Algebraic aspects of solving Ring-LWE, including ring-based improvements in the Blum-Kalai-Wasserman algorithm

KATHERINE E. STANGE

Abstract. We provide a reduction of the Ring-LWE problem to Ring-LWE problems in subrings, in the presence of samples of a restricted form (i.e. (a, b) such that a is restricted to a multiplicative coset of the subring). To create and exploit such restricted samples, we propose Ring-BKW, a version of the Blum-Kalai-Wasserman algorithm which respects the ring structure. Off-the-shelf BKW dimension reduction (including coded-BKW and sieving) can be used for the reduction phase. Its primary advantage is that there is no need for back-substitution, and the solving/hypothesis-testing phase can be parallelized. There is also the opportunity to exploit symmetry to reduce table sizes. The results apply to two-power cyclotomic Ring-LWE with parameters proposed for practical use (including all splitting types).

1. Introduction

Ring Learning with Errors (Ring-LWE) [19] [20], and Learning with Errors (LWE) [22] more generally, are leading candidates for post-quantum cryptography. The cryptographic hard problem (*Search Ring-LWE*) is formally similar to discrete logarithm problems, so that protocols can be transferred from the latter context to the former. But it also allows for new applications, such as homomorphic encryption [6]. Ring-LWE is also fortunate in having security reductions from other lattice problems.

Ring-LWE is distinguished from Learning with Errors (LWE) by the use of lattices from number fields. This injection of number-theoretical structure leads to performance improvements, but may add vulnerabilities. So far, the number-theoretical structure has been only weakly exploited for attacks. The ring structure plays a role in security when the error distribution is skewed [8] [9] [10] [13] [14], or the secret is chosen from a subring or other ring-related non-uniform distribution [5]. However, the best known attacks on parameters suggested for implementation are still generic attacks for LWE, e.g. [2]. The Blum-Kalai-Wasserman (BKW) algorithm is one

---

Date: May 13, 2019.

2010 *Mathematics Subject Classification.* Primary: 94A60, 11T71, 11R18.

Key words and phrases. Ring learning with errors, Ring-LWE, Blum-Kalai-Wasserman, post-quantum cryptography, cyclotomic field.

This research was supported by NSF-CAREER CNS-1652238 and NSF EAGER DMS-1643552.

such attack, which proceeds (in the first phase) combinatorially to create new samples in a linear subspace of the original problem, while controlling error expansion [3]. BKW has the drawback of requiring exponentially many samples, so lattice algorithms are generally preferred. Nevertheless, its performance has been of significant interest: for analysis and recent improvements, see [1] [12] [15] [16] [17] [18].

This paper focuses on two-power-cyclotomic unital (but equivalently, dual [11] [21]) Search Ring-LWE, with no restriction on the splitting behaviour of the prime q. The core of the paper is a reduction from higher-dimensional Ring-LWE problems with samples of a restricted form, to lower-dimensional Ring-LWE problems with the same error width, which is given in Theorem 5.2. The restricted form is as follows: samples (a, b) such that a lies in a cyclotomic subring, or a fixed multiplicative coset of such a subring. In the context of these theorems, it is natural to ask about creating samples of this restricted form using a ring variant of the Blum-Kalai-Wasserman algorithm.

One thus obtains a *Ring-BKW* algorithm, which uses the reduction phase of BKW, including all known speedups, to reduce the Ring-LWE problem to a subring. Then, the symmetry of the ring structure allows us to engineer an entire suite of subring problems in polynomially more time, whose solutions collectively solve original Ring-LWE problem, again in polynomial time. *Thus, the 'hypothesis testing' phase of BKW is parallelized, and the exponential 'back-substitution' phase is eliminated.* State-of-the-art off-the-shelf code for the BKW reduction phase and hypothesis testing phase may be used. In Section 8, we describe the Ring-BKW algorithm.

The paper also addresses the use of symmetry to reduce the table sizes in BKW, here termed *advanced keying* in Section 9. We also discuss a square-root speedup over exhaustive search (which may be used, for example, in hypothesis testing); see Corollary 5.3.

The key theoretical properties which are potentially advantageous (to an attacker) of Ring-LWE vs. plain LWE, are:

(1) Ring homomorphisms into smaller instances of the problem (the main tool of [9] [10] [13] [14]).
(2) The ability to *rotate* samples, e.g. replacing (a, b) with $(\zeta a, \zeta b)$ or $(a, \zeta b)$, which are different but related Ring-LWE samples (see notation in Section 2); these represent symmetries of the lattice.
(3) The existence of subrings as linear subspaces (which is important in [5]).
(4) More generally, the multiplicative structure of certain linear subspaces.
(5) In the case of 2-power cyclotomics, the orthogonality of the lattice of the ring of integers and the orthogonal nature of the trace.

For us, all five of these attributes play an important role. It is a secondary purpose of this paper to lay out these advantages in a clear manner, to facilitate future analysis of the security of ring aspects of Ring-LWE. See Section 4.

Finally, it is also a secondary purpose of this paper to provide a treatment of the Ring-LWE problem which is inviting to the mathematical community.

Code demonstrating the correctness of the algorithm is available at:
https://math.katestange.net/code/ring-bkw/.

Acknowledgements. First, I would like to thank the anonymous referees on an earlier draft of this paper, who pointed out an important simplification. Second, I would like to thank my mother, Ursula Stange, and my husband, Jonathan Wise, without whose childcare help in the face of snowstorms, viruses, cancellations and fender-benders, this paper simply would not have been completed. To mathematician moms (and dads) everywhere: take heart.

2. Background and Setup for Ring-LWE

It is typical to set notation for Ring-LWE as in, for example, [5]; here we briefly review this notation in our context, and define the Ring-LWE problems.

2.1. **Number field K and ring R.** Let K be a number field over the rationals, of degree n. Then K is equipped with a bilinear form given by a modification of the trace pairing,

$$\langle \alpha, \beta \rangle = \sum_{\sigma \mathbb{R}} \sigma(\alpha\beta) + \frac{1}{2} \sum_{\sigma \mathbb{C}} \mathrm{Re}(\sigma(\alpha)\overline{\sigma(\beta)}). \tag{1}$$

Here the sums are over real and complex embeddings, respectively (note that including both elements of each pair of conjugate complex embeddings necessitates the factor of $\frac{1}{2}$). This gives an isomorphism of $K_{\mathbb{R}} := \mathbb{R} \otimes_{\mathbb{Q}} K$ with $\mathbb{R}^n$, taking the pairing above to the standard inner product. We can also denote the norm by $||\mathbf{x}|| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$. One can also access this inner product using the Minkowski or canonical embedding, as the standard inner product.

The ring of integers R of K forms a lattice in $K_{\mathbb{R}}$.

