

SPOILER: Speculative Load Hazards Boost Rowhammer and Cache Attacks

Saad Islam¹, Ahmad Moghimi¹, Ida Bruhns², Moritz Krebbel², Berk Gulmezoglu¹, Thomas Eisenbarth^{1, 2},
and Berk Sunar¹

¹Worcester Polytechnic Institute, Worcester, MA, USA

²University of Lübeck, Lübeck, Germany

Abstract

Modern microarchitectures incorporate optimization techniques such as *speculative loads* and *store forwarding* to improve the memory bottleneck. The processor executes the load speculatively before the stores, and forwards the data of a preceding store to the load if there is a potential dependency. This enhances performance since the load does not have to wait for preceding stores to complete. However, the dependency prediction relies on partial address information, which may lead to false dependencies and stall hazards.

In this work, we are the first to show that the dependency resolution logic that serves the speculative load can be exploited to gain information about the physical page mappings. Microarchitectural side-channel attacks such as Rowhammer and cache attacks rely on the reverse engineering of the virtual-to-physical address mapping. We propose the SPOILER attack which exploits this leakage to speed up this reverse engineering by a factor of 256. Then, we show how this can improve the Prime+Probe attack by a 4096 factor speed up of the eviction set search, even from sandboxed environments like JavaScript. Finally, we improve the Rowhammer attack by showing how SPOILER helps to conduct DRAM row conflicts deterministically with up to 100% chance, and by demonstrating a double-sided Rowhammer attack with normal user's privilege. The later is due to the possibility of detecting contiguous memory pages using the SPOILER leakage.

1 Introduction

Microarchitectural attacks have evolved over the past decade from attacks on weak cryptographic implementations [6] to devastating attacks breaking through layers of defenses provided by the hardware and the Operating System (OS) [49]. These attacks can steal secrets such as cryptographic keys [5, 41] or keystrokes [30]. More advanced attacks can entirely subvert the OS memory isolation to read the memory content from more privileged security domains [32], and to bypass defense mechanisms such as Kernel Address Space Layout

Randomization (KASLR) [10, 16]. Rowhammer attacks can further break the data and code integrity by tampering with memory contents [26, 44]. While most of these attacks require local access and native code execution, various efforts have been successful in conducting them remotely [47] or from within a remotely accessible sandbox such as JavaScript [39].

Memory components such as DRAM [26] and cache [40] are not the only microarchitectural attack surfaces. *Spectre* attacks on the speculative branch prediction unit [27, 35] imply that side channels such as caches can be used as a primitive for more advanced attacks on speculative engines. Speculative engines predict the outcome of an operation before its completion, and they enable execution of the following dependent instructions ahead of time based on the prediction. As a result, the pipeline can maximize the instruction level parallelism and resource usage. In rare cases where the prediction is wrong, the pipeline needs to be flushed resulting in performance penalties. However, this approach suffers from a security weakness, in which an adversary can fool the predictor and introduce arbitrary mispredictions that leave microarchitectural footprints in the pipeline. These footprints can be collected through the cache side channel to steal secrets.

Modern processors feature further speculative behavior such as *memory disambiguation*, store forwarding and speculative loads [9]. A load operation can be executed speculatively before preceding store operations. During the speculative execution of the load, false dependencies may occur due to the unavailability of physical address information. These false dependencies need to be resolved to avoid computation on invalid data. The occurrence of false dependencies and their resolution depend on the actual implementation of the memory subsystem. Intel uses a proprietary memory disambiguation and *dependency resolution logic* in the processors to predict and resolve false dependencies that are related to the speculative load. In this work, we discover that the dependency resolution logic suffers from an unknown false dependency independent of the 4K aliasing [37, 46]. The discovered false dependency happens during the 1 MB aliasing of speculative memory accesses which is exploited to leak

information about physical page mappings.

The state-of-the-art microarchitectural attacks [22, 42] either rely on knowledge of physical addresses or are significantly eased by that knowledge. Yet, knowledge of the physical address space is only granted with root privileges. Cache attacks such as *Prime+Probe* on the Last-Level Cache (LLC) are challenging due to the unknown mapping of virtual addresses to cache sets and slices. Knowledge about the physical page mappings enables more attack opportunities using the Prime+Probe technique. Rowhammer [26] attacks require efficient access to rows within the same bank to induce fast row conflicts. To achieve this, an adversary needs to reverse engineer layers of abstraction from the virtual address space to DRAM cells. Availability of physical address information facilitates this reverse engineering process. In sandboxed environments, attacks are more limited, since in addition to the limited access to the address space, low-level instructions are also inaccessible [17]. Previous attacks assume special access privileges only granted through weak software configurations [22, 31, 52] to overcome some of these challenges. In contrast, SPOILER only relies on simple operations, `load` and `store`, to recover crucial physical address information, which in turn enables Rowhammer and cache attacks, by leaking information about physical pages without assuming any weak configuration or special privileges.

1.1 Our Contribution

We have discovered a novel microarchitectural leakage which reveals critical information about physical page mappings to user space processes. The leakage can be exploited by a limited set of instructions, which is visible in all Intel generations starting from the 1st generation of Intel Core processors, independent of the OS and also works from within virtual machines and sandboxed environments. In summary, this work:

1. exposes a previously unknown microarchitectural leakage stemming from the false dependency hazards during speculative load operations.
2. proposes an attack, SPOILER, to efficiently exploit this leakage to speed up the reverse engineering of virtual-to-physical mappings by a factor of 256 from both native and JavaScript environments.
3. demonstrates a novel eviction set search technique from JavaScript and compares its reliability and efficiency to existing approaches.
4. achieves efficient DRAM row conflicts and the first *double-sided Rowhammer* attack with normal user-level privilege using the contiguous memory detection capability of SPOILER.
5. explores how SPOILER can track nearby load operations from a more privileged security domain right after a context switch.

1.2 Related Work

Spectre attacks [27, 35] exploit vulnerabilities in the speculative branch prediction unit. Transient execution of instructions after a fault, as exploited by *Meltdown* [32] and *Foreshadow* [49], can leak memory content of protected environments. Similarly, transient behavior due to the lazy store/restore of the FPU and SIMD registers can leak register contents from other contexts [45]. New variants of both Meltdown and Spectre have been systematically analyzed [7]. The Speculative Store Bypass (SSB) vulnerability [13] is a variant of the Spectre attack and relies on the stale sensitive data in registers to be used as an address for speculative loads which may then allow the attacker to read this sensitive data. In contrast to previous attacks on speculative and transient behaviors, we discover a new leakage on the undocumented memory disambiguation and dependency resolution logic. SPOILER is **not** a Spectre attack. The root cause for SPOILER is a weakness in the address speculation of Intel’s proprietary implementation of the memory subsystem which directly leaks timing behavior due to physical address conflicts. Existing spectre mitigations would therefore not interfere with SPOILER.

The timing behavior of the 4K aliasing false dependency on Intel processors have been studied [3, 58]. *MemJam* [37] uses this behavior to perform a side-channel attack, and Sullivan et al. [46] demonstrate a covert channel. These works only mention the 4K aliasing as documented by Intel [21], and the authors conclude that the address aliasing check is a two stage approach: Firstly, it uses page offset for the initial guess. Secondly, it performs the final resolution based on the exact physical address. On the contrary, we discover that the undocumented *address resolution* logic performs additional partial address checks that lead to an unknown, but observable aliasing behavior based on the physical address.

Several microarchitectural attacks have been discovered to recover virtual address information and break KASLR by exploiting the Translation Lookaside Buffer (TLB) [19], Branch Target Buffer (BTB) [10] and Transactional Synchronization Extensions (TSX) [24]. Additionally, Gruss et al. [16] exploit the timing information obtained from the `prefetch` instruction to leak the physical address information. The main obstacle to this approach is that the `prefetch` instruction is not accessible in JavaScript, and it can be disabled in native sandboxed environments [59], whereas SPOILER is applicable to sandboxed environments including JavaScript.

Knowledge of the physical address enables adversaries to bypass OS protections [25] and ease other microarchitectural attacks [31]. For instance, the `procfs` filesystem exposes physical addresses [31], and *Huge pages* allocate contiguous physical memory [22, 33]. *Drammer* [52] exploits the Android *ION memory allocator* to access contiguous memory. However, access to the aforementioned primitives is restricted on most environments by default. We do not have any assumption about the OS and software configuration, and we exploit

a hardware leakage with minimum access rights to find virtual pages that have the same least significant 20 physical address bits. *GLitch* [11] detects contiguous physical pages by exploiting row conflicts through the GPU interface. In contrast, our attack does not rely on a specific integrated GPU configuration, and it is widely applicable to any system running on an Intel CPU. We use SPOILER to find continuous physical pages with a high probability and verify it by producing row conflicts. SPOILER is particularly helpful for attacks in sandboxed low-privilege environments such as JavaScript, where previous methods require a time-consuming brute forcing of the memory addresses [17, 39, 44].

