

TIKTAG: Breaking ARM’s Memory Tagging Extension with Speculative Execution

Juhee Kim

Seoul National University
kimjuhi96@snu.ac.kr

Jinbum Park

Samsung Research
jinb.park@samsung.com

Sihyeon Roh

Seoul National University
sihyeonroh@snu.ac.kr

Jaeyoung Chung

Seoul National University
jjy600901@snu.ac.kr

Youngjoo Lee

Seoul National University
youngjoo.lee@snu.ac.kr

Taesoo Kim

Samsung Research and
Georgia Institute of Technology
taesoo@gatech.edu

Byoungyoung Lee*

Seoul National University
byoungyoung@snu.ac.kr

Abstract—ARM Memory Tagging Extension (MTE) is a new hardware feature introduced in ARMv8.5-A architecture, aiming to detect memory corruption vulnerabilities. The low overhead of MTE makes it an attractive solution to mitigate memory corruption attacks in modern software systems and is considered the most promising path forward for improving C/C++ software security. This paper explores the potential security risks posed by speculative execution attacks against MTE. Specifically, this paper identifies new TIKTAG gadgets capable of leaking the MTE tags from arbitrary memory addresses through speculative execution. With TIKTAG gadgets, attackers can bypass the probabilistic defense of MTE, increasing the attack success rate by close to 100%. We demonstrate that TIKTAG gadgets can be used to bypass MTE-based mitigations in real-world systems, Google Chrome and the Linux kernel. Experimental results show that TIKTAG gadgets can successfully leak an MTE tag with a success rate higher than 95% in less than 4 seconds. We further propose new defense mechanisms to mitigate the security risks posed by TIKTAG gadgets.

1. Introduction

Memory corruption vulnerabilities present significant security threats to computing systems. Exploiting a memory corruption vulnerability, an attacker corrupts the data stored in a memory, hijacking the control flow or crafting the data of the victim. Such exploitation allows the attacker to execute arbitrary code, escalate its privilege, or leak security-sensitive data, critically harming the security of the computing system.

In response to these threats, ARM Memory Tagging Extension (MTE) has recently been proposed since ARMv8.5-A architecture, which is a new hardware extension to mitigate memory corruption attacks. Technically, MTE provides two hardware primitive operations, (i) *tag* and (ii) *tag check*. A tag operation assigns a tag to a memory location (i.e., a 4-bit tag to each 16-byte memory). Then a tag check operation is performed when accessing the memory, which compares

two tags, one embedded within the pointer to access the memory and the other associated with the memory location to-be-accessed. If these two tags are the same, the access is allowed. Otherwise, the CPU raises a fault.

Using MTE, various mitigation techniques can be developed depending on which tag is assigned or which memory regions are tagged. For instance, MTE-supported memory allocators, such as Android Scudo [4] and Chrome Partition-Alloc [13], assign a random tag for all dynamically allocated memory. Specifically, a memory allocator is modified to assign a random tag for each allocation. Then, a pointer to this allocated memory embeds the tag, and as the pointer is propagated, the tag is accordingly propagated together. When any dynamically allocated memory is accessed, a tag check operation is enforced. As the tags are randomly assigned at runtime, it is difficult for the attacker to correctly guess the tag. Thus tag check operation would statistically detect memory corruptions.

MTE introduces significant challenges for attackers to exploit the memory corruption vulnerability. This is because MTE-based solutions detect a violation behavior close to the root cause of spatial and temporal memory corruptions. Specifically, since MTE ensures that the relationship between a pointer and a memory location is not corrupted, it promptly detects the corruptions—i.e., MTE promptly detects the moment when out-of-bounds access takes place in a heap-overflow vulnerability or when a dangling pointer is dereferenced in use-after-free. This offers strong security advantages to MTE, particularly compared to popular mitigation techniques such as CFI [1, 55, 66], which does not detect memory corruption but detects control-flow hijack behavior (i.e., an exploitation behavior). For these reasons, MTE is considered *the most promising path forward for improving C/C++ software security* by many security experts [7, 48], since its first adoption with the Pixel 8 device in October 2023.

In this paper, we study if MTE provides the security assurance as it is promised. In particular, we analyzed if speculative execution attacks can be a security threat to breaking MTE. To summarize our results, we found that speculative

• *Corresponding author.

execution attacks are indeed possible against MTE, which severely harms the security assurance of MTE. We discovered two new gadgets, named TIKTAG-v1 and TIKTAG-v2, which can leak the MTE tag of an arbitrary memory address. Specifically, TIKTAG-v1 exploits the speculation shrinkage of the branch prediction and data prefetchers, and TIKTAG-v2 exploits the store-to-load forwarding behavior.

To demonstrate the exploitability of real-world MTE-based mitigations, we developed two real-world attacks having distinctive attack surfaces: Google Chrome and the Linux kernel. Our evaluation results show that TIKTAG gadgets can leak MTE tags with a success rate higher than 95% in less than 4 seconds.

Compared to the previous works on MTE side-channels [20, 40], we think this paper makes unique contributions for the following reasons.

First, Project Zero at Google reported that they were not able to find speculative tag leakage from the MTE mechanisms [40]. They concluded that speculative MTE check results do not induce distinguishable cache state differences between the tag check success and failure. In contrast, this paper discovers that tag checks indeed generate the cache state difference in speculative execution.

Second, this paper provides a new perspective on the root cause of the MTE tag leakage. Previously, StickyTags [20], discovered an MTE tag leakage gadget, which we named TIKTAG-v1 gadget, and suspected that its root cause is memory contention on spurious tag check faults. On the contrary, this paper conducts an in-depth analysis and reveals that speculation shrinkage in branch prediction and data prefetching logic are the root cause of the TIKTAG-v1 gadget.

Third, this paper reports new MTE tag leakage gadgets, specifically the variants of TIKTAG-v1 gadget and the new TIKTAG-v2 gadget, along with exploitation against Chrome and the Linux kernel. We further suggest that the exploitation of TIKTAG gadgets can be mitigated in a similar way to the existing transient execution attack mitigations [9, 29, 36, 49], while retaining the security benefits of using MTE.

At the time of writing, MTE is still in the early stages of adoption. Considering its strong security advantage, it is expected that a large number of MTE-based mitigations (e.g., sensitive data protection [28, 30] and data-flow integrity [10, 26, 41, 64]) is expected to be deployed in the near future on MTE-supporting devices (e.g., Android mobile phones). As such, the results of this paper, particularly in how TIKTAG gadgets are constructed and how MTE tags can be leaked, shed light on how MTE-based solutions should be designed or how CPU should be implemented at a micro-architectural level. We have open-sourced TIKTAG gadgets at <https://github.com/compsec-snu/tiktag> to help the community understand the MTE side-channel issues.

Responsible Disclosure. We reported MTE tag leakage gadgets to ARM in November 2023. ARM acknowledged and publicly disclosed the issue in December 2023 [34]. Another research group reported a similar issue to ARM and published their findings [20], which were conducted independently from our work. We reported the speculative vulnerabilities in Google Chrome V8 to the Chrome Security

Team in December 2023. They acknowledged the issues but decided not to fix the vulnerabilities because the V8 sandbox is not intended to guarantee the confidentiality of memory data and MTE tags. Since the Chrome browser currently does not enable its MTE-based defense by default, we agree with their decision to some extent. However, we think that browser security can be improved if MTE-based defenses are deployed with the countermeasures we suggest (§6.1.4). We also reported the MTE oracles in the Pixel 8 device to the Android Security Team in April 2024. Android Security Team acknowledged the issue as a hardware flaw of Pixel 8, decided to address the issue in Android’s MTE-based defense, and awarded a bounty reward for the report.

2. Background

2.1. Memory Tagging Extension

Memory Tagging Extension (MTE) [32] is a hardware extension to prevent memory corruption attacks, available since ARMv8.5-A architecture. MTE has been recently adopted by Pixel 8 [39] since October 2023. MTE assigns a 4-bit tag for each 16 bytes of memory and stores the tag in the unused upper bits of a pointer. During memory access, the tag in the pointer is checked against the tag assigned for the memory region. If the tags match, access is permitted; otherwise, the CPU raises a tag check fault (TCF).

MTE offers three modes—*synchronous*, *asynchronous*, and *asymmetric*—to balance performance and security. Synchronous mode provides the strongest security guarantee, where the tag check fault is synchronously raised at the faulting load/store instruction. Asynchronous mode offers the best performance, where the tag check fault is asynchronously raised at context switches. Asymmetric mode strikes a balance between performance and security, with load instructions operating in synchronous mode and store instructions in asynchronous mode.

Based on MTE, various mitigation schemes can be developed. *deterministic tagging* assigns a globally known tag to each allocation. This approach can deterministically isolate memory regions [31] or detect bounded spatial memory corruptions [20]. *random tagging*, on the other hand, assigns a random tag generated at allocation time. This approach probabilistically prevents spatial and temporal memory errors at per-allocation granularity, with a maximum detection rate of 15/16 (i.e., 1/16 chance of tag collision). Unlike deterministic tagging, random tagging does not reveal the tag information to attackers, requiring them to guess the tag to exploit memory corruption vulnerabilities. Consequently, random tagging is commonly adopted in real-world allocators (e.g., Android Scudo allocator [4], Chrome PartitionAlloc [13]) and Linux Hardware Tag-Based KASAN [24].

2.2. Speculative Execution Attack

A speculative execution attack is a class of attacks that exploit the CPU’s speculative behaviors to leak sensitive

information [22, 29, 36, 42, 70–72, 76]. Spectre [29] and Meltdown [36] are well-known speculative execution attacks, where the attacker speculatively executes the victim code to load data that cannot be accessed during the normal execution (e.g., out-of-bounds access). If the speculatively loaded data affects the cache, the attacker can infer its value by observing the cache state (e.g., cache hit/miss based on access latency). Such speculative information leakage typically requires two attacker’s capabilities: i) controlling the cache state by flushing or evicting cache sets before the victim accesses the data, and ii) measuring the time precisely enough to discern cache hits and misses. Recent studies have extended speculative execution attacks to bypass hardware security features such as Address Space Layout Randomization (ASLR) [16] and Pointer Authentication Code (PAC) [57].

3. Threat Model

We consider a threat model where the target system employs Memory Tagging Extension (MTE) [32] to prevent memory corruption. The allocator in the target system tags each allocation with a *random tag*, and the tag is checked on every memory access. We assume random tagging since it is architecturally designed to improve security [32] and commonly developed in real-world MTE-enabled systems (e.g., Android scudo allocator [4], Chrome PartitionAlloc [13], and Linux Hardware Tag-Based KASAN [24]).

We assume that the attacker possesses knowledge of the memory corruption vulnerabilities in the target system, and aims to exploit the vulnerabilities to gain unauthorized access to the system. From the attacker’s perspective, triggering the vulnerabilities imposes a high probability of crashing the target process with a tag check fault, which notifies the system administrator of the attack. We further detail the specific threat model in real-world attack scenarios (§6).

4. Finding Tag Leakage Gadgets

The security of MTE random tag assignment relies on the confidentiality of the tag information per memory address. If the attacker can learn the tag of a specific memory address, it can be used to bypass MTE—e.g., exploiting memory corruption only when the tag match is expected. In this section, we present our approach to discovering MTE tag leakage gadgets. We first introduce a template for an MTE tag leakage gadget (§4.1) and then present a template-based fuzzing to discover MTE tag leakage gadgets (§4.2).

4.1. Tag Leakage Template

We first designed a template for a speculative MTE tag leakage gadget, which allows the attacker to leak the tag of a given memory address through speculative execution (Figure 1). The motivation behind the template is to trigger MTE tag checks in a speculative context and observe the cache state after the speculative execution. If there is any

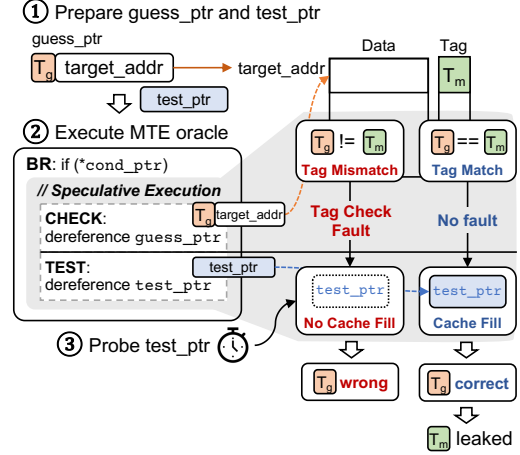


Figure 1: An MTE tag leakage template

difference between tag match and mismatch, the attackers can potentially leak the tag check results and infer the tag value. Since tag mismatch during speculative execution is not raised as an exception, such an attempt is not detected.

