

Yet Another MicroArchitectural Attack: Exploiting I-cache

Onur Aciicmez

Samsung Information Systems America, Samsung Electronics
95 West Plumeria Drive, San Jose, CA 95134, USA
onur.aciicmez@gmail.com,

Abstract. MicroArchitectural Attacks (MA), which can be considered as a special form of Side-Channel Analysis, exploit microarchitectural functionalities of processor implementations and can compromise the security of computational environments even in the presence of sophisticated protection mechanisms like virtualization and sandboxing. This newly evolving research area has attracted significant interest due to the broad application range and the potentials of these attacks. Cache Analysis and Branch Prediction Analysis were the only types of MA that had been known publicly. In this paper, we introduce Instruction Cache (I-Cache) as yet another source of MA and present our experimental results which clearly prove the practicality and danger of I-Cache Attacks.

Keywords: Instruction Cache, Modular Exponentiation, Montgomery Multiplication, RSA, Side Channel Analysis, MicroArchitectural Analysis.

1 Introduction

Side-channel cryptanalysis exploit the information leakage from execution time, power consumption or any such side-channels during the computation of cryptographic operations, c.f. [19, 20]. Cryptographic implementations leak sensitive information because of the physical properties and requirements of the cryptographic implementations and computational environments. Classical cryptography analyzes the cryptosystems as perfect mathematical models and ignores such physical requirements, thus fails to identify side-channel leakages. Therefore it is inevitable to utilize both classical cryptography and side-channel cryptanalysis in order to develop and implement secure systems and security architectures.

The initial focus of side-channel research was on smart card security. We discuss the reasons of this situation in the next section. For now, we just want to mention that side-channel analysis of computer systems started to attract more attention after Brumley and Boneh demonstrated a successful and practical remote timing attack on real applications over a local network [14]. Since then, we have seen increased research efforts on the security analysis of the daily life PC platforms from side-channel point of view. The most important recent advance in the field is the realization of MicroArchitectural Analysis.

MicroArchitectural Analysis (MA), which is a newly evolving area of side-channel cryptanalysis, studies the effects of common processor components and their functionalities on the security of software cryptosystems, c.f. [11, 2, 4, 3, 7, 29, 32]. As a natural consequence

of strictly throughput, performance, and “performance per watt” oriented goals of modern processor designs and also highly “time-to-market” driven business philosophy, the resulting products, i.e. commodity processor architectures in the market, lack a thorough security analysis. The main element that gave birth to MicroArchitectural Analysis area is indeed this particular gap between the current processor architectures and the ideal secure computing environment. The advances in MicroArchitectural Analysis field initiated a new research vector to identify, analyze, and mitigate the security vulnerabilities that are caused by the design and implementation of processor components.

All of these cited pure software MicroArchitectural attacks, including the one presented in this paper, can compromise security systems despite of sophisticated partitioning methods such as memory protection, sandboxing or even virtualization. The reason for the failure of these trust mechanisms is because these new attacks “*simply exploit deeper processor ingredients below the trust architecture boundary*” as stated in [2, 4]. The new security and virtualization technologies such as Intel’s LT and VT, AMD’s Pacifica, ARM’s Trustzone, software based virtualization mechanisms like those from VMWare are all potentially susceptible to MA attacks. We want to emphasize that so far there had not been any publicly known MA attack incidents on these systems. But we believe it is only a matter of time until they are shown to be compromised via MA.

It is crucial to identify every possible MicroArchitectural vulnerability in order to understand the real potential of MicroArchitectural Analysis and to develop more secure systems by employing appropriate software countermeasures and making required hardware changes to future architectures.

In this paper, we identify a new MicroArchitectural attack source, in other words, another processor component that causes security vulnerabilities. We show that Instruction Cache, which is used to reduce the average time to read instruction codes from main memory, can be exploited to extract sensitive information regarding the execution of a cryptosystem. Our results clearly show the practicality and danger of I-cache attacks and put I-cache into the list of known MicroArchitectural attack sources.

In the next section, we give an overview of MicroArchitectural Analysis including a brief history. We explain what an instruction cache is and how it works in Section 3. The basics of RSA cryptosystem and the details of OpenSSL’s RSA implementation are presented in Section 4. Sections 5 and 6 outline the underlying idea of instruction cache attacks and detail a sample attack on sliding window exponentiation of RSA implemented in OpenSSL. We also present our experimental results, which prove the practicality of instruction cache analysis concept, in Section 6. Then we point out a protected RSA implementation and conclude our paper in the last section.

