Evaluating the security of CRYSTALS-Dilithium in the quantum random oracle model

Kelsey A. Jackson*[1], Carl A. Miller†[1,2], and Daochen Wang‡[1,3]

[1]Joint Center for Quantum Information and Computer Science (QuICS), University of Maryland
[2]Computer Security Division, National Institute of Standards and Technology (NIST)
[3]Department of Computer Science, University of British Columbia

Abstract

In the wake of recent progress on quantum computing hardware, the National Institute of Standards and Technology (NIST) is standardizing cryptographic protocols that are resistant to attacks by quantum adversaries. The primary digital signature scheme that NIST has chosen is CRYSTALS-Dilithium. The hardness of this scheme is based on the hardness of three computational problems: Module Learning with Errors (MLWE), Module Short Integer Solution (MSIS), and SelfTargetMSIS. MLWE and MSIS have been well-studied and are widely believed to be secure. However, SelfTargetMSIS is novel and, though classically as hard as MSIS, its quantum hardness is unclear. In this paper, we provide the first proof of the hardness of SelfTargetMSIS via a reduction from MLWE in the Quantum Random Oracle Model (QROM). Our proof uses recently developed techniques in quantum reprogramming and rewinding. A central part of our approach is a proof that a certain hash function, derived from the MSIS problem, is collapsing. From this approach, we deduce a new security proof for Dilithium under appropriate parameter settings. Compared to the only other rigorous security proof for a variant of Dilithium, Dilithium-QROM, our proof has the advantage of being applicable under the condition $q = 1 \bmod 2n$, where q denotes the modulus and n the dimension of the underlying algebraic ring. This condition is part of the original Dilithium proposal and is crucial for the efficient implementation of the scheme. We provide new secure parameter sets for Dilithium under the condition $q = 1 \bmod 2n$, finding that our public key sizes and signature sizes are about $2.5\times$ to $2.8\times$ larger than those of Dilithium-QROM for the same security levels.

1 Introduction

Quantum computers are theoretically capable of breaking the underlying computational hardness assumptions for many existing cryptographic schemes. Therefore, it is vitally important to develop new cryptographic primitives and protocols that are resistant to quantum attacks. The goal of NIST's Post-Quantum Cryptography Standardization Project is to design a new generation of cryptographic schemes that are secure against quantum adversaries. In 2022, NIST selected three new digital signature schemes for standardization [Ala+22]: Falcon, SPHINCS+, and CRYSTALS-Dilithium. Of the three, CRYSTALS-Dilithium [Bai+21], or Dilithium in shorthand, was identified as the primary choice for post-quantum digital signing.

To practically implement post-quantum cryptography, users must be provided with not only assurance that a scheme is secure in a post-quantum setting, but also the means by which to judge parameter choices and thereby balance their own needs for security and efficiency. The goal of the current work is to provide rigorous assurance of the security of Dilithium as well as implementable parameter sets. A common model for the security of digital signatures is existential unforgeability against chosen message attacks, or EUF-CMA. In this setting, an adversary is allowed to make sequential queries to a signing oracle for the signature scheme, and then afterwards the adversary attempts to forge a signature for a new message. We work in the setting of *strong* existential unforgeability (sEUF-CMA) wherein we must also guard against the possibility that an

*kaj22475@umd.edu
†camiller@umd.edu
‡wdaochen@gmail.com

adversary could try to forge a new signature for one of the messages already signed by the oracle. (See Section 2 for details.)

Additionally, we utilize the quantum random oracle model (QROM) for hash functions. We recall that when a hash function $H : X \to Y$ is used as a subroutine in a digital signature scheme, the random oracle model (ROM) assumes that one can replace each instance of the function H with a black box that accepts inputs from X and returns outputs in Y according to a uniformly randomly chosen function from X to Y. (This model is useful because random functions are easier to work with in theory than actual hash functions.) The random oracle model needs to be refined in the quantum setting because queries to the hash function can be made in superposition: for any quantum state of the form $\sum_{x\in X} \alpha_x |x\rangle$, where $\forall x \in X, \alpha_x \in \mathbb{C}$, a quantum computer can efficiently prepare the superposed state $\sum_{x\in X} \alpha_x |x\rangle |H(x)\rangle$. The quantum random oracle model (QROM) therefore assumes that each use of the hash function can be simulated by a black box that accepts a quantum state supported on X and returns a quantum state supported on $X \times Y$ (computed by a truly random function from X to Y) [Bon+11]. While no efficient and truly random functions actually exist, the QROM is generally trusted and it enables the application of a number of useful proof techniques.

1.1 Known security results for Dilithium

CRYSTALS-Dilithium is based on arithmetic over the ring $R_q := \mathbb{Z}_q[X]/(X^n + 1)$, where q is an odd prime and n is a power of 2. Similar to other Dilithium literature, we generally leave the parameters q, n implicit. For any non-negative integer η, let $S_\eta \subseteq R_q$ denote the set of all polynomials with coefficients from $\{-\eta, -\eta + 1, \ldots, \eta\}$. The security analysis for Dilithium in [Bai+21] is based on three computational problems. The first two are standard problems (Definitions 1 and 2) but the third problem is non-standard (Definition 3).

The first problem is the Module Learning With Errors (MLWE) problem. Assuming that a matrix $A \in R_q^{m\times k}$ and short vectors $s_1 \in S_\eta^k$ and $s_2 \in S_\eta^m$ are chosen uniformly at random, the MLWE problem is to distinguish the matrix-vector pair $(A, t := As_1 + s_2)$ from a uniformly random matrix-vector pair.

Definition 1 (Module Learning with Errors (MLWE))**.** Let $m, k, \eta \in \mathbb{N}$. The advantage of an algorithm $\mathcal{A}$ for solving $\mathsf{MLWE}_{m,k,\eta}$ is defined as:

$$\begin{aligned}\mathrm{Adv}^{\mathsf{MLWE}}_{m,k,\eta}(\mathcal{A}) :=\big|&\Pr[b = 0 \mid A \leftarrow R_q^{m\times k},\ t \leftarrow R_q^m,\ b \leftarrow \mathcal{A}(A, t)] \\ &- \Pr[b = 0 \mid A \leftarrow R_q^{m\times k},\ (s_1, s_2) \leftarrow S_\eta^k \times S_\eta^m, t := As_1 + s_2,\ b \leftarrow \mathcal{A}(A, t)]\big|.\end{aligned} \tag{1}$$

Here, the notation $\mathcal{A}(x)$ denotes $\mathcal{A}$ taking input x. We note that the MLWE problem is often phrased in other contexts with the short vectors s_1 and s_2 coming from a Gaussian, rather than a uniform, distribution. The use of a uniform distribution is one of the particular features of CRYSTALS-Dilithium.

The second problem, MSIS, is concerned with finding short solutions to randomly chosen linear systems over R_q.

Definition 2 (Module Short Integer Solution (MSIS))**.** Let $m, k, \gamma \in \mathbb{N}$. The advantage of an algorithm $\mathcal{A}$ for solving $\mathsf{MSIS}_{m,k,\gamma}$ is defined as:

$$\mathrm{Adv}^{\mathsf{MSIS}}_{m,k,\gamma}(\mathcal{A}) := \Pr\big[[I_m|A]\cdot y = 0 \wedge 0 < \|y\|_\infty \le \gamma \mid A \leftarrow R_q^{m\times k},\ y \leftarrow \mathcal{A}(A)\big]. \tag{2}$$

The third problem is a more complex variant of MSIS that incorporates a hash function H.

Definition 3 (SelfTargetMSIS)**.** Let $\tau, m, k, \gamma \in \mathbb{N}$ and $H\colon \{0,1\}^* \to B_\tau$, where $B_\tau \subseteq R_q$ is the set of polynomials with exactly τ coefficients in $\{-1, 1\}$ and all remaining coefficients zero. The advantage of an algorithm $\mathcal{A}$ for solving $\mathsf{SelfTargetMSIS}_{H,\tau,m,k,\gamma}$ is defined as[1]:

$$\mathrm{Adv}^{\mathsf{SelfTargetMSIS}}_{H,\tau,m,k,\gamma}(\mathcal{A}) := \Pr\Big[H([I_m|A]\cdot y \parallel M) = y_{m+k} \wedge \|y\|_\infty \le \gamma \;\Big|\; A \leftarrow R_q^{m\times k},\ (y, M) \leftarrow \mathcal{A}^{|H\rangle}(A)\Big]. \tag{3}$$

[1] $\parallel$ denotes string concatenation. $\mathcal{A}^{|H\rangle}$ denotes $\mathcal{A}$ with quantum query access to H – a formal definition can be found in Definition 4.

The security guarantee for CRYSTALS-Dilithium is given in [KLS18, Section 4.5] by the inequality[2]

$$\mathrm{Adv}^{\mathsf{sEUF\text{-}CMA}}_{\mathsf{Dilithium}}(\mathcal{A}) \le \mathrm{Adv}^{\mathsf{MLWE}}_{k,l,\eta}(\mathcal{B}) + \mathrm{Adv}^{\mathsf{MSIS}}_{k,l,\zeta'}(\mathcal{D}) + \mathrm{Adv}^{\mathsf{SelfTargetMSIS}}_{H,\tau,k,l+1,\zeta}(\mathcal{C}), \tag{4}$$

where all terms on the right-hand side of the inequality depend on parameters that specify Dilithium, and sEUF-CMA stands for strong unforgeability under chosen message attacks. The interpretation of Eq. (4) is: if there exists a quantum algorithm $\mathcal{A}$ that attacks the sEUF-CMA-security of Dilithium, then there exist quantum algorithms $\mathcal{B}, \mathcal{C}, \mathcal{D}$ for MLWE, MSIS, and SelfTargetMSIS that have advantages satisfying Eq. (4) and run in time comparable to $\mathcal{A}$. Eq. (4) implies that breaking the sEUF-CMA security of Dilithium is at least as hard as solving one of the MLWE, MSIS, or SelfTargetMSIS problems. While MLWE and MSIS are known to be no harder than LWE and SIS, respectively, there are no known attacks taking advantage of their module structure so it is generally believed that they are as hard as their unstructured counterparts [LS15]. In turn, LWE and SIS are at least as hard as the (Gap) Shortest Vector Problem, which is the underlying hard problem of lattice cryptography [Ajt96; Reg09; Pei16].

However, the final problem, SelfTargetMSIS, is novel and so its difficulty is an open question. The problem is known to be as classically hard as MSIS since there exists a reduction from MSIS to SelfTargetMSIS in the ROM [KLS18; BN06]. The reduction uses the following "rewinding" argument. Any randomized algorithm can be specified by a deterministic circuit with auxiliary random bits. Therefore, given a randomized algorithm for SelfTargetMSIS, we can run its deterministic circuit with some randomly chosen bits to obtain one solution and then rewind and run it again using the same bits chosen from before, while at the same time reprogramming the random oracle at the query corresponding to the output of the first run, to obtain a second solution. Subtracting these two solutions to SelfTargetMSIS yields a solution to MSIS. However, the argument fails for the following reasons in the QROM (where a quantum algorithm can make queries in superposition to a quantum random oracle):

1. The randomness in a quantum algorithm includes the randomness of measurement outcomes. We cannot run a quantum algorithm twice and guarantee that the "random bits" will be the same in both runs because we cannot control measurement outcomes. More generally, we cannot rewind a quantum algorithm to a post-measurement state.
2. Since a quantum algorithm can make queries in superposition, it is no longer clear where to reprogram the random oracle.

Currently, the only explicit rigorous proof of Dilithium's security based on conventional hardness assumptions [KLS18] requires modifying the parameters such that $q = 5 \bmod 8$ and $2\gamma < \sqrt{q/2}$, where γ is a length upper bound on vectors corresponding to valid signatures. This ensures that all non-zero vectors in $S_{2\gamma}$ are invertible which equips Dilithium with a so-called "lossy mode". This variant is called Dilithium-QROM. [KLS18] then prove that a signature scheme with such a lossy mode is EUF-CMA. However, the Dilithium specification [Bai+21] uses a value of q satisfying $q = 1 \bmod 2n$ (for $n = 256$) which is incompatible with the assumption that $q = 5 \bmod 8$. The fact that $q = 1 \bmod 2n$ is central to claims about the speed of the algorithms in [Bai+21]: this condition implies that R_q is isomorphic to the direct product ring $\mathbb{Z}_q^{\times n}$ (or $\mathbb{Z}_q^n$ in shorthand) via the Number Theoretic Transform which allows for fast matrix multiplication over R_q. Therefore, it is highly desirable to find a security proof that works under the assumption that $q = 1 \bmod 2n$. Moreover, when $q = 5 \bmod 8$, the ring R_q is *structurally* different from when $q = 1 \bmod 2n$ since in the former case R_q is isomorphic to $\mathbb{F}_{q^{n/2}} \times \mathbb{F}_{q^{n/2}}$ [LN17]. Therefore, it may be imprudent to translate any claims of security in the case $q = 5 \bmod 8$ to the case $q = 1 \bmod 2n$.

