

Side-channel Assisted Existential Forgery Attack on Dilithium - A NIST PQC candidate

Prasanna Ravi¹, Mahabir Prasad Jhanwar², James Howe³, Anupam Chattopadhyay⁴, and Shivam Bhasin¹

¹ Temasek Laboratories, Nanyang Technological University, Singapore

² Ashoka University, India

³ University of Bristol, UK

⁴ School of Computer Science and Engineering
Nanyang Technological University, Singapore

prasanna.ravi@ntu.edu.sg mahavir.jhavar@ashoka.edu.in
james.howe@bristol.ac.uk anupam@ntu.edu.sg sbhasin@ntu.edu.sg

Abstract. The recent lattice-based signature scheme Dilithium, submitted as part of the CRYSTALS (Cryptographic Suite for Algebraic Lattices) package, is one of a number of strong candidates submitted for the NIST standardization process of post-quantum cryptography. The Dilithium signature scheme is based on the Fiat-Shamir paradigm and can be seen as a variant of the Bai-Galbraith scheme (BG) combined with several improvements from previous ancestor lattice-based schemes like GLP and BLISS signature schemes. One of the main features of Dilithium is the *compressed* public-key, which is a rounded version of the LWE instance. This implies that Dilithium is not breakable with the knowledge of only the secret or the error of the LWE instance, unlike its ancestor lattice-based signature schemes. In this paper, we investigate the security of Dilithium against a combination of side-channel and classical attacks. Side-channel attacks on schoolbook and optimised polynomial multiplication algorithms in the signing procedure are shown to extract the secret component of the LWE instance, which is just one among the multiple components of the secret-key of Dilithium. We then propose an alternative signing procedure, through which it is possible to forge signatures with only the extracted portion of the secret-key, without requiring the knowledge of all its elements. Thus showing that Dilithium too breaks on just knowing the secret portion of the LWE instance, similar to previous lattice-based schemes.

1 Introduction

Recently, NIST has called for proposals for standardization of post-quantum cryptographic schemes for public-key encryption, digital signatures, and key establishment protocols [18]. This initiative is partly driven by the onset of the era of practical and scalable quantum computers [3, 15, 26], which has motivated the cryptographic community to develop cryptographic schemes that are immune

to cryptanalytic efforts using quantum algorithms. Amongst the initial candidates from different types of post-quantum cryptography, lattice-based cryptography has emerged as the largest category in terms of submission numbers. NIST [19] state that “schemes that can be made resistant to side-channel attack at minimal cost are more desirable than those whose performance is severely hampered by any attempt to resist side-channel attacks. Thus, it has become essential to the standardisation process to analyse the implementation security of these proposals against active and passive side-channel attacks.”

Over the years, significant research has shown that lattice-based cryptographic algorithms [13, 22] have the performance required to replace currently used classical cryptographic primitives such as RSA and ECC [14]. Most of the efficient lattice-based signature schemes [9, 10, 13] are based on the variant of the Fiat-Shamir with Aborts framework, first proposed by Lyubashevsky [16]. The BLISS signature scheme [9] based on the same framework attracted significant attention from the scientific community, owing to its security and efficiency guarantees. However, its physical security came under strong scrutiny, owing to the number of implementation attacks [6, 11, 12, 20, 21], which targetted its modules such as Gaussian sampling and rejection sampling.

These attacks have lead the cryptographic community to look beyond the BLISS signature scheme and recently Ducas et al. proposed Dilithium [10, 17], a lattice-based signature scheme submitted to NIST for post-quantum standardisation, which removes the Gaussian sampling attack vector by employing uniform sampling instead. The signature scheme derives its hardness guarantees from hard problems on modular lattices, which has a lesser exploitable structure compared to ideal lattices, which is used in BLISS. The scheme also derives certain properties from other previously proposed signature schemes [2, 9, 13] to improve upon the sizes of public-key and signature. Certain decisions on the algorithm and parameters have been taken to make sure that some implementation attacks reported on BLISS cannot be easily applied on the Dilithium signature scheme.

Most, if not all, lattice-based signature schemes partition their secret-keys in the form of two polynomials; \mathbf{s}_1 and \mathbf{s}_2 , and are used to form a LWE instance $\mathbf{t} = \mathbf{a}\mathbf{s}_1 + \mathbf{s}_2$, which is declared as the public key. A scheme utilising this format thus breaks down with knowledge of either of the secrets, since the LWE instance becomes trivially solvable. While this is the case with Dilithium’s ancestral schemes like BLISS, GLP and BG, Dilithium publishes only a rounded version of the LWE instance as its public key. Thus, unlike its ancestral schemes, Dilithium does not trivially break-down with the knowledge of the secret or the error. In this paper, we demonstrate a side-channel assisted existential forgery attack on the Dilithium signature scheme wherein we mainly bring to light the possibility of forging signatures with only the partial knowledge of the secret-key \mathbf{s}_1 . We first demonstrate a side-channel attack on the polynomial multiplication component, to retrieve the secret of the LWE instance, which is only one of the many components of its secret-key, and further propose an alternate signing procedure that utilizes only the extracted portion of the secret-key (\mathbf{s}_1) to produce valid

signatures, thus showing that Dilithium too breaks with the knowledge of only the secret of the LWE instance \mathbf{s}_1 .

1.1 Contribution

In this work, we propose a side-channel assisted existential forgery attack on the Dilithium signature scheme.

- The side-channel attack is designed to retrieve partial secret-key ($\bar{\mathbf{s}}_1$) through a power analysis attack on the polynomial multiplier in the signing procedure.
- We further exploit certain properties of the Dilithium signature scheme to show that one can forge signatures with only knowledge of the partial secret-key. This procedure can also be viewed as an alternative signing procedure which involves lesser computations compared to the original scheme, but accompanied with a certain failure probability.

The idea of the attack over the polynomial multiplier was proposed by Espitau et al. [12]. We modified the attack to suit the specifications of Dilithium and exploit the leakage in a different operation of the polynomial multiplication. We acknowledge another parallel work of Bruinderink and Pessl [7] which also proposes a forgery attack with the knowledge of the partial secret-key, however we provide a simpler alternate forgery scheme with a concrete analysis of the properties of the Dilithium scheme that facilitated our attack.

2 Preliminaries

Let $q \in \mathbb{N}$ be a prime. For an integer r and an even positive integer α , we define centered reduction modulo q , denote as $r \pmod{\pm \alpha}$, to be the unique integer r_0 such that, $r \equiv r_0 \pmod{\alpha}$ and $\frac{\alpha}{2} < r_0 \leq \frac{\alpha}{2}$. The usual modulo reduction is denoted by $r \pmod{q}$. We let \mathbb{Z} denote the ring of integers, and \mathbb{Z}_q the ring of integers modulo q . The elements in \mathbb{Z}, \mathbb{Z}_q are denoted by regular font letters, viz. $a, b \in \mathbb{Z}$ or \mathbb{Z}_q . We let R_q denote the ring $\mathbb{Z}_q[X]/\langle X^n + 1 \rangle$, $q, n \in \mathbb{N}$. Bold lower-case letters represents elements in R_q , viz. $\mathbf{a}, \mathbf{b} \in R_q$. Column vectors with coefficients in R_q are denoted by bold lower-case letters with an overhead bar, viz. $\bar{\mathbf{a}}, \bar{\mathbf{b}} \in R_q^\ell$, $\ell \in \mathbb{N}$. Note that, elements $\mathbf{a} \in R_q$ are polynomials of degree at most $n - 1$ with coefficients in \mathbb{Z}_q and therefore can be represented as n -length vectors (a_0, \dots, a_{n-1}) , $a_i \in \mathbb{Z}_q$. By default, all vectors will be column vectors. Matrices with coefficients in R_q are denoted by bold upper-case letters, viz. $\mathbf{A} \in R_q^{k \times \ell}$, $k, \ell \in \mathbb{N}$. Transpose of a vector $\bar{\mathbf{a}}$ is denoted by $\bar{\mathbf{a}}^\top$. For $q, n, \ell \in \mathbb{N}$, R_q^ℓ is an R module with scalar multiplication $R \times R_q^\ell \rightarrow R_q^\ell$ defined as follows: for $\mathbf{a} \in R$, $\bar{\mathbf{b}} = (\mathbf{b}_0, \dots, \mathbf{b}_{n-1})^\top \in R_q^\ell$, $\mathbf{a}\bar{\mathbf{b}} = (\mathbf{a}\mathbf{b}_0, \dots, \mathbf{a}\mathbf{b}_{n-1})^\top$, where $\mathbf{a}\mathbf{b}_i$ is polynomial multiplication defined in the ring R_q . A definition for the length of elements in R_q is as follows. For an element $\mathbf{a} = (a_0, \dots, a_{n-1})$, define $\|\mathbf{a}\|_\infty = \max_{0 \leq i \leq n-1} \|a_i\|_\infty$, where $\|a_i\|_\infty = |a_i \pmod{\pm q}|$. Similarly, for $\bar{\mathbf{a}} = (\mathbf{a}_1, \dots, \mathbf{a}_\ell)^\top \in R_q^\ell$, we define $\|\bar{\mathbf{a}}\|_\infty = \max_{1 \leq i \leq \ell} \|\mathbf{a}_i\|_\infty$. Finally, for a $\eta \in \mathbb{N}$,

define $S_\eta = \{\mathbf{a} \in R_q \mid \|\mathbf{a}\|_\infty \leq \eta\}$. For a set X , we write $x \stackrel{\$}{\leftarrow} X$ to denote that x is chosen uniformly at random from X , and $x \stackrel{D}{\leftarrow} X$ when $D : X \rightarrow [0, 1]$ is a specific (non-uniform) probability distribution.

