


ANTS XIV
Proceedings of the Fourteenth
Algorithmic Number Theory Symposium

On the security of
the multivariate ring learning with errors problem

Carl Bootland, Wouter Castryck, and Frederik Vercauteren



On the security of the multivariate ring learning with errors problem

Carl Bootland, Wouter Castryck, and Frederik Vercauteren

The multivariate ring learning with errors (m -RLWE) problem was introduced in 2015 by Pedrouzo-Ulloa, Troncoso-Pastoriza and Pérez-González. Instead of working over a polynomial residue ring with one variable as in RLWE, it works over a polynomial residue ring in several variables. However, care must be taken when choosing the multivariate rings for use in cryptographic applications as they can be either weak or simply equivalent to univariate RLWE. For example, Pedrouzo-Ulloa et al. suggest using tensor products of cyclotomic rings, in particular power-of-two cyclotomic rings. They claim incorrectly that the security increases with the product of the individual degrees. We present simple methods to solve the search m -RLWE problem far more efficiently than was claimed in the previous literature by reducing the problem to the RLWE problem in dimension equal to the maximal degree of its components (and not the product) and where the noise increases with the square-root of the degree of the other components. Our methods utilise the fact that the defining cyclotomic polynomials share algebraically related roots. We use these methods to successfully attack the search variant of the m -RLWE problem for a set of parameters estimated to offer more than 2600 bits of security, and being equivalent to solving the bounded distance decoding problem in a highly structured lattice of dimension 16384, in less than two weeks of computation time or just a few hours if parallelized on 128 cores. Finally, we also show that optimizing module-LWE cryptosystems by introducing an extra ring structure as is common practice to optimize LWE, can result in a total breakdown of security.

1. Introduction

In concurrent and independent work, Stehlé et al. [22] and Lyubashevsky et al. [14] introduced ring variants of the learning with errors (LWE) problem. The problem in the former is known as the polynomial learning with errors (PLWE) problem while the latter is known as the ring learning with errors (RLWE) problem. The main advantage of using a ring variant over the original problem is that the schemes are much more efficient and the size of the public keys is significantly smaller. Later, a module variant was introduced in [4] where it is called the general learning with errors problem and captures both previous problems as extremes of a broader class of problems.

MSC2010: primary 11T71; secondary 11R18, 94A60.

Keywords: multivariate, ring learning with errors, cyclotomic rings, module-LWE.

For a ring R , free and of finite rank (as a module) over \mathbb{Z} , and positive integers n and q set $R_q = R/qR$. Samples from the module-LWE distribution are of the form (\mathbf{a}, b) where $\mathbf{a} \leftarrow R_q^n$ is uniformly sampled and $b = \langle \mathbf{a}, \mathbf{s} \rangle + e \pmod q$ where $e \leftarrow \chi$ is sampled from an error distribution and $\mathbf{s} \in R_q^n$ is the secret vector. LWE is the case when $R = \mathbb{Z}$ and the ring variant is when $n = 1$ but now the ring R can be thought of as a polynomial residue ring. Thus in going from LWE to its ring variant we replace the inner product of vectors by the product of polynomials (modulo some polynomial modulus). The module-LWE problem is used in cryptographic primitives such as the NIST submissions Saber [8] and Kyber [3].

As previously stated, module-LWE bridges the gap between LWE and RLWE, but is still not as efficient as RLWE. It is thus tempting to replace the inner product in module-LWE by a product of polynomials, just like RLWE, but where now the coefficients are from a polynomial residue ring (in an independent variable) rather than simply integers. This idea naturally leads to the multivariate ring learning with errors (m -RLWE) problem as introduced by Pedrouzo-Ulloa, Troncoso-Pastoriza and Pérez-González in a series of papers [18; 19; 20] between 2015 and 2017. Essentially this does to module-LWE what RLWE does to LWE — by adding more structure they are able to construct more efficient schemes with smaller key sizes.

Originally, only the simplest case of the problem in two variables was formulated. They define this problem in [18], which they call the bivariate RLWE (2-RLWE) problem using the ring $R_q[x, y] = \mathbb{Z}_q[x, y]/(f(x), g(y))$ as follows:

Problem 1.1. Given a bivariate polynomial residue ring $R_q[x, y]$ with $f(x) = x^{n_1} + 1$, $g(y) = y^{n_2} + 1$ and an error distribution $\chi[x, y]$ on $R_q[x, y]$ that generates small-norm random bivariate polynomials in $R_q[x, y]$,¹ 2-RLWE relies upon the computational indistinguishability between samples $(a_i, b_i = a_i \cdot s + e_i)$ and (a_i, u_i) where $a_i, u_i \leftarrow R_q[x, y]$ are chosen uniformly at random from the ring $R_q[x, y]$, and $s, e_i \leftarrow \chi[x, y]$ are drawn from the error distribution.

Although not explicitly stated in [18], f and g are taken to be two-power cyclotomics, i.e., n_1 and n_2 are powers of two.

The authors then construct a method for encrypted image processing whose security is based on the 2-RLWE problem. The sample parameters proposed for use are $n_1 = n_2 = 2^i$, $\lceil \log_2 q \rceil = 22 + 3i$ for $i = 7, 8, 9, 10$. Using the lower bound given in [13, Equation (5.2)] these instances are estimated to have bit security 2663, 10288, 38880 and 146675 respectively, though these parameters fall well outside the range of parameters for which the bound was derived, so these security levels are unlikely to be accurate; however, using the LWE-estimator of Albrecht et al. [1] gives even larger security estimates. Thus it is clear Pedrouzo-Ulloa et al. believe these parameter suggestions give a very high security level. However, in light of our attack, which we will see works in dimension $n_1 = n_2$, the LWE-estimator gives the estimated security levels as 32, 33, 35 and 98 bits respectively.

¹Technically, there is no norm on the ring $R_q[x, y]$ so this statement does not make mathematical sense. What is meant by $\chi[x, y]$ is to sample an element in $\mathbb{Z}[x, y]$ whose degree in x is at most $n_1 - 1$ and whose degree in y is at most $n_2 - 1$ and whose coefficient vector has small-norm, smallness being a function of q , and then reducing the polynomial modulo q , f and g .

Further, in [19], Pedrouzo-Ulloa et al. reformulate the m -RLWE problem in terms of the tensor product of number fields and consider the ring R now as the tensor product of the corresponding rings of integers. They proceed by generalising the security reductions of Lyubashevsky et al. from RLWE to standard problems on ideal lattices to the multivariate case, now reducing them to multivariate ideal lattice problems.

Finally, in [20], Pedrouzo-Ulloa et al. build upon the m -RLWE problem, this time again specialised to power-of-two cyclotomics, and give a number of useful multidimensional signal processing operations and optimizations for use with their m -RLWE based homomorphic encryption scheme.