1.2 Overview of main result

The main result of our paper is the first proof of the computational hardness of the SelfTargetMSIS problem, presented in Section 3. This hardness result implies a new security proof for Dilithium which, unlike the previous proof in [KLS18], applies to the case $q = 1 \bmod 2n$. Specifically, we reduce MLWE to

[2] Strictly speaking, there should be two other terms ($\mathrm{Adv}^{\mathsf{PR}}_{\mathsf{Sam}}(\mathcal{D})$ and $2^{-\alpha+1}$) on the right-hand side of Eq. (4). However, we ignore them in the introduction as it is easy to set parameters such that these terms are very small. We also mention that the original proof of this inequality uses a flawed analysis of Fiat-Shamir with aborts. The flaw was found and fixed in [Bar+23; Dev+23].

SelfTargetMSIS. By Eq. (4), our result implies that the security of Dilithium (with parameters that are not too far from the original parameters) can be based on the hardness of MLWE and MSIS.

Theorem 1 (Informal version of Theorem 2)**.** *Let $m, k, \tau, \gamma, \eta \in \mathbb{N}$. Suppose $q \geq 16$, $q = 1 \bmod 2n$, and $2\gamma\eta n(m+k) < \lfloor q/32 \rfloor$. If there exists an efficient quantum algorithm $\mathcal{A}$ that solves $\mathsf{SelfTargetMSIS}_{H,\tau,m,k,\gamma}$ with advantage ϵ, under the assumption that H is a random oracle, then there exists an efficient quantum algorithm for solving $\mathsf{MLWE}_{m+k,m,\eta}$ with advantage at least $\Omega(\epsilon^2/Q^4)$. Here, Q denotes the number of quantum queries $\mathcal{A}$ makes to H.*

We now give a high-level overview of the proof. The first step is to define two experiments: the *chosen-coordinate binding* experiment CCB and the *collapsing* experiment Collapse. These experiments are interactive protocols between a verifier and a prover. The protocols end with the verifier outputting a bit b. If $b = 1$, the prover is said to *win* the experiment. The reduction then proceeds in three steps: (i) reduce winning CCB to solving SelfTargetMSIS, (ii) reduce winning Collapse to winning CCB, and (iii) reduce solving MLWE to winning Collapse. Combining these steps together gives a reduction from MLWE to SelfTargetMSIS. The reduction can be illustrated as

$$\mathsf{SelfTargetMSIS} \overset{(i)}{\longleftarrow} \mathsf{CCB} \overset{(ii)}{\longleftarrow} \mathsf{Collapse} \overset{(iii)}{\longleftarrow} \mathsf{MLWE}, \tag{5}$$

where the left arrow means "reduces to".

Step (i): SelfTargetMSIS $\leftarrow$ CCB**.** In the CCB experiment, the prover is first given a uniformly random $A \in R_q^{m \times l}$ which it uses to send the verifier some $z \in R_q^m$, the verifier then sends the prover a challenge c chosen uniformly at random from B_τ, and finally the prover sends the verifier a response $y \in R_q^l$. The prover wins if $Ay = z$, $\|y\|_\infty \leq \gamma$, and the last coordinate of y is c.

We directly apply the main result of [DFM20] to reduce winning CCB when $l = m + k$ to solving $\mathsf{SelfTargetMSIS}_{H,\tau,m,k,\gamma}$ when H is a random oracle. In more detail, the result implies that an efficient algorithm that wins SelfTargetMSIS using Q queries with probability ϵ can be used to construct another efficient algorithm that wins CCB with probability at least $\Omega(\epsilon/Q^2)$.

Step (ii): CCB $\leftarrow$ Collapse**.** In the Collapse experiment, the prover is first given a uniformly random $A \in R_q^{m \times l}$ which it uses to send the verifier some $z \in R_q^m$ together with a quantum state that must be supported only on $y \in R_q^l$ such that $Ay = z$, $\|y\|_\infty \leq \gamma$. Then, the verifier samples a uniformly random bit b'. If $b' = 1$, the verifier measures the quantum state in the computational basis, otherwise, it does nothing. The verifier then returns the quantum state to the prover. The prover responds by sending a bit b' to the verifier and wins if $b' = b$. The advantage of the prover is $2p - 1$ where p is its winning probability.

By using techniques in [DS23; Unr16], we reduce winning Collapse to winning CCB. More specifically, we show that an efficient algorithm that wins CCB with advantage ϵ can be used to construct another efficient algorithm that wins Collapse with advantage at least $\epsilon(\epsilon - 1/|B_\tau|)$, which is roughly ϵ^2 since $1/|B_\tau|$ is very small for the τs we will consider. We generalize techniques in [DS23; Unr16] to work for challenge sets of size > 2, which is necessary since the challenge set in the CCB experiment, B_τ, generally has size > 2. The key idea of first applying the quantum algorithm for winning CCB to the uniform superposition of all challenges remains the same.

Step (iii): Collapse $\leftarrow$ MLWE**.** We build on techniques in [LMZ23; LZ19] to reduce winning Collapse to winning MLWE. More specifically, we show that an efficient algorithm that wins Collapse with advantage ϵ can be used to construct another efficient algorithm that solves $\mathsf{MLWE}_{l,m,\eta}$ with advantage at least $\epsilon/4$. Given a quantum state supported on $y \in R_q^l$ with $Ay = z$ and $\|y\|_\infty \leq \gamma$, as promised in the Collapse experiment, [LMZ23; LZ19] considers the following two measurements. Sample $b \in R_q^l$ from one of the two distributions defined in MLWE (see Eq. (1)), compute a rounded version of $b \cdot y$ in a separate register, and measure that register. When $n = 1$, [LMZ23] shows that the effect of the measurement in one case is close to the computational basis measurement and in the other case is close to doing nothing. Therefore, an algorithm for winning Collapse can be used to solve MLWE. Our work extends [LMZ23] to arbitrary n provided $q = 1 \bmod 2n$. The extension relies on the fact that each coefficient of $b \cdot \Delta$, where $0 \neq \Delta \in R_q$ and

b is chosen uniformly at random from R_q, is uniformly random in $\mathbb{Z}_q$. (This is despite the fact that $b \cdot \Delta$ is generally not uniformly random in R_q.) We establish this fact using the explicit form of the isomorphism between R_q and $\mathbb{Z}_q^n$ when $q = 1 \bmod 2n$.

Finally, in Section 4, we propose explicit sets of parameters using $q \approx 3 \times 10^{13}$ and $n = 256$, such that $q = 1 \bmod 2n$. These sets of parameters achieve different levels of security. We compare our sets of parameters with sets proposed by the Dilithium specifications [Bai+21] and the Dilithium-QROM construction of [KLS18]. Our parameter sets lead to larger public key and signature sizes for the same security levels. The advantage of our parameter sets is that they *provably* (in contrast to parameters proposed by Dilithium specifications) endow Dilithium with sEUF-CMA-security for a q satisfying $q = 1 \bmod 2n$ (in contrast to parameters proposed in Dilithium-QROM).

2 Preliminaries

$\mathbb{N}$ denotes the set of positive integers. For $k \in \mathbb{N}$, $[k]$ denotes the set $\{1, \dots, k\}$. An alphabet refers to a finite non-empty set. Given an alphabet S, the notation $s \leftarrow S$ denotes selecting an element s uniformly at random from S. Given two alphabets A and B, the notation B^A denotes the set of functions from A to B. We write the concatenation of arbitrary strings a, b as $a \parallel b$. Given matrices $A_1, \dots, A_n$ of the same height, $[A_1|A_2|...|A_n]$ denotes the matrix with the A_is placed side by side. log refers to the base-2 logarithm.

We always reserve the symbol q for an odd prime and n for a positive integer that is a power of 2. R_q denotes the ring $\mathbb{Z}_q[X]/(X^n + 1)$ (following the convention in other Dilithium literature [Bai+21; KLS18], we leave the n-dependence implicit). For $k \in \mathbb{N}$, a primitive kth root of unity in $\mathbb{Z}_q$ is an element $x \in \mathbb{Z}_q$ such that $x^k = 1$ and $x^j \neq 1$ for all $j \in [k-1]$; such elements exist if and only if $q = 1 \bmod k$. Given $r \in \mathbb{Z}_q$, we define $r \bmod^{\pm} q$ to be the unique element $r' \in \mathbb{Z}$ such that $-(q-1)/2 \leq r' \leq (q-1)/2$ and $r' = r \bmod q$. For any $r = a_0 + a_1 X + \cdots + a_{n-1} X^{n-1} \in R_q$, we define $|r|_i := |a_i \bmod^{\pm} q|$ for all $i \in \{0, 1, \dots, n-1\}$ and $\|r\|_\infty := \max_i |r|_i$. For $r \in R_q^m$, we define $\|r\|_\infty := \max_{i \in [m]} \|r_i\|_\infty$. For $\eta \in \mathbb{N}$, S_η denotes the set $\{r \in R_q \mid \|r\|_\infty \leq \eta\}$. For $\tau \in \mathbb{N}$, B_τ denotes the set $\{r \in R_q \mid \|r\|_\infty = 1, \|r\| = \sqrt{\tau}\}$ and so $|B_\tau| = 2^\tau \binom{n}{\tau}$.

2.1 Quantum computation

A (quantum) state, or density matrix, ρ on $\mathbb{C}^d$ is a positive semi-definite matrix in $\mathbb{C}^{d \times d}$ with trace 1. A pure state is a state of rank 1. Since a pure state can be uniquely written as $|\psi\rangle\langle\psi|$ where $|\psi\rangle \in \mathbb{C}^d$ and $\langle\psi| := |\psi\rangle^\dagger$, we usually refer to a pure state by just $|\psi\rangle$. A (projective) measurement is given by a set $\mathcal{P} = \{P_1, \dots, P_k\} \subseteq \mathbb{C}^{d \times d}$ such that $\sum_i P_i = 1$, and $\forall i, j \in [k]$, $P_i = P_i^\dagger$ and $P_i P_j = \delta_{i,j} P_i$. The effect of performing such a measurement on ρ is to produce the density matrix $\sum_{i=1}^k P_i \rho P_i$.

A register is either (i) an alphabet Σ or (ii) an m-tuple $X = (Y_1, \dots, Y_m)$ where $m \in \mathbb{N}$ and $Y_1, \dots, Y_m$ are alphabets.

Case (i) The size of the register is $|\Sigma|$, a density matrix on the register refers to a density matrix on $\mathbb{C}^{|\Sigma|}$, and the computational basis measurement on the register refers to the measurement $\{|x\rangle\langle x| \mid x \in \Sigma\}$, where $|x\rangle$ denotes the vector in $\mathbb{C}^\Sigma \cong \mathbb{C}^{|\Sigma|}$ that is 1 in the xth position and zero elsewhere.

Case (ii) The size of the register is $|Y_1| \times \cdots \times |Y_m|$, a density matrix on the register refers to a density matrix on $\mathbb{C}^{|Y_1|} \otimes \cdots \otimes \mathbb{C}^{|Y_m|}$, and the computational basis measurement on the register refers to the measurement $\{|y_1\rangle\langle y_1| \otimes \cdots \otimes |y_m\rangle\langle y_m| \mid y_1 \in Y_1, \dots, y_m \in Y_m\}$.

A quantum algorithm $\mathcal{A}$ is specified by a register $X = (Y_1, \dots, Y_m)$ where $|Y_i| = 2$ for all i and a sequence of elementary gates, i.e., $2^m \times 2^m$ unitary matrices that are of the form

$$S := \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad H := \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad \text{or} \quad \text{CNOT} := \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{6}$$

tensored with 2×2 identity matrices.[3] The unitary matrix U associated with $\mathcal{A}$ is the product of its elementary gates in sequence. The time complexity of $\mathcal{A}$, Time($\mathcal{A}$), is its number of elementary gates.

[3] When we later consider a quantum algorithm on a register of size $d \in \mathbb{N}$, we mean a quantum algorithm on a register $(Y_1, \dots, Y_m)$ where $|Y_i| = 2$ for all i and m is the smallest integer such that $2^m \geq d$.

To perform a computation given an input $x \in \{0,1\}^k$ where $k \leq m$, $\mathcal{A}$ applies U to the starting state $|\psi_0\rangle \coloneqq |x_1+1\rangle \otimes \cdots \otimes |x_k+1\rangle \otimes |1\rangle^{\otimes(m-k)}$ and measures all registers in the computational basis. We also need the definition of a quantum *query* algorithm.

Definition 4 (Quantum query algorithm)**.** Let $t \in \mathbb{N}$. A quantum query algorithm $\mathcal{A}$ using t queries is specified by registers X, Y, Z and a sequence of $t+1$ quantum algorithms $\mathcal{A}_0, \mathcal{A}_1, \ldots, \mathcal{A}_t$, each with register (X, Y, Z). The time complexity of $\mathcal{A}$, $\mathrm{Time}(\mathcal{A})$, is $t + \sum_{i=0}^{t} \mathrm{Time}(\mathcal{A}_i)$.

