Ideal-SVP is Hard for Small-Norm Uniform Prime Ideals

Joël Felderhoff[1], Alice Pellet-Mary[2], Damien Stehlé[3], and Benjamin Wesolowski[4]

[1] Inria Lyon, Lyon, France, ENS de Lyon and LIP, UMR 5668, Lyon, France, joel.felderhoff@ens-lyon.fr
[2] Univ. Bordeaux, CNRS, Inria, Bordeaux INP, IMB, Talence, France, alice.pellet-mary@math.u-bordeaux.fr
[3] ENS de Lyon and CryptoLab, Inc. damien.stehle@cryptolab.co.kr
[4] ENS de Lyon, CNRS, UMPA, UMR 5669, Lyon, France benjamin.wesolowski@ens-lyon.fr

Abstract. The presumed hardness of the Shortest Vector Problem for ideal lattices (Ideal-SVP) has been a fruitful assumption to understand other assumptions on algebraic lattices and as a security foundation of cryptosystems. Gentry [CRYPTO'10] proved that Ideal-SVP enjoys a worst-case to average-case reduction, where the average-case distribution is the uniform distribution over the set of *inverses of prime ideals* of small algebraic norm (below $d^{O(d)}$ for cyclotomic fields, where d refers to the field degree). De Boer et al. [CRYPTO'20] obtained another random self-reducibility result for an average-case distribution involving *integral ideals* of norm $2^{O(d^2)}$.

In this work, we show that Ideal-SVP for the uniform distribution over inverses of small-norm prime ideals reduces to Ideal-SVP for the uniform distribution over small-norm prime ideals. Combined with Gentry's reduction, this leads to a worst-case to average-case reduction for the uniform distribution over the set of *small-norm prime ideals*. Using the reduction from Pellet-Mary and Stehlé [ASIACRYPT'21], this notably leads to the first distribution over NTRU instances with a polynomial modulus whose hardness is supported by a worst-case lattice problem.

1 Introduction

Lattice-based cryptography is built upon on the hardness of a variety of computational problems related to the shortest vector problem (SVP), consisting in finding a shortest non-zero vector in a Euclidean lattice, possibly up to some approximation factor. As generic lattices typically lead to poor performance, cryptographic schemes often use so-called algebraic variants. The case of *ideal lattices* has attracted particular attention since the introduction of Ring-SIS [Mic02,LM06,PR06] and Ring-LWE [SSTX09,LPR10,PRS17]. These problems have both enabled the construction of very efficient cryptosystems, and are known to be at least as hard as finding short non-zero vectors in ideal lattices.

Let $K = \mathbb{Q}[X]/P(X)$ be a number field of degree d (i.e., $P(X) \in \mathbb{Z}[X]$ is an irreducible polynomial of degree d) and let $\mathcal{O}_K$ be its ring of integers. As K is naturally a Hermitian vector space (via the canonical embedding $\sigma : K \to \mathbb{C}^d$), any ideal in $\mathcal{O}_K$ is a Euclidean lattice. Such ideals, with the associated lattice structure, are *ideal lattices*. In this work, we focus on the id-HSVP problem (ideal Hermite Shortest Vector Problem), an approximate version of SVP for ideal lattices consisting in finding a non-zero-vector in the ideal lattice whose norm is within a given factor of the root determinant. Note that id-HSVP and SVP for ideal lattices are equivalent up to some small parameter losses.

In cryptographic applications, it is typically insufficient for a computational problem to be hard in the worst case: one needs instances to be hard *on average* for some distribution with non-trivial entropy. Such a guarantee can be provided by proving worst-case to average-case reductions: if there is an algorithm which performs well on random instances of problem A with non-negligible probability (i.e., for the average case), then there is an algorithm which performs well for any instance of problem B (i.e., for the worst case). When the two problems are the same (up to the approximation factor), we may refer to this property as random self-reducibility. Note that self-reducibility is associated not only to a problem, but also to a distribution on the instances, and the choice of distribution may be of critical importance.

The first random self-reducibility result for id-HSVP was proven by Gentry [Gen09a,Gen10], for supporting the security of the first fully homomorphic encryption scheme [Gen09a,Gen09b]. Gentry proved that id-HSVP in the worst case reduces to id-HSVP for *the inverse of a uniformly chosen prime ideal among those with algebraic norm in a prescribed interval* $[A, B]$. Note that Gentry states this result in terms of the Bounded Distance Decoding problem for the prime ideals themselves (not their inverses) – the two formulations are equivalent thanks to Regev's quantum reduction from SIVP in a lattice L to BDD in the dual lattice $L^\vee$ [Reg05]. Gentry's reduction enables interval boundaries A and B that have a bounded ratio and can be chosen as small as $\Delta_K^{O(1)} \cdot d^{O(d)}$, where Δ_K and d respectively refer to the field discriminant and degree.[5] A weaker result is proved in [Gen10], but it can be boosted as detailed in [Gen09a]. This reduction requires an ideal-factoring oracle (which can be implemented in quantum polynomial time using Shor's algorithm [Sho94]) but is otherwise polynomial-time, and introduces a loss in the approximation factor that is bounded as $\Delta_K^{O(1/d)} \cdot d^{O(1)}$.

A different average-case distribution is considered in [BDPW20]. The space of all ideal lattices, up to isometries, is itself an arithmetic-geometric object, the Arakelov class group, and comes with a natural notion of "uniform distribution". Mathematically, this distribution is convenient because of the rich theory surrounding it. Computationally, one cannot work directly with it: first because it is continuous, and second because we do not have a canonical way of representing the isometry class of a lattice. Thanks to an appropriate rounding procedure, de Boer et al. [BDPW20] introduced a distribution on ideals of

[5] For the sake of simplicity, we assume for the introduction that we are given a basis of the ring of integers $\mathcal{O}_K$ whose vectors have norms $\leq \Delta_K^{O(1/d)} \cdot d^{O(1)}$

norm $\Delta_K^{O(1)} \cdot 2^{O(d^2)}$ that mimics this continuous distribution and proved the self-reducibility of id-HSVP for this distribution.[6] This reduction is polynomial and also incurs an approximation factor loss of $\Delta_K^{O(1/d)} \cdot d^{O(1)}$. Note that the ideals of this distribution have much larger norms than those of Gentry's reduction. This unfortunately leads to cryptographic instances of larger sizes. In [PS21], the authors observed that the algebraic norm reached by the reduction from de Boer et al. can be decreased, but at the expense of a super-polynomial running-time.

The above discussion raises the following question:

Can we reduce id-HSVP *in the worst-case to* id-HSVP *for uniform prime ideals of norms bounded as* $\Delta_K^{O(1)} \cdot d^{O(d)}$*, in time polynomial in* $\Delta_K^{1/d}$ *and* d*?*

Contributions. We describe a new quantum self-reduction for id-HSVP. We prove that if $\mathcal{W}$ is a set of ideals and $\mathcal{W}^{-1}$ is the set of inverses of the ideals of $\mathcal{W}$, then solving id-HSVP for the uniform distribution over $\mathcal{W}^{-1}$ reduces to solving id-HSVP for the uniform distribution over $\mathcal{W}$ and to solving id-HSVP for a uniform ideal within those having their norms in a prescribed interval. Both the cost of the reduction and the loss in the approximation factor are polynomially bounded in the degree d and the root-discriminant $\Delta_K^{1/d}$ of the number field. The precise statement is provided in Theorem 5.1.

When specialized with $\mathcal{W}$ chosen as the set of prime ideals of algebraic norm $\Delta_K^{O(1)} \cdot d^{O(d)}$, our reduction implies that solving id-HSVP for the inverse of uniform primes ideals is no harder than solving it for uniform prime ideals (still for those of algebraic norm $\Delta_K^{O(1)} \cdot d^{O(d)}$). The success probability of this reduction is proportional to the proportion of prime ideals among all integral ideals of norm bounded by some $A = \mathrm{poly}(\Delta_K)$. Combined with Gentry's reduction [Gen09a], our work implies the random self-reducibility of id-HSVP for the uniform distribution over prime ideals. As Gentry's original reduction considers the bounded distance decoding problem, we present an adaptation to the shortest vector problem in Appendix C. Note that the polynomial dependency in the proportion of prime ideals may have a considerable impact on the cost of this reduction (there exists number fields for which the proportion of prime ideals is exponentially small in the degree).

This new reduction, along with the Karp reduction of [PS21], gives a new distribution over NTRU instances with modulus polynomial in d and $\Delta_K^{1/d}$ whose difficulty relies on the worst-case problem id-HSVP. To our knowledge this is the first time a distribution over NTRU instance with polynomial modulus is based on a worst-case problem, even though this distribution needs a factoring oracle to be sampled from.

Technical overview. We now give an overview of the average-case to average-case reduction for id-HSVP. Let $\mathcal{W}$ be a set of fractional ideals represented

[6] The bound on the norm is obtained by combining Lemma 4.1 and Theorem 4.5 from [BDPW20].

by their Hermite Normal Form. The goal of our reduction is to find (with non-negligible probability) a short non-zero vector in a given uniform element of $\mathcal{W}^{-1}$, given access to two oracles: $\mathcal{O}_{\mathcal{W}}$ which solves id-HSVP with non-negligible probability for a uniform element of $\mathcal{W}$, and $\mathcal{O}_{\mathcal{I}}$ which solves it with non-negligible probability for a uniform integral ideal with norm between A and $4A$, for $A = \Delta_K^{O(1/d)} \cdot d^{O(1)}$. In everything that follows we assume that we have a factoring oracle (for integers, or equivalently, for integral ideals). Such an oracle can be instantiated in quantum polynomial time with Shor's algorithm, or in sub-exponential time with the number field sieve algorithm.

Before diving into our contribution, let us explain a key idea developed in [Boe22, Chap. 6]. By *ideal of norm 1*, we mean a (replete) ideal[7] of the form $I/\mathcal{N}(I)^{1/d}$. The space of ideals of norm 1 has a natural notion of uniformity. Let B_r denote the ℓ_∞ ball of radius r. In [Boe22, Th. 6.21], it is proved that if J is sampled uniformly in the set of ideals of norm 1, and x is uniform in $B_r \cap J$, then the integral ideal $x \cdot J^{-1}$ is almost uniform in the set of integral ideals of norm less than r^d.

Now, our reduction follows the following structure. We are given a uniform $I \in \mathcal{W}$, and tasked with finding a short non-zero vector $v_{I^{-1}} \in I^{-1}$.

1. Find a short non-zero vector $v_I \in I$ with the oracle $\mathcal{O}_{\mathcal{W}}$.
2. Generate a uniform norm-1 ideal I', together with a short non-zero vector $v_{I'} \in I'$. The ideal $J = I' \cdot I/\mathcal{N}(I)^{1/d}$ is also uniform in the space of ideals of norm 1, and we can compute a short basis $\mathbf{B}_J$ of J thanks to the short non-zero vectors v_I and $v_{I'}$.
3. Sample $x \in B_r \cap J$; this uses our knowledge of the good basis $\mathbf{B}_J$. Hopefully, the integral ideal $\mathfrak{b} = x \cdot J^{-1}$ is almost uniform in the set of integral ideals of bounded norm.
4. Find a short non-zero vector $v_{\mathfrak{b}} \in \mathfrak{b}$ with the oracle $\mathcal{O}_{\mathcal{I}}$.
5. Return the vector $v_{I^{-1}} = x^{-1} \cdot v_{I'} \cdot v_{\mathfrak{b}} \cdot \mathcal{N}(I)^{-1/d} \in I^{-1}$.

One can check that $v_{I^{-1}} \in I^{-1}$, but is it short? Its factors are short by construction, except possibly x^{-1}. Indeed, the element x itself is bounded (it is in the set B_r), but its inverse may not be. To circumvent this issue, we would like to replace the ℓ_∞ ball B_r with another shape X which contains only *balanced* vectors (i.e., close to a vector of the form $\lambda \cdot (1, \ldots, 1)$), so that for any short $x \in X$, we have that x^{-1} is small. We prove that the result of [Boe22] holds for general sets X satisfying certain conditions. We consider a new shape $\mathcal{B}^{\eta}_{A,B}$ (see Figure 1 and Definition 4.1) that verifies the conditions, and contains only balanced elements. Now, replacing B_r with $\mathcal{B}^{\eta}_{A,B}$ in Step 3, we sample an element x such that x^{-1} is small, hence all the factors of $v_{I^{-1}}$ are small, and $v_{I^{-1}}$ is indeed a solution to id-HSVP in I^{-1}.

While Step 3 constitutes the main difficulty of the reduction, and the technical core of our paper, let us briefly comment on Step 2. We need to sample a uniform norm-1 ideal I', together with a short non-zero vector $v_{I'} \in I'$.

[7] A replete ideal is a subset of $K_{\mathbb{R}} := K \otimes_{\mathbb{Q}} \mathbb{R}$ of the form $\alpha \cdot I$ where $I \subseteq \mathcal{O}_K$ is an integral ideal of $\mathcal{O}_K$ and $\alpha \in K_{\mathbb{R}}^{\times}$ is invertible. More details can be found in the preliminaries.

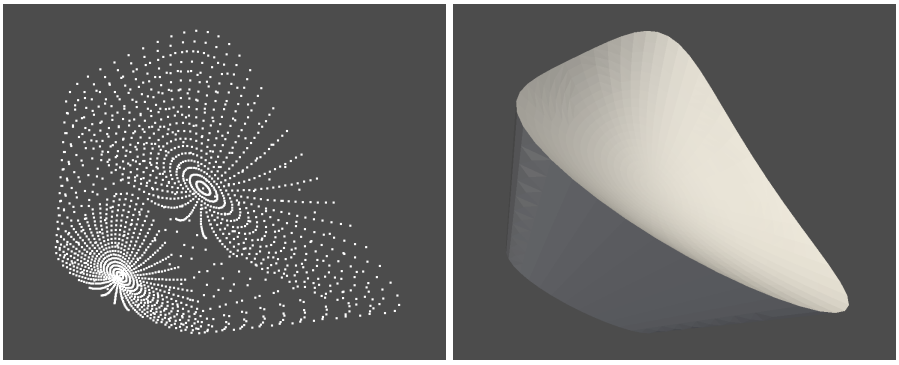

Fig. 1. A plot of $\mathcal{B}^{\eta}_{A,B}$ intersected with the subspace $K^{+}_{\mathbb{R}} := \{x \in K_{\mathbb{R}} \,|\, \sigma_i(x) \in \mathbb{R}_{>0} \text{ for all } i\}$. Here we have $(d_{\mathbb{R}}, d_{\mathbb{C}}) = (3, 0)$, $A = 20, B = 40$ and $\eta = \exp(1)$.

In [BDPW20], it is proven that if an ideal $\mathfrak{p}$ is sampled uniformly in the set of prime ideals with norm less than $(d^d \cdot \Delta_K)^c$ for some constant c, then, up to a small Gaussian factor, the ideal $\mathfrak{p}/\mathcal{N}(\mathfrak{p})^{1/d}$ is close to uniform in the set of norm-1 ideals. It is therefore sufficient to generate such a prime ideal $\mathfrak{p}$ together with a short element $v_{\mathfrak{p}} \in \mathfrak{p}$. The technique is extracted from [Gen09a, Chap. 17], and requires a factoring oracle. It first samples a small element $x \in \mathcal{O}_K$ with the Gaussian distribution. It then factors $(x) = \mathfrak{p}_1^{e_1} \cdot \ldots \cdot \mathfrak{p}_k^{e_k}$ and uniformly selects one of the factors $\mathfrak{p}_i$. Finalizing with a rejection sampling step, it can be proved that the chosen $\mathfrak{p}$ is almost uniform in the set of primes of norm $\lesssim \mathcal{N}(x)$.

We now have a reduction from id-HSVP for inverses of ideal of a set $\mathcal{W}$, to id-HSVP for ideals of $\mathcal{W}$ and id-HSVP for a uniform ideal of norm in some interval $[A, 4A]$ for A as small as $\Delta_K^{O(1))} \cdot d^{O(d)}$. This gives a trivial reduction from id-HSVP for a uniform ideal to id-HSVP for an uniform prime ideal, with a success probability decrease of a factor $O(1/\tilde{\rho}_A)$, where $1/\tilde{\rho}_A$ is the proportion of prime ideals among the set of all integral ideals of norm $\leq A$. We can now combine this last reduction with our main result (taking $\mathcal{W}$ to be the set of prime ideals of norm in $[A, 4A]$) in order to reduce id-HSVP for inverses of prime ideals to id-HSVP for prime ideals. This, combined with the worst-case to average-case reduction of [Gen09a] gives a worst-case to average-case reduction for id-HSVP where the average-case is the uniform distribution over prime ideals of norm in $[A, 4A]$.

Finally, note that a reduction from id-HSVP to NTRU was recently given in [PS21]. It transforms an integral ideal I into an NTRU instance of modulus polynomial larger than $\mathcal{N}(I)^{1/d}$. Our self-reduction (in contrast with the one from [Gen09a]) applies to integral ideals and can be composed with the one from [PS21]. The distribution of NTRU instances obtained by sampling a uniform prime ideal of norm in $[A, 4A]$ and applying [PS21, Alg. 4.1] is at least as difficult to solve as worst-case id-HSVP. By setting $A = \Delta_K^{O(1)} \cdot d^{O(d)}$, we obtain

an NTRU modulus bounded as $\Delta_K^{O(1/d)} \cdot d^{O(1)}$. Note that "overstretched NTRU" attacks [ABD16,CJL16,KF17] do not apply for this distribution as, among others, they require a much larger modulus.

Related works on the hardness of id-HSVP. On the upper bound front, it has been shown that id-HSVP is susceptible to better algorithms than the generic HSVP. Cramer et al. [CDPR16] described an algorithm for id-HSVP in cyclotomic fields for principal ideals with an approximation factor $\exp \widetilde{O}(\sqrt{d})$ in quantum polynomial time. It was later generalized to all ideals [CDW17] of cyclotomic fields and (with pre-processing) to all number fields [PHS19]. Note that in the present work, all our reductions feature polynomial losses on the approximation factor, and hence apply to id-HSVP for polynomial approximation factors, a regime that is not impacted by these algorithms. Still, families of easy instances for id-HSVP have been identified even for polynomial approximation factors [PXWC21,PML21,BEP22], specifically ideals stabilized by many field automorphisms. While these families are very sparse, their existence further motivates the study of different distributions of id-HSVP instances.

2 Preliminaries

The notation ln will refer to the base-e logarithm. For any function $f : X \to \mathbb{R}$ and $S \subseteq X$ with S countable, we define $f(S) := \sum_{x \in S} f(x)$.

We will let $\mathcal{G}(\mathbf{c}, s)$ denote the continuous Gaussian distribution of center $\mathbf{c}$ and of standard deviation s on some vector spaces that will always be specified. We will use both the statistical distance and Rényi divergence between distributions. Let D_1, D_2 be distributions over a countable set X. Their statistical distance is $\mathrm{SD}(D_1, D_2) = \sum_{x \in X} |D_1(x) - D_2(x)|/2$. For any event $E \subseteq X$, we have $D_2(E) \geq D_1(E) - \mathrm{SD}(D_1, D_2)$. If $\mathrm{Supp}(D_1) \subseteq \mathrm{Supp}(D_2)$, their Rényi divergence of infinite order is $\mathrm{RD}_\infty(D_1 \parallel D_2) = \max_{x \in \mathrm{Supp}(D_1)} D_1(x)/D_2(x)$. For any event $E \subseteq \mathrm{Supp}(D_1)$, we have $D_2(E) \geq D_1(E)/\mathrm{RD}_\infty(D_1 \parallel D_2)$.

When using oracles with a non-zero probability of failing, we assume without loss of generality that either the oracle returns a valid result or $\perp$ (as in our cases, the validity of the output can always be checked efficiently).

2.1 Lattices

Let $n \geq 1$. A lattice in $\mathbb{R}^n$ is a discrete additive subgroup of $\mathbb{R}^n$ of the form $L = \sum_{1 \leq i \leq k} \mathbf{b}_i \cdot \mathbb{Z}$ for some linearly independent $\mathbf{b}_i \in \mathbb{R}^n$ that are said to form a basis of L. A lattice is said to be full rank if any of its bases is full-rank in $\mathbb{R}^n$. If L is a lattice, we define its covering radius as $\mu(L) = \inf_{x \in \mathbb{R}^n} \mathrm{dist}(x, \mathcal{L})$ and its volume as $\mathrm{vol}(L) = \sqrt{\det(\mathbf{B} \cdot \mathbf{B}^T)}$, where $\mathbf{B} = (\mathbf{b}_1 \| \ldots \| \mathbf{b}_k)$ (this quantity is independent of the choice of basis $(\mathbf{b}_1 \| \ldots \| \mathbf{b}_k)$ of L). For any basis $(\mathbf{b}_1, \ldots, \mathbf{b}_k) \in \mathbb{R}^{n \times k}$ of a lattice, we define $\|\mathbf{B}\| = \max_i \|\mathbf{b}_i\|_2$ and let $(\mathbf{b}_1^\star, \ldots, \mathbf{b}_k^\star)$ denote its Gram-Schmidt vectors.

2.2 Algebraic number theory

We present here the number theoretic objects we will use throughout this work. For an in-depth introduction to the field, the reader is referred to [Coh96,Neu13]. Let K be a number field of degree $d \geq 2$ and discriminant Δ_K. Let $\mathcal{O}_K$ be its ring of integers.

Ideals. An ideal $\mathfrak{a} \subseteq \mathcal{O}_K$ is called an *integral* ideal. A *fractional ideal* of K is a discrete subset of K of the form $(x) = x\mathfrak{a}$, where $x \in K$ and $\mathfrak{a}$ is an integral ideal. Equivalently, a fractional ideal is a finitely generated $\mathcal{O}_K$-submodule of K. A fractional ideal of the form $x \cdot \mathcal{O}_K$ is called *principal*. In this work, we will take the convention that gothic letters (such as $\mathfrak{a}, \mathfrak{b}, \mathfrak{p}$) correspond to integral ideals, while upper-case letters (such as I, J) refer to ideals that are not necessarily integral.

