

A subfield lattice attack on overstretched NTRU assumptions

Cryptanalysis of some FHE and Graded Encoding Schemes

Martin Albrecht^{1*}, Shi Bai^{2**}, and Léo Ducas^{3***}

¹ Information Security Group, Royal Holloway, University of London.

`martin.albrecht@royalholloway.ac.uk`

² ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France.

`shih.bai@gmail.com`

³ Cryptology Group, CWI, Amsterdam, The Netherlands.

`ducas@cwi.nl`

Abstract. The subfield attack exploits the presence of a subfield to solve overstretched versions of the NTRU assumption: norming the public key h down to a subfield may lead to an easier lattice problem and any sufficiently good solution may be lifted to a short vector in the full NTRU-lattice. This approach was originally sketched in a paper of Gentry and Szydło at Eurocrypt’02 and there also attributed to Jonsson, Nguyen and Stern. However, because it does not apply for small moduli and hence NTRUencrypt, it seems to have been forgotten. In this work, we resurrect this approach, fill some gaps, analyze and generalize it to any subfields and apply it to more recent schemes. We show that for significantly larger moduli —a case we call overstretched— the subfield attack is applicable and asymptotically outperforms other known attacks.

This directly affects the asymptotic security of the bootstrappable homomorphic encryption schemes LTV and YASHE which rely on a mildly overstretched NTRU assumption: the subfield lattice attack runs in sub-exponential time $2^{O(\lambda/\log^{1/3}\lambda)}$ invalidating the security claim of $2^{\Theta(\lambda)}$. The effect is more dramatic on GGH-like Multilinear Maps: this attack can run in polynomial time without *encodings of zero* nor the *zero-testing parameter*, yet requiring an additional quantum step to recover the secret parameters exactly.

We also report on practical experiments. Running LLL in dimension 512 we obtain vectors that would have otherwise required running BKZ with block-size 130 in dimension 8192. Finally, we discuss concrete aspects of this attack, the condition on the modulus q to guarantee full immunity, discuss countermeasures and propose open questions.

* Supported by EPSRC grant EP/L018543/1 “Multilinear Maps in Cryptography”.

** Supported by ERC Starting Grant ERC-2013-StG-335086-LATTAC.

*** Supported by a grant from CWI from budget for public-private-partnerships and by a grant from NXP Semiconductors through the European Union’s H2020 Programme under grant agreement number ICT-645622 (PQCRYPTO) and ICT-644209 (HEAT).

Keywords: Subfield lattice attack, overstretched NTRU, FHE, Graded Encoding Schemes.

1 Introduction

Lattice-based cryptography relies on the presumed hardness of lattice problems such as the shortest vector problem (SVP) and its variants. For efficiency, many practical lattice-based cryptosystems are based on assumptions on structured lattices such as the NTRU lattice. Introduced by Hoffstein, Piper and Silverman [HPS96,HPS98], the NTRU assumption states that it is hard to find a short vector in the \mathcal{R} -module

$$\Lambda_h^q = \{(x, y) \in \mathcal{R}^2 \text{ s.t. } hx - y = 0 \bmod q\}$$

with the promise that a very short solution —the private key— (f, g) exists. The ring $\mathcal{R} = \mathbb{Z}[X]/(P(X))$ is a polynomial ring of rank n over \mathbb{Z} , typically a circular convolution ring ($P(X) = X^n - 1$) or the ring of integers in a cyclotomic number field ($P(X) = \Phi_m(X)$ and $n = \phi(m)$).

Following the pioneer scheme NTRUENCRYPT [HPS98], the NTRU assumption has been re-used in various cryptographic constructions such as signatures schemes [HHGP⁺03,DDLL13], fully homomorphic encryption [LTV12,BLLN13] and a candidate construction for cryptographic multi-linear maps [GGH13a,LSS14,ACLL15]. After two decades of cryptanalysis, the NTRUENCRYPT scheme remains essentially unbroken, and is one of the fastest candidates for the public-key cryptosystems in the post-quantum era.

Coppersmith and Shamir [CS97] noticed that recovering a short enough vector, may it be different from the actual secret key (f, g) , may be sufficient for an attack and claimed that the celebrated LLL algorithm of Lenstra, Lenstra and Lovász [LLL82] would lead to such an attack. However, it turned out [HPS98] that for sufficiently large dimension n , a much stronger lattice reduction is required and that the NTRUENCRYPT is asymptotically secure. Meanwhile, parameters have been updated to take account for progress in lattice reduction algorithms and potential quantum speed-ups [HPS⁺15].

Other types of attacks have been considered, such as Odlyzko’s meet-in-the-middle attack described in [HSW06]. In practice, the best known algorithm for attacking NTRU lattices is the combined lattice-reduction

The full version of the paper is available on <http://eprint.iacr.org/2016/127>.

and meet-in-the-middle attack of Howgrave-Graham [HG07]. Asymptotically, a slightly sub-exponential attack against the ternary-NTRU problem was proposed by Kirchner and Fouque [KF15], with a heuristic complexity $2^{\Theta(n/\log \log q)}$, which is to our knowledge the only sub-exponential attack when q is polynomial in n .

It is typically assumed that NTRU lattices are essentially as intractable as unstructured lattices with similar parameters⁴, but without the structure of \mathcal{R} -module.

In the present work, we consider the application of lattice reduction in a *subfield* to attack the NTRU assumption for large moduli q . This subfield lattice attack is asymptotically faster than the direct lattice attack as soon as q is super-polynomial, and may also be relevant for polynomially-sized q . We call the problem⁵ considered in this work “overstretched NTRU” to distinguish it from the original NTRU parameter choices, which remain secure.

Asymptotics. The subfield attack leads to solving overstretched NTRU instances in time complexity $\text{poly}(n) \cdot 2^{\Theta(\beta)}$ with $\beta/\log \beta = \Theta(n \log n / \log^2 q)$ when ever the relative degree parameter $r = \Theta(\log q / \log n)$ is greater than 1. In comparison, the direct lattice attack required setting $\beta/\log \beta = \Theta(n/\log q)$.

We are mostly concerned with overstretched NTRU assumptions when q is super-polynomial in n , in which case the best known attacks are already sub-exponential in n . For cryptographic relevance, we will therefore state all our asymptotics in terms of what was previously thought as the security parameter λ : given $q = q(\lambda)$ we constrain $n = n(\lambda)$ so that the previously best known attack requires exponential time $2^{\Theta(\lambda)}$. In this cryptographic metric, the subfield lattice attack is sub-exponential as soon as q is super-polynomial, and gets polynomial for larger parameters $q = 2^{\Theta(\lambda)} = 2^{\Theta(\sqrt{n})}$.

Our contribution. In this work, we resurrect⁶ the subfield lattice attack sketched in [GS02, Sec. 6], attributed to Gentry, Szydlo, Jonsson, Nguyen and Stern. It consists of norming down the secret key to a subfield, run-

⁴ Volume, dimension and length of unusually short vectors.

⁵ The NTRU problem has also been recently been referred to as DSPR (Decisional Small Polynomial Ratio), but we prefer its historical name for fair attribution of this invention.

⁶ A preliminary version of this work qualified the attack considered in this work as new. We are grateful to John Schanck for pointing us to this prior art.

ning lattice reduction in the subfield to solve a smaller, potentially easier lattice problem and lifting the solution back to the full field.

While the original sketch [GS02] only considered the maximal real subfield, we naturally generalize it to any subfield. We also spell out a different lifting step from arbitrary subfields and prove it applicable even if only an approximation of the normed-down key is found.

We then show that this algorithm solves the overstretched NTRU problem in sub-exponential time when the modulus q is quasi-polynomial in the security parameter λ and in polynomial time when the modulus q is super-exponential in λ (equivalently, $q = 2^{\tilde{\Theta}(\sqrt{n})}$). Applying this algorithm, we show that it gives a subexponential attack on parameter choices for NTRU-based FHE schemes [LTV12,BLLN13] which were believed secure previously. We also show that this algorithm enables new attacks on GGH-like graded encoding schemes [GGH13a,LSS14,ACLL15]. These attacks lead to subexponential classical and polynomial-time quantum attacks on GGH-like constructions but do not require encodings of zero nor do they use the zero-testing parameter in contrast to previous work [HJ15].

We also report on experimental results for the subfield lattice attack which show that the attack is meaningful in practice. Using LLL in dimension 512 we have obtained vectors that would have required running BKZ with block-size about 130 in dimension 8192. We refer the reader to the full version of this work for the experimental results.

Related work. As mentioned above, a variant of the attack considered in this work was sketched in [GS02]. Moreover, the Gentry-Szydlo algorithm from the same work, which allows to reconstruct an element a given the ideal (a) as well as the Gram element $a\bar{a}$, i.e. the norm $N_{\mathbb{K}/\mathbb{K}^+}(a)$ of a relatively to the real subfield, can be seen as a subfield attack. It lead to an attack of the NSS scheme [HPS01] in which the Gram element $a\bar{a}$ was leaked as the covariance of a certain function of the signatures. The Gentry-Szydlo algorithm was recently revisited [LS14].

This attack is very similar in spirit to an attack of Gentry [Gen01] against the NTRU-composite assumption which tackles NTRU problems over rings \mathcal{R} that can be written as direct products $\mathcal{R} \simeq \mathcal{R}_1 \times \mathcal{R}_2$. More specifically [Gen01] targets circulant convolution rings $\mathbb{Z}[X]/(X^n - 1) \simeq \mathbb{Z}[X]/(X^{n_1} - 1) \times \mathbb{Z}[X]/(X^{n_2} - 1)$ where $n = n_1 n_2$. Under such condition, there exists a projection $\pi : \mathcal{R} \rightarrow \mathcal{R}_1$ that is a ring homomorphism, and he showed that this projection could only increase the Euclidean length of secret polynomials by a factor $\sqrt{n_2}$. This makes this attack very powerful

(even when the modulus q is quite small). Because this projection is a ring homomorphism, this approach is not limited to NTRU and would also apply to Ring-SIS or Ring-LWE.

In some sense, the line of work by Lauter et al. [ELOS15,EHL14,CLS15] against skewed⁷ variants of Ring-LWE falls in this framework, with a direct factorization of the rings \mathcal{R} modulo q : $(\mathcal{R}/q\mathcal{R}) \simeq (\mathcal{R}_1/q\mathcal{R}_1) \times (\mathcal{R}_2/q\mathcal{R}_2)$. As already noted in [Gen01], this requires the —seemingly sporadic— property that the projection map $\pi_q : (\mathcal{R}/q\mathcal{R}) \rightarrow (\mathcal{R}_1/q\mathcal{R}_1)$ induces only a manageable geometric distortion. Similar ideas are being explored to attack schemes based on certain quasi-cyclic binary codes in work [Loi14,LJ14,HT15].

