

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 → 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~~

\$1,299

ORDER VIA PHONE: +555 12345





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




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

Cache miss




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

Cache miss




Request




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

Cache miss




Request

Response




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


Cache miss





Request

Response










Shared Memory


ATTACKER


flush
access





VICTIM


access

















FAIL



- Very short timings
- rdtsc instruction: “cycle-accurate” timestamps

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

- Do you measure what you *think* you measure?
- *Out-of-order* execution → what is really executed?

rdtsc

function()

[...]

rdtsc

rdtsc

[...]

rdtsc

function()

rdtsc

rdtsc

function()

[...]

- use pseudo-serializing instruction `rdtscp` (recent CPUs)

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

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


- use pseudo-serializing instruction `rdtscp` (recent CPUs)
- and/or use serializing instructions like `cuid`
- 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](#).














- 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



- **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?

Profiling Phase

- Preprocessing step to find exploitable addresses automatically

Profiling Phase

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

Profiling Phase

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

Profiling Phase

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

Profiling Phase

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

Exploitation Phase

- Monitor exploitable addresses


Attacker address space




Cache




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


Attacker address space



Cache



Victim address space



Repeat for higher accuracy


Attacker address space



Cache



Victim address space



Continue with next address


Attacker address space



Cache



Victim address space



Continue with next address

```
Terminal
File Edit View Search Terminal Help
% sleep 2; ./spy 300 7f05140a4000-7f051417b000 r-xp 0x20000 08:02 26
8050 /usr/lib/x86_64-linux-gnu/gedit/libgedit.so

Inkscape 0.48.5 (2017-03-22)
Inkscape 0.48.5 (2017-03-22)
Terminal
```

```
Terminal
File Edit View Search Terminal Help
shark% ./spy
```

```
Untitled Document 1
Open + Save - + x
1
I
Plain Text Tab Width: 2 Ln 1, Col 1 INS
```

ADDRESS

0x7c800

KEY

n



The diagram consists of the word 'ADDRESS' written vertically on the left. To its right is the hexadecimal value '0x7c800'. Further to the right is the word 'KEY' written vertically. Below 'KEY' is a small black square, and to its left is a small letter 'n'.



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






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



Distinguish `n` from other keys by monitoring `0x7c800`



Memory Address

Memory Address

--

Cache

Memory Address

--

Cache



Tag	Data


Memory Address


--	--


Cache


Tag	Data
















Problem: working on congruent addresses
















→ replacement policy











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


Attacker
address space





Cache





Victim
address space





























Pros: less restrictive

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

Pros: less restrictive

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

Pros: less restrictive

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

Pros: less restrictive

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







- Fast
- Virtual Tag is not unique (Context switches)





- Fast
- Virtual Tag is not unique (Context switches)
- Shared memory more than once in cache






- Slow (TLB lookup for index)
- Shared memory only once in cache!







- Slow (TLB lookup for index)
- Virtual Tag is not unique (Context switches)




- Slow (TLB lookup for index)
- Virtual Tag is not unique (Context switches)
- Shared memory more than once in cache




- Fast



- Fast
- 4 KiB pages: last 12 bits of VA and PA are equal



- Fast
- 4 KiB pages: last 12 bits of VA and PA are equal
- Using more bits is unpractical (like VIVT)



- Fast
- 4 KiB pages: last 12 bits of VA and PA are equal
- Using more bits is unpractical (like VIVT)

→ Cache size $\leq \# \text{ ways} \cdot \text{page size}$

- 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?




#1.2: Which addresses map to the same set?




- function H that maps slices is undocumented

#1.2: Which addresses map to the same set?



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

#1.2: Which addresses map to the same set?



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

#1.2: Which addresses map to the same set?

3 functions, depending on the number of cores

		Address bit																																
		3 7	3 6	3 5	3 4	3 3	3 2	3 1	3 0	2 9	2 8	2 7	2 6	2 5	2 4	2 3	2 2	2 1	2 0	1 9	1 8	1 7	1 6	1 5	1 4	1 3	1 2	1 1	1 0	0 9	0 8	0 7	0 6	
2 cores	o_0					$\oplus$			$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$	
4 cores	o_0					$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$
	o_1				$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$	
8 cores	o_0		$\oplus$	$\oplus$		$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$
	o_1	$\oplus$		$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$		$\oplus$		$\oplus$	$\oplus$	$\oplus$	$\oplus$	$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$	
	o_2	$\oplus$	$\oplus$	$\oplus$	$\oplus$			$\oplus$	$\oplus$			$\oplus$	$\oplus$			$\oplus$	$\oplus$		$\oplus$			$\oplus$		$\oplus$		$\oplus$		$\oplus$				$\oplus$		


- 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 → 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 → lowest 21 bits are the same as virtual address
- enough to compute the cache set












chip




chip





64k cells
1 capacitor,
1 transistor each



CPU wants to access row 1









CPU wants to access row 1

→ row 1 activated

→ row 1 copied to row buffer





CPU wants to access row 2



CPU wants to access row 2

→ row 2 activated






CPU wants to access row 2

→ row 2 activated

→ row 2 copied to row buffer




CPU wants to access row 2


→ row 2 activated

→ row 2 copied to row buffer

→ **slow** (row conflict)




