



CSI:Rowhammer

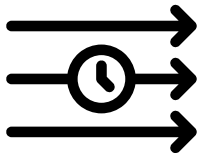
Closing the Case of Half-Double and Beyond

BlackHat Europe 2022

7th December 2022

Andreas Kogler
Graz University of Technology

Jonas Juffinger
Graz University of Technology



Since 2014



Numerous **exploits**



Mitigations
ineffektive



Something
fundamentally
better?

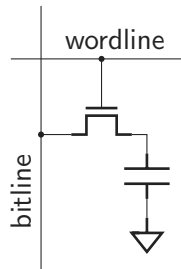
What if we could **transparently correct arbitrary bitflips** if we know that the OS can reconstruct the data from a different source?

Introduction

- **Large, cheap** and **energy efficient**

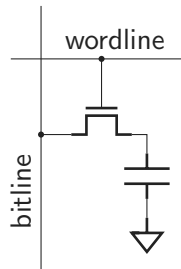
Dynamic Random Access Memory (DRAM)

- **Large, cheap** and **energy efficient**
- Cell: one transistor & one capacitor



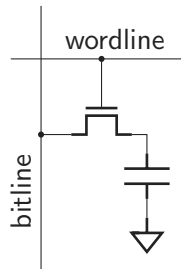
Dynamic Random Access Memory (DRAM)

- **Large, cheap** and **energy efficient**
- Cell: one transistor & one capacitor
- Capacitor loses charge over time



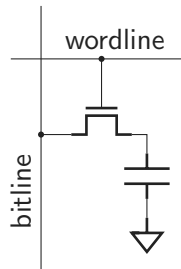
Dynamic Random Access Memory (DRAM)

- **Large, cheap** and **energy efficient**
- Cell: one transistor & one capacitor
- Capacitor loses charge over time
- Periodic refreshes **required**

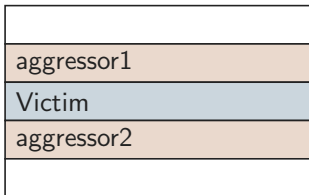


Dynamic Random Access Memory (DRAM)

- **Large, cheap** and **energy efficient**
- Cell: one transistor & one capacitor
- Capacitor loses charge over time
- Periodic refreshes **required**
- Organized in **rows**

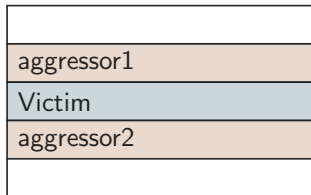


- Hardware fault of the DRAM [3]



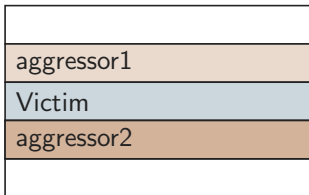
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



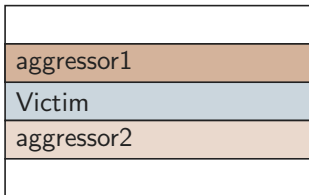
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



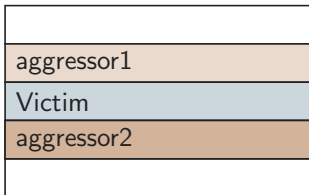
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



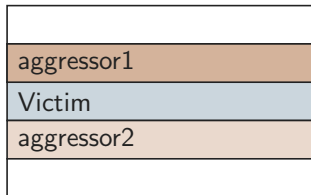
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



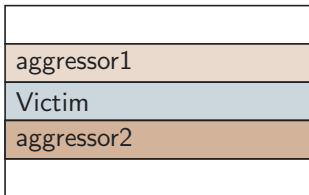
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



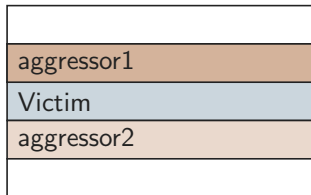
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



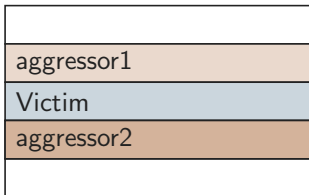
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```


- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



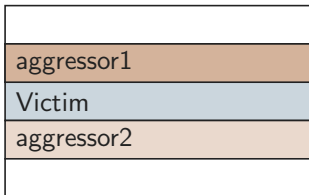
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



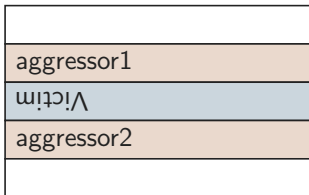
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



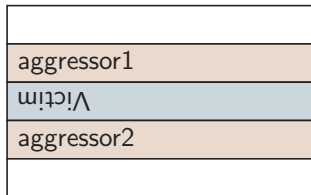
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows



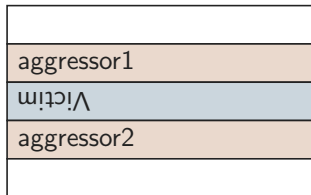
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows
- Worse with every new DRAM generation



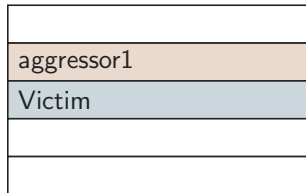
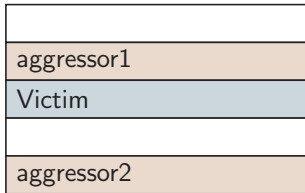
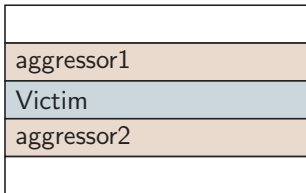
```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

- Hardware fault of the DRAM [3]
- Frequent accesses flip bits in neighboring rows
- Worse with every new DRAM generation
- Enables attacks, many countermeasures proposed



