

Rowhammer: From the Basics to Sophisticated New Variants

Moritz Lipp, Michael Schwarz

February 20, 2018

Graz University of Technology

- **Moritz Lipp**

- PhD Student @ Graz University of Technology
-  @mlqxyz
-  moritz.lipp@iaik.tugraz.at

- **Michael Schwarz**

- PhD Student @ Graz University of Technology
-  @misc0110
-  michael.schwarz@iaik.tugraz.at

- Security and privacy rely on secrets (unknown to attackers)

- Security and privacy rely on secrets (unknown to attackers)
- Secrets can leak through side channels

- Security and privacy rely on secrets (unknown to attackers)
- Secrets can leak through side channels
- Software-based attacks → no physical access

Software-based side-channel attacks leak secrets via...

- Algorithm runtimes
- Cache states
- DRAM states
- ...

- Rowhammer has a special rule among microarchitectural attacks



- Rowhammer has a special rule among microarchitectural attacks
- Does not leak secrets...

- Rowhammer has a special rule among microarchitectural attacks
- Does not leak secrets...
- ...but it can change memory contents

- Rowhammer has a special rule among microarchitectural attacks
- Does not leak secrets...
- ...but it can change memory contents
- A software-based fault attack on the DRAM



Background

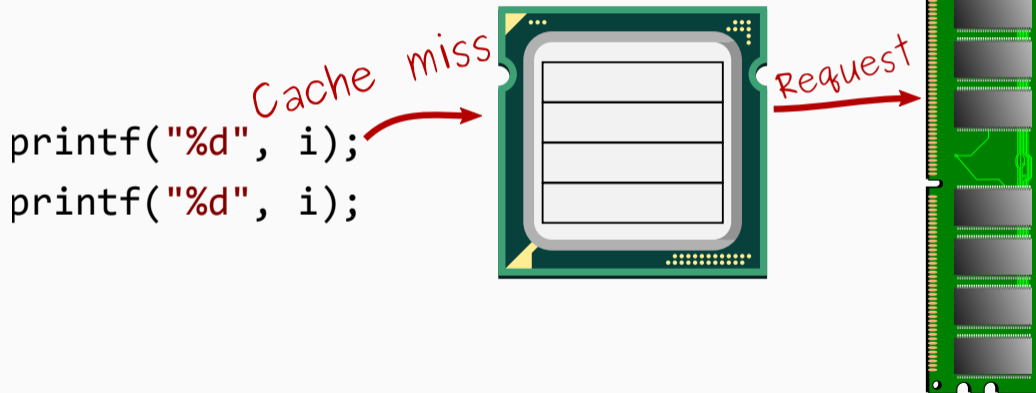
```
printf("%d", i);  
printf("%d", i);
```

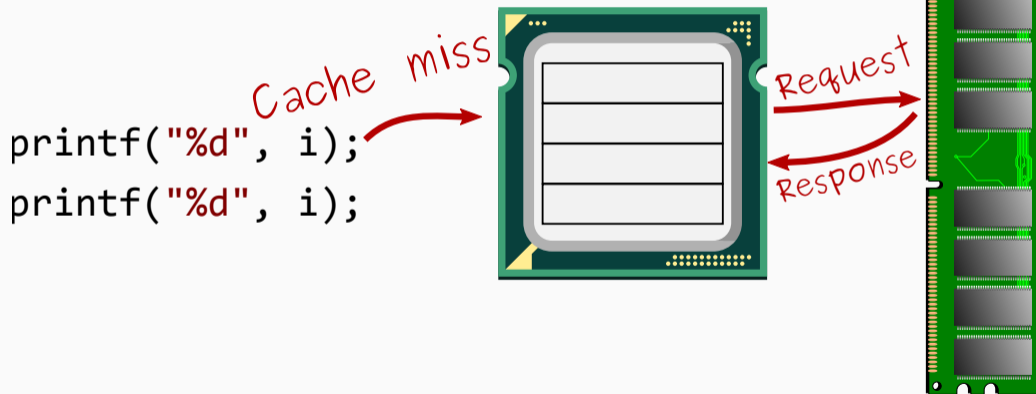



```
printf("%d", i);  
printf("%d", i);
```

