International Journal of Parallel Programming (2019) 47:538–570
https://doi.org/10.1007/s10766-018-0611-9

MemJam: A False Dependency Attack Against Constant-Time Crypto Implementations

Ahmad Moghimi[1] · **Jan Wichelmann**[2] · **Thomas Eisenbarth**[2] · **Berk Sunar**[1]

Received: 15 February 2018 / Accepted: 29 October 2018 / Published online: 9 November 2018

Abstract

Cache attacks exploit memory access patterns of cryptographic implementations. Constant-time implementation techniques have become an indispensable tool in fighting cache timing attacks. These techniques engineer the memory accesses of cryptographic operations to follow a uniform key independent pattern. However, the constant-time behavior is dependent on the underlying architecture, which can be highly complex and often incorporates unpublished features. The *CacheBleed* attack targets cache bank conflicts and thereby invalidates the assumption that microarchitectural side-channel adversaries can only observe memory with cache line granularity. In this work, we propose *MemJam*, which utilizes *4K Aliasing* to establish a side-channel attack that exploits false dependency of memory read-after-write events and provides a high quality intra cache line timing channel. As a proof of concept, we demonstrate the first key recovery attacks on constant-time implementations of all symmetric block ciphers supported in the current intel integrated performance primitives (Intel IPP) cryptographic library: triple DES, AES and SM4. Further, we demonstrate the first intra cache level timing attack on SGX by reproducing the AES key recovery results on an enclave that performs encryption using the aforementioned constant-time implementation of AES. Our results show that we can not only use this side channel to efficiently attack memory dependent cryptographic operations but also to bypass proposed protections. Compared to *CacheBleed*, which is limited to older processor generations, *MemJam* is the first intra cache level attack applicable to all major Intel processors including the latest generations and also applies to the SGX extension.

Keywords Side-channel attacks · False dependency · Microarchitectural side channels · 4K Aliasing · Cache · Timing attack

This is an extended version of the paper that was presented in part at the RSA Conference Cryptographers Track (CT-RSA 2018, Springer LNCS) [53].

✉ Ahmad Moghimi
amoghimi@wpi.edu

Extended author information available on the last page of the article

1 Introduction

In cryptographic implementations, timing channels can be introduced by key dependent operations, which can be exploited by local or remote adversaries [18,58]. Modern microarchitectures are complex and support various shared resources, and the operating system (OS) maximizes the resource sharing among concurrent tasks [52,63]. From a security standpoint, concurrent tasks with different permissions share the same hardware resources, and these resources can expose exploitable timing channels. A typical model for exploiting microarchitectural timing channels is for a spy process to cause resource contention with a victim process and to measure the timing of its own or of the victim operations [2,45,62,67]. The observed timing behavior give adversaries strong evidence on the victim's resource usage pattern, thus they leak critical runtime data. Among the shared resources, attacks on cache have received significant attention, and their practicality has been demonstrated in scenarios such as cloud computing [31,36,45,62,76,79]. A distinguishable feature of cache attacks is the ability to track memory accesses with high temporal and spatial resolution. Thus, they excel at exploiting cryptographic implementations with secret dependent memory accesses [11,35,58,68]. Examples of such vulnerable implementations include using S-Box tables [71], and efficient implementations of modular exponentiation and scalar multiplication [34,48].

The weakness of key dependent cache activities has motivated researchers and practitioners to protect cryptographic implementations against cache attacks [15,67]. The simplest approach is to minimize the memory footprint of lookup tables. Using a single 8-Bit S-Box in the advanced encryption standard (AES) rather than T-Tables makes cache attacks on AES inefficient in a noisy environment, since the adversary can only distinguish accesses between 4 different cache lines. Combining small tables with cache state normalization, i.e., loading all table entries into cache before each operation, defeats cache attacks in asynchronous mode, where the adversary is only able to perform one observation per operation [59]. More advanced side channels such as exploitation of the thread scheduler [33], cache attack on interrupted execution of Intel Software Guard eXtension (SGX) [54], performance degradation [6] and leakage of other microarchitectural resources [1,3] remind us of the importance of constant-time software implementations. One way to achieve constant-time memory behavior is the adoption of small tables in combination with accessing all cache lines on each lookup [67]. The overhead would be limited and is minimized by the parallelism we can achieve in modern processors. Another constant-time approach adopted by some public cryptographic schemes is interleaving the multipliers in memory known as scatter–gather technique [16].

Constant-time implementations have effectively eliminated the first generation of timing attacks that exploit obvious key dependent leakages. The common view is that performance penalty is the only downside which, once paid, there is no need to be further worried. However, this is far from the reality and constant-time implementations may actually give a false sense of security. A commonly overlooked fact is that constant-time implementations and related protections are relative to the underlying hardware [29]. In fact, there are major obstacles preventing us from obtaining true constant-time behavior. Processors constantly evolve with new microarchitectural fea-

tures rolled quietly with each new release and the variety of such subtle features makes comprehensive evaluation impossible. A great example is the cache bank conflicts attack on OpenSSL RSA scatter–gather implementation: it shows that adversaries with intra cache level resolution can successfully bypass constant-time techniques relied on cache-line granularity [77]. As a consequence, what might appear as a perfect constant-time implementation becomes insecure in the next processor release—or worse—an unrecognized behavior might be discovered, invalidating the earlier assumption.

1.1 Our Contribution

We propose an attack named *MemJam* by exploiting false dependency of memory read-after-write, and demonstrate key recovery against three different cryptographic implementations which are secure against cache attacks with experimental results on both regular and SGX environments. In summary:

- **False dependency attack** A side-channel attack on the false dependency of memory read-after-write. We show how to dramatically slow down the victim's accesses to specific memory blocks, and how this read latency can be exploited to recover low address bits of the victim's memory accesses.
- **Attack on all block ciphers in IPP** The attacks utilize intra cache level information on constant-time implementations of Triple DES, AES and SM4, chosen from Intel integrated performance primitives (Intel IPP), which is optimized for both security and speed and a default choice to be used in SGX enclaves.
- **Attack on SGX enclave** The first intra cache level attack against SGX Enclaves supported by key recovery results on the constant-time AES implementation. The aforementioned constant-time implementation of AES is part of the SGX SDK source code.
- **Countermeasures** Discussion of software-level and hardware-level countermeasures against *MemJam* including various constant-time implementation techniques. Bypasses of remarkable protections such as proposals based on constant-time techniques [16,67], static and runtime analysis [46,78] and new cache architecture [21,47,51,73].

1.2 Experimental Setup and Generic Assumptions

Our experimental setup is a Dell XPS 8920 desktop machine with an Intel(R) Core i7-7700 processor running Ubuntu 16.04. The Core i7-7700 has 4 hyper-threaded physical cores. Our only assumptions are that the attacker is able to co-locate on one of the logical processor pairs within the same physical core as the victim. In the cryptographic attacks, the attacker can measure the time of victim encryption. The attacker further knows which cryptographic implementation is used by the victim, but she does not need to have any knowledge of the victim's binary or the offset of the S-Box tables. We will discuss assumptions that are specific to the attack on SGX at Sect. 6.

2 Related Work

Side channels including power, electromagnetic and timing channels have been studied for a few decades [18,19,49]. Timing side channels can be constructed through the processor cache to perform key recovery attacks against cryptographic operations such as RSA [35], ECDSA [11], ElGamal [79], DES [68] and AES [45,58]. On multiprocessor systems, attacks on the shared LLC—a shared resource among all the cores—perform well even when attacker and victim reside in different cores [45]. Flush+Reload, Prime+Probe, Evict+Reload, and Flush+Flush are some of the proposed attack methodologies with different adversarial scenarios [31,58,76]. Performance degradation attacks can improve the channel resolution [6,33]. LLC attacks are highly practical in cloud, where an attacker can identify where a particular victim is located [62,79]. Despite the applicability of LLC attacks, attacks on core-private resources such as L1 cache are as important [1,13]. Attacks on SGX in a system level adversarial scenario are notable examples [50,54]. There are other shared resources, which can be utilized to construct timing channels [28]. Exploitation of Branch Target Buffer (BTB) leaks if a branch has been taken by a victim process [1,3,50]. Logical units within the processor can leak information about arithmetic operations [4,8]. *CacheBleed* proposes cache bank conflicts and false dependency of memory write-after-read as side channels with intra-cache granularity [77]. However, cache bank conflict leakage does not exist on current Intel processors, and we verify the authors' claim that the proposed write-after-read false dependency side channel does not allow efficient attacks. In a concurrent, but an independent contribution, *4K Aliasing* has been analyzed and used as a covert channel [66]: They show that false dependency side channel can be used to detect multi-tenancy in Infrastructure as a Service (IaaS) Clouds.

2.1 Defense

Software and hardware strategies have been proposed such as alternative lookup tables, data-independent memory access pattern, static or disabled cache, and cache state normalization to defend against cache attacks [67]. Scatter–gather techniques have been adopted by RSA and ECC implementations [16]. In particular, introducing redundancy and randomness to the S-Box tables for AES has been proposed [15]. A custom memory manager [80], relaxed inclusion caches [47] and solutions based on cache allocation technology (CAT) such as Catalyst [51] and vCat [73] are proposed to defend against LLC contention. Sanctum [21] is a new processor design with respect to cache attacks. Detection-based countermeasures have also been proposed using performance counters, which can be used to detect cache attacks in cloud environments [17,78]. MASCAT [46] is proposed to block cache attacks with code analysis techniques. CacheD [70] detects potential cache leakages in production software. Nonetheless, these proposals assume that the adversary cannot distinguish accesses within a cache line. That is, attacks with intra cache-line granularity are considered out-of-scope. Doychev et al. proposed the only software leakage detector that considers full address bits as its leakage model [25].

3 Background

3.1 Multitasking

The memory management subsystem shares the dynamic random-access memory (DRAM) among all concurrent tasks, in which a virtual memory region is allocated for each task transparent to the physical memory. Each task is able to use its entire virtual address space without meddling of memory accesses from others. Memory allocations are performed in pages, where each virtual memory page can be stored in a DRAM page with a virtual-to-physical page mapping. The logical processors are also shared among these tasks and each logical processor executes instructions from one task at a time, and switches to another task. Memory write and read instructions work with virtual addresses, and the virtual address is translated to the corresponding physical address to perform the memory operation. The OS is responsible for page directory management and virtual page allocation. The OS assists the processor to perform virtual-to-physical address translation by performing an expensive page walk. The processor saves the address translation results in a memory known as Translation Look-aside Buffer (TLB) to avoid the software overhead introduced by the OS. Intel microarchitecture follows a multi-stage pipeline and adopts different optimization techniques to maximize the parallelism and multitasking during the pipeline stages [39]. Among these techniques, hyper-threading allows each core to run multiple concurrent threads, and each thread shares all the core-private resources. As a result, if one resource is busy by a thread, other threads can consume the remaining available resources. Hyper-threading is abstracted to the software stack: OS and applications interact with the logical processors.

3.2 Cache Memory

DRAM memory is slow compared to the internal CPU components. Modern microarchitectures take advantage of a hierarchy of cache memories to fill the speed gap. Intel processors have two levels of core-private cache (L1, L2), and a Last Level Cache (LLC) shared among all cores. The closer the cache memory is to the processor, the faster, but also smaller it is compared to the next level cache. Cache memory is organized into different sets, and each set can store some number of cache lines. The cache line size, which is 64 byte, is the block size for all memory operations outside of the CPU. The higher bits of the physical address of each cache line is used to determine which set to store/load the cache line. When the processor tries to access a cache line, a cache hit or miss occurs respective of its existence in the relevant cache set. If a cache miss occurs, the memory line will be stored to all 3 levels of cache and to the determined sets. Reloads from the same address would be much faster when the memory line exists in cache. In a multicore system, the processor has to keep cache consistent among all levels. In Intel architecture, cache lines follow a write-back policy, i.e., if the data in L1 cache is overwritten, all other levels will be updated. The LLC is inclusive of L2 and L1 caches, which means that if a cache line in LLC is evicted, the corresponding L1 and L2 cache lines will also be evicted [39]. These policies help

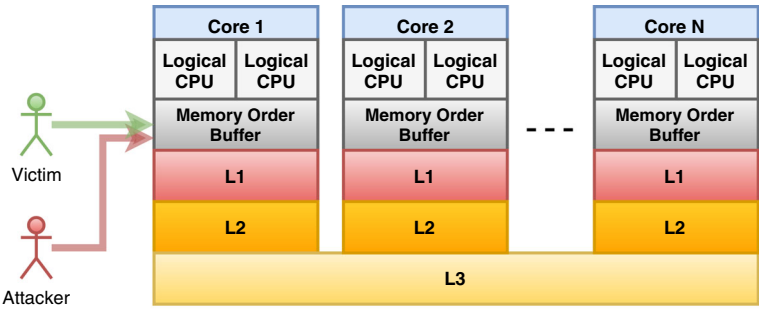

Fig. 1 Cache hierarchy of an Intel processor: the L3 cache is shared among available cores. Core-private caches such as L1 and L2 are shared between logical hyper-threading CPUs. Memory access dependencies are determined within the memory order buffer (MOB). An adversary who is co-located on the same core exploiting hyper-threading, can mount attacks on victims sharing the same resources

to avoid stale cached data where one processor reads invalid data mutated by another processor.

