

SLasH-DSA: Breaking SLH-DSA Using an Extensible End-To-End Rowhammer Framework

✉ Jeremy Boy
University of Luebeck

✉ Antoon Purnal^{*}
Google

✉ Anna Pätschke
University of Luebeck

✉ Luca Wilke[†]
Azure Research, Microsoft

✉ Thomas Eisenbarth
University of Luebeck

Abstract

As quantum computing advances, *Post-Quantum Cryptography* (PQC) schemes are adopted to replace classical algorithms. Among them is the *Stateless Hash-Based Digital Signature Algorithm* (SLH-DSA) that was recently standardized by NIST and is favored for its conservative security basis.

In this work, we present the first software-only universal forgery attack on SLH-DSA, leveraging Rowhammer-induced bit flips to corrupt the internal state and forge signatures. While prior work targeted embedded systems and required physical access, our attack is software-only, targeting commodity desktop and server hardware, significantly broadening the threat model. We demonstrate full end-to-end attacks against SLH-DSA in OpenSSL 3.5.1, achieving universal forgery for the SHAKE-128f (deterministic), SHA2-128s, and SHAKE-192f (randomized) parameter sets after one hour (deterministic) or eight hours (randomized) of hammering and post-processing ranging from minutes to an hour, and showing theoretical attack complexities for most parameter sets. Our post-processing is informed by a novel complexity analysis that, given a concrete set of faulty signatures, identifies the most promising computational path to pursue.

To enable the attack, we introduce SWAGE, a modular and extensible framework for implementing end-to-end Rowhammer-based fault attacks. SWAGE abstracts and automates key components of practical Rowhammer attacks. Unlike prior tooling, SWAGE is untangled from the attacked code, making it reusable and suitable for frictionless analysis of different targets. Our findings highlight that even theoretically sound PQC schemes can fail under real-world conditions, underscoring the need for additional implementation hardening or hardware defenses against Rowhammer.

1 Introduction

With the accelerating progress in practical quantum computing, the cryptographic community is undergoing a major transition toward the adoption of *Post-Quantum Cryptography* (PQC). In response, NIST has released several PQC standards, including FIPS 205 [32] that describes the *Stateless Hash-Based Digital Signature Algorithm* (SLH-DSA) signature scheme (known as SPHINCS⁺ [6] prior to standardization). However, the adoption of strong cryptographic primitives alone is not sufficient to achieve security in real-world deployments. Besides being cryptographically sound, the implementation must also be able to withstand the various (micro)architectural side-channel and hardware-level attacks bedeviling modern CPUs.

One prominent and persistent threat is the *Rowhammer* vulnerability. Since its discovery in 2014 [25], Rowhammer has evolved from a reliability issue into a powerful attack vector capable of inducing targeted bit flips in *Dynamic Random Access Memory* (DRAM) through frequent accesses to adjacent memory rows. Despite years of developing countermeasures, Rowhammer remains a viable threat even on the most recent DDR5 memory modules [10, 13, 21, 22]. Over the past decade, it was shown how to leverage Rowhammer to tamper with isolation, escalate privileges, and leak secret data [4, 12, 18, 28, 38, 42]. Rowhammer-based attacks can also operate remotely via the browser or network requests [17, 30].

As quantum-resilient schemes become part of the critical infrastructure, it becomes vital to understand their resilience against practical fault attacks. In this work, we examine the vulnerability of the NIST-standardized post-quantum signature scheme SLH-DSA to Rowhammer-based fault attacks. As a hash-based scheme, SLH-DSA derives its resistance against large-scale quantum computers solely from the security properties of hash functions, making it an attractive and conservative choice. Unlike other standardized hash-based signatures (XMSS and LMS) [11], for which state-keeping is essential to maintain security, SLH-DSA is stateless, enabling it to act as a drop-in replacement for existing signa-

^{*}Work partially conducted while at PQShield.

[†]Work partially conducted while at University of Luebeck.

ture schemes. While prior work demonstrated that physical fault injection can enable universal forgeries against SLH-DSA [14], we show for the first time that a similar attack is feasible entirely in software, requiring no physical access. This significantly broadens the attack surface.

We demonstrate attacks against the SLH-DSA implementation in OpenSSL 3.5.1 [33], showing practical forgery for the SHAKE-128f (deterministic), SHA2-128s and SHAKE-192f (randomized) parameter sets. We show that an attacker can universally forge signatures for the SLH-DSA-SHAKE-192f parameter set after eight hours of Rowhammer attack and 1417 seconds of post-processing. Besides practical forgery attacks, we show theoretical attack complexities for most parameter sets. For SLH-DSA’s deterministic signing mode, which is not the default setting in OpenSSL but specified and allowed by the FIPS 205 standard, we achieve universal forgery for the SLH-DSA-SHAKE-128f parameter set after one hour of Rowhammer and 151 seconds of post-processing.

A given collection of faulty signatures leaves the attacker with several potential SLH-DSA subcomponents to target during the post-processing phase. To select the target with the lowest expected runtime, we derive, for the first time, the offline attack complexity for a *concrete* set of faulty signatures.

Performing end-to-end Rowhammer attacks in practice is non-trivial and entails several low-level challenges: First, the attacker needs to reverse-engineer the mapping from physical addresses to DRAM addresses for the CPU that they are using. Based on that, they must find a hammering pattern that bypasses the *Target Row Refresh* (TRR) mitigation on the used DIMM module. Hammering patterns require a certain amount of physically contiguous memory. However, unprivileged Linux users do not have an API that guarantees the allocation of such memory. Thus, for a realistic attack, the adversary needs to resort to certain quirks of the Linux memory allocator and side-channels in the memory subsystem of the CPU. Finally, the attacker needs to ensure that the victim application places the attackable data structure in a memory location they can manipulate through their hammering pattern. While prior work [2, 7, 21, 28, 37] has addressed these steps, the provided artifacts either only solve parts of the attack chain or are closely coupled with the attacked application.

To provide a holistic solution, we introduce SWAGE, a modular and extensible framework for implementing end-to-end Rowhammer-based fault attacks. SWAGE integrates modules for each of the mentioned attacks steps and can easily be applied to different victim applications. In summary, our framework and the presented attack highlight the need to address Rowhammer and similar threats in the post-quantum era.

1.1 Our Contributions

- We showcase a realistic Rowhammer attack against all parameter sets in OpenSSL’s SLH-DSA implementation in both deterministic and randomized signing modes,

including the end-to-end forgery of signatures for NIST security levels 1 and 3.

- We provide a detailed vulnerability analysis of SLH-DSA in the presence of Rowhammer-based faults.
- We derive the offline complexity given concrete faulty signatures, identifying the most promising forgery target.
- We implement SWAGE, a modular Rowhammer attack framework to ease the implementation of Rowhammer attacks. We make our source code available on Github ¹.

1.2 Responsible Disclosure

We officially reported our findings to the OpenSSL security team on August 04, 2025. Our report was acknowledged on August 07, 2025. We were informed that fault attacks are outside of OpenSSL’s threat model, and permission was granted to release our findings.

2 Background

This section provides the technical context for understanding our Rowhammer-based attack on the SLH-DSA scheme. We begin by describing the relevant characteristics of modern DRAM enabling Rowhammer attacks. We then summarize the structure and operation of the SLH-DSA signature scheme, highlighting components relevant to our attack strategy.

2.1 Memory Management

Physical memory is implemented as an array of *Dynamic Random Access Memory* (DRAM) cells, arranged in rows and columns within banks, which are in turn grouped into ranks. Each bank includes a row buffer that contains the value of the most recently accessed row to speed up repeated read operations. A DRAM cell consists of a capacitor and a transistor; the capacitor’s charge stores the bit’s value, and the transistor provides access to the stored value. As the charge leaks over time, the DRAM cells must periodically be refreshed.

Modern operating systems (OS) use virtual memory to manage physical memory through a layer of indirection. Processes work with virtual addresses that are then translated to physical addresses by the *Memory Management Unit* (MMU) with the help of page tables. After the translation of the virtual address to a physical address, the physical address is mapped to a physical location in DRAM. This mapping is CPU-specific and proprietary but can be reverse-engineered [35, 44].

When a process requests memory, the OS needs to select the physical pages to back it. Linux relies on the buddy allocator strategy, which manages contiguous physical memory in buckets of increasing size. These can be split for smaller allocations and merged again when memory gets freed [7].

¹<https://github.com/UzL-ITS/SLasH-DSA>

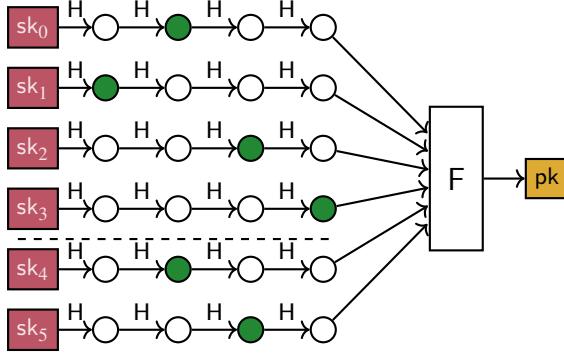


Figure 1: A WOTS⁺ signature for message $m = (1, 0, 2, 3)$ with checksum $c = (1, 2)$. H and F are hash functions, nodes represent hash function outputs. The nodes that are composed to produce the signature are shown in green. The public key is a hash of the chain endpoints. Here, chain length $w = 4$, message length $\ell_1 = 4$, and checksum length $\ell_2 = 2$.

2.2 Rowhammer

Rowhammer, discovered by Kim et al. [25], is a hardware vulnerability in DRAM that allows an attacker to flip bits in inaccessible *victim rows* by repeatedly accessing a set of neighboring *aggressor rows* that are under their control. After having mostly been considered a reliability issue at first, Rowhammer has been used to attack cryptographic schemes and circumvent OS isolation [4, 12, 18, 28, 38]. Manufacturers implemented various countermeasures such as *Target Row Refresh* (TRR) in DDR4 memory, which identifies unusual memory access patterns and then issues refresh commands for suspected *victim rows*. However, research has shown that more complex hammering patterns can overload the access tracking in TRR, re-enabling Rowhammer attacks [13, 21]. While the DDR5 standard has more principled mitigation [8, 23], recent work [22] reports to have found at least one vulnerable setup.

