

Efficient CPA Attack on Hardware Implementation of ML-DSA in Post-Quantum Root of Trust

Merve Karabulut

College of Engineering & Computer Science
Florida Atlantic University
mkarabulut2024@fau.edu

Reza Azarderakhsh

College of Engineering & Computer Science
Florida Atlantic University
razarderakhsh@fau.edu

Abstract—Side-channel attacks (SCA) pose a significant threat to cryptographic implementations, including those designed to withstand the computational power of quantum computers. This paper introduces the first side-channel attack on an industry-grade post-quantum cryptography implementation. Specifically, we present a Correlation Power Analysis (CPA) attack targeting the open-source hardware implementation of ML-DSA within a Silicon Root of Trust framework developed through a multi-party collaboration involving leading technology companies.

Our attack focuses on the modular reduction process that follows the Number Theoretic Transform-based polynomial point-wise multiplication. By exploiting side-channel leakage from a distinctive unique reduction algorithm and leveraging the zeroization mechanism used to securely erase sensitive information by clearing internal registers, we significantly enhance the attack’s efficacy. Our findings reveal that an adversary can extract the secret keys using only 10,000 power traces. With access to these keys, an attacker could forge signatures for certificate generation, thereby compromising the integrity of the root of trust. This work highlights the vulnerabilities of industry-standard root-of-trust systems to side-channel attacks. It underscores the urgent need for robust countermeasures to secure commercially deployed systems against such threats.

Index Terms—Side-channel Attacks, ML-DSA, Quantum-Resistant Cryptography, Correlation Power Analysis, Zeroization, Modular Reduction, Root-of-Trust.

I. INTRODUCTION

Quantum computing presents a significant challenge to current classical cryptosystems based on the computational difficulty of problems like integer factorization [1] and discrete logarithms [2], [3]. Shor’s algorithm demonstrates that these problems are vulnerable to quantum computers, which can tackle their complexity in polynomial time [4]. To address this threat, the National Institute of Standards and Technology (NIST) initiated the post-quantum cryptography (PQC) standardization process [5], evaluating multiple candidates for digital signatures and key encapsulation algorithms. CRYSTALS-Dilithium, Falcon, and SPHINCS+ were selected as standards for digital signatures, with CRYSTALS-Dilithium standing out for its robustness and adaptability. NIST has officially renamed CRYSTALS-Dilithium as the Modular Lattice Digital Signature Algorithm (ML-DSA), a term used consistently throughout this paper [6].

ML-DSA is a lattice-based digital signature scheme derived from the Fiat-Shamir paradigm [6], with security rooted in

the computational difficulty of finding short vectors in lattices. It combines strong security guarantees with operational efficiency [7]. However, like many cryptographic schemes, its implementations might be vulnerable to side-channel attacks [8]–[26], which exploit information leakage through power consumption, electromagnetic emissions, and other observable channels. Addressing these vulnerabilities is critical, as side-channel resilience is essential for secure use in real-world applications.

Given the critical nature of cryptographic operations, prior research focuses extensively on analyzing and exploiting side-channel vulnerabilities [8]–[22] to understand their implications and propose mitigations [27]–[31]. Most of these studies target software implementations [8]–[20], as they are more accessible for evaluation and frequently employed in early-stage algorithm testing. While post-quantum cryptography can be implemented in software, the computationally intensive structure of the ML-DSA algorithm makes hardware implementations the preferred choice for security-critical applications such as Root of Trust (RoT) [32]. Prior research on software implementations often focuses on reference designs that are unrepresentative of deployed systems.

This leaves a critical gap in pinpointing the vulnerabilities of hardware implementations, particularly within widely adopted cryptographic frameworks. In contrast to extensive research on software implementations, studies targeting hardware side-channel attacks remain comparatively limited [21], [22]. Existing works, such as those by Steffen *et al.* [21] and Wang *et al.* [22], focus primarily on standalone operations in ML-DSA or CRYSTALS-Dilithium hardware implementation, which are not aligned with NIST’s Federal Information Processing Standards (FIPS) 204 standardization [6]. Notably, none of these works target post-quantum cryptographic implementations used in industry-level applications.

This work addresses a critical gap by targeting a cryptographic Intellectual Property (IP) [34], an implementation of ML-DSA, within a Silicon Root of Trust framework [32]. This collaborative framework, developed by leading technology companies, is designed to enhance trust in modern System-on-Chip (SoC) platforms. **It is important to note that the implementation analyzed in this work represents an early, work-in-progress version, which continues to undergo updates and has not yet been finalized.** We present

TABLE I: Comparison of Hardware Side-Channel Attacks on CRYSTALS-Dilithium and ML-DSA

Work	Type	Implementation	Target Op.	Traces ($\times 10^3$)	Success Rate	Real-World App.
[21]	P	Dilithium [33]	NTT	350	94.2%	✗
[21]	P	Dilithium [33]	Decoding	50	93.2%	✗
[21]	NP	Dilithium [33]	Pointwise Mult.	66	–	✗
[22]	NP	Dilithium [33]	Pointwise Mult.	70	–	✗
This work	NP	ML-DSA [34]	Pointwise Mult.	10	99.99%	✓

P = Profiling, NP = Non-Profiling, Target Op. = Target Operation, Real-World App. = Real-World Application

the first side-channel attack on a post-quantum cryptography implementation used in an industry-standard setting. Our approach achieves a 99.99% success rate using only 10,000 power traces, significantly outperforming prior works. Table I highlights these advancements, comparing our results with previous hardware attacks on CRYSTALS-Dilithium and ML-DSA. For example, Steffen *et al.* achieved a 94.2% success rate in profiling attack on Number Theoretic Transform (NTT) operation using 350,000 traces and a 93.2% success rate in decoding with 50,000 traces. Similarly, Wang *et al.* conducted a CPA attack that required 70,000 traces.

Our non-profiling approach specifically targets the ML-DSA signing process, eliminating the need for identical devices or pre-generated templates. This attack effectively exposes vulnerabilities in the implementation of unprotected ML-DSA hardware, emphasizing the critical need for robust countermeasures in post-quantum cryptography implementations. The contributions of this paper are summarized as follows:

- We present the first side-channel attack targeting the hardware implementation of ML-DSA [34], an open-source design intended for adoption in hardware platforms by major technology companies such as the leading technology companies. Our findings reveal a critical vulnerability in a widely used hardware design.
- We target the NTT-based pointwise multiplication operation, identifying and exploiting previously unknown vulnerabilities. These vulnerabilities allow us to recover partial secret key coefficients, which are integral to the certificate generation process within the Silicon RoT.
- We analyze and identify a vulnerability in the reduction algorithm implemented in the studied hardware [34], designed specifically for the modulus $q = 8380417 = 2^{23} - 2^{13} + 1$. Unlike commonly used reduction techniques such as Barrett [35] or Montgomery [36] algorithms, this implementation leverages a particular structure of the modulus to optimize performance, inadvertently introducing exploitable side-channel leakage.
- We demonstrate the efficiency of our side-channel attack by successfully recovering partial secret key coefficients with a 99.99% success rate using only 10,000 power traces. This significantly improves on existing attacks, which need far more traces for comparable or lower success rates.

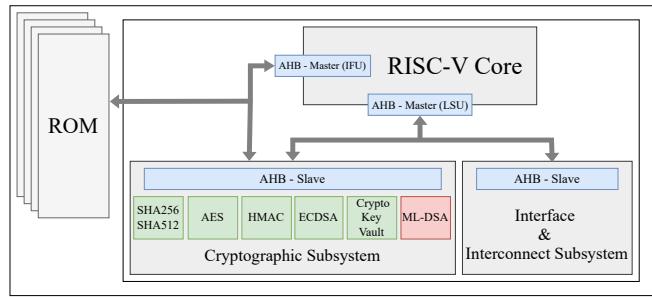


Fig. 1: High level of the RoT, which employs classical and PQC digital signatures for authentication [32].