```
1 for (i = 0; i < N; ++i) {  
2   *aggressor1;  
3   *aggressor2;  
4   flush(aggressor1);  
5   flush(aggressor2);  
6 }
```

Rowhammer Patterns



Mitigations focus on the **characteristics** of Rowhammer

Characteristics 1: Flips are Infrequent



- **ECC** Error Correcting Codes



- **ECC Error Correcting Codes**
 - Additional hardware redundancy



- **ECC** Error Correcting Codes
 - Additional hardware redundancy
 - Catch **low** number of flips in HW



- **ECC** Error Correcting Codes
 - Additional hardware redundancy
 - Catch **low** number of flips in HW
- Increase the **refresh rate**

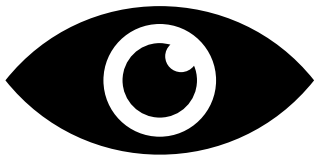


- **ECC** Error Correcting Codes
 - Additional hardware redundancy
 - Catch **low** number of flips in HW
- Increase the **refresh rate**
 - Reduce the attackers time window

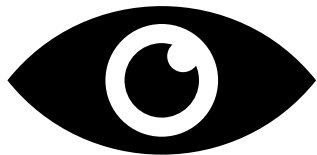


- **ECC** Error Correcting Codes
 - Additional hardware redundancy
 - Catch **low** number of flips in HW
- Increase the **refresh rate**
 - Reduce the attackers time window
 - Default *64ms*

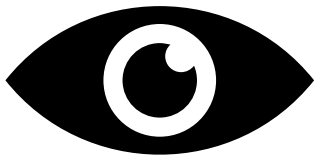
Characteristics 2: Attacks are Detectable



- Use Performance Monitoring Counters (PMCs)



- Use Performance Monitoring Counters (PMCs)
 - Monitor DRAM activity

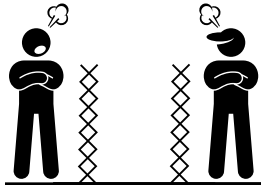


- Use Performance Monitoring Counters (PMCs)
 - Monitor DRAM activity
- Stop **suspicious** applications

Characteristics 3: Flips happen at Distance 1

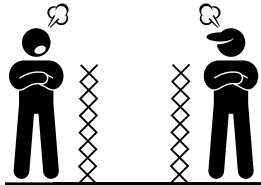


- Flips only happen in **neighboring** rows
- Mitigations based row layout
 - Guard rows
 - Targeted Row Refresh



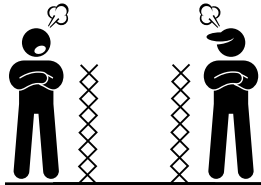
- Allocate **guard rows**

Agressor
Guard
Victim
Guard
Agressor



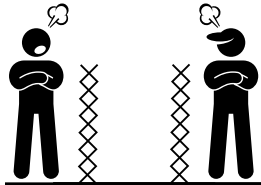
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



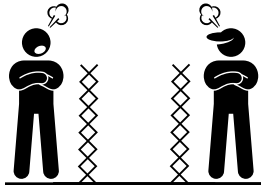
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



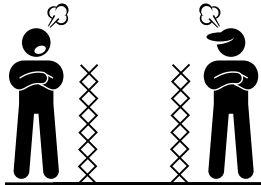
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



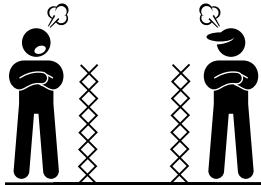
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



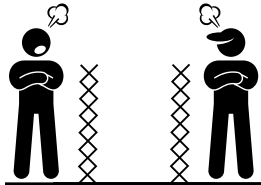
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



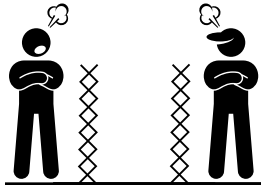
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



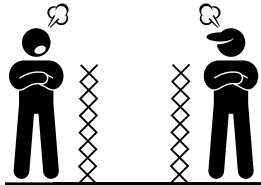
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



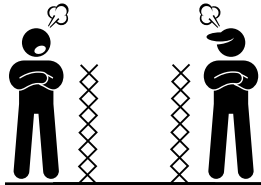
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



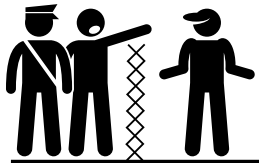
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



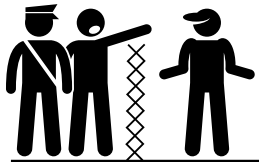
- Allocate **guard rows**
- Can not reach victim

Agressor
Guard
Victim
Guard
Agressor



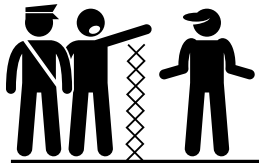
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



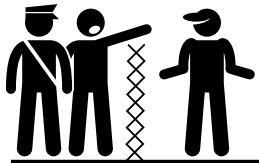
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



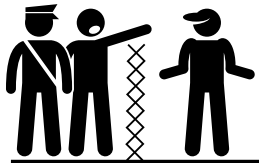
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



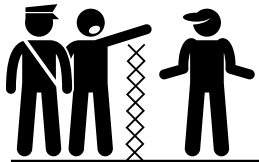
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



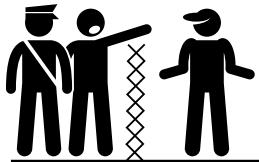
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



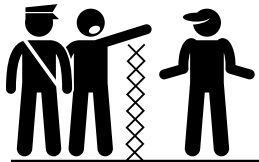
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



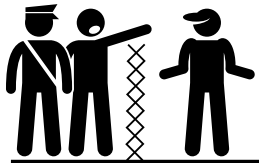
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



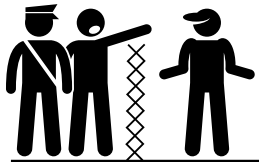
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



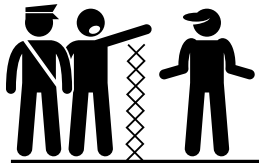
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



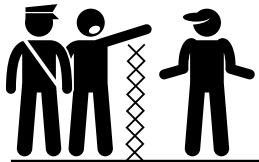
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



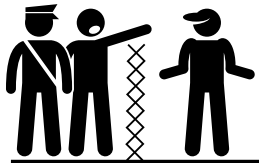
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



- Additional victim refreshes
- Standardized in LPDDR4x

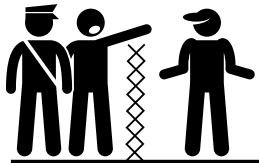
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)



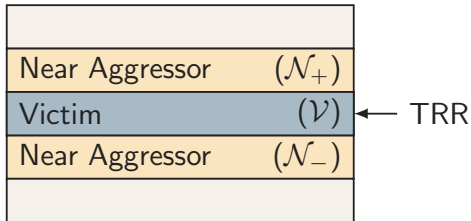
- Additional victim refreshes
- Standardized in LPDDR4x

Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)

Target Row Refresh (TRR)



- Additional victim refreshes
- Standardized in LPDDR4x



Let's start the cat and mouse game!



The Problem with Rowhammer Mitigations

- Bit Flips are infrequent - ECC, Increased Refresh Rate

- ~~Bit Flips are infrequent – ECC, Increased Refresh Rate [3]~~

- ~~Bit Flips are infrequent - ECC, Increased Refresh Rate [3]~~
- Detectable with Performance Counters - ANVIL, HexPADS

- ~~Bit Flips are infrequent - ECC, Increased Refresh Rate [3]~~
- ~~Detectable with Performance Counters - ANVIL, HexPADS [2]~~



- ~~Bit Flips are infrequent - ECC, Increased Refresh Rate [3]~~
- ~~Detectable with Performance Counters - ANVIL, HexPADS [2]~~
- Hammer Distance is 1 - TRR, ZebRAM, B-CATT

- ~~Bit Flips are infrequent - ECC, Increased Refresh Rate [3]~~
- ~~Detectable with Performance Counters - ANVIL, HexPADS [2]~~
- ~~Hammer Distance is 1 - TRR, ZebRAM, B-CATT [4]~~

Would perfect TRR fix Rowhammer?

The Half-Double Effect

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Distance-1**

- $(\mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
- *Classic* double-sided Rowhammer

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Distance-1**
 - $(\mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
 - *Classic* double-sided Rowhammer
- **First** flip after:
 - 18 000 hammers in 1.2 *ms*

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Distance-1**
 - $(\mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
 - *Classic* double-sided Rowhammer
- **First** flip after:
 - 18 000 hammers in 1.2 *ms*
 - ✓ **Within** the refresh window
 - ✗ **Mitigated** by TRR

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Distance-2**

- $(\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\infty}$
- Distance two double-sided Rowhammer

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Distance-2**
 - $(\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\infty}$
 - Distance two double-sided Rowhammer
- **First flip after:**
 - 4 000 000 hammers in 270 *ms*

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Distance-2**
 - $(\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\infty}$
 - Distance two double-sided Rowhammer
- **First** flip after:
 - 4 000 000 hammers in 270 *ms*
 - **X Not** within the refresh windows

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Half-Double**

- $((\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\beta} \rightarrow \mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
- **Many** distance-2 accesses with a **few** distance-1 accesses

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Half-Double**

- $((\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\beta} \rightarrow \mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
- **Many** distance-2 accesses with a **few** distance-1 accesses
- **First** flip after:
 - 296 960 hammers in 20 *ms*
 - Dilution $\beta = 57$ (5120 distance-1 accesses)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Half-Double**

- $((\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\beta} \rightarrow \mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
- **Many** distance-2 accesses with a **few** distance-1 accesses
- **First** flip after:
 - 296 960 hammers in 20 *ms*
 - Dilution $\beta = 57$ (5120 distance-1 accesses)
 - ✓ **Within** the refresh window

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Half-Double**

- $((\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\beta} \rightarrow \mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
- **Many** distance-2 accesses with a **few** distance-1 accesses
- **First** flip after:
 - 296 960 hammers in 20 *ms*
 - Dilution $\beta = 57$ (5120 distance-1 accesses)
 - ✓ **Within** the refresh window
 - ✓ **Assisted** by TRR

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

- **Half-Double**

- $((\mathcal{F}_+ \rightarrow \mathcal{F}_-)^{\beta} \rightarrow \mathcal{N}_+ \rightarrow \mathcal{N}_-)^{\infty}$
- **Many** distance-2 accesses with a **few** distance-1 accesses
- **First** flip after:
 - 296 960 hammers in 20 *ms*
 - Dilution $\beta = 57$ (5120 distance-1 accesses)
 - ✓ **Within** the refresh window
 - ✓ **Assisted** by TRR

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)



Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)



Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

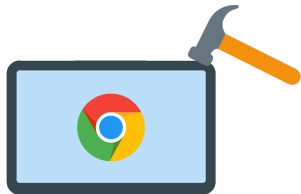
Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Far Aggressor	(\mathcal{F}_+)	
Near Aggressor	(\mathcal{N}_+)	← TRR
Victim	(\mathcal{V})	
Near Aggressor	(\mathcal{N}_-)	
Far Aggressor	(\mathcal{F}_-)	

Far Aggressor	(\mathcal{F}_+)	
Near Aggressor	(\mathcal{N}_+)	
Victim	(\mathcal{V})	
Near Aggressor	(\mathcal{N}_-)	← TRR
Far Aggressor	(\mathcal{F}_-)	

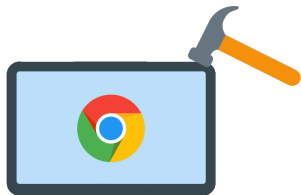
Far Aggressor	(\mathcal{F}_+)
Near Aggressor	(\mathcal{N}_+)
Victim	(\mathcal{V})
Near Aggressor	(\mathcal{N}_-)
Far Aggressor	(\mathcal{F}_-)

Exploitable in the Wild?



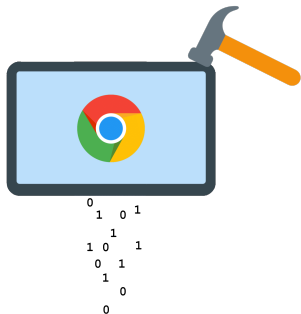
```
0
1 0 1
  1
1 0 1
  0 1
    1
    0
    0
```

- Target PFN in Page Table Entry [6]

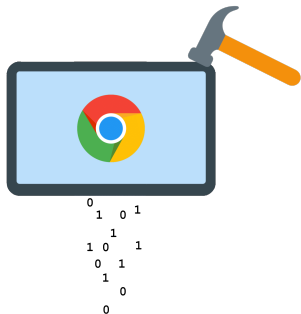