3.3 L1 Cache Bottlenecks

The L1 cache port has a limited bandwidth and simultaneous accesses will block each other. This bottleneck is critical in super-scalar multiprocessor systems. Older processor generations adopted multiple banks as a workaround to this problem [5], in which each bank can operate independently and serve one request at a time. While this partially solved the bandwidth limit, it creates the cache bank conflict phenomena where simultaneous accesses to the same bank will be blocked. Intel resolved the cache bank conflict issue with the Haswell generation [39]. Another bottleneck mentioned in various resources is due to the false dependency of memory addresses with the same cache set and offset [5,39]. Simultaneous read and write with addresses that are multiples of 4 kB is not possible, and they halt each other. The processor cannot determine the dependency from the virtual address, and addresses with the same last 12 bits have the chance to map to the same physical address. Such simultaneous access can happen between two logical processors and/or during the out-of-order execution, where there is a chance that a memory write/read might be dependent on a memory read/write with the same last 12 bits of address. Such dependencies cannot be determined on the fly, thus they cause latency (Fig. 1).

3.4 Cache Attacks

Cache attacks can be exploited by adversaries where they share system cache memory with benign users. In scenarios where the adversary can colocate with a victim on the same core, she can attack core-private resources such as L1 cache, e.g., OS adversaries [50,54]. In cloud environment, virtualization platforms allow sharing of logical processors to different VMs; however, attacks on the shared LLC have a higher impact, since LLC is shared across the cores. In cache timing attacks, the

attacker either measure the timing of the victim operations, e.g., *Evict+Time* [58] or the timing of his own memory accesses, e.g., *Prime+Probe* [45]. The attacker needs to have access to an accurate time resource such as the *RDTSC* instruction. In the basic form, attacks are performed by one observation per entire operation. In certain scenarios, these attacks can be improved by interrupting the victim and collecting information about the intermediate memory states. Side-channel attacks exploiting cache bank conflicts rely on synchronous resource contention. *CacheBleed* methodology is somewhat similar to Prime+Probe, where the attacker performs repeated operations, and measures it's own access time [77]. In a cache bank conflict attack, the adversary repeatedly performs simultaneous reads to the same cache bank and measures their completion time. A victim on a colocated logical processor who access the same cache bank would cause latency to the attacker's memory reads.

4 *MemJam*: Read-After-Write Attack

MemJam utilizes *false dependencies*. Data dependency occurs when an instruction refers to the data of a preceding instruction. In pipelined designs, hazards and pipeline stalls can occur from dependencies if the previous instruction has not finished. There are cases where false dependencies occur, i.e. the pipeline stalls even though there is no true dependency. Reasons for false dependencies are register reuse and limited address space for the Arithmetic Logic Unit (ALU). False dependencies degrade instruction level parallelism and cause overhead. The processor eliminates false dependencies arising from register reuse by a register renaming approach. However, there exist other false dependencies that need to be addressed during the software optimization, e.g. *Partial Register Stalls* [39,40].

In this work, we focus on a critical false dependency mentioned as *4K Aliasing* where data that is multiples of 4K apart in the address space is seen as dependent. 4K Aliasing happens due to virtual addressing of L1 cache, where data is accessed using virtual addresses, but tagged and stored using physical addresses. Multiple virtual addresses can refer to the same data with the same physical address and the determination of dependency for concurrent memory accesses, requires virtual address translation. Physical and virtual address share the last 12 bits, and any data accesses whose addresses differ in the last 12 bits (i.e. the distance is not 4k) cannot have a dependency. For the fairly rare remaining cases, address translation needs to be done before resolving the dependency, which causes latency. Note that the *granularity* of the potential dependency, i.e. whether two addresses are considered "same", depends also on the microarchitecture, as dependencies can occur at the *word* or cache *line* granularity (i.e. ignoring the last 2 or last 6 bits of the address, respectively). These rare false dependencies due to 4K Aliasing can be exploited to attack memory, since the attacker can deliberately process falsely dependent data by matching the last 12 bits of his own address with a security critical data inside a victim process.

4K Aliasing has been mentioned in various places as an optimization problem existing on all major Intel processors [5,39]. We verify the results of Yarom et al. [77], the only security related work regarding false dependencies, which exploited *write-after-read* dependencies. The resulting timing leakage by write stall after read is not sufficient to be used in any cryptographic attack. *MemJam* exploits a different channel due to the false dependency of *read-after-write*, which causes a higher latency and is thus simply observable. Intel Optimization Manual highlights the *read-after-write* performance overhead in various sections [39]. As described in Section 11.8, this hazard occurs when a memory write is closely followed by a read, and it causes the read to be reissued with a potential 5 cycles penalty.[1] In Section B.1.4 on memory bounds, write operations are treated under the store bound category. In contrast to load bounds, Top-down Microarchitecture Analysis Method (TMAM)[2] reports store bounds as fraction of cycles with low execution port utilization and small performance impact. These descriptions in various sections highlight that *read-after-write* stall is considered more critical than *write-after-read* stall.

4.1 Memory Dependency Fuzz Testing

We performed a set of experiments to evaluate the memory dependency behavior between two logical processors. In these experiments, we have thread $\mathcal{A}$ and $\mathcal{B}$ running on the *same* physical core, but on *different* logical processors, as shown in Fig. 2. Both threads perform memory operations; only thread $\mathcal{B}$ measures its timing and hence the timing impact of introduced false dependencies.

Read-after-read (**RaR**) In the first experiment, the two logical threads $\mathcal{A}$ and $\mathcal{B}$ read from the same shared cache and can potentially block each other. This experiment can reveal cache bank conflicts, as used by *CacheBleed* [77]. $\mathcal{B}$ uses Listing 1 to perform read measurements and $\mathcal{A}$ constantly reads from different memory offsets and tries to introduce conflicts. $\mathcal{A}$ reads from three different type of offsets: (**1**) Different cache line than $\mathcal{B}$, (**2**) same cache line, but different offset than $\mathcal{B}$, and (**3**) same cache line and same offset as $\mathcal{B}$. As depicted, there is no obvious difference between the histograms for three cases in Fig. 3a verifying the lack of cache bank conflicts on 7th generation CPUs.

Write-after-read (**WaR**) The histogram results for the second experiment on false dependency of write-after-read is shown in Fig. 3b, in which the cache line granularity is obvious. Thread $\mathcal{A}$ constantly reads from different type of memory offsets, while thread $\mathcal{B}$ uses Listing 2 to perform write measurements. The standard deviation for conflicted cache line (blue) and conflicted offset (red) between thread $\mathcal{A}$ and $\mathcal{B}$ is distinguishable from the green bar where there is no cache line conflict. This shows a high capacity cache granular behavior, but the slight difference between conflicted line and offset verifies the previous results stating a weak offset dependency [77].

[1] `LD_BLOCKS_PARTIAL.ADDRESS_ALIAS` Performance Monitoring Unit (PMU) event counts the number of times reads were blocked.

[2] Top-Down Characterization is a hierarchical organization of event-based metrics that identifies the dominant performance bottlenecks in an application.

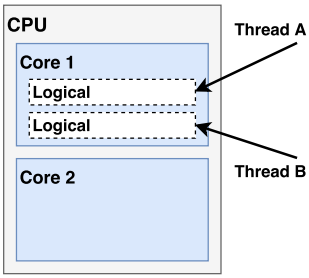

Fig. 2 Based on the attack model, thread *A* and *B* both run on the same core, and introduce and probe stall hazards

```
loop:
rdtscp;
mov % eax, (%r9);
movb 0 x 0000 (%r10), % al;
movb 0 x 1000 (%r10), % al;
movb 0 x 2000 (%r10), % al;
movb 0 x 3000 (%r10), % al;
movb 0 x 4000 (%r10), % al;
movb 0 x 5000 (%r10), % al;
movb 0 x 6000 (%r10), % al;
movb 0 x 7000 (%r10), % al;
add $4, %r9;
dec % r11;
jnz loop;
```

Listing 1 is used to probe 8 parallel reads. *r*9 points to a measurement buffer, and *r*11 is initialized with the probe count

```
loop:
rdtscp
mov % eax, (%r9);
movb % al, 0 x 0000 (%r10);
movb % al, 0 x 1000 (%r10);
movb % al, 0 x 2000 (%r10);
movb % al, 0 x 3000 (%r10);
movb % al, 0 x 4000 (%r10);
movb % al, 0 x 5000 (%r10);
movb % al, 0 x 6000 (%r10);
movb % al, 0 x 7000 (%r10);
add $4, %r9
dec % r11
jnz loop
```

Listing 2 is used to probe 8 parallel writes. *r*9 points to a measurement buffer, and *r*11 is initialized with the probe count

Read-after-write (**RaW**) Figure 3c shows an experiment on measuring false dependency of read-after-write, in which, thread $\mathcal{A}$ constantly writes to different memory offsets. Thread $\mathcal{B}$ uses Listing 1 to perform read measurements. Accesses to three different types of offsets are clearly distinguishable. The conflicted cache line accesses (blue) are distinguishable from non-conflicted accesses (green). More importantly, conflicted accesses to the same offset (red) are also distinguishable from conflicted cache line accesses, resulting in a side channel with intra cache-line granularity. There is an average of 2 cycle penalty if the same cache line has been accessed, and a 10 cycle penalty if the same offset has been accessed. Note that the word offsets in our platform have 4 byte granularity. From an adversarial standpoint, this means that an adversary

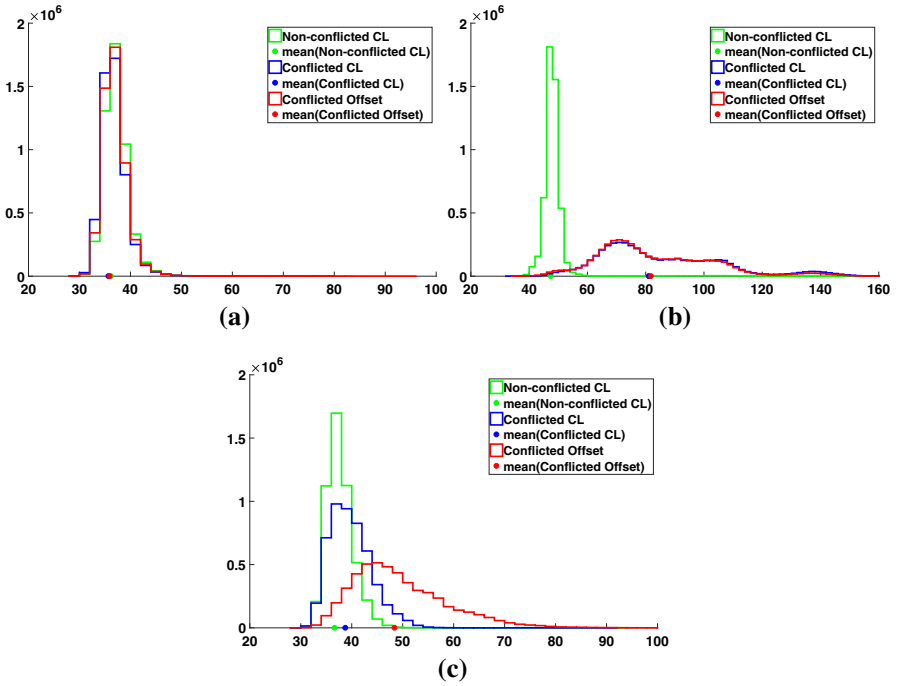

Fig. 3 Three different scenarios where different cache line (green), same cache line (blue) and same offset (red) have been accessed by two logical processors. Experiment (**c**) on RaW latency has distinguishable characteristics for the conflicted word offset (red), while (**a**), (**b**) feature nimble differences (Color figure online)

learns about bits 2–11 of the victim memory access, in which 4 bits (bits 2–5) are related to intra cache-line resolution, and thus goes beyond any other microarchitectural side channels known to exist on 6th and 7th generation Intel processors (Fig. 6).

