

Øzone: Efficient Execution with Zero Timing Leakage for Modern Microarchitectures

Zelalem Birhanu Aweke and Todd Austin
University of Michigan
Email: {zaweke, austin}@umich.edu

Abstract—Time variation during program execution can leak sensitive information. Time variations due to program control flow and hardware resource contention have been used to steal encryption keys in cipher implementations such as AES and RSA. A number of approaches to mitigate timing-based side-channel attacks have been proposed including cache partitioning, control-flow obfuscation and injecting timing noise into the outputs of code. While these techniques make timing-based side-channel attacks more difficult, they do not eliminate the risks. Prior techniques are either too specific or too expensive, and all leave remnants of the original timing side channel for later attackers to attempt to exploit.

In this work, we show that the state-of-the-art techniques in timing side-channel protection, which limit timing leakage but do not eliminate it, still have significant vulnerabilities to timing-based side-channel attacks. To provide a means for total protection from timing-based side-channel attacks, we develop **Ozone**, the first zero timing leakage execution resource for a modern microarchitecture. Code in **Ozone** executes under a special hardware thread that gains exclusive access to a single core’s resources for a fixed (and limited) number of cycles during which it cannot be interrupted. Memory access under **Ozone** thread execution is limited to pre-allocated cache lines that cannot be evicted, and all **Ozone** threads begin execution with a known fixed microarchitectural state. We evaluate **Ozone** using a number of security sensitive kernels that have previously been targets of timing side-channel attacks, and show that **Ozone** eliminates timing leakage with minimal performance overhead.

I. INTRODUCTION

Perhaps the most fruitful hardware security vulnerability has been timing-based side-channel attacks. A *timing-based side-channel attack*¹ is one in which the time it takes to perform specific operations reveals information about secrets within the system. There are three characteristics of modern systems that enable timing-based side-channel attacks:

1. Variable-latency operations: In an effort to improve the performance of a system, common operations are made to take less time to execute than uncommon ones; thus, it becomes possible to infer the constituency of instructions a program executes simply by examining its latency. Furthermore, control flow, speculation, caches, and a variety of other system features render execution latencies that are a function of the program’s code and data, thereby creating opportunities to learn about such things. Kocher used this characteristic to attack RSA authentication [11], after noting that early RSA implementations performed different amounts of work for “0” and “1” key bits.

2. Resource sharing: To keep costs in check, functional and storage resources in a modern microarchitecture are shared concurrently among threads and processes; thus, it becomes possible for an attacker program to create contention on these resources and examine, by virtue of its own performance, the extent to which the shared resources are utilized by the victim program. Examples of resource sharing timing side-channel attacks include L1 instruction cache based attacks [2], L1 data cache based attacks [20], branch predictor based attacks [1] and most recently last-level cache based attacks [17], [25].

¹Some taxonomies of side-channel attacks create distinct categories for cache and control flow attacks. Since these attacks are perpetrated to correlate secrets to the timing of the cache or control flow, we broadly address all of these attacks under the class of timing attacks.

3. Fine-grained performance monitoring: To facilitate debugging and code optimization, modern systems provide high-precision timing facilities, *e.g.*, Intel’s *rdtsc* user-level cycle counter read operation, and performance monitoring capabilities. While most timing side-channel attacks make use of high-precision timing facilities, performance counters (*e.g.*, cache miss counters) have also been used in the past [12].

A. Mitigating Timing Attacks

There is a large body of work that has proposed various techniques to mitigate timing-based side-channel attacks. We broadly note that these techniques work to address each of the above three system characteristics that leak timing information. Variable-latency operations can be mitigated through a variety of software and/or hardware measures to reduce the data-dependent latency of a program. For example, some efforts to reduce control flow dependence on data utilize if-conversion and predicated execution to eliminate control hammers² [7], [18], [21]. Other efforts work to lock down critical data into the cache [24] or a scratchpad memory [15] so as to reduce memory latency variability. While these techniques work to limit the degree of data-dependent timing variation in programs, they do not completely eliminate it since a wide range of microarchitectural features continue to perform data-dependent optimizations, leading to attacks via branch predictors [1] and cache conflicts [3], [5], among others.

To address resource sharing, techniques have proposed to isolate programs to their own dedicated CPU and caches, so as to eliminate potential competition for resources from other aggressor threads, *e.g.* [6]. While these techniques can eliminate aggressors at great cost, they still ultimately suffer from timing leakage due to variable-latency operations that remain on the single-threaded microarchitecture. As an example, Bernstein’s cache attack on AES [3] works by monitoring how AES-internal data cache conflicts vary with plaintext, thus making the aggressor in this case the AES code itself.

