## CopyCat: Controlled Instruction-Level Attacks on Enclaves

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





Intel Labs -Sept. 10 2020

#### OS/Hypervisor Security Model



Traditional Security Model

### Trusted Execution Environment (TEE) - Intel SGX

• Intel Software Guard eXtensions (SGX)



## Trusted Execution Environment (TEE) - Intel 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



Traditional Security Model

## Trusted Execution Environment (TEE) - Intel 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
- Protects against system level adversary

#### New Attacker Model:

Attacker gets full control over OS



Traditional Security Model

#### • 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



#### • 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



#### • 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
- µarch Side Channel
  - Constant-time Coding
  - Flushing and Isolating buffers
  - Probabilistic



#### • 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
- µarch Side Channel
  - Constant-time Coding
  - Flushing and Isolating buffers
  - Probabilistic
- Deterministic Attacks
  - Page Fault, A/D Bit, etc. (4kB Granularity)





[6] Evtyushkin, 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 S&P 2015.
[10] Wang, Wenhao, et al. "Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX." ACM CCS 2017.

• Malicious OS controls the interrupt handler

NOP	ADD	XOR	MUL	DIV	ADD	MUL	NOP	NOP	
Enclave Execution Thread Starts									Time

• Malicious OS controls the interrupt handler



• 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



- 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







- 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

Target Code Page									
CALL	ADD	XOR	MUL	PUSH	ADD	MUL	MOV	NOP	

Time

- 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



Page-table Attacks

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





 $\mathbf{c} = \mathbf{0}$ 

test/je

**c** = 1

test/je

C Code



31



call

mov

cal





C Code



C Code





# Crypto means Crpyptoattacks

## Binary Extended Euclidean Algorithm (BEEA)

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

1: **procedure** MODINV(*u*, modulus *v*)

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

return d<sub>i</sub>

## Binary Extended Euclidean Algorithm

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



1:	pro	<b>cedure</b> MODINV( <i>u</i> , modulus <i>v</i> )
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 <i>d</i> :

## CopyCat on WolfSSL

• Translate instruction Counts to Basic Block Transitions

11,3,8,5,4,4,13,11,3,8,5,4,4,8,11,3,8,11,3,8,13,4,3,3,8,11,3,11,5,4,4

• Translate instruction Counts to Basic Block Transitions

 11,3,8,5,4,4,13,11,3,8,5,4,4,8,11,3,8,11,3,8,13,4,3,3,8,11,3,11,5,4,4

 DDD
 8
 CSSS
 13
 DDD
 8
 CSSS
 8
 DDD
 8
 DASDD
 8
 DDD
 11
 CSSS

**Rule 1:** 
$$? \xrightarrow{11} ? \xrightarrow{3} ? = D \rightarrow D \rightarrow D$$
.  
**Rule 2:**  $? \xrightarrow{13} ? \xrightarrow{4} ? \xrightarrow{3} ? \xrightarrow{3} ? = D \rightarrow A \rightarrow S \rightarrow D \rightarrow D$ .  
**Rule 3:**  $? \xrightarrow{5} ? \xrightarrow{4} ? \xrightarrow{4} ? \xrightarrow{4} ? = C \rightarrow S \rightarrow S \rightarrow S$ .

#### CopyCat on WolfSSL

• Translate instruction Counts to Basic Block Transitions



**Rule 1:** 
$$? \xrightarrow{11} ? \xrightarrow{3} ? = D \rightarrow D \rightarrow D$$
.  
**Rule 2:**  $? \xrightarrow{13} ? \xrightarrow{4} ? \xrightarrow{3} ? \xrightarrow{3} ? = D \rightarrow A \rightarrow S \rightarrow D \rightarrow D$ .  
**Rule 3:**  $? \xrightarrow{5} ? \xrightarrow{4} ? \xrightarrow{4} ? = C \rightarrow S \rightarrow S \rightarrow S$ .  
**Rule 4:**  $S? \xrightarrow{13} ? = S2 \rightarrow v$ -loop.  
**Rule 5:**  $S? \xrightarrow{8} ? = S1 \rightarrow u$ -loop.