2 Overview and Brief History of MicroArchitectural Analysis

The initial focus of side-channel research was on smart card security. Smart cards store secret values and they are designed to protect and process these secrets. Therefore, there is a serious financial gain involved in cracking smart cards, as well as, analyzing them and

developing more secure smart card technologies, and this is one of the main reasons why smart cards had the initial focus of side-channel research. However, the recent advances and trends in microprocessor market, especially the development of microprocessor based security features (e.g. Intel’s LT and VT Technologies, AMD Pacifica), and also the recent promises from the Trusted Computing community indicate the security assurance of storing and processing secret values, establishing virtually separate execution environments, etc. on computer platforms. As an eventual result, the side-channel analysis of computer platforms has become a necessity.

Another reason of the high attention to side-channel analysis of smart cards is due to the ease of applying such attacks on smart cards. The measurements of side-channel information on smart cards are almost “noiseless”, which makes such attacks very practical. On the other hand, there are so many factors that affect such measurements on real commodity computer systems. These factors make it much more difficult to perform successful side-channel attacks on “real” computers within our daily life. Thus, the side-channel vulnerability of computer systems was not seriously considered to be harmful until 2003. This was changed when Brumley and Boneh demonstrated a successful and practical remote timing attack on real applications over a local network [14]. They simply adapted the attack principle introduced by Schindler in [34] and applied it to a real web server to show that side-channel attacks are a real danger not only to smart cards but also to widely used computer systems. Their work was significantly improved by Acicmez et. al. in 2005 [9].

We have seen increased research efforts on the side-channel analysis of commodity PC platforms for the last few years. Soon, it was realized that the functionality of some microprocessor components cause serious side-channel leakage. These efforts led to the development of MicroArchitectural Analysis area.

MicroArchitectural Attacks exploit the microarchitectural components of a processor to reveal cryptographic keys. The functionality of some processor components generates data dependent variations in execution time and power consumption during the execution of cryptosystems. These variations either directly gives the key value out during a single cipher execution (c.f. [2]) or leaks information which can be gathered during many executions and analyzed to compromise the system (c.f. [29, 11, 24]).

The actual roots of MA goes a long way back to [18, 36]. Although the security risks of processor components like cache were implicitly pointed out in these publications, concrete and widely applicable security attacks based upon processor functionalities have recently been worked out and immediately attracted significant public interest. There are currently two types of MA in the literature¹: Cache Analysis and Branch Prediction Analysis.

A cache-based attack, abbreviated to “cache attack” from here on, exploits the cache behavior of a cryptosystem by obtaining the execution time and/or power consumption variations generated via cache hits and misses. The cache vulnerability of computer systems

¹ There is, in fact, a new paper describing a recently discovered MA type [6]. The details of it were not publicly available by the time we wrote this current paper. Therefore we omit this attack here and prefer not to disclose the details.

has been known for a long time, c.f. [18, 19, 21], however actual realistic and practical cache attacks were not developed until recent years.

Cache analysis techniques enable an unprivileged process to attack another process, e.g., a cipher process, running in parallel on the same processor as done in [29, 24, 32]. Furthermore, some of the cache attacks can even be carried out remotely, e.g., over a local network [7].

The previous cache attacks are data-path attacks, i.e., exploit the data access patterns of a cipher. The memory accesses of software cryptosystems, especially S-box based ciphers like DES and AES, employ key-dependent table lookups, indices of which are simple functions of the key and the plaintext. Revealing these memory access patterns, i.e. lookup indices, via cache statistics and the knowledge of the processed message, e.g. in a known-text attack, make it relatively easy to break these ciphers.

A new group of MA attacks, called Branch Prediction Analysis (BPA), on the other hand, exploit instruction path of a cipher [2, 1, 4, 3]. In other words, an adversary can reveal the execution flow of a cipher using BPA, and if this execution flow is key dependent as in the case of RSA and ECC, then he can compromise the system. The most powerful BPA, which is called Simple Branch Prediction Analysis (SBPA), was shown to extract almost all of the RSA key bits during a single RSA operation. Immediately after it became public around the end of 2006, SBPA attracted very significant attention due to its implications. Aciğmez et. al. briefly outlined why SBPA endangers most of the current systems and detailed some techniques to show how SBPA can be used to break even “thought-to-be-side-channel-immune” systems in [2, 1, 5].

In this paper, we combine these two concepts: exploiting cache architecture and revealing instruction paths. We present several attacks that rely on instruction cache (I-cache) architecture of CPUs. All of the previous cache attacks exploit the side-channel leakage through data cache of a CPU. To the best of our knowledge, this is the first approach to exploit side-channel leakage due to I-cache architectures. Our experimental results indicate that I-cache analysis is as efficient as SBPA.

