

CacheBleed: A Timing Attack on OpenSSL Constant Time RSA

Yuval Yarom · Daniel Genkin · Nadia Heninger

the date of receipt and acceptance should be inserted later

Abstract The scatter-gather technique is a commonly implemented approach to prevent cache-based timing attacks. In this paper we show that scatter-gather is not constant time. We implement a cache timing attack against the scatter-gather implementation used in the modular exponentiation routine in OpenSSL version 1.0.2f. Our attack exploits cache-bank conflicts on the Sandy Bridge microarchitecture. We have tested the attack on an Intel Xeon E5-2430 processor. For 4096-bit RSA our attack can fully recover the private key after observing 16,000 decryptions.

Keywords side-channel attacks, cache attacks, cryptographic implementations, constant-time, RSA

1 Introduction

1.1 Overview

Side-channel attacks are a powerful method for breaking theoretically secure cryptographic primitives. Since the first works by Kocher [37], these attacks have been used extensively to break the security of numerous cryptographic implementations. At a high level, it is possible to distinguish between two types of side-channel attacks, based on the methods used by the attacker: hardware-based attacks, which monitor the leakage through measurements (usually using dedicated lab equipment) of physical phenomena such as electromagnetic radiation [48], power consumption [35, 36], or acoustic emanation [25], and software-

based attacks, which do not require additional equipment but rely instead on the attacker software running on or interacting with the target machine. Examples of the latter include timing attacks which measure timing variations of cryptographic operations [9, 18, 19] and cache attacks which observe cache access patterns [45, 46, 54].

Percival [46] published in 2005 a *cache attack*, which targeted the OpenSSL [44] 0.9.7c implementation of RSA. In this attack, the attacker and the victim programs are collocated on the same machine and processor, and thus share the same processor cache. The attack exploits the structure of the processor cache by observing minute timing variations due to cache contention. The cache consists of fixed-size *cache lines*. When a program accesses a memory address, the cache-line-sized block of memory that contains this address is stored in the cache and is available for future use. The attack traces the changes that the victim program execution makes in the cache and, from this trace, the attacker is able to recover the private key used for the decryption.

In order to implement the modular exponentiation routine required for performing RSA public and secret key operations, OpenSSL 0.9.7c uses a sliding-window exponentiation algorithm [13]. This algorithm precomputes some values, called *multipliers*, which are used throughout the exponentiation. The access pattern to these precomputed multipliers depends on the exponent, which, in the case of decryption and digital signature operations, should be kept secret. Because each multiplier occupies a different set of cache lines, Percival [46] was able to identify the accessed multipliers and from that recover the private key. To mitigate this attack, Intel implemented a countermeasure that changes the memory layout of the precomputed multipliers. The countermeasure, often called *scatter-gather*, interleaves the multipliers in memory to ensure that the same cache lines are accessed irrespective of the multiplier used [16]. While this countermeasure ensures that the same cache lines are always

Yuval Yarom
University of Adelaide and Data61, CSIRO

Daniel Genkin
University of Pennsylvania and University of Maryland

Nadia Heninger
University of Pennsylvania

accessed, the offsets of the accessed addresses within these cache lines depend on the multiplier used and, ultimately, on the private key.

In this work we investigate two microarchitectural effects that allow an adversary to partially recover offsets within accessed cache lines.

To facilitate concurrent access to the cache, the cache is often divided into multiple *cache banks*. While concurrent accesses to different cache banks can always be handled, on some processor models (such as Intel Sandy Bridge and Ivy Bridge microarchitectures) each cache bank can only handle a limited number of concurrent requests. Often a single request at a time. A *cache-bank conflict* occurs when too many requests are made concurrently to the same cache bank. In the case of a conflict, some of the conflicting requests are delayed. Both Bernstein [9] and Osvik et al. [45] have warned that accesses to different offsets within cache lines may leak information through timing variations due to cache-bank conflicts. While timing variations due to cache-bank conflicts are documented in the Intel Optimization Manual [32], no attack exploiting these has ever been published.

The second effect is a false dependency between read and write operations, which prevents simultaneous read and write operations from addresses that are spaced by a multiple of 4,096 bytes [22]. Bernstein and Schwabe [10] demonstrate that timing differences based on this false dependency can be measured within the same process. Yet, as in the case of cache bank conflicts, no attack exploiting this effect has ever been published.

In the absence of a demonstrated risk, Intel continued to contribute code that uses scatter-gather to OpenSSL [26, 27] and to recommend the use of the technique for side-channel mitigation [14, 15]. Consequently, the technique is in widespread use in the current versions of OpenSSL and its forks, such as LibreSSL [39] and BoringSSL [12]. It is also used in other cryptographic libraries, such as the Mozilla Network Security Services (NSS) [42].

1.2 Our Contribution

In this work we present CacheBleed, the first side-channel attack to systematically recover sub-cache-line information. In Section 3 we describe two variants of CacheBleed. The first creates contention on a cache bank and measures the timing variations due to conflicts and the second creates false dependencies and measures the timing variations they induce. Next, in Section 4 we use the cache-bank conflicts variant of CacheBleed in order to attack the scatter-gather implementation of OpenSSL’s modular exponentiation routine. Finally, in Section 5 we discuss the results of applying the false-dependency variant of CacheBleed to OpenSSL, showing non-exploitable information leaks through this channel.

Using the cache-bank conflicts variant of CacheBleed we are able to recover 60% of the secret exponent bits after observing 16,000 RSA decrypts or signing operations. To find the remaining bits we adapt the Heninger-Shacham algorithm [28] for the information we collect with CacheBleed. In order to achieve full key extraction, our attack requires about two CPU hours. Parallelizing across multiple CPUs, we achieved key extraction in only a few minutes. See Section 6 for a more complete discussion.

1.3 Targeted Software and Hardware

Software. In this paper we target the modular exponentiation operation as implemented in OpenSSL version 1.0.2f which was the latest version of OpenSSL prior to our disclosure to OpenSSL. As mentioned above, similar (and thus potentially vulnerable) code can be found in several forks of OpenSSL such as LibreSSL [39] and BoringSSL [12]. Other cryptographic libraries, such as the Mozilla Network Security Services (NSS) [42] use similar techniques and may be vulnerable as well.

Hardware. Our attacks exploit cache-bank conflicts present in Intel Sandy Bridge Processor family, and false dependencies present in Intel Haswell and Skylake processors, both of which do not experience cache-bank conflicts. We used three platforms for our experiments. These are summarized in Table 1. We exploited the cache-bank conflicts on the Sandy Bridge machine and the false dependency on the Haswell and the Skylake machines.

Table 1 Specifications of tested platforms.

	Sandy Br.	Haswell	Skylake
Processor	Xeon E5-2430	Xeon E5-2618L	i7-6770HQ
Cores	6	8	4
Frequency	2.20 GHz	2.30 GHz	2.60 GHz
OS	CentOS 6.7	CentOS 6.7	CentOS 7.2

Disclosure and Mitigation. We have reported our results to the developers of OpenSSL, LibreSSL, NSS, and BoringSSL. We worked with the OpenSSL developers to evaluate and deploy countermeasures to prevent the attacks described in this paper (CVE-2016-0702). These countermeasures were subsequently incorporated into OpenSSL 1.0.2g and BoringSSL. A fix for LibreSSL is included in the current development release (Version 2.5.0). For NSS, our attack was documented under Mozilla bug 1252035 and a countermeasure addressing potential exploitation was included in version 3.24.

2 Background

2.1 OpenSSL’s RSA Implementation

RSA [49] is a public-key cryptosystem which supports both encryption and digital signatures. To generate an RSA key pair, the user generates two prime numbers p, q and computes $N = pq$. Next, given a public exponent e (OpenSSL uses $e = 65537$), the user computes the secret exponent $d \equiv e^{-1} \pmod{\phi(N)}$. The public key is the integers e and N and the secret key is d and N . In textbook RSA encryption, a message m is encrypted by computing $m^e \pmod{N}$ and a ciphertext c is decrypted by computing $c^d \pmod{N}$.