In comparison, this work tackles NTRU when the ring \mathcal{R} equals $\mathcal{O}_{\mathbb{K}}$ (the ring of integer of a number field \mathbb{K}) and therefore cannot be a direct product; and when \mathbb{K} admits proper subfields. Due to the aforementioned attack of [Gen01], direct product rings are now avoided for lattice-based cryptography, and the typical choice is to use the ring of integers of a cyclotomic number field of the form $\mathcal{R} = \mathcal{O}_{\mathbb{Q}(\omega_m)} = \mathbb{Z}[\omega_m]$. This setting allows to argue worst-case hardness of certain problems (Ring-SIS [Mic02], Ideal-LWE [SSTX09], later improved and renamed to Ring-LWE [LPR10]). Yet all those number fields admit proper subfields (at least, the maximal real subfield). Instead of using a projection map π , this attack exploits a relative norm map $N_{\mathbb{K}/\mathbb{L}} : \mathcal{O}_{\mathbb{K}} \rightarrow \mathcal{O}_{\mathbb{L}}$, which is only a multiplicative map. This induces a significant yet manageable blow-up on the Euclidean length of secret polynomials and requires a large modulus q . This seems to also limit this attack to the NTRU setting.

Our work is also strongly inspired by the the logarithm-subfield strategy of Bernstein [Ber14], which anticipated other works towards a logarithm attack [CGS14,CDPR16]. While the presence of subfields was in the end not necessary for the recovery of short generators of principal ideals in cyclotomic rings, we show in this work that, indeed, the presence of proper subfields can be exploited in other specific set-ups.

Concurrently and independently to this work, Cheon, Jeong and Lee also investigated subfield attacks on GGH-like graded encoding schemes in work [CJL16]. The general approach is very similar to the one adopted in this work. In [CJL16], however, the trace map is utilised instead of the norm and the result is only presented for the case of powers-of-two cyclotomic rings. Despite using the trace map —which is linear— they

⁷ It was recently shown that these attacks were in fact made possible by an improper choice of a very skewed error distributions leading to several noise-free linear equations [CIV16,Pei16].

obtain a growth of the secret that is similar to ours: multiplicative. For example, when the relative degree of \mathbb{K} over \mathbb{L} is $r = 2$, the trace map $\text{Tr}_{\mathbb{K}/\mathbb{L}}$ sends g/f to $g/f + \bar{g}/\bar{f} = (g\bar{f} + \bar{g}f)/f\bar{f}$ where $\bar{\cdot}$ denotes the adequate automorphism. For comparison, the norm $N_{\mathbb{K}/\mathbb{L}}$ sends g/f to $g\bar{g}/f\bar{f}$. Using the norm map is therefore slightly better when both f, g have the same size (the numerator is smaller by a factor $\approx \sqrt{r}$); but the trace map could be very advantageous when $g \gg f$. Furthermore, Cheon, Jeong and Lee achieve better results for GGH-like graded encoding schemes by making use of the zero-testing parameter which leads to a polynomial-time classical attack for large levels of multilinearity κ .

Outline. Section 2 gives preliminaries on the geometry of NTRU lattices and a brief introduction of the lattice reduction algorithms. Section 3 then presents the subfield lattice attack with its asymptotic performance analyzed in Subsection 3.4. In Section 4, we apply this attack to the FHE and MLM constructions proposed in recent literature. In Section 5, we report experimental results for the subfield lattice attack. Finally, Section 6 presents the conclusions and suggests directions for future research.

Acknowledgments. We are grateful to Alice Silverberg, and to the participant of the Conference on Mathematics of Cryptography for enlightening talks and discussions. We thank Dan J. Bernstein, Ronald Cramer, Jeffrey Hoffstein, Hendrik W. Lenstra, John Schanck and Damien Stehlé for helpful discussions and comments.

We thank the PSMN (Pôle Scientifique de Modélisation Numérique, Lyon, France) for providing computing facilities.

2 Preliminaries

Vectors are presented in row vectors. The notation $[\cdot]_q$ denotes reduction modulo an integer q .

2.1 Number fields and subfields

We assume some familiarity with basic algebraic number theory. The reader may refer to [Sam70] for an introduction on the topic.

Let \mathbb{K} be a number field of degree $n = [\mathbb{K} : \mathbb{Q}]$ over \mathbb{Q} , and assume \mathbb{K} is a Galois extension of \mathbb{Q} with the Galois group G . The fundamental theorem of Galois Theory states an one-to-one correspondence between the subgroups G' of G and the subfields \mathbb{L} of \mathbb{K} with G' being the subgroup of G fixing \mathbb{L} . Let therefore \mathbb{L} be a subfield of \mathbb{K} and G' be the subgroup

of G fixing \mathbb{L} , and denote $n' = [\mathbb{L} : \mathbb{Q}]$, $r = [\mathbb{K} : \mathbb{L}]$ (so $r = n/n'$). The number fields \mathbb{K} , \mathbb{L} and therefore the degrees n , n' and relative degree r are fixed in the rest of this work.

The relative norm $N_{\mathbb{K}/\mathbb{L}} : \mathbb{K} \rightarrow \mathbb{L}$ (resp. relative trace $\text{Tr}_{\mathbb{K}/\mathbb{L}} : \mathbb{K} \rightarrow \mathbb{L}$) is a multiplicative (resp. an additive) map defined by

$$N_{\mathbb{K}/\mathbb{L}} : a \mapsto \prod_{\psi \in G'} \psi(a), \quad \text{resp.} \quad \text{Tr}_{\mathbb{K}/\mathbb{L}} : a \mapsto \sum_{\psi \in G'} \psi(a). \quad (1)$$

The canonical inclusion $\mathbb{L} \subset \mathbb{K}$ will be written explicitly as $L : \mathbb{L} \rightarrow \mathbb{K}$. The ring of integers of \mathbb{K} and \mathbb{L} are denoted by $\mathcal{O}_{\mathbb{K}}$ and $\mathcal{O}_{\mathbb{L}}$.

A number field of degree n admits n embeddings –i.e. field morphisms– to the complex numbers. Writing $\mathbb{K} = \mathbb{Q}(X)/(P(X))$ for some monic irreducible polynomial P , and letting $\alpha_1, \dots, \alpha_n \in \mathbb{C}$ be the distinct complex roots of P , each embedding $e_i : \mathbb{K} \rightarrow \mathbb{C}$ consists of evaluating $a \in \mathbb{K}$ at a root α_i , formally $e_i : a \mapsto a(\alpha_i)$. The Galois group acts by permutation on the set of embeddings.

Cyclotomic Number Field. We denote by ω_m an arbitrary primitive m -th root of unity. For cryptanalytic purposes, we are mostly interested in the case when $\mathbb{K} = \mathbb{Q}(\omega_m)$ is the m -th cyclotomic number field; But we may also want to instantiate the attack for subfields \mathbb{L} of \mathbb{K} that are not necessarily cyclotomic number fields.

The number field $\mathbb{L} = \mathbb{Q}(\omega_m)$ has degree $n = \phi(m)$, and has a Galois group isomorphic to \mathbb{Z}_m^* : explicitly $i \in \mathbb{Z}_m^*$ corresponds to the automorphism $\psi_i : \omega_m \mapsto \omega_m^i$. Any number field $\mathbb{Q}(\omega_{m'})$ for $m' | m$ is a subfield of $\mathbb{Q}(\omega_m)$, but there are other proper subfields. In particular, the maximal real subfield $\mathbb{Q}(\omega_m + \bar{\omega}_m)$ is a proper subfield of degree $n/2$, and more generally, $\mathbb{K} = \mathbb{Q}(\omega_m)$ admits a subfield of degree n' for any divisor $n' | n$.⁸

We recall (see [Was97], Theorem 2.6) that the ring of integers $\mathcal{O}_{\mathbb{K}}$ of $\mathbb{K} = \mathbb{Q}(\omega_m)$ is exactly $\mathbb{Z}[\omega_m]$.

2.2 Coprimality in $\mathcal{O}_{\mathbb{L}}$

To argue below that we can lift solutions in the subfield to the full field, we rely on two randomly chosen elements in $\mathcal{O}_{\mathbb{L}}$ being coprime. We use density results to estimate such probability. The density of coprime pairs of ideals [Sit10] and elements [FM14] in $\mathcal{O}_{\mathbb{L}}$ is $1/\zeta_{\mathbb{L}}(2)$ where $\zeta_{\mathbb{L}}$ denotes the Dedekind zeta function over \mathbb{K} .

⁸ For example, 7 is prime, so $\mathbb{Q}(\omega_7)$ admits no cyclotomic number fields as proper subfields, yet it admits two proper subfields: $\mathbb{Q}(\omega_7 + \bar{\omega}_7)$ of degree 3 and $\mathbb{Q}(\omega_7 + \omega_7^2 + \omega_7^4)$ of degree 2.

We consider $\zeta_{\mathbb{L}}$ for cyclotomic number fields $\mathbb{K} = \mathbb{Q}(\omega_m)$ where $m = p^k$ for some prime p . The next lemma shows that $\lim_{k \rightarrow \infty} \zeta_{\mathbb{L}}(s) = 1/(1 - p^{-s})$ for real $s > 3/2$.

Lemma 1. *Let \mathbb{L} be a cyclotomic number field $\mathbb{Q}(\omega_{m'})$ for $m' = p^k$. Then for any real $s > 3/2$ we have*

$$\lim_{k \rightarrow \infty} \zeta_{\mathbb{L}}(s) = 1/(1 - p^{-s}).$$

In particular $\lim_{k \rightarrow \infty} \zeta_{\mathbb{L}}(2) = 4/3$ for cyclotomic number fields of conductor $m' = 2^k$.

Proof. Please refer to the full version of this work for the proof. \square

Further, we numerically approximated $\zeta_{\mathbb{L}}^{-1}(2)$ for $\mathbb{L} = \mathbb{Q}[x]/(x^n + 1)$ for $n = 128$ and $n = 256$ by computing the first 2^{22} terms of the Dirichlet series of the Dedekind zeta function for \mathbb{L} and then evaluated the truncated series at 2. In both cases we get a density ≈ 0.75 .

We stress that our pairs f', g' are random elements obtained as relative norms $N_{\mathbb{K}/\mathbb{L}}(f), N_{\mathbb{K}/\mathbb{L}}(g)$ of random *short* f and g , and under the additional condition that f is invertible modulo q . However, our experiments indicate that $3/4$ is a good approximation of the actual probability of coprimality. Additionally, it seems that this requirement is an artifact of our proof, as experiments succeeded even when those elements had a common factor.

