

Ring-LWE in Polynomial Rings

Léo Ducas and Alain Durmus*

ENS, Dépt. Informatique, 45 rue d'Ulm, 75005 Paris, France.

Abstract. The Ring-LWE problem, introduced by Lyubashevsky, Peikert, and Regev (Eurocrypt 2010), has been steadily finding many uses in numerous cryptographic applications. Still, the Ring-LWE problem defined in [LPR10] involves the fractional ideal R^\vee , the dual of the ring R , which is the source of many theoretical and implementation technicalities. Until now, getting rid of R^\vee , required some relatively complex transformation that substantially increase the magnitude of the error polynomial and the practical complexity to sample it. It is only for rings $R = \mathbb{Z}[X]/(X^n + 1)$ where n a power of 2, that this transformation is simple and benign.

In this work we show that by applying a different, and much simpler transformation, one can transfer the results from [LPR10] into an “easy-to-use” Ring-LWE setting (*i.e.* without the dual ring R^\vee), with only a very slight increase in the magnitude of the noise coefficients. Additionally, we show that creating the correct noise distribution can also be simplified by generating a Gaussian distribution over a particular extension ring of R , and then performing a reduction modulo $f(X)$. In essence, our results show that one does not need to resort to using any algebraic structure that is more complicated than polynomial rings in order to fully utilize the hardness of the Ring-LWE problem as a building block for cryptographic applications.

1 Introduction

Since its recent introduction, the *Ring-LWE* problem [LPR10] has already been used as a building block for numerous cryptographic applications. In addition to its original functionality as the basis of efficient lattice-based cryptosystems [LPR10], it has since been used as a hardness assumption in the constructions of efficient signature schemes [MP11,Lyu11], fully-homomorphic encryption schemes [BV11b,BV11a,BGV11,GHS11], pseudo-random functions [BPR11], protocols for doing secure multi-party computation [DPSZ11,LATV11], and also gives an explanation for the hardness of the NTRU cryptosystem [SS11].

A very natural way in which one would like to be able to define the (decisional) Ring-LWE problem is as follows: for a polynomial ring $R_q = \mathbb{Z}_q[X]/(f(X))$ and a random polynomial $w \in R_q$, it is computationally hard to distinguish the uniform distribution over $R_q \times R_q$ from ordered pairs of the form $(a_i, a_i w + e_i)$, where a_i are uniformly distributed in R_q and e_i are polynomials in R whose coefficients are independently distributed Gaussians. Unfortunately, the results from

* This work was partially supported by the European Research Council.

[LPR10] do not directly imply that the above problem is hard based on the worst-case hardness of lattice problems, except in the one case when $f(X) = X^n + 1$ for n a power of 2, and thus most papers that use the Ring-LWE problem only use this one specific ring. The reason for this limitation is that the problem statement in [LPR10] requires w to be in the *dual* ring of R (which is a fractional ideal) and for the distribution of the noise to be a spherical Gaussian in the *embedding* representation of R . And it is only in the case that $R = \mathbb{Z}[X]/(X^n + 1)$ that the dual ring is simply a scaling of R (thus, one can simply multiply by the scaling and end up in R) and the embedding is just a rigid rotation and a scaling (thus the spherical Gaussian distribution is not affected by the transformation). For *all* other cyclotomic polynomials, while it is possible to transform the problem that was proved hard in [LPR10] to the one described above, the transformation between the polynomial and embedding representations involves multiplication by a skewed matrix, and the dual of R is a (possibly very) skewed fractional ideal of R . Therefore there is no obvious way to generate the noise directly in the ring R , nor work entirely in the ring R without utilizing a transformation that can substantially increase the magnitude of the error polynomials.

A natural question to ask at this point is whether there is ever a reason to use a ring other than $R_q = \mathbb{Z}_q[X]/(X^n + 1)$. While it's true that this ring has some very nice features, and we believe that it should be used whenever possible, there are situations where an alternative may be preferable. Since $X^n + 1$ is only irreducible when n is a power of 2, these polynomials are scarce. Thus it is conceivable that to achieve a certain security level, it may be advantageous to try to find a polynomial of some particular degree rather than round up to the next power of 2. A different, and a probably even stronger reason to use a different ring, is that other cyclotomic polynomials may have a more desirable structure for the task at hand. An example of this is the recent result of Gentry, Halevi, and Smart [GHS11] who show that there are particular cyclotomic polynomials that allow for much faster (at least asymptotically) instantiations of fully-homomorphic encryption. Their hardness assumption is that the Ring-LWE problem, instantiated with polynomial rings as in our description above, is a difficult problem. Using the result of our current paper, it can actually be shown that their scheme has tight connections to worst-case lattice problems (modulo a small change in the way the errors are generated, but this can be easily remedied).

1.1 Our Results

Our main result (Theorem. 2) essentially shows that for any cyclotomic polynomial $\Phi_m(X)$, one can work entirely in the ring $\mathbb{Z}[X]/(\Phi_m)$, and generate the noise distribution without resorting to complex embeddings.

Our analysis (Sect. 5) shows that for primes m (and even wider class) our simplification comes at almost no cost in term of algorithmic simplicity, tightness and efficiency compared to the scarce class of m that are powers of 2 as used for practical application in [LPR10]; thus increasing the density of usable $m < M$ from $\mathcal{O}(\log(M)/M)$ to $\mathcal{O}(1/\log(M))$.




Fig. 1. Mappings Between Different Representations (see Sect. f 2 or formal definitions. The polynomial Θ_m is defined to be $X^m - 1$ if m is odd, and $X^{m/2} + 1$ when m is even.

Our main result is a consequence of two theorems with surprisingly elementary proofs. The first theorem (see Section 4) states that every cyclotomic ring of integers $R = \mathbb{Z}[X]/(\Phi_m) \cong \mathbb{Z}[\zeta_m]$ contains mR^\vee , where R^\vee is its dual (if m is even, it actually contains $\frac{m}{2}R^\vee$). What this means is that one can scale everything that is in R^\vee by a factor of m (or $m/2$) and end up in the ring R . Similarly, if something were uniform, either statistically or computationally, modulo R^\vee , then m times it will be uniform modulo mR^\vee and thus uniform modulo R , since mR^\vee is an additive subgroup of R . This transformation is not completely tight (except in the case that $\Phi_m(X) = X^{m/2} + 1$) because we end up with something that is uniform modulo a subgroup of R , whereas we only use the randomness modulo R . This loss of tightness, however, is very small, resulting in the noise being at most $\sqrt{m/\phi(m)}$ “larger than necessary” (see the discussion after Theorem 2).

