

Security:


Can we afford to have it?

Can we afford not to have it?

Daniel Gruss

2023-10-09

Graz University of Technology

A man in a blue long-sleeved shirt is sitting at a black table outdoors on a brick-paved area. He is holding a black mug. On the table, there is a microphone on a stand, a black mug, and some papers. A white sign is attached to the front of the table with black text. The background shows a park-like setting with trees and a building.

side channel
= obtaining meta-data and
deriving secrets from it

CHANGE MY MIND



- Profiling cache utilization with performance counters?



- Profiling cache utilization with performance counters? → No






- Profiling cache utilization with performance counters? → No
- Observing cache utilization with performance counters and using it to infer a crypto key?




- Profiling cache utilization with performance counters? → No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes



- Profiling cache utilization with performance counters? → No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload?



- Profiling cache utilization with performance counters? → No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload? → No



- Profiling cache utilization with performance counters? → No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload? → No
- Measuring memory access latency with Flush+Reload and using it to infer keystroke timings?



- Profiling cache utilization with performance counters? → No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload? → No
- Measuring memory access latency with Flush+Reload and using it to infer keystroke timings? → Yes

A close-up portrait of Morpheus from the movie The Matrix. He is bald and has a serious expression. He is wearing dark sunglasses. The reflections in the sunglasses show Neo, Trinity, and Morpheus in a scene from the movie. The background is a blurred outdoor setting.