2.3 Euclidean geometry

The number field \mathbb{K} (or \mathbb{L}) is viewed as a Euclidean \mathbb{Q} -vector space by endowing it with the inner product

$$\langle a, b \rangle = \sum_e e(a)\bar{e}(b) \quad (2)$$

where e ranges over all the n (or n') embeddings $\mathbb{K} \rightarrow \mathbb{C}$. This defines a Euclidean norm denoted by $\|\cdot\|$. In addition to the Euclidean norm, we will make use of the operator norm $|\cdot|$ defined by:

$$|a| = \sup_{x \in \mathbb{K}^*} \|ax\|/\|x\|. \quad (3)$$

It is easy to check that the operator norm $|a|$ of a equals to the maximal absolute complex embedding of a :

$$|a| = \max_e |e(a)| \quad (4)$$

where e ranges over all the embeddings $e : \mathbb{K} \rightarrow \mathbb{C}$. We note that if $\omega \in \mathbb{K}$ is a root of unity, then $|\omega| = 1$. The operator's norm is sub-multiplicative: $|ab| \leq |a| |b|$, and we have the inequality $|a| \leq \|a\|$. The Euclidean norm and the operator norm are invariant under automorphisms $\psi : \mathbb{K} \mapsto \mathbb{K}$,

$$\|a\| = \|\psi(a)\|, \quad |a| = |\psi(a)| \quad (5)$$

since the group of automorphisms acts by permutation on the set of embeddings. One also verifies that $\|L(a)\|^2 = r\|a\|^2$ and $|L(a)| = |a|$ for all $a \in \mathbb{L}$. Additionally, the algebraic norm can be bounded in term of geometric norms:

$$N_{\mathbb{K}/\mathbb{Q}}(a) \leq |a|^n \leq \|a\|^n. \quad (6)$$

The inner product (and therefore the Euclidean norm) are extended in a coefficient-wise manner to vectors of \mathbb{K}^d : $\langle (a_1, \dots, a_d), (b_1, \dots, b_d) \rangle = \sum \langle a_i, b_i \rangle$.

Definition 1. A distribution \mathcal{D} over \mathbb{K}^d is said to be isotropic of variance $\sigma^2 \geq 0$ if, for any $y \in \mathbb{K}^d$ it hold that

$$\mathbb{E}_{x \leftarrow \mathcal{D}} [\langle x, y \rangle^2] = \sigma^2 \|y\|^2$$

where $\mathbb{E}[\cdot]$ denotes the expectation of a random variable.

Remark. In most theoretical work, the distributions of secrets or errors are spherical discrete Gaussian distribution over $\mathcal{O}_{\mathbb{K}}$ which are isotropic —up to negligible statistical distance. For simplicity, some practically oriented work instead chose random ternary coefficients. In the typical power-of-two case cyclotomic case, such distribution is isotropic of variance $2n/3$. Yet, for more general choices $\mathbb{K} = \mathbb{Q}(\omega_m)$, in the worse case (when m is composed of many small distinct prime factor), this may induce up to quasi-polynomial distortion $n^{\log(n)}$ (see [LPR10]). Such choice of set-up should only marginally affect our asymptotic results.

2.4 $\mathcal{O}_{\mathbb{K}}$ modules and lattices

To avoid confusion, we shall speak of the rank of $\mathcal{O}_{\mathbb{K}}$ -modules and of \mathbb{K} -vectors-spaces when $\mathbb{K} \neq \mathbb{Q}$, and restrict the term of dimension to \mathbb{Z} -modules and \mathbb{Q} -vector spaces.

The dimension $\dim(A)$ of a lattice A is the dimension over \mathbb{Q} of the \mathbb{Q} -vector space it spans⁹. We recall that the minimal distance of a lattice A

⁹ Or equivalently, the size of a minimal sets of \mathbb{Z} -generators, since \mathbb{Z} is a principal ideal domain.

is defined as $\lambda_1(A) = \min_{v \in A \setminus \{0\}} \|v\|$. Also, the volume of a lattice $\text{Vol}(A)$ is defined as the square root of the absolute determinant of the Gram matrix of any basis $\{b_1 \dots b_{\dim(A)}\}$ of A $\text{Vol}(A) = \sqrt{\det([\langle b_i, b_j \rangle]_{i,j})}$. For any set of \mathbb{Q} -linearly independent vectors $\{v_1, \dots, v_{\dim(A)}\} \subset A$, we have the inequality:

$$\text{Vol}(A) \leq \prod \|v_i\|. \quad (7)$$

The rank of an $\mathcal{O}_{\mathbb{K}}$ module $M \subset \mathbb{K}^d$ can be defined as the rank over \mathbb{K} of the \mathbb{K} vector-space it spans, but it does not necessarily equal the size of a minimal set of $\mathcal{O}_{\mathbb{K}}$ -generators¹⁰. The Euclidean vector space structure of \mathbb{K}^d allows to view any discrete $\mathcal{O}_{\mathbb{K}}$ -module $M \subset \mathbb{K}^d$ as a lattice. The discriminant $\Delta_{\mathbb{K}}$ of a number field relates to the volume of its ring of integers $\sqrt{|\Delta_{\mathbb{K}}|} = \text{Vol}(\mathcal{O}_{\mathbb{K}})$. More generally, we have the identity:

$$\text{Vol}(a\mathcal{O}_{\mathbb{K}}) = N_{\mathbb{K}/\mathbb{Q}}(a)\sqrt{|\Delta_{\mathbb{K}}|}. \quad (8)$$

This gives rise to a lower bound on the volume $\mathcal{O}_{\mathbb{K}}$ -modules of rank 1 in term of its minimal distance:

Lemma 2. *Let $M \subset \mathbb{K}^d$ be a discrete $\mathcal{O}_{\mathbb{K}}$ -module of rank 1. It follows that $\text{Vol}(M) \leq \lambda_1(M)^n \sqrt{|\Delta_{\mathbb{K}}|}$.*

Proof. Without loss of generality, we may assume that $d = 1$ (by constructing a \mathbb{K} -linear isometry $\iota : \text{Span}_{\mathbb{K}}(M) \rightarrow \mathbb{K} \otimes_{\mathbb{Q}} \mathbb{R}$). Let $a \in \mathbb{K} \otimes_{\mathbb{Q}} \mathbb{R}$ be a shortest vector of M , we have $M \supset a\mathcal{O}_{\mathbb{K}}$, therefore $\text{Vol}(M) \leq \text{Vol}(a\mathcal{O}_{\mathbb{K}}) = N_{\mathbb{K}/\mathbb{Q}}(a)\sqrt{|\Delta_{\mathbb{K}}|}$, and we conclude noting that $N_{\mathbb{K}/\mathbb{Q}}(a) \leq \|a\|^n$. \square

2.5 NTRU assumption

Let us first describe the NTRU problem as follows.

Definition 2 (NTRU problem, a.k.a. DSPR). *The NTRU problem is defined by four parameters: a ring \mathcal{R} (of rank n and endowed with an inner product), a modulus q , a distribution \mathcal{D} , and a target norm τ . Precisely, $\text{NTRU}(\mathcal{R}, q, \mathcal{D}, \tau)$ is the problem of, given $h = [gf^{-1}]_q$ (conditioned on f being invertible mod q) for $f, g \leftarrow \mathcal{D}$, finding a vector $(x, y) \in \mathcal{R}^2$ such that $(x, y) \neq (0, 0) \bmod q$ and of Euclidean norm less than $\tau\sqrt{2n}$ in the lattice*

$$\Lambda_h^q = \{(x, y) \in \mathcal{R}^2 \text{ s.t. } hx - y = 0 \bmod q\}. \quad (9)$$

¹⁰ Non-principal ideals of \mathbb{K} being a counter-example.

We may abuse notation and denote $\text{NTRU}(\mathcal{R}, q, \sigma, \tau)$ for $\text{NTRU}(\mathcal{R}, q, \mathcal{D}, \tau)$ where \mathcal{D} is any reasonable isotropic distribution of variance σ^2 .

Note that $\text{NTRU}(\mathcal{R}, q, \sigma, \sigma)$ is essentially the problem of recovering the secret key (f, g) . Yet, in many cases, solving $\text{NTRU}(\mathcal{R}, q, \sigma, \tau)$ for some $\tau > \sigma$ is enough to break NTRU-like cryptosystems.

The NTRU lattice Λ_h^q . The lattice Λ_h^q defined by the instance $h \leftarrow \text{NTRU}(\mathcal{O}_{\mathbb{K}}, q, \sigma, \tau)$ has dimension $2n$ and volume $\text{Vol}(\mathcal{R})^2 q^n$. Consequently, if h were to be uniformly random, the Gaussian heuristic predicts that the shortest vectors of Λ_h^q have norm $\text{Vol}(\mathcal{R})^{1/n} \sqrt{nq/\pi e}$. Therefore, whenever $\sigma < \text{Vol}(\mathcal{R})^{1/n} \sqrt{q/2\pi e}$, the lattice Λ_h^q admits an *unusually short vector*. This vector is not formally a unique shortest vector: for example, if $\mathbb{K} = \mathbb{Q}(\omega_m)$, $\mathcal{R} = \mathcal{O}_{\mathbb{K}}$, all rotations $(\omega_m^i f, \omega_m^i g)$ of that vector have the same norm.

Target parameter τ for attacks. Because no solution would be expected if h was uniformly random, note that solving $h \leftarrow \text{NTRU}(\mathcal{R}, q, \sigma, \tau)$ for $\tau < \text{Vol}(\mathcal{R})^{1/n} \sqrt{q/2\pi e}$ already constitutes a distinguishing attack on the NTRU problem. As we discuss in Section 4, solving NTRU for such τ would break the FHE scheme based on NTRU from [LTV12] and typical parameter choices for the scheme presented in [BLLN13].

2.6 Lattice reduction algorithms

Lattice reduction algorithms have been studied for many years in work such as [LLL82, Sch87, GN08, HPS11]. From a theoretical perspective, one of the best lattice reduction algorithm is the slide reduction algorithm from [GN08].

