

# Malware Guard Extension: Using SGX to Conceal Cache Attacks

**Michael Schwarz, Samuel Weiser, Daniel Gruss, Clémentine Maurice, Stefan Mangard**

July 6, 2017

Graz University of Technology

- We show that **malicious** SGX enclaves are possible

- We show that **malicious** SGX enclaves are possible
- We present methods to circumvent certain **limitations** of SGX

- We show that **malicious** SGX enclaves are possible
- We present methods to circumvent certain **limitations** of SGX
- We mount a cache attack on a RSA implementation within a **different** enclave

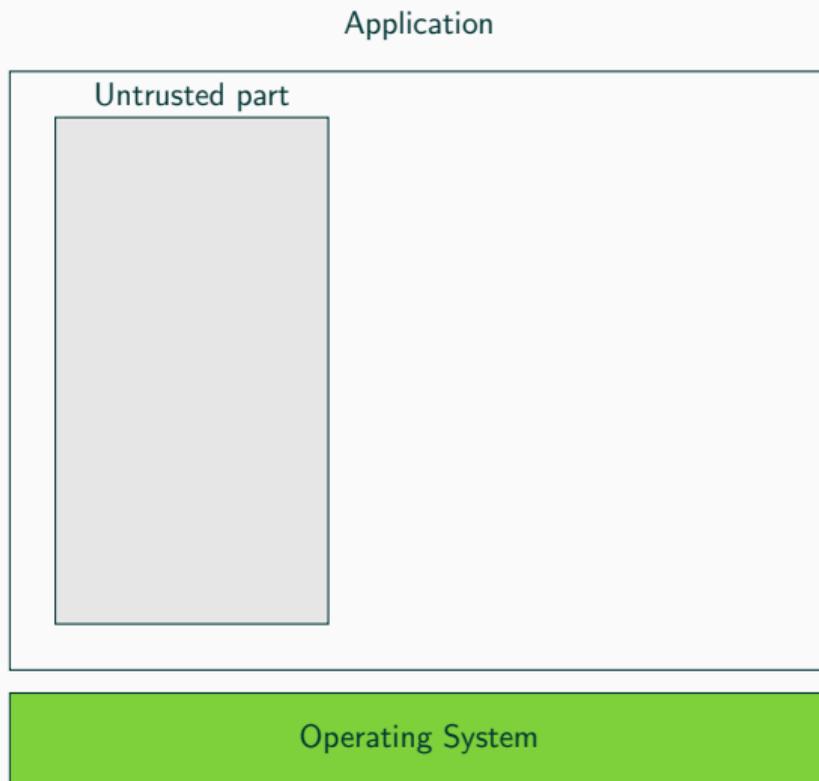
- We show that **malicious** SGX enclaves are possible
- We present methods to circumvent certain **limitations** of SGX
- We mount a cache attack on a RSA implementation within a **different** enclave
- We abuse SGX to **prevent detection** of the attack

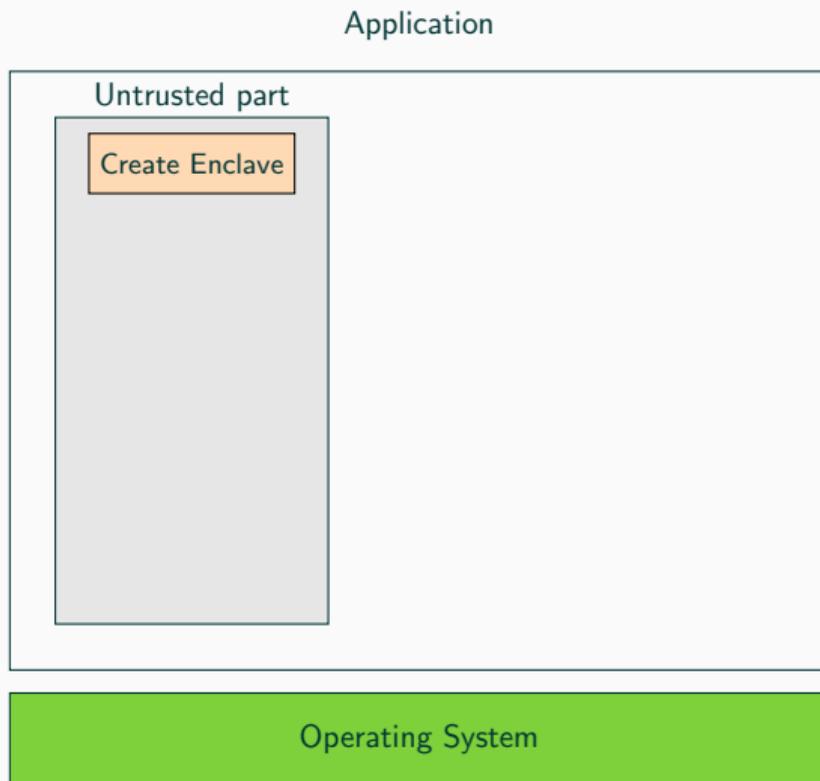
- We show that **malicious** SGX enclaves are possible
- We present methods to circumvent certain **limitations** of SGX
- We mount a cache attack on a RSA implementation within a **different** enclave
- We abuse SGX to **prevent detection** of the attack
- We show that the stealthy attack even works across **Docker** containers

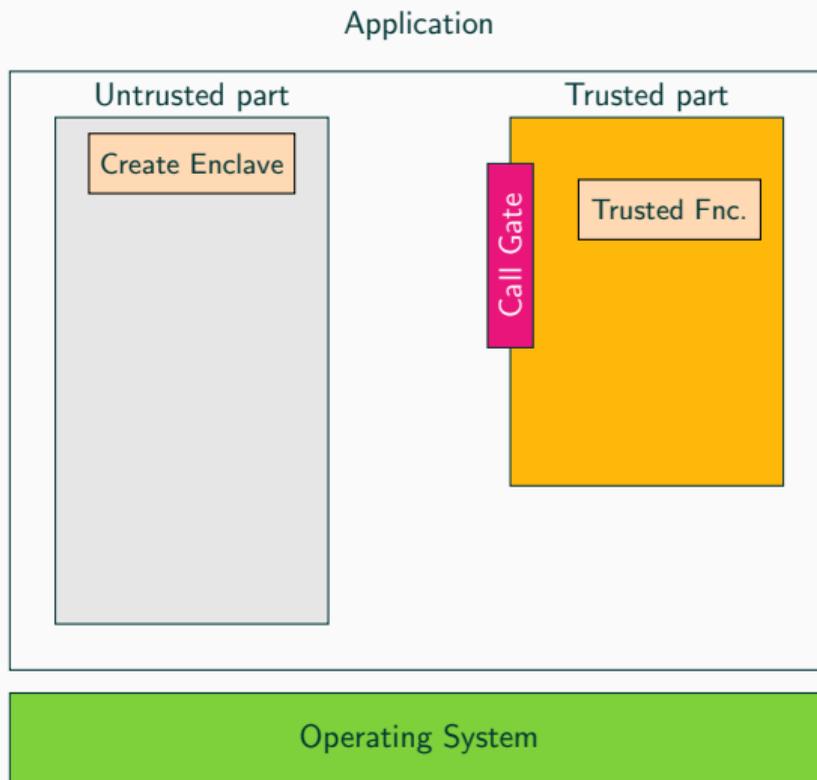
- We show that **malicious** SGX enclaves are possible
- We present methods to circumvent certain **limitations** of SGX
- We mount a cache attack on a RSA implementation within a **different** enclave
- We abuse SGX to **prevent detection** of the attack
- We show that the stealthy attack even works across **Docker** containers
- We discuss **countermeasures** to prevent such attacks

