Revisiting Multivariate Ring Learning with Errors and its Applications on Lattice-based Cryptography

Alberto Pedrouzo-Ulloa[1], Juan Ramón Troncoso-Pastoriza[2], Nicolas Gama[3], Mariya Georgieva[3], and Fernando Pérez-González[1]

[1] University of Vigo
{apedrouzo,fperez}@gts.uvigo.es
[2] EPFL
juan.troncoso-pastoriza@epfl.ch
[3] Inpher
{nicolas,mariya}@inpher.io

Abstract. The "Multivariate Ring Learning with Errors" problem was presented as a generalization of Ring Learning with Errors (RLWE), introducing efficiency improvements with respect to the RLWE counterpart thanks to its multivariate structure. Nevertheless, the recent attack presented by Bootland *et al.* has some important consequences on the security of the multivariate RLWE problem with "non-coprime" modular functions; this attack transforms instances of m-RLWE with power-of-two cyclotomic modular functions of degree $n = \prod_i n_i$ into a set of RLWE samples with dimension $\max_i \{n_i\}$. This is especially devastating for low-degree modular functions (e.g., $\Phi_4(x) = 1 + x^2$). In this work, we revisit the security of multivariate RLWE and propose new alternative instantiations of the problem that avoid the attack while still preserving the advantages of the multivariate structure, especially when using low-degree modular functions. Additionally, we show how to parameterize these instances in a secure and practical way, therefore enabling constructions and strategies based on m-RLWE that bring notable space and time efficiency improvements over current RLWE-based constructions.

Keywords: Tensor of Number Fields, Lattice Cryptography, Homomorphic Encryption, Ring Learning with Errors, Multivariate Rings, Hardness Assumptions

1 Introduction

Lattices have become a very promising tool for the development and improvement of new cryptographic constructions, notably those belonging to the field of homomorphic encryption. Instead of directly working with lattice assumptions, it is frequent to deal with assumptions whose underlying security can be based on the hardness of lattice problems. Among them, the family of LWE (Learning

with Errors) [77, 78] has become the preferred one due to its versatility. Lyubashevsky *et al.* [60, 58] proposed a variant of LWE called Ring-LWE (or RLWE), whose hardness can be reduced from hardness problems over ideal lattices (instead of the general ones used in the LWE version). RLWE has proven to be more practical than LWE, as the underlying primitives can be usually more efficient; e.g., RLWE enables a notable reduction in the size of the public and secret keys in public key cryptosystems.

The ring structure of RLWE enables homomorphic cryptography with a ring homomorphism supporting both addition and multiplication of ciphertexts. Among the possible polynomial rings used for this purpose, the most practical ones are those where the modular function is a cyclotomic polynomial of the form $1 + z^n$, with n a power of two. They present two advantages: (a) they enable efficient implementations of polynomial operations through fast radix algorithms of the NTT (Number Theoretic Transforms) [1, 50], and (b) the polynomial operations over the used ring correspond to basic blocks in practical applications in Computer Vision and Signal Processing [67, 71, 66], comprising, among others, linear convolutions, filtering, and linear transforms.

Recently, a multivariate version of RLWE (m-RLWE) was proposed as a means to efficiently deal with encrypted multidimensional structures, such as videos or images [68, 66, 72]. In this scenario, the use of a tensorial decomposition in "coprime" cyclotomic rings (see [61, 58, 60]) is a priori not applicable, as these structures require that the modular functions have the same form (e.g., $1 + z^n$). This is the context in which m-RLWE [68] was originally introduced.

Additionally, current hot problems in (fully) homomorphic encryption involve the optimization of elementary polynomial operations through fast transforms and, especially, the search for optimal strategies to execute homomorphic slot manipulations and trade off storage and computation needs for relinearization operations. These are fundamental blocks in homomorphic processing and in the implementation of the bootstrapping (see [48, 21, 25, 31]) primitives enabling fully homomorphic encryption. As we will show, m-RLWE can bring significant efficiency improvements in all of them (see Sections 8 and 9).

The use of the tensor of lattices and/or adding a multivariate structure to the involved rings has been the subject matter of several previous works, but with very different targets. We briefly survey here the closest ones: (a) In [51], the authors applied the standard tensor product of lattices to improve the hardness factor of the SVP problem under different assumptions. (b) In [61], the authors define an isomorphism between some cyclotomic fields and a tensor product of cyclotomic fields when the order m in $\Phi_m(z)$ can be factored into several (different) prime powers. (c) The "tensor" representation also appears in the definition of the GLWE problem (also called Module-LWE [55]) which was originally introduced in [17, 15]. In fact, analogously to LWE versus RLWE, the introduced multivariate RLWE problem can be seen as a ring version of the GLWE problem, by means of adding for a second time a ring structure into the module. (d) Finally, the FHEW fully homomorphic encryption scheme features [11] a ring tensoring for a speed-up of the homomorphic accumulator, and also bivari-

ate rings are used as a means to enhance the efficiency of polynomial products inside the refreshing procedure in [63].

It can be shown that the m-RLWE problem can be reduced from worst case discrete Gaussian Sampling (equivalent to SIVP) over the tensor of rings (see [70]). Unfortunately, a recent work [12] shows an effective attack against m-RLWE when the univariate subrings share common roots, therefore considerably lowering the security of the underlying problem. Hence, our main contribution in this work is to redefine the m-RLWE problem and find secure instantiations that preserve the efficient results on multivariate RLWE [72], by basing their security on a subset of RLWE on general number fields (see the recent work by Peikert *et al.* [75], that generalizes the RLWE problem to any modulus and any ring over number fields).

We now informally introduce the definition of m-RLWE, the attack by Bootland *et al.* [12], and the rationale of our solution, all exemplified in the bivariate case.

Bivariate RLWE Let $K_{(T)} = K_x \bigotimes K_y$ be the tensor product of 2 cyclotomic number fields of dimensions $n_x = \phi(m_x)$ and $n_y = \phi(m_y)$, and $R = \mathbb{Z}[x,y]/(\Phi_{m_x}(x), \Phi_{m_y}(y))$ the tensor of their corresponding ring of integers.

We define a Bivariate Ring LWE sample (see Definition 2 for the general formulation of m-RLWE) as the pair $(a, b = (a \cdot s)/q + e \bmod R^\vee)$, where $a \leftarrow R_q$ is uniformly random and $e \leftarrow \Psi$ comes from the error distribution Ψ.

Bootland *et al.*'s attack Choices of modular functions $f_x(x) = \Phi_{m_x}(x)$, $f_y(y) = \Phi_{m_y}(y)$ as $f_x(x) = x^{n_x} + 1$, $f_y(y) = y^{n_y} + 1$ have been proposed in [68], as this structure presents computational advantages and can be very beneficial for practical applications.

Bootland *et al.*'s attack is able to exploit common roots on the involved rings to factorize the multivariate RLWE samples into RLWE samples of smaller dimension. For example, consider that $n_x = n_y = n$; by applying the substitution $y \leftarrow x$, we obtain n RLWE samples of dimension n each, hence decreasing the n^2 lattice dimension of the original m-RLWE sample.

Secure multivariate RLWE instantiations Let $m = m_x m_y$ and $\gcd(m_x, m_y) = 1$; then, the m-th cyclotomic field $K = \mathbb{Q}(\zeta_m) \cong \mathbb{Q}[x]/(\Phi_m(x))$ (with ζ_m the m-th root of unity) is isomorphic (see Theorem 1) to the bivariate field

$$K \cong \mathbb{Q}[x,y]/(\Phi_{m_x}(x), \Phi_{m_y}(y)). \tag{1}$$

Consequently, by considering instantiations satisfying $\gcd(m_x, m_y) = 1$, the bivariate RLWE problem becomes equivalent to the equally sized RLWE problem. However, we would like to search for other instantiations where the modular functions can have a similar form and, if possible, the same degree.

By restricting ourselves to the most common scenario of "power-of-two" cyclotomics, modular functions of the form $\{x^{n_x} + d_x, y^{n_y} + d_y, z^{n_z} + d_z, \ldots\}$, could

avoid Bootland *et al.*'s attack for some parameters $\{n_x, d_x, n_y, d_y, n_z, d_z, \ldots\}$. E.g., the rings $\mathbb{Z}[x]/(x^{64}+1)$ and $\mathbb{Z}[y]/(y^{27}+5)$ do not have common roots, so trivial substitutions such as $x \to y$ cannot be applied. Additionally, whenever we reduce modulo q and work over R_q, we can impose (for the sake of efficiency) that both modular functions $x^{64}+1$ and $y^{27}+5$ factor in linear terms enabling the use of variants of the NTT. Additionally, slot encoding and slot manipulations are still possible in the plaintext ring by means of the pre-/post-processing, as presented in [67]. Analogously to the negayclic convolution, these pre-/post-processing steps preserve the properties of the NTT transform over a ring with an α-generalized convolution [64].

This seems to effectively avoid a substitution attack; however, there might be some small ideal divisor for which, modulo some particular q, the noise does not increase substantially, and we can distinguish the resulting sample from uniform. This attack has been extensively studied by Peikert in [73] and we will discuss it in Section 7.1.

The proposed solution The previous strategy preserves most of the advantages of the multivariate constructions while apparently avoiding the effects of Bootland *et al.*'s attack. However, the security of these instantiations is not based on any specific formulation of the RLWE problem, and there is no trivial way of parameterizing them. This raises the following questions:

1. *Can we find multivariate rings similar to* $\mathbb{Z}[x, y, \ldots]/(x^{n_x}+d_x, y^{n_y}+d_y, \ldots)$ *while (a) still preserving the aforementioned advantages of this structure, and (b) basing its security on the hardness of the RLWE problem (see Definition 7); i.e., without a decrease in the ring dimension due to Bootland's attack (see Theorem 3)?*
2. *If these multivariate rings exist, how can the values* $\{n_x, n_y, \ldots\}$, $\{d_x, d_y, \ldots\}$ *be chosen to easily define the ring of integers* R*, its dual* $R^\vee$ *and the basis?*

From this point forward, we focus on answering these two questions. To this aim, we identify number fields whose ring of integers (and their dual) have the sought structure (see Section 4). In particular, we divide this set of fields in two categories:

1. *Multiquadratic number fields* (see Section 5). These structures enable efficient radix-2 transforms for faster polynomial arithmetic (see Section 8).
2. *More general number fields with modular functions* $\{x^{n_x}+d_x,\ y^{n_y}+d_y,\ \ldots\}$ (see Section 6). These structures support all the signal processing applications described in [71], and the matrix operations introduced by the original MHEAAN scheme [26] (not based on coprime cyclotomic polynomials [28]) while preserving the equivalent RLWE security.

Rationale on the security of our solution The weakness of some m-RLWE instantiations is rooted on the existence of (small norm) zero divisors in the compositum field. For example, $\mathbb{Q}[x, y]/(x^2+1, y^2+1)$ has zero divisors as

$x + y$ (e.g., $(x + y)(x - y) = 0$), and hence m-RLWE samples defined on rings $\mathbb{Z}[x, y]/(x^2 + 1, y^2 + 1)$ can be easily factored, as the effective *degree* can be reduced with substitutions $\{x \to y, -x \to y\}$. Additionally, as these roots have *small norm*, the noise in the reduced samples is not increased enough to preserve security.

Instead of the previously proposed $\mathbb{Z}[x, y]/(x^2 + 1, y^2 + 1)$, we work with a bivariate ring with modular functions of the form $\{x^{n_x} + d_x, y^{n_y} + d_y\}$ (we use $\mathbb{Z}[x, y]/(x^2 + 1, y^2 + 3)$ as our example). The use of different modular functions avoids a trivial substitution attack. However, we need to rule out the possibility of (small norm) substitution attacks, such as the one from [12], modulo some q; even if they exist, finding them would require solving a hard subset-sum mod$q$ (knapsack) problem.

As there is a security reduction from ideal lattices to RLWE defined on general number fields [75], we search for the ring of integers of *multivariate number fields*. This gives us a way to find secure parameters for the used ring, and also the right error distribution to guarantee that the noise increase after a substitution modulo q is enough to preserve the required security [73]. To exemplify this rationale, we compare the differences between a bivariate cyclotomic ring (which can be seen as a univariate cyclotomic ring), and our proposed solution.

Consider the ring $\mathbb{Z}[z]/\Phi_{12}(z)$ with $\Phi_{12}(z) = z^4 - z^2 + 1$. There is an isomorphism with the bivariate ring $\mathbb{Z}[x, y]/(\Phi_4(x), \Phi_3(y))$ where $\Phi_4(x) = x^2 + 1$ and $\Phi_3(y) = y^2 + y + 1$. Therefore, our intuition is that if we found an effective substitution attack on our example ring $\mathbb{Z}[x, y]/(x^2 + 1, y^2 + 3)$, this would work analogously for the cyclotomic bivariate case $\mathbb{Z}[x, y]/(\Phi_4(x), \Phi_3(y))$. In particular, if we apply the transformation $T(y) = 2y + 1$ in the ring $\mathbb{Z}[y]/(y^2 + 3)$, we obtain $\mathbb{Z}[y]/y^2 + y + 1$, which is the mentioned cyclotomic ring with $\Phi_3(y)$. Consequently, for this particular case, it is clear that the existence of an attack in our example ring implies an attack to the bivariate cyclotomic ring.

For more general multivariate rings, we can apply a similar idea. In general, for a secure bivariate ring such as $\mathbb{Z}[x, y]/(x^{n_x} + d_x, y^{n_y} + d_y)$, we can search for a transformation $y \leftarrow T(y)$ where the new modular function can share at least some roots with $x^{n_x} + d_x$. If this transformation can be effectively applied, we could use it to attack multivariate cyclotomic rings.

Thus, this strengthens the belief that an attack on secure m-RLWE instantiations defined on a general number field should provide us with either an attack to RLWE on the product of prime-powers cyclotomic rings, and/or a better understanding on the weaknesses of general cyclotomic rings.

For a discussion on the practical security of RLWE on the proposed number fields we refer the reader to Section 7.1.

Division algebras and non-norm condition In [46], the authors propose an alternative variant of LWE over cyclic algebras, which they denote as CLWE (Cyclic Algebra LWE). The main difference with respect to RLWE relies on the fact that, instead of adding a ring structure, they incorporate into Module-LWE a cyclic algebra structure, constructing a non-commutative variant of LWE.

The security of CLWE is supported by the hardness of finding short vectors in certain structured lattices induced by ideals in a cyclic algebra $\mathcal{A}$. Additionally, they explicitly address Bootland *et al.*'s attack by means of the "non-norm" condition (see Proposition 1).

Let a cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ where K is a number field of degree k and L is a Galois extension of K of degree n such that $Gal(L/K) = \langle\theta\rangle$. For a non-zero $\gamma \in K$, the cyclic algebra is defined as $\mathcal{A} = L \oplus uL \oplus \ldots \oplus u^{n-1}L$ where $u \in \mathcal{A}$ and $u^n = \gamma$ (see Definition 9).

The *non-norm* condition on γ (see Proposition 1) avoids Bootland *et al.*'s attack by stating that the lowest power of γ which appears in $N_{L/K}(L)$ is γ^n, where $N_{L/K}$ represents the relative norm of L into K (see [46] for more details).

This defense against Bootland *et al.*'s attack also relies on avoiding the existence of zero divisors, which are needed for the attack to succeed. In our case, as we already work in a number field, we can adhere to the security conditions established by Peikert [73] to avoid this type of attacks.

It is worth mentioning that we see both approaches as potentially complementary, in such a way that the underlying field K considered in the (cyclic) division algebra (see Definition 9) could be one of the multivariate fields discussed in Sections 5 and 6.

Contributions The main contribution of this work is the definition and parameterization of secure instantiations of the multivariate Ring Learning With Errors problem [70, 71], supported by the extended reduction [75] of the original proof by Lyubashevsky *et al.* [60, 58]. The proposed instantiations address the vulnerability leveraged on Bootland's attack to m-RLWE [12], while still preserving all the efficiency improvements that m-RLWE brings. Moreover we show that is possible to securely instantiate the m-RLWE problem, because the canonical embedding of R has a polynomial skewness (λ_n/λ_1).

The applications of these secure instantiations are numerous, achieving improved space-time tradeoffs in the most critical lattice operations, and therefore enabling more efficient homomorphic processing and closing the gap to the realization of practical fully homomorphic encryption. This is the main list of applications:

- We introduce the $\boldsymbol{\alpha}$-generalized Walsh-Hadamard Transform as the basic block that can replace Number Theoretic Transforms in multivariate rings, achieving an improvement on the computational complexity of degree-n polynomial products by a factor $\log(n)$ in terms of elemental multiplications, with additional savings in memory usage (see Section 8).
- We enable net improvements in cryptographic primitives built on top of m-RLWE, such as efficient time and space computation of automorphisms, relinearizations, packing, unpacking and homomorphic slot manipulation, and, consequently, bootstrapping, improving on current achievable trade-offs in RLWE (see Section 9).
- We instantiate a simple cryptosystem based on m-RLWE (see Section 7.2), and exemplify with it the use of the multivariate structure of m-RLWE

to improve on complex number embedding, enabling fully packed complex numbers, compared to the exponentially decreasing packing ratio of current approaches [26, 28] (see Section 10). This enables applications in homomorphically encrypted approximate arithmetic, complex processing, and efficient multidimensional signal manipulation (see Appendix A).

Structure The rest of the paper is organized as follows: Section 2 describes Bootland *et al.*'s attack to multivariate RLWE. Section 3 introduces some algebraic number theory notions and the main definitions for the m-RLWE problem. Section 4 describes the followed strategy to achieve secure instantiations of multivariate RLWE, including the well-known tensor of "coprime" cyclotomic rings. Section 5 focuses on the analysis of multiquadratic rings. Section 6 studies a set of more general multivariate rings. Section 7 includes a discussion on the achieved resilience against known attacks together with example instantiations that showcase the practicality of multivariate RLWE, and discusses some practical applications. Finally, Section 11 draws some conclusions. Additionally, the included Section 8 particularizes the problem to rings enabling an $\boldsymbol{\alpha}$-generalized Walsh-Hadamard Transform, and compares its performance with fast NTT algorithms currently used in state-of-the-art RLWE cryptosystems. Section 9 introduces the strategies for homomorphic packing/unpacking and the space/time tradeoffs improving on current RLWE relinearization and bootstrapping operations.

2 Worst case security of multivariate RLWE

We first introduce the notation used in this section. Polynomials are denoted with regular lowercase letters, omitting the polynomial variable (i.e., a instead of $a(x)$) when there is no ambiguity.

We follow a recursive definition of multivariate modular rings: $R_q[x_1] = \mathbb{Z}_q[x_1]/f_1(x_1)$ denotes the polynomial ring in the variable x_1 modulo $f_1(x_1)$ with coefficients belonging to $\mathbb{Z}_q$. Analogously, $R_q[x_1, x_2] = (R_q[x_1])[x_2]/(f_2(x_2))$ is the bivariate polynomial ring with coefficients belonging to $\mathbb{Z}_q$ reduced modulo univariate $f_1(x_1)$ and $f_2(x_2)$. In general, $R_q[x_1, \ldots, x_l]$ (resp. $R[x_1, \ldots, x_l]$) represents the multivariate polynomial ring with coefficients in $\mathbb{Z}_q$ (resp. $\mathbb{Z}$) and the l modular functions $f_i(x_i)$ with $1 \leq i \leq l$. The polynomial a can also be denoted by a column vector $\boldsymbol{a}$ whose components are the corresponding polynomial coefficients.

For the sake of clarity, we present the definition of multivariate RLWE with power-of-two cyclotomic polynomials, as originally introduced in [68], but all the results in this section can be generalized to any cyclotomic function:

Definition 1 (multivariate RLWE with power-of-two modular functions as $x_i^{n_i} + 1$)**.** *Given a multivariate polynomial ring $R_q[x_1, \ldots, x_l]$ with $f_j(x_j) = 1 + x_j^{n_j}$ for $j = 1, \ldots, l$ where $n = \prod_j n_j$ (with all n_j a power of two) and an error distribution $\chi[x_1, \ldots, x_l] \in R_q[x_1, \ldots, x_l]$ that generates small-norm random multivariate polynomials in $R_q[x_1, \ldots, x_l]$, the multivariate polynomial*

RLWE problem relies upon the computational indistinguishability between samples $(a_i, b_i = a_i \cdot s + e_i)$ *and* (a_i, u_i)*, where* $a_i, u_i \leftarrow R_q[x_1, \ldots, x_l]$ *are chosen uniformly at random from the ring* $R_q[x_1, \ldots, x_l]$*;* $s, e_i \leftarrow \chi[x_1, \ldots, x_l]$ *are drawn from the error distribution.*

The original works of multivariate RLWE [68, 71] assume that the search and decision m-RLWE problems (see Definitions 3 and 4) in dimension $n = \prod_{i=1}^{m} n_i$ are as hard as the corresponding RLWE problems in dimension n. However, Bootland *et al.* [12] introduced an attack that can exploit modular functions that allow repeated "low-norm" roots in the multivariate ring. As a result, when the subrings of the tensor have common roots, this attack is able to factor the m-RLWE samples into RLWE samples of smaller dimension, hence reducing the security of these m-RLWE samples to that of solving a set of independent RLWE samples which are easiest to break. E.g., for the ring $\mathbb{Z}[x, y]/(x^{2n} + 1, y^n + 1)$, changes of variable $y \leftarrow x^{2i}$ with $i \in \mathbb{Z}_{2n}^*$ factors the m-RLWE sample into n different RLWE samples with rings of modular function $x^{2n} + 1$ and an increase in the error variance of n (maximum degree of $y^n + 1$).