Read-after-weak-Write (**RawW**) In this experiment on the read-after-write conflicts, we followed a less greedy strategy on the conflicting thread. Rather than constantly writing to the same offset, $\mathcal{A}$ executes write instructions to the same offset with some gaps filled with other memory accesses and instructions. As shown in Fig. 4, the channel dramatically became less effective. This tells us that causing read access penalty will be more effective with constant writes to the same offset without additional instruction. In this regard, we will use Listing 3 in our attack to achieve the maximum conflicts.

Read-after-write latency In the last experiment, we tested the delay of execution over a varying number of conflicting reads. We created a code stub that includes 64 memory read instructions and a random combination of instructions between memory reads to create a more realistic computation. The combination is chosen in a way to avoid unexpected halts and to maintain the parallelism of all read operations. We measure the execution time of this computation on $\mathcal{B}$, while $\mathcal{A}$ is writing to a conflicting offset. First, we measured the computation with 64 memory reads to addresses without conflicts. Our randomly generated code stub takes an average of 210 cycles to execute. On each

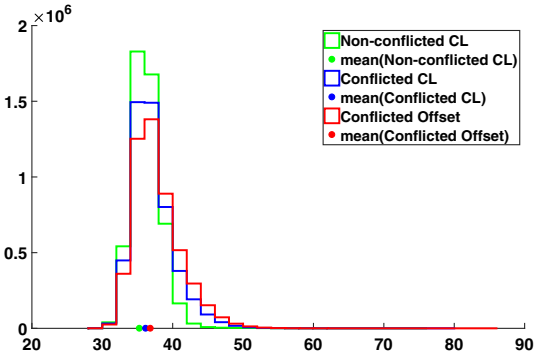

Fig. 4 RawW: compared to Fig. 3c, this shows a lower impact on access latency

step of the experiments, as shown in Fig. 5, we change some of the memory offsets to have the same last 12 bits of address as of $\mathcal{A}$'s conflicting write offset. We observe a growth on read accesses' latency by increasing the number of conflicting reads. Fig. 5 shows the results for a number of experiments. In all of them, the overall execution time of $\mathcal{B}$ is strongly dependent on the number of conflicting reads. Hence, we can use the RaW dependency to introduce strong timing behavior using bits 2–11 of a chosen target memory address.

```
mov %[target], %rax;
write_loop:
  .rept 100;
 movb $0, (%rax);
  .endr;
jmp write_loop;
```

Listing 3 Write Conflict Loop: Unnecessarily instructions are avoided to minimize usage of other processor units and to maximize the RaW conflict effect.

5 *MemJam* Correlation Attack

MemJam uses read-after-write false dependencies to introduce timing behavior to otherwise constant-time implementations. The resulting latency is then exploited using a correlation attack. *MemJam* proceeds with the following steps:

1. Attacker launches a process constantly writing to an address using Listing 3 where the last 12 bits match the virtual memory offset of a *critical* data that is read in the victim's process.
2. While the attacker's conflicting process is running, attacker queries the victim for encryption and records a ciphertext and execution time pair of the victim. Higher time infers more accesses to the *critical* offset.
3. Attacker repeats the previous step collecting ciphertext and time pairs.

The attack methodology resembles the *Evict+Time* strategy originally proposed by Tromer et al. [67], except that the attacker uses false dependencies rather than evictions to slow down the target *and* that the slowdown only applies to an 4-byte block of a cache line. Furthermore, *all* of the victim's accesses addresses with the same last 12 bits are slowed down while an eviction only slows the first memory access(es).

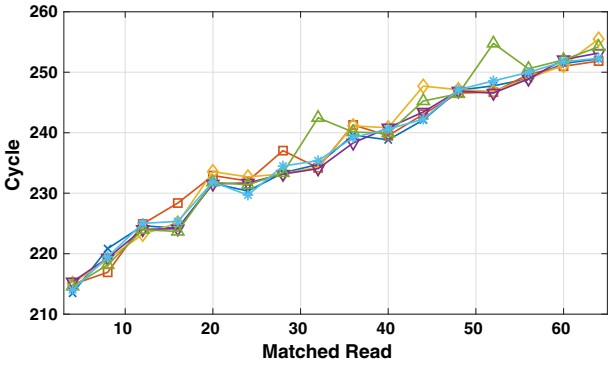

Fig. 5 The cycle count for mixed operations with RaW conflicts. More conflicts cause higher delay

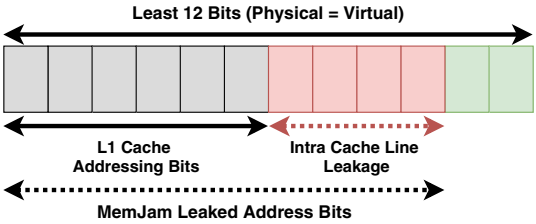

Fig. 6 Intra cache level leakage: *MemJam* latency is related to 10 address bits, in which 4 bits are intra cache level bits

Based on the intra cache level leakage in Fig. 6, we divide a 64 byte cache line into 4-byte blocks and hypothesize that the access counts to a block are correlated with the running time of victim, while the attacker jams memory reads to that block, i.e., the attacker expects to observe a higher time when there are more accesses by the victim to the targeted 4-byte block and lower time when there are lower number of accesses. Based on this hypothesis, we apply a classical correlation based side-channel approach [49] to attack implementations of three different block ciphers, namely Triple DES, AES and SM4. SM4 among AES, Triple DES, and RC4 are the only available symmetric ciphers as part of Intel's IPP crypto library [42].[3] Each implementation has optimizations to hinder cache attacks. In fact, the Triple DES and the AES implementations feature a constant cache profile and can thus be considered resistant to most microarchitectural attacks including cache attacks and high-resolution attacks as described in [54]. *MemJam* can still extract the keys from both implementations due to the intra cache-line spatial resolution as depicted in Fig. 6. We describe the targeted implementations next, as well as the correlation models we use to attack them.

5.1 Attack 1: IPP Constant-Time Triple DES

The *Triple DES* encryption algorithm [55] is an extension of the DES (Data Encryption Standard) algorithm. While Triple DES was recently declared as deprecated [56], mainly due to its insufficient block size [12] and the resulting attacks for larger amounts of encrypted data, it is still supported or even required in many applications: The

[3] Patents investigated by Intel verify the importance of SM4 [32,72,75].

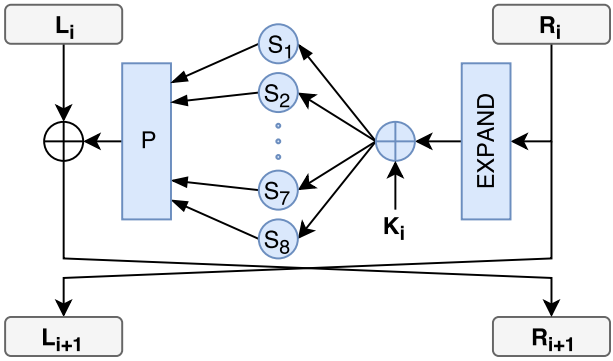

Fig. 7 Feistel function (blue) in round i of the DES algorithm: first the current right block R_i is expanded to 48 bits and XORed with the round key. Then this value is divided into eight 6-bit blocks, which are substituted by 4-bit blocks using eight different S-boxes. Finally the result is permuted and XORed with the current left block (Color figure online)

latest EMVCo[4] specification for payment systems permits its usage without further restrictions [27], and current TLS 1.2 standard contains Triple DES as a legacy cipher [23].

Given three different 56-bit keys $\mathcal{K}_1, \mathcal{K}_2, \mathcal{K}_3$ and a 64-bit plain text block M, Triple DES in Encrypt–Decrypt–Encrypt (EDE) mode calculates the cipher text C as

$$C := 3\mathrm{DES}_{\mathcal{K}_1,\mathcal{K}_2,\mathcal{K}_3}(M) = \mathrm{DES}_{\mathcal{K}_3}\left(\mathrm{DES}^{-1}_{\mathcal{K}_2}\left(\mathrm{DES}_{\mathcal{K}_1}(M)\right)\right).$$

DES itself is a Feistel network with 16 rounds. First the plain text M is permuted using an initial permutation $M' := \mathrm{IP}(M)$ and then divided into two 32-bit blocks $M' = L_0.R_0$. In round $i \in \{0, \ldots, 15\}$ the algorithm then calculates

$$L_{i+1} := R_i \qquad \text{and} \qquad R_{i+1} = L_i \oplus f(K_i, R_i)$$

for a given round key K_i. The ciphertext C is obtained by applying the inverse of the initial permutation to the last blocks:

$$C := \mathrm{IP}^{-1}\left(L_{16} \cdot R_{16}\right).$$

The Feistel function f (Fig. 7) takes a 48-bit round key K_i and the current right block R_i, and computes its output by doing the following steps:

1. Expand R_i to 48 bits by generating eight 6-bit blocks

$$B_{i,j} := R_i[4j - 1 \bmod 32].R_i[4j + 0] \ldots R_i[4j + 3].R_i[4j + 4 \bmod 32]$$

for $j \in \{0, \ldots, 7\}$.

[4] EMVCo is an industry consortium managing a payment system standard that was originally created by EuroPay, MasterCard and Visa (resulting in the EMV trademark). Current members include American Express, MasterCard, Visa and UnionPay [26].

2. Partition the round key to eight 6-bit blocks $K_i = K_{i,0} \dots K_{i,7}$ and set the substitution box inputs as
$$S_{i,j}^{\text{in}} := B_{i,j} \oplus K_{i,j}$$
for each $j \in \{0, \dots, 7\}$.
3. Use eight S-boxes $\text{S}_0, \dots \text{S}_7$ to convert the 6-bit inputs into 4-bit outputs:
$$S_{i,j}^{\text{out}} := \text{S}_j \left(S_{i,j}^{\text{in}} \right)$$
for each $j \in \{0, \dots, 7\}$.
4. Permute the S-Box outputs using a round permutation P to acquire the Feistel function output
$$\text{output} := \text{P} \left(S_{i,0}^{\text{out}} \dots S_{i,7}^{\text{out}} \right).$$

The round keys are generated using a schedule consisting of left shifts and permutations [55], we skip a deeper explanation here. Decryption works the same as encryption, except that the round keys are applied in reverse order.

Our target, the Triple DES implementation of Intel's Integrated Performance Primitives Crypto library, comes in various flavors where each is optimized for a specific instruction set, but they all have similar cache behavior: The central DES encryption/decryption function `Cipher_DES` first applies the initial permutation, that is implemented as a fixed number of bit operations without any memory accesses. The following 16 rounds are unrolled, each round has exactly $2 + 16$ memory accesses, where the first two memory accesses load the respective round key. The eight S-box inputs are processed consecutively; for each input (**1**) the substitution is performed (by reading from the fixed S-box array), and then (**2**) the 4-bit S-Box output is converted into its 32-bit permuted form (using another lookup table). Finally, these permuted outputs are XORed with each other to acquire the output of the Feistel function. Each S-box has $2^6 = 64$ 1-byte entries and therefore fits exactly into one cache line; the same applies to the permutation lookup table, which has $2^4 = 16$ 4-byte entries.

This implies that each cache line is accessed exactly once per round, leading to constant time cache behavior that prevents any attacks with cache-line granularity. To obtain data-dependent timing behavior, we used *MemJam* to induce false dependencies on the first four bytes of the first S-box, slowing down the read accesses to this offset. Since this gives us 4 bytes of resolution, we can deduce 4 bits of the respective S-box input, which correspond to 4 bits of the round key. A single observation consists of the resulting cipher text C_i and the amount of clock cycles T_i the Triple DES operation takes to execute. Using n of such measurements (with random plain texts), we can work ourselves into the cipher, starting from the last round.

Single-round attack on triple DES Each cipher text C_i consists of blocks $L_{16} = R_{15}$ and R_{16}, where the former directly gives us the eight 6-bit blocks $B_{15,0}, \dots B_{15,7}$. We guess the round key block $K_{15,0}$, and set

- $v[i] := 1$, if $S_{15,0}^{\text{in}} = B_{15,0} \oplus K_{15,0} = \cdot\cdot 0000$
- $v[i] := 0$, else

for a binary vector $v \in \{0, 1\}^n$.

