CacheLight: Defeating the CacheKit Attack

Mauricio Gutierrez
Arizona State University
maguti14@asu.edu

Ziming Zhao
Arizona State University
ziming.zhao@asu.edu

Adam Doupé
Arizona State University
doupe@asu.edu

Yan Shoshitaishvili
Arizona State University
yans@asu.edu

Gail-Joon Ahn
Arizona State University
gahn@asu.edu

ABSTRACT

To protect software systems from attacks, ARM introduced a hardware security extension known as TrustZone. TrustZone provides an isolated execution environment, which can be used to deploy various memory integrity and malware detection tools. However, a new type of rootkit, namely CacheKit, can exploit cache incoherency and cache locking mechanisms in TrustZone to hide itself from such inspections. Therefore, it is imperative to design a new approach to ensure the correct use of cache locking and prevent malicious code from being hidden in the cache.

In this paper, we present CacheLight, which leverages the TrustZone and Virtualization extensions of the ARM architecture to allow the system to continue to securely provide these hardware facilities to users while preventing attackers from exploiting them. CacheLight restricts the ability to lock the cache to the Secure World of the processor such that the Normal World can still request certain memory to be locked into the cache by the secure operating system (OS) through a Secure Monitor Call (SMC). This grants the secure OS the power to verify and validate the information that will be locked in the requested cache way thereby ensuring that any data that remains in the cache will not be inconsistent with what exists in main memory for inspection. Malicious attempts to hide data can be prevented and recovered for analysis while legitimate requests can still generate valid entries in the cache.

CCS CONCEPTS

• **Security and privacy → Embedded systems security**; *Tamper-proof and tamper-resistant designs*; *Hardware attacks and countermeasures*; Malware and its mitigation;

KEYWORDS

Hardware Assisted Security; TrustZone; Embedded Systems Security; Rootkit Defense; Cache Locking

ACM Reference Format:
Mauricio Gutierrez, Ziming Zhao, Adam Doupé, Yan Shoshitaishvili, and Gail-Joon Ahn. 2018. CacheLight: Defeating the CacheKit Attack. In *The Second Workshop on Attacks and Solutions in Hardware Security (ASHES'18), October 19, 2018, Toronto, ON, Canada.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3266444.3266449

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ASHES'18, October 19, 2018, Toronto, ON, Canada
© 2018 Association for Computing Machinery.
ACM ISBN 978-1-4503-5996-2/18/10…$15.00
https://doi.org/10.1145/3266444.3266449

1 INTRODUCTION

ARM introduced the TrustZone hardware security extensions, which provide an isolated Trusted Execution Environment (TEE). In TrustZone, the TEE, namely Secure World, is separated from the rich, non-secure Normal World at the hardware level. Secure World has privileged access to all system resources, while Normal World, where the rich OS and untrusted applications run, can only access those resources allocated to it.

Since the code in Secure World has the privilege to access Normal World memory and CPU registers, but not vice versa, system integrity checking and malware detection tools can be installed in the Secure World to monitor the potentially compromised Normal World. Therefore, there have been various research efforts into using the TEE provided by hardware to perform the detection of rootkits in memory [7, 26]. However, all these solutions depend on the ability of the high-privileged TEE to have access over the entire physical memory and the malicious code being present in memory. For example, Trustdump is a forensic toolkit that has a small piece of memory acquisition and integrity checking code stored in Secure World [24]. From the TEE it can attempt to verify physical memory regions of the Normal World, beyond the reach of the potentially corrupted system.

However, in order to avoid memory introspection tools running in the Secure World (SW), more advanced rootkits are exploiting the fact that Secure World does not have access to the Normal World (NW) cache [27]. In the ARM TrustZone architecture all the resources are tagged with an additional bit (the NS bit) that indicates which world they belong to. Although the Secure World has access to many of the Normal World resources, the presence of the additional NS bit means that the Secure World cannot access the Normal World cache, which is referred to as cache incoherence. CacheKit is a new type of rootkit, which exploits cache incoherence and cache locking mechanisms to hide itself from introspection tools running in the Secure World. By hiding the rootkit in Normal World cache, it avoids introspection from Secure World tools and therefore evades the primary method of detection used against rootkits.

These new stealthy attacks leave many real-time and embedded systems that offer cache locking mechanisms exposed and vulnerable. Therefore, a new defense approach is necessary to continue to provide these hardware facilities for timing and performance sensitive processes. Throughout this paper, we study possible solutions to determine the best approach to defending against such

attacks. After thorough research and experimentation, we design and implement CacheLight, a lightweight approach for preventing malicious abuse of cache locking mechanisms. This novel solution leverages both the TrustZone and Virtualization extensions in the ARM architecture to allow legitimate users to continue to utilize cache locking while giving the Trusted Execution Environment (TEE) the power to ensure system security by controlling and verifying use of said mechanisms.

2 BACKGROUND: ARM, CACHE AND TRUSTZONE

2.1 ARM Architecture

With the TrustZone Security Extensions enabled, an ARM processor has 9 modes of operations. The `usr` mode has privilege level 0 and it is where user space programs run. The `svc` mode has privilege level 1 and is where most parts of the kernel run.

At the system level view, the AMRv7 architecture has 16 core registers when Security Extensions are implemented; 13 general-purpose 32-bit registers `R0` - `R12`, the Stack Pointer (`SP`), Link Register (`LR`) and the Program Counter (`PC`). Additionally, the `Current Program Status Register` `CPSR` holds processor status and control information, including the current processor mode. The value of this register is copied and saved upon entry in the respective `Saved PSR` by all modes except User and System.

The ARM architecture supports coprocessors to extend the functionality of the ARM processor. Each coprocessor has its own set of registers. Coprocessor instructions provide access to sixteen coprocessors in the ARMv7 architecture described as `CP0` - `CP15`. `CP15` is called the System Control coprocessor and is reserved for the control and configuration of the ARM processor system, including architecture and feature identification [3].

2.2 ARM TrustZone

TrustZone is a set of hardware security extensions on the processor, memory, and peripherals that ensure complete system isolation for running secure code. The isolated environment is referred to as the Secure World while the Normal World is the name given to the other environment where the OS and other programs run. The Secure World is thought to oversee the Normal World as it has higher access privileges. This means that the Secure World can access most of the resources that belong to the Normal World. However, Normal World does not have access to any of the resources available to the Secure World.

