

Transient Execution Attacks: Still hARMful?

Barbara Gigerl (@barbarag2112), Claudio Canella (@cc0x1f)

May 17, 2019

Graz University of Technology



Barbara Gigerl

Master student @ Graz University of Technology

🐦 @barbarag2112

✉️ barbara.gigerl@student.tugraz.at



Claudio Canella

PhD student @ Graz University of Technology

🐦 @cc0x1f

✉ claudio.canella@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




- 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...






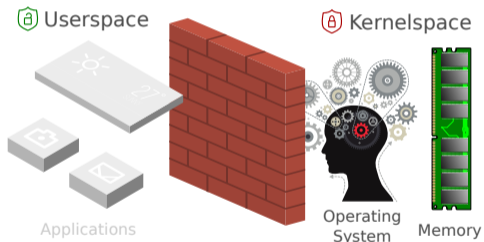
- ...and remove them if not needed?




- ...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...



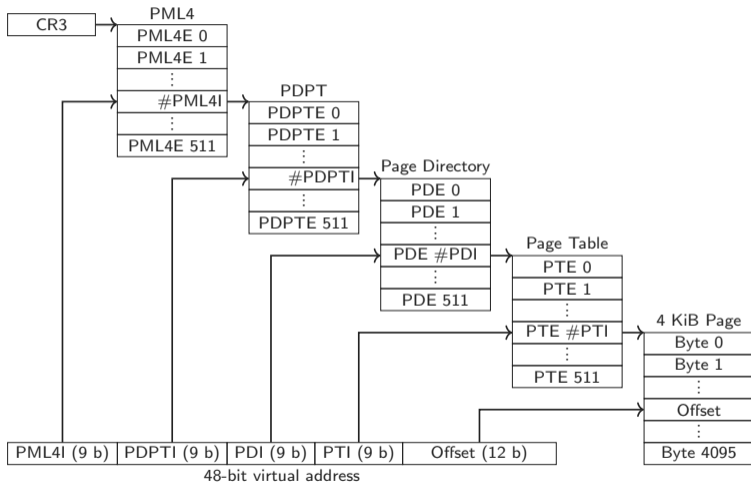
- CPU uses **virtual address spaces** to isolate processes



- 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



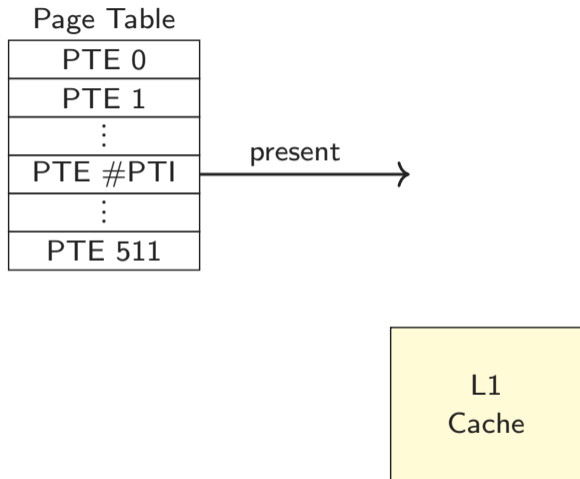
- An even **worse** bug → Foreshadow-NG/L1TF
- Exploitable from **VMs**
- Allows **leaking** data from the **L1** cache
- Same mechanism as Meltdown
- Just a **different bit** in the PTE