2 Background

2.1 Memory Management

The virtual memory manager shares the DRAM across all running tasks by assigning isolated virtual address spaces to each task. The assigned memory is allocated in pages, which are typically 4 kB each, and each virtual page will be stored as a physical page in DRAM through a virtual-to-physical page mapping. Memory instructions operate on virtual addresses, which are translated within the processor to the corresponding physical addresses. The page offset comprising of least significant 12 bits of the virtual address is not translated. The processor only translates the bits in the rest of the virtual address, the virtual page number. The OS is the reference for this translation, and the processor stores the translation results inside the TLB. As a result, repeated translations of the same address are performed more efficiently.

2.2 Cache Hierarchy

Modern processors incorporate multiple levels of caches to avoid the DRAM access latency. The cache memory on Intel processors is organized into sets and slices. Each set can store a certain number of lines, where the line size is 64 bytes. The 6 Least Significant Bits (LSBs) of the physical address are used to determine the offset within a line and the remaining bits are used to determine which set to store the cache line in. The number of physical address bits that are used for mapping is higher for the LLC, since it has a large number of sets, e.g., 8192 sets. Hence, the untranslated part of the virtual address bits which is the page offset, cannot be used to index the LLC sets. Instead, higher physical address bits are used. Further, each set of LLC is divided into multiple slices, one slice for each logical processor. The mapping of the physical addresses to the slices uses an undocumented function [23]. When the processor accesses a memory address, a cache hit or miss occurs. If a miss occurs in all cache levels, the memory line has to be fetched from DRAM. Accesses to the same memory address would be served from the cache unless other memory accesses evict that cache line. In addition, we can use the

`clflush` instruction, which follows the same memory access check as other memory operations, to evict our own cache lines from the entire cache hierarchy.

2.3 Prime+Probe Attack

In the Prime+Probe attack, the attacker first fills an entire cache set by accessing memory addresses that are mapped to the same set, an *eviction set*. Later, the attacker checks whether the victim program has displaced any entry in the cache set by accessing the eviction set again and measuring the execution time. If this is the case, the attacker can detect congruent addresses, since the displaced entries cause an increased access time. However, finding the eviction sets is difficult due to the unknown translation of virtual addresses to physical addresses. Since an unprivileged attacker has no access to hugepages [20] or the virtual-to-physical page mapping such as the `pagemap` file [31], knowledge about the physical address bits greatly speeds up the eviction set search. As mentioned in Section 2.1, the virtual to physical addresses are identical in the least 12 significant bits.

2.4 Rowhammer Attack

DRAM consists of multiple memory banks, and each bank is subdivided into rows. When the processor accesses a memory location, the corresponding row needs to be activated and loaded into the row buffer. If the processor accesses the same row again, it is called a row hit, and the request will be served from the row buffer. Otherwise, it is called a row conflict, and the previous row will be deactivated and copied back to the original row location, after which the new row is activated. DRAM cells leak charge over time and need to be refreshed periodically to maintain the data. A Rowhammer [26] attack causes cells of a victim row to leak faster by activating the neighboring rows repeatedly. If the refresh cycle fails to refresh the victim fast enough, that leads to bit flips. Once bit flips are found, they can be exploited by placing any security-critical data structure or code page at that particular location and triggering the bit flip again [15, 44, 57]. The Rowhammer attack requires fast access to the same DRAM cells by bypassing the CPU cache, e.g., using `clflush` [26]. Additionally, cache eviction based on an eviction set can also result in access to DRAM cells when `clflush` is not available [4, 17]. Efficiently building eviction sets may thus also enhance Rowhammer attacks. For a successful Rowhammer attack, it is essential to collocate multiple memory pages within the same bank and adjacent to each other. A number of physical address bits, depending on the hardware configuration, are used to map memory pages to banks [42]. Since the rows are generally placed sequentially within the banks, access to adjacent rows within the same bank can be achieved if we have access to contiguous physical pages.




Figure 1: The Memory Order Buffer includes circular buffers SDB, SAB and LB. SDB, SAB and PAB of the data cache (DCACHE) have the same number of entries. SAB may initially hold the virtual address and the partial physical address. MOB requests the TLB to translate the virtual address and update the PAB with the translated physical address.

2.5 Memory Order Buffer

The processor manages memory operations using the Memory Order Buffer (MOB). MOB is tightly coupled with the data cache. The MOB assures that memory operations are executed efficiently by following the Intel memory ordering rule [36] in which memory stores are executed in-order and memory loads can be executed out-of-order. These rules have been enforced to improve the efficiency of memory accesses, while guaranteeing their correct commitment. Figure 1 shows the MOB schematic according to Intel [1, 2]. The MOB includes circular buffers, *store buffer*¹ and *load buffer*. A store will be decoded into two micro ops to store the address and data, respectively, into to the store buffer. The store buffer enables the processor to continue executing other instructions before commitment of the stores. As a result, the pipeline does not have to stall for the stores to complete. This further enables the MOB to support out-of-order execution of the load.

Store forwarding is an optimization mechanism that sends the store data to a load if the load address matches any of the store buffer entries. This is a speculative process, since the MOB cannot determine the true dependency of the load on stores based on the store buffer. Intel’s implementation of the store buffer is undocumented, but a potential design suggests that it will only hold the virtual address, and it may include part of the physical address [1, 2, 28]. As a result, the processor may falsely forward the data, although the physical addresses do not match. The complete resolution will be delayed until the load commitment, since the MOB needs to ask the TLB for the complete physical address information, which is time consuming. Additionally, The data cache may hold the translated store addresses in a Physical Address Buffer (PAB) with equal number of entries as the store buffer.

¹Store buffer consists of *Store Address Buffer (SAB)* and *Store Data Buffer (SDB)*. For simplicity, we use *Store Buffer* to mention the logically combined SAB and SDB units.

3 Speculative Load Hazards

As we mentioned earlier, memory loads can be executed out-of-order and before the preceding memory stores. If one of the preceding stores modifies the content of a location in memory, the memory load address is referring to, out-of-order execution of the load will operate on stale data, which results in invalid execution of a program. This out-of-order execution of the memory load is speculative behavior, since there is no guarantee during the execution time of the load that the virtual addresses corresponding to the memory stores do not conflict with the load address after translation to physical addresses. Figure 2 demonstrates this effect on a hypothetical processor with 7 pipeline stages. As multiple stores may be blocked due to limited resources, the execution of the load and dependent instructions in the pipeline, the *load block*, will bypass the stores since the MOB assumes the load block to be independent of the stores. This speculative behavior improves the memory bottleneck by letting other instructions continue their execution. However, if the dependency of the load and preceding stores is not verified, the load block may be computed on incorrect data which is either falsely forwarded by store forwarding (false dependency), or loaded from a stale cache line (unresolved true dependency). If the processor detects a false dependency before committing the load, it has to flush the pipeline and re-execute the load block. This will cause observable performance penalties and timing behavior.

3.1 Dependency Resolution

Dependency checks and resolution occur in multiple stages depending on the availability of the address information in the store buffer. A load instruction needs to be checked against all preceding stores in the store buffer to avoid false dependencies and to ensure the correctness of the data. A potential design [18, 28]², suggests the following stages for the dependency check and resolution, as shown in Figure 3:

1. **Loosenet:** The first stage is the *loosenet* check where the page offsets of the load and stores are compared³. In case of a loosenet hit, the compared load and store may be dependent and the processor will proceed to the next check stage.
2. **Finenet:** The next stage, called *finenet*, uses upper address bits. The *finenet* can be implemented to check the upper virtual address bits [18], or the physical address tag [28]. Either way, it is an intermediate stage, and it is

²The implementation of the MOB used in Intel processors is unpublished and therefore we cannot be certain about the precise architecture. Our results agree with some of the possible designs that are described in the Intel patents

³According to `Ld_Blocks_Partial:Address_Alias` Hardware Performance Counter (HPC) event [21], *loosenet* is defined by Intel as the mechanism that only compare the page offsets.

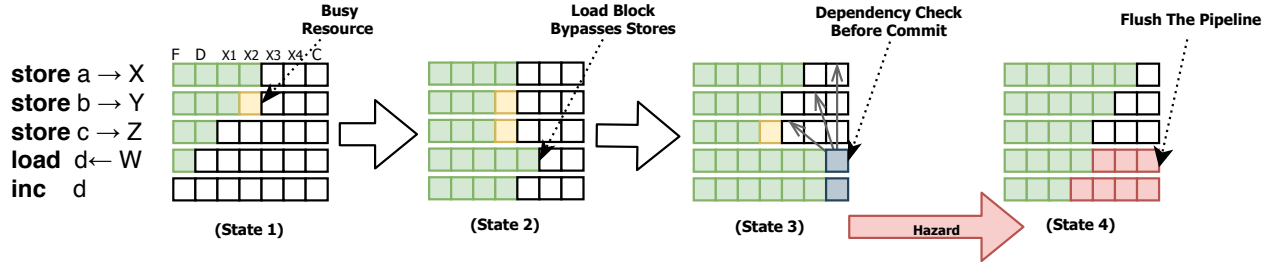


Figure 2: The speculative load is demonstrated on a hypothetical processor with pipeline stages: F = Fetch, D = Decode, X_{1-4} = Executions, and C = Commit. When the memory stores are blocked competing for resources (State 1), the load will bypass the stores (State 2). The load block including the dependent instructions will not be committed until the dependency of the address W versus X, Y, Z are resolved (State 3). In case of a dependency hazard (State 4), the pipeline is flushed and the load is restarted.