II. BACKGROUND

This section provides an overview of open-source implementation of ML-DSA [34] within the silicon RoT internal block [32], focusing on its parameter settings and operations involving the secret key, which are of interest to adversaries. Additionally, it introduces the adversarial capabilities and outlines the general CPA methods used to recover the secret key.

A. Notations

Let n and q be two integers, where $n = 256$ and $q = 8380417$. We define R_q as the polynomial ring $\mathbb{Z}[x]/(x^n + 1)$. The infinity norm $\|x\|_\infty$ denotes the maximum absolute value among all coefficients of a polynomial x .

Matrices are represented by bold capital letters, such as $\mathbf{A} \in R_q^{k \times \ell}$. Vectors are represented by bold lowercase letters, such as \mathbf{y} . For polynomial vector $\mathbf{s}_1 \in \mathbb{R}_q^\ell$, it is pseudorandomly sampled from the set of ℓ -dimensional polynomial vectors, where each coordinate polynomial has short coefficients within the range $[-\eta, \eta]$. Polynomials in the NTT domain [7] are denoted with a hat (e.g., \hat{c} , where $\hat{c} = \text{NTT}(c)$). This notation is transitive, so $\hat{\mathbf{s}}$ represents each polynomial in \mathbf{s} being individually transformed into the NTT domain. The symbol "o" denotes pointwise multiplication, and " \perp " indicates an invalid or undefined value. The set B_τ represents all possible challenge polynomials with exactly τ coefficients of ± 1 , and the rest 0. Lastly, H represents the SHAKE algorithm, as defined in [6]. This cryptographic hash function is used for key derivation and signature binding operations throughout the ML-DSA signature scheme.

B. Root of Trust

This hardware security module provides essential RoT capabilities designed for integration into SoC systems used in cloud platforms [32]. Figure 1 illustrates RoT block's high-level architecture, which consists of a RISC-V core, ROM, cryptographic subsystem, and an interface & interconnect subsystem, all interconnected via the AHB (Advanced High-Performance Bus) protocol. These components collaboratively enable secure device identity and hierarchical key derivation. This architecture ensures robust protection against firmware compromise while facilitating secure initialization and reliable communication within the SoC system.

TABLE II: Parameter sets of ML-DSA [6]

Security Level	ML-DSA-44	ML-DSA-65	ML-DSA-87
Parameter	Values		
q	8380417	8380417	8380417
d	13	13	13
τ	39	49	60
λ	128	192	256
γ_1	2^{17}	2^{19}	2^{19}
γ_2	$\frac{q-1}{88}$	$\frac{q-1}{32}$	$\frac{q-1}{32}$
(k, ℓ)	(4,4)	(6,5)	(8,7)
η	2	4	2
$\beta = \tau \cdot \eta$	78	196	120
ω	80	55	75

The ROM provides immutable storage for the RoT code, serving as the foundation for secure initialization and attestation. It is directly linked to the RISC-V core, which features two AHB master interfaces: the Instruction Fetch Unit (IFU) for fetching instructions and the Load/Store Unit (LSU) for data transfers. These interfaces interact with the cryptographic subsystem and the interface and interconnect subsystem via AHB slave connections. The cryptographic subsystem supports hashing (SHA-256/SHA-512), encryption (AES), keyed-hash authentication (HMAC), elliptic curve signatures (ECDSA), and a classical cryptographic key vault [32].

In addition to classical cryptography, the subsystem integrates ML-DSA, a post-quantum digital signature scheme, to provide quantum resistance. The hardware security module uses ML-DSA-87 in combination with ECC Secp384r1 to implement a dual-signature scheme for verifying firmware integrity, ensuring a security level of approximately 192 bits for ECDSA [37] and level 5 for ML-DSA. By securing services such as secure boot, firmware verification, and key management, ML-DSA is critical in maintaining the integrity of runtime (RT) code and RoT for Measurement (RTM) [32].

However, our results demonstrate that open-source ML-DSA implementation introduces vulnerabilities that are exploitable via side-channel attacks. Specifically, we show that an adversary can extract the ML-DSA secret keys used to establish SoC attestation. With these keys, an attacker could forge signatures, compromise firmware integrity, and ultimately undermine the RoT. This attack highlights the critical need for robust countermeasures to secure post-quantum cryptographic implementations within hardware security modules and similar platforms.

C. Configuring ML-DSA

To better understand the significance and vulnerabilities of open-source ML-DSA implementation in this framework, it is essential to examine ML-DSA's algorithm steps and security properties. ML-DSA is a lattice-based digital signature scheme secure against quantum computers [6]. Its security is based on the hardness of the Module Learning With Errors (Module-LWE) and Module Short Integer Solution (Module-MSIS) problems [38]. It has been selected by the NIST as the primary algorithm for quantum-secure digital signatures due to its efficiency, robust security properties, and scalability.

Algorithm 1 Key Generation [6]

Input: Seed ξ
Output: Public key pk and private key sk

- 1: $(\rho, \rho', K) \leftarrow H(\xi)$
- 2: $\mathbf{A} \in R_q^{k \times \ell} \leftarrow \text{ExpandA}(\rho)$
- 3: $(\mathbf{s}_1 \in \mathbb{R}_q^\ell, \mathbf{s}_2 \in \mathbb{R}_q^k) \leftarrow \text{ExpandS}(\rho')$ ▷ Generate secrets.
- 4: $\mathbf{t} \leftarrow \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$
- 5: $(\mathbf{t}_1, \mathbf{t}_0) \leftarrow \text{Power2Round}(\mathbf{t}, d)$
- 6: $pk \leftarrow \text{pkEncode}(\rho, \mathbf{t}_1)$
- 7: $tr \leftarrow H(pk)$
- 8: $sk \leftarrow \text{skEncode}(\rho, K, tr, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$ ▷ Encode the private key.
- 9: **return** (pk, sk)

Algorithm 2 Signature Generation [6]

Input: Secret key sk , message M
Output: Signature σ

- 1: $(\rho, K, tr, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0) \leftarrow \text{skDecode}(sk)$
- 2: $\mathbf{A} \in R_q^{k \times \ell} \leftarrow \text{ExpandA}(\rho)$
- 3: $\mu \leftarrow H(tr \| M)$
- 4: $\rho'' \leftarrow H(K \| \mu)$
- 5: $\kappa \leftarrow 0, (\mathbf{z}, \mathbf{h}) \leftarrow \perp$
- 6: **while** $(\mathbf{z}, \mathbf{h}) = \perp$ **do** ▷ Rejection sampling loop
- 7: $\mathbf{y} \leftarrow \text{ExpandMask}(\rho'', \kappa)$
- 8: $\mathbf{w} \leftarrow \mathbf{A}\mathbf{y}$
- 9: $\mathbf{w}_1 \leftarrow \text{HighBits}(\mathbf{w})$
- 10: $\tilde{c} \leftarrow H(\mu \| \mathbf{w}_1)$
- 11: $c \in B_\tau \leftarrow \text{SampleInBall}(\tilde{c})$
- 12: $\mathbf{z} \leftarrow \mathbf{y} + \boxed{cs_1}$ ▷ Target operation: pointwise multiplication
- 13: $r_0 \leftarrow \text{LowBits}(\mathbf{w} - cs_2)$
- 14: **if** $\|\mathbf{z}\|_\infty \geq \gamma_1 - \beta$ **or** $\|r_0\|_\infty \geq \gamma_2 - \beta$ **then**
- 15: $(\mathbf{z}, \mathbf{h}) \leftarrow \perp$
- 16: **else**
- 17: $\mathbf{h} \leftarrow \text{MakeHint}(-c\mathbf{t}_0, \mathbf{w} - cs_2 + c\mathbf{t}_0)$
- 18: **if** $\|c\mathbf{t}_0\|_\infty \geq \gamma_2$ **or** # of 1's in $\mathbf{h} > \omega$ **then**
- 19: $(\mathbf{z}, \mathbf{h}) \leftarrow \perp$
- 20: **end if**
- 21: **end if**
- 22: $\kappa \leftarrow \kappa + \ell$
- 23: **end while**
- 24: **return** $\sigma = (\tilde{c}, \mathbf{z}, \mathbf{h})$

