

CopyCat: Controlled Instruction-Level Attacks on Enclaves

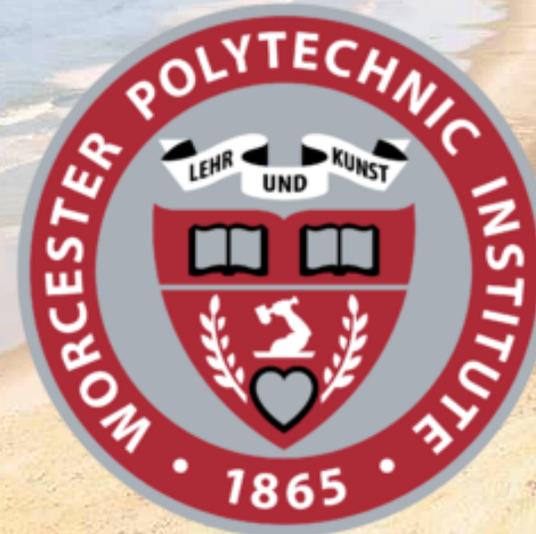
- Daniel Moghimi
- Jo Van Bulck
- Nadia Heninger
- Frank Piessens
- Berk Sunar

29TH USENIX
SECURITY SYMPOSIUM

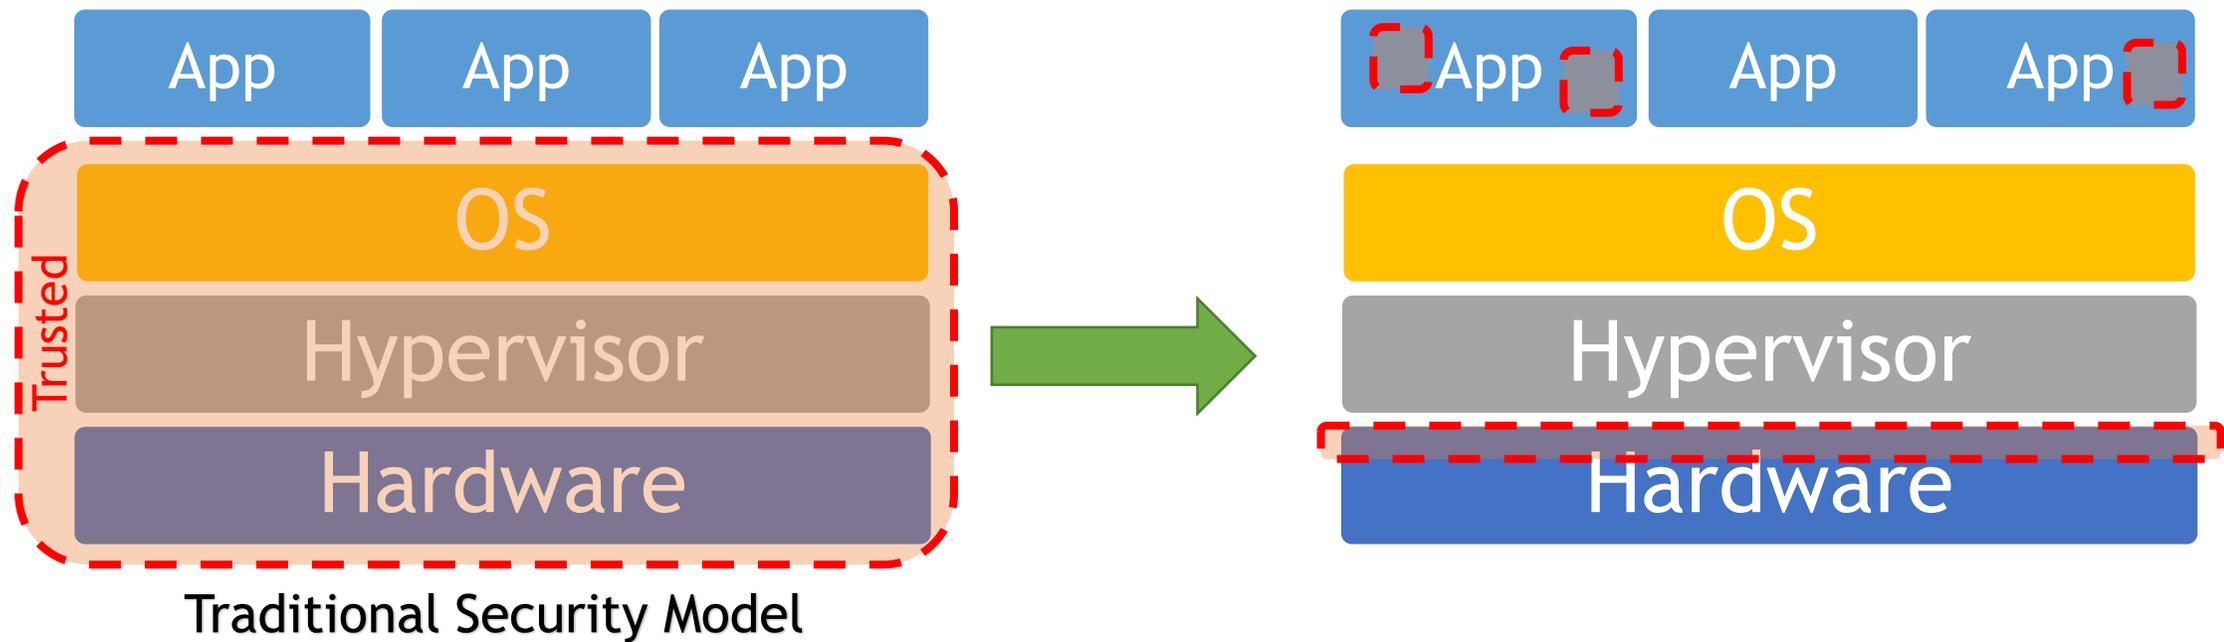
AUGUST 12-14, 2020

DistriNet

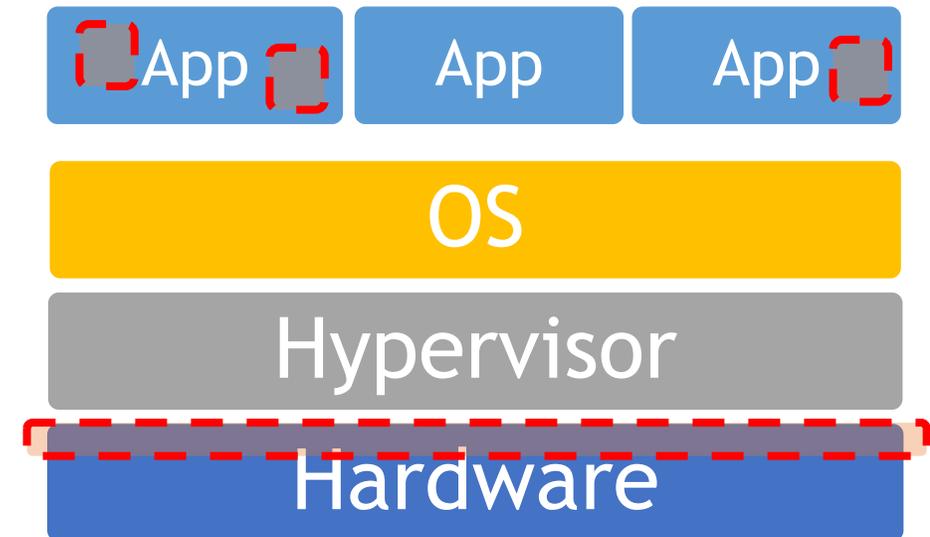
KU LEUVEN



- Intel Software Guard eXtensions (SGX)



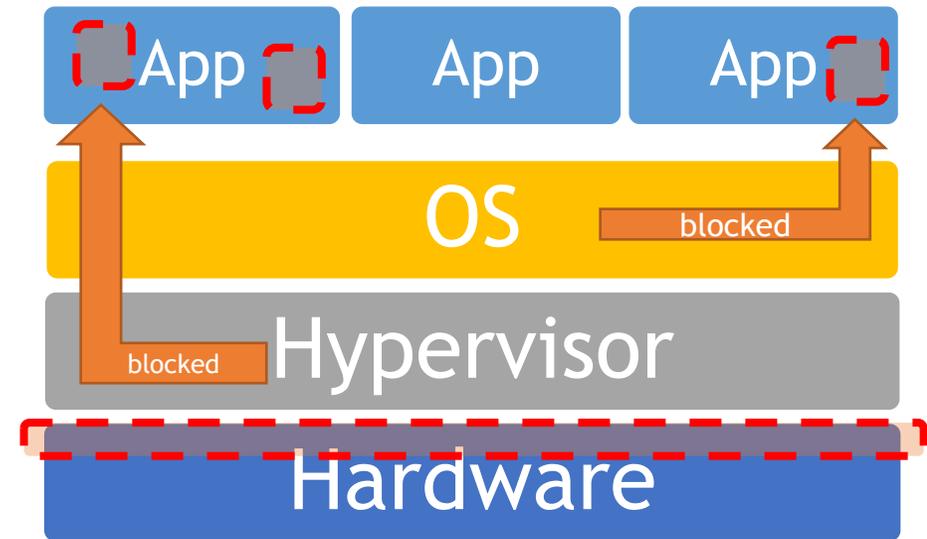
- Intel Software Guard eXtensions (SGX)
- **Enclave:** Hardware protected user-level software module
 - Mapped by the Operating System
 - Loaded by the user program
 - Authenticated and Encrypted by CPU



- Intel Software Guard eXtensions (SGX)
- **Enclave:** Hardware protected user-level software module
 - Mapped by the Operating System
 - Loaded by the user program
 - Authenticated and Encrypted by CPU
- Protects against system level adversary

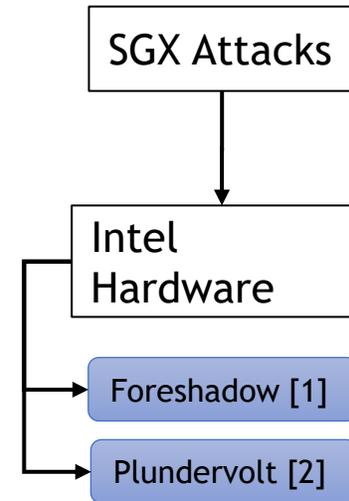
New Attacker Model:

