

# Introduction to Microarchitectural Attacks

**Daniel Gruss** 

June 18, 2019

Graz University of Technology

- security and privacy rely on secrets (unknown to attackers)
- secrets can leak through side channels

- security and privacy rely on secrets (unknown to attackers)
- secrets can leak through side channels
- software-based  $\rightarrow$  no physical access









# 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





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









www.tugraz.at



www.tugraz.at





www.tugraz.at

























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

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

- Do you measure what you think you measure?
- Out-of-order execution  $\rightarrow$  what is really executed?

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



## **Memory Hierarchy**




www.tugraz.at



• L1 and L2 are private



- L1 and L2 are private
- last-level cache:



- L1 and L2 are private
- last-level cache:
  - divided in slices



- L1 and L2 are private
- last-level cache:
  - divided in slices
  - shared across cores



- L1 and L2 are private
- last-level cache:
  - divided in slices
  - shared across cores
  - inclusive



• inclusive LLC: superset of L1 and L2





• inclusive LLC: superset of L1 and L2





• inclusive LLC: superset of L1 and L2



- **inclusive** LLC: superset of L1 and L2
- data evicted from the LLC is also evicted from L1 and L2





- inclusive LLC: superset of L1 and L2
- data evicted from the LLC is also evicted from L1 and L2
- a core can evict lines in the private L1 of another core

• Locate key-dependent memory accesses

- Locate key-dependent memory accesses
- How?

• Preprocessing step to find exploitable addresses automatically

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)
  - called "Cache Template"

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)
  - called "Cache Template"

- Preprocessing step to find exploitable addresses automatically
  - w.r.t. "events" (keystrokes, encryptions, ...)
  - called "Cache Template"

Exploitation Phase

• Monitor exploitable addresses





#### Victim address space



Cache is empty



Attacker triggers an event



Attacker checks one address for cache hits ("Reload")



Update number of cache hits per event



Attacker flushes shared memory





#### Victim address space



## Repeat for higher accuracy



### Continue with next address

Daniel Gruss — Graz University of Technology

Victim address space



## Continue with next address

Daniel Gruss — Graz University of Technology

Victim address space

	Terminal		- 0	×	Cinen 🕳	<b>4</b>	Untitled	Document 1	Saue	± :	 
File Edit View Search Terminal Help											
% sleep 2; ./spy 300 7f05 8050 ∎	5140a4000-7f051417b000 r-xp 0 /usr/lib/x86_64-linux-gnu/ge	r-xp 0x20000 08: gnu/gedit/libged	02 2 11t.s	2 26 t.so	1						
							I				
- Director son		SUM- 00 01 6010 F	1-14-2								
12				×							
File Edit View Search Terminal Help shark% ./spy []											
phome/gamer.ja.							Plain Text 👻	Tab Width: 2 🐱	Ln 1, Col 1		INS





## Profiling Phase: 1 Event, 1 Address

ADDRESS



# Example: Cache Hit Ratio for (0x7c800, n): 200 / 200



## **Profiling Phase: All Events, 1 Address**



## Example: Cache Hit Ratio for (0x7c800, u): 13 / 200

Daniel Gruss — Graz University of Technology

15

## **Profiling Phase: All Events, 1 Address**



Distinguish n from other keys by monitoring 0x7c800

## Profiling Phase: All Events, All Addresses



# Directly mapped cache

Memory Address

# Directly mapped cache




#### Memory Address



Tag	Data

#### Memory Address



Cache

Tag	Data

#### www.tugraz.at

# Directly mapped cache



#### Daniel Gruss — Graz University of Technology

2<sup>b</sup> bytes









## Problem: working on congruent addresses











 $\rightarrow$  replacement policy

## Flush+Reload



## Flush+Reload



#### www.tugraz.at

## Flush+Reload



# Flush+Reload



#### www.tugraz.at

## Flush+Reload



Pros: fine granularity (1 line)

Pros: fine granularity (1 line)

