On the Non-Existence of Short Vectors in Random Module Lattices *

Ngoc Khanh Nguyen

IBM Research – Zurich, Switzerland and Ruhr Universität Bochum, Germany
nkn@zurich.ibm.com

Abstract

Recently, Lyubashevsky & Seiler (Eurocrypt 2018) showed that small polynomials in the cyclotomic ring $\mathbb{Z}_q[X]/(X^n+1)$, where n is a power of two, are invertible under special congruence conditions on prime modulus q. This result has been used to prove certain security properties of lattice-based constructions against unbounded adversaries. Unfortunately, due to the special conditions, working over the corresponding cyclotomic ring does not allow for efficient use of the Number Theoretic Transform (NTT) algorithm for fast multiplication of polynomials and hence, the schemes become less practical.

In this paper, we present how to overcome this limitation by analysing zeroes in the Chinese Remainder (or NTT) representation of small polynomials. Concretely, we follow the proof techniques from Stehlé and Steinfeld (Eprint 2013/004) and provide upper bounds on the probabilities related to the (non)-existence of a short vector in a random module lattice with no assumptions on the prime modulus. Then, we apply these results, along with the generic framework by Kiltz et al. (Eurocrypt 2018), to a number of lattice-based Fiat-Shamir signatures so they can both enjoy tight security in the quantum random oracle model and support fast multiplication algorithms (at the cost of slightly larger public keys and signatures), such as the Bai-Galbraith signature scheme (CT-RSA 2014), Dilithium-QROM (Kiltz et al., Eurocrypt 2018) and qTESLA (Alkim et al., PQCrypto 2017). These techniques can also be applied to prove that recent commitment schemes by Baum et al. (SCN 2018) are statistically binding with no additional assumptions on q.

Keywords: Lattice-based cryptography, Fiat-Shamir signatures, module lattices, lossy identification schemes, provable security.

1 Introduction

Cryptography based on the hardness of lattice problems, such as Module-SIS or Module-LWE [PR06, LM06, LPR10], seems to be a very likely replacement for traditional cryptography after the eventual arrival of quantum computers. With the ongoing NIST PQC Standardization Process, we are closer to using quantum-resistant encryption schemes and digital signatures in real life. For additional efficiency, many practical lattice-based constructions work over *fully-splitting* polynomial rings $R_q := \mathbb{Z}_q[X]/(f(X))$ where $f(X) = X^n + 1$ is a cyclotomic polynomial, n is a power of two and the prime q is selected so that $f(X)$ splits completely into n linear factors modulo q. With such a choice of parameters, multiplication in the polynomial ring can be performed very quickly using the Number Theoretic Transform (NTT), e.g. [GLP12, ADPS16, SAB$^+$17, LDK$^+$17]. Indeed, one obtains a speed-up of about a factor of 5 by working over rings where $X^n + 1$ splits completely versus just 2 factors (for primes of size between 2^{20} and 2^{29} [LS18]). Moreover, the structure of fully-splitting rings allows us to perform various operations in parallel as well as conveniently cache and sample polynomials which also significantly improves efficiency of the schemes.

*This is the full version of [Ngu19] presented at Asiacrypt 2019.

Unfortunately, it is sometimes difficult to prove security of lattice-based constructions when working over fully-splitting polynomial rings [KLS18, BAA[+]17, BDL[+]18]. Usually, the reason is that these security proofs rely on the assumption that polynomials of small norm are invertible. Recently, Lyubashevsky and Seiler [LS18] (generalising [LN17]) showed that when n is a power of two and under certain conditions on prime modulus q, small elements of R_q are indeed invertible. The result, however, is meaningful only when $X^n + 1$ does not split into many factors modulo q (e.g. at most 32 for $n = 512$). Consequently, we cannot apply the standard NTT algorithm in such polynomial rings unless we drop the invertibility assumption [1].

One concrete example is the Dilithium-QROM signature scheme introduced by Kiltz et al. [KLS18] (which is the extended version of Dilithium [LDK[+]17] secure in the quantum random oracle model). Its security relies on the *lossy* properties of its underlying identification scheme. Indeed, the protocol satisfies lossy soundness, i.e. no unbounded adversary can impersonate the prover if the public key was generated from the lossy key generation algorithm [AFLT12]. We briefly recall the argument for Dilithium-QROM. Suppose that $\mathcal{A}$ is an adversary which tries to impersonate the prover. First, assume that $\mathcal{A}$ can only output a valid response for at most one challenge from the challenge space ChSet. Thus, the probability of success can be bounded by $1/|\mathsf{ChSet}|$ which should be negligible. A more interesting case is when we assume that $\mathcal{A}$ can respond correctly to two distinct challenges. Then, by combining verification equations and the fact that the public key is chosen from the lossy key generation, the adversary finds a short vector in a certain random module lattice. Thus, Kiltz et al. set the prime q which satisfies $q \equiv 5 \pmod 8$, so that the invertibility assumption holds. Finally, they prove that the probability of existence of a short vector in such a lattice is negligible for a suitable choice of parameters.

In this paper, we show how to circumvent the invertibility assumption, and combined with the generic framework by Kiltz et al. [KLS18], construct secure lattice-based signatures in the quantum random oracle model without any conditions on prime modulus q. Consequently, we can work over fully-splitting rings and at the same time, use the NTT algorithm for fast multiplication of polynomials. We apply our results to the second-round candidates of the NIST PQC Standardization Process. Namely, we improve the efficiency of Dilithium-QROM as well as qTESLA [BAA[+]17]. We also briefly explain how our techniques can be applied to recent lattice-based commitment schemes [BDL[+]18].

1.1 Our Contribution

Main results. The ultimate goal is to provide an upper bound on the probability of existence of a short vector in a random module lattice (see Theorem 1.1, formally Corollary 3.9) and other related probabilities (Theorem 3.8 and Theorem 3.10). Informally, it states that the probability, over the uniformly random matrix $\mathbf{A}$, that there exists a pair of vectors $(\mathbf{z}_1, \mathbf{z}_2)$, which consists of small polynomials in R_q and $\mathbf{z}_1 \neq \mathbf{0}$, such that $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}$ is small (for a suitable choice of parameters). In the context of Fiat-Shamir identification and signature schemes, $\mathbf{A}$ represents a public key matrix and $\mathbf{z}_1$ (and sometimes $\mathbf{z}_2$ as well) represents a difference of two signatures/responses. In order to prove these results, we apply techniques from Stehlé and Steinfeld [SS13] in the module setting. Our upper bound depends on the tail function $\mathcal{T}$. For readability, we hide the concrete formula for $\mathcal{T}$ here and we refer to the formal statement in Corollary 3.9.

We recall that a similar result was also presented by Kiltz et al. (e.g. Lemma 4.6 in [KLS18]) but they only consider the case when $q \equiv 5 \pmod 8$ so that invertibility properties can be applied [LN17, LS18]. Here, we generalise their result on how to bound that probability without any assumptions on the prime modulus q.

Theorem 1.1 (Informal). *Denote $S_\alpha := \{y \in R_q : ||y||_\infty \leq \alpha\}$ and let $\ell, k, \alpha_1, \alpha_2 \in \mathbb{N}$. Then*

$$\Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}}[\exists (\mathbf{z}_1, \mathbf{z}_2) \in S_{\alpha_1}^{\ell} \backslash \{\mathbf{0}\} \times S_{\alpha_2}^{k} : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}] \leq \frac{|S_{\alpha_1}|^{\ell} \cdot |S_{\alpha_2}|^{k}}{q^{nk}} + \mathcal{T}(q, \ell, k, \alpha_1, \alpha_2), \tag{1}$$

where $\mathcal{T}(q, \ell, k, \alpha_1, \alpha_2)$ is a function defined in Corollary 3.9.

[1] Lyubashevsky and Seiler [LS18] showed, however, how to combine the FFT algorithm and Karatsuba multiplication in order to multiply in partially-splitting rings faster.

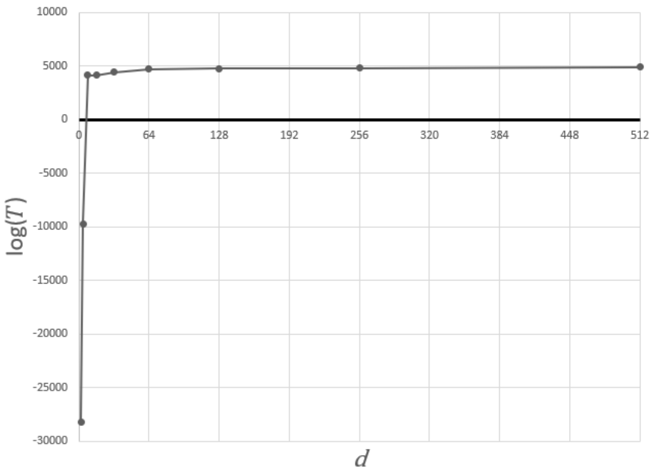

Figure 1: Let $(n, q, \ell, k, \alpha_1, \alpha_2) = (512, \approx 2^{45}, 4, 4, 1.8 \cdot 10^6, 3.6 \cdot 10^6)$. The graph presents values of $\log(\mathcal{T}(q, \ell, k, \alpha_1, \alpha_2))$ depending on the number of irreducible polynomials d that $X^n + 1$ splits into modulo q. One notes that for prime moduli $q \approx 2^{45}$ such that $d \in \{2, 4\}$, the value of $\mathcal{T}$ is sufficiently small, hence so is the right-hand side of Equation (1). On the other hand, values of $\mathcal{T}$ rocket for $d \geq 8$ and therefore q or dimensions (k, ℓ) of the matrix $\mathbf{A}$ must be increased in order to keep the upper bound in (1) small enough.

Figure 1 shows values of the tail function $\mathcal{T}$ for different prime moduli q. We observe that the more $f(x) = X^n + 1$ splits modulo q then the larger the value of $\mathcal{T}$. When $f(x)$ only splits into two factors, our upper bound is essentially equal to

$$\frac{|S_{\alpha_1}|^\ell \cdot |S_{\alpha_2}|^k}{q^{nk}}.$$

Indeed, in this case the value of $\mathcal{T}$ is negligible and hence, we obtain an upper bound identical to Kiltz et al. On the other hand, if we want to work over fully-splitting polynomial rings in order to apply the Number Theoretic Transform algorithm, we would have to increase q as well as the dimensions (k, ℓ) of the matrix $\mathbf{A}$ so that $\mathcal{T}(q, \ell, k, \alpha_1, \alpha_2)$ stays small. Unfortunately, this implies larger public key and signature size.

Key techniques. We provide an overview of the proof of Theorem 1.1. As mentioned before, we closely follow the proof strategy from [SS13]. For completeness, we recall the argument. Let d be the divisor of n such that

$$X^n + 1 \equiv \prod_{i=1}^{d} f_i(X) \pmod{q}$$

for distinct polynomials $f_i(X)$ of degree n/d that are irreducible in $\mathbb{Z}_q[X]$. In other words, $X^n + 1$ splits into d irreducible polynomials modulo q. The proof sketch goes as follows.

Step 1: We apply the union bound:

$$\begin{aligned}\Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}}&[\exists (\mathbf{z}_1, \mathbf{z}_2) \in S_{\alpha_1}^\ell \backslash \{\mathbf{0}\} \times S_{\alpha_2}^k : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}] \\ &\leq \sum_{(\mathbf{z}_1, \mathbf{z}_2) \in S_{\alpha_1}^\ell \backslash \{\mathbf{0}\} \times S_{\alpha_2}^k} \Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}].\end{aligned} \tag{2}$$

Step 2: We identify the subset Z of $S_{\alpha_1}^\ell \backslash \{\mathbf{0}\} \times S_{\alpha_2}^k$ which satisfies:

$$(\mathbf{z}_1, \mathbf{z}_2) \in Z \iff \Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}] > 0.$$

Hence, the probability in Equation (1) can be bounded by

$$\sum_{(\mathbf{z}_1, \mathbf{z}_2) \in Z} \Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}].$$

Step 3: Next, we propose a partitioning of the set Z into subsets $Z_0, Z_1, ..., Z_d$, i.e. $Z = \bigcup_{i=0}^{d} Z_i$. Then, we show that for each $(\mathbf{z}_1, \mathbf{z}_2) \in Z_i$, the probability

$$p_i := \Pr_{\mathbf{A} \leftarrow R_q^{k\times\ell}}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}]$$

is the same and we compute it. Thus, the probability in Equation (1) can now be bounded by:

$$\sum_{i=0}^{d} \sum_{(\mathbf{z}_1,\mathbf{z}_2)\in Z_i} p_i = \sum_{i=1}^{d} |Z_i| \cdot p_i$$

Step 4: We find an upper bound on $|Z_i|$.

Zero function. In this paper, we will consider zeroes in the "Chinese Remainder representation" [2] of polynomials in R_q. Formally, we define the following Zero function:

$$\mathsf{Zero}(y) := \{i : y \equiv 0 \pmod{(f_i(X), q)}\} \text{ and } \mathsf{Zero}(\mathbf{y}) := \bigcap_{j=1}^{k} \mathsf{Zero}(y_j),$$

where $y \in R_q$ and $\mathbf{y} = (y_1, ..., y_k) \in R_q^k$. Note that if y is invertible then $|\mathsf{Zero}(y)| = 0$. Lyubashevsky and Seiler [LS18] proved that whenever a non-zero y has small Euclidean norm then $|\mathsf{Zero}(y)| = 0$. Moreover, there is a relationship between the Euclidean norm of y and the size of set $\mathsf{Zero}(y)$ (see [SS13], proof of Lemma 3.2). In particular, the result implies that relatively small elements of R_q have only a few zeroes in the Chinese Remainder representation. This observation will be crucial for Steps 3 and 4.

Zero rows. Consider the equation $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}$ and let $j \in \mathsf{Zero}(\mathbf{z}_1)$. If we look at this equation modulo $(f_j(X), q)$ then we just end up with $\mathbf{z}_2 = \mathbf{0}$, i.e. $j \in \mathsf{Zero}(\mathbf{z}_2)$ and thus $j \in \mathsf{Zero}(\mathbf{z}_1||\mathbf{z}_2)$ where $||$ denotes usual concatenation of vectors. Consequently, $\mathsf{Zero}(\mathbf{z}_1) \subseteq \mathsf{Zero}(\mathbf{z}_1||\mathbf{z}_2)$. Clearly, we have $\mathsf{Zero}(\mathbf{z}_1||\mathbf{z}_2) \subseteq \mathsf{Zero}(\mathbf{z}_1)$ and therefore these two sets are equal. This implies that the subset Z introduced in Step 2 can be identified as:

$$Z = \{(\mathbf{z}_1, \mathbf{z}_2) : \mathsf{Zero}(\mathbf{z}_1) = \mathsf{Zero}(\mathbf{z}_1||\mathbf{z}_2)\}.$$

Define $Z_i = \{(\mathbf{z}_1, \mathbf{z}_2) : \mathsf{Zero}(\mathbf{z}_1) = \mathsf{Zero}(\mathbf{z}_1||\mathbf{z}_2) \wedge |\mathsf{Zero}(\mathbf{z}_1)| = i\} \subseteq Z$ (Step 3). Informally, we say that $(\mathbf{z}_1, \mathbf{z}_2) \in Z_i$ has i zero rows, since if we write down the components of $\mathbf{z}_1$ and $\mathbf{z}_2$ in the Chinese Remainder representation, in columns, then we get exactly i rows filled with zeroes.

For fixed $(\mathbf{z}_1, \mathbf{z}_2) \in Z_i$, we compute the probability p_i defined in Step 3 by counting the number of possible $\mathbf{A}$ which satisfy $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}$. This could be done by considering the equation modulo $(f_j(X), q)$ for all $j \notin \mathsf{Zero}(\mathbf{z}_1)$. Indeed, for such j there is a simple way to count all $\mathbf{A} \in (\mathbb{Z}_q[X]/(f_j(x)))^{k\times\ell}$ which satisfy $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}$ modulo $f_j(X)$. Concretely, one of the components of $\mathbf{z}_1$, say z_u, is going to be invertible modulo $f_j(X)$ and therefore all entries of $\mathbf{A}$ not related to z_u can be chosen arbitrarily. The rest, however, will be adjusted so that the equation holds. On the other hand, if $j \in \mathsf{Zero}(\mathbf{z}_1) = \mathsf{Zero}(\mathbf{z}_1||\mathbf{z}_2)$ then $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2$ is simply equal to $\mathbf{0}$ modulo $(f_j(X), q)$ for any $\mathbf{A}$. By applying Chinese Remainder Theorem, we obtain the total number of possible $\mathbf{A} \in R_q^{k\times\ell}$ which satisfy the equation above.

The only thing left is to provide an upper bound on $|Z_i|$ (Step 4). Firstly, we observe that if $(\mathbf{z}_1, \mathbf{z}_2) \in Z_i$ then clearly $|\mathsf{Zero}(z_j)| \geq i$ for $j = 1, ..., \ell$ where $\mathbf{z}_1 = (z_1, ..., z_\ell)$. Since each component of $\mathbf{z}_1 \in S_{\alpha_1}^{\ell} \backslash \{\mathbf{0}\}$ has infinity norm at most α_1, and assuming this value is relatively small, we get that each component of $\mathbf{z}_1$ has only a few zeroes in the Chinese Remainder representation (Lemma 3.2). Hence, for some larger values of i, we simply get $Z_i = \varnothing$. The second observation is that if $(\mathbf{z}_1, \mathbf{z}_2) \in Z_i$ and $\mathbf{y}_1, \mathbf{y}_2$ are vectors of some "small" polynomials then $(\mathbf{z}_1 + \mathbf{y}_1, \mathbf{z}_2 + \mathbf{y}_2)$ is likely not to have exactly i zero rows. For example, suppose that

$$\mathsf{Zero}(\mathbf{z}_1 + \mathbf{y}_1, \mathbf{z}_2 + \mathbf{y}_2) = \mathsf{Zero}(\mathbf{z}_1 + \mathbf{y}'_1, \mathbf{z}_2 + \mathbf{y}'_2)$$

for some other small $\mathbf{y}'_1, \mathbf{y}'_2$. This implies that $(\mathbf{y}_1 - \mathbf{y}'_1, \mathbf{y}_2 - \mathbf{y}'_2)$ has at least i zero rows. In particular, each component of $\mathbf{y}_1 - \mathbf{y}'_1$, say $\hat{y}_j$, has at least i zeroes in the Chinese Remainder representation. However,

[2] Alternatively, we call it "FFT/NTT representation" in the fully-splitting case.

we know that $\hat{y}_j$ is a polynomial of small norm by the choice of $\mathbf{y}_1$ and $\mathbf{y}'_1$. Therefore, $\hat{y}_j$ has only a few zeroes (by the observation above or Lemma 3.2). By picking sufficiently small $\mathbf{y}_1$ and $\mathbf{y}'_1$ we can make sure that each component $\hat{y}_j$ of $\mathbf{y}_1 - \mathbf{y}'_1$ has less than i zeroes. This would lead to a contradiction. In conclusion, our approach for bounding $|Z_i|$ is to, for each $(\mathbf{z}_1, \mathbf{z}_2) \in Z_i$, generate all pairs of form $(\mathbf{z}_1 + \mathbf{y}_1, \mathbf{z}_2 + \mathbf{y}_2) \notin Z_i$, for vectors of sufficiently small polynomials $\mathbf{y}_1, \mathbf{y}_2$, and applying the pigeonhole principle along with other simple counting arguments.