The instantiations of (multivariate) RLWE with cyclotomic rings where the different modular functions have "coprime" order are not affected by this attack, as they do not introduce these "common" roots (see Section 4.1).

We now give a more formal description of the attack, particularized for bivariate RLWE (2-RLWE) with power of two cyclotomics (Definition 1). Let $(a, b = as+e) \in R_q^2[x, y]$ and $R_q[x, y] = \mathbb{Z}_q[x, y]/\left(x^{n_x} + 1, y^{n_y} + 1\right)$ with $n_x \geq n_y$ and $k = \frac{n_x}{n_y}$ without loss of generality.

Now we define the map $\tilde{\Theta}$:

$$\begin{aligned} \tilde{\Theta} : \mathbb{Z}_q[x, y]/(x^{n_x} + 1, y^{n_y} + 1) &\rightarrow (\mathbb{Z}_q[x]/(x^{n_x} + 1))^{n_y} \\ a(x, y) &\rightarrow \left(a(x, x^k), a(x, x^{3k}), \ldots, a(x, x^{(2n_y-1)k})\right) \end{aligned}$$

This map is a ring homomorphism, and if q is odd it is also invertible (see [12]). This transforms the pair $(a, b) \in R_q[x, y]$ into $(\tilde{\Theta}(a), \tilde{\Theta}(b) \in R_q^{n_y}[x]$. If we denote each of the components by $\tilde{\Theta}_i$, for $i = 1, \ldots, n_y$, we have

$$\left(\tilde{\Theta}_i(a), \tilde{\Theta}_i(b) = \tilde{\Theta}_i(a)\tilde{\Theta}_i(s) + \tilde{\Theta}_i(e)\right) \in R_q^2[x], \tag{2}$$

for $i = 1, \ldots n_y$. This results in n_y different RLWE samples of dimension n_x and whose noise has a variance n_y times higher than the original 2-RLWE sample (the result of adding n_y independent variables).

The attack works by trying to break the obtained n_y RLWE samples. Once this is done, as the map is invertible, it is possible to reconstruct the original secret key with the different n_y smaller keys.

This attack can be generalized to an m-RLWE sample (Definition 1) by recursively applying "versions" of this map $(l - 1)$ times. This recursion converts an m-RLWE sample into $\frac{n}{n_1}$ RLWE samples (assuming, without loss of generality, that $n_1 \geq n_2 \geq \ldots \geq n_l$) with dimension n_1 and an error variance $\frac{n}{n_1}$ times higher.

3 Multivariate Ring Learning with Errors

This section revisits the main definitions from algebraic number theory and multivariate RLWE, including a generalized version of the multivariate polynomial RLWE problem which admits any type of cyclotomic polynomial as modular function. For the sake of clarity, we particularize to power-of-two modular cyclotomic functions (Definition 1) when exemplifying some of the results, but this does not affect to the generality of the discussion.

3.1 Algebraic Number Theory background

This section presents the fundamental concepts of lattices and algebraic number theory and extends them to the more general case of a tensor of number fields on which m-RLWE is based.

The Space $\boldsymbol{H_{(T)} = \bigotimes_i H_i}$ When dealing with cyclotomic fields, it is useful to work with the subspace $H \subseteq \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2}$ with $s_1 + 2s_2 = n$, where the tuple (s_1, s_2) is called the signature of the number field, and H satisfies

$$H = \{(x_1, \ldots, x_n) \in \mathbb{R}^{s_1} \times \mathbb{C}^{2s_2} \text{ such that } x_{s_1+s_2+j} = \bar{x}_{s_1+j}, \forall j \in [s_2]\} \subseteq \mathbb{C}^n \tag{3}$$

An orthonormal basis $\{\boldsymbol{h}_j\}_{j\in[n]}$ for H can be defined as

$$\boldsymbol{h}_j = \begin{cases} \boldsymbol{e}_j & \text{if } j \in [s_1] \\ \frac{1}{\sqrt{2}}(\boldsymbol{e}_j + \boldsymbol{e}_{j+s_2}) & \text{if } s_1 < j \leq s_1 + s_2 \\ \frac{\sqrt{-1}}{\sqrt{2}}(\boldsymbol{e}_{j-s_2} - \boldsymbol{e}_j) & \text{if } s_1 + s_2 < j \leq s_1 + 2s_2, \end{cases}$$

where $\boldsymbol{e}_j$ are the vectors of the standard basis in $\mathbb{R}^n$. Each element $a = \sum_{j\in[n]} a_j \boldsymbol{h}_j \in H$ (with $a_j \in \mathbb{R}$) has its own l_p norm. For our purposes, we define the subspace $H_{(T)} = \bigotimes_{i\in[l]} H_i$ as the tensor product of l subspaces H_i (each subspace H_i defined as in Eq. (3) but with $s_1 + 2s_2 = n_i$).

In particular, if we see each element belonging to each H_i as a different linear transformation, we are actually working with the Kronecker product of the subspaces H_i. We can therefore express an orthonormal basis for $H_{(T)}$ given by $\{\boldsymbol{h}_j\}_{j\in[n]}$ as the result of the Kronecker product of the original basis of each H_i, by defining any invertible mapping for j and $\{j_1, \ldots, j_l\}$, where $\boldsymbol{h}_j = \bigotimes_{i\in[l]} \boldsymbol{h}_{j_i}^{(i)}$ are the basis vectors for $H_{(T)}$, and $n = \prod_{i\in[l]} n_i$; each $\{\boldsymbol{h}_{j_i}^{(i)}\}_{j_i\in[n_i]}$ is the orthonormal basis of each $H_i \subseteq \mathbb{C}^{n_i}$ for $i \in [l]$.

Lattice background A lattice in our multivariate setting is defined as an additive subgroup of $H_{(T)}$. We only consider full rank lattices, obtained as the set of all integer linear combinations of a set of n linear independent basis vectors $\boldsymbol{B} = \{\boldsymbol{b}_1, \ldots, \boldsymbol{b}_n\} \subset H_{(T)}$

$$\Lambda = \mathcal{L}(B) = \left\{ \sum_{i\in[n]} z_i \boldsymbol{b}_i \text{ such that } \boldsymbol{z} \in \mathbb{Z}^n \right\}$$

The minimum distance $\lambda_1(\Lambda)$ of a lattice Λ for the norm $||.||$ is given by the length of the shortest non-zero lattice vector, that is, $\lambda_1(\Lambda) = \min_{\boldsymbol{x}\in\Lambda/\boldsymbol{x}\neq\boldsymbol{0}} ||\boldsymbol{x}||$.

The dual lattice of $\Lambda \subset H_{(T)}$ is defined as $\Lambda^* = \{\boldsymbol{x} \in H_{(T)} | \langle \Lambda, \boldsymbol{x}\rangle \subseteq \mathbb{Z}\}$ and it satisfies $(\Lambda^*)^* = \Lambda$.

Gaussian Measures The results on nonspherical Gaussian distributions presented in [58] can be extended to our case. Hence, we revisit here some of the concepts for Gaussian measures, adapted to our tensor setting.

We consider the Gaussian function $\rho_r : H_{(T)} \to (0, 1]$ with $r > 0$ as $\rho_r(\boldsymbol{x}) = \exp(-\pi||\boldsymbol{x}||^2/r^2)$. A continuous Gaussian probability distribution D_r can be obtained by normalizing the previous function to obtain a probability density function as $r^{-n}\rho_r(\boldsymbol{x})$. Extending this to the non spherical Gaussian case, we consider the vector $\boldsymbol{r} = \bigotimes_{i\in[l]} \boldsymbol{r}_i$ where $\boldsymbol{r} = (r_1, \ldots, r_n) \in (\mathbb{R}^+)^n$ and $\boldsymbol{r}_i = (r_{i,1}, \ldots, r_{i,n_i}) \in (\mathbb{R}^+)^{n_i}$ and whose components satisfy $r_{i,j+s_1+s_2} = r_{i,j+s_1}$. Finally, a sample from $D_{\boldsymbol{r}}$ is given by $\sum_{j\in[n]} x_j \boldsymbol{h}_j$ where each x_j is drawn independently from a Gaussian distribution D_{r_j} over $\mathbb{R}$; r_j equals $\prod_{i\in[l]} r_{i,j_i}$ (where l is the number of "unidimensional" spaces H_i in the tensor, that is $n = \prod_{i\in[l]} n_i$) and we are using any invertible mapping between $\{j\}_{j\in[n]}$ and $\{j_i\}_{j_i\in[n_i],i\in[l]}$.

3.2 Main Definitions for Multivariate Ring-LWE

Let $K_{(T)} = \bigotimes_{i\in[l]} K_i$ be the tensor product of l cyclotomic fields of dimension $n_i = \phi(m_i)$ each, and $R = \bigotimes_{i\in[l]} \mathcal{O}_{K_i}$ ($R^\vee = \bigotimes_{i\in[l]} \mathcal{O}^\vee_{K_i}$) the tensor of their corresponding (dual of the) ring of integers. We have the following definitions:

Definition 2 (Multivariate Ring LWE distribution)**.** *For $s \in R^\vee_q$ and an error distribution ψ over $K_{(T),\mathbb{R}}$, a sample from the m-RLWE distribution $A_{s,\psi}$ over*[4] *$R_q \times \mathbb{T}$ is generated by $a \leftarrow R_q$ uniformly at random, $e \leftarrow \psi$, and outputting $(a, b = (a \cdot s)/q + e \bmod R^\vee)$.*

Definition 3 (Multivariate Ring LWE, Search)**.** *Let Ψ be a family of distributions over $K_{(T),\mathbb{R}}$. m-RLWE$_{q,\Psi}$ denotes the search version of the m-RLWE problem. It is defined as follows: given access to arbitrarily many independent samples from $A_{s,\Psi}$ for some arbitrary $s \in R^\vee_q$ and $\psi \in \Psi$, find s.*

Definition 4 (Multivariate Ring LWE, Average-Case Decision)**.** *Let Υ be a distribution over a family of error distributions, each over $K_{(T),\mathbb{R}}$. The average-case decision version of the m-RLWE problem, denoted m-R-DLWE$_{q,\Upsilon}$, is to distinguish with non-negligible advantage between arbitrarily many independent*

[4] $\mathbb{T} = K_{(T),\mathbb{R}}/R^\vee$ and $K_{(T),\mathbb{R}} = K_{(T)} \otimes_{\mathbb{Q}} \mathbb{R}$.

samples from $A_{s,\psi}$, for a random choice of $(s,\psi) \leftarrow U(R_q^\vee) \times \Upsilon$,[5] and the same number of uniformly random and independent samples from $R_q \times \mathbb{T}$.

For an asymptotic treatment of the m-RLWE problems, we let $K_{(T)}$ come from an infinite sequence of tensors of number fields $\mathbb{K} = \{K_{(T),n}\}$ of increasing dimension n ($n = \prod_i \phi(m_i)$ is the number of basis elements that form the integral basis), and let q, Ψ, and Υ depend on n as well.

Error distributions We include here two definitions about the error distributions.

Definition 5 (extension of Lyubashevsky *et al.* [58], Definition 3.4)**.** *For a positive real $\alpha > 0$, the family $\Psi_{\leq\alpha}$ is the set of all elliptical Gaussian distributions $D_{\boldsymbol{r}}$ (over $K_{(T),\mathbb{R}}$), where each parameter $r_i \leq \alpha$ with $i \in [n]$.*

Definition 6 (extension of Lyubashevsky *et al.* [58], Definition 3.5)**.** *Let $K_{(T)} = \bigotimes_{i\in[l]} K_i$ where the K_i are the m_i-th cyclotomic number field having degree $n_i = \phi(m_i)$. For a positive real $\alpha > 0$, a distribution sampled from Υ_α is given by an elliptical Gaussian distribution $D_{\boldsymbol{r}}$ (over $K_{(T),\mathbb{R}}$) whose parameters are $r_j \in [n]$ using the unidimensional index (see Section 3.1), and each r_j satisfies $r_j^2 = \alpha^2(1+\sqrt{n}x_j)$ where different x_j,x_k that do not correspond to conjugate positions are chosen independently from the distribution $\Gamma(2,1)$.*

Practical applications [68, 67, 66] usually deal with variants of the problem:

- *discrete b*: Instead of working with an error distribution ψ over $K_{(T),\mathbb{R}}$, the m-RLWE distribution $A_{s,\chi}$ can use χ as a discrete error distribution over $R^\vee$, so that the element b belongs to $R_q^\vee$.
- *small key*: Instead of a uniform s, s can be a "short key" equivalently sampled from the error distribution (this is known as "normal form" in [61]), with equivalent security. Given a list of l m-RLWE samples, s can be substituted with the error e of any sample (a,b) whose term a is invertible in R_q, which occurs with constant probability by Claim 3.2 below.
- *power of 2 cyclotomic*: Instead of sampling a and s from R_q and $R_q^\vee$ respectively, both are usually sampled from R_q (this is usually known as the non-dual variant). In general, the works which consider s in R_q deal with cyclotomic fields where m_i is a power of two. It can be shown that for this particular type of cyclotomic fields both definitions are equivalent.
- *modulus switching*: The original definitions of the problem are presented with a prime modulus q that splits the space into small independent coordinates. With the same hardness guarantees, it is possible to modulus-switch to other compute-friendly modulus at the price of a slight increase of the error.

Lyubashevsky *et al.* [61] show that the variant of RLWE with discrete and short error (R-DLWE$_{q,\chi}$) is as hard as the original R-DLWE$_{q,\psi}$, by following the technique from [4]. These results can be adapted to our more general case as follows:

[5] $U(R_q^\vee)$ represents the uniform distribution over $R_q^\vee$.

Claim. The fraction of invertible elements in R_q is $\bigotimes_{i\in[l]} \mathcal{O}_{K_i}/\langle q\rangle$, for prime $q = 1 \bmod \phi(m_i)$ for all i is $(1-\frac{1}{q})^n$, with $n = \prod_i \phi(m_i)$. Thus, if $q \geq n$, this probability is constant.

Proof. Since R_q is in bijection with the ring $(\mathbb{Z}/q\mathbb{Z})^n$ via the tensor embedding mod q, so an element is invertible iff. its image does not contain any zero. Hence, there are $(q-1)^n$ invertible elements out of q^n. □

Pseudorandomness of m-RLWE: To show that the m-RLWE distribution is pseudorandom (that is, there exists a reduction from the search problem to the decision variant of the hardness problem) we rely on the results from [58], applied to the case of multivariate elements. The main needed properties are those related to the decomposition of $\langle q\rangle$ into n prime ideals ($q \equiv 1 \bmod \phi(m_i)$ for all i) and the use of the automorphisms that permute the prime ideals.

4 Proposed approach for secure multivariate rings

Despite the efficiency benefits of multivariate RLWE, its security can be much smaller than originally expected for those instances vulnerable to Bootland *et al.*' s attack [12]. This motivates us to redefine the set of instantiations that preserve the security in the tensor lattice dimension.

This section enumerates those secure instantiations of multivariate RLWE. With this in mind, we first briefly revise the choice of "coprime" order cyclotomics explicitly included in [61]. Afterwards, we discuss the possibility of using a more general set of number fields, enabling other multivariate rings that can be more convenient for practical applications.

4.1 Multivariate RLWE as a subset of RLWE

It is well known that for two cyclotomic number fields $\mathbb{Q}(\zeta_a)$ and $\mathbb{Q}(\zeta_b)$ with coprime orders $\gcd(a,b) = 1$, their product is the cyclotomic number field $\mathbb{Q}(\zeta_{ab})$ (see Lemma 11.8 in [33]). For convenience, we include an adapted version of this property using the polynomial representation of the cyclotomic number fields.

Theorem 1 (Tensorial decomposition of cyclotomic number fields, see Equation (1.1) **in [61]).** *The m-th cyclotomic field $K = \mathbb{Q}(\zeta_m) \cong \mathbb{Q}[x]/(\Phi_m(x))$ (with ζ_m the m-th root of unity) is isomorphic to the multivariate field*

$$K \cong \mathbb{Q}[x_1,\ldots,x_l]/(\Phi_{m_1}(x_1),\ldots,\Phi_{m_l}(x_l)), \tag{4}$$

where $m = \prod_i m_i$ is decomposed in its prime-power decomposition with $\gcd(m_j, m_k) = 1$ for all $j \neq k$.

This fact gives an alternative basis to the power basis $\{1, x, \ldots, x^{\phi(m)-1}\}$ for the ring of integers $R = \mathbb{Z}[x]/\Phi_m(x)$; this basis is the "powerful" basis of

K composed of elements $\prod_i x_i^{j_i}$ with $0 \leq j_i < \phi(m_i)$.[6] This "powerful" basis has some very nice properties [61] which make it more appealing than the more "conventional" power basis. Additionally the authors of [61] provide a detailed analysis on how the performance of ring operations can be improved by means of this multivariate structure.

Besides [61], the use of the multivariate structure in Eq. (4) has been exploited to enhance polynomial operations in both the HElib [47, 48] and the MHEAAN [28] libraries. This gives us a first approach to deal with multivariate instantiations which do not suffer a decrease on the underlying lattice dimension. However, this structure is not flexible enough to convey the same benefits that general multivariate structures can achieve; in particular, it cannot preserve the interesting structure of power-of-two cyclotomics $(1 + x^n)$.

4.2 More general RLWE instantiations

We look now beyond cyclotomics, into more general and flexible number fields and their parameterization. We first introduce the definitions of RLWE over any number field [75], and then give the intuition on the properties required to resist the Bootland *et al.*'s attack. A detailed discussion on the choice of good parameters and the security of RLWE on these number fields follows in Sections 5, 6 and 7.1.

RLWE over any number field Peikert *et al.* [75] have recently generalized the RLWE problem to *any number field*. Let K be a number field with ring of integers $R = \mathcal{O}_K$; let $R^\vee$ be the fractional codifferent ideal of K, and let $\mathbb{T} = K_\mathbb{R}/R^\vee$. Let $q \geq 2$ be a (rational) integer modulus, and for any fractional ideal I of K,[7] let $\mathcal{I}_q = \mathcal{I}/q\mathcal{I}$. We include now the relevant definitions of RLWE over any number field that we use in our formulation.

Definition 7 (Ring-LWE Distribution, Definition 2.14 in [75])**.** *For $s \in R_q^\vee$ and an error distribution ψ over $K_\mathbb{R}$, the $R - LWE$ distribution $A_{s,\psi}$ over $R_q \times \mathbb{T}$ is sampled by independently choosing a uniformly random $a \leftarrow R_q$ and an error term $e \leftarrow \psi$, and outputting $(a, b = (a \cdot s)/q + e \bmod R^\vee)$.*

Definition 8 (Ring-LWE, Average-Case Decision, Definition 2.15 in [75])**.** *Let Υ be a distribution over a family of error distributions, each over $K_\mathbb{R}$. The average-case Ring-LWE decision problem, denoted $R - LWE_{q,\Upsilon}$, is to distinguish (with non-negligible advantage) between independent samples from $A_{s,\psi}$ for a* random *choice of $(s, \psi) \leftarrow U(R_q^\vee) \times \Upsilon$, and the same number of uniformly random and independent samples from $R_q \times \mathbb{T}$.*

[6] This basis does not coincide with the power basis under the mentioned automorphism and considering the map $x^{\frac{m}{m_i}} \rightarrow x_i$ for $i = 1, \ldots, l$ (see [61]).

[7] For any fractional ideal $\mathcal{I} \subset K$ there is $a \in \mathcal{O}_K$ such that $a\mathcal{I} \subseteq \mathcal{O}_K$ is an integral ideal of $\mathcal{O}_K$.

Theorem 2 (Theorem 6.2 from [75]). *Let K be an arbitrary number field of degree n, $\mathcal{I}$ any fractional ideal of K, and $R = \mathcal{O}_K$. Let $\alpha = \alpha(n) \in (0,1)$, and let $q = q(n) \geq 2$ be an integer such that $\alpha q \geq 2 \cdot \omega(1)$. There is a polynomial-time quantum reduction from $K - DGS_\gamma$ to (average-case, decision) $R - LWE_{q,\Upsilon_\alpha}$, for any*

$$\gamma = \max\left\{\eta(\mathcal{I}) \cdot \sqrt{2}/\alpha \cdot \omega(1), \sqrt{2n}/\lambda_1(\mathcal{I}^\vee)\right\}.$$

Additionally, it is worth highlighting some observations on the choice of a particular number field in RLWE, as stated in [75]:

- The geometry of the dual ideal $R^\vee$ affects the error rate α (chosen to be smaller than the minimum distance $\lambda_1(R^\vee)$). As α decreases, worst-case hardness theorems give weaker guarantees (i.e., larger approximation factors), and known attacks on Ring-LWE become more efficient.
- A similar phenomenon arises for rings with large "expansion factors" (see [59]) which imposes smaller α for achieving correct decryption; hence, good rings for practical applications have small expansion factors.
- Besides the two previous relations, there is no practical evidence on which particular number field is better in terms of security.

Ad-hoc countermeasures to Bootland *et al.*'s attack Bootland's attack [12] shows that a reduced RLWE sample is at least as hard as an m-RLWE sample. To prove the converse, we can use an oracle for m-RLWE. With access to such oracle and a set of RWLE samples with different keys, we can construct an m-RWLE sample (with a slight increase in the noise variance) by means of the reverse map of Bootland *et al.*'s attack (i.e., $\tilde{\Theta}^{-1}$). Once this oracle returns the secret key of the m-RLWE sample, the original keys of the RLWE sample can be recovered by means of the map $\tilde{\Theta}$.

