# Making NTRUEncrypt and NTRUSign as Secure as Standard Worst-Case Problems over Ideal Lattices*

Damien Stehlé[1] and Ron Steinfeld[2]

[1] ENS de Lyon, Laboratoire LIP (U. Lyon, CNRS, ENS Lyon, INRIA, UCBL),
46 Allée d'Italie, 69364 Lyon Cedex 07, France.
damien.stehle@ens-lyon.fr – http://perso.ens-lyon.fr/damien.stehle
[2] Clayton School of Information Technology,
Monash University, Clayton, Australia.
ron.steinfeld@monash.edu – http://users.monash.edu.au/~rste/

**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 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 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 from discrete Gaussians, then the public key, which is their ratio, is statistically indistinguishable from uniform over its range. We also show how to rigorously extend the encryption secret key into a signature secret key. The security then follows from the already proven hardness of the R-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 [27]. 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 [27, 11]. At the ANTS'98 conference, the NTRU authors gave an improved presentation including a thorough assessment of its practical security against lattice attacks [28]. We refer to [24] for an up-to-date account on the past 15 years of security and performance analyses. Nowadays, NTRUEncrypt is commonly 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 [33]. It is also often considered as the most viable post-quantum public-key encryption (see, e.g., [58]). The authors of NTRUEncrypt also proposed a signature scheme based on a similar design. The history of NTRUSign started with NSS in 2001 [29]. Its development has been significantly more hectic and controversial, with a series of cryptanalyses and repairs (see, e.g., [20, 22, 31, 67, 49, 52] and the survey [24]).

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.

* Some of the results in this paper have been presented in preliminary form at Eurocrypt 2011 [64]. The results in this paper improve and significantly extend those in [64]; the most significant addition is the security analysis of a provably secure variant of NTRUSign.

It originated in 1996 with Ajtai's acclaimed worst-case to average-case reduction [3], leading to a collision-resistant hash function that is as hard to break as solving several natural worst-case 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 [59, 60]: LWE is both hard on the average (standard 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 [60, 54, 21, 8, 1], among others, and the surveys [47, 61]).

The main drawback of cryptography based on LWE and SIS lies in its limited efficiency. A key typically contains a random matrix over the ring $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ for a small $q$, whose dimensions are (at least) 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 [44] succeeded in restricting SIS to structured matrices while preserving a worst-case to average-case reduction. The worst-case problem is a restriction of a standard lattice problem to the specific family of cyclic lattices. The structure of Micciancio's matrices allows for an interpretation in terms of arithmetic in the ring $\mathbb{Z}_q[x]/(x^n - 1)$, where $n$ is the dimension of the worst-case lattices and $q$ is a small prime. Micciancio's construction leads to a family of pre-image resistant hash functions, with complexity quasi-linear in the security parameter $n$: The efficiency gain stems from the use of the discrete Fourier transform for multiplying polynomials. In two concurrent works, Peikert, Rosen, Lyubashevsky and Micciancio [57, 39] later suggested to change the ring to $\mathbb{Z}_q[x]/\Phi$ with a polynomial $\Phi$ 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, now often called the *Ideal Small Integer Solution* or *Ring Small Integer Solution* problem (R-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 [38] introduced an efficient digital signature provably as secure as R-SIS (in the random oracle model). Also in 2009, Stehlé, Steinfeld, Tanaka and Xagawa [65] introduced a structured (albeit somewhat restricted) variant of LWE, which they proved as hard as R-SIS (under a quantum reduction), and allowed for the design of an asymptotically efficient CPA-secure encryption scheme. In an independent and concurrent work, Lyubashevsky et al. [21] proposed a ring variant of LWE, called R-LWE, whose great flexibility allows for more natural (and efficient) cryptographic constructions.

**Our results.** The high efficiency and industrial standardization of `NTRUEncrypt` and `NTRUSign` strongly motivate a theoretically founded study of their security. Indeed, in the absence of such a study so far, their security has remained in doubt over the last 15 years since the initial NTRU publication. This work addresses this problem.

We propose a mild modification of `NTRUEncrypt` that is CPA-secure in the standard model, under the assumed quantum hardness of standard worst-case problems over ideal lattices (for $\Phi = x^n + 1$ with $n$ a power of 2); and we describe a variant of `NTRUSign` that is existentially unforgeable in the random oracle model, under the assumed classical hardness of the same problems over ideal lattices. The `NTRUEncrypt` modifications are summarized at the end of the introduction. The most substantial additional modification for `NTRUSign` is the use of a discrete Gaussian sampler [21, 55, 13] in the signing process, that ensures that no secret information is leaked while signing (thus preventing the learning attack from [52]). We also give the first rigorous analysis of the algorithm that extends an `NTRUEncrypt` secret key into an `NTRUSign` secret key.

We stress that our main goal in this paper is to provide, for the first time, a firm theoretical grounding for the security of the NTRU schemes, in the asymptotic sense. The practical instantiations of our schemes are likely to be significantly less efficient than the original schemes. However, several of our modifications incur negligible performance overheads over the original schemes, while bringing their security level closer to the provably secure schemes. For instance, the extra error term we add to the `NTRUEncrypt` scheme is a cheap way to address the lack of IND-CPA security of the original scheme.

**Overview of our techniques.** Our main technical contribution is the modification and analysis of the NTRU 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 the ring $\mathbb{Z}_q[x]/(x^n - 1)$ (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 R-SIS and R-LWE (we actually show a weaker distribution property, which still suffices for linking the security to R-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 may be of independent interest, is a new regularity result for the ring $R_q := \mathbb{Z}_q[x]/(x^n + 1)$ when the polynomial $x^n + 1$ with $n$ a power of 2 has $n$ factors modulo prime $q$: Given $a_1, \ldots, a_m$ uniform in $R_q$, we would like $\sum_{i \leq m} s_i a_i$ to be within exponentially small statistical distance to uniformity, with small random $s_i$'s and small $m$. Micciancio's regularity bound [44, Se. 4.1] (see also [65, 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 has been concurrently and independently obtained by Lyubashevsky et al. in [43]. An additional difficulty in the proof of public-key uniformity, which we handle via an inclusion-exclusion argument, is that we need the randomizers $s_i$ 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 [26, Se. 4] and in [25, 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 (most notably the co-primality of the two secret key polynomials over $\mathbb{Z}[x]/(x^n - 1)$), 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. Furthermore, the security of the signature scheme follows from the hardness of R-SIS, even with this additional rejection.

Finally, the cryptographic schemes are obtained from (structured variants of) the Gentry et al. [21] signature and dual encryption schemes, via an *inversion-based dimension reduction* of the R-SIS/R-LWE instances. We explain it in the case of R-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 $\sum_i s_i a_i = 0 \bmod q$. 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 R-SIS when $m = 2$.

**Comparison between `NTRUEncrypt` and its provably secure variant.** Let $R_{\text{NTRU}}$ denote 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, choose $p \in R_{\text{NTRU}}$ with small coefficients, co-prime with $q$ and such that the plaintext space $R_{\text{NTRU}}/p$ 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 $(f, g) \in R_{\text{NTRU}}^2$ that are sampled randomly with large prescribed proportions of zero coefficients, and with their other coefficients belonging to $\{-1, 1\}$. For improved decryption efficiency, one may choose $f$ as $f = 1 + p\overline{f}$ with $\overline{f}$ as described just above, so that $f = 1 \bmod p$. With high probability, we (heuristically) expect the polynomial $f$ to be 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 $pgs + fM$ (seen as an element of $R_{\text{NTRU}}$). The latter requires that $q$ is not too small.
- 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 decryption errors from occurring, by setting the parameters carefully.

In order to achieve IND-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 `NTRUEncrypt` scheme (which preserve its quasi-linear computation and space complexity):

1. We replace $R_{\text{NTRU}}$ 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 [41]). This also allows us to take $p = 2$.
3. We sample $f$ and $g$ from discrete Gaussians over $R$, 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 $f = p\overline{f} + 1$ with $\overline{f}$ 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 [5, Se. 3.1]).

These modifications may be expensive to implement in practice, because of the hidden constant factor overheads. However, they suggest several computationally inexpensive modifications to the original NTRUEncrypt design that bring it closer to the provably secure variant. The addition of a noise component $e$ in the encryption function (Modification 4) does not require a large increase of $q$ for ensuring decryption correctness, but allows thwarting a simple Chosen Plaintext Attack based on the following observation: If $C$ is an encryption of $M$ in the original NTRUEncrypt scheme, then the ring element $(C - M)/h \bmod q$ has small coefficients. Modification 3 is much more expensive to implement, as our analysis requires the standard deviation to be quite large, leading to secret key polynomials $f$ and $g$ with much bigger coefficients than in the original scheme. Then the modulus $q$ needs being significantly increased in order to enable decryption correctness. However, this modification may hint that taking $f$ and $g$ a little less small than in the original design may increase security. This would for example thwart the so-called hybrid attack on NTRU [30] and allow using a smaller $n$. A drawback of taking non-sparse polynomials $f$ and $g$ is that multiplications by $f$ and $g$ would become more costly. An alternative, suggested by Modification 2, is to take a modulus $q$ so that $x^n \pm 1$ has $n$ distinct linear factors modulo $q$: In that setup, the ring $R/q$ admits a natural and efficient Fast Fourier Transform. Finally, Modification 1 suggests replacing $x^n - 1$ by $x^n + 1$. The former has been shown insecure in the context of hash functions [56, Se. 4.1], although we actually do not know of any such attack in the context of NTRU.

**Related works.** Like NTRUEncrypt, Gentry's somewhat homomorphic scheme [18] also has ciphertexts consisting of a single ring element. It also admits a security proof under the assumed quantum hardness of standard worst-case problems over ideal lattices [19]. Our security analysis for the modified NTRUEncrypt scheme allows encrypting and decrypting $\Omega(n)$ plaintext bits for $\widetilde{O}(n)$ bit operations, while achieving security against $2^{g(n)}$-time attacks, for any $g(n) \leq o(n)$, assuming the worst-case hardness of $\mathcal{P}oly(n)$-Ideal-SVP against $2^{O(g(n))}$-time quantum algorithms. The latter assumption is believed to be valid for any $g(n) = o(n)$. Gentry's analysis from [19, 17] can be generalized to handle $2^{g(n)}$-time attacks while encrypting and decrypting $O(g(n))$ plaintext bits for $\widetilde{O}(n)$ bit operations, under the assumed hardness of $2^{\Omega(g(n))}$-Ideal-SVP against $2^{O(g(n))}$-time quantum algorithms. The latter assumption is known to be invalid when $g(n) = \widetilde{\Omega}(\sqrt{n})$ (using [62]), thus limiting the attacker's strength the analysis can handle. On the other hand, Gentry's scheme allows homomorphic additions and multiplications, whereas ours seems restricted to homomorphic additions.

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 does not follow immediately whether this proof remains meaningful in the case of quantum attackers. As pointed out in [7], one should be extremely cautious with the random oracle in a quantum setup. Fortunately, the security proof for our NTRUSign scheme falls in the class of 'history-free' reductions as defined in [7] and shown to imply security in the quantum-accessible random oracle model.

Similarly, the security of NAEP (the CCA-secure variant of NTRUEncrypt) relies on the random oracle (see [32]). Since the reduction from standard problems over ideal lattices to R-LWE is quantum, the security of NAEP remains open, both quantumly and classically.

We also mention a couple of works building upon some of the results of this paper, since its publication in a preliminary form in [64]. In [66], it is shown how to adapt the NTRUSign trapdoor key generation algorithm from the present paper to construct an NTRU-based lossy trapdoor function and use it to upgrade the IND-CPA security of the NTRUEncrypt scheme to chosen-ciphertext security (IND-CCA2) in the standard model, while preserving the same asymptotic efficiency, up to constant factors. An extension in another direction is given in [37], which shows how to modify our `NTRUEncrypt` variant to achieve a fully-homomorphic multi-key encryption scheme. For this, the scheme in [37] requires the secret key coefficients to be much smaller than the $O(\mathcal{P}oly(n) \cdot q^{1/2})$ value needed for our statistical uniformity bounds in this paper. The security of the scheme in [37] relies also, besides the hardness of R-LWE, on the assumed computational indistinguishability of the resulting public key from uniformity.

**Open problems.** Our study is restricted to the sequence of rings $\mathbb{Z}[x]/\Phi_n$ where $\Phi_n = x^n + 1$ with $n$ a power of 2. An obvious drawback is that this does not allow for much flexibility on the choice of $n$ (in the case of NTRU, the degree was assumed prime, which provides more freedom). The R-LWE problem is known to be hard when $\Phi_n$ is cyclotomic [41] (for an appropriate choice of modulus $q$). The R-SIS problem is known to be hard under even milder conditions on $\Phi_n$ (see [39, 56]). We chose to restrict ourselves to cyclotomic polynomials of order a power of 2 because it makes the description of the schemes simpler to follow. Our results are likely to hold for more general rings than those we considered. An interesting choice could be the cyclotomic rings of prime order (i.e., $\Phi_n = (x^n - 1)/(x - 1)$ with $n$ prime) as these are large subrings of the original NTRU rings and one might then be able to show that the hardness carries over to the NTRU rings.

Reducing the constant factor overheads of our provably secure schemes with respect to the original NTRU schemes, while preserving a proof with respect to standard problems, is a remaining interesting challenge. A related open question with additional applications (see [37]) is to prove the computational indistinguishability of the NTRU public key with secret key coefficients significantly smaller than $q^{1/2}$, assuming the hardness of a standard problem, such as R-LWE.

**Road-map.** In Section 2, we provide the necessary background material in elementary algebraic number theory and on the R-LWE and R-SIS problems. Section 3 is devoted to the description and security proof of the modified encryption scheme. Finally, we consider `NTRUSign` in Section 4.

**Notation.** If $q$ is a non-zero integer, we let $\mathbb{Z}_q$ denote the ring of integers modulo $q$, i.e., the set $\{0, \ldots, q-1\}$ with addition and multiplication modulo $q$. For a ring $(R, +, \times)$, we let $R^\times$ denote the set of invertible elements of $R$. If $q$ is a prime power, we let $\mathbb{F}_q$ denote the finite field with $q$ elements. If $z \in \mathbb{C}$, its real and imaginary parts will be denoted by $\Re(z)$ and $\Im(z)$ respectively. Vectors will be denoted in bold. If $\boldsymbol{x} \in \mathbb{R}^n$, then $\|\boldsymbol{x}\|$ denotes the Euclidean norm of $\boldsymbol{x}$. The inner product of two vectors $\boldsymbol{x}$ and $\boldsymbol{y}$ will be denoted by $\langle \boldsymbol{x}, \boldsymbol{y} \rangle$. We use ln to denote the natural logarithm.

The standard $n$-dimensional Gaussian function (resp. distribution) with center $\mathbf{0}$ and variance $\sigma$, will be denoted by $\rho_\sigma(\boldsymbol{x})$ (resp. $\nu_\sigma$), i.e., $\rho_\sigma(\boldsymbol{x}) = \exp(-\pi\|\boldsymbol{x}\|^2/\sigma^2)$ (resp. $\nu_\sigma(\boldsymbol{x}) = \rho_\sigma(\boldsymbol{x})/\sigma^n$). If $E$ is a finite set, we let $U(E)$ denote the uniform distribution over $E$. 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 $f(F)$. If $D_1$ and $D_2$ are two probability distributions over a discrete domain $E$, their statistical distance is $\varDelta(D_1; D_2) = \frac{1}{2}\sum_{x \in E} |D_1(x) - D_2(x)|$. We write $z \hookleftarrow D$ when the random variable $z$ is sampled from the distribution $D$.

