Middle-Product Learning With Errors

Miruna Roşca[1,2], Amin Sakzad[3], Damien Stehlé[1], and Ron Steinfeld[3]

[1] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France
[2] Bitdefender, Romania
[3] Faculty of Information Technology, Monash University, Australia

Abstract. We introduce a new variant MP-LWE of the Learning With Errors problem (LWE) making use of the Middle Product between polynomials modulo an integer q. We exhibit a reduction from the Polynomial-LWE problem (PLWE) parametrized by a polynomial f, to MP-LWE which is defined independently of any such f. The reduction only requires f to be monic with constant coefficient coprime with q. It incurs a noise growth proportional to the so-called expansion factor of f. We also describe a public-key encryption scheme with quasi-optimal asymptotic efficiency (the bit-sizes of the keys and the run-times of all involved algorithms are quasi-linear in the security parameter), which is secure against chosen plaintext attacks under the MP-LWE hardness assumption. The scheme is hence secure under the assumption that PLWE is hard for at least one polynomial f of degree n among a family of f's which is exponential in n.
Keywords. LWE, PLWE, public-key encryption.

1 Introduction

Lattice-based cryptography relies in great parts on the assumed hardness of two well-studied and closely related problems: the Small Integer Solution problem (SIS) introduced in [Ajt96] and the Learning With Errors problem (LWE) introduced in [Reg09]. They lead to numerous cryptographic constructions, are conjectured exponentially hard to solve even for quantum algorithms, and enjoy reductions from standard worst-case lattice problems such as finding a short non-zero vector in a lattice (ApproxSVP). However, the resulting cryptographic constructions suffer from large keys and/or rather inefficient algorithms. This is because the problems themselves involve large-dimensional random matrices over a ring $\mathbb{Z}_q$ (for some $q \geq 2$).

To obtain more efficient SIS-based primitives, Lyubashevsky and Micciancio [LM06], and Peikert and Rosen [PR06] introduced the Polynomial SIS problem (PSIS), inspired from [Mic07,HPS98].[4] $\mathsf{PSIS}^{(f)}$ can be

[4] The problem was called Ideal-SIS in [LM06], Cyclotomic-SIS in [PR06], and is now commonly referred to as Ring-SIS. We prefer to call it PSIS as it is not defined in

described in terms of elements of $\mathbb{Z}_q[x]/f$ for an integer $q \geq 2$ and a polynomial f that parametrizes the problem. Equivalently, it may be described as SIS where the uniform matrix is replaced by a structured matrix (the precise structure depends on f). PSIS allows the design of fast digital signatures, among other applications (see [Lyu09], for example).

This approach was extended to LWE by Stehlé *et al.* [SSTX09], who introduced and studied the (search version of) Polynomial-LWE problem (PLWE).[5] Lyubashevsky *et al.* [LPR13] introduced the Ring-LWE problem, which involves number fields rather than polynomials, and proposed a reduction from its search to decision versions, in the case of cyclotomic polynomials. (See also [EHL14,CLS15] for extensions to larger classes of fields of the Ring-LWE search to decision reduction.) Power-of-2 cyclotomic polynomials (for which PLWE and Ring-LWE match) have been exploited to design fast encryption schemes, among others (see [ADPS16], for example). Cryptographic schemes based on PLWE/Ring-LWE most often enjoy keys of $\widetilde{O}(\lambda)$ bit-sizes and algorithms with $\widetilde{O}(\lambda)$ runtime, where λ refers to the security parameter (i.e., all known attacks run in time $\geq 2^\lambda$) and the $\widetilde{O}(\cdot)$ notation hides poly-logarithmic factors.

Switching from unstructured SIS and LWE to their polynomial counterparts PSIS and PLWE has undeniable efficiency advantages. However, the security guarantees are severely degraded. PSIS and PLWE also enjoy reductions from worst-case lattice problems such as ApproxSVP, but these lattice problems, e.g., $\mathsf{ApproxSVP}^{(f)}$, are restricted to lattices that correspond to ideals of $\mathbb{Z}[x]/f$, where f is the polynomial that parametrizes PSIS and PLWE: under some conditions on f, there exists a reduction from $\mathsf{ApproxSVP}^{(f)}$ with small approximation factor, to $\mathsf{PSIS}^{(f)}$ and $\mathsf{PLWE}^{(f)}$ (see [LM06,PR06,SSTX09]). It is entirely possible that $\mathsf{PSIS}^{(f)}/\mathsf{PLWE}^{(f)}$ could be easy to solve for some polynomials f, and hard for others.[6] For instance, if f has a linear factor over the integers, then it is well-known that $\mathsf{PSIS}^{(f)}/\mathsf{PLWE}^{(f)}$ are computationally easy (we note that the reductions from $\mathsf{ApproxSVP}^{(f)}$ require f to be irreducible). Finding weak f's

terms of number fields but polynomial rings (as opposed to Ring-LWE), similarly to the Polynomial-LWE problem (PLWE) we consider in this work. It is possible to define a SIS variant of Ring-LWE, i.e., involving number fields: in the common case of power-of-2 cyclotomics, PSIS and Ring-SIS match (as do PLWE and Ring-LWE). In this work, we are interested in larger classes of polynomials, making the distinction important.

[5] It was originally called Ideal-LWE, by analogy to Ideal-SIS.

[6] We note that the stability of the polynomial rings under multiplication by x can be exploited to accelerate some known lattice algorithms by small polynomial factors, but we are interested here in more drastic weaknesses.

for PLWE has been investigated in [EHL14,ELOS15,CLS15,CLS16]. The attacks presented in this sequence of articles were used to identify such f's, but they only work for error distributions with small width relative to the geometry of the corresponding ring [CIV16b,CIV16a,Pei16]. In another sequence of works, Cramer *et al.* [CDPR16,CDW17] showed that $\mathsf{ApproxSVP}^{(f)}$ is easier for f a cyclotomic polynomial of prime-power conductor than for general lattices. More concretely, the authors of [CDW17] give a quantum polynomial-time algorithm for $\mathsf{ApproxSVP}^{(f)}$ with approximation factor $2^{\widetilde{O}(\sqrt{n})}$, where n is the degree of f. As a comparison, for such approximation factors and arbitrary lattices, the best known algorithms run in time $2^{\widetilde{O}(\sqrt{n})}$ (see [Sch87]). Finally, we note that the choice of non-cyclotomic polynomials in [BCLvV16] was motivated by such weaknesses. Even though the results in [CDPR16,CDW17] impact $\mathsf{ApproxSVP}^{(f)}$, it may be argued that it could have implications for $\mathsf{PLWE}^{(f)}$ as well, possibly even for lower approximation factors. On the other hand, it could be that similar weaknesses exist for $\mathsf{ApproxSVP}^{(f)}$ considered in [BCLvV16], although none is known at the moment. This lack of understanding of which f's correspond to hard $\mathsf{PLWE}^{(f)}$ problems motivates research into problems that are provably as hard as $\mathsf{PLWE}^{(f)}$ for the hardest f in a large class of polynomials, while preserving the computational efficiency advantages of PLWE. Our results are motivated by and make progress in this direction.