2.1 Underlying Lattice Problems

Definition 1 (MLWE). *The module learning with error (MLWE) assumption states, roughly, that the distribution $(\mathbf{A}, \bar{\mathbf{t}})$ and $(\mathbf{A}, \mathbf{A}\bar{\mathbf{s}}_1 + \bar{\mathbf{s}}_2)$ are computationally indistinguishable when $\mathbf{A} \stackrel{\$}{\leftarrow} R_q^{k \times \ell}$ ($k, \ell \in \mathbb{N}$), $\bar{\mathbf{t}} \stackrel{\$}{\leftarrow} R_q^k$, $\bar{\mathbf{s}}_1 \stackrel{\$}{\leftarrow} S_\eta^\ell$, and $\bar{\mathbf{s}}_2 \stackrel{\$}{\leftarrow} S_\eta^k$ ($\eta \in \mathbb{N}$). In particular, for positive integers k, ℓ, η , the advantage of an adversary \mathcal{A} in solving a decisional $\text{MLWE}_{k, \ell, \eta}$ problem instance over the ring R_q is*

$$\begin{aligned} \text{Adv}^{\mathcal{A}}.\text{MLWE}_{k, \ell, \eta} &= |\mathbb{P}[\mathcal{A}(\mathbf{A}, \bar{\mathbf{t}}) = 1 \mid \mathbf{A} \stackrel{\$}{\leftarrow} R_q^{k \times \ell}; \bar{\mathbf{t}} \stackrel{\$}{\leftarrow} R_q^k] \\ &\quad - \mathbb{P}[\mathcal{A}(\mathbf{A}, \mathbf{A}\bar{\mathbf{s}}_1 + \bar{\mathbf{s}}_2) = 1 \mid \mathbf{A} \stackrel{\$}{\leftarrow} R_q^{k \times \ell}; \bar{\mathbf{s}}_1 \stackrel{\$}{\leftarrow} S_\eta^\ell; \bar{\mathbf{s}}_2 \stackrel{\$}{\leftarrow} S_\eta^k]| \end{aligned}$$

and the MLWE hardness assumption says that the advantage $\text{Adv}^{\mathcal{A}}.\text{MLWE}_{k, \ell, \eta}$ for any probabilistic polynomial time (PPT) adversary \mathcal{A} is negligible.

Definition 2 (MSIS). *The MSIS hardness assumptions states that given $(\mathbf{A}, \bar{\mathbf{t}}, \beta)$, where $\mathbf{A} \stackrel{\$}{\leftarrow} R_q^{k \times \ell}$, $\bar{\mathbf{t}} \in R_q^k$, and $\beta \in \mathbb{N}$, it is hard to compute an $\bar{\mathbf{x}} \in R_q^{\ell+k}$ such that $\|\bar{\mathbf{x}}\|_\infty \leq \beta$ and $[\mathbf{A} \mid \mathbf{I}]\bar{\mathbf{x}} = \bar{\mathbf{t}}$, where $\mathbf{I} = \mathbf{I}_{k \times k}$, k order identity matrix. In particular, for positive integers k, ℓ, β , the advantage of an adversary \mathcal{A} in solving a $\text{MSIS}_{k, \ell, \beta}$ problem instance over the ring R_q is*

$$\text{Adv}^{\mathcal{A}}.\text{MSIS}_{k, \ell, \beta} = \mathbb{P}[\|\bar{\mathbf{x}}\|_\infty \leq \beta \wedge [\mathbf{A} \mid \mathbf{I}]\bar{\mathbf{x}} = \bar{\mathbf{t}} \mid \mathbf{A} \stackrel{\$}{\leftarrow} R_q^{k \times \ell}; \bar{\mathbf{t}} \in R_q^k; \bar{\mathbf{x}} \leftarrow \mathcal{A}(\mathbf{A}, \bar{\mathbf{t}}, \beta)]$$

and the MSIS hardness assumption says that the advantage $\text{Adv}^{\mathcal{A}}.\text{MSIS}_{k, \ell, \beta}$ for any probabilistic polynomial time (PPT) adversary \mathcal{A} is negligible.

2.2 Rounding Algorithms

This section describes several rounding procedures used in Dilithium. The purpose of them is to provide better compression of Dilithium's public-key and signature. Figure 1 lists the full details of these algorithms. For non-negative integers r, α , the procedure $r \pmod{\pm \alpha}$ outputs unique integer r_0 in $-\alpha/2 < r_0 \leq \alpha/2$ such that $r \equiv r_0 \pmod{\alpha}$. The procedure $\text{D}_q(r, \alpha)$ decomposes r into a pair of integers (r_1, r_0) . The procedures $\text{HB}_q(r, \alpha)$ and $\text{LB}_q(r, \alpha)$ extracts r_1 and r_0 respectively from $\text{D}_q(r, \alpha)$. The procedure $\text{MH}_q(u, r, \alpha)$ produces a bit $h \in \{0, 1\}$. The procedure $\text{UH}_q(h, r, \alpha)$ shows how to use h as a hint to derive $\text{HB}_q(u + r, \alpha)$ when z is not known and its correctness is stated in the following lemma.

Lemma 1 ([17]). *For $r, z \in \mathbb{Z}_q$ with $\|z\|_\infty \leq \alpha/2$, we have*

$$\text{UseHint}_q(\text{MakeHint}_q(z, r, \alpha), r, \alpha) = \text{HighBits}_q(r + z, \alpha) \quad (1)$$

<pre> 1 Function CenteredModulo $r \pmod{\pm \alpha}$: 2 Compute $r_0 = r \pmod{\alpha}$ 3 if $r_0 > \alpha/2$ then 4 $r_0 = r_0 - \alpha$ 5 end 6 return $r \pmod{\pm \alpha} = r_0$ 7 Function Decompose $D_q(r, \alpha)$: 8 Compute $r = r \pmod{q}$ 9 Compute $r_0 = r \pmod{\pm \alpha}$ 10 if $r - r_0 = q - 1$ then 11 $r_1 = 0$ 12 $r_0 = r_0 - 1$ 13 else 14 $r_1 = (r - r_0)/\alpha$ 15 end 16 return $D_q(r, \alpha) = (r_1, r_0)$ 17 Function HighBits $HB_q(r, \alpha)$: 18 Compute $(r_1, r_0) = D_q(r, \alpha)$ 19 return $HB_q(r, \alpha) = r_1$ 20 Function LowBits $LB_q(r, \alpha)$: 21 Compute $(r_1, r_0) = D_q(r, \alpha)$ 22 return $LB_q(r, \alpha) = r_0$ </pre>	<pre> 1 Function MakeHint $MH_q(u, r, \alpha)$: 2 Compute $r_1 = HB_q(r, \alpha)$ 3 Compute $v_1 = HB_q(u + r, \alpha)$ 4 if $r_1 = v_1$ then 5 $h = 0$ 6 else 7 $h = 1$ 8 end 9 return $MH_q(u, r, \alpha) = h$ 10 Function UseHint $UH_q(h, r, \alpha)$: 11 Compute $m = (q - 1)/\alpha$ 12 Compute $(r_1, r_0) = D_q(r, \alpha)$ 13 if $h = 1$ and $r_0 > 0$ then 14 $r_1 = (r_1 + 1) \pmod{m}$ 15 end 16 else if $h = 1$ and $r_0 \leq 0$ then 17 $r_1 = (r_1 - 1) \pmod{+m}$ 18 else 19 $r_1 = r_1$ 20 end 21 return $UH_q(h, r, \alpha) = r_1$ </pre>
--	--

Fig. 1: Rounding Algorithms

All these procedures extend naturally to ring elements in R_q . For example, if $\mathbf{r} = (r_0, \dots, r_{n-1}) \in R_q$ then $D_q(\mathbf{r}, \alpha) = (D_q(r_0), \dots, D_q(r_{n-1}))$. Finally the procedures are extended to vectors in R_q^ℓ and matrices in $R_q^{k \times \ell}$: for $\bar{\mathbf{a}} = (\mathbf{a}_1, \dots, \mathbf{a}_\ell)^\top \in R_q^\ell$, $HB_q(\bar{\mathbf{a}}, \alpha) = (HB_q(\mathbf{a}_1, \alpha), \dots, HB_q(\mathbf{a}_\ell, \alpha))$.

