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://www.ics.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 and that of its more recent digital signature counterpart. In the present work, we show how to modify NTRUEncrypt and NTRUSign to make them provably secure in the standard (resp. random oracle) model, under the assumed quantum (resp. classical) 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 of the encryption scheme 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 Ideal-SIS and R-LWE problems.

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

1 Introduction

The NTRU encryption scheme devised by Hoffstein, Pipher and Silverman, was first presented at the rump session of Crypto'96 [16]. 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 [5]. At the ANTS'98 conference, the NTRU authors gave an improved presentation including a thorough assessment of its practical security against lattice attacks [17]. 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 [21]. It is also often considered as the most viable post-quantum public-key encryption (see, e.g., [44]). The authors of NTRUEncrypt also proposed a signature scheme based on a similar design. The history of NTRUSign started with NSS in 2001 [18]. Its development has been significantly more hectic and controversial, with a series of cryptanalyses and repairs (see [9, 11, 19, 51, 33, 36] and the survey [13]).

In parallel to the break-and-repair development of the practically efficient NTRU schemes, 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 natural worst-case

* This is an extended version (Ver. 10, Oct. 2011) of a paper published at Eurocrypt 2011. It contains new results on a provably secure variant of the NTRUSign signature scheme.

problems defined over Euclidean 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 [45, 46]: 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 [46, 39, 10, 4, 1] among others, and the surveys [31, 47]).

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 [28] 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 structuredness 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: The efficiency gain stems from the use of the discrete Fourier transform. Peikert, Rosen, Lyubashevsky and Micciancio [43, 24] 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 the *Ideal Short Integer Solution* problem (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, Lyubashevsky [23] introduced an efficient digital signature provably as secure as Ideal-SIS (in the random oracle model). Also in 2009, Stehlé, Steinfeld, Tanaka and Xagawa [50] introduced a structured (albeit somewhat restricted) 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. The restrictions have recently been waived by Lyubashevsky, Peikert and Regev [27], who introduced a ring variant of LWE, called R-LWE, which allows for more natural cryptographic constructions.

Our Results. The efficiency of cryptography based on Ideal-SIS and R-LWE has been steadily converging towards that of the NTRU primitives. However, the most recent constructions remain computationally more expensive. As an illustration, Lyubashevsky's signature requires the transmission of at least 3 ring elements, and the ElGamal-type encryption scheme derived from [27] (see [41]) requires the transmission of at least 2 ring elements. On the other hand, both `NTRUSign` and `NTRUEncrypt` transmit a single ring element. We close this gap: We prove that (mild) modifications of `NTRUEncrypt` and `NTRUSign` are (CPA-)secure in the standard (resp. random oracle) model, under the assumed quantum (resp. classical) 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 at the end of the introduction. The most substantial additional modification for `NTRUSign` is the use of the fast discrete Gaussian sampler from [40] (which is faster than the one from [10]) in the signing process, which ensures that no secret information is leaked while signing (thus preventing the learning attack from [36]). Our construction also provides a collision-resistant hash very similar to those of [43, 24], which we call `NTRUHash`. We stress that our main goal in this paper is to provide a firm theoretical grounding for the security of the NTRU schemes, 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 technical contribution is the modification and analysis of the key generation algorithms.

In `NTRUEncrypt`, the secret key consists of two sparse polynomials of degrees $< n$ and coefficients in $\{-1, 0, 1\}$. 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 would be desirable to guarantee the latter property, in order to exploit the established hardness of Ideal-SIS and R-LWE (we actually show a weaker distribution property, which still suffices for linking the security to Ideal-SIS and R-LWE). For this purpose, 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 (also known as left-over hash lemma) 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. Micciancio's regularity bound [28, Se. 4.1] (see also [50, Le. 6]) does not suffice for our purposes: For $m = O(1)$, it bounds the distance to uniformity by a constant. To achieve the desired closeness to uniformity, we choose the a_i's uniform among the invertible elements of R_q and we sample the s_i's according to discrete Gaussians with small standard deviations ($\approx q^{1/m}$). A similar regularity bound could be obtained with an FFT-based technique recently developed by Lyubashevsky, Peikert and Regev [26]. 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.

For `NTRUSign`, the technique described in [15, Se. 4] and in [14, Se. 5] to extend an `NTRUEncrypt` secret key into an `NTRUSign` secret key is only heuristic. For instance, it samples an encryption secret key and rejects the sample until some desirable properties are satisfied (e.g., the co-primality of the two secret key polynomials over the rationals), but the security impact of this procedure is not carefully analyzed. We show that in our modified context, the rejection probability can be proven to be sufficiently away from 1, by relating it to the Dedekind zeta function of the cyclotomic fields under scope, and even with this additional rejection, the security of the signature scheme follows from the hardness of Ideal-SIS.

Finally, the cryptographic schemes are obtained from (structured variants of) the Gentry et al [10] signature and dual encryption schemes, via an *inversion-based dimension reduction* of the Ideal-SIS/R-LWE instances. We explain it in the case of Ideal-SIS: Given $(a_i)_{i \leq m}$ uniformly and independently chosen in R_q, find an $\boldsymbol{s} \in R^m \setminus \{\mathbf{0}\}$ with $R := \mathbb{Z}[x]/(x^n + 1)$ such that $\sum s_i a_i = 0 \bmod q$. If q is sufficiently large, the event "a_m invertible in R_q" occurs with non-negligible probability, so the average case hardness of the problem is essentially unchanged if we divide all a_i's by a_m. We can then remove $a_m = 1$ from the input, by making it implicit. This improvement is most dramatic for Ideal-SISwhen $m = 2$.

Brief comparison between 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 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. E.g, if q is chosen as above, one may take $p = 3$ or $p = x + 2$.

The NTRUEncrypt secret key is a pair of polynomials $(f, g) \in R^2_{\text{NTRU}}$ that are sampled randomly with large prescribed proportions of zeros, and with their other coefficients belonging to $\{-1, 1\}$. For improved decryption efficiency, one may choose f such that $f = 1 \bmod p$. 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 and reduce the result modulo q. If the ciphertext was properly generated, this gives $pgs + fM \bmod q$. Since the five involved ring elements have small coefficients, it can be expected that after reduction modulo q the obtained representative is exactly $pgs + fM$ (in R_{NTRU}).
- Reduce the result of the previous step 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 prevent them from occurring, by setting the parameters carefully.

In order to achieve CPA-security under the assumption that standard lattice problems are (quantumly) hard to solve in the worst-case for the family of ideal lattices, we make a few modifications to the original NTRU scheme (which preserve its quasi-linear computation 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 $q \leq \mathcal{P}oly(n)$ as a prime integer such that $f = x^n + 1$ splits into n distinct linear factors modulo q. This allows us to use the search to decision reduction for R-LWE with ring $R_q := R/q$ (see [27]). This also allows us to take $p = 2$.
3. We sample f and g from discrete Gaussians over the set of elements of R, rejecting the samples that are not invertible modulo 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]).

ROAD-MAP. In Section 2, we provide the necessary background material. In Section 3, we prove properties satisfied by some generalized families of random lattices, which eventually lead to the improved regularity bounds for the ring R_q mentioned above. Section 4 is devoted to the key generation algorithms of the modified NTRUEncrypt and NTRUSign schemes. We give the modified NTRU constructions in Section 5. Finally, Section 6 concludes with some open problems.

NOTATION. If q is a non-zero integer, we denote by $\mathbb{Z}_q$ the ring of integers modulo q, i.e., the set $\{0, \ldots, q-1\}$ with the addition and multiplication modulo q. Vectors will be denoted in bold. If $\boldsymbol{x} \in \mathbb{R}^n$, then $\|\boldsymbol{x}\|$ denotes the Euclidean norm of $\boldsymbol{x}$. If $z \in \mathbb{C}$, its real and imaginary parts will be denoted by $\Re(z)$ and $\Im(z)$ respectively. If q is a prime number, we denote by $\mathbb{Z}_q$ the field of integers modulo q. We make use of the Landau notations $O(\cdot), \widetilde{O}(\cdot), o(\cdot), \omega(\cdot), \Omega(\cdot), \widetilde{\Omega}(\cdot)$, $\Theta(\cdot)$. 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 $\text{Exp}(\mu)$ the exponential distribution on $\mathbb{R}$ with mean μ; its corresponding density is $f(x) = \frac{1}{\mu}\exp(-\frac{x}{\mu})$. If E is a finite set, we denote the uniform distribution over E by $U(E)$. If a function f over a countable domain E takes non-negative real values, its sum over an arbitrary $F \subseteq E$ will be denoted by $f(F)$. We say that a sequence of events E_n holds with overwhelming probability if $\Pr[\neg E_n] \leq f(n)$ for a negligible function f. If D_1 and D_2 are two probability distributions over a 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 Some Background Results on the Geometry of Numbers and in Algebraic Number Theory

We refer to [29] for an introduction on the computational aspects of lattices, and to [31] and [47] for detailed surveys on lattice-based cryptography.

2.1 Euclidean lattices

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 non-zero vector of L. 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, we define the k-th *successive minimum* $\lambda_k(L)$ for $k \leq n$ as the smallest r such that L contains at least k linearly independent vectors of norm $\leq r$. The *dual* of L is defined as $\widehat{L} = \{\boldsymbol{c} \in \mathbb{R}^n : \forall i, \langle \boldsymbol{c}, \boldsymbol{b}_i\rangle \in \mathbb{Z}\}$, which is also a lattice: Indeed, if $B = (\boldsymbol{b}_i)_i$ is a basis matrix of L, then B^{-T} is a basis matrix for $\widehat{L}$. This implies that $\widehat{\widehat{L}} = L$.

For a lattice $L \subseteq \mathbb{R}^n$, a real $\sigma > 0$ and a point $\boldsymbol{c} \in \mathbb{R}^n$, 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 $\mathbf{0}$. We extend the definition of $D_{L,\sigma,\boldsymbol{c}}$ to any subset M of 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 $\rho_{1/\sigma}(\widehat{L} \setminus \mathbf{0}) \leq \delta$. We will typically consider $\delta = 2^{-n}$. We will use the following results.

Lemma 2.1 ([30, 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 2.2 ([38, 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 2.3 ([30, Proof of 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 $\rho_{\sigma,\boldsymbol{c}}(L) = \frac{\sigma^n}{\det(L)}(1+\varepsilon)$, with $|\varepsilon| \leq \delta$. As a consequence, we have $\frac{\rho_{\sigma,\boldsymbol{c}}(L)}{\rho_\sigma(L)} \in \left[\frac{1-\delta}{1+\delta}, 1\right]$.*

Lemma 2.4 ([30, 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 2.5 ([10, Cor. 2.8]). *Let $L' \subseteq L \subseteq \mathbb{R}^n$ be two 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 2.6 ([42, Le. 2.11]). *For any full-rank lattice $L \subseteq \mathbb{R}^n$, $\boldsymbol{c} \in \mathbb{R}^n$, $\delta \in (0,1)$, $\sigma \geq 2\eta_\delta(L)$ and $\boldsymbol{b} \in L$, we have $D_{L,\sigma,\boldsymbol{c}}(\boldsymbol{b}) \leq \frac{1+\delta}{1-\delta} \cdot 2^{-n}$.*

Lemma 2.7 ([10, 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 the lattice dimension n.*

We will need the following result on one-dimensional projections of discrete Gaussians. Other results on these projections are known (see [30, Le. 4.2] and [38, Th. 5.2]), but do not seem to suffice for our needs.