1.2 Applications

Digital signatures. Kiltz et al. [KLS18] presented a generic framework for constructing secure Fiat-Shamir signatures in the quantum random oracle model (QROM). As a concrete instantiation, they introduced a new signature scheme Dilithium-QROM, which is a modification of the original Dilithium scheme [LDK+17], and is tightly based on the hardness of Module-LWE problem in the QROM. However, in order to obtain security of Dilithium-QROM, Kiltz et al. choose the prime modulus q to be congruent to 5 modulo 8. This assumption assures that the underlying polynomial ring $\mathbb{Z}_q[X]/(X^n + 1)$ splits into two subrings modulo q and invertibility results can be applied [LN17, LS18]. Unfortunately, polynomial multiplication algorithms in such rings are not efficient. We show how to apply our probability results to the security of Dilithium-QROM so that one can avoid such special assumptions on q (in particular, one could choose q so that R_q splits completely and NTT along with other optimisations can be applied). The only disadvantage is that, in order to keep the probabilities small, one should slightly increase the size of q and dimensions (k, ℓ). Unfortunately, this results in having both considerably larger public keys and signatures.

General results by Kiltz et al. can also be applied to obtain a security proof in the QROM for a number of existing Fiat-Shamir signature schemes similar to Dilithium such as the Bai-Galbraith scheme [BG14] (see Section 4) or qTESLA [BAA+17]. So far, security of the latter scheme in the quantum random oracle model is proven assuming a certain non-standard conjecture. However, one can also obtain it by applying the framework by Kiltz et al. and using our probability upper bounds. Consequently, one gets a tightly secure version of qTESLA in the QROM without any non-standard conjecture. We recall that our results allow this signature scheme to work over fully-splitting rings so that the use of NTT for polynomial multiplication is possible. However, as in the case of Dilithium-QROM, we would end up with larger public key and signature size compared to the original qTESLA (see Table 2).

Commitment schemes. Recently, Baum et al. [BDL+18] presented efficient commitment schemes from Module-SIS and Module-LWE. However, both their new statistically binding commitment scheme and their improved construction from [BKLP15] rely on the general invertibility result from [LS18], i.e. special congruence conditions on the prime modulus q. Our probability upper bounds can be applied to prove the statistically binding property of these constructions, and consequently, one could now consider working in fully-splitting rings. As before, we observe that choosing primes q such that $X^n + 1$ splits into many factors modulo q results in having both larger commitment and proof size.

1.3 Related Works

The first asymptotically-efficient lattice-based signature scheme using the "Fiat-Shamir with Aborts" paradigm was presented in [Lyu09] which is based on the Ring-SIS problem. Later on, Lyubashevsky [Lyu12] improved the scheme by basing it on the combination of Ring-SIS and Ring-LWE. Since then, many substantial improvements have been proposed [GLP12, BG14, LDK+17, BAA+17]. In the meantime, lossy identification schemes were introduced and used to construct secure digital signatures in the quantum random oracle model [AFLT12, Unr17, ABB+17, KLS18].

Invertibility of "small" polynomials [3] is an important property in the context of (approximate) zero-knowledge proofs based on lattices. For example, one usually needs the difference set $\mathcal{C} - \mathcal{C}$ to contain only invertible polynomials for extraction purposes [SSTX09, BKLP15] where $\mathcal{C}$ is a challenge set. Lyubashevsky and Neven [LN17] proved that if q is congruent to 5 modulo 8 then the polynomial ring $R_q = \mathbb{Z}_q[X]/(X^n + 1)$ splits into two subrings and elements of small infinity norm are indeed invertible. This result was generalised by Lyubashevsky and Seiler [LS18]. Concretely, they showed that if $q \equiv 2k + 1 \pmod{4k}$ for some k then $X^n + 1$ splits into k irreducible polynomials modulo q

[3] What we mean by "small" is that the polynomial has small infinity or Euclidean norm.

and also small elements in R_q are invertible. These results have been recently applied in the context of computing probabilities related to the security of lattice-based signatures and commitment schemes, e.g. [KLS18, BDL+18].

Existence of short vectors in random module lattices was first investigated independently by Lyubashevsky et al. [LPR13] and Stehlé and Steinfeld [SS13] in the context of regularity bounds for cyclotomic rings. Later on, these results were extended by Langlois and Stehlé [LS15] (to modules rather than rings), and more recently by Rosca et al. [RSW18] to all rings.

1.4 History

This paper appeared at Asiacrypt 2019 [Ngu19] and this is the full version. It contains additional background (Appendix A) as well as applications to Dilithium-QROM (Appendix B). Moreover, since publishing our paper, we were not fully aware that almost identical results have already been shown in [SS13, LS15] and thus we treated the content in Section 3 as our independent contribution. In this version, we give appropriate credit to the authors.

2 Preliminaries

For $n \in \mathbb{N}$, let $[n] := \{1, \ldots, n\}$. For a set S, $|S|$ is the cardinality of S, $\mathcal{P}(S)$ is the power set of S and $\mathcal{P}_i(S)$ is the set of all subsets of S of size i. If S is finite, we denote the sampling of a uniform random element x by $x \leftarrow S$, while we denote the sampling according to some distribution $\mathfrak{D}$ by $x \leftarrow \mathfrak{D}$. By $[\![B]\!]$ we denote the bit that is 1 if the Boolean statement B is true, and 0 otherwise.

Algorithms. Unless stated otherwise, we assume all our algorithms to be probabilistic. We denote by $y \leftarrow \mathsf{A}(x)$ the probabilistic computation of algorithm A on input x. If A is deterministic, we write $y := \mathsf{A}(x)$. The notation $y \in \mathsf{A}(x)$ is used to indicate all possible outcomes y of the probabilistic algorithm A on input x. We can make any probabilistic A deterministic by running it with fixed randomness. We write $y := \mathsf{A}(x; r)$ to indicate that A is run on input x with randomness r. The notation $\mathsf{A}(x) \Rightarrow y$ denotes the event that A on input x returns y. Eventually, we write $\mathrm{Time}(\mathsf{A})$ to denote the running time of A.

2.1 Cyclotomic Rings

Let n be a power of two. Denote R and R_q respectively to be the rings $\mathbb{Z}[X]/(X^n+1)$ and $\mathbb{Z}_q[X]/(X^n+1)$, for a prime q. We also set d to be the divisor of n such that

$$X^n + 1 \equiv \prod_{i=1}^{d} f_i(X) \pmod q$$

for distinct polynomials $f_i(X)$ of degree n/d that are irreducible in $\mathbb{Z}_q[X]$. Alternatively, we say that $X^n + 1$ splits into d polynomials modulo q. If $d = n$ then $X^n + 1$ *fully splits*. By default, all the equalities and congruences between ring elements in this paper are modulo q.

Regular font letters denote elements in R or R_q and bold lower-case letters represent column vectors with coefficients in R or R_q. Bold upper-case letters denote matrices. By default, all vectors are column vectors.

Modular reductions. For an even (resp. odd) positive integer α, we define $r' = r \bmod^{\pm} \alpha$ to be the unique element r' in the range $-\frac{\alpha}{2} < r' \leq \frac{\alpha}{2}$ (resp. $-\frac{\alpha-1}{2} \leq r' \leq \frac{\alpha-1}{2}$) such that $r' = r \bmod \alpha$. For any positive integer α, we define $r' = r \bmod^{+} \alpha$ to be the unique element r' in the range $0 \leq r' < \alpha$ such that $r' = r \bmod \alpha$. When the exact representation is not important, we simply write $r \bmod \alpha$.

Sizes of elements. For an element $w \in \mathbb{Z}_q$, we write $\|w\|_\infty$ to mean $|w \bmod^{\pm} q|$. Define the ℓ_∞ and ℓ_2 norms for $w = w_0 + w_1 X + \ldots + w_{n-1} X^{n-1} \in R$ as follows:

$$\|w\|_\infty = \max_i \|w_i\|_\infty, \quad \|w\| = \sqrt{\|w_0\|_\infty^2 + \ldots + \|w_{n-1}\|_\infty^2}.$$

Similarly, for $\mathbf{w} = (w_1, \ldots, w_k) \in R^k$, we define

$$\|\mathbf{w}\|_\infty = \max_i \|w_i\|_\infty, \quad \|\mathbf{w}\| = \sqrt{\|w_1\|^2 + \ldots + \|w_k\|^2}.$$

For a finite set $S \subseteq R^k$, however, we set

$$\|S\|_\infty = \max_{\mathbf{w}\in S} \|\mathbf{w}\|_\infty, \quad \|S\| = \max_{\mathbf{w}\in S} \|\mathbf{w}\|.$$

We write S_η to denote all elements $w \in R$ such that $\|w\|_\infty \leq \eta$.

Extracting high-order and low-order bits. To reduce the size of the public key, we need some algorithms that extract "higher-order" and "lower-order" bits of elements in $\mathbb{Z}_q$. The goal is that when given an arbitrary element $r \in \mathbb{Z}_q$ and another small element $z \in \mathbb{Z}_q$, we would like to be able to recover the higher order bits of $r + z$ without needing to store z. The algorithms are exactly as in [DLL+17, KLS18], and we repeat them for completeness in Figure 2. They are described as working on integers modulo q, but one can extend it to polynomials in R_q by simply being applied individually to each coefficient.

$\mathsf{Power2Round}_q(r, \delta)$

- $r := r \bmod^+ q$
- $r_0 := r \bmod^\pm 2^\delta$
- **return** $(r - r_0)/2^\delta$

$\mathsf{UseHint}_q(h, r, \alpha)$

- $m := (q-1)/\alpha$
- $(r_1, r_0) := \mathsf{Decompose}_q(r, \alpha)$
- **if** $h = 1$ and $r_0 > 0$ **return** $(r_1 + 1) \bmod^+ m$
- **if** $h = 1$ and $r_0 \leq 0$ **return** $(r_1 - 1) \bmod^+ m$
- **return** r_1

$\mathsf{MakeHint}_q(z, r, \alpha)$

- $r_1 := \mathsf{HighBits}_q(r, \alpha)$
- $v_1 := \mathsf{HighBits}_q(r + z, \alpha)$
- **return** $[\![r_1 \neq v_1]\!]$

$\mathsf{Decompose}_q(r, \alpha)$

- $r := r \bmod^+ q$
- $r_0 := r \bmod^\pm \alpha$
- **if** $r - r_0 = q - 1$
- **then** $r_1 := 0$; $r_0 := r_0 - 1$
- **else** $r_1 := (r - r_0)/\alpha$
- **return** (r_1, r_0)

$\mathsf{HighBits}_q(r, \alpha)$

- $(r_1, r_0) := \mathsf{Decompose}_q(r, \alpha)$
- **return** r_1

$\mathsf{LowBits}_q(r, \alpha)$

- $(r_1, r_0) := \mathsf{Decompose}_q(r, \alpha)$
- **return** r_0

Figure 2: Supporting algorithms for Dilithium and Dilithium-QROM [KLS18].

Lemma 2.1 *Suppose that q and α are positive integers satisfying $q > 2\alpha$, $q \equiv 1 \pmod{\alpha}$ and α even. Let $\mathbf{r}$ and $\mathbf{z}$ be vectors of elements in R_q where $\|\mathbf{z}\|_\infty \leq \alpha/2$, and let $\mathbf{h}, \mathbf{h}'$ be vectors of bits. Then the $\mathsf{HighBits}_q$, $\mathsf{MakeHint}_q$, and $\mathsf{UseHint}_q$ algorithms satisfy the following properties:*

1. $\mathsf{UseHint}_q(\mathsf{MakeHint}_q(\mathbf{z}, \mathbf{r}, \alpha), \mathbf{r}, \alpha) = \mathsf{HighBits}_q(\mathbf{r} + \mathbf{z}, \alpha)$.
2. *Let* $\mathbf{v}_1 = \mathsf{UseHint}_q(\mathbf{h}, \mathbf{r}, \alpha)$. *Then* $\|\mathbf{r} - \mathbf{v}_1 \cdot \alpha\|_\infty \leq \alpha + 1$.
3. *For any* $\mathbf{h}, \mathbf{h}'$, *if* $\mathsf{UseHint}_q(\mathbf{h}, \mathbf{r}, \alpha) = \mathsf{UseHint}_q(\mathbf{h}', \mathbf{r}, \alpha)$, *then* $\mathbf{h} = \mathbf{h}'$.

Lemma 2.2 *If $\|\mathbf{s}\|_\infty \leq \beta$ and $\|\mathsf{LowBits}_q(\mathbf{r}, \alpha)\|_\infty < \alpha/2 - \beta$, then*

$$\mathsf{HighBits}_q(\mathbf{r}, \alpha) = \mathsf{HighBits}_q(\mathbf{r} + \mathbf{s}, \alpha).$$

Ideal lattices. An integer lattice of dimension n is an additive subgroup of $\mathbb{Z}^n$. For simplicity, we only consider full-rank lattices. The determinant of a full-rank lattice Λ of dimension n is equal to the size of the quotient group $\mathbb{Z}^n/\Lambda$. We denote $\lambda_1(\Lambda) = \min_{\|\mathbf{w}\|\in\Lambda} \|\mathbf{w}\|$. We say that Λ is an *ideal lattice* in R if Λ is an ideal of R. There exists a lower bound on $\lambda_1(\Lambda)$ if Λ is an ideal lattice [LS18, PR07]. Assuming that n is a power of two, we get a simplified bound.

Lemma 2.3 ([LS18], Lemma 2.7). *If Λ is an ideal lattice in R, then $\lambda_1(\Lambda) \geq \det(\Lambda)^{1/n}$.*

The MLWE Assumption. For integers m, k, and a probability distribution $D : R_q \to [0, 1]$, we say that the advantage of algorithm A in solving the decisional $\mathsf{MLWE}_{m,k,D}$ problem over the ring R_q is

$$\begin{aligned} \mathrm{Adv}^{\mathsf{MLWE}}_{m,k,D} := \big| &\Pr[\mathsf{A}(\mathbf{A}, \mathbf{t}) \Rightarrow 1 \mid \mathbf{A} \leftarrow R_q^{m\times k}; \mathbf{t} \leftarrow R_q^m] \\ &- \Pr[\mathsf{A}(\mathbf{A}, \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2) \Rightarrow 1 \mid \mathbf{A} \leftarrow R_q^{m\times k}; \mathbf{s}_1 \leftarrow D^k; \mathbf{s}_2 \leftarrow D^m] \big| . \end{aligned}$$

The MLWE assumption states that the above advantage is negligible for all polynomial-time algorithms A. It was introduced in [LS15], and is a generalization of the LWE assumption from [Reg05]. The Ring-LWE assumption [LPR10] is a special case of MLWE where $k = 1$. Analogously to LWE and Ring-LWE, it was shown in [LS15] that solving the MLWE problem for certain parameters is as hard as solving certain worst-case problems in certain algebraic lattices.

3 Zeroes in the Chinese Remainder Representation

In this section, we present general results about existence of solutions $(\mathbf{A}, \mathbf{t}) \in R_q^{k\times\ell} \times R^k$ to the equation $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}$ (and other similar ones), for some $\mathbf{z}_1 \in R_q^\ell, \mathbf{z}_2 \in R_q^k, c \in R_q\backslash\{0\}$, and compute the probability of satisfying such equations for uniformly random $\mathbf{A}$ and $\mathbf{t}$. The results are crucial for security analysis of Fiat-Shamir signature schemes. For instance, security of Dilithium-QROM [KLS18] relies heavily on the assumption that c is invertible in R_q or $\mathbf{z}_1$ contains an invertible component. In such a case, the probability can be calculated straightforwardly. Hence, q is chosen so that $q \equiv 5 \pmod 8$ because then, polynomials in R_q of small (infinity) norm are proved to be invertible [LS18, LN17]. We circumvent these assumptions by following proof techniques from Stehlé and Steinfeld [SS13]. More specifically, we analyse "zeroes in the Chinese Remainder Representation" of $\mathbf{z}_1, \mathbf{z}_2$ and c in order to provide general upper bounds on the probabilities.

3.1 Zero Rows

We start by introducing the Zero function.

Definition 3.1 Let $y \in R_q$. We define a set

$$\mathsf{Zero}(y) := \{i \in [d] : y \equiv 0 \pmod{f_i(X)}\}.$$

For a vector $\mathbf{y} = (y_1, ..., y_k) \in R_q^k$, we set $\mathsf{Zero}(\mathbf{y}) := \bigcap_{j=1}^k \mathsf{Zero}(y_j)$ and similarly for multiple vectors $\mathbf{y}_1, ..., \mathbf{y}_\ell$ over R_q, $\mathsf{Zero}(\mathbf{y}_1, ..., \mathbf{y}_\ell) := \bigcap_{j=1}^\ell \mathsf{Zero}(\mathbf{y}_j)$.

Informally, we say that y has i zeroes in the Chinese Remainder Representation if $|\mathsf{Zero}(y)| = i$. One observes that $\mathsf{Zero}(y) = \varnothing$ if and only if y is invertible, by the Chinese Remainder Theorem. Also, $\mathsf{Zero}(y) = [d] \iff y = 0$.

Lyubashevsky and Seiler [LS18] showed that if $\|y\| < q^{1/d}$ then y is invertible. Obviously, it is not very interesting if d is large (e.g. $d = n$). Here, we consider the more general result wihch takes into account the number of zeroes in the Chinese Remainder Representation. We note that this was already shown in the proof of Lemma 3.2 in [SS13].

Lemma 3.2 *Let $y \in R_q$ such that $0 < ||y|| < q^{m/d}$ for some $m \in [d]$. Then, $|\mathsf{Zero}(y)| < m$.*

Proof. Suppose that $|\mathsf{Zero}(y)| \geq m$ and pick any $i_1, ..., i_m \in \mathsf{Zero}(y)$. Define the following set:

$$\Lambda = \{z \in R : \forall j \in [m], z \equiv 0 \pmod{f_{i_j}(X)}\}.$$

Firstly, note that Λ is an additive group and $y \in \Lambda$. Moreover, for any $z \in \Lambda$, we have $z \cdot X \in \Lambda$ since each $f_{i_j}(X)$ is a factor of $X^n + 1$ modulo q. Therefore, Λ is an ideal of R, and hence an ideal lattice in the ring R. Consider the Chinese Remainder representation modulo q of all the elements in Λ. Note that they have 0 in the coefficients corresponding to $f_{i_j}(X)$ for $j \in \{1, ..., m\}$ and arbitrary values everywhere else. This implies that $\det(\Lambda) = |\mathbb{Z}^n/\Lambda| = q^{nm/d}$. Hence, by Lemma 2.3 we have $\lambda_1(\Lambda) \geq q^{m/d}$. However, we know that $||y|| > 0$, thus y is non-zero. Eventually, we obtain $||y|| < q^{m/d} \leq \lambda_1(\Lambda) \leq ||y||$ which leads to contradiction. □

The lemma above implies that if a polynomial $y \in R_q$ is short enough, then it has only a few zeroes in the Chinese Remainder Representation (but is not necessarily invertible).