To divide the two worlds, `CP15` has a `Security Configuration Register` (`SCR`) with a non-secure (NS) bit that determines the security context of the processor. When the NS bit is set, this indicates that the processor is in the Normal World and when it is cleared, the processor is in Secure World. The two worlds are separated by adding this NS control bit to all system resources. Monitor Mode is responsible for handling world-switches. To enter Monitor Mode, the process must execute a Secure Monitor Call (SMC) instruction, at which point Monitor Mode handles the world switching and determines the necessary function or handler to execute [1].

2.3 ARM Cache

Caches can be divided into types based on whether the index and tag bits correspond to physical or virtual addresses. *Physically indexed, physically tagged (PIPT)* caches use the physical address for both the index and the tag. While this is simple and avoids problems with aliasing, it is also slow, as the physical address must be looked up in the translation tables before that address can be looked up in the cache. *Virtually indexed, physically tagged (VIPT)* caches use the virtual address for the index. They are faster than PIPT caches as the cache line can be looked up in parallel with the TLB translation, however the tag cannot be compared until the physical address is available. While they still face aliasing problems, the physical tag bits make sharing more manageable than with *Virtually indexed, virtually tagged (VIVT)* caches that use the virtual address for both the index and the tag. In general, the ARM L1 instruction cache is VIPT for greater speed and since instructions generally are not shared. However, the L1 data cache and any lower level caches are generally PIPT, with some options for implementing them as VIPT [4].

With the implementation of TrustZone technology, the processor cache is also extended with an additional tag bit. This tag bit is used to record the type of memory access made; secure or non-secure. This makes it so that the processor always knows which lines of cache belong to which world and cache flushing between world switches is not necessary. The two worlds can evict each other's lines in the cache as needed [3]. However, it is possible to prevent cache lines from being evicted by locking them. While this allows for performance optimization, it also gives rise to the cache incoherence which CacheKit exploits to hide from memory inspection tools. After the virtual address and physical address translation has occurred in the Translation Lookaside Buffer, cache stores the memory data of the physical address and keeps the NS bit to indicate the security state needed to access the entry. This NS bit is set by hardware and it is not directly accessible by system software [2].

3 BACKGROUND: CACHEKIT ATTACK

CacheKit is a new type of rootkit that exploits a cache incoherence and cacke locking mechanisms in the ARM TrustZone architecture [27]. Dividing the two worlds completely results in a cache incoherence where the contents of Secure and Normal World are different even if they map to the same physical address. CacheKit exploits the fact that even though the Secure World can access the memory of Normal World, the two worlds are separated such that they cannot gain access to the other's cache. There are three major steps in establishing CacheKit: Loading, Locking, and Hiding. First, a technique known as Cache-as-RAM is used to ensure that the rootkit is loaded only into cache of the Normal World where it can avoid detection from the Secure World [20, 28]. Then, the ARM hardware support settings are exploited to keep the code persistent in cache as long as possible. Finally, the translation tables are modified such that the malicious code in cache maps to unused I/O addresses in physical memory so that if cache content is flushed to RAM for inspection, the data is simply lost. This ensures that even if the rootkit were to be flushed into memory for any reason, any trace of the malicious code would be lost [27].

3.1 Loading

In the ARM architecture, the only way to read or write a cache line is to have the processor read from or write to virtual memory. Therefore, to load the rootkit into cache the CacheKit module must first enable caching on memory. This is done by setting the paging table memory attribute as WriteBack. The WriteBack configuration ensures that load register (LDR) instructions trigger a cache line fill. The key of cache loading is to ensure that data is loaded to cache and cache only, the rootkit should not be stored in RAM at any point as it would make it visible to memory inspection.

3.2 Locking

The Cortex-A8 processor allows system software to lock up to seven cache ways out of the total eight ways. This ability to maintain the code persistently in cache is essential as cache is very volatile. Once the rootkit is loaded, it should remain in memory as long as possible so an attacker can maintain control over the infected system. While using Cache-as-RAM enables us to better hide the malware, cache locking allows the code to survive long enough to make it useful. First, the cache corresponding to all the memory addresses to be locked in cache will need to be flushed out. Second, the cache way to be used is unlocked and the other ways are locked. This means that any cache fills made by the LDR and STR instructions will be made to the way that has been designated for the rootkit. Lastly, once the rootkit has been loaded into cache the way that has been reserved for the rootkit is locked and the rest of the cache is unlocked. Consequently, the implementation of Cache-as-RAM to store the rootkit gives it exceptional stealth as it can no longer be detected by memory inspection tools. However, there is a trade-off in storing the rootkit in cache. Cache is very volatile and even though cache lines can be locked with hardware control, they will still respond to cache maintenance instructions if they are called. Therefore, when a cache flush is called the contents of the rootkit will indeed be written out to memory.

3.3 Hiding

Having successfully loaded and stored the rootkit in cache, CacheKit must now address the two main issues of remaining concealed. The first problem is that even when locked, the cache lines are still responsive to cache maintenance instructions. As stated, ARMv7 provides various cache maintenance instructions that can cause the contents of the cache to be written out to memory. This is dangerous because if the malicious code is flushed to memory then it becomes detectable to memory inspection techniques. Similarly, the second problem has to do with writing back to memory and introspection from the normal world kernel. Detection methods that sequentially map each physical page into the kernel memory space would still be able to read the cache. Therefore, the idea of mapping the cache lines to unused physical I/O address space is proposed to resolve these issues with direct cache locking. While mapping the cache to I/O address space does mean that an attacker can maintain stealth by destroying the data since no backup memory exists, it also means that they lose the rootkit if it is ever flushed from cache. Since the effectiveness of a rootkit depends on how long it can remain hidden in the system this is a major trade-off between stealth and persistence. Therefore, for CacheKit to be most effective it is imperative to have an environment where data can persist in cache.

4 DEFEATING CACHEKIT ATTACKS: NAÏVE APPROACHES

In this section, we discuss several possible solutions to defeat CacheKit attacks.

4.1 Naïve Prevention

First, we focus on preventing rootkits from hiding in Normal World cache. According to the ARM Technical reference manual for the Cortex-A8 [1], the CL bit determines whether or not cache locking can be done in the Normal World. If CL is set (to 1) then cache locking is available in both worlds. However, if the CL bit is cleared (to 0) then cache locking is only available in the Secure World. Attempting to use registers associated with cache locking, namely the `L2 Lock Down Register`, results in an undefined instruction exception. Therefore, this provides a simple and direct solution to preventing CacheKit attacks by disabling the cache locking capabilities in the Normal World and only allowing trusted applications running in privileged SW levels to lock cache entries. While this would prevent such attacks entirely because the rootkits would no longer be able to remain persistent in cache, it also limits a lot of what NW users can accomplish in their applications.