3 Instruction Cache

Our software I-cache timing attack exploits the functionality of I-cache implemented in microprocessors. A high-frequency processor needs to retrieve data at a very high speed in order to utilize its functional resources. The latency of a main memory is not short enough to match this demand of high-speed data delivery. The gap between the memory and the processor speed has been continuously increasing for the last 3 decades as Moore’s Law holds. Common to all processors, the attempt to overcome the drawbacks of this gap is the employment of a special buffer called cache.

A cache is a small and fast storage area used by a CPU to reduce the average memory access time. It acts as a buffer between the main memory and the processor core and provides the processor fast and easy access to the most frequently used data (including instructions) without frequent external bus accesses.

Cache stores the copies of the most frequently used data. When the processor needs to read a location in main memory, it first checks to see if the data is already in the cache. If the data is already in the cache (a cache hit), the processor immediately uses this data instead of accessing the main memory, which has a significantly longer latency than a cache. Otherwise (a cache miss), the data is read from the memory and a copy of it is stored in the cache. This copy is expected to be used in the near future due to the temporal locality property.

A cache is partitioned into a number of non-overlapping fixed size blocks, called cache blocks or cache lines. The minimum amount of data that can be read from the main memory into a cache at once is called cache line or cache block size, i.e., each cache miss causes a cache block to be retrieved from a higher level memory. The reason why a block of data is transferred from the main memory to the cache instead of transferring only the data that is currently needed lies in spatial locality property. Since a cache is limited in size, storing new data in a cache mandates eviction of some of the previously stored data.

Before moving on to the next section, we want to mention two very important concepts that affect the functional behavior of a cache: the mapping strategy and the replacement policy. We will only give very brief information on these concepts in this paper. For further discussion on cache architectures and locality properties see [16, 33, 15].

Cache mapping strategy is the method of deciding where to store, and thus to search for, a data in a cache. Three main cache mapping strategies are direct, fully associative and set associative mapping. In a direct mapped cache, a particular data block can only be stored in a single certain location in the cache. On the contrary, a data block can be placed in potentially any location in a fully associative cache. The location of a particular placement is determined by the replacement policy. Set associative mapping is a blend of these two mapping strategies. Set associative caches are divided into a number of same size sets, called cache sets, and each set contains the same fixed number of cache lines. A data block can be stored only in a certain cache set (just like in a direct mapped cache), however it can be placed in any location inside this set (like in a fully associative cache). Again, the particular location of a data inside its cache set is determined by the replacement policy.

The replacement policy is the method of deciding which data block to evict from the cache in order to place the new one in. The ultimate goal is to choose the data that is most unlikely to be used in the near future. There are several cache replacement policies proposed in the literature (c.f. [16, 33]). In this document, we focus on a specific one: least-recently-used (LRU). It is the most commonly used policy and it picks the data that is least recently used among all of the candidate data blocks that can be evicted from the cache.

Many processors employ different caches for data and code segments of a process. The instruction cache is responsible for storing recently used instructions from the code segment and quickly delivering them to the processor core when the accessed instructions are in the I-cache. When a process starts executing a code block that is not in the cache, i.e., in case of a cache miss, the processor loads these instructions from main memory into the cache. This situation happens either at the initial execution of a function (i.e., a cold miss), or after a cache conflict (i.e., conflict miss). Since a cache is limited in size, several different code blocks share the same cache sets/lines. A cache conflict or collision is the situation that occurs when

an attempt is made to store two or more different data/code items at a cache location that can hold only one of them. In case of a cache conflict between different code blocks, they evict each other from the instruction cache when their executions are interleaved. In our I-cache attacks, we exploit this particular consequence of cache conflicts by creating intentional conflicts between the instructions of RSA cipher and a spy code and forcing the processor to evict the RSA instructions out of I-cache.

4 The concept of I-cache Attacks

Cryptosystems have data-dependent memory access patterns, which can be revealed by observing cache hit/miss statistics through side channels. Cache attacks rely on the cache hits and misses that occur during the encryption / decryption process of a cryptosystem. Even if the same instructions are executed for any particular (plaintext, cipherkey) pair, the cache behavior during the execution was shown to cause variations in the program execution time and power consumption. Cache attacks try to exploit such variations to narrow the exhaustive search space of cryptographic keys, c.f. [31, 39, 40, 7, 8, 29, 30, 11, 25, 22, 24, 32, 12, 13].

