On the hardness of the NTRU problem

Alice Pellet-Mary* and Damien Stehlé**

Abstract. The 25 year-old NTRU problem is an important computational assumption in public-key cryptography. However, from a reduction perspective, its relative hardness compared to other problems on Euclidean lattices is not well-understood. Its decision version reduces to the search Ring-LWE problem, but this only provides a hardness upper bound.

We provide two answers to the long-standing open problem of providing reduction-based evidence of the hardness of the NTRU problem. First, we reduce the worst-case approximate Shortest Vector Problem over ideal lattices to an average-case search variant of the NTRU problem. Second, we reduce another average-case search variant of the NTRU problem to the decision NTRU problem.

1 Introduction

In the NTRU encryption scheme [HPS98], the public key is an element h of a polynomial ring R_q that can be chosen as $\mathbb{Z}_q[x]/\Phi$ for some degree d monic irreducible polynomial Φ and some integer $q \geqslant 2$. This public key h is far from uniform in R_q, as it can be written as $h = f/g \bmod q$ where the secret key polynomials $f, g \in R = \mathbb{Z}[x]/\Phi$ have coefficients with small magnitudes compared to $\sqrt{q}$. In most concrete instantiations, such as the original scheme and the NTRU and NTRU Prime Round-3 candidates to the NIST post-quantum cryptography standardization project [CDH+20, BBC+20], the coefficients of f and g even belong to $\{-1, 0, 1\}$ and q grows as a small degree polynomial in d. As a result, the set of such h's is very sparse in R_q. The tasks of distinguishing h from uniform and recovering a sufficiently short pair (f, g) from h are respectively known as the decision and search variants of the NTRU problem.

The search NTRU problem can be solved with lattice reduction algorithms (such as [Sch87]), but to succeed for the most usual setting of $q \leqslant \mathrm{poly}(d)$, they require a computational effort growing as $\exp(O(d))$. In [KF15], Kirchner and Fouque described a heuristic algorithm with slightly subexponential cost $\exp(O(d/\log\log d))$ for the usual setting of $q \leqslant \mathrm{poly}(d)$ and $\|f\|_\infty, \|g\|_\infty \leqslant O(1)$. If the magnitude bound grows as $\Omega(\sqrt{d})$, then the cost of this algorithm is $\exp(O(d))$. In the completely different regime of very large q (but with $\|f\|$ and $\|g\|$ growing at a much smaller pace), recent works [ABD16, CJL16, KF17] have shown that the NTRU problem is significantly easier than previously thought. For example, one can recover appropriately distributed f, g with

* CNRS, Inria and Université de Bordeaux, alice.pellet-mary@math.u-bordeaux.fr
** ENS de Lyon and Institut Universitaire de France, damien.stehle@ens-lyon.fr

$\|f\|, \|g\| \leqslant \text{poly}(d)$ from h in quantum polynomial time when $q \geqslant \exp(\widetilde{\Omega}(\sqrt{d}))$. Prior to those works, it was believed that $q \geqslant \exp(\widetilde{\Omega}(d))$ was necessary for polynomial cost. This range of modulus q is far from the one used for the NTRU encryption scheme. However, NTRU instances with a large modulus q can occur in more advanced cryptographic constructions such as [LTV12] and [GGH13].

On the lower-bound front, it was shown in [SS11] for Φ a power-of-2 cyclotomic and extended in [WW18] to all cyclotomics that if f, g are Gaussian over R (restricted to elements that are invertible modulo q) with standard deviation that is a little larger than $\sqrt{q}$, then the distribution of $h = f/g \bmod q$ is within $2^{-\Omega(d)}$ statistical distance from the uniform distribution over invertible elements of R_q. This variant of decision NTRU is therefore vacuously hard. This parameter regime is relevant to the NTRU signature algorithm [HHP+03, SS13]. It also allows to obtain an NTRU-like public-key encryption scheme, but less efficient than with smaller secret key polynomials f, g.

Despite 25 years of study, little is known about the relationships between the NTRU problem variants and between them and other well-studied problems over Euclidean lattices. To our knowledge, the only exceptions are the direct reduction from decision NTRU to search NTRU and a reduction from decision NTRU to the search version of the Ring-LWE problem [SSTX09, LPR10], sketched in [Pei16, Se. 4.4.4]. Note that this only provides an *upper bound* to the hardness of the NTRU problem. Given this state of affairs, Peikert asked the following question in [Pei16, Se. 7.1]:

Is there a worst-case hardness reduction, or a search-to-decision reduction, for an NTRU-like problem?

Contributions. We provide positive answers to both components of the above question.

First, we give a reduction from the approximate Shortest Vector Problem restricted to ideal lattices (ideal-SVP) to a worst-case variant of the search NTRU problem. Combining the latter with the recent worst-case to average-case reduction for ideal-SVP from [dBDPW20] leads to a reduction from worst-case ideal-SVP to an average-case version of the search NTRU problem. The instance distribution is inherited from the distribution over ideal lattices considered in [dBDPW20]. We also show that this distribution over NTRU instances h can be efficiently sampled from, together with a corresponding trapdoor (f, g), if one has access to a quantum computer or if the modulus q is sufficiently large: this property allows to sample an NTRU encryption public key along with a corresponding secret key.

Second, we exhibit a reduction from another (average-case) variant of the search NTRU problem (see below) to the decision NTRU problem. The reduction works for a wide set of distributions for the search NTRU instances, and the decision NTRU instance distribution is directly derived from the considered search NTRU distribution. A sufficient condition on the search NTRU distribution is that it produces with overwhelming probability an h with trapdoor (f, g) such that f and g have balanced coefficients (in canonical embedding) and f

or g is coprime to q. This covers in particular the standard ternary distribution for f and g (i.e., $f, g \leftarrow \mathcal{U}(\{-1, 0, 1\}^d)$) provided we reject (f, g) when they are not balanced enough or not coprime to q (heuristically, this should happen with probability $\leqslant 1/2$). On the other hand, the choice of the decision NTRU distribution is much less flexible: even if we start with a ternary distribution for the search NTRU instances, it is very unlikely that the decision NTRU distribution we obtain is ternary. Similarly to the first reduction, we show that if the samples h from the search NTRU distribution can be efficiently sampled along with a corresponding trapdoor (f, g), then so can the samples from the resulting decision NTRU instance.

Technical overview. For the sake of simplicity, in the forthcoming discussion, we restrict ourselves to power-of-2 cyclotomic defining polynomials, i.e., $\Phi = x^d + 1$ for d a power of 2. In this case, the ring $R = \mathbb{Z}[x]/(x^d + 1)$ matches the ring of integers of the degree-d cyclotomic number field. Moreover, the coefficient embedding (which is the one usually considered in the NTRU literature) and the canonical embedding (used in this article) define the same geometry, up to scaling and rotation. (In the core of the paper, the results are presented for arbitrary number fields.)

To state the above contributions formally, we consider several variants of the NTRU problem. We say that $h \in R_q = \mathbb{Z}_q[x]/(x^d + 1)$ is an NTRU instance with gap γ if there exists $(f, g) \in R^2 \backslash \{(0, 0)\}$ such that $g \cdot h = f \bmod q$ and $\|f\|, \|g\| \leqslant \sqrt{q}/\gamma$. Note that writing $g \cdot h = f \bmod q$ rather than the more standard $h = f/g \bmod q$ allows one to consider g's that are not invertible modulo q and suffices for cryptographic applications. The norm $\|f\|$ is the Euclidean norm of the vector made of the coefficients of f, and the comparison to $\sqrt{q}$ is justified by the fact that for a uniformly chosen h, one expects the smallest such pair (f, g) to have Euclidean norm around $\sqrt{q}$, up to a small polynomial in d (in the core of the paper, we consider the Euclidean norm induced by the canonical embedding, which leads to a slightly different definition, differing by another $\sqrt{d}$ factor). In the literature, the bound on $\|f\|, \|g\|$ is often absolute rather than relative to $\sqrt{q}$: our definition variant stresses the distance to the uniform h regime. For a distribution $\mathcal{D}$ over NTRU instances with gap γ, the decision problem $(\mathcal{D}, \gamma, q)$-dNTRU consists in distinguishing between $\mathcal{D}$ and the uniform distribution over R_q. On the search NTRU side, the situation is more complex. We consider two variants of search NTRU, both of which with a worst-case and an average-case version. For $\gamma \geqslant \gamma'$, the worst-case vector NTRU problem $\text{wcNTRU}_{\text{vec}}$ consists, given as input an NTRU instance h with gap γ, in recovering $(f, g) \neq (0, 0)$ such that $g \cdot h = f \bmod q$ and $\|f\|, \|g\| \leqslant \sqrt{q}/\gamma'$. Note that if $h \in R_q$ has a trapdoor (f, g), then $(t \cdot f, t \cdot g)$ is another NTRU trapdoor of a possibly larger Euclidean norm, for any non-zero $t \in R$. The $\text{wcNTRU}_{\text{vec}}$ definition allows solutions whose norms are within an approximation factor γ/γ' from the norms of the promise. Even though there may be plenty of solutions of the form $(t \cdot f, t \cdot g)$ for $t \in R$, the pair ratio $h_{\mathbb{R}} = (tf)/(tg) = f/g$ over $K := \mathbb{Q}[x]/(x^d + 1)$ is an invariant. This motivates the definition of the worst-case module NTRU problem $\text{wcNTRU}_{\text{mod}}$, which consists in recovering $h_{\mathbb{R}}$

from h. This is equivalent to recovering the rank-1 submodule $(f,g)^T \cdot K \cap M_h$ of the rank-2 R-module $M_h = \{(f',g')^T \in R^2 : g' \cdot h = f' \bmod q\}$, hence justifying the name. The average-case counterparts to wcNTRU$_{\text{vec}}$ and wcNTRU$_{\text{mod}}$ are defined analogously.

We now sketch the reduction from ideal-SVP to wcNTRU$_{\text{vec}}$. Let us consider the worst-case variants, and the restriction of ideal-SVP to principal ideals with a known generator: we are given as input a generator z of a principal ideal $I = \langle z \rangle$ of R, and want to use a wcNTRU$_{\text{vec}}$ oracle to find a short non-zero vector in I. Any element $g \in I$ is of the form $g = z \cdot r$ for some $r \in R$. Consider a short non-zero $g \in I$. Multiplying it by q/z, we obtain that $g \cdot (q/z) = 0 \bmod q$. This already looks like an NTRU equation with a candidate q/z for h. But note that q/z is in $K = \mathbb{Q}[x]/(x^d+1)$ and has no a priori reason to belong to $R = \mathbb{Z}[x]/(x^d+1)$, whereas the element h of an NTRU instance must belong to R. To handle this difficulty, we can round q/z to R (coefficient-wise). This leads to $g \cdot \lfloor q/z \rceil = -g \cdot \{q/z\} \bmod q$, where both g and $f := -g \cdot \{q/z\}$ are small elements of R. We obtain the existence of a small pair $(f,g) \in R^2 \backslash \{(0,0)\}$ such that $g \cdot \lfloor q/z \rceil = f \bmod q$. We can then provide the element $h := \lfloor q/z \rceil$ to the wcNTRU$_{\text{vec}}$ oracle. The latter returns a pair $(f',g') \in R^2 \backslash \{(0,0)\}$ such that $g' \cdot \lfloor q/z \rceil = f' \bmod q$, and it can be proved that for any such sufficiently short pair, we have that g' is a short non-zero element of I. To handle possibly non-principal ideals (and also principal ideals with unknown generator), we rely on the 2-element representation of ideals.

If we forget polynomial factors and rely on a wcNTRU$_{\text{vec}}$ oracle with parameters q, γ and γ', the above allows to find γ_{svp} approximations to a shortest non-zero vector of an arbitrary ideal of volume $\leqslant N$ for $N^{1/d} \approx \sqrt{q}/\gamma$ and $\gamma_{\text{svp}} \approx \gamma/\gamma'$. Note that the reduction is worst-case to worst-case and handles bounded-volume ideals. To handle both limitations, we rely on the recent worst-case to average-case reduction for ideal-SVP from de Boer *et al* [dBDPW20]. By using the reduction with ideals from the average-case distribution from [dBDPW20], we obtain a reduction from worst-case ideal-SVP to average-case NTRU$_{\text{vec}}$. Further, the ideals from the average-case distribution from [dBDPW20] have volumes bounded as $\exp(O(d^2))$. This leads to q of the order of $\exp(O(d))$, which is significantly larger than in many applications. We refine the analysis of [dBDPW20] to show that by allowing the worst-case to average-case ideal-SVP reduction to run in time higher than polynomial in d, the average-case ideals from [dBDPW20] can be chosen with smaller volumes. The resulting NTRU modulus q is still slightly larger than polynomial, but it can be chosen as small as $d^{\omega(1)}$ if one considers sub-exponential time reductions.

We now provide an overview of our second main result, which is a reduction from average-case NTRU$_{\text{mod}}$ to dNTRU. This one is applicable for q larger than some moderate poly(d). At the core of the reduction is an NTRU rerandomization process. Assume we are given some $h \in R_q$ for which there exists a short pair $(f,g) \neq (0,0)$ with $g \cdot h = f \bmod q$. Now, for any $x_1, x_2 \in R$, we have $g \cdot (x_1 h + x_2) = x_1 f + x_2 g \bmod q$, which may be rewritten as $g \cdot h' = f' \bmod q$ with $h' = x_1 h + x_2$ and $f' = x_1 f + x_2 g$. Further, if x_1 and x_2 are short, then so

is f'. This hence gives a way to produce arbitrarily many NTRU samples with a common denominator g, from a single one. Our aim is to query the dNTRU oracle on many such samples, and gather relevant information to solve $\mathrm{NTRU}_{\mathrm{mod}}$. Concretely, we define the dNTRU distribution and show how to tweak the rerandomization process to be able to use the Oracle Hidden Center Problem (OHCP) framework from [PRS17]. At a high level, in the OHCP framework, one is given access to a decision oracle whose acceptance probabilities on a family of distributions $(\mathcal{D}_z)_{z\in\mathbb{C}}$ is a function of the distance $|z-c|$ to a hidden center $c \in \mathbb{C}$. Under some conditions on the oracle behaviour, there exists an efficient algorithm that recovers an arbitrarily accurate approximation $\widetilde{c}$ to c, by querying the OHCP oracle on samples from $\mathcal{D}_z$ for well-chosen values of z. Prior to this work, the OHCP framework has been used to provide a reduction from ideal-SVP to the decision version of Ring-LWE [PRS17], and a search to decision reduction for Ring-LWE [RSW18].

Let us now look more closely at the rerandomization of f. It was shown in [LSS14] that by sampling x_1 and x_2 from spherical Gaussians over R with standard deviation sufficiently above $\max(\|f\|, \|g\|)$, the distribution of $x_1 f + x_2 g$ is Gaussian over the ideal $\langle f\rangle + \langle g\rangle$ with a covariance matrix that is a function of f and g. This spherical Gaussian rerandomization defines our dNTRU distribution. We extend the proof of [LSS14] to show that if instead we sample x_1 and x_2 from correlated non-spherical Gaussians over R, then the distribution of $x_1 f + x_2 g$ is Gaussian over $\langle f\rangle + \langle g\rangle$ with a covariance matrix that can be made to depend solely on $|f(\zeta) - z\cdot g(\zeta)|$ for ζ an arbitrary complex root of $\Phi = x^d + 1$, and $z \in \mathbb{C}$ arbitrary. The center of the OHCP instance is $c = f(\zeta)/g(\zeta) = h_{\mathbb{R}}(\zeta)$ (recall that $h_{\mathbb{R}} = f/g$ belongs to $K = \mathbb{Q}[x]/(x^d+1)$). Using the dNTRU oracle within the OHCP framework hence allows us to recover an approximation to $h_{\mathbb{R}}(\zeta)$. In the applications from [PRS17, RSW18] of the OHCP framework, one recovers a vector $\mathbf{c}$ of OHCP centers from an approximation $\widetilde{\mathbf{c}}$ by observing that $\mathbf{c}$ belongs to a lattice: the exact center $\mathbf{c}$ can hence be obtained by simply rounding a sufficiently precise approximation $\widetilde{\mathbf{c}}$. In our case, we cannot proceed similarly, as $h_{\mathbb{R}}$ has rational coordinates. We instead show that the LLL algorithm [LLL82] can be used in a manner similar to [KLL84] to recover $h_{\mathbb{R}} = f/g \in K$ from a sufficiently precise approximation to $h_{\mathbb{R}}(\zeta)$, given an a priori upper bound to $\max(\|f\|, \|g\|)$.