Our second theorem (see Section 5) deals with the noise generation. In the Ring-LWE definition of [LPR10], the noise needs to be a spherical Gaussian in the *canonical embedding* representation of the ring $\mathbb{Q}[X]/(\Phi_m)$ (see Figure 1), and to convert it to the polynomial representation, one needs to perform transformation $\sigma^{-1} \circ T$, where σ^{-1} is the multiplication by the inverse of a complex Vandermonde matrix (and T is a multiplication by a very simple matrix). Ideally, one would like to avoid working with the complex numbers and generate the noise by simply drawing it from the ring $\mathbb{Q}[X]/(\Phi_m)$; but unfortunately this method does not lead to the correct distribution in the embedding representation. What we show is that an almost equally simple way of generating the noise does lead to the correct distribution. We consider the ring $\mathbb{Q}[X]/(\Theta_m)$, where $\Theta_m(X) = X^m - 1$ if m is odd, and $X^{m/2} + 1$ if m is even (notice that Φ_m is a factor of Θ_m). We then show that the transformation denoted by $T^{-1} \circ \gamma$ from $\mathbb{Q}[X]/(\Theta_m)$ to the embedding representation actually preserves the spherical Gaussian distribution! This means that one can sample in $\mathbb{Q}[X]/(\Theta_m)$ by picking each coefficient independently from a continuous Gaussian distribution (rounded to \mathbb{Q} , see details in 2), and it will be the correct distribution required by [LPR10]. Then to move the noise from $\mathbb{Q}[X]/(\Theta_m)$ to $\mathbb{Q}[X]/(\Phi_m)$, one simply performs the transformation β , which is just a reduction modulo Φ_m .

In addition to making our noise generation much simpler to implement, the reduction modulo Φ_m is also simpler to analyze than $\sigma^{-1} \circ T$. This allows us to make several improvements in constructions that use rings other than

$\mathbb{Z}[X]/(X^n + 1)$ (for rings $\mathbb{Z}[X]/(X^n + 1)$, the mapping β is just the identity, and so there is nothing to analyze). As realized in previous works that used ideal lattices (e.g. [LM06,Gen10,GHS11]), multiplication in polynomial rings increases the size of the coefficients by a factor that depends on the size of the coefficients in the multiplicands, and also on the ring itself, and the ring in which the coefficients grow the least is $\mathbb{Z}_q[X]/(X^n + 1)$. As a consequence, if one were to, for example, implement the encryption scheme from [LPR10] in the ring $\mathbb{Z}[X]/(\Phi_p)$ for some prime p , one would observe that the noise grows by a factor of approximately $\sqrt{2}$ larger than in the ring $\mathbb{Z}[X]/(X^n + 1)$. We show that by analyzing the noise in the ring $\mathbb{Q}[X]/(\Phi_p)$, one can actually remove some of the noise that is introduced by the reduction modulo Φ_m ; it seems that our strategy makes the coefficients grow only $(1 + o(1))$ times as much (see Section 6).

2 Preliminaries

Cyclotomic Ring Let ζ_m be a primitive m^{th} root of unity and the cyclotomic polynomial $\Phi_m(X) \in \mathbb{Q}[X]$ be its minimal monic polynomial. Thus m is the smallest integer for which $\zeta_m^m = 1$ and Φ_m is the rational polynomial with the smallest degree of which ζ_m is a root. It is known that $\Phi_m \in \mathbb{Z}[X]$ and the other roots of Φ_m (the conjugates of ζ_m) are the elements of the set $\{\zeta_m^k \mid k \in \mathbb{Z}_m^*\}$. Thus, Φ_m has degree $\phi(m)$, the totient of m . So, the number field $\mathbb{Q}(\zeta_m)$, which we will call the m^{th} cyclotomic field, has degree $\phi(m)$ and its power basis is $\{1, \zeta_m, \dots, \zeta_m^{\phi(m)-1}\}$.

Extension of the Cyclotomic Ring For a given each integer m we define the polynomial $\Theta_m(X)$ as $X^m - 1$ if m is odd, and $X^{m/2} + 1$ when m is even. It gives a natural ring extension $\mathbb{Z}[X]/(\Theta_m)$ of the cyclotomic ring $\mathbb{Z}[X]/(\Phi_m)$: as Φ_m is a factor of Θ_m , the reduction modulo Φ_m , noted β is a ring morphism (it preserve both sum and product). The power basis of $\mathbb{Z}[X]/(\Theta_m)$ is $\{1, \zeta_m, \dots, \zeta_m^{m-1}\}$ when m is odd and $\{1, \zeta_m, \dots, \zeta_m^{\frac{m}{2}-1}\}$ when m is even.

Ring of integers The ring of integers of $\mathbb{Q}(\zeta_m)$ is $\mathbb{Z}[\zeta_m] \cong \mathbb{Z}[X]/(\Phi_m)$. According to the following theorem from [Con09, Theorem 3.7], the dual (or co-different ideal) of $\mathbb{Z}[\zeta_m]$, denoted by $\mathbb{Z}[\zeta_m]^\vee$, is the fractional ideal $\frac{1}{\Phi'_m(\zeta_m)}\mathbb{Z}[\zeta_m]$, where Φ'_m is the derivative of Φ_m . While the dual has many nice properties and is extensively used in the proof of the hardness of Ring-LWE in [LPR10], in the current paper we only need its definition.

Embeddings of cyclotomic fields The field $\mathbb{Q}(\zeta_m) \simeq \mathbb{Q}[X]/(\Phi_m)$ has exactly $\phi(m)$ embeddings $(\sigma_k)_{k \in \mathbb{Z}_m^*}$, defined by $\sigma_k : x \mapsto x(\zeta_m^k)$, for $k \in \mathbb{Z}_m^*$. The canonical embedding $\sigma : \mathbb{Q}(\zeta_m) \rightarrow \mathbb{C}^{\phi(m)}$ is defined as the direct sum of all the embeddings : $\sigma(x) = \bigoplus_{k \in \mathbb{Z}_m^*} \sigma_k(x)$. Note that that for each $k \in \mathbb{Z}_m^*$ and any

$x \in \mathbb{Q}(\zeta_m)$, we have $\sigma_{-k}(x) = \overline{\sigma_k(x)}$. Thus for a proper indexation of \mathbb{Z}_m^* the image H of σ is the \mathbb{Q} vector space generated by the columns of $\sqrt{2} \cdot T$ where :

$$T = \frac{1}{\sqrt{2}} \begin{pmatrix} \text{Id}_{\phi(m)/2} & \mathbf{i} \text{Id}_{\phi(m)/2} \\ \text{Id}_{\phi(m)/2} & -\mathbf{i} \text{Id}_{\phi(m)/2} \end{pmatrix} \quad \text{with } \mathbf{i} = \sqrt{-1}$$

In other words, for any element $x \in \mathbb{Q}(\zeta_m)$, there exists a vector $v \in \mathbb{Q}^{\phi(m)}$ such that $\sigma(x) = \sqrt{2}Tv$, and vice versa. For the rest of the paper, we will consider the column vectors of T as the canonical basis for the embedding space H .

Gaussian Distributions By ψ_s we denote the Gaussian distribution with mean 0 and standard deviation s over \mathbb{R} ; and by ψ_s^d the spherical Gaussian distribution over \mathbb{R}^d of the vector (v_1, \dots, v_d) where each coordinate is drawn independently from ψ_s .

For our purpose, one would like the Gaussian distributions to be defined over \mathbb{Q} rather than \mathbb{R} , so that an element drawn from $\psi_s^{\phi(m)}$ may be seen as element of the field $\mathbb{Q}(\zeta_m)$. The theoretical solution to that issues is to work with the tensor product $\mathbb{Q}(\zeta_m) \otimes_{\mathbb{Q}} \mathbb{R}$ as done in [LPR10].

However, in practice elements needs to be represented finitely, typically using floating points numbers of a fixed mantissa. For simplicity we choose this solution: we consider that output of Gaussian distribution ψ_s^d are rounded off to rational numbers using a fine enough grid so that all our results go through except with a negligibly small probability.

