

# 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. We also present a method to exploit symmetry to reduce table sizes, samples needed, and runtime during the reduction phase. 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) [24] [25], and Learning with Errors (LWE) [27] 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 [8]. 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 [9] [10] [11] [15] [16], or the secret is chosen from a subring or other ring-related non-uniform distribution [7]. In the related NTRU cryptosystem, the norm and trace maps to subfields play a role in attacks [1, 12, 17, 23].

---

*Date:* July 2, 2021.

*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.

However, the best known attacks on Ring-LWE parameters suggested for implementation are still generic attacks for LWE, e.g. [3]. The Blum-Kalai-Wasserman (BKW) algorithm is one 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 [5]. BKW has the drawback of requiring exponentially many samples, unless sample amplification is used [21]. Nevertheless, its performance has been of significant interest: for analysis and recent improvements, see [2] [14] [18] [19] [20] [22]. (Note that sample amplification does not immediately transfer from the LWE to the Ring-LWE setting, at least if one wishes the amplified samples to have Ring-LWE, and not just underlying LWE, format; the analogue would be the ‘sample rotation’ described below.)

This paper focuses on two-power-cyclotomic unital (but equivalently, dual [13] [26]) 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 the 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 (Theorem 5.2).* State-of-the-art off-the-shelf code for the BKW reduction phase and hypothesis testing phase may be used. Note that the reduction phase of BKW is the dominant phase for runtime, and hypothesis testing is typically polynomial, but the now-eliminated back-substitution phase runs in time which is also exponential, but differs only by a smaller polynomial factor from the reduction phase; hence the overall runtime savings is a polynomial factor. 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. Compared to a BKW reduction phase completely blind to the ring structure, this reduces the table size and samples needed by a factor of the block size, as well as reducing runtime, but requires that block sizes be taken to be a (possibly varying) power of 2.

We also discuss a square-root speedup over exhaustive search (which may be used, for example, in hypothesis testing); see Corollary 5.3.

See Section 10 for more discussion of practical runtime.

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 [10] [11] [15] [16]).
- (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 (previously used in lattice sieving [6] [28]; more generally, manipulation of samples by multiplication was exploited in [4]).
- (3) The existence of subrings as linear subspaces (which is important in [7]).
- (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, [7]; 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,

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

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, and (1) is chosen in such a way that the isomorphism is exactly that arising from the *Minkowski* or *canonical* embedding of algebraic number theory. We can also denote the norm by  $\|\mathbf{x}\| = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ .

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

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

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  $\sqrt{\pi/2} \sim 1.25\dots$  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 ideal generated by  $q$  in  $R$ . 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  $R$ , 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 in the cryptographic setting), 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  $\psi$  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  $\psi$ , and outputting  $(a, b := as + e)$ .

Such outputs  $(a, b)$  are called *samples*, and in a cryptographic 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 samples 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  $\psi$  and secret distribution  $\varphi$ , is as follows: Given an error distribution  $\psi$  over  $R_q$  and a secret distribution  $\varphi$  over  $R_q$ , and some number of samples drawn from the distribution  $A_{s,\psi}$  for some fixed  $s$  drawn from  $\varphi$ , compute  $s$ .

**Definition 2.2.** The *decisional Ring-LWE problem*, for error distribution  $\psi$  and secret distribution  $\varphi$ , is as follows: Given an error distribution  $\psi$  over  $R_q$  and a secret distribution  $\varphi$  over  $R_q$ , distinguish with non-negligible advantage, between

- (1) samples drawn from the distribution  $A_{s,\psi}$  for some fixed  $s$  drawn from  $\varphi$ ; 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 inverse of the different ideal). However, in the case that  $K$  is a  $2^N$ -th cyclotomic field,  $R \cong 2^{N-1}R^\vee$  and this isomorphism is realized as a scaling in the canonical embedding, and thus preserves the error distribution up to scaling, so we can interchange the dual version with the simpler ‘unital’ version considered here [24].

Search-to-decision reductions are known in a variety of contexts [24]. 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), or by multiplication by  $a^{-1}$  in

the ring. By introducing a small error  $e$ , so we have  $(a, as+e)$ , multiplication by  $a^{-1}$  is no longer helpful, and 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 will now specialize to the 2-power cyclotomic case, fixing values for the variables

$$K, R, q, R_q, n$$

from the last section, and defining

$$m, \zeta := \zeta_m, \chi, \chi_0, E_{\chi_0}$$

for the 2-power cyclotomic case. Whenever we say refer to *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  $\zeta_m$  for a primitive  $m$ -th root of unity in  $R$  and for its image in quotients of this ring.