Theorem 1 ([GN08]). *There is an algorithm that, given $\epsilon > 0$, the basis B of a lattice L of dimension d , and performing at most*

$$\text{poly}(d, 1/\epsilon, \text{bitsize}(B))$$

many operations and calls to an SVP oracle in dimension β , outputs a vector $v \in L$ whose length satisfies the following bounds:

- *the approximation-factor bound:*

$$\|v\| \leq ((1 + \epsilon)\gamma_\beta)^{\frac{d-\beta}{\beta-1}} \cdot \lambda_1(L) \quad (10)$$

where $\lambda_1(L)$ is the length of a shortest vector in L and $\gamma_\beta \approx \beta$ is the β -dimensional Hermite constant.

– the Hermite-factor bound:

$$\|v\| \leq ((1 + \epsilon)\gamma_\beta)^{\frac{d-1}{2\beta-2}} \cdot \text{Vol}(L)^{1/d} \quad (11)$$

Alternatively, one may use the BKZ algorithm [Sch87] and its terminated variant [HPS11]. Similar to slide reduction, the terminated BKZ performs at most $\text{poly}(d, 1/\epsilon, \text{bitsize}(B))$ many operations and calls to an SVP oracle in dimension β ; and outputs a vector $v \in L$ whose length has order $\beta^{\Theta(n/\beta)} \cdot \text{Vol}(L)^{1/d}$. Using [Lov87, p. 25], the terminated BKZ also provides an algorithm to find an approximated shortest vector of length $\beta^{\Theta(n/\beta)} \cdot \lambda_1(L)$ in similar time.

It is well known [CN11] that in practice lattice reduction algorithms achieve much shorter results and are more efficient, but the approximation and Hermite factors remain of the order of $\beta^{\Theta(n/\beta)}$ asymptotically, for a computational cost in $\text{poly}(\lambda) \cdot 2^{\Theta(\beta)}$. We will use such estimate in the following analysis.

3 The subfield lattice attack

The subfield lattice attack works in three steps. First, we map the NTRU instance to the chosen subfield, then we apply lattice reduction, and finally we lift the solution to the full field. We first describe the three steps of the attacks in Sections 3.1, 3.2 and 3.3. In Section 3.4, we then analyze the asymptotic performances compared to direct reduction in the full field for cryptographically relevant asymptotic parameters.

We are given an instance $h \leftarrow \text{NTRU}(\mathcal{O}_{\mathbb{K}}, q, \sigma, \tau)$, and $(f, g) \in \mathcal{O}_{\mathbb{K}}$ is the associated secret. We wish to recover a short vector of Λ_h^q .

3.1 Norming down

We define $f' = N_{\mathbb{K}/\mathbb{L}}(f)$, $g' = N_{\mathbb{K}/\mathbb{L}}(g)$, and $h' = N_{\mathbb{K}/\mathbb{L}}(h)$. The subfield attack follows from the following observation: (f', g') is a vector of $\Lambda_{h'}^q$ and depending on the parameters it may be an unusually short one.

Lemma 3. *Let $f, g \in \mathcal{O}_{\mathbb{K}} \otimes_{\mathbb{Q}} \mathbb{R}$ be sampled from continuous spherical Gaussians of variance σ^2 . For any constant $c > 0$, there exists a constant C , such that,*

$$\|g'\| \leq (\sigma n^C)^r, \quad \|f'\| \leq (\sigma n^C)^r, \quad |f'| \leq (\sigma n^C)^r, \quad |f'^{-1}| \leq (n^C/\sigma)^r$$

except with probability $O(n^{-c})$.

Proof. For all embeddings $e : \mathbb{K} \mapsto \mathbb{C}$, it simultaneously holds that

$$\sigma/n^C \leq |e(f)| \leq \sigma n^C \quad (12)$$

except with polynomially small probability $O(n^{-c})$. Once this is established, the conclusion follows using the invariant $|\psi(a)| = |a|$ since $f' = \prod \psi(f)$, where ψ ranges over r automorphisms of \mathbb{K} .

To prove inequality (12), note that for each embedding e , the $\Re(e(f))$ and $\Im(e(f))$ follow a Gaussian distribution of parameter $\Theta(n)\sigma$. Classical tails inequality gives the upper bound $|e(f)| \leq \sigma n^C$. For the lower bound, we remark that the probability density function of a Gaussian of parameter $\Theta(n)\sigma$ is bounded by $1/(\Theta(n)\sigma)$. This implies that the probability that a sample falls in the range $\frac{1}{\Theta(n)\sigma}[-\epsilon, \epsilon]$ is less than 2ϵ . It remains to choose $\epsilon = \Theta(n^{-c-1})$ which gives the conclusion by the union-bound. \square

In this work, we assume that Lemma 3 holds also for all reasonable distributions considered in cryptographic constructions.

Heuristic 1 *For any m and any $f, g \in \mathcal{O}_{\mathbb{K}}$ with reasonable isotropic distribution of variance σ^2 , and any constant $c > 0$, there exists a constant C , such that,*

$$\|g'\| \leq (\sigma n^C)^r, \quad \|f'\| \leq (\sigma n^C)^r, \quad |f'| \leq (\sigma n^C)^r, \quad |f'^{-1}| \leq (n^C/\sigma)^r$$

except with probability $O(n^{-c})$.

3.2 Lattice reduction in the subfield

We now apply a lattice reduction algorithm with block-size β to the lattice $\Lambda_{h'}^q$, and according to the approximation factor bound (10) we obtain a vector $(x', y') \in \Lambda_{h'}^q$ of norm:

$$\|(x', y')\| \leq \beta^{\Theta(2n'/\beta)} \cdot \lambda_1(\Lambda_{h'}^q) \leq \beta^{\Theta(n/\beta r)} \cdot \|(f', g')\| \quad (13)$$

$$\leq \beta^{\Theta(n/\beta r)} \cdot (n\sigma)^{\Theta(r)}. \quad (14)$$

Next, we argue that if the vector (x', y') is short enough, then it must be an $\mathcal{O}_{\mathbb{K}}$ -multiple of (f', g') . In turn, this will allow us to lift (x', y') to a short vector in the full lattice Λ_h^q .

Theorem 2. *Let $f', g' \in \mathcal{O}_{\mathbb{L}}$ be such that $\langle f' \rangle$ and $\langle g' \rangle$ are coprime ideals and that $h'f' = g' \bmod q\mathcal{O}_{\mathbb{L}}$ for some $h' \in \mathcal{O}_{\mathbb{L}}$. If $(x', y') \in \Lambda_{h'}^q$ has length satisfying*

$$\|(x', y')\| < \frac{q}{\|(f', g')\|} \quad (15)$$

then $(x', y') = v(f', g')$ for some $v \in \mathcal{O}_{\mathbb{L}}$.

Proof. We first prove that that $B = \{(f', g'), (F', G')\}$ is a basis of the $\mathcal{O}_{\mathbb{L}}$ -module $\Lambda_{h'}^q$ for some $(F', G') \in \mathcal{O}_{\mathbb{L}}^2$. The argument is adapted from [HHGP⁺03], Section 4.1. By coprimality, there exists (F', G') such that $f'G' - g'F' = q \in \mathcal{O}_{\mathbb{L}}$. We note that:

$$\begin{aligned} f'(F', G') - F'(f', g') &= (0, q); \\ g'(F', G') - G'(f', g') &= (-q, 0); \\ [f'^{-1}]_q(f', g') &= (1, h') \bmod q. \end{aligned}$$

That is, the module M generated by B contains $q\mathcal{O}_{\mathbb{L}}^2$ and $(1, h')$: we have proved that $\Lambda_{h'}^q \subset M$. Because $\det_{\mathbb{L}}(B) = f'G' - g'F' = q = \det_{\mathbb{L}}(\{(1, h'), (0, q)\})$ we have $\text{Vol}(M) = |\Delta_{\mathbb{L}}|q^{n'} = \text{Vol}(\Lambda_{h'}^q)$ and therefore $M = \Lambda_{h'}^q$.

We denote $\Lambda = (f', g')\mathcal{O}_{\mathbb{L}}$ and Λ^* the projection of $(F', G')\mathcal{O}_{\mathbb{L}}$ orthogonally to Λ . Let s^* of length λ_1^* be a shortest vector of Λ^* . We will conclude using the fact that any vector of $\Lambda_{h'}^q$ of length less than λ_1^* must belong to the sublattice Λ . It remains to give an lower bound for λ_1^* .

We will rely on the identity $\text{Vol}(\Lambda) \cdot \text{Vol}(\Lambda^*) = \text{Vol}(\Lambda_{h'}^q) = |\Delta_{\mathbb{L}}|q^{n'}$. By Lemma 2, we have

$$\text{Vol}(\Lambda) \leq |\Delta_{\mathbb{L}}|^{1/2} \|(f', g')\|^{n'} \quad \text{and} \quad \text{Vol}(\Lambda^*) \leq |\Delta_{\mathbb{L}}|^{1/2} \|s^*\|^{n'}. \quad (16)$$

We deduce that $\lambda_1^* = \|s^*\| \geq q/\|(f', g')\|$. Therefore, the hypothesis (15) ensures that $\|(x', y')\| < \lambda_1^*$, and we conclude that $(x', y') \in \Lambda = (f', g')\mathcal{O}_{\mathbb{L}}$. \square

We note that according to Heuristic 1, the length condition of Theorem 2 are satisfied asymptotically when

$$\beta^{\Theta(n/\beta r)} \cdot (n\sigma)^{\Theta(r)} \leq q. \quad (17)$$

The probability of satisfying the coprimality condition for random f', g' is discussed in Section 2.2, where we argue it to be larger than a constant. On the other hand, experiments (cf. Section 5) show that the coprimality condition does not seems necessary in practice for the subfield lattice attack to succeed.

The partial conclusion is that, one may recover non-trivial information about f and g — namely, a small multiple of (f', g') — by solving an NTRU instance in a subfield. Depending on the parameters, this new problem is potentially easier since the dimension $n' = n/r$ of $\mathcal{O}_{\mathbb{L}}$ is significantly smaller than the dimension $2n$ of the full lattice Λ_h^q .

3.3 Lifting the short vector

It remains to lift the solution from the sub-ring $\mathcal{O}_{\mathbb{L}}$ to $\mathcal{O}_{\mathbb{K}}$. Simply compute the vector (x, y) where

$$x = L(x') \quad \text{and} \quad y = L(y') \cdot h / L(h') \bmod q \quad (18)$$

where $L : \mathbb{L} \rightarrow \mathbb{K}$ is the canonical inclusion map of $\mathbb{L} \subset \mathbb{K}$.