The parameters in Table II define the scheme for three security levels, with ML-DSA-87 offering the highest security level. Across all levels, the modulus $q = 8380417$ and the base ring $R_q = \mathbb{Z}_q[x]/(x^{256} + 1)$ remain constant, while parameters such as (k, ℓ) , τ , and η vary to balance efficiency and security. At higher levels, larger τ and (k, ℓ) values are used to enhance robustness, while η determines the coefficient range of secret polynomials. Derived parameters like β and ω further influence performance and compliance with constraints. The target design [34] is implemented with the ML-DSA-87 parameters: $(k, \ell) = (8, 7)$, $\tau = 60$, and $\eta = 2$. These settings provide the strongest cryptographic guarantees within the ML-DSA scheme, ensuring maximum security.

The ML-DSA scheme includes key generation, signature generation, and signature verification. This paper focuses on the key generation and signature generation, as these processes work with secret variables that may be vulnerable to side-channel attacks. In contrast, the signature verification function does not operate with secret keys and, therefore, is not relevant




Fig. 2: Illustration of the ML-DSA signing process, highlighting the multiplication of c and s_1 in red, which is the primary target for side-channel analysis.

to this attack. Consequently, this paper does not discuss the signature verification process.

The key generation procedure, as defined in [6] and detailed in Algorithm 1, begins with a seed ξ that generates all cryptographic parameters. Using a hash function H , ξ is expanded into ρ , ρ' , and K . ρ deterministically generates the public matrix $\mathbf{A} \in R_q^{k \times \ell}$ via `ExpandA`, while ρ' derives the secret polynomials $s_1 \in \mathbb{R}_q^\ell$ and $s_2 \in \mathbb{R}_q^k$. The public vector $\mathbf{t} = \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$, representing a Module-LWE instance, is split by `Power2Round` into high-order bits \mathbf{t}_1 (included in the public key) and low-order bits \mathbf{t}_0 (stored in the secret key). The public key is $pk = (\rho, \mathbf{t}_1)$, while the private key is $sk = (\rho, K, tr, s_1, s_2, \mathbf{t}_0)$, providing all components required for signing. This process outputs a compact public key for efficient transmission and storage and a detailed private key for signing operations.

The signature generation procedure, as outlined in Algorithm 2 and illustrated in Figure 2, begins with reconstructing the public matrix $\mathbf{A} \in R_q^{k \times \ell}$ using the `ExpandA` function and the seed ρ . Also, the secret key sk is decoded to retrieve its components: ρ , K , tr , s_1 , s_2 , and \mathbf{t}_0 . To bind the signature to the message, the concatenation of tr and the message M is hashed using H to produce μ . Another hash operation involving K and μ generates ρ'' , which seeds the pseudorandom masking polynomial \mathbf{y} via the `ExpandMask` function. The intermediate vector $\mathbf{w} = \mathbf{A}\mathbf{y}$ is then computed, followed by extracting its high-order bits \mathbf{w}_1 using the `HighBits` function. The challenge polynomial c is generated by hashing μ and \mathbf{w}_1 , and it is subsequently sampled with fixed size B_τ using `SampleInBall`. The core operation of the signing process is the computation of the response polynomial $\mathbf{z} = \mathbf{y} + c\mathbf{s}_1$, where $c\mathbf{s}_1$ represents the pointwise modular multiplication. The polynomial c and the secret polynomial s_1 are central to this operation, which is highlighted as the primary target in the proposed attack. This operation is crucial for linking the signature to the secret key, and its details are further discussed in Section III. The modular reduction applied during this step ensures that all coefficients remain within the defined modulus q . The low-order bits of $\mathbf{w} - c\mathbf{s}_2$ are extracted as r_0 , which are subjected to rejection sampling to ensure \mathbf{z} and r_0 meet security constraints. If the constraints are not satisfied, the process resets (\mathbf{z}, \mathbf{h}) and derives a new \mathbf{y} . After satisfying

the constraints, the hint polynomial \mathbf{h} is generated, which compresses low-order bits to enhance verification efficiency. The final signature, $\sigma = (\hat{c}, \mathbf{z}, \mathbf{h})$, binds the message M to the private key components, completing the signing process.

The ML-DSA scheme offers two signing variants: deterministic and hedged. The key distinction lies in the use of randomness during the generation of the commitment value \mathbf{y} . In the hedged variant, fresh randomness is introduced via a Random Bit Generator (RBG) to mitigate side-channel and fault attacks that exploit predictable patterns [6].

Incorporating randomness during the generation of \mathbf{y} would randomize the challenge polynomial c , disrupting the patterns exploited by side-channel analysis. This would significantly reduce the effectiveness of the proposed attack by obscuring the predictable relationship between power consumption and the processed data. However, the studied implementation [34] is configured to work with ML-DSA deterministic variant. Thus, this implementation creates an exploitable vulnerability.

The deterministic signing approach, adopted in the studied ML-DSA implementation [34], is designed to create stable and predictable cryptographic outputs tied to the device's identity. This approach ensures reproducibility and avoids generating different signatures for the same device under identical conditions, simplifying device identity management and certificate creation [39]. However, while this behavior supports specific design objectives, it inadvertently introduces vulnerabilities. Specifically, deterministic signing lacks the fresh randomness required to protect against side-channel attacks [6].

D. Threat Model

This work adheres to the standard assumptions in non-profiled power side-channel attacks, where the adversary has physical access to the device and captures power measurements during the execution of the signing operation. The adversary provides input messages and obtains the corresponding signatures, which are known and non-secret information.

In this context, the adversary targets the Root of Trust (RoT) during the certificate generation process. The adversary can provide different firmware payloads as input messages to the signing process, leveraging the system's functionality to generate corresponding certificates. Since the input firmware is known and updatable, the adversary can carefully control the inputs while capturing power traces during the certificate

generation operation. Using these power traces, the adversary exploits side-channel leakage to extract critical secrets, such as private keys, used in the signing process. This aligns with the standard non-profiled power side-channel attack assumptions.

For the proposed differential side-channel attack, we employ the Pearson correlation coefficient as a distinguisher [40]. Unlike profiling attacks, the adversary does not rely on pre-built templates. Instead, a classical CPA approach is executed to recover the secret key. The CPA process is carried out in the following four steps:

- Collect n power traces during the signing operation, each consisting of m -length power samples. Store these traces in an $n \times m$ matrix T .
- Generate an $n \times k$ intermediate value matrix V corresponding to all possible key guesses. The intermediate values are calculated based on the controlled input values and the key candidates from the overall key space.
- Use the power model to map the intermediate value matrix V to an $n \times k$ Hamming Distance matrix H , where each element $H_{i,j}$ represents the Hamming Distance corresponding to $V_{i,j}$.
- Compute the correlation coefficients between the Hamming Distance matrix H and the actual power consumption matrix T . Store these coefficients in a correlation matrix $R_{k \times m}$. The Pearson correlation coefficient is calculated as:

$$R_{i,j} = \frac{\sum_{x=1}^n (H_{x,i} - \bar{H}_i)(T_{x,j} - \bar{T}_j)}{\sqrt{\sum_{x=1}^n (H_{x,i} - \bar{H}_i)^2 \cdot \sum_{x=1}^n (T_{x,j} - \bar{T}_j)^2}}. \quad (1)$$

Following these steps, the adversary identifies the key guess that maximizes the correlation, thereby recovering the secret key without requiring prior profiling.

