MPSign: A Signature from Small-Secret Middle-Product Learning with Errors

Shi Bai[1], Dipayan Das[2], Ryo Hiromasa[3], Miruna Rosca[4,5], Amin Sakzad[6], Damien Stehlé[4,7], Ron Steinfeld[6], and Zhenfei Zhang[8]

[1] Department of Mathematical Sciences, Florida Atlantic University, USA.
[2] Department of Mathematics, National Institute of Technology, Durgapur, India.
[3] Mitsubishi Electric, Japan.
[4] Univ. Lyon, EnsL, UCBL, CNRS, Inria, LIP, F-69342 Lyon Cedex 07, France.
[5] Bitdefender, Romania.
[6] Faculty of Information Technology, Monash University, Australia.
[7] Institut Universitaire de France.
[8] Algorand, USA.

Abstract. We describe a digital signature scheme MPSign, whose security relies on the conjectured hardness of the Polynomial Learning With Errors problem (PLWE) for at least one defining polynomial within an exponential-size family (as a function of the security parameter). The proposed signature scheme follows the Fiat-Shamir framework and can be viewed as the Learning With Errors counterpart of the signature scheme described by Lyubashevsky at Asiacrypt 2016, whose security relies on the conjectured hardness of the Polynomial Short Integer Solution (PSIS) problem for at least one defining polynomial within an exponential-size family. As opposed to the latter, MPSign enjoys a security proof from PLWE that is tight in the quantum-access random oracle model.

The main ingredient is a reduction from PLWE for an arbitrary defining polynomial among exponentially many, to a variant of the Middle-Product Learning with Errors problem (MPLWE) that allows for secrets that are small compared to the working modulus. We present concrete parameters for MPSign using such small secrets, and show that they lead to significant savings in signature length over Lyubashevsky's Asiacrypt 2016 scheme (which uses larger secrets) at typical security levels. As an additional small contribution, and in contrast to MPSign (or MPLWE), we present an efficient key-recovery attack against Lyubashevsky's scheme (or the inhomogeneous PSIS problem), when it is used with sufficiently small secrets, showing the necessity of a lower bound on secret size for the security of that scheme.

1 Introduction

The Polynomial Short Integer Solution (PSIS) and Polynomial Learning With Errors (PLWE) were introduced as variants of the SIS and LWE problems leading to more efficient cryptographic constructions [LM06,PR06,SSTX09]. Let $n, m, q \geq$

2 and $f \in \mathbb{Z}[x]$ monic of degree n. A $\mathsf{PSIS}_{q,m}^{(f)}$ instance consists in m uniformly chosen elements $a_1, \ldots, a_m \in \mathbb{Z}_q[x]/f$, and the goal is to find $z_1, \ldots, z_m \in \mathbb{Z}[x]/f$ not all zero and with entries of small magnitudes such that $z_1 a_1 + \cdots + z_m a_m = 0 \bmod q$. A $\mathsf{PLWE}_q^{(f)}$ instance consists of oracle access to the uniform distribution over $\mathbb{Z}_q[x]/f \times \mathbb{Z}_q[x]/f$; or to oracle access to the distribution of $(a_i, a_i \cdot s + e_i)$, where a_i is uniform in $\mathbb{Z}_q[x]/f$, $e_i \in \mathbb{Z}[x]/f$ has random coefficients of small magnitudes, and the so-called secret $s \in \mathbb{Z}_q[x]/f$ is uniformly sampled but identical across all oracle calls. The goal is to distinguish between the two types of oracles.

For any fixed f, the hardness of $\mathsf{PSIS}^{(f)}$ and $\mathsf{PLWE}^{(f)}$ has been less investigated than that of SIS and LWE. In particular, it could be that $\mathsf{PSIS}^{(f)}$ and $\mathsf{PLWE}^{(f)}$ are easy, or easier, to solve for some defining polynomials f than for others. To mitigate such a risk, Lyubashevsky [Lyu16] introduced a variant of PSIS that is not parametrized by a specific polynomial f but only a degree n, and is at least as hard as $\mathsf{PSIS}^{(f)}$ for exponentially many polynomials f of degree n. We will let it be denoted by $\mathsf{PSIS}^{\emptyset}$. Further, Lyubashevsky designed a signature scheme whose security relies on the hardness of this new problem, and hence on the hardness of $\mathsf{PSIS}^{(f)}$ for at least one f among exponentially many. This signature scheme enjoys asymptotic efficiency, similar (up to a constant factor) to those based on $\mathsf{PSIS}^{(f)}$ for a fixed f. Later on, Rosca *et al.* [RSSS17] introduced an LWE counterpart of $\mathsf{PSIS}^{\emptyset}$: the Middle-Product Learning with Errors problem (MPLWE). Similarly to $\mathsf{PSIS}^{\emptyset}$, MPLWE is not parametrized by a specific polynomial f but only a degree n, and is at least as hard as $\mathsf{PLWE}^{(f)}$ for exponentially many polynomials f of degree n. To illustrate the cryptographic usefulness of MPLWE, Rosca *et al.* built a public-key encryption scheme whose IND-CPA security relies on the MPLWE hardness assumption. A more efficient encryption scheme and a key encapsulation mechanism ([SSZ17,SSZ19]) were later proposed as a submission to the NIST standardization process for post-quantum cryptography [NIS].

In [RSSS17], it was observed that several $\mathsf{LWE}/\mathsf{PLWE}^{(f)}$ techniques leading to more cryptographic functionalities do not easily extend to MPLWE, possibly limiting its cryptographic expressiveness. These include a polynomial leftover hash lemma, the construction of trapdoors for MPLWE that allow to recover the secret s, and the “HNF-ization” technique of [ACPS09] which would allow to prove hardness of MPLWE with small-magnitude secrets. The leftover hash lemma and trapdoor sampling questions were recently studied in [LVV19], with an application to identity-based encryption, though only for security against an adversary whose distinguishing advantage is non-negligible (as opposed to exponentially small). On the HNF-ization front, the main result of [RSSS17] was mis-interpreted in [Hir18] (see Theorem 1 within this reference), in that the latter work assumed that the hardness result of [RSSS17] was for secrets whose coefficients were distributed as those of noise terms (and hence of small magnitudes). The main result from [Hir18] was a signature scheme with security relying on MPLWE.

1.1 Contributions

In this work, we give a reduction from $\mathsf{PLWE}^{(f)}$ to a variant of MPLWE in which the secret has small-magnitude coefficients. The reduction works for a family of defining polynomials f that grows with the security parameter.

We then build an identification scheme which follows Schnorr's general framework [Sch89] and which can be upgraded to a signature scheme that is tightly secure in the quantum-access random oracle model (QROM), using [KLS18]. We show that MPSign is unforgeable against chosen message attacks ($\mathsf{UF\text{-}CMA}$), which means that no adversary may forge a signature on a message for which it has not seen a signature before. We did not manage to prove that there is no adversary who may forge a new signature on a previously signed message, i.e., that the scheme is strongly unforgeable against chosen message attacks ($\mathsf{UF\text{-}sCMA}$). Nevertheless, any $\mathsf{UF\text{-}CMA}$ secure signature can be upgraded to a $\mathsf{UF\text{-}sCMA}$ secure signature using a one-time $\mathsf{UF\text{-}sCMA}$ secure signature [Kat10]. Such a one-time signature can be achieved easily by a universal one-way hash function (by Lamport's one-time signature) [Kat10] or key collision resistant pseudo-random function (by Winternitz one-time signature) [BDE$^+$11].

We provide concrete parameters for MPSign corresponding to level 1 security of the NIST post-quantum standardization process (via the SVP core hardness methodology from [ADPS16]), which take into account our tight QROM security proof with respect to small secret MPLWE (rather than just taking in account the classical ROM security proof as, e.g., in the Dilithium scheme parameter selection [DKL$^+$18]). We also provide parameters that achieve similar security to those from [Lyu16], to allow for a reasonably fair comparison. The MPSign verification key is larger but its signature size is twice smaller.

Our MPSign signature length savings over the scheme of [Lyu16] arise mainly due to our use of much smaller secret key coordinates. Therefore, one could wonder the reducing the size of the secret key coordinates in the scheme of [Lyu16] would also give a secure signature scheme. As an additional small contribution, we show that the answer is negative by presenting a simple efficient key recovery attack on Lyubashevsky's scheme with sufficiently small secret coordinates. Our attack works (heuristically) when the underlying inhomogeneous variant of $\mathsf{PSIS}^{\emptyset}$ has a unique solution, and shows that a lower bound similar to that shown sufficient in the security proof of [Lyu16] is also *necessary* for the security of Lyubashevsky's scheme (and the underlying inhomogeneous $\mathsf{PSIS}^{\emptyset}$ problem) with small secret coordinates.

Finally, we provide a proof-of-concept implementation in Sage, publicly available at `https://github.com/pqc-ntrust/middle-product-LWE-signature`.

1.2 Comparison with prior works

Our signature construction is similar to the one in [Hir18]. However, the proof of the latter is incorrect: in its proof of high min-entropy of commitments (see [Hir18, Lemma 7]), it is assumed that the middle n coefficients of the product between a uniform $a \in \mathbb{Z}_q[x]$ of degree $< n$ and a fixed polynomial y of degree $\leq 2n$, are uniform. In fact, this distribution depends on the rank of a Hankel matrix associated to y and encoding the linear function from a to the considered coefficients of the product. This Hankel matrix can be of low rank and, when it is the case, the resulting distribution is uniform on a very small subset of the range. Interestingly, the distribution of these Hankel matrices (for a uniform y) was recently studied in [BBD$^+$19], in the context of proving hardness of an MPLWE variant with deterministic noise. We do not know how to fix the error from [Hir18]. As a result, we use a different identification scheme to be able to make our proofs go through. Concretely, the identification scheme from [Hir18] used the Bai-Galbraith [BG14] compression technique to decrease the signature size. We circumvent the difficulty by not using the Bai-Galbraith compression technique.

Lyubashevsky's signature from [Lyu16] can also be viewed as secure under the assumption that $\mathsf{PLWE}^{(f)}$ is hard for at least one f among exponentially many defining polynomials f, like ours. Indeed, it was proved secure under the assumption that $\mathsf{PSIS}^{\emptyset}$ is hard, it was proved that $\mathsf{PSIS}^{(f)}$ reduces to $\mathsf{PSIS}^{\emptyset}$ for exponentially many defining polynomials f, and $\mathsf{PLWE}^{(f)}$ (directly) reduces to $\mathsf{PSIS}^{(f)}$. Furthermore, MPLWE (both with small-magnitude secrets and uniform secrets) reduces to $\mathsf{PSIS}^{\emptyset}$, whereas the converse is unknown. Hence it seems that in terms of assumptions, Lyubashevsky's signature outperforms ours. However, the security proof from [Lyu16] only holds in the random oracle model, as opposed to ours which is tight in the quantum-access random oracle model (QROM). Recent techniques on Fiat-Shamir in the QROM [LZ19,DFMS19] might be applicable to [Lyu16], but they are not tight.

We now compare MPSign with LWE-based signature schemes and efficient lattice-based signature schemes such as those at Round 2 of the NIST post-quantum standardization process [NIS]: Dilithium [DKL$^+$18], Falcon [PFH$^+$19] and Tesla [BAA$^+$19]. Compared to LWE-based signatures, our proposal results in much smaller values for the sum of sizes of a signature and a public key, with much stronger security guarantees than the efficient schemes based on polynomial rings. For example, scaling Dilithium with NIST security level 1 parameters to LWE requires multiplying the public key size by the challenge dimension $n = 256$, since for an LWE adaptation of Dilithium, the public key would be a matrix with n columns instead of 1. For NIST security level 1, the public key and signature sizes sum would be above 300kB for an LWE adaptation of Dilithium, whereas the same quantity is 47KB for MPSign (see Table 2). Now, compared to the Dilithium, Falcon and Tesla NIST candidates, security guarantees are different. The security of Dilithium and Tesla relies on the module variants of PLWE and PSIS for a fixed polynomial [LS15]. In the case of Dilithium, the known security proof in the QROM is quite loose [LZ19], unless one relies on an ad hoc assumption like SelfTargetMSIS [KLS18]. Moreover, in the case of Dilithium, the SIS instance is in an extreme regime: the maximum infinity norm of the vectors to be found are below $q/2$, but their Euclidean norms may be above q. Currently, no reduction backs the assumption that SIS is intractable in that parameter regime. In Falcon, the public key is assumed pseudo-random, which is an adhoc version of the NTRU hardness assumption [HPS98]. Oppositely, the security of MPSign relies on the assumed PLWE hardness for at least one polynomial among exponentially many. Overall, MPSign is an intermediate risk-performance tradeoff between fixed-ring and LWE-based schemes.