2.2. **Gaussian distribution.** Having geometry (in particular a norm $||\cdot||^2$) on $K_{\mathbb{R}}$ allows us to define Gaussian distributions. For a *Gaussian parameter* $r > 0$, we write

$$\rho_r : K_{\mathbb{R}} \to (0, 1], \quad \rho_r(\mathbf{x}) = \exp(-\pi ||\mathbf{x}||^2 / r^2).$$

Normalizing this to obtain a probability distribution function $r^{-n}\rho_r$, we obtain the *continuous Gaussian probability distribution of width* r on $K_{\mathbb{R}}$, denoted D_r.

Note that, when considered with respect to an orthonormal basis, such a distribution is the sum of independent distributions in each coordinate, each having width r. In this paper, we are concerned exclusively with this case.

With this normalization, the variance is $r^2/2\pi$, and one standard deviation is $r/\sqrt{2\pi}$. It is a sum of independent Gaussians in each coordinate for which the range $[-r, r]$ corresponds to $\sqrt{\pi/2} \sim 1.25\ldots$ standard deviations.

In practice, the tails of the Gaussian may be cut off, so that the number of possible values in each coordinate is finite.

One may discretize a Gaussian distribution to obtain a distribution $\mathcal{D}_r$ on a lattice $\mathcal{L} \subset K_{\mathbb{R}}$. That is, one takes

$$\rho_r(\mathcal{L}) = \sum_{\lambda \in \mathcal{L}} \rho_r(\lambda)$$

and one samples element $\lambda \in \mathcal{L}$ with probability

$$\frac{\rho_r(\lambda)}{\rho_r(\mathcal{L})}.$$

If $\mathcal{L}$ has an orthonormal basis, then again this distribution consists of independent distributions on the coefficients of the basis.

2.3. **Prime q and quotient ring R_q.** Let qR be the extension to R of a rational prime ideal $q\mathbb{Z}$. The fundamental setting of the Ring-LWE problem is the ring $R_q := R/qR$.

Letting $q = \mathfrak{q}_1^{e_1} \cdots \mathfrak{q}_g^{e_g}$ be the unique decomposition of q into distinct prime ideals $\mathfrak{q}_i$ in K, the Chinese remainder theorem gives

$$R_q \cong \bigoplus_{i=1}^{g} R/\mathfrak{q}_i^{e_i}.$$

If q is unramified (which is typically the case), then $e_i = 1$ for all i. If K is Galois (also typically the case), then the Galois group acts transitively on the $\mathfrak{q}_i$ and they all have the same residue degree (the residue degree is the dimension of the quotient field $R/\mathfrak{q}_i$ as an $\mathbb{F}_q$-vector space).

2.4. **Ring-LWE distributions.** For any $s \in R_q$ (the *secret*), and any distribution ψ over R_q (the *error distribution*), we write $A_{s,\psi}$ for the associated *Ring-LWE distribution for secret s* over $R_q \times R_q$, given by sampling a uniformly over R_q, sampling e from ψ, and outputting $(a, b := as + e)$.

Such outputs (a, b) are called *samples*, and in a crytographic application, these are observed publicly, while the secret is not meant to be exposed.

For the error distribution, we wish to define a 'small' distribution on R_q, i.e. concentrated near the origin (in comparison to q, which is large). It is typical to choose for the error distribution a discretized Gaussian distribution as described above (considered post factum modulo qR). This is the context in which security reductions apply. In implementations, it is sometimes suggested to approximate this by a uniform distribution on a box around the origin, etc.

2.5. **Ring-LWE Problems.** The two fundamental Ring-LWE problems are (a) *search*: to compute the secret, upon observing sufficiently many samples; or (b) *decision*: to determine if the sample are hiding a secret at all, as opposed to being random noise. We state them more formally as follows.

Definition 2.1. The *search Ring-LWE problem*, for error distribution ψ and secret distribution φ, is as follows: Given an error distribution ψ over R_q and a secret distribution φ over R_q, and some number of samples drawn from the distribution $A_{s,\psi}$ for some fixed s drawn from φ, compute s.

Definition 2.2. The *decisional Ring-LWE problem*, for error distribution ψ and secret distribution φ, is as follows: Given an error distribution ψ over R_q and a secret distribution φ over R_q, distinguish with non-negligable advantage, between

(1) samples drawn from the distribution $A_{s,\psi}$ for some fixed s drawn from φ; and
(2) samples drawn uniformly from $R_q \times R_q$.

We remark that Ring-LWE is frequently defined in the context of the dual $R^\vee$ (the different ideal), but in the case that K is a two-power cyclotomic field, we have $R^\vee$ is isomorphic to R as an R-module, so we can interchange with the simpler 'unital' version considered here [19].

Search-to-decision reductions are known in a variety of contexts [19]. This paper concerns both problems, but especially the search problem.

The Ring-LWE problem is formally similar to the discrete logarithm problem, which could be phrased in terms of *samples* (a, a^s) in a finite field: given (a, a^s), find s. In the ring R_q, solving for s given (a, as) can be accomplished using linear algebra (Gaussian elimination). By introducing a small error e, so we have $(a, as + e)$, Gaussian elimination becomes useless, as it amplifies the errors to the point of washing out all useful information. From another perspective, the security stems from the fact that addition of an error value is somehow unpredictably mixing with respect to the multiplicative structure.

Another consequence of this setup is that given just one sample (a, b), one has as many solutions s to $b = as + e$ as there are possible values for e. In fact, the problem only has a unique solution once we have enough samples. If the samples are not Ring-LWE samples at all, then with sufficiently many samples, it becomes overwhelmingly likely that there are no values of s so that $b_i - a_i s$ is in the support of the error distribution for all samples s. If the samples are Ring-LWE, this is the point at which the true secret is the only solution, with overwhelming probability.

3. Specializing to 2-power cyclotomic Ring-LWE

We now specialize to the following situation, fixing the variables

$$K, R, q, R_q, m, n, \zeta := \zeta_m, \chi, \chi_0, E_{\chi_0}$$

for the 2-power cyclotomic case. Whenever we say *consider 2-power cyclotomic Ring-LWE*, we refer to all the conventions in this section.

3.1. **Ring R.** We let K and R be the $2n$-th cyclotomic field and ring of integers, respectively, where n is a power of two. This is of dimension n (note that $\varphi(2n) = n$), and can be presented as

$$R = \mathbb{Z}[\zeta_{2n}] = \mathbb{Z}[x]/(x^n + 1).$$

We will use the notation $m = 2n$ and ζ_m for a primitive m-th root of unity in R *and for its image in quotients of this ring.*

3.2. **The ζ-basis for R and its quotients.** A basis for R is

$$1, \zeta_m, \zeta_m^2, \dots, \zeta_m^{n-1}.$$

This will be called the *ζ-basis*. We have the relation $\zeta_m^n + 1 = 0$ in R and in all its quotients (this is the $2n$-th cyclotomic polynomial evaluated at ζ_m), but the minimal polynomial for ζ_m varies in these quotients, and may be a proper divisor of this cyclotomic polynomial. Nevertheless, in all quotients of R, we still obtain a ζ-basis, i.e. a power basis in terms of $\zeta := \zeta_m$.