RSA-CRT. RSA decryption is often implemented using the Chinese remainder theorem (CRT), which provides a speedup over exponentiation mod N . Instead of computing $c^d \pmod{N}$ directly, RSA-CRT splits the secret key d into two parts $d_p = d \pmod{p-1}$ and $d_q = d \pmod{q-1}$, and then computes two parts of the message as $m_p = c^{d_p} \pmod{p}$ and $m_q = c^{d_q} \pmod{q}$. The message m can then be recovered from m_p and m_q using Garner’s formula [23]:

$$h = (m_p - m_q)(q^{-1} \pmod{p}) \pmod{p}$$

and

$$m = m_q + hq.$$

Exponentiation algorithms. The main operation performed during RSA decryption is the modular exponentiation, that is, calculating $a^b \pmod{k}$ for some secret exponent b . Several algorithms for modular exponentiation have been suggested. In this work we are interested in the two algorithms that OpenSSL has used.

Fixed-Window Exponentiation. In the *fixed-window exponentiation* algorithm, also known as *m-ary exponentiation*, the n -bit exponent b is represented as an $\lceil n/w \rceil$ digit integer in base 2^w for some chosen *window size* w . That is, b is rewritten as $b = \sum_{i=0}^{\lceil n/w \rceil - 1} 2^{wi} \cdot b_i$ where $0 \leq b_i < 2^w$. The pseudocode in [Algorithm 1](#) demonstrates the fixed-window exponentiation algorithm. In the first step, the algorithm precomputes a set of *multipliers* $a_j = a^j \pmod{k}$ for $0 \leq j < 2^w$. It then scans the base 2^w representation of b from the most significant digit ($b_{\lceil n/w \rceil - 1}$) to the least significant (b_0). For each digit b_i it squares an intermediate result w times and then multiplies the intermediate result by a_{b_i} . Each of the square or multiply operations is followed by a modular reduction.

Sliding-Window Exponentiation. The *sliding-window algorithm* represents the exponent b as a sequence of digits b_i such that $b = \sum_{i=0}^{n-1} 2^i \cdot b_i$, with b_i being either 0 or an odd number $0 < b_i < 2^w$. The algorithm first precomputes $a_1, a_3, \dots, a_{2^w-1}$ as in the fixed-window case. It then scans the exponent from the most significant to the least significant

Algorithm 1: Fixed-window exponentiation

```

input : window size  $w$ , base  $a$ , modulus  $k$ ,  $n$ -bit exponent
           $b = \sum_{i=0}^{\lceil n/w \rceil} 2^{wi} \cdot b_i$ 
output:  $a^b \pmod{k}$ 

//Precomputation
 $a_0 \leftarrow 1$ 
for  $j = 1, \dots, 2^w - 1$  do
  |  $a_j \leftarrow a_{j-1} \cdot a \pmod{k}$ 
end

//Exponentiation
 $r \leftarrow 1$ 
for  $i = \lceil n/w \rceil - 1, \dots, 0$  do
  | for  $j = 1, \dots, w$  do
    |  $r \leftarrow r^2 \pmod{k}$ 
  | end
  |  $r \leftarrow r \cdot a_{b_i} \pmod{k}$ 
end
return  $r$ 

```

digit. For each digit, the algorithm squares the intermediate result. For non-zero digit b_i , it also multiplies the intermediate result by a_{b_i} .

The main advantages of the sliding-window algorithm over the fixed-window algorithm are that, for the same window size, sliding window needs to precompute half the number of multipliers, and that fewer multiplications are required during the exponentiation. The sliding-window algorithm, however, leaks the position of the non-zero multipliers to adversaries who can distinguish between squaring and multiplication operations. Furthermore, the number of squaring operations between consecutive multipliers may leak the values of some zero bits. Up to version 0.9.7c, OpenSSL used sliding-window exponentiation. As part of the mitigation of the Percival [46] cache attack, which exploits these leaks, OpenSSL changed their implementation to use the fixed-window exponentiation algorithm.

Since both algorithms precompute a set of multipliers and access them throughout the exponentiation, a side-channel attack that can discover which multiplier is used in the multiplication operations can recover the digits b_i and from them obtain the secret exponent b .

2.2 The Intel Cache Hierarchy

We now turn our attention to the cache hierarchy in modern Intel processors. The cache is a small, fast memory that exploits the temporal and spatial locality of memory accesses to bridge the speed gap between the faster CPU and slower memory. In the processors we are interested in, the cache hierarchy consists of three levels of caching. The top level, known as the *L1 cache*, is the closest to the execution core and is the smallest and the fastest cache. Each successive cache level is larger and slower than the preceding one, with the *last-level cache* (LLC) being the largest and slowest.

Cache Structure. The cache stores fixed-sized chunks of memory called *cache lines*. Each cache line holds 64 bytes of data that come from a 64-byte aligned *block* in memory. The cache is organized as multiple *cache sets*, each consisting of a fixed number of *ways*. A block of memory can be stored in any of the ways of a single cache set. For the higher cache levels, the mapping of memory blocks to cache sets is done by selecting a range of address bits. For the LLC, Intel uses an undisclosed hash function to map memory blocks to cache sets [34, 41, 55]. The L1 cache is divided into two sub caches: the *L1 data cache* (L1-D) which caches the data the program accesses, and the *L1 instruction cache* (L1-I) which caches the code the program executes. In multi-core processors, each of the cores has a dedicated L1 cache. However, multithreaded cores share the L1 cache between the two threads.

Cache Sizes. In the Intel Sandy Bridge microarchitecture, each of the L1-D and L1-I caches has 64 sets and 8 ways to a total capacity of $64 \cdot 8 \cdot 64 = 32,768$ bytes. The L2 cache has 512 sets and 8 ways, with a size of 256 KiB. The L2 cache is *unified*, storing both data and instructions. Like the L1 cache, each core has a dedicated L2 cache. The L3 cache, or the LLC, is shared by all of the cores of the processor. It has 2,048 sets *per core*, i.e. the LLC of a four core processor has 8,192 cache sets. The number of ways varies between processor models and ranges between 12 and 20. Hence the size of the LLC of a small dual core processor is 3 MiB, whereas the LLC of an 8-core processor can be in the order of 20 MiB. The Intel Xeon E5-2430 processor we used for many of our experiments is a 6-core processor with a 20-way LLC of size 15 MiB. More recent microarchitectures support more cores and more ways, yielding significantly larger LLCs.

Cache Lookup Policy. When the processor attempts to access data in memory, it first looks for the data in the L1 cache. In a *cache hit*, the data is found in the cache. Otherwise, in a *cache miss*, the processor searches for the data in the next level of the cache hierarchy. By measuring the time to access data, a process can distinguish cache hits from misses and identify whether the data was cached prior to the access.

2.3 Microarchitectural Side-Channel Attacks

In this section we review related works on microarchitectural side-channel timing attacks. These attacks exploit timing variations that are caused by contention on microarchitectural hardware resources in order to leak information on the usage of these resources, and indirectly on the internal operation of the victim. For a more complete coverage of microarchitectural timing attacks, see Ge et al. [24] and references therein.

Acıçmez and Seifert [6] distinguish between two types of channels: those that rely on a persistent state and those that exploit a transient state. Persistent-state channels exploit the limited storage space within the targeted microarchitectural resource. Transient-state channels, in contrast, exploit the limited bandwidth of the targeted element.