The previous studies in MicroArchitectural Analysis area, such as [29, 32, 24, 2, 4], initiated a new attack paradigm, which relies on simultaneous multi-threading / multi-tasking functionality of modern processors. Simultaneous multithreaded (SMT) processors have the capability of executing more than one execution thread simultaneously on the same physical processor. Multi-core processors also have the same capability. The main difference between multi-core and SMT is that expensive resources of a processor core (e.g., functional units, data and instruction cache, BTB) are shared between different threads in a SMT processor. The basic and simple resources are explicitly doubled to give the sense of two logical processors on a single physical processor core. Thus, this design technique enables the simultaneous execution of more than one process on the same physical processor core by taking advantage of thread-level parallelism, as if there were more than one processor [35, 37].

Single-threaded processor cores, on the other hand, execute only a single process/thread at any given time. However, the operating systems manage to distribute the processor time among all the active processes and give the users the feeling of a parallel, multi-threading execution. The OS basically decomposes the execution of each process into a series of short threads and schedules the execution of these threads with respect to each other.

Irrespectively of single-threaded or hardware-assisted multi-threaded, some processor resources are always shared among the active threads on the system, which enables one process to spy on another process, c.f. [29, 32, 24, 2, 4]. Although the memory protection mechanisms prevents a process to directly read other processes' data, the functionality of shared resources leak the so-called metadata, c.f. [29], and (e.g.) causes disclosure of the secret / private keys used in security systems.

In our I-cache attacks, we rely on the concept of executing a spy code, which keeps track of the changes in the state of I-cache, i.e., metadata, during the execution of a cipher process. A spy code / process can run simultaneously or quasi-parallel with the cipher process and

determine which instructions are executed by the cipher. It achieves this goal by spying on the cipher execution via observing the I-cache state transitions.

Assume that an adversary tries to understand whether a certain I-cache set is “touched” by the cipher, i.e., modified, during the execution of a part of cipher code. The spy allows the cipher to run and takes over the processor shortly before the execution of the “spied-on” part of cipher code. This task of pausing the cipher execution at a determined point, even though it sounds nice, is very tricky and requires very fine-crafted spy code. However, it is feasible and was successfully used in earlier studies on cache attacks by Neve et. al. in [24] to devise an attack on the last AES round, which is composed of a relatively small number of instructions. A similar idea is also presented in a recent paper [38]. Although [38] proposes to exploit OS scheduling mechanism to steal CPU cycles unfairly, the cheating idea and the source code can easily be adapted to MicroArchitectural Analysis attacks.

After the spy takes over, it ensures that this particular I-cache set does not contain any instructions from the cipher by executing a set of “dummy” instructions. These dummy instructions are not intended to perform any calculations or tasks other than filling some I-cache space. These dummy instructions shall fill completely and precisely this I-cache set, no more no less. During the execution of dummy instructions, the processor has to store them into the cache, which inevitably causes the eviction of the previous entries in that I-cache location. That way, the spy sets the state of this particular I-cache set to a known predetermined state and then it lets the cipher run the “spied-on” part of code.

The spy takes control of the processor after the execution of a relatively small number of cipher instructions. If the executed cipher instructions touch the I-cache set under observation, this will cause the eviction of some of these spy-owned dummy instructions from the I-cache. When the spy takes control of the processor, it re-executes the same dummy instructions but this time also measures their total execution time. If some of these dummy instructions are not in I-cache, which indicates the modification of this I-cache set due to cipher execution, then the measured execution time will take longer simply because the evicted instructions must be retrieved from the memory which has a significantly larger latency compared to the cache.

During the quasi-parallel execution of spy and cipher processes, a malicious spy routine can continuously interrupt the cipher execution with short intervals and apply the above basic technique to every single I-cache set each time it takes the control. On simultaneous multi-threading systems, the spy routine does not even need to interrupt the cipher execution and can observe it “on-the-fly” as done in [32, 1]. If the adversary can get measurements with high enough resolution, i.e., if he can estimate which cache sets are modified during the execution of which part of the cipher code, this will reveal the execution flow of the cipher. Therefore, such a spy routine has the potential of revealing the entire execution flow of the cipher on almost any processor architecture as long as there is an I-cache and its metadata is preserved during the transfer of processor time between different processes.

In the next section, we will outline a sample cache attack on OpenSSL’s sliding window exponentiation. We want to mention that we use OpenSSL’s SWE just as a case study to prove the concept of I-cache analysis. The actual application range of I-cache attacks is much

more broader than this simple case study. Our results from this case study show that an adversary can easily get a measurement resolution high enough to compromise very critical security systems.