3.3. **Prime q.** Let q be an odd prime, unramified in R.

3.4. **Ring R_q and further quotients.** We consider the quotient ring

$$R_q = R/qR \cong (\mathbb{Z}/q\mathbb{Z})[x]/(x^n + 1),$$

which is an $\mathbb{F}_q$-vector space of dimension n. We may use the same ζ-basis for this ring.

We may also consider further quotients $R/\mathfrak{a}$ for $\mathfrak{a} \mid qR$. We may also use a ζ-basis for these rings, although it may be of lower dimension over $\mathbb{F}_q$ (so fewer powers required). We have

$$R/\mathfrak{a} \cong \mathbb{F}_q[x]/(g(x))$$

where $g(x) \mid x^n + 1$. In particular, identifying $\zeta \in R/qR$ with its image in $R/\mathfrak{a}$, the latter has an $\mathbb{F}_q$-basis $1, \zeta, \zeta^2, \dots, \zeta^{\deg(g)-1}$.

3.5. **Error distribution χ, coefficient distribution χ_0 and coefficient support E_{χ_0}.** We will denote the error distribution by χ. If this error distribution is formed using independent identically distributed coefficients on the ζ-basis, with coefficient distribution χ_0 supported on a subset $E_{\chi_0} \subseteq \mathbb{F}_q$, then we say that χ is *formed on a ζ-basis with coefficients distributed according to* χ_0. This is true, for example, of a discrete Gaussian distribution on two-power cyclotomics, or a distribution formed by choosing coefficients uniformly from some subset of $\mathbb{F}_q$. For the former observation, the relevant fact is the following: *the power basis associated to ζ_m is orthonormal in the canonical embedding.* To see this, use (1) and observe that if ζ_m^a has order $2^\ell \geq 2$, then $-\overline{\zeta_m}^a$ does also, hence the real parts of the complex embeddings of roots of unity form a collection symmetrical about zero[1].

[1] This is an advantage of the awkward alteration of the trace pairing in (1).

For this paper, we will concern ourselves exclusively with this case.

3.6. **Secret distribution.** We will not make any particular assumption on the secret distribution. It may be taken to be uniform on R_q.

4. Key theoretical properties

In this section we highlight several key aspects of Ring-LWE absent in LWE.

4.1. **Ring homomorphisms.** If a Ring-LWE problem is presented in R_q, then for any $\mathfrak{a} \mid qR$, we have a ring homomorphism

$$\rho : R_q \to R/\mathfrak{a}.$$

This transports samples distributed according to $A_{s,\chi}$ to samples distributed according to $A_{\rho(s),\rho(\chi)}$.

In general, the effect of ρ on χ is problematic, i.e. it spreads out the error widely. As an illustration, we give a proposition governing the behaviour of ρ on χ in the 2-power cyclotomic case, when $q \equiv 1 \pmod 4$.

Proposition 4.1. *Suppose we are in the 2-power cyclotomic case, and $R/\mathfrak{a} \cong \mathbb{F}_{q^k}$, and $q \equiv 1 \pmod 4$. If, in R_q, the error distribution χ is formed on the ζ-basis in R_q with coefficients drawn from χ_0 on $\mathbb{F}_q$, then $\chi' := \rho(\chi)$ is formed on the ζ-basis in $\mathbb{F}_{q^k}$ with coefficients drawn from χ_0' on $\mathbb{F}_q$, where $\rho(\zeta_m^k) \in \mathbb{F}_q$ and*

$$\chi_0' = \sum_{i=0}^{n/k-1} \rho(\zeta_m^k)^i \chi_0.$$

Proof. Define $r = \operatorname{ord}_2(q-1)$, meaning that $2^r \mid q-1$ but $2^{r+1} \nmid q-1$. Since $q \equiv 1 \pmod 4$, we have $r \geq 2$. Furthermore, $q^i + 1 \equiv 2 \pmod 4$ for all i, so that $\operatorname{ord}_2(q^2-1) = \operatorname{ord}_2((q-1)(q+1)) = r+1$ and, by induction

$$\operatorname{ord}_2(q^{2^i} - 1) = r + i$$

for all $i \geq 1$. As k is defined as the embedding degree of the $2n$-th roots of unity, we obtain $k = \frac{2n}{2^r}$.

The element $\rho(\zeta_m^k)$ satisfies $\rho(\zeta_m^k)^{2n/k} = 1$ in $R/\mathfrak{a}$. Hence it is itself a primitive $2n/k$-th root of unity, i.e. 2^r-th root of unity. Hence $\rho(\zeta_m^k) \in \mathbb{F}_q$ by the definition of r.

The main statement now follows from the fact that $1, \zeta_m, \dots, \zeta_m^{k-1}$ is an $\mathbb{F}_q$-basis of $\mathbb{F}_{q^k}$, that $\rho(\zeta_m^k) \in \mathbb{F}_q$ and that for $0 \leq j < k$ and $0 \leq i < n/k$, we have

$$\rho(\zeta_m^{ik+j}) = \rho(\zeta_m^k)^i \rho(\zeta_m^j) = \rho(\zeta_m^k)^i \zeta_m^j.$$

□ □

For example, in the case that $k = n/2$, we obtain

$$\chi_0' = \chi_0 + \rho(\zeta_m^k)\chi_0.$$

This means the coefficients of χ' are chosen from a sum of two Gaussian distributions with different coefficients. This is worse than twice a single Gaussian. For, the latter is simply a wider Gaussian, and the size of its support grows by approximately $\sqrt{2}$. However, here the size of the support $E_{\chi'}$ is approximately the square of the size of E_χ. This is a symptom of the protective property of these ring homomorphisms: they transform the error badly. In fact, very quickly the image of a Gaussian error approaches uniform in the image ring as the dimension of the image ring decreases.

4.2. **Rotating samples.** The ring structure allows us to generate new (but not independent) samples from old.

Proposition 4.2. *Suppose χ is invariant under multiplication by ζ. Then if (a, b) is distributed according to $A_{s,\chi}$, then*

(1) $(\zeta a, \zeta b)$ is also distributed according to $A_{s,\chi}$,
(2) $(a, \zeta b)$ is distributed according to $A_{\zeta s,\chi}$.

In particular, in the 2-power cyclotomic case, a discrete Gaussian is invariant under multiplication by ζ_m and all its powers.