III. THE PROPOSED SIDE-CHANNEL ATTACK

ML-DSA is fundamentally based on the Module-LWE problem, which relies on the hardness of recovering secret polynomials from structured linear equations. The cryptographic strength of ML-DSA hinges on maintaining the secrecy of key polynomials such as s_1 , s_2 , and t_0 . These polynomials are interconnected through relationships derived from the ML-DSA construction.

One of the key relationships in ML-DSA can be represented by the equation $\mathbf{z} = \mathbf{y} + c\mathbf{s}_1$, as illustrated in Figure 2, where \mathbf{z} and c are known polynomial sets, and \mathbf{s}_1 is a secret polynomial. Before multiplication, the polynomials are transformed into the NTT domain to enable efficient computations, as shown in Figure 3. Extracting \mathbf{s}_1 enables the recovery of \mathbf{y} , which is used to generate \mathbf{w} . Once \mathbf{w} is known, the attacker can compute \mathbf{s}_2 using the outputs of the `MakeHint` function. Ultimately, with \mathbf{s}_1 and \mathbf{s}_2 , the attacker can solve $t = A\mathbf{s}_1 + \mathbf{s}_2$, effectively compromising the Module-LWE foundation of ML-DSA. Consequently, recovering \mathbf{s}_1 provides the attacker with the full secret key of ML-DSA, enabling them to forge arbitrary signatures.




Fig. 3: Pointwise multiplication of the polynomial vector \mathbf{s}_1 and the challenge polynomial c in the NTT domain after transformation. Here, \hat{c} and $\hat{\mathbf{s}}_1$ represent the NTT-transformed forms of c and \mathbf{s}_1 , respectively.

Recovering \mathbf{s}_1 using CPA requires an operation where \mathbf{s}_1 interacts with a known, controllable value, and the interaction must vary based on different inputs to facilitate side-channel analysis. Since the key generation process produces unique \mathbf{s}_1 values for each input set, it is not suitable for CPA. Instead, we focus on the signature generation process, where \mathbf{s}_1 interacts with a publicly known and attacker-controllable value. As specified in line 12 of Algorithm 2, \mathbf{s}_1 undergoes pointwise multiplication with c .

As illustrated in Figure 2, c is derived during the signature generation process from the `SampleInBall` function, which uses a hash value \tilde{c} computed as $H(\mu || \mathbf{w}_1)$. Although \mathbf{w} is a secret value, the high bits of \mathbf{w} , which is \mathbf{w}_1 , is public value. This sequence ensures that c is a publicly known and controllable polynomial, making it an ideal candidate for side-channel analysis. This interaction satisfies the criteria for side-channel analysis, making it a highly effective target for recovery.

IV. IMPLEMENTATION-LEVEL ANALYSIS

This section provides an in-depth examination of the target implementation of [34], with a focus on the NTT implementation block, which is utilized for both NTT transformations and pointwise operations such as modular multiplication, addition, and subtraction. Analyzing this block is critical as it reveals the detailed operation involving the challenge polynomial c and the secret polynomial \mathbf{s}_1 . Understanding this interaction is essential to pinpointing the vulnerabilities in the implementation and identifying the specific stages where sensitive information may leak through side-channel channels.

A. NTT in Target Implementation

In ML-DSA, the polynomials are represented in the NTT domain to optimize computational efficiency during polynomial arithmetic. The use of the NTT domain enables faster computations, particularly for polynomial multiplications. NTT transforms polynomials from their coefficient representation into a point-value representation, enabling polynomial multiplication as simple element-wise multiplication. This transformation relies on modular arithmetic and carefully selected roots of unity. After computations in this transformed domain, the inverse NTT (INTT) is applied to return the result

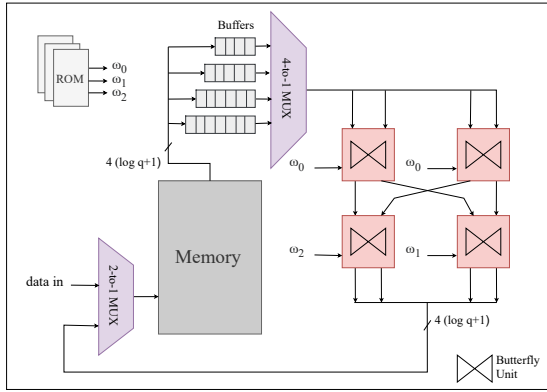


Fig. 4: NTT architecture in target implementation [34]. The design features a 2x2 butterfly architecture for optimized performance.

to the coefficient domain. The NTT architecture in the target implementation processes $n = 256$ coefficients per polynomial using eight iterative stages. These stages correspond to $\log_2 n = 8$.

Figure 4 illustrates the NTT architecture of the studied ML-DSA module, a pipelined design optimized for processing four coefficients per cycle. The inputs to the NTT are four coefficients read from memory and four precomputed n -th root of unity (ω) stored in ROM. The architecture features a 2x2 butterfly unit configuration with two stages of butterflies, each stage containing two butterfly units. This pipelined approach allows the initial stage to compute the first layer of the NTT transformation for four coefficients in each cycle while the second stage simultaneously processes the next layer. This design minimizes memory overhead by eliminating the need to write intermediate results back to memory between layers.

The memory system in the NTT design stores input coefficients, intermediate results, and transformed values, with each memory address holding four coefficients. The memory bus width is defined as $4 \cdot \log_2(q-1)$, where q is the prime modulus used in the NTT. During operation, coefficients are read in a specific sequence and buffered to synchronize with the pipeline. A 4-to-1 multiplexer (MUX) routes the coefficients to the butterfly units, and after processing, the results are directed back to memory via a 2-to-1 MUX. The 2-to-1 MUX enables selection between initializing the memory with polynomials in the non-NTT domain or passing intermediate NTT values from the butterfly units to continue with the next stage.

This NTT architecture also facilitates hardware resource sharing across other ML-DSA operations, such as pointwise multiplication and pointwise addition and subtraction. Each butterfly unit consists of a modular adder, a modular subtractor, and a modular multiplier, all of which operate with the prime modulus q . Instead of utilizing dedicated multipliers, adders, and subtractors for pointwise operations, the open-source ML-DSA implementation [34] design reuses the butterfly units within the NTT module to perform these tasks. This approach optimizes resource use but introduces shared vulnerabilities,




Fig. 5: Pointwise multiplication architecture in target design [34]. This design reuses the hardware utilized for the NTT block.

as all operations rely on the same hardware.

B. Pointwise Multiplication in Target Implementation

Pointwise multiplication in the target implementation performs element-wise multiplication of two polynomials in the NTT domain. As shown in Figure 5, the architecture reuses key NTT components, including memory, multiplexers, and pointwise multiplication cores instead of butterfly cores, to optimize resource usage while supporting the modular arithmetic required for cryptographic operations.

In this implementation, two memories are used to store the input polynomials, with each memory address holding four coefficients. The cores execute four modular multiplications in parallel, enabling the processing of multiple coefficients concurrently. A multiplexer is used to direct the coefficients to the appropriate processing units [34].