Finally, to mitigate the problems of fine-grained performance analysis, efforts have proposed time padding and obfuscation to obscure the true activities of a program, *e.g.*, [6], [21]. The former technique delays results from a program such that all results return after a fixed maximum delay. While this approach eliminates timing differences, it doesn’t eliminate resource contention based attacks. Obfuscation techniques work by injecting immense amounts of noise into the execution of the program, through the use of extraneous memory accesses and path executions, *e.g.*, [21]. While these techniques do attenuate timing leakage, they do so with enormous slowdown.

B. Eliminating Timing Attacks with Ozone

In this work, we build the first execution resource, called **Ozone**, that exhibits **zero timing leakage** on a modern microarchitecture. Our approach is low cost, and it draws upon earlier work in quelling control and data timing leakage, but

²A control hammer is a CFG construct where control diverges from a single point based on a predicate and then reconverges again to a single point. This construct is typically formed by IF and SWITCH statements.

also introduces new techniques to silence microarchitectural timing leakage. In the Ozone execution environment, vulnerable codes, such as an AES algorithm, will execute with a fixed (and well known) latency regardless of inputs. To achieve zero timing leakage, Ozone places restrictions on the code it can execute (e.g., no CFG hammocks), thus, Ozone targets small applications and/or security-sensitive portions of applications.

Specifically, Ozone achieves zero timing leakage using the following measures to eliminate each of the three properties that enable timing attacks: *i)* the Ozone compiler guarantees that the instruction trace of an Ozone thread has a fixed number of instructions executing on a fixed control path, making the thread's execution independent of input data, *ii)* Ozone codes execute under a special hardware thread that gains exclusive access to a single core's resources for a fixed (and limited) number of cycles during which it cannot be interrupted, *iii)* memory access is limited to pre-allocated cache lines in L1 instruction and data caches, and *iv)* all Ozone threads begin execution with a known fixed microarchitectural state.

The ultimate challenge for Ozone is to demonstrate that its restrictive execution environment is capable of running codes that are vulnerable to timing-based side-channel attacks. To this end, we show in this work that we can efficiently execute a wide range of known-vulnerable codes inside the Ozone execution environment. *To our knowledge, our benchmarks represent the largest collection of known side-channel vulnerable codes analyzed to date, by a wide margin.* Indeed, the breadth of the codes analyzed makes a strong case for the Ozone approach to stopping timing-based side-channel attacks.

Specifically, we make the following novel contributions:

- We begin by demonstrating the significant value of zero timing leakage over state-of-the-art mitigation techniques that only attenuate timing leakage. We show, using the decade-old Bernstein's cache attack, that state-of-the-art timing side-channel mitigations are half measures.
- We detail the first *low-cost implementation* of an execution mode for modern microarchitectures with *zero timing leakage*, via the Ozone execution environment. In addition, we demonstrate the utility of our approach by porting a broad array of side-channel vulnerable application components into the Ozone execution environment.
- We analyze the security and efficiency of the Ozone zero timing leakage execution capability in the context of a modern microarchitecture, by examining the performance of the Ozone execution capability on the Gem5 detailed microarchitecture simulator.

The remainder of this paper is organized as follows. In Section 2, we demonstrate why state-of-the-art approaches to attenuate timing leakage are not sufficient to protect program secrets. Section 3 discusses the requirements for zero timing leakage execution and details the Ozone architectural enhancements. Section 4 gives details about the implementation of Ozone, and performs security analysis of the Ozone execution resource. Section 5 examines related work, and finally we conclude in Section 6.

II. WHY ATTENUATING TIMING LEAKAGE IS NOT ENOUGH

A number of techniques have been proposed to mitigate timing-based side-channel attacks. One popular mitigation technique is isolating security sensitive code execution. This is done by allocating private resources to security sensitive applications [6], [16], [24]. In [16], [24] hardware and software mechanisms are used to partition caches into private regions to mitigate cache-based timing side-channel attacks. Another recent work by Braun *et. al.* [6] targets both control-flow and resource contention based side-channel attacks. The method

uses time padding through delay loops to account for time variations due to different program paths. To provide protection from cache-based timing attacks, the method reserves a core for the duration of execution of a security sensitive function, and it utilizes page coloring to reserve L3 cache resources.

While these techniques provide significant attenuation of timing leakage, they all still leak some timing information, in particular, due to their inability to fully control how microarchitectural state affects program timing. *To demonstrate that even small amounts of timing leakage creates significant vulnerabilities, we show that a state-of-the-art timing leakage mitigation technique is still susceptible to the decade-old Bernstein cache attack [3].* We also show in this section that the Ozone execution environment, detailed in the following section, is not vulnerable to Bernstein's cache attack (or any other timing-based attack).