We call these *rotated samples*. One could also rotate by other small values, e.g. $1 + \zeta_m$ in the 2-power cyclotomic case, at a small cost in changing the error distribution.

4.3. **Subrings and trace maps.** If considering Ring-LWE in R_q, where R is the ring of integers of a number field K, then any subfield $L \subseteq K$ gives rise to a subring $S \subseteq R$ (i.e., the ring of integers of L) and, modulo q, to a subring $S_q \subseteq R_q$. Then S_q is an $\mathbb{F}_q$-vector subspace of R_q, and R_q has a module structure over S_q. The dimensions of K over L, R over S and R_q over S_q agree.

There is a linear map $T := \mathrm{Tr}_{S_q}^{R_q} : R_q \to S_q$ satisfying the following relationship to the usual trace map from R to S:

$$\mathrm{Tr}_S^R(x) \bmod qS = \mathrm{Tr}_{S_q}^{R_q}(x \bmod qR).$$

To see this, remark that qS is elementwise fixed by the Galois group of K/L and qR is the extension of qS, so the Galois group takes qR to itself. Therefore the Galois group acts on R_q fixing S_q. Therefore we may define $\mathrm{Tr}_{S_q}^{R_q}(x)$ to be the sum of $\sigma(x)$ for σ in the Galois group of K/L, and the relationship above holds.

The ring R is always an S-module, but the reader is cautioned that in a general number field, R may not be a *free* module over S.

4.4. **Multiplicative cosets of subrings.** The set $a_0 S_q$, for any invertible $a_0 \in R_q$, is an $\mathbb{F}_q$-vector subspace of R_q of dimension equal to the dimension of S_q. Distinct such subspaces intersect only at subspaces of non-invertible elements of R_q, and R_q^* (the invertible elements of R_q) lie in the union of all such subspaces.

Let us write $A_{a_0 S_q, s, \chi}$ for the distribution on $a_0 S_q \times R_q$ given by choosing a uniformly in $a_0 S_q$, choosing e according to error distribution χ and outputting $(a, b := as + e)$.

Proposition 4.3. *If (a, b) is distributed according to $A_{a_0 S_q, s, \chi}$ where χ is invariant under multiplication by ζ, then*

(1) $(\zeta a, \zeta b)$ is distributed according to $A_{\zeta a_0 S_q, \zeta s, \chi}$, and
(2) $(a, \zeta b)$ is distributed according to $A_{a_0 S_q, \zeta s, \chi}$.

The multiplicative coset structure gives rise to another type of sample reduction, beyond ring homomorphism. We have

Proposition 4.4. *Suppose $s \in R_q$ is fixed. Define $T := Tr^{R_q}_{S_q}$, the trace map. Consider a collection of samples distributed according to $A_{a_0 S_q, s, \chi}$, where $a_0 \in R_q^*$ is fixed and satisfies $T(a_0) \neq 0$. Then T maps such samples to samples distributed according to $A_{s', T(\chi)}$ in S_q, where*

$$s' = \frac{T(a_0 s)}{T(a_0)}.$$

Proof. For $a = a_0 a' \in a_0 S_q$, since T is S_q-linear, we have

$$T(as) = a' T(a_0 s).$$

This implies that

$$(T(a), T(as + e)) = \left(a' T(a_0), a' T(a_0) \left(\frac{T(a_0 s)}{T(a_0)} \right) + T(e) \right)$$

This proves the proposition. $\square$ $\square$

4.5. **Trace maps for two-power cyclotomics.** The final piece to the puzzle is the behaviour of the trace map T in the previous section. In the case of the 2-power cyclotomics, the trace map is particularly well-behaved in terms of its effect on the error distribution. In fact, it takes very many of the basis elements ζ_m to zero. This is a feature of the orthogonality of the basis $1, \zeta_m, \ldots, \zeta_m^{n-1}$, and it may be proved with reference to basic algebraic number theory, as follows.

Using the notation of Section 4.3 in the case of the 2-power m-th cyclotomics K, let L be the k-th cyclotomic subfield. One may take $\zeta_k = \zeta_m^{m/k}$

and S_q has a basis $1, \zeta_k, \ldots, \zeta_k^{k/2-1}$ over $\mathbb{F}_q$. We collect terms to write

$$\begin{aligned} R_q &= \mathbb{Z} + \zeta_m \mathbb{Z} + \cdots + \zeta_m^{m/2-1}\mathbb{Z} \\ &= (\mathbb{Z} + \zeta_k \mathbb{Z} + \cdots + \zeta_k^{k/2-1}\mathbb{Z}) + \zeta_m(\mathbb{Z} + \zeta_k \mathbb{Z} + \cdots + \zeta_k^{k/2-1}\mathbb{Z}) \\ &\qquad\qquad + \cdots + \zeta_m^{m/k-1}(\mathbb{Z} + \zeta_k \mathbb{Z} + \cdots + \zeta_k^{k/2-1}\mathbb{Z}) \\ &= S_q + \zeta_m S_q + \cdots + \zeta_m^{m/k-1} S_q. \end{aligned}$$

In other words, R_q has a ζ-basis over S_q.

The elements of the Galois group of L/K are given by $\zeta_m \mapsto \zeta_m^a$ for $a \in (\mathbb{Z}/m\mathbb{Z})^*$ satisfying $a \equiv 1 \pmod{k}$, and so

$$\begin{aligned} \mathrm{Tr}_{S_q}^{R_q}(\zeta_m^i) &= \sum_{\substack{0 \le a < m \\ a \equiv 1 \pmod{k}}} \zeta_m^{ia} \\ &= \zeta_m^i \sum_{a=0}^{m/k-1} \zeta_m^{iak} \\ &= \begin{cases} 0 & i \not\equiv 0 \pmod{\frac{m}{k}} \\ \frac{m}{k}\zeta_m^i & i \equiv 0 \pmod{\frac{m}{k}} \end{cases} . \end{aligned}$$

In particular, for the trace to the index two subfield, we have:

$$\frac{1}{2} T_{S_q}^{R_q}(\zeta_m^i) = \begin{cases} \zeta_m^i & i \equiv 0 \pmod{2} \\ 0 & i \equiv 1 \pmod{2} \end{cases} .$$

This special case can be seen directly by observing that if i is even, then $\zeta_m^i \in S$, while if i is odd, then ζ_m^i is the square root of something in S, i.e. it satisfies the minimal polynomial $x^2 - \zeta_m^{2i}$, and hence has trace zero. An alternate proof of the general case then follows by application of the special case $\log_2(m/k)$ times.

In summary then, the trace map is much less problematic when applied to error distributions than a ring homomorphism is. The following proposition is now immediate.

