



The Promises and Pitfalls of Hardware-Assisted Security

Alexandra Dmitrienko Julius-Maximilians-Universität Würzburg

alexandra.dmitrienko@uni-wuerzburg.de

SEPTEMBER 9 - 13, 2019

CROSSING Summer School on Sustainable Security & Privacy

The Great Promise of Trusted Computing



Historical Overview: Deployed Systems

1970	1980		1990	2000		2010	
Cambridge CAP	VAX/V	MS		Trusted Platform Module (TPM)	PUFs		
			Java security architecture	Late launch	ו/TXT	TPM 2.0	
Protection rings						Intel SGX	
Reference monitor Secure		are-assisted boot					

Computer security Mobile security Smart card security

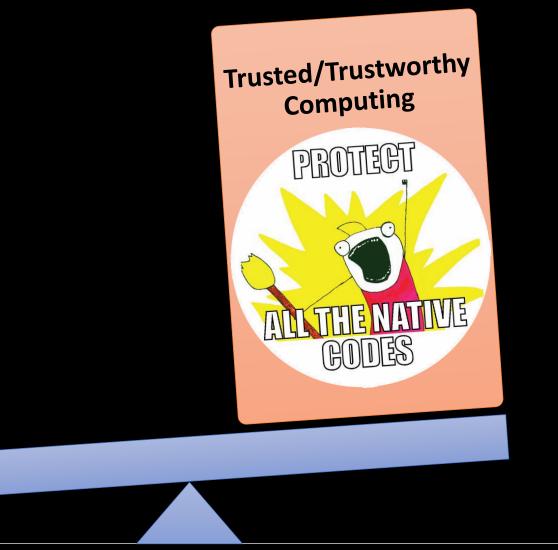
Historical Overview: Deployed Systems

1970	198	30 1990	2000	2010	
Cambridge CAP	VAX/V	'MS	Trusted Platform Module (TPM)		
	Simple sm cards	art Java security architecture	Late launch/TXT	TPM 2.0	
	Protectior	n rings		Intel	SGX
Reference	monitor	Hardware-assisted secure boot			
	J	Java Card platform			
					Computer security Mobile security Smart card security

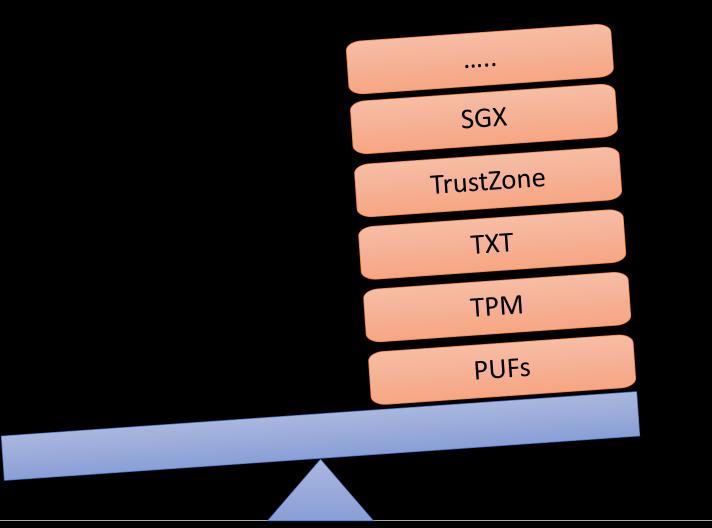
Historical Overview: Deployed Systems

1970	1980		1990	2000			2010		
Cambridge CAP	VAX/VMS				d Platform le (TPM)	PUFs	GP TEE	standards	
	Simple sm cards	art	Java security architecture		Late launc	h/TXT	TPM 2.	0	
	Protection	n rings	TI M-	Shield	ARM TrustZone	On-bo e Crede		Intel SGX	
Reterence monitor			ware-assisted e boot	Mobile hardware security architectures Mobile OS security architectures					
	J	ava Card platform							
						e Truste le (MTN			
					Wiedd		1)	Μ	omputer security obile security nart card security

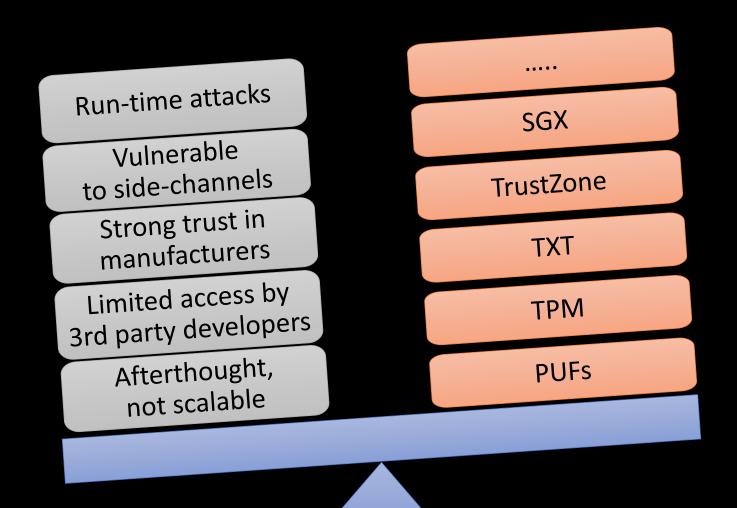
Trusted Computing under Attack



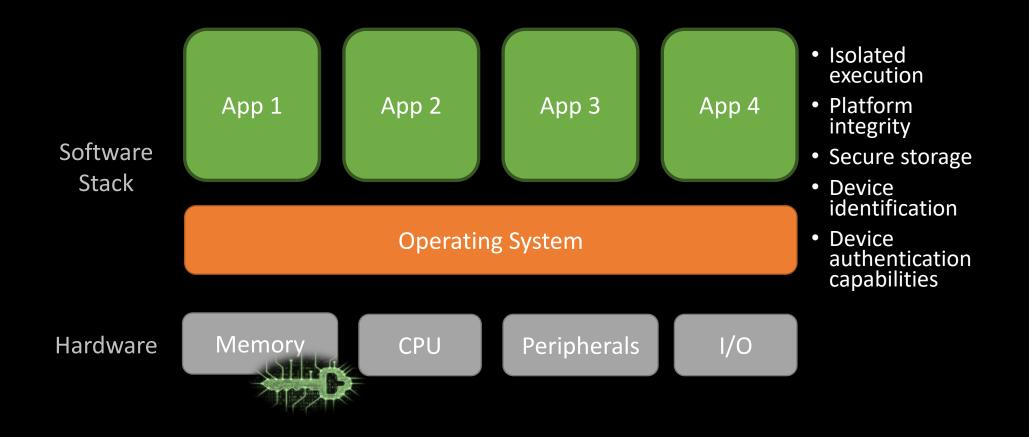
Trusted Computing under Attack

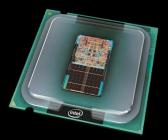


Trusted Computing under Attack

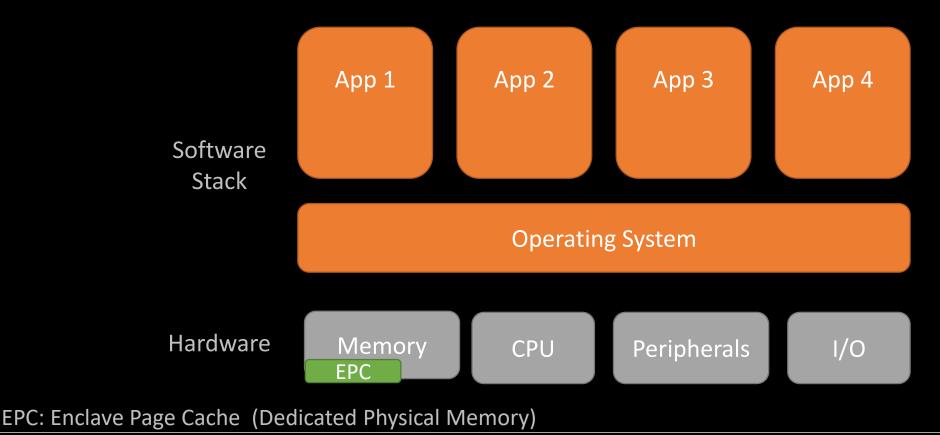


Goal: Self-Contained Security

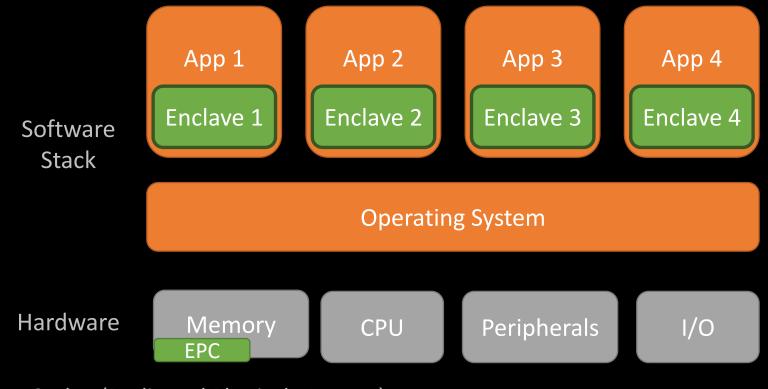




Intel SGX

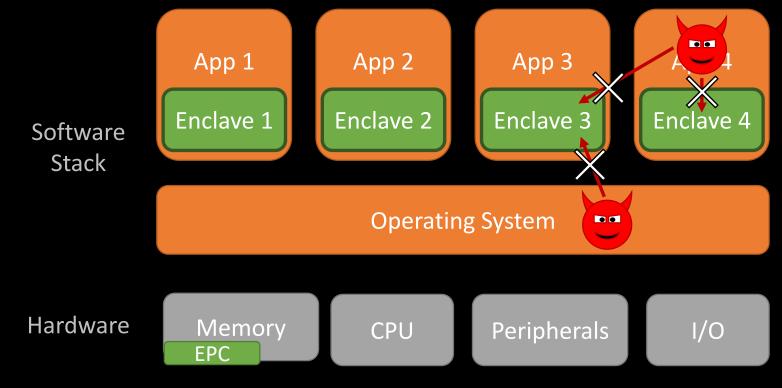


- OS creates and manages enclaves, allocates memory from Enclave Page Cache (EPC)
- OS maps physical to virtual memory, as well as loads data and code into enclave
- Trust assumptions: All software components untrusted



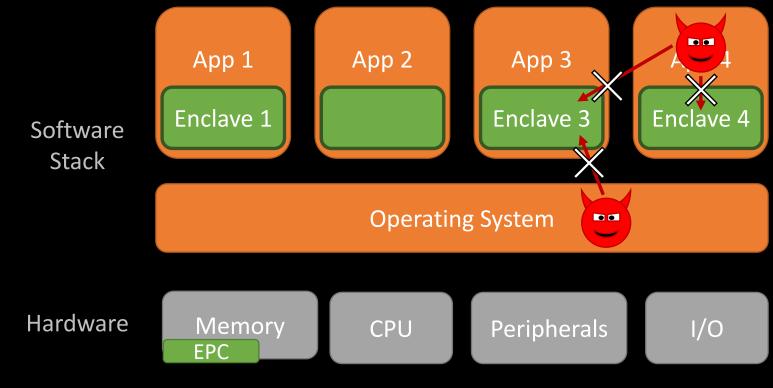
EPC: Enclave Page Cache (Dedicated Physical Memory)

- OS creates and manages enclaves, allocates memory from Enclave Page Cache (EPC)
- OS maps physical to virtual memory, as well as loads data and code into enclave
- Trust assumptions: All software components untrusted



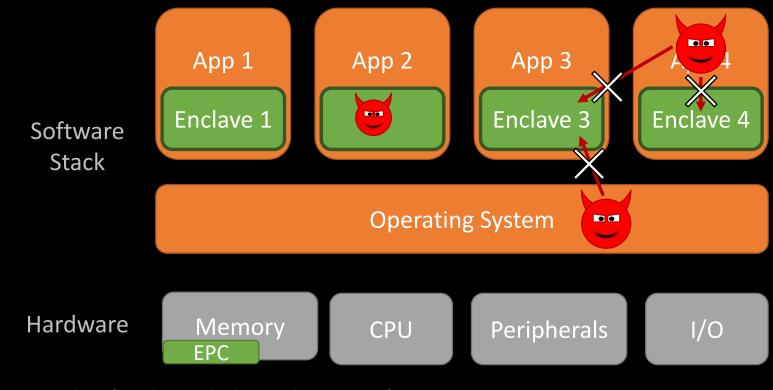
EPC: Enclave Page Cache (Dedicated Physical Memory)

 Asynchrones Enclave Exit (AEX): Enclaves interruptable, CPU saves/deletes context in CPU registers



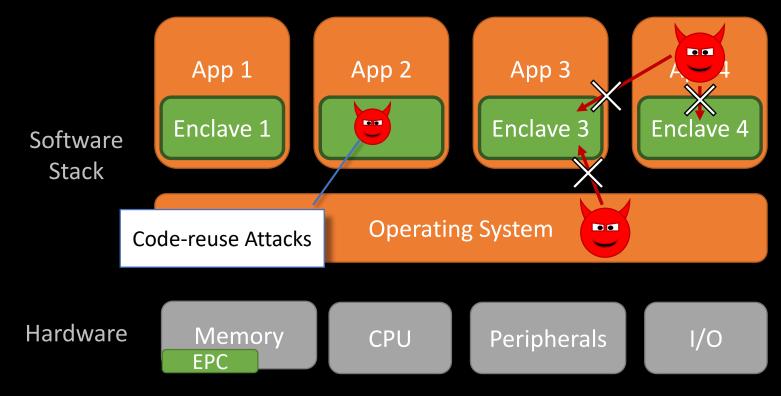
EPC: Enclave Page Cache (Dedicated Physical Memory)

 Asynchrones Enclave Exit (AEX): Enclaves interruptable, CPU saves/deletes context in CPU registers



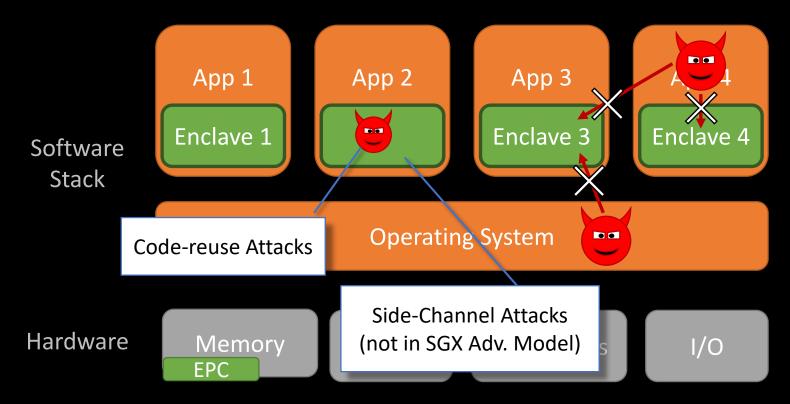
EPC: Enclave Page Cache (Dedicated Physical Memory)

 Asynchrones Enclave Exit (AEX): Enclaves interruptable, CPU saves/deletes context in CPU registers



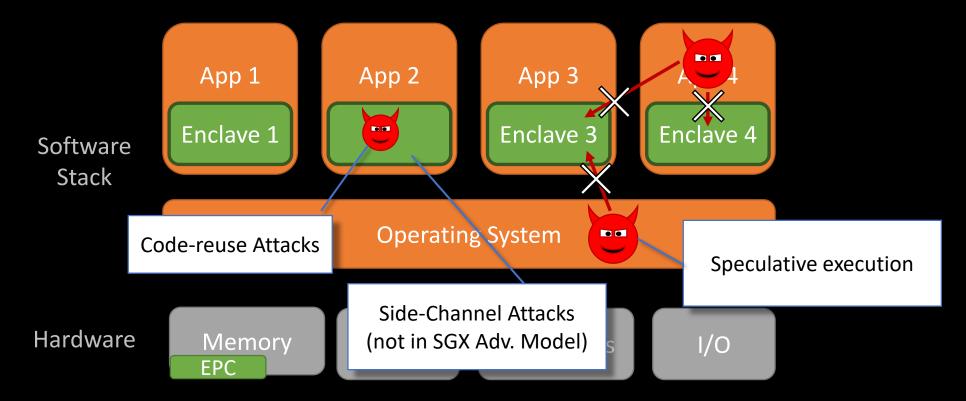
EPC: Enclave Page Cache (Dedicated Physical Memory)

 Asynchrones Enclave Exit (AEX): Enclaves interruptable, CPU saves/deletes context in CPU registers



EPC: Enclave Page Cache (Dedicated Physical Memory)

 Asynchrones Enclave Exit (AEX): Enclaves interruptable, CPU saves/deletes context in CPU registers



EPC: Enclave Page Cache (Dedicated Physical Memory)



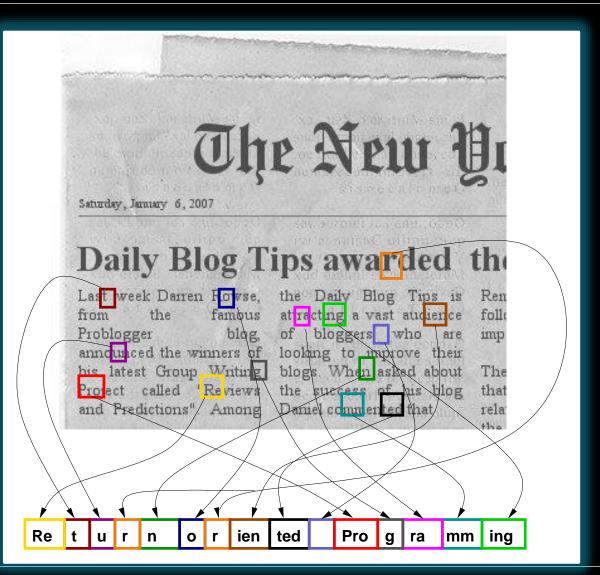
Code-reuse Attacks: Big Picture



Code-reuse Attacks: Big Picture



Code-reuse Attacks: Big Picture



Hacking in Darkness: ROP against Secure Enclaves

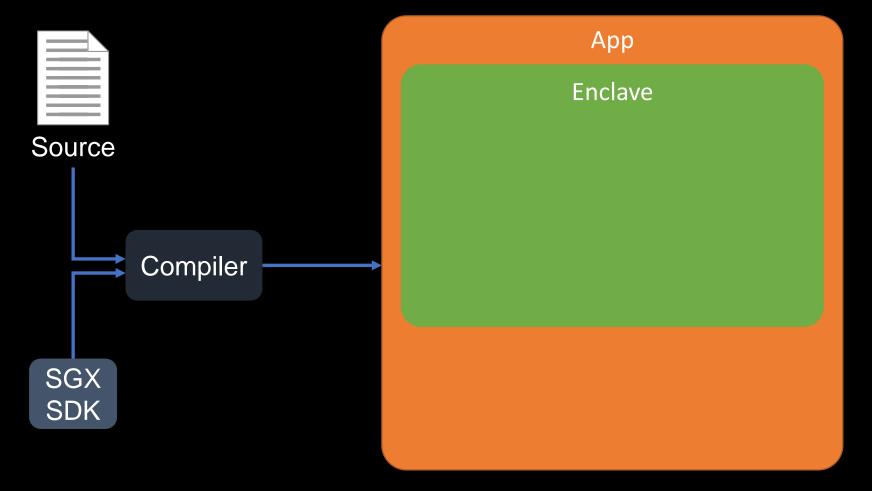
[Lee et al., USENIX Security 2017]

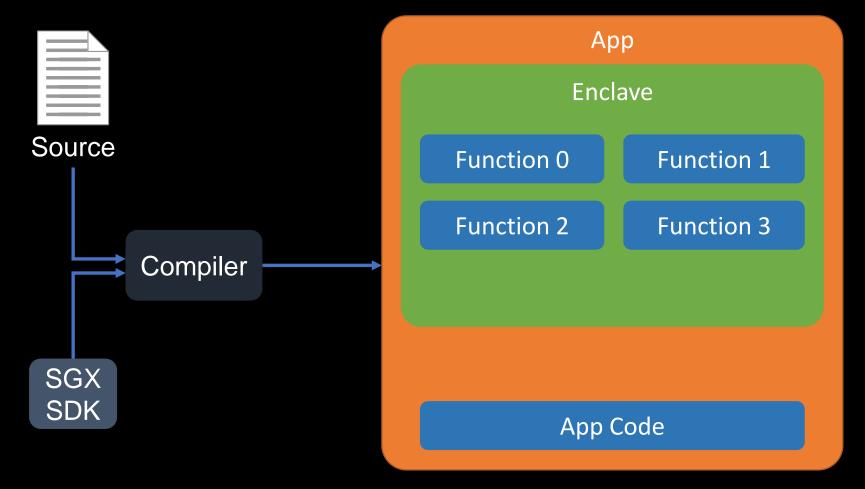
- Memory corruption attack against Intel SGX (Dark-ROP)
- Combines ROP techniques with oracles that inform about internal state of a victim enclave
- Requires kernel privileges
- Relies on running the target enclave multiple times and crashes to leak information
- Demonstrates how the security of SGX can be disarmed
 - Exfiltration of all memory contents from the enclave (code and data)
 - Bypassing the SGX attestation

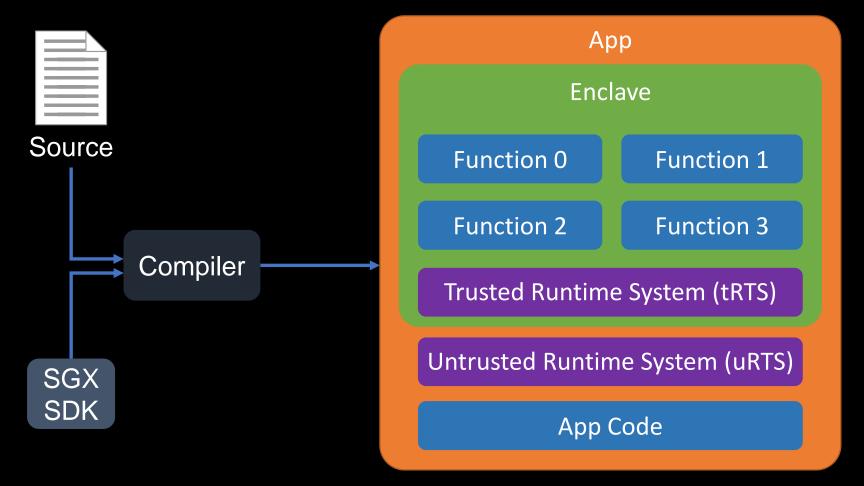
SGX-Shield: Randomization for SGX Enclaves