THERE IS NO NOISE

NOISE IS JUST SOMEONE ELSE'S DATA









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




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

Cache miss




Request




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


Cache miss

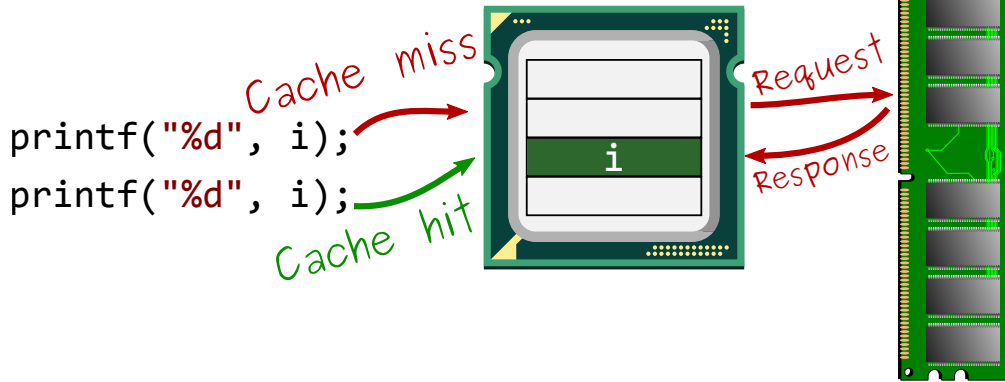


Request

Response








DRAM access,
slow

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

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


Cache miss

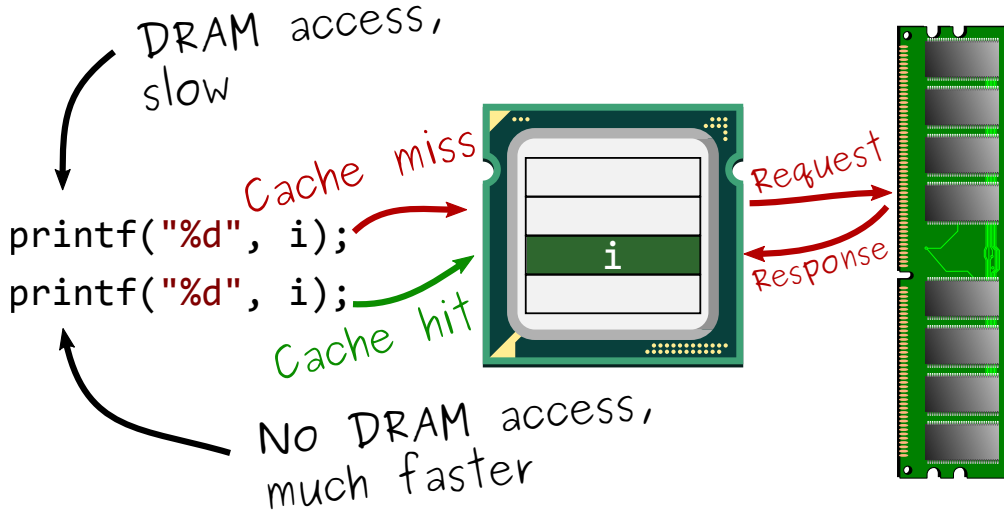
Cache hit





Request


Response





































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


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



- Add a **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**






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

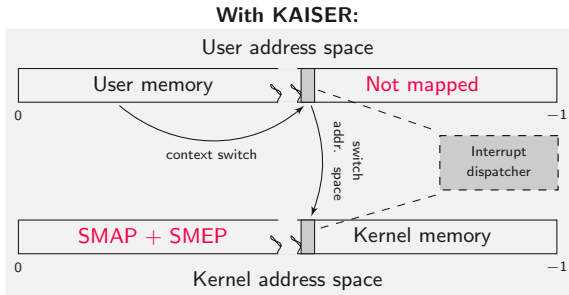

KAISER /'kAIZə/

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



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









PIZZA


SPECIAL RECIPES















»A table for 6 please«





Speculative Cooking



»A table for 6 please«





PIZZA

SPECIAL RECIPES



PIZZA


SPECIAL RECIPES