2 Preliminaries

The notations in this paper are almost verbatim from [RSSS17] to maintain consistency and facilitate comparison.

Let $q > 1$ be an integer. We let $\mathbb{Z}_q$ denote the ring of integers modulo q and by $\mathbb{Z}_{\leq q}$ the set $\{-q, \ldots, q\}$ of integers of absolute value less or equal to q. We will write $\mathbb{R}_q$ to denote the group $\mathbb{R}/q\mathbb{Z}$.

Let $n > 0$. For a ring R, we will use the notation $R^{<n}[x]$ to denote the set of all polynomials in $R[x]$ with degree less than n. This notation may be extended to any unstructured set S.

For any vector $a = (a_0, a_1, \ldots, a_{n-1})^T \in \mathbb{Z}^n$, we let $\overline{a}$ denote the reversed vector $(a_{n-1}, a_{n-2}, \ldots, a_0)^T \in \mathbb{Z}^n$ and we write $\|a\|_\infty := \max_i |a_i|$. When there is no ambiguity, we identify a polynomial with its vector of coefficients.

For any matrix $M \in \mathbb{R}^{m \times n}$, we let $\sigma_1(M) \geq \sigma_2(M) \geq \cdots \geq \sigma_n(M)$ denote its singular values. We use the notation $\|M\|$ to denote its largest singular value $\sigma_1(M)$ and we denote by $\mathbf{I}_m$ the $m \times m$ identity matrix.

For a distribution D on a set X, we denote by $x \overset{\$}{\leftarrow} D$ the choice of an element x according to D. For simplicity, when D is the uniform distribution on X, we use the notation $a \overset{\$}{\leftarrow} X$.

All logarithms used in this paper are in base 2.

2.1 Polynomials and matrices

For a polynomial $f \in \mathbb{Z}[x]$ of degree $m \geq 1$ and a polynomial $a \in \mathbb{Z}^{<k}[x]$, we make use of the following matrices:

- $\mathsf{Rot}_f^d(a)$: the $d \times m$ matrix whose i-th row is given by the coefficients of the polynomial $x^{i-1} \cdot a \bmod f$;
- M_f: the $m \times m$ matrix whose (i,j)-th element is the constant coefficient of the polynomial $x^{i+j-2} \bmod f$;
- M_f^d: the $d \times m$ matrix obtained by keeping only the first d rows of M_f;
- $\mathsf{Toep}^{d,k}(a)$: the $d \times (k+d-1)$ matrix whose i-th row is given by the coefficients of the polynomial $x^{i-1} \cdot a$.

Note that $\mathsf{Rot}_f^d(a) = \mathsf{Toep}^{d,k}(a) \cdot \mathsf{Rot}_f^{k+d-1}(1)$. Also, for any $a' \in \mathbb{Z}[x]$ such that $a' = a \bmod f$, we have that $\mathsf{Rot}_f^d(a) = \mathsf{Rot}_f^d(a')$.

The expansion factor of a polynomial $f \in \mathbb{Z}[x]$ of degree m is defined as:

$$\mathrm{EF}(f) = \max\left(\frac{\|g \bmod f\|_\infty}{\|g\|_\infty} : g \in \mathbb{Z}^{<2m-1}[x] \setminus \{0\}\right).$$

The following lemma provides bounds on the norms of the matrices M_f and $\mathsf{Rot}_f^d(1)$, in terms of $\mathrm{EF}(f)$. A bound on $\|M_f\|$ was first proved in [RSSS17, Le. 2.8] and improved later in [LVV19, Le. 9]. The bound on $\|\mathsf{Rot}_f^k(1)\|$ can be obtained by noticing that $\mathsf{Rot}_f^k(1)$ contains $\mathbf{I}_{\deg(f)}$ as a submatrix and all its other entries are bounded by $\mathrm{EF}(f)$.

Lemma 1. *Let $f \in \mathbb{Z}[x]$ and $k \geq \deg(f) \geq d$. Then*

1. $\|M_f^d\| \leq \sqrt{d} \cdot \mathrm{EF}(f)$
2. $\|\mathsf{Rot}_f^k(1)\|^2 \leq \deg(f) + (k - \deg(f)) \cdot \deg(f) \cdot \mathrm{EF}(f)^2$.

We now recall the middle-product of two polynomials and some of its elementary properties. Let us consider a pair of polynomials $(a, b) \in \mathbb{Z}^{<d_a}[x] \times \mathbb{Z}^{<d_b}[x]$. Multiplying the two polynomials, we get a polynomial in $\mathbb{Z}^{<d_a+d_b-1}[x]$. If $d_a + d_b - 1 = d + 2k$ for some integers d and k, then the middle-product of size d of a and b is obtained by multiplying a and b, then deleting the coefficients of x^i for $i \leq k-1$ and $i \geq k+d$ and dividing the remaining by x^k. Note that the middle-product is an additive homomorphism when either of its inputs is fixed.

Definition 1 (Middle-Product). *Let d_a, d_b, d, k be integers such that $d_a + d_b - 1 = d + 2k$. The middle-product $\odot_d$ is the map from $\mathbb{Z}^{<d_a}[x] \times \mathbb{Z}^{<d_b}[x]$ to $\mathbb{Z}^{<d}[x]$ defined as: $(a, b) \to a \odot_d b = \lfloor \frac{a \cdot b \mod x^{k+d}}{x^k} \rfloor$.*

Lemma 2 ([RSSS17, Le. 3.2]). *Let $d, k > 0$. For all $r \in \mathbb{Z}^{<k+1}[x]$, $a \in \mathbb{Z}^{<k+d}[x]$ and $b = r \odot_d a$, we have $\overline{b} = \mathsf{Toep}^{d,k+1}(r) \cdot \overline{a}$.*

Lemma 3 ([RSSS17, Le. 3.3]). *Let $d, k, n > 0$. For all $r \in \mathbb{Z}^{<k+1}[x], a \in \mathbb{Z}^{<n}[x]$ and $s \in \mathbb{Z}^{<n+d+k-1}[x]$, we have $r \odot_d (a \odot_{d+k} s) = (r \cdot a) \odot_d s$.*

2.2 Gaussian distributions

A symmetric matrix $\Sigma \in \mathbb{R}^{n \times n}$ is positive definite if $x^t \Sigma x > 0$ for every non-zero vector $x \in \mathbb{R}^n$. For any non-singular matrix $B \in \mathbb{R}^{n \times n}$, the matrix $\Sigma = BB^t$ is positive definite and we say that $B = \sqrt{\Sigma}$. Every positive definite matrix Σ has a square root $B = QD$, where $\Sigma = QD^2Q^t$ is the spectral decomposition of Σ. Note that the square root of a positive definite matrix is not unique ($B' = BH$ is also a square root of Σ for every orthogonal matrix $H \in \mathbb{R}^{n \times n}$). If $\Sigma \in \mathbb{R}^{n \times n}$ is a positive definite matrix, its inverse is also positive definite and, moreover, the set of positive definite matrices is closed under addition.

For a positive definite matrix $\Sigma \in \mathbb{R}^{n \times n}$, we define the Gaussian function on $\mathbb{R}^n$ of covariance matrix Σ as $\rho_\Sigma(x) = \exp(-\pi x^t \Sigma^{-1} x)$ for every $x \in \mathbb{R}^n$. The probability distribution whose density is proportional to ρ_Σ is called the Gaussian distribution and is denoted D_Σ. When $\Sigma = s^2 \cdot \mathbf{I}_n$, we use the notations ρ_s and D_s instead of ρ_Σ and D_Σ, respectively.

Given a (full-rank) lattice $\Lambda \subset \mathbb{R}^n$ we define $\rho_\Sigma(\Lambda) := \sum_{x \in \Lambda} \rho_\Sigma(x)$. Using this, we can now define the discrete Gaussian distribution over Λ of covariance parameter Σ as $D_{\Lambda,\Sigma}(x) = \rho_\Sigma(x)/\rho_\Sigma(\Lambda)$ for every $x \in \Lambda$. The dual of a lattice $\Lambda \subset \mathbb{R}^n$ is $\Lambda^* := \{y \in \mathbb{R}^n : \langle y, x\rangle \in \mathbb{Z} \text{ for every } x \in \Lambda\}$. For $\varepsilon > 0$, we define the smoothing parameter $\eta_\varepsilon(\Lambda)$ as the smallest $r > 0$ such that $\rho_{1/r}(\Lambda^* \setminus \{0\}) \leq \varepsilon$. If $\Lambda_1 \subseteq \Lambda_2$ are two lattices, we have that $\eta_\varepsilon(\Lambda_2) \leq \eta_\varepsilon(\Lambda_1)$. We will use the following standard results.

Lemma 4 ([MR04, Le. 3.3]). *For any full-rank lattice $\Lambda \subset \mathbb{R}^n$ and $\varepsilon > 0$, we have $\eta_\varepsilon(\Lambda) \leq \lambda_n(\Lambda) \cdot \sqrt{\ln(2n(1 + 1/\varepsilon))/\pi}$.*

Lemma 5 ([LPSS14, Le. 5]). *Let $\Sigma_1, \Sigma_2 \in \mathbb{R}^{n\times n}$ two covariance matrices and Λ_1, Λ_2 full-rank lattices in $\mathbb{R}^n$ such that $1 \geq \eta_\varepsilon((\Sigma_1^{-1} + \Sigma_2^{-1})^{1/2} \cdot (\Lambda_1 \cap \Lambda_2))$ for some $\varepsilon \in (0, 1/2)$. If $x_1 \overset{\$}{\leftarrow} D_{\Lambda_1,\Sigma_1}$ and $x_2 \overset{\$}{\leftarrow} D_{\Lambda_2,\Sigma_2}$, then the statistical distance between the distribution of $x_1 + x_2$ and $D_{\Lambda_1+\Lambda_2,\Sigma_1+\Sigma_2}$ is less than 4ε.*

Lemma 6 ([Ban95, Le. 2.10]). *For any full-rank lattice $\Lambda \subset \mathbb{R}^n$ and $\sigma > 0$, we have $\Pr_{x\leftarrow D_{\Lambda,\sigma}}(\|x\|_\infty > \sigma \cdot t) \leq 2n \cdot \exp(-\pi \cdot t^2)$.*

2.3 Polynomial and middle-product learning with errors

In this section we recall the formal definitions of PLWE and MPLWE and of the distributions they make use of.

Definition 2 (PLWE distribution). *Let f be a polynomial of degree m and $q \geq 2$. Let χ be a distribution over $\mathbb{Z}[x]/(f)$ and s a fixed element in $\mathbb{Z}_q[x]/(f)$. We define $\mathsf{P}_{q,\chi}(s)$ as the distribution obtained by sampling $a \overset{\$}{\leftarrow} \mathbb{Z}_q[x]/(f), e \overset{\$}{\leftarrow} \chi$, and returning $(a, b = a \cdot s + e) \in \mathbb{Z}_q[x]/(f) \times \mathbb{Z}_q[x]/(f)$.*

Definition 3 (PLWE). *Let f be a polynomial of degree m and $q \geq 2$. Let χ_1 and χ_2 be distributions over $\mathbb{Z}_q[x]/(f)$. The decision $\mathsf{PLWE}^{(f)}_{q,\chi_1,\chi_2}$ problem consists in distinguishing between arbitrarily many samples from $\mathsf{P}_{q,\chi_1}(s)$ and the same number of uniform samples in $\mathbb{Z}_q[x]/(f) \times \mathbb{Z}_q[x]/(f)$, with non-negligible probability over the choice of $s \overset{\$}{\leftarrow} \chi_2$.*

The hardness of PLWE was investigated in [SSTX09,LPR13,PRS17,RSW18], among others. Of particular importance to the present work, it was observed in [LPR13] that the reduction from uniform secret to small secret described in [ACPS09] in the context of LWE also applies to PLWE.

