

Hardware-Software Co-Design against Microarchitectural Attacks

Daniel Gruss

September 4, 2019

Graz University of Technology



National Geographic

side channel = obtaining meta-data and deriving secrets from it

CHANGE MY MIND





• Profiling cache utilization with performance counters?



- Profiling cache utilization with performance counters? \rightarrow No





- Profiling cache utilization with performance counters? \rightarrow No
- Observing cache utilization with performance counters and using it to infer a crypto key?



- Profiling cache utilization with performance counters? \rightarrow No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes



- Profiling cache utilization with performance counters? \rightarrow No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload?



- Profiling cache utilization with performance counters? \rightarrow No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload? \rightarrow No

6	

- Profiling cache utilization with performance counters? \rightarrow No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload? \rightarrow No
- Measuring memory access latency with Flush+Reload and using it to infer keystroke timings?

6	

- Profiling cache utilization with performance counters? \rightarrow No
- Observing cache utilization with performance counters and using it to infer a crypto key? → Yes
- Measuring memory access latency with Flush+Reload? \rightarrow No
- Measuring memory access latency with Flush+Reload and using it to infer keystroke timings? \rightarrow Yes



• traditional cache attacks (crypto, keys, etc)



- traditional cache attacks (crypto, keys, etc)
- actual misspeculation (e.g., branch misprediction)

www.tugraz.at



- traditional cache attacks (crypto, keys, etc)
- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc

www.tugraz.at



- traditional cache attacks (crypto, keys, etc)
- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc

www.tugraz.at



- traditional cache attacks (crypto, keys, etc)
- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc
- Let's avoid the term Speculative Side-Channel Attacks



- traditional cache attacks (crypto, keys, etc)
- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc
- Let's avoid the term Speculative Side-Channel Attacks
- Let's be more precise



- traditional cache attacks (crypto, keys, etc)
- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc
- Let's avoid the term Speculative Side-Channel Attacks
- Let's be more precise
- $\bullet \ \rightarrow$ then we can think about actual mitigations









But what about these two?









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

Cache Hits



www.tugraz.at

Cache Hits Cache Misses



12	Terminal			× Ciner	r⊈ ∓	Untitled	Document 1	Saue =	: _	+ ×
File Edit View Search Terminal Help								2335 3		
% sleep 2; ./spy 300 7f0 8050	f05140a4000-7f051417b000 r /usr/lib/x86_64-linux-g	r-xp 0x20000 08:02 26 -gnu/gedit/libgedit.so	02 20 111.50							
						I				
in econos Intelatebi 3.		CORS 14 03 2017 20 CORS 14 03 20 CORS 14 00 CORS 14 00	-44-26							
File Edit View Search Terminal Help shark% ./spy []										
onumeroaments;						PlainText 👻	Tab Widds: 2 👻	Ln 1, Col 1		INS

Cache Template Attack Demo









Data



 \rightarrow replacement policy















www.tugraz.at 📕



www.tugraz.at 📕

























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 (\rightarrow cross-VM)

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

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

Cons: coarser granularity (1 set)



www.tugraz.at



• 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

www.tugraz.at





• 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





S: total number of different addresses











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

¹Executed 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% 🗸	

¹Executed 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 🗡

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

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

¹Executed 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% 🗸	

 $^{^1\}mathsf{Executed}$ 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 🗸

 $^{^1\}mathsf{Executed}$ 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		

 $^{^1\}mathsf{Executed}$ 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



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)

Miss (intended)

Miss

Time in ns

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



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



Time in ns


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



Time in ns



Time in ns



Time in ns



Time in ns



Time in ns



Time in ns



Time in ns



Time in ns

Miss (intended)	Miss (intended)	н	Miss	Miss	
<i>P</i> -2-1	-1-17	((34 ac	cesses,	191ns)
		Π			

Time in ns

Daniel Gruss — Graz University of Technology

(intended)

(intended)

Miss (interceled)	Miss (intended)	Miss	Miss
(intended)	(intended)		
P_2_1	_1_17	(34 20)	-95595
1 -2-1	-1-11	(34 act	
Miss (intended)	Miss (intended)	ннынын	Miss H

Time in ns

Miss (intended)	Miss (intended)	Miss	Miss
(intendèd)	(intended)		
P_2_1	_1_17 ('34 ac	-96696
1 -2-1	-1-1/ (J4 aC	
Miss (intended)	Miss (intended)	нинини	Miss H H
(intended)	(intended)		

