

On error distributions in ring-based LWE

Wouter Castryck, Iliia Iliashenko and Frederik Vercauteren

ABSTRACT

Since its introduction in 2010 by Lyubashevsky, Peikert and Regev, the ring learning with errors problem (ring-LWE) has become a popular building block for cryptographic primitives, due to its great versatility and its hardness proof consisting of a (quantum) reduction from ideal lattice problems. But, for a given modulus q and degree n number field K , generating ring-LWE samples can be perceived as cumbersome, because the secret keys have to be taken from the reduction mod q of a certain fractional ideal $\mathcal{O}_K^\vee \subset K$ called the codifferent or ‘dual’, rather than from the ring of integers \mathcal{O}_K itself. This has led to various non-dual variants of ring-LWE, in which one compensates for the non-duality by scaling up the errors. We give a comparison of these versions, and revisit some unfortunate choices that have been made in the recent literature, one of which is scaling up by $|\Delta_K|^{1/2n}$ with Δ_K the discriminant of K . As a main result, we provide, for any $\varepsilon > 0$, a family of number fields K for which this variant of ring-LWE can be broken easily as soon as the errors are scaled up by $|\Delta_K|^{(1-\varepsilon)/n}$.

1. Introduction: ring-based versions of LWE

About a decade ago, Regev [19] proposed a new hard problem for use in public-key cryptography, namely, the learning with errors problem (LWE), which, informally stated, is about solving an approximate linear system

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_m \end{pmatrix} = A \cdot \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \\ \vdots \\ \epsilon_m \end{pmatrix}$$

for an unknown secret $\mathbf{s} = (s_1, s_2, \dots, s_n)$ over $\mathbb{Z}/q\mathbb{Z}$, with q some integer modulus. The entries of A are selected independently and uniformly at random in $\mathbb{Z}/q\mathbb{Z}$ and the ϵ_i are small error terms, obtained by sampling from a fixed Gaussian with mean 0 and standard deviation $\rho \geq \sqrt{2n/\pi}$, and reducing the outcome mod q . These errors are elements of $\mathbb{R}/q\mathbb{Z}$, but, in practice, they can be rounded to the nearest element of $\mathbb{Z}/q\mathbb{Z}$. To recover \mathbf{s} uniquely, the system has to be overdetermined: that is $m > n$. In fact, in Regev’s model, an attacker is allowed to ask for new equations indefinitely, in the (conjecturally vain) hope of learning more information about \mathbf{s} : hence the terminology learning with errors.

The LWE problem is being acclaimed for three reasons. Firstly, it enjoys a hardness proof in the form of a reduction from worst-case forms of certain well-established lattice problems [2, 17, 19], providing security guarantees that are lacking for classical hard problems such as

Received 21 February 2016.

2010 Mathematics Subject Classification 11T71 (primary), 11R04, 11R11, 11R18, 11T22 (secondary).

Contributed to the Twelfth Algorithmic Number Theory Symposium (ANTS-XII), Kaiserslautern, Germany, 29 August–2 September 2016.

This work was supported by the European Commission through the ICT programme under contracts H2020-ICT-2014-1 644209 HEAT and H2020-ICT-2014-1 645622 PQCRYPTO. The authors would like to thank the anonymous reviewers for their highly valuable comments, which led to several improvements.

integer factorization or discrete logarithm computation. Secondly, it seems that LWE would remain hard in a post-quantum world, unlike the classical problems [20]. Thirdly, LWE has proved to be very versatile for use in cryptography, enabling applications that were not known before, such as homomorphic encryption [1, 3]. Its major drawback, however, is that the key sizes of the resulting cryptosystems are impractically large because, typically, one needs the entire $(m \times n)$ -matrix A .

One idea to address this [14, 16] is to use a ring structure $R_q = \mathbb{Z}[x]/(q, f(x))$ for some monic degree n polynomial $f(x) \in \mathbb{Z}[x]$ through the isomorphism (a priori of modules)

$$\varphi : \left(\frac{\mathbb{Z}}{q\mathbb{Z}} \right)^n \rightarrow R_q : (s_1, s_2, \dots, s_n) \mapsto s_1 + s_2x + \dots + s_nx^{n-1}.$$

Each block of n rows of A is replaced by the matrix $A_{\mathbf{a}}$ of multiplication by a random ring element $\mathbf{a}(x)$, say, with respect to the polynomial basis $1, x, \dots, x^{n-1}$, in order to obtain an approximate linear system of the form

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \\ \vdots \\ \epsilon_n \end{pmatrix}, \tag{1.1}$$

which, using φ , can be rewritten as $\mathbf{b}(x) = \mathbf{a}(x) \cdot \mathbf{s}(x) + \boldsymbol{\epsilon}(x)$. When storing $\mathbf{a}(x)$ rather than $A_{\mathbf{a}}$, one gains a factor n , thereby addressing the key size issue. The general set-up above is called *ring-based LWE* (not to be confused with ring-LWE) and the terminology allows for any error distribution Ψ on R_q for the error terms $\boldsymbol{\epsilon}(x)$. In the remainder of the article, we will consider three variants, all of which sample the error terms from a linear transform of an n -dimensional spherical Gaussian. More precisely, in each case there exists a fixed matrix $T \in \mathbb{R}^{n \times n}$ such that $(\epsilon_1 \epsilon_2 \dots \epsilon_n)^t$ arises as the reduction modulo $q\mathbb{Z}^n$ of $T \cdot (e_1 e_2 \dots e_n)^t$, where each e_i is sampled independently from the same normal distribution $\mathcal{N}(0, \rho^2)$ with mean 0 and standard deviation ρ , which we think of as depending on n only, in a non-negligible and polynomially bounded way. Again, in practice, one can round ϵ_i to the nearest element of $\mathbb{Z}/q\mathbb{Z}$, but, for analytical reasons, it is convenient not to do this. The different choices for T are summarized in Table 1 and how these T arise is explained in detail in the next few paragraphs.

Poly-LWE can be considered the most straightforward generalization of LWE, in that each error ϵ_i is drawn independently from $\mathcal{N}(0, \rho^2)$. In particular, the matrix T is simply the identity matrix. For the sake of analogy, one could again impose $\rho \geq \sqrt{2n/\pi}$, although there is no theoretical basis for this. Indeed, restricting to multiplication matrices comes at the cost of giving up on the uniform randomness, thereby invalidating Regev’s hardness proof and, in fact, it is possible to cook up instances of the problem having certain flaws. For example, if $f(1) \equiv 0 \pmod q$, then $\mathbf{b}(1) \equiv \mathbf{a}(1) \cdot \mathbf{s}(1) + \mathbf{e}(1) \pmod q$, which can, in certain special cases, be exploited to obtain information about the secret [10], thereby mimicking an attack on early versions of NTRU that use arithmetic modulo $f(x) = x^n - 1$ (see [13]). This concern is partly addressed by restricting to irreducible $f(x) \in \mathbb{Z}[x]$, which we do from now on.