2.3 Dilithium

In the following, we recall the Dilithium signature scheme [10, 17]. The underlying approach of the scheme is based on the ‘‘Fiat-Shamir with Aborts’’ approach [16] and is an improved variant of the scheme proposed by Bai and Galbraith [2]. Dilithium depends on parameters $q, n, k, \ell, \eta, d, \gamma_1, \gamma_2, \beta, \omega$. Various constraints and recommended values for these parameters are provided in the specifications [17]. While reading the scheme below, the reader may keep in mind the following set of recommended parameters: ($q = 8380417 = 2^{23} - 2^{13} + 1, n = 256, k = 3, \ell = 2, \eta = 7, d = 14, \gamma_1 = (q - 1)/16, \gamma_2 = \gamma_1/2, \beta = 375, \omega = 64$). In addition, the scheme also uses:

- $H : \{0, 1\}^* \rightarrow B_h$, a hash function, where B_h denote the set of elements of R_q that have h coefficients that are either -1 or 1 and the rest are 0 ;

- $\text{ExpandA} : \{0, 1\}^{256} \rightarrow \{\mathbf{A} \mid \mathbf{A} \in R_q^{k \times \ell}\}$;
- $\text{ExpandMask} : \{0, 1\}^{256} \times \{0, 1\}^* \times \mathbb{N} \rightarrow S_{\gamma_1-1}^\ell$; and
- a collision resistant hash function $\text{CRH} : \{0, 1\}^* \rightarrow \{0, 1\}^{384}$.

For complete details of these function please refer to [17]. The key generation, signing and verification algorithms for Dilithium are presented in Figure 2.

- **Key Generation:** The key generation algorithm, KeyGen , begins by choosing $\rho, K \xleftarrow{\$} \{0, 1\}^{256}$. It then computes a matrix $\mathbf{A} \in R_q^{k \times \ell} = \text{ExpandA}(\rho)$. Next, it samples secret-key vectors $\bar{\mathbf{s}}_1 \in S_\eta^\ell$, $\bar{\mathbf{s}}_2 \in S_\eta^k$. It then computes $\bar{\mathbf{t}} = \mathbf{A}\bar{\mathbf{s}}_1 + \bar{\mathbf{s}}_2$ and decomposes it into $\text{HB}_q(\bar{\mathbf{t}}, 2^d) = \bar{\mathbf{t}}_1$ and $\text{LB}_q(\bar{\mathbf{t}}, 2^d) = \bar{\mathbf{t}}_0$. Finally, $\bar{\mathbf{s}}_1, \bar{\mathbf{s}}_2, \bar{\mathbf{t}}_0, K$ are set as part of the secret-key; and $\bar{\mathbf{t}}_1$ is set as part of public-key.
- **Signing:** The signing algorithm, Sign , proceeds by generating a masking vector $\bar{\mathbf{y}} \in S_{\gamma_1-1}^\ell$ using ExpandMask function. It then computes $\bar{\mathbf{w}} = \mathbf{A}\bar{\mathbf{y}}$ and sets $\bar{\mathbf{w}}_1$ to be the $\text{HB}_q(\bar{\mathbf{w}}, 2\gamma_2)$. The challenge, \mathbf{c} , is created as the hash of μ and $\bar{\mathbf{w}}_1$. The potential signature is then computed as $\bar{\mathbf{z}} = \bar{\mathbf{y}} + \mathbf{c}\bar{\mathbf{s}}_1$. To avoid the dependency of $\bar{\mathbf{z}}$ on the secret-key $\bar{\mathbf{s}}_1$, rejection sampling on $\bar{\mathbf{z}}$ is used unless $\|\bar{\mathbf{z}}\|_\infty < \gamma_1 - \beta$ (the probability of the later event is high as β is so chosen that $\|\mathbf{c}\bar{\mathbf{s}}_2\|_\infty \leq \beta$). The rejection sampling on $\bar{\mathbf{z}}$ also checks if $\text{LB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2) < \gamma_2 - \beta$ (note that, $\mathbf{A}\bar{\mathbf{z}} - \mathbf{c}\bar{\mathbf{t}} = \bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2$), this check is to ensure both security and correctness. In order to further proceed with the sampled $\bar{\mathbf{z}}$ it finally checks if $\text{HB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2) = \bar{\mathbf{w}}_1$. At this moment, the tuple $(\bar{\mathbf{z}}, \mathbf{c})$, together with $\bar{\mathbf{t}}$, can be used to verify if $c \stackrel{?}{=} H(\mu \parallel \text{HB}_q(\mathbf{A}\bar{\mathbf{z}} - \mathbf{c}\bar{\mathbf{t}}, 2\gamma_2))$. But the public-key of the verifier only has $\bar{\mathbf{t}}_1 = \text{HB}_q(\bar{\mathbf{t}}, 2^d)$. To overcome this problem, the signing algorithm provides a hint vector $\bar{\mathbf{h}}$ (in fact, a binary matrix) which will allow the verifier to compute $\bar{\mathbf{w}}_1$. This is done in Step 23 (Figure 2) with some other additional checks carried out in Step 24.
- **Verification:** The verifier computes $\bar{\mathbf{w}}'_1 = \text{UH}_q(\bar{\mathbf{h}}, \mathbf{A}\bar{\mathbf{z}} - \mathbf{c}\bar{\mathbf{t}}_1 \cdot 2^d, 2\gamma_2)$, and then accepts if $\|\bar{\mathbf{z}}\|_\infty < \gamma_1 - \beta$, $\mathbf{c} = H(\text{CRH}(\text{CRH}(\rho \parallel \bar{\mathbf{t}}_1) \parallel M), \bar{\mathbf{w}}'_1)$, and $\text{wt}(\bar{\mathbf{h}}) \leq \omega$. Let us look at why verification works, in particular as to why $\bar{\mathbf{w}}_1 = \text{UH}_q(\bar{\mathbf{h}}, \mathbf{A}\bar{\mathbf{z}} - \mathbf{c}\bar{\mathbf{t}}_1 \cdot 2^d, 2\gamma_2)$. Notice that $\mathbf{A}\bar{\mathbf{z}} - \mathbf{c}\bar{\mathbf{t}}_1 \cdot 2^d = \bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2 + \mathbf{c}\bar{\mathbf{t}}_0$. Therefore,

$$\begin{aligned}
\bar{\mathbf{w}}_1 &= \text{HB}_q(\bar{\mathbf{w}}, 2\gamma_2) \\
&\stackrel{*}{=} \text{HB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2) \\
&= \text{HB}_q((-\mathbf{c}\bar{\mathbf{t}}_0) + (\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2 + \mathbf{c}\bar{\mathbf{t}}_0), 2\gamma_2) \\
&\stackrel{**}{=} \text{UH}_q(\text{MH}_q(-\mathbf{c}\bar{\mathbf{t}}_0, \bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2 + \mathbf{c}\bar{\mathbf{t}}_0, 2\gamma_2), \bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2 + \mathbf{c}\bar{\mathbf{t}}_0, 2\gamma_2) \\
&= \text{UH}_q(\bar{\mathbf{h}}, \mathbf{A}\bar{\mathbf{z}} - \mathbf{c}\bar{\mathbf{t}}_1 \cdot 2^d, 2\gamma_2).
\end{aligned}$$

The * equality is true as follows: $\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2 = \bar{\mathbf{r}}_1 \times 2\gamma_2 + \text{LB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2)$ such that $\|\text{LB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2)\|_\infty \leq \gamma_2 - \beta$ (ensured during signing). This implies, $\bar{\mathbf{w}} = (\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2) + \mathbf{c}\bar{\mathbf{s}}_2 = \bar{\mathbf{r}}_1 \times 2\gamma_2 + \text{LB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2) + \mathbf{c}\bar{\mathbf{s}}_2$. But, as $\|\mathbf{c}\bar{\mathbf{s}}_2\|_\infty \leq \beta$, $\|\text{LB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2) + \mathbf{c}\bar{\mathbf{s}}_2\|_\infty \leq \|\text{LB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2)\|_\infty + \|\mathbf{c}\bar{\mathbf{s}}_2\|_\infty \leq \gamma_2 - \beta + \beta =$

γ_2 . Thus, from the definition of centered reduction modulo $2\gamma_2$, it follows that $\text{HB}_q(\bar{\mathbf{w}} - \mathbf{c}\bar{\mathbf{s}}_2, 2\gamma_2) = \text{HB}_q(\bar{\mathbf{w}}, 2\gamma_2) = \bar{\mathbf{w}}_1$. The ** equality follows from Lemma 1, as $\| -\mathbf{c}\bar{\mathbf{t}}_0 \|_\infty \leq 2^d \leq \gamma_2$.