We assume the attacker aims to leak the tag T_m assigned to `target_addr`. To achieve this, the attacker prepares two pointers: `guess_ptr` and `test_ptr` (①). `guess_ptr` points to `target_addr` while embedding a tag T_g —i.e., `guess_ptr = (Tg«56)|(target_addr & ~(0xff«56))`. `test_ptr` points to an attacker-accessible, uncached address with a valid tag.

Next, the attacker executes the template with `guess_ptr` and `test_ptr` (②). The template consists of three components in order: BR, CHECK, and TEST. BR encloses CHECK and TEST using a conditional branch, ensuring that CHECK and TEST are speculatively executed. In CHECK, the template executes a sequence of memory instructions to trigger MTE checks. In TEST, the template executes an instruction updating the cache status of `test_ptr`, observable by the attacker later.

Our hypothetical expectation from this template is as follows: The attacker first trains the branch predictor by executing the template with `cond_ptr` storing 1 and `guess_ptr` containing a valid address and tag. After training, the attacker executes the template with `cond_ptr` storing 0 and `guess_ptr` pointing to `target_addr` with a guessed tag, causing speculative execution of CHECK and TEST. If the MTE tag matches in CHECK, the CPU would continue to speculatively execute TEST, accessing `test_ptr` and filling its cache line. If the tags do not match, the CPU may halt the speculative execution of TEST, leaving the cache line of `test_ptr` unfilled. Consequently, the cache line of `test_ptr` would not be filled. After executing the template, the attacker can measure the access latency of `test_ptr` after execution, and distinguish the cache hit and miss, leaking the tag check result (③). The attacker can then brute-force the template executions with all possible T_g values to eventually leak the correct tag of `target_addr`.

Results. We tested the template on real-world ARMv8.5 devices, Google Pixel 8 and Pixel 8 pro. We varied the

number and type of memory instructions in CHECK and TEST, and observed the cache state of `test_ptr` after executing the template. As a result, we identified two speculative MTE leakage gadgets, TIKTAG-v1 (§5.1) and TIKTAG-v2 (§5.2) that leak the MTE tag of a given memory address in both Pixel 8 and Pixel 8 pro.

4.2. Tag Leakage Fuzzing

To automatically discover MTE tag leakage gadgets, we developed a fuzzer in a similar manner to the Spectre-v1 fuzzers [50]. The fuzzer generates test cases consisting of a sequence of assembly instructions for the speculatively executed blocks in the tag leakage template (i.e., CHECK and TEST). The fuzzer consists of the following steps:

Based on the template, the fuzzer first allocates memory for `cond_ptr`, `guess_ptr`, and `test_ptr`. `cond_ptr` and `guess_ptr` point to a fixed 128-byte memory region individually. `test_ptr` points to a variable 128-byte aligned address from a 4KB memory region initialized with random values. Then, the fuzzer randomly picks two registers to assign `cond_ptr` and `guess_ptr` from the available registers (i.e., `x0-x28`). The remaining registers hold a 128-byte aligned address within a 4KB memory region or a random value.

The fuzzer populates CHECK and TEST blocks using a predefined set of instructions (i.e., `ldr`, `str`, `eor`, `orr`, `nop`, `isb`) to reduce the search space. Given an initial test case, the fuzzer randomly mutates the test case by inserting, deleting, or replacing instructions to generate new test cases.

The fuzzer runs test cases in two phases: (i) a branch training phase, with `cond_ptr` storing true and `guess_ptr` containing a correct tag; and (ii) a speculative execution phase, with `cond_ptr` storing false and `guess_ptr` containing either a correct or wrong tag. The fuzzer executes each test case twice. The first execution runs the branch training phase and then the speculative execution phase with the correct tag. The second execution is the same as the first, but the only difference is to run the speculative execution phase with the wrong tag. After each execution, the fuzzer measures the access latency of a cache line and compares the cache state between the two executions. This process is repeated for each cache line of the 4KB memory region. If a notable difference is observed, the fuzzer considers the test case as a potential MTE tag leakage gadget.

Results. We developed the fuzzer and tested it on the same ARMv8.5 devices. As a result, we additionally identified variants of TIKTAG-v1 (§5.1) that utilize linked list traversal. The fuzzer was able to discover the gadgets within 1-2 hours of execution without any prior knowledge of them.

5. TIKTAG Gadgets

In this section, we present TIKTAG gadgets discovered by the tag leakage template and fuzzing (§4). Each gadget featuring unique memory access patterns leaks the MTE tag of a given memory address. TIKTAG-v1 (§5.1) exploits the speculation shrinkage in branch prediction and data

```
BR:    ldr cond, [cond_ptr]
      cbz cond, END
CHECK: ldr r0, [guess_ptr]
      ldr r0, [guess_ptr]
      ... ; >= 2 loads
      GAP: orr test_ptr, test_ptr, #0
      ... ; >= 10 cycles
TEST:  ldr r1, [test_ptr] ; INDEP_LD
      str r1, [test_ptr] ; INDEP_ST
      ldr r1, [test_ptr, r0] ; DEP_LD
      str r1, [test_ptr, r0] ; DEP_ST
END:
```

Figure 2: TIKTAG-v1 gadget




Figure 3: Cache hit rate of `test_ptr` after executing TIKTAG-v1. X-axis shows the length of GAP and Y-axis shows the cache hit rate.

prefetching, and TIKTAG-v2 (§5.2) leverages the blockage of store-to-load forwarding. We further analyze the root cause and propose mitigations at hardware and software levels.

5.1. TIKTAG-v1: Exploiting Speculation Shrinkage

During our experiment with the MTE tag leakage template (§4.1), we observed that multiple tag checks in CHECK influence the cache state of `test_ptr`. With a single tag check, `test_ptr` was always cached regardless of the tag check result. However, with two or more tag checks, `test_ptr` was cached on tag match, but not cached on tag mismatch.

In summary, we found that the following three conditions should hold for the template to leak the MTE tag: (i) CHECK should include at least two loads with `guess_ptr`; (ii) TEST should access `test_ptr` (either load or store, dependent or independent to CHECK); and (iii) CHECK should be close to BR (within 5 CPU cycles) while TEST should be far from CHECK (more than 10 CPU cycles away). If these conditions are met, `test_ptr` is cached on tag match and not cached on tag mismatch. If any of these is not met, `test_ptr` is either always cached or never cached. Based on this observation, we developed TIKTAG-v1, a gadget that leaks the MTE tag of any given memory address.

Gadget. Figure 2 illustrates the TIKTAG-v1 gadget. In BR, the CPU mispredicts the branch result and speculatively

executes CHECK. In CHECK, `guess_ptr` is dereferenced two or more times, triggering tag checks with Tg against Tm. GAP provides the time gap between BR and TEST, and then in TEST, `test_ptr` is accessed. GAP can be filled with various types of instructions, such as computational instructions (e.g., `orr`, `mul`) or memory instructions (e.g., `ldr`, `str`), as long as it provides more than 10 CPU cycles of the time gap.

Experimental Results. We found that TIKTAG-v1 is an effective MTE tag leakage gadget on the ARM Cortex-X3 (core 8 of Pixel 8 devices) in both MTE synchronous and asynchronous modes. We leveraged the physical CPU cycle counter (i.e., `PMCCNTR_EL0`) for time measurement, which was enabled by modifying the Linux kernel. An L1 cache hit is determined if the access latency is less than or equal to 35 CPU cycles. In a real-world setting, the virtual CPU cycle counter (i.e., `CNTVCT_EL0`) is available to the user space with lower resolution, which also can effectively observe the tag leakage behavior.

We experimented to measure the cache hit rate of `test_ptr` after executing TIKTAG-v1. To verify condition (i), we varied the number of loads in CHECK (i.e., `Len(CHECK)`) from 1, 2, 4, and 8. To verify condition (ii), we varied the types of memory access in TEST (i.e., independent/dependent, load/store). To verify condition (iii), we filled GAP with a sequence of `orr` instructions (where each `orr` is dependent on the previous one) and varied its length (i.e., `Len(GAP)`). Figure 3 shows the experimental results. The x-axis represents `Len(GAP)` and the y-axis represents the cache hit rate of `test_ptr` measured over 1,000 trials.

When `Len(CHECK)` was 1, the cache hit rate of `test_ptr` had no difference between tag match and mismatch. `test_ptr` was either always cached (i.e., load access) or cached until a certain threshold and then not cached (i.e., store access). When `Len(CHECK)` was 2 or more, depending on the tag check result, the cache hit rate differed, validating the condition (i). If the tag matched, the cache hit rate was the same as when `Len(CHECK)` was 1. If the tag mismatched, the cache hit rate dropped compared to the tag match (annotated with arrows). This difference was observed in all access types of TEST, validating the condition (ii). The cache hit rate drop was observed after about 10 `orr` instructions in GAP. Further experiments inserting lengthy instructions between BR and CHECK (§E) confirmed that CHECK should be close to BR, validating the condition (iii).

When the tag mismatched, the cache hit was periodic and the period got shorter as `Len(CHECK)` increased. When `Len(CHECK)` is 2, a cache hit occurred in 5 out of 6 trials (83%, arrows (a)). When `Len(CHECK)` is 4, a cache hit occurred in 2 out of 3 trials (66%, arrows b). When `Len(CHECK)` is 8, the pattern differed between load and store accesses (arrows c and d). For store access, a cache hit was observed every 3 trials (33%, arrows c). For load access, the pattern changed over iterations from all cache hits (0%) to a cache hit every 2 trials (50%) (arrows d).

A similar cache hit rate drop was observed when `guess_ptr` points to an unmapped address and generated speculative address translation faults. This further indicates

Test	Branch training phase			Speculative execution phase (tag match / mismatch)			Measure test_ptr	Root cause
	BR	CHECK	TEST	BR	CHECK	TEST		
Baseline	m1	m2	m3	m1	m2	m3	m3 hit	Not speculative execution
				m1	m2	m3	m3 miss (periodic)	
-SE	m1	m2	SB m3	m1	m2	SB m3	m3 hit	Data prefetching
				m1	m2	SB m3	m3 miss (periodic)	
-SE & -DP	m1	m2	SB m3	m1	m2	SB m4	m4 miss (always)	Speculative execution
				m1	m2	SB m4	m4 miss (always)	
-DP	m1	m2	m3	m1	m2	m4	m4 hit	
				m1	m2	m4	m4 miss (periodic)	

Figure 4: TIKTAG-v1 ablation study. SE: Speculative Execution, DP: Data Prefetch. m1 and m2 are memory addresses for `cond_ptr` and `target_addr`, respectively. m3 and m4 are memory address for `test_ptr`. The tests were conducted with `LEN(CHECK)=2` and `TEST=INDEP_LD`

that TIKTAG-v1 can be utilized as an address-probing gadget [18] useful for breaking ASLR.

Root Cause. Analyzing the gadget, we found that tag check results affect the CPU’s data prefetching behavior and the speculative execution. This refutes the previous studies on speculative MTE tag leakage [20, 40], which stated that tag check faults do not affect the speculation execution and did not state the impacts of the data prefetching.

In general, modern CPUs speculatively access memory in two cases: speculative execution [29] and data prefetching [6, 17]. To identify the root cause of TIKTAG-v1 in these two cases, we conducted an ablation study (Figure 4). First, we eliminated the effect of speculative execution by inserting a speculation barrier (i.e., `sb`) between CHECK and TEST. Second, we varied the memory access pattern between branch training and speculative execution phases to eliminate the effect of data prefetching.

In Baseline, no speculation barrier was inserted, and both branch training and speculative execution phases accessed the same addresses in order. In this case, `test_ptr` was cached on tag match, but not cached on tag mismatch. In -SE, a speculation barrier was inserted to prevent TEST from being speculatively executed. Here, the same cache state difference was observed, indicating that the difference in Baseline is not due to the speculative execution at least in this case.