Attacker gets full control over OS



- Intel's Responsibility

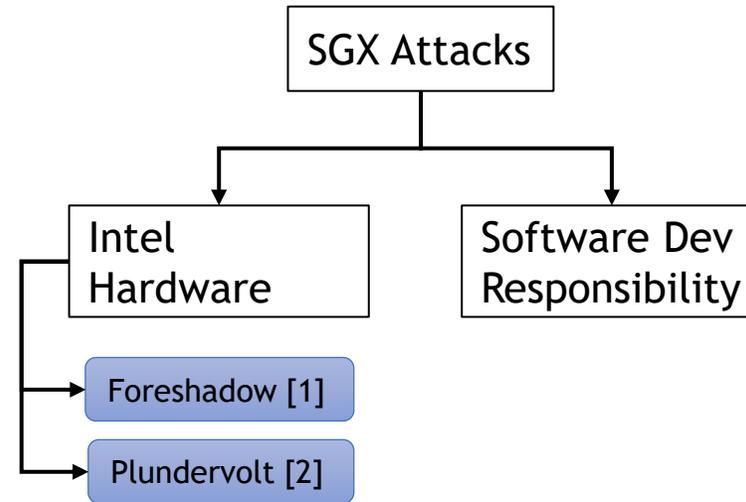
- Microcode Patches / Hardware mitigation
- TCB Recovery
 - Old Keys are Revoked
 - Remote attestation succeeds only with mitigation.



[1] Van Bulck et al. "Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution." USENIX Security 2018.
[2] Murdock et al. "Plundervolt: Software-based fault injection attacks against Intel SGX." IEEE S&P 2020.

• Intel's Responsibility

- Microcode Patches / Hardware mitigation
- TCB Recovery
 - Old Keys are Revoked
 - Remote attestation succeeds only with mitigation.



[1] Van Bulck et al. "Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution." USENIX Security 2018.

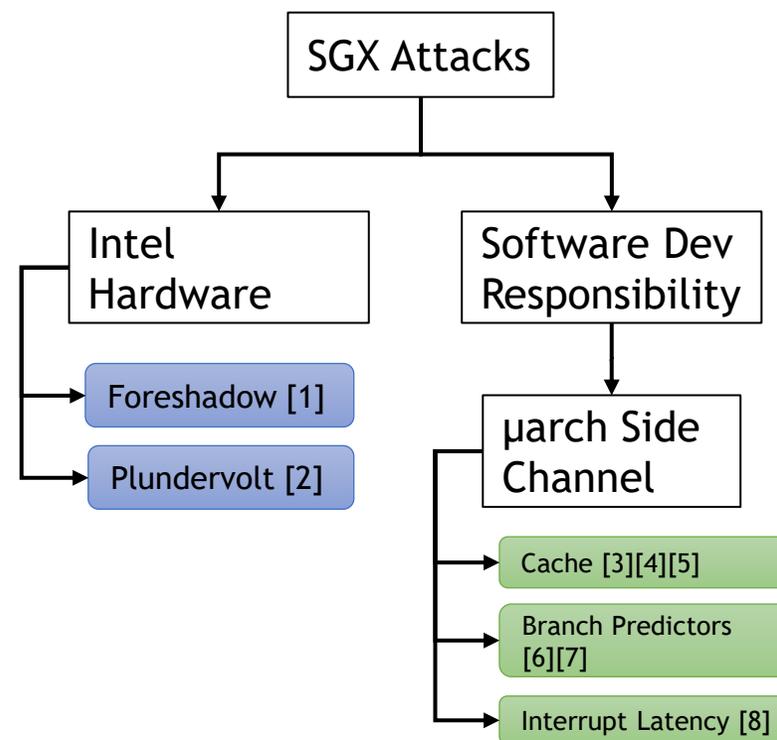
[2] Murdock et al. "Plundervolt: Software-based fault injection attacks against Intel SGX." IEEE S&P 2020.

• Intel's Responsibility

- Microcode Patches / Hardware mitigation
- TCB Recovery
 - Old Keys are Revoked
 - Remote attestation succeeds only with mitigation.
- Hyperthreading is out
 - Remote Attestation Warning

• μ arch Side Channel

- Constant-time Coding
- Flushing and Isolating buffers
- Probabilistic



[1] Van Bulck et al. "Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution." USENIX Security 2018.

[2] Murdock et al. "Plundervolt: Software-based fault injection attacks against Intel SGX." IEEE S&P 2020.

[3] Moghimi et al. "Cachezoom: How SGX amplifies the power of cache attacks." CHES 2017.

[4] Brasser et al. "Software grand exposure: {SGX} cache attacks are practical." USENIX WOOT 2017.

[5] Schwarz et al. "Malware guard extension: Using SGX to conceal cache attacks." DIMVA 2017.

[6] Evtushkin, Dmitry, et al. "Branchscope: A new side-channel attack on directional branch predictor." ACM SIGPLAN 2018.

[7] Lee, Sangho, et al. "Inferring fine-grained control flow inside {SGX} enclaves with branch shadowing." USENIX Security 2017.

[8] Van Bulck et al. "Nemesis: Studying microarchitectural timing leaks in rudimentary CPU interrupt logic." ACM CCS 2018.

• Intel's Responsibility

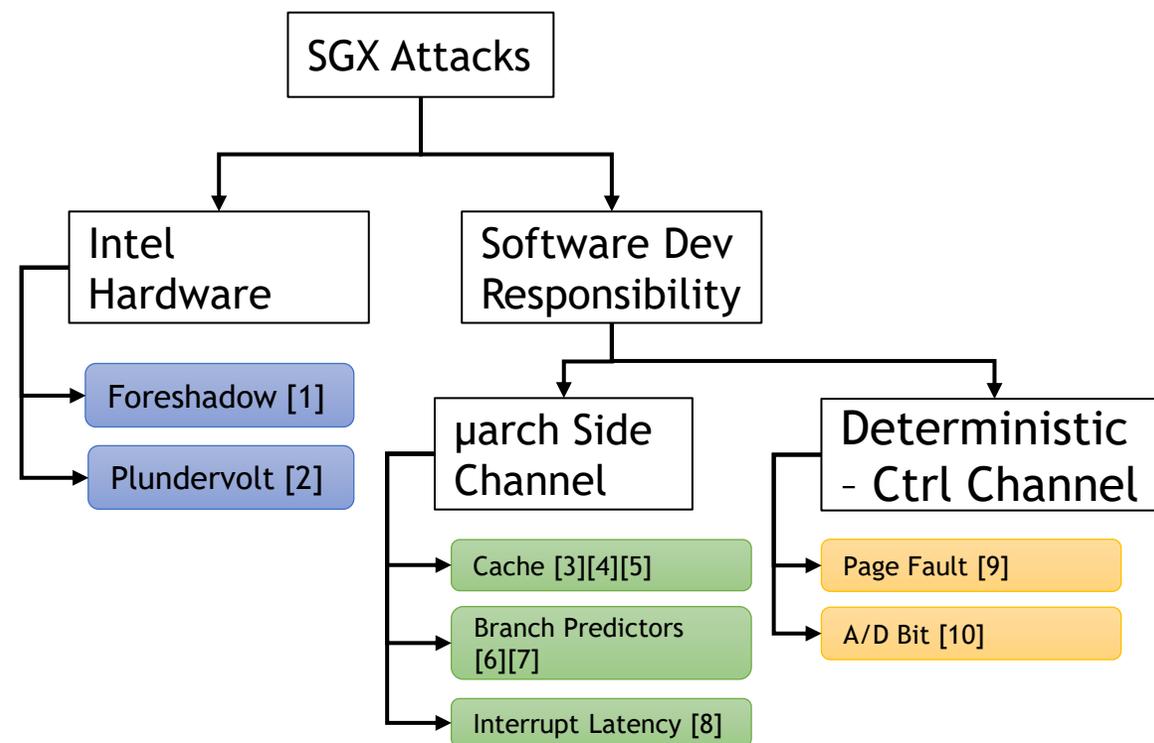
- Microcode Patches / Hardware mitigation
- TCB Recovery
 - Old Keys are Revoked
 - Remote attestation succeeds only with mitigation.
- Hyperthreading is out
 - Remote Attestation Warning

• μ arch Side Channel

- Constant-time Coding
- Flushing and Isolating buffers
- Probabilistic

• Deterministic Attacks

- Page Fault, A/D Bit, etc. (4kB Granularity)



[1] Van Bulck et al. "Foreshadow: Extracting the keys to the intel SGX kingdom with transient out-of-order execution." USENIX Security 2018.

[2] Murdock et al. "Plundervolt: Software-based fault injection attacks against Intel SGX." IEEE S&P 2020.

[3] Moghimi et al. "Cachezoom: How SGX amplifies the power of cache attacks." CHES 2017.

[4] Brasser et al. "Software grand exposure: {SGX} cache attacks are practical." USENIX WOOT 2017.

[5] Schwarz et al. "Malware guard extension: Using SGX to conceal cache attacks." DIMVA 2017.

[6] Evtushkin, Dmitry, et al. "Branchscope: A new side-channel attack on directional branch predictor." ACM SIGPLAN 2018.

[7] Lee, Sangho, et al. "Inferring fine-grained control flow inside {SGX} enclaves with branch shadowing." USENIX Security 2017.

[8] Van Bulck et al. "Nemesis: Studying microarchitectural timing leaks in rudimentary CPU interrupt logic." ACM CCS 2018.

[9] Xu et al. "Controlled-channel attacks: Deterministic side channels for untrusted operating systems." IEEE SGP 2015.

[10] Wang, Wenhao, et al. "Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX." ACM CCS 2017.