**3.2. The  $\zeta$ -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  *$\zeta$ -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  $\zeta_m$ ), but the minimal polynomial for  $\zeta_m$  varies in these quotients, and may be a proper divisor of this cyclotomic polynomial. Nevertheless, in all quotients of  $R$ , we still obtain a  $\zeta$ -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  $\zeta$ -basis for this ring (to be explicit, the images of the  $\zeta$  basis for  $R$  under the reduction modulo  $q$ ).

We may also consider further quotients  $R/\mathfrak{a}$  for  $\mathfrak{a} \mid qR$ . We may also use a  $\zeta$ -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  $\chi$ , coefficient distribution  $\chi_0$  and coefficient support  $E_{\chi_0}$ .** We will denote the error distribution by  $\chi$ . If this error distribution is formed using independent identically distributed coefficients on the  $\zeta$ -basis, with coefficient distribution  $\chi_0$  supported on a subset  $E_{\chi_0} \subseteq \mathbb{F}_q$ , then we say that  $\chi$  is *formed on the  $\zeta$ -basis with coefficients distributed according to  $\chi_0$* . 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  $\zeta_m$  is orthonormal (after scaling) in the canonical embedding*. To see this, use (1) and observe that if  $\zeta_m^a$  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. 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  $R_q$ . Note, however, that the method of [4, Section 3, Targeting  $e_i$ ] could be used to manipulate the samples so the secret can be taken from the error distribution, preserving the Ring-LWE structure of the samples.

#### 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

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

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

In general, the effect of  $\rho$  on  $\chi$  is problematic, i.e. it spreads out the error widely. As an illustration, we give a proposition governing the behaviour of  $\rho$  on  $\chi$  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  $\chi$  is formed on the  $\zeta$ -basis in  $R_q$  with coefficients drawn from  $\chi_0$  on  $\mathbb{F}_q$ , then  $\chi' := \rho(\chi)$  is formed on the  $\zeta$ -basis in  $\mathbb{F}_{q^k}$  with coefficients drawn from  $\chi'_0$  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  $r = \text{ord}_2(q-1)$ , 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  $\text{ord}_2(q^2 - 1) = \text{ord}_2((q-1)(q+1)) = r+1$  and, by induction

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

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  $\chi'$  are chosen from a sum of two Gaussian distributions with different coefficients. This is less controlled than twice a single Gaussian. For, twice a Gaussian is simply a wider Gaussian, and the size of its support grows by approximately  $\sqrt{2}$ . However, in an uneven linear combination the size of the support  $E_{\chi'}$  is approximately the square of the size of  $E_{\chi}$ . (To be explicit, since  $\chi_0$  is discrete,  $c\chi_0$  is “spaced out” into isolated spikes, and each spike of support is transformed into a small gaussian by the addition of  $\chi_0$  to form  $c\chi_0 + \chi_0$ .) This is a symptom of the protective property of these ring homomorphisms: they transform the error to something less amenable to attack. In fact, very quickly the image of a Gaussian error approaches uniform in the image ring as the dimension of the image ring decreases. And Ring-LWE samples with uniform error are informationless.

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

**Proposition 4.2.** *Suppose  $\chi$  is invariant under multiplication by  $\zeta$ . 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  $\zeta_m$  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. (This may allow for adapting the notion of sample amplification to the Ring-LWE case; see [21].)

**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  $T := \text{Tr}_{S_q}^{R_q} : R_q \rightarrow S_q$  satisfying the following relationship to the usual trace map from  $R$  to  $S$ :

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

To see this, remark that  $qS$  is elementwise fixed by the Galois group of  $K/L$  and  $qR$  is the extension of  $qS$  to  $R$ , so the Galois group takes  $qR$  to itself. Therefore the Galois group acts on  $R_q$  fixing  $S_q$ . Therefore we may define  $\text{Tr}_{S_q}^{R_q}(x)$  to be the sum of  $\sigma(x)$  for  $\sigma$  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 consisting 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  $\chi$  and outputting  $(a, b := as + e)$ .