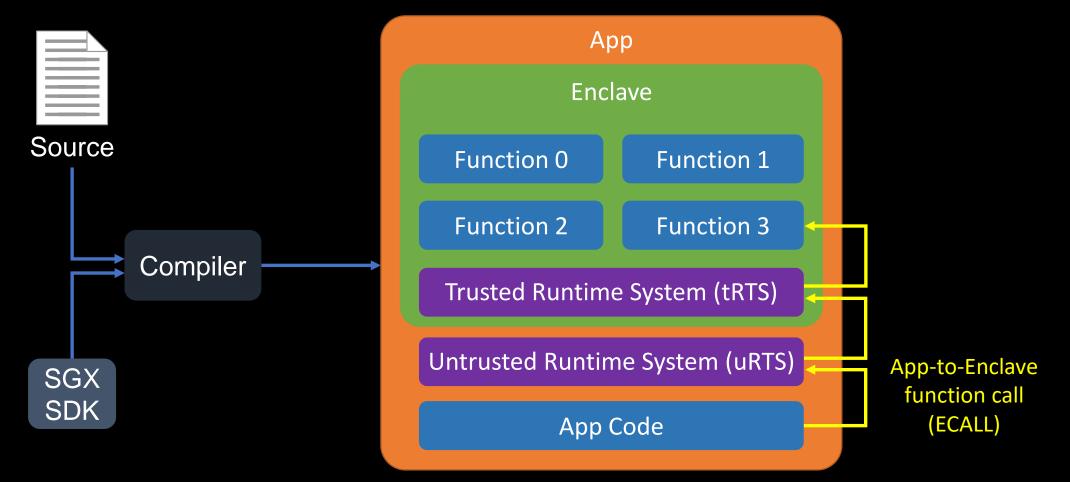
[Seo et al., NDSS 2017]

- Address Space Layout Randomization (ASLR) for SGX enclaves
- Effective against ROP, since it relies on addresses of code snippets (gadgets)
- Limited entropy due to limited memory space
- Still effective against Dark-ROP
 - Since an enclave will be re-randomized after the crash

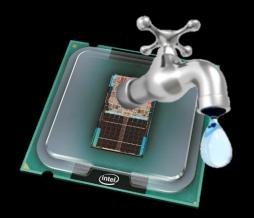




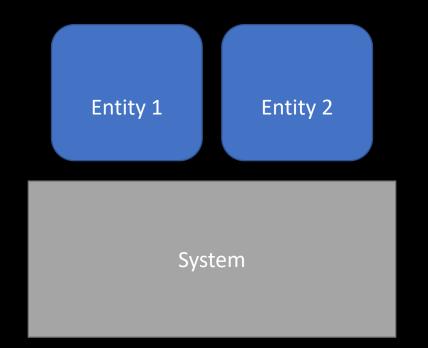


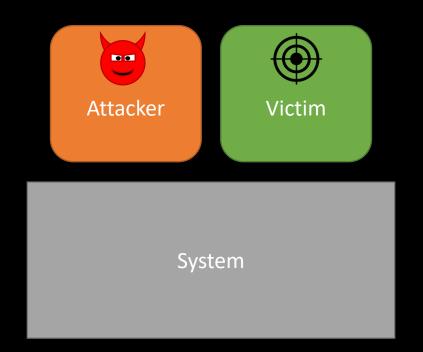


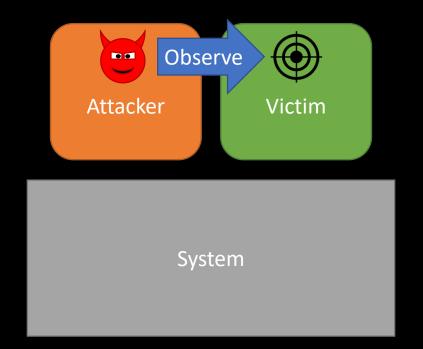
- tRTS is not randomized by SGX-Shield
- It cannot be randomized due to architectural specifics
 - E.g., enclave functions are invoked using fixed pre-defined entry points
- Contributions by Biondo et al.:
 - show that tRTS has enough gadgets to mount ROP
 - develop new techniques that do not require enclave crashes
 - new techniques do not require kernel privileges from an attacker



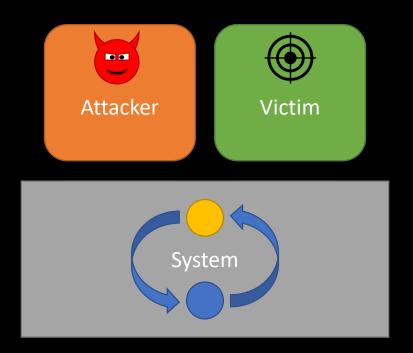
Leaky SGX

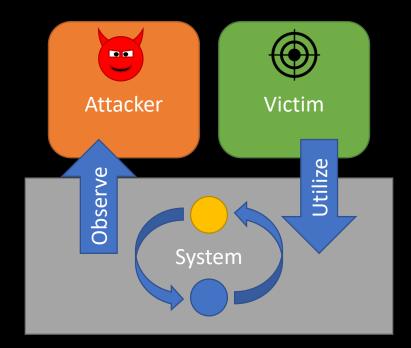


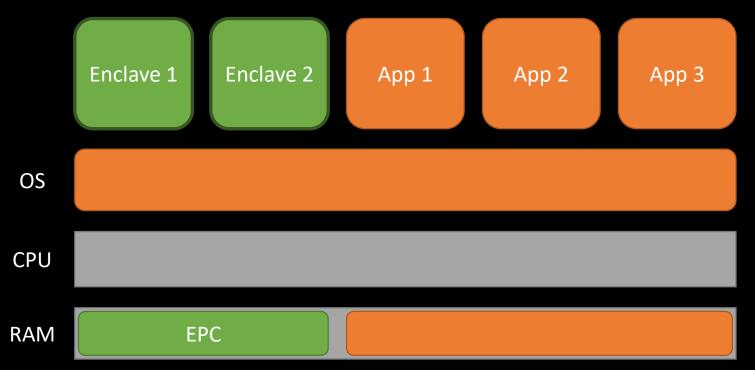






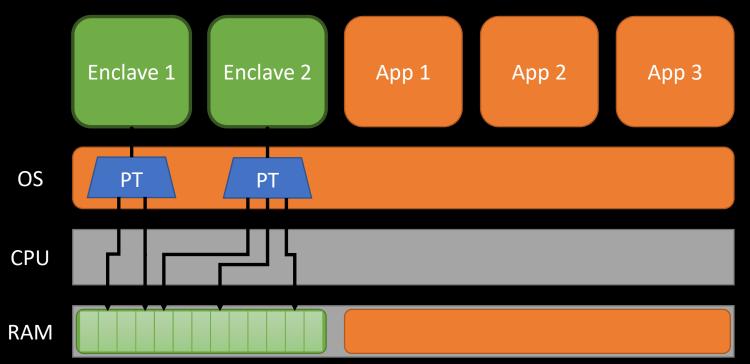






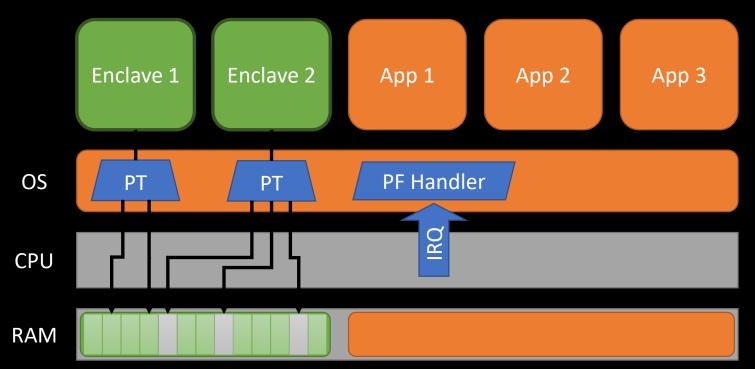
Granularity: page 4K, good for big data structures

EPC: Enclave Page CachePT: Page TablesPF: Page-Fault



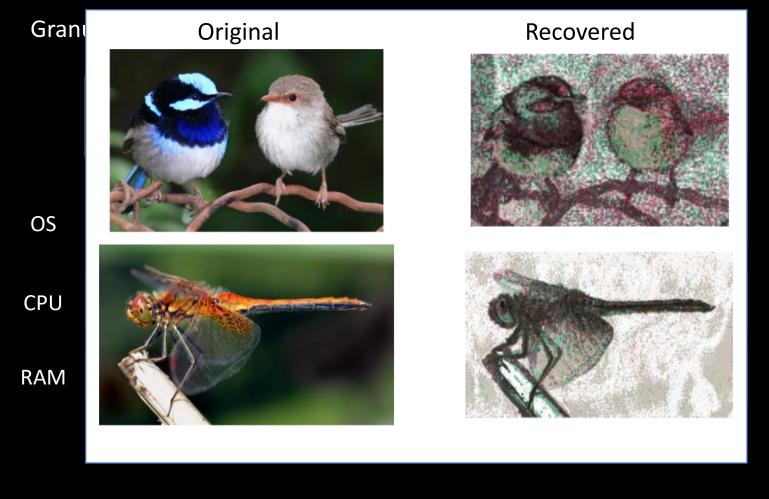
Granularity: page 4K, good for big data structures

EPC: Enclave Page Cache PT: Page Tables PF: Page-Fault

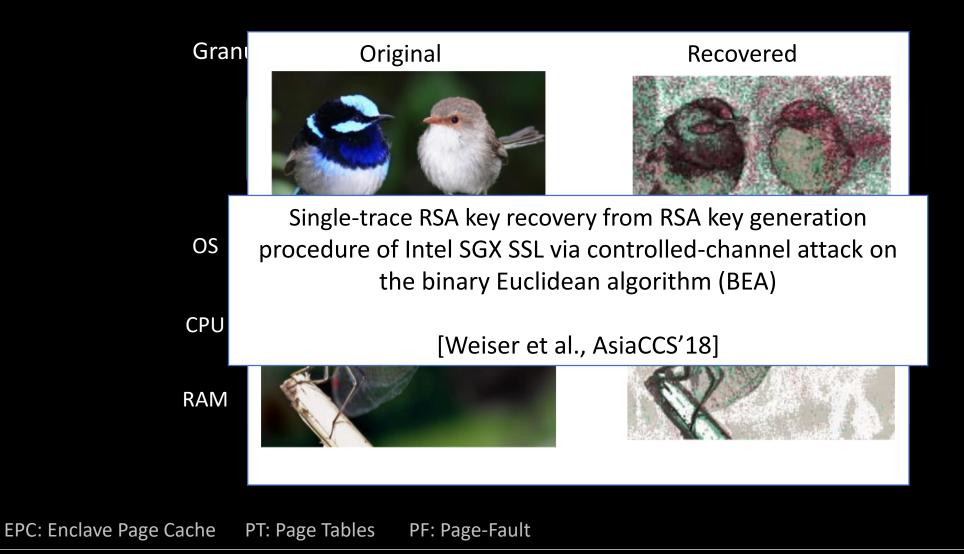


Granularity: page 4K, good for big data structures

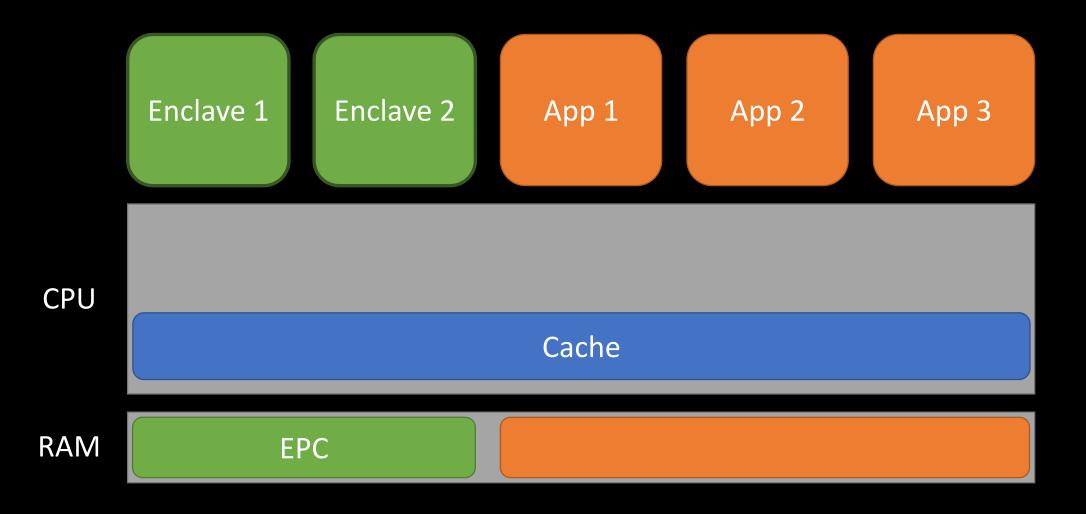
EPC: Enclave Page Cache PT: Page Tables PF: Page-Fault



