Making NTRU as Secure as Worst-Case Problems over Ideal Lattices

Damien Stehlé[1] and Ron Steinfeld[2]

[1] CNRS, Laboratoire LIP (U. Lyon, CNRS, ENS Lyon, INRIA, UCBL),
46 Allée d'Italie, 69364 Lyon Cedex 07, France.
damien.stehle@gmail.com – http://perso.ens-lyon.fr/damien.stehle

[2] Centre for Advanced Computing - Algorithms and Cryptography,
Department of Computing, Macquarie University, NSW 2109, Australia
ron.steinfeld@mq.edu.au – http://web.science.mq.edu.au/~rons

Abstract. NTRUEncrypt, proposed in 1996 by Hoffstein, Pipher and Silverman, is the fastest known lattice-based encryption scheme. Its moderate key-sizes, excellent asymptotic performance and conjectured resistance to quantum computers could make it a desirable alternative to factorisation and discrete-log based encryption schemes. However, since its introduction, doubts have regularly arisen on its security. In the present work, we show how to modify NTRUEncrypt to make it provably secure in the standard model, under the assumed quantum hardness of standard worst-case lattice problems, restricted to a family of lattices related to some cyclotomic fields. Our main contribution is to show that if the secret key polynomials are selected by rejection from discrete Gaussians, then the public key, which is their ratio, is statistically indistinguishable from uniform over its domain. The security then follows from the already proven hardness of the R-LWE problem.

Keywords. Lattice-based cryptography, NTRU, provable security.

1 Introduction

NTRUEncrypt, devised by Hoffstein, Pipher and Silverman, was first presented at the Crypto'96 rump session [14]. Although its description relies on arithmetic over the polynomial ring $\mathbb{Z}_q[x]/(x^n - 1)$ for n prime and q a small integer, it was quickly observed that breaking it could be expressed as a problem over Euclidean lattices [6]. At the ANTS'98 conference, the NTRU authors gave an improved presentation including a thorough assessment of its practical security against lattice attacks [15]. We refer to [13] for an up-to-date account on the past 15 years of security and performance analyses. Nowadays, NTRUEncrypt is generally considered as a reasonable alternative to the encryption schemes based on integer factorisation and discrete logarithm over finite fields and elliptic curves, as testified by its inclusion in the IEEE P1363 standard [18]. It is also often considered as the most viable post-quantum public-key encryption (see, e.g., [31]).

In parallel to a rising number of attacks and practical improvements on `NTRUEncrypt` the (mainly) theoretical field of provably secure lattice-based cryptography has steadily been developed. It originated in 1996 with Ajtai's acclaimed worst-case to average-case reduction [2], leading to a collision-resistant hash function that is as hard to break as solving several worst-case problems defined over lattices. Ajtai's average-case problem is now referred to as the *Small Integer Solution* problem (SIS). Another major breakthrough in this field was the introduction in 2005 of the *Learning with Errors* problem (LWE) by Regev [32]: LWE is both hard on the average (worst-case lattice problems quantumly reduce to it), and sufficiently flexible to allow for the design of cryptographic functions. In the last few years, many cryptographic schemes have been introduced that are provably as secure as LWE and SIS are hard (and thus provably secure, assuming the worst-case hardness of lattice problems). These include CPA and CCA secure encryption schemes, identity-based encryption schemes, digital signatures, *etc* (see [32, 29, 11, 5, 1] among others, and the surveys [25, 33]).

The main drawback of cryptography based on LWE and SIS is its limited efficiency. A key typically contains a random matrix defined over $\mathbb{Z}_q$ for a small q, whose dimension is linear in the security parameter; consequently, the space and time requirements seem bound to be at least quadratic with respect to the security parameter. In 2002, Micciancio [23] succeeded in restricting SIS to structured matrices while preserving a worst-case to average-case reduction. The worst-case problem is a restriction of a standard lattice problem to the specific family of cyclic lattices. The structure of Micciancio's matrices allows for an interpretation in terms of arithmetic in the ring $\mathbb{Z}_q[x]/(x^n - 1)$, where n is the dimension of the worst-case lattices and q is a small prime. Micciancio's construction leads to a family of pre-image resistant hash functions, with complexity quasi-linear in n. Peikert, Rosen, Lyubashevsky and Micciancio [30, 19] later suggested to change the ring to $\mathbb{Z}_q[x]/\Phi$ with a Φ that is irreducible over the rationals, sparse, and with small coefficients (e.g., $\Phi = x^n + 1$ for n a power of 2). The resulting hash function was proven collision-resistant under the assumed hardness of the modified average-case problem, called Ideal-SIS. The latter was itself proven at least as hard as the restrictions of standard worst-case lattice problems to a specific class of lattices (called ideal lattices). In 2009, Stehlé *et al.* [35] introduced a structured variant of LWE, which they proved as hard as Ideal-SIS (under a quantum reduction), and allowed for the design of an asymptotically efficient CPA-secure encryption scheme. In an independent concurrent work, Lyubashevsky *et al.* [21] proposed a ring variant of LWE, called R-LWE, whose great flexibility allows for more natural (and efficient) cryptographic constructions.

Our Results. The high efficiency and industrial standardization of `NTRUEncrypt` strongly motivate a theoretically founded study of its security. Indeed, in the absence of such a study so far, its security has remained in doubt over the last 15 years since its publication. In this paper, we address this problem. We prove that a mild modification of `NTRUEncrypt` is CPA-secure, assuming the quantum hardness of standard worst-case problems over ideal lattices (for $\Phi = x^n + 1$ with n a power of 2). The `NTRUEncrypt` modifications are summarized below. We stress that our main goal in this paper is to provide a firm theoretical grounding for the security of NTRUEncrypt in the asymptotic sense. We leave to future work the consideration of practical issues, in particular the selection of concrete parameters for given security levels. As for other lattice-based schemes, the latter requires evaluation of security against practical lattice reduction attacks, which is out of the scope of the current work.

Our main contribution is the modification and analysis of the key generation algorithm. The secret key consists of two sparse polynomials of degrees $< n$ and coefficients of magnitude at most c, for a small constant c (typically, $c \in \{2, 3\}$). The public key is their quotient in $\mathbb{Z}_q[x]/(x^n - 1)$ (the denominator is resampled if it is not invertible). A simple information-theoretic argument shows that the public key cannot be uniformly distributed in the whole ring. It may be possible to extend the results of [4] to show that it is "well-spread" in the ring, but it still would not suffice for showing its cryptographic pseudorandomness, which seems necessary for exploiting the established hardness of R-LWE. To achieve a public key distribution statistically close to uniform, we sample the secret key polynomials according to a discrete Gaussian with standard deviation $\approx q^{1/2}$. An essential ingredient, which could be of independent interest, is a new regularity result for the ring $R_q := \mathbb{Z}_q[x]/(x^n + 1)$ when the polynomial $x^n + 1$ (with n a power of 2) has n factors modulo prime q: given $a_1, \ldots, a_m$ uniform in R_q, we would like $\sum_{i \leq m} s_i a_i$ to be within exponentially small statistical distance to uniformity, with small random s_i's and small m. Note that a similar regularity bound can be obtained with an FFT-based technique recently developed by Lyubashevsky, Peikert and Regev [22]. An additional difficulty in the public-key 'uniformity' proof, which we handle via an inclusion-exclusion argument, is that we need the s_i's to be invertible in R_q (the denominator of the public key is one such s_i): we thus sample according to a discrete Gaussian, and reject the sample if it is not invertible.

Brief comparison of NTRUEncrypt and its provably secure variant

Let R_{NTRU} be the ring $\mathbb{Z}[x]/(x^n - 1)$ with n prime. Let q be a medium-size integer (typically, either a prime or a power of 2 of the same order of magnitude as n). Finally, let $p \in R_{\text{NTRU}}$ with small coefficients, co-prime with q and such that the plaintext space R_{NTRU}/p is large (typically, one may take $p \in \{2, 3\}$ or $p = x + 2$).

The NTRUEncrypt secret key is a pair of polynomials $(f, g) \in R_{\text{NTRU}}^2$ that are sampled randomly with large prescribed proportions of zeros, and with their other coefficients having small magnitude. For improved decryption efficiency, one may choose f such that $f = 1 \bmod p$ (a typical choice [17] is to choose g and F with coefficients in $\{0, 1\}$ and set $f = 1 + p \cdot F$). With high probability, the polynomial f is invertible modulo q and modulo p, and if that is the case, the public-key is $h = pg/f \bmod q$ (otherwise, the key generation process is restarted). To encrypt a message $M \in R_{\text{NTRU}}/p$, one samples a random element $s \in R_{\text{NTRU}}$ of small Euclidean norm and computes the ciphertext $C = hs + M \bmod q$. The following procedure allows the owner of the secret key to decrypt:

- Compute $fC \bmod q$. If C was properly generated, this gives $pgs + fM \bmod q$. Since p, g, s, f, M have small coefficients, it can be expected that after reduction modulo q the obtained representative is $pgs + fM$ (in R_{NTRU}).
- Reduce the latter modulo p. This should provide $fM \bmod p$.
- Multiply the result of the previous step by the inverse of f modulo p (this step becomes vacuous if $f = 1 \bmod p$).

Note that the encryption process is probabilistic, and that decryption errors can occur for some sets of parameters. However, it is possible to arbitrarily decrease the decryption error probability, and even to eliminate it completely.

In order to achieve CPA-security we make a few modifications to the original `NTRUEncrypt` (which preserve its quasi-linear time and space complexity):

1. We replace R_{NTRU} by $R = \mathbb{Z}[x]/(x^n + 1)$ with n a power of 2. We will exploit the irreducibility of $x^n + 1$ and the fact that R is the ring of integers of a cyclotomic number field.
2. We choose a prime $q \leq \mathcal{P}oly(n)$ such that $f = x^n + 1 \bmod q$ has n distinct linear factors (i.e., $q = 1 \bmod 2n$). This allows us to use the search to decision reduction for R-LWE with ring $R_q := R/q$ (see [21]), and also to take $p = 2$.
3. We sample f and g from discrete Gaussians over R, rejecting the samples that are not invertible in R_q. We show that $f/g \bmod q$ is essentially uniformly distributed over the set of invertible elements of R_q. We may also choose $f = pf' + 1$ with f' sampled from a discrete Gaussian, to simplify decryption.
4. We add a small error term e in the encryption: $C = hs + pe + M \bmod q$, with s and e sampled from the R-LWE error distribution. This allows us to derive CPA security from the hardness of a variant of R-LWE (which is similar to the variant of LWE from [3, Se. 3.1]).