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

- (1)  $(\zeta a, \zeta b)$  is distributed according to  $A_{\zeta a_0 S_q, s, \chi}$ , 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  $T := \text{Tr}_{S_q}^{R_q}$ , the trace map described above. Consider a collection of samples distributed according*

to  $A_{a_0 S_q, s, \chi}$ , where  $a_0 \in R_q^*$  is fixed and  $T(a_0)$  is invertible. 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.  $\square$

**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  $\zeta_m$  to zero. This is a feature of the orthogonality of the basis  $1, \zeta_m, \dots, \zeta_m^{n-1}$ , 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  $1, \zeta_k, \dots, \zeta_k^{k/2-1}$  over  $\mathbb{F}_q$ . We collect terms to write

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

In other words,  $R_q$  has a  $\zeta$ -basis over  $S_q$ .

The elements of the Galois group of  $K/L$  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

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

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

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

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  $\zeta_m^i$  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 preserves the error distribution up to small factors. The following proposition, which is now immediate, makes this explicit.

**Proposition 4.5.** *Suppose we are in the two-power cyclotomic case as in Section 3, where in particular  $R$  is the ring of integers of the  $m$ -th cyclotomics, with  $m$  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  $T := \text{Tr}_{S_q}^{R_q}$  for the trace map described in Section 4.3. Suppose that  $\chi$  is an error distribution formed on the  $\zeta$ -basis of  $R_q$  with coefficients chosen according to  $\chi_0$ . Then  $\frac{k}{m}T$  takes values in  $S_q$  and  $\frac{k}{m}T(\chi)$  is the error distribution formed on the  $\zeta$ -basis of  $S_q$  with coefficients from  $\chi_0$ .*

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  $\chi$ . Let  $a_0 \in R_q$  be a fixed invertible element. Let  $T := \text{Tr}_{S_q}^{R_q}$ , 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  $\zeta$ -invariant  $\chi$  and its resulting distribution  $\frac{k}{m}T(\zeta^i \chi)$  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  $\zeta^i$  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  $a' := aa_0^{-1} \in S_q$ .

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.** *Suppose  $R$  is the ring of  $m$ -th cyclotomic integers, for  $m = 2n$  a power of two, and  $S$  is the ring of  $k$ -th cyclotomic integers, where  $k \mid m$ , so the extension  $S \subseteq R$  is of degree  $m/k$ . Suppose that the rational prime  $q$  is unramified in  $R$ .*

*Consider a Ring-LWE instance in  $R_q$  with secret  $s$  and error distribution  $\chi$  which is invariant under multiplication by  $\zeta = \zeta_m$ , a primitive  $m$ -th root of unity. Let  $a_0 \in R_q$  be a fixed invertible element. Let  $T := \text{Tr}_{S_q}^{R_q}$  as defined in Section 4.3, and suppose that  $T(a_0)$  is invertible.*

*Suppose one obtains  $N$  samples  $(a, b)$  distributed according to  $A_{a_0 S_q, s, \chi}$  (notation from Section 4.4).*

*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. These  $m/k$  problems are independent in the sense that setting up any one of them does not require having solved any other one.*

*Furthermore, if  $\chi$  is formed on the  $\zeta$ -basis from coefficient distribution  $\chi_0$  on  $\mathbb{F}_q$  (see Section 3.5 for definition), then so is  $\frac{k}{m}T(\chi)$ .*

*Proof.* Set  $i = j$  in Proposition 5.1 for each  $j$  in the range  $j = 0, \dots, m/k - 1$ , 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

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

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  $\{\zeta^j\}_{j=0,\dots,m/k-1}$  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 carry out Gaussian elimination of dimension  $n = m/2$  over  $\mathbb{F}_q$ , which is polynomial in  $m$  and  $\log q$ .  $\square$

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  $\chi$  formed on a  $\zeta$ -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  $\chi$  is invariant under multiplication by  $\zeta$ . 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  $\chi$  on  $R_q$  gives an error distribution  $\frac{1}{2}T_{S_q}^{R_q}(\chi)$  on  $S_q$ . If  $\chi$  is formed on a  $\zeta$ -basis with coefficients supported in  $E_\chi \subsetneq \mathbb{F}_q$ , then  $\frac{1}{2}T_{S_q}^{R_q}(\chi)$  is formed on a  $\zeta$ -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 solution 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  $1/\sqrt{2}$ .  $\square$

## 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 [5]. 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*  $\beta$  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  $\beta$  in  $W$ . The cost of a reduction step is exponential in  $\beta$ , so we keep  $\beta$  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  $\beta$  is an upper limit on the number of blocks used overall. Each reduction into codimension  $\beta$  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 [2] [14] [18] [19] [20] [22].