2.3 SLH-DSA

SLH-DSA is a hash-based digital signature scheme standardized by NIST in 2024 [32]. The building blocks of SLH-DSA are again hash-based signature schemes: (1) the few-time signature scheme *Forest of Random Subsets* (FORS), and (2) a hierarchical organization of the many-time signature scheme *eXtended Merkle Signature Scheme* (XMSS). The XMSS trees, in turn, each manage a number of hash-based one-time signature scheme *Winternitz One-Time Signature Scheme*⁺ (WOTS⁺) instances. Although FORS is an essential component of SLH-DSA, its specification is not relevant to understand the attack described in this paper. For clarity, we omit its specification and refer to the FIPS 205 standard [32].

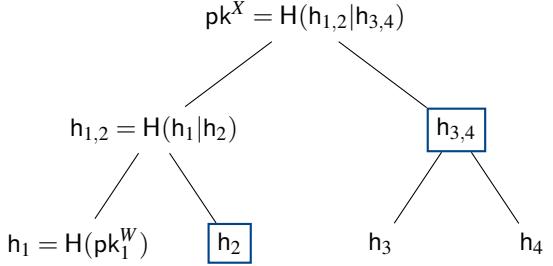


Figure 2: A Merkle tree of height 2 with $2^2 = 4$ leaves. The leaf h_i represents the hash of a WOTS⁺ key pk_i^W . A parent node $h_{i,j}$ represents the hash of the concatenation of its children h_i, h_j . The final hash is the root node pk^X and represents the public key. The nodes forming the authentication path $auth(pk_1^W)$ are highlighted with blue rectangles.

2.3.1 The Winternitz One-Time Signature Scheme⁺

The *Winternitz One-Time Signature Scheme*⁺ (WOTS⁺) is a hash-based one-time signature scheme. It consists of a series of hash chains, for which the starting points, together, can be considered to be equivalent to the WOTS⁺ secret key. To sign a message, a hash function is applied to each starting point a message-dependent number of times, as depicted in Figure 1. The message to be signed is appended with a checksum on this message, to avoid, e.g., an attacker to forge a signature on $m_1 = (1, 1, 2, 3)$ given the signature on $m_0 = (1, 0, 2, 3)$. A WOTS⁺ instance is characterized by the security parameter n , indicating the output size of the hash function, and the message block size lg_w , which determines the chain length $w = 2^{lg_w}$. Derived from lg_w and n are the number of message chains ℓ_1 and checksum chains ℓ_2 .

2.3.2 eXtended Merkle Signature Scheme (XMSS)

XMSS is a hash-based signature scheme that is able to sign more than one message by organizing a large number of WOTS⁺ key pairs as a Merkle tree (Figure 2). The WOTS⁺ public keys are the leaves of the Merkle tree. An XMSS signature on a message consists of the WOTS⁺ signature on the message, and the authentication path that allows a verifier to recompute the Merkle tree root from the WOTS⁺ public key that was used, confirming its legitimacy. Although XMSS can sign more than one message, it must be ensured that the same WOTS⁺ instance is not used to sign more than one message, i.e., XMSS is a *stateful* scheme. An XMSS instance is characterized by the security parameter n and tree height h' , implying $2^{h'}$ WOTS⁺ key pairs as leaves.

XMSS instances can be organized in a hierarchical fashion as *hypertrees*, or tree of trees. In this case, the Merkle tree root (and hence the public key) of XMSS level l is signed by one of the leaves of XMSS level $l + 1$.

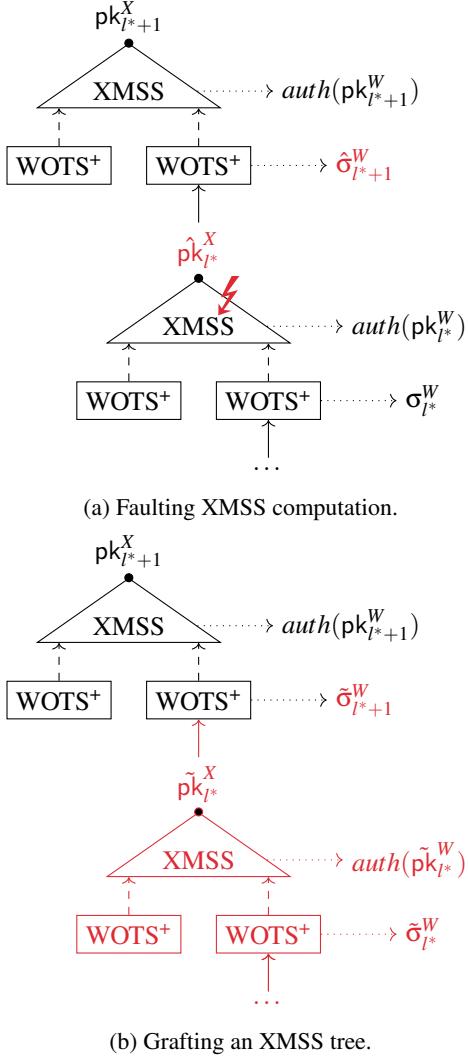


Figure 3: Grafting tree attack on SLH-DSA [14].

2.3.3 SLH-DSA

SLH-DSA uses FORS and XMSS as building blocks to construct a practical scheme that can sign many different messages and is *stateless*, i.e., SLH-DSA does not require keeping track of which secret keys were already used.

First, the message is signed using a FORS key pair. The FORS signature's public key is signed using an XMSS hypertree. The root node of the top-level hypertree corresponds to the public key of the scheme. The particular FORS key pair to use and the resulting path through the hypertree of XMSS trees is determined by a value R , which is the output of a hash of the secret key and the message to be signed (deterministic mode), or the secret key, the message, and additional user-provided entropy (randomized mode). An SLH-DSA signature consists of a FORS signature of the message and a series of XMSS signatures proving the FORS key's validity.

SLH-DSA is standardized with the hash functions SHA-2 and SHAKE. The standard describes both *small* and *fast* variants (trading off the number of XMSS levels and tree size) and in different security categories (with, respectively, 128-bit, 192-bit, and 256-bit hash output lengths). For the detailed specification of each variant, we refer to FIPS 205 [32].

2.4 Grafting Tree Attack

Under correct operation, the design of SLH-DSA ensures that each WOTS⁺ secret key is only used to sign *one single* public key root. This invariant is crucial to uphold, since the security of WOTS⁺ degrades rapidly as the number of signatures with the same secret key increases. The *grafting tree attack* [9] builds on this observation. By triggering faulty computations during signature generation, it causes the same WOTS⁺ secret key of a given XMSS level to sign multiple different public key roots of previous XMSS levels. Hence, it exposes additional WOTS⁺ secret values (cf. Figure 1), rendering the WOTS⁺ instance *compromised*. In particular, during the computation of an XMSS tree, the attacker injects faults such that the root of the tree gets corrupted (Figure 3a). Note that any fault during the entire XMSS computation leads to the corruption of its root. The faulty root is then signed by the next level in the hypertree using WOTS⁺. When the process is repeated, the attacker can gather multiple signatures with the same WOTS⁺ secret key.

The attacker can use a compromised WOTS⁺ instance as in Figure 3b. With enough signatures, the attacker can use a compromised WOTS⁺ instance (in layer $l^* + 1$) to sign a *grafted* (i.e., fully attacker-generated) XMSS tree (in layer l^*). There is some offline computation effort to arrive at a grafted tree for which the public key is compatible with the values that are leaked from the compromised WOTS⁺ instance. As each additional signature potentially reveals new WOTS⁺ chain values, obtaining more faulty signatures that visit the same WOTS⁺ key pair reduces the tree grafting complexity. The most straightforward version of the attack uses a valid (uncorrupted) signature in addition to the faulted signature(s), but prior work showed that this is not a strict requirement [14].

From the point of view of the signature verifier, the grafted tree appears valid, as it contains a valid signature from the WOTS⁺ key pair in layer $l^* + 1$. Since the grafted tree is fully attacker-controlled, including its own secret keys, an attacker uses a grafted tree to produce a forgery for any chosen message that visits the compromised WOTS⁺ node.

3 Threat Model

We adopt the standard Rowhammer attack model that differentiates between an offline phase and an online phase. During the offline phase, the attacker has unfettered access to a replicated target system to prepare the attack.

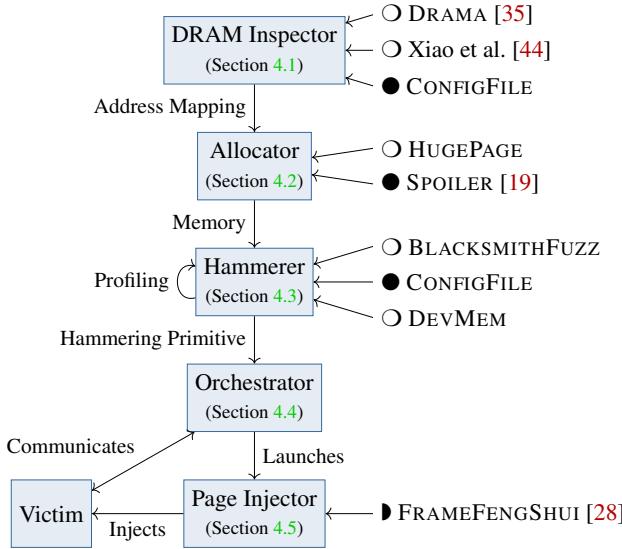


Figure 4: Flowchart showing the steps of a Rowhammer attack using SWAGE. \circ depicts offline phase usage of the strategy, \bullet stands for online usage and \blacktriangleright is a strategy used in both offline and online phase of an attack.

In the online phase, the attacker’s access is heavily restricted, as detailed in the following. The attacker can execute code (e.g., SWAGE) with user-level privileges on the target system. They have access to high-resolution timers like `rdtscp`.

The target system uses DDR4 memory, possibly with TRR mitigation. The attacker is able to co-locate their code with the victim on the same logical CPU core [2, 28, 36]. Finally, we assume that the attacker can influence when the victim process allocates the memory for the targeted data structure, as, e.g., enabled by the fork-on-connection strategy used in security-focused server processes like OpenSSH [28] or swapping due to memory pressure [36].

4 Swage: A Modular Universal End-to-End Rowhammer Attack Framework

Conducting Rowhammer attacks in practice is challenging: As the attack interacts with different aspects of the target system, many moving parts need to be considered. Though a number of tools have been proposed to facilitate Rowhammer attacks, they often cover only a specific aspect of the attack, such as finding reproducible hammering patterns or reverse-engineering the physical memory layout [7, 13, 21, 35]. Most of the time, the software design of these tools is tightly coupled to the evaluation scenario of the respective paper, making it difficult to reuse them in an end-to-end Rowhammer attack. We propose SWAGE, a novel modular end-to-end framework for Rowhammer attacks. SWAGE is implemented in Rust, C, and Python, and is designed to be modular and extensible. We

provide a SWAGE artifact as an open-source project and invite the community to contribute to its development.