Figure 3: The dependency check logic: *loosenet* initially checks the least 12 significant bits (page offset) and the *finenet* checks the upper address bits, related to the page number. The final dependency using the physical address matching might still fail due to partial physical address checks.

not the final dependency resolution. In case of a *finenet* hit, the processor blocks the load and/or forwards the store data, otherwise, the dependency resolution will go into the final stage.

- Physical Address Matching:** At the final stage, the physical addresses will be checked. Since this stage is the final chance to resolve potential false dependencies, we expect the full physical address to be checked. However, one possible design suggests that if the physical addresses are not available, the physical address matching returns true and continues with the store forwarding [18].

Since the page offset is identical between the virtual and physical address, *loosenet* can be performed as soon as the store is decoded. [2] suggests that the store buffer only holds bit 19 to 12 of the physical address. Although the PAB holds the full translated physical address, it is not clear in which stage this information can be available to the MOB. As a result, the *finenet* check may be implemented based on checking the partial physical address bits. As we verify later, the dependency resolution logic may fail to resolve the dependency at multiple intermediate stages due to unavailability of the full physical address.

4 The SPOILER Attack

As described in Section 3, speculative loads may face other aliasing conditions in addition to the 4K aliasing, due to the partial checks on the higher address bits. To confirm this, we design an experiment to observe timing behavior of a speculative load based on higher address bits. For this purpose, we propose Algorithm 1 that executes a speculative load after multiple stores and further make sure to fill the store buffer with addresses that cause 4K aliasing during the execution of the load. Having w as the window size, the algorithm iterates over a number of different memory pages, and for each page, it performs stores to that page and all previous w pages within a window. Since the size of the store buffer varies between different processor generations, we choose a big enough window ($w = 64$) to ensure that the load has 4K aliasing with the maximum number of entries in the store buffer and hence maximum potential conflicts. Following the stores, we measure the timing of a load operation from a different memory page, as defined by x . The time measurement is performed using the code in Listing 1. Since we want the load to be executed speculatively, we can not use a store fence such as `mfence` before the load. As a result, our measurements are an estimate of execution time for the *speculatively load* and nearby microarchitectural events. This may include a negligible portion of overhead for the execution of stores, and/or any delay due to the dependency resolution. If we iterate over a diverse set of addresses with different virtual and physical page numbers, but the same page offset, we should be able to monitor any discrepancy.

4.1 Speculative Dependency Analysis

In this section, we use Algorithm 1 and Hardware Performance Counters (HPC) to perform an empirical analysis of the dependency resolution logic. HPCs can keep track of low-level hardware-related events in the CPU. The counters are accessible via special purpose registers and can be used to analyze the performance of a program. They provide a

Algorithm 1 Address Aliasing

```
for  $p$  from  $w$  to  $\text{PAGE\_COUNT}$  do
  for  $i$  from  $w$  to 0 do
     $\text{data} \xrightarrow{\text{store}} \text{buffer}[(p-i) \times \text{PAGE\_SIZE}]$ 
  end for
   $t_1 = \text{rdtscp}()$ 
   $\text{data} \xleftarrow{\text{load}} \text{buffer}[x \times \text{PAGE\_SIZE}]$ 
   $t_2 = \text{rdtscp}()$ 
   $\text{measure}[p] \leftarrow t_2 - t_1$ 
end for
return  $\text{measure}$ 
```

```
rdtscp;
mov %eax, %esi;
mov (%rbx), %eax;
rdtscp;
mfence;
sub %esi, %eax;
```

Listing 1: Timing measurement of a speculative load.

powerful tool to detect microarchitectural components that cause bottlenecks. Software libraries such as Performance Application Programming Interface (PAPI) [48] simplifies programming and reading low-level HPC on Intel processors. Initially, we execute [Algorithm 1](#) for 1000 different virtual pages. [Figure 4](#) (black) shows the cycle count for each iteration with a set of 4 kB aliased store addresses. Interestingly, we observe multiple step-wise peaks with a very high latency. Then, we use PAPI to monitor 30 different performance counters listed in [Table 5](#) in appendix while running the same experiment. At each iteration, only one performance counter is monitored alongside the aforementioned timing measurement. After each speculative load, the performance counter value and the load time are both recorded. Finally, we obtain a list of the timings and performance counter value pairs.

To find any relation between the observed high latency and a particular event, we compute correlation coefficients between counters and the timing measurements. Since the latency only occurs in the small region of the trace where the timing increases, we only need to compute the correlation on these regions. When an increase of at least 200 clock cycles is detected, the next s values from timing and the HPC traces are used to calculate the correlations, where s is the number of steps from [Table 1](#).

As shown in [Figure 5](#), two events have a high correlation with the leakage: `Cycle_Activity:Stalls_Ldm_Pending` has the highest correlation of 0.985. This event shows the number of cycles for which the execution is stalled and no instructions are executed due to a pending load. `Ld_Blocks_Partial:Address_Alias` has an inverse correlation with the leakage ([Figure 4](#), red). This event counts




Figure 4: Step-wise peaks (black) with a very high latency can be observed on some of the virtual pages. Affected HPC events, `Cycle_Activity:Stalls_Ldm_Pending` (blue) and `Ld_Blocks_Partial:Address_Alias` (red).




Figure 5: Correlation with HPCs listed in [Table 5](#) of the appendix. `Ld_Blocks_Partial:Address_Alias` and `Cycle_Activity:Stalls_Ldm_Pending` (both in red) have strong positive and negative correlations, respectively.

the number of false dependencies in the MOB when loosenet resolves the 4K aliasing condition. Separately, `Exe_Activity:Bound_on_Stores` increases with more number of stores within the inner window loop in [Listing 1](#), but it does not have a correlation with the leakage. The reason behind this behavior is that the store buffer is full, and additional store operations are pending. However, since there is no correlation with the leakage, this shows that the timing behavior is not due to the stores delay. We also attempt to profile any existing counters related to the memory disambiguation. However, the events `Memory_Disambiguation.Success` and `Memory_Disambiguation.Reset` are not available on the modern architectures that are tested.

4.2 Leakage of the Physical Address Mapping

In this experiment, we evaluate whether the observed step-wise latency has any relationship with the physical page numbers by observing the pagemap file. As shown in [Figure 4](#) (black), the average execution time for a load is 200 cycles for each set of addresses in the store buffer. Among multiple executions of the same experiment, we can always observe step-wise peaks with a very high latency which appear once

CPU Model	Architecture	Steps	SB Size
Intel Core i7-8650U	Kaby Lake R	22	56
Intel Core i7-7700	Kaby Lake	22	56
Intel Core i5-6440HQ	Skylake	22	56
Intel Xeon E5-2640v3	Haswell	17	42
Intel Xeon E5-2670v2	Ivy Bridge EP	14	36
Intel Core i7-3770	Ivy Bridge	12	36
Intel Core i7-2670QM	Sandy Bridge	12	36
Intel Core i5-2400	Sandy Bridge	12	36
Intel Core i5 650	Nehalem	11	32
Intel Core2Duo T9400	Core	N/A	20
Qualcomm Kryo 280	ARMv8-A	N/A	*
AMD A6-4455M	Bulldozer	N/A	*

Table 1: 1 MB aliasing on various architectures: The tested AMD and ARM architectures, and Intel Core generation do not show similar effects. The Store Buffer (SB) sizes are gathered from Intel Manual [21] and *wikichip.org* [54–56].

in every 256 pages on average. The least 20 bits of physical address for the `load` address matches with the physical addresses of performing `stores` to virtual pages with the highest peak. Multiple repetitions of this experiment show that we always detect peaks with different virtual addresses, which have the matching least 20 bits of physical address. This observation clearly discovers the existence of 1 MB aliasing effect based on the physical addresses. This 1 MB aliasing leaks information about 8 bits of mapping that were unknown to the user space processes.

Matching this observation with the previously observed `Cycle_Activity:Stalls_Ldm_Pending` with a high correlation, the speculative load has been stalled to resolve the dependency with conflicting store buffer entries after the occurrence of a 1 MB aliased address. This observation verifies that the latency is due to the pending `load`. When the latency is at the highest point, `Ld_Blocks_Partial:Address_Alias` drops to zero, and it increments at each down step of the peak. This implies that the `loosenet` check does not resolve the rest of the store dependencies whenever there is a 1 MB aliased address in the store buffer.

4.3 Evaluation

In the previous experiment, the execution time of the `load` operation that is delayed by 1 MB aliasing decreases gradually in each iteration (Figure 6). The number of steps to reach the normal execution time is consistent among multiple experiments on the same processor. When the first store in the window loop accesses a memory address with the matching 1 MB aliased address, the latency is at its highest point, marked as “1” in Figure 6. As the window loop accesses this address later in the loop, it appears closer to the `load` with a lower latency like the steps marked as 5, 15 and 22. This observation matches the *carry chain algorithm* described by

Intel [18] where the aliasing check starts from the most recent `store`. As shown in Table 1, experimenting with various processor generations shows that the number of steps has a linear correlation with the size of the store buffer which is architecture dependent. While the leakage exists on all Intel Core processors starting from the first generation, the timing effect is higher for the more recent generations with a bigger store buffer size. The analyzed ARM and AMD processors do not show similar behaviour⁴.