Cons: restrictive

1. needs clflush instruction (not available e.g., in JS)

Pros: fine granularity (1 line)

Cons: restrictive

- 1. needs clflush instruction (not available e.g., in JS)
- 2. needs shared memory























1. no need for clflush instruction (not available e.g., in JS)

- 1. no need for clflush instruction (not available e.g., in JS)
- 2. no need for shared memory

- 1. no need for clflush instruction (not available e.g., in JS)
- 2. no need for shared memory

- 1. no need for clflush instruction (not available e.g., in JS)
- 2. no need for shared memory

Cons: coarser granularity (1 set)

• Paging: memory translated page-wise from virtual to physical
- Paging: memory translated page-wise from virtual to physical
- TLB (translation lookaside buffer) caches virtual to physical mapping

- Paging: memory translated page-wise from virtual to physical
- TLB (translation lookaside buffer) caches virtual to physical mapping
- TLB has some latency

- Paging: memory translated page-wise from virtual to physical
- TLB (translation lookaside buffer) caches virtual to physical mapping
- TLB has some latency
- Worst case for Cache: mapping not in TLB, need to load mapping from RAM

- Paging: memory translated page-wise from virtual to physical
- TLB (translation lookaside buffer) caches virtual to physical mapping
- TLB has some latency
- Worst case for Cache: mapping not in TLB, need to load mapping from RAM
- Solution: Use virtual addresses instead of physical addresses

• VIVT: Virtually indexed, virtually tagged

- VIVT: Virtually indexed, virtually tagged
- PIPT: Physically indexed, physically tagged

- VIVT: Virtually indexed, virtually tagged
- PIPT: Physically indexed, physically tagged
- PIVT: Physically indexed, virtually tagged

- VIVT: Virtually indexed, virtually tagged
- PIPT: Physically indexed, physically tagged
- PIVT: Physically indexed, virtually tagged
- VIPT: Virtually indexed, physically tagged







• Shared memory more than once in cache











• Shared memory more than once in cache







• Using more bits is unpractical (like VIVT)



- Using more bits is unpractical (like VIVT)
- $\rightarrow~{\sf Cache~size}\,\leq\,\#$  ways  $\cdot~{\sf page~size}$

• L1 caches: VIVT or VIPT

- L1 caches: VIVT or VIPT
- L2/L3 caches: PIPT

We need to evict cache lines without clflush or shared memory:

1. which addresses do we access to have congruent cache lines?

We need to evict cache lines without clflush or shared memory:

- 1. which addresses do we access to have congruent cache lines?
- 2. without any privilege?

We need to evict cache lines without clflush or shared memory:

- 1. which addresses do we access to have congruent cache lines?
- 2. without any privilege?
- 3. and in which order do we access them?



"LRU eviction":

• assume that cache uses LRU replacement



"LRU eviction":

- assume that cache uses LRU replacement
- accessing *n* addresses from the same cache set to evict an *n*-way set



"LRU eviction":

- assume that cache uses LRU replacement
- accessing *n* addresses from the same cache set to evict an *n*-way set
- eviction from last level  $\rightarrow$  from whole hierarchy (it's inclusive!)

# #1.2: Which addresses map to the same set?



# #1.2: Which addresses map to the same set?



# #1.2: Which addresses map to the same set?



• function H that maps slices is undocumented



- function H that maps slices is undocumented
- reverse-engineered by Maurice et al



- function H that maps slices is undocumented
- reverse-engineered by Maurice et al
- hash function basically an XOR of address bits

## 3 functions, depending on the number of cores

			Address bit																														
		3	3	3	3	3	3	3	3	2	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	0	0	0
		7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6
2 cores	<i>o</i> 0						$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$
4 cores	00					$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$
	$o_1$				$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$				$\oplus$									
	<i>o</i> 0		$\oplus$	$\oplus$		$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$
8 cores	$o_1$	$\oplus$		$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$				$\oplus$									
	02	$\oplus$	$\oplus$	$\oplus$	$\oplus$			$\oplus$			$\oplus$			$\oplus$	$\oplus$				$\oplus$														