Recently, Lyubashevsky [Lyu16] introduced a variant $R^{<n}$-SIS of SIS that is not parametrized by a polynomial f and which enjoys the following desirable properties. First, an efficient algorithm for $R^{<n}$-SIS with degree bound n leads to an efficient algorithm for $\mathsf{PSIS}^{(f)}$ for all f's in a family of polynomials of size exponential in n. Second, there exists a signature scheme which is secure under the assumption that $R^{<n}$-SIS is hard, involves keys of bit-size $\widetilde{O}(\lambda) = \widetilde{O}(n)$ and whose algorithms run in time $\widetilde{O}(\lambda)$. In this sense, $R^{<n}$-SIS can serve as an alternative cryptographic foundation that hedges against the risk that $\mathsf{PSIS}^{(f)}$ is easy to solve for some f (as long as it stays hard for some f in the family).

Our contributions. Our main contribution is the introduction of an LWE counterpart to Lyubashevsky's $R^{<n}$-SIS problem. Let $n, q \geq 2$. We let $\mathbb{Z}_q^{<n}[x]$ denote the set of polynomials with coefficients in $\mathbb{Z}_q$ and degree $< n$. For $a \in \mathbb{Z}_q^{<n}[x]$ and $s \in \mathbb{Z}_q^{<2n-1}[x]$, we let $a \odot_n s = \lfloor (a \cdot s \bmod x^{2n-1})/x^{n-1} \rfloor \in \mathbb{Z}_q^{<n}[x]$ denote the polynomial obtained by multiplying a and s and keeping only the middle n coefficients. Middle-Product LWE (MP-LWE), with parameters $n, q \geq 2$ and $\alpha \in (0, 1)$, consists

in distinguishing arbitrarily many samples (a_i, b_i) uniform in $\mathbb{Z}_q^{<n}[x] \times (\mathbb{R}/q\mathbb{Z})^{<n}[x]$, from the same number of samples (a_i, b_i) with a_i uniform in $\mathbb{Z}_q^{<n}[x]$ and $b_i = a_i \odot_n s + e_i$, where each coefficient of e_i is sampled from the Gaussian distribution of standard deviation $\alpha \cdot q$, and s is uniformly chosen in $\mathbb{Z}_q^{<2n-1}[x]$.

We give a reduction from (decision) $\mathsf{PLWE}^{(f)}$ to (decision) MP-LWE of parameter n, for every monic f of degree n whose constant coefficient is coprime with q. The noise parameter amplifies linearly with the so-called Expansion Factor of f, introduced in [LM06]. The noise parameter in MP-LWE can for example be set to handle all monic polynomials $f = x^n + g$ with constant coefficient coprime with q, $\deg g \leq n/2$ and $\|g\| \leq n^c$ for an arbitrary $c > 0$. For any c, this set of f's has exponential size in n. We note that similar restrictions involving the expansion factor appeared before in [LM06,SSTX09].

Finally, we describe a public-key encryption scheme that is IND-CPA secure under the MP-LWE hardness assumption, involves keys of bit-size $\widetilde{O}(\lambda)$ and whose algorithms run in time $\widetilde{O}(\lambda)$. The scheme is adapted from Regev's [Reg09]. Its correctness proof involves an associativity property of the middle product. To establish its security, we prove that a related hash function family involving middle products is universal, and apply a generalized version of the leftover hash lemma. The standard leftover hash lemma does not seem to suffice for our needs, as the first part of the ciphertext is not statistically close to uniform, contrarily to Regev's encryption scheme.

Open problems. Our reduction is from the decision version of $\mathsf{PLWE}^{(f)}$ to the decision version of MP-LWE. (It can be adapted to the search counterparts, but it is unclear how to use the hardness of search MP-LWE for cryptographic purposes.) Unfortunately, the hardness of decision $\mathsf{PLWE}^{(f)}$ is currently supported by the presumed hardness of $\mathsf{ApproxSVP}^{(f)}$ for very few polynomials f. Such reductions for larger classes of polynomials f would strengthen our confidence in the hardness of MP-LWE. A first strategy towards this goal would be to design a reduction from search $\mathsf{PLWE}^{(f)}$ to decision $\mathsf{PLWE}^{(f)}$ for larger classes of f's than currently handled (the reduction from [LPR13] requires f to be cyclotomic). This reduction could then be combined with the one from $\mathsf{ApproxSVP}^{(f)}$ to $\mathsf{PLWE}^{(f)}$ from [SSTX09], which only requires f to be irreducible with bounded expansion factor. A second strategy would be to reduce decision $\mathsf{Ring\text{-}LWE}^{(f)}$ to decision $\mathsf{PLWE}^{(f)}$ and rely on the new reduction from ApproxSVP restricted to ideals of the number field K_f to deci-

sion Ring-LWE$^{(f)}$ from [PRSD17]. Indeed, this new reduction is not restricted to cyclotomic polynomials.

We show the cryptographic relevance of MP-LWE by adapting Regev's encryption scheme to the middle-product algebraic setting. Adapting the dual-Regev scheme from [GPV08] does not seem straightforward. Indeed, it appears that we would need a leftover hash lemma for polynomials over $\mathbb{Z}_q[x]$ that are not folded modulo some polynomial f. The difficulty is that the constant coefficients of the polynomials are now "isolated", in the sense that the constant coefficient of a polynomial combination of polynomials only involves the constant coefficients of these polynomials. Hopefully, solving this difficulty would also enable the construction of a trapdoor for MP-LWE, similar to those that exist for LWE and SIS (see [MP12] and references therein). Independently, showing that the MP-LWE secret could be sampled from a small-norm distribution, as achieved for LWE in [ACPS09], may allow for a more efficient ElGamal-type encryption, similar to the one described in [LPR13].

Notations. We use the notation $U(X)$ for the uniform distribution over the set X. If D_1 and D_2 are two distributions over the same countable domain, we let $\Delta(D_1, D_2)$ denote their statistical distance. We let $\|\mathbf{b}\|$ and $\|\mathbf{b}\|_\infty$ denote the Euclidean and infinity norm of any vector $\mathbf{b}$ over the reals, respectively. Similarly, if b is a polynomial over the reals, we let $\|b\|$ denote the Euclidean norm of its coefficient vector. For a matrix $\mathbf{M}$ we let $\mathbf{M}_{i,j}$ denote its element in the i-th row and j-th column. We let $\|\mathbf{M}\|$ denote the largest singular value of $\mathbf{M}$.