EPC: Enclave Page Cache PT: Page Tables PF: Page-Fault

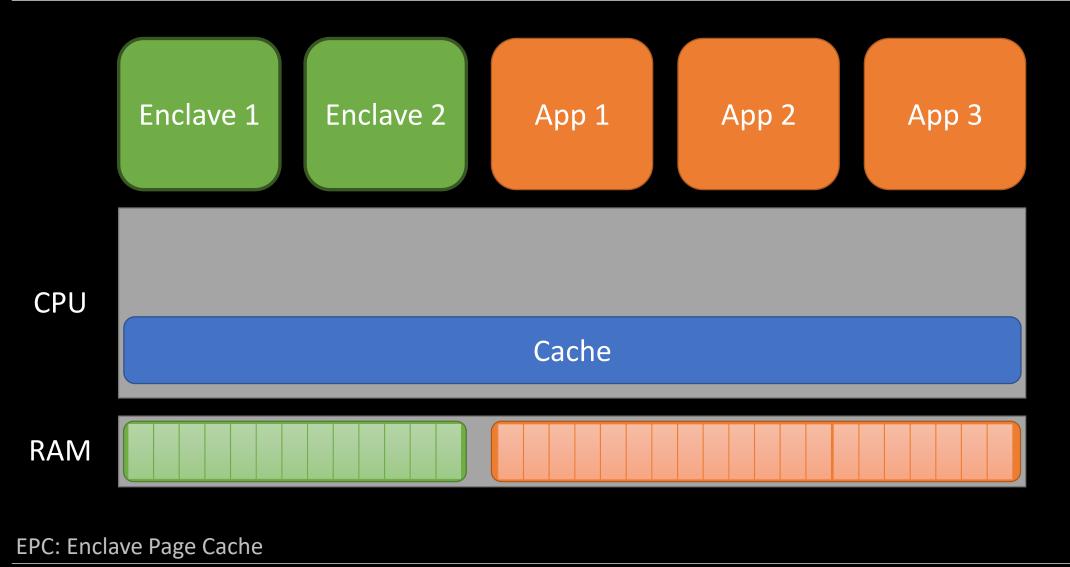


Cache Attacks on SGX: Hack in The Box

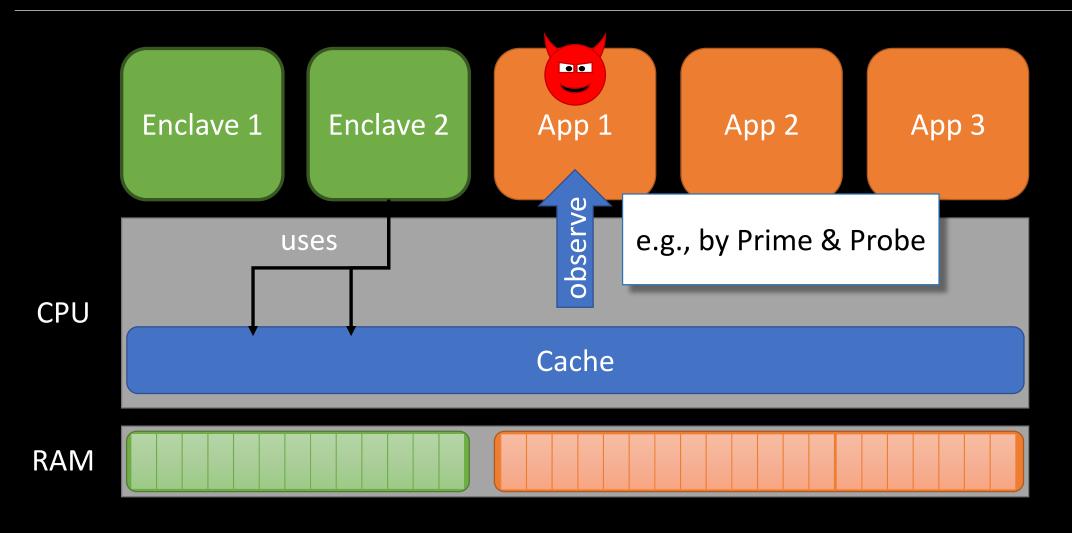


EPC: Enclave Page Cache

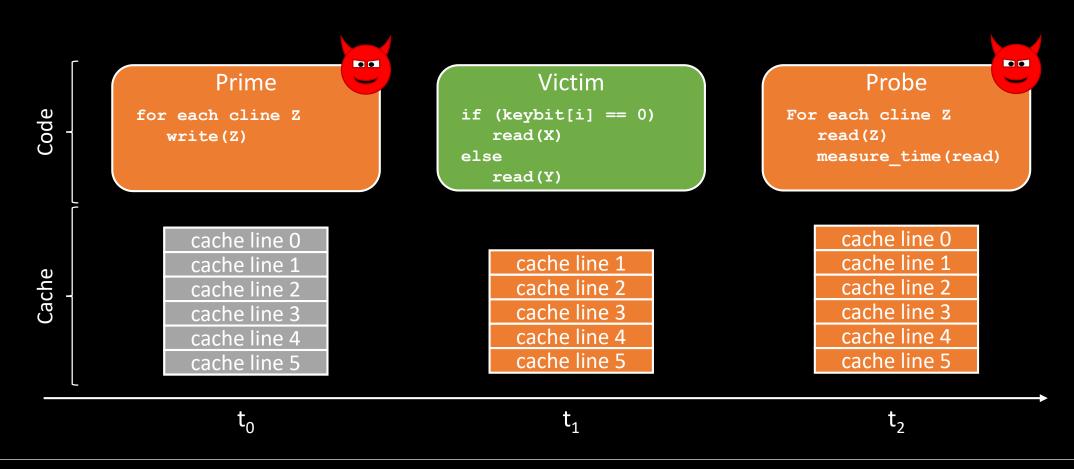
Cache Attacks on SGX: Hack in The Box

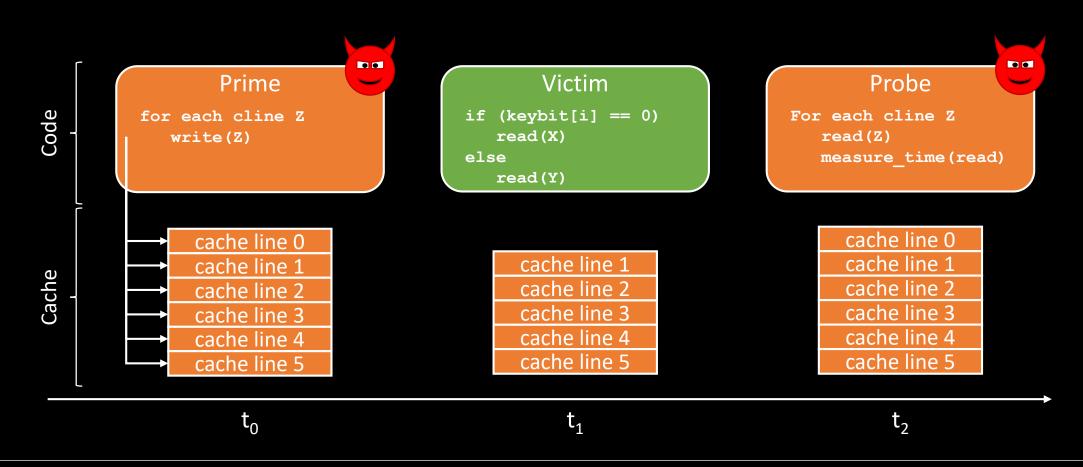


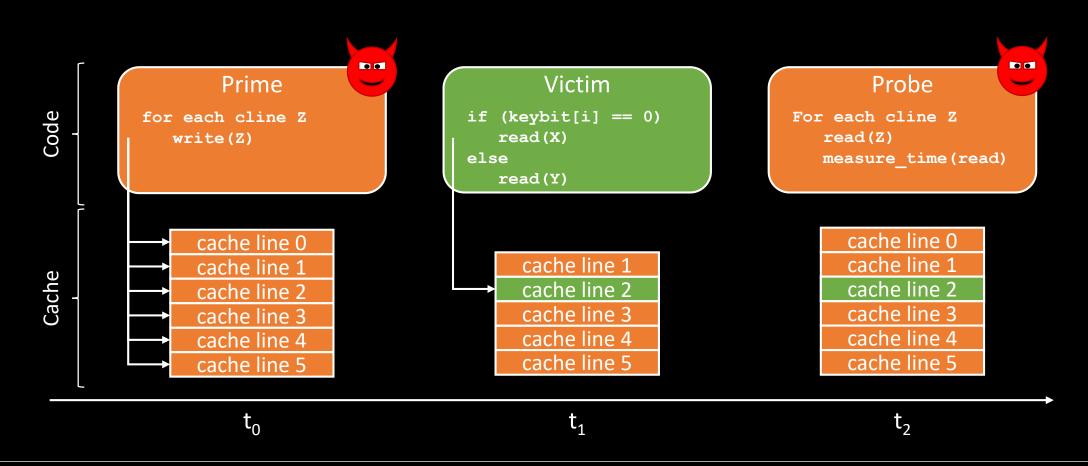
Cache Attacks on SGX: Hack in The Box

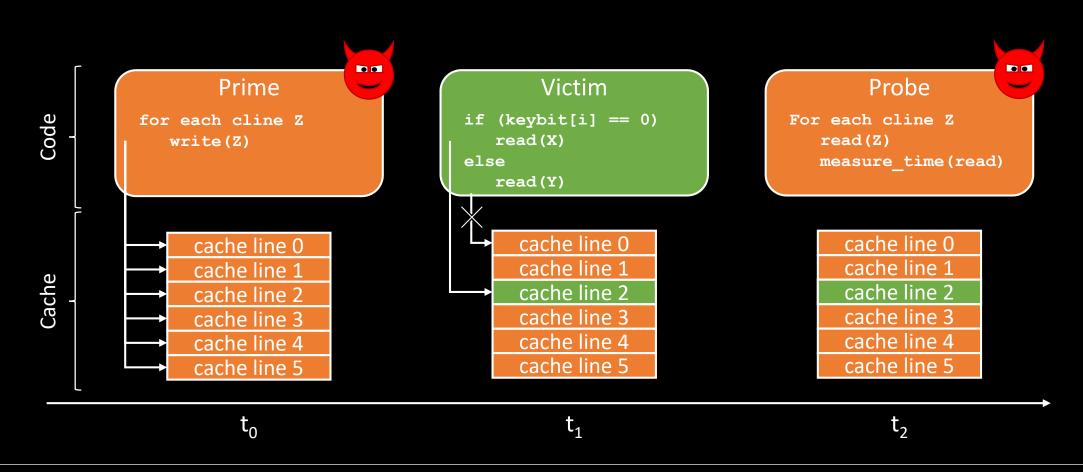


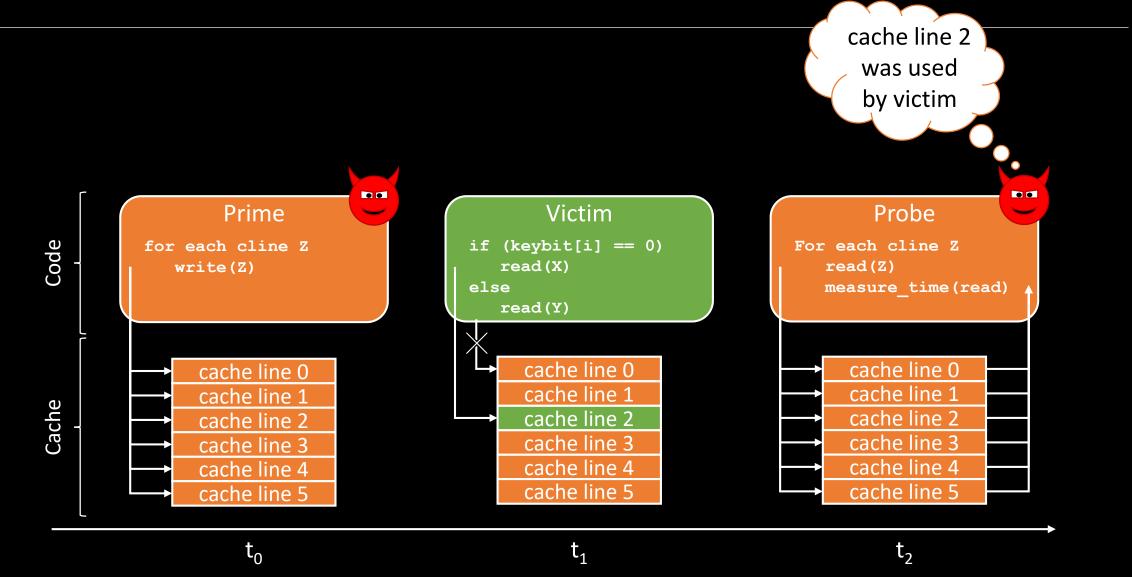
EPC: Enclave Page Cache







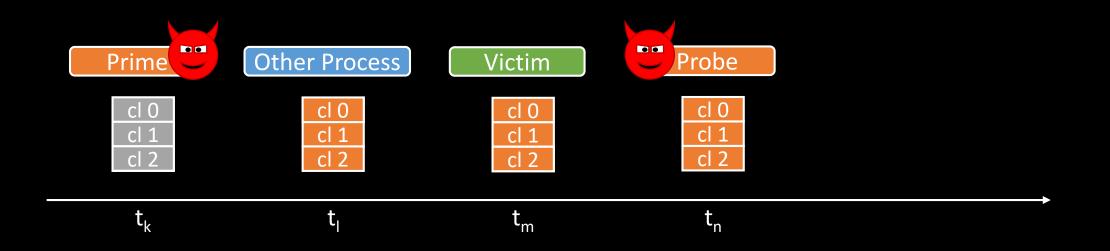




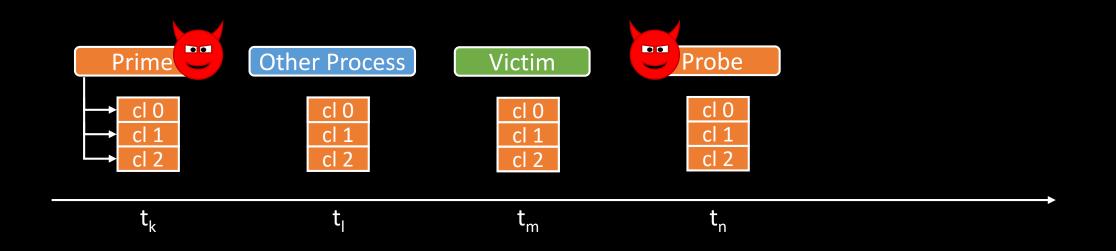
How to measure the time difference?

- #1: Time Stamp Counter (TSC)
 - Not precise enough to reliably distinguish the difference between L1 vs. L2 hits
 - Reading the time stamp counter by itself suffers from noise
- #2: Counting thread:
 - a thread that only performs a loop that constantly increments a value (basically a timer)
 - Slows down the victim, can be detected
- #3: Performance Monitoring Counter (PMC):
 - can be configured to count different events: executed cycles, cache hits or cache misses for the different caches, mis-predicted branches, etc.
 - Anti Side-channel Interference (ASCI) feature:
 - Can be configured to disable thread-specific performance monitoring of enclaves

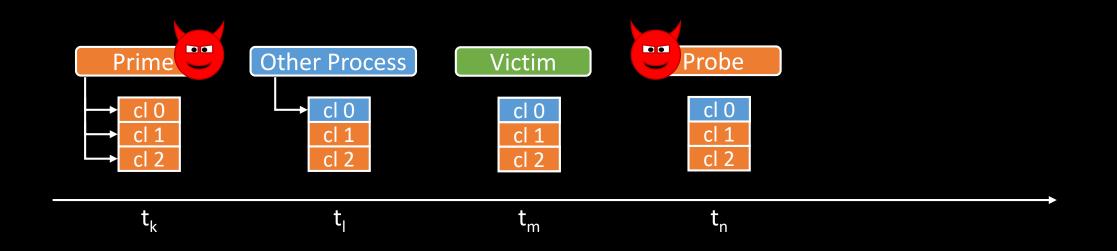
- Operating System and any other software running on the platform generate noise
- Even attacker's own code pollutes the cache



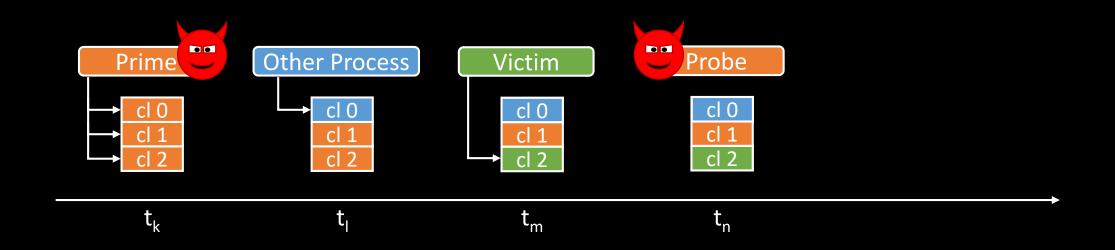
- Operating System and any other software running on the platform generate noise
- Even attacker's own code pollutes the cache



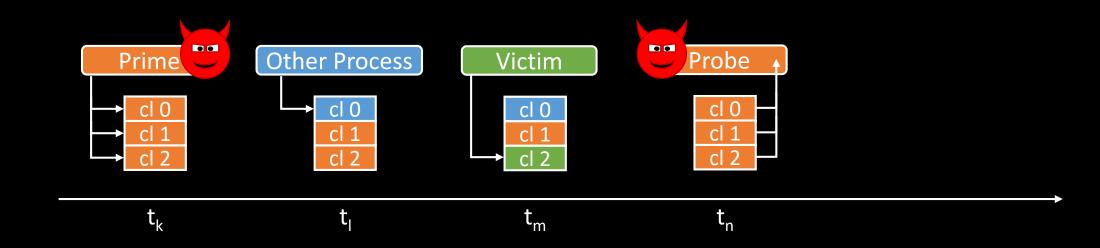
- Operating System and any other software running on the platform generate noise
- Even attacker's own code pollutes the cache



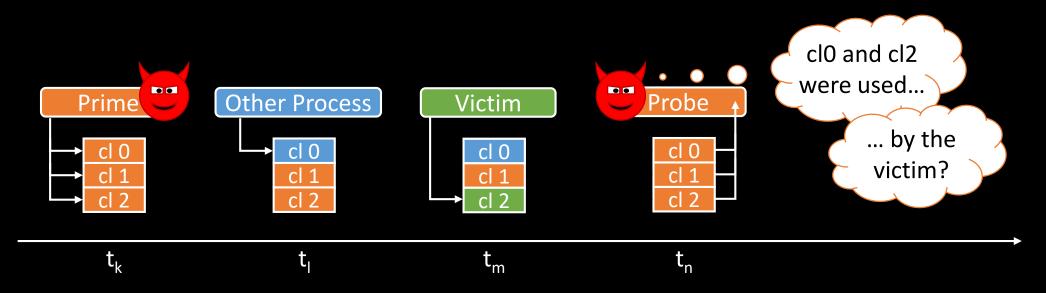
- Operating System and any other software running on the platform generate noise
- Even attacker's own code pollutes the cache

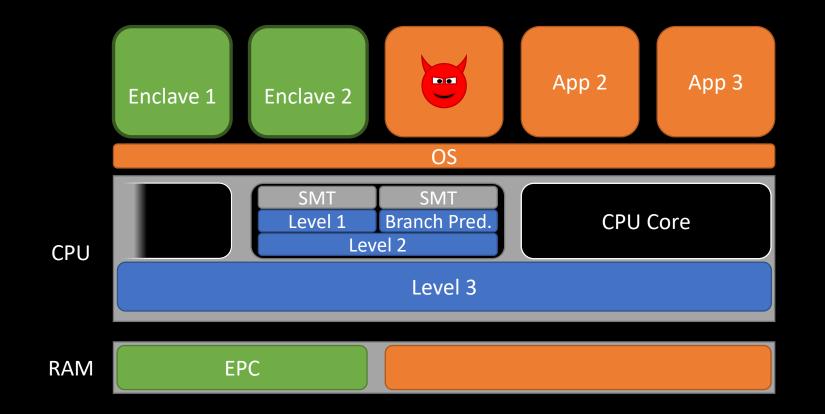


- Operating System and any other software running on the platform generate noise
- Even attacker's own code pollutes the cache

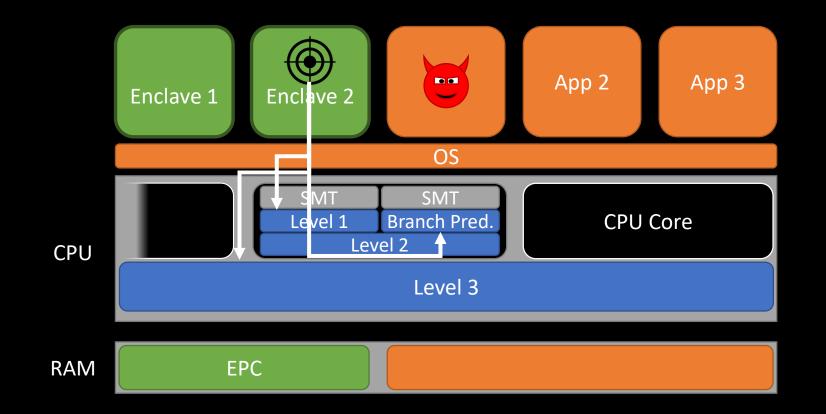


- Operating System and any other software running on the platform generate noise
- Even attacker's own code pollutes the cache

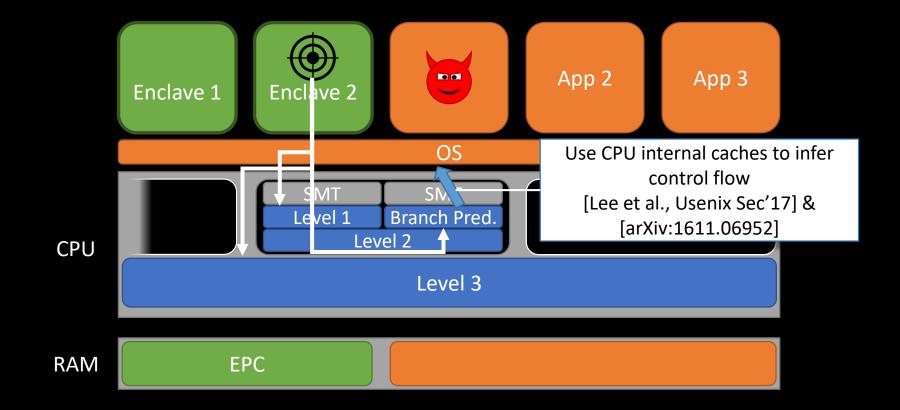




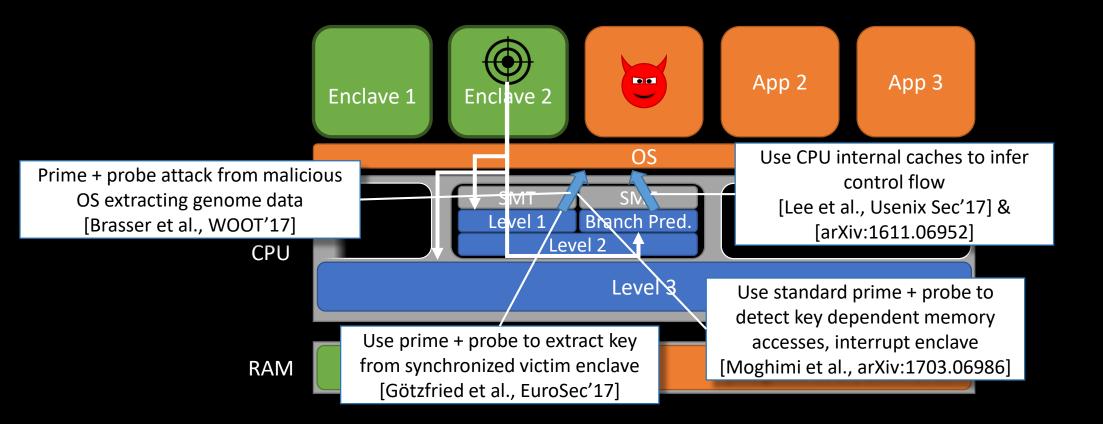
EPC: Enclave Page Cache SMT: Simultaneous Multithreading



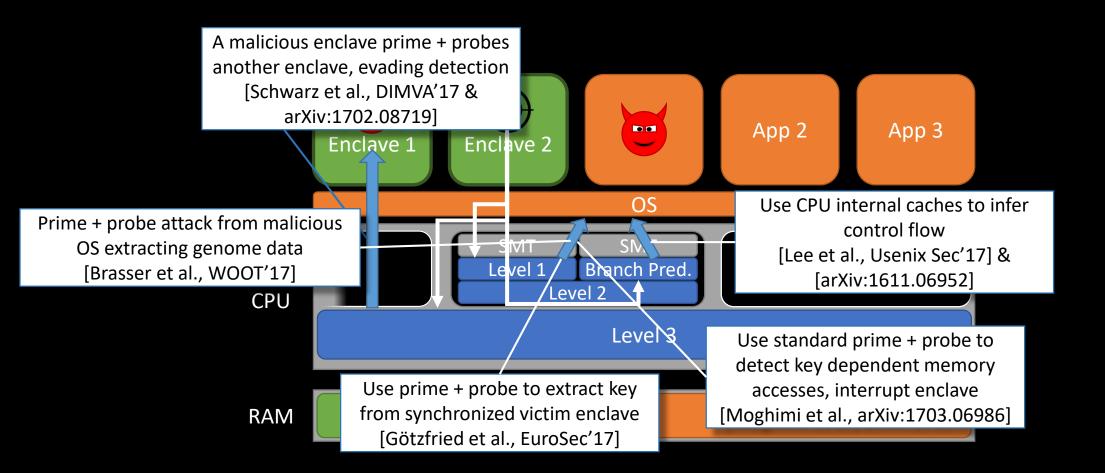
EPC: Enclave Page Cache SMT: Simultaneous Multithreading



EPC: Enclave Page Cache SMT: Simultaneous Multithreading



EPC: Enclave Page Cache SMT: Simultaneous Multithreading

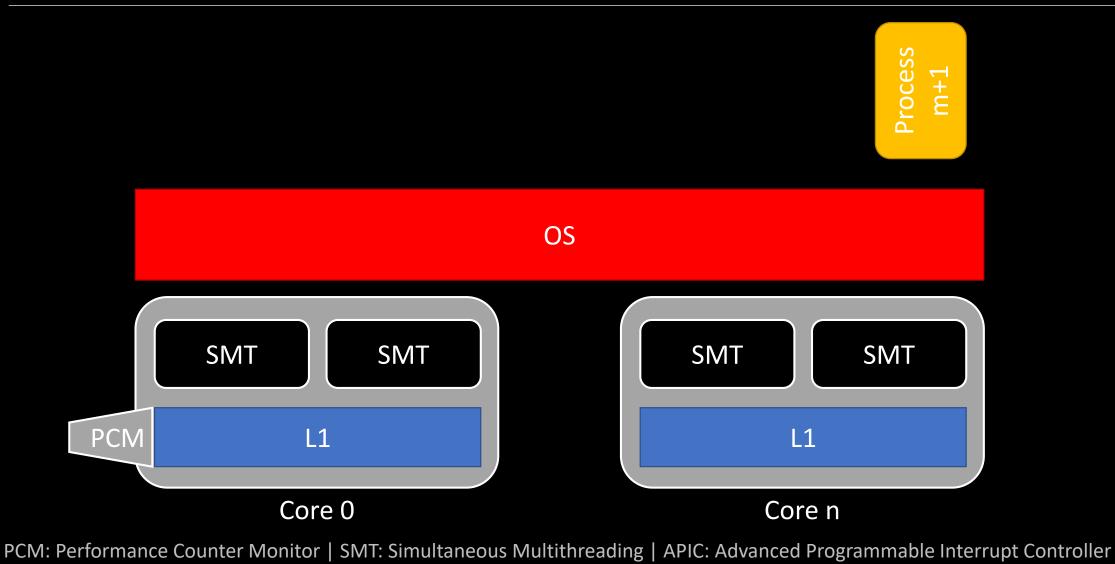


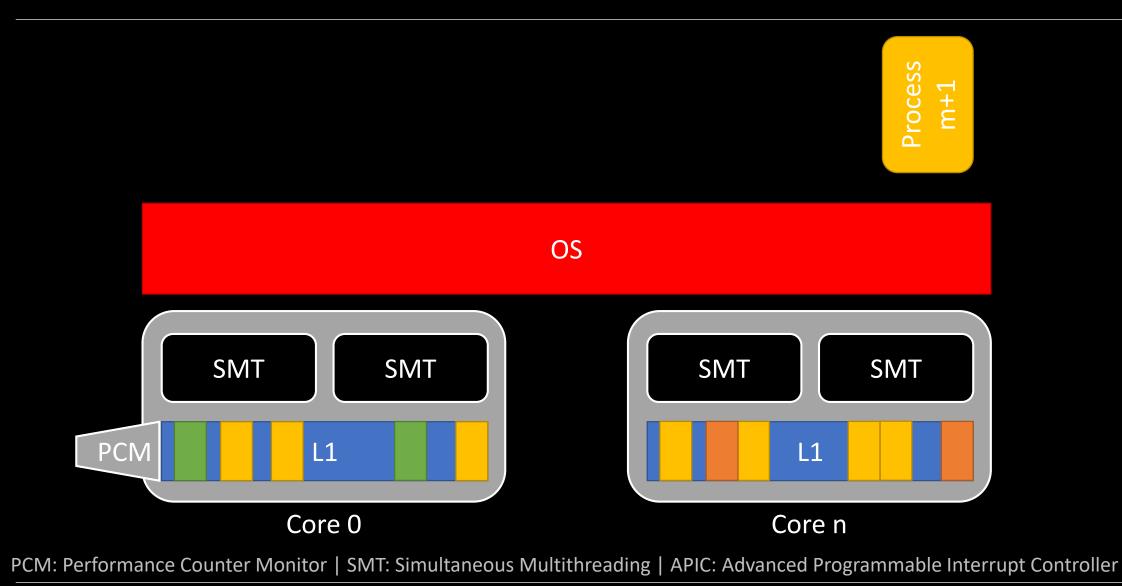
EPC: Enclave Page Cache SMT: Simultaneous Multithreading

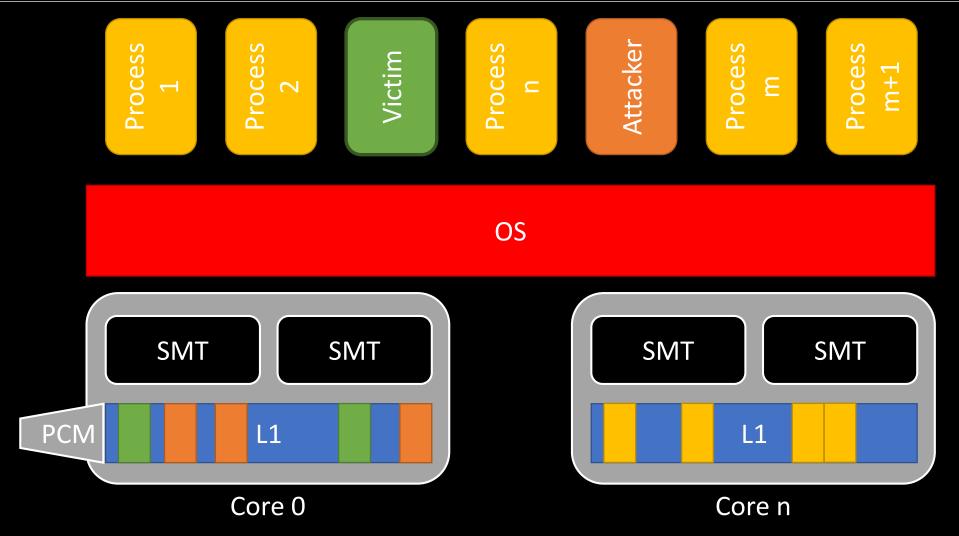
SGX Side-Channel Attacks Comparison

	Attack Type	Observed Cache	Interrupting Victim	Time Measurement	Attacker Code	Attacked Victim
Lee et al.	Branch Shadowing	BTB / LBR	Yes	Execution Timing	OS	RSA & SVM classifier
Moghimi et al.	Prime + Probe	L1(D)	Yes	TCS	OS	AES
Götzfried et al.	Prime + Probe	L1(D)	No	PCM	OS	AES
Our Attack	Prime + Probe	L1(D)	No	PCM	OS	RSA & Genome Sequencing
Schwarz et al.	Prime + Probe	L3	No	Counting Thread	Enclave	AES