Therefore, disabling cache locking may not be an option. Many modern processors feature cache locking mechanisms to allow for finer control of cache eviction policies and thereby improving the cache hit rate. This can have a considerable impact on the performance of a system if managed correctly [19]. For this reason, a wide variety of processors across different manufacturers and processor families offer this option. Some of ARM's Cortex and ARM11 processors allow for way locking in which locking is available at the granularity of ways of a set-associative cache. Line locking is a more fine-grained approach where it is possible to have a different number of locked lines in different sets of the cache; Intel's XScale, the ARM9 family, and BlackFin 5xx family processors support this kind of locking mechanism [18].

4.2 Naïve Detection

Given that it is most likely a part of a real-time or general embedded system that depends on cache locking to give critical processes timing predictability or the ability to meet stringent performance requirements, we cannot move this functionality to the Secure World. Not only would it affect the timing with the need for world switching but it would also greatly broaden the Trusted Code Base, defeating the purpose of having a small, isolated, trusted environment. Therefore, it is necessary to enable cache locking in the Normal World and thus we must be able to detect malicious attempts by a potentially compromised Normal World OS.

Originally, we had considered that in order to detect an attack, we would need not only the attempt to lock the cache but also the remapping of a virtual address to a physical address that does not correspond to system memory. That is, these two fundamental necessities of CacheKit would be enough to determine that a malicious user is attempting to hide code in the Normal World cache. CacheKit prevents its data from being retrieved for forensic

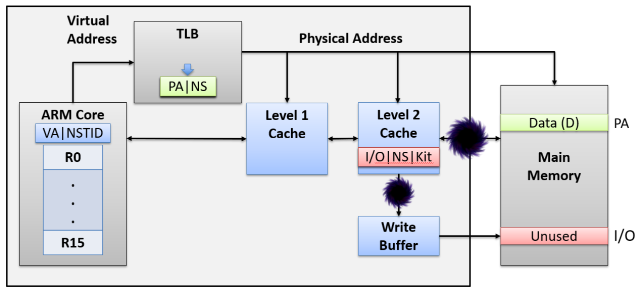

Figure 1: PIPT Caches prevent the retrieval of hidden data

purposes by mapping to the physical region of reserved unused I/O so that if the cache is ever flushed, the malicious code is simply lost and cannot be analyzed. However, these two functionalities together make for a very specific signature that can most definitely signal malicious activity. This is because while some time critical programs may need instructions or data to be locked in the cache, we would never need anything outside the System Memory Region to be locked.

The first step is to detect if Normal World has attempted to lock the cache by reading the value of the `L2 Lock Down Register`. Anything other than 0 would indicate that one or more cache ways are locked. If we are able to then find a page table entry that maps a virtual address to a physical address outside of the System Memory Region, we would be able to detect a rootkit hiding in the cache. This may require a page-table walk, because we have to check all entries to ensure completeness. While a page-table walk is expensive, we would only need to execute it when a lock is detected. In this manner, we would then be able to force a cache flush or invalidation to trigger a write back to memory. To ensure that we are able to retrieve the malicious code we also have to modify the virtual address that has an invalid mapping to a new, valid memory region that we have reserved from Secure World for memory inspection. However, we find that this approach will be ineffective in retrieving the malicious data in the cache because even if we are able to re-map the virtual address to a valid physical address, the L2 cache in the ARM Architecture is Physically Indexed, Physically Tagged. This means that once the re-mapped address has been loaded into the cache, it is already too late to retrieve the data. This approach would only work with a processor that supports both cache locking and VIVT caches. In Cortex-A8, the L2 cache is Physically Indexed, Physically Tagged (PIPT) [1]. This means that the tag and index bits stored in the cache to identify which memory address the cache line maps to are taken from the physical address and not the virtual address.

As shown in Figure 1, even if we are successful in re-mapping the virtual address to a valid physical address, once we flush the cache the write to memory will be done according to the tag and index bits stored with the respective cache lines. Therefore, once the rootkit has been stored in the cache with the modified physical address that points to reserved, unused I/O, any eviction will always be done to this location. CacheKit creates a "black-hole" situation where once it is in the cache, there is no retrieving it because there is no way to access the tag and index bits of the cache other than with perhaps a JTAG Debugger. Therefore, the only solution left is to develop a defense mechanism to prevent rootkits, or any malicious code, from being loaded into the cache in the first place. This makes a detection solution considerably difficult. However, by leveraging TrustZone technology, prevention is still possible.

5 DEFEATING CACHEKIT ATTACKS: CACHELIGHT

In this section, we present CacheLight, a defense mechanism for preventing malicious software from inhabiting the cache for any significant period of time. We leverage the TrustZone Security Extensions to ensure that even if an attacker is able to briefly compromise the Normal World kernel, they will not be able to leave behind any malicious code in the Normal World cache. The main idea is to expand on the most simple solution of disabling the ability to lock the cache by clearing the CL bit. However, we still want user programs that have time-sensitive operations in real-time and embedded systems to be able to make use of the cache locking mechanism. Therefore, we implement CacheLight as a Secure World kernel function that can be called from Normal World kernel to have the Secure OS perform the locking for the Non-Secure OS. This means that while users can still take advantage of the benefits of cache locking for legitimate purposes, the Secure OS can have better control over what gets locked in the cache and perform integrity checking to prevent malicious code from being loaded in. This is implemented as an SMC which the Normal World can call and the Secure World can then handle.

5.1 Workflow

The first step is to determine what arguments to pass with the SMC to Secure World. In order to service the call, the SMC Handler in Secure World must know:

1. OP Code: The operation code to determine what the SMC is for and properly handle the request.
2. VA: The virtual address that corresponds to the data that the NW wants to lock in cache.
3. LockDownReg: the `L2 Lockdown Register` value to be used. This determines which cache ways the Normal World wishes to lock.
4. Size: the size or amount of data to be loaded into the cache from the base address (VA).

With these four arguments we have all the information we need to load and lock the data into the cache on behalf of the Normal World. More importantly, the data must first be loaded into physical memory by the Normal World, which ensures that whatever is being loaded into the cache is consistent with what appears in RAM and therefore exposed to memory inspection tools.