TABLE 1. Noise distributions $T \cdot \mathcal{N}(0, \rho^2)^n$ in three instantiations of ring-based LWE, with B the (real) canonical embedding matrix, $A_{f'(x)}$ the matrix of multiplication by $f'(x)$ and $\lambda \in \mathbb{R}_{\geq 1}$.

| | Poly-LWE | Ring-LWE | SCG ring-based LWE |
|-----|------------------|--------------------------|------------------------|
| T | $I_{n \times n}$ | $A_{f'(x)} \cdot B^{-1}$ | $\lambda \cdot B^{-1}$ |

Ring-LWE was introduced by Lyubashevsky, Peikert and Regev in [16] and admits a hardness proof akin to the one for general LWE. The main difference is that the error terms are generated in a way that is canonical for the underlying number field K , defined by $f(x)$, and, in particular, does not depend on the choice of the defining polynomial $f(x)$ itself (unlike poly-LWE). For the purpose of this introduction, it suffices to think of ring-LWE samples as above, except that the spherical error vector $(e_1 \ e_2 \ \dots \ e_n)$ is being transformed in the following specific way

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + A_{f'(x)} \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}. \quad (1.2)$$

Here $B \in \mathbb{R}^{n \times n}$ is the Vandermonde matrix $(\alpha_i^{j-1})_{i,j}$ generated by the roots $\alpha_1, \alpha_2, \dots, \alpha_n \in \mathbb{C}$ of $f(x)$, turned into a real matrix using an easy unitary transformation. The factor B^{-1} expresses the fact that ring-LWE errors are actually generated in the codomain of the canonical embedding of the number field $K = \mathbb{Q}[x]/(f(x))$. On the other hand, $A_{f'(x)}$ is the matrix of multiplication by the derivative $f'(x)$ of our defining polynomial. It compensates for the fact that we sampled the secret $\mathbf{s}(x)$ from the reduction mod q of $R = \mathbb{Z}[x]/(f(x))$, rather than from the reduction mod q of a certain fractional ideal $R^\vee \subset K$, called the *dual* of R . It is convenient to think of $A_{f'(x)}$ as an *integral* matrix, that is, as the matrix of multiplication by $f'(x)$ in R with respect to the \mathbb{Z} -basis $1, x, \dots, x^{n-1}$, so that ring-LWE is just ring-based LWE with $T = A_{f'(x)} \cdot B^{-1}$.

The matrix $A_{f'(x)} \cdot B^{-1}$ transforms our spherical distribution into an ellipsoidal one. In particular, the errors in certain coordinates might be systematically much larger than those in others, and they may no longer be independent. But it is crucial to observe that the error coordinates are being *scaled up* on average, in the sense of the geometric mean. Indeed, one can show that $|\det A_{f'(x)}| = \Delta$ and that $|\det B| = \sqrt{\Delta}$, where Δ denotes the absolute value of the discriminant of $f(x)$. Thus

$$|\det(A_{f'(x)} \cdot B^{-1})| = \sqrt{\Delta},$$

meaning that, on average, the errors tend to grow by a factor $\Delta^{1/2n}$.

SCG ring-based LWE where SCG stands for scaled canonical Gaussian, was analyzed in a series of papers [5, 6, 11] in which it was called non-dual ring-LWE. In this version, one considers samples of the form

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + \lambda \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}, \quad (1.3)$$

where $\lambda \geq 1$ denotes a fixed real number. This variant basically replaces the matrix $A_{f'(x)}$ in ring-LWE by a scalar matrix. The authors called this variant non-dual ring-LWE since the matrix $A_{f'(x)}$ corresponds to the factor coming from working with the dual. However, we will avoid using this terminology, so as not to get confused with non-dual instantiations of actual ring-LWE, as in (1.2).

Note that one cannot simply remove $A_{f'(x)}$ (that is, take $\lambda = 1$), since the remaining factor B^{-1} has determinant $1/\sqrt{\Delta}$, which typically scales down the errors to a point where they become negligible, leading to *exact* equations in s_1, s_2, \dots, s_n that can be solved using linear algebra. This is, of course, highly undesirable and, to remedy this, the authors of [5, 6, 11] used $\lambda = \Delta^{1/2n}$ in order to undo the factor B^{-1} determinant wise.




FIGURE 1. Coordinate-wise error distributions for $f(x) = x^{256} + 8190$, $\rho = 8.35$, and $\lambda = \sqrt{\Delta}^{1/256}$.

This choice of scalar does indeed take back the errors to a reasonable size, but only on average. If the ellipsoidal distribution induced by B^{-1} is extremely skew, then there might be error coordinates that remain negligibly small after scaling. The following example, introduced in [11] and revisited in [4], illustrates this. For $f(x) = x^{256} + 8190$, the successive radii of the corresponding 256-dimensional ellipsoid go down geometrically, as illustrated in Figure 1. It turns out that with

$$\rho = 8.35 \quad \text{and} \quad \lambda = \sqrt{\Delta}^{1/256} \approx 1422.72,$$

the coordinates at the highest 45 indices become zero after rounding, with overwhelming probability. Thus each sample yields 45 exact equations in s_1, s_2, \dots, s_n , and about six samples suffice to recover the entire secret. If one would take $\lambda = 1$, the effect is even more pronounced since then over 240 errors are negligible and one only requires two samples. In general, for this attack to work it is enough that B^{-1} admits a very short \mathbb{Z} -linear combination of its rows. See [18] for a more thorough analysis of all this.

To us it seems more natural to take $\lambda = \Delta^{1/n}$: in this way one compensates determinant-wise for the removal of $A_{f'(x)}$. For this choice of scalar, we are unaware of any attacks on SCG ring-based LWE and it would be interesting to know whether a variant of the hardness proof of [16] applies here.

The main result of this article is that $\Delta^{1/n}$ is a lower bound for λ , in the following sense. For each $\varepsilon > 0$, we provide a family of irreducible polynomials $f(x) \in \mathbb{Z}[x]$ of increasing degree n , for which $O(n)$ SCG samples of the form (1.3) with $\lambda = \Delta^{(1-\varepsilon)/n}$ are sufficient to recover the entire secret using standard linear algebra.

In fact, as we will see, the analogous result also applies to ring-LWE, for the same families of polynomials. In other words, as soon as one scales down the right-most term in (1.2) by $\Delta^{\varepsilon/n}$, then the corresponding samples leak exact equations, again allowing one to find the entire secret easily. However, as suggested by a reviewer, in this case the statement admits an easier proof, based on the trivial fact that $1 \in R$.

The article is organized as follows. In § 2, we give a more formal introduction to ring-LWE, while § 3 is devoted to the SCG ring-based LWE version that was studied in [5, 6, 11]. Apart from providing more details, these descriptions will differ slightly from the one given in the introduction: instead of $\mathbb{Z}[x]/(f(x))$, we will work in the potentially larger ring of integers \mathcal{O}_K of K . Then, in § 4, we state a rigorous version of our tightness result on the scaling factor λ and provide a proof. Finally, in § 5, we make some additional comments from the point of view of Galois theory.

2. Ring-LWE set up formally

The actual ring-LWE problem is formulated using the ring of integers $R = \mathcal{O}_K$ of a given degree n number field K , which one considers along with a modulus $q \in \mathbb{Z}$. A central role is played by the codifferent R^\vee of K , which is defined as the inverse (fractional) ideal of the different ideal $\mathfrak{d} \subset R$. Alternatively, it can be viewed as the dual of R with respect to the trace pairing

$$R^\vee = \{x \in K \mid \text{Tr}_{K/\mathbb{Q}}(xR) \subset \mathbb{Z}\}. \tag{2.1}$$

The reductions of R and R^\vee modulo qR and qR^\vee , respectively, are denoted by R_q and R_q^\vee , respectively. The ring-LWE problem is then about finding a fixed secret $\mathbf{s} \in R_q^\vee$ from an arbitrary number of approximate equations of the form

$$\mathbf{b} = \mathbf{a} \cdot \mathbf{s} + \boldsymbol{\epsilon}, \tag{2.2}$$

where $\mathbf{a} \in R_q$ is chosen uniformly at random and $\boldsymbol{\epsilon}$ is a small error term sampled from a distribution that will be described in the next paragraph. Recall that everything is to be interpreted modulo qR^\vee . After agreeing upon a \mathbb{Z} -basis of R^\vee , this can be rewritten as

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \epsilon_1 \\ \vdots \\ \epsilon_n \end{pmatrix},$$

where the s_i are the coordinates of \mathbf{s} , the b_i are the coordinates of \mathbf{b} , the ϵ_i are the coordinates of $\boldsymbol{\epsilon}$ and $A_{\mathbf{a}}$ is the matrix of multiplication by \mathbf{a} with respect to the chosen \mathbb{Z} -basis, all considered modulo $q\mathbb{Z}^n$.

