


# How the Hardware undermines Software Security


**Daniel Gruss**


March 25, 2019

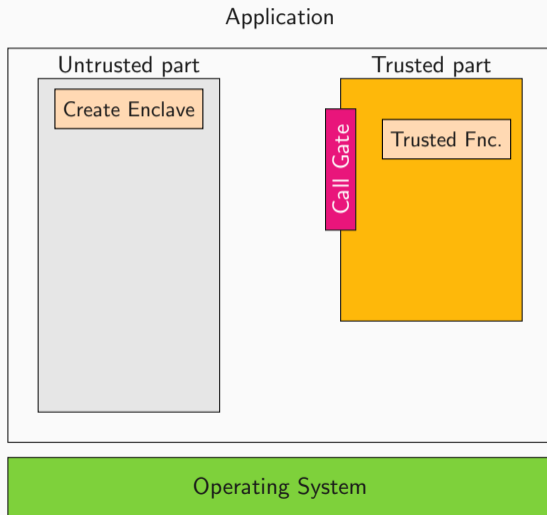
Graz University of Technology



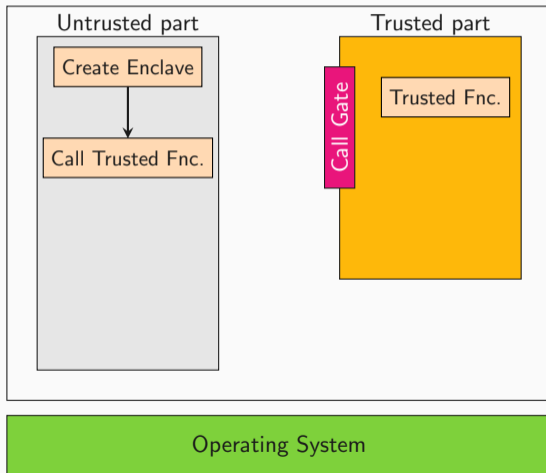
- Instruction-set extension
- Integrity and confidentiality of code and data in untrusted environments
- Run with user privileges and restricted, e.g., no system calls
- Run programs in enclaves using protected areas of memory