Once we have established the necessary parameters to pass to Secure World, we can detail the necessary work-flow, or steps, of the approach as shown in Figure 2. Secure World has all the necessary information to first verify that the request to lock memory is valid by ensuring only memory regions allocated to Normal World are being requested. Furthermore, all the data must first be loaded into main memory by the Normal World, and is therefore available to inspection, should Secure World find an invalid PA is given. Otherwise, it can then go through the process of servicing legitimate

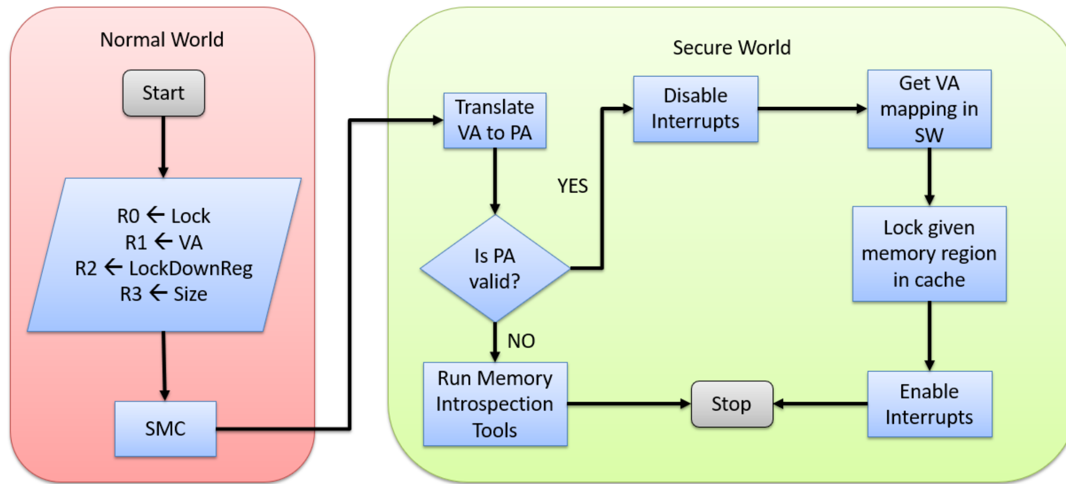

Figure 2: Overall Workflow of CacheLight

requests and locking the requested memory in cache. The steps for this case are detailed in the following sections.

5.2 Virtual to Physical Address Translation

Virtual address translation is provided by way of the Memory Management Unit. In ARMv7 the MMU works with the L1 and L2 memory system to translate virtual addresses to physical addresses and controls accesses to and from external memory. Therefore, it is enhanced with security extensions and multiprocessor extensions to provide address translation and access permission checks. Virtual-to-physical address mappings and memory attributes are defined in main memory as page tables; each World has its own set of page tables and TLB identifiers to remove the requirement for context switch TLB flushes [1].

Therefore, the first step in preventing rootkits from being loaded into cache is to translate the given VA to the PA. The ARMv7 Architecture provides VA to PA address translation operations using registers from `CP15`. There is a `Physical Address Register (PAR)` that holds the PA after a successful translation or the source of the abort after an unsuccessful translation. It is a read/write register banked in Secure and Non-secure states and accessible in privileged modes only [1]. The `PAR` bits [31:12] contain the physical address after a successful translation. Bits [11:0] are taken from the virtual address to obtain the complete PA. To make use of this register for our address translation we do VA to PA translation in the other Secure or Non-secure state. The purpose of the VA to PA translation in the other Secure or Non-secure state is to translate the address with the current virtual mapping in the Non-secure state while the core is in the Secure state [1]. To access the VA to PA translation in the other Secure or Non-secure state, we write CP15 c7 with Opcode2 set to 4 for privileged read permission. In this manner, the given virtual address will be translated using the Normal World's translation tables and we can retrieve a PA to work with.

5.3 Verifying Memory Contents

Once we have the physical address we are able to determine it is a valid address within the system memory. We can check if the PA maps to I/O memory region, in which case we know we should not allow such an address to be locked in the cache. Depending on the implementation, it is also possible to check that the PA is within the memory region allocated to the Normal World by the Secure World for a much better, tighter security check. At this point we can flush the cache and run the memory inspection tools of the Secure World to retrieve and analyze the malicious code that the attacker attempted to leave behind. Since in this scheme only the Secure World is allowed to perform cache locking, we are able to ensure that any memory that will remain in the cache does not differ to what is in RAM and thus we can be certain that there is no incoherence from what we are inspecting in the RAM and what the system is executing in the cache.

On the other hand, if the request is for a valid physical address then the secure world must be able to perform the loading and locking of the requested memory region into the specified cache way.

5.3.1 *Enabling and Disabling Interrupts.*

The first step is to disable interrupts so that we cannot be pre-empted. To lock the given memory, and only the given memory, in the given cache way(s) the locking process must not be interrupted. Therefore it is imperative that this code be within a non-preemtible critical section. Once the locking is complete, Secure World can enable interrupts again.

5.4 Mapping Normal World Memory to Secure World

With the virtualization extensions active in both worlds, it is necessary to map a virtual address in the Secure World to the given physical address from the Normal World. This means that the Normal World memory must be mapped in the Secure World as well. Memory that both worlds have access to is referred to as World-Shared Memory. World-Shared Memory is designated as non-secure