```
0
1 0 1
  1
1 0 1
  0 1
    1
      0
        0
```

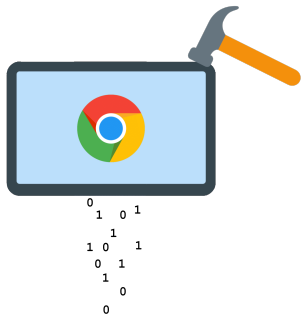
- Target PFN in Page Table Entry [6]
- **C1**: Allocation of Contiguous Memory



- Target PFN in Page Table Entry [6]
- **C1**: Allocation of Contiguous Memory
- **C2**: Alternative to Memory Templating



- Target PFN in Page Table Entry [6]
- **C1**: Allocation of Contiguous Memory
- **C2**: Alternative to Memory Templating
- **C3**: Memory Massaging



- Target PFN in Page Table Entry [6]
- **C1**: Allocation of Contiguous Memory
- **C2**: Alternative to Memory Templating
- **C3**: Memory Massaging
- **C4**: Bit-Flip Verification

- Corrupt page table entries can **kill** the attacker process

- Corrupt page table entries can **kill** the attacker process

```
if ( /*misprediction*/ ) {  
    access(probe + (*ptr & 1));  
}  
if ( is_cached(probe) ) {  
    // ptr[0-4] valid  
}
```

- **Verify** if address save to access
- **Spectre** gadget

- Corrupt page table entries can **kill** the attacker process

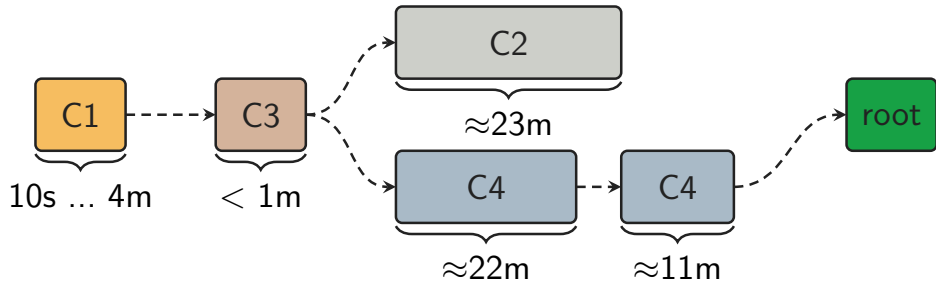
```
if ( /*misprediction*/ ) {  
    access(probe + (*ptr & 1));  
}  
if ( is_cached(probe) ) {  
    // ptr[0-4] valid  
}
```

- **Verify** if address save to access
- **Spectre** gadget
- Cached → accessible

- Corrupt page table entries can **kill** the attacker process

