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.
[2] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENSL, INRIA, UCBL), France.
[3] Cryptology Group, CWI, Amsterdam, The Netherlands.

Abstract. We exploit the presence of a subfield to solve the NTRU problem for large moduli q: norming-down the public key h 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.
We restrict ourselves to choices of dimensions $n(\lambda)$ and modulus $q(\lambda)$ that were previously thought to offer resistance against attacks in time exponential in the security parameter λ. For any super-polynomial $q(\lambda)$, the subfield attack can be made sub-exponential in λ, or even polynomial as $q(\lambda)$ gets larger.
The subfield lattice attack directly affects the asymptotic security of the bootstrappable homomorphic encryption schemes LTV and YASHE. It also makes GGH-like Multilinear Maps vulnerable to principal ideals attacks — therefore leading to a quantum break — and almost vulnerable to a statistical attack a-la Gentry-Szydlo. No *encodings of zero* nor *zero-testing parameter* are required.
We also provide meaningful practical experiments. Using just LLL in dimension 512 we obtain vectors that would have required running BKZ with block-size 130 in dimension 8192.
Finally, we discuss concrete aspects of this attack, the potential immunity of NTRUEncrypt and Bliss parameters, issue preliminary recommendations and suggest countermeasures.

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 assumption on structured lattices such as the NTRU lattice. Introduced by Hoffstein, Pipher 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 ($P = \Phi_m$, $n = \phi(m)$).

Following on 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] were the first to notice that recovering a short enough vector, potentially different from the actual secret key (f, g), may be sufficient for an attack and claimed

⋆ Supported by EPSRC grant EP/L018543/1 "Multilinear Maps in Cryptography". Email: martin.albrecht@royalholloway.ac.uk
⋆⋆ Supported by ERC Starting Grant ERC-2013-StG-335086-LATTAC. Email: shih.bai@gmail.com
⋆⋆⋆ This work has been 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). Email: ducas@cwi.nl

that the celebrated LLL algorithm of Lenstra, Lenstra and Lovász [LLL82] would lead to an attack. However, it turned out [HPS98] that much stronger lattice reduction is required and that the NTRUEncrypt scheme is asymptotically secure. Meanwhile, parameters have been updated to take account for progress in lattice reduction and potential quantum speed-ups [HPS+15].

Other attacks have been considered, such as Odlyzko's meet-in-the-middle attack described in [JHW06]. In practice, the best known algorithm for attacking NTRU lattices is the combined lattice-reduction and meet-in-the-middle attack of Howgrave-Graham [HG07]. Asymptotically, a slightly sub-exponential attack against ternary-NTRU 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.

As of today, those NTRU lattices remained essentially as intractable as lattices with similar parameters[4], but without the structure of $\mathcal{R}$-module. An exception —discussed below— is an attack of Gentry [Gen01] tackling the case of composite rings.

In the present work, we describe how to use lattice reduction in a *subfield* to attack the NTRU assumption for large moduli q. This subfield lattice attack is asymptotically faster than the previously known attacks as soon as q is super-polynomial, and may also be relevant for polynomially-sized q.

Asymptotics. We are mostly concerned with the NTRU assumption 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 $q = 2^{\tilde{\Theta}(\lambda)} = 2^{\tilde{\Theta}(\sqrt{n})}$.

Our contribution. We present a new subfield lattice attack which consists of norming down to a subfield, running lattice reduction to solve a smaller, easier lattice problem and lifting back up. We then show that the proposed algorithm solves the 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. We stress that our attacks 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 note that the behavior of the lattice reduction algorithms on the special instances considered in this work seems not to be captured by current lattice reduction models: we are yet unable to provide practical predictions for the hidden constants in our asymptotic results.

Previous work. Our work is very similar in spirit to an attack of Gentry [Gen01] against the NTRU-composite assumption. His attack 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 he 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 this condition, there

[4] Volume, dimension and length of unusually short vectors.

exists a projection $\pi : \mathcal{R} \to \mathcal{R}_1$ that is a ring morphism, and he shows that this projection can 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 morphism, this approach is not limited to NTRU, and would also apply to Ring-SIS or Ring-LWE.

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

In comparison, this work tackles NTRU when $\mathcal{R} = \mathcal{O}_\mathbb{K}$ is a the ring of integer of a number field $\mathbb{K}$ (and therefore can not be a direct product), and that $\mathbb{K}$ admit proper subfields. Due to Gentry's attack and others, direct product rings are now avoided for lattice-based cryptography, and the typical choice is to use the rings of integer of cyclotomic number fields $\mathcal{R} = \mathcal{O}_{\mathbb{Q}(\omega_m)} = \mathbb{Z}[\omega_m]$. This setting allows to argue worst-case hardness of certain problems (Ring-SIS [Mic02], Ring-LWE [LPR10]). Yet all those number fields admit proper subfields (at least, the maximal real subfield). Instead of a projection map π, we exploit a relative norm map $\mathrm{N}_{\mathbb{K}/\mathbb{L}} : \mathcal{O}_\mathbb{K} \to \mathcal{O}_\mathbb{L}$, 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 also resonates with the logarithm-subfield strategy of Bernstein [Ber14], which anticipated other works towards a logarithm attack [CGS14,CDPR15]. 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 specifics set-ups.

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; Subsection 3.4 analyzes its asymptotic performances. In Section 4, we apply our 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, Hendrik Lenstra and Damien Stehlé for helpful discussions and comments. Finally, we thank the PSMN (Pôle Scientifique de Modélisation Numérique, Lyon, France) for providing computing facilities.

2 Preliminaries

Vectors will be considered as row vectors. The notation $[\,\cdot\,]_q$ denotes reduction modulo an integer q.