• last-level cache is physically indexed
- last-level cache is physically indexed
- root privileges needed for physical addresses

- last-level cache is physically indexed
- root privileges needed for physical addresses
- use 2 MB pages  $\rightarrow$  lowest 21 bits are the same as virtual address

- last-level cache is physically indexed
- root privileges needed for physical addresses
- use 2 MB pages  $\rightarrow$  lowest 21 bits are the same as virtual address
- $\rightarrow\,$  enough to compute the cache set







www.tugraz.at



www.tugraz.at







64k cells 1 capacitor, 1 transitor each







CPU wants to access row 1



- $\ensuremath{\mathsf{CPU}}$  wants to access row 1
- ightarrow row 1 activated









DRAM bank

- $\ensuremath{\mathsf{CPU}}$  wants to access row 1
- ightarrow row 1 activated
- ightarrow row 1 copied to row buffer





CPU wants to access row 2



- CPU wants to access row 2
- $\rightarrow$  row 2 activated









DRAM bank

#### CPU wants to access row 2

 $\rightarrow$  row 2 activated

ightarrow row 2 copied to row buffer







- CPU wants to access row 2  $\rightarrow$  row 2 activated  $\rightarrow$  row 2 copied to row buffer
- $\rightarrow$  slow (row conflict)





CPU wants to access row 2-again







CPU wants to access row 2-again

ightarrow row 2 already in row buffer



### DRAM bank

#### CPU wants to access row 2—again

ightarrow row 2 already in row buffer







- CPU wants to access row 2—again
- $\rightarrow$  row 2 already in row buffer
- $\rightarrow$  fast (row hit)





## row buffer = cache

# **Timing difference**





• Cache set is determined by part of physical address



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



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



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



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





# **Physical Addresses**















**Physical Addresses** 



#### Daniel Gruss — Graz University of Technology

www.tugraz.at

40







www.tugraz.at




















www.tugraz.at

Daniel Gruss — Graz University of Technology















### Result on an Intel i5-6200U



Daniel Gruss — Graz University of Technology

www.tugraz.at

42



#### Daniel Gruss — Graz University of Technology



• LRU replacement policy: oldest entry first



- LRU replacement policy: oldest entry first
- timestamps for every cache line



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



- LRU replacement policy: oldest entry first
- timestamps for every cache line
- access updates timestamp



• no LRU replacement

Daniel Gruss — Graz University of Technology











• no LRU replacement

Daniel Gruss — Graz University of Technology




• no LRU replacement



• no LRU replacement



- no LRU replacement
- only 75% success rate on Haswell



- no LRU replacement
- only 75% success rate on Haswell
- $\bullet\,$  more accesses  $\rightarrow\,$  higher success rate, but too slow



 $\rightarrow$  fast and effective on Haswell: eviction rate  ${>}99.97\%$ 

• represent accesses as a sequence of numbers: 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4, ...

- represent accesses as a sequence of numbers: 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4, ...
- can be a long sequence

- represent accesses as a sequence of numbers: 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4, ...
- can be a long sequence
- all congruent addresses are indistinguishable w.r.t eviction strategy

- represent accesses as a sequence of numbers: 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4, ...
- can be a long sequence
- all congruent addresses are indistinguishable w.r.t eviction strategy
- $\rightarrow\,$  adding more unique addresses can increase eviction rate

- represent accesses as a sequence of numbers: 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4, ...
- can be a long sequence
- all congruent addresses are indistinguishable w.r.t eviction strategy
- $\rightarrow\,$  adding more unique addresses can increase eviction rate
- $\rightarrow~$  multiple accesses to one address can increase the eviction rate