## 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. The window size may be chosen at will, for example, and need not depend upon the ring structure. 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  $\zeta$ -basis of  $R_q$  over  $\mathbb{F}_q$ , namely

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

and then reorder it to produce a *prioritized basis*. The most important property we desire for our purposes is that if one of  $\zeta^i$  and  $\zeta^j$  has lower multiplicative order than the other, then it comes later than the other. One computationally convenient 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:

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

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. For example, coded-BKW, under the assumption the secret  $s$  is small, associates to each  $a$  a codeword  $c$  from a linear code. Then the sample  $(a, as + e)$  is replaced with  $(c, as + e)$ , which is a valid sample with a larger error, before it is fed to the BKW tables. The tables then have fewer rows because their rows are chosen from codewords. In sieving, imagine that one has stored the original  $a$  along with each new sample  $(c, as + e)$ . The difference between  $a$  and  $c$  measures the error inflation introduced by coding. A collision between  $(c, a_1s + e)$  and  $(c, a_2s + e)$  being passed to another table

has an  $a = a_1 - a_2$  that is not actually 0 in the first few entries, only small. Among the vectors being fed from one table to the next, one can pause to sieve them, creating vectors whose  $a$ 's are somewhat smaller. This reduces the error inflation introduced by the coding process.

The important thing is that, whatever technique is used, after reduction, one has obtained samples with  $a \in S_q = S/qS$  for some  $S$  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. [14]. 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.

---

The ring structure is not relevant in step (1); one uses BKW reduction as for any LWE problem (in particular, the window size can be chosen without regard to the ring structure). In fact, 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  $w$  and at least  $N$  samples.

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

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

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

*Proof.* We will use Algorithm 1. We will set  $f = w/\sigma$ . The time to run the reduction phase is  $t_{\mathcal{B}}(n, B, w/\sigma, N)$ . 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 (guaranteed to succeed by the choice of  $f$ ) takes time  $t_{\mathcal{R}}(B)$  each. Then reconstructing the secret (as in Theorem 5.2) again takes polynomial time.  $\square$

## 9. ADVANCED KEYING

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

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

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  $a_1 \pm \zeta^i a_2$  for some  $i$ . These are all perfectly useful collisions for the algorithm, if the error term is  $\zeta$ -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. (If one chooses to compare to running BKW without rotating samples at all, i.e. ring-blind, it will both reduce storage *and* require fewer samples.)

In the discussion that follows, the reduction algorithm described in Section 6 will be called *traditional BKW reduction* to distinguish it from the *advanced keying BKW reduction* proposed in this section. There are a variety of modern speedups and alternatives (such as coded-BKW and sieving) which could also be combined with advanced keying, but for purposes of clarity we will ignore these until later in this section. In particular, in traditional BKW reduction, when a collision is recorded, nothing is added to the current table, but the difference is passed to the next table. (Later, it will prove helpful to call this *one-difference* and compare it to *all-differences* where new samples are stored as well as passed on, to increase the number of collisions.)

Our proposal in this section is an analogue of the space-saving technique used in traditional BKW, wherein for each sample  $(a, b)$  we may derive two samples  $(a, b)$  and  $(-a, -b)$ : we choose one canonically (where the first non-zero coefficient of  $a$  is in  $\{1, \dots, \frac{q-1}{2}\}$ , 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 rows in the process. More precisely, the number of rows of the table for each block never exceeds  $(q^B - 1)/2$ , since the possible non-zero vectors come in pairs of which we store at most one. Furthermore, this is also a time efficiency issue. If instead one simply included  $(a, b)$  and  $(-a, -b)$  among the incoming samples, then without this trick, the collisions  $a_1 + a_2$  and  $-a_1 - a_2$  are both sent on to the next table, both are multiplied by  $-1$  thereafter, and we actually end up with repeat samples that must be weeded out at a later stage. For reference, traditional BKW reduction, with this space-saving technique, is given explicitly in Algorithm 3.

The fundamental observation is that the prioritized basis proposed in the last section is particularly well-suited to this type of strategy, because of the resulting ‘negacyclic permutation’ effect of multiplication by  $\zeta$ . It results in a savings of  $1/2B$  instead of  $1/2$  and is completely analogous to the trick above in both space and efficiency savings. It requires that the block size  $B$  be a power of 2.