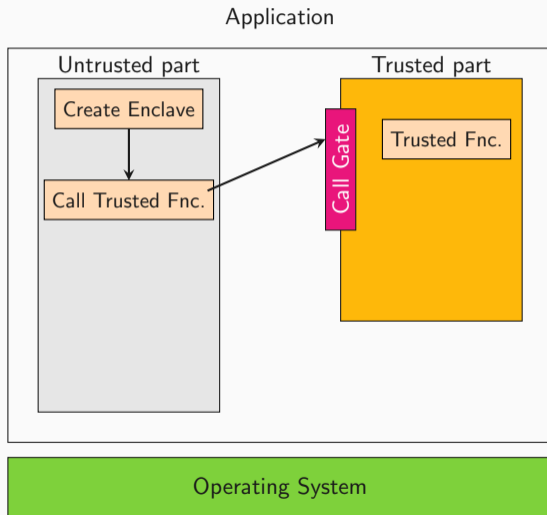


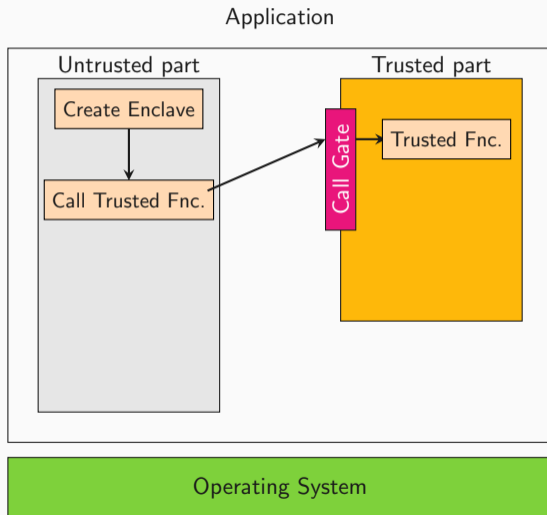




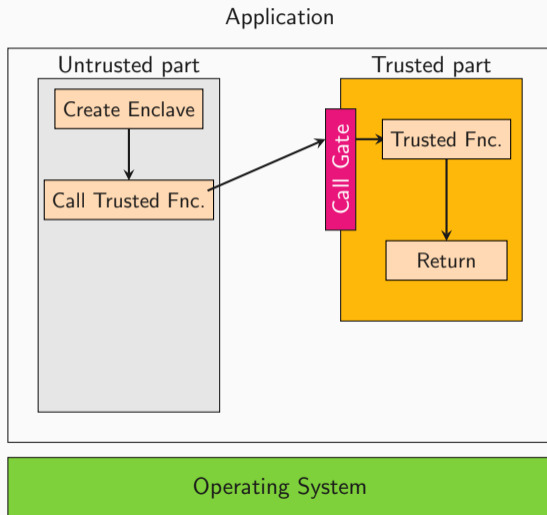
## Application

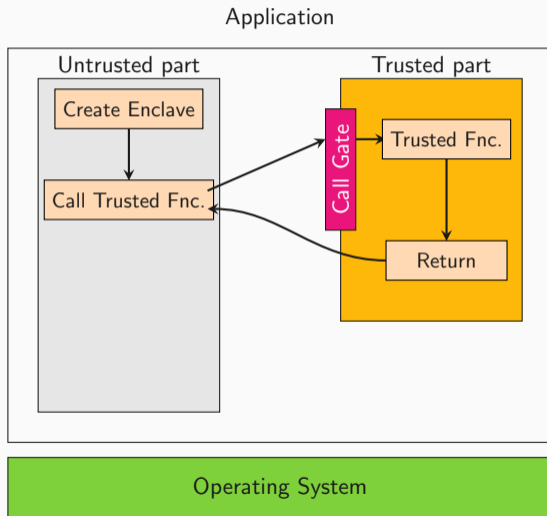


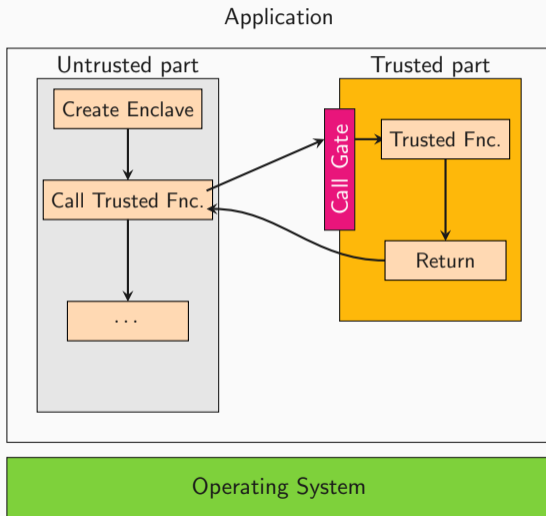













































1337 4242

## FOOD CACHE

**Revolutionary** concept!

Store your food at home,  
never go to the grocery store  
during cooking.

Can store **ALL** kinds of food.



ONLY TODAY INSTEAD OF ~~\$1,300~~

**\$1,299**

ORDER VIA PHONE: +555 12345





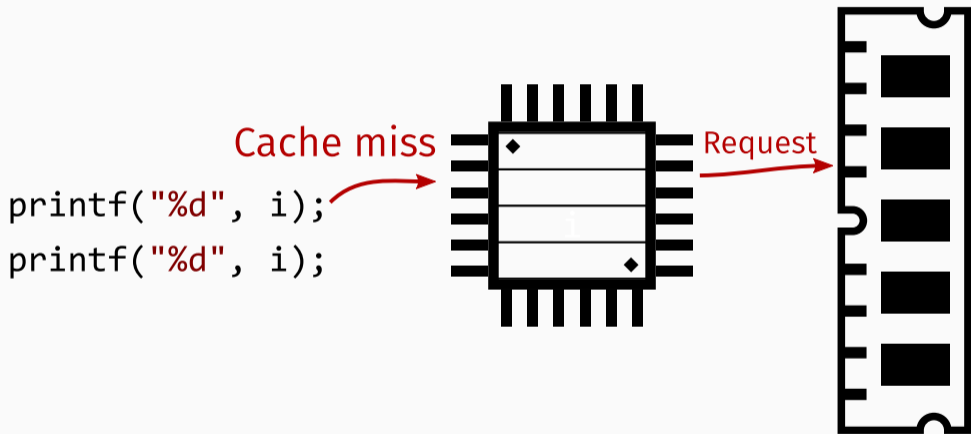
```
printf("%d", i);  
printf("%d", i);
```




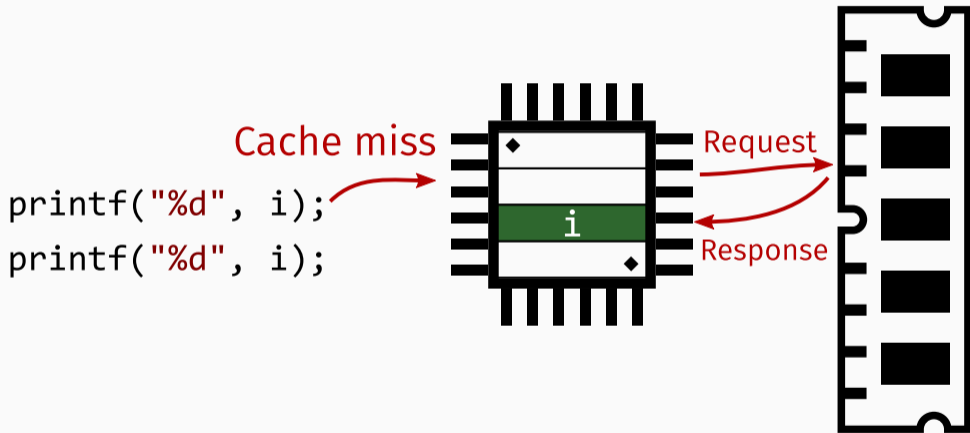
```
printf("%d", i);  
printf("%d", i);
```