- represent accesses as a sequence of numbers: 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4, ...
- can be a long sequence
- all congruent addresses are indistinguishable w.r.t eviction strategy
- $\rightarrow\,$  adding more unique addresses can increase eviction rate
- $\rightarrow~$  multiple accesses to one address can increase the eviction rate
  - $\bullet$  indistinguishable  $\rightarrow$  balanced number of accesses

S: total number of different addresses







• P-2-2-1-4  $\rightarrow$  1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4

• 
$$P-2-2-1-4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4$$

• 
$$P - 2 - 2 - 1 - 4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4'$$

Daniel Gruss — Graz University of Technology

48

• 
$$P - 2 - 2 - 1 - 4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4$$
  
 $D = 2$ 

• 
$$P-2-2-1-4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4$$
  
 $D=2$ 
 $C=2$ 

• 
$$P-2-2-1-4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4$$
  
 $L=1$   
 $D=2$   
 $C=2$ 

• 
$$P-2-2-1-4 \rightarrow (1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4)$$
  
 $L=1$   
 $D=2$   
 $C=2$ 

• P-1-1-1-4 ightarrow 1, 2, 3, 4 ightarrow LRU eviction with set size 4

strategy	# accesses	eviction rate	loop time
P-1-1-1-17	17		
P-1-1-20	20		

<sup>&</sup>lt;sup>1</sup>Executed in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-17	17	74.46% 🗡	
P-1-1-20	20	99.82% 🗸	

 $<sup>^1\</sup>mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡

<sup>&</sup>lt;sup>1</sup>Executed in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34		

 $<sup>^1\</sup>mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34	99.86% 🗸	

 $<sup>^1\</sup>mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34	99.86% 🗸	191 ns 🗸

 $<sup>^1\</sup>mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34	99.86% 🗸	191 ns 🗸
P-2-2-1-17	64		

 $^1\mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34	99.86% 🗸	191 ns 🗸
P-2-2-1-17	64	99.98% 🗸	

 $^1\mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34	99.86% 🗸	191 ns 🗸
P-2-2-1-17	64	99.98% 🗸	180 ns 🗸

 $^1\mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

strategy	# accesses	eviction rate	loop time
P-1-1-17	17	74.46% 🗡	307 ns 🗸
P-1-1-20	20	99.82% 🗸	934 ns 🗡
P-2-1-1-17	34	99.86% 🗸	191 ns 🗸
P-2-2-1-17	64	99.98% 🗸	180 ns 🗸

 $\rightarrow$  more accesses, smaller execution time?

 $<sup>^1\</sup>mathsf{Executed}$  in a loop, on a Haswell with a 16-way last-level cache

P-1-1-1-17 (17 accesses, 307ns)

P-2-1-1-17 (34 accesses, 191ns)

Time in ns

### P-1-1-17 (17 accesses, 307ns)



#### P-2-1-1-17 (34 accesses, 191ns)



Time in ns
Miss	Miss
(intended)	(intended)

#### P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н
--------------------	--------------------	---

## P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

# P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

# P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

# P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

# P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

# P-2-1-1-17 (34 accesses, 191ns)



Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	н	н	н	н	н	н	н
--------------------	--------------------	---	---	---	---	---	---	---

Time in ns

Miss (intended)	Miss (intended)	н	Miss
--------------------	--------------------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	ļ		1		н	н	н	н	н	
--------------------	--------------------	---	--	---	--	---	---	---	---	---	--

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss
--------------------	--------------------	---	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)		н	н	н	н	H	н	ŀ	н	Miss
------------------------------------	--	---	---	---	---	---	---	---	---	------

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss
--------------------	--------------------	---	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	н	н	н	н	н	н	н	н	Miss	н	
--------------------	--------------------	---	---	---	---	---	---	---	---	------	---	--

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss
--------------------	--------------------	---	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	н	н	41	E H	н	н	н	Miss	н	н	
------------------------------------	---	---	----	-----	---	---	---	------	---	---	--

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss
--------------------	--------------------	---	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	HHHHHHH Miss	нн	4
------------------------------------	--------------	----	---