2.4 Security of Dilithium

The security of Dilithium is derived from MLWE and MSIS problems. In particular, security against key-recovery attack under classical random oracle model is based on the hardness assumption of the MLWE problem; and the security against existential signature forgery is due to MSIS hardness assumption. The scheme’s security against strong signature forgery attack, under quantum random oracle is also discussed by the authors [17].

2.5 Side-Channel Attacks on Lattice-Based Cryptography

Lattice-based cryptography has gained significant attention from the perspective of SCA and fault attacks. Especially, predecesing lattice-based signature schemes like GLP [13], BLISS [9], and Ring-Tesla [1], which have been scrutinized through a number of active and passive implementation attacks, such as power analysis [20, 27], fault attacks [5, 11], branch-tracing [12], and cache-timing [6, 21]. These signature schemes opted to sample their error vectors from a Gaussian distribution and used rejection sampling to hide the information about the secret-key in the signature. Most of the side channel analysis research targetted the data dependent side-channel leakage from these Gaussian sampling and the rejection sampling components.

The attacks on the Gaussian samplers and rejection sampling techniques prompted the designers of the newer signature schemes, like Dilithium and qTESLA, to instead sample from uniform distributions during signing. This both simplifies the rejection step in addition to simplifying the sampling procedure, albeit with the cost of increased size of the signatures.

Lattice-based signature schemes have also been the target of a number of fault attacks [5, 7, 11]. Bindel et al. [5] proposed a variety of randomization, zeroing, and skipping fault attacks that lead to either key recovery or forging of signatures, applicable to BLISS, GLP, and Ring-TESLA. This was followed by a generic and much stronger fault attack on lattice-based signature schemes by Espitau *et al.* [11], based on loop abort faults. They were able to show that the signing step $\bar{\mathbf{z}} = \bar{\mathbf{s}}_1 \mathbf{c} + \bar{\mathbf{y}}$, can be converted into a solvable CVP instance, provided the masking polynomials are limited to low degrees using loop abort faults. The first differential style fault attacks on Dilithium and qTESLA, signatures that are potential NIST standards, have been reported by Bruinderink and Pessl [7], which targets the generation of the $\bar{\mathbf{z}}$ component of the signature and work with a fault model of injecting random faults using clock glitches on an ARM Cortex-M4F based microcontroller.

```

1 Function KeyGen( $q, n, k, \ell, \eta, d, \gamma_1, \gamma_2, \beta, \omega$ ):
2    $\rho \leftarrow \{0, 1\}^{256}$ 
3    $K \leftarrow \{0, 1\}^{256}$ 
4    $A \in R_q^{k \times \ell} := \text{ExpandA}(\rho)$ 
5    $(\bar{s}_1, \bar{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k$ 
6    $\bar{t} = A\bar{s}_1 + \bar{s}_2$ 
7    $(\bar{t}_1, \bar{t}_0) = D_q(\bar{t}, 2^d)$ 
8    $\text{tr} = \text{CRH}(\rho \| \bar{t}_1)$ 
   Return :  $\text{pk} = (\bar{t}_1, \rho, n, q, k, d, \ell, \gamma_1, \gamma_2, \beta, \omega), \text{sk} = (K, \text{tr}, \bar{s}_1, \bar{s}_2, \bar{t}_0, \text{pk})$ 

9 Function Sign( $\text{sk}, M$ ):
10   $A \in R_q^{k \times \ell} := \text{ExpandA}(\rho)$ 
11   $\mu = \text{CRH}(\text{tr} \| M)$ 
12   $\kappa = 0, (\bar{z}, \bar{h}) = \perp$ 
13  while  $(\bar{z}, \bar{h}) = \perp$  do
14     $\bar{y} \in S_{\gamma_1-1}^\ell := \text{ExpandMask}(K \| \mu \| \kappa)$ 
15     $\bar{w} = A\bar{y}$ 
16     $\bar{w}_1 = \text{HB}_q(\bar{w}, 2\gamma_2)$ 
17     $c \in B_{60} = H(\mu \| \bar{w}_1)$ 
18     $\bar{z} = \bar{y} + c\bar{s}_1$ 
19     $(\bar{r}_1, \bar{r}_0) := D_q(\bar{w} - c\bar{s}_2, 2\gamma_2)$ 
20    if  $\|\bar{z}\|_\infty \geq \gamma_1 - \beta$  or  $\|\bar{r}_0\|_\infty \geq \gamma_2 - \beta$  or  $\bar{r}_1 \neq \bar{w}_1$  then
21       $(\bar{z}, \bar{h}) = \perp$ 
22    else
23       $\bar{h} = \text{MH}_q(-c\bar{t}_0, \bar{w} - c\bar{s}_2 + c\bar{t}_0, 2\gamma_2)$ 
24      if  $\|c\bar{t}_0\|_\infty \geq \gamma_2$  or  $\text{wt}(\bar{h}) > \omega$  then
25         $(\bar{z}, \bar{h}) = \perp$ 
26      end
27     $\kappa = \kappa + 1$ 
28  end
   Return :  $\sigma = (\bar{z}, \bar{h}, c)$ 

29 Function Verify( $\text{pk}, M, \sigma = (\bar{z}, \bar{h}, c)$ ):
30   $A \in R_q^{k \times \ell} := \text{ExpandA}(\rho)$ 
31   $\mu = \text{CRH}(\text{CRH}(\rho \| \bar{t}_1) \| M)$ 
32   $\bar{w}_1 := \text{UH}_q(\bar{h}, A\bar{z} - c\bar{t}_1 \cdot 2^d, 2\gamma_2)$ 
33  if  $c = H(\mu, \bar{w}_1)$  and  $\|\bar{z}\|_\infty < \gamma_1 - \beta$  and  $\text{wt}(\bar{h}) \leq \omega$  then
34    return 1
35  else
36    return 0
37  end

```

Fig. 2: Dilithium Signature Scheme

3 Motivation

In this section, we motivate the idea of existential forgery attacks with only knowledge of the partial secret-key. We first investigate the key features that distinguish Dilithium from some of the previous practical lattice-based signature schemes like BLISS [9], GLP [13], and Ring-TESLA [1] signature schemes. We provide some intuition which lead to the idea of signature forgery using partial knowledge of the secret-key. Additionally, we also highlight the importance of the signature forgery technique from the perspective of a side-channel adversary.

3.1 A few observations on Dilithium

Although most of the features in Dilithium are similar to those in its predecessors, like GLP or BLISS, one stand out property that differentiates it is the public-key being a compressed version of the LWE instance. The reason for this is to compress the size of the public-key. As noted earlier, only the d higher-order bits (i.e., 14 out of the 23 bits) of every coefficient of $\bar{\mathbf{t}}$ are revealed as the public-key. Alternatively, this can be seen as an additional error of about $(\pm 2^{14})$ to every coefficient, thus making it a LWE instance with an additional error. With such an instance at our disposal, even retrieval of the complete $\bar{\mathbf{s}}_1$ does not result in a straight-forward retrieval of the error $\bar{\mathbf{s}}_2$. But, previous schemes like GLP, BLISS, and BG signature schemes would be completely broken with the knowledge of either the secret or the error of the LWE instance. Thus, preliminary observations suggest that the *compressed* public key offers increased protection against signature forgery though the authors do not claim any additional security due to the compression of the public-key [17].

Another difference is that the $\bar{\mathbf{z}}_2$ component of the signature, that directly depends on the value of $\bar{\mathbf{s}}_2$, previously present in schemes like GLP and BLISS, has been completely removed owing to the effort to decrease the signature size, an idea derived from the BG signature scheme. Thus, Dilithium only consists of three components, $(\bar{\mathbf{z}}, \mathbf{c})$ and the additional hint vector $\bar{\mathbf{h}}$. One can generate the signature components $\bar{\mathbf{z}}$ and \mathbf{c} with only the knowledge of $\bar{\mathbf{s}}_1$. As stated before, preliminary observations suggest that the $\bar{\mathbf{s}}_2$ component is required to generate the hint vectors. But, another parallel observation of the signing and verification procedures suggest that both the signer and verifier only commit over the higher-order bits of a particular intermediate value $(\bar{\mathbf{w}}_1)$. With $\bar{\mathbf{s}}_2$ being small and only used in the product of $\mathbf{c} \cdot \bar{\mathbf{s}}_2$, which is again a module with very small coefficients, it leads us to think if $\bar{\mathbf{s}}_2$ will have any observable impact over the committed value $(\bar{\mathbf{w}}_1)$, which also naturally leads us to think if it is possible to bypass the use of $\bar{\mathbf{s}}_2$ to forge signatures. A concrete analysis of this idea is followed in Section 4, where we show that it is indeed possible to forge signatures with only the knowledge of $\bar{\mathbf{s}}_1$.