For the security of their multivariate schemes, the authors claim and give a sketch proof in [20, Proposition 1] that the 2-RLWE problem above is equivalent to the RLWE problem in the ring $\mathbb{Z}_q[z]/(h(z))$ where $h(z) = z^{n_1 n_2} + 1$, however, as will become obvious, this is not true, as we can solve the 2-RLWE problem far more easily. The flaw is that while $\mathbb{Q}[z]/(h(z))$ certainly contains isomorphic copies of $\mathbb{Q}[x]/(f(x))$ and $\mathbb{Q}[y]/(g(y))$, it is not the smallest number field which does so. If we assume $n_1 \geq n_2$ then in this specific case, $\mathbb{Q}[x]/(f(x))$ itself has this property. This shows that we expect to be able to solve the 2-RLWE problem by solving $\max\{n_1, n_2\}$ dimensional problems, not dimension $n_1 n_2$. This logic can be made to work more generally with any cyclotomic fields, not just power of two cyclotomics, as detailed in Section 3A.

In this paper, we give a simple assessment of the security of the m -RLWE problem and present an efficient attack when the polynomial moduli are related in a certain way. The basic idea of the attack is to apply a number of “smallness”-preserving ring homomorphisms which reduce the problem to standard RLWE problems of much lower dimension and with a slightly larger error distribution. Solving the search variant in each case gives us enough information to recover the secret in the original m -RLWE problem. For example, for the 2-RLWE problem above with $n_1 \geq n_2$ the problem is reduced to n_2 instances of the RLWE problem in dimension n_1 , the same modulus q and with the noise growing only by a factor of $\sqrt{n_2}$. This attack shows that the stated hardness of the problem is much lower than had been previously asserted in the literature which claimed security equivalent to RLWE in dimension $n_1 n_2$.

We remark that shortly after our results appeared in an online preprint, Cheon, Kim and Yhee [7] used the m -RLWE problem in defining a generalisation of the HEAAN homomorphic encryption scheme suitable for approximate matrix arithmetic. They also pointed out our evaluation attack and hence used cyclotomic polynomials of coprime order. Furthermore, the original authors of m -RLWE, together with Gama and Georgieva suggested redefining the problem to instead use modular functions of the form $x^{n_1} + d_1, y^{n_2} + d_2, \dots$, where the d_i are small integers, in order to avoid our attack [17].

The remainder of the paper is organised as follows: in Section 2 we recall the required background and in Section 3 we define the m -RLWE problem and show that in many cases it is equivalent to the standard RLWE problem. In Section 4 we present our attack on the remaining cases of m -RLWE and the results of our implementation, and in Section 5 we remark that the standard optimization trick of going from LWE to RLWE, when applied to module-LWE, can result in a total breakdown of security. Finally, we conclude the paper in Section 6.

2. Preliminaries

Let $[n]$ denote the set $\{0, 1, 2, \dots, n-1\}$. For a commutative ring R and an element $r \in R$ we denote by (r) the principal ideal of R generated by r ; namely,

$$(r) = \{rs \mid s \in R\}.$$

For a finite set S we denote by $U(S)$ the uniform distribution on S .

2A. Subgaussians. We also require the notion of a subgaussian random variable. We follow the approach in [15, Section 2.3] and say that a random variable X over \mathbb{R} is subgaussian with parameter $s > 0$ if for all $t \in \mathbb{R}$ we have

$$\mathbb{E}(e^{2\pi t X}) \leq e^{\pi s^2 t^2}.$$

We also use the same notation for the probability distribution of X . It is a simple exercise to show that the sum of subgaussian distributions is also subgaussian:

Lemma 2.1. *Let $s_i \geq 0$ and suppose that we have independent and identically distributed random variables X_i which are subgaussian with parameter s_i . Define X to be the random variable that is the sum of the X_i and set $s = (\sum_i s_i^2)^{1/2}$, then X is subgaussian with parameter s .*

We can also apply Markov's inequality to the subgaussian random variable X with parameter s which shows that

$$\Pr(|X| \geq t) \leq 2e^{-\pi t^2/s^2}.$$

2B. RLWE and its variants. Here we also introduce the distinction between the so-called dual- and primal-RLWE problems as well as the polynomial RLWE problem, abbreviated to PLWE. The starting point for the first two problems is a number field K and its ring of integers \mathcal{O}_K and an integer modulus $q \geq 2$. Typically K is a cyclotomic number field but this need not be the case. Samples are of the form (a_i, b_i) where $b_i = a_i s + e_i$ and $a_i \in \mathcal{O}_K/q\mathcal{O}_K$ is sampled uniformly at random and e_i is sampled from an error distribution on $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R}$. The difference between the two cases is that in the dual-RLWE case the secret s is sampled from $\mathcal{O}_K^{\vee}/q\mathcal{O}_K^{\vee}$, with \mathcal{O}_K^{\vee} the fractional ideal dual to \mathcal{O}_K , while in the primal-RLWE case it is sampled from $\mathcal{O}_K/q\mathcal{O}_K$. Finally, in the PLWE case $a_i, s \in \mathbb{Z}_q[x]/(f)$ for some monic irreducible polynomial f and the error term is an element of $\mathbb{R}[x]/(f)$.

The actual problems come in two variants; a decision version where one has to determine whether the second component of the samples is computed according to the RLWE distribution or chosen randomly as in Problem 1.1, and a search version where one is asked to find the secret s .

It has been shown by Ducas and Durmus [9] for cyclotomic fields, and by Rosca, Stehlé, and Wallet [21] more generally, that one can reduce dual-RLWE to primal-RLWE with only a limited growth in the error term. Also in [21] they show that the reduction can be extended from primal-RLWE to PLWE. Since m -RLWE is defined to use exclusively cyclotomic rings, for simplicity, we will focus on the PLWE problem in this paper. Our attack is, however, more general and we explain the modifications needed to generalise this to the other more general RLWE problems where appropriate.

2C. Search RLWE as a BDD problem. In this section we recall a simple and well-known lattice attack on the search variant of the RLWE problem by considering it as a special case of the bounded distance decoding problem (BDD). The attack works given enough samples and is practical for low-dimensional problems.

Suppose we are given ℓ samples $\{(a_i, b_i)\}_{i \in [\ell]}$ from the PLWE distribution and suppose we are working in the ring $R = \mathbb{Z}_q[x]/(f(x))$, $\deg(f) = n$. Then we know that if s is the secret polynomial we have $b_i = sa_i + e_i$ for some e_i with small coefficients. We can rewrite this as a vector-matrix equation by replacing the elements of R by their (row) vector of coefficients (with respect to the standard power basis in x) which we denote in bold; if M_{a_i} is the matrix of multiplication by a_i then we have $\mathbf{b}_i = sM_{a_i} + \mathbf{e}_i$. Since s is the same for each sample we can concatenate all of the samples into one equation:

$$(\mathbf{b}_1 \cdots \mathbf{b}_\ell) = s(M_{a_1} \cdots M_{a_\ell}) + (\mathbf{e}_1 \cdots \mathbf{e}_\ell).$$

This is an instance of the bounded distance decoding (BDD) problem in the q -ary lattice \mathcal{L} spanned by the rows of $(M_{a_1} \cdots M_{a_\ell})$ (with entries taken as integers) and $qI_{n\ell}$; the target vector being $\mathbf{v} = (\mathbf{b}_1 \cdots \mathbf{b}_\ell)$. Any BDD-solver, such as Kannan's embedding technique [12] or Babai's nearest plane algorithm [2], can thus be used to solve search PLWE. In general, both the ring R and its dual R^\vee can be written as an integral lattice with a suitable choice of basis and the same approach can be taken to write the search problem as a BDD problem.