Pizza















































































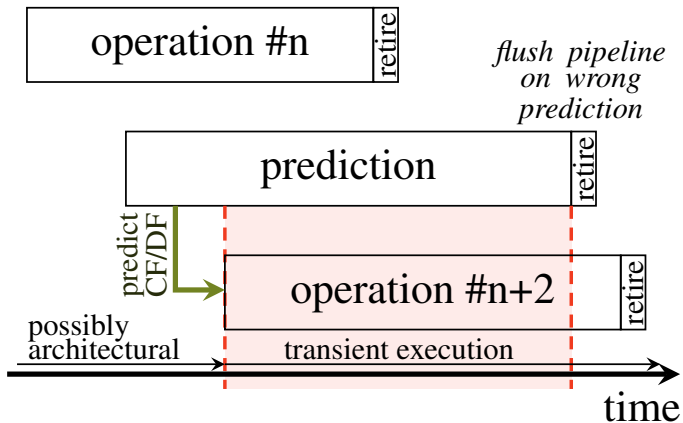


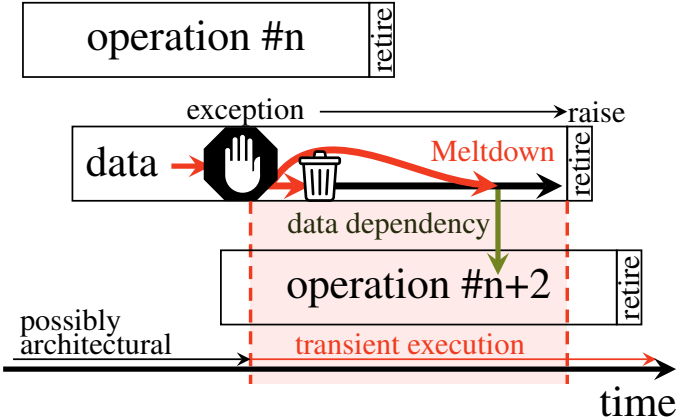












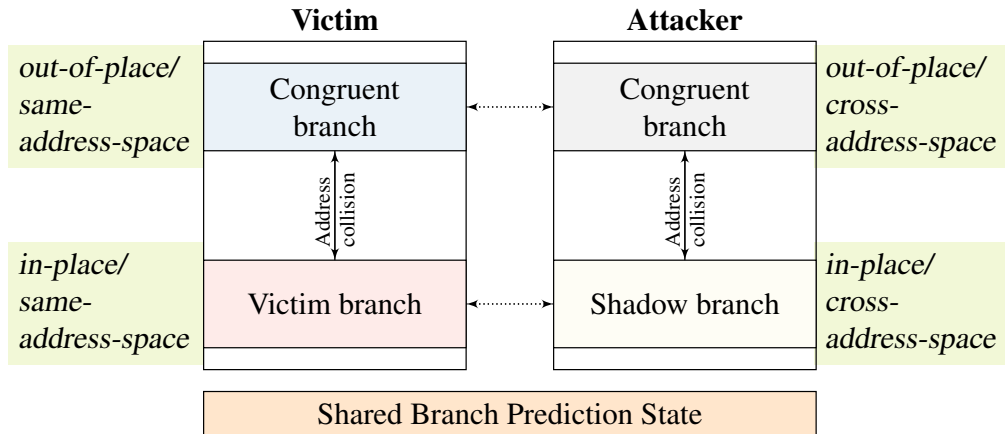


Table 1: Reported performance impacts of countermeasures

Defense \ Impact	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

Test - Mozilla Firefox (on lab02)

Test x +

file:///home/dgruss/rowhammerjs/rowhammer.html Search



320: 12
330: 9
340: 1
350: 0
360: 1
370: 2
380: 199
390: 76
400: 72
410: 231
420: 572
1250



[!] Found flip (254 != 255) at array index 340021386 when hammering indices 339881984 and 340156416



[!] Found flip (239 != 255) at array index 340022176 when hammering indices 339881984 and 340156416

[!] Found flip (191 != 255) at array index 340023138 when hammering indices 339881984 and 340156416



[!] Found flip (254 != 255) at array index 340025146 when hammering indices 339881984 and 340156416









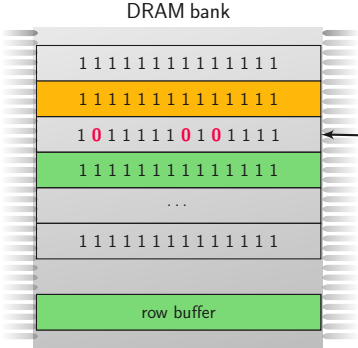
Cells leak faster upon proximate accesses → Rowhammer




Cells leak faster upon proximate accesses → Rowhammer




Cells leak faster upon proximate accesses → Rowhammer

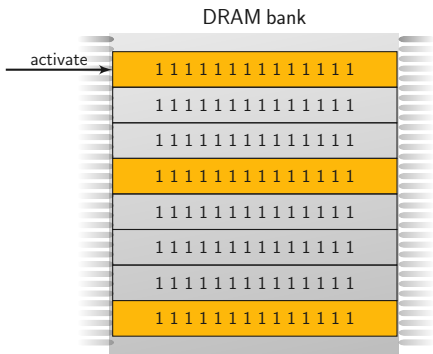



bit flips in row 2!