Figure 4 shows the modular architecture of SWAGE. Most modules have multiple instantiation options that allow choosing between easy prototyping of attacks by using root privileges or performing realistic attacks that only require user space access. In the following, we describe each module and its role in the attack workflow in more detail.

4.1 DRAM Inspector

In order to conduct a Rowhammer attack, the attacker needs information about the locality in memory to arrange the aggressor rows and the victim row for their attack. The DRAM INSPECTOR module provides information about the physical memory layout of the target system. When accessing a memory location, the memory controller applies a *memory mapping* function to the physical address. This function is specific to any given CPU generation and usually not disclosed by manufacturers. However, there exist side-channel attacks that can be used to reverse-engineer the physical memory layout. For this, SWAGE employs a software suite combining the outputs of both DRAMA [35] and Xiao et al. [44] for improved reliability. This step only needs to be done once per targeted CPU during the offline phase. Afterward, the mapping can be loaded from a configuration file.

4.2 Allocator

To construct the memory access patterns required for Rowhammer, the attacker needs to be able to allocate physically contiguous memory. The ALLOCATOR module provides a unified interface for allocating contiguous memory blocks. SWAGE comes with two implementations for allocating these blocks, supporting both attack prototyping with elevated privileges and realistic scenarios. For the former, the attacker can configure the Linux kernel to reserve memory for huge pages at boot time and then allocate them via the HUGEPAGE instantiation of the ALLOCATOR module to get access to vast amounts of physically contiguous memory as demonstrated by prior work [21].

For realistic scenarios, the attacker has access to the SPOILER instantiation. While the Linux allocator does not make any guarantees on physical contiguous memory for larger chunks of allocated memory, its use of the buddy allocator strategy makes it highly likely that some parts of the allocated memory are indeed physically contiguous [28]. Following [2], we use the Spoiler side-channel [19] to reveal all 1 MiB blocks of contiguous memory inside our allocated buffer. Note that their approach does not give information on the positioning of these 1 MiB blocks relative to each other. As discussed later, 1 MiB of contiguous memory may not always suffice to find a hammering pattern. Thus, we use the row buffer side-channel in combination with our already obtained

information about the DRAM addressing function to reveal which of the discovered 1 MiB blocks are contiguous and form 4 MiB blocks of contiguous memory.

4.3 Hammerer

The HAMMERER module is the core of SWAGE and implements strategies to fuzz and replay Rowhammer patterns. The BLACKSMITHFUZZ instantiation is a fork of the Blacksmith fuzzer [21]. It is intended to be used in conjunction with the HUGEPAGE memory allocation strategy to find reproducible hammering patterns during the offline phase of the attack. Our analysis shows that Blacksmith tends to generate patterns that require large amounts of physically contiguous memory, making them unsuitable in a realistic attack scenario. Thus, we added a novel component that breaks up the hammering pattern into chunks that only require 1 or 4 MiB of contiguous memory and re-evaluates their reproducibility. Rapid attack prototyping in the offline phase can be achieved with the DEV-MEM backend: Instead of performing Rowhammer, it allows to architecturally flip a bit in the targeted location by using the Linux kernel’s `/dev/mem` interface to access arbitrary physical memory. In the online phase, the CONFIGFILE backend is used to load and hammer a previously found pattern.

4.4 Orchestrator

The ORCHESTRATOR module implements the actual attack against a victim process. As input from the previous stages, it has access to a hammering primitive and a mapping to the memory page \mathcal{P} at which the hammering primitive will cause a bit flip. The details of the victim process are abstracted by a victim-specific wrapper that needs to implement four methods: `start`, `init`, `check`, `stop`. To start the attack workflow, the ORCHESTRATOR in parallel calls the `start` method and the page injector. The `start` method may be used to trigger the creation of the victim process or other initialization steps. The page injector ensures that the victim uses the vulnerable page \mathcal{P} to store the data that the attacker wants to manipulate via Rowhammer and will be discussed in Section 4.5. Afterward, the ORCHESTRATOR calls the `init` method to synchronize with the execution of the victim and start the Rowhammer attack. The synchronization step could, e.g., involve triggering the victim via a network request or the use of a cache attack. Next, the `check` method is called to ascertain the attack’s success, assuming the victim creates some kind of output or otherwise observable behavior. Finally, the `stop` method terminates the attack cycle.

4.5 Page Injection

Page injection refers to techniques to manipulate the Linux memory allocator to return a specific, attacker-chosen memory page upon the next allocation. In the context of Rowham-

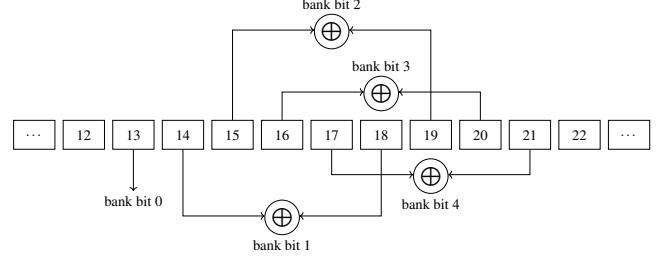


Figure 5: DRAM bank-selection mapping for the target system. Bank bit 0 is taken directly from physical bit 13, while bank bits 1-4 are computed as the XOR (\oplus) of bit pairs ($b_{14} \oplus b_{18}$, $b_{15} \oplus b_{19}$, $b_{16} \oplus b_{20}$, and $b_{17} \oplus b_{21}$), respectively. Dots indicate omitted neighboring physical address bits.

mer, this enables the attacker to ensure that the victim places the data structure *exactly* on the page that the attacker can introduce bit flips in. SWAGE implements a technique from prior work [2, 28] that exploits an optimization in the Linux kernel allocator. The optimization consists of a per-CPU list of memory pages that is consulted during memory allocation before going to the main memory allocator. This can be exploited to inject a target page into a victim process co-located on the attacker’s core. A recent novel page injection technique, RUBICON [7], combines page injection and eviction strategies to implement cross-core page injection.

5 Evaluation of SWAGE

During the offline part of a Rowhammer attack, the attacker collects information about the target. We evaluated the components of SWAGE, experimentally verifying the reliability of the attack primitives presented in Section 4. For evaluating an attack on SLH-DSA, we conducted our SWAGE experiments on a system with an Intel i5-6400 CPU and a single G.Skill AEGIS 16 GB DDR4-2133 DIMM, running Ubuntu 20.04.6 with Linux kernel version 6.8.12-generic with default settings. We target an implementation with attackable data residing on the stack; the concrete instantiations of such data are discussed in Section 6. In the following, we describe the results of our SWAGE analysis.

5.1 Reverse-Engineering the DRAM Address Mapping

For a Rowhammer attack, an attacker needs to determine the mapping of physical addresses to DRAM banks, rows, and columns. SWAGE uses the row conflict timing measurements to determine the physical mapping function running the algorithms from both DRAMA [35] and Xiao et al. [44] for improved reliability. The row mapping function for the target system is simple: the row index is determined by physical bits 18 to 29. Similarly, the column index is determined by

physical bits 0 to 12. Figure 5 shows the reverse-engineered bank mapping function for the target system. We verify the findings by checking the timing of pairs of addresses with same or different bank indices. As it shows no significant outlier behavior, the mapping function is considered correct.

5.2 Contiguous Memory Allocation

We evaluated the SPOILER implementation in SWAGE by repeatedly allocating a 4 MiB block using the ALLOCATOR module. We configure the allocator to search for contiguous memory blocks in a 2 GiB buffer. Repeating the allocation experiment 100 times, we measured the elapsed time for the memory allocation and the false-positive rate, i.e., the number of times the allocator falsely identifies a physically contiguous 4 MiB block via the SPOILER timing side-channel. On the target machine, we find that an allocation using SPOILER takes 41.6 seconds on average, with a false positive rate of 66%. However, the falsely identified 4 MiB memory blocks consist of either a pair of 2 MiB blocks, a quartet of 1 MiB blocks, or one 3 MiB block and one 1 MiB block. While technically non-contiguous, those *composite blocks* can still be used in a Rowhammer attack if they are consistent with the bank timing function, i.e., if they show the expected timing measurements due to row buffer conflicts.

5.3 Page Injection

Evaluating the page injection technique in the SWAGE PAGE INJECTOR, we find that the page injection is successful in most of the cases. This is in line with findings in previous work [4, 28]. When the page injection fails, the attacker restarts the online phase of the attack and tries again. Those failed page injection attempts most likely happen due to noise caused by allocations from different processes on the same CPU. As soon as the page injection is successful, the attacker can proceed to collect signatures from the signing server.

5.4 Reproducible Hammering Patterns

SWAGE employs BLACKSMITH [21] to find reproducible hammering patterns in the HAMMERER module. BLACKSMITH is a fuzzing-based software suite to find reproducible hammering patterns. It generates *non-uniform access patterns* by randomizing three domains: How often an aggressor row is activated (frequency), the time between the start of the hammering pattern and the first activation of an aggressor row (phase), and how often an aggressor row is activated back-to-back (amplitude). By using BLACKSMITH, SWAGE can effectively bypass TRR, enabling Rowhammer attacks in state of the art DDR4 memory modules. The hammering pattern from BLACKSMITH can be imported in SWAGE, where the pattern is just-in-time compiled to an assembly function to maximize the hammering performance.

Table 1: Reproducibility results of BLACKSMITHFUZZ. Each candidate pattern was tested over 10 rounds, each allowing up to 1000 pattern repetitions or terminating early upon the first encountered bit flip.

Pattern	Bit Flips		Retries		Time (s)	
	Avg.	Max	Avg.	Max	Avg.	Max
\mathcal{A}	1.3	2	1.4	4	1.22	3.54
\mathcal{B}	0.0	0	1000.0	1000	901.10	901.45
\mathcal{C}	0.0	0	1000.0	1000	883.49	885.65
\mathcal{D}	2.1	6	8.9	26	7.93	23.13
\mathcal{E}	1.1	2	271.6	780	236.13	678.13
\mathcal{F}	0.2	1	863.9	1000	774.79	896.88
\mathcal{G}	1.0	1	23.7	51	20.23	43.53
\mathcal{H}	1.1	2	10.8	34	9.52	29.97

We evaluated BLACKSMITHFUZZ on the target machine, summarizing the results of an eight hour fuzzing run in Table 1. The table shows reproducibility results for eight different hammering patterns showing at least one bit flip in the fuzzing run. We see from the table a varying degree of reproducibility: While patterns \mathcal{B} and \mathcal{C} are not reproducible at all, patterns \mathcal{A} , \mathcal{D} , and \mathcal{H} are highly reproducible. However, pattern \mathcal{D} usually causes bit flips in multiple locations in DRAM. While this can be a desired property for some attack scenarios, the rogue flips potentially lead to system instability. Therefore, for the SLash-DSA attack, we chose pattern \mathcal{G} due to its high reproducibility and controlled bit flipping behavior.