Lemma 7. *Let f be a polynomial of degree m and $q \geq m$ such that the factors of f modulo q are distinct. Let χ_1 and χ_2 be distributions over $\mathbb{Z}_q[x]/(f)$. Then there is a* ppt *reduction from $\mathsf{PLWE}^{(f)}_{q,\chi_1,\chi_2}$ to $\mathsf{PLWE}^{(f)}_{q,\chi_1,\chi_1}$.*

The condition on q ensures that a uniform element in $\mathbb{Z}_q/(f)$ is invertible with non-negligible probability.

Definition 4 (MPLWE distribution). *Let $n, d > 0$. Let χ be a distribution over $\mathbb{Z}^{<d}[x]$ and $s \in \mathbb{Z}_q^{n+d-1}[x]$. We define $\mathsf{MP}_{q,n,d,\chi}(s)$ as the distribution obtained by sampling $a \overset{\$}{\leftarrow} \mathbb{Z}_q^{<n}[x], e \overset{\$}{\leftarrow} \chi$, and returning $(a, b = a \odot_d s + e) \in \mathbb{Z}_q^{<n}[x] \times \mathbb{Z}_q^{<d}[x]$.*

Definition 5 (MPLWE)**.** *Let $n, d > 0$. Let χ_1 and χ_2 be distributions over $\mathbb{Z}_q^{<d}[x]$ and $\mathbb{Z}_q^{n+d-1}[x]$, respectively. The decision $\mathsf{MPLWE}_{q,n,d,\chi_1,\chi_2}$ problem consists in distinguishing between arbitrarily many samples from $\mathsf{MP}_{q,n,d,\chi_1}(s)$ and the same number of uniform samples in $\mathbb{Z}_q^{<n}[x] \times \mathbb{Z}_q^{<d}[x]$, with non-negligible probability over the choice of $s \overset{\$}{\leftarrow} \chi_2$.*

The PLWE (resp. MPLWE) assumption states that the advantage of any polynomial time algorithm trying to solve the PLWE (resp. MPLWE) problem is negligible. The main result in [RSSS17] is a reduction from a variant of $\mathsf{PLWE}^{(f)}$ (for exponentially many f's with respect to parameter n) for which the noise is drawn from a continuous distribution and the secret is uniformly distributed, to a variant of the MPLWE problem for which the noise distribution is also continuous and the secret is also uniformly distributed. In this work, we will be interested in discrete noise distributions and secret distributions taking small values compared to the modulus q. Compared to [RSSS17], discretizing the noise distribution can be achieved via routine techniques and is more convenient both for our proofs and application. Oppositely, having the secret distribution take small values compared to q requires a new idea.

2.4 Cryptographic definitions

Pseudorandom functions. We will use a pseudorandom function to transform an identification scheme to a deterministic signature scheme.

Definition 6. *A pseudorandom function* PRF *is an efficiently computable map* $\mathsf{PRF} : \mathcal{K} \times \{0,1\}^n \to \{0,1\}$ *where $\mathcal{K}$ is a finite key space and n, k are integers. For any quantum adversary A trying to distinguish the output of the* PRF *from a uniform output, we associate the advantage function*

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

where $\mathsf{RF} : \{0,1\}^n \to \{0,1\}$ *is a uniformly sampled function and A has only classical access to the oracles* $\mathsf{PRF}(K,\cdot)$ *and* $\mathsf{RF}(\cdot)$.

Identification schemes. We recall some basic security properties of particular identification schemes. We closely follow the notations used in [KLS18].

A canonical identification scheme is a protocol between two parties: a prover P and a verifier V. The prover sends a commitment W and the verifier selects a uniform challenge c and sends it to P. Upon receiving c, the prover sends back a response Z to the verifier. After it receives Z, the verifier makes a deterministic decision.

Definition 7 (**Canonical identification scheme**)**.** *A canonical identification scheme is a tuple of classical* ppt *algorithms* $\mathsf{ID} := (\mathsf{IGen}, \mathsf{P}, \mathsf{V})$.

- *The key generation algorithm* IGen *takes as input a security parameter* λ *(in unary) and returns the public and secret keys* $(\mathsf{pk}, \mathsf{sk})$*. The public key defines the set of challenges* ChSet*, the set of commitments* WSet*, and the set of responses* ZSet*.*
- *The prover algorithm* P *consists of two sub-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 inputs the secret key* sk*, a commitment* W*, a challenge* c*, and a state* St *and returns a response* $Z \in \mathsf{ZSet} \cup \{\perp\}$*, where* $\perp \notin \mathsf{ZSet}$ *is a special symbol indicating failure.*
- *The verifier algorithm* V *takes as inputs the public key* pk *and the conversation transcript* (W, c, Z) *and outputs* 1 *(acceptance) or* 0 *(rejection).*

If $Z = \perp$, then we set $(W, c, Z) = (\perp, \perp, \perp)$. The triple $(W, c, Z) \in \mathsf{WSet} \times \mathsf{ChSet} \times \mathsf{ZSet} \cup \{(\perp, \perp, \perp)\}$ generated in this way is called a *transcript*. Given the public key pk, the transcript is valid if $V(pk, W, c, Z) = 1$.

We say that ID has *correctness error* δ if for all public and secret keys generated by IGen, all possible transcripts in $\mathsf{WSet} \times \mathsf{ChSet} \times \mathsf{ZSet}$ with $Z \neq \perp$ are valid and the probability that a honestly generated transcript is $(\perp, \perp, \perp)$ is less than δ.

We say that the canonical identification scheme ID has α *bits of min-entropy* if

$$\Pr_{(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{IGen}(\lambda)} \left(H_\infty(W | (W, St) \leftarrow P_1(\mathsf{sk})) \geq \alpha\right) \geq 1 - 2^{-\alpha}.$$

We are interested in the following security properties.

Definition 8 (No-abort honest-verifier zero-knowledge). *A canonical identification scheme* ID *is* ε_{zk}*-perfect no-abort honest-verifier zero-knowledge (*ε_{zk}*-perfect* $\mathsf{na\text{-}HVZK}$*) if there exists a* ppt *algorithm* Sim *which given only the public key* pk *outputs* (W, c, Z) *such that the statistical distance between* $(W, c, Z) \leftarrow \mathsf{Sim}(\mathsf{pk})$ *and* $(W, c, Z) \leftarrow \mathsf{Trans}(\mathsf{pk})$ *is at most* ε_{zk} *and the element* c *from* $(W, c, Z) \leftarrow \mathsf{Sim}(\mathsf{pk})$ *follows a uniform distribution conditioned on* $c \neq \perp$*.*

Definition 9 (Lossiness). *A canonical identification scheme is lossy (and we call it* LID*) if there exists a lossy key generation algorithm* $\mathsf{LossyIGen}$ *that takes as input* λ *and returns a public key* pk_{ls} *and no secret key such that the public keys generated by* IGen *and* $\mathsf{LossyIGen}$ *are indistinguishable. In other words, for any quantum adversary* A*, the following quantity is negligible:*

$$\begin{aligned}\mathrm{Adv}_{\mathsf{ID}}^{loss}(A) :=& |\Pr(A(\mathsf{pk}_{ls}) = 1 | \mathsf{pk}_{ls} \leftarrow \mathsf{LossyIGen}(\lambda)) \\ &- \Pr(A(\mathsf{pk}) = 1 | (\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{IGen}(\lambda))|.\end{aligned}$$

Definition 10 (Lossy soundness). *A canonical identification scheme is* ε_{ls}*-lossy-sound if, for every quantum adversary* A*, the following probability that* A *could impersonate the prover is less than* ε_{ls}*:*

$$\Pr\left[\mathsf{V}(\mathsf{pk}_{ls}, W^*, c^*, Z^*) = 1 \;\middle|\; \begin{array}{l} \mathsf{pk}_{ls} \leftarrow \mathsf{LossyIGen}(\lambda); \\ (W^*, St) \leftarrow A(\mathsf{pk}_{ls}); \\ c^* \leftarrow \mathsf{ChSet}; Z^* \leftarrow A(St, c^*) \end{array}\right].$$

Digital signatures. We recall the definition of a digital signature.

Definition 11 (Digital signature). *A digital signature scheme* SIG *with correctness error* $\delta \geq 0$ *consists of a triple of* ppt *classical algorithms* $(\mathsf{G}, \mathsf{S}, \mathsf{V})$ *such that for every pair of outputs* $(\mathsf{sk}, \mathsf{vk})$ *of* $\mathsf{G}(1^\lambda)$ *and any message* M*,*

$$\Pr[\mathsf{V}(\mathsf{vk}, M, \mathsf{S}(\mathsf{sk}, M)) = 0] \leq \delta$$

where the probability is taken over the randomness of algorithms S *and* V*.*

The algorithm G *is called the key-generation algorithm,* S *is called the signing algorithm,* V *is called is the verification algorithm. The elements* sk *and* vk *are the signing and verification keys.*

Definition 12 (Unforgeability). *A signature scheme* $\mathsf{SIG} := (\mathsf{G}, \mathsf{S}, \mathsf{V})$ *is said to be unforgeable against one-per-message chosen message attack (*$\mathsf{UF\text{-}CMA_1}$*) in the quantum random oracle model if for every* ppt *quantum forger* $\mathcal{F}$ *having quantum access to the random oracle and classical access to the signing oracle, the probability that after seeing the public key and*

$$\{(M_1, \mathsf{S}(\mathsf{sk}, M_1)), \ldots, (M_Q, \mathsf{S}(\mathsf{sk}, M_Q))\}$$

for any Q *(*$Q = \mathsf{poly}(n)$*) adaptively chosen distinct messages* M_i *of its choice, forger* $\mathcal{F}$ *can produce* $M^* \notin \{M_i\}$ *and* σ^* *such that* $\mathsf{V}(\mathsf{vk}, M^*, \sigma^*) = 1$*, is negligibly small. The probability is taken over the randomness of* $\mathsf{G}, \mathsf{S}, \mathsf{V}$ *and* $\mathcal{F}$*, and is denoted by* $\mathrm{Adv}_{\mathsf{SIG}}^{\mathsf{UF\text{-}CMA_1}}(\mathcal{F})$*.*

One can extend this definition to the scenario where the attacker may have access to more than one signature for any of $\mathsf{poly}(n)$ adaptively chosen messages $\{M_i\}$. In that case, if no quantum adversary $\mathcal{F}$ can produce a valid signature for a message $M^* \notin \{M_i\}$, we say that the signature scheme is unforgeable against chosen message attack ($\mathsf{UF\text{-}CMA}$).

In the *strong* corresponding $\mathsf{UF\text{-}CMA}/\mathsf{UF\text{-}CMA_1}$ experiments, the adversary may return a forgery for a message which has already been queried to the signing oracle, but with a different signature.

As showed in [BPS16], a $\mathsf{UF\text{-}CMA_1}$ signature scheme can be combined with a pseudo-random function to obtain a signature scheme that is $\mathsf{UF\text{-}CMA}$, and the conversion is tight (further, the upgrade preserves strongness). As observed in [KLS18], this transformation still applies when the attacker is quantum and is given quantum access to the random oracle.