Work in progress and open problems

Our study is restricted to the sequence of rings $\mathbb{Z}[x]/\Phi_n$ with $\Phi_n = x^n + 1$ with n a power of 2. An obvious drawback is that this does not allow for much flexibility on the choice of n (in the case of NTRU, the degree was assumed prime, which provides more freedom). The R-LWE problem is known to be hard when Φ_n is cyclotomic [21]. We chose to restrict ourselves to cyclotomic polynomials of order a power of 2 because it makes the error generation of R-LWE more efficient, and the description of the schemes simpler to follow. Our results are likely to hold for more general rings than those we considered. An interesting choice could be the cyclotomic rings of prime order (i.e., $\Phi_n = (x^n - 1)/(x - 1)$ with n prime) as these are large subrings of the NTRU rings (and one might then be able to show that the hardness carries over to the NTRU rings).

An interesting open problem is to obtain a CCA secure variant of our scheme in the standard model, while maintaining its efficiency (within constant factors). The selection of concrete parameters based on practical security estimates for the worst-case SVP in ideal lattices or the average-case hardness of R-LWE/Ideal-SIS is also left as a future work.

The authors of NTRUEncrypt also proposed a signature scheme based on a similar design. The history of NTRUSign started with NSS in 2001 [16]. Its development has been significantly more hectic and controversial, with a series of cryptanalyses and repairs (see the survey [13]). In a work in progress, we construct a variant of NTRUSign with unforgeability related to the worst-case hardness of standard problems over ideal lattices, in the random oracle model. Our construction modifies the NTRUSign key generation and adapts the GPV signature scheme [11] to this setting.

Like NTRUEncrypt, Gentry's somewhat homomorphic scheme [9] also has ciphertexts consisting of a single ring element. It also admits a security proof under the assumed quantum hardness of standard worst-case problems over ideal lattices [10]. Our security analysis for the modified NTRUEncrypt scheme allows encrypting and decrypting $\Omega(n)$ plaintext bits for $\widetilde{O}(n)$ bit operations, while achieving security against $2^{g(n)}$-time attacks, for any $g(n)$ that is $\Omega(\log n)$ and $o(n)$, assuming the worst-case hardness of $\mathcal{P}oly(n)$-Ideal-SVP against $2^{O(g(n))}$-time quantum algorithms. The latter assumption is believed to be valid for any $g(n) = o(n)$. Gentry's analysis from [10, 8] can be generalized to handle $2^{g(n)}$-time attacks while encrypting and decrypting $O(g(n))$ plaintext bits for $\widetilde{O}(n)$ bit operations, under the assumed hardness of $2^{\Omega(g(n))}$-Ideal-SVP against $2^{O(g(n))}$-time quantum algorithms. The latter assumption is known to be invalid when $g(n) = \widetilde{\Omega}(\sqrt{n})$ (using [34]), thus limiting the attacker's strength the analysis can handle. On the other hand, Gentry's scheme allows homomorphic additions and multiplications, whereas ours seems restricted to additions. Our scheme and Gentry's seem to be closely related, and we leave to future work the further investigation of this relation.

Notation. We denote by $\rho_\sigma(\boldsymbol{x})$ (resp. ν_σ) the standard n-dimensional Gaussian function (resp. distribution) with center $\mathbf{0}$ and variance σ, i.e., $\rho_\sigma(\boldsymbol{x}) = \exp(-\pi\|\boldsymbol{x}\|^2/\sigma^2)$ (resp. $\nu_\sigma(\boldsymbol{x}) = \rho_\sigma(\boldsymbol{x})/\sigma^n$). We denote by $\mathrm{Exp}(\mu)$ the exponential distribution on $\mathbb{R}$ with mean μ and by $U(E)$ the uniform distribution over a finite set E . If D_1 and D_2 are two distributions on discrete domain E, their statistical distance is $\Delta(D_1; D_2) = \frac{1}{2}\sum_{x\in E}|D_1(x) - D_2(x)|$. We write $z \hookleftarrow D$ when the random variable z is sampled from the distribution D.

2 A Few Background Results

A (full-rank) *lattice* is a set of the form $L = \sum_{i\leq n}\mathbb{Z}\boldsymbol{b}_i$, where the $\boldsymbol{b}_i$'s are linearly independent vectors in $\mathbb{R}^n$. The integer n is called the *lattice dimension*, and the $\boldsymbol{b}_i$'s are called a *basis* of L. The *minimum* $\lambda_1(L)$ (resp. $\lambda_1^\infty(L)$) is the Euclidean (resp. infinity) norm of any shortest vector of $L \setminus \mathbf{0}$. If $B = (\boldsymbol{b}_i)_i$ is a basis matrix of L, the *fundamental parallelepiped* of B is the set $\mathcal{P}(B) = \{\sum_{i\leq n} c_i\boldsymbol{b}_i : c_i \in [0, 1)\}$. The volume $|\det B|$ of $\mathcal{P}(B)$ is an invariant of the lattice L which we denote by $\det L$. Minkowski's theorem states that $\lambda_1(L) \leq \sqrt{n}(\det L)^{1/n}$. More generally, the k-th *minimum* $\lambda_k(L)$ for $k \leq n$ is defined as the smallest r such that L contains $\geq k$ linearly independent vectors of norm $\leq r$. The *dual* of L is the lattice $\widehat{L} = \{\boldsymbol{c} \in \mathbb{R}^n : \forall i, \langle\boldsymbol{c}, \boldsymbol{b}_i\rangle \in \mathbb{Z}\}$.

For a lattice $L \subseteq \mathbb{R}^n$, $\sigma > 0$ and $\boldsymbol{c} \in \mathbb{R}^n$, we define the *lattice Gaussian distribution* of support L, deviation σ and center $\boldsymbol{c}$ by $D_{L,\sigma,\boldsymbol{c}}(\boldsymbol{b}) = \frac{\rho_{\sigma,\boldsymbol{c}}(\boldsymbol{b})}{\rho_{\sigma,\boldsymbol{c}}(L)}$, for any $\boldsymbol{b} \in L$. We will omit the subscript $\boldsymbol{c}$ when it is $\boldsymbol{0}$. We extend the definition of $D_{L,\sigma,\boldsymbol{c}}$ to any $M \subseteq L$ (not necessarily a sublattice), by setting $D_{M,\sigma,\boldsymbol{c}}(\boldsymbol{b}) = \frac{\rho_{\sigma,\boldsymbol{c}}(\boldsymbol{b})}{\rho_{\sigma,\boldsymbol{c}}(M)}$. For $\delta > 0$, we define the *smoothing parameter* $\eta_\delta(L)$ as the smallest $\sigma > 0$ such that $\rho_{1/\sigma}(\widehat{L} \setminus \boldsymbol{0}) \leq \delta$. It quantifies how large σ needs to be for $D_{L,\sigma,\boldsymbol{c}}$ to behave like a continuous Gaussian. We will typically consider $\delta = 2^{-n}$.

Lemma 1 ([24, Le. 3.3]). *For any full-rank lattice $L \subseteq \mathbb{R}^n$ and $\delta \in (0, 1)$, we have $\eta_\delta(L) \leq \sqrt{\ln(2n(1 + 1/\delta))/\pi} \cdot \lambda_n(L)$.*

Lemma 2 ([28, Le. 3.5]). *For any full-rank lattice $L \subseteq \mathbb{R}^n$ and $\delta \in (0, 1)$, we have $\eta_\delta(L) \leq \sqrt{\ln(2n(1 + 1/\delta))/\pi}/\lambda_1^\infty(\widehat{L})$.*

Lemma 3 ([24, Le. 4.4]). *For any full-rank lattice $L \subseteq \mathbb{R}^n$, $\boldsymbol{c} \in \mathbb{R}^n$, $\delta \in (0, 1)$ and $\sigma \geq \eta_\delta(L)$, we have $\Pr_{\boldsymbol{b} \hookleftarrow D_{L,\sigma,\boldsymbol{c}}}[\|\boldsymbol{b}\| \geq \sigma\sqrt{n}] \leq \frac{1+\delta}{1-\delta} 2^{-n}$.*

Lemma 4 ([11, Cor. 2.8]). *Let $L' \subseteq L \subseteq \mathbb{R}^n$ be full-rank lattices. For any $\boldsymbol{c} \in \mathbb{R}^n$, $\delta \in (0, 1/2)$ and $\sigma \geq \eta_\delta(L')$, we have $\Delta(D_{L,\sigma,\boldsymbol{c}} \bmod L'; U(L/L')) \leq 2\delta$.*

Lemma 5 ([11, Th. 4.1]). *There exists a polynomial-time algorithm that takes as input any basis $(\boldsymbol{b}_i)_i$ of any lattice $L \subseteq \mathbb{Z}^n$ and $\sigma = \omega(\sqrt{\log n}) \max \|\boldsymbol{b}_i\|$ (resp. $\sigma = \Omega(\sqrt{n}) \max \|\boldsymbol{b}_i\|$), and returns samples from a distribution whose statistical distance to $D_{L,\sigma}$ is negligible (resp. exponentially small) with respect to n.*

The most famous lattice problem is SVP. Given a basis of a lattice L, it aims at finding a shortest vector in $L \setminus \boldsymbol{0}$. It can be relaxed to γ-SVP by asking for a non-zero vector that is no longer than $\gamma(n)$ times a solution to SVP, for a prescribed function $\gamma(\cdot)$. It is believed that no subexponential quantum algorithm solves the computational variants of γ-SVP in the worst case, for any $\gamma \leq \mathcal{P}oly(n)$. The smallest γ which is known to be achievable in polynomial time is exponential, up to poly-logarithmic factors in the exponent ([34, 26]).

Ideal lattices and algebraic number theory

Ideal lattices. Let n a power of 2 and $\Phi = x^n + 1$ (which is irreducible over $\mathbb{Q}$). Let R be the ring $\mathbb{Z}[x]/\Phi$. An (integral) ideal I of R is a subset of R closed under addition and multiplication by arbitrary elements of R. By mapping polynomials to the vectors of their coefficients, we see that an ideal $I \neq 0$ corresponds to a full-rank sublattice of $\mathbb{Z}^n$: we can thus view I as both a lattice and an ideal. An *ideal lattice* for Φ is a sublattice of $\mathbb{Z}^n$ that corresponds to a non-zero ideal $I \subseteq R$. The *algebraic norm* $\mathcal{N}(I)$ is the cardinality of the additive group R/I. It is equal to $\det I$, where I is regarded as a lattice. Any non-zero ideal I of R satisfies $\lambda_n(I) = \lambda_1(I)$. In the following, an ideal lattice will implicitly refer to a Φ-ideal lattice.