3 The Main Result

In this section we give the main result of this paper. We describe a distribution over $R_q \times R_q$, where $R_q = \mathbb{Z}_q[X]/(\Phi_m)$ which is computationally indistinguishable from the uniform distribution over $R_q \times R_q$ based on the worst-case hardness of the approximate shortest vector problem in ideal lattices. The proof of our theorem will use results that we later prove in Sections 4 and 5 that will aid us in transforming the hard Ring-LWE problem defined in [LPR10] into one in which all operations are performed in polynomial rings.

Theorem 1 ([LPR10]). *Let m be integer, and q be a prime congruent to 1 modulo m . Let denote K be the number field $\mathbb{Q}(\zeta_m)$, $R = \mathbb{Z}[\zeta_m]$ be its ring of integers, R^\vee be the fractional ideal $\frac{1}{\Phi_m(\zeta_m)}\mathbb{Z}[\zeta_m]$, q be a prime congruent to $1 \pmod{m}$. Also, let k be any positive integer and $\alpha \in (0, 1)$ be a real number such that $\alpha q > \omega(\sqrt{\log m})$. If there exists an algorithm that can solve the decisional Ring-LWE problem, that is distinguish (with some advantage $1/\text{poly}(m)$) between k uniformly random samples drawn from $R/qR \times K/R^\vee$ and k samples $(a_i, \frac{a_i w}{q} + e_i) \in R/qR \times K/R^\vee$ where a_i are chosen uniformly at random from R/qR , w is chosen uniformly at random from R^\vee/qR^\vee , and the e_i are sampled in the embedding space H from the distribution $\psi_s^{\phi(m)}$ for $s = \alpha \cdot \left(\frac{\phi(m)k}{\log(\phi(m)k)} \right)^{1/4}$,*

then there exists a quantum algorithm that runs in time $O(q \cdot \text{poly}(m))$ that solves the approximate Shortest Vector Problem to within a factor $\tilde{O}(\sqrt{m}/\alpha)$ in any ideal of the ring $\mathbb{Z}[\zeta_m]$.

Before stating our main theorem, we believe that it would be helpful to first understand why everything turns out to be so simple and convenient when working with the ring of integers $\mathbb{Z}[\zeta_m]$ when m is a power of 2 (and not so convenient otherwise). If m is a power of 2, then $\Phi_m = x^n + 1$, where $n = m/2$, and therefore $\Phi' = nX^{n-1}$, and so $\Phi'(\zeta_m) = n\zeta_m^{n-1}$. The last equation implies that $n\zeta_m^{n-1}R^\vee = R$ (and since $\zeta_m^j R = R$ for any integer j , we have $nR^\vee = R$), which gives us a very simple way to remove the ring R^\vee and work entirely in the ring R . When given a sample $(a_i, \frac{a_i w}{q} + e_i) \in R/qR \times K/R^\vee$, we can simply multiply the second element of the ordered pair by n and get $(a_i, \frac{a_i w n}{q} + e_i n) \in R/qR \times K/nR^\vee$. Now we observe that since the e_i were chosen from the distribution $\psi_s^{\phi(m)}$, the ne_i are distributed according to $\psi_{ns}^{\phi(m)}$. And since w was chosen uniformly at random from R^\vee/qR^\vee , we have that nw is uniformly random in R/qR . Thus the problem of distinguishing uniformly random samples in $R/qR \times K/R$ from samples $(a_i, \frac{a_i w'}{q} + e'_i) \in R/qR \times K/R$ where a_i and w' are drawn uniformly from $R/(q)$ and the e'_i are drawn according to the distribution $\psi_{ns}^{\phi(m)}$ is exactly equivalent to the problem from Theorem 1. We now turn to how one would generate the errors e'_i directly in the ζ_m power basis, without first generating them in the embedding space and then doing the transformation. The main observation here is that the linear transformation $\sigma^{-1} \circ T$ (see Figure 1) from the embedding space H to the power basis representation turns out to be a multiplication by a scaled orthogonal matrix. Therefore, the spherical Gaussian distribution in H remains a spherical Gaussian distribution in the power basis representation, and can therefore be sampled directly in the latter domain.

On the other hand, if ζ_m is a primitive root of unity for any other m except a power of 2, then neither of the above-described conditions hold. It is still possible to multiply elements in R^\vee by $\Phi'(\zeta_m)$ in order to take them into R , but this transformation does not result in “nice” distributions in the power basis of R . It is known that there exist cyclotomic polynomials Φ_m whose coefficients are of the order of $m^{\log m}$, and thus Φ'_m also has coefficients of that magnitude. Therefore when multiplying an element by $\Phi'(\zeta_m)$, the coefficients of the product in the power basis will also very likely have such large coefficients, and thus the noise will increase by a super-polynomial factor. And even for simple cyclotomic polynomials such as Φ_p for some prime p , its derivative will have $\Omega(p)$ coefficients of size $\Omega(p)$, and so the multiplication by Φ'_p could increase the coefficients by a factor of p^2 . Additionally, if ζ_m is a primitive root of unity and m is not a power of 2, then the mapping $\sigma^{-1} \circ T$ from the embedding space H to the power basis representation is no longer an orthogonal linear map, and thus the spherical Gaussian distribution is no longer preserved.

Theorem 2 (Main Theorem). *Let m be an integer, and let R_q be the ring $\mathbb{Z}_q[X]/(\Phi_m)$ where q is a prime congruent to 1 modulo m . Also, let k be any positive integer, $\alpha \in (0, 1)$ be a real number such that $\alpha q > \omega(\sqrt{\log m})$, and*

define m' to be equal to m if m is odd and $m/2$ if m is even. If there is an algorithm that can solve the Ring-LWE problem, that is distinguish (with some advantage $1/\text{poly}(m)$) between k uniformly random samples drawn from $R_q \times R_q$ and k samples $(a_i, a_i w + e_i) \in R_q \times R_q$, where a_i and w are chosen uniformly at random from R_q and $e_i = \lceil e'_i \bmod \Phi_m \rceil$ with $e'_i \in \mathbb{Q}[X]/(\Theta_m)$ is distributed as $\psi_s^{m'}$ for $s = \sqrt{m'}\alpha q \left(\frac{\phi(m)k}{\log(\phi(m)k)}\right)^{1/4}$; then there exists a quantum algorithm that runs in time $O(q \cdot \text{poly}(m))$ that solves the approximate Shortest Vector Problem to within a factor $\tilde{O}(\sqrt{m}/\alpha)$ in any ideal of the ring $\mathbb{Z}[\zeta_m]$.

Before we give the proof of this theorem (which uses results from Sections 4 and 5), we would like to draw the reader's attention to several things.

First, we emphasize that the error distribution is generated by sampling a polynomial $g_0 + g_1 X + \dots + g_{d-1} X^{m'-1} \in \mathbb{Q}[X]/(\Theta_m)$ where g_i simply are independent Gaussian variables, then reducing modulo Φ_m , and only then rounding each coefficient to the nearest integer. While it would have been slightly more convenient to be able to round and then do a reduction modulo Φ_m , the two distributions are not equivalent.

Secondly, we point out that by using a Lemma similar to [ACPS09, Lemma 2], it can be shown that instead of choosing the secret w uniformly from R_q , it can be drawn from the same distribution as the error vectors e_i . The only consequence of this is that the value of k in the theorem increases by one.