```
if ( /*misprediction*/ ) {  
    access(probe + (*ptr & 1));  
}  
if ( is_cached(probe) ) {  
    // ptr[0-4] valid  
}
```

- **Verify** if address save to access
- **Spectre** gadget
- Cached → accessible
- **Suppresses** corruption faults



There
are a lot of
Rowhammer
exploits



We have
TRR



TRR
amplifies
Rowhammer



topfoto.com

~~ECC~~

~~ChipKill~~

~~TRR~~

~~PARA~~

~~PRA~~

~~Refresh Rate~~

~~MASCAT~~

~~HexPADS~~

~~ARMOR~~

~~NO OOM~~

~~B-CATT~~

~~G-CATT~~

~~ZebRAM~~

~~MemGuard~~

Rethinking Rowhammer Mitigations



General approach to data integrity
protection



- **General** approach to data integrity protection



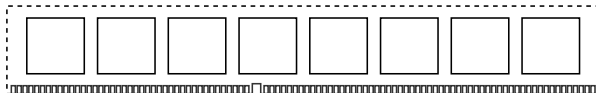
- **General** approach to data integrity protection
- Detect **all** data integrity failures with a MAC

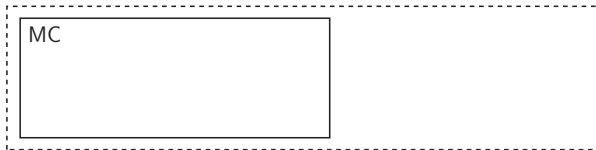
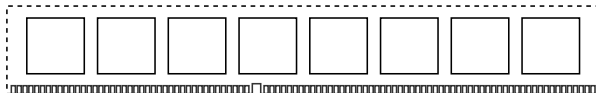


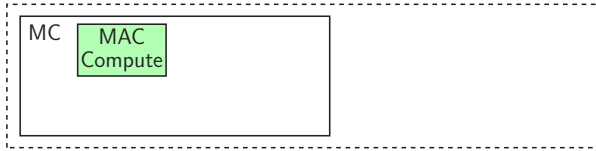
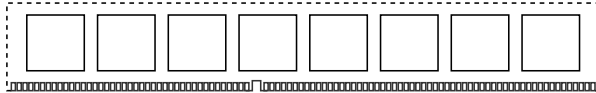
- **General** approach to data integrity protection
- Detect **all** data integrity failures with a MAC
- Best effort correction

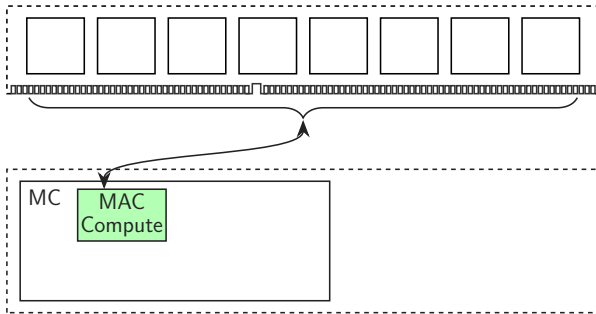


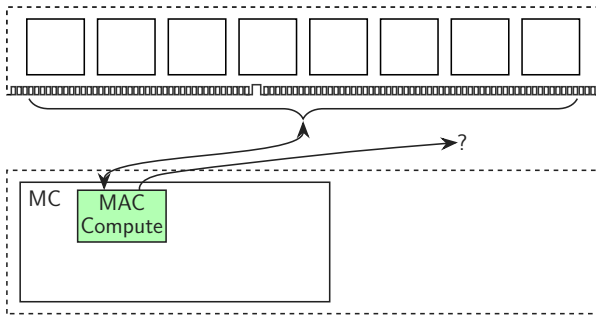
- **General** approach to data integrity protection
- Detect **all** data integrity failures with a MAC
- Best effort correction
- **All** Rowhammer attacks are DoS in the **worst case**

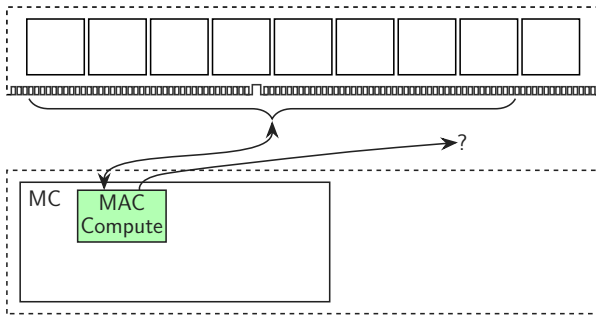


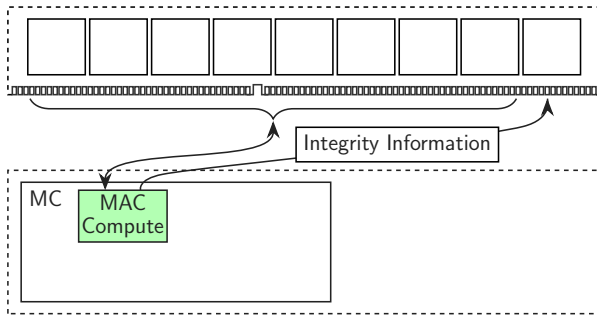


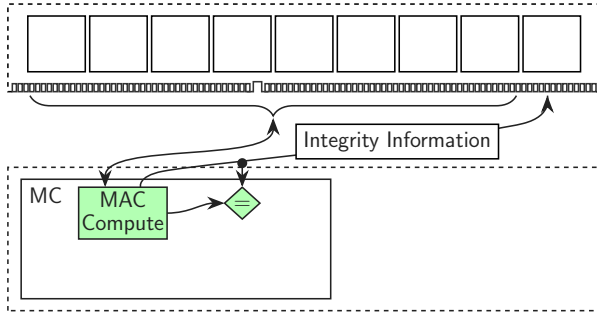


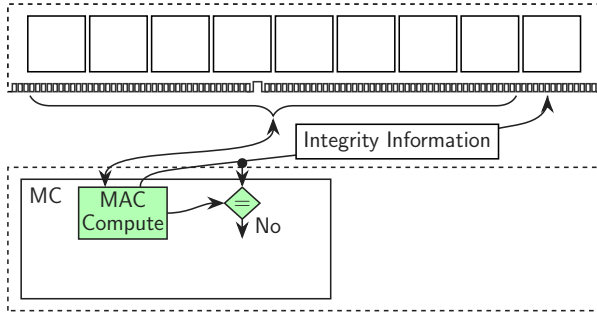


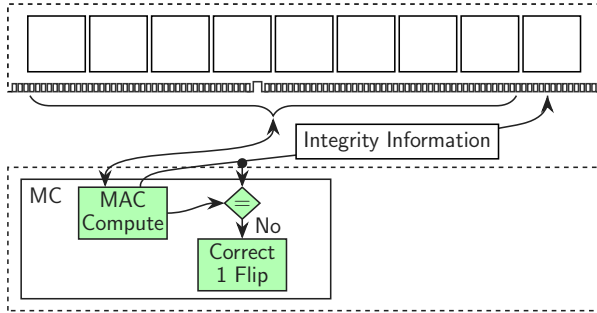


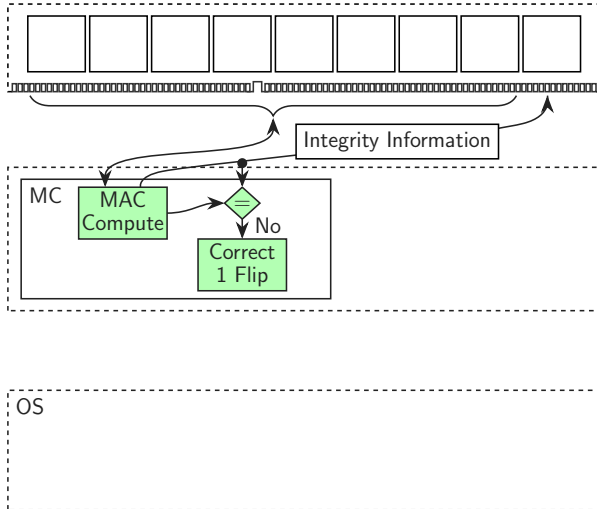


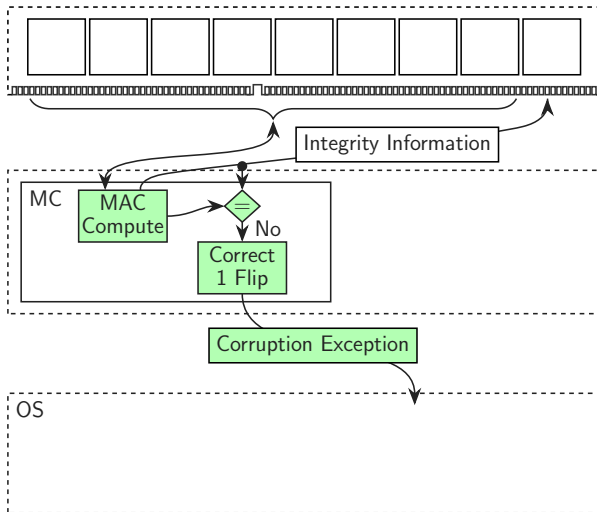


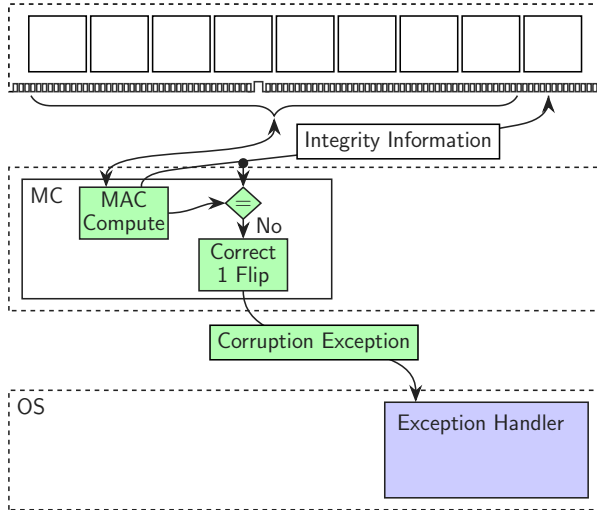


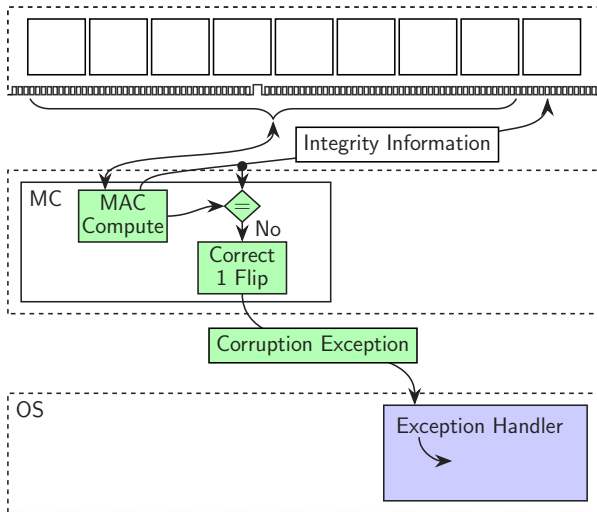


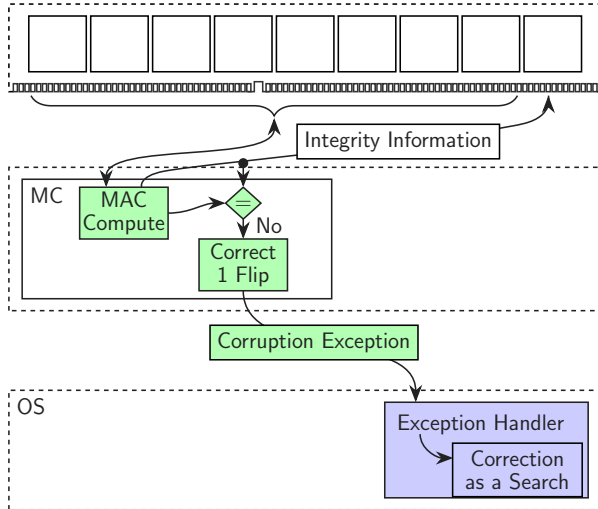


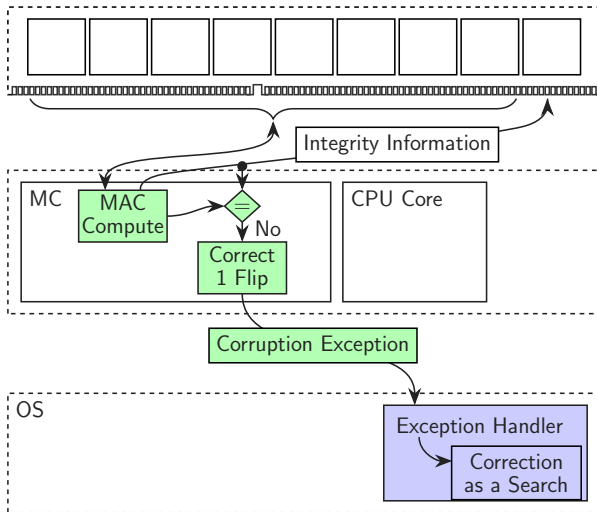


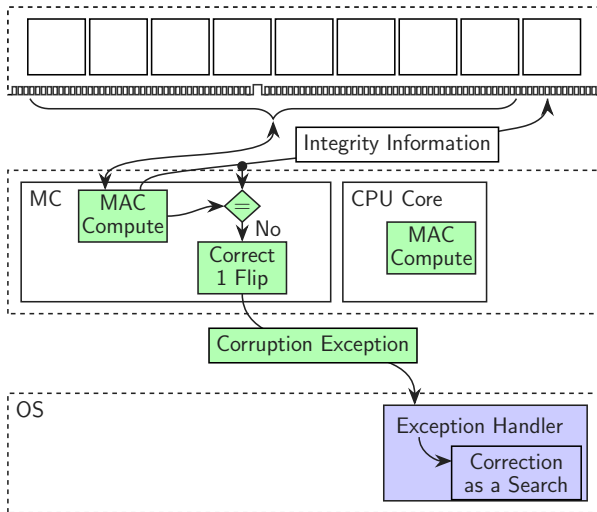


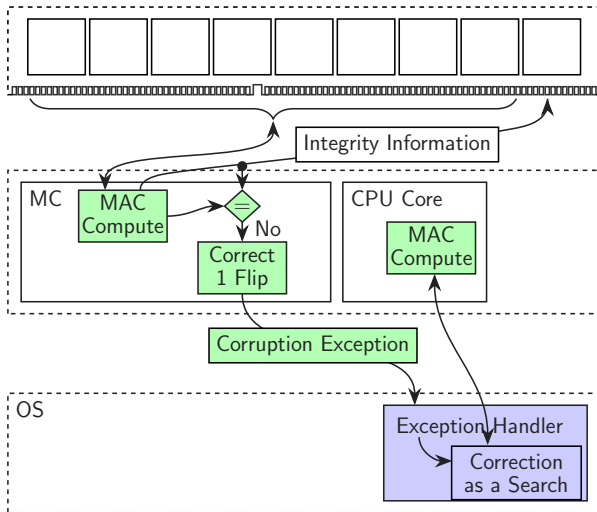


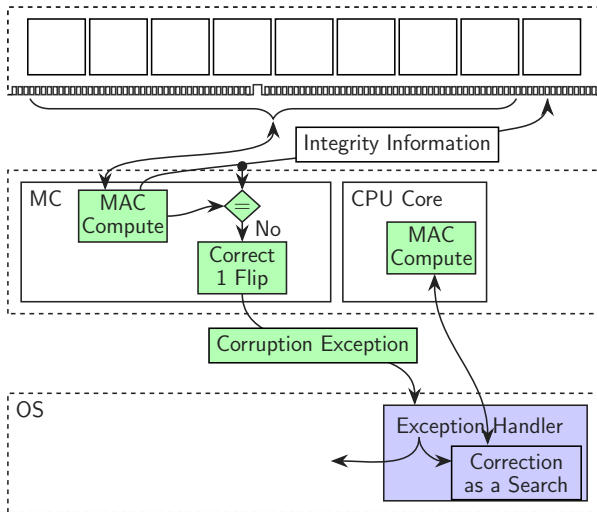


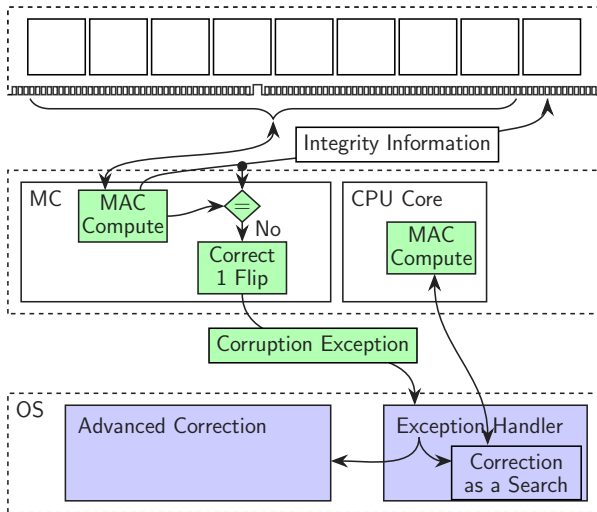


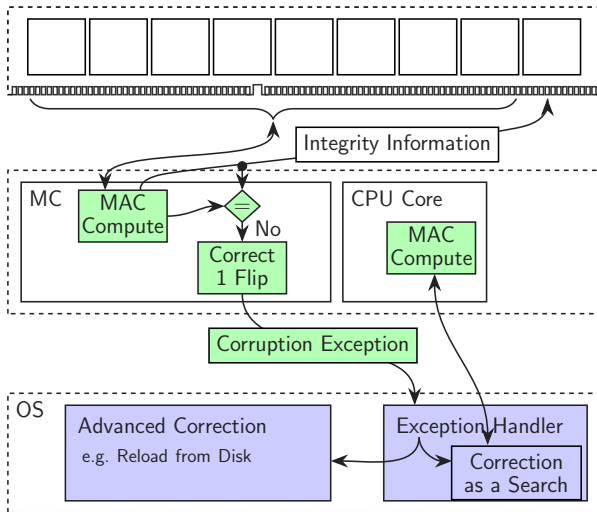


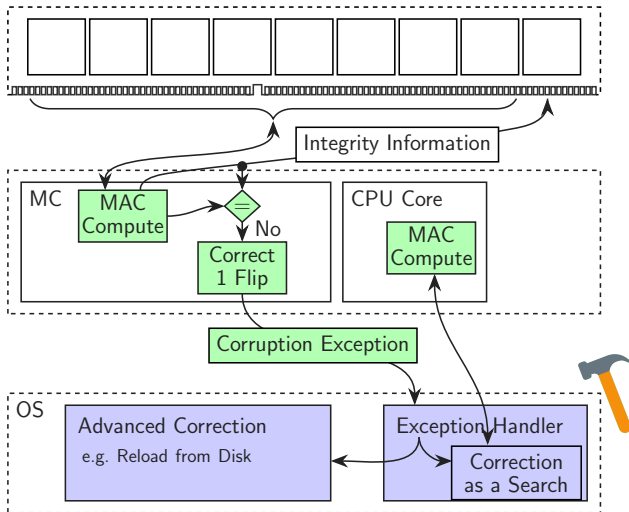


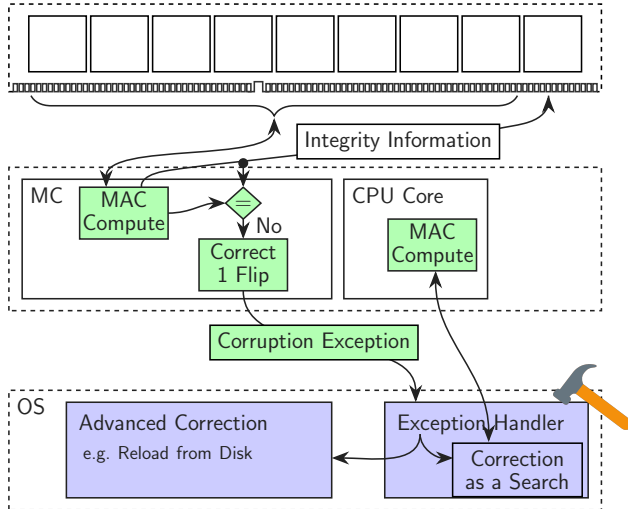


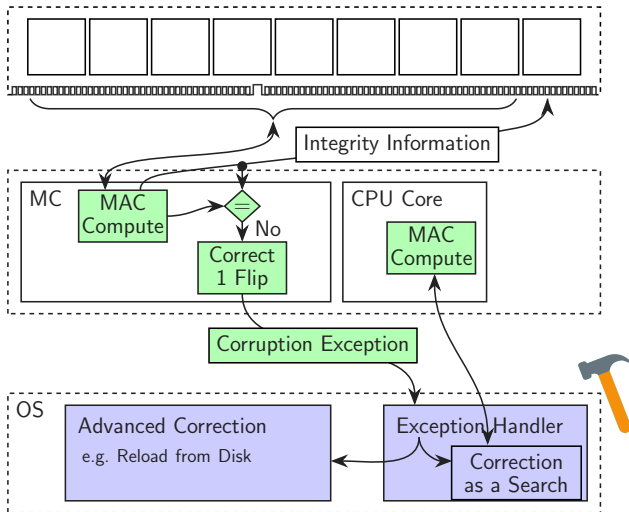


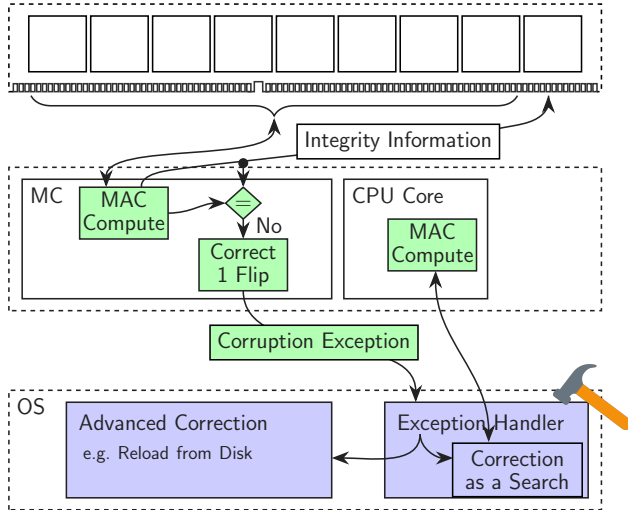


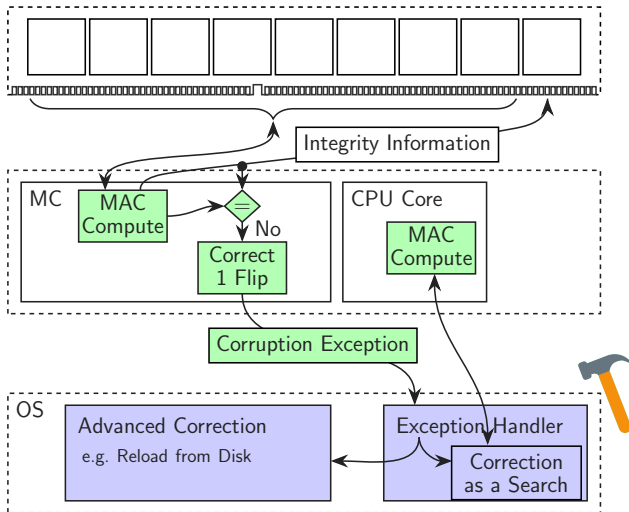


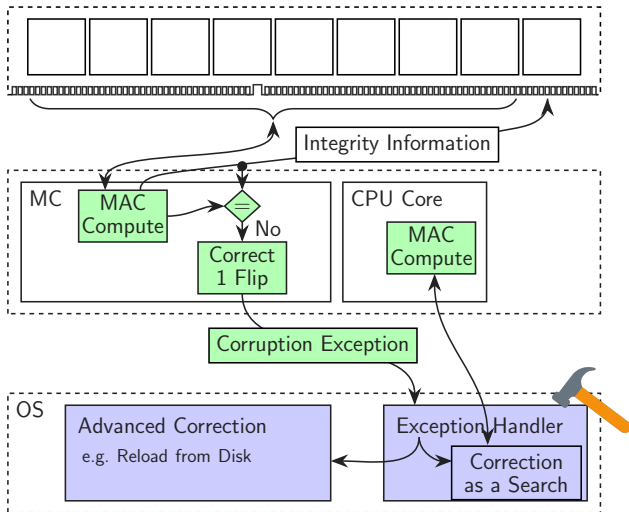


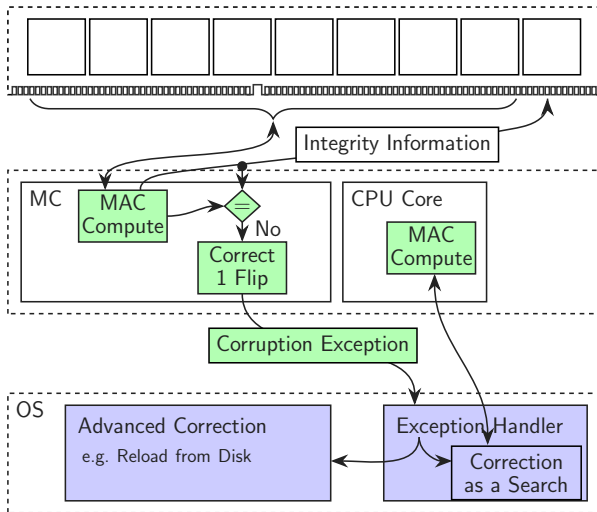


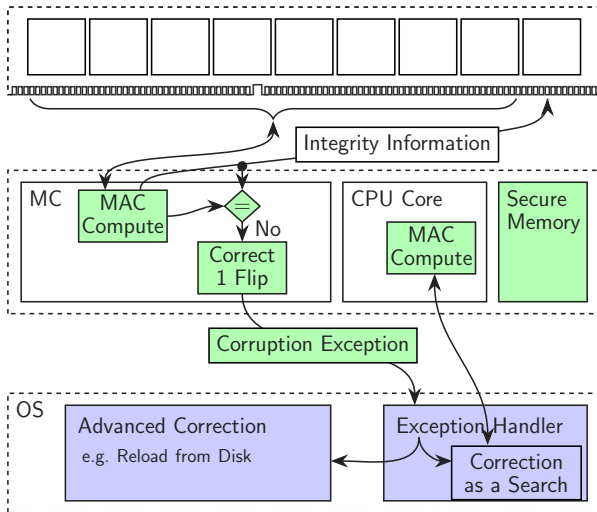


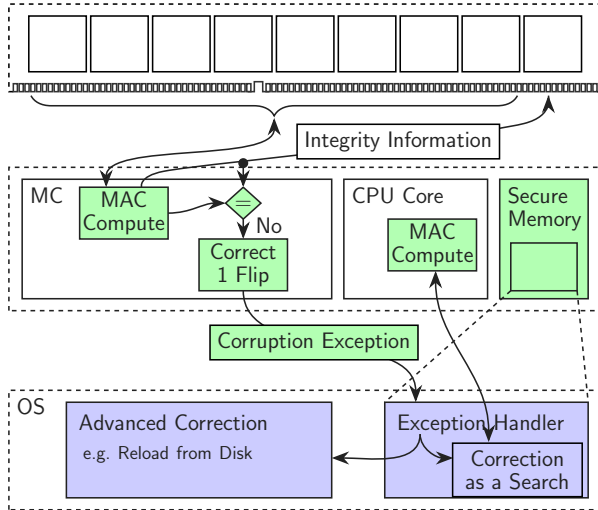


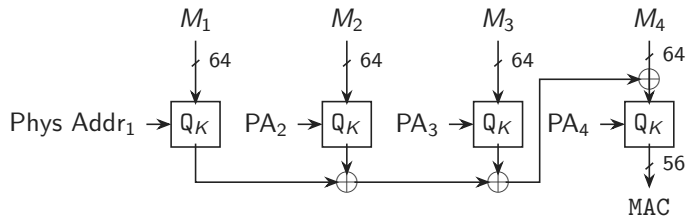


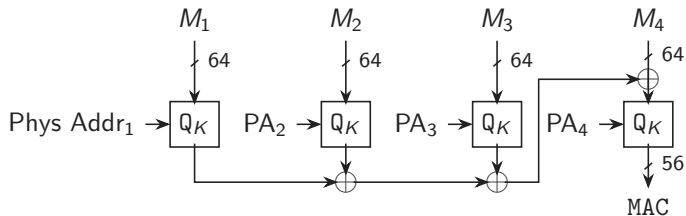




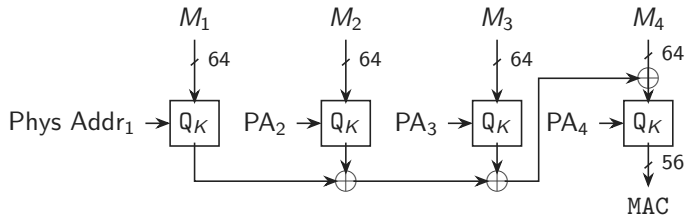




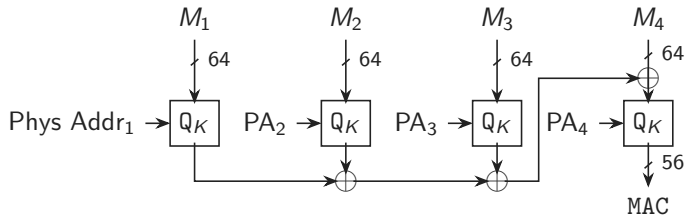




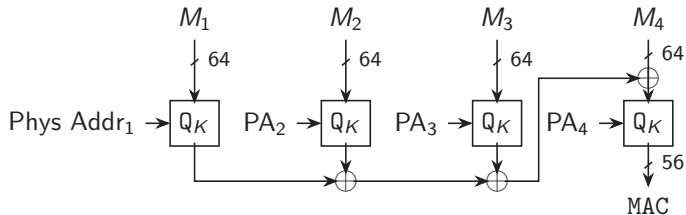
