On Removing Rejection Conditions in Practical Lattice-Based Signatures

Rouzbeh Behnia[1], Yilei Chen[2], and Daniel Masny[3]

[1] University of South Florida
[2] Tsinghua University
[3] Visa Research

Abstract. Digital signatures following the methodology of "Fiat-Shamir with Aborts", proposed by Lyubashevsky, are capable of achieving the smallest public-key and signature sizes among all the existing lattice signature schemes based on the hardness of the Ring-SIS and Ring-LWE problems. Since its introduction, several variants and optimizations have been proposed, and two of them (i.e., Dilithium and qTESLA) entered the second round of the NIST post-quantum cryptography standardization. This method of designing signatures relies on rejection sampling during the signing process. Rejection sampling is crucial for ensuring both the correctness and security of these signature schemes.
In this paper, we investigate the possibility of removing the two rejection conditions used both in Dilithium and qTESLA. First, we show that removing one of the rejection conditions is possible, and provide a variant of Lyubashevsky's signature with comparable parameters with Dilithium and qTESLA. Second, we give evidence on the difficulty of removing the other rejection condition, by showing that two very general approaches do not yield a signature scheme with correctness or security.

1 Introduction

The emergence of quantum computers has made the development of signatures with post-quantum security a necessity. A promising source of post-quantum hardness is computational intractability assumptions on lattices. Common assumptions are the hardness of learning with errors (LWE) problem, and the short integer solution (SIS) problem, which are both related to solving the shortest vector problem in a lattice [1,35].

The origin of lattice-based signatures can be traced back to the proposal of Goldreich et al. [21] and the NTRU signature scheme [24]. They use the "trapdoor approach". Namely, they let the public verification key and the secret signing key (the trapdoor) be a "bad" basis and a "good" basis of a lattice, respectively. However, the initial schemes were broken since signatures leaked information about the secret "good" basis. By obtaining sufficiently many signatures, the secret "good" basis could be completely recovered [20,33].

Lattice signatures following the trapdoor approach were fixed by the seminal work of Gentry et al. [19]. Their trapdoor mechanism allowed to produce a

signature securely without leaking the signing key following the full-domain hash paradigm. When relying on the NTRU problem, this approach leads to more efficient signatures. However, this approach has been less competitive for LWE or SIS, since it led to large key and signature sizes.

A different method for constructing digital signatures is through the Fiat-Shamir [17] transformation. This technique uses a random oracle to transform an interactive identification protocol to a digital signature, which is non-interactive. The challenge of constructing a lattice based identification protocol is that the security of LWE and SIS inherently rely on the fact that a solution does not only need to have a particular algebraic structure, but it also needs to be small. Finding a large solution is easy for a SIS instance as well as in case of an LWE instance when the noise term is treated as a part of the solution. This is a significant difference compared to the realm of cyclic groups and the discrete logarithm assumption, where efficient identification protocols exist, e.g., the Schnorr identification protocol [36].

A key principle of the Schnorr identification protocol is to rerandomize a discrete logarithm problem instance and expose the rerandomized solution as an evidence of authenticity, or in case of a digital signature, as a signature. This can be efficiently done using a uniform masking term. Unfortunately, this does not translate to the lattice realm, since the verification mechanism that checks the validity of a signature needs to ensure that a solution is small. A uniform masking term would make the signature large, hence, one would need to use a small masking term, which when applied in a straightforward fashion, would expose parts of the secret key.

In his paper "Fiat-Shamir with Aborts", Lyubashevsky [27] has overcome this obstacle. One of his key findings is the idea of aborting, in case information of the secret key is leaked. This process of rejection and resampling, i.e. rejection sampling, helps to ensure correctness and security and has led to a fruitful line of signature schemes based on the LWE and SIS problems [28,22,15,5,16,25,4].

Nevertheless, this does not lead to a smooth adaptation of signatures with additional properties, such as blind signatures [9], multi signatures [7] and threshold signatures [11,12]. Moreover, there is a concern of potential side-channel attacks. Countermeasures for such attacks have been studied [32]. Nevertheless, providing less attack surface to such attacks would be preferable.

While rejection sampling has provided a solution for efficient constructing lattice based signatures, it constitutes at the same time an obstacle. Rejection based lattice signatures have in common that their rejection mechanism achieves integrity and security by ensuring two rejection conditions. The focus of this work is to investigate the necessity of these two rejection conditions for Fiat-Shamir based lattice signatures. The question we want to answer is:

How crucial are the two rejection conditions for efficiency when applying the "Fiat-Shamir" paradigm to lattices?

1.1 Contributions

We show both positive and negative results on removing the rejection conditions in Fiat-Shamir based lattice signatures. Out of the two rejection conditions used both in Dilithium and qTESLA, we show that removing one of the rejection conditions is possible. As a result, we provide a variant of Lyubashevsky's signature with one rejection condition.

The variant of Fiat-Shamir based lattice signature we propose can be instantiated with comparable parameters with Dilithium and qTESLA in terms of security, public-key and signature sizes, and rejection rate. The key difference to the previous schemes is that the secret key and masking terms are sampled uniformly random over the base ring.

The remaining rejection condition in our signature scheme is used to ensure that the first message (the commitment of the masking term) in the 3 round lattice-based ID protocol can be recovered from the rest of the transactions. We show that this rejection conditions is unlikely to be removed when the scheme uses a polynomially large modulus. First, we adopt a recent result by Guo et al. [23], which states that there is no non-interactive reconciliation mechanism for lattice based key exchange protocols [14,34]. In our case this translates to the fact that there is no reconciliation mechanism without taking additional hints about the first message. We then take a step further and consider reconciliation mechanisms that takes hints. Indeed, the reconciliation mechanisms with hints used in lattice based key exchange protocols [14,34] can be adopted in our signature scheme to remove the rejection condition and provide correctness. Unfortunately both reconciliations mechanisms are not reusable under the same initial key exchange messages [18,13] which could be considered as a public key. Since a signature scheme needs to allow polynomially many signatures per public key, both reconciliation mechanisms are not sufficient for our purposes. Even further, we show an attack against a wide ranged class of potential reconciliation mechanisms.

1.2 Technical Overview

Let us recall the idea of "Fiat-Shamir with Aborts" from a more technical perspective. We present a simplified version of Bai and Galbraith's scheme [5] based on Lyubashevsky's "Fiat-Shamir with Aborts" paradigm [28], which is followed by Dilithium and qTESLA [16,4]. In the overview we assume the base ring is $\mathbb{Z}_q$. The final scheme is instantiated on a polynomial ring.

The public key pk consists of a uniform $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{y} \approx \mathbf{sA} \mod q$, where $\mathbf{s} \in \mathbb{Z}_q^{n \times n}$ is sampled from a distribution such that the norm of $\mathbf{s}$ is small (the typical choices are uniform or Gaussian with small standard deviations). Let $\lfloor \cdot \rfloor_p$ be the rounding function that drops the $\log q - \log p$ least significant bits.

To sign a message m, there are two steps to follow:

1. Sample a small $\mathbf{r} \in \mathbb{Z}_q^n$, compute $H(\lfloor \mathbf{r}^t \mathbf{A} \rfloor_p, \mathsf{pk}, \mathsf{m}) = \mathbf{c}$, where $\mathbf{c}$ is a small vector in $\mathbb{Z}_q^n$.

2. Compute $\mathbf{z}^t = \mathbf{r}^t + \mathbf{c}^t\mathbf{s} \in \mathbb{Z}_q^n$, then check the following two conditions. If they are satisfied, output signature $(\mathbf{z}, \mathbf{c})$; if they are not satisfied, restart from the first line.
   (a) $\mathbf{r}^t + \mathbf{c}^t\mathbf{s}$ is sufficiently small and does not leak the secret $\mathbf{s}$.
   (b) $\lfloor \mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{y} \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$.

The verification algorithm accepts a signature $(\mathbf{z}, \mathbf{c})$ if and only if $\mathbf{z}$ is sufficiently small and $H(\lfloor \mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{y} \rfloor_p, \mathsf{pk}, \mathsf{m}) = \mathbf{c}$.

For the security of the scheme, it is important that $\mathbf{z}$ is small, since only then would breaking the scheme lead to solving the SIS problem. But this presents a challenge, since in addition to $\mathbf{c}^t$ and $\mathbf{z}$, $\mathbf{r}$ needs to be small too. Therefore $\mathbf{r}$ might not completely hide the sensitive term $\mathbf{c}^t\mathbf{s}$ when publishing $\mathbf{z}$. A carefully tailored rejection sampling, i.e., Step 2.(a), resolves this issue.

The second rejection, i.e., Step 2.(b), seems to be important mostly for correctness. For correctness, the verification algorithm needs to be able to recover the same $\mathbf{c}$, by hashing $\lfloor \mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{y} \rfloor_p$ as the signing algorithm which hashes $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$. This is only the case when $\lfloor \mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{y} \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$. We will observe that this step is crucial for security as well.

Removing Rejection condition 2.(a). In the scheme proposed in this paper, the signing algorithm samples $\mathbf{r}$ uniformly at random from $\mathbb{Z}_q^n$, instead of sampling $\mathbf{r}$ with small norm. The signing algorithm then computes $H(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p, \mathsf{pk}, \mathsf{m}) = \mathbf{c}$, $\mathbf{z}^t = \mathbf{r}^t + \mathbf{c}^t\mathbf{s}$ and rejects if $\lfloor \mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{y} \rfloor_p \neq \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$. As a result, we are able to remove Step 2.(a) in the signing algorithm. Consequently, the verification algorithm no longer checks whether $\mathbf{z}$ is small.

The security of the scheme relies on the fact that the public key is indistinguishable from uniform based on the LWE assumption and for a uniform public key. Forging a signature in our scheme is related to finding a vector $\mathbf{r} \in \mathbb{Z}_q^n$, given a uniform $\mathbf{A} \in \mathbb{Z}_q^{n \times m}$ and $\mathbf{y} \in \mathbb{Z}_p^m$ such that $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p = \mathbf{y}^t$. Depending on the choice of parameters p, q, m and n, the hardness of this problem varies from trivial, computationally hard or even statistically hard (when $\mathbf{y}$ is uniformly random). The problem of finding $\mathbf{r}$, is identical to finding $\mathbf{r}$ and a sufficiently small noise term $\mathbf{e}$, which is related to the rounding function, such that $\mathbf{r}^t\mathbf{A} + \mathbf{e}^t = \tilde{\mathbf{y}}^t$, for $\tilde{\mathbf{y}} \in \mathbb{Z}_q^m$. We therefore refer to this problem as Bounded Distance Decoding. In our actual scheme, we use a more tailored problem that we denote as Adaptive BDD (ABDD). Like BDD, ABDD can be computationally or statistical hard. Our scheme also accomplishes a tight security proof and security in the quantum random oracle model (QROM) for the reasons pointed out by Unruh [37] and Kiltz, Lyubashevsky and Schaffner [25].

Let us remark that one can also choose q to be sufficiently larger than p such that the rounding function is more likely to round away the noise term affecting $\mathbf{y}$. The probability of a rejection can be made negligibly small, when choosing a super-polynomially large modulus q. But choosing a large modulus q makes the scheme inefficient.

Evidence on the difficulty of removing Rejection condition 2.(b). Compared to Rejection condition (a), Rejection condition (b) seems much harder to remove

without sacrificing the efficiency of the scheme. In fact, we show two general approaches of removing Rejection condition (b) fail.

