

# Attack Directories, Not Caches: Side-Channel Attacks in a Non-Inclusive World

Mengjia Yan, Read Sprabery, Bhargava Gopireddy, Christopher Fletcher, Roy Campbell, Josep Torrellas  
University of Illinois at Urbana Champaign  
{myan8, spraber2, gopired2, cwfletch, rhc, torrella}@illinois.edu

**Abstract**—Although clouds have strong virtual memory isolation guarantees, cache attacks stemming from shared caches have proved to be a large security problem. However, despite the past effectiveness of cache attacks, their viability has recently been called into question on modern systems, due to trends in cache hierarchy design moving away from inclusive cache hierarchies.

In this paper, we prove that cache attacks are still dangerous by designing the first cross-core Prime+Probe attack on a *sliced non-inclusive* last-level cache. This attack works with minimal assumptions: the adversary need not share any virtual memory with the victim nor be run on the same processor core. We also show the first high-bandwidth Evict+Reload attack on the same hardware. In both attacks, we leverage the microarchitecture of the coherence *directory*, and argue for targeting the directory to bootstrap conflict-based cache attacks. We demonstrate both attacks by extracting key bits during RSA operations in GnuPG on a state-of-the-art non-inclusive Intel Skylake-X server.

## I. INTRODUCTION

Cloud computing on shared machines is now ubiquitous. Cloud hypervisors share physical hardware resources between concurrent guest Virtual Machines (VMs), giving each VM the impression that it owns the entire cloud. On the business side, dynamically sharing hardware between tenants is essential to keep cloud computing economically viable. However, in this environment, an obvious challenge is security. Fortunately, researchers and industry have developed a suite of techniques—e.g., the hypervisor-OS privilege hierarchy and hardware security extensions such as Intel SGX [1]—to provide virtual memory isolation between VMs and between processes within a VM.

Unfortunately, virtual memory isolation is insufficient to maintain privacy in the cloud. The very fact that users share the same physical machine leads to *shared resource attacks*, whereby the adversary can infer sensitive information by monitoring how the victim uses available hardware [2]–[6]. Of these, cache attacks [2], [7], [8] are arguably the most popular and powerful, enabling an adversary to learn fine-grain information regarding a victim process’ memory access pattern—e.g., attacks can disclose encryption keys [8], user keystrokes [9], user web behavior [10], and more. Worse, these attacks can succeed even when the victim and adversary are run on different processor cores and do not share virtual memory by exploiting hardware characteristics of the last-level cache (LLC), which is shared across cores [8].

### A. Challenges for Current Cache Attacks

Despite their past successes, the viability of LLC cache attacks has been called into question on modern systems due to recent trends in processor design. To start with, many prior attacks [7], [9]–[11] can be mitigated out of the gate, as virtualized environments are now advised to disable shared virtual memory between VMs [12].

Without sharing virtual memory with a victim, the adversary must carefully consider the *cache hardware architecture* in mounting a successful attack. This is where problems arise. First, modern cache hierarchies are becoming *non-inclusive or exclusive*. Prior LLC attacks without shared virtual memory (e.g., [8]) rely on LLCs being *inclusive*, as this gives adversaries the ability to evict cache lines that are resident in the victim’s private caches. Non-inclusive cache behavior is significantly more complicated than that of inclusive caches (Section III). Second, modern LLCs are physically partitioned into multiple *slices*. Sliced LLCs notoriously complicate attacks, as the mapping between victim cache line address and cache slice is typically proprietary. Taken together, these challenges cause current LLC attacks to fail on modern systems (e.g., the Intel Skylake-X [13]).

Modern systems are moving to non-inclusive cache hierarchies due to the redundant storage that inclusive designs entail. Indeed, AMD servers have always used exclusive LLCs [11], and Intel servers are now moving to this design [13]. We expect the trend of non-inclusive caches to continue, as the cost of inclusive caches grows with core count (Section II).

### B. This Paper: Modernizing Cross-Core Cache Attacks

In this paper, we design a novel cross-core cache attack that surmounts all of the above challenges. Specifically, our attack does not require the victim and adversary to share cores or virtual memory, and succeeds on state-of-the-art *sliced non-inclusive caches*, such as those in Skylake-X [13]. Our key insight is that in a machine with non-inclusive cache hierarchies, we can still attack the directory structure. Directories are an essential part of modern cache hierarchies, as they maintain tracking information for each cache line resident in the cache hierarchy.<sup>1</sup> Since the directory must track *all* cache lines, and not just cache lines in the LLC, it offers an attack surface

<sup>1</sup>This should not be confused with the “directory protocol” used in multi-socket attacks that assume shared virtual memory between the adversary and victim [11].

similar to that of an inclusive cache. Indeed, our work suggests that conflict-based LLC attacks (on inclusive, non-inclusive or exclusive cache hierarchies) should target directories, not caches, as directories are a homogeneous resource across these different cache hierarchy designs.

**Contributions.** To summarize, this paper makes the following contributions:

1) We develop an algorithm to find groups of cache lines that completely fill a given set of a given slice in a non-inclusive LLC (called an *Eviction Set*). This modernizes prior work on Eviction Set creation, which only works for sliced inclusive LLCs.

2) Using our Eviction Sets, we reverse engineer the directory structure in Skylake-X, and identify vulnerabilities in directory design that can be leveraged by cache-based side channel attacks.

3) Based on our insights into the directory, we present two attacks. The first is a Prime+Probe attack on sliced non-inclusive LLCs. Our attack does not require the victim and adversary to share cores or virtual memory. The second attack is a novel, high-bandwidth Evict+Reload attack that uses multi-threaded adversaries to bypass non-inclusive cache replacement policies.

4) We use our two attacks to attack square-and-multiply RSA on the modern Intel Skylake-X server processor. Both of these attacks are firsts: although prior work implemented an Evict+Reload attack on non-inclusive LLCs, it cannot attack RSA due to its low-bandwidth. Finally, we construct efficient covert channels for sliced non-inclusive LLCs.

## II. BACKGROUND

### A. Memory Hierarchy and Basic Cache Structures

Modern high-performance processors contain multiple levels of caches that store data and instructions for fast access. The cache structures closer to the core, such as the L1, are the fastest, and are called higher-level caches. The ones farther away from the core and closer to main memory are slower, and are called lower-level caches. High-performance processors typically feature two levels of private caches (L1 and L2), followed by a shared L3 cache—also referred to as LLC for last-level cache.

The L1 cache is designed to be small (e.g., 32-64KB) and to respond very fast, typically within a few cycles. The L2 cache is slightly bigger (e.g., 256KB-1MB) and takes around 10-20 cycles. Finally, the LLC is designed to be large (e.g., several to tens of MBs) and has a latency of 40-60 cycles. The LLC latency is still much lower than the main memory access latency, which is on the order of 200-300 cycles.

A cache consists of the *data array*, which stores the data or code, and the *tag array*, which stores the high-order bits of the addresses of the data or code. The cache is organized in a number of *cache lines*, each one of size  $B$  bytes. The cache is typically set-associative, with  $S$  sets and  $W$  ways. A cache line occupies one way of a cache set. The set in which a cache line

belongs is determined by its address bits. A memory address is shown in Figure 1.




Fig. 1. Example of a memory address broken down into tag, index, and block offset bits. The actual bit field sizes correspond to the L2 and the LLC slice of the Intel Skylake-X system, as we reverse-engineered in Section V. We refer to the LLC slice set index as the LLC set index in this paper.

The lower  $\log_2 B$  bits indicate the *block offset* within a cache line. The next  $\log_2 S$  bits form the index of the set that the cache line belongs to. The remaining bits of the address form the *tag*. The tags of all the lines present in the cache are stored in the tag array. When a load or store request is issued by the core, the tag array of the L1 cache is checked to find out if the data is present in the cache. If it is a hit, the data is sent to the core. If it is a miss, the request is sent to the L2 cache. Similarly, if the request misses in L2 it is further sent to the LLC and then to main memory. Note that, generally, lower levels of the cache hierarchy have more sets than higher levels. In that case, cache lines that map to different LLC sets may map to the same L2 set, due to the pigeonhole principle.

### B. Multi-Core Cache Organization

The LLC in a modern multi-core is usually organized into as many slices (partitions) as the number of cores. Such an organization, shown in Figure 2, is helpful to keep the design modular and scalable. Each slice has an associativity of  $W_{\text{slice}}$  and contains  $S_{\text{slice}}$  sets.  $S_{\text{slice}}$  is  $1/N$  the total number of sets in the LLC, where  $N$  is the number of cores.

Core 0	LLC Slice 0	LLC Slice 4	Core 4
Core 1	LLC Slice 1	LLC Slice 5	Core 5
Core 2	LLC Slice 2	LLC Slice 6	Core 6
Core 3	LLC Slice 3	LLC Slice 7	Core 7

Fig. 2. Example of a sliced LLC design with 8 cores.

Processors often use an undocumented hash function to compute the slice ID to which a particular line address maps to. The hash function is designed to distribute the memory lines uniformly across all the slices. In the absence of knowledge about the hash function used, a given cache line can be present in any of the slices. Therefore, from an attacker’s perspective, the effective associativity of the LLC is  $N \times W_{\text{slice}}$ . The hash function used in Intel’s Sandybridge processor has been reconstructed in prior work [14], and found to be an xor of selected address bits. The slice hash function for the Skylake-X is more complex, as we find in Appendix B.

We now discuss two important cache design choices, and the trade offs behind them.

a) *Inclusiveness*: The LLC can be either *inclusive*, *exclusive*, or *non-inclusive* of the private caches. In an inclusive LLC, the cache lines in private L2 caches are also present in the LLC, whereas in an exclusive LLC, a cache line is never present in both the private L2 caches and in the LLC. Finally, in a non-inclusive LLC, a cache line in the private L2 caches may or may not be present in the LLC.

The inclusive design wastes chip area and power due to the replication of data. Typically, as the number of cores increases, the LLC size must increase, and hence the average LLC access latency increases [15], [16]. This suggests the use of large L2s, which minimize the number of LLC accesses and, therefore, improve performance. However, increasing the L2 size results in a higher waste of chip area in inclusive designs, due to the replication of data. The replication of data can be as high as the L2 capacity times the number of cores. Therefore, non-inclusive cache hierarchies have recently become more common. For example, the most recent server processors by Intel use non-inclusive caches [13], [15], [17]. AMD has always used non-inclusive L3s in their processors [11].

b) *Cache Coherence and Directories*: When multiple cores read from or write to the same cache line, the caches should be kept coherent to prevent the use of stale data. Therefore, each cache line is assigned a state to indicate whether it is shared, modified, invalid, etc. A few state bits are required to keep track of this per-line state in hardware in the cache tag array or directory.