Persistent-State Attacks. The PRIME+PROBE attack [45, 46] is an example of a persistent-state attack. The attack exploits the limited storage space in cache sets to identify the sets used for the victim’s data. The attacker preloads data to the cache and allows the victim to execute before measuring the time to access the preloaded data. When the victim accesses its data, it is loaded into the cache, replacing some of the attacker’s preloaded data. Accessing data that has been replaced will take longer than accessing data still in the cache. Thus the attacker can identify the cache sets that the victim has accessed. Persistent-state channels have targeted the L1 data cache [9, 17, 45, 46], the L1 instruction cache [1, 5, 56], the branch prediction buffer [3, 4], the last-level cache [31, 33, 40, 51, 54], and DRAM open rows [47]. The PRIME+PROBE attack was used to recover the accessed multipliers in the sliding-window exponentiation of OpenSSL 0.9.7c [46] and of GnuPG 1.4.18 [31, 40].

Transient-state Attacks. Transient-state channels have been investigated mostly within the context of covert channels, where a *Trojan* process tries to covertly exfiltrate information. The idea dates back to Lampson [38] who suggests that processes can leak information by modifying their CPU usage. Covert channels were also observed with shared bus contention [30, 53]. Acıçmez and Seifert [6] are the first to publish a side-channel attack based on a transient state. The attack monitors the usage of the multiplication functional unit in a hyperthreaded processor. Monitoring the unit allows an attacker to distinguish between the square and the multiply phases of modular exponentiation. The attack was tested on a victim running fixed-window exponentiation, so no secret information was obtained.

Another transient-state channel uses bus contention to leak side-channel information [52]. By monitoring the capacity of the memory bus allocated to the attacker, the attacker is able to distinguish the square and the multiply steps. Because the attack of [52] was only demonstrated in a simulator, the question of whether actual hardware leaks such high-resolution information is still open.

Attack Classification. It is common to also classify cryptographic side-channel attacks as *time-driven*, *trace-driven*, or *access-driven* based on the type of information that the attacker learns about a victim cipher [2, 43]. Time-driven attacks measure the overall execution time of the victim. In trace-driven attacks, the attacker learns the outcome of each of the victims memory accesses. Finally, in access-driven at-




Fig. 1 Conventional (left) vs. scatter-gather (right) memory layout.

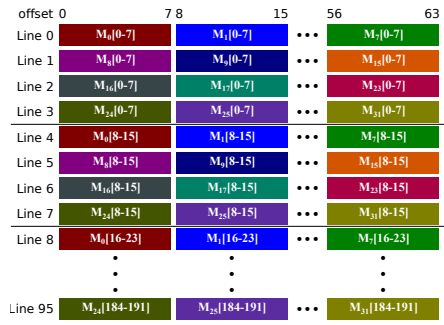


Fig. 2 The memory layout of the multipliers table in OpenSSL

tacks, the attacker learns partial information about the memory addresses accessed by victim.

2.4 Scatter-Gather Implementation

One of the countermeasures Intel recommends against side-channel attacks is to avoid secret-dependent memory access *at coarser than cache line granularity* [14, 15]. This approach is manifested in the patch Intel contributed to the OpenSSL project to mitigate the Percival [46] attack. The patch¹ changes the layout of the multipliers in memory. Instead of storing the data of each of the multipliers in consecutive bytes in memory, the new layout scatters each multiplier across multiple cache lines [16]. Before use, the fragments of the required multiplier are gathered to a single buffer which is used for the multiplication. Figure 1 contrasts the conventional memory layout of the multipliers with the layout used in the scatter-gather approach. This scatter-gather design ensures that the order of accessing cache lines when performing a multiplication is independent of the multiplier used.

Because Intel cache lines are 64 bytes long, the maximum number of multipliers that can be used with scatter-gather is 64. For large exponents, increasing the number of multipliers reduces the number of multiply operations performed during the exponentiations. Gopal et al. [26] suggest dividing the multipliers into 16-bit fragments rather than into bytes. This improves performance by allowing loads of two bytes in a single memory access, at the cost of reducing the maximum number of multipliers to 32. Gueron [27] recommends 32-bit fragments, thus reducing the number of multipliers to 16. He shows that the combined savings from the reduced number of memory accesses and the smaller cache footprint of the multipliers outweighs the performance loss due to the added multiplications required with less multipliers.

The OpenSSL Scatter-Gather Implementation. The implementation of exponentiation in the current version of

OpenSSL (1.0.2f) deviates slightly from the layout described above. For 2048-bit and 4096-bit key sizes the implementation uses a fixed-window algorithm with a window size of 5, requiring 32 multipliers. Instead of scattering the multipliers in each cache line, the multipliers are divided into 64-bit fragments, scattered across groups of four consecutive cache lines. (See Figure 2). That is, the table that stores the multipliers is divided into groups of four consecutive cache lines. Each group of four consecutive cache lines stores one 64-bit fragment of each multiplier. To avoid leaking information on the particular multiplier used in each multiplication, the gather process accesses all of the cache lines and uses a bit mask pattern to select the ones that contain fragments of the required multiplier. Furthermore, to avoid copying the multiplier data, the implementation combines the gather operation with the multiplication. This spreads the access to the scattered multiplier across the multiplication.

Key-Dependent Memory Accesses. Because the fragments of each multiplier are stored in a fixed offset within the cache lines, all of the scatter-gather implementations described above have memory accesses that depend on the multiplier used and thus on the secret key. For a pure scatter-gather approach, the multiplier is encoded in the low bits of the addresses accessed during the gather operation. For the case of OpenSSL’s implementation, only the three least significant bits of the multiplier number are encoded in the address while the other two bits are used as the index of the cache line within the group of four cache lines that contains the fragment.

We note that because these secret-dependent accesses are at a finer than cache line granularity, the scatter-gather approach has been considered secure against side-channel attacks [27].

2.5 Intel L1 Cache Banks

With the introduction of superscalar computing in Intel processors, cache bandwidth became a bottleneck for processor performance. To alleviate the issue, Intel introduced a cache

¹ <https://github.com/openssl/openssl/commit/46a643763de6d8e39ecf6f76fa79b4d04885aa59>

design consisting of multiple *banks* [7]. Each of the banks serves part of the cache line specified by the offset in the cache line. The banks can operate independently and serve requests concurrently. However, each bank can only serve one request at a time. When multiple accesses to the same bank are made concurrently, only one access is served, while the rest are delayed until the bank can handle them.

Fog [20] notes that cache-bank conflicts prevent instructions from executing simultaneously on Pentium processors. Delays due to cache-bank conflicts are also documented for other processor versions [21, 22, 32].

Both Bernstein [9] and Osvik et al. [45] mention that cache-bank conflicts cause timing variations and warn that these may result in a timing channel which may leak information about low address bits. Tromer et al. [50] note that while scatter-gather has no secret-dependent accesses to cache lines, it does have secret-dependent access to cache banks. However, although the risk of side-channel attacks based on cache-bank conflicts has been identified long ago, the effect was previously demonstrated using self-conflicts (i.e., between multiple concurrent memory accesses in the same process). Consequently, such information could only be used for a time-driven attack (which have never been demonstrated). This paper is the first empirical demonstration of a full fledged key extraction attack exploiting cache-bank conflicts between different processes.

2.6 False Dependency of Write after Read

Even the best cache design cannot guarantee that data is always cached when the program needs it. Consequently, some memory access operations must stall until the data is available. To avoid stalling the whole program when data is not immediately available, processors support *out-of-order* execution. When data for one instruction is not available, the processor can execute instructions following the stalled instruction with the stalled instruction being executed only when the data becomes available.

Processors that support out-of-order execution must ensure that changing the order of the instructions does not change the semantics of the program. For example, if one instruction loads data from a memory location and a later instruction stores data to the same location, the processor must ensure that the load instruction does not get the value stored by the subsequent store instruction. In such a case we say that there is a *dependency* between the load and the store instructions.