We loose the two least significant bits (written as "·") due to the 4-byte resolution of *MemJam*. Since the IPP implementation reverses the bit order of each block and round key, the least significant bits are written first. Maximizing the correlation

$$\mathrm{corr}(v, T)$$

between the binary vector v and the clock cycle count vector T over all possible round key blocks $K_{15,0}$ then gives us the four key bits $\mathcal{K}_3[2]$, $\mathcal{K}_3[21]$, $\mathcal{K}_3[36]$ and $\mathcal{K}_3[49]$, since the slow runs should be nearly uniformly distributed for wrong guesses.

Multi-round attack on triple DES To get the missing 52 key bits, we repeat the attack process in a similar fashion for round 14: The round key block $K_{14,0}$ that we are interested in gives us key bits $\mathcal{K}_3[9]$, $\mathcal{K}_3[28]$, $\mathcal{K}_3[31]$ and $\mathcal{K}_3[43]$, but we also need the last four bits of block $B_{14,0}$; for these, we have to partially calculate $L_{15} = R_{16} \oplus \mathrm{P}(S^{\mathtt{out}}_{15,0} \ldots S^{\mathtt{out}}_{15,7})$, which depends on $K_{15,1}$, $K_{15,4}$, $K_{15,5}$ and $K_{15,7}$, summing up to $4 \cdot 6 = 24$ additional key bits, of which two are already included in the round key $K_{14,0}$.

Repeating the same process for round 13 (where we need almost all key bits from round 15 to calculate the relevant S-boxes in round 14) yields another 21 bits of key $\mathcal{K}_3$. The remaining 5 key bits are derived from round 12. To obtain the remaining keys $\mathcal{K}_1$ and $\mathcal{K}_2$, we repeat the attack using cipher texts decrypted with $\mathcal{K}_3$.

To reduce the computational effort one can also take additional measurements on the other S-boxes, yielding up to 32 key bits in round 15; however, this also multiplies the amount of measurements, and one still needs to analyze prior rounds to retrieve the missing 24 key bits, although with greatly reduced time complexity. So, overall, we see that there is a trade-off between the amount of measurements and the computation time spent on the analysis.

Triple DES key recovery results on synthetic data To verify the correctness of our attack we first generated some synthetic data, where the timings were set as the amount of accesses to the first four bytes of the first S-box. In this noise free setting we needed less than 1000 observations to find 19 bits of the 14th round key, with a correlation of 0.201.

Triple DES key recovery results using *MemJam* The time needed for a successful attack primarily depends on the amount of measurements and the number of simultaneously guessed bits. The attacks on round 15 (4 key bits) and 12 (5 key bits) are negligible, but round 14 (26 key bits) needs $2^{26}n$ steps and round 13 (21 key bits) $2^{21}n$ steps; this corresponds to tens of hours of computation time per DES key. While this is significantly less than guessing all 56 bits at once, reducing the amount of measurements is still desirable. Figure 8 shows the correlations for different measurement counts, when guessing 14 key bits in round 14. Experiments showed that 250,000–300,000 measurements suffice to recover all three keys.

5.2 Attack 2: IPP Constant-Time AES

AES is a cipher based on substitution permutation network (SPN) with 10 rounds supporting 128-bit blocks and 128/192/256-bit keys [22]. The SubBytes is a security-

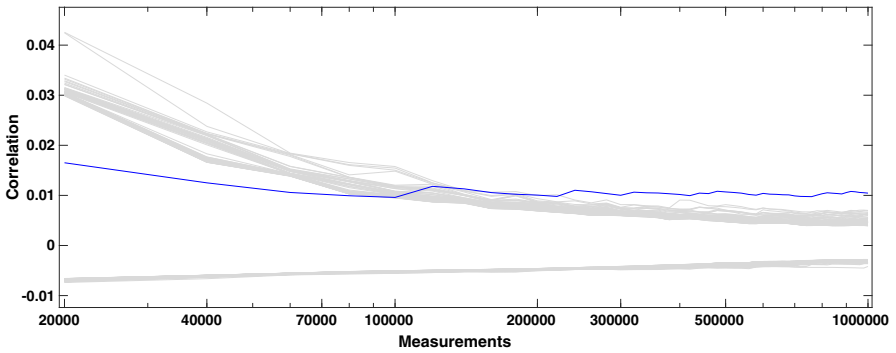

Fig. 8 The hundred highest and lowest timing correlations when guessing 14 key bits in round 14, depending on the amount of measurements (logarithmic scale). The correct key (blue) becomes distinguishable at around 250,000 measurements (Color figure online)

critical operation and the straightforward way to implement AES SubBytes operation efficiently in software is to use lookup tables. SubBytes operates on each byte of cipher state, and it maps an 8-bit input to an 8-bit output using a non-linear function. A precomputed 256 byte lookup table known as S-Box can be used to avoid recomputation. There are efficient implementations using T-Tables that output 32-bit states and combine SubBytes and MixColumns operations. T-Table implementations are highly vulnerable to cache attacks. During AES rounds, a state table is initiated with the plaintext, and it holds the intermediate state of the cipher. Round keys are mixed with states, which are critical S-Box inputs and the main source of leakage. Hence, even an adversary who can partially determine which entry of the S-Box has been accessed is able to learn some information about the key.

Among the efforts to make AES implementations more secure against cache attacks, `Safe2Encrypt_RIJ128` function from Intel IPP cryptographic library is noteworthy. This implementation is the only production-level AES software implementation that features true cache constant-time behavior and does not utilize hardware extensions such as AES-NI or SSSE3 instruction sets. This implementation is also part of the Linux SGX SDK [38] and can be used for production code if the SDK is compiled from the scratch, i.e., it does not use prebuilt binaries. We verified the match between the implementation in Intel IPP binary and SGX SDK source code through reverse engineering. This implementation follows a very simple direction: (**1**) it implements AES using 256 byte S-Box lookups without any optimization such as T-Tables, (**2**) instead of accessing a single byte of memory on each S-Box lookup, it fetches four values from the same vertical column of 4 different cache lines and saves them to a local cache aligned buffer, finally, (**3**) It performs the S-Box replacement by picking the correct S-Box entry from the local buffer. This implementation is depicted in Fig. 9. This implementation protects AES against any kind of cache attacks, as the attacker sees a constant cache access pattern: The S-Box table only occupies 4 cache lines, and on each SubBytes operation, all of them will sequentially be accessed. This implementation can be executed in less than 2000 cycles on a recent laptop processor.

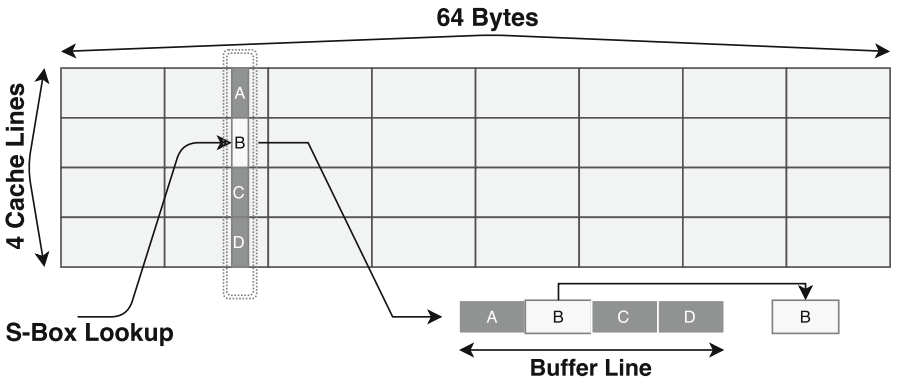

Fig. 9 Constant-time table lookup used by Intel IPP: each lookup preloads 4 values to a cache aligned buffer, thus it accesses all the 4 S-Box cache lines. The actual output will be chosen from the buffer using the high address bits

This is fast enough for many cryptographic applications, and it provides full protection against cache attacks, even if the attacker can interrupt the execution pipeline.

Based on *MemJam* 4-byte granular leakage channel, and the design of AES, we can create a simple correlation model to attack this implementation. The accessed table index of the last round for a given ciphertext byte c and key byte k is given as $index = S^{-1}(c \oplus k)$. We define matrix **A** for the access profile where each row corresponds to a known ciphertext, and each column indicates the number of accesses when $index < 4$. While we assume that the attacker causes slow-downs to the first 4-byte block of S-Box, we define matrix **L** for leakage where each row corresponds to a known ciphertext and each column indicates the victim's encryption time. Then our correlation attack is defined as the correlation between **A** and **L**, in which the higher the number of accesses, the higher the running time. Our results will verify that correlation is high, even though the implementation has dummy accesses to the monitored block. These can be ignored as noise, slightly reducing our maximum achievable correlation.

AES key recovery results on synthetic data We first verified the correctness of our correlation model on synthetic data using a noise free leakage (generated by PIN [41]). For each of the 16 key bytes using a vector that matches exactly to the number of accesses to the targeted block of S-Box for different ciphertexts, all the correct key bytes will have the highest correlation after 32,000 observations with the best and worst correlations of 0.046 and 0.029 respectively.

AES Key recovery results using (*MemJam*) Relying on the verification of Synthetic Data, we plugged in the real attack data vector, which consists of pairs of ciphertext and time measured through repeated encryption of unknown data blocks. Results on AES show that we can effectively exploit the timing information, and break the so-called constant-time implementation. The victim execution of AES encryption function takes about 1700 and 2000 cycles without and with an active thread on the logical processor pair, respectively. The target AES implementation performs 640 memory accesses to the S-Box, including dummy accesses. If the spy thread constantly writes to any address that collides with an S-Box block offset, the time will increase to a range

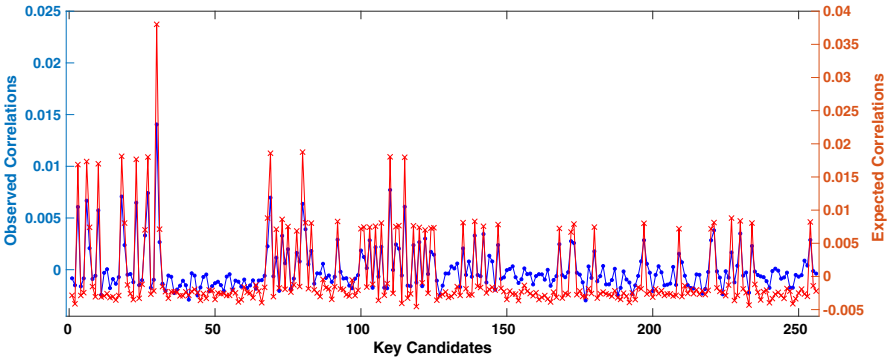

Fig. 10 Linearity of the number of accesses to the first block and the execution time of AES: the synthetic correlation and *MemJam* observed correlation show similar behavior with slight difference due to the added noise

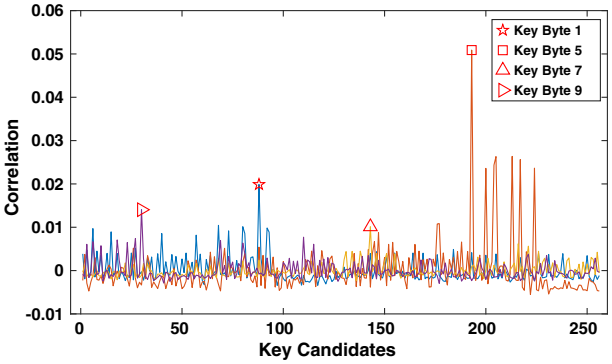

Fig. 11 Correlations for 4 key bytes using 2 million observations. Correct key byte candidates have the highest correlations

between 2000 and 2300 cycles. The observed variation in this range has a correlation with the number of accesses to that block. Figure 10 shows the linear relationship between the correlation of synthetic data and real attack data for one key byte after 2 million observations. Most of the possible key candidates for a target key byte have a matching peak and hill between the two observations. The highest correlation points in both cases declare the correct key byte (0.038 red, 0.014 blue). The quantitative difference is due to the expected noise in the real measurements.