2.1 Number fields and subfields

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 $\mathbb{Q}$, of 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 $\mathbb{K}$, 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 we have $r = n/n'$). The number fields $\mathbb{K}$, $\mathbb{L}$ and therefore the degrees n, n' and relative degree r are fixed in this rest of this work.

The relative norm $\mathrm{N}_{\mathbb{K}/\mathbb{L}} : \mathbb{K} \to \mathbb{L}$ (resp. relative trace $\mathrm{Tr}_{\mathbb{K}/\mathbb{L}} : \mathbb{K} \to \mathbb{L}$) is the multiplicative (resp. additive) map defined by

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

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

Cyclotomic Number Field We note ω_m an arbitrary primitive m-th root of unity. For cryptanalytic purposes, we are mostly interested in the case where $\mathbb{K} = \mathbb{Q}(\omega_m)$ is the m-th cyclotomic number field, but we want to instantiate our attack for subfields $\mathbb{L}$ of $\mathbb{K}$ that are not necessary cyclotomic number fields.

The number field $\mathbb{K} = \mathbb{Q}(\omega_m)$ 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 $n'|n$.[5]

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

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

Below, we will rely on two principal ideals in $\mathcal{O}_{\mathbb{K}}$ being coprime in some proofs. The density of coprime pairs of ideals [Sit10] and elements [FM14] in $\mathcal{O}_{\mathbb{K}}$ is $1/\zeta_{\mathbb{K}}(2)$ where $\zeta_{\mathbb{K}}$ denotes the Dedekind zeta function over $\mathbb{K}$. The next lemma shows that $\zeta_{\mathbb{K}}(2) \leq \zeta(2)$ where ζ is the Riemann zeta function for any $\mathbb{K}$.

Lemma 1. *If $\mathbb{K}$ is an extension of $\mathbb{L}$, then $\zeta_{\mathbb{K}}(s) \leq \zeta_{\mathbb{L}}(s)$ for any real $s > 1$. In particular*

$$\zeta_{\mathbb{K}}(2) \leq \zeta(2) = \pi^2/6$$

where ζ is the Riemann zeta function.

Proof. We have

$$\zeta_{\mathbb{K}}(s) = \prod_{P \subseteq \mathcal{O}_{\mathbb{K}}} \frac{1}{1 - (N_{\mathbb{K}/\mathbb{Q}}(P))^{-s}}$$

Each prime ideal P of $\mathbb{K}$ contains a prime ideal p that lies below in $\mathbb{L}$. The absolute norm of P is no smaller than that of p; and hence the claim follows. □

We have a lower bound $6/\pi^2$ for the density. Further, we numerically approximated $\zeta_{\mathbb{K}}^{-1}(2)$ for $\mathbb{K} = \mathbb{Q}[x]/(x^n + 1)$ for $n = 128$ and $n = 256$ by computing the first 2^{22} terms of the Dirichlet series of the Dedekind zeta function for $\mathbb{K}$ 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 principal ideals with short generators under the additional condition that f is invertible modulo q. However, our experiments indicate that we may heuristically use the density as discussed as the probability of our pairs being coprime.

[5] 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.

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

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

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

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

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

$$|a| = \max_e |e(a)| \tag{4}$$

where e ranges over all the embeddings $e : \mathbb{K} \to \mathbb{C}$. We note that if $\omega \in \mathbb{K}$ is a root of unity, then $|\omega| = 1$.

The Euclidean norm and the operator's norm are invariant under automorphisms $\psi : \mathbb{K} \mapsto \mathbb{K}$,

$$\|a\| = \|\psi(a)\|, \quad |a| = |\psi(a)| \tag{5}$$

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

$$\mathrm{N}_{\mathbb{K}/\mathbb{Q}}(a) \leq |a|^n \leq \|a\|^n. \tag{6}$$

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

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 works, 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 works instead chose random ternary coefficients. In the typical power-of-two case cyclotomic case, such distribution is isotropic of variance $n2/3$. 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 set-up choice 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 $\dim(\Lambda)$ of a lattice Λ is the dimension over $\mathbb{Q}$ of the $\mathbb{Q}$-vector space it spans.[6].We recall that the minimal distance of a lattice Λ is defined as $\lambda_1(\Lambda) = \min_{v \in \Lambda \setminus \{0\}} \|v\|$.

[6] Or equivalently, the size of a minimal sets of $\mathbb{Z}$-generators, since $\mathbb{Z}$ is a principal ideal domain

Also, the volume of a lattice $\mathrm{Vol}(\Lambda)$ is defined as square root of the absolute determinant of the Gram matrix of any basis $\{b_1 \dots b_{\dim(\Lambda)}\}$ of Λ:

$$\mathrm{Vol}(\Lambda) = \sqrt{\det([\langle b_i, b_j \rangle]_{i,j})} \tag{7}$$

For any set of $\mathbb{Q}$-linearly independent vectors $\{v_1, \dots, v_{\dim(\Lambda)}\} \subset \Lambda$, we have the inequality:

$$\mathrm{Vol}(\Lambda) \leq \prod \|v_i\| \tag{8}$$

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 is not necessary equal to the size of a minimal set of $\mathcal{O}_\mathbb{K}$-generators.[7]. 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 integer $\sqrt{|\Delta_\mathbb{K}|} = \mathrm{Vol}(\mathcal{O}_\mathbb{K})$. It is a positive integer. More generally, we have the identity:

$$\mathrm{Vol}(a\mathcal{O}_\mathbb{K}) = \mathrm{N}_{\mathbb{K}/\mathbb{Q}}(a)\sqrt{|\Delta_\mathbb{K}|} \tag{9}$$

This give 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. Then $\mathrm{Vol}(M) \leq \lambda_1(M)^n \sqrt{|\Delta_\mathbb{K}|}$.*

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

2.5 NTRU assumption

Let us first describe the NTRU problem as follows.