- Pipelined PMAC construction



- Pipelined PMAC construction
- QARMA_{5-64- σ_0} block cipher [1]



- Pipelined PMAC construction
- QARMA_{5-64- σ_0} block cipher [1]
- 256-bit data 5.13 ns



- Pipelined PMAC construction
- QARMA_{5-64- σ_0} block cipher [1]
- 256-bit data 5.13 ns
- 512-bit data 6.60 ns

Error Correction in Software





- Physical address in CR2



- **Physical address** in CR2
- **Load** corrupted data with `csi_load`



- **Physical address** in CR2
- **Load** corrupted data with `csi_load`
- **Compute** MAC with `csi_mac`



- **Physical address** in CR2
- **Load** corrupted data with `csi_load`
- **Compute** MAC with `csi_mac`
- **Write** back corrected data with `csi_xchg`



MACs **cannot** correct bit flips



- MACs **cannot** correct bit flips



- MACs **cannot** correct bit flips
- Brute force search with approximate equality



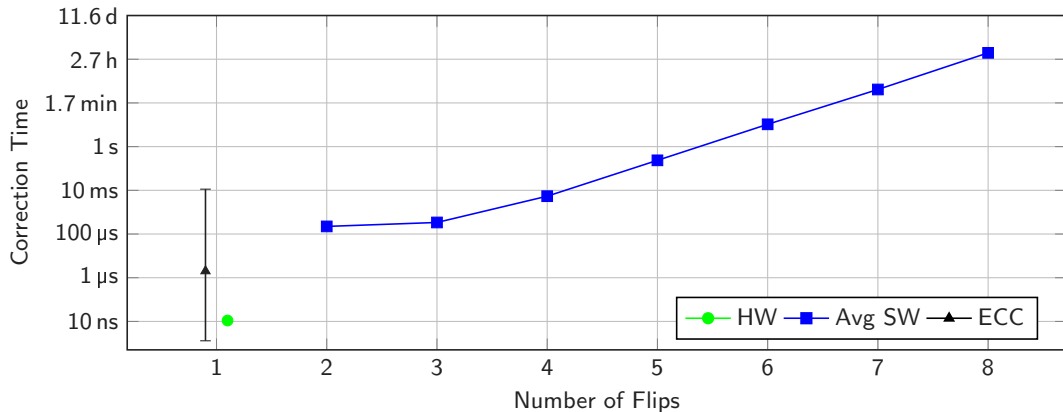
- MACs **cannot** correct bit flips
- Brute force search with approximate equality
0010110100101101 → 01011010

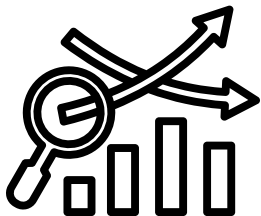


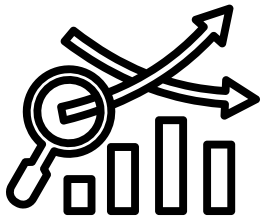
- MACs **cannot** correct bit flips
- Brute force search with approximate equality
0010110100101101 → 01011010
0010110100101101 → 0101**0**010 ✓



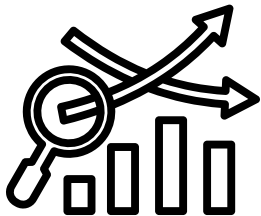
- MACs **cannot** correct bit flips
- Brute force search with approximate equality
0010110100101101 → 01011010
0010110100101101 → 0101**0**010 ✓
- Parity bits to **shrink** search space



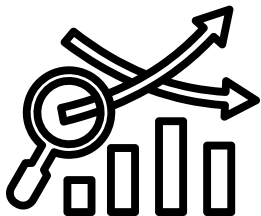




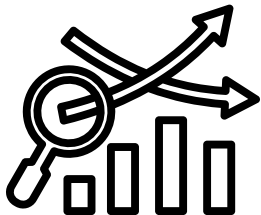
- Implemented CSI:Rowhammer in **gem5**



- Implemented CSI:Rowhammer in **gem5**
- Modified **Linux** kernel

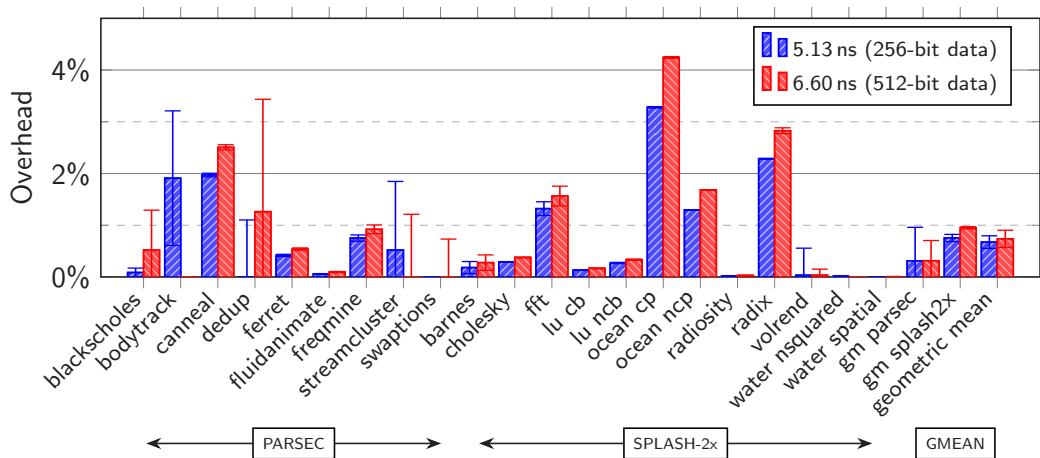


- Implemented CSI:Rowhammer in **gem5**
- Modified **Linux** kernel
- Evaluated **correct functionality**



- Implemented CSI:Rowhammer in **gem5**
- Modified **Linux** kernel
- Evaluated **correct functionality**
- Evaluated **performance overhead**

CSI:Rowhammer – Performance Overhead







- **Approximate Equality**

- **Approximate Equality**

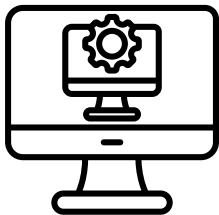
Data Flips	$\log_2(p)$	d	MAC Strength
5	26.0	3	41.2
6	31.5	2	45.4
7	38.8	1	50.2
8	42.4	0	56.0

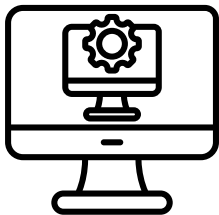
- **Approximate Equality**
- Silent Data Corruption rate less than once per 10^9 billion years.

Data Flips	$\log_2(p)$	d	MAC Strength
5	26.0	3	41.2
6	31.5	2	45.4
7	38.8	1	50.2
8	42.4	0	56.0

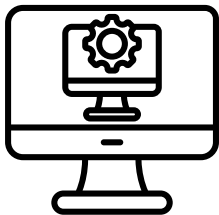
- **Approximate Equality**
- Silent Data Corruption rate less than once per 10^9 billion years.
