

Faulting Hardware from Software

Daniel Gruss

2020-09-13

Graz University of Technology



CTV

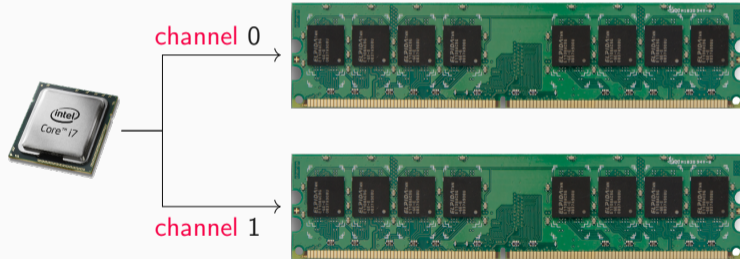
Test - Mozilla Firefox (on lab02)

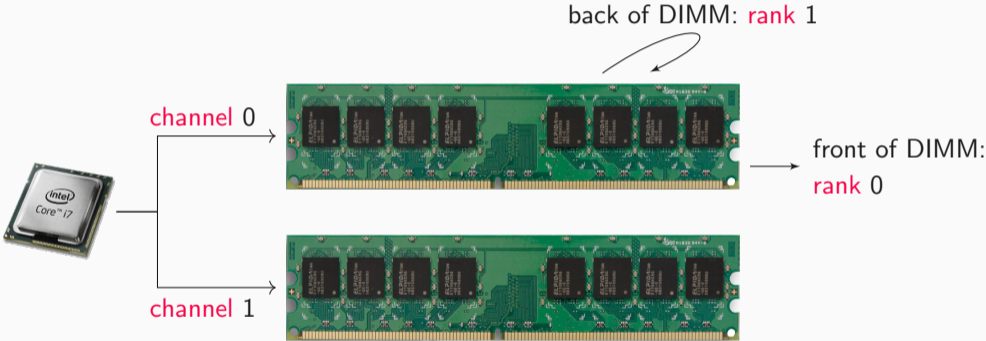
Test

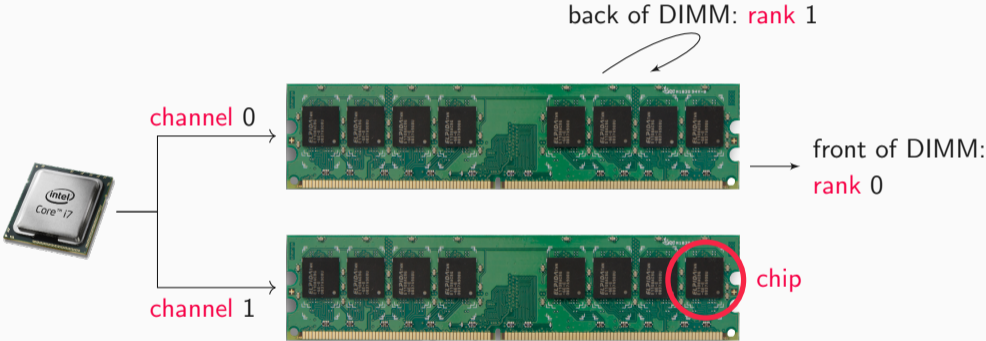
file:///home/dgruss/rowhammerjs/rowhammer.html

```
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
```

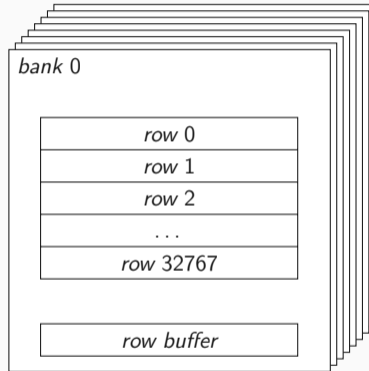


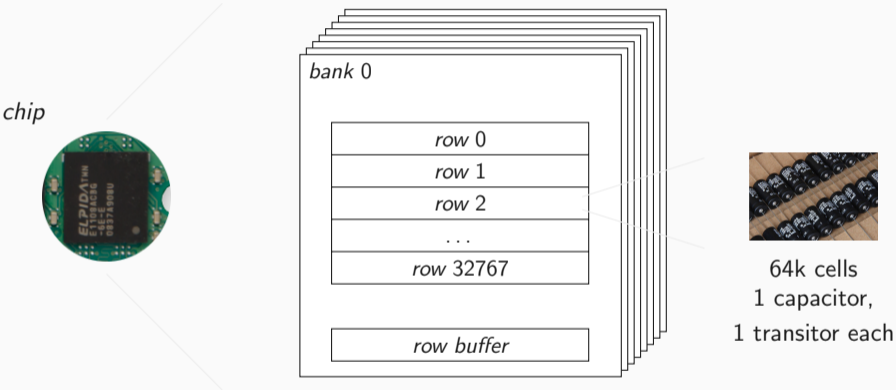


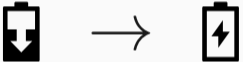
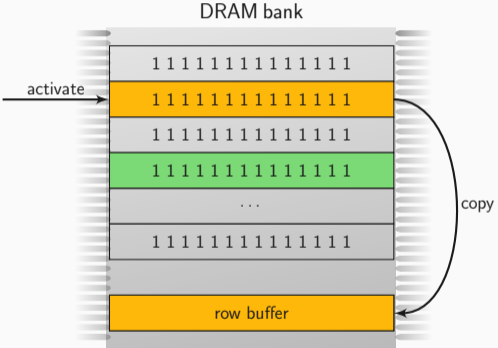


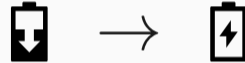
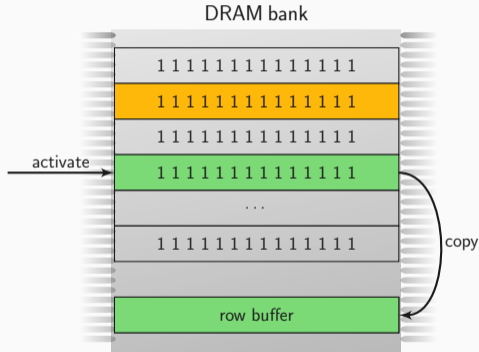


chip

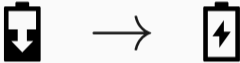
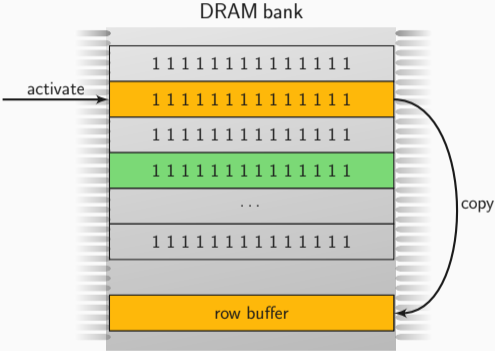




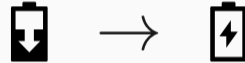
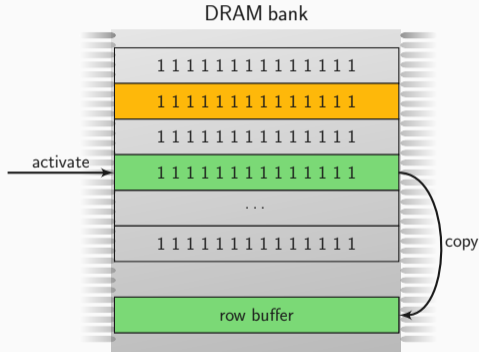




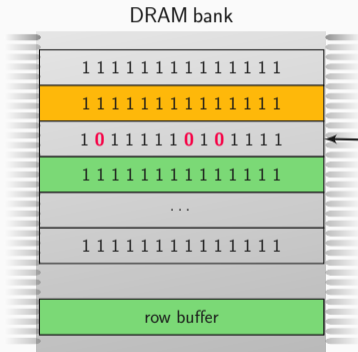
Cells leak faster upon proximate accesses → Rowhammer



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





- 85% affected [Kim+14] (see Figure)





- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]





- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



- First believed to be safe



- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



- First believed to be safe
- We showed bit flips [Pes+16]



- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



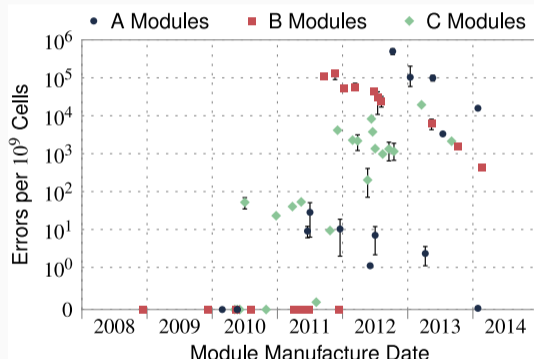
- First believed to be safe
- We showed bit flips [Pes+16]
- 67% affected [Lan16]



- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



- First believed to be safe
- We showed bit flips [Pes+16]
- 67% affected [Lan16]





BIT FLIPS

BIT FLIPS EVERYWHERE



Memory accesses must be

- **uncached**: reach DRAM
- **fast**: race against the next row refresh
- **targeted**: reach specific row

How do we get enough uncached accesses?







- `clflush` instruction → original paper [Kim+14]



- `clflush` instruction → original paper [Kim+14]
- cache eviction [GMM16; Awe+16]

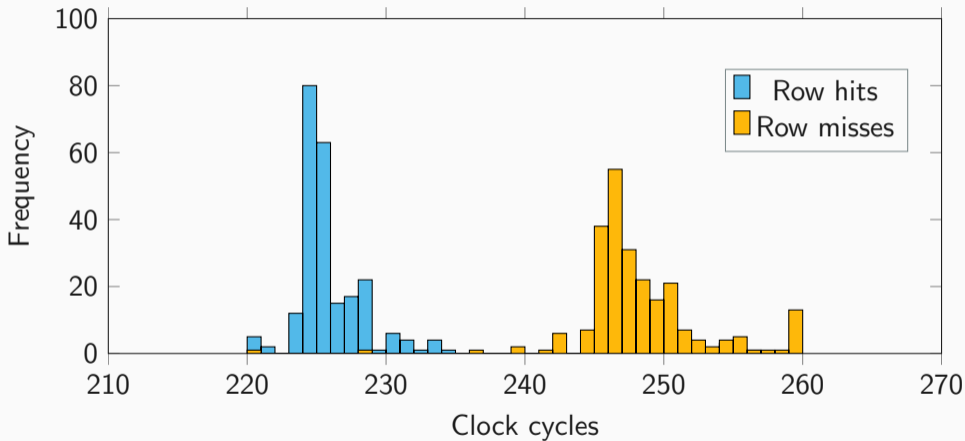


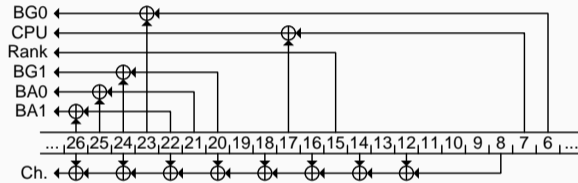
- `clflush` instruction → original paper [Kim+14]
- cache eviction [GMM16; Awe+16]
- non-temporal accesses [QS16]



- `clflush` instruction → original paper [Kim+14]
- cache eviction [GMM16; Awe+16]
- non-temporal accesses [QS16]
- uncached memory [Vee+16]

How do we target accesses?