Let U_i denote the unitary associated with $\mathcal{A}_i$, $\gamma \coloneqq |Y|$, and $\phi\colon Y \to \mathbb{Z}_\gamma$ be a bijection. Given $H\colon X \to Y$, let O^H denote the unitary matrix defined by $O^H |x\rangle |y\rangle |z\rangle = |x\rangle |\phi^{-1}(\phi(y) + \phi(H(x)))\rangle |z\rangle$ for all $(x, y, z) \in X \times Y \times Z$. Then:

1. $\mathcal{A}^{|H\rangle}$ denotes the algorithm with register (X, Y, Z) that computes as follows. Apply U_0 to the starting state $|\psi_0\rangle$. Then, for each $i = 1, \ldots, t$ in sequence, apply O^H then U_i. Finally, measure all registers in the computational basis.
2. $\mathcal{A}^H$ denotes the algorithm with register (X, Y, Z) that computes as follows. Apply U_0 to the starting state $|\psi_0\rangle$. Then, for each $i = 1, \ldots, t$ in sequence, measure register X in the computational basis and apply O^H then U_i. Finally, measure all registers in the computational basis.

In the definitions of $\mathcal{A}^{|H\rangle}$ and $\mathcal{A}^H$, we have described what it means for a quantum algorithm to make quantum and classical queries to a function H, respectively. Under this description, we can naturally define quantum query algorithms that make classical queries to one function and quantum queries to another. Such algorithms are relevant in the security definition of Dilithium as described in the next subsection.

2.2 Digital signature schemes

Let par be common system parameters shared by all participants.

Definition 5 (Digital signature scheme)**.** A digital signature scheme is defined by a triple of randomized algorithms $\mathsf{SIG} = (\mathrm{Gen}, \mathrm{Sign}, \mathrm{Ver})$ such that

1. The key generation algorithm $\mathrm{Gen}(\mathrm{par})$ outputs a public-key, secret-key pair (pk, sk) such that pk defines the message set MSet.
2. The signing algorithm $\mathrm{Sign}(sk, m)$, where $m \in \mathrm{MSet}$, outputs a signature σ.
3. The verification algorithm $\mathrm{Ver}(pk, m, \sigma)$ outputs a single bit $\{0, 1\}$.

We say SIG has correctness error $\gamma \geq 0$ if for all (pk, sk) in the support of $\mathrm{Gen}(\mathrm{par})$ and all $m \in \mathrm{MSet}$,

$$\Pr[\mathrm{Ver}(pk, m, \sigma) = 0 \mid \sigma \leftarrow \mathrm{Sign}(sk, m)] \leq \gamma. \tag{7}$$

Definition 6 (EUF-CMA and sEUF-CMA)**.** Let $\mathsf{SIG} = (\mathrm{Gen}, \mathrm{Sign}, \mathrm{Ver})$ be a signature scheme. Let $\mathcal{A}$ be a quantum query algorithm. Then

$$\mathrm{Adv}^{\mathsf{EUF\text{-}CMA}}_{\mathsf{SIG}}(\mathcal{A}) \coloneqq \Pr[\mathrm{Ver}(pk, m, \sigma) = 1,\ m \notin \mathrm{SignQ} \mid (pk, sk) \leftarrow \mathrm{Gen}(\mathrm{par}), (m, \sigma) \leftarrow \mathcal{A}^{\mathrm{Sign}(sk,\cdot)}(pk)],$$
$$\mathrm{Adv}^{\mathsf{sEUF\text{-}CMA}}_{\mathsf{SIG}}(\mathcal{A}) \coloneqq \Pr[\mathrm{Ver}(pk, m, \sigma) = 1,\ (m, \sigma) \notin \mathrm{SignQR} \mid (pk, sk) \leftarrow \mathrm{Gen}(\mathrm{par}), (m, \sigma) \leftarrow \mathcal{A}^{\mathrm{Sign}(sk,\cdot)}(pk)],$$

where SignQ is the set of queries made by $\mathcal{A}$ to $\mathrm{Sign}(sk, \cdot)$ and SignQR is the set of query-response pairs $\mathcal{A}$ sent to and received from $\mathrm{Sign}(sk, \cdot)$.

When par is a function of $\lambda \in \mathbb{N}$, we say that SIG is (s)EUF-CMA-secure if for every $\mathrm{poly}(\lambda)$-time quantum query algorithm $\mathcal{A}$, we have $\mathrm{Adv}^{\mathsf{(s)EUF\text{-}CMA}}_{\mathsf{SIG}}(\mathcal{A}) \leq \mathrm{negl}(\lambda)$.

In this paper, we use the definition of the Dilithium signature scheme as specified in [Bai+21]. In the concrete parameters section, Section 4, we adopt the same notation as in [Bai+21]. The definition of Dilithium involves a function $H\colon \{0,1\}^* \to B_\tau$ that is classically accessible by its Sign and Ver algorithms. In the definitions of EUF-CMA and sEUF-CMA security of Dilithium, we assume that the quantum algorithm $\mathcal{A}$ has classical query access to $\mathrm{Sign}(sk, \cdot)$ and quantum query access to H. Our proof of Dilithium's security will assume that H can be modeled by a random oracle.

2.3 Cryptographic problems and experiments

We now give the formal definitions of the chosen-coordinate binding and collapsing experiments mentioned in the introduction. More general versions of these definitions can be found in, e.g., [Unr12; DS23]. First, we define a "plain" version of $\mathsf{SelfTargetMSIS}$, where the input matrix is not given in Hermite Normal Form. First reducing $\mathsf{SelfTargetMSIS}$ from $\mathsf{Plain\text{-}SelfTargetMSIS}$ will be convenient later on.

Definition 7 ($\mathsf{Plain\text{-}SelfTargetMSIS}$)**.** Let $\tau, m, l, \gamma \in \mathbb{N}$ and $H\colon \{0,1\}^* \to B_\tau$. The advantage of solving $\mathsf{Plain\text{-}SelfTargetMSIS}_{H,\tau,m,l,\gamma}$ with a quantum query algorithm $\mathcal{A}$ for message $M \in \{0,1\}^*$ is defined as

$$\mathrm{Adv}^{\mathsf{Plain\text{-}SelfTargetMSIS}}_{H,\tau,m,l,\gamma}(\mathcal{A}) \coloneqq \Pr\Big[H(Ay \parallel M) = y_l \wedge \|y\|_\infty \le \gamma \;\Big|\; A \leftarrow R_q^{m\times l},\ (y, M) \leftarrow \mathcal{A}^{|H\rangle}(A)\Big]. \tag{8}$$

Definition 8 (Chosen-Coordinate Binding (CCB))**.** Let $\tau, m, l, \gamma \in \mathbb{N}$. The advantage of a quantum algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ for winning $\mathsf{CCB}_{\tau,m,l,\gamma}$, denoted $\mathrm{Adv}^{\mathsf{CCB}}_{\tau,m,k,\gamma}(\mathcal{A})$, is defined as the probability that the experiment below outputs 1.

> **Experiment** $\mathsf{CCB}_{\tau,m,l,\gamma}$.
> 1. Sample $A \leftarrow R_q^{m\times l}$.
> 2. $(z, T) \leftarrow \mathcal{A}_1(A)$, where $z \in R_q^m$ and T is an arbitrary register.
> 3. Sample $c \leftarrow B_\tau$.
> 4. $y \leftarrow \mathcal{A}_2(T, c)$, where $y \in R_q^l$.
> 5. Output 1 if $Ay = z$, $\|y\|_\infty \le \gamma$, and $y_l = c$.

When τ, m, l, γ are functions of $\lambda \in \mathbb{N}$, we say that the MSIS hash function is chosen-coordinate binding (CCB) if for every poly(λ)-time quantum algorithm $\mathcal{A}$, $\mathrm{Adv}^{\mathsf{CCB}}_{\tau,m,l,\gamma}(\mathcal{A}) \le 1/|B_\tau| + \mathrm{negl}(\lambda)$.

Definition 9 (Collapsing ($\mathsf{Collapse}$))**.** Let $m, l, \gamma \in \mathbb{N}$. The advantage of a quantum algorithm $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ for winning $\mathsf{Collapse}_{m,l,\gamma}$, denoted $\mathrm{Adv}^{\mathsf{Collapse}}_{m,l,\gamma}$, is defined as $2p - 1$ where p is the probability the experiment below outputs 1.

> **Experiment** $\mathsf{Collapse}_{m,l,\gamma}$.
> 1. Sample $A \leftarrow R_q^{m\times l}$.
> 2. $(Y, Z, T) \leftarrow \mathcal{A}_1(A)$, where Y is a register on R_q^l, Z is a register on R_q^m, and T is an arbitrary register.
> 3. Sample $b \leftarrow \{0,1\}$. If $b = 1$, measure Y in the computational basis.
> 4. $b' \leftarrow \mathcal{A}_2(Y, Z, T)$.
> 5. Output 1 if $b' = b$.

We say $\mathcal{A}$ is valid if the state on the register (Y, Z) output by $\mathcal{A}_1$ in step 2 is supported on elements $(y, z) \in R_q^l \times R_q^m$ such that $Ay = z$ and $\|y\|_\infty \le \gamma$. When m, l, γ are functions of $\lambda \in \mathbb{N}$, we say that the MSIS hash function is collapsing if for every poly(λ)-time quantum algorithm $\mathcal{A}$, $\mathrm{Adv}^{\mathsf{Collapse}}_{m,l,\gamma}(\mathcal{A}) \le 1/2 + \mathrm{negl}(\lambda)$.

3 Security proof for $\mathsf{SelfTargetMSIS}$

The main result of this subsection is the following theorem which follows from Propositions 1 to 4.

Theorem 2 ($\mathsf{SelfTargetMSIS}$ security)**.** *Let $m, k, \tau, \gamma, \eta \in \mathbb{N}$. Suppose $q \ge 16$, $q = 1 \bmod 2n$, and $2\gamma\eta n(m + k) < \lfloor q/32 \rfloor$. Suppose that there exists a quantum query algorithm $\mathcal{A}$ for solving $\mathsf{SelfTargetMSIS}_{H,\tau,m,k,\gamma}$ using Q queries and expected advantage ϵ over uniformly random $H\colon \{0,1\}^* \to B_\tau$. Then, for all $w \in \mathbb{N}$, there exists a quantum algorithm $\mathcal{B}$ that solves $\mathsf{MLWE}_{m+k,m,\eta}$ with advantage at least*

$$\frac{\epsilon - nq^{-k}}{4(2Q+1)^2}\left(\frac{\epsilon - nq^{-k}}{(2Q+1)^2} - \frac{1}{|B_\tau|}\right) - \frac{1}{4}\frac{1}{3^w}. \tag{9}$$

Moreover, $\mathrm{Time}(\mathcal{B}) \le \mathrm{Time}(\mathcal{A}) + \mathrm{poly}(\log|B_\tau|, w, n, \log q, m, k)$.

Assuming that the choice of parameters as functions of the security parameter λ is such that $nq^{-k} = \mathrm{negl}(\lambda)$, $1/|B_\tau| = \mathrm{negl}(\lambda)$, and $w = \mathrm{poly}(\lambda)$, Theorem 2 shows that the advantage of $\mathcal{B}$ is roughly $\Omega(\epsilon^2/Q^4)$.

The proof of Theorem 2 proceeds by the following sequence of reductions, which we have labeled by the number of the section in which they are proven:

$$\mathsf{SelfTargetMSIS} \xleftarrow{3.2} \mathsf{Plain\text{-}SelfTargetMSIS} \xleftarrow{3.3} \mathsf{CCB} \xleftarrow{3.4} \mathsf{Collapse} \xleftarrow{3.5} \mathsf{MLWE}.$$

First, we establish some properties of R_q that will be used in Sections 3.2 and 3.5.