Lemma 2.8. *For any full-rank lattice $L \subseteq \mathbb{R}^n$, $\boldsymbol{c} \in \mathbb{R}^n$, $\delta \in (0,1)$, $t \geq \sqrt{2\pi}$, unit vector $\boldsymbol{u} \in \mathbb{R}^n$ and $\sigma \geq \frac{t}{\sqrt{2\pi}} \cdot \eta_\delta(L)$, we have:*

$$\Pr_{\boldsymbol{b} \hookleftarrow D_{L,\sigma,\boldsymbol{c}}}\left[|\langle \boldsymbol{b}-\boldsymbol{c}, \boldsymbol{u}\rangle| \leq \frac{\sigma}{t}\right] \leq \frac{1+\delta}{1-\delta}\frac{\sqrt{2\pi \mathrm{e}}}{t}.$$

Similarly, if $\sigma \geq \eta_\delta(L)$, we have:

$$\Pr_{\boldsymbol{b} \hookleftarrow D_{L,\sigma,\boldsymbol{c}}}\left[|\langle \boldsymbol{b}-\boldsymbol{c}, \boldsymbol{u}\rangle| \geq t\sigma\right] \leq \frac{1+\delta}{1-\delta} t\sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi t^2}.$$

Proof. Let U be an orthonormal matrix whose first row is $\boldsymbol{u}$. We are interested in the random variable X that corresponds to the first component of the vector $\boldsymbol{b}' - \boldsymbol{c}'$ with $\boldsymbol{b}' \hookleftarrow D_{L',\sigma,\boldsymbol{c}'}$, $\boldsymbol{c}' = U\boldsymbol{c}$ and $L' = UL$. We have:

$$\Pr\left[|X| \leq \frac{\sigma}{t}\right] = \frac{(\rho_{\sigma,\boldsymbol{c}'} \cdot \mathbf{1}_{\sigma/t,\boldsymbol{c}'})(L')}{\rho_{\sigma,\boldsymbol{c}'}(L')},$$

where $\mathbf{1}_{\sigma/t,\boldsymbol{c}'}(\boldsymbol{x})$ with $\boldsymbol{x} \in \mathbb{R}^n$ is defined as 1 if $|x_1 - c'_1| \leq \sigma/t$ and 0 otherwise. We first estimate the denominator. We have $\eta_\delta(L') = \eta_\delta(L)$ and $\det(L') = \det(L)$. Therefore, thanks to Lemma 2.3, we have $\rho_{\sigma,\boldsymbol{c}'}(L') = \frac{\sigma^n}{\det(L)}(1+\varepsilon)$ with $|\varepsilon| \leq \delta$.

We now provide an upper bound for the numerator. For any $\boldsymbol{x} \in \mathbb{R}^n$, we have $\mathbf{1}_{\sigma/t,\boldsymbol{c}'}(\boldsymbol{x}) \leq \mathrm{e}^K \cdot \exp\left(-K\frac{|x_1-c'_1|^2}{\sigma^2/t^2}\right)$, where $K = \frac{1}{2} - \frac{\pi}{t^2} \geq 0$. As a consequence:

$$(\rho_{\sigma,\boldsymbol{c}'} \cdot \mathbf{1}_{\sigma/t,\boldsymbol{c}'})(L') \leq \mathrm{e}^K \cdot \rho_{\sigma,D\boldsymbol{c}'}(DL'),$$

where D is the diagonal matrix whose first coefficient is $\sqrt{1+Kt^2/\pi}$ and whose other diagonal coefficients are 1. It can be checked that $\eta_\delta(DL') \leq \sqrt{1+Kt^2/\pi} \cdot \eta_\delta(L')$ and $\det(DL') = \sqrt{1+Kt^2/\pi} \cdot \det(L')$. Using Lemma 2.3 once more provides the result.

The proof of the second statement is similar. We are interested in:

$$\Pr\left[|X| \geq \sigma t\right] = \frac{(\rho_{\sigma,\boldsymbol{c}'} \cdot \bar{\mathbf{1}}_{\sigma t,\boldsymbol{c}'})(L')}{\rho_{\sigma,\boldsymbol{c}'}(L')},$$

where X, L' and $\boldsymbol{c}'$ are defined as above, and $\bar{\mathbf{1}}_{\sigma t,\boldsymbol{c}'}(\boldsymbol{x})$ with $\boldsymbol{x} \in \mathbb{R}^n$ is defined as 1 if $|x_1 - c'_1| > \sigma t$ and 0 otherwise. The denominator is handled as above. For the numerator, note that for any $x \geq \sigma t$, we have $\exp(-\pi \frac{x^2}{\sigma^2}) \leq \sqrt{\mathrm{e}} \cdot \exp(-\pi t^2) \cdot \exp(-\frac{x^2}{2\sigma^2 t^2})$. This gives:

$$(\rho_{\sigma,\boldsymbol{c}'} \cdot \bar{\mathbf{1}}_{\sigma t,\boldsymbol{c}'})(L') \leq \sqrt{\mathrm{e}} \cdot \exp(-\pi t^2) \rho_{\sigma,D\boldsymbol{c}'}(DL'),$$

where D is the diagonal matrix whose first coefficient is $\frac{1}{t\sqrt{2\pi}}$ and whose other diagonal coefficients are 1. It can be checked that $\eta_\delta(DL') \leq \eta_\delta(L')$ and $\det(DL') = \frac{1}{t\sqrt{2\pi}} \cdot \det(L')$. Using Lemma 2.3 once more provides the result. ☐

2.2 Algebraic Number Theory and Lattices

Ideal lattices. Let $\Phi \in \mathbb{Z}[x]$ a monic degree n irreducible polynomial. Let R denote the polynomial ring $\mathbb{Z}[x]/\Phi$. Let I be an (integral) ideal of R, i.e., a subset of R that is closed under addition, and multiplication by arbitrary elements of R. By mapping polynomials to the vectors of their coefficients, we see that a non-zero ideal I 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 $I \subseteq \mathbb{Z}[x]/\Phi$. The *algebraic norm* of a non-zero ideal I is the cardinality of the additive group R/I, and is equal to $\det I$, where I is regarded as an ideal lattice. In the following, an ideal lattice will implicitly refer to a Φ-ideal lattice. For $v \in R$ we denote by $\|v\|$ its Euclidean norm (as a vector). We define the multiplicative *expansion factor* $\gamma_\times(R)$ by $\gamma_\times(R) = \max_{u,v \in R} \frac{\|u \times v\|}{\|u\| \cdot \|v\|}$. A typical choice is $\Phi = x^n + 1$ with n a power of 2, for which $\gamma_\times(R) = \sqrt{n}$ (see [8, p. 174]).

In this work, we will restrict ourselves to $\Phi = x^n + 1$ for n a power of 2. In this setup, any ideal I of R satisfies $\lambda_n(I) = \lambda_1(I)$. Since these Φ's respectively correspond to the $2n$-th cyclotomic polynomial, the ring R is exactly the maximal order (i.e., the ring of integers) of the corresponding cyclotomic number 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 for any element $\alpha \in R$, 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.

The following result is a consequence of Lemma 2.8.

Lemma 2.9. *For any non-zero ideal lattice $I \subseteq R$, $\boldsymbol{c} \in K$, $\delta \in (0, 1)$, $t \geq \sqrt{2\pi}$, $u \in K$ and $\sigma \geq \eta_\delta(I)$, we have*

$$\Pr_{b \hookleftarrow D_{I,\sigma,c}} \left[\|(b - c) \times u\| \geq t\sigma \|u\| \sqrt{n}\right] \leq \frac{1 + \delta}{1 - \delta} tn\sqrt{2\pi \mathrm{e}} \cdot \mathrm{e}^{-\pi t^2}.$$

Proof. A coefficient of $(b-c)\times u \in R$ can be seen as a scalar product between the coefficient vector of $b-c$ and a permutation of the coefficient vector of u. Therefore, by Lemma 2.8, the magnitude of each coefficient of $(b-c)\times u$ is $\geq t\sigma$ with probability $\leq \frac{1+\delta}{1-\delta} t\sqrt{2\pi \mathrm{e}}\cdot \mathrm{e}^{-\pi t^2}$. The union bound implies that all the coefficients magnitudes are $\leq t\sigma$ with probability $\geq 1-\frac{1+\delta}{1-\delta} nt\sqrt{2\pi\mathrm{e}}\cdot \mathrm{e}^{-\pi t^2}$. If that is the case, then $\|(b-c)\times u\| \leq t\sigma\sqrt{t}$, which completes the proof. □

For the analysis of the key generation of the signature scheme (in Subsection 4.3), we need the following result on the inverse (over $K = \mathbb{Q}[x]/(x^n+1)$) of a discrete Gaussian sample. If b is sampled from $D_{I,\sigma}$ for some ideal $I \subseteq R$, we expect $\|b\|$ to be proportional to σ. Since $b \cdot b^{-1} = 1$ over K, it is natural to expect $\|b^{-1}\|$ to be proportional to σ^{-1}.

Lemma 2.10. *Let n a power of 2, $\Phi = x^n+1$ and $R = \mathbb{Z}[x]/\Phi$. For any ideal $I \subseteq R$, $\delta \in (0,1)$, $t \geq \sqrt{2\pi}$ and $\sigma \geq \frac{t}{\sqrt{2\pi}}\cdot \eta_\delta(I)$, we have:*

$$\Pr_{b \hookleftarrow D_{I,\sigma}}\left[\|b^{-1}\| \geq \frac{t}{\sigma\sqrt{n/2}}\right] \leq \frac{1+\delta}{1-\delta}\frac{n\sqrt{2\pi\mathrm{e}}}{t}.$$

Proof. Let $(b^{(i)})_{i\leq n}$ (resp. $(b^{-(i)})_{i\leq n}$) be the complex embeddings of b (resp. b^{-1}). We have $b^{-(i)} = (b^{(i)})^{-1}$, for all i. We first show that it is unlikely that b has a small embedding. Wlog we consider $b^{(1)} = \sum_j b_j \zeta^j$ (where the b_j's are the coefficients of the polynomial b). We let $Re^2 = \sum_j \Re(\zeta^j)^2$ and $Im^2 = \sum_j \Im(\zeta^j)^2$. By applying Lemma 2.8 twice, we obtain:

$$\max\left(\Pr\left[|\Re b^{(1)}| \leq \frac{\sigma Re}{t}\right], \Pr\left[|\Im b^{(1)}| \leq \frac{\sigma Im}{t}\right]\right) \leq \frac{1+\delta}{1-\delta}\frac{\sqrt{2\pi\mathrm{e}}}{t}.$$

We have $Re^2 + Im^2 = n$, which implies that $\max(Re, Im) \geq \sqrt{n/2}$. Therefore:

$$\Pr\left[|b^{(1)}| \leq \frac{\sigma\sqrt{n/2}}{t}\right] \leq \frac{1+\delta}{1-\delta}\frac{\sqrt{2\pi\mathrm{e}}}{t}.$$

Now, the union bound implies that $\Pr[\exists i : |b^{(i)}| \leq \frac{\sigma\sqrt{n/2}}{t}] \leq \frac{1+\delta}{1-\delta}\frac{n\sqrt{2\pi\mathrm{e}}}{t}$. The latter event is exactly the same as $\max_i |b^{-(i)}| \geq \frac{t}{\sigma\sqrt{n/2}}$. Finally, the identity $\|b^{-1}\| \leq \max_i |b^{-(i)}|$ allows us to complete the proof. □

Dedekind Zeta function. We review some facts about the Dedekind zeta function (see, e.g., [35, Ch. VII]), which is used in the analysis of the modified `NTRUSign`. The Möbius function for ring R is a function from the ideals of R to $\{-1,0,1\}$ and is defined as follows: Let $I = \prod_{i=1}^r (J_i)^{e_i}$ denote the unique prime ideal factorization of I in R, where J_i are distinct prime ideals in R and $e_i \in \mathbb{Z}$ for $i \leq r$; Then $\mu(I) = 0$ if there exists i with $e_i \geq 2$, $\mu(I) = (-1)^r$ if $e_i = 1$ for all i and $\mu(R) = 1$. The Dedekind zeta function of the ring R is a function $\zeta_K : \mathbb{R} \to \mathbb{R}$ defined as

$$\zeta_K(s) = \sum_{I\subseteq R} \mathcal{N}(I)^{-s},$$

where the sum is over all ideals of R. The series $\zeta_K(s)$ converges for $s > 1$, and:

$$\zeta_K(s)^{-1} = \prod_{\text{prime } J\subseteq R}(1-\mathcal{N}(J)^{-s}) = \sum_{I\subseteq R}\mu(I)\cdot\mathcal{N}(I)^{-s},$$

where the product is over all *prime* ideals of R and the sum is over all ideals of R.