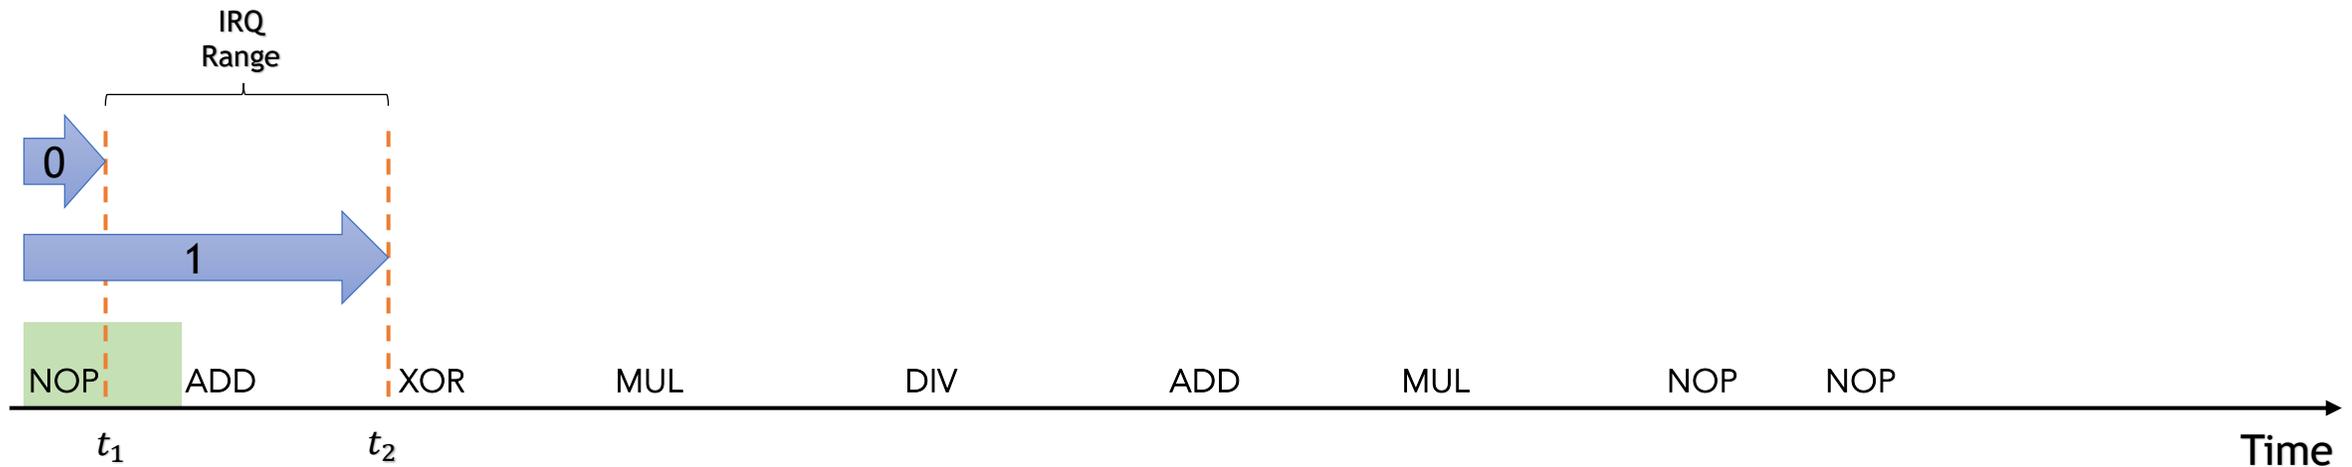
CopyCat Attack



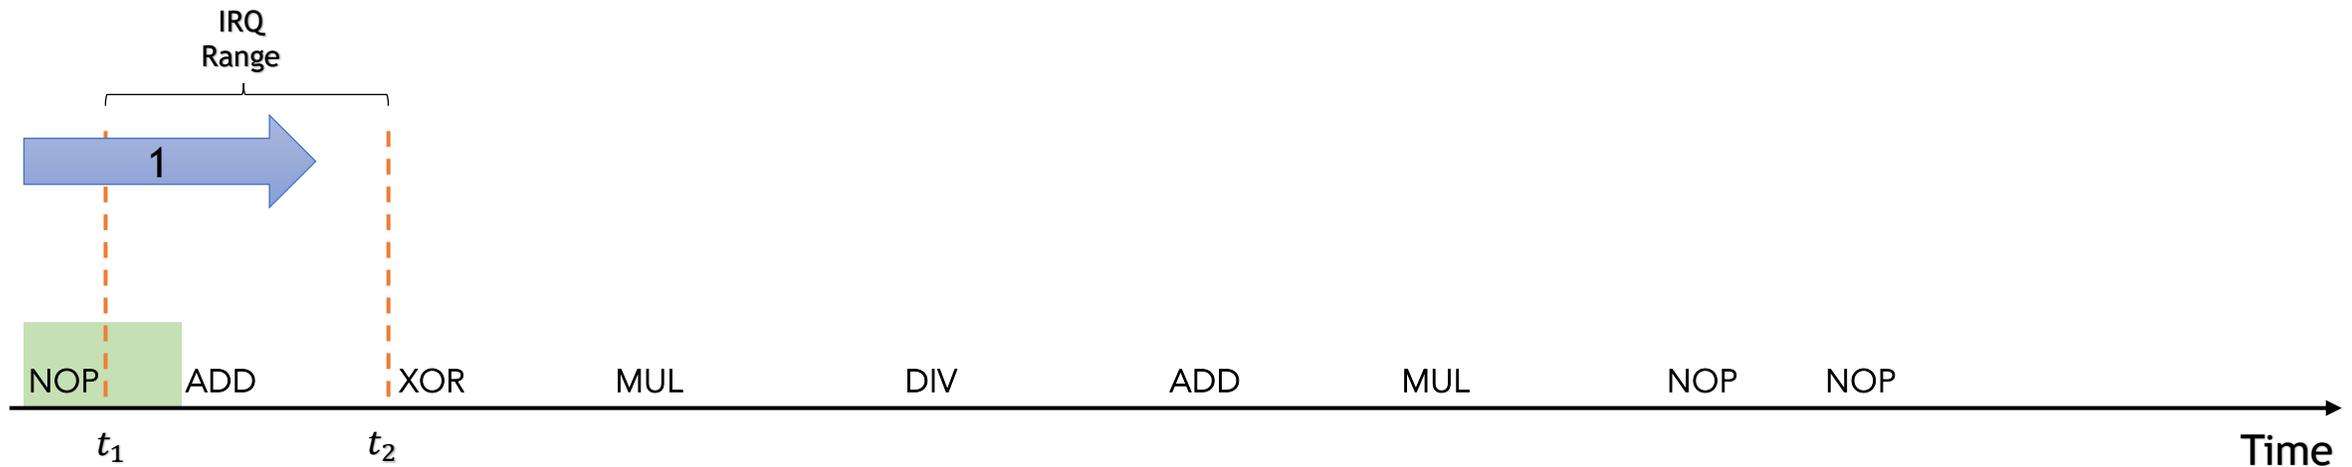
- Malicious OS controls the interrupt handler



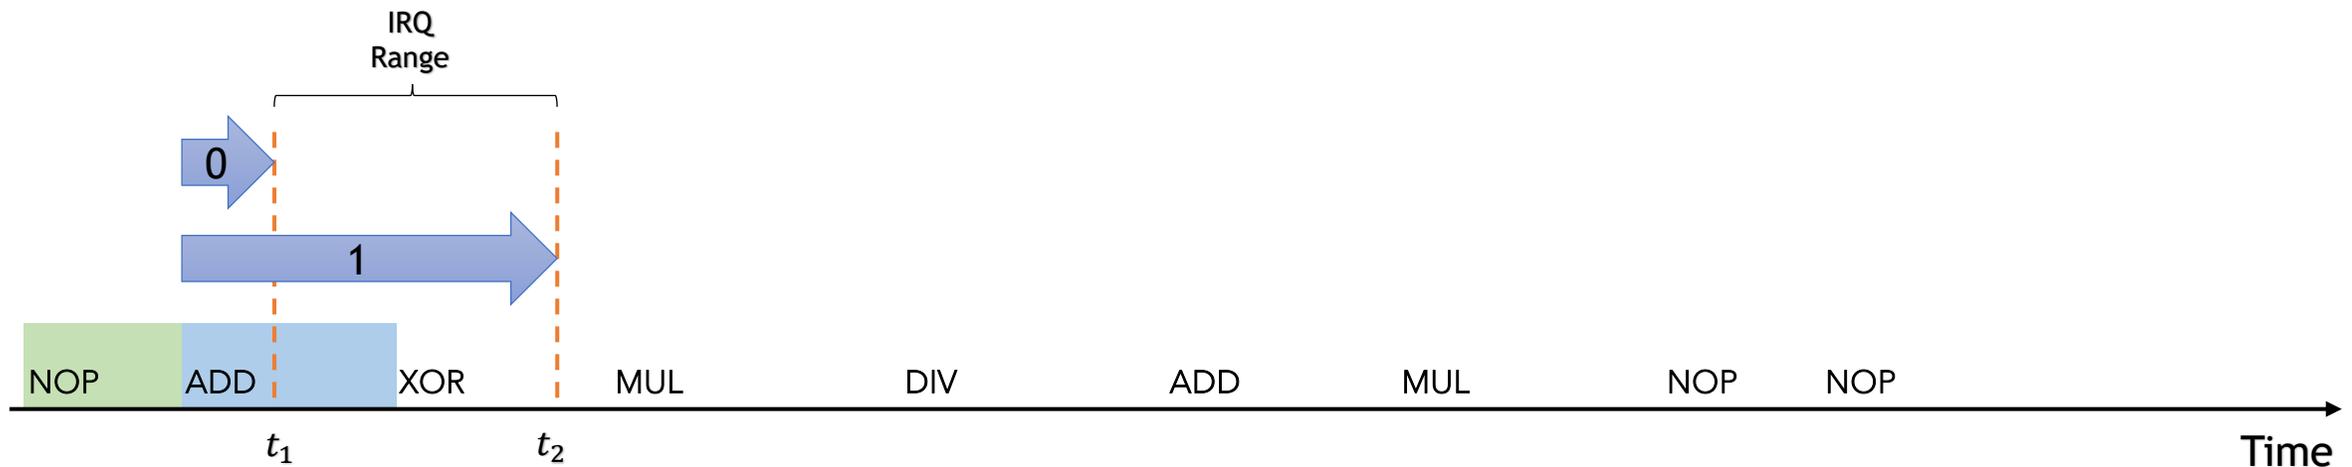
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions



