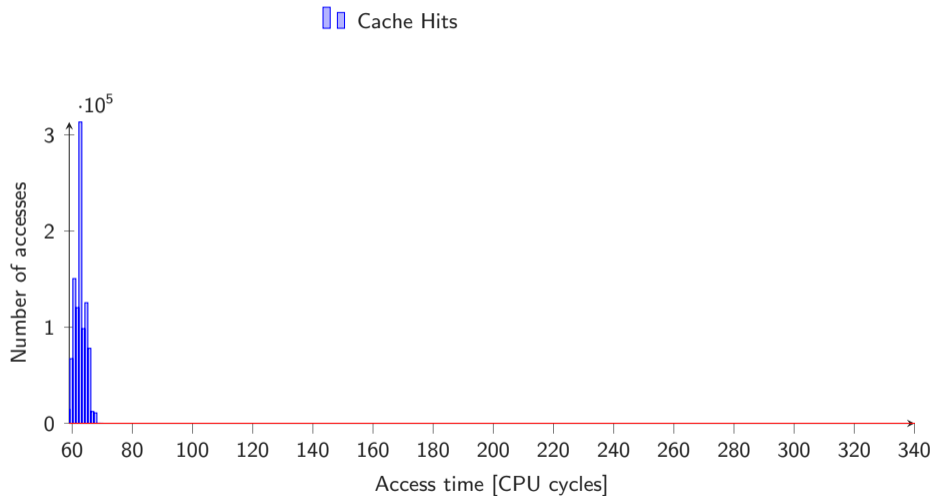












































```
char array[256 * 4096]; // 256 pages of memory
```




```
char array[256 * 4096]; // 256 pages of memory  
  
*(volatile char*) 0; // raise_exception();  
array[84 * 4096] = 0;
```

- Flush+Reload over all pages of the array






- 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**






- 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**



- 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**
- Give such instructions a name: **transient instructions**



- Add another **layer of indirection** to test

```
char array[256 * 4096]; // 256 pages of memory
```



- Add another **layer of indirection** to test

```
char array[256 * 4096]; // 256 pages of memory
```

```
// read kernel address (raises exception)
```

```
char data = *(char*) 0xffffffff81a000e0;
```

```
array[data * 4096] = 0;
```



- Add another **layer of indirection** to test

```
char array[256 * 4096]; // 256 pages of memory
```


```
// read kernel address (raises exception)
```

```
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 **too late**



- CPU uses data in out-of-order execution before permission check





**MELTDOWN**

- CPU uses data in out-of-order execution before permission check
- Meltdown can read any kernel address




**MELTDOWN**

- CPU uses data in out-of-order execution before permission check
- Meltdown can read any kernel address
- Physical memory is usually mapped in kernel




## MELTDOWN


- CPU uses data in out-of-order execution before permission check
  - Meltdown can read any kernel address
  - Physical memory is usually mapped in kernel
- Read arbitrary memory




- 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




- 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
- ...leaks from line fill buffer



- Kernel addresses in user space are a problem

- Kernel addresses in user space are a problem
- Why don't we take the kernel addresses...







- ...and remove them if not needed?




- ...and remove them if not needed?
- User accessible check in hardware is not reliable



## 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
- Looking closer at the check...