Cache miss

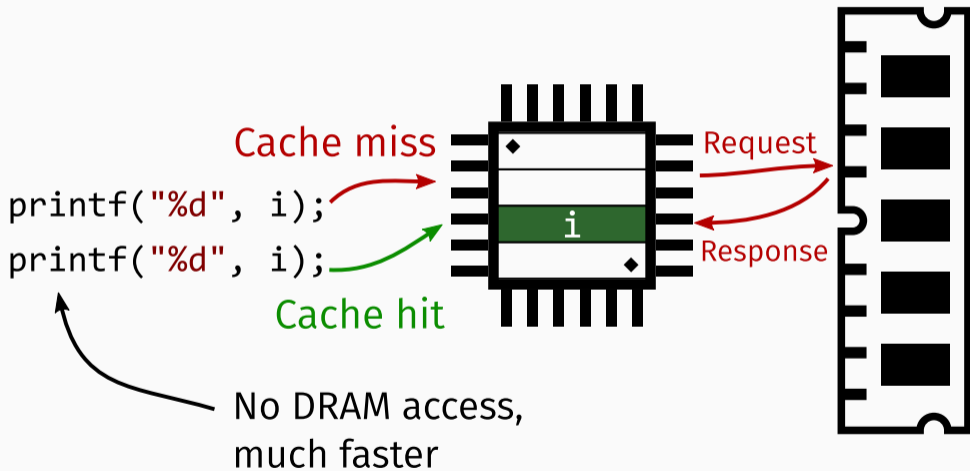











































**FAIL**



- Very short timings
- rdtsc instruction: “cycle-accurate” timestamps

```
[...]  
rdtsc  
function()  
rdtsc  
[...]
```

- Do you measure what you *think* you measure?
- *Out-of-order* execution → what is really executed?

rdtsc

function()

[...]

rdtsc

rdtsc

[...]

rdtsc

function()

rdtsc

rdtsc

function()

[...]

- use pseudo-serializing instruction `rdtscp` (recent CPUs)

- use pseudo-serializing instruction `rdtscp` (recent CPUs)
- and/or use serializing instructions like `cpuid`

- use pseudo-serializing instruction `rdtscp` (recent CPUs)
- and/or use serializing instructions like `cpuid`
- and/or use fences like `mfence`

- use pseudo-serializing instruction `rdtscp` (recent CPUs)
- and/or use serializing instructions like `cpuid`
- and/or use fences like `mfence`


Intel, *How to Benchmark Code Execution Times on Intel IA-32 and IA-64 Instruction Set Architectures White Paper*, December 2010.




AUGUST 22, 2018 BY BRUCE


# Intel Publishes Microcode Security Patches, No Benchmarking Or Comparison Allowed!


UPDATE: **Intel has resolved their microcode licensing issue which I complained about in this blog post.** The new license text is [here](#).
















- Flush+Reload has beautifully nice timings, right?




- Flush+Reload has beautifully nice timings, right?
- Well... steps of 2-4 cycles





- Flush+Reload has beautifully nice timings, right?
- Well... steps of 2-4 cycles
  - only 35-70 steps between hits and misses



- Flush+Reload has beautifully nice timings, right?
- Well... steps of 2-4 cycles
  - only 35-70 steps between hits and misses
- On some devices only 1-2 steps!



- We can build our own timer



- We can build our **own timer**
- Start a thread that continuously increments a global variable



- We can build our **own timer**
- Start a thread that continuously increments a global variable
- The global variable is our **timestamp**





**ARE YOU REALLY EXPECTING TO  
OUTPERFORM THE HARDWARE COUNTER?**

CPU cycles one increment takes

rdtsc  3

```
1 timestamp = rdtsc();
```



CPU cycles one increment takes

rdtsc  3

C

```
1 while(1) {  
2     timestamp++;  
3 }
```

CPU cycles one increment takes

rdtsc  3

C  4.7

```
1 while(1) {  
2   timestamp++;  
3 }
```

CPU cycles one increment takes

rdtsc 3

C 4.7

Assembly

```
1 mov &timestamp, %rcx
2 1: incl (%rcx)
3 jmp 1b
```

CPU cycles one increment takes



```
1 mov &timestamp, %rcx  
2 1: incl (%rcx)  
3 jmp 1b
```