Page Table

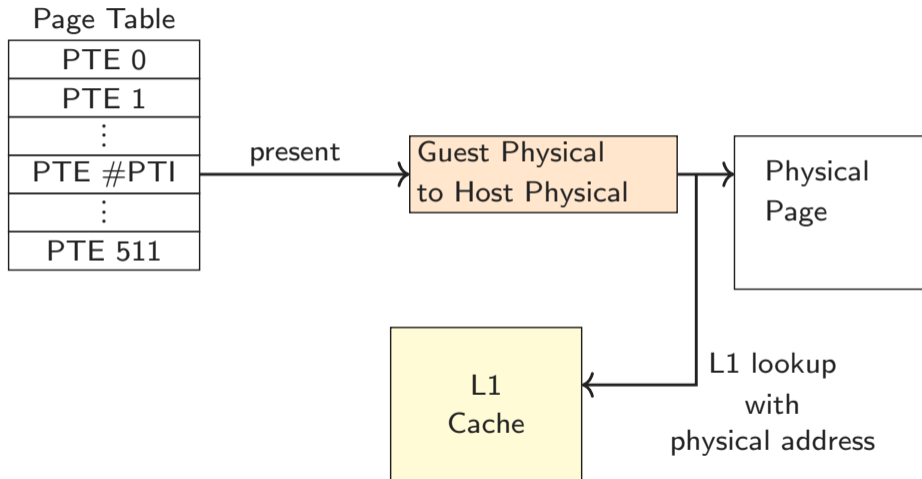
PTE 0
PTE 1
⋮
PTE #PTI
⋮
PTE 511

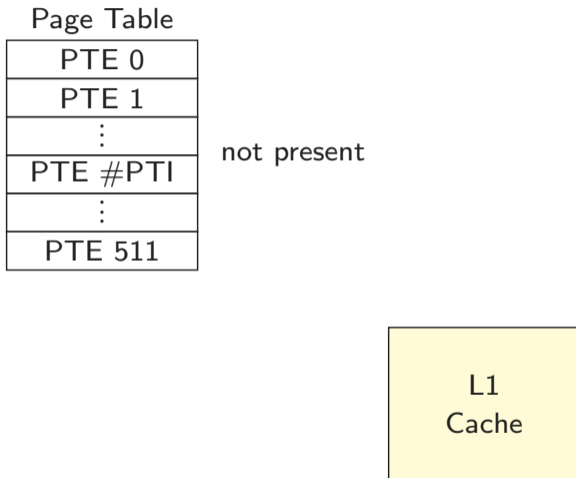



L1
Cache













- 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

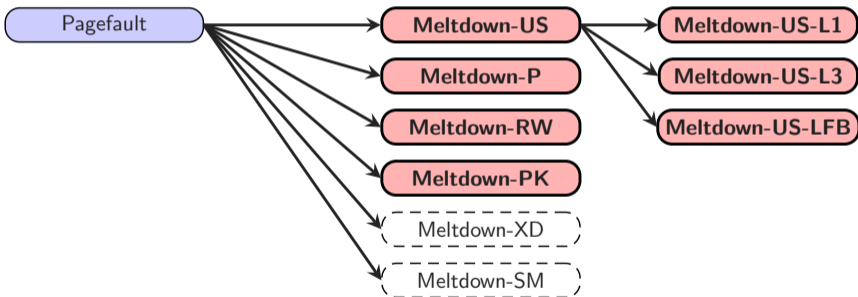
Pagefault








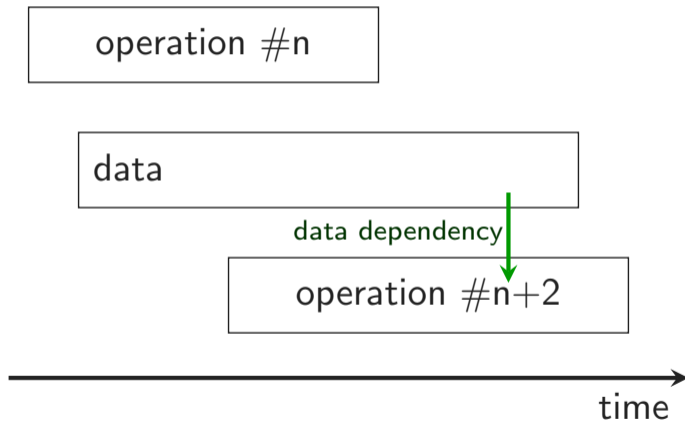


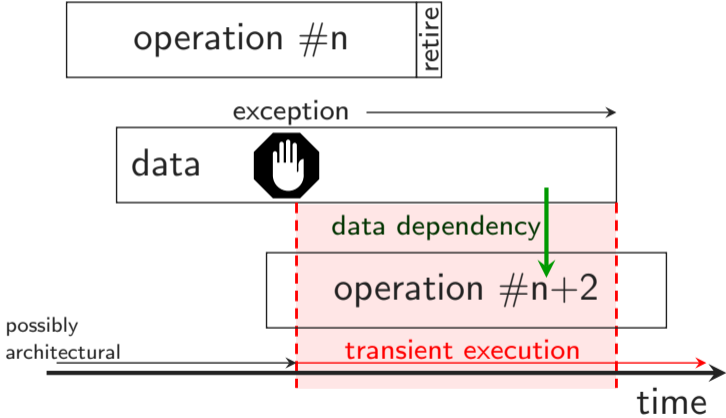


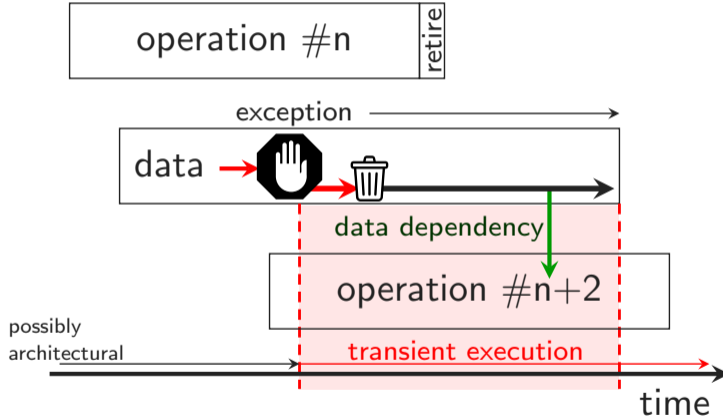
operation #n

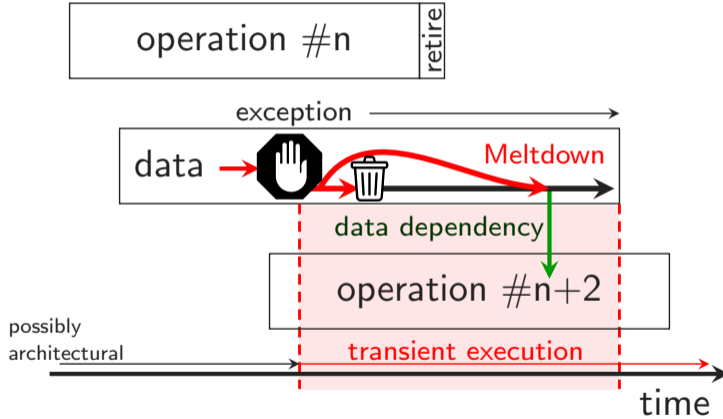


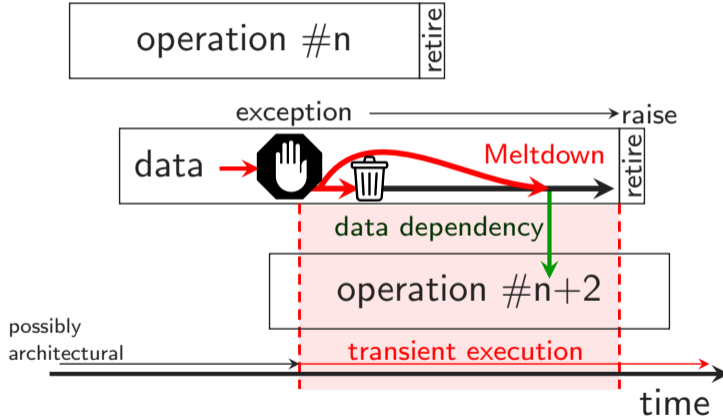










