

TPPD: Targeted Pseudo Partitioning based Defence for Cross-Core Covert Channel Attacks

Jaspinder Kaur, and Shirshendu Das *IEEE Member*

Abstract—Contemporary computing employs cache hierarchy to fill the speed gap between processors and main memories. In order to optimise system performance, Last Level Caches(LLC) are shared among all the cores. Cache sharing has made them an attractive surface for cross-core timing channel attacks. In these attacks, an attacker running on another core can exploit the access timing of the victim process to infiltrate the secret information. One such attack is called cross-core Covert Channel Attack (CCA). Timely detection and then prevention of cross-core CCA is critical for maintaining the integrity and security of users, especially in a shared computing environment. In this work, we have proposed an efficient cross-core CCA mitigation technique. We propose a way-wise cache partitioning on targeted sets, only for the processes suspected to be attackers. In this way, the performance impact on the entire LLC is minimised, and benign applications can utilise the LLC to its full capacity. We have used a cycle-accurate simulator (gem5) to analyse the performance of the proposed method and its security effectiveness. It has been successful in abolishing the cross-core covert timing channel attack with no significant performance impact on benign applications. It causes 23% less cache misses in comparison to existing partitioning based solutions while requiring $\approx 0.26\%$ storage overhead.

Index Terms—Cache Security, Timing Channel Attacks, Cache Partitioning, Covert Channel Attack (CCA), Last Level Cache (LLC).

I. INTRODUCTION

Modern multi-core processors have adopted a multi-level cache hierarchy to address the rising needs for high-performance computing. Caches are small-sized faster memories deployed to bridge the speed gap between processor and main memory. Most commercial processors available in the market are equipped with multiple levels of caches. The top levels of caches are private to each core, and the Last Level Cache (LLC) is shared among all the processing cores¹. Cache memories boost overall performance by allowing processors quicker access to data. While improving the system's overall performance, the LLCs have also become an attractive target surface for timing channel-based attacks [1] due to these reasons:

- 1) The significant time difference between a cache hit and miss provides an effective timing channel that is exploited to unveil the memory accesses of the targeted process [2], [3].

The authors are associated with the Department of CSE, Indian Institute of Technology Ropar, Punjab, India, 140001. Email: {2017csz0002@iitrpr.ac.in, shirshendu@iitrpr.ac.in}

¹All the cache memories in this paper are considered as set-associative cache. We used the term “set” and “way” of a set-associative cache without detailed explanation about them.

- 2) Shared nature of LLC allows an attacker process to interfere in cache occupancy of other processes. The attacker uses the timing channel to understand other processes' access patterns on the shared cache as described in Figure 1. These accesses are further analysed to reveal the underlying secret [4].

Any remote process, sharing LLC with the victim process, can mount these attacks without requiring special privileges or shared address space [5]. The victim and attacker process may run in separate cores, but they share the LLC space. Such attacks are called *cross-core attack*. These attacks become more threatening in shared computing environments such as the cloud, where multiple users from different security domains share the underlying LLC. Cache timing channel can be implemented in two forms: Side-Channel Attack(SCA) and Covert Channel Attack(CCA) [6], [7]. The private keys of cryptography algorithms like AES [8], RSA [9], and ECDSA [10] have been the primary target of cache-based SCAs. In a CCA, two suspicious processes, spy and trojan, communicate covertly by exploiting the cache timing channel [6]. The trojan runs on a different core and knows some secret information that it needs to send to the spy. The attacker uses CCA when no direct communication between trojan and spy is possible because of the security restrictions. It is crucial to detect and prevent these attacks effectively in order to preserve the confidentiality and integrity of the computer systems [11], [12]. This paper is based on proposing countermeasures for preventing cross-core CCA attacks. Section II-A discusses CCA in more detail.

The CCA attack that is discussed in this paper is cross-core attacks; based on the shared LLC used by the multi-core processors. Attack defence mechanisms based on not allowing cache sharing through LLC partitioning have been proposed previously [13], [14]. Cache partitioning techniques were initially proposed for improving the performance of the system by fairly dividing the LLC among the cores (or applications) [15]–[17]. Most of them were not proposed keeping security in mind. These partitioning techniques divide the LLC either way-wise or set-wise. The way-wise partitioning techniques are more prevalent where the ways of a set-associative LLC is partitioned among the cores (or applications) [15]. The partition can be either static or dynamic. The static partitioning techniques cannot change the partition during the execution, while the dynamic partitioning techniques can change the partition based on the requirement of the running applications. Since cache partitioning can separate the LLC space used by trojan and spy, the cache interference of two applications can

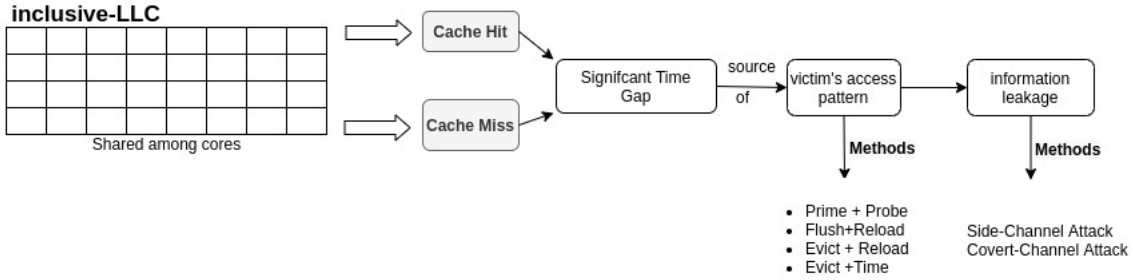


Fig. 1: The step-wise details of LLC based timing channel attacks

be prevented. However, there are two major challenges in using cache partitioning for preventing CCA:

- 1) The static partition-based solutions suffer from severe performance degradation as cache cannot be utilised fairly based on the dynamic behaviour of the system [15].
- 2) Dynamic partitioning techniques utilise the cache more efficiently and hence improve performance. However, the dynamic partitioning technique itself can be exploited to mount CCA as proposed in [18]. In this work, the authors have shown that if the attacker knows the logic of partition change, it can misuse this and change the partition as per the attacker wish. Thus, a dynamic partitioning based countermeasure can restrict the attack but indirectly opens another option for CCA.

Hence for preventing CCA, a partitioning technique is required, in which the attacker cannot change partition size. Also, the technique should not degrade the performance of the overall system.

There are some existing countermeasures proposed based on the static partition where the technique applies to all the sets within the LLC. In this paper, we call this type of countermeasures as non-targeted countermeasures. Such countermeasures cause higher performance degradation in the system. Our attack mitigation mechanism overcomes the drawbacks of the existing partitioning-based solutions. We do this by providing Targeted Pseudo Partitioning based Defence (TPPD). It is termed as “targeted” because the attack prevention mechanism is only applied to the cache sets suspected to be participating in CCA. The sets which participate in the timing channel attacks are termed as *targeted set*. The remaining cache sets can behave as before, thus minimising the effect of static partitioning on the performance. A practical attack detection technique is deployed that send the information regarding suspicious processes (trojan and spy) and cache sets involved. Attack detector based on conflict misses pattern [19] has been tested to identify the suspicious processes and sets involved.

The main contributions of this work are listed as following:

- 1) We propose an effective attack defense mechanism, TPPD, that dismantles the cross-core CCA with an insignificant effect on system performance.
- 2) The proposed TPPD creates the pseudo partition only on the targeted set and is applied only for the trojan and spy process.
- 3) Effect of TPPD on the performance of PARSEC benchmark [20] is tested experimentally in gem5 simulator

[21].

- 4) Experiments are conducted to test the effectiveness of TPPD in mitigating cross-core CCA.

The organization of the paper is as follows. Section II describes the background and related work. Section III describes the threat model taken under consideration, and Section IV discusses our proposed work in details. The experimental analysis are shown in Section V. Finally Section VI concludes the paper.

II. BACKGROUND AND RELATED WORK

This section describes the background information required to understand the proposed attack defence mechanism. In cache timing channel attack, the attacker process uses different attack methods to unveil the cache access pattern of the target process. In side-channel attacks, the target is an innocent victim process performing cache accesses as part of its underlying operation without any malign intentions. In a covert channel attack, there are two malign processes performing cache accesses with the intent of leaking secrets. An abstract overview of cache timing channel attacks is shown in Figure 1.

A. Cross-Core Covert Channel Attacks

In a covert timing channel attack, two suspicious processes: spy and trojan, are involved. Trojan has access to some critical information that it wants to transmit to spy, but this transmission is not allowed under system security policy. Trojan uses cache timing channel to leak this information to spy without being noticed. This type of attack is called Covert Channel Attack (CCA). The CCA is called cross-core when the spy and trojan execute on two different cores. In this case, both trojan and spy use the shared LLC to create a timing channel. Cache timing channel attack can be constructed using various attack techniques like Prime+Probe (P+P), Evict+Reload (E+R), Flush+Reload (F+R), and Evict+Time (E+T) [6]. These attack techniques, even though different, follow these basic three steps:

- Step 1: In the first step, the aim is to bring the shared LLC in a predictable state known to the attackers. It is done by either bringing some data in LLC or evicting it out by the spy.
- Step 2: In this step, the spy process sits idle, and the trojan process performs conditional memory accesses based on the bit to transmit (bit 1 or 0) .