In the first approach, we consider the possibility of constructing functions g and f that map $\mathbf{r}^t\mathbf{A}$ and $\mathbf{r}^t\mathbf{A}+\hat{\mathbf{e}}^t$ for any bounded error term $\hat{\mathbf{e}}$ to the same value, i.e. $g(\mathbf{r}^t\mathbf{A}) = f(\mathbf{r}^t\mathbf{A}+\hat{\mathbf{e}}^t)$, while ensuring that $g(\mathbf{r}^t\mathbf{A})$ serves as a commitment of $\mathbf{r}$, or at least preserves high min-entropy; then we can apply the hash function on $g(\mathbf{r}^t\mathbf{A})$ instead of $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$. However, one can show such functions g, f do not exist when the modulus q is polynomially large. The result follows by the one of Guo et al. [23], which shows a similar impossibility result for the lattice-based key exchange protocols [14,34].

In the second approach, we try to adapt the reconciliation mechanism used in lattice-based key-encapsulation mechanisms [14,34]. While the reconciliation mechanisms can be adapted in our signature scheme to provide correctness, we show that they leak information about the error term $\mathbf{e}$ in the public key. More generally, we rule out the possibility of achieving security when a string of the form of $\mathbf{r}^t\mathbf{A} + \hat{\mathbf{e}}^t$, where $\hat{\mathbf{e}}$ is bounded and independent of the challenge $\mathbf{c}$, is recovered from any potential reconciliation mechanism (no matter if the mechanism is the same from the ones in [14,34] or not).

We also discuss the generalizations and limitations of these two types of negative results.

Concrete parameters. For the variant of Fiat-Shamir lattice-based signature scheme we propose, we give a detailed analysis of the hardness of ABDD and appropriate parameter choices in Section 3 and Section 7. Even though, we are able to remove the rejection condition on $\mathbf{z}$ being small, the larger dimension m causes a significant loss of efficiency compared to Dilithium and qTESLA in the random oracle model (ROM). When considering tight security reductions, i.e. when ABDD is statistically hard, the efficiency, security and rejection rate are comparable to those of Dilithium-QROM and qTESLA-provable, for carefully chosen parameters. We give concrete parameter choices for our scheme and a comparison to Dilithium and qTESLA in Section 7.

Organization. In Section 4, we present the signature scheme with one rejection condition and analyze its correctness. In Section 5, we prove its security under the LWE and ABDD assumptions. In Section 6, we present the negative results of removing the other rejection condition. In Section 7, we instantiate our scheme with parameter choices for different levels of security and compare it with Dilithium [16] and qTESLA [4].

2 Preliminaries

Notations. We use κ to denote the security parameter and $v \xleftarrow{\$} S$ for a uniformly random variable v over set S. $\approx_s$ and $\approx_c$ denote statistically close and computationally indistinguishable, respectively. For $i \in \mathbb{N}$, we use $[i]$ to denote $\{1, \ldots, i\}$. For a modulus $q \in \mathbb{N}$, we denote $\mathbb{Z}/q\mathbb{Z}$ by $\mathbb{Z}_q$ and represent $\mathbb{Z}_q$ by $[-\lfloor \frac{q}{2} \rfloor, \lfloor \frac{q-1}{2} \rfloor]$.

Vectors are written as a bold lower-case letter (e.g. $\mathbf{v}$) and its i^{th} component is v_i. A matrix is written as a bold capital letter (e.g. $\mathbf{A}$) and its i^{th} column vector is $\mathbf{a}_i$. The ℓ_p-norm is $\|\mathbf{v}\|_p := (\sum v_i^p)^{1/p}$ and the infinity norm is $\|\mathbf{v}\|_\infty := \max_i\{|v_i|\}$. The length of a matrix is the norm of its longest column: $\|\mathbf{A}\|_p := \max_i \|\mathbf{a}_i\|_p$. For a random variable X, $H_\infty(X) := -\log(\max_x \Pr[X = x])$ is the min-entropy.

2.1 Lattices

An n-dimensional lattice Λ of rank $k \leq n$ is a discrete additive subgroup of $\mathbb{R}^n$. Given k linearly independent basis vectors $\mathbf{B} = \{\mathbf{b}_1, ..., \mathbf{b}_k \in \mathbb{R}^n\}$, the lattice generated by $\mathbf{B}$ is the following:

$$\Lambda(\mathbf{B}) = \Lambda(\mathbf{b}_1, ..., \mathbf{b}_k) = \left\{ \sum_{i=1}^{k} x_i \cdot \mathbf{b}_i, x_i \in \mathbb{Z} \right\}.$$

The following presents the decisional learning with errors (LWE) problem.

Definition 1 (Decisional learning with errors [35]). *For $n, m \in \mathbb{N}$ and modulus $q \geq 2$ and a noise distribution χ over $\mathbb{Z}$. An $\mathsf{LWE}_{n,m,q,\chi}$ sample is obtained from sampling $\mathbf{s} \xleftarrow{\$} \mathbb{Z}_q^n$, $\mathbf{A} \xleftarrow{\$} \mathbb{Z}_q^{n\times m}$, $\mathbf{e} \xleftarrow{\$} \chi^m$, and outputting $(\mathbf{A}, \mathbf{y}^t := \mathbf{s}^t\mathbf{A} + \mathbf{e}^t \mod q)$.*

We say $\mathsf{LWE}_{n,m,q,\chi}$ is (t, ϵ)-secure if for any ppt algorithm $\mathcal{A}$ with running time t,

$$|\Pr[\mathcal{A}^{\mathcal{O}_{\mathsf{LWE}}}(1^\kappa) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_{\mathsf{uniform}}}(1^\kappa) = 1]| \leq \epsilon,$$

where oracle $\mathcal{O}_{\mathsf{LWE}}$ outputs $\mathsf{LWE}_{n,m,q,\chi}$ samples and oracle $\mathcal{O}_{\mathsf{uniform}}$ outputs uniform samples over $\mathbb{Z}_q^{n\times m} \times \mathbb{Z}_q^m$.

Cyclotomic rings. Let $\phi_{2d}(x)$ be the $2d$-th cyclotomic polynomial where d is a power of two (i.e., $\phi_{2d}(x) := x^d + 1$), $R := \mathbb{Z}[x]/\phi_{2d}(x)$ be the associated cyclotomic ring, and $q \geq 2$ be an integer modulus so that $R_q := R/qR$ is the polynomial ring with coefficients modulo q. Usually q is taken as a prime with $q \equiv 1 \mod 2n$ in order to use the Number Theoretic Transform (NTT) [29].

Geometrically, we are concerned with the *coefficient embedding* where each polynomial $h(x) \in R$ is associated with its ordered vector of coefficients,

$$h(x) \mapsto \mathbf{h} = (h_1, \ldots, h_d) \in \mathbb{Z}^d. \tag{1}$$

Multiplication by a polynomial $g(x) = \sum_{i=1}^{d} g_i x^i \in R$ is given by the anti-cyclic matrix

$$\psi(g) := \begin{bmatrix} g_1 & -g_d & -g_{d-1} & \cdots & -g_2 \\ g_2 & g_1 & -g_d & \cdots & -g_3 \\ g_3 & g_2 & g_1 & \cdots & -g_4 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ g_d & g_{d-1} & g_{d-2} & \cdots & g_1 \end{bmatrix} \in \mathbb{Z}^{d\times d}.$$

Therefore, we have $\psi(g) \cdot \mathbf{h}$ as the coefficient embedding of $h(x) \cdot g(x) \in R$. In fact, the map $\psi : R \to \mathbb{Z}^{d\times d}$ taking a polynomial and mapping it to its anti-cyclic matrix is a (another) ring embedding, from R to $d \times d$ integer matrices. The same holds when we consider the quotient R_q, except now the previous is expressed via operations modulo q. An important feature of the coefficient embedding is that scalars in the ring, $\alpha \in \mathbb{Z} \subset R$ (or $\mathbb{Z}_q \subset R_q$), embed as scalar matrices, $\psi(\alpha) = \alpha \cdot \mathbf{I}_d \in \mathbb{Z}^{d\times d}$ (or $\mathbb{Z}_q^{d\times d}$). Further, vectors over R or R_q have an entry-wise coefficient embedding: $\mathbf{v} = (v_1, \ldots, v_l) \in R^l \mapsto (v_{1,1}, \ldots, v_{1,d}, \ldots, v_{l,1}, \ldots, v_{l,d}) \in \mathbb{Z}^{dl}$. We suggest the reader regularly traverse between the rings (R or R_q) and their coefficient embedding, in $\mathbb{Z}^d$ or $\mathbb{Z}_q^d$, as well as their representations as linear transformations, while reading this article. Lastly, our norm will be the normal Euclidean norm under the the coefficient embedding, $\|\mathbf{x}\| = \|\psi(\mathbf{x})\|_2$.

For the generality and modularity of the presentation, we present our main scheme under the module-LWE problem.

Definition 2 (Module-LWE [8,26]). *Let $l, k \in \mathbb{N}$ and modulus $q \geq 2$. Let $R := \mathbb{Z}[x]/(x^d + 1)$, where d is a power-of-two. Let χ be a noise distribution over R. A module-LWE sample is obtained from sampling $\mathbf{s} \xleftarrow{\$} R_q^l$, $\mathbf{A} \xleftarrow{\$} R_q^{l\times k}$, $\mathbf{e} \leftarrow \chi^k$, and outputting $(\mathbf{A}, \mathbf{y}^t := \mathbf{s}^t\mathbf{A} + \mathbf{e}^t \mod q)$.*

We say $\mathsf{MLWE}_{d,l,k,q,\chi}$ is (t, ϵ)-secure if for any ppt algorithm $\mathcal{A}$ with running time t,

$$|\Pr[\mathcal{A}^{\mathcal{O}_{\mathsf{MLWE}}}(1^\kappa) = 1] - \Pr[\mathcal{A}^{\mathcal{O}_{\mathsf{uniform}}}(1^\kappa) = 1]| \leq \epsilon,$$

where oracle $\mathcal{O}_{\mathsf{MLWE}}$ outputs $\mathsf{MLWE}_{d,l,k,q,\chi}$ samples and oracle $\mathcal{O}_{\mathsf{uniform}}$ outputs uniform samples over $R_q^{l\times k} \times R_q^k$.

In context of MLWE, we often use n to denote the effective dimension of the MLWE secret, which is $n := d\ell$ and m to denote the effective sample size, which is $m := dk$.

Challenge Space and Noise Growth. For our signature scheme, we define as previous works [5,16,25], a challenge space $\mathcal{C} \subset R_q$. The elements in $\mathcal{C}$ consist of entries in $\{-1, 0, 1\}$ with respect to their coefficient embedding in R_q of Hamming weight at most w_c.

For an appropriate choice of R_q, all elements in $\mathcal{C}$ as well as their differences are invertible. This is the case for e.g. $R_q := \mathbb{Z}_q[X]/(x^n + 1)$, where n is a power of two and q is a prime congruent 5 mod 8 [30].