We now introduce the notion of ZeroRows which will be crucial in proving the main theorem.

Definition 3.3 Let $k \in \mathbb{N}$ and $A \subseteq R_q^k$ be a non-empty set. Then, we write $\mathsf{ZeroRows}_i(A)$ to denote

$$\mathsf{ZeroRows}_i(A) := \{\mathbf{a} \in A : |\mathsf{Zero}(\mathbf{a})| = i\}.$$

We say that $\mathbf{a} \in \mathsf{ZeroRows}_i(A)$ has i zero rows.

Name ZeroRows comes from the fact that if $\mathbf{a} = (a_1, ..., a_k) \in \mathsf{ZeroRows}_i(A)$ and if we write down the Chinese Remainder Representation of $a_1, ..., a_k$ as column vectors [4] then we get exactly i rows filled only with zeroes.

The next result gives an upper bound on $\mathsf{ZeroRows}_i(S_\alpha^k)$ for fixed $i > 0, k$ and α. The key idea of the proof is as follows. For simplicity, consider $z' := z + X^j, z'' := z + X^\ell$ for some distinct $j, \ell \in [2n]$ and $z \in R_q$. To begin with, note that $\mathsf{Zero}(z') \cap \mathsf{Zero}(z'') = \varnothing$. Indeed, if there exists some $u \in \mathsf{Zero}(z') \cap \mathsf{Zero}(z'')$ then

$$z + X^\ell \equiv z'' \equiv 0 \equiv z' \equiv z + X^j \pmod{f_u(X)}.$$

Hence, we get a contradiction, since $X^j - X^\ell$ is invertible [BCK+14]. Therefore,

$$|\{z + X^j \in \mathsf{ZeroRows}_i(S_\alpha) : j \in [2n]\}| \leq \lfloor d/i \rfloor.$$

This is because if size of the set is strictly larger than d/i then, by definition of $\mathsf{ZeroRows}_i(S_\alpha)$ and the pigeonhole principle, we would have $\mathsf{Zero}(z + X^j) \cap \mathsf{Zero}(z + X^\ell) \neq \varnothing$ for some distinct j, ℓ. Thus, we end up with:

$$|\{z + X^j \notin \mathsf{ZeroRows}_i(S_\alpha) : j \in [2n]\}| \geq 2n - \lfloor d/i \rfloor.$$

Our main strategy is that for each $\mathbf{z} \in \mathsf{ZeroRows}_i(S_\alpha^k)$, we count all $\mathbf{z}'$ of form $\mathbf{z} + \mathbf{y}$ (where $\mathbf{y}$ is a somewhat small polynomial) such that $\mathbf{z}' \notin \mathsf{ZeroRows}_i(S_\alpha^k)$ similarly as above, and eventually, obtain an upper bound on $|\mathsf{ZeroRows}_i(S_\alpha^k)|$. The bound depends on the size of a set $W_i \subseteq R_q$, which satisfies the following property: for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u - v)| < i$ [5]. Later on, we show how to use our previous result, i.e. Lemma 3.2, to construct such sets.

Lemma 3.4 *Let $k, \alpha \in \mathbb{N}, i \in [d]$ and $W_i \subseteq R_q$ be a set of polynomials in R_q such that for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u - v)| < i$. Then,*

$$|\mathsf{ZeroRows}_i(S_\alpha^k)| \leq \frac{\binom{d}{i} \cdot |S_{\alpha + \|W_i\|_\infty}|^k}{|W_i|^k}.$$

Proof. Firstly, take any $\mathbf{z} = (z_1, ..., z_k) \in S_\alpha^k$ and define

$$\mathsf{Bad}(z_1, ..., z_k) := \{(z_1 + y_1, ..., z_k + y_k) \in \mathsf{ZeroRows}_i(S_\alpha^k) : y_1, ..., y_k \in W_i\}.$$

We claim that $|\mathsf{Bad}(z_1, ..., z_k)| \leq \binom{d}{i}$. Indeed, suppose $|\mathsf{Bad}(z_1, ..., z_k)| > \binom{d}{i}$ and define the function

$$F : \mathsf{Bad}(z_1, ..., z_k) \to \mathcal{P}_i([d]), (z_1', ..., z_k') \longmapsto \mathsf{Zero}(z_1', ..., z_k').$$

Note that F is well-defined by definition of Bad. Also, $|\mathsf{Bad}(z_1, ..., z_k)| > \binom{d}{i} = |\mathcal{P}_i([d])|$ implies that F is not injective. Hence,

$$F(z_1 + y_1, ..., z_k + y_k) = I = F(z_1 + y_1', ..., z_k + y_k')$$

for some set $I \in \mathcal{P}_i([d])$, $y_1, ..., y_k, y_1', ..., y_k' \in W_i$ and $y_j \neq y_j'$ for some index j. Take any $u \in I$. Then, $z_j + y_j \equiv 0 \equiv z_j + y_j' \pmod{f_u(X)}$ and consequently, $y_j - y_j' \equiv 0 \pmod{f_u(X)}$. Since we picked arbitrary $u \in I$, we proved that $|\mathsf{Zero}(y_j - y_j')| \geq i$. However, this leads to a contradiction by the definition of the set W_i.

Now, define a set

$$\mathsf{Good}(z_1, ..., z_k) := \{(z_1 + y_1, ..., z_k + y_k) \notin \mathsf{ZeroRows}_i(S_\alpha^k) : y_1, ..., y_k \in W_i\}.$$

[4] Namely, for each a_i we define a corresponding column vector $(a'_{i,1}, ..., a'_{i,d})$, where $a'_{i,j}$ is the element of $\mathbb{Z}_q[X]/(f_j(X))$, such that $a_i \equiv a'_{i,j} \pmod{f_j(X)}$, for $j \in [d]$.

[5] In the example above, W_1 is represented by the set $\{X^j : j \in [2n]\}$. Indeed, $|\mathsf{Zero}(X^j - X^k)| < 1$ for all distinct j, k.

Clearly, $|\mathsf{Good}(z_1,...,z_k)| = |W_i|^k - |\mathsf{Bad}(z_1,...,z_k)| \geq |W_i|^k - \binom{d}{i}$. Consider the following set

$$S = \bigcup_{(z_1,...,z_k)\in \mathsf{ZeroRows}_i(S_\alpha^k)} \mathsf{Good}(z_1,...,z_k).$$

One observes that $S \subseteq S^k_{\alpha+\|W_i\|_\infty} \backslash \mathsf{ZeroRows}_i(S_\alpha^k)$ by definition of Good, which gives us an upper bound on $|S|$. We are now interested in finding a lower bound for $|S|$. Let $(\hat{z}_1,...,\hat{z}_k)$ be an element of S and denote

$$\mathsf{COUNT}(\hat{z}_1,...,\hat{z}_k) := \{(z_1,...,z_k) \in \mathsf{ZeroRows}_i(S_\alpha^k) : (\hat{z}_1,...,\hat{z}_k) \in \mathsf{Good}(z_1,...,z_k)\}.$$

We claim that $|\mathsf{COUNT}(\hat{z}_1,...,\hat{z}_k)| \leq \binom{d}{i}$. Informally, this means that $(\hat{z}_1,...,\hat{z}_k)$ belongs to at most $\binom{d}{i}$ "good" sets (out of $|\mathsf{ZeroRows}_i(S_\alpha^k)|$). Just like before, assume that $|\mathsf{COUNT}(\hat{z}_1,...,\hat{z}_k)| > \binom{d}{i}$ and define a function

$$F : \mathsf{COUNT}(\hat{z}_1,...,\hat{z}_k) \to \mathcal{P}_i([d]), (z_1,...,z_k) \longmapsto \mathsf{Zero}(z_1,...,z_k).$$

Then,

$$F(z_1,...,z_k) = I = F(z_1',...,z_k')$$

for some set $I \in \mathcal{P}_i([d])$ and $z_1,...,z_k,z_1',...,z_k' \in S_\alpha$ such that there exists an index j which satisfies $z_j \neq z_j'$. Since $(\hat{z}_1,...,\hat{z}_k) \in \mathsf{Good}(z_1,...,z_k)$ and $(\hat{z}_1,...,\hat{z}_k) \in \mathsf{Good}(z_1',...,z_k')$, we have that $z_j + y_j = \hat{z}_j = z_j' + y_j'$ for some distinct $y_j, y_j' \in W_i$. Take any $u \in I$ and note that $z_j \equiv 0 \equiv z_j' \pmod{f_u(X)}$. Therefore,

$$y_j \equiv \hat{z}_j - z_j \equiv \hat{z}_j \equiv \hat{z}_j - z_j' \equiv y_j' \pmod{f_u(X)}.$$

Hence, $|\mathsf{Zero}(y_j - y_j')| \geq i$. Similarly as before, we observe that this leads to a contradiction by the definition of W_i. Thus, $|\mathsf{COUNT}(\hat{z}_1,...,\hat{z}_k)| \leq \binom{d}{i}$. This implies:

$$|S| \geq \frac{\sum_{\mathbf{z}\in \mathsf{ZeroRows}_i(S_\alpha^k)} |\mathsf{Good}(\mathbf{z})|}{\binom{d}{i}} \geq \frac{\sum_{\mathbf{z}\in \mathsf{ZeroRows}_i(S_\alpha^k)} |W_i|^k - \binom{d}{i}}{\binom{d}{i}}$$

Combining the lower bound as well as the upper bound for $|S|$ we get:

$$\begin{aligned} |S_{\alpha+\|W_i\|_\infty}|^k - |\mathsf{ZeroRows}_i(S_\alpha^k)| &\geq |S| \\ &\geq \frac{1}{\binom{d}{i}} |\mathsf{ZeroRows}_i(S_\alpha^k)| \cdot |W_i|^k - |\mathsf{ZeroRows}_i(S_\alpha^k)|. \end{aligned} \tag{3}$$

Therefore, $|\mathsf{ZeroRows}_i(S_\alpha^k)| \leq \frac{\binom{d}{i}\cdot|S_{\alpha+\|W_i\|_\infty}|^k}{|W_i|^k}$. □

We point out that the proof does not work for $i = 0$. In this case, we can use the obvious upper bound: $\mathsf{ZeroRows}_0(S_\alpha^k) \leq |S_\alpha^k|$.

Consider again the equation $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}$, where $\mathbf{A}$ and $\mathbf{t}$ are variables, and denote $\mathbf{A} = (a_{i,j})$ and $\mathbf{t} = (t_1,...,t_k)$. Clearly, we have $\mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2) \subseteq \mathsf{Zero}(\mathbf{z}_1, c)$. Suppose that $\mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2) \neq \mathsf{Zero}(\mathbf{z}_1, c)$. If we write $\mathbf{z}_1 = (z_1,...,z_\ell)$ and $\mathbf{z}_2 = (z_1',...,z_k')$, then there exist some i, j such that $i \in \mathsf{Zero}(\mathbf{z}_1, c)$ and $i \notin \mathsf{Zero}(\mathbf{z}_1, c, z_j')$. Note that

$$\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t} \implies a_{j,1}z_1 + ... + a_{j,\ell}z_\ell + z_j' = ct_j.$$

However,

$$0 \not\equiv z_j' \equiv a_{j,1}z_1 + ... + a_{j,\ell}z_\ell + z_j' \equiv ct_j \equiv 0 \pmod{f_i(X)},$$

which leads to a contradiction. Therefore, if $\mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2) \neq \mathsf{Zero}(\mathbf{z}_1, c)$ then we end up with no solutions. This motivates us to extend the $\mathsf{ZeroRows}$ function as follows.

Definition 3.5 Let $k \in \mathbb{N}$ and $A \subseteq R_q^k, B \subseteq R_q^\ell$ be non-empty sets. Then, we define $\mathsf{ZeroRows}_i(A; B)$ to be

$$\mathsf{ZeroRows}_i(A; B) := \{(\mathbf{a}, \mathbf{b}) \in A \times B : \mathsf{Zero}(\mathbf{a}, \mathbf{b}) = \mathsf{Zero}(\mathbf{a}) \wedge |\mathsf{Zero}(\mathbf{a})| = i\}.$$

Sometimes, we write $\mathsf{ZeroRows}_i(A_1, A_2; B)$ to denote $\mathsf{ZeroRows}_i(\bar{A}; B)$, where $\bar{A} = A_1 \times A_2$.

Using the same techniques as before, one can prove a similar result to Lemma 3.4 which is related to the modified $\mathsf{ZeroRows}$ function.

Lemma 3.6 *Let $k, \ell, \alpha_1, \alpha_2 \in \mathbb{N}, i \in [d]$ and $W_i \subseteq R_q$ be a set of polynomials in R_q such that for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u-v)| < i$. Take any set $\mathcal{D} \subseteq R_q \backslash \{\mathbf{0}\}$ and define e to be the largest integer which satisfies $\|\mathcal{D}\| \geq q^{e/d}$. Then,*

$$|\mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})| \leq \frac{\binom{e}{i} \cdot |S_{\alpha_1+\|W_i\|_\infty}|^\ell \cdot |S_{\alpha_2+\|W_i\|_\infty}|^k \cdot |\mathcal{D}|}{|W_i|^{\ell+k}}.$$

Proof. Since we follow the same strategy as in the proof of Lemma 3.4, we only provide a proof sketch. To begin with, take any $\mathbf{z}_1 = (z_1, ..., z_\ell) \in S^\ell_{\alpha_1}, c \in \mathcal{D}, \mathbf{z}_2 = (z'_1, ..., z'_k) \in S^k_{\alpha_2}$ and define

$$\mathsf{Bad}(\mathbf{z}_1, c, \mathbf{z}_2) := \{(\mathbf{z}_1 + \mathbf{y}, c, \mathbf{z}_2 + \mathbf{y}') \in \mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2}) : \mathbf{y} \in W^\ell_i, \mathbf{y}' \in W^k_i\}.$$

We point out that c stays still. Using the same technique as before, one can prove that $|\mathsf{Bad}(\mathbf{z}_1, c, \mathbf{z}_2)| \leq \binom{e}{i}$. Informally, this is because we only consider all subsets of $\mathsf{Zero}(c)$ (instead of $[d]$ like last time) of size i and c has at most e zeroes in the Chinese Remainder Representation (Lemma 3.2).

Now, we define a set

$$\mathsf{Good}(\mathbf{z}_1, c, \mathbf{z}_2) := \{(\mathbf{z}_1 + \mathbf{y}, c, \mathbf{z}_2 + \mathbf{y}') \notin \mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2}) : \mathbf{y} \in W^\ell_i, \mathbf{y}' \in W^k_i\}.$$

As before, we have $|\mathsf{Good}(\mathbf{z}_1, c, \mathbf{z}_2)| = |W_i|^{\ell+k} - |\mathsf{Bad}(\mathbf{z}_1, c, \mathbf{z}_2)| \geq |W_i|^{\ell+k} - \binom{e}{i}$. Consider the following set

$$S = \bigcup_{(\mathbf{z}_1, c, \mathbf{z}_2) \in \mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})} \mathsf{Good}(\mathbf{z}_1, c, \mathbf{z}_2).$$

We have that

$$S \subseteq S^\ell_{\alpha_1+\|W_i\|_\infty} \times \mathcal{D} \times S^k_{\alpha_2+\|W_i\|_\infty} \backslash \mathsf{ZeroRows}_i(S^k_\alpha)$$

by definition of Good. Let $(\mathbf{z}'_1, c, \mathbf{z}'_2)$ be an element of S and denote

$$\mathsf{COUNT}(\mathbf{z}'_1, c, \mathbf{z}'_2) := \{(\mathbf{z}_1, c, \mathbf{z}_2) \in \mathsf{ZeroRows}_i(S^k_\alpha) : (\mathbf{z}'_1, c, \mathbf{z}'_2) \in \mathsf{Good}(\mathbf{z}_1, c, \mathbf{z}_2)\}.$$

Similarly as before, we can show that $|\mathsf{COUNT}(\hat{z}_1, ..., \hat{z}_k)| \leq \binom{e}{i}$. Hence, we get:

$$\begin{aligned} |S| &\geq \frac{\sum_{(\mathbf{z}_1, c, \mathbf{z}_2) \in \mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})} |\mathsf{Good}(\mathbf{z}_1, c, \mathbf{z}_2)|}{\binom{e}{i}} \\ &\geq \frac{\sum_{(\mathbf{z}_1, c, \mathbf{z}_2) \in \mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})} |W_i|^{\ell+k} - \binom{e}{i}}{\binom{e}{i}} \end{aligned} \tag{4}$$

Combining the lower bound as well as upper bound for $|S|$ we get:

$$|S_{\alpha_1+\|W_i\|_\infty}|^\ell \cdot |\mathcal{D}| \cdot |S_{\alpha_2+\|W_i\|_\infty}|^k - |\mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})| \geq |S|,$$

and

$$|S| \geq \frac{1}{\binom{e}{i}} |\mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})| \cdot |W_i|^{\ell+k} - |\mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})|.$$

Therefore, $|\mathsf{ZeroRows}_i(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})| \leq \frac{\binom{e}{i} \cdot |S_{\alpha_1+\|W_i\|_\infty}|^\ell \cdot |S_{\alpha_2+\|W_i\|_\infty}|^k \cdot |\mathcal{D}|}{|W_i|^{\ell+k}}$. □

Again, we note that the lemma does not hold for $i = 0$. In this case, we use a simple bound: $|\mathsf{ZeroRows}_0(S^\ell_{\alpha_1}, \mathcal{D}; S^k_{\alpha_2})| \leq |S_{\alpha_1}|^\ell \cdot |S_{\alpha_2}|^k \cdot |\mathcal{D}|$.

In Lemma 3.6 we have an additional condition $0 \notin \mathcal{D}$. This is because otherwise we cannot define the integer e. Recall that e represents the maximal number of zeroes in the Chinese Remainder Representation that an element in $\mathcal{D}$ can have. Hence, in case $\mathcal{D} = \{0\}$, we can simply set $e = d$ and follow the strategy as in Lemma 3.6. Thus, we end up with the following corollary.