On first observation, the solution for dependencies like this may seem trivial: we only need to keep track of the addresses that currently-stalled load instructions try to read, and delay stores to these addresses until after the conflicting loads execute. However, in practice, we have a problem that

we do not always know that such conflicts exist. The main reason is that programs execute within a virtual address space, which maps (virtual) program addresses to (physical) memory addresses. Multiple virtual addresses may map to the same physical address. Hence, knowing only the virtual addresses is not enough to identify dependencies.

Virtual-to-physical memory mapping operates at a resolution of *pages*, which on Intel processors are 4,096 (2^{12}) bytes long. Consequently, the 12 least significant bits of a virtual address and of the matching physical address are the same. Hence, if the 12 least significant bits of two virtual addresses are different, we know that there is no dependency between these addresses and a store to one address does not have to stall until the load from the other address completes [29]. However, if the 12 least significant bits of two addresses are the same, a dependency *may* exist, and the store must wait until the processor resolves the physical addresses before we can determine if the dependency really exists. Thus, in the case of a *false dependency*, i.e. when a dependency does not exist, but the processor cannot determine that until after address resolution, store instructions take slightly longer than without the false dependency.

Bernstein and Schwabe [10] demonstrate timing variations due to false dependencies and suggest these can be used for cryptanalysis. However, as in the case of cache-bank conflicts, only self-conflicts have been demonstrated and these could only be used for a time-driven attack, had one been demonstrated.

3 The CacheBleed Attack

We now proceed to describe CacheBleed, the first side-channel attack to systematically recover access information at a granularity higher than a cache line. We present two variants of the attack. The first identifies the times at which a victim accesses data in a monitored cache bank by measuring the delays caused by contention on the cache bank. The second variant recovers similar information by measuring timing variations due to false dependencies. In our attack scenario, we assume that the victim and the attacker run concurrently on two hyperthreads of the same processor core. Thus, the victim and the attacker share the L1 data cache and the path to it. For the Sandy Bridge processor we use the attack variant that exploits cache-bank conflicts. On the Haswell and Skylake processors, which do not experience cache-bank conflicts, we use the variant exploiting false dependencies.

Exploiting cache-bank conflicts. Recall that the Sandy Bridge L1 data cache is divided into multiple banks and that the banks cannot handle concurrent load accesses. The attacker issues a large number of load accesses to a cache bank and measures the time to fulfill these accesses. If during the

attack the victim also accesses the same cache bank, the victim accesses will contend with the attacker for cache bank access, causing delays in the attack. Hence, when the victim accesses the monitored cache bank the attack will take longer than when the victim accesses other cache banks.

```

1  rdtscp
2  movq   %rax, %r10
3
4  addl   0x000(%r9), %eax
5  addl   0x040(%r9), %ecx
6  addl   0x080(%r9), %edx
7  addl   0x0c0(%r9), %edi
8  addl   0x100(%r9), %eax
9  addl   0x140(%r9), %ecx
10 addl   0x180(%r9), %edx
11 addl   0x1c0(%r9), %edi
.
.
.
256 addl   0xf00(%r9), %eax
257 addl   0xf40(%r9), %ecx
258 addl   0xf80(%r9), %edx
259 addl   0xfc0(%r9), %edi
260
261 rdtscp
262 subq   %r10, %rax

```

Listing 1 Cache-Bank Collision Attack Code

To implement CacheBleed we use the code in Listing 1. The bulk of the code (Lines 4–259) consists of 256 addl instructions that read data from addresses that are all in the same cache bank. (The cache bank is selected by the low bits of the memory address in register r9). We use four different destination registers to avoid contention on the registers themselves. Before starting the accesses, the code takes the value of the current cycle counter (Line 1) and stores it in register r10 (Line 2). After performing 256 accesses, the previously stored value of the cycle counter is subtracted from the current value, resulting in the number of cycles that passed during the attack.

Exploiting False Dependencies. Haswell processors and newer microarchitectures do not experience cache-bank conflicts [22]. For these architectures, we design an attack based on false dependencies. Recall that a false dependency occurs when the processor needs to handle concurrent load and store to addresses that have the same 12 least significant bits. In the case of a false dependency, the store operation is delayed until the physical addresses of both instructions are known.

To exploit false dependencies, the attacker issues a large number of store instructions, all with the same 12 least significant bits. If during the attack the victim loads data from an address that has the same 12 least significant bits, a false dependency will occur, delaying the store operations. The

attacker can identify the false dependency event through measuring the time it takes to perform the store operations.

```

1  rdtscp
2  movq   %rax, %r10
3
4  addl   %eax, 0x0000(%r9)
5  addl   %eax, 0x1000(%r9)
6  addl   %eax, 0x2000(%r9)
7  addl   %eax, 0x3000(%r9)
8  addl   %eax, 0x0000(%r9)
.
.
.
32  addl   %eax, 0x0000(%r9)
33  addl   %eax, 0x1000(%r9)
34  addl   %eax, 0x2000(%r9)
35  addl   %eax, 0x3000(%r9)
36
37  rdtscp
38  subq   %r10, %rax

```

Listing 2 Cache-Bank Collision Attack Code

The implementation (Listing 2) is very similar to the implementation of the cache-bank-conflicts attack. The only differences are that our code modifies memory, i.e. stores new data to it, that the affected addresses are spaced by 0x1000 (4,096) bytes from each other and that we use only 32 operations, instead of using 256.




Fig. 3 Distribution of time to read 256 entries from a cache bank on a Sandy Bridge Processor.




Fig. 4 Distribution of time to write 32 entries at an offset on a Skylake Processor.




Fig. 5 Distribution of time to write 32 entries at an offset on a Haswell Processor.

Measurements. We ran the attack code on the three platforms described in Table 1. We used the cache-bank conflicts attack on the Sandy Bridge platform and the false dependency attack on the other two platform. Figure 3, Figure 4 and Figure 5 show the histogram of the running times of the attack code under several scenarios.²

Scenario 1: Idle. In the first scenario, *idle hyperthread*, the attacker is the only program executing on the core. That is, one of the two hyperthreads executes the attack code while the other hyperthread is idle.

On the Sandy Bridge, the attack takes about 230 cycles. Completing 256 add operations in this time clearly shows that the Intel processor is superscalar and that the cache can handle more than one access in a CPU cycle.

On the Skylake and Haswell architectures we only perform 32 add-to-memory operations. We see that these operations are significantly slower than the add operations used in the Sandy Bridge attack. The main cause of the difference is that add-to-memory operations need to first load the data from the cache before they can compute the result and write it back.

Scenario 2: Pure Compute. The second scenario has a victim running a computation on the registers, without any memory access. As we can see, access in this scenario is slower than when there is no victim. Because the victim does not perform memory accesses, cache-bank conflicts and false dependencies cannot explain this slowdown. Hyperthreads, however, share most of the resources of the core, including the execution units, read and write buffers and the register allocation and renaming resources [22]. Contention on any of these resources can explain the slowdown we see when running a pure-compute victim.

Scenario 3: Pure Memory. At the other extreme is the *pure memory* victim, which continuously accesses the cache bank that the attacker monitors. As we can see, the attack code takes almost twice as long to run in this scenario. On the Sandy Bridge and the Skylake platforms, the distribution of attack times is completely distinct from any of the other scenarios. Hence identifying the victim in this scenario

² For clarity, the presented histograms show the envelope of the measured data.

is trivial. This scenario is, however, not realistic—programs usually perform some calculation.

Scenarios 4 and 5: Mixed Load. The last two scenarios aim to measure a slightly more realistic case. In this case, one in four victim operations is a memory access, where all of these memory accesses are to the same cache bank or page offset. In this scenario we measure both the case that the victim accesses the monitored offset (*mixed-load*) and when there is no contention between the victim and the attacker (*mixed-load-NC*). We see that the two scenarios are distinguishable, but there is some overlap between the two distributions. Consequently, a single measurement may be insufficient to distinguish between the two scenarios.

Observing the graphs, we notice that in the cache-bank conflicts attack, the difference between the distributions of the mixed-case scenario is quite large, with only a small overlap. In the case of the false dependencies attack, the overlap is quite small, particularly on the Haswell architecture.