- CPU uses **virtual address spaces** to isolate processes



- CPU uses **virtual address spaces** to isolate processes
- Physical memory is organized in **page frames**



- CPU uses **virtual address spaces** to isolate processes
- Physical memory is 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

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

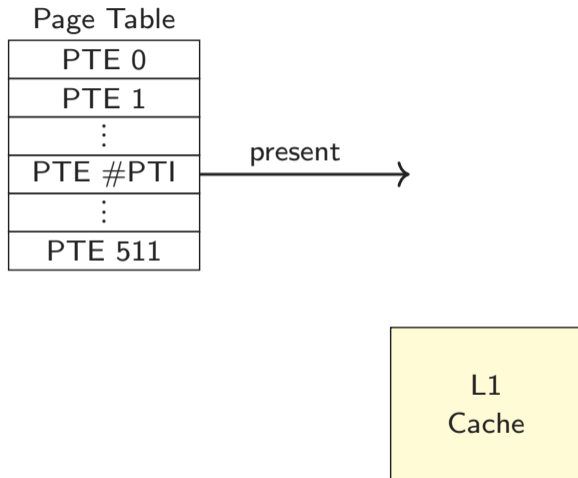



- An even **worse** bug → Foreshadow-NG/L1TF
- Exploitable from **VMs**
- Allows **leaking** data from the **L1** cache
- Same mechanism as Meltdown
- Just a **different bit** in the PTE

Page Table

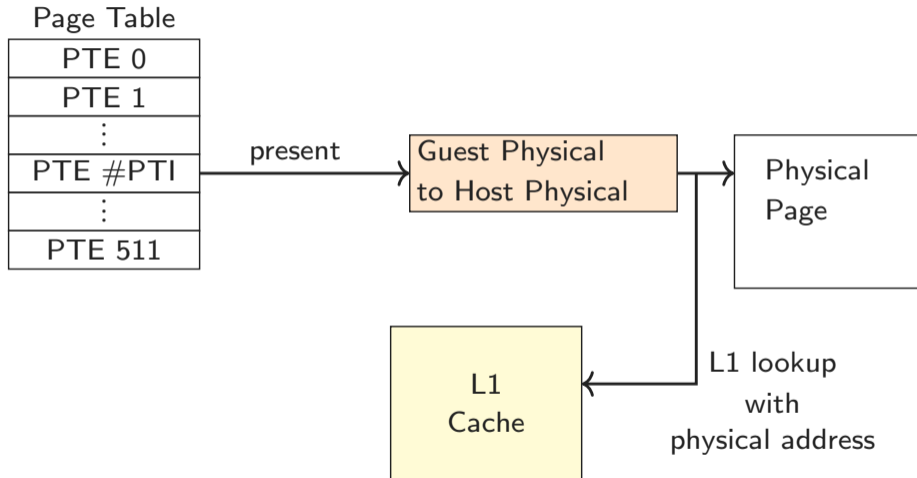
PTE 0
PTE 1
⋮
PTE #PTI
⋮
PTE 511

L1  
Cache








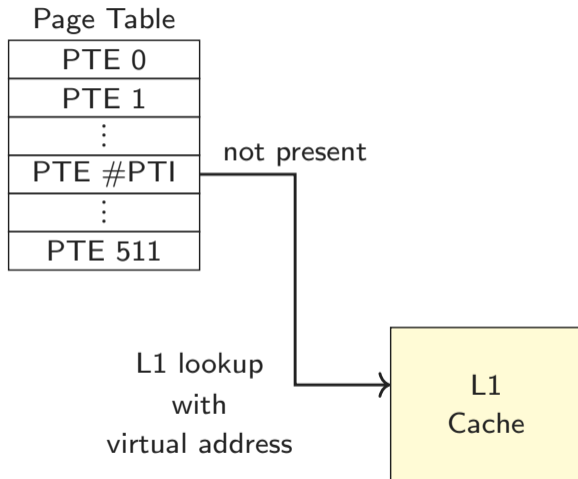


Page Table

PTE 0
PTE 1
⋮
PTE #PTI
⋮
PTE 511

not present







**Demo**  
Foreshadow-NG