Two samples will in practice uniquely define s , and the more samples one has, the better the chance of solving the problem. Since we will use the BDD-solver as a black box in our algorithm, we simply refer to the tool of Albrecht et al. [1] which can be used to estimate the running time of these algorithms.

3. The m -RLWE Problem

In [19] the authors define the multivariate RLWE distribution, in its dual formulation, in terms of a tensor product of number fields $K = \bigotimes_{i \in [m]} K_i$ where each K_i is a cyclotomic field; not necessarily distinct. The ring R used is now the tensor product, $R = \bigotimes_{i \in [m]} \mathcal{O}_{K_i}$, where \mathcal{O}_{K_i} is the ring of integers of the number field K_i . Further, one defines $\mathbb{T} := K_{\mathbb{R}}/R^\vee$ where R^\vee is the dual fractional ideal of R called the codifferent ideal. Finally, for an integer modulus $q \geq 2$, set $R_q = R/qR$ and $R_q^\vee = R^\vee/qR^\vee$.

Definition 3.1 (multivariate RLWE distribution). For $s \in R_q^\vee$ and an error distribution ψ over $K_{\mathbb{R}}$, a sample from the m -RLWE distribution $A_{s, \psi}$ over $R_q \times \mathbb{T}$ is generated by sampling $a \leftarrow R_q$ uniformly at random, $e \leftarrow \psi$, and outputting $(a, b = (a \cdot s)/q + e \bmod R^\vee)$.

One can then define the multivariate RLWE search and decision problems in the standard way.

Definition 3.2 (multivariate RLWE search problem). Let Ψ be a family of distributions over $K_{\mathbb{R}}$. Denote by m -RLWE $_{q, \Psi}$ the search version of the m -RLWE problem: given access to arbitrarily many independent samples from $A_{s, \psi}$ for some fixed uniformly random $s \in R_q^\vee$ and $\psi \in \Psi$, find s .

Definition 3.3 (multivariate RLWE decision problem). Let Γ be a distribution over a family of error distributions, each over $K_{\mathbb{R}}$. The average-case-decision version of the m -RLWE problem, denoted by

m -R-DLWE $_{q,\Gamma}$, is to distinguish with nonnegligible advantage between arbitrarily many independent samples from $A_{s,\psi}$, for a random choice of $(s, \psi) \leftarrow U(R_q^\vee) \times \Gamma$, and the same number of uniformly random and independent samples from $R_q \times \mathbb{T}$.

3A. Decomposition of m -RLWE and the compositum field. It is well known that the n -th cyclotomic ring (respectively, field) can be split into a tensor product of prime-power cyclotomic rings (respectively, fields), with these prime powers being those appearing in the factorisation of n . In the case of rings, if we denote the j -th cyclotomic polynomial by Φ_j , we have that if the prime power factorisation of n is $n = p_1^{e_1} \cdots p_m^{e_m}$ then,

$$\frac{\mathbb{Z}[x]}{(\Phi_n(x))} \cong \frac{\mathbb{Z}[x]}{(\Phi_{p_1^{e_1}}(x))} \otimes \cdots \otimes \frac{\mathbb{Z}[x]}{(\Phi_{p_m^{e_m}}(x))}.$$

If φ is the isomorphism from the right-hand side to the left, and we have an instance of the m -RLWE problem in the right-hand tensor product of rings modulo q then lifting the coefficients to \mathbb{Z} , applying φ and reducing modulo q will give an instance of the RLWE problem since $\varphi(q) = q$ and φ is a linear map when considering the rings as \mathbb{Z} -lattices. Furthermore, this map is “smallness”-preserving so the resulting error distribution is still a distribution of small elements, though possibly with some degradation in precisely how small. As a result we obtain the following observation.

Observation. The m -RLWE problem for cyclotomic fields with defining polynomials Φ_{n_i} is only distinct from the RLWE problem when the n_i are not all pairwise coprime.

Going back to the more general case of arbitrary number fields K_i , the way to view the problem is via the notion of the compositum of fields; in our case this is the smallest number field which contains isomorphic copies of each K_i . Then there is a natural algebra homomorphism from the tensor product of the K_i to the compositum; in fact, there can be many such homomorphisms: if we fix one then we can first apply any automorphisms of the K_i before applying this homomorphism to give the others.

We can then distinguish two cases. The first case is the so-called *linearly disjoint* case: the map is injective (and, as such, automatically bijective in our case) and so the tensor product and the compositum are isomorphic. We remark this is only true in terms of the number fields themselves and not the corresponding rings of integers. However, only when this map is not injective is the m -RLWE problem distinct from the RLWE problem and this is the crux of the flaw in the reduction from m -RLWE to RLWE given in [19]. Instead of having to solve a lattice problem in the tensor product of fields whose dimension is the product of the degrees of the defining polynomials, one can work in the compositum field where the lattice problem now has dimension the degree of the compositum as a number field which can be much smaller.

For well-behaved number fields, the natural linear map from the tensor product of the K_i to the compositum is again somewhat “smallness”-preserving. This means that the corresponding RLWE problems in the compositum field may still have small enough error polynomials to be able to mount an attack against them. We note that the m -RLWE problem was introduced to improve the efficiency of certain applications of somewhat homomorphic encryption; the number fields which can be used in these advanced cryptographic primitives are well-behaved in this sense.

Since the RLWE problem is widely deemed to be a hard problem in large dimensions, we will only be interested in the case when the fields K_i are not linearly disjoint. The simplest case of this for cyclotomic fields is when $m = 2$ and the two fields are prime-power cyclotomic fields for the same prime. In particular we will focus on the prime 2 as this is a very popular choice for efficiency reasons.

4. Attacks

4A. A distinguishing attack. Our attack is inspired by the “evaluation at one” attack and its variants on nonstandard decisional PLWE problems [10; 11; 5]. These attacks work if the defining polynomial f of the ring $R = \mathbb{Z}[x]/(f(x))$ has a small root modulo q , say $f(\theta) \equiv 0 \pmod{q}$. Then evaluation at $x = \theta$ is well defined and guessing the value of $s(\theta)$ one can test if $e(\theta) = b(\theta) - a(\theta)s(\theta)$ is distributed according to the error distribution evaluated at θ . This requires $e(\theta)$ to be distinguishable from uniform, which it is if $e(\theta)$ remains small enough; hence θ should also be small, e.g., $\theta = \pm 1$.