Figure 11 shows the correlation of 4 different key bytes after 2 million observations with the correct key bytes having the highest correlations. Our repeated experiments with different keys and ciphertexts show that 15 correct key bytes have the highest correlation ranks, and there is only the key byte at index 15 that has a high rank but not necessarily the highest. Figure 12 shows the key ranks over the number of observations. Key byte ranks take values between 1 and 256, where 1 means that the correct key byte is the most likely one. As it is shown, after only 200,000 observations, the key space is reduced to a computationally insecure space and a key can be found with an efficient key enumeration method [30]. After 2 million observations, all key bytes

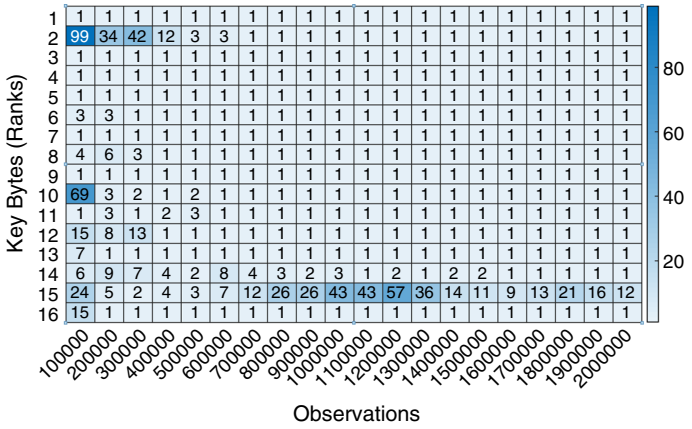

Fig. 12 The rank for correct key bytes are reduced with more observation. After 2 million observations, 15 out of 16 key bytes are recovered

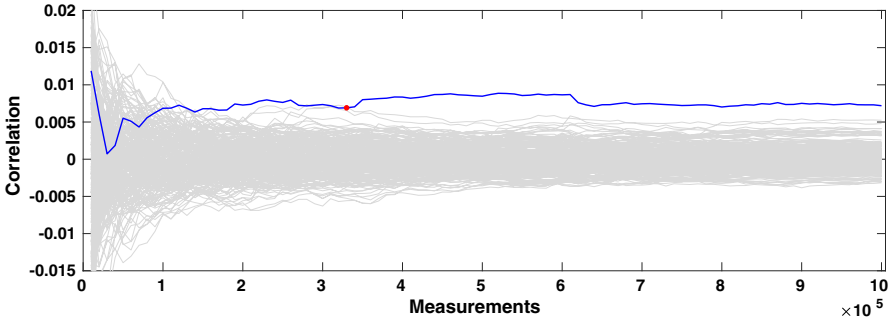

Fig. 13 The timing correlations for guessing one of the AES key bytes, depending on the amount of measurements. The correct key (blue) becomes distinguishable at around 65,000 measurements (Color figure online)

except one of them are recovered. For most of the key bytes, only tens of thousands of measurements is suffice to recover the correct key byte (Fig. 13). The non-optimized implementation of this attack processes 2 million observations in 5 min.

5.3 Attack 3: IPP Cache Protected SM4

SM4 (formerly SMS4) is a block cipher standardized by the Chinese government and the standard encryption for Wireless LAN Wired Authentication and Privacy Infrastructure (WAPI) [24]. SM4 features an unbalanced Feistel structure and supports 128-bit blocks and keys. SM4 is known to be secure and no relevant cryptanalytic attacks exist for the cipher. Figure 14 shows a schematic of one round of SM4. T1–T4 are 4×32-bit state variables of SM4. Within each round, the last three state variables and a 32-bit round key are mixed, and each byte of the output will be replaced by a non-linear S-Box value. After the non-linear layer, the combined 32-bit output of

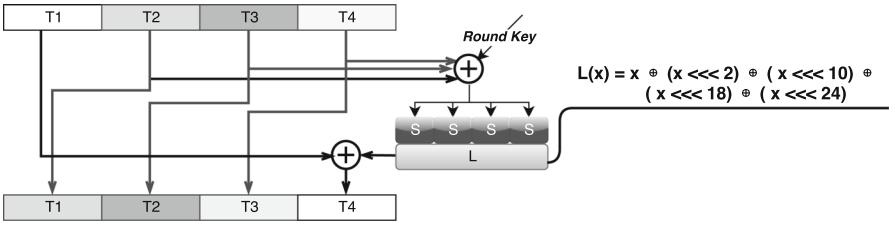

Fig. 14 SM4 Feistel structure: in each round, the last three words from the state buffer and the round key will be added. Each byte of the output will be replaced by S-Box lookup. The function L performs a linear bit permutation

S-Boxes x are diffused using the linear function L. The output of L is then mixed with the first 32-bit state variable to generate a new random 32-bit state value. The same operation is repeated for 32 rounds, and each time a new 32-bit state is generated as the next round T4 state. The current T2, T3, T4 are treated as T1, T2, and T3 for the next round. The final 16 bytes of the entire state after the last round produce the ciphertext. SM4 Key schedule produces 32×32-bit round keys from a 128-bit key. Since the key schedule is reversible, recovering 4 repeated round keys provides enough entropy to reproduce the cipher key.

All the SM4 operations except the S-Box lookup are performed in 32-bit word sizes. Hence, SM4 implementation is both simple and efficient on modern architectures. We chose the function `cpSMS4_Cipher` from Intel IPP Cryptography library. Our target is based on the straightforward cipher algorithm with addition of S-Box cache state normalization. We recovered this implementation through reverse engineering of Intel IPP binaries. The implementation preloads four values from different cache lines of S-Box before the first round, and it mixes them with some dummy variables, forcing the processor to fill the relevant cache lines with S-Box table. This cache prefetching mechanism protects SM4 against asynchronous cache attacks. On our experimental setup, the implementation runs in about 700 cycles, which informs us that this implementation maintain a high speed while secure against asynchronous attacks. Interrupted attacks that leak intermediate states would not be as simple, since the interruption need to happen faster than 700 cycles. We will further discuss the difficulty of correlating any cache-granular information, even if we assume the adversary can interrupt the encryption and perform some intermediate observations.

$$
\begin{aligned}
x_{32} &= c_1 \oplus c_2 \oplus c_3 \oplus k_{32} \\
d_2 &= c_1, d_3 = c_2, d_4 = c_3 \\
d_1 &= L\left(s\left(x_{32}^1\right), s\left(x_{32}^2\right), s\left(x_{32}^3\right), s\left(x_{32}^4\right)\right) \oplus c_4 \\
x_{31} &= d_1 \oplus d_2 \oplus d_3 \oplus k_{31} \\
e_2 &= d_1, e_3 = d_2, e_4 = d_3 \\
e_1 &= L\left(s\left(x_{31}^1\right), s\left(x_{31}^2\right), s\left(x_{31}^3\right), s\left(x_{31}^4\right)\right) \oplus d_4 r \\
x_{30} &= e_1 \oplus e_2 \oplus e_3 \oplus k_{30} \\
f_2 &= e_1, f_3 = e_2, f_4 = e_3 \\
f_1 &= L\left(s\left(x_{30}^1\right), s\left(x_{30}^2\right), s\left(x_{30}^3\right), s\left(x_{30}^4\right)\right) \oplus e_4 \\
x_{29} &= f_1 \oplus f_2 \oplus f_3 \oplus k_{29} \\
g_2 &= f_1, g_3 = f_2, g_4 = f_3 \\
g_1 &= L\left(s\left(x_{29}^1\right), s\left(x_{29}^2\right), s\left(x_{29}^3\right), s\left(x_{29}^4\right)\right) \oplus f_4 \\
x_{28} &= g_1 \oplus g_2 \oplus g_3 \oplus k_{28}
\end{aligned}
\tag{1}
$$

Single-round attack on SM4 We define c_1, c_2, c_3, c_4 as the four 32-bit words of a ciphertext and k_r as the secret round key for round r. We recursively follow the cipher structure from the last round with our ciphertext words as inputs, and write the last 5 rounds' relations as Eq. 1. In each round, x_r^i is the S-Box index, and i is the byte offset of the 32-bit word x_r. With a similar approach to the attack on AES, we define matrix **A** for the access profile, where each row corresponds to a known ciphertext, and each column indicates the number of accesses when $x_r^i < 4$. Then we define the matrix **L** for the observed timing leakage and the correlation between **A** and **L** similar to the AES attack. In contrast, S-Box indices in the AES attack are defined based on a non-linear inverse S-Box operation of key and ciphertext, which eventually maps to all possible key candidates. In SM4, the index x_r^i is defined before any non-linear operation. As a result, an attack capable of distinguishing accesses of 4 out of 256 S-Box entries reveals only 6 bits per key byte. In the mentioned relations, performing the attack using this model on x_{32}^i, recovers the 6 most significant bits of each key byte i for the last round key (Total of 24 out of the 32 bits).

Multi-round attack on SM4 The relationship for round 31 can be used not only to recover 6-bit key candidates of round 31, but also the remaining unknown 8 bits of entropy for round 32. This is due to the linear property of function L and the recursive nature of newly created state variables. After the attack on round 32, similar to the round key, we only have certainty about 24 bits of the new state variable d_1, but this information will be propagated as the input to round 31. The next round of attack for key byte of round 31 needs more computation to process an 8 bit of unknown key and 8 bit of unknown state (total of 16 bit), but this is computationally feasible, and the 8-bit key from round 32 with highest correlation can be recovered by attacking the S-Box indices in round 31. We recursively applied this model to each round resulting a correlation attack with the following steps, which gives us enough entropy to recover the key:

1. $x_{32} \rightarrow$ 24 bits of k_{32}.
2. $x_{31} \rightarrow$ 24 bits of k_{31} + 8 bits of k_{32}
3. $x_{30} \rightarrow$ 24 bits of k_{30} + 8 bits of k_{31}
4. $x_{29} \rightarrow$ 24 bits of k_{29} + 8 bits of k_{30}
5. $x_{28} \rightarrow$ 24 bits of k_{28} + 8 bits of k_{29}
6. Recover the key from $k_{32}, k_{31}, k_{30}, k_{29}$

SM4 key recovery results on synthetic data Our noise-free synthetic data shows that 3000 observations are enough to find all correct 6-bit and 8-bit round key candidates with the highest correlations. Even in an interrupted cache attack or without cache protection, targeting this implementation using a cache-granular information would be much harder and inefficient due to the lack of intra cache-line resolution. If we only distinguish the 64-byte cache lines out of a 256-byte S-Box, we only learn 4×2-bit (total of 8 bits) out of 32-bit round keys, and on each round, we need to solve 8 bits + 24 bits of uncertainty. Although solving 32-bit of uncertainty sounds possible for a noise-free data, it is computationally much harder in a practical noisy setting. Our intra cache line leakage can exploit SM4 efficiently in a known-ciphertext scenario, while the best efficient cache attack on SM4 requires chosen plaintexts [57].

SM4 key recovery results using *MemJam* The results on SM4 show even more effective key recovery against this implementation compared to AES. Figure 15 shows the correlation rate over measurements for one key byte in the first round of attack which 13,000 measurements is suffice to distinguish the correct 6-bit round key (blue) for this key byte. Figure 16 shows the correlation for 6-bit round keys after 5 rounds of repeated attack, and the correlation for 12-bit key candidates can be seen in Fig. 17. The attack expects assurance on the correct key candidates for each round of attack before proceeding to the next round due to the recursive structure of SM4. In our experiment using real measurement data, we have noticed that 40,000 observations are sufficient to have assurance of correct key candidates with the highest correlations. Our implementation of the attack can recover the correct 6-bit and 8-bit keys, and it takes about 5 min to recover the cipher key. In Fig. 17, we plotted the accumulated per byte correlations for all 8-bit candidates within each round of attack. During the computation of 6-bit candidates, the 8-bit candidates relate to 4 different state bytes. This accumulation greatly increases the result and the correct 8-bit key candidates have a very high aggregated correlation compared to the 6-bit candidates.