2 Background

In this section, we provide the background definitions and results that are necessary to present our contributions.

2.1 Probabilities

We will use the following variant of the leftover hash lemma. We recall that a (finite) family $\mathcal{H}$ of hash functions $h : X \to Y$ is universal if $\Pr_{h \hookleftarrow U(\mathcal{H})}[h(x_1) = h(x_2)] = 1/|Y|$, for all $x_1 \neq x_2 \in X$.

Lemma 2.1. *Let X, Y, Z denote finite sets. Let $\mathcal{H}$ be a universal family of hash functions $h : X \to Y$. Let $f : X \to Z$ be arbitrary. Then for any random variable T taking values in X, we have:*

$$\Delta(\ (h, h(T), f(T))\ ,\ (h, U(Y), f(T))\ )\ \leq\ \frac{1}{2} \cdot \sqrt{\gamma(T) \cdot |Y| \cdot |Z|},$$

where $\gamma(T) = \max_{t \in X} \Pr[T = t]$.

In the problems we will study, the so-called noise distributions will be Gaussian.

Definition 2.1 *We define the Gaussian function on* $\mathbb{R}^n$ *of covariance matrix* $\mathbf{\Sigma}$ *as* $\rho_{\mathbf{\Sigma}}(\mathbf{x}) := \exp(-\pi \cdot \mathbf{x}^T \mathbf{\Sigma}^{-1} \mathbf{x})$ *for every vector* $\mathbf{x} \in \mathbb{R}^n$*. The probability distribution whose density is proportional to* $\rho_{\mathbf{\Sigma}}$ *is called the Gaussian distribution and is denoted* $D_{\mathbf{\Sigma}}$*. When* $\mathbf{\Sigma} = s^2 \cdot \mathsf{Id}_n$*, we write* ρ_s *and* D_s *instead of* $\rho_{\mathbf{\Sigma}}$ *and* $D_{\mathbf{\Sigma}}$*, respectively.*

2.2 Polynomials and Structured Matrices

Let R be a ring. For $k > 0$, we let $R^{<k}[x]$ denote the set of polynomials in $R[x]$ of degree $< k$. Given a polynomial $a = a_0 + a_1 x + \cdots + a_{k-1} x^{k-1} \in R^{<k}[x]$ and some $j < k$, we use the following notations: $\mathbf{a} = (a_0, \ldots, a_{k-1})^T \in R^k$ and $\overline{\mathbf{a}} = (a_{k-1}, \ldots, a_0)^T \in R^k$. The latter notation is extended to the corresponding polynomial.

Definition 2.2. *Let* f *be a polynomial of degree* m*. For any* $d > 0$ *and any* $a \in R[x]$*, we let* $\mathsf{Rot}_f^d(a)$ *denote the matrix in* $R^{d \times m}$ *whose* i*-th row is given by the coefficients of the polynomial* $(x^{i-1} \cdot a) \bmod f$*, for any* $i = 1, \ldots, d$*. We will use the notation* $\mathsf{Rot}_f(a)$ *instead of* $\mathsf{Rot}_f^m(a)$*.*

Note that if $a' = a \bmod f$, then $\mathsf{Rot}_f^d(a) = \mathsf{Rot}_f^d(a')$ for any d. Note also that $\mathsf{Rot}_f(a \cdot b) = \mathsf{Rot}_f(a) \cdot \mathsf{Rot}_f(b)$ for any $a, b \in R[x]$.

Definition 2.3. *Let* f *be a polynomial of degree* m*. We define* $\mathbf{M}_f$ *as the (Hankel) matrix in* $R^{m \times m}$ *such that for any* $1 \leq i, j \leq m$*, the coefficient* $(\mathbf{M}_f)_{i,j}$ *is the constant coefficient of* $x^{i+j-2} \bmod f$*.*

Matrix $\mathbf{M}_f$ helps rewriting multiplication on the left by matrix $\mathsf{Rot}_f(a)$ as a multiplication on the right by $\mathbf{a}$.

Lemma 2.4. *For any* $a \in R^{<m}[x]$*, we have* $\mathsf{Rot}_f(a) \cdot (1, 0, \ldots, 0)^T = \mathbf{M}_f \cdot \mathbf{a}$*.*

Proof. First, the i-th coordinate of the left hand side is the constant coefficient of $x^{i-1} \cdot a \bmod f$. Second, the i-th entry of the right hand side is

$$((a_0 x^{i-1} \bmod f) \bmod x) + \cdots + ((a_{m-1} x^{m+i-2} \bmod f) \bmod x),$$

which can be re-written as $x^{i-1}(a_0 + \cdots + a_{m-1} x^{m-1} \bmod f) \bmod x = (x^{i-1} \cdot a \bmod f) \bmod x$. The latter is the constant coefficient of $x^{i-1} \cdot a \bmod f$. □

Definition 2.5. *For any $d, k > 0$ and $a \in R^{<k}[x]$, we let $\mathsf{Toep}^{d,k}(a)$ denote the matrix in $R^{d\times(k+d-1)}$ whose i-th row, for $i = 1, \ldots, d$, is given by the coefficients of $x^{i-1} \cdot a$.*

The following property will be useful in proving our main result.

Lemma 2.6. *For any $d, k > 0$ and any $a \in R^{<k}[x]$, we have $\mathsf{Rot}_f^d(a) = \mathsf{Toep}^{d,k}(a) \cdot \mathsf{Rot}_f^{k+d-1}(1)$.*

Proof. It is sufficient to prove that the rows of $\mathsf{Rot}_f^d(a)$ and $\mathsf{Toep}^{d,k}(a) \cdot \mathsf{Rot}_f^{k+d-1}(1)$ are equal. We just note that the i-th row of $\mathsf{Rot}_f^{k+d-1}(1)$ is $x^{i-1} \bmod f$, for $i = 1, \ldots, k+d$ and these will fill the gap in the definitions of $\mathsf{Rot}_f^d(a)$ and $\mathsf{Toep}^{d,k}(a)$. □

We now recall the definition of the expansion factor [LM06].

Definition 2.7. *Let $f \in \mathbb{Z}[x]$ of degree m. Then the expansion factor of f is defined as $\mathsf{EF}(f) = \max(\|g \bmod f\|_\infty / \|g\|_\infty : g \in \mathbb{Z}^{<2m-1}[x] \setminus \{0\})$.*

We remark that there are numerous polynomials with bounded expansion factor. One class of such polynomials [LM06] is the family of all $f = x^m + h$, for $h = \sum_{i \le m/2} h_i x^i$ and $\|\mathbf{h}\|_\infty \in \mathrm{poly}(m)$: we then have $\mathsf{EF}(f) \in \mathrm{poly}(m)$.