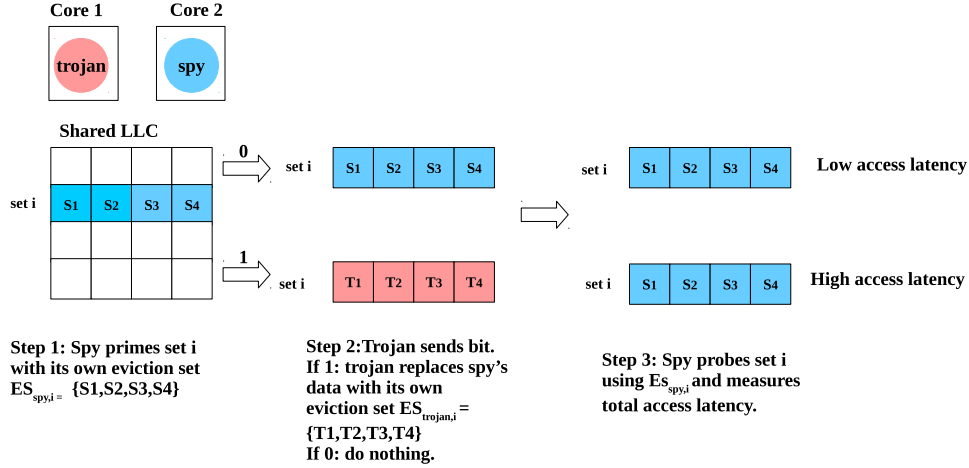


Fig. 2: Cross-Core covert channel attack using Prime+Probe technique

- Step 3: In this step, the spy process observes the changes made in the earlier known state of the cache by the trojan to detect the bit as 1 or 0.

Repeating the above mentioned steps, the trojan can send multiple bits to the spy.

CCA using Prime+Probe (P+P) technique is described in Figure 2. In this method, spy and trojan do not require any shared address space. Both these processes have their own Eviction Set (ES) [22]. ES is the group of unique block addresses mapping to the same set. The number of addresses in each ES is either equal to or more than the associativity of the underlying cache. In the prime phase, the spy uses the addresses present in its ES to fill the targeted set with its own blocks. After priming the cache, spy waits for a trojan to send the bit. If trojan wants to transmit bit 1, it replaces all of spy's block from the targeted set with its own blocks (from trojan's ES). For bit 0, it does not do anything; thus, spy's blocks stay in the set. After this, in the probe phase, spy accesses the addresses from its ES again. In this phase, spy faces high latency in case trojan has removed its blocks from the set; otherwise, less latency. Based on this latency, the spy unveils the cache access pattern of trojan and the secret bit transmitted. The rest of the discussions of this paper assume Prime+Probe (P+P) based CCA.

B. Existing Attack Mitigation Techniques

As discussed in Section II-A, cache sharing is the critical pre-requirement of cross core cache timing channel attacks. Prohibiting cache sharing across processes or different security domains is an effective solution. However, it is not feasible as the impact of a non-shared cache will lead to under-utilisation of cache capacity thus, impacting system performance significantly. Various works with different cache partitioning techniques that try to maintain cache utilisation have been proposed in recent years. Partitioning Locked cache (PLcache) [23] allows compiler and programmer to mark security critical lines, and these will be locked in the cache. In this architecture, each cache line is augmented with process id and a bit dictating whether that line is locked or not.

A locked line of a process cannot evict a locked line of a different process, while no unlocked line is allowed to replace a locked line. Thus, disabling inter-process and intra-process cache interference that has been identified as the root cause of cache access-based attacks. NonMonopolizable (NoMo) [24] cache allocated v number of ways of each set to an active thread. The allowed range for v is 1 to A/M , where A is the associativity of the cache, and M is the maximum number of threads allowed per processing core. Data of one thread cannot evict reserved lines of another thread; thus, an attacker can not know the victim's cache access.

NoMo and PLCache offer easy to implement static partition solutions, but the static partitioning can lead to not utilising cache to its full capacity. Security Dynamic Cache Partitioning (SecDCP) cache [25] addresses this issue by proposing a dynamic cache partition. SecDCP allocates a number of ways to different security classes based on their requirement. However, dynamic partitioning in itself can be exploited to create cache covert timing channel attack as described in [18]. A secure dynamic partitioning in terms of SCA is proposed in [13]. It proposes a secure dynamic cache partitioning cache called FairSDP, where partitioning size is determined by the cache usage of non-critical processes, thus not revealing the requirement of security-critical processes. However, this does not work against CCA as it involves two suspicious processes, and there is no security-critical process.

All the mitigation techniques discussed above apply some strict cache partitioning policy across all sets for all processes. Hence these techniques reduce the performance of the LLC as well as the entire system. Also, most of the existing countermeasures are non-targeted, as the mechanism is applied to all the cache sets and processes regardless of their participation in the attack. A targeted countermeasure can be less expensive as the necessary actions can be applied only in the sets currently under attack. An alternative solution to prevent cross-core CCA is the recently proposed randomised LLCs [26], [27]. In these works, the block-to-set mapping of set-associative cache changes after an epoch. However, these techniques suffer significant performance degradation because of the remapping

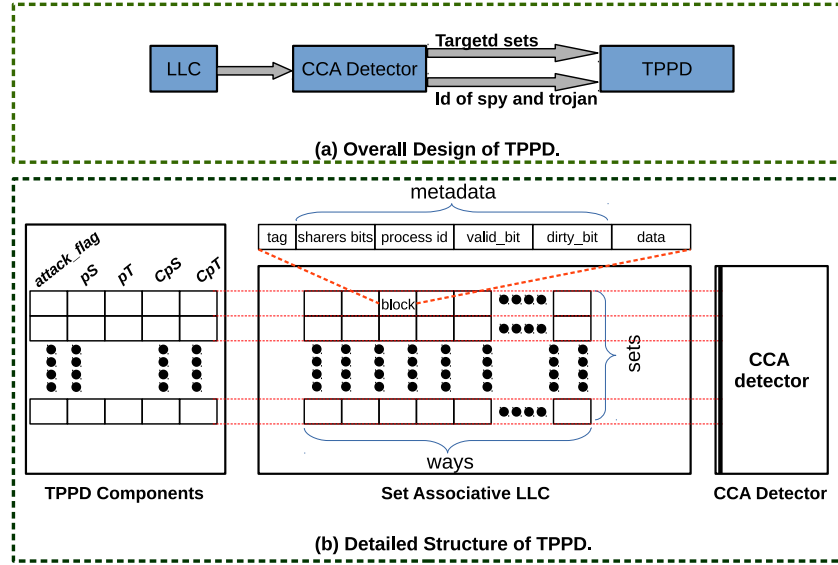


Fig. 3: Overall design of TPPD and its structure

[28]. The number of write-backs required to remap is high. Also, the policies may not be useful for recent replacement policies like Hawkeye [29]. Hence in this work, we have been motivated to propose a targeted countermeasure for the cross-core CCA.

III. THREAT MODEL

In this paper, we have assumed a chip multiprocessor (CMP) with two levels of cache memories. Each core in the CMP has its own private L1 cache, and all the cores share a common L2 cache as LLC. The LLC is inclusive, and all the cache memories are set-associative. The cores, L1 caches, and LLC are connected with some on-chip interconnects. The cross-core CCA attack discussed in this paper runs spy and trojan on two different cores. These two attacking processes (spy and trojan) are not using any shared address space. Hence no cache blocks can be shared among these two processes, and the eviction sets (c.f. Section II-A) of both the processes are also different. However, both spy and trojan share the LLC, and they can evict each other blocks from the LLC. The eviction set of both spy and trojan maps to the same targeted set. We propose TPPD to mitigate cross-core CCA. This attack can be based on Prime+Probe [6], Evict+Reload [6], or Evict+Time [6] methods, where spy and trojan processes rely on replacing each other's block to transmit bits. Trojan has a piece of secret information that spy cannot access directly due to underlying system security policies. Trojan transmits this information to spy through LLC based covert communication channel (details of CCA is already discussed in Section II-A). Both trojan and spy can successfully create their eviction sets as they are aware of the virtual address to LLC set mapping. Both attacker processes are unprivileged processes that do not have any kind of special privilege.