3.1 Properties of R_q

Lemma 1. *Suppose $q = 1 \bmod 2n$. Let w be a primitive $(2n)$-th root of unity in $\mathbb{Z}_q$. Then for all $m \in \mathbb{Z}$ such that $0 \neq |m| < n$, the following equation holds in $\mathbb{Z}_q$: $\sum_{j=0}^{n-1} w^{2mj} = 0$.*

Proof. Consider the following equation in $\mathbb{Z}_q$:

$$(1 - w^{2m}) \cdot \sum_{j=0}^{n-1} w^{2mj} = 1 - w^{2mn} = 0, \tag{10}$$

where the first equality uses a telescoping sum and the second uses $w^{2n} = 1$. But $1 - w^{2m} \neq 0$ since $0 \neq |m| < n$ and w is a primitive $(2n)$-th root of unity in $\mathbb{Z}_q$. Therefore, since $\mathbb{Z}_q$ is an integral domain when q is prime, $\sum_{j=0}^{n-1} w^{2mj} = 0$ as required. □

Lemma 2. *Suppose $q = 1 \bmod 2n$. Then, $R_q \cong \mathbb{Z}_q^n$ as algebras over $\mathbb{Z}_q$.*[4]

Proof. For q prime, the multiplicative group $\mathbb{Z}_q^*$ of non-zero elements in $\mathbb{Z}_q$ is cyclic. Let g be a generator of $\mathbb{Z}_q^*$. Let $w := g^{(q-1)/(2n)}$, which is well-defined since $q = 1 \bmod 2n$. Define the mapping $\phi \colon R_q \to \mathbb{Z}_q^n$ by:

$$\phi(p(x)) = \begin{pmatrix} 1 & w & \dots & w^{n-1} \\ 1 & w^3 & \dots & w^{3(n-1)} \\ \vdots & & \ddots & \vdots \\ 1 & w^{(2n-1)} & \dots & w^{(2n-1)(n-1)} \end{pmatrix} \begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{n-1} \end{pmatrix}, \tag{11}$$

where $p(x) := a_0 + a_1 x + \cdots + a_{n-1}x^{n-1}$. It is clear that ϕ is a linear map. To see that ϕ is homomorphic with respect to multiplication, observe that for any $\tilde{p}(x) \in \mathbb{Z}_q[x]$ such that $p(x) = \tilde{p}(x) \bmod (x^n + 1)$, we have

$$\phi(p(x)) = \big(\tilde{p}(w^1), \tilde{p}(w^3), \dots, \tilde{p}(w^{(2n-1)})\big), \tag{12}$$

since $(w^{2k-1})^n + 1 = 0$ in $\mathbb{Z}_q$ for all $k \in [n]$.

To see that ϕ is bijective, observe its explicit inverse $\phi' \colon \mathbb{Z}_q^n \to R_q$, defined by

$$\phi'(c_0, \dots, c_{n-1}) = a_0 + a_1 x + \cdots + a_{n-1}x^{n-1}, \quad \text{where} \tag{13}$$

$$\begin{pmatrix} a_0 \\ a_1 \\ \vdots \\ a_{n-1} \end{pmatrix} := n^{-1} \begin{pmatrix} 1 & 1 & \dots & 1 \\ w^{-1} & w^{-3} & \dots & w^{-(2n-1)} \\ \vdots & & \ddots & \vdots \\ w^{-(n-1)} & w^{-3(n-1)} & \dots & w^{-(2n-1)(n-1)} \end{pmatrix} \begin{pmatrix} c_0 \\ c_1 \\ \vdots \\ c_{n-1} \end{pmatrix} \tag{14}$$

and n^{-1} denotes the multiplicative inverse of n in $\mathbb{Z}_q$, which exists since $q = 1 \bmod 2n \implies n < q$.

Since w is a primitive $(2n)$-th root of unity in $\mathbb{Z}_q$, Lemma 1 implies that the matrices corresponding to ϕ and ϕ' multiply to the identity in $\mathbb{Z}_q$. Therefore, ϕ' is the inverse of ϕ. □

[4] To be clear, the algebra $\mathbb{Z}_q^n$ over $\mathbb{Z}_q$ refers to the set $\mathbb{Z}_q^n$ equipped with component-wise addition and multiplication, and scalar multiplication defined by $\alpha \cdot (c_0, \dots, c_{n-1}) := (\alpha c_0, \dots, \alpha c_{n-1})$, where $\alpha \in \mathbb{Z}_q$ and $(c_0, \dots, c_{n-1}) \in \mathbb{Z}_q^n$.

3.2 Reduction from Plain-SelfTargetMSIS to SelfTargetMSIS

Proposition 1. *Suppose $q = 1 \bmod 2n$. Let $m, k, \gamma, \tau \in \mathbb{N}$ and $H\colon \{0,1\}^* \to B_\tau$. Suppose that there exists a quantum query algorithm $\mathcal{A}$ using Q queries that solves $\mathsf{SelfTargetMSIS}_{H,\tau,m,k,\gamma}$ with advantage ϵ, then there exists a quantum query algorithm $\mathcal{B}$ using Q queries for solving $\mathsf{Plain\text{-}SelfTargetMSIS}_{H,\tau,m,m+k,\gamma}$ with advantage at least $\epsilon - n/q^k$. Moreover, $\mathrm{Time}(\mathcal{B}) \leq \mathrm{Time}(\mathcal{A}) + O(n\log(q)\cdot mk\min(m,k))$.*

Proof. The probability that a uniformly random $B \leftarrow \mathbb{Z}_q^{m\times(m+k)}$ has row-echelon form $[I_m|B']$ (i.e., rank m) is at least $(1-1/q^k)$. Therefore, by Lemma 2, the probability that a uniformly random $A \leftarrow R_q^{m\times(m+k)}$ does not have row-echelon form $[I_m|A']$ is at most $1-(1-1/q^k)^n \leq n/q^k$. When A has row-echelon form $[I_m|A]$, $\mathcal{B}$ first performs row reduction and then runs $\mathcal{A}$. Since the time to perform row reduction on A is $O(n\log(q)\cdot mk\min(m,k))$, the proposition follows. □

3.3 Reduction from CCB to Plain-SelfTargetMSIS

Let S, U, C, R be alphabets, $V\colon S\times U\times C\times R \to \{0,1\}$, and $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ be a quantum algorithm. We define the Σ-experiment by:

> **Σ-experiment.**
> 1. $s \leftarrow S$.
> 2. $(u, T) \leftarrow \mathcal{B}_1(s)$, where $u \in U$ and T is an arbitrary register.
> 3. $c \leftarrow C$.
> 4. $r \leftarrow \mathcal{B}_2(T, c)$.
> 5. Output 1 if $V(s, u, c, r) = 1$.

The advantage of $\mathcal{B}$ for winning the Σ-experiment is the probability of the experiment outputting 1.

In this subsection, we use the following theorem from [DFM20].

Theorem 3 (Measure-and-reprogram [DFM20, Theorem 2])**.** *Let $\mathcal{A}$ be a quantum query algorithm using Q queries that takes input $s \in S$ and outputs $u \in U$ and $r \in R$. Then, there exists a two-stage quantum algorithm $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ (not using any queries) such that the advantage of $\mathcal{B}$ in the Σ-experiment is at least*

$$\frac{1}{(2Q+1)^2}\Pr\Big[V(s,u,H(u),r)\ \Big|\ H \leftarrow C^U,\ s \leftarrow S,\ (u,r) \leftarrow \mathcal{A}^{|H\rangle}(s)\Big]. \tag{15}$$

Moreover, $\mathrm{Time}(\mathcal{B}_1) + \mathrm{Time}(\mathcal{B}_2) \leq \mathrm{Time}(\mathcal{A}) + Q$ *Here, $f \in C^U$ indicates a function $f : U \to C$.*

We remark that in the original statement of the theorem, $\mathrm{Time}(\mathcal{B}_1) + \mathrm{Time}(\mathcal{B}_2)$ is upper bounded by $\mathrm{Time}(\mathcal{A}) + \mathrm{poly}(Q, \log(|U|, \log(|C|))$. The second term accounts for the cost of instantiating Q queries to a $2(Q+1)$-wise independent hash function family from U to C. By the well-known Vandermonde matrix method (see, e.g., [Zha12, Section 6]), this cost can be upper bounded by $O(Q^2\cdot\log(|U|)\cdot\log(|C|))$. However, we follow the convention in [KLS18, Section 2.1] and equate this cost to Q under the fair assumption that $\mathcal{B}$, like $\mathcal{A}$, can also query a random oracle at unit cost.

Proposition 2. *Let $m, l, \gamma, \tau \in \mathbb{N}$. Let $H\colon \{0,1\}^* \to B_\tau$. Suppose there exists a quantum query algorithm $\mathcal{A}$ for solving $\mathsf{Plain\text{-}SelfTargetMSIS}_{H,\tau,m,l,\gamma}$ using Q queries with expected advantage ϵ over uniformly random H. Then there exists a quantum algorithm $\mathcal{B} = (\mathcal{B}_1, \mathcal{B}_2)$ for winning $\mathsf{CCB}_{\tau,m,l,\gamma}$ with advantage at least $\epsilon/(2Q+1)^2$. Moreover* $\mathrm{Time}(\mathcal{B}_1) + \mathrm{Time}(\mathcal{B}_2) \leq \mathrm{Time}(\mathcal{A}) + Q$.

Proof. The quantum query algorithm $\mathcal{A}$ for $\mathsf{Plain\text{-}SelfTargetMSIS}_{H,\tau,m,l,\gamma}$ takes input A and outputs (y, M). So there exists another quantum query algorithm $\mathcal{A}'$ using Q queries that outputs $((Ay \parallel M),\ y)$.

The first part of the proposition follows from applying Theorem 3 to $\mathcal{A}'$ with the following parameter settings which make the Σ-experiment identical to the $\mathsf{CCB}_{\tau,m,l,\gamma}$ experiment

1. Set $S = R_q^{m\times l}$, U to be the query space of $\mathcal{A}'$, $C = B_\tau$, and $R = R_q^l$.

2. Set $V\colon R_q^{m\times l}\times U\times B_\tau\times R_q^l\to\{0,1\}$ by

$$V(A,u,c,y)=\mathbb{1}[z=Ay,\ \|y\|_\infty\le\gamma,\ y_l=c], \tag{16}$$

where $u\in\{0,1\}^*$ is parsed as $u=(z\parallel M)$ with $z\in R_q^m$ and $M\in\{0,1\}^*$.

□

3.4 Reduction from Collapse to CCB

In this subsection, we will use the following lemma, which can be found as [DS23, Proposition 29].

Lemma 3. *Let P,Q be projectors in $\mathbb{C}^{d\times d}$ and ρ be a density matrix in $\mathbb{C}^d$ such that $\rho Q=\rho$. Then* $\mathrm{tr}(QP\rho P)\ge\mathrm{tr}(P\rho)^2$.

The following proposition is similar to [Unr16, Theorem 32] and [DS23, Theorem 28] except the size of the challenge set in the CCB experiment (in step 3 of Definition 8) is not restricted to being 2.

Proposition 3. *Let $m,l,\gamma,\tau\in\mathbb{N}$. Suppose that there exists a quantum algorithm $\mathcal{A}=(\mathcal{A}_1,\mathcal{A}_2)$ that succeeds in $\mathsf{CCB}_{\tau,m,l,\gamma}$ with advantage ϵ, then there exists a valid quantum algorithm $\mathcal{B}=(\mathcal{B}_1,\mathcal{B}_2)$ that succeeds in $\mathsf{Collapse}_{m,l,\gamma}$ with advantage at least $\epsilon(\epsilon-1/|B_\tau|)$. Moreover,* $\mathrm{Time}(\mathcal{B}_1)\le\mathrm{Time}(\mathcal{A}_1)+\mathrm{Time}(\mathcal{A}_2)+O(ml\log(q)\log(B_\tau))$ *and* $\mathrm{Time}(\mathcal{B}_2)\le\mathrm{Time}(\mathcal{A}_2)+O(\log(B_\tau))$.

Proof. We assume without loss of generality (wlog) that the arbitrary register in step 2 of the $\mathsf{CCB}_{\tau,m,l,\gamma}$ experiment (Definition 8) is of the form (Y,T'), where Y is a register on R_q^l and T' is an arbitrary register. We assume wlog that $\mathcal{A}_1$ prepares a state $|\phi\rangle$ on register (Y,Z,T'), where Z is a register on R_q^m, and measures Z in the computational basis to produce the z in step 2 of the $\mathsf{CCB}_{\tau,m,l,\gamma}$ experiment. We also assume wlog that $\mathcal{A}_2$ acts on its input register (Y,T',C), where C is a register on B_τ that contains the c from step 3 of the $\mathsf{CCB}_{\tau,m,l,\gamma}$ experiment, as follows:

1. Apply a unitary U of the form $\sum_{r\in B_\tau}U_r\otimes|r\rangle\langle r|$ on (Y,T',C).
2. Measure Y in the computational basis.

We proceed to construct $\mathcal{B}=(\mathcal{B}_1,\mathcal{B}_2)$ for the $\mathsf{Collapse}_{m,l,\gamma}$ experiment (Definition 9). We first construct $\mathcal{B}_1$, given input $A\in R_q^{m\times l}$, as follows:

1. Run $\mathcal{A}_1(A)$ to prepare state $|\phi\rangle$ on register (Y,Z,T').
2. Prepare state $|\psi\rangle:=|B_\tau|^{-1/2}\sum_{r\in B_\tau}|r\rangle$ on register C in time $O(\log(B_\tau))$. The current state on register (Y,Z,T',C) is $\sigma:=|\phi\rangle\langle\phi|\otimes|\psi\rangle\langle\psi|$. Apply U on register (Y,T',C) and then measure register (Y,Z,T',C) with the projective measurement $\{\Pi,1-\Pi\}$, where Π is defined by

$$\Pi:=\sum_{r\in B_\tau}\sum_{\substack{(y,z)\in R_q^l\times R_q^m:\\ \|y\|_\infty\le\gamma,\ Ay=z,\ y_l=r}}|y,z\rangle\langle y,z|\otimes 1_{T'}\otimes|r\rangle\langle r|\,. \tag{17}$$

   This measurement can be implemented by computing a bit indicating whether the constraints defining Π are satisfied into a separate register and then measuring that register, which takes time $O(ml\log(q)+\log(B_\tau))$.
3. Let B be a bit register. If Π is measured, set the bit stored in B to 1. If $(1-\Pi)$ is measured, replace the state on register (Y,Z) with $|0^l\rangle\otimes|0^m\rangle$, set the bit stored in B to 0. Then output the register (Y,Z,T',C,B).