Two types of hardware protocols are used to maintain cache coherence—*snoop-based* and *directory-based*. The snoop-based protocols rely on a centralized bus to order and broadcast the different messages and requests. As the number of cores is increased, the centralized bus quickly proves to be a bottleneck. Therefore, most modern processors use a directory-based protocol, which uses point-to-point communication. In a directory-based protocol, a *directory* structure is used to keep track of which cores contain a copy of a given line in their caches, and whether the line is dirty or clean in those caches.

In an inclusive LLC design, the directory information can be conveniently co-located with the tag array of the LLC slice. Since the LLC is inclusive of all the private caches, the directory state of all the cache lines in any private cache is present in such a directory. The hardware can obtain the list of sharer cores of a particular line by simply checking the line’s directory entry in the LLC. There is no need to query all the cores. However, the directory in a non-inclusive cache hierarchy design is more complicated, as we reverse engineered in Section V.

### C. Cache-based Side Channel Attacks

Cache-based side-channel attacks are a serious threat to secure computing, and have been demonstrated on a variety of platforms, from mobile devices [18] and desktop computers [2], [19] to server deployments [8], [10], [20]. Side-channel attacks bypass software isolation mechanisms and are difficult to detect. They can detect coarse-grained information such as when a user is typing [18] down to much more fine-grained

information such as a user’s behavior on the web [21], and even RSA [8], [22] and AES [19], [20], [23]–[25] encryption keys.

Cache-based side channel timing attacks leverage timing differences in memory accesses to deduce information about a victim workload. There are many such attacks (e.g., [2], [7]–[11], [18]–[21], [25]–[33]). In this paper, we refer to *target address* as the address which, when accessed, reveals information about victim behavior.

As an example of cache-based attack, consider the square-and-multiply exponentiation algorithm, which is widely used in many encryption algorithms such as RSA and ElGamal. Algorithm 1 shows an implementation. In the process of computing its output, the algorithm iterates over exponent bits from high to low. For each bit, it performs a *sqr* and a *mod* operation. Then, if the exponent bit is “1”, the algorithm performs a *mul* and a *mod* operation that are otherwise skipped. The target address can be the address of the line 6 in Algorithm 1. The instruction is only executed if the previous test on line 5 is true. If, at every iteration, an attacker can both evict the instruction from the cache, and probe to see if the victim has brought it back into the cache, then the full exponent is revealed. The key can then be derived from this exponent.

---

#### Algorithm 1: Square-and-multiply exponentiation.

---

**Input** : base  $b$ , modulo  $m$ , exponent  $e = (e_{n-1} \dots e_0)_2$   
**Output**:  $b^e \bmod m$

```

1  $r = 1$ 
2 for  $i = n - 1$  downto 0 do
3    $r = \text{sqr}(r)$ 
4    $r = \text{mod}(r, m)$ 
5   if  $e_i == 1$  then
6      $r = \text{mul}(r, b)$ 
7      $r = \text{mod}(r, m)$ 
8   end
9 end
10 return  $r$ 

```

---

In the first phase of a cache-based attack, the attacker first identifies the target address. This can be done using source code analysis or through a cache template attack [9]. In the second phase, the attacker gathers timing information to carry out the attack. In this phase, the attacker follows three steps:

- 1) Evict the target address from the resource in which it is resident.
- 2) Wait a time period during which the victim may access the target address.
- 3) Measure time variances to determine the location of the target address.

Three existing attacks highlight the three-step process outlined above. These are listed in order of difficulty, with the last being the most difficult for an attacker.

**Flush+Reload** [7], [33]: This attack is easiest thanks to the use of shared memory between the attacker and the victim. Shared memory is possible on many platforms due to page-deduplication and shared libraries [7]. The attacker can simply flush the target address using the `clflush` instruction (step 1). After waiting for a period of time (step 2), the attacker re-accesses the target address and measures the latency (step 3).

The attacker will expect a cache hit when accessing the memory flushed in step 1 if the victim has accessed the memory in the interval, and a miss otherwise. This attack is also referred to as a *flush-based* attack. This type of attack has serious limitations: it relies on `clflush` and cloud platforms are now advised turning sharing off, which disables this attack [34], [35].

**Evict+Reload** [18]: This attack also relies on shared memory for the reload operation, but does not flush or evict data using `clflush`, as not all architectures have a `clflush` instruction, and some defenses have suggested making `clflush` a privileged instruction [34] or disabling it all together [35]. In this case, the attacker uses cache conflicts to evict the target address (step 1). Specifically, the attacker accesses enough addresses mapped to the same cache set as the target address to evict the target address. The other steps are the same as Flush+Reload. Besides the additional complexity of creating such conflicts, this attack is otherwise the same as Flush+Reload, with an alternative flushing mechanism.

**Prime+Probe** [19]: This attack does not need shared memory. The attack steps are called prime, wait, and probe. In the prime step, the attacker evicts the target address from the cache by accessing a group of addresses mapped to the same cache set as the target address. During the wait step, the attacker waits. Finally, in the probe step, the attacker re-accesses the group of addresses used in the prime step, to measure victim activity. A cache miss in the probe step indicates that the victim has accessed the target address during the interval, and caused the eviction of one of the addresses accessed during the prime step from the cache.

Evict+Reload and Prime+Probe are referred to as *conflict-based* attacks due to fact that they exploit conflicts in cache structures.

#### D. Eviction Set

An *Eviction Set (EV)* is a collection of addresses that are all mapped to a specific cache set of a specific cache slice, and that are able to evict the current contents of the whole set in that slice. In a slice with  $W_{\text{slice}}$  ways, an eviction set must contain at least  $W_{\text{slice}}$  addresses to occupy all the ways and evict the complete contents of the set. We refer to *Eviction Addresses* as the addresses in an Eviction Set. In an inclusive LLC, both Evict+Reload and Prime+Probe use Eviction Sets to evict the target address from the private caches. Further, the probe operation in Prime+Probe measures the latency of accessing an Eviction Set to deduce the victim’s activity.

### III. THE CHALLENGE OF NON-INCLUSIVE CACHES

Previous cross-core cache side-channel attacks only work for inclusive cache hierarchies (e.g., [8], [34]). In a non-inclusive cache hierarchy, attackers must overcome the two main challenges that we describe next. In this discussion, we assume that the `clflush` instruction is disabled (Section II) [36] and that shared memory between attacker and victim has been disabled [12].

#### A. Lack of Visibility into the Victim’s Private Cache

In a non-inclusive cache hierarchy, an attacker running on a core seemingly cannot evict an address from another core’s private cache — i.e., it cannot create an *Inclusion Victim* in the second core’s cache. To see why, consider Figure 3, which shows a shared LLC and two private caches. The attacker runs on Cache 1 and the victim on Cache 0. The target line is shown in a light shade in Cache 0.

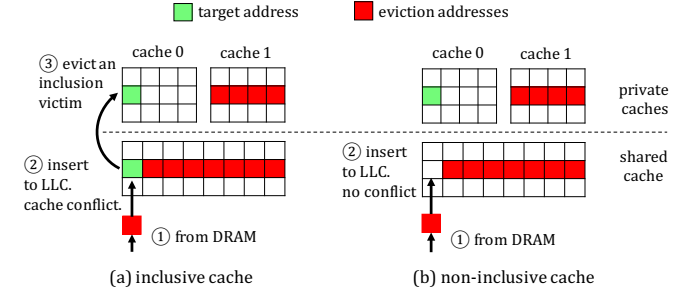


Fig. 3. Attempting to evict a target line from the victim’s private cache in inclusive (a) and non-inclusive (b) cache hierarchies.

Figure 3(a) shows an inclusive hierarchy. An LLC set contains lines from the attacker (in a dark shade) plus the target line from the victim (in a light shade). The attacker references an additional line that maps to the same LLC set. That line will evict the target line from the LLC, and because of inclusivity, also from private Cache 0, creating an inclusion victim. The ability to create these inclusion victims on another cache is what enables cross-core attacks.

Figure 3(b) shows a non-inclusive hierarchy. In this case, the target line is in the victim’s cache, and not in the LLC. Consequently, when the attacker references an additional line that maps to the same LLC set, there is no invalidation sent to Cache 0. The attacker has no way to create inclusion victims in the victim’s cache.

#### B. Eviction Set Construction is Hard

In a later section, we will show that we perform Prime+Probe and Evict+Reload attacks in non-inclusive cache hierarchies using an Eviction Set (EV). However, the algorithm used to create an EV in inclusive cache hierarchies [8] does not work for non-inclusive hierarchies. Creating an EV in non-inclusive hierarchies is harder. The reason is that it is less obvious what memory accesses are required to reliably evict the target line, which is currently in the private cache, from the entire cache hierarchy.

To see why, consider Figure 4, which shows a private cache and two slices of the shared LLC. Victim and attacker run on the same core. Figure 4(a) shows an inclusive hierarchy. The target line is also in one slice of the LLC. To evict the target from the cache hierarchy, the attacker only needs to reference enough lines to fill the relevant set in the corresponding slice of the LLC. This is because, as these lines fill the set, they will also fill the set in the private cache, and evict the target line from it. This is the EV, shown in a dark shade. The order and number of accesses to each of the lines in the EV required

to evict the target address is determined by the replacement algorithm used in the LLC slice.

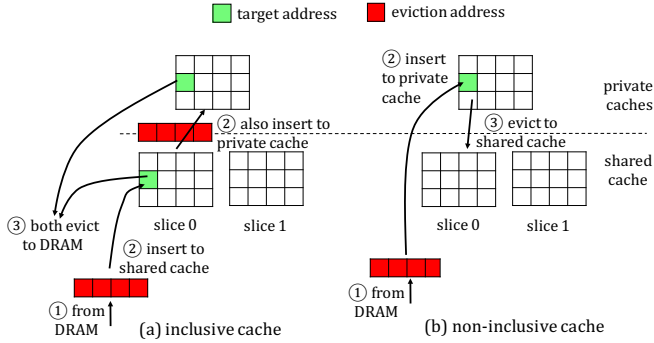


Fig. 4. Attempting to evict a target line in inclusive (a) and non-inclusive (b) cache hierarchies. Victim and attacker run on the same core.

Figure 4(b) shows a non-inclusive hierarchy. In this case, the target line is not in the LLC. As the core accesses the same dark cache lines as in Figure 4(a), the lines go first to the private cache, bypassing the LLC. The replacement algorithm used in the private cache will determine when the target line is evicted from the private cache into the LLC. When the target is evicted, it will go to one of the LLC slices, depending on the mapping of addresses to slices. Then, the core needs to evict enough lines into that LLC slice to create enough conflicts to evict the target line from the slice.

Overall, the order and number of accesses to each of the lines in the EV required to evict the target address is determined by multiple factors, including the replacement algorithm used in the private cache, the mapping of the line addresses to LLC slices, and the replacement algorithm used in the LLC slice. In the inclusive case, only the replacement algorithm in the LLC affects evictions.

