

### Software-based Microarchitectural Attacks

#### **Daniel Gruss**

July 8, 2018

Graz University of Technology







**Operating System** 



#### Application

Untrusted part
Create Enclave

**Operating System** 



#### Application

**Operating System** 



#### Application

**Operating System** 

Daniel Gruss — Graz University of Technology



Application

**Operating System** 

Daniel Gruss — Graz University of Technology

Untrusted part Trusted part Create Enclave Trusted Fnc. Call Call Trusted Fnc.

Application

**Operating System** 

Daniel Gruss — Graz University of Technology

Untrusted part Trusted part Create Enclave Trusted Fnc. Call Call Trusted Fnc. Return

Application

**Operating System** 

Daniel Gruss — Graz University of Technology



Application

**Operating System** 

Daniel Gruss — Graz University of Technology



Application

**Operating System** 

Daniel Gruss — Graz University of Technology

Untrusted part Trusted part Create Enclave Trusted Fnc. æ Call Trusted Fnc. Return **Operating System** 

Application

Daniel Gruss — Graz University of Technology



















# $\overrightarrow{B}$











#### **Protection from Side-Channel Attacks**

#### **Protection from Side-Channel Attacks**

Intel SGX does not provide explicit protection from side-channel attacks.

#### **Protection from Side-Channel Attacks**

Intel SGX does not provide explicit protection from side-channel attacks. It is the enclave developer's responsibility to address side-channel attack concerns.

## **CAN'T BREAK YOUR SIDE-CHANNEL PROTECTIONS**

# IF YOU DON'T HAVE ANY

imgflip.com





- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX



- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX

#### Teechain

 $\left[\ldots\right]$  We assume the TEE guarantees to hold





- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX

#### **Teechain**

[...] We assume the TEE guarantees to hold and do not consider side-channel attacks [5, 35, 46] on the TEE.





- Ledger SGX Enclave for blockchain applications
- BitPay Copay Bitcoin wallet
- Teechain payment channel using SGX

#### Teechain

[...] We assume the TEE guarantees to hold and do not consider side-channel attacks [5, 35, 46] on the TEE. Such attacks and their mitigations [36, 43] are outside the scope of this work. [...]

Raw Prime+Probe trace...<sup>1</sup>



<sup>1</sup>M. Schwarz, D. Gruss, S. Weiser, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

Daniel Gruss — Graz University of Technology

...processed with a simple moving average...<sup>2</sup>



<sup>2</sup>M. Schwarz, D. Gruss, S. Weiser, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

...allows to clearly see the bits of the exponent<sup>3</sup>



<sup>3</sup>M. Schwarz, D. Gruss, S. Weiser, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.











• Usually no physical access



- Usually no physical access
- Local code



- Usually no physical access
- Local code
- Co-located code


- Usually no physical access
- Local code
- Co-located code
- Different meanings of "remote"



- Usually no physical access
- Local code
- Co-located code
- Different meanings of "remote"
  - 1. Attacker controls code in browser sandbox (e.g., [Ore+15; GMM16])



- Usually no physical access
- Local code
- Co-located code
- Different meanings of "remote"
  - 1. Attacker controls code in browser sandbox (e.g., [Ore+15; GMM16])
  - 2. Attacker cannot control any code on the system



## TIMING IS EVERYTHING









### 1337 4242

#### FOOD CACHE

#### Revolutionary concept!

Store your food at home, never go to the grocery store during cooking.

Can store **ALL** kinds of food.

ONLY TODAY INSTEAD OF \$1,300



ORDER VIA PHONE: +555 12345



**CPU** Cache



\_

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

Daniel Gruss — Graz University of Technology















www.tugraz.at

**CPU** Cache



















Daniel Gruss — Graz University of Technology



Daniel Gruss — Graz University of Technology



• theoretical maximum accuracy of  $5.4 \cdot 10^{-44}$ s



- theoretical maximum accuracy of  $5.4 \cdot 10^{-44}$ s
- feasible today:  $850 \cdot 10^{-21}$ s



- theoretical maximum accuracy of  $5.4 \cdot 10^{-44}$ s
- feasible today:  $850 \cdot 10^{-21}$ s

Microarchitectural Attacks



- theoretical maximum accuracy of  $5.4 \cdot 10^{-44}$ s
- feasible today:  $850 \cdot 10^{-21}$ s

Microarchitectural Attacks

• often around nanoseconds



- theoretical maximum accuracy of  $5.4 \cdot 10^{-44}$ s
- feasible today:  $850 \cdot 10^{-21}$ s

Microarchitectural Attacks

- often around nanoseconds
- sometimes much lower



• in the range of multiple GHz



• in the range of multiple GHz

Microarchitectural Attacks



• in the range of multiple GHz

Microarchitectural Attacks

• usually varying frequency (depending on the attack)



• in the range of multiple GHz

Microarchitectural Attacks

- usually varying frequency (depending on the attack)
- between a few ns (< 1 GHz) and multiple seconds (< 1 Hz) (or even worse)