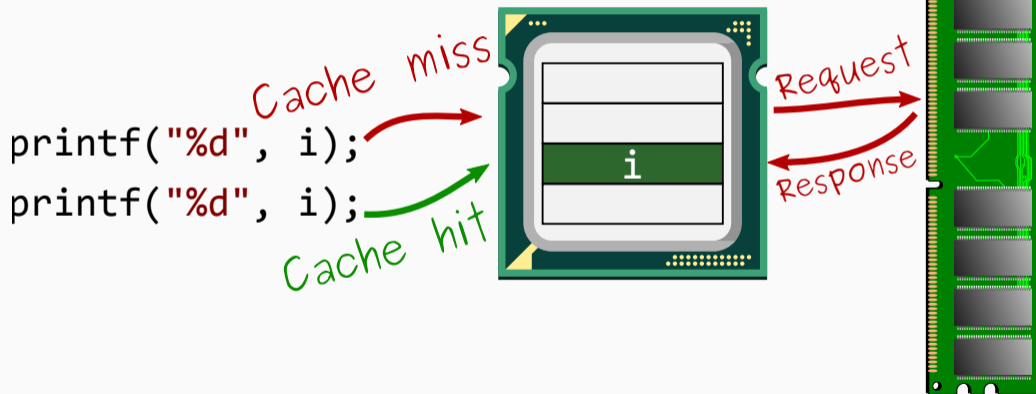
Cache miss




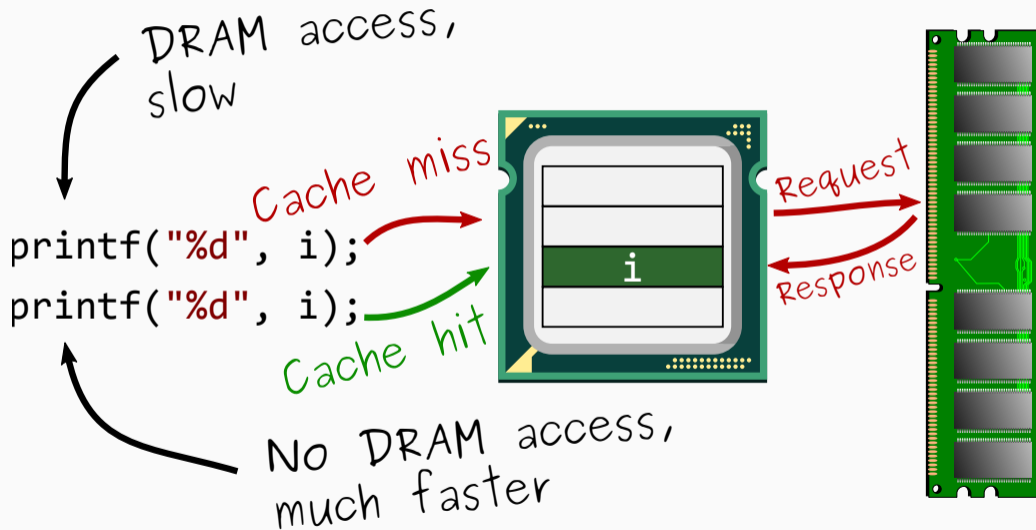


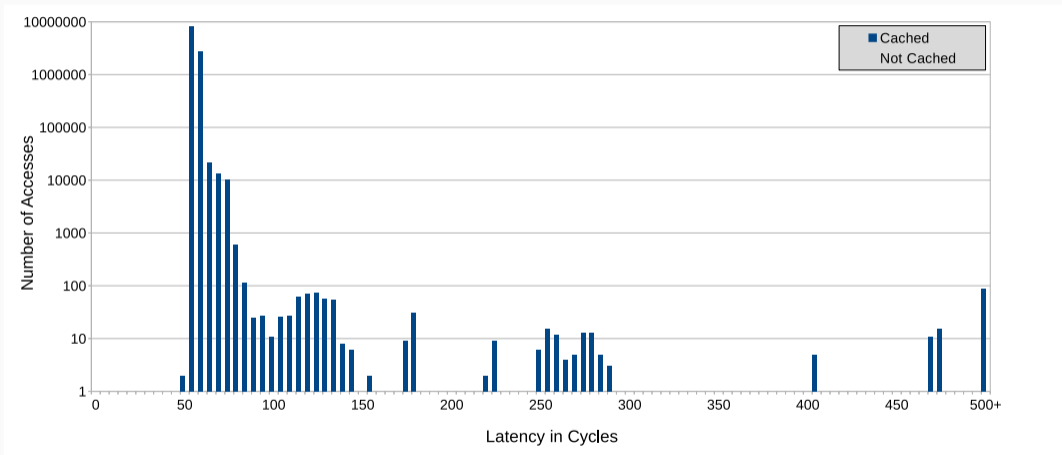


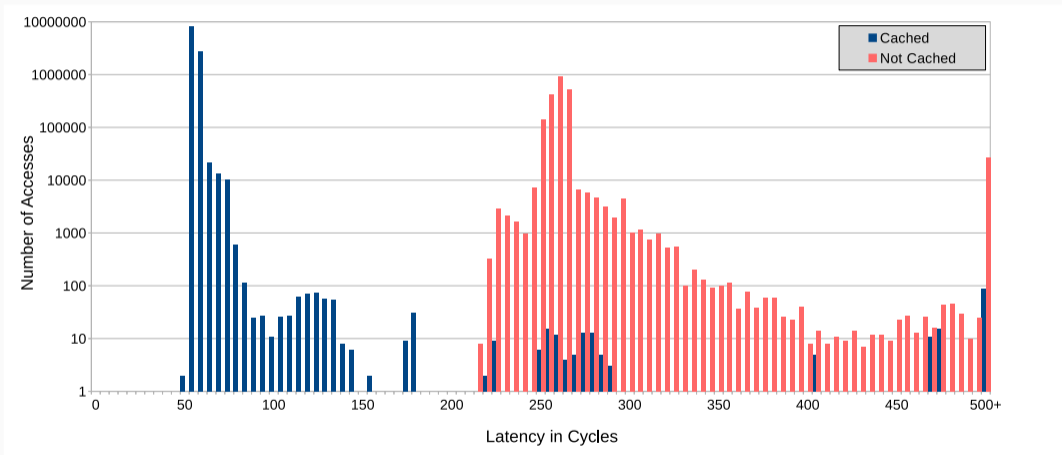













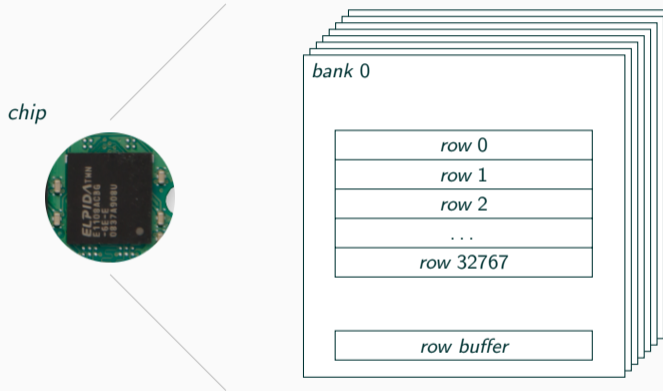


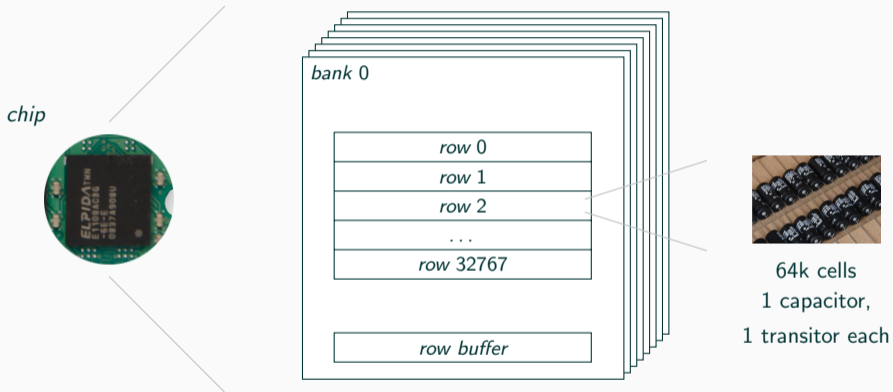













- DRAM internally is only capable of reading entire rows

- DRAM internally is only capable of reading entire rows
- Capacitors in cells discharge when reading them

- DRAM internally is only capable of reading entire rows
- Capacitors in cells discharge when reading them
- Bits are buffered when reading them from the cells


- DRAM internally is only capable of reading entire rows
- Capacitors in cells discharge when reading them
- Bits are buffered when reading them from the cells
- Then, bits are written back to the cells again