Corollary 3.7 *Let $k, \ell, \alpha_1, \alpha_2 \in \mathbb{N}, i \in [d]$ and $W_i \subseteq R_q$ be a set of polynomials in R_q such that for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u-v)| < i$. Then,*

$$|\mathsf{ZeroRows}_i(S^\ell_{\alpha_1}; S^k_{\alpha_2})| \leq \frac{\binom{d}{i} \cdot |S_{\alpha_1+\|W_i\|_\infty}|^\ell \cdot |S_{\alpha_2+\|W_i\|_\infty}|^k}{|W_i|^{\ell+k}}.$$

3.2 Computing Probabilities

We state and prove the main technical results of our paper. The first one provides an upper bound on the probability (over $\mathbf{A}$ and $\mathbf{t}$) of existence of $(\mathbf{z}_1, \mathbf{z}_2, c)$ which satisfy $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}$. This can be applied to the security analysis of the Bai-Galbraith scheme [BG14] or qTESLA [ABB+17, BAA+17]. The second one, however, considers a slightly different equation: $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta$ where $\mathbf{t}_1 = \mathsf{Power2Round}_q(\mathbf{t}, \delta)$ for some δ, and can be applied to the security analysis of Dilithium-QROM [KLS18]. In both cases, we closely follow the proof strategy in Lemma 3.2 of [SS13].

Theorem 3.8 *Let $\alpha_1, \alpha_2 \in \mathbb{N}$ and $\mathcal{D} \subseteq R_q \backslash \{0\}$. Also, for $i = 1, ..., d$, define $W_i \subseteq R_q$ to be a set of polynomials such that for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u - v)| < i$. Then*

$$\Pr_{\mathbf{A} \leftarrow R_q^{k \times \ell}, \mathbf{t} \leftarrow R_q^k}[\exists (\mathbf{z}_1, \mathbf{z}_2, c) \in S_{\alpha_1}^\ell \times S_{\alpha_2}^k \times \mathcal{D} : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] \leq \frac{|S_{\alpha_1}|^\ell \cdot |S_{\alpha_2}|^k \cdot |\mathcal{D}|}{q^{nk}} + \sum_{i=1}^{e} \frac{\binom{e}{i} \cdot |S_{\alpha_1 + \|W_i\|_\infty}|^\ell \cdot |S_{\alpha_2 + \|W_i\|_\infty}|^k \cdot |\mathcal{D}|}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}} \tag{5}$$

where e is the largest integer such that $||\mathcal{D}|| \geq q^{e/d}$.

Proof. Fix $\mathbf{z}_1 = (z_1, ..., z_\ell), \mathbf{z}_2 = (z'_1, ..., z'_k)$ and c. We first prove that

$$\mathsf{Zero}(\mathbf{z}_1, c) \neq \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2) \implies \Pr_{\mathbf{A} \leftarrow R^{k \times \ell}, \mathbf{t} \leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] = 0.$$

Suppose that $\mathsf{Zero}(\mathbf{z}_1, c) \neq \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2)$. Then, there exists some $i \in [d]$ such that $i \in \mathsf{Zero}(\mathbf{z}_1, c)$ and $i \notin \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2)$. This implies that there is some $j \in [k]$ so that $i \notin \mathsf{Zero}(\mathbf{z}_1, c, z'_j)$ (otherwise $i \in \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2)$). In particular, we have $z'_j \not\equiv 0 \pmod{f_i(X)}$. Denote $\mathbf{A} = (a_{i,j})$ and $\mathbf{t} = (t_1, ..., t_k)$ and note that

$$\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t} \implies a_{j,1}z_1 + ... + a_{j,\ell}z_\ell + z'_j = ct_j.$$

However,

$$0 \not\equiv z'_j \equiv a_{j,1}z_1 + ... + a_{j,\ell}z_\ell + z'_j \equiv ct_j \equiv 0 \pmod{f_i(X)},$$

contradiction.

Hence, there are no $\mathbf{A}, \mathbf{t}$ which satisfy $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}$. Thus, we only consider $(\mathbf{z}_1, c, \mathbf{z}_2)$ such that $\mathsf{Zero}(\mathbf{z}_1, c) = \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2)$, alternatively $(\mathbf{z}_1, c, \mathbf{z}_2) \in \mathsf{ZeroRows}_i(S_{\alpha_1}^\ell, \mathcal{D}; S_{\alpha_2}^k)$ for some $i \leq e$. We claim that

$$\Pr_{\mathbf{A} \leftarrow R^{k \times \ell}, \mathbf{t} \leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] = 1/q^{nk(1-i/d)}.$$

Note that we can write:

$$\Pr_{\mathbf{A} \leftarrow R^{k \times \ell}, \mathbf{t} \leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] = \prod_{i=1}^{k} \Pr_{a_{i,1}, ..., a_{i,\ell}, t_i \leftarrow R_q}[a_{i,1}z_1 + ... + a_{i,\ell}z_\ell + z'_i = c \cdot t_i].$$

Let us fix an index i and define

$$A = \{(a_1, ..., a_\ell, t) \in R_q^{\ell+1} : \sum_{j=1}^{\ell} a_j z_j + z'_i = c \cdot t\}.$$

We want to show that $|A| = q^{n(\ell + i/d)}$. Take any $u \in [d]$ and consider the set

$$A_u = \{(a_1, ..., a_\ell, t) \in (\mathbb{Z}_q[X]/(f_u(X)))^{\ell+1} : a_1z_1 + ... + a_\ell z_\ell + z'_i \equiv c \cdot t \pmod{f_u(X)}\}.$$

If $u \in \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2)$ then any $a_1, ..., a_\ell, t$ satisfy the equation, because

$$z_1 \equiv ... \equiv z_\ell \equiv z'_i \equiv c \equiv 0 \pmod{f_u(X)}.$$

Hence, $|A_u| = q^{(l+1) \cdot n/d}$. If $u \notin \mathsf{Zero}(\mathbf{z}_1, c, \mathbf{z}_2)$ then one of $z_1, ..., z_\ell, c$ is invertible modulo $(f_u(X), q)$, without loss of generality say z_j. Then, $a_1, ..., a_{j-1}, a_{j+1}, ..., a_\ell, c$ can be chosen arbitrarily and a_j is picked such that the equation is satisfied. Therefore, $|A_u| = q^{\ell \cdot n/d}$. Now, by the Chinese Remainder Theorem we have that

$$|A| = \prod_{u=1}^{d} |A_u| = q^{i\cdot(\ell+1)\cdot n/d + (d-i)\cdot \ell \cdot n/d} = q^{n(\ell + i/d)}.$$

Hence,

$$\Pr_{a_{i,1},\ldots,a_{i,\ell},t_i \leftarrow R_q}[a_{i,1}z_1 + \ldots + a_{i,\ell}z_\ell + z_i' = c\cdot t_i] = \frac{|A|}{q^{(\ell+1)\cdot n}} = 1/q^{n(1-i/d)}.$$

Eventually, we obtain $\Pr_{\mathbf{A}\leftarrow R^{k\times\ell},\mathbf{t}\leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] = 1/q^{nk(1-i/d)}$.

Now, we combine the observations above and Lemma 3.6. For clarity, set $Z_i = \mathsf{ZeroRows}_i(S_{\alpha_1}^\ell, \mathcal{D}; S_{\alpha_2}^k)$. Then,

$$\begin{aligned}
&\Pr_{\mathbf{A}\leftarrow R^{k\times\ell},\mathbf{t}\leftarrow R_q^k}[\exists(\mathbf{z}_1,\mathbf{z}_2,c)\in S_{\alpha_1}^\ell \times S_{\alpha_2}^k \times \mathcal{D} : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] \\
&\quad\le \sum_{\mathbf{z}_1\in S_{\alpha_1}^\ell, c\in\mathcal{D}, \mathbf{z}_2\in S_{\alpha_2}^k} \Pr_{\mathbf{A}\leftarrow R^{k\times\ell},\mathbf{t}\leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] \\
&\quad\le \sum_{i=0}^{e} \sum_{(\mathbf{z}_1,c,\mathbf{z}_2)\in Z_i} \Pr_{\mathbf{A}\leftarrow R^{k\times\ell},\mathbf{t}\leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}] \\
&\quad\le \sum_{i=0}^{e} \sum_{(\mathbf{z}_1,c,\mathbf{z}_2)\in Z_i} 1/q^{nk(1-i/d)} \\
&\quad\le \sum_{i=0}^{e} |Z_i|/q^{nk(1-i/d)} \\
&\quad\le \frac{|S_{\alpha_1}|^\ell \cdot |S_{\alpha_2}|^k \cdot |\mathcal{D}|}{q^{nk}} + \sum_{i=1}^{e} \frac{\binom{e}{i}\cdot |S_{\alpha_1 + \|W_i\|_\infty}|^\ell \cdot |S_{\alpha_2 + \|W_i\|_\infty}|^k \cdot |\mathcal{D}|}{|W_i|^{\ell+k}\cdot q^{nk(1-i/d)}}.
\end{aligned} \tag{6}$$

□

We can obtain a very similar result for $\mathcal{D} = \{0\}$ using Corollary 3.7. We just need to pick e to be the integer, such that any non-zero $(\mathbf{z}_1,\mathbf{z}_2) \in S_{\alpha_1}^\ell \times S_{\alpha_2}^k$ has at most e zero rows. Since each component of $\mathbf{z}_1$ has norm at most $\alpha_1\sqrt{n}$, we could choose the maximal e so that $\alpha_1\sqrt{n} \ge q^{e/d}$. We omit the proof since it is very similar to the one for Theorem 3.8.

Corollary 3.9 *Let $\alpha_1, \alpha_2 \in \mathbb{N}$. Also, for $i = 1,...,d$, define $W_i \subseteq R_q$ to be a set of polynomials such that for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u-v)| < i$. Then*

$$\Pr_{\mathbf{A}\leftarrow R_q^{k\times\ell}}[\exists(\mathbf{z}_1,\mathbf{z}_2)\in S_{\alpha_1}^\ell\backslash\{\mathbf{0}\} \times S_{\alpha_2}^k : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = \mathbf{0}] \le \frac{|S_{\alpha_1}|^\ell \cdot |S_{\alpha_2}|^k}{q^{nk}} + \sum_{i=1}^{e} \frac{\binom{d}{i}\cdot |S_{\alpha_1 + \|W_i\|_\infty}|^\ell \cdot |S_{\alpha_2 + \|W_i\|_\infty}|^k}{|W_i|^{\ell+k}\cdot q^{nk(1-i/d)}}, \tag{7}$$

where e is the largest integer such that $\alpha_1\sqrt{n} \ge q^{e/d}$.

The next theorem considers a modified equation $\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta$ where $\mathbf{t}_1 = \mathsf{Power2Round}_q(\mathbf{t},\delta)$ for some $\delta \in \mathbb{N}$. However, we need to take a slightly different approach in order to provide a reasonable upper bound for the probability due to the appearance of $\mathsf{Power2Round}_q$ function.

Theorem 3.10 *Let $\alpha_1, \alpha_2, \delta \in \mathbb{N}$ and $\mathcal{D} \subseteq R_q\backslash\{\mathbf{0}\}$. Also, for $i = 1,...,d$, define $W_i \subseteq R_q$ to be a set of polynomials such that for any two distinct $u, v \in W_i$, $|\mathsf{Zero}(u-v)| < i$. Then*

$$\begin{aligned}
&\Pr_{\mathbf{A}\leftarrow R_q^{k\times\ell},\mathbf{t}\leftarrow R_q^k}[\exists(\mathbf{z}_1,\mathbf{z}_2,c)\in S_{\alpha_1}^\ell \times S_{\alpha_2}^k \times \mathcal{D} : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1\cdot 2^\delta] \le \\
&\quad\le |D|\cdot|S_{\alpha_2}|^k \cdot ((\frac{2^\delta}{q^{(1-e_1/d)}})^{nk} + \frac{|S_{\alpha_1}|^\ell}{q^{nk}} + \sum_{i=1}^{e_2} \frac{\binom{d}{i}\cdot |S_{\alpha_1+\|W_i\|_\infty}|^\ell}{|W_i|^\ell \cdot q^{nk(1-i/d)}})
\end{aligned} \tag{8}$$

where $\mathbf{t}_1 = \mathsf{Power2Round}_q(\mathbf{t}, \delta)$ *and* e_1 *(resp.* e_2*) is the largest integer such that* $||\mathcal{D}|| \geq q^{e_1/d}$ *(resp.* $\alpha_1\sqrt{N} \geq q^{e_2/d}$*).*

Proof. **Case 1.** suppose that $\mathbf{z}_1 = 0$. Then, the probability becomes:

$$\Pr_{\mathbf{t}\leftarrow R_q^k}[\exists(\mathbf{z}_2, c) \in S_{\alpha_2}^k \times \mathcal{D} : \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta].$$

Fix $\mathbf{z}_2 = (z_1, ..., z_k)$, c and denote $\mathbf{t} = (t_1, ..., t_k)$. Consider the following probability:

$$\Pr_{\mathbf{t}\leftarrow R_q^k}[\mathbf{z}_2 = c\mathbf{t}] = \prod_{j=1}^{k} \Pr[z_j = ct_j].$$

By definition of e_1, we have $|\mathsf{Zero}(c)| \leq e_1$ by Lemma 3.2. Take arbitrary $j \in [k]$. We compute the maximal number of polynomials t_j satisfying $z_j = ct_j$. Define a set

$$T_u = \{t \in \mathbb{Z}_q[X]/(f_u(X)) : z_j \equiv ct \pmod{f_u(X)}\}.$$

Clearly, $|T_u| \leq q^{n/d}$. Let $u \notin \mathsf{Zero}(c)$. Then, c is invertible modulo $(f_u(X), q)$. Therefore, $|T_u| = 1$. By the Chinese Remainder Theorem, the number of polynomials t_j satisfying $z_j = ct_j$ is at most

$$\prod_{u=1}^{k} |T_u| \leq q^{|\mathsf{Zero}(c)|\cdot n/d} \leq q^{e_1\cdot n/d}.$$

Hence, we end up with

$$\Pr[z_j = ct_j] \leq \frac{q^{e_1\cdot n/d}}{q^n} = \frac{1}{q^{n(1-e_1/d)}}.$$

Thus:

$$\Pr_{\mathbf{t}\leftarrow R_q^k}[\mathbf{z}_2 = c\mathbf{t}] = \prod_{j=1}^{k} \Pr[z_j = ct_j] \leq \frac{1}{q^{nk(1-e_1/d)}}.$$

For $\mathbf{t} \in R_q^k$, the most frequent value of each coefficient of $\mathbf{t}_1$ occurs at most 2^δ times. Hence,

$$\Pr_{\mathbf{t}\leftarrow R_q^k}[\mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] \leq (\frac{2^\delta}{q^{(1-e_1/d)}})^{nk}.$$

Eventually, by the union bound we obtain:

$$\Pr_{\mathbf{t}\leftarrow R_q^k}[\exists(\mathbf{z}_2, c) \in S_{\alpha_2}^k \times \mathcal{D} : \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] \leq \sum_{\mathbf{z}_2\in S_{\alpha_2}^k, c\in\mathcal{D}} (\frac{2^\delta}{q^{(1-e_1/d)}})^{nk},$$

and the sum is equal to $|\mathcal{D}| \cdot |S_{\alpha_2}|^k \cdot (\frac{2^\delta}{q^{(1-e_1/d)}})^{nk}$.

Case 2. Suppose that $\mathbf{z} = (z_1, ..., z_\ell) \neq \mathbf{0}$ and fix $\mathbf{z}_2 = (z_1', ..., z_k')$ and c. Also, denote $\mathbf{A} = (a_{i,j})$, $\mathbf{t} = (t_1, ..., t_k)$ and $t_i' = \mathsf{Power2Round}_q(t_i, \delta)$ for $i \in [k]$. Then,

$$\Pr_{\mathbf{A}\leftarrow R^{k\times\ell}, \mathbf{t}\leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] = \prod_{i=1}^{k} \Pr_{a_{i,1},...,a_{i,\ell}, t_i\leftarrow R_q}[\sum_{j=1}^{\ell} a_{i,j}z_j + z_i' = c \cdot t_i' \cdot 2^\delta].$$

Let us fix an index i and consider the set

$$A_t = \{(a_1, ..., a_\ell) \in R_q^\ell : \sum_{j=1}^{\ell} a_j z_j + z_i' = c \cdot t' \cdot 2^\delta\},$$

where $t' = \mathsf{Power2Round}_q(t)$. We want to prove that $|A_t| \leq q^{n(\ell-1+m/d)}$, where $m = |\mathsf{Zero}(z_1, ..., z_\ell)|$. Define

$$A_t^u = \{(a_1, ..., a_\ell) \in (\mathbb{Z}_q[X]/(f_u(X)))^\ell : \sum_{j=1}^{\ell} a_j z_j \equiv c \cdot t' \cdot 2^\delta - z_i' \pmod{f_u(X)}\}.$$

Clearly, we have $|A_t^u| \le q^{\ell \cdot n/d}$. Consider $u \notin \mathsf{Zero}(z_1, ..., z_\ell)$. This means that z_w is invertible modulo $(f_u(X), q)$ for some $w \in [\ell]$. Hence, we can pick any possible values for $a_1, ..., a_{w-1}, a_{w+1}, ..., a_\ell$ and then adjust a_w so that it satisfies the equation. Note that for fixed $a_1, ..., a_{w-1}, a_{w+1}, ..., a_\ell$, there is exactly one such a_w. Thus, $|A_t^u| = q^{(\ell-1)\cdot n/d}$. By the Chinese Remainder Theorem, we get

$$|A_t| = \prod_{u=1}^{d} |A_t^u| \le q^{m \cdot n\ell/d} \cdot q^{(d-m)\cdot(\ell-1)n/d} = q^{n(\ell-1+m/d)}.$$

Since we consider uniform distribution for $a_{i,1}, ..., a_{i,\ell}, t_i$, we can conclude that:

$$\Pr_{a_{i,1},...,a_{i,\ell},t_i \leftarrow R_q}[\sum_{j=1}^{\ell} a_{i,j} z_j + z'_i = c \cdot t'_i \cdot 2^\delta] = \frac{\sum_{t_i \in R_q} A_{t_i}}{q^{\ell \cdot n} \cdot q^n} \le \frac{q^{n(\ell-1+m/d)}}{q^{\ell \cdot n}} = 1/q^{n(1-m/d)}.$$

Therefore, $\Pr_{\mathbf{A} \leftarrow R^{k\times\ell}, \mathbf{t} \leftarrow R_q^k}[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] \le 1/q^{nk(1-m/d)}$.