Recall from Theorem 2 that $(x', y') = v(f', g')$. We set $\tilde{f} = L(f')/f$, $\tilde{g} = L(g')/g$ and $\tilde{h} = L(h')/h$. Note that \tilde{f}, \tilde{g} and \tilde{h} are integers of \mathbb{K} . We rewrite

$$\begin{aligned} x &= L(v) \cdot \tilde{f} \cdot f \bmod q. \\ y &= L(v) \cdot L(g')/\tilde{h} = L(v) \cdot g\tilde{g}/\tilde{h} \bmod q \\ &= L(v) \cdot \tilde{f} \cdot g \bmod q. \end{aligned}$$

That is, under condition (17) we have found a short multiple of (f, g) :

$$\begin{aligned} (x, y) &= u \cdot (f, g) \in \Lambda_h^q \quad \text{with } u = L(v) \cdot \tilde{f} \in \mathcal{O}_{\mathbb{K}} \\ \|(x, y)\| &\leq |v| \cdot |f|^{r-1} \cdot \|(f, g)\| \\ &\leq |x'| \cdot |f'|^{-1} \cdot |f|^{r-1} \cdot \|(f, g)\| \\ &\leq \beta^{\Theta(n/\beta r)} \cdot (n\sigma)^{\Theta(r)}. \end{aligned}$$

The first inequality is established by writing \tilde{f} as the product of $r - 1$ many $\psi(f)$ where the ψ 's are automorphisms of \mathbb{K} . The second inequality decomposes $v = x'/f'$, and the last follows from Lemma 3 or Heuristic 1.

Not only we have found a short vector of Λ_h^q , but also have the guarantee that it is an $\mathcal{O}_{\mathbb{K}}$ -multiple of the secret key (f, g) . This second property will prove useful to mount attacks on the graded encoding schemes [GGH13a].

3.4 Asymptotic performance

For the subfield attack to be successful, we require

$$\sqrt{q} = \beta^{\Theta(2n/(\beta r))} \cdot \lambda_1(\Lambda_{h'}^q) = \beta^{\Theta(2n/(\beta r))} \cdot n^{\Theta(r)}$$

when $\sigma = \text{poly}(n)$. Hence, asymptotically we get

$$\frac{\beta}{\log \beta} = \Theta\left(\frac{4n}{r \log q - 2r^2 \log n}\right),$$

where we require $r \log q - 2r^2 \log n > 0$. Setting $r = 1$ roughly recovers the lattice attack in the full field. Setting $r = \log q / (4 \log n)$ minimizes the expression.

We illustrate the complexity for two extreme cases, where all parameters are expressed in term of a security parameter λ , and are such that the previously best known attack required time greater than 2^λ . Additionally, it is assumed that \mathbb{K} contains adequate subfields so that a subfield \mathbb{L} of the desired relative degree r exists. This condition is satisfied asymptotically for the typical choice $\mathbb{K} = \mathbb{Q}(\omega_{2^k})$.

In the first case, we set $q = 2^{\hat{\Theta}(\lambda)}$, and the subfield attack is polynomial in the security parameter. For the second case, we show that as soon as q gets super-polynomial, the subfield attack can be made sub-exponential.

Remark. Our analysis does not rule out that the attack may even be relevant even for polynomial gaps q/σ : it could be that it remains exponential but with a better constant than the direct attack.

Exponential and super-exponential q . We set:

$$n = \Theta(\lambda^2 \log^2 \lambda), \quad q = \exp(\Theta(\lambda \log^2 \lambda)), \quad \sigma = \text{poly}(\lambda). \quad (19)$$

Complexity of the direct lattice attack. With such parameters, using 2^λ operations, we argue that one may not find any vector shorter than $\lambda_1(q\mathcal{O}_{\mathbb{K}}) = q\sqrt{n}$. Indeed, one may run lattice reduction up to block-size $\beta = \Theta(\lambda)$. Either from approximation bound or Hermite bound, the vector found should not be shorter than:

$$\beta^{\Theta(n/\beta)} = \exp(\Theta(\lambda^2 \log^3(\lambda)/\lambda)) > \lambda_1(q\mathcal{O}_{\mathbb{K}}). \quad (20)$$

We verify that having such choice of super-quadratic n makes the Kirchner-Fouque [KF15] attack at least exponential in λ : $\exp(\Theta(n/\log \log q)) = \exp(\Theta(\lambda^2 \log^2(\lambda)/\log \lambda)) > \exp(\Theta(\lambda))$.

Complexity of the subfield attack. In contrast, the same parameters allow the subfield attack to recover a vector of norm less than \sqrt{q} in polynomial time: set $r = \Theta(\lambda)$ and $\beta = \Theta(\log \lambda)$. Then, the vector found will have norm

$$\beta^{\Theta(n/\beta r)} \cdot n^{\Theta(r)} = \exp\left(\Theta\left(\frac{\lambda^2 \log \lambda \log \log \lambda}{\lambda \log \lambda} + \lambda \log \lambda\right)\right) \quad (21)$$

$$= \exp(\Theta(\lambda \log \lambda \log \log \lambda)) < \sqrt{q}. \quad (22)$$

Similarly, setting $n = \Theta(\lambda^2)$, $q = \exp(\Theta(\lambda))$, $\beta = \Theta(\log^{1+\varepsilon} \lambda)$, $r = \Theta(\lambda / (\log \lambda \log \log \lambda))$ leads to a quasi-polynomial version of the subfield attack for exponential q .

Quasi-polynomial q . We set

$$n = \Theta(\lambda \log^\varepsilon \lambda \log \log(\lambda)), \quad q = \exp(\Theta(\log^{1+\varepsilon} \lambda)), \quad \sigma = \text{poly}(\lambda).$$

Complexity of the direct lattice attack. With such parameters, using 2^λ operations, we argue that one may not find any vector shorter than $\lambda_1(q\mathcal{O}_{\mathbb{K}}) = q\sqrt{n}$. Indeed, one may run lattice reduction up to block-size $\beta = \Theta(\lambda)$. Either from approximation bound or Hermite bound, the vector found should not be shorter than:

$$\beta^{\Theta(n/\beta)} = \exp(\Theta(\log^{1+\varepsilon} \lambda \log \log \lambda)) > \lambda_1(q\mathcal{O}_{\mathbb{K}}). \quad (23)$$

We verify that having such choice of super-linear n makes the Kirshner-Fouque [KF15] attack at least exponential in λ : $\exp(\Theta(n/\log \log q)) = \exp(\Theta(\lambda \log^\varepsilon \lambda \log \log \lambda / \log \log^{1+\varepsilon} \lambda)) > \exp(\Theta(\lambda))$.

Complexity of the subfield attack. In contrast, the same parameters allow the subfield attack to recover a vector of norm less than \sqrt{q} in sub-exponential time $\exp(\lambda / \log^{\varepsilon/3} \lambda)$: set $r = \Theta(\log^{2\varepsilon/3} \lambda)$ and $\beta = \Theta(\lambda / \log^{\varepsilon/3} \lambda)$. Then, the vector found will have norm

$$\begin{aligned} \beta^{\Theta(n/\beta r)} \cdot n^{\Theta(r)} &= \exp\left(\Theta\left(\frac{\log^{1+\frac{4}{3}\varepsilon}(\lambda) \log \log(\lambda)}{\log^{\frac{2}{3}\varepsilon}(\lambda)} + \log^{1+2/3\varepsilon}(\lambda)\right)\right) \\ &= \exp\left(\Theta\left(\log^{1+2/3\varepsilon}(\lambda) \log \log(\lambda)\right)\right) < \sqrt{q}. \end{aligned} \quad (24)$$

4 Applications

We apply this attack to the FHE and MLM constructions from the literature and show that it necessitates to increase parameters for these schemes to remain secure at level λ . In the cryptographic context, we typically have $\mathbb{K} = \mathbb{Q}(\omega_m)$, m a power of 2, and speak of the ring $\mathcal{R} = \mathbb{Z}_q[X]/(X^n + 1) \simeq \mathcal{O}_{\mathbb{K}}$ endowed with the canonical inner product of its coefficients vector. The ring isomorphism $\mu : \mathcal{R} \rightarrow \mathcal{O}_{\mathbb{K}}$ is a scaled isometry: $\|\mu(x)\| = \sqrt{n}\|x\|$. This normalization is quite convenient, for example $\|1_{\mathcal{R}}\| = 1$.

4.1 Fully Homomorphic Encryption

NTRU-like schemes are used to realise fully homomorphic encryption starting with the LTV scheme from [LTV12]; the scheme was optimized and implemented in [DHS15].

LTV is motivated by [SS11] which shows that under certain choices of parameters the security of an NTRU-like scheme can be reduced to security of Ring-LWE. That is, [SS11] shows that if f and g have norms $> \sqrt{q} \cdot \text{poly}(\lambda)$, then $h = [g/f]_q \in \mathbb{Z}_q[X]/(X^n + 1)$ — with n a power of two — is statistically indistinguishable from a uniformly sampled element. Note that under this choice of parameters the subfield lattice attack does not apply.

However, this choice of parameters rules out even performing one polynomial multiplication and hence the schemes in [LTV12,DHS15] are based on an additional assumption that $[g/f]_q$ is computationally indistinguishable from random even when f and g are small. This assumption — which essentially states that Decisional-NTRU is hard — is called the Decisional Small Polynomial Ratio assumption (DSPR) in [LTV12]. Note that this work shows that DSPR does not hold in the presence of subfields and an overstretched NTRU assumption.

LTV can evaluate circuits of depth $L = \mathcal{O}(n^\varepsilon / \log n)$ for $q = 2^{n^\varepsilon}$ with $\varepsilon \in (0, 1)$ and its decryption circuit can be implemented in depth $\mathcal{O}(\log \log q + \log n)$. This implies

$$\begin{aligned} \log(n^{\varepsilon+1}) &< n^\varepsilon / \log n, \\ \log(n^{\varepsilon+1}) &< \log q / \log n, \end{aligned}$$

i.e. that q must be super-polynomial in n to realise fully homomorphic encryption from LTV.

A scale-invariant variant of the scheme in [LTV12] called YASHE was proposed in [BLLN13]. This variant does not require the DSPR assumption by reducing the noise growth during multiplication. This allows f and g to be sampled from a sufficiently wide Gaussian, such that the reduction in [SS11] goes through. Sampling f and g this way allows to evaluate circuits of depth $L = \mathcal{O}(\log q / (\log \log q + \log n))$ [BLLN13, Theorem 2] for \mathbb{Z}_2 being the plaintext space.

On the other hand, setting the bounds on f, g to $\|f\|_\infty = \|g\|_\infty = B_{\text{key}} = 1$, the plaintext space to \mathbb{Z}_2 via $t = 2$, the multiplicative expansion factor of the ring to $\delta = n$ by assuming n is a power of two and $w = O(1)$, then the multiplicative expansion factor of YASHE is $\mathcal{O}(n^2)$. For correctness, it is required that the noise be less than $q/4$.