Next, in -SE & -DP, the memory access pattern was also varied in the speculative execution phase to prevent `test_ptr` from being prefetched. As a result, `test_ptr` was always not cached, verifying that the CPU failed to prefetch `test_ptr` due to the divergence in the access pattern. If we compare -SE and -SE & -DP, the cache state difference was observed only when data prefetching is enabled (-SE). Considering the CPU’s mechanism, such a difference seems to be due to data prefetching—i.e., the CPU prefetches data based on the previous access pattern, but skips it on tag mismatch.

Finally, in -DP, we removed the speculation barrier to re-enable the speculative execution of TEST while still varying

<pre>BR: ldr cond, [cond_ptr] CHECK: ldr r0, [guess_ptr] ldr r0, [guess_ptr] GAP: ldr ptr, [ptr] ; ptr1 ldr ptr, [ptr] ; ptr2 ldr ptr, [ptr] ; ptr3 TEST: ldr ptr, [ptr] ; ptr4 END:</pre>	<pre>BR: ldr cond, [cond_ptr] CHECK: ldr r0, [guess_ptr] ldr r0, [guess_ptr] GAP: ldr ptr, [ptr] ; ptr1 END: ldr ptr, [ptr] ; ptr2 ldr ptr, [ptr] ; ptr3 TEST: ldr ptr, [ptr] ; ptr4</pre>
(a) variant 1	(b) variant 2

Figure 5: TIKTAG-v1 gadget variants

the memory access pattern. In this case, the difference is observed again between tag match and mismatch. Comparing -DP and -SE & -DP, we can conclude that the speculative execution is also the root cause of the cache difference—i.e., the CPU halts speculative execution on tag check faults.

We suspect that the CPU optimizes performance by halting speculative execution and data prefetching on tag check faults. A relevant patent filed by ARM [8] explains that the CPU can reduce speculations on *wrong path events* [5], which are events indicating the possibility of branch misprediction, such as spurious invalid memory accesses. By detecting branch misprediction earlier, the CPU can save recovery time from wrong speculative execution and improve the data prefetch accuracy by not prefetching the wrong path-related data. Since these optimizations are beneficial in both MTE synchronous and asynchronous modes, we think that the tag leakage behaviors were observed in both MTE modes.

We also think there is a time window to detect wrong path events during speculative execution upon branch prediction, which seems to be 5 CPU cycles. As explained in the patent [8], the CPU may maintain speculation confidence values for speculative execution and data prefetching. We think the CPU reduces the confidence values on tag check faults, halts speculation if it drops below a certain threshold, and restores it to the initial level. This reasoning explains the periodic cache miss of `test_ptr` on tag mismatch, where the confidence value is repeatedly reduced below the threshold (i.e., cache miss) and restored (i.e., cache hit). In addition, when TEST is store access, the speculation barrier made `test_ptr` always not cached. This indicates that the CPU does not prefetch data for store access, thus the speculation window shrinkage is the only root cause in such cases.

Variants. Running the tag leakage fuzzer (§4.2), we discovered variants of TIKTAG-v1 leveraging linked list traversal (Figure 5). Before the gadget, a linked list of 4 instances is initialized, where each instance points to the next instance (i.e., `ptr0` to `ptr3`), and the cache line of the last instance (`ptr3`) is flushed. The gadget traverses the linked list by accessing `ptr0` to `ptr3` in order, where TEST accesses `ptr3` only if the branch result is true. After the gadget, the cache hit rate of `ptr3` is measured.

In the first variant (Figure 5a), TEST is located in the true branch of the conditional branch BR. In the second variant (Figure 5b), TEST is located out of the conditional branch path, where both the true and false branches merge. In both variants, if the tag matched in CHECK, `ptr3` was

```

; SLIDE (0..40 inst)
BR:    ldr cond, [cond_ptr]
      cbz cond, END
CHECK: str test_ptr, [guess_ptr]
      ; GAP (0..3 inst)
      ldr ptr, [guess_ptr]
TEST:  ldr r0, [ptr, r0] // DEP_LD
      str r0, [ptr, r0] // DEP_ST
END:
```

Figure 6: TIKTAG-v2: A speculative tag check fault blocks store-to-load forwarding

cached, but not cached if the tag mismatched, as in the original TIKTAG-v1 gadget (Figure 3). We think the root cause is the same as the original gadget—i.e., the CPU changes speculative execution and data prefetching behaviors on tag check faults. Moreover, the variants can be effective in realistic scenarios like linked list traversal, and the tag leakage can also be observed from the memory access outside the conditional branch’s scope.

Mitigation. TIKTAG-v1 exploits the speculation shrinkage on tag check faults in speculative execution and data prefetching. To prevent MTE tag leakage at the micro-architectural level, the CPU should not change the speculative execution or data prefetching behavior on tag check faults.

To prevent TIKTAG-v1 at the software level, the following two approaches can be used.

i) *Speculation barrier*: Speculation barrier can prevent the TIKTAG-v1 gadget from leaking the MTE tag. If TEST contains store access, placing a speculation barrier (i.e., `sb`) or instruction synchronization barrier (i.e., `isb`) makes `test_ptr` always not cached, preventing the tag leakage. If TEST contains load access, placing the barrier before TEST does not prevent tag leakage, but placing it before CHECK mitigates tag leakage, since speculative tag check faults are not raised.

ii) *Prevention of gadget construction*: TIKTAG-v1 requires that tag check faults are raised within a time window from the branch. By inserting a sequence of instructions before CHECK to extend the time window, the CPU does not reduce the speculations and `test_ptr` is always cached (§E).

We note that inserting additional barriers or instructions may introduce performance overhead, and their effectiveness should be evaluated in the context of the target software.

5.2. TIKTAG-v2: Exploiting Store-to-Load Forwarding

Inspired by Spectre-v4 [43] and LVI [71] attacks, we experimented with the MTE tag leakage template to trigger store-to-load forwarding behavior [65]. As a result, we discovered that store-to-load forwarding behavior differs on tag check result if the following conditions hold: (i) CHECK triggers store-to-load forwarding, and (ii) TEST accesses memory dependent on the forwarded value. If the tag matches in CHECK, the memory accessed in TEST is cached; otherwise, it is not cached. Based on this observation, we developed TIKTAG-v2.




Figure 7: Cache hit rate of `test_ptr` after executing TIKTAG-v2

Gadget. Figure 6 illustrates the TIKTAG-v2 gadget. The initial setting follows the MTE tag leakage template (§4.1), where the CPU mispredicts the branch `BR` and speculatively executes `CHECK` and `TEST`. At `CHECK`, the CPU triggers store-to-load forwarding behavior by storing `test_ptr` at `guess_ptr` and immediately loading the value from `guess_ptr` as `ptr`. If `CHECK` succeeds the tag checks (i.e., `Tg` matches `Tm`), the CPU forwards `test_ptr` to `ptr` and speculatively accesses `test_ptr` in `TEST`. If `CHECK` fails the tag checks (i.e., `Tg` mismatches `Tm`), the CPU blocks `test_ptr` from being forwarded to `ptr` and does not access `test_ptr`.

Experimental Results. TIKTAG-v2 showed its effectiveness as an MTE tag leakage gadget in ARM Cortex-A715 (core 4-7 of Pixel 8). We identified one requirement for TIKTAG-v2 to exhibit the tag leakage behavior: the store and load instructions in `CHECK` should be executed within 5 instructions. If this requirement is met, the cache hit rate of `test_ptr` after TIKTAG-v2 exhibited a notable difference between tag match and mismatch (Figure 7). Otherwise, the store-to-load forwarding always succeeded and the CPU forwarded `test_ptr` to `ptr`, and `test_ptr` was always cached.

To verify the requirement, we conducted experiments by inserting two instruction sequences, `SLIDE` and `GAP` (Figure 6). Both sequences consist of bitwise OR operations (`orr`), each dependent on the previous one, while not changing register or memory states. `SLIDE` is added before the gadget to control the alignment of store-to-load forwarding in the CPU pipeline. `GAP` is added between the store and load instructions in `CHECK` to control the distance between them. The results varying `SLIDE` from 0 to 40 and `GAP` from 0 to 4 are shown in Figure 7. In each subfigure, the x-axis represents the length of `GAP`, and the y-axis represents the cache hit rate of `test_ptr` after the gadget, measured over 1000 trials.

On tag match, the cache hit rate of `test_ptr` was near 100% in all conditions. However, on tag mismatch, the hit rate dropped on every 5 instructions in `SLIDE`. With `Len(GAP)=0` (i.e., no instruction), the hit rate drop was most frequent,

occurring in 4 out of `SLIDE` length. With `Len(GAP)=1, 2, 3`, the hit rate drop occurred in 3, 2, and 1 times every 5 instructions in `SLIDE`, respectively. With `Len(GAP)=4`, no hit rate drop was observed.

Similarly to TIKTAG-v1, the blockage of store-to-load forwarding was not specific to the MTE tag check fault, but was also observed with address translation fault, thus TIKTAG-v2 can also be utilized as an address-probing gadget.

Root Cause. The root cause of TIKTAG-v2 is likely due to the CPU preventing store-to-load forwarding on tag check faults. The CPU detects the store-to-load dependency utilizing internal buffers that log memory access information, such as Load-Store Queue (LSQ), and forwards the data if the dependency is detected. Although there is no documentation detailing the store-to-load forwarding mechanism on tag check faults, a relevant patent filed by ARM [2] provides a hint on the possible explanation.

The patent suggests that if the store-to-load dependency is detected, the load instruction can skip the tag check and the CPU can always forward the data. If so, store-to-load forwarding would not leak the tag check result (i.e., data is forwarded both on tag match and mismatch), as observed when `Len(GAP)` is 4 or more. When `Len(GAP)` is less than 4, however, the store-to-load succeeded on tag match and failed on tag mismatch.

We suspect that the CPU performs the tag check for the load instruction if the store-to-load dependency is not detected, and the CPU blocks the forwarding on tag check faults to prevent meltdown-like attacks [36, 42]. Considering the affected core (i.e., Cortex-A715) dispatches 5 instructions in a cycle [59], it is likely that the CPU cannot detect the dependency if the store and load instructions are executed in the same cycle, since the store information is not yet written to the internal buffers. If `Len(GAP)` is 4 or more, the store and load instructions are executed in the different cycles, and the CPU can detect the dependency. Therefore, the CPU skips the tag check and always forwards the data from the store to load instructions. If `Len(GAP)` is less than 4, the store and load instructions are executed in the same cycle, and the CPU fails to detect the dependency and performs the tag check for the load instruction. In this case, the forwarding is blocked on tag check faults.

Mitigation. To prevent tag leakage in TIKTAG-v2 at the micro-architectural level, the CPU should be designed to either always allow or always block the store-to-load forwarding regardless of the tag check result. Always blocking the store-to-load forwarding may raise a performance issue. Always allowing forwarding would not introduce significant performance overhead, but may introduce spectre/meltdown-like data leakage issues. In the future, hardware-based solutions can be developed to prevent micro-architectural data leakage issues [37, 75, 78].

At the software level, the following mitigations can be applied to prevent TIKTAG-v2.

i) *Speculation barrier:* If a speculation barrier is inserted before `CHECK`, the CPU does not speculatively execute store-to-load forwarding on both tag match and mismatch. Thus,

test_ptr is not cached regardless of the tag check result.

ii) *Prevention of gadget construction* : If the store and load instructions in CHECK are not executed within 5 instructions, the store-to-load forwarding is always allowed on both cases, making test_ptr cached always. The potential gadgets can be modified to have more than 5 instructions between the store and load instructions in CHECK, by adding dummy instructions (e.g., nop) or reordering the instructions.

As noted in §5.1, inserting additional barriers or instructions may introduce performance overhead, thus the performance impact should be evaluated.

6. Real-World Attacks

To demonstrate the exploitability of TIKTAG gadgets in MTE-based mitigation, this section develops two real-world attacks against Chrome and Linux kernel (Figure 9). There are several challenges to launching real-world attacks using TIKTAG gadgets. First, TIKTAG gadgets should be executed in the target address space, requiring the attacker to construct or find gadgets from the target system. Second, the attacker should control and observe the cache state to leak the tag check results. In the following, we demonstrate the real-world attacks using TIKTAG gadgets on two real-world systems: the Google Chrome browser (§6.1) and the Linux kernel (§6.2), and discuss the mitigation strategies.