Discussion. The two reductions put forward in this work provide some evidence towards supporting the conjectured hardness of the search vectorial NTRU problem and the decision NTRU problem. They may give the impression that the hardness of the NTRU problems lies somewhere between the hardness of the ideal-SVP and that of Ring-LWE. This is however neglecting the fact that there are several NTRU problem variants, and it is unclear whether they are computationally equivalent. In particular, the reductions are incompatible, in that the first one reduces to $\mathrm{NTRU}_{\mathrm{vec}}$ and the second one from $\mathrm{NTRU}_{\mathrm{mod}}$. $\mathrm{NTRU}_{\mathrm{mod}}$ reduces to $\mathrm{NTRU}_{\mathrm{vec}}$, but it is a reduction from $\mathrm{NTRU}_{\mathrm{vec}}$ to $\mathrm{NTRU}_{\mathrm{mod}}$ that we would need to obtain a chain of reductions from ideal-SVP to Ring-LWE via the computationally equivalent NTRU problems. Note that if we assume that

ideal-SVP is easy, then these problems are computationally equivalent (see Subsection 3.4), but the reduction from ideal-SVP to NTRU_{vec} becomes vacuous. In fact, it seems that NTRU_{vec} and NTRU_{mod} could even be of different natures: when attempting to solve NTRU_{vec} using an NTRU_{mod} oracle, it is unclear how to make the approximation factor γ/γ' appear, as NTRU_{mod} is only parametrized by the promise gap γ. Better understanding the differences between the NTRU variants seems important to better capture the NTRU hardness. In this direction, note that the known attacks specific to NTRU [ABD16, CJL16, KF17] are mostly relevant for NTRU_{mod}: they can also be used to solve NTRU_{vec}, but the quality of the solution obtained for NTRU_{vec} is the same as the one we would obtain by running the attack to solve NTRU_{mod}, and then running an ideal-SVP solver on the dense rank-1 sub-module to obtain a somehow short vector.

Despite the apparent uncomposability of our two reductions, it would be interesting to have NTRU instance distributions that are compatible with both of them. The second reduction is very permissive with respect to the NTRU_{mod} instance distribution, but the latter still has to satisfy some properties (see Definition 5.1). In particular, the canonical embedding of f and g should be bounded from below and above, and the ideal $\langle f \rangle + \langle g \rangle$ should be coprime with $\langle q \rangle$. We note that in the reduction from ideal-SVP to $\text{wcNTRU}_{\text{vec}}$, the element g is an element of the ideal-SVP instance ideal, which could be chosen Gaussian. Using standard properties of lattice Gaussians, it is not unlikely that one can prove the desired property on its canonical embedding. There seems to be less flexibility in the choice of $f = -g \cdot \{q/z\}$. However, one could replace the deterministic rounding by a Gaussian rounding, to then use a similar approach as the one for g. Concerning the co-primality with $\langle g \rangle$, one could hope to use an inclusion-exclusion argument for Gaussian sums like the one in [SS11].

Concerning the hardness of the NTRU problems relatively to ideal-SVP and Ring-LWE, note that the state of the art suggests that ideal-SVP might be strictly easier than Ring-LWE, as ideal-SVP is known to reduce to Ring-LWE [SSTX09, LPR10, PRS17] but no reduction from Ring-LWE to ideal-SVP is known. In fact, Ring-LWE seems less related to ideal-SVP than to finding two short linearly independent vectors in rank-2 modules over R (SIVP): for an appropriate parametrisation, Ring-LWE reduces to the latter problem [LS15, Se. 5] and, although for some other parametrisation, the latter problem reduces to Ring-LWE (by combining [LS15, Se. 4] and [AD17]). From a lattice perspective, NTRU is a generalization of the unique Shortest Vector Problem to rank-2 modules. At this stage, it is unclear whether its complexity matches the one of ideal-SVP (i.e., SVP for rank-1 modules) or the one of SIVP restricted to rank-2 modules. It could also be strictly in between.

2 Preliminaries

The notations log and ln respectively denote the logarithms in bases 2 and e. For n an integer, we let $[n]$ denote the set $\{1, 2, \ldots, n\}$. Vectors and matrices are denoted with bold lower-case and upper-case letters, respectively. The statistical

distance between two distributions D_1 and D_2 with compatible countable supports is defined as $\mathrm{dist}(D_1, D_2) = \frac{1}{2}\sum_x |D_1(x) - D_2(x)|$. We write $D_1 \approx_\varepsilon D_2$ if $\mathrm{dist}(D_1, D_2) \leqslant \varepsilon$ for some $\varepsilon > 0$. If X is a finite set, then we let $U(X)$ denote the uniform distribution over X. If $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^m$ are linearly independent vectors, then the notation $(\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_n)$ refers to their Gram-Schmidt orthogonalization. The notation $|||\cdot|||$ refers to the matrix norm induced by the Euclidean norm. Finally, we define $\widetilde{O}(d^t)$ as $O(d^t \operatorname{poly}(\log d))$ for any $t \geqslant 0$ including $t = 0$.

2.1 Euclidean lattices

A lattice $L \subset \mathbb{R}^m$ is a set of the form $L = \mathbf{B} \cdot \mathbb{Z}^{m \times n}$ for some full column-rank matrix $\mathbf{B} \in \mathbb{R}^{m \times n}$ (for some $m \geqslant n \geqslant 1$). The columns of $\mathbf{B}$ are said to form a basis of L. For $i \in [n]$, the ith lattice minimum is defined as $\lambda_i(L) = \min(r : \dim L \cap \mathcal{B}(r) \geqslant i)$, where $\mathcal{B}(r)$ denotes the closed ball of $\mathbb{R}^m$ of radius r. The determinant $\det(L)$ is defined as $\sqrt{\det(\mathbf{B}^T\mathbf{B})}$, which is independent of the particular choice of basis $\mathbf{B}$ of L. Minkowski's (second) theorem states that $\prod_{i\in[n]} \lambda_i(L) \leqslant \sqrt{n}^n \cdot \det(L)$.

In this article, we will be interested in the ideal Hermite Shortest vector problem. We first recall below the definition of the Hermite Shortest Vector Problem (HSVP) for arbitrary lattices, and we will instantiate it for ideal lattices in Section 2.4.

Definition 2.1 (γ-HSVP)**.** *Let $\gamma \geqslant 1$. Given as input a lattice $L \subset \mathbb{Q}^n$ (represented by an arbitrary $\mathbb{Z}$-basis), the* γ-HSVP *problem asks to find a vector $\mathbf{w} \in L\backslash\{\mathbf{0}\}$ such that $\|\mathbf{w}\| \leqslant \gamma \cdot \sqrt{n} \cdot \det(L)^{1/n}$.*

By Minkowski's theorem, this problem is well-defined for any $\gamma \geqslant 1$.

2.2 Discrete Gaussian distributions

Let $\mathbf{S} \in \mathrm{GL}_n(\mathbb{R})$ be an invertible matrix. The Gaussian density function with parameter $\mathbf{S}$ is defined over $\mathbb{R}^n$ by

$$\rho_{\mathbf{S}}(\mathbf{x}) = \mathrm{e}^{-\pi\|\mathbf{S}^{-1}\mathbf{x}\|^2}.$$

When the matrix $\mathbf{S}$ is diagonal with diagonal coefficients all equal to some $\sigma > 0$, we also use the notation $\rho_\sigma = \rho_{\mathbf{S}}$. Let $L \subset \mathbb{R}^n$ be a full rank lattice, and $\mathbf{c} \in \mathbb{R}^n$. The discrete Gaussian distribution $D_{L,\mathbf{S},\mathbf{c}}$ over L with center $\mathbf{c}$ and parameter $\mathbf{S}$ is the distribution for which the probability of any $\mathbf{v} \in L$ is $\rho_{\mathbf{S}}(\mathbf{v}-\mathbf{c})/\rho_{\mathbf{S}}(L-\mathbf{c})$, where $\rho_{\mathbf{S}}(T) = \sum_{t\in T} \rho_{\mathbf{S}}(t)$ for any countable $T \subset \mathbb{R}^n$. Again, we will use the notation $D_{L,\sigma,\mathbf{c}}$ when $\mathbf{S} = \mathrm{diag}(\sigma)$ for some $\sigma > 0$. When $\mathbf{c} = \mathbf{0}$, we omit the subscript $\mathbf{c}$.

If $L \subset \mathbb{R}^n$ is a lattice, its smoothing parameter $\eta_\varepsilon(L)$ is defined as the smallest $\sigma > 0$ such that $\rho_{1/\sigma}(L^\star\backslash\{\mathbf{0}\}) \leqslant \varepsilon$, where $L^\star = \{\mathbf{c} \in \mathrm{span}(L) : \forall \mathbf{b} \in L : \langle \mathbf{b}^\star, \mathbf{b}\rangle \in \mathbb{Z}\}$ is the dual of L. For any n-dimensional lattice L and $\varepsilon > 0$, we have the following upper bound on the smoothing parameter (see [MR07, Le. 3.3]):

$$\eta_\varepsilon(L) \leqslant \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} \cdot \lambda_n(L). \tag{2.1}$$

The following Lemma (adapted from [GPV08, Th. 4.1]) shows that one can efficiently sample (bounded) elements from a distribution that is statistically close to a discrete Gaussian distribution. A proof can be found in the full version.

Lemma 2.2. *There exists a* ppt *algorithm that takes as input a basis* $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ *of an* n*-dimensional lattice* L*, a parameter* $\sigma \geqslant \sqrt{n} \cdot \max_i \|\mathbf{b}_i\|$ *and a center* $\mathbf{c} \in \mathrm{Span}(L)$ *and outputs a sample from a distribution* $\widetilde{D}_{\mathbf{B},\sigma,\mathbf{c}}$ *such that*

- $D_{L,\sigma,\mathbf{c}} \approx_{2^{-\Omega(n)}} \widetilde{D}_{\mathbf{B},\sigma,\mathbf{c}}$*;*
- *for all* $\mathbf{v} \leftarrow \widetilde{D}_{B,\sigma,\mathbf{c}}$*, it holds that* $\|\mathbf{v} - \mathbf{c}\| \leqslant \sqrt{n} \cdot \sigma$ *and* $\mathbf{v} \neq \mathbf{0}$*.*

The following lemma bounds the statistical distance between two discrete Gaussian distributions over the same lattice L, depending on the distance between their centers and their parameter matrices. Similar results were already present in previous works, such as in [Reg09, Claim 2.2] for 1-dimensional continuous Gaussian distributions, and in the proof of [dBDPW20, Th. 4.4] for the case of ideal lattices with specific parameters and centers. Since the following precise statement seems new, we provide a proof in the full version for the sake of completeness.

Lemma 2.3. *Let* $L \subset \mathbb{R}^n$ *be a full rank lattice,* $\mathbf{S}_1, \mathbf{S}_2 \in \mathrm{GL}_n(R)$ *be two invertible matrices and* $\mathbf{c}_1, \mathbf{c}_2 \in \mathbb{R}^n$ *be two vectors. If* $\eta_{1/2}(\mathbf{S}_1^{-1}L), \eta_{1/2}(\mathbf{S}_2^{-1}L) \leqslant 1/2$*, then it holds that*

$$\mathrm{dist}\left(D_{L,\mathbf{S}_1,\mathbf{c}_1}, D_{L,\mathbf{S}_2,\mathbf{c}_2}\right) \leqslant 4\sqrt{n} \cdot \left(\sqrt{\left\|\mathbf{S}_2^{-1}\mathbf{S}_1 - \mathbf{I}_n\right\|} + \sqrt{\|\mathbf{S}_2^{-1}(\mathbf{c}_1 - \mathbf{c}_2)\|}\right).$$

The next lemma states that a lattice Gaussian distribution with sufficiently large standard deviation is almost uniform when reduced modulo a sublattice.

Lemma 2.4 ([GPV08, Cor. 2.8]). *Let* $L_1 \subseteq L_2$ *be two lattices of rank* n*. If* $1 \geqslant \eta_\varepsilon(L_1)$ *for some* $\varepsilon < 1/2$*, then* $(D_{L_2,1} \bmod L_1) \approx_{2\varepsilon} U(L_2 \bmod L_1)$*.*

2.3 Number fields

Let K be a number field of degree $d \geqslant 2$ and $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R}$. We let R denote its ring of integers. We identify any element of K with its canonical embedding vector $\sigma : x \mapsto (\sigma_1(x), \ldots, \sigma_d(x))^T \in \mathbb{C}^d$. This leads to an identification of $K_\mathbb{R}$ with $\{\mathbf{y} \in \mathbb{C}^d : \forall i \in [r_\mathbb{R}], y_i \in \mathbb{R} \text{ and } \forall i \in [r_\mathbb{C}], \overline{y_{r_\mathbb{R}+2(i+1)}} = y_{r_\mathbb{R}+2i+1}\}$, where $r_\mathbb{R}$ and $r_\mathbb{C}$ respectively denote the number of real and complex embeddings. Via this identification, the set $K_\mathbb{R}$ is a real vector subspace of dimension d embedded in $\mathbb{C}^d$. In the following, for any element $x \in R, K$ or $K_\mathbb{R}$, we will let $\|x\|$ (resp. $\|x\|_\infty$) denote the Hermitian norm (resp. infinity norm) of the vector $\sigma(x) \in \mathbb{C}^d$. The set $\sigma(R)$ is a lattice, and the absolute field discriminant Δ_K is defined as $\Delta_K = |\det(\sigma(R))^2|$.[1] We have $\Delta_K \geqslant (\pi/4)^d \cdot (d^d/d!)^2$, which implies that we have $d = O(\log \Delta_K)$, for Δ_K growing to infinity.

[1] Note that in order to avoid having absolute values everywhere in the rest of the article, we define Δ_K as the absolute value of the discriminant of K.

The (absolute value of the) algebraic norm of $x \in K_{\mathbb{R}}$ is defined as $\mathcal{N}(x) = \prod_i |\sigma_i(x)|$. Any non-zero element $r \in R$ has algebraic norm $\geqslant 1$, which implies in particular that $\|r\|_\infty \geqslant 1$.

In this work, we assume that we know a monic polynomial $\Phi \in \mathbb{Z}[X]$ defining K and a $\mathbb{Z}$-basis $(r_1, \ldots, r_d)$ of R, where the r_i's are represented by polynomials modulo Φ (of degree $< d$) with rational coefficients. Let $D_\Phi > 0$ be the smallest integer such that $D_\Phi \cdot r_i$ has integral coefficients for all i (i.e., D_Φ is the common denominator to all the r_i polynomials), then the bit-size of D_Φ is polynomial in d and $\|\Phi\|$, where $\|\Phi\|$ is the Euclidean norm of the vector of coefficients of Φ (see for instance [Sut16, Se. 12.4]).

We will assume that this basis has been LLL-reduced [LLL82]. We define $\delta_K = \max_i \|r_i\|_\infty$. Since $\|r\|_\infty \geqslant 1$ for all $r \in R \backslash \{0\}$, we know that $\delta_K \geqslant 1$. Using Minkowski's second theorem and the LLL-reducedness of $(r_1, \ldots, r_d)$, we have that $\delta_K \leqslant \Delta_K^{O(1)}$. In the case of cyclotomic number fields, taking the power basis gives $\delta_K = 1$. For an element $x = \sum_i x_i r_i \in K_{\mathbb{R}}$, define $\lfloor x \rceil = \sum_i \lfloor x_i \rceil r_i$. We will also use the notation $\{x\} = x - \lfloor x \rceil$. It holds that $\|\{x\}\|_\infty \leqslant d/2 \cdot \delta_K$, and hence that $\|\{x\}\| \leqslant d^{3/2} \cdot \delta_K$.