For any fractional ideals I, J, we define the product $I \cdot J$ as the ideal generated by all products $a \cdot b$ for $a \in I, b \in J$ and the inverse I^{-1} as the ideal $I^{-1} = \{x \in K, xI \subseteq \mathcal{O}_K\}$. An integral ideal $\mathfrak{p}$ is said to be prime if there do not exist $\mathfrak{a}$ and $\mathfrak{b}$ integral and distinct from $\mathfrak{p}$ such that $\mathfrak{p} = \mathfrak{a}\mathfrak{b}$. These properties give the set of fractional ideals a group structure, the quotient group of fractional ideals of K by principal ideals is the class group of K, it is denoted Cl_K and is finite. We define the algebraic norm of an integral ideal $\mathfrak{a}$ by $\mathcal{N}(\mathfrak{a}) = [\mathcal{O}_K : \mathfrak{a}]$. We have $\mathcal{N}(\mathfrak{a}\mathfrak{b}) = \mathcal{N}(\mathfrak{a})\mathcal{N}(\mathfrak{b})$ for all integral ideals $\mathfrak{a}$ and $\mathfrak{b}$. If I is a fractional ideal, there exists an integer N such that $N \cdot I$ is integral, and we define $\mathcal{N}(I) = \mathcal{N}(N \cdot I)/N^d$ (this is independent of the choice of N). The multiplicativity property of the norm carries over to fractional ideals. For any set $\mathcal{W}$ of fractional ideals we define $\mathcal{W}^{-1} = \left\{I^{-1}, I \in \mathcal{W}\right\}$. For any $2 \leq A \leq B$, we define $\mathcal{I}_{A,B}$ the set of integral ideal with norm in $[A, B]$ and $\mathcal{P}_{A,B}$ the set of prime ideals with norm in $[A, B]$.

Embedding and ideal lattices. The canonical embedding $\sigma : K \to \mathbb{C}^d$ is defined as $x \mapsto (\sigma_1(x), \ldots, \sigma_d(x))$, where the σ_i's are the complex embeddings of K, ordered so the $d_\mathbb{R}$ ones with values in $\mathbb{R}$ come first, and $\sigma_i = \overline{\sigma_{d_\mathbb{C}+i}}$ for all $d_\mathbb{R} < i \leq d_\mathbb{R} + d_\mathbb{C}$ (note that $d = d_\mathbb{R} + 2d_\mathbb{C}$). We define $K_\mathbb{R} = K \otimes_\mathbb{Q} \mathbb{R}$, which is a ring containing K. The complex embeddings σ_i and the canonical embedding σ are extended to $K_\mathbb{R}$, and we have that $\sigma(K_\mathbb{R})$ is the set of $\mathbf{x} \in \mathbb{C}^d$ such that $x_i \in \mathbb{R}$ for $1 \leq i \leq d_\mathbb{R}$ and $x_{d_\mathbb{C}+i} = \overline{x_i}$ for $d_\mathbb{R} < i \leq d_\mathbb{R} + d_\mathbb{C}$. This is a real vector space of dimension d and a ring where addition and multiplication are performed coordinate-wise. The canonical embedding allows us to view any element x of K (and of $K_\mathbb{R}$) as a vector in $\mathbb{C}^d$, and to define $\|x\| = \|\sigma(x)\|$. The (absolute) algebraic norm of $x \in K_\mathbb{R}$ is defined as $\mathcal{N}(x) = \prod_{1 \leq i \leq d} |\sigma_i(x)|$. We have $\mathcal{N}(x \cdot \mathcal{O}_K) = \mathcal{N}(x)$.

We define $K_\mathbb{R}^0$ the set of norm-1 elements of $K_\mathbb{R}$. A replete ideal is a subset of $K_\mathbb{R}$ of the form $x \cdot \mathfrak{a}$, where $\mathfrak{a}$ is an integral ideal and $x \in K_\mathbb{R}^\times$ is invertible (we exclude divisors of zero). With this notation, if $\mathfrak{a}$ is principal, we call $x \cdot \mathfrak{a}$ a principal replete ideal. Multiplication and inversion are extended to the set of replete ideals by $(x \cdot \mathfrak{a}) \cdot (y \cdot \mathfrak{b}) = (xy) \cdot (\mathfrak{a}\mathfrak{b})$ and $(x \cdot \mathfrak{a})^{-1} = x^{-1} \cdot \mathfrak{a}^{-1}$. If we

remove the zero ideal, then this set is a (multiplicative) group. We also extend the algebraic norm to replete ideals by $\mathcal{N}(x \cdot \mathfrak{a}) = \mathcal{N}(x) \cdot \mathcal{N}(\mathfrak{a})$.

Every non-zero replete ideal $x\mathfrak{a}$ corresponds to a full-rank lattice $\sigma(x\mathfrak{a})$. By abuse of notation, we identify $x\mathfrak{a}$ and $\sigma(x\mathfrak{a})$. We have $\mathrm{vol}(x\mathfrak{a}) = \sqrt{\Delta_K} \cdot \mathcal{N}(x\mathfrak{a})$, and the covering radius in ℓ_∞ norm of $x\mathfrak{a}$ is bounded from above by:

$$\mu_\infty(x\mathfrak{a}) \leq d \cdot \lambda_d^{(\infty)}(x\mathfrak{a}) \leq d \cdot \lambda_1^{(\infty)}(x\mathfrak{a}) \cdot \lambda_d^{\infty}(\mathcal{O}_K) \leq d \cdot \Delta_K^{3/(2d)} \cdot \mathcal{N}(x\mathfrak{a})^{1/d}, \quad (1)$$

where we bounded $\lambda_1(x\mathfrak{a})$ by $\Delta_K^{1/(2d)} \cdot \mathcal{N}(x\mathfrak{a})^{1/d}$ using Minkowski's theorem and $\lambda_d^{\infty}(\mathcal{O}_K)$ by $\Delta_K^{1/d}$ using [BST+20, Th. 3.1] (adapted to the ℓ_∞ norm in [Boe22, Th. A.4]). For an (integral / fractional / replete) ideal, we call the corresponding image an (integral / fractional / replete) ideal lattices (with respect to K). We define idLat^0 as the set of replete ideal lattices of norm 1. This is a compact subgroup of idLat, and it admits a uniform distribution $\mathcal{U}(\mathrm{idLat}^0)$.

Lemma 2.1. *Let J be a replete ideal, then*

$$\Pr_{I \leftarrow \mathcal{U}(\mathrm{idLat}^0)} \left(\exists x \in K_{\mathbb{R}}^{\times} : \ J = I \cdot (x) \right) = \frac{1}{|\mathrm{Cl}_K|}$$

Proof. For any replete ideal $I = (x) \cdot \mathfrak{a}$ with $\mathfrak{a}$ integral, we define $[I] = [\mathfrak{a}] \in \mathrm{Cl}_K$. The value of $[I]$ does not depend on the choices of $\mathfrak{a}$ and x. The function $I \mapsto [I]$ for $I \in \mathrm{idLat}^0$ is a surjective morphism whose kernel is the set of principal replete ideals of norm 1 in $K_{\mathbb{R}}$. The lemma states that if I is sampled from $\mathcal{U}(\mathrm{idLat}^0)$, then the probability that it belongs to a fixed coset of Cl_K is $|\mathrm{Cl}_K|^{-1}$, which follows directly from the fact that $[\cdot]$ is a surjective morphism. □

We define the logarithmic embedding of $K_{\mathbb{R}}$, by taking the natural logarithm of every embedding of an element:

$$\begin{aligned} \mathrm{Ln} : K_{\mathbb{R}}^{\times} &\longrightarrow \mathrm{Ln}(K_{\mathbb{R}}) = \mathbb{R}^d \\ x &\longmapsto (\ln |\sigma_i(x)|)_i \end{aligned}$$

The following lemma is a standard result on the logarithmic embedding. The first statement is a rewriting of the equality $\mathrm{Ln}(x) = 0$ and the second one is a consequence of Dirichlet's unit theorem.

Lemma 2.2. *The function* Ln *has kernel $\left\{x \in K_{\mathbb{R}}^{\times} : \forall i, |x_i| = 1\right\}$, whose intersection with $\mathcal{O}_K$ is the set μ_K of roots of unity of K.*

By relying on random walks in the Arakelov class group of K, de Boer et al. [BDPW20] proposed an efficient algorithm to sample from $\mathcal{U}(\mathrm{idLat}^0)$. We give here a simplified version of this result borrowed from [FPS22], for a single-step random walk.

Lemma 2.3 ([FPS22, Le. 2.4]). *Let $A \geq 2$, and D the distribution over* idLat^0 *of*

$$I = u \cdot \mathrm{Exp}(\zeta) \cdot \mathfrak{p} \cdot \mathcal{N}(\mathfrak{p})^{-1/d},$$

for $\mathfrak{p}$ *uniform in* $\mathcal{P}_{0,A}$*,* u *uniform in* $\left\{x \in K_{\mathbb{R}}^{\times} : \forall i \leq d, |x_i| = 1\right\}$ *and* ζ *sampled according to* $\mathcal{G}(0, d^{-3/2})$ *in* $\mathrm{span}(\mathrm{Ln}(\mathcal{O}_K^{\times}))$ *conditioned on* $\|\zeta\| \leq 1/d$*. Then there exists an absolute constant* $c > 1$ *such that if* $A \geq (d^d \cdot \Delta_K)^c$*, then*

$$\mathrm{SD}(D,\mathcal{U}(\mathrm{idLat}^0)) = 2^{-\Omega(d)}.$$

Balanced elements. For the reductions presented in this article, it will sometimes be convenient to use balanced elements of $K_{\mathbb{R}}$, i.e., elements whose ℓ_∞ norm and the one of their inverse are not far from the geometric mean of their coordinates: in other terms they do not have an exceptionally small or large coordinate in comparison to the others. This property is convenient as it implies that multiplying an ideal by one of these elements will not change its geometry significantly, in particular if x is balanced and v is small in the ideal $x \cdot I$, then $x^{-1} \cdot v$ will be small in I. The formal definition is as follows.

Definition 2.4. *Let* $\eta > 1$*. An element* x *in* $K_{\mathbb{R}}$ *is said to be* η*-balanced if*

$$\|x\|_\infty \leq \eta \cdot |\mathcal{N}(x)|^{\frac{1}{d}} \quad \text{and} \quad \|x^{-1}\|_\infty \leq \eta \cdot |\mathcal{N}(x)|^{-\frac{1}{d}}.$$

Density of prime ideals. For any $A \geq 1$, we let $\tilde{\rho}_A$ denote the inverse of the proportion of prime ideals among all integral ideals of K of norm $\leq A$, i.e.,

$$\tilde{\rho}_A := \frac{|\{\mathfrak{a} \subset \mathcal{O}_K \,|\, \mathcal{N}(\mathfrak{a}) \leq A\}|}{|\{\mathfrak{p} \subset \mathcal{O}_K \text{ prime} \,|\, \mathcal{N}(\mathfrak{p}) \leq A\}|}.$$

In this article, we will be interested in $\tilde{\rho}_A$ for values of A of the order of $\mathrm{poly}(\Delta_K)$. Unfortunately, we are not aware of estimates for $\tilde{\rho}_A$ when A is this "small". However, it is known that when the number field K is fixed and A tends to infinity, it holds that

$$\tilde{\rho}_A \underset{A\to\infty}{\sim} \rho_K \cdot \ln(A),$$

where ρ_K is the residue of the Dedekind zeta function at 1. This comes from the fact that $|\{\mathfrak{p} \subset \mathcal{O}_K \text{ prime} \,|\, \mathcal{N}(\mathfrak{p}) \leq A\}| \sim A/\ln(A)$ (see [BS96, Th. 8.7.4]), and that $|\{\mathfrak{a} \subset \mathcal{O}_K \,|\, \mathcal{N}(\mathfrak{a}) \leq A\}| \sim \rho_K \cdot A$ (see [Web08]). The quantity ρ_K is known to be $\mathrm{poly}(\log \Delta_K)$ for some number fields such as cyclotomic fields (under ERH, see [Boe22, Th. A.5]), but there also exist families of fields in which ρ_K is exponential in the degree and $\Delta_K^{1/(2d)}$ (e.g., for some multi-quadratic number fields).

2.3 Algorithmic problems

Representing field elements and ideals. We assume that we know a $\mathbb{Z}$-basis $\mathbf{B}_{\mathcal{O}_K}$ of $\mathcal{O}_K$, and that it is LLL-reduced with respect to the geometry induced by σ (in some cases, a much better basis could be known). We define $\delta_K := \|\mathbf{B}_{\mathcal{O}_K}\|$. Since $\mathbf{B}_{\mathcal{O}_K}$ is LLL-reduced, we have that $\delta_K \leq 2^d \cdot \lambda_d(\mathcal{O}_K) = O(2^d \cdot \Delta_K^{1/d})$ from [BST+20, Th. 3.1], which implies that $\log \delta_K = O(\log \Delta_K)$.

Elements of K will be represented as vectors of $\mathbb{Q}^d$, corresponding to their coordinates in the basis $\mathbf{B}_K$. Fractional ideals of K will be represented by a $\mathbb{Z}$-basis, i.e., d elements of K generating the ideal (each element being represented as a vector of $\mathbb{Q}^d$ as described above). The bases we obtain for a fractional ideal I are in $\mathbb{Q}^{d\times d}$, so they admit a Hermite Normal Form (HNF), which provides a canonical representation for I. When replete ideals are used in algorithms, they will be represented by an arbitrary basis with size polynomial in the log of their norm and in $\log \Delta_K$ (with a polynomial number of bits of precision).

Algorithmic problems in ideals. In this article, we will consider the Hermite shortest vector problem in ideals, as well as related algorithmic problems.

Definition 2.5. *Let $\gamma \geq 1$. The ideal Hermite Shortest Vector Problem* id-HSVP$_\gamma$ *asks, given as input a fractional ideal I represented by its HNF basis, to find a non-zero element $x \in I$ such that $\|x\| \leq \gamma \cdot \mathrm{vol}(I)^{1/d}$. For a finite set X of fractional ideals, the average-case variant* X-avg-id-HSVP$_\gamma$ *asks to solve* id-HSVP$_\gamma$ *when the input ideal I is uniformly sampled in X. The success probability of an algorithm $\mathcal{A}$ when solving* X-avg-id-HSVP$_\gamma$ *is defined as*

$$\Pr_{I \hookleftarrow X}\left[x \in I \text{ and } \|x\| \leq \gamma \cdot \mathrm{vol}(I)^{1/d} \mid \mathcal{A}(I) = x\right],$$

where the randomness is taken over the choice of I and the possible internal randomness of $\mathcal{A}$.

The problem inv-HSVP$_\gamma$ *is* id-HSVP$_\gamma$ *restricted to inverses of integral ideal lattices.*

The problems inv-HSVP$_\gamma$ and id-HSVP$_\gamma$ are equivalent under Karp reductions, without any loss in the approximation factor, as shown in the following lemma (the other direction follows from the definition).

Lemma 2.6 (Folklore). *For any $\gamma \geq 1$, there is a Karp polynomial-time reduction from* id-HSVP$_\gamma$ *to* inv-HSVP$_\gamma$.

Proof. Let I be a fractional ideal for which we want to solve the id-HSVP$_\gamma$ problem. We will show that there exists $x \in \mathbb{Q}$ such that $xI = \mathfrak{a}^{-1}$ is the inverse of an integral ideal $\mathfrak{a} \subseteq \mathcal{O}_K$. If such an element x can be computed efficiently, then the reduction simply computes x, then compute $\mathfrak{a}^{-1} = xI$ and runs the inv-HSVP$_\gamma$ solver on $\mathfrak{a}^{-1}$ (which is a valid input for inv-HSVP). Since multiplication by $x \in \mathbb{Q}$ consists in scaling the lattice corresponding to I, then a solution to id-HSVP$_\gamma$ in xI provides a solution to id-HSVP$_\gamma$ in I (by multiplying it by x^{-1}). Note that the reduction preserves the approximation factor γ.

Let us then show that such an x exists and can be computed in polynomial time. Write $I = \mathfrak{a}\mathfrak{b}^{-1}$, with $\mathfrak{a}, \mathfrak{b} \subseteq \mathcal{O}_K$ integral ideals, and define $x = \mathcal{N}(\mathfrak{a})^{-1}$. Note that such $\mathfrak{a}$, $\mathfrak{b}$ and x can be computed in polynomial time from I (we do not require that $\mathfrak{a}$ and $\mathfrak{b}$ are coprime, so the choice we make is not unique). Let us show that for such x, it holds that $(xI)^{-1} \subseteq \mathcal{O}_K$ is an integral ideal. By definition, we have $(xI)^{-1} = \mathcal{N}(\mathfrak{a}) \cdot \mathfrak{a}^{-1}\mathfrak{b}$. Since $\mathfrak{a}$ is integral, it holds that $\mathcal{N}(\mathfrak{a}) \cdot \mathcal{O}_K \subseteq \mathfrak{a}$.

Indeed, note that the group $\mathcal{O}_K/\mathfrak{a}$ has cardinality $\mathcal{N}(\mathfrak{a})$. Lagrange's theorem then gives that any element of $\mathcal{O}_K/\mathfrak{a}$ has order dividing $\mathcal{N}(\mathfrak{a})$, i.e., for any $x \in \mathcal{O}_K$, we have $\mathcal{N}(\mathfrak{a}) \cdot x \in \mathfrak{a}$. We hence obtain that the ideal $\mathcal{N}(\mathfrak{a}) \cdot \mathfrak{a}^{-1} \subseteq \mathcal{O}_K$ is integral. Since $\mathfrak{b}$ is integral by construction, this proves that $(xI)^{-1}$ is integral. □

2.4 Algorithms on ideals

We will often manipulate ideals and their basis. We will use the following results on how to derive a short basis from a full-rank set of vectors.

Lemma 2.7 (Corollary of [MG02, Le. 7.1]). *There exists a polynomial time algorithm that takes as input a basis $\mathbf{B}$ of an n-dimensional lattice L and a set of n linearly independent vectors $\mathbf{s}_1, \cdots, \mathbf{s}_n \in L$ and outputs a new basis $\mathbf{C}$ of L such that $\|\mathbf{C}^\star\| \leq \max_i \|\mathbf{s}_i^*\|$ and $\|\mathbf{C}\| \leq \sqrt{n} \cdot \max_i \|\mathbf{s}_i\|$.*

We will use Lemma 2.7 to perform arithmetic over ideals while bounding the sizes of the outputs.

Lemma 2.8. *There exist polynomial-time algorithms* `InvertIdeal`, `ReduceIdeal` *and* `MultiplyIdeals` *with the following specifications.*

- `InvertIdeal` *takes as input an integral ideal $\mathfrak{a}$ and outputs a basis $\mathbf{B}$ of $\mathfrak{a}^{-1}$ such that $\|\mathbf{B}^\star\| \leq \delta_K$ and $\|\mathbf{B}\| \leq \sqrt{d} \cdot \delta_K$.*
- `ReduceIdeal` *takes as input a basis $\mathbf{B}$ of an ideal $I \subset K_\mathbb{R}$ and a vector $v \in I \setminus \{0\}$ and returns a basis $\mathbf{B}_I$ of I such that $\|\mathbf{B}_I^\star\| \leq \delta_K \cdot \|v\|$ and $\|\mathbf{B}_I\| \leq \sqrt{d} \cdot \delta_K \cdot \|v\|$.*
- `MultiplyIdeals` *takes as input bases $\mathbf{B}_I$ and $\mathbf{B}_J$ of two ideals $I, J \subseteq K_\mathbb{R}$ and output $\mathbf{B}_{IJ}$ a basis of $I \cdot J$ such that $\|\mathbf{B}_{IJ}^\star\| \leq \|\mathbf{B}_I\| \cdot \|\mathbf{B}_J\|$ and $\|\mathbf{B}_{IJ}\| \leq \sqrt{d} \cdot \|\mathbf{B}_I\| \cdot \|\mathbf{B}_J\|$.*

Proof. `InvertIdeal` starts by computing a basis $\mathbf{B}$ of $\mathfrak{a}^{-1}$, which can be done in polynomial time from a representation of $\mathfrak{a}$ by generators. Then, the algorithm runs the algorithm from Lemma 2.7 with input the basis $\mathbf{B}$ of $\mathfrak{a}^{-1}$ and the vectors of the known basis $\mathbf{B}_{\mathcal{O}_K}$ of $\mathcal{O}_K$ (in the role of the short vectors $\mathbf{s}_i$). Note that since $\mathfrak{a}$ is integral, we have that $\mathcal{O}_K \subseteq \mathfrak{a}^{-1}$, and hence the vectors of $\mathbf{B}_{\mathcal{O}_K}$ are indeed in $\mathfrak{a}^{-1}$. Also, the euclidean norm of those vectors is bounded from above by δ_K, by definition. We conclude by using Lemma 2.7.

For `ReduceIdeal`, note that the set $v \cdot \mathbf{B}_{\mathcal{O}_K}$ is a free subset of I whose vectors have norms $\leq \|v\| \cdot \delta_K$. We can then define `ReduceIdeal` as the application of Lemma 2.7 with input $\mathbf{B}, v \cdot \mathbf{B}_{\mathcal{O}_K}$.

Let $\mathbf{B}_I = (b_i^{(I)})_i, \mathbf{B}_J = (b_i^{(J)})_i$ be the inputs to `MultiplyIdeals`. Then the set $(b_i^{(I)} \cdot b_j^{(J)})_{i,j}$ generates IJ and has size d^2, this implies that there exists a $\mathbb{Q}$-free family $(r_i)_{i=1,\ldots d}$ inside it, which can be found in polynomial time and verifies $\max_i \|r_i\| \leq \|\mathbf{B}_I\| \cdot \|\mathbf{B}_J\|$. Further, a basis $\mathbf{B}$ of IJ can be found in polynomial time. We then apply Lemma 2.7 with input $\mathbf{B}, (r_i)_i$. □

For $I = \mathcal{O}_K$, the following lemma states that one can quantumly and efficiently sample a random prime ideal together with a short element in it, hence the name. We give a proof based on [PS21] but note that a similar statement was already given as [Gen09a, Th. 16.3.4, Le. 17.2.1] (see also [Gen10, Se. 3.3]).

Lemma 2.9 (Adapted from [PS21, Lemma C.1]). *There exists an algorithm* `SampleWithTrap` *that on input integers $2 \leq A < B$, a real $\varepsilon \in (0,1)$ and a basis $\mathbf{B}_I$ of a fractional ideal I, samples a pair $(\mathfrak{p}, w)$ such that*

1. *the distribution of $\mathfrak{p}$ is within statistical distance ε from the uniform distribution over $\mathcal{P}_{A,B}$;*
2. *the element w belongs to $I \cdot \mathfrak{p} \setminus \{0\}$;*
3. *we have $\|w\| \leq 2\sqrt{4d + \ln(24B/\varepsilon)} \cdot s$ with $s = \max(s_{sample}, s_{smooth})$ and*
   - $s_{sample} = \sqrt{d} \cdot \|\mathbf{B}_I^\star\|$.
   - $s_{smooth} = (\Delta_K \cdot B \cdot \mathcal{N}(I))^{1/d} \cdot \sqrt{\ln(24B/\varepsilon)}$.