In practice, even this mixed-load scenario is not particularly realistic. Typical programs will access memory in multiple cache banks. Hence the differences between measurement distributions may be much smaller than those presented in Figure 3. In the next section we show how we overcome this limitation and correctly identify a small bias in the cache-bank access patterns of the victim.

4 Attacking the OpenSSL Modular Exponentiation Implementation

To demonstrate the technique in a real scenario, we use CacheBleed to attack the implementation of the RSA decryption in version 1.0.2f of OpenSSL.³ This section describes how we exploit cache-bank conflicts on the Sandy Bridge processor to attack OpenSSL. The next section discusses the application of the false-dependencies attack.

The implementation in OpenSSL uses a fixed-window exponentiation with $w = 5$. As discussed in Section 2.4 OpenSSL uses a combination of the scatter-gather technique with masking for side-channel attack protection. Recall that the multipliers are divided into 64-bit fragments. These fragments are scattered into 8 *bins* along the cache lines such that the three least significant bits of the multiplier select the bin. The fragments of a multiplier are stored in groups of four consecutive cache lines. The two most significant bits of the multiplier select the cache line out of the four in which the fragments of the multiplier are stored. See Figure 2. The multiplication code selects the bin to read using the least significant bits of the multiplier. It then reads

³ Following the disclosure of our attack, OpenSSL released a fix for the problem. See Section 7. Version 1.0.2f is, therefore, the latest version vulnerable to the attack described here.




Fig. 6 Measurement trace of OpenSSL RSA decryption




Fig. 7 Measurement trace of OpenSSL RSA decryption—detailed view

a fragment from the selected bin in each of the four cache lines and uses masking to select the fragment of the required multiplier. Because the multiplication code needs to access the multiplier throughout the multiplication, the cache banks of the bin containing the multiplier are accessed more often than other cache banks. We use CacheBleed to identify the bin and, consequently, to find the three least significant bits of the multiplier.

Identifying Exponentiations. We begin by demonstrating that it is possible to identify the exponentiation operations using cache-bank conflicts. Indeed, using the code in Listing 1, we create a sequence of measurements of cache-bank conflicts. As mentioned in Section 3, the difference between the distributions of measurements in similar scenarios may be very small. Consequently, a single measurement is unlikely to be sufficient for identifying the bin used in each multiplication. To distinguish the distributions, we create multiple sequences and average the measurements at each trace point to get a trace of the average measurement time. Figure 6 shows the traces of measurements of two bins, each averaged over 1,000 decryptions using a 4096-bit key.

The figure clearly shows the two exponentiations executed as part of the RSA-CRT calculation. Another interesting feature is that the measurements for the two bins differ by about 4 cycles. The difference is the result of the OpenSSL modular reduction algorithm, which accesses even bins more often than odd bins. Consequently, there is more contention on even bins, and measurements on even bins take slightly longer than those on odd bins.

Identifying Multiplication Operations. Next, we show that it is also possible to identify the individual multiplication operations performed during the modular exponentiation operation. Indeed, Figure 7 shows a small section of the traces of the odd bins. In these traces, we can clearly see the multiplication operations (marked with arrows) as well as the spikes for each of the squaring and modular reduction operations. Recall that the OpenSSL exponentiation repeatedly calculate sequences of five modular squaring and reduction operations followed by a modular multiplication.

Identifying Multiplier Values. Note that in the second and fourth multiplications, the measurements in the trace of bin 3 (yellow) take slightly longer than the measurements of the other bins. This indicates that the three least significant digits of the multiplier used in these multiplications are 011. Similarly, the spike in the green trace observed during the third multiplication indicates that the three least significant bits of the multiplier used are 001. This corresponds to the ground truth where the multipliers used in the traced section are 2, 11, 1, 11.

As we can see, we can extract the multipliers from the trace. However, there are some practical challenges that complicate both the generation of the traces and their analysis. We now discuss these issues.

Aligning CacheBleed Measurement Sequences for Averaging. Recall that the traces shown in Figure 7 are generated by averaging the sequences of CacheBleed measurements over 1,000 decryptions. When averaging, we need to ensure that the sequences align with each other. That is, we must ensure that each measurement is taken in the same relative time in each multiplication.

To ensure that the sequences are aligned, we use the FLUSH+RELOAD attack [54] to find the start of the exponentiation. Once found, we start collecting enough CacheBleed measurements to cover the whole exponentiation. FLUSH+RELOAD has a resolution of about 500 cycles, ensuring that the sequences start within 500 cycles, or up to two measurements, of each other.

Relative Clock Drift. Aligning the CacheBleed sequences at the start of the exponentiation does not result in a clean signal. This is because both the victim and the attacker are user processes, and they may be interrupted by the operating system. The most common interruption is due to timer interrupts, which on Linux-based operating systems happen every millisecond. Since each modular exponentiation in the calculation of a 4096-bit RSA-CRT decryption takes 5 ms, we experience 5 to 6 timer interrupts during the exponentiation. Timer interrupts can be easily identified because serving them takes over 5,000 cycles, whereas non-interrupted measurements take around 300 cycles. Consequently, if a measurement takes more than 1,000 cycles, we assume that it was interrupted and therefore discard it.




Fig. 8 CacheBleed average trace towards the end of the exponentiation




Fig. 9 Measurement trace after a lowpass filter

The attacker, however, does not have exact information on the interrupts that affect the victim, resulting in clock drift between the attacker and the victim. As we progress through the exponentiation, the signal we capture becomes more noisy. **Figure 8** shows the signal towards the end of the exponentiation. As we can see, the multiplications are barely visible.

To reduce noise, we pass the signal through a low-pass filter, which removes high frequencies from the signal and highlights the behavior at the resolution of one multiplication. **Figure 9** shows the result of passing the above trace through the filter. It is possible to clearly identify three multiplications, using bins 7, 5 and 1.

Aligning Traces of Multiple Bins. As discussed above, measurements in even bins are slower on average than measurements in odd bins. This creates two problems. The first is that we need to normalize the traces before comparing them to find the multiplier. The second problem is that we use the measurements as a virtual clock. Consequently, when we measure over a fixed period of time, traces of even bins will be shorter, i.e., have fewer measurements, than traces of odd bins. This creates a clock shift between traces belonging to even bins and traces belonging to odd bins, which increases as the exponentiation progresses. In order to normalize the trace length, we remove element 0 of the frequency domain. This effectively subtracts the trace’s average from each trace measurement, thereby making all the traces be at the same length.




Fig. 10 The frequency spectrum of a trace




Fig. 11 Normalized resampled traces

We then find the frequency of multiplications in the trace by looking at the frequency domain of the trace. **Figure 10** shows the frequency spectrum of two of the traces. For a 4096-bit key, OpenSSL performs two exponentiations with 2048-bit exponents. With a window size of 5, there are $2048/5 \approx 410$ multiplications. As we can see, there is a spike around the frequency 410 matching the number of multiplications. Using the frequency extracted from the trace, rather than the expected number of multiplications, allows us to better adjust to the effects of noise at the start and end of the exponentiation which might otherwise result in a loss of some multiplications.

Partial Key Extraction. We used CacheBleed to collect 16 traces, one for each of the 8 bins in each of the two exponentiations. Each trace is the average of 1,000 sequences of measurements, totalling 16,000 decryption operations. **Figure 11** shows a sample of the analyzed traces, i.e. after averaging, passing through a low-pass filter, normalizing the signal and resampling. As we can see, the used bins are clearly visible in the figure.

We managed to recover the three least significant bits of almost all of the multipliers. Due to noise at the start and the end of the exponentiations, we are missing one or two of the leading and trailing multiplications of each exponentiation. In **Section 6**, we show that the information we obtain about the three least significant bits of almost all of the multipliers is enough for key extraction.