For a rational $x = x_1/x_2$ with $x_1, x_2 \in \mathbb{Z}$ and $\gcd(x_1, x_2) = 1$, we define $\mathrm{size}(x) = 1 + \log|x_1| + \log|x_2|$. For an element $x = \sum_i x_i r_i \in K$, we define $\mathrm{size}(x) = \sum_i \mathrm{size}(x_i)$. The following lemma shows that if we have a sufficiently precise approximation to an embedding of $x \in K$, then one can recover x exactly. This seems folklore, but as we were unable to find a proof, we provide one in the full version. The result and the proof strategy are mentioned in [Coh00, Se. 6.2.4] in the context of quadratic fields and in Roblot's PhD thesis [Rob97] (just after Lemma 2.14). But both references are very brief on the topic. We note that a detailed study was done on a $\mathfrak{p}$-adic counterpart in [Bel04a].

Lemma 2.5. *Let $k \leqslant d$ arbitrary. There exists an algorithm that, given $\widetilde{y}$ such that $|\widetilde{y} - \sigma_k(x)| \leqslant 2^{-p}$ for some $x \in K$ and some $p \geqslant \mathrm{poly}(d, \log \delta_K, \log \|\Phi\|, \mathrm{size}(x))$, recovers x as a rational linear combination of the basis $(r_1, \ldots, r_d)$ of R in* ppt *with respect to p.*

2.4 Ideals and Modules

Ideals. An integral ideal I is a subset of R that is stable by addition and by multiplication with any element of R. A fractional ideal is a subset of K of the form $x \cdot I$ for some $x \in K$ and some integral ideal $I \subseteq R$. We write $\langle z \rangle$ the principal (fractional) ideal generated by $z \in K$. Using the canonical embedding, any non-zero fractional ideal of K is identified to a d-dimensional lattice, called ideal lattice. The algebraic norm of an integral ideal $I \subseteq R$ is defined by $\mathcal{N}(I) = |R/I|$. We extend the notation to a fractional ideal xI with $x \in K$ and I an integral ideal, by setting $\mathcal{N}(xI) = \mathcal{N}(x) \cdot \mathcal{N}(I)$. For a non-zero fractional ideal $I = I_1/2$ with $I_1, I_2 \subseteq R$ and $\gcd(I_1, I_2) = R$, we define the quantity $\mathrm{size}(I) := \log(\mathcal{N}(I_1)) + \log(\mathcal{N}(I_2))$.

Two-element representation of an ideal. Any fractional ideal I can be generated by only two elements, i.e., there exist $x, y \in K$ such that $I = \langle x \rangle + \langle y \rangle$ (see, e.g., [Coh95, Prop. 4.7.7]). In fact, for any $x \in I \backslash \{0\}$, there exists $y \in I$ such that $I = \langle x \rangle + \langle y \rangle$. The lemma below states that computing such a y, given as input (I, x), can be done in probabilistic polynomial time.

Lemma 2.6 (Adapted from [Bel04b, Alg. 6.15] and [FS10, Th. 3]). *There exists a probabilistic algorithm taking a fractional ideal $I \subset K$ and a non-zero $x \in I$ as inputs, computing $y \in I$ such that $I = \langle x \rangle + \langle y \rangle$, and whose run-time is polynomial in* $\mathrm{size}(x)$, $\mathrm{size}(I)$ *and* $\log(\Delta_K)$.

Proof. Wlog, we can restrict the study to non-zero integral ideals. The algorithm is the same as the one given in [FS10, Fig. 1], except that in Step 1, the element x_1 is chosen to be x, rather than the first vector of a reduced basis. The correctness proof is unchanged. The upper bounds on the bit-sizes of the elements appearing during the algorithm execution do change, but one can check that all these bit-sizes stay polynomial in $\mathrm{size}(x)$, as well as the other quantities related to I and K that were already present in [FS10] (which are all polynomial in $\mathrm{size}(I)$ and $\log \Delta_K$). So overall, the run-time remains polynomial in $\mathrm{size}(x)$, $\mathrm{size}(I)$ and $\log \Delta_K$. □

Algorithmic problems over ideal lattices. The ideal-HSVP (or id-HSVP for short) problem is the HSVP problem restricted to lattices that are (fractional) ideal lattices. Using the fact that for an ideal lattice $I \subset K$ we have $\det(I) = \sqrt{|\Delta_K|} \cdot \mathcal{N}(I)$, the problem admits the following equivalent formulation.

Definition 2.7 (γ-id-HSVP). *Let $\gamma \geqslant 1$. Given as input a non-zero fractional ideal $I \subset K$ (represented by an arbitrary $\mathbb{Z}$-basis), the γ-id-*HSVP *problem asks to find an element $\mathbf{w} \in I \backslash \{\mathbf{0}\}$ such that $\|\mathbf{w}\| \leqslant \gamma \cdot \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$.*

Observe that γ-id-HSVP is equivalent to γ'-SVP in ideal lattices, up to polynomial losses $\leqslant \sqrt{d} \cdot \Delta_K^{1/(2d)}$ in the approximation factors γ and γ', thanks to the inequalities

$$\mathcal{N}(I)^{1/d} \leqslant \lambda_1(I) \leqslant \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d},$$

which hold for any non-zero fractional ideal I. The approximation factor loss is polynomial when $\Delta_K^{1/(2d)} \leqslant \mathrm{poly}(d)$.

If $\gamma = \exp(\widetilde{O}(d^\alpha))$ for $\alpha \in [0, 1]$, then Id-HSVP can be solved using lattice reduction algorithms [Sch87], in time $\exp(\widetilde{O}(d^{1-\alpha}))$. In [CDW21], Cramer, Ducas and Wesolowski obtained a heuristic quantum polynomial-time algorithm for $\gamma = \exp(\widetilde{O}(d^{1/2}))$ for cyclotomic fields. In [PHS19], Pellet-Mary, Hanrot and Stehlé gave a quantum heuristic algorithm for $\gamma = \exp(\widetilde{O}((\log \Delta_K)^{\alpha+1})/d)$ running in time $\exp(\widetilde{O}((\log \Delta_K)^{1-2\alpha}))$ for any field K, where $\alpha \in [0, 1/2]$ is arbitrary. They also propose a classical variant of their algorithm, achieving the same approximation factor γ in time $\exp(\widetilde{O}((\log \Delta_K)^{\max(2/3, 1-2\alpha)}))$ for any field K; and in time $\exp(\widetilde{O}(d^{\max(1/2, 1-2\alpha)}))$ for cyclotomic fields. Both the classical and the quantum algorithms require an advice depending only on the field K, whose bit-length is bounded as $\exp(\widetilde{O}((\log \Delta_K)^{1-2\alpha}))$.

Smoothing ideals. The following lemma from [PRS17] provides a sufficient condition for a diagonal matrix $\mathbf{S}$ to be above the smoothing parameter of an ideal lattice.

Lemma 2.8 ([PRS17, Le. 6.9]). *Let $I \subset K$ be a fractional ideal and $\mathbf{S} \in \mathbb{R}^{d\times d}$ be a diagonal matrix with positive diagonal coefficients. Assume that*

$$c := (\prod_i S_{ii})^{1/d} \cdot (\mathcal{N}(I)\Delta_K)^{-1/d} \geqslant 1,$$

then $1 \geqslant \eta_\varepsilon(\mathbf{S}^{-1}I)$, where $\varepsilon = \exp(-c^2 d)$.

Modules. For $\ell \geqslant k \geqslant 1$, a rank-$k$ module $M \subset K_{\mathbb{R}}^{\ell}$ is a set of the form $M = \mathbf{b}_1 I_1 + \ldots + \mathbf{b}_k I_k$ for some non-zero ideals $(I_i)_i$ and some $K_{\mathbb{R}}$-linearly independent vectors $(\mathbf{b}_i)_i$ (i.e., if $\sum_i y_i \mathbf{b}_i = \mathbf{0}$, then all y_i's must be 0). The tuple $((I_i, \mathbf{b}_i))_i$ is called a pseudo-basis of M. If M admits a pseudo-basis for which all the I_i's are equal to R, then M is called free. We define $\det(M)$ as the determinant of M when identified with a kd-dimensional lattice via the canonical embedding σ. For any pseudo-basis $((I_i, \mathbf{b}_i))_i$ of M, we have

$$\det(M)^2 = \Delta_K^k \cdot \mathcal{N}\Big(\det_{K_{\mathbb{R}}}(\overline{\mathbf{B}}^T \mathbf{B}) \prod_i I_i^2\Big), \tag{2.2}$$

where $\det_{K_{\mathbb{R}}}$ is the determinant of a square matrix over $K_{\mathbb{R}}$.

2.5 Oracle Hidden Center Problem

In the search to decision reduction from Section 5, we will make use of the OHCP technique from [PRS17]. The proof of Proposition 2.10 is provided in the full version.

Definition 2.9 (Oracle Hidden Center Problem [PRS17, Def. 4.3]). *Let $\varepsilon, \delta \in (0,1)$ and $\beta \geqslant 1$. An OHCP instance consists in a scale parameter $D > 0$ and a randomized oracle $\mathcal{O} : \mathbb{R}^k \times \mathbb{R}_{\geqslant 0} \to \{0,1\}$ such that, for all $\mathbf{z} \in \mathbb{R}^k$ with $\|\mathbf{z} - \mathbf{z}^*\| \leqslant \beta D$ and $t \in \mathbb{R}_{\geqslant 0}$, it holds that $\Pr(\mathcal{O}(\mathbf{z}, t) = 1) = p(t + \log\|\mathbf{z} - \mathbf{z}^*\|)$, where $\mathbf{z}^* \in \mathbb{R}^k$ is an unknown center satisfying $\delta D \leqslant \|\mathbf{z}^*\| \leqslant D$ and $p(\cdot)$ is an unknown function. The goal of the OHCP is to recover $\tilde{\mathbf{z}} \in \mathbb{R}^k$ such that $\|\tilde{\mathbf{z}} - \mathbf{z}^*\| \leqslant \varepsilon D$.*

Proposition 2.10 (Adapted from [PRS17, Prop. 4.4]). *There exists an algorithm that takes as input a parameter $\kappa \geqslant 20\log(k+1)$, the scaling parameter D and the oracle $\mathcal{O}$ of a $(\exp(-\kappa), \exp(-\kappa), 1 + 1/\kappa)$-OHCP instance in dimension k, and solves it with probability $\geqslant 1 - \exp(-\kappa)$, in time $\mathrm{poly}(\kappa, k)$, provided the oracle $\mathcal{O}$ satisfies the extra following conditions. For some $p_\infty \in [0,1]$ and $t^* \geqslant 0$ we have*

1. *$p(s^*) - p_\infty \geqslant 1/\kappa$;*
2. *$|p(t) - p_\infty| \leqslant 2\exp(-t/\kappa)$ for any $t \geqslant 0$;*
3. *for any $t_1, t_2 \geqslant 0$, it holds that $|p(t_1) - p(t_2)| \leqslant \kappa\sqrt{|t_1 - t_2|}$;*

where $p(t) = \Pr(\mathcal{O}(\mathbf{0}, t) = 1)$.

3 Different variants of the NTRU problem

In this section, we define the three variants of the NTRU problem that we will consider in this work.

3.1 NTRU instances

We first define NTRU instances, which will be the inputs to the NTRU problem variants. We also consider the less standard case of tuple NTRU instances, which has also been considered in cryptographic constructions (see, e.g., the variant of the candidate cryptographic multilinear map from [GGH13] proposed in [LSS14, Se. 6]). All definitions of this section readily extend to the tuple setting, in a manner that is consistent with the second part of Definition 3.1.

Definition 3.1 ((γ, q)-NTRU instance). *Let $q \geqslant 2$ an integer and $\gamma > 0$ a real number. A (γ, q)-*NTRU *instance is an element $h \in R_q$ such that there exists $(f, g) \in R^2 \backslash \{(0, 0)\}$ with $g \cdot h = f \bmod q$ and $\|f\|, \|g\| \leqslant \sqrt{q}/\gamma$. The pair (f, g) is called a trapdoor of the* NTRU *instance h.*

*For $t \geqslant 1$ and γ and q as above, a (γ, q, t)-tuple-*NTRU *instance is a tuple $(h_i)_{i \leqslant t} \in R_q$ such that there exists $((f_i)_{i \leqslant t}, g) \in R^{t+1} \backslash \{(0, \ldots, 0)\}$ with $g \cdot h_i = f_i \bmod q$ and $\max_i \|f_i\|, \|g\| \leqslant \sqrt{q}/\gamma$.*

For a uniform h in R_q, we will see below that the expected norm of a smallest trapdoor (f, g) is of the order of $\sqrt{q}$ (up to factors depending on the field). Hence, the quantity γ of an NTRU instance measures the gap between the size of a short trapdoor of h and the size of a smallest trapdoor of h we would have expected if h was uniform modulo q. Note also that any (γ, q)-NTRU instance is also a (γ', q)-NTRU instance for any $\gamma' \leqslant \gamma$ (the quantity γ is only a lower bound on the promised gap).

We now consider distributions over NTRU instances. To be useful for constructing cryptosystems, these distributions must be efficiently samplable and we also need to be able to sample, together with the NTRU instance h, a trapdoor (f, g) for h. This motivates the following definition.

Definition 3.2 ($(\mathcal{D}, \gamma, q)$-NTRU setup). *Let $q \geqslant 2$, $\gamma > 0$ and $\mathcal{D}$ a distribution over (γ, q)-*NTRU *instances. A $(\mathcal{D}, \gamma, q)$-*NTRU *setup is a* ppt *algorithm (with respect to $\log q$ and $\log \Delta_K$) sampling triples $(h, f, g) \in R_q \times R^2$ such that*

- *the marginal distribution of h is $\mathcal{D}$,*
- *$(f, g) \neq (0, 0)$ and $\|f\|, \|g\| \leqslant \sqrt{q}/\gamma$,*
- *$g \cdot h = f \bmod q$.*

It was shown in [SS11] that for power-of-2 cyclotomic fields, there exists a $(\mathcal{D}, \gamma, q)$-NTRU setup with $\mathcal{D} \approx_{2^{-\Omega(d)}} U(R_q^\times)$ for any prime $q \geqslant 5$ and some $\gamma = 1/\operatorname{poly}(d)$. This was extended to any cyclotomic field in [WW18]. In such cases, the decision NTRU problem introduced below is information-theoretically hard, if we replace $U(R_q)$ by $U(R_q^\times)$. In this work, we rather focus on the case of $\gamma \geqslant 1$.

3.2 Decision NTRU problem

We can now define the decision variant of the NTRU problem.

Definition 3.3 ($(\mathcal{D},\gamma,q)$-dNTRU). *Let $q \geqslant 2$, $\gamma \geqslant 1$ and $\mathcal{D}$ a distribution over (γ,q)-NTRU instances. The $(\mathcal{D},\gamma,q)$ decisional* NTRU *problem ($(\mathcal{D},\gamma,q)$-*dNTRU *for short) asks to distinguish between samples from $\mathcal{D}$ and from $U(R_q)$. The advantage of an algorithm $\mathcal{A}$ against the $(\mathcal{D},\gamma,q)$-*dNTRU *problem is defined as*

$$\mathrm{Adv}(\mathcal{A}) := \Big| \Pr_{h\leftarrow\mathcal{D}}\big(\mathcal{A}(h)=1\big) - \Pr_{u\leftarrow U(R_q)}\big(\mathcal{A}(u)=1\big)\Big|,$$

where the probabilities are also over the internal randomness of $\mathcal{A}$.

A reduction from dNTRU to sRLWE is sketched in [Pei16, Se. 4.4.4].

3.3 Search NTRU problems

We consider two different search variants for the NTRU problem. The first one consists in finding a trapdoor (f,g) for an NTRU instance h such that $\|f\|$ and $\|g\|$ are as small as possible, whereas the second variant only asks to recover any multiple (xf,xg) (with $x\in K$) of a small trapdoor (f,g). We explain below why both variants may be of interest. Further, for both variants, the definition comes with worst-case and average-case flavours.