3.2 Analysis from the perspective of a side-channel adversary

Analysis of previous works on implementation security of lattice-based signatures [5, 7, 11, 12] suggest that most of the attacks only concentrated on recovering

\bar{s}_1 , as the retrieval of \bar{s}_2 from the completely available LWE instance directly followed. Thus, the step of calculating the \bar{z} component of the signature, which involves the only computation using \bar{s}_1 in the signing procedure, has been greatly scrutinized from a SCA and fault attack perspective to a great extent. A similar attack can also be mounted on Dilithium to retrieve \bar{s}_1 (Step 19 of signing in Figure 2) provided the implementation is not sufficiently hardened against the previously known attacks. Upon recovering \bar{s}_1 , there are two possible methods that an adversary can follow to forge signatures.

The first technique is to try to recover the other secret \bar{s}_2 using statistical analysis. With the knowledge of \bar{s}_1 , the adversary can trivially calculate the \bar{w} component for all the valid signatures. Since we also know that all the signatures output adhere to the condition $\text{HB}_q(\bar{w} - \mathbf{c}\bar{s}_2, 2\gamma_2) \leq \gamma_2 - \beta$, it is possible to derive certain conditions of the possible values of \bar{s}_2 . But as stated by Bruinderink and Pessl [7], it would require a large number of signatures and a very high computational effort, thus increasing the attacker’s complexity. Additionally, an adversary cannot use this technique in an ephemeral key setting as they cannot expect to observe suitably many signatures to be generated using the same secret-key.

Additionally, the authors of Dilithium do not assume any secrecy with respect to the lower order bits of $\bar{\mathbf{t}}$ (i.e, $\bar{\mathbf{t}}_0$). In fact, the security proof of Dilithium is sketched under the assumption that the whole of $\bar{\mathbf{t}}$ is known to the adversary. Thus, it might indeed be possible that the whole of $\bar{\mathbf{t}}$ leaks as part of the signature and observations of sufficiently many signatures might lead to the recovery of the complete LWE instance, $\bar{\mathbf{t}}$. But again, we expect the number of signatures and the computational effort to be very high, which cannot be expected in a practical SCA setting.

The other technique the SCA adversary can attempt is to recover the other secret, \bar{s}_2 or $\bar{\mathbf{t}}_0$, through side channels. There have been previous works on SCA and fault attacks over the polynomial multiplication operation $\mathbf{c}\bar{s}_i$, for $i = \{1, 2\}$ [12, 27]. Since both the products are computed during the signing procedure, it is possible to mount similar SCA attacks to recover both \bar{s}_1 and \bar{s}_2 and also $\bar{\mathbf{t}}_0$.

But, in a practical SCA setting, it might not be always possible to retrieve the secrets completely and might pose a significant challenge owing to the requirement for the attacker to either analyse a large number of observations [12] for a given secret-key or inject a large number of successful faults [5]. Thus, an adversary might have to brute-force over the remaining coefficients of the attacked component (\bar{s}_2 or $\bar{\mathbf{t}}_0$) to completely recover the key.

We analysed this case, with the recommended parameter set of Dilithium, to estimate the brute-force complexity. Figure 3 shows the brute-force complexity of finding the remaining coefficients of both the secrets, \bar{s}_1 and \bar{s}_2 plotted against the success rates in retrieving coefficients. It shows that the attack’s complexity increases substantially, compared to when brute-forcing just one of the secrets. Thus, we can see that even for a success rate in retrieving the coefficients that is as high as 98%, the brute-force complexity is well over the 128 bit security level.

Even brute-forcing over one of the secrets with very high success rates of over 97% also crosses the 128-bit security level. Thus, in a practical setting, trying to retrieve both \bar{s}_1 and \bar{s}_2 might not be feasible even if success rate is very high and only very small percentage of coefficients are left unrecovered.

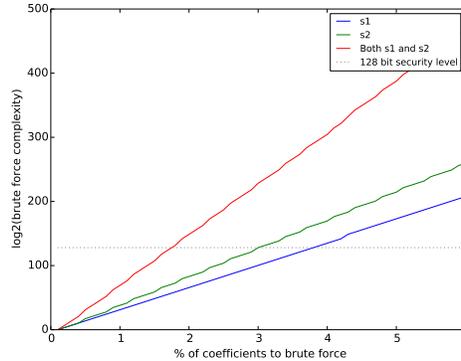


Fig. 3: Brute-force complexity versus the % of coefficients yet to be found.

Alternatively, the adversary can try to retrieve the remaining coefficients of \bar{s}_1 using algorithms that solve lattice problems. A similar technique was proposed by Bindel et al. [5], who propose a hybrid approach to reduce the number of fault attacks to retrieve a certain number of coefficients, leaving the remaining coefficients to be retrieved using LWE solvers. This technique reduces the attacker’s complexity in retrieving the whole of \bar{s}_1 . But, even if the whole of \bar{s}_1 is retrieved, the remaining coefficients of \bar{s}_2 still remain unknown. Referring to Figure 3, we again see that brute-forcing over remaining coefficients of one of the secrets requires a very high computation cost, even with a very high success rate.

Given all the above scenarios from an SCA perspective, we try to motivate the importance of our forgery signature scheme which requires only the partial knowledge of the secret-key. We do acknowledge another very recent parallel work by Bruinderink and Pessl [7], which proposes a different strategy to forge signatures in Dilithium with only the knowledge of \bar{s}_1 . However, the research lacks a clear analysis of the underlying factors that enabled signature forgery, with only the knowledge of \bar{s}_1 . However, we propose an alternate simpler forgery signature scheme backed up by a careful analysis along with a novel EM-based side channel attack strategy to recover \bar{s}_1 .

4 Existential Forgery Assuming Partial Secret-Key Knowledge

We now present an existential forgery attack on Dilithium assuming the knowledge of partial secret-key, \bar{s}_1 . For the following attack, refer to the function `Sign` in Figure 2. We consider an adversary \mathcal{A} , who has knowledge of \bar{s}_1 , and whose goal is to produce a valid signature of a new message. Execution of Lines 10 through 18 require knowledge of K (Line 14) and \bar{s}_1 (Line 18) the partial secret-key. The security of Dilithium is unaffected if K is made public, indeed K is only used to deterministically generate randomness \bar{y} . Thus, \mathcal{A} proceeds through Lines 10 to 18, choosing \bar{y} uniformly at random from $S_{\gamma_1-1}^\ell$ in Line 14 and computing \bar{z} in Line 18 using \bar{s}_1 .

Signature verification requires knowledge of \bar{w}_1 . In the specifications [17], it is proven that the $\mathbb{P}[\bar{w}_1 = \text{HB}_q(\bar{w} - \mathbf{c}\bar{s}_2, 2\gamma_2)] \approx e^{-n\beta(\frac{\ell}{\gamma_1} + \frac{k}{\gamma_2})}$, very close to 1. Now $\bar{w} - \mathbf{c}\bar{s}_2 = (-\mathbf{c}\bar{t}_0) + (\bar{w} - \mathbf{c}\bar{s}_2 + \mathbf{c}\bar{t}_0)$. Writing $\bar{u} = -\mathbf{c}\bar{t}_0$, $\bar{r} = \bar{w} - \mathbf{c}\bar{s}_2 + \mathbf{c}\bar{t}_0$, and the fact that $\mathbb{P}[\|\bar{u}\|_\infty \leq \gamma_2] \approx 1$ (indeed, $\|-\mathbf{c}\bar{t}_0\|_\infty \leq \|\bar{t}_0\|_\infty < 2^d < \gamma_2$) together ensure (Lemma 1) that \bar{w}_1 can be computed as follows:

$$\begin{aligned} \text{UH}_q(\text{MH}_q(\bar{u}, \bar{r}, 2\gamma_2), \bar{r}, 2\gamma_2) &= \text{HB}_q(\bar{u} + \bar{r}, 2\gamma_2) \\ &= \text{HB}_q(\bar{w} - \mathbf{c}\bar{s}_2, 2\gamma_2) \\ &= \bar{w}_1. \end{aligned}$$

But computation of $\text{MH}_q(\bar{u}, \bar{r}, 2\gamma_2)$ requires knowledge of $\bar{u} = -\mathbf{c}\bar{t}_0$, and therefore subsequently the partial secret-key \bar{t}_0 . *The problem for \mathcal{A} is now to compute hint matrix $\bar{h} = \text{MH}_q(\bar{u}, \bar{r}, 2\gamma_2)$ without the knowledge of \bar{t}_0 .* In the following we first show that UH_q function could be inverted to produce the correct hint. In particular $\text{UH}_q(h, r, \alpha)$, given $q, \alpha, r \in \mathbb{Z}_q$ and $\text{HB}_q(u + r, \alpha)$, is invertible on h if $\|u\|_\infty \leq \alpha/2$. Lemma 2 summarizes this fact.