Let $\rho_s : \mathbb{R}^n \to \mathbb{R}$ be the Gaussian function defined by

$$\rho_s(\mathbf{x}) := \exp\left(-\pi \frac{\|\mathbf{x}\|^2}{s^2}\right)$$

The discrete Gaussian distribution for Gaussian parameter and a lattice $\Lambda \subset \mathbb{R}^n$ is defined by

$$D_{\Lambda,s}(\mathbf{x}) \propto \begin{cases} \rho_s(\mathbf{x}) & \text{if } \mathbf{x} \in \Lambda \\ 0 & \text{otherwise} \end{cases}$$

In the most general case, we consider χ over R_q to be a spherical discrete Gaussian distribution for Gaussian parameter $s := \alpha q$, i.e. each entry of the coefficient embedding representation of a ring element sampled from χ is sampled independently from the discrete Gaussian distribution $D_{\mathbb{Z},\alpha q}$. For a discrete Gaussian with parameter αq over $\mathbb{Z}$ [31], it holds that

$$\Pr_{e \xleftarrow{\$} D_{\mathbb{Z},\alpha q}} [|e| \geq t] \leq 3 \exp\left(-\pi \frac{t^2}{(\alpha q)^2}\right).$$

Hence, for $\eta \geq \alpha q \log \kappa$, we have $\Pr[|e| \geq \eta] \leq \mathsf{negl}$. By the generalized Hoeffding bound, we obtain

$$\Pr_{c \xleftarrow{\$} \mathcal{C}, \mathbf{e} \xleftarrow{\$} [-\eta,\eta]^d} \left[|\sum_{i=1}^{d} c_i e_i| \geq t \right] \leq 2 \exp\left(-\frac{t^2}{2 w_c \eta^2}\right),$$

where w_c is the Hamming weight of c. Hence, for $b_e \geq \eta \sqrt{w_c} \log \kappa$, we have $\Pr\left[|\sum c_i e_i| \geq b_e\right] \leq \mathsf{negl}$ as well as $\Pr[\|c\mathbf{e}\|_\infty \geq b_e] \leq \mathsf{negl}$. Note that a trivial bound yields $|\sum c_i e_i| \leq \eta w_c$.

Rounding over R_q. For $p, q \in \mathbb{N}$, $q > p$, the rounding operation $\lfloor a \rceil_p : \mathbb{Z}_q \to \mathbb{Z}_p$ is defined by multiplying a by p/q and rounding the result to the nearest smaller integer, i.e. $\lfloor \frac{p}{q} a \rfloor$. In a ring R, we apply the rounding operation component wise with respect to the coefficient embedding.

Further, for a noise distribution χ and dimension m, we define the set $\mathcal{B} \subset R_q$ by all elements $h(x) \in R$ that have coefficients $(h_0, \ldots, h_{d-1})$ within the interval $[0, b_r]$ in $\mathbb{Z}$, where $b_r := \lfloor \frac{q}{p} \rfloor - 1$ bounds the error term caused by the rounding function $\lfloor \cdot \rceil_p$. Here, for any $a \in \mathbb{Z}_p$, the rounding function maps all values in the interval $[\lfloor \frac{aq}{p} \rfloor, b_r + \lfloor \frac{aq}{p} \rfloor]$ in $\mathbb{Z}_q$ to a.

2.2 Digital Signatures

The following presents syntax and security definition of a digital signature scheme.

Definition 3 (Digital Signatures). *A* digital signature *scheme for a messages space* M *is a triplet of ppt algorithms* $(\mathsf{KGen}, \mathsf{Sign}, \mathsf{Verify})$ *with the following syntax*

KGen**:** *Takes as input* 1^κ *and outputs a key pair* $(\mathsf{pk}, \mathsf{sk})$.
Sign**:** *Takes as input* sk*, a message* $\mathsf{m} \in \mathsf{M}$ *and outputs a signature* σ.
Verify**:** *Takes as input* pk*, a message* $\mathsf{m} \in \mathsf{M}$*, a signature* σ *and outputs* 1 *if* σ *is a valid signature under* pk *for message* m*. Otherwise, it outputs* 0.

For correctness, for any $\mathsf{m} \in \mathsf{M}$*, we require that* $\mathsf{Verify}(\mathsf{pk}, \mathsf{m}, \sigma) = 1$*, where* $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}(1^\kappa)$*,* $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, \mathsf{m})$.

Definition 4 (Existential Unforgeability under Chosen Message Attacks (UF-CMA) Security). *A signature scheme* SGN *is* (t, ϵ, q_S, q_H)-UF-CMA *secure (existentially unforgeable under chosen message attacks) if for all algorithms* $\mathcal{A}$ *running in time at most* t *and making at most* q_S *queries to the signing oracle and* q_H *queries to the random oracle,*

$$\Pr\left[\begin{array}{l|l}\mathsf{Verify}(\mathsf{pk}, \mathsf{m}^*, \sigma^*) = 1 & (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{KGen}(1^\kappa) \\ \wedge\ \mathsf{m}^* \notin \{\mathsf{m}_i \mid i \in [q_S]\} & (\mathsf{m}^*, \sigma^*) \leftarrow \mathcal{A}^{\mathcal{O}_H, \mathsf{Sign}(\mathsf{sk}, \cdot)}(\mathsf{pk})\end{array}\right] \leq \epsilon,$$

where for $i \in [q_S]$*, on the* i*-th query* m_i *the signing oracle* $\mathsf{Sign}(\mathsf{sk}, \cdot)$ *returns* $\sigma_i \leftarrow \mathsf{Sign}(\mathsf{sk}, \mathsf{m}_i)$ *to* $\mathcal{A}$ *and* $\mathcal{O}_H$ *denotes query access to a random oracle.*

3 Adaptive Bounded Distance Decoding

The security of our signature scheme requires the hardness of the adaptive bounded distance decoding problem, defined as follows.

Definition 5 (Bounded Distance Decoding (BDD)). *Let* $q, l, k \in \mathbb{N}$ *and ring* R*. Bounded distance decoding for a tolerance set* $\mathcal{B} \subset R_q$*, and dimension* l *is* (t, ϵ)*-hard if for any ppt algorithm* $\mathcal{A}$ *with running time* t

$$\Pr\left[\begin{array}{l|l}\exists \mathbf{e} \in \mathcal{B}^k \ s.t. & \mathbf{A} \xleftarrow{\$} R_q^{l \times k}, \mathbf{y} \xleftarrow{\$} R_q^k \\ \mathbf{y}^t - \mathbf{z}^t\mathbf{A} = \mathbf{e}^t & \mathbf{z} \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{y})\end{array}\right] \leq \epsilon.$$

Definition 6 (Adaptive Bounded Distance Decoding (ABDD)). *Let* $p, q, l, k \in \mathbb{N}$ *and ring* R*. Adaptive bounded distance decoding for dimension* l *and challenge set* $\mathcal{C} \subset R$ *is* (t, ϵ)*-hard if for any ppt algorithm* $\mathcal{A}$ *with running time* t

$$\Pr\left[\mathbf{w}^t = \lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p \left| \begin{array}{l} \mathbf{A} \xleftarrow{\$} R_q^{l \times k}, \mathbf{y} \xleftarrow{\$} R_q^k \\ (\mathbf{w}, \mathsf{st}) \leftarrow \mathcal{A}(\mathbf{A}, \mathbf{y}) \\ c \xleftarrow{\$} \mathcal{C} \\ \mathbf{z} \leftarrow \mathcal{A}(c, \mathsf{st}) \end{array} \right.\right] \leq \epsilon,$$

where st *is a state of algorithm* $\mathcal{A}$*.*

For some parameter choices, the BDD and ABDD problems are statistically hard, meaning that with overwhelming probability, the problems are hard even for computationally unbounded adversaries. For some other parameter choices, the BDD and ABDD problems are conjectured to be computationally hard.

In both the computational and the statistical parameter regimes, we can reduce bounded distance decoding by a simple rewinding argument to adaptive bounded distance decoding.

Lemma 1. *Let there be a classical, i.e. non-quantum, algorithm* $\mathcal{A}$ *that* (t, ϵ)*-breaks* ABDD *for parameters* p, q, l, k*, ring* R *and challenge set* $\mathcal{C}$*, where all differences of two elements are invertible. Then there is an algorithm* $\mathcal{A}'$ *that* (t', ϵ')*-breaks* BDD *for parameters* q, l, k, d *and tolerance set* $\mathcal{B} \subseteq R_q$ *such that for all* $e' \in \mathcal{B}'$*,* $\psi(e') \in [-b_r, b_r]^{d \times d}$*,* $b_r := \lfloor \frac{q}{p} \rfloor - 1$*, where* $\epsilon' \geq \epsilon^2 - 1/|\mathcal{C}|$*,* $t' \leq 2t$*.*

Proof. Algorithm $\mathcal{A}'$ receives a BDD challenge $(\mathbf{A}, \mathbf{y})$ and needs to find a $\mathbf{z}$ such that $\mathbf{y}^t - \mathbf{z}^t\mathbf{A} = \mathbf{e}^t$ with $\mathbf{e} \in \mathcal{B}^k$. $\mathcal{A}'$ samples two challenges $c_1, c_2 \xleftarrow{\$} \mathcal{C}$. Unless $c_1 = c_2$, which occurs with a probability at most of $1/|\mathcal{C}|$, $(c_1 - c_2)$ is invertible and $\mathcal{A}'$ defines $\tilde{\mathbf{y}} := (c_1 - c_2)^{-1}\mathbf{y}$.

$\mathcal{A}'$ invokes $\mathcal{A}$ on input $\mathbf{A}, \tilde{\mathbf{y}}$ to receive a response $\mathbf{w}$. $\mathcal{A}'$ also stores $\mathcal{A}$'s current state st such that $\mathcal{A}$' can run $\mathcal{A}$ twice on the same state (rewinding). $\mathcal{A}'$ runs $\mathcal{A}$ with state st on input c_1 to receive solution $\mathbf{z}_1$ and on state st and input c_2 to receive $\mathbf{z}_2$. $\mathcal{A}'$ outputs $\mathbf{z} := \mathbf{z}_1 - \mathbf{z}_2$.

According to Jensen's inequality $\mathbf{z}_1$ and $\mathbf{z}_2$ will be correct with a probability of at least ϵ^2. Correct $\mathbf{z}_1$ and $\mathbf{z}_2$ have the form

$$\mathbf{w}^t = \mathbf{z}_1^t\mathbf{A} - c_1\mathbf{y}^t + \mathbf{e}_1^t = \mathbf{z}_2^t\mathbf{A} - c_2\mathbf{y}^t + \mathbf{e}_2^t,$$

where $\mathbf{e}_1, \mathbf{e}_2$ have coefficients in interval $[0, b_r]$ in $\mathbb{Z}$. In that case,

$$(c_1 - c_2)\tilde{\mathbf{y}}^t - (\mathbf{z}_1^t - \mathbf{z}_2^t)\mathbf{A} = \mathbf{e}_1^t - \mathbf{e}_2^t.$$

By the definition of $\mathbf{z} = \mathbf{z}_1 - \mathbf{z}_2$ and $\mathbf{y} = (c_1 - c_2)\tilde{\mathbf{y}}$, $\mathbf{y}^t - \mathbf{z}^t\mathbf{A} = \mathbf{e}_1^t - \mathbf{e}_2^t$. Therefore, $\mathbf{e} := \mathbf{e}_1 - \mathbf{e}_2 \in \mathcal{B}^k$. □

To make the ABDD problem statistically hard, we first provide a parameter setting under which the BDD problem under the reduction in Lemma 1 is statistically hard. We give a bound with simple proof, but this bound is very loose in terms of the minimum size of dimension k. We see this more as a proof of concept and use a heuristic bound that directly bounds the hardness of ABDD when determining our choices of parameters.