- DRAM internally is only capable of reading entire rows
- Capacitors in cells discharge when reading them
- Bits are buffered when reading them from the cells
- Then, bits are written back to the cells again
- → Row buffer




CPU reads row 1:

→ activate and copy to row buffer









CPU reads row 2:

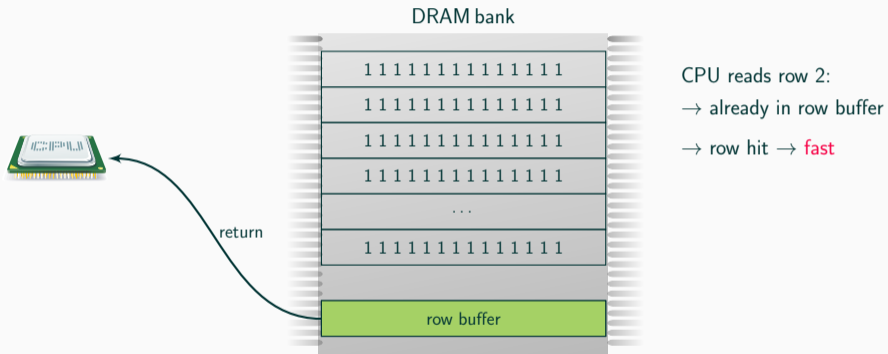
→ activate and copy to row buffer

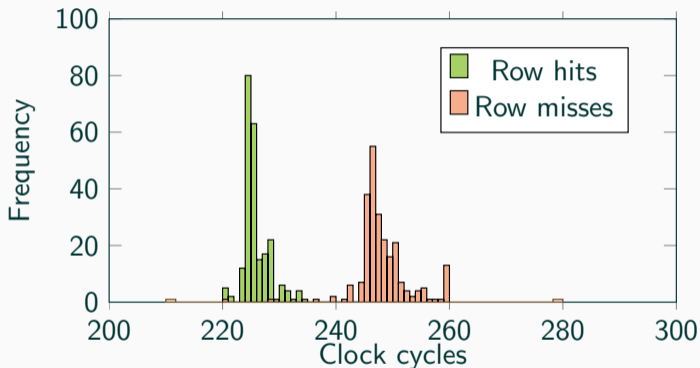
→ row conflict → **slow**






CPU reads row 2:
→ already in row buffer
→ row hit → **fast**






Row hits (≈ 225 cycles) and row conflicts (≈ 247 cycles)


Rowhammer




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



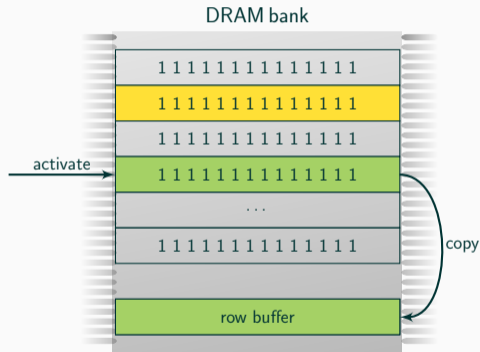
- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer



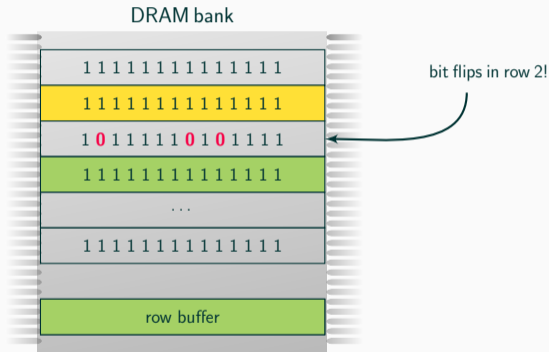
- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → Rowhammer