Lemma 2.8. *For $f \in \mathbb{Z}[x]$, we have $\|\mathbf{M}_f\| \le \deg(f) \cdot \mathsf{EF}(f)$.*

Proof. By definitions of $\mathbf{M}_f$ and $\mathsf{EF}(f)$, we have that $|(\mathbf{M}_f)_{i,j}| \le \mathsf{EF}(f)$, for $1 \le i, j \le m$. Therefore, the largest singular value of $\mathbf{M}_f$ is bounded from above by $m \cdot \mathsf{EF}(f)$. □

2.3 The Polynomial Learning With Errors Problem (PLWE)

We first define the distribution the PLWE problem is based on. For the rest of this paper, we will use the notation $\mathbb{R}_q := \mathbb{R}/q\mathbb{Z}$.

Definition 2.9 (P distribution). *Let $q \ge 2$, $m > 0$, f a polynomial of degree m, χ a distribution over $\mathbb{R}[x]/f$. Given $s \in \mathbb{Z}_q[x]/f$, we define the distribution $\mathsf{P}_{q,\chi}^{(f)}(s)$ over $\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ obtained by sampling $a \hookleftarrow U(\mathbb{Z}_q[x]/f)$, $e \hookleftarrow \chi$ and returning $(a, b = a \cdot s + e)$.*

Definition 2.10 (PLWE). *Let $q \ge 2$, $m > 0$, f a polynomial of degree m, χ a distribution over $\mathbb{R}[x]/f$. The (decision) $\mathsf{PLWE}_{q,\chi}^{(f)}$ consists in distinguishing between arbitrarily many samples from $\mathsf{P}_{q,\chi}^{(f)}(s)$ and the same number of samples from $U(\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f)$, with non-negligible probability over the choices of $s \hookleftarrow U(\mathbb{Z}_q[x]/f)$.*

One can also define a search variant of $\mathsf{PLWE}_{q,\chi}^{(f)}$, which would consist in computing $s \in \mathbb{Z}_q[x]/f$ from arbitrarily many samples from $\mathsf{P}_{q,\chi}^{(f)}(s)$.

3 The Middle-Product Learning With Errors Problem

We first recall the definition of the middle product of two polynomials and some of its properties.

3.1 The Middle-Product

Let R be a ring. Assume we multiply two polynomials a and b of degrees $< d_a$ and $< d_b$, respectively. Assume that $d_a + d_b - 1 = d + 2k$ for some integers d and k. Then the middle-product of size d of a and b is obtained by multiplying a and b, deleting the (left) coefficients of $1, x, \ldots, x^{k-1}$, deleting the (right) coefficients of $x^{k+d}, x^{k+d+1}, \ldots, x^{d+2k-1}$, and dividing what remains (the middle) by x^k.

Definition 3.1. *Let d_a, d_b, d, k be integers such that $d_a + d_b - 1 = d + 2k$. The middle-product $\odot_d : R^{<d_a}[x] \times R^{<d_b}[x] \to R^{<d}[x]$ is the map:*

$$(a, b) \mapsto a \odot_d b = \left\lfloor \frac{(a \cdot b) \bmod x^{k+d}}{x^k} \right\rfloor .$$

We use the same notation $\odot_d$ for every d_a, d_b such that $d_a + d_b - 1 - d$ is non-negative and even.

The middle-product of polynomials is used in computer algebra to accelerate computations in polynomial rings (see, e.g., [Sho99,HQZ04]). As it is part of the output of polynomial multiplication, it can be computed with a number of ring additions and multiplications that is quasi-linear number in $d_a + d_b$. Faster algorithms exist [HQZ04].

The (reversed) coefficient vector of the middle-product of two polynomials is in fact equal to the product of the Toeplitz matrix associated to one polynomial by the (reversed) coefficient vector of the second polynomial.

Lemma 3.2. *Let $d, k > 0$. Let $r \in R^{<k+1}[x]$ and $a \in R^{<k+d}[x]$ and $b = r \odot_d a$. Then $\overline{\mathbf{b}} = \mathsf{Toep}^{d,k+1}(r) \cdot \overline{\mathbf{a}}$. In other words, we have $\mathbf{b} = \overline{\mathsf{Toep}^{d,k+1}(r) \cdot \overline{\mathbf{a}}}$.*

Proof. We first note that $\mathsf{Toep}^{d,2k+d}(r \cdot a) = \mathsf{Toep}^{d,k+1}(r) \cdot \mathsf{Toep}^{k+d,k+d}(a)$. Thus, by definition of the middle-product, we have that the coefficients of b appear in the first row of $\mathsf{Toep}(r \cdot a)$, namely $b_i = \mathsf{Toep}^{d,2k+d}(r \cdot a)_{1,k+i+1}$ for $i < d$. But since $\mathsf{Toep}(r \cdot a)$ is constant along its diagonals, we also have that b appear (in reversed order) in the $(k+d)$-th column of $\mathsf{Toep}^{d,2k+d}(r \cdot a)$, namely $b_i = \mathsf{Toep}^{d,2k+d}(r \cdot a)_{d-i,k+d}$ for $i < d$. Therefore, vector $\overline{\mathbf{b}}$ is the $(k+d)$-th column of $\mathsf{Toep}^{d,2k+d}(r \cdot a)$, which is equal to $\mathsf{Toep}^{d,k+1}(r) \cdot \mathbf{a}'$, where $\mathbf{a}'$ is the $(k+d)$-th column of $\mathsf{Toep}^{k+d,k+d}(a)$. Since $\mathsf{Toep}^{k+d,k+d}(a)$ is constant along its diagonals, its first row is equal to its reversed $(k+d)$-th column, so $\mathbf{a}' = \overline{\mathbf{a}}$, as required. □

The middle-product is an additive homomorphism when either of its inputs is fixed. As a consequence of the associativity of matrix multiplication and Lemma 3.2, the middle-product satisfies the following associativity property.

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

Proof. Note first that the degree bounds match. Now, by Lemma 3.2, the vector associated to the reverse of $r \odot_d (a \odot_{d+k} s)$ is $\mathsf{Toep}^{d,k+1}(r) \cdot (\mathsf{Toep}^{d+k,n}(a) \cdot \overline{\mathbf{s}})$. Similarly, the vector associated to the reverse of $(r \cdot a) \odot_d s$ is $\mathsf{Toep}^{d,k+n}(r \cdot a) \cdot \overline{\mathbf{s}}$. The result follows from observing that $\mathsf{Toep}^{d,k+1}(r) \cdot \mathsf{Toep}^{d+k,n}(a) = \mathsf{Toep}^{d,k+n}(r \cdot a)$. □

3.2 Middle-Product Learning With Errors

Before stating MP-LWE, we first introduce a distribution its definition relies on.