Note that evaluation at θ is equivalent to reduction modulo the ideal generated by $x - \theta$ and on further reduction by q the ring is nontrivial if and only if $f(\theta)$ and q are not coprime. To stand any chance of distinguishing though, $f(\theta)$ and q should have a large common factor so that the quotient ring is not too small; this is the case when $f(\theta) \equiv 0 \pmod{q}$. More generally, for the attack to succeed we really only need that $\mathbb{Z}[x]/(f(x), q, x - \theta) = \mathbb{Z}/(f(\theta), q)$ is large enough to distinguish the distribution of $e(\theta)$ from uniform.

In our setting, the ring R is equal to $\mathbb{Z}[x, y]/(f(x), g(y))$ so we look for an ideal \mathcal{I} of R such that \mathcal{I} and (q) are not coprime. In particular, viewing R as

$$\frac{\frac{\mathbb{Z}[x]}{(f(x))}[y]}{(g(y))}$$

we can try to find a root of $g(y)$ modulo q in the ring $\mathbb{Z}[x]/(f(x))$. If such a root $\theta(x)$ exists, we can try to distinguish between $e(x, \theta(x))$ of the form $b(x, \theta(x)) - a(x, \theta(x))s(x, \theta(x))$, hence coming from genuine m -RLWE samples, and $e(x, \theta(x))$ coming from uniformly random samples.

Example 4.1. As a small example let us take $f(x) = x^4 + 1$ and $g(y) = y^2 + 1$. We look for a solution to $y^2 + 1 \equiv 0 \pmod{q}$ in the ring $\mathbb{Z}[x]/(x^4 + 1)$. It is easy to see that a solution is $y = x^2$; hence we have found a root. Thus the mapping $a(x, y) \mapsto a(x, x^2)$ is a ring homomorphism from $\mathbb{Z}[x, y]/(x^4 + 1, y^2 + 1)$ to $\mathbb{Z}[x]/(x^4 + 1)$. The error polynomials will be sampled coefficient-wise with respect to the standard power basis $x^i y^j$ which we use throughout this paper. Thus writing $e(x, y) = \sum_{i=0}^3 \sum_{j=0}^1 e_{i,j} x^i y^j$ we see that under this homomorphism the error polynomial $e(x, y)$ is mapped to

$$\sum_{i=0}^3 \sum_{j=0}^1 e_{i,j} x^{i+2j} = (e_{0,0} - e_{2,1}) + (e_{1,0} - e_{3,1})x + (e_{2,0} + e_{0,1})x^2 + (e_{3,0} + e_{1,1})x^3.$$

We thus see that the image of the error polynomial also has small coefficients as they are just a signed sum of two of the original coefficients. In particular, the coefficients of the error term are distinguishable from

random elements modulo q for large enough q . This means a distinguishing attack can be successfully mounted against the decisional m -RLWE problem in this setting.

We can in fact go a step further in the above example as $y = -x^2$ is another solution to $y^2 + 1 \equiv 0 \pmod{q}$. This may not seem to add much but using this second solution we can perform an attack on the search variant of the problem making the attack much more powerful. More generally, having multiple roots may make a direct attack on the search variant feasible. This will be demonstrated in practice in the next section.

4B. Multiple roots. Take the example of the 2-RLWE problem of Problem 1.1 with $f(x) = x^{n_1} + 1$ and $g(y) = y^{n_2} + 1$ for n_1 and n_2 powers of two so that without loss of generality we can assume that $n_2 \mid n_1$ and let $k = n_1/n_2$. Here we have many roots of $g(y)$ in $\mathbb{Z}[x]/(f(x))$ even before reducing modulo q . Namely we have $g(x^{(2i+1)k}) = 0$ for $i \in [n_2]$ and each of the roots is distinct. We can thus define the map

$$\begin{aligned} \Theta : \mathbb{Z}[x, y]/(f(x), g(y)) &\rightarrow (\mathbb{Z}[x]/(f(x)))^{n_2}, \\ a(x, y) &\mapsto (a(x, x^k), a(x, x^{3k}), \dots, a(x, x^{(2n_2-1)k})). \end{aligned}$$

This map is essentially the canonical embedding of $\mathbb{Z}[y]/(y^{n_2} + 1)$ where, instead of mapping into $\mathbb{Z}[e^{\pi i/n_2}]^{n_2} \subset \mathbb{C}^{n_2}$, each component maps into the ring of integers of the compositum of fields which is isomorphic to $\mathbb{Z}[x]/(x^{n_1} + 1)$ in our case. Thus we see that Θ is a ring homomorphism. We denote by Θ_i the i -th component of Θ which is again a ring homomorphism.

Just like the canonical embedding, the map Θ is injective. Write $a(x, y) = \sum_{j=0}^{n_2-1} a_j(x)y^j$ and let \mathbf{a} be the vector of coefficients with respect to the power basis in y : $\mathbf{a} = (a_0(x), \dots, a_{n_2-1}(x))$. Then,

$$\Theta(a(x, y)) = \mathbf{a} \begin{pmatrix} 1 & 1 & \dots & 1 \\ x^k & x^{3k} & \dots & x^{(2n_2-1)k} \\ x^{2k} & x^{6k} & \dots & x^{(2n_2-1)2k} \\ \vdots & \vdots & \ddots & \vdots \\ x^{(n_2-1)k} & x^{3(n_2-1)k} & \dots & x^{(2n_2-1)(n_2-1)k} \end{pmatrix}.$$

This matrix is a Vandermonde matrix and thus has determinant $\prod_{0 \leq i < j < n_2} (x^{(2j+1)k} - x^{(2i+1)k})$ which is nonzero as the $x^{(2i+1)k}$ are distinct for $i \in [n_2]$. Hence Θ is injective and can thus be inverted. Further, for $n_2 > 2$, the absolute value of this determinant is a square root of the discriminant of the number field $\mathbb{Q}(e^{\pi i/n_2})$. It is well known (see, for example, [24, Proposition 2.1]) that the discriminant is $n_2^{n_2}$ so the determinant is one of $\pm n_2^{n_2/2}$. Hence for odd q the corresponding map Θ modulo q which we denote by $\bar{\Theta}$ is also invertible; here we mean the map

$$\begin{aligned} \bar{\Theta} : \mathbb{Z}_q[x, y]/(f(x), g(y)) &\rightarrow (\mathbb{Z}_q[x]/(f(x)))^{n_2}, \\ a(x, y) &\mapsto (a(x, x^k), a(x, x^{3k}), \dots, a(x, x^{(2n_2-1)k})). \end{aligned}$$