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



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




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

DDR3:

- Kim et al.: 110/129 modules from 3 vendors, all but 3 since mid-2011
- Seaborn and Dullien: 15/29 laptops

DDR4 believed to be safe:


- we showed bit flips (Pessl et al. 2016)





Prevalence, by Kim et al. 2014


1. Flush cache via `clflush` instruction → original paper (Kim et al. 2014)
2. Cache eviction (Gruss, Maurice, et al. 2016; Aweke et al. 2016)
3. Non-temporal accesses (Qiao et al. 2016)
4. Uncached memory (Veen et al. 2016)

























































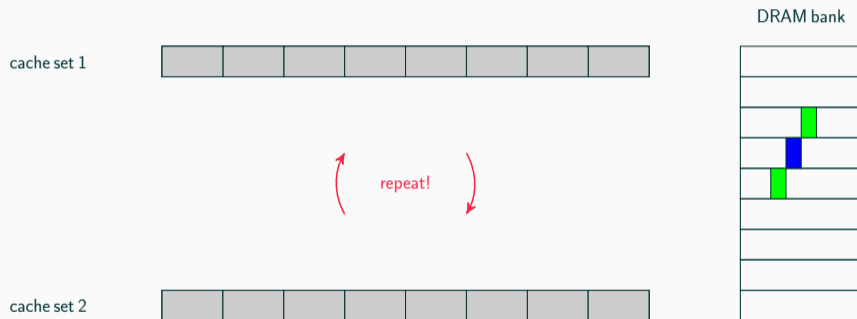
















cache set 1



cache set 2




DRAM bank

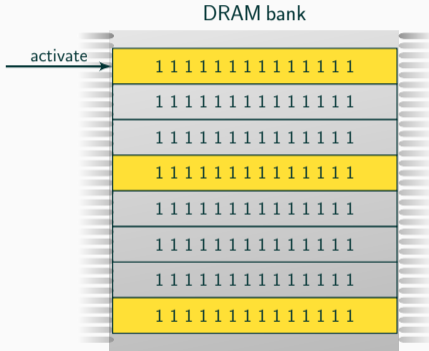


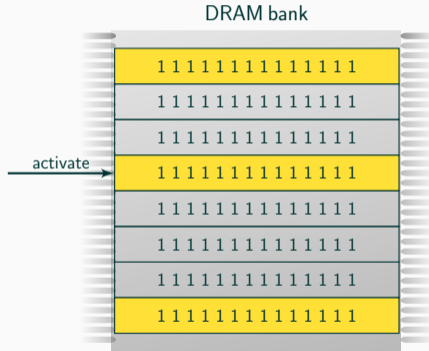
bit flip!




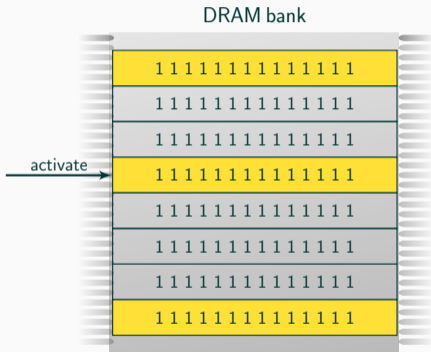
- There are three different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row
- #3: Hammer only one row next to victim row

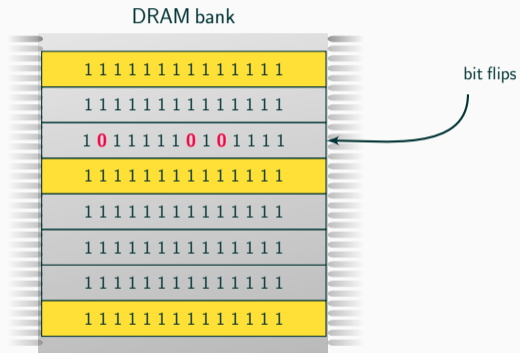

















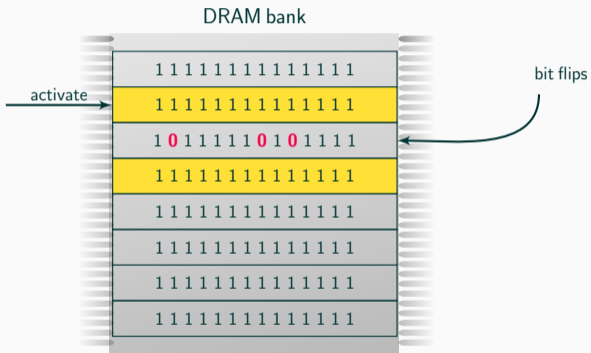






























- We proposed one-location hammering

- We proposed one-location hammering
- The only technique which accesses only one row

- We proposed one-location hammering
- The only technique which accesses only one row
- Circumvents all countermeasures which expect accesses to multiple rows

- We proposed one-location hammering
- The only technique which accesses only one row
- Circumvents all countermeasures which expect accesses to multiple rows
- Why does it work?