We can therefore express the security of m-RLWE in terms of RLWE, but the decrease of the involved dimension considerably reduces the applicability of the problem with "non-coprime" modular functions. The security of $\prod_{j\neq k}\phi(\gcd(m_j, m_k))$ independent RLWE samples with dimension $\frac{\prod_{i\in[l]}\phi(m_i)}{\prod_{j\neq k}\phi(\gcd(m_j,m_k))}$ could be reduced to that of one m-RLWE sample (according to Definition 2) with dimensions $\{\phi(m_1), \ldots, \phi(m_l)\}$:

Theorem 3 ($\tilde{\Theta}^{-1}$ transform from [12]). *Let L independent univariate RLWE samples $(a_i, b_i) \in R_q \times \mathbb{T}$ for $i \in [L]$ and dimension n. We can transform (this transformation is invertible when q is prime) these L samples by means of the (inverse) of Bootland* et al.*'s attack into one m-RLWE sample with l dimensions $\{\phi(m_1), \ldots, \phi(m_l)\}$ (see Definition 2) satisfying $L = \prod_{j\neq k}\phi(\gcd(m_j, m_k))$ and having for the RLWE sample $n = \frac{\prod_{i\in[l]}\phi(m_i)}{L}$. This transformation slightly increases the variance of the error distribution by a factor L.*

The decrease in the lattice dimension by a factor $L = \prod_{j\neq k}\phi(\gcd(m_j, m_k))$ brings about the question of whether we can *modify some of the multivariate RLWE constructions* where $L > 1$ to *avoid Bootland et al.'s attack.*

Followed strategy By considering instantiations satisfying $\gcd(m_j, m_k) = 1$ for all $j \neq k$, we straightforwardly go back again to the RLWE problem. However, we would like to find other instantiations where the modular functions can have a similar form and degree. We will hence focus on modular functions as follows: $\{x^{n_x} + d_x, y^{n_y} + d_y, z^{n_z} + d_z, \ldots\}$, which can avoid Bootland's attack for certain parameters, while enabling NTT-like fast transforms and preserving the advantages of the originally introduced m-RLWE constructions.

However, the security of these instantiations is not based on any specific formulation of the RLWE problem, so we do not have a clear way of choosing the right parameters. In the next two sections, we focus on number fields satisfying Definition 7 and whose ring of integers (and their dual) has the aforementioned structure. In particular, we focus on multiquadratic number fields (Section 5) and more general multivariate rings (Section 6). Before delving into these two cases, we briefly discuss an alternative secure approach based on division algebras.

Cyclic division algebras (over multivariate number fields) A recent and independent work [46] presents a different solution based on cyclic algebras. To explicitly avoid Bootland *et al.*'s attack, they make use of the "non-norm" condition [81], which has the role of avoiding the existence of zero divisors. For the sake of completeness, we include the relevant definitions.

Definition 9 (Cyclic Algebra, Definition 9 in [46])**.** *Let K be a number field with degree k, and let L be a Galois extension of K of degree n such that the Galois group of L over K is cyclic of degree n, $Gal(L/K) = \langle\theta\rangle$. For a non-zero $\gamma \in K$ we define the resulting cyclic algebra*

$$\mathcal{A} = (L/K, \theta, \gamma) := L \oplus uL \oplus \ldots \oplus u^{n-1}L,$$

where $u \in \mathcal{A}$ is some auxiliary generating element of $\mathcal{A}$ satisfying the additional relations $xu = u\theta(x)$ for all $x \in L$ and $u^n = \gamma$ with n the degree of the algebra $\mathcal{A}$. Such algebra $\mathcal{A}$ is called a division algebra if every element $a \in \mathcal{A}$ has an inverse $a^{-1} \in \mathcal{A}$ such that $aa^{-1} = 1$.

The *non-norm* condition for γ (Proposition 1) states that the lowest power of γ that appears in $N_{L/K}(L)$ is γ^n. In [46], the authors prove that this condition on γ avoids the existence of the map $\tilde{\Theta}$ used in Bootland *et al.*'s attack [12].

Proposition 1 ("Non-Norm" Condition, adapted Proposition 3.5 from [81])**.** *The cyclic algebra $\mathcal{A} = (L/K, \theta, \gamma)$ of degree n is a division algebra if and only if the smallest positive $t \in \mathbb{Z}$ such that γ^t is the norm of some element of L is n. The element γ is referred as the non-norm element.*

It is worth noting that these results seem to be complementary to ours, and different division algebras could arise from considering our proposed multivariate number fields (see Sections 5 and 6) as the underlying field K in Definition 9.

5 Multiquadratic Rings

Let $K = \mathbb{Q}(\sqrt{d_i})$ be a field with prime d_i (hence squarefree) satisfying $d_i = 1 \bmod 4$; its ring of integers is $\mathcal{O}_K = \mathbb{Z}\left[\frac{1+\sqrt{d_i}}{2}\right]$ with basis $\{1, \frac{1+\sqrt{d_i}}{2}\}$ and discriminant $\Delta_K = d_i$, then we can also represent $\mathcal{O}_K$ as a polynomial ring $\mathbb{Z}[x]/x^2 - x + \frac{1-d_i}{4}$ ($\mathcal{O}_K$ is free of rank 2), according to (see Proposition 2):

Proposition 2 (Proposition 2.24 from [82])**.** *Let $K = \mathbb{Q}(\sqrt{d})$ be a quadratic field with d a squarefree integer. If $d \equiv 2, 3 (\mathrm{mod}\, 4)$, then $\mathcal{O}_K = \mathbb{Z}\left[\sqrt{d}\right] \simeq \mathbb{Z}[x]/(x^2 - d)$ and $\mathcal{O}_K$ is free of rank 2 over $\mathbb{Z}$ with basis $\{1, \sqrt{d}\}$. If $d \equiv 1 (\mathrm{mod} 4)$, then $\mathcal{O}_K = \mathbb{Z}\left[\frac{1+\sqrt{d}}{2}\right] \simeq \mathbb{Z}[x]/(x^2 - x + \frac{1-d}{4})$ and $\mathcal{O}_K$ is free of rank 2 over $\mathbb{Z}$ with basis $\{1, \frac{1+\sqrt{d}}{2}\}$.*

Let us extend the field to $\mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_l})$ (a multiquadratic field), with $\gcd(d_1, \ldots, d_l) = 1$, but still sticking to the case $d_i = 1 \bmod 4$. Taking into account that $\mathcal{O}_K \mathcal{O}_{K'} = \mathcal{O}_F$ when $\gcd(\Delta_K, \Delta_{K'}) = 1$, where F is the compositum over $\mathbb{Q}$ (see [62]) of two subfields $K = \mathbb{Q}(\sqrt{d_1})$ and $K' = \mathbb{Q}(\sqrt{d_2})$ (see [32]), we have that $\mathcal{O}_F = \mathbb{Z}\left[\frac{1+\sqrt{d_1}}{2}, \frac{1+\sqrt{d_2}}{2}\right]$. This can be generalized to the case of a field with l "coprime" squares, whose resulting ring of integers is

$$\mathcal{O}_K = \mathbb{Z}\left[\frac{1+\sqrt{d_1}}{2}\right] \cdot \ldots \cdot \mathbb{Z}\left[\frac{1+\sqrt{d_l}}{2}\right]. \tag{5}$$

Therefore, as all d_i are different primes, the discriminants of $\mathbb{Q}(\sqrt{d_i})$ are also coprime, which implies that the ring of integers can be expressed as the product of the respective univariate rings of integers.

However, the definition of RLWE (see Definition 8) works on the dual of the ring of integers, due to its geometric properties. The dual can be obtained through Theorem 4:

Theorem 4 (Theorem 3.7 **from [34])**. *Let $K = \mathbb{Q}(\alpha)$ and let $f(T)$ be the minimal polynomial of α in $\mathbb{Q}[T]$. Write*

$$f(T) = (T - \alpha)(c_0(\alpha) + c_1(\alpha)T + \ldots + c_{n-1}(\alpha)T^{n-1}),\ c_i(\alpha) \in K.$$

The dual basis to $\{1, \alpha, \ldots, \alpha^{n-1}\}$ relative to the trace product is

$$\left\{\frac{c_0(\alpha)}{f'(\alpha)}, \frac{c_1(\alpha)}{f'(\alpha)}, \ldots, \frac{c_{n-1}(\alpha)}{f'(\alpha)}\right\}.$$

In particular, if $K = \mathbb{Q}(\alpha)$ and $\alpha \in \mathcal{O}_K$ then

$$(\mathbb{Z} + \mathbb{Z}\alpha + \ldots + \mathbb{Z}\alpha^{n-1})^\vee = \frac{1}{f'(\alpha)}(\mathbb{Z} + \mathbb{Z}\alpha + \ldots + \mathbb{Z}\alpha^{n-1}).$$

Particularized to the quadratic case, Theorem 4 says that whenever the ring of integers has a power basis, the basis of the dual is

$$\left\{1, \frac{1+\sqrt{d_i}}{2}\right\}^\vee = \left\{\frac{1}{f'(\alpha)}, \frac{1}{f'(\alpha)}\frac{1+\sqrt{d_i}}{2}\right\}, \tag{6}$$

where $f(x) = x^2 - x + \frac{1-d}{4}$ and $\alpha = \frac{1+\sqrt{d}}{2}$, so $f'(x) = 2x - 1$; evaluated at $x = \alpha = \frac{1+\sqrt{d_i}}{2}$, it satisfies $f'(\alpha) = \sqrt{d_i}$.

As dual commutes tensoring, this result can be straightforwardly extended to the compositum case with several d_i. Additionally, we see that we can go from the dual to $\mathcal{O}_K$ by just scaling with $\sqrt{d_i}$ (or multiplying with the polynomial $2x - 1$).

Following our requirements, we need a ring of the form $\mathbb{Z}[x_1, \ldots, x_l]/(x_1^2 - d_1, \ldots, x_l^2 - d_l)$, which is an *order* of the field $\mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_l})$, but not necessarily its ring of integers and a Dedekind domain.[8] However, we can only base its security on RLWE defined on a number field of the form $\mathbb{Q}(\sqrt{d_1}, \ldots, \sqrt{d_l})$ (see Definition 7) and its ring of integers satisfying $\mathbb{Z}[x_1, \ldots, x_l]/(x_1^2 - x_1 + \frac{1-d_1}{4}, \ldots, x_l^2 - x_l + \frac{1-d_l}{4})$. We will therefore show that we can define an invertible map modulo q from the ring $\mathcal{O}_K$ (and its dual $\mathcal{O}_K^\vee$) to the ring $\mathbb{Z}[x_1, \ldots, x_l]/(x_1^2 - d_1, \ldots, x_l^2 - d_l)$, while still basing its security on the original RLWE formulation from Definition 7. Additionally, this map does not significantly increase the noise; in fact, it also decorrelates it in the coefficient domain, enabling direct sampling of the noise in the coefficient representation with an independent error distribution.

The map, applied to each variable x_i, works as follows:

- We apply the change of variable $x \to \frac{x+1}{2}$.
- We multiply the sample by a factor 2.

This mapping can be applied whenever the inverse of 2 exists modulo q. The multiplication by 2 is applied afterwards to avoid the potentially high distortion introduced by the factor $\frac{1}{2}$ into the noise.

Canonical Embedding Let $K = \mathbb{Q}(\sqrt{d})$, and note that $\frac{1}{2x-1}$ evaluated at $x = \frac{1+\sqrt{d}}{2}$ equals $\frac{1}{\sqrt{d}}$. We define the Embedding map $\mathcal{E}$ going from $\mathcal{O}_K^\vee \cong \frac{1}{\sqrt{d}}\mathbb{Z}[x]/x^2 - x + \frac{1-d}{4}$ to $\mathbb{C}^2$, as the substitutions $\{x \leftarrow \frac{1+\sqrt{d}}{2}, \sqrt{d} \leftarrow \sqrt{d}\}$ and $\{x \leftarrow \frac{1-\sqrt{d}}{2}, \sqrt{d} \leftarrow -\sqrt{d}\}$. This gives this transformation matrix for $\mathcal{E}$

$$\frac{1}{\sqrt{d}}\begin{pmatrix} 1 & \frac{1+\sqrt{d}}{2} \\ -1 & \frac{\sqrt{d}-1}{2} \end{pmatrix}. \tag{7}$$

[8] A recent work [10] discusses the hardness of a generalization of Ring-LWE called Order-LWE, which can be leveraged to have more freedom in the choice of the multivariate rings (see Section 9 for more details on the advantages of Order-LWE). We also refer the reader to [74] for a recent study on the connections between several algebraic LWE variants.

The inverse map $\mathcal{E}^{-1}$ is defined as the product with the matrix

$$\begin{pmatrix} \frac{\sqrt{d}-1}{2} & -\frac{1+\sqrt{d}}{2} \\ 1 & 1 \end{pmatrix}. \tag{8}$$

Sampling the error directly in the coefficient domain Finally, if we define the noise in the embedding of the dual ring as two independent Gaussian variables $e_0, e_1 \in \chi$ with variance σ^2, we have in the ring $\frac{1}{x}\mathbb{Z}[x]/x^2 - d$ after following the whole "processing chain":

$$\frac{1}{x}\left(\underbrace{(e_0+e_1)}_{2\sigma^2}x + \underbrace{\sqrt{d}(e_0-e_1)}_{2d\sigma^2}\right) \bmod x^2 - d.$$

Hence, the noise is not correlated in the coefficient domain and we can easily sample the error distribution considering an appropriate variance per coefficient.

For simplicity, we have focused on a quadratic field, but the embedding can be extended to the multiquadratic case by means of the Kronecker product.

Multiquadratic RLWE Let us define the multiquadratic version of m-RLWE, where all the modular functions have the form $f_i(x_i) = d_i + x_i^2$, as

Definition 10 (multivariate polynomial RLWE with quadratic modular functions)**.** *Given a multivariate polynomial ring $R_q^\vee[x_1, \ldots, x_l]$ with $f_j(x_j) = d_j + x_j^2$ for $j = 1, \ldots, l$ where $l = \log_2 n$ (with n a power of two) and an error distribution $\chi[x_1, \ldots, x_l] \in R_q^\vee[x_1, \ldots, x_l]$ that generates small-norm random multivariate polynomials in $R_q^\vee[x_1, \ldots, x_l]$, the multivariate polynomial RLWE relies upon the computational indistinguishability between samples $(a_i, b_i = a_i \cdot s + e_i)$ and (a_i, u_i), where $a_i \leftarrow R_q[x_1, \ldots, x_l]$, $u_i \leftarrow R_q^\vee[x_1, \ldots, x_l]$ are chosen uniformly at random from the rings $R_q[x_1, \ldots, x_l]$ and $R_q^\vee[x_1, \ldots, x_l]$; and $s, e_i \leftarrow \chi[x_1, \ldots, x_l]$ are drawn from the error distribution (see Section 5).*

The security reduction from Theorem 2 applies to this particular version of the m-RLWE problem whenever $-d_i = 1 \bmod 4$ and $\gcd(\Delta_K, \Delta_{K'}) = 1$. Section 7.1 gives further insights on the security and practicality of the chosen parameterization, and exemplifies it with a concrete instantiation. In particular, Proposition 6 gives a sufficient condition to consider the problem secure against known attacks.[9]

Comparison with Gaussian integers We now compare the multiquadratic RLWE with the particular case of power-of-two cyclotomics m-RLWE (see Definition 1) where all the used modular functions have the same form $f_i(x_i) = 1 + x_i^2$:

[9] It is worth mentioning that even when the Principal Ideal Problem is easy in multiquadratics [7], to the best of our knowledge, this has not been proven enough to solve RLWE.

Definition 11 (multivariate polynomial RLWE with $\Phi_4(\cdot)$ as modular functions)**.** *Given a multivariate polynomial ring $R_q[x_1, \ldots, x_l]$ with $f_j(x_j) = 1 + x_j^2$ for $j = 1, \ldots, l$ where $l = \log_2 n$ (with n a power of two) and an error distribution $\chi[x_1, \ldots, x_l] \in R_q[x_1, \ldots, x_l]$ that generates small-norm random multivariate polynomials in $R_q[x_1, \ldots, x_l]$, the multivariate polynomial RLWE relies upon the computational indistinguishability between samples $(a_i, b_i = a_i \cdot s + e_i)$ and (a_i, u_i), where $a_i, u_i \leftarrow R_q[x_1, \ldots, x_l]$ are chosen uniformly at random from the ring $R_q[x_1, \ldots, x_l]$; and $s, e_i \leftarrow \chi[x_1, \ldots, x_l]$ are drawn from the error distribution.*

The comparison of our secure multiquadratic RLWE samples with RLWE samples from Definition 11 is specially relevant, as the latter are severely affected by Bootland *et al.*'s attack. Samples from Definition 11 can be reduced to a dimension of 2, by applying the map $\tilde{\Theta}$ a total of $(\log_2 n - 1)$ times, yielding $n/2$ RLWE samples with $f(x) = 1 + x^2$ and error variance $n/2$ times higher than the original m-RLWE sample; this can be very easily solved. Consequently, despite of the efficiency of the polynomial operations on the rings instantiated according to Definition 11, they are not valid for cryptographic applications. Meanwhile, the samples from a secure instantiation of multiquadratic RLWE (Definition 10) preserve the lattice dimension n and withstand Bootland's attack.

Another advantage of the multiquadratic RLWE problem is that it also enables very efficient polynomial operations, without decreasing security. In particular, it is possible to apply a variant of the Fast Walsh-Hadamard transform (over finite rings instead of the usual real numbers), featuring a convolution property that relates the coefficient-wise representation with the transformed domain. This transform can be very efficiently computed with FFT-like algorithms (specifically, a variant of the Fast Walsh-Hadamard transform) whose computational cost is only $\mathcal{O}(n \log n)$ additions and $\mathcal{O}(n)$ products, hence considerably speeding up practical implementations. For more details, we refer the reader to Section 8, where we show how the well-known asymptotic cost of $\mathcal{O}(n \log n)$ for cyclotomic rings with polynomials of n coefficients can be improved by a factor of $\log n$ in terms of elemental multiplications.

6 More general multivariate rings

Let us consider now general fields $\mathbb{Q}(a_1^{1/n}, \ldots, a_l^{1/n})$, for which the a_i are square-free and coprime, but for simplicity we will assume that they are independent primes. The results shown in the previous section for multiquadratics cannot be straightforwardly generalized to these fields, as the individual univariate fields $\mathbb{Q}(a_i^{1/n})$ can easily have common factors in their discriminants (i.e., be non-coprime), in such a way that finding a basis for the multivariate ring of integers is not trivial.[10]

[10] We refer the reader to Section 9 for a discussion on the advantages that Order-LWE [10] brings about with respect to RLWE when choosing a basis for the ring of integers.

We explain the followed path that leads to our definition of valid, secure and easily parameterizable multivariate rings. We start by choosing number fields whose ring of integers $\mathcal{O}_K$ can be represented as $\mathbb{Z}[x]/x^n + ax + b$, that is, as polynomial rings whose modular function has the form $x^n + ax + b$. For this to be a valid ring $\mathcal{O}_K$ for K, it has to be irreducible over $\mathbb{Q}$, for which we use Eisensntein's criterion:

Proposition 3 (Eisenstein's criterion [6])**.** *The polynomial $p(x) = a_n x^n + a_{n-1}x^{n-1} + \ldots + a_1 x + a_0$, where $a_i \in \mathbb{Z}$ for all $i = 0, \ldots, n$ and $a_n \neq 0$ (which means that the degree of $p(x)$ is n) is irreducible if some prime number p divides all coefficients $a_0, \ldots, a_{n-1}$, but not the leading coefficient a_n and, moreover, p^2 does not divide the constant term a_0.*

Therefore, we impose the following two conditions on $f(x) = x^n + ax + b$:

- Both a and b have to be divisible by a prime p and not by p^2 (Eisenstein's criterion).
- If we choose b as a prime, a has to be divisible by b.

Now, we can compute the discriminant for this number field by resorting to [79, Chapter 2.7]:

Proposition 4 (An example of the calculation of a discriminant [79])**.** *For the calculation of Δ_K in a number field $K = \mathbb{Q}(x)$ being a extension of finite degree n of $\mathbb{Q}$ and $f(x) = x^n + ax + b$ the minimal polynomial of x over $\mathbb{Q}$, we obtain*

$$\Delta_K = (-1)^{\frac{n(n-1)}{2}}(n^n b^{n-1} + (-1)^{n-1}(n-1)^{n-1}a^n). \tag{9}$$

For $n = 2$ (respectively, 3) we rediscover the well-known expressions $a^2 - 4b$ (respectively, $-27b^2 - 4a^3$).

Theorem 5 (Theorem 8.11 **from [33]).** *For $\mathbb{Z}$-lattices $\mathcal{L}' \subset \mathcal{L}$ inside K, $[\mathcal{L}' : \mathcal{L}]^2 < \infty$ and*

$$disc_{\mathbb{Z}}(\mathcal{L}') = [\mathcal{L}' : \mathcal{L}]^2 \cdot disc_{\mathbb{Z}}(\mathcal{L}).$$

In particular, if $\mathcal{L}' \subset \mathcal{O}_K$ and the integer $disc_{\mathbb{Z}}(\mathcal{L}') \in \mathbb{Z} - \{0\}$ is squarefree then $[\mathcal{O}_K : \mathcal{L}'] = 1$; i.e., $\mathcal{L}' = \mathcal{O}_K$.

If we choose values for a and b such that Δ_K is squarefree, Theorem 5 guarantees that the ring of integers has a power basis of the form $\{1, \alpha, \alpha^2, \ldots\}$, with α a root of $x^n + ax + b$. Consequently, $\mathbb{Z}[x]/x^n + ax + b$ is a valid ring of integers.