Cells leak faster upon proximate accesses → Rowhammer



















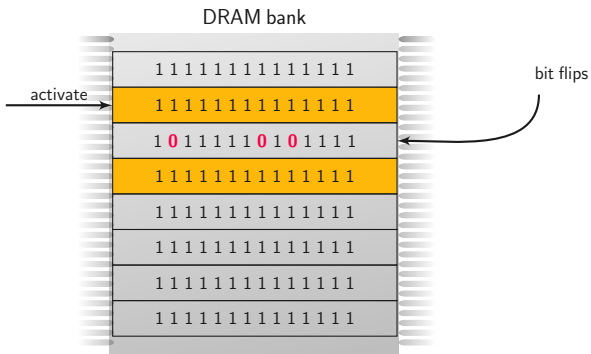















**HAMMERING
TWO ROWS**





**HAMMERING
TWO ROWS**





**HAMMERING
A SINGLE ROW**













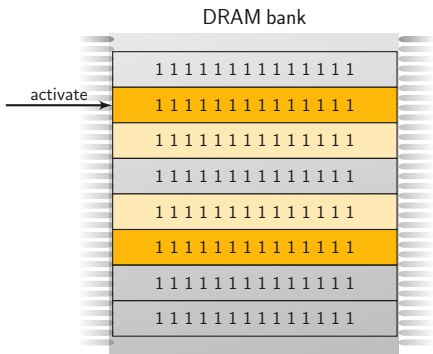














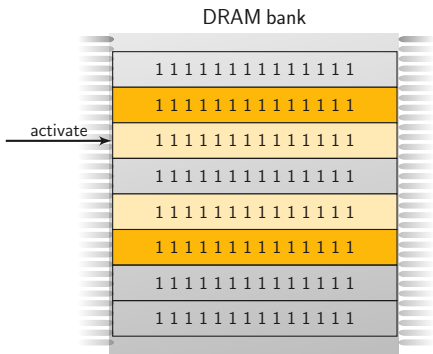


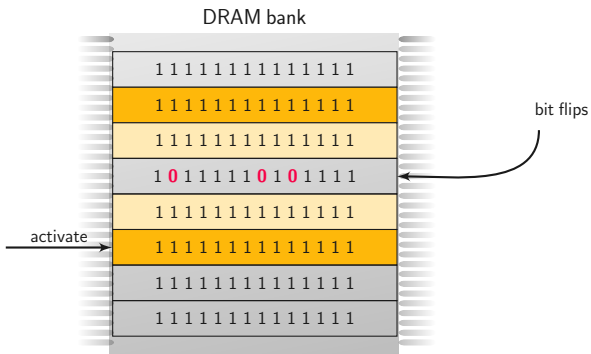





























```
uint64_t multiplier = 0x1122334455667788;
uint64_t correct    = 0xdeadbeef * multiplier;
uint64_t var        = 0xdeadbeef * multiplier;

while (var == correct)
{
    var = 0xdeadbeef * multiplier;
}
uint64_t flipped_bits = var ^ correct;
```

```
do
{
    i++;
    plaintext = <randomly generated>

    result1 = aes128_enc(plaintext);
    result2 = aes128_enc(plaintext);
} while (vec_equal_128(result1,result2) && i<iterations);
```




- Should be related to **undervolting**



- Should be related to **undervolting**
- From protected TEE **vaults**



- Should be related to **undervolting**
- From protected TEE **vaults**
- Steal





- Should be related to **undervolting**
- From protected TEE **vaults**
- Steal, corrupt



- Should be related to **undervolting**
- From protected TEE **vaults**
- Steal, corrupt, **plunder**, ...


















ENERGY FORECAST

20.9% of projected electricity demand

- Networks
- Production of ICT
- Consumer devices
- Data centres



0.09%

0.40%

There are alternatives

There are alternatives to security!

How expensive is security?

RACE FOR A VACCINE

WE WANT TO
BE FIRST

WE WANT TO
BE FIRST!

WE WANT
TO BE
FIRST!

WE WANT TO
BE FIRST!



RACE TO ACT ON CLIMATE

YOU GO
FIRST!

WHY SHOULD
WE BE FIRST

NO, YOU
GO FIRST

THEY SHOULD
GO FIRST!



Garino



FUNCTIONALITY




SECURITY


Moore's Law: The number of transistors on microchips has doubled every two years.

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important for other aspects of technological progress in computing - such as processing speed or the price of computers.

Transistor count



Data source: Wikipedia (wikipedia.org/wiki/Transistor_count)
OurWorldinData.org - Research and data to make progress against the world's largest problems.



THE
END
IS
NEAR

OF MOORE'S LAW

Before

System Information Benchmark your Current Settings

Advanced Tuning

Cache

Other

Stress Test

Benchmarking

Profiles

App-Profile Pairing

Current Score

XTU: 1921 Marks

Maximum Processor Frequency: 4.15 GHz

Highest CPU Temperature: 96 °C

Share Online

Run XTU Benchmark

After

System Information Benchmark your Current Settings

Advanced Tuning

Cache

Other

Stress Test

Benchmarking

Profiles

App-Profile Pairing

Current Score

XTU: 2122 Marks

Maximum Processor Frequency: 4.13 GHz


Highest CPU Temperature: 95 °C

Compare Online

Run XTU Benchmark

I7-9750H

CPU Undervolting



Huge difference!!!