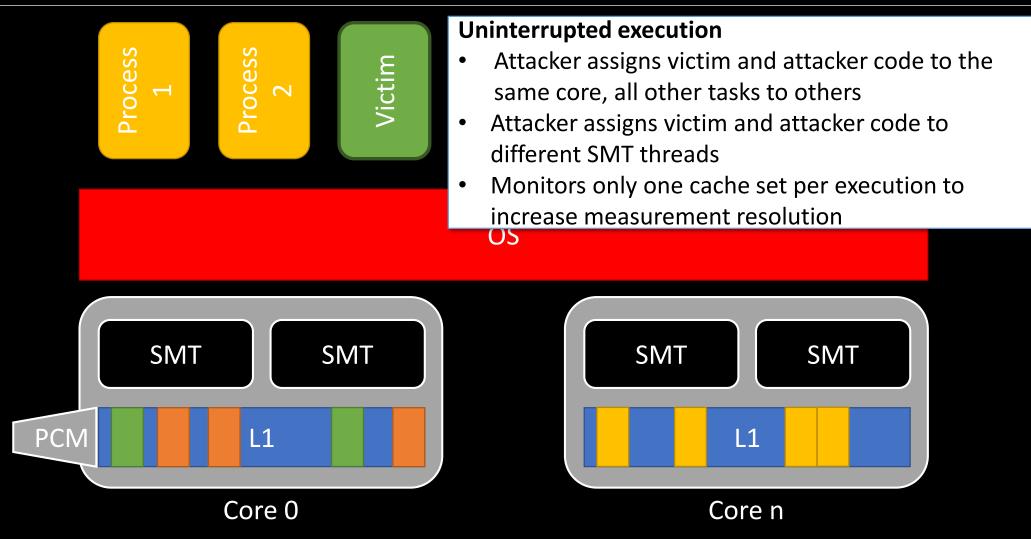
PCM: Performance Counter Monitor BTB: Branch Target Buffer LBR: Last Branch Record TSC: Time Stamp Counter



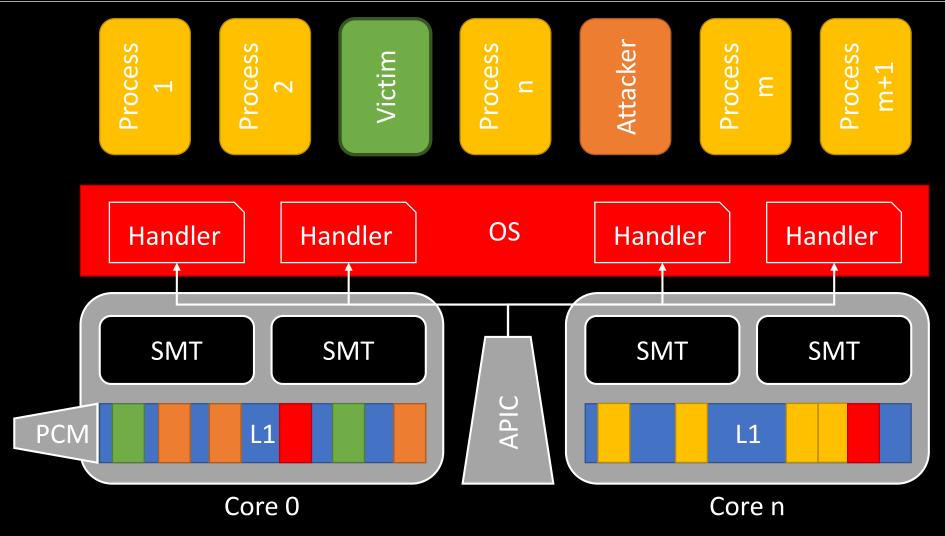




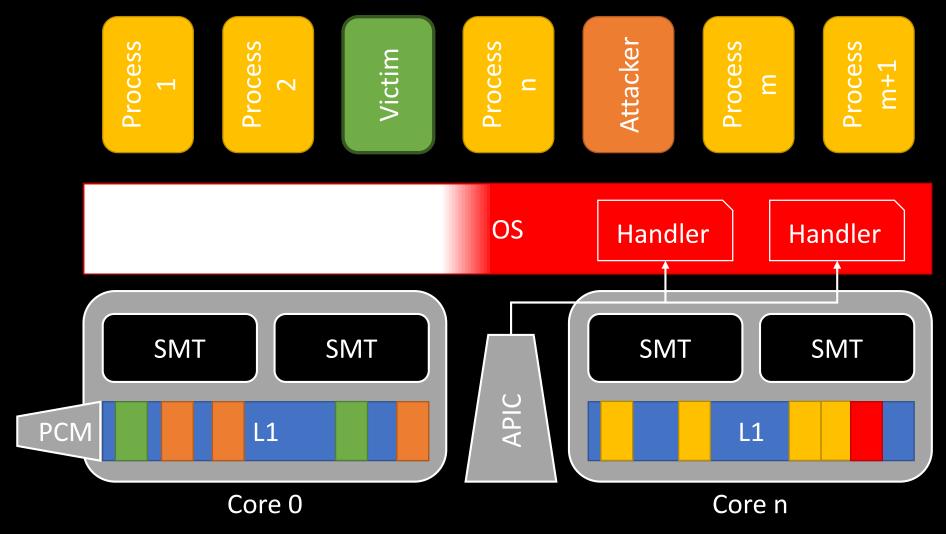
PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller



PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller



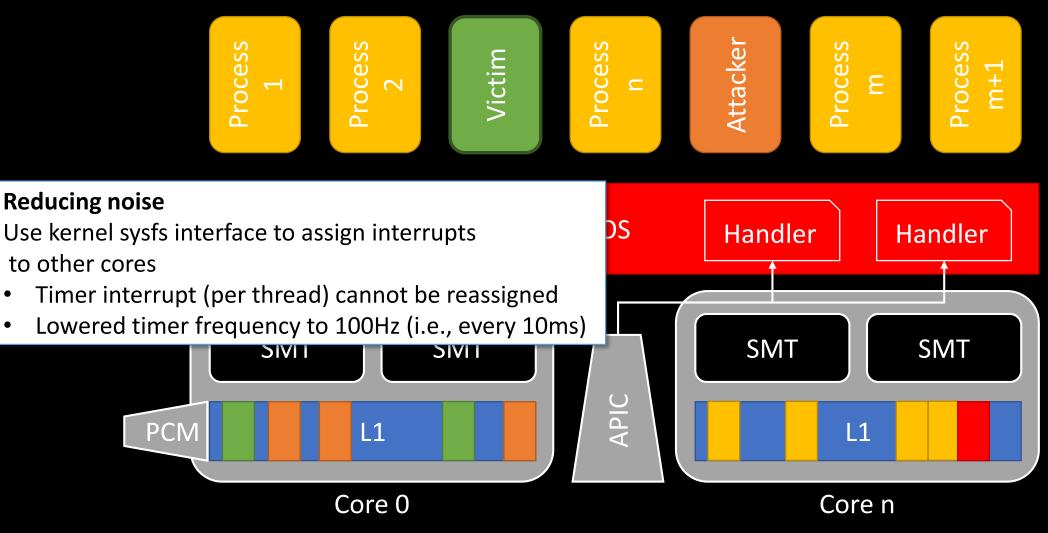
PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller



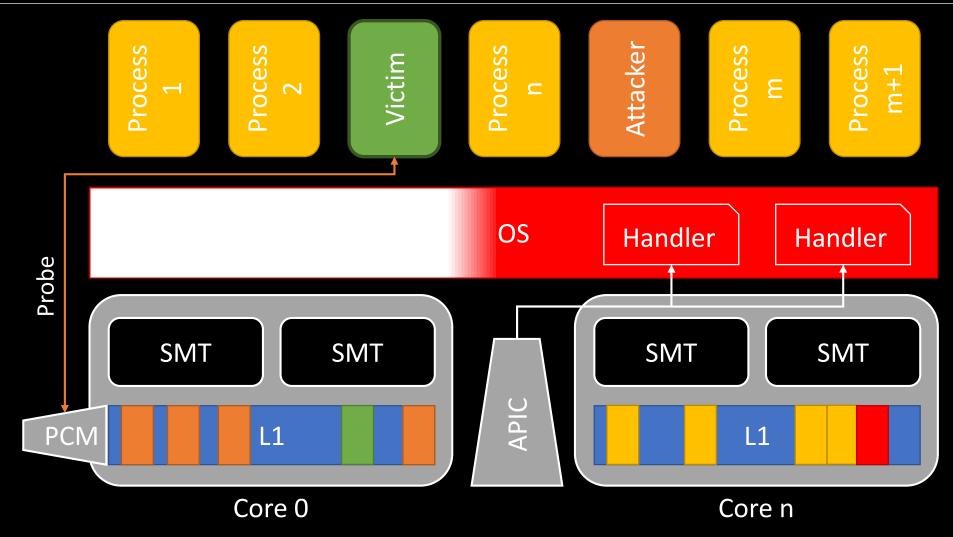
PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller

٠

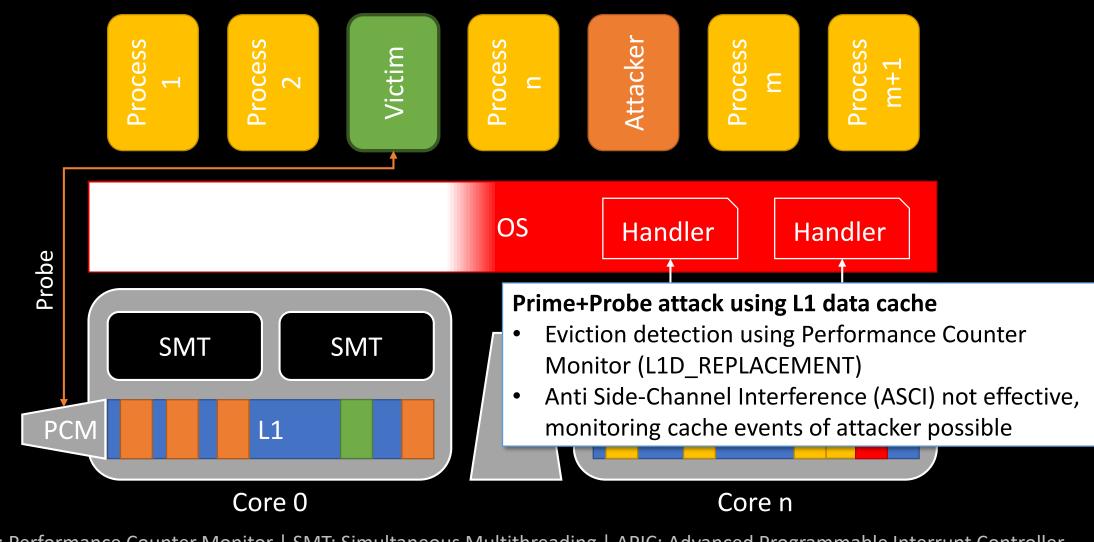
ullet



PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller



PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller



PCM: Performance Counter Monitor | SMT: Simultaneous Multithreading | APIC: Advanced Programmable Interrupt Controller

Our Attack Use-Cases



[arXiv:1702.07521]

[Brasser et al., WOOT 2017]



- Attacking RSA implementation from the Intel IIP crypto library in the Intel SGX SDK
- Extracting 2048-bit RSA decryption key
- Attacking open source k-mer analysis tool PRIMEX [Lexa et al., Bioinformatics 2003]
- Extracting genome sequences

Extracting RSA decryption key

RSA Key Exfiltration: Victim Enclave

• RSA Decryption: $m = c^d \pmod{N}$

Algorithm 1 Fixed-window exponentiation **Input:** $a, e, N \in \mathbb{N}$ **Output:** $x \leftarrow a^e \mod N$ 1: Precompute $q[i] \leftarrow a^i$ for $1 \le i \le 2^k$ 2: Let $e = (e_j, e_{j-1}, \dots, e_1, e_0)$ 3: Initialize $x \leftarrow e_i$ 4: for $i \leftarrow j - 1$ down to 0 do $x \leftarrow x^{2^k} \mod N$ 5: 6: **if** $e_i \neq 0$ **then** $x \leftarrow g[e_i] \cdot x \mod N$ 7: end if 8: 9: end for

RSA Key Exfiltration: Victim Enclave

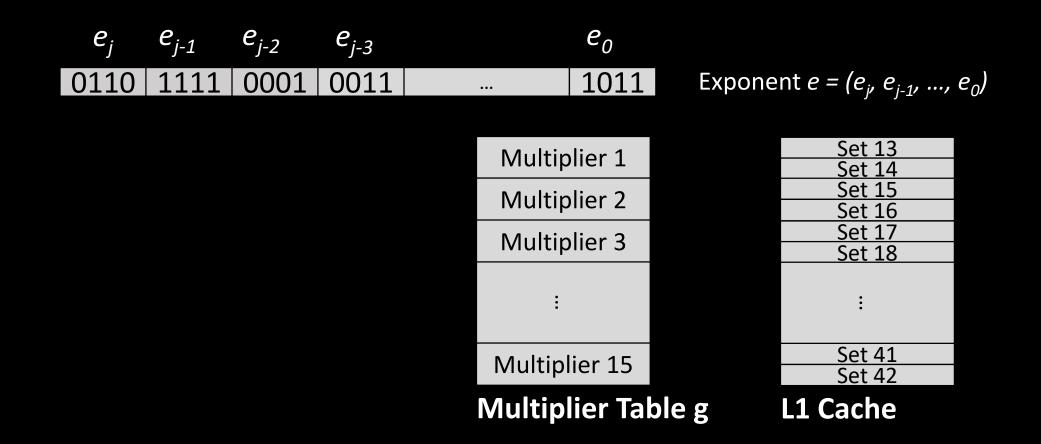
• RSA Decryption: $m = c^d \pmod{N}$

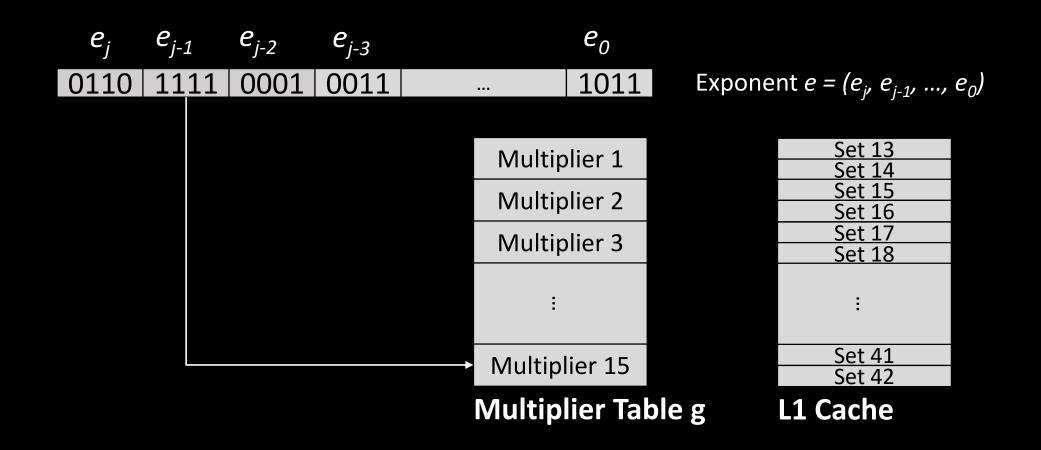
Algorithm 1 Fixed-window exponentiation Input: $a, e, N \in \mathbb{N}$ **Output:** $x \leftarrow a^e \mod N$ 1: Precompute $q[i] \leftarrow a^i$ for $1 \le i \le 2^k$ 2: Let $e = (e_j, e_{j-1}, \dots, e_1, e_0)$ 3: Initialize $x \leftarrow e_i$ 4: for $i \leftarrow j - 1$ down to 0 do $x \leftarrow x^{2^k} \mod N$ 5: 6: **if** $e_i \neq 0$ **then** $x \leftarrow g[e_i] \cdot x \mod N$ 7: end if 8: 9: end for

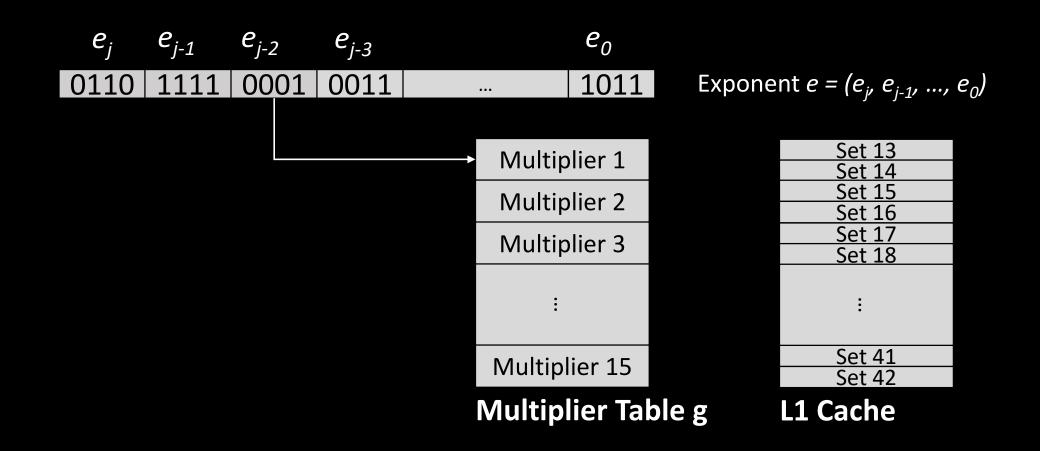
RSA Key Exfiltration: Victim Enclave

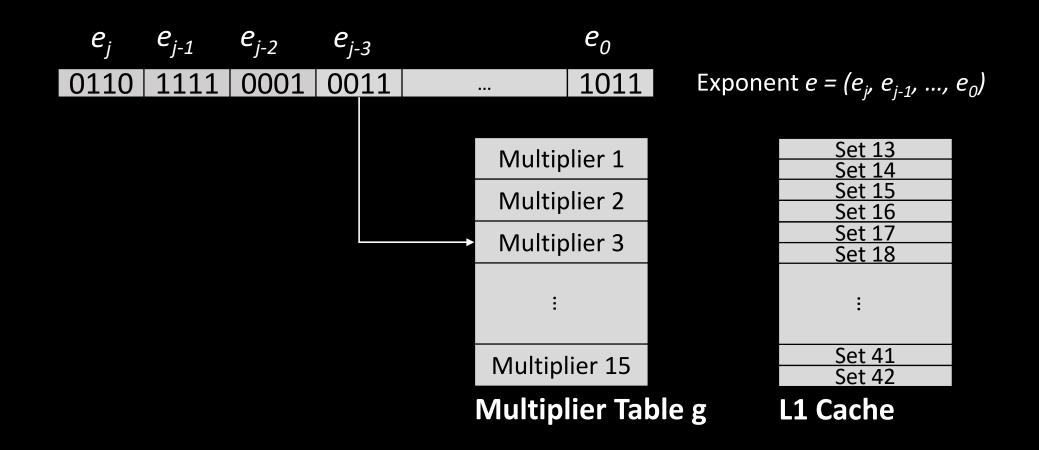
• RSA Decryption: $m = c^d \pmod{N}$

Algorithm 1 Fixed-window exponentiation Input: $a, e, N \in \mathbb{N}$ **Output:** $x \leftarrow a^e \mod N$ 1: Precompute $q[i] \leftarrow a^i$ for $1 \le i \le 2^k$ 2: Let $e = (e_j, e_{j-1}, \dots, e_1, e_0)$ 3: Initialize $x \leftarrow e_i$ 4: for $i \leftarrow j - 1$ down to 0 do $x \leftarrow x^{2^k} \mod N$ 5: 6: **if** $e_i \neq 0$ **then** $x \leftarrow g[e_i] \cdot x \mod N$ 7: end if 8: Secret-dependent memory access! 9: end for









- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

Each colored dot represents a multiplier access candidate, 15 monitoring rounds

- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

																	_																		-			-							-							
		•••				•	•	••	•) 😐	•	••		•	••	•	•	•		•	••	•	• •	•	• •	•	••	•		•	•		• •			•	• •			•	•	•	0		••	••	• •	•	• •	•••
• • • •			• • •		•	•			• •		•			• •	•	•	•		٠	•	•	•	• •	•	• •	•	• •	•	••	• •	•	•	۲	•	• •			•	•	• •	•	٠	•	• •		• •			• •	•	• •	
•••			• • •		•	•		•	•	•			•	• •	•	•	•	•	• •		•	•					• •	•	• •		•	• •			•			•	•	• •	•	• •			•	• •	•	•	• •	•		•••
					•	•	•	••	• •		• •	•		•						•	•	•	•		•	•	•	•	• •	• •	•		• •	•	•	•			•	• •				•	••	•			••			
			• • •			•			• •		• •	•	•			•			•	•		•	•		•							•		•	••	•				•				•				•	•		•	
			• •						-																												•							-								
			_	_		-	-			_			-	_	-	-		-	_		_		_	-	-	-		_	_		-	_		_	_	-	-	-		-	-				_		_		-	-		
		-				-	-	-	-		-	-	-	-	-	-		-					_					-			-			_	-	-	_	-	-	-	-		-	-	-		-	_	-	_		
				-	_	_	-	-			-	-	-							-		-	-	-		-		-						-	-	-		-		-	-		-			-		-		_		-
							-		-	-			-	-	-	-			_			-			-			-		-		-		-	-	-		-		-	-	-				-		-	-	-		
(() •				-			-				-	-	-	_	-	-		-		-		-	-				_	-		_		_		-	-	-	-	-	-	-	-						-	-	-	-		-
• • •				-	-		-	-	_			-	-		_		_			-		-	-				-							-		-			-	-	-	-				-	-	-		-	-	
(() • •	• • •	• ••	• • •	••	• • •	•	•	٠	•	•	• •	••	•	• •	•	•	•	٠	٠	••	٠		•		• •	•		•		•	•	•		•	•	•	• •	•			•		••	• •		• •		• •	••	• •	•	•
	• •	• •	• •	٠	•	••		•	•	• •	• •)	٠		•					•	•	•		•	• •		• •			•		••	•		•			•	•	•		• •				•			• •		•	• • •