Definition 3.4 ($(\mathcal{D},\gamma,\gamma',q)$-NTRU$_{\mathrm{vec}}$ and (γ,γ',q)-wcNTRU$_{\mathrm{vec}}$). *Let $q\geqslant 2$, $\gamma\geqslant\gamma'>0$ and $\mathcal{D}$ a distribution over (γ,q)-NTRU instances. The $(\mathcal{D},\gamma,\gamma',q)$ average-case search* NTRU *vector problem ($(\mathcal{D},\gamma,\gamma',q)$-NTRU$_{\mathrm{vec}}$ for short) asks, given as input some h sampled from $\mathcal{D}$, to compute a pair $(f,g)\in R^2\backslash\{(0,0)\}$ such that $g\cdot h = f \bmod q$ and $\|f\|,\|g\|\leqslant\sqrt{q}/\gamma'$. The advantage of an algorithm $\mathcal{A}$ against the $(\mathcal{D},\gamma,\gamma',q)$-NTRU$_{\mathrm{vec}}$ problem is defined as*

$$\mathrm{Adv}(\mathcal{A}) = \Pr_{h\leftarrow\mathcal{D}}\left(\mathcal{A}(h)=(f,g) \text{ with } \begin{vmatrix} g\cdot h = f \bmod q \\ (f,g)\neq(0,0) \\ \|f\|,\|g\|\leqslant\sqrt{q}/\gamma' \end{vmatrix}\right),$$

where the probability is also over the internal randomness of $\mathcal{A}$.

The (γ,γ',q) worst-case search NTRU *vector problem ((γ,γ',q)-wcNTRU$_{\mathrm{vec}}$ for short) asks, given as input a (γ,q)-NTRU instance h, to compute a pair $(f,g)\in R^2\backslash\{(0,0)\}$ such that $g\cdot h = f \bmod q$ and $\|f\|,\|g\|\leqslant\sqrt{q}/\gamma'$.*

Before describing the second search variant of the NTRU problem, we prove the following lemma, which states that all short trapdoors (f,g) of an NTRU instance h are K-multiples of one another.

Lemma 3.5. *Let $q\geqslant 2$, $\gamma>\sqrt{2}$ and h be a (γ,q)-NTRU instance. Then, for all trapdoors $(f,g),(f',g')\in R^2\backslash\{(0,0)\}$ with $\|f\|,\|g\|,\|f'\|,\|g'\|\leqslant\sqrt{q}/\gamma$ and $g\cdot h = f \bmod q$, $g'\cdot h = f' \bmod q$, it holds that $(f,g)=x\cdot(f',g')$ for some $x\in K$.*

Equivalently, there exists a unique field element $h_K\in K$ such that, for all trapdoors $(f,g)\in R^2\backslash\{(0,0)\}$ with $\|f\|,\|g\|\leqslant\sqrt{q}/\gamma$ and $g\cdot h = f \bmod q$, it holds that $f/g = h_K$ (where the division is performed in K and not modulo q).

Proof. Let (f, g) and (f', g') be as in the lemma statement. Then

$$g' \cdot f = g' \cdot (g \cdot h) = g \cdot (g' \cdot h) = g \cdot f' \bmod q.$$

This implies that $g'f - gf' \in qR$. Moreover, we know that $\|g'f - gf'\| \leqslant \|g'\| \cdot \|f\| + \|g\| \cdot \|f'\| \leqslant 2q/\gamma^2 < q$ by assumption on γ. Since any non-zero element of R has euclidean norm at least 1, we conclude that all non-zero elements of qR have norm at least q, and so $g'f - gf' = 0$ in K as desired. The equivalent formulation follows immediately by taking $h_K = f/g$ for any short trapdoor (f, g). Note that g must be invertible in K because otherwise $g = 0$, which implies that $f \in qR$ and so f cannot satisfy $\|f\| \leqslant \sqrt{q}/\gamma$. □

We now describe our second search variant of the NTRU problem. Since we have seen in Lemma 3.5 that recovering a K-multiple of a short trapdoor is equivalent to recovering the (unique) element h_K, we will use this second approach in the description of the problem.

Definition 3.6 ($(\mathcal{D}, \gamma, q)$-NTRU$_{\text{mod}}$ and (γ, q)-wcNTRU$_{\text{mod}}$). *Let $q \geqslant 2$, $\gamma > \sqrt{2}$ and $\mathcal{D}$ a distribution over (γ, q)-NTRU instances. The $(\mathcal{D}, \gamma, q)$ search* NTRU *module problem ($(\mathcal{D}, \gamma, q)$-NTRU$_{\text{mod}}$ for short) asks, given as input an* NTRU *instance h sampled from $\mathcal{D}$, to recover the unique field element $h_K \in K$ associated to h (as defined in Lemma 3.5). The advantage of an algorithm $\mathcal{A}$ against the $(\mathcal{D}, \gamma, q)$-NTRU$_{\text{mod}}$ problem is defined as*

$$\text{Adv}(\mathcal{A}) = \Pr_{h \leftarrow \mathcal{D}} \Big(\mathcal{A}(h) = h_K \Big),$$

where the probability is also over the internal randomness of $\mathcal{A}$.

The (γ, q) worst-case search NTRU *module problem ((γ, q)-wcNTRU$_{\text{mod}}$ for short) asks, given as input a (γ, q)-NTRU instance h, to recover the unique field element $h_K \in K$ associated to h.*

We note that NTRU$_{\text{mod}}$ is definitionally convenient in that the quantity h_K that we are looking for is unique. In NTRU$_{\text{vec}}$, on the contrary, the short trapdoor (f, g) that we are looking for is far from being unique: it can always be multiplied by small elements of R to obtain other trapdoors.

Given a (γ, q)-NTRU instance h, one can construct the following free rank-2 module M_h:

$$M_h := \begin{pmatrix} 1 & 0 \\ h & q \end{pmatrix} \cdot R^2 = \left\{ (g, f)^T \in R^2 \,|\, g \cdot h = f \bmod q \right\}.$$

This module is called the NTRU-module associated to h. As a lattice, it has determinant $\det M_h = \Delta_K \cdot q^d$ and dimension $2d$. If it were a generic lattice with such determinant and dimension, we would heuristically expect that its minimum is $\Theta(\sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \sqrt{q})$. However, since h is a (γ, q)-NTRU instance with $\gamma > \sqrt{2}$, we know that there exists an unexpectedly short vector $(g, f)^T$ in the module M_h. This short vector is not unique, any small multiple $(rg, rf)^T$

with $r \in R$ small is also a short vector of M_h. However, Lemma 3.5 implies that the module spanned by all these short vectors has rank 1 and is unique. Moreover, since this module contains unexpectedly short vectors, it will have an unexpectedly small volume. Summing up, the rank-2 module M_h has multiple unexpectedly short vectors and a unique unexpectedly short rank-1 sub-module. $\mathrm{NTRU}_{\mathrm{vec}}$ asks to find any of the unexpectedly short non-zero vectors of M_h, whereas $\mathrm{NTRU}_{\mathrm{mod}}$ asks to recover the unique short rank-1 sub-module (hence the names "NTRU vector" and "NTRU module").

3.4 Elementary relations between the different NTRU problems

$\mathrm{NTRU}_{\mathrm{mod}}$ and $\mathrm{NTRU}_{\mathrm{vec}}$ respectively reduce to their worst-case counterparts. The proof of the following lemma is similarly direct.

Lemma 3.7. *Let $q \geqslant 2$, $\gamma \geqslant \gamma' > \sqrt{2}$. Then there exists a* ppt *reduction from (γ, q)-wcNTRU$_{\mathrm{mod}}$ to (γ, γ', q)-wcNTRU$_{\mathrm{vec}}$. In the average-case setup, the reduction preserves the distribution of instances.*

If one assumes that ideal-HSVP is easy, then the latter admits a converse result. The proof of the following lemma is available in the full version.

Lemma 3.8. *Let $q \geqslant 2$, $\gamma \geqslant \gamma' > \sqrt{2}$ and $\varepsilon > 0$. Then there exists a* ppt *reduction from $(\gamma, \gamma_{\mathrm{vec}}, q)$-wcNTRU$_{\mathrm{vec}}$ to (γ, q)-wcNTRU$_{\mathrm{mod}}$ and γ_{hsvp}-id-HSVP, where*

$$\gamma_{\mathrm{vec}} = \frac{1}{(1+\varepsilon)\sqrt{2}\Delta_K^{1/(2d)}} \cdot \frac{\gamma}{\gamma_{\mathrm{hsvp}}}.$$

In the average-case setup, the $\mathrm{NTRU}_{\mathrm{mod}}$ *and* $\mathrm{NTRU}_{\mathrm{vec}}$ *instance distributions are identical.*

To reduce dNTRU to $\mathrm{NTRU}_{\mathrm{mod}}$, it suffices to show that for a uniform h, we do not expect an unexpectedly short non-zero vector (or short rank-1 submodule) in M_h. The proof of the following lemma is available in the full version.

Lemma 3.9. *Let $q \geqslant 2$ be a prime that does not divide Δ_K, $\gamma > 16 \cdot \sqrt{d} \cdot \Delta_K^{1/(2d)}$ and $\mathcal{D}$ a distribution over (γ, q)-NTRU instances. Then there exists a* ppt *reduction from $(\mathcal{D}, \gamma, q)$-dNTRU to $(\mathcal{D}, \gamma, q)$-NTRU$_{\mathrm{mod}}$. Further, the reduction makes a single call to the* $\mathrm{NTRU}_{\mathrm{mod}}$ *oracle, and if the advantage of the* $\mathrm{NTRU}_{\mathrm{mod}}$ *solver is ε, then the advantage of the resulting* dNTRU *solver is $\geqslant \varepsilon - 2^{-d}$.*

The objective of the next two sections is to (partly) complete the picture by giving two more sophisticated reductions: a reduction from id-HSVP to $\mathrm{NTRU}_{\mathrm{vec}}$ and a reduction from $\mathrm{NTRU}_{\mathrm{mod}}$ to dNTRU.

4 Reduction from ideal-HSVP to $\mathrm{NTRU}_{\mathrm{vec}}$

This section is devoted to reducing worst-case id-HSVP to average-case $\mathrm{NTRU}_{\mathrm{vec}}$. For this purpose, we first exhibit a Karp reduction from worst-case id-HSVP to $\mathrm{wcNTRU}_{\mathrm{vec}}$. This reduction is then enhanced by using the worst-case to average-case reduction for id-HSVP from [dBDPW20], resulting in a reduction from worst-case id-HSVP to average-case $\mathrm{NTRU}_{\mathrm{vec}}$, where the $\mathrm{NTRU}_{\mathrm{vec}}$ average-case distribution is defined as the distribution obtained by applying the worst-case to worst-case reduction to the distribution on ideals from [dBDPW20]. In the process, we improve the reduction of [dBDPW20] to better suit our needs. We extend it to regimes in which it is not polynomial-time anymore, but allows to reach smaller values for the NTRU modulus q, and we show that it allows to sample from the average-case id-HSVP distribution along with a short non-zero element of the ideal (provided q is sufficiently large, or we have access to a quantum computer). The latter is important to allow to sample from the average-case distribution over NTRU instances, along with a trapdoor.

4.1 Transforming an ideal lattice into an NTRU module

In this section, we show how to efficiently 'embed' an ideal lattice into an NTRU module such that any sufficiently short vector of the NTRU module provides a short vector of the embedded ideal lattice. We first give an efficient reduction from ideal-HSVP to worst-case vectorial NTRU.

Theorem 4.1. *Let $q \geqslant 2$ and $\gamma \geqslant \gamma' > 0$ with $\gamma \cdot \gamma' \cdot \sqrt{d} > 1$. Let $\gamma_{\mathrm{hsvp}} = 4d\delta_K \cdot \gamma/\gamma'$. There is a* ppt *(Karp) reduction from γ_{hsvp}-id-HSVP to (γ, γ', q)-$\mathrm{wcNTRU}_{\mathrm{vec}}$ for ideals $I \subseteq R$ satisfying $\mathcal{N}(I) \in [N/2^d, N]$, with*

$$N = \left\lfloor \left(\frac{\sqrt{q}}{\gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{\frac{1}{2d}}} \right)^d \right\rfloor.$$

Note that the reduction is restricted to integral ideals of bounded norms. The lower bound is not restrictive: given a non-zero integral ideal I such that $\mathcal{N}(I) \leqslant N$, we can scale it to the non-zero integral ideal $I' = \lfloor (N/\mathcal{N}(I))^{1/d} \rfloor \cdot I$, which satisfies $\mathcal{N}(I') \in [N/2^d, N]$ and for which a γ_{hsvp}-id-HSVP solution directly leads to a γ_{hsvp}-id-HSVP solution for I. Concerning the upper bound restriction, the id-HSVP worst-case to average-case reduction from [dBDPW20] (as refined in Subsection 4.2) shows that we can wlog focus on integral ideals I of norms $N \approx 2^{d^{1+\alpha}}$ for some $\alpha \in (0, 1]$. This impacts the choice of the NTRU modulus q.

Let us now focus on the problem parameters. If we put aside factors that depend only on the number field, we can set $N^{1/d} \approx \sqrt{q}/\gamma$, and we then obtain that $\gamma_{\mathrm{hsvp}} \approx \gamma/\gamma'$. This means that the approximation factor (which is γ/γ' in the NTRU case) stays roughly the same, and that the root determinant of the NTRU module is γ times larger than the one of the ideal lattice.

Algorithm 4.1 Transforming an ideal lattice into an NTRU instance

Input: A $\mathbb{Z}$-basis of a non-zero ideal $I \subseteq R$ and a modulus q.
Output: An NTRU instance h.

1: Compute $z \in K$ such that $I = R \cap \langle z \rangle$ (see Lemma 4.2).
2: Let $h = \lfloor q/z \rceil \bmod q \in R_q$.
3: **return** h

The transformation that embeds an ideal lattice into an NTRU module is described in Algorithm 4.1. In Lemma 4.3, we show some properties of Algorithm 4.1, which will be used to prove Theorem 4.1.

Lemma 4.2. *There exists a* ppt *algorithm (in* $\mathrm{size}(I)$ *and* $\log \Delta_K$*) which, given a non-zero integral ideal* I *as input, computes* $z \in K$ *such that* $I = R \cap \langle z \rangle$*.*

Proof. If $I = 0$, then the algorithm returns $z = 0$. If $I = R$, it returns $z = 1$. We now assume that I is neither 0 nor R. Since $I \subseteq R$, it holds that $1 \in I^{-1}$. Let $y \in I^{-1}$ be the output of the algorithm of Lemma 2.6, given $(I^{-1}, 1)$ as input: we have $I^{-1} = \langle 1 \rangle + \langle y \rangle$. Note that $I \neq R$ implies that $y \neq 0$. We then define $z = 1/y$, which fulfills our needs as $J_1 \cap J_2 = (J_1^{-1} + J_2^{-1})^{-1}$ for any non-zero fractional ideals J_1 and J_2. □

When using Lemma 4.2 in Algorithm 4.1, the element z is necessarily non-zero, as I is non-zero. The analysis of Algorithm 4.1 follows the intuition provided by the case of principal ideals (with a known generator) described in the introduction.

Lemma 4.3. *Let* $q \geqslant 2$ *and* $I \subseteq R$ *a non-zero integral ideal. On input* (I, q)*, Algorithm 4.1 outputs* $h \in R_q$ *such that*

- *there exists a pair* $(f, g) \in R^2 \backslash \{(0,0)\}$ *with* $g \cdot h = f \bmod q$ *and* $\|f\|, \|g\| \leqslant d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$*;*
- *for any pair* $(f', g') \in R^2 \backslash \{(0,0)\}$ *with* $g' \cdot h = f' \bmod q$ *and* $\|f'\|_\infty, \|g'\|_\infty < q/(d \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d})$*, we have* $g' \in I \backslash \{0\}$*.*

Moreover, Algorithm 4.1 runs in time polynomial in $\mathrm{size}(I)$*,* $\log q$ *and* $\log \Delta_K$*.*

Proof. The run-time of the algorithm follows from Lemma 4.2. For the proofs of the two main statements, we consider $g \in I \backslash \{0\}$ with minimal infinity norm. By Minkowski's bound, we have that $\|g\|_\infty \leqslant \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$.