5 A Case Study: Attack on OpenSSL’s Sliding Window

OpenSSL takes advantage of the difference between multiprecision multiplication and square operations to improve the performance of its RSA implementation. During a montgomery operation, OpenSSL first calls either multiplication or square functions from BIGNUM library and then reduces the result to the modulus via montgomery reduction function. In case of sliding window exponentiation, this technique causes key-dependent sequence of multiplication and square function calls. The current version of OpenSSL employs either sliding window or fixed window exponentiation depending on the user’s choice.

A practical method to reveal the multiplication/square operation sequence is the following. The spy function can evict the instructions of multiplication function (or square function, resp.) and measure the execution time of its own dummy instructions as described in the previous section. The higher execution time in spy measurements indicate the execution of this multiplication (square, resp.) function. Therefore, the spy can determine when the cipher calls this particular function, which also directly reveals the multiplication/square operation sequence. The spy function can perform the attack either “on-the-fly” on simultaneous multi-threading systems (c.f. [32, 1]) or via exploiting OS-scheduling (c.f. [24]).

Our attack scenario is the following. A “protected” crypto process executes the RSA signing/decryption process and also a spy process is executed simultaneously or quasi-parallel with the cipher and it continuously does the following:

1. continuously executes a number of dummy instructions, and
2. measures the overall execution time of all of these instructions

in such a way that these dummy instructions precisely evicts the instructions of BIGNUM multiprecision multiplication function from I-cache.

Assume that the multiplication function instructions span from logical address A to B . Due to the properties of cache architectures, an instruction block, i.e., continuous consecutive instructions, must span from logical address A_1 to B_1 to map to the same I-cache sets with the multiplication function, where the least significant parts of A and A_1 and also B and B_1 must be equal. To completely evict the multiplication function from I-cache, the spy routine has to execute a number of different such instruction blocks and this number needs to be equal or greater than the number of associativity of the I-cache. Clearly, implementing this spy routine requires the knowledge of logical address space of the multiplication function and the details of the I-cache. We want to mention that it is easy to learn the properties of an I-cache either from the manufacturer specs or by using simple benchmarks as explained in [41].

To validate our aforementioned I-cache analysis strategy, we performed some practical experiments. We compiled the RSA decryption function of OpenSSL (version 0.9.8d) with




Fig. 1. Experimental Results

the choice of SWE Exponentiation. We disassembled the executable file to see the logical addresses of BIGNUM multiprecision multiplication function instructions. GNU Project debugger (i.e. gdb) has two functions, “info line” and “disas”, that we used for this task. Then we implemented our spy routine based on these logical addresses and the parameters of the I-cache architecture on our platform. Then we carried out our attack by letting the spy routine run and make measurements with relatively short intervals during the execution of RSA signing operation.

After analyzing the spy measurements, we ended up with the results shown in Figure 1. These timing measurements are taken during *a single RSA operation* under a random 1024-bit key. It is clear in the figure that one can observe the operation sequence of RSA via I-cache analysis. These results clearly indicate that such I-cache attacks are feasible and as dangerous and serious as Simple Branch Prediction Analysis, which attracted very significant attention immediately after its publication because of the several very serious security vulnerabilities it implied. The same vulnerabilities sketched out in [2, 1] can also be exploited via I-cache analysis.

6 Discussions

Looking closely to this graph, one can also distinguish different phases of sliding window exponentiation:

- A table initialization phase of the first exponentiation
- B the first exponentiation phase
- C table initialization phase of the second exponentiation
- D the second exponentiation phase

E the remaining RSA operations including CRT combination

If an adversary runs a stand-alone spy process on a machine, he also needs to ensure that the measurement results from the spy indeed correspond to the cipher process. Since there are possibly many other processes running on the same machine, he needs to distinguish when the spy process is measuring the cipher execution but not another process. In an experimental setup, we can set the system to ensure this synchronization by (e.g.) assigning high priorities to spy and cipher processes. However, in a real attack, where the adversary does not have privileges to configure such parameters, this synchronization issue becomes problematic. As seen in our results, different RSA phases are distinguishable, which can enable an adversary to understand whether the spy process is spying on a cipher or another process and thus overcome the synchronization problem.

We could reveal the operation sequence of sliding window exponentiation almost completely during a single RSA operation via I-cache analysis. Revealing the multiplication / square sequence of an RSA with binary exponentiation algorithm, which was the case in [1], would directly give the value of d out. In our "proof-of-concept" attack on sliding window exponentiation, the multiplication / square sequence reveals around 200 "scattered" bits of each 512-bit exponents, c.f., [32]. It is not clear today whether the knowledge of 200 bits scattered over 512-bit exponent is sufficient to break RSA. In other words, we do not know any methods that can directly leverage this information to factorize the public modulus. However, the application range of I-cache analysis is definitely not limited to SWE. [1] and [5] discuss several situations that are potentially vulnerable to SBPA, which are also valid for I-cache analysis. For example, leakage of just a few secret bits of the respective ephemeral keys leads to a total break of (EC)DSA, c.f., [17, 26, 27].