 <https://github.com/IAIK/drama>







- 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



- 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
- Alternatively: Build a PUF [Ana+18]

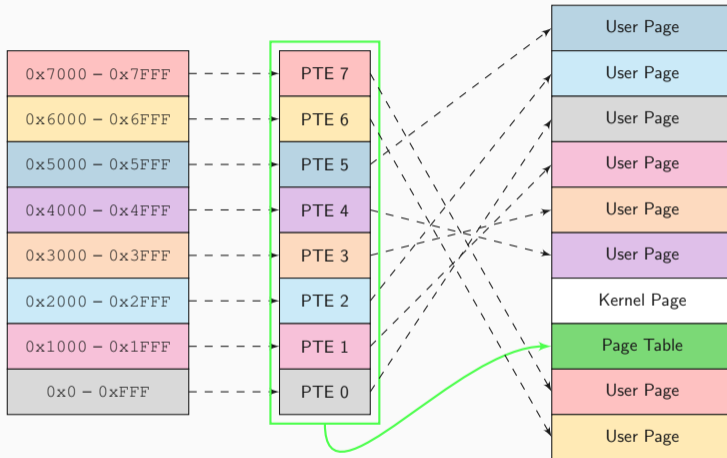
P	RW	US	WT	UC	R	D	S	G		
										X

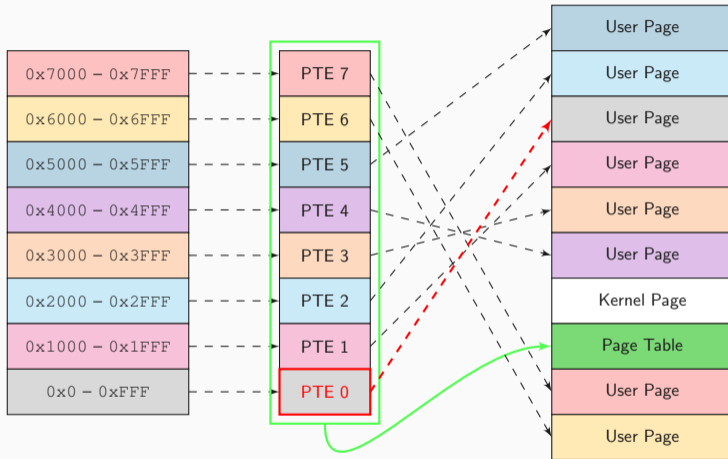
P	RW	US	WT	UC	R	D	S	G	Ignored	
Ignored										
Ignored										
Ignored				Ignored						X

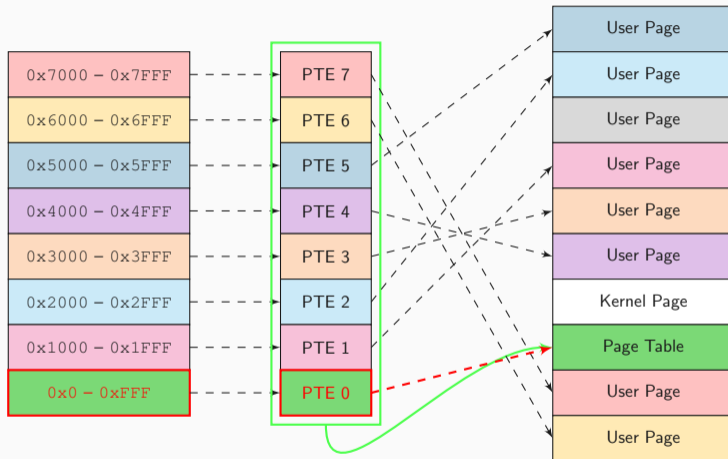
P	RW	US	WT	UC	R	D	S	G	Ignored	
Physical Page Number										
									Ignored	X

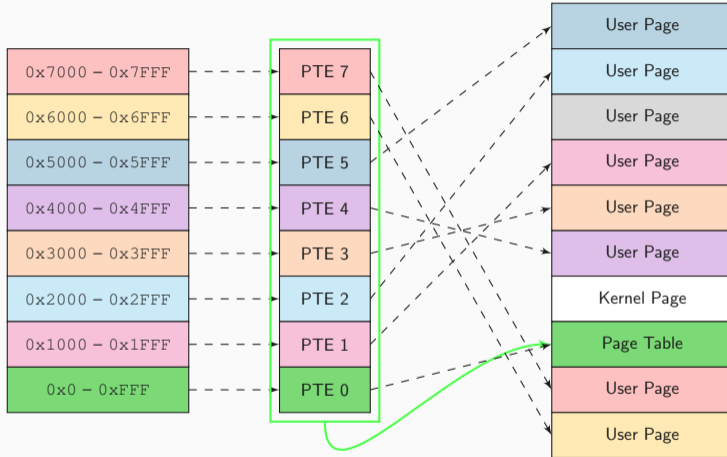
P	RW	US	WT	UC	R	D	S	G	Ignored	
Physical Page Number										
									Ignored	X

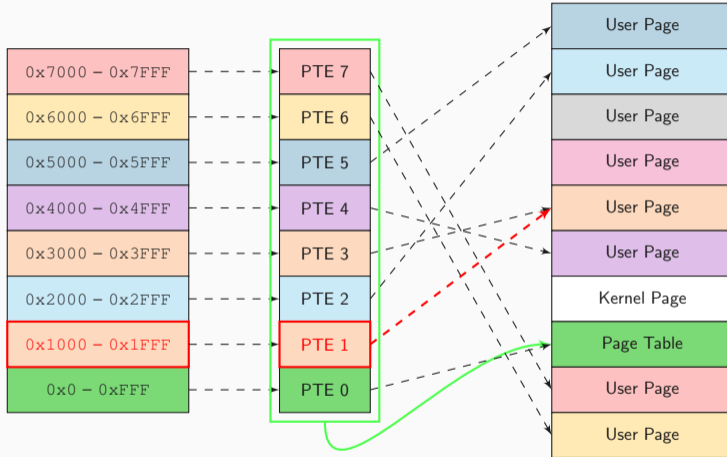
Each 4 KB page table consists of 512 such entries