By including more "univariate" subrings, $\mathbb{Z}[x_1, \ldots, x_l]/(x_1^n + a_1 x + b_1, \ldots, x_l^n + a_l x + b_l)$ becomes a valid ring of integers when all the discriminants are coprime [32]. Therefore, this is a feasible strategy to define RLWE over a multivariate ring, as the product of univariate rings with modular functions $x^n + a_i x + b_i$.[11]

[11] To define the dual $\mathcal{O}_K^\vee$ we can make use of Theorem 4 which states that whenever the ring of integers has a power basis, the basis of the dual is the same basis, scaled by $\frac{1}{f'(\alpha)} = \frac{1}{n\alpha^{n-1}+a}$, where α is a root of $f(x)$.

Finding valid parameters for $f(x) = x^n + ax + b$: Unfortunately, the two previous conditions (Eisenstein's criterion from Proposition 3 and Theorem 5) cannot be satisfied at the same time

- To satisfy the Eisenstein's criterion, b and a have to be divisible by at least a prime p (i.e., $\gcd(a, b) = u \cdot p$ for some $u \in \mathbb{Z}$), this introduces a factor p^{n-1} in Δ_K (see Equation (9)), in such a way that Δ_K is not squarefree and not satisfying $[\mathcal{O}_K : \mathcal{L}'] = 1$ in Theorem 5.
  We could still work with these multivariate rings provided that their discriminants are coprime, but it seems that there is no straightforward way to determine the "powerful" basis of the ring of integers: starting from Proposition 4, it is known that $\mathbb{Z}[\alpha] \subseteq \mathcal{O}_K \subseteq \frac{1}{\Delta_K}\mathbb{Z}[\alpha]$ where $f(\alpha) = 0$.
- Additionally, Eisenstein's criterion is a sufficient but *not necessary* condition for irreducibility of the modular functions. Without the imposed restrictions, we could search for squarefree and coprime discriminants, but we would have to verify the irreducibility of the involved functions case-by-case. Nevertheless, this is not impossible to find, as it is known that monogenic fields are not scarce [40]; in fact, for random polynomials f, it has been conjectured that $\mathbb{Z}[x]/f(x)$ of degree ≥ 4 is a ring of integers with probability $\gtrsim 0.307$ [53].

Transformation based on Modulus Switching Let us assume that we have found valid (monogenic) $x_i^n + a_i x_i + b_i$ functions defining the ring of integers $\mathbb{Z}[x_i]/x_i^n + a_i x_i + b_i$; they do not yet feature the desired $x^n + d$ form.

In order to achieve this, we consider a map from the original RLWE samples to RLWE samples modulo q, that removes the term ax if q divides a. It is worth noting that this transformation is nothing but a modulus switching to q, and if it were possible to break RLWE modulo q, the original secret key could be recovered or at least the indistinguishability assumption could be broken.

The trick relies on all the modular functions having the form $f_i(x_i) = x_i^n + \underbrace{a_i' q}_{a_i} x_i + b_i$. Hence, a reduction modulo q converts the modular functions into $f_i(x_i) = x_i^n + b_i$. We show the effect of this transformation on the ring $q\mathcal{O}_K^\vee$ for the univariate case (it extends to the multivariate case, as dual commutes tensoring):

- $\mathcal{O}_K^\vee$ is defined as $\frac{1}{f'(\alpha)}\mathcal{O}_K$; under the polynomial ring $\mathbb{Z}[x]/x^n + a_i' q x + b_i$, this implies that the dual is $\frac{1}{nx^{n-1} + a_i' q}\mathbb{Z}[x]/x^n + a_i' q x + b_i$.
- After reducing modulo q, we obtain $\frac{1}{nx^{n-1}}\mathbb{Z}_q[x]/x^n + b_i$; considering that x has inverse modulo q, we can multiply numerator and denominator by x to obtain $\frac{x}{nx^n} = \frac{x}{-nb_i}$.
- The factor $\frac{1}{-nb_i}$ can be removed by just a scaling (moving to the ring of integers $\mathcal{O}_K$), so we can directly work on $\mathbb{Z}_q[x]/x^n + b_i$. This gives a "basis" $\{b_i, x, x^2, \ldots, x^{n-1}\}$ (or a basis $\{\frac{1}{n}, \frac{x}{nb_i}, \frac{x^2}{nb_i}, \ldots, \frac{x^{n-1}}{nb_i}\}$ without scaling).

Decodability of the transformed $x^n + ax + b$: Elias *et al.* [40] use an heuristic perturbation method to bound the spectral norm of the canonical embedding with $f(x) = x^n + ax + b$. As the condition number is stable for most of the random perturbations of the canonical embedding matrix associated to $x^n + 1$, they conjecture that many f functions have a bounded spectral norm in terms of a and b; therefore, we can consider that the spectral norm $s_1(N_f)$ (N_f represents the inverse of the canonical embedding matrix) is likely bounded by $\sqrt{\max(a, b)} \cdot \det(N_f)^{1/n}$ [19]. Consequently, the same arguments about noise behavior in [19, 73] still apply, and in order to guarantee the prevalence of the security reduction (see Proposition 6), the noise wraps around modulo q in some of the polynomial coefficients ($\max(a, b) \approx q$). This is due to the large q factor introduced in $f(x)$, which requires the use of a high error variance, rendering some of the polynomial coefficients modulo q useless. This makes these RLWE samples harder to use for cryptographic applications.

Valid and practical parameterizations for Multivariate Rings The previous solutions to parameterize multivariate rings with modular functions $x^n + d$ are not satisfactory, as (a) the search of valid univariate rings is not easy to handle (due to the impossibility to use Eisenstein's criterion) and (b) the obtained samples are not practical for cryptographic applications due to their high noise in some polynomial coefficients.

Here we follow a slightly different approach, releasing the condition on equal-degree modular functions; that is, we consider multivariate rings as $\mathbb{Z}[x_1, \ldots, x_l]/(x_1^{n_1} + d_1, \ldots, x_l^{n_l} + d_l)$. Again, to simplify the explanation we only consider an univariate ring with modular function $x^n + d$, but all the results can be analogously extended to the multivariate case (see Section 5) by requiring coprime discriminants.

First, for $f(x) = x^n + d$, Equation (9) simplifies to $\Delta_K = (-1)^{\frac{n(n-1)}{2}} n^n d^{n-1}$. Let d be a prime number and $n = u^m$ a prime power. Then,

- $f(x)$ is an irreducible polynomial over $\mathbb{Q}$ by the Eisenstein's criterion (Propostion 3).
- $f(x)$ is monogenic for d and n satisfying the following Proposition 5.

Proposition 5 (Adapted Proposition 3 from [40])**.** *Let n be a power of a prime u. If d is squarefree and u^2 does not divide $(-1)^n(d^{n-1} + 1)d$, then the polynomials $x^n + d$ are monogenic.*

Proposition 5 shows that $f(x)$ can be monogenic even when its discriminant is not squarefree. If $f(x)$ satisfies Proposition 5, we have $\mathcal{O}_K = \mathbb{Z}[x]/x^n + d$ and $\mathcal{O}_K^\vee = \frac{1}{nx^{n-1}}\mathbb{Z}[x]/x^n + d$.

In order to extend these results to multivariate rings $\mathbb{Z}[x_1, \ldots, x_l]/(x_1^{n_1} + d_1, \ldots, x_l^{n_l} + d_l)$, we only have to consider functions $\{x_1^{n_1} + d_1, \ldots, x_l^{n_l} + d_l\}$ satisfying Proposition 5 and having coprime discriminants. This basically means that all the d_i and n_i are respectively different primes and power primes.

Analogously to the *multiquadratic rings* in Section 5, we can directly map the error distribution in the coefficient domain. In particular, for the ring $\frac{1}{n_i x^{n_i-1}}\mathbb{Z}[x_i]/x_i^{n_i}+d_i$, the parameter for the error distribution in the $(j-1)$-th coefficient $(1 \leq j \leq n_i)$ is given by $\sqrt{n_i} d_i^{\frac{n_i-j}{n_i}} r$, where r is the parameter of an independent spherical error distribution in the embedding domain [19]. This extends to multivariate rings by means of the Kronecker product. As the resulting embedding matrix is the Kronecker product of the embedding matrices associated to each univariate ring, the singular values are the result of the Kronecker product of the singular values for each univariate embedding matrix.

Finally, we introduce the definition of multivariate RLWE with the proposed modular functions $f_i(x_i) = d_i + x_i^{n_i}$:

Definition 12 (multivariate RLWE with modular functions as $x_i^{n_i}+d_i$)**.** *Given a multivariate polynomial ring $R_q[x_1,\ldots,x_l]$ with $f_j(x_j) = d_j + x_j^{n_j}$ for $j = 1,\ldots,l$ where $n = \prod_j n_j$ (where all n_j are prime powers) and an error distribution $\chi[x_1,\ldots,x_l] \in R_q^\vee[x_1,\ldots,x_l]$ that generates small-norm random multivariate polynomials in $R_q^\vee[x_1,\ldots,x_l]$, the multivariate polynomial RLWE relies upon the computational indistinguishability between samples $(a_i, b_i = a_i \cdot s + e_i)$ and (a_i, u_i), where $a_i \leftarrow R_q[x_1,\ldots,l]$, $u_i \leftarrow R_q^\vee[x_1,\ldots,l]$ are chosen uniformly at random from the rings $R_q[x_1,\ldots,x_l]$ and $R_q^\vee[x_1,\ldots,x_l]$; $s, e_i \leftarrow \chi[x_1,\ldots,x_l]$ are drawn from the error distribution.*

For the ring $R^\vee[x_1,\ldots,x_l]$, we define $\chi[x_1,\ldots,x_l]$ as the distribution generating polynomials belonging to $R^\vee[x_1,\ldots,x_l]$ and whose parameter per coefficient satisfies $r\prod_{i\in[l]}\sqrt{n_i} d_i^{\frac{n_i-j_i}{n_i}}$, where $1 \leq j_i \leq n_i$ and $1 \leq i \leq l$, and hence represents the parameter for the coefficient associated to the monomial $x_1^{j_1-1}\cdot\ldots\cdot x_l^{j_l-1}$.

Some examples of valid parameters: In order to show the feasibility of the proposed parameterization, we exemplify it with some practical use cases for bivariate RLWE; we will consider $n_1 = 2^{11} = 2048$ and $n_2 = 3^7 = 2187$, and $d_1 = 5$, $d_2 = 7$, for which we prove that they meet the conditions of Proposition 5

- $2^2 = 4$ does not divide $5^{2047}+1$, or equivalently, $5^{2047}+1 \neq 0 \bmod 4$. We have $5^{2047}+1 \bmod 4 = 1^{2047}+1 = 2 \neq 0$.
- $3^2 = 9$ does not divide $7^{3^7-1}+1$, or equivalently, $7^{3^7-1}+1 \neq 0 \bmod 9$. We have $7^{3^7-1}+1 = 7^{-1}7^{3^7}+1 = 7^{-1}7^{3^7 \bmod 6}+1 = 7^2+1 = 50 = 5 \bmod 9 \neq 0$.

Consequently, with this choice of parameters we can work on the number field $K = \mathbb{Q}((-5)^{1/2048}, (-7)^{1/2187})$, with $\mathcal{O}_K = \mathbb{Z}[x,y]/(x^{2048}+5, y^{2187}+7)$ and $\mathcal{O}_K^\vee = \frac{1}{4478976x^{2047}y^{2186}}\mathcal{O}_K$.

As for the example mentioned in the introduction, with functions $x^{64}+1$ and $y^{27}+5$, we can also verify that

- $x^{64}+1$ is the $\Phi_{128}(x)$ power-of-two cyclomic, hence it is monogenic.
- $y^{27}+5$ is monogenic by Proposition 5, as $3^2 = 9$ does not divide 5 or $5^{26}+1$.

Additionally, as both discriminants are coprime, the product is directly the corresponding ring of integers.

7 Security of multivariate RLWE and example instantiations

This section includes a discussion on several aspects of the proposed solutions in this work, namely their security, the geometric interpretation of the problem, and the feasibility of the proposed parameterizations. With this purpose, we enumerate the known attacks in the literature and include an example instantiation of a simple bivariate RLWE scheme. We refer to next Sections and the Appendix for a description of the applications enabled by our constructions.

7.1 Resilience against known attacks

The formulation proposed in this work involves working with rings whose modular function is x^n+d or, more generally, x^n+ax+b. Some particular instantiations of these rings have already been studied in the literature and we can find specific attacks to "variants" of the RLWE problem (e.g., PLWE together with non-dual and dual RLWE versions) defined over them.

In general, the known attacks can be divided in two main types [73]:

- Attacks using a reduction modulo an *ideal divisor* $\mathfrak{q}$ of the modulus qR [39, 54, 40, 22–24]. These attacks try to distinguish between the error distribution and the uniform distribution modulo an *ideal divisor*.
- A reduction to *errorless* LWE [19] which exploits the relation between RLWE and LWE. Expressing RLWE in its LWE form, the error term of some of the equations can be removed by means of a rounding operation, and linear algebra can be used to search for the secret key.

All these attacks have been generalized and studied in depth by Peikert in [73], where he concludes that *all the concrete insecure RLWE instantiations made use of error distributions which were insufficiently well spread relative to the rings*, meaning that none of the vulnerable instantiations satisfy the conditions from Theorem 2 to have worst-case hardness. In [73], Peikert also gives sufficient conditions to make RLWE secure against the previous attacks. We summarize the main relevant results for our constructions.

Proposition 6 (Invulnerability condition from [73])**.** *Let $\psi = D_r$ (see Definition 7) be a spherical Gaussian error distribution over $K_{\mathbb{R}}$ for some $r > 0$; a sufficient condition for invulnerability to the attacks from [39, 54, 40, 22, 19, 73, 23] is*

$$r \geq 2. \tag{10}$$

The validity of Proposition 6 to resist the previous attacks is shown in the following two theorems: Theorem 6 (for the attack based on reduction modulo an ideal divisor) and Theorem 7 (for the attack based on errorless LWE).

Theorem 6 (Theorem 5.2 **from [73])**. *Given a Ring-LWE sample $(a, b = s \cdot a + e) \in R_q \times K_{\mathbb{R}}/qR^\vee$ where $e \leftarrow D_r$ is transformed into n LWE samples*

$(\boldsymbol{A}_a, \boldsymbol{b} = \boldsymbol{s}^T\boldsymbol{A}_a + \boldsymbol{e}^T)$, *where* $\boldsymbol{b} \in (\mathbb{R}/q\mathbb{Z})^n$ *and* $\boldsymbol{e} \in \mathbb{R}^n$ *are respectively the coefficient vectors of* $b \in K_\mathbb{R}/qR^\vee$ *and* $e \in K_\mathbb{R}$ *(with respect to the chosen basis of* $R^\vee$*), and* $\boldsymbol{A}_a \in \mathbb{Z}_q^{n\times n}$ *is the matrix of multiplication by* $a \in R_q$ *with any element of* $R_q^\vee$ *(with respect to the chosen bases of* $R, R^\vee$*). Then, for any* $\mathbb{Z}$*-basis* $B^\vee = (b_j^\vee)$ *of* $R^\vee$ *used above, each entry of* $\boldsymbol{e}$ *is a continuous Gaussian of parameter at least* $r\sqrt{n} \geq 2\sqrt{n}$ *(which is the required lower bound to apply the worst-case hardness theorems for plain-LWE).*

Theorem 7 (Theorem 5.1 from [73]). *Let* $\mathfrak{q} \subseteq R$ *be any ideal of norm* $N(\mathfrak{q}) \leq 2^n$*, and let the error parameter* $r \geq 2$ *satisfy condition* (10)*. Then the reduced error distribution* $D_r \bmod \mathfrak{q}R^\vee$ *is within statistical distance* 2^{-2n} *of uniform over* $K_\mathbb{R}/\mathfrak{q}R^\vee$*.*

7.2 Geometric interpretation and examples of multivariate RLWE

In this section, we give a high level overview of how to instantiate a secure multivariate RLWE sample from Definition 12, exemplifying it in the bivariate case (all rings are defined over variables x, y, omitted when unambiguous).

This example can also be used as a means to showcase complex numbers packing into slots, obtaining a net improvement on the number of available slots per ciphertext when comparing to the recent results in [26] (see Section 10). For the sake of clarity, we introduce a simple SHE scheme which enables homomorphic additions and multiplications without taking into account some of the more advanced techniques typically considered in the literature (see Appendix C for a brief explanation of the proposed optimizations).

A multivariate RLWE sample For simplicity, we consider a bivariate RLWE sample $(a, b = a \cdot s + e) \in R_q \times R_q^\vee$, where $a \in R_q[x, y]$, $s \in R_q^\vee[x, y]$ and $e \leftarrow \chi[x, y] \in R^\vee[x, y]$. We can use a uniformly random s or follow conventional approaches where s is a small key (see Section 3).

Geometry of R*, its dual* $R^\vee$ *and an example for* $\{x^2+3, y^2-5\}$ To easily illustrate the geometry of R and $R^\vee$, we use a simple example $R = \mathbb{Z}[x, y]/(x^2+3, y^2-5)$. By means of the canonical embedding, we know that the substitutions $\{x \leftarrow \pm\sqrt{-3}, y \leftarrow \pm\sqrt{5}\}$ yield the four different *slots* in the embedding domain.

This clearly shows that $\lambda_1(R) \leq \sqrt{n} = 2$ by the embedding of 1, and we can also obtain the embedding of the elements x, y and xy. The term xy can be used to obtain an upper-bound for $\lambda_4(R)$, such that $\lambda_4(R) \leq 2\sqrt{15}$.

This is generalizable to any multiquadratic with $l = \log_2 n$ variables, by considering the embedding of 1 and $\prod_{i\in[l]} x_i$, obtaining $\lambda_1(R) \leq \sqrt{n}$ and $\lambda_n(R) \leq \sqrt{n}\prod_{i\in[l]}\sqrt{d_i}$. As the l-th prime is asymptotically $p_l \sim l \log l$, a worst-case for $l = \log_2 n$ is $d_l^l \sim l^l(\log l)^l = (\log_2 n)^{\log_2 n}(\log_2 \log_2 n)^{\log_2 n}$. Combining the two previous expressions we have that $\lambda_n(R)$ (and hence also the ratio $\frac{\lambda_n R}{\lambda_1(R)}$) is polynomially upper-bounded by n.

These bounds are straightforwardly extended to the dual $R^\vee$ by taking into account the corresponding "tweak" factor. For the multiquadratic scenario, the dual only suffers a scaling by the square roots of the d_i terms (R is sparser than the dual $R^\vee$). However, considering higher degrees in the modular functions $x_i^{n_i} + d_i$, the tweak factor can turn the noise in the non-dual version of RLWE into highly non-spherical.

A very detailed analysis of these effects (including also some enlightening visual examples) can be found in [73].

Parameters' choice We show now how to select correct parameters $\{n_x, n_y, d_x, d_y\}$ satisfying the conditions established in Sections 5 and 6 for valid number fields.

As a brief summary, and focusing on $n_x, n_y > 2$, this mainly implies that: (1) the discriminants of $K_x = \mathbb{Q}[x]/x^{n_x} + d_x$ and $K_y = \mathbb{Q}[y]/y^{n_y} + d_y$ are coprime, i.e., $\gcd(\Delta_{K_x}, \Delta_{K_y}) = 1$, and (2) n_x, n_y are prime powers satisfying Proposition 5.

This enables the definition of $\mathcal{O}_K = R = \mathbb{Z}[x,y]/(x^{n_x} + d_x, y^{n_y} + d_y)$ as the ring of integers. Analogously, the dual is $\mathcal{O}_K^\vee = \frac{1}{n_x n_y x^{n_x-1} y^{n_y-1}} \mathbb{Z}[x,y]/(x^{n_x} + d_x, y^{n_y} + d_y)$ (see Section 6 for some particular choices).

In this bivariate case, the error distribution $\chi[x,y]$ samples polynomials in $\mathcal{O}_K^\vee$ whose coefficients are independently sampled from Gaussian distributions with different standard deviations. In particular, σ is equal to $r\sqrt{n} d_x^{\frac{n_x - j_x}{n_x}} d_y^{\frac{n_y - j_y}{n_y}}$ for the coefficient associated to the monomial $x^{j_x-1} y^{j_y-1}$ with $1 \leq j_x \leq n_x$ and $1 \leq j_y \leq n_y$.

Working on $q\mathcal{O}_K$ As it is usually done with power-of-two cyclotomics, we can directly transform the dual into the ring of integers by means of a scaling. If we have $\mathcal{O}_K^\vee = \frac{1}{n_x n_y x^{n_x-1} y^{n_y-1}} \mathbb{Z}[x,y]/(x^{n_x} + d_x, y^{n_y} + d_y)$, we can first multiply the dual by $\frac{xy}{xy}$, to see the simplified relation $\frac{xy}{xy}\mathcal{O}_K^\vee = \frac{xy}{n d_x d_y} \mathcal{O}_K$.