IV. OUR PROPOSAL

Our proposed targeted defense mechanism, TPPD, works against LLC based covert channel attacks (cross-core CCA) by

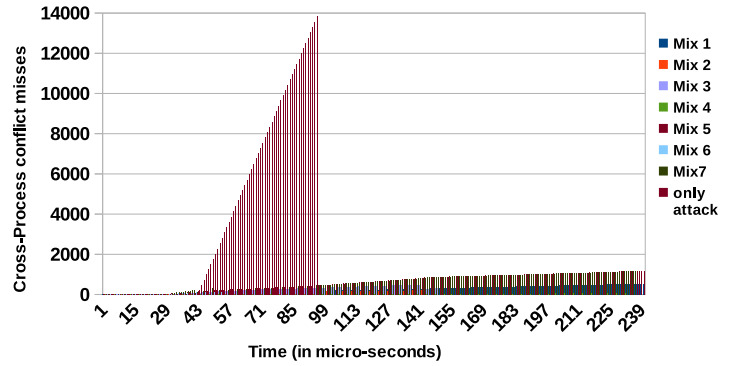


Fig. 4: Cross-process conflict misses per unit of time(.10 seconds).

creating way-wise pseudo partitioning between the suspicious process pair (spy and trojan). This partition technique used in TPPD cannot be changed by the attacker process as per its requirement. Also, the partition is only applied on the sets suspected to participate in the covert channel communication. The purpose of using the term “pseudo” is discussed later in this section. TPPD decreases the cache interference between the trojan and spy on the targeted sets, thus disturbing the access latency pattern observed by the spy to interpret the bit sent. TPPD receives information regarding suspicious process ids (spy and trojan) and the targeted sets from a CCA detector module; then applies the proposed pseudo-partitioning on these sets for the suspicious processes. The CCA detection module is discussed in Section IV-A. The overall design of TPPD is shown in Figure 3(a). In this section, we describe TPPD in detail, along with how it can be implemented with minimal modification in architectural design.

A. Cross-Core Covert Channel Attack Detector

TPPD works efficiently by applying attack mitigation mechanisms only for suspicious process pairs on the LLC sets suspected to be used for cache covert channel attacks. The countermeasure proposed in TPPD depends on an efficient CCA detection mechanism. In this section, we describe how existing attack detection techniques can be modified or extended to identify suspicious processes and cache sets.

1) *Conflict Misses Pattern based Detector*: A cross-core CCA detection mechanism based on conflict misses pattern is proposed in [19]. It is based on the observation that set under attack have a higher number of conflict misses than innocent processes, and these misses form a ping pong pattern because spy and trojan processes evict each other’s data regularly to communicate covertly using cache timing channel attacks. This work proposed a two-step detection mechanism in which, firstly, sets with low conflict misses are filtered out. In the second step, sets observing ping pong pattern of conflict misses between two processes are declared suspicious along with processes participating in this pattern. However, we observed that observing total cross-process misses per unit of time is sufficient in order to identify sets and processes used for mounting CCA. We have also observed the same experimentally as shown in Figure 4. The figure shows the cross-process conflict misses on an LLC set. The experiment is performed on different mixes of benign processes (Mix- x) from Parsec benchmarks (described in Table I) and also on an attacker process pair (only-attack). The details about the experiment setup, benchmarks and creating attack process pair are discussed in Section V. From the figure, it can be observed that the attacker process pair suffers high cross-process conflict misses when compared to any mix of benign benchmarks for the duration of the attack.

2) *Cache Occupancy based Detector*: Another way to detect the cache sets and the processes involved in covert channel communication is through tracking cache occupancy, process wise [30]. This detection method is based on the observation that cache occupancy of the suspicious process have a high negative correlation as these two keep evicting each other’s block frequently. This method requires collecting cache occupancy traces and then further analysing these traces to determine occupancy correlation between processes. It was observed that benign processes do not have such a high negative correlation. This method did not require any additional storage overhead and was able to detect suspicious processes with high accuracy. However, this method was able to disclose the identity of suspicious processes with high accuracy but not the cache sets on which the attack was mounted as cache occupancy was measured for the whole LLC. If cache occupancy trace is collected for each set individually, then suspicious sets can be observed as well.

3) *The detection mechanism used for TPPD*: In this work, we have used the concept of conflict-miss pattern-based CCA detector to detect the attack. However, the cache occupancy-based detector can also be used to detect the attack. Since our work is on detection based countermeasure, we have used

TABLE I: PARSEC benchmark mixes used.

Mixes	Benchmarks Involved
Mix 1	blacksholes + canneal
Mix 2	blacksholes + dedup
Mix 3	blacksholes + fluidanimate
Mix 4	blacksholes + freqmine
Mix 5	canneal + dedup
Mix 6	canneal + fluidanimate
Mix 7	freqmine + fluidanimate

the existing detection techniques. Figure 3(a) shows that the outcome of the CCA detector is a pair of attacking processes (trojan and spy) and the targeted sets.

B. Targeted Pseudo Partitioning based Defense (TPPD)

When the CCA detector detects an attack, it sends information about the processes and sets involved in the CCA. TPPD enables pseudo partitioning on the suspected sets for the suspected process pair (trojan and spy). Two non-overlapping way-wise partitions are created and assigned to trojan and spy. The spy and trojan are unable to replace each other’s data if the block is present in this allocated partition. This restriction, however, exists solely for identified malicious pair; other benign processes can access this cache set as before without any constraint. Because of this reason, the proposed partitioning is called pseudo as it is invisible to all processes except attackers. As our defense mechanism is targeted, it only affects the cache sets and the processes involved in CCA, thus not impacting the system’s performance significantly.

The proposed TPPD can handle multiple CCA attacks on different cache sets simultaneously. However, in this section, we have assumed only one pair of attacker processes (spy and trojan) and only one targeted set. For the rest of this section, we assume that the detector detects pS as a spy and pT as a trojan. Here pS and pT are the process id of spy and trojan, respectively. Also, we assume that the detected targeted set is s' . TPPD can be implemented on top of any replacement policy with minor modifications. A replacement policy has three major modules: (a) insertion, (b) promotion and (c) eviction [16]. To implement TPPD, minimal changes are required in the eviction module of a replacement policy. TPPD expects two victim selection approaches on the targeted set s' . We call it as “dual victim policy”.

- a) The default victim selection of a replacement policy. Let us represent the victim block selected through this method as $V(s')$, where s' is the targeted set.
- b) Victim block selection excluding a particular process p . Let us represent the victim block selected through this method as $V_x(s', p)$. Here the victim block is selected from the targeted set s' , excluding the blocks owned by process p .

This dual victim policy can be implemented on most of the existing replacement policies [29], [31] with minor modifications. The concept of TPPD, generic to any replacement policy (those can be modified for dual victim selection), is discussed next.

Consider that for a cache block b , $O(b)$ represents the owner process to which the block belongs. The dual eviction policy is only required for the targeted set s' , detected by the CCA detector as discussed in Section IV-A. For other sets, $s; s \neq s'$, the victim block is always $V(s)$. Since the attacker can target any of the cache sets, the facility of dual victim selection must be present in all the sets. However, it is used only for the targeted set. To implement the dual eviction policy, we need to maintain two counters for each set $CpS(i)$ and $CpT(i)$ where i is the set number. $CpS(s')$ and $CpT(s')$ are used to count the number of spy and trojan blocks respectively present in s' . When pT (trojan) block needs to replace an existing block of pS (spy) from s' , TPPD sets some additional conditions. If $O(V(s'))$ is pS (spy) and $CpS(s')$ is less than a threshold, pT cannot replace $V(s')$. In that case, the victim block is selected by $V_x(s', pS)$. This policy restricts the trojan to remove all of the spy blocks from s' . A similar policy is also applied when a spy tries to evict a block of trojan from s' . The ping-ponging pattern of conflict misses between the spy and trojan is interrupted because continuous eviction of each other's data is restricted. In this way, a spy cannot observe the cache access latency pattern to distinguish bit 0 and 1, resulting in noisy covert channel communication. The value of the threshold depends on the underlying replacement policy. In the rest of this paper, we use LRU to implement TPPD, and an appropriate threshold value to inhibit attack is also discussed.

1) *Structure of LLC using TPPD*: Figure 3(b) shows the structure of the LLC using TPPD. The additional components required for TPPD are shown in the two sides of the set-associative LLC. These additional components are divided into two major parts: the CCA Detector and TPPD Components. For TPPD, each set maintains a tuple, $(attack_flag, pS, pT, CpS, CpT)$. Here the attribute, $attack_flag$ is a single bit of information that indicates whether the set is under attack or not. The other attributes are already defined in this section. The attributes excluding $attack_flag$, are only required for a suspicious set. When the CCA detector detects a CCA attack, the corresponding attributes of the attacker set are updated in the TPPD component. As mentioned in Section IV-A the CCA detection module is a cross-process conflict-miss pattern-based detector as proposed in [19].

2) *Engagement and disengagement with dual victim policy*: Since the CCA attack happens during execution, in the beginning, all the sets are non-suspicious. A set engages with dual victim policy only when it has been detected as targeted set by the CCA detector. Once a process pair and a set s' has been detected as suspicious, one option is to continue with the proposed dual victim policy (in s') till the termination of these processes. However, the option may reduce the performance of these processes in false-positive cases. The second option is to periodically check the set status in the CCA detector. If any time the cross-process misses between spy and trojan reduces in s' , the set may again reset as non-suspicious. Though the chances of such false-positive cases in the CCA detector are very few, we have only used the first option in this paper.