We now prove the existence of f such that (f, g) is a short trapdoor for h. By multiplying g with h, we obtain

$$g \cdot h = g \cdot \lfloor q/z \rceil = g \cdot q/z + f,$$

with $f := -g \cdot \{q/z\}$. Since $g \in I$ and $z^{-1} \in I^{-1}$ (because $I \subseteq \langle z \rangle$), we have that $g \cdot q/z \in qR$. This implies that $f \in R$ and $gh = f \bmod q$, as desired. Let us now compute an upper bound on the norm of f (we already know that $\|g\| \leqslant$

$\sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$). We know from the preliminaries that $\|\{q/z\}\|_\infty \leqslant d/2 \cdot \delta_K$, from which we obtain:

$$\|f\| \leqslant \|g\| \cdot (d \cdot \delta_K) \leqslant d^{3/2} \cdot \delta_K \cdot \Delta_K^{\frac{1}{2d}} \cdot \mathcal{N}(I)^{\frac{1}{d}}.$$

Let us now prove the second property of the lemma. Let $(g', f') \in R^2 \backslash \{(0,0)\}$ be such that $g' \cdot h = f' \bmod q$ and

$$\|f'\|_\infty, \|g'\|_\infty < \frac{q}{d \cdot \delta_K \cdot \Delta_K^{\frac{1}{2d}} \cdot \mathcal{N}(I)^{\frac{1}{d}}}.$$

We first show that $g' \neq 0$. Assume by contradiction that $g' = 0$. Then $f' = 0 \bmod q$, i.e., $f' \in qR$. But any non-zero element of qR has infinity norm $\geqslant q$ (using the fact that any non-zero element of R has infinity norm $\geqslant 1$). Since we know that $\|f'\|_\infty < q$, we conclude that $f' = 0$, which contradicts the assumption that $(f', g') \neq (0,0)$.

We now show that $g' \in I$. Recall that z is such that $I = R \cap \langle z \rangle$. Since we already know that $g' \in R$, it suffices to prove that $g' \in \langle z \rangle$, i.e., that $g'/z \in R$. By definition of h, we have:

$$g' \cdot q/z = g' \cdot h + g' \cdot \{q/z\} = f' + g' \cdot \{q/z\} + q \cdot r,$$

for some $r \in R$. Multiplying this equation by g/q (recall that g is a shortest non-zero vector of I for the infinity norm), we obtain

$$g' \cdot g/z = (f' + g' \cdot \{q/z\}) \cdot g/q + g \cdot r.$$

We have seen that $g/z \in R$, so that both terms $g' \cdot g/z$ and $g \cdot r$ are in R. We hence have that the term $(f' + g' \cdot \{q/z\}) \cdot g/q$ must also belong to R. Further, we know that

$$\begin{aligned}\|(f' + g' \cdot \{q/z\}) \cdot g/q\|_\infty &\leqslant (\|f'\|_\infty + \|g'\|_\infty \cdot \|\{q/z\}\|_\infty) \cdot \|g\|_\infty/q \\ &\leqslant \max(\|f'\|_\infty, \|g'\|_\infty) \cdot (1 + d/2 \cdot \delta_K) \cdot \Delta_K^{\frac{1}{2d}} \cdot \mathcal{N}(I)^{\frac{1}{d}}/q.\end{aligned}$$

By assumption, the above is < 1. Since no non-zero element of R has infinity norm < 1, we conclude that $f' + g' \cdot \{q/z\} = 0$. This implies that $g' \cdot q/z = q \cdot r$. Dividing this equality by q, we obtain that $g'/z \in R$, as desired. □

We are now ready to prove Theorem 4.1.

Proof (Theorem 4.1). The reduction consists in calling Algorithm 4.1 on I and q to obtain some $h \in R_q$, then calling the wcNTRU$_{\text{vec}}$ oracle on h and returning the oracle output.

Let $I \subseteq R$ be a γ_{hsvp}-id-HSVP instance satisfying $\mathcal{N}(I) \in [N/2^d, N]$, with N as in the theorem statement. The first statement of Lemma 4.3 ensures that the element h computed by the reduction is a valid (γ, γ', q)-wcNTRU$_{\text{vec}}$ instance. The wcNTRU$_{\text{vec}}$ oracle hence outputs a pair $(f', g') \in R^2 \backslash \{(0,0)\}$ such that $g' \cdot h = f' \bmod q$ and $\|f'\|, \|g'\| \leqslant \sqrt{q}/\gamma'$. By the parameter conditions,

the assumption of the second statement of Lemma 4.3 holds. We hence have that $g' \in I \backslash \{0\}$. Further, by definition of N, the lower bound on $\mathcal{N}(I)$ and definition of γ_{hsvp}, we have

$$\|g'\| \leqslant \frac{\sqrt{q}}{\gamma'} \leqslant \frac{2^{1/d} \cdot N^{\frac{1}{d}} \cdot \gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{\frac{1}{2d}}}{\gamma'} \leqslant \gamma_{\mathrm{hsvp}} \cdot \sqrt{d} \cdot \Delta_K^{\frac{1}{2d}} \cdot \mathcal{N}(I)^{\frac{1}{d}}.$$

Note that we used the inequality $\lfloor x \rceil \geqslant x/2$, which holds for any $x \geqslant 1$. □

4.2 From worst-case id-HSVP to average-case id-HSVP

In [dBDPW20], the authors gave a worst-case to average-case reduction for id-HSVP, for a certain average-case distribution of ideals. We adapt [dBDPW20, Th. 4.5] to Theorem 4.4 below, so that it better fits with our application. We explain in the full version how to adapt the proof.

Theorem 4.4 (Adapted from [dBDPW20, Th. 4.5], ERH). *Let K a number field of degree d and $N \geqslant (12d^{1.5} \log d \cdot \delta_K \cdot \Delta_K^{1/(2d)})^d$ an integer. Let $\gamma > 0$. There exist $\gamma' = \gamma \cdot O(d^{1.5} \Delta_K^{1/d})$, a distribution $\mathcal{D}_N^{\text{id-HSVP}}$ over non-zero integral ideals of K of norm $\leqslant N$ and a reduction:*

- *from worst-case γ'-id-HSVP for all fractional ideals of K,*
- *to average-case γ-id-HSVP for integral ideals distributed from $\mathcal{D}_N^{\text{id-HSVP}}$.*

The reduction decreases the success probability by at most $2^{-\Omega(d)}$, makes a single call to the average-case γ-id-HSVP oracle, and runs in time $T_\beta^{\text{id-HSVP}} + \mathrm{poly}(\log N, \mathrm{size}(I), \log \Delta_K)$ where

- *I is the input (worst-case) ideal;*
- *$T_\beta^{\text{id-HSVP}}$ is the time needed to solve id-HSVP with approximation factor $2^{d/\beta}$ and*

$$\beta = \left\lceil \frac{d}{\log\left(N^{1/d}/(6d^{1.5} \log d \cdot \delta_K \cdot \Delta_K^{1/(2d)})\right)} \right\rceil.$$

Moreover, there exist $N_0 = \mathrm{poly}(\Delta_K^{1/d}, \delta_K, d)^d$ and a ppt *algorithm $\mathcal{A}$ (with respect to $\log N$ and $\log \Delta_K$) such that, for all $N \geqslant N_0$, algorithm $\mathcal{A}$ samples pairs (J, w) such that:*

- *the ideal J is a non-zero integral ideal of norm $\leqslant N$;*
- *the distribution $\widetilde{\mathcal{D}}_N^{\text{id-HSVP}}$ of J satisfies $\widetilde{\mathcal{D}}_N^{\text{id-HSVP}} \approx_{2^{-\Omega(d)}} \mathcal{D}_N^{\text{id-HSVP}}$;*
- *the element $w \in J \backslash \{0\}$ satisfies $\|w\| \leqslant \mathrm{poly}(d, \delta_K, \Delta_K^{1/d}, 2^{\sqrt{\log \Delta_K + d \log d}}) \cdot \mathcal{N}(J)^{1/d}$.*

If we have access to a factoring oracle or if $N \geqslant N_0' = N_0 \cdot 2^{O(d\sqrt{\log \Delta_K + d \log d})}$, then we can reduce the size of w down to $\|w\| \leqslant \mathrm{poly}(d, \delta_K, \Delta_K^{1/d}) \cdot \mathcal{N}(J)^{1/d}$.

Note that even though the reduction relies on a worst-case id-HSVP solver, the latter is with an approximation factor $2^{d/\beta}$ which is typically much larger than γ'. This implies that $T_\beta^{\text{id-HSVP}}$ is expected to be much smaller than the time needed to solve γ'-id-HSVP. Assume that $\Delta_K^{1/(2d)}$ and δ_K are both $\text{poly}(d)$ and that we use the lattice reduction algorithm from [Sch87] with block size β to solve $2^{d/\beta}$-id-HSVP. It runs in time $T_\beta^{\text{id-HSVP}} = 2^{O(\beta)}$ (up to a $\text{poly}(\log N, \log \Delta_K)$ factor). Then, it can be seen that the reduction is polynomial-time when $N = 2^{\Omega(d^2)}$; it becomes more expensive when N is below this bound; and it ends up being $2^{O(d)}$ when $N \approx \text{poly}(d)^d$. The run-time of the reduction can be improved using id-HSVP algorithms such as those mentioned in Subsection 2.3. In all cases, we note that one can sample ideals J from $\mathcal{D}_N^{\text{id-HSVP}}$, together with a short vector of J in quantum polynomial time even for small N, and in classical polynomial time for larger N's (of the order of $2^{O(d^{1.5}\sqrt{\log d})}$ if $\Delta_K^{1/(2d)}$ and δ_K are both $\text{poly}(d)$).

All the ingredients for the proof of Theorem 4.4 are present in [dBDPW20], however the latter only considered the case of $N \geqslant (2^d \cdot 6d^{1.5} \log d \cdot \Delta_K^{1/(2d)} \cdot \delta_K)^d$, since this is the range of parameters for which the reduction runs in polynomial time. The generalization to smaller N and larger run-time is relatively immediate and is provided in the full version. A further difference with [dBDPW20] is that the distribution $\mathcal{D}_N^{\text{id-HSVP}}$ in [dBDPW20] is over the inverses of integral ideals (see [dBDPW20, Le. 4.1]) whereas here it is more convenient to have a distribution over integral ideals. Finally, we also explain in the full version how to sample ideals from $\mathcal{D}_N^{\text{id-HSVP}}$ with a somehow short vector.

4.3 An average-case distribution of NTRU instances

In this subsection, we define a distribution $\mathcal{D}_{q,\gamma}^{\text{NTRU}}$ over (γ, q)-NTRU instances. This distribution is defined as the one being produced by Algorithm 4.2. In fact, Algorithm 4.2 actually provides a $(\widetilde{\gamma}, q)$-NTRU setup for some $\widetilde{\gamma} \geqslant \gamma$, i.e., the instance h can be sampled along with a trapdoor (f, g) that may be a little larger than a shortest one.

Algorithm 4.2 Sampling h from $\mathcal{D}_{q,\gamma}^{\text{NTRU}}$ together with a trapdoor

Input: An integer $q \geqslant 2$ and a real $\gamma \geqslant 1$
Output: A triple $(h, f, g) \in R_q \times R^2$.

1: Let $N = \left\lfloor \left(\frac{\sqrt{q}}{\gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}} \right)^d \right\rfloor$.
2: Sample I from $\widetilde{\mathcal{D}}_N^{\text{id-HSVP}}$ with $v \in I \backslash \{0\}$ such that $\|v\| \leqslant \text{poly}(d, \delta_K, \Delta_K^{1/d}) \cdot \mathcal{N}(I)^{1/d}$ (see Theorem 4.4).
3: Let $I' = \lfloor (N/\mathcal{N}(I))^{1/d} \rfloor \cdot I$ and $v' = \lfloor (N/\mathcal{N}(I))^{1/d} \rfloor \cdot v$.
4: Run Algorithm 4.1 on I'; let $h \in R_q$ be the output and z as in Algorithm 4.1.
5: Compute $g = v'$ and $f = -g \cdot \{q/z\}$.
6: **return** (h, f, g).

Lemma 4.5. *There exist* $\Gamma = \mathrm{poly}(d, \delta_K, \Delta_K^{1/d})$ *and* $\Gamma' = \Gamma \cdot 2^{O(\sqrt{\log \Delta_K + d \log d})}$ *such that if* $\sqrt{q}/\gamma \geqslant \Gamma$ *(resp.* $\sqrt{q}/\gamma \geqslant \Gamma'$*), then Algorithm 4.2 runs in quantum (resp. classical) polynomial time (with respect to* $\log q$ *and* $\log \Delta_K$*).*

Proof. Let $\Gamma = 2d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot N_0^{1/d}$ (resp. $\Gamma' = 2d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot (N_0')^{1/d}$), where N_0 (resp. N_0') is as in the second part of Theorem 4.4. Note that we have $\Gamma = \mathrm{poly}(d, \delta_K, \Delta_K^{1/d})$ (resp. $\Gamma' = \Gamma \cdot 2^{O(\sqrt{\log \Delta_K + d \log d})}$) as desired. Moreover, by definition of N and using the fact that $\sqrt{q}/\gamma \geqslant \Gamma$ (resp. $\sqrt{q}/\gamma \geqslant \Gamma'$), we have $N \geqslant N_0$ (resp. $N \geqslant N_0'$). Hence, by Theorem 4.4, one can sample (I, v) in Step 2 in quantum (resp. classical) time $\mathrm{poly}(\log N, \log \Delta_K) = \mathrm{poly}(\log \Delta_K, \log q)$.

By Theorem 4.4, we also know that the ideal I is non-zero and satisfies $\mathcal{N}(I) \leqslant N$, hence $\lfloor (N/\mathcal{N}(I))^{1/d} \rfloor \neq 0$. Therefore, the ideal I' computed at Step 3 is also non-zero, and v' is a non-zero element of I'. Thanks to Lemma 4.3, we know that Algorithm 4.1 can be run on I' in time $\mathrm{poly}(\mathrm{size}(I'), \log q, \log \Delta_K)$. Since I' is integral and $\mathcal{N}(I') \leqslant N \leqslant q^d$, we conclude that $\mathrm{size}(I') \leqslant \mathrm{poly}(\log q, \log \Delta_K)$. Finally, computing f using the formula $-g \cdot \{q/z\}$ can also be done in time $\mathrm{poly}(\log q, \log \Delta_K)$, since the rounding operation in R is efficient. □

Now that it is established that Algorithm 4.2 terminates, we can formally define $\mathcal{D}_{\gamma,q}^{\mathrm{NTRU}}$ as the distribution produced by the algorithm.

Definition 4.6 (Distribution $\mathcal{D}_{q,\gamma}^{\mathrm{NTRU}}$). *Let* q, γ *as in Algorithm 4.2. The distribution* $\mathcal{D}_{\gamma,q}^{\mathrm{NTRU}}$ *over* R_q *is defined as the distribution of the element* h *produced by Algorithm 4.2 on input* (q, γ)*.*

Lemma 4.7. *The support of the distribution* $\mathcal{D}_{q,\gamma}^{\mathrm{NTRU}}$ *is contained in the set of* (γ, q)-NTRU *instances.*

Proof. Let h be computed by Algorithm 4.2 on input (q, γ). By the first property of Lemma 4.3, there exists a trapdoor $(f^\star, g^\star) \neq (0, 0)$ for h, with $\|f^\star\|, \|g^\star\| \leqslant d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I')^{1/d}$. We have $\mathcal{N}(I') = \lfloor (N/\mathcal{N}(I))^{1/d} \rfloor^d \cdot \mathcal{N}(I) \leqslant N$. Using the definition of N, we conclude that $\|f^\star\|, \|g^\star\| \leqslant \sqrt{q}/\gamma$. □

Algorithm 4.2 gives a way to sample from $\mathcal{D}_{q,\gamma}^{\mathrm{NTRU}}$ together with a trapdoor.