- 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

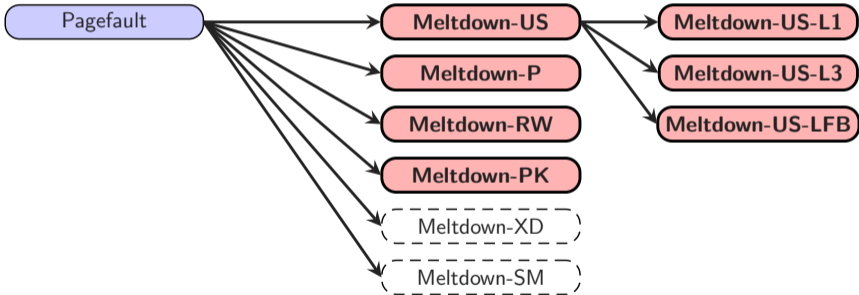
Pagefault










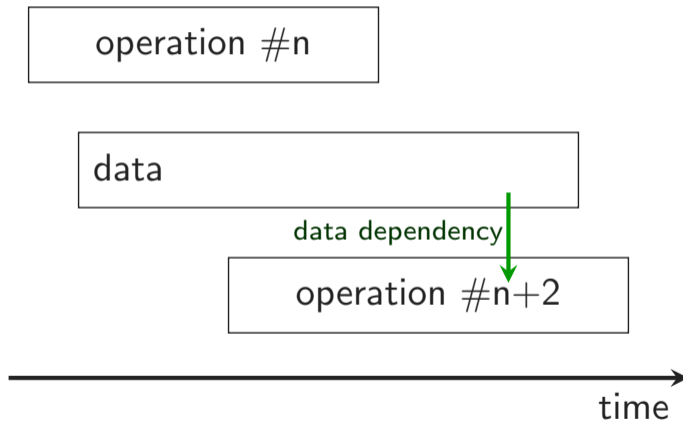


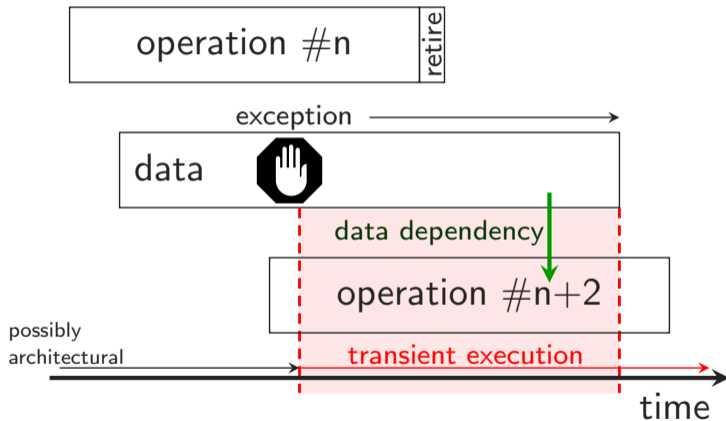
operation #n

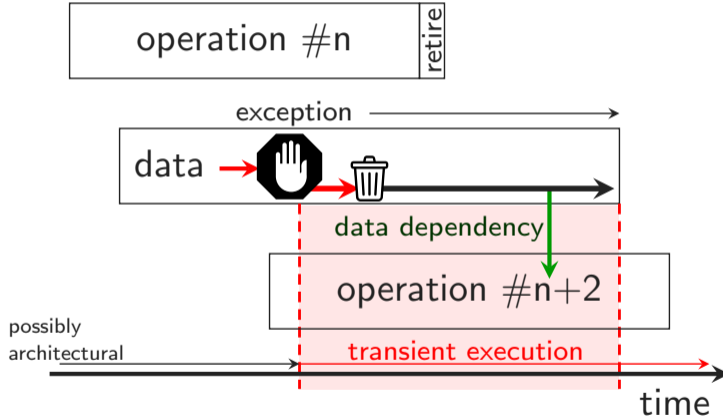


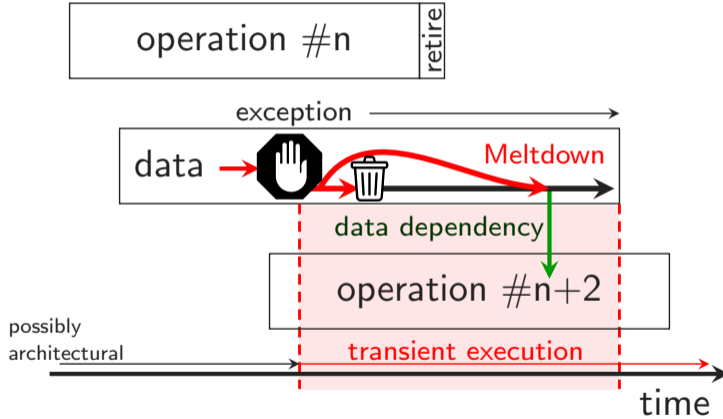


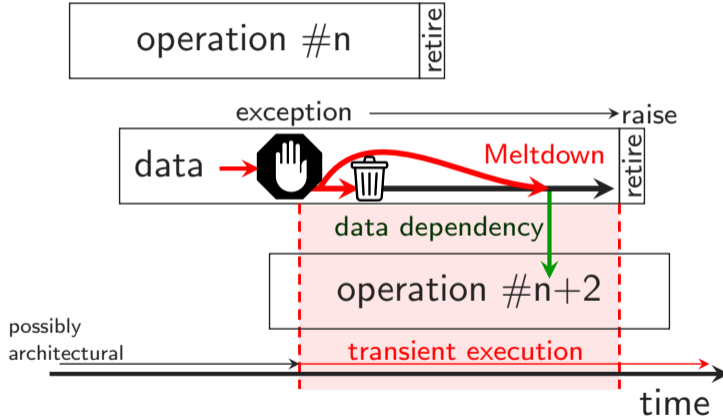












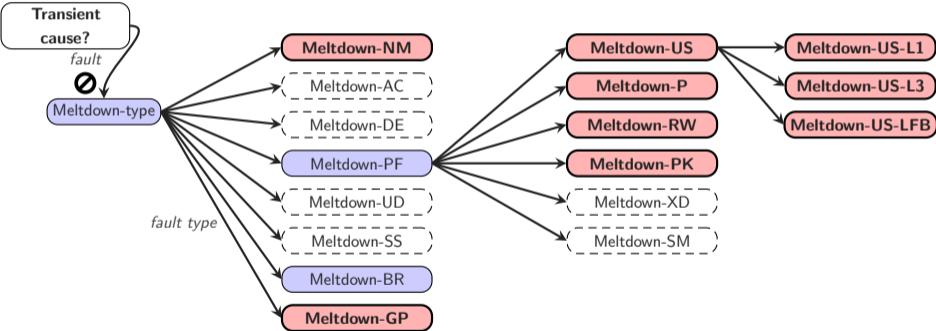


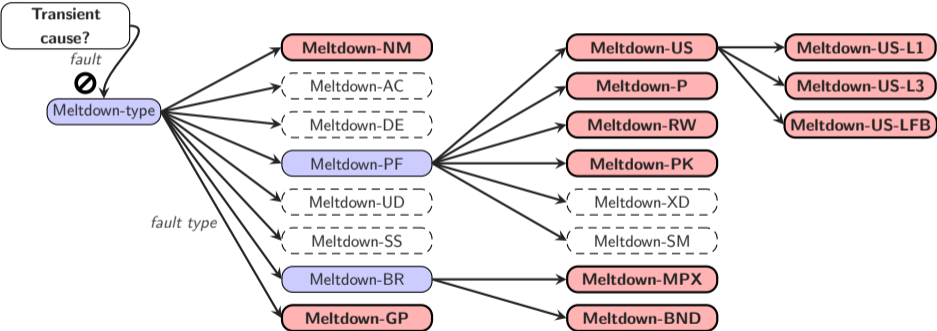
Transient  
cause?