Proposition 4.5. *Let R be the ring of integers of the m-th cyclotomics, where m is a power of two. Let S be the subring of integers of the k-th cyclotomics (hence k is also a power of two). Write $T := Tr_{S_q}^{R_q}$ for the trace map. Suppose that χ is an error distribution formed on the ζ-basis of R_q with coefficients chosen according to χ_0. Then $\frac{k}{m}T$ takes values in S_q and $\frac{k}{m}T(\chi)$ is the error distribution formed on the ζ-basis of S_q with coefficients from χ_0.*

The efficacy of the trace map with respect to the error distribution is due to its being an orthogonal projection to the space spanned by a subset of an orthonormal basis.

5. Reducing to a smaller ring

We demonstrate that if one can find sufficiently many samples whose a values are restricted to a fixed multiplicative coset of a subring, then we can reduce the Ring-LWE problem to multiple independent Ring-LWE instances in the subring, without error inflation.

For this section, we are in the two-power cyclotomic case. Let R be the ring of m-th cyclotomic integers, where m is a power of two (which have dimension n, where $m = 2n$) and S be the ring of k-th cyclotomic integers, where $k \mid m$. Then we have an extension of rings, $S \subseteq R$ of degree m/k. Suppose that the rational prime q is unramified in R.

Proposition 5.1. *Consider a Ring-LWE instance in R_q with secret s and error distribution χ. Let $a_0 \in R_q$ be a fixed invertible element. Let $T := \mathrm{Tr}^{R_q}_{S_q}$, and suppose that $T(a_0)$ is invertible.*

Let i be an integer. Then in time polynomial in n and $\log q$, one can reduce a Ring-LWE sample from distribution $A_{a_0 S_q, s, \chi}$ to a Ring-LWE sample in S_q drawn according to secret

$$\frac{T(a_0 \zeta^i s)}{T(a_0)}$$

and error distribution $\frac{k}{m} T(\zeta^i \chi) \subseteq S_q$.

In particular, by Proposition 4.5, coefficient distributions of a ζ-invariant χ and its resulting distribution $\frac{d}{n} T(\zeta^i \chi)$ are of the same size; it is in this sense that the errors do not inflate.

Proof. Consider the sample (a, b) where $b = as + e$. Multiplying the second coordinate of the sample by ζ^i and taking the trace $\frac{k}{m} T$, we obtain as in Proposition 4.4, a sample

$$\begin{aligned}
&\left(\frac{k}{m} T(a), \frac{k}{m} T(\zeta^i b)\right) \\
&\left(\frac{k}{m} T(a), \frac{k}{m} T(a \zeta^i s + \zeta^i e)\right) \\
&= \left(a' \frac{k}{m} T(a_0), a' \frac{k}{m} T(a_0) \cdot \left(\frac{T(a_0 \zeta^i s)}{T(a_0)}\right) + \frac{k}{m} T(\zeta^i e)\right),
\end{aligned}$$

where $a' := a a_0^{-1} \in S_q$.

Multiplication in the ring, and taking the trace, are polynomial in the ring size. $\square$

The following is the main theorem of the paper.

Theorem 5.2. *Consider a Ring-LWE instance in R_q with secret s and error distribution χ which is invariant under multiplication by ζ. Let $a_0 \in R_q$ be a fixed invertible element. Let $T := \mathrm{Tr}^{R_q}_{S_q}$, and suppose that $T(a_0)$ is invertible.*

Suppose one obtains N samples (a, b) distributed according to $A_{a_0 S_q, s, \chi}$. Then in time linear in the number of samples N, and polynomial in n and $\log q$, one can reduce the computation of the secret $s \in R_q$ to the solution of m/k Search Ring-LWE problems in S_q with error distribution $\frac{k}{m}T(\chi)$, having N samples each.

If χ is formed on the ζ-basis from coefficient distribution χ_0 on $\mathbb{F}_q$, then so is $\frac{k}{m}T(\chi)$.

Proof. Set $i = j$ in Proposition 5.1 for each j in the range $j = 0, \ldots, m/k-1$, to obtain N samples having secret

$$c_j := \frac{T(a_0 \zeta^j s)}{T(a_0)}.$$

Using an oracle that solves Search Ring-LWE in S_q, obtain c_j.

Collecting all the values c_j, we have a linear system of m/k equations over S_q, whose indeterminates are the coefficients of s (expressed in terms of a basis for R_q over S_q), of the form

$$T(a_0 \zeta^j s) = c_j T(a_0), \quad j = 0, \ldots, m/k - 1.$$

The linear equations are independent provided that $\{a_0 \zeta^j\}$ is a set of S_q-independent vectors in R_q. We saw above that $\{\zeta^j\}_{j=0,\ldots,m/k-1}$ is a basis for R_q over S_q. Thus independence is guaranteed by the fact that a_0 is invertible. Note that we can consider this system to consist of n independent linear equations over $\mathbb{F}_q$. The system can be solved by Gaussian elimination to recover s.

All the field operations concerned are polynomial in the size of the ring. We must apply the trace to N samples m/k times, and we must solve Gaussian elimination of dimension $n = m/2$ over $\mathbb{F}_q$, which is polynomial in m and $\log q$. $\square$ $\square$

As a small corollary, note that in any small Ring-LWE situation where exhaustive search may apply, it is equally possible to use the above for a square-root speedup, provided many samples are available. As an example, if we have a coefficient distribution with support not including all of $\mathbb{F}_q$, then the following statement demonstrates the approach.

Corollary 5.3. *Consider a Ring-LWE instance in R_q with secret s and error distribution χ formed on a ζ-basis with coefficient distribution having support strictly smaller than $\mathbb{F}_q$.*

There is an algorithm to solve this problem, with success probability $1/2$, in time and number of samples $q^{n/2}$ times factors polynomial in $n \log q$, using space polynomial in $n \log q$.

Proof. Note that the hypotheses guarantee χ is invariant under multiplication by ζ. Let S_q be the ring of index two in R_q (i.e. n-th roots of unity). Collect samples, discarding all but those with $a \in S_q$. In time $O(Nq^{n/2})$ we can accumulate N samples with $a \in S_q$. Apply Theorem 5.2 to reduce to two Ring-LWE problems in S_q with N samples each. The error distribution χ on R_q gives an error distribution $\frac{1}{2}T^{R_q}_{S_q}(\chi)$ on S_q. If χ is formed on a ζ-basis with coefficients supported in $E_\chi \subsetneq \mathbb{F}_q$, then $\frac{1}{2}T^{R_q}_{S_q}(\chi)$ is formed on a ζ-basis with coefficients supported in $E_\chi \subsetneq \mathbb{F}_q$. Therefore, if the number of samples is sufficient, the reduced Ring-LWE problems are solvable using exhaustive search through possible s values.