Finally, analogously to the $x^n + 1$ functions, we can scale the (a,b) sample by $n = n_x n_y$ and also $d_x d_y$. This gives us a sample $(a(x,y), b'(x,y) = n d_x d_y x y b(x,y)) \in R_q^2$. Consequently, we can directly work on the ring of integers with $(a, b = as + e) \in R_q^2$ where $a \leftarrow R_q$, $s \leftarrow R_q$ (or also $s \leftarrow \chi[x,y]$) and $e \leftarrow \chi[x,y]$. After the multiplication with the monomial xy, the error distribution $\chi[x,y]$ generates independent coefficients from a Gaussian distribution of $\sigma = r\sqrt{n} d_x^{\frac{n_x - j_x}{n_x}} d_y^{\frac{n_y - j_y}{n_y}}$ for $1 < j_x \leq n_x$ and $1 < j_y \leq n_y$, $\sigma = r\sqrt{n} d_x^{\frac{2n_x - j_x}{n_x}} d_y^{\frac{n_y - j_y}{n_y}}$ for $j_x = 1$ and $1 < j_y \leq n_y$, $\sigma = r\sqrt{n} d_x^{\frac{n_x - j_x}{n_x}} d_y^{\frac{2n_y - j_y}{n_y}}$ for $1 < j_x \leq n_x$ and $j_y = 1$ while $\sigma = r\sqrt{n} d_x^{\frac{2n_x - 1}{n_x}} d_y^{\frac{2n_y - 1}{n_y}}$ for $j_x = j_y = 1$.

SHE over Multivariate Rings The basic example cryptosystem described in Table 1 follows the structure of the SHE version introduced in [16] and implemented in [56]. The main difference relies on the fact that our polynomial elements belong to the multivariate rings $R[x,y]$, $R_t[x,y]$ and $R_q[x,y]$ (see Definition 12),

Table 1. Parameters and Primitives of a Somewhat Homomorphic Cryptosystem based on a multivariate version of RLWE (see [68, 71])

Parameters		
Let $R_t[x,y]$ be the cleartext ring and $R_q[x,y]$ the ciphertext ring. The noise distribution $\chi[x,y]$ in $R_q[x,y]$ takes its coefficients from a spherically-symmetric truncated i.i.d Gaussian $\mathcal{N}(\mathbf{0}, r^2\boldsymbol{J}^2)$. q is an integer satisfying $t < q$ and is relatively prime to t. All the previous parameters are chosen in terms of the security parameter λ where $n = 2^{\lfloor \log \lambda \rceil - 1}$		
Example SHE Cryptographic Primitives		
SH.KeyGen	Process	$s, e \leftarrow \chi[x,y]$, $a_1 \leftarrow R_q[x,y]$; $sk = s$ and $pk = (a_0 = -(a_1 s + te), a_1)$
SH.Enc	Input	$pk = (a_0, a_1)$ and $m \in R_t[x,y]$
	Process	$u, f, g \leftarrow \chi[x,y]$ and the fresh ciphertext is $\boldsymbol{c} = (c_0, c_1) = (a_0 u + tg + m, a_1 u + tf)$
SH.Dec	Input	sk and $\boldsymbol{c} = (c_0, c_1, \ldots, c_{\gamma-1})$
	Process	$m = \left(\left(\sum_{i=0}^{\gamma-1} c_i s^i\right) \mod q\right) \mod t$
SH.Add	Input	$\boldsymbol{c} = (c_0, \ldots, c_{\beta-1})$ and $\boldsymbol{c'} = (c'_0, \ldots, c'_{\gamma-1})$
	Process	$\boldsymbol{c}_{add} = (c_0 + c'_0, \ldots, c_{\max(\beta,\gamma)-1} + c'_{\max(\beta,\gamma)-1})$
SH.Mult	Input	$\boldsymbol{c} = (c_0, \ldots, c_{\beta-1})$ and $\boldsymbol{c'} = (c'_0, \ldots, c'_{\gamma-1})$
	Process	Using a symbolic variable v their product $\boldsymbol{c''}$ can be obtained from the relation $\left(\sum_{i=0}^{\beta-1} c_i v^i\right) \cdot \left(\sum_{i=0}^{\gamma-1} c'_i v^i\right) = \sum_{i=0}^{\beta+\gamma-2} c''_i v^i$

contrarily to the traditional univariate version $\mathbb{Z}[x]/1+x^n$ and its analogous rings modulo t and q. In Table 1 the diagonal of $\boldsymbol{J}$ has the corresponding standard deviations of χ normalized by r (i.e., σ/r) for each coefficient of the bivariate polynomials.

In particular, our plaintext ring R_t is basically a bivariate polynomial $R_t[x,y] = \mathbb{Z}_t[x,y]/(x^{n_x} + d_x, y^{n_y} + d_y)$ which is encoded as a sub-module of $\mathbb{T} = K_{\mathbb{R}}/R^{\vee}$ (see Definition 7). Our example is based on the scheme introduced in [16], but other choices are possible, and we briefly discuss them in the Appendix C. Regarding the achieved noise bounds, they are analogous to the computations from [16] by taking into account the expansion factor of the involved rings.

The additional variables of the multivariate structure can bring about some significant advantages: more efficient polynomial operations (see Section 8), better space/efficiency tradeoffs when working with automorphisms (see Section 9), and can also be very useful when working with multidimensional structures (see Appendix A and also the works [68, 71, 26] for more details on practical applications). In particular, in [26, 28] the authors present a library called MHEAAN, based on multivariate RLWE, which is optimized to perform homomorphic matrix operations.

Correctness and Security The condition for correct decryption is that the effective noise $||(\sum_{i=0}^{\gamma-1} c_i s^i) \bmod q)||_{\infty}$ remains smaller than $q/2$. Let us consider a simplified version of Theorem 2 from [16] where only the effect of noise is taken into account, and let $\max\{\sigma\}$ be the maximum standard deviation of the polynomials sampled from $\chi[x,y]$. Let M be the maximum coefficient of the evaluated degree-D polynomial; if $M(t \max\{\sigma\} d_x d_y n\sqrt{n})^D$ is smaller than $q/2$, the scheme of Table 1 can evaluate degree-D multivariate polynomials over el-

ements which belong to $R_t[x, y]$. We could also consider a more tight empirical condition for q, as stated in [56].

Regarding the security of this SHE scheme, it relies on the indistiguishability assumption of the polynomial multivariate version of RLWE (with *adequately chosen secure parameters* $\chi[x, y]$, $\{n_x, d_x, n_y, d_y\}$ and q) featured in Definition 12; breaking this assumption implies, as stated in Theorem 2, the existence of a quantum algorithm which solves short vector problems over ideal lattices. For a practical estimation of the bit security, we can apply the LWE security estimator developed by Albrecht el al. [3, 2] to the cryptosystems built on multivariate RLWE and also the estimates included in the standards document [20] for a general random lattice with the same dimension ($n = \prod n_i$). This is plausible, analogously to what it is typically done with ideal lattices, as a secure instantiation of m-RLWE works with full-rank lattices, for which no substantially faster attacks are known than for general lattices.

8 Multiquadratic Rings with Fast Walsh Hadamard Transforms

This section focuses on improving the cost of the underlying polynomial operations for cryptographic primitives based on RLWE, especially polynomial products (convolutions). We show how the well-known asymptotic cost of $\mathcal{O}(n \log n)$ for cyclotomic rings with polynomials of n coefficients can be improved by a factor of $\log n$ in terms of elemental multiplications when working on m-RLWE (or RLWE over a multivariate number field). To this aim, we particularize the multivariate version to degree-2 polynomials and introduce an ($\boldsymbol{\alpha}$-generalized) variant of the Walsh-Hadamard transform (over finite rings instead of the usual real numbers), featuring a convolution property that relates the coefficient-wise representation with the transformed domain. This transform can be very efficiently computed with FFT algorithms (specifically, with a variant of the Fast Walsh-Hadamard transform) whose computational cost is only $\mathcal{O}(n \log n)$ additions, hence being much more amenable for a practical implementation. It is worth noting that the effect of the efficiency improvement brought about by our approach goes beyond somewhat homomorphic encryption schemes (including also the NTRU setting [52, 8]), also enhancing any cryptographic primitives involving polynomial multiplications, e.g., the candidates of the NIST Post-Quantum challenge [2].

We start by introducing the notation used in this section. Polynomials are denoted with regular lowercase letters, omitting the polynomial variable (e.g., a instead of $a(x)$) when there is no ambiguity. We follow a recursive definition of multivariate modular rings: $R_q[x_1] = \mathbb{Z}_q[x_1]/f_1(x_1)$ denotes the polynomial ring in the variable x_1 modulo $f_1(x_1)$ with coefficients belonging to $\mathbb{Z}_q$. Analogously, $R_q[x_1, x_2] = (R_q[x_1])[x_2]/(f_2(x_2))$ is the bivariate polynomial ring with coefficients belonging to $\mathbb{Z}_q$ reduced modulo $f_1(x_1)$ and $f_2(x_2)$. In general, $R_q[x_1, \ldots, x_l]$ (resp. $R[x_1, \ldots, x_l]$) represents the multivariate polynomial ring with coefficients in $\mathbb{Z}_q$ (resp. $\mathbb{Z}$) and the l modular functions $f_i(x_i)$ with

$1 \leq i \leq l$. The polynomial a can also be denoted by a column vector $\boldsymbol{a}$ whose components are the corresponding polynomial coefficients.

For this section, we deal with a specific version of m-RLWE where all the used modular functions have the same form $f_i(x_i) = d_i + x_i^2$ (see Definition 10).

The security reduction from Theorem 2 applies to this particular version of the m-RLWE problem. To this aim, parameteres d_i have to be chosen as indicated in the beginning of Section 5. Additionally, Proposition 6 gives a sufficient condition to make the problem secure against the attacks described in Section 7.1.

After defining the specific version of the problem, we introduce the ($\boldsymbol{\alpha}$-generalized) Hadamard transform, that we apply to reach the aforementioned performance gains on polynomial convolutions.

8.1 Faster polynomial arithmetic over multivariate rings

The Hadamard transform over real numbers is usually applied by means of the recursion

$$\boldsymbol{H}_i = \frac{1}{\sqrt{2}} \begin{pmatrix} \boldsymbol{H}_{i-1} & \boldsymbol{H}_{i-1} \\ \boldsymbol{H}_{i-1} & -\boldsymbol{H}_{i-1} \end{pmatrix}, \tag{11}$$

where $i \in \mathbb{N}$ and $\boldsymbol{H}_0 = 1$.

It can be seen that the matrix $\boldsymbol{H}_i$ with $i \geq 1$ is equivalent to the Kronecker product of i DFT (Discrete Fourier Transform) matrices of size 2 ($\boldsymbol{H}_1$ equals the DFT matrix of size 2); that is, it can be seen as a $\underbrace{2 \times 2 \times \cdots \times 2}_{i \text{ times}}$-DFT transform (defined over i dimensions of length 2 each).

Analogously to the DFT, the Walsh Hadamard Transform (WHT) of size n possesses a particular fast algorithm called FWHT (Fast Walsh Hadamard Transform) which can be very efficiently computed with no products and with a cost of $\mathcal{O}(n \log n)$ additions and subtractions (see [42, 83]). Hence, when working over rings satisfying a convolution property with the Hadamard transform, it is possible to efficiently compute the multiplication of elements from these rings with a cost of $\mathcal{O}(n)$ elemental multiplications.

Security reasons prevent us from directly working over rings satisfying this convolution property with the Walsh Hadamard transform (that is, multivariate rings whose modular functions are $f(x_i) = x_i^2 - 1$), as they can be easily factored into $(x_i - 1)(x_i + 1)$ over $\mathbb{Z}$. Therefore, we resort to the type of multivariate rings presented in Definition 10 and apply an ($\boldsymbol{\alpha}$-generalized) variant of the WHT.

α-generalized convolutions An α-generalized convolution[12] corresponds to the ring operation defined over polynomials belonging to $\mathbb{Z}_q[z]/1 - \alpha z^n$. Figure 1 shows the realization of an α-generalized convolution between vectors of length N (with $l = 0, \ldots, N-1$), by means of a cyclic convolution combined with an element-wise pre/post-processing applied before/after [64, 67].

[12] For example, with $\alpha = -1$ we have a negacyclic convolution. In the literature, this convolution operation is also called negative wrapped convolution.

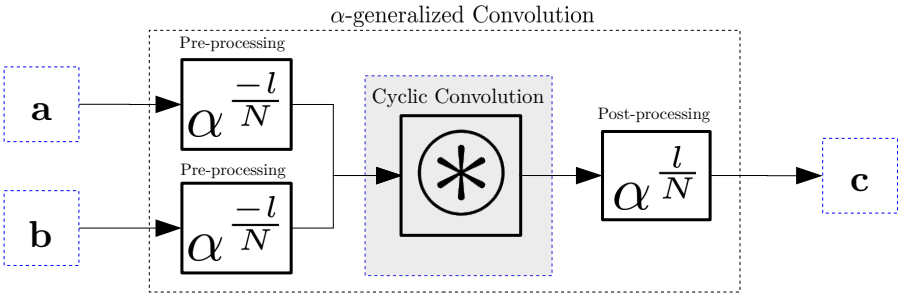

Fig. 1. Block diagram for the implementation of an α-generalized convolution.

As the cyclic convolution can be efficiently computed by means of standard fast algorithms, this means that an α-generalized convolution can be implemented with only a light overhead ($\mathcal{O}(n)$ scalar products). [13]

$\boldsymbol{\alpha}$-generalized Walsh Hadamard transform We are mainly interested in modular functions with the form $x_i^2 + d_i$. We can rewrite $1 - \alpha x^n$ as $-\alpha((-\alpha)^{-1} + x^n)$. Hence for $x_i^2 + d_i$ we have $d_i = (-\alpha_i)^{-1} = -\alpha_i^{-1}$. For this particular type of polynomial rings we can define the following direct and inverse transforms:

$$\boldsymbol{W}_1 = \boldsymbol{H}_1 \begin{pmatrix} 1 & 0 \\ 0 & (\alpha_1)^{-1/2} \end{pmatrix}, \quad \text{and} \quad \boldsymbol{W}_1^{-1} = 2^{-1} \begin{pmatrix} 1 & 0 \\ 0 & (\alpha_1)^{1/2} \end{pmatrix} \boldsymbol{H}_1,$$

where the square-roots $(\alpha_i)^{\frac{1}{2}}$ and $(\alpha_i)^{\frac{-1}{2}}$ have to exist in R_q for all i (see Definition 10). Equivalently, if q is an odd prime, we can make use of the Legendre symbol $\left(\frac{-d \bmod p}{p}\right)$ to check when the multivariate ring factors into linear terms. To this aim we need that $\left(\frac{-d_i \bmod q}{q}\right) = 1$ for a prime q and for all i. We also redefine $\boldsymbol{H}_1 = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ without taking into account the normalizing factor $\frac{1}{2}$.

Therefore, now we can extend this definition to multivariate rings with modular functions of the form $x_i^2 + d_i$: we consider the Kronecker product of the matrices $\boldsymbol{W}_1$ and $\boldsymbol{W}_1^{-1}$ as $\boldsymbol{W}_i = \bigotimes_{j \in [i]} \boldsymbol{W}_1$ and $\boldsymbol{W}_i^{-1} = \bigotimes_{j \in [i]} \boldsymbol{W}_1^{-1}$, arriving to the following expression:

$$\boldsymbol{W}_i = \boldsymbol{H}_i \left(\bigotimes_{j \in [i]} \begin{pmatrix} 1 & 0 \\ 0 & (\alpha_j)^{-1/2} \end{pmatrix} \right), \quad \text{and} \quad \boldsymbol{W}_i^{-1} = 2^{-i} \left(\bigotimes_{j \in [i]} \begin{pmatrix} 1 & 0 \\ 0 & (\alpha_j)^{1/2} \end{pmatrix} \right) \boldsymbol{H}_i,$$

where the normalizing factors are again left outside $\boldsymbol{H}_i$.

Consequently, if we define the vector $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_l)^T$, when working over the multivariate ring $R_q[x_1, \ldots, x_l]$ with $f_j(x_j) = d_j + x_j^2$ for $j = 1, \ldots, l$ we can

[13] It is common to include these additional scalar products inside the butterflies of the FFT algorithms to further enhance the efficiency.

use the transforms $\boldsymbol{W}_l$ and $\boldsymbol{W}_l^{-1}$ analogously to the use of negacyclic NTTs in the univariate RLWE. Both $\boldsymbol{W}_l$ and $\boldsymbol{W}_l^{-1}$ transforms can be efficiently computed in $\mathcal{O}(n)$ (where $n = 2^l$) elemental multiplications thanks to the FWHT. This enables the computation of the $\boldsymbol{H}_l$ matrix multiplications with only $\mathcal{O}(n \log n)$ additions and subtractions and no products, which brings a net improvement with respect to the analogous and wide-spread radix implementation of the NTT.

Implementation of the Fast Walsh-Hadamard Transform (FWHT) Algorithm 1 shows a pseudocode implementation of the (cyclic) FWHT (Fast Walsh-Hadamard Transform) implementation (see [42, 83]), computing the Hadamard transform of a length-n vector $\boldsymbol{a}$. It can be easily seen that this algorithm requires a total of $n \log_2 n$ additions (specifically, $\frac{n \log_2 n}{2}$ additions and $\frac{n \log_2 n}{2}$ subtractions), instead of the n^2 additions/subtractions required when directly applying the matrix multiplication.

Algorithm 1 Fast Walsh-Hadamard Transform ($\boldsymbol{H}_i\boldsymbol{a}$ with $i \geq 1$)

procedure FastWalsh-HadamardTransform($\boldsymbol{a}$)
Input:
    $\boldsymbol{a}$ such that $\texttt{length}(\boldsymbol{a}) = n = 2^i$ and $i \geq 1$
Algorithm for $\texttt{FWHT}(\boldsymbol{a})$ *(computing* $\boldsymbol{H}_i\boldsymbol{a}$*)*:
    $\texttt{depth} = 1$;
    for $j = 0$ until $\log_2 n - 1$ **do**
        $\texttt{scale} = 2 * \texttt{depth}$;
        for $k = 0$ until $\left\lfloor \frac{\texttt{length}(\boldsymbol{a})-1}{\texttt{scale}} \right\rfloor$ **do**
            for $l = \texttt{scale} * k$ until $\texttt{scale} * k + \texttt{depth} - 1$ **do**
                $\texttt{ac} = \boldsymbol{a}[l]$;
                $\boldsymbol{a}[l] = \boldsymbol{a}[l] + \boldsymbol{a}[l + \texttt{depth}]$;
                $\boldsymbol{a}[l + \texttt{depth}] = \texttt{ac} - \boldsymbol{a}[l + \texttt{depth}]$;
        $\texttt{depth} = 2 * \texttt{depth}$;
Output:
    $\boldsymbol{a} \leftarrow \boldsymbol{H}_i\boldsymbol{a}$

Finally, the $\boldsymbol{\alpha}$-generalized version of the direct (inverse) FWHT can be defined by adding a right (left) product by a diagonal matrix, so that the total cost for the negacyclic FWHT on a length-n vector is n elemental multiplications and $n \log_2 n$ additions.

Implementation and evaluation Polynomial multiplications are the main bottleneck of lattice cryptography, as they are the most time-consuming basic blocks of any cryptographic algorithm, from encryption/decryption to relinearization and bootstrapping. The replacement of the traditional NTTs by FWHT by transitioning from cryptographic constructions built on univariate RLWE to the proposed multivariate version can bring a considerable improvement in terms of computational efficiency. To showcase the achieved gains, we have implemented

Algorithm 1 in C++ and compared it with one of the currently most efficient radix-2 implementations of the NTT [50]; this is the algorithm featured in the NFLlib, one of the fastest lattice-based cryptographic libraries available for homomorphic encryption. NFL also off-loads the complexity of the bit-reversal operation to the INTT, in order to speed up the NTT, and makes use of SSE and AVX2 optimizations to further enhance the obtained performance. Figure 2 shows the comparison of the obtained run times for a wide range of practical values of n (vector size or polynomial degree), when using our FWHT implementations, including an SSE/AVX2 vectorized version. It can be seen that we obtain a reduction to about 45-50% of the time of the NTT (38-43% of the INTT) in the non-vectorized implementation of the FWHT with respect to the fast NTT of NFLlib, while the vectorized one further reduces this figure to 22-24% (19-22% of the INTT). Finally, it is worth noting that the memory consumption of the FWHT is much lower, as it does not need to store the tables of the twiddle factors.

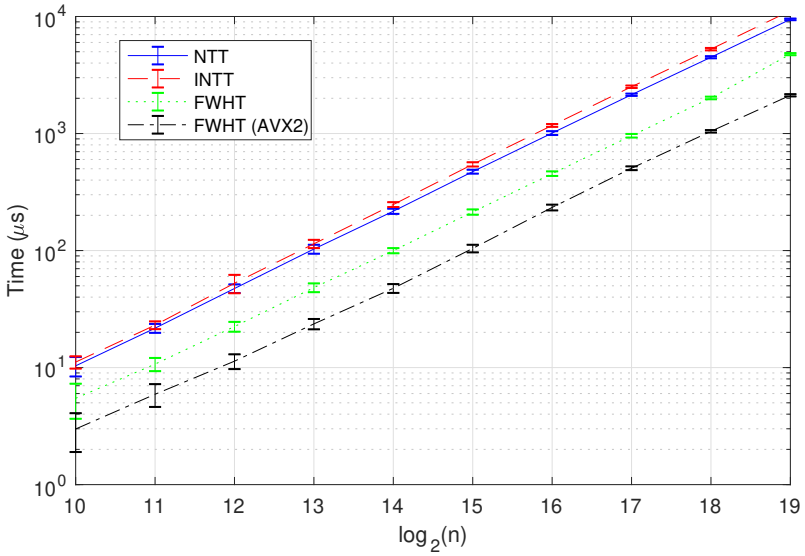

Fig. 2. Runtimes of the proposed FWHT compared to the NTT/INTT from [50].