We note that, in non-inclusive cache hierarchies, the replacement algorithms in the private caches and LLC slices can be quite sophisticated. What specific line is chosen as a replacement victim depends not only on the number and order of accesses to the lines, but also on the coherence state of the lines in the cache as well. Specifically, we have empirically observed that the replacement algorithm in the LLC slices tries to minimize the eviction of lines that are present in multiple private caches. This particular heuristic affects the ability to create an effective EV for Evict+Reload attacks, where lines are shared between attacker and victim.

### C. Attack Overview

We address these two challenges in two novel ways. To handle the difficulty of creating EVs, we propose a novel way to create EVs for non-inclusive caches in Section IV. Using EVs and other techniques, we reverse engineer the Skylake-X directory structure (Section V). This process reveals key insights into directory entry replacement policies and inclusivity properties. In particular, we derive conditions for when an attacker is able to use the directory to create inclusion victims in the private caches of a non-inclusive cache hierarchy. Based on our reverse engineering results, and the new EV

construction methodology, we design effective “Prime+Probe” and “Evict+Reload” attacks in non-inclusive caches hierarchies in Section VI.

## IV. CONSTRUCTING EVICTION SETS

In this section, we present an EV construction algorithm for non-inclusive caches. Recall that an EV is a collection of memory addresses that fully occupy a specific cache set of a specific LLC slice. This is a core primitive that we use to reverse engineer the directory (Section V) and later complete our attacks (Section VI).

Liu et al. [8] proposed an EV construction algorithm for sliced inclusive caches. However, it does not work for non-inclusive caches. The reason was discussed in Section III: line eviction from the cache hierarchy is less predictable because it depends on the L2 replacement algorithm, the mapping of line addresses to LLC slices, and the replacement algorithm used in the LLC slices. We fix these issues by developing a new implementation for `check_conflict`, an important subroutine used in Liu et al., that works on non-inclusive caches. We present evaluation results to demonstrate the effectiveness of our algorithm in Section VII-A.

### A. The Role of `check_conflict` in EV Construction Algorithm

The EV construction algorithm by Liu et al. [8] uses a function that we call `check_conflict(Address  $x$ , Collection  $U$ )` (called `probe` in [8]), shown in Algorithm 2. This function checks if the addresses in Collection  $U$  conflict with  $x$  in the LLC. The function should return `true` if  $U$  contains  $W_{\text{slice}}$  or more addresses, which are mapped to the same slice and set as  $x$ . The function should return `false` otherwise. The EV construction algorithm works only if this function has very low false positive and false negative rates.

---

#### Algorithm 2: Baseline `check_conflict` for inclusive caches.

---

```

1 Function check_conflict ( $x, U$ ):
2   access  $x$ 
3   for each  $addr$  in  $U$  do
4     | access  $addr$ 
5   end
6    $t$  = measure time of accessing  $x$ 
7   return  $t \geq LLC\_miss\_threshold$ 
8 end

```

---

To see why these requirements are important, consider how `check_conflict` is used in the EV construction algorithm. The high-level idea is to start with a collection  $U$  known to conflict with  $x$  in an LLC slice. Then, one removes an address  $y$  from the collection  $U$  and obtains a new collection  $U' = U - y$ . If the conflict with  $x$  disappears when checking against  $U'$ , then we know that  $y$  must contribute to the conflict. In such a case,  $y$  is considered to be in the EV for  $x$ . Clearly, the operation needs low false positive and false negative rates to precisely observe the disappearance of conflicts. Appendix A provides more details on how the algorithm is used, and why a high-accuracy implementation is important.

## B. New `check_conflict` Function

We first discuss why the `check_conflict` function designed by Liu et al. [8] has a high false negative rate when applied naïvely to non-inclusive caches. We then show how the function can be modified to work in non-inclusive caches. In the following discussion, we assume that all the addresses in  $U$  have the same LLC set index bits as  $x$ .

**Baseline `check_conflict` [8].** In Algorithm 2, the base function first accesses the target address  $x$ , ensuring that the line is cached. It then accesses all the addresses in  $U$ . If a later access to line  $x$  takes a short time, it means that the line is still cached. Otherwise, it means that the line has been evicted out of the cache due to cache conflicts caused by  $U$ . Thus, the access latency can be used to determine whether  $U$  contains enough addresses to evict  $x$ .

When applied to non-inclusive caches, this function has a high false negative rate. Specifically, when  $U$  contains enough addresses to evict  $x$ , the function is supposed to return `true`, but it returns `false`. To see how the false negatives happen, consider a minimal  $U$ , which has exactly  $W_{\text{slice}}$  addresses mapped to the same LLC slice as  $x$ . On non-inclusive caches, when accessing  $U$ , some of these  $W_{\text{slice}}$  lines may never be evicted from the L2 into the LLC. Hence, these addresses do not have a chance to conflict with  $x$  in the LLC, and  $x$  is not evicted, resulting in a false negative. Moreover, since the replacement algorithm of L2 is neither LRU nor pseudo-LRU, simply accessing  $U$  multiple times, does not guarantee a small false negative rate, as we validate in Section VII-A.

**Naïve flushL2 function.** To reduce the false negative rate, we need to flush all the lines in  $U$  from the L2 to the LLC. It would be convenient if we had a special instruction to do so, but such an instruction does not exist in x86. Hence, we leverage L2 conflicts to achieve the flush effect.

We create an extra collection of addresses, called  $L2\_occupy\_set$ , which contains  $W_{L2}$  addresses mapped to the same L2 set as  $U$ . When accessed,  $L2\_occupy\_set$  forces all lines in  $U$  to be evicted to the LLC. Our modified `check_conflict` function is shown in Algorithm 3. After accessing  $x$  and all the addresses in  $U$  as in the base function (line 2-5), the addresses in  $L2\_occupy\_set$  are accessed (line 6-8). In this way, every line in  $U$  gets a chance to conflict with  $x$  in LLC, and we can significantly reduce the false negative rate.

---

### Algorithm 3: New `check_conflict` for non-inclusive caches.

---

```

1 Function check_conflict ( $x, U$ ):
2   access  $x$ 
3   for each  $addr$  in  $U$  do
4     | access  $addr$ 
5   end
6   for each  $addr$  in  $L2\_occupy\_set$  do
7     | // the flushL2 operation: to evict  $U$  from L2 to LLC
8     | access  $addr$ 
9   end
10   $t$  = measure time of accessing  $x$ 
11  return  $t \geq LLC\_miss\_threshold$ 

```

---

However, this naïve flushL2 approach has a high false positive rate. A false positive can occur when  $U$  does not contain enough addresses to evict  $x$  from the LLC slice, but then some addresses in  $L2\_occupy\_set$  end up getting evicted to the LLC, and they evict  $x$  from the LLC. In this case, the function is supposed to return `false`, but it returns `true`.

**Reliable flushL2 function.** To reduce the false positive rate in the naïve flushL2 function, we need to make sure accesses to  $L2\_occupy\_set$  do not interfere with the conflicts between  $U$  and  $x$  in the LLC. We can achieve this by leveraging the one-to-many set mapping relationship between L2s and LLCs.

For a reliable design, we select  $L2\_occupy\_set$  such that its addresses are mapped to the same L2 set as addresses in  $U$ , but to a *different* LLC set than used by addresses in  $U$  (and  $x$ ). As mentioned before, upper level caches like the L2 contain fewer cache sets than lower level caches like the LLC (Section II-A). For example, in Skylake-X, the L2 has 1024 sets, while an LLC slice has 2048 sets. Correspondingly, the L2 uses 10 bits (bits 6-15) from the physical address as the set index, while the LLC slice uses 11 bits (bits 6-16). Therefore,  $L2\_occupy\_set$  can be constructed by simply flipping bit 16 of  $W_{L2}$  addresses in  $U$ . Such addresses can be used to evict  $U$  from the L2 but do not conflict with  $U$  in the LLC.

In summary, we designed a reliable `check_conflict` function with both low false positive rate and low false negative rate, which can be effectively used in the EV construction algorithm for non-inclusive caches. We evaluate the effectiveness of the function in Section VII-A. For independent interest, we use our EV creation routine to partially reverse engineer the Skylake-X slice hash function in Appendix B.

## V. REVERSE ENGINEERING THE DIRECTORY STRUCTURE IN INTEL SKYLAKE-X

In this section, we leverage the Eviction Sets (EVs) we constructed in Section IV to verify the existence of the directory structure in an 8-core Intel Core i7-7820X processor, which uses the Intel Skylake-X series microarchitecture. We also provide detailed information about the directory's associativity, inclusivity, replacement policies (for both private and shared data), and interactions with the non-inclusive caches. These insights will be used to complete our attack in Section VI. Skylake-X is a server processor for cloud computing and datacenters. A comparison of the cache parameters in this processor with previous Skylake series processors is listed in Table I.

	Skylake-S	Skylake-X/Skylake-SP
L1-I	32KB, 8-way	32KB, 8-way
L1-D	32KB, 8-way	32KB, 8-way
L2	256KB/core 16-way, inclusive	1MB/core 16-way, inclusive
LLC	2MB/core 16-way, inclusive	1.375MB/core 11-way, non-inclusive

TABLE I  
CACHE STRUCTURES IN INTEL PROCESSORS.

Relative to the older Skylake-S processor, the Skylake-X/SP LLC is non-inclusive and, correspondingly, Skylake-X/SP can




Fig. 5. Access latency for a single cache line when the line is in different locations in Intel Skylake-X cache hierarchy.

support larger L2 caches relative to the LLC. The L2 in Skylake-X/SP grows to 1 MB per core, which is 4 times larger than before, while the LLC size shrinks from 2 MB per core to 1.375 MB per core. The associativity in the LLC is also reduced from 16-way to 11-way.

#### A. Timing Characteristics of Cache Access Latencies

We first conduct a detailed analysis of the timing characteristics of the cache access latencies on Skylake-X. This information can be used to infer the location of a specified cache line, and is useful in reverse engineering the directory structure.

For each cache location, we measure the access latency by using the `rdtsc` instruction to count the cycles for one access. We use the `lfence` instruction to make sure we get the timestamp counter after the memory access is complete as suggested in [37]. Thus, all the latencies presented include delays introduced by the execution of `lfence`.

Figure 5 shows the distribution of cache access latencies for different locations in the non-inclusive cache. We have binned latencies in the figure at the granularity of 25 cycles. Note that each latency is for accessing a single cache line. From the figure, L1 and L2 access latencies are short, both below 50 cycles. An LLC access takes around 100 cycles, and a DRAM latency is around 350 cycles.

A remote L2 access occurs when a thread accesses a line that is currently in another core’s L2. From the figure, a remote L2 access takes around 200 cycles, which is shorter than the DRAM latency. We leverage the difference between the remote L2 access latency and the DRAM latency in the “Evict+Reload” attack to infer victim’s accesses.