Definition 3.4 (MP distribution). *Let $n, d > 0$, $q \geq 2$, and χ a distribution over $\mathbb{R}^{<d}[x]$. For $s \in \mathbb{Z}_q^{<n+d-1}[x]$, we define the distribution $\mathsf{MP}_{q,n,d,\chi}(s)$ over $\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$ as the one obtained by: sampling $a \hookleftarrow U(\mathbb{Z}_q^{<n}[x])$, $e \hookleftarrow \chi$ and returning $(a, b = a \odot_d s + e)$.*

Definition 3.5 (MP-LWE). *Let $n, d > 0$, $q \geq 2$, and a distribution χ over $\mathbb{R}^{<d}[x]$. The (decision) $\mathsf{MP\text{-}LWE}_{n,d,q,\chi}$ consists in distinguishing between arbitrarily many samples from $\mathsf{MP}_{q,n,d,\chi}(s)$ and the same number of samples from $U(\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x])$, with non-negligible probability over the choices of $s \hookleftarrow U(\mathbb{Z}_q^{<n+d-1}[x])$.*

It is possible to define a search variant of $\mathsf{MP\text{-}LWE}_{q,n,d,\chi}$, which would consist in computing $s \in \mathbb{Z}_q^{<n+d-1}[x]$ from arbitrarily many samples from $\mathsf{MP}_{q,n,d,\chi}(s)$.

Note that $\mathsf{MP\text{-}LWE}_{q,n,d,\chi}$ can also be viewed as a variant of LWE, in which the samples are correlated. Thanks to Lemma 3.2, it can indeed be restated as follows. Given many samples $(\mathsf{Toep}^{d,n}(a_i), \overline{\mathbf{b}}_i) \in \mathbb{Z}_q^{d\times(n+d-1)} \times \mathbb{R}_q^d$ for uniformly chosen $a_i \in \mathbb{Z}_q^{<n}[x]$, decide if the vectors $\overline{\mathbf{b}}_i$ are uniformly sampled in $\mathbb{R}_q^d$ or are of the form $\overline{\mathbf{b}}_i = \mathsf{Toep}^{d,n}(a_i)\cdot\overline{\mathbf{s}}+\overline{\mathbf{e}}_i$ for some common $s \hookleftarrow U(\mathbb{Z}_q^{<n+d-1}[x])$ and $e_i \hookleftarrow \chi$.

3.3 Hardness of MP-LWE

The following reduction from PLWE to $\mathsf{MP\text{-}LWE}$ is our main result.

Theorem 3.6. *Let $n, d > 0$, $q \geq 2$, and $\alpha \in (0,1)$. For $S > 0$, we let $\mathcal{F}(S,d,n)$ denote the set of polynomials $f \in \mathbb{Z}[x]$ that are monic, have constant coefficient coprime with q, have degree m in $[d,n]$ and that satisfy $\mathsf{EF}(f) < S$. Then there exists a ppt reduction from $\mathsf{PLWE}^{(f)}_{q,D_{\alpha\cdot q}}$ for any $f \in \mathcal{F}(S,d,n)$ to $\mathsf{MP\text{-}LWE}_{q,n,d,D_{\alpha'\cdot q}}$ with $\alpha' = \alpha d S$.*

Proof. We first reduce $\mathsf{PLWE}^{(f)}$ to a variant of $\mathsf{MP\text{-}LWE}$ whose only dependence on f lies in the noise distribution (see Lemma 3.7 below). Then we remove the latter dependence, by adding a compensating Gaussian distribution (see Lemma 3.8 below). The bound on the magnitude of matrix $\mathbf{M}_f$ from Lemma 2.8 for $\chi = D_{\alpha\cdot q}$ implies that

$$\|\mathbf{\Sigma}_0\| = \alpha q \|\mathbf{J}\cdot\mathbf{M}_f^d\| = \alpha q \|\mathbf{M}_f^d\| \leq \alpha q d \mathsf{EF}(f) < \alpha q d S.$$

Hence, taking $\alpha' q = \alpha q d S$ completes the proof. □

Lemma 3.7. *Let $n, d > 0$, $q \geq 2$, and χ a distribution over $\mathbb{R}^{<d}[x]$. Then there exists a ppt reduction from $\mathsf{PLWE}^{(f)}_{q,\chi}$ for any monic $f \in \mathbb{Z}[x]$ with constant coefficient coprime with q and degree $m \in [d,n]$, to $\mathsf{MP\text{-}LWE}_{q,n,d,\mathbf{J}\cdot\mathbf{M}_f^d\cdot\chi}$. Here, matrix $\mathbf{M}_f^d$ is the one obtained by keeping only the first d rows of $\mathbf{M}_f$, and $\mathbf{J} \in \mathbb{Z}^{d\times d}$ is the one with 1's on the anti-diagonal and 0's everywhere else.*

Proof. We describe below an efficient randomized mapping ϕ that takes as input a pair $(a_i, b_i) \in \mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ and maps it to a pair $(a_i', b_i') \in \mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x]$, such that ϕ maps $U(\mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f)$ to $U(\mathbb{Z}_q^{<n}[x] \times \mathbb{R}_q^{<d}[x])$ and $\mathsf{P}^{(f)}_{q,\chi}(s)$ to $\mathsf{MP}_{q,n,d,\chi'}(s')$, for some s' that depends on s and some χ' that depends on χ and f.

The reduction is then as follows:

- Sample $t \hookleftarrow U(\mathbb{Z}_q^{<n+d-1}[x])$.
- Each time the $\mathsf{MP\text{-}LWE}$ oracle requests a new sample, ask for a fresh PLWE sample (a_i, b_i), compute $(a'_i, b'_i) = \phi(a_i, b_i)$ and give $(a'_i, b'_i) + (0, a'_i \odot_d t)$ to the $\mathsf{MP\text{-}LWE}$ oracle.
- When $\mathsf{MP\text{-}LWE}$ terminates, return its output.

Assuming ϕ satisfies the specifications above, the reduction maps uniform samples to uniform samples, and $\mathsf{P}_{q,\chi}^{(f)}(s)$ samples for a uniform s that is common to all samples to $\mathsf{MP}_{q,n,d,\mathbf{J}\cdot\mathbf{M}_f^d\cdot\chi}(s'+t)$ samples for a uniform $s'+t$ that is common to all samples.

We now describe ϕ. Let $(a_i, b_i) \in \mathbb{Z}_q[x]/f \times \mathbb{R}_q[x]/f$ be an input pair. Let m denote the degree of f. We sample $r_i \hookleftarrow U(\mathbb{Z}_q^{<n-m}[x])$ and set $\phi(a_i, b_i) = (a'_i, b'_i)$ with:

$$a'_i = a_i + f \cdot r_i \in \mathbb{Z}_q^{<n}[x] \ , \quad \overline{\mathbf{b}'_i} = \mathbf{M}_f^d \cdot \mathbf{b}_i \in \mathbb{R}_q^{<d}[x].$$

As a_i and r_i are uniformly distributed in $\mathbb{Z}_q^{<m}[x]$ and $\mathbb{Z}_q^{<n-m}[x]$ respectively, the polynomial a'_i is uniformly distributed in $\mathbb{Z}_q^{<n}[x]$ (we refer to [Lyu16, Lemma 2.10] for a fully detailed proof). Here, we use the assumption that f is monic.