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

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

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



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



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



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



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



- Single-trace Attack during DSA signing:  $k_{inv} = k^{-1} \mod n$ 
  - Iterative over the entire recovered trace with n as input  $\rightarrow k_{inv}$
  - Plug  $k_{inv}$  in  $s_1 = k_1^{-1}(h r_1, x) \mod n \rightarrow \text{get private key } x$
- Single-trace Attack during RSA Key Generation:  $q_{inv} = q^{-1} \mod p$ 
  - We know that  $\mathbf{p}.\mathbf{q} = \mathbf{N}$ , and  $\mathbf{N}$  is public
  - Branch and prune Algorithm with the help of the recovered trace
- Single-trace Attack during RSA Key Generation:  $d = e^{-1} \mod \lambda(N)$ 
  - Similar attack but instead use  $\lambda(N) = \frac{(p-1)(q-1)}{2^i}$
  - Only 81% of the keys have the above property
  - It works even on a hardcoded and big value for e, i.e.  $e \neq 65537$

## CopyCat on WolfSSL - Cryptanalysis Results

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

## CopyCat - Bypassing ECDSA Timing Countermeasure

```
int wc ecc mulmod ex(mp int* k, ecc point *G, ecc point *R, mp int* a, mp int
          * modulus, int map, void* heap) { ...
2 for (;;) {
  if (--bitcnt == 0) { /* grab next digit as required */
    if (digidx == -1) {
       break:
 5
    buf = get_digit(k, digidx);
    bitcnt = (int)DIGIT_BIT;
    --digidx;
Q
10
  i = (buf >> (DIGIT BIT - 1)) \& 1; /* grab the next msb from the multiplicand */
11
  buf <<= 1;
12
  if (mode == 0) {
13
    mode = i; /* timing resistant – dummy operations */
14
    err = ecc_projective_add_point(M[1], M[2], M[2], a, modulus, mp);...
15
    err = ecc_projective_dbl_point(M[2], M[3], a, modulus, mp);...
16
17
  }...
  err = ecc_projective_add_point(M[0], M[1], M[i^1], a, modulus, mp);...
18
  err = ecc_projective_dbl_point(M[2], M[2], a, modulus, mp);...
19
  } /* end for */...}
20
```

Table 2: Minimum number of signature samples for each bias class to reach 100% recovery success for the lattice-based key recovery on wc\_ecc\_mulmod\_ex of ECDSA, with lattice reduction time L-TIME and trace collection time T-TIME.

LZBS	DIM	L-TIME	SIGNATURES	IRQs	<b>T-TIME</b>
4	75	30 sec	1,200	3.9M	13.3 sec
5	58	5 sec	1,856	6.0M	20.4 sec
6	46	3 sec	2,944	9.6M	33.7 sec
7	42	2 sec	5,376	17.5M	1 min

### How about other Crypto libraries?

- 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)
- There is still lots of other cases of micro leakages due to usage of branches, e.g. Intel IPP Crypto lehmer's GCD with optimizations

	Operation (Subroutine)	Implementation	Secret Branch	Exploitable	$\textbf{Computation} \rightarrow \textbf{Vulnerable Callers}$	Single-Trace Attack
	Scalar Multiply (wc_ecc_mulmod_ex)	Montgomery Ladder w/ Branches		· · · · · · · · · · · · · · · · · · ·	$(k \times G) \rightarrow wc\_ecc\_sign\_hash$	×
.cl	Greatest Common Divisor (fp_gcd)	Euclidean (Divisions)	~	×	N/A	N/A
WolfSSE	Modular Inverse (fp_invmod)	BEEA	~	v	$(k^{-1} \mod n) \to \operatorname{wc}_{DsaSign}$ $(q^{-1} \mod p) \to \operatorname{wc}_{MakeRsaKey}$ $(e^{-1} \mod \Lambda(N)) \to \operatorname{wc}_{MakeRsaKey}$	
	Greatest Common Divisor (mpi_gcd)	Euclidean (Divisions)	~	×	N/A	N/A
Libgcrypt	Modular Inverse (mpi_invm)	Modified BEEA [43, Vol II, §4.5.2]	· · · ·	· · · ·	$ \begin{array}{l} (k^{-1} \mod n) \to \{ \texttt{dsa,elgamal} \}.\texttt{c::sign,gcry_ecc\_ecdsa\_sign} \\ (q^{-1} \mod p) \to \texttt{generate}_\{\texttt{std,fips,x931}\} \\ (e^{-1} \mod \Lambda(N)) \to \texttt{generate}_\{\texttt{std,fips,x931}\} \end{array} $	
OpenSSL	Greatest Common Divisor (BN_gcd) Modular Inverse (BN_mod_inverse_no_branch)	BEEA BEEA w/ Branches	<mark>×</mark>		$gcd(q-1,p-1) \rightarrow RSA_X931\_derive\_ex$ N/A	<b>v</b> N/A
TOP CTYPTO	Greatest Common Divisor (ippsGcd_BN)	Modified Lehmer's GCD	~	?	$gcd(q-1,e)  ightarrow  ext{cpIsCoPrime} \ gcd(p-1,q-1)  ightarrow  ext{isValidPriv1_rsa}$	N/A N/A
W.	Modular Inverse (cpModInv_BNU)	Euclidean (Divisions)		×	N/A	<u></u> N/A