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

- **Open-page policy:** Keep row opened and buffered
 - Beneficial for memory access latency, power consumption, bank utilization for a low number of memory accesses

- **Open-page policy:** Keep row opened and buffered
 - Beneficial for memory access latency, power consumption, bank utilization for a low number of memory accesses
- **Close-page policy:** Immediately close row, ready to open a new row

- **Open-page policy:** Keep row opened and buffered
 - Beneficial for memory access latency, power consumption, bank utilization for a low number of memory accesses
- **Close-page policy:** Immediately close row, ready to open a new row
 - Achieve better performance when a high number of memory accesses is performed


- **Open-page policy:** Keep row opened and buffered
 - Beneficial for memory access latency, power consumption, bank utilization for a low number of memory accesses
- **Close-page policy:** Immediately close row, ready to open a new row
 - Achieve better performance when a high number of memory accesses is performed
 - Perform better on multi-core systems David et al. 2011

- 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


- Distribution of bit flips over 4kb-aligned memory regions
- Test each technique for 8 hours
- Scanned for bit flips after every hammering attempt
 - Hammering a random location of more than 100 000 randomly-chosen 4kB pages



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

Exploits

- 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

| P | RW | US | WT | UC | R | D | S | G | | |
|------------|----|----|----|---------|---|---|---|---|--|---|
| [Redacted] | | | | | | | | | | |
| [Redacted] | | | | | | | | | | |
| [Redacted] | | | | [Green] | | | | | | X |

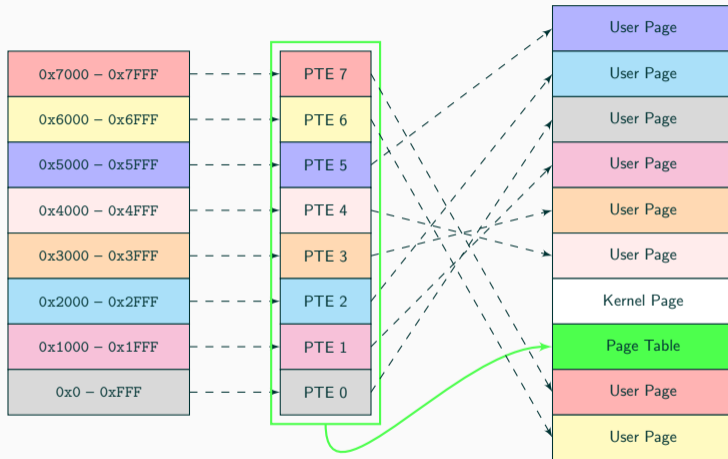
| | | | | | | | | | | |
|---|----|----|----|---------|---|---|---|---|---------|---|
| P | RW | US | WT | UC | R | D | S | G | Ignored | |
| | | | | | | | | | | |
| | | | | | | | | | | |
| | | | | Ignored | | | | | | X |

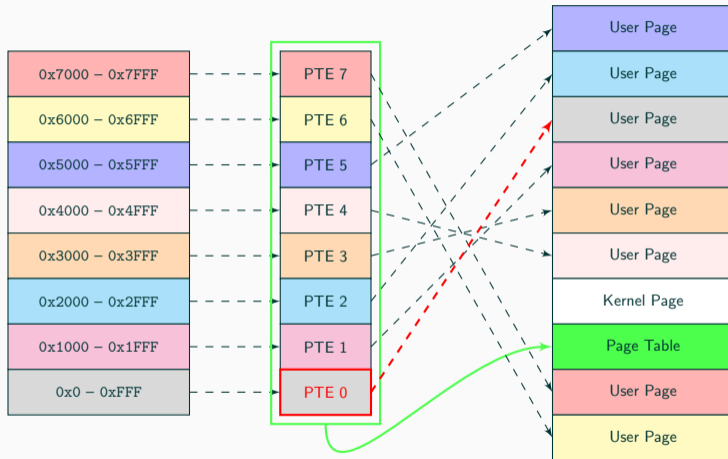
| | | | | | | | | | | |
|----------------------|----|----|----|----|---|---|---|---|---------|---|
| P | RW | US | WT | UC | R | D | S | G | Ignored | |
| Physical Page Number | | | | | | | | | | |
| | | | | | | | | | Ignored | X |

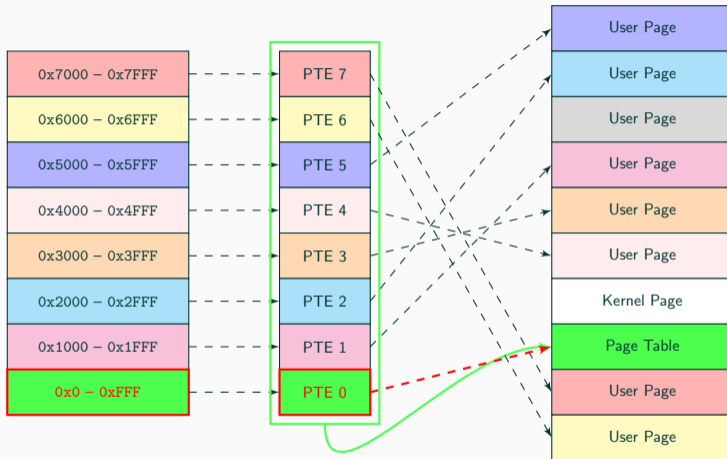
| | | | | | | | | | | |
|-----------------------------|----|----|----|----|---|---|---|---|---------|---|
| P | RW | US | WT | UC | R | D | S | G | Ignored | |
| Physical Page Number | | | | | | | | | | |
| | | | | | | | | | Ignored | X |