Now we can apply the union bound. First of all, note that if $i > e_2$ then $\mathsf{ZeroRows}_i(S_{\alpha_1}^\ell \backslash \{\mathbf{0}\}) = \varnothing$ by Lemma 3.2. Hence,

$$S_{\alpha_1}^\ell \backslash \{\mathbf{0}\} = \bigcup_{i=0}^{d} \mathsf{ZeroRows}_i(S_{\alpha_1}^\ell \backslash \{\mathbf{0}\}) = \bigcup_{i=0}^{e_2} \mathsf{ZeroRows}_i(S_{\alpha_1}^\ell \backslash \{\mathbf{0}\}).$$

For simplicity, denote $Z_i = \mathsf{ZeroRows}_i(S_{\alpha_1}^\ell \backslash \{\mathbf{0}\})$. Then,

$$\begin{aligned}
\Pr[\exists (\mathbf{z}_1, \mathbf{z}_2, c) \in S_{\alpha_1}^\ell \times S_{\alpha_2}^k \times \mathcal{D} : \mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] \\
&\le \sum_{\mathbf{z}_1 \in S_{\alpha_1}^\ell \backslash\{0\}, \mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \Pr[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] \\
&\le \sum_{\mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \sum_{i=0}^{e_2} \sum_{\mathbf{z}_1 \in Z_i} \Pr[\mathbf{A}\mathbf{z}_1 + \mathbf{z}_2 = c\mathbf{t}_1 \cdot 2^\delta] \\
&\le \sum_{\mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \sum_{i=0}^{e_2} \sum_{\mathbf{z}_1 \in Z_i} 1/q^{nk(1-i/d)} \\
&\le \sum_{\mathbf{z}_2 \in S_{\alpha_2}^k, c \in \mathcal{D}} \sum_{i=0}^{e_2} |Z_i|/q^{nk(1-i/d)}.
\end{aligned} \tag{9}$$

By Lemma 3.4, $|Z_i| \le \frac{\binom{d}{i} \cdot |S_{\alpha_1 + ||W_i||_\infty}|^\ell}{|W_i|^\ell}$. Also, we have $|Z_0| \le |S_{\alpha_1}|^\ell$. Therefore, we can bound the probability above by:

$$|\mathcal{D}| \cdot |S_{\alpha_2}|^k \cdot (|S_{\alpha_1}|^\ell / q^{nk} + \sum_{i=1}^{e_2} \frac{\binom{d}{i} \cdot |S_{\alpha_1 + \|W_i\|_\infty}|^\ell}{|W_i|^\ell \cdot q^{nk(1-i/d)}}). \tag{10}$$

The theorem now follows from combining the two cases. □

3.3 Constructing W_i

All the probability results presented in the previous subsection depend on the sizes of sets W_i. Recall that a set W_i satisfies a condition that for any two distinct $u, v \in W_i$, we have $|\mathsf{Zero}(u - v)| < i$. Based on the upper bounds obtained above, we would like to construct large sets W_i but with small infinity norm $||W_i||_\infty$.

Let us start by constructing W_1. We choose

$$W_1 := \{X^i : i \in [2n]\}.$$

Clearly, $X^i - X^j \in R_q$ is invertible, for $i \ne j$, so $|\mathsf{Zero}(X^i - X^j)| = 0 < 1$. Also, $|W_1| = 2n$ and $||W_1||_\infty = 1$.

Now, let us fix $i \geq 2$. The main idea is to set W_i to be a subset of $S = \{u \in R_q : ||u|| < \frac{1}{2} q^{i/d}\}$, i.e $|| W_i|| < \frac{1}{2} q^{i/d}$. Note that if we pick two distinct $u, v \in S$, then $0 < ||u - v|| < q^{i/d}$ by the triangle inequality. Hence, by Lemma 3.2 we get that $|\mathsf{Zero}(u - v)| < i$. Therefore, any subset of S will satisfy the condition for W_i [6].

If $t := \left\lfloor \frac{q^{i/d}}{2} \right\rfloor$ is smaller than $\sqrt{n}$ then we set

$$W_i := \{\sum_{j=1}^{t^2} \epsilon_j \cdot X^{\alpha_j} \in R_q : \epsilon_1, ..., \epsilon_{t^2} \in \{-1, 0, 1\}, \{\alpha_1, ..., \alpha_{t^2}\} \in \mathcal{P}_{t^2}([n])\}.$$

Then, $|| W_i||_\infty = 1, || W_i|| = t < \frac{1}{2} q^{i/d}$ and

$$| W_i| = \sum_{j=0}^{t^2} \binom{n}{j} \cdot 2^j.$$

Suppose that $t \geq \sqrt{n}$. In this case, we provide two constructions of W_i and in the experiments we choose the one that minimises the overall probability.

1. Set $W_i := S$. Then, $|| W_i||_\infty = \left\lfloor \frac{1}{2} q^{i/d} \right\rfloor$ and $| W_i| \geq V_n(\frac{1}{2} q^{i/d} - \sqrt{n})$ [7] where $V_N(r)$ is the volume of an n-dimensional ball of radius r.
2. Set $W_i := S_{\lfloor t/\sqrt{n} \rfloor}$. Clearly, we have the following properties: $W_i \subseteq S$, $|| W_i||_\infty = \lfloor t/\sqrt{n} \rfloor$ and $| W_i| = (2 \lfloor t/\sqrt{n} \rfloor + 1)^n$.

4 Applications to the Bai-Galbraith Scheme

We present a slightly modified version of Bai-Galbraith scheme [BG14] whose security is based on MLWE in the quantum random oracle model. First, we construct the corresponding lossy identification protocol [8]. Results from the previous section will be used to prove security properties of this ID scheme. Then, using the main result of [KLS18], we obtain the secure signature scheme in the QROM. Note that identical techniques can be applied to other closely related signature schemes, such as qTESLA [ABB+17, BAA+17] or the original scheme [BG14]. We focus on the modified scheme because it is actually a simpler version of Dilithium-QROM. Since the highly-optimised version of Dilithium-QROM can be somewhat overwhelming to readers who are not already comfortable with such constructions, we consider its simplified version here.

4.1 The Identification Protocol

The algorithms for identification protocol $\mathsf{ID} = (\mathsf{IGen}, \mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ are described in Figure 3 with the concrete parameters $\mathsf{par} = (q, d, n, k, \ell, \gamma, \gamma', \eta, \beta)$ given later in Table 1 and Table 2.

We want the challenge space in these ID and signature schemes to be a subset of the ring R, have size a little larger than 2^{256}, and consist of polynomials with small norms. In this paper, we set the dimension n of the ring R to be equal to 512. Hence, let us define the following challenge set:

$$\mathsf{ChSet} := \{c \in R \mid \|c\|_\infty = 1 \text{ and } \|c\| = \sqrt{46}\}. \tag{11}$$

Hence, ChSet consists of elements in R with $-1/0/1$ coefficients that have exactly 46 non-zero coefficients. The size of this set is $\binom{n}{46} \cdot 2^{46}$, which for $n = 512$ is greater than 2^{265}.

Key Generation. The key generation starts with choosing a random 256-bit seed ρ and expanding into a matrix $\mathbf{A} \in R_q^{k \times \ell}$ by an extendable output function Sam, i.e. a function on bit strings in which the output

[6] Note that this technique can also be used for W_1 as long as $q^{1/d}$ is large enough.

[7] This can be proven similarly as in [BDL+18] by putting a box of side-length 1 centered on every integer point and checking that the ball is completely covered by these boxes.

[8] For readers not familiar with definitions of lossy and canonical identification schemes, we provide all necessary background in Appendix A.

$\underline{\mathsf{IGen}(\mathsf{par})}$

$\rho \leftarrow \{0,1\}^{256}$

$\mathbf{A} \leftarrow R_q^{k\times\ell} := \mathsf{Sam}(\rho)$

$(\mathbf{s}_1, \mathbf{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k$

$\mathbf{t} := \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$

$pk = (\rho, \mathbf{t})$

$sk = (\rho, \mathbf{s}_1, \mathbf{s}_2)$

return (pk, sk)

$\underline{\mathsf{P}_1(sk)}$

$\mathbf{A} \leftarrow R_q^{k\times\ell} := \mathsf{Sam}(\rho)$

$\mathbf{y} \leftarrow S_{\gamma'-1}^\ell$

$\mathbf{w} := \mathbf{A}\mathbf{y}$

$\mathbf{w}_1 := \mathsf{HighBits}_q(\mathbf{w}, 2\gamma)$

return $(W = \mathbf{w}_1, St = (\mathbf{w}, \mathbf{y}))$

$\underline{\mathsf{P}_2(sk, W = \mathbf{w}_1, c, St = (\mathbf{w}, \mathbf{y}))}$

$\mathbf{z} := \mathbf{y} + c\mathbf{s}_1$

if $\|\mathbf{z}\|_\infty \geq \gamma' - \beta$ **or** $\|\mathsf{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma)\|_\infty \geq \gamma - \beta$

  then $\mathbf{z} := \perp$

else return $Z = \mathbf{z}$

$\underline{\mathsf{V}(pk, W = \mathbf{w}_1, c, Z = \mathbf{z})}$

return $[\![\|\mathbf{z}\|_\infty < \gamma' - \beta]\!]$ **and** $[\![\mathbf{w}_1 = \mathsf{HighBits}_q(\mathbf{A}\mathbf{z} - c\mathbf{t}, 2\gamma)]\!]$

Figure 3: Modified Bai-Galbraith identification protocol.

can be extended to any desired length, modeled as a random oracle. The secret keys $(\mathbf{s}_1, \mathbf{s}_2) \in S_\eta^\ell \times S_\eta^k$ have uniformly random coefficients between $-\eta$ and η (inclusively). The value $\mathbf{t} = \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$ is then computed. The public key needed for verification is $(\rho, \mathbf{t})$ and the secret key is $(\rho, \mathbf{s}_1, \mathbf{s}_2)$.

Protocol Execution. The prover starts the identification protocol by reconstructing $\mathbf{A}$ from the random seed ρ. The next step has the prover sample $\mathbf{y} \leftarrow S_{\gamma'-1}^\ell$ and then compute $\mathbf{w} = \mathbf{A}\mathbf{y}$. He then writes $\mathbf{w} = 2\gamma \cdot \mathbf{w}_1 + \mathbf{w}_0$, with $\mathbf{w}_0$ between $-\gamma$ and γ (inclusively), and then sends $\mathbf{w}_1$ to the verifier.

The set ChSet is defined as in Equation (11), and $\mathsf{ZSet} = S_{\gamma'-\beta-1}^\ell \times \{0,1\}^k$. The set of commitments WSet is defined as $\mathsf{WSet} = \{\mathbf{w}_1 \ : \ \exists \mathbf{y} \in S_{\gamma'-1}^\ell \text{ s.t. } \mathbf{w}_1 = \mathsf{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma)\}$.

The verifier generates a random challenge $c \leftarrow \mathsf{ChSet}$ and sends it to the prover. The prover computes $\mathbf{z} = \mathbf{y} + c\mathbf{s}$. If $\mathbf{z} \notin S_{\gamma'-\beta-1}^\ell$, then the prover sets his response to $\perp$. He also replies with $\perp$ if $\mathsf{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma) \notin S_{\gamma-\beta-1}^k$. Eventually, the verifier checks whether $\|\mathbf{z}\|_\infty < \gamma' - \beta$ and that $\mathbf{A}\mathbf{z} - c\mathbf{t}$.

4.2 Security Analysis

We omit proofs of correctness and non-abort honest verifier zero-knowledge properties since they have already been analysed in the previous works [BG14, DLL+17, KLS18, ABB+17]. Instead, we focus on lossyness, min entropy and computational unique response. We recall that sets W_i are introduced in Section 3.3.

Lemma 4.1 *If $\beta \geq \max_{s\in S_\eta, c\in\mathsf{ChSet}} ||cs||_\infty$, then* ID *is perfectly* naHVZK *and has correctness error $\nu \approx 1 - \exp(-\beta n \cdot (k/\gamma + \ell/\gamma'))$.*

Lossyness. Let us consider the scheme in which the public key is generated uniformly at random (Figure 4), rather than as in IGen of Figure 3. It is enough to show that even if the prover is computationally unbounded, he only has approximately a $1/|\mathsf{ChSet}|$ probability of making the verifier accept during each run of the identification scheme.

$\underline{\mathsf{LossyIGen}(\mathsf{par})}$

$\rho \leftarrow \{0,1\}^{256}$; $\mathbf{A} \leftarrow R_q^{k\times\ell} := \mathsf{Sam}(\rho)$

$\mathbf{t} \leftarrow R_q^k$

return $pk = (\rho, \mathbf{t})$

Figure 4: The lossy instance generator $\mathsf{LossyIGen}$.

Since the output of $\mathsf{LossyIGen}$ is uniformly random over $R_q^{k\times\ell} \times R_q^k$ and the output of IGen in Figure 3 is $(\mathbf{A}, \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2)$ where $\mathbf{A} \leftarrow R_q^{k\times\ell}$ and $(\mathbf{s}_1, \mathbf{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k$, we get that

$$\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{LOSS}}(\mathsf{A}) = \mathrm{Adv}_{k,\ell,D}^{\mathsf{MLWE}}(\mathsf{A}),$$

where D is the uniform distribution over S_η.

Lemma 4.2 *Let e_ℓ be the largest integer which satisfies $q^{e_\ell/d} \leq 2\sqrt{46}$. Then,* ID *has $\varepsilon_{\mathsf{ls}}$-lossy soundness, where*

$$\varepsilon_{\mathsf{ls}} \leq \frac{1}{|\mathsf{ChSet}|} + \frac{|S_{2(\gamma'-\beta-1)}|^\ell \cdot |S_{4\gamma+2}|^k \cdot |\mathsf{ChSet}|^2}{q^{nk}} + \sum_{i=1}^{e_\ell} \frac{\binom{e_\ell}{i} \cdot |S_{2(\gamma'-\beta-1)+\|W_i\|_\infty}|^\ell \cdot |S_{4\gamma+2+\|W_i\|_\infty}|^k \cdot |\mathsf{ChSet}|^2}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}. \tag{12}$$

Proof. Consider an unbounded adversary C that is executed in game $\mathsf{LOSSY\text{-}IMP}$ of Figure 5.

```
GAME LOSSY-IMP:
pk_ls := (ρ, t) ← LossyIGen(par)
(w1, St) ← C(pk_ls)
c ← ChSet
z ← C(St, c)
return [[w1 = HighBits_q(Az − ct, 2γ)]] and [[||z||∞ < γ' − β]]
```

Figure 5: The lossy impersonation game $\mathsf{LOSSY\text{-}IMP}$.

Assume that for some $\mathbf{w}_1$, there exist two $c \neq c' \in \mathsf{ChSet}$ and two $\mathbf{z}, \mathbf{z}'$ that lead to C winning, i.e. $\|\mathbf{z}\|_\infty, \|\mathbf{z}'\|_\infty < \gamma' - \beta$ and

$$\begin{aligned} \mathbf{w}_1 &= \mathsf{HighBits}_q(\mathbf{A}\mathbf{z} - \mathbf{t}c, 2\gamma), \\ \mathbf{w}_1 &= \mathsf{HighBits}_q(\mathbf{A}\mathbf{z}' - \mathbf{t}c', 2\gamma). \end{aligned}$$

By Lemma 2.1, we know that this implies

$$\begin{aligned} \|\mathbf{A}\mathbf{z} - \mathbf{t}c - \mathbf{w}_1 \cdot 2\gamma\|_\infty &\leq 2\gamma + 1, \\ \|\mathbf{A}\mathbf{z}' - \mathbf{t}c' - \mathbf{w}_1 \cdot 2\gamma\|_\infty &\leq 2\gamma + 1. \end{aligned}$$

By the triangle inequality, we have that

$$\|\mathbf{A}(\mathbf{z} - \mathbf{z}') - \mathbf{t} \cdot (c - c')\|_\infty \leq 4\gamma + 2\,,$$

which can be rewritten as

$$\mathbf{A}(\mathbf{z} - \mathbf{z}') + \mathbf{u} = \mathbf{t} \cdot (c - c') \tag{13}$$

for some $\mathbf{u}$ such that $\|\mathbf{u}\|_\infty \leq 4\gamma + 2$ (and $\|\mathbf{z} - \mathbf{z}'\|_\infty \leq 2(\gamma' - \beta - 1)$).

If $\mathbf{A} \leftarrow R_q^{k \times \ell}$ and $\mathbf{t} \leftarrow R_q^k$, then, by Theorem 3.8, we have that Equation (13) is satisfied with probability less than

$$\frac{|S_{2(\gamma'-\beta-1)}|^\ell \cdot |S_{4\gamma+2}|^k \cdot |\mathcal{D}|}{q^{nk}} + \sum_{i=1}^{e_\ell} \frac{\binom{e_\ell}{i} \cdot |S_{2(\gamma'-\beta-1)+\|W_i\|_\infty}|^\ell \cdot |S_{4\gamma+2+\|W_i\|_\infty}|^k \cdot |\mathcal{D}|}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}},$$

where $\mathcal{D} := \{c - c' : c, c' \in \mathsf{ChSet}\} \backslash \{0\}$ and sets W_i's are defined in Section 3.3.

Thus, except with the above probability, for every $\mathbf{w}_1$, there is at most one possible c that allows C to win. In other words, except with the above probability, C has at most a $1/|\mathsf{ChSet}|$ chance of winning. ☐

Note that we do not make any assumptions on the prime q. However, small d (e.g. $d = 2$ for $q \equiv 3$ or $5 \pmod 8$) implies small e_ℓ. As a consequence, the smaller d we choose, then the probability above also decreases.

Min-entropy. Now, we prove that the $\mathbf{w}_1$ sent by the honest prover in the first step is extremely likely to be distinct for every run of the protocol.