Lemma 2.11. *Let $K_n = \mathbb{Q}[x]/\Phi_n$, for n a power of 2. Then we have $\zeta_{K_n}(2) = O(1)$, and for $\varepsilon \in (0,1)$, we have $\zeta_{K_n}(1+\varepsilon) \leq 2\exp(2 \cdot (\varepsilon(1-\varepsilon))^{-1} \cdot n^{1-\varepsilon})$.*

Proof. Let $R = \mathbb{Z}[x]/\Phi$. For a (rational) prime p, we let $\pi_K(p)$ denote the number of prime ideals contained in R having norm a power of p, i.e., dividing the principal ideal $\langle p \rangle \subseteq R$. We recall that by Dedekind's theorem, $\pi_K(p)$ is the number of distinct irreducible factors of $\Phi = x^n + 1$ over $\mathbb{Z}_p$, so $\pi_K(p) \leq \min(n, p)$. Also, since K is a normal extension of $\mathbb{Q}$ with Δ_K a power of 2, all the prime ideals above $p > 2$ have identical norm $p^{n/\pi_K(p)}$ (see, e.g., [34, Ch. 4]). Using this, we have, for s $>$ 1:

$$\begin{aligned}
\zeta_K(s) &= \prod_{\text{prime } p} \prod_{\text{prime } J | \langle p \rangle} (1 - \mathcal{N}(J)^{-s})^{-1} \\
&= \frac{2^s}{2^s - 1} \prod_{\text{prime } p > 2} (1 - p^{-sn/\pi_K(p)})^{-\pi_K(p)} \\
&\leq \frac{2^s}{2^s - 1} \prod_{\text{prime } p,\ 2 < p \leq n} (1 - p^{-sn/p})^{-p} \cdot \prod_{\text{prime } p > n} (1 - p^{-s})^{-n}.
\end{aligned}$$

We used the fact that for any fixed $x \in (0,1)$, the function $t \mapsto (1 - x^{-1/t})^{-t}$ is non-decreasing for $t > 0$.

We first deal with the case $s = 2$, where we have:

$$\begin{aligned}
\zeta_K(2) &\leq \frac{4}{3} \prod_{\text{prime } p,\ 2 < p \leq n/2} (1 - p^{-4})^{-p} \cdot \prod_{\text{prime } p,\ n/2 < p \leq n} (1 - p^{-2})^{-p} \cdot \prod_{\text{prime } p > n} (1 - p^{-2})^{-n} \\
&\leq \frac{4}{3} \exp\left(\sum_{\text{prime } p,\ 2 < p \leq n} (p^{-3} + p^{-7}) + \sum_{\text{prime } p,\ n/2 < p \leq n} p^{-1} + n \sum_{\text{prime } p > n} (p^{-2} + p^{-4}) \right),
\end{aligned}$$

where we used the inequality $\ln(1 - x) \geq -x - x^2$, for $x \in [0, 1/3]$. We now show that each one of these sums is $O(1)$. We have:

$$\sum_{\text{prime } p \leq n} p^{-3} \leq \int_1^n x^{-3} dx \leq 1/2.$$

Similarly, we have $\sum_{p \leq n} p^{-7} \leq 1/6$, $\sum_{p > n} p^{-2} \leq n^{-1}$ and $\sum_{p > n} p^{-4} \leq n^{-3}/3$. It remains to bound $\sum_{n/2 < p \leq n} p^{-1}$. It is proved in [52, Th. 9, p. 16] that $\sum_{p \leq x} p^{-1} = \log\log x + c + O(1/\log x)$, for some constant c. We thus obtain that:

$$\sum_{\text{prime } p,\ n/2 < p \leq n} p^{-1} \leq \log \frac{\log n}{\log(n/2)} + O\left(\frac{1}{\log n}\right) = \log\left(1 + \frac{\log 2}{\log(n/2)}\right) + O\left(\frac{1}{\log n}\right) = O\left(\frac{1}{\log n}\right).$$

We now consider the case $s = 1 + \varepsilon$. We have:

$$\begin{aligned}
\zeta_K(1+\varepsilon) &\leq 2 \prod_{\text{prime } p,\ 2 < p \leq n} (1 - p^{-(1+\varepsilon)n/p})^{-p} \cdot \prod_{\text{prime } p > n} (1 - p^{-(1+\varepsilon)})^{-n} \\
&\leq 2\exp\left(\sum_{\text{prime } p,\ 2 < p \leq n} (p^{-(1+\varepsilon)\frac{n}{p}+1} + p^{-2(1+\varepsilon)\frac{n}{p}+1}) + n \cdot \sum_{\text{prime } p > n} (p^{-(1+\varepsilon)} + p^{-2(1+\varepsilon)}) \right).
\end{aligned}$$

where we again used the inequality $\ln(1-x) \geq -x - x^2$, for $x \in [0, 1/3]$. The first sum above is bounded as:

$$2 \cdot \sum_{\text{prime } 2 < p \leq n} p^{-\varepsilon} \leq 2 \int_2^n x^{-\varepsilon} dx \leq 2 \frac{n^{1-\varepsilon}}{1-\varepsilon}.$$

Similarly, the second sum above is bounded as $2n \cdot \sum_{p>n} p^{-(1+\varepsilon)} \leq 2 \cdot \varepsilon^{-1} \cdot n^{1-\varepsilon}$. This gives the claimed bound on $\zeta_K(1+\varepsilon)$. □

In our study of the Dedekind zeta function (to be used for analyzing the key generation algorithm of NTRU), we use the following bound.

Lemma 2.12. *Let $N \geq 1$ and $\varepsilon \in (0,1)$. The number $H(N)$ of ideals $I \subseteq R_n$ satisfying $\mathcal{N}(I) \leq N$ is bounded as $H(N) \leq 2 \exp(2 \cdot (\varepsilon(1-\varepsilon))^{-1} \cdot n^{1-\varepsilon}) \cdot N^{1+\varepsilon}$.*

Proof. For $k \geq 1$, let $M(k)$ denote the number of ideals of R_n of norm exactly k. Observe that for $s > 1$, we have $\zeta_K(s) = \sum_{I \subseteq R} \mathcal{N}(I)^{-s} = \sum_{k \geq 1} M(k) \cdot k^{-s} \geq \sum_{k \leq N} M(k) \cdot k^{-s}$. Using $\sum_{k \leq N} M(k) \cdot k^{-s} \geq \sum_{k \leq N} M(k) \cdot N^{-s} = H(N) \cdot N^{-s}$, we obtain that $H(N) \leq \zeta_K(s) \cdot N^s$. Setting $s = 1 + \varepsilon$ and applying Lemma 2.11 completes the proof. □

The value $\zeta_{\mathbb{Q}}(2) = \pi^2/6$ is famous because its inverse is the probability that two "random" integers are co-prime. The next lemma considers the generalization of that fact to K_n.

Lemma 2.13. *Assume that $\sigma \geq n^{1.5} \ln^5 n$. Then, for n sufficiently large:*

$$\Pr_{f,g \hookleftarrow D_{R,\sigma}} [\langle f, g \rangle \neq R] \leq 1 - \frac{1}{2\zeta_K(2)} + 2^{-n+1}.$$

Proof. By Lemma 2.4, we have:

$$\begin{aligned} \Pr[\langle f, g \rangle \neq R] &\leq \Pr[\langle f, g \rangle \neq R \ \& \ \|f\|, \|g\| \leq \sqrt{n}\sigma] + \Pr[\|f\| > \sqrt{n}\sigma \text{ or } \|g\| > \sqrt{n}\sigma] \\ &\leq \Pr[\langle f, g \rangle \neq R \ \& \ \|f\|, \|g\| \leq \sqrt{n}\sigma] + 2^{-n+1}. \end{aligned}$$

We bound $\Pr[\langle f, g \rangle \neq R \ \& \ \|f\|, \|g\| \leq \sqrt{n}\sigma]$ by using an argument inspired by [49]. Since any ideal I containing the principal ideal $\langle f \rangle$ has norm $\mathcal{N}(I) \leq \mathcal{N}(\langle f \rangle)$, the condition $\|f\| \leq \sqrt{n}\sigma$ implies $\mathcal{N}(I) \leq \mathcal{N}(\langle f \rangle) \leq (\sqrt{n}\sigma)^n$. Therefore, we have $\Pr[\langle f, g \rangle \neq R \ \& \ \|f\|, \|g\| \leq \sqrt{n}\sigma] \leq 1 - p$, with

$$p := D_{\mathbb{Z}^{2n},\sigma}\Bigg(\mathbb{Z}^{2n} \setminus \bigcup_{\substack{\text{prime } I \subseteq R \\ \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} I \times I\Bigg) = \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} \mu(I) \cdot D_{\mathbb{Z}^n,\sigma}(I)^2,$$

where in the second equality, we used the inclusion-exclusion principle (and μ is the Möbius function for ring R). Recall that $\zeta_K(2)^{-1} = \sum_{I \subseteq R} \mu(I) \cdot \mathcal{N}(I)^{-2}$. We now show that $\left|p - \frac{1}{\zeta_K(2)}\right| \leq \frac{1}{2\zeta_K(2)}$. This implies $p \geq \frac{1}{2\zeta_K(2)}$, as required. We have:

$$\left|p - \frac{1}{\zeta_K(2)}\right| \leq \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} \left|D_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| + \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) > (\sqrt{n}\sigma)^n}} \mathcal{N}(I)^{-2}.$$

To bound the first sum, we recall that for any (even fractional) ideal I, we have $\lambda_n(I) = \lambda_1(I) \leq \sqrt{n}\mathcal{N}(I)^{\frac{1}{n}}$, so for any $\delta \in (0, 1/2)$, the smoothing parameter $\eta_\delta(I)$ is smaller than $B_\delta \cdot \mathcal{N}(I)^{\frac{1}{n}}$, where $B_\delta = \sqrt{n \ln(2n(1 + 1/\delta))/\pi}$ (by Lemma 2.1). It follows from Lemma 2.3 that $\left|D_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| \leq 16\delta/\mathcal{N}(I)^2$ if$\mathcal{N}(I) \leq (\sigma/B_\delta)^n$ and $I \subseteq R$. Assume now that $(\sigma/B_\delta)^n < \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n$, and let $k = \left\lceil \frac{\mathcal{N}(I)^{\frac{1}{n}}}{\sigma/B_\delta} \right\rceil$. Since $I \subseteq \frac{1}{k} \cdot I$, we have $D_{\mathbb{Z}^n,\sigma}(I) \leq D_{\mathbb{Z}^n,\sigma}(\frac{1}{k} \cdot I)$. Also, by the choice of k, we have $\eta_\delta(\frac{1}{k} \cdot I) = \frac{1}{k}\eta_\delta(I) \leq \sigma$. Now, $D_{\mathbb{Z}^n,\sigma}(\frac{1}{k} \cdot I) = \frac{\rho_\sigma(\frac{1}{k} \cdot I \cap \mathbb{Z}^n)}{\rho_\sigma(\mathbb{Z}^n)} \leq \frac{\rho_\sigma(\frac{1}{k} \cdot I)}{\rho_\sigma(\mathbb{Z}^n)} \leq (\frac{2B_\delta}{\sigma})^n \frac{1+\delta}{1-\delta}$, where in the last inequality we applied Lemma 2.3 twice, assuming $\sigma \geq \eta_\delta(\mathbb{Z}^n)$, and using $\det(\frac{1}{k} \cdot I) = \frac{1}{k^n} \cdot \mathcal{N}(I) \geq (\frac{\sigma}{2B_\delta})^n$. Therefore, $D_{\mathbb{Z}^n,\sigma}(I)^2 \leq (\frac{2B_\delta}{\sigma})^{2n} \frac{(1+\delta)^2}{(1-\delta)^2}$. Finally, assuming that $\sigma \geq 2B_\delta$ and $\delta = 2^{-7}$, we obtain:

$$\begin{aligned}
\sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} \left|D_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| &\leq \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sigma/B_\delta)^n}} \left|D_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| + \sum_{\substack{I \subseteq R \\ (\sigma/B_\delta)^n < \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} \left|D_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| \\
&\leq 16\delta \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sigma/B_\delta)^n}} \frac{1}{\mathcal{N}(I)^2} \; + \; 2 \cdot H((\sqrt{n}\sigma)^n) \cdot \left(\frac{2B_\delta}{\sigma}\right)^{2n} \\
&\leq \frac{\zeta_K(2)}{8} \; + \; 2 \cdot H((\sqrt{n}\sigma)^n) \cdot \left(\frac{2B_\delta}{\sigma}\right)^{2n},
\end{aligned}$$

where $H(N)$ is the number of (integral) ideals of R of norm $\leq N$. From Lemma 2.12 with $\varepsilon = \frac{\log\log n}{\log n}$, we know that $H(N) \leq 2\exp(4n) \cdot N^{1+\varepsilon}$. Taking $\sigma \geq n^{1.5} \ln^5 n$ provides $H\left((\sqrt{n}\sigma)^n\right) \cdot \left(\frac{2B_\delta}{\sigma}\right)^{2n} \leq \frac{1}{16\zeta_K(2)}$, for sufficiently large n. Overall, the first sum is $\leq \frac{1}{4\zeta_K(2)}$ for n sufficiently large.