The Bernstein cache attack works by repeatedly running an AES encryption algorithm while only changing a single byte of input plain text and then inferring key information based on how that single change affects cache performance. As such, there is no external adversary to build protections against, instead, the AES code's cache experiences interference from itself (in the form of capacity and conflict misses), which results in timing variations that expose key information.

Fig. 1. Equations for First Round AES Encryption.

Legend : $K = 16$ byte key, $n = 16$ byte plain text

$$s[i] = K[i] \oplus n[i], \text{ where } i = 0 \text{ to } 15$$

$$t_0 = T_0[s[0]] \oplus T_1[s[5]] \oplus T_2[s[10]] \oplus T_3[s[15]] \oplus x_0$$

$$t_1 = T_0[s[4]] \oplus T_1[s[9]] \oplus T_2[s[14]] \oplus T_3[s[3]] \oplus x_1$$

$$t_2 = T_0[s[8]] \oplus T_1[s[13]] \oplus T_2[s[2]] \oplus T_3[s[7]] \oplus x_2$$

$$t_3 = T_0[s[12]] \oplus T_1[s[1]] \oplus T_2[s[6]] \oplus T_3[s[11]] \oplus x_3$$

Specifically, the attack carefully manipulates indexing into the AES T-tables. Figure 1 shows the equations for the first round of AES encryption as implemented in the OpenSSL cryptographic library [19]. The implementation uses four T-tables ($T_0 - T_3$). Accesses to these tables are derived from the AES key. In the equations, (K) is the encryption key and (n) is an input data to be encrypted. The table indexes are created by XOR'ing a key byte ($K[i]$) with a plain text byte ($n[i]$).

To demonstrate how even attenuated timing leakage leaves programs vulnerable, we perpetrated Bernstein's cache attack on an AES algorithm with state-of-the-art protections [6] while running on a simulated processor model. The AES program's protections and cache configuration are similar to the one's used in [6]. The program is protected such that it contains no input-dependent control, including a fixed number of loop iterations. In addition, the AES algorithm starts execution with a flushed cache and reset branch predictor state. The simulated caches are a 32KB L1 data cache, a 32KB instruction cache and a 256KB L2 shared cache. We used the Gem5 microarchitectural simulator for our experiments [4]. The AES encryption process is the only running process and the time measurements are total execution times in cycles.

Figure 2 shows relative total execution time for AES encryption as a function of $s[0] = (K[0] \oplus n[0])$. Each point in the graph is obtained by averaging the total execution time for 2^{10} 128-bit keys with a fixed value of $n[0]$. The total execution times are given relative to the average execution time for all values of $n[0]$. On the graph, we can see that there is a slight variation in execution time across different values of $n[0]$. *This timing variation was sufficient to allow us to fully recover the AES secret key.* For example, from the graph, the maximum




Fig. 2. Relative Execution Time of AES Encryption: The graph shows the sorted relative total execution times for AES encryption as the value of an input data byte, $n[0]$, varies from 0 to 255, for a state-of-the-art protection mechanism [6] and Ozone.




Fig. 3. The Ozone Compiler: The Ozone compiler, built as a back-end phase of LLVM, ensures that the code marked for Ozone execution *i*) has input-independent control flow via if-conversion and control verification, and *ii*) only accesses scratchpad memory.

execution time is observed when $n[0]$ is 6. Therefore we can conclude that, for any unknown key byte $K[0]$, the maximum execution time is observed at $K[0] \oplus n[0] = 6$. By choosing different values for $n[0]$ and measuring total execution time, we could infer the value of $K[0]$. We then repeated this process on the remaining key bytes to fully recover the AES secret key.

Figure 2 also shows that the Ozone execution, running with the same inputs as the baseline experiment, executes precisely the same number of cycles regardless of changes to the input. This property hold for both changes in the plain text and the secret key values. Consequently, there is zero timing leakage for the Ozone executions, and it is not possible to implement Bernstein’s attack on an Ozone execution.

To summarize, in this section we showed that current mitigation techniques can only limit timing side channels due to the inherent variable latencies associated with microarchitectural features such as caches. We also showed that even highly attenuated timing leakage can still reveal sensitive information. As such, to fully stem the vulnerabilities of timing-based side-channel attacks, an architecture that eliminates all sources of timing variation is required.

III. THE OZONE ARCHITECTURE

The goal of Ozone is to eliminate all sources of timing side channels and achieve zero timing leakage. This is achieved by ensuring that the execution of Ozone code (*i.e.*, the portion of code that we want protected from timing-based side-channel attacks) executes the same number of cycles regardless of its inputs. Moreover, the code’s run time cannot change due to the activities of other threads and processes. By doing this, an attacker will not be able to gain any input-related information about code executing inside the Ozone execution resource. In this section we will describe how the Ozone architecture achieves zero timing leakage.