As for the error distribution, the main role is played by the canonical embedding

$$\sigma : K \rightarrow \mathbb{C}^n : \alpha \mapsto (\sigma_1(\alpha), \dots, \sigma_n(\alpha)),$$

where $\sigma_1, \dots, \sigma_s$ are the real ring monomorphisms from K to \mathbb{R} and $\sigma_{s+1}, \dots, \sigma_{s+2t}$ are the complex ring monomorphisms from K to \mathbb{C} (so that $n = s + 2t$), ordered such that $\sigma_{s+i} = \tau \circ \sigma_{s+t+i}$ for $i = 1, \dots, t$, where $\tau : \mathbb{C} \rightarrow \mathbb{C} : z \mapsto \bar{z}$ denotes complex conjugation. Thus σ takes values in

$$H = \{(z_1, \dots, z_n) \in \mathbb{C}^n \mid z_1, \dots, z_s \in \mathbb{R} \text{ and } \bar{z}_{s+i} = z_{s+t+i} \text{ for } i = 1, \dots, t\},$$

which, when equipped with the Hermitian inner product $\langle \cdot, \cdot \rangle$, is seen to be isomorphic to the standard inner product space \mathbb{R}^n by considering the basis given by the columns of the unitary matrix

$$U = \begin{pmatrix} I_{s \times s} & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} I_{t \times t} & \frac{\mathbf{i}}{\sqrt{2}} I_{t \times t} \\ 0 & \frac{1}{\sqrt{2}} I_{t \times t} & -\frac{\mathbf{i}}{\sqrt{2}} I_{t \times t} \end{pmatrix}.$$

It is well known that, under this identification of H with \mathbb{R}^n , the image $\sigma(I)$ of a fractional ideal $I \subset K$ is a lattice of rank n and that $\sigma(R^\vee)$ is the complex conjugate of the dual lattice

$$\sigma(R)^* := \{\alpha \in H \mid \langle \alpha, \sigma(R) \rangle \subset \mathbb{Z}\},$$

as is immediate from (2.1); more generally $\sigma(I)^* = \tau(\sigma(I^\vee))$, where $I^\vee = (\partial I)^{-1}$. Now consider a spherical Gaussian on \mathbb{R}^n , say, with distribution function

$$\Gamma_r^n(\mathbf{x}) = \frac{1}{r^n} \exp\left(-\pi \frac{\|\mathbf{x}\|^2}{r^2}\right),$$

where we note that $\Gamma_r^1 = \mathcal{N}(0, r^2/2\pi)$ and that

$$\Gamma_r^n = \Gamma_r^1 \times \Gamma_r^1 \times \dots \times \Gamma_r^1.$$

We view Γ_r^n as a distribution on H through the above identification with \mathbb{R}^n . Pulling it back along the canonical embedding and reducing it mod qR^\vee results in a distribution Ψ_r on the torus

$$(R^\vee \otimes_{\mathbb{Z}} \mathbb{R})/qR^\vee,$$

from which the errors are to be sampled.

We can now formulate the ring-LWE problem precisely. Let $\mathfrak{U}(R_q)$ and $\mathfrak{U}(R_q^\vee)$ denote the uniform distributions on R_q and R_q^\vee , respectively. For $\mathbf{s} \in R_q^\vee$ and $r \in \mathbb{R}_{>0}$, we let $A_{\mathbf{s},r}$ be the distribution over

$$R_q \times (R_q^\vee \otimes_{\mathbb{Z}} \mathbb{R})/qR^\vee$$

obtained by sampling $\mathbf{a} \leftarrow \mathfrak{U}(R_q)$, $\epsilon \leftarrow \Psi_r$ and returning $(\mathbf{a}, \mathbf{a} \cdot \mathbf{s} + \epsilon)$.

DEFINITION 1 (Ring-LWE over a number field K with error parameter r). For a random but fixed choice of $\mathbf{s} \leftarrow \mathfrak{U}(R_q^\vee)$, the (search) *ring-LWE* problem is to recover \mathbf{s} with non-negligible probability from arbitrarily many independent samples from $A_{\mathbf{s},r}$.

Here it is understood that $r \geq 2\omega(\sqrt{\log n})$ for some superlinear function $\omega = \omega(n)$. It may seem surprising that this bound is less restrictive than in standard LWE, where one assumes that $r = \sqrt{2\pi}\rho \geq 2\sqrt{n}$. But this is only superficial: the lattice of possible products $\mathbf{a} \cdot \mathbf{s}$ is much denser because \mathbf{s} was sampled from R_q^\vee and, relative to this, the ring-LWE bound is considerably larger.

In their seminal paper [16], Lyubashevsky, Peikert and Regev provided the following hardness result. They actually deal with a slight variant called the ring-LWE $_{\leq r}$ problem, where each sample is taken from $A_{\mathbf{s},\mathbf{r}}$ for some arbitrary fixed \mathbf{r} taken from

$$\{(r_1, \dots, r_n) \in (\mathbb{R}^+)^n \mid r_i \leq r \text{ for all } i = 1, \dots, s \text{ and } r_{s+i} = r_{s+t+i} \leq r \text{ for all } i = 1, \dots, t\}.$$

The distribution $A_{\mathbf{s},\mathbf{r}}$ is defined in roughly the same way as $A_{\mathbf{s},r}$, the main difference being that the spherical Gaussian Γ_r^n is to be replaced by the ellipsoidal Gaussian $\Gamma_{r_1}^1 \times \Gamma_{r_2}^1 \times \dots \times \Gamma_{r_n}^1$. If we think of the error width $r \geq 2\omega(\sqrt{\log n})$ and the modulus $q \geq 2$ as quantities that vary with n , then the hardness result [16, Theorem 4.1] gives the following theorem.

THEOREM 2.1. *For some negligible $\varepsilon = \varepsilon(n)$ there is a probabilistic polynomial-time quantum reduction from DGS_γ to ring-LWE $_{\leq r}$, where*

$$\gamma : I \mapsto \max\{\eta_\varepsilon(I) \cdot (\sqrt{2}q/r) \cdot \omega(\sqrt{\log n}), \sqrt{2n}/\lambda_1(I^\vee)\}.$$

Here $\eta_\varepsilon(I)$ is the smoothing parameter of $\sigma(I)$ with threshold ε and $\lambda_1(I^\vee)$ is the length of a shortest vector of $\sigma(I^\vee)$.

The statement involves the discrete Gaussian sampling problem DGS_γ , which is about producing samples from a spherical Gaussian in H with parameter r' , restricted to the lattice $\sigma(I)$, for any given non-zero ideal $I \subset R$ and any $r' \geq \gamma(I)$. As discussed in [16], there are easy reductions from standard lattice problems to the discrete Gaussian sampling problem.

3. SCG ring-based LWE

To allow for a common framework for poly-LWE and ring-LWE, from now on we restrict ourselves to number fields K for which the different ideal \mathfrak{d} is principal, say, generated by

$\theta \in R$, so that $R^\vee = R/\theta$. This restriction is mainly for convenience: in general, one can replace θ in the discussion below by a so-called *tweaking* factor (see [7, 18]). But principality holds in most cases of interest. For instance, if K is monogenic, meaning that the ring of integers R is of the form $\mathbb{Z}[x]/(f(x))$, then one can take $\theta = f'(x)$. More generally, ∂ is principal if and only if R is a so-called complete intersection, that is, of the form $\mathbb{Z}[x_1, x_2, \dots, x_n]/(f_1, f_2, \dots, f_n)$, in which case one can take $\theta = |(\partial f_i / \partial x_j)_{i,j}|$ (see [9]).