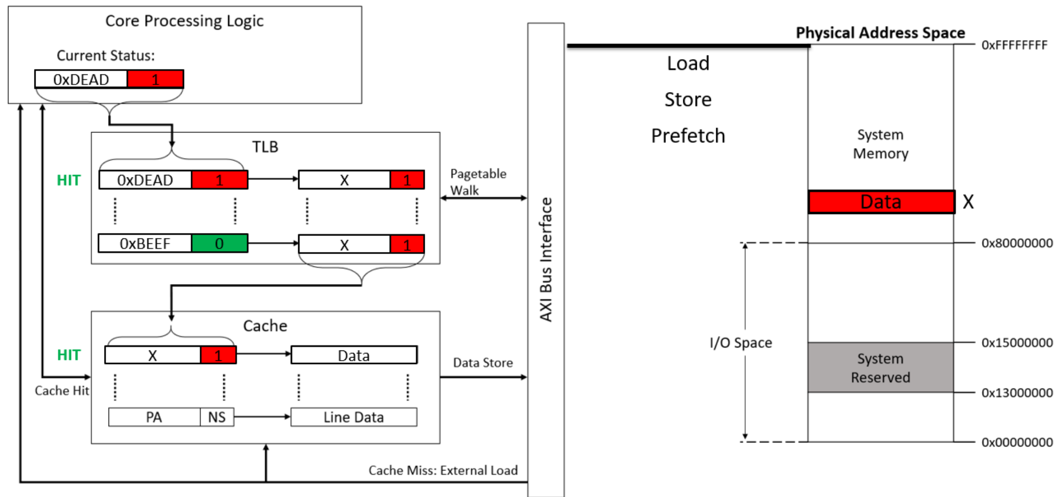

Figure 3: Memory System for an ARM Core

since NW can only make non-secure accesses while SW can make both secure and non-secure accesses. Now that valid VA entries exist for Normal World physical addresses, Secure World can use these to load the requested memory region and thereby create the cache line entires on behalf of the Normal World. Given the PA that the Normal World would like to lock, a helper function is defined to build the offset within the physical Normal World memory and add that offset to the virtual start address of the mapping within the Secure World kernel, returning the Secure World VA for the specified PA. Since memory is shared, this VA can be used to create cache entries tagged as Non-Secure that will generate cache hits when the Normal World attempts to access said memory, effectively allowing the Secure World to create locked cache entries for the Normal World.

5.5 World-Shared Memory

As mentioned in Section 2, all addresses are tagged with an additional NS bit to indicate what mode the entry belongs to. Normal World is only able to access entries tagged with an NS bit equal to 1. Therefore, when we load the data into the cache we have to ensure that it is tagged with NS = 1 so that Normal World will have a cache hit when it tries to do a look-up in the TLB. Figure 3 shows how the memory system of a theoretical ARM processor might handle the state associated with Security Extensions when accessing the memory system [2].

First, the core processing logic attempts a data load, a data store, or an instruction prefetch. The hardware passes the Virtual Address (VA) and the current world (Non-Secure Table Identifier, or NSTID) to the TLB to enable it to perform address translation. The NSTID is passed by the hardware and depends on the current state of the processor, not the NS bit. In this case the NSTID will represent a secure world entry [2].

Then, the TLB loads the physical address and the NS-bit associated with the VA and NSTID it was passed, performing a page-table walk and forcing NS=1 if NSTID=1 if necessary. The translation tables are responsible for keeping track of the NS bit and whether the VA resolves to a secure or non-secure physical address [3]. The TLB then passes this information to the cache to perform the actual data or instruction access. If the Secure World has mapped the non-secure memory containing the data the Normal World wants to lock in its translation tables, then the Secure World can directly access the non-secure cache lines. Therefore, a Normal World application can pass data to the Secure World though any level in the cache hierarchy. This enables a high performance system in comparison to solutions that require cached data to be flushed out of the cache and in to external memory [2].

Therefore, when the loading and locking of memory is performed, the resulting cache entires will be tagged with an NS bit of 1 rather than zero. This means that when the Normal World attempts to access this address that it requested Secure World to load into the cache for it, it will generate a cache hit. Normal World will find a cache entry that matches the PA translation with an NS bit of 1, indicating that this memory is designated for Normal World use and will return the data from that cache line. This also means that Secure World can gain access to the Normal World cache if it creates the necessary translation table entries by implementing World-Shared Memory.

5.6 Locking NW Memory Into Cache From SW

Finally, we can now perform the loading and locking of the memory in the Secure World exactly as it would have been done in the Normal World. We pass the virtual address, the value to use for the `L2 Lockdown Register`, and the size of the data to load as arguments. Therefore, Secure World has all the information required to load and lock data to the requested cache ways. By implementing World-Shared Memory, the Secure World can issue non-secure memory accesses that create non-secure cache entires for the Normal World. When control is returned to the Normal World, the legitimate process will be able to access the time-critical section locked in the cache as intended. Therefore, this mechanism gives

Table 1: Comparison of Different Defense Approaches.

Approach	Category	Pros	Cons
CacheKit Detection	Detection	- Flag potential malware and flush the cache	- Cannot verify existence of rootkit - Cannot retrieve code for analysis or forensic purposes
Disabling Cache Locking in NW	Prevention	- Simple - Direct - No Overhead - Done upon system initialization	- Cache locking mechanisms not available to normal user - Processor might have been selected for those specific features
CacheLight	Prevention	- Lightweight - Overhead in set-up, not execution time - Prevents loss of malicious code for analysis - Control over what is locked in cached - NW supports cache locking	- Additional world-switch overhead in set-up time - Can still load malicious code into cache

the Secure World a lot of power and control over what gets loaded into the cache, allowing it to ensure that there is no incoherency with what is stored in memory and what is stored in Normal World cache.

5.7 Comparing Approaches

To the best of our knowledge, there have been no other defense mechanisms proposed against these kinds of attacks. Therefore, to evaluate the best solution, we provide a comparison of the three approaches introduced here in Table 1. We conclude that, while there is room for further work, the most robust and effective solution is the CacheLight approach. In the following chapters we implement and evaluate CacheLight in terms of security and performance.

6 CACHELIGHT IMPLEMENTATION

6.1 Genode: A Secure World OS

To defend against this new type of attack we first replicate the original prototype scenario on the i.MX53 Development Board. Once functionality was verified, we moved to deploy the defense environment. The first step is to have a Trusted Execution Environment (TEE) in the Secure World. After careful consideration we decided to use the open-source Genode Project which supports the i.MX53. Genode provides support for various Normal World operating systems, hardware platforms, and scenarios. The scenario we choose was designed for the i.MX53 with TrustZone hardware capabilities. The Linux kernel was modified such that certain kernel functions could handle world switches using an SMC call to the Secure World. The CacheKit attack module can then be cross-compiled for the new target kernel in the Normal World and deployed on the new Genode environment [16].

6.2 Building and Deploying The Environment

Using Genode 18.02, we build the Genode TrustZone Virtual Machine Monitor (TZ-VMM) scenario for the i.MX53 board. This build generates a bootable image for the target board as specified in the configuration file. The next step is to obtain the necessary bootloader provided by Genode Labs and generate the boot image. Finally, a bootable SD card is prepared for the board using both the boot image and the Genode image generated earlier. Now we have the Genode TZ-VMM running on the i.MX53 development board. The details of using and understanding Genode and all of the different scenarios can be found in their Foundations Book [12].