We make use of the Landau notations $O(\cdot), \widetilde{O}(\cdot), o(\cdot), \omega(\cdot), \varOmega(\cdot), \widetilde{\varOmega}(\cdot), \varTheta(\cdot)$. A function $f(n)$ is said negligible if $f(n) = n^{-\omega(1)}$. 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 $f$.

## 2 Reminders on Euclidean lattices and in algebraic number theory

We refer to [45] and [4, 50, 51] for introductions to the computational aspects of lattices and to algebraic number theory respectively, and to [47, 61] 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 $L$, denoted 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 any $k \leq n$ as the smallest $r$ such that $L$ contains at least $k$ linearly independent vectors of norm $\leq r$. The *dual lattice* of $L$ is defined as $\widehat{L} = \{\boldsymbol{c} \in \mathbb{R}^n : \forall i, \langle \boldsymbol{c}, \boldsymbol{b}_i \rangle \in \mathbb{Z}\}$.

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 $\sigma$ 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 $\mathbf{0}$. 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 use the following results.

**Lemma 2.1 ([53, Le. 3.5],[46, Le. 3.3]).** *For any full-rank lattice $L \subseteq \mathbb{R}^n$ and $\delta \in (0,1)$, we have $\eta_\delta(L) \leq \sqrt{\frac{\ln(2n(1+1/\delta))}{\pi}} \cdot \min\left(\lambda_n(L), 1/\lambda_1^\infty(\widehat{L})\right)$.*

**Lemma 2.2 ([46, 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 2.3 ([46, 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 $\Pr_{\boldsymbol{b} \hookleftarrow D_{L,\sigma,\boldsymbol{c}}}[\|\boldsymbol{b}\| \geq \sigma\sqrt{n}] \leq \frac{1+\delta}{1-\delta} \cdot 2^{-n}$.*

**Lemma 2.4 ([21, 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 2.5 ([56, 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 2.6 ([21, 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 $\sigma = \omega(\sqrt{\ln n}) \max \|\boldsymbol{b}_i\|$, and returns samples from a distribution whose statistical distance to $D_{L,\sigma}$ is negligible with respect to $n$.*

We will need the following result on one-dimensional projections of discrete Gaussians. Other results on these projections are known (see [46, Le. 4.2] and [53, Cor. 5.3]), but do not seem to suffice for our needs. The second half of Lemma 2.7 below is akin to [53, Cor. 5.3], but, to the extent of our knowledge, the first half is new.

**Lemma 2.7.** *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:*

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

*Proof.* Let $U$ be an orthonormal matrix whose first row is $\boldsymbol{u}^T$. 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 2.2, 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 $K := \frac{1}{2} - \frac{\pi}{t^2} \in [0, 1/2]$. 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 $\det(DL') = \sqrt{1+Kt^2/\pi} \cdot \det(L')$. Lemma 2.2 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:

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

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 2.2 once more provides the result. □

### 2.2 Algebraic number theory and lattices

**Ideal lattices.** Let $\Phi \in \mathbb{Z}[x]$ be 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 $R$. For elements $r_1 \ldots, r_k$ of $R$, we let $\langle r_1, \ldots, r_k\rangle$ denote the minimal ideal of $R$ containing these elements, and we say that $r_1, \ldots, r_k$ *generate* this ideal. 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 $\Phi$ 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 $\Phi$-ideal lattice. For $v \in R$ we let $\|v\|$ denote its Euclidean norm (as a vector).

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 this $\Phi$ corresponds 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 let $(\sigma_i)_{i\le n}$ denote the canonical complex embeddings: We can choose $\sigma_i : P \mapsto P(\zeta^{2i+1})$ for $i \le n$. For any $\alpha$ in $\mathbb{Q}[\zeta]$, we define its $T_2$-norm by $T_2(\alpha)^2 = \sum_{i\le n} |\sigma_i(\alpha)|^2$ and its algebraic norm by $\mathcal{N}(\alpha) = \prod_{i\le n} |\sigma_i(\alpha)|$. The arithmetic-geometric inequality gives $\mathcal{N}(\alpha)^{2/n} \le \frac{1}{n} T_2(\alpha)^2$. Also, for the specific cyclotomic fields we are considering, the polynomial norm (the norm of the coefficient vector of $\alpha$ 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 $\alpha$. For the sake of simplicity, we will try to use the polynomial terminology wherever possible (and we refer to [41, 43] for a more mathematical exposition).

The following result is a consequence of Lemma 2.7.

**Lemma 2.8.** *For any non-zero ideal lattice $I \subseteq R$, $c \in K$, $\delta \in (0,1)$, $t \ge \sqrt{2\pi}$, $u \in K$ and $\sigma \ge \eta_\delta(I)$, we have*

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

*Proof.* A coefficient of $(b-c) \times u \in R$ can be viewed as an inner product between the coefficient vectors of $b - c$ and of some $u'$ obtained by permuting the coefficients of $u$ and multiplying them by $\pm 1$. Therefore, by Lemma 2.7, the magnitude of each coefficient of $(b-c) \times u$ is $\ge t\sigma\|u'\|$ with probability $\le \frac{1+\delta}{1-\delta} t\sqrt{2\pi\mathrm{e}} \cdot \mathrm{e}^{-\pi t^2}$. The equality $\|u'\| = \|u\|$ and the union bound imply that all the magnitudes of the coefficients are $\le t\sigma\|u\|$ with probability $\ge 1 - \frac{1+\delta}{1-\delta} nt\sqrt{2\pi\mathrm{e}} \cdot \mathrm{e}^{-\pi t^2}$. If that is the case, then $\|(b-c) \times u\| \le t\sigma\|u\|\sqrt{n}$, which completes the proof. □

**On the reduction of the ring modulo $q$.** Let $q$ be a prime integer and $R_q := R/qR = \mathbb{Z}_q[x]/\Phi$. Because of the choice of $\Phi = x^n + 1$ with $n$ a power of 2, the factorisation of $\Phi$ modulo $q$ is always of the form $\Phi = \prod_{i\le k_q} \Phi_i$, where all the $\Phi_i$'s are irreducible modulo $q$ and share the same degree $d_q = n/k_q$. The number of factors $k_q$ is a power of 2 that can range from 2 (if $q = 3 \bmod 8$) to $n$ (if $q = 1 \bmod 2n$). The Chinese Remainder Theorem provides a ring isomorphism between $R_q$ and $(\mathbb{F}_{q^{d_q}})^{k_q}$:

$$a \ \mapsto \ \left(a \bmod \Phi_1, \ldots, a \bmod \Phi_{k_q}\right).$$

Both extreme situations can prove interesting. Choosing $q$ such that $\Phi$ has exactly $n$ distinct linear factors modulo $q$ allows for faster implementations, as the ring $R_q$ then admits a natural FFT: Multiplication of elements of $R_q$ can be performed within $O(n \ln n)$ additions and multiplications in $\mathbb{F}_q$ (see [16, Ch. 8], [40, Se. 2.1]). Oppositely, choosing $q$ such that $\Phi$ has only two irreducible factors modulo $q$ makes the ring $R_q$ behave very similarly to a field (it has very few zero divisors). For example, this choice allows for proving statistical uniformity of the revised NTRU public key for smaller values of $q$, and to have the security of the schemes rely on weaker assumptions. For both choices of $q$, Dirichlet's theorem on arithmetic progressions implies that infinitely such primes exist. Furthermore, Linnik's theorem asserts that the smallest such prime is $\le \mathcal{P}oly(n)$. For our particular choice of $n$ (a power of 2), the smallest such primes are known to be $O(n^{2.5})$, and, after some $\mathcal{P}oly(n)$ threshold, these primes are quite frequent (see [34, 15]).

**Module $q$-ary lattices.** We call an $m$-dimensional lattice that contains $q\mathbb{Z}^m$ a $q$-ary lattice.

An $R$-module is a set of the form $M = \sum_{i\le d} R\boldsymbol{b}_i \subseteq K^m$. If the $\boldsymbol{b}_i$'s are $K$-linearly independent, we call them an $R$-basis of $M$. Note that contrarily to lattices, some $R$-modules may not admit an $R$-basis (we refer the reader to [10, Ch. 1] and [14] for alternative compact representations). 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. Since these lattices are $q$-ary, we call them module $q$-ary lattices.

In [55], Peikert described a significantly faster algorithm than the discrete Gaussian sampler from [21], 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 2.9 (Adapted from [55]).** *There exists a $\widetilde{O}(nm)$-time off-line/on-line algorithm that takes as input an $R$-basis $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m$ of a module $q$-ary lattice $L \subseteq R^m$, with $q = \mathcal{P}oly(n)$, $\boldsymbol{c} \in \mathbb{Q}^{mn}$ and $\sigma = \omega(\sqrt{mn \ln n}) \max \|\boldsymbol{b}_i\|$, and returns samples from a distribution whose statistical distance to $D_{L,\sigma,\boldsymbol{c}}$ is negligible with respect to $n$. The complexity bound holds assuming pre-computations (off-line) are performed using $q$, $\sigma$ and $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m$, but not $\boldsymbol{c}$.*

Recently, Ducas and Nguyen [13] showed how to perform the pre-computations of Lemma 2.9 in expected time $\widetilde{O}(mn)$.

### 2.3 Computational problems

**The Shortest Vector Problem.** The most famous algorithmic problem on lattices 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 $\gamma$-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. $\gamma$-Ideal-SVP), which is implicitly parameterized by a sequence of polynomials $\Phi$ of growing degrees. No algorithm is known to perform non-negligibly better for ($\gamma$-)Ideal-SVP than for ($\gamma$-)SVP. It is believed that no subexponential quantum algorithm solves the computational variants of $\gamma$-SVP or $\gamma$-Ideal-SVP in the worst case, for any $\gamma$ that is polynomial in the dimension. The smallest $\gamma$ which is known to be achievable in polynomial time is exponential, up to poly-logarithmic factors in the exponent [36, 62, 48].

**The Small Integer Solution problem over Rings.** R-SIS was introduced in [39, 56], as an average-case variant of $\gamma$-SVP in module $q$-ary lattices.

**Definition 2.1.** *The Ring Small Integer Solution problem with parameters $q, m, \beta$ and $\Phi$ (R-SIS$^{\Phi}_{q,m,\beta}$) is as follows: Given $m$ 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 R-SIS is related to the worst-case hardness of Ideal-SVP, as follows. The result is adapted from [39], using tools from [41].

**Theorem 2.1 (Adapted from [39]).** *Let $n = 2^k$, $\Phi = x^n + 1$ and $\varepsilon > 0$. Let $m, q > 0$ such that $q \geq \beta\sqrt{n} \cdot \omega(\ln n)$ and $m, \ln q \leq \mathcal{P}oly(n)$. A polynomial-time algorithm solving R-SIS$^{\Phi}_{q,m,\beta}$ with non-negligible probability can be used to solve $\gamma$-Ideal-SVP in polynomial-time with $\gamma \geq \beta\sqrt{n} \cdot \omega(\sqrt{\ln n})$.*

**The Learning With Errors problem over Rings.** For $s \in R_q$ and $\psi$ 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 $\psi$. The Ring Learning With Errors problem (R-LWE) was introduced by Lyubashevsky et al. in [41] and shown hard for specific error distributions $\psi$ closely related to Gaussians.

**Definition 2.2.** *Let $\Gamma$ be a distribution over a family of distributions on $R$. The Ring Learning With Errors Problem with parameters $q, \Gamma$ and $\Phi$ ($\text{R-LWE}^{\Phi}_{q,\Gamma}$) is as follows. Let $\psi$ be sampled from $\Gamma$ 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 non-negligible over the randomness of the input, the randomness of the samples and the internal randomness of the algorithm.*

Note that this definition differs from the one of [41] in the following respects: We use the polynomial representation (which is handled by applying the complex FFT to the error term); we use $R_q$ rather than $R_q^*$ (for our choice of $\Phi$, we have $R_q^* = \frac{1}{n} R_q$); and the noise distributions are discrete.

R-LWE can be interpreted as a problem over module $q$-ary 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 assessing whether the vector $\boldsymbol{b}$ is generated uniformly modulo the (module) lattice $L(\boldsymbol{a})$ or around the origin according to some Gaussian-like distribution and then reduced modulo the lattice.

**Theorem 2.2 (Adapted from [41]).** *Assume that $\alpha q = \omega(n\sqrt{\ln n})$ with $\alpha \in (0,1)$ and $q = \mathcal{P}oly(n)$ prime with $q = 1 \bmod 2n$. Consider the distribution $\overline{\Gamma}_\alpha$ defined below in this section. There exists a randomized polynomial-time quantum reduction from $\gamma$-Ideal-SVP to $\text{R-LWE}_{q,\overline{\Gamma}_\alpha}$, denoted by $\text{R-LWE}_{q,\alpha}$ in the sequel, with $\gamma = \omega(n^{1.5} \ln n)/\alpha$.*

**Variants of R-LWE.** For $s \in R_q$ and $\psi$ 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 $\psi$. When $q = \Omega(n)$, the probability for a uniform element of $R_q$ of being invertible is non-negligible, and thus R-LWE remains hard even when $A_{s,\psi}$ and $U(R_q \times R_q)$ are respectively replaced by $A^{\times}_{s,\psi}$ and $U(R_q^{\times} \times R_q)$. We call $\text{R-LWE}^{\times}$ the latter variant.

Furthermore, as explained in [5, Le. 2], the nonce $s$ can also be chosen from the error distribution without incurring any security reduction. We call $\text{R-LWE}^{\times}_{\text{HNF}}$ the corresponding modification of R-LWE. We recall the argument, for completeness. Assume an algorithm $\mathcal{A}$ can solve $\text{R-LWE}^{\times}_{\text{HNF}}$. We use $\mathcal{A}$ to solve $\text{R-LWE}^{\times}$. 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 $A^{\times}_{s,\psi}$ to $A^{\times}_{-e_1,\psi}$, and $U(R_q^{\times} \times R_q)$ to itself.

We remark that a simpler variant of R-LWE with fixed number of samples and fixed spherical noise distribution is proven hard in [42]. However, we chose not to use this simpler variant in this work since its proven hardness involves a larger Ideal-SVP approximation factor $\gamma$ than the variant of R-LWE considered in the theorem above. The simplified variant offers a different trade-off between the underlying hardness assumption and the cost of sampling noise vectors.

**Noise definition and noise generation.** We now describe the distribution $\overline{\Gamma}_\alpha$. It is somewhat tedious 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. Alternative R-LWE noise generation algorithms are described in [43, 12].