# Background

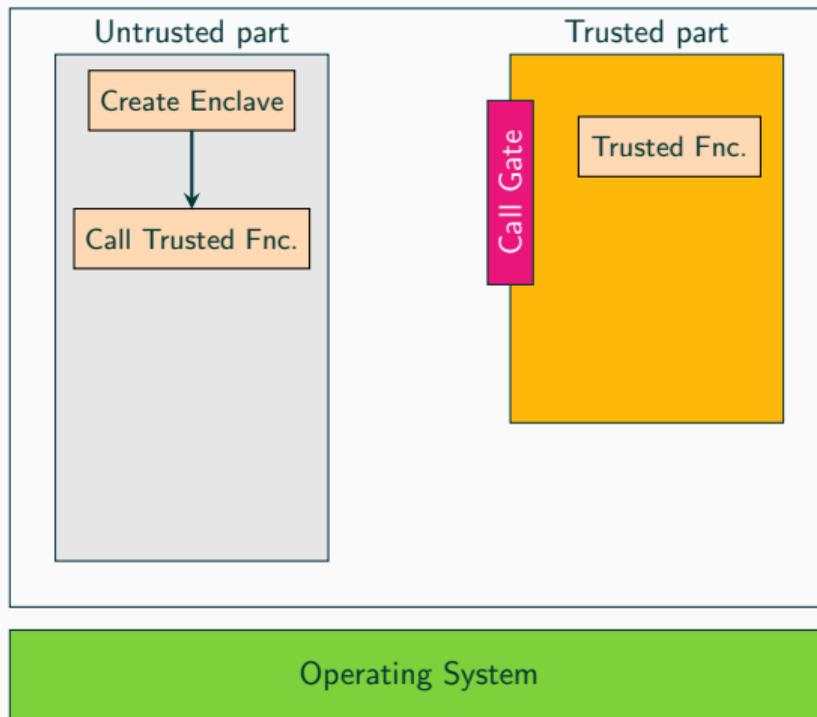
---

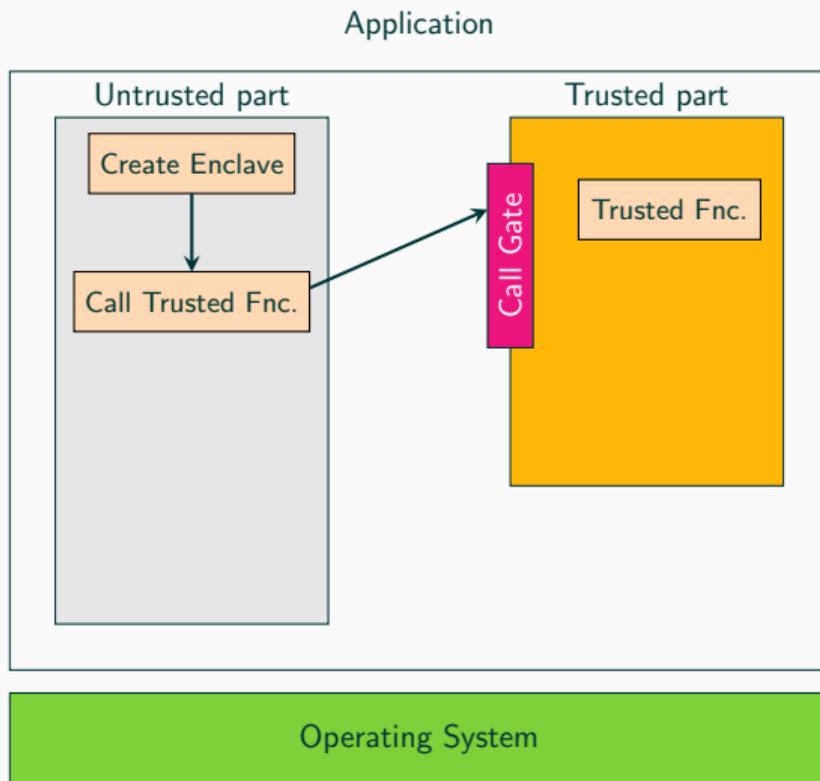


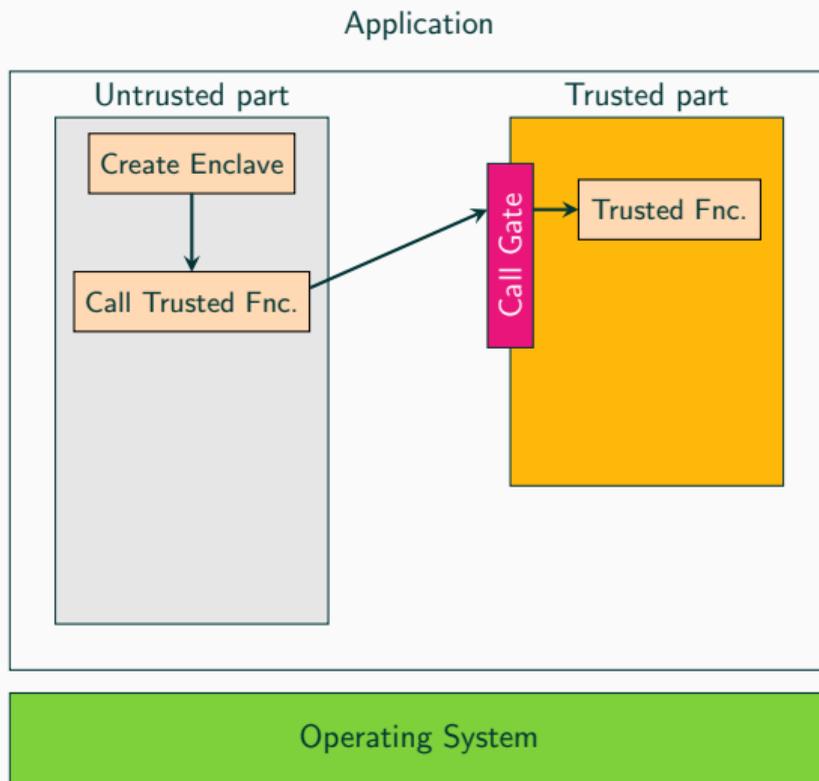


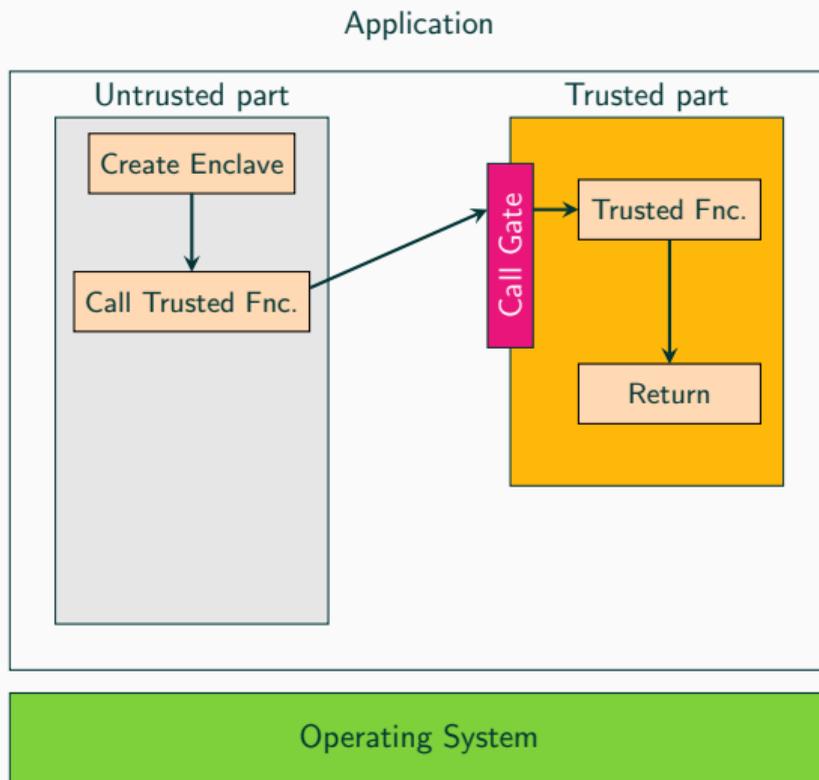


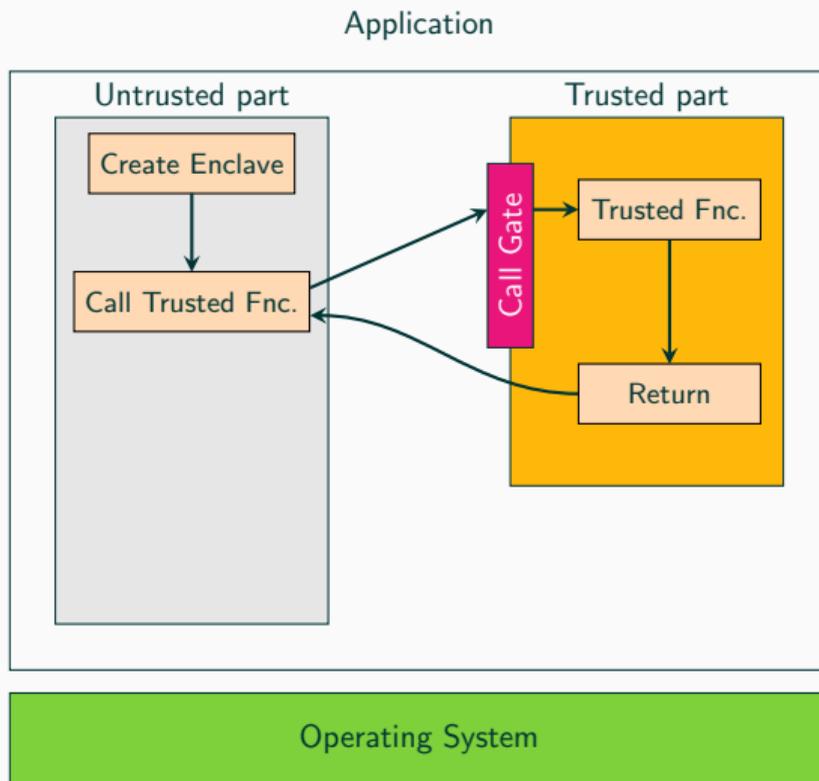
## Application

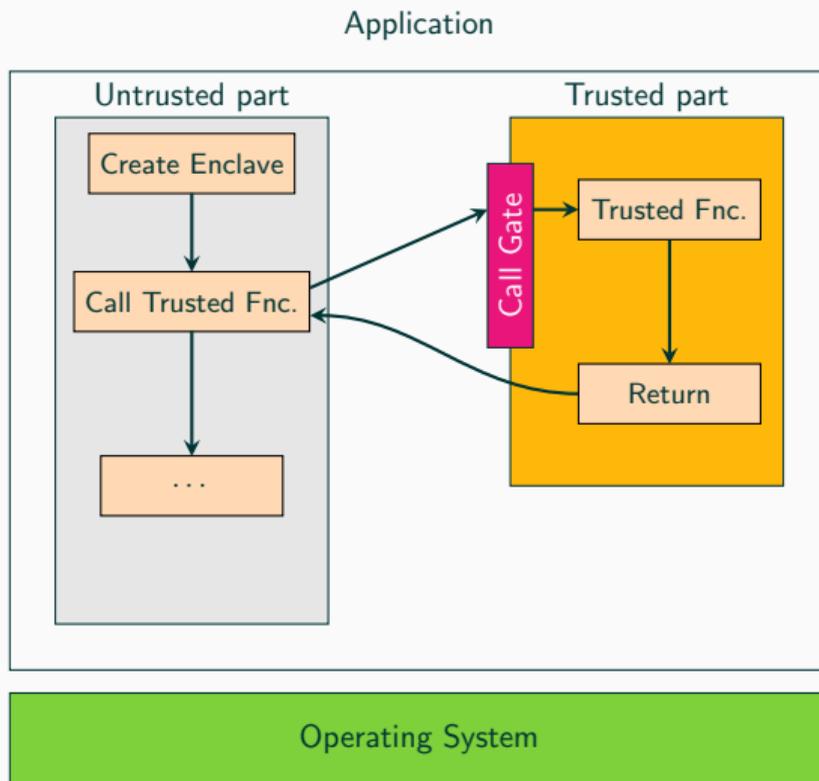


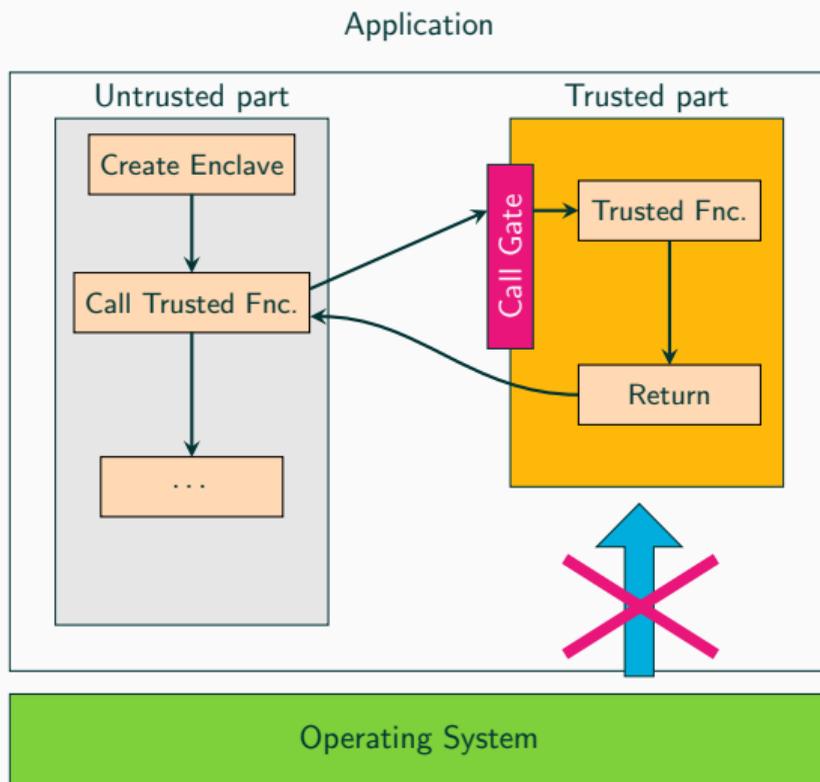












$$M = C^d \bmod n$$

$$M = C^d \pmod n$$

1 1 0 0 1 1 0 ...

Result = C

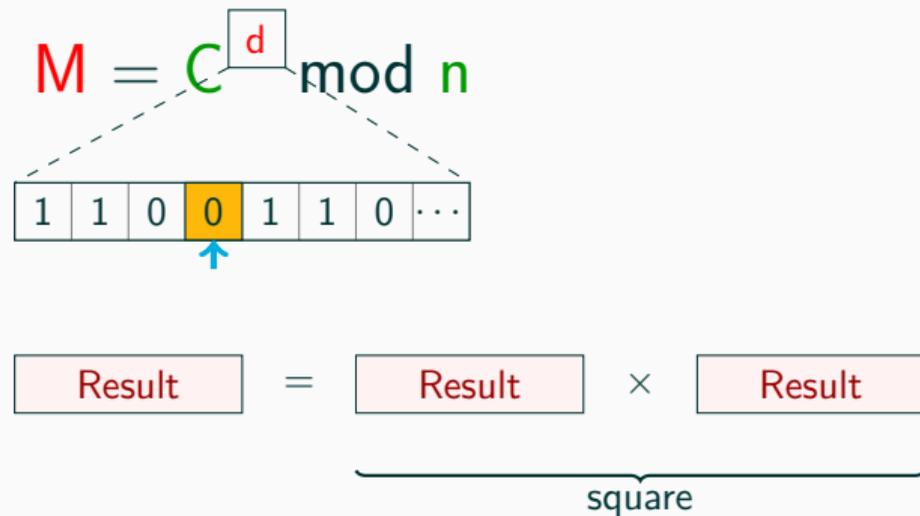
$$M = C^d \pmod n$$

1 1 0 0 1 1 0 ...

$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}} \times \underbrace{C}_{\text{multiply}}$$

$$M = C^d \bmod n$$

$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}}$$



$$M = C^d \pmod n$$

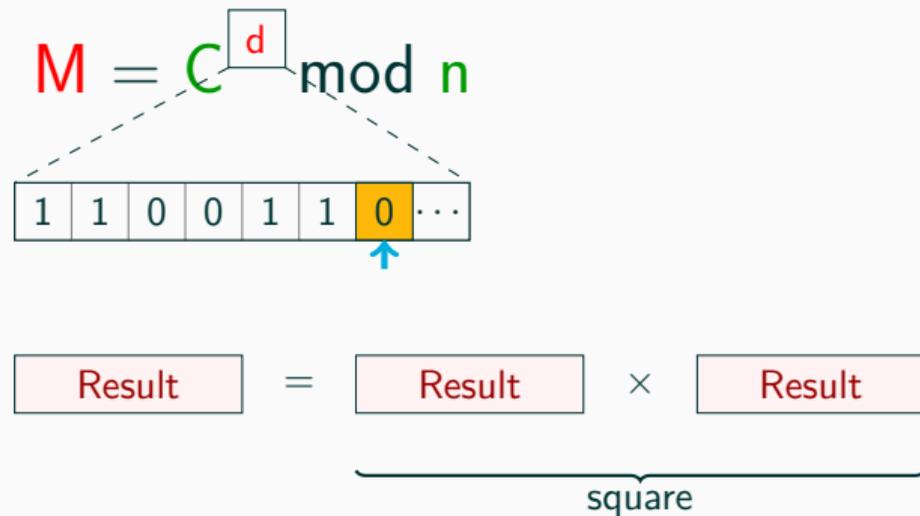
1 1 0 0 1 1 0 ...

$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}} \times \underbrace{C}_{\text{multiply}}$$

$$M = C^d \pmod n$$

1 1 0 0 1 1 0 ...

$$\text{Result} = \underbrace{\text{Result} \times \text{Result}}_{\text{square}} \times \underbrace{C}_{\text{multiply}}$$



Prime+Probe...

## Prime+Probe...

- exploits the **timing difference** when accessing...

## Prime+Probe...

- exploits the **timing difference** when accessing...
  - cached data (fast)

## Prime+Probe...

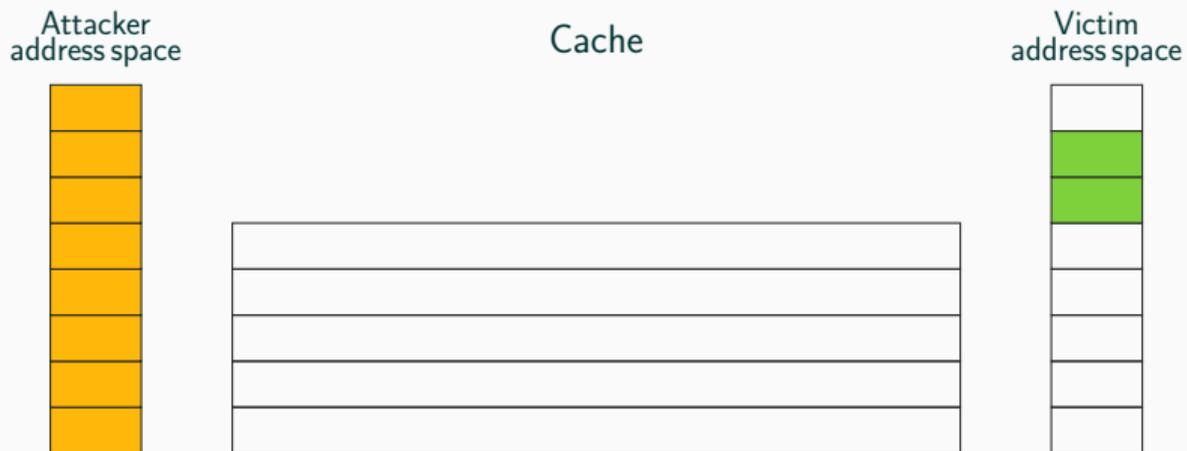
- exploits the **timing difference** when accessing...
  - cached data (fast)
  - uncached data (slow)

## Prime+Probe...

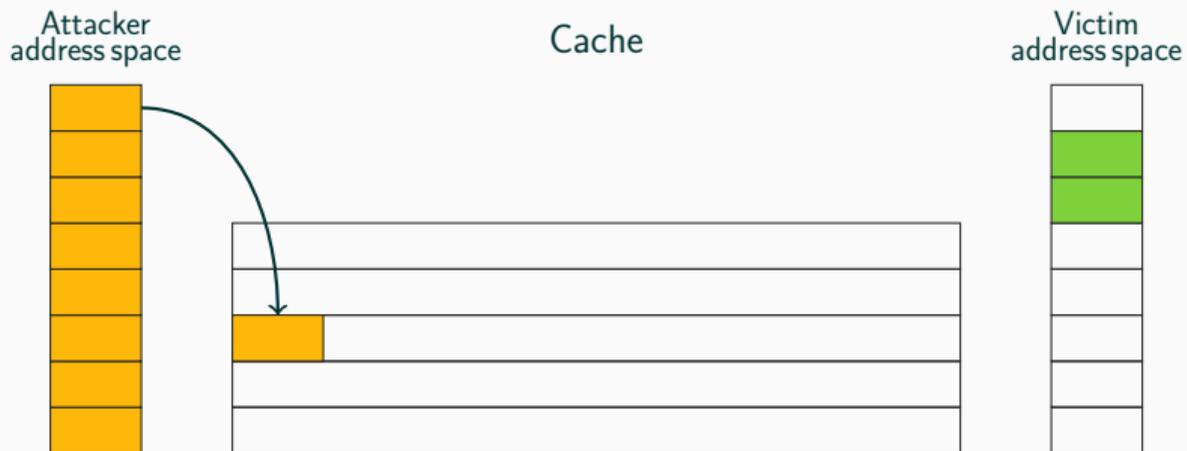
- exploits the **timing difference** when accessing...
  - cached data (fast)
  - uncached data (slow)
- is applied to one cache set

## Prime+Probe...

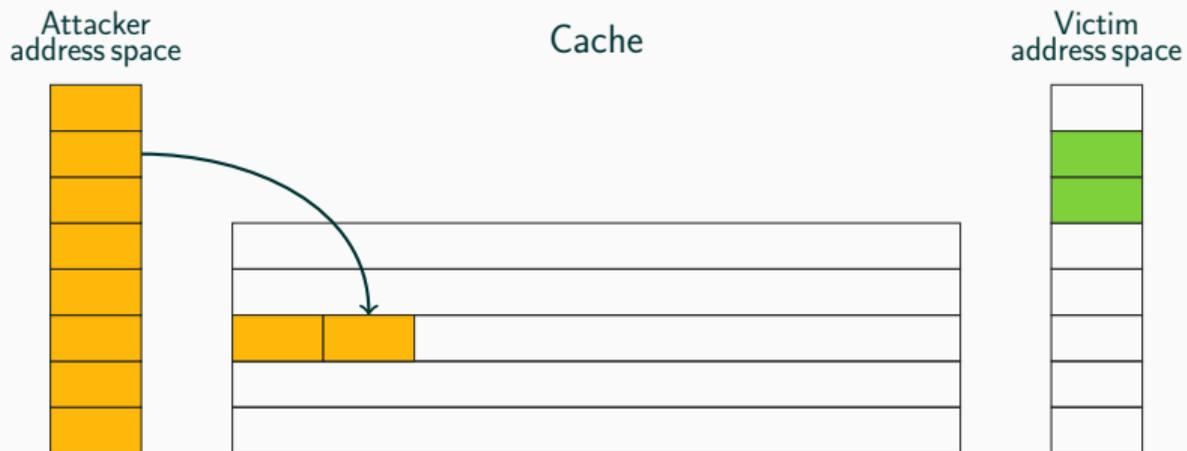
- exploits the **timing difference** when accessing...
  - cached data (fast)
  - uncached data (slow)
- is applied to one cache set
- works **across CPU cores** as the last-level cache is shared



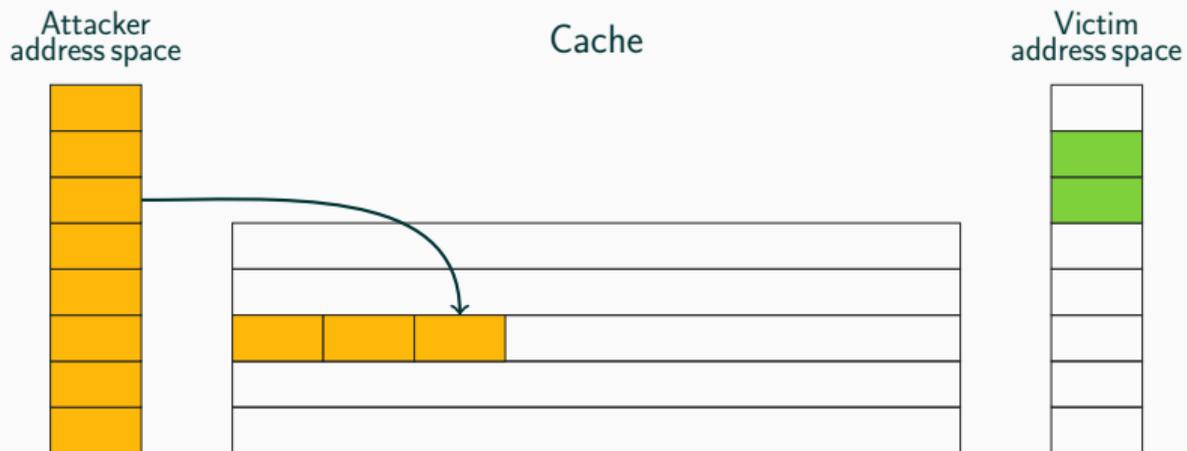
**Step 0:** Attacker fills the cache (prime)



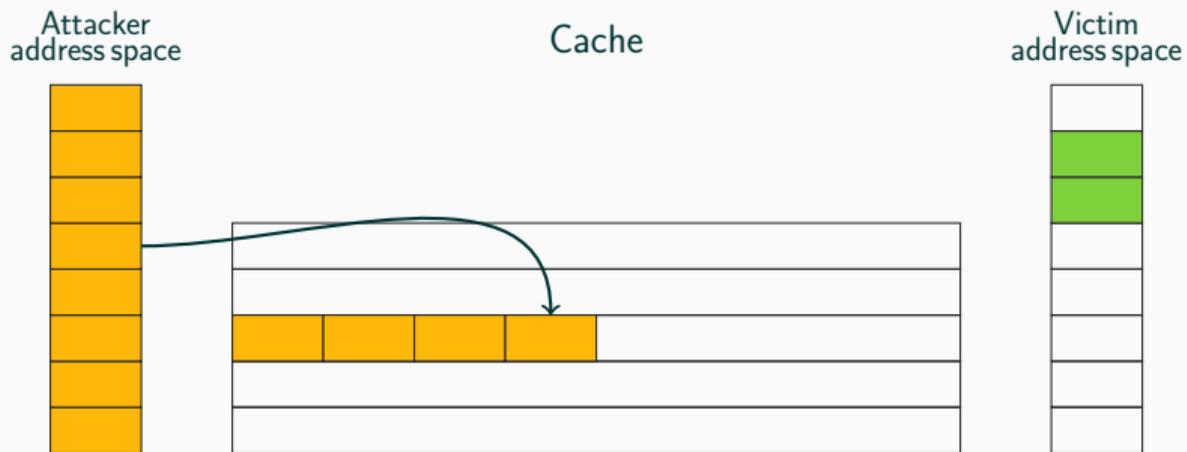
**Step 0:** Attacker fills the cache (prime)



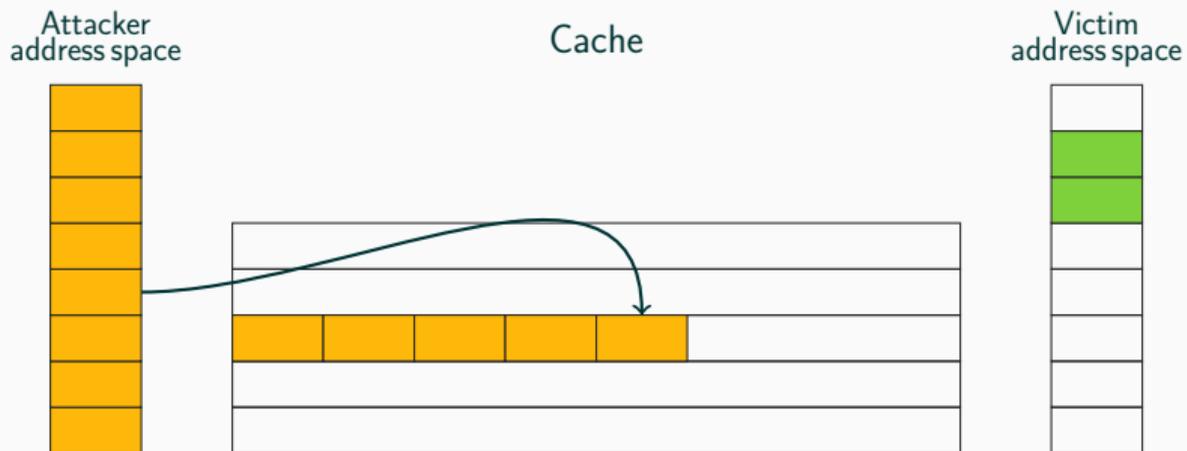
**Step 0:** Attacker fills the cache (prime)