Lemma 2. *Let $\mathcal{B}'$ be a subset of R_q such that for all $e' \in \mathcal{B}'$, $\psi(e') \in [-b_r, b_r]^{d\times d}$. Let $\mathcal{C}$ be a subset of R_q such that all the $c \in \mathcal{C}$ are invertible in R_q and $\psi(c) \in [-2, 2]^{d\times d}$. If $5^d \cdot (2b_r + 1)^{dk} \cdot q^{d(l-k)}$ is a negligible function, then the following statement holds.*

$$\Pr_{\mathbf{A}\xleftarrow{\$}R_q^{l\times k}, \mathbf{y}\xleftarrow{\$}R_q^k}\left[\ \exists\ \mathbf{z} \in R_q^l, \mathbf{e}' \in {\mathcal{B}'}^k, c' \in \mathcal{C}'\ \textit{s.t.}\ \mathbf{A}^t\mathbf{z} - c'\mathbf{y} = \mathbf{e}' \pmod q\right] \leq \mathsf{negl}(\kappa)$$

Proof. Since c is invertible in R_q, rewrite the equation as $\mathbf{y} = (\mathbf{A}^t\mathbf{z} - \mathbf{e}') \cdot c^{-1} \pmod q$. Then for all $\mathbf{A} \in R^{l\times k}$, the number of possible vectors on the right side of the equation is at most $N := 5^d \cdot (2b_r + 1)^{dk} \cdot q^{d\ell}$. Therefore, if $N/q^{dk} = 5^d \cdot (2b_r + 1)^{dk} \cdot q^{d(l-k)} \in \mathsf{negl}(\kappa)$, the equation holds. □

4 Proposed Scheme

In our scheme, we choose a random oracle-like hash function $H : R_p^k \times R_q^{l\times k} \times R_q^k \times \mathsf{M} \to \mathcal{C}$, that hashes a rounded vector of ring elements, the public key and a message to our challenge set $\mathcal{C}$. Further, we define a set Good that is used to determine whether a signature is safe to publish, i.e. the acceptance condition, as well as whether it satisfies correctness. We choose b_e such that $\|c\mathbf{e}\|_\infty \leq b_e$ with overwhelming probability over the choice of $\mathbf{e} \leftarrow \chi^k$ and $c \xleftarrow{\$} \mathcal{C}$.

Definition 7 (Rejection Condition, Set Good). *Let h denote the coefficient embeddingdefined in Eqn. (1) . For parameters $p, q, d, k, b_e \in \mathbb{N}$, $q > p$ and ring $R := \mathbb{Z}[x]/(x^d + 1)$, we define the set $\mathsf{Good} \subset R_q^k$ as follows. $(h_i(x))_{i\in[k]} \in R_q^k$, if for all $j \in [k], i \in [d]$,*

$$h_{j,i} \in \mathbb{Z}_q \setminus \bigcup_{\ell=-\lfloor p/2 \rfloor}^{\lfloor p/2 \rfloor} \left[-b_e + \left\lfloor \frac{\ell q}{p} \right\rfloor, b_e + \left\lfloor \frac{\ell q}{p} \right\rfloor \right].$$

In Figure 1, we depict our scheme. The key generation algorithm samples the secret $\mathbf{s}$ from a uniform distribution and the error $\mathbf{e}$ from the noise distribution χ. To sign a message $\mathsf{m} \in \mathsf{M}$, we first sample a one-time masking term $\mathbf{r}$ from uniform from the same domain as the secret key, and compute $c := H(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p, \mathsf{pk}, \mathsf{m})$. The process is restarted if $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \notin \mathsf{Good}$. For an eligible $\mathbf{r}$, i.e., $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$, the signer computes $\mathbf{z} := \mathbf{r} + c\mathbf{s} \pmod q$ and outputs the signature as $\sigma := (\mathbf{z}, c)$. The verification of our scheme checks if $c = H(\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p, \mathsf{pk}, \mathsf{m})$.

The condition $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$ implies that $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$. Hence our scheme is correct. In fact, our signature scheme is perfectly correct, though during the signing process, there might be many rejections. Nevertheless, for a acceptance probability of ρ_r, the signing process is expected to terminate in time $t \propto 1/\rho_r$. More precisely, the signing process terminates with an overwhelming probability within κ/ρ_r rejections.

We give a more formal treatment of correctness in the following.

$(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KGen}(1^\kappa)$:	$\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, \mathsf{m})$:	$\{0, 1\} \leftarrow \mathsf{Verify}(\mathsf{pk}, \sigma, \mathsf{m})$:
$\mathbf{s} \xleftarrow{\$} R_q^l$, $\mathbf{e} \leftarrow \chi^k$,	Repeat until $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$	Parse $\sigma = (\mathbf{z}, c)$
$\mathbf{A} \xleftarrow{\$} R_q^{l\times k}$	$\quad \mathbf{r}^t \xleftarrow{\$} R_q^l$,	$\mathbf{w} := \lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p$
$\mathbf{y}^t := \mathbf{s}^t\mathbf{A} + \mathbf{e}^t \pmod q$	$\quad c := H(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p, \mathsf{pk}, \mathsf{m})$	if $c = H(\mathbf{w}, \mathsf{pk}, \mathsf{m})$
$\mathsf{sk} := \mathbf{s}$, $\mathsf{pk} := (\mathbf{A}, \mathbf{y})$	$\quad \mathbf{z} := \mathbf{r} + c\mathbf{s} \pmod q$	then Return 1
Return $(\mathsf{pk}, \mathsf{sk})$	Return $\sigma := (\mathbf{z}, c)$	else Return 0

Fig. 1. Proposed signature scheme

5 Correctness and Security Analysis

Lemma 3 (Correctness and Termination). *The signature scheme in Figure 1 is perfectly correct and has a heuristic acceptance rate of*

$$\rho_r := \left(\frac{b_r - 2b_e - 1}{b_r} \right)^{dk},$$

where $b_r := \lfloor \frac{q}{p} \rfloor - 1$.

Proof. Let $(\mathbf{z}, c)$ be the output of $\mathsf{Sign}(\mathsf{sk}, \mathsf{m})$ for $\mathsf{m} \in \mathsf{M}$ and $(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{KGen}(1^\kappa)$. By the acceptance condition, $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$ always holds. By the definition of set Good, the coefficient of each entry of $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t$ have a distance larger than b_e from the rounding borders (namely, $\left\lfloor \frac{\ell q}{p} \right\rfloor$ for $\ell = -\lfloor p/2 \rfloor, ..., \lfloor p/2 \rfloor$) caused by rounding function $\lfloor \cdot \rfloor_p$. Hence, $\mathbf{r}^t\mathbf{A}$ rounds to the same value as $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t$, i.e. $\lfloor \mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$.

By the definition of $\mathbf{z}$ and public key $(\mathbf{A}, \mathbf{y})$, $\mathbf{z}^t\mathbf{A} - c\mathbf{y}^t = \mathbf{r}^t\mathbf{A} - c\mathbf{e}^t$. Further, by the definition of c, $c = H(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p, \mathsf{pk}, \mathsf{m})$. Thus, the verification check $c = H(\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p, \mathsf{pk}, \mathsf{m})$ passes and Verify returns 1.

For the acceptance rate ρ_r, we need to compute the probability over $\mathbf{r} \xleftarrow{\$} R_q^\ell$ and $c \xleftarrow{\$} \mathcal{C}$ that $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$. With overwhelming probability, $\|c\mathbf{e}\|_\infty \leq b_e$. The probability that a random element u in $\mathbb{Z}_q$ falls in the bad region excluded in Good is

$$\Pr_{u \xleftarrow{\$} \mathbb{Z}_q} \left[u \in \bigcup_{\ell=-\lfloor p/2 \rfloor}^{\lfloor p/2 \rfloor} \left[-b_e + \left\lfloor \frac{\ell q}{p} \right\rfloor, b_e + \left\lfloor \frac{\ell q}{p} \right\rfloor \right] \right] \leq \frac{2b_e + 1}{b_r},$$

For the claimed heuristic bound in the lemma statement, we use the heuristic that the coefficients of $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t$ fall independently in the bad region. □

In Theorem 1 below, we prove that our signature scheme is EU-CMA in the ROM.

Theorem 1. *Let LWE be* $(t_{\mathsf{LWE}}, \epsilon_{\mathsf{LWE}})$*-hard,* ABDD *be* $(t_{\mathsf{ABDD}}, \epsilon_{\mathsf{ABDD}})$*-hard and* $H_\infty(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p \mid \mathbf{A}) \geq \xi$*. Then, the signature scheme in Figure 1 is* $(t_\mathcal{A}, \epsilon_\mathcal{A}, q_S, q_H)$*-*UF-CMA *secure in the programmable random oracle model, where* $t_\mathcal{A} \approx t_{\mathsf{LWE}} + t_{\mathsf{ABDD}}$ *and* $\epsilon_\mathcal{A} \leq \epsilon_{\mathsf{LWE}} + q_H\epsilon_{\mathsf{ABDD}} + q_S 2^{-\kappa} + \kappa^2\rho_r^{-2}q_S^2 2^{-\xi} + 2\kappa\rho_r^{-1}q_S q_H 2^{-\xi}$.

Proof. On a high level, we prove this theorem in two hybrids. In the first hybrid, we exploit the programmability of the random oracle to respond to signature queries without knowing the secret key. This step of faithfully simulating signatures without knowing the secret key crucially relies on the rejection sampling condition.

During the second hybrid, the public key of our signature scheme is replaced with uniform randomness. In this hybrid, there will be no secret key that allows to sign messages and, furthermore, it is infeasible for an adversary who cannot program the random oracle to forge signatures.

In the following, we define the two hybrids and show that: 1) by a statistical argument, simulated signatures are identically distributed as signatures created by the signing algorithm with access to the secret key, i.e. every algorithm has the same advantage in the UF-CMA game and hybrid 1; 2) there is no algorithm that has a different advantage in hybrid 1 and hybrid 2, unless it implicitly breaks the LWE assumption; 3) there is no algorithm that can forge a signature in hybrid 2, unless it implicitly breaks the ABDD assumption.

To summarize, this proves the theorem statement. The detailed description of the hybrids and the UF-CMA game are depicted in Figure 2.

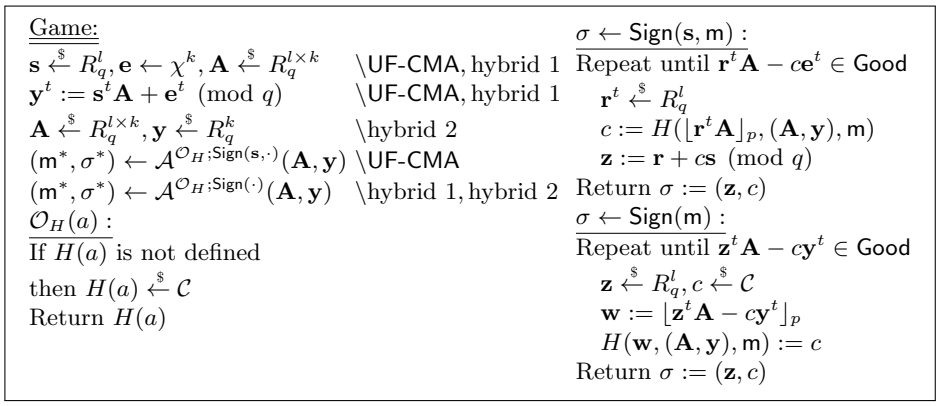

Fig. 2. UF-CMA security game and hybrids to prove Theorem 1

We start the formal argument with showing that any adversary that is successful in the UF-CMA game is also successful in hybrid 1.