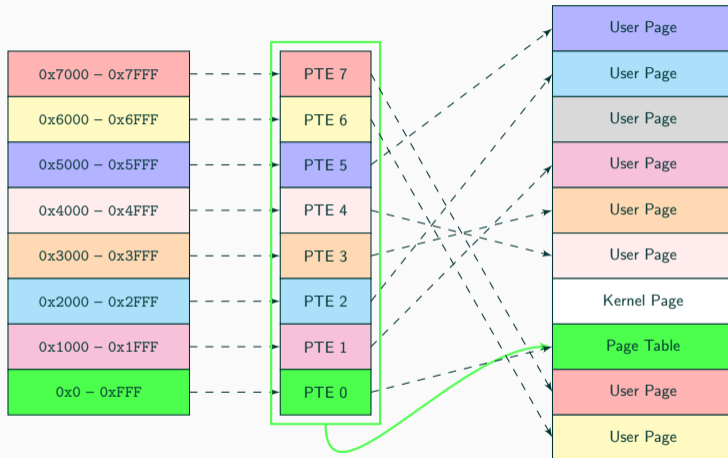
Each 4 KB page table consists of 512 such entries

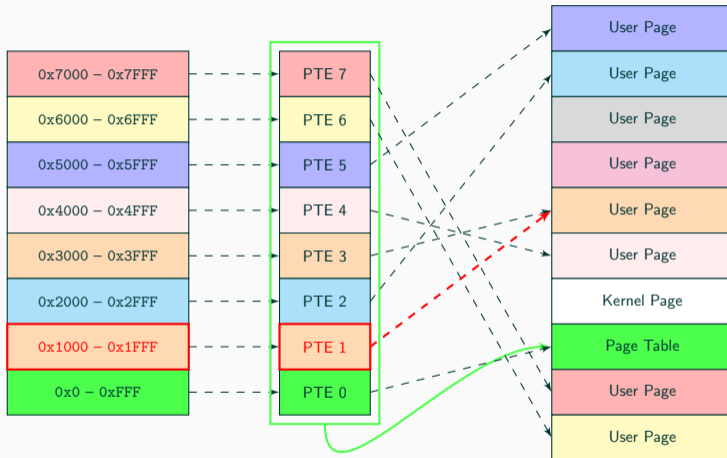
Getting root by targeting the kernel

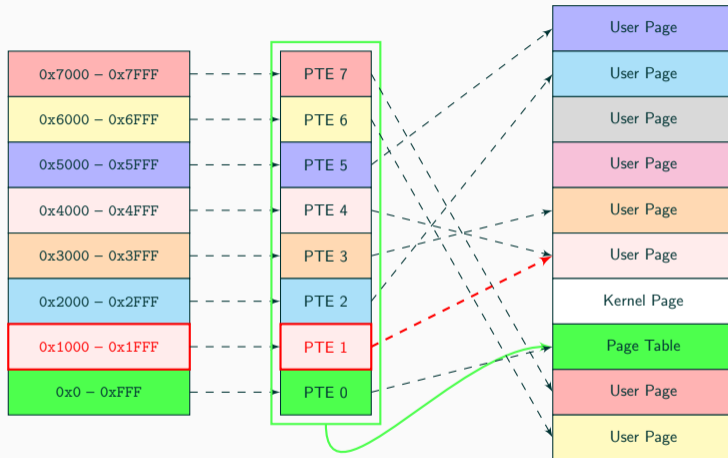


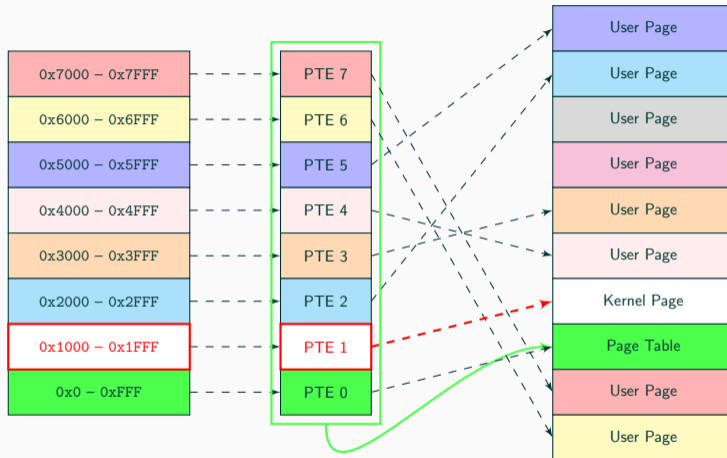


















Hammering memory locations in different rows




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




Release bit flip page



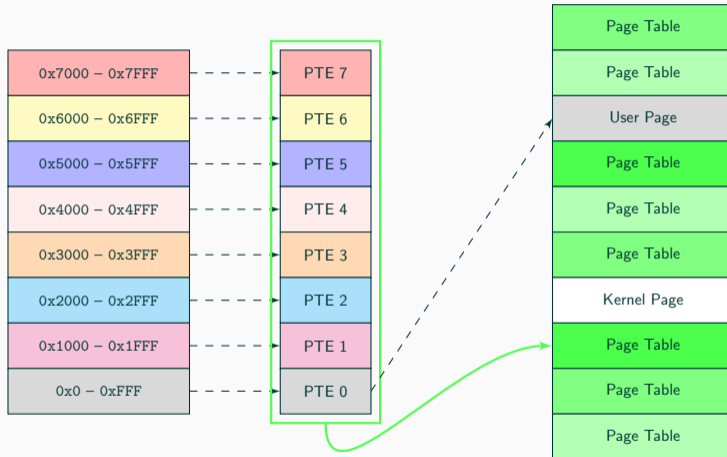
Release bit flip page

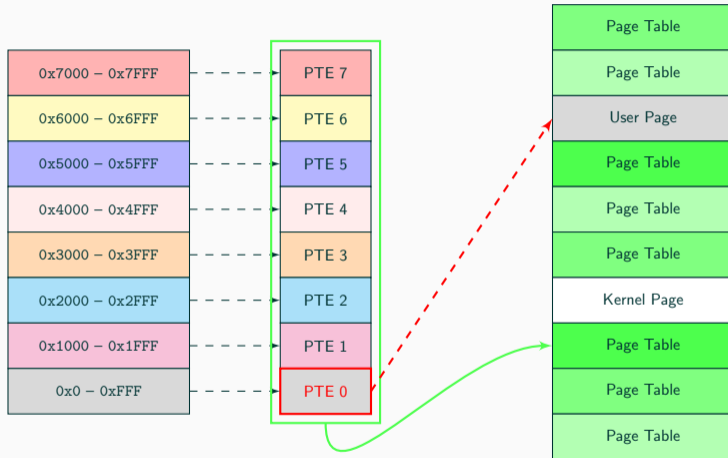


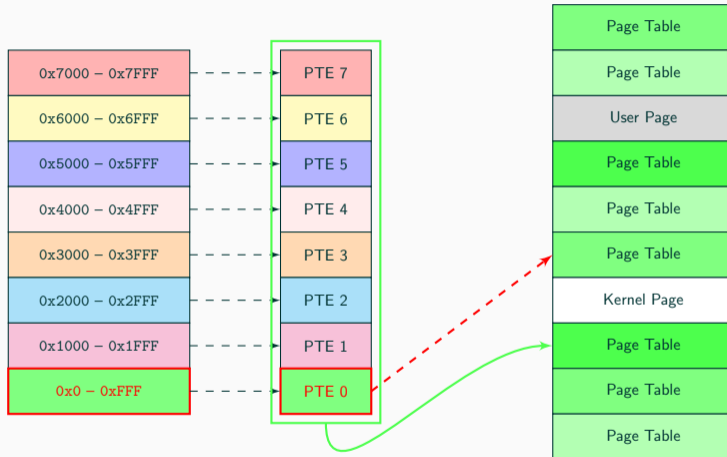
Place page table on bit flip page



Place page table on bit flip page







1. Scan for flips
2. Exhaust or massage memory to place a page table at target location
3. Gain access to your own page table → kernel privileges

- Idea from Seaborn et al. 2015

- Idea from Seaborn et al. 2015
- Same idea applied in several other works:
 - Rowhammer.js (Gruss, Maurice, et al. 2016)
 - One bit flips, one cloud flops (Xiao et al. 2016)
 - Drammer (Veen et al. 2016)

What if we cannot target kernel pages?

- 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`

- Target sudo, as it is the easiest to exploit

- Target sudo, as it is the easiest to exploit
- Use Rowhammer to circumvent password check

- Target sudo, as it is the easiest to exploit
- Use Rowhammer to circumvent password check
- Changing program logic with bit flips

- Target sudo, as it is the easiest to exploit
- Use Rowhammer to circumvent password check
- Changing program logic with bit flips
- What happens if we induce bit flips in instructions?

















- Conditional jumps are not the only targets

- Conditional jumps are not the only targets
- Other targets include

- Conditional jumps are not the only targets
- Other targets include
 - Comparisons
 - Addresses of memory loads/stores
 - Address calculations
 - ...

- Conditional jumps are not the only targets
- Other targets include
 - Comparisons
 - Addresses of memory loads/stores
 - Address calculations
 - ...
- Manual analysis of `sudo` revealed 29 possible bitflips

- Conditional jumps are not the only targets
- Other targets include
 - Comparisons
 - Addresses of memory loads/stores