6.3 Deploying the CacheKit Attack

We replicated CacheKit on the i.MX53 development board, which features a single ARM Cortex A8 processor. The project is implemented as a kernel module that once installed on the board can use the cache manipulation tools in ARM to successfully load, lock, and conceal the rootkit in cache. Next, we must be able to cross-compile the CacheKit attack module against the new modified target kernel running in the Normal World. Unfortunately, the modified kernel is provided as a pre-compiled image in the Genode repository. Therefore, we obtain the source code of the modified kernel and create our own build. This then gave us a target to cross-compile the CacheKit attack module against to be able to deploy it on the new environment.

6.4 Deploying the CacheLight Defense

The Normal World kernel was modified such that certain kernel functions could handle world switches using a Secure Monitor Call (SMC) to the Secure World. According to their documentation, there were six different SMC's added to the Linux kernel [15]. However, upon closer inspection of the current release, Genode 18.02, there are actually only 4 SMC's that are implemented. Furthermore, they do not follow the ARM SMC Calling Convention [5], but rather implement their own, simpler version of handling SMC's. In Genode, an SMC is still used to generate a synchronous exception that is handled by Secure Monitor code running in Exception Level 3 (EL3). The Secure Monitor Mode handles the world switching and then hands off the handling of the exception to the Virtual Machine Monitor running as an unprivileged user-level component in Genode. However, instead of using the ARM Calling Convention, since there are only four different calls implemented, the function to perform is simply passed as an argument in register `R0`. Registers `R1-R10` are then used to pass any necessary arguments to the Secure World. Upon a world-switch the Monitor will save the state of the Virtual Machine, including these registers, so that the Secure World can have access to them. The VMM then simply checks the value in R0 to determine what function the Normal World needs it to perform for the SMC and expects any relevant arguments in the

Table 2: SMC Implementation in Modified Linux

SMC Number	Function
0	FRAMEBUFFER
1	INPUT
2	SERIAL
3	BLOCK
4	LOCK

other registers [15]. This makes it fairly simple for us to add and define our own, fifth, SMC for a request to Secure World to lock memory in the cache. The table below shows the different SMC's that Genode had already implemented with the addition of our own "LOCK" SMC with code 4.

As shown in Table 2 our SMC has a function code of 4, this is passed through register `R0` and is used to determine what function the Secure World needs to perform and the respective arguments in the remaining registers. In our case, we pass the following arguments:

(1) R1 is the virtual address to be loaded and locked in the cache.
(2) R2 contains the lock down register value to be used. This determines which cache ways the Normal World wishes to lock
(3) R3 the size or amount of data to be loaded into the cache from the base address in R1

With these three arguments we have all the information we need to load and lock the data into the cache on behalf of the Normal World. More importantly, we have the location of the data before it gets stored in the cache which allows us to regulate what can be locked into the Normal World cache. Once the request is handled by Genode, it restores the VM's state and switches back to Normal World operations.

After we have implemented an SMC for CacheLight in the Genode core, we need to define what happens once it is called by the Linux kernel. The SMC would be called whenever the Normal World has some code that needs to be locked in the cache. The Normal World would load the operation code, virtual address, value for the lock down register, and size of data to be loaded into the respective registers and perform an SMC. Once Genode reaches the exception handling for the call, we need to be able to create an entry in the cache with the data requested by the Normal World that is accessible by Normal World. Overall, the entire process is implemented as roughly 200 lines of C and assembly code in the Secure World kernel which is a minimal increase in the Trusted Code Base (TCB) of the Secure World.

Genode utilizes the ARM Virtualization Extensions for the TZ-VMM scenario [17]. Therefore, to establish World-Shared Memory, the Genode bootstrap process is modified to map the RAM region that is later allocated to the Normal World VM to a designated virtual address space in the Secure World. Once defined, the mapping can be added in the Platform constructor of the bootstrap process to create the page table entries in the Secure World for the Normal World RAM space, effectively creating World-Shared Memory. The entries are created with the necessary flags, among them the NS bit that indicates the virtual addresses resolve to non-secure physical space.

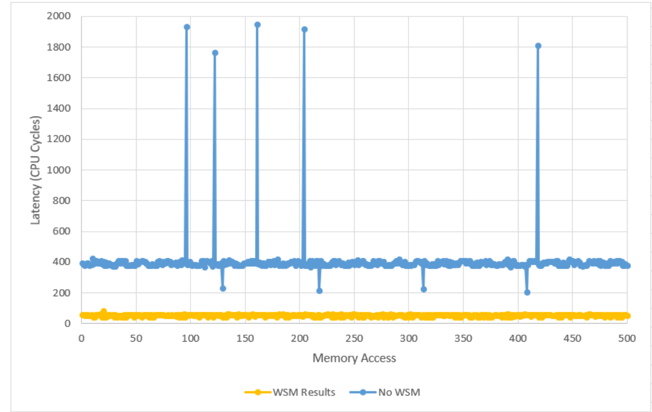

Figure 4: World-Shared Memory Allows Secure World to Create Non-Secure Cache Entries for Normal World

Once we deploy CacheLight, we are able to detect whenever the CacheKit module is attempting to lock malicious code into the cache. By doing security checks on the PA being locked into cache, CacheLight can detect anomalous behavior and if needed deploy inspection tools. If the attempt to lock the cache is malicious, it can then flush the caches and run memory introspection tools to determine the nature of the attack and retrieve any relevant data for forensic analysis. On the other hand, if the request is determined to be legitimate, CacheLight can service it by taking advantage of world-shared memory.

7 EVALUATION

In this section we verify the ability of the two worlds to share data through any level of the cache hierarchy using World-Shared memory as well as the performance impact of CacheLight.

7.1 Effects of World-Shared Memory

For CacheLight to work, Secure World must be able to create cache entries tagged with a non-secure NS bit that Normal World can then access. While Normal World is forced by hardware to only make non-secure memory accesses, Secure World has the option to perform both secure and non-secure accesses. The type of access made is determined by the NS bit in the first-level Page table descriptor [3]. This NS bit defines the physical address space, Secure or Non-secure, for all of the Large pages and Small pages of memory described by that table. For accesses from Secure state, it determines whether the access is to Secure or Non-secure memory. However, it is ignored by accesses from Non-secure state [2]. In a cacheable secure memory access, the linefill is requested using secure access and the data is tagged in the cache as secure data. In a cacheable non-secure memory access, the linefill is requested using non-secure access and if no security error is received, the data is tagged in cache as non-secure data. Therefore, the NS bit that is used to tag the cache records the type of memory access that was made, allowing Secure World to cause allocation of non-secure lines into the cache.