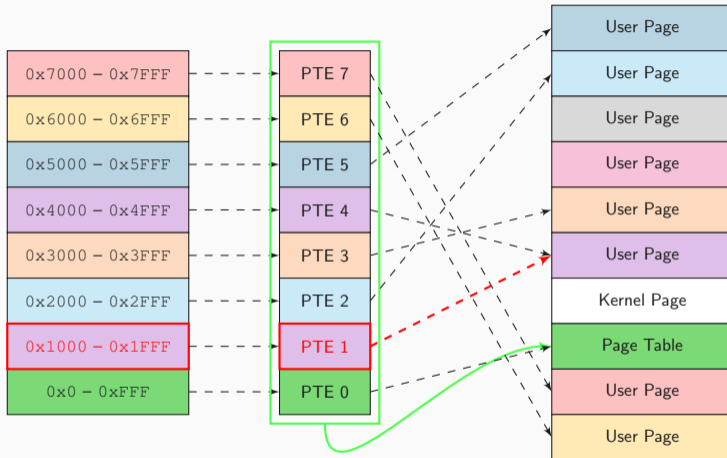


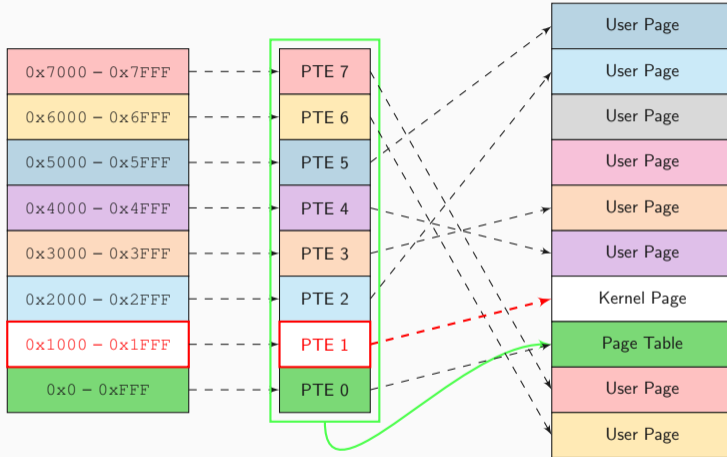


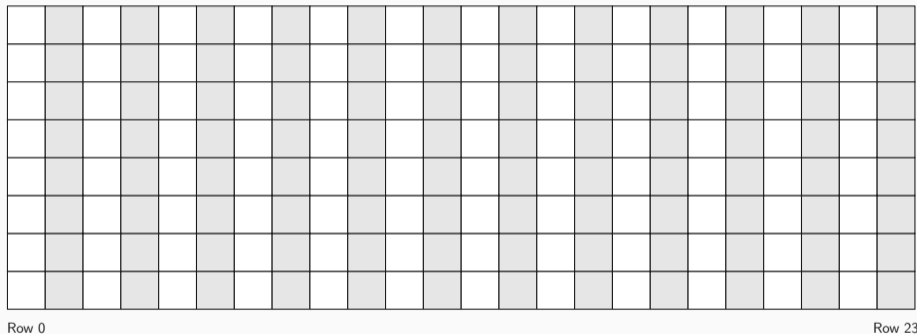




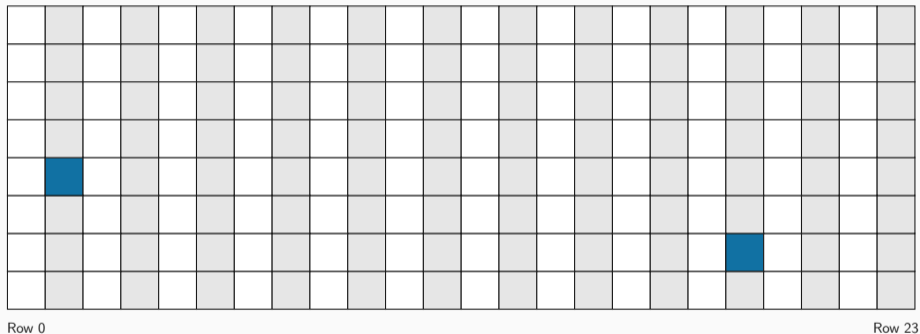




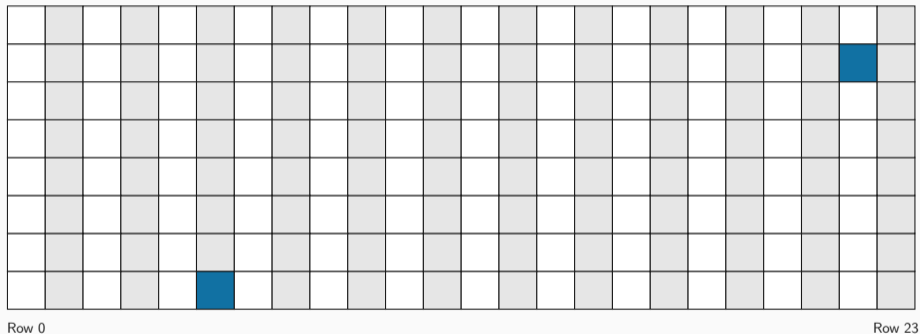




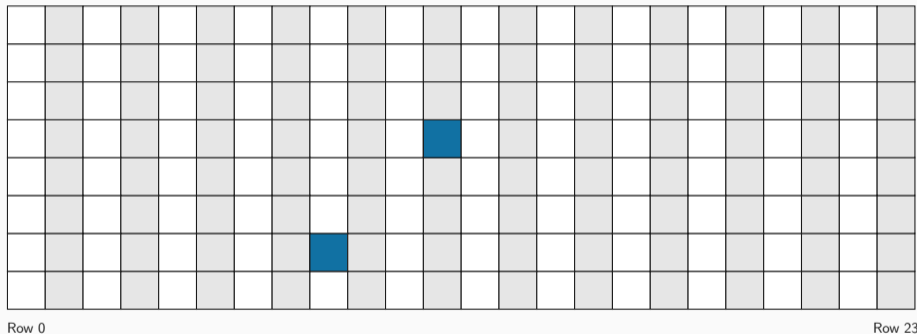
Hammering memory locations in different rows



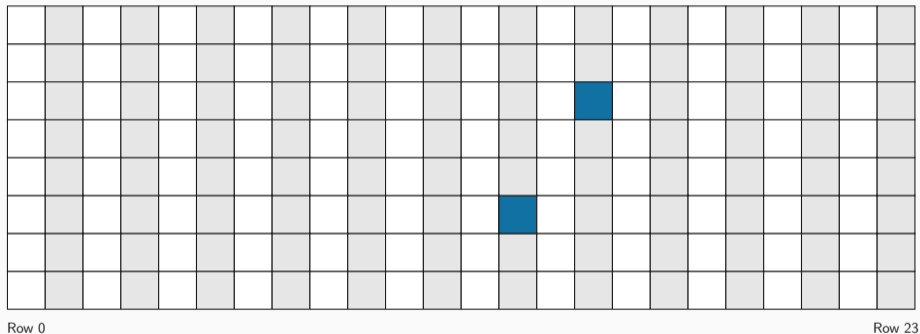
Hammering memory locations in different rows



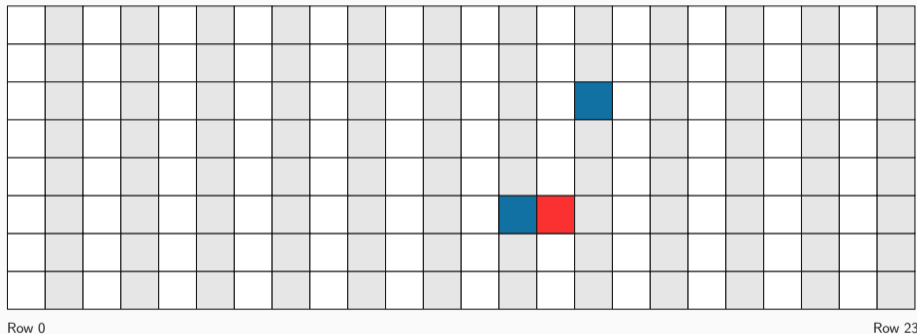
Hammering memory locations in different rows



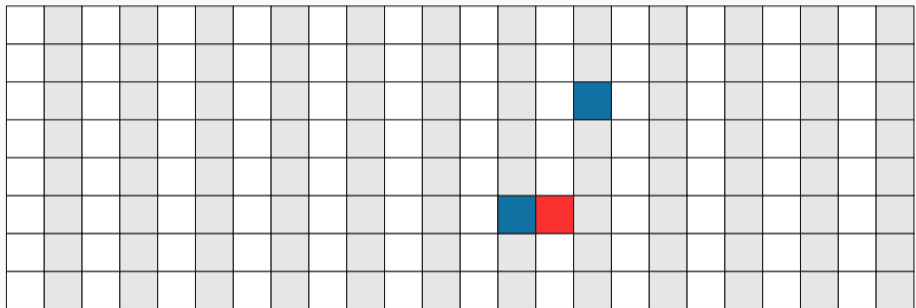
Hammering memory locations in different rows



Hammering memory locations in different rows

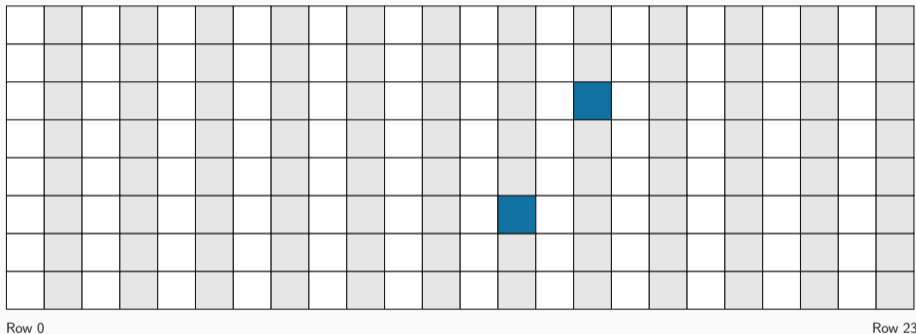


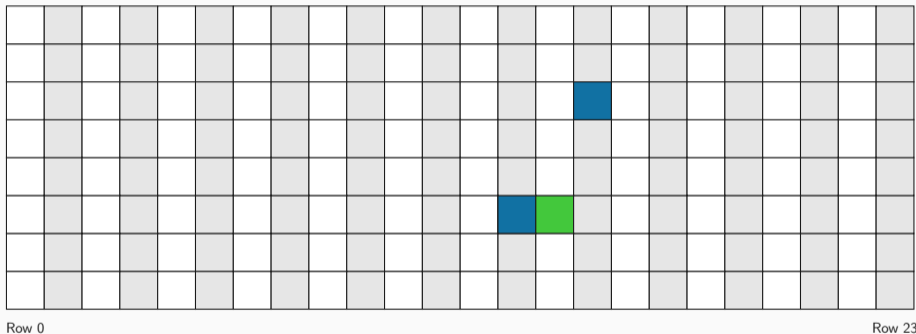
Hammering memory locations in different rows

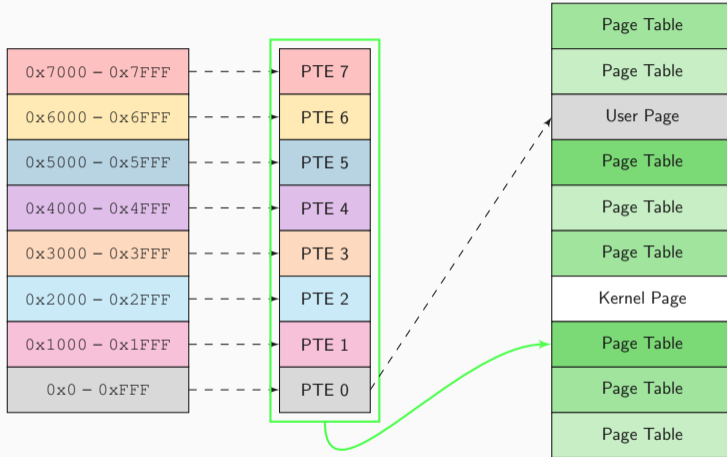


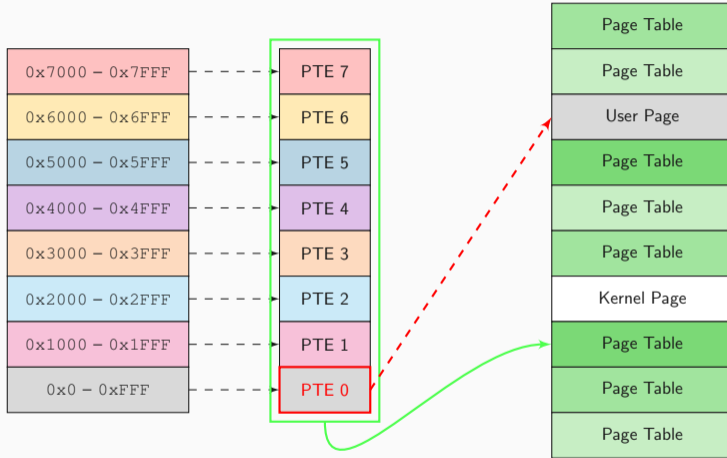
Row 0

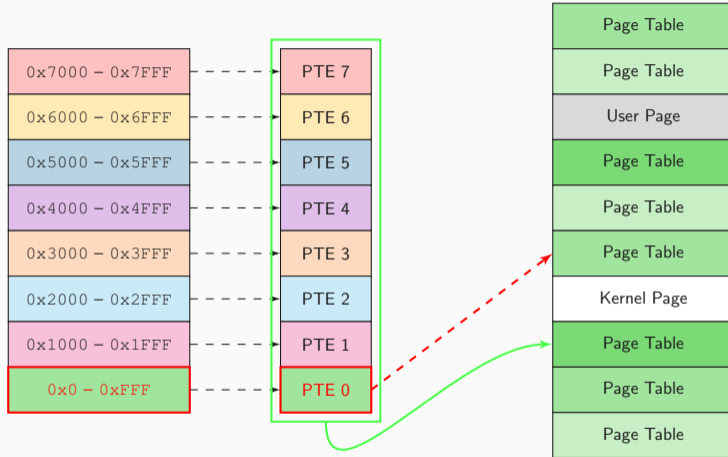
Row 23

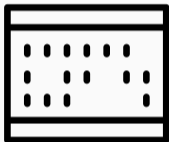


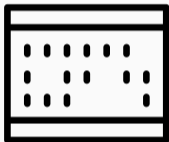


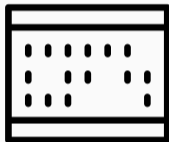




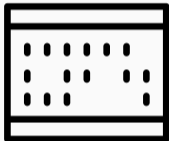




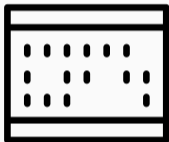




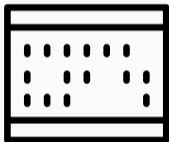
- Idea from [SD15]