We now bound the second sum, as follows:

$$\begin{aligned}
\sum_{\substack{I \subseteq R \\ \mathcal{N}(I) > (\sqrt{n}\sigma)^n}} \mathcal{N}(I)^{-2} = \sum_{k > (\sqrt{n}\sigma)^n} \frac{H(k) - H(k-1)}{k^2} &= \sum_{k > \lfloor(\sqrt{n}\sigma)^n\rfloor} \frac{H(k)}{k^2} - \sum_{k \geq \lfloor(\sqrt{n}\sigma)^n\rfloor} \frac{H(k)}{(k+1)^2} \\
&\leq \sum_{k > (\sqrt{n}\sigma)^n} H(k) \left(\frac{1}{k^2} - \frac{1}{(k+1)^2}\right) \\
&\leq 2\exp(4n) \cdot \sum_{k \geq (\sqrt{n}\sigma)^n} \frac{2k+1}{k^{1-\varepsilon}(k+1)^2},
\end{aligned}$$

where we used the bound on $H(k)$ from Lemma 2.12. Now, notice that the summand is $\leq \frac{2}{k^{2-\varepsilon}}$, which allows us to bound the second sum by $O(\exp(4n) \cdot (\sqrt{n}\sigma)^{-(1-\varepsilon)n}) = o(1)$, so the latter is $\leq \frac{1}{4\zeta_K(2)}$ for sufficiently large n, which completes the proof. □

MODULE q-ARY LATTICES. Let q be a prime number, and R_q be $R/qR = \mathbb{Z}_q[x]/\Phi$. In the present work, we consider a q that splits Φ into n distinct linear factors: $\Phi = \prod_{i \leq n} \Phi_i = \prod_{i \leq n}(x - \phi_i) \bmod q$. This is equivalent to assuming that the prime number q satisfies $q = 1 \bmod n$. 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 the upper bound (the current record seems to be $O(n^{5.2})$, see [53]). 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 actually provides a natural fast Discrete Fourier Transform, and thus multiplication of elements of R_q can be performed within $n \log n$ additions and multiplications modulo q (see [7, Ch. 8], [25, Se. 2.1]).

Let $\boldsymbol{a} \in R_q^m$. We define the following families of R-modules:

$$\boldsymbol{a}^{\perp} := \{(t_1, \ldots, t_m) \in R^m : \sum_i t_i a_i = 0 \bmod q\},$$
$$L(\boldsymbol{a}) := \{(t_1, \ldots, t_m) \in R^m : \exists s \in R_q, \forall i, t_i = a_i \cdot s \bmod q\}.$$

These modules correspond to mn-dimensional integer lattices, via the mapping of an element of R^m to the concatenation of the coefficient vectors.

Recently, Peikert [40] showed how to significantly improve on the efficiency of the Gaussian sampling algorithm from [10], in the case of q-ary lattices, and even further for module q-ary lattices. In the following adaptation, we bound Peikert's $s_1(B)$ by $\sqrt{n} \max \|\boldsymbol{b}_i\|$ (using the Cauchy-Schwarz inequality).

Lemma 2.14 (Adapted from [40]). *There exists a $\widetilde{O}(nm)$-time off-line/on-line algorithm that takes as input any R-basis $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m$ of a module q-ary lattice $L \subseteq R^m$, with $q = \mathcal{P}oly(n)$, $\mathbf{c} \in \mathbb{Q}^{mn}$ and $\sigma = \omega(\sqrt{mn \log n}) \max \|\boldsymbol{b}_i\|$ (resp. $\sigma = \Omega(\sqrt{m}n) \max \|\boldsymbol{b}_i\|$), and returns samples from a distribution whose statistical distance to $D_{L,\sigma,\mathbf{c}}$ is negligible (resp. exponentially small) with respect to n. The complexity bound holds assuming pre-computations (off-line) are performed using q, σ and $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m$, but not $\mathbf{c}$.*

2.3 The Shortest Vector, Ideal-SIS and R-LWE Problems

THE SHORTEST VECTOR PROBLEM. The most famous lattice problem is SVP. Given a basis of a lattice L, it aims at finding a shortest vector in $L \setminus \{\mathbf{0}\}$. 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 $\gamma(\cdot)$. If we restrict the set of input lattices to ideal lattices, we obtain the problem Ideal-SVP (resp. γ-Ideal-SVP), which is implicitly parameterized by a sequence of polynomials Φ of growing degrees. No algorithm is known to perform non-negligibly better for (γ-)Ideal-SVP than for (γ-)SVP. It is believed that no subexponential quantum algorithm solves the computational variants of γ-SVP or γ-Ideal-SVP in the worst case, for any γ that is polynomial in the dimension. The smallest γ which is known to be achievable in polynomial time is exponential, up to poly-logarithmic factors in the exponent ([22, 48, 32]).

THE IDEAL SMALL INTEGER SOLUTION PROBLEM. Ideal-SIS is an average-case variant of γ-SVP in certain structured lattices.

Definition 2.1. *The Ideal Small Integer Solution problem with parameters q, m, β and Φ (Ideal-SIS$^{\Phi}_{q,m,\beta}$) is as follows: Given n, and m polynomials $a_1, \ldots, a_m$ chosen uniformly and independently in R_q, find $\boldsymbol{t} \in \boldsymbol{a}^{\perp} \setminus \mathbf{0}$ such that $\|\boldsymbol{t}\| \leq \beta$.*

The average-case hardness of Ideal-SIS is related to the worst-case hardness of Ideal-SVP, as follows.

Theorem 2.1 (Adapted from [24]). *Let $n = 2^k$, $\Phi = x^n + 1$ and $\varepsilon > 0$. Let $m \leq \mathcal{P}oly(n)$ and $q = \Omega(\beta m^2 n (\log n)^{1/2+\varepsilon})$ be integers. A polynomial-time (resp. subexponential-time) algorithm solving* Ideal-SIS$^{\Phi}_{q,m,\beta}$ *with probability $1/\mathcal{P}oly(n)$ (resp. $2^{-o(n)}$) can be used to solve γ-*Ideal-SVP *in polynomial-time (resp. subexponential-time) with $\gamma = O(\beta m n^{1/2} (\log n)^{1+\varepsilon})$ (resp. $\gamma = O(\beta m n^{1.5}\sqrt{\log n})$).*

The Ring Learning With Errors 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 uniformly chosen in R_q and e sampled independently from ψ. The Ring Learning With Errors problem (R-LWE) was introduced by Lyubashevsky et al. in [27] and shown hard for specific error distributions ψ. These are slightly technical to define, but for the present work, the important facts to be remembered are that the samples are small (with probability exponentially close to 1), and can be obtained in quasi-linear time.

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 [27], 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 these correspond to complex discrete Fourier transforms. These matrix-vector multiplications can be performed in $O(n \log n)$ complex-valued arithmetic operations with the Cooley-Tukey FFT. Moreover, they are 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 $\sigma_i = \sqrt{\alpha^2 q^2 + x_i}$ with the x_i's sampled independently from the distribution $\mathrm{Exp}(n\alpha^2 q^2)$.

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). Finally, by taking $r = 1$ in the result below, we obtain that with probability $\geq 1 - n^{-\omega(1)}$, any sample from $\overline{\Upsilon}_\alpha$ in R has norm $\leq \alpha q \sqrt{n}\omega(\log n)$.

Lemma 2.15. *Let $y, r \in R$, with r fixed and y sampled from $\overline{\Upsilon}_\alpha$. Then*

$$\Pr\left[\|yr\| \geq \alpha q \sqrt{n}\omega(\log n) \cdot \|r\|\right] \leq n^{-\omega(1)} \quad \text{and} \quad \Pr\left[\|yr\|_\infty \geq 4\alpha q \omega(\log n) \cdot \|r\|\right] \leq n^{-\omega(1)}.$$

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 be sampled from Υ_α. The involved $\boldsymbol{\sigma}$ satisfies $\sigma_k = \sqrt{\alpha^2 q^2 + x_k}$, with the x_k's sampled independently from the distribution $\mathrm{Exp}(n\alpha^2 q^2)$. We have $\max \sigma_k \leq \alpha q \sqrt{n}\omega(\sqrt{\log n})$ with probability $1 - n^{-\omega(1)}$. The field element $y \in K$ is sampled from $\rho'_{\boldsymbol{\sigma}}$, and actually derived from a sample $\boldsymbol{z}$

from $\rho_{\boldsymbol{\sigma}}$. 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 yr is $(y^{(k)}r^{(k)})_k$. We have $\max_k |y^{(k)}r^{(k)}| \leq \alpha q\sqrt{n}\omega(\log n) \cdot |r^{(k)}|$, with probability $1 - n^{-\omega(1)}$. We thus obtain $\|yr\| = \frac{1}{\sqrt{n}}T_2(yr) \leq \alpha q\omega(\log n) \cdot T_2(r) = \alpha q\sqrt{n}\omega(\log n) \cdot \|r\|$.

We now prove the second statement. The coefficient in x^j of yr is

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

The ith summand of the last sum follows a normal law of mean 0 and standard deviation $\leq 2|r^{(i)}|\max\sigma_k$. Therefore, the coefficient in x^j of yr follows a normal law of standard deviation $\leq \frac{4}{n}T_2(r)\max\sigma_k$, which is $\leq \frac{4}{\sqrt{n}}\alpha q\omega(\sqrt{\log n}) \cdot T_2(r)$ with probability $1 - n^{-\omega(1)}$. This completes the proof. □

We now define our adaptation of R-LWE.

Definition 2.2. *The Ring Learning With Errors Problem with parameters q, α and Φ (R-LWE$^{\Phi}_{q,\alpha}$) is as follows. Let ψ be sampled from $\overline{\Upsilon}_\alpha$ and s be chosen uniformly in 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 the distribution $A_{s,\psi}$ or $U(R_q \times R_q)$. 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.*

R-LWE can be interpreted as a problem over q-ary module lattices. Let m be the number of samples asked to the oracle, and let $(a_i, b_i)_{i\leq m}$ be the samples. Then solving R-LWE consists in telling whether the vector $\boldsymbol{b}$ is generated uniformly modulo the (module) lattice $L(\boldsymbol{a})$ or sampled around the origin according to some Gaussian-like distribution derived from $\overline{\Upsilon}_\alpha$ and then reduced modulo the lattice.