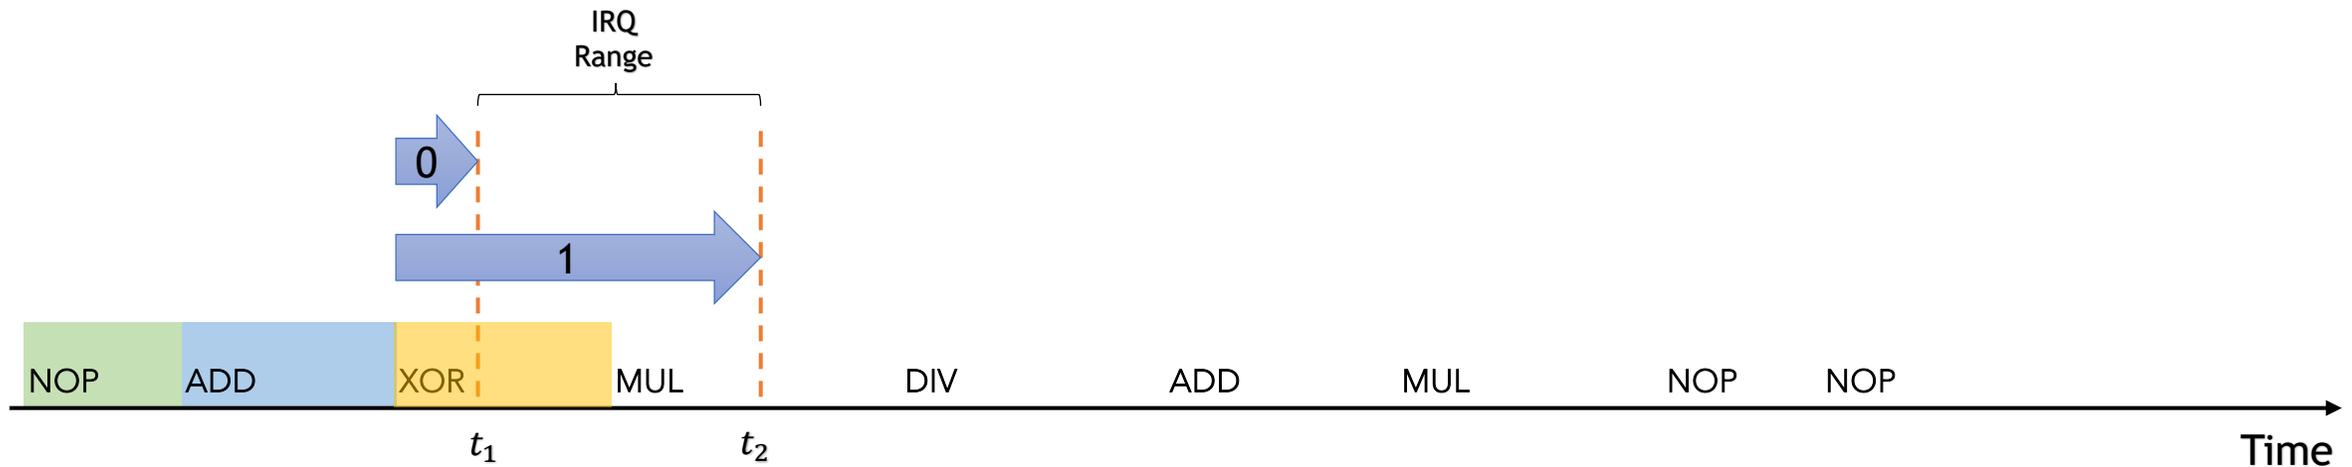
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions



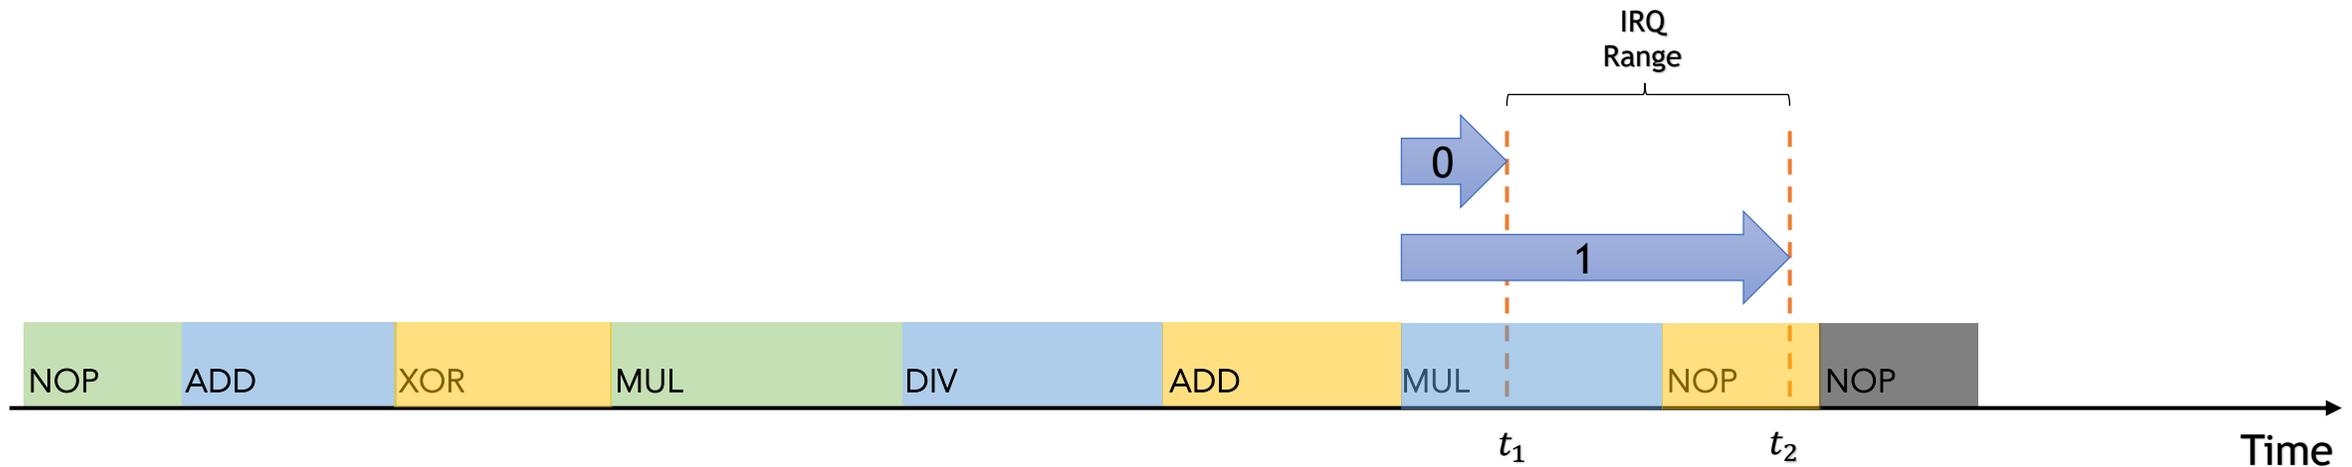
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions



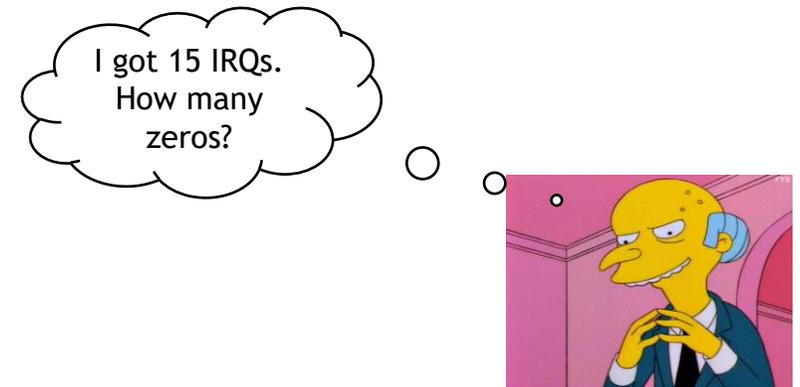
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions



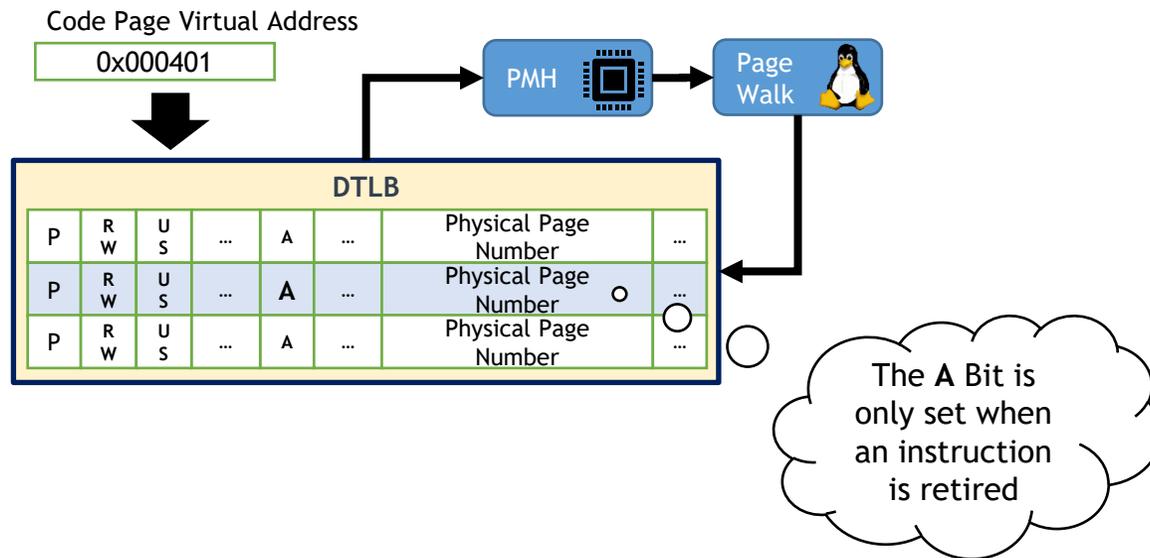
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions



- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions



- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions
- Filtering Zeros out: Clear the A bit before, Check the A bit after

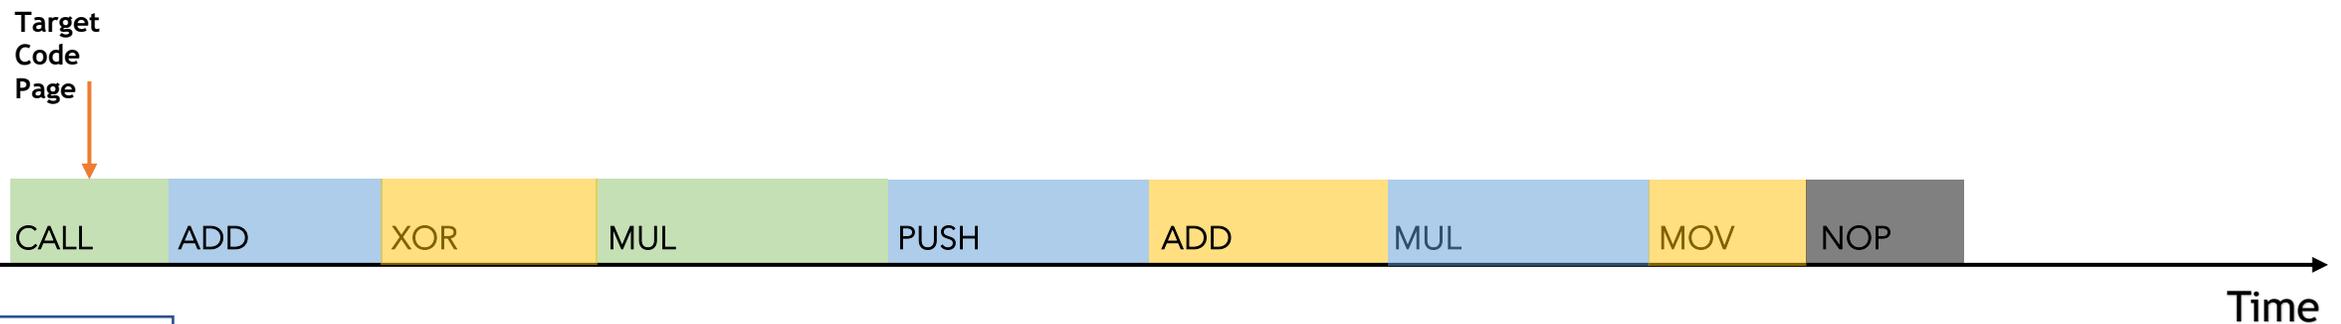


I got 15 IRQs.
How many
zeros?

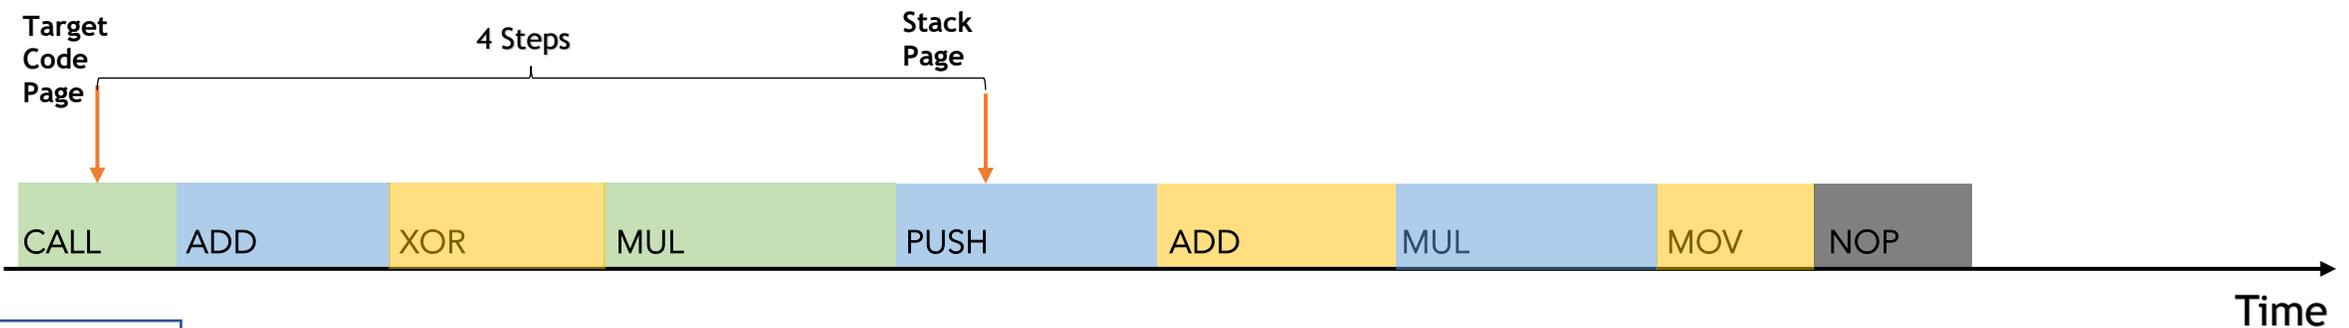


- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions
- Filtering Zeros out: Clear the A bit before, Check the A bit after
- Deterministic Instruction Counting

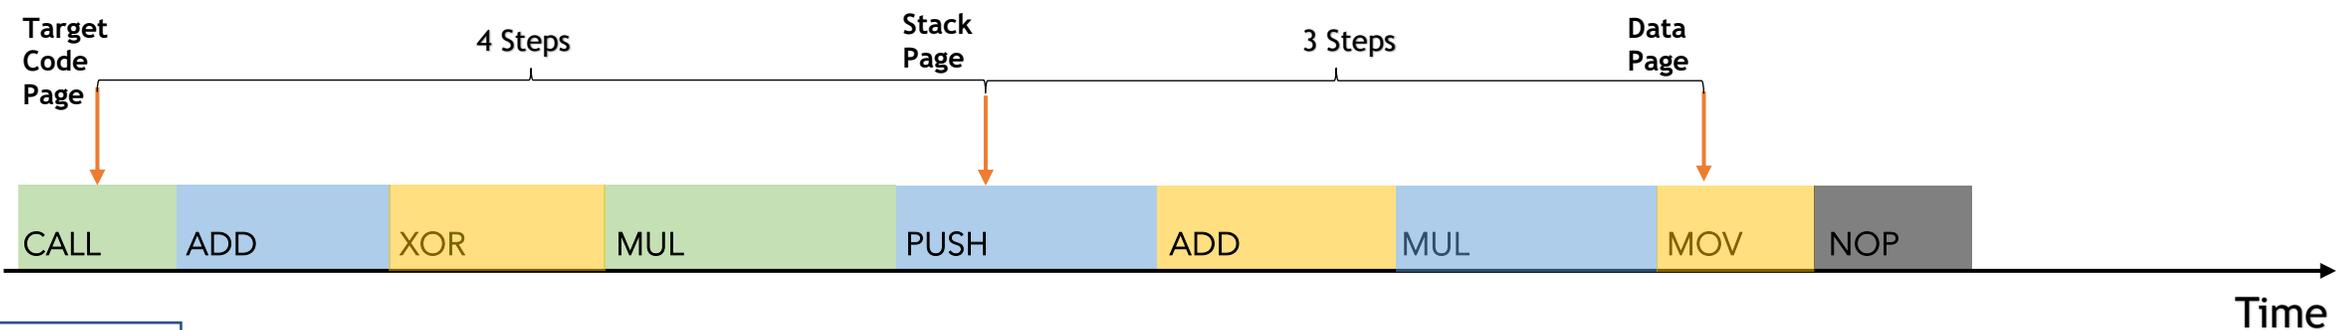
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions
- Filtering Zeros out: Clear the A bit before, Check the A bit after
- Deterministic Instruction Counting
- Counting from start to end is not useful.
 - A Secondary oracle
 - Page table attack as a deterministic secondary oracle