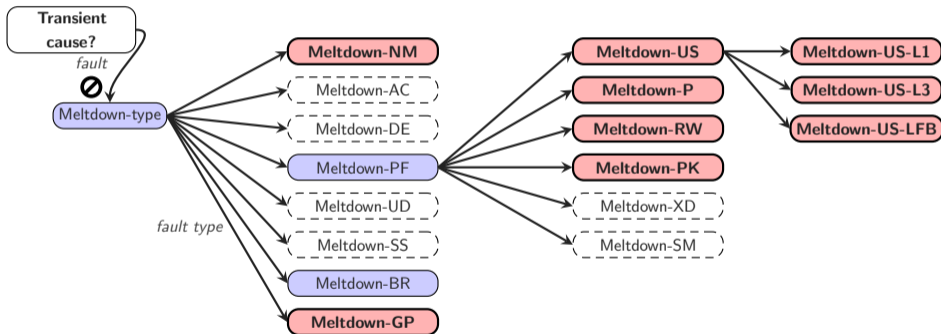


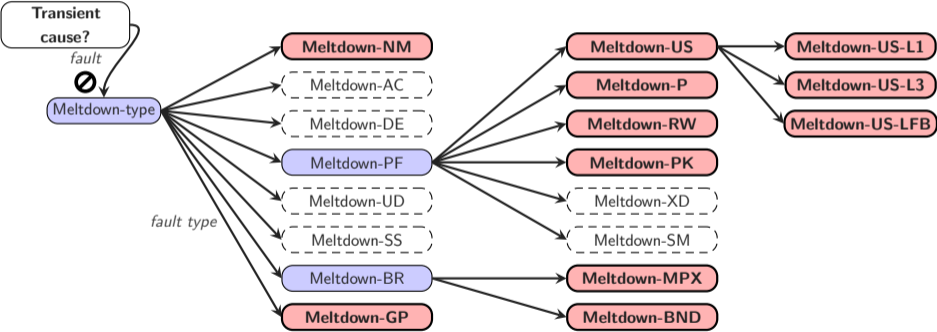


Transient
cause?











- 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
- 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

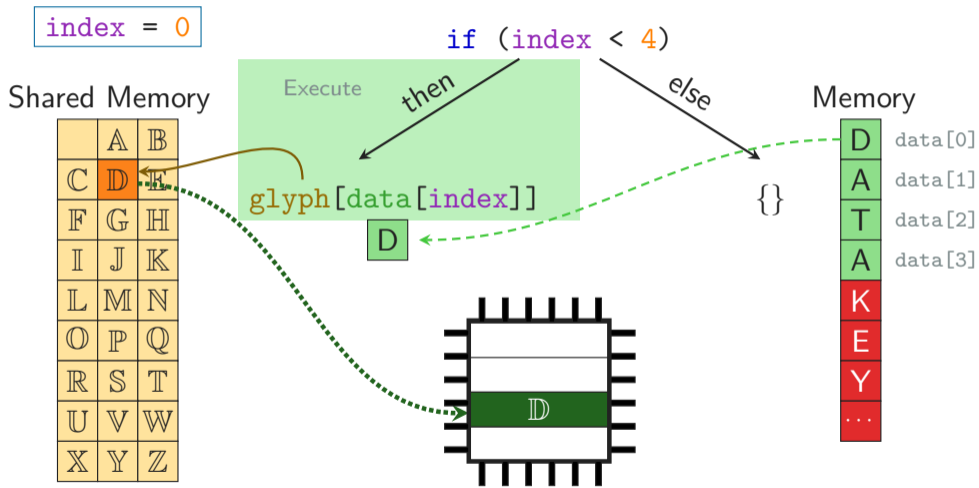


Memory

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

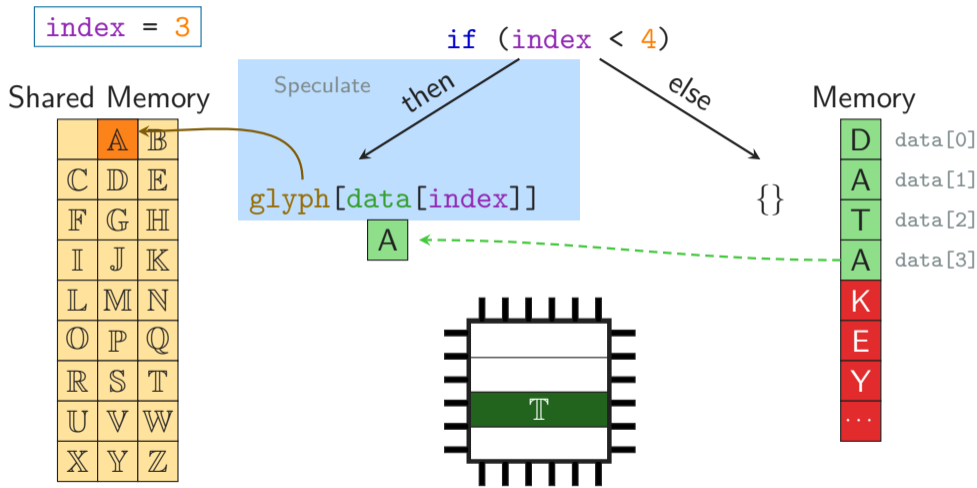
	A	B
C	D	E
F	G	H
I	J	K
L	M	N
O	P	Q
R	S	T
U	V	W
X	Y	Z




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	
...	