In the following experiments, we rely on the access latency to infer the location of a cache line.

#### B. Existence of the Sliced Directory

The following experiment is designed to verify the existence of a directory and its structure. In each round of the experiment, a single thread conducts the following three steps in order.

- 1) Access target cache line  $x$ .
- 2) Access  $N$  cache lines. In the “reference” setup, these cache lines have the same LLC set index bits as  $x$ , and can be mapped to different LLC slices. In a “sameEV” setup, these lines are from the same EV as  $x$ , and are therefore mapped to the same LLC slice (Section IV).

- 3) Access the target cache line  $x$  again while measuring the access latency.

Generally, step 2 is repeated multiple times to bypass the noise due to the cache replacement policy. The access latency for step 3 is shown in Figure 6, as a function of the number of addresses in the EV.




Fig. 6. Target line access latency as a function of the number of addresses in the EV in experiments to verify the existence of the directory.

According to the timing characteristics in Figure 5, we know that latencies around 40, 100 or 300 cycles indicate that the target line is in local L2, LLC or DRAM, respectively. In the “reference” configuration, if 16 or more lines are accessed in step 2, the target line is evicted from L2 to LLC. These evictions are caused by L2 conflicts because the L2 associativity ( $W_{L2}$ ) is 16. We start to observe that the target line is evicted to DRAM when more than 75 addresses are accessed in step 2. This number is less than 104 ( $W_{L2} + N_{slice} \times W_{slice}$ ) since the slice hash function is imperfect.

In the “sameEV” setup, we observe L2 misses when 12 cache lines are accessed in step 2, before reaching the L2 associativity. Moreover, the target line is evicted out of the LLC when 21 lines are accessed, even though the L2 cache and one LLC slice should be able hold up to 27 lines ( $W_{L2} + W_{LLC}$ ). The 12 versus 16 (L2 case) and 21 versus 27 (LLC case) indicate that there exists some bottleneck, other than the L2 and LLC slice associativity. This indicates the presence of some set-associative on-chip structure, and that conflicts in this structure can cause L2 and LLC evictions. We further verify the inclusivity property of this structure in later sub-sections.

We know the structure is sliced and looked-up using the LLC slice hash function. Notice that addresses conflict in this structure only if they are from the same LLC slice as in the “sameEV” configuration. Addresses from different LLC slices do not cause conflicts in this structure, as seen in the “reference” configuration.

We can deduce the structure’s associativity for private cache lines from the following experiment. The associativity seen for L2 lines is 12, and the associativity seen for L2 and LLC lines is 21. We can check the number of sets in the structure as follows. Since addresses that belong to the same EV can have different values for the 17-48 bits, we know that the bits above the 16th bit are not used in the set index. Furthermore, we found that addresses with different  $b_{16}$  never conflict, i.e. they are mapped to different sets. Thus, the set index for this structure uses bits 6-16 from the physical address. We conclude that each slice of this structure has the same number of sets as an LLC slice, namely 2048.

1. There exists a set-associative structure that operates alongside the non-inclusive caches. Contention on this structure can interfere with cache line states in both the L2 and LLC.
2. The structure is sliced and is looked up using the LLC slice hash function. Each slice has the same number of sets as one LLC slice.
3. The associativity of the structure for L2 lines is 12, and the associativity of the structure for L2 and LLC lines is 21.

### C. Inclusivity and Associativity for Private Cache Lines

We use the term private cache line to refer to a line which has been accessed by a single core. The term shared cache line is used to refer to a line which has been accessed by multiple cores. We observed that the non-inclusive LLC cache in Skylake-X behaves differently towards private and shared cache lines, resulting in different conflicting patterns.

We conduct a two-thread experiment as follows to reverse engineer the inclusivity of the set-associative structure and the cache for private lines. The two threads are pinned to different cores.

- 1) Thread A accesses target line  $x$ .
- 2) Thread B accesses  $N$  cache lines. Different lines are selected for “reference” and “sameEV” setups as in the previous experiment.
- 3) Thread A accesses target line  $x$  again and measures the access latency.

The access latencies in step 3 are shown in Figure 7, as a function of the number of addresses in the EV.




Fig. 7. Target line access latency as a function of the number of EV addresses in experiments to analyze the inclusive property for private cache lines.

In the “reference” setup, the target line in thread A’s private L2 is never evicted. Due to the non-inclusive property of the

LLC, thread B is unable to interfere in thread A’s private cache state, and hence loses any visibility into thread A’s private cache accesses. However, in the “sameEV” setup, the target line is evicted from the L2 to LLC when 12 lines are accessed by thread B, and it is further evicted to DRAM when thread B accesses 21 lines.

This experiment shows that the structure is shared by all the cores, and conflicts on this structure can interfere with L2 and LLC cache states. In particular, the “sameEV” configuration shows that we *can create inclusion victims across cores*, since lines in the L2 can be evicted due to the contention on the structure. We can safely conclude that the structure is inclusive to all the lines in the cache hierarchy, including L2 and LLC. This characteristic can be leveraged by an attacker to gain visibility into a victim’s private cache and build “Prime+Probe” attacks. Moreover, the experiment also confirms the same associativity across cores as last experiment.

4. The structure is shared by all cores.
5. The structure is inclusive, and contention on the structure can cause inclusion victims.

### D. Inclusivity and Associativity for Shared Cache Lines

The following experiment is designed to reverse engineer the inclusivity and associativity of the structure for shared cache lines. We use 2 or 3 threads in different modes.

- 1) Thread A and B both access the target cache line  $x$  to ensure the line has been marked as “shared” by the processor.
- 2) In “1evictor” mode, thread B accesses  $N$  cache lines; in the “2evictors” mode, thread B and C access  $N$  cache lines to put those lines into the “shared” state. In both modes, different cache lines are selected for “reference” and “sameEV” setups as in previous experiments.
- 3) Thread A accesses the target line  $x$  again and measures access latency.

To summarize: in the “1evictor” mode, only  $x$  is in the shared state; in the “2evictors” mode, both  $x$  and the  $N$  eviction lines are in the shared state. We show access latencies for step 3 in Figure 8.




Fig. 8. Target line access latency as a function of the number of EV addresses in experiments to analyze the inclusive property for shared cache lines.

In the “1evictor\_sameEV” setup, the shared target line is never evicted out of thread A’s private L2. However, we showed that it is enough to cause remote L2 evictions of *private* lines by accessing 12 addresses from the same EV as in Figure 7. Comparing the two cases, we can infer that the cache coherence



state—namely whether a line is shared or not—plays a role in the cache line replacement policy. The replacement policy prefers not to evict shared cache lines.

In the “2evictors\_sameEV” setup, threads B and C are able to evict the target line out of LLC by accessing 11 shared lines, while in the “2evictors\_ref” setup, we begin to observe stable LLC misses when around 85 lines are accessed. The characteristics in “2evictors\_sameEV” indicates the associativity of the inclusive structure for shared cache lines is 11. Moreover, this experiment indicates how an attacker can use shared eviction lines to evict a shared target line out of the cache hierarchy, which we will leverage to build stable and efficient “Evict+Reload” attacks.

6. The cache replacement policy takes into account the coherence state and prefers not to evict cache lines which have been accessed by multiple cores.
7. The associativity of the inclusive structure for shared cache lines is 11.

### E. Putting It All Together: the Directory Structure

We can infer the inclusive structure is a directory for two reasons. First, Intel has used a directory-based coherence protocol since Nehalem [38].<sup>2</sup> To support the directory-based protocol, structures to store the presence information for all the lines in the cache hierarchy are required. Thus the directory, if it exists, must be inclusive like the structure we found. Second, the replacement policy in the inclusive structure takes coherence state into account. As discussed in Section II, the coherence state is maintained by the directory. This again supports our hypothesis that the inclusive structure is the directory. We use the term “directory” to refer to the inclusive structure in the rest of the paper.

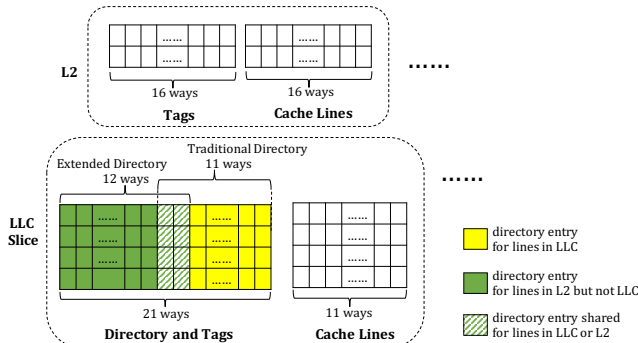


Fig. 9. Reverse engineered directory structure.

Figure 9 shows the possible structure of the directory in one LLC slice. From Section V-B: the directory is sliced, is looked up using the same LLC slice hash function, and has the same number of sets as the LLC. Such a design makes sense. An LLC slice can co-locate with its directory, and cache line look-ups can be performed locally without introducing extra network traffic.

<sup>2</sup>Sometimes, they call the structure to store directory information as “core valid bits”.

From Section V-B, each “directory slice” has 21 ways in total (denoted  $W_{dir} = 21$ ) for all the lines in the cache, including the L2 and LLC. From Section V-B and V-C, there are maximally 12 ways can be used for lines present in the L2 but not in LLC, which we call “Extended Directory” (ED). We denote ED associativity as  $W_{ED} = 12$ . From public documentation (Table I), we know that the LLC slice not including the ED, space we call the “Traditional Directory”, is 11-way set associative. Or,  $W_{slice} = 11$ . One may expect  $W_{dir} = W_{ED} + W_{slice}$  but  $21 < 12 + 11$ . From this mismatch, we infer that 2 ways in each directory slice are dynamically shared between the traditional directory and the extended directory.

Finally, the migration between the extended directory and traditional directory operates as follows. An ED conflict causes an entry to be migrated from the extended directory to the traditional directory, and the corresponding cache line to be evicted from an L2 to the LLC. A conflict in the traditional directory causes an entry to be removed from the whole directory slice, which causes the corresponding cache line to be evicted out of the entire cache hierarchy. For private lines, 21 addresses are needed to cause a traditional directory conflict. From Section V-B, the private lines will take up the ED and traditional directory, after which we see conflicts. For shared lines, 11 addresses are needed to cause conflicts. We found that shared lines, after being accessed a sufficient number of times, migrate from the ED to the traditional directory. In this case, the line lives in multiple L2s and in the traditional directory at the same time.<sup>3</sup> This is likely a performance optimization as LLC hits are faster than remote L2 hits (Figure 5). Thus, heavily shared lines should be accessible from the LLC.

The directory structure described above perfectly matches our reverse engineering results. Even though the actual implementation may use a slightly variant of design, our interpretation of the structure is helpful in understanding the directory and cache interactions and designing the attacks.