For $\boldsymbol{\sigma} \in \mathbb{R}^n$ with positive coordinates, we define the elliptical 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 [41], 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 $V \in \mathbb{C}^{n \times n}$ with upper half equal to $\frac{1}{n}\left(\zeta^{-(2j+1)k}\right)_{0 \leq j < n/2, 0 \leq k < n}$ and bottom half equal to the complex conjugate of the upper half. These matrix multiplications can be performed using complex discrete Fourier transforms, i.e., with $O(n \ln n)$ complex-valued arithmetic operations with the Cooley-Tukey FFT. Moreover, they are numerically extremely stable: If all operations are performed with a numerical precision of $p = \Omega(\ln n)$ bits, then the computed output vector $fl(\boldsymbol{y})$ satisfies $\|fl(\boldsymbol{y}) - \boldsymbol{y}\| \leq C \cdot (\ln n) \cdot 2^{-p} \cdot \|\boldsymbol{y}\|$, where $C$ is some absolute constant and $\boldsymbol{y}$ is the vector that would be obtained with exact computations. We refer to [23, Ch. 24] 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 $\sigma_i = \sigma_{i+n/2} = \alpha q \sqrt{1 + \sqrt{n} x_i}$ with the $x_i$'s sampled independently from the distribution $\Gamma(2, 1)$ for $i \leq n/2$. The distribution $\Gamma(2, 1)$ has density $x \exp(-x)$ for $x \geq 0$ and zero for $x < 0$.

Apart from a scaling factor and the choice of the polynomial representation, our R-LWE variant differs from that of [41] in that we round to $R$ using a rejection. The R-LWE problem 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.

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 run-time is bounded by a quantity that follows a geometric law of parameter $< 1$. Furthermore, in 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 Euclidean norm $\leq \alpha q n^{1/4} \omega(\sqrt{\ln n})$. The following statement improves on a bound given in Lemma 6 of the Eurocrypt proceedings paper presenting an earlier version of our results, that exploits the narrower $\Gamma(2, 1)$ distribution of the $x_i$'s. It also fixes a couple of mistakes in [64, Le. 6].

**Lemma 2.10.** *Let $y, r \in R$, with $r$ fixed and $y$ sampled from $\overline{\Upsilon}_\alpha$, with $\alpha q \geq n^{1/4}$. Then*

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

*Proof.* We define $\Upsilon_\alpha$ 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 $\Upsilon_\alpha$ instead of $\overline{\Upsilon}_\alpha$.

Let $y$ be sampled from $\Upsilon_\alpha$. The involved $\boldsymbol{\sigma}$ satisfies $\sigma_k = \sigma_{k+n/2} = \alpha q \sqrt{1 + \sqrt{n} x_k}$, with the $x_k$'s sampled independently from the distribution $\Gamma(2, 1)$. Let $(r^{(k)})_k$ be the embedding vector of $r$. Multiplying $y$ by $r$ is the same as sampling from $\rho_{\sigma'}$ with $\sigma'_k = \sigma'_{k+n/2} = \sigma_k |r^{(k)}|$ (see [42], and also [35, Le. 9] for a proof). We have $\sigma'_k \leq \alpha q n^{1/4} \omega(\sqrt{\ln n}) \cdot |r^{(k)}|$ for all $k \leq n$, with probability at least $1 - n^{-\omega(1)}$.

In order to obtain the coefficients of $yr$, it suffices to apply the matrices $\frac{1}{\sqrt{2}}\begin{pmatrix}1 & 1\\ i & -i\end{pmatrix} \otimes \mathrm{Id}_{n/2} \in \mathbb{C}^{n\times n}$ and $V$ to the row vector of the Gaussian samples. The magnitude of each entry of the matrix product being $\leq O(1/n)$, the coefficients of the polynomial $yr$ are distributed as statistically independent (one-dimensional) Gaussians of standard deviations $\leq \alpha q n^{-3/4}\omega(\sqrt{\ln n}) \cdot T_2(r) = \alpha q n^{-1/4}\omega(\sqrt{\ln n}) \cdot \|r\|$. The Euclidean norm of the resulting $n$-dimensional vector is $\leq \alpha q n^{1/4}\omega(\sqrt{\ln n}) \cdot \|r\|$ with probability $\geq 1 - n^{-\omega(1)}$. To complete the proof, observe that all the coordinates are $\leq \alpha q n^{-1/4}\omega(\ln n) \cdot \|r\|$ with probability $\geq 1 - n^{-\omega(1)}$. The additional rounding error $O(\sqrt{n})$ only changes the hidden constant factor in the $\omega(\ln n)$ factor, thanks to the condition $\alpha q \geq n^{1/4}$. □

## 3 A provably secure variant of NTRUEncrypt

In NTRUEncrypt, the public key $h$ is the ratio of the randomly generated secret key polynomials $f$ and $g$, whose coefficients have small magnitudes. In order to derive the IND-CPA security of the revised scheme from the hardness of R-LWE, we ensure that the distribution of $h$ is statistically very close to uniform over $R_q^\times$. (Computational indistinguishability fromuniformity would actually suffice, but we do not know how to achieve it based on standard lattice assumptions.) For this purpose, we sample $f$ and $g$ from the distribution $D_\sigma^\times$, obtained by sampling from $D_{\mathbb{Z}^n,\sigma}$ and rejecting if the sample (interpreted as an element of $R$) is not invertible modulo $q$. We will eventually choose $\sigma \approx n^c q^{1/2}$ for some small constant $c$.

The proof that the ratio $g/f$ is close to uniform when $f, g \hookleftarrow D_\sigma^\times$ proceeds in several steps. We aim at bounding the quantity $\sum_{a\in R_q^\times} |\Pr_{f,g}[g/f = a] - |R_q^\times|^{-1}|$ by some small amount $\varepsilon$. To do that, we show that with overwhelming probability over the choice of $a$, each term $|\Pr_{f,g}[g/f = a] - |R_q^\times|^{-1}|$ is $< |R_q^\times|^{-1} \cdot \varepsilon$. This is equivalent to showing that for the overwhelming majority of the pairs $(a_1, a_2) \in (R_q^\times)^2$, the quantity $|\Pr_{f,g}[fa_1 + ga_2 = 0] - |R_q^\times|^{-1}|$ is $< |R_q^\times|^{-1} \cdot \varepsilon$.

The latter statement can be seen as a consequence of a regularity bound for $(a_1, \ldots, a_m, \sum_i t_i a_i)$ with $m = 2$. More precisely, we prove a small bound $< |R_q^\times|^{-1} \cdot \varepsilon$ on the statistical distance $\varDelta$ to uniformity over $(R_q^\times)^m \times R_q$ of the distribution of $(a_1, \ldots, a_m, \sum_i t_i a_i)$ where the $a_i$'s are sampled uniformly and independently in $R_q^\times$ and the $t_i$'s are independently sampled from $D_\sigma^\times$. We need an unusually small bound on the statistical distance $\varDelta$, because we eventually sum this bound over $|R_q^\times|$ to obtain the uniformity of the public key $h$. A similar strong regularity was independently used by Agrawal et al. in [2, Th. 3] in the context of (non-structured) SIS/LWE for proving the security of an identity-based encryption scheme.

Another unusual facet of our regularity bound is the fact that the support of the $t_i$'s is not a lattice. We circumvent this difficulty by writing the support as the lattice $\mathbb{Z}^n$ minus the union of the lattices $L_{\varPhi_i} = \{x \in R : \varPhi_i | (x \bmod q)\}$ corresponding to the ideals $\langle q, \varPhi_i\rangle$ of $R$. (Recall that the $\varPhi_i$'s are the irreducible factors of $\varPhi$ modulo $q$). This observation leads us to obtain the desired regularity bound by combining regularity bounds for the $t_i$'s sampled in lattices, with an inclusion-exclusion technique (Theorem 3.1).

The remainder of the proof is more classical. The uniformity of $\sum_i t_i a_i$ for the $t_i$'s sampled from a lattice Gaussian is obtained by proving uniformity of the vector $\boldsymbol{t}$ made of the $t_i$'s taken modulo the kernel of the map $\boldsymbol{t} \mapsto \sum t_i a_i \bmod q$. Note that this kernel is a lattice. As $\boldsymbol{t}$ follows a lattice Gaussian distribution, uniformity modulo the kernel follows by studying the smoothing parameter of the kernel lattice and using Lemma 2.4. The latter is the purpose of Subsection 3.1.

The bounds in this section improve and generalize the results presented in [64]. In particular, they show that, for a given desired closeness to uniformity of $h = g/f$, using a modulus $q$ such that $x^n + 1$ splits into $k_q = O(1)$ irreducible factors allows to reduce the required standard deviation $\sigma$ for $f, g$ by a factor $\approx \sqrt{n}$, versus the case $k_q = n$ studied in [64].

### 3.1 New results on random module $q$-ary lattices

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

We generalize the definitions of the $\boldsymbol{a}^\perp$ and $L(\boldsymbol{a})$ lattices to incorporate the ideals of $R_q$. Let $\Phi = \prod_{i \le k_q} \Phi_i$ be the factorisation as a product of irreducible factors modulo $q$. Recall that the $\Phi_i$'s share the same degree $d_q = n/k_q$. The ideals of $R_q$ are of the form

$$I_S := \left(\prod_{i \in S} \Phi_i\right) \cdot R_q = \left\{ a \in R_q : \forall i \in S, a = 0 \bmod \Phi_i \right\}, \quad \text{with } S \subseteq \{1, \ldots, k_q\}.$$

We also define $L_S$ as the lattice corresponding to the ideal $\langle q, \prod_{i \in S} \Phi_i \rangle$ of $R$. More explicitly, we have $L_S = \{x \in R : (x \bmod q) \in I_S\}$.

For $\boldsymbol{a} \in R_q^m$ and $S \subseteq \{1, \ldots, k_q\}$, we define the following families of $R$-modules:

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

where $S$ is an arbitrary subset of $\{1, \ldots, k_q\}$. Note that $\boldsymbol{a}^\perp(I_S)$ is the intersection of $\boldsymbol{a}^\perp$ with the Cartesian product of $m$ copies of $L_S$. Also, if $S = \emptyset$ (resp. $S = \{1, \ldots, n\}$), then we have $\boldsymbol{a}^\perp(I_S) = \boldsymbol{a}^\perp$ (resp. $L(\boldsymbol{a}, I_S) = L(\boldsymbol{a})$).

We now describe an automorphism of $R$ that will help us exhibit the duality between the modules above. In the ring $R$, we have $x^{-1} = -x^{n-1}$. Therefore, mapping $a(x) \in R$ to $a^\star(x) = a(x^{-1}) \in R$ provides ring automorphism. This map induces a bijection from the set of factors $\Phi_i$ to itself. It has the following useful matrix interpretation: If we let $A$ denote the $n \times n$ matrix having as its $i$-th row the coefficient vector of $x^i \cdot a(x)$ for $i = 0, \ldots, n-1$, then $a^\star(x)$ has coefficient vector the first column of $A$. For an ideal $I_S = (\prod_{i \in S} \Phi_i) \cdot R_q$ of $R$, we let $I_S^\star$ denote the ideal $(\prod_{i \in S} \Phi_i^\star) \cdot R_q$.

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

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

*Proof.* We first prove that $\frac{1}{q} L(\boldsymbol{a}^\star, I_{\overline{S}}^\star) \subseteq \widehat{\boldsymbol{a}^\perp(I_S)}$. Let $(t_1, \ldots, t_m) \in \boldsymbol{a}^\perp(I_S)$ and $(u_1, \ldots, u_m) \in L(\boldsymbol{a}^\star, I_{\overline{S}}^\star)$. Write $t_i = \sum_{j<n} t_{i,j} x^j$ and $u_i = \sum_{j<n} u_{i,j} x^j$ for any $i \le m$. Our goal is to show that $\sum_{i \le m, j \le n} t_{i,j} u_{i,j} = 0 \bmod q$. This is equivalent to showing that the constant coefficient of the polynomial $\sum_{i \le m} t_i u_i^\star$ is 0 modulo $q$. It thus suffices to show that $\langle \boldsymbol{t}, \boldsymbol{u}^\star \rangle = 0 \bmod q$. By definition

of the $u_i$'s, there exists $s \in R_q$ such that $(u_i \bmod q) = a_i^\star \cdot s + b_i$ for some $b_i \in I_{\overline{S}}^\star$. We have the following, modulo $q$:

$$\langle \boldsymbol{t}, \boldsymbol{u}^\star \rangle = s^\star \cdot \langle \boldsymbol{t}, \boldsymbol{a} \rangle + \langle \boldsymbol{t}, \boldsymbol{b}^\star \rangle = 0,$$

where we used that $\langle \boldsymbol{t}, \boldsymbol{a} \rangle = 0 \bmod q$ by definition of $\boldsymbol{t}$ and $\langle \boldsymbol{t}, \boldsymbol{b}^\star \rangle = 0 \bmod q$ because $(t_i \bmod q) \in I_S$ and $b_i^\star \in I_{\bar{S}}$ for each $i \leq m$. This provides the desired inclusion.

The reverse inclusion $\frac{1}{q} L(\boldsymbol{a}^\star, I_{\overline{S}}^\star) \supseteq \widehat{\boldsymbol{a}^\perp(I_S)}$ is equivalent, by duality, to $L\widehat{(\boldsymbol{a}^\star, I_{\overline{S}}^\star})\subseteq \frac{1}{q} \boldsymbol{a}^\perp(I_S)$. To show the latter, it suffices to consider the elements of $L(\boldsymbol{a}^\star, I_{\overline{S}})$ corresponding to $s = 1$. □

We now 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 the Minkowski upper bound $\det(L(\boldsymbol{a}, I_S))^{\frac{1}{mn}} = q^{(1-\frac{1}{m})\frac{|S|}{k_q}}$ on $\lambda_1^\infty(L(\boldsymbol{a}, I_S))$. (We have $\det(L(\boldsymbol{a}, I_S)) = q^{(m-1)|S|d_q}$ because there are $q^{|S|d_q + m(n-|S|d_q)}$ points of $L(\boldsymbol{a}, I_S)$ in the cube $[0, q-1]^{mn}$.) We provide two lower bounds. The first lower bound is useful for all parameter settings and matches the Minkowski upper bound up to a factor $\frac{1}{\sqrt{n}} q^{-\varepsilon}$ for an arbitrarily small constant $\varepsilon > 0$. The second bound is specific to the case $|S| = k_q$ and matches the Minkowski bound up to a factor $q^{-k_q \cdot \varepsilon}$, thus improving on the first bound by a factor $\approx \sqrt{n}$ in the case $k_q = O(1)$ (which was not treated in [64]). Even in the case $k_q = n$, the first bound improves on the bound given in [64], by using a point counting bound based on the minima of the ideals of $R_q$.

