

Hardware-Software Co-Design to Eliminate Cache Leakage

Lukas Giner, Daniel Gruss

June 11, 2019

Graz University of Technology



Lukas Giner

PhD student @ Graz University of Technology

- ♥ @redrabbyte
- ☑ lukas.giner@iaik.tugraz.at



Daniel Gruss

Asst. Prof. @ Graz University of Technology

- 🕑 @lavados
- ☑ daniel.gruss@iaik.tugraz.at

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? \rightarrow 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? \rightarrow Yes
- Measuring memory access latency with Flush+Reload? \rightarrow No

$\mathbf{}$	

- 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?

$\boldsymbol{\wedge}$	

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

www.tugraz.at

• actual misspeculation (e.g., branch misprediction)



• traditional cache attacks (crypto, keys, etc)

www.tugraz.at

- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc



• traditional cache attacks (crypto, keys, etc)

www.tugraz.at

- actual misspeculation (e.g., branch misprediction)
- Meltdown, Foreshadow, ZombieLoad, etc



• traditional cache attacks (crypto, keys, etc)

www.tugraz.at

- 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)

www.tugraz.at

- 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)

www.tugraz.at

- 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?

• CPU Caches are a shared resource





CPU Caches are a shared resource
 → there's probably a side-channel



- CPU Caches are a shared resource
 → there's probably a side-channel
- Cache side-channels are currently very powerful



- CPU Caches are a shared resource
 → there's probably a side-channel
- Cache side-channels are currently very powerful
- At the moment: simple, fast & reliable



- CPU Caches are a shared resource
 → there's probably a side-channel
- Cache side-channels are currently very powerful
- At the moment: simple, fast & reliable
- They enable a large number of attacks, e.g.:



- CPU Caches are a shared resource
 → there's probably a side-channel
- Cache side-channels are currently very powerful
- At the moment: simple, fast & reliable
- They enable a large number of attacks, e.g.: "direct" cache attacks,



- CPU Caches are a shared resource
 → there's probably a side-channel
- Cache side-channels are currently very powerful
- At the moment: simple, fast & reliable
- They enable a large number of attacks, e.g.: "direct" cache attacks, Spectre & Meltdown,



- CPU Caches are a shared resource
 → there's probably a side-channel
- Cache side-channels are currently very powerful
- At the moment: simple, fast & reliable
- They enable a large number of attacks, e.g.: "direct" cache attacks, Spectre & Meltdown, Zombieload, ...

How do caches work?



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



CPU Cache





CPU Cache



www.tugraz.at




www.tugraz.at 📕



Lukas Giner, Daniel Gruss — Graz University of Technology

www.tugraz.at









Cache Hits



www.tugraz.at

Cache Hits Cache Misses













 \rightarrow replacement policy























ScatterCache

if cache attacks are simple because the mapping to sets is simple ...

if cache attacks are simple because the mapping to sets is simple ...

instead of this:



if cache attacks are simple because the mapping to sets is simple \ldots









ScatterCache - Hardware



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

ScatterCache - Hardware



- Index Derivation Function (IDF) takes an address and returns a cache set
- Depends on hardware key K and optional Security Domain ID (SDID)

ScatterCache - Hardware



- Index Derivation Function (IDF) takes an address and returns a cache set
- Depends on hardware key K and optional Security Domain ID (SDID)
- → unique combination of cache lines for each address



www.tugraz.at



• ScatterCache requires no software support, default SDID = 0



- ScatterCache requires no software support, default SDID = 0
- But OS support enables security domains



- ScatterCache requires no software support, default SDID = 0
- But OS support enables security domains
 - ightarrow shared read-only pages can be private in the cache!



- ScatterCache requires no software support, default SDID = 0
- But OS support enables security domains \rightarrow shared read-only pages can be private in the cache!
- OS can define SDID per process and separate user space and kernel space



- ScatterCache requires no software support, default SDID = 0
- But OS support enables security domains
 → shared read-only pages can be private in the cache!
- OS can define SDID per process and separate user space and kernel space
- Process can request distinct SDIDs for memory ranges
• Unshared memory has no shared cache lines

Unshared memory has no shared cache lines
 → Flush+Reload, Flush+Flush and Evict+Reload are not possible

- Unshared memory has no shared cache lines
 → Flush+Reload, Flush+Flush and Evict+Reload are not possible
- Shared, read-only memory is like unshared memory, given OS support. Without OS support, eviction-based attacks are hindered

- Unshared memory has no shared cache lines
 → Flush+Reload, Flush+Flush and Evict+Reload are not possible
- Shared, read-only memory is like unshared memory, given OS support. Without OS support, eviction-based attacks are hindered
- Shared, writable memory can't be separated, eviction-based attacks are hindered