### F. The Root Cause of the Vulnerability

The directory in non-inclusive caches is *inclusive*, since it needs to keep information for *all* the cache lines that are present in the cache hierarchy. This inclusivity can be leveraged to build cache-based attacks. An attacker can create conflicts in the directory to enforce cache line evictions from a victim’s private cache, and create inclusion victims.

Considering the usage of directories, the directory capacity ( $W_{dir} \times S_{dir}$ ) should be large enough to hold information for all the cache lines. In this case, how can it be possible to cause a directory conflict before causing cache conflicts? This section answers this question by analyzing the root cause of the vulnerability in cache and directory design.

The root cause is that *the directory associativity is smaller than the sum of the cache associativity which it is supposed to support*. More specifically, a directory conflict is possible if

<sup>3</sup>This is consistent with the cache being non-inclusive. Non-inclusive means the cache is not always inclusive but may be for certain lines in certain circumstances.

the following conditions are satisfied, where ED indicates the directory entries used by L2 cache lines.

$$W_{\text{ED}} < W_{\text{L2}} \times N_{\text{L2}}$$

or  $W_{\text{dir}} < W_{\text{L2}} \times N_{\text{L2}} + W_{\text{slice}}$

where  $N_{\text{L2}}$  is the number of L2s in the system (usually equal to the number of cores).

We find these conditions commonly exist in directory designs for performance purposes. First, as the number of cores increase on the chip, hardware architects try to avoid unified directory designs, since a unified directory suffers from contention and introduces very high performance overhead. In a distributed directory design, each directory slice is only responsible to handle part of the coherence traffic, and it is unnecessary to make the associativity of each slice as high as the total L2 associativity. Second, hardware architects try to limit the associativity of directories, since a high-associativity design increases look-up latency as well as consumes more energy and area. For example, an 8-core Intel Skylake-X processor with 16-way L2s would require a 128-way ED slice to avoid ED conflicts, which is very expensive to build. Moreover, considering there are 8 ED slices, the total ED capacity will be 8 times the total L2 capacity, which is a significant waste of resources.

We found that the above conditions do not hold in several AMD processors. At the same time, we found that the corresponding memory and coherence operations are much slower on AMD machines. This is evidence supporting our claim that sliced directories are important structures in high performance processor design. See Appendix C for more details.

## VI. ATTACK STRATEGIES

Leveraging the EV construction algorithm in Section IV and our reverse engineering results of the directory structure in Section V, we can now build effective side channel attacks in *non-inclusive* cache hierarchies. In this section, we first present our Prime+Probe attack targeting the *Extended Directory* (ED). We then show an innovative multi-threaded Evict+Reload attack to achieve fine-grained monitoring granularity. Finally, for completeness, we provide a brief discussion on Flush+Reload attacks.

### A. Prime+Probe

We propose a Prime+Probe attack targeting the directory structure in a non-inclusive cache hierarchy. To the best of our knowledge, this is the first Prime+Probe attack on non-inclusive caches. We first present our customized Prime+Probe attack on Skylake-X, and then discuss how to generalize the attack to other vulnerable platforms.

On Intel Skylake-X, an attacker can leverage the inclusivity of the ED to gain visibility to a victim’s private cache state. Before the attack, the attacker uses the EV construction algorithm in Section IV to construct an EV, which is mapped to the same LLC slice and set as the target address. Since the ED is sliced and looked-up using the LLC slice hash function,

it follows that EVs are mapped to the same ED set and slice as the target address. Thus, EVs can be used in the prime and probe steps.

Then our Prime+Probe attack follows the same steps as a general Prime+Probe attack. In the prime step, the attacker accesses  $W_{\text{ED}}$  EV lines (Section V-E) to occupy all the ways within the ED set, and evicts the target line from the victim’s L2 to the LLC. During the wait interval, if the victim accesses the target line, it causes an ED conflict and one of the EV addresses will be evicted from the attacker’s private cache to the LLC. In the probe step, when the attacker accesses the EV addresses again, it will observe the LLC access. Alternatively, if the victim does not access the target line in the wait interval, the attacker will observe only L2 hits in the probe step. After the probe step, the ED set is fully occupied by EV addresses, and can be used as the prime step for the next attack iteration.

**Attack granularity.** The attack granularity is determined by the time per attack iteration, which is composed of wait time and probe time. The more efficient the probe operation is, the finer granularity an attack can achieve.

In our ED-based Prime+Probe attack, the probe time is very short. The attacker only needs to distinguish between local L2 latency and LLC latency, which is much shorter than the probe time in inclusive cache attacks, where the attacker needs to distinguish between LLC latency and DRAM latency.

**Generalizing the attack.** The attack above is customized for Intel Skylake-X. We now discuss how to generalize the attack to other vulnerable platforms which satisfy the conditions discussed in Section V-F.

First, a special characteristic within the Skylake-X is that the ED associativity is smaller than the L2 associativity ( $W_{\text{ED}} \leq W_{\text{L2}}$ ), which allows us to trigger ED conflicts using a single attacker thread. If this condition is not satisfied, we can still mount a Prime+Probe attack with *multiple* attacker threads, running on different cores, as long as  $W_{\text{ED}} \leq W_{\text{L2}} \times (N_{\text{L2}} - 1)$ . For example, consider the case that  $W_{\text{ED}} = W_{\text{L2}} \times (N_{\text{L2}} - 1)$ . The attacker can use  $(N_{\text{L2}} - 1)$  threads running on all the cores except for the victim’s core, where each thread accesses  $W_{\text{L2}}$  addresses to occupy the ED set.

Second, the directory in Skylake-X uses the same hash function as the LLC, thus we can directly use the EVs constructed for LLC slices to create directory conflicts. If the sliced directory uses a different hash function, which is rare for performance reasons, the attack should still work but will need a new EV construction algorithm for the directory.

### B. Evict+Reload

On non-inclusive caches, an attacker could leverage the directory’s inclusivity to build Evict+Reload attacks using a similar approach as in Prime+Probe. However, the evict operation in Evict+Reload is more challenging than the prime operation since the target line is shared by the attacker and the victim. As we showed in Section V-D, the cache replacement policy takes into account the coherency state—namely that the target line is *shared*—and prefers not to evict the directory entries for shared lines.

We propose an innovative multi-threaded Evict+Reload attack that can achieve fine-grained monitoring granularity. Specifically, the attacker consists of two threads: an attacker thread which executes the evict and reload operations, and a helper thread to assist evicting the shared target line. The attacker thread and the helper thread share all the eviction addresses. The main idea is that the helper thread switches the eviction addresses to the *shared* coherence state, similar to the “2evictors\_sameEV” setup in Section V-D. This brings eviction addresses to the same replacement priority as the target address in the directory. Throughout the entire attack, the helper thread is running in the background and continuously accessing the addresses in the EV in order to keep those addresses in the cache. By using helper thread, the attacker thread only needs to access the EV addresses a limited number of times in the evict step to ensure eviction of the target line from LLC. If the victim accesses the target line during the wait interval, the attacker will see a remote L2 access latency in the reload step; otherwise, it will observe a DRAM access latency.

### C. Flush+Reload

A Flush+Reload attack on non-inclusive caches follows the same procedure as the one on inclusive caches. This process has been referred to as Invalidate+Transfer [11] on AMD’s non-inclusive caches. We evaluate this attack on Intel’s platform for completeness, though it is not necessary to leverage our new attack channel on directories.

The attacker uses the `clflush` instruction to evict the target address from all levels of the cache hierarchy to DRAM. If during the wait interval, the victim accesses the target address as discussed in Section II-C, the line will be brought into the victim’s local L2, which is the attacker’s remote L2. In the measurement phase, the attacker reloads the target address and measures the access latency. The latency is then compared to the DRAM access latency and remote L2 latency to figure out the location of target address, thus deducing the victim’s behavior.

## VII. EVALUATION

### A. Effectiveness of `check_conflict` Functions

We evaluate the effectiveness of the `check_conflict` function by measuring the false positive rates and the false negative rates. We consider three designs, the baseline function proposed by Liu et al. (“no\_flushL2”), and two modified functions discussed in Section IV, i.e. “flushL2\_naive” and “flushL2\_reliable”.

We obtained 8 EVs and confirmed the correctness by checking their conflicting behaviors as in Section V. To measure the false positive rate, we set the argument  $U$  as a collection of addresses with 10 addresses ( $< W_{\text{slice}}$ ) from the same EV as  $x$ , and 5 addresses from each of the other EVs. We then count the number of times when the function mistakenly returns *true*. To measure the false negative rate, an extra address from the same EV as  $x$  is added to the collection  $U$ , so that  $U$  contains 11 eviction addresses ( $= W_{\text{slice}}$ ). Similarly we count the number of times when the function mistakenly returns *false*. In each

implementation, the eviction operation (line 3-5 of Algorithm 2, line 3-8 in Algorithm 3) is repeated multiple times. Figure 10 shows how false positive rate and false negative rate change with the number of eviction operations performed.

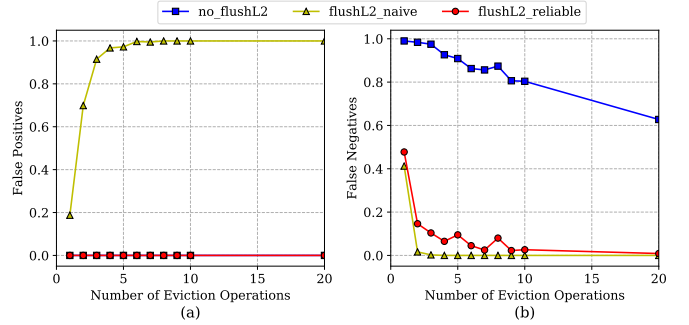


Fig. 10. Comparing effectiveness of different `check_conflict` functions.

In Figure 10(a), both “no\_flushL2” and “flushL2\_reliable” have no false positives. “flushL2\_naive” has a much higher false positive rate due to the extra conflicts introduced by the `L2_occupy_set`. In Figure 10(b), both “flushL2\_naive” and “flushL2\_reliable” can achieve very low false negative rate when eviction operations are repeated around 10 times. The false negative rate of the “no\_flushL2” approach stays high even though the evictions are performed 20 times. In conclusion, our reliable flushL2 approach in `check_conflict` function is effective and can achieve both low false negative rate and false positive rate.

### B. Extended Directory Timing Characteristics

As we leverage ED conflicts to construct our Prime+Probe attack, it is very important to understand their timing impact on cache access latencies, as shown in Figure 11. The figure shows the access latency of a number of addresses from the same EV. In the “no\_EDconf” case, we simply measure the latency of EV accesses. In the “1\_EDconf” case, between two measurements, we use a different thread on another core to issue one access to the same ED set to cause one ED conflict. Thus, the latency in “no\_EDconf” is the expected probe latency with no victim accesses during wait intervals, while the “1\_EDconf” latency corresponds to the expected probe latency when victim accesses the target line.