**Lemma 3.2.** *Let $n \geq 8$ be a power of 2 and $q \geq 5$. Assume that $\Phi = x^n + 1$ splits into $k_q$ distinct irreducible factors modulo $q$, each of degree $d_q = n/k_q$. Then, for $m \geq 2$ and $\varepsilon > 0$, we have*

$$\lambda_1^\infty(L(\boldsymbol{a}, I_S)) \geq \begin{cases} \frac{1}{\sqrt{n}} q^{(1-\frac{1}{m})\frac{|S|}{k_q} - \varepsilon} & \text{for any } 0 \leq |S| \leq k_q \\ q^{1-\frac{1}{m} - k_q \cdot \varepsilon} & \text{for } |S| = k_q \end{cases}$$

*except with probability $\leq 2^{4mn} q^{-\varepsilon mn}$ over the uniformly random choice of $\boldsymbol{a}$ in $(R_q^\times)^m$.*

*Proof.* By the Chinese Remainder Theorem, we know that $R_q$ (resp. $R_q^\times$) is isomomorphic to $(\mathbb{F}_{q^{d_q}})^{k_q}$ (resp. $(\mathbb{F}_{q^{d_q}}^\times)^{k_q}$) via the isomorphism $t \mapsto (t \bmod \Phi_i)_{i \leq k_q}$. Let $\Phi_S = \prod_{i \in S} \Phi_i$: it is a degree $|S|d_q$ 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 $< B$. We bound $p$ from above by using 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 $\gcd(s, \Phi_S) = \gcd(t_i, \Phi_S)$ (up to multiplication by an element of $\mathbb{F}_{q^{d_q}}^\times$): 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 $\gcd(s, \Phi_S) = \gcd(t_i, \Phi_S) = \Phi_{S'}$ for some $S' \subseteq S$ of cardinality $0 \leq k \leq |S|$. For any $j \in S'$, we have $t_i = a_i s = 0 \bmod \Phi_j$ regardless of the value of $a_i \bmod \Phi_j$, whereas 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 $\mathbb{F}_{q^{d_q}}^\times$. So, overall, there are $(q^{d_q} - 1)^{k_q + k - |S|}$ distinct $a_i$'s in $R_q^\times$ such that $t_i = a_i s \bmod I_S$. This leads to $p_i(t_i, s) = (q^{d_q} - 1)^{k - |S|}$.

So far, we have shown that the probability $p$ can be bounded from above by:

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

For $|S'| = k$, let $N(B,k)$ denote the number of $t \in R_q$ such that $\|t\|_\infty < B$ and $t = \Phi_{S'} t'$ for some $t' \in R_q$ of degree $< n - kd_q = n(1 - k/k_q)$. We consider two upper bounds for $N(B,k)$, from which we get the claimed bounds on $\lambda_1^\infty(L(\boldsymbol{a}, I_S))$.

As our first bound for $N(B,k)$, with $B = \frac{1}{\sqrt{n}} \cdot q^\beta$, we claim that $N(B,k) \leq 2^{2n} q^{(\beta - k/k_q)n}$ for $k < \beta \cdot k_q$ and $N(B,k) = 0$ for $k \geq \beta \cdot k_q$. For this, we observe that $N(B,k)$ is the number of points of the lattice $I_{S'} + q\mathbb{Z}^n = \langle \Phi_{S'}, q \rangle$ in the hypercube $C(2B)$ of sidelength $2B$, where a hypercube of sidelength $\ell$ is defined by $C(\ell) = \{\boldsymbol{v} \in \mathbb{R}^n : \|\boldsymbol{v}\|_\infty < \ell/2\}$. Let $\lambda := \lambda_1^\infty(I_{S'} + q\mathbb{Z}^n)$. If we center a hypercube $C(\lambda)$ of sidelength $\lambda$ on each of the $N(B,k)$ points of $I'_S + q\mathbb{Z}^n$ in $C(2B)$, the resulting $N(B,k)$ hypercubes do not intersect, and yet are all contained within the enlarged hypercube $C(2B + \lambda)$. It follows that $N(B,k) \leq \frac{\mathrm{vol}(C(2B+\lambda))}{\mathrm{vol}(C(\lambda))} = (\frac{2B}{\lambda} + 1)^n$. To derive a lower bound on $\lambda$, note that for any $t \in I_{S'}$, we have $\mathcal{N}(t) = \mathcal{N}(\langle t \rangle) \geq \mathcal{N}(\langle \Phi_{S'}, q \rangle) = q^{kd_q}$, where the inequality is because the ideal $\langle t \rangle$ is a sub-ideal of $\langle \Phi_{S'}, q \rangle$, and the last equality is because $\deg \Phi_{S'} = kd_q$. It follows from the arithmetic-geometric inequality that $\|t\| = \frac{1}{\sqrt{n}} T_2(t) \geq \mathcal{N}(t)^{1/n} \geq q^{k/k_q}$. By equivalence of norms, we conclude that $\|t\|_\infty \geq \lambda \geq \frac{1}{\sqrt{n}} q^{k/k_q}$. Hence, using $B = \frac{1}{\sqrt{n}} q^\beta$, for $k \geq \beta \cdot k_q$, we have $\lambda \geq B$ so $N(B,k) = 0$, while for $k < \beta \cdot k_q$, we have $N(B,k) \leq (\frac{2B}{\lambda} + 1)^n \leq (2q^{\beta - k/k_q} + 1)^n \leq 2^{2n} q^{(\beta - k/k_q)n}$, as claimed.

As our second bound for $N(B,k)$, we claim that $N(B,k) \leq (2B)^{n - kd_q} = (2B)^{n(1 - k/k_q)}$. Indeed, since the degree of $\Phi_{S'}$ is $kd_q$, the vector $\overline{t}$ formed by the $n - kd_q$ low-order coefficients of $t = \Phi_{S'} t'$ is related to the vector $\overline{t'}$ formed by the $n - kd_q$ low-order coefficients of $t'$ by a lower triangular $(n - kd_q) \times (n - kd_q)$ matrix whose diagonal coefficients are equal to the non-zero constant coefficient of $\Phi_{S'}$. Hence this matrix is non-singular modulo $q$ and the mapping from $\overline{t'}$ to $\overline{t}$ is one-to-one. This provides the claim.

Using the fact that the number of subsets of $S$ is $2^{|S|}$, and the fact that the number of $s \in R_q/I_S$ divisible by $\Phi_{S'}$ is $q^{d_q(|S| - k)}$, the above upper bound on $p$ implies that

$$p \ \leq \ 2^{(m+1)|S|} \cdot \max_{0 \leq k \leq |S|} \frac{N(B,k)^m}{q^{(m-1)(|S| - k)d_q}}.$$

Using our first bound on $N(B,k)$ with $B = \frac{1}{\sqrt{n}} q^\beta$, we get

$$p \leq 2^{(m+1)(|S| + 2n)} \cdot \max_{0 \leq k < \beta \cdot k_q} q^{n\left(m(\beta - \frac{k}{k_q}) - (m-1)\frac{|S| - k}{k_q}\right)}.$$

Viewed as a function of $k$, the exponent in the right hand side is maximized for $k = 0$. It then has the value $-mn\varepsilon$, when $\beta = (1 - \frac{1}{m})\frac{|S|}{k_q} - \varepsilon$. This gives the first claimed bound on $\lambda_1^\infty(L(\boldsymbol{a}, I_S))$.

In the case $|S| = k_q$, using our second bound on $N(B,k)$ with $B = q^\beta$, and noting that $N(B, k_q) = 0$, we get

$$p \ \leq \ 2^{(m+1)(|S| + 2n)} \cdot \max_{0 \leq k < k_q} q^{n((1-\beta)m - 1)\left(\frac{k}{k_q} - 1\right)} \ = \ 2^{(m+1)(|S| + 2n)} \cdot q^{-\frac{n}{k_q}((1-\beta)m - 1)},$$

where the last equality holds for any $\beta \leq 1 - \frac{1}{m}$. Using $\beta = 1 - \frac{1}{m} - k_q\varepsilon$ gives the second claimed bound on $\lambda_1^\infty(L(\boldsymbol{a}, I_S))$. □

In our analysis of the distribution of the NTRU key $g/f$ with $k_q = O(1)$, we will also use a lower bound on $\lambda_1(\boldsymbol{a}^\perp(I_S))$. As in Lemma 3.2, we give two bounds, although in this case our application only needs the first bound.

**Lemma 3.3.** *Let $n \geq 8$ be a power of 2 and $q \geq 5$. Assume that $\Phi = x^n + 1$ splits into $k_q$ distinct irreducible factors modulo $q$, each of degree $d_q = n/k_q$. Then, for $m \geq 2$ and $\varepsilon > 0$, we have*

$$\lambda_1^\infty(\boldsymbol{a}^\perp(I_S)) \geq \begin{cases} \frac{1}{\sqrt{n}} q^{\frac{1}{m}+(1-\frac{1}{m})\frac{|S|}{k_q}-\varepsilon} & \text{for any } 0 \leq |S| \leq k_q \\ q^{\frac{1}{m}-k_q\cdot\varepsilon} & \text{for } |S| = 0 \end{cases}$$

*except with probability $\leq 2^{4n}q^{-\varepsilon mn}$ over the uniformly random choice of $\boldsymbol{a}$ in $(R_q^\times)^m$.*

*Proof.* We proceed analogously to the proof of Lemma 3.2.

Let $p$ denote the probability (over $\boldsymbol{a}$) that $L(\boldsymbol{a}^\perp(I_S))$ contains a non-zero vector $\boldsymbol{t}$ of infinity norm $< B$. We bound $p$ from above by using the union bound, summing the probabilities $p(\boldsymbol{t}) = \Pr_{\boldsymbol{a}}[\sum_{i\leq m} a_i t_i = 0 \bmod q]$ over all possible values for $\boldsymbol{t}$ of infinity norm $< B$ and $t_i \in I_S$ for $i = 1, \ldots, m$. By the Chinese Remainder Theorem, we have $p(\boldsymbol{t}) = \prod_{j\leq k_q} p_j(\boldsymbol{t})$, where $p_j(\boldsymbol{t}) = \Pr_{\boldsymbol{a}}[\sum_{i\leq m} a_i t_i = 0 \bmod \Phi_j]$. Let $\Phi_S = \prod_{i\in S}\Phi_i$, $\Phi_{\bar{S}} = \prod_{i\in\bar{S}}\Phi_i$ and $\Phi_{S'} = \gcd(t_1, \ldots, t_m, \Phi_{\bar{S}}) = \prod_{i\in S'}\Phi_i$ for some $S' \subseteq \bar{S}$ of cardinality $0 \leq k \leq |\bar{S}|$. For any $j \in S \cup S'$, we have $\sum_{i\leq m} t_i a_i = 0 \bmod \Phi_j$ regardless of the value of $a_i \bmod \Phi_j$. For any $j \in \bar{S}\backslash S'$, there exists $i \leq m$ such that $t_i \neq 0 \bmod \Phi_j$ so that for any choice of $\{a_j\}_{j\neq i}$, there is a unique value of $a_i \bmod \Phi_j$ such that $\sum_{i\leq m} t_i a_i = 0 \bmod \Phi_j$; It follows that $p_j(\boldsymbol{t}) = \frac{1}{q^{d_q}-1}$. As a consequence, we have $p(\boldsymbol{t}) = \frac{1}{(q^{d_q}-1)^{|\bar{S}|-k}}$, and:

$$p \leq \sum_{0\leq k\leq|\bar{S}|} \sum_{\substack{S' \subseteq \bar{S} \\ |S'| = k}} \sum_{\substack{\boldsymbol{t} \in (R_q)^m \\ \forall i, 0 < \|t_i\|_\infty < B \\ \forall i, \Phi_S \cdot \Phi_{S'} | t_i}} \frac{1}{(q^{d_q}-1)^{|\bar{S}|-k}}.$$

For $S'$ with $|S'| = k$, let $N(B, k)$ denote the number of $t \in R_q$ such that $\|t\|_\infty < B$ and $t = \Phi_S\Phi_{S'}t'$ for some $t' \in R_q$ of degree $< n(1 - (k + |S|)/k_q)$. Exactly as in the proof of Lemma 3.2, we derive two upper bounds for $N(B, k)$, from which we get the claimed bounds on $\lambda_1^\infty(L(\boldsymbol{a}, I_S))$. The first upper bound, with $B = \frac{1}{\sqrt{n}}q^\beta$, shows that $N(B, k) = 0$ for $k \geq \beta \cdot k_q - |S|$, while $N(B, k) \leq 2^{2n}q^{(\beta-(|S|+k)/k_q)n}$ for $k < \beta\cdot k_q - |S|$. The second bound is $N(B, k) \leq (2B)^{n(1-(|S|+k)/k_q)}$.

The first bound on $N(B, k)$ with $B = \frac{1}{\sqrt{n}}q^\beta$, leads to

$$p \leq 2^{2|\bar{S}|+2n} \cdot \max_{0\leq k<\beta\cdot k_q} q^{n\left(m(\beta-\frac{|S|+k}{k_q})-\frac{|\bar{S}|-k}{k_q}\right)}.$$

Viewed as a function of $k$, the exponent in the right hand side is maximized for $k = 0$. It then has the value $-mn\varepsilon$, when $\beta = \frac{1}{m} + (1 - \frac{1}{m})\frac{|S|}{k_q} - \varepsilon$. This gives the first claimed bound.

In the case $|S| = 0$, using our second bound on $N(B, k)$ with $B = q^\beta$, and noting that $N(B, k_q) = 0$, we get

$$p \ \leq \ 2^{2|\bar{S}|+n} \cdot \max_{0\leq k<k_q} q^{n(1-m\beta)\left(\frac{k}{k_q}-1\right)} \ = \ 2^{2|\bar{S}|+n} \cdot q^{n(1-m\beta)\left(1-\frac{1}{k_q}\right)}.$$

where the last equality holds for any $\beta \leq \frac{1}{m}$. Using $\beta = \frac{1}{m} - k_q \cdot \varepsilon$ gives the second claimed bound. □

### 3.2 Regularity bounds for ring $\boldsymbol{R_q}$

We now study the closeness to uniformity of the 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 [44], 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 [44, 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 [44] 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 $m \ln d = \Omega(n)$, which inevitably leads to inefficient cryptographic functions. When the $a_i$'s are chosen uniformly from the whole ring $R_q$ with $q = 1 \bmod 2n$, 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 [21, 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 R-SIS trapdoor generation of [65, Se. 3], thus extending the latter to a fully splitting $q$.

**Theorem 3.1.** *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into $k_q$ irreducible 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 $\sigma \geq \ln(2mn(1 + 1/\delta))/\pi \cdot \min(\sqrt{n} \cdot q^{\frac{1}{m}+\varepsilon}, q^{\frac{1}{m}+k_q\varepsilon})$. Then for all except a fraction $\leq 2^{4mn} q^{-\varepsilon mn}$ of $\boldsymbol{a} \in (R_q^\times)^m$, we have $\eta_\delta(\boldsymbol{a}^\perp) \leq \sqrt{\ln(2mn(1 + 1/\delta))/\pi} \cdot \min(\sqrt{n} \cdot q^{\frac{1}{m}+\varepsilon}, q^{\frac{1}{m}+k_q\cdot\varepsilon})$, and the distance to uniformity of $\sum_{i\leq m} t_i a_i$ is $\leq 2\delta$. As a consequence:*

$$\Delta\left[\Big(a_1, \ldots, a_m, \sum_{i\leq m} t_i a_i\Big);\ U\Big((R_q^\times)^m \times R_q\Big)\right] \leq 2\delta + 2^{4mn} q^{-\varepsilon mn}.$$

*Proof.* 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 $\leq 2^{4mn} q^{-\varepsilon mn}$ 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 $\boldsymbol{t} \bmod \boldsymbol{a}^\perp$. In the following, since it is needed for our analysis of the NTRU key generation algorithm (see Theorem 3.2 in Section 3.3) we actually study the distance to uniformity of $\boldsymbol{t} \bmod \boldsymbol{a}^\perp(I_S)$ for any $S \subseteq \{1, \ldots, k_q\}$. By Lemma 2.4, we have $\Delta_{\boldsymbol{a}} \leq 2\delta$ if $\sigma$ is greater than the smoothing parameter $\eta_\delta(\boldsymbol{a}^\perp(I_S))$ of $\boldsymbol{a}^\perp(I_S) \subseteq \mathbb{Z}^{mn}$. To bound $\eta_\delta(\boldsymbol{a}^\perp(I_S))$ from above, we apply Lemma 2.1, which reduces the task to bounding the minimum of the dual lattice from below. By Lemma 3.1, the latter lattice is $\widehat{\boldsymbol{a}^\perp(I_S)} = \frac{1}{q} \cdot L(\boldsymbol{a}^\star, I_{\overline{S}}^\star)$ (where $\boldsymbol{a}^\star \in (R_q^\times)^m$ is in one-to-one correspondence with $\boldsymbol{a}$), and the latter task has been addressed by Lemma 3.2. Hence, we obtain the following result as a direct consequence of Lemmata 2.1, 2.4, 3.1 and 3.2.