By restricting SVP (resp. γ-SVP) to instances that are ideal lattices, we obtain Ideal-SVP (resp. γ-Ideal-SVP). The latter is implicitly parameterized by the sequence of polynomials $\Phi_n = x^n + 1$, where n is restricted to powers of 2. No algorithm is known to perform non-negligibly better for (γ-)Ideal-SVP than for (γ-)SVP.

Properties of the ring R. For $v \in R$ we denote by $\|v\|$ its Euclidean norm (as a vector). We define the multiplicative *expansion factor* $\gamma_\times(R)$ by $\gamma_\times(R) = \max_{u,v\in R} \frac{\|u\times v\|}{\|u\|\cdot\|v\|}$. For our choice of Φ, we have $\gamma_\times(R) = \sqrt{n}$ (see [9, p. 174]).

Since Φ is the $2n$-th cyclotomic polynomial, the ring R is exactly the maximal order (i.e., the ring of integers) of the cyclotomic field $\mathbb{Q}[\zeta] \cong \mathbb{Q}[x]/\Phi =: K$, where $\zeta \in \mathbb{C}$ is a primitive $2n$-th root of unity. We denote by $(\sigma_i)_{i\leq n}$ the canonical complex embeddings: We can choose $\sigma_i : P \mapsto P(\zeta^{2i+1})$ for $i \leq n$. For any α in $\mathbb{Q}[\zeta]$, we define its T_2-norm by $T_2(\alpha)^2 = \sum_{i\leq n} |\sigma_i(\alpha)|^2$ and its algebraic norm by $\mathcal{N}(\alpha) = \prod_{i\leq n} |\sigma_i(\alpha)|$. The arithmetic-geometric inequality gives $\mathcal{N}(\alpha)^{2/n} \leq \frac{1}{n} T_2(\alpha)^2$. Also, for the particular cyclotomic fields we are considering, the polynomial norm (the norm of the coefficient vector of α when expressed as an element of K) satisfies $\|\alpha\| = \frac{1}{\sqrt{n}} T_2(\alpha)$. We also use the fact that for any $\alpha \in R$, we have $|\mathcal{N}(\alpha)| = \det\langle\alpha\rangle$, where $\langle\alpha\rangle$ is the ideal of R generated by α. For simplicity, we will try to use the polynomial terminology wherever possible.

Let q be a prime number such that Φ has n distinct linear factors modulo q (i.e., $q = 1 \bmod 2n$): $\Phi = \prod_{i\leq n} \Phi_i = \prod_{i\leq n}(x - \phi_i) \bmod q$. Let $R_q = R/qR = \mathbb{Z}_q[x]/\Phi$. Dirichlet's theorem on arithmetic progressions implies that infinitely such primes exist. Furthermore, Linnik's theorem asserts that the smallest such q is $\mathcal{P}oly(n)$, and much effort has been spent to decrease this bound (the current record seems to be $O(n^{5.2})$, see [36]). Furthermore, we can write ϕ_i as r^i, where r is a primitive $(2n)$-th root of unity modulo q. This implies that the Chinese Remainder Theorem in R_q provides a natural fast Discrete Fourier Transform, and thus multiplication of elements of R_q can be performed within $O(n \log n)$ additions and multiplications modulo q (see [7, Ch. 8], [20, Se. 2.1]).

The R-LWE problem

For $s \in R_q$ and ψ a distribution in R_q, we define $A_{s,\psi}$ as the distribution obtained by sampling the pair $(a, as{+}e)$ with $(a, e) \hookleftarrow U(R_q)\times\psi$. The Ring Learning With Errors problem (R-LWE) was introduced by Lyubashevsky *et al.* [21] and shown hard for specific error distributions ψ. These are slightly technical to define (see below), but for the present work, the important facts to be remembered are that the samples are small (see Lemma 6), and can be obtained in quasi-linear time.

The error distributions ψ that we use are an adaptation of those introduced in [21]. They are sampled from a family of distributions $\overline{\Upsilon}_\alpha$ that we now define. For $\boldsymbol{\sigma} \in \mathbb{R}^n$ with positive coordinates, we define the ellipsoidal Gaussian $\rho_{\boldsymbol{\sigma}}$ as the row vector of independent Gaussians $(\rho_{\sigma_1}, \ldots, \rho_{\sigma_n})$, where $\sigma_i = \sigma_{i+n/2}$ for $1 \leq i \leq n/2$. As we want to define R-LWE in the polynomial expression of R rather than with the so-called "space H" of [21], we apply a matrix transformation to the latter Gaussians. We define a sample from $\rho'_{\boldsymbol{\sigma}}$ as a sample from $\rho_{\boldsymbol{\sigma}}$, multiplied first (from the right) by $\frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ i & -i\end{pmatrix}\otimes\mathrm{Id}_{n/2}\in\mathbb{C}^{n\times n}$, and second by $V=\frac{1}{n}\left(\zeta^{-(2j+1)k}\right)_{0\leq j,k<n}$. Note that vector multiplication by matrix V corresponds to a complex discrete Fourier transform, and can be performed in $O(n\log n)$ complex-valued arithmetic operations with the Cooley-Tukey FFT. Moreover, it is numerically extremely stable: if all operations are performed with a precision of $p=\Omega(\log n)$ bits, then the computed output vector $fl(\boldsymbol{y})$ satisfies $\|fl(\boldsymbol{y})-\boldsymbol{y}\|\leq C\cdot(\log n)\cdot 2^{-p}\cdot\|\boldsymbol{y}\|$, where C is some absolute constant and $\boldsymbol{y}$ is the vector that would be obtained with exact computations. We refer to [12, Se. 24.1] for details. We now define a sample from $\overline{\rho'}_{\boldsymbol{\sigma}}$ as follows: compute a sample from $\rho'_{\boldsymbol{\sigma}}$ with absolute error $<1/n^2$; if it is within distance $1/n^2$ of the middle of two consecutive integers, then restart; otherwise, round it to a closest integer and then reduce it modulo q. Finally, a distribution sampled from $\overline{\Upsilon}_\alpha$ for $\alpha\geq 0$ is defined as $\overline{\rho'}_{\boldsymbol{\sigma}}$, where $\sigma_i=\sqrt{\alpha^2q^2+x_i}$ with the x_i's sampled independently from the distribution $\mathrm{Exp}(n\alpha^2q^2)$.

Sampling from $\rho'_{\boldsymbol{\sigma}}$ can be performed in time $\widetilde{O}(n)$. Sampling from $\overline{\Upsilon}_\alpha$ can also be performed in expected time $\widetilde{O}(n)$, and the running-time is bounded by a quantity that follows a geometric law of parameter <1. Furthermore, in all our cryptographic applications, one could pre-compute such samples off-line (i.e., before the message M to be processed is known).

Lemma 6. *Assume that $\alpha q\geq\sqrt{n}$. For any $r\in R$, we have $\Pr_{y\hookleftarrow\overline{\Upsilon}_\alpha}[\|yr\|_\infty\geq\alpha q\omega(\log n)\cdot\|r\|]\leq n^{-\omega(1)}$.*

Proof. We define Υ_α exactly as $\overline{\Upsilon}_\alpha$, but without the rejection step from $\rho'_{\boldsymbol{\sigma}}$ to $\overline{\rho'}_\sigma$. Because of the bound on the rejection probability, it suffices to prove the result with Υ_α instead of $\overline{\Upsilon}_\alpha$. Let $y\hookleftarrow\Upsilon_\alpha$. The involved $\boldsymbol{\sigma}$ satisfies $\sigma_k=\sqrt{\alpha^2q^2+x_k}$, with the x_k's sampled independently from the distribution $\mathrm{Exp}(n\alpha^2q^2)$. We have $\max\sigma_k\leq\alpha q\sqrt{n}\omega(\sqrt{\log n})$ with probability $1-n^{-\omega(1)}$. We write $y=y'+\eta$, where the field element $y'\in K$ is sampled from $\rho'_{\boldsymbol{\sigma}}$, and actually derived from a sample $\boldsymbol{z}$ from $\rho_{\boldsymbol{\sigma}}$, and $\eta\in K$ is the error in rounding $y'\in K$ to $y\in R$, with $\|\eta\|_\infty\leq 1/2$. Then $\|yr\|_\infty\leq\|y'r\|_\infty+\|\eta r\|_\infty$. Using the Schwartz inequality, the second term can be bounded as $\|\eta r\|_\infty\leq\frac{\sqrt{n}}{2}\|r\|$. We now bound the first term. The embedding vector of y' has the following shape:

$$\frac{1}{\sqrt{2}}(z_1+iz_{n/2+1},\ldots,z_{n/2}+iz_n,z_1-iz_{n/2+1},\ldots,z_{n/2}-iz_n).$$

Let $(r^{(k)})_k$ be the embedding vector of r. Then the embedding vector of $y'r$ is $(y'^{(k)}r^{(k)})_k$. The coefficient in x^j of $y'r$ is

$$\begin{aligned}\frac{1}{n}\sum_{0\leq k<n}\zeta^{-(2j+1)k}y^{(k)}r^{(k)}&=\frac{2}{n}\Re\left(\sum_{0\leq k<n/2}\zeta^{-(2j+1)k}y^{(k)}r^{(k)}\right)\\&=\frac{\sqrt{2}}{n}\sum_{0\leq k<n/2}\Re\left((\zeta^{-(2j+1)k}r^{(k)})(z_{k+1}+iz_{n/2+k+1})\right).\end{aligned}$$

The kth summand of the last sum follows a normal law of mean 0 and standard deviation $|r^{(k)}|\sigma_k$. Therefore, the coefficient in x^j of yr follows a normal law of standard deviation $\leq \frac{1}{n}T_2(r)\max\sigma_k$, which is $\leq \frac{1}{\sqrt{n}}\alpha q\omega(\sqrt{\log n})\cdot T_2(r) = \alpha q\omega(\log n)\cdot\|r\|$ with probability $1-n^{-\omega(1)}$. Using $\alpha q \geq \sqrt{n}$, we get $\|yr\|_\infty + \|\eta r\|_\infty \leq \alpha q\omega(\log n)\cdot\|r\|$ with probability $1-n^{-\omega(1)}$, as claimed. □

We now define our adaptation of R-LWE.