Unfortunately, the discussions in [1] and [5] are not comprehensive in terms of covering the potential threats due to MicroArchitectural analysis. We have discovered a vulnerability, which is not mentioned in [1, 5], in OpenSSL library that can be exploited via I-cache analysis. A new patch was already prepared by OpenSSL team to fix this vulnerability and it will be available in future version of OpenSSL soon. Our objective in this paper is to describe and prove only the concept of I-cache analysis. Our results and details on this OpenSSL vulnerability will follow in subsequent publications.

Comparison of I-cache Analysis to Data Cache and Branch Prediction Analysis

Data cache attacks try to reveal the data-access patterns of cryptosystems. On the other hand, we reveal the instruction flow of cryptosystems in I-cache analysis. The cryptosystem implementations with fixed instruction flow, which is usually the case for block ciphers like AES, are not vulnerable to I-cache and Branch prediction analysis whilst data cache attacks can exploit the table lookups of such ciphers. It is also possible to determine the execution flow of a cipher (e.g. RSA) by analyzing the data access patterns as done in [32]. However, implementations can avoid this threat by carefully handling the data structures. For example, OpenSSL changed to way it handles the RSA structures to avoid data cache attacks like [32]. Even when the data structures are handled in a special way, I-cache analysis can compromise

the implementations if the execution flow remains key-dependent. Similarly, data cache attacks can be applied on implementations with fixed execution flow if the data access patterns are key-dependent. Therefore, both data and instruction cache analysis must be considered during the design and implementation of security critical systems.

The basic difference between I-cache and Branch prediction analysis is the following. Branch prediction analysis presented in [1, 3] specifically targets conditional branches. A conditional branch controls the execution of different instruction paths. Thus, the outcome of a conditional branch, which can be observed via BPA, leaks the instruction path to an adversary. However, using conditional branch is only one way to implement execution flow control. There are other techniques, which may be protected against BPA, to conditionally alter the execution flow without the use of conditional branches. In this sense, I-cache analysis is broader than BPA because it reveals the execution flow regardless of how execution flow control is implemented. For example, [10] proposes to use indirect jumps instead of conditional branches as a countermeasure to BPA, which is still vulnerable to I-cache analysis.

7 Conclusions

We have showed that a major processor component, Instruction Cache (I-cache), causes serious security vulnerabilities and can be used in a side-channel attack as a source of information leakage. The special side-channel area that exploits processor components is called MicroArchitectural Analysis (MA) and there are currently two types of MA in the literature: Cache Analysis and Branch Prediction Analysis. Our contribution in this paper is to introduce I-cache Analysis as yet another MA type.

We have presented a simple pure software-based I-cache attack on OpenSSL’s RSA implementation as a “proof of concept” to show that I-cache attacks have the potential to reveal the execution flow of cryptosystems like RSA, which can lead to a complete break if the cryptosystem is implemented with key-dependent execution flow. I-cache analysis, just like cache and branch prediction analysis, can compromise security systems even in the presence of security mechanisms like sandboxing and virtualization because all of these attacks exploit deep processor functionalities which are below the trust architecture boundary of these security mechanisms.

I-cache attacks are instruction-path attacks unlike the previous cache attacks, which exploit the data-path of a cipher execution. Branch Prediction Analysis is the first instance of MA that reveals the instruction paths of the cryptosystems. Similar to Simple Branch Prediction Analysis, which is the most powerful Branch Prediction Analysis variant, I-cache analysis has the potential to reveal the complete operation sequence of a cryptosystem during a single execution. Simple Branch Prediction Analysis attracted very significant attention immediately after its publication because of the several very serious security vulnerabilities it implied. The same vulnerabilities pointed out in [2, 1] can also be exploited via I-cache analysis.

It is extremely important and urgent to identify every possible MicroArchitectural vulnerability in order to understand the real potential of MicroArchitectural Analysis and to develop

more secure systems by developing and employing appropriate software countermeasures and possibly making required hardware changes to future processors. It is advisable to avoid key-dependent instruction paths in cryptographic software implementations as much as possible. A method was already outlined in [5] to protect RSA against instruction path attacks like Branch Prediction and I-cache Analysis. We believe, such techniques and countermeasures shall be developed for other cryptosystems and employed in cryptographic applications / libraries.