5 OpenSSL and the False-Dependencies attack

As discussed earlier, Haswell and newer microarchitectures do not experience cache-bank conflicts. They do, however, experience false dependencies, which are affected by low-

address-bits. In this section we discuss applying the false-dependencies variant of the CacheBleed attack to OpenSSL 1.0.2f. We implemented this attack on the Skylake machine because, as we have seen in Section 3, it shows a slightly larger difference between measurement distributions for similar test scenarios.

The main challenge with deploying the false-dependencies attack is that our “signal” is much weaker. The difference between the measurement distributions is smaller than in the case of cache-bank conflicts, requiring more samples to identify the correct scenario. Moreover, for a false dependency, the addresses the attacker uses must match the 12 least significant bits of the addresses used by the victim, whereas for a cache-bank conflict only 6 bits should match. Consequently we can expect the number of false dependencies experienced to be 64 times lower than the number of cache-bank conflicts. Thus, to be able to observe if a (false) dependency induced stall has occurred we need many more samples than are required for the cache-bank conflicts attack.




Fig. 12 False-Dependencies attack Measurement trace of OpenSSL RSA decryption

We use the code in Listing 2 to collect sequences of measurements of false dependencies during the computation of OpenSSL exponentiation. Averaging over 100,000 decryptions we get the trace in Figure 12. As we can see, the general structure of the OpenSSL decryption is clearly visible. We note that at this scale, traces for all of the bins completely overlap. (Compare with Figure 6 where even and odd bins show different measurements).

If we increase the number of decryptions to 500,000 and zoom in, we can see the individual squaring and multiplication operations. (Figure 13). Observing the individual multiplication operations, we can see that the traces of the even bins are slightly (0.1 cycle) higher than those of the odd bins. This demonstrates that the attack can distinguish between different activity levels at the different bins.

Unfortunately, our attack is unable to identify the bins that contain the individual multipliers. We speculate that a significant increase in the number of samples we average over might allow us to recover the multiplier values as well. However, such an attack would require several months of constant decryptions, and is, therefore less practical. We

note, however, that our attack demonstrates that the Skylake microarchitecture experiences timing variations that depend on low address bits. Further advances in attack techniques may allow better recovery of such variations and pose a significant security threat.

6 Recovering the RSA Private Key

Successfully carrying out the attack in the Section 4 for a 4096-bit modulus allowed us to learn the three least significant bits of every window of five bits for the Chinese remainder theorem coefficients $d_p = d \bmod (p - 1)$ and $d_q = d \bmod (q - 1)$. In this section, we describe how to use this knowledge to recover the full private RSA key. We use the techniques of Heninger and Shacham [28] and İnci et al. [31].

Solving For The Modular Multipliers. We have partial knowledge of the bits of d_p and d_q , where each satisfies the relation $ed_p = 1 + k_p(p - 1)$ and $ed_q = 1 + k_q(q - 1)$ for positive integers $k_p, k_q < e$. In the common case of $e = 65537$, this leaves us with at most 2^{32} possibilities for pairs of k_p, k_q to test. Following [31], the k_p and k_q are related, so we only need to search 65,537 possible values of k_p .

We start by rearranging the relations on d_p and d_q to obtain $ed_p - 1 - k_p = k_p p$ and $ed_q - 1 - k_q = k_q q$. Multiplying these together, we obtain the relation

$$(ed_p - 1 + k_p)(ed_q - 1 + k_q) = k_p k_q N. \quad (1)$$

Reducing modulo e yields $(k_p - 1)(k_q - 1) \equiv k_p k_q N \bmod e$.

Thus, given a value for k_p we can solve for the unique value of $k_q \bmod e$. We do not have enough information about d_p and d_q to deduce further information, so we must test all e values of k_p .

Branch and Prune Algorithm. For each candidate k_p and k_q , we will use Equation 1 to iteratively solve for d_p and d_q starting from the least or most significant bits, branching to generate multiple potential solutions when bits are unknown and pruning potential solutions when known bits contradict a given solution. In contrast to [28], the bits we know are not randomly distributed. Instead, they are synchronized to the three least significant bits of every five, with one or two full windows of five missing at the least and most significant positions of each exponent. This makes our analysis much simpler: when a bit of d_p and d_q is unknown at a location i , we branch to generate two new solutions. When a bit of d_p and d_q is known at a particular location, using the same heuristic assumption as in [28], an incorrect solution will fail to match the known bit of d_p and d_q with probability 0.5. When k_p and k_q are correct, we expect our algorithm to generate four new solutions for every pair of unknown bits, and prune these to a single correct solution at every string of

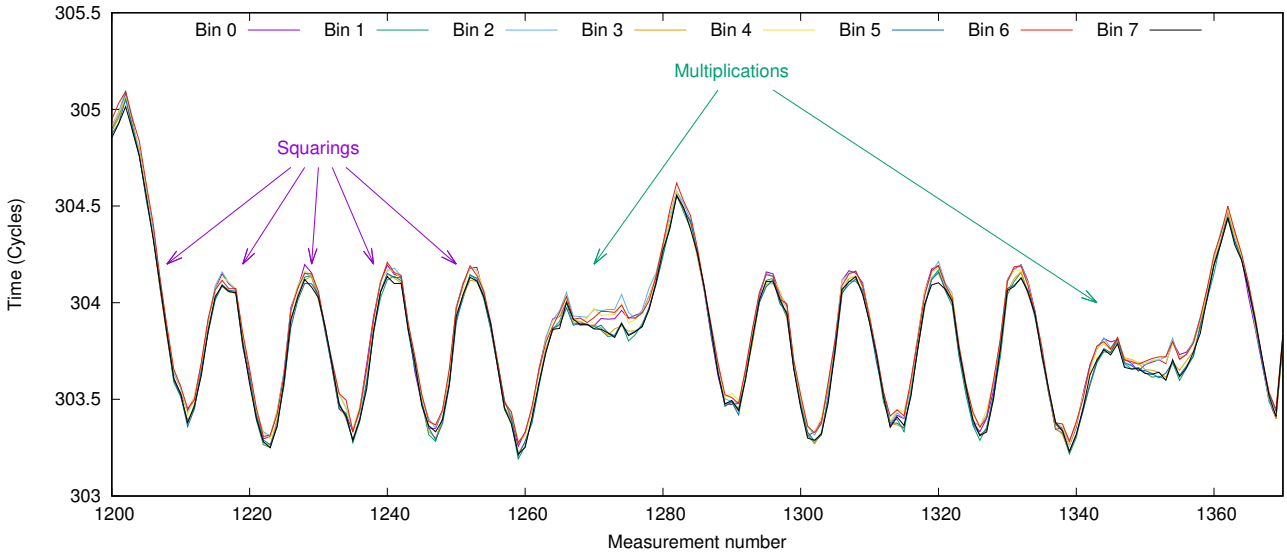


Fig. 13 False-Dependencies attack Measurement trace of OpenSSL RSA decryption — Detailed View

three known bits. When k_p and k_q are incorrect, we expect no solutions to remain after a few steps.

Empirical Results. We tested key recovery on the output of our attack run on a 4096-bit key, which correctly recovered the three least significant bits of every window of five, but missed the two least significant windows and one most significant window for both d_p and d_q . We implemented this algorithm in Sage and ran it on a Cisco UCS Server with two 2.30GHz Intel E5-2699 processors and 128 GiB of RAM. For the correct values of k_p and k_q , our branch-and-prune implementation recovered the full key in 1 second on a single core after examining 6,093 candidate partial solutions, and took about 160 ms to eliminate an incorrect candidate pair of k_p and k_q after examining an average of 1,500 candidate partial solutions. A full search of all 65,537 candidate pairs of k_p and k_q parallelized across 36 hyperthreaded cores took 3.5 minutes. We assumed the positions of the missing windows at the most and least significant bits were known. If the relative positions are unknown, searching over more possible offsets would increase the total search time by a factor of 9.