Definition 1. *The Ring Learning With Errors Problem with parameters q,α and Φ (R-LWE$^{\Phi}_{q,\alpha}$) is as follows. Let $\psi \hookleftarrow \overline{\Upsilon}_\alpha$ and $s \hookleftarrow U(R_q)$. Given access to an oracle $\mathcal{O}$ that produces samples in $R_q \times R_q$, distinguish whether $\mathcal{O}$ outputs samples from $A_{s,\psi}$ or from $U(R_q \times R_q)$. The distinguishing advantage should be $1/\mathcal{P}oly(n)$ (resp. $2^{-o(n)}$) over the randomness of the input, the randomness of the samples and the internal randomness of the algorithm.*

The following theorem indicates that R-LWE is hard, assuming that the worst-case γ-Ideal-SVP cannot be efficiently solved using quantum computers, for small γ. It was recently improved by Lyubashevsky *et al.* [22]: if the number of samples that can be asked to the oracle $\mathcal{O}$ is bounded by a constant (which is the case in our application), then the result also holds with simpler errors than $e \hookleftarrow \psi \hookleftarrow \overline{\Upsilon}_\alpha$, and with an even smaller Ideal-SVP approximation factor γ. This should allow to both simplify the modified `NTRUEncrypt` and to strengthen its security guarantee.

Theorem 1 (Adapted from [21]). *Assume that $\alpha q = \omega(n\sqrt{\log n})$ (resp. $\Omega(n^{1.5})$) with $\alpha \in (0,1)$ and $q = \mathcal{P}oly(n)$. There exists a randomized polynomial-time (resp. subexponential) quantum reduction from γ-*Ideal-SVP *to* R-LWE$_{q,\alpha}$*, with $\gamma = \omega(n^{1.5}\log n)/\alpha$ (resp. $\Omega(n^{2.5})/\alpha$).*

The differences with [21] in the above formulation are the use of the polynomial representation (handled by applying the complex FFT to the error term), the use of R_q rather than $R_q^\vee := R^\vee/q$ where $R^\vee$ is the codifferent (here we have $R_q^\vee = \frac{1}{n}R_q$), and the truncation of the error to closest integer if it is far from the middle of two consecutive integers. The new variant remains hard because a sample passes the rejection step with non-negligible probability, and the rounding can be performed on the oracle samples obliviously to the actual error.

Variants of R-LWE. For $s \in R_q$ and ψ a distribution in R_q, we define $A^\times_{s,\psi}$ as the distribution obtained by sampling the pair $(a, as+e)$ with $(a,e) \hookleftarrow U(R_q^\times)\times\psi$, where $R_q^\times$ is the set of invertible elements of R_q. When $q = \Omega(n)$, the probability for a uniform element of R_q of being invertible is non-negligible, and thus R-LWE remains hard even when $A_{s,\psi}$ and $U(R_q \times R_q)$ are respectively replaced by $A^\times_{s,\psi}$ and $U(R_q^\times \times R_q)$. We call R-LWE$^\times$ the latter variant.

Furthermore, similarly to [3, Le. 2] and as explained in [22], the nonce s can also be chosen from the error distribution without incurring any security loss. We call R-LWE$^\times_{\mathrm{HNF}}$ the corresponding modification of R-LWE. We recall the argument, for completeness. Assume an algorithm $\mathcal{A}$ can solve R-LWE$^\times_{\mathrm{HNF}}$. We use $\mathcal{A}$ to solve R-LWE$^\times$. The principle is to transform samples $((a_i,b_i))_i$

into samples $((a_1^{-1}a_i, b_i - a_1^{-1}b_1a_i))_i$, where inversion is performed in $R_q^\times$. This transformation maps $A^\times_{s,\psi}$ to $A^\times_{-e_1,\psi}$, and $U(R_q^\times \times R_q)$ to itself.

3 New Results on Module q-ary Lattices

In this section, we present strong regularity bounds for the ring R_q. For this purpose, we first study two families of R-modules.

3.1 Duality results for some module lattices

Let $\boldsymbol{a} \in R_q^m$. We define the following families of R-modules, for I an arbitrary ideal of R_q:

$$\boldsymbol{a}^\perp(I) := \{(t_1,\dots,t_m) \in R^m : \forall i, (t_i \bmod q) \in I \ \text{ and } \sum_i t_i a_i = 0 \bmod q\},$$

$$L(\boldsymbol{a}, I) := \{(t_1,\dots,t_m) \in R^m : \exists s \in R_q, \forall i, (t_i \bmod q) = a_i \cdot s \bmod I\}.$$

We also define $\boldsymbol{a}^\perp$ and $L(\boldsymbol{a})$ as $\boldsymbol{a}^\perp(R_q)$ and $L(\boldsymbol{a}, \langle 0\rangle)$ respectively. The ideals of R_q are of the form $I_S := \prod_{i\in S}(x - \phi_i)\cdot R_q = \{a \in R_q : \forall i \in S, a(\phi_i) = 0\}$, where S is any subset of $\{1,\dots,n\}$ (the ϕ_i's are the roots of Φ modulo q). We define $I_S^\times = \prod_{i\in S}(x - \phi_i^{-1})\cdot R_q$.

Lemma 7. *Let $S \subseteq \{1,\dots,n\}$ and $\boldsymbol{a} \in R_q^m$. Let $\overline{S} = \{1,\dots,n\} \setminus S$ and $\boldsymbol{a}^\times \in R_q^m$ be defined by $a_i^\times = a_i(x^{-1})$. Then (considering both sets as mn-dimensional lattices by identifying R and $\mathbb{Z}^n$):*

$$\widehat{\boldsymbol{a}^\perp(I_S)} = \frac{1}{q} L(\boldsymbol{a}^\times, I^\times_{\overline{S}}).$$

Proof. We first prove that $\frac{1}{q}L(\boldsymbol{a}^*, I^\times_{\overline{S}}) \subseteq \widehat{\boldsymbol{a}^\perp(I_S)}$. Let $(t_1,\dots,t_m) \in \boldsymbol{a}^\perp(I_S)$ and $(t'_1,\dots,t'_m) \in L(\boldsymbol{a}^*, I_{\overline{S}})$. Write $t_i = \sum_{j<n} t_{i,j}x^j$ and $t'_i = \sum_{j<n} t'_{i,j}x^j$ for any $i \le m$. Our goal is to show that $\sum_{i\le m, j\le n} t_{i,j}t'_{i,j} = 0 \bmod q$. This is equivalent to showing that the constant coefficient of the polynomial $\sum_{i\le m} t_i(x)t'_i(x^{-1})$ is 0 modulo q. It thus suffices to show that $\sum_{i\le m} t_i(x)t'_i(x^{-1}) \bmod q = 0$ (in R_q). By definition of the t'_i's, there exists $s \in R_q$ such that $(t'_i \bmod q) = a_i^\times \cdot s + b'_i$ for some $b'_i \in I^\times_{\overline{S}}$. We have the following, modulo q:

$$\sum_{i\le m} t_i(x)t'_i(x^{-1}) = s(x^{-1}) \cdot \sum_{i\le m} t_i(x)a_i(x) + \sum_{i\le m} t_i(x)b'_i(x^{-1}).$$

Both sums in the right hand side evaluate to 0 in R_q, which provides the desired inclusion.

To complete the proof, it suffices to show that $L\widehat{(\boldsymbol{a}^\times, I^\times_{\overline{S}})} \subseteq \frac{1}{q}\boldsymbol{a}^\perp(I_S)$. It can be seen by considering the elements of $L(\boldsymbol{a}^\times, I_{\overline{S}})$ corresponding to $s = 1$. □

3.2 On the absence of unusually short vectors in $L(\boldsymbol{a}, I_S)$

We show that for $\boldsymbol{a} \hookleftarrow U((R_q^\times)^m)$, the lattice $L(\boldsymbol{a}, I_S)$ is extremely unlikely to contain unusually short vectors for the infinity norm, i.e., much shorter than guaranteed by the Minkowski upper bound $\det(L(\boldsymbol{a}, I_S))^{\frac{1}{mn}} = q^{(1-\frac{1}{m})\frac{|S|}{n}}$ (we have $\det(L(\boldsymbol{a}, I_S)) = q^{(m-1)|S|}$ because there are $q^{n+(m-1)(n-|S|)}$ points of $L(\boldsymbol{a}, I_S)$ in the cube $[0, q-1]^{mn}$). Note that our lower bound approaches the Minkowski bound as $\frac{|S|}{n}$ approaches 1, but becomes progressively looser as $\frac{|S|}{n}$ drops towards $\approx 1 - \frac{1}{m}$. Fortunately, for our applications, we will be using this bound with $\frac{|S|}{n} = 1 - \varepsilon$ for some small ε, where the bound is close to being tight.

Lemma 8. *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 5$. Then, for any $S \subseteq \{1, \ldots, n\}$, $m \geq 2$ and $\varepsilon > 0$, we have $\lambda_1^\infty(L(\boldsymbol{a}, I_S)) \geq \frac{1}{\sqrt{n}} q^\beta$, with:*

$$\begin{aligned}\beta &:= 1 - \frac{1}{m} + \frac{1 - \sqrt{1 + 4m(m-1)\left(1 - \frac{|S|}{n}\right) + 4m\varepsilon}}{2m} \\ &\geq 1 - \frac{1}{m} - \varepsilon - (m-1)\left(1 - \frac{|S|}{n}\right),\end{aligned}$$

except with probability $\leq 2^n (q-1)^{-\varepsilon n}$ over the uniformly random choice of $\boldsymbol{a}$ in $(R_q^\times)^m$.

Proof. Recall that $\Phi = \prod_{i \leq n} \Phi_i$ for distinct linear factors Φ_i. By the Chinese Remainder Theorem, we know that R_q (resp. $R_q^\times$) is isomomorphic to $(\mathbb{Z}_q)^n$ (resp. $(\mathbb{Z}_q^\times)^n$) via the isomorphism $t \mapsto (t \bmod \Phi_i)_{i \leq m}$. Let $g_{I_S} = \prod_{i \in S} \Phi_i$: it is a degree $|S|$ generator of I_S.

Let p denote the probability (over the randomness of $\boldsymbol{a}$) that $L(\boldsymbol{a}, I_S)$ contains a non-zero vector $\boldsymbol{t}$ of infinity norm $< B$, where $B = \frac{1}{\sqrt{n}} q^\beta$. We upper bound p by the union bound, summing the probabilities $p(\boldsymbol{t}, s) = \Pr_{\boldsymbol{a}}[\forall i, t_i = a_i s \bmod I_S]$ over all possible values for $\boldsymbol{t}$ of infinity norm $< B$ and $s \in R_q / I_S$. Since the a_i's are independent, we have $p(\boldsymbol{t}, s) = \prod_{i \leq m} p_i(t_i, s)$, where $p_i(t_i, s) = \Pr_{a_i}[t_i = a_i s \bmod I_S]$.