Time in ns



Time in ns

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

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

Miss Miss (intended) (intended)	нинини и	Miss H H H H
------------------------------------	----------	--------------

Time in ns

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

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

Miss Miss (intended) (intended)	нннннннн Miss	нннн
------------------------------------	---------------	------

Time in ns

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

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

Miss Miss (intended) (intended)	нынынын Місс	ннннн
------------------------------------	--------------	-------

Time in ns

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

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

Miss Miss (intended) (intended)	ННННННН Міза	нннннн
------------------------------------	--------------	--------

Time in ns

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

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

Miss Miss (intended) (intended)	НИНИНИИ Miss	нннннн
------------------------------------	--------------	--------

Time in ns

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

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

Miss (intended)	Miss (intended)	ннннннн	Miss 9	нннннн	Miss
--------------------	--------------------	---------	--------	--------	------

Time in ns

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

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

Miss Miss (intended) (intended)	нныныны	Miss ННННННН	Miss
------------------------------------	---------	--------------	------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended)	Miss (intended)	ннннннн	Miss	ннинини	Miss	,
--------------------	--------------------	---------	------	---------	------	---

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended)	Miss (intended)	ннинини	Miss	нннннн	Miss	нн
--------------------	--------------------	---------	------	--------	------	----

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	HHHHHHH Miss	. нининии м	liss HHH
------------------------------------	--------------	-------------	----------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	нининини	Miss НИНИНИИ	Miss H H H H
------------------------------------	----------	--------------	--------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	нининин Маа	НИНИНИИ Miss	нннн
------------------------------------	-------------	--------------	------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	нининин м	ss ННИНИНИИ	Miss HHHHHH
------------------------------------	-----------	-------------	-------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	ннынынын Miss	нинини	Miss НННННН
------------------------------------	---------------	--------	-------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	нннннннн Miss	нининии	Miss НИНИНИ
------------------------------------	---------------	---------	-------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) H H H H H H Miss	HHHHHHH Miss	нининин _{Miss}
---------------------------------------	--------------	-------------------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended)	Miss НИНИИИИ Miss	ннинини Маа
----------------------	-------------------	-------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended)	Miss (intended)	ннинини	Miss НИНИНИИ	Miss ИНИНИНИ	Miss H H
--------------------	--------------------	---------	--------------	--------------	----------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended) HHHHHH	ss HHHHHHHH Miss	НИННИНИ Miss	ннн
---------------------------	------------------	--------------	-----

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	ннинини	Miss	нинини	Miss ННННННН	Miss HHHH
------------------------------------	---------	------	--------	--------------	-----------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss Miss (intended) (intended)	ныныны	Miss НИННИНИ	Miss HHHHHHHH	Miss ННННН
------------------------------------	--------	--------------	---------------	------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended)	Miss (intended)	нининии	Miss ННИННИИ	Miss ННННННН	Miss НИНИ
--------------------	--------------------	---------	--------------	--------------	-----------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended)	Miss (intended) HHHHHH	Miss	нынынын	Miss HHHHHHHH	Miss HHHHH
--------------------	---------------------------	------	---------	---------------	------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended)	Miss (intended) HHHHHH	Miss	нынынын	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	H Miss	Miss	Miss
--------------------	--------------------	--------	------	------	--------	------	------	--------	------	------

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

Miss (intended)	Miss (intended) HHHHHH	Miss	нынынын	Miss HHHHHHHH	Miss HHHHH
--------------------	---------------------------	------	---------	---------------	------------

Time in ns

Daniel Gruss — Graz University of Technology



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

Miss (intended)	Miss (intended)	нныныны	Miss	ныныныны	Miss	нининии	Miss HHHHH
--------------------	--------------------	---------	------	----------	------	---------	------------

Time in ns

Daniel Gruss — Graz University of Technology

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

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

Miss (intended) (Miss (intended)	нининини	Miss H H H	HHHH Miss	нннннн Міза	. нынын
----------------------	--------------------	----------	------------	-----------	-------------	---------

Time in ns

	Miss (intende	Miss d) (intended)	H Miss	Miss	Miss	H Miss	Miss	Miss	H Miss	Miss	Miss	4 Miss	Miss
--	------------------	-----------------------	--------	------	------	--------	------	------	--------	------	------	--------	------

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

Miss (intended)	Miss (intended) HHHHHH	Miss	нынынын	Miss HHHHHHHH	Miss HHHHH
--------------------	---------------------------	------	---------	---------------	------------

Time in ns





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. [...]

www.tugraz.at

Raw Prime+Probe trace...²



 $^2\mathsf{M}.$ Schwarz, D. Gruss, S. Weiser, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.