9 Slot manipulation in multivariate rings

In this section we introduce the main improvements that m-RLWE brings to slot manipulation when packing several plaintext inputs into a ciphertext, with applications in relinearization and bootstrapping operations. Packing into slots [80]

helps to take advantage of the available space in the plaintext ring, therefore improving cipher expansion. The use of this packing strategy also enables working with homomorphic "slot"-wise additions and multiplications, i.e., SIMD (Single Instruction, Multiple Data) operations with encrypted data. This is usually combined with a mechanism to efficiently move and exchange the plaintext contents across slots, by taking advantage of the properties of the available automorphisms in the used ring. In general, in the ring $R_t = \mathbb{Z}_t[z]/\Phi_m(z)$, we can define a set of automorphisms $\phi(m)$ as different transformations $\rho_i : R_t \rightarrow R_t$ with $i \in \mathbb{Z}_m^*$, which apply a change of variable $z \rightarrow z^i$ over the elements in R_t.

Current lattice-based homomorphic cryptosystems leverage automorphisms to perform linear transformations across plaintext slots. Whereas applying an automorphism is a very efficient operation, it produces a ciphertext encrypted under a different secret key, and consequently, a switching key operation is needed to recover a ciphertext under the original secret key. This switching key process has two main drawbacks [49]: (a) a notable computational overhead and (b) an increase in the memory requirements due to the need of adding additional public information ("switching key/relinearization" matrices, a.k.a. evaluation keys).

In general, there is a tradeoff between these two dimensions: when the number of evaluation keys increases, the involved switching key runtime decreases, and conversely, when the number of keys is reduced, a chain of several switching key operations is needed, hence increasing the runtime. In a recent work [49], Halevi and Shoup explore several strategies to optimize this tradeoff, claiming improvements of even 75 times faster runtimes than those of their previous implementation, together with a reduction of up to a half in the required memory space to store the evaluation keys.

This section focuses on two different aspects: (1) We show how the introduced multivariate rings over the RLWE problem (see Sections 5 and 6) enable considerable improvements in the efficiency of the homomorphic packing/unpacking into slots, therefore greatly improving essential blocks for homomorphic encryption, such as bootstrapping, and (2) we analyze the structure of the available set of automorphisms on these rings, also showing that our solution can improve on both the runtime and the memory requirements with respect to the state of the art [49].

It is worth highlighting that some of the exemplified solutions in this section are sketched out with negacyclic rings. For completeness, in Section 9.4 we give some insights on how to extend these results to the more general multivariate rings showcased in this manuscript.

9.1 Efficient Slot Packing/Unpacking

The homomorphic packing/unpacking of plaintext values into slots is one of the most important examples of the evaluation of linear transformations on the ciphertexts, bootstrapping being one of the most representatives applications [48, 21, 25]. The way current cryptosystems implement this packing/unpacking is by means of a decomposition of the matrix multiplication into element-wise products between the different diagonals of the matrix and different rotated

versions of the ciphertext (hence by adding the result of a set of multiplications between plaintexts and rotated ciphertexts).

The main bottleneck of this process is the number of switching key matrices required to rotate the ciphertexts. Working with n slots, a total of $n-1$ rotations, hence $n-1$ switching key matrices, is required in the worst case. Available strategies to reduce this number of matrices come at the cost of also increasing the runtimes per automorphism/switching key operation.

Thanks to the introduced $(\log_2 n)$-RLWE, we break the need of a number of rotations (automorphisms/switching key operations) equal to the number of slots, and *we enable homomorphically packing/unpacking operations with a single switching key operation.* This is mainly due to the structure that the multivariate rings from Definition 10 present, which enables a much more efficient algorithm to compute the slot packing/unpacking, as we show next.

Consider a plaintext ring $R_t[x_1, \ldots, x_i]$, then the required matrices for packing and unpacking are respectively:

$$\boldsymbol{V}_i = 2^{-i} \underbrace{\left(\bigotimes_{j \in [i]} \begin{pmatrix} 1 & 0 \\ 0 & (-1)^{1/2} \end{pmatrix} \right)}_{\boldsymbol{B}} \boldsymbol{H}_i \quad \text{and} \quad \boldsymbol{V}_i^{-1} = \boldsymbol{H}_i \underbrace{\left(\bigotimes_{j \in [i]} \begin{pmatrix} 1 & 0 \\ 0 & (-1)^{-1/2} \end{pmatrix} \right)}_{\boldsymbol{B}^{-1}},$$

where $\beta = (-1)^{1/2}$ is the 4-th root of unity over the plaintext modulo t. Instead of directly applying these linear transformations (following the conventional approach), we resort to the NTT pre-/post-processing presented in [67], where the authors show how a DFT/NTT transform can be expressed in terms of element-wise products (NTT and a one-stage pre-/post-processing) and a negacyclic convolution. We show this process step by step, by computing first $\boldsymbol{H}_i$ and then $\boldsymbol{B}$ (resp. $\boldsymbol{B}^{-1}$).

$\boldsymbol{H}_i$ evaluation Adapting the results from [67] to the structure of our particular rings, it can be seen that the $\boldsymbol{H}_i$ matrix can be homomorphically evaluated by means of an automorphism and a negacyclic convolution with an all ones vector. That is, if we have encrypted a polynomial $a \in R_t[x_1, \ldots, x_i]$, let us define a polynomial $\mathbf{1}(x_1, \ldots, x_i) = \prod_{j \in [i]}(1 + x_j)$, such that the result of the multiplication

$$\mathbf{1}(x_1, \ldots, x_i) a(-x_1, \ldots, -x_i) \in R_t[x_1, \ldots, x_i]$$

is a polynomial whose coefficients correspond to the cyclic Hadamard transform.

$\boldsymbol{B}^{-1}$ and $\boldsymbol{B}$ evaluation The mentioned pre-/post-processing corresponds to the main diagonal of the matrices $\boldsymbol{B}^{-1}$ and $\boldsymbol{B}$, which comprise only four different values: $\{1, -1, \beta^{-1}, -\beta^{-1}\}$ for $\boldsymbol{B}^{-1}$ and $\{1, -1, \beta, -\beta\}$ for $\boldsymbol{B}$. This element-wise multiplication can be performed homomorphically over the encrypted polynomial coefficients through a change of variable in the ciphertext's polynomials: (1)

$\{x_j \to \beta^{-1}x_j\}_{j\in[i]}$ to calculate the $\boldsymbol{B}^{-1}$ matrix multiplication, and (2) $\{x_j \to \beta x_j\}_{j\in[i]}$ for the $\boldsymbol{B}$ matrix multiplication.[14]

Finally, we only need a relinearization/key switching operation to recover the original secret key after the two changes of variables $\{x_j \to -x_j\}_{j\in[i]}$ and $\{x_j \to \beta x_j\}_{j\in[i]}$ for packing (respectively $\{x_j \to \beta^{-1}x_j\}_{j\in[i]}$ and $\{x_j \to -x_j\}_{j\in[i]}$ for unpacking).

9.2 Automorphisms in Multiquadratic Rings and their Hypercube Structure

We show now how m-RLWE improves on the tradeoffs between space and computational cost when dealing with automorphisms, with respect to the univariate version.

Let $\mathbb{A}[z]/1+z^2$ be a polynomial ring as the one described by Definition 10, and α be an element $\alpha \in \mathbb{A}[z]/1+z^2$; then, we denote by $\theta_i^{(z)}(\alpha) \in \mathbb{A}[z]/1+z^2$ the transformation over α which applies the change of variable $z \to z^i$ with $i \in \mathbb{Z}_4^*$. For these particular rings, both transformations are, respectively, the identity $z \to z$ and the negation $z \to -z$. Reducing modulo t (the modulo of the plaintext ring), the effect of the latter transformation over the slots would be equivalent to a block shift where each block is composed by one half of the total slots. This shift is graphically illustrated in Figure 3, where ψ is the 4-th root of unity modulo t (i.e., $\psi^4 \equiv 1 \bmod t$), and the two blocks of slots encoded respectively in $\alpha(\psi)$ and $\alpha(\psi^3)$ get shifted by applying $z \to -z$.[15]

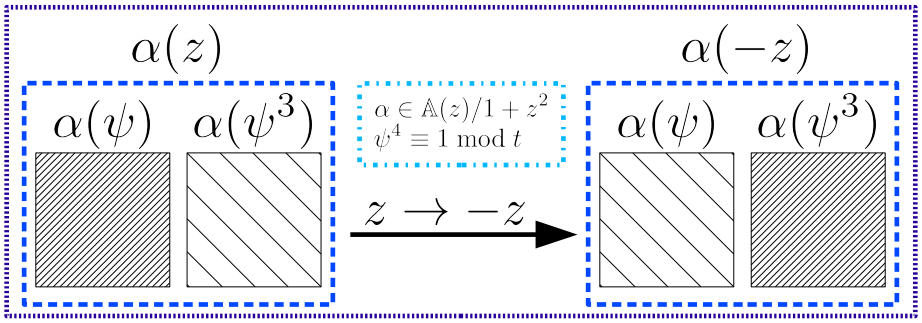

Fig. 3. Representation of the rotation between two blocks of slots encoded in α.

Going back to the notation $R_t[x_1,\ldots,x_l]$ with $f_j(x_j) = 1+x_j^2$ for our ring, we can then apply combinations of these two transformations with the different variables x_j for $j \in [l]$. Analogously to [49], this gives a multidimensional structure

[14] Making use of the decomposition of the formulation of the Bluestein FFT algorithm from [67], we can implement this change of variable by means of a homomorphic negacyclic convolution with NTT/INTT(diag($\boldsymbol{B}$)) and NTT/INTT(diag($\boldsymbol{B}^{-1}$)).

[15] With rings $\mathbb{A}[z]/d+z^2$ we have similar automorphisms $\{z \to z\}$ and $\{z \to -z\}$.

on the automorphisms group considering the composition of transformations

$$\theta_{i_1,\ldots,i_l}(\alpha) = \theta_{i_1}^{(x_1)}(\theta_{i_2}^{(x_2)}(\ldots\theta_{i_l}^{(x_l)}(\alpha)\ldots)) \in R_t[x_1,\ldots,x_l],$$

where $\alpha \in R_t[x_1,\ldots,x_l]$, $t \equiv 1 \bmod 4$ and $i_1,\ldots,i_l \in \mathbb{Z}_4^*$.

This multidimensional structure of the automorphisms group can be seen as an l-tuple with 2 different values per component (which gives a total of 2^l different automorphisms). Hence, similarly to the shift property of a multidimensional DFT [65], this group satisfies both the abelian and sharply transitive properties required to perform any type of permutation [45].

Logarithmic increase in space and computational cost (Strategy 1) The effect of each of the automorphisms over the slots can be visually represented as a hypercube with as many dimensions as independent variables the rings have, that is, with a total of $\log_2 n$ dimensions. As a graphical example, Figure 4 shows the slot structure corresponding to a multivariate ring with 7 independent variables; in this case, each different vertex of the hypercube represents one of the $n = 128$ available slots, where the allowed transitions between vertices depend on the chosen strategy, as we describe next.

In case of storing n switching key matrices (corresponding to all the automorphisms), any vertex transition will be allowed through one single switching key operation. However, it is possible to store less switching key matrices (which, combined, represent the whole set of automorphisms), hence increasing the number of subsequent automorphisms/switching key operations for transitioning from one vertex to another.

Due to the specific structure of our multivariate rings, we propose an optimal strategy with $\log_2 n$ switching key matrices, each one corresponding to a different transformation $x_i \to -x_i$; with the additional advantage that these transformations are their own inverses. Following this strategy, we can also see the different slots (vertices in Figure 4) as a binary vector of length $\log_2 n$, where the available operations are bit-wise XOR operations with vectors

$$\{(1,0,\ldots,0),(0,1,0,\ldots,0),\ldots,(0,\ldots,0,1)\}$$

belonging to the standard basis of dimension $\log_2 n$. In the example of Figure 4 (with $\log_2 n = 7$), this method would be equivalent to working with 7 independent vectors (of the standard basis) enabling only movements between vertices in the dimension associated to the vector.

It can be seen that with this strategy the farthest slot to a given one is always the slot represented as its ones' complement, i.e., the opposite vertex. This implies a total of $\log_2 n$ automorphisms/switching key operations. Hence, in the worst case we have an increase in the computational cost by a factor of $\log_2 n$ when storing $\log_2 n$ switching key matrices and working with n slots. This is a considerable reduction in the memory requirements when compared to the approximately $\mathcal{O}(D)$ and $\mathcal{O}(\sqrt{D})$ factors considered by Halevi and Shoup [49] when working with D slots (in one dimension).

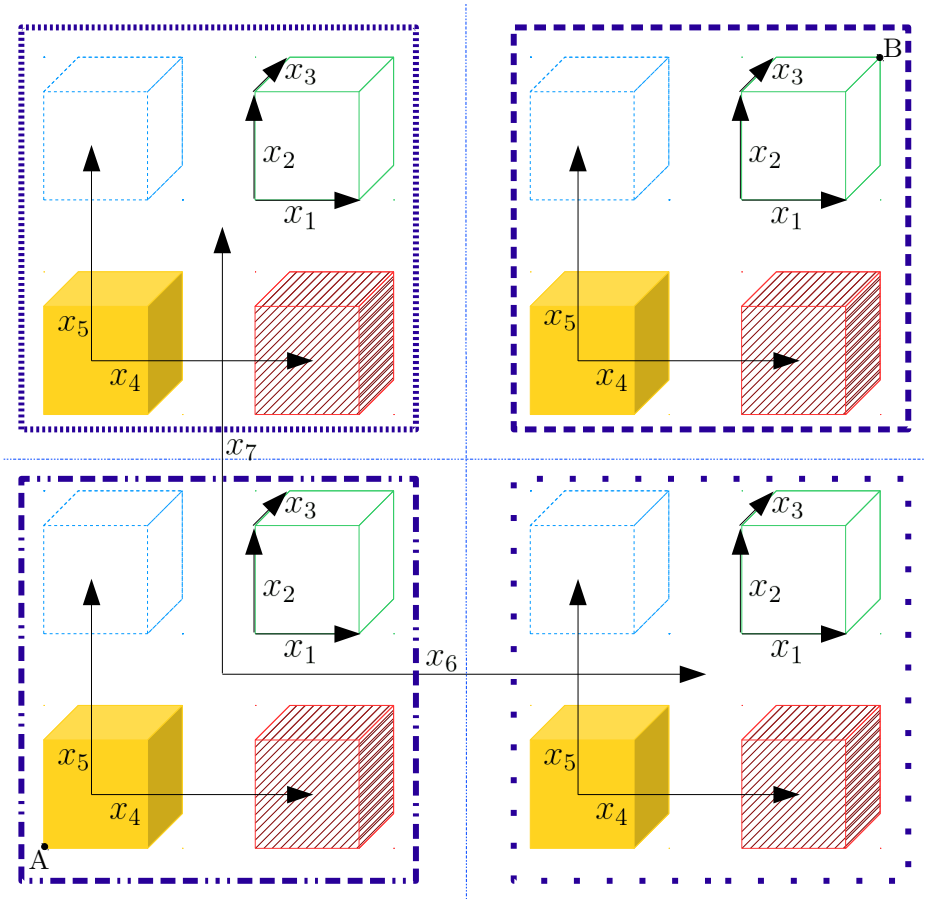

Fig. 4. Representation of the hypercube structure of the group of automorphisms available in the multivariate polynomial RLWE with $\Phi_4(\cdot)$ as modular function and considering 7 independent variables $\{x_1, \ldots, x_7\}$.

As a quick comparison, for the practical values reported in [49], i.e., $n = \phi(m) = 16384$, our strategy achieves an increase factor of 14 on the computational cost, which is not considerably higher than their results, but with huge savings in storage for our case: we store only 14 matrices, compared to the 51 matrices and 3 automorphisms/switching key operations achieved by [49] for a similar value of $\phi(m) = 15004$ and one dimension with $D = 682$ following a baby-step/giant-step strategy (see Appendix B).

Finally, it must be noted that when applying a switching key, noise constraints force the need of decomposing the coefficients of the involved polyno-

mials in some specific base.[16] As this decomposition does not straightforwardly commute with the NTT/INTT (or CRT over the polynomial modular function) representation, the inverse and direct transforms have to be applied over the polynomials. Our setting in multivariate rings with FWHT enables a reduction on complexity for these transforms by a factor of $\mathcal{O}(\log n)$ in terms of elemental products; i.e., *this yields a net gain factor of* $\log n$ *in storage while keeping the same order of (multiplicative) computational complexity.*

Efficiency/space tradeoffs In practical scenarios, the tradeoff between used memory and computational cost might require a different balance with less space efficiency than the $\log_2 n$ achieved by the described strategy. Consequently, we also cover two additional strategies which lead to an improvement of the computational cost by a factor of 2.

Strategy 2: Our first approach adds to the previous $\log_2 n$ matrices those which are associated to "diagonal" vectors in our hypercube representation of the autormorphisms (see Figure 4); that is, we work with automorphisms $\{x_i \to x_i^{l_i}, x_j \to x_j^{l_j}\}$ where $l_i, l_j \in \mathbb{Z}_4^*$ and $i, j \in [\log_2 n]$, being $i \neq j$. Going back again to the binary representation of the slots, the additional automorphisms could be seen as the result of all pairwise XOR operations of different vectors of the standard basis of length $\log_2 n$.

The number of needed switching key matrices is therefore increased to

$$\binom{1 + \log_2 n}{2} = \frac{(1 + \log_2 n) \log_2 n}{2}.$$

In order to calculate the associated computational cost for this strategy, we resort to induction, working first with the odd natural numbers, and afterwards with the even natural numbers. Let the multivariate ring $R_t[x_1, \ldots, x_l]$ with $f_i(x_i) = 1 + x_i^2$ where $i = 1, \ldots, l$ and $l = \log_2 n$, if we consider only the odd values of l we have:

- For $l = 1$, any transition can be applied with only one automorphism/relinearization operation.
- Assuming that l variables require k automorphisms/relinearization operations, it can be shown that adding two variables (i.e., $l + 2$), $k + 1$ automorphisms/relinearization operations are needed. We can graphically see this by resorting to the binary representation: moving between any two slots implies, in the worst case (consider one vector and its ones' complement), one additional XOR operation.
- Therefore, by induction, odd values of l require $\lceil \frac{l}{2} \rceil$ automorphisms/relinearization operations.

[16] This is true unless we resort to the strategy of Bajard *et al.* [5] which takes advantage of the CRT decomposition over the polynomial coefficients. However, this strategy cannot be applied always, as it requires a highly composite modulo with primes of an adequate machine size (see [1]).

The argument is analogous for even l. First, we consider $l = 2$, where with only one automorphism/relinearization operation is enough to move between any of the slots. Next, the same reasoning as before could be applied between l and $l+2$ variables, resulting in a total of $\frac{l}{2}$ automorphisms/relinearization operations for l variables.

Taking into account both results, this strategy yields an increase in the number of automorphisms/switching key operations by a factor of $\lceil \frac{\log_2 n}{2} \rceil$. Hence, we can reduce by a half the computational cost compared to our previous strategy, with a quadratic increase in the memory requirements of $\frac{(1+\log_2 n)\log_2 n}{2}$ instead of $\log_2 n$. For instance, with $n = 16384$ this would give an increase in cost by a factor of 7 and a total of 105 stored matrices.

Strategy 3: The incurred increase in space requirements by Strategy 2 might not be acceptable for certain applications; therefore, our next approach preserves the cost improvement, but achieving a negligible increase in the number of required matrices: $1 + \log_2 n$ matrices instead of $\mathcal{O}((\log n)^2)$.

The idea behind this approach is adding to the switching key matrices for transformations of the form $\{x_i \to -x_i\}$ for $i = 1, \ldots, \log_2 n$ the following one

$$\{x_1 \to -x_1, \ldots, x_{\log_2 n} \to -x_{\log_2 n}\}.$$

As a graphical explanation, let us consider again the binary representation of the slots: in addition to working with those XOR operations with vectors belonging to the standard basis of length $\log_2 n$, now we can also apply the ones' complement of every "slot" in one operation (e.g., in Figure 4 we could directly move with one automorphism/switching key operation from point A to point B).

Therefore, the worst case automorphism requiring $l = \lceil \frac{\log_2 n}{2} \rceil$ matrices with our first strategy can now be computed with just one matrix. Moreover, as we know that $l - \lceil \frac{l}{2} \rceil \leq \lceil \frac{l}{2} \rceil$ for any $l \in \mathbb{N}$, then the farthest slot position can be achieved by only $\lceil \frac{l}{2} \rceil = \lceil \frac{\log_2 n}{2} \rceil$ automorphisms. Consequently, we can see that with $1 + \log_2 n$ matrices, we only need a maximum of $\lceil \frac{\log_2 n}{2} \rceil$ automorphism/switching key operations. For instance, with $n = 16384$ this would give an increase in cost by a factor of 7 and a total of 15 matrices in terms of use of memory.

9.3 Automorphisms in Multivariate Power-of-Two Cyclotomic Rings

It can be useful to expand Definition 10 to also cover more general multivariate rings, which can be leveraged by some applications (see Appendix A). Most of these applications consider a general multivariate ring as the R and R_q, where each of the modular functions can be defined as different power-of-two cyclotomic polynomials $f_i(x_i) = x_i^{n_i} + 1$.[17]

[17] Analogously to the procedure we followed with multiquadratics in Section 9.2, we exemplify these results with power-of-two cyclotomics. They can be similarly ex-

In this section the discussed efficiency/space tradeoffs achievable with automorphisms on the FWHT-enabled rings will be expanded to these rings (at the cost of lacking the faster FFT algorithms for the negacyclic Hadamard transform).