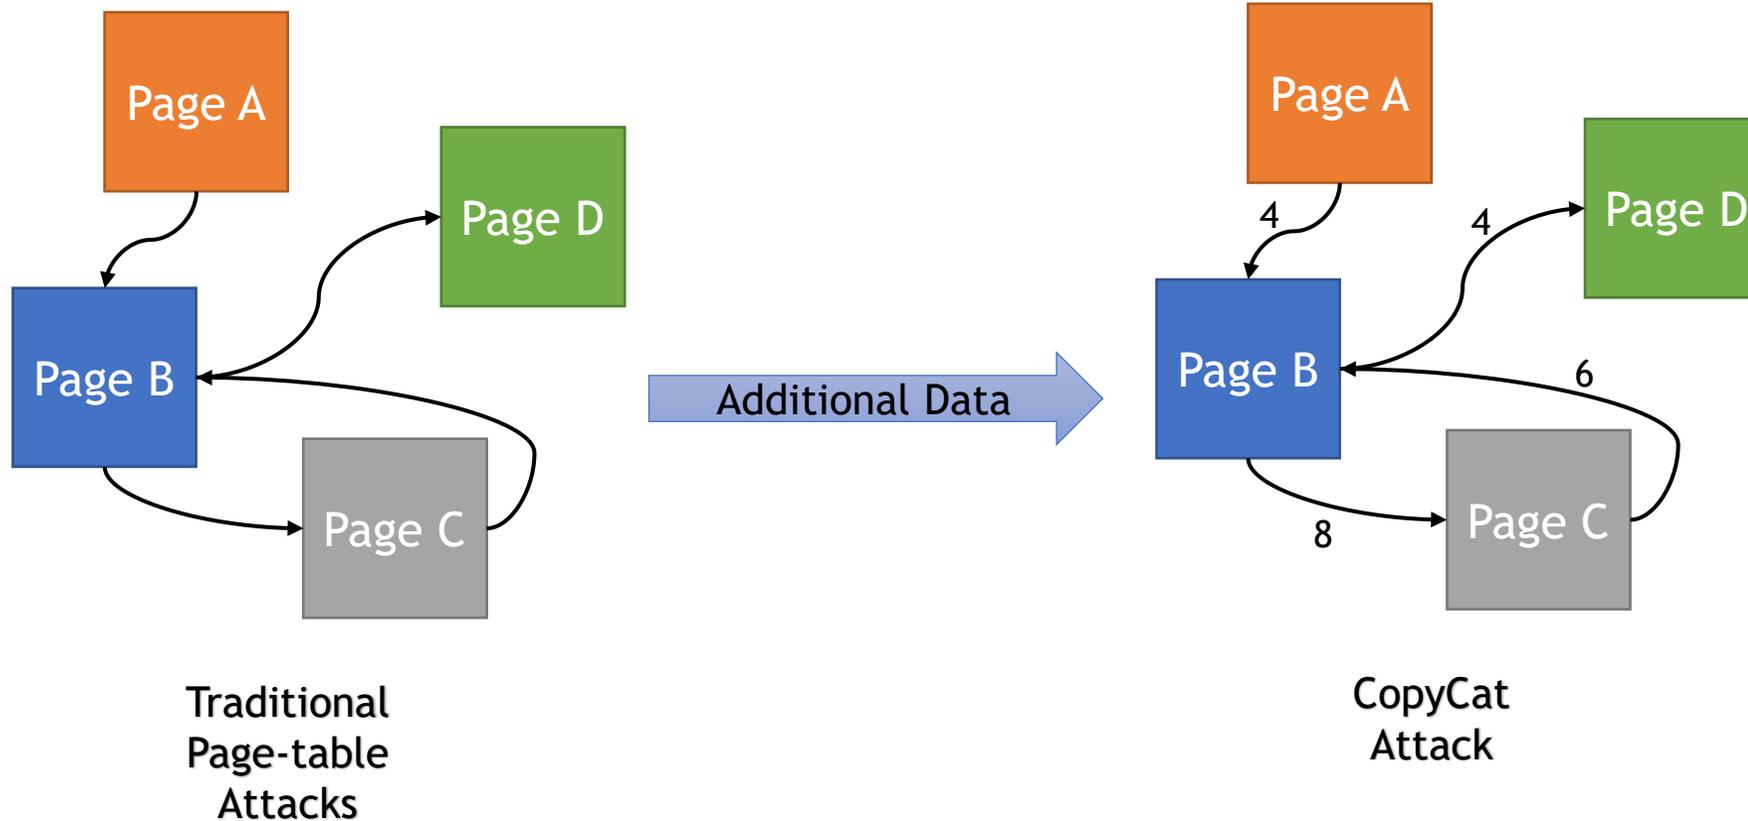
- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions
- Filtering Zeros out: Clear the A bit before, Check the A bit after
- Deterministic Instruction Counting
- Counting from start to end is not useful.
 - A Secondary oracle
 - Page table attack as a deterministic secondary oracle



- Malicious OS controls the interrupt handler
- A threshold to execute 1 or 0 instructions
- Filtering Zeros out: Clear the A bit before, Check the A bit after
- Deterministic Instruction Counting
- Counting from start to end is not useful.
 - A Secondary oracle
 - Page table attack as a deterministic secondary oracle

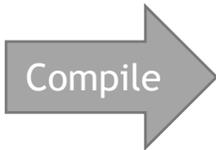


- Previous Controlled Channel attacks leak Page Access Patterns
- CopyCat additionally leaks number of instructions per page



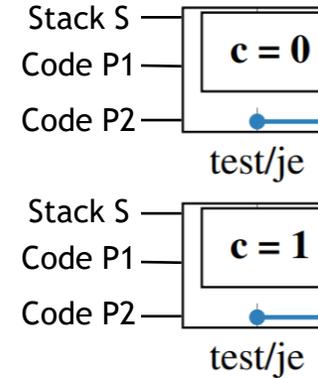
```
if(c == 0) {  
    r = add(r, d);  
}  
else {  
    r = add(r, s);  
}
```

C Code



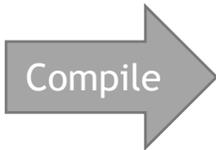
Compile

```
test %eax, %eax  
je label  
mov %edx, %esi  
label:  
call add  
mov %eax, -0xc(%rbp)
```

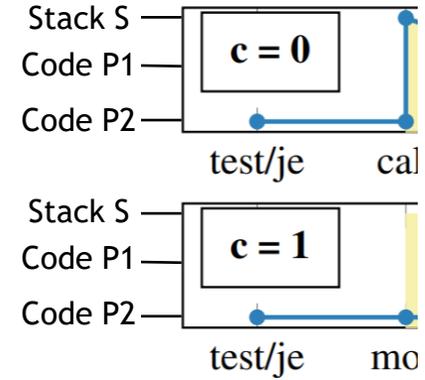


```
if(c == 0) {  
    r = add(r, d);  
}  
else {  
    r = add(r, s);  
}
```

C Code

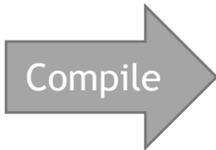


```
test %eax, %eax  
je label  
mov %edx, %esi  
label:  
call add  
mov %eax, -0xc(%rbp)
```

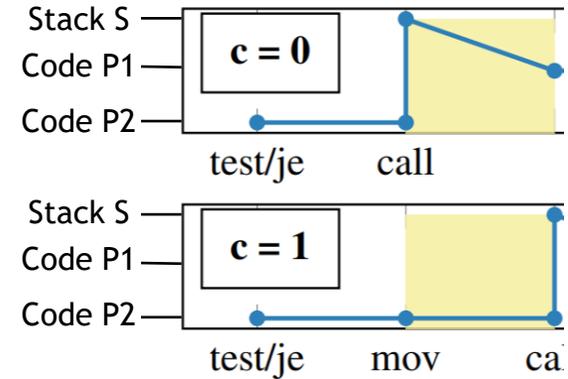


```
if(c == 0) {  
    r = add(r, d);  
}  
else {  
    r = add(r, s);  
}
```

C Code



```
test %eax, %eax  
je label  
mov %edx, %esi  
label:  
call add  
mov %eax, -0xc(%rbp)
```

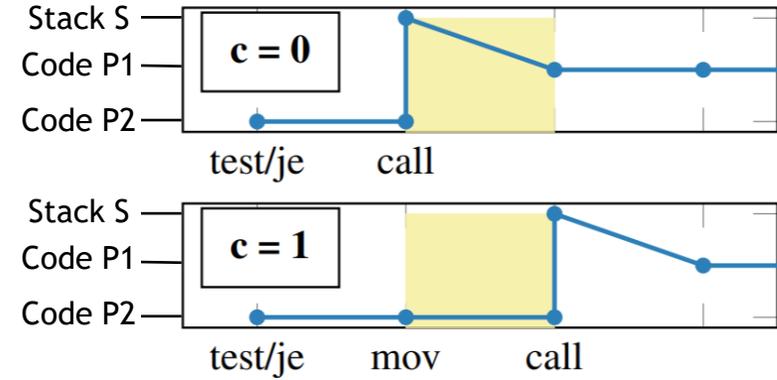


```
if(c == 0) {  
    r = add(r, d);  
}  
else {  
    r = add(r, s);  
}
```

C Code

Compile

```
test %eax, %eax  
je label  
mov %edx, %esi  
label:  
call add  
mov %eax, -0xc(%rbp)
```



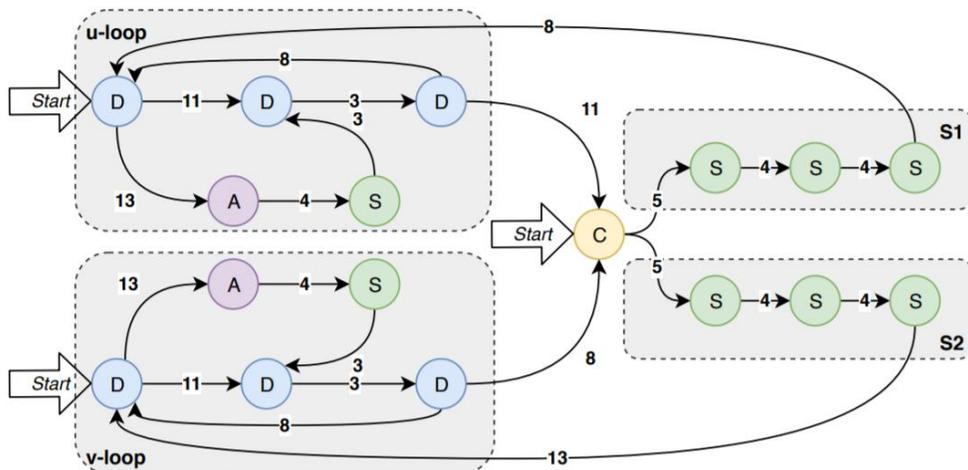


**Crypto means
Crpyptoattacks**

- Previous attacks only leak some of the branches w/ some noise

```
1: procedure MODINV( $u$ , modulus  $v$ )
2:    $b_i \leftarrow 0$   $d_i \leftarrow 1$ ,  $u_i \leftarrow u$ ,  $v_i = v$ ,
3:   while isEven( $u_i$ ) do
4:      $u_i \leftarrow u_i/2$ 
5:     if isOdd( $b_i$ ) then
6:        $b_i \leftarrow b_i - u$ 
7:      $b_i \leftarrow b_i/2$ 
8:   while isEven( $v_i$ ) do
9:      $v_i \leftarrow v_i/2$ 
10:    if isOdd( $d_i$ ) then
11:       $d_i \leftarrow d_i - u$ 
12:     $d_i \leftarrow d_i/2$ 
13:    if  $u_i > v_i$  then
14:       $u_i \leftarrow u_i - v_i$ ,  $b_i \leftarrow b_i - d_i$ 
15:    else
16:       $v_i \leftarrow v_i - u_i$ ,  $d_i \leftarrow d_i - b_i$ 
17:    return  $d_i$ 
```