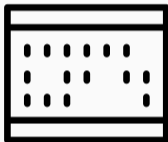
- Idea from [SD15]
- Same idea applied in several other works:



- Idea from [SD15]
- Same idea applied in several other works:
 - Rowhammer.js [GMM16]



- Idea from [SD15]
- Same idea applied in several other works:
 - Rowhammer.js [GMM16]
 - One bit flips, one cloud flops [Xia+16]



- Idea from [SD15]
- Same idea applied in several other works:
 - Rowhammer.js [GMM16]
 - One bit flips, one cloud flops [Xia+16]
 - Drammer [Vee+16]





- Scan entire physical memory (very fast) and:





- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]
 - Read keys [Xia+16]



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]
 - Read keys [Xia+16]
 - Corrupt signatures [BM16; Pod+18]



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]
 - Read keys [Xia+16]
 - Corrupt signatures [BM16; Pod+18]
 - Modify certificates



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]
 - Read keys [Xia+16]
 - Corrupt signatures [BM16; Pod+18]
 - Modify certificates
 - Configurations



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]
 - Read keys [Xia+16]
 - Corrupt signatures [BM16; Pod+18]
 - Modify certificates
 - Configurations
 - etc.



- Scan entire physical memory (very fast) and:
 - Modify binary pages executed in root privileges [Xia+16]
 - Modify credential structs [Vee+16]
 - Read keys [Xia+16]
 - Corrupt signatures [BM16; Pod+18]
 - Modify certificates
 - Configurations
 - etc.
- pages are pretty unique: 32768 bits per page



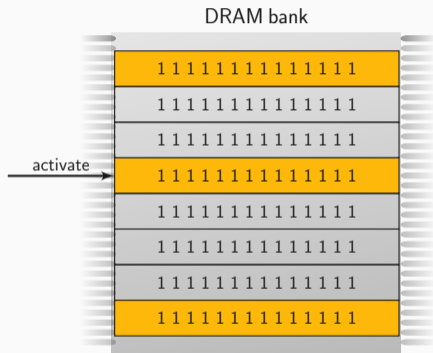
- There are two different hammering techniques

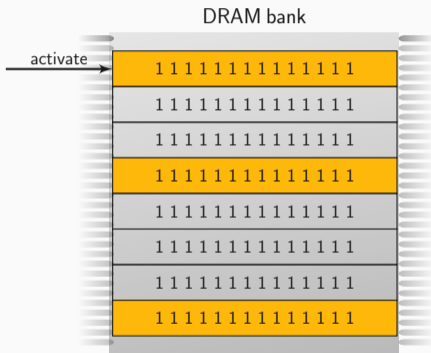


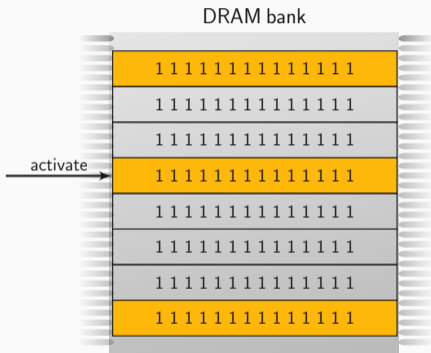
- There are two different hammering techniques
- #1: Hammer one row next to victim row and other random rows

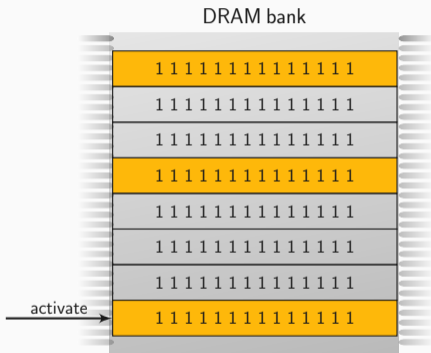


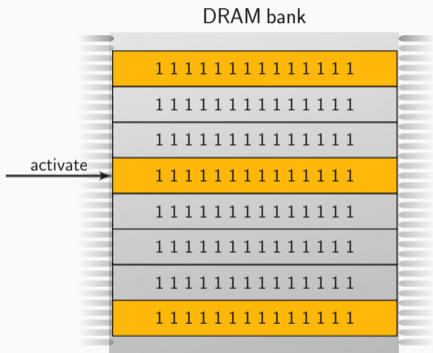
- There are two different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row