Theorem 2.2 (Adapted from [27]). *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 main differences in the above formulation of the result from [27] are the use of the polynomial representation (which is handled by applying the complex FFT to the error term), the use of R_q rather than $R_q^\times$ (here we have $R_q^\times = \frac{1}{n}R_q$, and the truncation of the error to closest integer when the latter is away from the middle of two consecutive integers). The new variant remains hard because a sample passes the rejection step with non-negligible probability, and because 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 uniformly chosen in $R_q^\times$ and e sampled independently from ψ, 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 $A_{s,\psi}^\times$ and $U(R_q^\times \times R_q)$. We call R-LWE$^\times$ the latter variant.

Furthermore, similarly to [3, Le. 2] and as mentioned in [41, Sl. 8], the nonce s can also be chosen from the error distribution without incurring any security reduction. We call R-LWE$_{\text{HNF}}^\times$ the corresponding modification of R-LWE. We recall the argument, for completeness. Assume an algorithm $\mathcal{A}$ can solve R-LWE$_{\text{HNF}}^\times$. We use $\mathcal{A}$ to solve R-LWE$^\times$. The principle is to transform samples $((a_i, b_i))_i$ into samples $((a_1^{-1}a_i, b_i - a_1^{-1}b_1 a_i))_i$, where inversion is performed in $R_q^\times$. This transformation maps $A_{s,\psi}^\times$ to $A_{-e_1,\psi}^\times$, and $U(R_q^\times \times R_q)$ to itself.

3 New Results on Module q-ary Lattices

In the present section, we exploit the duality between variants of the $\boldsymbol{a}^\perp$ and $L(\boldsymbol{a})$ lattices to obtain improved regularity bounds over the ring R_q.

3.1 Duality results for generalized module q-ary lattices

We generalize the definitions of the $\boldsymbol{a}^\perp$ and $L(\boldsymbol{a})$ lattices to incorporate the ideals of R_q, as this will be useful for key generation procedures of the modified NTRU schemes (in Section 4). 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 $\{1, \ldots, n\}$. For any $I_S = \prod_{i\in S}(x - \phi_i) \cdot R_q$, we define $I_S^\times = \prod_{i\in S}(x - \phi_i^{-1}) \cdot R_q$.

Let $\boldsymbol{a} \in R_q^m$. We define the following families of R-modules:

$$\boldsymbol{a}^\perp(I_S) := \{(t_1, \ldots, t_m) \in R^m : \forall i, (t_i \bmod q) \in I_S \ \text{ and } \ \sum_i t_i a_i = 0 \bmod q\},$$
$$L(\boldsymbol{a}, I_S) := \{(t_1, \ldots, t_m) \in R^m : \exists s \in R_q, \forall i, (t_i \bmod q) = a_i \cdot s \bmod I_S\},$$

where S is an arbitrary subset of $\{1, \ldots, n\}$. If $S = \emptyset$ (resp. $S = \{1, \ldots, n\}$), then we recover $\boldsymbol{a}^\perp$ (resp. $L(\boldsymbol{a})$).

Lemma 3.1. *Let $S \subseteq \{1, \ldots, n\}$ and $\boldsymbol{a} \in R_q^m$. Let $\overline{S}$ be the complement of S and $\boldsymbol{a}^\times \in R_q^m$ be defined by $a_i^\times = a_i(x^{-1})$. Then (considering both sets are considered as mn-dimensional lattices):*

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

Proof. We first prove that $\frac{1}{q}L(\boldsymbol{a}^*, I_{\overline{S}}^\times) \subseteq \widehat{\boldsymbol{a}^\perp(I_S)}$. Let $(t_1, \ldots, t_m) \in \boldsymbol{a}^\perp(I_S)$ and $(t_1', \ldots, t_m') \in L(\boldsymbol{a}^*, I_{\overline{S}})$. Write $t_i = \sum_{j<n} t_{i,j} x^j$ and $t_i' = \sum_{j<n} t_{i,j}' x^j$ for any $i \leq m$. Our goal is to show that $\sum_{i\leq m, j\leq n} t_{i,j} t_{i,j}' = 0 \bmod q$. This is equivalent to showing that the constant coefficient of the polynomial $\sum_{i\leq m} t_i(x) t_i'(x^{-1})$ is 0 modulo q. It thus suffices to show that $\sum_{i\leq m} t_i(x) t_i'(x^{-1}) \bmod q = 0$ (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 $b_i' \in I_{\overline{S}}^\times$. We have the following, modulo q:

$$\sum_{i\leq m} t_i(x) t_i'(x^{-1}) = s(x^{-1}) \cdot \sum_{i\leq m} t_i(x) a_i(x) + \sum_{i\leq m} t_i(x) b_i'(x^{-1}).$$

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 $\widehat{L(\boldsymbol{a}^\times, I_{\overline{S}}^\times)} \subseteq \frac{1}{q}\boldsymbol{a}^\perp(I_S)$. 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 a uniformly chosen $\boldsymbol{a} \in (R_q^\times)^m$, 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 $\det(L(\boldsymbol{a}, I_S))^{\frac{1}{mn}} = q^{\frac{|S|}{n} - \frac{1}{m}}$.

Lemma 3.2. *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 5$. Then, for any $S \subseteq \{1, \dots, n\}$, $m \geq 2$ and $\varepsilon > 0$, we have $\lambda_1^\infty(L(\boldsymbol{a}, I_S)) \geq \frac{1}{\sqrt{n}} q^\beta$, with:*

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

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:

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

For $h = \prod_{i \in S'} \Phi_i$ 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 subideal 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 $h = \prod_{i \in S'} \Phi_i$ 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 of small height. Similarly to [28], 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 is very small, namely $m = 2$.

The regularity result in [28, 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 of height $\leq d$ for some $d < q$. In this case, the regularity bound from [28] 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 [10, 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 [50, Se. 3], thus extending the latter to a fully splitting q. It also shows that the encryption scheme from [27] can be shown secure against subexponential (quantum) attackers, assuming the subexponential (quantum) hardness of standard worst-case problems over ideal lattices.

Theorem 3.1. *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:*

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

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 uniform bound $\Delta_{\boldsymbol{a}} \leq 2\delta$, for all except a fraction $\leq (q-1)^{-\varepsilon n}$ 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. 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 $\boldsymbol{t} \bmod \boldsymbol{a}^\perp$. By Lemma 2.5, 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 2.2, 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 2.2, 2.5, 3.1 and 3.2.

Lemma 3.3. *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:*

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

Theorem 3.1 follows by taking $S = \emptyset$ and $\boldsymbol{c} = \boldsymbol{0}$.

4 Revised key generation algorithms for the NTRU schemes

We now use the results of the previous section on modular q-ary lattices to derive key generation algorithms for the NTRU schemes, where the generated public keys follow distributions for which Ideal-SVP is known to reduce to R-LWE and Ideal-SIS.

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 2.7), 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 2.7. 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 $R_q^\times$.

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`.

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

Lemma 4.1. *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 $\Pr_{f' \hookleftarrow D_{\mathbb{Z}^n,\sigma}}[(p \cdot f' + a \bmod q) \notin R_q^\times] \leq n(1/q + 2\delta)$.*

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 $k \leq n$. The result then follows from the Chinese Remainder Theorem and the union bound. We have $\mathcal{N}(I) = q$, so that $\lambda_1(I) \leq \sqrt{n} q^{1/n}$, by Minkowski's theorem. Since I is an ideal of R, we have $\lambda_n(I) = \lambda_1(I)$, and Lemma 2.1 gives that $\sigma \geq \eta_\delta(I)$. Lemma 2.5 then shows that $f \bmod I$ is within distance $\leq 2\delta$ to uniformity on R/I, so we have $p \cdot f' + a = 0 \bmod I$ (or, equivalently, $f' = -a/p \bmod I$) with probability $\leq 1/q + 2\delta$, 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 4.2. *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 8n$. Let $\sigma \geq \sqrt{2n \ln(6n)/\pi} \cdot q^{1/n}$. The secret key polynomials f, g returned by the algorithm of Fig. 1 satisfy, with probability $\geq 1 - 2^{-n+3}$:*

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

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

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 2.4 and 4.1. □

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 $\sigma \geq \mathcal{P}oly(n) \cdot q^{1/2+\varepsilon}$ for an arbitrary $\varepsilon > 0$, but restricted (by rejection) to $R_q^\times - p^{-1}$ and $R_q^\times$, respectively. We denote by $D_{\sigma,z}^\times$ 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 $g/f \bmod q$ computed by the algorithm of Fig. 1. Since $p \in R^\times_q$, multiplication by p is a bijection of R_q, and thus the statistical closeness to uniformity carries over to the public key $h = pg/f$.

Theorem 4.1. *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 $\varepsilon > 0$ and $\sigma \geq 2n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+2\varepsilon}$. Let $p \in R^\times_q$, $y_i \in R_q$ and $z_i = -y_i p^{-1} \bmod q$ for $i \in \{1, 2\}$. Then*

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

Proof. For $a \in R^\times_q$, we define $Pr_a = \Pr_{f_1,f_2}[(y_1 + pf_1)/(y_2 + pf_2) = a]$, where $f_i \hookleftarrow D^\times_{\sigma,z_i}$ 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 $a \in R^\times_q$. This directly gives the claimed bound. The fraction of $a \in R^\times_q$ such that $|Pr_a - (q-1)^{-n}| \leq \varepsilon'$ is equal to the fraction of $\boldsymbol{a} = (a_1, a_2) \in (R^\times_q)^2$ 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 $R^\times_q$), and $-a_2/a_1$ is uniformly random in $R^\times_q$ when $\boldsymbol{a} \hookleftarrow U((R^\times_q)^2)$.

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 $\boldsymbol{a}^{\perp\times} = \boldsymbol{a}^{\perp} \cap (R^\times_q + q\mathbb{Z}^n)^2$. Therefore:

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

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 $-a_1/a_2 \in R^\times_q$, 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 $R^\times_q + q\mathbb{Z}^n = \mathbb{Z}^n \setminus \bigcup_{S \subseteq \{1,\ldots,n\}, S \neq \emptyset} (I_S + q\mathbb{Z}^n)$. Using the inclusion-exclusion principle, we obtain:

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

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

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

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

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

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 first notice 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 $|S| \le \varepsilon n$, we apply Lemma 3.3 with $m = 2$. Since $|S|/n + \varepsilon \le 2\varepsilon$, the Lemma 3.3 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 $\boldsymbol{a} \in (R_q^\times)^2$:

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

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 2.1 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 2.5, 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:

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

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

4.3 NTRUSign's key generation algorithm

Our new key generation for `NTRUSign` is given in Fig. 2. It is inspired from the algorithm contained in [15, Se. 4] and described in more details in [14, Se. 5]. The vector (f, g) produced by the `NTRUEncrypt` key generation algorithm is a short vector in the R-module generated by the rows of the matrix $\begin{bmatrix} 1 & g/f \\ 0 & q \end{bmatrix}$. The goal of the algorithm of Fig. 2 is to extend this vector (f, g) into a short basis $\begin{bmatrix} f & g \\ F & G \end{bmatrix}$ of the module.

Because of the rejection tests, the output public key h may not be uniformly distributed in $R_q^\times$, as it was previously. Uniformity is important for us to eventually be able to use Theorem 2.1 to prove the security of the signature scheme. In fact, as we will show in Subsection 5.2, it suffices that the combined rejection probabilities of Steps 3, 4 and 7 is non-negligibly away from 1.

By Lemma 2.13, when no rejection is performed in Steps 1–3, the rejection probability of Step 4 is (assuming that $\sigma \ge n^{3/2} \ln^5 n$ and that n is sufficiently large):

$$\Pr_{f,g \hookleftarrow D_{R,\sigma}}[\langle f, g \rangle \ne R] \le 1 - \frac{1}{2\zeta_K(2)} + 2^{-n+1}.$$

Inputs: $n, q \in \mathbb{Z}$, $\sigma \in \mathbb{R}$.
Output: A key pair $(sk, pk) \in R^{2\times 2} \times R_q^\times$.
1. Sample f from $D_{\mathbb{Z}^n,\sigma}$; 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. If $\|f\| > \sqrt{n} \cdot \sigma$ or $\|g\| > \sqrt{n} \cdot \sigma$, restart.
4. If $\langle f, g\rangle \neq R$, restart.
5. Compute $F_1, G_1 \in R$ such that $fG_1 - gF_1 = 1$; $F_q := qF_1$, $G_q := qG_1$.
6. Use Babai's nearest plane algorithm to approximate (F_q, G_q) by an integer linear combination of $(f, g), (xf, xg), \ldots, (x^{n-1}f, x^{n-1}g)$. Let $(F, G) \in R^2$ be the output, such that there exists $k \in R$ with $(F, G) = (F_q, G_q) - k(f, g)$.
7. If $\|(F, G)\| > n\sigma$, restart.
8. Return secret key $sk = \begin{bmatrix} f & g \\ F & G \end{bmatrix}$ and public key $pk = h = g/f \in R_q^\times$.