Fig. 11. Prime/Probe Latency

A high-resolution low-noise Prime+Probe attack requires the probe operation to be efficient and clearly distinguishable. From Figure 11,  $W_{\text{ED}}$  (12) is the optimal number of probe addresses we should use in Prime+Probe. First, the impact

of ED conflicts is large and clearly observable. The timing difference between no ED conflicts and a single ED conflict is around 90 cycles. Second, accessing 12 addresses takes very short time, around 230 cycles with a ED conflict. With such efficient prime/probe operation, we can do fine-grained monitoring. It is also feasible to use 13-15 addresses, but it is not optimal due to the longer access latency and slightly larger variance.

### C. Covert Channel on Extended Directories

We demonstrate a covert channel that utilizes the extended directory between two different processes. One process serves as the sender and the other as the receiver. The sender and receiver run on separate cores, and each utilizes 7 addresses that are mapped to the same LLC slice. Together there are 14 addresses, which are enough to cause measurable ED conflicts.




Fig. 12. The upper plot shows receiver’s access latencies on a slice not being used for the covert channel, while the lower one shows the one used in the covert channel. Sender transmits sequence “101010...”.

Since we have not reverse engineered the slice hash function, the sender and the receiver cannot directly negotiate which slice to use. Before communication, the receiver scans all slices to find the one the sender is using. The sender transmits a bit “0” by idling for 5000 cycles, and keeps accessing the 7 addresses for 5000 cycles to transmit a bit “1”. The receiver decodes the two states by taking latency samples every 1000 cycles. On our 3.6GHz machine, it takes 5000 cycles to transmit one bit, thus the bandwidth is 0.2Mbit/s. With a better protocol than we are using, the bandwidth can be further improved.

Figure 12 shows the results of our reliable covert communication channel. The upper plot shows the latencies the receiver observes when accessing the wrong slice which are all low as they are reflective of L1 cache hits, around 80 cycles. It is clear to see the sender’s message of “101010” in the lower plot. The receiver observes a  $\sim 400$  cycle latency due to ED conflicts when decoding a “1” bit, which is easily differentiated from the  $\sim 80$  cycle L1 hits for a “0” bit.

### D. Side Channel Attacks on the Square-and-Multiply Exponentiation Algorithm

We evaluate the effectiveness of our side channel attacks on the square-and-multiply exponentiation vulnerability in GnuPG 1.4.13. The implementation is similar to the one presented in Algorithm 1 in Section II. As discussed before, a victim’s accesses on function `sqr` and `mul` can leak

the value of exponent. In GnuPG, these two functions are implemented recursively, thus the target address identifying each function will be accessed multiple times throughout the execution of either operation. We show how this algorithm remains vulnerable on non-inclusive caches by attacking it with Prime+Probe, Evict+Reload and Flush+Reload attacks.

1) *Flush+Reload*: We evaluate a cross-core Flush+Reload attack on this new platform for completeness. The victim and the attacker run on separate cores. The flush and reload operations are used on the addresses located at the entry of the `sqr` and `mul` functions. We use a wait time of 2000 cycles between the flush and reload.




Fig. 13. Access latencies measured in the reload operation in Flush+Reload. A sequence of “1001110101” can be deduced as part of the exponent.

Figure 13 shows the time measurement of the reload operation for 100 samples. A low latency reload operation, less than 250 cycles, indicates the victim has accessed the target address during the wait interval. A high latency, around 350 cycles, means the victim has not accessed the target address. According to the algorithm, an access on `sqr` followed by an access on `mul` indicates a bit “1”, and two consecutive accesses on `sqr` without `mul` accesses in the between indicate a bit “0”. From Figure 13, we can see that each `sqr` operation completes after 3 samples, or about 6000 cycles. Leveraging this information, the attacker is able to deduce part of the exponent as “1001110101”.

In Flush+Reload, errors stem from times when the attacker’s flush operation overlaps with victim accesses. Such occurrences cause lost bits.

2) *Prime+Probe*: In our Prime+Probe attacks, we use 12 probe addresses from an eviction set for the target address, and use 500 cycles as the attacker wait interval. We are able to monitor with such small granularity due to the efficient probe operation on the ED. We only monitor one target address, i.e. the address located at the entry of `mul` function.




Fig. 14. Access latencies measured in the probe operation in Prime+Probe. A sequence of “01010111011001” can be deduced as part of the exponent.

Figure 14 shows the access latencies measured in the probe operation as results of our Prime+Probe attack for 400 samples. If there is no victim access of the target address, the probe operation will see L2 hits for all the probe addresses without ED conflicts, taking around 160 cycles. Otherwise, if the victim accesses the target address, ED conflicts will be observed, resulting in long access latency, around 230 cycles. We do not track victim accesses on the `sqr` function; this the same approach taken in [8]. Instead, the number of square operations can be deduced from the length of the interval between two consecutive multiply operations. The attacker can deduce a sequence of “01010111011001” as part of the exponent from Figure 14.

In Prime+Probe attacks, most errors stem from the time it takes the hardware to recover from the ED “busy” state, which may lead to observing a multiply operation for more samples than it actually executed. Moreover, variances due to ED conflicts also cause the wrong number of square operations to be estimated.




Fig. 15. Access latencies measured in the reload operation in Evict+Reload. A sequence of “010101110110101” can be deduced as part of the exponent.

3) *Evict+Reload*: Our novel Evict+Reload attack utilizes 1 attacker thread and 2 helper threads. The 2 helper threads access an EV shared with the attacker thread (Section VI-B). The attacker thread accesses 16 addresses mapped to the same L2 set as the target line 6 times and then access the EV once. We tested multiple eviction approaches, and found this method is highly reliable, and also very efficient, only taking around 1200 cycles. We monitor both the square and multiply operations and use 4000 cycles as the wait interval.

Figure 15 shows the access latencies measured in the reload step as the results for the Evict+Reload attack for 100 samples. The figure can be interpreted in the same way as the one for Flush+Reload, and the attacker can decode the part of the exponent as sequence “010101110110101”.

Compared to Flush+Reload, the Evict+Reload attack on non-inclusive caches tend to suffer more errors. Since the evict operation takes longer than the flush operation, the probability that the evict step overlaps with the victim’s access is higher.

## VIII. RELATED WORK

There has been a variety of cache-based side-channel attacks in literature. We start by reviewing the attacks most closely related to our attack, those on non-inclusive caches. We then briefly discuss side-channel attacks on inclusive caches.

### A. Attacks on Non-Inclusive Caches

There are two known attacks on non-inclusive caches that require page-sharing [11], [18]. ARMageddon [18] leverages Evict+Reload to attack a non-inclusive ARM LLC. Irazoqui et al. [11] leverage Flush+Reload to attack a non-inclusive AMD dual-socket machine. First, both works rely on shared virtual memory. Second, neither of these works address the complexities stemming from sliced caches. Thus, our result is more general.

ARMageddon’s [18] usage of Evict+Reload on non-inclusive caches is much slower than our Evict+Reload attack as it must access many more addresses for the evict phase. Their method will be even slower for larger caches. For example, ARMageddon attacks L1 caches that are at most 32KB while their shared L2 cache is at most 2048KB. On the Skylake-X we attack, the L2 is 1MB and the LLC is 11MB. Despite having larger caches, our attack succeeds and with finer granularity than ARMageddon. Additionally, we overcome issues related to sliced caches that are not present on ARM architectures.

### B. Attacks on Inclusive Caches

Same-core side-channel attacks [27], [39], [40] leverage hyper-threading to co-locate victims and attackers on the same core, and exploit cache timing differences between L1 and L2 cache accesses. Others exploit operating system scheduler to achieve core-co-residency, overcoming the need for hyper-threading [19]. Cross-core attacks are more difficult as timing information comes from a much larger LLC which increases noise as it is shared across many cores. Yarum et al. [7] proposed a cross-core, cross-vm Flush+Reload attack on LLC by leveraging shared memory stemming from memory deduplication. Liu et al. [8] proposed a practical Prime+Probe attack on inclusive LLC, which does not rely on shared memory, as we discussed earlier.

## IX. COUNTERMEASURES

**Hardware-based.** Our attack causes conflicts on the limited number of directory entries (including ED) to create inclusion victims. One approach to prevent the attack is to eliminate contention for directory entries. This can be realized in a few different ways, some of which introduce severe performance degradation. First, we can increase the associativity of the extended directory in each LLC slice to equal the maximum possible entries in a set i.e.  $N_{L2} \times W_{L2}$ . In a sliced directory design, there would be  $N_{ED} \times S_{ED} \times N_{L2} \times W_{L2}$ . This results in a large amount of wasted area on the chip. Second, we can build a centralized directory structure to store the status of cache lines in all private caches. However, such centralized structure is not scalable and will be a serious performance bottleneck. Third, we can store the directory information along with the private L2 caches (e.g., a snoopy design). Such a design has scalability issues as core count increases. Indeed, it is folklore in the architecture community that that snoopy designs don’t scale beyond 20 cores.

Beyond general architectural changes, we can also prevent the attack by applying several side channel-prevention techniques that have been used for inclusive caches [34], [41], [42]. For example, the directory entry replacement policy can be modified to mimic SHARP [34], which prevents the creation of inclusion victims in the LLC. By preventing the replacement of directory entries occupied by a different core than the requesting core, the proposed attack could be prevented. Alternatively, one can partition the directory entries to every core in a manner similar to Intel CAT and prevent conflicts from other cores.

**Software-based.** Software only approaches to cache side-channel defenses suffer from a variety of drawbacks. Some of these use cache-coloring techniques [43]–[47] or constant time program transformation [48] which incur potentially large overheads, or costly application level changes. Compiler level defenses are transparent to developers, but incur large runtime overhead [49]. Other transparent techniques focus on kernel level changes but remain probabilistic [50], [51]. Nomad [52] is a probabilistic defense that operates at the cloud scheduler level to keep two tenants from being co-scheduled on the same host for long periods. It is challenging to mount a probabilistic defense against fine-grained attacks such as the one presented in Section VI.

## X. CONCLUSION

In this paper, we identify the directory as the unifying structure across different cache hierarchies on which to mount an LLC conflict-based side-channel attack. Based on this insight, we present two attacks. The first one is a Prime+Probe attack on non-inclusive LLCs. Our attack does not require the victim and adversary to share cores or virtual memory, and succeeds on state-of-the-art non-inclusive sliced caches such as those in Skylake-X [13]. The second attack is a novel, high-bandwidth Evict+Reload attack that uses multi-threaded adversaries to bypass non-inclusive cache replacement policies. We attack square-and-multiply RSA on the modern Intel Skylake-X processor, using both of our attacks. Moreover, to demonstrate all of the above, we conduct an extensive study to reverse engineer the directory structure in Intel Skylake-X processor. We develop a new eviction set construction methodology to find groups of cache lines that completely fill a given set of a given slice in a non-inclusive LLC.