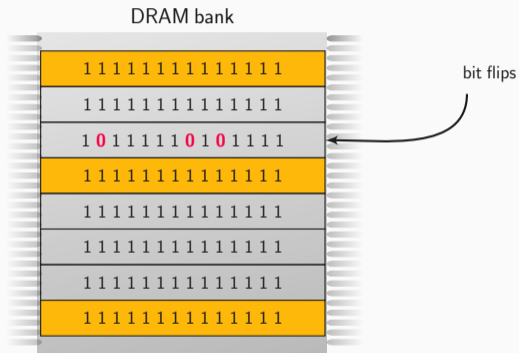


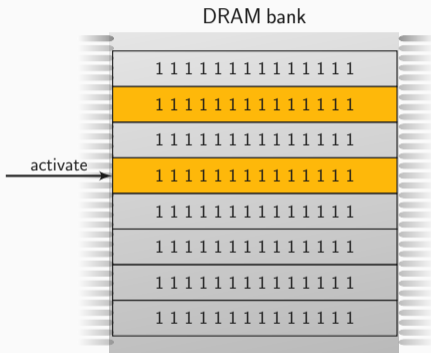


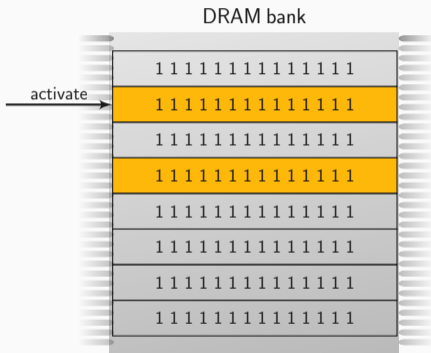


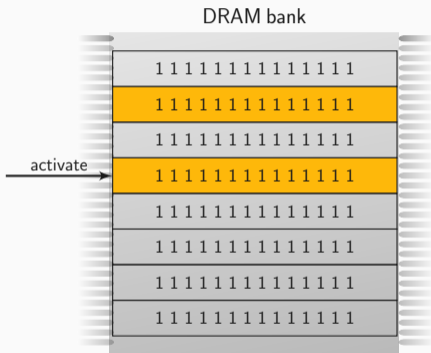


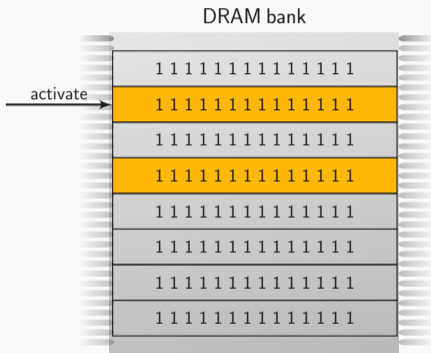


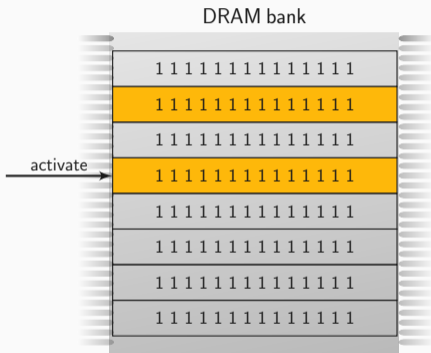


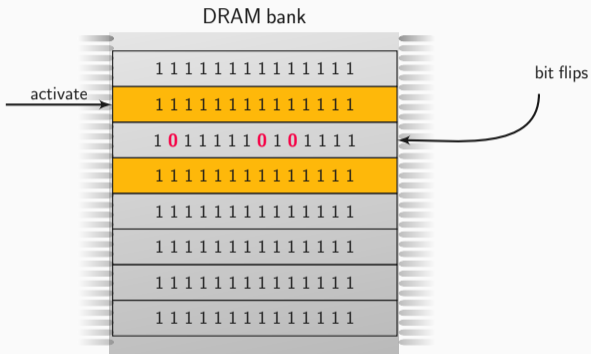














**HAMMERING
TWO ROWS**



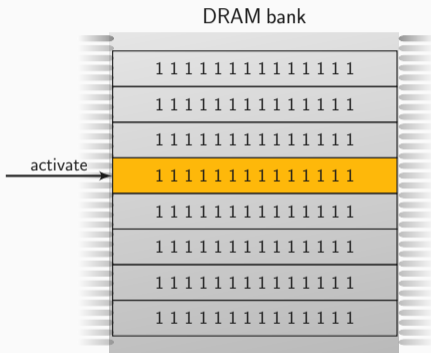
**HAMMERING
TWO ROWS**

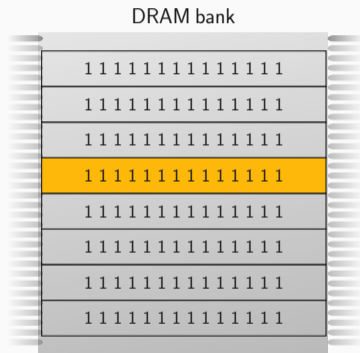


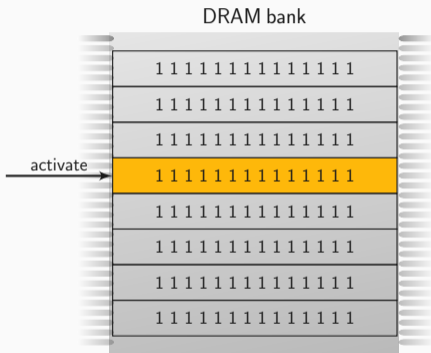
**HAMMERING
A SINGLE ROW**

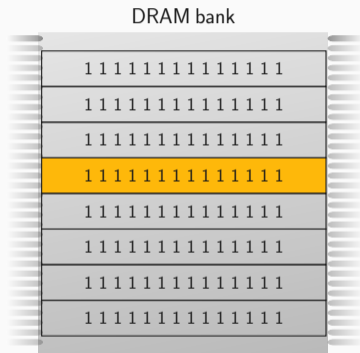


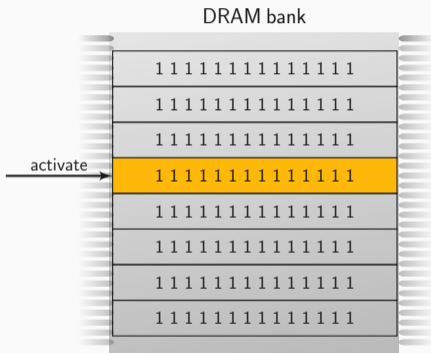
- There are **three** different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row
- **#3: Hammer only one row next to victim row**

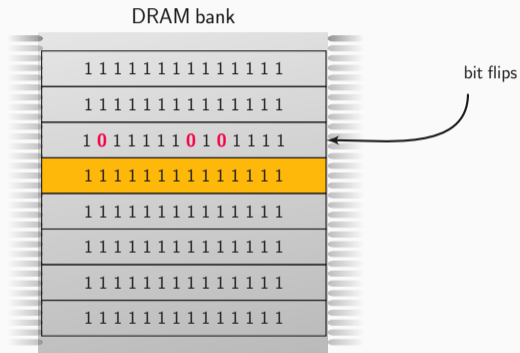






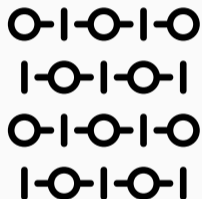




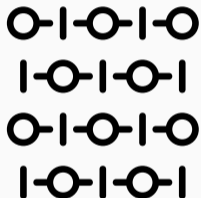


File Edit View Search Terminal Help

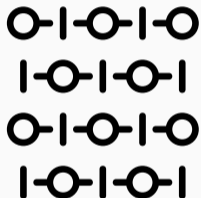
```
dgruss@lab05 ~/flipfloyd (git) -[master] % make
g++ -std=c++11 -O3 -o rowhammer rowhammer.cc
dgruss@lab05 ~/flipfloyd (git) -[master] % ./rowhammer 13
Allocating memory... 90%
```



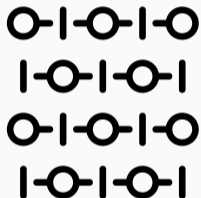
- **Open-page policy:** Keep row opened and buffered



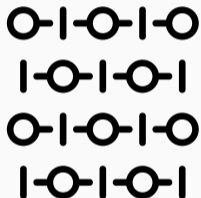
- **Open-page policy:** Keep row opened and buffered
 - Low latency for subsequent accesses to same row
 - High latency for accesses to any other row



- **Open-page policy:** Keep row opened and buffered
 - Low latency for subsequent accesses to same row
 - High latency for accesses to any other row
- **Close-page policy:** Immediately close row, ready to open a new row



- **Open-page policy:** Keep row opened and buffered
 - Low latency for subsequent accesses to same row
 - High latency for accesses to any other row
- **Close-page policy:** Immediately close row, ready to open a new row
 - Medium latency for accesses to any row



- **Open-page policy:** Keep row opened and buffered
 - Low latency for subsequent accesses to same row
 - High latency for accesses to any other row
- **Close-page policy:** Immediately close row, ready to open a new row
 - Medium latency for accesses to any row
 - Perform better on multi-core systems [Dav+11]



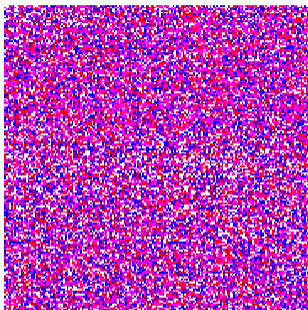
- Policies that preemptively close rows, would allow one-location hammering



- Policies that preemptively close rows, would allow one-location hammering
- We observed close-page policies on desktop computers



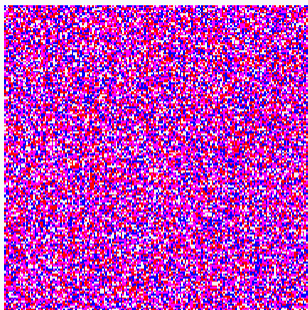
- Policies that **preemptively close rows**, would **allow one-location hammering**
- We observed close-page policies on desktop computers
- Mobile devices (e.g., laptops) seem to use mostly open-page policies



Double-sided

77.0 % bit offsets

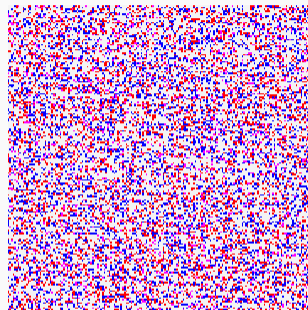
51.7 % 0→1 bit flips



Single-sided

78.5 % bit offsets

54.1 % 0→1 bit flips



One-location

36.5 % bit offsets

51.6 % 0→1 bit flips

What if we cannot target kernel pages?

What if we cannot target kernel pages?

Opcode Flipping



- 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
- Implicitly: e.g., ping or mount



- Many applications perform actions as root
- They can be used by unprivileged users as well
- Implicitly: e.g., ping or mount
- Explicitly: `sudo`



- Many applications perform actions as root
- They can be used by unprivileged users as well
- Implicitly: e.g., ping or mount
- Explicitly: `sudo`
- Target `sudo` (easy to exploit)













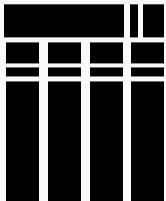




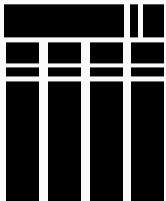
How to get the target virtual page to the target physical location?

How to get the target virtual page to the target physical location?

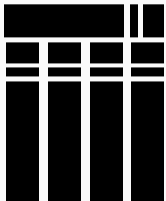
Memory Waylaying



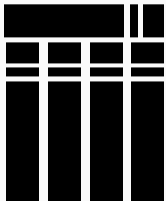
- If a binary is loaded the first time, it is loaded to the memory



- 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



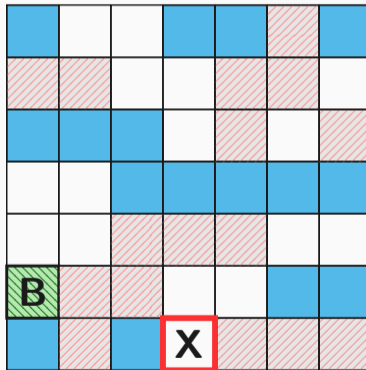
- If a binary is loaded the first time, it is loaded to the memory
- It stays in memory (in the page cache) even after execution
- Only evicted if page cache is full
- Page cache is huge - usually all unused memory