Without loss of generality, we can rewrite our sample (2.2) as

$$\mathbf{a} \cdot \frac{\mathbf{s}}{\theta} = \frac{\mathbf{b}}{\theta} + \epsilon,$$

where now $\mathbf{s} \in R_q$ and ϵ is sampled from Ψ_r . Multiplying by θ then gives $\mathbf{b} = \mathbf{a} \cdot \mathbf{s} + \theta \cdot \epsilon$. After fixing a \mathbb{Z} -basis $\alpha_1, \alpha_2, \dots, \alpha_n$ of R , we obtain

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + A_\theta \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}, \tag{3.1}$$

where the s_i are the coordinates of \mathbf{s} , the b_i are the coordinates of \mathbf{b} , $A_{\mathbf{a}}$ is the matrix of multiplication by \mathbf{a} , A_θ is the matrix of multiplication by θ and $B = U \cdot \Sigma$ with Σ the matrix of the canonical embedding σ , all expressed with respect to the basis $\alpha_1, \alpha_2, \dots, \alpha_n$. Note that Σ is just the complex matrix having $\sigma(\alpha_1), \sigma(\alpha_2), \dots, \sigma(\alpha_n)$ as its columns. The e_i are sampled independently from the univariate Gaussian Γ_r^1 . The formula (3.1) is to be considered modulo q , but note that, in the case of the subexpression $B^{-1} \cdot (e_1 e_2 \dots e_n)^t$, it only makes sense to do so *after* elaborating the product. On average, the factor $A_\theta \cdot B^{-1}$ causes the errors to expand, because $|\det A_\theta| = \Delta$ and $|\det B| = |\det \Sigma| = \text{covol}(\sigma(R)) = \sqrt{|\Delta|}$, where $\Delta = |\Delta_K|$ is the absolute value of the discriminant of K (see [12]).

REMARK 1. In the monogenic case, we can take $\theta = f'(x)$ and work with respect to the basis $1, x, \dots, x^{n-1}$. For these choices, we exactly recover (1.2) and we enter the discussion from the introduction. Note that $\Delta = |\text{disc } f(x)|$, in this case.

Taking another generator θ of ∂ boils down to replacing the right-most term in (3.1), that is, the vector of coordinates of the error term $\theta \cdot \epsilon$, by

$$M \cdot A_\theta \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}$$

for some matrix $M \in \text{GL}_n(\mathbb{Z})$. The same remark applies to switching to another basis of R , in which case M arises as the corresponding matrix of base change. In particular, if for one choice of basis a certain error coordinate is negligible, then for another choice of basis a certain non-trivial \mathbb{Z} -linear combination of the error coordinates will be negligible; the converse is also true.

EXAMPLE 1. Let $\beta_1, \beta_2, \dots, \beta_n$ be the basis of R^\vee that is dual to our given basis $\alpha_1, \alpha_2, \dots, \alpha_n$ of R with respect to the trace pairing. In other words, $\sigma(\beta_1), \sigma(\beta_2), \dots, \sigma(\beta_n) \in \mathbb{C}^n$ are the columns of the conjugate transpose $\tau(\Sigma)^t$ of Σ . But then $\theta \cdot \beta_1, \theta \cdot \beta_2, \dots, \theta \cdot \beta_n$ is also a basis of R , so we can change bases. In this case, one verifies that the matrix of base

change M is $B^t \cdot B \cdot A_\theta^{-1}$, with respect to which our ring-LWE samples become

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + B^t \cdot \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}. \tag{3.2}$$

If we express the ring-LWE samples directly in terms of the basis $\beta_1, \beta_2, \dots, \beta_n$ of R^\vee , then we would find the same formula. Thus, in some sense, (3.2) is more in the actual spirit of [16] than (3.1), but it is less suitable for discussing the SCG ring-based LWE version from [5, 6, 11]. \square

Recall, from the introduction, that the SCG ring-based LWE from [5, 6, 11] leaves out the multiplication by θ step, and compensates for it by a scalar. Formally, one considers samples of the form

$$\mathbf{b} = \mathbf{a} \cdot \mathbf{s} + \lambda \cdot \boldsymbol{\epsilon},$$

where \mathbf{s} is now taken from R_q rather than R_q^\vee . As before, let $\mathbf{a} \leftarrow \mathcal{U}(R_q)$ and $\boldsymbol{\epsilon} \leftarrow \Psi_r$, and let $\lambda \geq 1$ be a fixed real scalar. Let $A_{\mathbf{s},r}^\lambda$ be the resulting distribution over

$$R_q \times (R_q \otimes_{\mathbb{Z}} \mathbb{R})/qR$$

returning $(\mathbf{a}, \mathbf{a} \cdot \mathbf{s} + \lambda \cdot \boldsymbol{\epsilon})$.

DEFINITION 2 (SCG ring-based LWE with scalar λ). For a random but fixed choice of $\mathbf{s} \leftarrow \mathcal{U}(R_q)$, the problem is to recover \mathbf{s} with non-negligible probability from an arbitrary number of independent samples from $A_{\mathbf{s},r}^\lambda$.

When expressed with respect to a basis $\alpha_1, \alpha_2, \dots, \alpha_n$ of R , such a sample converts into an expression of the form

$$\begin{pmatrix} b_1 \\ b_2 \\ \vdots \\ b_n \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \\ \vdots \\ s_n \end{pmatrix} + \lambda \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{pmatrix}. \tag{3.3}$$

Equivalently, one can also just remove the scalar λ and sample the errors e_i from $\Gamma_{\lambda \cdot r}$ instead of Γ_r . Here too, switching to another basis amounts to multiplying the right-most factor from the left with a matrix $M \in \text{GL}_n(\mathbb{Z})$.

As mentioned in the introduction, the authors of [5, 6, 11] took $\lambda = \Delta^{1/2n}$, while to us the most natural choice of scalar seems $\lambda = |\Delta|^{1/n}$, in order to compensate determinant-wise for the removal of A_θ . It would be interesting to know whether the latter choice allows for a hardness statement similar to Theorem 2.1. If A_θ happens to be a scalar matrix itself, then both problems are, of course, equivalent. For instance, this is the case if K is the 2^m th cyclotomic field for some $m \geq 2$, where one can take $\lambda = \theta = 2^{m-1} = n$.

EXAMPLE 2. To illustrate these different flavors of ring-based LWE, we analyze a simple example that will act as one of the building blocks in our main theorem. Let $d \equiv 1 \pmod 4$ be a positive squarefree integer and consider the real quadratic field $K = \mathbb{Q}(\sqrt{d})$. It has discriminant d and its ring of integers $R = \mathbb{Z}[(1+\sqrt{d})/2]$ admits the integral basis $1, (1+\sqrt{d})/2$. The different ideal \mathfrak{d} is the principal ideal generated by $\theta = \sqrt{d}$. With respect to this basis

$$A_\theta = \begin{pmatrix} -1 & \frac{-1+d}{2} \\ 2 & 1 \end{pmatrix}, \quad \Sigma^{-1} = \frac{1}{\sqrt{d}} \begin{pmatrix} \frac{-1+\sqrt{d}}{2} & \frac{1+\sqrt{d}}{2} \\ 1 & -1 \end{pmatrix}, \quad U = I_{2 \times 2}.$$

So a ring-LWE sample reads

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} + A_{\theta} \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} + \begin{pmatrix} \frac{-1 + \sqrt{d}}{2} & \frac{-1 - \sqrt{d}}{2} \\ 1 & 1 \end{pmatrix} \cdot \begin{pmatrix} e_1 \\ e_2 \end{pmatrix},$$

while a SCG ring-based LWE sample reads

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} + \sqrt{d} \cdot B^{-1} \cdot \begin{pmatrix} e_1 \\ e_2 \end{pmatrix} = A_{\mathbf{a}} \cdot \begin{pmatrix} s_1 \\ s_2 \end{pmatrix} + \begin{pmatrix} \frac{-1 + \sqrt{d}}{2} & \frac{1 + \sqrt{d}}{2} \\ 1 & -1 \end{pmatrix} \cdot \begin{pmatrix} e_1 \\ e_2 \end{pmatrix},$$

for scaling factor $\lambda = |\Delta|^{1/n} = \sqrt{d}$. □

For the sake of completeness, let us conclude with the setting where one considers (3.1) with the entire matrix product $A_{\theta} \cdot B^{-1}$ replaced by a real scalar $\lambda \geq 1$. In the monogenic case $R = \mathbb{Z}[x]/(f(x))$, where one takes $\lambda = 1$ and works with respect to the basis $1, x, \dots, x^{n-1}$, one recovers the poly-LWE problem from the introduction. Note that in order to compensate for the removal of $A_{\theta} \cdot B^{-1}$ determinant-wise, it is more natural to take $\lambda = \Delta^{1/2n}$ (here too it would be interesting to know whether the resulting problem enjoys a hardness proof). The more aggressive choice for $\lambda = 1$ may be motivated by the error bound in Regev’s original work on LWE [19], where there is no number field at play, and by NTRU, where the errors are bounded by a small constant. Taking smaller errors has advantages towards the efficiency of the resulting cryptosystems, but the security risks of doing so are not fully understood.