Each colored dot represents a multiplier access candidate, 15 monitoring rounds

- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

Each colored dot represents a multiplier access candidate, 15 monitoring rounds

- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

		••

Each colored dot represents a multiplier access candidate, 15 monitoring rounds

- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

Each colored dot represents a multiplier access candidate, 15 monitoring rounds

Time

- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

Each colored dot represents a multiplier access candidate, 15 monitoring rounds

Time

- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

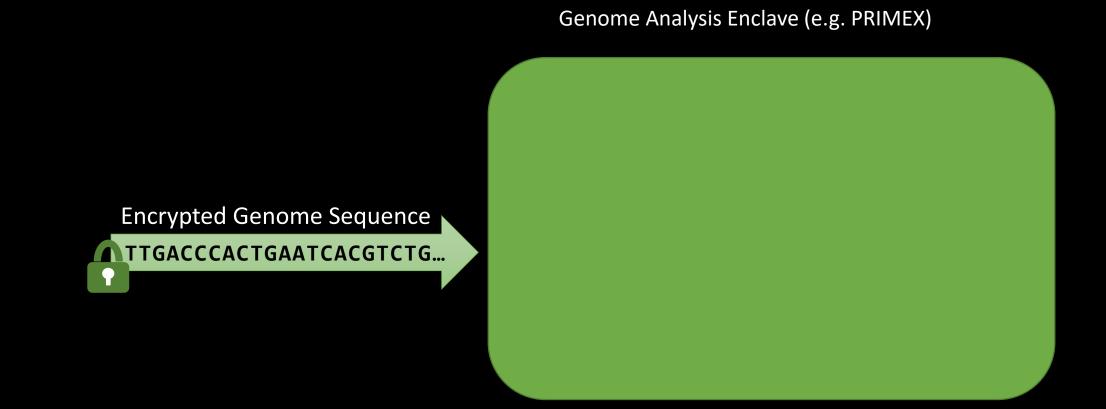
Each colored dot represents a multiplier access candidate, 15 monitoring rounds

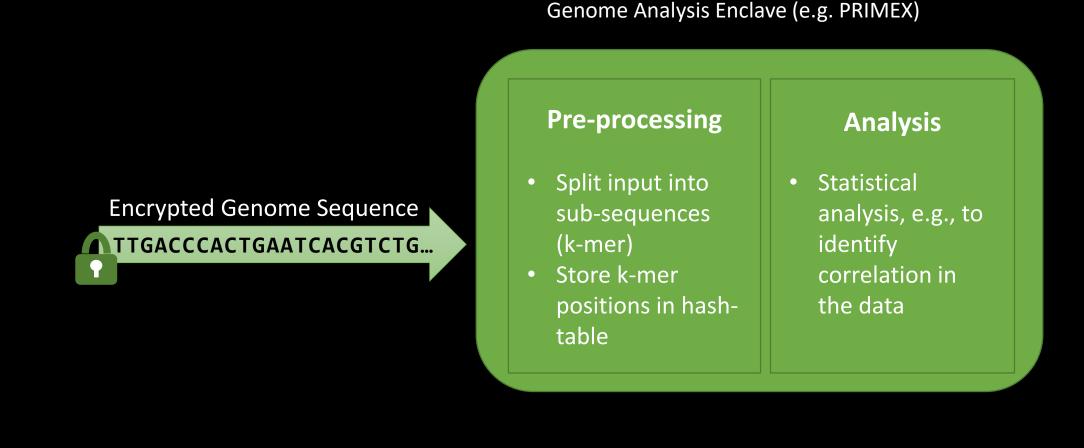
- 2048-bit Chinese Remainder Theorem RSA key
- Only 300 decryptions to leak 70% of key bits
- Enough to recover key [Heninger et. al., CRYPTO'09]

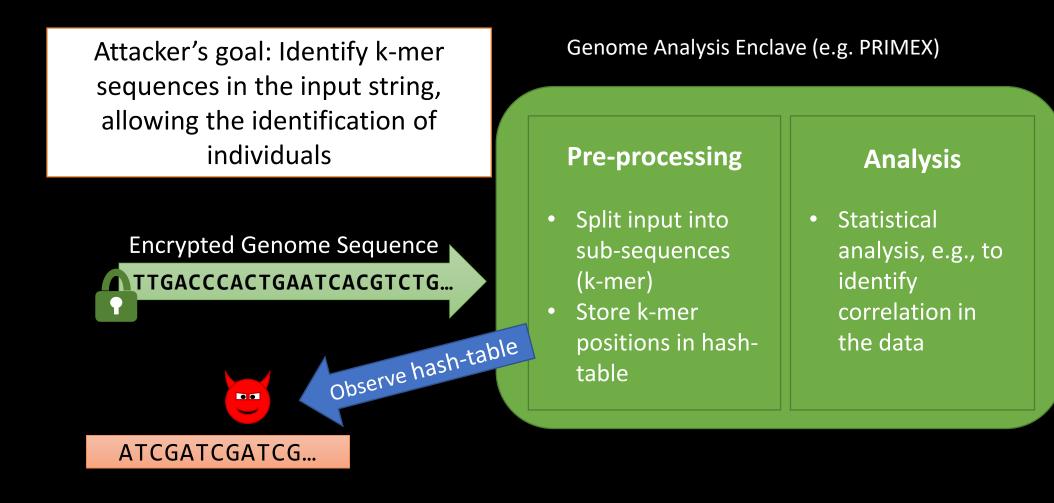
Each colored dot represents a multiplier access candidate, 15 monitoring rounds

Genome Analysis Enclave (e.g. PRIMEX)









Some Basics on Human Genomes

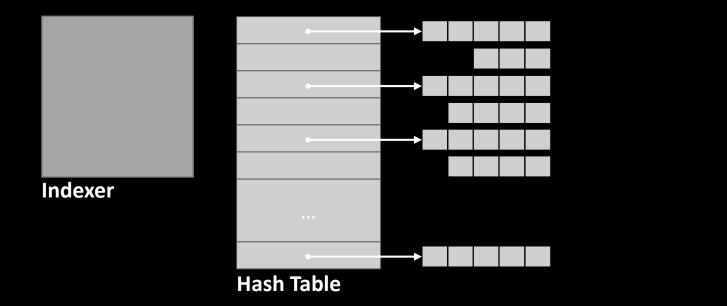
TTGACCCACTGAATCACGTCTGACCGCGCGTACGCGG TCACTTGCGGTGCCGTTTTCTTTGTTACCGACGACCG ACCAGCGACAGCCACCGCGCGCTCACTGCCACCAAAA GAGTCATATCGATCGATCGATCGATCGATCGATCGAT CGATCGATCGATCGATCGATCGATCGATCGATCATCA CAGCCGACCAGTTTCTGGAACGTTCCCGATACTGGAA CGGTCCTAATGCAGTATCCCACCCTCCTTCCATCGAC GCCAGTCGAATCACGCCGCCAGCCACCGTCCGCCAGC CGGCCAGAATACCGATGACTCGGCGGTCTCGTGTCGG TGCCGGCCTCGCAGCCATTGTACTGGCCCTGGCCGCA GTGTCGGCTGCCGCTCCGATTGCCGGGGCGCAGTCCG CCGGCAGCGGTGCGGTCTCAGTCACCATCGGCGACGT GGACGTCTCGCCTGCGAACCCAACCACGGGCACGCAG GTGTTGATCACCCCGTCGATCAACAACTCCGGATCGG CAAGCGGGTCCGCGCGCGTCAACGAGGTCACGCTGCG CGGCGACGGTCTCCTCGCAACGGAAGACAGCCTGGGG

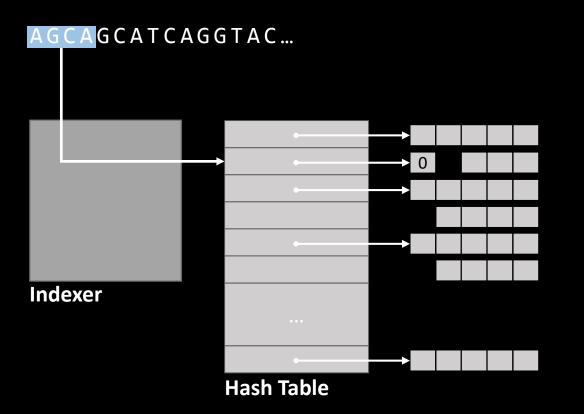
Some Basics on Human Genomes

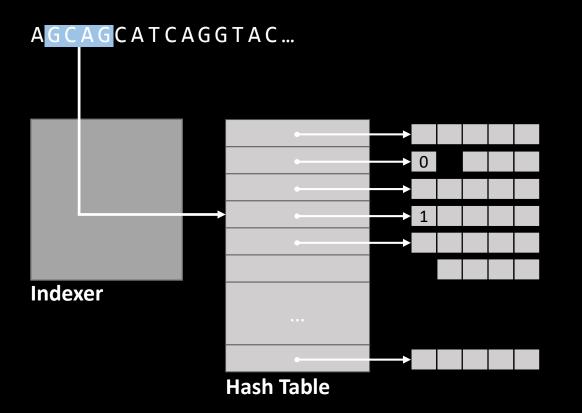
- Nucleobases
 - Adenine (A)
 - Cytosine (C)
 - Guanine (G)
 - Thymine (T)
- Microsatellite
 - Forensic analysis
 - Genetic fingerprinting
 - Kinship analysis

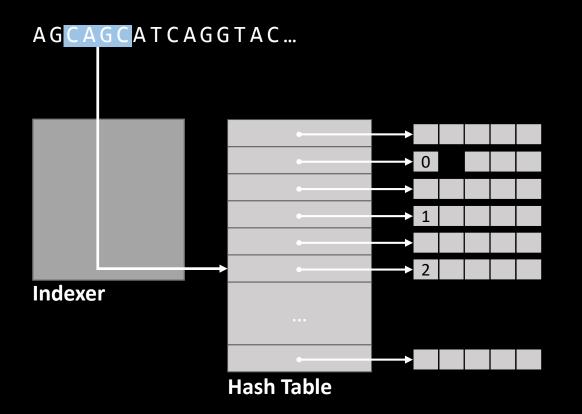
TTGACCCACTGAATCACGTCTGACCGCGCGTACGCGG TCACTTGCGGTGCCGTTTTCTTTGTTACCGACGACCG ACCAGCGACAGCCACCGCGCGCTCACTGCCACCAAAA GAGTCATATCGATCGATCGATCGATCGATCGATCGAT CGATCGATCGATCGATCGATCGATCGATCGATCATCA CAGCCGACCAGTTTCTGGAACGTTCCCGATACTGGAA CGGTCCTAATGCAGTATCCCACCCTCCTTCCATCGAC GCCAGTCGAATCACGCCGCCAGCCACCGTCCGCCAGC CGGCCAGAATACCGATGACTCGGCGGTCTCGTGTCGG TGCCGGCCTCGCAGCCATTGTACTGGCCCTGGCCGCA GTGTCGGCTGCCGCTCCGATTGCCGGGGCGCAGTCCG CCGGCAGCGGTGCGGTCTCAGTCACCATCGGCGACGT GGACGTCTCGCCTGCGAACCCAACCACGGGCACGCAG GTGTTGATCACCCCGTCGATCAACAACTCCGGATCGG CAAGCGGGTCCGCGCGCGTCAACGAGGTCACGCTGCG CGGCGACGGTCTCCTCGCAACGGAAGACAGCCTGGGG

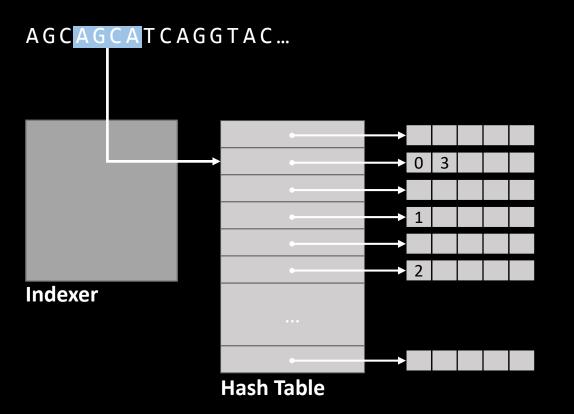
AGCAGCATCAGGTAC...

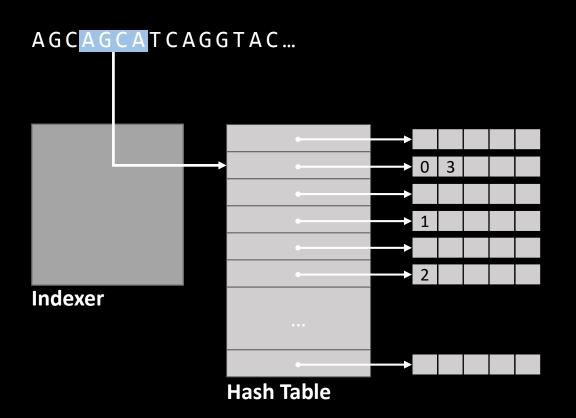




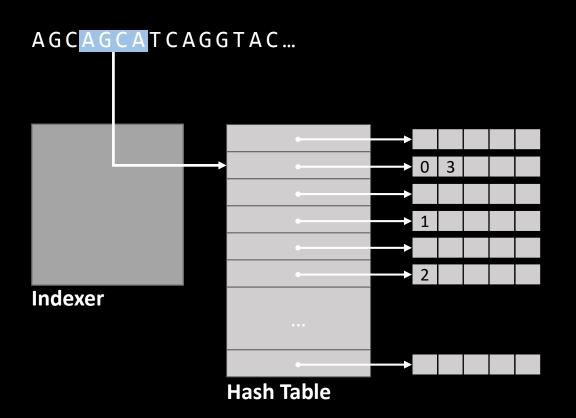




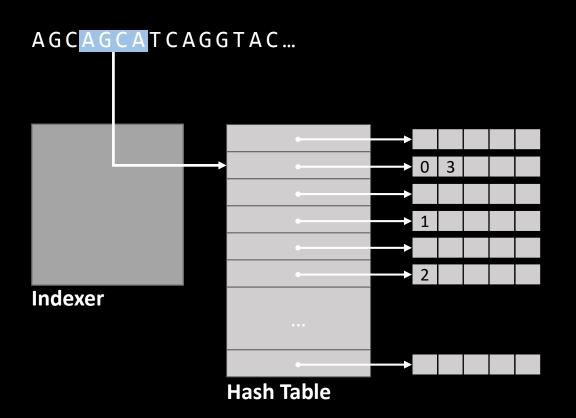




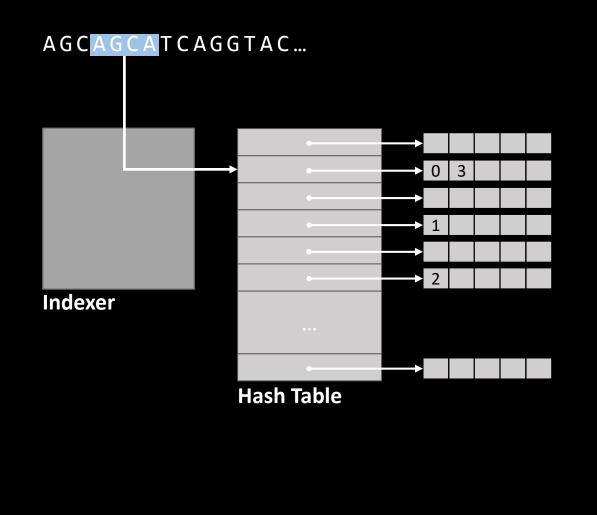
- Hash table access pattern
 - Hash table entry 8 bytes
 - Cache line size 64 bytes
 - Collisions
- Genome unstructured
- Microsatellites structured



- Hash table access pattern
 - Hash table entry 8 bytes
 - Cache line size 64 bytes
 - Collisions
- Genome unstructured
- Microsatellites structured



- Hash table access pattern
 - Hash table entry 8 bytes
 - Cache line size 64 bytes
 - Collisions
- Genome unstructured
- Microsatellites structured

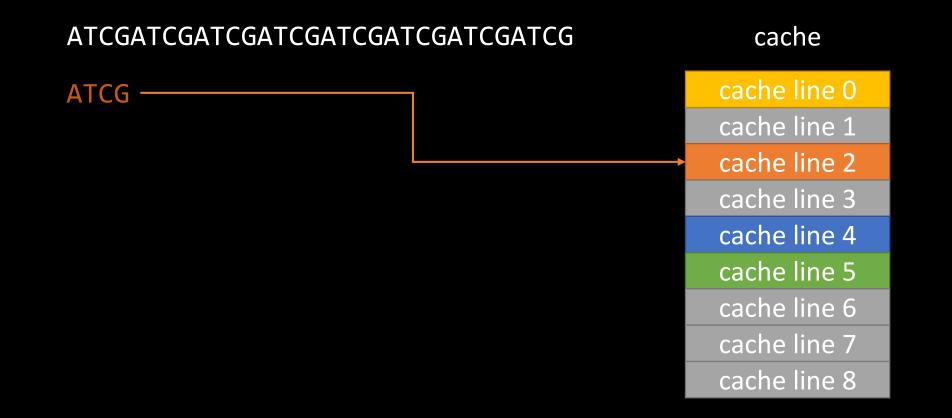


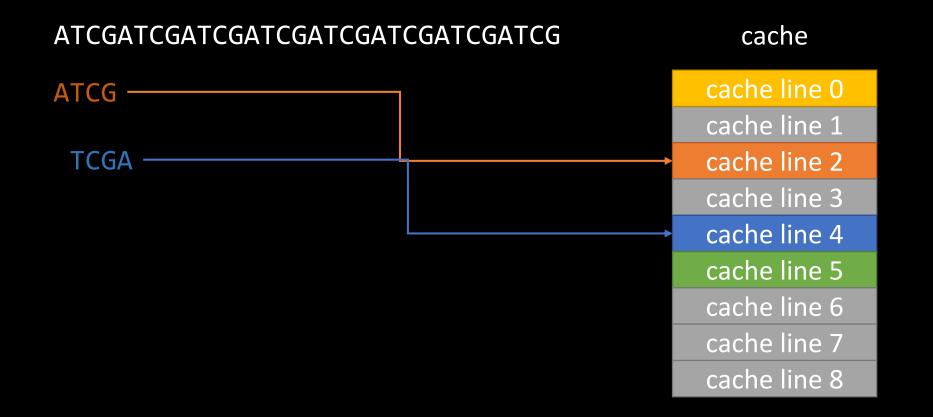
- Hash table access pattern
 - Hash table entry 8 bytes
 - Cache line size 64 bytes
 - Collisions
- Genome unstructured

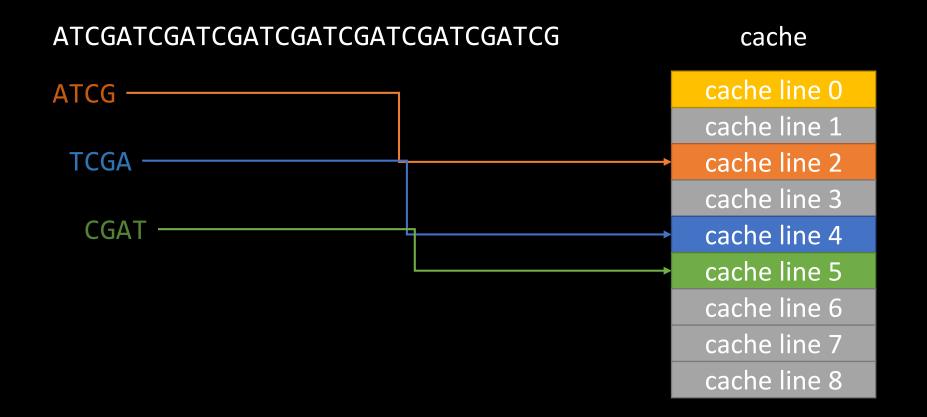
Microsatellites structured

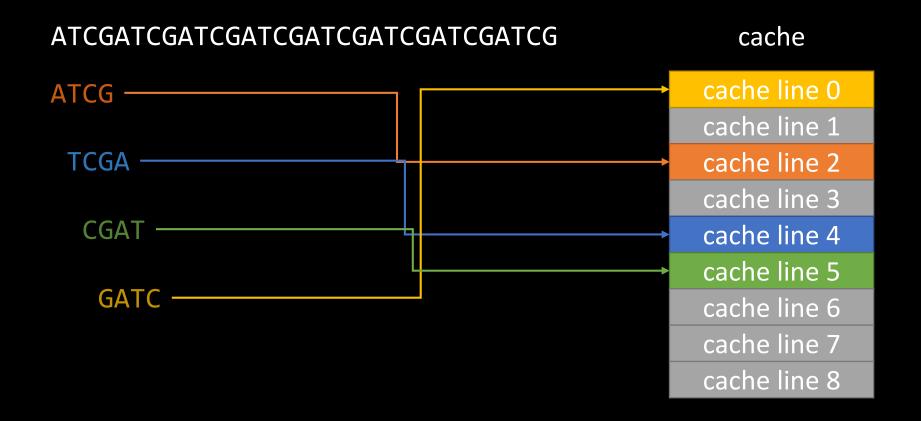
ATCGATCGATCGATCGATCGATCGATCGATCG



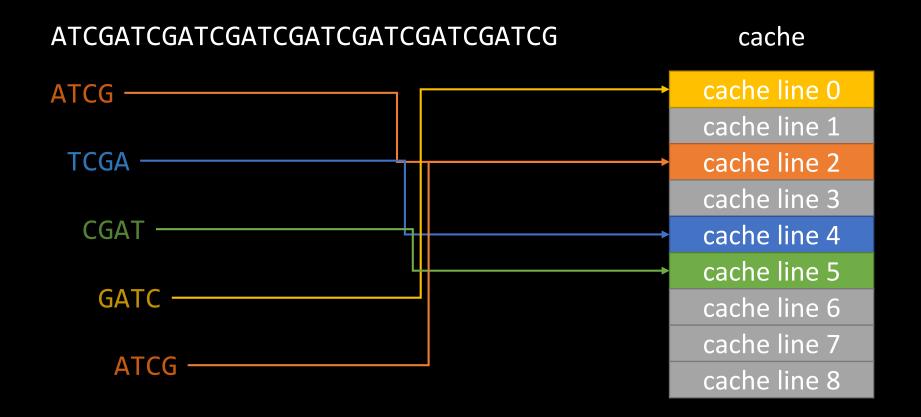




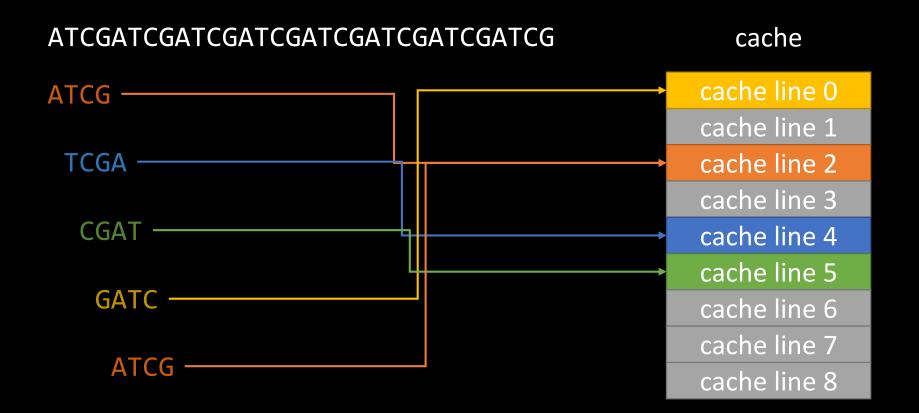




Microsatellites and Processed k-mers



Microsatellites and Processed k-mers



The microsatellite will activate cache lines 2, 4, 5 and 0 repeatedly

SEPTEMBER 9 – 13, 2019 CROSSING Summer School on Sustainable Security & Privacy

Genome Sequencing Attack Results

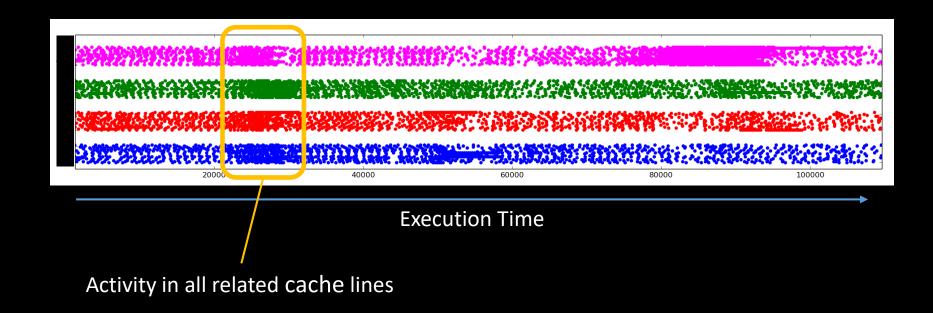
- Monitor cache lines associated to satellite
- High activity in cache lines reveal occurrence of satellite in input string