Lemma 4. *Let there be an algorithm that (t, ϵ, q_S, q_H) breaks the* UF-CMA *security and $H_\infty(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p \mid \mathbf{A}) \geq \xi$. Then, there is also an algorithm that $(t', \epsilon', q'_S, q'_H)$ forges a signature in hybrid 1 for $t' \approx t$, $\epsilon' \geq \epsilon - q_S 2^{-\kappa} - \kappa^2 \rho_r^{-2} q_S^2 2^{-\xi} - 2\kappa\rho_r^{-1} q_S q_H 2^{-\xi}$, $q'_S = q_S$, and $q'_H = q_H$.*

Proof. The difference between the UF-CMA game and hybrid 1 is how signing queries are answered. In the UF-CMA game, one first samples $\mathbf{r} \xleftarrow{\$} R_q^l$, computes $c = H(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p, \mathsf{pk}, \mathsf{m})$, rejects if $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \notin \mathsf{Good}$ and then computes $\mathbf{z} = \mathbf{r} + c\mathbf{s}$. In hybrid 1, one samples first $\mathbf{z} \xleftarrow{\$} R_q^l$, $c \xleftarrow{\$} \mathcal{C}$, rejects if $\mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \notin \mathsf{Good}$ and finally programs the random oracle H on point $(\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p, (\mathbf{A}, \mathbf{y}), \mathsf{m})$ to be equal to c. In the following, we show that created signatures $(\mathbf{z}, c)$ have the same distribution.

As a first intermediate step, we want to show that the generated signatures $(\mathbf{z}, c)$ before the rejection have the same distribution in game UF-CMA and hybrid 1. This can only be the case if the reprogramming step of the oracle does not fail. Except with probability $q_S 2^{-\kappa}$, there are at most κ/ρ_r many reprogrammings per signature for all signature queries. The amount of defined points of the random oracle within hybrid 1 is upper bounded by $\kappa\rho_r^{-1} q_S + q_H$. At each reprogramming step, $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$ has at least min-entropy ξ given $\mathbf{A}$. Hence the probability that the random oracle is already defined for partial input $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$ is at most $(\kappa\rho_r^{-1} q_S + q_H) 2^{-\xi}$. Since there are at most $\kappa\rho_r^{-1} q_S$ reprogramming steps, the probability that reprogramming the random oracle fails in hybrid 1 is upper bounded by $q_S 2^{-\kappa} + \kappa\rho_r^{-1} q_S (\kappa\rho_r^{-1} q_S + q_H) 2^{-\xi}$. For the remaining parts of the proof, we assume that the reprogramming does not fail.

The challenge c in game UF-CMA is the output of the random oracle on input $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p, (\mathbf{A}, \mathbf{y}), \mathsf{m}$ and therefore uniformly distributed. In hybrid 1, c is sampled

uniformly at random and it is programmed to be the output of the oracle on input $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p, (\mathbf{A}, \mathbf{y}), \mathsf{m}$. Under the premise that $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$, c has the same distriubtion in game UF-CMA and hybrid 1.

We focus now on showing that $\mathbf{z}$ has the same distribution. In game UF-CMA, $\mathbf{z} := \mathbf{r} + c\mathbf{s}$, where $\mathbf{r} \xleftarrow{\$} R_q^l$. In hybrid 1, $\mathbf{z} \xleftarrow{\$} R_q^l$ and we can define $\mathbf{r} := \mathbf{z} - c\mathbf{s}$. Therefore in hybrid 1, $\mathbf{r}$ is also uniform and $\mathbf{z}$ is determined by $\mathbf{r}$, $\mathbf{s}$ and $\mathbf{c}$ as in UF-CMA.

It is left to show that the premise $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$ is implied by the rejection condition and that the rejection condition does not introduce any difference between the signature disitribution in game UF-CMA and hybrid 1. The latter is easy to show. $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$ is identical with $\mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \in \mathsf{Good}$, because $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t = \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t$. Obviously, we could replace the rejection condition in the orignial scheme with the publicly verifiable condition $\mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \in \mathsf{Good}$ that we need in hybrid 1. The only reason against it is a slight performance gain due to the fact that $\mathbf{r}^t\mathbf{A}$ has already been computed when evaluating the random oracle.

By the same argument as used for correctness (see Lemma 3), $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$ implies $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$. Therefore all signatures obtained by the adversary, i.e. that pass the rejection condition, have the same distribution in hybrid 1 and game UF-CMA.

All other signatures, i.e. the once that trigger the rejection condition, remain hidden and an adversary could at most observe a reprogrammed challenge. This might be a problem, because there could be a slight bias in the random oracle since the partial input $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p$ might be biased with the output c (which disappears for not rejected signatures where $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$). But in order to detect this bias, he would need to guess $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p$ which has at least min-entropy ξ for the same reason why $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$ has at least min-entropy ξ. The ability of an adversary to detect this bias is upper bounded by $\kappa\rho_r^{-1}q_S q_H 2^{-\xi}$. □

Lemma 5. *Let there be an algorithm that (t, ϵ, q_S, q_H) forges a signature in hybrid 1 and let LWE be $(t_{\mathsf{LWE}}, \epsilon_{\mathsf{LWE}})$-secure. Then, there is also an algorithm that $(t', \epsilon', q'_S, q'_H)$ forges a signature in hybrid 2 for $t' \approx t + t_{\mathsf{LWE}}$, $\epsilon' \geq \epsilon - \epsilon_{\mathsf{LWE}}$, $q'_S = q_S$, and $q'_H = q_H$.*

Proof. The lemma follows from a straightforward reduction to LWE. The only difference between hybrid 1 and hybrid 2 is the distribution of the public key $(\mathbf{A}, \mathbf{y})$. In hybrid 1, it is LWE distributed, while uniform in hybrid 2. If there is an algorithm that ϵ forges in hybrid 1 and ϵ' forges in hybrid 2, then LWE can be told apart from uniform with advantage $|\epsilon - \epsilon'|$, i.e. $\epsilon_{\mathsf{LWE}} \geq |\epsilon - \epsilon'|$. □

Lemma 6. *Let there be an algorithm that (t, ϵ, q_S, q_H) forges a signature in hybrid 2. Then, there is also an algorithm that $(t_{\mathsf{ABDD}}, \epsilon_{\mathsf{ABDD}})$ solves ABDD for $t_{\mathsf{ABDD}} \approx t$, $\epsilon_{\mathsf{ABDD}} \geq \frac{1}{q_H}\epsilon$.*

Proof. We prove the lemma by embedding an ABDD challenge in hybrid 2 such that if an algorithm forges successfully, it solves the ABDD problem. This is

straight forward. We use the ABDD challenge $(\mathbf{A}, \mathbf{y})$ as a public key in hybrid 2. We guess a random oracle query for point $(\mathbf{w}, (\mathbf{A}, \mathbf{y}), \mathsf{m}^*)$ to request a challenge c for query $\mathbf{w}^* = \mathbf{w}$ from the ABDD challenger. We program the random oracle by setting $H(\mathbf{w}, (\mathbf{A}, \mathbf{y}), \mathsf{m}^*) = c$. With a probability of $\frac{1}{q_H}$, the forgery will be for this c and message m^* thereby a valid signature $(\mathbf{z}, c)$ contains a valid ABDD solution $\mathbf{z}$. □

□

For applicability of Theorem 1, we need to show that $\xi \leq H_\infty(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p \mid \mathbf{A})$ is sufficiently large. Technically, it would be sufficient to show that it is hard for any efficient adversary to compute $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$, given $\mathbf{A}$. This would be sufficient, since it only needs to be hard for an efficient adversary to guess the points where the random oracle is going to be programmed during the simulation. Though, using computational intractability is not necessary.

Instead, we use a similar approach as used by Bai and Galbraith [5, Lemma 3], relying on the fact that the public key component $\mathbf{A}$ has at least one invertible ring element. Unlike [5], we do not need to rely on a Gaussian heuristic, since in our case $\mathbf{r}$ is chosen uniformly at random, which leads to a very simple analysis.

Lemma 7. *For any $\mathbf{A} \in R_q^{l \times k}$ with an invertible entry $a_{i,j} \in R_q$,*

$$H_\infty(\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p \mid \mathbf{A}) \geq d \log p,$$

where $\lfloor \mathbf{r}^t\mathbf{A} \rfloor_p \in \mathbb{Z}_p^m$.

Proof. Since $a_{i,j}$ is invertible,

$$H_\infty(r_i a_{i,j} \mid \mathbf{A}) = H_\infty(r_i) = d \log q.$$

The rounding function causes to lose $\log(q/p)$ entropy at each of the d coefficients of $r_i \in R_q$ with respect to the coefficient embedding. □

6 The Difficulty of Removing the Remaining Rejection Condition

In the signature scheme presented in Section 4, we have removed one rejection condition used in Dilithium [16] and qTESLA [4]. The other rejection condition which checks if $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$ is left to ensure that $\lfloor \mathbf{z}^t\mathbf{A} - c\mathbf{y}^t \rfloor_p = \lfloor \mathbf{r}^t\mathbf{A} \rfloor_p$, i.e., the rounding of $\mathbf{r}^t\mathbf{A} \pmod q$ is recoverable from $\mathbf{z}^t\mathbf{A} - c\mathbf{y}^t = \mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \pmod q$.

An interesting question is whether we can further remove this rejection condition while maintaining the efficiency of the scheme. By "maintaining the efficiency of the scheme", let us remind the readers that in theory, there are ways of achieving a Fiat-Shamir type of lattice signature without using rejection sampling, such as setting the modulus q to be super-polynomially large, or creating many independent commitments and challenges and allowing the third message in the ID protocol to answer only a few of them. However, such methods significantly increase the sizes of public-keys and signatures. Instead, we focus on methods that remove the rejection condition with a potentially low cost.

6.1 Impossibility of Extracting Consistent Values from Commitments with Errors

Suppose there are functions g, f that map $\mathbf{r}^t\mathbf{A}$ and $\mathbf{r}^t\mathbf{A}+\hat{\mathbf{e}}^t$ for any bounded error term $\hat{\mathbf{e}}$ to the same value, and make sure that $g(\mathbf{r}^t\mathbf{A})$ serves as a commitment of $\mathbf{r}$, or at least preserves high min-entropy; then we can apply the hash function on $g(\mathbf{r}^t\mathbf{A})$ instead of $\lfloor \mathbf{r}^t\mathbf{A} \rceil_p$.

However, when q is polynomial, no balanced functions g, f are able to guarantee that $g(\mathbf{r}^t\mathbf{A}) = f(\mathbf{r}^t\mathbf{A} + \hat{\mathbf{e}}^t)$ with probability $1 - \mathsf{negl}$. Here a boolean function is called balanced if it outputs 0 and 1 with almost the same probability over a random input from the domain. The result follows a recent result of Guo et al. [23], which shows a similar impossibility result for the lattice-based key exchange protocols [14,34].

The following corollary is implicit from [23, Theorem 1].

Corollary 1. *Let $m, q \geq 2$, χ be a symmetric distribution over $\mathbb{Z}_q$ such that for any $a \in \mathbb{Z}_q \backslash 0$, it holds that $\Pr_{x\leftarrow\chi}[ax = 0] \leq 9/10$, and $\Pr_{x\leftarrow\chi}[ax = q/2] \leq 9/10$. Consider the joint distribution of $(\mathbf{x}, \mathbf{y})$ where $\mathbf{x} \leftarrow U(\mathbb{Z}_q^m)$, $\mathbf{y} = \mathbf{x} + \mathbf{e}$ where $\mathbf{e} \leftarrow \chi^m$. Then, for any balanced function $f, g : \mathbb{Z}_q^m \rightarrow \{0, 1\}$, it holds that*

$$\Pr_{(\mathbf{x},\mathbf{y})}[f(\mathbf{x}) = g(\mathbf{y})] \leq 1 - \Omega(1/q^2)$$

In our application, $\mathbf{x} = \mathbf{r}^t\mathbf{A} \pmod q$. $\mathbf{x}$ is uniform over R_q since $\mathbf{A}$ is sampled uniformly random over R_q.