From identification schemes to digital signatures: Fiat-Shamir. The Fiat-Shamir heuristic is a technique to convert an identification scheme $\mathsf{ID} := (\mathsf{IGen}, \mathsf{P}, \mathsf{V})$ to a digital signature scheme $\mathsf{SIG} := (\mathsf{G} = \mathsf{IGen}, \mathsf{S}, \overline{\mathsf{V}})$ in the random oracle model (ROM).

The main result in [KLS18] is a security statement of the signature scheme obtained via the Fiat-Shamir transformation in the setup where the adversary has quantum access to the random oracle, but classical access to the signing oracle.

```
$\mathsf{S}$ $(\mathsf{sk}, M)$
$i = 0$
while $Z = \perp$ and $i \leq k_m$ do
    $i = i + 1$
    $(W, St) := \mathsf{P}_1(\mathsf{sk})$
    $c := H(W \| M)$
    $Z := \mathsf{P}_2(\mathsf{sk}, W, c, St)$
end while
if $Z = \perp$ then
    $\sigma = \perp$
else
    $\sigma = (W, Z)$
end if
output $\sigma$
```

```
$\overline{\mathsf{V}}$ $(\mathsf{pk}, M, \sigma)$
$c := H(W \| M)$
output $\mathsf{V}(\mathsf{pk}, W, c, Z) \in \{0, 1\}$
```

Fig. 1: The signature SIG obtained via Fiat-Shamir transform

Theorem 1 ([KLS18, Th. 3.1]). *Consider an identification scheme* ID *which is lossy,* ε_{zk}*-perfect* $\mathsf{na\text{-}HVZK}$*, has* α *bits of entropy and is* ε_{ls}*-lossy sound and the signature scheme* SIG *obtained by applying the Fiat-Shamir transform to the identification scheme* ID*, as in Figure 1.*

For any quantum adversary A *against* $\mathsf{UF\text{-}CMA}_1$ *security that issues at most* Q_H *quantum queries to the random oracle and* Q_S *classical signing queries, there exists a quantum adversary* B *against* ID *such that*

$$\mathrm{Adv}_{\mathsf{SIG}}^{\mathsf{UF\text{-}CMA}_1}(A) \leq \mathrm{Adv}_{\mathsf{ID}}^{loss}(B) + 8(Q_H + 1)^2 \cdot \varepsilon_{ls} + k_m Q_S \cdot \varepsilon_{zk} + 2^{-\alpha+1}.$$

and Time(B)=Time(A)+$k_m Q_H$*.*

Moreover, if we de-randomize the signature scheme in Figure 1 by using a pseudo-random function PRF*, then for any quantum adversary* A *against* $\mathsf{UF\text{-}CMA}$ *security that issues at most* Q_H *quantum queries to the random oracle and* Q_S *classical signing queries, there exists a quantum adversary* B *against* ID *and a quantum adversary* C *against the* PRF *such that*

$$\mathrm{Adv}_{\mathsf{DSIG}}^{\mathsf{UF\text{-}CMA}}(A) \leq \mathrm{Adv}_{\mathsf{ID}}^{loss}(B) + 8(Q_H + 1)^2 \cdot \varepsilon_{ls} + k_m Q_S \cdot \varepsilon_{zk} + 2^{-\alpha+1} + \mathrm{Adv}_{\mathsf{PRF}}^{\mathsf{PR}}(C).$$

The de-randomized version of the signature scheme $\mathsf{DSIG} := (\mathsf{IGen}, \mathsf{DS}, \overline{\mathsf{V}})$ obtained from Fiat-Shamir transformation is given in Figure 2. Here, the PRF key K is also a part of the secret key in the signature scheme.

3 Hardness of middle-product LWE with small secrets

As mentioned earlier, a main obstacle towards building a signature scheme directly from MPLWE with the Fiat-Shamir with aborts methodology is the need of smaller secrets. In this section, we show that MPLWE remains at least as hard as

$\mathsf{DS}\ ((\mathsf{sk}, K), M)$

```
$i = 0$
while $Z = \perp$ and $i \leq k_m$ do
    $i = i + 1$
    $(W, St) := \mathsf{P}_1\,(\mathsf{sk}; \mathsf{PRF}_K(0\|i\|M))$
    $c := H(W\|M)$
    $Z := \mathsf{P}_2\,(\mathsf{sk}, W, c, St; \mathsf{PRF}_K(1\|i\|M))$
end while
if $Z = \perp$ then
    $\sigma = \perp$
else
    $\sigma = (W, Z)$
end if
output $\sigma$
```

$\overline{\mathsf{V}}\ (\mathsf{pk}, M, \sigma)$

```
$c := H(W\|M)$
output $\mathsf{V}(\mathsf{pk}, W, c, Z) \in \{0, 1\}$
```

Fig. 2: The de-randomized signature DSIG obtained via Fiat-Shamir transform

PLWE for numerous parametrizing polynomials f, when the secret s is sampled from a specific distribution χ_s that produces small secrets with overwhelming probability.

Let $q \geq 2$, $n \geq d > 0$, $T > 0$ and $k := n + d - 1$. By $J_i \in \mathbb{Z}^{i \times i}$ we denote the matrix with 1's on the anti-diagonal and 0's everywhere else. Let $\mathcal{E}(T, d, n)$ denote the set of all monic polynomials $g(x) \in \mathbb{Z}[x]$ with constant coefficient coprime to q, degree $m \in [d, n]$, and $\sigma_m(M_f) \geq T$.

Theorem 2. *For any polynomial $f \in \mathcal{E}(T, d, n)$ and $1 \geq \alpha \geq \frac{2\sqrt{n}}{qT}$, there is a* ppt *reduction from $\mathsf{PLWE}^{(f)}_{q,D_{\mathbb{Z}^m,\alpha q},D_{\mathbb{Z}^m,\alpha q}}$ to $\mathsf{MPLWE}_{q,n,d,D_{\mathbb{Z}^d,\alpha'' q},D_{\mathbb{Z}^k,\alpha' q}}$, where $\alpha' = \alpha n\sqrt{2n} \cdot \mathrm{EF}(f)^2$ and $\alpha'' = \alpha\sqrt{2d} \cdot \mathrm{EF}(f)$.*

Proof. We first reduce $\mathsf{PLWE}^{(f)}$ to a variant of MPLWE where the dependence on f lies both in the secret and error distributions. Using the same idea as in [RSSS17, Le. 3.7] except for the fact that now we do not rerandomize the secret to make it uniform, we know that there is a ppt reduction from $\mathsf{PLWE}^{(f)}_{q,\chi_e,\chi_s}$ to $\mathsf{MPLWE}_{q,n,d,\chi'_e,\chi'_s}$ where $\chi'_e = J_d \cdot M_f^d \cdot \chi_e$ and $\chi'_s = J_{n+d-1} \cdot \mathsf{Rot}_f^{d+n-1}(1) \cdot M_f \cdot \chi_s$. We now make the following notations: $B_s := J_k \cdot \mathsf{Rot}_f^k(1) \cdot M_f \cdot \alpha q \mathbf{I}_m$ and $B_e := J_d \cdot M_f^d \cdot \alpha q \mathbf{I}_m$, and $\Sigma_s := B_s \cdot B_s^t \in \mathbb{R}^{k \times k}$ and $\Sigma_e := B_e \cdot B_e^t \in \mathbb{R}^{d \times d}$, respectively. This means that there is a a ppt reduction from $\mathsf{PLWE}^{(f)}_{q,D_{\mathbb{Z}^m,\alpha q},D_{\mathbb{Z}^m,\alpha q}}$ to $\mathsf{MPLWE}_{q,n,d,D_{\mathbb{Z}^d,\Sigma_e},D_{\mathbb{Z}^k,\Sigma_s}}$. We now have, using Lemma 1, that

$$\begin{aligned}
\|\Sigma_s\| &\leq (\alpha q)^2 \cdot \|\mathsf{Rot}_f^{d+n-1}(1)\|^2 \cdot \|M_f\|^2 \\
&\leq (\alpha q)^2 \cdot \big(m + (d + n - 1 - m) \cdot m \cdot \mathrm{EF}(f)^2\big)\, m \cdot \mathrm{EF}(f)^2 \\
&\leq (\alpha q)^2 \cdot (n + (n-1) \cdot n \cdot \mathrm{EF}(f)^2) n \cdot \mathrm{EF}(f)^2 \\
&\leq (\alpha q)^2 \cdot n^3 \cdot \mathrm{EF}(f)^4 < (\alpha' q)^2
\end{aligned}$$

and

$$\|\Sigma_e\| \leq (\alpha q)^2 \cdot \|M_f^d\|^2 \leq d \cdot (\alpha q \cdot \mathrm{EF}(f))^2 < (\alpha'' q)^2.$$

Since $\|\Sigma_s\| < (\alpha' q)^2$ and $\|\Sigma_e\| < (\alpha'' q)^2$, there exist two symmetric positive definite matrices Σ'_s and Σ'_e such that $\Sigma_s + \Sigma'_s = (\alpha' q)^2 \mathbf{I}_k$ and $\Sigma_e + \Sigma'_e = (\alpha'' q)^2 \mathbf{I}_d$. We now replace the rerandomization to uniform of the reduction of [RSSS17, Le. 3.7] by a rerandomization to a Gaussian distribution. We first sample $t \xleftarrow{\$} D_{\mathbb{Z}^k, \Sigma'_s}$. For any $\mathsf{MPLWE}_{q,n,d,D_{\mathbb{Z}^d,\Sigma_e},D_{\mathbb{Z}^k,\Sigma_s}}$ sample (a_i, b_i), we sample $e' \xleftarrow{\$} D_{\mathbb{Z}^d, \Sigma'_e}$ and output $(a'_i, b'_i) = (a_i, b_i + a_i \odot_d t + e'_i)$. If (a_i, b_i) is uniform, so is (a'_i, b'_i). If $b_i = a_i \odot_d s + e_i$, then

$$b'_i = a_i \odot_d s + e_i + a_i \odot_d t + e'_i = a_i \odot_d (s + t) + (e_i + e'_i).$$

The matrices Σ_s, Σ'_s, Σ_e and Σ'_e are all symmetric, so they are in particular orthogonally diagonalizable. Moreover, since Σ_s and Σ'_s (resp. Σ_e and Σ'_e) commute, it means that Σ_s and Σ'_s (resp. Σ_e and Σ'_e) are simultaneously diagonalizable. We can hence write $\Sigma_s = U D_s U^t$ and $\Sigma'_s = U D'_s U^t$ for two diagonal matrices D_s and D'_s such that $(\alpha' q)^2 \mathbf{I}_k = D_s + D'_s$ and an orthogonal matrix $U \in \mathbb{R}^{k \times k}$. Similarly, we can write $\Sigma_e = V D_e V^t$ and $\Sigma'_e = V D'_e V^t$, where D_e and D'_e are diagonal, $D_e + D'_e = (\alpha'' q)^2 \mathbf{I}_d$ and $V \in \mathbb{R}^{d \times d}$ is orthogonal. Since the smoothing parameter is invariant to rotations, we can write

$$\begin{aligned}
\eta_{2^{-k}}(\sqrt{\Sigma_s^{-1} + \Sigma_s'^{-1}} \cdot \mathbb{Z}^k) &= \eta_{2^{-k}}(\sqrt{U(D_s^{-1} + D_s'^{-1})U^t} \cdot \mathbb{Z}^k) \\
&= \eta_{2^{-k}}(U\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k) \\
&= \eta_{2^{-k}}(\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k).
\end{aligned}$$

Using Lemma 4, we have that

$$\eta_{2^{-k}}(\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k) \leq \max_i \sqrt{1/\sigma_i(\Sigma_s) + 1/((\alpha' q)^2 - \sigma_i(\Sigma_s))} \cdot \sqrt{k+1}.$$

We showed that $\sigma_1(\Sigma_s) \leq (\alpha q)^2 \sigma_1(M_f)^2 \sigma_1(\mathsf{Rot}_f^{d+n-1}(1))^2 \leq (\alpha' q)^2/2$, which means that $(\alpha' q)^2 - \sigma_i(\Sigma_s) \geq \sigma_i(\Sigma_s)$ for any $i \leq k$ and thus $1/\sigma_i(\Sigma_s) + 1/(\alpha' q)^2 - \sigma_i(\Sigma_s) \leq 2/\sigma_i(\Sigma_s) \leq 2/\sigma_k(\Sigma_s)$ for any $i \leq k$.