7 Mitigation

Countermeasures for the CacheBleed attack can operate at the hardware, the system or the software level. Increasing the bandwidth of the cache banks, so that conflicts do not occur, offers some protection. Next, since our false-dependencies attack cannot recover keys, no effective attack is currently known for newer microarchitectures which do not suffer from cache-bank conflicts [22, 32] (such as Nehalem and Haswell). However, this does not rule out future

key extraction attacks based on false dependencies. Eliminating timing variations due to false dependencies can solve the issue, however this solution will only apply to future processors, with the vulnerabilities remaining in processors currently in the market.

Disabling Hyperthreading. The simplest countermeasure at the system level is to disable hyperthreading. Disabling hyperthreading, or only allowing hyperthreading between processes within the same protection domain, prevents any concurrent access to the cache banks and eliminates any conflicts. Unlike attacks on persistent state, which may be applicable when a core is time-shared, the transient state that CacheBleed exploits is not preserved during a context switch. Hence the core can be time-shared between non-trusting processes. The limited security of hyperthreading has already been identified [6]. We recommend that hyperthreading be disabled even on processors that are not vulnerable to CacheBleed for security-critical scenarios where untrusted users share processors.

We note that AMD processors before the Zen microarchitecture [8] do not support hyperthreading and, as such, are not vulnerable to our attack techniques. Since AMD processors using Zen microarchitecture are not commercially available at the time of writing this paper, we leave the task of attacking them for future work.

Constant-Time Implementations. At the software level, the best countermeasure is to use a *constant-time* implementation, i.e. one that does not have secret-dependent branches or memory accesses. A common technique for implementing constant-time table lookup is to use a combination of arithmetic and bitwise operations to generate a mask that depends on the secret value. The whole table is then accessed and the mask is used to select the required table en-

try. Mozilla’s countermeasure for CacheBleed uses this approach.

Modifying Memory Accesses. Rather than using to a constant-time implementation, OpenSSL mitigates CacheBleed through a combination of two changes. The first change is to use 128-bit memory accesses, effectively halving the number of bins used. The second change is to modify the memory access pattern during the gathering process so that the software accesses a different offset in each of the four cache lines.

Combining the four different offsets with the 128-bit accesses means that when gathering a multiplier fragment, OpenSSL accesses all 16 of the cache banks. The order of accessing the cache banks depends on the value of the multiplier, so the design leaks secret key information to adversaries that can recover the order of the accesses. We note, however, that our attack does not have the resolution required to determine the order of successive memory accesses and that we are not currently aware of any technique for exploiting this leak.

Furthermore, using 128-bit memory accesses means that the potential leakage created by the order of accessing the cache banks is only two bits for each multiplier, or 40% of the bits of the exponents for the 5-bit windows used by OpenSSL for both 2048 and 4096-bit exponents. Our key recovery technique will produce exponentially many solutions in this case: heuristically, we expect it to branch to produce two solutions for each multiplier. In this case, the attacker could use the branch-and-prune method to produce exponentially many candidates up to half the length of each Chinese remainder theorem exponent d_p or d_q , and then use the method of Blömer and May [11] to recover the remaining half in polynomial time. Thus even if an adversary is able to exploit the leak, full key recovery may only be feasible for very small keys without further algorithmic improvements.

While we are not aware of a practical exploit of the leak in the OpenSSL code, we believe that leaving a known timing channel is an undue risk. We have conveyed information about the leak and our concerns to the OpenSSL development team.

8 Conclusions

In this work, we presented CacheBleed, the first timing attack to recover low address bits from secret-dependent memory accesses. We have demonstrated that the attack is effective against state-of-the-art cryptographic software, widely thought to be immune to timing attacks. Our attack requires a strong adversarial model, including a requirement for co-residency within the same execution core, a large number of trace acquisitions and a limited number of vulnerable microarchitecture. It, nevertheless, demonstrates that

previously speculated vulnerabilities are a real threat which can be exploited for key extraction.

The timing variations that underlie this attack and the risk associated with them have been known for over a decade. Osvik et al. [45] warn that “*Cache bank collisions (e.g., in Athlon 64 processors) likewise cause timing to be affected by low address bits.*” Bernstein [9] mentions that “*For example, the Pentium 1 has similar cache-bank conflicts.*” A specific warning about the cache-bank conflicts and the scatter-gather technique appears in Footnote 38 of Tromer et al. [50].

Our research illustrates the risk to users when cryptographic software developers dismiss a widely hypothesized potential attack merely because no proof-of-concept has yet been demonstrated. This is the prevailing approach for security vulnerabilities, but we believe that for cryptographic vulnerabilities, this approach is risky, and developers should be proactive in closing potential vulnerabilities even in the absence of a fully practical attack. To that end we observe that OpenSSL’s decision to use an ad-hoc mitigation techniques, instead of deploying a constant-time implementation, continues to follow such a risky approach.

Acknowledgements

We would like to thank Daniel J. Bernstein for suggesting the name CacheBleed and for helpful comments.

This material is based upon work supported by the U.S. National Science Foundation under Grants No. CNS-1408734, CNS-1505799, and CNS-1513671, a gift from Cisco, the Blavatnik Interdisciplinary Cyber Research Center, the Check Point Institute for Information Security, a Google Faculty Research Award, the Israeli Centers of Research Excellence I-CORE program (center 4/11), the Leona M. & Harry B. Helmsley Charitable Trust, and by NATO’s Public Diplomacy Division in the Framework of “Science for Peace”.

References

1. Onur Aciğmez. Yet another microarchitectural attack: exploiting I-cache. In *CSAW*, Fairfax, VA, US, 2007.
2. Onur Aciğmez and Çetin Kaya Koç. Microarchitectural attacks and countermeasures. In *Cryptographic Engineering*, pages 475–504. 2009.
3. Onur Aciğmez, Shay Gueron, and Jean-Pierre Seifert. New branch prediction vulnerabilities in OpenSSL and necessary software countermeasures. In *11th IMA International Conference on Cryptography and Coding*, pages 185–203, Cirencester, UK, 2007.