Figure 4 shows the effects of implementing World-Shared Memory to be able to create and lock non-secure cache lines from the

Table 3: Performance Impact of CacheLight When Locking One Cache Way

	CPU Cycles	Nanoseconds	Percent of Total
Total Time	1,930,060	1,930.06	100
Locking Time	654,733	654.73	33.92
World-Switch Overhead	1,272,500	1,272.50	65.93
CacheLight Overhead	2,827	2.83	0.15

Secure World. As shown, if there is no WSM then Secure World generates a secure memory access and the cache entries created and locked are tagged as secure. If the Normal World attempts to access that memory, it gets a miss in the cache and has to load that same address but with a non-secure tag which it would be unable to lock in the cache. However, with WSM, the Secure World can create and lock cache entries tagged as non-secure. Then, when the Normal World attempts to access that memory, it will get a hit in the cache and can use the locked entires. Therefore, it is possible to restrict cache locking capabilities to the Secure World for security measures while still allowing legitimate Normal World programs to request locked non-secure cache entries from Secure World.

7.2 Performance Evaluation

To defend against malware that attempts to evade memory introspection tools, CacheLight incurs some overhead in the cache locking process. In a typical system, a Normal World privileged program can directly perform the locking, however, with CacheLight there are additional security checks and World-Switches that must occur. Therefore, to measure the performance impact of CacheLight the approach is split into three parts. First, there is the essential and necessary locking and loading of the cache which is code that runs regardless of whether CacheLight is implemented or not. In addition, there is the overhead of the security checks and address translations done by CacheLight. Finally, that leaves the overhead of switching to and from Secure World to lock the cache. Timing analysis of these three parts is performed using the available ARM Performance Monitoring Unit. Since, the overhead portions are constant with respect to the size of the data being locked, analysis is performed for locking a single cache way of size 32 kB.

As shown in Table 3, the World-Switch to and from Secure World is the main source of overhead for CacheLight. In this experiment, with the i.MX53 running Genode, it makes up almost 66% of the time taken for CacheLight to process a valid lock request for a single cache way. However, this is unavoidable as switching between worlds is necessary for the implementation. This can only be addressed by how the ARM architecture handles a World-Switch or the efficiency of the Genode monitor code for switching between worlds. Furthermore, while it is considerable, it is a one-time cost of setting up the locked cache lines and therefore should occur scarcely. On the other hand, the actual overhead of the security measures of CacheLight is minimal.

8 RELATED WORK

Overtime, rootkits and the methods for detecting and defending against them have been in an evolutionary race of hide and seek. Originally, persistent rootkits needed to modify nonvolatile storage to survive system power cycles which meant that file integrity checking tools could effectively detect them [14]. Therefore, rootkits switched to reside only in the operating system kernel memory to defeat this storage-based detection. To detect this new type of rootkit, defenders acquire the system memory using a dedicated secure coprocessor [21] or physical hardware [8]. In their search to acquire higher root privileges, attackers have developed several different rootkits. Virtual machine based rootkits (VMBR) insert a customized malicious hypervisor beneath the currently running operating system [22]. Firmware based rootkits infect the firmware on I/O devices [13] or the system BIOS [11].

As their counterpart, new hardware and software rootkit detectors with higher privilege are also proposed. Hypervisors are commonly used to introspect the untrusted operating system [6]. However, vulnerabilities are frequently found in the hypervisors. This gave rise to the use of hardware features, such as security extensions in various processors. In this paper, we cover ARM TrustZone, however, each developer has their own set of hardware security extensions. For example, AMD has SVM [10] and Intel provides TXT [25]. These hardware features provide a TEE with guaranteed isolation and the highest privilege that defenders can claim as their own. CacheKit seeks to exploit weaknesses in these new hardware defense features to evade detection by the OS kernel. Another proposed design that has much the same goals is Shadow Walker, as it exploits the I-TLB and D-TLB coherency problem in the Intel architecture to hide the rootkits [23]. There is also an implementation called Cloaker that can hide its presence by locking the page translation it altered in the translation look-aside buffer [9]. However, none of these have ever been able to avoid the last-level defense of memory inspection. Thus, we note that CacheKit provides a new level of stealth as it is able to evade physical memory inspection by hiding in the Normal World cache.

CacheKit is the only one able to evade the physical memory check. Furthermore, by mapping the cache lines to unused I/O space we ensure that the malicious code cannot be examined. In the worst case, it is evicted and then destroyed before it can be inspected and analyzed. To the best of our knowledge, there have been no proposed defenses against this new kind of attack, leaving many ARM systems vulnerable and exposed.

9 FUTURE WORK

While this initial implementation of CacheLight shows great potential, there are further applications and checks for which CacheKit could be leveraged to tighten security. For example, limiting or regulating how much of the cache a request can lock. Furthermore, there is the case where a rootkit is loaded, but not locked, in the cache with a remapped address so the attempt fails. However, this attempt would not be retrievable for analysis, therefore, in such cases valuable forensic data is lost about the attack and attacker. Perhaps there could be work done to expand the robustness of the approach in such cases. Finally, there is still more work that

can be done on the response after an attack is detected and verified. Looking deeper into the what inspection tools to run in case CacheLight flags potential malicious code and the appropriate response mechanisms to employ if the tool does find malicious code in memory.

10 CONCLUSION

In this paper we present CacheLight, a lightweight approach to preventing malicious use of cache locking mechanisms while allowing time-critical applications to legitimately utilize them to ensure execution times in embedded and real-time systems. CacheLight allows the Normal World to perform cache locking through requesting it as a service from the Secure World. All that is needed is a minimal increase in the Trusted Code Base to handle a new SMC, which the OS running in the TEE can then validate and verify to prevent any malicious code from being hidden in the cache.

Upon world switch, the Secure World can now handle and verify the validity of any cache lock request to ensure that any data that will persist in the cache not only maps to a valid address in memory but is also consistent with what is present in main memory; effectively bringing the contents of the cache to light. Additionally, because the Secure World does not hand control back to Normal World after verifying the address, but rather performs the loading and locking on behalf of the Normal World, the attacker cannot bypass the security checks by passing different addresses in the arguments. Should CacheLight find that the attempt to lock the cache is malicious, it can then flush the caches and run memory introspection tools to determine the nature of the attack and retrieve any relevant data for forensic analysis. On the other hand, if the request is determined to be legitimate, CacheLight can service it by taking advantage of World-Shared Memory.