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	ННИНИНИИ Miss	н	н	н	н	
------------------------------------	---------------	---	---	---	---	--

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss
------------------------------------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	ннннннн Маа	нннн
------------------------------------	-------------	------

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss
------------------------------------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	ННННННН Miss	нынын
------------------------------------	--------------	-------

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss
------------------------------------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	н	н	н	н	н	н	H	н		Miss	H	-	01	H	•	ŀ		
------------------------------------	---	---	---	---	---	---	---	---	--	------	---	---	----	---	---	---	--	--

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss
------------------------------------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)		н	ŀ	-	,	4	•		H		н		-	н		Miss	н	-			н	н	ŀ	4	I	
------------------------------------	--	---	---	---	---	---	---	--	---	--	---	--	---	---	--	------	---	---	--	--	---	---	---	---	---	--

Time in ns

Miss Miss (intended) (intende	) H	Miss	Miss	Miss	
----------------------------------	-----	------	------	------	--

#### P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	HHHHHHHH Miss	а НИИННИИ Miss
------------------------------------	---------------	----------------

Time in ns

Miss (intended)	Miss (intended)	l Miss	Miss	Miss	н	Miss
--------------------	--------------------	--------	------	------	---	------

#### P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	ннинини	Miss	нннннн	f Miss
--------------------	--------------------	---------	------	--------	--------

Time in ns

Miss (intended)	Miss (intended)	l Miss	Miss	Miss	н	Miss
--------------------	--------------------	--------	------	------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intender	нннннн	Miss H H	ннннн	Miss
-----------------------------------	--------	----------	-------	------

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss	н	Miss
------------------------------------	---	------	------	------	---	------

#### P-2-1-1-17 (34 accesses, 191ns)

Miss M (intended) (inter	iss nded) нинининин	Miss H	нинини	Miss	н
-----------------------------	------------------------	--------	--------	------	---

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss	н	Miss
------------------------------------	---	------	------	------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended) (i	Miss intended) HHHHHHHHH	Miss H H H	ннынн	Miss H F	нн
-----------------------	-----------------------------	------------	-------	----------	----

Time in ns

Miss (intended)	Miss (intended)	l Miss	Miss	Miss	н	Miss
--------------------	--------------------	--------	------	------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Time in ns

Miss (intended)	Miss (intended)	l Miss	Miss	Miss	н	Miss
--------------------	--------------------	--------	------	------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss	н	Miss
------------------------------------	---	------	------	------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	нинини	Miss	нынынын Miss	ныныны
------------------------------------	--------	------	--------------	--------

Time in ns

Miss Miss (intended) (intended)	н	Miss	Miss	Miss	н	Miss	Miss
------------------------------------	---	------	------	------	---	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	ННННННН Miss	нинини	Miss ННННННН
------------------------------------	--------------	--------	--------------

Time in ns

Miss Miss (intended) (intended)	ŀ	Miss	Miss	Miss	н	Miss	Miss
------------------------------------	---	------	------	------	---	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	ННННННН Miss	нининии	Miss ННННННН
------------------------------------	--------------	---------	--------------

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------

# P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нининии	Miss	ымынымы	Miss ННННН	H Miss
--------------------	--------------------	---------	------	---------	------------	--------

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended) (	Miss (intended)	н Місс НІННИНИ	Miss HHHHHHH Miss I
----------------------	--------------------	----------------	---------------------

Time in ns

Daniel Gruss — Graz University of Technology

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нннннн Miss	инынныны м	fiss нининини	Miss HH
--------------------	--------------------	-------------	------------	---------------	---------

Time in ns

Daniel Gruss — Graz University of Technology

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss Miss (intended) (intended)	ННННННН Miss	нныннын Міза	ННННННН Miss P	нн
------------------------------------	--------------	--------------	----------------	----