Further, if b_i is uniformly distributed, then so is its coefficient vector $\mathbf{b}_i$, and so is $\mathbf{M}_f^d \cdot \mathbf{b}_i$. Indeed, as the constant coefficient is coprime with q, matrix $\mathbf{M}_f$ is invertible modulo q (reordering its columns makes it triangular, with diagonal coefficients all equal to the constant coefficient of f).

Now, assume that $b_i = a_i \cdot s + e_i$, for some $s \in \mathbb{Z}_q[x]/f$ and $e_i \hookleftarrow \chi$. Thanks to Subsection 2.2, we know that $\mathsf{Rot}_f(b_i) = \mathsf{Rot}_f(a_i) \cdot \mathsf{Rot}_f(s) + \mathsf{Rot}_f(e_i)$, and, by taking the first columns and d first rows, we have

$$\begin{aligned}
\mathbf{M}_f^d \cdot \mathbf{b}_i &= \mathsf{Rot}_f^d(a_i) \cdot \mathbf{M}_f \cdot \mathbf{s} + \mathbf{M}_f^d \cdot \mathbf{e}_i \\
&= \mathsf{Rot}_f^d(a'_i) \cdot \mathbf{M}_f \cdot \mathbf{s} + \mathbf{M}_f^d \cdot \mathbf{e}_i \\
&= \mathsf{Toep}^{d,n}(a'_i) \cdot \mathsf{Rot}_f^{d+n-1}(1) \cdot \mathbf{M}_f \cdot \mathbf{s} + \mathbf{M}_f^d \cdot \mathbf{e}_i \\
&= \mathsf{Toep}^{d,n}(a'_i) \cdot \overline{\mathbf{s}'} + \mathbf{M}_f^d \cdot \mathbf{e}_i,
\end{aligned}$$

where $\mathbf{s}' = \overline{\mathsf{Rot}_f^{d+n-1}(1) \cdot \mathbf{M}_f \cdot \mathbf{s}}$. Since $\mathbf{b}'_i = \overline{\mathbf{M}_f^d \cdot \mathbf{b}_i} = \overline{\mathsf{Toep}(a'_i) \cdot \overline{\mathbf{s}'}} + \overline{\mathbf{M}_f^d \cdot \mathbf{e}_i}$, we get that $\mathbf{e}'_i = \overline{\mathbf{M}_f^d \cdot \mathbf{e}_i}$, which makes the distribution in $\mathsf{MP\text{-}LWE}$ equals to the claimed $\mathbf{J} \cdot \mathbf{M}_f^d \cdot \chi$. This completes the proof. □

We now remove the dependence in f of the noise distribution.

Lemma 3.8. *Let $n, d > 0$, $q \geq 2$. Let $\sigma' > 0$. Let $\mathbf{\Sigma}_0 \in \mathbb{R}^{d \times d}$ be symmetric definite positive matrix with $\|\mathbf{\Sigma}_0\| < \sigma'$. Then there exists a* ppt *reduction from* $\mathsf{MP\text{-}LWE}_{q,n,d,D_{\mathbf{\Sigma}_0}}$ *to* $\mathsf{MP\text{-}LWE}_{q,n,d,D_{\sigma' \cdot \mathsf{Id}_d}}$*, where* Id_d *denotes the d-dimensional identity matrix.*

Proof. The reduction is as follows. We first note that, there exists a positive definite matrix $\mathbf{\Sigma}'$, such that $\mathbf{\Sigma}_0 + \mathbf{\Sigma}' = \sigma' \cdot \mathsf{Id}_d$. The positive definiteness is guaranteed by fact that $\|\mathbf{\Sigma}_0\| < \sigma'$. Then, for any $\mathsf{MP\text{-}LWE}_{q,n,d,D_{\mathbf{\Sigma}_0}}$ input sample (a_i, b_i), we sample $e'_i \hookleftarrow D_{\mathbf{\Sigma}'}$ and compute $(a'_i, b'_i) = (a_i, b_i + e'_i)$.

Observe that the reduction maps uniform samples to uniform samples, and $\mathsf{MP}_{q,n,d,D_{\mathbf{\Sigma}_0}}(s)$ samples to $\mathsf{MP}_{q,n,d,D_{\sigma' \cdot \mathsf{Id}_d}}(s)$ samples. This completes the proof. □

4 Public-Key Encryption from MP-LWE

We now describe a public key encryption scheme that is IND-CPA secure, under the MP-LWE hardness assumption. The scheme is an adaptation of Regev's from [Reg09]. It relies on parameters $q, n, d, t \geq 2$ with q odd, and a noise rate $\alpha \in (0, 1)$. We let $\chi = \lfloor D_{\alpha q} \rceil$ denote the distribution over $\mathbb{Z}^{<d+k}[x]$ where each coefficient is sampled from $D_{\alpha \cdot q}$ and then rounded to nearest integer. The plaintext space is $\{0,1\}^{<d}[x]$, while the ciphertext space is $\mathbb{Z}_q^{<k+n}[x] \times \mathbb{Z}_q^{<d}[x]$.

$\mathsf{KeyGen}(1^\lambda)$**.** Sample $s \hookleftarrow U(\mathbb{Z}_q^{<n+d+k-1}[x])$. For every $i \leq t$, sample $a_i \hookleftarrow U(\mathbb{Z}_q^{<n}[x])$, $e_i \hookleftarrow \chi$ and compute $b_i = a_i \odot_{d+k} s + 2 \cdot e_i \in \mathbb{Z}_q^{<d+k}[x]$. Return the secret key $\mathsf{sk} := s$ and the public key $pk := (a_i, b_i)_{i \leq t}$.

$\mathsf{Encrypt}(\mathsf{pk} = (a_i, b_i)_{i \leq t}, \mu)$**.** For $i \leq t$, sample $r_i \hookleftarrow U(\{0,1\}^{<k+1}[x])$, and return $c = (c_1, c_2)$ with:

$$c_1 = \sum_{i \leq t} r_i \cdot a_i \ , \quad c_2 = \mu + \sum_{i \leq t} r_i \odot_d b_i.$$

$\mathsf{Decrypt}(\mathsf{sk} = s, c)$**.** Return the plaintext $\mu' = (c_2 - c_1 \odot_d s \bmod q) \bmod 2$.