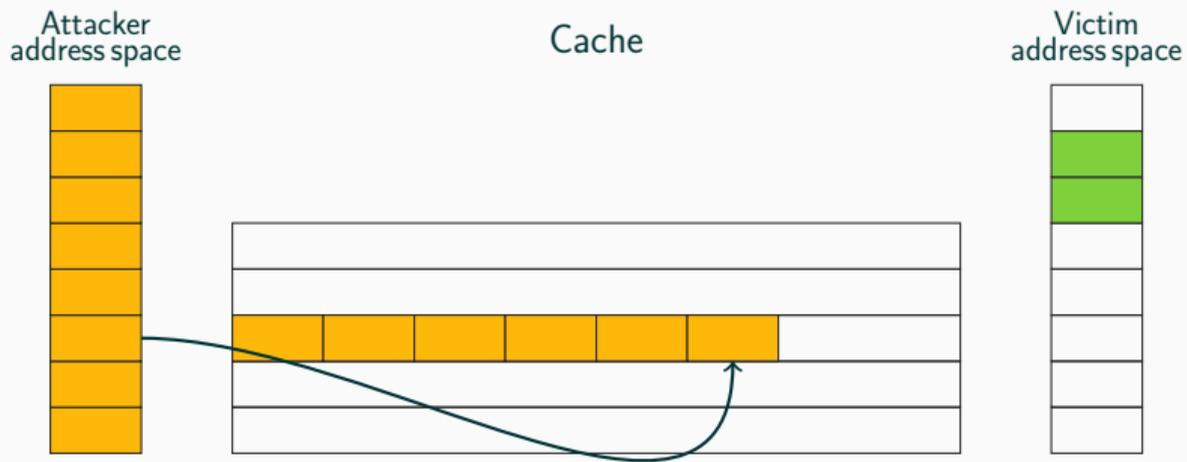
**Step 0:** Attacker fills the cache (prime)



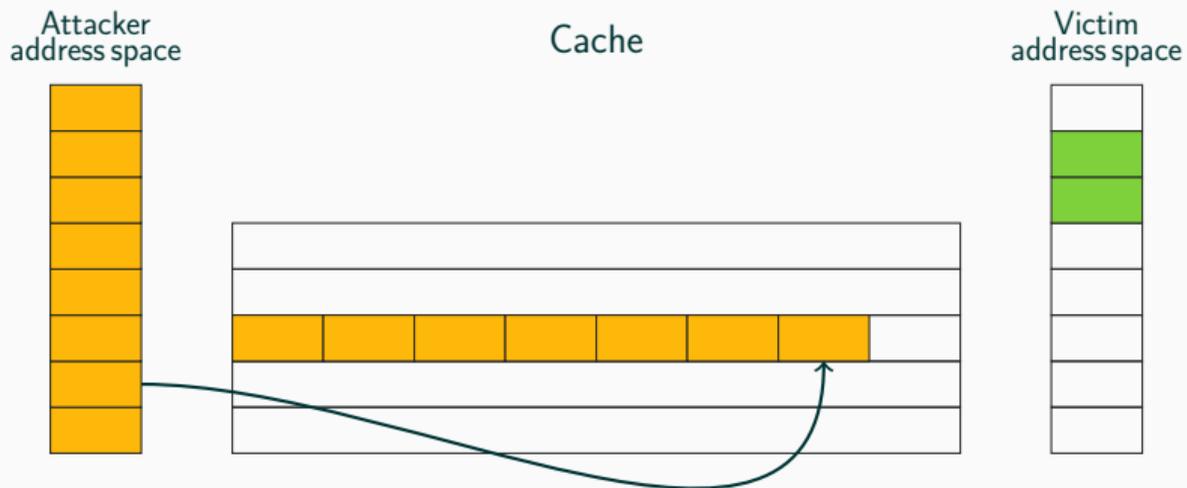
**Step 0:** Attacker fills the cache (prime)



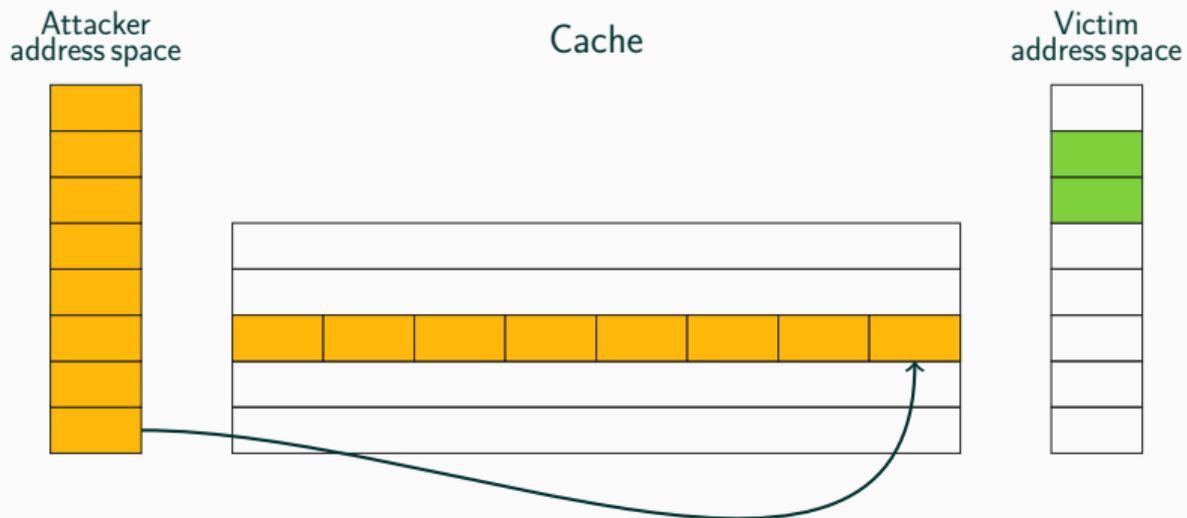
**Step 0:** Attacker fills the cache (prime)



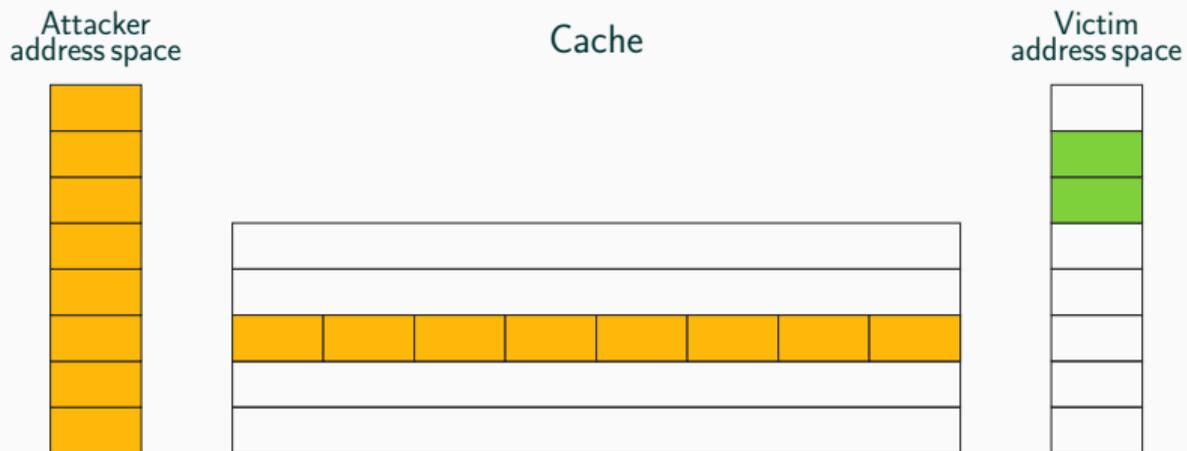
**Step 0:** Attacker fills the cache (prime)



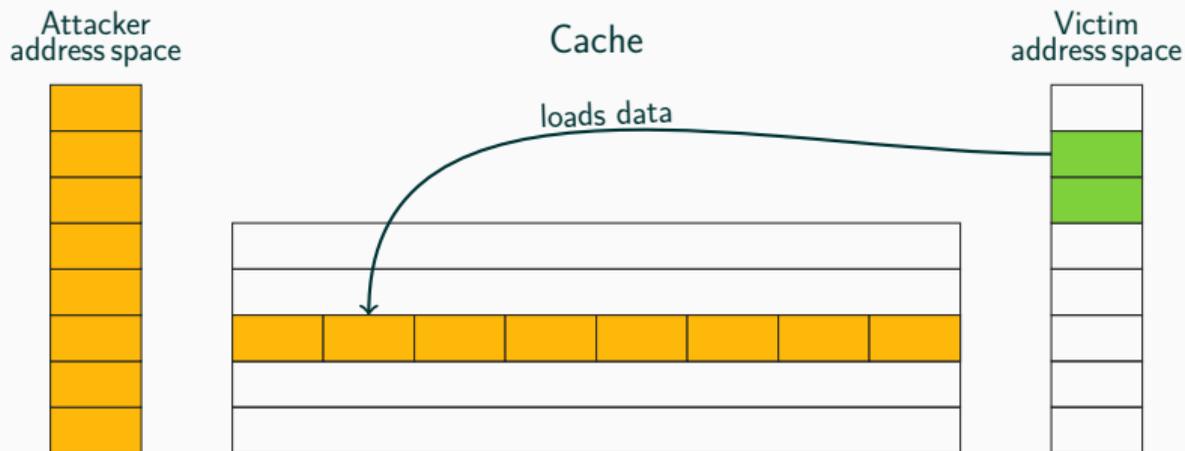
**Step 0:** Attacker fills the cache (prime)



**Step 0:** Attacker fills the cache (prime)

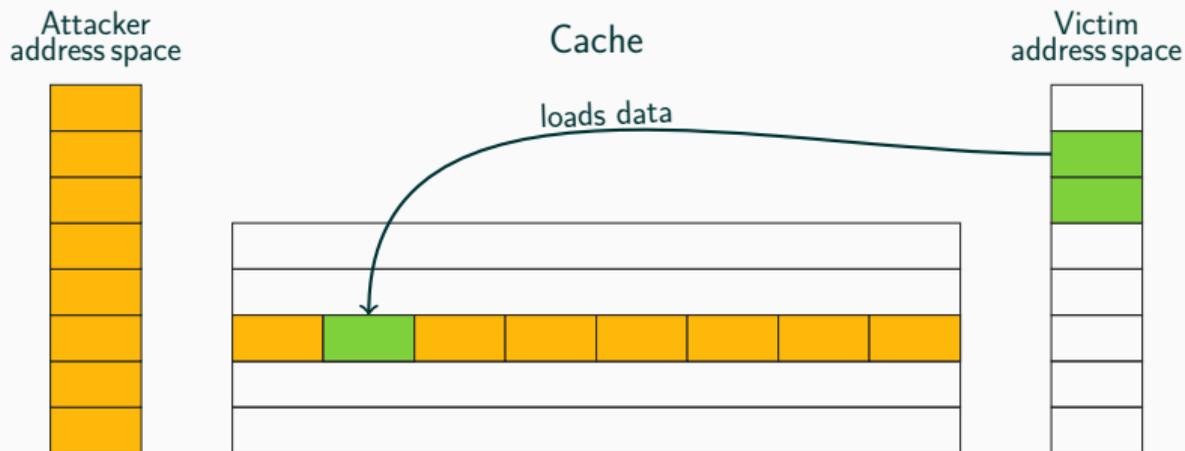


**Step 0:** Attacker fills the cache (prime)



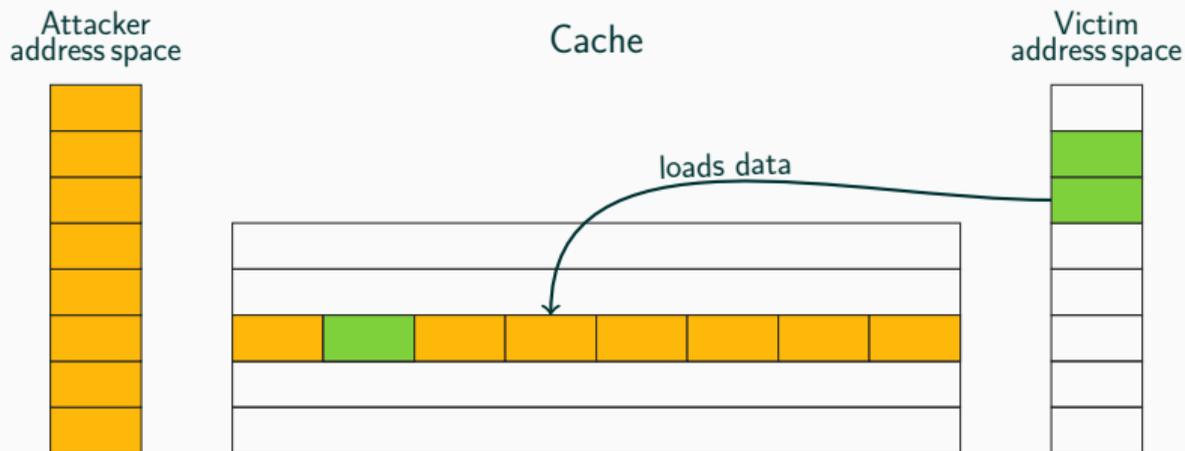
**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data



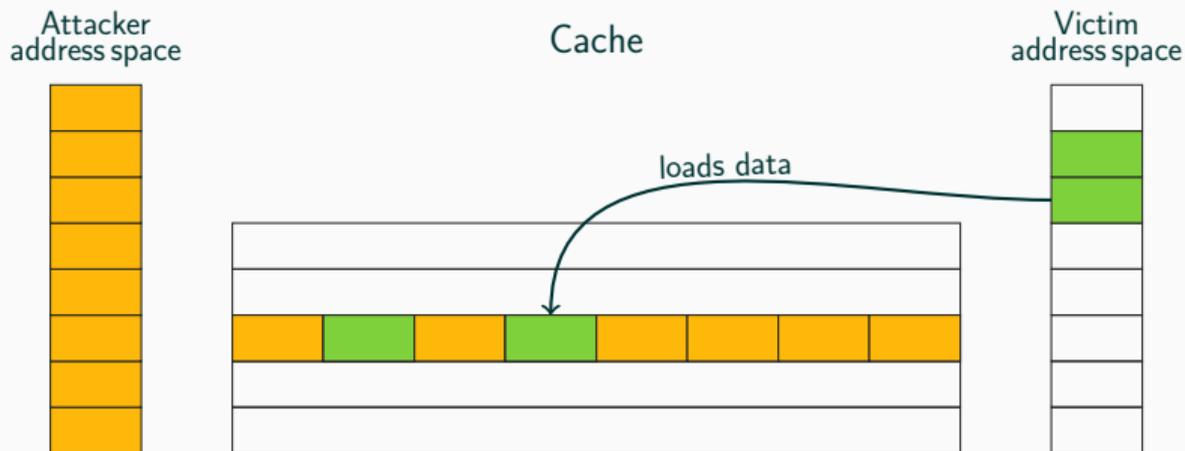
**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data



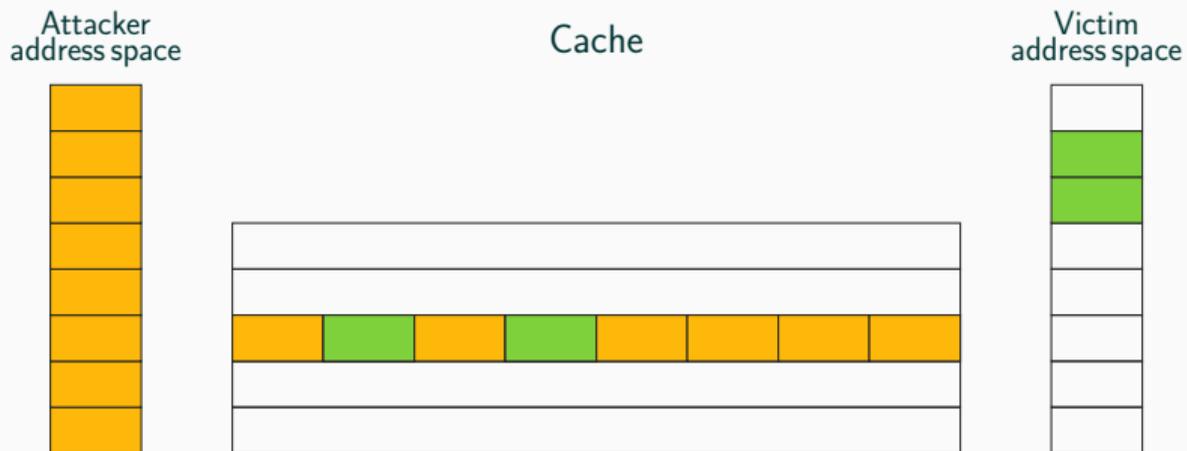
**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data