**Lemma 3.4.** *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into $k_q$ irreducible factors modulo prime $q \geq 5$. Let $S \subseteq \{1, \ldots, k_q\}$, $m \geq 2$, $\varepsilon > 0$, $\delta \in (0, 1/2)$, $\boldsymbol{c} \in \mathbb{R}^{mn}$ and $\boldsymbol{t} \hookleftarrow D_{\mathbb{Z}^{mn},\sigma,\boldsymbol{c}}$, with*

$$\sigma \geq \begin{cases} \sqrt{n \ln(2mn(1+1/\delta))/\pi} \cdot q^{1-(1-\frac{1}{m})(1-\frac{|S|}{k_q})+\varepsilon} & \text{for any } 0 \leq |S| \leq k_q \\ \sqrt{\ln(2mn(1+1/\delta))/\pi} \cdot q^{\frac{1}{m}+k_q \cdot \varepsilon} & \text{for } |S| = 0. \end{cases}$$

*Then for all except a fraction $\leq 2^{4mn} q^{-\varepsilon mn}$ 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 $\boldsymbol{c} = \boldsymbol{0}$. □

### 3.3 Revised key generation algorithm for the NTRUEncrypt

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 R-SIS.

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 2.6), 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 negligible, the impact on our results is also negligible. Furthermore, it can be checked that our conditions on standard deviations are much stronger than the one in Lemma 2.6. From now on, we will assume we have a perfect discrete Gaussian sampler.

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

*Proof.* We are to bound the probability that $p \cdot f' + a$ belongs to $I := \langle q, \Phi_k \rangle$ by $q^{-n/k_q} + 2\delta$, for any $k \leq k_q$. The result then follows from the Chinese Remainder Theorem and the union bound. We have $\mathcal{N}(I) = q^{n/k_q}$, so that $\lambda_1(I) \leq \sqrt{n} q^{1/k_q}$, 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 2.4 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 $\leq q^{-n/k_q} + 2\delta$, 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 3.6.** *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into $k_q$ irreducible factors modulo prime $q \geq 8n$. Let $\sigma \geq \sqrt{n \ln n} \cdot q^{1/k_q}$. The secret key polynomials $f, g$ returned by the algorithm of Fig. 1 satisfy, with probability $\geq 1 - 2^{-n+3}$:*

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

*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 acceptance probability. The result follows by combining Lemmata 2.3 and 3.5. □

In the algorithm of Fig. 1, the polynomials $f'$ and $g$ are independently sampled from the discrete Gaussian distribution $D_{\mathbb{Z}^n,\sigma}$ 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 3.2 to show that the statistical closeness to uniformity of a quotient of two distributions $(z + p \cdot D_{\sigma,y}^\times)$ 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 $p \in R_q^\times$, multiplication by $p$ induces a bijection of $R_q$, and thus the statistical closeness to uniformity carries over to the public key $h = pg/f$. The following theorem gives two bounds, whose usefulness depends on the number of irreducible factors $k_q$ in the factorization of $x^n + 1$ modulo $q$. The first bound is most useful for large $k_q = \Omega(n)$, while the second bound is better for small $k_q = O(1)$, allowing a smaller $\sigma$ by a factor $\approx \sqrt{n}$ versus the first bound.

**Theorem 3.2.** *Let $n \geq 8$ be a power of 2 such that $\Phi = x^n + 1$ splits into $k_q$ irreducible factors modulo prime $q \geq 5$. Let $0 < \varepsilon' < 1/3$, $y_i \in R_q$ and $z_i = -y_i p^{-1} \bmod q$ for $i \in \{1, 2\}$. Then*

$$\Delta\left[\frac{y_1 + p \cdot D_{\sigma,z_1}^\times}{y_2 + p \cdot D_{\sigma,z_2}^\times} \bmod q \,;\, U\left(R_q^\times\right)\right] \leq \begin{cases} 2^{10n} q^{-\frac{\lfloor \varepsilon' k_q \rfloor}{k_q} \cdot n} & \textit{if } \sigma \geq n \cdot \sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+\varepsilon'} \\ 2^{10n} q^{-\varepsilon' n} & \textit{if } \sigma \geq \sqrt{n \ln(8nq)} \cdot q^{\frac{1+k_q\varepsilon'}{2}} \textit{ and } q \geq n^{\frac{k_q}{1-2k_q\varepsilon'}}. \end{cases}$$

*Proof.* For $a \in R_q^\times$, we define $Pr_a = \Pr_{f_1,f_2}[(y_1 + pf_1)/(y_2 + pf_2) = a]$, where $f_i \hookleftarrow D_{\sigma,z_i}^\times$ for $i \in \{1, 2\}$. We are to show that $|Pr_a - |R_q^\times|^{-1}| \leq 2^{2n+5} q^{-n\lfloor \varepsilon' k_q \rfloor / k_q} \cdot |R_q^\times|^{-1} =: \varepsilon''$ (resp. $\leq 2^{6n+4} q^{-\varepsilon' n} \cdot |R_q^\times|^{-1}$). This directly gives the claimed bounds. The fraction of $a \in R_q^\times$ such that $|Pr_a - |R_q^\times|^{-1}| \leq \varepsilon''$ is equal to the fraction of $\boldsymbol{a} = (a_1, a_2) \in (R_q^\times)^2$ such that $|Pr_{\boldsymbol{a}} - |R_q^\times|^{-1}| \leq \varepsilon''$, where $Pr_{\boldsymbol{a}} = \Pr_{f_1,f_2}[a_1 f_1 + a_2 f_2 = a_1 z_1 + a_2 z_2]$. This is because $a_1 f_1 + a_2 f_2 = a_1 z_1 + a_2 z_2$ is equivalent to $(y_1 + pf_1)/(y_2 + pf_2) = -a_2/a_1$ (in $R_q^\times$), and $-a_2/a_1$ is uniformly random in $R_q^\times$ when $\boldsymbol{a} \hookleftarrow U((R_q^\times)^2)$.

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

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

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

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

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

In the rest of the proof, we show that, except for a fraction $\leq 2^{9n} q^{-\varepsilon' n}$ of $\boldsymbol{a} \in (R_q^{\times})^2$:

$$D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp\times}) = (1 + \delta_0)|R_q^{\times}| q^{-2n},$$
$$\forall i \in \{1, 2\} : \ D_{\mathbb{Z}^n,\sigma}(z_i + R_q^{\times} + q\mathbb{Z}^n) = (1 + \delta_i)|R_q^{\times}| q^{-n}.$$

where $|\delta_i| \leq 2^{2n+2} q^{-n\lfloor \varepsilon' k_q \rfloor / k_q}$ (resp. $|\delta_i| \leq 2^{6n+1} q^{-\varepsilon' n}$) for $i \in \{0, 1, 2\}$. The bounds on $|Pr_{\boldsymbol{a}} - |R_q^{\times}|^{-1}|$ follow by a routine computation.

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

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

To get our first our bound, we proceed as follows. For the terms of (1) with $|S| \leq \varepsilon' k_q$, we apply the first bound of Lemma 3.4 with $m = 2$ and $\varepsilon = \varepsilon'/2$. The assumption of Lemma 3.4 on $\sigma$ holds, with $\delta := q^{-n(1+\lfloor \varepsilon' k_q \rfloor / k_q)}$. Further, we have $\det(\boldsymbol{a}^{\perp}(I_S)) = q^{n(1+|S|/k_q)}$: Indeed, since $\boldsymbol{a} \in (R_q^{\times})^2$, there are $q^{n(1-|S|/k_q)}$ elements of $\boldsymbol{a}^{\perp}(I_S)$ in $[0, q-1]^{2n}$. We conclude that $|D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^{\perp}(I_S)) - q^{-n(1+|S|/k_q)}| \leq 2\delta$, for all except a fraction $\leq 2^{8n} q^{-\varepsilon' n}$ 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 $|S| > \varepsilon' k_q$, we choose $S' \subseteq S$ with $|S'| = \lfloor \varepsilon' k_q \rfloor$. Then we have $\boldsymbol{a}^{\perp}(I_S) \subseteq \boldsymbol{a}^{\perp}(I_{S'})$ and hence $D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^{\perp}(I_S)) \leq D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^{\perp}(I_{S'}))$. By using with $S'$ the above result for small $|S|$, we obtain $D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^{\perp}(I_S)) \leq 2\delta + q^{-n(1+\lfloor \varepsilon' k_q \rfloor / k_q)}$.

Overall, we have, except possibly for a fraction $\leq 2^{9n} q^{-\varepsilon' n}$ of $\boldsymbol{a} \in (R_q^{\times})^2$:

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

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

For our second bound, we argue as follows. For the term of (1) with $|S| = 0$, we apply the second bound in Lemma 3.4 with $\varepsilon = \varepsilon'/2$. By the choice of $\sigma$, the Lemma 3.4 assumption on $\sigma$ holds, with $\delta := q^{-2n}$. We have $|R/\boldsymbol{a}^{\perp}(I_S)| = \det(\boldsymbol{a}^{\perp}(I_S)) = q^n$ and hence $|D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^{\perp}(I_S)) - q^{-n}| \leq 2\delta$, for all except a fraction $\leq 2^{8n} q^{-\varepsilon' n}$ of $\boldsymbol{a} \in (R_q^{\times})^2$.

For the terms of (1) with $|S| \geq 1$, unlike the argument above, we cannot choose for $|S| = 1$ an $I'_S$ with $S' \subseteq S$ and $\det \boldsymbol{a}^{\perp}(I_{S'}) \approx q^{(1+\varepsilon)n}$: such an ideal $I_{S'}$ does not exist, as the only possible choice for $S'$ is the empty set, which gives $\det \boldsymbol{a}^{\perp}(I_{S'}) = q^n$, and the latter is too small. Instead, we proceed as follows. Let $L' = N \cdot \mathbb{Z}^{2n}$, where $N = \lceil \frac{1}{4} q^{1/2+\varepsilon'/2} \rceil$. Note that $\det L' = N^{2n} \geq 2^{-4n} q^{(1+\varepsilon')n}$, and since $\lambda_{2n}(L') = N \leq \frac{1}{2} q^{1/2+\varepsilon'/2}$, we have by Lemma 2.1 with $\delta = q^{-2n}$ that $\eta_\delta(L') \leq \sqrt{n \ln(8nq)} q^{1/2+\varepsilon'/2}$. Hence, by Lemma 2.4 and the choice of $\sigma$, we have $D_{\mathbb{Z}^{2n},\sigma}(L') \leq 2^{4n} q^{-(1+\varepsilon')n} + 2\delta$. To use the last bound, we now show that, for $|S| \geq 1$, we have $D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp}(I_S)) \leq D_{\mathbb{Z}^{2n},\sigma}(L')$. For this, we use a rounding process $\phi : \mathbb{Z}^{2n} \to L'$ to map $\boldsymbol{z} + \boldsymbol{a}^{\perp}(I_S)$ onto a subset of $L'$ such that the following two properties hold:

1. The map $\phi$ is one-to-one on $\boldsymbol{z} + \boldsymbol{a}^{\perp}(I_S)$,
2. For each $\boldsymbol{v} \in \mathbb{Z}^{2n}$, we have $\|\phi(\boldsymbol{v})\| \leq \|\boldsymbol{v}\|$.

Since $D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{w}) \geq D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{v})$ for any $\boldsymbol{v}, \boldsymbol{w} \in \mathbb{Z}^{2n}$ with $\|\boldsymbol{w}\| \leq \|\boldsymbol{v}\|$, property 2 of $\phi$ implies that $D_{\mathbb{Z}^{2n},\sigma}(\boldsymbol{z} + \boldsymbol{a}^{\perp}(I_S)) \leq \sum_{\boldsymbol{v} \in \boldsymbol{z}+\boldsymbol{a}^{\perp}(I_S)}(\phi(\boldsymbol{v}))$, and by property 1 of $\phi$, the points $\{\phi(\boldsymbol{v})\}_{\boldsymbol{v} \in \boldsymbol{z}+\boldsymbol{a}^{\perp}(I_S)}$ are *distinct* points of $L'$, so that $\sum_{\boldsymbol{v} \in \boldsymbol{z}+\boldsymbol{a}^{\perp}(I_S)}(\phi(\boldsymbol{v})) \leq D_{\mathbb{Z}^{2n},\sigma}(L')$, as required. It remains to define $\phi$ and show that it has both properties. For $\boldsymbol{v} \in \mathbb{Z}^{2n}$, let $\phi(\boldsymbol{v})$ round each coordinate $v_i$ of $\boldsymbol{v}$ to the nearest multiple of $N$ which is less than or equal to $|v_i|$ in absolute value, i.e., $\phi(\boldsymbol{v}) = (v'_1, \ldots, v'_{2n})$ with $v'_i = \lfloor \frac{|v_i|}{N} \rfloor \cdot N \cdot sign(v_i)$. Since $|v'_i| \leq |v_i|$, property 2 of $\phi$ is clearly satisfied. To show property 1, note that $\|\phi(\boldsymbol{v}) - \boldsymbol{v}\|_\infty < N$ for all $\boldsymbol{v}$ in $\mathbb{Z}^{2n}$. Suppose towards a contradiction that $\phi$ is not one-to-one on $\boldsymbol{z} + \boldsymbol{a}^{\perp}(I_S)$. Then there exist two vectors $\boldsymbol{v}_1 \neq \boldsymbol{v}_2$ in $\boldsymbol{z} + \boldsymbol{a}^{\perp}(I_S)$ with $\phi(\boldsymbol{v}_1) = \phi(\boldsymbol{v}_2) = \boldsymbol{v}$. A triangle inequality then gives that $\boldsymbol{v}_1 - \boldsymbol{v}_2$ is a non-zero vector of $\boldsymbol{a}^{\perp}(I_S)$ with $\|\boldsymbol{v}_1 - \boldsymbol{v}_2\| < 2N \leq q^{1/2+\varepsilon'/2}$. However, by the first bound of Lemma 3.3 with $m = 2$, $|S| = 1$, and $\varepsilon = \varepsilon'/2$, we have $\lambda_1^\infty(\boldsymbol{a}^{\perp}(I_S)) \geq \frac{1}{\sqrt{n}} q^{\frac{1}{2} + \frac{1}{2k_q} - \frac{\varepsilon'}{2}}$, except for a fraction $\leq 2^{4n} q^{-\varepsilon' n}$ of $\boldsymbol{a} \in (R_q^{\times})^2$. By the condition on $q$, this gives a contradiction, so $\phi$ has property 1, except for a fraction $\leq 2^{4n} q^{-\varepsilon' n}$ of $\boldsymbol{a} \in (R_q^{\times})^2$. We conclude that for the terms with $|S| \geq 1$, we have $D_{\mathbb{Z}^{2n},\sigma,-\boldsymbol{z}}(\boldsymbol{a}^{\perp}(I_S)) \leq 2^{4n+1} q^{-(1+\varepsilon')n}$. Hence, similarly to the first bound, we obtain our second bound $|\delta_0| \leq \frac{q^{2n}}{(q^{n/k_q}-1)^{k_q}} 2^{5n+1} q^{-(1+\varepsilon')n} \leq 2^{6n+1} q^{-\varepsilon' n}$.