Hence, to evaluate a circuit of depth L , YASHE requires $q/4 > \mathcal{O}(n^{2L})$ or $L = \mathcal{O}(\log q / \log n)$ under this choice of parameters. As a consequence, YASHE is usually instantiated with f and g very short, cf. [LN14].

Following [BV11, Lemma 4.5], Appendix H of [BLLN13] shows that YASHE is bootstrappable if it can evaluate depth $L = \mathcal{O}(\log \log q + \log n)$ circuits. For $\|f\|_\infty = \|g\|_\infty = B_{key} = 1$ this implies

$$\begin{aligned}\log \log q + \log(n) &< \log q / \log n, \\ \log(n \log q) &< \log q / \log n,\end{aligned}$$

i.e. q must be super-polynomial in n for YASHE to provide fully homomorphic encryption.

To establish a target size, recall that NTRU-like encryption of a binary message $\mu \in \mathbb{Z}_2$ is given by $c = h \cdot e_1 + e_2 + \mu \lfloor q/2 \rfloor$ for random errors of variance ς^2 . To decrypt from a solution (F, G) to the instance $h \leftarrow \text{NTRU}(\mathcal{R}, q, \sigma, \tau)$, simply compute $Fc = G \cdot e_1 + F \cdot e_2 + F \cdot \mu \lfloor q/2 \rfloor$. The error term $G \cdot e_1 + F \cdot e_2$ will have entries of magnitudes $\varsigma \tau \sqrt{n}$ which we require to be $< q/2$ to decrypt correctly. Hence, we require $F, G < q/(2\varsigma\sqrt{n})$. In [LTV12, BLLN13] like in other FHE schemes, ς is chosen to be bounded by a very small, constant value.

In [CS15] several Ring-based FHE schemes are compared. For comparability amongst the considered schemes and performance, the authors chose the coefficients of f, g from $\{-1, 0, 1\}$ with the additional guarantee that only 64 coefficients are non-zero in f or g . Then, to establish hardness they assume that an adversary who can find an element $< q$ in a q -ary lattice with dimension m and volume q^n wins for all schemes considered. Now, to achieve security against lattice attacks, the root Hermite factor δ_0 in $q = \delta_0^m q^{n/m}$ should be small enough, where “small enough” depends on which prediction for lattice reduction is used. In [DHS15] the same approach is used to pick parameters, but for a slightly smaller target norm of $q/4$.

The attack presented in this work results in a subexponential attack in the security parameter λ for LTV and YASHE, if L is sufficiently large to enable fully homomorphic encryption and if n is chosen to be minimal such that a lattice attack on the full field does not succeed. Set

$$q = \exp(\Theta((\epsilon + 1) \log^2 n))$$

to satisfy correctness. Now, to rule out lattice attacks on the full field set $n = \Theta(\lambda \log \lambda \log^2 \lambda)$. Hence, for $\beta = \lambda$ we have

$$\begin{aligned}\beta^{\Theta(n/\beta)} &> \sqrt{q}, \\ \Theta(\log^2 \lambda \log \log^2 \lambda) &> \Theta(\log^2 \lambda).\end{aligned}$$

For the subfield attack, pick $\beta = \Theta(\lambda / \log^{1/3} \lambda)$ and $r = \Theta(\log^{2/3} \lambda)$ and we get

$$\begin{aligned} \beta^{\Theta(n/\beta r)} \cdot n^{\Theta(r)} &< \sqrt{q}, \\ \Theta\left(\log^{\frac{5}{3}} \lambda \log \log^2 \lambda\right) &< \Theta(\log^2 \lambda). \end{aligned}$$

4.2 Graded Encoding Schemes

In [GGH13a] a candidate construction for graded encoding schemes approximating multilinear maps was proposed. The GGH construction was improved in [LSS14] and implemented and improved further in [ACLL15]. In these schemes, short elements $m_i \in \mathbb{Z}[X]/(X^n + 1)$ are encoded as $[(r_i \cdot g + m_i)/z]_q \in \mathcal{R}/q\mathcal{R}$ for some r_i, g with norms of size $\text{poly}(\lambda)$ and some random z . For correctness, the latest improvements [ACLL15] require a modulus $q = \text{poly}(\lambda)^\kappa$, where κ is the multi-linearity level. The subfield attack is therefore applicable in sub-exponential time for any $\kappa = \log^\epsilon \lambda$, according to Section 3.4, and would become polynomial for $\kappa > \Theta(\lambda \log \lambda)$. In practice, the fact that the constants in the exponent $q = \lambda^{\Theta(\kappa)}$ is quite large could make this attack quite powerful even for small degrees of multi-linearity.

While initially these constructions permitted the inclusion of encodings of zero ($m_i = 0$) to achieve multilinear maps, it was shown that these encodings break security [HJ15]. Without such encodings, the construction still serves as building-block for realizing Indistinguishability Obfuscation [GGH⁺13b].

To estimate parameters, [ACLL15] proceeds as follows¹¹. Given encodings $x_0 = [(r_0 \cdot g + m_0)/z]_q$ and $x_1 = [(r_1 \cdot g + m_1)/z]_q$ for unknown $m_0, m_1 \neq 0$ we may consider the NTRU lattice A_h^q where $h = [x_0/x_1]_q$. This lattice contains a short vector $(r_0 \cdot g + m_0, r_1 \cdot g + m_1)$. In [ACLL15] all elements of norm $\approx \|r_0 \cdot g + m_0\| = \sigma_1^*$ are considered “interesting” and recovering any such element is considered an attack. This is motivated by the fact that if an attacker recovers $r_0 \cdot g + m_0$ exactly, then it can recover z . This completely breaks the scheme.

The subfield lattice attack does not yield the vector $(r_0 \cdot g + m_0, r_1 \cdot g + m_1)$ exactly but only a relatively small multiple of it $u(r_0 \cdot g + m_0, r_1 \cdot g + m_1)$. We provide two approaches to completely break the scheme from this small multiple. The first approach consists of solving a principal ideal problem and leads to a quantum polynomial-time and classical subexponential attack. The second approach relies on a statistical leak using the

¹¹ The attack is attributed to Steven Galbraith in [ACLL15].

Gentry-Szydło algorithm [GS02,LS14], but is just outside reach with our current tools [GGH13a]. This approach is arguably worrisome, and the authors of [GGH13a] spent significant efforts to rule this approach out completely.

We remark that unlike previous cryptanalysis advances of multi-linear maps [HJ15] this attack does not rely either on the zero testing parameter, neither on encodings of zero. Our cryptanalytic result therefore impacts all applications of multilinear maps, from multi-party key exchange to jigsaw puzzles and Indistinguishability Obfuscation [GGH⁺13b]. For completeness, we note that the CLT construction [CLT13] of Graded Encoding Schemes is also subject to a quantum polynomial-time attack, because it relies on the hardness of factoring large integers.

The principal ideal problem and short generator recovery. The problem of recovering a short principal ideal generator from any generator received a lot of attention recently, and a series of works has lead to subexponential classical and polynomial-time quantum attacks against principal ideal lattices [EHKS14,CGS14,CDPR16,BS16]. Precisely, given the ideal $\mathfrak{I} = \langle g \rangle$, Biasse and Song [BS16] showed how to recover an arbitrary generator ug of \mathfrak{I} in quantum polynomial time, extending the recent breakthrough of Eisenträger et al. [EHKS14] on quantum algorithms over large degree number fields. Such results were conjectured already in a note of Cambell et al. [CGS14], where a classical polynomial time algorithm is also suggested to recover the original g from ug (namely, LLL in the log-unit lattice). The correctness of a similar algorithm was formally established using analytical number theory by Cramer et al. [CDPR16].

In combination with this subfield lattice attack, this directly implies a polynomial quantum attack. Indeed, the subfield lattice attack allows to recover $u(r_0 \cdot g + m_0)$ for some relatively short u . Repeating this attack several time, and obtaining $u(r_0 \cdot g + m_0)$ for various u eventually leads to the reconstruction of the ideal $\langle r_0 \cdot g + m_0 \rangle$. Because $r_0 \cdot g + m_0$ follows exactly a discrete Gaussian distribution, the approach sketched above can be applied, and reveals $r_0 \cdot g + m_0$ exactly, and therefore z .

In conclusion, for any degree of multi-linearity κ , the subfield attack can be complemented with a quantum polynomial step to a complete break. Alternatively, when $\kappa = O(\lambda^c)$ for any $c < 1/2$, — leading according to the previous best known attacks to a choice of dimension $n = \tilde{O}(\lambda^{1+c})$ — the $2^{\tilde{O}(n^{2/3})}$ algorithms of Biasse and Biasse and Fiecker [Bia14,BF14] combined lead to a classical attack in time sub-exponential in λ .

The statistical attack. This attack consists in recovering $u\bar{u}$ and $\langle u \rangle$ and using the Gentry-Szydlo algorithm [GS02,LS14] to recover u .

To recover $\langle u \rangle$, note that we are given $u(a_0, a_1)$. We will assume that $\langle a_0 \rangle, \langle a_1 \rangle$ are coprime with constant probability, cf. Section 2.2. Under this assumption, $\langle u \rangle$ can be recovered as $\langle u \rangle = \langle ua_0 \rangle + \langle ua_1 \rangle$.¹²

To recover more information on u , we can compute $ua_0 \cdot [x_i/x_0]_q = ua_i$ for other $i > 1$, and the equation hold over \mathcal{R} because u and a_i are small. For $i > 1$, a_i is independent of u and follows a spherical Gaussian of parameter σ . It follows that the variance of ua_i leaks $u\bar{u}$: $\mathbb{E}[ua_i \cdot \bar{u}\bar{a}_i] = \sigma^2 u\bar{u}$.

Given polynomially many samples x_i one can therefore recover $u\bar{u}$ up to a $1 + 1/\text{poly}(\lambda)$ approximation factor. The original attack of Gentry-Szydlo algorithm [GS02,LS14] requires the exact knowledge of $u\bar{u}$ that could be obtained by rounding when u has poly-sized coefficient. However, the u provided by the subfield lattice attack is much larger. In [GGH13a] this algorithm is revisited and extended to when $u\bar{u}$ is only known up to a $1 + (\log n)^{-\Theta(\log n)}$ approximation factor.

In conclusion, with the current algorithmic tools this approach is asymptotically inapplicable if we assume only a polynomial number of available samples, but only barely so. This raises the question of how to improve the tolerance of the Gentry-Szydlo algorithm¹³. Yet, because $(\log n)^{\Theta(\log n)}$ is arguably not so large, it is unclear whether this approach is really infeasible in practice.