Write  $\mathbf{a} \in \mathbb{F}_q^n$  for the vector of coefficients of  $a$  in the prioritized basis. The action of  $\zeta^h$  (taking  $a$  to  $\zeta^h a$ ) on such a vector permutes the entries, and swaps the sign on some of them (since  $\zeta^n = -1$ ). Suppose  $h$  is exactly divisible by  $2^\ell$  (i.e.  $\text{ord}_2(h) = \ell$ ). With regards to the permutation only (ignoring the signs), the permutation has the property that it stabilizes each consecutive block of length  $n/2^\ell$  throughout (that is, it permutes each block individually). For fixed  $\ell$ , there are exactly  $n/2^\ell$  such integers  $h$  (note that  $h$  is taken modulo  $2^n$ , for  $h = 2^n$  results in the identity permutation). The following consequence is key:

**Property 1.** *Let  $B \mid n$  denote block size. Then applying  $\zeta^{n/B}$  preserves the property that  $\mathbf{a}$  has first block (or series of any number of first blocks) consisting of zero entries.*

This property will allow us to rotate samples by any of the  $B$  quantities  $1, \zeta^{n/B}, \zeta^{2n/B}, \dots, \zeta^{(B-1)n/B}$  during BKW reduction with block size  $B$ .

Next, one must specify a *canonical choice of representative* from the set of possible rotations  $\{a, \zeta^{n/B}a, \dots, \zeta^{(B-1)n/B}a\}$ , depending only on the first non-zero block of entries, up to an overall sign. A possible canonical choice is the ordering which has smallest first entry (in absolute value), together with some tie-breaking conventions, e.g. smallest second entry, etc., and if all entries are equal in absolute value, then some appropriate convention on sign changes between  $\mathbf{a}$  and  $|\mathbf{a}|$ , etc. However, any ordering of the possible length- $B$  vectors modulo overall sign, will do. It is not possible to break

a tie if the first  $B$  entries of the two rotations actually *agree* up to overall sign under one of the rotations. However, in this case we have found a “self-match,” meaning that two of the rotations have a difference which has all zero in the block under consideration, and so at most one of the two rotations need be stored, and the difference is sent to the following block, as with any collision, as in a traditional BKW algorithm.

The advanced keying BKW reduction is given in Algorithm 3, and for comparison purposes, the traditional BKW reduction using all rotations of each sample is given in Algorithm 2.

---

**Algorithm 2** Traditional BKW Reduction Phase

---

```

1: Create empty Tables 1 through  $n/B$ .
2: for each initially available sample  $(a, b)$  do
3:   for  $j = 0$  to  $n - 1$  do
4:     Rotate the sample by  $\zeta^j$ , to obtain  $(a_1, b_1)$ .
5:     Send the sample  $(a_1, b_1)$  to Table 1.
6:   end for
7: end for
8: for each sample  $(a, b)$  sent to Table  $i$ ,  $i < n/B$  do
9:   if  $a$  has all 0 entries in block  $i$  then
10:    send sample  $(a, b)$  on to Table  $i + 1$ 
11:   end if
12:   Multiply by  $-1$  if necessary to ensure the first non-zero coefficient of
     $a_1$  is in the range 1 to  $(q + 1)/2$ .
13:   if a collision is found (i.e. a sample  $(a_0, b_0)$  already exists in the table
    having the same first  $i$  blocks of size  $B$ ) then
14:     Subtract  $(a_1, b_1)$  from  $(a_0, b_0)$  to obtain a new sample whose first  $i$ 
    blocks of size  $B$  are zero
15:     Send the result to Table  $i + 1$ .
16:   else
17:     Store the associated sample in Table  $i$ .
18:   end if
19: end for

```

---

Correctness of Algorithm 3 is a consequence of Property 1. Furthermore, Algorithms 2 and 3 catch the same collisions in the following heuristic sense. For each collision  $\zeta^i a_1 - \zeta^j a_2$ , there will be another collision at  $\zeta^{i+k} a_1 - \zeta^{j+k} a_2$  for any  $k \equiv 0 \pmod{n/B}$ . In Algorithm 2, all  $B$  of these collisions are passed on to the next table after storing  $B$  new rows in the current table. But any one of the samples sent on can generate the others via rotation, so only one of them is actually needed at the next table. In Algorithm 3, only one of them is stored and only one is sent onward (but only one is needed). However, there is some difference in the final output because we are only keeping one sample per row, and the order of input samples to a given table may differ, resulting in a different table entry. If one uses the *all-differences*