Handling (2). For the bounds on $\delta_1$ and $\delta_2$, we use a similar argument. Let $i \in \{1, 2\}$. The $z_i$ term can be handled like the $\boldsymbol{z}$ term of (1). Therefore, in this case we need a good bound on $D_{\mathbb{Z}^n,\sigma,-z_i}(I_S + q\mathbb{Z}^n)$. By Lemma 2.4 this reduces to finding a good bound on the smoothing parameter of the ideal lattice $L_S = I_S + q\mathbb{Z}^n$. For this, we first observe that $L_S = \boldsymbol{a}^{\perp}(I_S)$ in the special case $m = 1$ and $a_1(x) = \prod_{i \in \bar{S}} \Phi_i(x)$, where $\bar{S}$ denotes the complement of $S$. Therefore, by Lemma 3.1, the dual lattice $\widehat{L_S} = \frac{1}{q} L(a_1^\star, I_{\bar{S}}^\star) = \frac{1}{q} L_{\bar{S}'}$ is also a (scaled) ideal lattice, for some $\bar{S}' \subseteq \{1, \ldots, k_q\}$ with $|\bar{S}'| = |\bar{S}|$, where we have used the fact that the mapping sending $a_1(x) = \prod_{i \in \bar{S}} \Phi_i(x)$ to $a_1^\star(x)$ induces a bijection on the factors $\Phi_i(x)$. Since $\det L_{\bar{S}'} = q^{n|\bar{S}|/k_q}$, we have by Minkowski's theorem that $\lambda_1^\infty(L_{\bar{S}'}) \leq q^{|\bar{S}|/k_q}$. Moreover, since $I_S + q\mathbb{Z}^n$ is an ideal lattice, Lemma 2.1 gives that $\eta_\delta(I_S + q\mathbb{Z}^n) \leq \frac{1}{q} \sqrt{\ln(2n(1 + 1/\delta))/\pi} \cdot \lambda_1^\infty(L_{\bar{S}'}) \leq \sqrt{n \ln(4nq)} q^{|S|/k_q} \leq \sigma$, for $\delta := q^{-n/2}$, assuming $|S| \leq k_q/2$. Hence, for a term of (2) with $|S| \leq k_q/2$, by Lemma 2.4, we have $|D_{\mathbb{Z}^n,\sigma,-z_i}(I_S + q\mathbb{Z}^n) - q^{-n|S|/k_q}| \leq 2\delta$.

For a term of (2) with $|S| > k_q/2$, we choose $S' \subseteq S$ with $|S'| = \lfloor k_q/2 \rfloor \geq k_q/3$ for $k_q \geq 2$. By using with $S'$ the above result for small $|S|$, we obtain $D_{\mathbb{Z}^n,\sigma,-z_i}(I_S + q\mathbb{Z}^n) \leq D_{\mathbb{Z}^n,\sigma,-z_i}(I_{S'} + q\mathbb{Z}^n) \leq 2\delta + q^{-n/3}$.

Overall, we have:

$$\left| D_{\mathbb{Z}^n,\sigma}(z_i + R_q^\times + q\mathbb{Z}^n) - \sum_{k=0}^{k_q} (-1)^k \binom{k_q}{k} q^{-k} \right| \leq 2^{n+1}\delta + 2 \sum_{k=\lceil k_q/2 \rceil}^{k_q} \binom{k_q}{k} q^{-n/3} \leq 2^{n+1}(\delta + q^{-n/3}),$$

which leads to the desired bound on $\delta_i$. This completes the proof of the theorem. □

### 3.4 A revised NTRUEncrypt scheme

In this section we present the provably secure variant of the NTRUEncrypt scheme. We define the scheme NTRUEncrypt with parameters $n, q, p, \alpha, \sigma$ 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 $\mathcal{P}$ as $\sum_{0 \leq i < n} \varepsilon_i x^i p$, 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 $\leq \frac{1}{2}\sqrt{\deg(p) + 1} \cdot \|p\|$. The parameter $\alpha$ is the R-LWE noise distribution parameter. Finally, the parameter $\sigma$ is the standard deviation of the discrete Gaussian distribution used in the key generation process (see Section 3.3).

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

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

The correctness conditions for the scheme are summarized below.

**Lemma 3.7.** *If* $\deg p \leq 1$ $\omega(n^{0.25} \ln n)\alpha\|p\|^2\sigma < 1$, *and* $\alpha q \geq n^{0.75}$, *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 3.6, we know that with probability $\geq 1 - 2^{-n+3}$ both $f$ and $g$ have Euclidean norms $\leq 4\sqrt{n}\|p\|\sigma$ if $\deg p \leq 1$. This implies that $\|pf\|, \|pg\| \leq 8\sqrt{n}\|p\|^2\sigma$, with probability $\geq 1 - 2^{-n+3}$. From Lemma 2.10, both $pfs$ and $pge$ have infinity norms $\leq 8\alpha q n^{0.25}\omega(\ln n) \cdot \|p\|^2\sigma$ with probability $1 - n^{-\omega(1)}$. Independently, we have:

$$\|fM\|_\infty \leq \|fM\| \leq \sqrt{n}\|f\|\|M\| \leq 4n\|p\|^2\sigma.$$

Since $\alpha q \geq n^{0.75}$, we conclude that $\|C''\|_\infty \leq 20\alpha q n^{0.25}\omega(\ln n) \cdot \|p\|^2\sigma$, with probability $1 - n^{-\omega(1)}$. □

The security of the scheme follows by a elementary reduction from the decisional R-LWE$^\times_{\text{HNF}}$, exploiting the uniformity of the public key in $R_q^\times$ (Theorem 3.2), and the invertibility of $p$ in $R_q$.

**Lemma 3.8.** *Suppose that $n$ is a power of 2 such that $\Phi = x^n + 1$ splits into $n$ linear factors modulo prime $q \geq 5$. Let $\varepsilon \in (0, 1/3)$, $\delta > 0$, $p \in R_q^\times$ and $\sigma \geq n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+\varepsilon}$. 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$^\times_{\text{HNF}}$ *with parameters $q$ and $\alpha$ 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$^\times_{\text{HNF}}$ that runs as follows, given oracle $\mathcal{O}$ that samples from either $U(R_q^\times \times R_q)$ or $A^\times_{s,\psi}$ for some previously chosen $s \hookleftarrow \psi$ and $\psi \hookleftarrow \overline{\Upsilon}_\alpha$. Algorithm $\mathcal{B}$ first calls 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 3.2, 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 $\psi$, 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 $A^\times_{s,\psi}$, 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 3.7 and 3.8 with Theorem 2.2 we obtain our main result.

**Theorem 3.3.** *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 $q^{\frac{1}{2}-\varepsilon} = \omega(n^{2.25}\ln^2 n)\|p\|^2$, with $\varepsilon = \omega(1/n)$ and $\varepsilon < 1/3$ and $p \in R_q^\times$ with $\deg(p) \leq 1$. Let $\sigma = n\sqrt{\ln(8nq)} \cdot q^{\frac{1}{2}+\varepsilon}$ and $\alpha^{-1} = \omega(n^{0.25}\ln n)\|p\|^2\sigma$. If there exists an IND-CPA attack against* `NTRUEncrypt` *which runs in time $\mathcal{P}oly(n)$ and has success probability $1/2 + 1/\mathcal{P}oly(n)$, then there exists a $\mathcal{P}oly(n)$-time quantum algorithm for $\gamma$-*Ideal-SVP *with $\gamma = \omega(n^{2.75}\ln^{2.5} n)\|p\|^2 q^{\frac{1}{2}+\varepsilon}$. Moreover, the decryption algorithm succeeds with probability $1 - n^{-\omega(1)}$.*

Overall, by choosing $\varepsilon = 1/(\ln n)$, the smallest $q$ for which the analysis holds is $\widetilde{\Omega}(n^{4.5})$, and the smallest $\gamma$ that can be obtained is $\widetilde{O}(n^5)$. Finally, we observe that our proof can be readily adapted to offer security against sub-exponential attackers, under the assumption that Ideal-SVP cannot be solved in quantum sub-exponential time for some polynomial approximation factor $\gamma$.

## 4 A provably secure variant of `NTRUSign`

In this section, we present our provably secure variant of the `NTRUSign` signature scheme. The key generation algorithm for our scheme extends the `NTRUEncrypt` secret key $(f, g)$ by computing another (linearly independent over the ring $R$) ‘short’ pair $(F, G)$ satisfying $F_1 h - G_1 = 0 \bmod q$, such that a full short basis matrix $M = \begin{bmatrix} f & g \\ F & G \end{bmatrix}$ is obtained for the NTRU $R$-module $L$ generated by the rows of the matrix $\begin{bmatrix} 1 & h \\ 0 & q \end{bmatrix}$. The method we use for generating $(F, G)$ is a variant of the original `NTRUSign` key generation algorithm in [25, 26].

Since the determinant of the lattice corresponding to the module $L$ is $q^n$, to make $M$ a basis of $L$, it suffices for the module vector $(F, G)$ to satisfy the determinant condition $fG - gF = q$. The main idea in [25, 26] for generating $(F, G)$ is the observation that if $f, g$ are coprime over $R$, then it is easy to compute $(F_1, G_1) \in R^2$ such that $fG_1 - gF_1 = 1$, and this vector can be easily lifted to a module vector $(F_q, G_q) = q \cdot (F_1, G_1)$ such that $M$ is a basis for $L$. Although $(F_q, G_q)$ is not short, thanks to the coprimality condition $fG_1 - gF_1 = 1$, it can be easily made short by length reduction, i.e., by subtracting from it a multiple $\approx qF_1/f$ of the given vector $(f, g)$ to get a vector $(F, G) \approx (F_q, G_q) - (qF_1/f) \cdot (f, g) = (0, q \cdot (G_1 - gF_1/f)) = (0, q/f)$, where $\|q/f\| \approx q/\|f\|$ is short. However, this procedure fails if $f$ and $g$ are not coprime over $R$, and this undesirable event is dealt with in the key generation procedure by rejecting $(f, g)$ and resampling new random candidates for $(f, g)$ until the coprimality condition holds. Since this rejection probability $p$ contributes a deterioration factor $\frac{1}{1-p}$ in the expected key generation time, and also to the security reduction cost (with respect to the close-to-uniform distribution of the public key $g/f$ when $(f, g)$ is sampled without rejection), it is important to bound $1 - p$ from below by a non-negligible function.

In [25, 26], the key generation algorithm, and in particular, the coprimality probability $1 - p$, are not rigorously analyzed. Here, we rigourously bound the coprimality probability $1 - p$ when $f$ and $g$ are independently sampled from a discrete Gaussian distribution $D_{\mathbb{Z}^n, \sigma}$ over $R$. Our argument is based on a generalization of the classical analysis of the probability $1 - p$ that two “random” integers are coprime, which gives the asymptotic value $1 - p = \prod_q (1 - 1/q^2) = \zeta(2)^{-1}$, where $\zeta(2) = \prod_q \frac{1}{1 - q^{-2}} = \frac{\pi^2}{6}$ is Riemann’s zeta function evaluated at 2, and the products run over all prime integers $q$. Our generalization of this analysis to the ring $R$ leads us to study the value of $\zeta_K(2)$, where $\zeta_K(2) = \prod_J \frac{1}{1 - \mathcal{N}(J)^{-2}}$ is the Dedekind zeta function for $K = \mathbb{Q}[x]/(x^n + 1)$, evaluated at 2, and the product now runs over all prime ideals $J$ of $R = \mathbb{Z}[x]/(x^n + 1)$. We show that $\zeta_K(2) = O(1)$ and, using some additional results on $\zeta_K$, that $1 - p \geq \frac{1}{2\zeta_K(2)} - o(1)$, so the acceptance probability $1 - p$ is in fact lower bounded by a constant.

As a further improvement on the key generation algorithm in [25, 26], we apply Babai’s nearest plane algorithm [6] to reduce the length the extended vector $(F, G)$, rather than applying Babai’s roundoff method as described above. This allows us to save a $\approx \sqrt{n}$ factor in the norm of $(F, G)$.

### 4.1 Additional results on ideal lattices

For the analysis of the key generation of the signature scheme (in Subsection 4.2), 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 $\sigma$. Since $b \cdot b^{-1} = 1$ over $K$, it is reasonable to expect $\|b^{-1}\|$ to be proportional to $\sigma^{-1}$.

**Lemma 4.1.** *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 2.7 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 now review some facts about the Dedekind zeta function (see, e.g., [51, Ch. VII]). 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 $I \neq 0$ strictly contained in $R$, where the $J_i$'s are distinct prime ideals in $R$ and $e_i$ is a positive integer 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 $i$. We extend the definition to $I = R$ by setting $\mu(R) = 1$. The Dedekind zeta function of the ring $R$ of integers of $K$ is the function $\zeta_K : \mathbb{R} \to \mathbb{R}$ defined by

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

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

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

where the product is over all *prime* ideals of $R$ and the sum is over all non-zero *integral* ideals of $R$.

**Lemma 4.2.** *Let $K_n = \mathbb{Q}[x]/\Phi_n$, for $n \geq 4$ 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 prime integer $p$, 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 $\Delta_K$ a power of 2, all the prime ideals above $p > 2$ have identical norm $p^{n/\pi_K(p)}$ (see, e.g., [50, 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 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:

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

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 [68, Th. 9, p. 16] that $\sum_{p\leq x} p^{-1} = \ln\ln x + c + O(1/\ln x)$, for some constant $c$. We thus obtain that:

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

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

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

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:

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

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

In our study of the Dedekind zeta function, we use the following bound.