The modular multiplication cores implement a reduction algorithm. Pointwise multiplication enables concurrent processing and ensures results remain within the finite field defined by $q = 8,380,417$ through modular arithmetic. Modular reduction maintains all coefficients within the range $[0, q)$ after multiplying two 23-bit coefficients and obtaining a 46-bit product.

C. Internal Structure of Modular Reduction

These modular multiplication cores use a unique reduction algorithm. It is specifically optimized for the prime modulus $q = 8,380,417$, leveraging the modulus's advantageous structure to enhance throughput [34]. Unlike conventional reduction methods, this design implements a fully pipelined architecture that utilizes shifter and adder logic to perform modular reduction efficiently and with a constant execution time. The reduction circuit processes 46-bit intermediate products, which result from the multiplication of 23-bit coefficients during pointwise polynomial multiplication.

We explain the modular reduction process using the multiplication of \hat{s}_1 the secret polynomial and \hat{c} the challenge polynomial as an example. These specific variables are chosen for illustration because they represent the focus of our side-channel analysis, being the first point of interaction between the secret key coefficients and deterministically generated, known values.

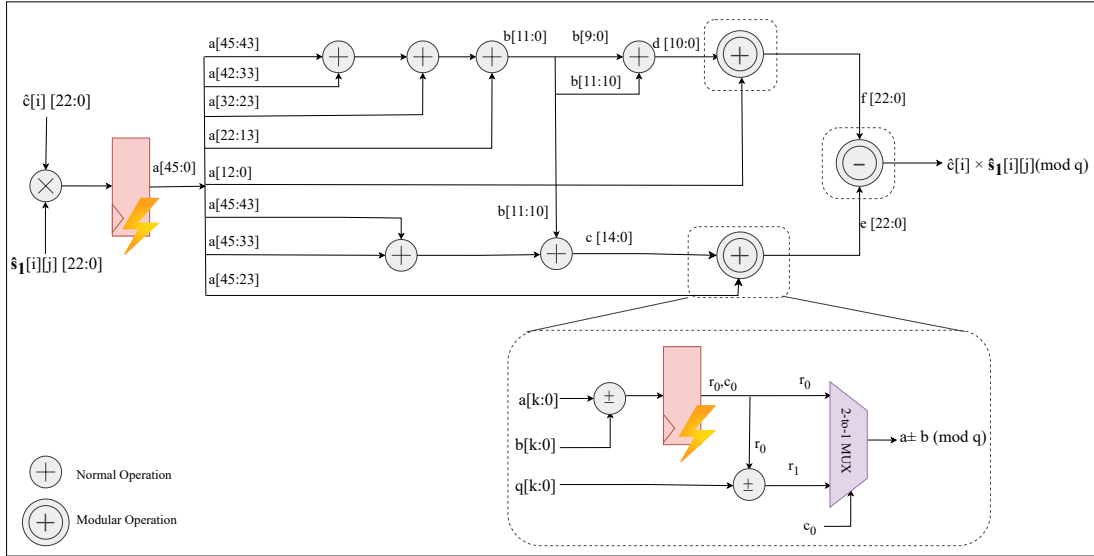


Fig. 6: The internal structure of the modular reduction circuit, as adapted from [34], is optimized for the prime modulus q . This design is particularly relevant for operations involving the multiplication of the NTT-transformed private polynomial \hat{s}_1 and the NTT-transformed challenge polynomial \hat{c} , where $\hat{s}_1[i][j]$ represents the j -th coefficient of the i -th polynomial in \hat{s}_1 , and $\hat{c}[i]$ is the i -th coefficient of \hat{c} .

The modular reduction begins with the computation of the 46-bit intermediate product:

$$a = \hat{s}_1[i][j] \cdot \hat{c}[i],$$

where $\hat{s}_1[i][j]$ is the j -th coefficient of the i -th polynomial in \hat{s}_1 , and $\hat{c}[i]$ is the i -th coefficient of \hat{c} . The modulus q is defined as:

$$q = 8,380,417 = 2^{23} - 2^{13} + 1,$$

which allows the reduction to leverage the equivalence:

$$2^{23} \equiv 2^{13} - 1 \pmod{q}.$$

The reduction algorithm decomposes the product a into bit segments:

$$a = 2^{13}a_{45:23} - a_{45:23} + a_{22:0},$$

where $a_{45:23}$ and $a_{22:0}$ represent the higher and lower bit ranges, respectively. To streamline computation, the reduction is defined as:

$$a = 2^{13}b - (a_{45:43} + a_{45:33} + a_{45:23}) + a_{12:0},$$

where:

$$b = a_{45:43} + a_{42:33} + a_{32:23} + a_{22:13}, \quad b < 2^{12}.$$

Further reduction of b is computed as:

$$2^{13}b = 2^{13}d - b_{11:10},$$

with:

$$d_{10:0} = b_{11:10} + b_{9:0}.$$

In the other stage, the 15-bit value $c_{14:0}$ is computed as:

$$c_{14:0} = a_{45:43} + a_{45:33}.$$

Additionally, the 23-bit value $e_{22:0}$ is computed as:

$$e_{22:0} = a_{45:23} + c_{14:0}.$$

Finally, the modular reduction result is:

$$a = d_{10:0} + a_{12:0} - e_{22:0},$$

where:

$$e_{22:0} = (a_{45:43} + a_{45:33} + a_{45:23}) + b_{11:10}.$$

D. Exploiting Modular Reduction Registers

Figure 6 exhibits the implementation details of the modular reduction circuit. This circuit features a three-stage pipelined process. In the first stage, the intermediate product $a = \hat{s}_1[i][j] \cdot \hat{c}[i]$ is computed during pointwise multiplication and buffered with 46-bit registers. In the second stage, the reduction circuit performs two modular addition operations in parallel, and these operations' intermediate values are stored in modular addition registers. Finally, the third stage performs a modular subtraction to ensure the result ranges within $[0, q)$. Throughout these stages, all intermediate values are stored in registers, forming potential leakage points that can be exploited through side-channel analysis.

The attack methodology exploits the deterministic behavior of the target hardware design, as implemented in the deterministic version of the IP [34], to target specific registers. It is worth noting that the IP also supports a non-deterministic configuration, which may impact the applicability of this methodology. To conduct the attack, different input messages are provided to the ML-DSA implementation [34], which computes the corresponding challenge polynomial c . This polynomial then interacts with the secret key s_1 through pointwise multiplication. Each coefficient of s_1 and c is multiplied



Fig. 7: Experimental setup for capturing power traces.

and subsequently processed by the reduction circuit, making the registers in this pipeline prime candidates for leakage exploitation.

The first vulnerability arises in the initial set of registers, which store the 46-bit intermediate product from the pointwise multiplication. By controlling the input message, the attacker can determine the exact value of the corresponding c coefficient. Using this knowledge, a hypothetical table for CPA is generated for each possible s_1 coefficient. Since \hat{s}_1 ranges between 0 and $q - 1$ after NTT transform, the attacker generates $q - 1$ hypothetical guesses. CPA is then performed as described in the II-D, correlating observed power traces with these hypothetical values. While this approach can reveal the coefficients of s_1 , it might be prone to false positives due to the linear nature of multiplication [41].

To enhance the attack and eliminate these false positives, we use the next set of registers in the modular addition stage. These registers store intermediate values computed by the addition as part of the modular reduction process. The same s_1 hypothetical guesses are reused, but the intermediate values in the CPA are adjusted to reflect the modular addition operation. By targeting this stage, the attack eliminates the false positives encountered from the first register set.