Definition 2 (NTRU problem). *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, $\mathsf{NTRU}(\mathcal{R}, q, \mathcal{D}, \tau)$ is the problem of, given $h = gf^{-1} \bmod q$ (conditioned on f being invertible $\bmod q$) 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*

$$\Lambda_h^q = \{(x, y) \in \mathcal{R}^2 \ \textit{s.t.}\ hx - y = 0 \bmod q\}. \tag{10}$$

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

Note that $\mathsf{NTRU}(\mathcal{R}, q, \sigma, \sigma)$ is essentially the problem of recovering the secret key (f, g). Yet, in many cases, solving $\mathsf{NTRU}(\mathcal{R}, q, \sigma, \tau)$ 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 $h \leftarrow \mathsf{NTRU}(\mathcal{O}_\mathbb{K}, q, \sigma, \tau)$ has dimension $2n$ and volume $\mathrm{Vol}(\mathcal{R})^2 q^n$. Consequently, if h were to be uniformly random, the Gaussian heuristic predicts that the shortest vectors of Λ_h^q have norm $\mathrm{Vol}(\mathcal{R})^{1/n}\sqrt{nq/\pi e}$. Therefore, whenever $\sigma < \mathrm{Vol}(\mathcal{R})^{1/n}\sqrt{q/2\pi e}$, 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.

[7] Non-principal ideals of $\mathbb{K}$ being a counter-example

Target parameter τ for attacks. Because no solution would be expected if h was uniformly random, solving $h \leftarrow \mathsf{NTRU}(\mathcal{R}, q, \sigma, \tau)$ for $\tau < \mathrm{Vol}(\mathcal{R})^{1/n}\sqrt{q/\pi e}$ already constitutes a distinguishing attack on the NTRU problem. The problem of distinguishing h from uniform is also known as the Decisional Small Polynomial Ratio problem [LTV12]. 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

Theoretically, one of the best lattice-reduction algorithm beyond LLL [LLL82] is the slide algorithm [Sch87,GN08].

Theorem 1 (from [GN08]). *There is an algorithm that, given $\epsilon > 0$, the basis B of a lattice L of dimension d, and making at most $\mathrm{poly}(d, 1/\epsilon, \mathrm{bitsize}(B))$ many operations and calls to an SVP oracle in dimension β, outputs a vector $v \in L$ whose length verify both following bounds:*

- *the approximation-factor bound:*

$$\|v\| \leq ((1+\epsilon)\gamma_\beta)^{\frac{d-\beta}{\beta-1}} \cdot \lambda_1(L) \tag{11}$$

  where $\lambda_1(L)$ is the length of a shortest vector in L.
- *the Hermite-factor bound:*

$$\|v\| \leq ((1+\epsilon)\gamma_\beta)^{\frac{d-1}{2\beta-2}} \cdot \mathrm{Vol}(L)^{1/d} \tag{12}$$

where $\gamma_\beta \approx \beta$ is the β-dimensional Hermite constant.

Alternatively, one may use BKZ with early termination, and a similar Hermite-factor inequality may be proved [HS07]. However, we are not aware of any proof of a similar approximation factor is known unless we leave BKZ running for super-polynomial time.

It is well known [CN11] that in practice lattice reduction algorithms achieve much shorter results and are more efficient, but the factors remains of the order of $\beta^{\Theta(n/\beta)}$, for a computational cost in $\mathsf{poly}(\lambda) \cdot 2^{\Theta(\beta)}$.

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 Subsections 3.1, 3.2 and 3.3. We then analyze in Subsection 3.4 the asymptotic performances compared to direct reduction in the full field for cryptographically relevant asymptotic parameters.

We are given an instance $h \leftarrow \mathsf{NTRU}(\mathcal{O}_\mathbb{K}, q, \sigma, \tau)$, 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 $f' = \mathrm{N}_{\mathbb{K}/\mathbb{L}}(f)$, $g' = \mathrm{N}_{\mathbb{K}/\mathbb{L}}(g)$, and $h' = \mathrm{N}_{\mathbb{K}/\mathbb{L}}(h)$. 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,*

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

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

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

$$\sigma/n^C \leq |e(f)| \leq \sigma n^C \tag{13}$$

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 (13), 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. □

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,*

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

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

Remark. A more precise reasonable guess of $\|f'\|$ could be $\sigma^r n'$, yet our experiments show that this is an over-approximation, and f' tends to be significantly shorter for large r.

3.2 Lattice reduction in the subfield

We now apply a lattice reduction algorithm with block-size β to the lattice $\Lambda^q_{h'}$, and according to the approximation factor bound (11) we obtain a vector $(x', y') \in \Lambda^q_{h'}$ of norm:

$$\|(x', y')\| \leq \beta^{\Theta(2n'/\beta)} \cdot \lambda_1(\Lambda^q_{h'}) \tag{14}$$

$$\leq \beta^{\Theta(n/\beta r)} \cdot \|(f', g')\| \tag{15}$$

$$\leq \beta^{\Theta(n/\beta r)} \cdot (n\sigma)^{\Theta(r)} \tag{16}$$

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^q_h.

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 $(x', y') \in \Lambda^q_{h'}$ has length verifying*

$$\|(x', y')\| < \frac{q}{\|(f', g')\|}, \tag{17}$$

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:

1. $f'(F', G') - F'(f', g') = (0, q)$
2. $g'(F', G') - G'(f', g') = (-q, 0)$
3. $[f^{-1}]_q(f', g') = (1, h') \bmod q$.