The inverse mapping from the image of Θ (or $\bar{\Theta}$ if it exists) is given by multiplying by the inverse of the Vandermonde matrix on the right. If we denote the Vandermonde matrix by $T = (T_{i,j})_{i,j \in [n_2]}$ then

its inverse is given by $U = (U_{i,j})_{i,j \in [n_2]}$ where $U_{i,j} = \frac{1}{n_2} x^{-2jk} T_{j,n_2-i} = \frac{1}{n_2} x^{-j(2i+1)k}$ where the indices are taken modulo n_2 . To see this we compute

$$\begin{aligned} (TU)_{i,j} &= \sum_{m=0}^{n_2-1} T_{i,m} U_{m,j} = \sum_{m=0}^{n_2-1} x^{i(2m+1)k} \frac{1}{n_2} x^{-j(2m+1)k} \\ &= \frac{1}{n_2} \sum_{m=0}^{n_2-1} x^{(i-j)(2m+1)k} = \delta_{i,j}. \end{aligned}$$

We now look at how large the coefficients of the t -th component of $\Theta(e(x, y))$, denoted $\Theta_t(e(x, y))$, are if $e(x, y)$ is sampled from the m -RLWE error distribution. We suppose that this error distribution has coefficients, with respect to the basis $x^i y^j$, sampled independently from a distribution that is subgaussian with parameter σ so writing $e(x, y) = \sum_{i=0}^{n_2-1} \sum_{j=0}^{n_1-1} e_{i,j} x^i y^j$, each $e_{i,j}$ is an independent subgaussian random variable with parameter σ . Then applying Θ_t for some $t \in [n_2]$ gives

$$\Theta_t(e(x, y)) = \sum_{i=0}^{n_2-1} \sum_{j=0}^{n_1-1} e_{i,j} x^{j+i(2t+1)k} = \sum_{l=0}^{n_1-1} \left(\sum_{i=0}^{n_2-1} (-1)^{q_{i,l}} e_{i,r_{i,l}} \right) x^l,$$

where we define $q_{i,l}$ and $r_{i,l}$ as the quotient and remainder of $l - i(2t+1)k$ on division by n_1 (which depends on t): $l - i(2t+1)k = q_{i,l} n_1 + r_{i,l}$ with $r_{i,l} \in [n_1]$. This can be seen by rewriting j as $j = l - i(2t+1)k \pmod{n_1}$ for some $l \in [n_1]$ (for each i separately) and noting that as j runs over $[n_1]$ so does l , after which one swaps the order of summation.

Thus we see that the coefficients of $\Theta_t(e(x, y))$ are the sum of n_2 subgaussians with parameter σ and so are themselves subgaussian with parameter $\sqrt{n_2} \sigma$.

4C. Our attack. Here we present a simple attack on the 2-RLWE problem. It combines both the simple lattice attack and the distinguishing attack. We stress that the attack is much more powerful than the distinguishing attack alone as firstly it solves a search rather than a decisional problem and secondly there is no need for any guessing during the attack. We point out that our attack has a strong similarity to Nussbaumer's algorithm for fast convolution [16].

We start with a number of samples $\{(a_j(x, y), b_j(x, y))\}_{j \in [\ell]}$ where

$$b_j(x, y) = a_j(x, y)s(x, y) + e_j(x, y).$$

The attack starts by evaluating the map $\bar{\Theta}$ on each sample; we define $\alpha_{i,j}(x) := \bar{\Theta}_i(a_j(x, y))$ and $\beta_{i,j}(x) := \bar{\Theta}_i(b_j(x, y))$. We note that since $\bar{\Theta}$ is a ring homomorphism we have, on defining $\epsilon_{i,j}(x) := \bar{\Theta}_i(e_j(x, y))$ and $\sigma_i(x) := \bar{\Theta}_i(s(x, y))$, that

$$\beta_{i,j}(x) = \alpha_{i,j}(x)\sigma_i(x) + \epsilon_{i,j}(x) \quad \text{for } i \in [n_2], j \in [\ell].$$

Our first goal is to find the $\sigma_i(x)$ and to do this we use the simple lattice attack from Section 2C since for a fixed i the samples $(\alpha_{i,j}(x), \beta_{i,j}(x))$ follow an $\text{RLWE}_{q, \sqrt{n_2} \Psi}$ distribution. This means we need to simply solve n_2 instances of an RLWE problem in dimension n_1 with noise distribution that is $\sqrt{n_2}$ times

		n_1											
		4		8		16		32		64		128	
instances		100		100		100		10		1		1	
block size		30		30		30		30		10		10	
		ℓ	p	ℓ	p	ℓ	p	ℓ	p	ℓ	p	ℓ	p
n_2	4	2	13	2	13	2	13	2	13	2	15	2	21
		3	9	3	10	3	10	3	11	3	13	3	20
	8			2	13	2	13	2	14	2	17	2	22
				3	10	3	10	3	11	3	15	3	20
	16					2	14	2	15	2	18	2	23
						3	11	3	12	3	16	3	22
	32							2	15	2	19	2	24
								3	12	3	17	3	22
	64									2	20	2	31
										3	18	3	24

Table 1. The number of samples $\ell \leq 3$ and the minimal $p \in \mathbb{N}$, $p \approx \log_2(q)$, for which our attack succeeded in each of the stated number of attempts for the stated block size, given n_1 , n_2 and $q = 2^p + 1$, and where the secret polynomial is sampled uniformly at random in R_q .

wider than for the m -RLWE problem; each instance is independent so can be solved in parallel. If this succeeds we have computed the image of $s(x, y)$ under $\bar{\Theta}$ and since $\bar{\Theta}$ is invertible for odd q we can compute $s(x, y)$ and solve the 2-RLWE problem.

4D. Implementation results. We implemented and tested our attack in SageMath [23], using the NTL library for lattice reduction. We tested our attack on the smallest parameter set given in [18], namely for $n_1 = n_2 = 128$ and q being the smallest prime larger than 2^{42} . The secret polynomial is sampled from the error distribution which samples coefficients independently from a discrete Gaussian with $\sigma = 8/\sqrt{2\pi} \approx 3.19$ (the default in SEAL [6]), larger than the stated $\sigma = 1$ in [18]. We were able to successfully recover the secret polynomial with just one sample using BKZ reduction with block size 10 to solve the BDD problem instances. This clearly shows that the estimated security level of over 2500 bits is a significant overestimate. We can see from the estimates given by the LWE estimator [1] that the parameter sets with $n_1 = n_2 = 256$ and $n_1 = n_2 = 512$ also offer little to no security (33 and 35 bits, respectively) while that for $n_1 = n_2 = 1024$ offers at most 98 bits.