...processed with a simple moving average...³



³M. 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 exponent⁴



⁴M. Schwarz, D. Gruss, S. Weiser, C. Maurice, and S. Mangard. Malware Guard Extension: Using SGX to Conceal Cache Attacks. In: DIMVA. 2017.



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





ScatterCache





• Secure design for n-way set associative caches



- Secure design for n-way set associative caches
 - $\bullet \ \ \mathsf{Addresses} \to \mathsf{cache} \ \mathsf{sets}$
 - Increases number of possible cache sets
 - IDs for security domains



- Secure design for n-way set associative caches
 - Addresses \rightarrow ??? \rightarrow cache sets
 - Increases number of possible cache sets
 - IDs for security domains



- Secure design for n-way set associative caches
 - Addresses \rightarrow ??? \rightarrow cache sets
 - Increases number of possible cache sets
 - IDs for security domains
- Established/existing concepts
 - Skewed caches / slices
 - Low latency cryptography (e.g., QARMA-64 [Ava17])











Daniel Gruss — Graz University of Technology



Daniel Gruss — Graz University of Technology







Daniel Gruss — Graz University of Technology



 Index Derivation Function (IDF) takes an address and returns a cache set

- Depends on hardware key and optional Security Domain ID (SDID)
- → Unique combination of cache lines for each address



 Index Derivation Function (IDF) takes an address and returns a cache set

- Depends on hardware key and optional Security Domain ID (SDID)
- → Unique combination of cache lines for each address



512 KiB (32 B lines), $n_{ways} = 8$, $b_{indices} = 11$ $\rightarrow 2^{96.7}$ sets Index Derivation Function (IDF) takes an address and returns a cache set

- Depends on hardware key and optional Security Domain ID (SDID)
- → Unique combination of cache lines for each address



$$\binom{n_{ways} \cdot 2^{b_{indices}} + n_{ways} - 1}{n_{ways}}$$
 possible cache sets

512 KiB (32 B lines), $n_{ways} = 8$, $b_{indices} = 11$ $\rightarrow 2^{96.7}$ sets Index Derivation Function (IDF) takes an address and returns a cache set

www.tugraz.at

- Depends on hardware key and optional Security Domain ID (SDID)
- → Unique combination of cache lines for each address
- Potential index collisions
- One n_{ways} multi-port memory



We want something that is closer to a traditional cache!
We want something that is closer to a traditional cache!

instead of this:



We want something that is closer to a traditional cache!

instead of this:

let's do this:





- Skewed cache [Sez93] (*i.e.*, traditional cache + more addressing logic) and an IDF
- Like cache slices

Daniel Gruss — Graz University of Technology



 $2^{b_{indices} \cdot n_{ways}}$ possible cache sets

- Skewed cache [Sez93] (*i.e.*, traditional cache + more addressing logic) and an IDF
- Like cache slices

Daniel Gruss — Graz University of Technology



 $2^{b_{indices} \cdot n_{ways}}$ possible cache sets

512 KiB (32 B lines),
$$n_{ways} = 8$$
, $b_{indices} = 11$
 $ightarrow 2^{88}$ sets

- Skewed cache [Sez93] (*i.e.*, traditional cache + more addressing logic) and an IDF
- Like cache slices

Daniel Gruss — Graz University of Technology



 $2^{b_{indices} \cdot n_{ways}}$ possible cache sets

512 KiB (32 B lines),
$$n_{ways} = 8$$
, $b_{indices} = 11$
 $ightarrow 2^{88}$ sets

- Skewed cache [Sez93] (*i.e.*, traditional cache + more addressing logic) and an IDF
- Like cache slices
- Random replacement policy (for now)



- Inputs: cache line address, SDID, key
- Outputs: *n_{ways}* indices (*b_{indices}* bits)