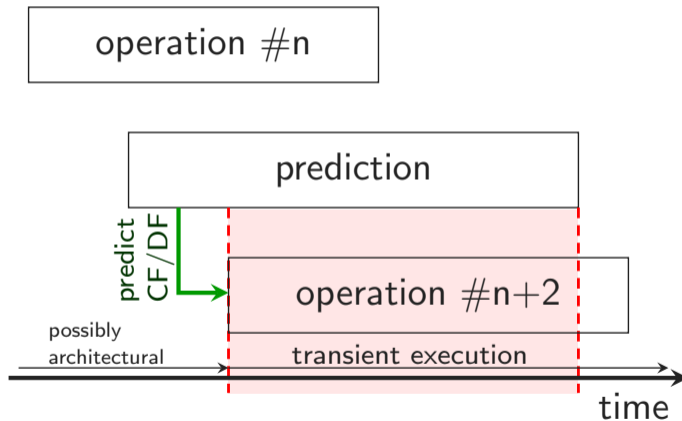


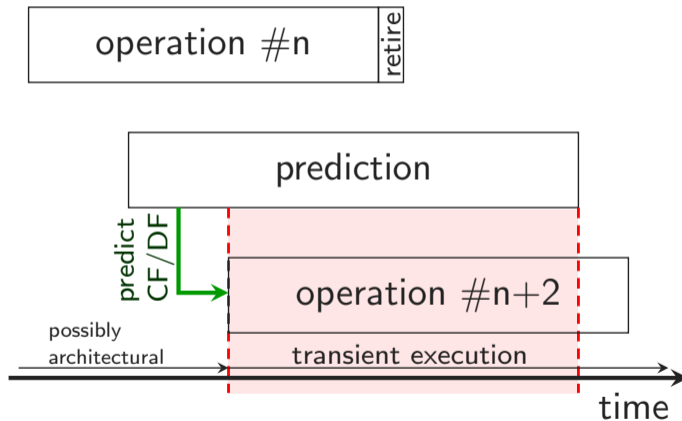
operation #n

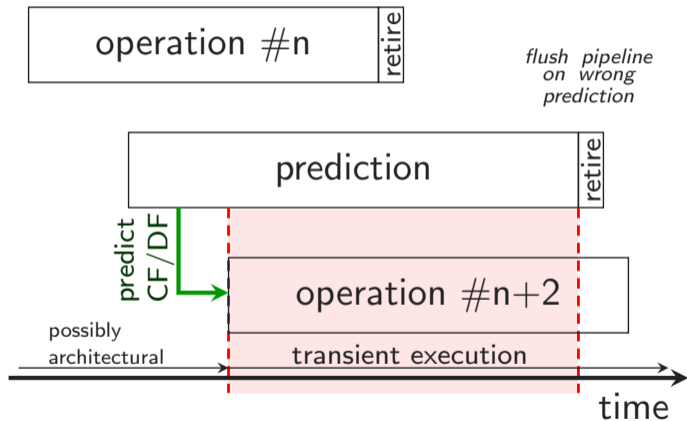


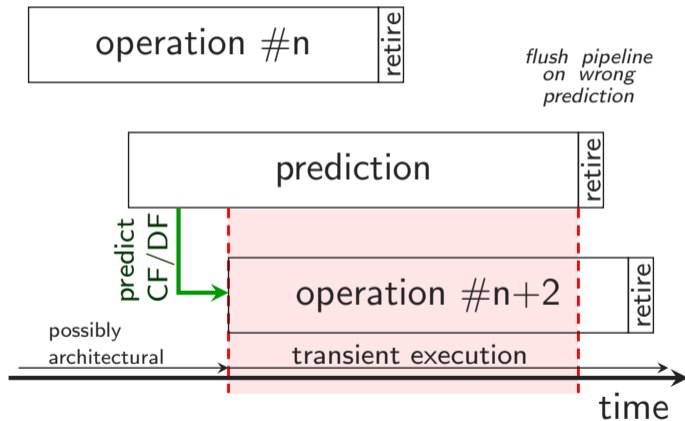














- Many predictors in modern CPUs




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




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



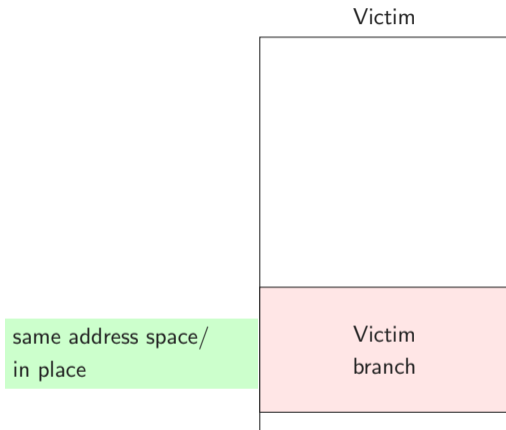
- Many predictors in modern CPUs
 - Branch taken/not taken (PHT)
 - Call/Jump destination (BTB)
 - Function return destination (RSB)




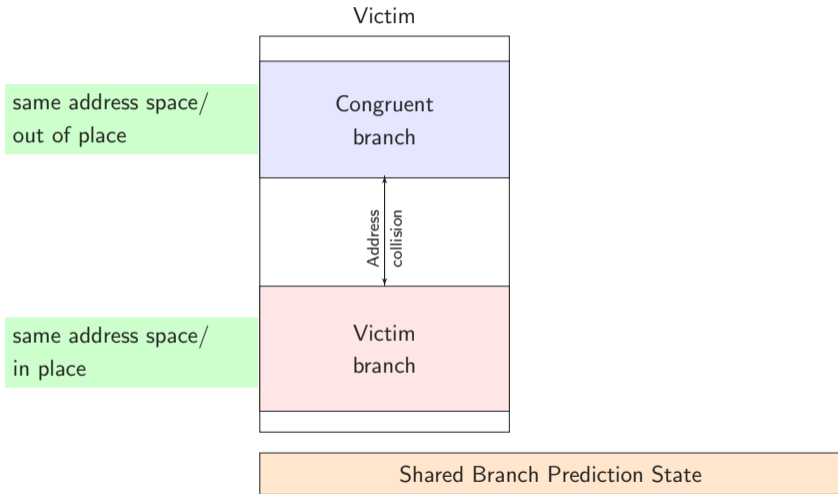
- **Many predictors** in modern CPUs
 - **Branch** taken/not taken (PHT)
 - **Call/Jump** destination (BTB)
 - Function **return** destination (RSB)
 - **Load** matches previous store (STL)