A. Threat Model

In this work, we model the adversary as an unprivileged process that uses time variation in program execution to extract sensitive information. We assume the adversary has access to fine-grained timing information and microarchitectural statistics such as CPU cycle count and cache miss rates through performance monitoring counters. We assume the hardware, hypervisor and operating system are trusted so that the adversary doesn’t have direct read-write access to memory that is not allocated to it. The system compiler is not trusted, instead a small trusted verifier is used to verify Ozone code generated by the compiler. Finally, we assume the attacker will only pursue timing-based side-channel attacks, thus we consider other side channels, such as power and electromagnetics, beyond the scope of this work.

B. Removing Input-Related Code Variability

The first step in achieving zero leakage is eliminating all sources of variable latency within the Ozone code. Assuming fixed-latency instructions, this can be achieved by dynamically executing the same trace of instructions independent of input data. To guarantee the same dynamic trace, Ozone code control flow must be independent of its inputs, which means *i*) control hammers (*e.g.*, IF statements) are not allowed, and *ii*) loops must execute a constant fixed number of iterations.

Control hammers are removed by the Ozone compiler via if-conversion of CFG hammers, shown in Figure 3. With if-conversion, all operations (*except stores*) on either side of a conditional statement are executed regardless of the condition. Store operations need special treatment because they must avoid the potential side-effects of the false-predicate code. Stores are handled using a special *conditional assignment function*, CMOV. Listings 1 to 4 show two examples of if-conversions that are common. Listing 1 shows a control hammer with stores in both control paths. In this example, prior to if-conversion, one of the two operations in the *if* and *else* bodies are evaluated based on the condition $input == 0$. After if-conversion (as shown in Listing 2), both operations are evaluated and a conditional assignment function is used to select the final value to store. Listing 3 shows an unbalanced control hammer (*if-then construct*), where only one path performs a store operation. In this case, the conditional branch is eliminated by evaluating the statement inside the *if* body and by conditionally storing the new result or the old value of y based on the condition (as shown in Listing 4).

```
if (input == 0) {
    y = a2;
} else {
    y = a2 * b;
}
```

Listing 1. If-Then-Else Conversion - Original Code

```
tv = a2;
fv = a2 * b;
pred = (input == 0);
y = CMOV(pred, tv, fv);
```

Listing 2. If-converted Code with Fixed Control

```
if (input == 0) {
    y = a * b;
}
```

Listing 3. If-Then Conversion - Original Code

```
tv = a * b;
pred = (input == 0);
y = CMOV(pred, tv, y);
```

Listing 4. If-converted Code with Fixed Control

```
CMOV (pred, tv, fv) {
    mov    tv,%rax //Load true value
    cmp   $0,pred //Check if false
    cmov  fv,%rax //Load false value?
    return %rax } //Return selected value
```

Listing 5. Conditional Store Function for x86




Fig. 4. The Ozone Architectural Extensions: The Ozone execution resource ensures zero timing leakage for execution on a modern microarchitecture using a few minor new components (green) and extensions of existing units (green/blue) in an existing core design.

The conditional assignment function (CMOV) used in the if-conversions takes three inputs: a predicate value (*pred*), a true value (*tv*) and a false value (*fv*), and returns either the true value or the false value based on the value of the predicate. It uses conditional move instructions such as the *CMOVcc* instruction from the x86 architecture [10] to eliminate conditional branches. Listing 5 shows an example implementation of our conditional store for the x86 architecture.

After if-conversion by the Ozone compiler, all operations are performed regardless of their predicate conditions. In doing so, there can be cases where unchecked operations could produce exceptions such as division-by-zero. The Ozone compiler relies on the programmer to make sure that such faulty operations do not occur. Since stores in the Ozone execution environment always target pre-allocated cache lines (detailed below), these stores are not subject to page faults, thus making it trivial to ensure fault-free code.

Figure 3 summarizes the Ozone compilation process. The input to the compiler is a programmer-annotated source code. Programmers annotate functions they want protected from timing-based side-channel attacks. In the first stage of the compilation process, conditional branches are eliminated by if-conversion. If-conversion is done on all annotated functions and functions called from annotated functions. Then, the control verification stage verifies that there are no conditional branch instructions in the Ozone code except for loop conditions with a fixed number of iterations and no early exits through *break* and *continue* statements. Finally, all verified Ozone code and associated data is placed in separate code and data sections that map to scratchpad memory. At the end of compilation, code with a fixed control path and data accesses to a scratchpad memory is produced; with support of the Ozone microarchitecture, this code will execute with a fixed latency regardless of its inputs.

C. Eliminating Timing Leakage with Ozone

Even after eliminating input-related control flow in Ozone code, the code will still be vulnerable to timing-based side-channel attacks on a traditional microarchitecture due to contention caused by resource sharing.