Using the bound on the smallest singular value of M_f, we now get that $\sigma_k(\Sigma_s) \geq (\alpha q)^2 \sigma_m(M_f)^2 \sigma_m(\mathsf{Rot}_f^{n+d-1}(1))^2 \geq (\alpha q)^2 \cdot T^2$, which guarantees that

$$\eta_{2^{-k}}(\sqrt{D_s^{-1} + D_s'^{-1}} \cdot \mathbb{Z}^k) \leq \sqrt{\frac{2}{(\alpha q)^2 \cdot T^2}} \cdot \sqrt{k+1} \leq 1$$

for $\alpha \geq \frac{2\sqrt{n}}{q \cdot T}$. As a consequence, using Lemma 5, the statistical distance between the distribution of $s + t$ and $D_{\mathbb{Z}^k, \alpha' q}$ is $< 4 \cdot 2^{-d} = 4\varepsilon$ as $k > d$.

Similarly, we have $\eta_{2^{-d}}(\sqrt{\Sigma_e^{-1} + \Sigma_e'^{-1}} \cdot \mathbb{Z}^d) \leq 1$ and the statistical distance between the distribution of $e_i + e'_i$ and $D_{\mathbb{Z}^d, \alpha'' q}$ is also $\leq 4\varepsilon$. This completes the proof. □

We notice that in contrast with the reduction from [RSSS17], the above reduction requires a lower bound on the noise parameter α which is used in order to approximate the distribution of the sum of two random discrete variables as in Lemma 5. The following result provides a concrete exponentially large family of polynomials f for which we manage to bound from below the smallest singular value of the matrix M_f.

Lemma 8. *Let $f = x^m + P(x) \in \mathbb{Z}[x]$ with $m \geq 2$ and $\deg(P) \leq m/2$. Then $\sigma_m(M_f) \geq \frac{1}{2+\sqrt{m}\cdot \mathrm{EF}(f)}$.*

Proof. By reordering the rows of M_f, the singular values stay the same and we can view M_f as a block of four matrices $D_1 \in \mathbb{Z}^{\lfloor m/2 \rfloor \times \lfloor m/2 \rfloor}$, $D_2 \in \mathbb{Z}^{\lceil m/2 \rceil \times \lceil m/2 \rceil}$, $\mathbf{0} \in \mathbb{Z}^{\lceil m/2 \rceil \times \lfloor m/2 \rfloor}$ and $T \in \mathbb{Z}^{\lfloor m/2 \rfloor \times \lceil m/2 \rceil}$ in the following way:

$$M_f = \left[\begin{array}{c|c} D_1 & T \\ \hline \mathbf{0} & D_2 \end{array}\right].$$

The matrices D_1 and D_2 are diagonal, $\mathbf{0}$ is the all-0 matrix and T is an upper triangular matrix. We now use the definition $\sigma_m(M_f) = \min(\|M_f \cdot y\|_2 : y \in \mathbb{R}^m, \|y\|_2 = 1)$. Let $y \in \mathbb{R}^m$ such that $\sigma_m(M_f) = \|M_f \cdot y\|_2$ and $\|y\|_2 = 1$. The vector y can be written as $y = (y_0^t | y_1^t)^t$, with $y_0 \in \mathbb{R}^{\lfloor m/2 \rfloor}$ and $y_1 \in \mathbb{R}^{\lceil m/2 \rceil}$. On the one hand, we have:

$$\begin{aligned} \|M_f \cdot y\|_2 \geq \|D_1 \cdot y_0 + T \cdot y_1\|_2 &\geq \|D_1 \cdot y_0\|_2 - \|T \cdot y_1\|_2 \\ &\geq \|y_0\|_2 - \|T\| \cdot \|y_1\|_2 \\ &\geq \|y\|_2 - \|y_1\|_2 - \|M_f\| \cdot \|y_1\|_2 \\ &\geq 1 - (1 + \sqrt{m} \cdot \mathrm{EF}(f)) \cdot \|y_1\|_2, \end{aligned}$$

where the last inequality is by Lemma 1. On the other hand, we also have

$$\|M_f \cdot y\|_2 \geq \|D_2 \cdot y_1\|_2 \geq \|y_1\|_2.$$

This provides the bound

$$\sigma_m(M_f) \geq \max\left(1 - (1 + \sqrt{m} \cdot \mathrm{EF}(f)) \cdot \|y_1\|_2, \|y_1\|_2\right) \geq \frac{1}{2 + \sqrt{m} \cdot \mathrm{EF}(f)}.$$

□

An elementary computation shows that for any polynomial as in the above Lemma 8, we have $\mathrm{EF}(f) \leq \frac{3}{4} m^2 \|P\|_\infty^2$ (see also [LM06, Se. 3.1] for a similar but more general statement). This implies the following corollary of Theorem 2.

Corollary 1. *Fix $S > 0$. For any degree $m \geq 2$ polynomial $f = x^m + P(x) \in \mathbb{Z}[x]$ with constant coefficient coprime with q such that $\deg(P) \leq m/2$ and $\|P\|_\infty^2 \leq 4S/3$ and any $1 \geq \alpha \geq 2\sqrt{n} \cdot (2 + \sqrt{n}S)/q$ there is a* ppt *reduction from $\mathsf{PLWE}^{(f)}_{q, D_{\mathbb{Z}^m, \alpha q}, D_{\mathbb{Z}^m, \alpha q}}$ to $\mathsf{MPLWE}_{q,n,d,D_{\mathbb{Z}^d, \alpha'' q}, D_{\mathbb{Z}^k, \alpha' q}}$, where $\alpha' = \alpha n \sqrt{2n} \cdot S^2$ and $\alpha'' = \alpha \sqrt{2d} \cdot S$.*

4 An attack on Inhomogeneous $\mathsf{PSIS}^{\emptyset}$ with small secrets

In contrast to our hardness result for MPLWE with small secret coordinates shown in the previous section, here we show a simple efficient attack on the Inhomogeneous $\mathsf{PSIS}^{\emptyset}$ problem from [Lyu16] with sufficiently small secret coordinates (such that it has a unique solution). Our algorithm gives a key recovery attack against a small secret variant of the signature scheme of [Lyu16], and shows that a lower bound on the size of the secret key coordinates similar to that in the security proof of [Lyu16] is *necessary* for the security of that signature scheme. MPSign achieves lower signature size than [Lyu16], by using small secret coordinates. The attack presented below shows that a similar improvement in signature size *cannot* be securely achieved in [Lyu16], stressing an MPSign advantage over the approach of [Lyu16].

We recall the definition of the Inhomogeneous $\mathsf{PSIS}^{\emptyset}$ problem (which we denote by I-$\mathsf{PSIS}^{\emptyset}$) from [Lyu16]. The hardness of that problem underlies the security of the key generation algorithm in the signature scheme of [Lyu16]. We note that our definition below is the 'exact' case of the 'approximate' definition in [Lyu16] (with the parameters of [Lyu16, Def. 3.3] set as $c = 1$, $s = \beta$ and $d_1 = d_2 = d$). This restriction makes our attack even stronger since a solution to the exact problem is also a solution to the 'approximate' problem.

Definition 13 (I-$\mathsf{PSIS}^{\emptyset}$). *Let $n, d > 0$. An instance of the I-$\mathsf{PSIS}^{\emptyset}_{q,n,d,k,\beta}$ problem consists of $(a_1, \ldots, a_k, t)$, where $a_i \xleftarrow{\$} \mathbb{Z}_q^{<n}[x]$ for $i = 1, \ldots, k$ and $t = \sum_{i=1}^{k} a_i \cdot s_i \in \mathbb{Z}_q^{<n+d-1}[x]$, where $s_i \xleftarrow{\$} [-\beta, \beta]^{<d}[x]$ for $i = 1, \ldots, k$. A solution to the problem is k elements $(s'_1, \ldots, s'_k)$ with $s'_i \in [-\beta, \beta]^{<d}[x]$ for $i = 1, \ldots, k$ such that*

$$\sum_{i=1}^{k} a_i \cdot s'_i = t.$$

Note that the public key of the signature scheme of [Lyu16] consists of an instance of I-$\mathsf{PSIS}^{\emptyset}$, and a solution is a valid secret key.

Our attack on I-$\mathsf{PSIS}^{\emptyset}$ works in the case where $s_1, \ldots, s_k$ is the unique solution, and consists of a simple greedy algorithm that exploits the zero triangles in the Toeplitz matrices associated with the polynomials a_i, to reduce the problem to a sequence of k-dimensional knapsack subproblems: for each $r < d$, we recover the k-tuple of coefficients of x^r in the polynomials $s_i(x)$ for $i = 1, \ldots, k$. When k is small (as is the case for efficient parameter sets), the attack is efficient.

In more detail, let $t(x) = \sum_{i=1}^{k} a_i(x) \cdot s_i(x) \in \mathbb{Z}_q^{<n+d-1}[x]$ be the target polynomial in an instance of I-$\mathsf{PSIS}^{\emptyset}$. We denote by t_r, $a_{i,r}$ and $s_{i,r}$ the coefficient of x^r in the polynomials $t(x), a_i(x), s_i(x)$, respectively. We observe that for any $r = 0, \ldots, d-1$, the coefficient t_r depends only on the coefficients of x^j for $j \le r$ of the s_i's, namely we have

$$t_r = \sum_{i=1}^{k} \sum_{j=0}^{r} a_{i,j} \cdot s_{i,r-j} = \sum_{i=1}^{k} a_{i,0} \cdot s_{i,r} + \sum_{i=1}^{k} \sum_{j=1}^{r} a_{i,j} \cdot s_{i,r-j}. \qquad (1)$$

Given an instance $(a_1, \dots, a_k, t)$ of the $\mathsf{I\text{-}PSIS}^{\emptyset}_{q,n,d,k,\beta}$ problem, our algorithm works as follows:

```
1 For $r = 0, \dots, d-1$:
    (a) Find *some* vector $s'_{*,r} := (s'_{1,r}, \dots, s'_{k,r}) \in [-\beta, \beta]^k$ such that
        $$t_r = \sum_{i=1}^{k} a_{i,0} \cdot s'_{i,r} + \sum_{i=1}^{k} \sum_{j=1}^{r} a_{i,j} \cdot s'_{i,r-j}. \qquad (2)$$
    (b) If no such vector $s'_{*,r}$ exists, return $\perp$.
2 Return $(s'_1, \dots s'_k)$, where $s'_i = \sum_{j=0}^{d-1} s'_{i,j} x^j$ for $i = 1, \dots, k$.
```

Lemma 9. *Suppose q is prime. With probability $\geq 1 - (4\beta + 1)^k/q$ over the choice of $a_1, \dots, a_k$, the solution $(s'_1, \dots, s'_k) = (s_1, \dots, s_k)$ to the I-$\mathsf{PSIS}^{\emptyset}_{q,n,d,k,\beta}$ problem is unique, and the above algorithm returns this solution in time $(2\beta + 1)^k \cdot \mathsf{poly}(n, d, \log q)$.*