6.2 Evidence on the Difficulty of Adapting the Reconciliation Mechanism

We also tried to adapt the reconciliation mechanism used in lattice-based key-encapsulation mechanisms [14,34]. The reconciliation mechanisms can be adapted to our signature scheme to provide correctness when removing all rejection conditions. Nevertheless, we show that they would leak information about the error term $\mathbf{e}$ in the public key. Therefore, this attempt fails for security reason.

Let us first recall the reconciliation mechanisms [14,34] used in the lattice-based key-exchange. Abstractly, the reconciliation mechanism uses two functions $\mathsf{hint} : \mathbb{Z}_q \rightarrow \{0, 1\}$ and $g : \mathbb{Z}_q \times \{0, 1\} \rightarrow \{0, 1\}$ such that, given $v \in \mathbb{Z}_q$, $\mathsf{hint}(v)$ is the hint bit of v, the reconciliation function $g(v + e, \mathsf{hint}(v))$ is equal to $\lfloor v \rceil$ whenever e is bounded (this is used to ensure correctness). Furthermore, if v is uniformly random from $\mathbb{Z}_q$, then $\lfloor v \rceil$ is uniformly random given $\mathsf{hint}(v)$ (this is used to ensure security in the key-exchange scheme).

The simple adaption of the reconciliation mechanism and the attack. Let us apply the reconciliation mechanism on the string $\mathbf{r}^t\mathbf{A}$ generated in the signature scheme. Suppose the first message in the three round protocol contains $\lfloor \mathbf{r}^t\mathbf{A} \rceil_p$ and $\mathsf{hint}(\mathbf{r}^t\mathbf{A})$, then we hash on $\lfloor \mathbf{r}^t\mathbf{A} \rceil_p$ and $\mathsf{hint}(\mathbf{r}^t\mathbf{A})$, or $\lfloor \mathbf{r}^t\mathbf{A} \rceil_p$ only. In the

final signature we give out $\mathbf{z}$, c, and $\mathsf{hint}(\mathbf{r}^t\mathbf{A})$. This at least gives a signature scheme with correctness due to the correctness of the reconciliation mechanism.

Let us run a security analysis over the modified scheme. To prove unforgeability, we need to simulate $\lfloor \mathbf{r}^t\mathbf{A} \rceil_p$ and $\mathsf{hint}(\mathbf{r}^t\mathbf{A})$. However it seems that the simulator will only be able to get $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t$, from which it seems tricky to simulate $\lfloor \mathbf{r}^t\mathbf{A} \rceil_p$ and the hint $\mathsf{hint}(\mathbf{r}^t\mathbf{A})$.

In fact, there is an explicit attack for the scheme. Suppose the adversary makes N signature queries. For the i^{th} query, the reconciliation mechanism guarantees to recover $\lfloor \mathbf{r}_i^t\mathbf{A} \rceil_p$ and $\mathbf{r}_i^t\mathbf{A} - c_i\mathbf{e}^t \pmod q$, for $i \in [N]$; it is then possible to recover the error term $\mathbf{e}$ in the public key when N is sufficiently large.

Here is the algorithm. From $\lfloor \mathbf{r}_i^t\mathbf{A} \rceil_p$ we can obtain

$$\mathbf{r}_i^t\mathbf{A} + \hat{\mathbf{e}}_i^t \pmod q \tag{2}$$

where $\hat{\mathbf{e}}_i$ denotes the error caused by rounding; it is bounded and independent of the error term $\mathbf{e}$ in the public key. Subtracting Eqn (2) by $\mathbf{r}_i^t\mathbf{A} - c_i\mathbf{e}^t \pmod q$ gives $\hat{\mathbf{e}}_i + c_i\mathbf{e}$ over the base ring R without mod q. We can view c_i, $\mathbf{e} \cdot c_i + \hat{\mathbf{e}}_i$ as "LWE samples without mod q", where the error vector $\mathbf{e}$ in the public key is treated as the secret term. The "LWE without mod q" problem can be solved in polynomial time given polynomially many samples when the variance of the error is polynomially bounded (an explicit analysis is given in [6]).

The same attack applies if an additional error term $\tilde{\mathbf{e}}$ independent of $\mathbf{e}$ is injected in the commitment, namely let the commitment be $\lfloor \mathbf{r}^t\mathbf{A} + \tilde{\mathbf{e}}^t \rceil_p$ and the hint be $\mathsf{hint}(\mathbf{r}^t\mathbf{A} + \tilde{\mathbf{e}}^t)$.

Generalizing the statistical attack to handle dependent noise. Suppose we modify the first attempt by giving out $\lfloor \mathbf{r}^t\mathbf{A} + E(\mathbf{e}) \rceil_p$ and $\mathsf{hint}(\mathbf{r}^t\mathbf{A} + E(\mathbf{e}))$ in the first message, where $E : R_q^{1\times k} \to R^{1\times k}$ is a possibly randomized function. As long as E outputs small vectors, then the correctness of the scheme still holds.

Following the attack mentioned previously, we can recover $\mathbf{r}_i^t\mathbf{A} + E(\mathbf{e}) + \hat{\mathbf{e}}_i^t$ where $\hat{\mathbf{e}}_i$ denotes the error caused by rounding. From there we can get

$$c_i, \quad \mathbf{e} \cdot c_i + E(\mathbf{e}) + \hat{\mathbf{e}}_i \tag{3}$$

without mod q. But $E(\mathbf{e})$ is a possibly non-linear function on $\mathbf{e}$, so the equation above does not give samples in the form of "LWE without mod q".

However, we show that information about the error term in the public key can still be extracted by statistical analysis as long as sufficient amount of samples of the form in Eqn. (3) are available.

Definition 8 (Hoeffding Bound). *Let $X_1, \ldots, X_n$ be independent random variables bounded in interval $[a_i, b_i]$, then for any positive $t \in \mathbb{R}$,*

$$\Pr[|\frac{1}{n}\sum X_i - \mathbb{E}[\frac{1}{n}\sum X_i]| \geq t] \leq 2e^{-\frac{2n^2t^2}{\sum_i (b_i - a_i)}}.$$

For clarity, we present the theorem in the form where the signature scheme is instantiated over the base ring $\mathbb{Z}_q$, or we can think of it as the coeffcient

embedding of the scheme over R_q. Notice that the attack runs polynomial in noise and rounding bound B. Hence if B is superpolynomial, the attack will not be efficient. But this would require a superpolynomial modulus q.

Theorem 2. *Let* SGN *be a signature scheme with public key* $\mathbf{A} \in \mathbb{Z}_q^{n\times m}, \mathbf{Y} = \mathbf{SA}+\mathbf{E} \in \mathbb{Z}_q^{n\times m}$*, based on the Fiat-Shamir transform of an identification scheme where first message contains* $\mathbf{r}^t\mathbf{A} + E(\mathbf{A}, \mathbf{r}, \mathbf{S}, \mathbf{E}) \in \mathbb{Z}_q^{1\times m}$*, second message contains* $\mathbf{c} \in \{0,1\}^n$ *and third message contains* $\mathbf{z} = \mathbf{r} + \mathbf{cS} \in \mathbb{Z}_q^{1\times n}$*, where* E *is an arbitrary randomized function with range* $[-B, B]^{1\times m}$*,* $2B \leq q$*. Then for any reconciliation function and hint that allows to recover* $\mathbf{r}^t\mathbf{A} + E(\mathbf{A}, \mathbf{r}, \mathbf{S}, \mathbf{E})$*, there is a key recovery attack against* SGN *given* $3B\kappa$ *signature queries and time polynomial in* κ *and* B*.*

Proof. We describe an successful attack algorithm $\mathcal{A}$ as follows. The fact $\mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{Y} = \mathbf{r}^t\mathbf{A} - \mathbf{c}^t\mathbf{E}$ allows $\mathcal{A}$ to compute

$$\begin{aligned} F(\mathbf{r}, \mathbf{c})_{\mathbf{A},\mathbf{S},\mathbf{E}} &:= \mathbf{r}^t\mathbf{A} + E(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) - (\mathbf{z}^t\mathbf{A} - \mathbf{c}^t\mathbf{Y}) \\ &= E(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) + \mathbf{c}^t\mathbf{E} = \begin{pmatrix} E_1(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) \\ \vdots \\ E_m(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) \end{pmatrix}^t + \mathbf{c}^t\mathbf{E} \end{aligned}$$

$\mathcal{A}$ uses an estimator to estimate the mean of several random variables. To get a good estimator for the mean μ of a random variable $X \in [-B, B]$, it computes $\mu' := \frac{1}{B\kappa}\sum x_i$, where for $i \in [B\kappa]$, $x_i \leftarrow X$. By the Hoeffding bound,

$$\Pr[|\mu' - \mu| \geq \frac{1}{3}] \leq 2e^{-\frac{2}{18}\kappa}.$$

$\mathcal{A}$ picks $B\kappa$ signatures for which $c_1 = 0$ and computes an estimate of the mean $\mu_{E_1,\mathbf{c}\setminus c_1}$ of the first entry of $F(\mathbf{r}, \mathbf{c})_{\mathbf{A},\mathbf{S},\mathbf{E}}$.

$$E_1(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) + \sum_i c_i\mathbf{E}_{1,i} = E_1(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) + \sum_{i\neq 1} c_i\mathbf{E}_{1,i}.$$

Here, we think of $\mathbf{r}, \mathbf{c}$ as the source of independence in each signature sample.

It then computes an estimation of the mean $\mu_{E_1,\mathbf{c},c_1}$ of $E_1(\mathbf{A}, \mathbf{S}, \mathbf{E}, \mathbf{r}) + \sum_i c_i\mathbf{E}_i$ for $B\kappa$ samples with $c_1 = 1$. With overwhelming probability the estimate matches the actual mean and thus $\mathcal{A}$ recovers $\mathbf{E}_{1,1}$ correctly by computing $\mathbf{E}_{1,1} = \lfloor\mu_{E_1,\mathbf{c},c_1}\rceil - \lfloor\mu_{E_1,\mathbf{c}\setminus c_1}\rceil$. It repeats this to recover the first row of $\mathbf{E}$ which allows to recover the first row of $\mathbf{S}$ as well. Repeating this for each entry allows to recover the whole secret $\mathbf{S}$. By Chernoff bound, for each $b \in \{0,1\}$ at least $B\kappa$ out of $3B\kappa$ many random signatures will correspond to challenge $\mathbf{c}_1 = b$ except negligible probability. By a union bound over all n entries of $\mathbf{c}$, this will hold for all entries of $\mathbf{c}$. Therefore, with overwhelming probability, $3B\kappa$ random signature queries are sufficient for the attack. □

6.3 Discussions of the Possible Generalizations and Limitations

Let us conclude this section with a summary of the possible generalizations and limitations of our negative results.

For the first negative result, we are not able to rule out the possibility that the function g depends on the public matrix $\mathbf{A}$. In [23], the authors are able to rule out such a possibility for the lattice-based key exchange. The setting in the signature scheme seems to be different. In fact, given that in our setting, using reconciliation mechanisms with hints has already provide a scheme with correctness, we do not attempt to rule out the possibility of achieving correctness.

For the second negative result, we are not able to rule out the possibility of using the reconciliation mechanism when the underlying commitment, on a string $\mathbf{r}$, is not of the form of $\mathbf{r}^t\mathbf{A}$. However, changing the structure of the commitment in the first round of the ID protocol seems to require a significant change in the lattice-based commitment protocol.