5.5 Reproducibility of Hammering Patterns With Block Splitting

As discussed in Section 4.3, SWAGE needs to adapt a hammering pattern found with BLACKSMITHFUZZ for use with distributed memory blocks. Block splitting allows the attacker to split the aggressor mapping of an access pattern into smaller memory blocks eliminating the need for vast amounts of physically contiguous memory. For the splitting process, the 1 GiB huge page of physically contiguous memory used by BLACKSMITHFUZZ is divided into 4 MiB chunks and corresponding 4 MiB memory regions are allocated – using SPOILER in this case – on the target system. This process preserves the local properties of the patterns within each 4 MiB chunks but destroys the global pattern, as the individual 4 MiB blocks are (w.h.p.) not adjacent to each other in physical memory. However, as gathering a large amount of 4 MiB blocks using the SPOILER approach is time-consuming, we also limit the size of the patterns that the fuzzer is allowed to generate. For the pattern \mathcal{G} that we selected in the preceding subsection, we require a total of ten 4 MiB chunks, each containing a subset of the aggressors. The reproducibility experiment confirms that the pattern remains effective even when applied to these smaller memory blocks.

5.6 Combined SWAGE Validation

In a last offline reproducibility experiment, the attacker deallocates the target page before initiating the hammering process, and then allocates it in a *dummy process*. This dummy process reserves a large stack-based array, fills it with a predetermined value, and then reads the array to verify whether the contents of the target page match the expected value before rewriting the array. This approach mimics the online phase of the attack, where a victim process is anticipated to interact with the target page. The results of this assessment demonstrate that the pattern maintains the reproducibility outlined in Table 1, underscoring its resilience to page walks, page table flushes, and other memory management operations carried out by the operating system. With that, SWAGE has successfully determined a reproducible access pattern and aggressor mapping that can then be used to induce bit flips in the target page.

6 SLasH-DSA Attack

We now present the SLasH-DSA attack, a novel end-to-end Rowhammer attack against SLH-DSA. We take previous work by Genêt [14] as a starting point for our discussion and refer to their work for further insights into complexity considerations not covered in this work.

Our work extends Genêt’s attack in several key aspects: While Genêt demonstrated the tree grafting attack using hardware-based clock glitching on embedded systems, we show for the first time that the attack is feasible with software-only Rowhammer on commodity hardware. In his practical experiments, Genêt targeted the SHAKE-256s parameter set in a modified SPHINCS⁺ reference implementation, blindly injecting clock glitches while computing XMSS trees and WOTS⁺ signatures. In contrast, we target all parameter sets in OpenSSL’s implementation of the now-standardized SLH-DSA, specifically attacking the `lnode` buffer in the `xmss_node` function during XMSS tree computation (Section 6.1). This allows us to exploit the extended lifetime of `lnode` on the stack, providing larger temporal windows for Rowhammer-induced bit flips. Beyond the attack vector and target, we introduce a novel exact complexity analysis for concrete compromised WOTS⁺ instances, enabling optimal candidate selection, whereas Genêt analyzed the average case complexity for uniformly distributed faults.

We first describe the fault analysis of the OpenSSL SLH-DSA implementation, then continue with the signature generation phase, where we leverage the SWAGE framework to carry out the attack, and finally describe the post-processing phase with the tree grafting attack.

Algorithm 1: `xmss_node`(`sk, i, z, pk, ADRS`) [32]

Input: Secret seed $sk = \mathbf{SK}.\text{seed}$, target node index i , target node height z , public seed pk , address $ADRS$

Output: n -byte root *node*

```

1 if  $z = 0$  then
2    $ADRS.\text{setTypeAndClear}(\text{WOTS\_HASH})$ 
3    $ADRS.\text{setKeyPairAddress}(i)$ 
4    $node \leftarrow \text{wots\_pkGen}(sk, pk, ADRS)$ 
5 else
6    $lnode \leftarrow \text{xmss\_node}(sk, 2i, z - 1, pk, ADRS)$ 
7    $rnode \leftarrow \text{xmss\_node}(sk, 2i + 1, z - 1, pk, ADRS)$ 
8    $ADRS.\text{setTypeAndClear}(\text{TREE})$ 
9    $ADRS.\text{setTreeHeight}(z)$ 
10   $ADRS.\text{setTreeIndex}(i)$ 
11   $node \leftarrow H(pk, ADRS, l \parallel r)$ 
12 return node

```

6.1 Fault Analysis of the OpenSSL SLH-DSA Implementation

In the hypertree phase of the SLH-DSA signature generation, `xmss_node` (Algorithm 1) is called repeatedly to compute the root nodes of the Merkle trees identified by addresses $ADRS$, a public seed pk , and a secret sk . In OpenSSL, `xmss_node` is implemented in file `crypto/slh_dsa/slh_xmss.c`. It takes as input the secret and public seed, the start index, the target height, and the $ADRS$ of the current tree. The function computes the root node of the Merkle tree bottom up, starting at the left-most WOTS⁺ node and working its way up to the root, closely matching the recursive depth-first search algorithm from the standard. The function recursively calls itself with the subtree corresponding to the current node’s children until the bottom of the tree is reached. Once at the bottom, the two leftmost WOTS⁺ public keys are computed from the secret seed, and the resulting public keys are stored in `lnode` and `rnode`. After returning control flow to the parent, the siblings in `lnode` and `rnode` are concatenated, hashed, and passed to the grandparent. This procedure is repeated until the top of the tree, where the root node acts as the scheme’s public key.

We analyze the implementation of `xmss_node` in OpenSSL to find suitable memory regions for our Rowhammer attack. Specifically, we are interested in high spatial suitability, i.e., the buffers in question are large enough to provide flexibility in the exact bit to be flipped, and high temporal suitability, i.e., the memory location provides a large enough temporal window in which the bit flip is able to occur. For spatial suitability, we observe that `lnode` holds the n -byte-output of hash function H . When computing the topmost `lnode` in a tree, the algorithm descends into $h' - 1$ levels, each maintaining its own `lnode`. In total, this makes a memory region of $n \cdot (h' - 1)$ bytes. Note that the spatial suitability increases with higher

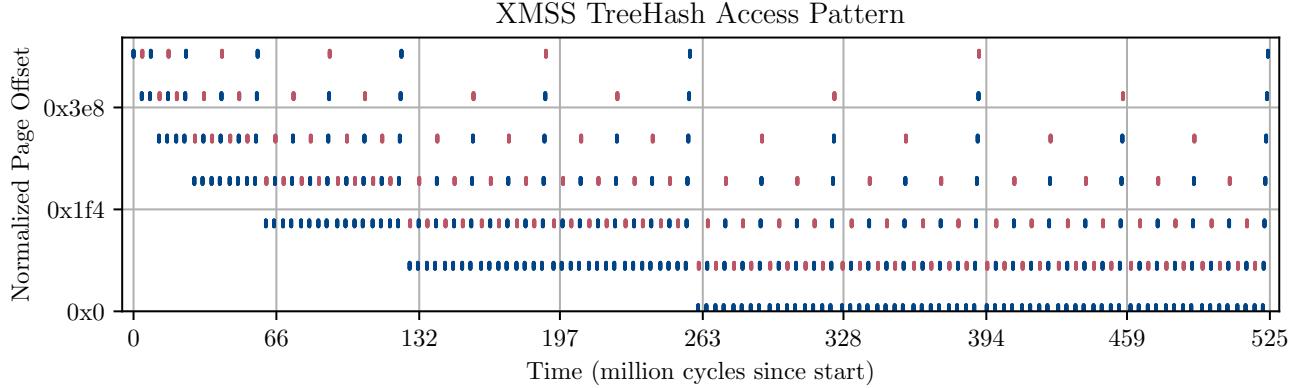


Figure 6: Access patterns for `lnode` in the `ossl_slh_xmss_node` function in the OpenSSL SLH-DSA implementation. The figure shows all accesses to the `lnode` buffers across all tree levels during generation of an XMSS public key with the SHA2-256s parameter set. Write accesses are shown in red, read accesses in blue.

security levels, as n is 32 bytes for security category five, but only 16 bytes for security category one.

For the temporal suitability of `lnode`, we assume the signing party currently executes `xmss_node` to compute the public key pk^X of an XMSS tree and just compute `lnode` in line 6 at the topmost layer of the tree. From the pseudocode, we see that the left subtree of the current node is processed before the right one, similar to a depth-first search. Meanwhile, the topmost `lnode` remains on the stack for a significant portion algorithm’s runtime, which makes it a prime target for a Rowhammer attack. We see that the temporal suitability decreases at lower levels in the XMSS tree. Therefore, the larger the size of the XMSS tree, the lower the temporal precision of inducing the bit flip needs to be. This makes the “fast” variants of SLH-DSA less temporally suitable. In an attack, this issue can be tackled by applying performance degradation techniques discussed in Section 6.2.

Figure 6 shows the results of an experimental verification of this observation, showing the access patterns of `lnode` while computing an XMSS tree in an instrumented build of OpenSSL. After a warm-up phase for the recursive calls to saturate the layers of the tree, the topmost page offsets in the figure correspond to the accesses to `lnode` in the top layer, and lower page offsets are accesses in lower layers. We observe that lower offsets in the array are written to and read from frequently, the higher offsets are accessed only few times per layer. We conclude that `lnode` is a both spatially and temporally suitable target for a Rowhammer attack.

The information gathered in the fault analysis of OpenSSL’s SLH-DSA implementation guides the implementation of the SLash-DSA orchestrator module in SWAGE. This module incorporates attack-specific code and framework configuration for SLash-DSA, while allowing to reuse SWAGE’s core components described in Section 4.

6.2 Performance Degradation & Signature Collection

After a fault analysis of OpenSSL, the attacker starts the online phase of the attack, where they apply memory massaging techniques to inject a target page into the victim process. They then hammer the target page while continuously collecting signatures from the signing server until a timeout is reached.