Time in ns

Daniel Gruss — Graz University of Technology

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	ннинини	Miss	ннинини	Miss HHHHHHHH	Miss HHHH
--------------------	--------------------	---------	------	---------	---------------	-----------

Time in ns

Daniel Gruss — Graz University of Technology

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нннннн	Miss HHHHHHHH	Miss ИННИНИИ	Miss HHHHH
--------------------	--------------------	--------	---------------	--------------	------------

Time in ns

Daniel Gruss — Graz University of Technology

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss	н
--------------------	--------------------	---	------	------	------	---	------	------	------	---

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	ныны	Miss ННИНИНИ	Miss HHHHHHHH	Miss HHHHH
--------------------	--------------------	------	--------------	---------------	------------

Time in ns

Daniel Gruss — Graz University of Technology

Miss (intended)	Miss (intended)	H Miss	Miss	Miss	H Miss	Miss	Miss	н	Miss
--------------------	--------------------	--------	------	------	--------	------	------	---	------

## P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нннннн	Miss HHHHHHHH	Miss ИННИНИИ	Miss HHHHH
--------------------	--------------------	--------	---------------	--------------	------------

Time in ns
Miss (intended)	Miss (intended)	H Miss	Miss	Miss	н	Miss	Miss	Miss	н	Miss	Miss	
--------------------	--------------------	--------	------	------	---	------	------	------	---	------	------	--

#### P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нныныны	Miss HHHHH	(HH Miss	нинини	Miss HHHHH
--------------------	--------------------	---------	------------	----------	--------	------------

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss	н	Miss	Miss	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------	---	------	------	------

#### P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нныныны	Miss HHHHH	(HH Miss	нинини	Miss HHHHH
--------------------	--------------------	---------	------------	----------	--------	------------

Time in ns

Miss (intended)	Miss (intended)	H Miss	Miss	Miss	H Miss	Miss	Miss	H Miss	Miss	Miss	H
--------------------	--------------------	--------	------	------	--------	------	------	--------	------	------	---

#### P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нныныны	Miss HHHHH	(HH Miss	нинини	Miss HHHHH
--------------------	--------------------	---------	------------	----------	--------	------------

Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss	Miss	н	Miss	Miss	Miss	н	Miss	Miss	Miss	н	Miss
--------------------	--------------------	---	------	------	------	---	------	------	------	---	------	------	------	---	------

#### P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нныныны	Miss HHHHH	(HH Miss	нинини	Miss HHHHH
--------------------	--------------------	---------	------------	----------	--------	------------

Time in ns

Miss (intended)	Miss (intended)	H Miss	Miss	Miss	H Miss	Miss	Miss	H Miss	Miss	Miss	H Miss	Miss
--------------------	--------------------	--------	------	------	--------	------	------	--------	------	------	--------	------

#### P-2-1-1-17 (34 accesses, 191ns)

Miss (intended)	Miss (intended)	нннннн	Miss HHHHHHHH	Miss ИННИНИИ	Miss HHHHH
--------------------	--------------------	--------	---------------	--------------	------------

Time in ns

# HELLO FROM THE OTHER SIDE (DEMO): VIDEO STREAMING OVER CACHE COVERT CHANNEL





DRAM bank

- Cells leak  $\rightarrow$  repetitive refresh necessary
- Maximum interval between refreshes to guarantee data integrity
- Cells leak faster upon proximate accesses  $\rightarrow$ Rowhammer



- Cells leak → repetitive refresh necessary
- Maximum interval between refreshes to guarantee data integrity
- Cells leak faster upon proximate accesses  $\rightarrow$  Rowhammer



- Cells leak → repetitive refresh necessary
  - Maximum interval between refreshes to guarantee data integrity
  - Cells leak faster upon proximate accesses  $\rightarrow$  Rowhammer



- Cells leak → repetitive refresh necessary
- Maximum interval between refreshes to guarantee data integrity
- Cells leak faster upon proximate accesses  $\rightarrow$  Rowhammer