CPU cycles one increment takes

rdtsc 3


C 4.7

Assembly 4.67

Optimized


```
1 mov &timestamp, %rcx
2 1: inc %rax
3 mov %rax, (%rcx)
4 jmp 1b
```

CPU cycles one increment takes

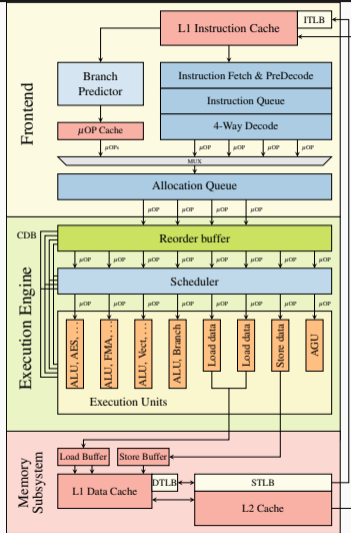


```
1 mov &timestamp, %rcx
2 1: inc %rax
3 mov %rax, (%rcx)
4 jmp 1b
```

WHY?



Meme Heaven







## Protection from Side-Channel Attacks

## Protection from Side-Channel Attacks

Intel SGX does not provide explicit protection from side-channel attacks.

## Protection from Side-Channel Attacks

Intel SGX does not provide explicit protection from side-channel attacks. It is the enclave developer's responsibility to address side-channel attack concerns.

**CAN'T BREAK YOUR SIDE-CHANNEL PROTECTIONS**

**IF YOU DON'T HAVE ANY**



- **Ledger SGX Enclave** for blockchain applications
- **BitPay Copay** Bitcoin wallet
- **Teechain** payment channel using SGX



- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX

### Teechain

[...] We assume the TEE guarantees to hold



- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX

### Teechain

[...] We assume the TEE guarantees to hold and do not consider side-channel attacks [5, 35, 46] on the TEE.






- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX

### Teechain

[...] We assume the TEE guarantees to hold and do not consider side-channel attacks [5, 35, 46] on the TEE. Such attacks and their mitigations [36, 43] are outside the scope of this work. [...]


Raw Prime+Probe trace...<sup>1</sup>



---

<sup>1</sup>Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.