Ozone eliminates contention by *i*) gaining exclusive access to the core's execution units during execution of Ozone code, *ii*) using pre-allocated cache lines from the L1 instruction and data caches and *iii*) ensuring all Ozone threads start execution with a known fixed microarchitectural state.

Figure 4 shows the microarchitecture of a processor with an Ozone hardware thread. The processor components are divided into three groups. The first group of resources is exclusively used by the Ozone thread (shown in green in Figure 4) including the Ozone architectural registers and the watch-dog timer (WDT). The second group, shown as half-green and

half-blue, is used by the Ozone execution resource exclusively for the lifetime of an Ozone thread. These resources include portions of L1 instruction and data caches that are locked to the Ozone thread and the Ozone branch predictor. In the processor each L1 instruction and data cache line is augmented with a single *lock bit*. When the lock bit is set, the corresponding line is locked to the current running Ozone thread and can not be evicted, even by the Ozone thread. Furthermore, the caches in the processor are exclusive so that an attacker can not attempt to indirectly evict data from the L1 caches by evicting the last-level cache lines. Once all Ozone threads are destroyed, the cache lines used by the Ozone threads are flushed, allowing their reuse by the rest of the system. The final group of resources, shown in blue in the figure, are explicitly not used by Ozone threads as contention on these resources could enable timing attacks. These off-limit resources include the main branch predictor, caches, other thread states, and DRAM.

An application that wants to use the Ozone resource, first creates an Ozone thread by specifying the resources it requires (*ozone_thread_create*). The resources requested include instruction and data memory sizes (including maximum stack space required for Ozone code execution), the fixed number of cycles required to execute the Ozone code, and a pointer to the Ozone code to load into the L1 instruction cache. Following the creation of a new Ozone thread, the OS allocates instruction and data memory space (including stack space) from the L1 instruction and data caches, zeros out the allocated cache lines, copies all read-only data and the Ozone code to the cache lines, and sets the lock bits of the lines. The OS then returns a handle to the Ozone thread, which is used by the application to invoke the thread. Once the Ozone thread starts executing, the locked cache lines can not be evicted by other threads or the Ozone thread itself. Furthermore the Ozone thread is not allowed to use new cache lines other than the pre-allocated ones. In addition to guaranteeing constant time execution, this ensures that other processes cannot learn about memory locations accessed by the Ozone thread.

When the Ozone thread is invoked by the main program (via *ozone_thread_invoke*), the processor switches to the Ozone thread context. This thread switch forces a flush of the processor pipeline (to eliminate reservation station contention with the previous thread), and then enables the Ozone branch predictor with a fixed initial state. If the Ozone branch predictor is a static predictor (*e.g.*, predict always-taken), this predictor initialization step can be skipped. The hardware context is then switched to the Ozone thread and execution begins. A watchdog timer (WDT) in the processor is started that keeps track of the expected number of execution cycles for the Ozone thread invocation. When the WDT expires, the timer interrupts the Ozone thread, stopping execution. If execution completed in the same cycle the WDT expires, the result of the execution is returned, otherwise, the Ozone thread is terminated as it is running too long or not long enough.

IV. EXPERIMENTAL EVALUATION

A. Ozone Implementation

The Ozone compiler is composed of if-converter and control verifier stages. The compiler stages are implemented as IR passes on the x86-targeted LLVM compiler [13]. The if-conversion is implemented by replacing stores by a call to the conditional assignment function as detailed in Section 3.

The Ozone microarchitecture is modeled using the Gem5 microarchitectural simulator [4]. Table I lists the configuration for the baseline out-of-order core. Ozone is integrated into the simulator as a special thread context, with the state detailed

TABLE I
BASELINE MICROARCHITECTURE CONFIGURATION.

L1 Instruction Cache Size	32kB
L1 Data Cache Size	32kB
L2 Unified Cache Size	256kB
Cache Associativity	4-way
Issue Type	Out-of-order
Issue Width	8
Branch Predictor Type	Tournament
Branch Predictor Size	56kB

TABLE II
ANALYZED BENCHMARKS WITH SCRATCHPAD MEMORY REQUIREMENTS.
The table also lists references to attacks on each benchmark application.

Benchmark	Instruction Memory Size	Data Memory Size	Attack Examples
AES-CBC	6.627K	12.16K	[3], [5], [23]
AES-XTS	6.382K	12.16K	[3], [5], [23]
GDK-Keymap	1.692K	6.728K	[9]
RSA	13.909K	13.311K	[11], [25], [26]
SHA512	3.308K	4.288K	[14]

in Section 3. Additionally, an always-taken predictor is implemented for stateless prediction of Ozone code branches, and instruction and data scratchpads are integrated into the simulator at a fixed address.