Let $T:=(T',C,B)$. We construct $\mathcal{B}_2$, given input register (Y,Z,T), as follows:

1. If B contains 0, output a uniformly random bit $b'\in\{0,1\}$.
2. Else apply $U^\dagger$ on register (Y,T',C). Then measure C with the projective measurement $\{|\psi\rangle\langle\psi|\,,1-|\psi\rangle\langle\psi|\}$ using (the inverse of) the preparation circuit for $|\psi\rangle$ in time $O(\log(|B_\tau|)$. If the outcome is $|\psi\rangle\langle\psi|$, output 0; else output 1.

It is clear that $\mathcal{B}$ is valid by definition. Moreover,

$$\text{Time}(\mathcal{B}_1) \le \text{Time}(\mathcal{A}_1) + \text{Time}(\mathcal{A}_2) + O(ml \log(q) \log(B_\tau)), \tag{18}$$

$$\text{Time}(\mathcal{B}_2) \le \text{Time}(\mathcal{A}_2) + O(\log(B_\tau)). \tag{19}$$

We proceed to lower bound the success probability of $\mathcal{B}$. We analyze the probabilities of the following disjoint cases corresponding to $\mathcal{B}$ being successful.

1. Case 1: In this case, $1 - \Pi$ is measured and $b' = b$. The probability that $1 - \Pi$ is measured is $(1 - \epsilon)$. Conditioned on $1 - \Pi$ being measured, b' is a uniformly random bit so the probability $b' = b$ is $1/2$. Therefore, the overall probability of this case is $(1 - \epsilon)/2$.
2. Case 2: In this case, Π is measured, $b = 1$, and then $1 - |\psi\rangle\langle\psi|$ is measured. The probability that Π is measured is ϵ and the probability that $b = 1$ is $1/2$. We now condition on these two events happening. Since $b = 1$, the state of register C in the input to $\mathcal{B}_2$ is a mixture of states of the form $|r\rangle\langle r|$ where $r \in B_\tau$. This is because $b = 1$ means that register Y is measured in the computational basis and conditioned on Π being measured, the C register is also measured in the computational basis (see the form of Π in Eq. (17)). Therefore, the probability of $\mathcal{B}_2$ measuring $|\psi\rangle\langle\psi|$ is $1/|B_\tau|$. Thus, the overall probability of this case is $\epsilon \cdot (1/2) \cdot (1 - 1/|B_\tau|)$.
3. Case 3: In this case, Π is measured, $b = 0$, and then $|\psi\rangle\langle\psi|$ is measured. The probability that $b = 0$ is $1/2$. Conditioned on $b = 0$, Lemma 3, applied with projectors $|\psi\rangle\langle\psi|$ and $U^\dagger \Pi U$ and state σ, shows that the probability of measuring Π and then $|\psi\rangle\langle\psi|$ is least ϵ^2. Therefore, the overall probability of this case is at least $\epsilon^2/2$.

Summing up the probabilities of the above cases, we see that the success probability of $\mathcal{B}$ is at least

$$\frac{1-\epsilon}{2} + \frac{\epsilon}{2}\Big(1 - \frac{1}{|B_\tau|}\Big) + \frac{\epsilon^2}{2} = \frac{1}{2} + \frac{\epsilon}{2}\Big(\epsilon - \frac{1}{|B_\tau|}\Big). \tag{20}$$

Therefore, the advantage of $\mathcal{B}$ is at least $\epsilon(\epsilon - 1/|B_\tau|)$, as required. ☐

3.5 Reduction from MLWE to Collapse

The proof structure of the main result of this subsection, Proposition 4, follows [LMZ23, Theorem 1]. We need to modify a number of aspects of their proof since it applies to the SIS hash function whereas here we consider its module variant, i.e., the MSIS hash function.

We will use a rounding function $\lfloor \cdot \rceil_t \colon \mathbb{Z}_q \to \{0, 1, \dots, t-1\}$, where $t \in \mathbb{N}$, that is defined as follows. For $j \in \{0, 1, \dots, t-1\}$, define

$$I_j \coloneqq \begin{cases} \{j\lfloor q/t \rfloor, j\lfloor q/t \rfloor + 1, \dots, j\lfloor q/t \rfloor + \lfloor q/t \rfloor - 1\} & \text{if } j \in \{0, 1, \dots, t-2\}, \\ \{(t-1)\lfloor q/t \rfloor, (t-1)\lfloor q/t \rfloor + 1, \dots, q-1\} & \text{if } j = t-1. \end{cases} \tag{21}$$

(Note that I_j contains exactly $\lfloor q/t \rfloor$ elements for $j \in \{0, 1, \dots, t-2\}$ and at least $\lfloor q/t \rfloor$ elements for $j = t-1$ with the constraint that $q/t \le |I_{t-1}| \le q/t + t - 1$.) Then, for $a \in \mathbb{Z}_q$, define $\lfloor a \rceil_t$ to be the unique $j \in \{0, 1, \dots, t-1\}$ such that $a \in I_j$.

We will also use the following convenient notation. Let Y and Z be registers and $f : Y \to Z$. The measurement $y \mapsto f(y)$ on register Y refers to the measurement implemented by computing $f(y)$ into a separate register Z, measuring Z in the computational basis, and discarding the result.

Finally, we will use the following lemma.

Lemma 4. *Let $0 \ne \Delta \in R_q^l$ and $\alpha \in \{0, \dots, n-1\}$. If $b \leftarrow R_q^l$, then $(b \cdot \Delta)_\alpha$ is uniformly distributed in $\mathbb{Z}_q$.*

Proof. Writing $b = (b_1, \dots, b_l)$ and $\Delta = (\Delta_1, \dots, \Delta_l)$, we have

$$(b \cdot \Delta)_\alpha = (b_1 \Delta_1)_\alpha + \dots + (b_l \Delta_l)_\alpha. \tag{22}$$

Since $\Delta \ne 0$, there exists an $i \in [l]$ such that $\Delta_i \ne 0$. To prove the lemma, it suffices to prove that $(b_i \Delta_i)_a$ is uniformly distributed in $\mathbb{Z}_q$.

Let ϕ, ϕ' be as defined in the proof of Lemma 2. Write $\phi(\Delta_i) = (c_0, \ldots, c_{n-1}) \in \mathbb{Z}_q$. Since $\Delta_i \neq 0$ there exists $j \in \{0, \ldots, n-1\}$ such that $c_j \neq 0$. Since b_i is a uniformly random element of R_q, $\phi(b_i)$ is a uniformly random element of $\mathbb{Z}_q^n$. Therefore, the distribution of $(b_i\Delta_i)_\alpha = \phi'(\phi(b_i)\phi(\Delta_i))_\alpha$ (where we used Lemma 2 for the equality) is the same as the distribution of

$$\phi'(d_0c_0, \ldots, d_{n-1}c_{n-1})_\alpha, \quad \text{where} \quad d_0, \ldots, d_{n-1} \leftarrow \mathbb{Z}_q. \tag{23}$$

By the linearity of ϕ',

$$\phi'(d_0c_0, \ldots, d_{n-1}c_{n-1})_\alpha = d_jc_j\phi'(e_j)_\alpha + \sum_{j' \neq j} d_{j'}c_{j'}\phi'(e_{j'})_\alpha, \tag{24}$$

where e_j denotes the jth standard basis vector of $\mathbb{Z}_q$.

But $\phi'(e_j)_\alpha = n^{-1} \cdot w^{-(2j+1)\alpha} \neq 0$ (see Lemma 2). Therefore $d_jc_j\phi'(e_j)_\alpha$ is uniformly distributed in $\mathbb{Z}_q$ if $d_j \leftarrow \mathbb{Z}_q$. Hence $(b_i\Delta_i)_\alpha$ is uniformly distributed in $\mathbb{Z}_q$ as required. □

The main result of this subsection is the following proposition.

Proposition 4. *Let $m, l, \gamma, \eta \in \mathbb{N}$. Suppose $q \geq 16$ and $2\gamma\eta nl < \lfloor q/32 \rfloor$. Suppose there exists a quantum algorithm $\mathcal{A}$ that succeeds in $\mathsf{Collapse}_{m,l,\gamma}$ with advantage ϵ. Then, for all $w \in \mathbb{N}$, there exists a quantum algorithm $\mathcal{B}$ that solves $\mathsf{MLWE}_{l,m,\eta}$ with advantage at least $(\epsilon - 3^{-w})/4$. Moreover,* $\mathrm{Time}(\mathcal{B}) \leq \mathrm{Time}(\mathcal{A}) + \mathrm{poly}(w)$.

Before proving this proposition, we first prove two lemmas. Let Y be a register on R_q^l and $A \in R_q^{m \times l}$. For $t \in \mathbb{N}$, we define the following measurements on Y:

- M_0: computational basis measurement.
- M_1^t: sample $e_1 \leftarrow S_\eta^m$, $e_2 \leftarrow S_\eta^l$, set $b \coloneqq e_1^\top A + e_2^\top \in R_q^l$, sample $s \leftarrow R_q$, then perform measurement $y \mapsto \lfloor (b \cdot y + s)_0 \rceil_t$.
- M_2^t: sample $b \leftarrow R_q^l$, $s \leftarrow R_q$, then perform measurement $y \mapsto \lfloor (b \cdot y + s)_0 \rceil_t$.

Lemma 5. *Let $t \in \mathbb{N}$ be such that $2\gamma\eta nl < \lfloor q/t \rfloor$. For all $y, y' \in R_q^l$ with $Ay = Ay'$ and $\|y'\|_\infty, \|y\|_\infty \leq \gamma$,*

$$M_1^t(|y\rangle\langle y'|) = \left(1 - \frac{t}{q} \cdot \mathrm{E}\left[|e \cdot (y - y')|_0 \;\middle|\; e \leftarrow S_\eta^l\right]\right) |y\rangle\langle y'| \,. \tag{25}$$

Proof. We have

$$M_1^t(|y\rangle\langle y'|) = \Pr\left[\lfloor (b \cdot y + s)_0 \rceil_t = \lfloor (b \cdot y' + s)_0 \rceil_t \;\middle|\; e_1 \leftarrow S_\eta^m, e_2 \leftarrow S_\eta^l, b \coloneqq e_1^\top A + e_2^\top, s \leftarrow R_q\right] \cdot |y\rangle\langle y'| \tag{26}$$

Writing $z \coloneqq Ay = Ay'$, we have

$$b \cdot y + s = (e_1 \cdot z + s) + e_2 \cdot y \quad \text{and} \quad b \cdot y' + s = (e_1 \cdot z + s) + e_2 \cdot y'. \tag{27}$$

The result follows by observing that $|e_2 \cdot (y - y')|_0 \leq \|e_2\|_\infty \cdot \|y - y'\|_\infty \cdot nl \leq 2\gamma\eta nl < \lfloor q/t \rfloor$ and $(e_1 \cdot z + s)$ is a uniformly random element of R_q. □

Lemma 6. *Let $t \in \mathbb{N}$ be such that $t^2 \leq q$. Then there exists $0 \leq p_t \leq 2/t$ such that for all $y, y' \in R_q^l$ with $y' \neq y$, we have*

$$M_2^t(|y\rangle\langle y|) = |y\rangle\langle y| \quad \textit{and} \quad M_2^t(|y\rangle\langle y'|) = p_t \, |y\rangle\langle y'| \,. \tag{28}$$

Proof. The first equality is clearly true. For the second, observe that

$$M_2^t(|y\rangle\langle y'|) = \Pr[\lfloor (b \cdot y + s)_0 \rceil_t = \lfloor (b \cdot y' + s)_0 \rceil_t \mid b \leftarrow R_q^l, s \leftarrow R_q]. \tag{29}$$

Write $y' = y + \Delta$ for some $0 \neq \Delta \in R_q^l$. Then, $(b \cdot \Delta)_0$ is uniformly distributed in $\mathbb{Z}_q$ by Lemma 4. Therefore, writing $p_t := \Pr[\lfloor u \rceil_t = \lfloor u + v \rceil_t \mid u, v \leftarrow \mathbb{Z}_q]$, we have

$$\begin{aligned}\Pr[\lfloor (b \cdot y + s)_0 \rceil_t =& \lfloor (b \cdot y' + s)_0 \rceil_t \mid b \leftarrow R_q^l, s \leftarrow R_q] \\ =& p_t = 1 - \Big(\frac{(t-1)\lfloor q/t \rfloor}{q} \cdot \frac{q - \lfloor q/t \rfloor}{q} + \frac{|I_{t-1}|}{q} \cdot \frac{q - |I_{t-1}|}{q} \Big) \leq \frac{1}{t} + \frac{t}{q} \leq \frac{2}{t},\end{aligned} \tag{30}$$

where the last inequality uses $t^2 \leq q$. □

Combining Lemmas 5 and 6 gives the following corollary.