Lemma 2. *Let $u \in \mathbb{Z}_q$ with $\|u\|_\infty \leq \alpha/2$, where $\alpha > 0$. Then, given any $r \in \mathbb{Z}_q$ and $\text{HB}_q(u + r, \alpha)$, it is easy to compute $\text{MH}_q(u, r, \alpha)$.*

Proof: The algorithm in Figure 4, given below, illustrates this claim. The correctness of the algorithm follows immediately from Lemma 1 and Lemma 3.

Lemma 3 ([17]). *Let $r \in \mathbb{Z}_q$ and $h, h' \in \{0, 1\}$. If $\text{UH}_q(h, r, \alpha) = \text{UH}_q(h', r, \alpha)$, then $h = h'$.*

Thus, in order to compute \bar{h} , \mathcal{A} must ensure that it has access to $\text{HB}_q(\bar{u} + \bar{r}, 2\gamma_2) = \text{HB}_q(\bar{w} - \mathbf{c}\bar{s}_2)$, and $\bar{r} = \bar{w} - \mathbf{c}\bar{s}_2 + \mathbf{c}\bar{t}_0$. Computing \bar{r} is easy as $\bar{w} - \mathbf{c}\bar{s}_2 + \mathbf{c}\bar{t}_0 = \mathbf{A}\bar{z} - \mathbf{c}\bar{t}_1 2^d$. To compute $\text{HB}_q(\bar{w} - \mathbf{c}\bar{s}_2)$ all \mathcal{A} has access to is \bar{w}_1 . But it is known that $\text{HB}_q(\bar{w} - \mathbf{c}\bar{s}_2, 2\gamma_2) = \bar{w}_1$ provided $\|\text{LB}_q(\bar{w} - \mathbf{c}\bar{s}_2, 2\gamma_2)\|_\infty \leq \gamma_2 - \beta$. A valid signer could ensure this check, but for \mathcal{A} it is not possible as it requires knowledge of \bar{s}_2 . But, this is not a problem as it is proven in [17] that, for a $\bar{y} \xleftarrow{\$} S_{\gamma_1}^\ell$, $\mathbb{P}[\|\text{LB}_q(\bar{w} - \mathbf{c}\bar{s}_2, 2\gamma_2)\|_\infty \leq \gamma_2 - \beta] \approx e^{-\frac{n\beta k}{\gamma_2}}$, very close to 1. We now finally present complete details of our existential forgery attack in Figure 5.

```

input :  $q, r, \alpha, \theta = \text{HB}_q(u + r, \alpha)$ 
output :  $\text{MH}_q(u, r, \alpha)$ 
1 Set  $h = 0$ ;
2 Compute  $\phi = \text{UH}_q(h, r, \alpha)$ ;
3 if  $\phi = \theta$  then
4 |   return  $h$ 
5 else
6 |   return 1
7 end

```

Fig. 4: Inverting UH_q for the hint bit h .

It is important to note that the attack works for the recommended parameter sets of Dilithium.

We implemented our forgery signing procedure by modifying the reference implementation of Dilithium that was submitted to the NIST standardization process. The results were obtained on an Intel Core-i5 (Haswell) processor running at 2.6 GHz with Turbo Boost and Hyperthreading and compiled with gcc-4.2.1 without modifying the compiler flags set for the reference implementation. We obtained an average signing time of about 0.3253 msec for our forgery signing procedure which is about 2.67 times faster than the signing time for the original signing procedure of Dilithium which runs at an average time of 0.8689 msec. The improved speed is due to the reduced number of operations and also the removal of rejection conditions over the norms of \bar{r}_0 and \bar{ct}_0 . We have uploaded the C code of our forgery signing procedure in the github link mentioned below⁵.

It is natural for our forgery scheme to produce invalid signatures since the conditional checks on the norm of \bar{r}_0 and \bar{ct}_0 are skipped. We attempted to empirically compute the failure probability of our forgery signing procedure. We could run our signature scheme for a total of 2^{28} times on the same platform with all the signatures verified correctly. This along with its increased signing rate leads us to hypothesize if our forgery signing procedure can be used as an alternative signing procedure with improved performance, provided it does not leak any information about the secret key. Concrete analysis of the same is left for future work.

4.1 Implications of the forgery signature scheme

Thus, we have shown that it is indeed possible to forge signatures, with only the knowledge of \bar{s}_1 , by successfully bypassing the use of the other elements of the secret-key. We do not claim any major break of the Dilithium signature scheme, but show that the knowledge of just one component, \bar{s}_1 , of the secret-key is enough to forge signatures which will be verified correctly with very high probability. But, it is important to state that the secret-key components; \bar{s}_1, \bar{s}_2 ,

⁵ https://github.com/jameshoweee/dilithium_forgery

```

input : public-key  $pk = (q, \rho, \bar{\ell}_1)$ , Partial secret-key =  $s_1$ , A message  $M$ 
output : A forged Dilithium signature

1  $A \sim R_q^{k \times \ell} := \text{Sam}(\rho)$ ;
2  $\mu = H(H(\rho \| \bar{\ell}_1) \| M)$ ;
3  $\bar{y} \xleftarrow{\$} S_{\gamma_1 - 1}^\ell$ ;
4  $\bar{w} = A\bar{y}$ ;
5  $\bar{w}_1 = \text{HB}_q(\bar{w}, 2\gamma_2)$ ;
6  $c = H(\mu, \bar{w}_1)$ ;
7  $\bar{z} = \bar{y} + c\bar{s}_1$ ;

8  $\bar{h} = \begin{bmatrix} h_{10} & \dots & h_{1(n-1)} \\ h_{20} & \dots & h_{2(n-1)} \\ \vdots & \vdots & \vdots \\ h_{k0} & \dots & h_{k(n-1)} \end{bmatrix} = \begin{bmatrix} 0 & \dots & 0 \\ 0 & \dots & 0 \\ \vdots & \vdots & \vdots \\ 0 & \dots & 0 \end{bmatrix}$ ;

9  $\theta = \begin{bmatrix} \theta_{10} & \dots & \theta_{1(n-1)} \\ \theta_{20} & \dots & \theta_{2(n-1)} \\ \vdots & \vdots & \vdots \\ \theta_{k0} & \dots & \theta_{k(n-1)} \end{bmatrix} = \begin{bmatrix} \text{UH}_q(h_{10}, [A\bar{z} - c\bar{\ell}_1 2^d]_{10}, 2\gamma_2) \dots \text{UH}_q(h_{1(n-1)}, [A\bar{z} - c\bar{\ell}_1 2^d]_{1(n-1)}, 2\gamma_2) \\ \text{UH}_q(h_{20}, [A\bar{z} - c\bar{\ell}_1 2^d]_{20}, 2\gamma_2) \dots \text{UH}_q(h_{2(n-1)}, [A\bar{z} - c\bar{\ell}_1 2^d]_{2(n-1)}, 2\gamma_2) \\ \vdots \\ \text{UH}_q(h_{k0}, [A\bar{z} - c\bar{\ell}_1 2^d]_{k0}, 2\gamma_2) \dots \text{UH}_q(h_{k(n-1)}, [A\bar{z} - c\bar{\ell}_1 2^d]_{k(n-1)}, 2\gamma_2) \end{bmatrix}$ ;

10 for  $i = 1$  to  $k$  do
11 |   for  $j = 0$  to  $n - 1$  do
12 | |   if  $\theta_{ij} \neq [\bar{w}_1]_{ij}$  then
13 | | |   Set  $h_{ij} = 1$ 
14 | |   end
15 |   end
16 end
17 if  $\text{UH}_q(\bar{h}, A\bar{z} - c\bar{\ell}_1 2^d, 2\gamma_2) \neq \bar{w}_1$  or  $\|\bar{z}\|_\infty \geq \gamma_1 - \beta$  or  $\text{wt}(\bar{h}) > \omega$  then
18 |   Go to 3
19 else
20 |   return  $(\sigma = (\bar{z}, \bar{h}, c))$ 
21 end

```

Fig. 5: Forgery(pk, \bar{s}_1, M)

and \bar{t}_0 , are related to one another and it is public knowledge that knowing any of the two aforementioned components will result in a direct break of the signature scheme. In this research, we have shown that just the knowledge of \bar{s}_1 is sufficient for an existential forgery attack.