Wlog we can assume that $\gcd(s, g_{I_S}) = \gcd(t_i, g_{I_S})$ (up to multiplication by an element of $\mathbb{Z}_q^\times$): If this is not the case, there exists $j \leq n$ such that either $t_i \bmod \Phi_j = 0$ and $s \bmod \Phi_j \neq 0$, or $t_i \bmod \Phi_j \neq 0$ and $s \bmod \Phi_j = 0$; In both cases, we have $p_i(t_i, s) = 0$ because $a_i \in R_q^\times$. We now assume that $\gcd(s, g_{I_S}) = \gcd(t_i, g_{I_S}) = \prod_{i \in S'} \Phi_i$ for some $S' \subseteq S$ of size $0 \leq d \leq |S|$. For any $j \in S'$, we have $t_i = a_i s = 0 \bmod \Phi_j$ regardless of the value of $a_i \bmod \Phi_j$, while for $j \in S \setminus S'$, we have $s \neq 0 \bmod \Phi_j$ and there exists a unique value of $a_i \bmod \Phi_j$ such that $t_i = a_i s \bmod \Phi_j$. Moreover for any $j \notin S$, the value of $a_i \bmod \Phi_j$ can be arbitrary in $\mathbb{Z}_q^\times$. So, overall, there are $(q-1)^{d+n-|S|}$ differents a_i's in $R_q^\times$ such that $t_i = a_i s \bmod I_S$. This leads to $p_i(t_i, s) = (q-1)^{d-|S|}$.

So far, we have showed that the probability p can be upper bounded by:

$$p \;\leq\; \sum_{0\leq d\leq |S|} \;\; \sum_{\substack{h=\prod_{i\in S'}\varPhi_i \\ S'\subseteq S \\ |S'|=d}} \;\; \sum_{\substack{s\in R_q/I_S \\ h|s}} \;\; \sum_{\substack{\boldsymbol{t}\in (R_q)^m \\ \forall i,\, 0<\|t_i\|_\infty<B \\ \forall i,\, h|t_i}} \;\; \prod_{i\leq m} \; (q-1)^{d-|S|}.$$

For $h = \prod_{i\in S'} \varPhi_i$ of degree d, let $N(B,d)$ denote the number of $t \in R_q$ such that $\|t\|_\infty < B$ and $t = ht'$ for some $t' \in R_q$ of degree $< n-d$. We consider two bounds for $N(B,d)$ depending on d.

Suppose that $d \geq \beta\cdot n$. Then we claim that $N(B,d) = 0$. Indeed, any $t = ht'$ for some $t' \in R_q$ belongs to the ideal $\langle h, q\rangle$ of R generated by h and q. For any non-zero $t \in \langle h,q\rangle$, we have $\mathcal{N}(t) = \mathcal{N}(\langle t\rangle) \geq \mathcal{N}(\langle h,q\rangle) = q^d$, where the inequality is because the ideal $\langle t\rangle$ is a full-rank sub-ideal of $\langle h,q\rangle$, and the last equality is because $\deg h = d$. It follows from the arithmetic-geometric inequality that $\|t\| = \frac{1}{\sqrt{n}} T_2(t) \geq \mathcal{N}(t)^{1/n} \geq q^{d/n}$. By equivalence of norms, we conclude that $\|t\|_\infty \geq \lambda_1^\infty(\langle h,q\rangle) \geq \frac{1}{\sqrt{n}} q^{d/n}$. We see that $d/n \geq \beta$ implies that $\|t\|_\infty \geq B$, so that $N(B,d) = 0$.

Suppose now that $d < \beta\cdot n$. Then we claim that $N(B,d) \leq (2B)^{n-d}$. Indeed, since the degree of h is d, the vector $\overline{t}$ formed by the $n-d$ low-order coefficients of t is related to the vector $\overline{t'}$ formed by the $n-d$ low-order coefficients of t' by a lower triangular $(n-d)\times(n-d)$ matrix whose diagonal coefficients are equal to 1. Hence this matrix is non-singular modulo q so the mapping from $\overline{t'}$ to $\overline{t}$ is one-to-one. This provides the claim.

Using the above bounds on $N(B,d)$, the fact that the number of subsets of S of cardinality d is $\leq 2^d$, and the fact that the number of $s \in R_q/I_S$ divisible by $h = \prod_{i\in S'} \varPhi_i$ is $q^{|S|-d}$, the above bound on p implies

$$p \;\leq\; 2^n \max_{d\leq\beta\cdot n} \frac{(2B)^{m(n-d)}}{(q-1)^{(m-1)(|S|-d)}}.$$

With our choice of B, we have $2B \leq (q-1)^\beta$ (this is implied by $n \geq 8, q \geq 5$ and $\beta \leq 1$). A straightforward computation then leads to the claimed upper bound on p. □

3.3 Improved regularity bounds

We now study the uniformity of distribution of $(m+1)$-tuples from $(R_q^\times)^m \times R_q$ of the form $(a_1, \ldots, a_m, \sum_{i\leq m} t_i a_i)$, where the a_i's are independent and uniformly random in $R_q^\times$, and the t_i's are chosen from some distribution on R_q concentrated on elements with small coefficients. Similarly to [23], we call the distance of the latter distribution to the uniform distribution on $(R_q^\times)^m \times R_q$ the *regularity* of the generalized knapsack function $(t_i)_{i\leq m} \mapsto \sum_{i\leq m} t_i a_i$. For our NTRU application we are particularly interested in the case where $m = 2$.

The regularity result in [23, Se. 4.1] applies when the a_i's are uniformly random in the whole ring R_q, and the t_i's are uniformly random on the subset of elements of R_q with coefficients of magnitude $\leq d$ for some $d < q$. In this case, the regularity bound from [23] is $\Omega(\sqrt{nq/d^m})$. Unfortunately, this bound is non-negligible for small m and q, e.g., for $m = O(1)$ and $q = \mathcal{P}oly(n)$. To make it exponentially small in n, one needs to set $m \log d = \Omega(n)$, which inevitably leads to inefficient cryptographic functions. When the a_i's are chosen uniformly from the whole ring R_q, the actual regularity is not much better than this undesirable regularity bound. This is because R_q contains n proper ideals of size $q^{n-1} = |R_q|/q$, and the probability $\approx n/q^m$ that all of the a_i's fall into one such ideal (which causes $\sum t_i a_i$ to also be trapped in the proper ideal) is non-negligible for small m. To circumvent this problem, we restrict the a_i's to be uniform in $R_q^\times$, and we choose the t_i's from a discrete Gaussian distribution. We show a regularity bound exponentially small in n even for $m = O(1)$, by using an argument similar to that used in [11, Se. 5.1] for unstructured generalized knapsacks, based on the *smoothing parameter* of the underlying lattices. Note that the new regularity result can be used within the Ideal-SIS trapdoor generation of [35, Se. 3], thus extending the latter to a fully splitting q. It also shows that the encryption scheme from [21] can be shown secure against subexponential (quantum) attackers, assuming the subexponential (quantum) hardness of standard worst-case problems over ideal lattices.

Theorem 2. *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 5$. Let $m \geq 2$, $\varepsilon > 0$, $\delta \in (0, 1/2)$ and $\boldsymbol{t} \hookleftarrow D_{\mathbb{Z}^{mn},\sigma}$, with $\sigma \geq \sqrt{n \ln(2mn(1 + 1/\delta))/\pi} \cdot q^{\frac{1}{m}+\varepsilon}$. Then for all except a fraction $\leq 2^n(q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^\times)^m$, we have $\eta_\delta(\boldsymbol{a}^\perp) \leq \sqrt{n \ln(2mn(1 + 1/\delta))/\pi} \cdot q^{\frac{1}{m}+\varepsilon}$, and the distance to uniformity of $\sum_{i \leq m} t_i a_i$ is $\leq 2\delta$. As a consequence:*

$$\Delta\left[\left(a_1, \ldots, a_m, \sum_{i \leq m} t_i a_i\right);\ U\left((R_q^\times)^m \times R_q\right)\right] \leq 2\delta + 2^n(q-1)^{-\varepsilon n}.$$

When using this result, one is typically interested in taking a small constant $\varepsilon > 0$, because it allows to lower the standard deviation σ and thus the required amount of randomness. Then a tiny δ should be chosen (e.g., $\delta \approx 2^n(q-1)^{-\varepsilon n}$), as it drastically lowers the statistical distance upper bound, without strengthening the standard deviation requirement much.

For each $\boldsymbol{a} \in (R_q^\times)^m$, let $D_{\boldsymbol{a}}$ denote the distribution of $\sum_{i \leq m} t_i a_i$ where $\boldsymbol{t}$ is sampled from $D_{\mathbb{Z}^{mn},\sigma}$. Note that the above statistical distance is exactly $\frac{1}{|R_q^\times|^m} \sum_{\boldsymbol{a} \in (R_q^\times)^m} \Delta_{\boldsymbol{a}}$, where $\Delta_{\boldsymbol{a}}$ is the distance to uniformity of $D_{\boldsymbol{a}}$. To prove the theorem, it therefore suffices to show a distance bound $\Delta_{\boldsymbol{a}} \leq 2\delta$, for all except a fraction $\leq 2^n(q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^\times)^m$.

Now, the mapping $\boldsymbol{t} \mapsto \sum_i t_i a_i$ induces an isomorphism from the quotient group $\mathbb{Z}^{mn}/\boldsymbol{a}^\perp$ to its range (note that $\boldsymbol{a}^\perp$ is the kernel of $\boldsymbol{t} \mapsto \sum_i t_i a_i$). The latter is R_q, thanks to the invertibility of the a_i's. Therefore, the statistical distance $\Delta_{\boldsymbol{a}}$ is equal to the distance to uniformity of $\boldsymbol{t}$ mod $\boldsymbol{a}^\perp$. By Lemma 4, we have $\Delta_{\boldsymbol{a}} \leq 2\delta$ if σ is greater than the smoothing parameter $\eta_\delta(\boldsymbol{a}^\perp)$ of $\boldsymbol{a}^\perp \subseteq \mathbb{Z}^{mn}$. To upper bound $\eta_\delta(\boldsymbol{a}^\perp)$, we apply Lemma 2, which reduces the task to lower bounding the minimum of the dual lattice $\widehat{\boldsymbol{a}^{\perp}} = \frac{1}{q} \cdot L(\boldsymbol{a}^{\times})$, where $\boldsymbol{a}^{\times} \in (R_q^{\times})^m$ is in one-to-one correspondence with $\boldsymbol{a}$.