B. Benchmark Applications

The benchmark applications used in our evaluation constitute security sensitive applications that have previously been subject to side-channel attacks. They are adapted directly from widely used libraries: OpenSSL [19], GDK [8] and glibc [22]. *Our intent with this benchmark suite is to demonstrate the broad utility of the Ozone execution environment, thus, we have ported into it all of the codes we could find that have been previously attacked with timing-based side channels.* The benchmark applications include AES encryption, RSA decryption, SHA512 hash algorithm, and a key-mapping function (GDK-Keymap). Table II lists references to timing side-channel attacks for each of the benchmark applications.

Table II shows the instruction and data memory requirements for the benchmarks, including the maximum stack space required. As can be observed from the table, the programs and their data can easily fit into L1 caches. For these codes, a 32kB L1 instruction cache and a 64kB L1 data cache will hold all of these algorithms simultaneously (and eliminate the need for context switch saves), and a 16kB L1 instruction and data caches will comfortably hold any single application.

C. Performance and Security Evaluation

We evaluated the security and performance of our benchmark codes running on the baseline out-of-order microarchitecture and in the Ozone execution resource across many random inputs. The inputs were varied in the following manner: AES-CBC encrypts with 1024 random keys with random 128-bit inputs, AES-XTS encrypts with 1024 random keys with random 256-bit inputs, GDK-keymap maps all possible 784 keyboard codes, RSA decrypts a random 128-byte message with 1024 random 1024-bit keys, and SHA512 hashes random inputs from length 1 to 128 bytes.

Figure 5 shows the results of the security analysis of the Ozone execution resource. The graphs show the performance of the baseline microarchitecture (in blue) and that of the Ozone execution (in orange) across a wide range of inputs, with the results sorted and plotted from fastest execution to

TABLE III
COST ASSESSMENT. This table estimates the bit-area cost of the Ozone execution resource.

Addition	Cost (in bits)
Ozone mode bit	1
Ozone thread context	80
L1 instruction cache lock bits	512 lines x 1 bit
L1 data cache lock bits	512 lines x 1 bit
Expected Total Cost: 1k bits	

the slowest. For each of the baseline experiments, the branch predictor and caches are flushed prior to the start of execution.

It is interesting to note the timing leakage characteristics of the baseline out-of-order executions. In all cases, there is evidence of the performance of the program changing as the inputs are varied. As expected, however, the Ozone executions exhibit zero timing leakage, by demonstrating a fixed execution latency across all inputs. The Ozone executions are also immune to aggressor thread perturbations since they utilize static always-taken branch speculation and use pre-allocated locked cache lines, all the while maintaining exclusive access to hardware resources for the duration of the execution.

For most of the codes, exclusive access to the resources (which also disables interrupts for the core) is a short duration, on the order of 500-7500 cycles. The one outlier is RSA, which requires more than 9.5M cycles. This would likely be too long to completely ignore interrupts and exceptions. Thus, a system integrating Ozone execution would likely dedicate it to only a subset of cores, leaving the remaining cores in the system to provide timely response to interrupts and exceptions.

The performance of the Ozone execution compares favorably to the baseline out-of-order executions. In more than half of the cases (3 of 5), the Ozone execution was more efficient than the baseline execution. This is possible due to faster L1 cache accesses (*i.e.*, no cache misses), more responsive branch prediction (*i.e.*, the baseline tournament predictor takes much longer to warm up compared to Ozone's stateless always-taken predictor), and fewer difficult-to-predict if-statements (which are removed during if-conversion). Conversely, the executions of *RSA* and *SHA512* were less efficient. In the case of *RSA*, this was primarily due to significant overheads incurred during if-conversion, which requires that both the true and false code of an IF statement be executed for all occurrences. In the case of *SHA512*, the Ozone implementation pads out the password being hashed to always 128 bytes, thus, it performs the worst-case length hash for all inputs. This extra work is unavoidable if one wants to conceal password length.

D. Ozone Resource Cost Assessment

Table III estimates the bit-area costs of implementing the Ozone execution resource. The table shows the various components necessary to implement Ozone threads and their cost in bits. As shown in the table, the Ozone resource requirements are quite modest, requiring only about 1k bits for our baseline microarchitecture.