A third comment is that just as in Theorem 1, the $\left(\frac{\phi(m)k}{\log(\phi(m)k)}\right)^{1/4}$ term in the standard deviation of the error is a consequence of converting elliptic distributions into spherical ones in [LPR10]. It is unclear whether having this term is actually necessary for hardness or whether the elliptical distributions in [LPR10] are an artifact of the proof, and so in practice it may be enough to just sample with standard deviation $\sqrt{m'}\alpha q$. Fortunately, most constructions involving Ring-LWE only require a small (usually a constant or a logarithmic) number of samples, and so for theoretical applications when one does not care too much about small polynomial factors, this term does not cause too much trouble.

The final comment that we would like to make is about the "tightness" of our reduction. It is natural to wonder whether our transformation from Ring-LWE in the domain in Theorem 1 to the one in the domain in Theorem 2 is tight, in the sense that one did not need to add more noise than necessary in order to obtain pseudo-randomness in $R_q \times R_q$. We now give an intuition for why the transformation is actually rather tight. Ignoring the $\left(\frac{\phi(m)k}{\log(\phi(m)k)}\right)^{1/4}$ term, which is a possibly removable artifact carried over from Theorem 1, the required noise in our new theorem is $\sqrt{m'}\alpha q$, where there is a requirement that $\alpha q > \omega(\sqrt{\log m})$. Thus the noise must have standard deviation at least $\omega(\sqrt{m' \log m})$. This is almost tight because by the result of Arora and Ge [AG11], if the standard deviation were $o(\sqrt{\phi(m)})$, then the Ring-LWE problem could be solved in sub-exponential time $2^{o(\phi(m))}$, which would then imply that the Shortest Vector Problem could be solved in sub-exponential time as well. And since $\sqrt{m'}/\phi(m) =$

$O(\sqrt{\log \log m})$, this is essentially the maximum tightness factor that we lose during our reduction.

Proof of Theorem 2 To prove the theorem, we will show how one can transform the samples from Theorem 1 into samples from the ring $R_q \times R_q$. Given samples of the form $(a_i, a_i w/q + e_i) \in \mathbb{R}_q \times \mathbb{Q}(\zeta_m)/R^\vee$ where a_i are chosen uniformly at random from R_q , w is chosen uniformly at random from R_q^\vee , and the e_i are sampled from the distribution $\psi_s^{\phi(m)}$ in the embedding space H , we scale the second element of each ordered pair by a factor of $m'q$ to obtain elements $(a_i, a_i w m' + q m' e_i) = (a_i, a_i w' + e'_i) \in R_q \times \mathbb{Q}(\zeta_m)/q m' R^\vee$ where w' is distributed uniformly at random in $m' R^\vee / m' q R^\vee$, and e'_i are sampled from the distribution $\psi_{s m' q}^{\phi(m)}$. Since we did nothing but scaling at this point, it is clear that distinguishing these ordered pairs from uniform ones in $R_q \times \mathbb{Q}(\zeta_m)/q m' R^\vee$ is as hard as the original problem from Theorem 1. We now apply Theorem 3 which states that $m' R^\vee \subseteq R$ to conclude that if we reduce the second entry of the ordered pairs modulo qR to obtain elements $(a_i, a_i w' + e'_i) \in R_q \times \mathbb{Q}(\zeta_m)/qR$ where w' is distributed uniformly at random in $m' R^\vee / qR$, and e'_i are sampled from the distribution $\psi_{s m' q}^{\phi(m)}$, the distinguishing problem is at least as difficult as before.

We now make the observation that instead of choosing w' uniformly at random from $m' R^\vee / qR$, we can choose it from R/qR without making the problem any easier. The reason is that given a pair $(a_i, a_i w' + e'_i)$, we can choose a uniformly random $w'' \in R/qR$ and output $(a_i, a_i w' + a_i w'' + e'_i) = (a_i, a_i (w' + w'') + e'_i)$, and the secret $w' + w''$ is uniform in R/qR . We can also observe that if we consider the element $a_i w' + e'_i$ in the power-basis representation and round each coefficient to the nearest integer, it is equivalent to only rounding the error term e'_i to the nearest integer because the product $a_i w'$ already has integer coefficients. Thus the problem of distinguishing rounded elements $(a_i, a_i s + \lceil e'_i \rceil) \in R_q \times R_q$ from random elements in $R_q \times R_q$ is at least as difficult as the problem from Theorem 1. The last thing we need to address is the noise generation. Currently, the e'_i are generated from the distribution $\psi_{s m' q}^{\phi(m)}$ in the embedding space H . Theorem 5 states that to obtain such a distribution, it is equivalent to sample the distribution $g_0 + g_1 X + \dots + g_{m'-1} X^{m'-1} \in \mathbb{Q}[X]/(\Theta_m)$ where each g_i is a normally distributed random variable with mean 0 and standard deviation that is $\sqrt{m'}$ times smaller than that required in the distribution in the embedding space H . And this is exactly the distribution from which the errors come from in the statement of our Theorem. \square

4 Mapping $\mathbb{Z}[\zeta_m]^\vee$ to $\mathbb{Z}[\zeta_m]$

In this section we prove that the element $\frac{m'}{\Phi_m(\zeta_m)}$, for $m' = m$ when m is odd and $m/2$ when it is even, is an element of the ring $\mathbb{Z}[\zeta_m]$, which implies that the ring $\mathbb{Z}[\zeta_m]$ contains $m' \mathbb{Z}[\zeta_m]^\vee$.

Theorem 3. *For $R = \mathbb{Z}[\zeta_m]$, we have $m' R^\vee \subseteq R$, where $m' = m$ if m is odd and $m/2$ if m is even.*

Proof: Let $\Theta_m(X)$ be the polynomial $X^m - 1$, if m is odd, and $X^{m/2} + 1$ if m is even. Then it is easily seen that $\Phi_m(X)$ is a factor of $\Theta_m(X)$, and we can write $\Theta_m(X) = \Phi_m(X)g(X)$ for some polynomial $g(X) \in \mathbb{Z}[X]$. By taking the derivative of both sides, we obtain the equation

$$m'X^{m'-1} = \Phi'_m(X)g(X) + \Phi_m(X)g'(X),$$

or equivalently,

$$m'X^{m'} = X\Phi'_m(X)g(X) + X\Phi_m(X)g'(X).$$

Evaluating both sides at ζ_m , we obtain

$$\pm m' = \zeta_m \Phi'_m(\zeta_m)g(\zeta_m) + \zeta_m \Phi_m(\zeta_m)g'(\zeta_m) = \zeta_m \Phi'_m(\zeta_m)g(\zeta_m)$$

since $\zeta_m^{m'} = 1$ when $m' = m$ and -1 when $m' = m/2$, and $\Phi_m(\zeta_m) = 0$. Now, using the definition that $R^\vee = \frac{1}{\Phi'_m(\zeta_m)}R$, we obtain

$$m'R^\vee = \frac{m'}{\Phi'_m(\zeta_m)}R = \pm \zeta_m g(\zeta_m)R \subseteq R,$$

where the last inclusion is true because $g(X) \in \mathbb{Z}[X]$, and so $g(\zeta_m) \in R$. \square

We get that if we multiply the different ideal by m' , we find a set included in the ring of integer. In fact, we prove in Appendix B that m' is the smallest integer which verifies this property. It mainly comes from the fact that m' is the radical of the finite group R^\vee/R , namely the least common multiple of orders of the elements in this group. And we get eventually this following characterization:

Theorem 4. *A integer k is such that $kR^\vee \subset R$ if and only if m' divides k .*

5 Geometry and Error Sampling

For the rest of the paper, let $m' \in \mathbb{Z}$ denote $m/2$ if m is even, and m if m is odd.