Corollary 1. *Let $t, d \in \mathbb{N}$ be such that $2\gamma\eta nl < \lfloor q/(td) \rfloor$ and $t^2 \leq q$. Let ρ be a density matrix on register Y. Suppose there exists $z \in R_q^m$ such that ρ is supported on $\{y \in R_q^l \mid Ay = z, \|y\|_\infty \leq \gamma\}$. Then*

$$M_1^t(\rho) = \frac{1}{d} M_1^{td}(\rho) + \Big(1 - \frac{1}{d}\Big)\rho, \tag{31}$$

$$M_2^t(\rho) = \frac{1}{d} M_0(M_1^{td}(\rho)) + \Big(1 - \frac{1}{d} - p_t\Big) M_0(\rho) + p_t \rho, \tag{32}$$

where p_t is as defined in Lemma 6.

Proof. The first equality is immediate. The second equality follows from the observation that $M_0(M_1^{td}(\rho)) = M_1^{td}(M_0(\rho))$ since M_0 and M_1 both act on ρ by entry-wise multiplication. □

Given the above lemmas, Proposition 4 follows from the proof of [LMZ23, Theorem 1]. The high-level idea of the proof is that M_1^t is close to the identity operation while M_2^t is close to M_0. Therefore, if the identity operation can be efficiently distinguished from M_0, then M_1^t and M_2^t can be efficiently distinguished, which solves the MLWE problem. For completeness, we give the details below.

Proof of Proposition 4. Let $t := 4$ and $d := 8$ so that $g := 1 - 1/d - p_t \geq 3/8$ and $dg \geq 3$, where p_t is as defined in Lemma 6. Let $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ be a valid algorithm for the $\mathsf{Collapse}_{m,l,\gamma}$ experiment (Definition 9) with advantage ϵ.

Fix $w \in \mathbb{N}$ and $A \in R_q^{m \times l}$. Let $T := \sum_{j=0}^{w-1} (dg)^{-j}$ and let $\mathcal{B}$ be the quantum algorithm defined on input $b \in R_q^l$ as follows:

1. Create state ρ on register (Y, Z, T) by running $\mathcal{A}_1(A)$.
2. Sample $j \in \{0, 1, \ldots, w-1\}$ with probability $(dg)^{-j}/T$.
3. Apply M_1^{td} to ρ on the Y register for j times. Call the resulting state ρ_j.
4. Sample $s \leftarrow R_q$ and apply the measurement $x \mapsto \lfloor (b \cdot x + s)_0 \rceil_t$ to ρ_j on the Y register to give state ρ'_j.
5. Compute bit $b' \in \{0, 1\}$ by running $\mathcal{A}_2(\rho'_j)$.
6. Output b' if j is even and $1 - b'$ if j is odd.

For $j \in \{0, 1, \ldots, w-1\}$, let ϵ_j denote the signed distinguishing advantage of $\mathcal{A}_2$ on inputs ρ_j versus $M_0(\rho_j)$, i.e., $\epsilon_j := \Pr[\mathcal{A}_2(\rho_j) = 0] - \Pr[\mathcal{A}_2(M_0(\rho_j)) = 0]$, and let δ_j denote the signed distinguishing advantage of $\mathcal{A}_2$ on inputs $M_1^t(\rho_j)$ versus $M_2^t(\rho_j)$. Then the signed distinguishing advantage of $\mathcal{B}$ on input distributions $\big[e_1 \leftarrow S_\eta^m,\ e_2 \leftarrow S_\eta^l,\ b := e_1^\top A + e_2^\top\big]$ versus $\big[b^\top \leftarrow R_q^l\big]$ is

$$\delta := \frac{1}{T} \sum_{j=0}^{w-1} (-dg)^{-j} \delta_j, \tag{33}$$

because $\rho'_j = M_1^t(\rho_j)$ if b is sampled according to $\big[e_1 \leftarrow S_\eta^m,\ e_2 \leftarrow S_\eta^l,\ b := e_1^\top A\ + e_2^\top\big]$ and $\rho'_j = M_2^t(\rho_j)$ if b is sampled according to $\big[b^\top \leftarrow R_q^l\big]$.

By Corollary 1 (which applies by the assumptions in the proposition and the validity of $\mathcal{A}$), we have $\delta_j = \frac{1}{d}\epsilon_{j+1} + g\epsilon_j$ for all $j \in \{0, 1, \ldots, w-2\}$. Therefore,

$$\epsilon_i(-dg)^{-i} = \epsilon_0 - \frac{1}{g}\sum_{j=0}^{i-1}(-dg)^{-j}\delta_j \quad \text{for all } i \in \{0, 1, \ldots, w-1\}. \tag{34}$$

Then,

$$\delta = \frac{g}{T}(\epsilon_0 - \epsilon_w(-dg)^{-w}). \tag{35}$$

We now unfix $A \in R_q^{m\times l}$ and take the expectation of Eq. (35) over $A \leftarrow R_q^{m\times l}$ to see that

$$|\mathrm{E}_A[\delta]| = \frac{g}{T}\Big|\mathrm{E}_A[\epsilon_0 - \epsilon_w(-dg)^{-w}]\Big| \geq \Big(g - \frac{1}{d}\Big)(\epsilon - (dg)^{-w}) \geq \frac{1}{4}\Big(\epsilon - \frac{1}{3^w}\Big), \tag{36}$$

where the first inequality uses $T \leq dg/(dg-1)$, $|\epsilon_w| \leq 1$, and $\epsilon = |\mathrm{E}_A[\epsilon_0]|$.

Since $\mathrm{Time}(\mathcal{B}) = \mathrm{Time}(\mathcal{A}) + \mathrm{poly}(w)$ and $|\mathrm{E}_A[\delta]|$ is the advantage of $\mathcal{B}$ for solving $\mathsf{MLWE}_{l,m,\eta}$, the proposition follows. □

4 Concrete parameters

In this section, we describe how to adjust the parameter settings of Dilithium using Theorem 2 to achieve different levels of security defined by NIST in the relevant Federal Information Processing Standards (FIPS) [NIST23, Appendix A]. We will use the same notation as in the Dilithium specification, [Bai+21]. [Bai+21] specifies Dilithium in terms of the following variables

$$q, n, k, l, H, \tau, d, \tau, \gamma_1, \gamma_2, \eta, \beta. \tag{37}$$

The variables q and n specify the ring R_q as before. The variables k, l are associated with sizes of matrices over R_q. H is the hash function used in Dilithium and τ is such that B_τ is the codomain of H. For conciseness, we will not explain the variables $d, \gamma_1, \gamma_2, \eta, \beta$, and refer the reader to [Bai+21] for their definitions.

The security analysis of CRYSTALS-Dilithium in [KLS18] leads to [KLS18, Eqs. (10) and (11)] which shows the following. Given a quantum query algorithm $\mathcal{A}$ for breaking the sEUF-CMA-security of Dilithium, there exist quantum algorithms $\mathcal{B}, \mathcal{D}, \mathcal{E}$ and quantum query algorithm $\mathcal{C}$ such that $\mathrm{Time}(\mathcal{B}) = \mathrm{Time}(\mathcal{C}) = \mathrm{Time}(\mathcal{A})$ and $\mathrm{Time}(\mathcal{D}) \approx \mathrm{Time}(\mathcal{A})$ with

$$\mathrm{Adv}^{\mathsf{sEUF\text{-}CMA}}_{\mathsf{Dilithium}}(\mathcal{A}) \leq 2^{-\alpha+1} + \mathrm{Adv}^{\mathsf{MLWE}}_{k,l,\eta}(\mathcal{B}) + \mathrm{Adv}^{\mathsf{SelfTargetMSIS}}_{H,\tau,k,l+1,\zeta}(\mathcal{C}) + \mathrm{Adv}^{\mathsf{MSIS}}_{k,l,\zeta'}(\mathcal{D}) + \mathrm{Adv}^{\mathsf{PR}}_{\mathsf{Sam}}(\mathcal{E}), \tag{38}$$

where ζ, ζ' are functions of parameters $\gamma_1, \gamma_2, \beta, d, \tau$ defined as follows:

$$\zeta \coloneqq \max(\gamma_1 - \beta, 2\gamma_2 + 1 + 2^{d-1}\tau) \quad \text{and} \quad \zeta' \coloneqq \max(2(\gamma_1 - \beta), 4\gamma_2 + 2). \tag{39}$$

$\mathrm{Adv}^{\mathsf{PR}}_{\mathsf{Sam}}(\mathcal{E})$ is the advantage of any algorithm distinguishing between the pseudorandom function used by Dilithium and a randomly selected function; and α is a min-entropy term that can be bounded using [KLS18, Lemma C.1 of ePrint version] by

$$\alpha \geq \min\left(-n\log\left(\frac{2\gamma_1+1}{2\gamma_2-1}\right), -kl\log(n/q)\right). \tag{40}$$

In the QROM, we can construct an optimal pseudorandom function using a random oracle such that $\mathrm{Adv}^{\mathsf{PR}}_{\mathsf{Sam}}(\mathcal{E})$ is asymptotically negligible and can be neglected.

Theorem 2 shows that the hardness of SelfTargetMSIS in the QROM is at least that of MLWE. Therefore, Theorem 2 and Eq. (38) rigorously imply the asymptotic result that, under suitable choices of parameters as functions of the security parameter λ, if there are no $\mathrm{poly}(\lambda)$-time quantum algorithms that solve MLWE or MSIS then there is no $\mathrm{poly}(\lambda)$-time quantum algorithm that breaks the sEUF-CMA security of Dilithium.

This is a very positive sign for the security of Dilithium as MSIS and MLWE are far better-studied problems and there is substantial support for the assumption that they are hard problems.

We proceed to give concrete estimates of the Core-SVP security of Dilithium under several choices of parameters using Theorem 2 and Eq. (38). These estimates rely on some heuristic assumptions that we will clearly state. We remark that the concrete security estimates appearing in [KLS18; Bai+21] use similar heuristic assumptions.

We begin by dividing both sides of Eq. (38) by Time($\mathcal{A}$). Using Time($\mathcal{B}$) = Time($\mathcal{C}$) = Time($\mathcal{A}$), assuming the approximation in Time($\mathcal{D}$) $\approx$ Time($\mathcal{A}$) can be replaced by equality, and using the "Our Work" parameters in Tables 2 and 3 for which $\alpha \geq 257$, we obtain

$$\frac{\mathrm{Adv}^{\mathsf{sEUF\text{-}CMA}}_{\mathsf{Dilithium}}(\mathcal{A})}{\mathrm{Time}(\mathcal{A})} \leq 2^{-256} + \frac{\mathrm{Adv}^{\mathsf{MLWE}}_{k,l,\eta}(\mathcal{B})}{\mathrm{Time}(\mathcal{B})} + \frac{\mathrm{Adv}^{\mathsf{SelfTargetMSIS}}_{H,\tau,k,l+1,\zeta}(\mathcal{C})}{\mathrm{Time}(\mathcal{C})} + \frac{\mathrm{Adv}^{\mathsf{MSIS}}_{k,l,\zeta'}(\mathcal{D})}{\mathrm{Time}(\mathcal{D})}. \tag{41}$$

By Theorem 2, for any $\eta' \in \mathbb{N}$ with $\eta' < \lfloor q/32 \rfloor / (2\zeta n(k+l+1))$, there exists a quantum algorithm $\mathcal{C}'$ for $\mathsf{MLWE}_{k+l+1,k,\eta'}$, such that

$$\frac{\mathrm{Adv}^{\mathsf{sEUF\text{-}CMA}}_{\mathsf{Dilithium}}(\mathcal{A})}{\mathrm{Time}(\mathcal{A})} \leq 2^{-256} + \frac{\mathrm{Adv}^{\mathsf{MLWE}}_{k,l,\eta}(\mathcal{B})}{\mathrm{Time}(\mathcal{B})} + \frac{8Q^2\sqrt{\mathrm{Adv}^{\mathsf{MLWE}}_{k+l+1,k,\eta'}(\mathcal{C}')}}{\mathrm{Time}(\mathcal{C})} + \frac{\mathrm{Adv}^{\mathsf{MSIS}}_{k,l,\zeta'}(\mathcal{D})}{\mathrm{Time}(\mathcal{D})}, \tag{42}$$

where Q is the number of queries $\mathcal{C}$ uses and we assumed that Eq. (9) is well-approximated by $\epsilon^2/(64Q^4)$, in particular, that τ is sufficiently large.