In our case, we need N large enough so that a Ring-LWE problem in S_q with N samples has a unique solutions with probability $1/\sqrt{2}$. Although N depends upon $|E_\chi|$, for the worst case $|E_\chi| = q-1$, N is still polynomial in $n \log q$. Solve the reduced problems by exhaustive search, which takes time $O(q^{n/2})$ and each succeeds with probability $1/\sqrt{2}$. □ □

6. Background on the Blum-Kalai-Wasserman algorithm

First, we will give a very brief overview of the Blum-Kalai-Wasserman (BKW) algorithm in the context of LWE [3]. It is a combinatorial algorithm in which samples are collected and stored so as to facilitate the creation of new samples, as iterated sums and differences of established ones. The goal is to create new samples for which a is restricted to a linear subspace. This is the *reduction phase* of the full BKW algorithm.

In BKW, after reduction, there is a *hypothesis testing phase*, in which one solves a lower-dimensional Ring-LWE problem (that given by restricting a to the subspace) by exhaustive search over possible secrets. And then there is a *back-substitution phase*, where the small piece of the secret recovered in hypothesis testing is used to rework the problem to prepare the next small piece for hypothesis testing.

One can think of BKW as a sort of controlled Gaussian elimination on a matrix whose rows are samples, in which one wants to obtain as much simplification as possible using just one sum or difference of rows. By keeping the coefficients of the linear combinations small, we prevent the error 'blow-up' that occurs with regular Gaussian elimination. The cost is in needing many more matrix rows (samples) in order to be able to choose good linear combinations. The back-substitution phase is analogous to the eponymous phase of Gaussian elimination, with the recovered portion of the secret taking the role of the free variable. From another point of view, BKW reduction is a sort of iterated birthday attack, in which one searches for and exploits collisions which eliminate entries of the vectors, reducing to a subspace, where one searches again for collisions, and so on.

Now let us be more precise. During the reduction phase, only the a-value of a sample matters, considered as a vector in a vector space V, and the goal is to create samples with $a \in W$, a linear subspace of V. Suppose, for the sake of explanation, that W is defined by the first r coefficients of its vectors being 0. One generates an ordered list of the first r entries of all the vectors a which are observed. Whenever a new vector a is observed, it is compared to the ordered list. If it is not already present, it is added. Otherwise, we have discovered two samples (a, b) and (a', b') for which $(a - a', b - b')$ is a new sample for which $a - a'$ lies in W. The penalty is that the error distribution of these new samples is widened. We begin a new table of such vectors as they are generated. In this way, we produce a large number of samples in a smaller subspace at the cost of inflating the error widths.

Instead of performing this reduction all at once, one chooses an appropriate *block size* B for BKW (which is fixed throughout in the naïve implementation), which is to say, the codimension of W as a subspace of V. Once we have produced enough samples in W, we can use these to perform another BKW reduction to a subspace $W' \subseteq W$ of codimension B in W. The cost of a reduction step is exponential in B, so we keep B as small as possible. We perform block reductions until the samples are all taken from a small enough subspace to run an exhaustive search or other strategy to finish off the problem. The limiting factor on shrinking B is an upper limit on the number of blocks used overall. Each reduction into codimension B has a cost in error-inflation. We have a limit on the total error inflation (because hypothesis testing will fail if the error is so inflated as to appear uniform), which limits the total number of blocks.

The BKW algorithm has been improved in recent years, including using coding theory to reduce the number of values that need to be stored and compared, sieving at each step, allowing the block size to vary, using the Fourier transform to speed up hypothesis testing; see [1] [12] [15] [16] [17] [18].

7. Reduction using BKW

In this section, we address the problem of finding sufficiently many samples (a, b) having a from an appropriate subring $S_q \subseteq R_q$, so that Theorem 5.2 will apply. For this, we use the reduction phase of the BKW algorithm. We emphasize that it is possible, once the samples have been given in an appropriate basis, to use an off-the-shelf BKW reduction algorithm, including coded BKW with sieving etc., for the reduction phase. Then, Theorem 5.2, which is polynomial time, replaces all the other phases of BKW.

The only adaptor necessary to connect BKW to Theorem 5.2 is an attention to the basis used. In order to perform the reduction, we begin with the ζ-basis of R_q over $\mathbb{F}_q$, namely

$$1, \zeta, \zeta^2, \ldots, \zeta^{n-1},$$

and then reorder it to produce a *prioritized basis.* The most important property for our purposes is that:

(1) if one of ζ^i and ζ^j has lower multiplicative order than the other, then it comes later than the other

One computationally convient way to accomplish this is to take the bit-reversal permutation on n elements (i.e. a maps to b if the binary representation of a in $\log_2(n)$ bits, read backwards, is b), then reserve the order. For concreteness, the prioritized basis (in part) is as follows:

$$\zeta_m^{n-1}, \zeta_m^{\frac{n}{2}-1}, \zeta_m^{\frac{3n}{4}-1}, \zeta_m^{\frac{n}{4}-1}, \ldots, \zeta_m^{\frac{3n}{4}}, \zeta_m^{\frac{n}{4}}, \zeta_m^{\frac{n}{2}}, 1.$$

Using any type of BKW reduction, one now reduces, with respect to this basis. To be precise, one seeks to eliminate the earlier coefficients of the elements a, as expressed in this basis. At the end, at most the last 2^k coefficients are non-zero, for some small k. For example, one may reduce until only the last 1, 2, 4 or 8 coefficients are possibly non-zero. The varying block sizes during the reduction algorithm itself need not respect any restrictions, and improvements such as coded-BKW with sieving, may be used.

After reduction, one has obtained samples with $a \in S_q$ for S of dimension 2^k. One then applies Theorem 5.2.

8. The Ring-BKW algorithm

In this section we summarize the Ring-BKW algorithm for completeness. In short, one uses an off-the-shelf BKW reduction algorithm on samples with respect to a particular choice of basis, then applies Theorem 5.2. The important point is that the back-substitution phase of BKW is no longer needed, and the hypothesis-testing phase can be parallelized. The hypothesis-testing phase can also be off-the-shelf, including recent improvements using the Fourier transform etc. [12]. However, we will elaborate somewhat.

Ring-BKW algorithm. Choose a subring $S \subseteq R$ of dimension B over $\mathbb{Z}$ (corresponding to a lower-degree 2-power cyclotomic field), to which we wish to reduce. Define R_q and S_q as before. The Ring-BKW Algorithm is given as Algorithm 1.

Algorithm 1 Ring-BKW Algorithm

(1) Run BKW Reduction (as in Section 7 above with prioritized basis) on the values a until all samples (a, b) have $a \in S_q$.
(2) Use Theorem 5.2 to create samples from n/B different Ring-LWE problems in S_q.
(3) Solve these Ring-LWE problems using any method of choice.
(4) Use Theorem 5.2 to recover the secret s in polynomial time from these solutions.