Performance Degradation. If available, OpenSSL utilizes hardware acceleration and vector instructions when signing messages with SLH-DSA. This increase in performance is most notable for the lower security levels of SHA-2 and *fast* variants of SLH-DSA. To still allow Rowhammer-induced bit flips against these variants, the signing server is slowed down using a *performance degradation attack*, artificially slowing down the execution of a target program, a technique also employed by related work [1, 4]. In SLash-DSA, the attacker occupies the victim’s core by using the `stress` tool, which spawns a parameter-dependent number of CPU intensive worker threads. Depending on the parameter set, they use as little as 15 workers for the “small” variants of SLH-DSA and 63 workers for the “fast” variants. In scenarios where this amount of additional CPU load is unacceptable, performance degradation attacks like HYPERDEGRADE [3] or CONTROLLED PREEMPTION [45] pose viable alternatives.

Collecting Signatures. The attacker now collects signatures from the signing server while concurrently hammering the target memory region. Due to the performance degradation attack, signing times are significantly prolonged. In our evaluation, depending on the parameter set, signature collection rates range from approximately 95 to 1917 signatures per hour. Given the spatial locality of Rowhammer-induced bit flips within DRAM rows, the attacker seeks corrupted sig-

natures where faults affect the same WOTS⁺ instance across multiple signature attempts, as these enable the tree grafting attack described below. The SWAGE framework orchestrates this process, coordinating the Hammerer module’s memory access patterns with signature collection until the configured timeout is reached.

After collecting signatures from the victim process, the attacker performs the grafting tree attack in an offline post-processing phase. In that phase, the attacker searches for XMSS trees matching the compromised WOTS⁺ instances (cf. Section 2.3) they collected from the victim (see Figure 3). The methodology for post-processing described below was first introduced by Castelnovi et al. [9].

6.3 Identifying Secret WOTS⁺ Values

In this step, the attacker identifies WOTS⁺ secret values from the collected signatures. First, the WOTS⁺ signatures (Figure 1) are extracted from the collected signatures and grouped by **ADRS**. The attacker then identifies the signatures for which the corresponding addresses map to the same WOTS⁺ instance. Subsequently, they calculate the exposed WOTS⁺ secret values for each **ADRS**. Let $(\sigma^W, \hat{\sigma}^W)$ be two distinct WOTS⁺ signatures for the same WOTS⁺ instance at **ADRS** with target layer l^* , and let m_i be the message chunk corresponding to the message signed by σ^W in chain i . The attacker identifies WOTS⁺ secret values in $\hat{\sigma}^W = (\hat{\sigma}^{(0)}, \dots, \hat{\sigma}^{(\ell-1)})$ using the exhaustive search described by Castelnovi et al. [9]:

1. For each $\hat{\sigma}^{(i)} \in \hat{\sigma}$: find $1 \leq k < m_i$ such that

$$\text{chain}\left(\hat{\sigma}^{(i)}, k, w - 1 - k, \text{pk}, \text{ADRS}\right) = p_i.$$

2. If such k exists, then $\hat{\sigma}^{(i)}$ exposes the WOTS⁺ secret values $k, \dots, m_i - 1$ of chain i . If no k leads to a match, the WOTS⁺ signature is discarded.

Complexity: Identifying secret WOTS⁺ values only yields a negligible computational cost, ranging from $2^{8.04}$ to $2^{12.98}$ hashes, depending on the parameter set.

6.4 Tree Grafting

After identifying the WOTS⁺ secret values, the attacker performs the grafting tree attack. In this step, the attacker tries to find an XMSS public key that can be signed using the targeted WOTS⁺ instance. This involves an exhaustive search for a secret **SK.seed** that produces an XMSS public key pk^X . However, only a limited number of XMSS trees can be signed, as the attacker can only sign messages that are compatible with the exposed chain secrets. Let $(\hat{\theta}_0, \hat{\theta}_1, \dots, \hat{\theta}_{\ell-1})$ be the exposed secret elements of the targeted WOTS⁺ instance, corresponding to message chunks $(\hat{m}_0, \hat{m}_1, \dots, \hat{m}_{\ell-1})$. Following the approach by Castelnovi et al. [9], the algorithm to find a suitable XMSS tree works as follows:

1. Create a new XMSS tree with public key $\tilde{\text{pk}}^X$ from a secret **SK.seed** drawn uniformly at random.
2. Split $\tilde{\text{pk}}^X$ and its WOTS⁺ checksum into chunks $(\tilde{r}_0, \dots, \tilde{r}_{\ell-1})$ of size w bits.
3. If $\tilde{r}_i \geq \hat{m}_i$ for all $0 \leq i < \ell$, return the grafted **SK.seed**. Repeat from step 1 otherwise.

Complexity: The hash complexity of tree grafting may impose significant computational cost, scaling with the number of signatures for a compromised WOTS⁺ instance and the number of compromised WOTS⁺ secret values. In our evaluation (Section 8), the observed grafting complexity ranges from $2^{27.98}$ hashes (SHAKE-128f deterministic) to $2^{130.46}$ hashes (SHAKE-128s deterministic) for parameter sets where faults were collected, with successful forgeries achieved for complexities up to $2^{32.95}$ (SHA2-128s randomized). The total attack complexity, combining grafting and seeking, determines the practical feasibility of the forgery. In Section 7, we describe how to determine the actual complexity of attacking specific faulted instances.

6.5 Path Seeking

The randomization value R , together with the message m to be signed, determines the path taken through the hypertree. This randomization value is determined by the signing party, and can therefore be chosen by the attacker. For the universal forgery, they have to find a value R' such that the signing visits the compromised WOTS⁺ instance and the grafted XMSS tree. Finding such a value R' boils down to searching an n byte number such that the $h - h'l^*$ most significant bits are equal to the index of the grafted subtree.

Complexity: Path seeking yields cost for each forged signature. Each attempt at finding a valid R' requires a single hash function call, and the attacker on average needs $2^{h - h'l^*}$ hash function calls to find a suitable R' .

7 An Exact Solution For Grafting Complexity

Despite previous analysis on the theoretical grafting complexity for a given SLH-DSA instance, the literature currently lacks a method for estimating the grafting complexity for a concrete compromised WOTS⁺ key pair after the collection of faulty signatures. While Genêt provides a sound estimate for the average grafting complexity for uniformly distributed WOTS⁺ secrets [14], in practice, an attacker has to choose from a concrete set of WOTS⁺ instances to graft a tree for. However, for concrete WOTS⁺ instances, there can be a wide variance in grafting complexity even for the same number of observed collisions. Our analysis closes this gap, allowing an attacker not only to estimate the runtime for the offline grafting phase, but also to rank compromised WOTS⁺ key pairs based on the expected forgery runtime. We introduce a novel

combinatorial technique to precisely determine the grafting success probability for a given set of exposed WOTS⁺ secret values, enabling attackers an improved estimation of the total attack complexity. This approach extends Genêt’s work by providing exact complexity analysis for concrete instances rather than average-case estimates.

7.1 Notation

Let $w \in \mathbb{Z}_+$ be the Winternitz parameter. For a WOTS⁺ signature σ of a message m , we denote the message chain indices in σ corresponding to message block $i \in \{0, 1, \dots, \ell_1 - 1\}$ by $m_i \in \{0, \dots, w - 1\}$. Similarly, denote the checksum chain indices corresponding to the checksum c of m by $c_i \in \{0, \dots, w - 1\}$ for $i \in \{0, \dots, \ell_2 - 1\}$. We denote a full WOTS⁺ chain by $\theta = (m_0, m_1, \dots, m_{\ell_1-1}, c_0, c_1, \dots, c_{\ell_2-1})$.

For chain i , we call the preimages $0 \leq \tilde{\theta}_i < \theta_i$ of θ_i the *WOTS⁺ secret values*. If the attacker obtains two or more WOTS⁺ signatures for distinct messages such that at least one secret value is exposed, we call this instance *compromised*.

7.2 A Combinatorial Solution For Grafting Complexity

There is a wide variance in the resulting grafting complexity even for the same number of collisions in a WOTS⁺ instance. Therefore, we now derive the expected runtime for forging a signature using a specific key pair. To exploit a compromised WOTS⁺ instance, the attacker combines the secret values exposed by signatures from that instance, collecting the minimal chain secret values for each chain. They then use the exposed chain secrets to graft and sign an XMSS tree at the next layer. However, they have to find a tree matching the exposed chain secrets. Since the sampled XMSS trees are distributed uniformly at random, the attacker can now compute the probability that a random message can be signed given the exposed chain secrets, a proxy for the post-processing *grafting complexity*. As a WOTS⁺ signature is only valid if message and checksum match, they have to find the number of valid message-checksum combinations. Enumerating all messages and filtering for the ones reachable with the available checksums is infeasible, because there are w^{ℓ_1} possible messages, ranging from 2^{128} to 2^{256} depending on the parameter set. However, since $\ell_2 = 3$ for all parameter sets, they can enumerate the reachable checksums and compute the number of valid messages for those checksums. We first formulate this challenge as a combinatorial problem and then introduce an implementation based on dynamic programming. For that, we denote the integer representation of a checksum c regarding a base w by $\text{to_int}(c, w) = \sum_{i=0}^{\ell_2} c_i \cdot w^i$.

Chain capacities: We define the per-component capacities of message chain $i = 0, \dots, \ell_1 - 1$ as $k_i = w - 1 - m_i$. These capacities describe the number of times each message chain can be iterated until the chain’s end at step $w - 1$ is reached.

Enumeration of checksums: Let \mathcal{T} be the set of checksums computable by continuing checksum chains until the chain’s end or the maximum checksum $\ell_1 \cdot (w - 1)$ is reached. More formally, $\mathcal{T} \subseteq \{0, 1, \dots, w - 1\}^{\ell_2}$ consist of the tuples $\tau = (\tau_0, \dots, \tau_{\ell_2-1})$ satisfying for all $0 \leq i \leq \ell_2 - 1$:

$$c_i \leq \tau_i < w \wedge \text{to_int}(\tau, w) \leq \ell_1 \cdot (w - 1).$$

Counting signable compositions: For each valid tuple $\tau \in \mathcal{T}$, define the capacity to be distributed across chains as $\kappa(\tau) = \text{to_int}(\tau, w) - \text{to_int}(\tau, w)$. We now want to determine the number of ways the capacity can be distributed under the constraints $\mathcal{K} = \{k_i \mid 0 \leq i < \ell_1\}$, i.e., the number of compositions $x = (x_0, x_1, \dots, x_{\ell_1})$ satisfying $\sum_{i=0}^{\ell_1} x_i = \kappa(\tau, \mathcal{K})$ such that $0 \leq x_i \leq k_i$. Generating solutions to this system is known as the *second-order restricted weak integer composition* [34] problem and, in general, there are exponentially many solutions to this system. However, we are only interested in the number of compositions $N(\kappa(\tau, \mathcal{K}))$ satisfying this system, which can be computed in time $O(\tau \cdot \ell_1)$ and memory $O(\tau)$ by the dynamic programming based algorithm shown in Algorithm 2.