To obtain the correct distribution of the error polynomials in the Ring-LWE problem in Theorem 1, we want the noise distribution over $\mathbb{Q}[X]/(\Phi_m)$ to map to a spherical Gaussian in the embedding space H . This is not a problem if the map $T^{-1} \circ \sigma$ is a scaled-orthonormal map, which is the case when m is a power of two. For a general m , a natural solution would be to generate the noise in the space H and then map it to $\mathbb{Q}[X]/(\Phi_m)$, however this requires dealing with the inverse Vandermonde matrix of σ^{-1} , making the noise generation much less efficient.

To overcome this technical issue, we use the ring extension $\mathbb{Q}[X]/(\Theta_m)$ and show that it is the natural ring for the error generation. First unlike $\mathbb{Q}[X]/(\Phi_m)$ the canonical embedding from this ring preserves sphericity of Gaussian distributions: thus one just needs to sample a spherical Gaussian in this extension then reduce modulo ϕ_m .

Theorem 5 (Geometry of $T^{-1} \circ \sigma \circ \beta$). *Let $v \in \mathbb{Q}[X]/(\Theta_m)$ be a random variable distributed as $\psi_s^{m'}$ in the power basis. Then the distribution of $(T^{-1} \circ \sigma \circ \beta)(v)$, seen in the canonical basis of H is the spherical Gaussian $\psi_s^{\phi(m)}/\sqrt{m'}$.*

Secondly, for a large class of integers m the reduction modulo Φ_m has a very simple and sparse matrix representation in the power basis. The knowledge of this matrix representation simplifies the geometric analysis of the error and products of errors, leading to some better theoretical bounds for correct decryption (see lemma 7), detailed below.

5.1 Analysis of β , the reduction modulo Φ_m

First, if B is very sparse and structured, this reduction can be implemented in a very simple ad-hoc way, while having better practical running time than general quasi-linear reduction algorithms. We will show that it is the case when $m = 2^k p$ for a any prime p , and also when $m = 2^k p' p$ if p' is a small prime.

Secondly, error distributions in the $\mathbb{Q}[X]/(\Phi_m)$ representation depend on the geometry of B , and thus the norms of B have an impact on the relation between m, s and q : the smaller the norms are, the smaller q one may choose while ensuring correct decryption. In particular, for any $e \in \mathbb{Q}[X]/(\Theta_m)$ we have: $\|\beta(e)\|_\infty \leq \|B\|_1 \|e\|_\infty$, which is related to the expansion factor inequality [LM06]. One may indeed only deal only with the expansion factor of Φ_m , and bound the error preimage in $\mathbb{Q}[X]/(\Theta_m)$. As described later, the main part of the error that needs to be dealt with for decryption has the form $ab + cd$ where a, b, c, d are drawn according to $\beta(\psi_s^{m'})$. Considering the tailcut function $\mathcal{E}(\tau) = \tau e^{1/2 - \tau^2/2}$ we have the following fact:

Fact 6 (Error Bound in the Extension Ring $\mathbb{Q}[X]/(\Theta_m)$) *Let $a, b, c, d \in \mathbb{Q}[X]/(\Theta_m)$ be distributed as $\psi_s^{m'}$. Then, $\|ab + cd\|_\infty \leq \sqrt{2m'} \tau \tau' s^2$ except with probability less than $n\mathcal{E}(\tau) + \mathcal{E}(\tau')^{2m'}$.*

Since β is ring morphism, preserving products as well as sums, this translate to $\mathbb{Q}[X]/(\Phi_m)$:

$$\|\beta(a)\beta(b) + \beta(c)\beta(d)\|_\infty = \|\beta(ab + cd)\|_\infty \leq \|B\|_1 \sqrt{2m'} \tau \tau' s^2.$$

However, the exact knowledge of B , together with the knowledge of the error distribution may lead to better bounds. While there is no simple explicit formula for B in general, some specific values of m makes B very simple. Obviously, when m is a power of two, B is the identity since $\Theta_m = \Phi_m$. When $m = 2^k p$ we have:

$$B = \begin{pmatrix} & & & -1 \\ & & & \vdots \\ & & & -1 \\ \text{Id}_{p-1} & & & \end{pmatrix} \text{ if } k = 0; \quad B = \begin{pmatrix} & & & -1 \\ & & & 1 \\ & & & \vdots \\ & & & -1 \\ \text{Id}_{p-1} & & & 1 \end{pmatrix} \otimes \text{Id}_{2^{k-1}} \text{ otherwise}$$

In that case, a better bound can be proved, replacing the constant $\|B\|_1 = 2$ by $\|B\|_2 = \sqrt{2}$.

Fact 7 (Error Bound for $m = 2^k p$) *Let p be a prime number, k a positive integer and assume $m = 2^k p$. Let $a, b, c, d \in \mathbb{Q}[X]/(\Theta_m)$ be distributed as $\psi_s^{m'}$. Then, $\|\beta(ab + cd)\|_\infty \leq 2\sqrt{m'} \tau \tau' s^2$, except with probability less than $m' \left(\mathcal{E}(\tau) + 3\mathcal{E}(\tau')^{2\lfloor m'/3 \rfloor} \right)$.*

The proof is available in the full version of this article. This statement raises the interesting question of whether it can be generalized to other values m , *i.e.* can we replace $\|B\|_1$ by $\|B\|_2$ (while keeping the exponent of $\mathcal{E}(\tau')$ big enough)? While such constant $\|B\|_2$ applies to Gaussian errors, its not clear if it applies in general for products of Gaussians.

Other polynomials Φ_m For general values of m the coefficients of B may be much bigger, and can even grow exponentially in m for product of many primes. Few is known about the behavior of the coefficients of Φ_m in terms of the prime decomposition of m , however Lam and Leung proved in [LL96] that Φ_{pq} for two primes p and q have its coefficient in $\{-1, 0, 1\}$.

A generalization of their proof gives a more detailed behavior:

Theorem 8. *If m is on the form $m = 2^k pq$ where p, q are two odd primes and $k \in \mathbb{N}$, B has coefficients in $\{-1, 0, 1\}$ and $\|B\|_1 = 2 \min(p, q)$.*

The proof will be detailed in the full version of this article.

Improved Decryption Additionally, the explicit knowledge of B can suggest strategies to improve the tolerance of the decryption algorithm. Such an idea is described when $m = p$ is a prime integer in section 6.4. It seems to improve the tolerance, replacing the $\|B\|_2 = \sqrt{2}$ factor by ≈ 1.16 for dimension $m \approx 500$; and seems to be $1 + o(1)$ when the dimension grows. With that improvement the tolerance loss compared to the encryption scheme based on Φ_{2^k} can becomes marginal.

6 Ring-LWE encryption Scheme

In this section we present an application example of our result, that is an adaptation of the [LPR10] scheme to general polynomial Φ_m , and sketch strategies to improve the decryption rate.