- Rowhammer second preimage after one year: $9.75 \cdot 10^{-5}\%$

Data Flips	$\log_2(p)$	d	MAC Strength
5	26.0	3	41.2
6	31.5	2	45.4
7	38.8	1	50.2
8	42.4	0	56.0

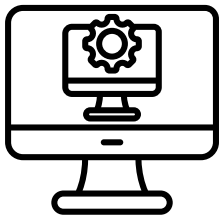




- Only **host support** required

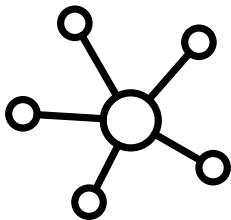


- Only **host support** required
- Guest support optional

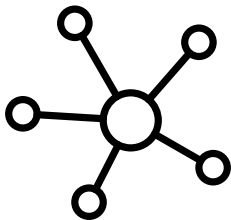


- Only **host support** required
- Guest support optional
- Makes upgrading easy

The Possibilities

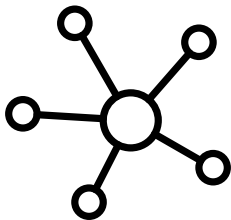


Correction in **software** opens up many new **possibilities**



Correction in **software** opens up many new **possibilities**

- Data correction as a service (DCaaS)



Correction in **software** opens up many new **possibilities**

- Data correction as a service (DCaaS)
- Correction by ChatGPT™

CSI:Rowhammer – Correction by ChatGPT

MA

Find the error in "Testmessoge" so that it has the following SHA512 hash
"9167bf29600fea2cb1ba4e21e2f311bf5c2641f0cd5cb5bec948177264ff
78e2bc346216427b75e75fe67e842ef3085ac303b8111a37adeaf977b2b
269ada636"



CSI:Rowhammer – Correction by ChatGPT



It is not possible for the string "Testmessoge" to have the SHA512 hash "9167bf29600fea2cb1ba4e21e2f311bf5c2641f0cd5cb5bec948177264ff78e2bc346216427b75e75fe67e842ef3085ac303b8111a37adeaf977b2b269ada636". This is because the SHA512 hash function is a one-way function, meaning that it is not possible to determine the input that produced a given hash. Therefore, it is not possible to determine the error in the input string "Testmessoge" based on the given hash value.



CSI:Rowhammer – Correction by ChatGPT

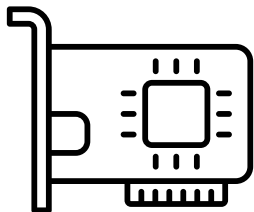
MA



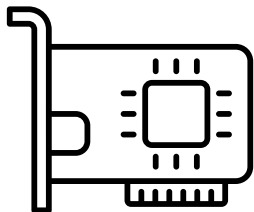
Can just brute force the correct answer for me?



This content may violate our [content policy](#). If you believe this to be in error, please [submit your feedback](#) — your input will aid our research in this area.

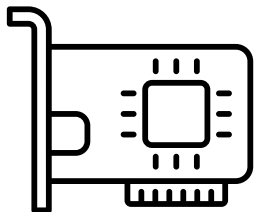


Hardware acceleration



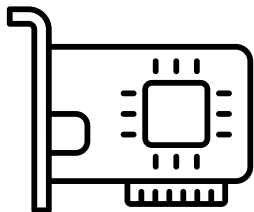
Hardware acceleration

- Ice Lake SP Xeon Server die size: 628 mm²



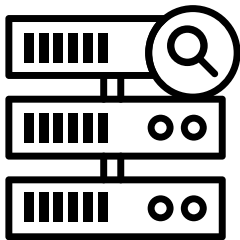
Hardware acceleration

- Ice Lake SP Xeon Server die size: 628 mm²
- Fits roughly 500 000 QARMA blocks

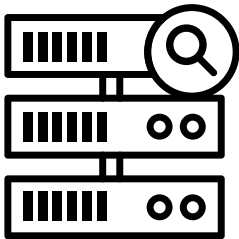


Hardware acceleration

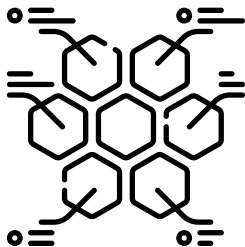
- Ice Lake SP Xeon Server die size: 628 mm²
- Fits roughly 500 000 QARMA blocks
- 8 bit flips correction: 5.44 h → 40 ms



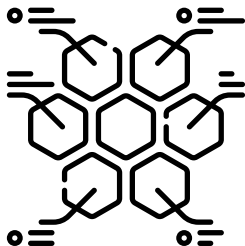
- Multi-bit flips not happening often



- Multi-bit flips not happening often
- **Specialized** data correction nodes