CPU wants to access row 2—again




CPU wants to access row 2—again

→ row 2 already in row buffer




CPU wants to access row 2—again
→ row 2 already in row buffer




CPU wants to access row 2—again


→ row 2 already in row buffer

→ **fast** (row hit)




row buffer = cache







- 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



8 kB row x in BG0 (1) and channel (1)

	Page #2	Page #3	Page #4	Page #5	Page #6	Page #7	Page #8
--	---------	---------	---------	---------	---------	---------	---------

8 kB row x in BG0 (0) and channel (1)


	Page #2	Page #3	Page #4	Page #5	Page #6	Page #7	Page #8
--	---------	---------	---------	---------	---------	---------	---------

8 kB row x in BG0 (1) and channel (0)

	Page #2	Page #3	Page #4	Page #5	Page #6	Page #7	Page #8
--	---------	---------	---------	---------	---------	---------	---------


8 kB row x in BG0 (0) and channel (0)


	Page #2	Page #3	Page #4	Page #5	Page #6	Page #7	Page #8
--	---------	---------	---------	---------	---------	---------	---------
























































Result on an Intel i5-6200U



“LRU eviction” memory accesses

cache set



“LRU eviction” memory accesses

cache set



- LRU replacement policy: oldest entry first


“LRU eviction” memory accesses

cache set

2	5	8	1	7	6	3	4
---	---	---	---	---	---	---	---


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

“LRU eviction” memory accesses



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

“LRU eviction” memory accesses




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

“LRU eviction” memory accesses



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

“LRU eviction” memory accesses




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

“LRU eviction” memory accesses




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

“LRU eviction” memory accesses




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

“LRU eviction” memory accesses



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

“LRU eviction” memory accesses



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


“LRU eviction” memory accesses

cache set

2	5	8	1	7	6	3	4
---	---	---	---	---	---	---	---


- no LRU replacement

“LRU eviction” memory accesses




- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses



- no LRU replacement

“LRU eviction” memory accesses

cache set

12	5	8	11	7	6	14	16
----	---	---	----	---	---	----	----


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

“LRU eviction” memory accesses

cache set

12	5	8	11	7	6	14	16
----	---	---	----	---	---	----	----

- no LRU replacement
- only 75% success rate on Haswell
- more accesses → higher success rate, but **too slow**



→ 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
- 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
- adding more **unique** addresses can increase eviction rate
- **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
- adding more **unique** addresses can increase eviction rate
- **multiple** accesses to one address can increase the eviction rate
- indistinguishable → **balanced** number of accesses

Write eviction strategies as: *P-C-D-L-S*

```
for (s = 0; s <= S - D; s += L)
  for (c = 0; c <= C; c += 1)
    for (d = 0; d <= D; d += 1)
      *a[s+d];
```

Write eviction strategies as: *P-C-D-L-S*

S: total number of different addresses
(= set size)



```
for (s = 0; s <= S - D; s += L)
  for (c = 0; c <= C; c += 1)
    for (d = 0; d <= D; d += 1)
      *a[s+d];
```

Write eviction strategies as: $P-C-D-L-S$

S : total number of different addresses
(= set size)

D : different addresses per inner access
loop



```
for (s = 0; s <= S - D; s += L)
  for (c = 0; c <= C; c += 1)
    for (d = 0; d <= D; d += 1)
      *a[s+d];
```

Write eviction strategies as: *P-C-D-L-S*

S: total number of different addresses
(= set size)

D: different addresses per inner access

```
for (s = 0; s <= S - D; s += L)
  for (c = 0; c <= C; c += 1)
    for (d = 0; d <= D; d += 1)
      *a[s+d];
```

L: step size of the inner access
loop

Write eviction strategies as: *P-C-D-L-S*

S: total number of different addresses
(= set size)

D: different addresses per inner access
loop

```
for (s = 0; s <= S - D; s += L)
  for (c = 0; c <= C; c += 1)
    for (d = 0; d <= D; d += 1)
      *a[s+d];
```

L: step size of the inner access
loop

C: number of repetitions of the inner
access loop

```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```

```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```

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


```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```

- $P-2-2-1-4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4$  $S = 4$


```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```

- $P-2-2-1-4 \rightarrow (1, 2), (1, 2), (2, 3), (2, 3), (3, 4), (3, 4)$ $\xleftarrow{S=4}$


```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```



```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```



```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```