Indud
Physical Side Channels

• in the range of multiple GHz

Microarchitectural Attacks

- usually varying frequency (depending on the attack)
- between a few ns (< 1 GHz) and multiple seconds (< 1 Hz) (or even worse)
- strongly dependent on the specific attack

Imiliar

Physical Side Channels

• in the range of multiple GHz

Microarchitectural Attacks

- usually varying frequency (depending on the attack)
- between a few ns (< 1 GHz) and multiple seconds (< 1 Hz) (or even worse)
- strongly dependent on the specific attack
  - device under test = measurement device

Imimi

Physical Side Channels

• in the range of multiple GHz

Microarchitectural Attacks

- usually varying frequency (depending on the attack)
- between a few ns (< 1 GHz) and multiple seconds (< 1 Hz) (or even worse)
- strongly dependent on the specific attack
  - device under test = measurement device
  - observer effect



device under test = measurement device

- measuring time takes some time
- limits the resolution
- measuring cache hits/misses manipulates the cache state
- virtually all measurements are destructive









• Race condition between attacker and victim (observer effect)



- Race condition between attacker and victim (observer effect)
- Speculative execution



- Race condition between attacker and victim (observer effect)
- Speculative execution
- Prefetching



- Race condition between attacker and victim (observer effect)
- Speculative execution
- Prefetching
- ...





- Race condition between attacker and victim (observer effect)
- Speculative execution
- Prefetching
- ...
- $\rightarrow$  Typically >99.99% precision and recall





# **Measuring Processor Operations**

- Very short timings
- rdtsc instruction: "cycle-accurate" timestamps

[...] rdtsc function() rdtsc [...]

- Do you measure what you think you measure?
- $\mathit{Out-of-order}\xspace$  execution  $\to$  what is really executed

rdtsc	rdtsc	rdtsc
function()	[]	rdtsc
[]	rdtsc	<pre>function()</pre>
rdtsc	function()	[]



• use pseudo-serializing instruction rdtscp (recent CPUs)

- use pseudo-serializing instruction rdtscp (recent CPUs)
- and/or use serializing instructions like cpuid

- use pseudo-serializing instruction rdtscp (recent CPUs)
- and/or use serializing instructions like cpuid
- and/or use fences like mfence

- use pseudo-serializing instruction rdtscp (recent CPUs)
- and/or use serializing instructions like cpuid
- and/or use fences like mfence

Intel, How to Benchmark Code Execution Times on Intel IA-32 and IA-64 Instruction Set Architectures White Paper, December 2010.

AUGUST 22, 2018 BY BRUCE

Intel Publishes Microcode Security

Patches, No Benchmarking Or

Comparison Allowed!

UPDATE: Intel has resolved their microcode licensing issue which I complained about in this blog post. The new license text is here.



### Cache Hits



www.tugraz.at

#### Cache Hits Cache Misses





## **Temporal Component**





## **Temporal Component**







• Flush+Reload had beautifully nice timings, right?



- Flush+Reload had beautifully nice timings, right?
- Well... steps of 2-4 cycles



- Flush+Reload had beautifully nice timings, right?
- Well... steps of 2-4 cycles
  - only 35-70 steps between hits and misses



- Flush+Reload had beautifully nice timings, right?
- Well... steps of 2-4 cycles
  - only 35-70 steps between hits and misses
- On some devices only 1-2 steps!



• We can build our own timer [Lip+16; Sch+17]



- We can build our own timer [Lip+16; Sch+17]
- Start a thread that continuously increments a global variable



- We can build our own timer [Lip+16; Sch+17]
- Start a thread that continuously increments a global variable
- The global variable is our timestamp





www.tugraz.at

CPU cycles one increment takes



1 timestamp = rdtsc();






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





- $_1$  mov &timestamp ,  $\ensuremath{\% rcx}$
- 2 1: inc %rax
- з mov %rax, (%rcx)
- 4 **jmp** 1b



















• physical: different offsets on the chip



- physical: different offsets on the chip
- microarchitectural:



- physical: different offsets on the chip
- microarchitectural:
  - different microarchitectural elements



- physical: different offsets on the chip
- microarchitectural:
  - different microarchitectural elements
  - more significant: huge virtual adress space



- physical: different offsets on the chip
- microarchitectural:
  - different microarchitectural elements
  - more significant: huge virtual adress space
  - 2<sup>48</sup> different virtual memory locations



- physical: different offsets on the chip
- microarchitectural:
  - different microarchitectural elements
  - more significant: huge virtual adress space
  - 2<sup>48</sup> different virtual memory locations
  - the location is often (part of) the secret

E .	Terminal	-	• •	×	Open 🔻	+	Untitled	i Document 1	Save	= .	- 4	- ×
File Edit View Search Terminal Help												
% sleep 2; ./spy 300 7f0 8050	05140a4000-7f051417b000 /usr/lib/x86_64-linux	-7f051417b000 r-xp 0x20000 0B:0 /x86_64-linux-gnu/gedit/libged	02 26 it.so	5								
📩 () ican car				-								
E [prefetch]			0	×								
File Edit Mew Search Terminal Help shark% ./spy []												
phome/danievjs:							Plain Text 👻	Təb Width: 2 👻 🛛 L	n 1, Col 1			NS

# **Cache Template**<sup>4</sup>



<sup>4</sup>D. Gruss, R. Spreitzer, and S. Mangard. Cache Template Attacks: Automating Attacks on Inclusive Last-Level Caches. In: USENIX Security Symposium. 2015.