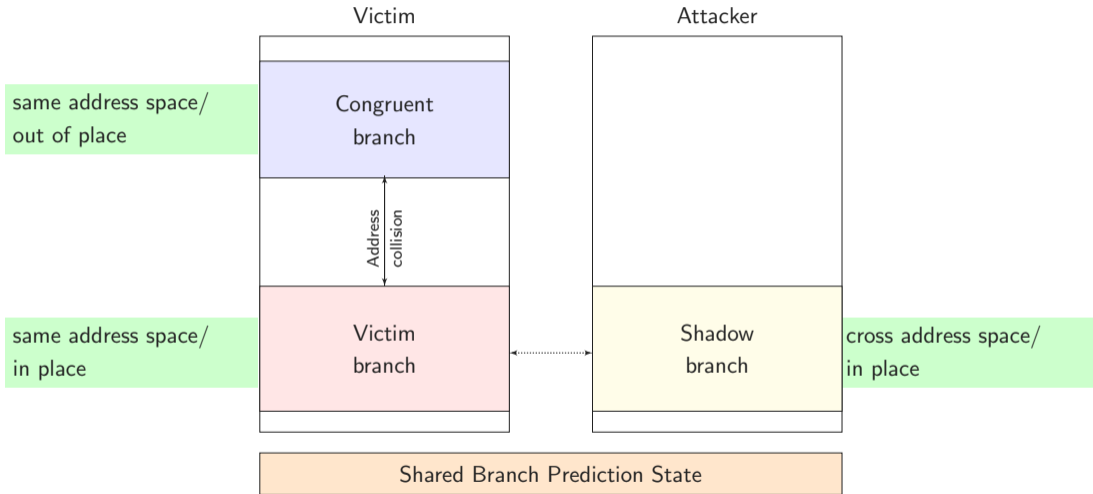
- **Many predictors** in modern CPUs
 - **Branch** taken/not taken (PHT)
 - **Call/Jump** destination (BTB)
 - Function **return** destination (RSB)
 - **Load** matches previous store (STL)
- Most are even **shared** among processes

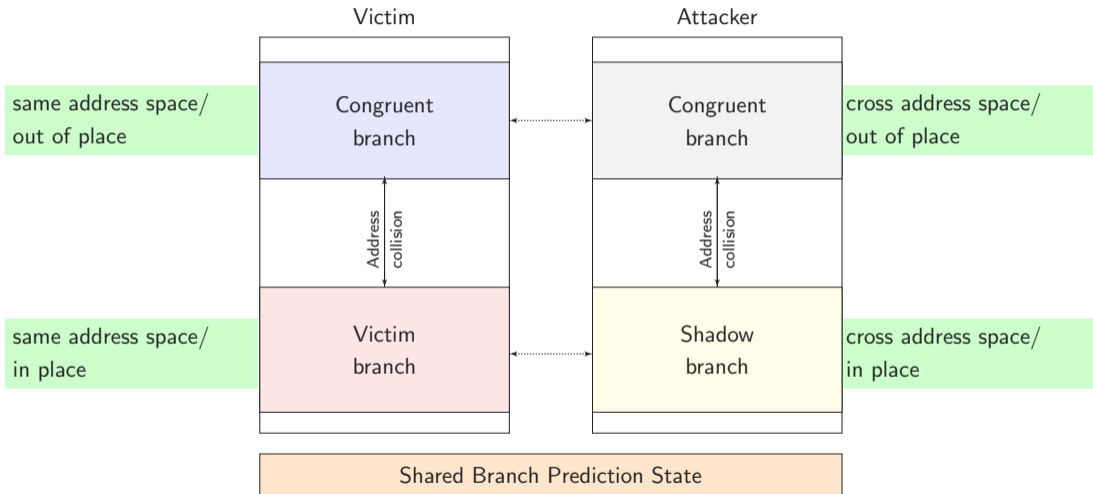














- Questions:



- Questions:
 - Size of speculation window?



- Questions:
 - **Size** of speculation window?
 - **Best** working gadget?



- Questions:
 - **Size** of speculation window?
 - **Best** working gadget?
- Experiment Setup:



- Questions:
 - **Size** of speculation window?
 - **Best** working gadget?
- Experiment Setup:
 - ARM Cortex A-57



- Questions:
 - **Size** of speculation window?
 - **Best** working gadget?
- Experiment Setup:
 - ARM Cortex A-57
 - **Spectre-PHT** with F+R




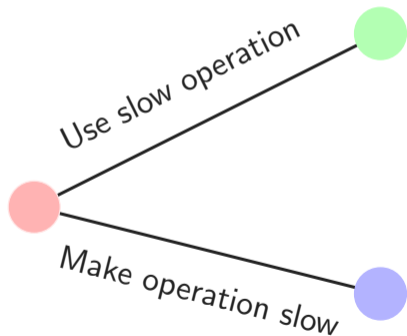
- Questions:
 - **Size** of speculation window?
 - **Best** working gadget?
- Experiment Setup:
 - ARM Cortex A-57
 - **Spectre-PHT** with F+R
 - 1 run: leak 1 byte




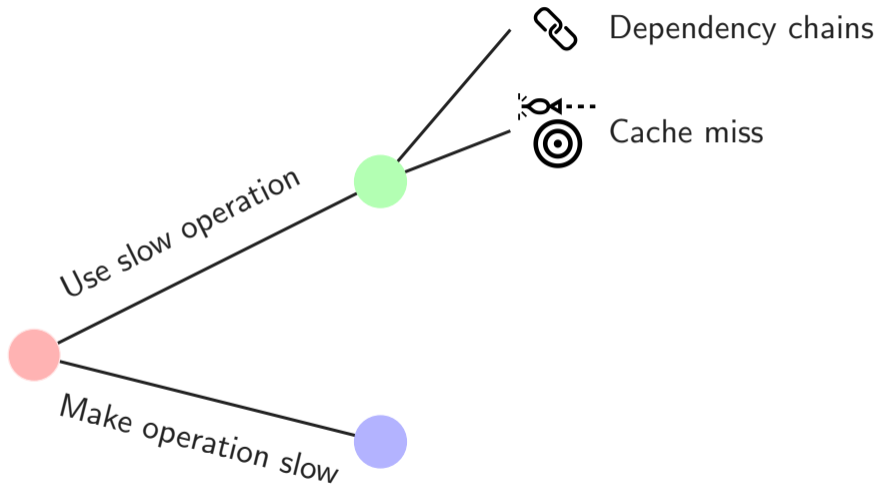
- Questions:
 - **Size** of speculation window?
 - **Best** working gadget?
- Experiment Setup:
 - ARM Cortex A-57
 - **Spectre-PHT** with F+R
 - 1 run: leak 1 byte
 - 1 test: do n runs in m different processes