- 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



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









`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

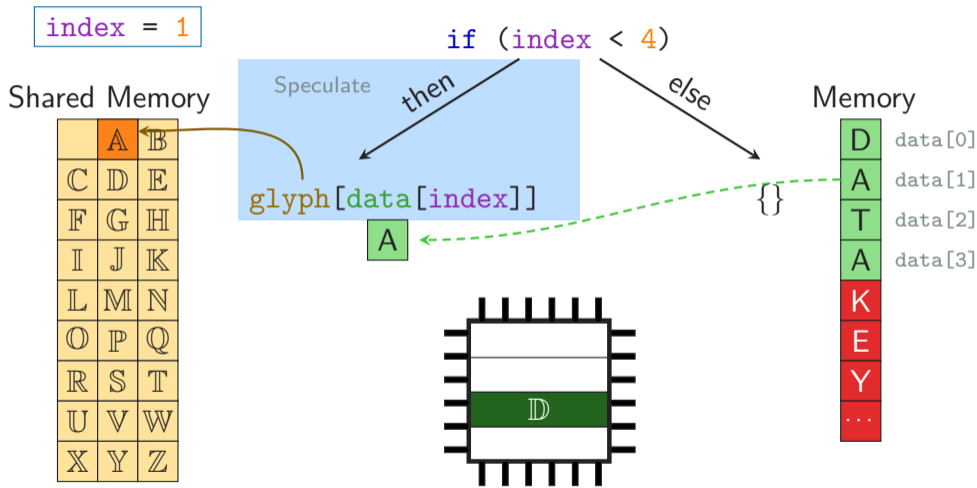



Memory

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








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

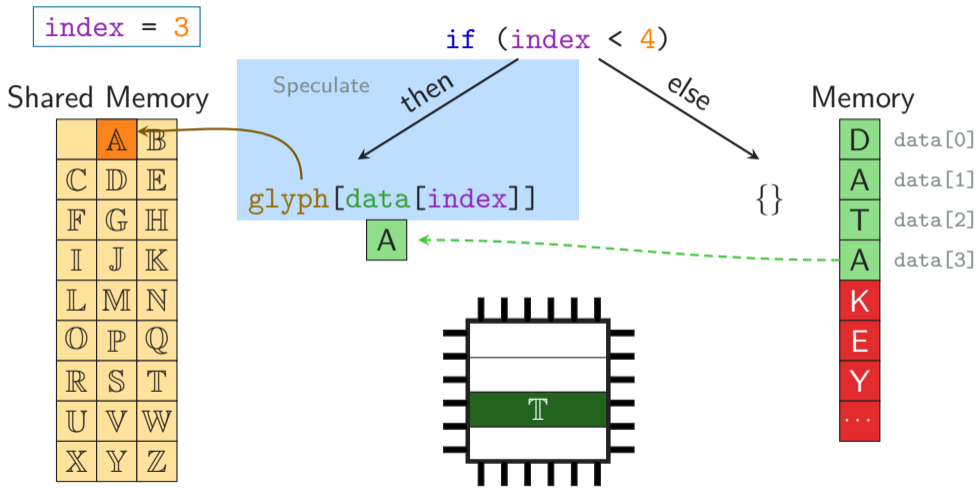



Memory

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

A








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

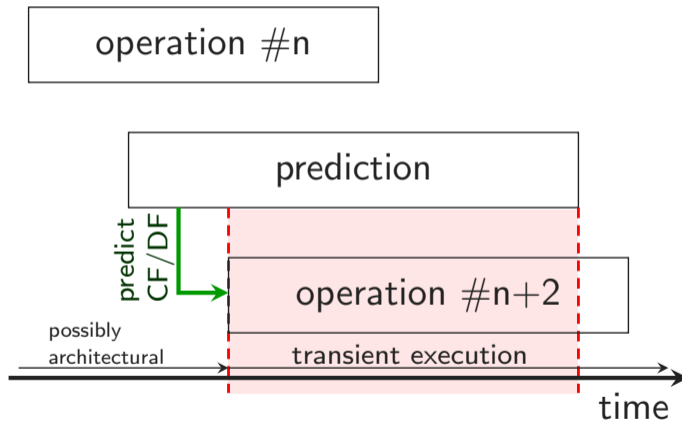
D	data[0]
A	data[1]
T	data[2]
A	data[3]
K	
E	
Y	
...	