A natural question arises as to whether the knowledge of \bar{s}_2 only will be sufficient to forge signatures. But, according to the claim of Bai and Galbraith [2], the signature does not send any information about the error \bar{s}_2 and thus does prove only the knowledge of \bar{s}_1 , but not \bar{s}_2 . Thus with only the knowledge of \bar{s}_2 , it should be impossible to forge signatures given the fact that \bar{s}_1 and \bar{s}_2 are independent. Thus, we show that among all the elements of the secret key $(\rho, K, tr, \bar{s}_1, \bar{s}_2, \bar{t}_0)$, just knowing \bar{s}_1 is sufficient to forge signatures, thus showing that the secrecy of \bar{s}_1 is pivotal for the security of Dilithium.

5 Retrieval of \bar{s}_1

To complete our existential forgery attack on Dilithium, we demonstrate a power side-channel attack to retrieve \bar{s}_1 by targetting the polynomial multiplication operation $c\bar{s}_1$ during the signing procedure. Recent public implementations of Dilithium utilise a number theoretic transform (NTT) for polynomial multiplication. Side-channel vulnerability of NTT-based polynomial multiplications is already known [27]. Designers can use other alternatives for computing multiplication, but the choice of the multiplication architecture does not necessarily protect against SCA. There are several works that acknowledge the fact that both the schoolbook and the sparse polynomial multiplier might yield better performance in certain corner cases over area-constrained devices, due to its simplified control logic [13, 23, 24]. Additionally, since the signature component has to be generated in the normal domain, use of the schoolbook or the sparse polynomial multiplier might yield better performance results compared to the use of the NTT polynomial multiplier. The sparse polynomial multiplier has been used in the reference implementation of another candidate lattice-based signature scheme, qTESLA [4], which further motivates us analyse the schoolbook polynomial multiplier and its efficient variant.

We use power analysis to demonstrate two types of attacks to recover \bar{s}_1 , targetting the polynomial multiplication operation $\bar{s}_1 \times c$, that is Step 19 of the signing procedure in Figure 2. SCA on two multiplier architectures are considered. The first one is a conventional DPA style attack on a schoolbook polynomial multiplier, while the second is a two stage horizontal and vertical DPA style attack on the efficient sparse polynomial multiplier protected with the countermeasure of randomized shuffling of non-zero indices of c . The latter attack is based on implementation by Pöppelmann et al. [25] on 8-bit AVR microcontrollers. Both the schoolbook and sparse polynomial multiplication algorithms have been used for design of compact multipliers in a number of reported implementation works on both hardware and software platforms [8, 24].

We would like to clarify that $\bar{s}_1 \in R_q^\ell$, while $c \in B_{60}$ and hence can be assumed to be an element of R_q . The operation of $\bar{s}_1 \times c$ will be repeated ℓ

times with each polynomial of module \bar{s}_1 . But for notational convenience in this section, $\bar{s}_1 \times \mathbf{c}$ is used to denote a single polynomial multiplication.

5.1 Attacking the schoolbook polynomial multiplier

The operands of the target polynomial multiplication operation ($\bar{s}_1 \times \mathbf{c}$) are \bar{s}_1 , which has small coefficients in $[-\eta, \eta]$, and \mathbf{c} , which is a sparse polynomial with coefficients in $\{0, \pm 1\}$. If a schoolbook algorithm is used for the polynomial multiplier, then $\mathbf{c}(x) = \mathbf{a}(x) \cdot \mathbf{b}(x)$ in an ideal lattice setting can be described by the equation below as follows:

$$\mathbf{a} \cdot \mathbf{b} = \left[\sum_{i=0}^{n-1} \sum_{j=0}^{n-1} \mathbf{a}_i \cdot \mathbf{b}_j x^{i+j} \right] \bmod \langle x^n + 1 \rangle = \sum_{i=0}^{n-1} \sum_{j=0}^{n-1} (-1)^{\lfloor \frac{i+j}{n} \rfloor} \mathbf{a}_i \cdot \mathbf{b}_j x^{i+j \bmod n} \quad (2)$$

The intermediate result $\mathbf{a}_i \cdot \mathbf{b}_j$ is targeted (seen in Equation 2), wherein one of the operands \mathbf{c} is known while the other operand \mathbf{s}_1 is unknown. This is an ideal setting to mount DPA attacks using a divide-and-conquer approach, and the secret-key can be retrieved one coefficient at a time. Each coefficient of \mathbf{s}_1 is multiplied with all the coefficients of \mathbf{c} and observing one signature provides 256 points on the trace for the attack. Hence, information about a particular coefficient can also be retrieved from multiple points on a single trace, in addition to getting information about the same coefficient by observing multiple signatures.

5.2 Attacking the sparse polynomial multiplier

The sparse polynomial multiplier implementation we use [25] is a much more optimized algorithm, which takes advantage of the fact that one of the operands is a sparse polynomial. Only the indices of the non-zero coefficients of \mathbf{c} are stored along with their sign. This restricts the multiplication of the coefficients of \bar{s}_1 to only those that are non-zero, thus saving computational time. As a preliminary countermeasure, the non-zero coefficients of \mathbf{c} are also shuffled and processed in a random order for each multiplication operation. Figure 6 describes the code corresponding to the computation of the sparse polynomial multiplier. Though \mathbf{c} is still known, the random shuffling of its coefficients prevents from mounting a straightforward DPA attack. An attack on a sparse polynomial multiplier was reported by Espitau et al. [12]. We use a similar technique to recover \bar{s}_1 , albeit with modifications to suit Dilithium, while also targeting a different source of leakage compared to the original attack. The first step of the attack works by retrieving the order in which \mathbf{c} is processed during every polynomial multiplication. This is done by performing a horizontal DPA attack on Step 5 of Figure 6. Since, one of the inputs (q) is known, a hypothesis over all the possible values of $\mathbf{c}[j]$ can be constructed to perform a horizontal style DPA attack using information available from a single trace, as the coefficients of \mathbf{c} are shuffled after each multiplication operation. Once the order of indices for each polynomial multiplication is revealed, this information can be utilized to retrieve coefficients of \bar{s}_1 using a DPA attack over computation in Step 9, since the value of val is dependent on $\mathbf{c}[j]$.

```

1 for (q=0; q<N; q++){
2     res[q] = 0;
3     for (j=0; j<N; j++){
4         int8_t val = 1;
5         int16_t i = (q - c[j]);
6         if (i < 0){
7             i += N;
8             val = -val;}
9         val *= s[i];
10        res[q] += val;}
11    }

```

Fig. 6: Code snippet of optimised sparse polynomial multiplier [22]

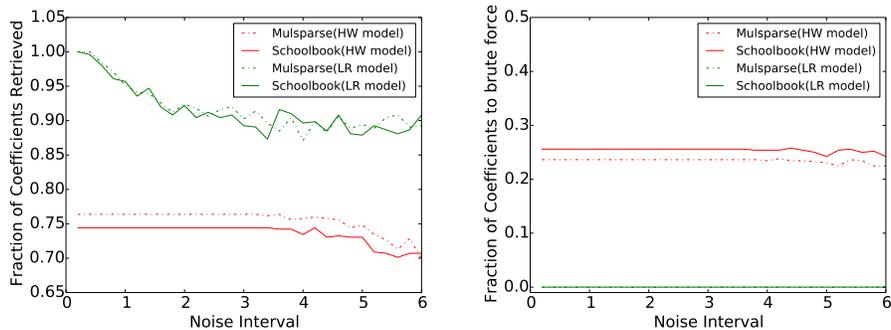
5.3 Experimental results

The above mentioned attacks were performed in a simulated setting, where uniform noise (over a given zero centered interval) is added to the traces, to test the robustness of the attack in a noisy environment. On assuming an 8-bit Hamming weight (HW) model leakage, there is a very limited number of observable hamming weights on the trace, given the fact that s_1 is small and coefficients of $c \in \{0, 1\}$. In addition to this, the distance between the various observable hamming weights is also very small (at least for values with the same sign). One can easily distinguish between whether the correct guess is a positive or negative value (due to the leading ones in the negative value). But distinguishing between different positive values is very tricky, as there is very little information that helps the attacker in distinguishing different key hypotheses. We used the least square error as the statistical distinguisher for both attacks, owing to the decreased number of observable hamming weights. We observed that some of the coefficients could not be retrieved even in a perfect noise free setting. For example, pairs $(3, 5)$ and $(-3, -5)$ are indistinguishable from each other, since they have the same hamming weights over 8 bits. For $\eta = 6$ and $N = 256$ (N being the degree of polynomial), the possible values of the coefficient are in $[-6, 6]$, and thus we cannot distinguish four out of the possible eleven cases. This boils down to a brute-force complexity of around $\approx 2^{79}$, on average, for each polynomial in the secret-key \bar{s}_1 , which renders the attack impractical. A work by Bindel et al. [5] showed that it is enough to reveal only 67% of coefficients in order to reveal the complete secret polynomial \bar{s}_1 for BLISS. It is possible that a similar technique could be used for s_1 in Dilithium to recover the remaining coefficients.