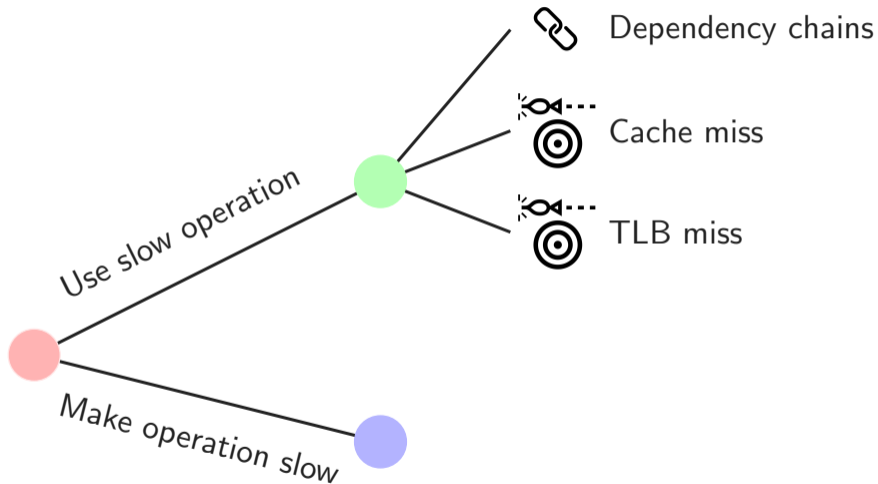


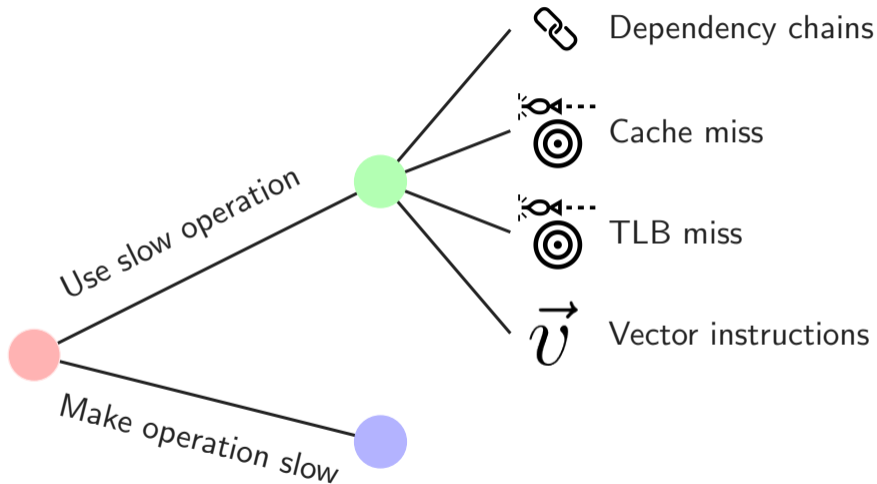


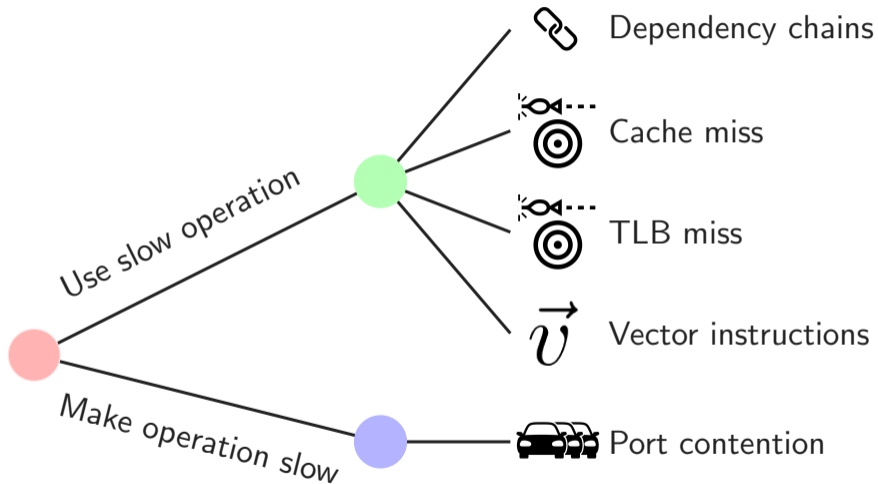


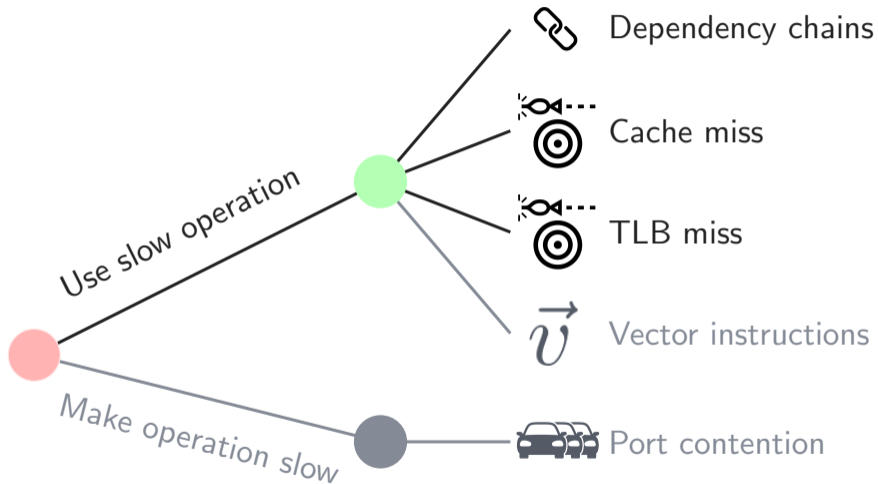


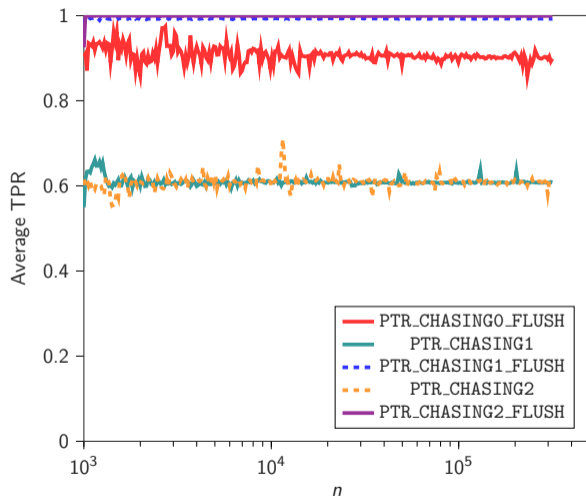


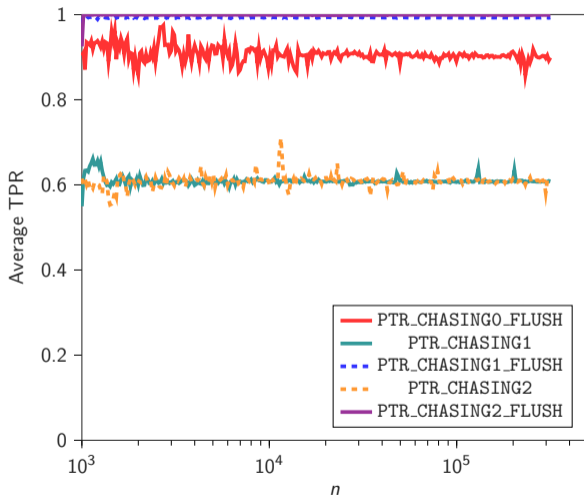












→ $n = 100000$

```
asm volatile("mov x3, #0")
```




```
asm volatile("mov x3, #0"  
"mul x3, x3, x3"  
"mul x3, x3, x3"  
...)
```


```
asm volatile("mov x3, #0"  
"mul x3, x3, x3"  
"mul x3, x3, x3"  
...  
"mov %0, x3"  
: "=r"(res));
```

```
asm volatile("mov x3, #0"  
"mul x3, x3, x3"  
"mul x3, x3, x3"  
...  
"mov %0, x3"  
: "=r"(res));  
if((x + res) < len)  
    oracle[data[x]*4096];
```

```
asm volatile("mov x3, #0"  
"mul x3, x3, x3"  
"mul x3, x3, x3"  
...  
"mov %0, x3"  
: "=r"(res));  
if((x + res) < len)  
oracle[data[x]*4096];
```



```
asm volatile("mov x3, #0"  
"mul x3, x3, x3"  
"mul x3, x3, x3"  
...  
"mov %0, x3"  
: "=r"(res));  
if((x + res) < len)  
    oracle[data[x]*4096];
```



Maximal speculation window size: 57 instructions at a chain length of 22

```
unsigned char value = 0;
char* ptr = &value;
char** ptr2 = &ptr;
char*** ptr3 = &ptr2;
...
```

```
unsigned char value = 0;  
char* ptr = &value;  
char** ptr2 = &ptr;  
char*** ptr3 = &ptr2;  
...
```

```
flush(ptr);  
flush(ptr2);  
flush(ptr3);  
...
```

```
unsigned char value = 0;
char* ptr = &value;
char** ptr2 = &ptr;
char*** ptr3 = &ptr2;
...

flush(ptr);
flush(ptr2);
flush(ptr3);
...


if((x + ***ptr3) < len)
    oracle[data[x]*4096];
```



```
unsigned char value = 0;
char* ptr = &value;
char** ptr2 = &ptr;
char*** ptr3 = &ptr2;