Furthermore, if the algorithm is given access to an oracle factoring integral ideals of norm smaller than $(2\sqrt{4d + \ln(24B/\varepsilon)} \cdot s)^d \cdot \mathcal{N}(I)^{-1}$, then the algorithm runs in expected time polynomial in $B/|\mathcal{P}_{A,B}|, B/A, \log \Delta_K, \log B, \log(1/\varepsilon)$ and in the size of I.

The proof is available in Appendix A. Note that we will use this result with $\varepsilon = \exp(-d)$ in order to simplify computations and subsequently omit this input.

Factoring ideals. Factoring an integral ideal $\mathfrak{a}$ in $\mathcal{O}_K$ can be done by factoring the algebraic norm $\mathcal{N}(\mathfrak{a})$ of $\mathfrak{a}$ over the integers; computing, for all the prime factors $p \,|\, \mathcal{N}(\mathfrak{a})$, the set of prime ideals whose norm is a power of p (there are at most d of those); and testing for each of these prime if they divide $\mathfrak{a}$. Factoring $\mathcal{N}(\mathfrak{a})$ can be performed quantumly in time polynomial in $\log \mathcal{N}(\mathfrak{a})$ (using Shor's algorithm [Sho94]). Computing the set of prime ideals of norm a given prime integer p can be performed classically in time polynomial in $\log p$ and $\log \Delta_K$ using Buchmann-Lenstra's algorithm [BL94], described in details in [Coh96, Sec. 6.2.5]. Finally, testing whether a prime ideal $\mathfrak{p}$ divides $\mathfrak{a}$ can be done in time polynomial in the bit-sizes of $\mathfrak{p}$ and $\mathfrak{a}$. Overall, factoring ideals can be done in quantum-polynomial time (using Shor's algorithm) or in classical sub-exponential time (using the Number Field Sieve).

2.5 Worst-case to average-case reduction for inverse of primes

In [Gen09a, Ch. 16 & 17], Gentry described a self-reduction for a variant of the bounded distance decoding problem, from worst-case ideals to prime ideals taken uniformly at random with their norm in some interval $[A, B]$ (for a suitable choice of A and B). This reduction can be adapted to the shortest vector problem (instead of the bounded distance decoding problem), but it requires to take the inverse of the ideals, implying that the average-case distribution we obtain is over the inverses of prime ideals uniformly chosen in the interval $[A, B]$. Below, we state the result of Gentry's reduction adapted to SVP, and provide a proof in Appendix C for the sake of completeness.

Theorem 2.10 (Adapted from [Gen09a, Ch. 16 & 17]). *There exist some* $C_{1,K} = \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ *and* $C_{2,K} = \mathrm{poly}(\log \Delta_K, \delta_K)$ *such that the following holds. Let* $\gamma_{avg} \in [1, 2^d]$*,* $A \geq C_{1,K}^d \cdot \gamma_{avg}^d$ *satisfying* $A \leq (\Delta_K)^{d^{O(1)}}$ *and* $\gamma = A^{1/d} \cdot C_{2,K}$*. Then*

$$\text{id-HSVP}_\gamma \quad \textit{reduces to} \quad \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{avg}}.$$

The reduction is probabilistic and, assuming it has access to an oracle factoring integral ideals whose norms have bit-size $\mathrm{poly}(\log \Delta_K)$*, it runs in expected time polynomial in its input size,* $\log \Delta_K$ *and* $1/\delta$*, where* δ *is the success probability of the* $\mathcal{P}_{A,4A}^{-1}$-avg-id-HSVP$_{\gamma_{avg}}$ *oracle.*[8]

3 Self-Reducibility of id-HSVP to Inverses

Let $\mathcal{W}$ be a set of fractional ideals. In this section, we provide a framework for reducing id-HSVP for the uniform distribution over $\mathcal{W}$ to id-HSVP for the uniform distribution over $\mathcal{W}^{-1} = \{I^{-1} : I \in \mathcal{W}\}$. The reduction, provided in Theorem 3.4 relies on three oracles (beyond the one for id-HSVP for $\mathcal{U}(\mathcal{W})$). The first one factors integral ideals, and can be instantiated with a quantum polynomial-time algorithm. The second one samples from $I \cap X$, where I is an arbitrary norm-1 replete ideal and X is a well-chosen set: this oracle will be instantiated in Section 4. The last one finds short non-zero vectors in integral ideals uniformly distributed within those having their norms in a prescribed interval. Overall, this will lead to a quantum polynomial-time reduction from $\mathcal{W}^{-1}$-avg-id-HSVP to $\mathcal{W}$-avg-id-HSVP and $\mathcal{I}_{A,4A}$-avg-id-HSVP for a well-chosen A.

The reduction is built in several steps. First, we show how to map a uniform norm-1 replete ideal to an integral ideal uniform among those with norms in $[A, 4A]$, using a new approach introduced in [Boe22, Sec. 6]. This is parametrized by a set X that will be instantiated in Section 4. The second step gives a way to randomize an arbitrary ideal to an integral ideal uniform among those with norms in $[A, 4A]$, along with a hint that allows to map a short vector of the resulting ideal to a short vector in the inverse of the input ideal. Finally, this allows to describe the reduction.

3.1 From a uniform norm-1 ideal to a uniform integral ideal

In this subsection, we present a way to sample uniformly among integral ideals whose norms belong to a prescribed interval. Given a compact set X verifying certain properties, our sampler takes as input a uniform ideal $I \in \mathrm{idLat}^0$, samples a point uniformly in $I \cap X$ and output $(x) \cdot I^{-1} \subseteq \mathcal{O}_K$. It holds that if X is well-designed, then the output distribution is close to the uniform distribution

[8] The choice of $4A$ for the upper bound on the norm of the ideals is not a strict requirement of this theorem. We instantiated the theorem with this value in order to simplify its statement.

over the set of integral ideals in terms of Rényi divergence. Our sampler generalizes [Boe22, Th. 6.9], where the set X is assumed to be the ℓ_∞ ball. This new degree of freedom will allow us (in Section 4) to choose a set X whose points are balanced, which will be essential for the proof of Theorem 5.1. Note that we do not use the Arakelov ray divisor formalism to state our results: those of [Boe22, Sec. 6] are stated with respect to a modulus $\mathfrak{m} \subseteq \mathcal{O}_K$ and here we take $\mathfrak{m} = \mathcal{O}_K$.

Definition 3.1. *Let $X \subset K_\mathbb{R}$. We say that X is compact and invariant by complex rotations if the following hold:*

- *$\sigma(X)$ is a compact subset of $\mathbb{C}^d$;*
- *for any $\zeta = (\zeta_1, \ldots, \zeta_d) \in \sigma(K_\mathbb{R})$ with $|\zeta_1| = \cdots = |\zeta_d| = 1$, it holds that $\sigma^{-1}(\zeta) \cdot X = \{\sigma^{-1}(\zeta) \cdot x \,|\, x \in X\} \subseteq X$.*

We consider the `IdealRound` algorithm (Algorithm 3.1), whose output distribution generalizes the distribution presented in [Boe22, Th. 6.9]. It is parametrized by an arbitrary compact set $X \subset K_\mathbb{R}$, takes as input a norm-1 replete ideal (i.e., an element of idLat^0) and returns an integral ideal. We define $\mathcal{D}_{\mathrm{Ideal}}(X)$ as the distribution $\texttt{IdealRound}_X(\mathcal{U}(\mathrm{idLat}^0))$. For the moment, we are not interested in the efficiency of $\texttt{IdealRound}_X$, but only in the relationship between $\mathcal{D}_{\mathrm{Ideal}}(X)$ and the uniform distribution over ideals with norms belonging to an interval. This is the purpose of the following result.

Algorithm 3.1 `IdealRound`

Input: $I \in \mathrm{idLat}^0$.
Parameter: $X \subset K_\mathbb{R}$ compact.
Output: An integral ideal $\mathfrak{a}$.
1: Sample $x \leftarrow \mathcal{U}(I \bigcap X)$.
2: Return $\mathfrak{a} = (x) \cdot I^{-1}$.

Lemma 3.2. *For any $t \in \mathbb{R}$, let $H_t = \{x \in \mathrm{Ln}\, K_\mathbb{R} \,|\, \sum_i x_i = t\}$. Let X be a compact subset of $K_\mathbb{R}$ invariant by complex rotations (as per Definition 3.1) and $B > A > 2$. Assume that:*

- *There exist some real numbers $C \geq 1$ and $C' > 0$ such that we have $|I \cap X| \in C' \cdot [1, C]$ for any $I \in \mathrm{idLat}^0$;*
- *there exists $C'' \in \mathbb{R}$ such that for any $t \in [\ln(A), \ln(B)]$ we have*

$$\mathrm{vol}\left(\mathrm{Ln}(X) \bigcap H_t\right) = C'';$$

- *for any $t \notin [\ln(A), \ln(B)]$, we have $\mathrm{vol}(\mathrm{Ln}(X) \cap H_t) = 0$.*

Then the support of $\mathcal{D}_{\mathrm{Ideal}}(X)$ is contained in $\mathcal{I}_{A,B}$ and

$$\mathrm{RD}_\infty(\mathcal{U}(\mathcal{I}_{A,B}) \,\|\, \mathcal{D}_{\mathrm{Ideal}}(X)) \leq C.$$

We now comment the conditions of Lemma 3.2. The second and third conditions state that, when embedded in $\mathrm{Ln}(K_{\mathbb{R}})$ the set $\mathrm{Ln}(X)$ should be contained between the two hyperplanes $H_{\log(A)}$ and $H_{\log(B)}$, and that between those hyperplanes, the slices $\mathrm{Ln}(X) \cap H_t$ should have constant volume. Those conditions will yield the bounds on the norm of the output ideal. The first condition states that for any norm-1 replete ideal I, the number of points in $X \cap I$ should be non-zero and almost independent of I. Conditions 1 and 2 will imply the near-uniformity of the output distribution. The proof below is adapted from [Boe22, Th. 6.9].

Proof. Fix an integral ideal $\mathfrak{b}$ and a norm-1 replete ideal I. We are going to compute bounds on

$$p_{I,\mathfrak{b}} = \Pr_x\left((x) \cdot I^{-1} = \mathfrak{b}\right) = \Pr_x\left((x) = I \cdot \mathfrak{b}\right) = \Pr_x\left(x \text{ generates } I \cdot \mathfrak{b}\right),$$

where the randomness is over $x \leftarrow \mathcal{U}(I \cap X)$. For an ideal J, we define $G_J = \{x \in K_{\mathbb{R}} : (x) = J\}$ as the set of generators of J (if J is not principal, it is the empty set). Note that $G_{I\cdot\mathfrak{b}} = \{x \in K_{\mathbb{R}} : (x) = I \cdot \mathfrak{b}\} \subseteq I$. We have

$$p_{I,\mathfrak{b}} = \frac{|G_{I\cdot\mathfrak{b}} \cap X|}{|I \cap X|} \in \left|G_{I\cdot\mathfrak{b}} \bigcap X\right| \cdot C'^{-1} \cdot [C^{-1}, 1],$$

where the inclusion follows from the first assumption of the lemma. For any I that is not in the class of $\mathfrak{b}^{-1}$ modulo principal ideals, we have that $G_{I\cdot\mathfrak{b}}$ is empty, since $I \cdot \mathfrak{b}$ is not principal. Let $[\mathfrak{b}^{-1}]^0$ be the set of all norm-1 replete ideals of the form $(\alpha) \cdot \mathfrak{b}^{-1}$ for some $\alpha \in K_{\mathbb{R}}$ (i.e., the coset of $\mathfrak{b}^{-1}$ in idLat^0 modulo principal ideals). Let $I_0 = \mathcal{N}(\mathfrak{b})^{1/d} \cdot \mathfrak{b}^{-1}$, which belongs to $[\mathfrak{b}^{-1}]^0$. There is a bijection between $K_{\mathbb{R}}^0/\mathcal{O}_K^\times$ and $[\mathfrak{b}^{-1}]^0$ given by $u \mapsto (u) \cdot I_0$. This implies that

$$\begin{aligned}\mathop{\mathbb{E}}_{I\leftarrow\mathcal{U}(\mathrm{idLat}^0)}\left(\left|G_{I\cdot\mathfrak{b}} \bigcap X\right|\right) &= \Pr_{I\leftarrow\mathcal{U}(\mathrm{idLat}^0)}(I \in [\mathfrak{b}^{-1}]^0) \cdot \mathop{\mathbb{E}}_u\left(\left|G_{\mathcal{N}(\mathfrak{b})^{1/d}\cdot(u)} \bigcap X\right|\right) \\ &= \frac{1}{|\mathrm{Cl}_K|} \cdot \mathop{\mathbb{E}}_u\left(\left|G_{\mathcal{N}(\mathfrak{b})^{1/d}\cdot(u)} \bigcap X\right|\right),\end{aligned}$$

where $u \leftarrow \mathcal{U}(K_{\mathbb{R}}^0/\mathcal{O}_K^\times)$ and the second equality comes from Lemma 2.1. Let μ_K be the set of roots of unity in K. Using the fact that the Ln function is $|\mu_K|$-to-1 when its input is restricted to generators of a principal replete ideal I, and that X is invariant by complex rotations, we have:

$$\forall I \in \mathrm{idLat}^0 : \ \left|G_I \bigcap X\right| = |\mu_K| \cdot \left|\mathrm{Ln}(G_I) \bigcap \mathrm{Ln}(X)\right|.$$

In our context, this implies that for any $u \in K_{\mathbb{R}}^0$,

$$\begin{aligned}\left|G_{\mathcal{N}(\mathfrak{b})^{1/d}\cdot(u)} \bigcap X\right| &= |\mu_K| \cdot \left|\mathrm{Ln}(X) \bigcap \left\{\mathrm{Ln}(x) : x = v \cdot u \cdot \mathcal{N}(\mathfrak{b})^{1/d}, v \in \mathcal{O}_K^\times\right\}\right| \\ &= |\mu_K| \cdot \left|\mathrm{Ln}(X) \bigcap (\Lambda_K + \mathrm{Ln}(u) + \mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d}))\right| \\ &= |\mu_K| \cdot \left|(\mathrm{Ln}(X) - \mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \bigcap (\Lambda_K + \mathrm{Ln}(u))\right|,\end{aligned}$$

where $\Lambda_K = \mathrm{Ln}\,\mathcal{O}_K^\times$. Note that Λ_K is full rank in H_0, and that $\mathrm{Ln}(u) \in H_0$ for any $u \in K_{\mathbb{R}}^0$. Moreover, the vector $\mathrm{Ln}(u)$ is uniform in H_0/Λ_K when u is uniform in $K_{\mathbb{R}}^0/\mathcal{O}_K^\times$. We are hence considering a uniform lattice shift and, for any measurable set $\mathcal{S} \subseteq H_0$, we have:

$$\mathop{\mathbb{E}}_u\Big(\big|(\Lambda_K + \mathrm{Ln}(u)) \bigcap \mathcal{S}\big|\Big) = \frac{\mathrm{Vol}(\mathcal{S})}{\mathrm{Vol}(\Lambda_K)}.$$

Applying this to the set $\mathcal{S} = (\mathrm{Ln}(X) - \mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \cap H_0$, we obtain

$$\mathop{\mathbb{E}}_u\left(\left|G_{\mathcal{N}(\mathfrak{b})^{1/d}\cdot(u)} \bigcap X\right|\right) = |\mu_K| \cdot \frac{\mathrm{Vol}((\mathrm{Ln}(X) - \mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \bigcap H_0)}{\mathrm{Vol}(\Lambda_K)}.$$

Observe that by definition of H_t for $t \in \mathbb{R}$, it holds that

$$(\mathrm{Ln}(X) - \mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})) \bigcap H_0 = \left(\mathrm{Ln}(X) \bigcap H_{\ln \mathcal{N}(\mathfrak{b})}\right) - \mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d}).$$

Since shifting by $\mathrm{Ln}(\mathcal{N}(\mathfrak{b})^{1/d})$ does not change the volume, we obtain

$$\mathop{\mathbb{E}}_u\left(\left|G_{\mathcal{N}(\mathfrak{b})^{1/d}\cdot(u)} \bigcap X\right|\right) = |\mu_K| \cdot \frac{\mathrm{Vol}(\mathrm{Ln}(X) \bigcap H_{\ln \mathcal{N}(\mathfrak{b})})}{\mathrm{Vol}(\Lambda_K)}.$$

Recall from the second and third assumptions that

$$\mathrm{Vol}\left(\mathrm{Ln}(X) \bigcap H_{\ln \mathcal{N}(\mathfrak{b})}\right) = \begin{cases} C'' & \text{if } \ln \mathcal{N}(\mathfrak{b}) \in [\ln A, \ln B], \\ 0 & \text{otherwise.} \end{cases}$$

Let $p = C'' \cdot |\mu_K|/(C' \cdot |\mathrm{Cl}_K| \cdot \mathrm{Vol}(\Lambda_K))$. Combining everything, this proves that

$$p_{\mathfrak{b}} := \mathop{\mathbb{E}}_{I \leftarrow \mathcal{U}(\mathrm{idLat}^0)}(\mathfrak{p}_{\mathfrak{b},I}) \in \begin{cases} p \cdot [C^{-1}, 1] & \text{if } \mathcal{N}(\mathfrak{b}) \in [A, B], \\ \{0\} & \text{otherwise.} \end{cases}$$

Observe that $p_{\mathfrak{b}}$ is equal to $\mathcal{D}_{\mathrm{Ideal}}(X)(\mathfrak{b})$, the probability of the ideal $\mathfrak{b}$ for the distribution $\mathcal{D}_{\mathrm{Ideal}}(X)$. The equation above then means that $\mathcal{D}_{\mathrm{Ideal}}(X)$ outputs ideals with norm in $[A, B]$ with probability essentially equal to p (up to a factor C), and other ideals with probability 0. We quantify this using the Rényi divergence. As $1 = \sum_{\mathfrak{b} \in \mathcal{I}_{A,B}} p_{\mathfrak{b}} \in p \cdot |\mathcal{I}_{A,B}| \cdot [C^{-1}, 1]$, we have that $p \in |\mathcal{I}_{A,B}|^{-1} \cdot [1, C]$, and hence:

$$\forall \mathfrak{b} \in \mathcal{I}_{A,B} : \ \frac{p_{\mathfrak{b}}}{\mathcal{U}(\mathcal{I}_{A,B})(\mathfrak{b})} \in [C^{-1}, C],$$

hence $\mathrm{RD}_\infty(\mathcal{U}(\mathcal{I}_{A,B}) \,\|\, \mathcal{D}_{\mathrm{Ideal}}(X)) \leq C$, which complete the proof. □

3.2 From an arbitrary ideal to a uniform integral ideal

Below, we give an algorithm, $\texttt{RandomizeIdeal}_{A,X}$ (see Algorithm 3.2), which on input an arbitrary ideal I, returns a uniform integral ideal $\mathfrak{b}$ and a short non-zero

vector $y \in \mathfrak{b}^{-1} \cdot I^{-1}$. The algorithm is parametrized by an integer A and a set X satisfying the conditions of Lemma 3.2. $\texttt{RandomizeIdeal}_{A,X}$ starts by sampling a uniform norm-1 ideal J, i.e., with distribution equal to $\mathcal{U}(\text{idLat}^0)$, along with a small element v_J in it, using the `SampleWithTrap` algorithm. Since $\mathcal{U}(\text{idLat}^0)$ is the Haar distribution on a compact group, the ideal $I' = J \cdot (I/\mathcal{N}(I)^{1/d})$ is also uniform. We then use `IdealRound` to map $\mathcal{U}(\text{idLat}^0)$ to the uniform distribution over integral ideals with norms in $[A, 4A]$. In more details, a uniform point x in $I' \cap X$ is sampled and Lemma 3.2 implies that $\mathfrak{b} := x \cdot I'^{-1}$ is almost uniform, and $v_J \cdot x^{-1}$ is a small element in $\mathfrak{b}^{-1} \cdot \mathcal{N}(I)^{1/d} \cdot I^{-1}$ if x is balanced. We note that Steps 7 and 8 below are exactly the `IdealRound` algorithm applied to the ideal I'. However, we cannot call this algorithm in a blackbox way, as we need to know the intermediate value x for Step 9 of the algorithm.

Algorithm 3.2 `RandomizeIdeal`

Input: A basis $\mathbf{B}_I$ of an ideal I.
Parameters: A integer and $X \subset K_{\mathbb{R}} \setminus \{0\}$ compact.
Oracles: $\mathcal{F}$ for factoring integral ideals, $\mathcal{S}$ for sampling from $\mathcal{U}(I \cap X)$ for $I \in \text{idLat}^0$.
Output: $\mathfrak{b}$ an integral ideal, $y \in \mathfrak{b}^{-1} \cdot I^{-1} \setminus \{0\}$.

1: Sample $(\mathfrak{q}, v_\mathfrak{q}) \leftarrow \texttt{SampleWithTrap}_{A,4A}(\mathbf{B}_{\mathcal{O}_K})$, using $\mathcal{F}$.
2: Sample $\zeta \leftarrow \mathcal{G}(0, d^{-3/2})$ in $\text{span}(\text{Ln}(\mathcal{O}_K^\times))$ conditioned on $\|\zeta\| \le 1/d$.
3: Sample u uniform in $\left\{x \in K_{\mathbb{R}}^\times : \forall i \le d, |x_i| = 1\right\}$.
4: Let $J = u \cdot \text{Exp}(\zeta) \cdot \mathcal{N}(\mathfrak{q})^{-1/d} \cdot \mathfrak{q}$ and $v_J = u \cdot \exp(\zeta) \cdot \mathcal{N}(\mathfrak{q})^{-1/d} \cdot v_\mathfrak{q}$.
5: Compute $\mathbf{B}_J = \texttt{ReduceIdeal}(J, v_J)$.
6: Let $I' = J \cdot I \cdot \mathcal{N}(I)^{-1/d}$ and $\mathbf{B}_{I'} = \texttt{MultiplyIdeals}(\mathbf{B}_J,\ \mathcal{N}(I)^{-1/d} \cdot \mathbf{B}_I)$.
7: Sample $x \leftarrow \mathcal{U}(I' \cap X)$, using $\mathcal{S}$.
8: Let $\mathfrak{b} = x \cdot I'^{-1}$.
9: Let $y = x^{-1} \cdot \mathcal{N}(I)^{-1/d} \cdot v_J$.
10: Return $(\mathfrak{b}, y)$.