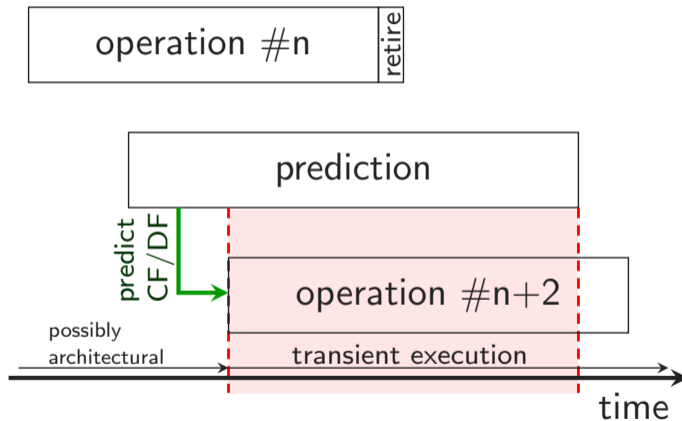
operation #n

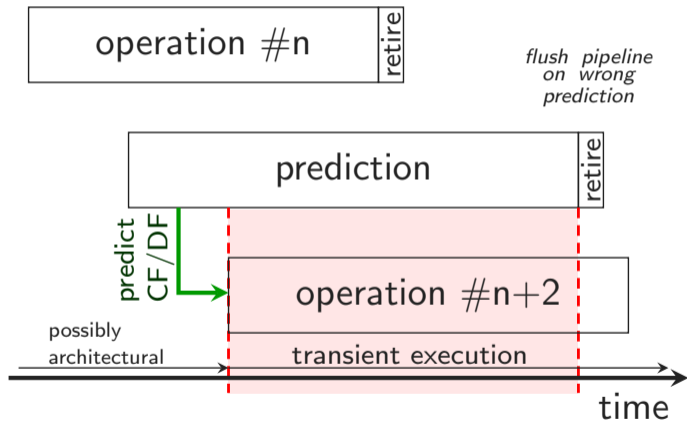




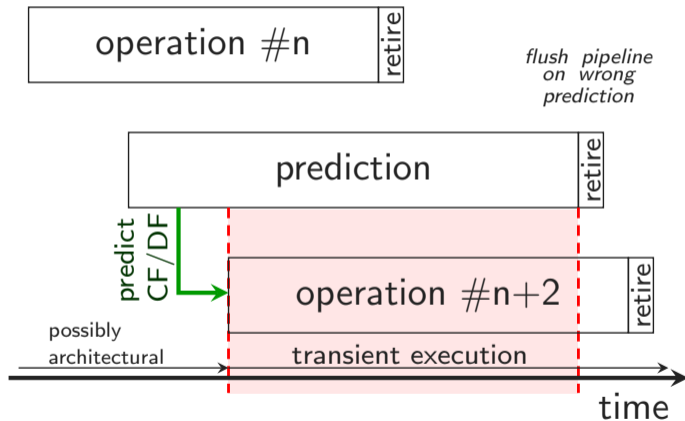















- Many predictors in modern CPUs



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



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



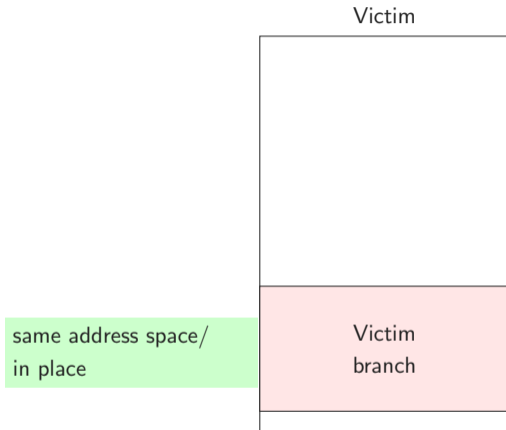
- **Many predictors** in modern CPUs
  - **Branch** taken/not taken (PHT)
  - **Call/Jump** destination (BTB)
  - Function **return** destination (RSB)




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

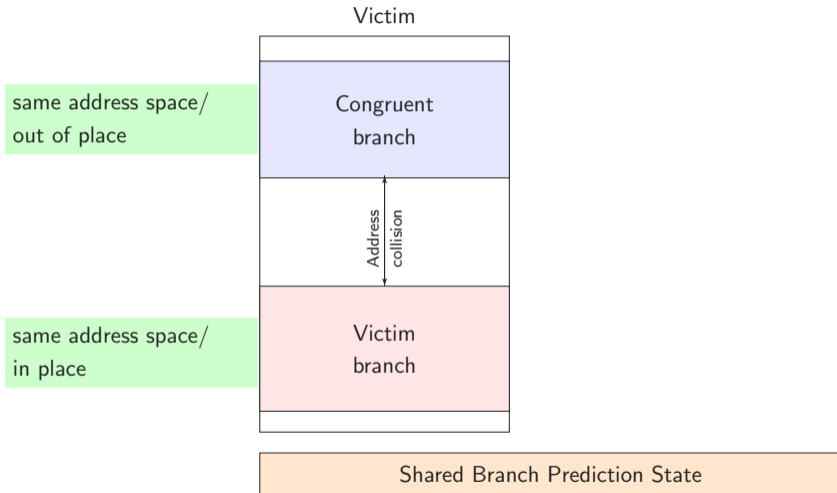


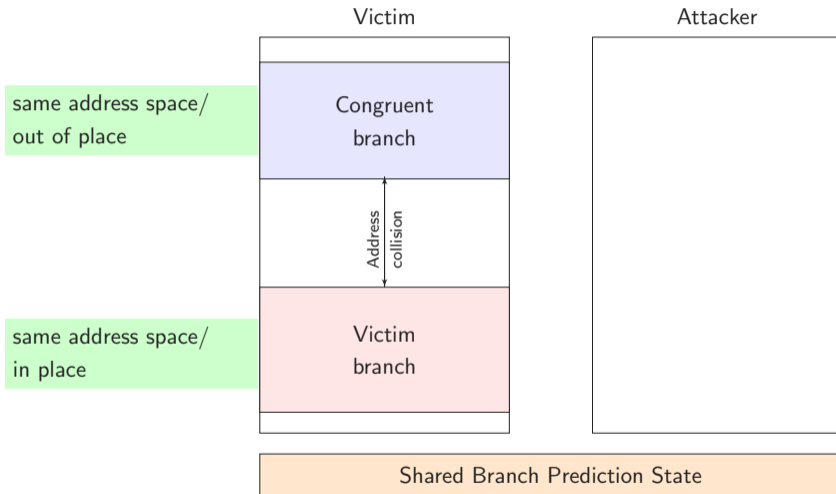
- **Many predictors** in modern CPUs
  - **Branch** taken/not taken (PHT)
  - **Call/Jump** destination (BTB)
  - Function **return** destination (RSB)