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 $\mathrm{Vol}(M) = |\Delta_\mathbb{L}|q^{n'} = \mathrm{Vol}(\Lambda_{h'}^q)$, 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 $\mathrm{Vol}(\Lambda) \cdot \mathrm{Vol}(\Lambda^*) = \mathrm{Vol}(\Lambda_{h'}^q) = |\Delta_\mathbb{L}|q^{n'}$. By Lemma 2, we have

$$\mathrm{Vol}(\Lambda) \leq |\Delta_\mathbb{L}|^{1/2}\|(f', g')\|^{n'} \quad \text{and} \quad \mathrm{Vol}(\Lambda^*) \leq |\Delta_\mathbb{L}|^{1/2}\|s^*\|^{n'} \tag{18}$$

We deduce that $\lambda_1^* = \|s^*\| \geq \frac{q}{\|(f', g')\|}$. Therefore, the hypothesis (17) ensures that $\|(x', y')\| < \lambda_1^*$, and we conclude that $(x', y') \in \Lambda = (f', g')\mathcal{O}_\mathbb{L}$. □

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

$$\beta^{\Theta(n/\beta r)} \cdot (n\sigma)^{\Theta(r)} \leq q. \tag{19}$$

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 co-primality 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 as 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

$$x = L(x') \quad \text{and} \quad y = L(y') \cdot h/L(h') \bmod q. \tag{20}$$

We set $\tilde{f} = L(f')/f$, $\tilde{g} = L(g')/g$ and $\tilde{h} = L(h')/h$ and 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 (19) we have found a short multiple of (f, g):

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

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

We demonstrate the complexity of the subfield attack for two extreme cases. In both cases, 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 enough subfield so that a subfield $\mathbb{L}$ of the desired relative degree r exists . This condition is verified asymptotically for the typical choice $\mathbb{K} = \mathbb{Q}(\omega_{2^k})$.

In the first case, we set $q = 2^{\tilde{\Theta}(\lambda)}$, and the subfield attack is polynomial in the security parameter. For the second case, we show that as soon as the gap 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:

$$n = \Theta(\lambda^2 \log^2 \lambda), \quad q = \exp(\Theta(\lambda \log^2 \lambda)), \quad \sigma = \mathsf{poly}(\lambda)\,. \tag{21}$$

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 hermit bound, the vector found should not be shorter than:

$$\beta^{\Theta(n/\beta)} = \exp\left(\Theta(\lambda^2 \log^3 \lambda/\lambda)\right) > \lambda_1(q\mathcal{O}_\mathbb{K}). \tag{22}$$

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 allows 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

$$\beta^{\Theta(n/\beta r)} \cdot n^{\Theta(r)} = \exp\left(\Theta\left(\frac{\lambda^2 \log\lambda \log\log\lambda}{\lambda \log\lambda} + \lambda\log\lambda\right)\right) \tag{23}$$

$$= \exp\left(\Theta(\lambda \log\lambda \log\log\lambda)\right) < \sqrt{q}. \tag{24}$$

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

Quasi-polynomial q. We set

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

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 hermit bound, the vector found should not be shorter than:

$$\beta^{\Theta(n/\beta)} = \exp\left(\Theta\left(\log(\lambda)^{1+\varepsilon}\log\log(\lambda)\right)\right) > \lambda_1(q\mathcal{O}_\mathbb{K}). \tag{25}$$

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

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

$$\begin{aligned}\beta^{\Theta(n/\beta r)}\cdot n^{\Theta(r)} &= \exp\left(\Theta\left(\frac{\log^{1+\frac{4}{3}\epsilon}(\lambda)\log\log(\lambda)}{\log^{\frac{2}{3}\epsilon}(\lambda)} + \log^{1+2/3\,\epsilon}(\lambda)\right)\right)\\ &= \exp\left(\Theta\left(\log^{1+2/3\,\varepsilon}(\lambda)\log\log(\lambda)\right)\right) < \sqrt{q}.\end{aligned} \tag{26}$$

4 Applications

We apply our attack to FHE and MLM constructions from the literature. To match the definitions of rings and lengths often used in this litterature, we restrict our discussion to cyclotomic fields $\mathbb{K} = \mathbb{Q}(\omega_m)$, 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 cannonical inner product of its coefficients vector. The ring isomorphism $\mu : \mathcal{R} \to \mathcal{O}_\mathbb{K}$ 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 optimised and implemented in [DHS15].

LTV is motivated by [SS11] which shows that under certain choices of parameter 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 $> \sqrt{q}\cdot\mathsf{poly}(\lambda)$, then $h = [f/g]_q \in \mathbb{Z}_q[X]/(X^n+1)$ — 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 $[f/g]_q$ 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 for all choices of parameters.

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

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

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 away with the need for 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 $L = \mathcal{O}\left(\frac{\log(q)}{\log\log(q)+\log(n)}\right)$ [BLLN13, Theorem 2] for $\mathbb{Z}_2$ being the plaintext space.

On the other hand, setting the bounds on f, g to $\|f\|_\infty = \|g\|_\infty = B_{key} = 1$, 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 $\mathcal{O}\left(n^2\right)$. For correctness, it is required that the noise is $< q/4$. Hence, to evaluate a circuit of depth L, YASHE requires $q/4 > \mathcal{O}\left(n^{2L}\right)$ or $L = \mathcal{O}\left(\frac{\log(q)}{\log(n)}\right)$ 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 bootstrapable if it can evaluate circuits of depth $L = \mathcal{O}(\log(\log(q)) + \log(n))$. For $\|f\|_\infty = \|g\|_\infty = B_{key} = 1$ this implies

$$\log\log(q) + \log(n) < \log(q)/\log(n),$$
$$\log(n\log(q)) < \log(q)/\log(n),$$

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 $h \leftarrow \mathsf{NTRU}(\mathcal{R}, q, \sigma, \tau)$, 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 choose 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 which 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 big 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

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

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

$$\beta^{\Theta(n/\beta)} > \sqrt{q}$$
$$\Theta\left(\log^2(\lambda)\log\log^2(\lambda)\right) > \Theta\left(\log^2(\lambda)\right)$$

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

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

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 $\mathsf{poly}(\lambda)$ and some random z. For correctness, the latest improvements [ACLL15] require a modulus $q = \mathsf{poly}(\lambda)^\kappa$, 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 follows[8]. 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 Λ_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 $\approx \|r_0 \cdot g + m_0\| = \sigma_1^\star$ 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 quantum polynomial-time attack. The second approach relies on a statistical leak using the Gentry-Szydlo 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 impact 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 also is suspect 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 [EHKS14,CGS14,CDPR15,BS16] has lead to subexponential classical and polynomial-time quantum attacks against principal ideal lattices. 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 Eisentrager 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. [CDPR15], when g is sampled according to a spherical Gaussian (may it be discrete or continuous).