Example parameters are $n \geq \lambda$, $k = d = n/2$, $q = \Theta(n^{5/2+c}\sqrt{\log n})$, $t = \Theta(\log n)$ and $\alpha = \Theta(1/n\sqrt{\log n})$, for $c > 0$ arbitrary. For these parameters, the scheme is correct (by Lemma 4.1) and secure under $\mathsf{MP\text{-}LWE}_{q,n,n,D_{\alpha q}}$ (by Lemma 4.3). These parameters allow to rely on the assumed hardness of $\mathsf{PLWE}^{(f)}_{q,D_{\beta \cdot q}}$ via Theorem 3.6, for $\beta = \Omega(\sqrt{n}/q)$ (hence preventing attacks *à la* [AG11]) and for any f monic of degree n,

with constant coefficient coprime with q and expansion factor $\leq n^c$. Finally, note that the scheme encrypts and decrypts n plaintext bits in time $\widetilde{O}(n)$, and the key pair has bit-length $\widetilde{O}(n)$.

Correctness follows from Lemma 3.3 and the proof of correctness of Regev's encryption scheme.

Lemma 4.1 (Correctness). *Assume that $\alpha < 1/(16\sqrt{\lambda tk})$ and $q \geq 16t(k+1)$. With probability $\geq 1 - d \cdot 2^{-\Omega(\lambda)}$ over the randomness of $(\mathsf{sk}, \mathsf{pk}) \hookleftarrow \mathsf{KeyGen}$, for all plaintext μ and with probability 1 over the randomness of $\mathsf{Encrypt}$, we have $\mathsf{Decrypt}(\mathsf{sk}, \mathsf{Encrypt}(\mathsf{pk}, \mu)) = \mu$.*

Proof. Assume that (c_1, c_2) is an encryption of μ under pk. Then we have, modulo q:

$$\begin{aligned} c_2 - c_1 \odot_d s &= \mu + \sum_{i \leq t} r_i \odot_d b_i - (\sum_{i \leq t} r_i \cdot a_i) \odot_d s \\ &= \mu + \sum_{i \leq t} \left(r_i \odot_d (a_i \ \odot_{d+k} s + 2 \cdot e_i) - (r_i \cdot a_i) \odot_d s \right) \\ &= \mu + 2 \sum_{i \leq t} r_i \odot_d e_i, \end{aligned}$$

where the last equality follows from Lemma 3.3. If $\|\mu + 2 \cdot \sum_{i \leq t} r_i \odot_d e_i\|_\infty < q/2$, then centered reduction modulo q of $c_2 - c_1 \odot_d s$ gives us $\mu + 2 \cdot \sum_{i \leq t} r_i \odot_d e_i$ (over the integers). Reducing modulo 2 then provides μ.

Now, each coefficient of $\sum_{i \leq t} r_i \odot_d e_i$ can be viewed as an inner product between a binary vector of dimension $t(k+1)$ and a vector sampled from $\lfloor D_{\alpha q} \rceil^{t(k+1)}$. Each coefficient individually has magnitude $\leq \alpha q \sqrt{\lambda t(k+1)} + t(k+1)$ with probability $\geq 1 - 2^{-\Omega(\lambda)}$, because of the Gaussian tail bound and the triangle inequality. By the union bound and triangular inequality, we obtain that $\|\mu + 2 \cdot \sum_{i \leq t} r_i \odot_d e_i\|_\infty < 2\alpha q \sqrt{t\lambda(k+1)} + 2t(k+1) + 1$ with probability $\geq 1 - d \cdot 2^{-\Omega(\lambda)}$. □

The security proof is adapted from that of Regev's encryption scheme from [Reg09], with a subtlety in the application of the leftover hash lemma. In Regev's scheme, if the public key is replaced by uniformly random elements, then the leftover hash lemma guarantees that the joint distribution of the public key and the encryption of an arbitrary plaintext is within exponentially small statistical distance from uniform. This property does not hold in our case: indeed, if $a_1, \ldots, a_t$ all have constant coefficient equal to 0 (this event occurs with a probability $1/q^t$, which is not exponentially small for our parameters), then so does $\sum_i r_i a_i$. However, we can show that the second component c_2 of the ciphertext is sta-

tistically close to uniform, given the view of the first component c_1. This suffices, as the plaintext is embedded in the second ciphertext component.

We first prove that the hash function family coming into play in the security proof is universal.

Lemma 4.2. *Let $q, k, d \geq 2$. For $(b_i)_i \in (\mathbb{Z}_q^{<d+k}[x])^t$, we let $h_{(b_i)_i}$ denote the map that sends $(r_i)_{i\leq t} \in (\{0,1\}^{<k+1}[x])^t$ to $\sum_{i\leq t} r_i \odot_d b_i \in \mathbb{Z}_q^{<d}[x]$. Then the hash function family $(h_{(b_i)_i})_{(b_i)_i}$ is universal.*

Proof. Our aim is to show that for $r_1, \ldots, r_t$ not all 0, we have

$$\Pr_{(b_i)_i,(b'_i)_i}\left[\sum_{i\leq t} r_i \odot_d b_i = \sum_{i\leq t} r_i \odot_d b'_i\right] = q^{-d}.$$

W.l.o.g. we may assume that $r_1 \neq 0$. By linearity, it suffices to prove that for all $y \in \mathbb{Z}_q^{<d}[x]$,

$$\Pr_{b_1}\left[r_1 \odot_d b_1 = y\right] = q^{-d}.$$

Let j be minimal such that the coefficient in x^j of r_1 is non-zero (i.e., equal to 1 as r_1 is binary). Then the equation $r_1 \odot_d b_1 = y$ restricted to entries $j+1$ to $j+d$ is a triangular linear system in the coefficients of b_1 with diagonal coefficients equal to 1. The map $b_1 \mapsto r_1 \odot_d b_1$ restricted to these coefficients of b_1 is hence a bijection. This gives the equality above. □

Lemma 4.3 (Security). *Assume that $t \geq (2\cdot\lambda+(k+d+n)\cdot\log q)/(k+1)$. Then the scheme above is* IND-CPA *secure, under the* MP-LWE$_{q,n,d+k,D_{\alpha q}}$ *hardness assumption.*

Proof. The IND-CPA security experiment is as follows. The challenger $\mathcal{C}$ samples a bit $b \hookleftarrow \{0,1\}$ and $(\mathsf{sk}, \mathsf{pk}) \hookleftarrow \mathsf{KeyGen}(1^\lambda)$; it gives pk to adversary $\mathcal{A}$ who sends back two plaintexts $\mu_0 \neq \mu_1$; the challenger computes $c \hookleftarrow \mathsf{Encrypt}(\mathsf{pk}, \mu_b)$ and sends it to $\mathcal{A}$, who outputs a bit b'. The scheme is secure if no ppt adversary $\mathcal{A}$ outputs $b' = b$ more probability that is non-negligibly away from $1/2$.