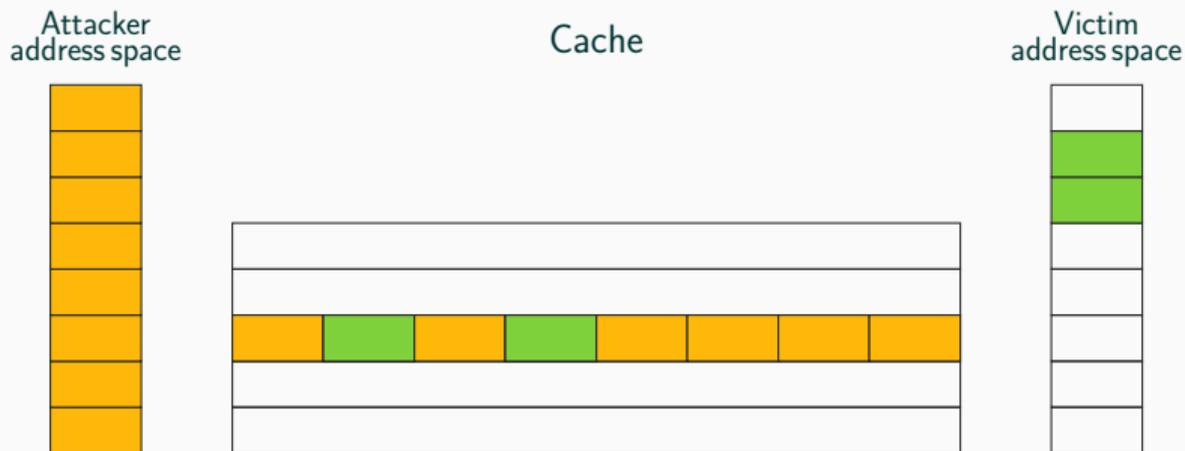
**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data



**Step 0:** Attacker fills the cache (prime)

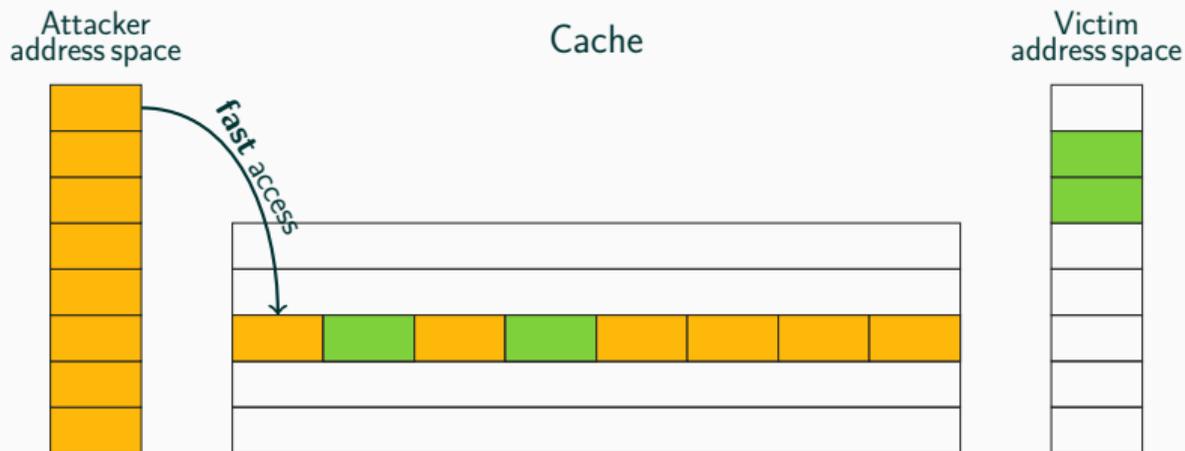
**Step 1:** Victim evicts cache lines by accessing own data



**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data

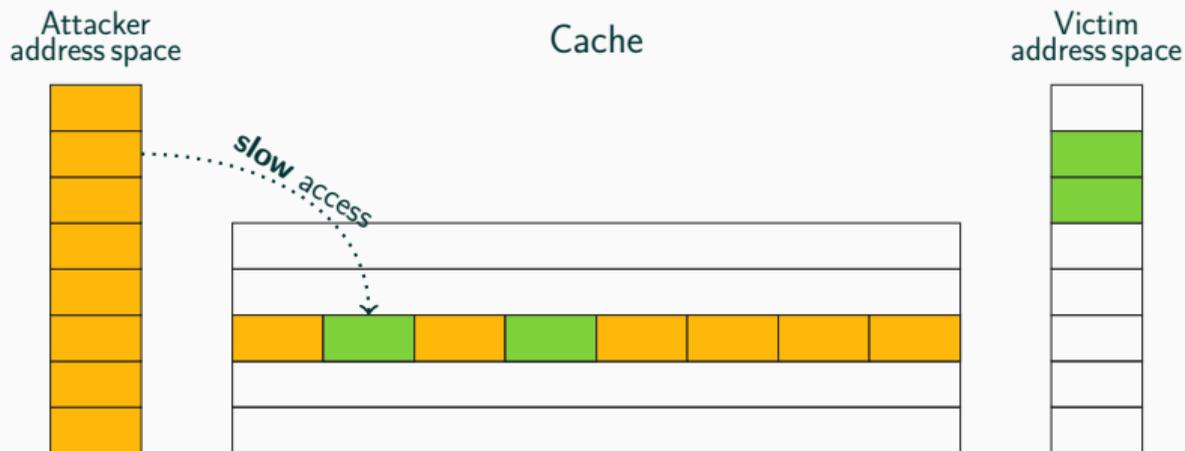
**Step 2:** Attacker probes data to determine if the set was accessed



**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data

**Step 2:** Attacker probes data to determine if the set was accessed



**Step 0:** Attacker fills the cache (prime)

**Step 1:** Victim evicts cache lines by accessing own data

**Step 2:** Attacker probes data to determine if the set was accessed

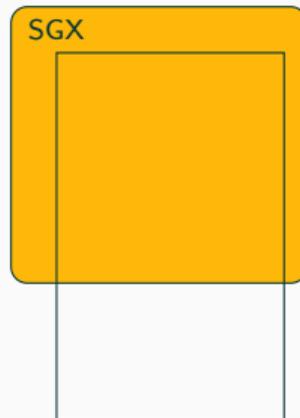
# Attack

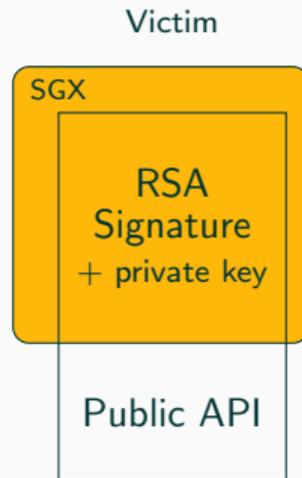
---

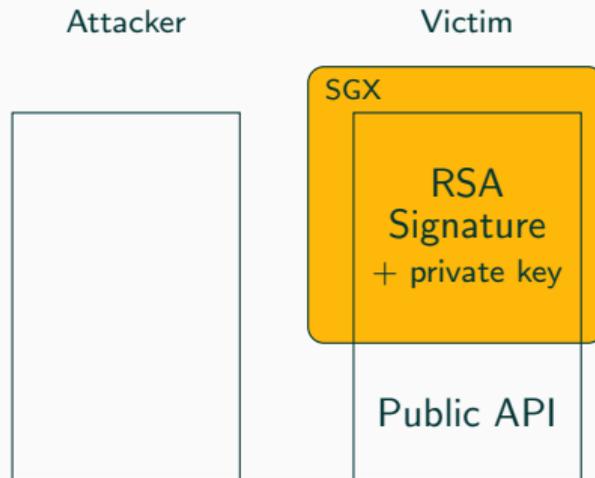
Victim

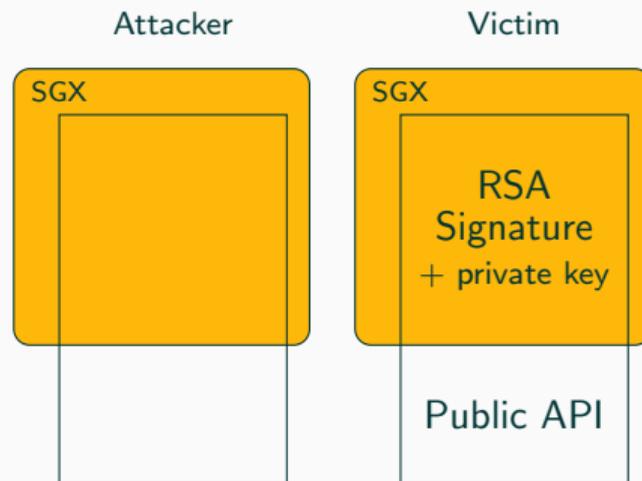


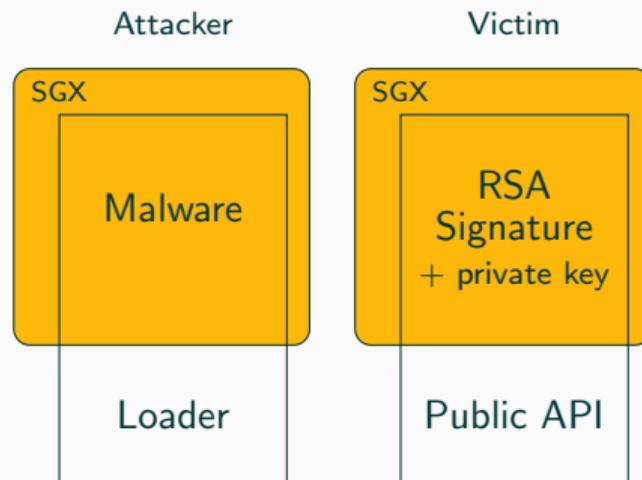
Victim

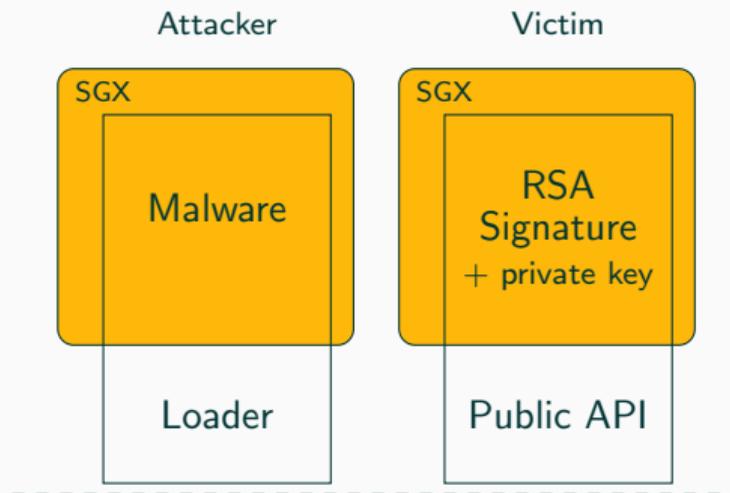


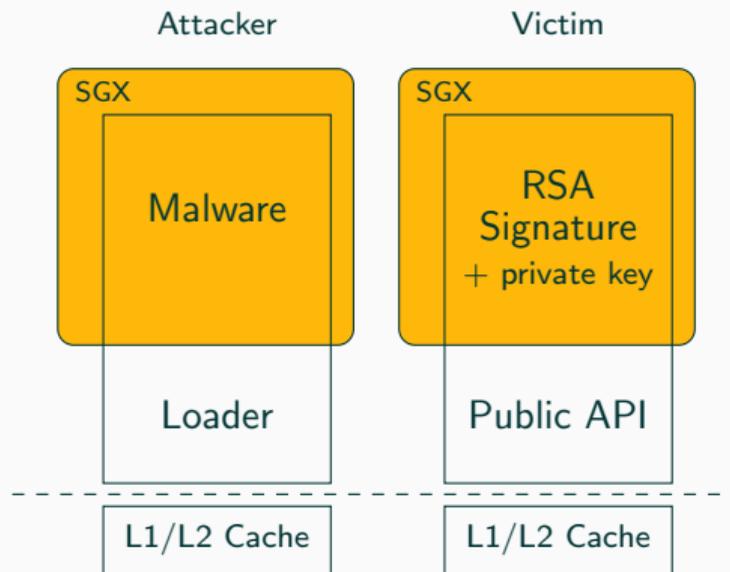


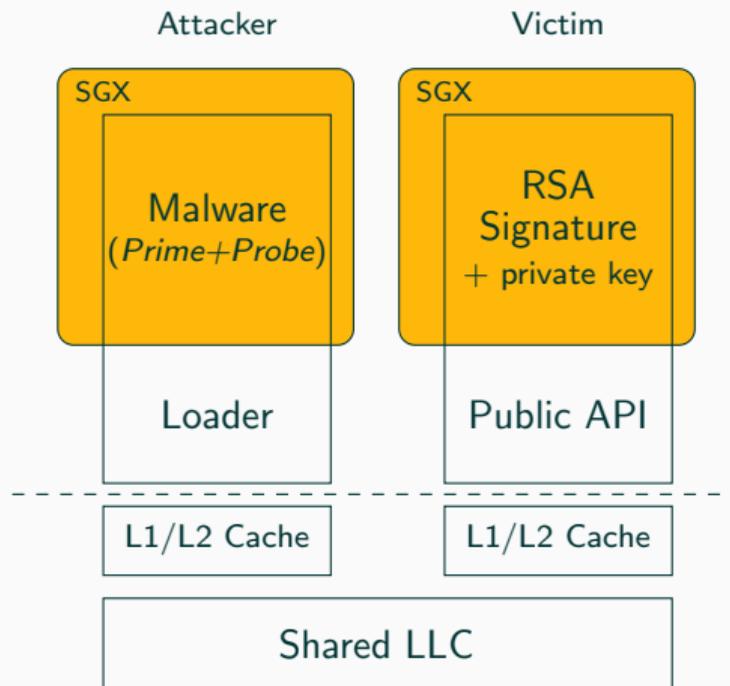














- No access to **high-precision timer** (`rdtsc`)



- No access to **high-precision timer** (`rdtsc`)
- No **syscalls**



- No access to **high-precision timer** (`rdtsc`)
- No **syscalls**
- No **shared memory**



- No access to **high-precision timer** (`rdtsc`)
- No **syscalls**
- No **shared memory**
- No **physical addresses**



- No access to **high-precision timer** (`rdtsc`)
- No **syscalls**
- No **shared memory**
- No **physical addresses**
- No 2 MB **large pages**



- We have to build our own timer



- We have to build our **own timer**
- Timer resolution must be in the order of cycles



- We have to build our **own timer**
- Timer resolution must be in the order of cycles
- Start a thread that continuously increments a global variable



- We have to build our **own timer**
- Timer resolution must be in the order of cycles
- Start a thread that continuously increments a global variable
- The global variable is our **timestamp**

CPU cycles one increment takes

```
rdtsc 1
```

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

CPU cycles one increment takes

rdtsc 1

C 4.7

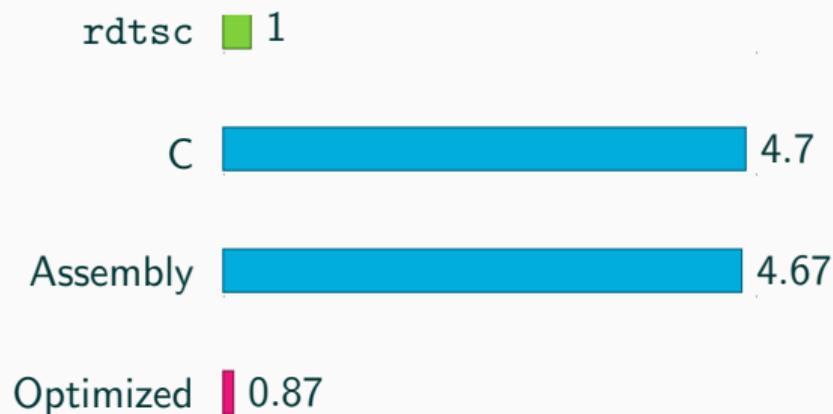
```
1 while(1) {  
2   timestamp++;  
3 }
```

CPU cycles one increment takes



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

CPU cycles one increment takes



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



- **Cache set** is determined by part of physical address



- **Cache set** is determined by part of physical address
- We have no knowledge of **physical addresses**



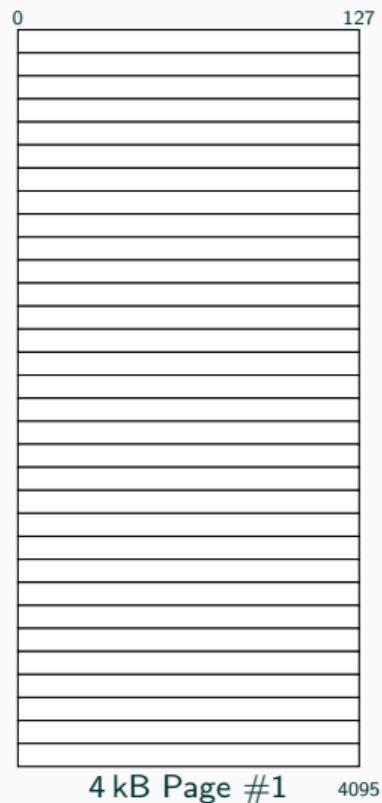
- **Cache set** is determined by part of physical address
- We have no knowledge of **physical addresses**
- Use the **DRAM mapping** reverse engineered by Pessl et al.