6.1. Attacking Chrome Browser

A web browser is a primary attack surface for web-based attacks as it processes untrusted web content, such as JavaScript and HTML. We first overview the threat model (§6.1.1) and provide a TIKTAG gadget constructed in the V8 JavaScript engine (§6.1.2). Then, we demonstrate the effectiveness of TIKTAG gadgets in exploiting the browser (§6.1.3) and discuss the mitigation strategies (§6.1.4).

6.1.1. Threat Model. We follow the typical threat model of Chrome browser attacks, where the attacker aims to exploit memory corruption vulnerabilities in the renderer process. We assume the victim user visits the attacker-controlled website, which serves a malicious webpage. The webpage includes crafted HTML and JavaScript, which exploit memory corruption vulnerabilities in the victim’s renderer process. We assume all Chrome’s state-of-the-art mitigation techniques are in place, including ASLR [16], CFI [12], site isolation [56], and V8 sandbox [60]. Additionally, as an orthogonal defense, we assume that the renderer process enables random MTE tagging in PartitionAlloc [13].

6.1.2. Constructing TIKTAG Gadget. In the V8 JavaScript environment, TIKTAG-v2 was successfully constructed and leaked the MTE tags of any memory address. However, we didn’t find a constructible TIKTAG-v1 gadget, since the tight timing constraint between BR and CHECK was not feasible in our speculative V8 sandbox escape technique (§A).

```

1 slow = new Uint8Array(64);
2 victim = new Float64Array(64);
3 probe = new Uint8Array(512);
4 PROBE_OFFSET = 0;
5
6 function TikTag_v2(idx) {
7   // BR
8   if (!slow[0]) // cond_ptr
9     return 0;
10
11   // CHECK
12   victim[idx] = PROBE_OFFSET; // store to guess_ptr
13   val = victim[idx];          // load from guess_ptr
14   // TEST
15   // if Tg == Tm, probe[PROBE_OFFSET] is cached
16   // if Tg != Tm, probe[PROBE_OFFSET] is not cached
17   return probe[val];
18 }

```

(a) JavaScript TIKTAG-v2 gadget

```

1 int TikTag_v2(double idx) {
2   // BR
3   if (!slow[0])
4     return 0;
5
6   // CHECK
7   // victim bound check
8   if (idx < 0 || (uint64_t)idx > victim.length)
9     return undefined;
10  victim.buffer[idx] = PROBE_OFFSET; // store
11  double val = victim.buffer[idx]; // load
12
13  // TEST
14  // probe bound check
15  if (val < 0 || val >= probe.length)
16    return undefined;
17  return probe.buffer[val];
18 }

```

(b) Pseudo C code of JIT-optimized JavaScript TIKTAG-v2 gadget

Figure 8: V8 TIKTAG-v2 gadget in JavaScript

V8 TikTag-v2 Gadget. Figure 8 is the TIKTAG-v2 gadget constructed in the V8 JavaScript engine and its pseudo-C code after JIT compilation. With this gadget, the attacker can learn whether the guessed tag Tg matches with the tag Tm assigned to target_addr. The attacker prepares three arrays, slow, victim, probe, and an idx value. slow is a Uint8Array with a length of 64 and is accessed in BR to trigger the branch misprediction. victim is a Float64Array with length 64, which is accessed to trigger store-to-load forwarding. probe is a Uint8Array with length 512, and is accessed in TEST to leak the tag check result. A Number type idx value is used in out-of-bounds access of victim. idx value is chosen such that victim[idx] points to target_addr with a guessed tag Tg (i.e., (Tg«56)|target_addr). To speculatively access the target_addr outside the V8 sandbox, we leveraged the speculative V8 sandbox escape technique we discovered during our research, which we detail in §A.

Line 8 of Figure 8a is the BR block of the TIKTAG-v2 gadget, triggering branch misprediction with slow[0]. Line 12-13 is the CHECK block, which performs the store-to-load forwarding with victim[idx], accessing target_addr with a guessed tag Tg. When this code is JIT-compiled (Figure 8b), a bound check is performed, comparing idx against victim.length. If idx is an out-of-bounds index, the code

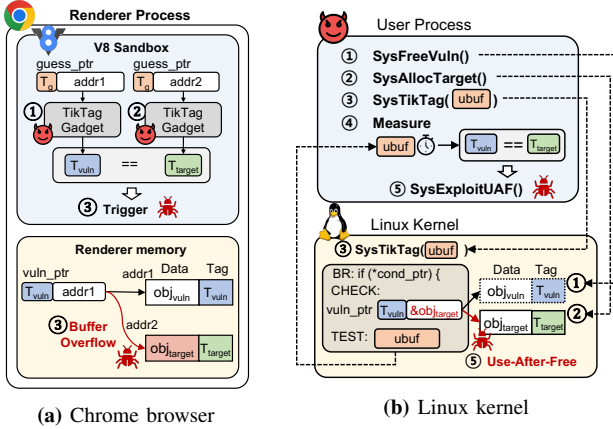


Figure 9: MTE bypass attacks

returns undefined, but if `victim.length` field takes a long time to be loaded, the CPU speculatively executes the following store and load instructions. After that, line 17 implements the TEST block, which accesses the probe with the forwarded value `val` as an index. Again, a bound check on `val` against the length of probe is preceded, but this check succeeds as `PROBE_OFFSET` is smaller than the length of probe array. As a result, `probe[PROBE_OFFSET]` is cached only when the store-to-load forwarding succeeds, which is the case when `Tg` matches `Tm`.

6.1.3. Chrome MTE bypass attack. Figure 9a illustrates the overall MTE bypass attack on the Chrome browser with arbitrary tag leakage primitive of TIKTAG gadgets. We assume a buffer overflow vulnerability in the renderer process, where exploiting a temporal vulnerability (e.g., use-after-free) is largely the same. The vulnerability overflows a pointer (i.e., `vuln_ptr`) to a vulnerable object (i.e., `objvuln`), corrupting the adjacent object (i.e., `objtarget`). With PartitionAlloc’s MTE enforcement, two objects have different tags with a 14/15 probability. To avoid raising an exception, the attacker needs to ensure that the tags of `objvuln` and `objtarget` are the same. TIKTAG-v2 can be utilized to leak the tag of `objvuln` (1) and `objtarget` (2). If both leaked tags are the same, the attacker exploits the vulnerability, which would not raise a tag check fault (3). Otherwise, the attacker frees and re-allocates `objtarget` and goes back to the first step until the tags match.

Triggering Cache Side-Channel. To successfully exploit a TIKTAG gadget, the attacker needs to satisfy the following requirements: i) branch training, ii) cache control, and iii) cache measurement. All three requirements can be met in JavaScript. First, the attacker can train the branch predictor by running the gadget with non-zero `slow[0]` and in-bounds `idx`, and trigger the branch misprediction in BR with zero value in `slow[0]` and out-of-bounds `idx`. Second, the attacker can evict the cache lines of `slow[0]`, `victim.length`, and `probe[PROBE_OFFSET]` with JavaScript cache eviction techniques [3, 19, 74]. Third, the attacker can measure the cache status of `probe[PROBE_OFFSET]` with a high-resolution timer based on `SharedArrayBuffer` [14, 62].

Exploiting Memory Corruption Vulnerabilities. Given the leaked MTE tags, the attacker can exploit spatial and temporal memory corruption vulnerabilities in the renderer. The attack strategy is largely the same as the traditional memory corruption attacks but should ensure that the vulnerability does not raise a tag check fault utilizing the leaked tags. We further detail the attack strategy in §C.

6.1.4. Mitigation. To mitigate the TIKTAG gadget-based MTE bypass attacks in the browser renderer process, the following mitigations can be employed:

i) *Speculative execution-aware sandbox:* To stop attackers from launching TIKTAG-based attacks from a sandboxed environment like V8 sandbox, the sandbox can be fortified by preventing any speculative memory access beyond the sandbox’s memory region. While modern web browsers employ a sandbox to isolate untrusted web contents from the renderer, they often overlook speculative paths. For instance, Chrome V8 sandbox [60] and Safari Webkit sandbox [54] do not completely mediate the speculative paths [25]. Based on current pointer compression techniques [68], speculative paths can be restricted to the sandbox region by masking out the high bits of the pointers.

ii) *Speculation barrier:* As suggested in §5, placing a speculation barrier after BR for potential TIKTAG gadgets can prevent speculative tag leakage attacks. However, this mitigation may not be applicable in the performance-critical browser environment, as it may introduce significant performance overhead.

iii) *Prevention of gadget construction:* As suggested in §5.2, the TIKTAG-v2 gadget can be mitigated by padding instructions between store and load instructions. A TIKTAG-v1 gadget, although we have not found an exploitable one, can be mitigated by padding instructions between a branch and memory accesses, as described in §5.1.

6.2. Attacking the Linux Kernel

The Linux kernel on ARM is widely used for mobile devices, servers, and IoT devices, making it an attractive attack target. Exploiting a memory corruption vulnerability in the kernel can escalate the user’s privilege, and thus MTE is a promising protection mechanism for the Linux kernel. TIKTAG-based attacks against the Linux kernel pose unique challenges different from the browser attack (§6.1). This is because the attacker’s address space is isolated from the kernel’s address space where the gadget will be executed. In the following, we first overview the threat model of the Linux kernel (§6.2.1) and provide a proof-of-concept TIKTAG gadget we discovered in the Linux kernel (§6.2.2). Finally, we demonstrate the effectiveness of TIKTAG gadgets in exploiting Linux kernel vulnerabilities (§6.2.3).

6.2.1. Threat Model. The threat model here is largely the same as that of typical privilege escalation attacks against the kernel. Specifically, we focus on the ARM-based Android Linux kernel, hardened with default kernel protections (e.g., KASLR, SMEP, SMAP, and CFI). We further assume the

```

1 static ssize_t snd_timer_user_read(struct file *file,
2   char __user *buffer, size_t count, loff_t *offset) {
3   ...
4   switch (tu->tread) {
5   default:
6     return -ENOTSUPP;
7     break;
8   }
9   // BR: speculation branch with cond_ptr (tu)
10  switch (tu->tread) {
11    case TREAD_FORMAT_TIME32: // branch prediction destination
12      // CHECK: dereference guess_ptr (tread) with 4 loads
13      tread32 = (struct snd_timer_tread32) {
14        .event = tread->event,
15        .tstamp_sec = tread->tstamp_sec,
16        .tstamp_nsec = tread->tstamp_nsec,
17        .val = tread->val,
18      };
19      // TEST: dereference test_ptr (buffer)
20      if (copy_to_user(buffer, &tread32, sizeof(tread32)))
21        err = -EFAULT;
22      break;
23    default: // correct branch destination
24      err = -ENOTSUPP;
25      break;
26  }
27 }

```

Figure 10: TIKTAG-v1 gadget in `snd_timer_user_read()`. `-/+` denotes the code changes to make the gadget exploitable.

kernel is hardened with an MTE random tagging solution, similar to the production-ready MTE solutions, Scudo [4]. To be specific, each memory object is randomly tagged, and a random tag is assigned when an object is freed, thereby preventing both spatial and temporal memory corruptions.

The attacker is capable of running an unprivileged process and aims to escalate their privilege by exploiting memory corruption vulnerabilities in the kernel. It is assumed that the attacker knows kernel memory corruption vulnerabilities but does not know any MTE tag of the kernel memory. Triggering memory corruption between kernel objects with mismatching tags would raise a tag check fault, which is undesirable for real-world exploits.

One critical challenge in this attack is that the gadget should be constructed by reusing the existing kernel code and executed by the system calls that the attacker can invoke. As the ARMv8 architecture separates user and kernel page tables, user space gadgets cannot speculatively access the kernel memory. This setup is very different from the threat model of attacking the browser (§6.1), which leveraged the attacker-provided code to construct the gadget. We excluded the eBPF-based gadget construction either [15, 27], because eBPF is not available for the unprivileged Android process [33].

6.2.2. Kernel TikTag Gadget. As described in §4.1, TIKTAG gadgets should meet several requirements, and each requirement entails challenges in the kernel environment.