Also by Theorem 2, we have Time($\mathcal{C}'$) is at most Time($\mathcal{C}$) plus polynomial terms. Heuristically assuming that we can neglect the polynomial terms and using $Q \leq$ Time($\mathcal{C}$), we obtain

$$\frac{\mathrm{Adv}^{\mathsf{sEUF\text{-}CMA}}_{\mathsf{Dilithium}}(\mathcal{A})}{\mathrm{Time}(\mathcal{A})} \leq 2^{-256} + \frac{\mathrm{Adv}^{\mathsf{MLWE}}_{k,l,\eta}(\mathcal{B})}{\mathrm{Time}(\mathcal{B})} + 8Q^{3/2}\sqrt{\frac{\mathrm{Adv}^{\mathsf{MLWE}}_{k+l+1,k,\eta'}(\mathcal{C}')}{\mathrm{Time}(\mathcal{C}')}} + \frac{\mathrm{Adv}^{\mathsf{MSIS}}_{k,l,\zeta'}(\mathcal{D})}{\mathrm{Time}(\mathcal{D})}. \tag{43}$$

Now, for NIST security level $l \in [5]$, we upper bound Q by B_l, where B_l is given in Table 1.

NIST Security Level (SLl)	SL1	SL2	SL3	SL4	SL5
Upper bound on Q (B_l)	2^{64}	2^{86}	2^{96}	2^{128}	2^{128}

Table 1: Upper bounds on Q for NIST security levels 1 to 5. These numbers are based on [NIST23, Appendix A] together the following well-known query complexity results if we model the block ciphers and hash functions used in [NIST23, Appendix A] as random functions. Given a random function $f\colon [N] \to [N]$, the number of queries to f needed to find a preimage of 1 is $\Theta(N^{1/2})$ [Gro96] and the number of queries to f needed to find a collision, i.e., $i \neq j$ such that $f(i) = f(j)$, is $\Theta(N^{1/3})$ [Zha15]. (We ignore the constants hidden in the Θ-notation; more detailed analysis is possible, see, e.g., [Jaq+20].)

From the third term on the right-hand side of Eq. (43), we see that the Quantum Core-SVP security of SelfTargetMSIS can be estimated by

$$\frac{z}{2} - \frac{3}{2}\log(B_l) - 3, \tag{44}$$

where z is the Quantum Core-SVP security of the associated MLWE problem.

Having reduced the sEUF-NMA security of Dilithium to the security of standard lattice problems MLWE and MSIS, we proceed to estimate their security. Following the analysis in the Dilithium specifications [Bai+21], we perform our security estimates via the Core-SVP methodology introduced in [Alk+16]. In the Core-SVP methodology, we consider attacks using the Block Khorkine-Zolotarev (BKZ) algorithm [SE94; CN11]. The BKZ algorithm with block size $\mu \in \mathbb{N}$ works by making a small number of calls to an SVP solver on μ-dimensional lattices. The Core-SVP methodology conservatively assumes that the run-time of the BKZ algorithm is equal to the cost of a single run of the SVP solver at its core. The latter cost is then estimated as $2^{0.265\mu}$ since this is the cost of the best quantum SVP solver [Bai+21, Section C.1] due to Laarhoven [Laa16, Section 14.2.10]. Therefore, to estimate the security of an MLWE or MSIS problem, it suffices to estimate the smallest $\mu \in \mathbb{N}$ such that BKZ with block-size μ can solve the problem. Then we say 0.265μ is the *Quantum Core-SVP* security of the problem.

To describe how the block-size can be estimated, it is convenient to define the function $\delta \colon \mathbb{N} \to \mathbb{R}$,

$$\delta(\mu) := \left(\frac{(\mu\pi)^{1/\mu}\mu}{2\pi e}\right)^{\frac{1}{2(\mu-1)}}. \tag{45}$$

Concrete security of MLWE. Our security analysis of MLWE generally follows the $\mathsf{Dilithium}$ specifications, [Bai+21, Appendix C.2]. For $a, b, \epsilon \in \mathbb{N}$, we first follow [Bai+21, Appendix C.2] and assume that $\mathsf{MLWE}_{a,b,\epsilon}$ is as hard as the Learning With Errors problem $\mathsf{LWE}_{na,nb,\epsilon}$ — for $a', b' \in \mathbb{N}$, $\mathsf{LWE}_{a',b',\epsilon}$ is defined to be the same as $\mathsf{MLWE}_{a,b,\epsilon}$ with n set to 1 so that $R_q = \mathbb{Z}_q$.

Then, as done in [Bai+21, Appendix C.2], we follow the security analysis in [Alk+16]. [Alk+16] considers two attacks based on the BKZ algorithm, known as the primal attack and dual attack. The block-size is then taken as the minimum of the block-sizes for the primal and dual attacks. These attacks are analyzed as follows.

1. Primal attack [Alk+16, Section 6.3]. Let $d := na + nb + 1$. Then to solve $\mathsf{LWE}_{na,nb,\epsilon}$, we set the BKZ block-size μ to be equal to the smallest integer ≥ 50 such that[5]

$$\xi\sqrt{\mu} \leq \delta(\mu)^{2\mu-d} \cdot q^{na/d}. \tag{46}$$

2. Dual attack [Alk+16, Section 6.4]. Let $d' := na + nb$. Then to solve $\mathsf{LWE}_{na,nb,\epsilon}$, we set the BKZ block-size μ to be equal to the smallest integer ≥ 50 such that

$$-2\pi^2\tau(\mu)^2 \geq \ln(2^{-0.2075\mu/2}), \tag{47}$$

   where $\tau(\mu) := \delta(\mu)^{d'-1} q^{nb/d'} \epsilon/q$.

Concrete security of MSIS. Our security analysis of MSIS uses heuristics in the $\mathsf{Dilithium}$ specifications [Bai+21, Appendix C.3] and [Lyu12] (which is in turn based on [MR09]).[6] For $a, b, \xi \in \mathbb{N}$, we first follow [Bai+21, Appendix C.3] and assume that $\mathsf{MSIS}_{a,b,\xi}$ is as hard as the Short Integer Solutions problem $\mathsf{SIS}_{na,nb,\xi}$ — for $a', b' \in \mathbb{N}$, $\mathsf{SIS}_{a',b',\xi}$ is defined to be the same as $\mathsf{MSIS}_{a',b',\xi}$ with n set to 1 so that $R_q = \mathbb{Z}_q$. Following [Lyu12], we estimate the security of $\mathsf{SIS}_{na,nb,\xi}$, by considering the attack that uses the BKZ algorithm with block-size μ to find a short non-zero vector in the lattice

$$L(A) := \{y \in \mathbb{Z}^{na+nb} \mid [I_{na} \mid A] \cdot y = 0 \bmod q\}, \tag{48}$$

where $A \leftarrow \mathbb{Z}_q^{na \times nb}$. Following [Lyu12, Eq. (3) of ePrint version], the BKZ algorithm is expected to find a vector $v \in L(A)$ of Euclidean length[7]

$$2^{2\sqrt{na\log(q)\log(\delta(\mu))}}. \tag{49}$$

We assume that the entries of v have the same magnitudes since a similar assumption is made in [Bai+21, Appendix C.3]. Then, to solve $\mathsf{SIS}_{na,nb,\xi}$, we set the BKZ block-size μ to be the smallest positive integer such that

$$\frac{1}{\sqrt{na+nb}} \cdot 2^{2\sqrt{na\log(q)\log(\delta(\mu))}} \leq \xi. \tag{50}$$

To set $\mathsf{Dilithium}$ parameters, we also require $q = 1 \bmod 2\gamma_2$, $q > 4\gamma_2$ (see [Bai+21, Lemma 1]), and $\beta = \tau\eta$ (see [Bai+21, Table 2]). Moreover, we set parameters to minimize the following metrics [KLS18]:

[5] In [Alk+16, Section 6.3], the exponent on $\delta(\mu)$ is given as $2\mu - d - 1$, but it is correct to $2\mu - d$ by [Alb+17, Section 3.2]. There can be spurious solutions with $0 < \mu < 50$ for which the approximations leading to the inequality break down.

[6] We were unable to completely reuse the analysis in [Bai+21, Section C.3] as it is not completely described. Comparing the estimates for μ obtained by the method here with that in [Bai+21, Table 1] (also reproduced in Table 2), we find our estimates are consistently around 4/5 times that given in [Bai+21, Table 1]. Therefore, our estimates underestimate the security of MSIS compared to [Bai+21].

[7] Compared to [Lyu12, Eq. (3) of ePrint version], we do not take the min of Eq. (49) with q since "trivial" vectors of the form q times a standard basis vector have too large of an infinity-norm to be a solution to $\mathsf{SIS}_{na,nb,\xi}$ when $\xi < q$, as will be the case for our parameter choices.

1. the public key size in bytes, $(nk(\lceil \log(q) \rceil - d) + 256)/8$,
2. the signature size in bytes, $(nl(\lceil \log(2\gamma_2) \rceil) + nk + \tau(\log(n) + 1))/8$,
3. the expected number of repeats to sign a message, $\exp\left(n\beta\left(\frac{l}{\gamma_1} + \frac{k}{\gamma_2}\right)\right)$.

In Tables 2 and 3, we give parameter sets achieving different levels of security that we calculated using the methodology described above. In both tables, we use

$$q_0 := 29996224302593 = 2^9(218107)(268613) + 1. \tag{51}$$

In particular, $q_0 = 1 \bmod 2n$.

In Table 2, we compare our parameters with those of the Dilithium specifications [Bai+21]. For the same security levels, we see that our public key sizes are about $12\times$ to $15\times$ that of Dilithium and our signature sizes are about $5\times$ to $6\times$ that of Dilithium. We stress that the main advantage of our parameters compared to Dilithium is that ours are based on rigorous reductions from hard lattice problems, whereas Dilithium's are based on highly heuristic reductions. In particular, the heuristic reduction from SelfTargetMSIS to (a variant of) MSIS given in [Bai+21, End of Section 6.2.1] has been recently challenged [Wan+22].

In Table 3, we compare our parameters with those of Dilithium-QROM, another Dilithium-based scheme with rigorous reductions from hard lattice problems. For the same security levels, we see that our public key and signature sizes are about $2.5\times$ to $2.8\times$ that of Dilithium-QROM. We stress that the main advantage of our parameters compared to Dilithium-QROM is that ours have $q = 1 \bmod 2n$ which is crucial for the efficient implementation of the scheme. More specifically, when $q = 1 \bmod 2n$, multiplying two elements in R_q can be performed using the Number Theoretic Transform in time $O(n \log(q))$ (compared to the naive cost of $O(n^2 \log(q))$). In contrast, Dilithium-QROM uses $q = 5 \bmod 8$ which is incompatible with $q = 1 \bmod 2n$ as $n > 2$ is a power of 2.

The main reason why we need to increase the public key and signature sizes is due to the loss in the reduction from MLWE to SelfTargetMSIS, as stated in Theorem 2. More specifically, when we calculate the Quantum Core-SVP numbers for the SelfTargetMSIS-based MLWE problem, we use Eq. (44) which considerably lowers security. Eq. (44) is derived from Theorem 2. We do not know if the loss is inherent or if our reduction could be tightened. Deciding which is the case is the main open question of our work.

	Dilithium [Bai+21]			Our Work		
	SL2	SL3	SL5	SL2	SL3	SL5
q	$2^{23} - 8191$	$2^{23} - 8191$	$2^{23} - 8191$	q_0	q_0	q_0
n	256	256	256	256	256	256
(k, l)	(4, 4)	(6, 5)	(22, 7)	(20, 20)	(25, 25)	(33, 33)
d	13	13	13	14	14	14
τ	39	49	60	45	45	45
γ_1	2^{17}	2^{19}	2^{19}	460800	520000	520000
γ_2	95232	261888	261888	1074452	1074452	1074452
ζ	350209	724481	769537	2517545	2517545	2517545
ζ'	380930	1048184	1048336	4297810	4297810	4297810
η	2	4	2	2	2	2
η'	N/A	N/A	N/A	8	7	5
pk size (bytes)	1312	1952	2592	19872	24832	32768
σ size (bytes)	2420	3293	4595	14757	18437	24325
Expected Repeats	4.25	5.10	3.85	4.17	5.17	8.76
SIS BKZ Block-Size	423	638	909	4476	5787	7942
Quantum Core-SVP	112	169	241	1186	1533	2104
LWE BKZ block-Size	423	624	863	1457, 1859	1916, 2416	2679, 3302
Quantum Core-SVP	112	165	229	386, 114	507, 173	709, 242

Table 2: We give parameter sets that achieve the same security levels as those proposed in the Dilithium specifications [Bai+21]. q_0 is defined in Eq. (51). When there are two comma-separated numbers in a cell, the first number is associated with the first MLWE problem in Eq. (43) and the second number is associated with the second, SelfTargetMSIS-based, MLWE problem in Eq. (43).