- **Cache set** is determined by part of physical address
- We have no knowledge of **physical addresses**
- Use the **DRAM mapping** reverse engineered by Pessl et al.
- Exploit timing differences to find DRAM **row borders**



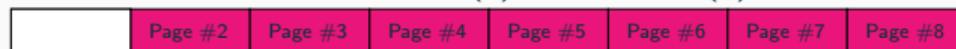
- **Cache set** is determined by part of physical address
- We have no knowledge of **physical addresses**
- Use the **DRAM mapping** reverse engineered by Pessl et al.
- Exploit timing differences to find DRAM **row borders**
- The 18 LSBs are '0' at a row border



8 kB row  $x$  in BG0 (1) and channel (1)



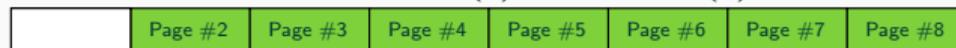
8 kB row  $x$  in BG0 (0) and channel (1)

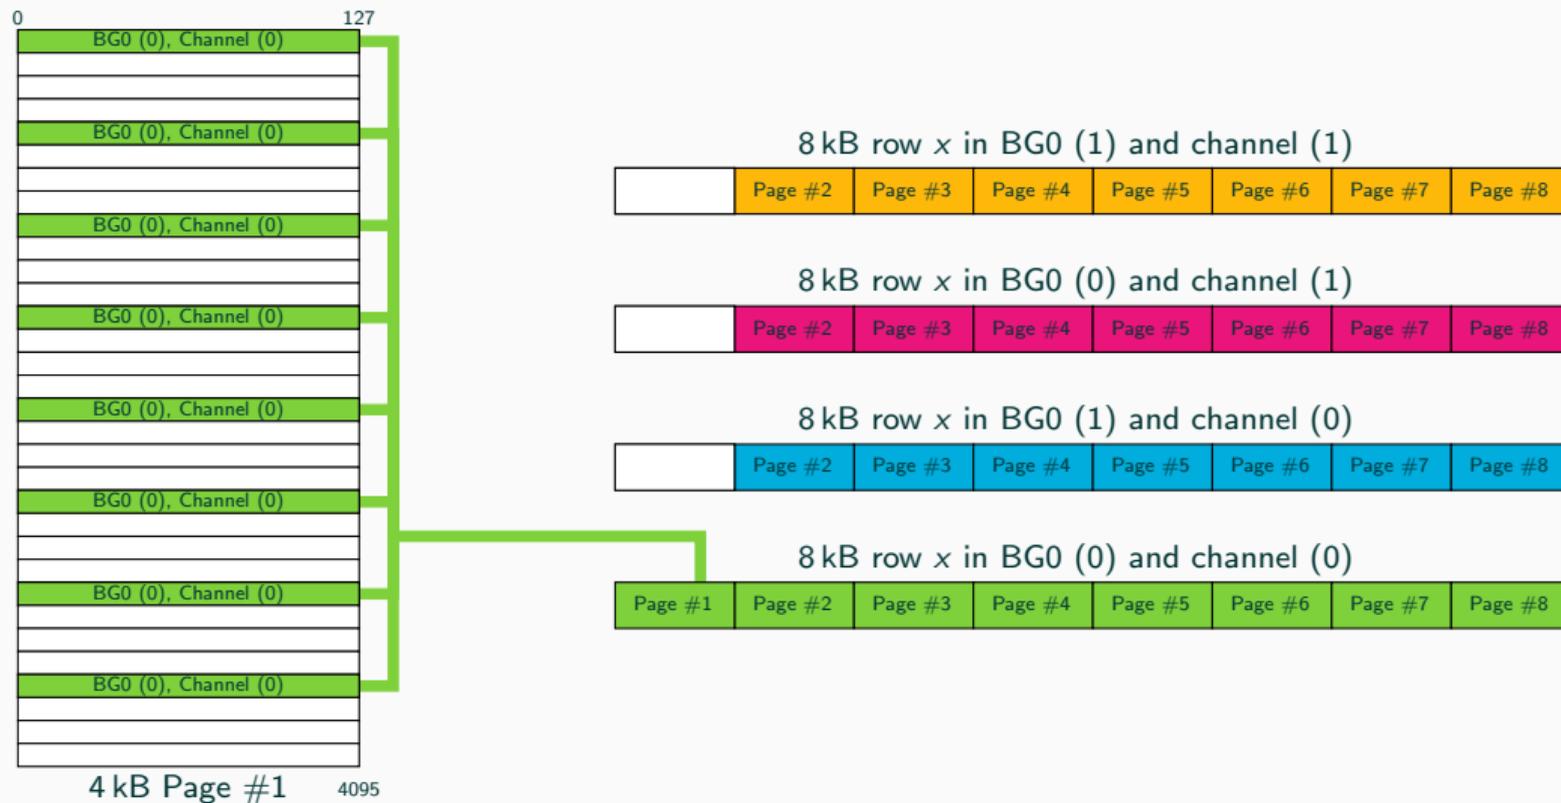


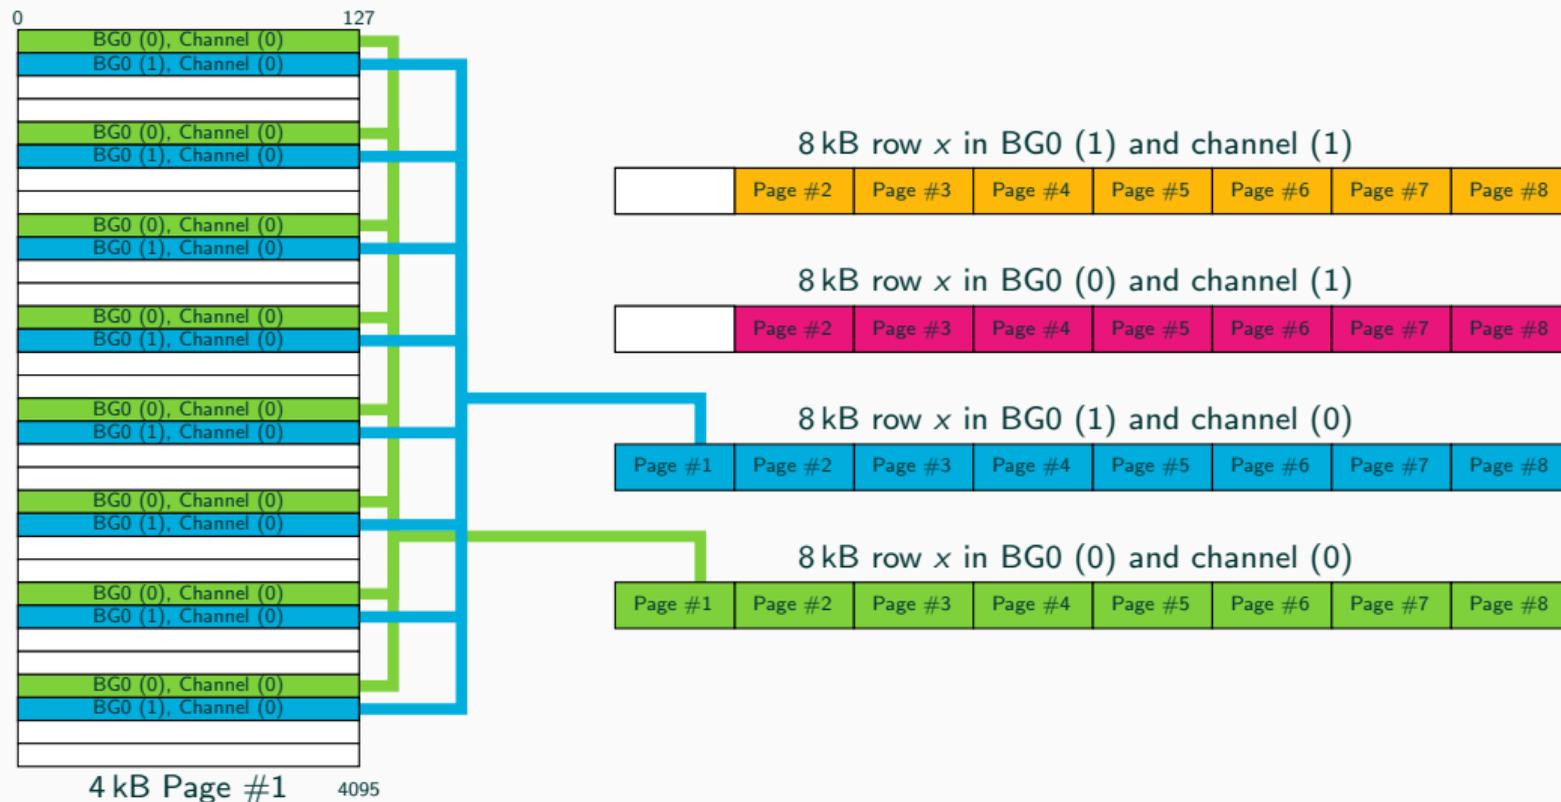
8 kB row  $x$  in BG0 (1) and channel (0)

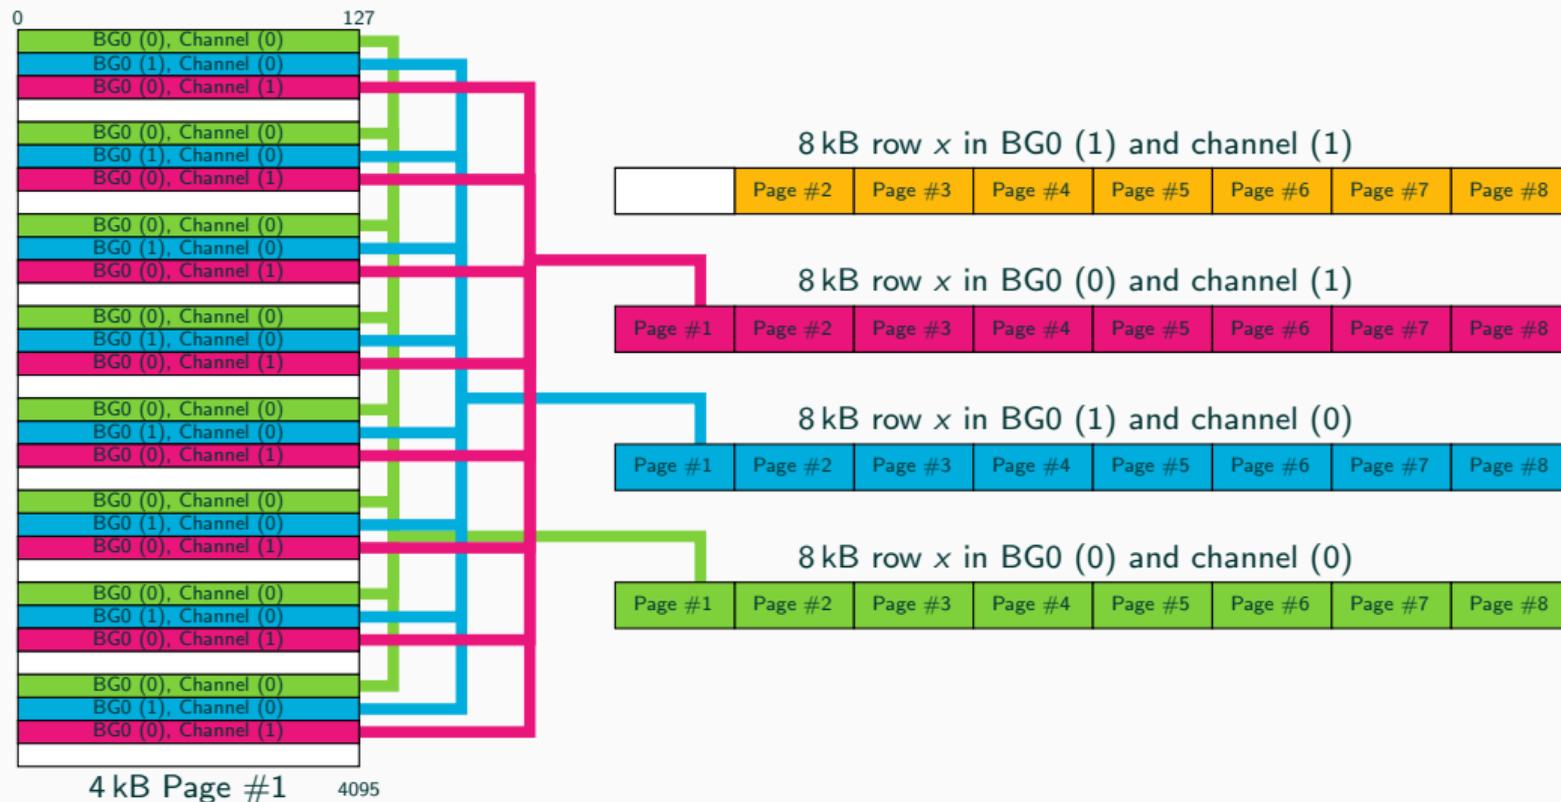


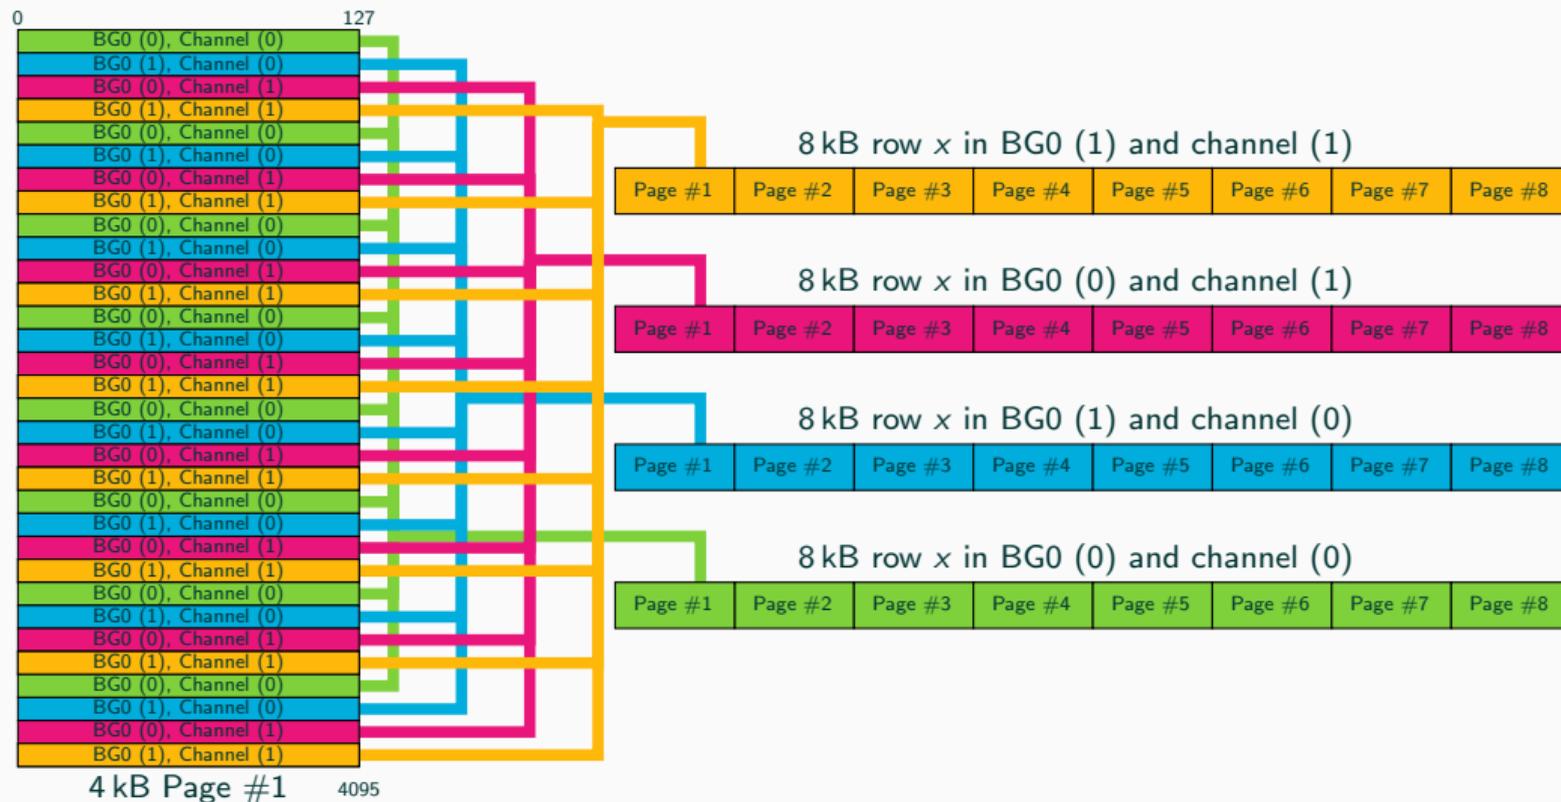
8 kB row  $x$  in BG0 (0) and channel (0)