Let us also remark that it is impossible to rule out a lattice-based signature scheme with polynomial modulus without rejection sampling, given the presence of the signature scheme based on lattice trapdoor [19]. But of course, this would require significant changes to our protocol or similar rejection sampling based protocols like [27,28,22,15,5,16,25,4].

7 Parameter and Comparison

We provide concrete parameters for the signatures in two settings. In the first setting, the parameters are set so that the adaptive bounded distance decoding problem (ABDD, cf. Definition 6) is hard even for a computationally unbounded adversary. In the second setting, the parameters are set so that the ABDD problem is computationally hard. Compared to the first setting, the second setting gives schemes with smaller public keys and less number of repetitions, while relying on one more computational assumptions.

It is reasonable to compare our first setting with Dilithium-QROM [25] and the "provable" version of qTESLA [4], since all of them set the scheme in the "lossy" mode, so that the schemes have tight security proofs from (Ring or Module)LWE in the QROM [25]; then compare our second setting with Dilithium [16], which also requires additional computational assumptions in addition to Module LWE. We do not compare with the "heuristic" version of qTESLA [4] since there are bugs in those parameter estimations, and the qTESLA team decided to drop the "heuristic" parameters in the second round of NIST PQC standardization.

Let us recall the notations. Let $R := \mathbb{Z}[x]/(x^d + 1)$, where d is a power-of-two. Let q be the bigger modulus, p be the smaller modulus after rounding. Let $b_r := \lfloor \frac{q}{p} \rfloor - 1$. When setting the concrete parameters, we assume $b_r + 1$ is a power of 2. Then rounding a number $a \in \mathbb{Z}_q$ to $\mathbb{Z}_p$ is effectively done by dropping the $\log(b_r+1)$ least significant bits when $a \in [0, \lfloor q/2 - 1 \rfloor]$, or dropping the $\log(b_r+1)$ least significant bits then subtracting by 1 when $a \in [-\lfloor q/2 \rfloor, 0)$. Let $\mathbf{A} \in R_q^{l \times k}$. Let $n = d \cdot l$, $m = d \cdot k$. Let $\mathcal{C} \subseteq R$ denote the space of the challenge c.

7.1 Parameters with Statistical Hardness of ABDD

We emphasize that under statistical hardness of ABDD, the underlying identification protocol of our signature scheme is statistically sound and honest verifier zero-knowledge. By a result of Unruh, this implies security of our signature scheme in the QROM.

Theorem 3 ([37], Theorem 1). *Assume that an identification protocol has completeness, unpredictable commitments, honest-verifier zero-knowledge and statistical soundness. Then the Fiat-Shamir transformed protocol has completeness, zero-knowledge and weak simulation soundness in the QROM.*

Alternatively, we can also rely on [25, Theorem 3.2, 3.3 and 3.4] for the same tightness bounds obtained for Dilithium in the QROM setting.[4]

To determine the concrete parameter setting where the statistical hardness of ABDD holds, we assume the Gaussian heuristic in the similar fashion of e.g. [5]. This heuristic implies the statistical hardness of ABDD for significantly smaller parameters than Lemma 2.

Theorem 4. *Let $\mathcal{B}$ be the set of all elements $h(x) \in R$ that have coefficients $(h_1, \ldots, h_d)$ within the interval $[0, b_r]$ in $\mathbb{Z}$. If $q^{n-m} \cdot |\mathcal{B}|^m$ is a negligible function, then, under Gaussian heuristic, for all but negligibly many $\mathbf{A} \in R_q^{l \times k}$ and $\mathbf{y} \in R_q^k$, the following statement holds. For all $\mathbf{v} \in R_q^k$,*

$$\Pr_{c \in \mathcal{C}} \left[\exists\ \mathbf{z} \in R_q^l, \mathbf{e}' \in \mathcal{B}^k \text{ such that } \mathbf{A}^t\mathbf{z} - c\mathbf{y} = \mathbf{v} + \mathbf{e}' \pmod q \right] \leq \mathsf{negl}(\kappa)$$

Proof. First we note that m has to be bigger than n (i.e. $k > l$). Otherwise, suppose $k = l$ and $\mathbf{A}$ is invertible, then $\mathbf{z}$ exists for any $c \in \mathcal{C}$.

To understand how small m can be, consider the following lattice Λ:

$$\Lambda := \left\{ \mathbf{t} \in \mathbb{Z}^{n+h+m} \mid [\hat{\mathbf{A}}^t \mid \hat{\mathbf{Y}}^t \mid \mathbf{I}_m] \cdot \mathbf{t} = \mathbf{0} \pmod q \right\},$$

where $\hat{\mathbf{A}} \in \mathbb{Z}^{n \times m}$, $\hat{\mathbf{Y}} \in \mathbb{Z}^{d \times m}$ denote the matrices obtained from taking the coefficient embedding of $\mathbf{A}$ and $\mathbf{y}$.

The determinant of Λ is q^m. Define a set $\mathcal{S} = [-q/2, q/2]^n \times [-1, 1]^d \times [0, b_r]^m$, which is a convex subset of $\mathbb{R}^{n+d+m}$. If we assume Gaussian heuristic, then for any $\mathbf{v} \in \mathbb{Z}^m$, fix a vector $\mathbf{t} \in \mathbb{Z}^{n+h+m}$ such that $[\hat{\mathbf{A}}^t \mid \hat{\mathbf{Y}}^t \mid \mathbf{I}_m] \cdot \mathbf{t} = \mathbf{v} \pmod q$, the number of elements in $\Lambda + \mathbf{t} \cap \mathcal{S}$ is expected to be $N := q^n \cdot |\mathcal{C}| \cdot |\mathcal{B}|^m / q^m$.

Therefore, if $N/|\mathcal{C}|$ is negligible, i.e. $q^n \cdot |\mathcal{B}|^m / q^m$ is negligible, then the probability statement in the theorem holds. □

[4] Theorem 3.2, 3.3 require a non-standard version for honest-verifier zero-knowledge that is called no-abort HVZK. A simulator for naHVZK does not output transcripts that lead to an abort. On the downside the challenge of accepting transcripts must be uniform. We observe that both theorems hold in case of HZVK simulators that do output commitment and challenge of aborting transcripts but do not necessarily have uniform challenges in accepting transcripts. In the security proofs, the QROM just needs to be also "programmed" for challenges that lead to a rejection.

Table 1. Parameters of our scheme assuming the statistical hardness of ABDD

Parameters	1	2	3	4	5	6
$\log_2 q$	23	45	30	28	27	31
d	256	512	1024	1024	1024	512
l	3	4	1	1	1	3
k	14	8	5	4	4	9
$n = d \cdot l$	768	2048	1024	1024	1024	1536
$m = d \cdot k$	3584	4096	5120	4096	4096	4608
η $(= \|\mathbf{e}\|_\infty)$	6	7	20	10	6	3
w = the weight of c	60	46	36	36	36	46
$\log_2(b_r)$	18	20	24	21	20	21
b_e	360	322	720	360	216	138
Expected repetitions	19097	12.4	1.55	4.08	5.41	1.83
LWE security	122.8	165.1	139.4	140.2	138.1	170.0
Public key size (bytes)	10336	23072	19232	14368	13856	17888
Signature size (bytes)	2247	11589	3972.5	3716.5	3588.5	6021.8

Now we can determine the concrete parameters according to the bound given in Theorem 4. The parameters are determined in the following order: we first pick q, d, l, b_r, then choose m so that

$$m \geq \frac{\kappa + n \log q}{\log q - \log_2 b_r}$$

The probability that $\mathbf{r}^t\mathbf{A} - c\mathbf{e}^t \in \mathsf{Good}$ is lower bounded by

$$\left(\frac{b_r - 2b_e - 1}{b_r}\right)^m \tag{4}$$

where $b_e \geq \|c\mathbf{e}\|_\infty$.

The public key is composed of the 256-bit seed to generate $\mathbf{A}$, and $\mathbf{y}$ which is of size $m \log_2 q$. The signature is composed of a ring vector of size $n \log_2 q$ plus $d + w$ bits to store the challenge c, where w is the weight of c. Currently, we do not adopt the public-key size optimization technique from [16], i.e., by dropping bits from the $\mathbf{y}$ term in the public key and adding hint bits in the signatures. It is reasonable to expect that over fixed q, d, l, k, b_e and b_r values, the optimization leads to smaller public keys, slightly bigger signatures, and slightly less security (since the ABDD problem would be slightly easier to solve given a vector $\mathbf{y}$ without the least significant bits).

In Table 1, we provide 6 sets of concrete parameters. Parameter sets 1 and 2 follow the q, d, l, η, w, b_e values chosen in [16] and [25], then derive m, expected repetition, LWE hardness, and the rest of the parameters. The parameter sets 3 to 6 choose the q, d, l, η, w, b_e values that are more suitable for our scheme.

In Table 2, we compare the parameters of our scheme with the parameters of Dilithium-QROM and qTESLA-provable. We only list the classical security estimation since under the commonly used LWE security estimation

Table 2. Comparison with Dilithium-QROM and qTESLA-provable.

Parameters	Classical security	PK size	Sign size	Exp. repetitions
Dilithium-QROM standard	140	7712	5696	4.3
qTESLA-p standard	140	14880	2592	3.45*
Ours standard-I	138.1	13856	3588.5	5.41
Ours standard-II	140.2	14368	3716.5	4.08
Ours standard-III	139.4	19232	3972.5	1.55
Dilithium-QROM high	175	9632	7098	2.2
qTESLA-p high	279	38432	5664	3.84*
Ours high	170.0	17888	6021.8	1.83

$*$: The expected repetition numbers reported in [4] are obtained by experiments, which are smaller than the estimated numbers derived from Formula (4).

model [3,2], schemes with 140-bit classical security are expected to have 128-bit quantum security (similarly for the higher security level). Compared to Dilithium-QROM, we achieve smaller signatures, similar rejection rates, but bigger public keys, since they use an extra public-key optimization technique. Compared to qTESLA-provable, under similar rejection rates, our public key sizes are smaller, but the signature sizes are larger.

7.2 Parameters with Computational Hardness of ABDD

We can further reduce parameter m, the width of the public key, so as to reduce the public key size and the number of rejections. The saving comes with a cost of additionally assuming the computational hardness of the ABDD problem. In the rest of the section, we first explain our model of estimating the cost of solving the ABDD problem using the BKZ algorithm, then provide the parameters of our signature scheme.

To explain our estimation model of solving BDD in general, we assume all the matrices are defined over $\mathbb{Z}$ in this paragraph. The best way we know to solve ABDD is to treat the entire $\mathbf{w}^* + c\mathbf{y}$ as the target vector for the standard bounded distance decoding problem (cf. Definition 5), and then use BKZ to solve the BDD problem.

Let $\mathbf{A} \in \mathbb{Z}_q^{n\times m}$, $\mathbf{t} = \mathbf{A}^t\mathbf{z} + \mathbf{e} \pmod q$ be such a BDD instance. To express the basis of the integer lattice used in the attack, we write the BDD instance in its Hermite normal form. Let $\bar{\mathbf{A}}$ be the first $n \times n$ part of $\mathbf{A}$, assuming it is invertible. Then

$$\mathbf{t}^t = \mathbf{z}^t\mathbf{A}+\mathbf{e}^t+q\mathbf{k}^t = \mathbf{z}^t[\bar{\mathbf{A}}, \mathbf{A}']+\mathbf{e}^t+q[\bar{\mathbf{k}}^t \mid \mathbf{k}'^t] = (\mathbf{z}^t\bar{\mathbf{A}})[\mathbf{I} \mid \bar{\mathbf{A}}^{-1}\cdot\mathbf{A}']+\mathbf{e}^t+q[\bar{\mathbf{k}}^t \mid \mathbf{k}'^t],$$

where $\mathbf{k} \in \mathbb{Z}^m$, $\bar{\mathbf{k}}$, $\mathbf{k}'$ are the top n and bottom $m-n$ entires of $\mathbf{k}$.