6.1 Definition

We consider m to be our main security parameter, and we assume it grows in an unbounded set of integer S such that $\|B\|_1$ is polynomially bounded : $\|B\|_1 \leq O(m^b)$ for some $b \geq 0$. For example, we can take $b = 0$ for the set $S = \{2^k p | k \in \mathbb{Z}, p \text{ is prime}\}$, while $S = \{2^k pq | k \in \mathbb{Z}, p, q \text{ are prime}\}$ gives $b = 1/2$.

Choose some small $\epsilon \in (0, 1/4)$, and set other parameters to grow as follow : the modulus $q = \Theta(m^{2+b+2\epsilon})$, and the standard deviation $s = \Theta(m^{3/4+\epsilon})$. Our encryption scheme is as follows :

- **Gen**(1^m) : Sample $w, e_1 \leftarrow \psi_s^{m'}$, and a uniformly in R_q . Set $\bar{w} = \lfloor \beta(w) \rfloor$ and $\bar{e}_1 = \lfloor \beta(e_1) \rfloor$. The private key $\bar{w} = \lfloor \beta(w) \rfloor \in R_q$ and the public key is (a, \bar{t}) where $\bar{t} = a\bar{w} + \bar{e}_1 \pmod q \in R_q$
- **Encrypt**($\bar{t} \in R_q, \mu \in \{0, 1\}^{\phi(m)}$) : To encrypt the message μ under the public key \bar{t} , draw $r, e_2, e_3 \leftarrow \psi_s^{m'}$, and set $\bar{r} = \lfloor \beta(r) \rfloor$, $\bar{e}_2 = \lfloor \beta(e_2) \rfloor$ and $\bar{e}_3 = \lfloor \beta(e_3) \rfloor$. Output $(u, v) \in R_q \times R_q$ where $u = a\bar{r} + \bar{e}_2 \pmod q$ and $v = \bar{t}\bar{r} + \bar{e}_3 + \mu \lfloor q/2 \rfloor \pmod q$.
- **Decrypt**($\bar{w} \in R_q, (u, v) \in R_q \times R_q$) : To decrypt (u, v) with the private key \bar{w} , compute $d = v - u\bar{w} \in R_q$, and decrypt the i -th bit μ_i as 0 if $d_i \in [-q/4; q/4]$, and as 1 otherwise.

6.2 Security

We prove semantic security based on the hardness of the approximate Shortest Vector Problem to within a factor $\tilde{O}(m^{5/2+b+\epsilon})$. For any constant number of samples k set :

$$\alpha^{-1} = \frac{\sqrt{m'}q}{s} \left(\frac{\phi(m')k}{\log(\phi(m')k)} \right)^{1/4} = O(m^{2+b+\epsilon}).$$

To fulfill the condition of theorem 2, we verify that :

$$\alpha q = s \frac{\log(\phi(m)k)^{1/4}}{\sqrt{m'}\phi(m)k^{1/4}} \geq \Theta(m^\epsilon) > \omega(\sqrt{\log m}), \quad \text{since } \frac{m}{\phi(m)} = O(\log \log m).$$

Note that we use the main theorem 2 in its modified form that replaces the uniform distribution of the secret w by the same Gaussian distribution as the error (see the discussion under the statement of theorem 2).

First, the public key distribution $(\bar{a}, \bar{t} = a\bar{w} + \bar{e}_1)$ follows the distribution defined in theorem 2 relatively to w which is distributed according to a Gaussian distribution. Thus for $k = 2$, our theorem states that this public key (\bar{a}, \bar{t}) is indistinguishable from the uniform distribution over $R_q \times R_q$.

We can now assume that (\bar{a}, \bar{t}) is uniformly random, in which case $(a, u = a\bar{r} + \bar{e}_2)$ and $(\bar{t}, v' = \bar{t}\bar{r} + \bar{e}_2)$ are two samples following the distribution of theorem 2, where \bar{r} is once again Gaussian. Using theorem 2 with $k = 3$, we deduce that (a, u) and (\bar{t}, v') are also indistinguishable from random, so is $v = v' + \mu \lfloor q/2 \rfloor$. That concludes the security proof.

6.3 Correctness

During decryption, we get :

$$\begin{aligned} d &= v - u\bar{w} = (\bar{a}\bar{w} + e_1)\bar{r} + \bar{e}_3 + \mu \lfloor q/2 \rfloor - (\bar{a}\bar{r} + \bar{e}_2) \pmod q \\ &= \bar{e}_1\bar{r} + \bar{e}_3 + \bar{e}_2\bar{w} + \mu \lfloor q/2 \rfloor \pmod q \end{aligned}$$

Thus, the decryption will be correct if $\|\bar{e}\|_\infty < q/4$ where $\bar{e} = \bar{e}_1\bar{r} + \bar{e}_3 + \bar{e}_2\bar{w}$. First, note that the rounding operations have a limited effect on the final result of the error \bar{e} : the difference between that computation with and without rounding is bounded by $\tilde{O}(B_1 m' s) = \tilde{O}(m^{7/4+b+\epsilon})$, this negligible compared to $q = \Theta(m^{2+b+2\epsilon})$. Similarly, one can neglect the contribution of \bar{e}_3 since $\|\bar{e}_3\|_\infty \leq \tilde{O}(\|B\|_1 s) = \tilde{O}(m^{3/4+b+\epsilon})$.

According to lemma 6, we have that $\|\bar{e}\|_\infty \leq \tilde{O}(\|B\|_1 \sqrt{m} s^2) \leq \tilde{O}(m^{2+b+\epsilon})$ except with negligible probability. On the other hand, q grows as $\Theta(m^{2+b+2\epsilon})$, thus, decryption is correct with overwhelming probability for large enough values of m .

6.4 Practical Improvements

For applications, any tricks to decrease the minimal value of q while preserving correct decryption might be worthwhile. We hereby presents two independent ideas.

Recovering approximation of the error preimage $e' \in \mathbb{Z}[X]/(\Theta_m)$ This first idea concerns the decryption algorithm. For simplicity, we restrict our attention to $m = p$ a prime. In this case we have for each index $i \leq p - 2$ that $\bar{e}_i = e'_i - e'_{p-1}$ where $e' = e_1 r + e_3 + e_2 w \in \mathbb{Z}[X]/(\Theta_m)$. Thus if we recover a good approximation x of e'_{p-1} we might reduce the error by adding x to each coordinate. Without warping modulo q , an approximation of e'_{p-1} may be recovered as the average $\frac{-1}{p-1} \sum_{i=0}^{p-2} \bar{e}_i$; the error should be less than $\approx \tau s^2$, using the heuristic that e' behave like a spherical Gaussian.

However, we need to consider \bar{e}_i modulo $q/2$ to get rid of the message. Our heuristic algorithm proceeds as follow : for a certain constant $\alpha \in (0, 1)$, find (one of) the smallest interval $[a, b]$ such that for at least $\alpha(p - 1)$ indexes $i \in [p - 1]$ verifies $\bar{e}_i \in ([a, b] \bmod q/2)$. Consider $a_i \in \mathbb{Z}$ as the unique integer representing $\bar{e}_i \in \mathbb{Z}_{q/2}$ in $[a, b]$, compute the average x of those a_i , and output the smallest representative of $-[x]$ modulo $q/2$ as an approximation of e'_{p-1} . Note that this algorithm can be implemented in quasi-linear time, by sorting the values e_i .