ScatterCache - Selecting the IDF



- Inputs: cache line address, SDID, key
- Outputs: *n_{ways}* indices (*b_{indices}* bits)
- Existing concepts and crypto primitives



ScatterCache - Selecting the IDF

- Inputs: cache line address, SDID, key
- Outputs: *n_{ways}* indices (*b_{indices}* bits)
- Existing concepts and crypto primitives
- SCv1: hashing
 - Block ciphers (e.g., PRINCE [Bor+12])
 - Tweakable block ciphers (e.g., QARMA [Ava17])
 - Permutation-based primitives (e.g., Keccak-p [Ber+11])



ScatterCache - Selecting the IDF

- Inputs: cache line address, SDID, key
- Outputs: *n_{ways}* indices (*b_{indices}* bits)
- Existing concepts and crypto primitives
- SCv1: hashing
 - Block ciphers (e.g., PRINCE [Bor+12])
 - Tweakable block ciphers (e.g., QARMA [Ava17])
 - Permutation-based primitives (e.g., Keccak-p [Ber+11])
- SCv2: permutation
 - Prevents birthday-bound index collisions
 - No off-the-shelf primitives



• ScatterCache \rightarrow last-level cache



ScatterCache - System Integration

- ScatterCache \rightarrow last-level cache
- Hardware managed key
 - Randomly generated at boot time
 - Rekeying with full cache flush



ScatterCache - System Integration

- ScatterCache \rightarrow last-level cache
- Hardware managed key
 - Randomly generated at boot time
 - Rekeying with full cache flush
 - Potential for iterative rekeying
 - \rightarrow concurrent: CEASER-S @ISCA [Qur19]



ScatterCache - System Integration

- ScatterCache \rightarrow last-level cache
- Hardware managed key
 - Randomly generated at boot time
 - Rekeying with full cache flush
 - Potential for iterative rekeying
 - \rightarrow concurrent: CEASER-S @ISCA [Qur19]
- SDID management via page table (indirection)
 - x86: Page Attribute Tables (PATs)
 - ARM: Memory Attribute Indirection Register (MAIRs)





- no OS support? \rightarrow default SDID = 0
- OS support? \rightarrow page-wise security domains



- no OS support? \rightarrow default SDID = 0
- OS support? \rightarrow page-wise security domains
 - ightarrow shared read-only pages can be private in the cache!



- no OS support? \rightarrow default SDID = 0
- OS support? → page-wise security domains
 → shared read-only pages can be private in the cache!
- OS defines domains

(pages, processes, containers, VMs, ...)



- no OS support? \rightarrow default SDID = 0
- OS support? → page-wise security domains
 → shared read-only pages can be private in the cache!
- OS defines domains

(pages, processes, containers, VMs, ...)

• Software-based page *rekeying* by changing the SDID

- Unshared memory has no shared (physical) addresses
 - \rightarrow No Flush+Reload, Evict+Reload, Flush+Flush
 - \rightarrow Specialized Prime+Probe is possible

- Unshared memory has no shared (physical) addresses
 - \rightarrow No Flush+Reload, Evict+Reload, Flush+Flush
 - \rightarrow Specialized Prime+Probe is possible
- Shared, read-only memory
 - \rightarrow Like unshared memory given OS support
 - $\rightarrow~$ Otherwise, eviction-based attacks are hindered

- Unshared memory has no shared (physical) addresses
 - \rightarrow No Flush+Reload, Evict+Reload, Flush+Flush
 - \rightarrow Specialized Prime+Probe is possible
- Shared, read-only memory
 - \rightarrow Like unshared memory given OS support
 - $\rightarrow~$ Otherwise, eviction-based attacks are hindered
- Shared, writable memory can't be separated
 - $\rightarrow\,$ Eviction-based attacks are hindered



- No end-to-end attack yet
 - $\rightarrow\,$ Simplified setting: perfect control, single access, no noise
 - \rightarrow Investigate the building blocks in simulation and analytically

- No end-to-end attack yet
 - $\rightarrow\,$ Simplified setting: perfect control, single access, no noise
 - \rightarrow Investigate the building blocks in simulation and analytically
- Finding congruent addresses ($n_{ways} = 8, b_{indices} = 11$)
 - $\bullet~\mbox{Full collisions}$ are unlikely $\rightarrow~\mbox{use partial collisions}$
 - Approach in the paper: $\approx 2^{25}$ profiled victim accesses

- No end-to-end attack yet
 - \rightarrow Simplified setting: perfect control, single access, no noise
 - \rightarrow Investigate the building blocks in simulation and analytically
- Finding congruent addresses $(n_{wavs} = 8, b_{indices} = 11)$
 - Full collisions are unlikely \rightarrow use partial collisions
 - Approach in the paper: $\approx 2^{25}$ profiled victim accesses
- Evicting one set with 99% needs 275 addresses

~?