Lemma 3.3. *Let $A \ge \max(\delta_K^d, d^d \Delta_K^c)$ for c as in Lemma 2.3. Let X be a compact subset of $K_{\mathbb{R}} \setminus \{0\}$ whose elements are η-balanced for some $\eta > 1$ and satisfy the assumptions of Lemma 3.2 for A and $B = 4A$. Assume that $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}| \le c'$ for some $c' < 1$. On input a basis $\mathbf{B}_I$ of an ideal I, $\texttt{RandomizeIdeal}_{A,X}$ runs in time polynomial in $\log A$, $\log \Delta_K$, $A/|\mathcal{P}_{A,4A}|$ and the size of its input, and returns $(\mathfrak{b}, y)$ satisfying*

$$\begin{aligned} &\mathfrak{b} \in \mathcal{I}_{A,4A}, \\ &y \in \mathfrak{b}^{-1} I^{-1} \setminus \{0\}, \\ &\|y\| \le 85 \cdot d \cdot \eta \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I\mathfrak{b})^{-1/d}. \end{aligned}$$

Finally, if D and $\mathcal{U}$ respectively denote the distribution of $\mathfrak{b}$ and the uniform distribution over $\mathcal{I}_{A,4A}$, then the following holds for any event $E \subseteq \mathcal{I}_{A,4A}$:

$$D(E) \ge \frac{\mathcal{U}(E)}{\Theta(1)} - 2^{-\Omega(d)}.$$

$$\begin{array}{ccccccccc} & & D_1 & \xleftrightarrow{\mathrm{SD}=2^{-\Omega(d)}} & D_2 & \xrightarrow{\mathrm{RD}_\infty=O(1)} & D_3 & \xleftrightarrow{\mathrm{SD}=2^{-\Omega(d)}} & D_4 \\ \texttt{IdealRound}(\cdot)\downarrow & & \downarrow & & \downarrow & & \downarrow & & \\ D = & \widetilde{D_1} & \xleftrightarrow{\mathrm{SD}=2^{-\Omega(d)}} & \widetilde{D_2} & \xrightarrow{\mathrm{RD}_\infty=O(1)} & \widetilde{D_3} & \xleftrightarrow{\mathrm{SD}=2^{-\Omega(d)}} & \widetilde{D_4} & \xrightarrow{\mathrm{RD}_\infty=O(1)} \mathcal{U}(\mathcal{I}_{A,4A}) \end{array}$$

Fig. 2. Relations between the distributions of the proof of Lemma 3.3.

Proof. We first bound the Euclidean norms of the variables occurring during the execution of the algorithm. By Lemma 2.9 and the assumption that $A \geq \delta_K^d$, we have that $0 < \|v_{\mathfrak{q}}\| \leq 51 \cdot d \cdot (A\Delta_K)^{1/d}$. Now, note that $\|u\|_\infty = 1$, $\|\exp(\zeta)\|_\infty \leq \exp(1/2)$ and $\mathcal{N}(\mathfrak{q})^{-1/d} \leq A^{-1/d}$. We then have $\|v_J\| \leq 85 \cdot d \cdot \Delta_K^{1/d}$ (and $v_J \neq 0$). Then, by Lemma 2.8, we have $0 < \|\mathbf{B}_J\| \leq 85 \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/d}$ and

$$\|\mathbf{B}_{I'}\| \leq 85 \cdot d^2 \cdot \delta_K \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I)^{-1/d} \cdot \|\mathbf{B}_I\|.$$

As elements of X are non-zero and η-balanced, we have that $\|x^{-1}\|_\infty \leq \eta \cdot \mathcal{N}(x)^{-1/d}$. Also, note that since $\mathcal{N}(I') = 1$, we have $\mathcal{N}(\mathfrak{b}) = \mathcal{N}(x)$. As a result, we obtain that $y \neq 0$ and:

$$\begin{aligned} \|y\| &\leq \mathcal{N}(I)^{-1/d} \cdot \|x^{-1}\|_\infty \cdot \|v_J\| \\ &\leq \mathcal{N}(I)^{-1/d} \cdot \eta \cdot \mathcal{N}(x)^{-1/d} \cdot 85 \cdot d \cdot \Delta_K^{1/d} \\ &= 85 \cdot d \cdot \eta \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I\mathfrak{b})^{-1/d}. \end{aligned}$$

The latter and the fact that $\mathcal{N}(\mathfrak{b}) = \mathcal{N}(x)$ belongs to $[A, 4A]$ (by assumption on X) provide the first statement on the output.

The previous computations show that every quantity manipulated by the algorithm has size polynomial in $\log A$, $\log \Delta_K$ and the bit-size of the input. Note that $\texttt{SampleWithTrap}_{A,4A}$ runs in polynomial time in $A/|\mathcal{P}_{A,4A}|$. The overall running time is then polynomial in $\log A$, $\log \Delta_K$, $A/|\mathcal{P}_{A,4A}|$ and the size of the input.

We now analyze the distribution of $\mathfrak{b}$. For this purpose, we define the following distributions (see also Figure 2):

- D_1 is the distribution of J at Step 4;
- D_2 is the distribution $u \cdot \exp(\zeta) \cdot \mathfrak{q} \cdot \mathcal{N}(\mathfrak{q})^{-1/d}$ where $\mathfrak{q}$ is uniform in $\mathcal{P}_{A,4A}$, and u, ζ are sampled as in Steps 2 and 3;
- D_3 is the same as D_2 but with $\mathfrak{q}$ uniform in $\mathcal{P}_{0,4A}$;
- D_4 is $\mathcal{U}(\mathrm{idLat}^0)$.

Note that we have the following relationships between the D_i's:

- $\mathrm{SD}(D_1, D_2) = 2^{-\Omega(d)}$, thanks to Lemma 2.9 and the data processing inequality;

- $\mathrm{RD}_\infty(D_3 \,\|\, D_2) = \Theta(1)$, thanks to the assumption on $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}|$;
- $\mathrm{SD}(D_3, D_4) = 2^{-\Omega(d)}$ thanks to Lemma 2.3.

We also define $\widetilde{D_i}$ (for $i \leq 4$) as the distribution of $\mathfrak{b}$ obtained by sampling J from D_i, setting $I' = J \cdot I \cdot \mathcal{N}(I)^{-1/d}$, sampling x from $\mathcal{U}(I' \cap X)$ and returning $x \cdot I'^{-1}$. Note that $\widetilde{D_1}$ is D and that $\widetilde{D_4}$ is $\mathcal{D}_{\mathrm{Ideal}}(X)$. Indeed, as $\mathcal{U}(\mathrm{idLat}^0)$ is invariant by multiplication by a fixed norm-1 replete ideal, the ideal $I' = J \cdot I \cdot \mathcal{N}(I)^{-1/d}$ is then distributed from $\mathcal{U}(\mathrm{idLat}^0)$. The data-processing inequalities of the statistical distance and Rényi divergence imply that the above relations also hold for $\widetilde{D_i}$ in place of D_i, for all i. Furthermore, by choice of X, the Rényi divergence from $\mathcal{U}(\mathcal{I}_{A,4A})$ to $\widetilde{D_4}$ is equal to $\Theta(1)$.

Using the probability preservation properties of the statistical distance and Rényi divergence, we obtain that for any event $E \subseteq \mathcal{I}_{A,4A}$, we have:

$$\widetilde{D_1}(E) \geq \frac{\mathcal{U}(E) - 2^{-\Omega(d)}}{\Theta(1)} - 2^{-\Omega(d)} = \frac{\mathcal{U}(E)}{\Theta(1)} - 2^{-\Omega(d)}$$

which completes the proof. □

3.3 From ideal to their inverses

Let $\mathcal{W}$ be a set of fractional ideals. Below, we reduce $\mathcal{W}^{-1}$-avg-id-HSVP to $\mathcal{W}$-avg-id-HSVP and $\mathcal{I}_{A,4A}$-avg-id-HSVP for some appropriate integer A and approximation factors. Recall that $\mathcal{W}^{-1}$ refers to the set $\{I^{-1}, I \in \mathcal{W}\}$.

The reduction is described as an algorithm, $\texttt{InverseToIntegral}^{\mathcal{W}}_{A,X}$ (Algorithm 3.3), which takes as input the inverse I^{-1} of an integral ideal $I \in \mathcal{W}$ and returns a short non-zero element of I^{-1}. It is parametrized by an integer A and a compact set X satisfying the conditions of Lemma 3.2. It relies on four oracles: oracle $\mathcal{O}_{\mathcal{W}}$ for solving $\mathcal{W}$-avg-id-HSVP, oracle $\mathcal{O}_{\mathcal{I}}$ for $\mathcal{I}_{A,4A}$-avg-id-HSVP, oracle $\mathcal{F}$ for factoring integral ideals; and oracle $\mathcal{S}$ for sampling from $I \cap X$ for $I \in \mathrm{idLat}^0$. Recall that $\mathcal{F}$ can be instantiated as a quantum polynomial time algorithm. An instantiation of oracle $\mathcal{S}$ will be provided in Section 4, based on the design of a nice set X for Lemma 3.2. The reduction first uses $\mathcal{O}_{\mathcal{W}}$ on the inverse I of its input, which gives a short non-zero vector $v_I \in I$. Then $\texttt{RandomizeIdeal}_{A,X}$ (introduced in the previous subsection) is invoked to randomize I into a uniform integral ideal $\mathfrak{b}$ with norm in $[A, 4A]$. $\texttt{RandomizeIdeal}_{A,X}$ also returns a short non-zero $y_{(I\mathfrak{b})^{-1}}$ in $(I\mathfrak{b})^{-1}$. Then $\mathcal{O}_{\mathcal{I}}$ is invoked on $\mathfrak{b}$ and returns a short non-zero $v_{\mathfrak{b}}$ in $\mathfrak{b}$. The reduction finally outputs $v_{\mathfrak{b}} \cdot y_{(I\mathfrak{b})^{-1}} \in I^{-1}$ that is short and non-zero.

The astute reader will notice that, in the above description, the vector v_I and hence the oracle $\mathcal{O}_{\mathcal{W}}$ do not seem to be used in the subsequent steps. In fact, we will be able to instantiate $\mathcal{S}$ only if given a short basis of I (see Lemma 4.9). The approximation factor reached by $\mathcal{O}_{\mathcal{W}}$ will lead to a lower bound condition on A: the smaller the approximation factor, the smaller the lower bound on A.

Theorem 3.4. *Let $\mathcal{W}$ be a finite set of fractional ideals. Let $\gamma_{\mathcal{W}}, \gamma_{\mathcal{I}} \geq 1$ and $A \geq \max(\delta_K^d, d^d \Delta_K^c)$ for c as in Lemma 2.3. Let X be a compact subset of $K_{\mathbb{R}} \setminus \{0\}$*

Algorithm 3.3 $\texttt{InverseToIntegral}^{\mathcal{W}}$

Input: I^{-1} with $I \in \mathcal{W}$.
Parameters: A integer and $X \subset K_{\mathbb{R}} \setminus \{0\}$ compact.
Oracles: $\mathcal{O}_{\mathcal{W}}$ for $\mathcal{W}$-avg-id-HSVP$_{\gamma_{\mathcal{W}}}$, $\mathcal{O}_{\mathcal{I}}$ for $\mathcal{I}_{A,4A}$-avg-id-HSVP$_{\gamma_{\mathcal{I}}}$,
$\mathcal{F}$ for factoring integral ideals and $\mathcal{S}$ for sampling from $\mathcal{U}(I \cap X)$ for $I \in \text{idLat}^0$.
Output: $x \in I^{-1} \setminus \{0\}$.

1: Compute $v_I \leftarrow \mathcal{O}_{\mathcal{W}}(I)$.
2: If $v_I = \bot$, then return $\bot$.
3: Compute $\mathbf{B}_I = \texttt{ReduceIdeal}(I, v_I)$.
4: Sample $(\mathfrak{b}, y_{(I\mathfrak{b})^{-1}}) \leftarrow \texttt{RandomizeIdeal}_{A,X}(\mathbf{B}_I)$, using $\mathcal{F}$ and $\mathcal{S}$.
5: Compute $v_{\mathfrak{b}} \leftarrow \mathcal{O}_{\mathcal{I}}(\mathfrak{b})$.
6: If $v_{\mathfrak{b}} = \bot$, then return $\bot$.
7: Return $v_{\mathfrak{b}} \cdot y_{(I\mathfrak{b})^{-1}}$.

whose elements are η-balanced for some $\eta > 1$ and satisfy the assumptions of Lemma 3.2 for A and $B = 4A$. Assume that $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}| \leq c'$ for some constant $c' < 1$. Let $\mathcal{O}_{\mathcal{W}}$ an oracle for $\mathcal{W}$-avg-id-HSVP$_{\gamma_{\mathcal{W}}}$ with success probability $\varepsilon_{\mathcal{W}}$ and $\mathcal{O}_{\mathcal{I}}$ an oracle for $\mathcal{I}_{A,4A}$-avg-id-HSVP$_{\gamma_{\mathcal{I}}}$ with success probability $\varepsilon_{\mathcal{I}}$.

When given access to $\mathcal{O}_{\mathcal{W}}$, $\mathcal{O}_{\mathcal{I}}$, an integral ideal-factoring oracle $\mathcal{F}$ and an oracle $\mathcal{S}$ for sampling from $\mathcal{U}(I \cap X)$ for $I \in \text{idLat}^0$, $\texttt{InverseToIntegral}^{\mathcal{W}}_{A,X}$ runs in expected time polynomial in $\log A, \log \Delta_K, A/|\mathcal{P}_{A,4A}|$ and the size of its input. Further, if its input I is such that I is distributed from $\mathcal{U}(\mathcal{W})$, it outputs $x \neq \bot$ with probability $\geq \varepsilon_{\mathcal{I}} \cdot (\varepsilon_{\mathcal{W}}/\Theta(1) - 2^{-\Omega(d)})$. If $x \neq \bot$, then we have

$$x \in I^{-1} \setminus \{0\} \quad \text{and} \quad \|x\| \leq \gamma' \cdot \text{Vol}(I^{-1})^{1/d},$$

for $\gamma' = 85 \cdot \gamma_{\mathcal{I}} \cdot \Delta_K^{1/d} \cdot d \cdot \eta$.

Proof. Assume first that neither v_I nor $v_{\mathfrak{b}}$ is equal to $\bot$. As the assumptions of Lemma 3.3 are satisfied, we have $y_{(I\mathfrak{b})^{-1}} \in (I\mathfrak{b})^{-1} \setminus \{0\}$ and

$$\|y_{(I\mathfrak{b})^{-1}}\| \leq 85 \cdot d \cdot \eta \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I\mathfrak{b})^{-1/d}.$$

Now, by assumption, we have that $v_{\mathfrak{b}} \in \mathfrak{b} \setminus \{0\}$ satisfies $\|v_{\mathfrak{b}}\| \leq \gamma_{\mathcal{I}} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{b})^{1/d}$. We then obtain that $x = v_{\mathfrak{b}} \cdot y_{(I\mathfrak{b})^{-1}} \in I^{-1}$ is non-zero and satisfies:

$$\|x\| \leq \|v_{\mathfrak{b}}\| \cdot \left\|y_{(I\mathfrak{b})^{-1}}\right\| \leq \gamma_{\mathcal{I}} \cdot \Delta_K^{3/(2d)} \cdot 85 \cdot d \cdot \eta \cdot \mathcal{N}(I)^{-1/d}.$$

Towards completing the proof, not that the algorithm succeeds if and only if neither v_I nor $v_{\mathfrak{b}}$ is equal to $\bot$. The probability that v_I is not $\bot$ is exactly $\varepsilon_{\mathcal{I}}$. Using Lemma 3.3 with the event E set to $\mathcal{O}_{\mathcal{I}}(\mathfrak{b})$ succeeding, we obtain that $v_{\mathfrak{b}}$ is not $\bot$ with probability $\geq \varepsilon_{\mathcal{W}}/\Theta(1) - 2^{-\Omega(d)}$. Note that the second probability is over the internal randomness of $\texttt{RandomizeIdeal}_{A,X}(\mathbf{B}_I)$. □

4 The sampling set

Lemma 3.2 states that if a compact X satisfies a certain number of conditions, then the output distribution of $\texttt{IdealRound}_X$ resembles the uniform distribution over integral ideals whose norms belong to a prescribed interval. In this subsection, we show that the set $\mathcal{B}^{\eta}_{A,B}$ defined below satisfies those constraints. We will later also use the fact that its elements are η-balanced. An instantiation of the set $\mathcal{B}^{\eta}_{A,B}$ can be visualized in Figure 1.

Definition 4.1. *Let $B > A > 0$ and $\eta > 1$. We define the set:*

$$\mathcal{B}^{\eta}_{A,B} = \left\{ x \in K_{\mathbb{R}}^{\times} \,\middle|\, \mathcal{N}(x) \in [A, B], \quad \left\| \mathrm{Ln}\left(\frac{x}{\mathcal{N}(x)^{1/d}} \right) \right\|_2 \leq \ln(\eta) \right\}.$$

The purpose of this section is to prove the following theorem.

Theorem 4.2. *Let $A, B, \eta, \delta > 0$ satisfying $A^{1/d} \geq d^3 \cdot \eta \cdot \max(\Delta_K^{3/(2d)}, \delta)$, $B/A \geq 4$ and $\eta \geq \mathrm{e}$. The set $\mathcal{B}^{\eta}_{A,B}$ is compact and invariant by complex rotations, satisfies the conditions of Lemma 3.2 and its elements are η-balanced. Further, there exists an algorithm $\texttt{SampleUniform}^{\eta}_{A,B}$ that, given as input a basis $\mathbf{B}_I$ of a norm-1 replete ideal satisfying $\|\mathbf{B}_I^{\star}\| \leq \delta$, samples uniformly in $I \cap \mathcal{B}^{\eta}_{A,B}$ and whose expected running time is polynomial in $\log B$, d and B/A.*

4.1 Volume of the set $\mathcal{B}^{\eta}_{A,B}$

Before proving that the assumptions of Lemma 3.2 are satisfied by $\mathcal{B}^{\eta}_{A,B}$, we first study its volume and its approximate invariance under translation.

Lemma 4.3. *For any $B > A > 0$ and $\eta > 1$, we have*

$$\mathrm{Vol}\left(\mathcal{B}^{\eta}_{A,B}\right) = \frac{2^{d_{\mathbb{R}}} \cdot (2\sqrt{2}\pi)^{d_{\mathbb{C}}} \cdot V_{d_{\mathbb{R}}+d_{\mathbb{C}}-1}}{\sqrt{d}} \cdot (B - A) \cdot (\ln \eta)^{d_{\mathbb{R}}+d_{\mathbb{C}}-1},$$

where V_n is the volume of the n-dimensional unit ℓ_2 hyperball for any $n \geq 1$.

The computation of the volume proceeds by a change of variable, between $\mathbb{R}^d$ and $\sigma(K_{\mathbb{R}})$. The relevant aspect of the volume formula for the present work is the linear dependency in $(B - A) \cdot (\ln \eta)^{d_{\mathbb{R}}+d_{\mathbb{C}}-1}$.

Proof. Let $\mathbf{e}_j$ denote the j-th elementary unit vector, we fix $\mathbf{C} = (\mathbf{c}_1, \ldots, \mathbf{c}_d) \in \mathbb{C}^{d \times d}$ an orthonormal $\mathbb{R}$-basis of $\sigma(K_{\mathbb{R}}) \subset \mathbb{C}^d$ defined by

$$\begin{array}{lll}
\mathbf{c}_j & = \mathbf{e}_j & \text{for } 1 \leq j \leq d_{\mathbb{R}}, \\
\mathbf{c}_{d_{\mathbb{R}}+j} & = 1/\sqrt{2} \cdot (\mathbf{e}_{d_{\mathbb{R}}+j} + \mathbf{e}_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}) & \text{for } 1 \leq j \leq d_{\mathbb{C}}, \\
\mathbf{c}_{d_{\mathbb{R}}+d_{\mathbb{C}}+j} & = i/\sqrt{2} \cdot (\mathbf{e}_{d_{\mathbb{R}}+j} - \mathbf{e}_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}) & \text{for } 1 \leq j \leq d_{\mathbb{C}}.
\end{array}$$

We let ϕ be the isomorphism sending an element of $\sigma(K_{\mathbb{R}})$ to its coordinates in the basis $\mathbf{C}$. Since $\mathbf{C}$ is orthonormal, the map ϕ preserves the geometry. In

particular, the volume of $\mathcal{B}_{A,B}^{\eta}$ is the same as the volume of $\phi(\mathcal{B}_{A,B}^{\eta})$ (which is a d-dimensional object in $\mathbb{R}^d$).

In order to compute this volume, we first introduce the set

$$\mathcal{B}_{A,B}^{\eta\,+} = \left\{ x \in \mathcal{B}_{A,B}^{\eta} \,\middle|\, \sigma_i(x) > 0 \text{ for all } 1 \leq i \leq d_{\mathbb{R}} \right\},$$

i.e., the elements of $\mathcal{B}_{A,B}^{\eta}$ whose real embeddings are all positive. Since $\mathcal{B}_{A,B}^{\eta}$ is invariant by complex rotations, the set $\mathcal{B}_{A,B}^{\eta}$ is the union of $2^{d_{\mathbb{R}}}$ distinct copies of $\mathcal{B}_{A,B}^{\eta\,+}$, hence we can focus on computing the volume of the latter. In order to compute this volume, we will exhibit a function F transforming a "nice box" of $\mathbb{R}^d$ into the set $\phi(\mathcal{B}_{A,B}^{\eta\,+})$. We will then use this function to perform a change of variable and compute the volume of $\phi(\mathcal{B}_{A,B}^{\eta\,+})$ from the volume of the nice box.