Lemma 4.3 *Let* e_m *be the largest integer which satisfies* $q^{e_m/d} \leq 2\gamma'\sqrt{n}$*. Then the identification scheme* ID *in Figure 3 has*

$$\alpha > \log\left(\min\left\{\frac{1}{M}, (2\gamma'-1)^{n\ell}\right\}\right)$$

bits of min-entropy, where

$$M := \frac{|S_{2\gamma'}|^\ell \cdot |S_{2\gamma}|^k}{q^{nk}} + \sum_{i=1}^{e_m} \frac{\binom{d}{i} \cdot |S_{2\gamma'+\|W_i\|_\infty}|^\ell \cdot |S_{2\gamma+\|W_i\|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}.$$

Proof. We claim that

$$\Pr_{\mathbf{A}\leftarrow R_q^{k\times\ell}}[\exists\, \mathbf{y} \neq \mathbf{y}' \in S_{\gamma'-1}^\ell \text{ s.t. } \mathsf{HighBits}_q(\mathbf{Ay}, 2\gamma) = \mathsf{HighBits}_q(\mathbf{Ay}', 2\gamma)]$$
$$\leq \frac{|S_{2\gamma'}|^\ell \cdot |S_{2\gamma}|^k}{q^{nk}} + \sum_{i=1}^{e_m} \frac{\binom{d}{i} \cdot |S_{2\gamma'+\|W_i\|_\infty}|^\ell \cdot |S_{2\gamma+\|W_i\|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}. \quad (14)$$

Indeed, if we write

$$\mathsf{Decompose}_q(\mathbf{Ay}, 2\gamma) = (\mathbf{w}_1, \mathbf{w}_0) \text{ and } \mathsf{Decompose}_q(\mathbf{Ay}', 2\gamma) = (\mathbf{w}_1', \mathbf{w}_0'),$$

then $\mathsf{HighBits}_q(\mathbf{Ay}, 2\gamma) = \mathsf{HighBits}_q(\mathbf{Ay}', 2\gamma)$ implies that $\mathbf{Ay} = \mathbf{w}_1 \cdot 2\gamma + \mathbf{w}_0$ and $\mathbf{Ay}' = \mathbf{w}_1' \cdot 2\gamma + \mathbf{w}_0'$ with $\mathbf{w}_1 = \mathbf{w}_1'$ and $\|\mathbf{w}_0\|_\infty, \|\mathbf{w}_0'\|_\infty \leq \gamma$. Hence,

$$\mathbf{A}(\mathbf{y} - \mathbf{y}') - (\mathbf{w}_0 - \mathbf{w}_0') = \mathbf{0} \quad (15)$$

where

$$\|\mathbf{y} - \mathbf{y}'\|_\infty < 2\gamma', \|\mathbf{w}_0 - \mathbf{w}_0'\|_\infty \leq 2\gamma.$$

Corollary 3.9 shows that the probability over the choice of $\mathbf{A} \leftarrow R_q^{k\times\ell}$, that there exist two non-zero elements of norm less than 2γ and $2\gamma'$, respectively, which satisfy Equation (15) is at most

$$\frac{|S_{2\gamma'}|^\ell \cdot |S_{2\gamma}|^k}{q^{nk}} + \sum_{i=1}^{e_m} \frac{\binom{d}{i} \cdot |S_{2\gamma'+\|W_i\|_\infty}|^\ell \cdot |S_{2\gamma+\|W_i\|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}} = M.$$

This proves Equation (14).

Now, we know that with probability at least $1 - M$ over the choice of $\mathbf{A} \leftarrow R_q^{k\times\ell}$, each $W = \mathsf{HighBits}_q(\mathbf{Ay}, 2\gamma)$ has exactly a $\frac{1}{\left|S_{\gamma'-1}^\ell\right|} = (2\gamma'-1)^{-n\ell}$ probability of being output. Thus, the claim in the lemma follows directly from the definition. □

Computational Unique Response. Here, we show the Computational Unique Response (CUR) property required for strong-unforgeability of the signature scheme.

Lemma 4.4 *Let* e_c *be the largest integer such that* $q^{e_c/d} \leq 2(\gamma' - \beta)\sqrt{n}$*. Then*

$$\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}) \leq \frac{|S_{2(\gamma'-\beta)}|^\ell \cdot |S_{4\gamma+2}|^k}{q^{nk}} + \sum_{i=1}^{e_c} \frac{\binom{d}{i} \cdot |S_{2(\gamma'-\beta)+\|W_i\|_\infty}|^\ell \cdot |S_{4\gamma+2+\|W_i\|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}$$

for all (even unbounded) adversaries A.

Proof. Let $(W, c, Z) = (\mathbf{w}_1, c, \mathbf{z})$ be any valid transcript and suppose A is able to generate a valid $Z' = \mathbf{z}' \neq Z$ such that $\mathsf{V}(pk = (\mathbf{A}, \mathbf{t}), \mathbf{w}_1, c, \mathbf{z}') = 1$. Thus, we have

$$\mathbf{w}_1 = \mathsf{UseHint}_q(\mathbf{h}, \mathbf{Az} - c\mathbf{t}, 2\gamma) \text{ and } \mathbf{w}_1 = \mathsf{UseHint}_q(\mathbf{h}', \mathbf{Az}' - c\mathbf{t}, 2\gamma).$$

The above two equations imply (by Lemma 2.1) that

$$\|\mathbf{Az} - c\mathbf{t} - \mathbf{w}_1 \cdot 2\gamma\|_\infty \leq 2\gamma + 1 \text{ and } \|\mathbf{Az}' - c\mathbf{t} - \mathbf{w}_1 \cdot 2\gamma\|_\infty \leq 2\gamma + 1.$$

q	d	γ
$2^{44} - 17043$	2	592493
$2^{44} - 8583$	4	593431
$2^{44} - 13743$	8	305156
$2^{44} - 7583$	16	282832
$2^{44} - 1599$	32	285978
$2^{45} - 36991$	64	364254
$2^{45} - 58111$	128	353952
$2^{45} - 511$	256	360620
$2^{45} - 23551$	512	359769

Table 1: Prime moduli q for each possible value of d. We used the main result of [LS18] for finding q. For each case, we also provide values γ such that $2\gamma | q-1$. Just like in [KLS18], we set $\gamma' = \gamma$.

By the triangle inequality, we have

$$\mathbf{A}(\mathbf{z}-\mathbf{z}') + \mathbf{u} = \mathbf{0}$$

for some $\mathbf{u}$ such that $\|\mathbf{u}\| \leq 4\gamma + 2$ and $\|\mathbf{z}-\mathbf{z'}\| < 2(\gamma' - \beta)$. Hence, by Corollary 3.9, the probability over the choice of $\mathbf{A} \leftarrow R_q^{k\times\ell}$, that there exist such $\mathbf{v}, \mathbf{u}$ is at most

$$\frac{|S_{2(\gamma'-\beta)}|^\ell \cdot |S_{4\gamma+2}|^k}{q^{nk}} + \sum_{i=1}^{e_c} \frac{\binom{d}{i} \cdot |S_{2(\gamma'-\beta)+\| W_i\|_\infty}|^\ell \cdot |S_{4\gamma+2+\| W_i\|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}.$$

□

4.3 Concrete Parameteres

In this subsection, we instantiate the modified Bai-Galbraith mBG signature scheme obtained by the Fiat-Shamir transformation from ID with concrete parameters (Table 1 and Table 2). We consider nine different instantiations of mBG for all possible $d \in \{2^i : i \in [9]\}$.

For each value of d, we have selected parameters (e.g. prime modulus q and γ) such that the ID scheme satisfies the following security properties: (i) $\varepsilon_{\mathsf{zk}} = 0$, (ii) the scheme has more than 2845 bits of min-entropy, i.e. $\alpha > 2845$, (iii) $\varepsilon_{\mathsf{ls}} \leq 2^{-264}$, (iv) $\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{C}) \leq 2^{-288}$. Following the steps in [KLS18], one can prove security of the modified Bai-Galbraith scheme in the quantum random oracle model (see Appendix A).

We compare the nine different instantiations of the modified Bai-Galbraith scheme (Table 2) with respect to recommended parameters in Table 2. Firstly, observe that for $d \leq 4$, we pick $q \approx 2^{44}$. In this case, we end up with public key and signature size 11.29kB and 5.69kB respectively.

The situation changes for $d = 8$. Interestingly, if one keeps the same parameters as for $d = 4$ then one still gets $\varepsilon_{\mathsf{ls}} \leq 2^{-264}$, hence the lossyness property is still preserved. The problem is, however, that the advantage $\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A})$ gets extremely big. Concretely, for parameters above we have $\log(\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A})) \approx 3483$. We found out that one of the compounds in the sum is actually dominating (see Lemma 4.4). Namely, we get:

$$\log\left(\frac{\binom{8}{1} \cdot |S_{2(\gamma'-\beta)+\| W_1\|_\infty}|^\ell \cdot |S_{4\gamma+2+\| W_1\|_\infty}|^k}{|W_1|^{\ell+k} \cdot q^{nk(1-1/8)}}\right) \approx 3483.$$

We believe the reason for it being so large is because for $d = 8$, $i = 1$ and $q \approx 2^{44}$ we have $t := \left\lfloor \frac{q^{i/d}}{2\sqrt{n}} \right\rfloor = 1$ (introduced in Section 3.3). Hence, W_1 has only 3^{512} elements. As a consequence, the value above is still big. Thus, a natural way to solve this issue would be to increase q. Unfortunately, in order to keep the MLWE problem hard, this would imply increasing the size of secret keys, i.e. η. Hence, β would also get bigger, so in order to keep the repetition rate $1/(1-\nu)$ small, we would have to increase the value of γ (and γ'). In this case, probabilities related to the security of ID, e.g. $\varepsilon_{\mathsf{ls}}, \log(\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}))$, would get considerably bigger, so one would need to consider larger q again and eventually, we would end up in a

d	2	4	8	16	32	64	128	256	512
n	512	512	512	512	512	512	512	512	512
(k, ℓ) (dimensions of $\mathbf{A}$)	$(4,4)$	$(4,4)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$
# of $\pm 1's$ in $c \in \mathsf{ChSet}$	46	46	46	46	46	46	46	46	46
η (max. coeff. of $\mathbf{s}_1, \mathbf{s}_2$)	5	5	2	2	2	2	2	2	2
$\beta(= \eta \cdot (\#\text{of 1's in } c))$	230	230	92	92	92	92	92	92	92
e_ℓ (lossyness)	0	0	0	1	2	5	10	21	42
e_c (CUR)	1	2	4	8	17	34	68	136	272
e_m (min-entropy)	1	2	4	8	17	34	68	136	272
$\log(\varepsilon_{\mathsf{ls}})$	-264	-264	-264	-264	-264	-264	-264	-264	-264
$\log(\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}))$	-1326	-1317	-592	-924	-288	-799	-986	-766	-677
α	3373	3363	3149	3481	2845	3356	3543	3324	3235
pk size (kilobytes)	11.29	11.29	14.11	14.11	14.11	14.43	14.43	14.43	14.43
sig size (kilobytes)	5.69	5.69	6.76	6.76	6.76	6.76	6.76	6.76	6.76
Exp. Repeats $\frac{1}{1-\nu}$	4.94	4.93	4.68	5.29	5.19	3.64	3.78	3.69	3.70
BKZ block-size to break LWE	480	480	600	600	600	585	585	585	585
Best known classical bit-cost	140	140	175	175	175	171	171	171	171
Best known quantum bit-cost	127	127	159	159	159	155	155	155	155

Table 2: Parameters for the modified Bai-Galbraith scheme. Recall that ν is the maximum coefficient of secret keys $\mathbf{s}_1, \mathbf{s}_2$ and $\beta = \nu \cdot (\#\text{ of } \pm 1\text{'s in } c \in \mathsf{ChSet})$. On the other hand, variables $e_\ell, e_c, e_m, \alpha, \varepsilon_{\mathsf{ls}}, \mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}), \nu$ are defined in Section 4.2.

vicious circle. We avoid that by increasing dimensions $(k, \ell) = (5, 5)$ of the matrix $\mathbf{A}$. Unfortunately, this comes at a price of larger public key (14.11kB) and signature (6.76kB) sizes. In order to minimise such costs, we decrease the size of secret keys $\eta = 2$ and thus, we select smaller values for γ. As before, we choose $q \approx 2^{44}$. We pick almost identical parameters for $d = 16$ and $d = 32$.

Next, we consider $d \geq 64$. If we choose the parameters as for $d = 32$ then the lossyness probability $\varepsilon_{\mathsf{ls}}$ is no longer small and therefore, we need to increase the $q \approx 2^{45}$. We observe that the new parameters still provide much more than 128 bits of security for MLWE. The public key gets slightly larger (14.43kB) and the signature size stays the same as before.

In order to maintain security of the Bai-Galbraith scheme in the quantum random oracle model for bigger d (i.e. $d \geq 256$), we need to increase both dimensions (k, ℓ) of the matrix $\mathbf{A}$ as well as the prime modulus q. This results in having 3.13kB larger public key and 1.07kB signature sizes than for $d = 2$. We remark that security parameters were chosen such that the expected number of repetitions of the protocol $1/(1 - \nu)$ is at most six. Indeed, admitting small repetition rate as well as supporting the use of the Number Theoretic Transform, efficient caching and polynomial sampling assures us that the protocol can be performed very efficiently.

Acknowledgments

The author would like to thank Vadim Lyubashevsky for fruitful discussions and anonymous reviewers for their useful comments. Also, the author would like to thank Damien Stehlé for additional comments.

This work was supported by the SNSF ERC Transfer Grant CRETP2-166734 FELICITY.

References

[ABB+17] Erdem Alkim, Nina Bindel, Johannes A. Buchmann, Özgür Dagdelen, Edward Eaton, Gus Gutoski, Juliane Krämer, and Filip Pawlega. Revisiting TESLA in the quantum random oracle model. In *Post-Quantum Cryptography - 8th International Workshop, PQCrypto 2017, Utrecht, The Netherlands, June 26-28, 2017, Proceedings*, pages 143–162, 2017.

[ADPS16] Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In Thorsten Holz and Stefan Savage, editors, *USENIX Security 2016*, pages 327–343. USENIX Association, August 2016.

[AFLT12] Michel Abdalla, Pierre-Alain Fouque, Vadim Lyubashevsky, and Mehdi Tibouchi. Tightly-secure signatures from lossy identification schemes. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 572–590. Springer, Heidelberg, April 2012.

[BAA+17] 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`.

[BCK+14] Fabrice Benhamouda, Jan Camenisch, Stephan Krenn, Vadim Lyubashevsky, and Gregory Neven. Better zero-knowledge proofs for lattice encryption and their application to group signatures. In Palash Sarkar and Tetsu Iwata, editors, *ASIACRYPT 2014, Part I*, volume 8873 of *LNCS*, pages 551–572. Springer, Heidelberg, December 2014.

[BDF+11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 41–69. Springer, Heidelberg, December 2011.

[BDL+18] Carsten Baum, Ivan Damgård, Vadim Lyubashevsky, Sabine Oechsner, and Chris Peikert. More efficient commitments from structured lattice assumptions. In Dario Catalano and Roberto De Prisco, editors, *SCN 18*, volume 11035 of *LNCS*, pages 368–385. Springer, Heidelberg, September 2018.

[BG14] Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In Josh Benaloh, editor, *CT-RSA 2014*, volume 8366 of *LNCS*, pages 28–47. Springer, Heidelberg, February 2014.

[BKLP15] Fabrice Benhamouda, Stephan Krenn, Vadim Lyubashevsky, and Krzysztof Pietrzak. Efficient zero-knowledge proofs for commitments from learning with errors over rings. In Günther Pernul, Peter Y. A. Ryan, and Edgar R. Weippl, editors, *ESORICS 2015, Part I*, volume 9326 of *LNCS*, pages 305–325. Springer, Heidelberg, September 2015.

[BPS16] Mihir Bellare, Bertram Poettering, and Douglas Stebila. From identification to signatures, tightly: A framework and generic transforms. In Jung Hee Cheon and Tsuyoshi Takagi, editors, *ASIACRYPT 2016, Part II*, volume 10032 of *LNCS*, pages 435–464. Springer, Heidelberg, December 2016.

[DLL+17] Léo Ducas, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - Dilithium: Digital signatures from module lattices. *IACR Cryptology ePrint Archive*, 2017:633, 2017. To appear in TCHES 2018.

[GLP12] Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In Emmanuel Prouff and Patrick Schaumont, editors, *CHES 2012*, volume 7428 of *LNCS*, pages 530–547. Springer, Heidelberg, September 2012.

[KLS18] Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 552–586. Springer, Heidelberg, April / May 2018.

[LDK+17] Vadim Lyubashevsky, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-DILITHIUM. Technical report, National Institute of Standards and Technology, 2017. available at `https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions`.

[LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact Knapsacks are collision resistant. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *ICALP 2006, Part II*, volume 4052 of *LNCS*, pages 144–155. Springer, Heidelberg, July 2006.

[LN17] Vadim Lyubashevsky and Gregory Neven. One-shot verifiable encryption from lattices. In Jean-Sébastien Coron and Jesper Buus Nielsen, editors, *EUROCRYPT 2017, Part I*, volume 10210 of *LNCS*, pages 293–323. Springer, Heidelberg, April / May 2017.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May / June 2010.

[LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A toolkit for ring-LWE cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 35–54. Springer, Heidelberg, May 2013.

[LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.

[LS18] Vadim Lyubashevsky and Gregor Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 204–224. Springer, Heidelberg, April / May 2018.

[Lyu09] Vadim Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 598–616. Springer, Heidelberg, December 2009.

[Lyu12] Vadim Lyubashevsky. Lattice signatures without trapdoors. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 738–755. Springer, Heidelberg, April 2012.

[Ngu19] Ngoc Khanh Nguyen. On the non-existence of short vectors in random module lattices. In *ASIACRYPT 2019, Part II*, LNCS, pages 121–150. Springer, Heidelberg, December 2019.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In Shai Halevi and Tal Rabin, editors, *TCC 2006*, volume 3876 of *LNCS*, pages 145–166. Springer, Heidelberg, March 2006.

[PR07] Chris Peikert and Alon Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In David S. Johnson and Uriel Feige, editors, *39th ACM STOC*, pages 478–487. ACM Press, June 2007.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-LWE and polynomial-LWE problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part I*, volume 10820 of *LNCS*, pages 146–173. Springer, Heidelberg, April / May 2018.

[SAB$^+$17] Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, and Damien Stehlé. CRYSTALS-KYBER. Technical report, National Institute of Standards and Technology, 2017. available at `https://csrc.nist.gov/projects/post-quantum-cryptography/round-1-submissions`.

[SS13] Damien Stehlé and Ron Steinfeld. Making NTRUEncrypt and NTRUSign as secure as standard worst-case problems over ideal lattices. Cryptology ePrint Archive, Report 2013/004, 2013. `http://eprint.iacr.org/2013/004`.

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, Heidelberg, December 2009.

[Unr17] Dominique Unruh. Post-quantum security of Fiat-Shamir. In Tsuyoshi Takagi and Thomas Peyrin, editors, *ASIACRYPT 2017, Part I*, volume 10624 of *LNCS*, pages 65–95. Springer, Heidelberg, December 2017.

[Zha12] Mark Zhandry. Secure identity-based encryption in the quantum random oracle model. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 758–775. Springer, Heidelberg, August 2012.

A Background

A.1 Quantum Adversaries

We focus on security games in the quantum random-oracle model (QROM). Here, *quantum adversaries* are given quantum access to the random oracles involved, and classical access to all other oracles (e.g., the signing oracle). For a quantum adversary A and an oracle O, we write $\mathsf{A}^{|\mathrm{O}\rangle}$ (resp. A^{O}) to denote that O is quantum-accessible (resp. accessed classically) by A. For more background on QROM and quantum adversaries we refer to [KLS18, BDF$^+$11, Zha12].

A.2 Pseudorandom Functions

A pseudorandom function PRF is a mapping $\mathsf{PRF} : \mathcal{K} \times \{0,1\}^m \to \{0,1\}^k$, where $\mathcal{K}$ is a finite key space and m, k are integers. To a quantum adversary A and PRF we associate the advantage function

$$\mathrm{Adv}^{\mathsf{PR}}_{\mathsf{PRF}}(\mathsf{A}) := \left| \Pr[\mathsf{A}^{\mathsf{PRF}(K,\cdot)} \Rightarrow 1 \mid K \leftarrow \mathcal{K}] - \Pr[\mathsf{A}^{\mathsf{RF}(\cdot)} \Rightarrow 1] \right|,$$

where $\mathsf{RF} : \{0,1\}^n \to \{0,1\}^k$ is a perfect random function. We note that while adversary A is quantum, it only gets classical access to the oracles $\mathsf{PRF}(K,\cdot)$ and $\mathsf{RF}(\cdot)$.