4. Main theorem

THEOREM 4.1. *Let $\rho : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be in $\text{poly}(n)$, let $(q_n)_{n \in \mathbb{N}}$ be any sequence of integer moduli and let $\varepsilon \in \mathbb{R}_{>0}$ be fixed. Then there exists a family of number fields $(K_{\ell})_{\ell \in \mathbb{N}}$ such that:*

- each K_{ℓ} is Galois over \mathbb{Q} ;
- the degree $n_{\ell} := [K_{\ell} : \mathbb{Q}]$ tends to infinity as ℓ does; and
- over K_{ℓ} , the SCG ring-based LWE problem with scalar $|\Delta_{K_{\ell}}|^{(1-\varepsilon)/n_{\ell}}$, error parameter $r = \rho(n_{\ell})$ and modulus $q_{n_{\ell}}$ can be solved in time $\text{poly}(n_{\ell} \cdot \log q_{n_{\ell}})$, using $O(n_{\ell})$ samples.

The same statement is true for the actual ring-LWE as soon as one scales down the errors by a factor of $|\Delta_{K_{\ell}}|^{\varepsilon/n_{\ell}}$.

REMARK 2. We certainly do not claim that all number fields become vulnerable after scaling inappropriately: the fields K_{ℓ} that will be constructed below are very special in the sense that the lattices $\sigma(\mathcal{O}_{K_{\ell}})$ and $\sigma(\mathcal{O}_{K_{\ell}}^{\vee})$ are extremely ‘skew’, that is, they have widely varying successive minima. In particular, our findings do not seem to apply to cyclotomic number fields, which are the main candidates for making their way to every-day cryptography. Therefore the practical impact of Theorem 4.1 is limited.

Proof of Theorem 4.1. Fix an $\ell \geq 2$ and pick prime numbers p_1, \dots, p_{ℓ} congruent to 1 mod 4 such that

$$m_{\ell} := p_1 p_2 \dots p_{\ell} \geq (2\sqrt{n_{\ell}} \rho(n_{\ell}) \sqrt{\log n_{\ell}})^{2/\varepsilon}. \tag{4.1}$$

For each p_i , consider the corresponding quadratic field $K_{\ell,i} = \mathbb{Q}(\sqrt{p_i})$. It has discriminant p_i and ring of integers $R_{\ell,i} = \mathbb{Z}[(1 + \sqrt{p_i})/2]$, which we equip with the basis $\alpha_{i,1} = 1, \alpha_{i,2} = (1 + \sqrt{p_i})/2$. We will analyze ring-LWE and SCG ring-based LWE in the field compositum

$$K_{\ell} = \mathbb{Q}(\sqrt{p_1}, \sqrt{p_2}, \dots, \sqrt{p_{\ell}}) \cong K_{\ell,1} \otimes_{\mathbb{Q}} K_{\ell,2} \otimes_{\mathbb{Q}} \dots \otimes_{\mathbb{Q}} K_{\ell,\ell},$$

which is clearly of degree $n_\ell := 2^\ell$. Because the discriminants p_i of $\mathbb{Q}(\sqrt{p_i})$ are mutually coprime, this tensor structure carries over to the integral elements [21, Theorem 2.6]: that is, the ring R_ℓ of integers in K_ℓ reads

$$R_\ell = \mathbb{Z}[(1 + \sqrt{p_1})/2, (1 + \sqrt{p_2})/2, \dots, (1 + \sqrt{p_\ell})/2] \cong R_{\ell,1} \otimes_{\mathbb{Z}} R_{\ell,2} \otimes_{\mathbb{Z}} \dots \otimes_{\mathbb{Z}} R_{\ell,\ell}.$$

Please do not confuse this notation with our previous notation R_q for the reduction of R mod q (in fact, the modulus will not play an important role in the current proof). Note that R_ℓ is a complete intersection, so the different ideal $\partial_\ell \subset R_\ell$ is generated by $\theta_\ell = \sqrt{p_1}\sqrt{p_2}\dots\sqrt{p_\ell} = \sqrt{m_\ell}$. Therefore the codifferent reads

$$R_\ell^\vee = \frac{1}{\sqrt{m_\ell}}\mathbb{Z}[(1 + \sqrt{p_1})/2, (1 + \sqrt{p_2})/2, \dots, (1 + \sqrt{p_\ell})/2] \cong R_{\ell,1}^\vee \otimes_{\mathbb{Z}} R_{\ell,2}^\vee \otimes_{\mathbb{Z}} \dots \otimes_{\mathbb{Z}} R_{\ell,\ell}^\vee,$$

that is, it is again naturally compatible with the tensor structure of K_ℓ .

We begin with the actual ring-LWE, where we assume that the samples are expressed with respect to the product basis

$$\{\alpha_{1,i_1}\alpha_{2,i_2}\dots\alpha_{\ell,i_\ell}\}_{i_\ell \in \{1,2\}^\ell}, \tag{4.2}$$

where i_ℓ abbreviates $(i_1, i_2, \dots, i_\ell)$. With respect to this basis, a ring-LWE sample reads

$$(b_\ell)_i^t = \mathbf{A}_\mathbf{a} \cdot (s_\ell)_i^t + A_{\theta_\ell} \cdot B^{-1} \cdot (e_\ell)_i^t. \tag{4.3}$$

Here $A_\mathbf{a}$ and A_{θ_ℓ} are the matrices of multiplication by \mathbf{a} and $\theta_\ell = \sqrt{m_\ell}$, respectively, and $B^{-1} = \Sigma^{-1}$ is the inverse of the canonical embedding matrix (note that $U = I_{n_\ell \times n_\ell}$ because K_ℓ is totally real). We think of the e_ℓ as being sampled independently from Γ_r^1 with $r = \rho(n_\ell)/|\Delta_{K_\ell}|^{\varepsilon/n_\ell}$, and the whole expression is considered modulo q_{n_ℓ} .

Because we work with respect to the product basis, the matrix $A_{\theta_\ell} \cdot B^{-1}$ arises as the Kronecker product of the corresponding matrices for the quadratic fields $K_{\ell,i}$, which, by Example 2, are given by

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

Note that

$$(0 \ 1) \cdot \begin{pmatrix} \frac{-1 + \sqrt{d}}{2} & \frac{-1 - \sqrt{d}}{2} \\ 1 & 1 \end{pmatrix} = (1 \ 1), \tag{4.4}$$

so, through the Kronecker product, we find that

$$(0 \ 0 \ \dots \ 1) \cdot A_{\theta_\ell} \cdot B^{-1} = (1 \ 1 \ \dots \ 1),$$

where the row vector on the left has 0 everywhere, except at index $i_\ell = (2, 2, \dots, 2)$, where it has a 1.