• Managed by operating system



- Managed by operating system
- Buffers pages in RAM for faster accesses



- Managed by operating system
- Buffers pages in RAM for faster accesses
- State of pages is tracked:



- Managed by operating system
- Buffers pages in RAM for faster accesses
- State of pages is tracked:
  - No write access  $\rightarrow$  clean  $\rightarrow$  no write back



- Managed by operating system
- Buffers pages in RAM for faster accesses
- State of pages is tracked:
  - No write access  $\rightarrow$  clean  $\rightarrow$  no write back
  - Write access  $\rightarrow$  dirty  $\rightarrow$  write back



Page Cache

- Managed by operating system
- Buffers pages in RAM for faster accesses
- State of pages is tracked:
  - No write access  $\rightarrow$  clean  $\rightarrow$  no write back
  - Write access  $\rightarrow$  dirty  $\rightarrow$  write back
- Implemented by all major operating systems



# Page Cache Attacks





RAM

eviction#1
eviction#2
eviction#3
eviction#4
eviction#5
foo.so#1
foo.so#2
foo.so#3
foo.so#4



Address space





Address space



റട





www.tugraz.at

# Page Cache Attacks





RAM





Address space





Address space



റട













www.tugraz.at















Address space





Address space



Ô٢

www.tugraz.at



Daniel Gruss — Graz University of Technology

Victim




Daniel Gruss — Graz University of Technology





• Takes virtual memory range, returns vector





- Takes virtual memory range, returns vector
- Indicates presence of queried pages in page cache





- Takes virtual memory range, returns vector
- Indicates presence of queried pages in page cache QueryWorkingSetEx (465.91 ns)





- Takes virtual memory range, returns vector
- Indicates presence of queried pages in page cache

QueryWorkingSetEx (465.91 ns)

• Takes process handle + virtual memory address, returns struct





- Takes virtual memory range, returns vector
- Indicates presence of queried pages in page cache

QueryWorkingSetEx (465.91 ns)

- Takes process handle + virtual memory address, returns struct
- Exposes attributes of queried page ...





- Takes virtual memory range, returns vector
- Indicates presence of queried pages in page cache

QueryWorkingSetEx (465.91 ns)

- Takes process handle + virtual memory address, returns struct
- Exposes attributes of queried page ...
- ... presence in working set

#### **Observe Page Cache State**



#### mincore $(2.04 \, \mu s)$

- Takes virtual memory range, returns vector
- Indicates presence of queried pages in page cache QueryWorkingSetEx (465.91 ns)
  - Takes process handle + virtual memory address, returns struct
  - Exposes attributes of queried page ...
  - ... presence in working set
  - ... number of working sets containing page (ShareCount)



• Necessary for detecting multiple accesses



- Necessary for detecting multiple accesses
- Bottleneck of side channel



- Necessary for detecting multiple accesses
- Bottleneck of side channel
- Linux: eviction (takes 149 ms)
- Windows: VirtualUnlock if possible (takes  $17.69 \, \mu s$ )
- Windows: SetProcessWorkingSetSize + eviction (takes 4.48 ms)









• We want the performance optimizations

Daniel Gruss — Graz University of Technology



- We want the performance optimizations
- Many side-channel attacks exploit intended behavior



- We want the performance optimizations
- Many side-channel attacks exploit intended behavior
- Often a trade-off between security and performance



- We want the performance optimizations
- Many side-channel attacks exploit intended behavior
- Often a trade-off between security and performance
- Every optimization is potentially a side channel



• We won't get rid of side channels



- We won't get rid of side channels
- More optimizations  $\rightarrow$  more side channels



- We won't get rid of side channels
- More optimizations  $\rightarrow$  more side channels
- But: low hanging fruits will disappear



# Software-based Microarchitectural Attacks

#### **Daniel Gruss**

July 8, 2018

Graz University of Technology

## References

D. Gruss, C. Maurice, and S. Mangard. Rowhammer.js: A Remote Software-Induced Fault Attack in JavaScript. In: DIMVA. 2016.
D. Gruss, R. Spreitzer, and S. Mangard. Cache Template Attacks: Automating Attacks on Inclusive Last-Level Caches. In: USENIX Security Symposium. 2015.
M. Lipp, D. Gruss, R. Spreitzer, C. Maurice, and S. Mangard. ARMageddon: Cache Attacks on Mobile Devices. In: USENIX Security Symposium. 2016.
Y. Oren, V. P. Kemerlis, S. Sethumadhavan, and A. D. Keromytis. The Spy in the Sandbox: Practical Cache Attacks in JavaScript and their Implications. In: CCS. 2015.

M. Schwarz, D. Gruss, S. Weiser, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.