Defining the function F. Let H be the $(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)$-dimensional subspace of $\mathbb{R}^d$ spanned by $\mathrm{Ln}(\mathcal{O}_K^{\times})$, i.e.,

$$H = \left\{ x \in \mathbb{R}^d \,\middle|\, \sum_{j \leq d} x_j = 0 \wedge \forall d_{\mathbb{R}} < j \leq d_{\mathbb{R}} + d_{\mathbb{C}} : x_{j+d_{\mathbb{C}}} = x_j \right\},$$

and let $\mathbf{B} = (b_{i,j}) \in \mathbb{R}^{d \times (d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}$ be any orthonormal basis of H. We define the following function f from $\mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times \mathbb{R}^{d_{\mathbb{C}}}$ to $\sigma(K_{\mathbb{R}})$ as:

$$f(N, \mathbf{z}, \boldsymbol{\theta}) = \exp(N/d) \cdot \exp\left(\mathbf{B} \cdot \mathbf{z} + i\hat{\boldsymbol{\theta}}\right),$$

where the second function exp is applied coordinate-wise to the vector $\mathbf{B} \cdot \mathbf{z} + i\hat{\boldsymbol{\theta}}$, and where $\hat{\boldsymbol{\theta}} = (\mathbf{0}^{d_{\mathbb{R}}} \mid \boldsymbol{\theta}^T \mid -\boldsymbol{\theta}^T)^T \in \mathbb{R}^d$. Note that f is injective on the set $\mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times [0, 2\pi)^{d_{\mathbb{C}}}$, and that its image indeed lies in $\sigma(K_{\mathbb{R}})$ (it even lies in the subset of $\sigma(K_{\mathbb{R}})$ whose first $d_{\mathbb{R}}$ coordinates are positive).

In order to obtain a transformation from $\mathbb{R}^d$ to itself, we compose the above function f with the function ϕ, and we obtain $F = \phi \circ f : \mathbb{R}^d \to \mathbb{R}^d$, which is injective on $\mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times [0, 2\pi)^{d_{\mathbb{C}}}$. Moreover, by letting $\mathrm{Ball}^{(n)}(R)$ denote the Euclidean ball of radius R in $\mathbb{R}^n$, we have that

$$\phi(\mathcal{B}_{A,B}^{\eta\,+}) = F\Big([\ln A, \ln B] \times \mathrm{Ball}^{(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}} \Big).$$

Indeed, let $(N, \mathbf{z}, \boldsymbol{\theta}) \in \mathbb{R} \times \mathbb{R}^{d_{\mathbb{R}} + d_{\mathbb{C}} - 1} \times [0, 2\pi)^{d_{\mathbb{C}}}$ and let $x = \sigma^{-1}(f(N, \mathbf{z}, \boldsymbol{\theta})) \in K_{\mathbb{R}}$. Then $\mathcal{N}(x) = \exp(N)$ (because $\mathbf{B} \cdot \mathbf{z}$ belongs to H, so the sum of its coordinates is zero) and $\mathrm{Ln}(x/\mathcal{N}(x)^{1/d}) = \mathbf{B} \cdot \mathbf{z}$, which implies that $\left\| \mathrm{Ln}(x/\mathcal{N}(x)^{1/d}) \right\|_2 = \|\mathbf{z}\|_2$ since $\mathbf{B}$ is orthonormal. The inclusion from right to left follows from these two observations and the definition of $\mathcal{B}_{A,B}^{\eta}$. For the inclusion from left to right, it suffices to observe that a pre-image of $x \in \mathcal{B}_{A,B}^{\eta\,+}$ is obtained by taking $N = \ln(\mathcal{N}(x))$, $\mathbf{z}$ equal to the coordinates of $\mathrm{Ln}(x/\mathcal{N}(x)^{1/d})$ in basis $\mathbf{B}$, and $\boldsymbol{\theta}$ equal to the arguments of $\sigma_i(x)$.

The set $[\ln A, \ln B] \times \mathrm{Ball}^{(d_{\mathbb{R}} + d_{\mathbb{C}} - 1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}$ is the "nice set" we mentioned above. To compute the volume of $\mathcal{B}_{A,B}^{\eta\,+}$, we will change variables, using

the function F, in order to transform $\phi(\mathcal{B}^{\eta,+}_{A,B})$ into this nice set. In order to perform this change of variables, we first compute the Jacobian matrix of F and its determinant.

Computing the Jacobian matrix of F. For $1 \leq i \leq d$, let F_i be the function corresponding to the i-th coordinate of F (note that F_i goes from $\mathbb{R}^d$ to $\mathbb{R}$). By definition of F and choice of $\mathbf{C}$ (which defines ϕ), one can check that the following holds, for $(N, \mathbf{z}, \boldsymbol{\theta}) \in \mathbb{R}^d$:

$$\begin{aligned}
F_i(N, \mathbf{z}, \boldsymbol{\theta}) &= \exp\Big(N/d + \textstyle\sum_j b_{i,j} z_j\Big) && \text{for } 1 \leq i \leq d_{\mathbb{R}},\\
F_{d_{\mathbb{R}}+i}(N, \mathbf{z}, \boldsymbol{\theta}) &= \sqrt{2} \cdot \exp\Big(N/d + \textstyle\sum_j b_{d_{\mathbb{R}}+i,j} z_j\Big) \cdot \cos(\theta_i) && \text{for } 1 \leq i \leq d_{\mathbb{C}},\\
F_{d_{\mathbb{R}}+d_{\mathbb{C}}+i}(N, \mathbf{z}, \boldsymbol{\theta}) &= \sqrt{2} \cdot \exp\Big(N/d + \textstyle\sum_j b_{d_{\mathbb{R}}+d_{\mathbb{C}}+i,j} z_j\Big) \cdot \sin(\theta_i) && \text{for } 1 \leq i \leq d_{\mathbb{C}}.
\end{aligned}$$

Note that here, we used the fact that for all $1 \leq j \leq d_{\mathbb{R}} + d_{\mathbb{C}} - 1$ and all $1 \leq i \leq d_{\mathbb{C}}$, it holds that $b_{d_{\mathbb{R}}+i,j} = b_{d_{\mathbb{R}}+d_{\mathbb{C}}+i,j}$ since the columns of $\mathbf{B}$ are in H. We then obtain that:

$$\begin{aligned}
\partial_N(F_i(N, \mathbf{z}, \boldsymbol{\theta})) &= F_i(N, \mathbf{z}, \boldsymbol{\theta}) \cdot \tfrac{1}{d} && \text{for } i \leq d,\\
\partial_{z_j}(F_i(N, \mathbf{z}, \boldsymbol{\theta})) &= F_i(N, \mathbf{z}, \boldsymbol{\theta}) \cdot b_{i,j} && \text{for } i \leq d \text{ and } j \leq d_{\mathbb{R}} + d_{\mathbb{C}} - 1,\\
\partial_{\theta_j}(F_{d_{\mathbb{R}}+j}(N, \mathbf{z}, \boldsymbol{\theta})) &= F_{d_{\mathbb{R}}+j}(N, \mathbf{z}, \boldsymbol{\theta}) \cdot \tfrac{-\sin\theta_j}{\cos\theta_j} && \text{for } j \leq d_{\mathbb{C}},\\
\partial_{\theta_j}(F_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}(N, \mathbf{z}, \boldsymbol{\theta})) &= F_{d_{\mathbb{R}}+d_{\mathbb{C}}+j}(N, \mathbf{z}, \boldsymbol{\theta}) \cdot \tfrac{\cos\theta_j}{\sin\theta_j} && \text{for } j \leq d_{\mathbb{C}},\\
\partial_{\theta_j}(F_i(N, \mathbf{z}, \boldsymbol{\theta})) &= 0 && \text{else.}
\end{aligned}$$

In short, the Jacobian matrix of F is $D_F(N, \mathbf{z}, \boldsymbol{\theta}) = \mathrm{diag}_i(F_i(N, \mathbf{z}, \boldsymbol{\theta})) \cdot \mathbf{M}$, where

$$\mathbf{M} = \left(\begin{array}{c|c|c} \frac{1}{d} & & \mathbf{0} \\ \vdots & \mathbf{B} & \mathrm{diag}_i\Big(\frac{-\sin\theta_i}{\cos\theta_i}\Big) \\ \frac{1}{d} & & \mathrm{diag}_i\Big(\frac{\cos\theta_i}{\sin\theta_i}\Big) \end{array}\right).$$

Computing the Jacobian determinant. We now compute the determinant of the matrix $D_F(N, \mathbf{z}, \boldsymbol{\theta})$. First, we have

$$\det\Big(\mathrm{diag}_i\Big(F_i(N, \mathbf{z}, \theta)\Big)\Big) = \exp(N) \cdot (\sqrt{2})^{2d_{\mathbb{C}}} \cdot \prod_{i \leq d_{\mathbb{C}}} \big(\sin\theta_i \cdot \cos\theta_i\big),$$

where we used again the fact that the sum of the coordinates of $\mathbf{B} \cdot \mathbf{z}$ is zero. We now focus on $\mathbf{M}$. Let $\widehat{\mathbf{M}}_{i,j}$ be the matrix $\mathbf{M}$ where we have removed the i-th line and j-th column. Observe that for $i \leq d_{\mathbb{C}}$, the $(d_{\mathbb{R}} + i)$-th row of $\mathbf{M}$ and the $(d_{\mathbb{R}} + d_{\mathbb{C}} + i)$-th row of $\mathbf{M}$ coincide except in the $(d_{\mathbb{R}} + d_{\mathbb{C}} + i)$-th column. So developing the determinant of $\mathbf{M}$ on the $(d_{\mathbb{R}} + d_{\mathbb{C}} + i)$-th column leads to

$$\begin{aligned}
|\det \mathbf{M}| &= \left|\frac{\sin\theta_i}{\cos\theta_i} + \frac{\cos\theta_i}{\sin\theta_i}\right| \cdot \left|\det \widehat{\mathbf{M}}_{d_{\mathbb{R}}+d_{\mathbb{C}}+i, d_{\mathbb{R}}+d_{\mathbb{C}}+i}\right| \\
&= \frac{1}{|\sin\theta_i \cdot \cos\theta_i|} \cdot \left|\det \widehat{\mathbf{M}}_{d_{\mathbb{R}}+d_{\mathbb{C}}+i, d_{\mathbb{R}}+d_{\mathbb{C}}+i}\right|.
\end{aligned}$$

Here, to check that the signs are correct, we observe that we can permute the rows of $\mathbf{M}$ without changing the absolute value of the determinant, and move the row with index $d_{\mathbb{R}}+i$ to position $d_{\mathbb{R}}+d_{\mathbb{C}}+i+1$ (so that it follows directly the row $d_{\mathbb{R}}+d_{\mathbb{C}}+i$). This ensures that both minor matrices are the same, and that the signs are opposite when we develop according to the $(d_{\mathbb{R}}+d_{\mathbb{R}}+i)$-th column. Repeating the process on the $d_{\mathbb{C}}$ last columns of M, we obtain that

$$|\det \mathbf{M}| = \Big(\prod_i \frac{1}{|\sin\theta_i \cdot \cos\theta_i|}\Big) \cdot |\det \widehat{\mathbf{M}}|,$$

where $\widehat{\mathbf{M}}$ is the top-left square sub-matrix of $\mathbf{M}$ of dimension $d_{\mathbb{R}}+d_{\mathbb{C}}$. Let $\mathbf{B}_0 \in \mathbb{R}^{d_{\mathbb{R}}\times(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}$ and $\mathbf{B}_1 \in \mathbb{R}^{d_{\mathbb{C}}\times(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}$ be sub-blocks of the matrix $\mathbf{B}$ such that $\mathbf{B} = (\mathbf{B}_0^T|\mathbf{B}_1^T|\mathbf{B}_1^T)^T$ (recall that $\mathbf{B}$ is an arbitrary orthonormal basis of H). Then all the entries of the first column of $\widehat{\mathbf{M}}$ are equal to $1/d$ and the remaining $d_{\mathbb{R}}+d_{\mathbb{C}}-1$ are $(\mathbf{B}_0^T|\mathbf{B}_1^T)^T$. Let us consider the following distortion of $\widehat{\mathbf{M}}$:

$$\mathbf{N} = \left(\begin{array}{c|c} \frac{1}{d}\cdot\mathbf{1}_{d_{\mathbb{R}}} & \mathbf{B}_0 \\ \hline \frac{\sqrt{2}}{d}\cdot\mathbf{1}_{d_{\mathbb{C}}} & \sqrt{2}\cdot\mathbf{B}_1 \end{array}\right),$$

where $\mathbf{1}_k$ refers to the k-dimensional all-1 vector. Then $\det\widehat{\mathbf{M}} = \sqrt{2}^{-d_{\mathbb{C}}}\cdot\det\mathbf{N}$. Furthermore, note that $\mathbf{N}^T\cdot\mathbf{N} = \mathrm{diag}(1/d,1,\ldots,1)$, because the columns of $\mathbf{B}$ are orthonormal and in H (so the sums of their coordinates are zero). This gives us that $|\det\mathbf{N}| = 1/\sqrt{d}$. Unrolling the above, we obtain

$$|\det\mathbf{M}| = \frac{1}{\sqrt{d}\cdot\sqrt{2}^{d_{\mathbb{C}}}\cdot\prod_i|\sin\theta_i\cdot\cos\theta_i|},$$

and

$$|\det(D_F(N,\mathbf{z},\boldsymbol{\theta}))| = \frac{\exp(N)\cdot\sqrt{2}^{d_{\mathbb{C}}}}{\sqrt{d}}.$$

Change of variables. We finally perform the change of variables using the function F to compute the volume of $\mathcal{B}^{\eta+}_{A,B}$ (recall that $\mathrm{vol}(\mathcal{B}^{\eta}_{A,B}) = 2^{d_{\mathbb{R}}}\cdot\mathrm{vol}(\mathcal{B}^{\eta+}_{A,B})$). Letting $\mathbb{1}_S(\cdot)$ denote the indicator function of a set S, we have

$$\begin{aligned}
\mathrm{vol}(\mathcal{B}^{\eta+}_{A,B}) &= \int_{\mathbf{x}\in\mathbb{R}^d} \mathbb{1}_{\phi(\mathcal{B}^{\eta+}_{A,B})}(\mathbf{x})\,\mathrm{d}\mathbf{x} \\
&= \int_{\substack{N\in[\ln A,\ln B]\\ \mathbf{z}\in\mathrm{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln\eta)\\ \boldsymbol{\theta}\in[0,2\pi)^{d_{\mathbb{C}}}}} |\det(D_F(N,\mathbf{z},\theta))|\,\mathrm{d}\boldsymbol{\theta}\,\mathrm{d}\mathbf{z}\,\mathrm{d}N \\
&= \frac{\sqrt{2}^{d_{\mathbb{C}}}}{\sqrt{d}}\cdot\int_{N\in[\ln A,\ln B]}\exp(N)\,\mathrm{d}N\cdot\int_{\mathbf{z}\in\mathrm{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln\eta)}\mathrm{d}\mathbf{z}\cdot\int_{\boldsymbol{\theta}\in[0,2\pi)^{d_{\mathbb{C}}}}\mathrm{d}\boldsymbol{\theta} \\
&= \frac{\sqrt{2}^{d_{\mathbb{C}}}}{\sqrt{d}}\cdot(B-A)\cdot V_{d_{\mathbb{R}}+d_{\mathbb{C}}-1}\cdot(\ln\eta)^{d_{\mathbb{R}}+d_{\mathbb{C}}-1}\cdot(2\pi)^{d_{\mathbb{C}}},
\end{aligned}$$

as desired. □

The proof of Lemma 4.3 gives us the volume of the set $\mathcal{B}^{\eta}_{A,B}$, but it also a way to sample uniformly in it.

Lemma 4.4. *There exists a probabilistic algorithm that samples from $\mathcal{U}(\mathcal{B}^{\eta}_{A,B})$ for any $B > A > 0$ and $\eta > 1$. The expected running time of this algorithm is polynomial in $\log B$, d (the degree of K) and B/A.*

Proof. Let ϕ and F be the same functions as in the proof of Lemma 4.3. Recall that

$$F\Big([\ln A, \ln B] \times \mathrm{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}\Big) = \phi(\mathcal{B}^{\eta+}_{A,B}), \tag{2}$$

and that F is injective on this set. It can be observed from their definitions that F, ϕ and ϕ^{-1} can be computed in time polynomial in d.

Note that if we can sample from $\mathcal{U}(\phi(\mathcal{B}^{\eta+}_{A,B}))$ in time T, then we can sample from $\mathcal{U}(\mathcal{B}^{\eta}_{A,B})$ in time $T + \mathrm{poly}(d)$. Indeed, it suffices to sample x from $\mathcal{U}(\phi(\mathcal{B}^{\eta+}_{A,B}))$; compute $\phi^{-1}(x)$ (which can be done in time $\mathrm{poly}(d)$); sample uniform signs $(\varepsilon_i)_i \in \{-1, 1\}^{d_{\mathbb{R}}}$; and output $\sigma^{-1}\big((\varepsilon_1, \ldots, \varepsilon_{d_{\mathbb{R}}}, 1, \ldots 1)\big) \cdot \phi^{-1}(x)$. In the rest of this proof, we then focus on sampling the random variable $\mathcal{U}(\phi(\mathcal{B}^{\eta+}_{A,B}))$.

Let Y be a random variable distributed over $[\ln A, \ln B] \times \mathrm{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}$ with density probability $f_Y(N, \mathbf{z}, \theta)$ proportional to $|\det(D_F(N, \mathbf{z}, \theta))|$, i.e., proportional to $\exp(N)$. From Equation (2) above, we know that $F(Y)$ is distributed as $\mathcal{U}(\phi(\mathcal{B}^{\eta+}_{A,B}))$. We are then reduced to sampling such a random variable Y.

Note that the domain of Y is a "nice box": $[\ln A, \ln B] \times \mathrm{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta) \times [0, 2\pi)^{d_{\mathbb{C}}}$. In this domain, we can sample a uniformly random variable Z in time $\mathrm{poly}(d)$ (to sample from the ball $\mathrm{Ball}^{(d_{\mathbb{R}}+d_{\mathbb{C}}-1)}(\ln \eta)$, one can sample a Gaussian element and then renormalize it inside the ball). To obtain a sample from Y, we then keep Z with probability $\exp(Z_N)/B$, where Z_N is the first coordinate of Z. Note that the rejection probability is indeed between 0 and 1 since $Z_N \leq \ln(B)$.

It only remains to estimate the cost of the rejection step. Since $Z_N \geq \ln(A)$, the probability of keeping Z is at least A/B, and so the expected number of rejections before acceptance is bounded from above by B/A. □

4.2 Properties of the set $\mathcal{B}^{\eta}_{A,B}$

The goal of this subsection is to prove that the set $\mathcal{B}^{\eta}_{A,B}$ satisfies the properties needed to apply Lemma 3.2.

Lemma 4.5. *For any $B > A > 0$ and $\eta > 1$, the set $\mathcal{B}^{\eta}_{A,B}$ is compact, invariant by complex rotations and its elements are η-balanced.*

Proof. Compactness follows from the fact that $\mathcal{B}^{\eta}_{A,B}$ is closed and contained in the ball in infinity norm with radius $\eta \cdot B^{1/d}$. Invariance by complex rotations follows from the fact that both $\mathcal{N}(\cdot)$ and $\mathrm{Ln}(\cdot)$ are invariant by complex rotations

(i.e., we have $\mathcal{N}(\zeta x) = \mathcal{N}(x)$ and $\mathrm{Ln}(\zeta x) = \mathrm{Ln}(x)$ if $x \in K_{\mathbb{R}}$ and $\zeta \in K_{\mathbb{R}}$ is such that $|\sigma_i(\zeta)| = 1$ for all i's). Let $x \in \mathcal{B}^{\eta}_{A,B}$, we have that

$$\left\| \frac{x}{\mathcal{N}(x)^{1/d}} \right\|_{\infty} = \exp\left(\left\| \mathrm{Ln}\left(\frac{x}{\mathcal{N}(x)^{1/d}} \right) \right\|_{\infty} \right) \leq \exp\left(\left\| \mathrm{Ln}(\frac{x}{\mathcal{N}(x)^{1/d}}) \right\|_{2} \right) \leq \eta.$$

The same holds for $\mathcal{N}(x)^{1/d}/x$ since $\left\|\mathrm{Ln}\left(x/\mathcal{N}(x)^{1/d}\right)\right\|_{\infty} = \left\|\mathrm{Ln}\left(\mathcal{N}(x)^{1/d}/x\right)\right\|_{\infty}$, which proves that x is η-balanced. □

We now prove that the slices $\mathrm{Ln}(\mathcal{B}^{\eta}_{A,B}) \cap H_t$ are empty when $t \notin [\ln A, \ln B]$ and have constant volume otherwise.

Lemma 4.6. *Let $B > A > 0$ and $\eta > 1$. For $t \in \mathbb{R}$, we define $H_t = \{x \in \mathrm{Ln}\, K_{\mathbb{R}} | \sum_i x_i = t\}$. Then $\mathrm{Ln}(\mathcal{B}^{\eta}_{A,B}) \cap H_t = \emptyset$ for $t \notin [\ln A, \ln B]$, and the volume of $\mathrm{Ln}(\mathcal{B}^{\eta}_{A,B}) \cap H_t$ is constant for $t \notin [\ln A, \ln B]$.*

Proof. By definition of $\mathcal{B}^{\eta}_{A,B}$, we have that

$$\mathrm{Ln}\left(\mathcal{B}^{\eta}_{A,B}\right) = \left\{ \mathbf{x} \in \mathrm{Ln}(K_{\mathbb{R}}) : \sum_{i \leq d} x_i \in [\ln A, \ln B], \left\| x - \left(\sum_{i \leq d} x_i \right) \cdot \mathbf{1}_d \right\|_2 \leq \ln \eta \right\},$$

where $\mathbf{1}_d$ refers to the d-dimensional all-1 vector. The intersection with H_t is the empty set if $t \notin [\ln A, \ln B]$. Otherwise, it is the ball centered in $t \cdot \mathbf{1}$ with radius $\ln(\eta)$, whose volume do not depend on t. □

At this stage, only the first condition of Lemma 3.2 remains to be proved. We start by an auxiliary lemma, where we prove that if we shift the set $\mathcal{B}^{\eta}_{A,B}$ by some small vector, then the resulting set is included in another slightly larger set $\mathcal{B}^{\eta'}_{A',B'}$. The parameter f in the lemma below quantifies how small the shift vector needs to be, as a function of the parameters A and η. For the rest of the article, one can think of f as being of the order of $\mathrm{poly}(d)$.