Putting it all together: Using $N(\kappa(\tau, \mathcal{K}))$ for all $\tau \in \mathcal{T}$, we can compute the probability that a random XMSS tree can be signed with a given compromised WOTS⁺ instance by taking the fraction of signable messages over the number of all messages: $\mathbb{P}(\text{Signable}) = \frac{\sum_{\tau \in \mathcal{T}} N(\kappa(\tau, \mathcal{K}))}{w^{\ell_1}}$. Modeling tree grafting as a Bernoulli process with success probability $\mathbb{P}(\text{Signable})$, we can, on average, expect the tree grafting to succeed after $\mathbb{C}(\text{Grafting}) = \frac{1}{\mathbb{P}(\text{Signable})}$ attempts. Multiplying $\mathbb{C}(\text{Grafting})$ with the number of hash operations required to compute an XMSS tree yields the grafting hash complexity.

8 Evaluation of SLasH-DSA

In this section, we evaluate the practical feasibility and effectiveness of our Rowhammer-based attack against SLH-DSA signature schemes. We leverage the SWAGE framework introduced in Section 4 to orchestrate the end-to-end attack: based on our fault analysis in Section 6.1, we implemented an SLH-DSA-specific ORCHESTRATOR module for SWAGE, and use SWAGE’s core components to perform the attack: SWAGE handles DRAM address reverse-engineering, hammering pattern execution, page injection to co-locate victim data with aggressor rows, and signature collection during the online phase. Taking advantage of SWAGE’s modular architecture, this allows us to focus on the target-specific parts of the attack, relying on established attack primitives implemented in the SWAGE core. We demonstrate the ability to successfully compromise the deterministic and randomized variants of SLH-DSA across different parameter sets. Our evaluation quantifies the computational complexity of the grafting and seeking phases, and analyzes the trade-offs between different attack strategies. The results show that our attack can achieve

practical signature forgery within reasonable time bounds for some SLH-DSA parameter configurations, taking only a few minutes to an hour on commodity hardware, while showing theoretical attack complexities for most parameter sets.

8.1 Experimental Setup

We conducted our experiments on a system with an Intel i5-6400 CPU and a single G.Skill AEGIS 16 GB DDR4-2133 DIMM, running Ubuntu 20.04.6 with Linux kernel version 6.8.12-generic with default settings. We evaluated the SLH-DSA implementation in OpenSSL version 3.5.1, released on July 1st, 2025. The signing process was implemented via a small signing server accepting signing requests via standard input and returning signatures to standard output. The signing server is configurable with all standardized SLH-DSA parameter sets with both random and deterministic signing. Both the signing server and OpenSSL were built with gcc version 11.4.0 with all `-O3` optimizations enabled. The post-processing is performed on an Intel Xeon Gold 5415+ with 128 GB RAM, running Ubuntu 20.04.6 with Linux kernel version 5.11.0-generic with default settings.

While the experimental setup we chose demonstrates the feasibility of the attack, newer hardware incorporates novel Rowhammer defenses, and newer software versions may include additional countermeasures. Future work should investigate the attack’s effectiveness on more recent and varying system configurations.

8.2 Attack Against Deterministic SLH-DSA

At the start of the signing procedure, the SLH-DSA signing algorithm determines the FORS instances for signing the message and the path taken through the hypertree. The FORS instances and path through the hypertree depend on the message to be signed, the **PK.seed** and **PK.root**, and an optional random value R . The randomization of the signing process is optional, and is set to **PK.seed** when deterministic mode is selected. We start the experiments with an attack against the deterministic variant of SLH-DSA. Table 2a shows the attack complexity after one hour of hammering. The table shows the number of signatures collected and the number of observed faults (i.e., signatures not matching the expected output from the signing server). For the most exposed WOTS⁺ instance, it reports the layer this instance was encountered, the grafting and seeking complexities in number of hashes and, if applicable, the post-processing time for a successful forgery. Our results highlight the increased susceptibility of deterministic SLH-DSA, as deterministic signing maximizes the number of possible collisions due to the reuse of FORS, XMSS, and WOTS⁺ instances throughout the signing procedure.

8.3 Attack Against Randomized SLH-DSA

In contrast to the deterministic variants, the randomized SLH-DSA signature scheme incorporates an additional random value R , which further diversifies the signing process even for signing the same message repeatedly. This randomization significantly complicates the attack methodology, as it disrupts the reuse of specific cryptographic instances that we exploit in the deterministic mode. In this subsection, we outline our approach to adapting collision-based strategies to randomized signing. We analyze the impact of the extra randomness on both the grafting and path seeking phases, and discuss the increased computational challenges introduced by randomization. Our experiments provide insights into the effectiveness and limitations of the attack under these conditions.

For deterministic signing, a fixed message always results in the same path through the hypertree, whereas for randomized signing, the path is determined using R . Due to the randomized path, it is unlikely that XMSS trees at low hypertree levels will be reused, and we cannot expect to find colliding WOTS⁺ instances for our attack on these levels. This is especially true for the “small” parameter sets of the scheme: For example, in the 256s parameter set with XMSS tree height $h' = 8$, there are 256 WOTS⁺ instances in the topmost layer 16, and 256^{15-l} at layer $0 < l < 16$. The “fast” variants, however, stand out with small XMSS trees and more hypertree layers. For example, the 192f parameter set has XMSS tree height $h' = 3$, resulting in 8 WOTS⁺ instances in the topmost layer 21, and 8^{20-l} at layer $0 < l < 21$. This makes WOTS⁺ collisions on layers other than the topmost layer feasible, increasing the attack surface of our attack against randomized signing. Table 2b shows the attack cost for a one-shot forgery attack, i.e., forging a valid signature for one attacker-chosen message, after eight hours of hammering the randomized variants of SLH-DSA for all parameter sets. While the number of faulted signatures serve as a first proxy for the success of an attack, the fault characteristic is important for the attack: for example, faulting the hash function’s internal constants will not produce useful faults for the attacker. We presume that the increased fault rate for some parameter sets, such as SHAKE-192s, where we observe a significant number of faulted signatures not corresponding to improved attack complexities, are due to this effect. We hypothesize however, that longer hammering time in general improves attack capabilities.

The experiments against both deterministic and randomized signing were conducted while applying performance degradation as described in Section 6.2. The experiments show that the end-to-end attack is feasible after only a few hours of hammering for some parameter sets, resulting in a post-processing time for both deterministic and randomized signing of only a few minutes. The grafting step complexity is calculated by applying our novel, exact solution method from Section 7. Weighing grafting vs. path seeking is a promising optimization for multi-shot signing, as the cost for path seek-

Table 2: Best total attack complexity for one-shot forgery after (a) one hour of hammering against deterministic signing and (b) eight hours of hammering against randomized signing with all parameter sets. We consider the attack complexity to be the sum of grafting and seeking for a one-shot forgery attack, i.e., signing one attacker-selected message. Grafting and Seeking columns show hash complexity (cf. Section 7), time is in seconds. Empty cells (–) indicate parameter sets where insufficient faulted signatures were collected to carry out an attack within a reasonable amount of time.

		(a)						(b)							
		Signatures		Faulted	Layer	Grafting	Seeking	Signatures		Faulted	Layer	Grafting	Seeking	Time (s)	
SHA2	128s	187	1	3	$2^{93.37}$	2^{36}	–	SHA2	128s	1503	37	6	$2^{32.95}$	2^9	3781
	128f	1917	1	6	$2^{30.66}$	2^{48}	–		128f	7659	1	–	–	–	–
	192s	105	2	5	$2^{52.32}$	2^{18}	–		192s	843	50	6	$2^{41.91}$	2^9	–
	192f	567	7	12	$2^{42.09}$	2^{30}	–		192f	4617	6	19	$2^{40.84}$	2^9	–
	256s	115	3	6	$2^{53.85}$	2^{16}	–		256s	949	102	7	$2^{51.59}$	2^8	–
	256f	275	3	4	$2^{50.18}$	2^{52}	–		256f	2279	26	15	$2^{52.65}$	2^8	–
SHAKE	128s	163	2	4	$2^{130.46}$	2^{27}	–		128s	1323	3	–	–	–	–
	128f	1683	3	17	$2^{27.98}$	2^{15}	151		128f	13479	33	18	$2^{85.85}$	2^{12}	–
	192s	95	1	–	–	–	–		192s	761	360	6	$2^{36.42}$	2^9	–
	192f	517	10	21	$2^{36.56}$	2^3	–		192f	4137	209	21	$2^{31.29}$	2^3	1417
	256s	105	5	3	$2^{53.20}$	2^{40}	–		256s	869	26	7	$2^{52.77}$	2^8	–
	256f	253	10	15	$2^{44.53}$	2^8	–		256f	2035	34	15	$2^{52.31}$	2^8	–

ing must be spent on each message, while grafting is only needed once.

For unsuccessful forgeries in Tables 2a and 2b, we identify two bottlenecks: for the SHAKE-192s (deterministic), SHA2-128f, and SHAKE-128s (randomized) parameter sets, insufficient faulted signatures were collected during the hammering campaigns, resulting in no compromised WOTS⁺ instances. Besides, for parameter sets with compromised WOTS⁺ instances but high grafting complexities, the offline post-processing phase is computationally infeasible on commodity hardware (e.g., $2^{85.85}$ hashes for SHAKE-128f randomized). However, using a more vulnerable DIMM with higher bit flip rates could improve fault collection during the online phase, potentially making attacks feasible for additional parameter sets. This is in contrast to Genêt’s hardware-based approach [14], where glitch injection enables attacks in all five trials of his practical attack against the SHAKE-256s parameter set with 1024 collected signatures per trial, the attack being restricted solely by the offline post-processing.

9 Related Work

In this section, we discuss prior work relevant to our study, including grafting tree attacks and fault analysis techniques, Rowhammer-based attacks that target PQC schemes, and existing countermeasures against Rowhammer vulnerabilities.