The following result is a direct consequence of Lemmata 2, 4, 7 and 8. Theorem 2 follows by taking $S = \emptyset$ and $\boldsymbol{c} = \boldsymbol{0}$.

Lemma 9. *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 5$. Let $S \subseteq \{1, \ldots, n\}$, $m \geq 2$, $\varepsilon > 0$, $\delta \in (0, 1/2)$, $\boldsymbol{c} \in \mathbb{R}^{mn}$ and $\boldsymbol{t} \hookleftarrow D_{\mathbb{Z}^{mn},\sigma,\boldsymbol{c}}$, with*

$$\sigma \geq \sqrt{n \ln(2mn(1 + 1/\delta))/\pi} \cdot q^{\frac{1}{m} + (m-1)\frac{|S|}{n} + \varepsilon}.$$

Then for all except a fraction $\leq 2^n (q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^{\times})^m$, we have:

$$\Delta\Big[\boldsymbol{t} \bmod \boldsymbol{a}^{\perp}(I_S);\ U(R/\boldsymbol{a}^{\perp}(I_S))\Big] \leq 2\delta.$$

4 A revised key generation algorithm

We now use the results of the previous section on modular q-ary lattices to derive a key generation algorithm for `NTRUEncrypt`, where the generated public key follows a distribution for which Ideal-SVP reduces to R-LWE.

4.1 `NTRUEncrypt`'s key generation algorithm

The new key generation algorithm for `NTRUEncrypt` is given in Fig. 1. The secret key polynomials f and g are generated by using the Gentry *et al.* sampler of discrete Gaussians (see Lemma 5), and by rejecting so that the output polynomials are invertible modulo q. The Gentry *et al.* sampler may not exactly sample from discrete Gaussians, but since the statistical distance can be made exponentially small, the impact on our results is also exponentially small. Furthermore, it can be checked that our conditions on standard deviations are much stronger than the one in Lemma 5. From now on, we will assume we have a perfect discrete Gaussian sampler.

By choosing a large enough standard deviation σ, we can apply the results of the previous section and obtain the (quasi-)uniformity of the public key. We sample f of the form $p \cdot f' + 1$ so that it has inverse 1 modulo p, making the decryption process of `NTRUEncrypt` more efficient (as in the original `NTRUEncrypt` scheme). We remark that the rejection condition on f at Step 1 is equivalent to the condition $(f' \bmod q) \notin R_q^{\times} - p^{-1}$, where p^{-1} is the inverse of p in $R_q^{\times}$.

The following result ensures that for some appropriate choice of parameters, the key generation algorithm terminates in expected polynomial time.

Lemma 10. *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 5$. Let $\sigma \geq \sqrt{n \ln(2n(1 + 1/\delta))/\pi} \cdot q^{1/n}$, for an arbitrary $\delta \in (0, 1/2)$. Let $a \in R$ and $p \in R_q^{\times}$. Then $\Pr_{f' \hookleftarrow D_{\mathbb{Z}^n,\sigma}}[(p \cdot f' + a \bmod q) \notin R_q^{\times}] \leq n(1/q + 2\delta)$.*

```
Inputs: $n, q \in \mathbb{Z}$, $p \in R_q^\times$, $\sigma \in \mathbb{R}$.
Output: A key pair $(sk, pk) \in R \times R_q^\times$.
 1. Sample $f'$ from $D_{\mathbb{Z}^n,\sigma}$; let $f = p \cdot f' + 1$; if $(f \bmod q) \notin R_q^\times$, resample.
 2. Sample $g$ from $D_{\mathbb{Z}^n,\sigma}$; if $(g \bmod q) \notin R_q^\times$, resample.
 3. Return secret key $sk = f$ and public key $pk = h = pg/f \in R_q^\times$.
```

Fig. 1. Revised Key Generation Algorithm for `NTRUEncrypt`.

Proof. We are to bound the probability that $p \cdot f' + a$ belongs to $I := \langle q, \Phi_k \rangle$ by $1/q + 2\delta$, for any $k \le n$. The result then follows from the Chinese Remainder Theorem and the union bound. We have $\mathcal{N}(I) = q$, so that $\lambda_1(I) \le \sqrt{n} q^{1/n}$, by Minkowski's theorem. Since I is an ideal of R, we have $\lambda_n(I) = \lambda_1(I)$, and Lemma 1 gives that $\sigma \ge \eta_\delta(I)$. Lemma 4 then shows that $f \bmod I$ is within distance $\le 2\delta$ to uniformity on R/I, so we have $p \cdot f' + a = 0 \bmod I$ (or, equivalently, $f' = -a/p \bmod I$) with probability $\le 1/q + 2\delta$, as required. □

As a consequence of the above bound on the rejection probability, we have the following result, which ensures that the generated secret key is small.

Lemma 11. *Let $n \ge 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \ge 8n$. Let $\sigma \ge \sqrt{2n \ln(6n)/\pi} \cdot q^{1/n}$. The secret key polynomials f, g returned by the algorithm of Fig. 1 satisfy, with probability $\ge 1 - 2^{-n+3}$:*

$$\|f\| \le 2n\|p\|\sigma \quad \text{and} \quad \|g\| \le \sqrt{n}\sigma.$$

If $\deg p \le 1$, then $\|f\| \le 4\sqrt{n}\|p\|\sigma$ with probability $\ge 1 - 2^{-n+3}$.

Proof. The probability under scope is lower than the probability of the same event without rejection, divided by the rejection probability. The result follows by combining Lemmata 3 and 10. □

4.2 Public key uniformity

In the algorithm of Fig. 1, the polynomials f' and g are independently sampled from the discrete Gaussian distribution $D_{\mathbb{Z}^n,\sigma}$ with $\sigma \ge \mathcal{P}oly(n) \cdot q^{1/2+\varepsilon}$ for an arbitrary $\varepsilon > 0$, but restricted (by rejection) to $R_q^\times - p^{-1}$ and $R_q^\times$, respectively. We denote by $D^\times_{\sigma,z}$ the discrete Gaussian $D_{\mathbb{Z}^n,\sigma}$ restricted to $R_q^\times + z$.

Here we apply the result of Section 3 to show that the statistical closeness to uniformity of a quotient of two distributions $(z + p \cdot D^\times_{\sigma,y})$ for $z \in R_q$ and $y = -zp^{-1} \bmod q$. This includes the case of the public key $h = pg/f \bmod q$ computed by the algorithm of Fig. 1.

Theorem 3. *Let $n \ge 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \ge 5$. Let $\varepsilon > 0$ and $\sigma \ge 2n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+2\varepsilon}$. Let $p \in R_q^\times$, $y_i \in R_q$ and $z_i = -y_i p^{-1} \bmod q$ for $i \in \{1, 2\}$. Then*

$$\Delta\left[\frac{y_1 + p \cdot D^\times_{\sigma,z_1}}{y_2 + p \cdot D^\times_{\sigma,z_2}} \bmod q \;;\; U\left(R_q^\times\right)\right] \;\le\; 2^{3n} q^{-\lfloor \varepsilon n \rfloor}.$$

Proof. For $a \in R_q^\times$, we define $Pr_a = \Pr_{f_1,f_2}[(y_1+pf_1)/(y_2+pf_2) = a]$, where $f_i \hookleftarrow D_{\sigma,z_i}^\times$ for $i \in \{1,2\}$. We are to show that $|Pr_a - (q-1)^{-n}| \leq 2^{2n+5}q^{-\lfloor \varepsilon n \rfloor} \cdot (q-1)^{-n} =: \varepsilon'$ for all except a fraction $\leq 2^{2n}(q-1)^{-\varepsilon n}$ of $a \in R_q^\times$. This directly gives the claimed bound. The fraction of $a \in R_q^\times$ such that $|Pr_a - (q-1)^{-n}| \leq \varepsilon'$ is equal to the fraction of $\boldsymbol{a} = (a_1, a_2) \in (R_q^\times)^2$ such that $|Pr_{\boldsymbol{a}} - (q-1)^{-n}| \leq \varepsilon'$, where $Pr_{\boldsymbol{a}} = \Pr_{f_1,f_2}[a_1 f_1 + a_2 f_2 = a_1 z_1 + a_2 z_2]$. This is because $a_1 f_1 + a_2 f_2 = a_1 z_1 + a_2 z_2$ is equivalent to $(y_1+pf_1)/(y_2+pf_2) = -a_2/a_1$ (in $R_q^\times$), and $-a_2/a_1$ is uniformly random in $R_q^\times$ when $\boldsymbol{a} \hookleftarrow U((R_q^\times)^2)$.

We observe that $(f_1, f_2) = (z_1, z_2) =: \boldsymbol{z}$ satisfies $a_1 f_1 + a_2 f_2 = a_1 z_1 + a_2 z_2$, and hence the set of solutions $(f_1, f_2) \in R$ to the latter equation is $\boldsymbol{z} + \boldsymbol{a}^{\perp\times}$, where $\boldsymbol{a}^{\perp\times} = \boldsymbol{a}^\perp \cap (R_q^\times + q\mathbb{Z}^n)^2$. Therefore:

$$Pr_{\boldsymbol{a}} = \frac{D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp\times})}{D_{\mathbb{Z}^n,\sigma}(z_1 + R_q^\times + q\mathbb{Z}^n) \cdot D_{\mathbb{Z}^n,\sigma}(z_2 + R_q^\times + q\mathbb{Z}^n)}.$$

We now use the fact that for any $\boldsymbol{t} \in \boldsymbol{a}^\perp$ we have $t_2 = -t_1 a_1/a_2$ so, since $-a_1/a_2 \in R_q^\times$, the ring elements t_1 and t_2 must belong to the *same* ideal I_S of R_q for some $S \subseteq \{1, \ldots, n\}$. It follows that $\boldsymbol{a}^{\perp\times} = \boldsymbol{a}^\perp \setminus \bigcup_{S \subseteq \{1,\ldots,n\}, S \neq \emptyset} \boldsymbol{a}^\perp(I_S)$. Similarly, we have $R_q^\times + q\mathbb{Z}^n = \mathbb{Z}^n \setminus \bigcup_{S \subseteq \{1,\ldots,n\}, S \neq \emptyset}(I_S + q\mathbb{Z}^n)$. Using the inclusion-exclusion principle, we obtain:

$$D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp\times}) = \sum_{S \subseteq \{1,\ldots,n\}} (-1)^{|S|} \cdot D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^\perp(I_S)), \tag{1}$$

$$\forall i \in \{1,2\} : D_{\mathbb{Z}^n,\sigma}(z_i + R_q^\times + q\mathbb{Z}^n) = \sum_{S \subseteq \{1,\ldots,n\}} (-1)^{|S|} \cdot D_{\mathbb{Z}^n,\sigma}(z_i + I_S + q\mathbb{Z}^n). \tag{2}$$

In the rest of the proof, we show that, except for a fraction $\leq 2^{2n}(q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^\times)^2$:

$$D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp\times}) = (1 + \delta_0) \cdot \frac{(q-1)^n}{q^{2n}},$$

$$\forall i \in \{1,2\} : \; D_{\mathbb{Z}^n,\sigma}(z_i + R_q^\times + q\mathbb{Z}^n) = (1 + \delta_i) \cdot \frac{(q-1)^n}{q^n}.$$

where $|\delta_i| \leq 2^{2n+2}q^{-\lfloor \varepsilon n \rfloor}$ for $i \in \{0,1,2\}$. The bound on $|Pr_{\boldsymbol{a}} - (q-1)^{-n}|$ follows by a routine computation.