	1	1		1
	医对射射的分割的	i in the second s		
		an and a second	in an	
			anne a sta	
20000	40000	60000	80000	100000

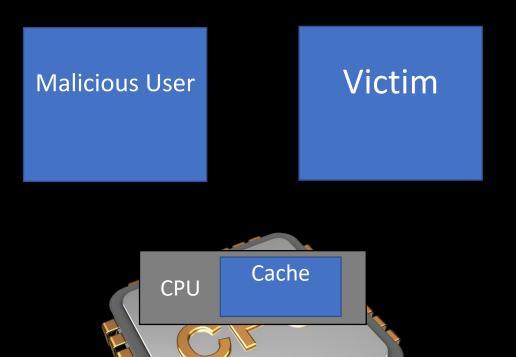
Execution Time

Genome Sequencing Attack Results

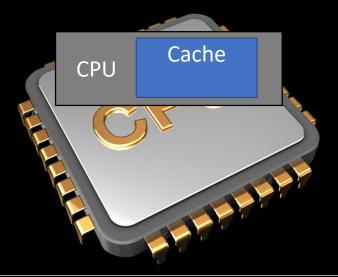
- Monitor cache lines associated to satellite
- High activity in cache lines reveal occurrence of satellite in input string

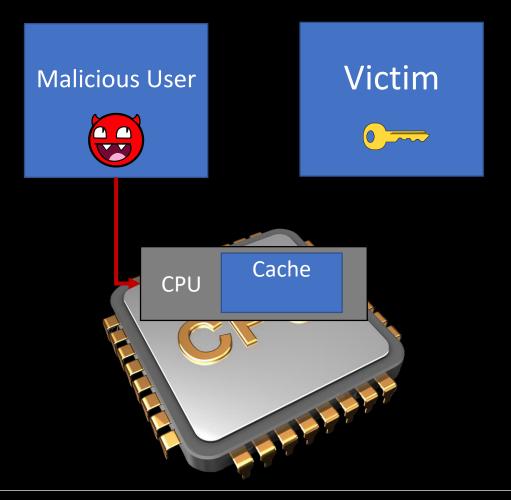


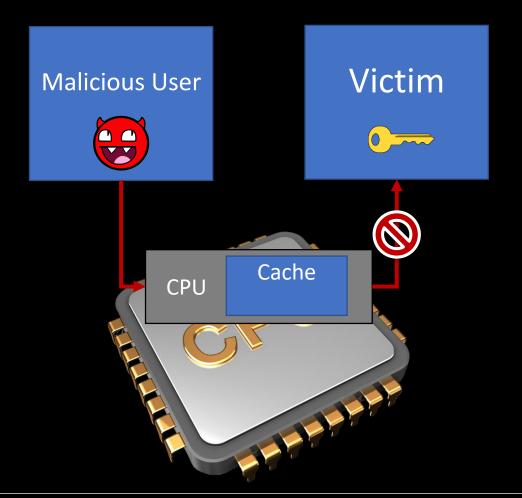
Speculative Execution Attacks

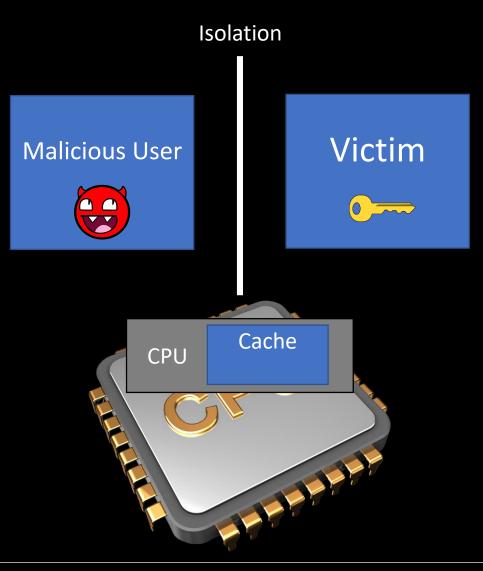


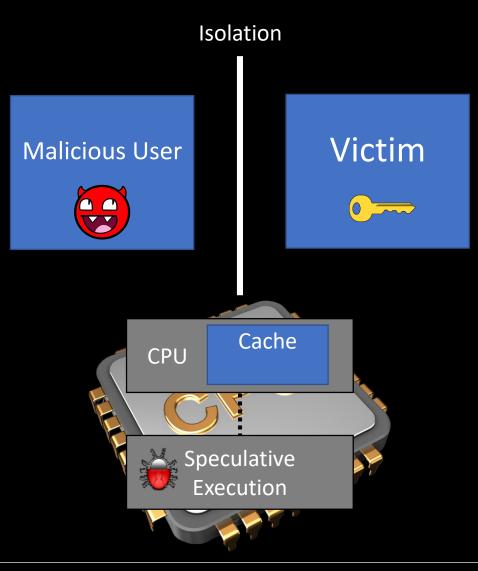


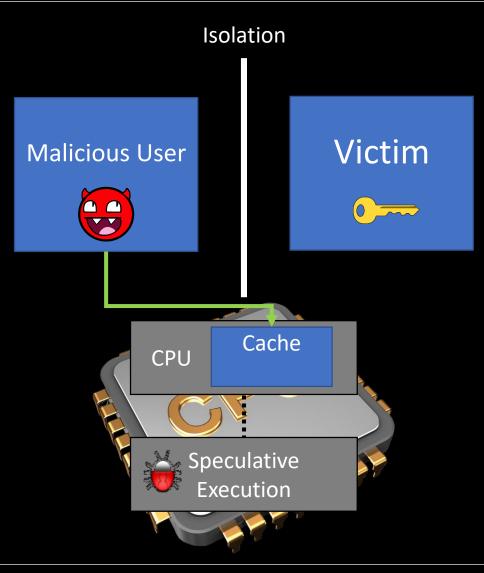












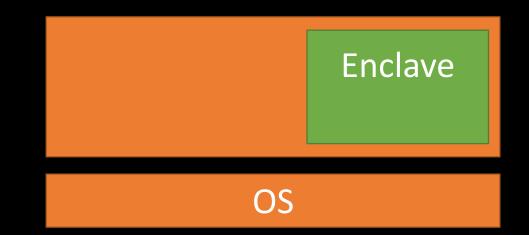
Meltdown

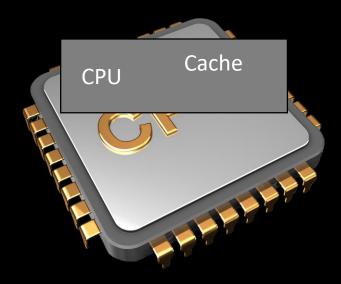
Exploits speculative execution bug

- attacker can read arbitrary physical memory (including kernel memory) from an unprivileged user process
- this can be used, e.g., to break kernel ASLR from unprivileged process
- or, to extract secrets from Intel SGX enclaves!

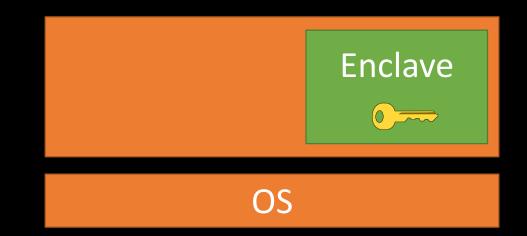


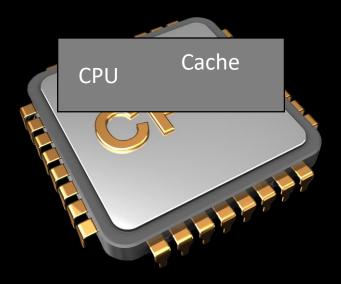
- Foreshadow [Van Bulck, USENIX Security 2018]
 - Extract long-term secrets from Intel Launching and Quoting Enclaves
 - Speculative access only possible for data in L1 cache
- Implications
 - Attacker can bypass vetting of enclaves by Intel
 - Attacker can forge local and remote attestations sent to other enclaves and to remote parties



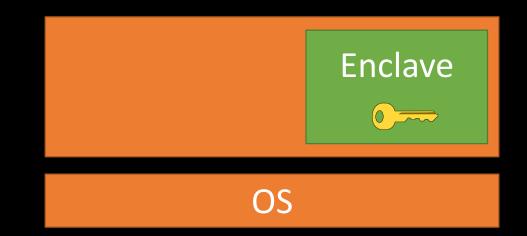


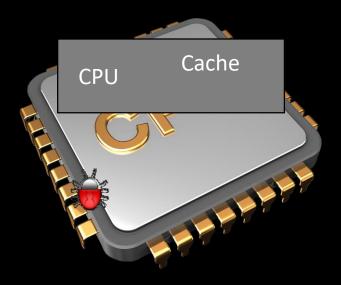
- Foreshadow [Van Bulck, USENIX Security 2018]
 - Extract long-term secrets from Intel Launching and Quoting Enclaves
 - Speculative access only possible for data in L1 cache
- Implications
 - Attacker can bypass vetting of enclaves by Intel
 - Attacker can forge local and remote attestations sent to other enclaves and to remote parties



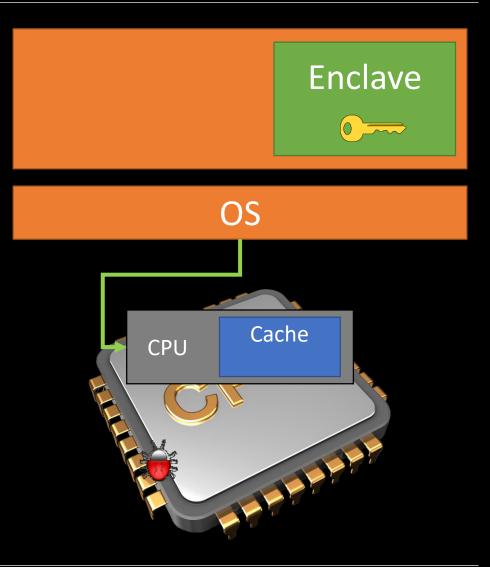


- Foreshadow [Van Bulck, USENIX Security 2018]
 - Extract long-term secrets from Intel Launching and Quoting Enclaves
 - Speculative access only possible for data in L1 cache
- Implications
 - Attacker can bypass vetting of enclaves by Intel
 - Attacker can forge local and remote attestations sent to other enclaves and to remote parties

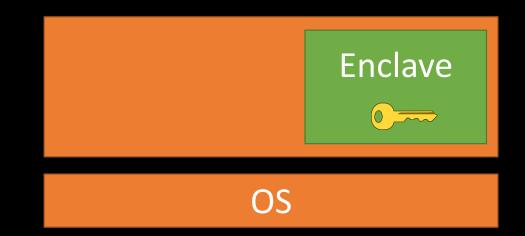


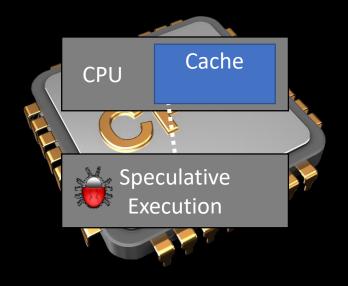


- Foreshadow [Van Bulck, USENIX Security 2018]
 - Extract long-term secrets from Intel Launching and Quoting Enclaves
 - Speculative access only possible for data in L1 cache
- Implications
 - Attacker can bypass vetting of enclaves by Intel
 - Attacker can forge local and remote attestations sent to other enclaves and to remote parties

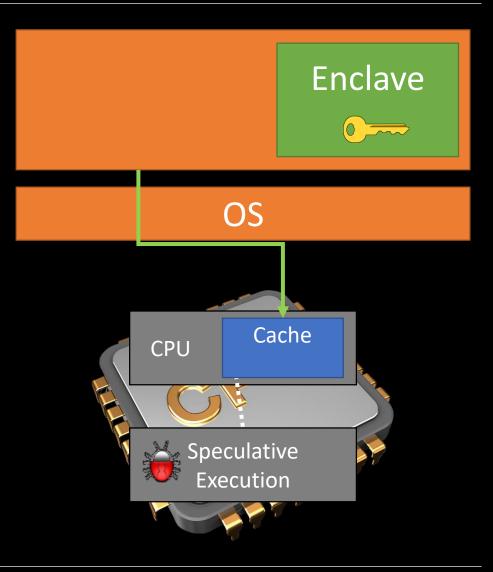


- Foreshadow [Van Bulck, USENIX Security 2018]
 - Extract long-term secrets from Intel Launching and Quoting Enclaves
 - Speculative access only possible for data in L1 cache
- Implications
 - Attacker can bypass vetting of enclaves by Intel
 - Attacker can forge local and remote attestations sent to other enclaves and to remote parties





- Foreshadow [Van Bulck, USENIX Security 2018]
 - Extract long-term secrets from Intel Launching and Quoting Enclaves
 - Speculative access only possible for data in L1 cache
- Implications
 - Attacker can bypass vetting of enclaves by Intel
 - Attacker can forge local and remote attestations sent to other enclaves and to remote parties



How to Get Enclave Data into L1 Cache?

- Run enclave and interrupt when target data was used
 - The enclave's usage of the target data brings it into the cache
- Use SGX paging mechanism
 - OS can swap in/out pages of enclaves
 - When an enclave page is swapped in, its content is loaded into L1 cache
 - Malicious OS can run attack without even running the enclave

Defenses Against Foreshadow

- Flush L1 cache on enclave exit
 - Provided via microcode update
 - Only effective without hyperthreading
- Include hyperthreading configuration in attestation report
 - "[...] the Intel SGX attestation will indicate whether hyperthreading has been enabled by the BIOS." [Intel*]
- Renew SGX keys
 - "The microcode update changes the Security Version Number (SVN) associated with the Intel SGX implementation and provides enclaves on the platform with new sealing and attestation keys." [Intel*]

Defenses Against Foreshadow

- Flush L1 cache on enclave exit
 - Provided via microcode update
 - Only effective without hyperthreading
- Include hyperthreading configuration in attestation report
 - "[...] the Intel SGX attestation will indicate whether hyperthreading has been enabled by the BIOS." [Intel*]
- Renew SGX keys
 - "The microcode update changes the Security Version Number (SVN) associated with the Intel SGX implementation and provides enclaves on the platform with new sealing and attestation keys." [Intel*]

* https://software.intel.com/security-software-guidance/software-guidance/l1-terminal-fault

Alternative Solutions?

Ohne Meltdown-Lücke: Chinesische x86-Prozessoren KX-5000 vorgestellt, Angriff auf AMDs ZEN 2 mit KX-7000 geplant



(Bild: Zhaoxin)

Zhaoxins neue Prozessoren der Serie KX-5000 haben bis zu acht Kerne und sollen nicht für die Meltdown-Lücke anfällig sein. Die übernächste Generation KX-7000 soll es bereits mit AMDs künftigen Zen-2-Prozessoren aufnehmen können.

Alternative Solutions?

Ohne Meltdown-Lücke: Chinesische x86-Prozessoren KX-5000 vorgestellt, Angriff auf AMDs ZEN 2 mit KX-7000 geplant



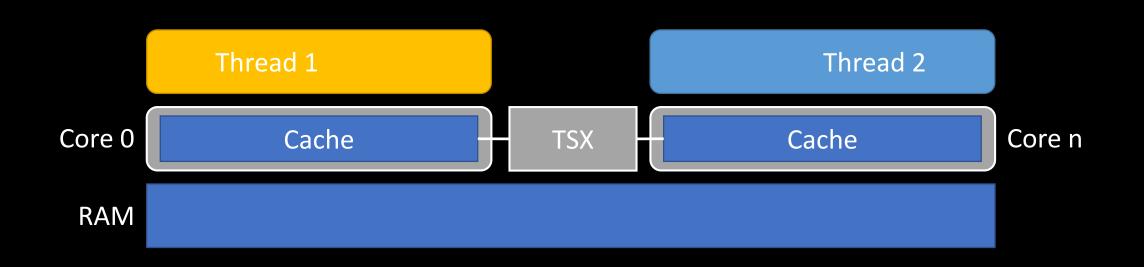
(Bild: Zhaoxin)

Zhaoxins neue Prozessoren der Serie KX-5000 haben bis zu acht Kerne und sollen nicht für die Meltdown-Lücke anfällig sein. Die übernächste Generation KX-7000 soll es bereits mit AMDs künftigen Zen-2-Prozessoren aufnehmen können.

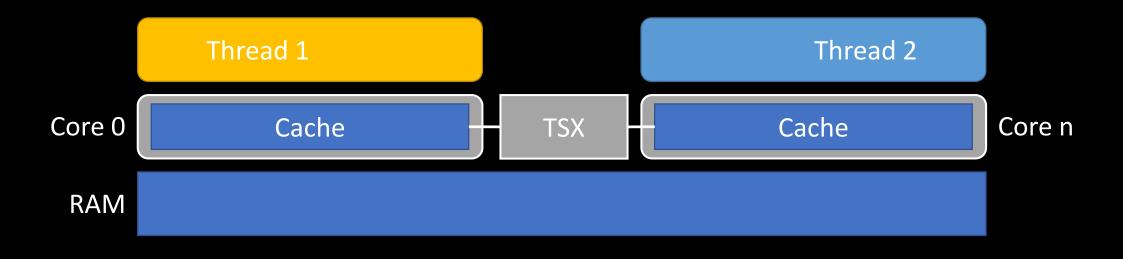
Buy Chinese Quality Chips, not cheap American copies!



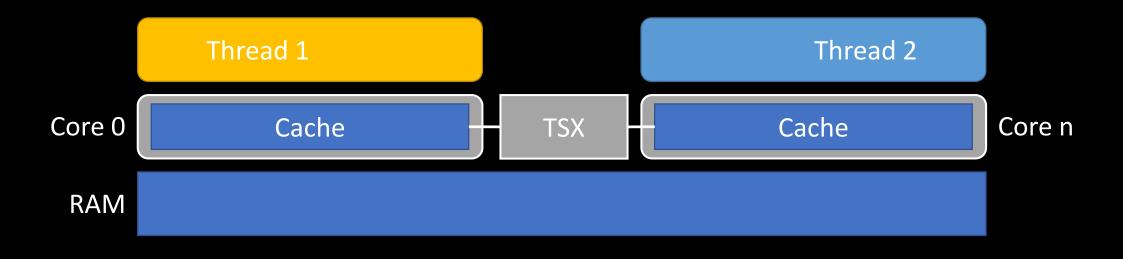
Side-Channel Defenses Using TSX



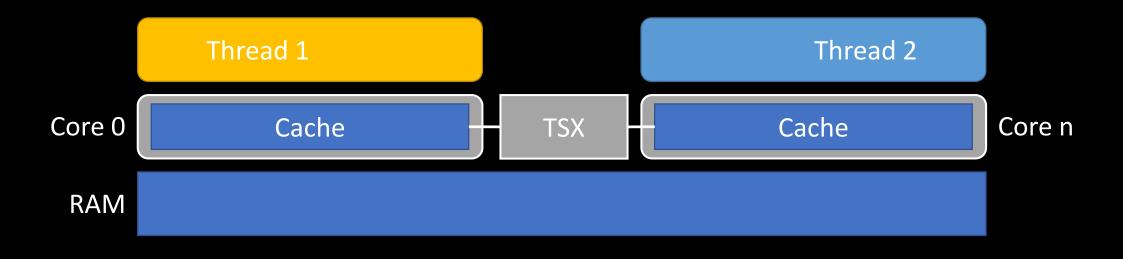
- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



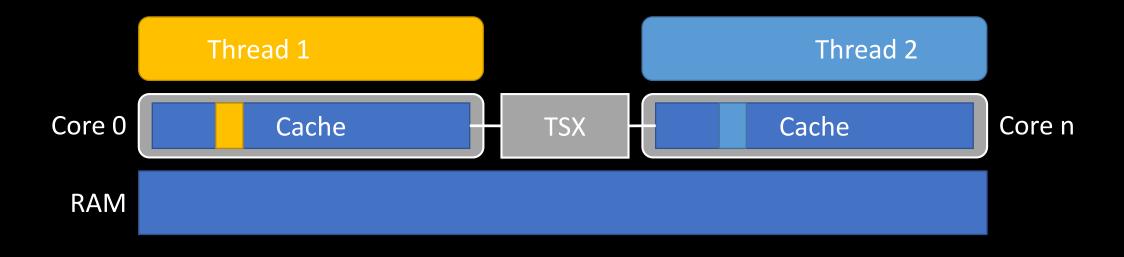
- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



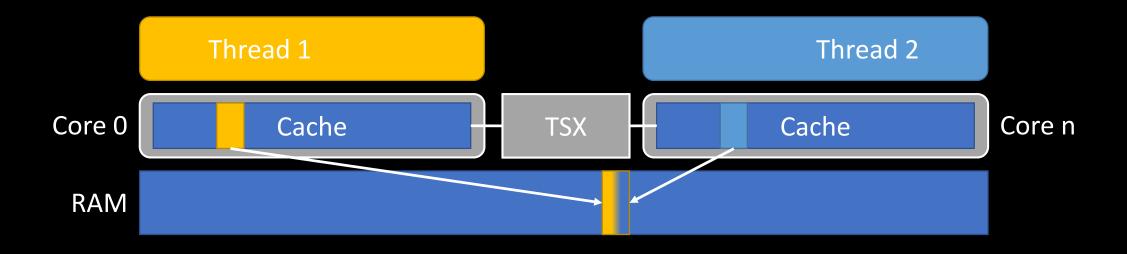
- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