A.3 Canonical Identification Schemes

A canonical identification scheme ID is a three-move protocol of the form depicted in Figure 6. The prover's first message W is called *commitment*, the verifier selects a uniform *challenge* c from set ChSet, and, upon receiving a *response* Z from the prover, makes a deterministic decision.

Definition A.1 (Canonical Identification Scheme). A canonical identification scheme ID is defined as a tuple of algorithms $\mathsf{ID} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$.

- The key generation algorithm IGen takes system parameters par as input and returns public and secret key (pk, sk). We assume that pk defines ChSet (the set of challenges), WSet (the set of commitments), and ZSet (the set of responses).
- The prover algorithm $\mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2)$ is split into two algorithms. P_1 takes as input the secret key sk and returns a commitment $W \in \mathsf{WSet}$ and a state St; P_2 takes as input the secret key sk, a commitment W, a challenge c, and a state St and returns a response $Z \in \mathsf{ZSet} \cup \{\bot\}$, where $\bot \notin \mathsf{ZSet}$ is a special symbol indicating failure.
- The verifier algorithm V takes the public key pk and the conversation transcript as input and outputs a *deterministic decision*, 1 (acceptance) or 0 (rejection).

Prover $\mathsf{P}(sk)$		Verifier $\mathsf{V}(pk)$
$(W, St) \leftarrow \mathsf{P}_1(sk)$	$\xrightarrow{\quad W \quad}$	
	$\xleftarrow{\quad c \quad}$	$c \leftarrow \mathsf{ChSet}$
$Z \leftarrow \mathsf{P}_2(sk, W, c, St)$	$\xrightarrow{\quad Z \quad}$	
		$d = \mathsf{V}(pk, W, c, Z) \in \{0,1\}$

Figure 6: A canonical identification scheme and its transcript (W, c, Z).

A *transcript* is a three-tuple $(W, c, Z) \in \mathsf{WSet} \times \mathsf{ChSet} \times \mathsf{ZSet} \cup \{\bot, \bot, \bot\}$. It is called *valid* (with respect to public-key pk) if $\mathsf{V}(pk, W, c, Z) = 1$. In Figure 7 we also define a transcript oracle Trans that returns a real interaction (W, c, Z) between prover and verifier as depicted in Figure 6, with the important convention that the transcript is defined as $(\bot, \bot, \bot)$ if $Z = \bot$.

Algorithm $\mathsf{Trans}(sk)$:
$(W, St) \leftarrow \mathsf{P}_1(sk)$
$c \leftarrow \mathsf{ChSet}$
$Z \leftarrow \mathsf{P}_2(sk, W, c, St)$
if $Z = \bot$ then return $(\bot, \bot, \bot)$
return (W, c, Z)

Figure 7: An honestly generated transcript (W, c, Z) output by the transcript oracle $\mathsf{Trans}(sk)$.

Definition A.2 (Correctness Error). Identification scheme ID has correctness error δ if for all $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$ the following holds:

- All possible transcripts (W, c, Z) satisfying $Z \neq \bot$ are valid, i.e., for all $(W, St) \in \mathsf{P}_1(sk)$, all $c \in \mathsf{ChSet}$ and all $Z \in \mathsf{P}_2(sk, W, c, St)$ with $Z \neq \bot$, we have $\mathsf{V}(pk, W, c, Z) = 1$.
- The probability that an honestly generated transcript (W, c, Z) contains $Z = \bot$ is bounded by δ, i.e., $\Pr[Z = \bot \mid (W, c, Z) \leftarrow \mathsf{Trans}(sk)] \leq \delta$.

Definition A.3 We call ID *commitment-recoverable*, if for any $(pk, sk) \in \mathsf{IGen}(\mathsf{par})$, $c \in \mathsf{ChSet}$, and $Z \in \mathsf{ZSet}$, there exists a unique $W \in \mathsf{WSet}$ such that $\mathsf{V}(pk, W, c, Z) = 1$. This unique W can be publicly computed using a commitment recovery algorithm as $W := \mathsf{Rec}(pk, c, Z)$.

We recall no-abort honest-verifier zero-knowledge, a weak variant of honest-verifier zero-knowledge that requires the transcript (as generated by $\mathsf{Trans}(sk)$) to be publicly simulatable, conditioned on $Z \neq \bot$.

Definition A.4 (No-Abort Honest-verifier Zero-knowledge). A canonical identification scheme ID is said to be $\varepsilon_{\mathsf{zk}}$-perfect naHVZK (no-abort honest-verifier zero-knowledge) if there exists an algorithm Sim that, given only the public key pk, outputs (W, c, Z) such that the following conditions hold:

- The distribution of $(W, c, Z) \leftarrow \mathsf{Sim}(pk)$ has statistical distance at most $\varepsilon_{\mathsf{zk}}$ from $(W', c', Z') \leftarrow \mathsf{Trans}(sk)$, where Trans is defined in Figure 7.
- The distribution of c from $(W, c, Z) \leftarrow \mathsf{Sim}(pk)$ conditioned on $c \neq \bot$ is uniform random in ChSet.

Note that if ID is commitment-recoverable, then we can abandon the W in the output of Trans and Sim since W can be publicly computed from (c, Z).

Definition A.5 (Min-Entropy). If the most likely value of a random variable W that is chosen from a discrete distribution D occurs with probability $2^{-\alpha}$, then we say that min-entropy$(W \mid W \leftarrow D) = \alpha$. We say that a canonical identification scheme ID has α *bits of min-entropy*, if

$$\Pr_{(pk,sk)\leftarrow \mathsf{IGen}(\mathsf{par})}[\text{min-entropy}(W \mid (W, St) \leftarrow \mathsf{P}_1(sk)) \geq \alpha] \geq 1 - 2^{-\alpha}.$$

In other words, except with probability $2^{-\alpha}$ over the choice of (pk, sk), the min-entropy of W will be at least α.

GAME LOSSY-IMP:

$pk_{\mathsf{ls}} \leftarrow \mathsf{LossyIGen}(\mathsf{par})$
$(W^*, St) \leftarrow \mathsf{C}(pk_{\mathsf{ls}})$
$c^* \leftarrow \mathsf{ChSet}$
$Z^* \leftarrow \mathsf{C}(St, c^*)$
return $[\![\mathsf{V}(pk_{\mathsf{ls}}, W^*, c^*, Z^*)]\!]$

Figure 8: The lossy impersonation game LOSSY-IMP.

We recall the computational unique response (CUR) property which states that it is computationally difficult to come up with (W, c, Z, Z') such that $\mathsf{V}(pk, W, c, Z) = \mathsf{V}(pk, W, c, Z') = 1$ and $Z' \neq Z$.

Definition A.6 (Computational Unique Response). To an adversary A we associate the advantage function

$$\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}) := \Pr\left[\begin{matrix}\mathsf{V}(pk, W, c, Z) = 1 \\ \mathsf{V}(pk, W, c, Z') = 1 \wedge Z \neq Z'\end{matrix} \middle| \begin{matrix}(pk, sk) \leftarrow \mathsf{IGen}(\mathsf{par}); \\ (W, c, Z, Z') \leftarrow \mathsf{A}(pk)\end{matrix}\right].$$

Lossy Identification schemes. We now define lossy identification schemes [AFLT12, KLS18].

Definition A.7 An identification scheme $\mathsf{ID} = (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ is lossy if there exists a lossy key generation algorithm LossyIGen that takes system parameters par as input and returns public key pk_{ls} (and no secret key sk).

We refer to $\mathsf{LID} = (\mathsf{IGen}, \mathsf{LossyIGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ as a lossy identification scheme. Let us define the LOSS *advantage function of a quantum adversary* A *against* ID as

$$\mathrm{Adv}_{\mathsf{LID}}^{\mathsf{LOSS}}(\mathsf{A}) := \left|\Pr[\mathsf{A}(pk_{\mathsf{ls}}) \Rightarrow 1 \mid pk_{\mathsf{ls}} \leftarrow \mathsf{LossyIGen}(\mathsf{par})] - \Pr[\mathsf{A}(pk) \Rightarrow 1 \mid (pk, sk) \leftarrow \mathsf{IGen}(\mathsf{par})]\right|.$$

We say that ID has $\varepsilon_{\mathsf{ls}}$-lossy soundness if for every (possibly unbounded, quantum) adversary C, $\Pr[\mathsf{LOSSY\text{-}IMP}^{\mathsf{C}} \Rightarrow 1] \leq \varepsilon_{\mathsf{ls}}$, where game LOSSY-IMP is defined in Figure 8.

Since C is unbounded, we can upper bound $\Pr[\mathsf{LOSSY\text{-}IMP}^{\mathsf{C}} \Rightarrow 1]$ as

$$\Pr[\mathsf{LOSSY\text{-}IMP}^{\mathsf{C}} \Rightarrow 1] \leq \mathbf{E}\left[\max_{W \in \mathsf{WSet}}\left(\Pr_{c \leftarrow \mathsf{ChSet}}[\exists Z \in \mathsf{ZSet} : \mathsf{V}(pk_{\mathsf{ls}}, W, c, Z) = 1]\right)\right], \tag{16}$$

where the expectation is taken over $pk_{\mathsf{ls}} \leftarrow \mathsf{LossyIGen}(\mathsf{par})$. W remark that the equality in Equation (16) is achieved for the "optimal" adversary C which on the "easiest" commitment $W \in \mathsf{WSet}$ and a random challenge $c \leftarrow \mathsf{ChSet}$ finds a response $Z \in \mathsf{ZSet}$ that the verifier accepts.

A.4 Digital Signatures

We define syntax and security of a digital signature scheme. Let par be common system parameters shared among all participants.

Definition A.8 (Digital Signature). A digital signature scheme SIG is defined as a triple of algorithms $\mathsf{SIG} = (\mathsf{Gen}, \mathsf{Sign}, \mathsf{Ver})$.

- The key generation algorithm Gen(par) returns the public and secret keys (pk, sk). We assume that pk defines the message space MSet.
- The signing algorithm $\mathsf{Sign}(sk, M)$ returns a signature σ.
- The deterministic verification algorithm $\mathsf{Ver}(pk, M, \sigma)$ returns 1 (accept) or 0 (reject).

Signature scheme SIG has correctness error γ if for all $(pk, sk) \in \mathsf{Gen}(\mathsf{par})$, all messages $M \in \mathsf{MSet}$, we have $\Pr[\mathsf{Ver}(pk, M, \mathsf{Sign}(sk, M)) = 0] \leq \gamma$.

Security. We define the UF-CMA (unforgeability against chosen-message attack), $\mathsf{UF\text{-}CMA}_1$ (unforgeability against one-per-message chosen-message attack), and UF-NMA (unforgeability against no-message attack) advantage functions of a quantum adversary A against SIG as $\mathrm{Adv}_{\mathsf{SIG}}^{\mathsf{UF\text{-}CMA}}(\mathsf{A}) := \Pr[\mathsf{UF\text{-}CMA}^{\mathsf{A}} \Rightarrow 1]$, $\mathrm{Adv}_{\mathsf{SIG}}^{\mathsf{UF\text{-}CMA}_1}(\mathsf{A}) := \Pr[\mathsf{UF\text{-}CMA}_1{}^{\mathsf{A}} \Rightarrow 1]$, and $\mathrm{Adv}_{\mathsf{SIG}}^{\mathsf{UF\text{-}NMA}}(\mathsf{A}) := \Pr[\mathsf{UF\text{-}NMA}^{\mathsf{A}} \Rightarrow 1]$, where the games

GAMES UF-CMA/UF-CMA$_1$/UF-NMA:		$\textsc{Sign}(M)$	$\textsc{Sign}_1(M)$
$(pk, sk) \leftarrow \mathsf{Gen}(\mathsf{par})$		$\mathcal{M} = \mathcal{M} \cup \{M\}$	**if** $M \in \mathcal{M}$ then return $\bot$
$(M^*, \sigma^*) \leftarrow \mathsf{A}^{\textsc{Sign}}(pk)$	//UF-CMA	$\sigma \leftarrow \mathsf{Sign}(sk, M)$	$\mathcal{M} = \mathcal{M} \cup \{M\}$
$(M^*, \sigma^*) \leftarrow \mathsf{A}^{\textsc{Sign}_1}(pk)$	//UF-CMA$_1$	**return** σ	$\sigma \leftarrow \mathsf{Sign}(sk, M)$
$(M^*, \sigma^*) \leftarrow \mathsf{A}(pk)$	//UF-NMA		**return** σ
return $[\![M^* \notin \mathcal{M}]\!] \wedge \mathsf{Ver}(pk, M^*, \sigma^*)$			

Figure 9: Games UF-CMA, UF-CMA$_1$, and UF-NMA.

UF-CMA, UF-CMA$_1$, and UF-NMA are given in Figure 9. We also consider *strong* unforgeability where the adversary may return a forgery on a message previously queried to the signing oracle, but with a different signature. In the corresponding experiments sUF-CMA and sUF-CMA$_1$, the set $\mathcal{M}$ contains tuples (M, σ) and for the winning condition it is checked that $(M^*, \sigma^*) \notin \mathcal{M}$.

Any UF-CMA$_1$ (sUF-CMA$_1$) secure signature scheme can be combined with a pseudo-random function PRF to obtain an UF-CMA (sUF-CMA) secure signature scheme by defining $\mathsf{Sign}'((sk, K), M) := \mathsf{Sign}(sk, M; \mathsf{PRF}_K(M))$, where K is a secret PRF key which is part of the secret key. This construction is well known in the classical setting [BPS16], and the same proof works in the quantum setting. Here PRF only has to provide security against quantum adversaries where the access to PRF is classical.

A.5 Fiat-Shamir Signatures in the QROM

For completeness, we recall the generic framework for constructing tight reductions in the quantum random oracle model from underlying hard problems to Fiat-Shamir signatures by Kiltz et al. [KLS18].

Let $\mathsf{ID} := (\mathsf{IGen}, \mathsf{P}, \mathsf{ChSet}, \mathsf{V})$ be a canonical identification scheme, let κ_m be a positive integer, and let $\mathsf{H} : \{0,1\}^* \to \mathsf{ChSet}$ be a hash function. The following signature scheme $\mathsf{SIG} := (\mathsf{Gen} = \mathsf{IGen}, \mathsf{Sign}, \mathsf{Ver})$ is obtained by the Fiat-Shamir transformation with aborts $\mathsf{FS}[\mathsf{ID}, \mathsf{H}, \kappa_m]$ [Lyu09].

$\mathsf{Sign}(sk, M)$	$\mathsf{Ver}(pk, M, \sigma)$
$\kappa := 0$	Parse $\sigma = (W, Z) \in \mathsf{WSet} \times \mathsf{ZSet}$
while $Z = \bot$ and $\kappa \le \kappa_m$ do	$c = \mathsf{H}(W \parallel M)$
$\quad \kappa := \kappa + 1$	**return** $\mathsf{V}(pk, W, c, Z) \in \{0, 1\}$
$\quad (W, St) \leftarrow \mathsf{P}_1(sk)$	
$\quad c = H(W \parallel M)$	
$\quad Z \leftarrow \mathsf{P}_2(sk, W, c, St)$	
if $Z = \bot$ **return** $\sigma = \bot$	
return $\sigma = (W, Z)$	

We make the convention that if $\sigma = (W, Z)$ is not in $\mathsf{WSet} \times \mathsf{ZSet}$, then $\mathsf{Ver}(pk, M, \sigma)$ returns 0 (reject). Clearly, if ID has correctness error δ, then SIG has correctness error $\gamma = \delta^{\kappa_m}$.

Define $\mathsf{SIG} := \mathsf{FS}[\mathsf{ID}, \mathsf{H}, \kappa_m]$ in the QROM. Then, the main result of [KLS18] is the following.

Theorem A.9 *Let* ID *be a lossy,* $\varepsilon_{\mathsf{zk}}$*-perfect* naHVZK *identification scheme which has* α *bits of min entropy, and is* $\varepsilon_{\mathsf{ls}}$*-lossy sound. Then, for any quantum adversary* A *against* UF-CMA$_1$ *(*sUF-CMA$_1$*) security that issues at most* Q_{H} *queries to the quantum random oracle* $|\mathsf{H}\rangle$ *and* Q_S *classical queries to the signing oracle* $\textsc{Sign}_1$*, there exists a quantum adversary* B *(and a quantum adversary* C *against* CUR*) such that*

$$\mathrm{Adv}^{\mathsf{UF\text{-}CMA}_1}_{\mathsf{SIG}}(\mathsf{A}) \le \mathrm{Adv}^{\mathsf{LOSS}}_{\mathsf{ID}}(\mathsf{B}) + 8(Q_{\mathsf{H}} + 1)^2 \cdot \varepsilon_{\mathsf{ls}} + \kappa_m Q_S \cdot \varepsilon_{\mathsf{zk}} + 2^{-\alpha+1} \ ,$$

$$\mathrm{Adv}^{\mathsf{sUF\text{-}CMA}_1}_{\mathsf{SIG}}(\mathsf{A}) \le \mathrm{Adv}^{\mathsf{LOSS}}_{\mathsf{ID}}(\mathsf{B}) + 8(Q_{\mathsf{H}} + 1)^2 \cdot \varepsilon_{\mathsf{ls}} + \kappa_m Q_S \cdot \varepsilon_{\mathsf{zk}} + 2^{-\alpha+1} + \mathrm{Adv}^{\mathsf{CUR}}_{\mathsf{ID}}(\mathsf{C}) \ ,$$

and $\mathrm{Time}(\mathsf{B}) = \mathrm{Time}(\mathsf{C}) = \mathrm{Time}(\mathsf{A}) + \kappa_m Q_{\mathsf{H}} \approx \mathrm{Time}(\mathsf{A})$.

B Applications to Dilithium-QROM

In this section, we recall the Dilithium-QROM signature scheme introduced by Kiltz et al. [KLS18] and take a new look at its security properties. We apply results from Section 3, so that Dilithium-QROM supports not only prime moduli q satisfying $q \equiv 5 \pmod 8$ but also any other primes.