Fig. 2. Revised Key Generation Algorithm for `NTRUSign`.

We now consider the rejection probability of Step 7 (without rejection in Steps 1–2).

Lemma 4.3. *Assume that $\sigma = \Omega(\sqrt{\log n})$. Then, as n grows to infinity:*

$$\Pr_{f,g \hookleftarrow D_{R,\sigma}^\times}\left[\|(F,G)\|^2 \leq \frac{1}{2}n^2\sigma^2 + \frac{q^2 \cdot \omega(n)}{\sigma^2} \;\middle|\; \langle f, g\rangle = R\right] = o(1),$$

where F and G are as defined in Steps 5 and 6 of the algorithm of Figure 2.

Proof. As we use Babai's nearest-plane algorithm, we have:

$$\|(F,G)\|^2 = \|(F_q, G_q)^*\|^2 + \|(e_f, e_g)\|^2,$$

where $(F_q, G_q)^*$ is the projection of (F_q, G_q) orthogonally to the K-span of (f, g). (this can also be interpreted as the projection of (F_q, G_q) orthogonally to the $\mathbb{Q}$-span of $(f, g), (xf, xg), \ldots, (x^{n-1}f, x^{n-1}g)$), and (e_f, e_g) is the rounding error of Babai's nearest plane algorithm, in rounding $(F_q, G_q) - (F_q, G_q)^*$ to a close point in the lattice $L(f, g)$ defined as the $\mathbb{Z}$-span of $(f, g), (xf, xg), \ldots, (x^{n-1}f, x^{n-1}g)$.

Since $\|(F_q, G_q)^*\| = \min_{k\in K} \|(F_q - kf, G_q - kg)\|$, to obtain an upper bound on $\|(F_q, G_q)^*\|$, it suffices to find a $k \in R$ such that $\|(F_q - kf, G_q - kg)\|$ is small. From the equation $fG_q - gF_q = q$, we obtain $G_q = qf^{-1} + g(f^{-1}F_q)$ (where inversion takes place in K). Taking $k := f^{-1}F_q$ gives $\|(F, G)^*\| \leq \|(0, qf^{-1})\| \leq q\|f^{-1}\|$. From Lemma 2.10 with "$t = \omega(n)$", we have that $\|f^{-1}\| \geq \frac{\omega(\sqrt{n})}{\sigma}$ with probability $\leq o(1)$, so $\|(F_q, G_q)^*\| \leq \frac{q\omega(\sqrt{n})}{\sigma}$, except with probability $o(1)$.

To upper bound $\|(e_f, e_g)\|$, note that $\|(e_f, e_g)\| \leq \frac{\sqrt{n}}{2}\max_i \|(x^i f, x^i g)\| = \frac{\sqrt{n}}{2}\|(f, g)\|$. By Lemma 2.4, we have $\|(f, g)\| \leq \sqrt{2n}\sigma$ with probability $\geq 1 - o(1)$, so $\|(e_f, e_g)\| \leq \frac{n\sigma}{\sqrt{2}}$, except with probability $o(1)$. This completes the proof. □

We can now analyze the overall rejection probability of the revised `NTRUSign` key generation algorithm.

Lemma 4.4. *Assume that $\sigma = \omega(\max(\frac{q^{\frac{1}{2}}}{n^{\frac{1}{4}}}, n^{1.5}\log^5 n))$ and $q \geq 128\zeta_K(2)n$. Then if n is sufficiently large, the combined rejection probability of Steps 3, 4 and 7 of the algorithm of Fig. 2 (i.e., when f and g are independently sampled from $D_\sigma^\times$) is $\leq 1 - c$, for some absolute constant $c > 0$.*

Proof. For $i \in \{3,4,7\}$, we denote by p_i the rejection probability of the test in Step i, i.e.:

- p_3 is the probability that $\|f\| > \sqrt{n}\sigma$ or $\|g\| > \sqrt{n}\sigma$, with $f, g \hookleftarrow D^{\times}_{R,\sigma}$.
- p_4 is the probability that $\langle f, g\rangle \neq R$ and $\|f\|, \|g\| \leq \sqrt{n}\sigma$, with $f, g \hookleftarrow D^{\times}_{R,\sigma}$.
- p_7 is the probability that $\max(\|F\|, \|G\|) > n\sigma$, $\langle f, g\rangle = R$ and $\|f\|, \|g\| \leq \sqrt{n}\sigma$, with $f, g \hookleftarrow D^{\times}_{R,\sigma}$.

For $i \in \{3,4,7\}$, let p'_i be defined exactly as p_i except that f and g are independently sampled from $D_{R,\sigma}$ rather than $D^{\times}_{R,\sigma}$. Let p_1 be the probability of rejection of f at Step 1. By the union bound, the probability of rejecting f or g at Steps 1 or 2 is $\leq 2p_1$. Hence for $i \in \{3,4,7\}$, we have $p_i \leq p'_i/(1-2p_1)$.

The rejection probability p_1 has already been studied in Subsection 4.1. Indeed, by Lemma 4.1 and the choice of σ and q, we have $p_1 \leq \frac{1}{32\zeta_K(2)}$. Lemmata 2.1 and 2.4 and the choice of σ imply that $p'_3 \leq 2^{-n+2}$. Finally, from Lemmata 2.13 and 4.3, we have that $p'_4 \leq 1 - \frac{1}{2\zeta_K(2)} + o(1)$ and $p'_7 = o(1)$. Recall from Lemma 2.11 that $\zeta_K(2) = O(1)$ when n grows to infinity, so for a large enough n, we have $p'_3+p'_4+p'_7 \leq 1 - \frac{1}{4\zeta_K(2)}$ and the total rejection probability $p_3+p_4+p_7 \leq \frac{p'_3+p'_4+p'_7}{1-2p_1} \leq 1 - \frac{1}{8\zeta_K(2)}$, as required. □

We can now conclude this section, with a correctness and efficiency statement for the revised `NTRUSign` key generation algorithm.

Theorem 4.2. *Let n be a power of 2 such that $\Phi = x^n + 1$ splits into n linear factors modulo prime $q \geq 128\zeta_K(2)n$. Let $\varepsilon \in (0, 1/2)$ and $\sigma \geq \max(2n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+2\varepsilon}, \omega(n^{1.5}\log^5 n))$. Then the algorithm of Fig. 2 terminates in expected polynomial time, and the output matrix $\begin{bmatrix} f & g \\ F & G \end{bmatrix}$ is an R-basis of the R-module spanned by the rows of $\begin{bmatrix} 1 & h \\ 0 & q \end{bmatrix}$. Furthermore, we have $\|(f,g)\| \leq 2\sqrt{n}\sigma$, and $\|(F,G)\| \leq n\sigma$. Finally, if n is sufficiently large, the distribution of the returned h is rejected with probability $c < 1$ for some absolute constant c from a distribution whose statistical distance from $U(R_q^{\times})$ is $\leq 2^{3n}q^{-\lfloor \varepsilon n \rfloor}$.*

Proof. The first statement is provided by Lemma 4.4. For the second statement, we refer to [15, Th. 1]. The norm inequalities are obvious from the description of the algorithm. Finally, the last statement is provided by Theorem 4.1 and Lemma 4.4. □

5 Cryptographic functions

Using our new results above, we describe in this section NTRU-like public-key encryption and digital signature schemes for which we can provide security proofs under worst-case hardness assumptions.In all constructions, we use $\Phi = x^n + 1$ with $n \geq 8$ a power of 2, $R = \mathbb{Z}[x]/\Phi$ and $R_q = R/qR$ with $q \geq 5$ prime such that $\Phi = \prod_{k=1}^{n} \Phi_k$ in R_q with distinct Φ_k's.

5.1 A revised `NTRUEncrypt` scheme

In this section we present the provably secure variant of the `NTRUEncrypt` scheme. We define the scheme `NTRUEncrypt` 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. 3. The encryption scheme `NTRUEncrypt`$(n, q, p, \sigma, \alpha)$.

The correctness conditions for the scheme are summarized below.

Lemma 5.1. *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 4.2, 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 2.15, both pfs and pge have infinity norms $\le 8\alpha q n^{1.5}\omega(\log n) \cdot \|p\|^2\sigma$ (resp. $32\alpha q\sqrt{n}\omega(\log n) \cdot \|p\|^2\sigma$), with probability $1 - n^{-\omega(1)}$. Independently, we have:

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

Since $\alpha q \ge \sqrt{n}$, we conclude that $\|C''\|_\infty \le (18+2\deg(p))\cdot\alpha q n^{1.5}\omega(\log n)\cdot\|p\|^2\sigma$ (resp. $72\alpha q n^{0.5}\omega(\log n)\cdot\|p\|^2\sigma$), with probability $1 - n^{-\omega(1)}$. □

The security of the scheme follows by a elementary reduction from the decisional R-LWE$^\times_{\mathrm{HNF}}$, exploiting the uniformity of the public key in $R_q^\times$ (Theorem 4.1), and the invertibility of p in R_q.

Lemma 5.2. *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 $A_{s,\psi}^{\times}$ for some previously chosen $s \hookleftarrow \psi$ and $\psi \hookleftarrow \overline{\Upsilon}_{\alpha}$. Algorithm $\mathcal{B}$ first calls oracle $\mathcal{O}$ to get a sample (h', C') from $R_q^{\times} \times R_q$. Then, algorithm $\mathcal{B}$ runs algorithm $\mathcal{A}$ 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 algorithm $\mathcal{A}$ 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 4.1, 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 sampled from ψ, the ciphertext C given to $\mathcal{A}$ has exactly the right distribution as in the IND-CPA attack. Overall, if $\mathcal{O}$ outputs samples from $A_{s,\psi}^{\times}$, then $\mathcal{A}$ succeeds and $\mathcal{B}$ returns 1 with probability $\geq 1/2 + \delta - q^{-\Omega(n)}$.

On the other hand, if oracle $\mathcal{O}$ 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 in this case, algorithm $\mathcal{B}$ outputs 1 with probability 1/2. The claimed advantage of $\mathcal{B}$ now follows. ☐

By combining Lemmata 5.1 and 5.2 with Theorem 2.2 we obtain our main result.

Theorem 5.1. *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 5.1 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. $\widetilde{O}(n^7)$).

5.2 A revised `NTRUSign` scheme

In this section we present a provably secure variant of `NTRUSign` (in the random oracle model). The scheme is an efficient variant of the GPV signature [10], where efficiency is improved both by using the ring structure (to reduce computation and storage from $\widetilde{O}(n^2)$ to $\widetilde{O}(n)$), and the NTRU key to reduce the key length and signature to a single ring element.

Collision-Resistant Preimage Sampleable Functions. We recall that the GPV signature [10] is built from a general cryptographic primitive introduced in [10] and called *Collision-Resistant Preimage Sampleable Functions* (CRPSF), which we recall.