References

1. O. Aciğmez, Ç. K. Koç, and J.-P. Seifert. On The Power of Simple Branch Prediction Analysis. *2007 ACM Symposium on InformAtion, Computer and Communications Security (ASIACCS'07)*, R. Deng and P. Samarati, editors, pages 312-320, ACM Press, 2007.
2. O. Aciğmez, Ç. K. Koç, and J.-P. Seifert. On The Power of Simple Branch Prediction Analysis. Cryptology ePrint Archive, Report 2006/351, October 2006.
3. O. Aciğmez, Ç. K. Koç, and J.-P. Seifert. Predicting Secret Keys via Branch Prediction. *Topics in Cryptology — CT-RSA 2007, The Cryptographers' Track at the RSA Conference 2007*, M. Abe, editor, pages 225-242, Springer-Verlag, Lecture Notes in Computer Science **series 4377**, 2007.
4. O. Aciğmez, J.-P. Seifert, and Ç. K. Koç. Predicting Secret Keys via Branch Prediction. Cryptology ePrint Archive, Report 2006/288, August 2006.
5. O. Aciğmez, S. Gueron, and J.-P. Seifert. New Branch Prediction Vulnerabilities in OpenSSL and Necessary Software Countermeasures. Cryptology ePrint Archive, Report 2007/039, February 2007.
6. O. Aciğmez and J.-P. Seifert. Cheap Hardware Parallelism Implies Cheap Security. *4th Workshop on Fault Diagnosis and Tolerance in Cryptography — FDTC 2007*, Austria, September 10, 2007, to appear.
7. O. Aciğmez, W. Schindler, and Ç. K. Koç. Cache Based Remote Timing Attack on the AES. *Topics in Cryptology — CT-RSA 2007, The Cryptographers' Track at the RSA Conference 2007*, M. Abe, editor, pages 271-286, Springer-Verlag, Lecture Notes in Computer Science **series 4377**, 2007.
8. O. Aciğmez and Ç. K. Koç. Trace-Driven Cache Attacks on AES. Cryptology ePrint Archive, Report 2006/138, April 2006.
9. O. Aciğmez, W. Schindler, Ç. K. Koç. Improving Brumley and Boneh Timing Attack on Unprotected SSL Implementations. *Proceedings of the 12th ACM Conference on Computer and Communications Security*, C. Meadows and P. Syverson, editors, pages 139-146, ACM Press, 2005.
10. G. Agosta, L. Breveglieri, I. Koren, G. Pelosi and M. Sykora. Countermeasures Against Branch Target Buffer Attacks. *4th Workshop on Fault Diagnosis and Tolerance in Cryptography (FDTC'07)*, to appear.
11. D. J. Bernstein. Cache-timing attacks on AES. Technical Report, 37 pages, April 2005.
12. G. Bertoni, V. Zaccaria, L. Breveglieri, M. Monchiero, G. Palermo. AES Power Attack Based on Induced Cache Miss and Countermeasure. *International Symposium on Information Technology: Coding and Computing - ITCC 2005*, volume 1, pages 4-6, 2005.
13. J. Bonneau and I. Mironov. Cache-Collision Timing Attacks against AES. *Cryptographic Hardware and Embedded Systems — CHES 2006*, L. Goubin and M. Matsui, editors, pages 201-215, Springer-Verlag, Lecture Notes in Computer Science **series 4249**, 2006.
14. D. Brumley and D. Boneh. Remote Timing Attacks are Practical. *Proceedings of the 12th Usenix Security Symposium*, pages 1-14, 2003.
15. P. Genua. A Cache Primer. Technical Report, Freescale Semiconductor Inc., 16 pages, 2004.
16. J. Handy. *The Cache Memory Book*. 2nd edition, Morgan Kaufmann, 1998.
17. N. A. Howgrave-Graham and N. P. Smart. Lattice Attacks on Digital Signature Schemes. *Design, Codes and Cryptography*, volume 23, pages 283-290, 2001.
18. W. M. Hu. Lattice scheduling and covert channels. *Proceedings of the IEEE Symposium on Security and Privacy*, pages 52-61, IEEE Computer Society, 1992.