4. Onur Aciçmez, Çetin Kaya Koç, and Jean-Pierre Seifert. Predicting secret keys via branch prediction. In *2007 CT-RSA*, pages 225–242, 2007.
5. Onur Aciçmez, Billy Bob Brumley, and Philipp Grabher. New results on instruction cache attacks. In *CHES*, Santa Barbara, CA, US, 2010.
6. Onur Aciçmez and Jean-Pierre Seifert. Cheap hardware parallelism implies cheap security. In *Fourth International Workshop on Fault Diagnosis and Tolerance in Cryptography*, pages 80–91, Vienna, AT, 2007.
7. Donald B. Alpert, Mustafiz R. Choudhury, and Jack D. Mills. Interleaved cache for multiple accesses per clock cycle in a microprocessor. US Patent 5559986, Sep 1996.
8. AMD. <http://www.amd.com/en-gb/innovations/software-technologies/zen-cpu>.
9. Daniel J. Bernstein. Cache-timing attacks on AES, 2005. Preprint available at <http://cr.yp.to/papers.html#cachetiming>.
10. Daniel J. Bernstein and Peter Schwabe. A word of warning. CHES’13 Rump Session, Aug 2013.
11. Johannes Blömer and Alexander May. New partial key exposure attacks on RSA. In *Advances in Cryptology—CRYPTO 2003: 23rd Annual International Cryptology Conference*, pages 27–43. Berlin, Heidelberg, 2003.
12. BoringSSL. <https://boringssl.googlesource.com/boringssl/>.
13. Jurjen Bos and Matthijs Coster. Addition chain heuristics. In *CRYPTO’89*, pages 400–407, Santa Barbara, CA, US, Aug 1989.
14. Ernie Brickell. Technologies to improve platform security. CHES’11 Invited Talk, Sep 2011. URL http://www.iacr.org/workshops/ches/ches2011/presentations/Invited%201/CHES2011_Invited_1.pdf.
15. Ernie Brickell. The impact of cryptography on platform security. CT-RSA’12 Invited Talk, Feb 2012. URL http://www.rsaconference.com/writable/presentations/file_upload/cryp-106.pdf.
16. Ernie Brickell, Gary Graunke, and Jean-Pierre Seifert. Mitigating cache/timing based side-channels in AES and RSA software implementations. RSA Conference 2006 session DEV-203, Feb 2006.
17. Billy Bob Brumley and Risto M. Hakala. Cache-timing template attacks. In *15th ASIACRYPT*, pages 667–684, Tokyo, JP, Dec 2009.
18. Billy Bob Brumley and Nicola Tuveri. Remote timing attacks are still practical. In *16th ESORICS*, Leuven, BE, 2011.
19. David Brumley and Dan Boneh. Remote timing attacks are practical. In *12th USENIX Security*, pages 1–14, Washington, DC, US, 2003.
20. Agner Fog. How to optimize for the Pentium processor. <https://notendur.hi.is/hh/kennsla/sti/h96/pentopt.txt>, Aug 1996.
21. Agner Fog. How to optimize for the Pentium family of microprocessors. <https://cr.yp.to/2005-590/fog.pdf>, Apr 2004.
22. Agner Fog. The microarchitecture of Intel, AMD and VIA CPUs: An optimization guide for assembly programmers and compiler makers. <http://www.agner.org/optimize/microarchitecture.pdf>, Jan 2016.
23. Harvey L. Garner. The residue number system. *IRE Trans. Electron. Computers*, EC-8(2):140–147, Jun 1959.
24. Qian Ge, Yuval Yarom, David Cock, and Gernot Heiser. A survey of microarchitectural timing attacks and countermeasures on contemporary hardware. *J. Cryptographic Engin.* (to appear).
25. Daniel Genkin, Adi Shamir, and Eran Tromer. RSA key extraction via low-bandwidth acoustic cryptanalysis. In *CRYPTO 2014*, pages 444–461, Santa Barbara, CA, US, Aug 2014.
26. Vinodh Gopal, James Guilford, Erdinc Ozturk, Wajdi Feghali, Gil Wolrich, and Martin Dixon. Fast and constant-time implementation of modular exponentiation. In *Embedded Systems and Communications Security*, Niagara Falls, NY, US, Sep 2009.
27. Shay Gueron. Efficient software implementations of modular exponentiation. *J. Cryptographic Engin.*, 2(1): 31–43, May 2012.
28. Nadia Heninger and Hovav Shacham. Reconstructing RSA private keys from random key bits. In *CRYPTO 2009*, pages 1–17, Santa Barbara, CA, US, Aug 2009.
29. Sebastian Hily, Zhongying Zhang, and Per Hammarlund. Resolving false dependencies of speculative load instructions. U.S. Patent 7,603,527, Oct 2009.
30. Wei-Ming Hu. Reducing timing channels with fuzzy time. In *1991 Comp. Soc. Symp. Research Security & Privacy*, pages 8–20, Oakland, CA, US, 1991.
31. Mehmet Sinan İnci, Berk Gülmezoğlu, Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. Seriously, get off my cloud! Cross-VM RSA key recovery in a public cloud. IACR Cryptology ePrint Archive, Report 2015/898, Sep 2015.
32. Intel 64 & IA-32 AORM. *Intel 64 and IA-32 Architectures Optimization Reference Manual*. Intel Corporation, Apr 2012.
33. Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. S&A: A shared cache attack that works across cores and defies VM sandboxing – and its application to AES. In *S&P*, San Jose, CA, US, May 2015.
34. Gorka Irazoqui, Thomas Eisenbarth, and Berk Sunar. Systematic reverse engineering of cache slice selection in Intel processors. In *Euromicro Conference on Digital*

- System Design*, Funchal, Madeira, Portugal, Aug 2015.
35. Paul Kocher, Joshua Jaffe, and Benjamin Jun. Differential power analysis. In *CRYPTO*, volume 1666 of *LNCSS*, pages 388–397, 1999.
 36. Paul Kocher, Joshua Jaffe, Benjamin Jun, and Pankaj Rohatgi. Introduction to differential power analysis. *J. Cryptographic Engin.*, 1:5–27, Apr 2011.
 37. Paul C Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In *16th Annual International Cryptology Conference on Advances in Cryptology*, pages 104–113. Springer, 1996.
 38. Butler W. Lampson. A note on the confinement problem. *CACM*, 16:613–615, 1973.
 39. LibreSSL Project. <https://www.libressl.org>.
 40. Fangfei Liu, Yuval Yarom, Qian Ge, Gernot Heiser, and Ruby B Lee. Last-level cache side-channel attacks are practical. In *S&P*, pages 605–622, San Jose, CA, US, May 2015.
 41. Clémentine Maurice, Nicolas Le Scouarnec, Christoph Neumann, Olivier Heen, and Aurélien Francillon. Reverse engineering Intel last-level cache complex addressing using performance counters. In *RAID*, Kyoto, Japan, Nov 2015.
 42. Mozilla. Network security services. <https://developer.mozilla.org/en-US/docs/Mozilla/Projects/NSS>.
 43. Michael Neve and Jean-Pierre Seifert. Advances on access-driven cache attacks on AES. In *13th International Workshop on Selected Areas in Cryptography*, Montreal, CA, 2006.
 44. OpenSSL Project. <https://openssl.org>.
 45. Dag Arne Osvik, Adi Shamir, and Eran Tromer. Cache attacks and countermeasures: The case of AES. In *2006 CT-RSA*, 2006.
 46. Colin Percival. Cache missing for fun and profit. In *BSDCan 2005*, Ottawa, CA, 2005.
 47. Peter Pessl, Daniel Gruss, Clémentine Maurice, Michael Schwarz, and Stefan Mangard. Reverse engineering Intel DRAM addressing and exploitation. *arXiv preprint arXiv:1511.08756*, 2015.
 48. Jean-Jacques Quisquater and David Samyde. Electromagnetic analysis (EMA): Measures and countermeasures for smart cards. In *E-Smart'01*, pages 200–210, Cannes, FR, Sep 2001.
 49. Ronald L. Rivest, Adi Shamir, and Leonard Adleman. A method for obtaining digital signatures and public-key cryptosystems. *CACM*, 21:120–126, 1978.
 50. Eran Tromer, Dag Arne Osvik, and Adi Shamir. Efficient cache attacks on AES, and countermeasures. *Journal of Cryptology*, 23(1):37–71, 2010.
 51. Joop van de Pol, Nigel P. Smart, and Yuval Yarom. Just a little bit more. In *2015 CT-RSA*, pages 3–21, San Francisco, CA, USA, Apr 2015.
 52. Yao Wang and G Edward Suh. Efficient timing channel protection for on-chip networks. In *6th NoCS*, pages 142–151, Lyngby, Denmark, 2012.
 53. Zhenyu Wu, Zhang Xu, and Haining Wang. Whispers in the hyper-space: High-speed covert channel attacks in the cloud. In *21st USENIX Security*, Bellevue, WA, US, 2012.
 54. Yuval Yarom and Katrina Falkner. FLUSH+RELOAD: a high resolution, low noise, L3 cache side-channel attack. In *23rd USENIX Security*, pages 719–732, San Diego, CA, US, 2014.
 55. Yuval Yarom, Qian Ge, Fangfei Liu, Ruby B. Lee, and Gernot Heiser. Mapping the Intel last-level cache. <http://eprint.iacr.org/>, Sep 2015.
 56. Yinqian Zhang, Ari Juels, Michael K. Reiter, and Thomas Ristenpart. Cross-VM side channels and their use to extract private keys. In *19th CCS*, pages 305–316, Raleigh, NC, US, Oct 2012.