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

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

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 $n = \tilde{\Theta}(\lambda^{1+c})$— the $2^{\tilde{O}(n^{2/3})}$ algorithm of Biasse [Bia14] leads a classical attack in time sub-exponential in λ.

The statistical attack. This attacks consists in recovering $u\bar{u}$ and $\langle u \rangle$ and use 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 $\langle u \rangle = \langle ua_0 \rangle + \langle ua_1 \rangle$.[9]

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 a independent of u and follows a spherical Gaussian of parameter σ. It follows that the variance of ua_i leaks $u\bar{u}$: $\mathbb{E}[ua_i \cdot \overline{ua_i}] = \sigma^2 u\bar{u}$.

Given polynomially many samples x_i on can therefore recover $u\bar{u}$ up to a $1 + \frac{1}{\mathsf{poly}(\lambda)}$ 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, but unfortunately 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 algorithm[10]. 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

We report on the experiments we performed. As in the previous section, this report consider the ring $\mathcal{R} = \mathbb{Z}_q[X]/(X^n + 1) \simeq \mathcal{O}_\mathbb{K}$ for n a power of 2, and endowed with the cannonical inner product of its coefficients vector: Euclidean lengths are scaled so that $\|1_\mathcal{R}\| = 1$.

We chose q to be the first prime greater than 2^k for integers k in certain range, with the additional constraint that the field of order q should have an $2n$-th root of unity to allow the application of the number theoretic transform (NTT). In each case, the secret (f, g) was chosen as a uniform ternary vector, which, in the power of two case is an isotropic distribution of variance $\sigma^2 = 2/3$.

[9] 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 $\langle u \rangle$.

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

There are two trials for each set of parameters. We used LLL[11] for the lattice reduction step in the subfield case. For comparison, we also provide the prediction of the required BKZ block-size for a full field attack (ffa).

Instance		
	$\lfloor \log_2 q \rfloor$	Modulus bitsize
	$\log_2 \|\|(f', g')\|\|$	Euclidean length of the secret in the subfield
LLL	$\log_2 \|\|(x', y')\|\|$	Euclidean length of LLL's output in the subfield
in the	α	Tentative root approximation factor $\left(\frac{\|\|(x',y')\|\|}{\|\|(f',g')\|\|}\right)^{1/2n'}$
subfield	$\exists v$?	Do we have $(x', y') = v(f', g')$ for some $v \in \mathcal{O}_{\mathbb{L}}$?
Lifted	$\log_2 \|\|(x, y)\|\|$	euclidean length of vector found by lifting to the full field
solution	Success	Is the attack successful, i.e. do we have $\|\|(x, y)\|\| < q^{3/4}$?
BKZ in the	δ (ffa)	Root-hermite factor required for the ffa, with target length q
full field	β (ffa)	Block size to reach root hermite factor δ

Table 1. Explanation of reported parameters.

Our experimental data are summarized in Tables 2, 3 and 4, corresponding to parameter sets $(n, n') = (2^{11}, 2^7)$, $(n, n') = (2^{11}, 2^8)$ and $(n, n') = (2^{12}, 2^8)$ respectively.

Remark. In several cases, the value v such that $(x', y') = v(f', g')$ exists in $\mathbb{L}$, but is only a half integer: $2v \in \mathcal{O}_{\mathbb{L}}$, yet $v \notin \mathcal{O}_{\mathbb{L}}$. Those exceptions are marked with a asterisk (Yes*) in the "$\exists v$?" column. Those exceptions happened only when both $\mathrm{N}_{\mathbb{Q}}(f')$ and $\mathrm{N}_{\mathbb{Q}}(g')$ where even: the coprimality conditions of Theorem 2 was not verified, precisely, both norms had 2 as a common factor, and therefore $\langle 1 + \omega_{2n'} \rangle$ was a common factor [12]. Note that this nevertheless lead to a successful lift.

6 Conclusions

Practicality of the attack. The largest instance we were able to break in practice with our limited resources is for the set of parameter $n = 2^{12}, q \approx 2^{190}$. Choosing a relative degree $r = 16$, the attack required to run LLL in 512 dimensions, which took 120 hours, single-threaded, using Sage [Dev15] and Fplll [ABC+]. The direct lattice reduction attack, according to root-hermite-factor based predictions [CN11], should have required running BKZ with block-size ≈ 130, and in 8192 dimensions, which is hardly feasible with the current state-of-the art [CN11] (requiring more than 2^{70} CPU cycles). We conclude that the attack 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,LSS14,ACLL15].

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. Unlike the Hermite-factor regime, this case could very well be influenced by the presence of many short vectors rather than just a few. Our preliminary experiments exhibited undocumented behavior, and a careful study is required.

[11] More precisely, we used Fplll [ABC+] packaged in Sage [Dev15]

[12] the prime 2 totally ramifies in $\mathbb{L} = \mathbb{Q}(\omega_{2^t})$: $\langle 1 + \omega_{2n'} \rangle^{n'} = \langle 2 \rangle$

The second issue is that we do not know the actual size of the shortest vector of $\Lambda^q_{h'}$, all we know is that it is shorter than (f', g'). In several cases (Table 2) 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 the shortest vector for small relative degree r as it is most surely the shortest in the full field (i.e. when $r = 1$).