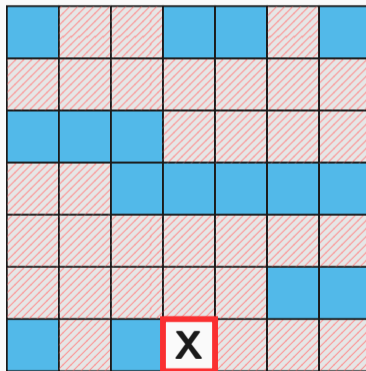
MEMORY WAYLAYING

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

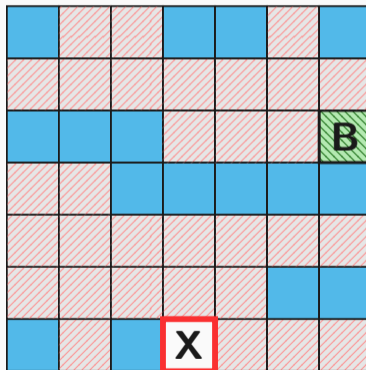
(1) Start



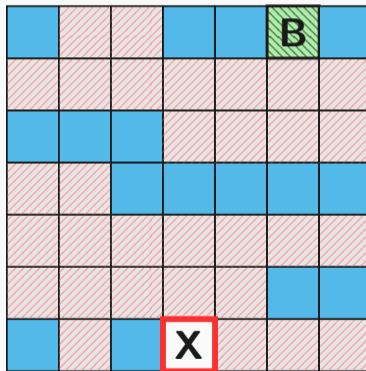
(2) Evict Page Cache



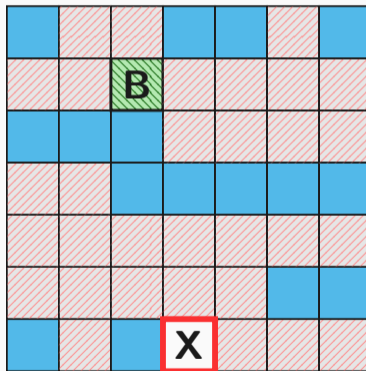
(3) Access Binary



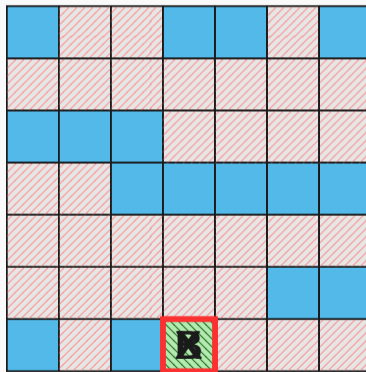
(4) Evict + Access



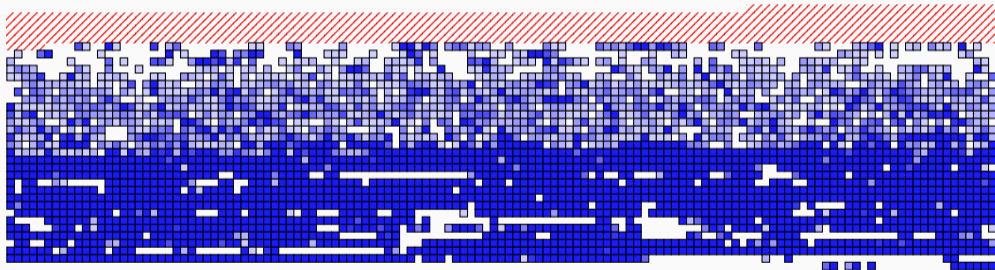
(5) Evict + Access



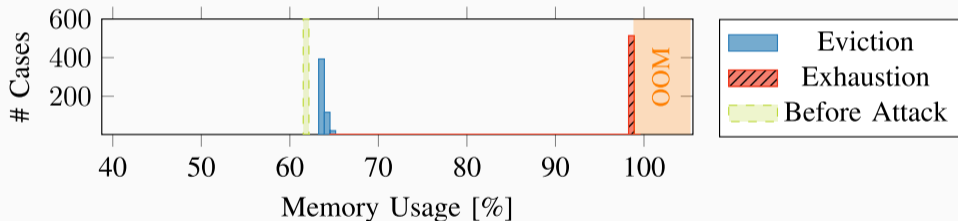
(6) Stop if target reached



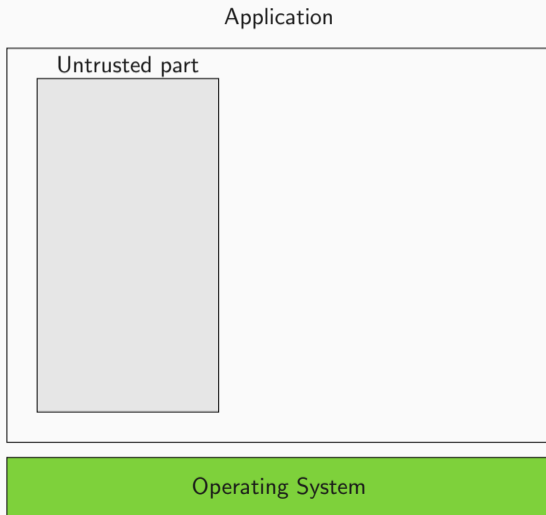
- New pages cover most of the physical memory

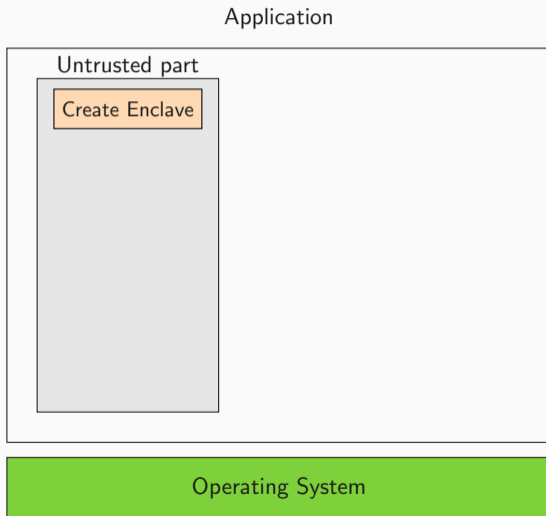


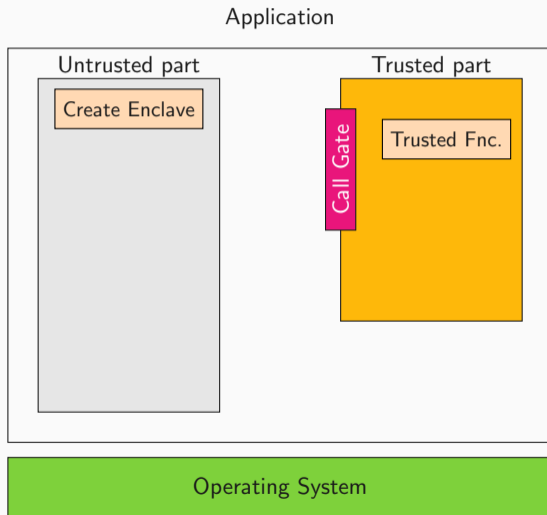
- Great advantage over memory massaging: only negligible memory footprint

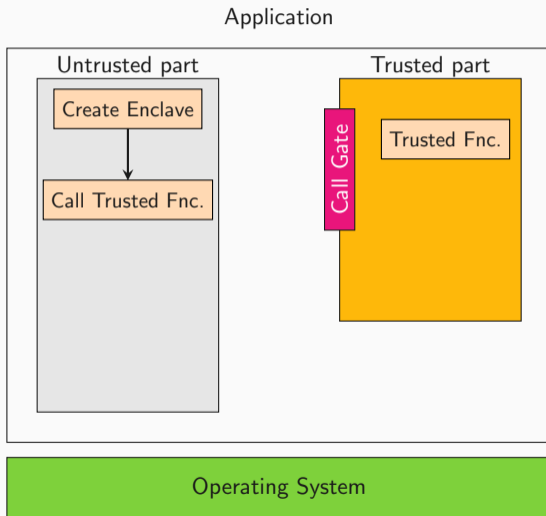


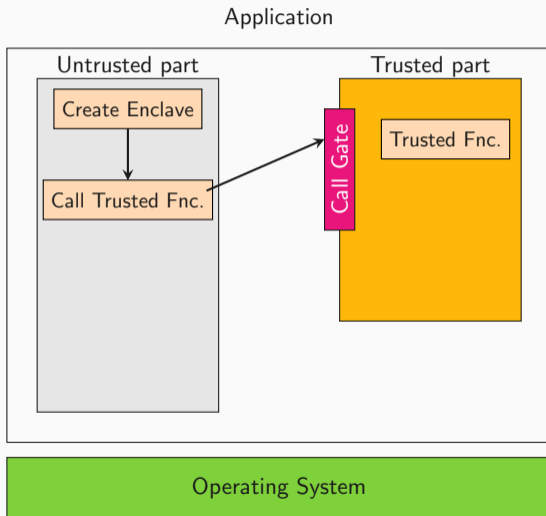
Rowhammer + SGX = Cheap Denial of Service