Now, consider the variant of the experiment above, in which $\mathcal{C}$ does not run $(\mathsf{sk}, \mathsf{pk}) \hookleftarrow \mathsf{KeyGen}(1^\lambda)$ but instead samples $\mathsf{pk} = (a_i, b_i)_i$ uniformly. Under the MP-LWE hardness assumption, the probabilities that $\mathcal{A}$ outputs $b' = b$ in both experiments are negligibly close. The reduction from MP-LWE to distinguishing the first and second experiments consists in multiplying by 2 (which is co-prime to q) and rounding the real samples given by an MP-LWE oracle to the nearest integer modulo q. The latter

maps MP-LWE with real noise to MP-LWE with rounded real noise (and uniform MP-LWE over the reals modulo q to a uniform MP-LWE over the integers modulo q).

We consider a third experiment, in which $\mathcal{C}$ also samples $\mathsf{pk} = (a_i, b_i)_i$, and additionally does not compute $c \hookleftarrow \mathsf{Encrypt}(\mathsf{pk}, \mu_b)$ before sending it to $\mathcal{A}$, but instead computes $c = (c_1, c_2)$ as follows. For $i \leq t$, it samples $r_i \hookleftarrow U(\{0,1\}^{<k+1}[x])$, $u \hookleftarrow U(\mathbb{Z}_q^{<d}[x])$, and sets:

$$c_1 = \sum_{i \leq t} r_i \cdot a_i \ , \quad c_2 = u.$$

Note that in this game, the view of $\mathcal{A}$ is independent of b, and hence the probability that it outputs $b' = b$ is exactly $1/2$. We argue below that the distributions of $((a_i, b_i)_i, c_1, c_2)$ in this new experiment and the latter one are within exponentially small statistical distance. The combination of these two facts provides the result.

It remains to prove that

$$\Delta\Big(((a_i, b_i)_i, \sum_{i \leq t} r_i \cdot a_i, \sum_{i \leq t} r_i \odot_d b_i) \ , \quad ((a_i, b_i)_i, \sum_{i \leq t} r_i \cdot a_i, u)\Big) \leq 2^{-\lambda},$$

where the a_i's, b_i's, r_i's and u are uniformly sampled in $\mathbb{Z}_q^{<n}[x]$, $\mathbb{Z}_q^{<d+k}[x]$, $U(\{0,1\}^{<k+1}[x])$ and $\mathbb{Z}_q^{<d}[x]$, respectively. By Lemma 4.2, the hash function family $h_{(b_i)_i}$ is universal. Further, the quantity $\sum_{i \leq t} r_i \cdot a_i$ belongs to $\mathbb{Z}_q^{<k+n}$, of cardinality q^{k+n}. Hence, by the Generalized Leftover Hash Lemma (see Lemma 2.1), the statistical distance above is bounded from above by $(2^{-(k+1)\cdot t} \cdot q^{k+d+n})^{1/2}/2$. □

Acknowledgments. We thank Guillaume Hanrot for helpful discussions. This work has been supported in part by ERC Starting Grant ERC-2013-StG-335086-LATTAC, and the Australian Research Council Discovery Grant DP150100285.

References

[ACPS09] B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proc. of CRYPTO*, volume 5677 of *LNCS*, pages 595–618. Springer, 2009.

[ADPS16] E. Alkim, L. Ducas, T. Pöppelmann, and P. Schwabe. Post-quantum key exchange - A new hope. In *Proc. of USENIX*, pages 327–343, 2016.

[AG11] S. Arora and R. Ge. New algorithms for learning in presence of errors. In *Proc. of ICALP*, pages 403–415. Springer, 2011.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proc. of STOC*, pages 99–108. ACM, 1996.

[BCLvV16] D. J. Bernstein, C. Chuengsatiansup, T. Lange, and C. van Vredendaal. NTRU Prime. Cryptology ePrint Archive, 2016. `http://eprint.iacr.org/2016/461`.

[CDPR16] R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Proc. of EUROCRYPT*. Springer, 2016.

[CDW17] R. Cramer, L. Ducas, and B. Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In *Proc. of EUROCRYPT*. Springer, 2017.

[CIV16a] W. Castryck, I. Iliashenko, and F. Vercauteren. On the tightness of the error bound in Ring-LWE. *LMS J Comput Math*, 2016.

[CIV16b] W. Castryck, I. Iliashenko, and F. Vercauteren. Provably weak instances of Ring-LWE revisited. In *Proc. of EUROCRYPT*, pages 147–167. Springer, 2016.

[CLS15] H. Chen, K. Lauter, and K. E. Stange. Attacks on search RLWE. 2015. To appear in SIAM Journal on Applied Algebra and Geometry. Available at `http://eprint.iacr.org/2015/971`.

[CLS16] H. Chen, K. Lauter, and K. E. Stange. Vulnerable Galois RLWE families and improved attacks. In *Proc. of SAC*. Springer, 2016.

[EHL14] K. Eisenträger, S. Hallgren, and K. Lauter. Weak instances of PLWE. In *Proc. of SAC*. Springer, 2014.

[ELOS15] Y. Elias, K. E. Lauter, E. Ozman, and K. E. Stange. Provably weak instances of Ring-LWE. In *Proc. of CRYPTO*. Springer, 2015.

[GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.

[HPS98] J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: A ring-based public key cryptosystem. In *Proc. of ANTS*, pages 267–288. Springer, 1998.

[HQZ04] G. Hanrot, M. Quercia, and P. Zimmermann. The middle product algorithm I. *Appl. Algebra Engrg. Comm. Comput.*, 14(6):415–438, 2004.

[LM06] V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. of ICALP*, pages 144–155. Springer, 2006.

[LPR13] V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. *J. ACM*, 60(6):43, 2013.

[Lyu09] V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *Proc. of ASIACRYPT*, pages 598–616. Springer, 2009.

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

[Mic07] D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007.

[MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *Proc. of EUROCRYPT*, pages 700–718. Springer, 2012.

[Pei16] C. Peikert. How not to instantiate Ring-LWE. In *Proc. of SCN*, pages 411–430. Springer, 2016.

[PR06] C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proc. of TCC*, pages 145–166. Springer, 2006.

[PRSD17] C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of Ring-LWE for any ring and modulus. In *Proc. of STOC*. ACM, 2017.

[Reg09] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

[Sch87] C. P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53(2-3):201–224, 1987.

[Sho99] V. Shoup. Efficient computation of minimal polynomials in algebraic extensions of finite fields. In *Proc. of ISSAC*, pages 53–58. ACM, 1999.

[SSTX09] D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *Proc. of ASIACRYPT*, pages 617–635. Springer, 2009.