V. CONCLUSION

In conclusion, this work pushes the state-of-the-art in timing-based side-channel resistant execution forward with the Ozone execution resource. Unlike previous work, Ozone implements a *zero timing leakage* execution capability, and it does so with low area cost and significantly less performance impact than previous non-zero timing-leakage proposals (no more than a 30% performance loss for a modern microarchitecture). Our approach is to map a carefully prepared input-independent code sequence, built with the Ozone compiler, to a microarchitecture with stateless Ozone branch predictors and

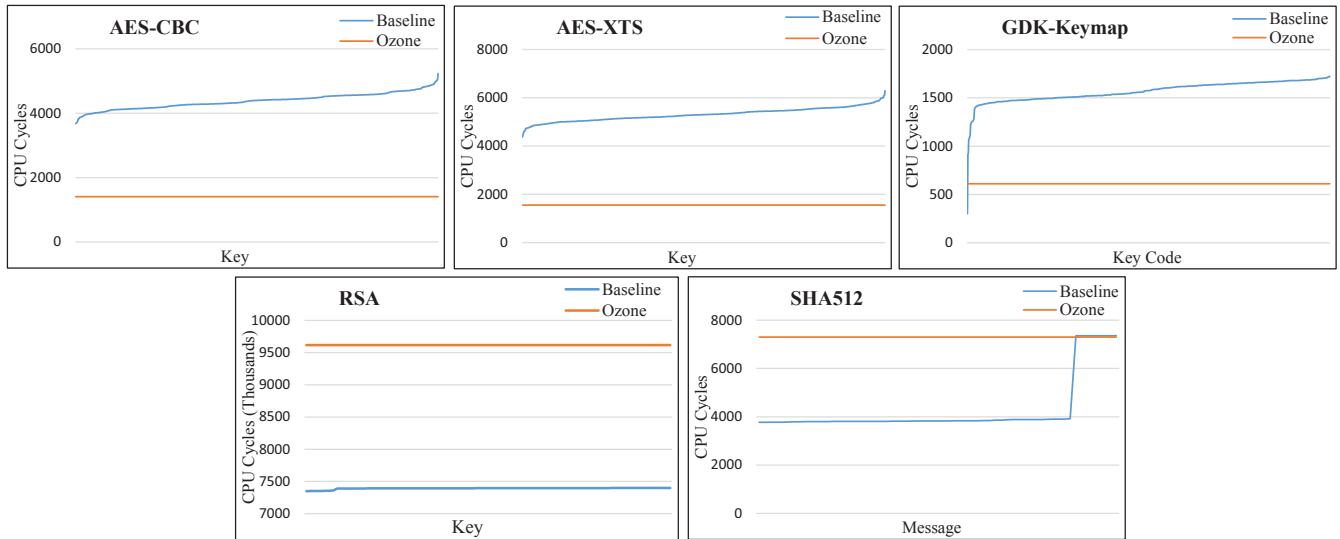


Fig. 5. Security Analysis of Ozone Executions: Each graph shows for each of the five benchmarks the performance variability across a wide range of inputs. The baseline architecture performance (in blue) is shown for varied inputs sorted from fastest to slowest. The Ozone lines are always horizontal, meaning that there is zero timing leakage as the timing is fixed for all tested inputs.

pre-allocated instruction and data cache lines. In this effort, we examine five benchmarks pulled from the literature as being particularly susceptible to side-channel attacks. The Ozone design allows these codes to efficiently execute with a fixed number of cycles, regardless of input, on even a complex microarchitecture. By eliminating all sources of timing variation in the Ozone environment, codes ported to it can rest assured that their executions do not leak secrets via timing channels.

Looking ahead, we see a great opportunity to address additional side channels with the Ozone execution resource. In particular, if the Ozone execution resource were moved into a physical co-processor, it could be designed with circuits to minimize power side channels, while still retaining its capability to execute code with zero timing leakage.

VI. ACKNOWLEDGMENTS

This work was supported in part by C-FAR, one of the six STARnet Centers, sponsored by MARCO and DARPA.