Lemma 4.7. *Let $B > A > 0$, $\eta > 1$ and $v \in K_{\mathbb{R}}$. Assume that $A^{1/d} \geq \eta \cdot f \cdot \|v\|_{\infty}$ for some $f > 1$. Then*

$$\mathcal{B}^{\eta}_{A,B} + v \ \subset \ \mathcal{B}^{\eta'}_{A',B'}$$

with $A' = A \cdot (1 - 1/f)^d$, $B' = B \cdot (1 + 1/f)^d$ and $\eta' = \eta \cdot \exp(2\sqrt{d}/(f-1))$.

Proof. Let $x \in \mathcal{B}^{\eta}_{A,B}$, we are going to show that $x + v \in \mathcal{B}^{\eta'}_{A',B'}$. The definition of $\mathcal{B}^{\eta}_{A,B}$ and the fact that $A^{1/d} \geq \eta \cdot f \cdot \|v\|_{\infty}$ imply that we have, for every i,

$$\frac{|v_i|}{|x_i|} \leq \frac{\|v\|_{\infty}}{|x_i|} \leq \frac{\|v\|_{\infty} \cdot \eta}{\mathcal{N}(x)^{1/d}} \leq \frac{\|v\|_{\infty} \cdot \eta}{A^{1/d}} \leq \frac{1}{f}.$$

The triangle inequality then gives that $|x_i + v_i| > 0$ for all i, and hence that $x + v \in K_{\mathbb{R}}^{\times}$. Further, note that

$$\frac{\mathcal{N}(x+v)}{\mathcal{N}(x)} = \prod_i \left| \frac{x_i + v_i}{x_i} \right| = \prod_i \left| 1 + \frac{v_i}{x_i} \right|.$$

Since $|v_i/x_i| \leq 1/f$ holds for all i, this implies that $\mathcal{N}(x+v)/\mathcal{N}(x) \in [(1-1/f)^d, (1+1/f)^d]$, which in turn shows that $\mathcal{N}(x+v) \in [A', B']$.

Towards completing the proof, note that

$$\begin{aligned}
\left\| \mathrm{Ln}\left(\frac{x+v}{\mathcal{N}(x+v)^{1/d}}\right) - \mathrm{Ln}\left(\frac{x}{\mathcal{N}(x)^{1/d}}\right) \right\| &= \left\| \mathrm{Ln}\left(1+\frac{v}{x}\right) - \frac{1}{d}\ln\left(\frac{\mathcal{N}(x+v)}{\mathcal{N}(x)}\right)\cdot \mathbf{1} \right\| \\
&\leq \left\| \mathrm{Ln}\left(1+\frac{v}{x}\right) \right\| + \frac{1}{\sqrt{d}} \cdot \left| \ln\left(\frac{\mathcal{N}(x+v)}{\mathcal{N}(x)}\right) \right| \\
&\leq \sqrt{d} \cdot \frac{\|v/x\|_\infty}{1-\|v/x\|_\infty} + \sqrt{d} \cdot \frac{1/f}{1-1/f} \\
&\leq \frac{2\sqrt{d}}{f-1},
\end{aligned}$$

where we used the fact that

$$|\ln(1+y)| = \max\left(\ln(1+y), \ln\left(1+\frac{-y}{1+y}\right)\right) \leq \frac{|y|}{1-|y|},$$

for any $y \in (-1,1)$. This implies that

$$\left\| \mathrm{Ln}\left(\frac{x+v}{\mathcal{N}(x+v)^{1/d}}\right) \right\| \leq \ln(\eta) + \frac{2\sqrt{d}}{f-1} = \ln(\eta').$$

We conclude that $x+v$ belongs to $\mathcal{B}^{\eta'}_{A',B'}$. □

We are now ready to prove that the first condition of Lemma 3.2 is satisfied. To count the number of points of the ideal lattice I that belong to $\mathcal{B}^{\eta}_{A,B}$, we tile the space with shifts of a fundamental domain of the lattice (concretely, the Voronoi cell for the ℓ_∞ norm). Using Lemma 4.7, we show that the union of Voronoi cells corresponding to elements of $I \cap \mathcal{B}^{\eta}_{A,B}$ contains a smaller version $\mathcal{B}^{\eta_0}_{A_0,B_0}$ of the set, and is contained in a larger version $\mathcal{B}^{\eta_1}_{A_1,B_1}$. By carefully choosing parameters, we can ensure that the ratio of volumes of these two sets is bounded from above by a constant. In the lemma statement, note that C' is independent of the ideal I, but may depend on the other parameters, such as A, B, η and K. This proof is an adaptation of [Boe22, Le. 6.13] with $\mathcal{B}^{\eta}_{A,B}$ instead of the ℓ_∞ ball.

Lemma 4.8. *Let A, B, η verifying $A^{1/d} \geq \eta \cdot d^3 \cdot \Delta_K^{3/(2d)}$, $B/A \geq 4$ and $\eta \geq \mathrm{e}$. There exists $C' > 0$ such that for any replete ideal $I \in \mathrm{idLat}^0$, we have*

$$\left| I \bigcap \mathcal{B}^{\eta}_{A,B} \right| \in C' \cdot [1, 340].$$

Proof. Let I be a norm-1 ideal, and let $V_\infty(I)$ be its ℓ_∞-norm Voronoi cell, i.e., $V_\infty(I) = \{y \in K_\mathbb{R} : \forall x \in I \setminus \{0\}, \|y+x\|_\infty \geq \|y\|_\infty\}$. We let $\mu_\infty(I)$ denote the (ℓ_∞-norm) radius of $V_\infty(I)$. By (1), we have that $\mu_\infty(I) \leq d \cdot \Delta_K^{3/(2d)}$. As

a consequence, Lemma 4.7 instantiated with $f = d^2$ gives that since $A^{1/d} \geq \eta \cdot d^3 \cdot \Delta_K^{3/(2d)}$

$$\mathcal{B}^{\eta}_{A,B} + V_\infty(I) \subset \mathcal{B}^{\eta_1}_{A_1,B_1},$$

with $A_1 = A \cdot (1 - 1/d^2)^d$, $B_1 = B \cdot (1 + 1/d^2)^d$ and $\eta_1 = \eta \cdot \exp(2\sqrt{d}/(d^2 - 1))$. Recall that we assumed that $d \geq 2$ in all the article, so that we have $f > 1$ as needed for Lemma 4.7.

Let $A_0 = A \cdot (1 - 1/d^2)^{-d}$, $B_0 = B \cdot (1 + 1/d^2)^{-d}$ and $\eta_0 = \eta \cdot \exp(-2\sqrt{d}/(d^2 - 1))$. From the lower bound on η (and $d \geq 2$), one can check that $\eta_0 > 1$. Moreover, we have that $B_0/A_0 \geq 1/3 \cdot B/A \geq 4/3$ and hence that $B_0 > A_0$. Finally, from $A_0 \geq A$ and $\eta_0 \leq \eta$, we obtain that $A_0^{1/d} \geq \eta_0 \cdot f \cdot \mu_\infty(I)$ with $f = d^2$. This implies that we can apply Lemma 4.7 again on A_0, B_0, η_0 and $f = d^2$ and we obtain:

$$\mathcal{B}^{\eta_0}_{A_0,B_0} + V_\infty(I) \subset \mathcal{B}^{\eta}_{A,B}.$$

Note that for any $x \in \mathcal{B}^{\eta_0}_{A_0,B_0}$, there exists some (not necessarily unique) $\ell_x \in I$ such that $x - \ell_x \in V_\infty(I)$. This implies that $\ell_x \in (\mathcal{B}^{\eta_0}_{A_0,B_0} + V_\infty(I)) \cap I$. Therefore, we have

$$\mathcal{B}^{\eta_0}_{A_0,B_0} \subseteq \bigcup_{\ell \in (\mathcal{B}^{\eta_0}_{A_0,B_0} + V_\infty(I)) \cap I} \ell + V_\infty(I) \subseteq \bigcup_{\ell \in \mathcal{B}^{\eta}_{A,B} \cap I} \ell + V_\infty(I).$$

The above union is made of sets that are disjoints except for volume-0 intersections, so we have

$$\begin{aligned} \mathrm{Vol}(\mathcal{B}^{\eta_0}_{A_0,B_0}) \leq \mathrm{Vol}\Big(\bigcup_{\ell \in \mathcal{B}^{\eta}_{A,B} \cap I} \ell + V_\infty(I) \Big) &= \Big| \mathcal{B}^{\eta}_{A,B} \bigcap I \Big| \cdot \mathrm{Vol}(V_\infty(I)) \\ &= \Big| \mathcal{B}^{\eta}_{A,B} \bigcap I \Big| \cdot \sqrt{\Delta_K}. \end{aligned}$$

Similarly, we have:

$$\Big| \mathcal{B}^{\eta}_{A,B} \bigcap I \Big| \cdot \sqrt{\Delta_K} \leq \mathrm{Vol}(\mathcal{B}^{\eta_1}_{A_1,B_1}).$$

This gives us

$$\Big| \mathcal{B}^{\eta}_{A,B} \bigcap I \Big| \in C' \cdot \left[1, \frac{\mathrm{Vol}(\mathcal{B}^{\eta_1}_{A_1,B_1})}{\mathrm{Vol}(\mathcal{B}^{\eta_0}_{A_0,B_0})} \right],$$

where $C' = \mathrm{Vol}(\mathcal{B}^{\eta_0}_{A_0,B_0})/\sqrt{\Delta_K} > 0$. It remains to bound the right boundary of the interval. By using Lemma 4.3, we obtain that

$$\frac{\mathrm{Vol}(\mathcal{B}^{\eta_1}_{A_1,B_1})}{\mathrm{Vol}(\mathcal{B}^{\eta_0}_{A_0,B_0})} = \frac{(B_1 - A_1) \cdot (\ln \eta_1)^{d_\mathbb{R} + d_\mathbb{C} - 1}}{(B_0 - A_0) \cdot (\ln \eta_0)^{d_\mathbb{R} + d_\mathbb{C} - 1}} \leq \frac{B_1 - A_1}{B_0 - A_0} \cdot \left(\frac{\ln \eta_1}{\ln \eta_0} \right)^{d-1}.$$

Recall that we have already seen that $B_0/A_0 \geq 4/3$. This implies that

$$\frac{B_1 - A_1}{B_0 - A_0} \leq \frac{B_1}{B_0 - A_0} = \left(1 + \frac{1}{d^2}\right)^{2d} \cdot \frac{1}{1 - (A_0/B_0)} \leq \frac{5}{2} \cdot \frac{1}{1 - 3/4} = 10.$$

Using the fact that $\ln \eta \geq 1$, we also have:

$$\left(\frac{\ln \eta_1}{\ln \eta_0}\right)^{d-1} = \left(\frac{\ln \eta + 2\sqrt{d}/(d^2-1)}{\ln \eta - 2\sqrt{d}/(d^2-1)}\right)^{d-1} \leq \left(\frac{1 + 2\sqrt{d}/(d^2-1)}{1 - 2\sqrt{d}/(d^2-1)}\right)^{d-1} \leq 34.$$

This completes the proof. □

4.3 Sampling uniform ideal elements in $\mathcal{B}^{\eta}_{A,B}$

We now show how to uniformly sample in $I \cap \mathcal{B}^{\eta}_{A,B}$, where I is a norm-1 ideal. For this purpose, $\texttt{SampleUniform}^{\eta}_{A,B}$ (Algorithm 4.1) uniformly samples in a larger $\mathcal{B}^{\eta_1}_{A_1,B_1}$ (using Lemma 4.4) and deterministically round to I using Babai's nearest plane algorithm [Bab86, Th. 3.1]. The sample is kept if it belongs to $\mathcal{B}^{\eta}_{A,B}$.

Algorithm 4.1 $\texttt{SampleUniform}^{\eta}_{A,B}$

Input: $\mathbf{B}_I$ a basis of an ideal $I \in \text{idLat}^0$.
Output: $x \in I \bigcap \mathcal{B}^{\eta}_{A,B}$.

1: Let $A_1 = A \cdot (1 - 1/d^2)^d$, $B_1 = B \cdot (1 + 1/d^2)^d$ and $\eta_1 = \eta \cdot \exp(2\sqrt{d}/(d^2-1))$.
2: **repeat**
3: Sample $y \leftarrow \mathcal{U}(\mathcal{B}^{\eta'}_{A',B'})$.
4: Run Babai's nearest plane algorithm on $(\mathbf{B}_I, y)$; let $x \in I$ be the output.
5: **until** $x \in \mathcal{B}^{\eta}_{A,B}$.
6: Return x.

Lemma 4.9. *Let A, B, η with $B/A \geq 4$ and $\eta \geq$ e. Let $I \in \text{idLat}^0$ given by a basis $\mathbf{B}_I$ and $\delta = \|\mathbf{B}_I^\star\|$. Assume that $A^{1/d} \geq d^{2.5} \cdot \eta \cdot \delta$. Then $\texttt{SampleUniform}^{\eta}_{A,B}$ samples uniformly in $I \cap \mathcal{B}^{\eta}_{A,B}$ and its expected running time is polynomial in $\log B$, d and B/A.*

Proof. Let $\mathcal{P}(\mathbf{B}_I) = \mathbf{B}_I^\star \cdot (-1/2, 1/2]^d$ be the rounding cell of Babai's nearest plane algorithm. In order to prove that the output distribution is uniform, it suffices to prove that for any point $x \in I \cap \mathcal{B}^{\eta}_{A,B}$, we have $\mathcal{P}(\mathbf{B}_I) + x \subset \mathcal{B}^{\eta_1}_{A_1,B_1}$. The definition of the nearest-plane algorithm's rounding cell implies that the ℓ_∞ norm of vectors in $\mathcal{P}(\mathbf{B}_I)$ is at most $\sqrt{d}\delta$. The definitions of A_1, B_1, η_1 and Lemma 4.7 (with $f = d^2$) allow us to conclude.

The running time follows from Lemma 4.4 and from bounding the probability that after Step 4, we have $x \notin \mathcal{B}^{\eta}_{A,B}$. This occurs if $y \notin \cup_{x \in \mathcal{B}^{\eta}_{A,B} \cap I} (x + \mathcal{P}(\mathbf{B}_I))$. As in the proof of Lemma 4.8, we have that:

$$\mathcal{B}^{\eta_0}_{A_0,B_0} \subset \sum_{x \in \mathcal{B}^{\eta}_{A,B} \cap I} x + \mathcal{P}(\mathbf{B}_I),$$

where $A_0 = A \cdot (1+1/d^2)^d$, $B_0 = B \cdot (1-1/d^2)^d$ and $\eta_0 = \eta \cdot \exp(-2\sqrt{d}/(d^2-1))$. The probability of exiting the loop is then bounded from below by

$$\frac{\mathrm{Vol}(\mathcal{B}^{\eta_0}_{A_0,B_0})}{\mathrm{Vol}(\mathcal{B}^{\eta_1}_{A_1,B_1})} \geq \frac{B_0 - A_0}{B_1 - A_1} \cdot \frac{(\ln \eta_0)^{d-1}}{(\ln \eta_1)^{d-1}} \geq \Omega(1),$$

where the inequalities are as in the proof of Lemma 4.8. □

5 Wrapping up

We combine Theorems 3.4 and 4.2 to obtain the main result from this work. To simplify the statement, we instantiate the integral ideal-factoring oracle with a quantum polynomial-time algorithm, and use the Extended Riemann Hypothesis. The latter allows us to bound $|\mathcal{P}_{0,A}|/|\mathcal{P}_{0,4A}|$ by a constant that is < 1 when $A \geq (\log \Delta_K)^{\Omega(1)}$ and $A/|\mathcal{P}_{A,4A}|$ by $O(\ln A)$ (see [BS96, Th. 8.7.4]).

Theorem 5.1 (ERH). *There exists $C_K = (d\delta_K \Delta_K^{1/d})^{O(1)}$ such that the following holds. Let $\mathcal{W}$ be a finite set of fractional ideals. Let $\gamma_{\mathcal{W}}, \gamma_{\mathcal{I}} \geq 1$ and A with $A^{1/d} \geq C_K \cdot \gamma_{\mathcal{W}}$. Let $\mathcal{O}_{\mathcal{W}}$ an oracle for $\mathcal{W}$-avg-id-HSVP$_{\gamma_{\mathcal{W}}}$ with success probability $\varepsilon_{\mathcal{W}}$ and $\mathcal{O}_{\mathcal{I}}$ an oracle for $\mathcal{I}_{A,4A}$-avg-id-HSVP$_{\gamma_{\mathcal{I}}}$ with success probability $\varepsilon_{\mathcal{I}}$.*

There exists a quantum algorithm making one call to $\mathcal{O}_{\mathcal{W}}$ and one call to $\mathcal{O}_{\mathcal{I}}$ whose running time is polynomial in $\log A$, $\log \Delta_K$ and the size of its input, such that the following holds. Given as input $I \sim \mathcal{U}(\mathcal{W})$, it outputs $x \in I^{-1} \setminus \{0\}$ with probability $\geq \varepsilon_{\mathcal{I}} \cdot (\varepsilon_{\mathcal{W}}/\Theta(1) - 2^{-\Omega(d)})$ such that

$$\|x\| \leq \gamma' \cdot \mathrm{Vol}(I^{-1})^{1/d} \quad \text{with} \quad \gamma' = 232 \cdot d \cdot \Delta_K^{1/d} \cdot \gamma_{\mathcal{I}}.$$

Proof. The algorithm is $\texttt{InverseToIntegral}^{\mathcal{W}}$ (Algorithm 3.3) instantiated with the set $\mathcal{B}^{\eta}_{A,B}$ with $B = 4A$ and $\eta = \mathrm{e}$.

Note that at Step 3 of $\texttt{InverseToIntegral}^{\mathcal{W}}$, we have $\|\mathbf{B}_I\| \leq \delta_K \cdot \|v_I\|$ (by Lemma 2.8). By definition of $\mathcal{O}_{\mathcal{W}}$, this implies that $\|\mathbf{B}_I\| \leq \delta_K \cdot \gamma_{\mathcal{W}} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(I)^{1/d}$. $\texttt{InverseToIntegral}^{\mathcal{W}}$ then calls `RandomizeIdeal` (Algorithm 3.2), which at its Step 6 computes a basis $\mathbf{B}_J$ of an ideal J that was showed in the proof of Lemma 3.3 to satisfy:

$$\begin{aligned}\|\mathbf{B}_J\| &\leq 85 \cdot d^2 \cdot \delta_K \cdot \Delta_K^{1/d} \cdot \mathcal{N}(I)^{-1/d} \cdot \|\mathbf{B}_I\| \\ &\leq 85 \cdot d^2 \cdot \delta_K^2 \cdot \Delta_K^{3/(2d)} \cdot \gamma_{\mathcal{W}} =: \delta.\end{aligned}$$

The result follows from Theorems 3.4 and 4.2, using δ as above. □

As a corollary, we obtain a quantum reduction from $\mathcal{I}^{-1}_{A,4A}$-avg-id-HSVP$_{\gamma'}$ to $\mathcal{I}_{A,4A}$-avg-id-HSVP$_{\gamma}$ and from $\mathcal{P}^{-1}_{A,4A}$-avg-id-HSVP$_{\gamma'}$ to $\mathcal{P}_{A,4A}$-avg-id-HSVP$_{\gamma}$ if $A^{1/d} \geq (d\delta_K \Delta_K^{1/d})^{\Omega(1)} \cdot \gamma$ and $\gamma' = O(d\Delta_K^{1/d}) \cdot \gamma$. Note that in the case of prime ideals, the success probability decreases with $\tilde{\rho}_A$ (the inverse of the proportion of prime ideals among all ideals of norm $\leq A$), which may or may not be small depending on the choice of the field K. This dependency arises from hoping that a uniform integral ideal is prime.

Corollary 5.2 (ERH). *There exists* $C_K = (d\delta_K \Delta_K^{1/d})^{O(1)}$ *such that the following holds. Let* $\gamma \geq 1$ *and* A *with* $A^{1/d} \geq C_K \cdot \gamma$. *Let* $\mathcal{O}$ *an oracle for* $\mathcal{I}_{A,4A}$-avg-id-HSVP$_\gamma$ *with success probability* $\varepsilon \geq 2^{-\Omega(d)}$.

There exists a quantum algorithm making two calls to $\mathcal{O}$ *whose running time is polynomial in* $\log A$, $\log \Delta_K$ *and the size of its input, such that, given as input* $\mathfrak{a} \sim \mathcal{U}(\mathcal{I}_{A,4A})$, *it outputs* $x \in \mathfrak{a}^{-1} \setminus \{0\}$ *with probability* $\Omega(\varepsilon^2)$ *with*

$$\|x\| \leq \gamma' \cdot \mathrm{Vol}(\mathfrak{a}^{-1})^{1/d} \quad \textit{with} \quad \gamma' = 232 \cdot d \cdot \Delta_K^{1/d} \cdot \gamma.$$

Corollary 5.3 (ERH). *There exists* $C_K = (d\delta_K \Delta_K^{1/d})^{O(1)}$ *such that the following holds. Let* $\gamma \geq 1$ *and* A *with* $A^{1/d} \geq C_K \cdot \gamma$. *Let* $\mathcal{O}$ *an oracle for* $\mathcal{P}_{A,4A}$-avg-id-HSVP$_\gamma$ *with success probability* $\varepsilon \geq 2^{-\Omega(d)}$.

There exists a quantum algorithm making two calls to $\mathcal{O}$ *whose running time is polynomial in* $\log A$, $\log \Delta_K$ *and the size of its input, such that, given as input* $\mathfrak{p} \sim \mathcal{U}(\mathcal{P}_{A,4A})$, *it outputs* $x \in \mathfrak{p}^{-1} \setminus \{0\}$ *with probability* $\Omega(\varepsilon^2/\tilde{\rho}_A)$ *with*

$$\|x\| \leq \gamma' \cdot \mathrm{Vol}(\mathfrak{p}^{-1})^{1/d} \quad \textit{with} \quad \gamma' = 232 \cdot d \cdot \Delta_K^{1/d} \cdot \gamma.$$

Combining Corollary 5.3 with Theorem 2.10, we obtain a quantum worst-case to average-case reduction for ideal-HSVP, where the average-case distribution is the uniform distribution over prime ideals with norm in some interval $[A, 4A]$.

Corollary 5.4 (ERH). *Let* $\gamma \geq 1$. *There exists some* $\gamma' = \gamma \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ *and* $A = \gamma^d \cdot \mathrm{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ *such that*

$$\text{id-HSVP}_{\gamma'} \quad \textit{reduces to} \quad \mathcal{P}_{A,4A}\text{-avg-id-HSVP}_\gamma.$$

The reduction is quantum and runs in expected time polynomial in its input size, $\log \Delta_K$, $1/\tilde{\rho}_A$ *and* $1/\varepsilon$, *where* ε *is the success probability of the oracle solving* $\mathcal{P}_{A,4A}$-avg-id-HSVP$_\gamma$.