First, in BR, a branch misprediction should be triggered with `cond_ptr`, which should be controllable from the user space. Since recent AArch64 processors isolate branch prediction training between the user and kernel [33], the branch training needs to be performed from the kernel space. Second, in CHECK, `guess_ptr` should be dereferenced. `guess_ptr` should be crafted from the user space such that it embeds a guess tag (Tg) and points to the kernel address

(i.e., `target_addr`) to leak the tag (Tm). Unlike the browser JavaScript environment (§6.1), user-provided data is heavily sanitized in system calls, so it is difficult to create an arbitrary kernel pointer. For instance, `access_ok()` ensures that the user-provided pointer points to the user space, and the `array_index_nospec` macro prevents speculative out-of-bounds access with the user-provided index. Thus, `guess_ptr` should be an existing kernel pointer, specifically the vulnerable pointer that causes memory corruption. For instance, a dangling pointer in use-after-free (UAF) or an out-of-bounds pointer in buffer overflow can be used. Lastly, in TEST, `test_ptr` should be dereferenced, and `test_ptr` should be accessible from the user space. To ease the cache state measurement, `test_ptr` should be a user space pointer provided through a system call argument.

Discovered Gadgets. We manually analyzed the source code of the Linux kernel to find the TIKTAG gadget meeting the aforementioned requirements. As a result, we found one potentially exploitable TIKTAG-v1 gadget in `snd_timer_user_read()` (Figure 10). This gadget fulfills the requirements of TIKTAG-v1 (§5.1). At line 10 (i.e., BR), the switch statement triggers branch misprediction with a user-controllable value `tu->tread` (i.e., `cond_ptr`). At lines 14-17 (i.e., CHECK), `tread` (i.e., `guess_ptr`) is dereferenced by four load instructions. `tread` points to a struct `snd_timer_tread64` object that the attacker can arbitrarily allocate and free. If a temporal vulnerability transforms `tread` into a dangling pointer, it can be used as a `guess_ptr`. At line 20, (i.e., TEST), a user space pointer `buffer` (i.e., `test_ptr`) is dereferenced in `copy_to_user`.

As this gadget is not directly reachable from the user space, we made a slight modification to the kernel code; we removed the early return for the default case at line 6. This ensures that the buffer is only accessed in the speculative path to observe the cache state difference due to speculative execution. Although this modification is not realistic in a real-world scenario, it demonstrates the potential exploitability of the gadget if similar code changes are made.

We discovered several more potentially exploitable gadgets, but we were not able to observe the cache state difference between the tag match and mismatch. Still, we think there is strong potential for exploiting those gadgets. Launching TIKTAG-based attacks involves complex and sensitive engineering, and thus we were not able to experiment with all possible cases. Especially, TIKTAG-v1 relies on the speculation shrinkage on wrong path events, which may also include address translation faults or other exceptions in the branch misprediction path. As system calls involve complex control flows, the speculation shrinkage may not be triggered as expected. In addition, several gadgets may become exploitable when kernel code changes. For instance, a TIKTAG-v1 gadget in `ip6mr_ioctl()` did not exhibit an MTE tag leakage behavior when called from its system call path (i.e., `ioctl`). However, the gadget had tag leakage when it was ported to other syscalls (e.g., `write`) with a simple control flow.

6.2.3. Kernel MTE bypass attack. Figure 9b illustrates the MTE bypass attacks on the Linux kernel. Taking a use-after-free vulnerability as an example, we assume the attacker has identified a corresponding TIKTAG gadget, `SysTikTagUAF()`, capable of leaking the tag check result of the dangling pointer created by the vulnerability. For instance, the TIKTAG-v1 gadget in `snd_timer_user_read()` (Figure 10) can leak the tag check result of `tread`, which can become a dangling pointer by a use-after-free or double-free vulnerability.

The attack proceeds as follows: First, the attacker frees a kernel object (i.e., `obj_vuln`) and leaves its pointer (i.e., `vuln_ptr`) as a dangling pointer (①). Next, the attacker allocates another kernel object (i.e., `obj_target`) at the address of `obj_vuln` with `SysAllocTarget()` (②). Then, the attacker invokes `SysTikTag()` with a user space buffer (i.e., `ubuf`) (③), and leaks the tag check result (i.e., `Tm == Tg`) by measuring the access latency of `ubuf` (④). If the tags match, the attacker triggers `SysExploitUAF()`, a system call that exploits the use-after-free vulnerability (⑤). Otherwise, the attacker re-allocates `obj_target` until the tags match.

Triggering Cache Side-Channel. As in §6.1.3, a successful TIKTAG gadget exploitation requires i) branch training, ii) cache control, and iii) cache measurement. For branch training, the attacker can train the branch predictor and trigger speculation with user-controlled branch conditions from the user space. For cache control, the attacker can flush the user space buffer (i.e., `ubuf`), while the kernel memory address can be evicted by cache line bouncing [23]. For cache measurement, the access latency of `ubuf` can be measured with the virtual counter (i.e., `CNTVCT_EL0`) or a memory counter-based timer (i.e., near CPU cycle resolution).

Exploiting Memory Corruption Vulnerabilities. TIKTAG gadgets enable bypassing MTE and exploiting kernel memory corruption vulnerabilities. The attacker can invoke the TIKTAG gadget in the kernel to speculatively trigger the memory corruption and obtain the tag check result. Then, the attacker can obtain the tag check result, and trigger the memory corruption only if the tags match. We detail the Linux kernel MTE bypass attack process in §D.

6.2.4. Mitigation. To mitigate TIKTAG gadget in the Linux kernel, the kernel developers should consider the following mitigations:

i) *Speculation barrier:* Speculation barriers can effectively mitigate TIKTAG-v1 gadget in the Linux kernel. To prevent attackers from leaking the tag check result through the user space buffer, kernel functions that access user space addresses, such as `copy_to_user` and `copy_from_user`, can be hardened with speculation barriers. As described in §5.1, leaking tag check results with store access can be mitigated by placing a speculation barrier before the store access (i.e., `TEST`). For instance, to mitigate the gadgets leveraging `copy_to_user`, a speculation barrier can be inserted before the `copy_to_user` invocation. For gadgets utilizing load access to the user space buffer, the barriers mitigate the gadgets if inserted between the branch and the kernel memory access (i.e., `CHECK`). For instance, to mitigate the gadgets leveraging `copy_from_user`, the kernel developers

TABLE 1: MTE schemes in Android and Chrome allocators.

Allocator	Scudo	PartitionAlloc
Usage	Android	Chrome
Allocation (new)	Random (0x0-0xf)*	Random (0x1-0xf)
Allocation (reuse)	None	None
Release	Random(0x0-0xf)*	Increment

* OddEvenTags support.

should carefully analyze the kernel code base to find the pattern of the conditional branch, kernel memory access, and `copy_from_user()`, and insert a speculation barrier between the branch and the kernel memory access.

ii) *Prevention of gadget construction:* To eliminate potential TIKTAG gadgets in the Linux kernel, the kernel source code can be analyzed and patched. As TIKTAG gadgets can also be constructed by compiler optimizations, a binary analysis can be conducted. For each discovered gadget, instructions can be reordered or additional instructions can be inserted to prevent the gadget construction, following the mitigation strategies in §5.1 and §5.2.

7. Evaluation

In this section, we evaluate the TIKTAG gadgets and MTE bypass exploits in two MTE-enabled systems, the Chrome browser (§7.1) and the Linux kernel (§7.2). All experiments were conducted on the Google Pixel 8 devices.

7.1. Chrome Browser Tag Leakage

We evaluated the TIKTAG-v2 gadget in the V8 JavaScript engine in two environments: i) the standalone V8 JavaScript engine, and ii) the Chromium application. The V8 JavaScript engine runs as an independent process, reducing the interference from the Android platform. The Chromium application runs as an Android application, subject to the Android’s application management such as process scheduling and thermal throttling. The experiments were conducted with the V8 v12.1.10 and Chromium v119.0.6022.0 release build.

We leveraged MTE random tagging schemes provided by the underlying allocators (Table 1). The standalone V8 used the Scudo allocator [4] (i.e., Android default allocator), which supports 16 random tags for random tagging and offers the `OddEvenTags` option. When `OddEvenTags` is enabled, Scudo alternates odd and even random tags for neighboring objects, preventing linear overflow (i.e., `OVERFLOW_TUNING`). When `OddEvenTags` is disabled, Scudo utilizes all 16 random tags for every object to maximize tag entropy for use-after-free detection (i.e., `UAF_TUNING`). By default, `OddEvenTags` is enabled, while we evaluate both settings. Upon releasing an object, Scudo sets a new random tag that does not collide with the previous one.

`PartitionAlloc` (i.e., Chrome default allocator) utilizes 15 random tags and reserves the tag 0x0 for unallocated memory. When releasing an object, `PartitionAlloc` increments the tag by one, making the tag of the re-allocated memory address predictable. However, in real-world exploits, it is challenging

TABLE 2: Results of MTE bypass exploits against the V8 JavaScript engine

Vuln.	OddEvenTags	Accuracy	Time (s)
Tag Leakage	0 (UAF_TUNING)	100/100 (100%)	3.04
Linear Overflow	0 (UAF_TUNING)	98/100 (98%)	13.52
Linear Overflow	1 (OVERFLOW_TUNING)	0/100 (0%)	N/A
Use-After-Free	0 (UAF_TUNING)	97/100 (97%)	10.66
Use-After-Free	1 (OVERFLOW_TUNING)	98/100 (98%)	6.09

TABLE 3: Results of MTE bypass exploits against the Chromium application


Vuln.	Accuracy	Time (s)
Tag Leakage	95/100 (95%)	2.54
Linear Overflow	97/100 (97%)	16.11
Use-After-Free	95/100 (95%)	21.90

to precisely control the number of releases for a specific address, thus we assume the attacker still needs to leak the tag after each allocation.

For the evaluation, we constructed the TIKTAG-v2 gadget in JavaScript (Figure 6) and developed MTE bypass attacks as described in §6.1.3. These attacks exploit artificial vulnerabilities designed to mimic real-world renderer vulnerabilities, specifically linear overflow [47] and use-after-free [45]. We developed custom JavaScript APIs to allocate, free, locate, and access the renderer object to manipulate the memory layout and trigger the vulnerabilities. It’s worth noting that our evaluation shows the best-case performance of MTE bypass attacks since real-world renderer exploits involve additional overheads in triggering the vulnerabilities and controlling the memory layout.

V8 JavaScript Engine. In the standalone V8 JavaScript engine, we evaluated the tag leakage of the TIKTAG-v2 gadget with *cache eviction* and a *memory-based timer*. For cache eviction, we used an L1 index-based random eviction set, 500 elements for `slow[0]` and `probe[PROBE_OFFSET]`, 300 elements for `victim.length`. The eviction performance of the random eviction set varies on each run, so we repeated the same test 5 times and listed the best result. The random eviction can be optimized with eviction set algorithms [74]. We used a memory counter-based timer with a custom worker thread incrementing a counter, which is equivalent to the `SharedArrayBuffer` timer [62]. For all possible tag guesses (i.e., `0x0-0xf`), we measured the access latency of `probe[PROBE_OFFSET]` after the gadget 256 times and determined the guessed tag with the minimum average access latency as the correct tag.

Table 2 summarizes the MTE bypass exploit results in V8. For a single tag leakage, the gadget was successful in all 100 runs (100%), with an average elapsed time of 3.04 seconds. MTE bypass exploits were evaluated over 100 runs for each vulnerability and `OddEvenTags` configuration (i.e., disabled (0) and enabled (1)). We excluded linear overflow exploit with `OddEvenTags` enabled, since the memory corruption is always detected with spatially adjacent objects tagged with different tags and the attack would always fail. The results demonstrate that the attacks were successful in over 97% of



(a) TIKTAG-v1 (user space)

(b) TIKTAG-v2 (kernel space)

Figure 11: Kernel context evaluation of TIKTAG gadgets

the runs, with an average elapsed time of 6 to 13 seconds. In use-after-free exploits, enabling `OddEvenTags` decreased the average elapsed time by around 40%, due to the decrease in tag entropy from 16 to 8, doubling the chance of tag collision between the temporally adjacent objects.