 - Address calculations
 - ...
- Manual analysis of `sudo` revealed 29 possible bitflips
- They all somehow skipped the password check

- Remaining problem: maneuver the target binary page to a vulnerable physical page

- Remaining problem: maneuver the target binary page to a vulnerable physical page
- Not as easy as with page tables

- Remaining problem: maneuver the target binary page to a vulnerable physical page
- Not as easy as with page tables
- Many page tables can simply be sprayed over memory

- Remaining problem: maneuver the target binary page to a vulnerable physical page
- Not as easy as with page tables
- Many page tables can simply be sprayed over memory
- There is only one copy of the binary in memory

- 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

- We propose **memory waylaying** to move a victim page to a target page

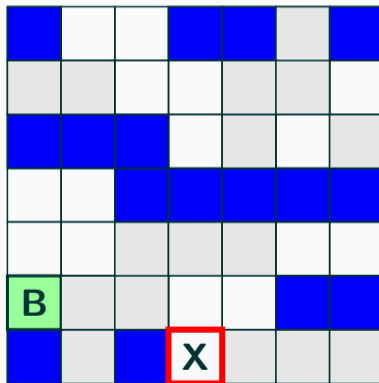
- We propose **memory waylaying** to move a victim page to a target page
- By filling the page cache with executable pages, we evict the victim binary

- We propose **memory waylaying** to move a victim page to a target page
- By filling the page cache with executable pages, we evict the victim binary
- Using the `mincore` function, we can check if the victim binary is evicted


- We propose **memory waylaying** to move a victim page to a target page
- By filling the page cache with executable pages, we evict the victim binary
- Using the `mincore` function, we can check if the victim binary is evicted
- Loading the victim binary results in a new physical page

- We propose **memory waylaying** to move a victim page to a target page
- By filling the page cache with executable pages, we evict the victim binary
- Using the `mincore` function, we can check if the victim binary is evicted
- Loading the victim binary results in a new physical page
- Continue until it is at the target page


(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




- Modify binary pages executed in root privileges (Xiao et al. 2016)
- Modify credential structs (Veen et al. 2016)
- Read keys (Xiao et al. 2016)
- Corrupt RSA signatures (Bhattacharya et al. 2016)
- Modify certificates
- Configurations
- etc.