Proof. We assume without loss of generality that $\gamma \leq 2^d$, since otherwise we can solve id-HSVP$_{\gamma'}$ in polynomial time using the LLL algorithm, for $\gamma' = \gamma \cdot \sqrt{d}$. We also assume that the success probability ε of the oracle solving $\mathcal{P}_{A,4A}$-avg-id-HSVP$_\gamma$ is $\geq 2^{-\Omega(d)}$, since otherwise one can run an exact SVP solver in time $1/\varepsilon$.

Let $C'_{1,K}$ be the max of the $C_{1,K}$ from Theorem 2.10 and the C_K from Corollary 5.3. Then $C'_{1,K} = \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ since both quantities are. Let $A = \left(C'_{1,K} \cdot (232d \cdot \Delta_K^{1/d}) \cdot \gamma\right)^d$. One can check that $A = \gamma^d \cdot \mathrm{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ as desired. Let also $\gamma' = A^{1/d} \cdot C_{2,K} = \gamma \cdot (232d \cdot \Delta_K^{1/d}) \cdot C'_{1,K} \cdot C_{2,K}$, where $C_{2,K}$ is as in Theorem 2.10. Similarly, one can check that $\gamma' = \gamma \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ as desired. Finally, let $\gamma_{\mathrm{avg}} = 232d \cdot \Delta_K^{1/d} \cdot \gamma$.

Note that A, γ and ε satisfy the conditions from Corollary 5.3. So there is a quantum reduction

$$\text{from} \quad \mathcal{P}^{-1}_{A,4A}\text{-avg-id-HSVP}_{\gamma_{\mathrm{avg}}} \quad \text{to} \quad \mathcal{P}_{A,4A}\text{-avg-id-HSVP}_\gamma,$$

which succeeds with probability $\delta = \Omega(\varepsilon^2/\tilde{\rho}_A)$ and runs in time polynomial in $\log \Delta_K$. Now, observe that γ_{avg}, A and γ' satisfy the conditions from Theorem 2.10, so there is a quantum reduction

$$\text{from} \quad \text{id-HSVP}_{\gamma'} \quad \text{to} \quad \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}},$$

which runs in expected time polynomial in its input size, $\log \Delta_K$ and $1/\delta$. Combining both reductions and instantiating with the lower bound on δ completes the proof. □

6 NTRU with polynomial modulus

The main result of this section is Corollary 6.3. It gives a distribution over NTRU instances with small modulus q that is hard on average, under the worst-case id-HSVP hardness assumption.

Definition 6.1 ([PS21, Def. 3.1 and 3.4]). *Let $\gamma \geq \gamma' \geq 1$ be real numbers. Let $q \geq 2$ be an integer.*

- *A (γ, q)-NTRU instance is an element $h \in \mathcal{O}_K/q\mathcal{O}_K$ such that there exists $(f, g) \in \mathcal{O}_K \setminus \{(0,0)\}$ verifying $f = h \cdot g \bmod q$ and $\|f\|, \|g\| \leq \sqrt{q}/\gamma$.*
- *The (γ, γ', q)-NTRU problem asks, given a (γ, q)-NTRU instance h, to find $(f, g) \in \mathcal{O}_K \setminus \{(0,0)\}$ verifying $f = h \cdot g \bmod q$ and $\|f\|, \|g\| \leq \sqrt{q}/\gamma'$.*
- *Let D be a distribution over (γ, q)-NTRU instances. The (D, γ, γ', q)-NTRU problem asks to solve (γ, γ', q)-NTRU for an instance sampled from D, with non-negligible probability (over the choice of the instance and the internal randomness of the algorithm).*

Note that in this work we are only interested in the vector version of NTRU from [PS21]. We let `IdealToNTRU` denote [PS21, Alg. 4.1]. It takes as input a basis of an integral ideal $\mathfrak{a}$ and a modulus q and outputs an instance of (γ, q)-NTRU whose solution is related to a short non-zero vector of $\mathfrak{a}$. The following result is a consequence of [PS21, Le. 4.3], whose proof is very similar to [PS21, Th. 4.1]. We provide a proof for the sake of completeness.

Theorem 6.2 (Adapted from [PS21, Th. 4.1]). *Let $\gamma \geq \gamma' \geq 1$ be real numbers, $q \geq 2$ be an integer, and*

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

Let $\mathfrak{a}$ be an integral ideal of norm in $[N, 2^{d+2} \cdot N]$ and $h = \texttt{IdealToNTRU}(\mathfrak{a}, q)$. Then h is a (γ, q)-NTRU instance. If (f, g) is a solution to (γ, γ', q)-NTRU on instance h, then g is a solution to γ_{HSVP}-id-HSVP for instance $\mathfrak{a}$, where $\gamma_{\text{HSVP}} = \gamma/\gamma' \cdot 4d^{1.5} \cdot \delta_K$.

Further, `IdealToNTRU` runs in time polynomial in its input size and in $\log \Delta_K$.

Note that the statement is void if $2^{d+2} \cdot N < 1$ (no integral ideal has norm in $(0,1)$): an extra parameter constraint is implicily required for it to be meaningful.

Proof. The running time of `IdealToNTRU` is stated in [PS21, Le. 4.3].

By [PS21, Le. 4.3], there exists $(f,g) \in {\mathcal{O}_K}^2 \setminus \{(0,0)\}$ such that $g \cdot h = f \bmod q$ and $\|f\|, \|g\| \leq d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d}$. (Note that δ_K in the present work is an upper bound on the quantity δ_K from [PS21].) Using $\mathcal{N}(\mathfrak{a}) \leq 2^{d+2} \cdot N$ and the definition of N, this gives that h is a (γ, q)-NTRU instance.

Assume now that $(f,g) \in \mathcal{O}_K \setminus \{(0,0)\}$ is a solution to (γ', γ, q)-NTRU for instance h. Then we have

$$\begin{aligned}\|f\|, \|g\| \leq \frac{\sqrt{q}}{\gamma'} &\leq \frac{q}{d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot (2^{d+2} \cdot N)^{1/d}} \\ &\leq \frac{q}{d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d}},\end{aligned}$$

where the second inequality comes the definition of N, and the third one comes from the assumption $\mathcal{N}(\mathfrak{a}) \leq 2^{d+2} \cdot N$. By [PS21, Le. 4.3], we obtain that $g \in \mathfrak{a} \setminus \{0\}$. Finally, the fact that g is a solution to γ_{HSVP} follows from

$$\begin{aligned}\|g\| \leq \frac{\sqrt{q}}{\gamma'} &= \frac{2^{(d+2)/d} \cdot N^{1/d} \cdot \gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}}{\gamma'} \\ &\leq \frac{4 \cdot \gamma \cdot d^{1.5} \cdot \delta_K}{\gamma'} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d},\end{aligned}$$

where the last inequality follows from the inequalities $\mathcal{N}(\mathfrak{a}) \geq N$ and $d \geq 2$. □

For $A, q \geq 2$, we define $D^{A,q}_{\mathrm{NTRU}} = \mathtt{IdealToNTRU}(\mathcal{U}(\mathcal{P}_{A,4A}), q)$. Theorem 6.2 implies a polynomial-time reduction from $\mathcal{I}_{A,4A}$-avg-id-HSVP to $(D^{A,q}_{\mathrm{NTRU}}, \gamma, \gamma', q)$-NTRU for well chosen γ, γ', A and q. Combining Corollary 5.4 and Theorem 6.2 give the following result.

Corollary 6.3 (ERH). *Let $\gamma \geq \gamma' \geq 1$. There exists an integer $q = (\gamma^4/\gamma'^2) \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$, and real numbers $\gamma_{\mathrm{HSVP}} = (\gamma/\gamma') \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ and $A = (\gamma/\gamma')^d \cdot \mathrm{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ such that*

$$\text{id-HSVP}_{\gamma_{\mathrm{HSVP}}} \quad \textit{reduces to} \quad (D^{A,q}_{\mathrm{NTRU}}, \gamma, \gamma', q)\text{-NTRU}.$$

The reduction is quantum and runs in expected time polynomial in its input size, $\log q$, $\log \Delta_K$, $1/\tilde{\rho}_A$ and $1/\varepsilon$, where ε is the success probability of the oracle solving $(D^{A,q}_{\mathrm{NTRU}}, \gamma, \gamma', q)$-NTRU.

Proof. Without loss of generality, we can assume that $\gamma/\gamma' \leq 2^d$, since otherwise we have a polynomial time algorithm solving id-HSVP$_{\gamma_{\mathrm{HSVP}}}$ for $\gamma_{\mathrm{HSVP}} = \gamma/\gamma'$. Let $\Gamma = (\gamma/\gamma') \cdot 4d^{1.5} \cdot \delta_K$. Let $A = \Gamma^d \cdot \mathrm{poly}(\Delta_K, (\log \Delta_K)^d, \delta_K^d)$ be as in Corollary 5.4, with "$\gamma = \Gamma$". Similarly, let $\gamma_{\mathrm{HSVP}} = \Gamma \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ be the quantity γ' from Corollary 5.4, with "$\gamma = \Gamma$". Finally, let $X = \gamma \cdot 2 \cdot (4A)^{1/d} \cdot$

$d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}$ and $q = \lfloor X^2 \rfloor$. Note that $q \geq X^2/4$. Note that $\gamma_{\mathrm{HSVP}} = (\gamma/\gamma') \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$ and that $q = (\gamma^4/\gamma'^2) \cdot \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$.

Let $N = \frac{1}{2^{d+2}} \cdot \left(\frac{\sqrt{q}}{\gamma \cdot d^{1.5} \cdot \delta_K \cdot \Delta_K^{1/(2d)}}\right)^d$ be as in Theorem 6.2. Using the fact that $X/2 \leq \sqrt{q} \leq X$ and the definition of X, we have that $[A, 4A] \subseteq [N, 2^{d+2} \cdot N]$. Hence, the support of the distribution $\mathcal{U}(\mathcal{P}_{A,4A})$ is contained in the set of integral ideals with norm in $[N, 2^{d+2} \cdot N]$.

Recall that $D_{\mathrm{NTRU}}^{A,q}$ is the distribution $\texttt{IdealToNTRU}(\mathcal{U}(\mathcal{P}_{A,4A}), q)$. By Theorem 6.2, there is a reduction from $\mathcal{P}_{A,4A}$-avg-id-HSVP$_\Gamma$ to $(D_{\mathrm{NTRU}}^{A,q}, \gamma, \gamma', q)$-NTRU, which runs in time polynomial in $\log q$, $\log \Delta_K$ and $\log A = \mathrm{poly}(\log \Delta_K)$ (since $\gamma/\gamma' \leq 2^d$) and preserves the the success probability of the algorithm. Moreover, from Corollary 5.4, id-HSVP$_{\gamma_{\mathrm{HSVP}}}$ reduces to $\mathcal{P}_{A,4A}$-avg-id-HSVP$_\Gamma$, which is quantum and runs in expected time polynomial in its input size, $\log \Delta_K$, $1/\tilde{\rho}_A$ and $1/\varepsilon$. Combining both reductions gives the desired result. □

Note that the distribution $D_{\mathrm{NTRU}}^{A,q}$ can be sampled from along with a trapdoor by running `SampleWithTrap` with appropriate parameters (in order to generate an ideal from $\mathcal{U}(\mathcal{P}_{A,4A})$ together with a short non-zero vector in it), and then running the `IdealToNTRU` algorithm. This, however, requires an access to a factoring oracle (for the `SampleWithTrap` algorithm). Finding a classical algorithm to efficiently sample from $D_{\mathrm{NTRU}}^{A,q}$ with a trapdoor is an interesting open problem.

Acknowledgments. The authors thank Koen de Boer, Guillaume Hanrot, Aurel Page and Noah Stephens-Davidowitz for helpful discussions. Joël Felderhoff is funded by the Direction Générale de l'Armement (Pôle de Recherche CYBER). The authors were supported by the CHARM ANR-NSF grant (ANR-21-CE94-0003) and by the PEPR quantique France 2030 programme (ANR-22-PETQ-0008).

References

ABD16. M. Albrecht, S. Bai, and L. Ducas. A subfield lattice attack on overstretched NTRU assumptions. In *CRYPTO*, 2016.

Bab86. L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 1986.

Ban93. W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Math. Ann.*, 1993.

BDPW20. K. de Boer, L. Ducas, A. Pellet-Mary, and B. Wesolowski. Random self-reducibility of Ideal-SVP via Arakelov random walks. In *CRYPTO*, 2020.

BEP22. K. Boudgoust, Gachon E., and A. Pellet-Mary. Some easy instances of Ideal-SVP and implications on the partial Vandermonde knapsack problem. In *CRYPTO*, 2022.

BL94. J. A. Buchmann and H. W. Lenstra. Computing maximal orders and factoring over $\mathbb{Z}_p$. *Preprint*, 1994.

BLP+13. Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, 2013.

Boe22. K. de Boer. *Random Walks on Arakelov Class Groups*. PhD thesis, Leiden University, 2022. Available on request from the author.
BS96. E. Bach and J. O. Shallit. *Algorithmic Number Theory: Efficient Algorithms*. MIT Press, 1996.
BST+20. M. Bhargava, A. Shankar, T. Taniguchi, F. Thorne, J. Tsimerman, and Y. Zhao. Bounds on 2-torsion in class groups of number fields and integral points on elliptic curves. *Journal of the AMS*, 2020.
CDPR16. R. Cramer, L. Ducas, C. Peikert, and O. Regev. Recovering short generators of principal ideals in cyclotomic rings. In *EUROCRYPT 2016*, 2016.
CDW17. R. Cramer, L. Ducas, and B. Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In *EUROCRYPT*, 2017.
CJL16. J. H. Cheon, J. Jeong, and C. Lee. An algorithm for NTRU problems and cryptanalysis of the GGH multilinear map without a low-level encoding of zero. *LMS Journal of Computation and Mathematics*, 2016.
Coh96. H. Cohen. *A course in computational algebraic number theory*. Springer, 1996.
FPS22. J. Felderhoff, A. Pellet-Mary, and D. Stehlé. On module unique-SVP and NTRU. In *ASIACRYPT*, 2022.
Gen09a. C. Gentry. *A Fully Homomorphic Encryption Scheme*. PhD thesis, Stanford University, 2009.
Gen09b. C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, 2009.
Gen10. C. Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *CRYPTO*, 2010.
GPV08. C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
KF17. P. Kirchner and P.-A. Fouque. Revisiting lattice attacks on overstretched NTRU parameters. In *EUROCRYPT*, 2017.
LM06. V. Lyubashevsky and D. Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP*, 2006.
LPR10. V. Lyubashevsky, C. Peikert, and O. Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, 2010.
MG02. D. Micciancio and S. Goldwasser. *Complexity of lattice problems: a cryptographic perspective*. Springer, 2002.
Mic02. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *FOCS*, 2002.
MR07. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007.
Neu13. J. Neukirch. *Algebraic number theory*. Springer, 2013.
PHS19. A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-SVP in ideal lattices with pre-processing. In *EUROCRYPT*, 2019.
PML21. C. Porter, A. Mendelsohn, and C. Ling. Subfield algorithms for Ideal- and Module-SVP based on the decomposition group. *IACR Cryptol. ePrint Arch.*, 2021.
PR06. C. Peikert and A. Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.
PRS17. C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *STOC*, 2017.
PS21. A. Pellet-Mary and D. Stehlé. On the hardness of the NTRU problem. In *ASIACRYPT*, 2021.

PXWC21. Y. Pan, J. Xu, N. Wadleigh, and Q. Cheng. On the ideal shortest vector problem over random rational primes. In *EUROCRYPT*, 2021.

Reg05. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, 2005.

Sho94. P. W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *FOCS*, 1994.

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

Web08. H. Weber. Lehrbuch der algebra, vol. ii. *Vieweg und Sohn, Braunschweig*, 1908.

A Additional preliminaries

For $s > 0$ and $\mathbf{x} \in \mathbb{R}^n$ we define $\rho_s(\mathbf{x}) = \exp\left(-\pi \cdot \|\mathbf{x}\|^2/s^2\right)$. Given an n-dimensional lattice L, a vector $\mathbf{u} \in \mathbb{R}^n$ and a parameter $s > 0$, we define the discrete Gaussian distribution $D_{L,s,\mathbf{u}}$ over L with center parameter $\mathbf{u}$ and standard deviation parameter s by $D_{L,s,\mathbf{u}}(\mathbf{x}) := \rho_s(\mathbf{x}-\mathbf{u})/\rho_s(L-\mathbf{u})$ for all $\mathbf{x} \in L$.

For any $\varepsilon > 0$ and n-dimensional lattice L, the smoothing parameter $\eta_\varepsilon(L)$ is defined as the smallest $s > 0$ such that $\rho_{1/s}(L^*\backslash\{0\}) \leq \varepsilon$, where $L^* = \{x \in \mathrm{span}_{\mathbb{R}}(L) \mid \langle x, \ell\rangle \in \mathbb{Z} \text{ for all } \ell \in L\}$. In the case of an ideal lattice I, we have (see [PRS17, Lemma 6.9]), for any $\varepsilon \in (0,1)$:

$$\eta_\varepsilon(I) \leq \Delta_K^{1/d} \cdot \mathcal{N}(I)^{1/d} \cdot \max\left(1, \sqrt{\frac{\ln(1/\varepsilon)}{d}}\right). \tag{3}$$

For any basis $\mathbf{B}$ of an n-dimensional lattice L, we have the following upper bound (obtained by instantiating [GPV08, Le. 3.1] with $\varepsilon = 1/2$ and observing that $\ln(6n/\pi) \leq n$ for all $n \geq 1$).

$$\eta_{1/2}(L) \leq \sqrt{n} \cdot \max_i \|\mathbf{b}_i^*\|. \tag{4}$$

We will use the following results on lattice Gaussians.

Lemma A.1 (Proof of [MR07, Le. 4.4]). *Let L be a rank n lattice, $\mathbf{u} \in \mathrm{span}(L)$ and $s \geq \eta_\varepsilon(L)$ for some $\varepsilon > 0$. Then it holds that*

$$\rho_s(L+\mathbf{u}) \in [1-\varepsilon, 1+\varepsilon] \cdot \frac{s^n}{\mathrm{vol}(L)}.$$

Corollary A.2. *Let I be a fractional ideal, $\mathfrak{a}$ be an integral ideal and $\varepsilon \in (0,1)$. Let $\mathbf{u} \in \mathrm{span}(I)$ and $s \geq \Delta_K^{1/d} \cdot \mathcal{N}(\mathfrak{a} \cdot I)^{1/d} \cdot \sqrt{\ln(3/\varepsilon)}$. Then*

$$D_{I,s,\mathbf{u}}(\mathfrak{a} \cdot I) \in [1-\varepsilon, 1+\varepsilon] \cdot \mathcal{N}(\mathfrak{a})^{-1}.$$

Proof. Note that since $\mathfrak{a}$ is integral, then $\mathfrak{a} \cdot I$ is a sub-lattice of I and $D_{I,s,\mathbf{u}}(\mathfrak{a} \cdot I)$ is well-defined. By (3) and the lower bound on s, we have $s \geq \eta_{\varepsilon'}(\mathfrak{a} \cdot I) \geq \eta_{\varepsilon'}(I)$, for $\varepsilon' = \varepsilon/3$. We can thus apply Lemma A.1 to both lattices I and $\mathfrak{a} \cdot I$. We obtain $D_{I,s,\mathbf{u}}(\mathfrak{a} \cdot I) = \rho_s(a \cdot I - \mathbf{u})/\rho_s(I-\mathbf{u}) \in [\frac{1-\varepsilon'}{1+\varepsilon'}, \frac{1+\varepsilon'}{1-\varepsilon'}] \cdot \mathrm{vol}(I)/\mathrm{vol}(\mathfrak{a} \cdot I)$. We conclude using the fact that $\mathrm{vol}(\mathfrak{a} \cdot I)/\mathrm{vol}(I) = \mathcal{N}(\mathfrak{a})$ and the choice of ε'. □

Lemma A.3 ([Reg05, Claim 3.8]). *For any $\varepsilon > 0$, lattice L, center $\mathbf{u} \in \mathrm{span}_{\mathbb{R}}(L)$ and parameter $s \geq \eta_\varepsilon(L)$, it holds that*

$$\rho_s(L+\mathbf{u}) \in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1\right] \cdot \rho_s(L).$$

Lemma A.4 ([Ban93, Le. 1.5]). *For any $c > 1/\sqrt{2\pi}$, any n-dimensional lattice L and any $\mathbf{u} \in \mathrm{span}(L)$, we have $\rho_1((L-\mathbf{u})\backslash c\sqrt{n}\mathcal{B}) \leq 2C^n\rho_1(L)$, where $\mathcal{B}$ denotes the Euclidean ball of radius 1 and $C = c\sqrt{2\pi\mathrm{e}} \cdot \mathrm{e}^{-\pi c^2} < 1$.*

Corollary A.5. *Let L be a lattice of rank n, $\mathbf{B}$ be any basis of L, $\mathbf{u} \in \mathrm{span}(L)$ and $s \geq \sqrt{n} \cdot \max_i \|\mathbf{b}_i^*\|$. For any $\varepsilon \in (0,1]$, it holds that $\Pr_{\mathbf{x} \leftarrow D_{L,s,\mathbf{u}}}(\|\mathbf{x}-\mathbf{u}\| \geq s \cdot \sqrt{\ln(1/\varepsilon)+4n}) \leq \varepsilon$.*

Proof. Without loss of generality, we can scale everything so that $s = 1$. Let us define $c := \sqrt{(1/n) \cdot \ln(1/\varepsilon) + 4}$. Then, we have

$$\Pr_{\mathbf{x} \leftarrow D_{L,1,\mathbf{u}}}\left(\|\mathbf{x}-\mathbf{u}\| \geq \sqrt{\ln(1/\varepsilon)+4n}\right) = \frac{\rho_1\left((L-\mathbf{u}) \setminus c\sqrt{n}\mathcal{B}\right)}{\rho_1(L-\mathbf{u})}.$$

Since $c \geq \sqrt{4} > 1/\sqrt{2\pi}$, we can apply Lemma A.4 to bound the numerator from above. In order to simplify the computations, we use the fact that $c \cdot \mathrm{e}^{-\pi c^2} \leq \mathrm{e}^{-c^2}$ for all $c > 1/\sqrt{2\pi}$. Then we see that $6 \cdot C^n \leq 6^n \cdot C^n \leq \mathrm{e}^{-\ln(1/\varepsilon)+(\ln(\sqrt{2\pi \mathrm{e}})+\ln(6)-4)n} \leq \mathrm{e}^{-\ln(1/\varepsilon)} = \varepsilon$. Using Lemma A.4, we hence obtain the bound

$$\rho_1\left((L-\mathbf{u}) \setminus c\sqrt{n}\mathcal{B}\right) \leq \varepsilon/3 \cdot \rho_1(L).$$

Let us now bound the quantity $\rho_1(L-\mathbf{u})$ from below. Using Lemma A.3 with $\varepsilon = 1/2$ (observe that $s \geq \eta_{1/2}(L)$ by Equation (4)), we see that $\rho_1(L-\mathbf{u}) \geq 1/3 \cdot \rho_1(L)$. Combining both inequalities provides the desired result. □

Lemma A.6 ([BLP+13, Le. 2.3]). *There is a probabilistic polynomial-time algorithm that, given a basis $\mathbf{B}$ of an n-dimensional lattice L, a center $\mathbf{u} \in \mathrm{span}_{\mathbb{R}}(L)$, and a parameter $s \geq \sqrt{\ln(2n+4)/\pi} \cdot \|\mathbf{B}^*\|$, outputs a sample distributed according to $D_{L,s,\mathbf{u}}$.*

The following lemma is adapted from [GPV08, Th. 4.1] and [PS21, Le. 2.2]. We will notably be interested in values of ε that are $2^{-\omega(n)}$, which is not captured in the typical variants of this statement. For completeness, a proof is provided.