## REFERENCES

- [1] Intel, “Intel Software Guard Extensions Programming Reference,” [software.intel.com/sites/default/files/329298-001.pdf](https://software.intel.com/sites/default/files/329298-001.pdf), 2013.
- [2] D. A. Osvik, A. Shamir, and E. Tromer, “Cache attacks and countermeasures: the case of aes,” in *Topics in Cryptology—CT-RSA 2006*. Springer, 2006, pp. 1–20.
- [3] O. Acıçmez, Ç. K. Koç, and J.-P. Seifert, “Predicting secret keys via branch prediction,” in *Cryptographers Track at the RSA Conference*. Springer, 2007, pp. 225–242.
- [4] Y. Wang, A. Ferraiuolo, and G. E. Suh, “Timing channel protection for a shared memory controller,” in *HPCA’14*.
- [5] M. Andryscio, D. Kohlbrenner, K. Mowery, R. Jhala, S. Lerner, and H. Shacham, “On subnormal floating point and abnormal timing,” in *Proceedings of the 2015 IEEE Symposium on Security and Privacy*, ser. SP ’15. Washington, DC, USA: IEEE Computer Society, 2015, pp. 623–639. [Online]. Available: <http://dx.doi.org/10.1109/SP.2015.44>
- [6] W. Wang, G. Chen, X. Pan, Y. Zhang, X. Wang, V. Bindshaedler, H. Tang, and C. A. Gunter, “Leaky cauldron on the dark land: Understanding memory side-channel hazards in SGX,” *CoRR*, vol. abs/1705.07289, 2017. [Online]. Available: <http://arxiv.org/abs/1705.07289>
- [7] Y. Yarom and K. Falkner, “Flush+reload: a high resolution, low noise, 13 cache side-channel attack,” in *23rd USENIX Security Symposium (USENIX Security)*, 2014, pp. 719–732.
- [8] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, “Last-level cache side-channel attacks are practical,” in *IEEE Symposium on Security and Privacy*, 2015, pp. 605–622.
- [9] D. Gruss, R. Spreitzer, and S. Mangard, “Cache template attacks: Automating attacks on inclusive last-level caches,” in *24th USENIX Security Symposium (USENIX Security)*, 2015, pp. 897–912.
- [10] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, “Cross-tenant side-channel attacks in paas clouds,” in *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2014, pp. 990–1003.
- [11] G. Irazoqui, T. Eisenbarth, and B. Sunar, “Cross processor cache attacks,” in *Proceedings of the 11th ACM on Asia Conference on Computer and Communications Security*. ACM, 2016, pp. 353–364.
- [12] “Transparent page sharing: new default setting.” <http://blogs.vmware.com/security/2014/10>, 2014.
- [13] Intel, “6th gen intel core x-series processor family datasheet - 7800x, 7820x, 7900x,” 2017, [online; accessed 30-Nov-2017]. [Online]. Available: <https://www.intel.com/content/www/us/en/products/processors/core/6th-gen-x-series-datasheet-vol-1.html>
- [14] R. Hund, C. Willems, and T. Holz, “Practical timing side channel attacks against kernel space aslr,” in *IEEE Symposium on Security and Privacy*. IEEE, 2013, pp. 191–205.
- [15] A. Jaleel, J. Nuzman, A. Moga, S. C. Steely, and J. Emer, “High performing cache hierarchies for server workloads: Relaxing inclusion to capture the latency benefits of exclusive caches,” in *2015 IEEE 21st International Symposium on High Performance Computer Architecture (HPCA)*, Feb 2015, pp. 343–353.
- [16] L. Zhao, R. Iyer, S. Makineni, D. Newell, and L. Cheng, “Ncid: a non-inclusive cache, inclusive directory architecture for flexible and efficient cache hierarchies,” in *Proceedings of the 7th ACM international conference on Computing frontiers*. ACM, 2010, pp. 121–130.
- [17] I. D. zone, “Intel xeon processor scalable family technical overview,” 2017, [online; accessed 29-Nov-2017]. [Online]. Available: <https://software.intel.com/en-us/articles/intel-xeon-processor-scalable-family-technical-overview>
- [18] M. Lipp, D. Gruss, R. Spreitzer, C. Maurice, and S. Mangard, “Armedddon: Cache attacks on mobile devices,” in *25th USENIX Security Symposium (USENIX Security 16)*. USENIX Association.
- [19] M. Neve and J.-P. Seifert, “Advances on access-driven cache attacks on aes,” in *Selected Areas in Cryptography*. Springer, 2006, pp. 147–162.
- [20] G. Irazoqui, T. Eisenbarth, and B. Sunar, “S\$A: A shared cache attack that works across cores and defies vm sandboxing—and its application to aes,” in *IEEE Symposium on Security and Privacy*. IEEE, 2015, pp. 591–604.
- [21] Y. Oren, V. P. Kemerlis, S. Sethumadhavan, and A. D. Keromytis, “The spy in the sandbox: Practical cache attacks in javascript and their implications,” in *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*. ACM, 2015, pp. 1406–1418.

- [22] D. M. Gordon, "A survey of fast exponentiation methods," *Journal of algorithms*, vol. 27, no. 1, pp. 129–146, 1998.
- [23] D. Gullasch, E. Bangerter, and S. Krenn, "Cache games—bringing access-based cache attacks on aes to practice," in *IEEE Symposium on Security and Privacy*. IEEE, 2011, pp. 490–505.
- [24] B. Gülmezoğlu, M. S. Inci, G. Irazoqui, T. Eisenbarth, and B. Sunar, "A faster and more realistic flush+ reload attack on aes," in *International Workshop on Constructive Side-Channel Analysis and Secure Design*. Springer, 2015, pp. 111–126.
- [25] J. Bonneau and I. Mironov, "Cache-collision timing attacks against aes," in *International Workshop on Cryptographic Hardware and Embedded Systems*. Springer, 2006, pp. 201–215.
- [26] C. Percival, "Cache Missing for Fun and Profit," <http://www.daemonology.net/papers/htt.pdf>, Oct. 2005, accessed on 22 April, 2017.
- [27] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "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, pp. 199–212.
- [28] Y. Xu, M. Bailey, F. Jahanian, K. Joshi, M. Hiltunen, and R. Schlichting, "An exploration of L2 cache covert channels in virtualized environments," in *Proceedings of the 3rd ACM workshop on Cloud computing security workshop*. ACM, 2011, pp. 29–40.
- [29] Z. Wu, Z. Xu, and H. Wang, "Whispers in the hyper-space: high-speed covert channel attacks in the cloud," in *USENIX Security symposium*, 2012, pp. 159–173.
- [30] C. Maurice, C. Neumann, O. Heen, and A. Francillon, "C5: cross-cores cache covert channel," in *Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2015, pp. 46–64.
- [31] T. Hornby, "Side-channel attacks on everyday applications: Distinguishing inputs with flush+reload," in *BackHat 2016*.
- [32] G. Irazoqui, M. S. Inci, T. Eisenbarth, and B. Sunar, "Know thy neighbor: Crypto library detection in cloud," *Proceedings on Privacy Enhancing Technologies*, vol. 2015, no. 1, pp. 25–40, 2015.
- [33] D. Gruss, C. Maurice, and K. Wagner, "Flush+ flush: A stealthier last-level cache attack," *arXiv preprint arXiv:1511.04594*, 2015.
- [34] M. Yan, B. Gopireddy, T. Shull, and J. Torrellas, "Secure hierarchy-aware cache replacement policy (sharp): Defending against cache-based side channel attacks," in *Proceedings of the 44th Annual International Symposium on Computer Architecture*. ACM, 2017, pp. 347–360.
- [35] C. D. N. Client., "Security: Disallow the x86 "clflush" instruction due to dram "rowhammer" problem," 2014 (accessed 2016-10-17). [Online]. Available: <https://bugs.chromium.org/p/nativeclient/issues/detail?id=3944>
- [36] Intel, "The intel 64 and ia-32 architectures software developer's manual," vol. 2A: Instruction Set Reference A-Z, no. 325383, 2016.
- [37] G. Paoloni, "How to benchmark code execution times on intel ia-32 and ia-64 instruction set architectures," 2010.
- [38] R. Singhal, "Inside intel next generation nehalem microarchitecture," in *Hot Chips*, vol. 20, 2008.
- [39] D. J. Bernstein, "Cache-timing attacks on aes," 2005.
- [40] E. Tromer, D. A. Osvik, and A. Shamir, "Efficient cache attacks on aes, and countermeasures," *Journal of Cryptology*, vol. 23, no. 1, pp. 37–71, 2010.
- [41] F. Liu, Q. Ge, Y. Yarom, F. Mckeen, C. Rozas, G. Heiser, and R. B. Lee, "Catalyst: Defeating last-level cache side channel attacks in cloud computing," in *22nd IEEE Symposium on High Performance Computer Architecture*, 2016.
- [42] Y. Wang, A. Ferraiuolo, D. Zhang, A. C. Myers, and G. E. Suh, "Secdcp: secure dynamic cache partitioning for efficient timing channel protection," in *Proceedings of the 53rd Annual Design Automation Conference*. ACM, 2016, p. 74.
- [43] B. Rodrigues, F. M. Quintão Pereira, and D. F. Aranha, "Sparse representation of implicit flows with applications to side-channel detection," in *Proceedings of the 25th International Conference on Compiler Construction*. ACM, 2016, pp. 110–120.
- [44] T. Kim, M. Peinado, and G. Mainar-Ruiz, "Stealthmem: system-level protection against cache-based side channel attacks in the cloud," in *USENIX Security symposium*, 2012, pp. 189–204.
- [45] J. Shi, X. Song, H. Chen, and B. Zang, "Limiting cache-based side-channel in multi-tenant cloud using dynamic page coloring," in *2011 IEEE/FIP 41st International Conference on Dependable Systems and Networks Workshops (DSN-W)*. IEEE, 2011, pp. 194–.
- [46] Y. Ye, R. West, Z. Cheng, and Y. Li, "Coloris: a dynamic cache partitioning system using page coloring," in *Proceedings of the 23rd international conference on Parallel architectures and compilation*. ACM, 2014, pp. 381–392.
- [47] H. Raj, R. Nathuji, A. Singh, and P. England, "Resource management for isolation enhanced cloud services," in *Proceedings of the 2009 ACM workshop on Cloud computing security*. ACM, 2009, pp. 77–84.
- [48] A. Rane, C. Lin, and M. Tiwari, "Raccoon: Closing digital side-channels through obfuscated execution," in *24th USENIX Security Symposium (USENIX Security 15)*. Washington, D.C.: USENIX Association, 2015, pp. 431–446. [Online]. Available: <https://www.usenix.org/conference/usenixsecurity15/technical-sessions/presentation/rane>
- [49] B. Coppens, I. Verbauwhede, K. De Bosschere, and B. De Sutter, "Practical mitigations for timing-based side-channel attacks on modern x86 processors," in *Security and Privacy, 2009 30th IEEE Symposium on*. IEEE, 2009, pp. 45–60.
- [50] Z. Zhou, M. K. Reiter, and Y. Zhang, "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, pp. 871–882.
- [51] V. Varadarajan, T. Ristenpart, and M. Swift, "Scheduler-based defenses against cross-vm side-channels," in *23rd USENIX Security Symposium (USENIX Security 14)*, 2014, pp. 687–702.
- [52] S.-J. Moon, Y. Sekar, and M. K. Reiter, "Nomad: Mitigating arbitrary cloud side channels via provider-assisted migration," in *Proceedings of the 22nd acm sigsac conference on computer and communications security*. ACM, 2015, pp. 1595–1606.
- [53] C. Maurice, N. Le Scouarnec, C. Neumann, O. Heen, and A. Francillon, "Reverse engineering intel last-level cache complex addressing using performance counters," in *Research in Attacks, Intrusions, and Defenses*. Springer, 2015, pp. 48–65.
- [54] G. Irazoqui, T. Eisenbarth, and B. Sunar, "Systematic reverse engineering of cache slice selection in intel processors," in *Euromicro Conference on Digital System Design (DSD)*. IEEE, 2015, pp. 629–636.
- [55] W. V. Quine, "The problem of simplifying truth functions," *The American Mathematical Monthly*, vol. 59, no. 8, pp. 521–531, 1952. [Online]. Available: <http://www.jstor.org/stable/2308219>
- [56] D. Molka, D. Hackenberg, and R. Schöne, "Main memory and cache performance of intel sandy bridge and amd bulldozer," in *Proceedings of the Workshop on Memory Systems Performance and Correctness*, ser. MSPC '14. New York, NY, USA: ACM, 2014, pp. 4:1–4:10. [Online]. Available: <http://doi.acm.org/10.1145/2618128.2618129>
- [57] M. Clark, "A new x86 core architecture for the next generation of computing," in *Hot Chips 28 Symposium (HCS), 2016 IEEE*. IEEE, 2016, pp. 1–19.
- [58] T. Singh, S. Rangarajan, D. John, C. Henrion, S. Southard, H. McIntyre, A. Novak, S. Kosonocky, R. Jotwani, A. Schaefer et al., "3.2 zen: A next-generation high-performance x86 core," in *Solid-State Circuits Conference (ISSCC), 2017 IEEE International*. IEEE, 2017, pp. 52–53.