---

**Algorithm 3** Advanced Keying BKW Reduction Phase

---

```
1: Create empty Tables 1 through  $n/B$ .
2: for each initially available sample  $(a, b)$  do
3:   for  $j = 0$  to  $n/B - 1$  do
4:     Rotate the sample by  $\zeta^j$ , to obtain  $(a_1, b_1)$ .
5:     Send the sample  $(a_1, b_1)$  to Table 1.
6:   end for
7: end for
8: for each sample  $(a, b)$  sent to Table  $i$ ,  $i < n/B$  do
9:   if  $a$  has all 0 entries in block  $i$  then
10:    send sample  $(a, b)$  on to Table  $i + 1$ 
11:   end if
12:   From  $a, \zeta^{n/B}a, \dots, \zeta^{(B-1)n/B}a$ , choose a canonical representative.
13:   for every sample  $(a_1, b_1)$  corresponding to a canonical representative
14:     do
15:       Multiply by  $-1$  if necessary to ensure the first non-zero coefficient
16:       of  $a_2$  is in the range 1 to  $(q + 1)/2$ .
17:       if a collision is found (i.e. a sample  $(a_0, b_0)$  already exists in the
18:       table having the same first  $i$  blocks of size  $B$ ) then
19:         Subtract  $(a_1, b_1)$  from  $(a_0, b_0)$  to obtain a new sample whose first
20:          $i$  blocks of size  $B$  are zero
21:         Send the result to Table  $i + 1$ .
22:       else
23:         Store the associated sample in Table  $i$ .
24:       end if
25:     end for
26:   end for
27: end for
```

---

variation, this difference disappears and the output of the two algorithms will be the same.

The following is immediate from Algorithm 3.

**Proposition 9.1.** *Each table in Algorithm 3 has at most  $\frac{q^B - 1}{2B}$  rows in total.*

Finally, we will remark again that BKW reduction improvements for LWE, such as coded-BKW and sieving, may also be adapted to use the advanced keying demonstrated here, *provided* block sizes can be maintained to be powers of 2 (varying them is ok). As some modern algorithms vary block size, this may be an impediment. The naïve way to do this would be to code samples first, then choose a canonical rotation of each codeword. Perhaps better, one could also code each rotation and choose the one with smallest error, which may introduce a significant improvement to the error inflation, depending on the choice of code. (Note that, for those familiar with coded-BKW, the notion of advanced keying is not so different than

coding, as it provides a sort of 'codeword' for each sample, without an error inflation.)

Algorithms 2 and 3, as well as a completely ring-blind version of BKW reduction were coded in Python in Sage Mathematics Software for comparison purposes. Some example results are given in Table 9. In short, the advanced keying did reduce table sizes and samples needed as described, and had a faster overall runtime. A few remarks are in order:

- (1) The experiments were chosen to represent a range of small parameter sets, where timings were in the range of seconds or minutes on a Lenovo X1 laptop.
- (2) After parameters were chosen, the number of samples was chosen to be a round number where the final table began to have a few samples on average; the timing therefore roughly represents the time until the final table begins to populate.
- (3) To compare meaningfully, the ring-blind algorithm uses  $n$  times as many initial samples, which is equal to the total number of rotations of incoming samples for the other algorithms. The fact that the final table is populated but not full in all cases is evidence that the number of samples needed by Algorithms 3 and 2 is  $1/n$  of those needed naïvely.
- (4) For some smaller parameter sets, we also tested a version of the algorithm (labelled AD = 'All Differences') in which every sample encountered is stored (so each row of the table can contain multiple samples) and every difference is passed on (i.e. the new sample is compared to everything already in its row). The purpose of this is to demonstrate that the advanced keying will still find the same number of samples. However, the AD version is significantly slower in all cases, so it was only implemented for some of the smaller parameter sets in the table.
- (5) Algorithms 2 and 3 are pseudocode; the implementation necessarily addressed details not covered in the pseudocode presentation. For example, some moderate attention was given to efficiency in the rotation of samples. For example, when only certain coefficients of the rotation were needed, only those were computed.