**Lemma 4.3.** *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 $k$. Note 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 4.2 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 4.4.** *Assume that $\sigma \geq 7n^{1.5}\ln^{1.5} n$.Then, for $n$ a sufficiently large power of 2:*

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

*Proof.* By Lemma 2.3, we have:

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

We bound $\Pr[\langle f,g\rangle \neq R \ \wedge \ \|f\|,\|g\| \leq \sqrt{n}\sigma]$ by using an argument adapted from [63]. 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 $\mathcal{N}(I) \leq \mathcal{N}(\langle f\rangle) \leq (\sqrt{n}\sigma)^n$. Thus, we have $\Pr[\langle f,g\rangle \neq R \ \wedge \ \|f\|,\|g\| \leq \sqrt{n}\sigma] \leq 1-p$, with:

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

where in the second equality, we used the inclusion-exclusion principle (and $\mu$ is the Möbius function for ring $R$), and $D^T_{\sigma,\mathbb{Z}^n}$ denotes the truncation of $D_{\sigma,\mathbb{Z}^n}$ to the ball $B_n(\sqrt{n}\sigma)$ of radius $\sqrt{n}\sigma$, i.e. $D^T_{\sigma,\mathbb{Z}^n}(\boldsymbol{x}) = D_{\sigma,\mathbb{Z}^n}(\boldsymbol{x})$ if $\boldsymbol{x} \in B_n(\sqrt{n}\sigma)$ and $D^T_{\sigma,\mathbb{Z}^n}(\boldsymbol{x}) = 0$ otherwise.Recall that $\zeta_K(2)^{-1} = \sum \mu(I) \cdot \mathcal{N}(I)^{-2}$, where the sum is over all ideals $I \subseteq R$.We now show that $\left|p - \zeta_K(2)^{-1}\right| \leq (2\zeta_K(2))^{-1}$. This implies $p \geq (2\zeta_K(2))^{-1}$, as required. We have:

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

To bound the first sum, we recall that for any (even fractional) ideal $I$, we have $\lambda_n(I) = \lambda_1(I) \leq \sqrt{n}\mathcal{N}(I)^{1/n}$, so. for any $\delta \in (0,1/2)$, the smoothing parameter $\eta_\delta(I)$ is no greater than $B_\delta \cdot \mathcal{N}(I)^{1/n}$, where $B_\delta = \sqrt{n\ln(2n(1+1/\delta))/\pi}$ (by Lemma 2.1). It follows from Lemma 2.2 that $\left|D_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| \leq 18\delta/\mathcal{N}(I)^2$ if $\mathcal{N}(I) \leq (\sigma/B_\delta)^n$ and $I \subseteq R$. We have $|D_{\mathbb{Z}^n,\sigma}(I) - D^T_{\mathbb{Z}^n,\sigma}(I)| = D_{\mathbb{Z}^n,\sigma}(I \setminus B_n(\sqrt{n}\sigma)) = D_{I,\sigma}(I \setminus B_n(\sqrt{n}\sigma)) \cdot D_{\mathbb{Z}^n,\sigma}(I) \leq 2^{-n+2} \cdot D_{\mathbb{Z}^n,\sigma}(I)$, where in the last inequality we applied Lemma 2.3. We conclude that $\left|D^T_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| \leq (18\delta + 2^{-n+5})/\mathcal{N}(I)^2$ for $I \subseteq R$ of norm $\leq (\sigma/B_\delta)^n$. Assume now that $(\sigma/B_\delta)^n < \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n$, and let $k = \left\lceil \frac{\mathcal{N}(I)^{1/n}}{\sigma/B_\delta} \right\rceil$. Since

$I \subseteq \frac{1}{k} \cdot I$, we have $D^T_{\mathbb{Z}^n,\sigma}(I) \leq D^T_{\mathbb{Z}^n,\sigma}(\frac{1}{k} \cdot I)$. Also, by the choice of $k$, we have $\eta_\delta(\frac{1}{k} \cdot I) = \frac{1}{k}\eta_\delta(I) \leq \sigma$. Now:

$$D^T_{\mathbb{Z}^n,\sigma}\left(\frac{1}{k} \cdot I\right) \leq D_{\mathbb{Z}^n,\sigma}\left(\frac{1}{k} \cdot I\right) = \frac{\rho_\sigma(\frac{1}{k} \cdot I \cap \mathbb{Z}^n)}{\rho_\sigma(\mathbb{Z}^n)} \leq \frac{\rho_\sigma(\frac{1}{k} \cdot I)}{\rho_\sigma(\mathbb{Z}^n)} \leq \left(\frac{2B_\delta}{\sigma}\right)^n \frac{1+\delta}{1-\delta},$$

where in the last inequality we applied Lemma 2.2 twice, using $\sigma \geq \eta_\delta(\mathbb{Z}^n)$ and $\det(\frac{1}{k} \cdot I) = \frac{1}{k^n} \cdot \mathcal{N}(I) \geq (\frac{\sigma}{2B_\delta})^n$. Therefore, we have $D^T_{\mathbb{Z}^n,\sigma}(I)^2 \leq (\frac{2B_\delta}{\sigma})^{2n}(\frac{1+\delta}{1-\delta})^2$. Finally, assuming that $\sigma \geq 2B_\delta$ and $\delta = \frac{1}{160\zeta_K(2)^2}$, we obtain:

$$\begin{aligned}
\sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} \left|D^T_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| &\leq \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sigma/B_\delta)^n}} \left|D^T_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| + \sum_{\substack{I \subseteq R \\ (\sigma/B_\delta)^n < \mathcal{N}(I) \leq (\sqrt{n}\sigma)^n}} \left|D^T_{\mathbb{Z}^n,\sigma}(I)^2 - \mathcal{N}(I)^{-2}\right| \\
&\leq (18\delta + 2^{-n+5}) \ \cdot \sum_{\substack{I \subseteq R \\ \mathcal{N}(I) \leq (\sigma/B_\delta)^n}} \mathcal{N}(I)^{-2} \ + \ 2 \cdot H((\sqrt{n}\sigma)^n) \cdot \left(\frac{2B_\delta}{\sigma}\right)^{2n} \\
&< \frac{1}{8\zeta_K(2)} \ + \ 2 \cdot H((\sqrt{n}\sigma)^n) \cdot \left(\frac{2B_\delta}{\sigma}\right)^{2n} + o(1),
\end{aligned}$$

where in the last inequality we used the choice of $\delta$ and the fact that $\sum_{I \subseteq R, N(I) \leq (\sigma\sqrt{n})^n} N(I)^{-2} \leq \sum_{I \subseteq R} N(I)^{-2} = \zeta_K(2)$. Recall that $H(N)$ is the number of (integral) ideals of $R$ of norm $\leq N$. From Lemma 4.3 with $\varepsilon = \frac{\ln \ln n}{\ln n}$, we know that $H(N) \leq 2\exp(\frac{4n}{\ln \ln n}) \cdot N^{1+\varepsilon}$. Taking $\sigma \geq 7n^{1.5} \ln^{1.5} n$ provides $H\left((\sqrt{n}\sigma)^n\right) \cdot \left(\frac{2B_\delta}{\sigma}\right)^{2n} \leq \frac{1}{8\zeta_K(2)}$, for sufficiently large $n$, using $\zeta_K(2) = O(1)$ from Lemma 4.2. Overall, the first sum is $\leq \frac{1}{4\zeta_K(2)}$ for $n$ sufficiently large.

We now bound the second sum, as follows:

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

where we used the bound on $H(k)$ from Lemma 4.3 with $\varepsilon = \frac{\ln \ln n}{\ln n}$. Now, notice that the summand is $\leq \frac{2}{k^{2-\varepsilon}}$, which allows us to bound the second sum by $O(\exp(\frac{4n}{\ln \ln n}) \cdot (\sqrt{n}\sigma)^{-(1-\varepsilon)n}) = o(1)$, so the latter is $\leq (4\zeta_K(2))^{-1}$ for sufficiently large $n$, which completes the proof. □