...processed with a simple moving average...<sup>1</sup>



---


<sup>1</sup>Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

...allows to clearly see the bits of the exponent<sup>1</sup>



<sup>1</sup>Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

**YOU CAN'T DO THAT!**



**THAT'S AGAINST THE RULES!**



**Back to Work**

6. Cook everything until  
vegetables are soft

6. Add ground beef  
and mix for 10 minutes

7. *Serve with cooked  
and peeled potatoes*







Wait for an hour



Wait for an hour



LATENCY

1. Wash and cut  
vegetables

2. Pick the basil leaves  
and set aside

3. Heat 2 tablespoons of  
oil in a pan

4. Fry vegetables until  
golden and softened



Dependency


1. Wash and cut vegetables

2. Pick the basil leaves and set aside

3. Heat 2 tablespoons of oil in a pan

4. Fry vegetables until golden and softened

Parallelize



```
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


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);
```





```
*(volatile char*) 0;  
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**



- Out-of-order instructions **leave microarchitectural traces**



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



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



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



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

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



- 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**

**I SHIT YOU NOT**

**THERE WAS KERNEL MEMORY ALL  
OVER THE TERMINAL**

e01d8130: 20 75 73 65 64 20 77 69 74 68 20 61 75 74 68 6f | used with autho  
e01d8140: 72 69 7a 61 74 69 6f 6e 20 66 72 6f 6d 0a 20 53 | rization from. S  
e01d8150: 69 6c 69 63 6f 6e 20 47 72 61 70 68 69 63 73 2c | ilicon Graphics,  
e01d8160: 20 49 6e 63 2e 20 20 48 6f 77 65 76 65 72 2c 20 | Inc. However,  
e01d8170: 74 68 65 20 61 75 74 68 6f 72 73 20 6d 61 6b 65 | the authors make  
e01d8180: 20 6e 6f 20 63 6c 61 69 6d 20 74 68 61 74 20 4d | no claim that M  
e01d8190: 65 73 61 0a 20 69 73 20 69 6e 20 61 6e 79 20 77 | esa. is in any w  
e01d81a0: 61 79 20 61 20 63 6f 6d 70 61 74 69 62 6c 65 20 | ay a compatible  
e01d81b0: 72 65 70 6c 61 63 65 6d 65 6e 74 20 66 6f 72 20 | replacement for  
e01d81c0: 4f 70 65 6e 47 4c 20 6f 72 20 61 73 73 6f 63 69 | OpenGL or associ  
e01d81d0: 61 74 65 64 20 77 69 74 68 0a 20 53 69 6c 69 63 | ated with. Silic  
e01d81e0: 6f 6e 20 47 72 61 70 68 69 63 73 2c 20 49 6e 63 | on Graphics, Inc  
e01d81f0: 2e 0a 20 2e 0a 20 54 68 69 73 20 76 65 72 73 69 | .. .. This versi  
e01d8200: 6f 6e 20 6f 66 20 4d 65 73 61 20 70 72 6f 76 69 | on of Mesa provi  
e01d8210: 64 65 73 20 47 4c 58 20 61 6e 64 20 44 52 49 20 | des GLX and DRI  
e01d8220: 63 61 70 61 62 69 6c 69 74 69 65 73 3a 20 69 74 | capabilities: it  
e01d8230: 20 69 73 20 63 61 70 61 62 6c 65 20 6f 66 0a 20 | is capable of.  
e01d8240: 62 6f 74 68 20 64 69 72 65 63 74 20 61 6e 64 20 | both direct and  
e01d8250: 69 6e 64 69 72 65 63 74 20 72 65 6e 64 65 72 69 | indirect renderi  
e01d8260: 6e 67 2e 20 20 46 6f 72 20 64 69 72 65 63 74 20 | ng. For direct  
e01d8270: 72 65 6e 64 65 72 69 6e 67 2c 20 69 74 20 63 61 | rendering, it ca  
e01d8280: 6e 20 75 73 65 20 44 52 49 0a 20 6d 6f 64 75 6c | n use DRI. modul  
e01d8290: 65 73 20 66 72 6f 6d 20 74 68 65 20 6c 69 62 67 | es from the libg






**K**ernel **A**ddress **I**solation to have **S**ide channels **E**fficiently **R**emoved

**KAISER** /'kAIZə/

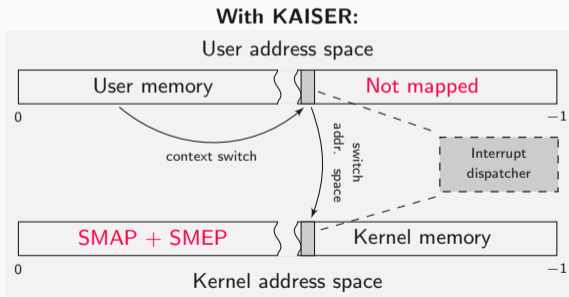