In Table 1 we report on a run of our attack with $n_1 \geq n_2$ and q of the form $2^p + 1$ for $p \in \mathbb{N}$. The secret polynomial s we try to find is chosen uniformly at random from $\mathbb{Z}_q[x, y]/(x^{n_1} + 1, y^{n_2} + 1)$ so the minimum number of 2-RLWE samples possible to recover s is two. We give the minimum q of the stated form for which the attack succeeded in a fixed number of consecutive instances with the stated number of samples; here we used the embedding approach combined with BKZ reduction to attempt to solve the BDD instances. Further, the coefficients of the error polynomials were sampled independently using a discrete Gaussian sampler with $\sigma = 3.19$. The results are heuristic as we only attempted to solve

		n_1											
		4		8		16		32		64		128	
instances		100		100		100		10		1		1	
block size		30		30		30		30		10		10	
		ℓ	p	ℓ	p	ℓ	p	ℓ	p	ℓ	p	ℓ	p
n_2	4	1	11	1	12	1	12	1	13	1	14	1	22
		2	9	2	9	2	10	2	11	2	13	2	20
	8			1	13	1	13	1	14	1	15	1	22
				2	10	2	10	2	11	2	14	2	21
	16					1	14	1	14	1	17	1	22
						2	11	2	12	2	15	2	21
	32							1	15	1	18	1	23
								2	12	2	16	2	22
	64									1	20	1	25
										2	17	2	23

Table 2. The number of samples $\ell \leq 2$ and the minimal $p \in \mathbb{N}$, $p \approx \log_2(q)$ for which our attack succeeded in the stated number of instances and with the stated block size, given n_1 , n_2 and $q = 2^p + 1$, and where the secret polynomial is sampled coefficient-wise with each coefficient uniformly random in $\{-1, 0, 1\}$.

a limited number of instances for each choice of n_1 , n_2 and q . It is certainly possible to find the secret for smaller q by increasing the block size used, and in specific instances this may not even be necessary.

In Table 2 we performed the same attack but this time with the coefficients of the secret polynomial taken from the uniform distribution on $\{-1, 0, 1\}$; hence a successful attack is possible with only one sample. While the case of the secret being sampled from the error distribution, as in the proposed image processing scheme of [20], can be viewed as having an extra sample $(1, 0 = 1 \cdot s - s)$ whose error is $-s$, it is often the case in practical applications of somewhat homomorphic encryption that the secret is sampled from this narrower distribution to get the most efficiency out of the scheme. It is therefore interesting to see how this choice affects our attack.

4E. The case of the general m -RLWE problem. The previous subsection showed that the 2-RLWE problem can be readily attacked with the combination of an evaluation attack and simple lattice reduction techniques. More generally, if the defining polynomials of the 2-RLWE problem are both p -th power cyclotomic polynomials of degree $\phi(p^{r_i})$, where ϕ is the Euler-totient function, then our attack straightforwardly applies to this case with the caveat that $\bar{\Theta}$ must be invertible modulo q which holds if q is coprime with $\phi(p^{r_2}) = p^{r_2-1}(p-1)$. We remark that if $h = \gcd(q, \phi(p^{r_2}))$ and $\phi(p^{r_2})$ are small, it is possible to compute all possible preimages of $\bar{\Theta}$ and test each of them in turn to determine the correct value of the secret, however this rather quickly becomes prohibitively expensive the larger h and r_2 become as there are $h^{\phi(p^{r_2})}$ possibilities to check.

Increasing the value of m when each of the defining polynomials is a p -th power cyclotomic polynomial of degree $n_i = \phi(p^{r_i})$ increases the difficulty of the problem since the error grows by a multiplicative

factor of $\sqrt{\prod_{i=j+1}^m n_i}$ in a lattice of dimension $\prod_{i=1}^j n_i$ for some $1 \leq j \leq m$; here we can choose the order of the n_i which best suits the attack. We therefore see that a trade-off can be made in choosing j : if $j = 1$ means the error is already too large for the lattice reduction attack to succeed, we can choose a larger j at the cost of having to perform lattice reduction in a lattice of larger dimension. In this way, taking large m offers some security but at a loss of efficiency if such a large m is not needed specifically for the application in mind.

When instantiating m -RLWE with an arbitrary tensor product of number fields we again wish to find an analogue for the map Θ . This will consist of algebra homomorphisms from $K = \bigotimes_{i \in [m]} K_i$ to the compositum field, which we denote by L . These algebra homomorphisms can naturally be extended to maps from $K_{\mathbb{R}}$ to $L \otimes_{\mathbb{Q}} \mathbb{R}$ which fix q so we can evaluate them on the components of samples from the m -RLWE distribution.

In the case that all the number fields K_i are Galois extensions then there are exactly

$$n := \prod_{i \in [m]} [K_i : \mathbb{Q}] = \prod_{i \in [m]} n_i$$

such algebra homomorphisms from K to L . Since all of the K_i are Galois, so is L ; if we define $N := |\text{Aut}(L)| = [L : \mathbb{Q}]$ as the number of automorphisms of L then up to automorphism in L there are $k := n/N$ distinct algebra homomorphisms which we denote by $\Theta = (\Theta_i)_{i \in [k]}$.

Again, Θ is injective so can be inverted; however for the attack to work we need $\bar{\Theta}$ to be invertible, that is, Θ to be invertible modulo q . Further, we also require Θ to map the error distribution ψ over $K_{\mathbb{R}}$ to elements of $L_{\mathbb{R}}$ which have small coefficients with respect to a known basis for L as a \mathbb{Q} -vector space. If these conditions are met then we can carry out the same attack of applying $\bar{\Theta}$ to the m -RLWE samples, solving k instances of the reduced problem in a lattice of dimension N and applying $\bar{\Theta}^{-1}$ to recover the secret.

To summarise the requirements for the full attack, we require for the number fields K_i to be Galois, for the map $\bar{\Theta}$ to be invertible and for Θ to map small elements to small elements. Nevertheless, if either of the first two conditions are not met it may still be possible to recover partial information about the secret using our approach.

5. The dangers of optimizing module based cryptosystems

We take the example of Kyber [3] which, when reduced to its simplest form, has a public key which is a module-LWE sample where the secret s is a small element of the module R_q^k where $R = \mathbb{Z}[x]/(x^n + 1)$ with n a power of two. Such a public key is then a pair (A, \mathbf{b}) with A a $k \times k$ matrix whose entries are chosen uniformly at random from R_q and $\mathbf{b} \in R_q^k$ with $\mathbf{b} = A\mathbf{s} + \mathbf{e}$ for some small error element $\mathbf{e} \in R_q^k$. This means a public key consists of $k(k+1)$ elements of R_q . One might be tempted to use a structured matrix, such as a negacyclic one, instead of a uniformly random one; after all this is essentially how one goes from LWE to its ring based counterpart RLWE and with our current understanding this latter optimization only incurs a negligible deterioration in security.

Let us fix some parameters and observe what happens. The suggested “paranoid” parameters from [3] are to take $k = 4$ and $n = 256$ and $q = 6781$ which gives a (post-quantum) security level of 218 bits, the largest given by the authors. Taking the matrix A to be negacyclic, that is a matrix of the form

$$\begin{pmatrix} a_0 & -a_{k-1} & -a_{k-2} & \cdots & -a_1 \\ a_1 & a_0 & -a_{k-1} & \cdots & -a_2 \\ a_2 & a_1 & a_0 & \cdots & -a_3 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{k-1} & a_{k-2} & a_{k-3} & \cdots & a_0 \end{pmatrix},$$

means that only 5 elements of R_q are needed to define the public key instead of 20. Further, as shown below, the scheme can be interpreted as adding a ring structure on top of R_q in a new variable y satisfying $y^4 + 1$ and replacing matrix multiplication by ring multiplication. Hence, we are in the m -RLWE setting and working in the tensor product of two power-of-two cyclotomic fields of degrees 256 and 4, respectively.