3) *Maintaining Process ID*: The replacement policy of TPPD needs the process id of a block stored in the LLC. Figure 3(b) shows it as a metadata of each LLC block. The existing partitioning techniques [15], [16] as well as some replacement policies [29], [32], [33] need process id for each block. Hence the overhead of maintaining process id in LLC as metadata cannot be considered as the overhead of TPPD. In case we assume only one process can run in a core, then instead of process id, we can also use the core id. Storing core id takes fewer bits than storing process id. This policy is used in some existing works [15], [16]. However, in a practical design, each block must store the process id. The additional components required for TPPD (as shown in Figure 3(b)) also need to store process id. We have considered all the additional components as an overhead of TPPD. In Section V-D, we have calculated the hardware overhead of TPPD, both assuming core id and process id as TPPD components.

Terminology used in Algorithms:

- *EvictionVictim(i, p)*: Finds the the victim block to be replaced from set i .
 - i : Set number of the victim block.
 - p : Owner process of the incoming block. Or the proceess responsible for triggering the cache replacement.
- *isUnderAttack(i)*: Returns true if the set i is under CCA attack, otherwise false. The value is provided by status bit maintained in additional TPPD structure.
- *getTrojan(i), getSpy(i)*: Returns the suspicious process pair(spy and trojan) on set i . The values are provided by the two suspicious process identifier maintained in TPPD structure.
- *getLRUVictim(i)*: Returns $V(i)$ from the set i . Or returning the victim block as per LRU replacement policy.
- *checkOwner(i, w, p)*: Returns true if process p is the owner of the block w in set i , otherwise false.
- *Owner(w)*: Returns owner process of block w . Also represented as $O(w)$.
- *$pS[i]$ and $pT[i]$* : Process identifier of spy and trojan respectively.
- *$CpS[i]$ and $CpT[i]$* : The counters to keep track of Spy and Trojan's blocks respectively in a suspicious set i . This is also defined in Section IV-B.
- *updateCounter(i, p_{in}, p_{out})*: Updates $CpS(i)$ or $CpT(i)$ for the set i . Here w is the incoming block. Here p_{in} and p_{out} are process id of incoming block and victim block respectively.
- *findVictimExcept(i, p)*: Returns $V_x(i, p)$, as discussed in Algorithm 1.
- *th_s and th_t* : Threshold partition size of spy and trojan respectively.

C. TPPD for LRU Replacement Policy

Algorithm 1 shows an implementation of $V_x(s', p)$ for LRU replacement policy. The alternative victim ($V_x(s', p)$) in LRU

Algorithm 1: Modified Eviction Policy for LRU

Input: sId : Cache set index. $omitP$: Process whose block not to evict.

Output: A victim block not belonging to process $omitP$.

```

1 Function findVictimExcept ( $sId, omitP$ ):
2   /*Assuming  $cache[row, assoc]$  is the
   set-associative LLC where  $row$  is the total sets and
    $assoc$  is the associativity.  $victim\_index$  and
    $max\_age$  are two variables.*/
3    $k = 0$ 
4   while  $k \neq assoc$  do
5     if  $checkOwner(sId, cache[sID, k], omitP)$  is
       FALSE then
6        $victim\_index = k$ 
7        $max\_age = cache[sID, k].age$ 
8       break
9     end
10     $k++$ 
11  end
12  while  $k \neq assoc$  do
13    if ( $max\_age < cache[sId, k].age$ ) then
14      if
          $checkOwner(sId, cache[sID, k], omitP)$ 
         is FALSE then
15         $victim\_index = k$ 
16         $max\_age = cache[sID, k].age$ 
17      end
18    end
19     $k++$ 
20  end
21  return  $cache[sId, victim\_index]$ 
22 end

```

can be selected in two ways. The first method is selecting a random block not belonging to p , and the second method is to choose the oldest block that does not belong to p . The algorithm describes the second method. All the experimental analyses in this paper use the second method. The complete mechanism of TPPD in terms of LRU replacement policy is discussed in Algorithm 2. The terminology used in this algorithm is described in the box. As mentioned in Section IV-B, we have used s' to represent the targeted set, but to write an algorithm, we have used i to represent any LLC set. The set can be either suspicious or non-suspicious. Here $0 \leq i < N$, and N is the total sets of LLC.

The main function in this algorithm is $EvictionVictim(i, p)$, which selects the appropriate victim block from the LLC as per the requirement of TPPD. The function is called every time a process p requests for a block and has a conflict miss. Process p is called an incoming process, and the block it requested as an incoming block. Initially, an original eviction victim w is selected per the LRU replacement policy (Line 2). Then it checks if the requested set is under attack or not (Line 3). If not, the original victim block w is returned (Line 4). However, additional checks are

Algorithm 2: Modified LRU Replacement Policy for implementing TPPD

Input: i : set index, p : process requesting replacement.

Output: w : way index of block to be evicted from set i

```

1 Function EvictionVictim ( $i, p$ ):
2    $w = getLRUVictim(i)$ 
3   if  $isUnderAttack(i)$  is FALSE then
4     return  $w$  // Return  $V(i)$ .
5   else
6     /* The set  $i$  is under attack. */
7      $pT = getTrojan(i)$ ;  $pS = getSpy(i)$ 
8     if  $checkOwner(i, w, pS)$  is TRUE then
9       // If the owner of  $w$  is spy.
10       $p_w = pS$ 
11    else if  $checkOwner(i, w, pT)$  is TRUE then
12      // If owner of  $w$  is trojan.
13       $p_w = pT$ 
14    else
15       $p_w = -1$  // if owner of  $w$  is an
16      innocent process.
17    end
18    if ( $(p! = pS) \& (p! = pT) \mid \mid (p == p_w)$ ) then
19      /* Incoming block from
20      innocent process or both
21      incoming and victim from same
22      process. */
23       $updateCounter(i, p, p_w)$ ; return  $w$ 
24    else if ( $(p_w == pS) \& (CpS[i] <$ 
25       $th_s) \mid \mid ((p_w == pT) \& (CpT[i] < th_t))$ )
26    then
27       $w'' = findVictimExcept(i, p_w)$ 
28      if  $checkOwner(i, w'', p)$  is FALSE then
29         $updateCounter(i, p, -1)$ 
30      end
31      return  $w''$ 
32    else
33       $updateCounter(i, p, p_w)$ ; return  $w$ 
34    end
35  end
36 Def  $updateCounter(i, p_{in}, p_{out})$ :
37   /*  $p_{in}$  is the process of incoming
38   block and  $p_{out}$  is the process of
39   outgoing (or victim) block. */
40  if  $p_{in}! = p_{out}$  then
41    if  $p_{out} == pS$  then
42       $CpS[i] --$ 
43    else if  $p_{out} == pT$  then
44       $CpT[i] --$ 
45    end
46    if  $p_{in} == pS$  then
47       $CpS[i] ++$ 
48    else if  $p_{in} == pT$  then
49       $CpT[i] ++$ 
50    end
51  end
52 end

```

required if the request is for the suspicious set. Everything mentioned from Line 5 onward in this function is for a suspicious set. First, the owner process p_w of the victim block w is determined (Line 6 to Line 13). The p_w can be either spy (pS), trojan (pT), or an innocent process. If the incoming process p is innocent, it can replace any block from the set. Hence, in that case, w is evicted irrespective of p_w and the corresponding counter ($CsT[i]$ or $CpT[i]$ based on p_w) needs to be updated. Line 14 and 15 do the same. The function $updateCounter(i, p_{in}, p_{out})$ is used for this which is discussed later. Similarly, the original victim w is evicted when the incoming block and victim block belong to the same process but without an update in counters.

The condition mentioned in Line 16 is true only when these three conditions are true: (a) p and p_w are not the same, (b) the incoming process p is either trojan or spy, and (c) the value of counter corresponding to p_w is less than the threshold. In this case the alternative victim block w'' is selected by calling function $findVictimExcept(i, p)$ (Line 17). This function is defined in Algorithm 1. New victim block w'' can be from the incoming process p or from an innocent process. The counter update is required only in the latter case (Line 18-20). In case the condition written in Line 16 is false then the original LRU victim w is selected as victim block (Line 23), and counters are updated. Function $updateCounter(i, p_{in}, p_{out})$ keeps spy and trojan updated based on the id of incoming process (p_{in}) and process of victim block (p_{out}). Note that no counter updates are required if p_{in} and p_{out} are the same. Also, counter updates are not required in non-suspicious sets. Since the counter maintains the total number of blocks available in the set for trojan and spy, they may need to be incremented or decremented during an eviction. From the algorithm, it can be observed that TPPD applies dual victim policy only to the suspicious sets. Also, the need for selecting the alternative victim is only required when trojan and spy tries to remove each other's block. Section V shows experimental analysis on different threshold values (th_s and th_t) that can be used for this algorithm.