## APPENDIX

### A. The Eviction Set Construction Algorithm

The complete EV construction algorithm we used for non-inclusive cache is a slightly modified version of the algorithm proposed by Liu et al. [8], as shown in Algorithm 4.

The `find_EV(Collection CS)` function above takes a collection of addresses, which we call a Candidate Set ( $CS$ ) as input, and outputs one eviction set ( $EV$ ) for one slice as output. For the algorithm to work, it is required that all the addresses in  $CS$  have the same LLC set index bits, and  $CS$  contains more than  $W_{\text{slice}}$  addresses for each slice. Such  $CS$  can be easily obtained by using a large number of addresses. To find eviction sets for all the slices within  $CS$ , we need to run the function the same number of times as the number of slices.

The function initializes  $EV$  as an empty set and selects a random address `test_addr` in the Candidate Set (line 2-3). It then tries to construct an eviction set containing all

---

**Algorithm 4:** Constructing an eviction set.

---

```
Input : candidate set  $CS$ ,  $LLC\_miss\_threshold$ 
Output :  $EV$ 
1 Function  $find\_EV(CS)$ :
2    $EV = \{\}$ 
3    $test\_addr =$  get a random addr from  $CS$ 
4    $CS' = CS - test\_addr$ 
5   // make sure there are enough addresses to conflict with  $test\_addr$ 
6   if  $check\_conflict(test\_addr, CS') \neq false$  then
7     return fail
8   end
9   for each addr in  $CS'$  do
10    if  $check\_conflict(test\_addr, CS' - addr) \neq false$  then
11      // if conflict disappears, we know  $addr$  contributes to
12      // the conflict, and it should be in the  $EV$ 
13      insert  $addr$  to  $EV$ 
14    else
15       $CS' = CS' - addr$ 
16    end
17  end
18  for each addr in  $CS$  do
19    if  $check\_conflict(test\_addr, EV) = true$  then
20      insert  $addr$  to  $EV$ 
21    end
22  end
23  return  $EV$ 
24 end
```

---

the addresses which are mapped to the same slice and set as  $test\_addr$  from  $CS$ . First, it creates a new set  $CS'$  by removing  $test\_addr$  from  $CS$  (line 4), and then performs a sanity check to make sure  $CS'$  contains enough addresses to evict  $test\_addr$  out of LLC using  $check\_conflict$  (line 5-7).

The loop (line 8-14) performs the bulk of the work, checking whether an address is mapped to the same slice as  $test\_addr$ . Since  $CS'$  conflicts with  $test\_addr$ , if removing an address  $addr$  causes the conflict disappear, we know that  $addr$  contributes to the conflict, and  $addr$  should be added to the  $EV$  (line 9-10). Such addresses are kept in  $CS'$ . Addresses which are not strictly necessary to cause conflicts with  $test\_addr$  are removed from  $CS'$  (line 12) and  $CS'$  should still conflict with  $test\_addr$  after the remove operations. After the loop, we can obtain a minimal eviction set with exactly  $W_{slice}$  number of addresses in  $EV$ .

It is possible that there are more than  $W_{slice}$  addresses from the same slice as  $test\_addr$  which have been conservatively removed in the loop. We use an extra loop (line 15-19) to find these addresses by iteratively checking each address in the original Candidate Set ( $CS$ ) to determine whether it conflicts with the obtained  $EV$ .

The  $check\_conflict$  function is extensively used in this algorithm. On line 9, the function is used to test whether removing an address from  $U$  can cause LLC conflicts disappear. It is the operation that requires the function have both low false positive rate and false negative rate, as discussed in Section IV.

### B. Slice Hash Function

Based on our  $EV$  construction results, we are able to reverse engineer part of the slice hash function in the Intel Skylake-X processor. Our goal here is to show that the slice hash function

is not a simple XOR operation of selected physical address bits. This behavior is significantly different from the one exhibited by previous Intel processors such as SandyBridge and IvyBridge. Considering all of the previous works on reverse-engineering slice hash functions [53], [54] rely on the usage of a simple XOR hash function, our results identify the need for more advanced reverse-engineering approaches.

We now briefly discuss how to get the partial hash function. We select 128 addresses with the same LLC set index bits to form a Candidate Set ( $CS$ ). The 6-16 bits of these addresses are set to the same value, while the bits 17-23 are varied. The goal is to reverse engineer how the bits 17-23 affect the output of the slice hash function.

First, we run  $find\_EV$  on the 128-address  $CS$  and obtain 8  $EV$ s. Each  $EV$  is mapped to one cache slice. Second, we try to figure out the slice id for each  $EV$ . Since the Skylake-X processor uses a mesh network-on-chip to connect L2s and LLC slices [13], a local LLC slice access takes shorter time than a remote slice access. We check the access latency of each  $EV$  from each core. We then get the id of the core from which the access latency is the lowest, and assign the core id to the  $EV$ . Finally, we use the Quine-McCluskey solver [55] to get the simplified boolean functions from the input bits to the slice id as below.  $o_i$  is the  $i$ th bit in the slice id, and  $b_i$  is the  $i$ th bit in the physical address.

$$\begin{aligned} o_2 &= b'_{23}b'_{19} + b'_{22}b'_{19} + b_{23}b_{22}b_{19} \\ o_1 &= (b_{23} + b_{22})(b_{20} \oplus b_{19} \oplus b_{18} \oplus b_{17}) + \\ &\quad b'_{23}b'_{22}(b_{20} \oplus b_{19} \oplus b_{18} \oplus b_{17})' \\ o_0 &= b'_{22}(b_{19} \oplus b_{18}) + b_{22}(b_{23} \oplus b_{21} \oplus b_{19} \oplus b_{18})' \\ &\quad \text{where } \{b_{63} \dots b_{24}\} = 0x810 \end{aligned}$$

These functions can not be further reduced to a simple XOR function. According to our observations, some of the higher bits (bits 24-63) also affect the hash function, which we have not fully reverse engineered.

### C. Attacking AMD Non-Inclusive Caches

	AMD Piledriver	AMD Zen (4-core CCX)
L1-I	64KB/2cores, 2-way	64KB, 4-way
L1-D	16KB, 4-way	32KB, 8-way
L2	2MB/2cores, 16-way, inclusive	512KB, 8-way, inclusive
LLC	8MB/8cores 64-way, non-inclusive	2MB/core 16-way, non-inclusive

TABLE II  
CACHE STRUCTURES IN AMD PROCESSORS

We tried to reverse engineer the non-inclusive cache in an 8-core AMD FX-8320 processor, which uses the AMD Piledriver microarchitecture. The cache parameters in this processor are listed in Table II.

We found the L2 caches in this processor are inclusive and shared by two cores. We verified that previous inclusive cache attacks work well, if the attacker and the victim are located on neighboring cores and share the same L2.

To see whether the non-inclusive LLC is vulnerable to cache attacks or not, we tried all the reverse engineering experiments



in Section V to detect the existence of directories. We did not observe extra conflicts besides cache conflicts. It is possible that the processor uses a snoop-based protocol [56], in which case there is no directory. It is also possible that the processor uses a unified and high-associativity directory design, and the unified directory associativity is at least as high as the total cache associativity. Thus the conditions in Section V-F do not hold. Note that in the latter case, the directory for L2 lines needs to have  $\geq 64$  ways.

We also tried to evaluate our attacks on an 8-core Ryzen 1700 processor, which uses the latest AMD Zen microarchitecture. The cache parameters are also listed in Table II. The processor consists of 2 Core Complexes (CCX). A CCX is a module containing 4 cores, which can connect to other CCX modules via Infinity Fabric [57], [58]. We did not observe extra conflicts other than the cache conflicts on this processor either. Given the small number of cores on each die and low L2 associativity (8 on AMD CCX compared to 16 on Intel Skylake-X), we hypothesize this processor either uses snoop-based protocol or a 32-way global directory for L2 lines.

**Performance implications for AMD designs.** We measured the remote L2 access time for the Piledriver processor and found it took a similar amount of time as a DRAM hit, which is significantly slower than the corresponding operation in the Intel design (Figure 5). This backs up our claim that sliced directories are important structures in high performance processors. For the Ryzen processor, we have a similar result: a cross-CCX access takes a similar amount of time as a DRAM access. Skylake-X/Skylake-SP can support up to 28 cores. Since each CCX is only 4 cores, constructing a similarly provisioned Ryzen system can mean most cross-core accesses turn into cross-CCX accesses.