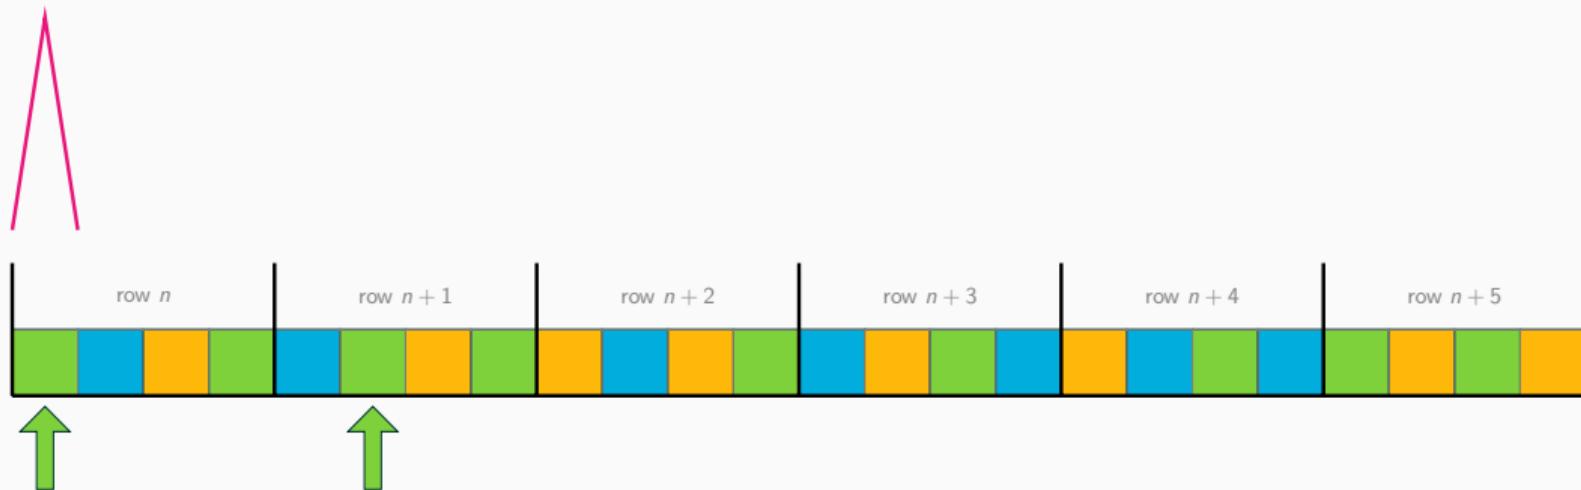


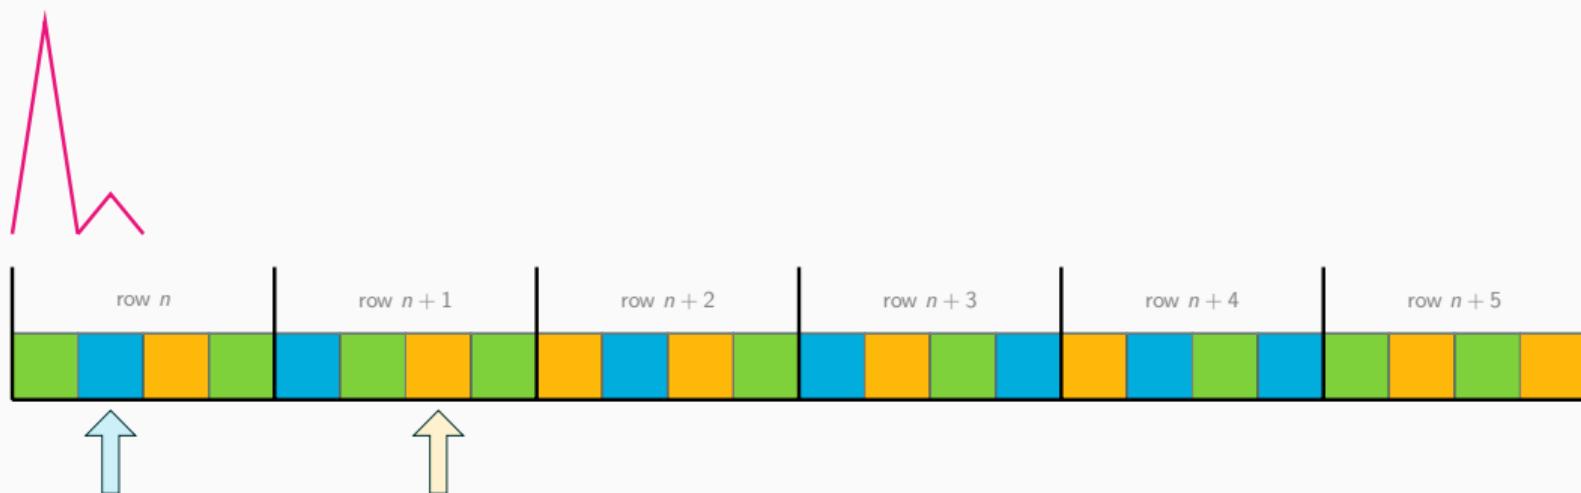


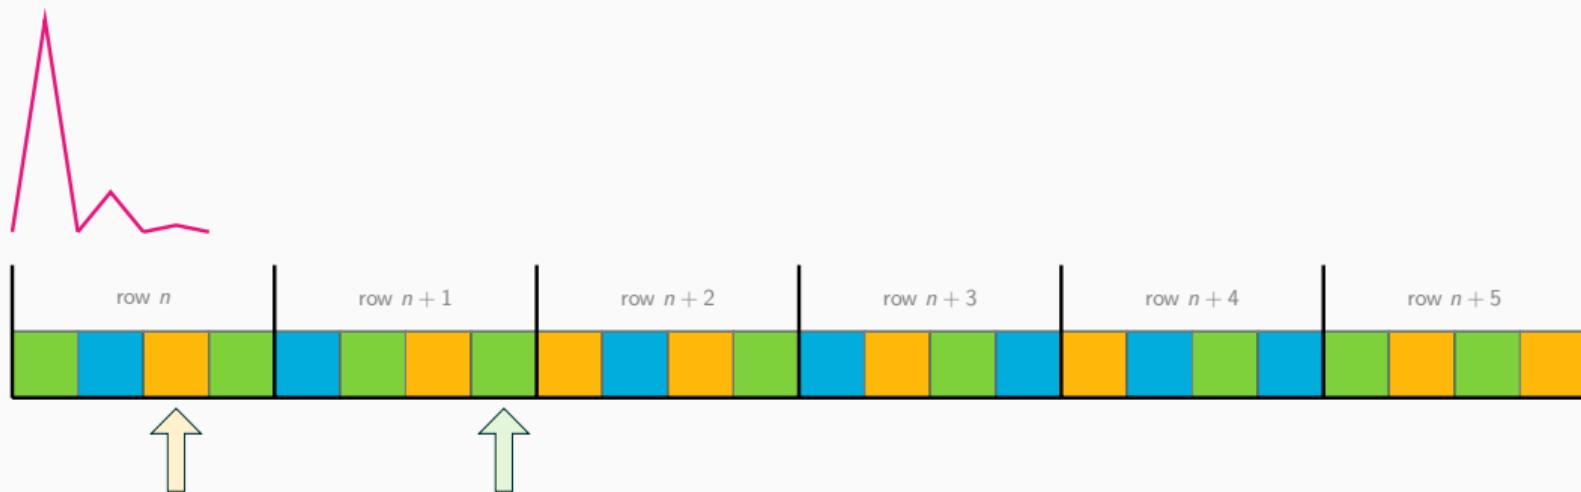


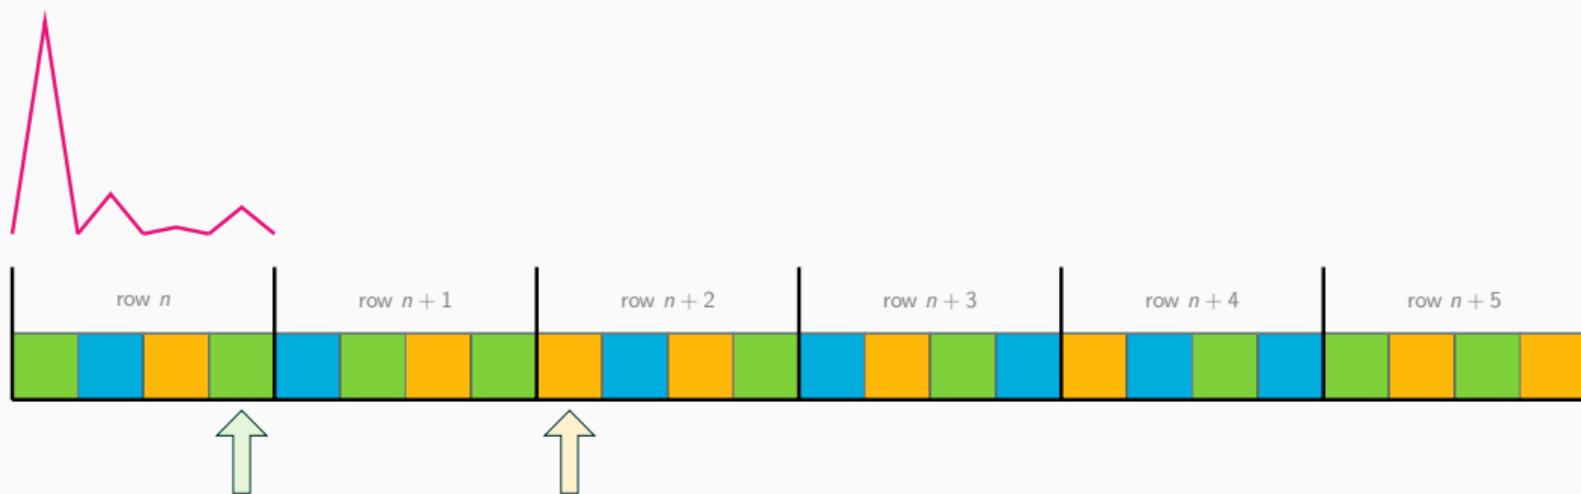


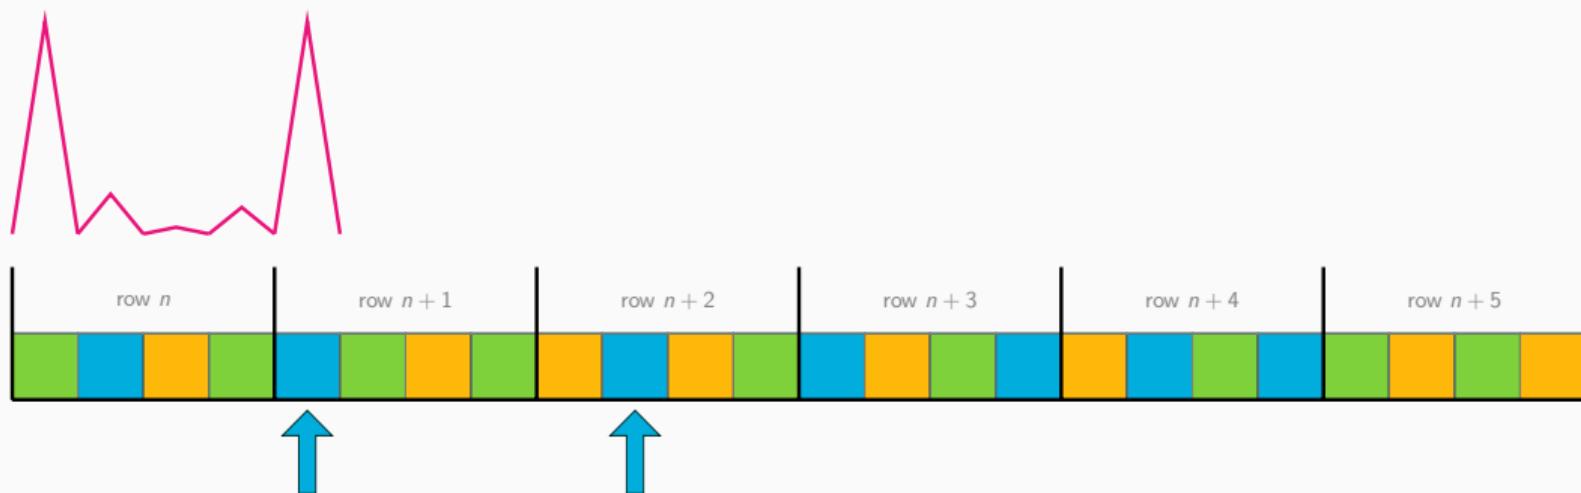


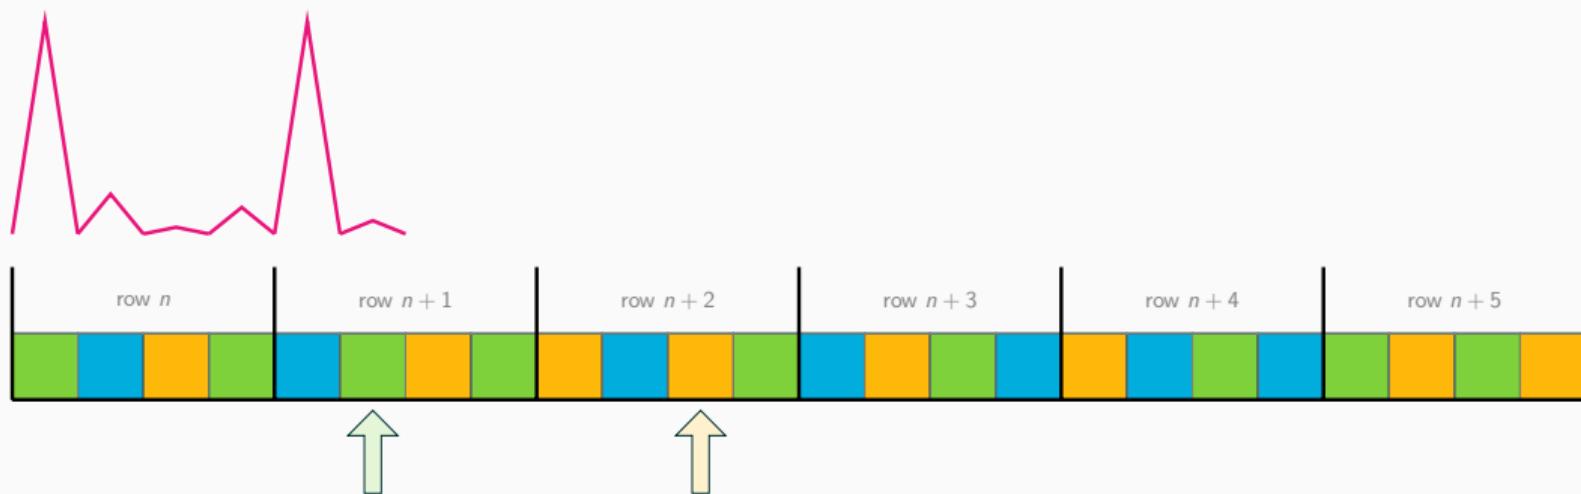






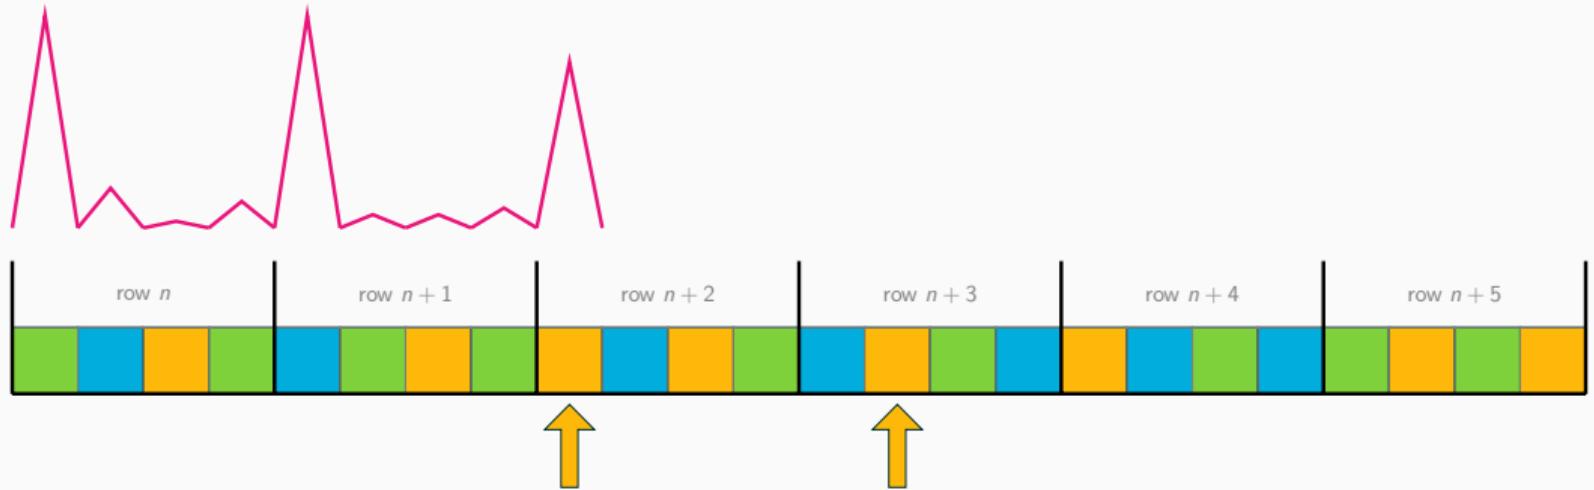


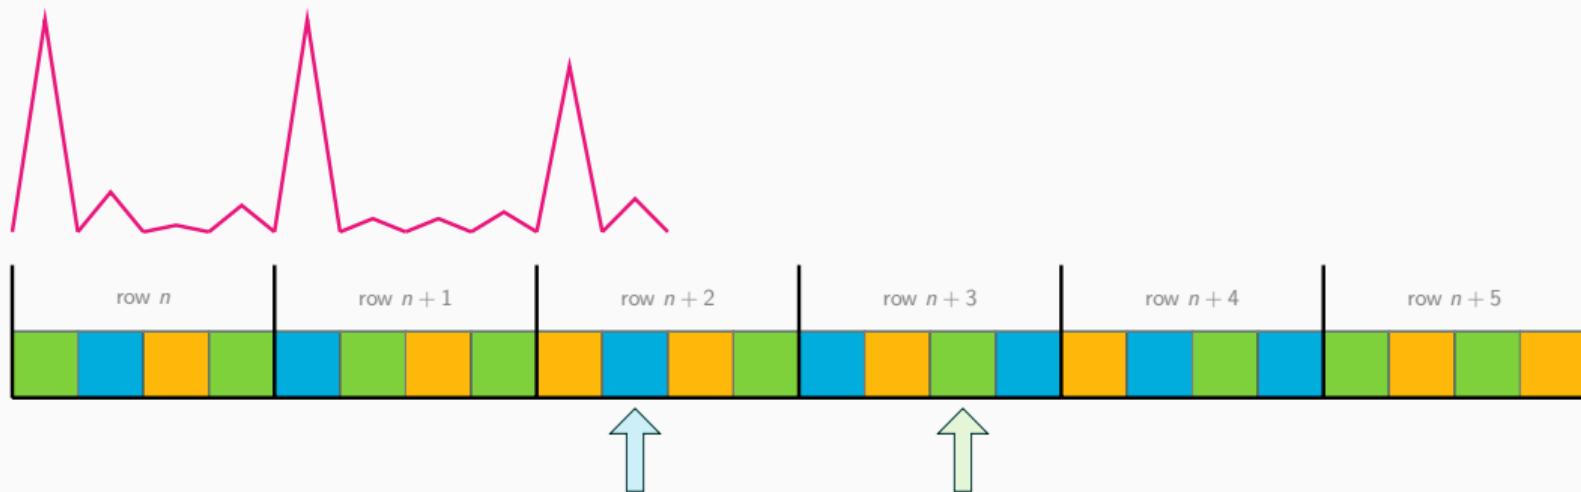




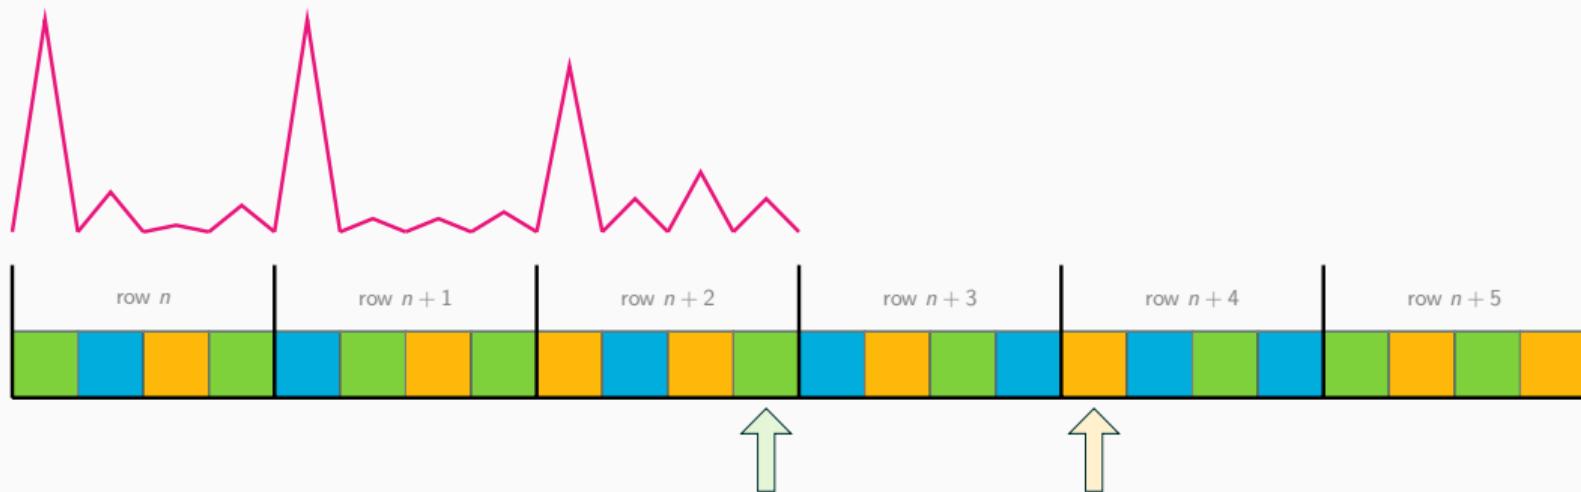












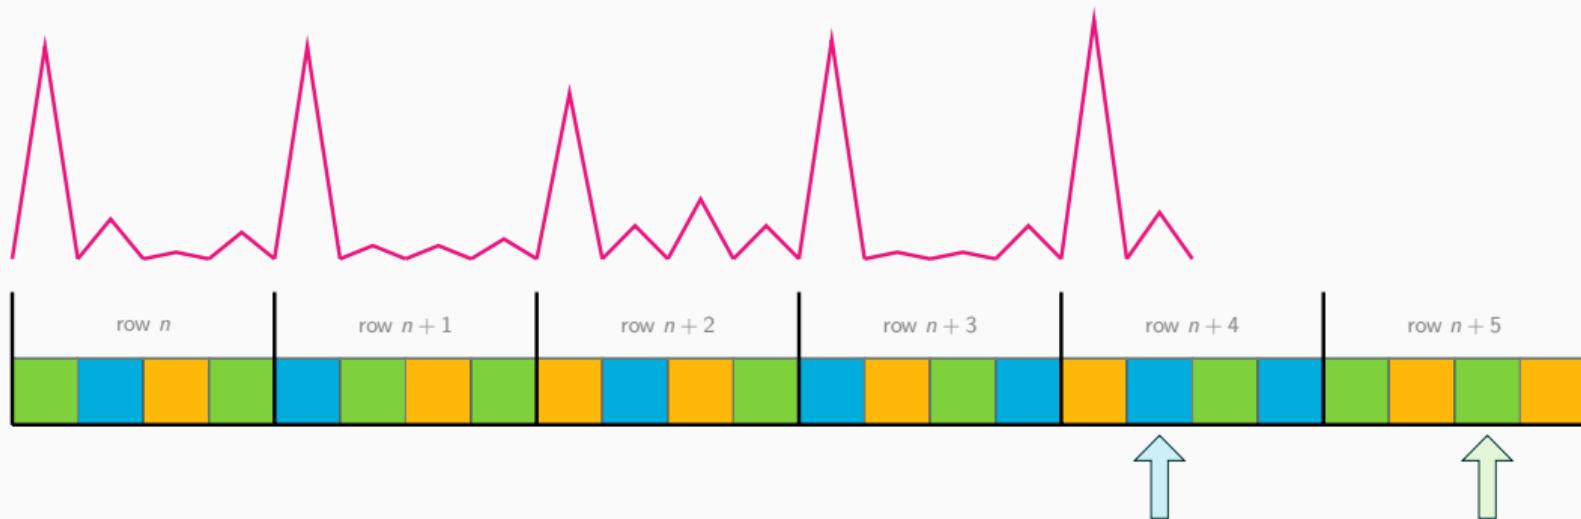


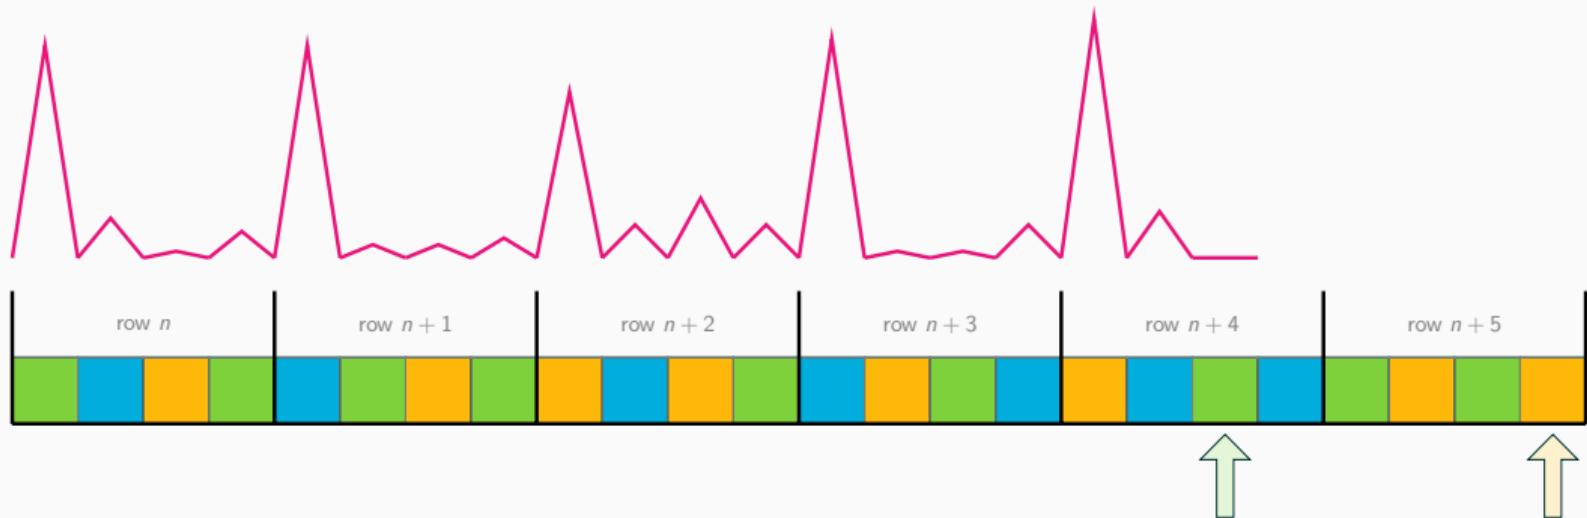




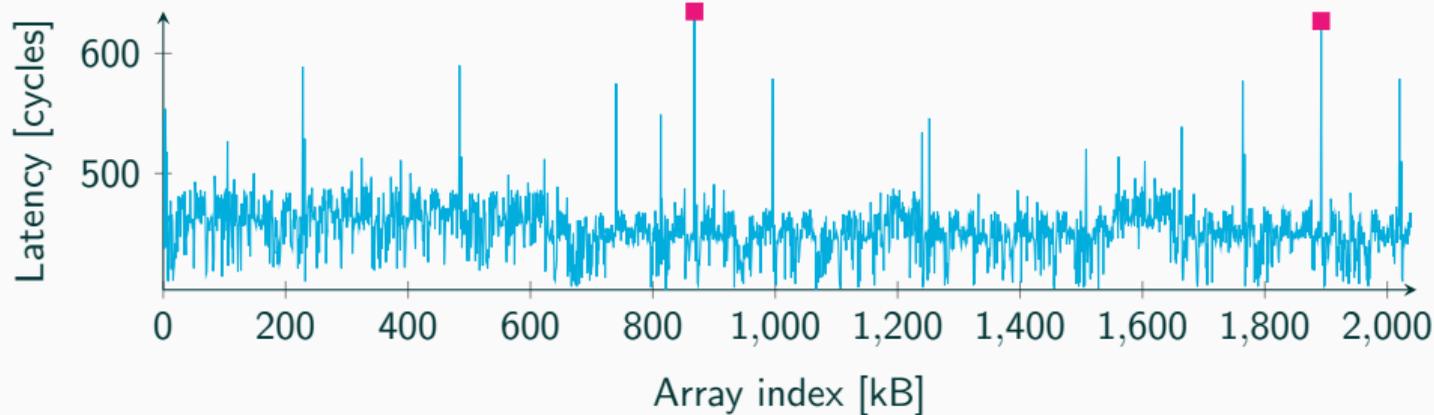


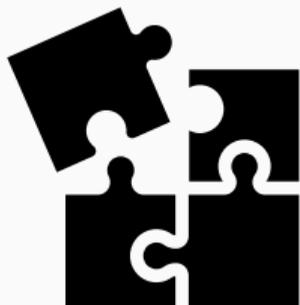




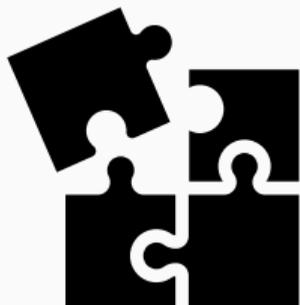


Result on an Intel i5-6200U

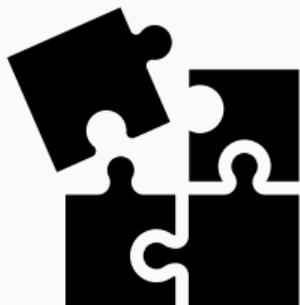




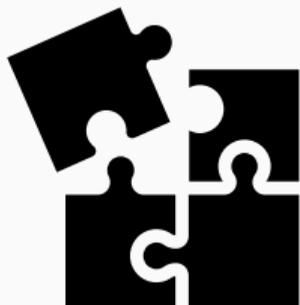
1. Use the **counting primitive** to measure DRAM accesses



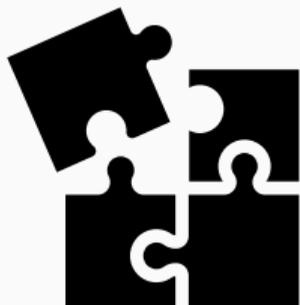
1. Use the **counting primitive** to measure DRAM accesses
2. Through the DRAM side channel, determine the **row borders**