	Dilithium-QROM [KLS18]		Our Work	
	recc.	very high	recc.	very high
q	$2^{45} - 21283$	$2^{45} - 21283$	q_0	q_0
n	512	512	256	256
(k, l)	(4, 4)	(5, 5)	(22, 22)	(24, 24)
d	15	15	14	14
τ	46	46	45	45
γ_1	905679	905679	506880	506880
γ_2	905679	905679	1074452	1074452
ζ	2565023	2565023	2517545	2517545
ζ'	3622718	3622718	4297810	4297810
η	7	3	2	2
η'	N/A	N/A	8	7
pk size (bytes)	7712	9632	21856	23840
σ size (bytes)	5690	7098	16229	17701
Expected Repeats	4.29	2.18	4.36	4.98
SIS BKZ Block-Size	N/A	N/A	4996	5522
Quantum Core-SVP	N/A	N/A	1323	1463
LWE BKZ Block-Size	480	600	1639, 2087	1823, 2300
Quantum Core-SVP	127	159	434, 129	483, 157

Table 3: We give parameter sets that achieve the same security levels as those proposed in Dilithium-QROM [KLS18]. q_0 is defined in Eq. (51). In the "Our Work" columns, we assume Q is bounded by 2^{96}, which corresponds to NIST Security Level 3. When there are two comma-separated numbers in a cell, the first number is associated with the first MLWE problem in Eq. (43) and the second number is associated with the second, SelfTargetMSIS-based, MLWE problem in Eq. (43).

Acknowledgments

This work is supported by the National Institute of Standards and Technology (NIST) and the Joint Center for Quantum Information and Computer Science (QuICS) at the University of Maryland. This research paper is partly a contribution of the U. S. federal government and is not subject to copyright in the United States. We thank Marcel Dall'Agnol, Jiahui Liu, Yi-Kai Liu, and Ray Perlner for helpful feedback and correspondence. We thank Amin Shiraz Gilani for his involvement during the early stages of this project.

References

[Ajt96] M. Ajtai. "Generating Hard Instances of Lattice Problems". In: *Proceedings of the 28th ACM Symposium on the Theory of Computing (STOC)*. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 99–108. ISBN: 0897917855. DOI: 10.1145/237814.237838 (Cited on page 3).

[Ala+22] Gorjan Alagic, Daniel Apon, David Cooper, Quynh Dang, Thinh Dang, John Kelsey, Jacob Lichtinger, Carl Miller, Dustin Moody, Rene Peralta, Ray Perlner, Angela Robinson, Daniel Smith-Tone, and Yi-Kai Liu. "Status report on the third round of the NIST post-quantum cryptography standardization process". In: *US Department of Commerce, NIST* (2022). DOI: 10.6028/NIST.IR.8413-upd1 (Cited on page 1).

[Alb+17] Martin R. Albrecht, Florian Göpfert, Fernando Virdia, and Thomas Wunderer. "Revisiting the Expected Cost of Solving uSVP and Applications to LWE". In: *Advances in Cryptology – ASIACRYPT 2017*. Ed. by Tsuyoshi Takagi and Thomas Peyrin. Cham: Springer International Publishing, 2017, pp. 297–322. ISBN: 978-3-319-70694-8. DOI: 10.1007/978-3-319-70694-8_11 (Cited on page 16).

[Alk+16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. "Post-quantum Key Exchange—A New Hope". In: *25th USENIX Security Symposium (USENIX Security 16)*. Austin, TX: USENIX Association, 2016, pp. 327–343. ISBN: 978-1-931971-32-4 (Cited on pages 15, 16).

[Bai+21] Shi Bai, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. "CRYSTALS-Dilithium: Algorithm Specifications and Supporting Documentation, Round 3.1". In: *NIST: Post Quantum Cryptography, Round 3 Finalists: Digital Signature Algorithms* (2021). https://pq-crystals.org/dilithium/ (Cited on pages 1–3, 5, 6, 14–17).

[Bar+23] Manuel Barbosa, Gilles Barthe, Christian Doczkal, Jelle Don, Serge Fehr, Benjamin Grégoire, Yu-Hsuan Huang, Andreas Hülsing, Yi Lee, and Xiaodi Wu. "Fixing and Mechanizing the Security Proof of Fiat-Shamir with Aborts and Dilithium". In: *Advances in Cryptology – CRYPTO 2023*. Ed. by Helena Handschuh and Anna Lysyanskaya. Cham: Springer Nature Switzerland, 2023, pp. 358–389. ISBN: 978-3-031-38554-4. DOI: 10.1007/978-3-031-38554-4_12 (Cited on page 3).

[BN06] Mihir Bellare and Gregory Neven. "Multi-Signatures in the Plain Public-Key Model and a General Forking Lemma". In: *Proceedings of the 13th ACM Conference on Computer and Communications Security*. CCS '06. Alexandria, Virginia, USA: Association for Computing Machinery, 2006, pp. 390–399. ISBN: 1595935185. DOI: 10.1145/1180405.1180453 (Cited on page 3).

[Bon+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. "Random Oracles in a Quantum World". In: *Advances in Cryptology – ASIACRYPT 2011*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 41–69. ISBN: 978-3-642-25385-0. DOI: 10.1007/978-3-642-25385-0_3 (Cited on page 2).

[CN11] Yuanmi Chen and Phong Q. Nguyen. "BKZ 2.0: Better Lattice Security Estimates". In: *Advances in Cryptology – ASIACRYPT 2011*. Ed. by Dong Hoon Lee and Xiaoyun Wang. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 1–20. ISBN: 978-3-642-25385-0. DOI: 10.1007/978-3-642-25385-0_1 (Cited on page 15).

[Dev+23] Julien Devevey, Pouria Fallahpour, Alain Passelègue, and Damien Stehlé. "A Detailed Analysis of Fiat-Shamir with Aborts". In: *Advances in Cryptology – CRYPTO 2023*. Ed. by Helena Handschuh and Anna Lysyanskaya. Cham: Springer Nature Switzerland, 2023, pp. 327–357. ISBN: 978-3-031-38554-4. DOI: 10.1007/978-3-031-38554-4_11 (Cited on page 3).

[DFM20] Jelle Don, Serge Fehr, and Christian Majenz. "The Measure-and-Reprogram Technique 2.0: Multi-round Fiat-Shamir and More". In: *Advances in Cryptology – CRYPTO 2020*. Ed. by Daniele Micciancio and Thomas Ristenpart. Cham: Springer International Publishing, 2020, pp. 602–631. ISBN: 978-3-030-56877-1. DOI: 10.1007/978-3-030-56877-1_21 (Cited on pages 4, 9).

[DS23] Marcel Dall'Agnol and Nicholas Spooner. "On the Necessity of Collapsing for Post-Quantum and Quantum Commitments". In: *18th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC 2023)*. Ed. by Omar Fawzi and Michael Walter. Vol. 266. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023, 2:1–2:23. ISBN: 978-3-95977-283-9. DOI: 10.4230/LIPIcs.TQC.2023.2 (Cited on pages 4, 7, 10).

[Gro96] Lov K. Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the 28th ACM Symposium on the Theory of Computing (STOC)*. 1996, pp. 212–219. DOI: 10.1145/237814.237866 (Cited on page 15).

[Jaq+20] Samuel Jaques, Michael Naehrig, Martin Roetteler, and Fernando Virdia. "Implementing Grover Oracles for Quantum Key Search on AES and LowMC". In: *Advances in Cryptology – EUROCRYPT 2020*. Ed. by Anne Canteaut and Yuval Ishai. Cham: Springer International Publishing, 2020, pp. 280–310. ISBN: 978-3-030-45724-2. DOI: 10.1007/978-3-030-45724-2_10 (Cited on page 15).

[KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. "A Concrete Treatment of Fiat-Shamir Signatures in the Quantum Random-Oracle Model". In: *Advances in Cryptology – EUROCRYPT 2018*. Ed. by Jesper Buus Nielsen and Vincent Rijmen. Cham: Springer International Publishing, 2018, pp. 552–586. ISBN: 978-3-319-78372-7. DOI: 10.1007/978-3-319-78372-7_18 (Cited on pages 3, 5, 9, 14–16, 18).

[Laa16] Thijs Laarhoven. "Search problems in cryptography: from fingerprinting to lattice sieving". English. PhD Thesis. Mathematics and Computer Science, 2016. ISBN: 978-90-386-4021-1 (Cited on page 15).

[LMZ23] Jiahui Liu, Hart Montgomery, and Mark Zhandry. "Another Round of Breaking and Making Quantum Money: How To Not Build It From Lattices, And More". In: *Advances in Cryptology – EUROCRYPT 2023*. Ed. by Carmit Hazay and Martijn Stam. Cham: Springer Nature Switzerland, 2023, pp. 611–638. ISBN: 978-3-031-30545-0. DOI: 10.1007/978-3-031-30545-0_21 (Cited on pages 4, 11, 13).

[LN17] Vadim Lyubashevsky and Gregory Neven. "One-Shot Verifiable Encryption from Lattices". In: *Advances in Cryptology – EUROCRYPT 2017*. Ed. by Jean-Sébastien Coron and Jesper Buus Nielsen. Cham: Springer International Publishing, 2017, pp. 293–323. ISBN: 978-3-319-56620-7. DOI: 10.1007/978-3-319-56620-7_11 (Cited on page 3).

[LS15] Adeline Langlois and Damien Stehlé. "Worst-Case to Average-Case Reductions for Module Lattices". In: *Designs, Codes and Cryptography* 75 (2015), pp. 565–599. DOI: 10.1007/s10623-014-9938-4 (Cited on page 3).

[Lyu12] Vadim Lyubashevsky. "Lattice Signatures without Trapdoors". In: *Advances in Cryptology – EUROCRYPT 2012*. Ed. by David Pointcheval and Thomas Johansson. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 738–755. ISBN: 978-3-642-29011-4. DOI: 10.1007/978-3-642-29011-4_43 (Cited on page 16).

[LZ19] Qipeng Liu and Mark Zhandry. "Revisiting Post-quantum Fiat-Shamir". In: *Advances in Cryptology – CRYPTO 2019*. Ed. by Alexandra Boldyreva and Daniele Micciancio. Cham: Springer International Publishing, 2019, pp. 326–355. ISBN: 978-3-030-26951-7. DOI: 10.1007/978-3-030-26951-7_12 (Cited on page 4).

[MR09] Daniele Micciancio and Oded Regev. "Lattice-based Cryptography". In: *Post-Quantum Cryptography*. Ed. by Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 147–191. ISBN: 978-3-540-88702-7. DOI: 10.1007/978-3-540-88702-7_5 (Cited on page 16).

[NIST23] National Institute of Standards and Technology. *Module-Lattice-Based Digital Signature Standard*. Tech. rep. Federal Information Processing Standards Publications (FIPS PUBS) 204. Washington, D.C.: U.S. Department of Commerce, 2023. DOI: 10.6028/NIST.FIPS.204.ipd (Cited on pages 14, 15).

[Pei16] Chris Peikert. "A Decade of Lattice Cryptography". In: *Found. Trends Theor. Comput. Sci.* 10.4 (2016), pp. 283–424. ISSN: 1551-305X. DOI: 10.1561/0400000074 (Cited on page 3).

[Reg09] Oded Regev. "On Lattices, Learning with Errors, Random Linear Codes, and Cryptography". In: *J. ACM* 56.6 (2009). ISSN: 0004-5411. DOI: 10.1145/1568318.1568324 (Cited on page 3).

[SE94] C. P. Schnorr and M. Euchner. "Lattice basis reduction: Improved practical algorithms and solving subset sum problems". In: *Mathematical Programming* 66.1 (1994), pp. 181–199. DOI: 10.1007/BF01581144 (Cited on page 15).

[Unr12] Dominique Unruh. "Quantum Proofs of Knowledge". In: *Advances in Cryptology – EUROCRYPT 2012*. Ed. by David Pointcheval and Thomas Johansson. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 135–152. ISBN: 978-3-642-29011-4. DOI: 10.1007/978-3-642-29011-4_10 (Cited on page 7).

[Unr16] Dominique Unruh. "Collapse-Binding Quantum Commitments Without Random Oracles". In: *Advances in Cryptology – ASIACRYPT 2016*. Ed. by Jung Hee Cheon and Tsuyoshi Takagi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 166–195. ISBN: 978-3-662-53890-6. DOI: 10.1007/978-3-662-53890-6_6 (Cited on pages 4, 10).

[Wan+22] Geng Wang, Wenwen Xia, Gongyu Shi, Ming Wan, Yuncong Zhang, and Dawu Gu. *Revisiting the Concrete Hardness of SelfTargetMSIS in CRYSTALS-Dilithium*. Cryptology ePrint Archive, Paper 2022/1601. iacr:2022/1601. 2022 (Cited on page 17).

[Zha12] Mark Zhandry. “Secure Identity-Based Encryption in the Quantum Random Oracle Model”. In: *Advances in Cryptology – CRYPTO 2012*. Ed. by Reihaneh Safavi-Naini and Ran Canetti. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 758–775. ISBN: 978-3-642-32009-5. DOI: 10.1007/978-3-642-32009-5_44 (Cited on page 9).

[Zha15] Mark Zhandry. “A note on the quantum collision and set equality problems”. In: *Quantum Information and Computation* 15.7-8 (2015), pp. 557–567. ISSN: 1533-7146 (Cited on page 15).