Some experimental observations:

- (1) The table sizes observed in Algorithm 3 are very close to  $1/B$  of the number observed in Algorithm 2, as expected.
- (2) The faster runtime of Algorithm 3 is a result of the fact that fewer samples are handled ( $1/B$  as many are fed to the first table compared to Algorithm 2), although they must be handled in more detail, so the speedup is less than a  $1/B$  factor.
- (3) Algorithms 2 and 3 use the exact same starting data, and it is reassuring that the reduced sample counts are similar, and the same in the AD version.

- (4) Algorithm 2 tends to find more samples than Algorithm 3. The difference is in which matches are found when more than two samples collide in a row in the table, and therefore is more pronounced as the number of rows grows.

## 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 of hypothesis testing can be solved in parallel.

Albrecht et al. computed the runtime for BKW [2]. This work has been rendered out of date by many of the modern speedups mentioned in the introduction, but it is likely safe to say a few things that still hold true about modern BKW runtimes. First, the reduction phase is the dominant cost. Second, however, the backsubstitution phase differs from the reduction phase by a polynomial factor, so eliminating it can be expected to give a polynomial factor speedup.

Advanced keying also offers a visible benefit when compared to a ring-blind implementation of BKW. For, compared to a ring-blind implementation, table sizes are reduced to  $1/B$  of their former size and the number of samples used is reduced to approximately  $1/n$  as many. Each sample must be treated rather more carefully however: it is rotated and a canonical choice made. However, experiments still indicate increasing runtime gains with dimension, even against traditional BKW with every sample rotated before beginning. Nevertheless, advanced keying requires block sizes to be a power of 2, and therefore may or may not be useful or extendable in view of the changing block sizes sometimes employed in BKW reduction.

The Ring-LWE Challenges [13] are in the form of *Tweaked Ring-LWE*, which refers to dual Ring-LWE transferred to the unital version (see [13, §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., Bai, S., Ducas, L.: A subfield lattice attack on overstretched NTRU assumptions: cryptanalysis of some FHE and graded encoding schemes. In: Advances in cryptology—CRYPTO 2016. Part I, Lecture Notes in Comput. Sci., vol. 9814, pp. 153–178. Springer, Berlin (2016), [https://doi.org/10.1007/978-3-662-53018-4\\_6](https://doi.org/10.1007/978-3-662-53018-4_6)
- [2] 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>
- [3] Alkim, E., Ducas, L., Pöppelmann, T., Schwabe, P.: Post-quantum key exchange—a new hope. In: 25th USENIX Security Symposium (USENIX Security 16).

$n = 2^3, B = 2^2, q = 211$	Ring-blind	Algorithm 2	Algorithm 3
Initial Samples	$4000 \cdot 2^3$	4000	4000
OD Table Size	31999	31996	7999
OD Reduced Samples	1	1	1
OD Runtime	1.43 s	1.80 s	2.42 s
AD Table Size	31999	31996	7999
AD Reduced Samples	1	1	1
AD Runtime	1.73 s	1.92 s	2.46 s
$n = 2^4, B = 2^2, q = 17$	Ring-blind	Algorithm 2	Algorithm 3
Initial Samples	$2000 \cdot 2^4$	2000	2000
OD Table Size	31985	31988	7997
OD Reduced Samples	15	3	3
OD Runtime	6.13 s	6.24 s	4.69 s
AD Table Size	36448	36623	9163
AD Reduced Samples	81	21	21
AD Runtime	8.43 s	9.73 s	6.27 s
$n = 2^5, B = 2^2, q = 7$	Ring-blind	Algorithm 2	Algorithm 3
Samples	$200 \cdot 2^5$	200	200
OD Table Size	6368	6386	1596
OD Reduced Samples	31	13	4
OD Runtime	7.39 s	8.07 s	4.23 s
$n = 2^6, B = 2^3, q = 3$	Ring-blind	Algorithm 2	Algorithm 3
Samples	$250 \cdot 2^6$	250	250
OD Table Size	15988	15993	1998
OD Reduced Samples	12	7	2
OD Runtime	27.0 s	29.0 s	10.2 s