We first remind the corresponding definition of multivariate RLWE with power-of-two cyclotomic polynomials (see Definition 1):

Definition (multivariate RLWE with power-of-two modular functions as $x_i^{n_i}+1$ (Definition 1))**.** *Given a multivariate polynomial ring $R_q[x_1,\ldots,x_l]$ with $f_j(x_j)=1+x_j^{n_j}$ for $j=1,\ldots,l$ where $n=\prod_j n_j$ (with all n_j a power of two) and an error distribution $\chi[x_1,\ldots,x_l]\in R_q[x_1,\ldots,x_l]$ that generates small-norm random multivariate polynomials in $R_q[x_1,\ldots,x_l]$, the multivariate polynomial RLWE relies upon the computational indistinguishability between samples $(a_i,b_i=a_i\cdot s+e_i)$ and (a_i,u_i), where $a_i,u_i\leftarrow R_q[x_1,\ldots,l]$ are chosen uniformly at random from the ring $R_q[x_1,\ldots,x_l]$; $s,e_i\leftarrow\chi[x_1,\ldots,x_l]$ are drawn from the error distribution.*

Tradeoffs in the size/efficiency of automorphisms We consider the ring R introduced in Definition 1; particularly, we work with $R_t[x_1,\ldots,x_l]$ where $t\equiv 1 \bmod 2n_i$ for $i=1,\ldots,l$. Analogously to our derivation in Section 9.2, when working with an element $\alpha\in R_t[x_1,\ldots,x_l]$, we have the transformations

$$\theta_{i_1,\ldots,i_l}(\alpha)=\theta_{i_1}^{(x_1)}(\theta_{i_2}^{(x_2)}(\ldots\theta_{i_l}^{(x_l)}(\alpha)\ldots))\in R_t[x_1,\ldots,x_l],$$

now with $i_j\in\mathbb{Z}_{2n_j}^*$ for all j.

This multidimensional structure can be seen again as an l-tuple, where each component has n_i different values, hence giving a total of $n=\prod_{i=1}^l n_i$ different automorphisms.

Strategy 4: Our main strategy works with n_i-1 matrices for each variable x_i, where each switching key matrix will correspond to an automorphism $\{x_i\rightarrow x_i^{l_i}\}$ for $l_i\in\mathbb{Z}_{2n_i}^*$ (except $\{x_i\rightarrow x_i\}$) and $i=1,\ldots,l$. This strategy yields a total of $\sum_{i=1}^l n_i-l$ matrices with a computational cost of l automorphism/switching key operations. Let us assume that all the matrices for every "univariate" change of variable have to be stored. However, the number of required matrices per "univariate" change of variable could be further improved [49] (that is, we could work with subsets $\mathbb{A}_i\in\mathbb{Z}_{2n_i}^*$ in such a way that the corresponding automorphisms would be $\{x_i\rightarrow x_i^{l_i}\}$ for $l_i\in\mathbb{A}_i$ and $i=1,\ldots,l$). [18]

We consider those $n_i=n^{\frac{1}{l}}$ for $i=1,\ldots,l$ (hence being all n_i equal). This gives us several tradeoffs depending on l and n where we have $l(n^{\frac{1}{l}}-1)$ matrices and an increase in the computational cost by a factor of l. Table 2 shows the number of required matrices and the increase in computational cost for $n=$

tended to more general rings of the form $x_i^{n_i}+d_i$. We refer the reader to Section 9.4 for more details.

[18] For a brief summary of Halevi and Shoup full and baby-step/giant-step strategies, see Appendix B.

Table 2. Practical space/efficiency tradeoffs of automorphisms for $n = 16384$

l	2	3	4	5	6	7
# Matrices	256	80	52	36	34	28
# Calls to switching key (worst-case)	2	3	4	5	6	7

Table 3. Space/efficiency tradeoffs of automorphisms

Strategy	# Matrices	# Calls to switching key (worst-case)
Strategy 1 from Section 9.2	$\log_2 n$	$\log_2 n$
Strategy 2 from Section 9.2	$\frac{(1+\log_2 n)\log_2 n}{2}$	$\lceil \frac{\log_2 n}{2} \rceil$
Strategy 3 from Section 9.2	$1 + \log_2 n$	$\lceil \frac{\log_2 n}{2} \rceil$
Strategy 4 from Section 9.3	$\approx n^{\frac{1}{l}} l - l$	l
Strategy 4 (general) from Section 9.3	$\sum_{i=1}^{l} n_i - l$	l

16384 and several values of l. As $n^{\frac{1}{l}}$ is not always a valid value (that is, a power of two), the choice of n_i can be optimized to achieve the smallest possible number of automorphisms ($\sum n_i$) such that $n = \prod n_i$.

Conversely, Table 3 summarizes the different tradeoffs we have presented in this section.

9.4 On the applicability to more general multivariate rings

It is worth noting that all the solutions exemplified above (Sections 9.2 and 9.3) are sketched out with negacyclic rings. In this section, we give some insights on how to extend these results to the more general multivariate rings showcased in this manuscript. To this aim, we resort to the generalized pre-/post-processing presented in [67], together with the decomposition of the NTT/INTT transforms into a chain of automorphisms and convolution operations.

An alternative set of modular functions Bernstein *et al* [8] propose a different non-cyclotomic ring. The authors argue that with cyclotomic rings it is easy to have non-trivial ring homomorphisms (as the modular function usually splits in linear factors to perform FFT algorithms) and a relatively small Galois group. Consequently, the authors propose rings of the form $\mathbb{Z}_q[x]/(f_p(x))$, with an irreducible modular function[19] $f_p(x) = x^p - x - 1$ and p prime, where the Galois group is the permutation group S_p with $p!$ elements, and the modulo q is inert in the ring. Hence, $\mathbb{Z}_q[x]/(x^p - x - 1)$ is indeed a finite field.

These modular functions are also interesting for our purposes, but for very different reasons. Let $K = \mathbb{Q}(\alpha)$ be a number field with α one of the roots of $x^n - x - 1$. We know that [36] modular functions $f_n(x) = x^n - x - 1$ with

[19] See [36] for more details on the properties exhibited by functions of the form $f_n(x) = x^n - x - 1$.

$n \geq 2$ are irreducible, and for $2 \leq n \leq 100$ the discriminant of $f_n(x)$ is squarefree. According to Theorem 5, this means that K is monogenic and $\mathcal{O}_K = \mathbb{Z}[x]/f_n(x)$.

Now, from Proposition 4, we have

$$\Delta_K = (-1)^{\frac{n(n-1)}{2}}(n^n(-1)^{n-1} - (n-1)^{n-1}),$$

so it is straightforward to find coprime discriminants for different values of n.

For example, the discriminants of $\{f_i(x)\}_{i=2,\ldots,7}$ are coprime. Therefore, we can define a multivariate RLWE sample over the ring of integers

$$\mathcal{O}_K = \mathbb{Z}[x_1, \ldots, x_7]/(f_2(x_2), \ldots, f_7(x_7))$$

for a multivariate number field of degree 5040 and 6 dimensions. In general, this gives an easy way to find multivariate number fields with many variables and a small expansion factor.

Operations over these rings are not as efficient as the ones with modular function $x^n - d$, but still acceptable; i.e., *in the worst case*, multiplications modulo $x^n - x - 1$ can be decomposed in multiplications modulo $x^n - x$ and $x^n - 1$, hence requiring two parallel efficient "cyclic" convolutions, and afterwards, adding the obtained results.

Automorphisms for more general multivariate rings The multivariate rings introduced in Section 6 are, in general, *separable but non-Galois field extensions*. This implies that the number of available automorphisms is strictly smaller than the degree of the extension (see Corollary 1).

Corollary 1 (Corollary 4.3 from [35])**.** *If L/K is a finite extension that is either inseparable or not normal then*

$$|Aut(L/K)| < [L : K],$$

being $[L : K]$ the degree of the field extension.

Fortunately, this is not a problem in practice as we can make use of Theorem 8 to extend the mentioned *separable* multivariate number fields in Section 6 to a Galois extension, where we have $\mathrm{Gal}(L/K) = \mathrm{Aut}(L/K) = [L : K]$; hence, automorphisms similar to the case of power-of-two cyclotomics (see Section 9.3) can still be applied.

Theorem 8 (Theorem 4.8 **from [35]).** *Every finite separable extension of a field can be enlarged to a finite Galois extension of the field. In particular, every finite extension of a field with characteristic* 0 *can be enlarged to a finite Galois extension.*

A toy example for a prime-degree field extension Consider the number field $\mathbb{Q}(d^{\frac{1}{p}})$ (with $d > 1$ and $d \in \mathbb{N}$) isomorphic to the polynomial ring $\mathbb{Q}[x]/(x^p - d)$ and satisfying the conditions from Section 6 (Proposition 5). We know that the roots of $x^p - d$ are $\{d^{\frac{1}{p}}, \zeta_p d^{\frac{1}{p}}, \ldots, \zeta_p^{p-1} d^{\frac{1}{p}}\}$. These roots are separable, but $\mathbb{Q}(d^{\frac{1}{p}})$

is not the corresponding splitting field, and hence $\mathbb{Q}(d^{\frac{1}{d}})$ is not a Galois field extension over the rationals $\mathbb{Q}$.

Even so, we know from Theorem 8 that this field can be extended to a Galois field where we have a Galois automorphism group which enables "rotations" of the slots. It suffices to add the root ζ_p by means of a symbolic variable y over the cyclotomic polynomial $\Phi_p(y) = \sum_{i=0}^{p-1} y^i$, i.e., we enlarge the number field (see Theorem 8) to have $\mathbb{Q}(d^{\frac{1}{p}}, \zeta_p)$ with d and p different primes.

For this extended number field and considering a polynomial representation with $\mathbb{Q}[x,y]/(x^p - d, \Phi_p(y))$ (thanks to the field isomorphism $d^{\frac{1}{p}} \to x, \zeta_p \to y$), we have the chain of transformations $\{x \to xy^i, y \to y^j\}$ with $i \in \mathbb{Z}_p$ and $j \in \mathbb{Z}_p^*$, which enables homomorphic "rotation" of the slots.

As an example, consider the polynomial $a(x) = \sum_{i=0}^{p-1} a_i x^i \bmod x^p - d$. We apply the change of variable $x \to xy$

$$\begin{aligned} a(x) &= \sum_{i=0}^{p-1} a_i x^i \\ &= \sum_{i=0}^{p-1} a_i x^i y^i \\ &= a_{p-1} y^{p-1} x^{p-1} + \sum_{i=0}^{p-2} a_i y^i x^i. \end{aligned}$$

Consider now the following relation given by $\Phi_p(y)$

$$y^{p-1} = -\sum_{i=0}^{p-2} y^i.$$

By applying it, we have:

$$a_{p-1} y^{p-1} x^{p-1} + \sum_{i=0}^{p-2} a_i y^i x^i = -a_{p-1} x^{p-1} \sum_{i=0}^{p-2} y^i + \sum_{i=0}^{p-2} a_i y^i x^i.$$

It is worth noting that the ring $\mathbb{Z}[x,y]/(x^p - d, \Phi_p(y))$ is not, in general, the ring of integers of the field $\mathbb{Q}(d^{\frac{1}{p}}, \zeta_p)$, but instead a subring of its ring of integers. This can be easily seen by inspecting the discriminants of $x^p - d$ and $\Phi_p(y)$ which are, respectively, $(-1)^{\frac{p(p-1)}{2}} p^p (-d)^{p-1}$ and p^{p-2}. As they are not coprime we cannot assert that the ring of integers of $\mathbb{Q}(d^{\frac{1}{p}}, \zeta_p)$ is the product of $\mathbb{Z}[x]/(x^p - d)$ and $\mathbb{Z}[y]/(\Phi_p(y))$.[20]

Consequently, when working with rings following Definition 12 in Section 6, if we want to (1) base the security on RLWE over a general number field and also (2) make use of the automorphisms, the reduction from Theorem 2 implies a loss

[20] If $x^p - d$ satisfies the conditions established in Proposition 5, $\mathbb{Z}[x]/(x^p - d)$ is the ring of integers of $\mathbb{Q}(d^{\frac{1}{p}})$.

in the lattice dimensionality; in the previous example of $\mathbb{Z}[x,y]/(x^p - d, \Phi_p(y))$, we end up working with a ring of degree $p(p-1)$, but being the original RLWE sample defined over a number field of degree p. Nevertheless, we can avoid this loss by basing the security in a generalization of RLWE called Order-LWE.

A much wider set of ring choices with Order-LWE Bolboceanu *et al.* [10] propose a generalization of RLWE which, instead of considering the ring of integers $\mathcal{O}_K$ and its dual $\mathcal{O}_K^\vee$, relies on the subrings called orders $\mathcal{O}$ and their corresponding duals $\mathcal{O}^\vee$ to define the underlying ideal lattices.

For a number field K of degree n, an order $\mathcal{O}$ in K is a subring of $\mathcal{O}_K$ containing a $\mathbb{Q}$-basis of full-rank n of K such that $\mathcal{O} \otimes_{\mathbb{Z}} \mathbb{Q} = K$. The ring of integers is the maximal order of K.

Order-LWE also presents worst-case hardness with respect to short vector problems, but in the invertible-ideal lattices of the considered order [10].

This result enables a relaxation of many of the restrictions imposed for the rings in Sections 5 and 6, by directly basing their hardness on Order-LWE. The previous example with the field $\mathbb{Q}(d^{\frac{1}{p}}, \zeta_p)$ and *order* $\mathbb{Z}(d^{\frac{1}{p}}, \zeta_p)$ can base its hardness on a lattice of dimension $p(p-1)$ by considering Order-LWE.

The use of the modular function $\Phi_p(y)$ seems to contradict our initial requirements regarding the desired form of the modular function (see Section 1). However, for efficient polynomial products we can substitute $\Phi_p(y)$ by $y^p - 1$ by just multiplying both polynomial elements and modular function with the term $y - 1$.

We plan to extend our results and optimizations to the corresponding relaxations offered by Order-LWE. In this direction, this work provides a wide set of concrete ring instantiations which could be considered to analyze the hardness of Order-LWE.

10 Improving on the packing capacity of complex numbers

We have addressed packing of integer numbers in Section 9, but complex numbers are more difficult to efficiently pack. Nevertheless, we can also leverage the multivariate structure to represent the complex arithmetic in a much more efficient way than previous recent approaches. Knowing that a total of $n/2$ complex slots can be packed over the ring $\mathbb{Z}[z]/1 + z^n$, Cheon *et al.* [27, 26] expand these results to the bivariate case $\mathbb{Z}[x,y]/(1+x^{n_x}, 1+y^{n_y})$, packing a total of $\frac{n_x}{2}\frac{n_y}{2} = \frac{n}{4}$ complex slots. Generalizing this strategy to l dimensions, packing is restricted to $\frac{n}{2^l}$ complex slots (where $n = \prod_{i=1}^{l} n_i$) when working over multivariate rings as $\mathbb{Z}[x_1, \ldots, x_l]/(1 + x_1^{n_1}, \ldots 1 + x_l^{n_l})$.[21] Consequently, this strategy leaves a huge gap of unused potential slots when transitioning to a multivariate ring.

[21] While this strategy was introduced for a weak instance of multivariate RLWE (i.e., vulnerable to Bootland *et al.*'s attack), a similar approach works for rings following Definition 12.

Nevertheless, it is possible to achieve the same number of complex slots as the univariate counterpart (that is, $n/2$ complex slots), effectively substituting the multivariate complex embedding map (as used in [26]) by its univariate version. Let us consider the ring $\mathbb{Z}[x_1, \ldots, x_l]/(d_1 + x_1^{n_1}, \ldots, d_l + x_l^{n_l})$, and choose one of the l independent variables to work with the canonical embedding map, x_1 without loss of generality. If we have a total of $n/2$ complex numbers to pack in one multivariate polynomial plaintext, we organize them as a set of $\frac{n}{n_1}$ complex vectors with length $n_1/2$. For each complex vector we use the encoding from [27], defined as the composition of the inverse of the complex embedding map and a discretization. This yields $\frac{n}{n_1}$ polynomials belonging to the ring $\mathbb{A} = \mathbb{Z}_t[x_1]/d_1 + x_1^{n_1}$.

Coming back to the multivariate ring representation, we can consider the new message as a polynomial in the ring $\mathbb{Z}_t[x_1, \ldots, x_l]/(d_1 + x_1^{n_1}, \ldots, d_l + x_l^{n_l})$. Hence, we gather all the polynomials in $\mathbb{A}$ as the different coefficients of the ring $\mathbb{A}[x_2, \ldots, x_l]/(d_2 + x_2^{n_2}, \ldots, d_l + x_l^{n_l})$, and we define encoding/decoding matrices working over $d_i + x_i^{n_i}$ modular functions (i.e., α-generalized INTTs/NTTs over t, see Section 8) for $i = 2, \ldots, l$, considering the identity matrix $\boldsymbol{I}_{n_1}$ of size $n_1 \times n_1$ for x_1 and the modular function $d_1 + x_1^{n_1}$. Using the vector representation of the plaintext polynomial, the encoding/decoding is performed by means of one matrix multiplication which can be efficiently realized with FFT-like algorithms.

This method can pack a total of $n/2$ complex slots while preserving the properties for the automorphisms (whenever we enlarge the number field to a Galois extension, see Section 9.4) and also removing the gap of the method used in [26], where the fraction of used slots decreases exponentially with the number of dimensions.

Finally, it is worth looking at the case where the considered multivariate rings are those from Definition 10 in Section 5. In this case, the modular functions have the form $d_i + x_i^2$, so the variable x_1 can directly represent the imaginary unit, therefore perfectly mapping the complex arithmetic without the need of applying the canonical embedding map over the polynomials in $\mathbb{A}$.

11 Conclusions

This work addresses the main security flaw of the multivariate RLWE problem revealed by Bootland *et al.* For this purpose, we have defined and parameterized practical and secure instantiations of the multivariate Ring Learning With Errors problem, supported by the extended reduction of the original proof by Lyubashevsky *et al.* [60, 58]. The proposed instantiations are resilient against Bootland's attack to m-RLWE [12], while still preserving all the efficiency improvements that m-RLWE brings. We have shown how to find practical parameters for the proposed instantiations to make them both secure and usable, therefore enabling improved space-time tradeoffs in many practical applications, comprising the most critical fundamental lattice operations (faster polynomial multiplications through $\boldsymbol{\alpha}$-generalized Walsh-Hadamard Transforms), efficient cryptographic operations such as computation of automorphisms, relin-

earizations, packing, unpacking and homomorphic slot manipulation, and, consequently, bootstrapping, and optimization of high level applications in encrypted approximate arithmetic, complex processing, and efficient multidimensional signal manipulation.

These contributions, combined, showcase the power and versatility of secure instantiations of the multivariate RLWE problem, and open up new research paths and strategies for realizing efficient (fully) homomorphic encryption.

Acknowledgments

GPSC is funded by the Agencia Estatal de Investigación (Spain) and the European Regional Development Fund (ERDF) under projects WINTER (TEC2016-76409-C2-2-R) and COMONSENS (TEC2015-69648-REDC). Also funded by the Xunta de Galicia and the European Union (European Regional Development Fund - ERDF) under projects Agrupación Estratéxica Consolidada de Galicia accreditation 2016-2019 and Grupo de Referencia ED431C2017/53, and also by the FPI grant (BES-2014-069018). EPFL is funded in part by the grant #2017-201 (DPPH) of the Swiss PHRT and by the grant #2018-522 (MedCo) of the Swiss PHRT and SPHN.