We concur with the decision made in [GGH13a], to attempt to rule out such an attack by design even if it is not yet known how to fully exploit it.

5 Experimental Verification

Please refer to the full version of this work for experiments.

6 Conclusions

Practicality of the attack. The largest instance we broke in practice is for the set of parameter $n = 2^{12}$ and $q \approx 2^{190}$. Choosing a relative degree

¹² Note that the subfield lattice attack may be tweaked to obtain a triplet $u(a_0, a_1, a_2)$ (or more) increasing the probability to recover $\langle u \rangle$.

¹³ Asymptotically, the natural idea of replacing LLL by slightly stronger lattice reduction does not seem to help, but should help in practice. The quasi-polynomial factor relates to a number theoretic heuristic. See Section 7.6 of [GGH13a].

$r = 16$, the attack required to run LLL in dimension 512, which took about 120 hours, single-threaded, using SAGE [Dev15] and FPLLL [ABC⁺]. The direct, full field lattice reduction attack, according to root-Hermite-factor based predictions [CN11], would have required running BKZ in block-size ≈ 130 , and in dimension 8192, which is hardly feasible with the current state-of-the art [CN11] (requiring more than 2^{70} CPU cycles). We conclude that the subfield attack proposed in this work is not only theoretical but also practical.

Obstructions to concrete predictions. We are currently unable to predict precisely how a given set of parameters would be affected, for example to predict the power of this attack against concrete parameter choices of NTRU-based FHE [LTV12,BLLN13] and Multilinear Maps [GGH13a].

There are two issues for those predictions. The first issue is that we make use of LLL/BKZ in the approximation-factor regime, not in the Hermite-factor regime. While the behavior of LLL/BKZ is quite well modeled in the latter regime, we are not aware of precise models for the former for NTRU lattices. Unlike the Hermite-factor regime, this case could very well be influenced by the presence of many short vectors rather than just a few.

The second issue is that we do not know the actual size of the shortest vector of $\Lambda_{h'}^q$: all we know is that it is no larger than (f', g') . In several cases in the experiments we found vectors $(x', y') = v(f', g')$ that were actually shorter than (f', g') —the tentative root-approximation factor α is less than 1. One may expect that (f', g') may still be (or close to) the shortest vector for small relative degree r as it is the shortest with high probability in the full field (i.e. when $r = 1$).

Immunity of NTRU encryption and BLISS signature schemes. If q is small enough, then the attacks should become inapplicable, even with the smallest possible relative dimension $r = 2$. Precisely, if (f', g') is not an unusually short vector of $\Lambda_{h'}^q$, then there is little hope that any lattice reduction strategy would lead to information on this vector. Quantitatively, this perfect immunity happens when $\|(f', g')\| \approx \sqrt{2} \cdot \sigma^2 \cdot n' > \sqrt{n'q/\pi e}$. This was the case of the old parameter of NTRU as discussed in [Gen01], which lead this attack being discarded. This is not the case of all the parameters of NTRUENCRYPT [HPS⁺15] and BLISS [DDL13], for which (f', g') is sometime unusually short vector, but not by a very large factor. Numerical values are given in Table 1.

When the vulnerability factor F is less than 1, the parameters achieve perfect immunity. When F is greater than 1, the subfield attack consist

Table 1: Vulnerability factor for some parameters of NTRUENCRYPT [HPS⁺15] and BLISS [DDL13].

Scheme	n	q	σ	$\sqrt{n'q/\pi e} / (\sqrt{2}\sigma^2 n') = F$			
NTRU-743	743	2048	0.82	298.7	/	349.8	= 0.85
NTRU-401	401	2048	0.82	219.6	/	189.5	= 1.16
BLISS-I	512	12289	0.55	607.0	/	108.6	= 5.59
BLISS-IV	512	12289	0.83	607.0	/	249.8	= 2.43

informally of solving “unusual-SVP” in dimension $2n' = n$, where the unusually short solutions are a factor F shorter than predicted by the Gaussian Heuristic.

According to this table, NTRU-743 should be perfectly immune to the subfield lattice attacks. For other parameters, it seems likely, despite imperfect immunity, that the subfield lattice attack will be more costly than the full attack, but calls for further study, especially for BLISS-I.

Note that the perfect immunity to this attack is achieved asymptotically around $\sigma \approx \Theta(q^{1/4})$, parameter for which h does not have enough entropy to be statistically close to random. For comparison, it was shown that for $\sigma = \omega(q^{1/2})$, h is statistically close to uniform [SS11]. We note that $\sigma > \Theta(q^{1/4})$ could provide enough entropy for the normed-down public key h' to be almost uniform. It would be interesting to see if the proof of [SS11] can be adapted to h' .

Recommendations. Even if credible predictions were to be made, we strongly discourage basing a cryptographic scheme on a set-up to which this attack is applicable. Indeed, it is quite likely that the performance of the attack may be improved in several ways. For example, after having found several subfield solutions $(x', y') = v(f', g')$, it is possible to run a lattice reduction algorithm in the lattice $(f', g') \cdot \mathcal{O}_{\mathbb{L}}$ of dimension n' rather than $2n'$ to obtain significantly shorter vectors. Additionally, the lifting step may also be improved in the case where $\mathcal{O}_{\mathbb{L}}$ is a real subfield using the Gentry-Syrdlo algorithm [GS02,LS14] to obtain shorter vector in the full field (i.e. recovering x from $N_{\mathbb{K}/\mathbb{L}}(x)$). More generally, one may recover x from $N_{\mathbb{K}/\mathbb{L}}(x)$ even when \mathbb{L} isn’t the real subfield of \mathbb{K} : assuming (x) is prime, it can be recovered as a factor of $N_{\mathbb{K}/\mathbb{L}}(x)$, which then leads to x via a short generator recovery; as mentioned before, both steps

are now known to be classically sub-exponential or even polynomial for quantum computers [Bia14,EHKS14,CGS14,BS16,CDPR16].

Evaluating concrete security against regular lattice attacks is already a difficult exercise, and leaving open additional algebraic and statistical attack opportunities will only make security assessment intractable. We therefore recommend that this set-up —NTRU assumption, presence of subfields, large modulus— be considered insecure.

Designing Immune Rings. We believe that our work further motivates the design and the study of number fields without subfields to fit for the lattice-based cryptographic purposes, as already recommended in [Ber14]. Even for assumptions that are not directly affected by this attack (Ring-SIS [Mic02], Ideal-LWE [SSTX09], Ring-LWE [LPR10]), it could be considered desirable to have efficient fallback options ready to use, in case subfields induce other unforeseen weaknesses. While this work does not suggest an immediate threat to the Ring-SIS and Ring-LWE, such a precaution is not unreasonable.

An interesting option has been suggested in [Ber14] to use rings of the form $\mathbb{Z}[X]/(X^p - X - 1)$. The design rationale seems to be that $\mathbb{Q}[X]/(X^p - X - 1)$ has a reasonable expansion factor¹⁴ which is often needed for the correctness in cryptographic schemes, but is a non Galois extension with a very large Galois group for its splitting field, which is intended to hinder algebraic handles. In particular it contains no proper subfields. This leads to the design of the NTRUPrime encryption scheme [BCLvV16]. We note that the security of this scheme is not supported by a worst-case hardness argument. If such an argument is desired then we note that the *search version* of Ideal/Ring-LWE is supported by worst-case hardness for *other choices of number field*, and this is actually sufficient to achieve provable CPA-secure encryption, as already proved by Stehlé, Steinfeld, Tanka and Xagawa [SSTX09].

Open Problems. Another natural option would be to choose p as a safe prime¹⁵ and to work with the ring of integer of the *totally real* number field $\mathbb{K} = \mathbb{Q}(\zeta_p + \bar{\zeta}_p)$. The field remains Galois, and its automorphism group may still allow a quantum worst-case (Ideal-SVP) to average-case (Ring-LWE) reduction a-la [LPR10] thanks to a generalization of the search to decision step presented in [CLS15]. Nevertheless the Galois group has

¹⁴ Multiplication of two small elements remains reasonably small.

¹⁵ A safe prime p is an odd prime such that $(p - 1)/2$ is also a prime. The terminology relates to weaknesses in RSA and Discrete Logarithm Problem introduced by the smoothness of $p - 1$ [Pol74].

prime order $(p-1)/2$, it has no proper subgroups, and \mathbb{K} has no proper subfields.

But working with $\mathbb{K} = \mathbb{Q}(\zeta_p + \bar{\zeta}_p)$ has a drawback: the class number $h(\mathbb{K}) = h_p^+$ seems quite small (see [Was97, Table 4 pp. 421]), and this makes the worst-case SIVP problem solvable in quantum polynomial time for approximation factors $2^{\tilde{O}(\sqrt{n})}$ as proved in [CDPR16,BS16]: the reduction of [LPR10] is vacuous for such parameters.

This raises the question of whether NTRU and Ring-LWE are actually strictly harder than SIVP in the underlying number field, whether algorithms for SIVP in \mathbb{K} can be lifted to modules over \mathbb{K} as used in NTRU, Ideal-LWE or Ring-LWE. In this regard, overstretched NTRU, and Ideal/Ring-LWE with large approximation factors over the ring $\mathbb{Z}(\zeta_p + \bar{\zeta}_p)$ are very interesting cryptanalytic target: despite those rings not being used in any proposed schemes so far, such an attack will teach us a great deal on the asymptotic security of ideal-lattice based cryptography.

References

- ABC⁺. M. Albrecht, S. Bai, D. Cadé, X. Pujol, and D. Stehlé. fpLLL-4.0, a floating-point LLL implementation. Available at <https://github.com/dstehle/fp111>.
- ACLL15. Martin R. Albrecht, Catalin Cocis, Fabien Laguillaumie, and Adeline Langlois. Implementing candidate graded encoding schemes from ideal lattices. In Tetsu Iwata and Jung Hee Cheon, editors, *ASIACRYPT 2015, Part II*, volume 9453 of *LNCS*, pages 752–775. Springer, Heidelberg, November / December 2015.
- BCLvV16. Daniel J. Bernstein, Chitchanok Chuengsatiansup, Tanja Lange, and Christine van Vredendaal. Ntru prime. Cryptology ePrint Archive, Report 2016/461, 2016. <http://eprint.iacr.org/>.
- Ber14. Dan Bernstein. A subfield-logarithm attack against ideal lattices. <http://blog.cr.yp.to/20140213-ideal.html>, February 2014.
- BF14. Jean-François Biasse and Claus Fieker. Subexponential class group and unit group computation in large degree number fields. *LMS J. Comput. Math.*, 17(suppl. A):385–403, 2014.
- Bia14. Jean-François Biasse. Subexponential time relations in the class group of large degree number fields. *Adv. Math. Commun.*, 8(4):407–425, 2014.
- BLLN13. Joppe W. Bos, Kristin Lauter, Jake Loftus, and Michael Naehrig. Improved security for a ring-based fully homomorphic encryption scheme. In Martijn Stam, editor, *14th IMA International Conference on Cryptography and Coding*, volume 8308 of *LNCS*, pages 45–64. Springer, Heidelberg, December 2013.
- BS16. Jean-François Biasse and Fang Song. Efficient quantum algorithms for computing class groups and solving the principal ideal problem in arbitrary degree number fields. In *27th ACM-SIAM Symposium on Discrete Algorithms (SODA’16)*, 2016.