However, any reduction algorithm to obtain values $a \in S_q$ will do as well. The following theorem relates any reduction algorithm to the solution of Search Ring-LWE. For the following, we consider Gaussian error with a well-defined width; an *expansion factor* refers to a multiplicative factor on the width.

Theorem 8.1. *Suppose that $\mathcal{B}$ is an algorithm which, given a Ring-LWE problem of dimension n over $\mathbb{F}_q$, produces N Ring-LWE samples of dimension B with error expansion factor of f, in time $t_{\mathcal{B}}(n, B, f, N)$, and using $r_{\mathcal{B}}(n, B, f, N)$ original samples.*

Suppose that $\mathcal{R}$ is an algorithm which solves Ring-LWE in dimension B over $\mathbb{F}_q$ in time $t_{\mathcal{R}}(B)$, given error width less than or equal to s and at least N samples.

Then, there is an algorithm $\mathcal{A}$ which solves Ring-LWE in R_q having width σ in time

$$t_{\mathcal{B}}(n, B, s/\sigma, N) + \frac{n}{B} t_{\mathcal{R}}(B) + N \cdot (\textit{time polynomial in } n \log q),$$

using $r_{\mathcal{B}}(n, B, s/\sigma, N)$ samples.

Proof. Algorithm 1 will succeed provided $f \leq s/\sigma$, so that the reduced Ring-LWE problems are solvable. Therefore the time to run the reduction phase is $t_{\mathcal{B}}(n, B, s/\sigma, N)$. The time to create the smaller Ring-LWE problems is linear in N and polynomial in $n \log q$ from Theorem 5.2. Solving the $\frac{n}{B}$ smaller Ring-LWE problems takes time $t_{\mathcal{R}}(B)$ each. Then reconstructing the secret (as in Theorem 5.2) again takes polynomial time. □ □

9. Advanced Keying Speedup

In the previous section, one uses BKW on LWE to perform reduction. Given a Ring-LWE sample, there are in fact n rotated samples one could feed into the reduction:

$$(a, b), (\zeta a, \zeta b), \ldots, (\zeta^{n-1} a, \zeta^{n-1} b).$$

Naïvely, one may include them all, or include the first one. Probably the best course of action is to include them all, to increase the number of collisions located amongst the available samples (since the number of samples needed is the downside to BKW in general). By including all rotations, one catches all collisions of the form $a_1 \pm \zeta^k a_2 = 0$. These are all perfectly useful collisions for the algorithm, if the error term is ζ-invariant. In this section we propose a space-saving approach based on symmetries, which is equivalent, in terms of collisions obtained per sample, to storing all rotations of the samples.

Our proposal in this section is an analogue of the space-saving technique used in regular BKW, wherein for each sample (a, b) we may derive two samples (a, b) and $(-a, -b)$: we choose one canonically (where the first coefficient of a is in $\left\{0, \ldots, \frac{q-1}{2}\right\}$, say), and save only this one. By doing so, we will catch all collisions between samples where their sum *or their difference* vanishes, and save half the table space in the process, since one need only store vectors whose first entry is chosen amongst $\frac{q+1}{2}$ possibilities, instead of q.

The fundamental observation is that the prioritized basis proposed in the last section is particularly well-suited to this type of strategy, because of its relationship with the 'negacyclic permutation' effect of multiplication by ζ. Consider the values

$$a, \zeta a, \zeta^2 a, \ldots, \zeta^{n-1} a.$$

Write $\mathbf{a}_i$ for the vector of coefficients of $\zeta^i a$ in the prioritized basis (which are elements of $\mathbb{F}_q$). Write $|\mathbf{a}_i|$ for the same list, but with each coefficient/entry replaced with the absolute value of its minimal representative modulo q (minimal representative means representative closest to zero). In particular, there are $\frac{q+1}{2}$ possible values for each entry.

Consider a general element a (i.e. consider the coefficients to be distinct variables) and write $|\mathbf{a}_i|_j$ for the j-th entry of $|\mathbf{a}_i|$. Let B divide n, a power of 2. Thus this improvement, at least in the description provided here, depends somewhat sensitively on block size. There are several useful observations in this case:

(1) The entries of $|\mathbf{a}_i|$ consist of a permutation of the entries of $|\mathbf{a}| = |\mathbf{a}_0|$.
(2) The blocks of length B, $|\mathbf{a}_i|_{0\leq j\leq B-1}$, for $i = 0, \ldots, n/B$ consist of disjoint subsets of the coefficients of $|\mathbf{a}|$.
(3) If $i \equiv i' \pmod{n/B}$, then the sequence $|\mathbf{a}_{i'}|_{0\leq j\leq B-1}$ is a permutation of the sequence $|\mathbf{a}_i|_{0\leq j\leq B-1}$.

(Similar statements apply to subsequent blocks of length B also.)

In view of these facts, one has the following strategy. Suppose one is searching for collisions amongst the initial blocks of length B. Then, provided one has a means of choosing a canonical representative amongst the blocks $|\mathbf{a}_i|_{0\leq j\leq B-1}$ for $i \equiv i_0 \pmod{n/B}$ (which are all reorderings of one another), one need only store one of these B different rotations. Hence the table size is reduced from $\frac{q^B-1}{2}$ (i.e. storing entries of length B with first non-zero entry chosen to be 'positive', equivalently in the range $0, \ldots, \frac{q-1}{2}$) to at most

$$\frac{q^B}{2B} + \frac{n^2 q^{n/B}}{2B},$$

obtained by storing the canonical choice amongst the B choices of $i \equiv i_0 \pmod{n/B}$, with the same convention on first entry.

A possible canonical choice is the ordering which has smallest first entry, together with some tie-breaking conventions, e.g. smallest second entry, etc., and if all entries are equal, then some appropriate convention on sign changes between $\mathbf{a}$ and $|\mathbf{a}|$, etc. It is not possible to break a tie if the first B entries of $\mathbf{a}_i$ and $\mathbf{a}_j$ actually agree under one of the rotations, which accounts for the correction factor. To make this precise, observe that what we are actually checking for is something in the kernel of $Tr^{R_q}_{S_q}\left((\zeta^i - \zeta^j)a\right)$ for $S \subset R$ of dimension B, and $i \not\equiv j$ considered modulo n/B. Each such kernel has size $q^{n/B}$. For any of these $\frac{n}{2B}\left(\frac{n}{B} - 1\right) q^{n/B}$ possibly problematic examples, we may wish to include all B rotations at most.

This method provides some advantage in table size if there is at least one sample for which we can store fewer than all possible rotations. However, the loose bound above indicates that the advantage is asymptotically a table-size factor of $1/B$ provided that $B > n/B$.