Definition 5.1 (CRPSF). *A CRPSF is specified by three probabilistic polynomial-time algorithms* (TrapGen, SampleDom, SamplePre) *such that:*

1. Generating a Function with Trapdoor: *Given a security parameter n,* $\mathsf{TrapGen}(1^n)$ *returns (a, t), where a is the description of an efficiently computable function $f_a : \mathcal{D}_n \to \mathcal{R}_n$ (for some efficiently recognizable domain $\mathcal{D}_n$ and range $\mathcal{R}_n$), and t is a trapdoor string for f_a. In the following, we fix some pair (a, t) returned by* $\mathsf{TrapGen}(1^n)$*. Note that the following properties need only hold for with probability negligibly (resp. exponentially) close to 1 over the choice of (a, t) output by* $\mathsf{TrapGen}(1^n)$.
2. Domain Sampling with Uniform Output: *Given a security parameter n,* $\mathsf{SampleDom}(1^n)$ *returns x sampled from a distribution over $\mathcal{D}_n$ such that the statistical distance between $f_a(x)$ and the uniform distribution over $\mathcal{R}_n$ is negligible (resp. exponentially small).*
3. Preimage Sampling with Trapdoor: *Given any $y \in \mathcal{R}_n$,* $\mathsf{SamplePre}(t, y)$ *outputs x such that $f_a(x) = y$ and the distribution of x is within a negligible (resp. exponentially small) distance to the conditional distribution of $x' \hookleftarrow$* $\mathsf{SampleDom}(1^n)$ *given $f_a(x') = y$.*
4. Preimage Min-Entropy: *For each $y \in \mathcal{R}_n$, the conditional min-entropy of $x \hookleftarrow$* $\mathsf{SampleDom}(1^n)$ *given $f_a(x) = y$ is $\omega(\log n)$ (resp. $\Omega(n)$).*
5. Collision-Resistance without Trapdoor: *For any probabilistic polynomial-time (resp. subexponential-time) algorithm* F*, the probability that* $\mathsf{F}(1^n, a)$ *outputs distinct $x, x' \in \mathcal{D}_n$ such that $f_a(x) = f_a(x')$ is negligible (resp. exponentially small), where the probability is taken over the choice of a and the random coins of* F.

Our CRPSF construction $\mathtt{NTRUPSF}(n, q, \sigma, s)$ is shown in Fig. 4. The parameters n and q defining the rings R and R_q are as above. The parameter σ is the width of the discrete Gaussian distribution used in the `NTRUSign` key generation process, while s is the width of the Gaussian used in the preimage sampling.

- **Generating a Function with Trapdoor** – $\mathsf{TrapGen}(1^n, q, \sigma)$: Run the `NTRUSign` key generation algorithm from Fig. 2, using n, q, σ as inputs. It returns an NTRU key $h = g/f \in R_q^\times$ and a trapdoor R-basis $sk = \begin{bmatrix} f & g \\ F & G \end{bmatrix}$ for the R-module $h^\perp = \{(z_1, z_2) \in R^2 : z_2 = hz_1 \bmod q\}$. The key h defines function $f_h(z_1, z_2) = hz_1 - z_2 \in R_q$ with domain $\mathcal{D}_n = \{\boldsymbol{z} \in R^2 : \|\boldsymbol{z}\| \leq s\sqrt{2n}\}$ and range $\mathcal{R}_n = R_q$. The trapdoor string for f_h is sk.
- **Domain Sampling with Uniform Output** – $\mathsf{SampleDom}(1^n, q, s)$: Sample $\boldsymbol{z}$ from $D_{\mathbb{Z}^{2n},s}$; if $\|\boldsymbol{z}\| > \sqrt{2n}s$, resample.
- **Preimage Sampling with Trapdoor** – $\mathsf{SamplePre}(B, \boldsymbol{y})$: To find a preimage in $\mathcal{D}_n$ for target $t \in R_q$ under f_h using the trapdoor sk, note that $\boldsymbol{c} = (1, h - t)$ is a preimage of t under f_h (not necessarily in $\mathcal{D}_n$). Sample $\boldsymbol{z}$ from $D_{h^\perp + \boldsymbol{c}, s}$, using trapdoor basis sk for $h^\perp$ and the algorithm of Lemma 2.14. Return $\boldsymbol{z}$.

Fig. 4. Construction of CRPSF primitive $\mathtt{NTRUPSF}(n, q, \sigma, s)$.

Theorem 5.2. *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 $q^{\frac{1}{2}-\varepsilon} = \omega(n^{4.5} \log^{1.5+\varepsilon'} n)$ (resp. $q^{\frac{1}{2}-\varepsilon} = \Omega(n^5 \log^{1+\varepsilon'} n)$), for some arbitrary $\varepsilon, \varepsilon' > 0$. Let $\sigma = 2n\sqrt{\ln(8nq)}q^{\frac{1}{2}+\varepsilon}$ and $s = \omega(n^2\sqrt{\log n} \cdot \sigma)$ (resp. $s = \Omega(n^{2.5} \cdot \sigma)$). Then the construction* $\mathtt{NTRUPSF}(n, q, \sigma, s)$ *from Fig. 4 is a CRPSF secure against $\mathcal{P}oly(n)$ (resp. $2^{o(n)}$) time algorithms, assuming the hardness of* γ-Ideal-SVP *against $\mathcal{P}oly(n)$ (resp. $2^{o(n)}$) time algorithms, with $\gamma = O(n \log^{1+\varepsilon'} n \cdot s)$ (resp. $\gamma = O(n^2\sqrt{\log n} \cdot s)$).*

Proof. The sets $\mathcal{D}_n$ and $\mathcal{R}_n$ are easily recognizable. Observe that the choice of s implies $s \geq \max(\sqrt{n}, \eta_{1/2}(\mathbb{Z}^{2n}))$, so by Lemmata 2.4 and 2.7, the distribution of $\boldsymbol{z} = (z_1, z_2)$ returned by SampleDom is within statistical distance $O(2^{-n})$ of $D_{\mathbb{Z}^{2n},s}$. To show Property 2 of Definition 5.1, we apply Theorem 3.1 with $\delta = n^{-\omega(1)}$ (resp. $\delta = 2^{-\Omega(n)}$) to conclude that thanks to the choice of s, except for a fraction $\leq 2^n(q-1)^{-\varepsilon n}$ of $(a_1, a_2) \in (R_q^\times)^2$, we have $\Delta(a_1 z_1 - a_2 z_2; U(R_q)) \leq 2\delta$ with $(z_1, z_2) \hookleftarrow D_{\mathbb{Z}^{2n},s}$. Since the mapping $\phi : x \mapsto a_2^{-1}x$ is a bijection of R_q, we have $\Delta(a_1 z_1 - a_2 z_2; U(R_q)) = \Delta(a_1 a_2^{-1} z_1 - z_2; U(R_q))$ for each a_1, a_2. Moreover, since $h = a_2^{-1}a_1$ is uniformly random in $R_q^\times$ when a_1 and a_2 are independently so, we get $\Delta(hz_1 - z_2; U(R_q)) \leq 2\delta$ with $(z_1, z_2) \hookleftarrow D_{\mathbb{Z}^{2n},s}$ except for a fraction $\leq 2^n(q-1)^{-\varepsilon n}$ of $h \in R_q^\times$. Finally, by Theorem 4.2, the distribution D_h of $h = g/f$ generated by TrapGen is obtained by rejection with constant rejection probability $c < 1$ from a distribution within statistical distance $2^{3n}q^{-\lfloor \varepsilon n/2 \rfloor}$ of $U(R_q^\times)$. It follows that $\Delta(hz_1 - z_2; U(R_q)) \leq 2\delta$ with $(z_1, z_2) \hookleftarrow D_{\mathbb{Z}^{2n},s}$ except with probability $\leq \frac{1}{1-c} \cdot (2^n(q-1)^{-\varepsilon n} + 2^{3n}q^{-\lfloor \varepsilon n/2 \rfloor}) = q^{-\Omega(n)}$ over the choice of the public key h, as required.

To show Property 3 of Definition 5.1, we first observe that, for any fixed $t \in R_q$, the conditional distribution of $\boldsymbol{z} \hookleftarrow D_{\mathbb{Z}^{2n},s}$ given $f_h(\boldsymbol{z}) = hz_1 - z_2 = t$ is exactly $F(\boldsymbol{z}) = \frac{\rho_s(\boldsymbol{z})}{\rho_s(h^\perp + \boldsymbol{c})} = D_{h^\perp + \boldsymbol{c}, s}$, where $\boldsymbol{c} = (1, h - t)$ is a preimage of t under f_h. Therefore, Property 3 follows from Lemma 2.14, the bound $\|sk\| \leq 2n^{1.5}\sigma$ from Theorem 4.2, and the choice of $s = \omega(n^2\sqrt{\log n} \cdot \sigma)$ (resp. $\Omega(n^{2.5} \cdot \sigma)$).

To show Property 4 of Definition 5.1, observe that the conditional preimage distribution is $D_{h^\perp + \boldsymbol{c}, s} = D_{h^\perp, s, -\boldsymbol{c}} + \boldsymbol{c}$, where $\boldsymbol{c} = (1, h - t)$, so it suffices to lower bound the min-entropy of $D_{h^\perp, s, -\boldsymbol{c}}$. By Lemma 2.6, the latter min-entropy is $\Omega(n)$ if the condition $s \geq 2\eta_{1/2}(h^\perp)$ is satisfied. Theorem 3.1 shows that for all except a fraction $\leq 2^n(q-1)^{-\varepsilon n}$ of $\boldsymbol{a} \in (R_q^\times)^2$, we have $\eta_{1/2}(\boldsymbol{a}^\perp) \leq \sqrt{\frac{n \ln(12n)}{\pi}} q^{\frac{1}{2}+\varepsilon}$. Since $\boldsymbol{a}^\perp = h^\perp$ with $h = a_2^{-1}a_1$, it follows that for all except a fraction $\leq 2^n(q-1)^{-\varepsilon n} = q^{-\Omega(n)}$ of $h \in R_q^\times$, we have $\eta_{1/2}(h^\perp) \leq \sqrt{\frac{n \ln(12n)}{\pi}} q^{\frac{1}{2}+\varepsilon}$. By the choice of s, the condition $s \geq 2\eta_{1/2}(h^\perp)$ is satisfied. By Theorem 4.2, the condition is satisfied except with probability $\frac{q^{-\Omega(n)}}{1-c} = q^{-\Omega(n)}$ over the choice of the public key h, as required.

Finally, we show Property 5 of Definition 5.1. Let $\mathcal{A}$ be a collision-finding algorithm for NTRUPSF with run-time $T = \mathcal{P}oly(n)$ (resp. $T = 2^{o(n)}$), and success probability $\delta = 1/\mathcal{P}oly(n)$ (resp. $\delta = 2^{-o(n)}$) over the choice of the public key h and the randomness of $\mathcal{A}$. By Theorem 4.2, the success probability of $\mathcal{A}$ over the choice of $h \hookleftarrow U(R_q^\times)$ and the randomness of $\mathcal{A}$ is at least $\delta' = (1-c)\delta - 2^{3n}q^{-\lfloor \varepsilon n/2 \rfloor}$. Note that we have $\delta' = 1/\mathcal{P}oly(n)$ (resp. $\delta' = 2^{-o(n)}$). We construct an algorithm $\mathcal{A}'$ for Ideal-SIS$_{q,2,\beta}$ with $\beta = 2\sqrt{2n}s$ that works as follows on input $(a_1, a_2) \hookleftarrow U(R_q^2)$. If $(a_1, a_2) \notin (R_q^\times)^2$, it aborts. Else, $\mathcal{A}'$ runs $\mathcal{A}$ on input $h = a_2^{-1}a_1$. If $\mathcal{A}$ succeeds, it outputs $(z_1, z_2) \neq (z_1', z_2')$ with $\|(z_1, z_2)\|, \|(z_1', z_2')\| \leq \sqrt{2n}s$ such that $a_1(z_1 - z_1') + a_2(z_2' - z_2) = 0$, and then $\mathcal{A}'$ returns $\boldsymbol{w} = (z_1 - z_1', z_2' - z_2)$. Note that $0 < \|\boldsymbol{w}\| \leq 2\sqrt{2n}s$, as required. Conditioned on $(a_1, a_2) \in (R_q^\times)^2$, the distribution of h given to $\mathcal{A}$ is $U(R_q^\times)$ and thus $\mathcal{A}$ succeeds with probability $\geq \delta'$. Since $(a_1, a_2) \in (R_q^\times)^2$ with probability $\geq 1 - 2n/q = \Omega(1)$, it follows that $\mathcal{A}'$ succeeds probability $\geq (1 - 2n/q)\delta' = 1/\mathcal{P}oly(n)$ (resp. $2^{-o(n)}$). Applying Theorem 2.1 using the choice of $q = \Omega(\beta n \log^{0.5+\varepsilon'} n)$, we obtain a $\mathcal{P}oly(n)$ (resp. $2^{o(n)}$) time algorithm for γ-Ideal-SVP with the claimed γ. □

The revised NTRUSign scheme. Given the NTRUPSF construction above, the revised NTRUSign follows the GPV 'Probabilistic Full Domain Hash' construction and is shown in Fig. 5. Besides the NTRUPSF parameters, it has an additional parameter k that indicates the randomizer length. Note that the GPV signature obtained directly from NTRUPSF has signatures on a message M consisting of two 'short' ring elements (σ_1, σ_2) and a randomizer $r \in \{0,1\}^k$ satisfying $h\sigma_1 - \sigma_2 = \mathcal{H}(r, M)$, where $\mathcal{H}$ is the random oracle. To reduce signature length, our `NTRUSign` variant eliminates σ_2 from the signature, since it can be easily recovered during verification from the remaining information.