Let $\mathbf{A}'' := \bar{\mathbf{A}}^{-1} \cdot \mathbf{A}'$. Let

$$\mathbf{B} := \begin{pmatrix} \mathbf{I}_n & \mathbf{0} \\ \mathbf{A}''^t & q\mathbf{I}_{m-n} \end{pmatrix}.$$

Table 3. Parameters of our scheme assuming the computational hardness of ABDD.

Parameters	1	2	3	4	5	6
$\log_2 q$	23	45	30	28	27	31
d	256	512	1024	1024	1024	512
l	3	4	1	1	1	3
k	6	5	2	2	2	5
$n = d \cdot l$	768	2048	1024	1024	1024	1536
$m = d \cdot k$	1536	2560	2048	2048	2048	2560
η $(= \|\mathbf{e}\|_\infty)$	6	7	20	10	6	3
$w =$ the weight of c	60	46	36	36	36	46
$\log_2(b_r)$	18	20	24	21	20	20
b_e	360	322	720	360	216	138
Expected repetitions	68.3	4.82	1.19	2.02	2.33	1.96
BKZ approx factor δ	1.0053	1.0045	1.0049	1.0042	1.0041	1.0035
BKZ block-size β	258	332	290	370	388	492
ABDD hardness	103.2	126.8	115.1	138.4	142.6	173.1
LWE security	122.8	165.1	139.4	140.2	138.1	170.0
Public key size (bytes)	4448	14432	7712	7200	6944	9952
Signature size (bytes)	2247	11589	3972.5	3716.5	3588.5	6021.8

Then $\mathbf{B} \cdot \begin{pmatrix} \bar{\mathbf{A}}^t \mathbf{z} + q\bar{\mathbf{k}} \\ \mathbf{k}' \end{pmatrix} - \mathbf{t} = -\mathbf{e}$.

So the problem can be solved by running an approximate-CVP solver on a given basis $\mathbf{B}$ and target $\mathbf{t}$, or running an approximate-SVP solver on $\mathbf{B'} := \begin{pmatrix} \mathbf{B} & \mathbf{t} \\ \mathbf{0} & 1 \end{pmatrix}$.

One of the most common cost model for estimating the running time of BKZ for solving approximate-SVP is the following. Let h be the dimension of the lattice ($h(\mathbf{B}') = m + 1$). Let δ be the approximation factor. Following [3,2], we use sieving as the SVP oracle with time complexity $2^{0.292\beta+16.4}$ in the block size β. BKZ is expected to return a vector of length $\delta^h \det^{1/h}$ for a lattice of dimension h. Therefore, we found the smallest block size β in achieving the needed δ corresponding to the length of $\mathbf{e}$, which can be obtained from $\frac{\|\mathbf{e}\|_2}{\det^{1/h}} = \delta^h$.

Finally, we used the heuristic $\delta \approx (\frac{\beta}{2\pi e}(\pi\beta)^{1/\beta})^{\frac{1}{2(\beta-1)}}$ to determine the relation between β and δ, and we set the total time complexity of BKZ with a block-size of β and dimension h as $8h \cdot \text{time}(SVP) = 8h \cdot 2^{0.292\beta+16.4}$ [10,3].

In Table 3, we provide the concrete parameters. In each of the 6 sets of parameters, the values of q, d, l, η and w follow the same choices of from Table 1. We then choose smaller values for k, so that the computational hardnesses of the ABDD problem match the hardnesses of breaking the LWE instance in the public key.

In Table 4 we compare the parameters of our scheme with the parameters of Dilithium [16]. Compared to Dilithium, the sizes of signatures and public keys of our scheme are larger. The public key is inherently larger in our scheme since m

Table 4. Comparison with Dilithium.

Parameters	Classical security	PK size	Sign size	Exp. repetitions
Dilithium standard	138	1472	2701	6.6
Ours standard-I	138.4	7200	3716.5	2.02
Ours standard-II	138.1	6944	3972.5	2.33
Dilithium high	174	1760	3366	4.3
Ours high	170.0	9952	6021.8	1.96

has to be larger than n for the hardness of ABDD to hold, whereas Dilithium can choose m smaller than n and base the security of the signature on the hardness of SIS. Even adding the public-key optimization technique from [16] is not likely to make the public key of our scheme smaller than Dilithium.

References

1. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108, 1996.
2. Martin R. Albrecht, Benjamin R. Curtis, Amit Deo, Alex Davidson, Rachel Player, Eamonn W. Postlethwaite, Fernando Virdia, and Thomas Wunderer. Estimate all the {LWE, NTRU} schemes! In *Security and Cryptography for Networks - 11th International Conference, SCN 2018, Amalfi, Italy, September 5-7, 2018, Proceedings*, pages 351–367, 2018.
3. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. *J. Mathematical Cryptology*, 9(3):169–203, 2015.
4. Erdem Alkim, Paulo S. L. M. Barreto, Nina Bindel, Patrick Longa, and Jefferson E. Ricardini. The lattice-based digital signature scheme qtesla. *IACR Cryptology ePrint Archive*, 2019:85, 2019.
5. Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, volume 8366 of *Lecture Notes in Computer Science*, pages 28–47. Springer, 2014.
6. Jonathan Bootle, Claire Delaplace, Thomas Espitau, Pierre-Alain Fouque, and Mehdi Tibouchi. LWE without modular reduction and improved side-channel attacks against BLISS. In *ASIACRYPT (1)*, volume 11272 of *Lecture Notes in Computer Science*, pages 494–524. Springer, 2018.
7. C. BOYD. Digital multisignatures. *Cryptography and Coding*, pages 241–246, 1986.
8. Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In *ITCS*, pages 309–325. ACM, 2012.
9. David Chaum. Blind signatures for untraceable payments. In *Advances in Cryptology: Proceedings of CRYPTO '82, Santa Barbara, California, USA, August 23-25, 1982.*, pages 199–203, 1982.
10. Yuanmi Chen. *Réduction de réseau et sécurité concréte du chiffrement complétement homomorphe*. PhD thesis, Paris 7, 2013.
11. Yvo Desmedt. Society and group oriented cryptography: A new concept. In *Advances in Cryptology - CRYPTO '87, A Conference on the Theory and Applications of Cryptographic Techniques, Santa Barbara, California, USA, August 16-20, 1987, Proceedings*, pages 120–127, 1987.

12. Yvo Desmedt and Yair Frankel. Threshold cryptosystems. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, pages 307–315, 1989.
13. Jintai Ding, Scott R. Fluhrer, and Saraswathy RV. Complete attack on RLWE key exchange with reused keys, without signal leakage. In *Information Security and Privacy - 23rd Australasian Conference, ACISP 2018, Wollongong, NSW, Australia, July 11-13, 2018, Proceedings*, volume 10946 of *Lecture Notes in Computer Science*, pages 467–486. Springer, 2018.
14. Jintai Ding, Xiang Xie, and Xiaodong Lin. A simple provably secure key exchange scheme based on the learning with errors problem. Cryptology ePrint Archive, Report 2012/688, 2012.
15. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO (1)*, volume 8042 of *Lecture Notes in Computer Science*, pages 40–56. Springer, 2013.
16. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystals-dilithium: A lattice-based digital signature scheme. *IACR Trans. Cryptogr. Hardw. Embed. Syst.*, 2018(1):238–268, 2018.
17. Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
18. Scott R. Fluhrer. Cryptanalysis of ring-lwe based key exchange with key share reuse. *IACR Cryptol. ePrint Arch.*, 2016:85, 2016.
19. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206, 2008.
20. Craig Gentry and Mike Szydlo. Cryptanalysis of the revised ntru signature scheme. In *International Conference on the Theory and Applications of Cryptographic Techniques*, pages 299–320. Springer, 2002.
21. Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Public-key cryptosystems from lattice reduction problems. In Burton S. Kaliski Jr., editor, *Advances in Cryptology - CRYPTO '97, 17th Annual International Cryptology Conference, Santa Barbara, California, USA, August 17-21, 1997, Proceedings*, volume 1294 of *LNCS*, pages 112–131. Springer, 1997.
22. Tim Güneysu, Vadim Lyubashevsky, and Thomas Pöppelmann. Practical lattice-based cryptography: A signature scheme for embedded systems. In *Cryptographic Hardware and Embedded Systems - CHES 2012 - 14th International Workshop, Leuven, Belgium, September 9-12, 2012. Proceedings*, pages 530–547, 2012.
23. Siyao Guo, Pritish Kamath, Alon Rosen, and Katerina Sotiraki. Limits on the efficiency of (ring) LWE based non-interactive key exchange. In *Public Key Cryptography (1)*, volume 12110 of *Lecture Notes in Computer Science*, pages 374–395. Springer, 2020.
24. Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. NTRUSIGN: digital signatures using the NTRU lattice. In *CT-RSA*, volume 2612 of *Lecture Notes in Computer Science*, pages 122–140. Springer, 2003.
25. Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of fiat-shamir signatures in the quantum random-oracle model. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 552–586. Springer, 2018.
26. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptography*, 75(3):565–599, 2015.

27. Vadim Lyubashevsky. Fiat-shamir with aborts: Applications to lattice and factoring-based signatures. In *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 598–616. Springer, 2009.
28. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *EUROCRYPT*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.
29. Vadim Lyubashevsky, Daniele Micciancio, Chris Peikert, and Alon Rosen. SWIFFT: A modest proposal for FFT hashing. In *Fast Software Encryption, 15th International Workshop, FSE 2008, Lausanne, Switzerland, February 10-13, 2008, Revised Selected Papers*, pages 54–72, 2008.
30. Vadim Lyubashevsky and Gregory Neven. One-shot verifiable encryption from lattices. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 293–323. Springer, 2017.
31. Daniele Micciancio and Michael Walter. Gaussian sampling over the integers: Efficient, generic, constant-time. In *Advances in Cryptology - CRYPTO 2017 - 37th Annual International Cryptology Conference, Santa Barbara, CA, USA, August 20-24, 2017, Proceedings, Part II*, pages 455–485, 2017.
32. Vincent Migliore, Benoît Gérard, Mehdi Tibouchi, and Pierre-Alain Fouque. Masking dilithium - efficient implementation and side-channel evaluation. In *Applied Cryptography and Network Security - 17th International Conference, ACNS 2019, Bogota, Colombia, June 5-7, 2019, Proceedings*, pages 344–362, 2019.
33. Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. In *Advances in Cryptology - EUROCRYPT 2006, 25th Annual International Conference on the Theory and Applications of Cryptographic Techniques, St. Petersburg, Russia, May 28 - June 1, 2006, Proceedings*, pages 271–288, 2006.
34. Chris Peikert. Lattice cryptography for the internet. In *PQCrypto*, volume 8772 of *Lecture Notes in Computer Science*, pages 197–219. Springer, 2014.
35. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
36. Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *Advances in Cryptology - CRYPTO '89, 9th Annual International Cryptology Conference, Santa Barbara, California, USA, August 20-24, 1989, Proceedings*, pages 239–252, 1989.
37. Dominique Unruh. Post-quantum security of fiat-shamir. In *Advances in Cryptology - ASIACRYPT 2017 - 23rd International Conference on the Theory and Applications of Cryptology and Information Security, Hong Kong, China, December 3-7, 2017, Proceedings, Part I*, pages 65–95. Springer, 2017.