Our experiments indicates that such a strategy decrease the $\sqrt{2} \approx 1.41$ factor to ≈ 1.16 for $m = 503$ and $\alpha = 0.9$, and keeps decreasing when the dimension grows. We conjecture that it asymptotically decrease as $1 + o(1)$. Similar idea should apply to $m = 2^k p$. While this suggest that the quality loss compared to cryptosystem based on the Φ_{2^k} polynomial can be almost reduced to nothing, implementing such a error recovery strategy would require more study.

Rejection during Key Generation The second idea consist of modifying the key generation algorithm **Gen** so that the couple (s, e_1) is rejected whenever $\|(s|e_1)\| \geq \sqrt{2m'}\tau''s$, where τ'' is chosen such that $\mathcal{E}(\tau'')^{2n} \leq 1/2$; only half of them are rejected, thus the advantage of the adversary is no more than doubled. For $m' \geq 500$ this improves our bound by $\tau'/\tau'' \approx 1.4/1.05$. The same idea

applies when using the tight bound of lemma 7 by rejecting (\bar{s}, \bar{e}_1) depending on $\|B \cdot \text{Circ}(\bar{s})\|^2 + \|B \cdot \text{Circ}(\bar{e}_1)\|^2$.

References

- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai, *Fast cryptographic primitives and circular-secure encryption based on hard learning problems*, CRYPTO, 2009, pp. 595–618.
- [AG11] Sanjeev Arora and Rong Ge, *New algorithms for learning in presence of errors*, ICALP (1), 2011, pp. 403–415.
- [BGV11] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan, *Fully homomorphic encryption without bootstrapping*, Cryptology ePrint Archive, Report 2011/277, 2011, To appear at ITCS 2012.
- [BPR11] Abhishek Banerjee, Chris Peikert, and Alon Rosen, *Pseudorandom functions and lattices*, Cryptology ePrint Archive, Report 2011/401, 2011, To appear in Eurocrypt 2012.
- [BV11a] Zvika Brakerski and Vinod Vaikuntanathan, *Efficient fully homomorphic encryption from (standard) LWE*, FOCS, 2011.
- [BV11b] ———, *Fully homomorphic encryption from ring-lwe and security for key dependent messages*, CRYPTO, 2011, pp. 505–524.
- [Con09] K. Conrad, *The different ideal*, 2009, Available at <http://www.math.uconn.edu/~kconrad/blurbs/>.
- [DPSZ11] I. Damgard, V. Pastro, N.P. Smart, and S. Zakarias, *Multiparty computation from somewhat homomorphic encryption*, Cryptology ePrint Archive, Report 2011/535, 2011.
- [Gen10] Craig Gentry, *Toward basing fully homomorphic encryption on worst-case hardness*, CRYPTO, 2010, pp. 116–137.
- [GHS11] Craig Gentry, Shai Halevi, and Nigel P. Smart, *Fully homomorphic encryption with polylog overhead*, Cryptology ePrint Archive, Report 2011/566, 2011, To appear in Eurocrypt 2012.
- [LATV11] Adriana Lopez-Alt, Eran Tromer, and Vinod Vaikuntanathan, *Cloud-assisted multiparty computation from fully homomorphic encryption*, Cryptology ePrint Archive, Report 2011/663, 2011.
- [LL96] T. Y. Lam and K. H. Leung, *On the cyclotomic polynomial $\phi_{pq}(x)$* , The American Mathematical Monthly **103** (Aug. - Sep., 1996), no. 7, 562–564.
- [LM06] Vadim Lyubashevsky and Daniele Micciancio, *Generalized compact knapsacks are collision resistant*, ICALP (2), 2006, pp. 144–155.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev, *On ideal lattices and learning with errors over rings*, EUROCRYPT, 2010, pp. 1–23.
- [Lyu11] Vadim Lyubashevsky, *Lattice signatures without trapdoors*, Cryptology ePrint Archive, Report 2011/537, 2011, To appear in Eurocrypt 2012.
- [MP11] Daniele Micciancio and Chris Peikert, *Trapdoors for lattices: Simpler, tighter, faster, smaller*, Cryptology ePrint Archive, Report 2011/501, 2011, To appear in Eurocrypt 2012.
- [SS11] Damien Stehlé and Ron Steinfeld, *Making NTRU as secure as worst-case problems over ideal lattices*, EUROCRYPT, 2011, pp. 27–47.
- [Ste05] William Stein, *Introduction to algebraic number theory*, 2005, Available at <http://wstein.org/courses/>.
- [Was97] Lawrence C. Washington, *Introduction to cyclotomic fields*, Graduate Texts in Mathematics, vol. 83, Springer-Verlag, New York, 1997.

A proof of Theorem 5

Proof: We proceed by considering $G \in \mathbb{C}^{\phi(m) \times m'}$, the matrix representing the linear map γ from the power basis of $\mathbb{Z}[X]/(\Theta_m)$ to the canonical basis of $\mathbb{C}^{\phi(m)}$ and will show $G\overline{G}^t = m' \text{Id}_{\phi(m)}$. Also note that T^{-1} is hermitian, that is $T^{-1} = \overline{T}^t$, and $T^{-1} \circ \gamma$ is a real linear map. Thus $E = T^{-1}G = \overline{E}$ so $EE^t = E\overline{E}^t = T^{-1}G\overline{G}^tT = m' \text{Id}_{\phi(m)}$.

This last equation implies that if a random variable $v \in \mathbb{Q}[X]/(\Theta_m)$ has covariance $s^2 \cdot \text{Id}_{m'}$ then the covariance of $(T^{-1} \circ \gamma)(v)$ is $s \cdot E \cdot \text{Id}_{m'} \cdot \overline{E}^t = s^2 m' \cdot \text{Id}_n$: the distribution of $(T^{-1} \circ \gamma)(v)$ is the spherical Gaussian $\psi_{s\sqrt{m'}}^{\phi(m)}$.

Now we show that $G\overline{G}^t = m' \text{Id}_{\phi(m)}$: let $g_{i,j}$ for $(i,j) \in [m'] \times \mathbb{Z}_m^*$ denotes the coefficients of G , that is $g_{i,j} = \sigma_j(X^i) = \zeta_m^{ij}$. Let $c_{i,j}$ for $i,j \in \mathbb{Z}_m^*$ denote the coefficients of $C = G \cdot \overline{G}^t$. For all $i,j \in \mathbb{Z}_m^*$ we have:

$$c_{i,j} = \sum_{k \in [m']} \zeta_m^{ik} \overline{\zeta_m^{jk}} = \sum_{k \in [m']} (\zeta_m^{i-j})^k = \begin{cases} m' & \text{if } i = j, \text{ since } \zeta_m^{i-j} = 1 \\ 0 & \text{otherwise, since } \zeta_m^{i-j} \neq 1 \text{ is a } m\text{-th root of unity} \\ & \text{(or an } m'\text{-th root when } m \text{ is even)} \end{cases}$$

□

B Proof of Theorem 4

First of all, we remind some facts about free abelian groups of finite rank directly apply to $\mathbb{Z}[X]/(\Phi_m)$ and its different ideal. For conciseness, we will note R for the ring $\mathbb{Z}[X]/(\Phi_m) \cong \mathbb{Z}[\zeta_m]$ in all this section.

Definition 9. Let G be a group and I a set. We says that a family of element $(e_i)_{i \in I}$ of G is a basis of G is every element of G can be written uniquely as a finite linear combination with integer coefficients of elements of this family. If I is finite, this cardinal is called the rank of G .