#### Responsible Disclosure

- WolfSSL fixed the issues in 4.3.0 and 4.4.0
  - Blinding for  $k^{-1} \mod n$  and  $e^{-1} \mod \lambda(N)$
  - Alternate formulation for  $q^{-1} \mod p$ :  $q^{p-2} \mod p$
  - 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]

## Interrupt Driven Attacks and Single Stepping

- Amplifying Transient Execution Attacks
  - Foreshadow, ZombieLoad, LVI, CrossTalk
- Amplifying Microarchitectural Side Channels
  - CacheZoom, BranchScope, Branch Shadowing, Bluethunder, etc.
- Interrupt Latency as a Side Channel
  - Nemesis, Frontal Attack
- CopyCat: Deterministic Instruction
   Counting as a Side Channel



Title	Publication details	Source code	SGX-Step features used
CrossTalk: Speculative Data Leaks Across Cores Are Real	S&P21	-	Single-stepping, page fault
Frontal Attack: Leaking Control-Flow in SGX via the CPU Frontend	arXiv20	-	Single-stepping interrupt latency, PTE A/D
From A to Z: Projective coordinates leakage in the wild	CHES20	-	Page fault
LVI: Hijacking Transient Execution through Microarchitectural Load Value Injection	S&P20	link	Single-stepping, page-table manipulation
CopyCat: Controlled Instruction-Level Attacks on Enclaves	USEC20	-	Single-stepping, page fault, PTE A/D
When one vulnerable primitive turns viral: Novel single- trace attacks on ECDSA and RSA	CHES20	-	Single-stepping, page fault, PTE A/D
Big Numbers - Big Troubles: Systematically Analyzing Nonce Leakage in (EC)DSA Implementations	USEC20	-	Page fault
Plundervolt: Software-based Fault Injection Attacks against Intel SGX	S&P20	link	Privileged interrupt/call gates, MSR
Bluethunder: A 2-level Directional Predictor Based Side- Channel Attack against SGX	CHES20	-	Single-stepping
Fallout: Leaking Data on Meltdown-resistant CPUs	CCS19	-	PTE A/D
A Tale of Two Worlds: Assessing the Vulnerability of Enclave Shielding Runtimes	CCS19	link	Single-stepping, page fault, PTE A/D
ZombieLoad: Cross-Privilege-Boundary Data Sampling	CCS19	link	Single-stepping, zero- stepping, page-table manipulation
SPOILER: Speculative Load Hazards Boost Rowhammer and Cache Attacks	USEC19	-	Single-stepping interrupt latency
Nemesis: Studying Microarchitectural Timing Leaks in Rudimentary CPU Interrupt Logic	CCS18	link	Single-stepping interrupt latency, page fault, PTE A/D
Foreshadow: Extracting the Keys to the Intel SGX Kingdom with Transient Out-of-Order Execution	USEC18	link	Single-stepping, zero- stepping, page-table manipulation
Single Trace Attack Against RSA Key Generation in Intel SGX SSL	AsiaCCS18	-	Page fault
Off-Limits: Abusing Legacy x86 Memory Segmentation to Spy on Enclaved Execution	ESSoS18	link	Single-stepping, IA32 segmentation, page fault
SGX-Step: A Practical Attack Framework for Precise	SysTEX17	link	Single-stepping, page fault,

	Attack	Code/Data	Granularity	Noise
μ-arch contention	DRAM row buffer conflicts [74] PRIME+PROBE cache conflicts [15, 30, 47, 58] Read-after-write false dependencies [46] Branch prediction history buffers [24, 34, 44] Interrupt latency [71] Port contention [3]	Code + data Code + data Data Code Code + data Code	<ul> <li>✗ Low (1-8 KiB)</li> <li>✗ Med (64-512 B cache line/set)</li> <li>✓ High (4 B)</li> <li>✓ High (branch instruction)</li> <li>✓ High (instruction latency class)</li> <li>✓ High (µ-op execution port)</li> </ul>	<ul> <li>× High</li> <li>~ Med</li> <li>× High</li> <li>~ Low</li> <li>× High</li> <li>× High</li> <li>× High</li> </ul>
Ctrl channel	Page faults [80] and page table A/D bits [72, 74] IA-32 segmentation faults [29] Page table FLUSH+RELOAD [72] COPYCAT	Code + data Code + data Code + data <b>Code</b>	<ul> <li>X Low (4 KiB)</li> <li>X Low/high (4 KiB; 1 B for enclaves ≤ 1 MiB)</li> <li>X Low (32 KiB)</li> <li>✓ High (instruction)</li> </ul>	<ul> <li>Deterministic</li> <li>Deterministic</li> <li>Low</li> <li>Deterministic</li> </ul>