Lemma 4.8. *Let* q, γ *as in Algorithm 4.2 and* Γ *(resp.* Γ'*) as in Lemma 4.5. If* $\sqrt{q}/\gamma \geqslant \Gamma$ *(resp.* $\sqrt{q}/\gamma \geqslant \Gamma'$*), then there exist* $\widetilde{\gamma} = \gamma/\mathrm{poly}(d, \delta_K, \Delta_K^{1/d})$ *such that Algorithm 4.2 is a* $(\mathcal{D}_{q,\gamma}^{\mathrm{NTRU}}, \widetilde{\gamma}, q)$-NTRU *quantum (resp. classical) setup.*

Proof. We have already seen in Lemma 4.5 that Algorithm 4.2 is quantum (resp. classical) ppt. We have seen in Lemma 4.7 that $\mathcal{D}$ is a distribution over (γ, q)-NTRU instances. It is hence a distribution over $(\widetilde{\gamma}, q)$-NTRU instances, as $\widetilde{\gamma} \leqslant \gamma$. We now show that the sampled pair $(f, g) \neq (0, 0)$ satisfies $g \cdot h = f \bmod q$ and $\|f\|, \|g\| \leqslant \sqrt{q}/\gamma \cdot \mathrm{poly}(d, \delta_K, \Delta_K^{1/d})$.

We have already seen that $g = v'$ is non-zero. Moreover, by definitions of $f = -g \cdot \{q/z\}$ and $h = \lfloor q/z \rceil$, is holds that $f = g \cdot h \bmod q$ (see the proof of

Lemma 4.3). Further, we have (successively using Theorem 4.4, the definition of I' and the definition of N):

$$\begin{aligned}\|g\| = \|v'\| &\leqslant \mathrm{poly}(d, \delta_K, \Delta_K^{1/d}) \cdot \mathcal{N}(I')^{1/d} \leqslant \mathrm{poly}(d, \delta_K, \Delta_K^{1/d}) \cdot N^{1/d} \\ &\leqslant \mathrm{poly}(d, \delta_K, \Delta_K^{1/d}) \cdot \frac{\sqrt{q}}{\gamma}.\end{aligned}$$

Moreover, by definition of f, we know that $\|f\| \leqslant \|g\| \cdot (d \cdot \delta_K)$. Hence, there exists some $\widetilde{\gamma} = \gamma / \mathrm{poly}(d, \delta_K, \Delta_K^{1/d})$ such that $\|f\|, \|g\| \leqslant \sqrt{q}/\widetilde{\gamma}$, as desired. □

4.4 From average-case id-HSVP to average-case NTRU

By combining the results from Subsections 4.1 and 4.3, we obtain that, for well-chosen parameters, average-case id-HSVP for distribution $\mathcal{D}_N^{\text{id-HSVP}}$ reduces to average-case $\mathrm{NTRU}_{\mathrm{vec}}$ for distribution $\mathcal{D}_{q,\gamma}^{\mathrm{NTRU}}$. The proof of Theorem 4.9 is available in the full version. This theorem can in turn be combined with Theorem 4.4 to obtain a reduction from worst-case id-HSVP to average-case $\mathrm{NTRU}_{\mathrm{vec}}$.

Theorem 4.9. *Let $q \geqslant 2$, $\gamma \geqslant 1$ and $\gamma' > 0$ such that $\gamma \cdot \gamma' \cdot \sqrt{d} > 1$ and $\sqrt{q}/\gamma \geqslant 13 \cdot d^3 \log d \cdot \delta_K^2 \cdot \Delta_K^{1/d}$. Define:*

$$N = \left\lfloor \left(\frac{\sqrt{q}}{\gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}} \right)^d \right\rfloor \quad \text{and} \quad \gamma_{\mathrm{hsvp}} = \frac{\gamma}{\gamma'} \cdot 4d\delta_K.$$

There is a ppt *reduction (with respect to $\log \Delta_K$ and $\log q$) from average-case γ_{hsvp}-id-HSVP for ideals sampled from $\widetilde{\mathcal{D}}_N^{\text{id-HSVP}}$ to $(\mathcal{D}_{q,\gamma}^{\mathrm{NTRU}}, \gamma, \gamma', q)$-$\mathrm{NTRU}_{\mathrm{vec}}$. The reduction makes a single call to the $\mathrm{NTRU}_{\mathrm{vec}}$ oracle and preserves the success probability.*

5 A search to decision reduction for NTRU

In this section, we provide a reduction from average-case search-$\mathrm{NTRU}_{\mathrm{mod}}$ with distribution $\mathcal{D}^{\mathrm{s}}$ to average-case dec-NTRU with distribution $\mathcal{D}^{\mathrm{d}}$. The distribution $\mathcal{D}^{\mathrm{s}}$ can be chosen from a large class of distributions (it only has to be bounded and to have an invertible denominator, as per Definition 5.1 below) and the distribution $\mathcal{D}^{\mathrm{d}}$ is a function of $\mathcal{D}^{\mathrm{s}}$. Moreover, we show that if the distribution $\mathcal{D}^{\mathrm{s}}$ enjoys an NTRU setup, then so does $\mathcal{D}^{\mathrm{d}}$.

5.1 Choice of the distributions

We start by describing a property of distributions that we will need for our search to decision reduction. We also describe the distribution $\mathcal{D}^{\mathrm{d}}$ as a function of $\mathcal{D}^{\mathrm{s}}$, and explain how one can sample h with a trapdoor (f, g) from $\mathcal{D}^{\mathrm{d}}$, provided there is an efficient algorithm doing it for $\mathcal{D}^{\mathrm{s}}$.

Definition 5.1 (Well-behaved elements and distributions). *Let $q \geqslant 2$ be an integer and $B > 1$ be a real number. An element $h \in R_q$ is said to be B-well-behaved if there exists $f, g \in R$ such that $gh = f \bmod q$; $\langle f \rangle + \langle g \rangle + \langle q \rangle = R$; and for all $1 \leqslant i \leqslant d$ we have $1/B \leqslant |\sigma_i(f)|, |\sigma_i(g)| \leqslant B$.*

A distribution $\mathcal{D}$ over R_q is said to be (B, ε)-well-behaved for some $\varepsilon \geqslant 0$ if the probability that $h \leftarrow \mathcal{D}$ is B-well-behaved is $\geqslant 1 - \varepsilon$.

Observe that any $(B, 0)$-well-behaved distribution over R_q is a distribution over (γ, q)-NTRU instances, where $\gamma = \sqrt{q}/(B\sqrt{d})$. Observe also that the condition $\langle f \rangle + \langle g \rangle + \langle q \rangle = R$ is equivalent to asking that g is invertible modulo q. Indeed, since $gh = f \bmod q$, then any prime factor dividing both $\langle g \rangle$ and $\langle q \rangle$ would also be a prime factor of $\langle f \rangle$, contradicting the coprimality condition. Let us now define a randomized mapping ϕ_B over R_q.

Definition 5.2 (Function ϕ_B). *Let $q \geqslant 2$ and $B > 1$. We define the randomized mapping ϕ_B over R_q as follows*

$$\begin{aligned} \phi_B : \quad & R_q \to R_q \\ & h \mapsto xh + y \bmod q \quad \text{where } x, y \leftarrow D_{R, 2Bd\delta_K}. \end{aligned}$$

We extend ϕ_B to distributions over R_q: for a distribution $\mathcal{D}$, we let $\phi_B(\mathcal{D})$ be the distribution over R_q obtained by sampling $h \leftarrow \mathcal{D}$ and then outputting $\phi_B(h)$.

Finally, we show that if $\mathcal{D}$ enjoys an NTRU setup, then so does $\phi_B(\mathcal{D})$.

Lemma 5.3. *Let $B \geqslant 1$, $q \geqslant 2$, $\gamma > 0$ and $\mathcal{D}$ a distribution over (γ, q)-NTRU instances. If there exists a $(\mathcal{D}, \gamma, q)$-NTRU setup, then there exists a $(\mathcal{D}', \gamma', q)$-NTRU setup where $\mathcal{D}'$ is a distribution over R_q such that $\mathcal{D}' \approx_{2^{-\Omega(d)}} \phi_B(\mathcal{D})$ and $\gamma' = \gamma/(4Bd^{1.5}\delta_K)$.*

Proof. Let $\mathcal{A}$ be a ppt algorithm (with respect to $\log q$ and $\log \Delta_K$) sampling triples $(h, f, g) \in R_q \times R^2$ such that the marginal distribution of h is $\mathcal{D}$, $(f, g) \neq (0, 0)$, $\|f\|, \|g\| \leqslant \sqrt{q}/\gamma$ and $g \cdot h = f \bmod q$.

We consider the following algorithm $\mathcal{B}$:

- run $\mathcal{A}$; let (h, f, g) be the output;
- use the algorithm from Lemma 2.2 with parameters $\sigma = 2Bd\delta_K$ and $\mathbf{c} = \mathbf{0}$ to sample x and y (using the basis $(r_1, \ldots, r_d)$ of R);
- return $(h', f', g') = (xh + y, xf + yg, g)$.

Note that $\mathcal{B}$ is ppt and that (f', g') is non-zero and satisfies $g' \cdot h' = f' \bmod q$. By Lemma 2.2, we also have

$$\|f'\| \leqslant 2Bd^{1.5}\delta_K \cdot (\|f\| + \|g\|) \leqslant 4Bd^{1.5}\delta_K \cdot \frac{\sqrt{q}}{\gamma}.$$

Finally, as the residual distribution of h is $\mathcal{D}$, Lemma 2.2 also implies that the residual distribution of h' is within statistical distance $2^{-\Omega(d)}$ from $\phi_B(\mathcal{D})$. □

We can now state the main result of this section: a reduction from $\mathrm{NTRU}_{\mathrm{mod}}$ to dNTRU, for well-chosen distributions. This theorem follows from Lemmas 2.5, 5.6 and 5.7, which are stated and proved in the following subsections. The proof of Theorem 5.4 is provided in the full version.

Theorem 5.4. *Let $q \geqslant 2$, $B \in (1, q]$, $\varepsilon \geqslant 0$ and $\mathcal{D}^{\mathrm{s}}$ be a (B, ε)-well-behaved distribution over R_q. Assume that $\log q, \log \Delta_K, \log \|\Phi\| \leqslant 2^{o(d)}$ (recall that Φ is a defining polynomial of K). Define $\gamma' := \frac{\sqrt{q}}{4B^2 d^2 \delta_K}$ and assume that $\gamma \geqslant 1$. Let $\mathcal{A}$ be an algorithm solving $(\phi_B(\mathcal{D}^{\mathrm{s}}), \gamma', q)$-dNTRU with advantage $\mathrm{Adv}(\mathcal{A}) \geqslant 2^{-o(d)}$. Then, there exists an algorithm $\mathcal{B}$ solving $(\mathcal{D}^{\mathrm{s}}, \gamma, q)$-$\mathrm{NTRU}_{\mathrm{mod}}$ with $\gamma = \sqrt{q}/(B\sqrt{d})$ and advantage $\mathrm{Adv}(\mathcal{B}) \geqslant (\mathrm{Adv}(\mathcal{A}) - 2\varepsilon)/4$. Algorithm $\mathcal{B}$ is* ppt *(with respect to $\log q$, $\log \Delta_K$, $\log \|\Phi\|$ and $\mathrm{Adv}(\mathcal{A})^{-1}$) and makes (possibly that many) oracle queries to $\mathcal{A}$.*

Observe that up to polynomial factors depending on the number field K, we have $\gamma \approx \sqrt{q}/B$ and $\gamma' \approx \sqrt{q}/B^2$. Said differently, the Euclidean norm of the short trapdoor is squared when we go from $\mathcal{D}^{\mathrm{s}}$ (which has short trapdoors of size roughly B) to $\phi_B(\mathcal{D}^{\mathrm{s}})$ (which has short trapdoors of size roughly B^2). Hence, one should consider $B \leqslant q^{1/4}$ for the dNTRU instances to have short trapdoors of norm $\geqslant \sqrt{q}$.

5.2 Creating new NTRU instances

In this section, we give a lemma which will allow us to rerandomize an NTRU instance h so that the distribution of the new NTRU instance depends on $c_1\sigma_1(f) + c_2\sigma_1(g)$ for some parameters c_1 and c_2 that we can customize. This lemma will be used to prove Lemma 5.7, in the next subsection.

Lemma 5.5. *Let $(f, g) \in R^2 \backslash \{(0,0)\}$ and $I = \langle f \rangle + \langle g \rangle$. Let $c_1, c_2 \in \sigma_1(K_{\mathbb{R}})$ (which is either $\mathbb{R}$ or $\mathbb{C}$), $s_0 > 0$ and $s \geqslant \sqrt{d}\delta_K \cdot (\|f\| + \|g\|)$.*

Given $t \in \sigma_1(K_{\mathbb{R}})$, we define $\psi(t) \in K_{\mathbb{R}}$ as $(t, 0, \ldots, 0)^T \in K_{\mathbb{R}}$ if σ_1 is a real embedding and as $(t/\sqrt{2}, \overline{t}/\sqrt{2}, 0, \ldots, 0)^T \in K_{\mathbb{R}}$ if σ_1 is a complex embedding with $\sigma_2 = \overline{\sigma_1}$.[2]

Let $\mathcal{D}$ be the output distribution of the following algorithm:

- *sample $c_0 \leftarrow D_{\sigma_1(K_{\mathbb{R}}), s_0, 0}$;*
- *sample $x \leftarrow D_{R, s, \psi(c_0 \cdot c_1)}$ and $y \leftarrow D_{R, s, \psi(c_0 \cdot c_2)}$;*
- *return $x \cdot f + y \cdot g \in I$.*

Then it holds that $\mathcal{D} \approx_{2^{-\Omega(d)}} D_{I, \mathbf{S}, 0}$, where $\mathbf{S}$ is a diagonal matrix with

$$
\begin{aligned}
S_{11} &= \sqrt{s_0^2 \cdot |c_1\sigma_1(f) + c_2\sigma_1(g)|^2 + s^2 \cdot (|\sigma_1(f)|^2 + |\sigma_1(g)|^2)} \\
S_{22} &= \begin{cases} S_{11} & \text{if } \sigma_1 \text{ is a complex embedding} \\ s \cdot \sqrt{|\sigma_2(f)|^2 + |\sigma_2(g)|^2} & \text{if } \sigma_1 \text{ is a real embedding} \end{cases} \\
S_{ii} &= s \cdot \sqrt{|\sigma_i(f)|^2 + |\sigma_i(g)|^2} \qquad \text{for } \ i \geqslant 3.
\end{aligned}
$$

The above can be obtained by combining the convolution result of [Pei10, Th. 3.1] and the discrete Gaussian leftover hash lemma from [LSS14, Th. 5.1]. Unfortunately, the statements of [Pei10, Th. 3.1] and [LSS14, Th. 5.1] do not

[2] The scaling by a factor $1/\sqrt{2}$ in the complex case ensures that the norm of $\psi(t)$ is still equal to $|t|$, which allows simpler expressions.

exactly match what we need (in particular, non-zero centers are not considered in [LSS14, Th. 5.1] and the convolution result of [Pei10, Th. 3.1] does not consider c_0 being sampled from a smaller space and extended with zeros). In the full version, we prove some slight variants of these results, in order to prove Lemma 5.5.

Observe that by taking $s = 2Bd\delta_K$ and $c_1 = c_2 = 0$, then the distribution of $x \cdot f + y \cdot g$ is exactly the distribution of the numerator of $\phi_B(h)$, over the randomness of ϕ_B (i.e., when h, f and g are fixed). Note that for Lemma 5.5 to be applicable, we need $s = 2Bd\delta_K \geqslant \sqrt{d}\delta_K \cdot (\|f\| + \|g\|)$, which holds true if $\|f\|_\infty, \|g\|_\infty \leqslant B$. This is the source of the 'standard deviation squaring' in Theorem 5.4. Finally, note that by using the lemma multiple times with the same h, we obtain tuple NTRU instances (as defined in Definition 3.1), implying that the dNTRU and NTRU_{vec} problem variants reduce to their tuple counterparts (under proper parametrization).

5.3 Using the OHCP framework

We now prove two lemmas for the core of the proof of Theorem 5.4. Lemma 5.6 essentially states that when sampling h from $\mathcal{D}^{\text{s}}$, then one should get a "good" h with non-negligible probability. Lemma 5.7 then shows that when h is "good", it is possible to recover a very accurate approximation of $\sigma_1(h_K)$ using the dNTRU oracle. Combining these two lemmas with Lemma 2.5 (which states that one can recover an element $x \in K$ exactly from a sufficiently good approximation of $\sigma_1(x)$) then yields Theorem 5.4 (whose proof is provided in the full version).