- Previous attacks only leak some of the branches w/ some noise
- CopyCat synchronously leaks all the branches wo/ any noise



```

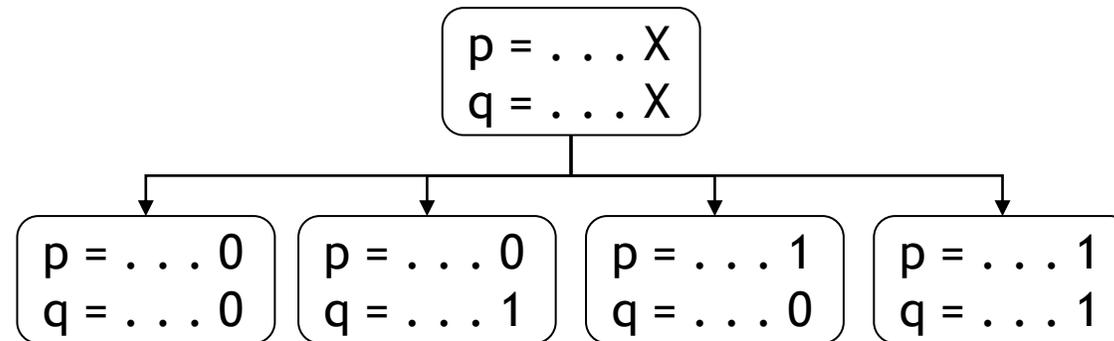
1: procedure MODINV( $u$ , modulus  $v$ )
2:    $b_i \leftarrow 0$   $d_i \leftarrow 1$ ,  $u_i \leftarrow u$ ,  $v_i = v$ ,
3:   while isEven( $u_i$ ) do
4:      $u_i \leftarrow u_i/2$ 
5:     if isOdd( $b_i$ ) then
6:        $b_i \leftarrow b_i - u$ 
7:      $b_i \leftarrow b_i/2$ 
8:   while isEven( $v_i$ ) do
9:      $v_i \leftarrow v_i/2$ 
10:    if isOdd( $d_i$ ) then
11:       $d_i \leftarrow d_i - u$ 
12:     $d_i \leftarrow d_i/2$ 
13:    if  $u_i > v_i$  then
14:       $u_i \leftarrow u_i - v_i$ ,  $b_i \leftarrow b_i - d_i$ 
15:    else
16:       $v_i \leftarrow v_i - u_i$ ,  $d_i \leftarrow d_i - b_i$ 
17:    return  $d_i$ 

```

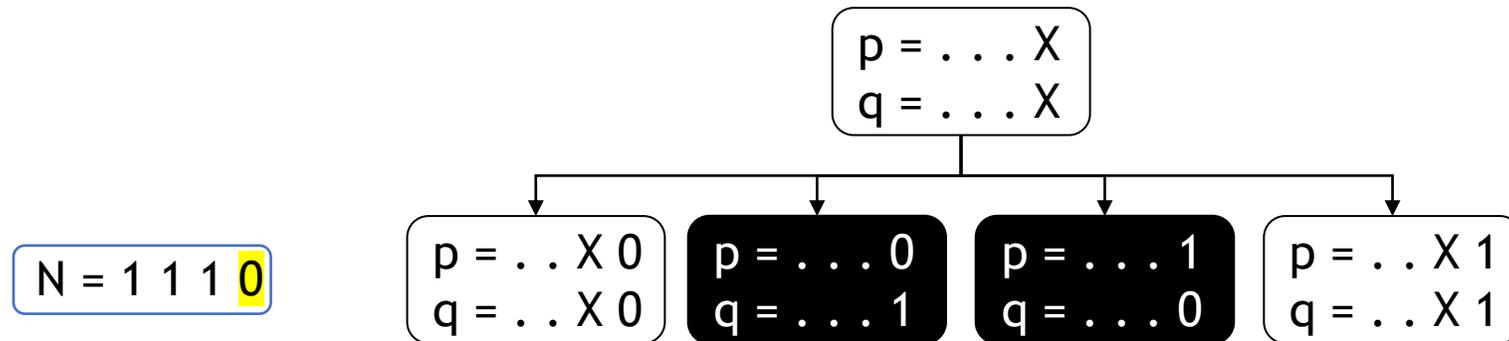
- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x

- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \bmod p$
 - We know that $p \cdot q = N$

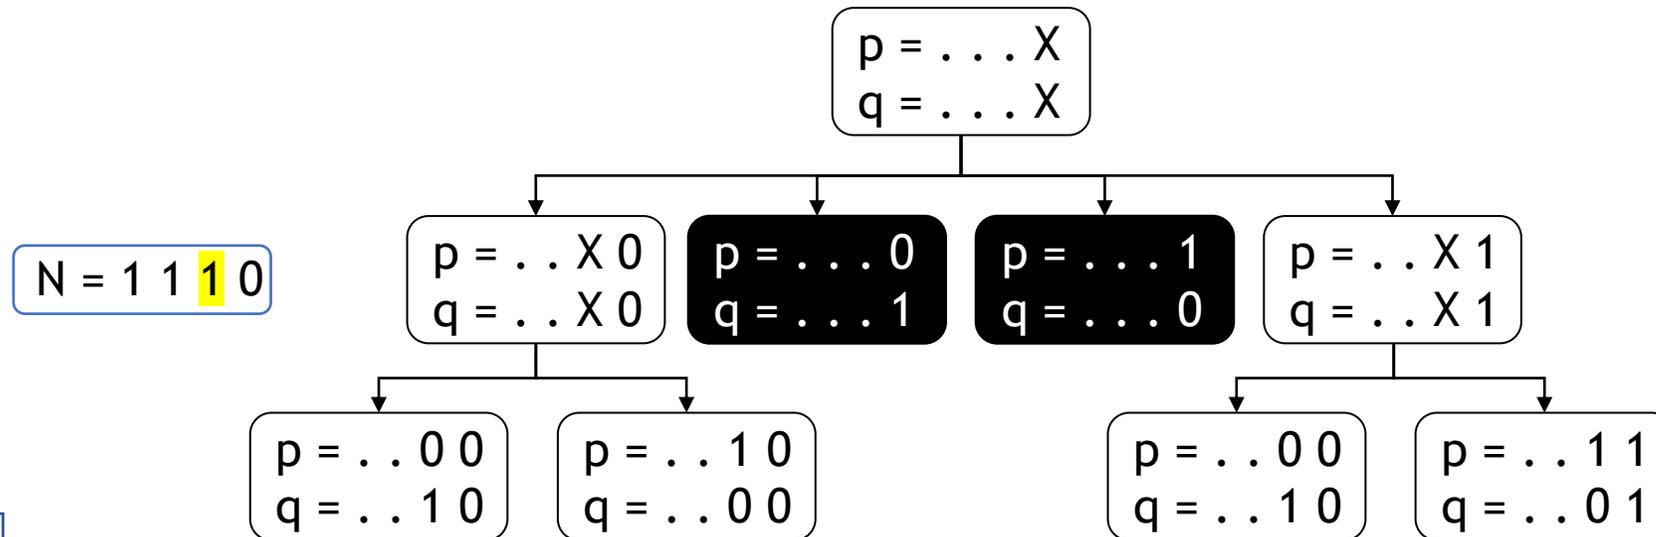
- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \bmod p$
 - We know that $p \cdot q = N$
 - Branch and prune Algorithm with the help of the recovered trace



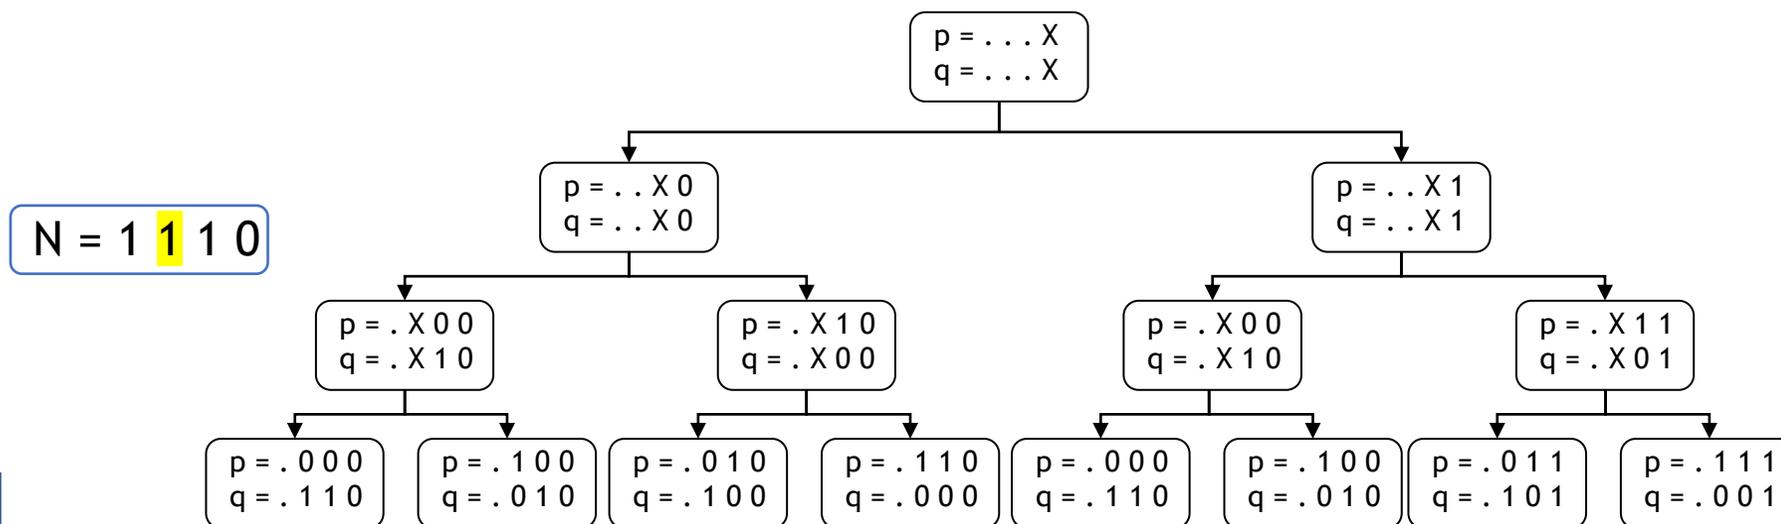
- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \bmod p$
 - We know that $p \cdot q = N$, and N is public
 - Branch and prune Algorithm with the help of the recovered trace