- Cells leak → repetitive refresh necessary
  - Maximum interval between refreshes to guarantee data integrity
  - Cells leak faster upon proximate accesses  $\rightarrow$  Rowhammer



- Cells leak → repetitive refresh necessary
- Maximum interval between refreshes to guarantee data integrity
- Cells leak faster upon proximate accesses  $\rightarrow$  Rowhammer

• There are two different hammering techniques

- There are two different hammering techniques
- #1: Hammer one row next to victim row and other random rows

- There are two different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row

- There are three different hammering techniques
- #1: Hammer one row next to victim row and other random rows
- #2: Hammer two rows neighboring victim row
- #3: Hammer only one row next to victim row



## #1 - Single-sided hammering









## #1 - Single-sided hammering





### #2 - Double-sided hammering





### #2 - Double-sided hammering





### #2 - Double-sided hammering





## DRAM bank \_

#### Daniel Gruss — Graz University of Technology



## DRAM bank \_

#### Daniel Gruss — Graz University of Technology










• They are not random  $\rightarrow$  highly reproducible flip pattern!



- They are not random  $\rightarrow$  highly reproducible flip pattern!
  - 1. Choose a data structure that you can place at arbitrary memory locations



- They are not random  $\rightarrow$  highly reproducible flip pattern!
  - 1. Choose a data structure that you can place at arbitrary memory locations
  - 2. Scan for "good" flips



- They are not random  $\rightarrow$  highly reproducible flip pattern!
  - 1. Choose a data structure that you can place at arbitrary memory locations
  - 2. Scan for "good" flips
  - 3. Place data structure there



- They are not random  $\rightarrow$  highly reproducible flip pattern!
  - 1. Choose a data structure that you can place at arbitrary memory locations
  - 2. Scan for "good" flips
  - 3. Place data structure there
  - 4. Trigger bit flip again







• Many applications perform actions as root



• Many applications perform actions as root



- Many applications perform actions as root
- They can be used by unprivileged users as well



- Many applications perform actions as root
- They can be used by unprivileged users as well



- Many applications perform actions as root
- They can be used by unprivileged users as well
- sudo

























• 85% affected [Kim+14] (see Figure)





- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]





- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



• First believed to be safe



- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



- First believed to be safe
- We showed bit flips [Pes+16]



- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]



- First believed to be safe
- We showed bit flips [Pes+16]
- 67% affected [Lan16]



- 85% affected [Kim+14] (see Figure)
- 52% affected [SD15]

# 

- First believed to be safe
- We showed bit flips [Pes+16]
- 67% affected [Lan16]





	1111111	1111111
	1111111	1010011
	0000000	0000000
activate >	1111111	1000000
	1111111	1111111
	1111111	1111111
	1111111	1111111
	1111111	1111111
		-

#### DRAM bank

		_	-
	1111111	1111111	
activate >	1111111	1010011	
	0000000	0000000	
	1111111	1000000	
	1111111	1111111	
	1111111	1111111	
	1111111	1111111	
	1111111	1111111	
			-

#### DRAM bank

	1111111	1111111
	1111111	1010011
	0000000	0000000
activate >	1111111	1000000
	1111111	1111111
	1111111	1111111
	1111111	1111111
	1111111	1111111
		-

#### DRAM bank

		_	-
	1111111	1111111	
activate >	1111111	1010011	
	0000000	0000000	
	1111111	1000000	
	1111111	1111111	
	1111111	1111111	
	1111111	1111111	
	1111111	1111111	
			-

#### DRAM bank

	1111111	1111111
	1111111	1010011
	0000000	0000000
activate >	1111111	1000000
	1111111	1111111
	1111111	1111111
	1111111	1111111
	1111111	1111111
		-

#### DRAM bank
RAMbleed





• We want the performance optimizations



- 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
- $\bullet\,$  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



# Introduction to Microarchitectural Attacks

# **Daniel Gruss**

June 18, 2019

Graz University of Technology