- No end-to-end attack yet
 - $\rightarrow\,$ Simplified setting: perfect control, single access, no noise
 - \rightarrow Investigate the building blocks in simulation and analytically
- Finding congruent addresses ($n_{ways} = 8, b_{indices} = 11$)
 - $\bullet\,$ Full collisions are unlikely \rightarrow use partial collisions
 - Approach in the paper: $\approx 2^{25}$ profiled victim accesses
- Evicting one set with 99% needs 275 addresses
- Two Prime+Probe variants ($n_{ways} = 8, b_{indices} = 12$)
 - 99% confidence: 35 to 152 victim accesses (repetitions)
 - Between 9870 and 1216 congruent addresses

â

- No end-to-end attack yet
 - $\rightarrow\,$ Simplified setting: perfect control, single access, no noise
 - \rightarrow Investigate the building blocks in simulation and analytically
- Finding congruent addresses ($n_{ways} = 8, b_{indices} = 11$)
 - $\bullet\,$ Full collisions are unlikely \rightarrow use partial collisions
 - Approach in the paper: $\approx 2^{25}$ profiled victim accesses
- Evicting one set with 99% needs 275 addresses
- Two Prime+Probe variants ($n_{ways} = 8, b_{indices} = 12$)
 - 99% confidence: 35 to 152 victim accesses (repetitions)
 - Between 9870 and 1216 congruent addresses
- Noise? Leakage analysis?

Daniel Gruss — Graz University of Technology

- No end-to-end attack yet
 - $\rightarrow\,$ Simplified setting: perfect control, single access, no noise
 - \rightarrow Investigate the building blocks in simulation and analytically
- Finding congruent addresses ($n_{ways} = 8, b_{indices} = 11$)
 - $\bullet\,$ Full collisions are unlikely \rightarrow use partial collisions
 - Approach in the paper: $\approx 2^{25}$ profiled victim accesses
 - Generalized by Purnal and Verbauwhede [PV19]: $\approx 2^{10}$
- Evicting one set with 99% needs 275 addresses
- Two Prime+Probe variants ($n_{ways} = 8, b_{indices} = 12$)
 - 99% confidence: 35 to 152 victim accesses (repetitions)
 - Between 9870 and 1216 congruent addresses
- Noise? Leakage analysis?

Daniel Gruss — Graz University of Technology



- Micro benchmarks using the gem5 full system simulator (ARM)
 - Poky Linux from Yocto 2.5 (kernel version 4.14.67)
 - GAP, MiBench, Imbench, scimark2
- SPEC CPU 2017 on custom cache simulator



- Micro benchmarks using the gem5 full system simulator (ARM)
 - Poky Linux from Yocto 2.5 (kernel version 4.14.67)
 - GAP, MiBench, Imbench, scimark2
- SPEC CPU 2017 on custom cache simulator
- Cache hit rate \geq set-associative caches (random replacement)
- 2% 4% below LRU on micro benchmarks, 0% 2% for SPEC

all these cache attacks over the past decades...



7. Serve with cooked and peeled potatoes







Wait for an hour



Wait for an hour

LATENCY

1. Wash and cut vegetables

2. Pick the basil leaves and set aside

3. Heat 2 tablespoons of oil in a pan

4. Fry vegetables until golden and softened


1. Wash and cut vegetables

Parallelize

2. Pick the basil leaves and set aside

3. Heat 2 tablespoons of oil in a pan

4. Fry vegetables until golden and softened



















segfault at ffffffff81a000e0 ip
0000000000400535
sp 00007ffce4a80610 error 5 in reader

• Kernel addresses are not accessible





segfault at ffffffff81a000e0 ip 0000000000400535 sp 00007ffce4a80610 error 5 in reader

- Kernel addresses are not accessible
- Are privilege checks also done when executing instructions out of order?