1. Use the **counting primitive** to measure DRAM accesses
2. Through the DRAM side channel, determine the **row borders**
3. Row borders have the 18 LSBs set to '0' → maps to **cache set '0'**



1. Use the **counting primitive** to measure DRAM accesses
2. Through the DRAM side channel, determine the **row borders**
3. Row borders have the 18 LSBs set to '0' → maps to **cache set '0'**
4. Build the **eviction set** for the Prime+Probe attack



1. Use the **counting primitive** to measure DRAM accesses
2. Through the DRAM side channel, determine the **row borders**
3. Row borders have the 18 LSBs set to '0' → maps to **cache set '0'**
4. Build the **eviction set** for the Prime+Probe attack
5. Mount **Prime+Probe** on the buffer containing the multiplier

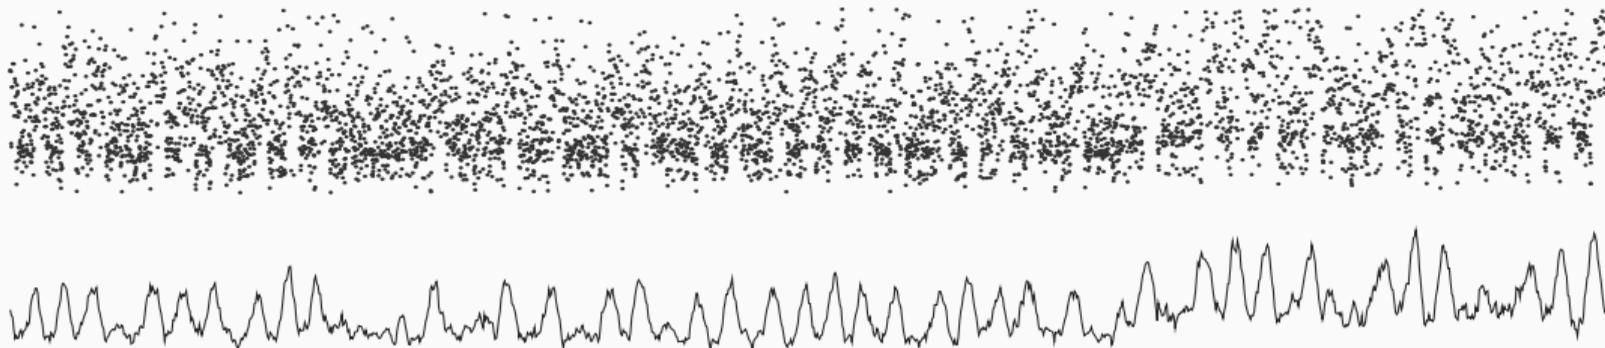
# Results

---

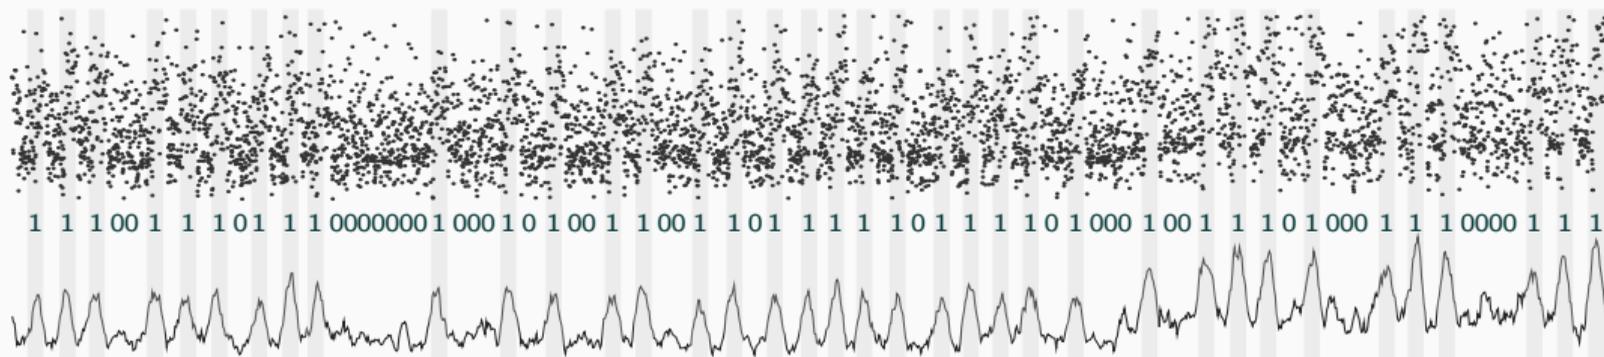
Raw Prime+Probe trace...



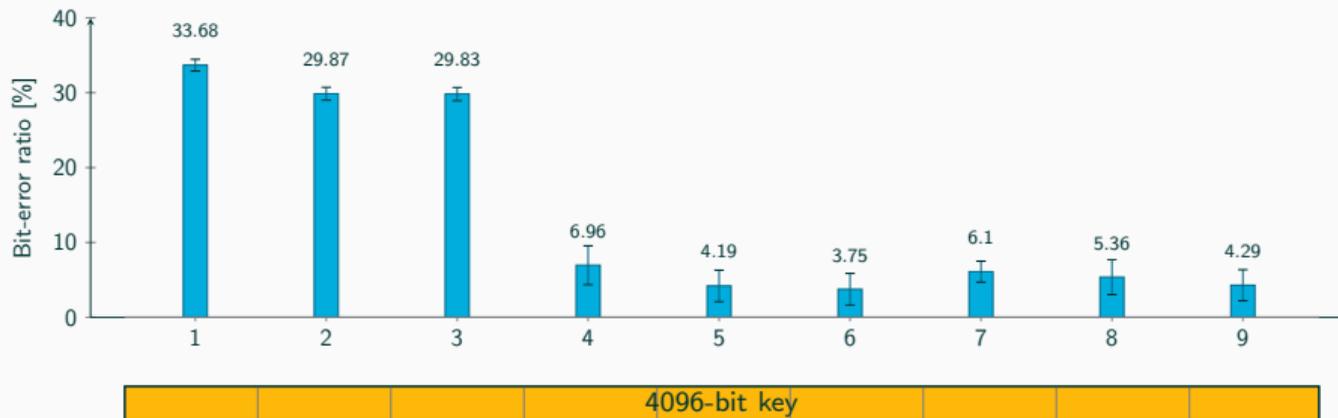
...processed with a simple moving average...