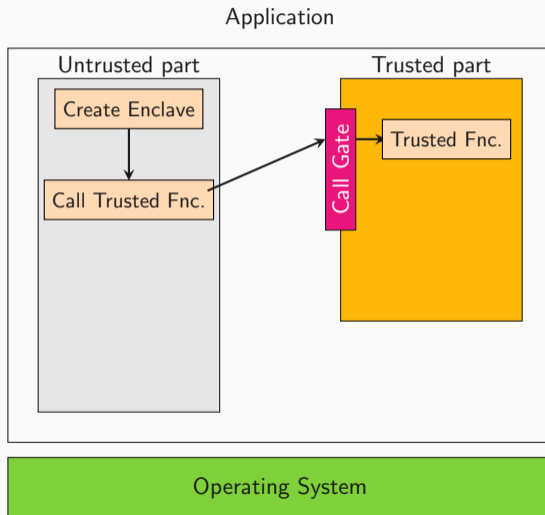


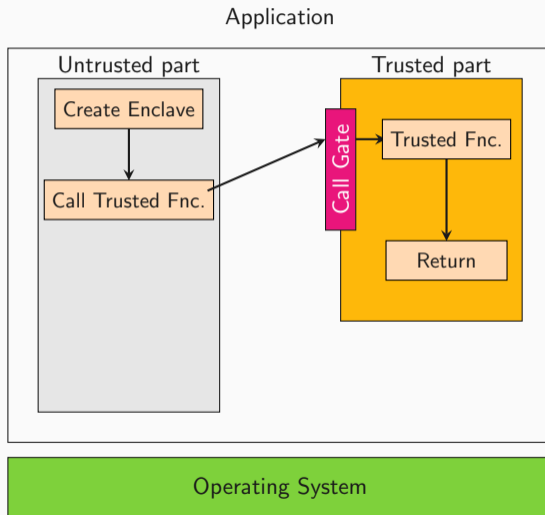


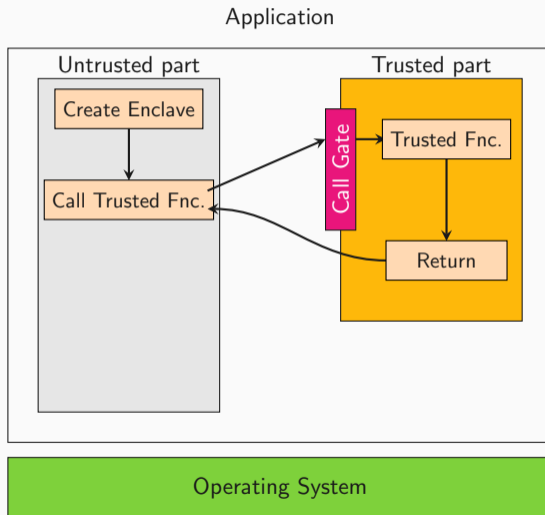


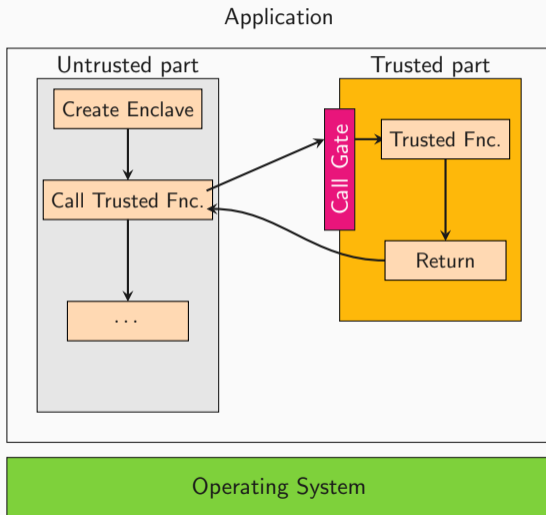


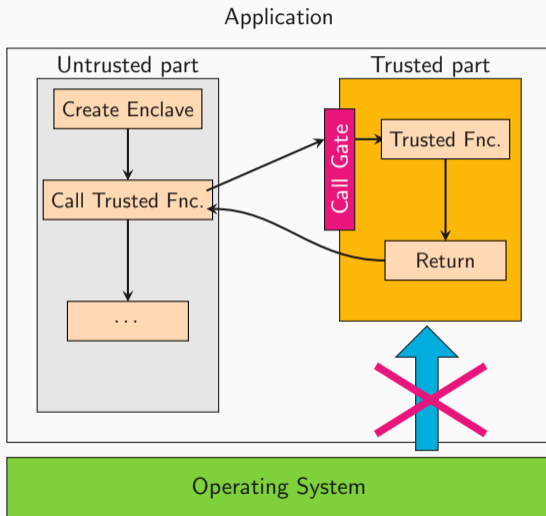




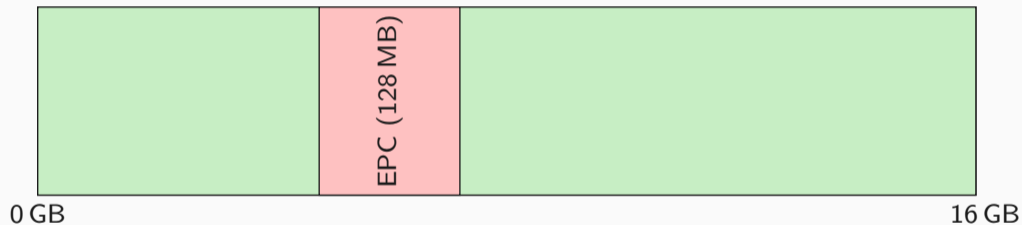


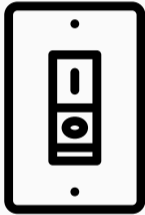




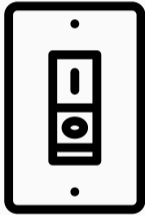




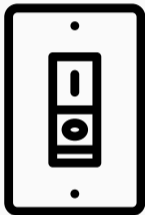




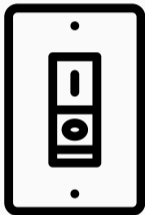
- What happens if a bit flips in the EPC?



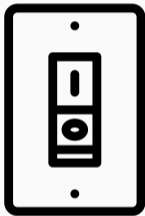
- What happens if a bit flips in the EPC?
- Integrity check will fail!



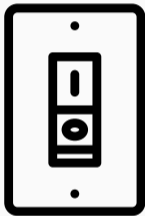
- What happens if a bit flips in the EPC?
 - Integrity check will fail!
- Locks up the memory controller



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



- What happens if a bit flips in the 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 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

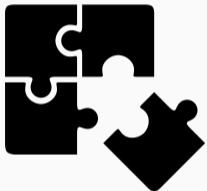


- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts
- ... including co-located tenants

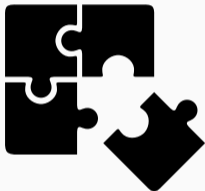


- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts
- ... including co-located tenants
- **Denial-of-Service Attacks in the Cloud** [Gru+18; Jan+17]

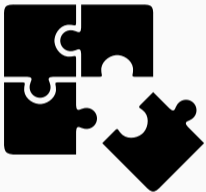
**SGX + One-location Hammering + Opcode Flipping =
Undetectable Exploit**



- SGX protects software from malicious environments



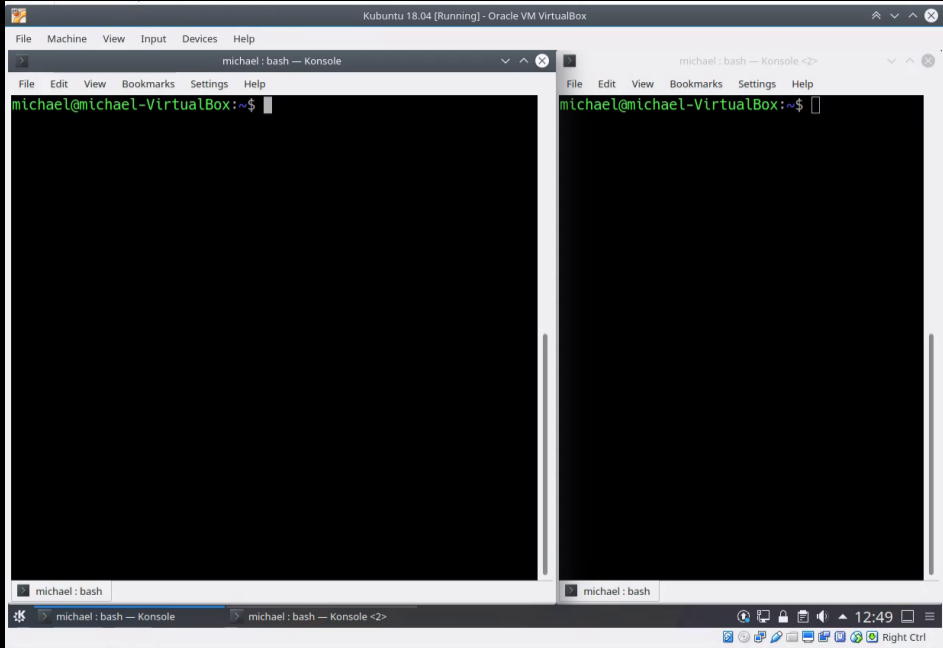
- SGX protects software from malicious environments
- Thwarts static and dynamic (= performance counters) analysis

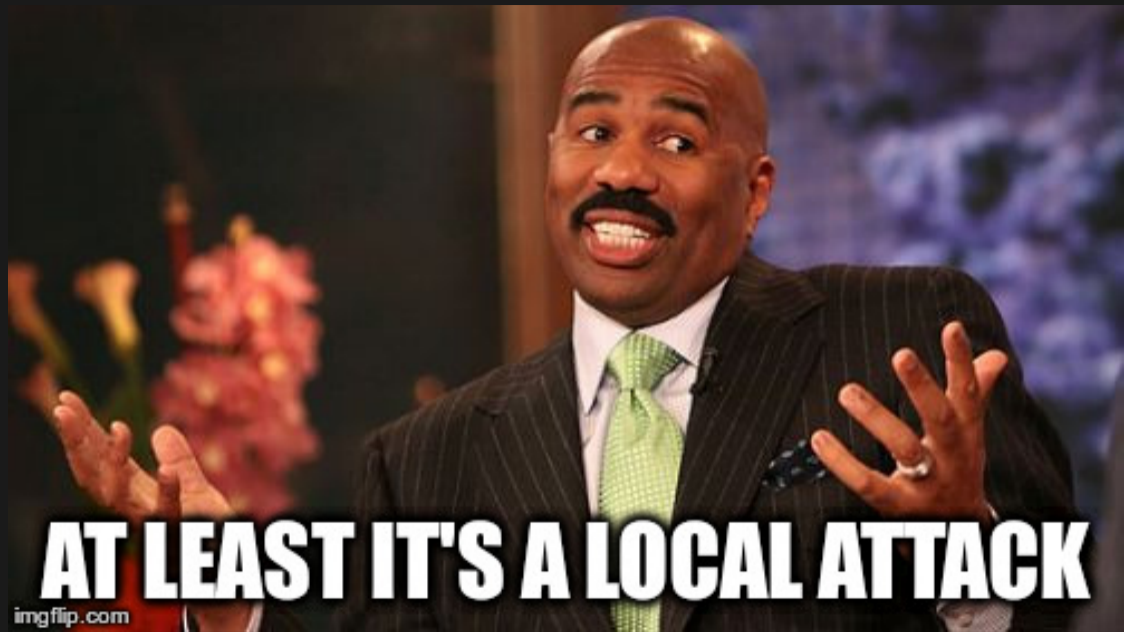


- SGX protects software from malicious environments
- Thwarts static and dynamic (= performance counters) analysis
- Hammering from SGX defeats countermeasures relying on this



STEALTH LEVEL: EXPERT





- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
= 671 875 accesses per second

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
- = 671 875 accesses per second
- Network packets access memory location up to 6 times (depending on kernel)

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
= 671 875 accesses per second
- Network packets access memory location up to 6 times
(depending on kernel)
→ 111 979 packets per second

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
- = 671 875 accesses per second
- Network packets access memory location up to 6 times (depending on kernel)
- 111 979 packets per second
- Network packets are a least 64 B

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
= 671 875 accesses per second
- Network packets access memory location up to 6 times
(depending on kernel)
→ 111 979 packets per second
- Network packets are at least 64 B
= 7 166 656 B/s = 7 MB/s = 57 Mb/s

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
= 671 875 accesses per second
- Network packets access memory location up to 6 times
(depending on kernel)
→ 111 979 packets per second
- Network packets are at least 64 B
= 7 166 656 B/s = 7 MB/s = 57 Mb/s
→ Doable on “modern” networks [Lip+17]

- $\geq 43\,000$ hammering attempts (within 64 ms) for a bit flip [GMM16]
= 671 875 accesses per second
- Network packets access memory location up to 6 times
(depending on kernel)
→ 111 979 packets per second
- Network packets are at least 64 B
= 7 166 656 B/s = 7 MB/s = 57 Mb/s
→ Doable on “modern” networks [Lip+17]







- RMClock



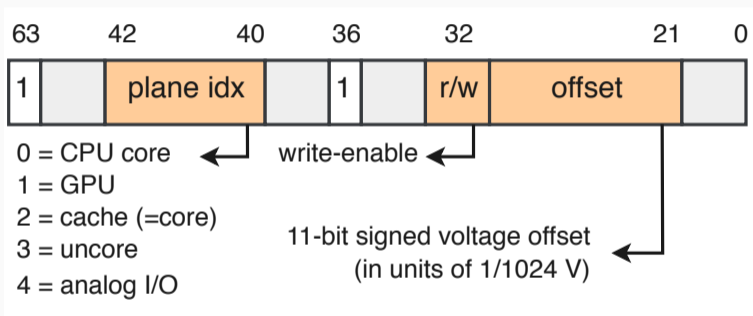
- RMClock
- Throttlestop



- RMClock
- Throttlestop
- linux-intel-undervolt



- RMClock
- Throttlestop
- linux-intel-undervolt
- Adrian Tang's PhD dissertation