  - **Load** matches previous store (STL)
- Most are even **shared** among processes

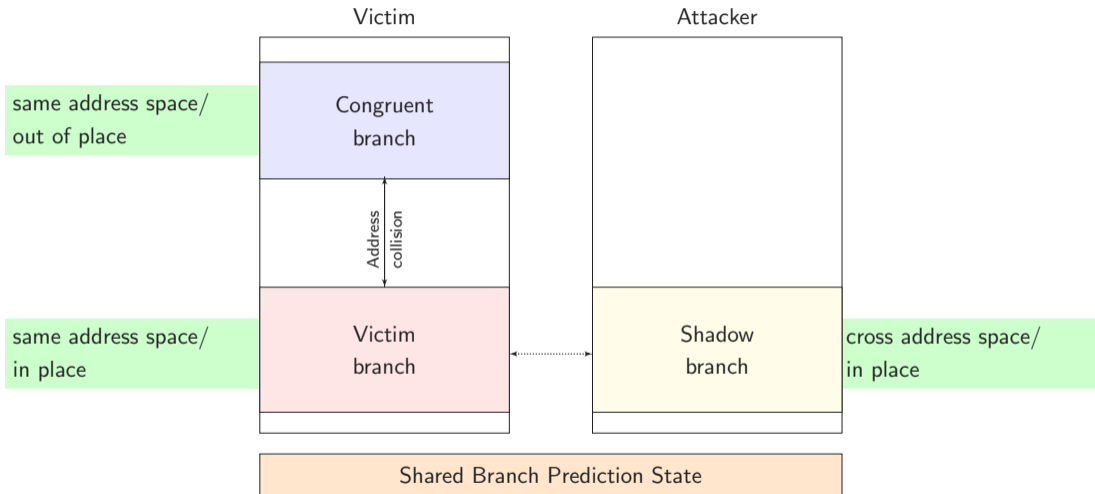


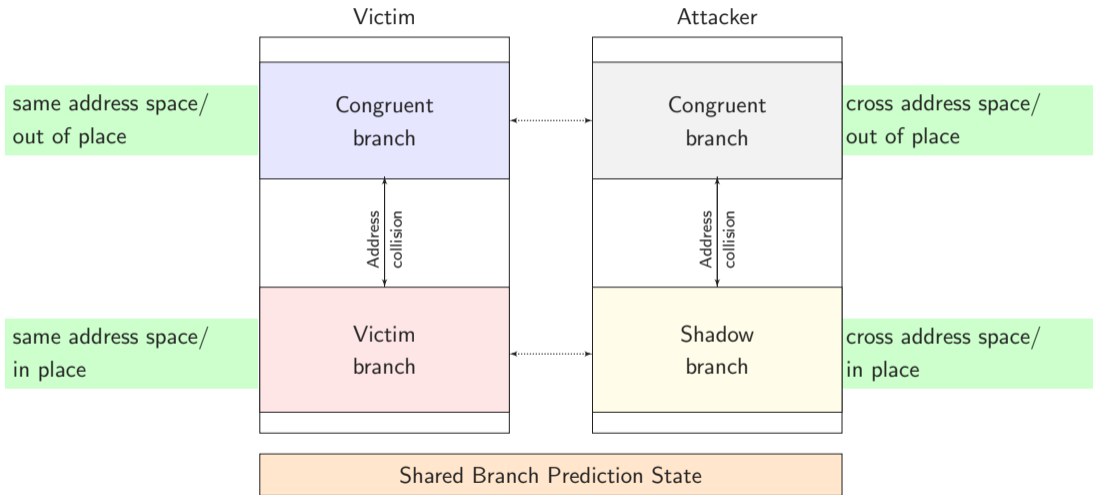
















Transient  
cause?













- 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	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 (◇).

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).



Retpoline (compiler extension)



## Retpoline (compiler extension)




```
    push <call_target>
    call 1f

2:                ; speculation continues here
    lfence        ; speculation barrier
    jmp 2b        ; endless loop

1:
    lea 8(%rsp), %rsp ; restore stack pointer
    ret          ; the call to <call_target>
```

→ Always predict to enter an endless loop