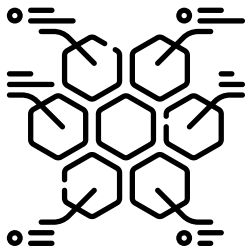


Kernel data with **known** structures



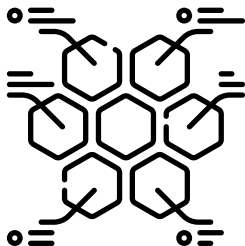
Kernel data with **known** structures

- Kernel code



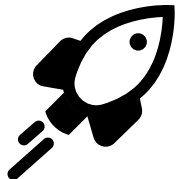
Kernel data with **known** structures

- Kernel code
- Page tables

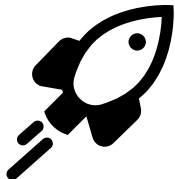


Kernel data with **known** structures

- Kernel code
- Page tables
- Task structs

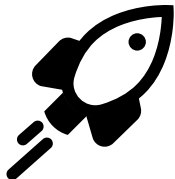


Performance impact caused by MAC computation



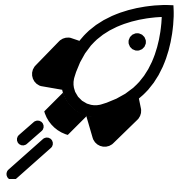
Performance impact caused by MAC computation

- Transient execution with unverified data



Performance impact caused by MAC computation

- Transient execution with unverified data
- Can be made **secure** [7, 5]




Performance impact caused by MAC computation

- Transient execution with unverified data
 - Can be made **secure** [7, 5]
- **No** performance impact


- Rowhammer mitigations assuming characteristics are broken.


- Rowhammer mitigations assuming characteristics are broken.
- Software-based correction opens up a lot of flexibility.

- Rowhammer mitigations assuming characteristics are broken.
- Software-based correction opens up a lot of flexibility.
- **Open Source** 

`https://github.com/IAIK/halfdouble`

`https://github.com/IAIK/csirowhammer`

- Rowhammer mitigations assuming characteristics are broken.
- Software-based correction opens up a lot of flexibility.
- **Open Source** 
<https://github.com/IAIK/halfdouble>
<https://github.com/IAIK/csirowhammer>
- **Many more details** in the papers

- Rowhammer mitigations assuming characteristics are broken.
 - Software-based correction opens up a lot of flexibility.
 - **Open Source** 
- <https://github.com/IAIK/halfdouble>
- <https://github.com/IAIK/csirowhammer>
- Many more **details** in the papers



Half-Double Rowhammer



CSI:Rowhammer

- [1] Roberto Avanzi. The QARMA Block Cipher Family: Almost MDS Matrices Over Rings With Zero Divisors, Nearly Symmetric Even-Mansour Constructions With Non-Involutory Central Rounds, and Search Heuristics for Low-Latency S-Boxes. In: *IACR Transactions on Symmetric Cryptology 2017.1* (2017), pp. 4–44.
- [2] Daniel Gruss, Moritz Lipp, Michael Schwarz, Daniel Genkin, Jonas Juffinger, Sioli O’Connell, Wolfgang Schoechl, and Yuval Yarom. Another Flip in the Wall of Rowhammer Defenses. In: *S&P*. 2018.
- [3] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, D607onghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors. In: *ACM SIGARCH Computer Architecture News 42.3* (2014), pp. 361–372.