Notations For two integers k and n , the predicate $k|n$ denotes that k divides n . Also, let n be an integer and p a prime numbers. We define the order of n at p , denoted by $\text{ord}_p(n)$, as the positive integer α such that $p^\alpha | n$ and $p^{\alpha+1}$ does not divide n . It is the exponent of p in the prime decomposition of n if $p|n$, and 0 otherwise.

Fact 10 There exists a basis $(e_i)_{1 \leq i \leq \phi(m)}$ for R^\vee and $\phi(m)$ positive integer $(b_i)_{1 \leq i \leq \phi(m)}$ such that $(b_i e_i)_{1 \leq i \leq \phi(m)}$ is a basis for R . And moreover, $\forall i \leq \phi(m) - 1, b_i | b_{i+1}$.

Proof: All elements of R can be uniquely written as a linear combination with integer coefficients of $(\zeta_m^i)_{0 \leq i \leq \phi(m)-1}$, as this family is rationally independent. In a similar way, we have that all elements of R^\vee can be uniquely written as a linear combination with integer coefficients of $(\frac{\zeta_m^i}{\Phi_n'(\zeta_m)})_{0 \leq i \leq \phi(m)-1}$. Since $R \subset$

R^\vee , we can write for all i , ζ_m^i in the latter family : $\forall i \in [0, \phi(m) - 1]$, $\zeta_m^i = \sum_{j=0}^{\phi(m)-1} a_{i,j} \frac{\zeta_m^j}{\phi_m(\zeta_m)}$. We end with a square matrix $A = (a_{i,j})_{0 \leq i, j \leq \phi(m)-1}$ of dimension $\phi(m)$ with integer coefficients. And we consider its Smith normal form (Proposition 2.1.5 in [Ste05]): namely, there exists two matrix U and V with integer coefficients of dimension $\phi(m)$, invertible as integer matrices, such that $UAV = D$, where D is a diagonal matrix with positive integer coefficients on the form :

$$\begin{pmatrix} b_1 & 0 & \cdots & 0 \\ 0 & \ddots & & \\ & & b_r & \\ \vdots & & & 0 \\ 0 & \cdots & & \ddots \\ 0 & \cdots & & & 0 \end{pmatrix}$$

And, we have $b_i | b_{i+1}$ for any $i < r$. Besides, in our case, $r = \phi(m)$ and $\forall i < r, b_i \neq 0$. Indeed, let's notice that A is a change-of-basis matrix for two \mathbb{Q} -basis of $\mathbb{Q}(\zeta_m)$, then invertible, and then its determinant is non-zero. But $\det(D) = \det(UAV) = \det(U) \det(A) \det(V) = \det(A)$, because U and V are invertible as integer matrices, and thus have their determinant equal to 1.

This decomposition of A gives us a basis $(e_i)_{1 \leq i \leq \phi(m)}$ for R^\vee and $\phi(m)$ integer $(b_i)_{1 \leq i \leq \phi(m)}$ such that $(b_i e_i)_{1 \leq i \leq \phi(m)}$ is a basis for R , where $b_i | b_{i+1}$ for any $i \leq \phi(m) - 1$. \square

Thanks to this result, we can state two immediate consequences.

Fact 11 *With the notation of the previous Fact 10, an integer k is such that $kR^\vee \subset R$ if and only if $b_{\phi(m)} | k$. Therefore, $b_{\phi(m)} | m'$.*

Moreover we have the following equality:

$$\prod_{i=1}^{\phi(m)} b_i = m^{\phi(m)} \prod_{\substack{p/m \\ p \text{ prime number}}} p^{\phi(m)/(p-1)}$$

Proof: First, an integer k is such that $kR^\vee \subset R$ if and only if

$$\forall i \leq \phi(m) \quad ke_i \in R \tag{1}$$

or equivalently, if and only if for all i , there exist $c_i \in \mathbb{N}$ such that $ke_i = a_i b_i e_i$. This can be rewritten as

$$\forall i, \quad b_i | k.$$

By Fact 10, all the b_i 's divides $b_{\phi(m)}$, so condition (1) is equivalent to: $b_{\phi(m)} | k$.

The Fact 10 gives us the cardinality of the finite group R^\vee/R , called the index of R in R^\vee and denoted $[R^\vee : R] : [R^\vee : R] = \prod_{i=1}^{\phi(m)} b_i$.

Yet, this index is known and in fact equals to the absolute value of the discriminant of the cyclotomic field (Theorem 4.6 in [Con09]), whom we know an exact expression (Proposition 2.7 of [Was97]):

$$\text{disc}(\mathbb{Q}(\zeta_m)) = (-1)^{\frac{\phi(m)}{2}} m^{\phi(m)} \prod_{\substack{p/m \\ p \text{ prime number}}} p^{\phi(m)/(p-1)}$$

□

Using previous facts, we may now prove our main results:

Theorem 12. *With the notation above, $m' = b_{\phi(m)}$ and a integer k is such that $kR^\vee \subset R$ if and only if $m'|k$.*

Proof: First, we prove that $m'|b_{\phi(m)}$. To prove this, we work on the prime factors of m' . More precisely, we show that for all prime $p|m'$, $\text{ord}_p(m') \leq \text{ord}_p(b_{\phi(m)})$, it immediately follows that $m'|b_{\phi(m)}$. Let p a prime factor of m' different from 2. Then by definition of m' , $\text{ord}_p(m') = \text{ord}_p(m)$.

We proceed by assuming that $\text{ord}_p(m) > \text{ord}_p(b_{\phi(m)})$, and show that it is absurd. From Fact 10 we have $b_i|b_{i+1}$ for all $i < \phi(m)$. Thus: $\text{ord}_p(m) - 1 \geq \text{ord}_p(b_i)$ and summing over all i we get:

$$(\text{ord}_p(m) - 1)\phi(m) \geq \sum_{i=1}^{\phi(m)} \text{ord}_p(b_i). \quad (2)$$

Fact 11 tells us that:

$$\prod_{i=1}^{\phi(m)} b_i = m^{\phi(m)} \prod_{\substack{p/m \\ p \text{ prime number}}} p^{\phi(m)/(p-1)}$$

and therefore

$$\begin{aligned} \text{ord}_p \left(\prod_{i=1}^{\phi(m)} b_i \right) &= \text{ord}_p(m)\phi(m) - \frac{\phi(m)}{p-1} \\ \sum_{i=1}^{\phi(m)} \text{ord}_p(b_i) &= \phi(m) \cdot \left(\text{ord}_p(m) - \frac{1}{p-1} \right) \end{aligned}$$

Combining with the inequality (2) we deduce:

$$(\text{ord}_p(m) - 1)\phi(m) \geq \phi(m) \left(\text{ord}_p(m) - \frac{1}{p-1} \right)$$

which is absurd since $p > 2$. We actually get that $\text{ord}_p(m') \leq \text{ord}_p(b_{\phi(m)})$, if p is prime factor of m' different from 2.

If 2 is a prime factor of m' , the same reasoning is similar, starting with $\text{ord}_2(m') = \text{ord}_2(m) - 1$.

Therefore $m' | b_{\phi(m)}$. And the Theorem 3 with the Fact 11 tell us that $b_{\phi(m)} | m'$. Thus $m' = b_{\phi(m)}$. And by the Fact 11 again, a integer k is such that $kR^\vee \subset R$ if and only if $m' | k$. \square