Proof. It follows from (1) that the solution $(s'_1, \dots, s'_k) = (s_1, \dots, s_k)$ satisfies (2) for each r and hence can be output by the algorithm. Now suppose, towards a contradiction, that the algorithm outputs $\perp$ or a different solution $(s'_1, \dots, s'_k) \neq (s_1, \dots, s_k)$. Then let $r^* \geq 0$ denote the *least* iteration r of the algorithm where the solution $s'_{*,r^*} := (s'_{1,r^*}, \dots, s'_{k,r^*})$ to (2) for $r = r^*$ is not equal to $s_{*,r^*} := (s_{1,r^*}, \dots, s_{k,r})$. From (2), we have

$$t_{r^*} = \sum_{i=1}^{k} a_{i,0} \cdot s'_{i,r^*} + \sum_{i=1}^{k} \sum_{j=1}^{r} a_{i,j} \cdot s_{i,r^*-j} = \sum_{i=1}^{k} a_{i,0} \cdot s_{i,r^*} + \sum_{i=1}^{k} \sum_{j=1}^{r} a_{i,j} \cdot s_{i,r^*-j},$$

and hence

$$\sum_{i=1}^{k} a_{i,0} \cdot (s_{i,r^*} - s'_{i,r^*}) = 0.$$

As a consequence, the vector $v^* := (s_{1,r^*} - s'_{1,r^*}, \dots, s_{k,r^*} - s'_{k,r^*}) \neq 0$ satisfies $\sum_{i=1}^{k} a_{i,0} v^*_i = 0$, and $v^* \in [-2\beta, 2\beta]^k$. We claim that such a non-zero vector v^* exists with probability at most $(4\beta + 1)^k/q$ over the uniform choice of the $a_{i,0}$'s. Indeed, since q is prime, the probability that a fixed non-zero vector $v \in [-2\beta, 2\beta]^k$ satisfies $\sum_{i=1}^{k} a_{i,0} v_i = 0$ is $1/q$. A union bound over all $\leq (4\beta + 1)^k$ non-zero vectors in $[-2\beta, 2\beta]^k$ provides the claim. Therefore, the algorithm outputs the unique solution $(s'_1, \dots, s'_k) = (s_1, \dots, s_k)$ with probability at least $1 - (4\beta + 1)^k/q$. The run-time follows since Step 1(a) in the algorithm can be implemented by an exhaustive search through all $(2\beta + 1)^k$ possible values for $s'_{*,r}$. □

We observe that the run-time can be reduced to $2^{O(k)} \cdot \mathsf{poly}(n, d, \log q)$ using a lattice closest vector algorithm to solve the k-dimensional knapsack problems.

By Lemma 9, our algorithm for $\mathsf{I\text{-}PSIS}^{\emptyset}_{q,n,d,k,\beta}$ succeeds with high probability when β is at least slightly smaller than $q^{1/k}/4$, and runs in polynomial time when $k = O(1)$, even for very high degrees n and d. In comparison, the hardness reduction for $\mathsf{I\text{-}PSIS}^{\emptyset}_{q,n,d,k,\beta}$ in [Lyu16, Le. 3.4] requires the lower bound $\beta > 2^{\lambda/(kd)-1} \cdot q^{1/k\cdot(1+n/d)}$ (where λ denotes the security parameter and is such that the success probability of the $\mathsf{I\text{-}PSIS}^{\emptyset}$ attacker handled by the reduction is $> 2^{-\lambda}$). Our attack gives an efficient key recovery attack against the signature scheme of [Lyu16] with small secrets β. For instance, the recommended parameters of the latter scheme have $k = 6$ and $q \approx 2^{30}$ and $\beta \approx 2^{11.5}$, but $\beta < 2^3$ will suffice for our attack to succeed. Moreover, heuristically, we expect that our algorithm will succeed with even larger β corresponding to a unique solution. The run-time is likely in practice to be in the order of minutes on a typical laptop [9], using LLL lattice reduction for solving the 6-dimensional knapsack instances; even a brute-force search of each knapsack instance would take in the order of only $(2\beta)^k < 2^{30}$ arithmetic operations. For the above parameters, our LLL-based implementation solved 7 out of 10 (resp. 2 out of 10) instances with $\beta = 7$ (resp. $\beta = 8$), taking about 3 minutes on a 3.1GHz Intel Core i5 CPU.

5 A signature scheme based on small secrets MPLWE

In this section, we build an identification scheme based on the middle-product learning with errors with small secrets assumption. Then, we show that Theorem 1 is applicable to our construction by checking all the theorem assumptions, as in [KLS18]. As a consequence, by the Fiat-Shamir transformation, we obtain a digital signature scheme that is secure under the middle-product learning with errors with small secrets assumption in the quantum random oracle model.

5.1 The identification scheme

We first present in Figure 3 an identification scheme which makes use of the middle-product of polynomials.

We use an extendable output function *Sam*, i.e., a function on bit strings in which the output can be extended to any required length. If we want the deterministic output y of *Sam* on input x to be uniformly distributed on the set S, we write $y \overset{\$}{\leftarrow} S := Sam(x)$.

The key generation starts by choosing a random string ρ and expanding it into a uniform polynomial $a \in \mathbb{Z}_q^{<n}[x]$ using the function *Sam*. The public key consists of a sample (a, b) drawn from the $\mathsf{MP}_{q,n,d+k,\chi}(s)$ distribution, where both the secret s and the error e follow a Gaussian distribution of parameter $\alpha' q$, respectively $\alpha'' q$.

In the first step of the protocol, the prover chooses two polynomials y_1 and y_2 whose coefficients are bounded in absolute value by a', respectively a'', and sends to the verifier the polynomial $w = a \odot_d y_1 + y_2$. The verifier chooses a random challenge from the challenge space

$$D_H := \{c \in \{0, 1, -1\}^{<k+1}[x] \text{ with } \|c\|_1 = \kappa\}$$

[9] https://github.com/pqc-ntrust/middle-product-LWE-signature

and sends it back to the prover. The challenge space consists of polynomials of small norms and the parameter κ is chosen such that the cardinality of the challenge space is large. The prover now applies rejection in order to make sure that his answer doesn't leak information about the secret key. Concretely, the prover computes $z_1 = c \odot_{n+d-1} s + y_1$ and $z_2 = c \odot_d e + y_1$ and checks if $\|z_1\|_\infty \leq A'$ and $\|z_2\|_\infty \leq A''$. If so, it accepts to send his answer (z_1, z_2) to the verifier. Otherwise, it aborts. We provide concrete parameters with which our scheme can be instantiated in practice in the next section.

```
IGen
$\rho \xleftarrow{\$} \{0,1\}^{256}$
$a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
$s \xleftarrow{\$} D_{\mathbb{Z}^{n+d+k-1},\alpha' q}$
$e \xleftarrow{\$} D_{\mathbb{Z}^{d+k},\alpha'' q}$
$b = a \odot_{d+k} s + e$
$\mathsf{pk} = (\rho, b)$
$\mathsf{sk} = (\rho, s, e)$
output $(\mathsf{pk}, \mathsf{sk})$
```

```
$\mathsf{P}_1$ $(\mathsf{sk})$
$y_1 \xleftarrow{\$} \mathbb{Z}^{<n+d-1}_{\leq a'}[x]$
$y_2 \xleftarrow{\$} \mathbb{Z}^{<d}_{\leq a''}[x]$
$w = a \odot_d y_1 + y_2$
output $W = w, St = (w, y_1, y_2)$
```

```
$\mathsf{P}_2(\mathsf{sk}, W = w, c, St = (w, y_1, y_2))$
$z_1 = c \odot_{n+d-1} s + y_1$
$z_2 = c \odot_d e + y_2$
if $\|z_1\|_\infty > A'$ or $\|z_2\|_\infty > A''$ then
    $(z_1, z_2) = \perp$
end if
output $Z = (z_1, z_2)$
```

```
$\mathsf{V}$ $(\mathsf{pk}, W = w, c, Z = (z_1, z_2))$
$a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
if $w = a \odot_d z_1 + z_2 - c \odot_d b$, $\|z_1\|_\infty \leq A'$ and $\|z_2\|_\infty \leq A''$ then
    output 1
else
    output 0
end if
```

Fig. 3: The identification scheme $(\mathsf{IGen}, \mathsf{V}, \mathsf{P} = (\mathsf{P}_1, \mathsf{P}_2))$

Lemma 10. *If $A' + \|c \odot_{n+d-1} s\|_\infty \leq a'$ and $A'' + \|c \odot_d e\|_\infty \leq a''$, then the identification scheme is perfectly* na-HVZK*, i.e., its transcripts are publicly simulatable and $\varepsilon_{zk} = 0$.*

Proof. Figure 4 (left) shows how to generate a real transcript using the secret key sk, and Figure 4 (right) shows how to simulate a transcript using only the public key pk. The identification scheme is perfectly na-HVZK if every pair of polynomials $(z_1, z_2) \in \mathbb{Z}^{<n+d-1}_{\leq A'}[x] \times \mathbb{Z}^{<d}_{\leq A''}[x]$ has the same probability to be generated in the Trans algorithm as in the Sim algorithm. This is indeed the case: our choice of parameters guarantees that $z_1 - c \odot_{n+d-1} s \in \mathbb{Z}^{<n+d-1}_{\leq a'}[x]$

and $z_2 - c \odot_d e \in \mathbb{Z}^{<d}_{\leq a''}[x]$ and moreover, for any secret key (s, e) and any pair (z_1, z_2), we have that

$$\begin{aligned}&\Pr(z_1 = c \odot_{n+d-1} s + y_1 | y_1 \overset{\$}{\leftarrow} \mathbb{Z}^{<n+d-1}_{\leq a'}[x])\\ &\qquad\qquad = \Pr(y_1 = z_1 - c \odot_{n+d-1} s | y_1 \overset{\$}{\leftarrow} \mathbb{Z}^{<n+d-1}_{\leq a'}[x])\end{aligned}$$

and

$$\Pr(z_2 = c \odot_d e + y_2 | y_2 \overset{\$}{\leftarrow} \mathbb{Z}^{<d}_{\leq a''}[x]) = \Pr(y_2 = z_2 - c \odot_d s | y_2 \overset{\$}{\leftarrow} \mathbb{Z}^{<d}_{\leq a''}[x]).$$

As a consequence, the probability of producing z_1 and z_2 in Trans such that $\|z_1\|_\infty \leq A'$ and $\|z_2\|_\infty \leq A''$ and not returning $\perp$ is $(\frac{2A'+1}{2a'+1})^{n+d-1}(\frac{2A''+1}{2a''+1})^d$, which means that the outputs of Trans and Sim have the same distribution. □

```
Trans (sk)
$a \overset{\$}{\leftarrow} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
$y_1 \overset{\$}{\leftarrow} \mathbb{Z}^{<n+d-1}_{\leq a'}[x]$
$y_2 \overset{\$}{\leftarrow} \mathbb{Z}^{<d}_{\leq a''}[x]$
$w = a \odot_d y_1 + y_2$
$c \overset{\$}{\leftarrow} D_H$
$z_1 = c \odot_{n+d-1} s + y_1$
$z_2 = c \odot_d e + y_2$
if $\|z_1\|_\infty > A'$ or $\|z_2\|_\infty > A''$ then
    output $\perp$
else
    output $(z_1, z_2, c)$
end if
```

```
Sim (pk)
$a \overset{\$}{\leftarrow} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
with probability
  $1 - (\frac{2A'+1}{2a'+1})^{n+d-1}(\frac{2A''+1}{2a''+1})^d$
output $\perp$
$c \overset{\$}{\leftarrow} D_H$
$z_1 \overset{\$}{\leftarrow} \mathbb{Z}^{<n+d-1}_{\leq A'}[x]$
$z_2 \overset{\$}{\leftarrow} \mathbb{Z}^{<d}_{\leq A''}[x]$
output $(z_1, z_2, c)$
```

Fig. 4: The transcript Trans and the simulation Sim algorithms

Lemma 11. *The scheme has correctness error* $\delta = 1 - (\frac{2A'+1}{2a'+1})^{n+d-1}(\frac{2A''+1}{2a''+1})^d$.