Chromium Application. In the Chromium application setting, we evaluated the TIKTAG-v2 gadget with *cache flushing* and a *SharedArrayBuffer-based timer*. Unlike V8, random eviction did not effectively evict cache lines, so we manually flushed the cache lines with `dc civac` instruction. We attribute this to the aggressive resource management of Android, which can be addressed in the future with cache eviction algorithms tailored for mobile applications. To measure the cache eviction set overhead, we included the cache eviction set traversals in all experiments, using the same cache eviction configuration of the V8 experiments. We measured access latency with a `SharedArrayBuffer`-based timer as suggested by web browser speculative execution studies [3, 19]. The MTE bypass exploits experiments were conducted in the same manner as the V8 experiments.

Table 3 shows the MTE bypass exploit results in the Chromium application. The tag leakage of the TIKTAG-v2 gadget in the Chromium application was successful in 95% of 100 runs, with an average elapsed time of 2.54 seconds. With the MTE bypass exploits, success rates were over 95% for both vulnerability types, with an average elapsed time of 16.11 and 21.90 seconds for linear overflow and use-after-free, respectively.

7.2. Linux Kernel Tag Leakage

The experiments were conducted on the Android 14 kernel v5.15 using the default configuration. We used 15 random tags (i.e., `0x0-0xe`) for kernel objects, as tag `0xf` is commonly reserved for the access-all tag in the Linux kernel [38]. The cache line eviction of kernel address `cond_ptr` to trigger the speculative execution was achieved by cache line bouncing [23] from the user space. For cache measurement, we utilized the virtual counter (i.e., `CNTVCT_EL0`) to determine the cache hit or miss with the

TABLE 4: Results of MTE bypass exploits against the Linux kernel

Vuln.	Accuracy	Time (s)
Tag Leakage (artificial)	100/100 (100%)	0.12
Tag Leakage (snd_timer_user_read())	100/100 (100%)	3.38
Buffer Overflow (artificial)	100/100 (100%)	0.18
Use-After-Free (snd_timer_user_read())	97/100 (97%)	6.86

threshold 1.0, which is accessible from the user space. As the virtual counter has a lower resolution (24.5MHz) than the CPU cycle frequency (2.4-2.9 GHz), the accuracy of the cache hit rate is lower than the physical CPU counter-based measurements in §5. The access time was measured in the user space or kernel space, depending on the experiment.

Kernel Context Evaluation. We first evaluated whether TIKTAG gadgets can leak MTE tags in the Linux kernel context (Figure 11). We created custom system calls containing TIKTAG-v1 (Figure 2) and TIKTAG-v2 (Figure 6) gadgets and executed them by calling the system calls from the user space. In CHECK, we accessed the `guess_ptr` that holds either the correct or wrong tag `Tg`. In TEST, `test_ptr` pointed to either a kernel address or a user space address, depending on whether the cache state difference was measured in the kernel or user space. When we leveraged a user space address as `test_ptr`, we passed a user buffer pointer to the kernel space as a system call argument and accessed the pointer in TEST using `copy_to_user()`. The user space address was flushed in the user space before the system call invocation, and the cache state was measured after the system call returned. When we used a kernel address as `test_ptr`, the cache flush and measurement were performed in the kernel. Each experiment measured the access time over 1000 runs.

When executing TIKTAG-v1 in the kernel context, the MTE tag leakage was feasible in both the kernel and user space, where the user space measurement results are shown in Figure 11a. Compared to the user space gadget evaluation (Figure 3), the kernel context required more loads in CHECK to distinguish the cache state difference. Specifically, the cache state difference was discernible from 4 loads in the kernel context, while the user space context required only 2 loads. This can be attributed to the noises from the kernel to the user space context switch overhead, such that the cache hit rates of the tag match cases were lower (i.e., under 90%) than the user space gadget evaluation (i.e., 100%).

When executing the TIKTAG-v2 gadget in the kernel space, MTE tag leakage was observed only in the kernel space (Figure 11b). When we measured the access latency of `test_ptr` in the user space, the gadget did not exhibit a cache state difference. Although the TIKTAG-v2 gadget might not be directly exploitable in the user space, cache state amplification techniques [19, 77] could be utilized to make it observable from the user space.

Kernel MTE Bypass Exploit. We evaluated MTE bypass exploits in the Linux kernel with two TIKTAG-v1 gadgets: an artificial TIKTAG-v1 gadget with 8 loads in CHECK (i.e., artificial) and a real-world TIKTAG-v1 gadget in `snd_timer_user_read()` (Figure 10). The artificial gadget

evaluates the best-case performance of MTE bypass attacks, while the `snd_timer_user_read()` gadget demonstrates real-world exploit performance. Both gadgets were triggered by invoking the system call containing the gadget from the user space, leveraging a user space address as `test_ptr`, and measuring the access latency of `test_ptr` in user space.

We conducted a tag leakage attack and MTE bypass attack for each gadget. For the MTE bypass attack, we synthesized a buffer overflow vulnerability. Each gadget dereferenced the vulnerable pointer (i.e., `guess_ptr`) to trigger tag checks; an out-of-bounds pointer and a dangling pointer for the buffer overflow and use-after-free exploits, respectively. The exploit methodology followed the process described in §D.

Table 4 summarizes the MTE bypass exploit results. For a single tag leakage, the gadgets successfully leaked the correct tag in all 100 runs (100%), with an average elapsed time of 0.12 seconds in the artificial gadget, and 3.38 seconds in the `snd_timer_user_read()` gadget. The MTE bypass exploit for the artificial TIKTAG-v1 gadget was successful in all 100 runs (100%), with an average elapsed time of 0.18 seconds. Regarding the MTE bypass exploit for the `snd_timer_user_read()` gadget, the success rate was 97% with an average elapsed time of 6.86 seconds. As the `snd_timer_user_read()` gadget involves complex kernel function calls and memory accesses, the performance of the MTE bypass exploit is slightly lower compared to the artificial gadget. Nevertheless, it still demonstrates a high success rate within a reasonable time frame.

8. Related work

MTE Security Analysis. Partap et al. [53] analyzed the software-level MTE support in real-world memory allocators. Google Project Zero [40] explored speculative execution attacks against MTE hardware for the first time. StickyTags [20] identified an MTE tag leakage gadget (which is similar to TIKTAG-v1) and proposed a deterministic tagging-based defense that does not utilize random tags due to the potential tag leakage. Compared to StickyTags, our work identified a new type of MTE tag leakage gadget, TIKTAG-v2, and analyzed the root cause of both TIKTAG-v1 and TIKTAG-v2 gadgets. We also demonstrated the real-world exploitation of TIKTAG gadgets in Google Chrome and the Linux kernel and proposed new defense mechanisms to mitigate the security risks posed by TIKTAG gadgets. While StickyTags proposed deterministic tagging due to the potential tag leakage, our work focuses on hardening the random tagging-based MTE defense, which are developed by major vendors including Google [39], the Linux kernel [24], and secure operating systems [21, 52, 67].

Speculative Attacks on Protection Mechanisms. Speculative probing [18] suggested that speculative execution can be used to probe address mappings and bypass address space layout randomization (ASLR). PACMAN [58] identified speculative gadgets that leak Pointer Authentication Code (PAC). PACMAN and TIKTAG share similarities in bypassing hardware security features by exploiting speculative

execution, but there are some key differences. PACMAN leverages software-observable behavior, where an error bit set in the pointer cannot be translated and accessed, resulting in a cache miss. In contrast, TIKTAG exploits black-box hardware logic, where the tag check fault blocks certain CPU operations, indirectly causing a cache miss. This distinction makes it more challenging to identify TIKTAG templates that cause MTE tag leakage (i.e., TIKTAG-v1 and TIKTAG-v2) and to further analyze their root causes, compared to PACMAN's PAC leakage gadgets. From the defense perspective, ARMv8.6 FEAT_FPAC mitigates PACMAN attacks by authentication and memory access, allowing all memory accesses regardless of the authentication result [35]. MTE tag leakage can be mitigated by preventing tag check faults from affecting memory accesses.

Transient Execution Attacks. Transient execution attacks exploit micro-architectural behaviors to leak secret information. Researchers have analyzed various micro-architectural implementations including speculative execution [29, 36, 70, 76], memory disambiguation prediction [22, 42, 43], and CPU internal buffers [71, 72]. Recent attacks exploited data prefetching behaviors to leak secret information or construct covert channels [11, 61, 63, 73]. Compared to these attacks, we identified for the first time that data prefetching behaviors can also be exploited to leak hardware exceptions, such as tag check faults (§5.1).

9. Conclusion

This paper explores the potential security risks posed by speculative execution attacks against ARM Memory Tagging Extension (MTE). We identify new MTE oracles, TIKTAG-v1 and TIKTAG v2, capable of leaking MTE tags from arbitrary memory addresses. TIKTAG gadgets can bypass MTE-based defense in real-world systems, including Google Chrome and the Linux kernel. Our findings provide significant insights into the design and deployment of both memory tagging-based hardware and software defenses.

Acknowledgment

This work was supported by SAMSUNG Research, Samsung Electronics Co., Ltd. This work was supported by Institute for Information & communications Technology Promotion (IITP) grant funded by the Korea government (MSIP) (No.2020-0-01840, Analysis on technique of accessing and acquiring user data in smartphone). This work was supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2023-00209093). The Institute of Engineering Research (IOER) and Automation and Systems Research Institute (ASRI) at Seoul National University provided research facilities for this work.

References

[1] M. Abadi, M. Budiu, U. Erlingsson, and J. Ligatti. Control-flow integrity principles, implementations, and applications. In *Proceedings*

of the ACM Transactions on Information and System Security, Nov. 2009.

[2] R. Abhishek, K. M. Bruce, and A. P. TONNERRE. Skipping tag check for tag-checked load operation, 2020. <https://patents.google.com/patent/US11221951/>.

[3] A. Agarwal, S. O'Connell, J. Kim, S. Yehezkel, D. Genkin, E. Ronen, and Y. Yarom. Spook.js: Attacking chrome strict site isolation via speculative execution. In *Proceedings of the 42nd IEEE Symposium on Security and Privacy (Oakland)*, San Francisco, CA, May 2022.

[4] android. external/scudo/standalone/combined.h, 2024. https://android.googlesource.com/platform/external/scudo/+refs/tags/android-platform-15.0.0_r1/standalone/combined.h.

[5] D. N. Armstrong, H. Kim, O. Mutlu, and Y. N. Patt. Wrong path events: Exploiting unusual and illegal program behavior for early misprediction detection and recovery. In *Proceedings of the 37th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Portland, OR, Dec. 2004.

[6] M. Bakhshalipour, P. Lotfi-Kamran, and H. Sarbazi-Azad. Domino temporal data prefetcher. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 131–142. IEEE, 2018.

[7] G. S. Blog. Mte - the promising path forward for memory safety, 2023. <https://security.googleblog.com/2023/11/mte-promising-path-forward-for-memory.html>.

[8] L. CAI, K. Nathella, J. Lee, and S. Dam. Prefetch mechanism for a cache structure, 2020. <https://patents.google.com/patent/US11526356B2/>.

[9] C. Carruth. Speculative load hardening, 2018. <https://lvm.org/docs/SpeculativeLoadHardening.html>.

[10] M. Castro, M. Costa, and T. Harris. Securing software by enforcing data-flow integrity. In *Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Seattle, WA, Nov. 2006.

[11] Y. Chen, L. Pei, and T. E. Carlson. Afterimage: Leaking control flow data and tracking load operations via the hardware prefetcher. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Vancouver, Canada, Apr. 2023.

[12] Chromium. [pac] enable armv8.3 pac (pointer authentication code), 2022.

[13] chromium. base/allocator/partition_allocator/partition_bucket.cc, 2024. https://chromium.googlesource.com/chromium/src/+refs/tags/131.0.6772.1/base/allocator/partition_allocator/src/partition_alloc/partition_bucket.cc.

[14] M. W. Docs. Sharedarraybuffer, 2024. https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Global_Objects/SharedArrayBuffer.

[15] eBPF. ebpf documentation, 2024. <https://ebpf.io/what-is-ebpf/>.

[16] J. Edge. Kernel address space layout randomization, 2013. <https://lwn.net/Articles/569635/>.

[17] J. W. Fu, J. H. Patel, and B. L. Janssens. Stride directed prefetching in scalar processors. In *Proceedings of the 25th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Portland, OR, Dec. 1992.

[18] E. Göktas, K. Razavi, G. Portokalidis, H. Bos, and C. Giuffrida. Speculative probing: Hacking blind in the spectre era. In *Proceedings of the 27th ACM Conference on Computer and Communications Security (CCS)*, Virtual, Nov. 2020.