D. How does TPPD effectively mitigate covert channel attack?

In this part, we give a theoretical analysis of how TPPD dismantles LLC-based cross-core covert channel attacks with minimal impact on system performance. Our observations and their explanations are described as follows:

Observation 1: The communication between the spy and the trojan of the covert channel attack is degraded when the suspicious process pair is prevented from accessing all the ways of a targeted cache set.

Explanation: A low and high cache access latency pattern observed by a spy to determine the bit transmitted by trojan is at the foundation of a successful covert channel attack. Trojan establishes this pattern by removing all of the spy's data in one case (e.g., bit 1) or none at all for another bit (bit 0). In TPPD design, cache interference between spy and trojan is decreased by restricting them from evicting each other's block when the current number of blocks of these processes falls below the threshold. As a result, the trojan is unable to replace all of the

spy's data for sending bit 1, and the spy cannot fully prime the cache set by replacing trojan's data with its. In this way, the spy's easily identifiable access latency pattern is obscured.

Observation 2: A targeted defence mechanism that affects only suspected sets and processes limit the overall system performance degradation.

Explanation: The main feature of our proposed defense mechanism is that it is a targeted countermeasure. It affects only processes and cache sets detected to be participating in covert channel attack by the CCA detector. Modified replacement policy to implement TPPD design is activated on suspicious sets, and it restricts only conflict misses occurring between the suspicious processes. Original replacement policy is implemented for remaining sets not participating in covert channel attack. On suspicious sets, trojan and spy processes have restrictions on evicting each other's block. The innocent processes have access to the entire cache set and can evict block of any process as determined by the original replacement policy. However, when suspicious processes are restricted from evicting each other's block, a new block is selected for eviction. This can lead to premature eviction of benign process blocks, thus slightly impacting overall system performance.

E. Calculating Performance of TPPD

TPPD can use different threshold values for th_s and th_t . These values can be either the same or different for spy and trojan. In the rest of the discussion, we have considered the same values ($th_s = th_t$) for both spy and trojan. The TPPD- z is considered as a configuration having $th_s = th_t = z$. The minimum value of z is 1, and the maximum is $A/2$, where A is the associativity of the LLC. The value of z can be changed dynamically, but for most of our experiments, we have fixed the value of z throughout the execution.

We have measured the isolated performance impact on benign applications through the following steps:

- When only the attacker processes are executing in the system, the impact of TPPD- z in terms of LLC misses ($DIFF_z$) is calculated by differentiating number of misses in LLC having TPPD- z and no defense.

$$DIFF_z = Misses_z - Misses_0 \quad \forall z \in (1, A/2) \quad (1)$$

Here $Misses_z$ means total misses on LLC while TPPD- z is active. $Misses_0$ means total LLC misses when no defense mechanism is implemented in LLC.

- In the next scenario, a Parsec benchmark [34] is run parallel to the attack processes. The difference in LLC misses ($DIFF'_{z,b}$) encountered in TPPD- z ($Misses_{z,b}$) and without defense ($Misses_{0,b}$) is measured for benchmark b .

$$DIFF'_{z,b} = Misses_{z,b} - Misses_{0,b} \quad \forall b \in PARSEC \quad (2)$$

Here $Misses_{z,b}$ means total LLC misses while TPPD- z is active. $Misses_{0,b}$ means total LLC misses without a defence mechanism.

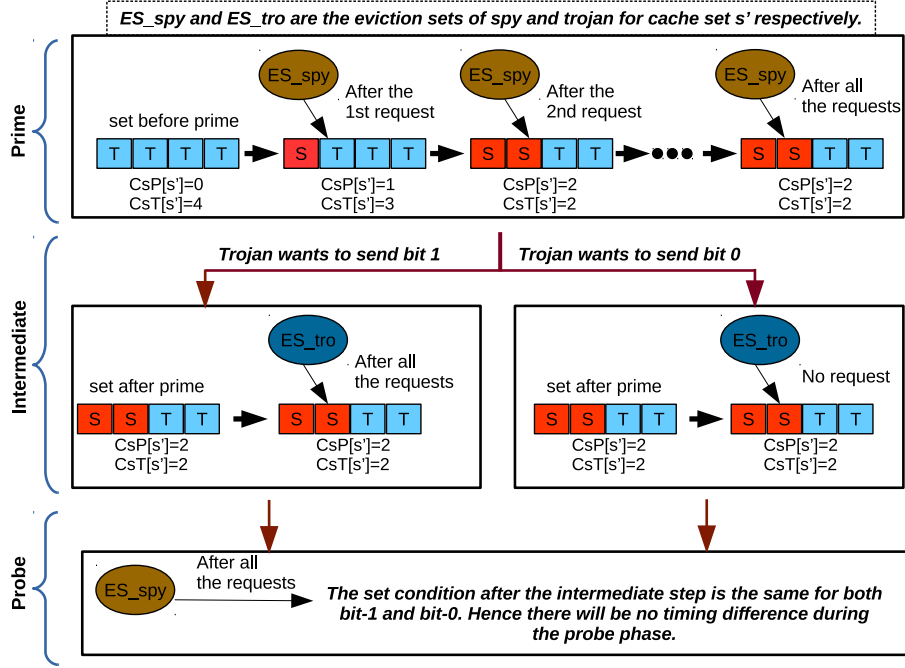


Fig. 5: An example to demonstrate the working of TPPD after detecting a CCA on set s' for pS (as spy) and pT (as trojan). We have assumed that the set s' contains all the blocks of trojan during the detection. The prime process starts with this assumption. The TPPD configuration considered here is TPPD-2. The figure shows that the Prime+Probe method cannot perform CCA once TPPD-2 becomes active.

- The isolated impact of TPPD- z on a benign application ($D_{z,b}$) and average impact ($AvgD_z$) is calculated as described in Equation 3 and 4 respectively.

$$D_{z,b} = DIFF'_{z,b} - DIFF_z \quad (3)$$

$$AvgD_z = \sum_{b \in PARSEC} \frac{D_{z,b}}{T_n} \quad (4)$$

Here, T_n is the total number of PARSEC applications considered.

These equations are used in Section V for performance analysis. The higher positive difference ($AvgD_z$) indicates more significant performance degradation of benign processes. A difference close to zero will signify that there has been no effect on these processes.

F. An example of TPPD

An example discussing how our proposed approach deteriorates covert channel communication on LLC is shown in Figure 5. Consider the Prime+Probe based CCA attack mounted on set s' of LLC. After attack detector detects this attack; it sends the set id s' along with process ids pS and pT to TPPD. When TPPD-2 is deployed initially, we assumed the entire cache set s' contains trojan's block (trojan is sending 1) as shown in Step 2 of Figure 2. This assumption is considered without any loss of generality and a better understanding of TPPD working.

During the prime phase, the spy requests for blocks from its eviction set. From the figure, it can be observed that at the

end of all the block requests of the prime phase, the spy could not be able to remove all the trojan's block from the set. This is because of the dual victim policy of TPPD-2. In its first two accesses, the spy is able to replace trojan's block as intended. However, for the next spy access, if the original LRU victim spy block is evicted, the number of spy's blocks present in the set falls below the threshold of 2, which is not allowed by TPPD-2. In this case, as per the dual victim policy of TPPD, an alternative replacement victim that does not belong to the trojan process is selected using $V_x(s', pT)$. After the prime phase, in the case of sending bit 1, the trojan removes all the blocks of spy by requesting block from its own eviction set. However, because of TPPD-2, the trojan cannot remove all the spy's blocks. From the figure, it can be observed that, after the intermediate phase, the status of the set is the same for both bit 1 and bit 0. In the probe phase, when spy probes set s' for the blocks from its eviction set, it encounters 2 misses and 2 hits in both cases of a bit 1 and bit 0.

G. Implementing TPPD on other Replacement Policies

The proposed TPPD can be implemented using any replacement policy where dual victim policy can be implemented. Algorithm 2 can be modified for that. The internal mechanism of the functions called from Line 2 and Line 21 of this algorithm needs to be changed as per the requirement of the replacement policy. In this work, we have not explored the TPPD for any other replacement policy except LRU, but that can be explored without any major changes. The proposed TPPD can be considered as a countermeasure for cross-core

TABLE II: System configuration for the experimental setup

Baseline Architecture	
Simulator	gem5 [21]
Architecture	4 cores each at 2.0 GHz
Coherence Protocol	MESI two level
Level 1 Cache	Inclusive L1I/L1D Private, 64 KB, 4-way, 2 cycles latency, 64B blocks
Level 2 Cache (LLC)	Inclusive Shared, 2MB 8-way, 18 cycles latency, 64B blocks
Main Memory	DRAM, 250-cycle latency
Replacement Policy	Least Recently Used(LRU)

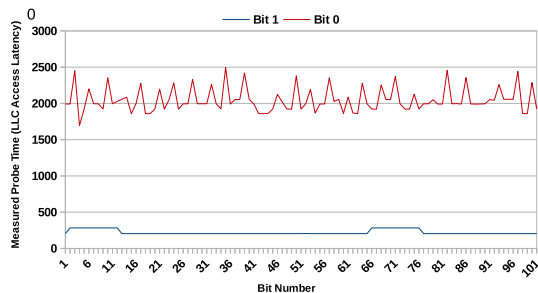


Fig. 6: Cache access latency observed by spy for bit 0 and 1 in LLC with no TPPD

CCA that can make any existing replacement policy secure from CCA.