19. P. C. Kocher. Timing Attacks on Implementations of Diffie–Hellman, RSA, DSS, and Other Systems. *Advances in Cryptology - CRYPTO '96*, N. Koblitz, editor, pages 104–113, Springer-Verlag, Lecture Notes in Computer Science **series 1109**, 1996.
20. P. C. Kocher, J. Jaffe, B. Jun. Differential Power Analysis. *Advances in Cryptology - CRYPTO '99*, M. Wiener, editor, pages 388–397, Springer-Verlag, Lecture Notes in Computer Science **series 1666**, 1999.
21. J. Kelsey, B. Schneier, D. Wagner, C. Hall. Side Channel Cryptanalysis of Product Ciphers. *Journal of Computer Security*, volume 8, pages 141–158, 2000.
22. C. Lauradoux. Collision attacks on processors with cache and countermeasures. *Western European Workshop on Research in Cryptology — WEWoRC 2005*, C. Wolf, S. Lucks, and P.-W. Yau, editors, pages 76–85, 2005.
23. A. J. Menezes, P. van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*, CRC Press, New York, 1997.
24. M. Neve and J.-P. Seifert. Advances on Access-driven Cache Attacks on AES. *Selected Areas of Cryptography — SAC'06*, to appear.
25. M. Neve, J.-P. Seifert, Z. Wang. A refined look at Bernstein's AES side-channel analysis. *Proceedings of ACM Symposium on Information, Computer and Communications Security — ASIACCS'06*, page 369, Taipei, Taiwan, March 21–24, 2006.
26. P. Q. Nguyen and I. E. Shparlinski. The Insecurity of the Digital Signature Algorithm with Partially Known Nonces. *Journal of Cryptology*, volume 15, no. 3, pages 151–176, Springer, 2002.
27. P. Q. Nguyen and I. E. Shparlinski. The Insecurity of the Elliptic Curve Digital Signature Algorithm with Partially Known Nonces. *Design, Codes and Cryptography*, volume 30, pages 201–217, 2003.
28. Openssl: the open-source toolkit for ssl/tls. Available online at: <http://www.openssl.org/>.
29. D. A. Osvik, A. Shamir, and E. Tromer. Cache Attacks and Countermeasures: The Case of AES. *Topics in Cryptology — CT-RSA 2006, The Cryptographers' Track at the RSA Conference 2006*, D. Pointcheval, editor, pages 1–20, Springer-Verlag, Lecture Notes in Computer Science **series 3860**, 2006
30. D. A. Osvik, A. Shamir, and E. Tromer. Other People's Cache: Hyper Attacks on HyperThreaded Processors. Presentation available at: <http://www.wisdom.weizmann.ac.il/~tromer/>.
31. D. Page. Theoretical Use of Cache Memory as a Cryptanalytic Side-Channel. Technical Report CSTR-02-003, Department of Computer Science, University of Bristol, June 2002.
32. C. Percival. Cache missing for fun and profit. *BSDCan 2005*, Ottawa, 2005.
33. D. Patterson and J. Hennessy. *Computer Architecture: A Quantitative Approach*. 4th edition, Morgan Kaufmann, 2006.
34. W. Schindler. A Timing Attack against RSA with the Chinese Remainder Theorem. *Cryptographic Hardware and Embedded Systems — CHES 2000*, Ç.K. Koç and C. Paar, editors, pages 110–125, Springer-Verlag, Lecture Notes in Computer Science **series 1965**, 2000.
35. J. Shen and M. Lipasti. *Modern Processor Design: Fundamentals of Superscalar Processors*. McGraw-Hill, 2005.
36. O. Sibert, P. A. Porras, and R. Lindell. The Intel 80x86 Processor Architecture: Pitfalls for Secure Systems. *IEEE Symposium on Security and Privacy*, pages 211–223, 1995.
37. A. Silberschatz, G. Gagne, and P. B. Galvin. *Operating system concepts*. 7th edition, John Wiley and Sons, 2005.
38. D. Tsafir, Y. Etsion and D. G. Feitelson. Secretly Monopolizing the CPU Without Superuser Privileges. 16th USENIX Security Symposium, April 2007.
39. Y. Tsunoo, T. Saito, T. Suzaki, M. Shigeri, H. Miyauchi. Cryptanalysis of DES Implemented on Computers with Cache. *Cryptographic Hardware and Embedded Systems — CHES 2003*, C. D. Walter, Ç. K. Koç, and C. Paar, editors, pages 62–76, Springer-Verlag, Lecture Notes in Computer Science **series 2779**, 2003.
40. Y. Tsunoo, E. Tsujihara, K. Minematsu, H. Miyauchi. Cryptanalysis of Block Ciphers Implemented on Computers with Cache. *ISITA 2002*, 2002.
41. K. Yotov, S. Jackson, T. Steele, K. Pingali and P. Stodghill. Automatic Measurement of Instruction Cache Capacity. 18th International Workshop on Languages and Compilers for Parallel Computing (LCPC), 2005.