## Retpoline (compiler extension)




```
    push <call_target>
    call 1f
2:                ; speculation continues here
    lfence        ; speculation barrier
    jmp 2b         ; endless loop
1:
    lea 8(%rsp), %rsp ; restore stack pointer
    ret           ; the call to <call_target>
```

→ Always predict to enter an endless loop


- What if someone decides to fix the wrong prediction?




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



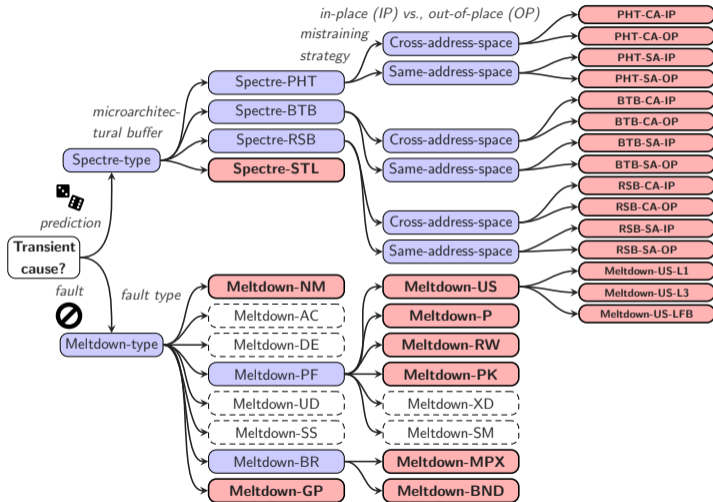
- Current mitigations are either **incomplete or cost performance**
- More **research** required




- 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**

# Exploiting the Microarchitecture: Transient Execution Attacks

**Michael Schwarz (@misc0110)**

April 11, 2019

Graz University of Technology