• Some do not work when hyper-threadnig is disabled (Strong TCB of Intel SGX)

	Attack	Code/Data	Granularity	Noise
μ-arch contention	DRAM row buffer conflicts [74] PRIME+PROBE cache conflicts [15, 30, 47, 58] Read-after-write false dependencies [46] Branch prediction history buffers [24, 34, 44] Interrupt latency [71] Port contention [3]	Code + data Code + data Data Code Code + data Code	<ul> <li>✗ Low (1-8 KiB)</li> <li>✗ Med (64-512 B cache line/set)</li> <li>✓ High (4 B)</li> <li>✓ High (branch instruction)</li> <li>✓ High (instruction latency class)</li> <li>✓ High (µ-op execution port)</li> </ul>	<ul> <li>× High</li> <li>~ Med</li> <li>× High</li> <li>~ Low</li> <li>× High</li> <li>× High</li> <li>× High</li> </ul>
Ctrl channel	Page faults [80] and page table A/D bits [72, 74] IA-32 segmentation faults [29] Page table FLUSH+RELOAD [72] COPYCAT	Code + data Code + data Code + data <b>Code</b>	<ul> <li>X Low (4 KiB)</li> <li>X Low/high (4 KiB; 1 B for enclaves ≤ 1 MiB)</li> <li>X Low (32 KiB)</li> <li>✓ High (instruction)</li> </ul>	<ul> <li>Deterministic</li> <li>Deterministic</li> <li>Low</li> <li>Deterministic</li> </ul>

- Some do not work when hyper-threadnig is disabled (Strong TCB of Intel SGX)
- Some can be mitigated by **flushing/isolating** microarchitectural buffers.

	Attack	Code/Data	Granularity	Noise
µ-arch contention	DRAM row buffer conflicts [74] PRIME+PROBE cache conflicts [15, 30, 47, 58] Read-after-write false dependencies [46] Branch prediction history buffers [24, 34, 44] Interrupt latency [71] Port contention [3]	Code + data Code + data Data Code Code + data Code	<ul> <li>✗ Low (1-8 KiB)</li> <li>✗ Med (64-512 B cache line/set)</li> <li>✓ High (4 B)</li> <li>✓ High (branch instruction)</li> <li>✓ High (instruction latency class)</li> <li>✓ High (µ-op execution port)</li> </ul>	<ul> <li>× High</li> <li>~ Med</li> <li>× High</li> <li>~ Low</li> <li>× High</li> <li>× High</li> <li>× High</li> </ul>
Ctrl channel	Page faults [80] and page table A/D bits [72, 74] IA-32 segmentation faults [29] Page table FLUSH+RELOAD [72] COPYCAT	Code + data Code + data Code + data <b>Code</b>	<ul> <li>X Low (4 KiB)</li> <li>X Low/high (4 KiB; 1 B for enclaves ≤ 1 MiB)</li> <li>X Low (32 KiB)</li> <li>✓ High (instruction)</li> </ul>	<ul> <li>Deterministic</li> <li>Deterministic</li> <li>Low</li> <li>Deterministic</li> </ul>

- Some do not work when hyper-threadnig is disabled (Strong TCB of Intel SGX)
- Some can be mitigated by **flushing/isolating** microarchitectural buffers.
- Some only apply to legacy enclave (32-bit)

	Attack	Code/Data	Granularity	Noise
µ-arch contention	DRAM row buffer conflicts [74] PRIME+PROBE cache conflicts [15, 30, 47, 58] Read-after-write false dependencies [46] Branch prediction history buffers [24, 34, 44] Interrupt latency [71] Port contention [3]	Code + data Code + data Data Code Code + data Code	<ul> <li>✗ Low (1-8 KiB)</li> <li>✗ Med (64-512 B cache line/set)</li> <li>✓ High (4 B)</li> <li>✓ High (branch instruction)</li> <li>✓ High (instruction latency class)</li> <li>✓ High (µ-op execution port)</li> </ul>	<ul> <li>× High</li> <li>~ Med</li> <li>× High</li> <li>~ Low</li> <li>× High</li> <li>× High</li> <li>× High</li> </ul>
Ctrl channel	Page faults [80] and page table A/D bits [72, 74] IA-32 segmentation faults [29] Page table FLUSH+RELOAD [72] COPYCAT	Code + data Code + data Code + data <b>Code</b>	<ul> <li>X Low (4 KiB)</li> <li>X Low/high (4 KiB; 1 B for enclaves ≤ 1 MiB)</li> <li>X Low (32 KiB)</li> <li>✓ High (instruction)</li> </ul>	<ul> <li>Deterministic</li> <li>Deterministic</li> <li>Low</li> <li>Deterministic</li> </ul>