V. EXPERIMENTS AND RESULTS

To perform the experimental analysis, TPPD is implemented on a full system cycle-accurate simulator, gem5 [21]. We consider a 4-core system setup with two levels of the cache hierarchy. Each core has its private L1 cache, and the L2 cache is considered a shared LLC. Other parameters of the setup are shown in Table II. To simulate the CCA attack, we have developed our own trojan and spy applications. These applications are written in C++ and can be executed on gem5. We considered a generic Prime+Probe attack as described in Figure 2 in this application. The ability to perform CCA attack by these two applications have experimented first on gem5. Parsec benchmarks [20] are used to measure the performance of the innocent applications in the presence of attacker applications. From the 4-cores, two cores are assigned to the attacker processes (spy and trojan). A multi-threaded Parsec benchmark is binded to the other two cores. We have considered all the Parsec benchmarks as innocent applications. To measure the worst-case performance impact, the attacking applications are developed such that once the attack starts, it continues during the execution of the system. Different configurations of TPPD (TPPD- z) are considered in order to fully analyse the performance behaviour of TPPD. The value of z varies from 1 to 4 (8 is the associativity of the LLC). TPPD-0 means baseline design with no defence mechanism.

A. Security Analysis

We assess the effectiveness of TPPD against cross-core based covert channel attacks on LLC. As described in Section

IV-D, TPPD reduces the difference in probe time observed by the spy for bit 0 and 1. This difference is substantially more significant when no defense mechanism is active, as evident by Figure 6. Here bit 0 represents cache miss and bit 1 as a cache hit. TPPD reduces this difference such that the information received by the spy becomes noisy. This difference becomes close in TPPD-1, TPPD-2 and TPPD-3 but does not overlap as represented in Figure 7a, 7b and 7c respectively. These TPPD configurations are effective against attacks based on fine-grained information, i.e., observing the number of cache accesses by victim applications in the targeted set. However, in a low-speed covert channel attack, we considered these configurations would not be effective, especially when there is low system noise. However, TPPD-4 was able to completely obfuscate the receiver’s access time pattern used to identify bit 0 and 1 as illustrated in Figure 7d. Hence, TPPD- $A/2$ is the recommended configuration to prevent most CCA attacks. Here A is the associativity of the LLC.

B. Performance Analysis

While deconstructing cache timing channel-based attacks, the attack defense mechanism must guarantee that system performance does not suffer. The Figure 8 represent the total LLC misses encountered in different benchmark applications for TPPD- z and TPPD-0. The upper part of the figure shows the actual values (in millions), while the bottom part shows the normalised LLC misses. The values are normalised w.r.t. TPPD-0. As mentioned above, the benchmarks are combined with the attacker applications during the execution. However, in the figure, the benchmark named as ”only-attack” is not a Parsec benchmark. This is designed only with the combination of trojan and spy. To analyse the performance impact of TPPD on innocent applications, it is important to first analyse the performance of ”only-attack”. From the figure, it can be observed that, in the case of ”only-attack”, if TPPD is applied to prevent CCA, the LLC misses are increased by 17% as compared to TPPD-0 (no defence). All the configurations of TPPD show almost similar results. However, for the rest of the benchmarks, the increment is less. On average, while a Parsec benchmark executes with two attacker applications, which are constantly trying to covertly communicate, the performance degradation in TPPD are 05%, 07%, 06% and 09% for TPP-1, TPPD-2, TPPD-3, and TPPD-4 respectively. The degradation in TPPD-4 is more because a higher threshold value may evict some more innocent blocks from the cache. A similar result can be seen in Figure 9 for Misses per Thousand Instructions (MPKI).

The important point here is that the misses are higher when TPPD applies only to the attacker set in ”only-attack”. However, the miss overhead is reduced when the cores mix attack applications with some innocent Parsec benchmarks. It means that the miss overhead of TPPD is higher only for the attacker applications. The overhead is less for innocent applications. The two main reasons for this are:

- TPPD is a targeted CCA countermeasure. Hence from the non-targeted set, the victim is selected per the original replacement policy.

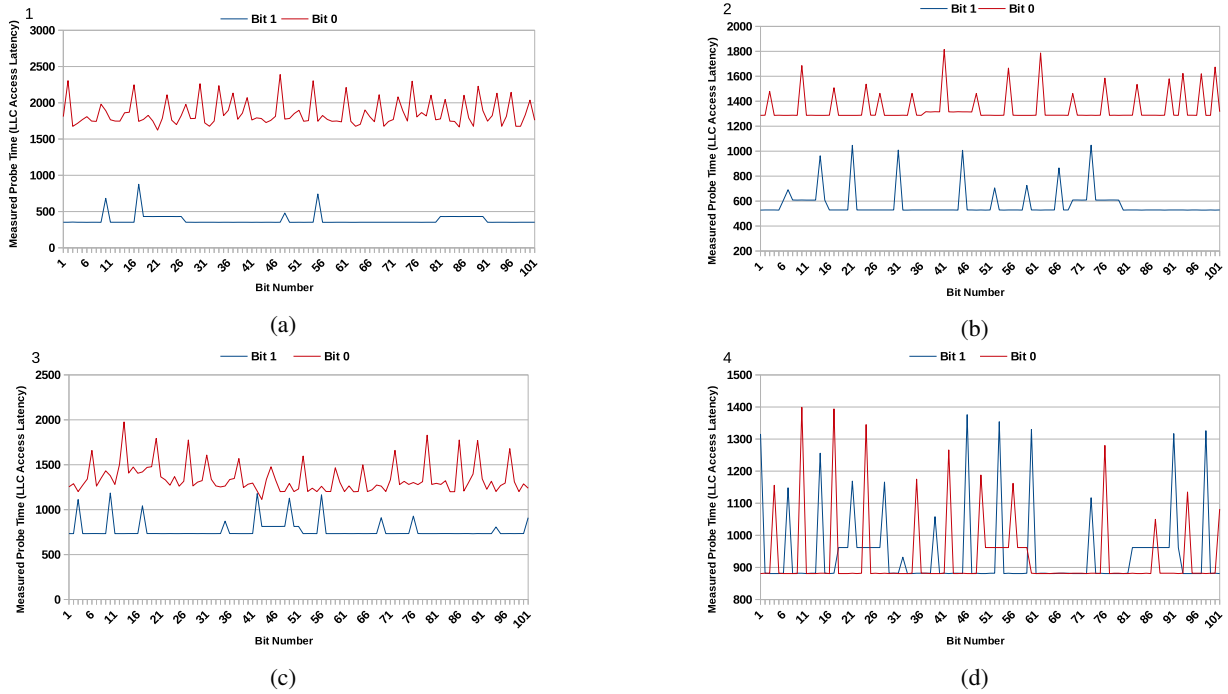


Fig. 7: Cache access latency for bit 0 and bit 1 in (a) TPPD-1, (b) TPPD-2, (c) TPPD-3, and (d) TPPD-4

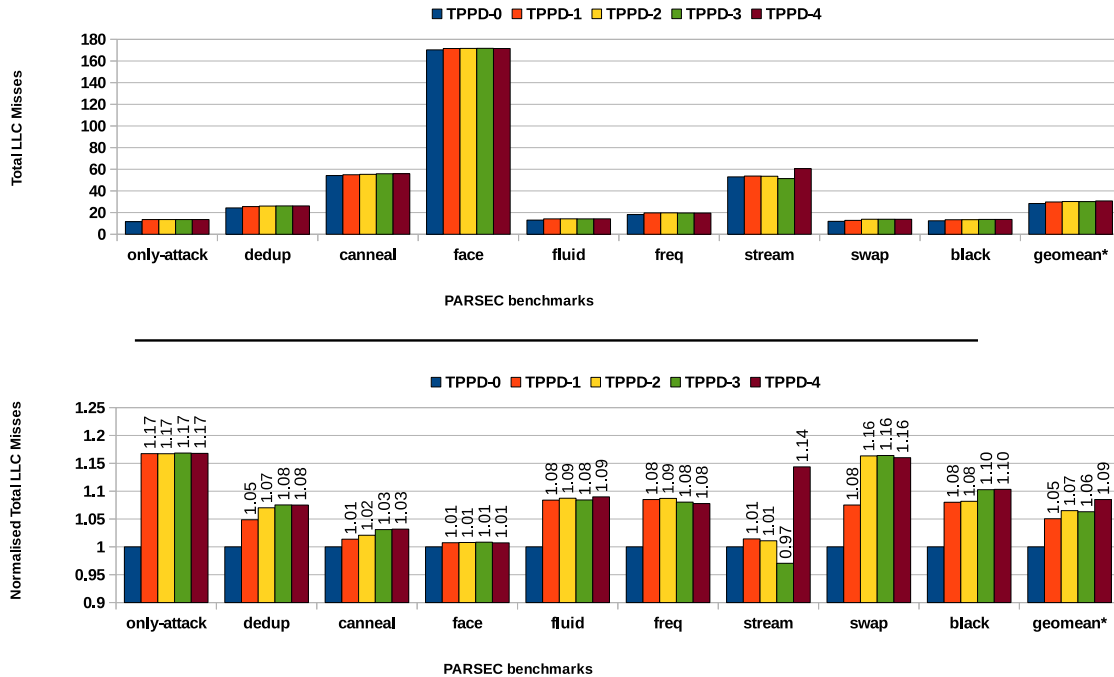


Fig. 8: Total number of misses in LLC for PARSEC benchmarks on different TPPD configurations. *The calculation of geomean is excluding “only-attack”. The lower part of the figure shows the LLC misses normalised to TPPD-0.