REFERENCES

- [1] Onur Aciicmez, Çetin Kaya Koç, and Jean-Pierre Seifert. On the Power of Simple Branch Prediction Analysis. In *Proceedings of the 2Nd ACM Symposium on Information, Computer and Communications Security, ASIACCS '07*, pages 312–320, New York, NY, USA, 2007. ACM.
- [2] Onur Aciicmez and Werner Schindler. A Vulnerability in RSA Implementations Due to Instruction Cache Analysis and Its Demonstration on OpenSSL. In *Proceedings of the 2008 The Cryptographers' Track at the RSA Conference on Topics in Cryptology, CT-RSA'08*, pages 256–273, Berlin, Heidelberg, 2008. Springer-Verlag.
- [3] D. Bernstein. Cache-timing Attacks on AES. preprint, 2005.
- [4] Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood. The Gem5 Simulator. *SIGARCH Comput. Archit. News*, 39(2):1–7, August 2011.
- [5] Joseph Bonneau and Ilya Mironov. Cache-collision Timing Attacks Against AES. In *Proceedings of the 8th International Conference on Cryptographic Hardware and Embedded Systems*.
- [6] B. A. Braun, S. Jana, and D. Boneh. Robust and Efficient Elimination of Cache and Timing Side Channels. *ArXiv e-prints*, May 2015.
- [7] B. Coppens, I. Verbauwhede, K. D. Bosschere, and B. D. Sutter. Practical Mitigations for Timing-Based Side-Channel Attacks on Modern x86 Processors. In *Security and Privacy, 2009 30th IEEE Symposium on*, pages 45–60, May 2009.
- [8] GDK 3 Reference Manual. <https://developer.gnome.org/gdk3/>.
- [9] Daniel Gruss, Raphael Spreitzer, and Stefan Mangard. Cache Template Attacks: Automating Attacks on Inclusive Last-level Caches. In *Proceedings of the 24th USENIX Conference on Security Symposium, SEC'15*, pages 897–912, Berkeley, CA, USA, 2015. USENIX Association.
- [10] Intel Inc. *Intel® 64 and IA-32 Architectures Software Developer's Manual, Volume 3 (3A, 3B & 3C): System Programming Guide*. June 2015.
- [11] Paul C. Kocher. Timing Attacks on Implementations of Diffie-Hellman, RSA, DSS, and Other Systems. In *Proceedings of the 16th Annual International Cryptology Conference on Advances in Cryptology, CRYPTO '96*, pages 104–113, London, UK, UK, 1996. Springer-Verlag.
- [12] L., A. Georges, and I. Verbauwhede. Exploiting Hardware Performance Counters. In *Fault Diagnosis and Tolerance in Cryptography, 2008. FDTC '08. 5th Workshop on*, pages 59–67, Aug 2008.
- [13] L., A. Georges, and Vikram Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the 2004 International Symposium on Code Generation and Optimization (CGO'04)*, Palo Alto, California, Mar 2004.
- [14] N. Lawson. Timing Attack in Google Keyczar library. *rdist Blog*, May 2009.
- [15] C. Liu, A. Harris, M. Maas, M. Hicks, M. Tiwari, and E. Shi. GhostRider: A Hardware-Software System for Memory Trace Oblivious Computation. *SIGPLAN Not.*, 50(4):87–101, March 2015.
- [16] F. Liu, Q. Ge, Y. Yarom, F. Mckeen, C. Rozas, G. Heiser, and R. B. Lee. CATalyst: Defeating Last-level Cache Side Channel Attacks in Cloud Computing. In *2016 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 406–418, March 2016.
- [17] F. Liu, Y. Yarom, Q. Ge, G. Heiser, and R. B. Lee. Last-Level Cache Side-Channel Attacks are Practical. In *Security and Privacy (SP), 2015 IEEE Symposium on*, pages 605–622, May 2015.
- [18] David Molnar, Matt Piotrowski, David Schultz, and David Wagner. The Program Counter Security Model: Automatic Detection and Removal of Control-flow Side Channel Attacks. In *Proceedings of the 8th International Conference on Information Security and Cryptology, ICISC'05*, pages 156–168, Berlin, Heidelberg, 2006. Springer-Verlag.
- [19] OpenSSL. <https://www.openssl.org/>.
- [20] Dag Arne Osvik, Adi Shamir, and Eran Tromer. Cache Attacks and Countermeasures: The Case of AES. In *Proceedings of the 2006 The Cryptographers' Track at the RSA Conference on Topics in Cryptology, CT-RSA'06*, pages 1–20, Berlin, Heidelberg, 2006. Springer-Verlag.
- [21] Ashay Rane, Calvin Lin, and Mohit Tiwari. Raccoon: Closing Digital Side-channels Through Obfuscated Execution. In *Proceedings of the 24th USENIX Conference on Security Symposium, SEC'15*, pages 431–446, Berkeley, CA, USA, 2015. USENIX Association.
- [22] The GNU C Library. <https://www.gnu.org/software/libc/>.
- [23] Eran Tromer, Dag Arne Osvik, and Adi Shamir. Efficient Cache Attacks on AES, and Countermeasures. *J. Cryptol.*, 23(2):37–71, January 2010.
- [24] Zhonghong Wang and Ruby B. Lee. New Cache Designs for Thwarting Software Cache-based Side Channel Attacks. In *Proceedings of the 34th Annual International Symposium on Computer Architecture, ISCA '07*, pages 494–505, New York, NY, USA, 2007. ACM.
- [25] Y. Yarom and K. Falkner. FLUSH+RELOAD: A High Resolution, Low Noise, L3 Cache Side-channel Attack. In *Proceedings of the 23rd USENIX Conference on Security Symposium, SEC'14*, pages 719–732, Berkeley, CA, USA, 2014. USENIX Association.
- [26] Y. Yarom, D. Genkin, and N. Heninger. CacheBleed: A Timing Attack on OpenSSL Constant Time RSA. In *Cryptographic Hardware and Embedded Systems - CHES 2016 - 18th International Conference, Santa Barbara, CA, 2016*, pages 346–367, 2016.