...

flush(ptr);
flush(ptr2);
flush(ptr3);
...


if((x + ***ptr3) < len)
    oracle[data[x]*4096];
```



```
unsigned char value = 0;
char* ptr = &value;
char** ptr2 = &ptr;
char*** ptr3 = &ptr2;
...

flush(ptr);
flush(ptr2);
flush(ptr3);
...

if((x + ***ptr3) < len)
    oracle[data[x]*4096];
```



Maximal speculation window size: 57 instructions at a chain length of 1

```
char* addr = malloc(...);  
*addr = 0;  
...
```

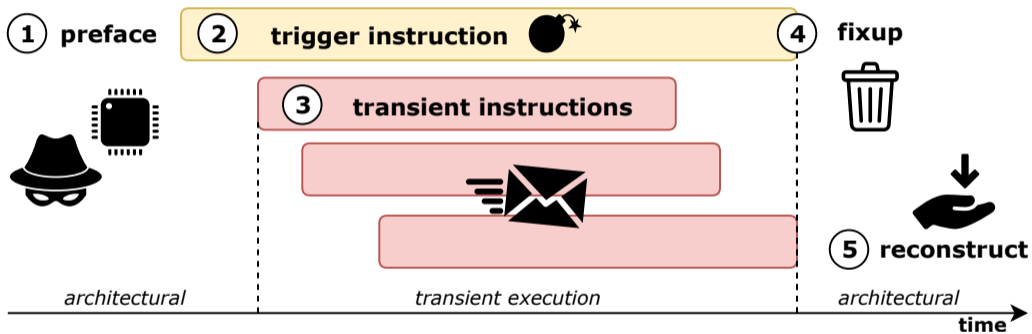
```
char* addr = malloc(...);  
*addr = 0;  
...  
flush_tlb();
```

```
char* addr = malloc(...);
*addr = 0;
...
flush_tlb();
if((x + *addr) < len)
    oracle[data[x] * 4096];
```

```
char* addr = malloc(...);  
*addr = 0;  
...  
flush_tlb();  
if((x + *addr) < len)  
    oracle[data[x] * 4096];
```

Average TPR **0.97**

Speculation window size **57** instructions



Attack

Covert channel

Meltdown-US/RW/GP/NM/PK

Meltdown-P

Meltdown-BR

Spectre-PHT

Spectre-BTB/RSB

Spectre-STL

NetSpectre


Attack	1. Preface
Covert channel	● Flush/Prime/Evict
Meltdown-US/RW/GP/NM/PK	● (Exception suppression)
Meltdown-P	○ (L1 prefetch)
Meltdown-BR	-
Spectre-PHT	● PHT poisoning
Spectre-BTB/RSB	● BTB/RSB poisoning
Spectre-STL	-
NetSpectre	○ Thrash/reset


Attack	1. Preface	2. Trigger example
Covert channel	● Flush/Prime/Evict	-
Meltdown-US/RW/GP/NM/PK	● (Exception suppression)	● <i>mov/rdmsr</i> /FPU
Meltdown-P	○ (L1 prefetch)	● <i>mov</i>
Meltdown-BR	-	○ <i>bound/bndclu</i>
Spectre-PHT	● PHT poisoning	○ <i>jz</i>
Spectre-BTB/RSB	● BTB/RSB poisoning	○ <i>call/jmp/ret</i>
Spectre-STL	-	○ <i>mov</i>
NetSpectre	○ Thrash/reset	○ <i>jz</i>