### 4.2 A revised NTRUSign key generation algorithm

The revised key generation for NTRUSign is given in Fig. 3. It is inspired from the algorithm of [26, Se. 4] and described in more details in [25, Se. 5]. The vector $(f, g)$ produced by the NTRUEncrypt

```
Inputs: $n, q \in \mathbb{Z}$, $\sigma > 0$.
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 $fG_1 - gF_1 = 1$, e.g., using a Hermite Normal Form algorithm (see [9, Ch. 2.4]); set $F_q := qF_1$ and $G_q := qG_1$.
6. Use Babai's nearest plane algorithm [6] 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. 3.** Revised key generation algorithm for NTRUSign.

key generation algorithm is a short vector in the $R$-module generated by the rows of the matrix $\begin{bmatrix} 1 & h \\ 0 & q \end{bmatrix}$ with $h = g/f \bmod q$. The algorithm of Fig. 3 extends $(f, g)$ into a short module basis $\begin{bmatrix} f & g \\ F & G \end{bmatrix}$.

Because of the rejection tests, the output public key $h$ may not be uniformly distributed in $R_q^\times$, as it was the case for NTRUEncrypt. Uniformity is important for us to be able to eventually rely on Theorem 2.1 to prove the security of the signature scheme. In fact, as we will show in Subsection 4.3, it suffices that the combined rejection probabilities of Steps 3, 4 and 7 is non-negligibly away from 1.

By Lemma 4.4, when no rejection is performed in Steps 1–3, the rejection probability of Step 4 is (assuming that $\sigma \geq 7n^{1.5} \ln^{1.5} n$ and that $n$ is a sufficiently large power of 2):

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

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

**Lemma 4.5.** *Assume that $\sigma \geq 7n^{1.5} \ln^{1.5} n$. Then, as $n$ grows to infinity:*

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

*where $F$ and $G$ are as defined in Steps 5 and 6 of the algorithm of Figure 3.*

*Proof.* We decompose $(F, G)$ as $(F, G) = (F_q, G_q)^* + (e_f, e_g)$, where $(F_q, G_q)^*$ is the projection of $(F_q, G_q)$ orthogonally to the $K$-span of $(f, g)$ (which can also be viewed as the projection of $(F_q, G_q)$ orthogonally to the $\mathbb{Q}$-span of $(f, g), (xf, xg), \ldots, (x^{n-1}f, x^{n-1}g)$). We have:

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

As we use Babai's nearest-plane algorithm, the vector $(e_f, e_g)$ 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 $\|(F,G)^*\| \leq \|(0, qf^{-1})\| = q\|f^{-1}\|$. From Lemma 4.1 with $t = \omega(n)$ and $I = R$, we get:

$$\Pr_{f \hookleftarrow D_{R,\sigma}}\left[\|f^{-1}\| \geq \frac{\omega(\sqrt{n})}{\sigma}\right] \leq o(1).$$

This remains the case when the event is conditioned on $\langle f, g\rangle = R$, as, by Lemma 4.4, the probability that $\langle f, g\rangle = R$ is bounded from below by a constant. Overall, we have that $\|(F_q, G_q)^*\| \leq \frac{q\omega(\sqrt{n})}{\sigma}$ holds except with probability $\leq o(1)$.

To 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 2.3, we have $\|(f,g)\| \leq \sqrt{2n}\sigma$ with probability $\geq 1 - o(1)$, when $f, g \hookleftarrow D_{R,\sigma}$. For the same reason as above, this remains the case when conditioning on $\langle f, g\rangle = R$. Overall, we have $\|(e_f, e_g)\| \leq \frac{n\sigma}{\sqrt{2}}$, except with probability $\leq o(1)$. This completes the proof. □

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

**Lemma 4.6.** *Assume that $q \geq 64n\zeta_K(2)$ and $\sigma = \omega(\max(\sqrt{n \ln n} \cdot q^{1/k_q}, q^{1/2}n^{-1/4}, n^{3/2}\ln^{3/2} n))$, where $k_q$ the number of irreducible factors of $x^n + 1$ modulo $q$. Then if $n$ is sufficiently large, the combined rejection probability of Steps 3, 4 and 7 of the algorithm of Fig. 3 (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 let $p_i$ denote 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 $f, g \hookleftarrow D_{R,\sigma}^\times$.
- $p_4$ is the probability that $\langle f, g\rangle \neq R$ and $\|f\|, \|g\| \leq \sqrt{n}\sigma$, with $f, g \hookleftarrow D_{R,\sigma}^\times$.
- $p_7$ is the probability that $(\|F, G\|) > n\sigma$, $\langle f, g\rangle = R$ and $\|f\|, \|g\| \leq \sqrt{n}\sigma$, with $f, g \hookleftarrow D_{R,\sigma}^\times$.

For $i \in \{3, 4, 7\}$, we define $p_i'$ as $p_i$ except that $f$ and $g$ are independently sampled from $D_{R,\sigma}$ rather than $D_{R,\sigma}^\times$. 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 3.3. Indeed, by Lemma 3.5 and the choice of $\sigma$ and $q$, we have $p_1 \leq \frac{1}{32\zeta_K(2)}$. Lemmata 2.1 and 2.3 and the choice of $\sigma$ imply that $p_3' \leq 2^{-n+2}$. Finally, from the choice of $\sigma$ and Lemmata 4.4 and 4.5, we have that $p_4' \leq 1 - \frac{1}{2\zeta_K(2)} + o(1)$ and $p_7' = o(1)$. Recall from Lemma 4.2 that $\zeta_K(2) = O(1)$ when $n$ grows to infinity. Therefore, 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 satisfies $p_3 + p_4 + p_7 \leq \frac{p_3' + p_4' + p_7'}{1 - 2p_1} \leq 1 - \frac{1}{8\zeta_K(2)}$, as required. □

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

**Theorem 4.1.** *Let $n$ be a power of 2 such that $\Phi = x^n + 1$ splits into $k_q \in \{2, n\}$ irreducible factors modulo prime $q \geq 64\zeta_K(2)n$. Let $\varepsilon \in (0, 1/3)$ and $\sigma \geq \max(n\sqrt{\ln(8nq)} \cdot q^{1/2+\varepsilon}, \omega(n^{3/2}\ln^{3/2} n))$ if $k_q = n$, or $\sigma \geq \max(\sqrt{n\ln(8nq)} \cdot q^{1/2+\varepsilon}, \omega(n^{3/2}\ln^{3/2} n))$ if $k_q = 2$. Then the algorithm of Fig. 3 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 $\begin{bmatrix} 1 & h \\ 0 & q \end{bmatrix}$ with $h = g/f \bmod q$. 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 $\leq 2^{10n}q^{-\lfloor \varepsilon n \rfloor}$.*

*Proof.* The first statement is provided by Lemma 4.6. For the second statement, we refer to [26, Th. 1]. The norm inequalities are obvious from the description of the algorithm. Finally, the last statement is provided by Theorem 3.2 and Lemma 4.6. □

### 4.3 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 instanciation of the Gentry et al. signature [21], 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 Gentry et al. signature is built from a general cryptographic primitive introduced in [21] and called *Collision-Resistant Preimage Sampleable Functions* (CRPSF), which we recall.

**Definition 4.1 (CRPSF).** *A CRPSF is specified by three probabilistic polynomial-time algorithms* $(\mathsf{TrapGen}, \mathsf{SampleDom}, \mathsf{SamplePre})$ *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 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.*
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 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 $\omega(\ln n)$.*
5. Collision-Resistance without Trapdoor: *For any probabilistic polynomial-time algorithm $\mathsf{F}$, 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, where the probability is taken over the choice of $a$ and the random coins of $\mathsf{F}$.*

Our CRPSF construction $\mathtt{NTRUPSF}(n,q,\sigma,s)$ is shown in Fig. 4. The parameters $n$ and $q$ define the rings $R$ and $R_q$. The parameter $\sigma$ 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.

**Theorem 4.2.** *Suppose $n$ is a power of $2$ such that $\Phi = x^n + 1$ splits into $k_q \in \{2, n\}$ irreducible factors modulo prime $q = \mathcal{P}oly(n)$, with $\sigma = n\sqrt{\ln(8nq)} \cdot q^{1/2+\varepsilon}$ and $q^{1/2-\varepsilon} = \widetilde{\Omega}(n^{7/2})$ if $k_q = n$, or $\sigma = \sqrt{n\ln(8nq)} \cdot q^{1/2+\varepsilon}$ and $q^{1/2-\varepsilon} = \widetilde{\Omega}(n^3)$ if $k_q = 2$, for some fixed $\varepsilon \in (0, \frac{\ln n}{\ln q})$. Let $s = \widetilde{\Omega}(n^{3/2}\sigma)$. Then the construction $\mathtt{NTRUPSF}(n,q,\sigma,s)$ from Fig. 4 is a CRPSF secure against*

- **Generating a Function with Trapdoor** – $\mathsf{TrapGen}(1^n, q, \sigma)$: Run the `NTRUSign` key generation algorithm from Fig. 3, using $n, q, \sigma$ 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 $\mathcal{D}_n = \{\boldsymbol{z} \in R^2 : \|\boldsymbol{z}\| \le s\sqrt{2n}\}$ 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** – $\mathsf{SamplePre}(sk, t)$: 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 2.9. Return $\boldsymbol{z}$.

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

*$\mathcal{P}oly(n)$ time algorithms, assuming the hardness of* $\gamma$-Ideal-SVP *against $\mathcal{P}oly(n)$ time algorithms, with $\gamma = \widetilde{O}(n \cdot s)$.*

*Proof.* The sets $\mathcal{D}_n$ and $\mathcal{R}_n$ are easily recognizable. Observe that $s \ge \max(\sqrt{n}, \eta_{1/2}(\mathbb{Z}^{2n}))$, so by Lemmata 2.3 and 2.6, the distribution of $\boldsymbol{z} = (z_1, z_2)$ returned by $\mathsf{SampleDom}$ is within negligible statistical distance of $D_{\mathbb{Z}^{2n},s}$. To show Property 2 of Definition 4.1, we apply Theorem 3.1 with $\delta = n^{-\omega(1)}$ to conclude that thanks to the choice of $s$, except for a fraction $\le 2^{8n}q^{-2\varepsilon n}$ of $(a_1, a_2) \in (R_q^\times)^2$, we have $\Delta(a_1z_1 - a_2z_2; U(R_q)) \le 2\delta$ 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 $\Delta(a_1z_1 - a_2z_2; U(R_q)) = \Delta(a_1a_2^{-1}z_1 - z_2; U(R_q))$ 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 $\Delta(hz_1 - z_2; U(R_q)) \le 2\delta$ with $(z_1, z_2) \hookleftarrow D_{\mathbb{Z}^{2n},s}$ except for a fraction $\le 2^{8n}q^{-2\varepsilon n}$ of $h \in R_q^\times$. Finally, by Theorem 4.1, the distribution $D_h$ of $h = g/f$ generated by $\mathsf{TrapGen}$ is obtained by rejection with constant rejection probability $c < 1$ from a distribution within statistical distance $2^{10n}q^{-\lfloor \varepsilon n \rfloor}$ of $U(R_q^\times)$. It follows that $\Delta(hz_1 - z_2; U(R_q)) \le 2\delta$ with $(z_1, z_2) \hookleftarrow D_{\mathbb{Z}^{2n},s}$ except with probability $\le \frac{1}{1-c} \cdot (2^{8n}q^{-2\varepsilon n} + 2^{10n}q^{-\lfloor \varepsilon n \rfloor}) = q^{-\Omega(n)}$ over the choice of the public key $h$, as required.

To show Property 3 of Definition 4.1, 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 $F(\boldsymbol{z}) = \frac{\rho_s(\boldsymbol{z})}{\rho_s(h^\perp + \boldsymbol{c})} = D_{h^\perp + \boldsymbol{c},s}(\boldsymbol{z})$, where $\boldsymbol{c} = (1, h - t)$. Therefore, Property 3 follows from Lemma 2.9, the upper bound $n\sigma$ on the trapdoor basis norm from Theorem 4.1, and the choice of $s = \omega(n^{3/2}\sqrt{\ln n} \cdot \sigma)$.

To show Property 4 of Definition 4.1, note 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 bound the min-entropy of $D_{h^\perp,s,-\boldsymbol{c}}$ from below. By Lemma 2.5, the latter min-entropy is $\Omega(n)$ if the condition $s \ge 2\eta_{1/2}(h^\perp)$ is satisfied. Theorem 3.1 shows that for all except a fraction $\le 2^{8n}q^{-\varepsilon n} = q^{-\Omega(n)}$ of $\boldsymbol{a} \in (R_q^\times)^2$, we have $\eta_{1/2}(\boldsymbol{a}^\perp) = \widetilde{O}(\sqrt{n}q^{\frac{1}{2}+\varepsilon})$. Since $\boldsymbol{a}^\perp = h^\perp$ with $h = a_2^{-1}a_1$, it follows that for all except a fraction $\le q^{-\Omega(n)}$ of $h \in R_q^\times$, we have $\eta_{1/2}(h^\perp) \le \widetilde{O}(\sqrt{n}q^{\frac{1}{2}+\varepsilon})$. By the choice of $s$, the condition $s \ge 2\eta_{1/2}(h^\perp)$ is satisfied. By Theorem 4.1, 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 4.1. Let $\mathcal{A}$ be a collision-finding algorithm for `NTRUPSF` with run-time $T = \mathcal{P}oly(n)$ and success probability $\delta = 1/\mathcal{P}oly(n)$ over the choice of the public key $h$ and the randomness of $\mathcal{A}$. By Theorem 4.1, 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 $\delta' = (1 - c)\delta - 2^{10n}q^{-\lfloor \varepsilon n \rfloor}$. Note that we have $\delta' = 1/\mathcal{P}oly(n)$. We construct an algorithm $\mathcal{A}'$ for R-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 $\geq (1 - 2n/q)\delta' = 1/\mathcal{P}oly(n)$. Applying Theorem 2.1 using the choice of $q = \widetilde{\Omega}(\beta\sqrt{n})$, we obtain a $\mathcal{P}oly(n)$ time algorithm for $\gamma$-Ideal-SVP with the claimed $\gamma$. □

**The revised `NTRUSign` scheme.** Given the `NTRUPSF` construction above, the revised `NTRUSign` follows the Gentry et al. '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 applying the Gentry et al. construction directly on `NTRUPSF` results in signatures on a message $M$ consisting of two ring elements $(\sigma_1, \sigma_2)$ 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 the signature length, our `NTRUSign` variant eliminates $\sigma_2$ 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`$(n, q, \sigma, s)$ 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 $(r, \sigma_1)$.
- **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`$(n, q, \sigma, s, k)$ from the `NTRUPSF` primitive in Fig. 4.

Since $\sigma_2$ is easily computed from $\sigma_1$, $r$ and the public information, the security of `NTRUSign` is equivalent to that of the Gentry et al. signature obtained from `NTRUPSF`, which in turn has been shown in [21, Prop. 6.2] to follow from the security of the underlying CRPSF. Combining with Theorem 4.2, we obtain our second main result.

**Corollary 4.1.** *Let $\varepsilon, n, q, \sigma, s$ satisfy the conditions in Theorem 4.2, and let $k = \omega(\ln n)$. Then, assuming the random oracle model for $\mathcal{H}$, the signature scheme* `NTRUSign`$(n, q, \sigma, s, k)$ *from Fig. 5 is strongly existentially unforgeable against a chosen message attack with $\mathcal{P}oly(n)$ run-time and $1/\mathcal{P}oly(n)$ success probability, assuming the hardness of* $\gamma$-Ideal-SVP *against $\mathcal{P}oly(n)$ time algorithms, with $\gamma = \widetilde{O}(n \cdot s)$.*

Note that if $\mathcal{H}$ runs in quasi-linear time, then so does the verification algorithm. Also, if precomputations are performed, then so does the signing algorithm (see [55, 13]). The amortized cost per signed bit is then $\widetilde{O}(1)$. Finally, we remark that the smallest $q$ and $\gamma$ that can be chosen in Theorem 4.2 and Corollary 4.1 are $\widetilde{\Omega}(n^{6/(1-2\varepsilon)})$ if $k_q = 2$ and $\widetilde{\Omega}(n^{7/(1-2\varepsilon)})$ if $k_q = n$. Finally, we observe that our proof can be readily adapted to offer security against sub-exponential attackers (in the random oracle model), under the assumption that Ideal-SVP cannot be solved in sub-exponential time for some polynomial approximation factor $\gamma$.

**Acknowledgements.** We thank Guillaume Hanrot, Vadim Lyubashevsky, Chris Peikert, Oded Regev, Igor Shparlinski, Joseph Silverman and Frederik Vercauteren for helpful discussions. We thank Jakob Breier for pointing out an error in the R-LWE noise sampler of the Eurocrypt proceedings version of this article. 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 (ARC) under Discovery Grant DP0987734, and partly by a Macquarie University Research Fellowship. Both authors were partly supported by ARC Discovery Grant DP110100628.

## References

1. S. Agrawal, D. Boneh, and X. Boyen. Efficient lattice (H)IBE in the standard model. In *Proc. of EUROCRYPT*, volume 6110 of *LNCS*, pages 553–572. Springer, 2010. Full version available from the authors upon request.
2. S. Agrawal, D. Boneh, and X. Boyen. Lattice basis delegation in fixed dimension and shorter-ciphertext hierarchical IBE. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 98–115. Springer, 2010.
3. M. Ajtai. Generating hard instances of lattice problems (extended abstract). In *Proc. of STOC*, pages 99–108. ACM, 1996.
4. S. Alaca and K.S. Williams. *Introductory Algebraic Number Theory.* Cambridge University Press, 2003.
5. 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.
6. L. Babai. On Lovász lattice reduction and the nearest lattice point problem. *Combinatorica*, 6:1–13, 1986.
7. D. Boneh, Ö. Dagdelen, M. Fischlin, A. Lehmann, C. Schaffner, and M. Zhandry. Random oracles in a quantum world. In *Proc. of ASIACRYPT*, pages 41–69, 2011.
8. 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.
9. H. Cohen. *A Course in Computational Algebraic Number Theory, 2nd edition.* Springer, 1995.
10. H. Cohen. *Advanced topics in computational number theory.* Springer, 2000.
11. D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *Proc. of Eurocrypt*, volume 1233 of *LNCS*, pages 52–61. Springer, 1997.
12. L. Ducas and A. Durmus. Ring-LWE in polynomial rings. In *Public Key Cryptography*, volume 7293 of *LNCS*, pages 34–51. Springer, 2012.
13. L. Ducas and P. Q. Nguyen. Faster gaussian lattice sampling using lazy floating-point arithmetic. In *Proc. of ASIACRYPT*, volume xxxx of *LNCS*, pages xxx–xxx. Springer, 2012.
14. C. Fieker and D. Stehlé. Short bases of lattices over number fields. In *Proc. of ANTS-IX*, volume 6197 of *LNCS*, pages 157–173. Springer, 2010.
15. E. Fogels. On the zeros of L-functions. *Acta Arith*, 11:67–96, 1965.
16. J. von zur Gathen and J. Gerhardt. *Modern Computer Algebra, 2nd edition.* Cambridge University Press, 2003.
17. C. Gentry. *A fully homomorphic encryption scheme.* PhD thesis, Stanford University, 2009. Manuscript available at `http://crypto.stanford.edu/craig`.
18. C. Gentry. Fully homomorphic encryption using ideal lattices. In *Proc. of STOC*, pages 169–178. ACM, 2009.
19. C. Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 116–137. Springer, 2010.
20. 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.
21. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *Proc. of STOC*, pages 197–206. ACM, 2008. Full version available at `http://eprint.iacr.org/2007/432.pdf`.
22. 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.
23. N. Higham. *Accuracy and Stability of Numerical Algorithms, 2nd edition.* SIAM, 2002.
24. J. Hoffstein, N. Howgrave-Graham, J. Pipher, and W. Whyte. Practical lattice-based cryptography: `NTRUEncrypt` and `NTRUSign`, 2009. Chapter of [?].
25. 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 [26]. Available at `http://www.securityinnovation.com/cryptolab/articles.shtml`.
26. 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.

27. 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.
28. 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.
29. 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.
30. N. Howgrave-Graham. A hybrid lattice-reduction and meet-in-the-middle attack against NTRU. In *Proc. of CRYPTO*, volume 4622 of *LNCS*, pages 150–169. Springer, 2007.
31. 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.
32. 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`.
33. IEEE P1363. Standard specifications for public-key cryptography. `http://grouper.ieee.org/groups/1363/`.
34. H. Iwaniec. On zeros of Dirichlet's L series. *Invent. Math.*, 23:97–104, 1974.
35. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *IACR Cryptology ePrint Archive*, 2012:090.
36. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math. Ann*, 261:515–534, 1982.
37. A. López-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *Proc. of STOC*, pages 1219–1234, 2012.
38. 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.
39. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *Proc. of ICALP*, volume 4052 of *LNCS*, pages 144–155. Springer, 2006.
40. 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.
41. 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.
42. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings, 2011. Draft for the extended version of [41], dated 24/04/2012.
43. V. Lyubashevsky, C. Peikert, and O. Regev. A toolkit for Ring-LWE cryptography, 2011. Draft, personal comunication.
44. D. Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Comput. Complexity*, 16(4):365–411, 2007.
45. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Kluwer Academic Press, 2002.
46. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput*, 37(1):267–302, 2007.
47. 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.
48. 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.
49. 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.
50. R. A. Mollin. *Algebraic Number Theory*. Chapman and Hall/CRC Press, 1999.
51. J. Neukirch. *Algebraic number theory*, volume 322 of *Grundlehren der Mathematischen Wissenschaften*. Springer, Berlin, 1999.
52. P. Q. Nguyen and O. Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *Journal of Cryptology*, 22(2):139–160, 2009.
53. C. Peikert. Limits on the hardness of lattice problems in $\ell_p$ norms. *Comput. Complexity*, 2(17):300–351, 2008.
54. C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *Proc. of STOC*, pages 333–342. ACM, 2009.
55. C. Peikert. An efficient and parallel gaussian sampler for lattices. In *Proc. of CRYPTO*, volume 6223 of *LNCS*, pages 80–97. Springer, 2010.
56. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *Proc. of TCC*, volume 3876 of *LNCS*, pages 145–166. Springer, 2006.

57. 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.
58. R. A. Perlner and D. A. Cooper. Quantum resistant public key cryptography: a survey. In *Proc. of IDtrust*, pages 85–93. ACM, 2009.
59. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *Proc. of STOC*, pages 84–93. ACM, 2005.
60. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.
61. O. Regev. The learning with errors problem, 2010. Invited survey in CCC 2010, available at `http://www.cs.tau.ac.il/~odedr/`.
62. C. P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theor. Comput. Science*, 53:201–224, 1987.
63. B. D. Sittinger. The probability that random algebraic integers are relatively $r$-prime. *Journal of Number Theory*, 130:164–171, 2010.
64. D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *Proc. of EUROCRYPT*, pages 27–47, 2011.
65. 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.
66. R. Steinfeld, S. Ling, J. Pieprzyk, C. Tartary, and H. Wang. NTRUCCA: How to Strengthen NTRUEncrypt to Chosen-Ciphertext Security in the Standard Model. In *Proc. of PKC*, pages 353–371, 2012.
67. 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.
68. G. Tenenbaum. *Introduction to Analytic and Probabilistic Number Theory.* Number 46 in Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1995.