As our time measurement for speculative load suggests, it is not possible to reason whether the high timing is due to a very slow `load` or commitment of store operations. If the step-wise delay matches the store buffer entries, this delay may be either due to the the dependency resolution logic performing a pipeline flush and restart of the `load` for each 4 kB aliased entry starting from the 1 MB aliased entry, or due to the `load` waiting for all the remaining `stores` to commit because of an unresolved hazard. To explore this further, we perform an additional experiment with all store addresses replaced with non-aliased addresses except for one. This experiment shows that the peak disappears if there is only a single 4 kB and 1 MB aliased address in the store buffer.

Lastly, we run the same experiments on a shuffled set of virtual addresses to assure that the contiguous virtual addresses may not affect the observed leakage. Our experiment with the shuffled virtual addresses exactly match the same step-wise behavior suggesting that the upper bits in virtual addresses do not affect the leakage behavior, and the leakage is solely due to the aliasing on physical address bits.

4.3.1 Comparison of Address Aliasing Scenarios

We further test other address combinations to compare additional address aliasing scenarios using Algorithm 1. As shown by Figure 7, when `stores` and the `load` access different cache sets without aliasing, the `load` is executed in 30 cycles, which is the typical timing for an L1 data cache load. When the `stores` have different memory addresses with the same page offset, but the `load` has a different offset, the `load` takes 100 cycles to execute. This shows that even memory addresses in the store buffer having 4K Aliasing conditions with each other that are totally unrelated to the speculative load create a memory bottleneck for the `load`. In the next scenario, 4K aliasing between the `load` and all `stores`, the average load time is about 200 cycles. While the aforementioned 4K aliasing scenarios may leak cross domain information about memory accesses (Section 7), the most interesting scenario is the 1 MB aliasing which takes more than 1200 cycles for the highest point in the peak. For simplicity, we refer to the 1 MB aliased address as *aliased address*, in the rest of the paper.

⁴We use `rdtscp` for Intel and AMD processors and the `clock_gettime` for ARM processors to perform the time measurements.




Figure 6: Step-wise peaks with 22 steps and a high latency can be observed on some of the pages (*Core i7-8650U* processor).




Figure 7: Histogram of the measurement for the speculative load with various store addresses. Load will be fast, 30 cycles, without any dependency. If there exists 4K aliasing only between the stores, the average is 100. The average is 200 when there is 4K aliasing of load and stores. The 1 MB aliasing has a distinctive high latency.

4.4 Discussion

4.4.1 The Curious Case of Memory Disambiguation

The processor uses an additional speculative engine, called the *memory disambiguator* [9, 29], to predict memory false dependencies and reduce the chance of their occurrences. The main idea is to predict if a load is independent of preceding stores and proceed with the execution of the load by ignoring the store buffer. The predictor uses a hash table that is indexed with the address of the load, and each entry of the hash table has a saturating counter. If the pre-commitment dependency resolution does not detect false dependencies, the counter is incremented, otherwise it will be reset to zero. After multiple successful executions of the same load instruction, the predictor assumes that the load is safe to execute. Every time the counter resets to zero, the next iteration of the load will be blocked to be checked against the store buffer entries. Mispredictions result in performance overhead due to pipeline flushes. To avoid repeated mispredictions, a watchdog mechanism monitors the success rate of the prediction, and it can temporarily disable the memory disambiguator.

In the experiment described above, the predictor of the memory disambiguator should go into a stable state after the first few iterations, since the memory load is always truly independent of any aliased store. Hence the saturating counter for the target speculative load address passes the threshold, and it never resets due to a false prediction. As a result, the memory disambiguator should always fetch the data into the cache without any access to the store buffer. However, since the memory disambiguation performs speculation, the dependency resolution at some point verifies the prediction. The misprediction watchdog is also supposed to only disable the memory disambiguator when the misprediction rate is high, but in this case we should have a high prediction rate. Accordingly, the observed leakage occurs after the disambiguation and during the last stages of dependency resolution, i.e., the memory disambiguator only performs prediction on the 4K aliasing at the initial loosenet check, and it cannot protect the pipeline from 1 MB aliasing that appears at a later stage.

4.4.2 Hyperthreading Effect

Similar to the 4K Aliasing [37, 46], we empirically test whether the 1 MB aliasing can be used as a covert/side channel through logical processors. Our observation shows that when we run our experiments on two logical processors on the same physical core, the number of steps in the peaks is exactly halved. This matches the description by Intel [21] where it is stated that the store buffer is split between the logical processors. As a result, the 1 MB aliasing effect is not visible and exploitable across logical cores. [28] suggests that loosenet checks mask out the stores on the opposite thread.

5 SPOILER from JavaScript

Microarchitectural attacks from JavaScript have a high impact as drive-by attacks in the browser can be accomplished without any privilege or physical proximity. In such attacks, collocation is automatically granted by the fact that the browser loads a website with malicious embedded JavaScript code.




Figure 8: Reverse engineering physical page mappings in JavaScript. The markers point to addresses having same 20 bits of physical addresses being part of the same eviction set.

The browsers provide a sandbox where some instructions like `clflush` and `prefetch` and file systems such as `procfs` are inaccessible, limiting the opportunity for attack. Genkin et al. [12] showed that side-channel attacks inside a browser can be performed more efficiently and with greater portability through the use of *WebAssembly*. Yet, *WebAssembly* introduces an additional abstraction layer, i.e. it emulates a 32-bit environment that translates the internal addresses to virtual addresses of the host process (the browser). *WebAssembly* only uses addresses of the emulated environment and similar to JavaScript, it does not have direct access to the virtual addresses. Using *SPOILER* from JavaScript opens the opportunity to puncture these abstraction layers and to obtain physical address information directly. Figure 8 shows the address search in JavaScript using *SPOILER*. Compared to native implementations, we replace the `rdtscp` measurement with a timer, based on a shared array buffer. We cannot use any fence instruction such as `lfence`, and as a result, there remains some negligible noise in the JavaScript implementation. However, the aliased addresses can still be clearly seen, and we can use this information to improve the state-of-the-art eviction set creation for both Rowhammer and cache attacks.

5.1 Efficient Eviction Set Finding

We use the algorithm proposed in [12]. It is a slight improvement to the former state-of-the-art brute force method [39] and consists of three phases:

- *expand*: A large pool of addresses P is allocated with the last twelve bits of all addresses being zero. A random address is picked as a witness t and tested against a candidate set C . If t is not evicted by C , it is added to C and a new witness will be picked. As soon as t gets evicted by C , C forms an eviction set for t .
- *contract*: Addresses are subsequently removed from the eviction set. If the set still evicts t , the next address is removed. If it does not evict t anymore, the removed address is added back to the eviction set. At the end of

this phase, we have a minimal eviction set of the size of the set associativity.

- *collect*: All addresses mapping to the already found eviction set are removed from P by testing if they are evicted by the found set. After finding 128 initial cache sets, this approach utilizes the linearity property of the cache: For each found eviction set, the bits 6-11 are enumerated instead. This provides 63 more eviction sets for each found set, leading to full cache coverage.

We test this approach on an Intel Core i7-4770 with four physical cores and a shared 8MB 16-way L3 cache. The approach yields an 80% accuracy rate to find all 8192 eviction sets when starting with a pool of 4096 pages. The entire eviction set creation process takes an average of 46s. We improve the algorithm by 1) using the addresses removed from the eviction set in the contract phase as a new candidate set and 2) removing more than one address at a time from the eviction set during the contract phase. The improved eviction set creation process takes 35s on average.

As we described above, we can use *SPOILER* to find addresses which have the same least 20 bits of the physical address. Therefore, we do not have to check a large address pool and just the aliased addresses detected by *SPOILER* can be used. To achieve this, we start again with a pool of addresses where the offset bits are identical. We then perform a aliased address search to form a new pool of addresses where the last twenty address bits match. The eviction set finding algorithm is then carried out on this address pool.