$\underline{\mathsf{IGen}(\mathsf{par})}$

$\rho \leftarrow \{0,1\}^{256}$
$\mathbf{A} \leftarrow R_q^{k\times\ell} := \mathsf{Sam}(\rho)$
$(\mathbf{s}_1, \mathbf{s}_2) \leftarrow S_\eta^\ell \times S_\eta^k$
$\mathbf{t} := \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$
$\mathbf{t}_1 := \mathsf{Power2Round}_q(\mathbf{t}, \delta)$
$\mathbf{t}_0 := \mathbf{t} - \mathbf{t}_1 \cdot 2^\delta$
$pk = (\rho, \mathbf{t}_1, \boxed{\mathbf{t}_0})$
$sk = (\rho, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$
return (pk, sk)

$\underline{\mathsf{P}_1(sk)}$

$\mathbf{A} \leftarrow R_q^{k\times\ell} := \mathsf{Sam}(\rho)$
$\mathbf{y} \leftarrow S_{\gamma'-1}^\ell$
$\mathbf{w} := \mathbf{A}\mathbf{y}$
$\mathbf{w}_1 := \mathsf{HighBits}_q(\mathbf{w}, 2\gamma)$
return $(W = \mathbf{w}_1, St = (\mathbf{w}, \mathbf{y}))$

$\underline{\mathsf{P}_2(sk, W = \mathbf{w}_1, c, St = (\mathbf{w}, \mathbf{y}))}$

$\mathbf{z} := \mathbf{y} + c\mathbf{s}_1$
if $\|\mathbf{z}\|_\infty \geq \gamma' - \beta$ **or** $\|\mathsf{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma)\|_\infty \geq \gamma - \beta$
  then $(\mathbf{z}, \mathbf{h}) := \perp$
else $\mathbf{h} := \mathsf{MakeHint}_q(-c\mathbf{t}_0, \mathbf{w} - c\mathbf{s}_2 + c\mathbf{t}_0, 2\gamma)$
return $Z = (\mathbf{z}, \mathbf{h})$

$\underline{\mathsf{V}(pk, W = \mathbf{w}_1, c, Z = (\mathbf{z}, \mathbf{h}))}$

return $[\![\|\mathbf{z}\|_\infty < \gamma' - \beta]\!]$ **and** $[\![\mathbf{w}_1 = \mathsf{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^\delta, 2\gamma)]\!]$

Figure 10: Dilithium-QROM identification scheme [KLS18]. We point out that the $\mathbf{t}_0$ part of the public key is assumed to be known by the adversary in the security proofs, but is not needed by the verifier for verification.

The main difference between the modified Bai-Galbraith scheme presented in Section 4, and Dilithium-QROM is "removing" the low order bits from $\mathbf{t}$ and using Lemmas 2.1 and 2.2. Consequently, they significantly reduce the size of the public key. This issue affects our security results since we now have to apply Theorem 3.10 (which is kind of a weaker bound) instead of Theorem 3.8.

B.1 Identification Protocol

The algorithms for identification protocol $\mathsf{ID} = (\mathsf{IGen}, \mathsf{P}_1, \mathsf{P}_2, \mathsf{V})$ are described in Figure 10 with the concrete parameters $\mathsf{par} = (q, d, n, k, \ell, \delta, \gamma, \gamma', \eta, \beta)$ given in Table 3 and Table 4. We set the challenge space as in Section 4, i.e.

$$\mathsf{ChSet} := \{c \in R \mid \|c\|_\infty = 1 \text{ and } \|c\| = \sqrt{46}\}.$$

Key Generation. As before, the key generation starts by choosing a random 256-bit seed ρ and expanding into a matrix $\mathbf{A} \in R_q^{k\times\ell}$ by an extendable output function Sam modeled as a random oracle. The secret keys $(\mathbf{s}_1, \mathbf{s}_2) \in S_\eta^\ell \times S_\eta^k$ have uniformly random coefficients between $-\eta$ and η. Then, the value $\mathbf{t} = \mathbf{A}\mathbf{s}_1 + \mathbf{s}_2$ is computed. The public key that is needed for verification is now $(\rho, \mathbf{t}_1)$ with $\mathbf{t}_1$ output by the $\mathsf{Power2Round}_q(\mathbf{t}, \delta)$ algorithm in Figure 2 (we have $\mathbf{t} = \mathbf{t}_1 \cdot 2^\delta + \mathbf{t}_0$ for some small $\mathbf{t}_0$), while the secret key is $(\rho, \mathbf{s}_1, \mathbf{s}_2, \mathbf{t}_0)$.

Even though the verifier never needs the value $\mathbf{t}_0$ (and thus it does not need to be included in the public key of the actual scheme), we do need this value in order to simulate transcripts (for the non-abort honest verifier zero-knowledge part). Hence, the security of our scheme is based on the fact that the adversary gets $\mathbf{t}_1$ and $\mathbf{t}_0$, whereas in reality he only gets $\mathbf{t}_1$.

The set ChSet is defined as in Equation (11), and $\mathsf{ZSet} = S_{\gamma'-\beta-1}^\ell \times \{0,1\}^k$ and the set of commitments WSet is defined as $\mathsf{WSet} = \{\mathbf{w}_1 \ : \ \exists \mathbf{y} \in S_{\gamma'-1}^\ell \text{ s.t. } \mathbf{w}_1 = \mathsf{HighBits}_q(\mathbf{A}\mathbf{y}, 2\gamma)\}$.

Protocol Execution. Similarly as before, the prover starts the identification protocol by reconstructing $\mathbf{A}$ from the random seed ρ. Then, he samples $\mathbf{y} \leftarrow S_{\gamma'-1}^\ell$ and later computes $\mathbf{w} = \mathbf{A}\mathbf{y}$. Next, the prover writes $\mathbf{w} = 2\gamma \cdot \mathbf{w}_1 + \mathbf{w}_0$, with $\mathbf{w}_0$ between $-\gamma$ and γ, and then sends $\mathbf{w}_1$ to the verifier. The verifier generates a random challenge $c \leftarrow \mathsf{ChSet}$ and sends it to the prover. The prover computes $\mathbf{z} = \mathbf{y} + c\mathbf{s}$. If $\mathbf{z} \notin S_{\gamma'-\beta-1}^\ell$, then the prover sets his response to $\perp$. He also replies with $\perp$ if $\mathsf{LowBits}_q(\mathbf{w} - c\mathbf{s}_2, 2\gamma) \notin S_{\gamma-\beta-1}^k$. Then, he sends $\mathbf{z}$ as well as a "hint" $\mathbf{h}$ which will allow the verifier to compute $\mathsf{HighBits}_q(\mathbf{A}\mathbf{z} - c\mathbf{t}, 2\gamma)$.

Eventually, the verifier checks whether $\|\mathbf{z}\|_\infty < \gamma' - \beta$ and that $\mathbf{A}\mathbf{z} - c\mathbf{t}_1 \cdot 2^\delta$ together with the hint $\mathbf{h}$ allow him to reconstruct $\mathbf{w}_1$.

B.2 Security Analysis

We omit some aspects of security analysis for Dilithium-QROM since they are identical to the proofs in [KLS18]. In particular, we skip the proof of correctness and non-abort honest verifier zero-knowledge properties.

Lemma B.1 *If* $\beta \geq \max_{s \in S_\eta, c \in \mathsf{ChSet}} ||cs||_\infty$*, then* ID *is perfectly* naHVZK *and has correctness error* $\nu \approx 1 - \exp(-\beta n \cdot (k/\gamma + \ell/\gamma'))$*.*

Lossyness, min entropy and computation unique response properties follow using methods from Section 4 (or [KLS18]). Therefore, we only provide proof sketch below. As before, we use the definitions of sets W_i from Section 3.3.

Lemma B.2 *Let* e_ℓ *(resp.* e'_ℓ*) be the largest integer which satisfies* $q^{e_\ell/d} \leq 2\sqrt{46}$ *(resp.* $q^{e'_l/d} \leq 2(\gamma' - \beta - 1)\sqrt{n}$*). Then,* ID *has* $\varepsilon_{\mathsf{ls}}$*-lossy soundness, where*

$$\begin{aligned}\varepsilon_{\mathsf{ls}} \leq \frac{1}{|\mathsf{ChSet}|} &+ |\mathsf{ChSet}|^2 \cdot ((\frac{(8\gamma+5)\cdot 2^\delta}{q^{(1-e_l/d)}})^{nk} + (\frac{(8\gamma+5)^k(4\gamma'-4\beta-3)^\ell}{q^k})^n) \\ &+ |\mathsf{ChSet}|^2 \cdot (\sum_{i=1}^{e'_l} \frac{\binom{d}{i}}{|W_i|^\ell} \cdot (\frac{(8\gamma+5)^k \cdot (4\gamma'-4\beta-3+2||W_i||_\infty)^\ell}{q^{k(1-i/d)}})^n)\end{aligned} \tag{17}$$

Proof. Consider an unbounded adversary C that is executed in game LOSSY-IMP of Figure 11 and suppose that for some $\mathbf{w}_1$, there exist two $c \neq c' \in \mathsf{ChSet}$ and two $(\mathbf{z}, \mathbf{h}), (\mathbf{z}', \mathbf{h}')$ that lead to C winning.

```
GAME LOSSY-IMP:
01 pk_ls := (ρ, t_1, t_0) ← LossyIGen(par)
02 (w_1, St) ← C(pk_ls)
03 c ← ChSet
04 (z, h) ← C(St, c)
05 return [[w_1 = UseHint_q(h, Az − c t_1 · 2^δ, 2γ)]] and [[||z||_∞ < γ' − β]]
```

Figure 11: The lossy impersonation game LOSSY-IMP [KLS18].

Thus, $\|\mathbf{z}\|_\infty, \|\mathbf{z}'\|_\infty < \gamma' - \beta$ and

$$\begin{aligned}\mathbf{w}_1 &= \mathsf{UseHint}_q(\mathbf{h}, \mathbf{A}\mathbf{z} - \mathbf{t}_1 c \cdot 2^\delta, 2\gamma), \\ \mathbf{w}_1 &= \mathsf{UseHint}_q(\mathbf{h}', \mathbf{A}\mathbf{z}' - \mathbf{t}_1 c' \cdot 2^\delta, 2\gamma).\end{aligned}$$

By Lemma 2.1, we have

$$\begin{aligned}\|\mathbf{A}\mathbf{z} - \mathbf{t}_1 c \cdot 2^\delta - \mathbf{w}_1 \cdot 2\gamma\|_\infty &\leq 2\gamma + 1, \\ \|\mathbf{A}\mathbf{z}' - \mathbf{t}_1 c' \cdot 2^\delta - \mathbf{w}_1 \cdot 2\gamma\|_\infty &\leq 2\gamma + 1.\end{aligned}$$

Therefore,

$$\|\mathbf{A}(\mathbf{z} - \mathbf{z}') - \mathbf{t}_1 \cdot 2^\delta \cdot (c - c')\|_\infty \leq 4\gamma + 2\,,$$

which can be rewritten as

$$\mathbf{A}(\mathbf{z} - \mathbf{z}') + \mathbf{u} = \mathbf{t}_1 \cdot 2^\delta \cdot (c - c') \tag{18}$$

for some $\mathbf{u}$ such that $\|\mathbf{u}\|_\infty \leq 4\gamma + 2$ (and $\|\mathbf{z} - \mathbf{z}'\|_\infty \leq 2(\gamma' - \beta - 1)$).

If $\mathbf{A} \leftarrow R_q^{k \times \ell}$ and $\mathbf{t} \leftarrow R_q^k$, then Theorem 3.10 tells us that Equation (18) is satisfied with probability less than

$$\begin{aligned}&|\mathcal{D}|\cdot((\frac{(8\gamma+5)\cdot 2^\delta}{q^{(1-e_l/d)}})^{nk} + (\frac{(8\gamma+5)^k(4\gamma'-4\beta-3)^\ell}{q^k})^n) + \\ &|\mathcal{D}| \cdot (\sum_{i=1}^{e'_l} \frac{\binom{d}{i}}{|W_i|^\ell} \cdot (\frac{(8\gamma+5)^k \cdot (4\gamma'-4\beta-3+2||W_i||_\infty)^\ell}{q^{k(1-i/d)}})^n),\end{aligned} \tag{19}$$

q	d	γ
$2^{45} - 21283$	2	905679
$2^{47} - 12535$	4	328911
$2^{47} - 4591$	8	326472
$2^{47} - 2271$	16	326704
$2^{47} - 8767$	32	320520
$2^{47} - 16255$	64	329226
$2^{47} - 12031$	128	322944
$2^{47} - 5631$	256	307232
$2^{47} - 23551$	512	285891

Table 3: Prime moduli q for each possible value of d. For each case, we also provide values γ such that $2\gamma | q-1$. Just like in [KLS18] and in Section 4, we set $\gamma' = \gamma$.

where $\mathcal{D} := \{c - c' : c, c' \in \mathsf{ChSet}\} \setminus \{0\}$. Thus, except with the above probability, for every $\mathbf{w}_1$, there is at most one possible c that allows C to win. Hence, except with the above probability, C has at most a $1/|\mathsf{ChSet}|$ chance of winning.

□

We omit proofs of the following lemmas since they are identical to the ones for Lemma 4.3 and Lemma 4.4.

Lemma B.3 *Let e_m be the largest integer which satisfies $q^{e_m/d} \leq 2\gamma'\sqrt{n}$. Then the identification scheme* ID *in Figure 10 has*

$$\alpha > \log\left(\min\left\{\frac{1}{M}, (2\gamma' - 1)^{n\ell}\right\}\right)$$

bits of min-entropy, where

$$M := \frac{|S_{2\gamma'}|^\ell \cdot |S_{2\gamma}|^k}{q^{nk}} + \sum_{i=1}^{e_m} \frac{\binom{d}{i} \cdot |S_{2\gamma' + \| W_i \|_\infty}|^\ell \cdot |S_{2\gamma + \| W_i \|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}.$$

Lemma B.4 *Let e_c be the largest integer such that $q^{e_c/d} \leq 2(\gamma' - \beta)\sqrt{n}$. Then*

$$\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}) \leq \frac{|S_{2(\gamma'-\beta)}|^\ell \cdot |S_{4\gamma+2}|^k}{q^{nk}} + \sum_{i=1}^{e_c} \frac{\binom{d}{i} \cdot |S_{2(\gamma'-\beta) + \| W_i \|_\infty}|^\ell \cdot |S_{4\gamma+2+\| W_i \|_\infty}|^k}{|W_i|^{\ell+k} \cdot q^{nk(1-i/d)}}$$

for all (even unbounded) adversaries A.

B.3 Concrete Parameters

Recall that Dilithium-QROM is obtained by applying the Fiat-Shamir transform on the ID scheme and using Sam as a pseudorandom function. Kiltz et al. give concrete parameters (see Table 4, $d = 2$) for Dilithium-QROM which provides 128 bits of quantum security (using similar argument as in Section 4.3). In particular, they set $q \equiv 5 \pmod 8$ such that they can apply the main result from [LS18]. We propose new instantiations of Dilithium-QROM for $d > 2$ with concrete parameters in Table 3 and Table 4 along with their security properties.

We observe that already for $d = 4$ we pick much larger modulus q and dimensions (k, ℓ) than in [KLS18]. The reason behind it is that the bound in Theorem 3.10 (consequently, Lemma B.2) is not tight due to avoiding problems related to the Power2Round function. Hence, once e'_l gets slightly bigger, then $\varepsilon_{\mathsf{ls}}$ rockets. We observe that increasing the prime modulus only does not solve the issue. Indeed, having larger q implies less secure MLWE problem [9]. Obviously, picking lower γ and γ' results in getting much

[9]For $q \approx 2^{45}$, the best known quantum bit-cost drops to only 119.

d	2	4	8	16	32	64	128	256	512
n	512	512	512	512	512	512	512	512	512
(k, ℓ) (dimensions of $\mathbf{A}$)	$(4,4)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$	$(5,5)$
# of $\pm 1's$ in $c \in \mathsf{ChSet}$	46	46	46	46	46	46	46	46	46
δ (dropped bits)	15	13	13	13	13	13	13	13	13
η (max. coeff. of $\mathbf{s}_1, \mathbf{s}_2$)	7	2	2	2	2	2	2	2	2
$\beta(= \eta \cdot (\#\text{of 1's in } c))$	322	92	92	92	92	92	92	92	92
e_ℓ (lossyness)	0	0	0	1	2	5	10	20	40
e'_ℓ (lossyness)	0	2	4	8	16	32	64	129	257
e_c (CUR)	0	2	4	8	16	32	64	129	257
e_m (min-entropy)	0	2	4	8	16	32	64	129	257
$\log(\varepsilon_{\mathsf{ls}})$	−264	−264	−264	−264	−264	−264	−264	−264	−264
$\log(\mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}))$	−865	−13685	−7447	−7244	−7381	−6641	−6759	−7077	−7508
α	2913	16244	10006	9803	9940	9200	9318	9636	10067
pk size (kilobytes)	7.71	10.91	10.91	10.91	10.91	10.91	10.91	10.91	10.91
sig size (kilobytes)	5.69	6.76	6.76	6.76	6.76	6.76	6.76	6.76	6.76
Exp. Repeats $\frac{1}{1-\nu}$	4.3	4.19	4.23	4.23	4.35	4.19	4.3	4.63	5.19
BKZ block-size to break LWE	480	550	550	550	550	550	550	550	550
Best known classical bit-cost	140	160	160	160	160	160	160	160	160
Best known quantum bit-cost	127	145	145	145	145	145	145	145	145

Table 4: Parameters for the Dilithium-QROM [KLS18] scheme for different values of $d \in \{2^i : i \in [9]\}$. Variables $e_l, e'_l, e_c, e_m, \alpha, \varepsilon_{\mathsf{ls}}, \mathrm{Adv}_{\mathsf{ID}}^{\mathsf{CUR}}(\mathsf{A}), \nu$ are defined in Section B.2.

higher number of repetitions [10]. Therefore, we must choose larger dimensions $(k, \ell) = (5, 5)$ of the matrix $\mathbf{A}$ in order to both keep $\varepsilon_{\mathsf{ls}}$ small and have 128 bit quantum security for MLWE. We also pick $q = 2^{47}$, $\gamma \approx 3 \cdot 10^5$ and $\delta = 13$ (dropped bits from $\mathbf{t}$). This results in getting larger public keys and signature sizes, namely 10.91kB and 6.76kB respectively.

Similarly for all $d \geq 8$, we select $q \approx 2^{47}$ so that $\varepsilon_{\mathsf{ls}} \leq 2^{-264}$. For larger d, we slightly lower γ in order to maintain the lossyness property. This, however, only marginally affects the runtime of this scheme. In all cases, we would need to repeat the protocol at most (around) five times. This observation combined with the support of the Number Theoretic Transform algorithm for $d = n = 512$ makes sure that the protocol can be executed efficiently.

Sizes of signatures as well as the public keys are the same for each $d \geq 4$ since all γ's we picked are close to each other. Unfortunately, they are respectively 1.07kB and 3.2kB larger than in the original Dilithium-QROM scheme.

[10] We recall that we need the condition $2^{\delta-1} \cdot \kappa < \gamma$ in order to apply Lemma 2.1, where κ is the maximal number of non-zero coefficients of a polynomial in ChSet.