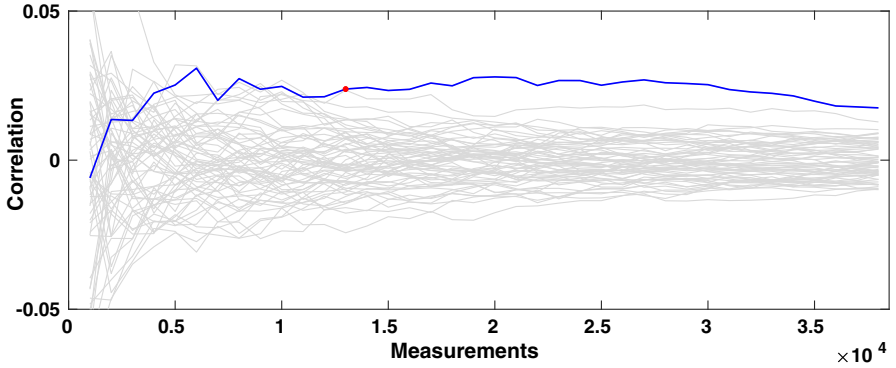

Fig. 15 The timing correlations for guessing one of the SM4 key bytes in a single round attack, depending on the amount of measurements. The correct key (blue) becomes distinguishable at around 13,000 measurements (Color figure online)

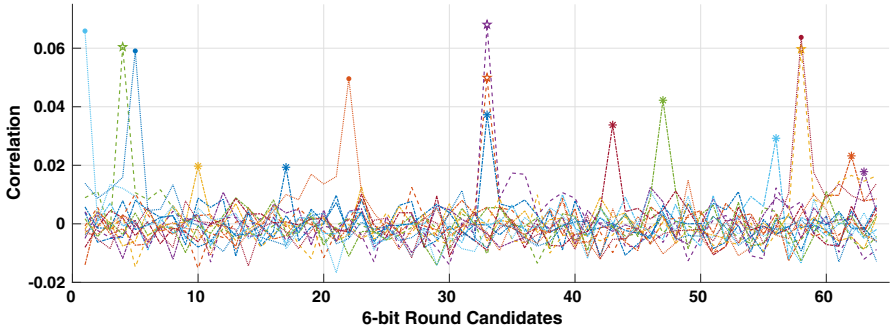

Fig. 16 Correlations for SM4 6-bit keys of the last 4 32-bit round key recovered through 5 rounds of attack using 40,000 observations

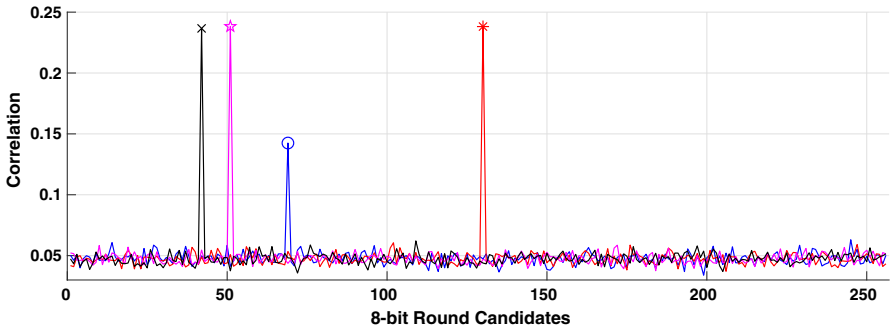

Fig. 17 The accumulated correlations for SM4 8-bit keys after 5 rounds using 40,000 observations. Each correct candidate has the highest correlation

6 *MemJam*ing SGX Enclave

Intel SGX is a trusted execution environment (TEE) extension released as part of Skylake processor generation [38]. The main goal of SGX is to protect runtime data and computation from system and physical adversaries. Having said that, SGX must remain secure in the presence of malicious OS, thus modification of OS resources for facilitation of side-channel attacks is relevant and within the considered threat model. Previous works demonstrate high resolution attacks with 4 kB page [69,74] and 64 B cache line granularity [14,54]. Intel declared microarchitectural leakages out of scope for SGX, thus pushing the burden of writing leakage free constant-time code onto enclave developers. Indeed, Intel follows this design paradigm and ensures constant cache-line accesses for its AES implementation, making it resistant to *all* previously known microarchitectural attacks in SGX.

In this section, we verify that *MemJam* is also applicable to SGX enclaves, as there is no fundamental microarchitectural changes to resist against memory false dependencies. We repeat the key recovery results against Intel's constant-time AES implementation after moving it into an SGX enclave. The results verify the exploitability of intra cache level channels against SGX secure enclaves. In fact, the attack can be reproduced in a straightforward manner. The only difference is a slower key recovery due to the increased measurement noise resulting from the enclave context switch.

6.1 SGX Enclave Experimental Setup and Assumptions

Following the threat model of *CacheZoom* [50,54], we assume that the system adversary has control over various OS resources. Please note that SGX was exactly designed to thwart the threat of such adversaries. The adversary uses its OS-level privileges to decrease the setup noise: We isolate one of the physical cores from the rest of the running tasks, and dedicate its logical processors to *MemJam* write conflict thread and the victim enclave. We further disable all the non-maskable interrupts on the target physical core and configure the CPU power and frequency scaling to maintain a constant frequency. We assume that the adversary can measure the execution time of

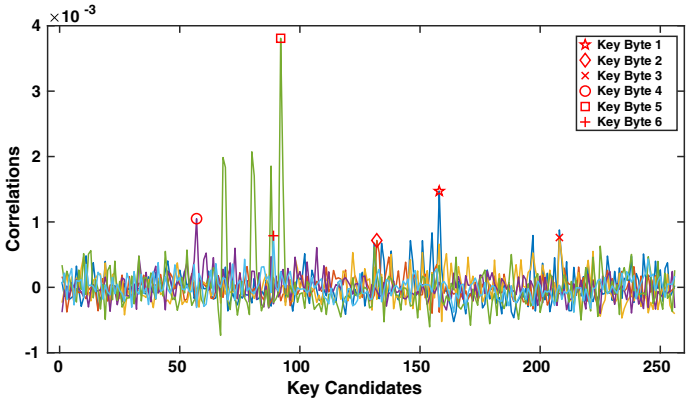

Fig. 18 Correlations for 6 key bytes using 5 million observations. All of the correct candidates have the highest correlations

an enclave interface that performs encryption, and the enclave interface only returns the ciphertext to the insecure environment. Both plaintexts and the secret encryption key are generated at runtime using *RDRAND* instruction, and they never leave the secure runtime environment of SGX enclave. The *RDTSC* instruction cannot be used inside an enclave. The attacker uses it right before the call to the enclave interface and again right after the enclave exit. As a result, the entire execution of the enclave interface, including the AES encryption, is measured. As before, an active thread causing read-after-write conflicts to the first four bytes of the AES S-Box is executed on the neighboring virtual processor of the SGX thread.

6.2 AES Key Recovery Results on SGX

Execution of the same AES encryption function as Sect. 5.2 inside an SGX enclave interface takes an average of 14,600 cycles with an active thread causing read-after-write conflicts to the first four bytes of the AES S-Box. The additional overhead is caused by the enclave context switch, which significantly increases the noise of the timing channel due to the variable timing behavior. Having that, this experiment shows a more practical timing behavior where adversaries cannot time the exact encryption operation, and they have to measure the time for a batch of operations. This not only shows that SGX is vulnerable to the *MemJam* attack, but it also demonstrates that *MemJam* is applicable in a realistic scenario. Figure 18 shows the key correlation results using 50 million timed encryptions in SGX, collected in 10 different time frames. We filtered outliers, i.e. measurements with high noise by only considering samples that are in the range of 2000 cycles of the mean. Among the 50 million samples, 93% pass the filtering, and we only calculated the correlations for the remaining traces. Figure 19 shows that we can successfully recover 14 out of 16 key bytes, revealing sufficient information for key recovery after 20 million observations.

These results show that even cryptographic libraries designed by experts that are fully aware of current attacks and of the leakage behavior of the target device may

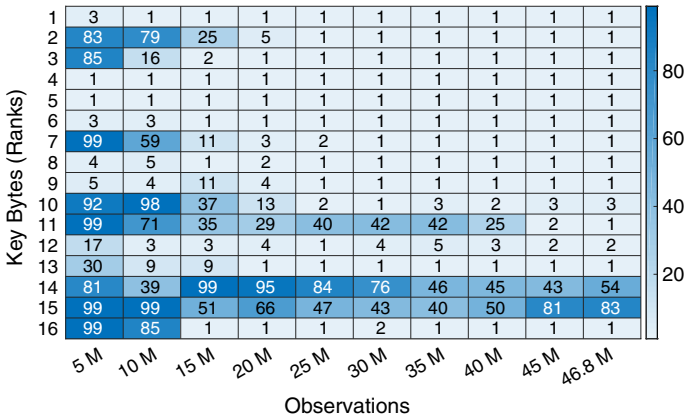

Fig. 19 The rank for correct key bytes with respect to the number of observations. Using the entire data set, after filtering the outliers, we can recover 14 out of 16 key bytes

fail at writing non-exploitable code. Modern microarchitectures are so complex that assumptions such as *constant cache line profiles* result in unexploitable constant-time implementations are seemingly impossible to fulfill.

7 Countermeasures to Memory Leakages

In this section, we focus on countermeasures that are relevant to memory-related side channels and their applicability to defend against *MemJam* attack. First, we discuss techniques to identify and protect weak software implementations. Then, we discuss proposed hardware defense mechanisms and attack detection methods.

7.1 Software Level Countermeasures

Constant-time implementation techniques are a known remedy to prevent memory leakage and have already seen some adoptions by researchers and practitioners. Indeed, the analyzed implementations of IPP all use some measures to prevent memory leakages, though only at cache-line granularity. `Cipher_DES` and `Safe2Encrypt_RIJ128` go further by ensuring an entirely uniform access pattern at cache-line granularity. Other crypto libraries have also adopted similar techniques, e.g. OpenSSL uses the scatter–gather technique [16] for their implementation of RSA [77].

Such hardened constant-time implementations are usually designed by experts who have knowledge of the underlying architecture and side-channel domain. Attacks such as *MemJam* show that uniform cache access pattern, cache state normalization [67] and scatter–gather technique [16] fail to protect cryptographic implementations. Bitsliced software implementations are secure against memory-related side channels and can be applied to cryptographic schemes such as DES and AES [10]. However, this limits the choice of efficient cryptographic schemes that are dependent on precomputed tables.

Table access masking is another technique that has been adopted by some of the cryptographic libraries to defend against intra-cache line leakage: The table lookup operation visits every single element of a table and uses an index mask to discard the irrelevant values.

Researchers have also proposed tools to automate the generation of code lacking memory leakages: `Raccoon` [60] enforces constant-time control flow, but stops at cache-line granularity and makes use of ORAM, which can be very costly. `Escort` [61] and `EncLang` [65] also transform code to constant-time representation in compilation phase. `EncLang` stops at page-level granularity and requires adoption of a new programming language. Escort is not focused on efficient protection against memory side channels and only focuses on arithmetic operations. That is, while these tools address memory leakages, they would need further fine-tuning to also address high-resolution attacks like *MemJam* with 4-byte spatial granularity. Limiting the resolution to cache-level granularity still leaves the door open for CacheBleed [77] and *MemJam* attacks. An alternative to generating robust code is to just ensure that code does not feature memory leakage by using analysis tools to verify constant-time properties: `MASCAT` [46] is a static code analysis tool, and `CacheD` [70] is a dynamic symbolic execution analyzer to detect cache leakages in software implementations. On the same direction, Langley's `ctgrind` and `ct-verif` [7] propose compiler-level verification techniques. Although these identification techniques can be extended to support an intra-cache line leakage model, there is only one proposal that practically considers this sensitive leakage model [25].

7.2 Hardware Level Countermeasures

Known Hardware solutions to defend against cache attacks generally ignore leakages through false dependency. Relaxed inclusion cache is a secure counterpart to the inclusive LLC which only aims to defend LLC contention [47]. Solutions such as CacheBar [80], Catalyst [51] and vCat [73] which isolate the LLC between different security domains cannot be scaled to thwart the *MemJam* attack, which exploits leakage in the L1 cache. Sanctum [21] is a secure processor design that uses page coloring to isolate cache. Further, they flush the L1 and TLB cache during context switch from/to secure enclaves. However, the effect of hyper-threading and false dependency has not been covered in such a design. A temporary workaround to defend against this attack is to disable hyper-threading. Ozone [9], as a zero timing leakage processor, aims to defend against such leakages by allocating a constant computational resource to one execution thread per core ignoring the hyper-threading model.