1. [german] Emperor, ruler of an empire
2. largest penguin, emperor penguin




**K**ernel **A**ddress **I**solation to have **S**ide channels **E**fficiently **R**emoved














- Our patch
- Adopted in Linux




- Our patch
- Adopted in Linux



- Adopted in Windows




- Our patch
- Adopted in Linux




- Adopted in Windows



- Adopted in OSX/iOS



- Our patch
- Adopted in Linux



- Adopted in Windows



- Adopted in OSX/iOS

→ **now in every computer**




**PIZZA**


*SPECIAL RECIPES*

















*»A table for 6 please«*





# Speculative Cooking





»A table for 6 please«





**PIZZA**

*SPECIAL RECIPES*



**PIZZA**


*SPECIAL RECIPES*


**Pizza**

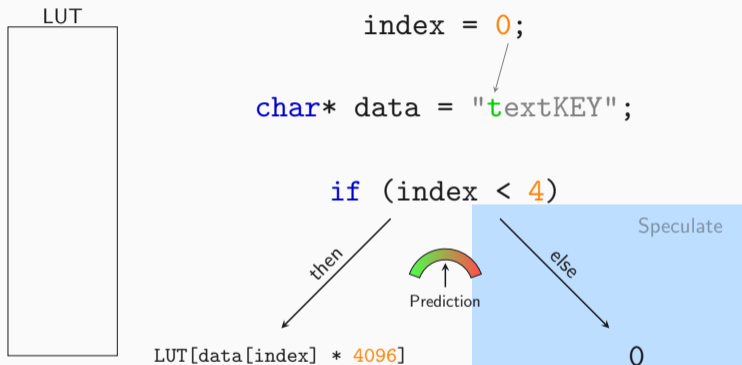




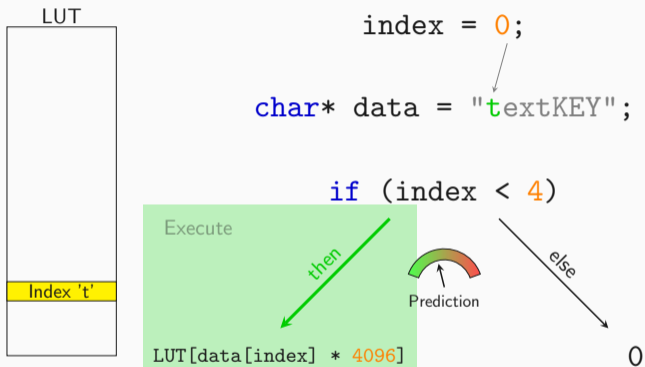






















































































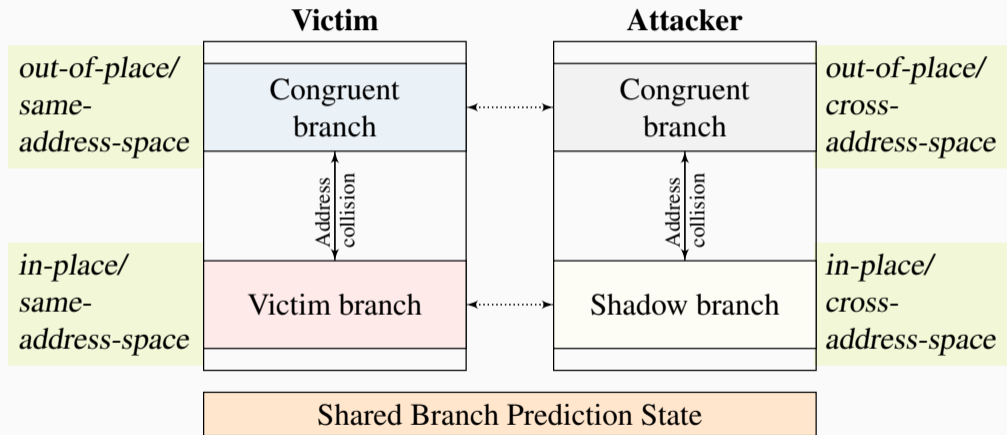


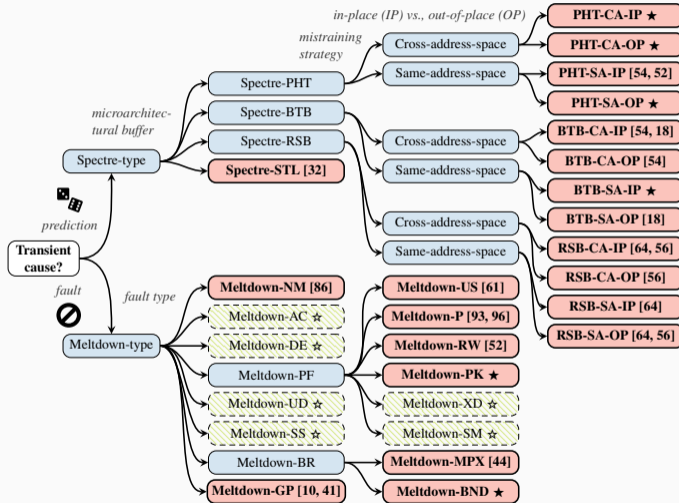














**BLOCKCHAIN**

**H** HD  
HISTORY.COM



## Computer Architecture Today

Informing the broad computing community about current activities, advances and future directions in computer architecture.

### Let's Keep it to Ourselves: Don't Disclose Vulnerabilities

by Gus Uht on Jan 31, 2019 | Tags: Opinion, Security



#### CONTRIBUTE

Editor: Alvin R. Lebeck

Associate Editor: Vijay Janapa Reddi

Contribute to Computer  
Architecture Today

**Table 1:** Spectre-type defenses and what they mitigate.

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	
		Attack																			
Intel	Spectre-PHT	□	□	□	◇	◇	◇	●	●	●	○	◇	◇	◇	◇	●	■	●	■	◇	
	Spectre-BTB	□	□	□	◇	●	◇	◇	◇	◇	◇	◇	●	●	◇	◇	■	●	◇	◇	
	Spectre-RSB	□	□	□	●	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	●	◇	◇	
	Spectre-STL	□	□	□	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	●	■	●	
ARM	Spectre-PHT	□	□	□	◇	◇	◇	●	●	●	○	◇	◇	◇	◇	◇	●	■	●	■	◇
	Spectre-BTB	□	□	□	◇	●	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	●	◇	◇	
	Spectre-RSB	□	□	□	●	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	●	◇	◇	
	Spectre-STL	□	□	□	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	●	■	●	
AMD	Spectre-PHT	□	□	□	◇	◇	◇	●	●	●	○	◇	◇	◇	◇	◇	●	■	●	■	◇
	Spectre-BTB	□	□	□	◇	●	◇	◇	◇	◇	◇	■	■	■	◇	■	●	◇	◇		
	Spectre-RSB	□	□	□	●	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	◇	■	●	◇	◇	
	Spectre-STL	□	□	□	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	◇	■	●	■	●	





**Table 2:** Reported performance impacts of countermeasures


<b>Defense</b> \ <b>Impact</b>	Performance Loss	Benchmark
InvisiSpec	22%	SPEC
SafeSpec	3% (improvement)	SPEC2017 on MARSSx86
DAWG	2–12%, 1–15%	PARSEC, GAPBS
RSB Stuffing	no reports	
Retpoline	5–10%	real-world workload servers
Site Isolation	only memory overhead	
SLH	36.4%, 29%	Google microbenchmark suite
YSNB	60%	Phoenix
IBRS	20–30%	two sysbench 1.0.11 benchmarks
STIPB	30– 50%	Rodinia OpenMP, DaCapo
IBPB	no individual reports	
Serialization	62%, 74.8%	Google microbenchmark suite
SSBD/SSBB	2–8%	SYSmark®2014 SE & SPEC integer
KAISER/KPTI	0–2.6%	system call rates
L1TF mitigations	-3–31%	various SPEC

What if we want to modify data?











chip




chip




64k cells  
1 capacitor,  
1 transistor each

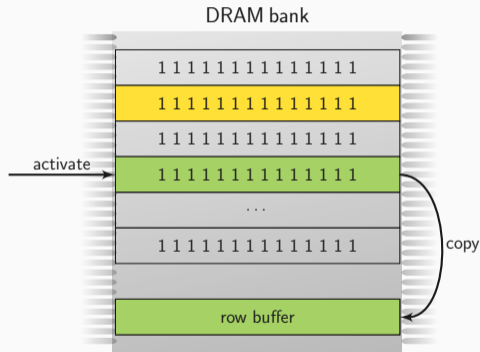





- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer




- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer




- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer



- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer



- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer



- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer


- There are two different hammering techniques

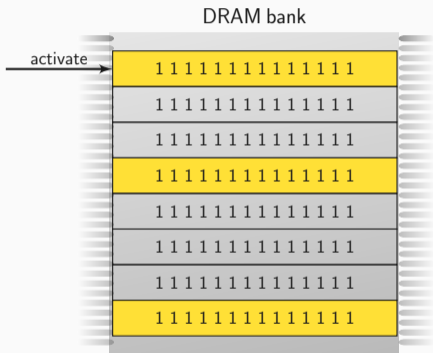
- There are two different hammering techniques
- #1: Hammer one row next to victim row and other random rows




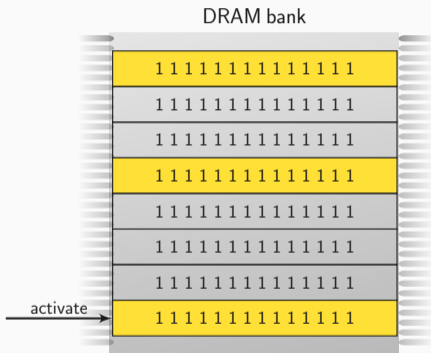
- There are two different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row


- There are **three** different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row
- #3: Hammer only one row next to victim row





















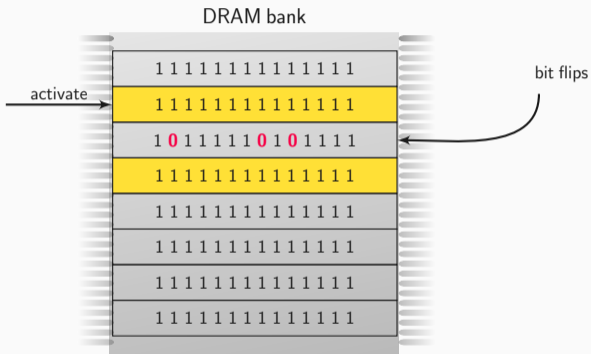








































- They are not random → highly reproducible flip pattern!





- They are not random → highly reproducible flip pattern!
  1. Choose a data structure that you can place at arbitrary memory locations



- They are not random → highly reproducible flip pattern!
  1. Choose a data structure that you can place at arbitrary memory locations
  2. Scan for “good” flips





- They are not random → highly reproducible flip pattern!
  1. Choose a data structure that you can place at arbitrary memory locations
  2. Scan for “good” flips
  3. Place data structure there




- They are not random → highly reproducible flip pattern!
  1. Choose a data structure that you can place at arbitrary memory locations
  2. Scan for “good” flips
  3. Place data structure there
  4. Trigger bit flip again











- Many applications perform actions as root







- Many applications perform actions as root



- Many applications perform actions as root
- They can be used by unprivileged users as well



- Many applications perform actions as root
- They can be used by unprivileged users as well



- Many applications perform actions as root
- They can be used by unprivileged users as well
- `sudo`

























- If a binary is loaded the first time, it is loaded to the memory



- If a binary is loaded the first time, it is loaded to the memory
- It stays in memory (in the page cache) even after execution



- If a binary is loaded the first time, it is loaded to the memory
- It stays in memory (in the page cache) even after execution
- Only evicted if page cache is full



- If a binary is loaded the first time, it is loaded to the memory
- It stays in memory (in the page cache) even after execution
- Only evicted if page cache is full
- Page cache is huge - usually all unused memory







# MEMORY WAYLAYING

Wait for the right moment, and then hit it with a bit flip!

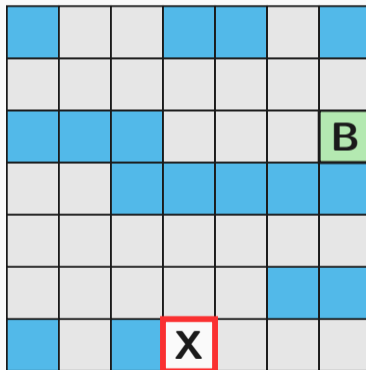
(1) Start




## (2) Evict Page Cache




## (3) Access Binary



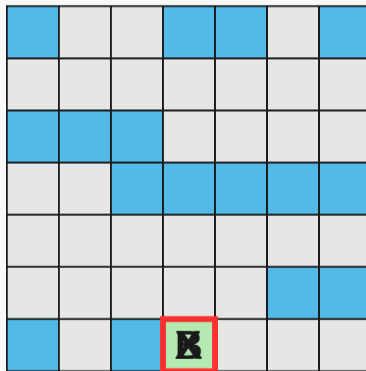
## (4) Evict + Access




(5) Evict + Access



(6) Stop if target reached




- New pages cover most of the physical memory







- Great advantage over memory massaging: only negligible memory footprint




**Rowhammer + SGX = Cheap Denial of Service**




- What happens if a bit flips in the SGX EPC?




- What happens if a bit flips in the SGX EPC?
- Integrity check will fail!




- What happens if a bit flips in the SGX EPC?
  - Integrity check will fail!
- Locks up the memory controller



- What happens if a bit flips in the SGX EPC?
- Integrity check will fail!
- Locks up the memory controller
- Not a single further memory access!



- What happens if a bit flips in the SGX EPC?
- Integrity check will fail!