[19] Google. Spectre, 2021. <https://leaky.page>.

[20] F. Gorter, T. Kroes, H. Bos, and C. Giuffrida. Sticky tags: Efficient and deterministic spatial memory error mitigation using persistent memory tags. In *Proceedings of the 43rd IEEE Symposium on Security and Privacy (Oakland)*, San Francisco, CA, May 2024.

[21] GrapheneOS. hardened_malloc, 2024. https://github.com/GrapheneOS/hardened_malloc.

[22] S. Islam, A. Moghimi, I. Bruhns, M. Krebbel, B. Gulmezoglu, T. Eisenbarth, and B. Sunar. {SPOILER}: Speculative load hazards boost rowhammer and cache attacks. In *Proceedings of the 28th USENIX Security Symposium (Security)*, Santa Clara, CA, Aug. 2019.

[23] P. Z. Jann Horn. Reading privileged memory with a side-channel, 2018. <https://googleprojectzero.blogspot.com/2018/01/>

reading-privileged-memory-with-side.html.

- [24] T. L. Kernel. Hardware tag-based kasan, 2023. <https://docs.kernel.org/dev-tools/kasan.html#hardware-tag-based-kasan>.
- [25] J. Kim, S. van Schaik, D. Genkin, and Y. Yarom. ileakage: Browser-based timerless speculative execution attacks on apple devices. In *Proceedings of the 30th ACM Conference on Computer and Communications Security (CCS)*, Copenhagen, Denmark, Nov. 2023.
- [26] J. Kim, J. Park, Y. Lee, C. Song, T. Kim, and B. Lee. Petal: Ensuring access control integrity against data-only attacks on linux. In *Proceedings of the 31st ACM Conference on Computer and Communications Security (CCS)*, Salt Lake City, UT, Oct. 2024.
- [27] O. Kirzner and A. Morrison. An analysis of speculative type confusion vulnerabilities in the wild. In *Proceedings of the 30th USENIX Security Symposium (Security)*, Virtual, Aug. 2021.
- [28] S. Knox. Real-time kernel protection (rkp), 2016.
- [29] P. Kocher, J. Horn, A. Fogh, D. Genkin, D. Gruss, W. Haas, M. Hamburg, M. Lipp, S. Mangard, T. Prescher, M. Schwarz, and Y. Yarom. Spectre attacks: Exploiting speculative execution. In *Proceedings of the 40th IEEE Symposium on Security and Privacy (Oakland)*, San Jose, CA, May 2019.
- [30] G. Li, H. Zhang, J. Zhou, W. Shen, Y. Sui, and Z. Qian. A hybrid alias analysis and its application to global variable protection in the linux kernel. In *Proceedings of the 32nd USENIX Security Symposium (Security)*, Anaheim, CA, Aug. 2023.
- [31] H. Liljestrand, C. Chinea, R. Denis-Courmont, J.-E. Ekberg, and N. Asokan. Color my world: Deterministic tagging for memory safety. *arXiv preprint arXiv:2204.03781*, 2022.
- [32] A. Limited. Memory tagging extension, 2019. https://developer.arm.com/-/media/Arm%20Developer%20Community/PDF/Arm_Memory_Tagging_Extension_Whitepaper.pdf.
- [33] A. Limited. Cache speculation side-channels v2.5, 2020. <https://developer.arm.com/documentation/102816/0205/>.
- [34] A. Limited. Speculative oracles on memory tagging, 2024. <https://developer.arm.com/documentation/109544/latest>.
- [35] A. Limited. Pacman security vulnerability, 2024. <https://developer.arm.com/documentation/ka005109/latest/>.
- [36] M. Lipp, M. Schwarz, D. Gruss, T. Prescher, W. Haas, A. Fogh, J. Horn, S. Mangard, P. Kocher, D. Genkin, Y. Yarom, and M. Hamburg. Meltdown: Reading kernel memory from user space. In *Proceedings of the 27th USENIX Security Symposium (Security)*, Baltimore, MD, Aug. 2018.
- [37] K. Loughlin, I. Neal, J. Ma, E. Tsai, O. Weisse, S. Narayanasamy, and B. Kasicki. {DOLMA}: Securing speculation with the principle of transient {Non-Observability}. In *Proceedings of the 30th USENIX Security Symposium (Security)*, Virtual, Aug. 2021.
- [38] G. P. Z. Mark Brand. Mte as implemented, part 3: The kernel, 2023. <https://googleprojectzero.blogspot.com/2023/08/mte-as-implemented-part-3-kernel.html>.
- [39] G. P. Z. Mark Brand. First handset with mte on the market, 2023. <https://googleprojectzero.blogspot.com/2023/11/first-handset-with-mte-on-market.html>.
- [40] G. P. Z. Mark Brand. Mte as implemented, part 1: Implementation testing, 2023. <https://googleprojectzero.blogspot.com/2023/08/mte-as-implemented-part-1.html>.
- [41] D. McKee, Y. Giannaris, C. O. Perez, H. Shrobe, M. Payer, H. Okhravi, and N. Burow. Preventing kernel hacks with hakc. In *Proceedings of the 2022 Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Apr. 2022.
- [42] M. Minkin, D. Moghimi, M. Lipp, M. Schwarz, J. Van Bulck, D. Genkin, D. Gruss, F. Piessens, B. Sunar, and Y. Yarom. Fallout: Leaking data on meltdown-resistant cpus. In *Proceedings of the 26th ACM Conference on Computer and Communications Security (CCS)*, London, UK, Nov. 2019.
- [43] MITRE. Cve-2018-3639, 2018. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2018-3639>.
- [44] MITRE. CVE-2019-2215., 2019. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2019-2215>.
- [45] Mitre. Cve-2020-6449, 2020. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2020-6449>.
- [46] Mitre. Cve-2022-0185, 2022. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2022-0185>.
- [47] Mitre. Cve-2023-5217, 2023. <https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2023-5217>.
- [48] K. Mitsunami. Delivering enhanced security through memory tagging extension, 2021. <https://community.arm.com/arm-community-blogs/b/architectures-and-processors-blog/posts/enhanced-security-through-mte>.
- [49] S. Narayan, C. Disselkoen, D. Moghimi, S. Cauligi, E. Johnson, Z. Gang, A. Vahldiek-Oberwagner, R. Sahita, H. Shacham, D. Tullsen, et al. Swivel: Hardening webassembly against spectre. In *Proceedings of the 30th USENIX Security Symposium (Security)*, Virtual, Aug. 2021.
- [50] O. Oleksenko, C. Fetzer, B. Köpf, and M. Silberstein. Revizor: Testing black-box cpus against speculation contracts. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Lausanne, Switzerland, Feb.-Mar. 2022.
- [51] J. Olšan. Cortex-x3: the new fastest core from arm (architecture analysis), 2022. <https://fuse.wikichip.org/news/6855/arm-unveils-next-gen-flagship-core-cortex-x3/>.
- [52] OP-TEE. libutils: add mte support in malloc() and friends, 2022. https://github.com/OP-TEE/optee_os/commit/08a5c4f9ae421384e52b87107283181e3fdd0f56.
- [53] A. Partap and D. Boneh. Memory tagging: A memory efficient design, 2022.
- [54] phakeobj. Gigacage, 2021. <https://phakeobj.netlify.app/posts/gigacage/>.
- [55] A. O. S. Project. Control flow integrity, 2022. <https://source.android.com/docs/security/test/cfi>.
- [56] T. C. Projects. Site isolation, 2024. <https://www.chromium.org/Home/chromium-security/site-isolation/>.
- [57] Qualcomm. Pointer authentication on armv8.3, 2017. <https://www.qualcomm.com/media/documents/files/whitepaper-pointer-authentication-on-armv8-3.pdf>.
- [58] J. Ravichandran, W. T. Na, J. Lang, and M. Yan. Pacman: attacking arm pointer authentication with speculative execution. In *Proceedings of the 49th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, New York, USA, June 2022.
- [59] N. L. Rocco. Arm-kerne 2022: Cortex-a715 und cortex-a510 refresh: Effizienz im fokus, 2022. <https://www.computerbase.de/2022-06/arm-cortex-x3-a715-a510-refresh/3/>.
- [60] saelo. V8 sandbox, 2023. <https://docs.google.com/document/d/1FM4FQmIhEqPG8uGp5o9A-mnPB5BOeScZYpkHjo0KKA8/edit?usp=sharing>.
- [61] T. Schlüter, A. Choudhari, L. Hetterich, L. Trampert, H. Nemati, A. Ibrahim, M. Schwarz, C. Rossow, and N. O. Tippenhauer. Fetch-bench: Systematic identification and characterization of proprietary prefetchers. In *Proceedings of the 30th ACM Conference on Computer and Communications Security (CCS)*, Copenhagen, Denmark, Nov. 2023.
- [62] M. Schwarz, C. Maurice, D. Gruss, and S. Mangard. Fantastic timers and where to find them: High-resolution microarchitectural attacks in javascript. In *Financial Cryptography and Data Security: 21st International Conference, FC 2017, Sliema, Malta, April 3-7, 2017, Revised Selected Papers 21*, pages 247–267. Springer, 2017.
- [63] Y. Shin, H. C. Kim, D. Kwon, J. H. Jeong, and J. Hur. Unveiling hardware-based data prefetcher, a hidden source of information leakage. In *Proceedings of the 25th ACM Conference on Computer and Communications Security (CCS)*, Tronto, Canada, Oct. 2018.
- [64] C. Song, B. Lee, K. Lu, W. Harris, T. Kim, and W. Lee. Enforcing kernel security invariants with data flow integrity. In *Proceedings of the 2016 Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, CA, Feb. 2016.
- [65] S. S. Stone, K. M. Woley, and M. I. Frank. Address-indexed memory disambiguation and store-to-load forwarding. In *Proceedings of the 38th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Barcelona, Spain, Dec. 2005.
- [66] C. Tice, T. Roeder, P. Collingbourne, S. Checkoway, Ú. Erlingsson, L. Lozano, and G. Pike. Enforcing forward-edge control-flow integrity in {GCC} & {LLVM}. In *Proceedings of the 23rd USENIX Security Symposium (Security)*, San Diego, CA, Aug. 2014.
- [67] Trustonic. Armv9-a: How our kinibi 600 trusted os utilises mte and ff-a features to create state-of-the-art tees, 2023.

<https://www.trustonic.com/technical-articles/armv9-a-how-our-kinibi-600-trusted-os-utilises-mte-and-ff-a-features-to-create-state-of-the-art-tees/>.

- [68] V8. Pointer compression in v8, 2020. <https://v8.dev/blog/pointer-compression>.
- [69] V8. v8/include/v8-internal.h, 2023. <https://github.com/v8/v8/blob/7161638e5ead74bf84a52d27e69ebda26fbd2416/include/v8-internal.h#L235C11-L235C11>.
- [70] J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx. Foreshadow: Extracting the keys to the intel {SGX} kingdom with transient {Out-of-Order} execution. In *Proceedings of the 27th USENIX Security Symposium (Security)*, Baltimore, MD, Aug. 2018.
- [71] J. Van Bulck, D. Moghimi, M. Schwarz, M. Lippi, M. Minkin, D. Genkin, Y. Yarom, B. Sunar, D. Gruss, and F. Piessens. Lvi: Hijacking transient execution through microarchitectural load value injection. In *Proceedings of the 41st IEEE Symposium on Security and Privacy (Oakland)*, Virtual, USA, May 2020.
- [72] S. Van Schaik, A. Milburn, S. Österlund, P. Frigo, G. Maisuradze, K. Razavi, H. Bos, and C. Giuffrida. Ridl: Rogue in-flight data load. In *Proceedings of the 40th IEEE Symposium on Security and Privacy (Oakland)*, San Jose, CA, May 2019.
- [73] J. R. S. Vicarte, M. Flanders, R. Paccagnella, G. Garrett-Grossman, A. Morrison, C. W. Fletcher, and D. Kohlbrenner. Augury: Using data memory-dependent prefetchers to leak data at rest. In *Proceedings of the 42nd IEEE Symposium on Security and Privacy (Oakland)*, San Francisco, CA, May 2022.
- [74] P. Vila, B. Köpf, and J. F. Morales. Theory and practice of finding eviction sets. In *Proceedings of the 40th IEEE Symposium on Security and Privacy (Oakland)*, San Jose, CA, May 2019.
- [75] O. Weisse, I. Neal, K. Loughlin, T. F. Wenisch, and B. Kasikci. Nda: Preventing speculative execution attacks at their source. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Columbus, Ohio, Oct. 2019.
- [76] J. Wikner and K. Razavi. {RETBLEED}: Arbitrary speculative code execution with return instructions. In *Proceedings of the 3125 USENIX Security Symposium (Security)*, Boston, MA, Aug. 2022.
- [77] H. Xiao and S. Ainsworth. Hacky racers: Exploiting instruction-level parallelism to generate stealthy fine-grained timers. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Vancouver, Canada, Apr. 2023.
- [78] J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas, and C. W. Fletcher. Speculative taint tracking (stt) a comprehensive protection for speculatively accessed data. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Columbus, Ohio, Oct. 2019.