Zeroization in Target Design. A key aspect of this design is its zeroization method, implemented to prevent data recovery. Zeroization is a method of securely erasing cryptographic keys, critical security parameters, and electronically stored data by altering or clearing storage contents to prevent data recovery [42]. The RoT enforces the ML-DSA module in the cryptographic IP to perform zeroization after the key generation and signing processes. All internal registers are cleared using a software-triggered mechanism, where a single-cycle pulse on the hardware interface zeroes the first register before the next cryptographic operation begins. This ensures that the first pointwise multiplication core always starts with a clean state by removing residual data from previous operations and preventing the retention of sensitive information across cryptographic stages. While the zeroization mechanism in the

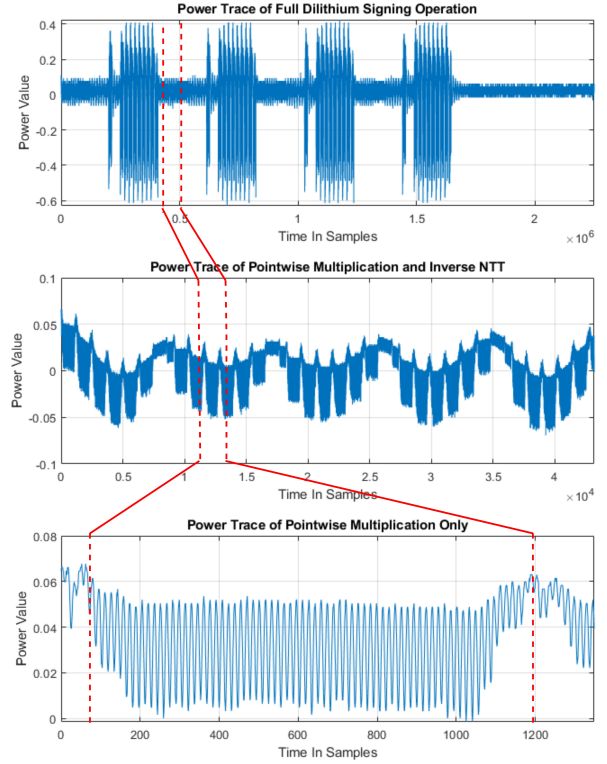


Fig. 8: Power traces of the ML-DSA signing operation. The top trace shows the full signing operation with multiple rejection loops. The middle trace isolates the NTT and pointwise multiplication stages, while the bottom trace zooms in on the initial pointwise multiplication.

studied ML-DSA module [34] is implemented as per FIPS 140-2 standards [42] to enhance security, it inadvertently aids side-channel analysis. By resetting all internal registers to zero before the start of key generation and signing operations, the mechanism simplifies attacks. The predictable initial state of the registers reduces the complexity of side-channel analysis, as attackers can leverage this knowledge to calculate the Hamming distance more easily and extract sensitive information. This dual effect highlights the importance of carefully balancing security measures to avoid unintentionally introducing vulnerabilities.

In the Hamming Distance leakage model, power leakage correlates with the bitwise XOR of a register's current and previous states. To exploit this model, the attacker needs to know or guess the initial state of the register before making hypotheses about subsequent states. Without zeroization, the initial state of each register would be unknown, requiring the attacker to simultaneously guess both the initial state and the next state of the register. This significantly complicates the attack, as it doubles the size of the hypothetical guess table by introducing an additional dimension of uncertainty for each coefficient.

With zeroization, the registers are guaranteed to start from a known state of zero per captured execution. This eliminates




Fig. 9: Results of the proposed side-channel analysis targeting two registers in the hardware implementation of ML-DSA. The first row shows the attack results for the multiplication result register, including the correlation plot (a), Pearson correlation evolution (b), and the rank of the correct key, which reaches the first position (c). The second row presents similar results for the modular arithmetic operation register, highlighting the correlation plot (d), Pearson correlation evolution (e), and the rank of the correct key (f).

the need to guess the initial state, reducing the guess table size by half and simplifying the analysis. The attacker can focus solely on guessing the next state of the registers based on hypothetical values of the secret coefficients s_1 . This reduction in complexity allows the attack to be performed more efficiently, as the Hamming Distance model aligns with the deterministic transitions introduced by the zeroization process.

V. EVALUATION RESULT

This section provides an overview of the experimental setup used to evaluate the proposed side-channel attack, including the hardware and software configuration for capturing and analyzing power traces. It also presents the results of leakage detection and correlation analysis, demonstrating the effectiveness of the attack in exploiting vulnerabilities in the ML-DSA hardware implementation.

A. Experimental Setup

The power traces of the signing execution in ML-DSA were collected using the ChipWhisperer-Lite and the CW305 Artix-7 FPGA target board [43] for side-channel evaluation. The experimental setup, shown in Figure 7, includes a Low-Pass Filter, a PicoScope 6404E oscilloscope [44], and the CW305 FPGA development board, which is equipped with a Xilinx Artix-7 chip. The FPGA was programmed with the open-source ML-DSA implementation¹, an implementation of ML-DSA [34], configured to operate at security level 5.

The oscilloscope samples power traces at a frequency of 156.25 MHz, corresponding to a sampling interval of 6.4

¹GitHub commit version: <https://github.com/chipsalliance/adams-bridge/tree/1f84699209b3cf77d5658e6b3101d6fec8ed0daa>

ns. Two probes are connected to the victim device: Channel A records the power traces with a noise filter applied, while Channel B captures the trigger signal. The FPGA board interfaces with the ChipWhisperer capture board for data acquisition and synchronization. The CPA analysis was conducted on a computer equipped with an Intel(R) Xeon(R) E5-1620 CPU running at 3.50 GHz and 32 GB of memory. The CPA program was implemented in Python 3.10.

Figure 8 illustrates the power traces captured during the ML-DSA signing process. The top subplot shows the complete signing operation, including multiple iterations of the rejection loop caused by out-of-bound signature values. The middle subplot focuses on the computational stages, including the INTT and pointwise multiplication. The bottom subplot zooms in on the pointwise multiplication step, showing the power transitions recorded during this critical operation. These power traces form the basis for side-channel evaluation in our experimental setup.

B. Visualizing Leakage: Evolution of Correlation

To further investigate the vulnerability of the target implementation, we conducted a CPA attack. The results of the side-channel analysis are presented in Figure 9, which visualizes the correlation evolution, Pearson correlation coefficients, and the rank of the correct key for two critical attack points. These figures provide a comprehensive overview of the leakage patterns and highlight the effectiveness of the proposed CPA methodology in exploiting these vulnerabilities.

The first attack targets the register holding the 46-bit intermediate result of the multiplication between the 23-bit coefficients of the secret polynomial vector $s_1[i]$ and the chal-

lenge polynomial $\hat{c}[i]$. As depicted in Figure 9a, the correlation plot illustrates the leakage caused by state transitions in this register. The attack leverages the Hamming Distance leakage model to correlate the power consumption with transitions in the register, enabling the adversary to distinguish the correct key from incorrect guesses. After processing 10,000 traces, the correlation associated with the correct key surpasses all other candidates. This observation is further supported by the Pearson correlation evolution in Figure 9b, which shows the correct key consistently demonstrating a progressively stronger and statistically significant correlation (99.99%) compared to incorrect guesses. The rank of the correct key, shown in Figure 9c, converges rapidly, demonstrating the success of the attack in recovering the secret polynomial coefficients. The effectiveness of this attack is attributed to the predictable nature of the pointwise multiplication operation and the leakage arising from state transitions within the register storing the multiplication result.