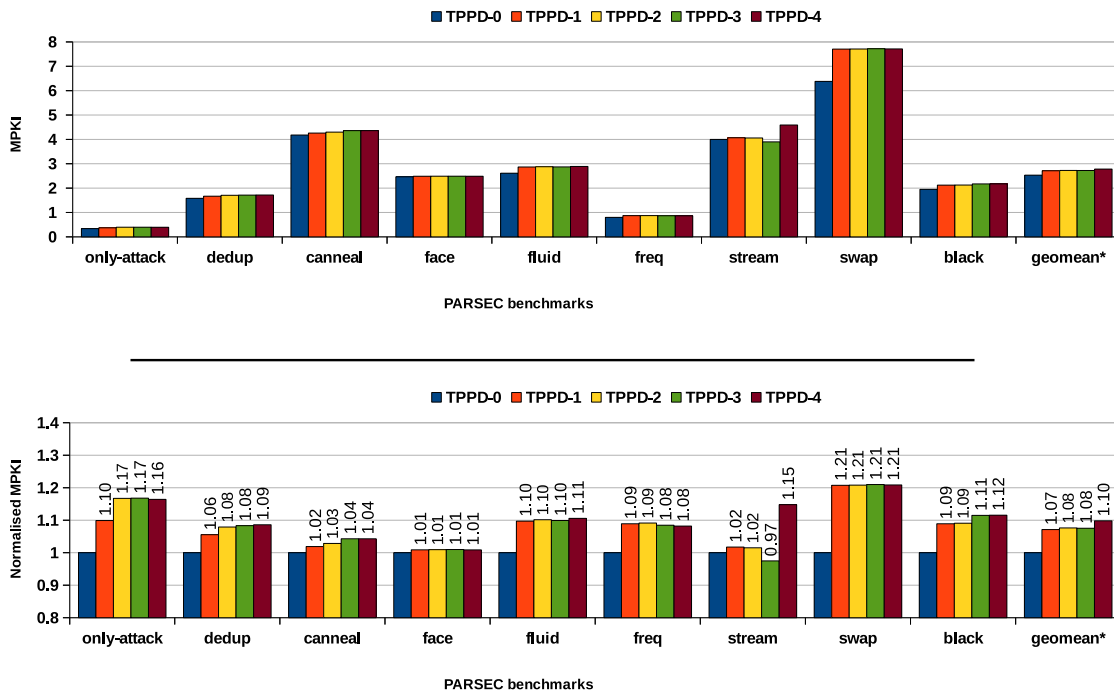


Fig. 9: Total MPKI in LLC for PARSEC benchmarks on different TPPD configurations. *The calculation of geomean is excluding “only-attack”. The lower part of the figure shows the MPKI normalised to TPPD-0.




Fig. 10: Isolated LLC misses for PARSEC benchmarks with different TPPD configurations measured using Equation 3

- In the targeted set, the innocent blocks has no restrictions, and they can remove any other block from the set.

Because of these two reasons, TPPD becomes an efficient countermeasure for cross-core covert channel attacks.

Figure 8 and 9 only shows the combined miss overhead when the innocent applications are running with attacker applications. To understand the isolated impact of TPPD on innocent applications, we have performed experiments based on Equation 3 and Equation 4. As evident from Figure 10, the increment in total LLC misses (compared to TPPD-0) when benchmark runs alongside attack process is always less than the misses when just attack runs in LLC augmented with TPPD- z . In this figure, a less number (large negative number) means less overhead on an application. On average, the applications face the lowest overhead in TPPD-1. However, the miss overhead in TPPD-4 (which is the recommended configuration) is also less.




Fig. 11: Comparison of MPKI for NOMO-2 vs TPPD-4 when PARSEC benchmark are executed in parallel to CCA. Here *geomean is calculated excluding “only-attack”.

C. Comparison with Existing Partitioning based Secure Cache Architectures

Existing partitioning-based approaches such as PLCache [23], SecVerilog Cache [35], and SecDCP [25] can combat cache-based side-channel attacks in which an attacker process attempts to reveal the cache access pattern of an innocent victim process in order to leak its secret, such as the private

key of cryptography algorithms. These secure cache designs rely on a reliable operating system, compiler, or programmer to recognise security-critical processes or data that could be vulnerable to cache timing channel attacks. PLCache pre-loads and lock the security critical data in the cache and prevents its eviction by unlocked data. In SecVerilog and SecDCP, processes are assigned to different security classes, and cache interference between these security domains is restricted. These security measures are ineffective in the presence of covert timing channel attacks where both processes are malicious.

The static cache partitioning based secure cache design: NOMO [24] can defend against covert channel attacks by completely isolating ways of spy and trojan. In our setup, NOMO-2, in which two ways were fixed per process, was able to dismantle the covert channel attack. However, significant performance degradation was observed in the case of NOMO-2 in comparison to our design TPPD-4 as shown in Figure 11 in terms of MPKI. It can be observed from the figure that NOMO shows less miss overhead over TPPD-4 when executed for “only-attack”. However, when innocent applications are executed with attacker applications, the overhead increases significantly. The main reason for that is the non-targeted nature of NOMO. The restricted eviction mechanism is uniformly applied to all the sets. Hence, the overhead of misses increases in non-targeted sets also. Since the attackers mainly requested for the blocks mapped to the targeted set, the other set mostly served innocent applications. Uniform applications of restriction eviction policy impact the miss rate of innocent application in non-targeted sets. On average, NOMO-2 had 23% higher MPKI than TPPD-4.

D. Storage and Latency Overhead

The structure of the LLC having TPPD as a countermeasure is already discussed in Section IV-B1. Figure 3(b) shows the two additional components required for TPPD: (a) TPPD Components, and (b) CCA Detector. In this section, we have discussed the storage and latency overhead of the TPPD components. As mentioned in Section IV-A, we have used the existing CCA detector technique [19] to detect the CCA attack. Each cache line is augmented with its owner process id to identify the cache block’s owner process, as seen in Figure 3(b). An alternative technique followed by some existing works [15], [16] is to maintain core id instead of process id along with each cache line. However, the overhead of storing process id along with the cache line cannot be considered as the overhead of TPPD. The issues with maintaining process id along with the cache lines are already discussed in Section IV-B3.

Figure 3(b) shows all the counters used in TPPD Components. The first counter, *attack_flag*, is a one-bit counter used to indicate whether or not the set is under attack. The next two fields (*pS* and *pT*) specify the attacking process pair if the set is under attack. The last two fields are spy and trojan counters, each with a size of $\log_2(A)$ bits to identify the number of spy and trojan blocks currently in the set. Here, A is the associativity of the LLC. Thus, the total storage overhead can

TABLE III: Storage overhead calculation of TPPD as per the LLC configuration mentioned in Table II. We have assumed 16 bits to represent a process id as the maximum process id possible in Linux is 32768.

Storage Overhead	y as core id		y as process id	
	per set	per LLC	per set	per LLC
status bit	1	512 Bytes	1	512 Bytes
suspicious process1	2 bits	$2 \times 4096 = 1 \text{ KB}$	16 bits	$16 \times 4096 = 8 \text{ KB}$
suspicious process2	2 bits	$2 \times 4096 = 1 \text{ KB}$	16 bits	$16 \times 4096 = 8 \text{ KB}$
spy Counter	3 bits	$3 \times 4096 = 1.5 \text{ KB}$	3 bits	$3 \times 4096 = 1.5 \text{ KB}$
trojan Counter	3 bits	$3 \times 4096 = 1.5 \text{ KB}$	3 bits	$3 \times 4096 = 1.5 \text{ KB}$
Total Overhead	11 bits	5.5 KB	39 bits	19.5 KB

be determined as $N \times (1 + 2(y + \log_2 A))$. Here N is the total number of sets in LLC, and y is the bits required to represent a suspicious process. In our work, y represents the cores in the system as a single process is binded per core making core id enough for identifying attacking pair. Table III shows the storage overhead of TPPD as per the LLC configurations mentioned in Table II. It can be observed from Table III that the total storage required for TPPD is 5.5 KB , i.e. 0.26% of total LLC size. However, in real systems, multiple processes may run on the same core, making core id not sufficient for uniquely identifying processes. The table shows that when process id is used instead of core id, the total size of this additional structure is 20.5 KB , i.e. $\approx 1\%$ of total LLC size.