Lemma 5.6. *Let $q \geqslant 2$, $B \in (1, q]$, $\varepsilon \geqslant 0$ and $\mathcal{D}^{\text{s}}$ be a (B, ε)-well-behaved distribution over R_q. Let $\mathcal{A}$ be an algorithm solving $(\phi_B(\mathcal{D}^{\text{s}}), \gamma, q)$-dNTRU for some $\gamma \geqslant 1$. Then, there exists a set $H \subset R_q$ such that every h in H is B-well-behaved; $\Pr_{h \leftarrow \mathcal{D}^{\text{s}}}(h \in H) \geqslant \text{Adv}(\mathcal{A})/2 - \varepsilon$; and for all $h \in H$*

$$\left| \Pr\left(\mathcal{A}(\phi_B(h)) = 1\right) - \Pr\left(\mathcal{A}(u) = 1\right) \right| \geqslant \text{Adv}(\mathcal{A})/2,$$

where the probabilities are taken over the internal randomness of $\mathcal{A}$, the randomness of ϕ_B and the random choice of $u \leftarrow U(R_q)$ (but not over the choice of h).

Proof. There exists $H_0 \subset R_q$ of weight $\geqslant \text{Adv}(\mathcal{A})/2$ under $\mathcal{D}^{\text{s}}$ such that for all $h \in H_0$, the advantage of $\mathcal{A}$ on $\phi_B(h)$ is at least $\text{Adv}(\mathcal{A})/2$. We define H as the subset of the h's in H_0 that are B-well-behaved. The result follows from the definition of (B, ε)-well-behavedness and the union bound. □

Lemma 5.7. *Let $q \geqslant 2$, $B \in (1, q]$, $\varepsilon \geqslant 0$ and $\mathcal{D}^{\text{s}}$ be a (B, ε)-well-behaved distribution over R_q. Let $\mathcal{D}^{\text{d}} = \phi_B(\mathcal{D}^{\text{s}})$. Let $\mathcal{A}$ and H as in Lemma 5.6. Assume that $\text{Adv}(\mathcal{A})^{-1}, \log q, \log \Delta_K \leqslant 2^{o(d)}$. Then, there exists a probabilistic algorithm $\mathcal{B}$ that, given an integer $\ell \leqslant 2^{o(d)}$ and any $h \in H$, recovers $\sigma_1(h_K)$ with ℓ bits of absolute precision[3] with probability $\geqslant 1 - 2^{-\Omega(d)}$ (where h_K is defined as in*

[3] The term "absolute precision" refers here to $|\tilde{x} - x| \leqslant 2^{-\ell}$, as opposed to the "relative precision" which would be $\frac{|\tilde{x}-x|}{|x|} \leqslant 2^{-\ell}$.

Lemma 3.5). Moreover, algorithm $\mathcal{B}$ runs in time polynomial in $\ell, \mathrm{Adv}(\mathcal{A})^{-1}, \log q$ and $\log \Delta_K$ and makes (possibly that many) oracle queries to $\mathcal{A}$.

Proof. In order to prove the lemma, we will express our problem as an instance of the Oracle Hidden Center Problem (see Definition 2.9) and then use Proposition 2.10 to conclude.

Let $h \in H$ be fixed once and for all, and given to $\mathcal{B}$. Let us also fix some (unknown) $(f, g) \in R^2$ such that $g \cdot h = f \bmod q$; g is invertible modulo q; and $|\sigma_i(f)|, |\sigma_i(g)| \in [1/B, B]$ for all embeddings σ_i (we know that such f and g exist since h is B-well-behaved by definition of H). We write $I = \langle f \rangle + \langle g \rangle$, which is also fixed once and for all (and is coprime to $\langle q \rangle$).

Let $k = 1$ if σ_1 is a real embedding and $k = 2$ if σ_1 is a complex embedding. In the following, we will identify $\mathbb{R}^k$ with $\sigma_1(K_{\mathbb{R}})$. Note that in both cases, the Euclidean norm of a vector in $\mathbb{R}^k$ corresponds to the absolute value of the element seen in $\mathbb{R}$ or $\mathbb{C}$.

In order to fit the OHCP framework, we need to describe a randomized oracle $\mathcal{O}$ that takes as input a pair $(z, t) \in \mathbb{R}^k \times \mathbb{R}^{\geqslant 0}$ and outputs 0 or 1 such that $\Pr_{\mathcal{O}}(\mathcal{O}(z, t) = 1) = P(t + \ln|z - \sigma_1(h_K)|)$, for some (unknown) function P (that may depend on h). In other words, we want that the acceptance probability of the oracle $\mathcal{O}$ only depends on $t + \ln|z - \sigma_1(h_K)|$ (when t and z vary).

We start by considering an oracle $\mathcal{O}^{\text{ideal}}$ that we do not know how to implement efficiently, but which is more convenient for the analysis. We will later replace it by an oracle $\mathcal{O}^{\text{approx}}$ that can be implemented efficiently and whose behavior is very close to the one of $\mathcal{O}^{\text{ideal}}$. Oracle $\mathcal{O}^{\text{ideal}}$ is as follows. On input $(z, t) \in \mathbb{R}^k \times \mathbb{R}^{\geqslant 0}$, it first samples $f' \leftarrow D_{I,\mathbf{S}}$, where $\mathbf{S}$ is a diagonal matrix with

$$
\begin{aligned}
S_{11} &= \sqrt{\exp(t-d)^2|\sigma_1(f) - z\sigma_1(g)|^2 + 4B^2d^2\delta_K^2(|\sigma_1(f)|^2 + |\sigma_1(g)|^2)} \\
S_{22} &= \begin{cases} S_{11} & \text{if } \sigma_1 \text{ is a complex embedding} \\ 2Bd\delta_K\sqrt{(|\sigma_2(f)|^2 + |\sigma_2(g)|^2)} & \text{if } \sigma_1 \text{ is a real embedding} \end{cases} \\
S_{ii} &= 2Bd\delta_K\sqrt{(|\sigma_i(f)|^2 + |\sigma_i(g)|^2)} \qquad \text{if } i \geqslant 3.
\end{aligned}
$$

The astute reader will observe that sampling such an f' may be difficult: this is why we will later introduce $\mathcal{O}^{\text{approx}}$. Oracle $\mathcal{O}^{\text{ideal}}$ then defines $h' = f'/g \bmod q$ (recall that g is invertible modulo q) and returns $\mathcal{A}(h')$.

Note that z and t only appear in S_{11} (and $S_{22} = S_{11}$ if σ_1 is a complex embedding). Since $|\sigma_1(f) - z\sigma_1(g)|/|\sigma_1(g)| = |\sigma_1(h_K) - z|$, we obtain that the success probability of the oracle depends only on $t + \ln|z - \sigma_1(h_K)|$ when t and z vary, as required (recall that h, f and g are fixed once and for all).

In Claim 5.8 below, we show that the oracle $\mathcal{O}^{\text{ideal}}$ satisfies all the desired conditions to be an OHCP oracle and the conditions of Proposition 2.10. This will imply that one can efficiently recover an approximation of $\sigma_1(h_K)$ by using the oracle $\mathcal{O}^{\text{ideal}}$ as a black box.

Claim 5.8. There exist a parameter $\kappa_0 = \mathrm{poly}(\mathrm{Adv}(\mathcal{A})^{-1}, \log q, \log \Delta_K)$ and an algorithm $\mathcal{B}'$ that takes as input any parameter $\kappa \geqslant \kappa_0$ and outputs $\widetilde{\sigma_1(h_K)} \in$

$\sigma_1(K_\mathbb{R})$ such that $|\widetilde{\sigma_1(h_K)} - \sigma_1(h_K)| \leqslant B^2 \cdot \exp(-\kappa)$ with probability $\geqslant 1 - \exp(-\kappa)$. Algorithm $\mathcal{B}'$ runs in time $\mathrm{poly}(\kappa)$ and makes (possibly that many) oracle queries to the OHCP oracle $\mathcal{O}^{\mathrm{ideal}}$ described above.

The difficulty with algorithm $\mathcal{B}'$ from Claim 5.8 is that it makes oracle calls to $\mathcal{O}^{\mathrm{ideal}}$, which we do not know how to run in polynomial time given only access to h and $\mathcal{A}$ (in order to run $\mathcal{O}^{\mathrm{ideal}}$ efficiently, we would probably need to know f and g). To handle this difficulty, we describe another oracle $\mathcal{O}^{\mathrm{approx}}$, whose behavior is very close to the one of $\mathcal{O}^{\mathrm{ideal}}$, but which can be run efficiently.

On input $(z, t) \in \mathbb{R}^k \times \mathbb{R}^{\geqslant 0}$, the randomized oracle $\mathcal{O}^{\mathrm{approx}}$ proceeds as follows. It first samples c_0 in $\mathbb{R}^k$ from the continuous Gaussian distribution $D_{\mathbb{R}^k, \exp(t-d), \mathbf{0}}$; it then defines $\mathbf{c}_1 = \psi(c_0) \in K_\mathbb{R}$ and $\mathbf{c}_2 = \psi(-c_0 \cdot z) \in K_\mathbb{R}$ (where ψ is as defined in Lemma 5.5); the oracle then samples $x \leftarrow \widetilde{D}_{R, 2Bd \cdot \delta_K, \mathbf{c}_1}$ and $y \leftarrow \widetilde{D}_{R, 2Bd \cdot \delta_K, \mathbf{c}_2}$ (see Lemma 2.2); finally, the oracle runs $\mathcal{A}$ on input $\hat{h} = x \cdot h + y \bmod q$, and outputs $\mathcal{A}(\hat{h})$.

Oracle $\mathcal{O}^{\mathrm{approx}}$ can indeed be run in polynomial time from h. Let us now write $\hat{f} = x \cdot f + y \cdot g$, so that $\hat{h} = \hat{f}/g \bmod q$. Observe that $\Pr(\mathcal{O}^{\mathrm{approx}}(z, t) = 1) = \Pr(\mathcal{A}(\hat{h}) = 1)$, and $\Pr(\mathcal{O}^{\mathrm{ideal}}(z, t) = 1) = \Pr(\mathcal{A}(h') = 1)$, where $\hat{h}$ and h' are two random variables. So $|\Pr(\mathcal{O}^{\mathrm{approx}}(z, t) = 1) - \Pr(\mathcal{O}^{\mathrm{ideal}}(z, t) = 1)| \leqslant \mathrm{dist}(\hat{h}, h')$. Since g is fixed, we have $\mathrm{dist}(\hat{h}, h') = \mathrm{dist}(\hat{f}, f')$, and we obtain that

$$|\Pr(\mathcal{O}^{\mathrm{approx}}(z, t) = 1) - \Pr(\mathcal{O}^{\mathrm{ideal}}(z, t) = 1)| \leqslant \mathrm{dist}(\hat{f}, f') \leqslant 2^{-\Omega(d)}.$$

The last inequality comes from Lemma 5.5 and Lemma 2.2.

To conclude, algorithm $\mathcal{B}$ is obtained by taking algorithm $\mathcal{B}'$ of Claim 5.8, but replacing its oracle calls to $\mathcal{O}^{\mathrm{ideal}}$ by oracle calls to $\mathcal{O}^{\mathrm{approx}}$, and taking $\kappa = \max(\kappa_0, d, \ell + 2\ln(B))$. By assumption on $\log q, \mathrm{Adv}(\mathcal{A}), \ell$ and $\log \Delta_K$, we know that $\kappa \leqslant 2^{o(d)}$ (recall that $B \leqslant q$), so that algorithm $\mathcal{B}$ makes at most $2^{o(d)}$ oracle calls to $\mathcal{O}^{\mathrm{approx}}$. Hence, we obtain that

$$|\Pr(\mathcal{B} \text{ succeeds}) - \Pr(\mathcal{B}' \text{ succeeds})| \leqslant 2^{o(d)} \cdot 2^{-\Omega(d)} = 2^{-\Omega(d)}.$$

This completes the proof of Lemma 5.7. □

Proof (Claim 5.8). First, we need to check that the oracle $\mathcal{O}^{\mathrm{ideal}}$ is a valid OHCP oracle. Let us write $z^* = \sigma_1(h_K)$. Since $\sigma_1(h_K) = \sigma_1(f)/\sigma_1(g)$, we know by choice of f and g that $\|z^*\| \in [1/B^2, B^2]$. Hence, the oracle $\mathcal{O}^{\mathrm{ideal}}$ and scale parameter $D = B^2$ form a valid instance of the $(\varepsilon, \delta, \beta)$-OHCP problem (cf Definition 2.9), for any $\varepsilon \in (0, 1)$, any $\delta \in (0, 1/B^4]$ and any $\beta \geqslant 1$.

We will show below that for all $\kappa \geqslant \kappa_0$ with

$$\kappa_0 := \max\left(4\,\mathrm{Adv}(\mathcal{A})^{-1},\, 8d(1 + \ln(q\Delta_K^{1/d})),\, 4\ln(B)\right),$$

the OHCP oracle satisfies the conditions of Proposition 2.10, with

$$p_\infty = \Pr_{u \leftarrow U(R_q)}(\mathcal{A}(u) = 1) \quad \text{and} \quad s^* = 0.$$

More formally, letting $p(t)$ denote $\Pr(\mathcal{O}^{\mathrm{ideal}}(\mathbf{0}, t) = 1)$ as in Proposition 2.10, we prove that

1. $p(s^*) - p_\infty \geqslant 1/\kappa$;
2. $|p(t) - p_\infty| \leqslant 2\exp(-t/\kappa)$ for any $t \geqslant 0$;
3. for any $t_1, t_2 \geqslant 0$, it holds that $|p(t_1) - p(t_2)| \leqslant \kappa\sqrt{|t_1 - t_2|}$.

Using Proposition 2.10, we the conclude that there exists an algorithm $\mathcal{B}'$ solving the $(\exp(-\kappa), \exp(-\kappa), 1 + 1/\kappa)$-OHCP problem in time $\mathrm{poly}(\kappa)$ by making oracle calls to $\mathcal{O}^{\mathrm{ideal}}$. Thanks to the condition $\kappa \geqslant 4\ln(B)$, it holds that $\exp(-\kappa) \leqslant 1/B^4$ is a valid choice of δ. Moreover, using the fact that $B \leqslant q$, we see that $\kappa_0 = \mathrm{poly}(\mathrm{Adv}(\mathcal{A})^{-1}, \log q, \log \Delta_K)$, which proves Claim 5.8. We now proceed to prove the three properties above.

Property 1. We want to show that $p(s^*)$ is very close to $\Pr(\mathcal{A}(\phi_B(h)) = 1)$, which will allow us to conclude with Lemma 5.6. Observe that by definition of the OHCP oracle $\mathcal{O}^{\mathrm{ideal}}$, we know that $p(s^*) = \Pr(\mathcal{A}(h') = 1)$, where $h' = f'/g \bmod q$. So in order to bound the difference between $\Pr\left(\mathcal{A}(\phi_B(h) = 1\right)$ and $p(s^*)$, it suffices to bound the statistical distance between the two random variables $\phi_B(h)$ and h', which is equivalent to bounding $\mathrm{dist}(g \cdot \phi_B(h), f')$ (i.e., it suffices to consider the numerator since the denominator is g in both cases).