The analysis was extended to a second attack point: the modular addition operation within the reduction circuit to evaluate the generality and robustness of the CPA methodology. This stage processes the intermediate multiplication result through modular arithmetic, temporarily storing intermediate values in registers. The internal structure of the modular reduction submodule, as illustrated in Figure 6, consists of modular addition and subtraction operations that follow the same flow and structure. These consistent computational patterns introduce predictable leakage, making the modular addition registers an equally viable target for side-channel analysis. Figure 9d demonstrates the correlation plot for this attack point, where the correct key is clearly distinguishable from other guesses. Similarly, the Pearson correlation evolution in Figure 9e demonstrates that the correct key progressively achieves higher and statistically significant correlation values (99.99%) compared to incorrect guesses as the number of traces increases. The rank of the correct key, presented in Figure 9f, converges rapidly, further validating the success of this attack point.

The effectiveness of these attacks is further facilitated by the constrained key guessing space in the hardware implementation [34], where modular reduction is applied after each butterfly operation during the NTT computation. Specifically, the key space in hardware is restricted to $[0, q)$, which is significantly smaller compared to the broader key range in the reference software implementation of ML-DSA [45], spanning $[-\eta - 8(q - 1), \eta + 8(q - 1)]$. This smaller key space reduces the complexity of the attack, making it easier to pinpoint the correct key.

The results from both attack points conclusively demonstrate that the open-source implementation of ML-DSA [34], an implementation of ML-DSA, remains vulnerable to side-channel leakage, particularly in the modular arithmetic stages. The CPA methodology employed in this paper effectively exploits predictable state transitions and the repetitive nature of modular arithmetic operations, emphasizing the necessity of robust countermeasures in hardware implementations of post-

quantum cryptographic algorithms.

VI. DISCUSSION

Our attack uncovers critical vulnerabilities in the targeted cryptographic implementation [34], specifically targeting the NTT-based polynomial pointwise multiplication. At the time of this study, the implementation lacked side-channel countermeasures such as masking [27]–[30] or shuffling [31], leaving it highly susceptible to side-channel leakage. While our findings highlight these vulnerabilities, the evaluation and implementation of countermeasures fall outside the scope of this work. Instead, our primary objective is to provide a detailed analysis of the identified weaknesses to inform the development of more secure cryptographic designs.

This work focuses on pinpointing vulnerabilities in the ML-DSA hardware implementation [34], particularly those arising in the NTT-based pointwise multiplication. This study emphasizes the critical need to address potential vulnerabilities in hardware implementations to prevent exploitation. By identifying these risks and communicating them to the broader cryptographic and hardware security communities, we aim to encourage proactive measures to mitigate side-channel attacks before they can be leveraged in practice.

It is important to note that the vulnerabilities identified in this work may not be confined to the NTT-based pointwise multiplication. Other components of the design [34], such as encoding, decoding, and NTT, could also leak sensitive information. Future research should explore these areas to ensure comprehensive protection against side-channel attacks and uncover additional weaknesses that may exist in the design.

Additionally, advanced techniques such as machine learning-based side-channel analysis [13]–[20] present a promising avenue for future research. These methods can identify complex or previously hidden leakage patterns, offering deeper insights into hardware vulnerabilities and informing the development of more robust countermeasures. Machine learning-based approaches also enable more generalized evaluations across different cryptographic implementations, enhancing their utility in securing hardware designs.

This study underscores the importance of systematically analyzing and addressing side-channel vulnerabilities in cryptographic implementations. As cryptographic algorithms transition from theoretical constructs to practical hardware deployments, ensuring resilience against both mathematical and physical attacks becomes imperative. By exposing these vulnerabilities and offering directions for future research, this work seeks to bridge the gap between theoretical security and practical robustness, advancing the security of hardware cryptographic designs in real-world applications.

VII. CONCLUSION

This paper exposes vulnerabilities in the hardware implementation of ML-DSA [34], specifically within an open-source Silicon RoT internal block. Using CPA, we demonstrated that critical cryptographic operations—modular reduction during

NTT-based polynomial pointwise multiplication—are vulnerable to side-channel leakage. Our analysis shows that an adversary can extract secret keys with as few as 10,000 power traces, posing a significant threat to firmware verification and RoT integrity. These findings underscore the need for robust countermeasures to secure post-quantum cryptographic implementations and highlight the importance of addressing advanced attack methodologies to mitigate emerging risks.