7.3 Attack Detection

Another approach is to detect attacks while they are happening and then react accordingly. Proposed methods to detect cache attacks at runtime utilize hardware performance counters [17,78] and transactional synchronization extensions (TSX) [20,64] to detect abnormal microarchitectural behavior. Defense mechanisms based on performance counters that monitor cache activities such as the number of cache

misses are incapable of detecting *MemJam*, as the attack does not introduce any irregular cache activity. Although one might argue using other performance counters for detection, its practicality is debatable. A monitoring agent needs to occupy an active thread and actively evaluate the number of memory read stalls. In our experiments, performing 50 million observations takes less than a minute. If such a detector exists, it has to monitor with a higher frequency than the attack, otherwise, it will be outperformed before any detection of suspicious behavior is possible. The effect of read-after-write hazards to the TSX has not been explored. However, we believe using TSX as a detector with low false-positives would not be practical, since the read-after-write hazards are common phenomena and TSX could fail due to other issues [39]. In the SGX world, detection of unexpected interrupts, as proposed in the literature [20], does not apply to *MemJam*.

7.4 Preventing *MemJam*

The `Cipher_DES` and `Safe2Encrypt_RIJ128` implementations have been designed to achieve a constant cache access profile by ensuring that the same cache lines are accessed every time regardless of the processed data. The 4-byte spatial resolution of *MemJam*, however, thwarts this countermeasure by providing intra cache-line resolution. One approach to restore security and protect against *MemJam* is to apply constant memory accesses with a 4-byte granularity. That would require accessing every fourth byte of the table for each memory lookup for the purpose of maintaining a uniform memory footprint. At that point, it might be easier to just do a *true* constant time implementation and access *all* entries each time, resting assured that there is no other effect somewhere hidden in the microarchitecture resulting in a leak with byte granularity.

The best remedy are hardware based implementations, e.g., AES-NI or hardware assisted implementations, e.g., SIMD-based bit-sliced implementations of AES or SM4. If available, such performant, yet constant-time instruction set extensions should exclusively be used to protect the targeted implementation in an efficient manner. For ciphers where such hardware support is not available, a true constant-time implementation e.g. based on bit-slicing seems to be the best, albeit slow, alternative. Intel IPP has different variants optimized for various generations of Intel instruction sets [43]. Intel IPP features different implementations of AES as well as SM4 in these variants. A list of these variants and implementations are given in Table 1. As shown, the software-only variant of each of the analyzed ciphers is vulnerable to *MemJam*.

8 Applicability of *MemJam*

MemJam exploits false dependencies of memory read-after-writes (4K Aliasing), which was turned into a cache-based timing attack with a 4-byte spatial resolution. This makes *MemJam* similar to *CacheBleed*, which also provides a 4 byte granularity [77]. Consequently, any countermeasures aimed at providing uniform accesses at cache-line granularity do not work, as discussed in Sect. 7. For *MemJam* to work, the

Table 1 DES, SM4 and AES implementations in all variants of Intel IPP library version 2018 [43]

Implementation	Function name	l9 n0 y8 k0 e9	m7 mx	n8	Linux SGX SDK
DES constant-time	Cipher_DES	✓	✓	✓	N/A
AES-NI	Encrypt_RIJ128_AES_NI	✓	×	×	✓ (prebuilt)
AES Bitsliced	SafeEncrypt_RIJ128	✓	×	✓	✓ (prebuilt)
AES constant-time	Safe2Encrypt_RIJ128	×	✓	×	✓ (source)
SM4 Bitsliced & AES-NI	cpSMS4_ECB_aesni	✓	×	×	N/A
SM4 cache normalized	cpSMS4_Cipher	✓	✓	✓	N/A

The variants will be merged at linker stage, each variant is optimized for a different generation of the Intel instruction set [37]. Developers can statically link specific variants with single processor static linking mode [43]

Table 2 Intel processor families and availability of the leakage channels. Major Intel processors suffer from 4K Aliasing, and are vulnerable to *MemJam* [5]

Release	Family	Cache bank conflicts	4K Aliasing
2006	Core	✓	✓
2008	Nehalem	×	✓
2011	Sandy bridge	✓	✓
2013	Silvermont, Haswell, Broadwell	×	✓
2015	Skylake	×	✓
2016	KabyLake	×	✓

false dependencies need to have an impact on the read after a false conflicting write. Table 2 highlights the availability of the cache bank conflicts and the 4K Aliasing leakage source: While bank conflicts are limited to few CPU generations, excluding all supporting SGX, 4K Aliasing is present in all Intel CPUs released in the last 10 years. Thus, *MemJam* applies to virtually all Intel CPUs that feature hyperthreading.

An adversary performing the *MemJam* attack also does not need to know about the offset of an S-Box in the binary, since she can simply scan the 10-bits address entropy by introducing conflicts to different offsets and measuring the timing of the victim. In such a scenario, we assume that the S-Box table is aligned with the cache line size, since an unaligned S-Box in memory is already vulnerable to cache attacks [44,58]. During the processing of an uniformly random input, each S-Box operation in an implementation such as `Safe2Encrypt_RIJ128` accesses the first word column of the table with a probability of $1/16$. Among 160 S-Box operations, an average of 10 memory accesses to the first S-Box is likely. While an attacker is causing RaW conflicts on increasing offsets, she can locate the S-Box offset as soon as she sees a timing behavior. This is important when it comes to obfuscated binaries or scenarios, where the offset of the S-Box is unknown.

As shown in Table 1, all block cipher implementations of IPP feature at least one vulnerable variant. In cases where there is an implementation based on the AES-NI instruction set (or SSSE3 respectively), the library falls back to the basic version at runtime if the instruction set extensions are not available. The usability of this depends

on the compilation and runtime configuration. Developers are allowed to statically link to a more risky variant [37], and they need to assure not to use the vulnerable versions during linking. These ciphers should be avoided in cases where the hardware does not provide support, e.g., Core and Nehalem do not support AES-NI; also AES-NI can be disabled in some BIOS. For Triple DES, IPP gives only one implementation option: the vulnerable one studied in this work. Thus, for applications that demand the use of Triple DES (and there are still many such applications, as discussed in Sect. 5.1), there is no secure alternative available in IPP. This highlights that current hardware support for cryptographic primitives is restricted and if any cipher without explicit hardware support is required, this limitation may endanger the provided security. *MemJam* is another piece of evidence that modern microarchitectures are too complex and constant-time implementations cannot simply be trusted, as assumptions about the underlying system often turn out to be wrong.

9 Conclusion

This work proposes *MemJam*, a new side-channel attack based on false dependencies. For the first time, we discovered new aspects of this side channel and its capabilities, and show how to extract secrets from modern cryptographic implementations. *MemJam* uses false read-after-write dependencies to slow down accesses of the victim to a particular 4-byte memory block *within* a cache line. The resulting latency of otherwise constant-time implementations was exploited with state-of-the art timing side-channel analysis techniques. We showed how to apply the attack to recent implementations of Triple DES, AES and SM4, as found in Intel IPP. According to the available resources, the source of the leakage for the *MemJam* attack is present in all Intel CPU families released in the last 10 years [5,39], including newest generation CPUs. Our results also show that *MemJam* is a viable intra cache level attack applicable to SGX enclaves. Prior to *MemJam*, it might have seemed reasonable to design SGX enclaves under the paradigm that constant cache line accesses result in leakage-free code. However, the increased 4-byte intra cache-line granularity of *MemJam* shows that only code with true constant-time properties, i.e. constant execution flow and constant memory accesses can be expected to have no remaining leakage on modern microarchitectures.

Acknowledgements This work is supported by the National Science Foundation, under Grant CNS-1618837.

Compliance with ethical standards

Responsible disclosure We have informed the Intel Product Security Incident Response Team of our findings. They have acknowledged the receipt and confirmed a work-in-progress patch for IPP library. Here is the time line for the responsible disclosure process: (1) **08/02/2017:** We informed our findings to the Intel Product Security Incident Response Team (Intel PSIRT). (2) **08/04/2017:** Intel PSIRT acknowledged the receipt. (3) **11/07/2017:** `Safe2Encrypt_RIJ128` was removed from the SGX SDK. (4) **11/17/2017:** Intel PSIRT confirmed a work-in-progress patch for IPP library (CVE-2017-5737). (5) **05/10/2018:** Intel PSIRT published an update for IPP library (CVE-2018-3691).

References

1. Acıiçmez, O., Brumley, B.B., Grabher, P.: New results on instruction cache attacks. In: International Workshop on Cryptographic Hardware and Embedded Systems. Springer (2010)
2. Acıiçmez, O., Gueron, S., Seifert, J.P.: New Branch Prediction Vulnerabilities in OpenSSL and Necessary Software Countermeasures. In: Galbraith, S.D. (eds.) Cryptography and Coding. Cryptography and Coding 2007. Lecture Notes in Computer Science, vol. 4887, pp. 185–203. Springer, Berlin, Heidelberg (2007)
3. Acıiçmez, O., Koç, Ç.K., Seifert, J.P.: Predicting secret keys via branch prediction. In: Cryptographers Track at the RSA Conference. Springer (2007)
4. Aciicmez, O., Seifert, J.P.: Cheap hardware parallelism implies cheap security. In: FDTC 2007. Workshop on Fault Diagnosis and Tolerance in Cryptography, 2007. IEEE (2007)
5. Agner: The microarchitecture of Intel, AMD and VIA CPUs: An optimization guide for assembly programmers and compiler makers. http://www.agner.org/optimize/microarchitecture.pdf
6. Allan, T., Brumley, B.B., Falkner, K., van de Pol, J., Yarom, Y.: Amplifying side channels through performance degradation. In: Annual Computer Security Applications Conference (ACSAC) (2016)
7. Almeida, J.B., Barbosa, M., Barthe, G., Dupressoir, F., Emmi, M.: Verifying constant-time implementations. In: USENIX Security Symposium, pp. 53–70 (2016)
8. Andrysco, M., Kohlbrenner, D., Mowery, K., Jhala, R., Lerner, S., Shacham, H.: On subnormal floating point and abnormal timing. In: 2015 IEEE Symposium on Security and Privacy (SP). IEEE (2015)
9. Aweke, Z.B., Austin, T.: Ozone: efficient execution with zero timing leakage for modern microarchitectures. Preprint. arXiv:1703.07706 (2017)
10. BearSSL: BearSSL constant-time crypto. https://www.bearssl.org/constanttime.html
11. Benger, N., Van De Pol, J., Smart, N.P., Yarom, Y.: Ooh Aah... just a little bit: a small amount of side channel can go a long way. In: International Workshop on Cryptographic Hardware and Embedded Systems. Springer (2014)
12. Bhargavan, K., Leurent, G.: On the practical (in-) security of 64-bit block ciphers: collision attacks on HTTP over TLS and OpenVPN. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security, pp. 456–467. ACM (2016)
13. Bonneau, J., Mironov, I.: Cache-collision timing attacks against AES. In: International Workshop on Cryptographic Hardware and Embedded Systems. Springer (2006)
14. Brasser, F., Müller, U., Dmitrienko, A., Kostiainen, K., Capkun, S., Sadeghi, A.R.: Software grand exposure: SGX cache attacks are practical. In: 11th USENIX Workshop on Offensive Technologies (WOOT 17). USENIX Association, Vancouver (2017). https://www.usenix.org/conference/woot17/workshop-program/presentation/brasser
15. Brickell, E., Graunke, G., Neve, M., Seifert, J.P.: Software mitigations to hedge AES against cache-based software side channel vulnerabilities. In: IACR Cryptology ePrint Archive (2006)
16. Brickell, E., Graunke, G., Seifert, J.P.: Mitigating cache/timing based side-channels in AES and RSA software implementations. In: RSA Conference 2006 Session DEV-203 (2006)
17. Briongos, S., Irazoqui, G., Malagón, P., Eisenbarth, T.: CacheShield: protecting legacy processes against cache attacks. Preprint. arXiv:1709.01795 (2017)
18. Brumley, D., Boneh, D.: Remote timing attacks are practical. Comput. Netw. **48**(5), 701–716 (2005)
19. Carluccio, D.: Electromagnetic side channel analysis for embedded crypto devices. Master's Thesis, Ruhr Universität Bochum (2005)
20. Chen, S., Zhang, X., Reiter, M.K., Zhang, Y.: Detecting privileged side-channel attacks in shielded execution with Déjá Vu. In: Proceedings of the 2017 ACM on Asia Conference on Computer and Communications Security. ACM (2017)
21. Costan, V., Lebedev, I.A., Devadas, S.: Sanctum: minimal hardware extensions for strong software isolation. In: USENIX Security Symposium (2016)
22. Daemen, J., Rijmen, V.: The Design of Rijndael: AES—The Advanced Encryption Standard. Springer, Berlin (2013)
23. Dierks, T., Rescorla, E.: The transport layer security (TLS) protocol version 1.2. RFC 5246 (2008). https://www.ietf.org/rfc/rfc5246.txt
24. Diffie, W., Ledin, G.: SMS4 encryption algorithm for wireless networks. IACR Cryptology ePrint Archive (2008)