Immunity of NTRU encryption and BLISS signature schemes? If q is small enough, then our attacks should become inapplicable, even with the smallest possible relative dimension $r = 2$. Precisely, if (f', g') is not an unusually short vector of $\Lambda^q_{h'}$, then there is little hope that any lattice reduction strategy would lead to information on this vector. Quantitatively, this total immunity happens when $\|(f', g')\| \approx \sigma^2 \cdot n' > \sqrt{n'q/\pi e}$. This is unfortunately not the case of the parameters of NTRUEncrypt [HPS+15] and Bliss [DDLL13], for which (f', g') is an unusually short vector, but not by a very large factor (ranging from 2 to 10). It is plausible, especially for NTRUEncrypt, that this is close enough to total immunity to make the sub-field attack more costly than the full attack, but calls for further study.

Note that the immunity to our 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 extended.

Recommendations. Even if credible predictions were to be made, we strongly discourage basing a scheme on a set-up where this attack applies. 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 lattice reduction in the lattice $(f', g')\mathcal{O}_{\mathbb{L}}$ 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-Syzdlo algorithm [GS02,LS14] to obtain shorter vector in the full field (i.e. recovering a from $\mathrm{N}(a)$). More generally, the lifting step may be improved by considering the relative norm equation problem [FJP97]. One may recover a from $\mathrm{N}_{\mathbb{K}/\mathbb{L}}(a)$ using ideal factorization problem, followed by a recovery of short generator of principal ideals step; as mentionned before, those problems are now known to be classically sub-exponential [Bia14,CDPR15] or even polynomial for quantum computers [EHKS14,BS16].

Evaluating concrete security against regular lattice attacks is already a difficult exercise, and leaving open additional algebraic and statistical attack surfaces 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 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], Ring-LWE [LPR10]), it could be considered desirable to have efficient fall-back options ready to use, in case subfields induces other unforeseen weaknesses. While this work does not suggest an immediate threat to Ring-SIS and Ring-LWE, such a precaution is not unreasonable.

A worthy option was suggested in [Ber14]: rings of the form $\mathbb{Z}[X]/(X^p - X - 1)$. We are unfortunately unaware of a detailed study of that ring for lattice-based cryptography purposes. It has been remarked that the total absence of non-trivial automorphisms could be quite problematic for the batch-efficiency of certain FHE schemes as HElib [HS14]. Similarly, the alternative scheme FHEW [DM15] would suffer from the absence of roots of unity.

Another interesting option is to choose p as a safe prime[13] 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 because the Galois group has prime order $\frac{p-1}{2}$, it has no proper subgroups, and $\mathbb{K}$ has no proper subfields. In practice, such rings may perform decently well, since, for example, the fast Fourier transform can benefit from a two-fold acceleration when the Fourier coefficients are all reals.

[13] A safe prime p is an odd prime, such that $\frac{p-1}{2}$ is also prime. The terminology relates to weaknesses in RSA and Discrete-logarithm introduced by the smoothness of $p-1$ [Pol74].

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/fplll`.

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.

Ber14. Dan Bernstein. A subfield-logarithm attack against ideal lattices. `http://blog.cr.yp.to/20140213-ideal.html`, Febuary 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. J.-F. Biasse and F. Song. A polynomial time quantum algorithm for computing class groups and solving the principal ideal problem in arbitrary degree number fields. `http://www.lix.polytechnique.fr/Labo/Jean-Francois.Biasse/`, 2016. In preparation.

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.

CDPR15. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. Cryptology ePrint Archive, Report 2015/313, 2015. `http://eprint.iacr.org/2015/313`.

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_CRYPTO/S07_Systems_and_Attacks/S07_Groves_Annex.pdf`.

CLS15. Hao Chen, Kristin Lauter, and Katherine E. Stange. Attacks on search rlwe. Cryptology ePrint Archive, Report 2015/971, 2015. `http://eprint.iacr.org/`.

CLT13. Jean-Sébastien Coron, Tancrède 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/`.

DDLL13. Léo Ducas, Alain Durmus, Tancrède 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.

DM15. Léo Ducas and Daniele Micciancio. FHEW: Bootstrapping homomorphic encryption in less than a second. In Elisabeth Oswald and Marc Fischlin, editors, *EUROCRYPT 2015, Part I*, volume 9056 of *LNCS*, pages 617–640. Springer, Heidelberg, April 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.

FJP97. Claus Fieker, Andreas Jurk, and M Pohst. On solving relative norm equations in algebraic number fields. *Mathematics of Computation of the American Mathematical Society*, 66(217):399–410, 1997.

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 Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 182–194. Springer, Heidelberg, May 2001.

GG14. Juan A. Garay and Rosario Gennaro, editors. *CRYPTO 2014, Part I*, volume 8616 of *LNCS*. Springer, Heidelberg, August 2014.

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 Menezes [Men07], pages 150–169.

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.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, pages 267–288, 1998.

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`.

HS07. Guillaume Hanrot and Damien Stehlé. Improved analysis of kannan's shortest lattice vector algorithm. In Menezes [Men07], pages 170–186.

HS14. Shai Halevi and Victor Shoup. Algorithms in HElib. In Garay and Gennaro [GG14], pages 554–571.

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.

JHW06. Joseph H. Silverman Jeffrey Hoffstein 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.

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. *Desings 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. Tancrède 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.

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 Garay and Gennaro [GG14], pages 280–296.

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.

Men07. Alfred Menezes, editor. *CRYPTO 2007*, volume 4622 of *LNCS*. Springer, Heidelberg, August 2007.

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.

Pol74. John M Pollard. Theorems on factorization and primality testing. In *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 76, pages 521–528. Cambridge Univ Press, 1974.