Formally, we can define the negacyclic module-LWE problem as follows. Let $R = \mathbb{Z}[x]/(x^n + 1)$ with n a power of two and let $q \geq 2$ and k be positive integers. Let s be an element of R_q^k and χ a distribution of small elements in R_q^k . A sample from the negacyclic module-LWE distribution with secret s is of the form $(A, b = As + e)$, where $A \in R_q^{k \times k}$ is a negacyclic matrix and $e \leftarrow \chi$. The negacyclic module-LWE decision problem is to decide whether a given set of samples of the form $(A_i, b_i) \in R_q^{k \times k} \times R_q^k$, with each A_i a negacyclic matrix are sampled from the negacyclic module-LWE distribution or with each b_i sampled uniformly at random from R_q^k instead. The negacyclic module-LWE search problem is, given samples from the negacyclic module-LWE distribution with secret s , to recover s .

Given a negacyclic matrix $A \in R_q^{k \times k}$ whose first column is $(a_0, \dots, a_{k-1})^T$, we can write $a(y) = \sum_{i=0}^{k-1} a_i y^i$ so that the equality $b = As$ is equivalent to

$$b(y) = a(y)s(y) \pmod{y^k + 1},$$

where $b(y) = \sum_{i=0}^{k-1} b_i y^i$ and $s(y) = \sum_{i=0}^{k-1} s_i y^i$ with the b_i and s_i the coordinates of the vectors b and s , respectively. We therefore see that the negacyclic module-LWE problem is equivalent to the m -RLWE problem in the ring $\mathbb{Z}[x, y]/(x^n + 1, y^k + 1)$.

Returning to our example of a structured Kyber variant, we can thus apply our attack with $n_1 = 256$ and $n_2 = 4$ which shows that we can recover s by solving four RLWE problems in dimension 256 from one sample where the error distribution has variance twice that of the original error distribution. Using the LWE-estimator [1], we find that this basic version of a structured Kyber offers at most 107 bits of security, essentially halving the bit security when compared to the original version of Kyber without any additional structure. Thus there is a large difference in terms of security between going from LWE to RLWE and going from module-LWE to m -RLWE if one is not careful.

We note this structured Kyber would also be weak with the “light” parameter set where $k = 2$, but for the standard parameters where $k = 3$ the above attack does not apply as 3 is not a power of two;

that is, $x^3 + 1$ has no roots in a power-of-two cyclotomic field. This again shows the subtlety of the problem of trying to optimize module-LWE. Care needs to be taken in choosing which method and for which parameters such an optimization can be applied without severely damaging the security of the problem.

6. Conclusion

In this paper we reconsidered the m -RLWE problem and its security. We showed that, with a combination of simple evaluation and lattice attacks, the security of the m -RLWE problem was dramatically less than had been previously estimated in the literature. We would therefore not recommend using 2-RLWE for values of n_1 or n_2 less than those used in standard RLWE based schemes for cryptographic purposes. More generally, we conclude that the m -RLWE problem using number fields with a small degree compositum field is insecure. Finally, this paper should also serve as a warning to implementers of module-LWE based cryptosystems to not blindly apply the standard optimization trick that is used to transform LWE into RLWE.

Acknowledgements

This work was supported in part by the Research Council KU Leuven grants C14/18/067 and STG/17/019, by CyberSecurity Research Flanders with reference number VR20192203 as well as by the Research Foundation Flanders (FWO) through WOG Coding Theory and Cryptography. Bootland was supported by a PhD fellowship of the Research Foundation Flanders (FWO).

References

- [1] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of Learning with Errors. *Journal of Mathematical Cryptology*, 9(3):169–203, 2015.
- [2] László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [3] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS - Kyber: A CCA-Secure Module-Lattice-Based KEM. In *2018 IEEE European Symposium on Security and Privacy*, pages 353–367, 2018.
- [4] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. Fully Homomorphic Encryption without Bootstrapping. Cryptology ePrint Archive, Report 2011/277, 2011. <https://eprint.iacr.org/2011/277>.
- [5] Wouter Castryck, Iliia Iliashenko, and Frederik Vercauteren. Provably Weak Instances of Ring-LWE Revisited. In *EUROCRYPT 2016*, pages 147–167. Springer-Verlag, 2016.
- [6] Hao Chen, Kim Laine, and Rachel Player. Simple Encrypted Arithmetic Library - SEAL v2.1. In Michael Brenner, Kurt Rohloff, Joseph Bonneau, Andrew Miller, Peter Y.A. Ryan, Vanessa Teague, Andrea Bracciali, Massimiliano Sala, Federico Pintore, and Markus Jakobsson, editors, *Financial Cryptography and Data Security*, pages 3–18. Springer International Publishing, 2017.
- [7] Jung Hee Cheon, Andrey Kim, and Donggeon Yhee. Multi-dimensional packing for HEAAN for approximate matrix arithmetics. Cryptology ePrint Archive, Report 2018/1245, 2018. <https://eprint.iacr.org/2018/1245>.
- [8] Jan-Pieter D’Anvers, Angshuman Karmakar, Sujoy Sinha Roy, and Frederik Vercauteren. Saber: Module-LWR Based Key Exchange, CPA-Secure Encryption and CCA-Secure KEM. In Antoine Joux, Abderrahmane Nitaj, and Tajjeeddine Rachidi, editors, *AFRICACRYPT 2018*, pages 282–305. Springer International Publishing, 2018.