Thus, given a ring-LWE sample (4.3), we can multiply both sides from the left by the row vector $(0 \ 0 \ \dots \ 1)$ in order to end up with a single linear equation in the secret $\mathbf{s} = (s_\ell)_i$, perturbed by an error of the form

$$(1 \ 1 \ \dots \ 1) \cdot (e_\ell)_i^t,$$

which behaves as if it were sampled from a univariate Gaussian $\Gamma_{r'}^1$, with $r' = \sqrt{n_\ell} \cdot r$. Now our primes p_i have been chosen in such a way that this error is likely to be negligible. More precisely, our bound (4.1) on m_ℓ implies that

$$r' = \frac{\sqrt{n_\ell} \cdot \rho(n_\ell)}{|\Delta_{K_\ell}|^{\varepsilon/n_\ell}} = \frac{\sqrt{n_\ell} \cdot \rho(n_\ell)}{\sqrt{m_\ell}^\varepsilon} \leq \frac{1}{2\sqrt{\log n_\ell}},$$

whose absolute value is less than $1/2$ with overwhelming probability, so a mere rounding results in an *exact* linear equation in the secret. In fact, by the lemma below, with very high probability, we can successfully repeat this during n_ℓ consecutive rounds, to end up with an exact linear system of n_ℓ equations in the n_ℓ unknowns s_ℓ . This system is likely to have full rank (if not, we can simply query a few more samples), so that the secret can be recovered using standard linear algebra over $\mathbb{Z}/q_{n_\ell}\mathbb{Z}$. This concludes the proof in the case of a proper ring-LWE.

To obtain the analogous result for a SCG ring-based LWE, using scaling factor $|\Delta_{K_\ell}|^{(1-\varepsilon)/n_\ell}$, one repeats the previous reasoning with $A_{\theta_\ell} \cdot B^{-1}$ replaced by $|\Delta_{K_\ell}|^{1/n} \cdot B^{-1}$. The analog of (4.4) reads

$$\begin{pmatrix} 0 & 1 \end{pmatrix} \cdot \begin{pmatrix} \frac{-1 + \sqrt{d}}{2} & \frac{1 + \sqrt{d}}{2} \\ 1 & -1 \end{pmatrix} = \begin{pmatrix} 1 & -1 \end{pmatrix},$$

leading to

$$\begin{pmatrix} 0 & 0 & \dots & 1 \end{pmatrix} \cdot |\Delta_{K_\ell}|^{1/n} \cdot B^{-1} = \begin{pmatrix} (-1)^{\eta(\iota)} \end{pmatrix}_\iota,$$

where $\eta(\iota)$ denotes the number of twos appearing in $\iota \in \{1, 2\}^\ell$. The right-hand side is again a norm $\sqrt{n_\ell}$ vector, which is the main ingredient needed for the rest of the proof to apply. \square

LEMMA 4.2. *Let P_n denote the probability that n independent samples from the univariate Gaussian $\Gamma_{1/2\sqrt{\log n}}^1$ are all at most $1/2$ in absolute value. Then $P_n \rightarrow 1$ as $n \rightarrow \infty$.*

Proof. Write $r = 1/2\sqrt{\log n}$ and let z be sampled from Γ_r^1 . Then P_n equals

$$\left(1 - 2P\left(z > \frac{1}{2}\right)\right)^n = \left(1 - \frac{2}{r} \int_{1/2}^\infty \exp\left(-\pi \frac{x^2}{r^2}\right)\right)^n \geq \left(1 - \frac{2}{r} \int_{1/2}^\infty 2x \exp\left(-\pi \frac{x^2}{r^2}\right)\right)^n,$$

so

$$P_n \geq \left(1 - \frac{\exp(-\pi \log n)}{\pi \sqrt{\log n}}\right)^n,$$

where the right-hand side is seen to converge to one using l'Hôpital's rule. \square

REMARK 3. The fields K_ℓ that were constructed in the above proof are totally real, but this is not essential. Indeed, if we also allow primes $p_i \equiv 3 \pmod 4$ and, instead, consider the field

$$K_\ell = \mathbb{Q}(\sqrt{p_1^*}, \sqrt{p_2^*}, \dots, \sqrt{p_\ell^*}),$$

where

$$p_i^* = (-1)^{(p_i-1)/2} p_i,$$

then the same conclusions would have followed.

As was pointed out to us by a reviewer, the part of Theorem 4.1 that deals with actual ring-LWE admits an easier and more broadly applicable proof. Just pick number fields having large enough discriminants, such as the ones constructed in the above proof, and apply the following observation.

THEOREM 4.3. *Let $(K_n)_{n \in \mathbb{N}}$ be a family of number fields of increasing degree n , let $\rho : \mathbb{N} \rightarrow \mathbb{R}_{>0}$ be in $\text{poly}(n)$ and let $(q_n)_{n \in \mathbb{N}}$ be any sequence of integer moduli. Then the ring-LWE problem in K_n with error parameter $r = \rho(n)$ can be solved in time $\text{poly}(n \cdot \log q_n)$, using $O(n)$ samples, as soon as the errors are being scaled down by at least $2\rho(n)\sqrt{n \log n}$.*

Proof. This is based on the simple fact that $1 \in R$, which implies that $\sigma(R)$ always contains the vector $(1, 1, \dots, 1)$. Thus there always exists a \mathbb{Z} -linear combination of the column vectors of $\Sigma = B \cdot U^{-1}$ having norm \sqrt{n} . As a consequence, the same \mathbb{Z} -linear combination of the rows of B^t is of said norm, meaning that, given a ring-LWE sample as in (3.2), one can extract from it a linear equation in the coordinates of the secret s_1, s_2, \dots, s_n that is perturbed by an error sampled from $\Gamma_{\rho(n) \cdot \sqrt{n}}^1$. As soon as one scales this down by a factor of size at least $2\rho(n)\sqrt{n \log n}$, by the previous lemma, about $O(n)$ samples suffice to recover \mathbf{s} . \square

5. A cyclotomic point of view

The fields K_ℓ constructed in the previous section are abelian; more precisely they are Galois with Galois group

$$\text{Gal}(K_\ell/\mathbb{Q}) \cong C_2 \times C_2 \times \dots \times C_2,$$

where C_2 denotes the group of order two. So, by the Kronecker–Weber theorem, it should be a subfield of some cyclotomic field. The following lemma shows that it is a subfield of $K := \mathbb{Q}(\zeta_{m_\ell})$. We identify the Galois group $\text{Gal}(K/\mathbb{Q})$ with $G := (\mathbb{Z}/(m_\ell))^\times$, where $a \in G$ acts on K as $\zeta_{m_\ell} \mapsto \zeta_{m_\ell}^a$.

LEMMA 5.1. *Let G^2 be the subgroup of squares in G . Then K_ℓ is the subfield of K fixed by G^2 .*

Proof. Denote the subfield of K fixed by G^2 as K^{G^2} . For each $c \in G/G^2$ consider

$$w_c = \text{Tr}_{K/K^{G^2}}(\zeta_{m_\ell}^c) = \sum_{h \in G^2} \zeta_{m_\ell}^{hc} \in K^{G^2}.$$

By the Chinese remainder theorem (CRT) we have the isomorphism

$$G \cong \mathbb{F}_{p_1}^\times \times \mathbb{F}_{p_2}^\times \times \dots \times \mathbb{F}_{p_\ell}^\times,$$

according to which the w_c can be decomposed as

$$w_c = \sum_{h \in G^2} \zeta_{m_\ell}^{hc} = \sum_{\substack{h_1 \in (\mathbb{F}_{p_1}^\times)^2 \\ h_\ell \in (\mathbb{F}_{p_\ell}^\times)^2}} \zeta_{p_1}^{h_1 c} \zeta_{p_2}^{h_2 c} \dots \zeta_{p_\ell}^{h_\ell c} = \prod_{i=1}^\ell \sum_{h \in (\mathbb{F}_{p_i}^\times)^2} \zeta_{p_i}^{hc}. \tag{5.1}$$

Every sum in the last product is a so-called Gaussian period, where the exponents run through either the quadratic residues or the quadratic non-residues modulo p_i . As all p_i are congruent to 1 modulo 4, such sums result in

$$\beta_{i,1} := \frac{-1 + \sqrt{p_i}}{2}, \quad \text{respectively, } \beta_{i,-1} := \frac{-1 - \sqrt{p_i}}{2}$$