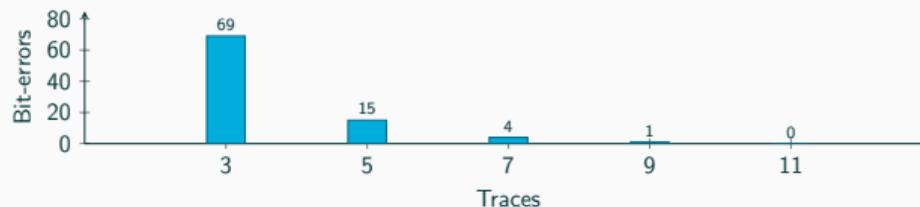
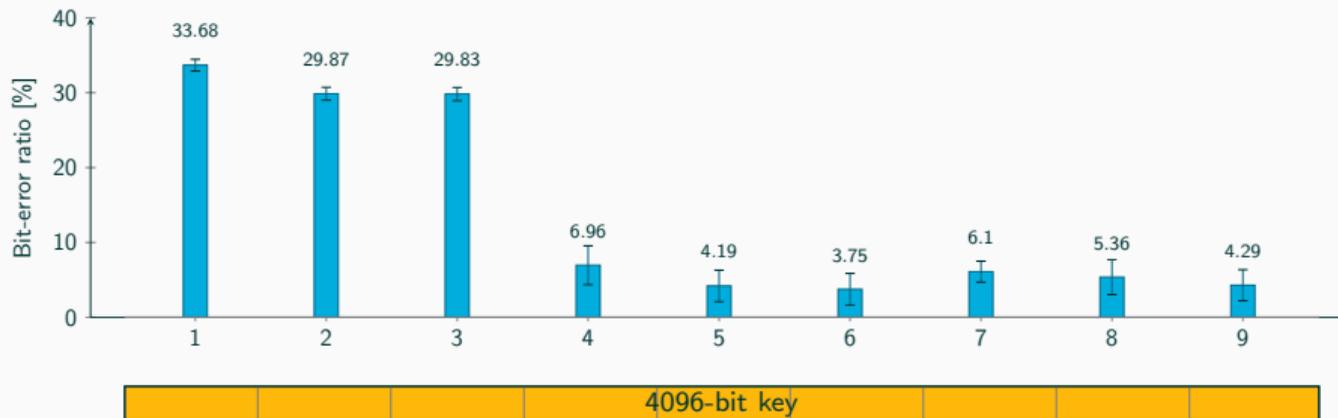
...allows to clearly see the bits of the exponent



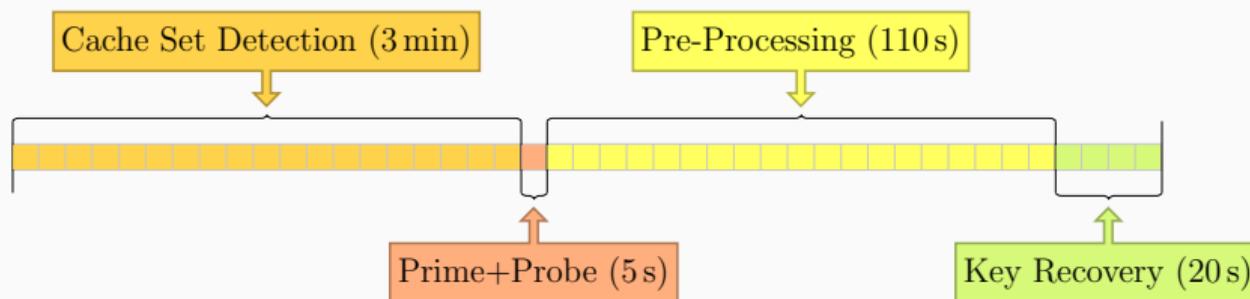
Error probability depends on which cache set of the key we attack

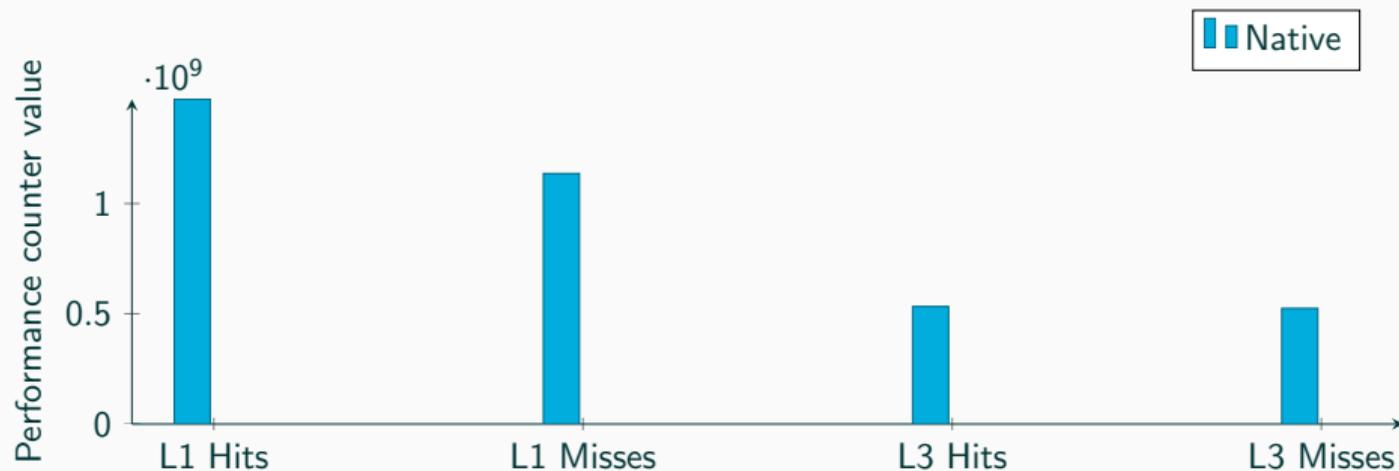


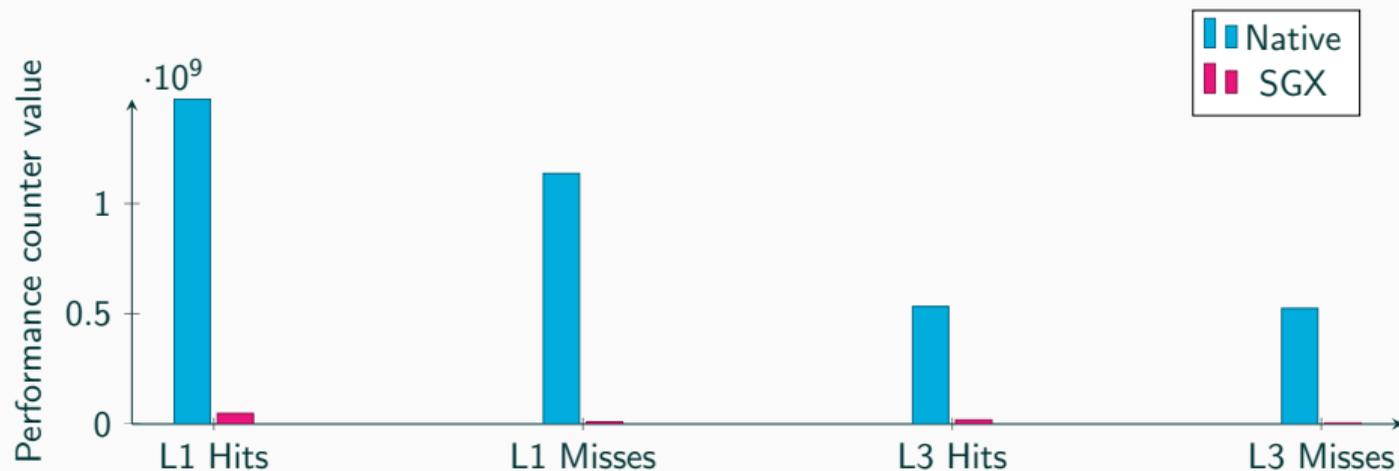
Error probability depends on which cache set of the key we attack

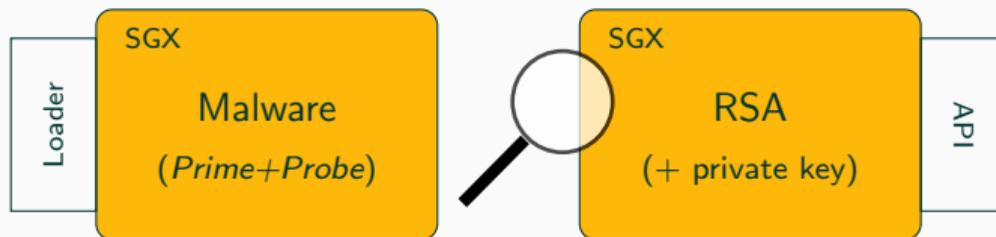


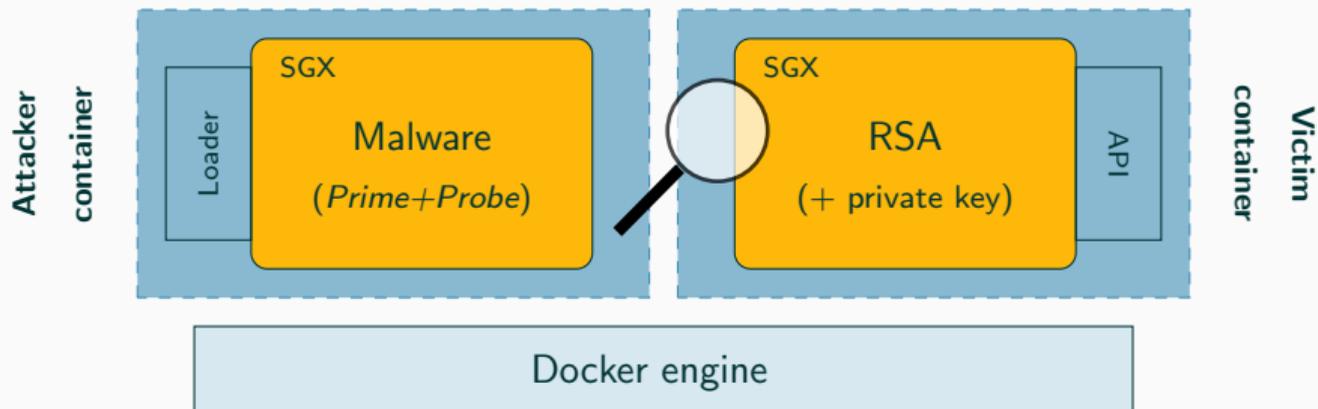
Full recovery of a 4096-bit RSA key in approximately 5 minutes

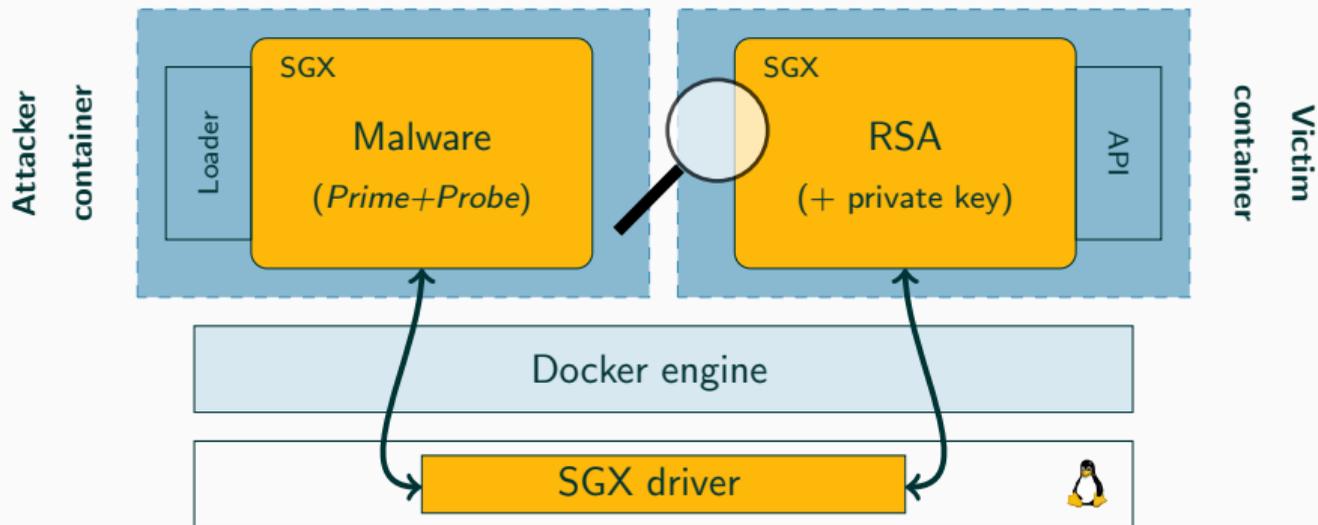






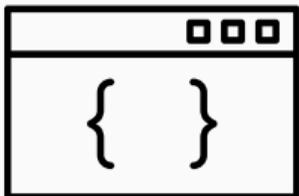




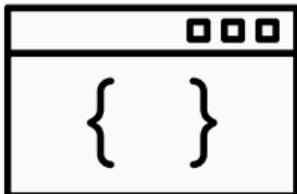


## Countermeasures

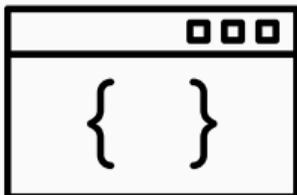
---



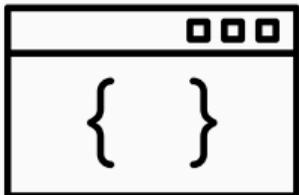
- Cache attacks can be prevented on **source level**



- Cache attacks can be prevented on **source level**
- Use **side-channel resistant** crypto implementations



- Cache attacks can be prevented on **source level**
- Use **side-channel resistant** crypto implementations
- **Exponent blinding** for RSA prevents multi-trace attacks



- Cache attacks can be prevented on **source level**
- Use **side-channel resistant** crypto implementations
- **Exponent blinding** for RSA prevents multi-trace attacks
- **Bit-sliced** implementations are not vulnerable to cache attacks



- Trusting the operating system weakens SGX threat model



- Trusting the operating system weakens SGX threat model
- Method for the operating system to inspect enclave code



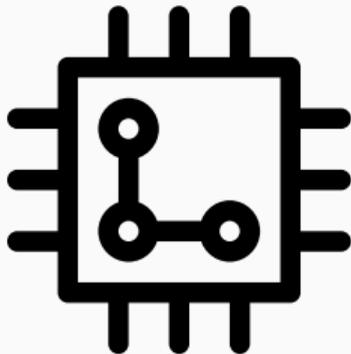
- Trusting the operating system weakens SGX threat model
- Method for the operating system to **inspect enclave code**
- Re-enable certain **performance counters**, such as L3 hits/misses



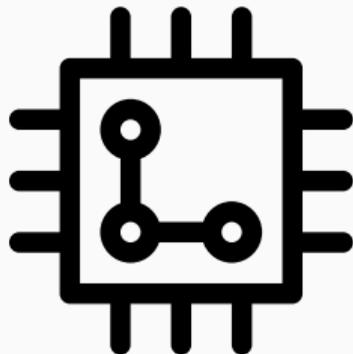
- Trusting the operating system weakens SGX threat model
- Method for the operating system to **inspect enclave code**
- Re-enable certain **performance counters**, such as L3 hits/misses
- **Enclave coloring** to prevent cross-enclave attacks



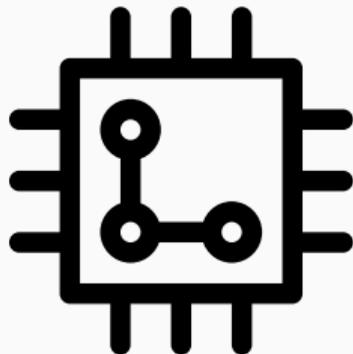
- Trusting the operating system weakens SGX threat model
- Method for the operating system to **inspect enclave code**
- Re-enable certain **performance counters**, such as L3 hits/misses
- **Enclave coloring** to prevent cross-enclave attacks
- **Heap randomization** to randomize cache sets



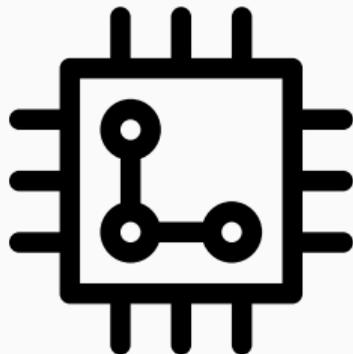
- Intel could prevent attacks by changing the hardware



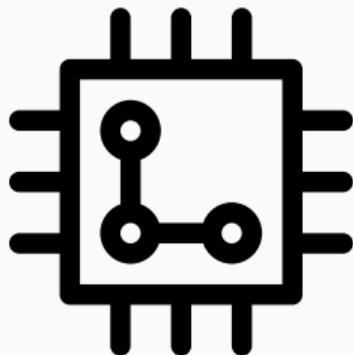
- Intel could prevent attacks by changing the hardware
- Combine Cache Allocation Technology (**CAT**) with SGX



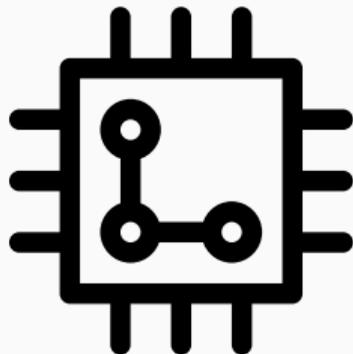
- Intel could prevent attacks by changing the hardware
- Combine Cache Allocation Technology (**CAT**) with SGX
  - Instead of controlling CAT from the OS, combine it with eenter



- Intel could prevent attacks by changing the hardware
- Combine Cache Allocation Technology (CAT) with SGX
  - Instead of controlling CAT from the OS, combine it with eenter
  - Entering an enclave would automatically activate CAT for this core



- Intel could prevent attacks by changing the hardware
- Combine Cache Allocation Technology (CAT) with SGX
  - Instead of controlling CAT from the OS, combine it with eenter
  - Entering an enclave would automatically activate CAT for this core
  - L3 is then isolated from all other enclaves and applications



- Intel could prevent attacks by changing the hardware
- Combine Cache Allocation Technology (**CAT**) with SGX
  - Instead of controlling CAT from the OS, combine it with eenter
  - Entering an enclave would automatically activate CAT for this core
  - L3 is then isolated from all other enclaves and applications
- Provide a non-shared **secure memory** element which is not cached

## Conclusion

---

- We showed that **attacks** can be mounted from **within SGX** enclaves

- We showed that **attacks** can be mounted from **within SGX** enclaves
- We presented new techniques for side-channel attacks, including a timing measurement with the currently **highest resolution**

- We showed that **attacks** can be mounted from **within SGX** enclaves
- We presented new techniques for side-channel attacks, including a timing measurement with the currently **highest resolution**
- Our end-to-end attack recovered 96% of a 4096-bit RSA key from a single trace, and the full key using only 11 traces

- We showed that **attacks** can be mounted from **within SGX** enclaves
- We presented new techniques for side-channel attacks, including a timing measurement with the currently **highest resolution**
- Our end-to-end attack recovered 96% of a 4096-bit RSA key from a single trace, and the full key using only 11 traces
- SGX allows to completely **hide an attack** from state-of-the-art detection techniques

- We showed that **attacks** can be mounted from **within SGX** enclaves
- We presented new techniques for side-channel attacks, including a timing measurement with the currently **highest resolution**
- Our end-to-end attack recovered 96% of a 4096-bit RSA key from a single trace, and the full key using only 11 traces
- SGX allows to completely **hide an attack** from state-of-the-art detection techniques
- The attack showed that SGX is **not a magic solution** to make software safe

**Thank you!**

# Malware Guard Extension: Using SGX to Conceal Cache Attacks

**Michael Schwarz, Samuel Weiser, Daniel Gruss, Clémentine Maurice, Stefan Mangard**

July 6, 2017

Graz University of Technology