**IF MY SYSTEMS RAN UNDERVOLTED
TOTALLY FINE FOR 10 YEARS**




**WHY DO WE
WASTE 40% ENERGY?**




Why are problems like Rowhammer not solved already?




... create bad incentives.






... create bad incentives.

- A “bit” more reliability




... create bad incentives.

- A “bit” more reliability
- Why not higher or dynamic refresh rates everywhere (e.g. TRR, PARA, ...)?




... create bad incentives.

- A “bit” more reliability
- Why not higher or dynamic refresh rates everywhere (e.g. TRR, PARA, ...)?
 - “just a few more targeted refreshes”




... create bad incentives.


- A “bit” more reliability
- Why not higher or dynamic refresh rates everywhere (e.g. TRR, PARA, ...)?
 - “just a few more targeted refreshes”
- Why not ECC everywhere?




... create bad incentives.

- A “bit” more reliability
 - Why not higher or dynamic refresh rates everywhere (e.g. TRR, PARA, ...)?
 - “just a few more targeted refreshes”
 - Why not ECC everywhere?
- What incentives does it create?






Fundamental problem: we assume what is still reliable




Fundamental problem: we assume what is still reliable

- Refreshing x times per second is fine




Fundamental problem: we assume what is still reliable

- Refreshing x times per second is fine
- Normal usage, no adversary



Fundamental problem: we assume what is still reliable

- Refreshing x times per second is fine
- Normal usage, no adversary
- Assume there won't be more than n bit errors



Fundamental problem: we assume what is still reliable

- Refreshing x times per second is fine
- Normal usage, no adversary
- Assume there won't be more than n bit errors

→ How far can we go with x while staying below n bit errors?



AFTER ALL THESE REFRESHES, WHY NOT

**WHY SHOULDN'T I OPTIMIZE
PERFORMANCE TO THE ABSOLUTE LIMIT?**

imgflip.com



Mobile vendors since 2018: let's add ECC by default



Mobile vendors since 2018: let's add ECC by default

- ECC memory → fewer bit flips + more security



Mobile vendors since 2018: let's add ECC by default

- ECC memory → fewer bit flips + more security

Also vendors:



Mobile vendors since 2018: let's add ECC by default

- ECC memory → fewer bit flips + more security

Also vendors:

- Let's squeeze out the last bit of efficiency for battery runtime until just before bit flips occur







- You never know how far is still safe



- You never know how far is still safe
- “safe” / “reliable” changes over time




- You never know how far is still safe
- “safe” / “reliable” changes over time
- Adversary is intelligent and improves attacks over time

Security vs Reliability


Security vs Reliability

YOU CAN'T DO THAT!




THAT'S AGAINST THE RULES!

Security for Efficiency?




Make bit flips degrade performance **without** impacting security




Make bit flips degrade performance **without** impacting security

- Cryptographic MAC



Make bit flips degrade performance **without** impacting security

- Cryptographic MAC
- Detect **any** number of bit flips



Make bit flips degrade performance **without** impacting security

- Cryptographic MAC
- Detect **any** number of bit flips
- Correction by **brute-force** search for correct data

# Errors	# MAC Comp.	Avg Duration
1	17	11 ns
2	771	3.68 μ s
3	33 800	124 μ s
4	1.51×10^6	6.65 ms
5	6.91×10^7	261 ms
6	3.07×10^9	12.8 s
7	1.21×10^{11}	9.11 min
8	5.72×10^{12}	6.11 h







- Silent data corruption less than once per 10^9 billion years



- Silent data corruption less than once per 10^9 billion years
- Second preimage after hammering for one year: $9.75 \cdot 10^{-5} \%$



- Silent data corruption less than once per 10^9 billion years
- Second preimage after hammering for one year: $9.75 \cdot 10^{-5} \%$
- Erroneous correction of 8-bit errors: 0.0161 %

Security:

Can we afford to have it?

Can we afford not to have it?

Daniel Gruss

2023-10-09

Graz University of Technology