- BV11. Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011.
- CDPR16. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part II*, pages 559–585. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- CG13. Ran Canetti and Juan A. Garay, editors. *CRYPTO 2013, Part I*, volume 8042 of *LNCS*. Springer, Heidelberg, August 2013.
- CGS14. Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: A cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014. Available at http://docbox.etsi.org/Workshop/2014/201410_CRYPT0/S07_Systems_and_Attacks/S07_Groves_Annex.pdf.
- CIV16. Wouter Castryck, Ilia Iliashenko, and Frederik Vercauteren. Provably weak instances of ring-lwe revisited. In Marc Fischlin and Jean-Sébastien Coron, editors, *Advances in Cryptology – EUROCRYPT 2016: 35th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Vienna, Austria, May 8–12, 2016, Proceedings, Part I*, pages 147–167. Springer Berlin Heidelberg, Berlin, Heidelberg, 2016.
- CJL16. Jung Hee Cheon, Jinhyuck Jeong, and Changmin Lee. An algorithm for ntru problems and cryptanalysis of the ggh multilinear map without an encoding of zero. Cryptology ePrint Archive, Report 2016/139, 2016. <http://eprint.iacr.org/>.
- CLS15. Hao Chen, Kristin Lauter, and Katherine E. Stange. Attacks on search RLWE. Cryptology ePrint Archive, Report 2015/971, 2015. <http://eprint.iacr.org/2015/971>.
- CLT13. Jean-Sébastien Coron, Tancrede Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In Canetti and Garay [CG13], pages 476–493.
- CN11. Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Dong Hoon Lee and Xiaoyun Wang, editors, *ASIACRYPT 2011*, volume 7073 of *LNCS*, pages 1–20. Springer, Heidelberg, December 2011.
- CS97. Don Coppersmith and Adi Shamir. Lattice attacks on NTRU. In Walter Fumy, editor, *EUROCRYPT’97*, volume 1233 of *LNCS*, pages 52–61. Springer, Heidelberg, May 1997.
- CS15. Anamaria Costache and Nigel P. Smart. Which ring based somewhat homomorphic encryption scheme is best? Cryptology ePrint Archive, Report 2015/889, 2015. <http://eprint.iacr.org/2015/889>.
- DDLL13. Léo Ducas, Alain Durmus, Tancrede Lepoint, and Vadim Lyubashevsky. Lattice signatures and bimodal gaussians. In Canetti and Garay [CG13], pages 40–56.
- Dev15. The Sage Developers. *Sage Mathematics Software*, 2015. <http://www.sagemath.org>.
- DHS15. Yarkın Doröz, Yin Hu, and Berk Sunar. Homomorphic aes evaluation using the modified ltv scheme. *Designs, Codes and Cryptography*, pages 1–26, 2015.

- EHKS14. Kirsten Eisenträger, Sean Hallgren, Alexei Kitaev, and Fang Song. A quantum algorithm for computing the unit group of an arbitrary degree number field. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing*, pages 293–302. ACM, 2014.
- EHL14. Kirsten Eisenträger, Sean Hallgren, and Kristin E. Lauter. Weak instances of PLWE. In Antoine Joux and Amr M. Youssef, editors, *SAC 2014*, volume 8781 of *LNCS*, pages 183–194. Springer, Heidelberg, August 2014.
- ELOS15. Yara Elias, Kristin E. Lauter, Ekin Ozman, and Katherine E. Stange. Provably weak instances of ring-LWE. In Gennaro and Robshaw [GR15], pages 63–92.
- FM14. Andrea Ferraguti and Giacomo Micheli. On the Mertens–Cesàro theorem for number fields. *Bulletin of the Australian Mathematical Society*, pages 1–12, 2014.
- Gen01. Craig Gentry. Key recovery and message attacks on NTRU-composite. In Pfitzmann [Pfi01], pages 182–194.
- GGH13a. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In Thomas Johansson and Phong Q. Nguyen, editors, *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 1–17. Springer, Heidelberg, May 2013.
- GGH⁺13b. Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *54th FOCS*, pages 40–49. IEEE Computer Society Press, October 2013.
- GN08. Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within Mordell’s inequality. In Richard E. Ladner and Cynthia Dwork, editors, *40th ACM STOC*, pages 207–216. ACM Press, May 2008.
- GR15. Rosario Gennaro and Matthew J. B. Robshaw, editors. *CRYPTO 2015, Part I*, volume 9215 of *LNCS*. Springer, Heidelberg, August 2015.
- GS02. Craig Gentry and Michael Szydlo. Cryptanalysis of the revised NTRU signature scheme. In Lars R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 299–320. Springer, Heidelberg, April / May 2002.
- HG07. Nick Howgrave-Graham. A hybrid lattice-reduction and meet-in-the-middle attack against NTRU. In Alfred Menezes, editor, *CRYPTO 2007*, volume 4622 of *LNCS*, pages 150–169. Springer, Heidelberg, August 2007.
- HHGP⁺03. Jeffrey Hoffstein, Nick Howgrave-Graham, Jill Pipher, Joseph H. Silverman, and William Whyte. NTRUSIGN: Digital signatures using the NTRU lattice. In Marc Joye, editor, *CT-RSA 2003*, volume 2612 of *LNCS*, pages 122–140. Springer, Heidelberg, April 2003.
- HJ15. Yupu Hu and Huiwen Jia. Cryptanalysis of GGH map. Cryptology ePrint Archive, Report 2015/301, 2015. <http://eprint.iacr.org/2015/301>.
- HPS96. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A new high speed public key cryptosystem, 1996. Draft Distributed at Crypto’96, available at <http://web.securityinnovation.com/hubfs/files/ntru-orig.pdf>.
- HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288, 1998.
- HPS01. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NSS: An NTRU lattice-based signature scheme. In Pfitzmann [Pfi01], pages 211–228.
- HPS11. Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In Phillip Rogaway, editor,

- CRYPTO 2011*, volume 6841 of *LNCS*, pages 447–464. Springer, Heidelberg, August 2011.
- HPS⁺15. Jeff Hoffstein, Jill Pipher, John M. Schanck, Joseph H. Silverman, William Whyte, and Zhenfei Zhang. Choosing parameters for NTRUEncrypt. Cryptology ePrint Archive, Report 2015/708, 2015. <http://eprint.iacr.org/2015/708>.
- HSW06. Jeffrey Hoffstein, Joseph H. Silverman, and William Whyte. Meet-in-the-middle attack on an ntru private key, 2006. Technical report, NTRU Cryptosystems, July 2006. Report #04, available at <http://www.ntru.com>.
- HT15. Adrien Hauteville and Jean-Pierre Tillich. New algorithms for decoding in the rank metric and an attack on the LRPC cryptosystem. In *IEEE International Symposium on Information Theory, ISIT 2015*, pages 2747–2751, 2015.
- KF15. Paul Kirchner and Pierre-Alain Fouque. An improved BKW algorithm for LWE with applications to cryptography and lattices. In Gennaro and Robshaw [GR15], pages 43–62.
- LJ14. Carl Löndahl and Thomas Johansson. Improved algorithms for finding low-weight polynomial multiples in $f_2[x]$ and some cryptographic applications. *Designs Codes and Cryptography*, 73(2):625–640, 2014.
- LLL82. Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.
- LN14. Tancrede Lepoint and Michael Naehrig. A comparison of the homomorphic encryption schemes FV and YASHE. In David Pointcheval and Damien Vergnaud, editors, *AFRICACRYPT 14*, volume 8469 of *LNCS*, pages 318–335. Springer, Heidelberg, May 2014.
- Loi14. Pierre Loidreau. On cellular codes and their cryptographic applications. In *ACCT, Fourteenth International Workshop on Algebraic and Combinatorial Coding Theory*, pages 234–239, 2014.
- Lov87. L. Lovasz. *An Algorithmic Theory of Numbers, Graphs and Convexity*. CBMS-NSF Regional Conference Series in Applied Mathematics. Society for Industrial and Applied Mathematics, 1987.
- LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May 2010.
- LS14. H. W. Lenstra and A. Silverberg. Revisiting the Gentry-Szydlo algorithm. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 280–296. Springer, Heidelberg, August 2014.
- LSS14. Adeline Langlois, Damien Stehlé, and Ron Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 239–256. Springer, Heidelberg, May 2014.
- LTV12. Adriana López-Alt, Eran Tromer, and Vinod Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In Howard J. Karloff and Toniann Pitassi, editors, *44th ACM STOC*, pages 1219–1234. ACM Press, May 2012.
- Mic02. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *43rd FOCS*, pages 356–365. IEEE Computer Society Press, November 2002.

- Pei16. Chris Peikert. How (not) to instantiate ring-lwe. Cryptology ePrint Archive, Report 2016/351, 2016. <http://eprint.iacr.org/>.
- Pfi01. Birgit Pfitzmann, editor. *EUROCRYPT 2001*, volume 2045 of *LNCS*. Springer, Heidelberg, May 2001.
- Pol74. John M Pollard. Theorems on factorization and primality testing. *Mathematical Proceedings of the Cambridge Philosophical Society*, 76(03):521–528, 1974.
- Sam70. Pierre Samuel. *Algebraic Theory of Numbers*. Hermann, Paris, 1970.
- Sch87. Claus-Peter Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.
- Sit10. Brian D Sittinger. The probability that random algebraic integers are relatively r-prime. *Journal of Number Theory*, 130(1):164–171, 2010.
- SS11. Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In Kenneth G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 27–47. Springer, Heidelberg, May 2011.
- SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 617–635. Springer, Heidelberg, December 2009.
- Was97. L.C. Washington. *Introduction to Cyclotomic Fields*. Graduate Texts in Mathematics. Springer New York, 1997.