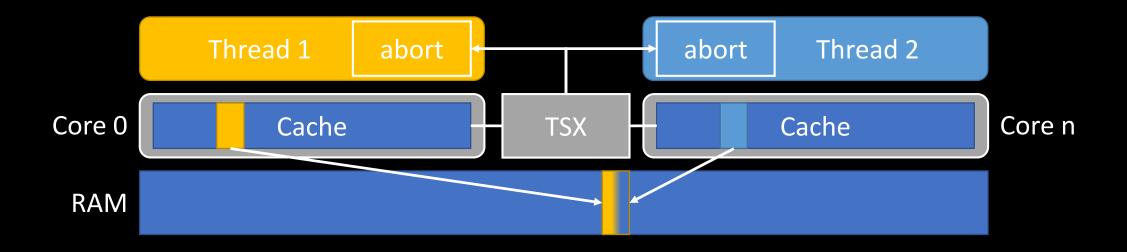
- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



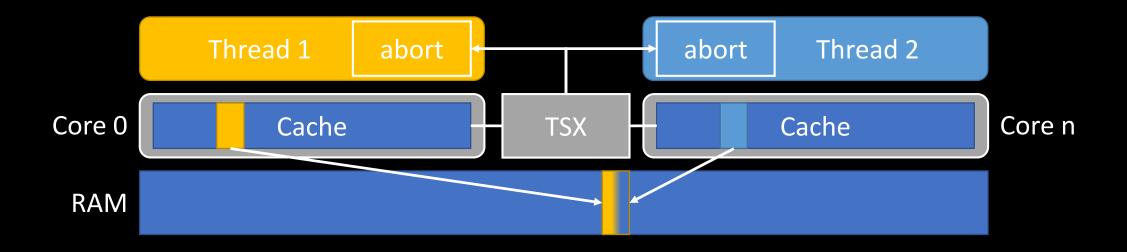
- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



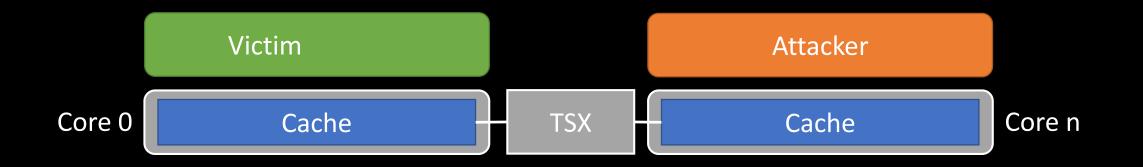
- Intel implementation of Hardware Transactional Memory (HTM)
- Designed for high-performance concurrency
- Allows synchronous memory transactions
- TSX is **not** available on all SGX-enable processors



SGX Specific Side-Channel Defenses Using TSX

Detecting enclave's interruption

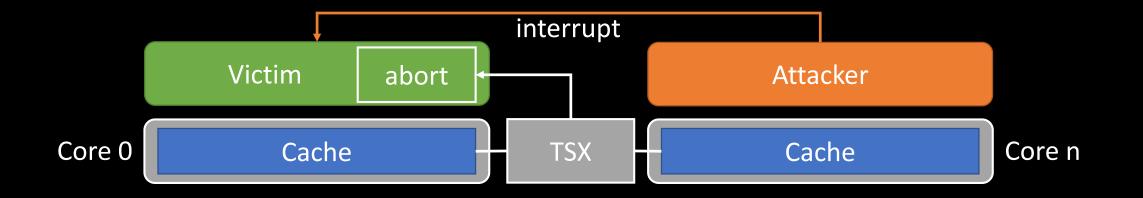
- Frequent interrupts evidence for side-channel attack
- T-SGX: Uses TSX feature to detect enclave interrupt [Shih et al., NDSS'17]
- Déjà Vu : Uses TSX to detect enclave slowdown [Chen et al., AsiaCCS'17]



SGX Specific Side-Channel Defenses Using TSX

Detecting enclave's interruption

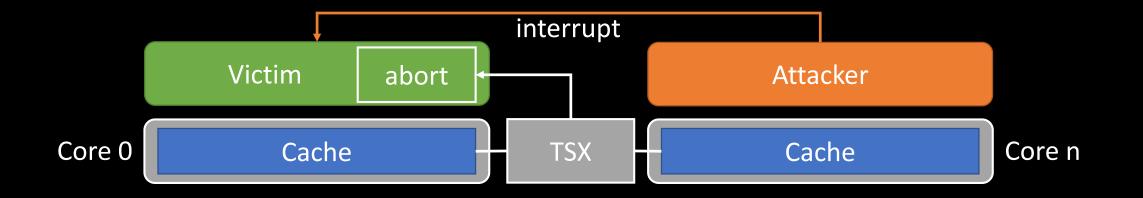
- Frequent interrupts evidence for side-channel attack
- T-SGX: Uses TSX feature to detect enclave interrupt [Shih et al., NDSS'17]
- Déjà Vu : Uses TSX to detect enclave slowdown [Chen et al., AsiaCCS'17]



SGX Specific Side-Channel Defenses Using TSX

Detecting enclave's interruption

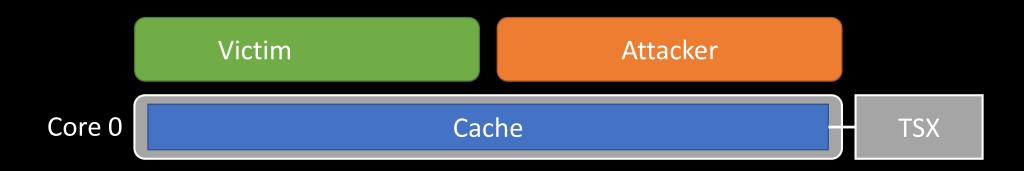
- Frequent interrupts evidence for side-channel attack
- T-SGX: Uses TSX feature to detect enclave interrupt [Shih et al., NDSS'17]
- Déjà Vu : Uses TSX to detect enclave slowdown [Chen et al., AsiaCCS'17]



SGX Specific Side-Channel Defenses Using TSX

Detecting cache evictions

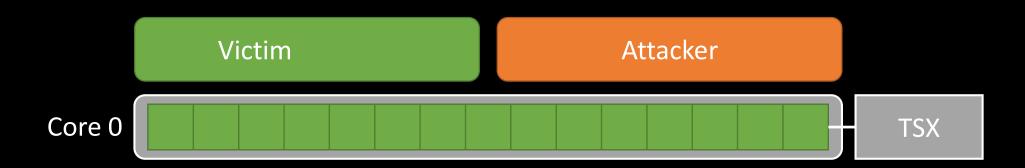
- Eviction of the victim's cache entries could lead to information leakage
- Cloak: Prime cache before accessing sensitive data [Schuster et al., USENIX 2017]



SGX Specific Side-Channel Defenses Using TSX

Detecting cache evictions

- Eviction of the victim's cache entries could lead to information leakage
- Cloak: Prime cache before accessing sensitive data [Schuster et al., USENIX 2017]



SGX Specific Side-Channel Defenses Using TSX

Detecting cache evictions

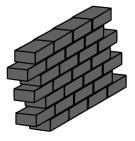
- Eviction of the victim's cache entries could lead to information leakage
- Cloak: Prime cache before accessing sensitive data [Schuster et al., USENIX 2017]



Temporal cache isolation

Temporal cache isolation

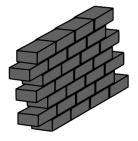
Cache partitioning / coloring







Cache partitioning / coloring

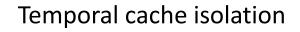


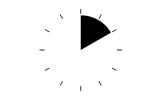




Temporal Cache Isolation

- Flush on each context switch
- Ineffective on SMT-enabled systems where caches are shared contemporaneously
- E.g., [Costan et al., USENIX Sec'16]







49



SMT: Simultaneous Multithreading

SEPTEMBER 9 – 13, 2019 CROSSING Summer School on Sustainable Security & Privacy

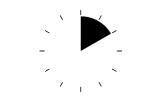
Temporal Cache Isolation

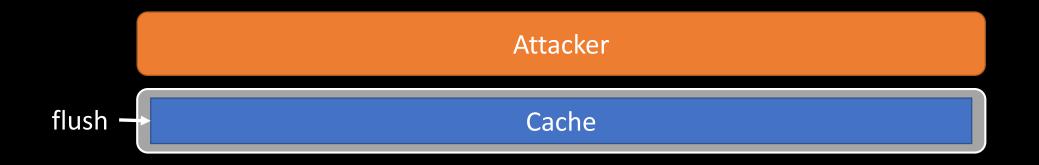
Flush on each context switch

SMT: Simultaneous Multithreading

- Ineffective on SMT-enabled systems where caches are shared contemporaneously
- E.g., [Costan et al., USENIX Sec'16]





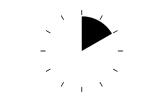


SEPTEMBER 9 – 13, 2019 CROSSING Summer School on Sustainable Security & Privacy

Temporal Cache Isolation

- Flush on each context switch
- Ineffective on SMT-enabled systems where caches are shared contemporaneously
- E.g., [Costan et al., USENIX Sec'16]







SMT: Simultaneous Multithreading

Cache Partitioning / Coloring

- Reduces the amount of cache available to individual software
- E.g., [Domnister et al., TACO'12]

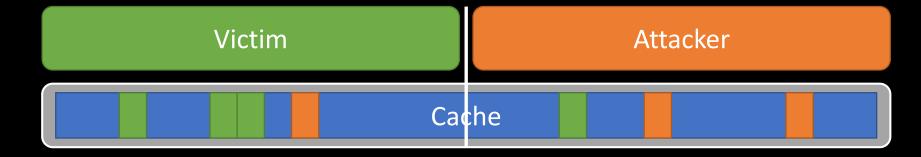




Cache partitioning / coloring

Cache Partitioning / Coloring

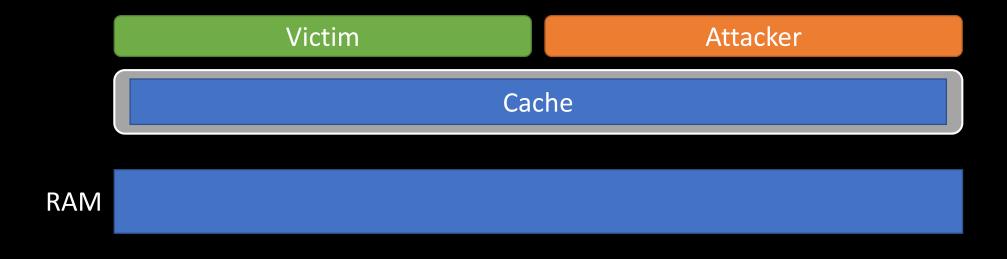
- Reduces the amount of cache available to individual software
- E.g., [Domnister et al., TACO'12]

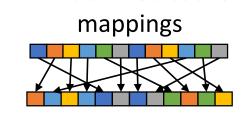




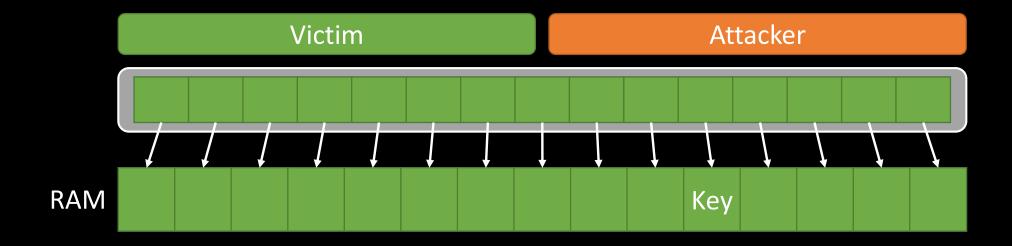
Cache partitioning / coloring

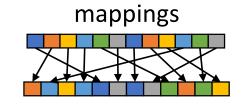
- Adversary cannot link cache observation with memory locations
- Frequency analysis or predictable access patterns can reveal randomization secret
- E.g., [Wang et al., ISCA'07]



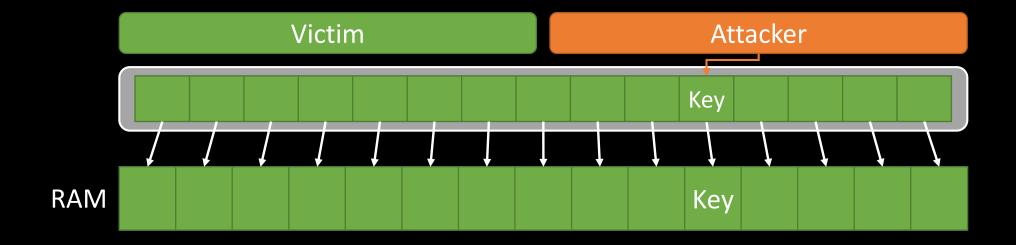


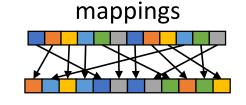
- Adversary cannot link cache observation with memory locations
- Frequency analysis or predictable access patterns can reveal randomization secret
- E.g., [Wang et al., ISCA'07]





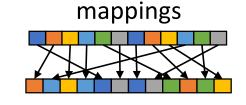
- Adversary cannot link cache observation with memory locations
- Frequency analysis or predictable access patterns can reveal randomization secret
- E.g., [Wang et al., ISCA'07]



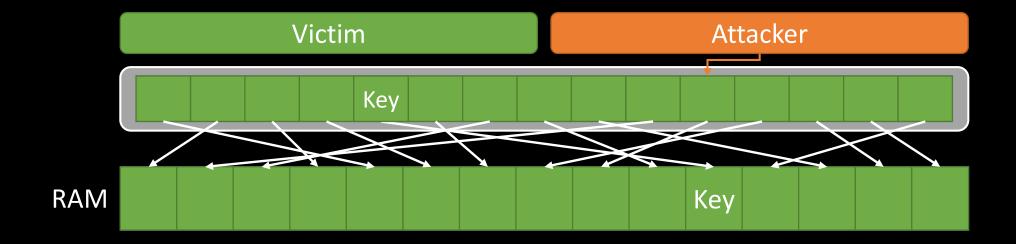


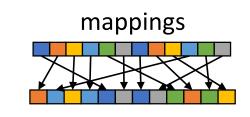
- Adversary cannot link cache observation with memory locations
- Frequency analysis or predictable access patterns can reveal randomization secret
- E.g., [Wang et al., ISCA'07]





- Adversary cannot link cache observation with memory locations
- Frequency analysis or predictable access patterns can reveal randomization secret
- E.g., [Wang et al., ISCA'07]



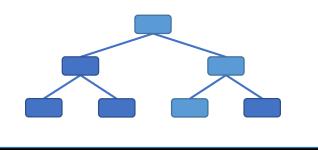


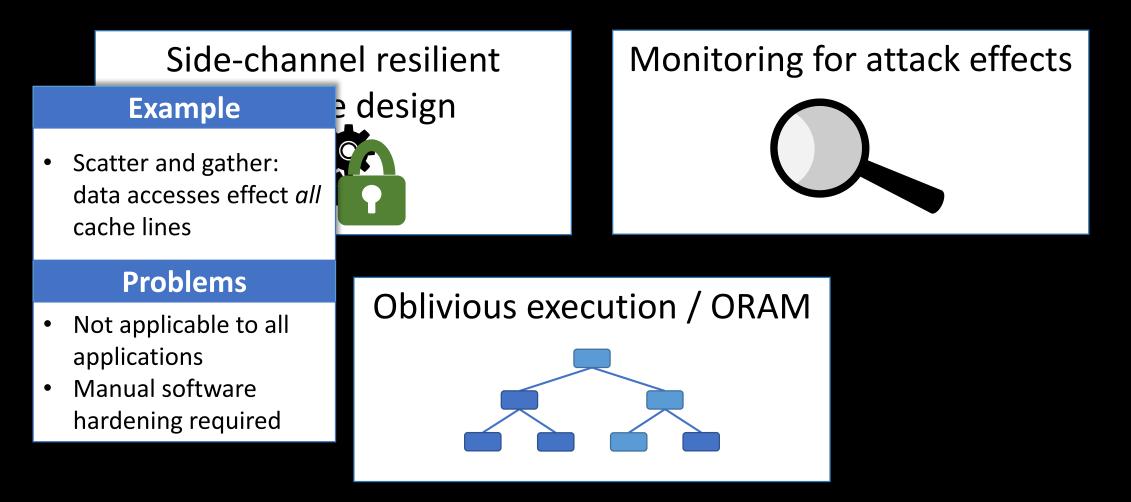


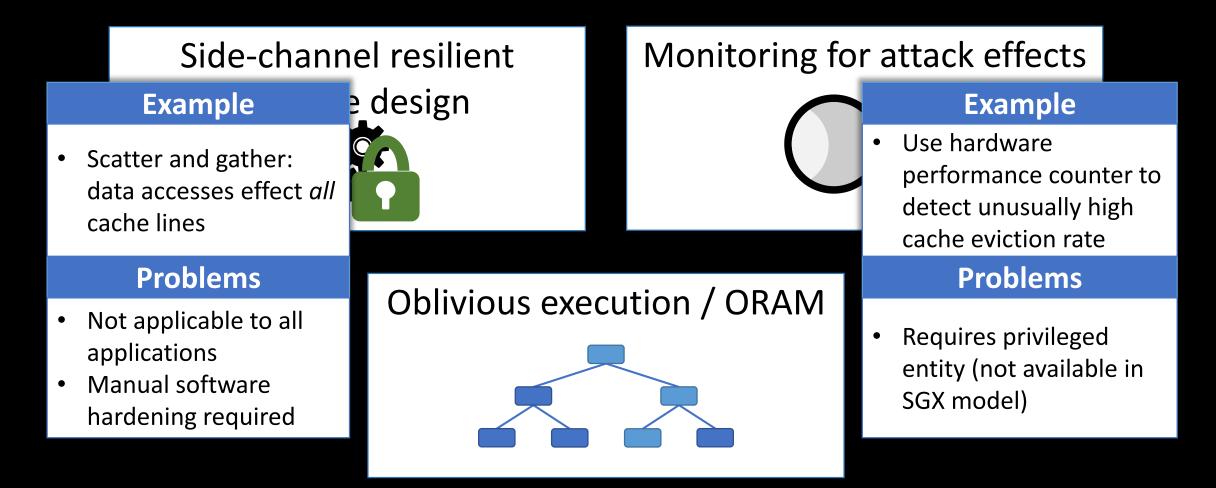
Monitoring for attack effects

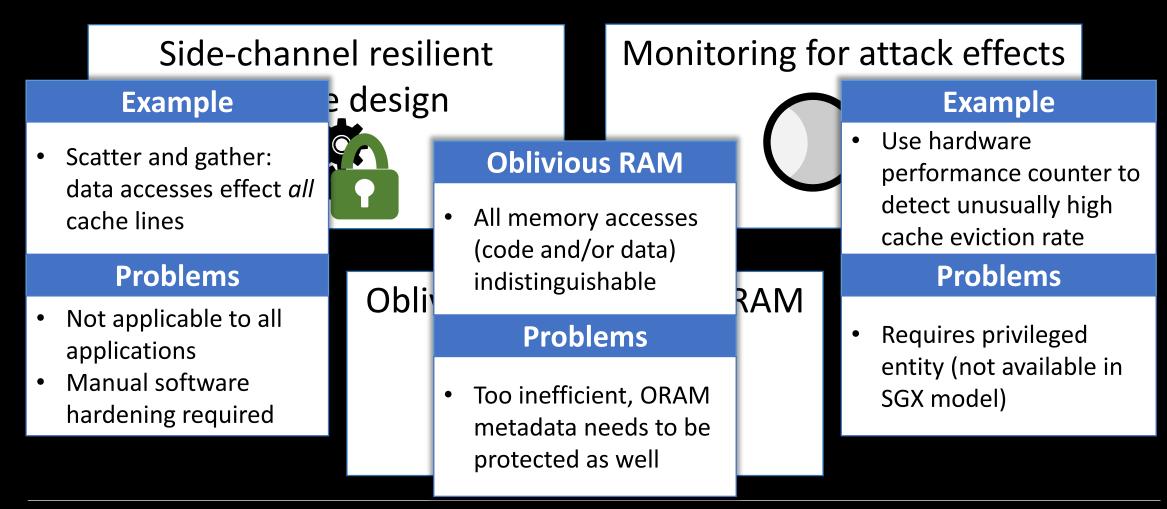










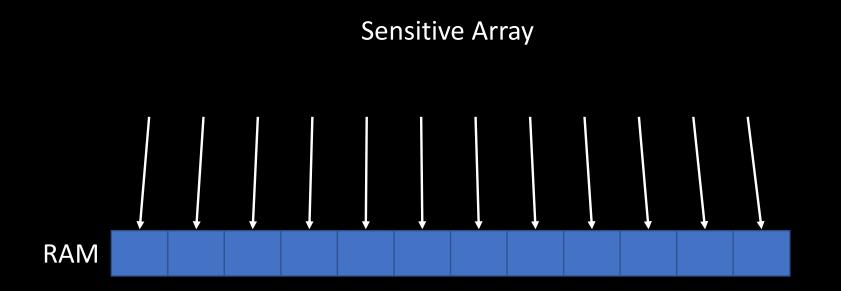


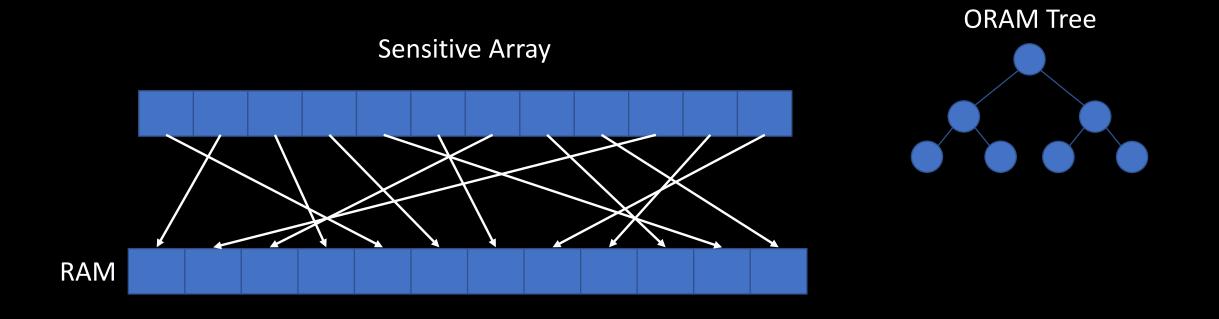
Our Recent Work: DR.SGX: Automated and Adjustable Side-Channel Protection for SGX using Data Location Randomization