9.1 Grafting Attacks and Countermeasures

Castelnovi et al. [9] were the first to present a theoretical analysis of the *grafting tree attack* including computational feasibility. Genêt et al. implemented the grafting tree attack against SPHINCS [15]. Genêt later presented a theoretical analysis of the grafting tree attack against SPHINCS⁺ along with a practical attack implementation [14]. The fault injection method in those works is a hardware-supported clock glitching attack focused on embedded devices.

Additionally, Genêt [14] examined the *layer caching* and *branch caching* countermeasures against grafting tree attacks. The layer caching countermeasure involves pre-computing upper layers of the SPHINCS⁺ hypertree, storing them alongside the public key. This reduces the number of attackable hash function calls, forcing the attacker to find collisions in the lower layers, decreasing collision probability and increasing path seeking complexity for successful faults. This countermeasure shows varying degrees of effectiveness, depending on the number of cached layers. With the *branch caching* countermeasure, on the other hand, the signing party stores all encountered WOTS⁺ signatures and public keys in a fixed-size cache. However, it was shown that this countermeasure is ineffective in protecting SLH-DSA, as cache misses and therefore recomputations of WOTS⁺ instances are inevitable with reasonably sized caches. To summarize, the best current solution to protect the scheme is *redundancy*, where (parts of)

the signature is verified by duplicate computation [5, 9].

9.2 Rowhammer Attacks against Post-Quantum Schemes

Rowhammer has also shown effective as a *fault injection method against other post-quantum schemes*. In PQHammer [4], Amer et al. describe end-to-end key recovery attacks based on Rowhammer against the post-quantum key encapsulation mechanisms (KEMs) BIKE and CRYSTALS-Kyber (ML-KEM), and against the lattice-based signature scheme CRYSTALS-Dilithium (ML-DSA). In Crowhammer [18], Haidar et al. show how to attack the post-quantum signature scheme Falcon to perform a Rowhammer-based key recovery attack. This line of research highlights the need to include Rowhammer attacks against post-quantum cryptography in the threat models. A starting point for this are tools like HammerTime [39], enabling profiling and simulation of Rowhammer attacks. While the profiling results indicate how many hammering experiments are required to flip a bit in a given offset range within a page, HammerTime does not cover all necessary steps for an end-to-end Rowhammer attack.

9.3 Fault Analysis

Orthogonal to exploiting memory fault vulnerabilities, another line of work *analyzes the vulnerability of schemes to fault attacks*. Achilles [29] systematically analyzes signature schemes by separating the fault model from the algorithm under analysis: A generalized signature scheme utilizes public and secret parameters combined with signing oracles that generate valid and faulty signatures. A fault model then specifies where to inject the faults, and in a post-processing, the secret key is recovered. The authors showcase the efficacy of Achilles in the key recovery domain with six different signature schemes including the post-quantum scheme CRYSTALS-Dilithium (ML-DSA). However, transitioning from key recovery to universal forgery of SLH-DSA is not straightforward as there is no obvious reduction from their attacker model for deriving secret values to forging valid signatures. Rainbow [31] is another automated fault analysis tool that emulates the target program with the Unicorn CPU emulator [40]. Rainbow allows to specify fault injection parameters for automated fault injection against a target program.

9.4 Rowhammer Countermeasures

There has been a variety of *Rowhammer countermeasures* implemented by manufacturers and proposed in scientific work. TRR, a DDR4 vendor-specific mechanism, tracks memory accesses and issues refresh commands on unusual access patterns, preventing simple Rowhammer attacks by preemptively restoring charge before bit flips can be induced. However, TRR is ineffective against advanced hammering pat-

terns [13, 21]. Hardware-supported *Error-Correcting Codes* (ECC) can detect and correct bit flips [16, 38], but increases DRAM cost and may be ineffective against certain Rowhammer attacks [10, 24] or introduce new attack vectors such as timing-based attack [28]. Even DDR5, while introducing new countermeasures [8, 23], has been found to be susceptible to Rowhammer attacks [20, 22].

OS-level mitigations like *guard rows* around sensitive memory pages [27, 41] prevent some Rowhammer attacks but introduce high memory and performance costs. They can furthermore be bypassed by attacks using non-adjacent aggressor rows [21, 26]. However, all of these countermeasures have individual limitations, are not universally applicable, or are not effective against all Rowhammer attacks.

10 Future Work

In this paper, we demonstrated the feasibility of a Rowhammer attack against SLH-DSA as an exemplary post-quantum signature scheme. We implement the SWAGE framework, enabling easy implementation of practical Rowhammer attacks. One remaining aspect requiring substantial manual effort with SWAGE is analyzing the binary and scheme to find memory locations susceptible to Rowhammer fault attacks. This could be improved by integrating SWAGE more closely with tools like Achilles or Rainbow (c.f. Section 9).

As shown in Table 2a and Table 2b, for some parameter sets, tree grafting and path seeking require significant CPU resources assuming a limited amount of hammering. Since both tree grafting and path seeking are heavily dependent on the hashing performance, a substantial speedup is to be expected with GPU acceleration. In 2025, Wang et al. presented CUSPX [43], a GPU-accelerated implementation of SPHINCS+. Preliminary benchmarks using their hash function implementations show a substantial increase in hash performance, achieving $2^{37.2}$ SHA-256 operations per second on an NVIDIA RTX 4090. This marks a 2^{14} -fold increase compared to our evaluation system. A GPU-accelerated implementation of the grafting and path seeking algorithms should speed up the attack significantly, potentially allowing to break some of the remaining parameters sets.

Moreover, for the attack presented in this paper, we assume that the attacker is core co-located with the victim for the page injection. In Rubicon [7], Bölcsei et al. demonstrate that page injection attacks also work across CPUs, widening the threat model to be considered for Rowhammer attacks and potentially making a wider range of systems vulnerable.

11 Conclusion

We presented SLasH-DSA, an end-to-end Rowhammer attack against SLH-DSA. We demonstrated this attack against OpenSSL to be feasible and effective in practice, allowing an

attacker to conduct a universal forgery attack against all NIST security levels of SLH-DSA by flipping bits in the signature generation process. In addition, we provided a novel complexity analysis for concrete faulted signatures that allows steering the post processing along the path of lowest complexity, speeding up the attack. This is, to the best of our knowledge, the first practical demonstration of a Rowhammer attack against SLH-DSA on a real-world system. To conduct the attack, we introduced SWAGE, a comprehensive open-source framework for performing Rowhammer attacks on real-world systems. SWAGE provides a modular interface for Rowhammer attacks, allowing users to focus on the attack logic rather than the underlying hardware and operating system details. To our knowledge, SWAGE is the first fully open-source end-to-end framework for Rowhammer attacks, providing a complete and extensible solution for attacks on real-world systems.

12 Acknowledgements

This work has been supported by BMFTR through the AnoMed project and by the BMBF project SASVI.

References

- [1] Andrew J. Adiletta, M. Caner Tol, Kemal Derya, Berk Sunar, and Saad Islam. LeapFrog: The Rowhammer Instruction Skip Attack. In *2025 IEEE 10th European Symposium on Security and Privacy (EuroS&P)*, pages 1067–1081. IEEE, 2025.
- [2] Andrew J. Adiletta, M. Caner Tol, Yarkin Doröz, and Berk Sunar. Mayhem: Targeted Corruption of Register and Stack Variables. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*, pages 467–482. Association for Computing Machinery, 2024.
- [3] Alejandro Cabrera Aldaya and Billy Bob Brumley. HyperDegrade: From GHz to MHz Effective CPU Frequencies. In Kevin R. B. Butler and Kurt Thomas, editors, *31st USENIX Security Symposium (USENIX Security 22)*, pages 2801–2818. USENIX Association, 2022.
- [4] Samy Amer, Yingchen Wang, Hunter Kippen, Thinh Dang, Daniel Genkin, Andrew Kwong, Alexander Nelson, and Arkady Yerukhimovich. PQ-Hammer: End-to-End Key Recovery Attacks on Post-Quantum Cryptography Using Rowhammer. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 3308–3323. IEEE, 2025.
- [5] Dorian Amiet, Lukas Leuenberger, Andreas Curiger, and Paul Zbinden. FPGA-based SPHINCS+ Implementations: Mind the Glitch. In *23rd Euromicro Conference on Digital System Design*, pages 229–237. IEEE, 2020.
- [6] Jean-Philippe Aumasson, Daniel J. Bernstein, Ward Beullens, Christoph Dobraunig, Maria Eichlseder, Scott Fluhrer, Stefan-Lukas Gazdag, Andreas Hülsing, Panos Kampanakis, Stefan Kölbl, Tanja Lange, Martin M. Lauridsen, Florian Mendel, Ruben Niederhagen, Christian Rechberger, Joost Rijneveld, and Peter Schwabe. SPHINCS+ Submission to the NIST Post-Quantum Project, v.3.1, 2022.
- [7] Matej Bölcseki, Patrick Jattke, Johannes Wikner, and Kaveh Razavi. Rubicon: Precise Microarchitectural Attacks with Page-Granular Massaging. In *2025 IEEE 10th European Symposium on Security and Privacy (EuroS&P)*, pages 395–410. IEEE, 2025.
- [8] Oguzhan Canpolat, A. Giray Yaglikçi, Geraldo F. Oliveira, Ataberk Olgun, Oguz Ergin, and Onur Mutlu. Understanding the Security Benefits and Overheads of Emerging Industry Solutions to DRAM Read Disturbance. In *Fourth Workshop on DRAM Security*. IEEE, 2024.
- [9] Laurent Castelnovi, Ange Martinelli, and Thomas Prest. Grafting Trees: A Fault Attack Against the SPHINCS Framework. In Tanja Lange and Rainer Steinwandt, editors, *Post-Quantum Cryptography*, pages 165–184. Springer International Publishing, 2018.
- [10] Lucian Cojocar, Kaveh Razavi, Cristiano Giuffrida, and Herbert Bos. Exploiting Correcting Codes: On the Effectiveness of ECC Memory Against Rowhammer Attacks. In *2019 IEEE Symposium on Security and Privacy (SP)*, pages 55–71. IEEE, 2019.
- [11] David A Cooper, Daniel C Apon, Quynh H Dang, Michael S Davidson, Morris J Dworkin, and Carl A Miller. Recommendation for Stateful Hash-Based Signature Schemes. *NIST Special Publication*, (208):800–208, 2020.
- [12] Andrea Di Dio, Mathé Hertogh, and Cristiano Giuffrida. Half Spectre, Full Exploit: Hardening Rowhammer Attacks with Half-Spectre Gadgets. In Marina Blanton, William Enck, and Cristina Nita-Rotaru, editors, *2025 IEEE Symposium on Security and Privacy (SP)*, pages 3583–3598. IEEE, 2025.
- [13] Pietro Frigo, Emanuele Vannacc, Hasan Hassan, Victor Van Der Veen, Onur Mutlu, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. TRRespass: Exploiting the Many Sides of Target Row Refresh. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 747–762. IEEE, 2020.
- [14] Aymeric Genêt. On Protecting SPHINCS+ Against Fault Attacks. In *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pages 80–114. IACR, 2023.