- Some do not work when hyper-threadnig is disabled (Strong TCB of Intel SGX)
- Some can be mitigated by **flushing/isolating** microarchitectural buffers.
- Some only apply to legacy enclave (32-bit)
- Some are limited to be applied synchronously.

	Attack	Code/Data	Granularity	Noise
μ-arch contention	DRAM row buffer conflicts [74] PRIME+PROBE cache conflicts [15, 30, 47, 58] Read-after-write false dependencies [46] Branch prediction history buffers [24, 34, 44] Interrupt latency [71] Port contention [3]	Code + data Code + data Data Code Code + data Code	<ul> <li>✗ Low (1-8 KiB)</li> <li>✗ Med (64-512 B cache line/set)</li> <li>✓ High (4 B)</li> <li>✓ High (branch instruction)</li> <li>✓ High (instruction latency class)</li> <li>✓ High (µ-op execution port)</li> </ul>	<ul> <li>✗ High</li> <li>∼ Med</li> <li>✗ High</li> <li>∼ Low</li> <li>✗ High</li> <li>✗ High</li> <li>✗ High</li> </ul>
Ctrl channel	Page faults [80] and page table A/D bits [72, 74] IA-32 segmentation faults [29] Page table FLUSH+RELOAD [72] COPYCAT	Code + data Code + data Code + data <b>Code</b>	<ul> <li>X Low (4 KiB)</li> <li>X Low/high (4 KiB; 1 B for enclaves ≤ 1 MiB)</li> <li>X Low (32 KiB)</li> <li>✓ High (instruction)</li> </ul>	<ul> <li>Deterministic</li> <li>Deterministic</li> <li>Low</li> <li>Deterministic</li> </ul>

## CopyCat and Macro-fusion

- Fused instructions are counted as one.
- Confirm/RE of the behavior of macro-fusion on Intel CPUs
- Macro-fusion is dependent on the program layout  $\rightarrow$  deterministic
  - The offset of a *cmp+branch* within a cache line
  - True when hyperthreading is disabled (Intel SGX TCB)

Macro-Fusibility								
Instruction	TEST	CMP	AND	ADD	SUB	INC	DEC	
J0/JN0	$\checkmark$	X	$\checkmark$	X	X	X	X	
JC/JB/JAE/JNB	$\checkmark$	$\checkmark$	√	$\checkmark$	✓	X	X	
JE/JZ/JNE/JNZ	$\checkmark$	$\checkmark$	✓	✓	✓	✓	√	
JNA/JBE/JA/JNBE	$\checkmark$	$\checkmark$	✓	$\checkmark$	✓	X	X	
JS/JNS/JP/JPE/JNP/JP0	$\checkmark$	X	√	X	X	X	X	
JL/JNGE/JGE/JNL/JLE/JNG/JG/JNLE	$\checkmark$	~	$\checkmark$	✓	$\checkmark$	✓	√	

https://en.wikichip.org/wiki/macro-operation\_fusion

- Instruction Level Granularity
  - Imbalance number of instructions
  - Leak the outcome of branches



- 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



- 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



- 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!



## Future Directions - Other TEE Models

- Virtual Machine TEE
  - AMD SEV
  - Intel TDX
- What are other ways to interrupt a TEE in the above models?
- What is the impact?
  - Guest OSS
  - Cryptographic Services
  - Other Applications





HOST-OS/VMM

## Future Directions - Non-cryptographic Application of Enclaves

- Data-dependent secret-processing applications
  - Confidential Deep Learning (
  - Trusted Database (EnclaveDB)
- Automated Leakage Analysis and Exploit Generation
  - Fuzzing and Taint Analysis
  - Dynamic Analysis



Pin?



#### Future Directions - Mitigation

- Compiler-based Solutions
  - Balancing secret-dependent branches with dummy instructions
- System-level Mitigation
  - Self-paging Enclave (Autarky)



**Figure 2.** Autarky enforces invocation of an enclave's selfpaging handler on each page fault.

## Questions?!



Daniel Moghimi @danielmgmi



https://github.com/j ovanbulck/sgx-step