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.

Was97. L.C. Washington. *Introduction to Cyclotomic Fields*. Graduate Texts in Mathematics. Springer New York, 1997.

Table 2. Experiment report. Parameters set $m = 2^{12}$ ($n = 2^{11}$), $r = 2^4$, $m' = 2^8$ ($n' = 2^7$).

Instance		Subfield LLL			Lifted		Fullfield BKZ	
$\lfloor \lg q \rfloor$	$\lg \|\|(f', g')\|\|$	$\lg \|\|(x', y')\|\|$	α (traf)	$\exists v$?	$\lg \|\|(x, y)\|\|$	Success	δ (ffa)	β (ffa)
180	81.16	82.21	1.0028	Yes	82.81	Yes	1.0153	11
	82.42	82.52	1.0003	Yes	82.95	Yes	1.0153	11
179	82.28	82.42	1.0004	Yes	82.76	Yes	1.0153	13
	82.90	82.92	1.0001	Yes	83.26	Yes	1.0153	13
178	81.93	82.74	1.0022	Yes	83.33	Yes	1.0152	14
	82.63	82.28	0.9990	Yes	82.88	Yes	1.0152	14
177	82.41	82.62	1.0006	Yes	83.50	Yes	1.0151	15
	83.35	82.48	0.9977	Yes	82.97	Yes	1.0151	15
176	81.97	82.62	1.0018	Yes	83.74	Yes	1.0150	16
	84.37	83.04	0.9964	Yes	83.58	Yes	1.0150	16
175	81.60	81.82	1.0006	Yes	82.63	Yes	1.0149	17
	80.94	81.84	1.0024	Yes	82.62	Yes	1.0149	17
174	83.85	82.76	0.9971	Yes	83.30	Yes	1.0148	18
	82.15	82.77	1.0017	Yes	83.47	Yes	1.0148	18
173	82.10	82.41	1.0008	Yes	83.15	Yes	1.0147	19
	82.20	82.56	1.0010	Yes	83.22	Yes	1.0147	19
172	82.23	82.15	0.9998	Yes	82.79	Yes	1.0147	20
	83.12	82.75	0.9990	Yes	83.33	Yes	1.0147	20
171	83.05	83.37	1.0009	Yes	84.11	Yes	1.0146	21
	83.00	83.03	1.0001	Yes	83.54	Yes	1.0146	21
170	84.24	83.02	0.9967	Yes	83.45	Yes	1.0145	22
	82.45	82.84	1.0011	Yes*	83.15	Yes	1.0145	22
169	83.31	82.82	0.9987	Yes	83.53	Yes	1.0144	23
	83.99	82.50	0.9960	Yes	83.44	Yes	1.0144	23
168	84.01	82.69	0.9965	Yes	83.32	Yes	1.0143	24
	82.91	82.13	0.9979	Yes	82.56	Yes	1.0143	24
167	83.33	82.66	0.9982	Yes	83.31	Yes	1.0142	25
	82.67	82.96	1.0008	Yes*	83.76	Yes	1.0142	25
166	82.88	82.38	0.9986	Yes	82.85	Yes	1.0141	26
	83.44	82.50	0.9975	Yes	82.87	Yes	1.0141	26
165	82.75	82.99	1.0006	Yes	83.50	Yes	1.0141	27
	82.74	82.55	0.9995	Yes	83.33	Yes	1.0141	27
164	82.43	89.67	1.0198	No	167.67	No	1.0140	28
	81.44	89.78	1.0228	No	167.73	No	1.0140	28
163	81.16	89.45	1.0227	No	166.69	No	1.0139	29
	84.57	89.25	1.0128	No	166.69	No	1.0139	29
162	82.60	88.73	1.0168	No	165.71	No	1.0138	30
	82.67	88.95	1.0172	No	165.71	No	1.0138	30
161	82.84	88.44	1.0153	No	164.70	No	1.0137	31
	81.97	88.20	1.0170	No	164.72	No	1.0137	31
160	80.82	87.73	1.0189	No	163.68	No	1.0136	32
	83.96	87.90	1.0107	No	163.72	No	1.0136	32

Each of this run took about 3.5 Hours. The predicted bitsize of $\|(f', g')\|$ was $\log_2(\sqrt{2}\sigma^r n^{r/2}/\sqrt{r}) = 89.82$.

Table 3. Experiment report. Parameters set $m = 2^{12}$ ($n = 2^{11}$), $r = 2^3$, $m' = 2^9$ ($n = 2^8$).