- **Key Generation** – $\mathsf{KeyGen}(1^n, q, \sigma, k)$: Run $\mathsf{TrapGen}(1^n, q, \sigma)$ of `NTRUPSF`(n, q, σ, s) to get key $h \in R_q^\times$ and trapdoor sk for function $f_h : \mathcal{D}_n \to \mathcal{R}_n$, where $\mathcal{D}_n = \{(z_1, z_2) \in R^2 : \|(z_1, z_2)\| \leq \sqrt{2n}s\}$, $\mathcal{R}_n = R_q$ and $f_h(z_1, z_2) = hz_1 - z_2$. Return the signer's public key h and secret key sk.
- **Signing Algorithm** – $\mathsf{Sign}(sk, M)$: Choose $r \hookleftarrow U(\{0,1\}^k)$, let $(\sigma_1, \sigma_2) := \mathsf{SamplePre}(sk, \mathcal{H}(r, M))$. Return (r, σ_1).
- **Verification Algorithm** – $\mathsf{Ver}(h, M, (r, \sigma_1))$: Compute $t = \mathcal{H}(r, M)$ and $\sigma_2 = h\sigma_1 - t$. Accept if $(\sigma_1, \sigma_2) \in \mathcal{D}_n$ and $r \in \{0,1\}^k$, else reject.

Fig. 5. Construction of `NTRUSign`(n, q, σ, s, k) from the `NTRUPSF` primitive in Fig. 4.

Since σ_2 is easily computed from σ_1 and the public information, the security of `NTRUSign` is equivalent to that of the GPV signature obtained from `NTRUPSF`, which in turn has been shown in [10, Prop. 6.2] to follow from the security of the underlying `NTRUPSF`. Combining with Theorem 5.2, we obtain our second main result.

Corollary 5.1. *Let $\varepsilon, \varepsilon', n, q, \sigma, s$ satisfy the conditions in Theorem 5.2, and let $k = \omega(\log n)$ (resp. $\Omega(n)$). Then, assuming the random oracle model for $\mathcal{H}$, the signature scheme* `NTRUSign`(n, q, σ, s, k) *from Fig. 5 is strongly existentially unforgeable against a chosen message attack with $\mathcal{P}oly(n)$ (resp. $2^{o(n)}$) run-time and $1/\mathcal{P}oly(n)$ (resp. $2^{-o(n)}$) success probability, assuming the hardness of* γ-Ideal-SVP *against $\mathcal{P}oly(n)$ (resp. $2^{o(n)}$) time algorithms, with $\gamma = O(n \log^{1+\varepsilon'} n \cdot s)$ (resp. $\gamma = O(n^2\sqrt{\log n} \cdot s)$).*

Note that if $\mathcal{H}$ runs in quasi-linear time, then so does the verification algorithm. Also, if pre-computations are performed, then so does the signing algorithm (see [40]). The amortized cost per signed bit is in both cases $\widetilde{O}(1)$. Finally, we remark that the smallest q that can be chosen in Theorem 5.2 and Corollary 5.1 is $\widetilde{\Omega}(n^9)$ (resp. $\widetilde{\Omega}(n^{10})$) for polynomially (resp. subexponentially) bounded attacks, and the smallest γ that can be obtained is $\widetilde{O}(n^{8.5})$ (resp. $\widetilde{O}(n^{10.5})$).

6 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 Ideal-SIS problem is known to be hard as soon as Φ_n is irreducible over the rationals, has small height and contains few coefficients (see [24]). The R-LWE problem is known to be hard when Φ_n is a cyclotomic polynomial (see [27]). We chose to restrict ourselves to cyclotomic polynomials of order a power of 2 for the sake of simplicity: 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 cyclotomic rings than those we considered. An interesting choice could be the cyclotomic polynomials of prime order (i.e., $\Phi_n = (x^n - 1)/(x - 1)$ with n prime) as the corresponding rings are large subrings of the NTRU rings (and one might then be able to show that the hardness carries over to the NTRU rings).

The modified NTRUSign can be shown hard to break for classical computers, in the random oracle model (assuming the worst-case hardness of standard lattice problems for ideal lattices). Because of the use of the random oracle, it is unclear whether this proof remains meaningful in the case of quantum attackers. As pointed out in [6], one should be extremely cautious with the random oracle in a quantum setup. Similarly, since the security of NAEP (the CCA-secure variant of NTRUEncrypt) relies on the random oracle (see [20]) and since the reduction from standard problems over ideal lattices to R-LWE is quantum, the security of NAEP remains open (both quantumly and classically).

Finally, 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 left as a future work.

Acknowledgements. We thank Guillaume Hanrot and Vadim Lyubashevsky, Chris Peikert, Oded Regev, Igor Shparlinski, Joseph Silverman and Frederik Vercauteren for helpful discussions. The first author was partly supported by the LaRedA ANR grant. The second author was supported by an Australian Research Fellowship (ARF) from the Australian Research Council under Discovery Grant DP0987734, and partly by a Macquarie University Research Fellowship.

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. 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.
5. D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *Proc. of Eurocrypt*, volume 1233 of *LNCS*, pages 52–61. Springer, 1997.
6. Ö. Dagdelen, M. Fischlin, A. Lehmann, and C. Schaffner. Random oracles in a quantum world. Cryptology ePrint Archive, Report 2010/428, 2010. http://eprint.iacr.org/2010/428.
7. J. von zur Gathen and J. Gerhardt. *Modern Computer Algebra, 2nd edition*. Cambridge University Press, 2003.
8. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proc. of STOC*, pages 169–178. ACM, 2009.
9. C. Gentry, J. Jonsson, J. Stern, and M. Szydlo. Cryptanalysis of the NTRU signature scheme (NSS) from Eurocrypt 2001. In *Proc. of Asiacrypt*, volume 2248 of *LNCS*, pages 1–20. Springer, 2001.
10. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008.
11. C. Gentry and M. Szydlo. Cryptanalysis of the revised NTRU signature scheme. In *Proc. of Eurocrypt*, volume 2332 of *LNCS*, pages 299–320. Springer, 2002.
12. N. Higham. *Accuracy and Stability of Numerical Algorithms, 2nd edition*. SIAM, 2002.
13. J. Hoffstein, N. Howgrave-Graham, J. Piphe, and W. Whyte. Practical lattice-based cryptography: NTRUEncrypt and ntrusign, 2009. Chapter of [37].
14. J. Hoffstein, N. A. Howgrave-Graham, J. Pipher, J. H. Silverman, and W. Whyte. NTRUSIGN: Digital signatures using the NTRU lattice, preliminary draft 2, dated april 2, 2002. Preliminary/extended version of [15]. Available at http://www.securityinnovation.com/cryptolab/articles.shtml.
15. J. Hoffstein, N. A. Howgrave-Graham, J. Pipher, J. H. Silverman, and W. Whyte. NTRUSIGN: Digital signatures using the NTRU lattice. In *Proc. of CT-RSA*, volume 2612 of *LNCS*. Springer, 2003.
16. 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.
17. 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.
18. 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.

19. N. Howgrave-Graham, P. Q. Nguyen, D. Pointcheval, J. Proos, J. H. Silverman, A. Singer, and W. Whyte. The impact of decryption failures on the security of NTRU encryption. In *Proc. of CRYPTO*, volume 2729 of *LNCS*, pages 226–246. Springer, 2003.
20. N. Howgrave-Graham, J. H. Silverman, A. Singer, and W. Whyte. NAEP: Provable security in the presence of decryption failures. Technical report, Cryptology ePrint Archive, 2003. `http://eprint.iacr.org/2003/172`.
21. IEEE P1363. Standard specifications for public-key cryptography. `http://grouper.ieee.org/groups/1363/`.
22. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann*, 261:515–534, 1982.
23. V. Lyubashevsky. Fiat-Shamir with aborts: Applications to lattice and factoring-based signatures. In *Proc. of Asiacrypt*, volume 5912 of *LNCS*, pages 598–616. Springer, 2009.
24. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. ICALP (2)*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
25. 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.
26. V. Lyubashevsky, C. Peikert, and O. Regev. FFT-based regularity for Ring-LWE, 2010. Personal communication.
27. 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.
28. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007.
29. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Kluwer Academic Press, 2002.
30. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007.
31. 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.
32. 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.
33. S. Min, G. Yamamoto, and K. Kim. Weak property of malleability in NTRUSign. In *Proc. of ACISP*, volume 3108 of *LNCS*, pages 379–390. Springer, 2004.
34. R. A. Mollin. *Algebraic Number Theory*. Chapman and Hall/CRC Press, 1999.
35. J. Neukirch. *Algebraic number theory*, volume 322 of *Grundlehren der Mathematischen Wissenschaften*. Springer, Berlin, 1999.
36. P. Q. Nguyen and O. Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *Journal of Cryptology*, 22(2):139–160, 2009.
37. P. Q. Nguyen and B. Vallée (editors). *The LLL Algorithm: Survey and Applications*. Information Security and Cryptography. Springer, 2009. Published after the LLL25 conference held in Caen in June 2007, in honour of the 25-th anniversary of the LLL algorithm.
38. C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Comput. Complexity*, 2(17):300–351, 2008.
39. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
40. C. Peikert. An efficient and parallel gaussian sampler for lattices. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 80–97. Springer, 2010.
41. C. Peikert. On ideal lattices and learning with errors over rings, 2010. Talk given at Eurocrypt'10, corresponding to [27], and available at `http://www.cc.gatech.edu/~cpeikert/`.
42. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proceedings of the 2006 Theory of Cryptography Conference (TCC)*, pages 145–166, 2006.
43. C. Peikert and A. Rosen. Lattices that admit logarithmic worst-case to average-case connection factors. In *Proc. of STOC*, pages 478–487. ACM, 2007.
44. R. A. Perlner and D. A. Cooper. Quantum resistant public key cryptography: a survey. In *Proc. of IDtrust*, pages 85–93. ACM, 2009.
45. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93. ACM, 2005.
46. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
47. O. Regev. The learning with errors problem, 2010. Invited survey in CCC 2010, available at `http://www.cs.tau.ac.il/~odedr/`.
48. C. P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theor. Comput. Science*, 53:201–224, 1987.

49. B. D. Sittinger. The probability that random algebraic integers are relatively r-prime. *Journal of Number Theory*, 130:164–171, 2010.
50. 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.
51. M. Szydlo. Hypercubic lattice reduction and analysis of GGH and NTRU signatures. In *Proc. of Eurocrypt*, volume 2656 of *LNCS*, pages 433–448. Springer, 2003.
52. G. Tenenbaum. *Introduction to Analytic and Probabilistic Number Theory*. Number 46 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1995.
53. T. Xylouris. On Linnik's constant, 2009. Available at `http://arxiv.org/abs/0906.2749` (in German).