Using Lemma 5.5 with $c_1 = c_2 = 0$ and $s = 2Bd\delta_K$, we know that the distribution of $g \cdot \phi_B(h)$ is within $2^{-\Omega(d)}$ statistical distance from $D_{I,\mathbf{S}_2,\mathbf{0}}$, where $\mathbf{S}_2$ is a diagonal matrix with i-th diagonal entry equal to $2Bd\delta_K \cdot \sqrt{|\sigma_i(f)|^2 + |\sigma_i(g)|^2}$. Moreover, by definition of $\mathcal{O}^{\mathrm{ideal}}$, the distribution of f' is $D_{I,\mathbf{S}_1,\mathbf{0}}$, where $\mathbf{S}_1$ is identical to $\mathbf{S}_2$, except for first diagonal coefficient (or first two diagonal coefficients if σ_1 is complex), which is equal to

$$\sqrt{(2Bd\delta_K)^2(|\sigma_1(f)|^2 + |\sigma_1(g)|^2) + \exp(-2d) \cdot |\sigma_1(f)|^2}.$$

We now apply Lemma 2.3 to show that these two Gaussian distributions are statistically close. We first check that $\eta_{1/2}(\mathbf{S}_i^{-1}I) \leqslant 1/2$, for $i \in \{1, 2\}$. We know from Equation (2.1) that

$$\begin{aligned}\eta_{1/2}(\mathbf{S}_i^{-1}I) &\leqslant \sqrt{\frac{\ln(2d(1+2))}{\pi}} \cdot \lambda_d(\mathbf{S}_i^{-1}I) \\ &\leqslant \sqrt{d} \cdot \lambda_d(\mathbf{S}_i^{-1}I)\end{aligned}$$

Recall that $I = \langle f \rangle + \langle g \rangle$, so that $f \in I$. Hence, we know that the $\mathbf{S}_i^{-1} \cdot f \cdot r_j$'s are linearly independent vectors of $\mathbf{S}_i^{-1} \cdot I$ (recall that the r_j's form a basis of R). For every j, it holds that $\|\mathbf{S}_i^{-1} \cdot f \cdot r_j\| \leqslant \delta_K \cdot \|\mathbf{S}_i^{-1} \cdot f\| \leqslant \delta_K \cdot \sqrt{d}/(2Bd\delta_K)$ (since every diagonal coefficient of $\mathbf{S}_i$ is no smaller than the corresponding coefficient of f multiplied by $2Bd\delta_K$). Hence, we conclude that $\lambda_d(\mathbf{S}_i^{-1}I) \leqslant 1/(2\sqrt{d})$ and that $\eta_{1/2}(\mathbf{S}_i^{-1}I) \leqslant 1/2$, as desired. We can apply Lemma 2.3 and we obtain that

$$\mathrm{dist}(D_{I,\mathbf{S}_1,\mathbf{0}}, D_{I,\mathbf{S}_2,\mathbf{0}}) \leqslant 4\sqrt{d} \cdot \sqrt{\left\|\left|\mathbf{S}_2^{-1}\mathbf{S}_1 - \mathbf{I}_d\right|\right\|}.$$

The matrix $\mathbf{S}_2^{-1}\mathbf{S}_1 - \mathbf{I}_d$ is zero, except for the top-left coefficient (or for the first two top-left coefficients if σ_1 is a complex embedding), which is equal to $\sqrt{1+\eta}-$

1 where $\eta = \exp(-2d) \cdot |\sigma_1(f)|^2/((2Bd\delta_K)^2 \cdot (|\sigma_1(f)|^2 + |\sigma_1(g)|^2))$. Since $\eta \leqslant \exp(-2d)$, we conclude that $|\sqrt{1+\eta} - 1| \leqslant \exp(-2d)$, and so $\left\|\mathbf{S}_2^{-1}\mathbf{S}_1 - \mathbf{I}_d\right\| \leqslant \exp(-2d)$ (or $\leqslant 2\exp(-2d)$ in case we had two non-zero coefficients). We finally obtain that $D_{I,\mathbf{S}_1,\mathbf{0}} \approx_{2^{-\Omega(d)}} D_{I,\mathbf{S}_2,\mathbf{0}}$, which in turn shows that

$$|p(s^*) - \Pr\left(\mathcal{A}(\phi_B(h) = 1\right)| \leqslant 2^{-\Omega(d)}.$$

Finally, since $h \in H$, we know from Lemma 5.6 that $|\Pr(\mathcal{A}(\phi_B(h) = 1) - p_\infty| \geqslant \mathrm{Adv}(\mathcal{A})/2$. Wlog, we can assume that $\Pr(\mathcal{A}(\phi_B(h) = 1) - p_\infty \geqslant 0$ (otherwise we can simply consider $\mathcal{A}' = 1 - \mathcal{A}$), from which we obtain that

$$p(s^*) - p_\infty \geqslant \mathrm{Adv}(\mathcal{A})/2 - 2^{-\Omega(d)} \geqslant \mathrm{Adv}(\mathcal{A})/4,$$

where the last inequality holds asymptotically when d tends to infinity, since we assumed that $1/\mathrm{Adv}(\mathcal{A}) \leqslant 2^{o(d)}$. By choice of κ, this implies that $p(s^*) - p_\infty \geqslant 1/\kappa$.

Property 2. To prove this second property, we want to show that when t is sufficiently large, then the distribution of $f' \bmod q$ (where f' is implicitly computed by the oracle $\mathcal{O}^{\mathrm{ideal}}$ as defined above) is statistically close to uniform in $R \bmod qR$. Recall that the support of f' is I, which may be a strict subset of R. However, we know that $I = \langle f\rangle + \langle g\rangle$ is coprime to $\langle q\rangle$. So if $\widetilde{f} \in I$ is uniform in $I/(qI)$, then $\widetilde{f} + qR$ is a uniform class of $R/(qR)$. Hence, it suffices to show that f' is statistically close to uniform in $I/(qI)$.

Recall that f' is sampled from the distribution $D_{I,\mathbf{S}}$, where $\mathbf{S}$ is a diagonal matrix with positive diagonal coefficients, with $S_{11} \geqslant \exp(t-d) \cdot |\sigma_1(f)|$ (we consider $z = 0$ here) and $S_{ii} \geqslant |\sigma_i(f)|$ for $i \geqslant 2$. Taking the product, we conclude that $\prod_i S_{ii} \geqslant \exp(t-d) \cdot \mathcal{N}(f)$. Let us call c the quantity $c = (\exp(t-d)\mathcal{N}(f)/(\mathcal{N}(qI)\cdot\Delta_K))^{1/d}$. Using Lemma 2.8, we know that when t is sufficiently large so that $c \geqslant 1$, then it holds that $1 \geqslant \eta_\varepsilon(\mathbf{S}^{-1}\cdot(qI))$ for $\varepsilon = \exp(-c^2 d)$. Moreover, applying Lemma 2.4 to $L_1 = \mathbf{S}^{-1} \cdot (qI)$ and $L_2 = \mathbf{S}^{-1} \cdot I$, we see that

$$\mathrm{dist}\left(D_{\mathbf{S}^{-1}\cdot I,1} \bmod \mathbf{S}^{-1} \cdot (qI), U(\mathbf{S}^{-1} \cdot I \bmod \mathbf{S}^{-1} \cdot (qI))\right) \leqslant 2\exp(-c^2 d).$$

Multiplying the outputs of these two distributions by $\mathbf{S}$, we finally obtain

$$\mathrm{dist}\left(D_{I,\mathbf{S}} \bmod qI, U(I \bmod qI)\right) \leqslant 2\exp(-c^2 d).$$

Using the fact that $c^2 \geqslant c$ (as $c \geqslant 1$), that $\exp(x) \geqslant x$ for all $x \in \mathbb{R}$, and that $\mathcal{N}(I) \leqslant \mathcal{N}(f)$, we obtain the upper bound

$$\begin{aligned} 2\exp(-c^2 d) \leqslant 2\exp(-cd) &\leqslant 2\exp\left(-e^{(t-d-\ln(q^d\Delta_K))/d} \cdot d\right) \\ &\leqslant 2\exp\left(-(t - d(1 + \ln(q\Delta_K^{1/d})))\right). \end{aligned}$$

If $t \geqslant 2d(1+\ln(q\Delta_K^{1/d}))$, then $(t - d(1+\ln(q\Delta_K^{1/d})) \geqslant t/2$ and $c \geqslant 1$, which implies that

$$|p(t) - p_\infty| \leqslant 2\exp(-t/2) \leqslant 2\exp(-t/\kappa).$$

For smaller t, note that $t \leqslant \kappa/2$. In this case, the upper bound $2\exp(-t/\kappa)$ is at least 1, and so the property is also satisfied.

Property 3. Let us fix some $t_1 \geqslant t_2 \geqslant 0$. We want to show that $|p(t_1) - p(t_2)| \leqslant \kappa \cdot \sqrt{|t_1 - t_2|}$. Observe first that since p takes values in $[0, 1]$ and $\kappa \geqslant 1$, then the condition is always satisfied when $|t_1 - t_2| \geqslant 1$. We will hence assume wlog that $0 \leqslant t_1 - t_2 \leqslant 1$.

We know from the definition of $\mathcal{O}^{\text{ideal}}$ that $|p(t_1) - p(t_2)| \leqslant \text{dist}(D_{I,\mathbf{S}_1}, D_{I,\mathbf{S}_2})$, where $\mathbf{S}_1$ and $\mathbf{S}_2$ are diagonal and equal, except for their for top-left coefficient (or two top-left coefficients if σ_1 is a complex embedding):

$$(S_1)_{11} = \sqrt{c + (\exp(t_1 - d)|\sigma_1(f)|)^2} \quad \text{and} \quad (S_2)_{11} = \sqrt{c + (\exp(t_2 - d)|\sigma_1(f)|)^2},$$

for some $c \geqslant 0$. As when proving Property 1, one can check that $\eta_{1/2}(\mathbf{S}_1^{-1} I)$, $\eta_{1/2}(\mathbf{S}_2^{-1} I) \leqslant 1/2$. Therefore, we can apply Lemma 2.3 to obtain that

$$\text{dist}\left(D_{I,\mathbf{S}_1}, D_{I,\mathbf{S}_2}\right) \leqslant 4\sqrt{d} \cdot \sqrt{\left\|\left\|\mathbf{S}_2^{-1}\mathbf{S}_1 - \mathbf{I}_d\right\|\right\|}.$$

Once again, the matrix $\mathbf{S}_2^{-1}\mathbf{S}_1 - \mathbf{I}_d$ is zero, except for its top-left coefficient (or two top-left coefficients) which is equal to

$$\sqrt{\frac{c + (\exp(t_1 - d)|\sigma_1(f)|)^2}{c + (\exp(t_2 - d)|\sigma_1(f)|)^2}} - 1 \leqslant \sqrt{\frac{(\exp(t_1 - d)|\sigma_1(f)|)^2}{(\exp(t_2 - d)|\sigma_1(f)|)^2}} - 1 = \exp(t_1 - t_2) - 1.$$

The first inequality comes from the fact that $t_1 \geqslant t_2$ (and c and $(\exp(t_2 - d)|\sigma_1(f)|)^2$ are non-negative). Finally, since $0 \leqslant t_1 - t_2 \leqslant 1$, we conclude that $\exp(t_1 - t_2) - 1 \leqslant 2|t_1 - t_2|$. This in turns implies that $|p(t_1) - p(t_2)| \leqslant 8\sqrt{d}\sqrt{|t_1 - t_2|} \leqslant \kappa\sqrt{|t_1 - t_2|}$, as desired. □

Acknowledgments. The authors thank Koen de Boer, Léo Ducas, Guillaume Hanrot, Miruna Rosca aux Adeline Roux-Langlois for insightful discussions. The first author was supported in part by CyberSecurity Research Flanders with reference number VR20- 192203 and by the Research Council KU Leuven grant C14/18/067 on Cryptanalysis of Post-quantum Cryptography. The second author was supported in part by European Union Horizon 2020 Research and Innovation Program Grant 780701 and BPI-France in the context of the national project RISQ (P141580).

References

ABD16. M. R. Albrecht, S. Bai, and L. Ducas. A subfield lattice attack on overstretched NTRU assumptions - cryptanalysis of some FHE and graded encoding schemes. In *CRYPTO*, 2016.

AD17. M. R. Albrecht and A. Deo. Large Modulus Ring-LWE ⩾ Module-LWE. In *ASIACRYPT*, 2017.

BBC+20. D. J. Bernstein, Brumley B. B., M.-S. Chen, C. Chuengsatiansup, T. Lange, A. Marotzke, B.-Y. Peng, N. Tuveri, C. van Vredendaal, and B.-Y. Yang. NTRU Prime round-3 candidate to the NIST post-quantum cryptography standardisation project, 2020. Available at https://ntruprime.cr.yp.to/.

Bel04a. K. Belabas. A relative van Hoeij algorithm over number fields. *J Symb Comput*, 37(5), 2004.

Bel04b. K. Belabas. Topics in computational algebraic number theory. *J. théorie des nombres de Bordeaux*, 16, 2004.

CDH+20. C. Chen, O. Danba, J. Hoffstein, A. Hülsing, J. Rijneveld, T. Saito, J. M. Schank, P. Schwabe, W. Whyte, K. Xagawa, T. Yamakawa, and Z. Zhang. NTRU round-3 candidate to the NIST post-quantum cryptography standardisation project, 2020. Available at https://ntru.org/.

CDW21. R. Cramer, L. Ducas, and B. Wesolowski. Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time. *J. ACM*, 68(2), 2021.

CJL16. J. H. Cheon, J. Jeong, and C. Lee. An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without an encoding of zero. *LMS J Comput Math*, 19(A), 2016.

Coh95. H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer, 1995.

Coh00. H. Cohen. *Advanced topics in computational number theory*. Springer, 2000.

dBDPW20. K. de Boer, L. Ducas, A. Pellet-Mary, and B. Wesolowski. Random self-reducibility of Ideal-SVP via Arakelov random walks. In *CRYPTO*, 2020.

FS10. C. Fieker and D. Stehlé. Short bases of lattices over number fields. In *ANTS*, 2010.

GGH13. S. Garg, C. Gentry, and S. Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, 2013.

GPV08. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.

HHP+03. J. Hoffstein, N. Howgrave-Graham, J. Pipher, J. H. Silverman, and W. Whyte. NTRUSIGN: digital signatures using the NTRU lattice. In *CT-RSA*, 2003.

HPS98. J. Hoffstein, J. Pipher, and J. H. Silverman. NTRU: a ring based public key cryptosystem. In *ANTS*, 1998.

KF15. P. Kirchner and P.-A. Fouque. An improved BKW algorithm for LWE with applications to cryptography and lattices. In *CRYPTO*, 2015.

KF17. P. Kirchner and P.-A. Fouque. Revisiting lattice attacks on overstretched NTRU parameters. In *EUROCRYPT*, 2017.

KLL84. R. Kannan, A. K. Lenstra, and L. Lovász. Polynomial factorization and nonrandomness of bits of algebraic and some transcendental numbers. In *STOC*, 1984.

LLL82. A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Math Ann*, 1982.

LPR10. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, 2010.

LS15. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Des Codes Cryptography*, 2015.

LSS14. A. Langlois, D. Stehlé, and R. Steinfeld. GGHLite: More efficient multilinear maps from ideal lattices. In *EUROCRYPT*, 2014.

LTV12. A. López-Alt, E. Tromer, and V. Vaikuntanathan. On-the-fly multiparty computation on the cloud via multikey fully homomorphic encryption. In *STOC*, 2012.

MR07. D. Micciancio and O. Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1), 2007.

Pei10. C. Peikert. An efficient and parallel gaussian sampler for lattices. In *CRYPTO*, 2010.

Pei16. C. Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4), 2016.

PHS19. A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-SVP in ideal lattices with pre-processing. In *EUROCRYPT*, 2019.

PRS17. C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *STOC*, 2017.

Reg09. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J ACM*, 2009.

Rob97. F.-X. Roblot. *Algorithmes de factorisation dans les extensions relatives et applications de la conjecture de Stark à la construction des corps de classes de rayon*. PhD thesis, Université Bordeaux 1, 1997. Available at http://math.univ-lyon1.fr/~roblot/resources/these.pdf.

RSW18. M. Rosca, D. Stehlé, and A. Wallet. On the ring-LWE and polynomial-LWE problems. In *EUROCRYPT*, 2018.

Sch87. C.-P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theoretical Computer Science*, 53, 1987.

SS11. D. Stehlé and R. Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, 2011.

SS13. D. Stehlé and R. Steinfeld. Making NTRUEncrypt and NTRUSign as secure as standard worst-case problems over ideal lattices. *IACR ePrint 2013/004*, 2013.

SSTX09. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, 2009.

Sut16. A. Sutherland. Lecture notes of Number Theory I, taught at MIT. Available at https://math.mit.edu/classes/18.785/2016fa/LectureNotes12.pdf, 2016.

WW18. Y. Wang and M. Wang. Provably secure NTRUEncrypt over any cyclotomic field. In *SAC*, 2018.