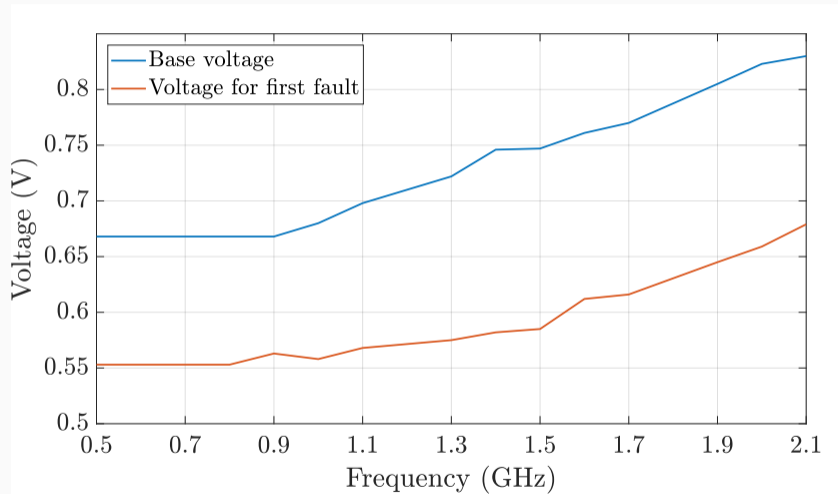

```
uint64_t multiplier = 0x1122334455667788;
uint64_t correct    = 0xdeadbeef * multiplier;
uint64_t var        = 0xdeadbeef * multiplier;

while (var == correct)
{
    var = 0xdeadbeef * multiplier;
}

uint64_t flipped_bits = var ^ correct;
```

bagger> █

I



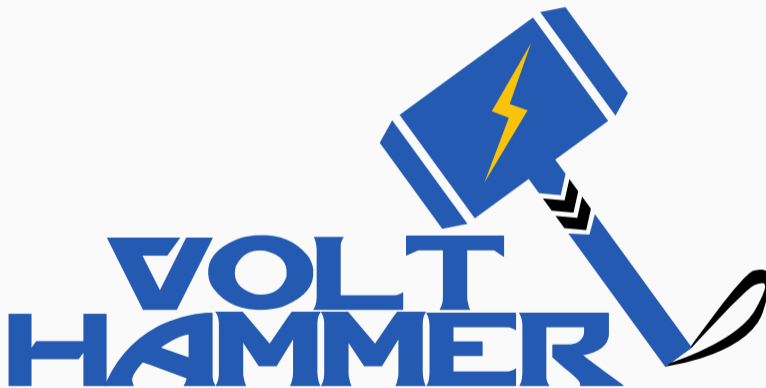


Aristotle Tzafalias

@Aristot73



thesis: given a cool name, researchers will be able to find a new vulnerability to match it, in a finite period of time; cool name-logo combination requirements have no measurable effect on aforementioned period.

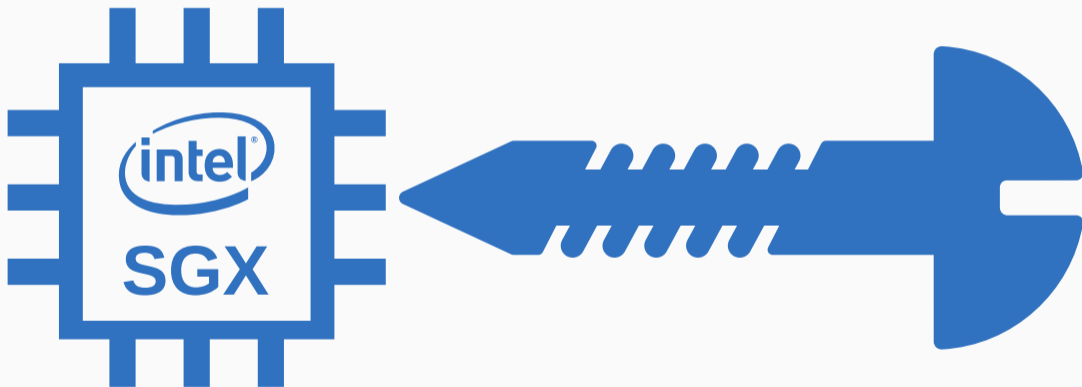


What about the integrity-protected SGX EPC?

[Userspace] Voltage: -261 mV

[Enclave] Multiplying 0xdeadbeef 0x1122334455667788

[Enclave] Multiply_it. result: 0.





- Public Key Crypto
- Encrypt/Sign messages
- Untrusted channel
- Encrypt/Verify messages with public key
- Decrypt/Sign messages with private key

$$n = p \times q$$

$$S = H^d \text{ mod } n$$

$$n = p \times q$$

$$S = H^d \bmod n$$

$$d_p = d \bmod p - 1$$

$$d_q = d \bmod q - 1$$

$$s_1 = H^{d_p} \bmod p$$

$$s_2 = H^{d_q} \bmod q$$

$$n = p \times q$$

$$S = ((q^{-1} \bmod p)(s_1 - s_2) \bmod p) \times q + s_2$$

$$d_p = d \bmod p - 1$$

$$d_q = d \bmod q - 1$$

$$s_1 = H^{d_p} \bmod p$$

$$s_2 = H^{d_q} \bmod q$$

$$n = p \times q$$

$$S = ((q^{-1} \bmod p)(s_1 - s_2) \bmod p) \times q + s_2$$

$$d_p = d \bmod p - 1$$

$$d_q = d \bmod q - 1$$

$$s_1 = H^{d_p} \bmod p$$

$$s_2 = H^{d_q} \bmod q$$

$$n = p \times q$$

$$S = ((q^{-1} \bmod p)(s_1 - s_2) \bmod p) \times q + s_2$$

$$S' = ((q^{-1} \bmod p)(s_1 - s_2) \bmod p) \times q + s_2$$

$$n = p \times q$$

$$S = \left(\overbrace{\text{something}}^{\text{length } p} \right) \times q + s_2$$

$$S' = \left(\overbrace{\text{something else}}^{\text{length } p} \right) \times q + s_2$$

$$n = p \times q$$

$$S = \left(\overbrace{\text{something}} \right) \times q + s_2$$

$$S' = \left(\overbrace{\text{something else}} \right) \times q + s_2$$

$$S - S' = \text{something} \times q$$

$$n = p \times q$$

$$S = \left(\overbrace{\hspace{10em}}^{\text{something}} \right) \times q + s_2$$

$$S' = \left(\overbrace{\hspace{10em}}^{\text{something else}} \right) \times q + s_2$$

$$S - S' = \hspace{10em} \text{something} \hspace{10em} \times q$$

$$q = \gcd(S - S', n)$$

```
uint8_t rsa_dec_ecall(int iterations)
{
    // Wait for first fault
    trigger_fault(iterations);

    // Actual decryption
    ippsRSA_Decrypt(ct, dec, pPrv, scratchBuffer);
}
```

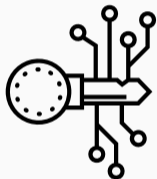
```
bagger> dog Enclave/encl
```

A close-up of Mufasa's face from Disney's 'The Lion King'. He has a serious, somewhat stern expression with a slight frown. The background is a bright blue sky with scattered white clouds.

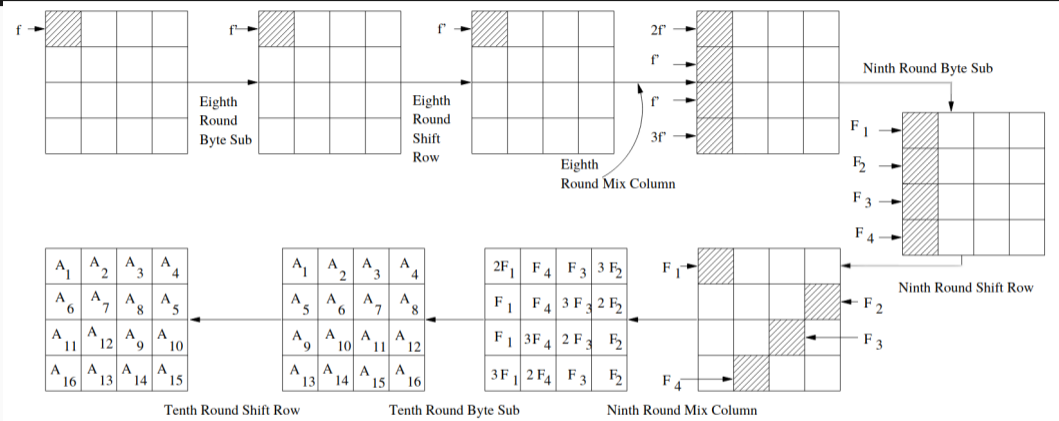
MUFARSA

**My Undervolting
Fault Attack on RSA**

**WHAT ELSE
CAN WE
BREAK?!**



- Symmetric Key Crypto
- Encrypt messages for transfer over public channel
- Encrypt data for (untrusted) storage
- 4×4 byte state
- 10 rounds: S-Box, ShiftRows, MixColumns, AddRoundKey



Michael Tunstall et al. Differential Fault Analysis of the Advanced Encryption Standard Using a Single Fault. In: International Workshop on Information Security Theory and Practices. 2011 [TMA11]

Instruction	Description
<code>AESENC</code>	Perform one round of an AES encryption flow
<code>AESENCLAST</code>	Perform the last round of an AES encryption flow
<code>AESDEC</code>	Perform one round of an AES decryption flow
<code>AESDECLAST</code>	Perform the last round of an AES decryption flow
<code>AESKEYGENASSIST</code>	Assist in AES round key generation
<code>AESIMC</code>	Assist in AES Inverse Mix Columns
<code>PCLMULQDQ</code>	Carryless multiply (CLMUL)


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

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

```
bagger> sudo ./aes-encrypt 100000 -262
```





```
struct_foo_t *foo = &arr[offset];  
foo->foo = enclave_secret;
```

```
struct_foo_t *foo = &arr[offset];  
foo->foo = encode_secret;
```



```
foo = arr + offset * 0x24
```

```
foo = arr + offset * 0x24
```



```
Creating enclave...  
==== Victim Enclave ====  
[pt.c] /dev/sgx-step opened!  
Enclave Base: 0x7f001a000000 ←  
Enclave Limit: 0x7f001c000000 ←  
EDBGRD: debug  
█
```

Voltage
0.584V

Undervolting
-235mV



- 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**, ...



Faulting Hardware from Software

Daniel Gruss

2020-09-13

Graz University of Technology