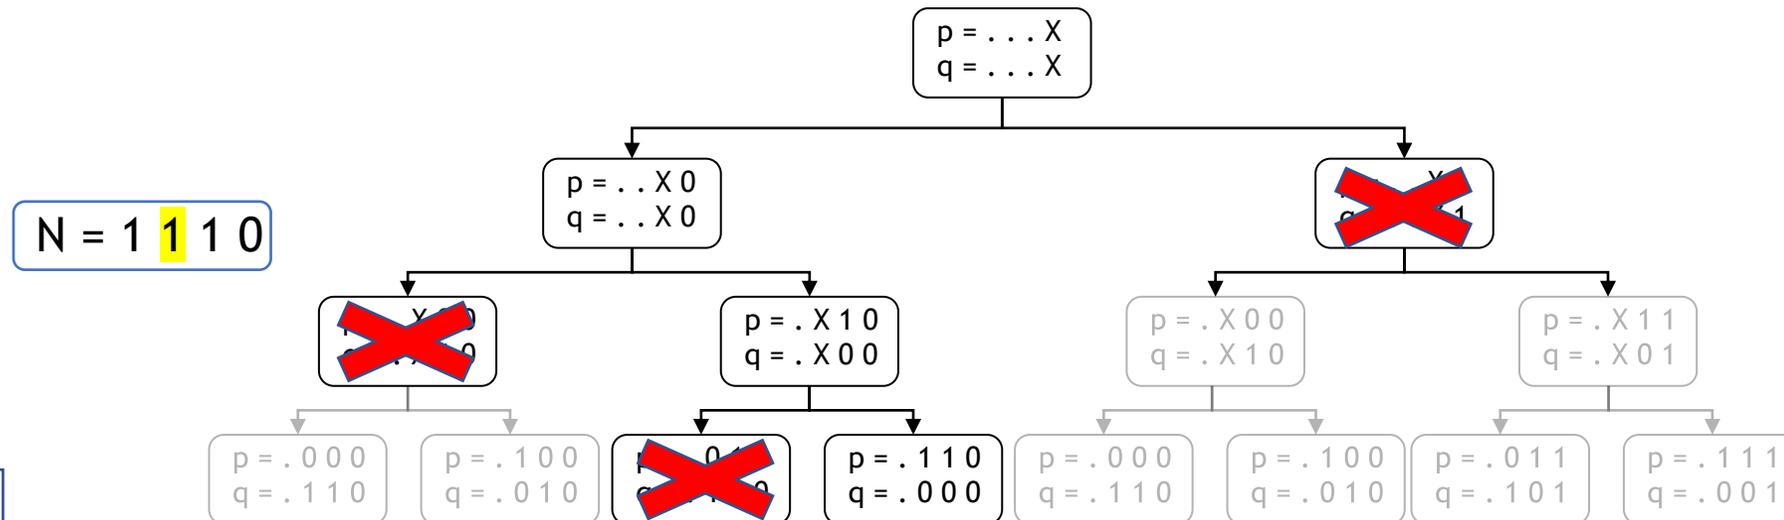
- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \bmod p$
 - We know that $p \cdot q = N$, and N is public
 - Branch and prune Algorithm with the help of the recovered trace



- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \bmod p$
 - We know that $p \cdot q = N$, and N is public
 - Branch and prune Algorithm with the help of the recovered trace



- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \text{ mod } n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \text{ mod } n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \text{ mod } p$
 - We know that $p \cdot q = N$, and N is public
 - Branch and prune Algorithm with the help of the recovered trace



- Single-trace Attack during DSA signing: $k_{inv} = k^{-1} \bmod n$
 - Iterative over the entire recovered trace with n as input $\rightarrow k_{inv}$
 - Plug k_{inv} in $s_1 = k_1^{-1}(h - r_1 \cdot x) \bmod n \rightarrow$ get private key x
- Single-trace Attack during RSA Key Generation: $q_{inv} = q^{-1} \bmod p$
 - We know that $p \cdot q = N$, and N is public
 - Branch and prune Algorithm with the help of the recovered trace
- Single-trace Attack during RSA Key Generation: $d = e^{-1} \bmod \lambda(N)$

- Executed each attack 100 times.
- DSA $k^{-1} \bmod n$
 - Average 22,000 IRQs
 - 75 ms to iterate over an average of 6,320 steps
- RSA $q^{-1} \bmod p$
 - Average 106490 IRQs
 - 365 ms to iterate over an average of 39,400 steps
- RSA $e^{-1} \bmod \lambda(N)$
 - $e^{-1} \bmod \lambda(N)$
 - Average 230,050 IRQs
 - 800ms to iterate over an average of 81,090 steps
- Experimental traces always match the leakage model in all experiments
→ Successful single-trace key recovery

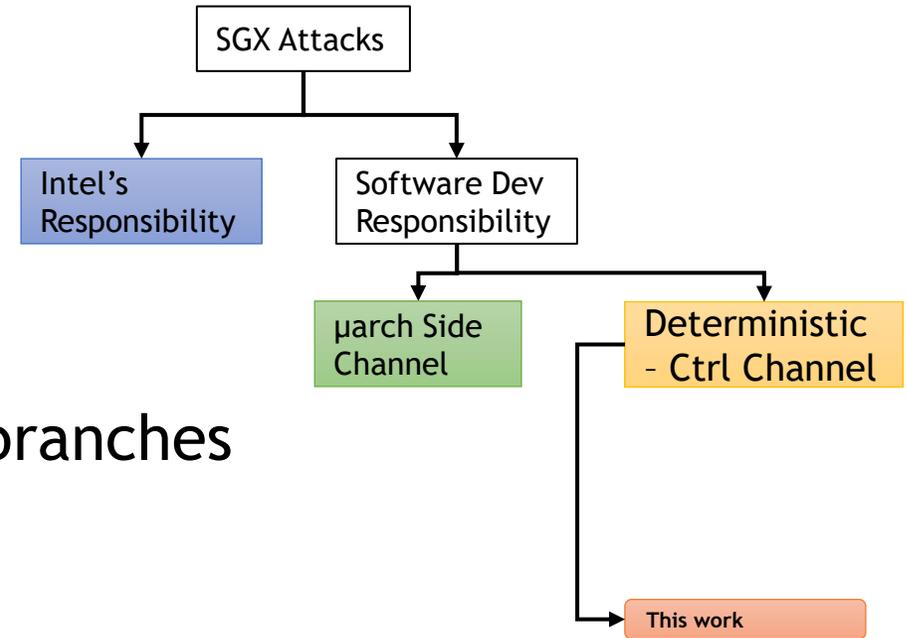
- Libgcrypt uses a variant of BEEA
 - Single trace attack on DSA, Elgamal, ECDSA, RSA Key generation
- OpenSSL uses BEEA for computing GCD
 - Single trace attack on RSA Key generation when computing $\gcd(q - 1, p - 1)$

	Operation (Subroutine)	Implementation	Secret Branch	Exploitable	Computation → Vulnerable Callers	Single-Trace Attack
WolfSSL	Scalar Multiply (wc_ecc_mulmod_ex)	Montgomery Ladder w/ Branches	✓	✓	$(k \times G) \rightarrow \text{wc_ecc_sign_hash}$	✗
	Greatest Common Divisor (fp_gcd)	Euclidean (Divisions)	✓	✗	N/A	N/A
	Modular Inverse (fp_invmod)	BEEA	✓	✓	$(k^{-1} \bmod n) \rightarrow \text{wc_DsaSign}$ $(q^{-1} \bmod p) \rightarrow \text{wc_MakeRsaKey}$ $(e^{-1} \bmod \Lambda(N)) \rightarrow \text{wc_MakeRsaKey}$	✓ ✓ ✓
Libgrypt	Greatest Common Divisor (mpi_gcd)	Euclidean (Divisions)	✓	✗	N/A	N/A
	Modular Inverse (mpi_invmod)	Modified BEEA [43, Vol II, §4.5.2]	✓	✓	$(k^{-1} \bmod n) \rightarrow \{\text{dsa, elgamal}\}.\text{c}::\text{sign_gcry_ecc_ecdsa_sign}$ $(q^{-1} \bmod p) \rightarrow \text{generate_}\{\text{std, fips, x931}\}$ $(e^{-1} \bmod \Lambda(N)) \rightarrow \text{generate_}\{\text{std, fips, x931}\}$	✓ ✓ ✓
OpenSSL	Greatest Common Divisor (BN_gcd)	BEEA	✓	✓	$\gcd(q - 1, p - 1) \rightarrow \text{RSA_X931_derive_ex}$	✓
	Modular Inverse (BN_mod_inverse_no_branch)	BEEA w/ Branches	✗	N/A	N/A	N/A
IPP Crypto	Greatest Common Divisor (ippsGcd_BN)	Modified Lehmer's GCD	✓	?	$\gcd(q - 1, e) \rightarrow \text{cpIsCoPrime}$	N/A
	Modular Inverse (cpModInv_BNU)	Euclidean (Divisions)	✓	✗	$\gcd(p - 1, q - 1) \rightarrow \text{isValidPriv1_rsa}$ N/A	N/A N/A

- WolfSSL fixed the issues in 4.3.0 and 4.4.0
 - Blinding for $k^{-1} \bmod n$ and $e^{-1} \bmod \lambda(N)$
 - Alternate formulation for $q^{-1} \bmod p$: $q^{p-2} \bmod p$
 - Using a constant-time (branchless) modular inverse [11]
- Libgcrypt fixed the issues in 1.8.6
 - Using a constant-time (branchless) modular inverse [11]
- OpenSSL fixed the issue in 1.1.1e
 - Using a constant-time (branchless) GCD algorithm [11]

[11] Bernstein, Daniel J., and Bo-Yin Yang. "Fast constant-time gcd computation and modular inversion." *CHES 2019*.

- **Instruction Level Granularity**
 - Imbalance number of instructions
 - Leak the outcome of branches
- **Fully Deterministic and reliable**
 - Millions of instructions tested
 - Attacks match the exact leakage model of branches
- **Easy to scale and replicate**
 - No reverse engineering of branches and microarchitectural components
 - Tracking all the branches synchronously
- **Branchless programming is hard!**





Daniel Moghimi

@danielmgmi



<https://github.com/jovanbulck/sgx-step>