REFERENCES

- [1] R. L. Rivest, A. Shamir, and L. Adleman, “A method for obtaining digital signatures and public-key cryptosystems,” *Communications of the ACM*, vol. 21, pp. 120–126, Feb. 1978.
- [2] W. Diffie and M. Hellman, “New directions in cryptography,” *IEEE Transactions on Information Theory*, vol. 22, no. 6, pp. 644–654, 1976.
- [3] N. Kobitz, “Elliptic curve cryptosystems,” *Mathematics of Computation*, vol. 48, no. 177, pp. 203–209, 1987.
- [4] P. W. Shor, “Algorithms for quantum computation: Discrete logarithms and factoring,” in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 124–134, 1994.
- [5] D. Moody, “Post-quantum cryptography standardization: Announcement and outline of NIST’s call for submissions,” in *International Conference on Post-Quantum Cryptography (PQCrypto)*, 2016.
- [6] National Institute of Standards and Technology, “FIPS 204: Module-lattice-based digital signature standard,” tech. rep., U.S. Department of Commerce, Aug. 2024. Accessed: 2024-11-16.
- [7] L. Ducas, T. Lepoint, V. Lyubashevsky, P. Schwabe, G. Seiler, and D. Stehlé, “CRYSTALS – dilithium: Digital signatures from module lattices,” *Cryptology ePrint Archive*, Paper 2017/633, 2017.
- [8] P. Ravi, M. P. Jhanwar, J. Howe, A. Chattopadhyay, and S. Bhasin, “Side-channel assisted existential forgery attack on dilithium - a NIST PQC candidate,” *Cryptology ePrint Archive*, Paper 2018/821, 2018.
- [9] A. Fournaris, C. Dimopoulos, and O. Koufopavlou, *Profiling Dilithium Digital Signature Traces for Correlation Differential Side-Channel Attacks*, pp. 281–294. Oct. 2020.
- [10] R. Primas, P. Pessl, and S. Mangard, “Single-trace side-channel attacks on masked lattice-based encryption,” *Cryptology ePrint Archive*, Paper 2017/594, 2017.
- [11] Z. Chen, E. Karabulut, A. Aysu, Y. Ma, and J. Jing, “An efficient non-profiled side-channel attack on the crystals-dilithium post-quantum signature,” in *2021 IEEE 39th International Conference on Computer Design (ICCD)*, pp. 583–590, 2021.
- [12] Z. Qiao, Y. Liu, Y. Zhou, J. Ming, C. Jin, and H. Li, “Practical public template attacks on crystals-dilithium with randomness leakages,” *IEEE Transactions on Information Forensics and Security*, vol. 18, pp. 1–14, 2023.
- [13] S. Marzougui, V. Ulitzsch, M. Tibouchi, and J.-P. Seifert, “Profiling side-channel attacks on dilithium: A small bit-fiddling leak breaks it all,” *Cryptology ePrint Archive*, Paper 2022/106, 2022.
- [14] P. Pessl and R. Primas, “More practical single-trace attacks on the number theoretic transform,” in *Progress in Cryptology—LATINCRYPT 2019: 6th International Conference on Cryptology and Information Security, Proceedings 6*, pp. 130–149, Springer, 2019.
- [15] M. J. Kannwischer, P. Pessl, and R. Primas, “Single-trace attacks on keccak,” *Cryptology ePrint Archive*, 2020.
- [16] J. Han, T. Lee, J. Kwon, J. Lee, I.-J. Kim, J. Cho, D.-G. Han, and B.-Y. Sim, “Single-trace attack on NIST round 3 candidate dilithium using machine learning-based profiling,” *IEEE Access*, vol. PP, pp. 1–1, 12 2021.
- [17] I.-J. Kim, T.-H. Lee, J. Han, B.-Y. Sim, and D.-G. Han, “Novel single-trace ML profiling attacks on NIST 3 round candidate dilithium,” *Cryptology ePrint Archive*, Paper 2020/1383, 2020.
- [18] R. Wang, K. Ngo, J. Gärtner, and E. Dubrova, “Single-trace side-channel attacks on CRYSTALS-dilithium: Myth or reality?,” *Cryptology ePrint Archive*, Paper 2023/1931, 2023.
- [19] A. Berzati, A. C. Viera, M. Chartouny, S. Madec, D. Vergnaud, and D. Vigilant, “Exploiting intermediate value leakage in dilithium: A template-based approach,” *Cryptology ePrint Archive*, Paper 2023/050, 2023.
- [20] O. Bronchain, M. Azouaoui, M. ElGhamrawy, J. Renes, and T. Schneider, “Exploiting small-norm polynomial multiplication with physical attacks: Application to CRYSTALS-dilithium,” *Cryptology ePrint Archive*, Paper 2023/1545, 2023.
- [21] H. Steffen, G. Land, L. Kogelheide, and T. Güneysu, “Breaking and protecting the crystal: Side-channel analysis of dilithium in hardware,” *Cryptology ePrint Archive*, Paper 2022/1410, 2022.
- [22] H. Wang, Y. Gao, Y. Liu, *et al.*, “In-depth correlation power analysis attacks on a hardware implementation of crystals-dilithium,” *Cybersecurity*, vol. 7, no. 21, 2024.
- [23] C. Mujdei, A. Beckers, J. M. B. Mera, A. Karmakar, L. Wouters, and I. Verbauwhede, “Side-channel analysis of lattice-based post-quantum cryptography: Exploiting polynomial multiplication,” *Cryptology ePrint Archive*, Paper 2022/474, 2022.
- [24] M. Hamburg, J. Hermelink, R. Primas, S. Samardjiska, T. Schamberger, S. Streit, E. Strieder, and C. van Vredendaal, “Chosen ciphertext k-trace attacks on masked CCA2 secure kyber,” *Cryptology ePrint Archive*, Paper 2021/956, 2021.
- [25] J.-P. D’Anvers, M. Tiepelt, F. Vercauteren, and I. Verbauwhede, “Timing attacks on error correcting codes in post-quantum schemes,” *Cryptology ePrint Archive*, Paper 2019/292, 2019.
- [26] G. Rajendran, P. Ravi, J.-P. D’Anvers, S. Bhasin, and A. Chattopadhyay, “Pushing the limits of generic side-channel attacks on LWE-based KEMs - parallel PC oracle attacks on kyber KEM and beyond,” *Cryptology ePrint Archive*, Paper 2022/931, 2022.
- [27] V. Migliore, B. Gérard, M. Tibouchi, and P.-A. Fouque, “Masking dilithium: Efficient implementation and side-channel evaluation,” *Cryptology ePrint Archive*, Paper 2019/394, 2019.
- [28] M. Azouaoui, O. Bronchain, G. Cassiers, C. Hoffmann, Y. Kuzovkova, J. Renes, M. Schönauer, T. Schneider, F.-X. Standaert, and C. van Vredendaal, “Protecting dilithium against leakage: Revisited sensitivity analysis and improved implementations,” *Cryptology ePrint Archive*, Paper 2022/1406, 2022.
- [29] J.-S. Coron, F. Gérard, M. Trannoy, and R. Zeitoun, “Improved gadgets for the high-order masking of dilithium,” *Cryptology ePrint Archive*, Paper 2023/896, 2023.
- [30] J.-S. Coron, F. Gérard, T. Lepoint, M. Trannoy, and R. Zeitoun, “Improved high-order masked generation of masking vector and rejection sampling in dilithium,” *Cryptology ePrint Archive*, Paper 2024/1149, 2024.
- [31] J. Lee, J. Han, S. Lee, J. Kwon, K.-H. Choi, J.-W. Huh, J. Cho, and D.-G. Han, “Systematization of shuffling countermeasures: With an application to crystals-dilithium,” *IEEE Access*, vol. 11, pp. 142862–142873, 2023.
- [32] C. A. Project, “Caliptra: A datacenter system-on-a-chip (soc) root-of-trust (rot),” 2024. Revision 2.0, Version 0.5. Available: <https://github.com/chipsalliance/Caliptra>, Accessed: 2024-10-25.
- [33] L. Beckwith, D. T. Nguyen, and K. Gaj, “High-performance hardware implementation of CRYSTALS-dilithium,” *Cryptology ePrint Archive*, Paper 2021/1451, 2021.
- [34] C. Alliance, “Adams-bridge,” 2024. Available: <https://github.com/chipsalliance/adams-bridge>, Accessed: 2024-10-25.
- [35] P. Barrett, “Implementing the rivest shamir and adleman public key encryption algorithm on a standard digital signal processor,” in *Advances in Cryptology — CRYPTO’ 86*, pp. 311–323, 1987.
- [36] P. L. Montgomery, “Modular multiplication without trial division,” *Mathematics of Computation*, vol. 44, pp. 519–521, 1985.
- [37] L. Chen, D. Moody, K. Randall, A. Regenscheid, and A. Robinson, “Recommendations for discrete logarithm-based cryptography: Elliptic curve domain parameters,” Special Publication (NIST SP) 800-186, National Institute of Standards and Technology (NIST), 2023.
- [38] A. Langlois and D. Stehlé, “Worst-case to average-case reductions for module lattices,” *Cryptology ePrint Archive*, Paper 2012/090, 2012. Accessed: 2024-12-04.
- [39] P. England, R. Aigner, K. Kane, A. Marochko, D. Mattoon, R. Spiger, S. Thom, and G. Zaverucha, “Device identity with dice and riot: Keys and certificates,” Tech. Rep. MSR-TR-2017-41, Microsoft Research, September 2017.
- [40] E. Brier, C. Clavier, and F. Olivier, “Correlation power analysis with a leakage model,” in *Cryptographic Hardware and Embedded Systems - CHES 2004, August 11-13, 2004. Proceedings*, vol. 3156 of *Lecture Notes in Computer Science*, pp. 16–29, 2004.
- [41] E. Karabulut and A. Aysu, “Falcon down: Breaking falcon post-quantum signature scheme through side-channel attacks,” *Cryptology ePrint Archive*, Paper 2021/772, 2021.

- [42] National Institute of Standards and Technology, “FIPS 140-2: Security requirements for cryptographic modules,” tech. rep., U.S. Department of Commerce, 2001. Change Notice 2, 12/3/2002. Superseded by FIPS 140-3.
- [43] C. O’Flynn and Z. D. Chen, “Chipwhisperer: An open-source platform for hardware embedded security research,” in *Constructive Side-Channel Analysis and Secure Design*, pp. 243–260, Springer International Publishing, 2014.
- [44] P. Technology, “Picoscope 6000e series datasheet,” 2024. Available: <https://www.picotech.com/download/datasheets/picoscope-6000e-series-data-sheet.pdf>, Accessed: 2024-11-28.
- [45] P. Team, “Crystals-dilithium: Post-quantum digital signature scheme,” 2024. Available: <https://github.com/pq-crystals/dilithium>, Accessed: 2024-11-28.