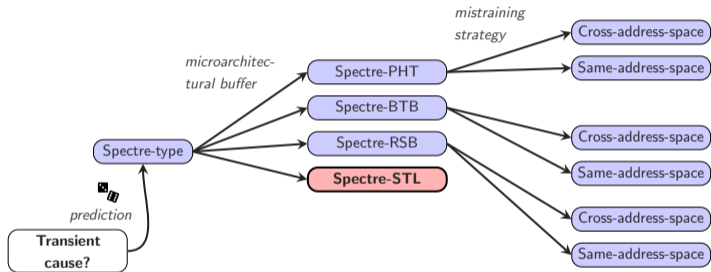
Attack	1. Preface	2. Trigger example	3. Transient
Covert channel	☑ Flush/Prime/Evict	-	☑ Load/AVX/Port/...
Meltdown-US/RW/GP/NM/PK	● (Exception suppression)	● <i>mov/rdmsr</i> /FPU	● Controlled encode
Meltdown-P	○ (L1 prefetch)	● <i>mov</i>	● Controlled encode
Meltdown-BR	-	○ <i>bound/bndclu</i>	○ Inadvertent leak
Spectre-PHT	☑ PHT poisoning	○ <i>jz</i>	○ Inadvertent leak
Spectre-BTB/RSB	☑ BTB/RSB poisoning	○ <i>call/jmp/ret</i>	○ ROP-style encode
Spectre-STL	-	○ <i>mov</i>	○ Inadvertent leak
NetSpectre	○ Thrash/reset	○ <i>jz</i>	○ Inadvertent leak

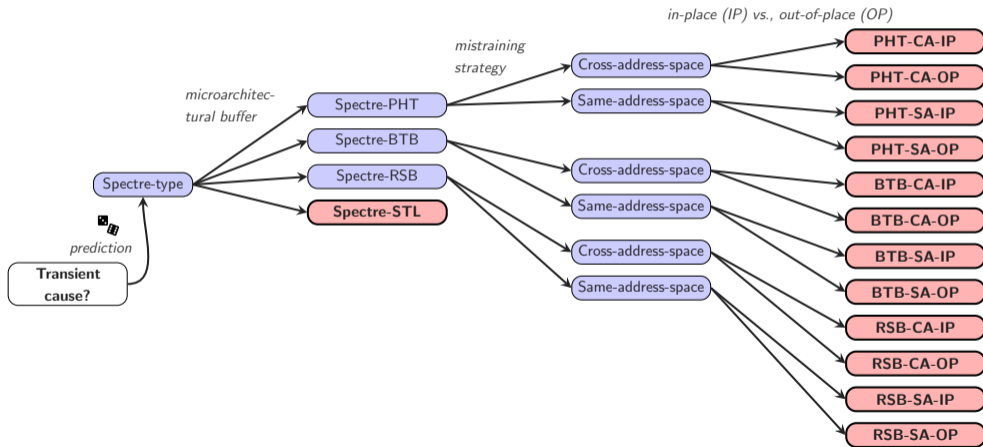
Attack	1. Preface	2. Trigger example	3. Transient	5. Reconstruction
Covert channel	⦿ Flush/Prime/Evict	-	⦿ Load/AVX/Port/...	⦿ Reload/Probe/Time
Meltdown-US/RW/GP/NM/PK	● (Exception suppression)	● <i>mov/rdmsr</i> /FPU	● Controlled encode	● Exception handling
Meltdown-P	○ (L1 prefetch)	● <i>mov</i>	● Controlled encode	& controlled decode
Meltdown-BR	-	○ <i>bound/bndclu</i>	○ Inadvertent leak	<i>same as above</i>
Spectre-PHT	⦿ PHT poisoning	○ <i>jz</i>	○ Inadvertent leak	● Controlled decode
Spectre-BTB/RSB	⦿ BTB/RSB poisoning	○ <i>call/jmp/ret</i>	○ ROP-style encode	● Controlled decode
Spectre-STL	-	○ <i>mov</i>	○ Inadvertent leak	● Controlled decode
NetSpectre	○ Thrash/reset	○ <i>jz</i>	○ Inadvertent leak	○ Inadvertent transmit

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	
ARM	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 (◇).

		Defense																	
Attack		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
ARM	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
		ARM	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 (◇).

		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
Attack																				
ARM	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 (◇).

		Defense																	
Attack		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
ARM	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 (◇).

		Defense																	
Attack		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
ARM	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
		ARM	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 (◇).

		Defense																	
Attack		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
ARM	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**




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



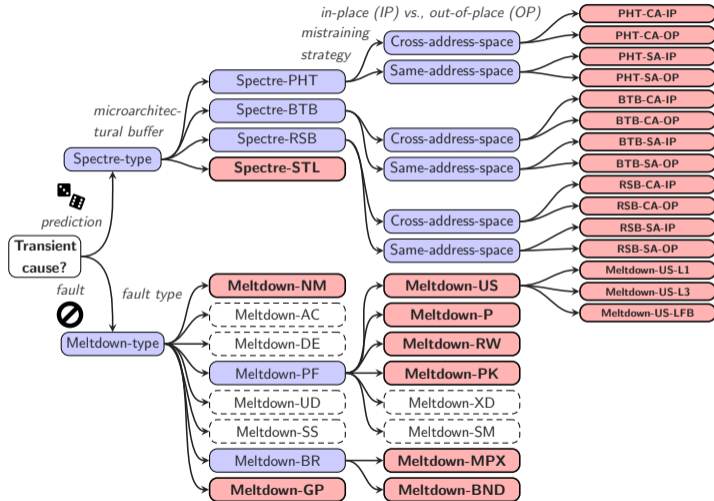
- Current mitigations are either **incomplete or cost performance**
→ More **research** required




- 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

Transient Execution Attacks: Still hARMful?

Barbara Gigerl (@barbarag2112), Claudio Canella (@cc0x1f)

May 17, 2019

Graz University of Technology