References

1. Aguilar-Melchor, C., Barrier, J., Guelton, S., Guinet, A., Killijian, M.O., Lepoint, T.: NFLlib: NTT-Based Fast Lattice Library, pp. 341–356. Springer (2016)
2. Albrecht, M.R., Curtis, B.R., Deo, A., Davidson, A., Player, R., Postlethwaite, E.W., Virdia, F., Wunderer, T.: Estimate all the LWE, NTRU schemes! (2018)
3. Albrecht, M.R., Player, R., Scott, S.: On the concrete hardness of Learning with Errors. J. Mathematical Cryptology 9(3), 169–203 (2015)
4. Applebaum, B., Cash, D., Peikert, C., Sahai, A.: Fast Cryptographic Primitives and Circular-Secure Encryption Based on Hard Learning Problems. In: CRYPTO. LNCS, vol. 5677, pp. 595–618. Springer (2009)
5. Bajard, J., Eynard, J., Hasan, M.A., Zucca, V.: A full RNS variant of FV like somewhat homomorphic encryption schemes. In: SAC. pp. 423–442 (2016)
6. Barile, M.: Eisenstein's Irreducibility Criterion. From MathWorld, A Wolfram Web Resource, created by Eric. W. Weisstein, `http://mathworld.wolfram.com/EisensteinsIrreducibilityCriterion.html`, accessed: 11 March 2019
7. Bauch, J., Bernstein, D.J., de Valence, H., Lange, T., van Vredendaal, C.: Short Generators Without Quantum Computers: The Case of Multiquadratics. In: EUROCRYPT. LNCS, vol. 10210, pp. 27–59 (2017)
8. Bernstein, D.J., Chuengsatiansup, C., Lange, T., van Vredendaal, C.: NTRU prime: Reducing attack surface at low cost. In: SAC. pp. 235–260 (2017)
9. Biasse, J., Ruiz, L.: FHEW with Efficient Multibit Bootstrapping. In: LATINCRYPT. LNCS, vol. 9230, pp. 119–135 (2015)
10. Bolboceanu, M., Brakerski, Z., Perlman, R., Sharma, D.: Order-lwe and the hardness of ring-lwe with entropic secrets. Crypt. ePrint Archive, Report 2018/494 (2018)
11. Bonnoron, G., Ducas, L., Fillinger, M.: Large FHE gates from tensored homomorphic accumulator. In: AFRICACRYPT. LNCS, vol. 10831, pp. 217–251 (2018)

12. Bootland, C., Castryck, W., Vercauteren, F.: On the security of the multivariate ring learning with errors problem. Crypt. ePrint Archive, Report 2018/966 (2018)
13. Bos, J., Lauter, K., Loftus, J., Naehrig, M.: Improved Security for a Ring-Based Fully Homomorphic Encryption Scheme. In: Cryptography and Coding, LNCS, vol. 8308, pp. 45–64. Springer (2013)
14. Brakerski, Z.: Fully Homomorphic Encryption without Modulus Switching from Classical GapSVP. In: CRYPTO, LNCS, vol. 7417, pp. 868–886. Springer (2012)
15. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (Leveled) Fully Homomorphic Encryption without Bootstrapping. ACM Trans. Comput. Theory 6(3), 13:1–13:36 (Jul 2014)
16. Brakerski, Z., Vaikuntanathan, V.: Fully Homomorphic Encryption from Ring-LWE and Security for Key Dependent Messages. In: CRYPTO, LNCS, vol. 6841. Springer (2011)
17. Brakerski, Z., Gentry, C., Vaikuntanathan, V.: (leveled) fully homomorphic encryption without bootstrapping. In: ITCS. pp. 309–325 (2012)
18. Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. In: IEEE FOCS. pp. 97–106 (2011)
19. Castryck, W., Iliashenko, I., Vercauteren, F.: Provably weak instances of ring-lwe revisited. In: EUROCRYPT. LNCS, vol. 9665, pp. 147–167 (2016)
20. Chase, M., Chen, H., Ding, J., Goldwasser, S., Gorbunov, S., Hoffstein, J., Lauter, K., Lokam, S., Moody, D., Morrison, T., Sahai, A., Vaikuntanathan, V.: Security of homomorphic encryption. Tech. rep., HomomorphicEncryption.org, Redmond WA (July 2017)
21. Chen, H., Han, K.: Homomorphic lower digits removal and improved FHE bootstrapping. In: EUROCRYPT. LNCS, vol. 10820, pp. 315–337 (2018)
22. Chen, H., Lauter, K.E., Stange, K.E.: Vulnerable Galois RLWE Families and Improved Attacks. Crypt. ePrint Archive, Report 2016/193 (2016)
23. Chen, H., Lauter, K.E., Stange, K.E.: Attacks on the search-rlwe problem with small errors. arXiv e-prints (2017)
24. Chen, H., Lauter, K.E., Stange, K.E.: Security considerations for galois non-dual RLWE families. arXiv e-prints (2017)
25. Cheon, J.H., Han, K., Kim, A., Kim, M., Song, Y.: Bootstrapping for approximate homomorphic encryption. In: EUROCRYPT. LNCS, vol. 10820, pp. 360–384 (2018)
26. Cheon, J.H., Kim, A.: Homomorphic Encryption for Approximate Matrix Arithmetic. Crypt. ePrint Archive, Report 2018/565 (2018)
27. Cheon, J.H., Kim, A., Kim, M., Song, Y.S.: Homomorphic Encryption for Arithmetic of Approximate Numbers. In: ASIACRYPT. LNCS, vol. 10624, pp. 409–437 (2017)
28. Cheon, J.H., Kim, A., Yhee, D.: Multi-dimensional packing for HEAAN for approximate matrix arithmetics. Crypt. ePrint Archive, Report 2018/1245 (2018)
29. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Faster Fully Homomorphic Encryption: Bootstrapping in Less Than 0.1 Seconds. In: ASIACRYPT. LNCS, vol. 10031, pp. 3–33 (2016)
30. Chillotti, I., Gama, N., Georgieva, M., Izabachène, M.: Faster Packed Homomorphic Operations and Efficient Circuit Bootstrapping for TFHE. In: ASIACRYPT. LNCS, vol. 10624, pp. 377–408 (2017)
31. Chillotti, I., Gama, N., Georgieva, M., Izabachne, M.: TFHE: Fast Fully Homomorphic Encryption over the Torus. Crypt. ePrint Archive, Report 2018/421 (2018)
32. Conrad, B.: Math 154: Discrimant of Composite Fields. `http://math.stanford.edu/~conrad/154Page/handouts/disccomposite.pdf`, accessed: 11 March 2019

33. Conrad, B., Landesman, A.: Math 154: Algebraic Number Theory. http://math.stanford.edu/~conrad/154Page/handouts/undergraduate-number-theory.pdf, accessed: 11 March 2019
34. Conrad, K.: The Different Ideal. https://kconrad.math.uconn.edu/blurbs/gradnumthy/different.pdf, accessed: 11 March 2019
35. Conrad, K.: The Galois Correspondence. https://kconrad.math.uconn.edu/blurbs/galoistheory/galoiscorr.pdf, accessed: 11 March 2019
36. Conrad, K.: The Galois Group of $x^n - x - 1$ over $\mathbb{Q}$. https://kconrad.math.uconn.edu/blurbs/gradnumthy/galoisselmerpoly.pdf, accessed: 11 March 2019
37. Costache, A., Smart, N.P.: Which ring based somewhat homomorphic encryption scheme is best? In: CT-RSA. pp. 325–340 (2016)
38. Ducas, L., Micciancio, D.: FHEW: bootstrapping homomorphic encryption in less than a second. In: EUROCRYPT. LNCS, vol. 9056, pp. 617–640 (2015)
39. Eisenträger, K., Hallgren, S., Lauter, K.E.: Weak instances of PLWE. In: SAC. pp. 183–194 (2014)
40. Elias, Y., Lauter, K.E., Ozman, E., Stange, K.E.: Provably weak instances of ring-lwe. In: CRYPTO. LNCS, vol. 9215, pp. 63–92 (2015)
41. Fan, J., Vercauteren, F.: Somewhat Practical Fully Homomorphic Encryption. Crypt. ePrint Archive, Report 2012/144 (2012)
42. Fino, B.J., Algazi, V.R.: Unified Matrix Treatment of the Fast Walsh-Hadamard Transform. IEEE Transactions on Computers C-25(11), 1142–1146 (Nov 1976)
43. Gentry, C.: A fully homomorphic encryption scheme. Ph.D. thesis, Stanford University (2009), crypto.stanford.edu/craig
44. Gentry, C.: Fully homomorphic encryption using ideal lattices. In: ACM STOC. pp. 169–178 (2009)
45. Gentry, C., Halevi, S., Smart, N.P.: Fully homomorphic encryption with polylog overhead. Crypt. ePrint Archive, Report 2011/566 (2011)
46. Grover, C., Ling, C.: Structured module learning with errors from cyclic algebras. Crypt. ePrint Archive, Report 2019/680 (2019)
47. Halevi, S., Shoup, V.: Algorithms in helib. In: CRYPTO. pp. 554–571 (2014)
48. Halevi, S., Shoup, V.: Bootstrapping for helib. In: EUROCRYPT. LNCS, vol. 9056, pp. 641–670 (2015)
49. Halevi, S., Shoup, V.: Faster homomorphic linear transformations in helib. Crypt. ePrint Archive, Report 2018/244 (2018)
50. Harvey, D.: Faster arithmetic for number-theoretic transforms. J. Symb. Comput. 60, 113–119 (2014)
51. Haviv, I., Regev, O.: Tensor-based Hardness of the Shortest Vector Problem to within Almost Polynomial Factors. Theory of Computing 8(1), 513–531 (2012)
52. Hoffstein, J., Pipher, J., Silverman, J.H.: NTRU: A Ring-Based Public Key Cryptosystem. In: ANTS-III. pp. 267–288 (1998)
53. Kedlaya, K.S.: A construction of polynomials with squarefree discriminants. arXiv e-prints (Mar 2011)
54. Laine, K., Lauter, K.E.: Key recovery for LWE in polynomial time. Crypt. ePrint Archive, Report 2015/176 (2015)
55. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. Des. Codes Cryptography 75(3), 565–599 (2015)
56. Lauter, K., Naehrig, M., Vaikuntanathan, V.: Can homomorphic encryption be practical? In: ACM CCSW. pp. 113–124 (2011)
57. Lopez-Alt, A., Tromer, E., Vaikuntanathan, V.: On-the-Fly Multiparty Computation on the Cloud via Multikey Fully Homomorphic Encryption. Crypt. ePrint Archive, Report 2013/094 (2013)

58. Lyubashevsky, V., Peikert, C., Regev, O.: On Ideal Lattices and Learning with Errors over Rings. J. ACM 60(6), 43:1–43:35 (Nov 2013)
59. Lyubashevsky, V., Micciancio, D.: Generalized compact knapsacks are collision resistant. In: ICALP. pp. 144–155 (2006)
60. Lyubashevsky, V., Peikert, C., Regev, O.: On Ideal Lattices and Learning with Errors over Rings. In: EUROCRYPT. pp. 1–23. Springer (2010)
61. Lyubashevsky, V., Peikert, C., Regev, O.: A Toolkit for Ring-LWE Cryptography. In: EUROCRYPT. LNCS, vol. 7881, pp. 35–54. Springer (2013)
62. of Mathematics, E.: Compositum. `https://www.encyclopediaofmath.org/index.php/Compositum`, accessed: 11 March 2019
63. Micciancio, D., Sorrell, J.: Ring Packing and Amortized FHEW Bootstrapping. In: ICALP. pp. 100:1–100:14 (2018)
64. Murakami, H.: Generalization of the cyclic convolution system and its applications. In: IEEE ICASSP. vol. 6, pp. 3351–3353 (2000)
65. Nussbaumer, H.: Fast Fourier Transform and Convolution Algorithms. Springer-Verlag (1982)
66. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J.R., Pérez-González, F.: Image denoising in the encrypted domain. In: IEEE WIFS. pp. 1–6 (2016)
67. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J.R., Pérez-González, F.: Number theoretic transforms for secure signal processing. IEEE Transactions on Information Forensics and Security 12(5), 1125–1140 (May 2017)
68. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J., Pérez-González, F.: Multivariate Lattices for Encrypted Image Processing. In: IEEE ICASSP. pp. 1707–1711 (2015)
69. Pedrouzo-Ulloa, A., Masciopinto, M., Troncoso-Pastoriza, J.R., Pérez-González, F.: Camera Attribution Forensic Analyzer in the Encrypted Domain. In: IEEE WIFS. pp. 1–7 (2018)
70. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J.R., Pérez-González, F.: On Ring Learning with Errors over the Tensor Product of Number Fields. arXiv e-prints (2016)
71. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J.R., Pérez-González, F.: Multivariate Cryptosystems for Secure Processing of Multidimensional Signals. arXiv e-prints (2017)
72. Pedrouzo-Ulloa, A., Troncoso-Pastoriza, J.R., Pérez-González, F.: Revisiting Multivariate Lattices for Encrypted Signal Processing. In: ACM IH&MMSec. pp. 161–172 (2019)
73. Peikert, C.: How (Not) to Instantiate Ring-LWE. In: SCN. pp. 411–430 (2016)
74. Peikert, C., Pepin, Z.: Algebraically structured lwe, revisited. Crypt. ePrint Archive, Report 2019/878 (2019)
75. Peikert, C., Regev, O., Stephens-Davidowitz, N.: Pseudorandomness of ring-LWE for Any Ring and Modulus. In: ACM STOC. pp. 461–473 (2017)
76. Rathee, D., Mishra, P.K., Yasuda, M.: Faster PCA and Linear Regression through Hypercubes in HElib. Crypt. ePrint Archive, Report 2018/801 (2018)
77. Regev, O.: On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. In: ACM STOC. pp. 84–93 (2005)
78. Regev, O.: On Lattices, Learning with Errors, Random Linear Codes, and Cryptography. J. ACM 56(6), 34:1–34:40 (Sep 2009)
79. Samuel, P.: Algebraic Theory of Numbers. Dover Publications (2008)
80. Smart, N.P., Vercauteren, F.: Fully homomorphic SIMD operations. Des. Codes Cryptography 71(1), 57–81 (2014)
81. Vehkalahti, R., Hollanti, C., Lahtonen, J.T., Ranto, K.: On the Densest MIMO Lattices from Cyclic Division Algebras. IEEE Trans. Information Theory 55(8), 3751–3780 (2009)

82. Weston, T.: Algebraic Number Theory. https://www.math.wisc.edu/~mmwood/748Fall2016/weston.pdf, accessed: 11 March 2019
83. Yarlagadda, R.K.R., Hershey, J.E.: Hadamard Matrix Analysis and Synthesis: With Applications to Communications and Signal/Image Processing. Kluwer Academic Publishers, Norwell, MA, USA (1997)

Appendix

Summary This Appendix includes some of the possible efficiency improvements and applications which can be implemented under the multivariate rings proposed in this work.

Structure Appendix A summarizes some of the applications that our constructions enable. Appendix B revises the Full and Baby-step/giant-step strategies from [49, 76]. Appendix C introduces additional optimizations that can be applied on the example presented in Section 7.2.

A Applications for Signal Processing

For the sake of completeness, this section focuses on some of the Secure Signal Processing (SSP) applications that benefit from m-RLWE to process encrypted signals in a more efficient and secure way than under RLWE, showcasing the applicability of m-RLWE. *While these results were originally presented on weak instances of m-RLWE vulnerable to the Bootland* et al.*'s attack, they can be adapted to deal with those rings from Definition 12.*

Image filtering In image processing, filtering is one of the most common building blocks, and it can be seamlessly implemented as a cyclic multidimensional convolution. While RLWE-based cryptosystems support univariate convolutions, they need to encrypt each row or column of the image or filter separately in order to implement a 2D or 3D convolution between two encrypted images (or an image and a filter). Conversely, m-RLWE introduces a natural way to work with multidimensional linear operations, and it achieves a more compact representation of the data, as it can effectively cipher one signal value per coefficient of the encryption polynomial. As shown in [68], the time needed for an encrypted convolution with an m-RLWE-based cryptosystem is between one and two orders of magnitude faster than with its RLWE counterpart for common image sizes,[22] while the security of the former can be much higher (*whenever we work on a secure instantiation of multivariate RLWE*), due to the large degree of the multivariate polynomials.

Image denoising Another ubiquitous image processing operation is image denoising. This operation involves a linear (Wavelet) transform, a thresholding

[22] In [68], the authors implement an m-RLWE extension in C using GMP 6.0.0 and FLINT. For a filtering application with an image of size 1014×1014 and a filter of size 11×11, they achieve runtimes of 134 s with RLWE and 8 s with the extension.

non-linear operation applied to each sub-band, and an inverse transform. By resorting to 2-RLWE and a polynomial representation of the thresholding operation, it is possible to efficiently perform all these operations with a circuit of limited depth and without an intermediate decryption of the image [66]. This produces a denoised image of size 256×256 in a few minutes. *If the 2-RLWE scheme is not implemented in a weak instantiation*, the RLWE counterpart would require polynomials of large degree in each image dimension to achieve the same security level, which renders the computation several orders of magnitude slower than with a 2-RLWE cryptosystem.

Camera Forensic Analyzer 2-RLWE has also been used to build a secure camera forensic analyzer [69], which enables outsourcing the extraction and detection of the PRNU (Photoresponse Non-Uniformity) fingerprint on encrypted images. The efficient computation of the denoising block for images proves to be a fundamental piece for this task. Additionally, as the size of the involved images is usually quite large, by using *secure* bivariate rings, it is possible to very efficiently perform 2-dimensional correlations, while also taking advantage of this size to increase the underlying lattice dimension.

Increased flexibility in image processing Finally, it is worth noting that the additional degrees of freedom that m-RLWE introduces give more flexibility to cope with signals with different structures, which is plainly impossible with the regular RLWE. In [71, 72], mechanisms for converting across different signal structures and perform efficient block processing are shown. Hence, m-RLWE enables (a) better packing schemes by grouping image pixels in blocks (e.g., for encrypted JPEG de-/compression by using block Discrete Cosine Transforms), or video sequences in frames, (b) encrypted multi-dimensional transforms that can work on a block-by-block basis taking advantage of the large signal dimensionality to increase the cryptosystem security with respect to their RLWE counterpart, (c) the use of the extra-variables to encode additional information which can be used to homomorphically evaluate encrypted divisions in the signal values, (d) flexible changes of the signal structure to update the packing and organization of the blocks, in order to seamlessly enable different operations on different dimensions.

B Full and Baby-step/giant-step

In a recent paper [49, 76], Halevi and Shoup introduce several improvements on the operations with automorphisms and their associated switching key matrices, implemented in HElib. To this aim, they take advantage of the underlying algebraic structure that can be found on the group of automorphisms in RLWE. Specifically, they exploit the fact that these automorphisms can have a multidimensional structure [45] which depends on the group $\mathbb{Z}_m^*/\langle t \rangle$.

The HElib library considers a "basis" $g_1, \ldots, g_d \in \mathbb{Z}_m^*$ where each element has "order" $D_1, \ldots, D_d$, respectively (each D_i is a positive natural number). This basis induces the following representation for the elements belonging to $\mathbb{Z}_m^*/\langle t \rangle$:

$$\{g_1^{e_1} \ldots g_d^{e_d} : 0 \leq e_i < D_i, i = 1, \ldots, d\}.$$

Due to the existing bijection between the slots and vectors $(e_1, \ldots, e_d)$ now we can independently apply rotations[23] in each different "hypercolum" i (where $i = 1, \ldots, d$) by means of one (if the i-th hypercolumn is a good dimension) or two (if the i-hypercolumn is a bad dimension) automorphisms.[24]

Without exploiting this multidimensional structure, we would have to work with a total of $\phi(m)$ different matrices to represent all the available automorphisms in the ring R; in a practical scenario, $\phi(m)$ can easily be above one or two thousand. However, by taking advantage of the different dimensions, we could represent the different automorphisms with as many as $\sum_{i=1}^{d} D_i$ matrices, and roughly increasing the number of required switching key operations by a factor of d.

In [49], the authors describe two main strategies for working in each of these dimensions:

- Full strategy: D_i matrices are needed for a dimension i and produce a cost of one or two automorphisms/switching key operations depending on whether i is a good or bad dimension.
- Baby-step/giant-step: $g + \lceil D_i/g \rceil - 1$ (roughly $\mathcal{O}(\sqrt{D_i})$) matrices are needed for a dimension i where $g = \lceil \sqrt{D_i} \rceil$; this yields a cost of two or three automorphisms/switching key operations depending on whether i is a good or bad dimension.

The HElib library [49] works by default with the full strategy for those dimensions of length at most 50 and with the baby-step/giant-step for higher lengths.

As an example, in [49] the authors report runtimes for the parameters $m = 15709$ where $\phi(m) = 15004$, $r = 22$ and only one dimension with $D = 682$, hence working with 682 slots. With a full strategy and considering a good dimension we would have a total of 681 matrices; and 51 matrices with a baby-step/giant-step strategy (682 and 52 matrices considering a bad dimension).

C Further Optimizations

The scheme we have chosen to exemplify the use of multivariate rings with RLWE in Section 7.2 of the main article can be further optimized. We based our choice on the scheme introduced in [16] for simplicity and clarity, but many other options could be taken into account. For example, in [37] the authors provide a detailed comparison among four of the main variants which are currently used in the literature: BGV [17, 15], NTRU [57] and their corresponding scale-invariant versions [14] which are, respectively, FV [41] and YASHE [13].

The use of a scale-invariant version simply involves additional division and rounding operations over the polynomial coefficients; these operations can be seamlessly addressed when working with multivariate polynomials.

[23] A rotation by h in the dimension i is defined as a map from the slot associated with $(e_1, \ldots, e_i, \ldots, e_d)$ to the slot $(e_1, \ldots, e_i + h \bmod D_i, \ldots, e_d)$.

[24] We say that the i-th hypercolumn is a good dimension if the order of g_i in $\mathbb{Z}_m^*$ is D_i; otherwise it is considered a bad dimension.

The main optimizations which are considered for the comparison in [37] are modulus switching and key switching [18]. The first one has been used in RLWE to work with leveled SHE schemes [17, 15], and it requires a chain of decreasing moduli in such a way that, after each homomorphic multiplication, a switch to a smaller modulus is performed. The effect of this operation is a notable reduction in the noise increase after each multiplication. Similarly to scale-invariant schemes, the use of modulus switching requires division and rounding operations over the coefficients of the polynomials.

Regarding the key switching operation, its use removes the dependency between the number of polynomial elements in the ciphertexts and the depth of the evaluated circuits. It is also used when working with automorphisms, where it helps to recover the ciphertexts under the original secret key.

Both modulus and key switching can be extended to work with multivariate polynomials. Firstly, division and rounding can be directly applied over the coefficients of multivariate polynomials, and secondly, switching key matrices can be analogously generated with multivariate polynomials.

Finally, an additional "optimization" which we could incorporate is the use of bootstrapping to obtain a FHE scheme, hence removing the upper bound on the depth of the evaluated circuits. For this purpose, conventional procedures could be applied over the SHE scheme, mainly consisting of homomorphically evaluating the decryption circuit by having access to an encrypted version of the secret key.

After Gentry's seminal work [44, 43], different improvements on the use of bootstrapping have appeared in the literature, varying from the recryption of binary gates [38, 9, 29, 30, 11, 31] to the optimization of the depth of the decryption circuit for RLWE-based SHE schemes [48, 21, 25]. An interesting follow-up work would be to study the behavior of our multivariate scheme with these different approaches.