[15] Aymeric Genêt, Matthias J. Kannwischer, Hervé Pelletier, and Andrew McLauchlan. Practical Fault Injection Attacks on SPHINCS, 2018.

[16] Daniel Gruss, Moritz Lipp, Michael Schwarz, Daniel Genkin, Jonas Juffinger, Sioli O’Connell, Wolfgang Schoechl, and Yuval Yarom. Another Flip in the Wall of Rowhammer Defenses. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 245–261. IEEE, 2018.

[17] Daniel Gruss, Clémentine Maurice, and Stefan Mangard. Rowhammer.js: A Remote Software-Induced Fault Attack in Javascript. In Juan Caballero, Urko Zurutuza, and Ricardo J. Rodríguez, editors, *Detection of Intrusions and Malware, and Vulnerability Assessment*, Lecture Notes in Computer Science, pages 300–321. Springer International Publishing, 2016.

[18] Calvin Abou Haidar, Quentin Payet, and Mehdi Tibouchi. Crowhammer: Full Key Recovery Attack on Falcon with a Single Rowhammer Bit Flip. In Yael Tuman Kalai and Seny F. Kamara, editors, *Advances in Cryptology — CRYPTO 2025*, pages 103–135. Springer Nature Switzerland, 2025.

[19] Saad Islam, Ahmad Moghimi, Ida Bruhns, Moritz Krebbel, Berk Gülmезoglu, Thomas Eisenbarth, and Berk Sunar. SPOILER: Speculative Load Hazards Boost Rowhammer and Cache Attacks. In *28th USENIX Security Symposium (USENIX Security 19)*, pages 621–637. USENIX Association, 2019.

[20] Patrick Jattke, Michele Marazzi, Flavien Solt, Max Wipfli, Stefan Gloor, and Kaveh Razavi. McSee: Evaluating Advanced Rowhammer Attacks and Defenses via Automated DRAM Traffic Analysis. In *34th USENIX Security Symposium (USENIX Security 25)*, pages 5621–5640. USENIX Association, 2025.

[21] Patrick Jattke, Victor Van Der Veen, Pietro Frigo, Stijn Gunter, and Kaveh Razavi. Blacksmith: Scalable Rowhammering in the Frequency Domain. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 716–734. IEEE, 2022.

[22] Patrick Jattke, Max Wipfli, Flavien Solt, Michele Marazzi, Matej Bölcseki, and Kaveh Razavi. ZenHammer: Rowhammer Attacks on AMD Zen-based Platforms. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, 2024.

[23] JEDEC. JESD79-5C: DDR5 SDRAM, 2024.

[24] Nureddin Kamadan, Walter Wang, Stephan van Schaik, Christina Garman, Daniel Genkin, and Yuval Yarom. ECC.fail: Mounting Rowhammer Attacks on DDR4 Servers with ECC Memory. In *34th USENIX Security Symposium (USENIX Security 25)*. USENIX Association, 2025.

[25] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors. In *2014 ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)*, pages 361–372. IEEE, 2014.

[26] Andreas Kogler, Jonas Juffinger, Salman Qazi, Yoongu Kim, Moritz Lipp, Nicolas Boichat, Eric Shiu, Matthias Nissler, and Daniel Gruss. Half-Double: Hammering From the Next Row Over. In *31st USENIX Security Symposium (USENIX Security 22)*, pages 3807–3824. USENIX Association, 2022.

[27] Radhesh Krishnan Konoth, Marco Oliverio, Andrei Tatar, Dennis Andriesse, Herbert Bos, Cristiano Giuffrida, and Kaveh Razavi. ZebRAM: Comprehensive and Compatible Software Protection Against Rowhammer Attacks. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 697–710. USENIX Association, 2018.

[28] Andrew Kwong, Daniel Genkin, Daniel Gruss, and Yuval Yarom. RAMBleed: Reading Bits in Memory Without Accessing Them. In *2020 IEEE Symposium on Security and Privacy (SP)*, pages 695–711. IEEE, 2020.

[29] Junkai Liang, Zhi Zhang, Xin Zhang, Qingni Shen, Yansong Gao, Xingliang Yuan, Haiyang Xue, Pengfei Wu, and Zhonghai Wu. Achilles: A Formal Framework of Leaking Secrets from Signature Schemes via Rowhammer. In *33rd USENIX Security Symposium (USENIX Security 25)*. USENIX Association, 2025.

[30] Moritz Lipp, Michael Schwarz, Lukas Raab, Lukas Lamster, Misiker Tadesse Aga, Clémentine Maurice, and Daniel Gruss. Nethammer: Inducing Rowhammer Faults through Network Requests. In *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS&PW)*, pages 710–719. IEEE, 2020.

[31] Alexandre Looss. Integrating Fault Injection In Development Workflows, 2022.

[32] NIST. Stateless Hash-Based Digital Signature Standard, 2024.

[33] OpenSSL. OpenSSL 3.5 Final Release - Live | OpenSSL Library, 2025.

[34] Daniel R. Page. Generalized Algorithm for Restricted Weak Composition Generation. *Journal of Mathematical Modelling and Algorithms in Operations Research*, (4):345–372, 2013.

[35] Peter Pessl, Daniel Gruss, Clémentine Maurice, Michael Schwarz, and Stefan Mangard. DRAMA: Exploiting DRAM Addressing for Cross-CPU Attacks. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 565–581. USENIX Association, 2016.

[36] Adnan Siraj Rakin, Md Hafizul Islam Chowdhury, Fan Yao, and Deliang Fan. DeepSteal: Advanced Model Extractions Leveraging Efficient Weight Stealing in Memories. In *2022 IEEE Symposium on Security and Privacy (SP)*, pages 1157–1174. IEEE, 2022.

[37] Kaveh Razavi, Ben Gras, Erik Bosman, Bart Preneel, Cristiano Giuffrida, and Herbert Bos. Flip Feng Shui: Hammering a Needle in the Software Stack. In Thorsten Holz and Stefan Savage, editors, *25th USENIX Security Symposium (USENIX Security 16)*, pages 1–18. USENIX Association, 2016.

[38] Mark Seaborn and Thomas Dullien. Project Zero: Exploiting the DRAM Rowhammer Bug to Gain Kernel Privileges, 2015.

[39] Andrei Tatar, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi. Defeating Software Mitigations Against Rowhammer: A Surgical Precision Hammer. In Michael Bailey, Thorsten Holz, Manolis Stamatogiannakis, and Sotiris Ioannidis, editors, *Research in Attacks, Intrusions, and Defenses*, pages 47–66. Springer International Publishing, 2018.

[40] Unicorn Team. Unicorn – The Ultimate CPU Emulator, 2015.

[41] Victor van der Veen, Martina Lindorfer, Yanick Fratantonio, Harikrishnan Padmanabha Pillai, Giovanni Vigna, Christopher Kruegel, Herbert Bos, and Kaveh Razavi. GuardION: Practical Mitigation of DMA-Based Rowhammer Attacks on ARM. In Cristiano Giuffrida, Sébastien Bardin, and Gregory Blanc, editors, *Detection of Intrusions and Malware, and Vulnerability Assessment - 15th International Conference, DIMVA 2018, Saclay, France, June 28-29, 2018, Proceedings*, Lecture Notes in Computer Science, pages 92–113. Springer, 2018.

[42] Yingchen Wang, Riccardo Paccagnella, Elizabeth Tang He, Hovav Shacham, Christopher W. Fletcher, and David Kohlbrenner. Hertzbleed: Turning Power Side-Channel Attacks into Remote Timing Attacks on x86. In Kevin R. B. Butler and Kurt Thomas, editors, *31st USENIX Security Symposium (USENIX Security 22)*, pages 679–697. USENIX Association, 2022.

[43] Ziheng Wang, Xiaoshe Dong, Heng Chen, Yan Kang, and Qiang Wang. CUSPX: Efficient GPU Implementations of Post-Quantum Signature SPHINCS+. *IEEE Transaction on Computers*, (1):15–28, 2025.

[44] Yuan Xiao, Xiaokuan Zhang, Yinqian Zhang, and Radu Teodorescu. One Bit Flips, One Cloud Flops: Cross-VM Row Hammer Attacks and Privilege Escalation. In *25th USENIX Security Symposium (USENIX Security 16)*, pages 19–35. USENIX Association, 2016.

[45] Yongye Zhu, Boru Chen, Zirui Neil Zhao, and Christopher W. Fletcher. Controlled Preemption: Amplifying Side-Channel Attacks from Userspace. In Lieven Eeckhout, Georgios Smaragdakis, Katai Liang, Adrian Sampson, Martha A. Kim, and Christopher J. Rossbach, editors, *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 162–177. Association for Computing Machinery, 2025.

13 Appendix

13.1 Number of Weak Compositions

The following DP-based algorithm computes the number of weak compositions $\sum_{i=0}^{\ell_1-1} x_i$ for a target sum τ respecting upper bounds k_i for $0 \leq i \leq \ell_1-1$:

Algorithm 2: WEAK-COMP($\tau, k_0, \dots, k_{\ell_1-1}$)

Input: Target sum τ , upper bounds k_0, \dots, k_{ℓ_1-1}
Output: Number of weak compositions
 $x = (x_0, \dots, x_{\ell_1-1})$ of τ respecting $0 \leq x_i \leq k_i$

```

1 Initialize
2   for  $t \leftarrow 0$  to  $\tau$  do
3      $dp[t] \leftarrow 0$ 
4    $dp[0] \leftarrow 1$            // one way to compose 0
5 for  $i \leftarrow 0$  to  $\ell_1 - 1$  do
6    $prefix \leftarrow 0$ 
7   for  $t \leftarrow 0$  to  $\tau$  do
8      $prefix \leftarrow prefix + dp[t]$ 
9     if  $t > k_i$  then  $prefix \leftarrow prefix - dp[t - (k_i + 1)]$ 
10     $new[t] \leftarrow prefix$ 
11    for  $t \leftarrow 0$  to  $\tau$  do
12       $dp[t] \leftarrow new[t]$ 
13 return  $dp[\tau]$ 

```