• Specialized Prime+Probe variants are still possible



• Specialized Prime+Probe variants are still possible

2

• But, overlap in more than 1 cache line is very unlikely \rightarrow Eviction is now probabilistic, $p = \frac{1}{n_{max}^2}$ to evict

- Specialized Prime+Probe variants are still possible
- But, overlap in more than 1 cache line is very unlikely \rightarrow Eviction is now probabilistic, $p = \frac{1}{p_{max}^2}$ to evict
- Evicting an address with 99% certainty needs 275 addresses for 8-way cache, instead of ≈ 8 for standard Prime+Probe

- Specialized Prime+Probe variants are still possible
- But, overlap in more than 1 cache line is very unlikely \rightarrow Eviction is now probabilistic, $p = \frac{1}{p_{max}^2}$ to evict
- Evicting an address with 99% certainty needs 275 addresses for 8-way cache, instead of \approx 8 for standard Prime+Probe
- Constructing this set requires $\approx 2^{25}$ profiled victim accesses, compared to less than 100 accesses for standard, noise-free Prime+Probe





• Macro benchmarks from SPEC CPU 2017 on custom cache simulator



- Macro benchmarks from SPEC CPU 2017 on custom cache simulator
- Cache hit rate always at or above levels of set-associative cache with random replacement



- Macro benchmarks from SPEC CPU 2017 on custom cache simulator
- Cache hit rate always at or above levels of set-associative cache with random replacement
- Typically 2% 4% below LRU on micro benchmarks, 0% 2% for 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













• Writing to unprotected memory exposes value to attackers



• Writing to unprotected memory exposes value to attackers \rightarrow Untaint register



- Writing to unprotected memory exposes value to attackers \rightarrow Untaint register
- Split stack into protected and unprotected half



- 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

- C. Canella, J. Van Bulck, M. Schwarz, M. Lipp, B. von Berg, P. Ortner, F. Piessens, D. Evtyushkin, and D. Gruss. A Systematic Evaluation of Transient Execution Attacks and Defenses. In: USENIX Security Symposium. 2019.
- D. Gruss, E. Kraft, T. Tiwari, M. Schwarz, A. Trachtenberg, J. Hennessey, A. Ionescu, and A. Fogh. Page Cache Attacks. In: CCS. 2019.
- M. Schwarz, C. Canella, L. Giner, and D. Gruss. Store-to-Leak Forwarding: Leaking Data on Meltdown-resistant CPUs. In: arXiv:1905.05725 (2019).
- M. Schwarz, M. Lipp, D. Moghimi, J. Van Bulck, J. Stecklina, T. Prescher, and D. Gruss. ZombieLoad: Cross-Privilege-Boundary Data Sampling. In: arXiv:1905.05726 (2019).
- M. Schwarz, R. Schilling, F. Kargl, M. Lipp, C. Canella, and D. Gruss. ConTExT: Leakage-Free Transient Execution. In: arXiv:1905.09100 (2019).
- M. Schwarz, S. Weiser, and D. Gruss. Practical Enclave Malware with Intel SGX. In: $\mathsf{DIMVA}.$ 2019.
- M. Werner, T. Unterluggauer, L. Giner, M. Schwarz, D. Gruss, and S. Mangard. ScatterCache: Thwarting Cache Attacks via Cache Set Randomization. In: USENIX Security Symposium. 2019.



• many attacks out there



- many attacks out there
- thorough defenses can defeat entire classes of attacks


- many attacks out there
- thorough defenses can defeat entire classes of attacks
- important to distinguish between different attacks



Hardware-Software Co-Design to Eliminate Cache Leakage

Lukas Giner, Daniel Gruss

June 11, 2019

Graz University of Technology







• Compiler Extension



- Compiler Extension
- Linux Patch



- Compiler Extension
- Linux Patch
- CPU Emulation in Bochs



- Compiler Extension
- Linux Patch
- CPU Emulation in Bochs
- Native via uncacheable memory (ConTExT-light)

Benchmark	SPEC Baseline	ConTExT	Overhead [%]
600.perlbench_s	7.03	6.86	+2.42
602.gcc_s	11.90	11.80	+0.84
605.mcf_s	9.06	9.16	-1.10
620.omnetpp_s	5.07	4.81	+5.13
623.xalancbmk_s	6.06	5.95	+1.82
625.×264_s	9.25	9.25	0.00
631.deepsjeng_s	5.26	5.22	+0.76
641.leela_s	4.71	4.64	+1.48
648.exchange2_s	would require Fortran runtime		
657.×z_s	12.10	12.10	0.00
Average			+1.26

 Table 1: Performance of the ConTExT split stack using the SPECspeed 2017 integer benchmark.