Handling (1). We note that, since $\boldsymbol{z} \in \mathbb{Z}^{2n}$, we have (for any $S \subseteq \{1, \ldots, n\}$):

$$D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^\perp(I_S)) = \frac{\rho_\sigma(\boldsymbol{z} + \boldsymbol{a}^\perp(I_S))}{\rho_\sigma(\mathbb{Z}^{2n})} = \frac{\rho_\sigma(\boldsymbol{z} + \boldsymbol{a}^\perp(I_S))}{\rho_\sigma(\boldsymbol{z} + \mathbb{Z}^{2n})} = D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^\perp(I_S)).$$

For the terms of (1) with $|S| \leq \varepsilon n$, we apply Lemma 9 with $m = 2$. Since $|S|/n + \varepsilon \leq 2\varepsilon$, the Lemma 9 assumption on σ holds, with $\delta := q^{-n-\lfloor \varepsilon n \rfloor}$. We have $|R/\boldsymbol{a}^\perp(I_S)| = \det(\boldsymbol{a}^\perp(I_S)) = q^{n+|S|}$: Indeed, since $\boldsymbol{a} \in (R_q^\times)^2$, there are $q^{n-|S|}$ elements of $\boldsymbol{a}^\perp(I_S)$ in $[0, q-1]^{2n}$. We conclude that $|D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^\perp(I_S)) -$

$q^{-n-|S|}| \le 2\delta$, for all except a fraction $\le 2^n(q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^\times)^2$ (possibly corresponding to a distinct subset of $(R_q^\times)^2$ for each possible S).

For a term of (1) with $|S| > \varepsilon n$, we choose $S' \subseteq S$ with $|S'| = \lfloor \varepsilon n \rfloor$. Then we have $\boldsymbol{a}^\perp(I_S) \subseteq \boldsymbol{a}^\perp(I_{S'})$ and hence $D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^\perp(I_S)) \le D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^\perp(I_{S'}))$. By using with S' the above result for small $|S|$, we obtain $D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^\perp(I_S)) \le 2\delta + q^{-n-\lfloor \varepsilon n \rfloor}$.

Overall, we have, except possibly for a fraction $\le 2^{2n}(q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^\times)^2$:

$$\begin{aligned}\left| D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp\times}) - \sum_{k=0}^{n} (-1)^k \binom{n}{k} q^{-n-k} \right| &\le 2^{n+1}\delta + 2 \sum_{k=\lceil \varepsilon n \rceil}^{n} \binom{n}{k} q^{-n-\lfloor \varepsilon n \rfloor} \\ &\le 2^{n+1}(\delta + q^{-n-\lfloor \varepsilon n \rfloor}).\end{aligned}$$

We conclude that $|\delta_0| \le \frac{q^{2n}}{(q-1)^n} 2^{n+1}(\delta + q^{-n-\lfloor \varepsilon n \rfloor}) \le 2^{2n+1}(\delta q^n + q^{-\lfloor \varepsilon n \rfloor})$, as required.

Handling (2). For the bounds on δ_1 and δ_2, we use a similar argument. Let $i \in \{1, 2\}$. The z_i term can be handled like like the $\boldsymbol{z}$ term of (1). We observe that for any $S \subseteq \{1, \ldots, n\}$, we have $\det(I_S + q\mathbb{Z}^n) = q^{|S|}$ and hence, by Minkowski's theorem, $\lambda_1(I_S + q\mathbb{Z}^n) \le \sqrt{n} \cdot q^{|S|/n}$. Moreover, since $I_S + q\mathbb{Z}^n$ is an ideal lattice, we have $\lambda_n(I_S + q\mathbb{Z}^n) = \lambda_1(I_S + q\mathbb{Z}^n) \le \sqrt{n} \cdot q^{|S|/n}$. Lemma 1 gives that $\sigma \ge \eta_\delta(I_S + q\mathbb{Z}^n)$ for any S such that $|S| \le n/2$, with $\delta := q^{-n/2}$. Therefore, by Lemma 4, for such an S, we have $|D_{\mathbb{Z}^n,\sigma,-z_i}(I_S + q\mathbb{Z}^n) - q^{-|S|}| \le 2\delta$.

For a term of (2) with $|S| > n/2$, we choose $S' \subseteq S$ with $|S'| = n/2$. By using with S' the above result for small $|S|$, we obtain $D_{\mathbb{Z}^n,\sigma,-z_i}(I_S + q\mathbb{Z}^n) \le D_{\mathbb{Z}^n,\sigma,-z_i}(I_{S'} + q\mathbb{Z}^n) \le 2\delta + q^{-n/2}$.

Overall, we have:

$$\begin{aligned}\left| D_{\mathbb{Z}^n,\sigma}(z_i + R_q^\times + q\mathbb{Z}^n) - \sum_{k=0}^{n} (-1)^k \binom{n}{k} q^{-k} \right| &\le 2^{n+1}\delta + 2 \sum_{k=n/2}^{n} \binom{n}{k} q^{-n/2} \\ &\le 2^{n+1}(\delta + q^{-n/2}),\end{aligned}$$

which leads to the desired bound on δ_i (using $\varepsilon < 1/2$). This completes the proof of the theorem. □

5 NTRUEncrypt Revisited

Using our new results above, we describe a modification of NTRUEncrypt for which we can provide a security proof under a worst-case hardness assumption. We use $\Phi = x^n + 1$ with $n \ge 8$ a power of 2, $R = \mathbb{Z}[x]/\Phi$ and $R_q = R/qR$ with $q \ge 5$ prime such that $\Phi = \prod_{k=1}^{n} \Phi_k$ in R_q with distinct Φ_k's.

We define our modified NTRUEncrypt scheme with parameters n, q, p, α, σ as follows. The parameters n and q define the rings R and R_q. The parameter $p \in R_q^\times$ defines the plaintext message space as $\mathcal{P} = R/pR$. It must be a polynomial with 'small' coefficients with respect to q, but at the same time we require $\mathcal{N}(p) =$

$|\mathcal{P}| = 2^{\Omega(n)}$ so that many bits can be encoded at once. Typical choices as used in the original `NTRUEncrypt` scheme are $p = 3$ and $p = x + 2$, but in our case, since q is prime, we may also choose $p = 2$. By reducing modulo the px^i's, we can write any element of p as $\sum_{0 \le i < n} \varepsilon_i x^i p$, with $\varepsilon_i \in (-1/2, 1/2]$. Using the fact that $R = \mathbb{Z}[x]/(x^n + 1)$, we can thus assume that any element of $\mathcal{P}$ is an element of R with infinity norm $\le (\deg(p)+1) \cdot \|p\|$. The parameter α is the R-LWE noise distribution parameter. Finally, the parameter σ is the standard deviation of the discrete Gaussian distribution used in the key generation process (see Section 4).

- **Key generation.** Use the algorithm of Fig. 1 and return $sk = f \in R_q^\times$ with $f = 1 \bmod p$, and $pk = h = pg/f \in R_q^\times$.
- **Encryption.** Given message $M \in \mathcal{P}$, set $s, e \hookleftarrow \overline{\Upsilon}_\alpha$ and return ciphertext $C = hs + pe + M \in R_q$.
- **Decryption.** Given ciphertext C and secret key f, compute $C' = f \cdot C \in R_q$ and return $C' \bmod p$.

Fig. 2. The encryption scheme `NTRUEncrypt`$(n, q, p, \sigma, \alpha)$.

The correctness conditions for the scheme are summarized below.

Lemma 12. *If $\omega(n^{1.5} \log n) \alpha \deg(p) \|p\|^2 \sigma < 1$ (resp. $\omega(n^{0.5} \log n) \alpha \|p\|^2 \sigma < 1$ if $\deg p \le 1$) and $\alpha q \ge n^{0.5}$, then the decryption algorithm of* `NTRUEncrypt` *recovers M with probability $1 - n^{-\omega(1)}$ over the choice of s, e, f, g.*

Proof. In the decryption algorithm, we have $C' = p \cdot (gs + ef) + fM \bmod q$. Let $C'' = p \cdot (gs + ef) + fM$ computed in R (not modulo q). If $\|C''\|_\infty < q/2$ then we have $C' = C''$ in R and hence, since $f = 1 \bmod p$, $C' \bmod p = C'' \bmod p = M \bmod p$, i.e., the decryption algorithm succeeds. It thus suffices to give an upper bound on the probability that $\|C''\|_\infty > q/2$.

From Lemma 11, we know that with probability $\ge 1 - 2^{-n+3}$ both f and g have Euclidean norms $\le 2n\|p\|\sigma$ (resp. $4\sqrt{n}\|p\|\sigma$ if $\deg p \le 1$). This implies that $\|pf\|, \|pg\| \le 2n^{1.5}\|p\|^2\sigma$ (resp. $8\sqrt{n}\|p\|^2\sigma$), with probability $\ge 1 - 2^{-n+3}$. From Lemma 6, both pfe and pgs have infinity norms $\le 2\alpha q n^{1.5} \omega(\log n) \cdot \|p\|^2\sigma$ (resp. $8\alpha q\sqrt{n}\omega(\log n) \cdot \|p\|^2\sigma$), with probability $1 - n^{-\omega(1)}$. Independently:

$$\|fM\|_\infty \le \|fM\| \le \sqrt{n}\|f\|\|M\| \le 2 \cdot (\deg(p) + 1) \cdot n^2 \|p\|^2 \sigma \quad (\text{resp. } 8n\|p\|^2\sigma).$$

Since $\alpha q \ge \sqrt{n}$, we conclude that $\|C''\|_\infty \le (6 + 2\deg(p)) \cdot \alpha q n^{1.5} \omega(\log n) \cdot \|p\|^2 \sigma$ (resp. $24\alpha q n^{0.5} \omega(\log n) \cdot \|p\|^2\sigma$), with probability $1 - n^{-\omega(1)}$. □

The security of the scheme follows by an elementary reduction from the decisional R-LWE$^\times_{\mathrm{HNF}}$, exploiting the uniformity of the public key in $R_q^\times$ (Theorem 3), and the invertibility of p in R_q.