5.1.1 Evaluation

The probability of finding a congruent address is $P(C) = 2^{\gamma-c-s}$, where c is the number of bits determining the cache set, γ is the number of bits attackers know, and s is the number of slices [53]. Since *SPOILER* allows us to control $\gamma \geq c$ bits, we are only left with uncertainty about a few address bits that influence the slice selection algorithm [23]. In theory, the eviction set search is sped up by a factor of 4096 by using aliased addresses in the pool, since on average one of 2^8 instead of one of 2^{20} addresses is a aliased address. Additionally, the address pool is much smaller, where 115 addresses are enough to find all the eviction sets. The resulting time gain is however smaller than the theoretical speed up, since the searching process for aliased addresses also takes a good amount of time. Table 2 shows the results. From each aliased address pool, 4 eviction sets can be found (corresponding to the 4 slices which are the only unknown part in the mapping). These can be enumerated again to form 63 more eviction sets since we still kept the bits 6-11 fixed. To accomplish full cache coverage, the aliased address pool has to be constructed 32 times. The *SPOILER* variant for finding eviction sets is more susceptible to system noise, which is why it needs more repetition (rounds) to get reliable values. On the other hand, it

Algorithm	Rounds	t_{total}	t_{AAS}	t_{ESS}	Success
Classic	3	46s	-	100%	80%
Improved	3	35s	-	100%	80%
AA	10	10s	54%	46%	67%
AA	20	12s	75%	25%	100%

Table 2: Comparison of different eviction set finding algorithms on an Intel Core i7-4770. Classic is the method from [12], improved is the same method with small improvements, *Aliased Address* (AA) uses SPOILER. t_{AAS} is the time percentage used for finding aliased addresses. t_{ESS} is the time percentage used for finding eviction sets.

is less prone to values deviating largely from the mean, which is a problem in the classic eviction set creation algorithm, that also does not succeed about one out of five times in our experiments. The unsuccessful attempts occur due to aborts if the algorithm takes much longer than statistically expected.

6 Rowhammer Attack using SPOILER

To perform a Rowhammer attack, the adversary needs to efficiently access DRAM rows adjacent to a victim row. In a single-sided Rowhammer attack, only one row is activated repeatedly to induce bit flips on one of the nearby rows. For this purpose, the attacker needs to make sure that multiple virtual pages co-locate on the same bank. The probability of co-locating on the same bank is low without the knowledge of physical addresses and their mapping to memory banks. In a double-sided Rowhammer attack, the attacker tries to access two different rows $n + 1$ and $n - 1$ to induce bit flips in the row n placed between them. While double-sided Rowhammer attacks induce bit flips faster due to the extra charge on the nearby cells of the victim row n , they further require access to contiguous memory pages. In this section, we show that SPOILER can help boosting both single and double-sided Rowhammer attacks by its additional 8-bit physical address information and resulting detection of contiguous memory.

6.1 DRAM Bank Co-location

DRAMA [42] reverse engineered the memory controller mapping. This requires elevated privileges to access physical addresses from the `pagemap` file. The authors have suggested that prefetch side-channel attacks [16] may be used to gain physical address information instead. SPOILER is an alternative way to obtain partial address information and is still feasible when the `prefetch` instruction is not available, e.g. in JavaScript. In our approach, we use SPOILER to detect aliased virtual memory addresses where the 20 LSBs of the physical addresses match. The memory controller uses these bits for mapping the physical addresses to the DRAM banks [42]. Even though the memory controller may use additional bits,

System Model	DRAM Configuration	# of Bits
Dell XPS-L702x (Sandy Bridge)	1 x (4GB 2Rx8)	21
	2 x (4GB 2Rx8)	22
Dell Inspiron-580 (Nehalem)	1 x (2GB 2Rx8) (b)	21
	2 x (2GB 2Rx8) (c)	22
	4 x (2GB 2Rx8) (d)	23
Dell Optiplex-7010 (Ivy Bridge)	1 x (2GB 1Rx8) (a)	19
	2 x (2GB 1Rx8)	20
	1 x (4GB 2Rx8) (e)	21
	2 x (4GB 2Rx8)	22

Table 3: Reverse engineering the DRAM memory mappings using DRAMA tool, # of Bits represents the number of physical address bits used for the bank, rank and channel [42].

the majority of the bits are known using SPOILER. An attacker can directly hammer such aliased addresses to perform a more efficient single-sided Rowhammer attack with a significantly increased probability of hitting the same bank. As shown in Table 3, we reverse engineer the DRAM mappings for different hardware configurations using the DRAMA tool, and only a few bits of physical address entropy beyond the 20 bits will remain unknown.

To verify if our aliased virtual addresses co-locate on the same bank, we use the row-conflict side channel as proposed in [11] (timings in the appendix, Section 10.2). We observe that whenever the number of physical address bits used by the memory controller to map data to physical memory is equal to or less than 20, we always hit the same bank. For each additional bit the memory controller uses, the probability of hitting the same bank is divided by 2 as there is one more bit of entropy. In general, we can formulate that our probability p to hit the same bank is $p = 1/2^n$, where n is the number of unknown physical address bits in the mapping. We experimentally verify the success rate for the setups listed in Table 3, as depicted in Figure 9.

In summary, SPOILER drastically improves the efficiency of finding addresses mapping to the same bank without administrative privilege or reverse engineering the memory controller mapping. Note that this approach works even in sandboxed environments such as JavaScript.

6.2 Contiguous Memory

For a double-sided Rowhammer attack, we need to hammer rows adjacent to the victim row in the same bank. This requires detecting contiguous memory pages in the allocated memory, since the rows are written to the banks sequentially. Without contiguous memory, the banks will be filled randomly and we will not be able to locate neighboring rows. We show that an attacker can use SPOILER to detect contiguous memory using 1 MB aliasing peaks. For this purpose, we compare the physical frame numbers to the SPOILER leakage for 10000




Figure 9: Bank co-location for various DRAM configurations (a), (b), (c) & (d) from Table 3. All aliased addresses, in the first plot, result in high access times and row conflicts for (a), where the 19 bits used by the memory controller collide. The remaining graphs depict the same experiments on configurations (b), (c) & (d) where an increasing number of bits are unknown. The regularity of the peaks shows that the allocated memory was contiguous, which is coincidental.

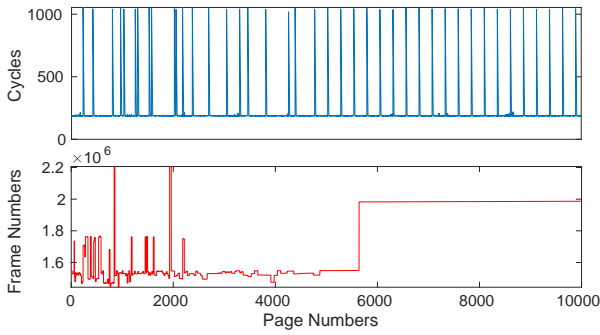


Figure 10: Relation between leakage peaks and the physical page numbers. The blue plot shows the leakage peaks from SPOILER. The red plot shows the decimal values of the physical frame numbers from the pagemap file. Once the peaks in the blue plot become regular, the red plot is linearly increasing, which shows contiguous memory allocation.

different virtual pages allocated using `malloc`. Figure 10 shows the relation between 1 MB aliasing peaks and physical page frame numbers. When the distance between the peaks is random, the trend of frame numbers also change randomly. After around 5000 pages, we observe that the frame numbers increase sequentially. The number of pages between the peaks remains constant at 256 where this distance comes from the 8 bits of physical address leakage due to 1 MB aliasing.

We also compare the accuracy of obtaining contiguous memory detected by SPOILER by analyzing the actual physical addresses from the pagemap file. By checking the difference between physical page numbers for each detected virtual page, we can determine the accuracy of our detection method: the success rate for finding contiguous memory is above 99% disregarding the availability of the contiguous pages. For de-

DRAM Model	Architecture	Flippy
M378B5273DH0-CK0	Ivy Bridge	✓
M378B5273DH0-CK0	Sandy Bridge	✓
M378B5773DH0-CH9	Sandy Bridge	✓
M378B5173EB0-CK0	Sandy Bridge	×
NT2GC64B88G0NF-CG	Sandy Bridge	×
KY996D-ELD	Sandy Bridge	×
M378B5773DH0-CH9	Nehalem	✓
NT4GC64B8HG0NS-CG	Sandy Bridge	×
HMA41GS6AFR8N-TF	Skylake	×

Table 4: DRAM Modules Susceptible to double-sided Rowhammer attacks using SPOILER




Figure 11: Amount of bit flips increases with the increase in number of hammerings. The timings do not include the time taken for reboots and 1 minute sleep time.

tailed experiment on the availability of the contiguous pages, see Section 10.3 in the appendix.

6.3 Double-Sided Rowhammer with SPOILER

As double-sided Rowhammer attacks are based on the assumption that rows within a bank are contiguous and based on the contiguous memory detection using SPOILER, we mount a practical double-sided Rowhammer attack on several DRAM modules without any root privileges. First, we use SPOILER to detect a suitable amount of contiguous memory. If enough contiguous memory is available in the system, SPOILER finds it, otherwise a double-sided Rowhammer attack will not be applicable. Second, we apply the row conflict side channel only to the located contiguous memory, and get a list of virtual addresses which are contiguously mapped within a bank. Finally, we start performing a double-sided Rowhammer attack by picking 3 consecutive addresses from our list. In our experiment, we empirically configure our Rowhammer attack combined with SPOILER to detect 10 MB of contiguous memory. As Table 4 shows, we could get bit flips on some of the available configurations with potential low quality cells.

Figure 11 shows the number of hammers compared to the amount of bit flips for configuration (e) in Table 3. We repeat this experiment 30 times for every measurement and

the results are then averaged out. On every experiment, the system is rebooted using a script because once the memory becomes fragmented, no more contiguous memory is available. The number of bit flips increases with more number of hammerings. Hammering for 500 million times is found to be an optimal number for this DRAM configuration, as the continuation of hammering is not increasing bit flips.