The operations of TPPD are performed in the background without affecting the critical execution path of the system.

VI. CONCLUSION

This paper proposes an effective and efficient mitigation mechanism, TPPD, for cross-core cache timing channel attacks. TPPD implements way-wise partitioning on the cache sets used for covert channel attacks but only for suspicious process pairs. However, remaining benign processes have unrestricted access to these and other sets, reducing the performance impact on system performance. It successfully abolishes LLC based covert communication between trojan and spy. Experiments have shown that it does not have any significant impact on the performance of benign applications (Parsec benchmark). The total storage overhead required for implementing TPPD design is approximately $\approx 0.26\%$ of LLC size. Compared to the existing partitioning based attack prevention mechanism NOMO, it caused 23% less LLC misses.

REFERENCES

- [1] C. Percival, “Cache Missing For Fun And Profit,” in *In Proc. of BSDCan*, 2005.
- [2] 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.
- [3] D. A. Osvik, A. Shamir, and E. Tromer, “Cache Attacks and Countermeasures: The Case of AES,” in *Cryptographers’ track at the RSA conference*, 2006, pp. 1–20.
- [4] D. J. Bernstein, “Cache-Timing Attacks on AES,” 2005.

- [5] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee, "Last-Level Cache Side-Channel Attacks are Practical," in *Symp. on Security and Privacy*, 2015, pp. 605–622.
- [6] J. Kaur and S. Das, "A survey on cache timing channel attacks for multicore processors," *Journal of Hardware and Systems Security*, pp. 1–21, 2021.
- [7] Y. Lyu and P. Mishra, "A survey of side-channel attacks on caches and countermeasures," *Journal of Hardware and Systems Security*, vol. 2, no. 1, pp. 33–50, 2018.
- [8] V. Rijmen and J. Daemen, "Advanced Encryption Standard," *Federal Information Processing Standards Publications, National Institute of Standards and Technology*, pp. 19–22, 2001.
- [9] R. L. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-key Cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.
- [10] D. Johnson, A. Menezes, and S. Vanstone, "The elliptic curve digital signature algorithm," *International journal of information security*, vol. 1, no. 1, pp. 36–63, 2001.
- [11] O. Aciğmez and Ç. K. Koç, "Trace-Driven Cache Attacks on AES (Short Paper)," in *Information and Communications Security*, 2006, pp. 112–121.
- [12] D. Gullasch, E. Bangerter, and S. Krenn, "Cache Games—Bringing Access-Based Cache Attacks on AES to Practice," in *IEEE Symp. on Security and Privacy*, 2011, pp. 490–505.
- [13] S. Sari, O. Demir, and G. Kucuk, "FairSDP: Fair and Secure Dynamic Cache Partitioning," in *4th Intl. Conf. on Computer Science and Engineering*, 2019, pp. 469–474.
- [14] F. Yao, H. Fang, M. Doroslovački, and G. Venkataramani, "COTSknighT: Practical Defense against Cache Timing Channel Attacks using Cache Monitoring and Partitioning Technologies," in *Hardware Oriented Security and Trust*.
- [15] M. K. Qureshi and Y. N. Patt, "Utility-based cache partitioning: A low-overhead, high-performance, runtime mechanism to partition shared caches," in *2006 39th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO'06)*. IEEE, 2006, pp. 423–432.
- [16] P. D. Halwe, S. Das, and H. K. Kapoor, "Towards a better cache utilization using controlled cache partitioning," in *2013 IEEE 11th International Conference on Dependable, Autonomic and Secure Computing*. IEEE, 2013, pp. 179–186.
- [17] C. Yang, L. Liu, K. Luo, S. Yin, and S. Wei, "Ciaccp: A correlation-and iteration-aware cache partitioning mechanism to improve performance of multiple coarse-grained reconfigurable arrays," *IEEE Transactions on Parallel and Distributed Systems*, vol. 28, no. 1, pp. 29–43, 2016.
- [18] A. Agarwal, J. Kaur, and S. Das, "Exploiting secrets by leveraging dynamic cache partitioning of last level cache," in *Design, Automation Test in Europe Conference Exhibition (DATE)*, 2021.
- [19] H. Fang, S. S. Dayapule, F. Yao, M. Doroslovački, and G. Venkataramani, "Product: Prefetch-obfuscator to Defend against Cache Timing Channels," *Intl. Journal of Parallel Programming*, vol. 47, no. 4, pp. 571–594, 2019.
- [20] C. Bienia, *Benchmarking modern multiprocessors*. Princeton University, 2011.
- [21] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti *et al.*, "The gem5 simulator," *ACM SIGARCH computer architecture news*, vol. 39, no. 2, pp. 1–7, 2011.
- [22] G. Irazoqui, T. Eisenbarth, and B. Sunar, "SSA: A Shared Cache Attack That Works across Cores and Defies VM Sandboxing – and Its Application to AES," in *IEEE Symp. on Security and Privacy*, 2015, pp. 591–604.
- [23] Z. Wang and R. B. Lee, "New Cache Designs for Thwarting Software Cache-based Side Channel Attacks," *SIGARCH Comput. Archit. News*, vol. 35, no. 2, pp. 494–505, 2007.
- [24] L. Domnitsler, A. Jaleel, J. Loew, N. Abu-Ghazaleh, and D. Ponomarev, "Non-Monopolizable Caches: Low-Complexity Mitigation of Cache Side Channel Attacks," *Transactions on Architecture and Code Optimization*, vol. 8, no. 4, pp. 1–21, 2012.
- [25] Y. Wang, A. Ferraiuolo, D. Zhang, A. C. Myers, and G. E. Suh, "Secdcp: Secure dynamic cache partitioning for efficient timing channel protection," in *2016 53rd ACM/EDAC/IEEE Design Automation Conference (DAC)*, 2016, pp. 1–6.
- [26] M. K. Qureshi, "Ceaser: Mitigating conflict-based cache attacks via encrypted-address and remapping," in *2018 51st Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, 2018, pp. 775–787.
- [27] —, "New attacks and defense for encrypted-address cache," in *Proceedings of the 46th International Symposium on Computer Architecture*, ser. ISCA '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 360–371.
- [28] P. Kumar, C. S. Yashavant, and B. Panda, "Damaru: A denial-of-service attack on randomized last-level caches," *IEEE Computer Architecture Letters*, vol. 20, no. 2, pp. 138–141, 2021.
- [29] A. Jain and C. Lin, "Back to the future: Leveraging belady's algorithm for improved cache replacement," in *Proceedings of the 43rd International Symposium on Computer Architecture*, ser. ISCA '16. IEEE Press, 2016, p. 78–89.
- [30] H. Fang, F. Yao, M. Doroslovački, and G. Venkataramani, "Negative correlation, non-linear filtering, and discovering of repetitiveness for cache timing channel detection," in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 2522–2526.
- [31] A. Jaleel, K. B. Theobald, S. C. Steely Jr, and J. Emer, "High performance cache replacement using re-reference interval prediction (rrip)," *ACM SIGARCH Computer Architecture News*, vol. 38, no. 3, pp. 60–71, 2010.
- [32] K. K. Dutta, P. N. Tanksale, and S. Das, "A fairness conscious cache replacement policy for last level cache," in *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2021, pp. 695–700.
- [33] T. S. Warrior, B. Anupama, and M. Mutyam, "An application-aware cache replacement policy for last-level caches," in *International Conference on Architecture of Computing Systems*. Springer, 2013, pp. 207–219.
- [34] C. Bienia and K. Li, "Parsec 2.0: A new benchmark suite for chip-multiprocessors," in *Proceedings of the 5th Annual Workshop on Modeling, Benchmarking and Simulation*, vol. 2011, 2009, p. 37.
- [35] D. Zhang, Y. Wang, G. E. Suh, and A. C. Myers, "A hardware design language for timing-sensitive information-flow security," *Acm Sigplan Notices*, vol. 50, no. 4, pp. 503–516, 2015.



Jaspinder Kaur A Ph.D scholar in the CSE department of Indian Institute of Technology Ropar. She has completed M.Tech (CSE) from the Department of Computer Engineering of Punjabi University Patiala, India. Her research interests include Computer Architecture, Cache Prefetching and Cache Security against timing channel attacks.



Shirshendu Das Received a Ph.D. degree (CSE) from the Indian Institute of Technology Guwahati, India, in 2016. Previously he did M.Tech (CSE) from the Indian Institute of Technology Guwahati, India. Presently he is an Assistant Professor in the Department of CSE, Indian Institute of Technology Ropar, Punjab, India. His area of research includes Computer Architecture, Network-on-Chip, and Hardware Security.