Lemma 13. *Suppose n is a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q = \omega(1)$. Let $\varepsilon, \delta > 0$, $p \in R_q^\times$ and $\sigma \ge 2n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+\varepsilon}$. If there exists an IND-CPA attack against* `NTRUEncrypt` *that runs in* *time T and has success probability $1/2+\delta$, then there exists an algorithm solving* R-LWE$^{\times}_{\mathrm{HNF}}$ *with parameters q and α that runs in time $T' = T + O(n)$ and has success probability $\delta' = \delta - q^{-\Omega(n)}$.*

Proof. Let $\mathcal{A}$ denote the given IND-CPA attack algorithm. We construct an algorithm $\mathcal{B}$ against R-LWE$^{\times}_{\mathrm{HNF}}$ that runs as follows, given oracle $\mathcal{O}$ that samples from either $U(R_q^{\times} \times R_q)$ or $A^{\times}_{s,\psi}$ for some previously chosen $s \hookleftarrow \psi$ and $\psi \hookleftarrow \overline{\Upsilon}_{\alpha}$. Algorithm $\mathcal{B}$ first calls $\mathcal{O}$ to get a sample (h', C') from $R_q^{\times} \times R_q$. Then, algorithm $\mathcal{B}$ runs $\mathcal{A}$ with public key $h = p \cdot h' \in R_q$. When $\mathcal{A}$ outputs challenge messages $M_0, M_1 \in \mathcal{P}$, algorithm $\mathcal{B}$ picks $b \hookleftarrow U(\{0,1\})$, computes the challenge ciphertext $C = p \cdot C' + M_b \in R_q$, and returns C to $\mathcal{A}$. Eventually, when $\mathcal{A}$ outputs its guess b' for b, algorithm $\mathcal{B}$ outputs 1 if $b' = b$ and 0 otherwise.

The h' used by $\mathcal{B}$ is uniformly random in $R_q^{\times}$, and therefore so is the public key h given to $\mathcal{A}$, thanks to the invertibility of p modulo q. Thus, by Theorem 3, the public key given to $\mathcal{A}$ is within statistical distance $q^{-\Omega(n)}$ of the public key distribution in the genuine attack. Moreover, since $C' = h \cdot s + e$ with $s, e \hookleftarrow \psi$, the ciphertext C given to $\mathcal{A}$ has the right distribution as in the IND-CPA attack. Overall, if $\mathcal{O}$ outputs samples from $A^{\times}_{s,\psi}$, then $\mathcal{A}$ succeeds and $\mathcal{B}$ returns 1 with probability $\geq 1/2 + \delta - q^{-\Omega(n)}$.

Now, if $\mathcal{O}$ outputs samples from $U(R_q^{\times} \times R_q)$, then, since $p \in R_q^{\times}$, the value of $p \cdot C'$ and hence C, is uniformly random in R_q and independent of b. It follows that $\mathcal{B}$ outputs 1 with probability 1/2. The claimed advantage of $\mathcal{B}$ follows. □

By combining Lemmata 12 and 13 with Theorem 1 we obtain our main result.

Theorem 4. *Suppose n is a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q = \mathcal{P}oly(n)$ such that $q^{\frac{1}{2}-\varepsilon} = \omega(n^{3.5} \log^2 n \deg(p) \|p\|^2)$ (resp. $q^{\frac{1}{2}-\varepsilon} = \omega(n^4 \log^{1.5} n \deg(p) \|p\|^2)$), for arbitrary $\varepsilon \in (0, 1/2)$ and $p \in R_q^{\times}$. Let $\sigma = 2n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+\varepsilon}$ and $\alpha^{-1} = \omega(n^{1.5} \log n \deg(p) \|p\|^2 \sigma)$. If there exists an IND-CPA attack against* `NTRUEncrypt`$(n, q, p, \sigma, \alpha)$ *which runs in time $T = \mathcal{P}oly(n)$ and has success probability $1/2 + 1/\mathcal{P}oly(n)$ (resp. time $T = 2^{o(n)}$ and success probability $1/2 + 2^{-o(n)}$), then there exists a $\mathcal{P}oly(n)$-time (resp. $2^{o(n)}$-time) quantum algorithm for* γ-Ideal-SVP *with $\gamma = O(n^4 \log^{2.5} n \deg(p) \|p\|^2 q^{\frac{1}{2}+\varepsilon})$ (resp. $\gamma = O(n^5 \log^{1.5} n \deg(p) \|p\|^2 q^{\frac{1}{2}+\varepsilon})$). Moreover, the decryption algorithm succeeds with probability $1 - n^{-\omega(1)}$ over the choice of the encryption randomness.*

In the case where $\deg p \leq 1$, the conditions on q for polynomial-time (resp. subexponential) attacks in Theorem 4 may be relaxed to $q^{\frac{1}{2}-\varepsilon} = \omega(n^{2.5} \log^2 n \cdot \|p\|^2)$ (resp. $q^{\frac{1}{2}-\varepsilon} = \omega(n^3 \log^{1.5} n \cdot \|p\|^2)$) and the resulting Ideal-SVP approximation factor may be improved to $\gamma = O(n^3 \log^{2.5} n \cdot \|p\|^2 q^{\frac{1}{2}+\varepsilon})$ (resp. $\gamma = O(n^4 \log^{1.5} n \cdot \|p\|^2 q^{\frac{1}{2}+\varepsilon})$). Overall, by choosing $\varepsilon = o(1)$, the smallest q for which the analysis holds is $\widetilde{\Omega}(n^5)$ (resp. $\widetilde{\Omega}(n^6)$), and the smallest γ that can be obtained is $\widetilde{O}(n^{5.5})$ (resp. $\widetilde{O}(n^7)$).

Acknowledgements. We thank G. Hanrot, V. Lyubashevsky, Khoa T. T. Nguyen, C. Peikert, O. Regev, I. Shparlinski, J. Silverman and F. Vercauteren for helpful discussions. The authors were partly supported by the LaRedA ANR grant, an Australian Research Fellowship (ARF) from the Australian Research Council under Discovery Grant DP0987734, a Macquarie University Research Fellowship, and ARC Discovery Grant DP110100628.

References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Proc. of Eurocrypt*, volume 6110 of *LNCS*, pages 553–572. Springer, 2010.
2. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proceedings of the 28th Symposium on the Theory of Computing (STOC 1996)*, pages 99–108. ACM, 1996.
3. B. Applebaum, D. Cash, C. Peikert, and A. Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proc. of CRYPTO*, volume 5677 of *LNCS*, pages 595–618. Springer, 2009.
4. W. D. Banks and I. E. Shparlinski. Distribution of inverses in polynomial rings. *Indagationes Mathematicae*, 12(3):303–315, 2001.
5. D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *Proc. of Eurocrypt*, volume 6110 of *LNCS*, pages 523–552. Springer, 2010.
6. D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *Proc. of Eurocrypt*, volume 1233 of *LNCS*, pages 52–61. Springer, 1997.
7. J. von zur Gathen and J. Gerhardt. *Modern Computer Algebra, 2nd edition.* Cambridge University Press, 2003.
8. C. Gentry. *A fully homomorphic encryption scheme.* PhD thesis, Stanford University, 2009. Manuscript available at `http://crypto.stanford.edu/craig`.
9. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proc. of STOC*, pages 169–178. ACM, 2009.
10. C. Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 116–137. Springer, 2010.
11. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
12. N. Higham. *Accuracy and Stability of Numerical Algorithms, 2nd edition.* SIAM, 2002.
13. J. Hoffstein, N. Howgrave-Graham, J. Pipher, and W. Whyte. Practical lattice-based cryptography: `NTRUEncrypt` and `NTRUSign`, 2009. Chapter of [27].
14. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a new high speed public key cryptosystem. Preprint; presented at the rump session of Crypto'96, 1996.
15. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a ring based public key cryptosystem. In *Proc. of ANTS*, volume 1423 of *LNCS*, pages 267–288. Springer, 1998.
16. J. Hoffstein, J. Pipher, and J. H. Silverman. NSS: An NTRU lattice-based signature scheme. In *Proc. of Eurocrypt*, volume 2045 of *LNCS*. Springer, 2001.
17. N. A. Howgrave-Graham, J. H. Silverman, and W. Whyte. Choosing parameter sets for NTRUEncrypt with NAEP and SVES-3. In *Proc. of CT-RSA*, volume 3376 of *LNCS*. Springer, 2005.
18. IEEE P1363. Standard specifications for public-key cryptography. `http://grouper.ieee.org/groups/1363/`.

19. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. ICALP*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
20. V. Lyubashevsky, D. Micciancio, C. Peikert, and A. Rosen. SWIFFT: A modest proposal for FFT hashing. In *Proc. of FSE*, volume 5086 of *LNCS*, pages 54–72. Springer, 2008.
21. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *Proc. of Eurocrypt*, volume 6110 of *LNCS*, pages 1–23. Springer, 2010.
22. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings, 2011. Draft for the extended version of [21], dated 01/02/2011.
23. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007.
24. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007.
25. D. Micciancio and O. Regev. Lattice-based cryptography. In *Post-Quantum Cryptography, D. J. Bernstein, J. Buchmann, E. Dahmen (Eds)*, pages 147–191. Springer, 2009.
26. D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on voronoi cell computations. In *Proc. of STOC*, pages 351–358. ACM, 2010.
27. P. Q. Nguyen and B. Vallée (editors). *The LLL Algorithm: Survey and Applications.* Information Security and Cryptography. Springer, 2009.
28. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Comput. Complexity*, 2(17):300–351, 2008.
29. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
30. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proc. of TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.
31. R. A. Perlner and D. A. Cooper. Quantum resistant public key cryptography: a survey. In *Proc. of IDtrust*, pages 85–93. ACM, 2009.
32. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
33. O. Regev. The learning with errors problem, 2010. Invited survey in CCC 2010, available at `http://www.cs.tau.ac.il/~odedr/`.
34. C. P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theor. Comput. Science*, 53:201–224, 1987.
35. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *Proc. of Asiacrypt*, volume 5912 of *LNCS*, pages 617–635. Springer, 2009.
36. T. Xylouris. On Linnik's constant, 2009. Available at `http://arxiv.org/abs/0906.2749` (in German).