- [4] **Andreas Kogler**, Jonas Juffinger, Salman Qazi, Yoongu Kim, Moritz Lipp, Nicolas Boichat, Eric Shiu, Mattias Nissler, and Daniel Gruss. Half-Double: Hammering From the Next Row Over. In: In major revision at USENIX Security Symposium. 2022.
- [5] Tamara Silbergleit Lehman, Andrew D. Hilton, and Benjamin C. Lee. PoisonIvy: Safe speculation for secure memory. In: MICRO. 2016.
- [6] Mark Seaborn and Thomas Dullien. Test DRAM for bit flips caused by the rowhammer problem. 2015. URL: <https://github.com/google/rowhammer-test>.
- [7] Weidong Shi and Hsien-Hsin S. Lee. Authentication Control Point and Its Implications For Secure Processor Design. In: MICRO. 2006.

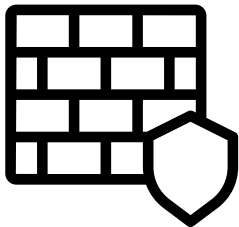
Additional Slides



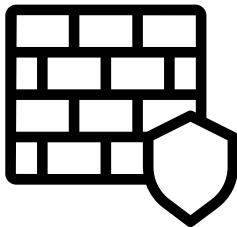
- Tested **13** DIMMs & devices
- **2** DIMMs affected
 - FPGA analysis
 - Exact numbers
- **5** out of **7** mobile devices affected
 - **Reversed** addressing
 - Unprivileged **flush**
 - Uncachable memory (10x)

Affected Devices - Flip Numbers

System	RAM	$N_{Hammers}$	$UC_{0 \rightarrow 1}$	$UC_{1 \rightarrow 0}$	$Flush_{0 \rightarrow 1}$	$Flush_{1 \rightarrow 0}$
Chromebook ₁	LPDDR4x	23 274	27	40	2	5
Chromebook ₂	LPDDR4x	23 586	235	2379	12	101
OnePlus 5T	LPDDR4x	25 687	2	30	1	24
Pixel 3	LPDDR4x	32 921	11	5	0	0
HTC U11	LPDDR4x	21 840	-	-	3	17

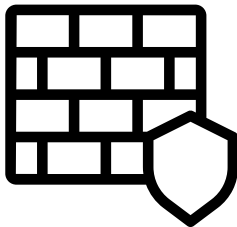


Protect exception handler against bit flips



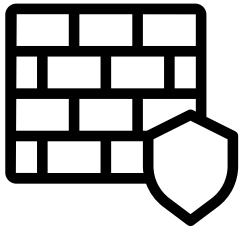
Protect exception handler against bit flips

- 1 Page IDT



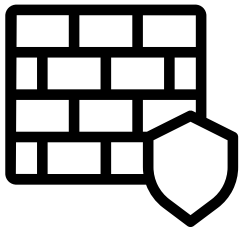
Protect exception handler against bit flips

- 1 Page IDT
- 1 Page GDT



Protect exception handler against bit flips

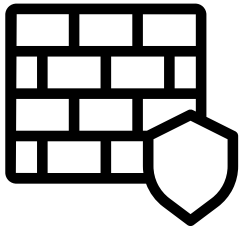
- 1 Page IDT
- 1 Page GDT
- 2 Pages exception handler + correction



Protect exception handler against bit flips

- 1 Page IDT
- 1 Page GDT
- 2 Pages exception handler + correction

→ Page table entries unprotected



Protect exception handler against bit flips

- 1 Page IDT
- 1 Page GDT
- 2 Pages exception handler + correction

→ Page table entries unprotected

- 4 lockable TLB entries

Two conditions

1. Bit must always be in **register**
 2. Bit must be **unique** for every process
- **CR3** register

```
1 corruption_exception_handler() {
2     if (has_nested_bit_set(CR3)) {
3         enable_interrupts();
4         error_correction_as_a_search(corruption_address);
5     } else {
6         set_nested_bit(CR3);
7         enable_interrupts();
8         advanced_error_correction(corruption_address);
9         clear_nested_bit(CR3);
10    }
11 }
```



1. Out of Order CPU core
2. MAC computation and check
3. New instructions
4. Modified Linux performing the correction as a search