Proof. First, we show that the verification procedure always accepts a honest transcript if $(z_1, z_2) \neq \perp$. Assume that $(z_1, z_2) \neq \perp$. It means that $\|z_1\|_\infty \leq A'$ and $\|z_2\|_\infty \leq A''$. Now we prove that

$$a \odot_d z_1 + z_2 - c \odot_d b = a \odot_d y_1 + y_2.$$

Because of Lemma 3, we have that

$$\begin{aligned}a \odot_d z_1 =& a \odot_d (c \odot_{n+d-1} s + y_1)\\ =& a \odot_d (c \odot_{n+d-1} s) + a \odot_d y_1\\ =& (a \cdot c) \odot_d s + a \odot_d y_1\end{aligned}$$

and

$$\begin{aligned}c \odot_d b =& c \odot_d (a \odot_{d+k} s + e)\\ =& c \odot_d (a \odot_{d+k} s) + c \odot_d e\\ =& (c \cdot a) \odot_d s + c \odot_d e.\end{aligned}$$

Overall, we obtain:

$$\begin{aligned}& a \odot_d z_1 + z_2 - c \odot_d b\\ &= ((a \cdot c) \odot_d s + a \odot_d y_1) + (c \odot_d e + y_2) - ((c \cdot a) \odot_d s + c \odot_d e)\\ &= a \odot_d y_1 + y_2.\end{aligned}$$

Since Sim outputs $\perp$ with the same probability as Trans, we know that the probability to have $(z_1, z_2) = \perp$ is exactly δ. □

Lemma 12. *The identification scheme* ID *is lossy.*

Proof. In the lossy key generation algorithm $\mathsf{LossyIGen}$ (Figure 5), we generate the public key (a, b) uniformly. The public keys generated by IGen and $\mathsf{LossyIGen}$ are indistinguishable by the MPLWE assumption. Indeed, for any quantum adversary A against ID, there exists an adversary B trying to distinguish MPLWE samples from uniform ones such that the loss advantage $\mathrm{Adv}_{\mathsf{ID}}^{loss}(A)$ is equal to the advantage of B. □

Lemma 13. *The identification scheme* ID *has* $d \cdot \log(2a''+1)$ *bits of min-entropy.*

Proof. Indeed, for every commitment ω, we have that:

$$\Pr_{a, y_1, y_2}(a \odot_d y_1 + y_2 = \omega) \leq \max_{a, y_1} \Pr_{y_2}(y_2 = \omega - a \odot_d y_1) \leq \frac{1}{(2a''+1)^d},$$

where the first probability is taken over the uniform choice of $a \in \mathbb{Z}_q^{<n}[x]$, $y_1 \in \mathbb{Z}_{\leq a'}^{<n+d-1}[x]$ and $y_2 \in \mathbb{Z}_{\leq a''}^{<d}[x]$. In the second one, the probability is taken over the uniform choice of $y_2 \in \mathbb{Z}_{\leq a''}^{<d}[x]$ and the maximum is taken over all $a \in \mathbb{Z}_q^{<n}[x]$ and $y_1 \in \mathbb{Z}_{\leq a'}^{<n+d-1}[x]$. □

Lemma 14. *The identification scheme* ID *is* ε_{ls}*-lossy-sound, where*

$$\varepsilon_{ls} \leq \frac{1}{|D_H|} + (4A'+1)^{n+d-1} \cdot (4A''+1)^d \cdot |D_H|^2 \cdot q^{-d}.$$

Proof. We show that relatively to a lossy key pk_{ls} generated by the $\mathsf{LossyIGen}$ algorithm in Figure 5, not even an unbounded quantum adversary can impersonate the prover. This reduces to the computation of the following probability taken over the uniform choice of $a \in \mathbb{Z}_q^{<n}[x]$, $b \in \mathbb{Z}_q^{<d+k}[x]$ and $c \in D_H$:

$$P := \Pr(\exists\ z_1 \in \mathbb{Z}_{\leq A'}^{<n+d-1}[x], z_2 \in \mathbb{Z}_{\leq A''}^{<d}[x] : a \odot_d z_1 + z_2 - c \odot_d b = w).$$

```
LossyIGen
    $\rho \xleftarrow{\$} \{0,1\}^{256}$
    $a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
    $b \xleftarrow{\$} \mathbb{Z}_q^{<d+k}[x]$
    output $\mathsf{pk}_{ls} = (a, b)$
```

Fig. 5: The LossyIGen algorithm

Let S denote the set of pairs (a, b) such that there exists at most one c for which there exist small z_1, z_2 such that $a \odot_d z_1 + z_2 - c \odot_d b = w$. We can write $P \leq P_1 + P_2$, where

$$P_1 = \Pr((a, b) \in S) \cdot \frac{1}{|D_H|} \leq \frac{1}{|D_H|}$$

and

$$\begin{aligned} P_2 &\leq \Pr((a, b) \notin S) \cdot 1 \\ &\leq \Pr(\exists\ c \neq c', z_1, z_2, z_1', z_2' : a \odot_d (z_1 - z_1') + z_2 - z_2' - (c - c') \odot_d b = 0) \\ &= \Pr(\exists\ e_c \in D_H - D_H \setminus \{0\}, e_1 \in \mathbb{Z}^{<n+d-1}_{\leq 2A'}, e_2 \in \mathbb{Z}^{<d}_{\leq 2A''} : \\ &\qquad\qquad a \odot_d e_1 + e_2 - e_c \odot_d b = 0), \end{aligned}$$

where a and b are uniformly sampled in $\mathbb{Z}_q^{<n}[x]$, respectively $\mathbb{Z}_q^{<d+k}[x]$, $c, c' \in D_H$, $z_1, z_1 \in \mathbb{Z}^{<n+d-1}_{\leq A'}[x]$, and $z_2, z_2' \in \mathbb{Z}^{<d}_{\leq A''}[x]$ and $D_H - D_H$ denotes the set $\{d - d' \mid d, d' \in D_H\}$.

Let us fix $(e_c \neq 0, e_1, e_2)$. The rank of $\mathsf{Toep}(e_c)$ is maximum for $e_c \neq 0$, which means that the function $b \mapsto e_c \odot_d b$ maps an element b from the uniform distribution on $\mathbb{Z}_q^{<d+k}[x]$ to an element b' from the uniform distribution on $\mathbb{Z}_q^{<d}[x]$. We can now write:

$$\Pr(a \odot_d e_1 + e_2 - e_c \odot_d b = 0) = \Pr(b' = a \odot_d e_1 + e_2) = q^{-d},$$

where the first probability is taken over the uniform choice of $a \in \mathbb{Z}_q^{<n}[x]$ and $b \in \mathbb{Z}_q^{<d+k}[x]$ and the second one is taken over the choice of $a \in \mathbb{Z}_q^{<n}[x]$ and $b' \in \mathbb{Z}_q^{<d}[x]$. We conclude that $P_2 \leq (4A' + 1)^{n+d-1} \cdot (4A'' + 1)^d \cdot |D_H|^2 \cdot q^{-d}$. □

5.2 The signature scheme

In Figure 6, we present our digital signature scheme which is obtained by the de-randomized Fiat-Shamir transform of the identification scheme ID. The correctness of the signature scheme follows (see [KLS18, p. 11]) from the correctness of the underlying identification scheme (Lemma 11). The scheme is UF-CMA secure in the quantum random oracle model, as discussed in Subsection 2.4.

The signature scheme relies on a hash function $H : \{0,1\}^* \to D_H$, which outputs elements with small norms and will be modelled by a random oracle in the security proof. We refer to [DDLL13] for an efficient method to construct such a hash function.

```
KeyGen
$\rho \xleftarrow{\$} \{0,1\}^{256}$
$a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
$s \xleftarrow{\$} D_{\mathbb{Z}^{n+d+k-1},\alpha' q}$
$e \xleftarrow{\$} D_{\mathbb{Z}^{d+k},\alpha'' q}$
$b = a \odot_{d+k} s + e$
$\mathsf{vk} = (b, \rho)$
$\mathsf{sk} = (s, e, K, \rho)$
output $(\mathsf{sk}, \mathsf{vk})$
```

```
Sign $(\mathsf{sk} = (s, e, K, \rho), M)$
$a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
$i = 0$
while $(z_1, z_2) = \perp$ and $i \leq k_m$ do
    $i = i + 1$
    $y_1 \xleftarrow{\$} \mathbb{Z}_{<a'}^{<n+d-1}[x] := Sam(K\|M\|i\|0)$
    $y_2 \xleftarrow{\$} \mathbb{Z}_{<a''}^{<d}[x] := Sam(K\|M\|i\|1)$
    $w = a \odot_d y_1 + y_2$
    $c := H(w\|M)$
    $z_1 = c \odot_{n+d-1} s + y_1$
    $z_2 = c \odot_d e + y_2$
    if $\|z_1\|_\infty > A'$ or $\|z_2\|_\infty > A''$ then
        $(z_1, z_2) = \perp$
    end if
end while
output $(z_1, z_2, c)$
```

```
Verify $(\mathsf{vk} = (b, \rho), M, (z_1, z_2, c))$
$a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$
$w = a \odot_d z_1 + z_2 - c \odot_d b$
if $c = H(w\|M)$, $\|z_1\|_\infty \leq A'$ and $\|z_2\|_\infty \leq A''$ then
    output 1
else
    output 0
end if
```

Fig. 6: The signature scheme

The key generation algorithm samples $a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x]$ using the extendable function Sam seeded with a 256-bit seed ρ, and then two small secret polynomials $s \xleftarrow{\$} D_{\mathbb{Z}^{n+d+k-1},\alpha' q}$ and $e \xleftarrow{\$} D_{\mathbb{Z}^{d+k},\alpha'' q}$. It outputs $(b = a \odot_{d+k} s + e, \rho)$ as the verification key vk and (s, e, K, ρ) as the signing key sk, K being a random key for the pseudorandom function $Sam(K\|\cdot)$ used in the signature algorithm.

To sign a message M, we first recompute $a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$, generate deterministic masking parameters $y_1 \xleftarrow{\$} \mathbb{Z}_{<a'}^{<n+d-1}[x] := Sam(K\|M\|i\|0)$ and $y_2 \xleftarrow{\$} \mathbb{Z}_{<a''}^{<d}[x] := Sam(K\|M\|i\|1)$, where i is the repetition index and compute $w = a \odot_d y_1 + y_2$. Then we compute $c := H(w\|M)$, $z_1 = c \odot_{n+d-1} s + y_1$ and $z_2 = c \odot_d e + y_2$. A potential signature is now (z_1, z_2, c). In order to make the signature pair (z_1, z_2) independent of the signing key, we perform rejection

sampling on potential signatures before outputting the right one. A potential signature (z_1, z_2, c) is output if both $\|z_1\|_\infty \leq A'$ and $\|z_2\|_\infty \leq A''$.

To check if (z_1, z_2, c) is a valid signature for a message M, we first recompute $a \xleftarrow{\$} \mathbb{Z}_q^{<n}[x] := Sam(\rho)$ and $w = a \odot_d z_1 + z_2 - c \odot_d b$ and we accept if $\|z_1\|_\infty \leq A'$, $\|z_2\|_\infty \leq A''$ and $c := H(w\|M)$.

6 Concrete parameters

In this section we give sample parameters with which our digital signature scheme can be instantiated. The choice of parameters takes into account the correctness error probability, the security and the efficiency of our scheme.

The signing acceptance probability is set to $p = 1/3$ as in [Lyu16] for a fair comparison.

The security proof of the scheme from [Lyu16] uses the random oracle model, while the security of our scheme, which is based on Theorem 1, holds in the more powerful quantum random oracle model.

In terms of efficiency, we focus on minimizing the size of a signature. Our signature size is $(n + d - 1)\lceil\log(A')\rceil + d\lceil\log(A'')\rceil + \kappa(\lceil\log(k+1)\rceil + 1)$ bits. The optimal value of d/n for minimizing the signature length is close to 0.5. As d/n reduces below 0.5, the signature dimension drops. Due to the lossiness condition, d/n and $\log q$ are inversely proportional, so we have to increase n to maintain security, which means that overall the signature length will increase. If d/n increases towards 1, $\log q$ reduces but the signature dimension increases and we cannot reduce the signature length.