10. In practice

It is evident that the runtime of Ring-BKW is expected to be better than that of standard BKW (in any of its current forms), since the reduction and hypothesis testing phases may be taken to be the same, but the backsubstitution phase is no longer required. Furthermore, the smaller Ring-LWE problems can be solved in parallel.

As a potential further improvement, advanced keying provides a table size reduction factor which is applicable when certain block sizes are powers of two. This may or may not be useful or extendable in view of the changing block sizes sometimes employed in BKW reduction.

The Ring-LWE Challenges [11] are in the form of *Tweaked Ring-LWE*, which refers to dual Ring-LWE transfered to the unital version (see [11, §2.3]), so that the parameter assumptions in this paper apply to the two-power cyclotomic challenges included therein. It would be very interesting to test these algorithms on those parameters, but it is beyond the scope of this paper.

References

[1] Albrecht, M.R., Cid, C., Faugère, J.C., Fitzpatrick, R., Perret, L.: On the complexity of the BKW algorithm on LWE. Des. Codes Cryptogr. 74(2), 325–354 (2015), https://doi.org/10.1007/s10623-013-9864-x

[2] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange—a new hope. In: 25th USENIX Security Symposium (USENIX Security 16). pp. 327–343. USENIX Association, Austin, TX (2016), https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim

[3] Blum, A., Kalai, A., Wasserman, H.: Noise-tolerant learning, the parity problem, and the statistical query model. J. ACM 50(4), 506–519 (2003), https://doi.org/10.1145/792538.792543

[4] Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors (extended abstract). In: STOC'13—Proceedings of the 2013 ACM Symposium on Theory of Computing, pp. 575–584. ACM, New York (2013), https://doi.org/10.1145/2488608.2488680

[5] Brakerski, Z., Perlman, R.: Order-lwe and the hardness of ring-lwe with entropic secrets. Cryptology ePrint Archive, Report 2018/494 (2018), https://eprint.iacr.org/2018/494

[6] Brakerski, Z., Vaikuntanathan, V.: Fully homomorphic encryption from ring-LWE and security for key dependent messages. In: Advances in cryptology—CRYPTO 2011, Lecture Notes in Comput. Sci., vol. 6841, pp. 505–524. Springer, Heidelberg (2011), https://doi.org/10.1007/978-3-642-22792-9_29

[7] Brakerski, Z., Vaikuntanathan, V.: Efficient fully homomorphic encryption from (standard) LWE. SIAM J. Comput. 43(2), 831–871 (2014), https://doi.org/10.1137/120868669

[8] Castryck, W., Iliashenko, I., Vercauteren, F.: Provably weak instances of ring-LWE revisited. In: Advances in cryptology—EUROCRYPT 2016. Part I, Lecture Notes in Comput. Sci., vol. 9665, pp. 147–167. Springer, Berlin (2016), https://doi.org/10.1007/978-3-662-49890-3_6

[9] Chen, H., Lauter, K., Stange, K.E.: Attacks on the search RLWE problem with small errors. SIAM J. Appl. Algebra Geom. 1(1), 665–682 (2017), https://doi.org/10.1137/16M1096566

[10] Chen, H., Lauter, K., Stange, K.E.: Security considerations for Galois non-dual RLWE families. In: Selected areas in cryptography—SAC 2016, Lecture Notes in Comput. Sci., vol. 10532, pp. 443–462. Springer, Cham (2017)

[11] Crockett, E., Peikert, C.: Challenges for ring-lwe. Cryptology ePrint Archive, Report 2016/782 (2016), https://eprint.iacr.org/2016/782
[12] Duc, A., Tramèr, F., Vaudenay, S.: Better algorithms for LWE and LWR. In: Advances in cryptology—EUROCRYPT 2015. Part I, Lecture Notes in Comput. Sci., vol. 9056, pp. 173–202. Springer, Heidelberg (2015), https://doi.org/10.1007/978-3-662-46800-5_8
[13] Eisenträger, K., Hallgren, S., Lauter, K.: Weak instances of plwe. In: Selected Areas in Cryptography–SAC 2014, pp. 183–194. Springer (2014)
[14] Elias, Y., Lauter, K., Ozman, E., Stange, K.: Provably weak instances of ring-lwe. In: Advances in Cryptology – CRYPTO 2015, Lecture Notes in Comput. Sci., vol. 9215, pp. 63–92. Springer, Heidelberg (2015)
[15] Guo, Q., Johansson, T., Må rtensson, E., Stankovski, P.: Coded-BKW with sieving. In: Advances in cryptology—ASIACRYPT 2017. Part I, Lecture Notes in Comput. Sci., vol. 10624, pp. 323–346. Springer, Cham (2017)
[16] Guo, Q., Johansson, T., Mårtensson, E., Stankovski, P.: On the asymptotics of solving the lwe problem using coded-bkw with sieving. IEEE Transactions on Information Theory (2019)
[17] Guo, Q., Johansson, T., Stankovski, P.: Coded-BKW: solving LWE using lattice codes. In: Advances in cryptology—CRYPTO 2015. Part I, Lecture Notes in Comput. Sci., vol. 9215, pp. 23–42. Springer, Heidelberg (2015), https://doi.org/10.1007/978-3-662-47989-6_2
[18] Kirchner, P., Fouque, P.A.: An improved BKW algorithm for LWE with applications to cryptography and lattices. In: Advances in cryptology—CRYPTO 2015. Part I, Lecture Notes in Comput. Sci., vol. 9215, pp. 43–62. Springer, Heidelberg (2015), https://doi.org/10.1007/978-3-662-47989-6_3
[19] Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Advances in cryptology—EUROCRYPT 2010, Lecture Notes in Comput. Sci., vol. 6110, pp. 1–23. Springer, Berlin (2010), https://doi.org/10.1007/978-3-642-13190-5_1
[20] Lyubashevsky, V., Peikert, C., Regev, O.: A toolkit for ring-LWE cryptography. In: Advances in cryptology—EUROCRYPT 2013, Lecture Notes in Comput. Sci., vol. 7881, pp. 35–54. Springer, Heidelberg (2013), https://doi.org/10.1007/978-3-642-38348-9_3
[21] Peikert, C.: How (not) to instantiate ring-LWE. In: Security and cryptography for networks, Lecture Notes in Comput. Sci., vol. 9841, pp. 411–430. Springer, [Cham] (2016), https://doi.org/10.1007/978-3-319-44618-9_22
[22] Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. In: STOC'05: Proceedings of the 37th Annual ACM Symposium on Theory of Computing, pp. 84–93. ACM, New York (2005), https://doi.org/10.1145/1060590.1060603

Department of Mathematics, University of Colorado, Campux Box 395, Boulder, Colorado 80309-0395

E-mail address: kstange@math.colorado.edu