7 Tracking Speculative Loads With SPOILER

Single-threaded attacks can be used to steal information from other security contexts running before/after the attacker code on the same thread [8, 38]. Example scenarios are I) context switches between processes of different users, or II) between a user process and a kernel thread, and III) Intel Software Guard eXtensions (SGX) secure enclaves [38, 51]. In such attacks, the adversary put the microarchitecture to a particular state, waits for the context switch and execution of the victim thread, and then tries to observe the microarchitectural state after the victim’s execution. We propose an attack where the adversary 1) fills the store buffer with arbitrary addresses, 2) issues the victim context switch and lets the victim perform a secret-dependent memory access, and 3) measures the execution time of the victim. Any correlation between the victim’s timing and the load address can leak secrets [58]. Due to the nature of SPOILER, the victim should access the memory while there are aliased addresses in the store buffer, i.e. if the stores are committed before the victim’s speculative load, there will be no dependency resolution hazard.

We first perform an analysis of the depth of the operations that can be executed between the stores and the load to investigate the viability of SPOILER. In this experiment, we repeat a number of instructions between stores and the load that are free from memory operations. Figure 12 shows the number of stall steps due to the dependency hazard with the added instructions. Although nop is not supposed to take any cycle, adding 4000 nop will diffuse the timing latency. Then, we test add and leal, which use the Arithmetic Logic Unit (ALU) and the Address Generation Unit (AGU), respectively. Figure 12 shows that only 1000 adds can be executed between the stores and load before the SPOILER effect is lost. Since each add typically takes about 1 cycle to execute, this roughly gives a 1000 cycle depth for SPOILER. Considering the observed depth, we discuss potential attacks that can track the speculative load in the following two scenarios.

7.1 SPOILER Context Switch

In this attack, we are interested in tracking a memory access in the privileged kernel environment after a context switch. First, we fill the store buffer with addresses that have the same page offset, and then execute a system call. During the execution of the system call, we expect to observe a delayed execution if a secret load address has aliasing with the stores. We




Figure 12: The depth of SPOILER leakage with respect to different instructions and execution units.




Figure 13: Execution time of mincore system call. When a kernel load address has aliasing with the attacker’s stores (red), the step-wise delay will appear. These timings are measured with Kernel Page Table Isolation disabled.

utilize SPOILER to iterate over various virtual pages, thus some of the pages have more noticeable latency due to the 1 MB aliasing. We analyze multiple syscalls with various execution times. For instance, Figure 13 shows the execution time for the mincore. In the first experiment (red), we fill the store buffer with addresses that have aliasing with a memory load operation in the kernel code space. The 1 MB aliasing delay with 7 steps suggests that we can track the address of a kernel memory load by the knowledge of our arbitrary filled store addresses. The blue line shows the timing when there is no aliasing between the target memory load and the attacker’s store. Surprisingly, only by filling the store buffer, the system call executes much slower: the normal execution time for mincore should be around 250 cycles (cyan). This proof of concept shows that SPOILER can be used to leak information from more privileged contexts, however this is limited only to loads that appear at the beginning of the next context.

7.2 Negative Result: SPOILER SGX

In this experiment, we try to combine SPOILER with the CacheZoom [38] approach to create a novel single-threaded side-channel attack against SGX enclaves with high temporal




Figure 14: The effect of SPOILER on TLB flush. The execution cycle always increases for 4 kB aliased addresses, except for some of the virtual pages inside in the the store buffer where we observe step-wise hills.

and spatial resolution (4-byte) [37]). We use *SGX-STEP* [50] to precisely interrupt every single instruction. *Nemesis* [51] shows that the interrupt handler context switch time is dependent on the execution time of the currently running instruction. On our test platform, *Core i7-8650U*, each context switch on an enclave takes about 12000 cycles to execute. If we fill the store buffer with memory addresses that match the page offset of a `load` inside the enclave in the interrupt handler, the context switch timing is increased to about 13500 cycles. While we cannot observe any correlation between the matched 4 kB or 1 MB aliased addresses, we do see unexpected periodic downward peaks with a similar step-wise behavior as SPOILER (Figure 14). We later reproduce a similar behavior by running SPOILER before an `ioctl` routine that flushes the TLB on each call. Intel SGX also performs an implicit TLB flush during each context switch. We can thus infer that the downward peaks occur due to the TLB flush, especially since the addresses for the downward peaks do not have any address correlation with the `load` address. This suggests that the TLB flush operation itself is affected by SPOILER. This effect eliminates the opportunity to observe any potential correlation due to the speculative load. As a result, we can not use SPOILER to track memory accesses inside an enclave. Further exploration of the root cause of the TLB flush effect can be carried out as a future work.

8 Mitigations

Software Mitigations The attack exploits the fact that when there is a `load` instruction after a number of `store` instructions, the physical address conflict causes a high timing behavior. This happens because of the speculatively executed `load` before all the `stores` are finished executing. There is no software mitigation that can completely erase this problem. While the timing behavior can be removed by inserting store fences between the `loads` and `stores`, this cannot be

enforced to the user’s code space, i.e., the user can always leak the physical address information. Another yet less robust approach is to execute other instructions between the `loads` and `stores` to decrease the depth of the attack. However, both of the approaches are only applicable to defend against attacks such as the one described in Section 7.

As for most attacks on JavaScript, removing accurate timers from the browser would be effective against SPOILER. Indeed, some timers have been removed or distorted by jitters as a response to attacks [32]. There is however a wide range of timers with varying precision available, and removing all of them seems impractical [11, 43].

Hardware Mitigations The hardware design for the memory disambiguator may be revised to prevent such physical address leakage, but modifying the speculative behavior may cause performance impacts. For instance, partial address comparison was a design choice for performance. Full address comparison may address this vulnerability, but will also impact performance. Moreover, hardware patches are difficult to be applied to legacy systems and take years to be deployed.

9 Conclusion

We introduced SPOILER, a novel approach for gaining physical address information by exploiting a new information leakage due to speculative execution. To exploit the leakage, we used the speculative load behavior after jamming the store buffer. SPOILER can be executed from user space and requires no special privileges. While speculative execution enables both SPOILER and Spectre and Meltdown, our newly found leakage stems from a completely different hardware unit, the Memory Order Buffer. We exploited the leakage to reveal information on the 8 least significant bits of the physical page number, which are critical for many microarchitectural attacks such as Rowhammer and cache attacks. We analyzed the causes of the discovered leakage in detail and showed how to exploit it to extract physical address information.

Further, we showed the impact of SPOILER by performing a highly targeted Rowhammer attack in a native user-level environment. We further demonstrated the applicability of SPOILER in sandboxed environments by constructing efficient eviction sets from JavaScript, an extremely restrictive environment that usually does not grant any access to physical addresses. Gaining even partial knowledge of the physical address will make new attack targets feasible in browsers even though JavaScript-enabled attacks are known to be difficult to realize in practice due to the limited nature of the JavaScript environment. Broadly put, the leakage described in this paper will enable attackers to perform existing attacks more efficiently, or to devise new attacks using the novel knowledge.

Responsible Disclosure We have informed the *Intel Product Security Incident Response Team* of our findings on December 1st, 2018, and they have acknowledged the receipt.

Acknowledgments

We thank Yuval Yarom for his valuable comments for improving the quality of this paper.

This work is supported by U.S. Department of State, Bureau of Educational and Cultural Affairs' Fulbright Program and National Science Foundation under grant CNS-1618837 and CNS-1814406.