(see [8]). One sees that $\{w_c\}_c$ is the product basis of K_ℓ obtained by equipping the $R_{\ell,i}$ with the \mathbb{Z} -bases $\beta_{i,1}, \beta_{i,-1}$ rather than $\alpha_{i,1}, \alpha_{i,2}$. In particular, the w_c generate K_ℓ , so $K_\ell \subset K^{G^2}$ and the lemma follows by comparing degrees. \square

As a byproduct of the above proof, we obtain that the w_c form a \mathbb{Z} -basis of R_ℓ , which is a special case of a more general statement [15, Proposition 6.1]. This kind of ‘trace basis’ is also used in the recent work on SCG ring-based LWE by Chen, Lauter and Stange [5], an example of which we will analyze later in this section. It is interesting to have a quick look at

our proof of Theorem 4.1, where now we express the samples with respect to the basis $\{w_c\}_c$, instead of (4.2). Here the factors in the Kronecker product decomposition of $A_{\theta_\ell} \cdot B^{-1}$ read

$$\left(\begin{array}{cc} \frac{-1-p_i}{2} & \frac{-1+p_i}{2} \\ \frac{1-p_i}{2} & \frac{1+p_i}{2} \end{array} \right) \cdot \frac{1}{\sqrt{p_i}} \left(\begin{array}{cc} \frac{1-\sqrt{p_i}}{2} & \frac{-1-\sqrt{p_i}}{2} \\ \frac{-1-\sqrt{p_i}}{2} & \frac{1-\sqrt{p_i}}{2} \end{array} \right) = \left(\begin{array}{cc} \frac{1-\sqrt{p_i}}{2} & \frac{1+\sqrt{p_i}}{2} \\ \frac{-1-\sqrt{p_i}}{2} & \frac{-1+\sqrt{p_i}}{2} \end{array} \right).$$

One sees that

$$(1 \ -1) \cdot \left(\begin{array}{cc} \frac{1-\sqrt{p_i}}{2} & \frac{1+\sqrt{p_i}}{2} \\ \frac{-1-\sqrt{p_i}}{2} & \frac{-1+\sqrt{p_i}}{2} \end{array} \right) = (1 \ 1).$$

So expanding the Kronecker product gives

$$(J(\iota))_\iota \cdot A_{\theta_\ell} \cdot B^{-1} = (1 \ 1 \ \dots \ 1), \tag{5.2}$$

where ι runs over all tuples $(i_1, i_2, \dots, i_\ell) \in \{1, -1\}^\ell$ and

$$J(\iota) = J(i_1, i_2, \dots, i_\ell) = \prod_{j=1}^{\ell} i_j$$

(this formula explains why we indexed the β_i by ± 1 rather than $1, 2$). The row vector $(1 \ 1 \ \dots \ 1)$ on the right-hand side of (5.2) has norm $\sqrt{m_\ell}$, so, as before, this can be used to obtain linear equations in the coordinates of the secret \mathbf{s} that carry negligible error terms, allowing one to recover \mathbf{s} by means of simple linear algebra.

REMARK 4. As before, the same claims apply to the SCG ring-based LWE and/or to the setting where we allow primes $p_i \equiv 3 \pmod{4}$, upon replacement of every appearance of $\sqrt{p_i}$ by $\sqrt{p_i^*}$.

REMARK 5. The letter J refers to the Jacobi symbol. Indeed, through the CRT,

$$G/G^2 \cong \frac{\mathbb{F}_{p_1}^\times}{(\mathbb{F}_{p_1}^\times)^2} \times \frac{\mathbb{F}_{p_2}^\times}{(\mathbb{F}_{p_2}^\times)^2} \times \dots \times \frac{\mathbb{F}_{p_\ell}^\times}{(\mathbb{F}_{p_\ell}^\times)^2} = \{\pm 1\} \times \{\pm 1\} \times \dots \times \{\pm 1\},$$

where, if $c \in G/G^2$ corresponds to $\iota = (i_1, i_2, \dots, i_\ell) \in \{1, -1\}^\ell$, then $w_c = \beta_{1,i_1} \beta_{2,i_2} \dots \beta_{\ell,i_\ell}$ and $J(\iota) = (c/m_\ell)$. Thus if we prefer to think of the rows and columns of the matrices A_θ and M as being indexed by $c \in G/G^2$ rather than $\iota \in \{1, -1\}^\ell$, then (5.2) becomes

$$\left(\left(\begin{array}{c} c \\ m_\ell \end{array} \right) \right)_c \cdot A_{\theta_\ell} \cdot M^{-1} = (1 \ 1 \ \dots \ 1),$$

an identity which we found remarkable at first sight.

To conclude this article, we note that, more generally, the presence of factors of the form $\mathbb{Z}[(1 + \sqrt{d})/2]$ for some $d \equiv 1 \pmod{4}$ may lead to unexpectedly short linear combinations of the rows of $A_\theta \cdot B^{-1}$ and B^{-1} and thus to weaker instances of ring-LWE and SCG ring-based LWE than one might hope, for an aggressive choice of scaling factor.

For instance, let us analyze the first example listed in [5, § 5.1]; the other examples admit a similar analysis. Here Chen *et al.* let $m = 2805 = 3 \cdot 5 \cdot 11 \cdot 17$ and they consider the fixed field $K^{G'}$ of $K = \mathbb{Q}(\zeta_m)$ under the action of

$$G' := \langle 1684, 1618 \rangle \subset G = \text{Gal}(K/\mathbb{Q}) = (\mathbb{Z}/(m))^\times.$$

Under the CRT decomposition $(\mathbb{Z}/(m))^\times \cong \mathbb{F}_3^\times \times \mathbb{F}_{11}^\times \times (\mathbb{Z}/(85))^\times$, this subgroup corresponds to $\{1\} \times \{1\} \times G'_{85}$, where G'_{85} denotes the index two subgroup of elements having Jacobi symbol 1. We again work with respect to the trace basis

$$w_c = \sum_{h \in G'} \zeta_m^{hc} = \zeta_3^c \cdot \zeta_{11}^c \cdot \sum_{h \in G'_{85}} \zeta_{85}^{hc}$$

where $c \in G/G'$. The latter sum equals $\beta_1 := (1 + \sqrt{85})/2$ or $\beta_{-1} := (1 - \sqrt{85})/2$, depending on whether $(c/85) = 1$ or not. As before, we conclude that the ring of integers equals

$$R := \mathcal{O}_{K^{G'}} = \mathbb{Z}[\zeta_3, \zeta_{11}, (1 + \sqrt{85})/2] \cong \mathbb{Z}[\zeta_3] \otimes_{\mathbb{Z}} \mathbb{Z}[\zeta_{11}] \otimes_{\mathbb{Z}} \mathbb{Z}[(1 + \sqrt{85})/2]$$

and that $\{w_c\}_c$ is the product basis

$$\{\zeta_3^i \zeta_{11}^j \beta_k\}_{\substack{i=1,2 \\ j=1,2,\dots,10 \\ k=1,-1}}$$

As in [5], let us have a look at SCG ring-based LWE with scaling factor $|\Delta|^{1/2n}$, where $\Delta = \Delta_{K^{G'}} = (-3) \cdot (-11^9) \cdot 85$ and $n = [K^{G'} : \mathbb{Q}] = 40$. Let M denote the matrix of the canonical embedding of $K^{G'}$ with respect to the above basis. Then the last Kronecker factor of $|\Delta|^{1/2n} \cdot M^{-1} = |\Delta|^{1/80} \cdot M^{-1}$ is given by

$$\frac{1}{\sqrt[4]{85}} \cdot \begin{pmatrix} \frac{1 + \sqrt{85}}{2} & \frac{-1 + \sqrt{85}}{2} \\ \frac{-1 + \sqrt{85}}{2} & \frac{1 + \sqrt{85}}{2} \end{pmatrix}.$$