• Adapted code



```
*(volatile char*)0;
array[84 * 4096] = 0; // unreachable
```





```
*(volatile char*)0;
array[84 * 4096] = 0; // unreachable
```

```
• Static code analyzer is not happy
```

1 warning: Dereference of null pointer

```
2 *(volatile char*)0;
```





• "Unreachable" code line was actually executed



www.tugraz.at



• Flush+Reload over all pages of the array



- "Unreachable" code line was actually executed
- Exception was only thrown afterwards





• Maybe there is no permission check in transient instructions...



- Maybe there is no permission check in transient instructions...
- ...or it is only done when commiting them

Let's hope this does not work...



- Maybe there is no permission check in transient instructions...
- ...or it is only done when commiting them
- Indirection through microarchitectural traces:

Let's hope this does not work...



- Maybe there is no permission check in transient instructions...
- ...or it is only done when commiting them
- Indirection through microarchitectural traces:

• Check whether any part of array is cached



• Flush+Reload over all pages of the array



• Index of cache hit reveals data



• Flush+Reload over all pages of the array



- Index of cache hit reveals data
- Permission check is in some cases not fast enough



















Speculative Cooking













Spectre-PHT (v1)





www.tugraz.at

Spectre-PHT (v1)



www.tugraz.at

Spectre-PHT (v1)



www.tugraz.at 📕

Spectre-PHT (v1)


Spectre-PHT (v1)





Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)





Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)





Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)





Daniel Gruss — Graz University of Technology

46

Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)





Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)





Spectre-PHT (v1)



Spectre-PHT (v1)



Spectre-PHT (v1)





Spectre-STL (v4): Ignore sanitizing write access and use unsanitized old value instead


































































Spectre v2







Spectre v2







Spectre-BTB (v2): mistrain BTB \rightarrow mispredict indirect jump/call

Spectre v2







Spectre-BTB (v2): mistrain BTB \rightarrow mispredict indirect jump/call Spectre-RSB (v5): mistrain RSB \rightarrow mispredict return • v1.1: Speculatively write to memory locations

⁵V. Kiriansky and C. Waldspurger. Speculative Buffer Overflows: Attacks and Defenses. In: arXiv:1807.03757 (2018).

- v1.1: Speculatively write to memory locations
- $\rightarrow\,$ Many more gadgets than previously anticipated n

 $^5\text{V}.$ Kiriansky and C. Waldspurger. Speculative Buffer Overflows: Attacks and Defenses. In: arXiv:1807.03757 (2018).

- v1.1: Speculatively write to memory locations
- $\rightarrow\,$ Many more gadgets than previously anticipated n
 - v1.2: Ignore writable bit

⁵V. Kiriansky and C. Waldspurger. Speculative Buffer Overflows: Attacks and Defenses. In: arXiv:1807.03757 (2018).

- v1.1: Speculatively write to memory locations
- $\rightarrow\,$ Many more gadgets than previously anticipated n
 - v1.2: Ignore writable bit
- $\rightarrow \,=\, {\sf Meltdown}{\sf -}{\sf RW}$

⁵V. Kiriansky and C. Waldspurger. Speculative Buffer Overflows: Attacks and Defenses. In: arXiv:1807.03757 (2018).



Shared Branch Prediction State

Classification Tree

www.tugraz.at



Mitigations?





JOIN BENEFIT - CONTRIBUTE - DISCOVER - BLOG

Computer Architecture Today

Informing the broad computing community about current activities, advances and future directions in computer architecture.

Let's Keep it to Ourselves: Don't Disclose Vulnerabilities

by Gus Uht on Jan 31, 2019 | Tags: Opinion, Security



CONTRIBUTE

Editor: Alvin R. Lebeck Associate Editor: Vijay Janapa Reddi

Contribute to Computer Architecture Today



Symbols show if an attack is mitigated (●), partially mitigated (●), not mitigated (○), theoretically mitigated (■), theoretically impeded (■), not theoretically impeded (□), or out of scope (◇).