But, it is known that real devices do not necessarily leak in a perfect HW model. Realistically speaking, each bit has a certain bias associated with it and hence each bit weighs differently in leakage. This model of leakage is known as the linear regression model, where leakage corresponding to each bit i is weighed with

a particular coefficient α_i [28]. Both attacks were repeated in a simulated linear regression leakage (LR) model. We found that an introduction of a very small bias that is random in $[-0.01, 0.01]$ to each bit, introduced enough information that resulted in the recovery of almost all coefficients. But to demonstrate a practical case, we use the α coefficients that were profiled using an AES S-box on an AVR microcontroller. An important point to note is that the α coefficients relate to the characteristics of the device and hence any non-sensitive algorithm can also be used for profiling the device. In the case of the LR leakage model, we observed that almost all coefficients were retrieved with very high confidence (with clear distinction between the correct and wrong key hypotheses). The different bias associated with each bit helped distinguish the correct hypothesis from the wrong ones much more easily compared to the HW leakage setting. We also added uniform noise in interval $[-x, x]$ to test the robustness of both the attacks in a noisy environment.

Figures 7a and 7b show the plots for the fraction of coefficients retrieved successfully and the fraction of coefficients to be brute-forced, respectively, against an increasing value of noise interval. Note that the value of x in the x -axis corresponds to uniform noise in the interval $[-x, x]$. It can be seen that the success rate for both the attacks in the HW model is around 75% with high confidence and a brute-force on 25% of the coefficients is required on average. Both attacks in the LR model retrieve almost all coefficients and have a very high success of 90% retrieval, even in presence of noise, with no coefficients to brute-force.



(a) Fraction of coefficients retrieved against uniform noise. (b) Fraction of coefficients to brute-force against uniform noise.

Fig. 7: Practical results for DPA attack against Noise

Comparing the results in Figure 7b with Figure 3, it can be seen that even with very low noise levels, the brute-force complexity can go beyond polynomial time. In such cases, the attacker runs at an advantage if only one of \bar{s}_1 or \bar{s}_2 has to be recovered.

6 Conclusion

In this work, we have demonstrated a side channel assisted forgery attack on the Dilithium signature scheme which is one of the leading candidates to the NIST call for standardization of post-quantum cryptography. Our attack proceeds in two stages, wherein the partial secret-key \bar{s}_1 is retrieved through a novel power analysis attack on the polynomial multiplier. We follow it up with a forgery signature scheme, wherein we show that it is indeed possible to forge signatures with only the knowledge of \bar{s}_1 while bypassing the use of the other parts of the secret-key. We again stress that we neither claim a major break of the Dilithium signature scheme nor any reduction in security of the same. But, we show that the Dilithium scheme breaks on only the knowledge of \bar{s}_1 , which is otherwise not very intuitive given the fact that only a fraction of the LWE instance is revealed as the public-key. Thus the secrecy of \bar{s}_1 is pivotal in ensuring the security of the Dilithium signature scheme.

References

1. Sedat Akleylek, Nina Bindel, Johannes Buchmann, Juliane Krämer, and Giorgia Azzurra Marson. An efficient lattice-based signature scheme with provably secure instantiation. In *International Conference on Cryptology in Africa*, pages 44–60. Springer, 2016.
2. Shi Bai and Steven D Galbraith. An Improved Compression Technique for Signatures Based on Learning with Errors. In *CT-RSA*, volume 8366, pages 28–47, 2014.
3. Rami Barends, Julian Kelly, Anthony Megrant, Andrzej Veitia, Daniel Sank, Evan Jeffrey, Ted C White, Josh Mutus, Austin G Fowler, Brooks Campbell, et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.
4. Nina Bindel, Sedat Akleylek, Erdem Alkim, Paulo S. L. M. Barreto, Johannes Buchmann, Edward Eaton, Gus Gutoski, Juliane Kramer, Patrick Longa, Harun Polat, Jefferson E. Ricardini, and Gustavo Zanon. qTESLA. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>.
5. Nina Bindel, Johannes Buchmann, and Juliane Krämer. Lattice-based signature schemes and their sensitivity to fault attacks. In *Fault Diagnosis and Tolerance in Cryptography (FDTC), 2016 Workshop on*, pages 63–77. IEEE, 2016.
6. Leon Groot Bruinderink, Andreas Hülsing, Tanja Lange, and Yuval Yarom. Flush, Gauss, and Reload—a cache attack on the BLISS lattice-based signature scheme. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 323–345. Springer, 2016.
7. Leon Groot Bruinderink and Peter Pessl. Differential Fault Attacks on Deterministic Lattice Signatures. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(3), 2018. <https://eprint.iacr.org/2018/355.pdf>.
8. Johannes Buchmann, Florian Göpfert, Tim Güneysu, Tobias Oder, and Thomas Pöppelmann. High-performance and lightweight lattice-based public-key encryption. In *Proceedings of the 2nd ACM International Workshop on IoT Privacy, Trust, and Security*, pages 2–9. ACM, 2016.

9. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal Gaussians. In *Advances in Cryptology–CRYPTO 2013*, pages 40–56. Springer, 2013.
10. Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehl. CRYSTALS-Dilithium: A Lattice-Based Digital Signature Scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(1):238–268, Feb. 2018.
11. Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Loop-abort faults on lattice-based Fiat-Shamir and hash-and-sign signatures. In *International Conference on Selected Areas in Cryptography*, pages 140–158. Springer, 2016.
12. Thomas Espitau, Pierre-Alain Fouque, Benoît Gérard, and Mehdi Tibouchi. Side-channel attacks on BLISS lattice-based signatures: Exploiting branch tracing against strongswan and electromagnetic emanations in microcontrollers. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1857–1874. ACM, 2017.
13. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 530–547. Springer, 2012.
14. Nils Gura, Arun Patel, Arvinderpal Wander, Hans Eberle, and Sheueling Chang Shantz. Comparing elliptic curve cryptography and RSA on 8-bit CPUs. In *International Conference on Cryptographic Hardware and Embedded Systems*, volume 4, pages 119–132. Springer, 2004.
15. TP Harty, DTC Allcock, C J Ballance, L Guidoni, HA Janacek, NM Linke, DN Stacey, and DM Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Physical review letters*, 113(22):220501, 2014.
16. Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 598–616. Springer, 2009.
17. Vadim Lyubashevsky, Leo Ducas, Eike Kiltz, Tancrede Lepoint, Peter Schwabe, Gregor Seiler, and Damien Stehle. CRYSTALS-Dilithium. Technical report, National Institute of Standards and Technology, 2017. available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions>.
18. NIST. Post-Quantum Crypto Project. <http://csrc.nist.gov/groups/ST/post-quantum-crypto/>, 2016.
19. NIST. Submission requirements and evaluation criteria for the post-quantum cryptography standardization process. <https://csrc.nist.gov/csrc/media/projects/post-quantum-cryptography/documents/call-for-proposals-final-dec-2016.pdf>, 2016.
20. Peter Pessl. Analyzing the shuffling side-channel countermeasure for lattice-based signatures. In *INDOCRYPT 2016*, pages 153–170. Springer, 2016.
21. Peter Pessl, Leon Groot Bruinderink, and Yuval Yarom. To BLISS-B or not to be: Attacking strongSwan’s Implementation of Post-Quantum Signatures. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security*, pages 1843–1855. ACM, 2017.
22. Thomas Pöppelmann, Léo Ducas, and Tim Güneysu. Enhanced lattice-based signatures on reconfigurable hardware. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 353–370. Springer, 2014.

23. Thomas Pöppelmann and Tim Güneysu. Towards efficient arithmetic for lattice-based cryptography on reconfigurable hardware. In *International Conference on Cryptology and Information Security in Latin America*, pages 139–158. Springer, 2012.
24. Thomas Pöppelmann and Tim Güneysu. Area optimization of lightweight lattice-based encryption on reconfigurable hardware. In *Circuits and Systems (ISCAS), 2014 IEEE International Symposium on*, pages 2796–2799. IEEE, 2014.
25. Thomas Pöppelmann, Tobias Oder, and Tim Güneysu. High-performance ideal lattice-based cryptography on 8-bit atxmega microcontrollers. In *International Conference on Cryptology and Information Security in Latin America*, pages 346–365. Springer, 2015.
26. John Preskill. Reliable quantum computers. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 454, pages 385–410. The Royal Society, 1998.
27. Robert Primas, Peter Pessl, and Stefan Mangard. Single-trace side-channel attacks on masked lattice-based encryption. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 513–533. Springer, 2017.
28. Werner Schindler, Kerstin Lemke, and Christof Paar. A stochastic model for differential side channel cryptanalysis. In *International Conference on Cryptographic Hardware and Embedded Systems*, pages 30–46. Springer, 2005.