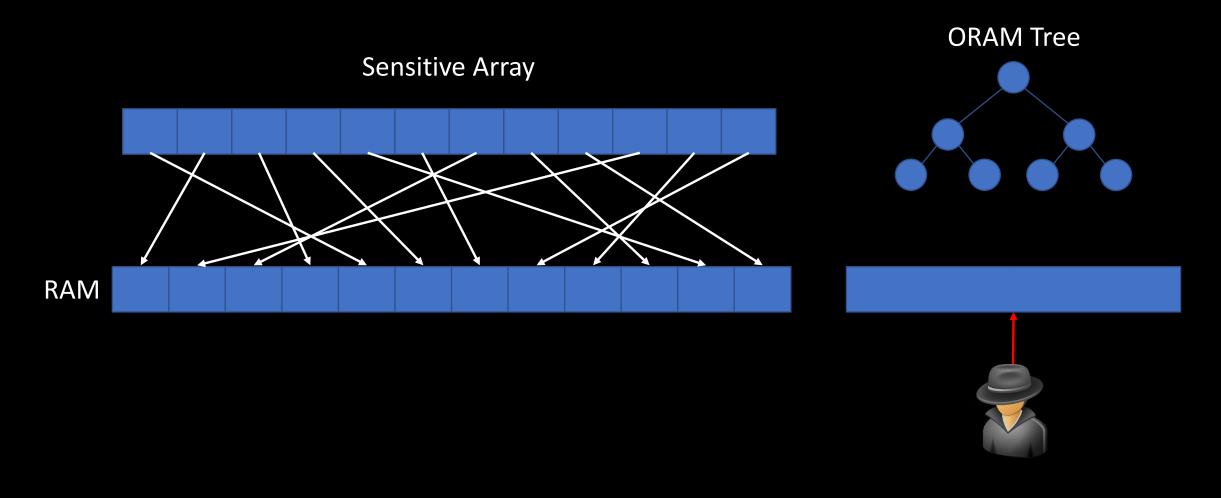
[Brasser et al., ACSAC 2019]

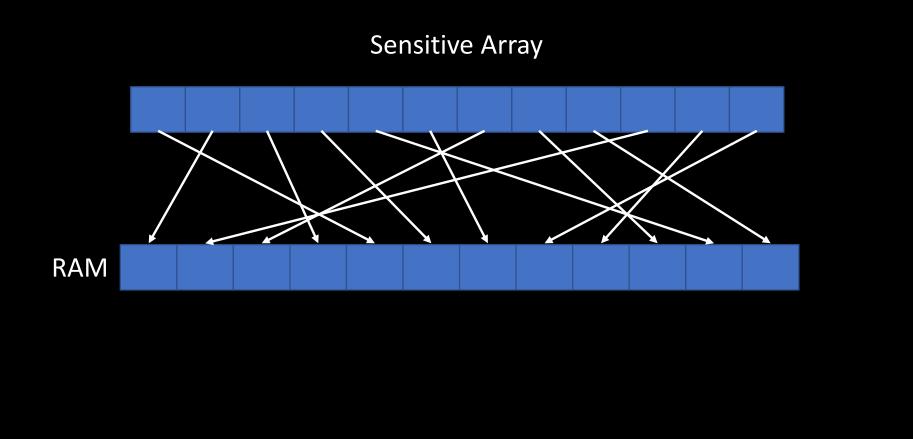
DR.SGX: Objective and Approach

- **Objective:** Similarly to ORAM, make memory accesses indistinguishable
 - but at a cheaper cost
 - without relying on meta-data that needs protection
- Approach: Runtime fine-grained data location randomization
 - format-preserving encryption to determine location of randomized data
 - only small constant-size metadata needed
 - compiler-based approach (no annotations needed)
 - gradual randomization, interleaved with enclave execution
 - configurable re-randomization rate







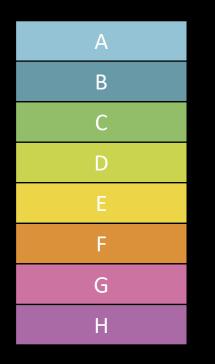




DR.SGX Re-randomization

FFX Format-Preserving Encryption scheme with AES as a block cipher

Initial layout



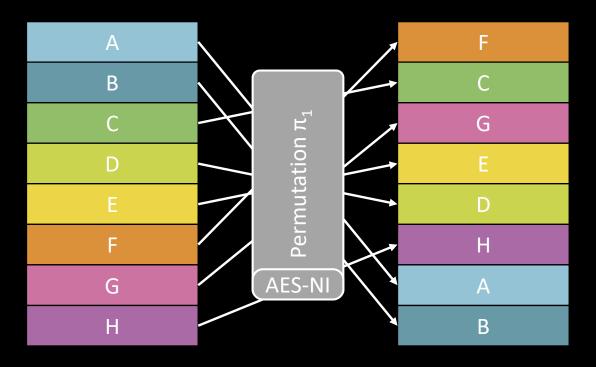
Time

DR.SGX Re-randomization

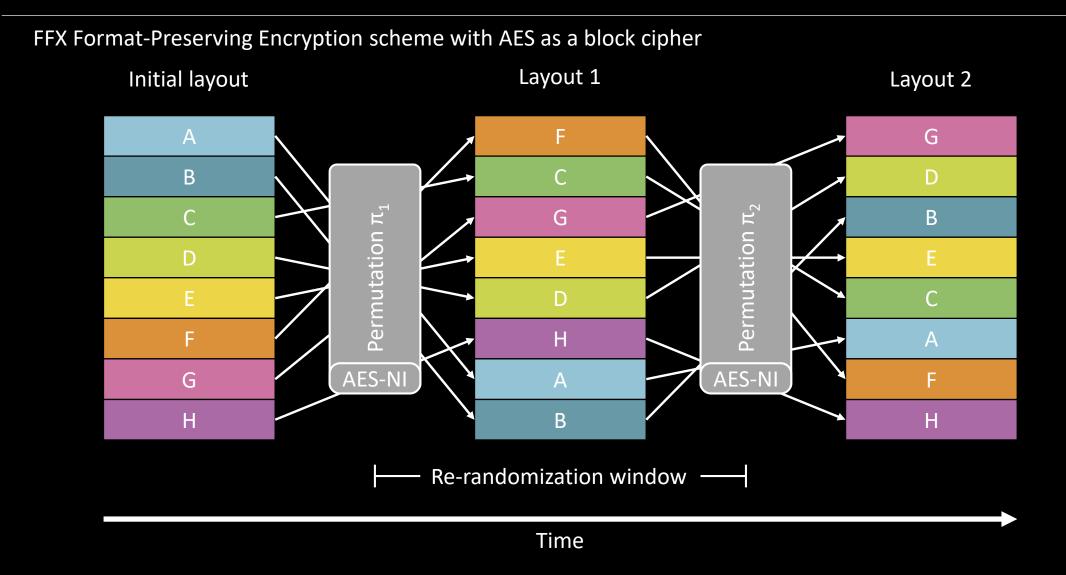
FFX Format-Preserving Encryption scheme with AES as a block cipher

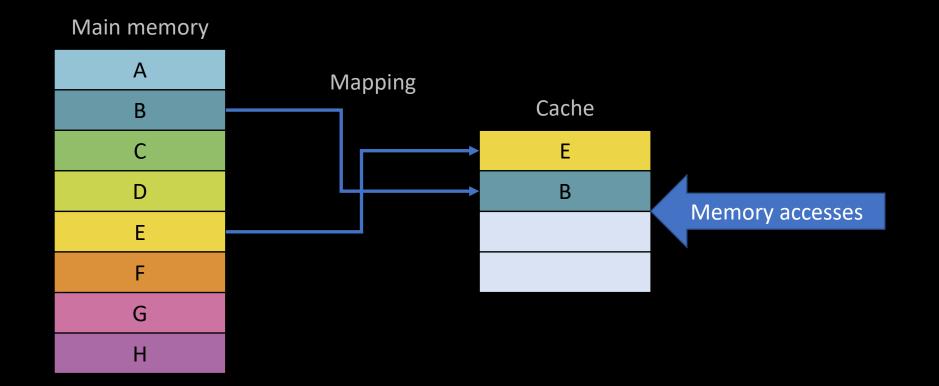
Initial layout

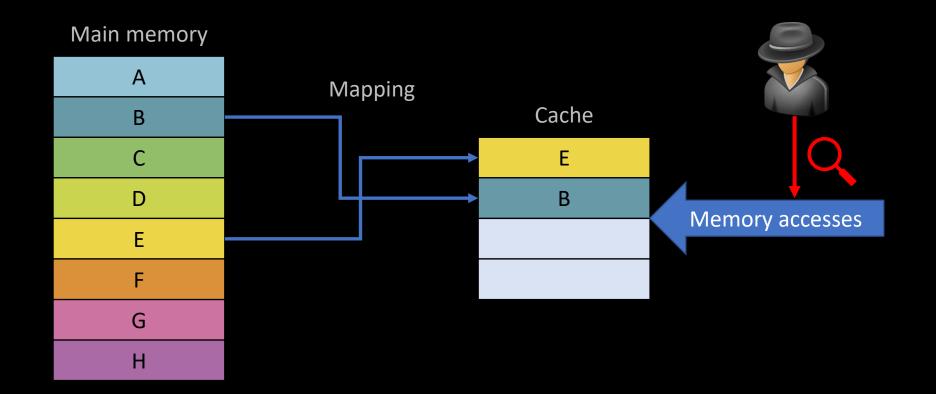
Layout 1

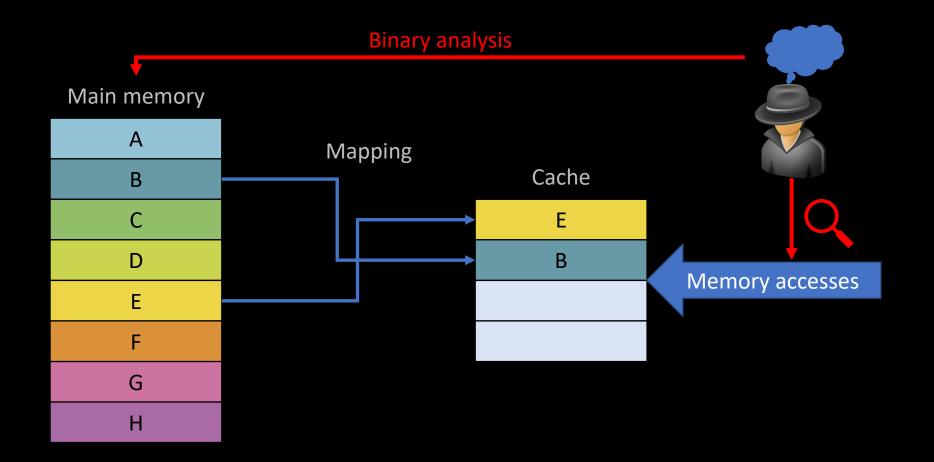


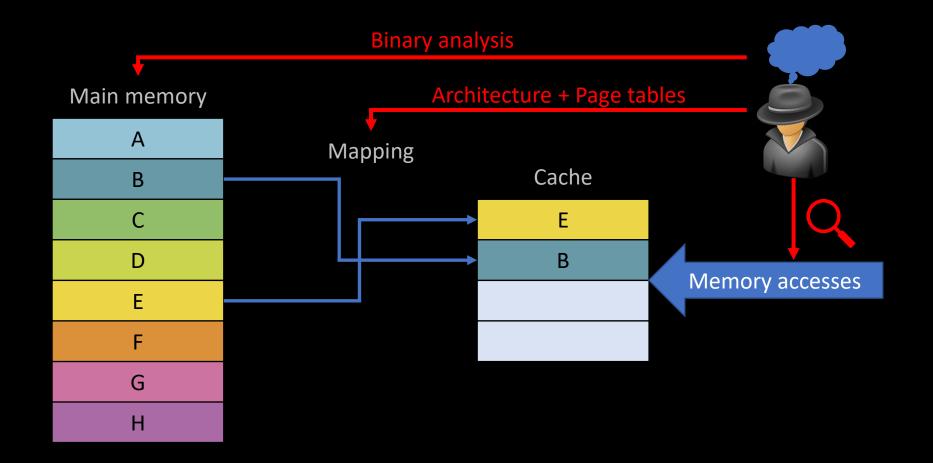
DR.SGX Re-randomization

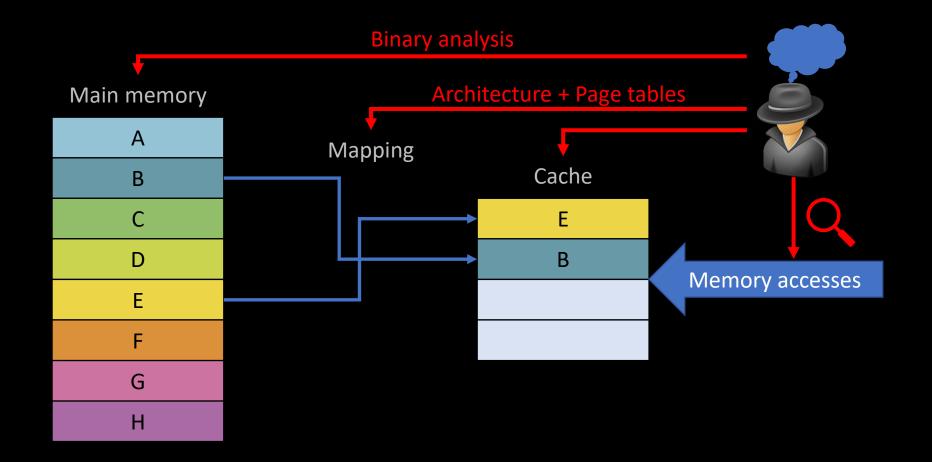


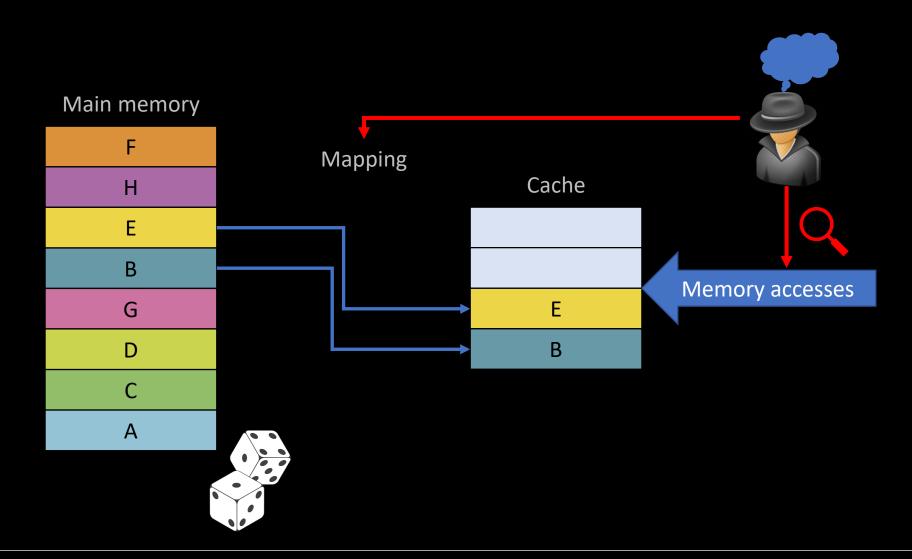


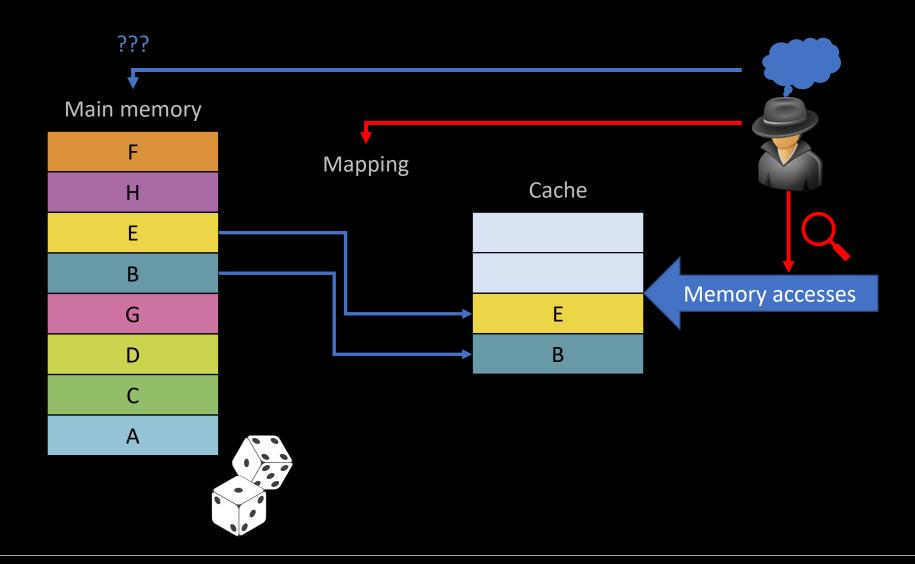


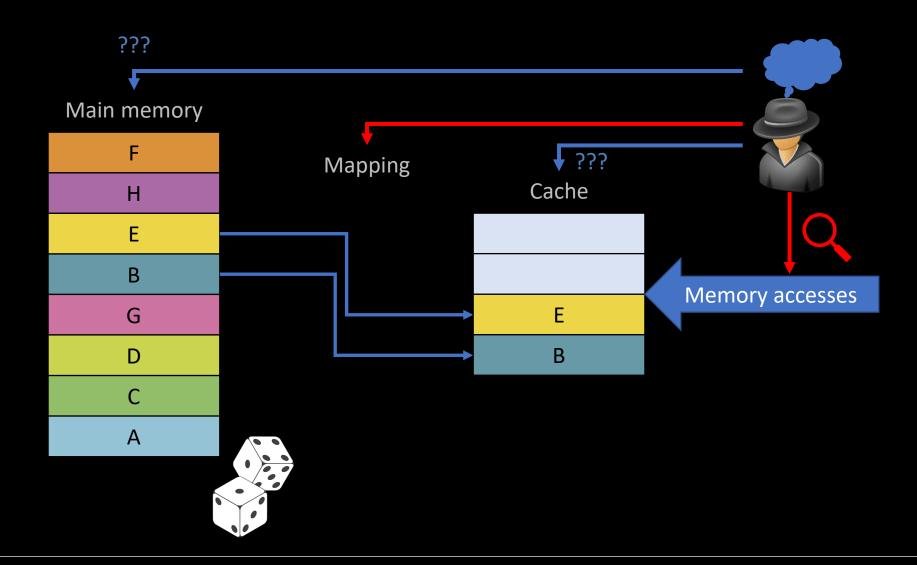






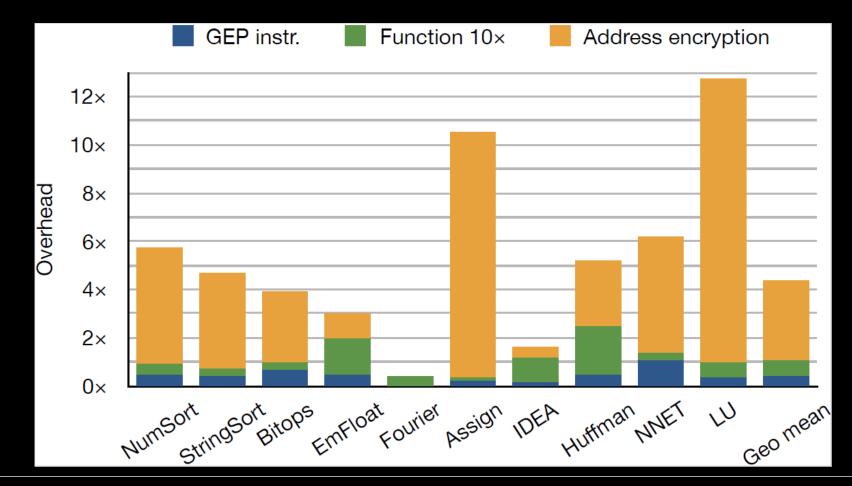






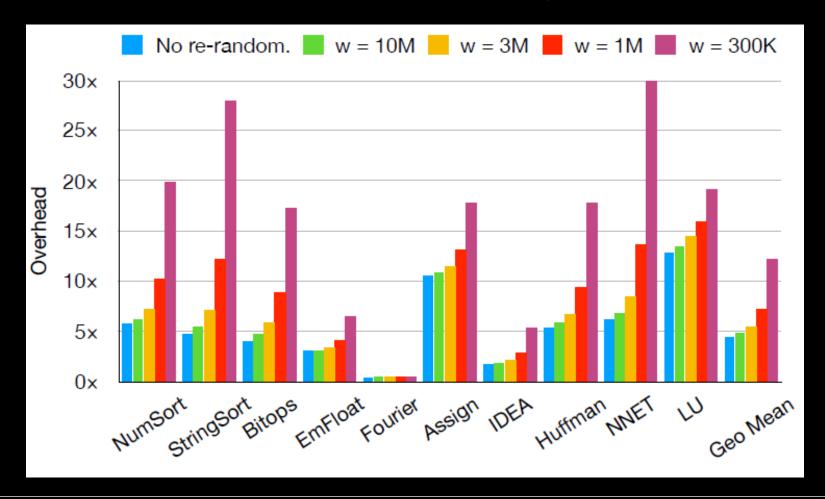
Performance Evaluation using Nbench

• Without runtime re-randomization (geometric mean about 4x)



Performance Evaluation using Nbench

• With different re-randomization windows (geometric mean up to 12x)



ORAM vs. Dr.SGX: Performance Comparison

- Obfuscuro [Ahmad et al., NDSS 2019]
 - Obfuscation engine on Intel SGX
 - Implements both, ORAM and oblivious execution
 - Performance overheads of 83x on average and up to 220x

• Dr. SGX

- Performance overhead 4x 12x
 - at least one order of magnitude lower than Obfuscuro
- Allows developers to balance between increased side-channel protection and the performance cost based on adjustable security parameter

Conclusion

- Great concepts suffer from implementation problems
- Intel SGX is no exception
- Side-channel attacks are a major threat to Intel SGX
 - Were deemed as 'too difficult' and were left out of the attacker model
 - Research has shown it otherwise
 - Attacks still can be improved through more automation
- Countermeasures
 - Range from specific protections against particular problems to generic solutions
 - Generic solutions, however, come at significant (prohibitive?) cost
 - There is a need for more efficient generic solutions