Rowhammer + SGX = Cheap Denial of Service

- SGX is an x86 instruction-set extension
- Integrity and Confidentiality of code and data in untrusted environments
- Run with user privileges and restricted, e.g., no system calls
- Execute programs in enclaves using protected areas of memory





- SGX explicitly protects against DRAM-based attacks
 - Cold-boot attacks
 - Memory bus snooping
 - Memory tampering attacks
- SGX EPC (Enclave Page Cache)
 - Region in DRAM
 - Cryptographically ensuring confidentiality, integrity and freshness of data
- Memory Encryption Engine (MEE)
 - Transparently encrypts memory

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

Sounds unsafe?

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

Sounds unsafe? **It is unsafe!**

- More and more services are deployed using SGX

- More and more services are deployed using SGX
- Cloud providers offer SGX support
 - Microsoft's Azure Confidential Computing

- More and more services are deployed using SGX
- Cloud providers offer SGX support
 - Microsoft's Azure Confidential Computing
- If a malicious enclave induces a bit flip, ...

- More and more services are deployed using SGX
- Cloud providers offer SGX support
 - Microsoft's Azure Confidential Computing
- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts

- More and more services are deployed using SGX
- Cloud providers offer SGX support
 - Microsoft's Azure Confidential Computing
- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts
- ... including co-located tenants

- More and more services are deployed using SGX
- Cloud providers offer SGX support
 - Microsoft's Azure Confidential Computing
- If a malicious enclave induces a bit flip, ...
- ... the entire machine halts
- ... including co-located tenants
- Denial-of-Service Attacks in the Cloud

**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:
 - MASCAT
 - ANVIL
 - HexPADS
 - Herath and Fogh
 - Gruss et al.
 - Zhang et al.
 - Chiappetta et al.

- Some countermeasures assume alternate access to at least two rows
- One-location hammering defeats countermeasures relying on this:
 - ANVIL
 - Corbet

- Some countermeasures assume that memory massaging requires a lot of memory
- Waylaying defeats countermeasures relying on this:
 - Gruss et al.
 - Van der Veen et al.

- Some countermeasures assume kernel data structure have to be targeted
- Opcode flipping defeats countermeasures relying on this:
 - G-CATT

| Defense Class | | Static Analysis | Performance Counters | Memory Access Pattern | Physical Proximity | Memory footprint |
|-------------------------------|--|-----------------|----------------------|-----------------------|--------------------|------------------|
| | | | | | | |
| Intel SGX | | ● | ● | ○ | ○ | ○ |
| One-location hammering | | ○ | ○ | ● | ○ | ○ |
| Opcode flipping | | ○ | ○ | ○ | ● | ○ |
| Memory waylaying | | ○ | ○ | ○ | ○ | ● |
| Defense class defeated | | ● | ● | ● | ● | ● |

- Many (academic) countermeasures were proposed to mitigate Rowhammer

- Many (academic) countermeasures were proposed to mitigate Rowhammer
- We showed that all of them can be circumvented (Gruss, Lipp, et al. 2018)

- Many (academic) countermeasures were proposed to mitigate Rowhammer
- We showed that all of them can be circumvented (Gruss, Lipp, et al. 2018)
- We cannot design countermeasures without completely understanding the attack

- Many (academic) countermeasures were proposed to mitigate Rowhammer
- We showed that all of them can be circumvented (Gruss, Lipp, et al. 2018)
- We cannot design countermeasures without completely understanding the attack
- Otherwise we only patch concrete exploits, but do not solve the problem

- We have to invest more into researching attacks

- We have to invest more into researching attacks
- There are still aspects of Rowhammer we do not fully understand

- We have to invest more into researching attacks
- There are still aspects of Rowhammer we do not fully understand
- However, this is required to design effective countermeasures

- We have to invest more into researching attacks
- There are still aspects of Rowhammer we do not fully understand
- However, this is required to design effective countermeasures
- Moreover, new features might introduce new attack vectors (e.g., SGX)

- We underestimated side-channel attacks for a long time

- We underestimated side-channel attacks for a long time
- Industry and customers have to reconsider priorities → focus more on security instead of performance

- We underestimated side-channel attacks for a long time
- Industry and customers have to reconsider priorities → focus more on security instead of performance
- Reliability issues (Rowhammer) can have security impacts

- We underestimated side-channel attacks for a long time
- Industry and customers have to reconsider priorities → focus more on security instead of performance
- Reliability issues (Rowhammer) can have security impacts
- More research is required to understand attacks to ultimately mitigate them

- Rowhammer attacks have a huge attack potential
- One-location hammering showed us that previous assumptions on how to trigger the Rowhammer bug are invalid
 - Keeping only one DRAM row open is sufficient
- Op-code flipping (in the sudo binary) allows to obtain root privileges
- Proposed countermeasures in software are ineffective

Rowhammer: From the Basics to Sophisticated New Variants

Moritz Lipp, Michael Schwarz

February 20, 2018

Graz University of Technology