www.tugraz.at 📕

Table 2: Reported performance impacts of countermeasures

Impact Defense	Performance Loss	Benchmark
InvisiSpec	22%	SPEC
SafeSpec	3% (improvement)	SPEC2017 on MARSSx86
DAWG	2-12%, 1-15%	PARSEC, GAPBS
RSB Stuffing	no reports	
Retpoline	5-10%	real-world workload servers
Site Isolation	only memory overhead	
SLH	36.4%, 29%	Google microbenchmark suite
YSNB	60%	Phoenix
IBRS	20-30%	two sysbench 1.0.11 benchmarks
STIPB	30- 50%	Rodinia OpenMP, DaCapo
IBPB	no individual reports	
Serialization	62%, 74.8%	Google microbenchmark suite
SSBD/SSBB	2-8%	SYSmark®2014 SE & SPEC integer
KAISER/KPTI	0-2.6%	system call rates
L1TF mitigations	-3-31%	various SPEC

ConTExT















• Mark secrets in source code



- Mark secrets in source code
- Propagate taint through memory hierarchy:



- Mark secrets in source code
- Propagate taint through memory hierarchy:
 - Pages



- Mark secrets in source code
- Propagate taint through memory hierarchy:
 - Pages
 - Cache Lines (in caches and buffers)



- Mark secrets in source code
- Propagate taint through memory hierarchy:
 - Pages
 - Cache Lines (in caches and buffers)
 - Registers





Unprotected























• Taint tracking is complicated and expensive





• Taint tracking is complicated and expensive





- Taint tracking is complicated and expensive ightarrow we don't do that
 - Software unintentionally losing taint?





- Taint tracking is complicated and expensive ightarrow we don't do that
 - Software unintentionally losing taint?



- Taint tracking is complicated and expensive ightarrow we don't do that
 - Software unintentionally losing taint? \rightarrow software problem
 - Variable copies?


- Taint tracking is complicated and expensive ightarrow we don't do that
 - Software unintentionally losing taint? \rightarrow software problem
 - Variable copies?



- Taint tracking is complicated and expensive \rightarrow we don't do that
- Software unintentionally losing taint? \rightarrow software problem
- \bullet Variable copies? \rightarrow compiler should avoid / warn
- Otherwise: always load/store from the same location



- Taint tracking is complicated and expensive \rightarrow we don't do that
- Software unintentionally losing taint? \rightarrow software problem
- \bullet Variable copies? \rightarrow compiler should avoid / warn
- Otherwise: always load/store from the same location



- $\bullet\,$ Taint tracking is complicated and expensive \rightarrow we don't do that
- \bullet Software unintentionally losing taint? \rightarrow software problem
- Variable copies? \rightarrow compiler should avoid / warn
- Otherwise: always load/store from the same location (compiler does that anyway)















 $\bullet\,$ Fully overwriting a tainted register $\rightarrow\,$ untaint



- $\bullet\,$ Fully overwriting a tainted register $\rightarrow\,$ untaint
- Writing to unprotected memory exposes value to attackers



- $\bullet\,$ Fully overwriting a tainted register $\rightarrow\,$ untaint
- Writing to unprotected memory exposes value to attackers
 - $\rightarrow \ {\sf Untaint \ register}$

Avoiding Taint Explosion



- $\bullet\,$ Fully overwriting a tainted register $\rightarrow\,$ untaint
- Writing to unprotected memory exposes value to attackers \rightarrow Untaint register
- Split stack into protected and unprotected half

Avoiding Taint Explosion



- $\bullet\,$ Fully overwriting a tainted register $\rightarrow\,$ untaint
- Writing to unprotected memory exposes value to attackers \rightarrow Untaint register
- Split stack into protected and unprotected half
- Stack spills of unprotected data \rightarrow stay unprotected as long as they stay in the cache

What about function arguments?



What about function arguments?





• Parameter declared secret? \rightarrow non-transient (over-tainting)



- Parameter declared secret? \rightarrow non-transient (over-tainting)
- Copy to non-secret local variable if necessary (e.g., based on arguments)







• Then you already have a side channel



- Then you already have a side channel
- $\rightarrow\,$ No need for a more complicated transient-execution attack



- Then you already have a side channel
- $\rightarrow\,$ No need for a more complicated transient-execution attack
- $\rightarrow\,$ Compiler issues a warning



• different classes of attacks need different classes of defenses



- different classes of attacks need different classes of defenses
- interesting opportunities for new defensive research

√ —
∽ —

- different classes of attacks need different classes of defenses
- interesting opportunities for new defensive research
- we need more cat-and-mouse game in system security (like in crypto algorithm design)



Hardware-Software Co-Design against Microarchitectural Attacks

Daniel Gruss

September 4, 2019

Graz University of Technology