```
for (s = 0; s <= S - D; s += L)
  for (c = 1; c <= C; c += 1)
    for (d = 1; d <= D; d += 1)
      *a[s+d];
```

- $P\text{-}2\text{-}2\text{-}1\text{-}4 \rightarrow 1, 2, 1, 2, 2, 3, 2, 3, 3, 4, 3, 4$

$L = 1$ $D = 2$ $C = 2$ $S = 4$
- $P\text{-}1\text{-}1\text{-}1\text{-}4 \rightarrow 1, 2, 3, 4 \rightarrow$ LRU eviction with set size 4

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17		
<i>P</i> -1-1-1-20	20		

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	
<i>P</i> -1-1-1-20	20	99.82% ✓	

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34		

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34	99.86% ✓	

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34	99.86% ✓	191 ns ✓

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34	99.86% ✓	191 ns ✓
<i>P</i> -2-2-1-17	64		

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34	99.86% ✓	191 ns ✓
<i>P</i> -2-2-1-17	64	99.98% ✓	

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34	99.86% ✓	191 ns ✓
<i>P</i> -2-2-1-17	64	99.98% ✓	180 ns ✓

¹Executed in a loop, on a Haswell with a 16-way last-level cache

We evaluated more than 10000 strategies...¹

strategy	# accesses	eviction rate	loop time
<i>P</i> -1-1-1-17	17	74.46% ✗	307 ns ✓
<i>P</i> -1-1-1-20	20	99.82% ✓	934 ns ✗
<i>P</i> -2-1-1-17	34	99.86% ✓	191 ns ✓
<i>P</i> -2-2-1-17	64	99.98% ✓	180 ns ✓

→ more accesses, smaller execution time?

¹Executed 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



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




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




Time in ns

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




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




Time in ns

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



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



Time in ns

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



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



Time in ns

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



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



Time in ns

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



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



Time in ns

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



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



Time in ns

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



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



Time in ns

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



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



Time in ns

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




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




Time in ns

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




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



Time in ns

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




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




Time in ns

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




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




Time in ns

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



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




Time in ns

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



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



Time in ns

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



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




Time in ns

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



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



Time in ns

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


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

Time in ns

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



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




Time in ns

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




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




Time in ns

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




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



Time in ns

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




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




Time in ns

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




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



Time in ns

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




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




Time in ns

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




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




Time in ns

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




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




Time in ns

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



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



Time in ns

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



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




Time in ns

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

Miss (intended)	Miss (intended)	[H][H][H][H][H]	Miss	[H][H][H][H][H]	Miss	[H][H][H][H][H]
---	---	---	--	---	--	---

Time in ns

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




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




Time in ns

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




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




Time in ns

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




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




Time in ns

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




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




Time in ns

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




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



Time in ns

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




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



Time in ns

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



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



Time in ns

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




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



Time in ns

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



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



Time in ns

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



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



Time in ns

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




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



Time in ns

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



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



Time in ns

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



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



Time in ns



```
File Edit View Bookmarks Settings Help
root@ip-172-31-31-32 ~ %
```


```
File Edit View Bookmarks Settings Help
root@ip-172-31-31-32 ~ %
```

```
File Edit View Bookmarks Settings Help
-- alert 0.4.5 on ip-172-31-31-32 --
[... ASCII art ...]
Active Interface: eth0
Interface Speed: unknown
Current RX Speed: 0.01 KB/s
Graph Top RX Speed: 0.98 KB/s
Overall Top RX Speed: 0.98 KB/s
Received Packets: 64
Bytes Received: 0.045 MB
Errors on Receiving: 0
Current TX Speed: 0.33 KB/s
Graph Top TX Speed: 2.64 KB/s
Overall Top TX Speed: 2.64 KB/s
Transmitted Packets: 61
Bytes Transmitted: 0.039 MB
Errors on Transmission: 0
```


```
File Edit View Bookmarks Settings Help
root@ip-172-31-31-32 ~ %
```

```
File Edit View Bookmarks Settings Help
-- alert 0.4.5 on ip-172-31-31-32 --
[... ASCII art ...]
Active Interface: eth0
Interface Speed: unknown
Current RX Speed: 0.05 KB/s
Graph Top RX Speed: 0.36 KB/s
Overall Top RX Speed: 0.36 KB/s
Received Packets: 36
Bytes Received: 0.003 MB
Errors on Receiving: 0
Current TX Speed: 0.29 KB/s
Graph Top TX Speed: 0.84 KB/s
Overall Top TX Speed: 0.84 KB/s
Transmitted Packets: 26
Bytes Transmitted: 0.010 MB
Errors on Transmission: 0
```


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




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




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




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



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



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




- Cells leak → repetitive **refresh** necessary
- Maximum interval between refreshes to guarantee **data integrity**
- Cells leak faster upon proximate accesses → 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
























































- They are not random → highly reproducible flip pattern!



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




- They are not random → 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 → 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 → 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]

DDR3

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

DDR4


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





A still from the movie Toy Story showing Woody and Buzz Lightyear. Woody is on the left, looking concerned. Buzz is on the right, wearing his green space suit and holding a purple lollipop in his raised right hand. The background is a simple room with a door and some toys on the floor.


BIT FLIPS


BIT FLIPS EVERYWHERE
















- 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
- More optimizations → more side channels



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

Introduction to Microarchitectural Attacks

Daniel Gruss

June 18, 2019

Graz University of Technology