References

- [1] Jeffery M Abramson, Haitham Akkary, Andrew F Glew, Glenn J Hinton, Kris G Konigsfeld, and Paul D Madland. Method and apparatus for performing a store operation, April 23 2002. US Patent 6,378,062.
- [2] Jeffrey M Abramson, Haitham Akkary, Andrew F Glew, Glenn J Hinton, Kris G Konigsfeld, Paul D Madland, David B Papworth, and Michael A Fetterman. Method and apparatus for dispatching and executing a load operation to memory, February 10 1998. US Patent 5,717,882.
- [3] agner. The microarchitecture of Intel, AMD and VIA CPUs: An optimization guide for assembly programmers and compiler makers.
- [4] Zelalem Birhanu Aweke, Salessawi Ferede Yitbarek, Rui Qiao, Reetuparna Das, Matthew Hicks, Yossi Oren, and Todd Austin. Anvil: Software-based protection against next-generation rowhammer attacks. *ACM SIGPLAN Notices*, 51(4):743–755, 2016.
- [5] Naomi Bengier, Joop van de Pol, Nigel P. Smart, and Yuval Yarom. “ooh aah... just a little bit” : A small amount of side channel can go a long way. In *Cryptographic Hardware and Embedded Systems – CHES 2014*, pages 75–92, Berlin, Heidelberg, 2014. Springer.
- [6] Daniel J Bernstein. Cache-timing attacks on aes, 2005.
- [7] Claudio Canella, Jo Van Bulck, Michael Schwarz, Moritz Lipp, Benjamin von Berg, Philipp Ortner, Frank Piessens, Dmitry Evtushkin, and Daniel Gruss. A systematic evaluation of transient execution attacks and defenses. *arXiv preprint arXiv:1811.05441*, 2018.
- [8] G. Chen, S. Chen, Y. Xiao, Y. Zhang, Z. Lin, and T. H. Lai. SgxPectre Attacks: Leaking Enclave Secrets via Speculative Execution. *ArXiv e-prints*, February 2018.
- [9] Jack Doweck. Inside intel® core microarchitecture. In *Hot Chips 18 Symposium (HCS), 2006 IEEE*, pages 1–35. IEEE, 2006.
- [10] Dmitry Evtushkin, Dmitry Ponomarev, and Nael Abu-Ghazaleh. Jump over aslr: Attacking branch predictors to bypass aslr. In *The 49th Annual IEEE/ACM International Symposium on Microarchitecture, MICRO-49*, pages 40:1–40:13, Piscataway, NJ, USA, 2016. IEEE Press.
- [11] Pietro Frigo, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. Grand pwning unit: Accelerating microarchitectural attacks with the gpu. In *Grand Pwning Unit: Accelerating Microarchitectural Attacks with the GPU*, page 0, Washington, DC, USA, 2018. IEEE, IEEE Computer Society.
- [12] Daniel Genkin, Lev Pachmanov, Eran Tromer, and Yuval Yarom. Drive-by key-extraction cache attacks from portable code. In Bart Preneel and Frederik Vercauteren, editors, *Applied Cryptography and Network Security*, pages 83–102, Cham, 2018. Springer International Publishing.
- [13] Google. speculative execution, variant 4: speculative store bypass . <https://bugs.chromium.org/p/project-zero/issues/detail?id=1528>. Accessed: 2018-11-23.
- [14] Mel Gorman. *Understanding the Linux Virtual Memory Manager*. Prentice Hall, London, 2004.
- [15] Daniel Gruss, Moritz Lipp, Michael Schwarz, Daniel Genkin, Jonas Juffinger, Sioli O’Connell, Wolfgang Schoecl, and Yuval Yarom. Another flip in the wall of rowhammer defenses. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 245–261. IEEE, 2018.
- [16] Daniel Gruss, Clémentine Maurice, Anders Fogh, Moritz Lipp, and Stefan Mangard. Prefetch side-channel attacks: Bypassing smap and kernel aslr. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, CCS ’16*, pages 368–379, New York, NY, USA, 2016. ACM.
- [17] Daniel Gruss, Clémentine Maurice, and Stefan Mangard. Rowhammer.js: A remote software-induced fault attack in javascript. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 300–321. Springer, 2016.
- [18] Sebastien Hily, Zhongying Zhang, and Per Hammarlund. Resolving false dependencies of speculative load instructions, October 13 2009. US Patent 7,603,527.
- [19] Ralf Hund, Carsten Willems, and Thorsten Holz. Practical timing side channel attacks against kernel space aslr. In *2013 IEEE Symposium on Security and Privacy*, pages 191–205. IEEE, 2013.
- [20] Mehmet Sinan İnci, Berk Gulmezoglu, Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. Cache attacks enable bulk key recovery on the cloud. In *Cryptographic Hardware and Embedded Systems – CHES 2016*, pages 368–388, Berlin, Heidelberg, 2016. Springer.
- [21] Intel. Intel® 64 and IA-32 Architectures Optimization Reference Manual.
- [22] Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. S\$A: A shared cache attack that works across cores and defies vm sandboxing – and its application to aes. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy, SP ’15*, pages 591–604, Washington, DC, USA, 2015. IEEE Computer Society.
- [23] Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. Systematic reverse engineering of cache slice selection in intel processors. In *2015 Euromicro Conference on Digital System Design (DSD)*, pages 629–636. IEEE, 2015.
- [24] Yeongjin Jang, Sangho Lee, and Taesoo Kim. Breaking kernel address space layout randomization with intel tsx. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, pages 380–392. ACM, 2016.
- [25] Vasileios P Kemerlis, Michalis Polychronakis, and Angelos D Keromytis. ret2dir: Rethinking kernel isolation. In *USENIX Security Symposium*, pages 957–972, 2014.
- [26] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping bits in memory without accessing them: An experimental study of dram disturbance errors. In *ACM SIGARCH Computer Architecture News*, volume 42, pages 361–372. IEEE Press, 2014.
- [27] Paul Kocher, Daniel Genkin, Daniel Gruss, Werner Haas, Mike Hamburg, Moritz Lipp, Stefan Mangard, Thomas Prescher, Michael Schwarz, and Yuval Yarom. Spectre attacks: Exploiting speculative execution. *ArXiv e-prints*, January 2018.
- [28] Steffen Kosinski, Fernando Latorre, Niranjana Cooray, Stanislaw Shwartsman, Ethan Kalifon, Varun Mohandru, Pedro Lopez, Tom Aviram-Rosenfeld, Jaroslav Topp, Li-Gao Zei, et al. Store forwarding for data caches, November 29 2016. US Patent 9,507,725.
- [29] Evgeni Krimer, Guillermo Savransky, Idan Mondjak, and Jacob Doweck. Counter-based memory disambiguation techniques for selectively predicting load/store conflicts, October 1 2013. US Patent 8,549,263.
- [30] Moritz Lipp, Daniel Gruss, Michael Schwarz, David Bidner, Clémentine Maurice, and Stefan Mangard. Practical keystroke timing attacks in sandboxed javascript. In *Computer Security – ESORICS 2017*, pages 191–209. Springer, 2017.
- [31] Moritz Lipp, Daniel Gruss, Raphael Spreitzer, Clémentine Maurice, and Stefan Mangard. Armageddon: Cache attacks on mobile devices. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 549–564, Austin, TX, 2016. USENIX Association.