25. Doychev, G., Köpf, B.: Rigorous analysis of software countermeasures against cache attacks. In: Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation (2017)
26. EMVCo: EMVCo overview. https://www.emvco.com/about/overview/
27. EMVCo: Integrated Circuit Card Specifications for Payment Systems—Book 2: Security and Key Management, Version 4.3 (2011)
28. Ge, Q., Yarom, Y., Cock, D., Heiser, G.: A Survey of microarchitectural timing attacks and countermeasures on contemporary hardware. IACR Cryptology ePrint Archive 2016/613 (2016)
29. Ge, Q., Yarom, Y., Li, F., Heiser, G.: Contemporary processors are leaky–and there is nothing you can do about it. The Computing Research Repository (2016)
30. Glowacz, C., Grosso, V., Poussier, R., Schueth, J., Standaert, F.X.: Simpler and more efficient rank estimation for side-channel security assessment. In: International Workshop on Fast Software Encryption, pp. 117–129. Springer (2015)
31. Gruss, D., Maurice, C., Wagner, K., Mangard, S.: Flush+Flush: A Fast and Stealthy Cache Attack. In: Caballero, J., Zurutuza, U., Rodríguez, R. (eds.) Detection of Intrusions and Malware, and Vulnerability Assessment. DIMVA 2016. Lecture Notes in Computer Science, vol. 9721, pp. 279–299. Springer, Cham (2016)
32. Gueron, S., Krasnov, V.: SM4 acceleration processors, methods, systems, and instructions. US Patent 9,513,913 (2016). https://www.google.com/patents/US9513913
33. Gullasch, D., Bangerter, E., Krenn, S.: Cache games-bringing access-based cache attacks on AES to practice. In: 2011 IEEE Symposium on Security and Privacy (SP). IEEE (2011)
34. Hankerson, D., López Hernandez, J., Menezes, A.: Software Implementation of Elliptic Curve Cryptography over Binary Fields. In: Koç, Ç.K., Paar, C. (eds.) Cryptographic Hardware and Embedded Systems – CHES 2000. CHES 2000. Lecture Notes in Computer Science, vol. 1965, pp. 1–24 . Springer, Berlin, Heidelberg (2000)
35. Inci, M.S., Gülmezoglu, B., Apecechea, G.I., Eisenbarth, T., Sunar, B.: Seriously, get off my cloud! cross-VM RSA Key Recovery in a Public Cloud. IACR Cryptology ePrint Archive (2015)
36. Inci, M.S., Gulmezoglu, B., Irazoqui, G., Eisenbarth, T., Sunar, B.: Cache attacks enable bulk key recovery on the cloud. In: International Conference on Cryptographic Hardware and Embedded Systems. Springer (2016)
37. Intel IPP linkage models-quick reference guide. https://intel.ly/2tGjLCw
38. Intel: intel(R) software guard extensions for Linux* OS. https://github.com/01org/linux-sgx
39. Intel: Intel 64 and IA-32 architectures optimization reference manual. https://www.intel.com/content/www/us/en/architecture-and-technology/64-ia-32-architectures-optimization-manual.html
40. Intel: Intel 64 and IA-32 architectures software developer manuals. https://software.intel.com/en-us/articles/intel-sdm
41. Intel: Pin, dynamic binary instrumentation tool. https://software.intel.com/en-us/articles/pin-a-dynamic-binary-instrumentation-tool
42. Symmetric cryptography primitive functions. https://software.intel.com/en-us/ipp-crypto-reference-symmetric-cryptography-primitive-functions
43. Understanding CPU dispatching in the intel IPP libraries. https://intel.ly/2MxXkWY
44. Irazoqui, G., Cong, K., Guo, X., Khattri, H., Kanuparthi, A., Eisenbarth, T., Sunar, B.: Did we learn from LLC side channel attacks? a cache leakage detection tool for crypto libraries. Preprint. arXiv:1709.01552 (2017)
45. Irazoqui, G., Eisenbarth, T., Sunar, B.: S$A: a shared cache attack that works across cores and defies VM sandboxing—and its application to AES. In: 2015 IEEE Symposium on Security and Privacy (SP) (2015)
46. Irazoqui, G., Eisenbarth, T., Sunar, B.: MASCAT: stopping microarchitectural attacks before execution. IACR Cryptology ePrint Archive (2016)
47. Kayaalp, M., Khasawneh, K.N., Esfeden, H.A., Elwell, J., Abu-Ghazaleh, N., Ponomarev, D., Jaleel, A.: RIC: relaxed inclusion caches for mitigating LLC side-channel attacks. In: Proceedings of the 54th Annual Design Automation Conference 2017. ACM (2017)
48. Koç, C.K.: Analysis of sliding window techniques for exponentiation. Comput. Math. Appl. **30**(10), 17–24 (1995)
49. Kocher, P., Jaffe, J., Jun, B., Rohatgi, P.: Introduction to differential power analysis. J. Cryptogr. Eng. **1**(1), 5–27 (2011)

50. Lee, S., Shih, M.W., Gera, P., Kim, T., Kim, H., Peinado, M.: Inferring fine-grained control flow inside SGX enclaves with branch shadowing. Preprint. arXiv:1611.06952 (2016)
51. Liu, F., Ge, Q., Yarom, Y., Mckeen, F., Rozas, C., Heiser, G., Lee, R.B.: Catalyst: defeating last-level cache side channel attacks in cloud computing. In: 2016 IEEE Symposium on High Performance Computer Architecture (HPCA) (2016)
52. Marr, D., Binns, F., Hill, D., Hinton, G., Koufaty, D., et al.: Hyper-threading technology in the netburst® microarchitecture. 14th Hot Chips (2002)
53. Moghimi, A., Eisenbarth, T., Sunar, B.: MemJam: a false dependency attack against constant-time crypto implementations. In: CT-RSA 2018. Springer (2018). arXiv:1711.08002
54. Moghimi, A., Irazoqui, G., Eisenbarth, T.: Cachezoom: how SGX amplifies the power of cache attacks. Preprint. arXiv:1703.06986 (2017)
55. National Institute of Standards and Technology: Federal Information Processing Standards (FIPS) Publication 46-3—Data Encryption Standard (DES) (1999). https://csrc.nist.gov/csrc/media/publications/fips/46/3/archive/1999-10-25/documents/fips46-3.pdf
56. National Institute of Standards and Technology: Update to current use and deprecation of TDEA (2017). https://csrc.nist.gov/News/2017/Update-to-Current-Use-and-Deprecation-of-TDEA
57. Nguyen, P.H., Rebeiro, C., Mukhopadhyay, D., Wang, H.: Improved differential cache attacks on SMS4. In: International Conference on Information Security and Cryptology, pp. 29–45. Springer (2012)
58. Osvik, D.A., Shamir, A., Tromer, E.: Cache attacks and countermeasures: the case of AES. In: Cryptographers Track at the RSA Conference (2006)
59. Page, D.: Defending against cache-based side-channel attacks. Inf. Secur. Tech. Rep. **8**(1), 30–44 (2003). https://doi.org/10.1016/S1363-4127(03)00104-3. http://www.sciencedirect.com/science/article/pii/S1363412703001043
60. Rane, A., Lin, C., Tiwari, M.: Raccoon: closing digital side-channels through obfuscated execution. In: USENIX Security Symposium, pp. 431–446 (2015)
61. Rane, A., Lin, C., Tiwari, M.: Secure, precise, and fast floating-point operations on x86 processors. In: USENIX Security Symposium, pp. 71–86 (2016)
62. Ristenpart, T., Tromer, E., Shacham, H., Savage, S.: Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In: Proceedings of the 16th ACM Conference on Computer and Communications Security. ACM (2009)
63. Schimmel, C.: UNIX Systems for Modern Architectures: Symmetric Multiprocessing and Caching for Kernel Programmers. Addison-Wesley, Reading (1994)
64. Shih, M.W., Lee, S., Kim, T., Peinado, M.: T-SGX: eradicating controlled-channel attacks against enclave programs. In: Proceedings of the 2017 Annual Network and Distributed System Security Symposium (NDSS), San Diego (2017)
65. Sinha, R., Rajamani, S., Seshia, S.A.: A compiler and verifier for page access oblivious computation. Technical Report, Technical Report UCB/EECS-2017-124, EECS Department, University of California, Berkeley (2017)
66. Sullivan, D., Arias, O., Meade, T., Jin, Y.: Microarchitectural minefields: 4K-aliasing covert channel and multi-tenant detection in IaaS clouds (2018)
67. Tromer, E., Osvik, D.A., Shamir, A.: Efficient cache attacks on AES, and countermeasures. J. Cryptol. **23**(1), 37–71 (2010)
68. Tsunoo, Y., Saito, T., Suzaki, T., Shigeri, M., Miyauchi, H.: Cryptanalysis of DES implemented on computers with cache. In: International Workshop on Cryptographic Hardware and Embedded Systems. Springer (2003)
69. Van Bulck, J., Weichbrodt, N., Kapitza, R., Piessens, F., Strackx, R.: Telling your secrets without page faults: stealthy page table-based attacks on enclaved execution. In: Proceedings of the 26th USENIX Security Symposium. USENIX Association (2017)
70. Wang, S., Wang, P., Liu, X., Zhang, D., Wu, D.: CacheD: identifying cache-based timing channels in production software. In: 26th USENIX Security Symposium (USENIX Security 17), pp. 235–252. USENIX Association, Vancouver (2017). https://www.usenix.org/conference/usenixsecurity17/technical-sessions/presentation/wang-shuai
71. Webster, A., Tavares, S.E.: On the design of S-boxes. In: Advances in Cryptology-CRYPTO'85: Proceedings. Springer (1986)
72. Wolrich, G., Gopal, V., Yap, K., Feghali, W.: SMS4 acceleration processors, methods, systems, and instructions (2016). https://www.google.com/patents/US9361106. US Patent 9,361,106

73. Xu, M., Thi, L., Phan, X., Choi, H.Y., Lee, I.: vCAT: Dynamic cache management using CAT virtualization. In: Real-Time and Embedded Technology and Applications Symposium (RTAS), 2017 IEEE. IEEE (2017)
74. Xu, Y., Cui, W., Peinado, M.: Controlled-channel attacks: deterministic side channels for untrusted operating systems. In: 2015 IEEE Symposium on Security and Privacy (SP), pp. 640–656. IEEE (2015)
75. Yap, K., Wolrich, G., Satpathy, S., Gulley, S., Gopal, V., Mathew, S., Feghali, W.: SMS4 acceleration hardware. US Patent 9,503,256 (2016). https://www.google.com/patents/US9503256
76. Yarom, Y., Falkner, K.: FLUSH+RELOAD: a high resolution, low noise, L3 cache side-channel attack. In: USENIX Security (2014)
77. Yarom, Y., Genkin, D., Heninger, N.: CacheBleed: a timing attack on OpenSSL constant-time RSA. J. Cryptogr. Eng. **7**(2), 99–112 (2017)
78. Zhang, T., Zhang, Y., Lee, R.B.: Cloudradar: a real-time side-channel attack detection system in clouds. In: International Symposium on Research in Attacks, Intrusions, and Defenses. Springer (2016)
79. Zhang, Y., Juels, A., Reiter, M.K., Ristenpart, T.: Cross-VM side channels and their use to extract private keys. In: Proceedings of the 2012 ACM Conference on Computer and Communications Security. ACM (2012)
80. Zhou, Z., Reiter, M.K., Zhang, Y.: A software approach to defeating side channels in last-level caches. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. ACM (2016)

Affiliations

Ahmad Moghimi[1] · Jan Wichelmann[2] · Thomas Eisenbarth[2] · Berk Sunar[1]

Jan Wichelmann
j.wichelmann@uni-luebeck.de

Thomas Eisenbarth
thomas.eisenbarth@uni-luebeck.de

Berk Sunar
sunar@wpi.edu

[1] Worcester Polytechnic Institute, Worcester, MA, USA

[2] University of Lübeck, Lübeck, Germany