- Locks up the memory controller
- Not a single further memory access!
- System halts immediately



- What happens if a bit flips in the SGX EPC?
- Integrity check will fail!
- Locks up the memory controller
- Not a single further memory access!
- System halts immediately





**SOUNDS UNSAFE?**

**IT IS UNSAFE!**



- If a malicious enclave induces a bit flip, ...



- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts



- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts
- ... including co-located tenants




- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts
- ... including co-located tenants
- **Denial-of-Service Attacks in the Cloud** [Gru+18; Jan+17]

**SGX + One-location Hammering + Opcode Flipping =  
Undetectable Exploit**




**STEALTH LEVEL: EXPERT**





Defense Class	<i>Static Analysis</i>	<i>Performance Counters</i>	<i>Memory Access Pattern</i>	<i>Physical Proximity</i>	<i>Memory footprint</i>
Bypass					
Intel SGX	●	●	○	○	○
One-location hammering	○	○	●	○	○
Opcode flipping	○	○	○	●	○
Memory waylaying	○	○	○	○	●
<b>Defense class defeated</b>	●	●	●	●	●



**AT LEAST IT'S A LOCAL ATTACK**

We have ignored microarchitectural attacks for many years:





We have ignored microarchitectural attacks for many years:

- attacks on crypto



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone





We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone → “not part of the threat model”



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone → “not part of the threat model”
- Rowhammer



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone → “not part of the threat model”
- Rowhammer → “only affects cheap sub-standard modules”



We have ignored microarchitectural attacks for many years:

- attacks on crypto → “software should be fixed”
- attacks on ASLR → “ASLR is broken anyway”
- attacks on SGX and TrustZone → “not part of the threat model”
- Rowhammer → “only affects cheap sub-standard modules”

→ for years we solely optimized for performance



- lower refresh rate = lower energy but more bit flips



- lower refresh rate = lower energy but more bit flips
- ECC memory → fewer bit flips



- lower refresh rate = lower energy but more bit flips
  - ECC memory → fewer bit flips
- it's an optimization problem



- lower refresh rate = lower energy but more bit flips
  - ECC memory → fewer bit flips
- it's an optimization problem
- what if “too aggressive” changes over time?





- lower refresh rate = lower energy but more bit flips
  - ECC memory → fewer bit flips
- it's an optimization problem
- what if “too aggressive” changes over time?
- difficult to optimize with an intelligent adversary



- new class of software-based attacks



- new class of software-based attacks
- many problems to solve around microarchitectural attacks and especially transient execution attacks



- new class of software-based attacks
- many problems to solve around microarchitectural attacks and especially transient execution attacks
- dedicate more time into identifying problems and not solely in mitigating known problems

# How the Hardware undermines Software Security

**Daniel Gruss**

March 25, 2019

Graz University of Technology

## References

---



Daniel Gruss et al. Another Flip in the Wall of Rowhammer Defenses. In: S&P. 2018.



Yeongjin Jang et al. SGX-Bomb: Locking Down the Processor via Rowhammer Attack. In: SysTEX. 2017.



Michael Schwarz et al. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.