TABLE 1. The term ‘Ring-blind’ refers to Algorithm 2 but with  $j = 0$  to 0 in line 3, i.e. without rotating any initial samples. A fixed list of samples was generated pseudorandomly for each experiment; ‘Initial Samples’ refers to how many were used from the beginning of the list. ‘Reduced Samples’ refers to the number of samples eventually contained in the last table. ‘Runtime’ refers to the wall time as measured in Sage Mathematics Software. ‘Table Size’ refers to the total number of rows stored not counting the final samples. ‘OD’ (One Difference) refers to the algorithms as presented in the paper. ‘AD’ (All Differences) refers to a modification in which every sample that matches a row is also stored in that row, and when a match is found, the differences with everything in the row are passed on.

pp. 327–343. USENIX Association, Austin, TX (2016), <https://www.usenix.org/conference/usenixsecurity16/technical-sessions/presentation/alkim>

[4] Bernstein, D.J., Lange, T.: Never trust a bunny. In: Radio Frequency Identification. Security and Privacy Issues. RFIDSec 2012., Lecture Notes in Comput.

Sci., vol. 7739, pp. 137–148. Springer, Berlin (2013), [https://doi.org/10.1007/978-3-662-49890-3\\_6](https://doi.org/10.1007/978-3-662-49890-3_6)

[5] 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>

[6] Bos, J.W., Naehrig, M., van de Pol, J.: Sieving for shortest vectors in ideal lattices: a practical perspective. *Int. J. Appl. Cryptogr.* 3(4), 313–329 (2017), <https://doi.org/colorado.idm.oclc.org/10.1504/IJACT.2017.089353>

[7] 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>

[8] 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](https://doi.org/10.1007/978-3-642-22792-9_29)

[9] 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](https://doi.org/10.1007/978-3-662-49890-3_6)

[10] 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>

[11] 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)

[12] Cheon, J.H., Jeong, J., Lee, C.: An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without a low-level encoding of zero. *LMS J. Comput. Math.* 19(suppl. A), 255–266 (2016), <https://doi.org/10.1112/S1461157016000371>

[13] Crockett, E., Peikert, C.: Challenges for ring-LWE. *Cryptology ePrint Archive, Report* 2016/782 (2016), <https://eprint.iacr.org/2016/782>

[14] 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](https://doi.org/10.1007/978-3-662-46800-5_8)

[15] Eisenträger, K., Hallgren, S., Lauter, K.: Weak instances of PLWE. In: *Selected Areas in Cryptography—SAC 2014*, pp. 183–194. Springer (2014)

[16] 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)

[17] Gentry, C., Szydlo, M.: Cryptanalysis of the revised NTRU signature scheme. In: *Advances in cryptology—EUROCRYPT 2002*, Lecture Notes in Comput. Sci., vol. 2332, pp. 299–320. Springer, Berlin (2002), [https://doi.org/10.1007/3-540-46035-7\\_20](https://doi.org/10.1007/3-540-46035-7_20)

[18] Guo, Q., Johansson, T., Mårtensson, 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)

[19] 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)

[20] 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](https://doi.org/10.1007/978-3-662-47989-6_2)

- [21] Herold, G., Kirshanova, E., May, A.: On the asymptotic complexity of solving LWE. *Des. Codes Cryptogr.* 86(1), 55–83 (2018), <https://doi-org.colorado.idm.oclc.org/10.1007/s10623-016-0326-0>
- [22] 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](https://doi.org/10.1007/978-3-662-47989-6_3)
- [23] Kirchner, P., Fouque, P.A.: Revisiting lattice attacks on overstretched NTRU parameters. In: *Advances in cryptology—EUROCRYPT 2017. Part I*, Lecture Notes in Comput. Sci., vol. 10210, pp. 3–26. Springer, Cham (2017)
- [24] 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](https://doi.org/10.1007/978-3-642-13190-5_1)
- [25] 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](https://doi.org/10.1007/978-3-642-38348-9_3)
- [26] 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](https://doi.org/10.1007/978-3-319-44618-9_22)
- [27] 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>
- [28] Schneider, M.: Sieving for shortest vectors in ideal lattices. In: *Progress in cryptology—AFRICACRYPT 2013*, Lecture Notes in Comput. Sci., vol. 7918, pp. 375–391. Springer, Heidelberg (2013), [https://doi-org.colorado.idm.oclc.org/10.1007/978-3-642-38553-7\\_22](https://doi-org.colorado.idm.oclc.org/10.1007/978-3-642-38553-7_22)

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF COLORADO, CAMPUS BOX 395,  
BOULDER, COLORADO 80309-0395

*Email address:* [kstange@math.colorado.edu](mailto:kstange@math.colorado.edu)