- [9] Léo Ducas and Alain Durmus. Ring-LWE in Polynomial Rings. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012*, pages 34–51. Springer Berlin Heidelberg, 2012.
- [10] Kirsten Eisenträger, Sean Hallgren, and Kristin Lauter. Weak Instances of PLWE. In Antoine Joux and Amr Youssef, editors, *SAC 2014*, pages 183–194. Springer International Publishing, 2014.
- [11] Yara Elias, Kristin E. Lauter, Ekin Özman, and Katherine E. Stange. Provably Weak Instances of Ring-LWE. In Rosario Gennaro and Matthew Robshaw, editors, *CRYPTO 2015*, pages 63–92. Springer Berlin Heidelberg, 2015.
- [12] Ravi Kannan. Minkowski’s Convex Body Theorem and Integer Programming. *Mathematics of Operations Research*, 12(3):415–440, 1987.
- [13] Richard Lindner and Chris Peikert. Better Key Sizes (and Attacks) for LWE-Based Encryption. In Aggelos Kiayias, editor, *CT-RSA 2011*, pages 319–339. Springer Berlin Heidelberg, 2011.
- [14] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On Ideal Lattices and Learning with Errors over Rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, pages 1–23. Springer Berlin Heidelberg, 2010.
- [15] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. A Toolkit for Ring-LWE Cryptography. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, pages 35–54. Springer Berlin Heidelberg, 2013.
- [16] H. Nussbaumer. Fast polynomial transform algorithms for digital convolution. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 28(2):205–215, 1980.
- [17] Alberto Pedrouzo-Ulloa, Juan Ramón Troncoso-Pastoriza, Nicolas Gama, Mariya Georgieva, and Fernando Pérez-González. Revisiting Multivariate Ring Learning with Errors and its Applications on Lattice-based Cryptography. Cryptology ePrint Archive, Report 2019/1109, 2019. <https://eprint.iacr.org/2019/1109>.
- [18] Alberto Pedrouzo-Ulloa, Juan Ramón Troncoso-Pastoriza, and Fernando Pérez-González. Multivariate lattices for encrypted image processing. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1707–1711. IEEE, 2015.
- [19] Alberto Pedrouzo-Ulloa, Juan Ramón Troncoso-Pastoriza, and Fernando Pérez-González. On Ring Learning with Errors over the Tensor Product of Number Fields. <https://arxiv.org/abs/1607.05244>, 2016.
- [20] Alberto Pedrouzo-Ulloa, Juan Ramón Troncoso-Pastoriza, and Fernando Pérez-González. Multivariate Cryptosystems for Secure Processing of Multidimensional Signals. <https://arxiv.org/abs/1712.00848>, 2017.
- [21] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the Ring-LWE and Polynomial-LWE Problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018*, pages 146–173. Springer International Publishing, 2018.
- [22] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, pages 617–635. Springer Berlin Heidelberg, 2009.
- [23] The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 7.5.1)*, 2017. <http://www.sagemath.org>.
- [24] Lawrence C. Washington. *Introduction to Cyclotomic Fields*. Graduate Texts in Mathematics. Springer New York, 1997.

Received 28 Feb 2020.

CARL BOOTLAND: carl.bootland@kuleuven.be
imec – COSIC, KU Leuven, Heverlee, Belgium

WOUTER CASTRYCK: wouter.castryck@kuleuven.be
imec – COSIC, KU Leuven, Leuven, Belgium

and

Department of Mathematics: Algebra and Geometry, Ghent University, Ghent, Belgium

FREDERIK VERCAUTEREN: frederik.vercauteren@kuleuven.be
imec – COSIC, KU Leuven, Heverlee, Belgium

VOLUME EDITORS

Stephen D. Galbraith
Mathematics Department
University of Auckland
New Zealand

<https://orcid.org/0000-0001-7114-8377>

The cover image is based on an illustration from the article “Supersingular curves with small noninteger endomorphisms”, by Jonathan Love and Dan Boneh (see p. 9).

The contents of this work are copyrighted by MSP or the respective authors. All rights reserved.

Electronic copies can be obtained free of charge from <http://msp.org/obs/4> and printed copies can be ordered from MSP (contact@msp.org).

The Open Book Series is a trademark of Mathematical Sciences Publishers.

ISSN: 2329-9061 (print), 2329-907X (electronic)

ISBN: 978-1-935107-07-1 (print), 978-1-935107-08-8 (electronic)

First published 2020.



MATHEMATICAL SCIENCES PUBLISHERS

798 Evans Hall #3840, c/o University of California, Berkeley CA 94720-3840
contact@msp.org

<http://msp.org>

Fourteenth Algorithmic Number Theory Symposium

The Algorithmic Number Theory Symposium (ANTS), held biennially since 1994, is the premier international forum for research in computational and algorithmic number theory. ANTS is devoted to algorithmic aspects of number theory, including elementary, algebraic, and analytic number theory, the geometry of numbers, arithmetic algebraic geometry, the theory of finite fields, and cryptography.

This volume is the proceedings of the fourteenth ANTS meeting, which took place 29 June to 4 July 2020 via video conference, the plans for holding it at the University of Auckland, New Zealand, having been disrupted by the COVID-19 pandemic. The volume contains revised and edited versions of 24 refereed papers and one invited paper presented at the conference.

TABLE OF CONTENTS

Commitment schemes and diophantine equations — José Felipe Voloch	1
Supersingular curves with small noninteger endomorphisms — Jonathan Love and Dan Boneh	7
Cubic post-critically finite polynomials defined over \mathbb{Q} — Jacqueline Anderson, Michelle Manes and Bella Tobin	23
Faster computation of isogenies of large prime degree — Daniel J. Bernstein, Luca De Feo, Antonin Leroux and Benjamin Smith	39
On the security of the multivariate ring learning with errors problem — Carl Bootland, Wouter Castryck and Frederik Vercauteren	57
Two-cover descent on plane quartics with rational bitangents — Nils Bruin and Daniel Lewis	73
Abelian surfaces with fixed 3-torsion — Frank Calegari, Shiva Chidambaram and David P. Roberts	91
Lifting low-gonal curves for use in Tuitman’s algorithm — Wouter Castryck and Floris Vermeulen	109
Simultaneous diagonalization of incomplete matrices and applications — Jean-Sébastien Coron, Luca Notarnicola and Gabor Wiese	127
Hypergeometric L -functions in average polynomial time — Edgar Costa, Kiran S. Kedlaya and David Roe	143
Genus 3 hyperelliptic curves with CM via Shimura reciprocity — Bogdan Adrian Dina and Sorina Ionica	161
A canonical form for positive definite matrices — Mathieu Dutour Sikirić, Anna Haensch, John Voight and Wessel P.J. van Woerden	179
Computing Igusa’s local zeta function of univariates in deterministic polynomial-time — Ashish Dwivedi and Nitin Saxena	197
Computing endomorphism rings of supersingular elliptic curves and connections to path-finding in isogeny graphs — Kirsten Eisenträger, Sean Hallgren, Chris Leonardi, Travis Morrison and Jennifer Park	215
New rank records for elliptic curves having rational torsion — Noam D. Elkies and Zev Klagsbrun	233
The nearest-colattice algorithm: Time-approximation tradeoff for approx-CVP — Thomas Espitau and Paul Kirchner	251
Cryptanalysis of the generalised Legendre pseudorandom function — Novak Kaluđerović, Thorsten Kleinjung and Dušan Kostić	267
Counting Richelot isogenies between superspecial abelian surfaces — Toshiyuki Katsura and Katsuyuki Takashima	283
Algorithms to enumerate superspecial Howe curves of genus 4 — Momonari Kudo, Shushi Harashita and Everett W. Howe	301
Divisor class group arithmetic on $C_{3,4}$ curves — Evan MacNeil, Michael J. Jacobson Jr. and Renate Scheidler	317
Reductions between short vector problems and simultaneous approximation — Daniel E. Martin	335
Computation of paramodular forms — Gustavo Rama and Gonzalo Tornaría	353
An algorithm and estimates for the Erdős–Selfridge function — Brianna Sorenson, Jonathan Sorenson and Jonathan Webster	371
Totally p -adic numbers of degree 3 — Emerald Stacy	387
Counting points on superelliptic curves in average polynomial time — Andrew V. Sutherland	403