- [32] Moritz Lipp, Michael Schwarz, Daniel Gruss, Thomas Prescher, Werner Haas, Anders Fogh, Jann Horn, Stefan Mangard, Paul Kocher, Daniel Genkin, Yuval Yarom, and Mike Hamburg. Meltdown: Reading kernel memory from user space. In *27th USENIX Security Symposium (USENIX Security 18)*, Baltimore, MD, 2018. USENIX Association.
- [33] Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B. Lee. Last-level cache side-channel attacks are practical. In *Proceedings of the 2015 IEEE Symposium on Security and Privacy*, SP '15, pages 605–622, Washington, DC, USA, 2015. IEEE Computer Society.
- [34] Errol L. Lloyd and Michael C. Loui. On the worst case performance of buddy systems. *Acta Informatica*, 22(4):451–473, Oct 1985.
- [35] Giorgi Maisuradze and Christian Rossow. ret2spec: Speculative execution using return stack buffers. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 2109–2122. ACM, 2018.
- [36] Intel 64 Architecture Memory Ordering White Paper. http://www.cs.cmu.edu/~410-f10/doc/Intel_Reordering_318147.pdf, 2008. Accessed: 2018-11-26.
- [37] Ahmad Moghimi, Thomas Eisenbarth, and Berk Sunar. Memjam: A false dependency attack against constant-time crypto implementations in SGX. In *Topics in Cryptology - CT-RSA 2018 - The Cryptographers' Track at the RSA Conference 2018, San Francisco, CA, USA, April 16-20, 2018, Proceedings*, pages 21–44, 2018.
- [38] Ahmad Moghimi, Gorka Irazoqui, and Thomas Eisenbarth. Cachezoom: How sgx amplifies the power of cache attacks. In *Cryptographic Hardware and Embedded Systems - CHES 2017*, pages 69–90. Springer, 2017.
- [39] Yossef Oren, Vasileios P. Kemerlis, Simha Sethumadhavan, and Angelos D. Keromytis. The spy in the sandbox: Practical cache attacks in javascript and their implications. In *Proceedings of the 22Nd ACM SIGSAC Conference on Computer and Communications Security*, CCS '15, pages 1406–1418, New York, NY, USA, 2015. ACM.
- [40] Colin Percival. Cache missing for fun and profit, 2005.
- [41] Cesar Pereida García, Billy Bob Brumley, and Yuval Yarom. "make sure dsa signing exponentiations really are constant-time". In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security*, CCS '16, pages 1639–1650, New York, NY, USA, 2016. ACM.
- [42] Peter Pessl, Daniel Gruss, Clémentine Maurice, Michael Schwarz, and Stefan Mangard. Drama: Exploiting dram addressing for cross-cpu attacks. In *USENIX Security Symposium*, pages 565–581, 2016.
- [43] Michael Schwarz, Clémentine Maurice, Daniel Gruss, and Stefan Mangard. Fantastic timers and where to find them: High-resolution microarchitectural attacks in javascript. In Aggelos Kiayias, editor, *Financial Cryptography and Data Security*, pages 247–267, Cham, 2017. Springer International Publishing.
- [44] Mark Seaborn and Thomas Dullien. Exploiting the dram rowhammer bug to gain kernel privileges. *Black Hat*, 15, 2015.
- [45] Julian Stecklina and Thomas Prescher. Lazyfp: Leaking fpu register state using microarchitectural side-channels. *arXiv preprint arXiv:1806.07480*, 2018.
- [46] Dean Sullivan, Orlando Arias, Travis Meade, and Yier Jin. Microarchitectural minefields: 4k-aliasing covert channel and multi-tenant detection in iaas clouds. In *Network and Distributed Systems Security (NDSS) Symposium*. The Internet Society, 2018.
- [47] Andrei Tatar, Radhesh Krishnan, Elias Athanasopoulos, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. Throwhammer: Rowhammer attacks over the network and defenses. In *2018 USENIX Annual Technical Conference (USENIX ATC 18)*, Boston, MA, 2018. USENIX Association.
- [48] Dan Terpstra, Heike Jagode, Haihang You, and Jack Dongarra. Collecting performance data with papi-c. In *Tools for High Performance Computing 2009*, pages 157–173. Springer, 2010.
- [49] Jo Van Bulck, Marina Minkin, Ofir Weisse, Daniel Genkin, Baris Kasikci, Frank Piessens, Mark Silberstein, Thomas F Wenisch, Yuval Yarom, and Raoul Strackx. Foreshadow: Extracting the keys to the intel sgx kingdom with transient out-of-order execution. In *Proceedings of the 27th USENIX Security Symposium*. USENIX Association, 2018.
- [50] Jo Van Bulck, Frank Piessens, and Raoul Strackx. Sgx-step: A practical attack framework for precise enclave execution control. In *Proceedings of the 2Nd Workshop on System Software for Trusted Execution*, SysTEX'17, pages 4:1–4:6, New York, NY, USA, 2017. ACM.
- [51] Jo Van Bulck, Frank Piessens, and Raoul Strackx. Nemesis: Studying microarchitectural timing leaks in rudimentary cpu interrupt logic. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 178–195. ACM, 2018.
- [52] Victor Van Der Veen, Yanick Fratantonio, Martina Lindorfer, Daniel Gruss, Clémentine Maurice, Giovanni Vigna, Herbert Bos, Kaveh Razavi, and Cristiano Giuffrida. Drammer: Deterministic rowhammer attacks on mobile platforms. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pages 1675–1689. ACM, 2016.
- [53] Pepe Vila, Boris Köpf, and José Francisco Morales. Theory and practice of finding eviction sets. *arXiv preprint arXiv:1810.01497*, 2018.
- [54] WikiChip. Ivy Bridge - Microarchitectures - Intel. [https://en.wikichip.org/wiki/intel/microarchitectures/ivy_bridge_\(client\)](https://en.wikichip.org/wiki/intel/microarchitectures/ivy_bridge_(client)). Accessed: 2019-02-05.
- [55] WikiChip. Kaby Lake - Microarchitectures - Intel. https://en.wikichip.org/wiki/intel/microarchitectures/kaby_lake. Accessed: 2019-02-05.
- [56] WikiChip. Skylake (client) - Microarchitectures - Intel. [https://en.wikichip.org/wiki/intel/microarchitectures/skylake_\(client\)](https://en.wikichip.org/wiki/intel/microarchitectures/skylake_(client)). Accessed: 2019-02-05.
- [57] Yuan Xiao, Xiaokuan Zhang, Yinqian Zhang, and Radu Teodorescu. One bit flips, one cloud flops: Cross-vm row hammer attacks and privilege escalation. In *USENIX Security Symposium*, pages 19–35, 2016.
- [58] Yuval Yarom, Daniel Genkin, and Nadia Heninger. CacheBleed: a timing attack on OpenSSL constant-time RSA. *Journal of Cryptographic Engineering*, 7(2):99–112, 2017.
- [59] Bennet Yee, David Sehr, Gregory Dardyk, J Bradley Chen, Robert Muth, Tavis Ormandy, Shiki Okasaka, Neha Narula, and Nicholas Fullagar. Native client: A sandbox for portable, untrusted x86 native code. In *Security and Privacy, 2009 30th IEEE Symposium on*, pages 79–93. IEEE, 2009.

10 Appendix

10.1 Tested Hardware Performance Counters

Counters	Correlation
UNHALTED_CORE_CYCLES	0.3077
UNHALTED_REFERENCE_CYCLES	0.1527
INSTRUCTION_RETIRED	0.2718
INSTRUCTIONS_RETIRED	0.2827
BRANCH_INSTRUCTIONS_RETIRED	0.3143
MISPREDICTED_BRANCH_RETIRED	0.0872
CYCLE_ACTIVITY:CYCLES_L2_PENDING	-0.0234
CYCLE_ACTIVITY:STALLS_LDM_PENDING	0.9819
CYCLE_ACTIVITY:CYCLES_NO_EXECUTE	0.2317
RESOURCE_STALLS:ROB	0
RESOURCE_STALLS:SB	-0.0506
RESOURCE_STALLS:RS	-0.0044
LD_BLOCKS_PARTIAL:ADDRESS_ALIAS	-0.9511
IDQ_UOPS_NOT_DELIVERED	-0.1455
IDQ:ALL_DSB_CYCLES_ANY_UOPS	0.0332
ILD_STALL:IQ_FULL	0.1021
ITLB_MISSES:MISS_CAUSES_A_WALK	0
TLB_FLUSH:STLB_THREAD	0
ICACHE:MISSES	0
ICACHE:IFETCH_STALL	0
LID:REPLACEMENT	0.3801
L2_DEMAND_RQSTS:WB_HIT	0.2436
LONGEST_LAT_CACHE:MISS	0.0633
CYCLE_ACTIVITY:CYCLES_LID_PENDING	-0.0080
LOCK_CYCLES:CACHE_LOCK_DURATION	0
LOAD_HIT_PRE:SW_PF	0
LOAD_HIT_PRE:HW_PF	0
MACHINE_CLEARS:CYCLES	0
OFFCORE_REQUESTS_BUFFER:SQ_FULL	0
OFFCORE_REQUESTS_DEMAND_DATA_RD	0.1765

Table 5: Counters profiled for correlation test

10.2 Row-conflict Side Channel

The row-conflict side channel retrieves the timing information of the CPU while doing direct accesses (using `clflush`) from the DRAM. A higher timing indicates that the two addresses are mapped to the same bank in the DRAM because reading an address from the same bank forces the rowbuffer to copy the previous contents back to the original row and then load the newly accessed data into the rowbuffer. Whereas, a low timings indicates that two addresses are not in the same bank (not sharing the same rowbuffer) and are loaded into separate rowbuffers. Figure 15 shows a wide gap (around 100 cycles) between row hits and row conflicts.

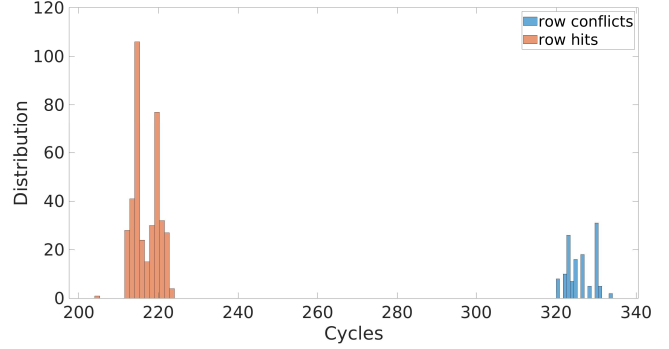


Figure 15: Timings for accessing the aliased virtual addresses (random addresses where 20 LSB of the physical address match). Row hits (orange) are clearly distinguishable from Row conflicts (blue).

10.3 Memory Utilization and Contiguity

The probability of obtaining contiguous memory changes depends on memory utilization of the system. We conduct an experiment to examine the effect of memory utilization on availability of contiguous memory. In this experiment, 1 GB memory is allocated. During the experiment, the memory utilization of the system is increased gradually from 20% to 90%. We measure the probability of getting the contiguous memory with two methods. The first one is checking the physical frame numbers from `pagemap` file to look for a 520 kB of contiguous memory. The second method is using `SPOILER` to find the 520 kB of contiguous memory. This 520 kB is required to get three consecutive rows with in a bank for a DRAM configuration having 256 kB row offset and 8 kB row size.

Figure 16 and Figure 17 show that when the memory has been fragmented after intense memory usage, it gets more difficult to allocate a contiguous chunk of memory required. Even decreasing the memory usage does not help to get a contiguous block of memory. Figure 17 depicts that after the memory utilization has been decreased from 70% to 60% and so on, there is not enough contiguous memory to mount a successful double-sided Rowhammer attack. Until the machine is restarted, the memory remains fragmented which makes a double-sided Rowhammer attack difficult, especially on targets like high-end servers where restarting is impractical.

The observed behavior can be explained by the *binary buddy allocator* which is responsible for the physical address allocation in the Linux OS [14]. This type of allocator is known to fragment memory significantly under certain circumstances [34]. The Linux OS uses a *SLAB/SLOB allocator* in order to circumvent the fragmentation problems. However, the allocator only serves the kernel directly. User space memory therefore still suffers from the fragmentation that the buddy allocator introduces. This also means that getting the contiguous memory required for a double-sided Rowhammer

attack becomes more difficult if the system under attack has been active for a while.




Figure 16: Finding Contiguous Memory of 520 kB with Increasing Memory Utilization. The overlap between the red and blue plot indicates the high accuracy of the contiguous memory detection capability of SPOILER as verified by the pagemap file.




Figure 17: Finding Contiguous Memory of 520 kB with Decreasing Memory Utilization.