Therefore, CacheLight can successfully prevent malicious code from hiding from SW introspection tools in the NW cache for any significant amount of time. Additionally, while we present a solution for the ARM architecture, the approach can be generalized to any architecture that employs the same execution separation idea. If the attack can be modified to a new architecture, then so can the defense. Moreover, CacheLight incurs the overhead of a world-switch for the set-up of the time-critical data. However, the initial setup of locking data in the cache is already expected to be expensive so that the performance and timing requirements can be met once the setup is done and the application running. CacheLight makes additional overhead to the setup process but not the execution of the time-critical process that requested the lock. Given that it provides security against an otherwise undetectable attack, the trade-off in setup time is extremely worthwhile.

REFERENCES

[1] ARM. 2006-2010. ARM Cortex-A8 Technical Reference Manual. http://infocenter.arm.com/help/topic/com.arm.doc.ddi0344k/DDI0344K_cortex_a8_r3p2_trm.pdf. (2006-2010).
[2] ARM. 2009. ARM Security Technology Building a Secure System using TrustZone Technology. http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.prd29-genc-009492c/index.html. (2009).
[3] ARM. 2012. ARM Architecture Reference Manual, ARMv7-A and ARMv7-R edition. http://infocenter.arm.com/help/index.jsp?topic=/com.arm.doc.ddi0406c/index.html. (2012).
[4] ARM. 2015. ARM Cortex-A Series Programmerś Guide for ARMv8-A. http://infocenter.arm.com/help/topic/com.arm.doc.den0024a/index.html. (2015).
[5] ARM. 2016. SMC CALLING CONVENTION System Software on ARM Platforms. http://infocenter.arm.com/help/topic/com.arm.doc.den0028b/ARM_DEN0028B_SMC_Calling_Convention.pdf. (2016).
[6] Ahmed M. Azab, Peng Ning, and Emre C. Sezer. 2009. A hypervisorbased integrity measurement agent. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*. ACM, Honolulu, Hawaii, 461–470.
[7] Ahmed M Azab, Peng Ning, Jitesh Shah, Quan Chen, Rohan Bhutkar, Guruprasad Ganesh, Jia Ma, and Wenbo Shen. 2014. Hypervision Across Worlds: Real-time Kernel Protection from the ARM TrustZone Secure World. In *Proceedings of the 21st ACM Conference on Computer and Communications Security (CCS)*. ACM, Scottsdale, Arizona, 90–102.
[8] Ellick Chan, Shivaram Venkataraman, Francis David, Amey Chaugule, and Roy Campbell. 2010. Forenscope: A framework for live forensics. In *Proceedings of the Annual Computer Security Applications Conference (ACSAC)*. ACM, Austin, Texas, 307–316.
[9] Francis M. David, Ellick M. Chan, Jefferty C. Carlyle, and Roy H. Campbell. 2008. Cloaker: Hardware supported rootkit concealment. In *Proceedings of the 29th IEEE Symposium on Security and Privacy (Oakland)*. IEEE, Oakland, CA, 296–310.
[10] Advanced Micro Devices. 2013. *Amd64 Architecture Programmer's Manual*. AMD.
[11] Shawn Embleton, Sherri Sparks, and Cliff Zou. 2013. Smm rootkit: a new breed of os independent malware. Security and Communication Networks. (2013).
[12] Norman Fenske. 2017. Genode Operating System Framework Foundations. (2017).
[13] John Heasman. 2007. Implementing and detecting a pci rootkit. BLackhat DC. (2007).
[14] Gene H. Kim and Eugene H. Spafford. 1994. The design and implementation of tripwire: A file system integrity checker. In *Proceedings of the 2nd ACM Conference on Computer and Communications Security (CCS)*. ACM, Fairfax, VA, 18–29.
[15] Genode Labs. 2017. An Exploration of ARM TrustZone Technology. Genode OS Documentation and Articles. (2017). https://genode.org/documentation/articles/trustzone
[16] Genode Labs. 2017. Genode: Operating System Framework. https://github.com/genodelabs/genode. (2017).
[17] Genode Labs. 2017. An in-depth look into the ARM virtualization extensions. Genode OS Documentation and Articles. (2017). https://genode.org/documentation/articles/arm_virtualization
[18] Y. Liang, T. Mitra, and L. Ju. 2015. Instruction Cache Locking Using Temporal Reuse Profile. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems* 34, 9 (Sept 2015), 1387–1400. https://doi.org/10.1109/TCAD.2015.2418320
[19] Tiantian Liu, Minming Li, and Chun Jason Xue. 2012. Instruction Cache Locking for Embedded Systems using Probability Profile. *Journal of Signal Processing Systems* 69, 2 (01 Nov 2012), 173–188.
[20] Y. Lu, L. Lo, G. R. Watson, and R. G. Minnich. 2006. CAR: Using Cache as RAM in LinuxBIOS. http://rere.qmqm.pl/âĹijmirq/cacheasramlb09142006.pdf. (2006).
[21] Nick L. Petroni, Timothy Fraser, Jesus Molina, and William A. Arbaugh. 2004. Copilot - a coprocessor-based kernel runtime integrity monitor. In *Proceedings of the 13th USENIX Security Symposium (Security)*. USENIX, San Diego, CA, 179–194.
[22] Joanna Rutkowska. 2006. Subverting vistatm kernel for fun and profit. Black Hat Briefings. (2006).
[23] Sherri Sparks and Jamie Butler. 2005. Shadow walker: Raising the bar for rootkit detection. *Black Hat Japan* (01 2005), 504–533.
[24] H. Sun, K. Sun, Y. Wang, J. Jing, and S. Jajodia. 2014. Trustdump: Reliable memory acquisition on smartphones. In *Proceedings of the 19th European Symposium on Research in Computer Security (ESORICS)*. Springer, Wroclaw, Poland, 202–218.
[25] Intel Trusted Execution Technology. 2016. *Intel Software Development Guide*. Intel.
[26] F. Zhang, J. Wang, K. Sun, and A. Stavrou. 2014. HyperCheck: A Hardware-AssistedIntegrity Monitor. *IEEE Transactions on Dependable and Secure Computing* 11, 4 (July 2014), 332–344.
[27] Ning Zhang, He Sun, Kun Sun, Wenjing Lou, and Y Thomas Hou. 2016. CacheKit: Evading memory introspection using cache incoherence. In *Proceedings of the 1st IEEE European Symposium on Security and Privacy*. IEEE, Saarbrücken, GERMANY, 337–352.
[28] V. J. Zimmer, M. A. Rothman, and S. M. Datta. 2004. Using a processor cache as ram during platform initialization. US Patent 20,040,103,272.. (2004).