Figure 12: Speculative V8 sandbox escape

```

1  ;; --- Optimized code ---
2  ;; optimization_id = 0
3  ;; kind = TURBOFAN
4  ;; name = TikTag_v2
5  ;; stack_slots = 6
6  ;; compiler = turbofan
7
8  ;; BR
9  ldrb w2, [x2] ;; x2 == &slow[0]
10 cbnz w2, #+0x64
11
12 add x2, x28, x2 ;; x2 == victim
13 ldur x3, [x2, #35] ;; victim.raw_length
14 lsr x3, x3, #29
15 ldur w4, [x2, #51] ;; victim.base_pointer
16 ldur x5, [x2, #43] ;; victim.external_pointer
17 add x5, x28, x5, lsr #27
18 fcvtzs x6, d0 ;; d0 == idx
19 cmn x6, #0x1 (1)
20 cinc x6, x6, vs
21 scvtf d1, x6
22 fcmp d1, d0 ;; check if idx is an integer
23 b.ne #+0x198
24 cmp x6, x3 ;; check if idx is in bounds
25 b.hs #+0x194
26
27 ;; CHECK
28 mov w3, w4
29 add x3, x3, x5
30 movz x7, #0x40 ;; PROBE_OFFSET
31 str w7, [x3, x6, lsl #2] ;; store
32 mov w3, w4
33 add x3, x3, x5
34 ldr w3, [x3, x6, lsl #2] ;; load
35
36 ;; TEST
37 movz x2, #0x400
38 movk x2, #0x39, lsl #32
39 ldrb w2, [x2, x3] ;; x2 == probe

```

Figure 13: Turbofan optimized assembly of TIKTAG-v2 gadget

Appendix A. Speculative V8 Sandbox Escape

The V8 sandbox [60] isolates JavaScript from the rest of the renderer by *pointer compression* and *guard regions* (Figure 12). Pointer compression enforces that all JavaScript pointers are compressed into a 32-bit or 35-bit index to represent a maximum 128 GB memory in Android, where memory addresses are calculated by adding the index to the sandbox base address. Guard regions are placed at the

```

1 function TikTagLeak(idx) {
2   for (Tg=0xf; Tg<=0xf; Tg++) {
3     res = TikTag_v2(idx | (Tg<<56)/8);
4     if (res == TAG_MATCH)
5       return Tg;
6   }
7   return -1;
8 }

```

(a) Tag Leakage

```

1 obj_vuln = AllocVuln();
2 obj_target = AllocTarget();
3 // obj_vuln.addr + offset
4 // = obj_target.addr
5 idxA = ObjToIdx(obj_vuln);
6 idxB = ObjToIdx(obj_target);
7
8 tagA = TikTagLeak(idxA);
9 tagB = TikTagLeak(idxB);
10
11 while (tagA != tagB) {
12   Free(obj_target);
13   obj_target = AllocTarget();
14   tagB = TikTagLeak(idxB);
15 }
16 // tagA == tagB
17 // corrupt obj_target
18 Exploit00B(obj_vuln, offset);

```

(b) Spatial bugs

```

1 obj_vuln = AllocVuln();
2 idx = ObjToIdx(obj_vuln);
3 tagA = TikTagLeak(idx);
4
5 Free(obj_vuln);
6 obj_target = AllocTarget();
7 // obj_vuln.addr == obj_target.addr
8 tagB = TikTagLeak(idx);
9
10 while (tagA != tagB) {
11   Free(obj_target);
12   obj_target = AllocTarget();
13   tagB = TikTagLeak(idx);
14 }
15 // tagA == tagB
16 // corrupt obj_target
17 ExploitUAF(obj_vuln);

```

(c) Temporal bugs

Figure 14: MTE bypass attacks against Chrome

start and end of the V8 sandbox with 32 GB size each, preventing potential out-of-bound access (i.e., worst case 32-bit index for the 8-byte element array). Thus, any native memory access instructions executing JavaScript are strictly limited to accessing the 4 GB boundary, and thus it cannot access the renderer’s memory.

Despite this strong memory isolation, we found the vulnerable case that the V8 sandbox does not guarantee complete isolation in the speculative paths. This vulnerable case is in the bound check implementation of TypedArray. TypedArray in JavaScript is an array type, supporting up to 32 GB array size [69]. The length of TypedArray is represented as 35 bits in V8, stored as a 64-bit field. TypedArray can be accessed using the index, which is a 64-bit value. Every TypedArray access using the index is preceded by a bound check, which compares the index against the length of the array. Specifically, given a 64-bit index, the bound check compares the whole 64-bit index against the 35-bit length. If the index is smaller than the length, the access is allowed; Otherwise, the access is denied.

The problem occurs when the bound check for the TypedArray access is speculatively executed. In this case, even if the index is larger than the array length, the access is still speculatively executed. This allows the untrusted JavaScript code to access beyond its restricted region, thereby escaping V8’s sandbox. To the best of our knowledge, this speculative sandbox escape vulnerability was first discovered by this work, and we accordingly reported this to the V8 security team in December 2023.

```

1 char ubuf[0x1000];
2 obj_vuln = SysAllocVuln();
3 obj_targ = SysAllocTarget();
4 oob_offset = GetOffset(obj_vuln, obj_targ);
5
6 while (1) {
7   SysTikTag00B(obj_vuln, oob_offset, ubuf);
8   if (Measure(ubuf) == TAG_MATCH) {
9     break;
10  }
11  SysFreeTarget(obj_target);
12  obj_target = SysAllocTarget();
13 }
14
15 // Now obj_vuln.tag == obj_target.tag
16 SysExploit00B(obj_vuln, oob_offset); // corrupt obj_target

```

(a) Spatial bugs

```

1 char ubuf[0x1000];
2 obj_vuln = SysAllocVuln();
3 SysFreeVuln(obj_vuln);
4 obj_target = SysAllocTarget();
5 // obj_vuln.addr == obj_target.addr
6
7 while (1) {
8   SysTikTagUAF(obj_vuln, ubuf);
9   if (Measure(ubuf) == TAG_MISMATCH) {
10    break;
11  }
12  SysFreeTarget(obj_target);
13  obj_target = SysAllocTarget();
14 }
15
16 // Now obj_vuln.tag == obj_target.tag
17 SysExploitUAF(obj_vuln); // corrupt obj_taret

```

(b) Temporal Bugs

Figure 15: Exploiting memory corruption bugs in the Linux kernel. Branch training and cache control are omitted for simplicity.

Appendix B. V8 Turbofan Optimized Assembly

Figure 13 shows the Turbofan optimized assembly of TIKTAG-v2 gadget in V8 JavaScript engine (Figure 8a). BR block contains the slow branch at line 9-10. CHECK block contains the PROBE_OFFSET store and load instructions at line 31-34, which are 2 machine instructions distance, satisfying the requirement of TIKTAG-v2 gadget (i.e., to be within the 5 instruction dispatch window). TEST block contains the load instruction using the forwarded value (i.e., val) at line 39.

Appendix C. Chrome V8 MTE Bypass Attack

Figure 14 shows examples of TIKTAG gadget-based MTE bypass attacks in the Chrome. A JavaScript function TikTagLeak() (Figure 14a) leaks the tag of the renderer memory with the TIKTAG-v2 gadget. This function brute-forces TikTag_v2() to check Tg against the tag Tm of target_addr. The assumption here is that the attacker implements ObjToIdx(), which returns idx, such that &victim[idx] points to the target_addr. With the leaked tags, the attacker can exploit spatial bugs (Figure 14b) and temporal bugs (Figure 14c) in the renderer without raising a tag check fault. The spatial bugs example is based on a




Figure 16: TIKTAG-v1 experiment by varying the BR and CHECK instruction distance (i.e., WINDOW)

linear buffer overflow vulnerability in the renderer process, CVE-2023-5217 [47], while the temporal bugs example is based on a use-after-free vulnerability in the renderer process, CVE-2020-6449 [45]. Both examples follow the same pattern: the attacker allocates renderer objects, leaks their tags, and triggers memory corruption only when the tags match.

Appendix D. Linux Kernel MTE Bypass Attack

Figure 15 demonstrates MTE bypass attacks using the TIKTAG gadgets in the Linux kernel. Figure 15a illustrates a spatial bug exploit, leveraging a kernel buffer overflow vulnerability (e.g., CVE-2022-0185 [46]), and Figure 15b illustrates a temporal bug exploit, leveraging a kernel use-after-free vulnerability (e.g., CVE-2019-2215 [44]). Functions prefixed with Sys denote system calls responsible for kernel object allocation, deallocation, or triggering the TIKTAG gadgets or memory corruption. Each vulnerable system call (i.e., SysExploit00B() and SysExploitUAF()) has their counterpart in TIKTAG gadget (i.e., SysTikTag00B() and SysTikTagUAF()) that leaks the tag check result through user space buffer (i.e., ubuf).

Appendix E. TIKTAG-v1 additional experimental results

Figure 16 shows the TIKTAG-v1 gadget experiment results by varying the number instructions (i.e., orr) between the BR and CHECK instructions (i.e., WINDOW). We experimented with the TIKTAG-v1 gadget in the same setting as in §5.1, but with the length of WINDOW varying from 0 to 30. We tested under two conditions: TEST containing a load instruction (Figure 16a) or a store instruction (Figure 16b). In both cases, the TIKTAG-v1 gadget showed cache hit rate difference when WINDOW is less than 25. The results were consistent when the orr instructions in WINDOW had no dependency, or had dependencies to CHECK, or even when they were replaced with nop instructions. Therefore, we think the time window between BR and CHECK requires around 25 instructions to trigger the tag check fault. As the affected core (i.e., Cortex-X3) has a 6-instruction dispatch window [51], we think the time for the TIKTAG-v1 to fetch the WINDOW instructions is around 5 cycles. We think that if CHECK triggers tag check faults within 5 cycles from BR, the CPU reduces the speculation execution and data prefetching, as described in §5.1.

Appendix F.

Meta-Review

The following meta-review was prepared by the program committee for the 2025 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

F.1. Summary

The paper presents two types of speculative execution attacks affecting ARM's MTE extension. The paper shows that ARM MTE exhibits a speculation shrinkage on an MTE tag mismatch, which can be exploited to brute force the tag value. The paper presents two gadgets that are capable of leaking the MTE tag of a given memory address. Finally, the paper presents two PoCs: for Chrome's v8 engine and the Linux kernel, for Google's Pixel 8 devices.

F.2. Scientific Contributions

- Identifies an Impactful Vulnerability
- Identifies an Impactful Vulnerability
- Provides a Valuable Step Forward in an Established Field

F.3. Reasons for Acceptance

- 1) The paper advances the understanding of speculative execution attacks on ARM's Memory Tagging Extension. It first performs a root-cause analysis of MTE tag leaks and experimentally shows that the cause of the tag leakage is due to speculation shrinkage leading to microarchitecture-level divergences, rather than the previously believed memory contention on spurious tag check faults. It discovers two variants of speculative execution attacks that affect ARM's MTE.
- 2) It presents end-to-end PoC attacks on the Linux kernel and Chrome's V8 engine, showing that these attacks are possible in the real world.
- 3) The prototype is open-sourced and available to other researchers, which advances research in the field of speculative execution.

F.4. Noteworthy Concerns

- 1) The reviewers were concerned about the cost of the proposed mitigations, which can be significantly high if implemented in software.