Instance		Subfield LLL			Lifted		Fullfield BKZ	
$\lfloor \lg q \rfloor$	$\lg \|\|(f', g')\|\|$	$\lg \|\|(x', y')\|\|$	α (traf)	$\exists v$?	$\lg \|\|(x, y)\|\|$	Success	δ (ffa)	β (ffa)
110	42.27	47.72	1.0074	Yes	49.20	Yes	1.0094	98
	41.85	47.55	1.0078	Yes	48.01	Yes	1.0094	98
109	42.15	47.64	1.0075	Yes	48.22	Yes	1.0093	100
	41.88	47.48	1.0076	Yes	47.93	Yes	1.0093	100
108	42.12	48.11	1.0081	Yes	48.71	Yes	1.0092	102
	42.04	48.13	1.0083	Yes	48.51	Yes	1.0092	102
107	42.28	47.89	1.0076	Yes	48.07	Yes	1.0091	104
	42.19	47.69	1.0075	Yes	48.21	Yes	1.0091	104
106	42.11	47.98	1.0080	Yes	48.46	Yes	1.0090	106
	42.15	48.01	1.0080	Yes	48.58	Yes	1.0090	106
105	41.53	47.52	1.0081	Yes*	47.94	Yes	1.0089	108
	41.73	47.53	1.0079	Yes	48.23	Yes	1.0089	108
104	42.18	47.94	1.0078	Yes	48.17	Yes	1.0088	110
	42.19	47.79	1.0076	Yes*	48.26	Yes	1.0088	110
103	42.67	47.89	1.0071	Yes	48.36	Yes	1.0088	112
	41.85	47.59	1.0078	Yes	47.94	Yes	1.0088	112
102	42.26	47.77	1.0075	Yes	48.52	Yes	1.0087	114
	41.72	47.52	1.0079	Yes	47.91	Yes	1.0087	114
101	41.77	47.72	1.0081	Yes	47.96	Yes	1.0086	117
	42.07	47.76	1.0077	Yes	48.26	Yes	1.0086	117
100	41.48	47.77	1.0085	Yes	48.16	Yes	1.0085	119
	42.14	47.71	1.0076	Yes	48.15	Yes	1.0085	119
99	41.83	47.67	1.0079	Yes	48.11	Yes	1.0084	121
	42.02	47.70	1.0077	Yes	48.03	Yes	1.0084	121
98	42.57	48.05	1.0074	Yes	48.42	Yes	1.0083	123
	41.74	47.88	1.0084	Yes	48.78	Yes	1.0083	123
97	42.60	47.80	1.0071	Yes	48.36	Yes	1.0082	126
	42.51	48.10	1.0076	Yes	48.47	Yes	1.0082	126
96	41.89	47.46	1.0076	Yes	48.01	Yes	1.0082	128
	41.87	48.09	1.0085	Yes	48.36	Yes	1.0082	128
95	42.25	47.75	1.0075	Yes	48.15	Yes	1.0081	131
	41.85	47.96	1.0083	Yes	48.59	Yes	1.0081	131
94	41.99	63.63	1.0297	No	97.71	No	1.0080	133
	42.57	63.32	1.0285	No	97.70	No	1.0080	133
93	41.87	62.75	1.0287	No	96.69	No	1.0079	136
	41.90	63.02	1.0290	No	96.69	No	1.0079	136
92	42.01	62.05	1.0275	No	95.70	No	1.0078	139
	42.79	62.12	1.0265	No	95.69	No	1.0078	139
91	42.10	62.08	1.0274	No	94.70	No	1.0077	141
	41.74	61.39	1.0270	No	94.69	No	1.0077	141
90	42.15	61.28	1.0262	No	93.73	No	1.0076	144
	42.07	61.08	1.0261	No	93.72	No	1.0076	144
89	41.86	60.54	1.0256	No	92.72	No	1.0076	147
	42.20	60.82	1.0255	No	92.70	No	1.0076	147

Each of this run took about 50 Hours. The predicted bitsize of $\|(f', g')\|$ was $\log_2(\sigma^r n^{r/2}/\sqrt{r}) = 48.66$.

Table 4. Experiment report. Parameters set $m = 2^{13}$ ($n = 2^{12}$), $r = 2^4$, $m' = 2^9$ ($n = 2^8$).

Instance		Subfield LLL			Lifted		Fullfield BKZ	
$\lfloor \lg q \rfloor$	$\lg \|\|(f', g')\|\|$	$\lg \|\|(x', y')\|\|$	α (traf)	$\exists v$?	$\lg \|\|(x, y)\|\|$	Success	δ (ffa)	β (ffa)
240	90.60	94.55	1.0054	Yes	95.13	Yes	1.0102	82
	90.78	94.67	1.0053	Yes	95.22	Yes	1.0102	82
235	91.16	95.06	1.0053	Yes	95.63	Yes	1.0100	86
	91.08	94.50	1.0046	Yes	95.17	Yes	1.0100	86
230	90.44	95.00	1.0062	Yes	95.70	Yes	1.0098	90
	90.58	94.62	1.0055	Yes	95.40	Yes	1.0098	90
225	91.57	95.56	1.0054	Yes*	96.28	Yes	1.0096	94
	90.19	94.68	1.0061	Yes	95.32	Yes	1.0096	94
220	90.62	95.01	1.0060	Yes	95.74	Yes	1.0094	98
	90.98	94.65	1.0050	Yes	95.34	Yes	1.0094	98
215	90.33	94.57	1.0057	Yes*	95.13	Yes	1.0091	103
	91.52	94.77	1.0044	Yes	95.26	Yes	1.0091	103
210	91.43	95.33	1.0053	Yes	95.81	Yes	1.0089	108
	90.48	94.73	1.0058	Yes	95.28	Yes	1.0089	108
205	91.59	94.64	1.0041	Yes*	95.04	Yes	1.0087	113
	92.93	94.50	1.0021	Yes	95.10	Yes	1.0087	113
200	90.44	94.57	1.0056	Yes	95.10	Yes	1.0085	119
	90.03	94.84	1.0065	Yes	95.51	Yes	1.0085	119
195	92.52	94.59	1.0028	Yes	95.37	Yes	1.0083	125
	92.60	94.74	1.0029	Yes	95.90	Yes	1.0083	125
190	90.27	94.57	1.0058	Yes	95.14	Yes	1.0081	131
	90.20	94.17	1.0054	Yes*	94.74	Yes	1.0081	131
185	91.02	108.99	1.0246	No	189.20	No	1.0079	137
	91.17	108.66	1.0240	No	189.22	No	1.0079	137
180	91.27	106.31	1.0206	No	184.20	No	1.0076	144
	91.29	106.39	1.0207	No	184.21	No	1.0076	144
175	90.08	103.93	1.0189	No	179.20	No	1.0074	151
	91.30	103.31	1.0164	No	179.21	No	1.0074	151

Each of this run took about 120 Hours. The predicted bitsize of $\|(f', g')\|$ was $\log_2(\sigma^r n^{r/2}/\sqrt{r}) = 97.82$.