So multiplying from the left by $(1 \ -1)$ leads to the row vector $(1 \ 1)/\sqrt[4]{85}$ of norm ≈ 0.4658 , which is ‘unexpectedly short’. The other Kronecker factors correspond to cyclotomic fields and have less surprising behavior. Here taking the first row (for instance) of each factor leads to norms $\sqrt{2}/\sqrt[4]{3} \approx 1.0746$ and $\sqrt{10}/\sqrt[20]{11^9} \approx 1.0750$, respectively. Thus multiplying $|\Delta|^{1/80} \cdot B^{-1}$ from the left by

$$(1, 0) \otimes (1, 0, 0, 0, 0, 0, 0, 0, 0) \otimes (1, -1)$$

yields a row vector of norm $\approx 1.0746 \cdot 1.0750 \cdot 0.4658 \approx 0.5381$. Since Chen *et al.* let $r = 1$, this results in a linear equation in the secret \mathbf{s} carrying an error term sampled from $\Gamma_{0.5381}^1$, roughly. By taking other rows of the cyclotomic parts one, in fact, finds twenty such independent equations. This is insufficient to break this concrete instance of SCG ring-based LWE using mere rounding (a substantial number of equations will carry an error that exceeds 1/2 in absolute value), but it is tight, so it provides an explanation as to why this was indirectly helpful to Chen *et al.* for successfully applying their χ^2 -analysis.

6. Conclusion

In this paper, we explained that *if* one wishes to set up SCG ring-based LWE in a degree n number field K , as was done in [5, 6, 11] in the context of potential attacks involving evaluation at one, then it is natural to scale up the errors by $|\Delta_K|^{1/n}$. More precisely, we proved that, for each $\varepsilon > 0$, scaling up by $|\Delta_K|^{(1-\varepsilon)/n}$ may indeed be insufficient, in the sense that there exist number fields for which the corresponding problem is easily broken. These observations also apply to the proper ring-LWE, in the sense that scaling down by $|\Delta_K|^{\varepsilon/n}$

leads to vulnerable families for any $\varepsilon > 0$. Some of our families implicitly exploit the structure of the Galois group, which raises the question of to what extent Galois theory can be used further in the analysis of the hardness of ring-LWE.

In any case, we stress that the families constructed in this paper are very special. In particular, it is unlikely that they will ever be used in a cryptographic context. Our main aim is to help delimit the room for flexibility when tweaking the parameters, or even the definition, of ring-LWE, as introduced in [16]. We refer the interested reader to the recent work of Peikert [18], in which a unifying framework is given.

References

1. Z. BRAKERSKI, C. GENTRY and V. VAIKUNTHANATHAN, ‘(Leveled) Fully homomorphic encryption without bootstrapping’, *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference – ITCS ’12* (ACM, New York, NY, 2012) 309–325.
2. Z. BRAKERSKI, A. LANGLOIS, C. PEIKERT, O. REGEV and D. STEHLÉ, ‘Classical hardness of learning with errors’, *ACM Symposium on the Theory of Computing – STOC ’13* (ACM, New York, NY, 2013) 575–584.
3. Z. BRAKERSKI and V. VAIKUNTHANATHAN, ‘Efficient fully homomorphic encryption from (standard) LWE’, *Proceedings of the 2011 IEEE 52nd Annual Symposium on Foundations of Computer Science – FOCS ’11* (IEEE, Washington, DC, 2011) 97–106.
4. W. CASTRYCK, I. ILIASHENKO and F. VERCAUTEREN, ‘Provably weak instances of Ring-LWE revisited’, *Advances in cryptology – EUROCRYPT 2016*, Lecture Notes in Computer Science 9665(1) (Springer, New York, NY, 2016) 147–167.
5. H. CHEN, K. LAUTER and K. STANGE, ‘Attacks on search RLWE’, Cryptology ePreprint Archive, Report 2015/971 2015.
6. H. CHEN, K. LAUTER and K. STANGE, ‘Vulnerable Galois RLWE families and improved attacks’, *Proceedings of Selected Areas in Cryptography (SAC 2016, St. John’s, Canada)*, Lecture Notes in Computer Science (Springer, New York, NY, to appear); Cryptology ePreprint Archive, Report 2016/193 2016.
7. E. CROCKETT and C. PEIKERT, ‘ $\Lambda \circ \lambda$: A functional library for lattice cryptography’, Cryptology ePreprint Archive, Report 2015/1134 2015.
8. H. DAVENPORT, *Multiplicative number theory*, 2nd edn, Graduate Texts in Mathematics 74 (Springer, New York, NY, 2000) (revised by H. Montgomery).
9. B. DE SMIT, ‘A differential criterion for complete intersections’, *Journées Arithmétiques 1995, Collect. Math.* 48 (1997) no. 1–2, 85–96.
10. K. EISENTRÄGER, S. HALLGREN and K. LAUTER, ‘Weak instances of PLWE’, *Selected areas in cryptography – SAC 2014*, Lecture Notes in Computer Science 8781 (Springer, New York, NY, 2014) 183–194.
11. Y. ELIAS, K. LAUTER, E. OZMAN and K. STANGE, ‘Provably weak instances of Ring-LWE’, *Advances in cryptology – CRYPTO ’15*, Lecture Notes in Computer Science 9215 (Springer, New York, NY, 2015) 63–92.
12. A. FRÖHLICH and M. TAYLOR, *Algebraic number theory*, Cambridge Studies in Advances Mathematics 27 (Cambridge University Press, Cambridge, 1991).
13. C. GENTRY, ‘Key recovery and message attacks on NTRU-Composite’, *EUROCRYPT ’01*, Lecture Notes in Computer Science 2045 (Springer, New York, NY, 2001) 182–194.
14. J. HOFFSTEIN, J. PIPHER and J. H. SILVERMAN, ‘NTRU: a ring-based public key cryptosystem’, *Proceedings of the Third International Symposium on Algorithmic Number Theory – ANTS-III*, Lecture Notes in Computer Science 1423 (Springer, New York, NY, 1998) 267–288.
15. H. JOHNSTON, ‘Notes on Galois modules’, *Notes accompanying the course ‘Galois Modules’ given in Cambridge* (2011), <https://www.dpms.cam.ac.uk/~hj31/GM.CourseNotes101.pdf> [accessed 22 July 2016].
16. V. LYUBASHEVSKY, C. PEIKERT and O. REGEV, ‘On ideal lattices and learning with errors over rings’, *J. ACM* 60 (2013) no. 6, article 43, 35.
17. C. PEIKERT, ‘Public-key cryptosystems from the worst-case shortest vector problem’, *ACM Symposium on the Theory of Computing – STOC ’09* (ACM, New York, NY, 2009) 333–342.
18. C. PEIKERT, ‘How (not) to instantiate Ring-LWE’, Cryptology ePrint Archive, Report 2016/351 2016.
19. O. REGEV, ‘On lattices, learning with errors, random linear codes, and cryptography’, *J. ACM* 56 (2009) no. 6, article 34, 40.
20. P. SHOR, ‘Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer’, *SIAM J. Comput.* 26 (1997) no. 5, 1484–1509.
21. L. WASHINGTON, *Introduction to cyclotomic fields*, Graduate Texts in Mathematics 83 (Springer, New York, NY, 1982).

Wouter Castryck
KU Leuven ESAT/COSIC and iMinds
Kasteelpark Arenberg 10
B-3001 Leuven-Heverlee
Belgium

and

Vakgroep Wiskunde
Universiteit Gent
Krijgslaan 281/S22
B-9000 Gent
Belgium

wouter.castryck@gmail.com

Frederik Vercauteren
KU Leuven ESAT/COSIC and iMinds
Kasteelpark Arenberg 10
B-3001 Leuven-Heverlee
Belgium

and

Open Security Research
Fangda 704
11 Kejinan 12th road
518000 Shenzhen
China

frederik.vercauteren@gmail.com

Ilia Iliashenko
KU Leuven ESAT/COSIC and iMinds
Kasteelpark Arenberg 10
B-3001 Leuven-Heverlee
Belgium

ilia.iliashenko@esat.kuleuven.be