The size of our public key (a, b) is $256 + (d + k)\lceil\log(q)\rceil$. Since for our lossiness property in the security proof we need a much larger q than the one used in [Lyu16], our public key becomes larger than the public key used in [Lyu16]. On the other hand, our scheme has significantly shorter signatures. Our savings in MPSign signature length over the scheme in [Lyu16] arise largely from the smaller secret key coordinates in MPSign. As our attack of Section 4 shows, such savings are not possible in the scheme of [Lyu16] due to the insecurity of $\mathrm{PSIS}^{\emptyset}$ with sufficiently small secret coordinates.

In order to set concrete parameters for our scheme achieving λ bits of security, we need to bound from above the advantage of any adversary trying to attack the UF-CMA security of MPSign in the quantum random oracle model by $2^{-\lambda}$. By Theorem 1 and Lemma 12, it is enough to bound Adv, $\mathrm{Adv}_{\mathsf{PRF}}^{PR}(C)$ and $2^{-d\log(2a'+1)+1}$ by $2^{-\lambda}/5$ and $8(Q_H + 1)^2 \cdot \varepsilon_{ls}$ by $2^{-\lambda+1}/5$, where the notations are those from Section 5 and Adv stands for the advantage of an adversary trying to solve the $\mathsf{MPLWE}_{q,n,d+k,\chi_1,\chi_2}$ problem, where both χ_1 and χ_2 are discrete Gaussians of parameters $\alpha'' q$, respectively $\alpha' q$. As it is standard in lattice-based cryptography, we further neglect the noise amplification in Theorem 2 and assume that the MPLWE problem with very small secret (with $\|s\|_\infty \approx 1$) is concretely as hard as the $\mathsf{PLWE}^{(f)}$ problem with very small secret. Indeed, there are no known attacks on the MPLWE with small secrets problem that exploit the very small secret when generic algebraic attacks on LWE are protected against (see, e.g., [AG11,ACF$^+$15a,ACF$^+$15b]). Since the discrete Gaussian distributions of the error and secret have small standard deviation, we assume that we can safely replace them by a corresponding centered binomial distribution, as has been done in many practical lattice-based encryption schemes (see [ADPS16,SSZ19,BDK$^+$19], among others).

We use [APS15] in order to estimate both the classical and quantum bit complexities of the primal attack against the $\mathsf{PLWE}^{(f)}$ problem associated to a polynomial f of maximum degree n from the family. The cost models we choose are bkz.sieve for classical security, respectively bkz.qsieve for quantum security.

We present in Table 1 a comparison between the efficiency of MPSign and the scheme described in [Lyu16]. For the same Hermite factor $\delta_0 = 1.005$ (driving the security level), by choosing $n = 2500$, $d = 1300$, $k = 512$ for our scheme, we manage to shorten the size of a signature by a factor of 2.1 and the size of the secret key by a factor of 11 at the cost of doubling the size of the public key.

	MPSign	[Lyu16]
public key size	19 KB	9.6 KB
secret key size	0.7 KB	8.8 KB
signature size	13 KB	27 KB
q	$\approx 2^{87}$	$\approx 2^{30}$

Table 1: Efficiency of MPSign

In the first column of Table 2, we provide concrete parameters for MPSign that satisfy both classical and quantum level 1 NIST requirements. Concretely, they achieve $\lambda \geq 143$ for classical adversaries and $\lambda \geq 130$ for quantum adversaries. The second column contains parameters for $\lambda = 89$ bits of quantum security, corresponding to a Hermite factor $\delta = 1.005$.

7 Implementation

We implemented MPSign in Sage (Python) as a proof-of-concept and the source code is publicly available.[10] For the experiments, we used a MacBook Pro with Intel i7-8559U CPU at 2.7 GHz. Turbo-boost and hyperthreading were both disabled. For a fair comparison, we also implemented the scheme from [Lyu16]. It is expected that both implementations are slower than if they were implemented with a system language (such as `C`) with an aim for optimization. Nonetheless, since both implementations use the same Gaussian sampler, the same *hash to challenge* function, and the same polynomial multiplication algorithm, we believe that the comparison is relatively fair.

We instantiate MPSign and the scheme from [Lyu16] with corresponding parameters achieving $\delta = 1.005$. (for MPSign these parameters may be found in

[10] `https://github.com/pqc-ntrust/middle-product-LWE-signature`

	$\lambda = 143$	$\lambda = 89$
n	3800	2500
d	1910	1300
k	512	512
q	$\approx 2^{90.9}$	$\approx 2^{87.3}$
κ	53	53
$\lvert D_H \rvert$	$\approx 2^{294}$	$\approx 2^{294}$
$\log A'$	≈ 21.0	≈ 20.4
$\log A''$	≈ 19.4	≈ 18.9
δ	1.004	1.005
$\alpha' q$	$2\sqrt{\pi}$	$2\sqrt{\pi}$
$\alpha'' q$	$2\sqrt{\pi}$	$2\sqrt{\pi}$
public key size	26.9 KB	19.5 KB
secret key size	1.06 KB	0.74 KB
signature size	20.1 KB	12.8 KB

Table 2: Sample parameters for MPSign

Table 2). In both benchmarks we iterated 1000 times, each time with a different seed and a different message to sign. The results of our comparison may be found in Table 3. The data are for the average cost in milliseconds. Our scheme is almost twice faster than the one from [Lyu16] in key generation and verification, and four times faster in signing. This is mainly due to the fact that the scheme from [Lyu16] requires scalar multiplications over vectors of polynomials, while our scheme involves a single middle-product (over a somewhat longer polynomial).

	[Lyu16]			MPSign		
	min	ave	max	min	ave	max
key generation	22.3	**25.9**	46.7	14.6	**16.3**	27.1
signing	111	**418**	5771	28.3	**99.6**	713
verification	15.0	**30.8**	53.0	16.3	**18.8**	28.6

Table 3: Performance comparison, in ms

Acknowledgments. The work of Shi Bai has been supported in part through NIST awards 60NANB18D216 and 60NANB18D217 and through NATO SPS Project G5448. The work of Damien Stehlé has been supported by BPI-France in the context of the national project RISQ (P141580) and by the European Union in the context of the PROMETHEUS project (Horizon 2020 Research and Innovation Program, grant 780701). Part of this work was done while Damien Stehlé was visiting the Simons Institute for the Theory of Computing. The work of Ron Steinfeld was supported in part by ARC Discovery Project grant DP180102199.

References

ACF$^+$15a. Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, Robert Fitzpatrick, and Ludovic Perret. Algebraic algorithms for LWE problems. *ACM Comm. Computer Algebra*, 49(2):62, 2015.

ACF$^+$15b. Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, Robert Fitzpatrick, and Ludovic Perret. On the complexity of the BKW algorithm on LWE. *Designs, Codes and Cryptography*, 74(2):325–354, 2015.

ACPS09. Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *CRYPTO*, pages 595–618. Springer, 2009.

ADPS16. Erdem Alkim, Léo Ducas, Thomas Pöppelmann, and Peter Schwabe. Post-quantum key exchange - A new hope. In *USENIX*, pages 327–343, 2016.

AG11. Sanjeev Arora and Rong Ge. New algorithms for learning in presence of errors. In *ICALP*, pages 403–415. Springer, 2011.

APS15. Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of LWE. *J. Mathematical Cryptology*, 9(3):169–203, 2015.

BAA$^+$19. 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: Algorithm specifications and supporting documentation. *NIST PQC round 2 submission document*, 2019.

Ban95. W. Banaszczyk. Inequalities for convex bodies and polar reciprocal lattices in $\mathbb{R}^n$. In *Discrete and Computational Geometry*, volume 13, pages 217–231. Springer, 1995.

BBD$^+$19. Shi Bai, Katharina Boudgoust, Dipayan Das, Adeline Roux-Langlois, Weiqiang Wen, and Zhenfei Zhang. Middle-product learning with rounding problem and its applications. In *ASIACRYPT*, pages 55–81. Springer, 2019.

BDE$^+$11. Johannes A. Buchmann, Erik Dahmen, Sarah Ereth, Andreas Hülsing, and Markus Rückert. On the security of the winternitz one-time signature scheme. In *AFRICACRYPT*, pages 363–378, 2011.

BDK$^+$19. Joppe W. Bos, Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, and Damien Stehlé. CRYSTALS - Kyber: a CCA-secure module-lattice-based KEM. In *Euro S P*, pages 353–367, 2019.

BG14. Shi Bai and Steven D. Galbraith. An improved compression technique for signatures based on learning with errors. In *CT-RSA*, pages 28–47, 2014.

BPS16. Mihir Bellare, Bertram Poettering, and Douglas Stebila. From identification to signatures, tightly: A framework and generic transforms. In *ASIACRYPT*, pages 435–464, 2016.

DDLL13. Léo Ducas, Alain Durmus, Tancrède Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In *CRYPTO*, pages 40–56. Springer, 2013.

DFMS19. Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. In *CRYPTO*, pages 356–383, 2019.

DKL$^+$18. Léo Ducas, Eike Kiltz, Tancrède Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - Dilithium: Digital signatures from module lattices. In *CHES*, pages 238–268, 2018.

Hir18. Ryo Hiromasa. Digital signatures from the middle-product LWE. In *ProvSec*, pages 239–257, 2018.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H Silverman. NTRU: A ring-based public key cryptosystem. In *International Algorithmic Number Theory Symposium*, pages 267–288. Springer, 1998.

Kat10. Jonathan Katz. *Digital Signatures.* Springer, 2010.

KLS18. Eike Kiltz, Vadim Lyubashevsky, and Christian Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. In *EUROCRYPT*, pages 552–586, 2018.

LM06. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP*, pages 144–155, 2006.

LPR13. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43:1–43:35, 2013.

LPSS14. San Ling, Duong Hieu Phan, Damien Stehlé, and Ron Steinfeld. Hardness of k-LWE and applications in traitor tracing. In *CRYPTO*, pages 315–334. Springer, 2014.

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

LVV19. Alex Lombardi, Vinod Vaikuntanathan, and Thuy Duong Vuong. Lattice trapdoors and IBE from middle-product LWE. In *TCC*, pages 24–54. Springer, 2019.

Lyu16. Vadim Lyubashevsky. Digital signatures based on the hardness of ideal lattice problems in all rings. In *ASIACRYPT*, pages 196–214. Springer, 2016.

LZ19. Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In *CRYPTO*, pages 326–355. Springer, 2019.

MR04. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. In *FOCS*, pages 372–381. IEEE, 2004.

NIS. NIST. Post-Quantum Cryptography - Round 1 Submissions. `https://csrc.nist.gov/Projects/Post-Quantum-Cryptography/Round-1-Submissions`.

PFH$^+$19. Thomas Prest, Pierre-Alain Fouque, Jeffrey Hoffstein, Paul Kirchner, Vadim Lyubashevsky, Thomas Pornin, Thomas Ricosset, Gregor Seiler, William Whyte, and Zhenfei Zhang. Falcon: Algorithm specifications and supporting documentation. *NIST PQC round 2 submission document*, 2019.

PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, pages 145–166. Springer, 2006.

PRS17. Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *STOC*, pages 461–473. ACM, 2017.

RSSS17. Miruna Rosca, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Middle-product learning with errors. In *CRYPTO*, pages 283–297. Springer, 2017.

RSW18. Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the ring-LWE and polynomial-LWE problems. In *EUROCRYPT*, pages 146–173. Springer, 2018.

Sch89. Claus-Peter Schnorr. Efficient identification and signatures for smart cards. In *CRYPTO*, pages 239–252. Springer, 1989.

SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, pages 617–635. Springer, 2009.

SSZ17. Ron Steinfeld, Amin Sakzad, and Raymond Kuo Zhao. Titanium: Proposal for a NIST Post-Quantum Public-key encryption and KEM standard, 2017.

SSZ19. Ron Steinfeld, Amin Sakzad, and Raymond Kuo Zhao. Practical MP-LWE-based encryption balancing security-risk versus efficiency. *Designs, Codes and Cryptography*, 87(12):2847–2884, 2019.