Lemma A.7. *There exists a probabilistic polynomial time algorithm that takes as input a basis $\mathbf{B}$ of an n-dimensional lattice L, an error bound $\varepsilon \in (0, 1/2]$, a parameter $s \geq \sqrt{n} \cdot \|\mathbf{B}^*\|$ and a center $\mathbf{u} \in \mathrm{span}(L)$ and outputs a sample from a distribution $\widetilde{D}_{\mathbf{B},s,\mathbf{u}}$ such that*

- $\mathrm{SD}(D_{L,s,\mathbf{u}}, \widetilde{D}_{\mathbf{B},s,\mathbf{u}}) \leq \varepsilon$;
- *for all $\mathbf{v} \leftarrow \widetilde{D}_{B,s,\mathbf{u}}$, it holds that $\|\mathbf{v}-\mathbf{u}\| < s \cdot \sqrt{\ln(1/\varepsilon)+4n}$.*

Proof. The algorithm from Lemma A.7 is obtained by running the algorithm from Lemma A.6 until the output $\mathbf{v}$ satisfies $\|\mathbf{v}-\mathbf{u}\| < s \cdot \sqrt{\ln(1/\varepsilon)+4n}$. From Corollary A.5, this event happens with probability at least $1-\varepsilon \geq 1/2$, hence the algorithm resamples at most twice on average, and the output distribution is within statistical distance $\leq \varepsilon$ from $D_{L,s,\mathbf{u}}$ (the distribution before rejection). Finally, note that $\sqrt{\ln(2n+4)/\pi} \leq \sqrt{n}$ for all $n \geq 1$, hence we can indeed apply Lemma A.6, and we conclude that the expected run time of the algorithm is polynomial. □

B Proof of Lemma 2.9

The `SampleWithTrap` algorithm is given below, as Algorithm B.1. It relies on an ideal-factoring oracle which can be implemented either in quantum polynomial time or in classical sub-exponential time. We prove the following statement, which can be viewed as a reformulation of Lemma 2.9. (Recall that factoring ideals reduces in polynomial time to factoring integers.)

Algorithm B.1 $\texttt{SampleWithTrap}_{A,B}$

Input: Integers $2 \leq A \leq B$, a real $\delta \in (0,1]$ and a basis $\mathbf{B}_I$ of a non-zero ideal I.
Oracle: $\mathcal{F}$ for factoring integral ideals.
Output: $(\mathfrak{p}, w)$ with $\mathfrak{p} \in \mathcal{P}_{A,B}$, and $w \in I\mathfrak{p}$.

Set $\varepsilon = \delta/(8B)$.
Set $M = \sqrt{4 + \ln(3/\varepsilon)/d}$.
Set $s = \max(\sqrt{d} \cdot \|\mathbf{B}_I^*\|,\ \Delta_K^{1/d} \cdot B^{1/d} \cdot \mathcal{N}(I)^{1/d} \cdot \sqrt{\ln(3/\varepsilon)})$.
Set $\mathbf{u} = Ms \cdot \mathbf{1}$ with $\mathbf{1} = (1, \ldots, 1)^T \in \mathbb{R}^d$.
Set $k_{\max} = d \cdot \log_A(2M \cdot \sqrt{d} \cdot \mathcal{N}(I)^{-1/d})$.

repeat
    Sample $w \leftarrow \widetilde{D}_{\mathbf{B}_I, s, \mathbf{u}}$ using Lemma A.7 with error bound ε.
    Compute $\mathfrak{a} = I^{-1} \cdot (w)$.
    Factor $\mathfrak{a}$ using $\mathcal{F}$ and let $\mathcal{S}$ be the set of distinct factors of $\mathfrak{a}$ in $\mathcal{P}_{A,B}$.
until $\mathcal{S} \neq \emptyset$.
Sample $\mathfrak{p}$ uniformly in $\mathcal{S}$.
With probability $1 - \frac{|\mathcal{S}| \cdot \mathcal{N}(\mathfrak{p})}{k_{\max} \cdot B}$, go to Step 6.

Return $(\mathfrak{p}, w)$

Lemma B.1. *Let $\mathcal{F}$ be an ideal-factoring oracle. Given as inputs two integers $2 \leq A < B$, a real $\delta \in (0,1]$ and the basis $\mathbf{B}_I$ of a non-zero ideal I,* `SampleWithTrap` *outputs $(\mathfrak{p}, w)$ such that*

- *the distribution of $\mathfrak{p}$ is at statistical distance δ from the uniform distribution on $\mathcal{P}_{A,B}$;*
- *the element w belongs to $I \cdot \mathfrak{p} \setminus \{0\}$ and verifies $\|w\| \leq 2s \cdot \sqrt{4d + \ln(24B/\delta)}$, where*

$$s = \max\left(\sqrt{d} \cdot \|\mathbf{B}_I^*\|,\ \Delta_K^{1/d} \cdot B^{1/d} \cdot \mathcal{N}(I)^{1/d} \cdot \sqrt{\ln(24B/\delta)}\right).$$

Furthermore, `SampleWithTrap` *runs in expected time polynomial in $B/|\mathcal{P}_{A,B}|$, B/A, $\log \Delta_K$, $\log B$, $\log(1/\delta)$ and in the size of its input.*

Proof. We first analyze the running time of `SampleWithTrap` and then its correctness.

Running time. Observe that every step of the algorithm can be performed in polynomial time. For Step 7, we use Lemma A.7, whose assumptions are indeed satisfied. We further observe that at Step 12, the rejection probability is always between 0 and 1, hence we can indeed reject with this probability. Note that we have $B \geq \mathcal{N}(\mathfrak{p})$ since $\mathfrak{p} \in \mathcal{P}_{A,B}$. Also, we have $|\mathcal{S}| \leq \log_A \mathcal{N}(\mathfrak{a})$. It hence suffices to show that for any non-zero ideal $\mathfrak{a}$ computed at Step 8, we have $\log_A \mathcal{N}(\mathfrak{a}) \leq k_{\max}$. From Lemma A.7, we know that $\|w - \mathbf{u}\| < s \cdot \sqrt{\ln(3/\varepsilon) + 4d}$. As $\|\mathbf{u}\| = \sqrt{d} \cdot M \cdot s = s \cdot \sqrt{\ln(3/\varepsilon) + 4d}$, we have $\|w\| \leq 2\|\mathbf{u}\| = 2M \cdot \sqrt{d} \cdot s$, which in turn implies that $\mathcal{N}(w) \leq \|w\|^d \leq (2M \cdot \sqrt{d} \cdot s)^d$. Hence, we conclude that $\mathcal{N}(\mathfrak{a}) \leq (2M \cdot \sqrt{d} \cdot s)^d \cdot \mathcal{N}(I)^{-1} = A^{k_{\max}}$. This shows that $(|\mathcal{S}| \cdot \mathcal{N}(\mathfrak{p}))/(k_{\max} \cdot B)$ belongs to $[0, 1]$, as desired.

We now study the probability of exiting the outer loop, from Step 6 to Step 12. It is bounded from below by $A/(k_{\max} B)$ (since we have $|\mathcal{S}| \geq 1$ when we exit the inner loop). Hence, the expected number of iterations of this loop is at most $k_{\max} \cdot B/A$. Since $A \geq 2$, then $k_{\max}$ is polynomial in $d = \mathrm{poly}(\log \Delta_K)$, $\log(s)$, $\log(M)$ and $\log \mathcal{N}(I^{-1})$. From the definition of s, one can check that $k_{\max}$ is polynomial in $\log \Delta_K$, $\log B$, $\log \log(1/\delta)$ and the size of the input.

It remains to bound from below the probability of exiting the inner loop, from Step 6 to Step 10. The proof of this statement is an adaptation of the proof of [Gen09a, Le. 15.2.3]. This probability can be written as:

$$\Pr_{w \leftarrow \widetilde{D}_{\mathbf{B}_I, s, \mathbf{u}}} \left(\exists \mathfrak{p} \in \mathcal{P}_{A,B} : \mathfrak{p} \text{ divides } I^{-1} \cdot (w) \right) = \sum_{w \in I} \mathbf{1}_W(w) \cdot \widetilde{D}_{\mathbf{B}_I, s, \mathbf{u}}(w) \tag{5}$$

where $W = \cup_{\mathfrak{p} \in \mathcal{P}_{A,B}} I \cdot \mathfrak{p}$ and $\mathbb{1}_W(\cdot)$ is the indicator function of W. For any $w \in I \setminus \{0\}$, we have

$$\mathbb{1}_W(w) \geq \frac{1}{\ln(\mathcal{N}(w \cdot I^{-1}))} \cdot \sum_{\substack{\mathfrak{p} \in \mathcal{P}_{A,B} \\ \mathfrak{p} \mid w \cdot I^{-1}}} \ln(\mathcal{N}(\mathfrak{p})).$$

Indeed, either $w \notin W$ and the sum on the right is empty, or $w \in W$ and the sum on the right is bounded from above by 1 (since the norm of the product of all the primes dividing $w \cdot I^{-1}$ is at most the norm of $w \cdot I^{-1}$ when $w \cdot I^{-1}$ is non-zero). Moreover, we have already seen that the algebraic norm of $\mathfrak{a} = w \cdot I^{-1}$ is at most $(2M \cdot \sqrt{d} \cdot s)^d \cdot \mathcal{N}(I^{-1})$, and by assumption we know that $\mathcal{N}(\mathfrak{p}) \geq A$ for all $\mathfrak{p} \in \mathcal{P}_{A,B}$. Hence, letting $\mathbb{1}_{I \cdot \mathfrak{p}}(\cdot)$ be the indicator function of $I \cdot \mathfrak{p}$, it holds that

$$\begin{aligned} \mathbb{1}_W(w) &\geq \frac{\ln(A)}{\ln\left((2M \cdot \sqrt{d} \cdot s)^d \cdot \mathcal{N}(I^{-1})\right)} \cdot \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbb{1}_{I \cdot \mathfrak{p}}(w) \\ &= \frac{1}{k_{\max}} \cdot \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbb{1}_{I \cdot \mathfrak{p}}(w), \end{aligned}$$

where $k_{\max}$ is defined as in Step 5.

Before returning to (5), note that $\widetilde{D}_{\mathbf{B},s,\mathbf{u}}(0) = 0$. Indeed, we have seen that $\|w - \mathbf{u}\| < M \cdot \sqrt{d} \cdot s$ and, by construction, we have $\|\mathbf{u}\| = M \cdot \sqrt{d} \cdot s$. Using the above, we obtain:

$$\begin{aligned}
\sum_{w \in I} \mathbb{1}_W(w) \cdot \widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(w) &= \sum_{w \in I \setminus \{0\}} \mathbb{1}_W(w) \cdot \widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(w) \\
&\geq \frac{1}{k_{\max}} \sum_{w \in I \setminus \{0\}} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbb{1}_{I \cdot \mathfrak{p}}(w) \cdot \widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(w) \\
&= \frac{1}{k_{\max}} \sum_{w \in I} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathbb{1}_{I \cdot \mathfrak{p}}(w) \cdot \widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(w) \\
&= \frac{1}{k_{\max}} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(I \cdot \mathfrak{p}).
\end{aligned}$$

From Lemma A.7, we know that $\mathrm{SD}(\widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}, D_{I,s,\mathbf{u}}) \leq \varepsilon$. Hence, it holds that $\widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(I \cdot \mathfrak{p}) \geq D_{I,s,\mathbf{u}}(I \cdot \mathfrak{p}) - \varepsilon$. Moreover, observe that by definition of s, it holds that for any $\mathfrak{p} \in \mathcal{P}_{A,B}$ we have $s \geq \mathcal{N}(I \cdot \mathfrak{p})^{1/d} \cdot \Delta_K^{1/d} \cdot \sqrt{\ln(3/\varepsilon)}$. Hence, we can apply Corollary A.2 and we obtain that $D_{I,s,\mathbf{u}}(\mathfrak{p} \cdot I) \geq (1-\varepsilon) \cdot \mathcal{N}(\mathfrak{p})^{-1} \geq 1/(2B)$. By choice of ε, we finally obtain

$$\widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}(I \cdot \mathfrak{p}) \geq \frac{1}{2B} - \varepsilon \geq \frac{1}{4B}.$$

Plugging this back in our lower bound on the probability to exit the inner loop, we have

$$\Pr_{w \leftarrow \widetilde{D}_{\mathbf{B}_I,s,\mathbf{u}}} (\exists \mathfrak{p} \in \mathcal{P}_{A,B} : \mathfrak{p} \text{ divides } I^{-1} \cdot (w)) \;\geq\; \frac{1}{k_{\max}} \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \frac{1}{4B} \;=\; \frac{|\mathcal{P}_{A,B}|}{4B \cdot k_{\max}}.$$

The expected number of iterations of the inner loop is then $\leq 4k_{\max} \cdot B/|\mathcal{P}_{A,B}|$.

Correctness. Let $(\mathfrak{p}, w)$ be the output of `SampleWithTrap` on input $(A, B, \delta, \mathbf{B}_I)$. By construction, we have $\mathfrak{p} \in \mathcal{P}_{A,B}$. Further, as $w \in I$ and $\mathfrak{p} | I^{-1} \cdot (w)$, we have that $w \in I \cdot \mathfrak{p}$. The bound on $\|w\|$ comes from the fact that $\|w\| \leq 2\|\mathbf{u}\| = 2M \cdot \sqrt{d} \cdot s$. It remains to prove that the distribution $\mathcal{D}$ of the ideal $\mathfrak{p}$ is within statistical distance $\leq \delta/2$ from uniform over $\mathcal{P}_{A,B}$.

Let us fix $\mathfrak{p} \in \mathcal{P}_{A,B}$ and compute $\mathcal{D}(\mathfrak{p})$. First, we compute the probability that $\mathfrak{p}$ is chosen at Step 11 of the algorithm. The distribution of the element w when exiting of the inner loop is $\widetilde{D}_{\mathbf{B},s,\mathbf{u}}$ conditioned on $w \in W = \cup_{\mathfrak{q} \in \mathcal{P}_{A,B}} I \cdot \mathfrak{q}$ (which is equivalent to $S \neq \emptyset$). Moreover, the ideal $\mathfrak{p}$ belongs to S if and only if $w \in I \cdot \mathfrak{q}$. So the probability that $\mathfrak{p}$ belongs to S in Step 11 is

$$\Pr(\mathfrak{p} \in S \text{ in Step 11}) = \frac{\widetilde{D}_{\mathbf{B},s,\mathbf{u}}(I \cdot \mathfrak{p})}{\widetilde{D}_{\mathbf{B},s,\mathbf{u}}(W)}.$$

Note that the quantity $\widetilde{D}_{\mathbf{B},s,\mathbf{u}}(W)$ is a fixed and independent of $\mathfrak{p}$ (and non-zero, since the algorithm terminates). In the rest of the computation, we will write it

p_0. After running Step 11, we obtain

$$\Pr\left(\mathfrak{p} \text{ is chosen in Step 11}\right) = \frac{\widetilde{D}_{\mathbf{B},s,\mathbf{u}}(I \cdot \mathfrak{p})}{|\mathcal{S}| \cdot p_0}.$$

By Lemma A.7, Corollary A.2 and the choice of s, we know that

$$\begin{aligned}\widetilde{D}_{\mathbf{B},s,\mathbf{u}}(I \cdot \mathfrak{p}) &\in [1-\varepsilon, 1+\varepsilon] \cdot \mathcal{N}(\mathfrak{p})^{-1} + [-\varepsilon, \varepsilon] \\ &\subseteq [1-\delta/4, 1+\delta/4] \cdot \mathcal{N}(\mathfrak{p})^{-1},\end{aligned}$$

where in the last inequality we used the fact that $\varepsilon = \delta/(8 \cdot B) \leq \delta/8 \cdot \mathcal{N}(\mathfrak{p})^{-1}$. Combining this with the equation above, we obtain that

$$\Pr\left(\mathfrak{p} \text{ is chosen in Step 11}\right) \in \left[1 - \frac{\delta}{4}, 1 + \frac{\delta}{4}\right] \cdot \frac{1}{\mathcal{N}(\mathfrak{p}) \cdot |\mathcal{S}| \cdot p_0}.$$

Finally, because of the rejection sampling in Step 12, we have

$$\begin{aligned}\Pr\left(\mathfrak{p} \text{ is selected after Step 12}\right) &\in \left[1 - \frac{\delta}{4}, 1 + \frac{\delta}{4}\right] \cdot \frac{1}{\mathcal{N}(\mathfrak{p}) \cdot |\mathcal{S}| \cdot p_0} \cdot \frac{|\mathcal{S}| \cdot \mathcal{N}(\mathfrak{p})}{k_{\max} \cdot B} \cdot \frac{1}{p_0'} \\ &= \left[1 - \frac{\delta}{4}, 1 + \frac{\delta}{4}\right] \cdot \frac{1}{k_{\max} \cdot B \cdot p_0 \cdot p_0'},\end{aligned}$$

where p_0' is the probability (over the random choice of w, the random choice of $\mathfrak{p}$ and the rejection probability of Step 12) that one exists the outer loop.

Overall, we have just proven that there exists some quantity C such that for any $\mathfrak{p} \in \mathcal{P}_{A,B}$, it holds that $\mathcal{D}(\mathfrak{p}) \in [1-\delta/4, 1+\delta/4] \cdot C$. Since $\sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \mathcal{D}(\mathfrak{p}) = 1$, it must be that $C \in \left[\frac{1}{1+\delta/4}, \frac{1}{1-\delta/4}\right] \cdot \frac{1}{|\mathcal{P}_{A,B}|}$. It implies that for all $\mathfrak{p} \in \mathcal{P}_{A,B}$,

$$\left|\mathcal{D}(\mathfrak{p}) - \frac{1}{|\mathcal{P}_{A,B}|}\right| \leq \max\left(1 - \frac{1-\delta/4}{1+\delta/4}, \frac{1+\delta/4}{1-\delta/4} - 1\right) \cdot \frac{1}{|\mathcal{P}_{A,B}|} \leq \frac{\delta}{|\mathcal{P}_{A,B}|}.$$

The statistical distance between $\mathcal{D}$ and the uniform distribution satisfies

$$\begin{aligned}\mathrm{SD}\left(\mathcal{D}, \mathcal{U}(\mathcal{P}_{A,B})\right) &= \frac{1}{2} \cdot \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \left|\mathcal{D}(\mathfrak{p}) - \frac{1}{|\mathcal{P}_{A,B}|}\right| \\ &\leq \frac{\delta}{2} \cdot \sum_{\mathfrak{p} \in \mathcal{P}_{A,B}} \frac{1}{|\mathcal{P}_{A,B}|} = \frac{\delta}{2}.\end{aligned}$$

This completes the proof. □

C Proof of Theorem 2.10

In this section, we provide a proof of Gentry's reduction for SVP, as stated in Theorem 2.10. The proof is similar to the one provided in Gentry's thesis [Gen09a], but we instantiate it directly with the shortest vector problem, instead of the variant of the bounded distance decoding problem used in [Gen09a].

C.1 Balanced-ideal-HSVP

In the proof, we will make use of balanced elements (as defined in Definition 2.4). We introduce the problem ideal-balanced-HSVP, and give a (folklore) proof that this problem is equivalent to id-HSVP (up to a polynomial loss in the approximation factor). The balanced version of id-HSVP will be more convenient to use in the following proof.

Definition C.1. *Let $\eta > 1$ and $\gamma \geq 1$. The problem* id-BHSVP$^{\eta}_{\gamma}$ *asks, given as input a fractional ideal I, to find a non-zero element $x \in I$ such that $\|x\| \leq \gamma \cdot \mathrm{vol}(I)^{1/d}$ and x is η-balanced. The problem* inv-BHSVP$^{\eta}_{\gamma}$ *is the problem* id-BHSVP$^{\eta}_{\gamma}$ *restricted to inverses of integral lattices.*

We describe in Algorithm C.1 a polynomial-time reduction from id-BHSVP to id-HSVP, which relies on Babai's nearest plane algorithm [Bab86].

Algorithm C.1 `BalanceElement`

Input: The HNF of a fractional ideal I, an element $x \in I$ and $M > 0$.
Output: $y \in I$.

1: Let $s = \sqrt{d} \cdot \delta_K \cdot \|x\|_\infty$.
2: Let $\mathbf{B}_I = \texttt{ReduceIdeal}(I, x)$.
3: Let $\mathbf{t} = s\sqrt{d}(M+1)/2 \cdot \mathbf{1}$ with $\mathbf{1} = (1, \ldots, 1) \in K_{\mathbb{R}}$.
4: Run Babai's nearest plane algorithm on $(\mathbf{B}_I, \mathbf{t})$; let $y \in I$ be the output.
5: Return y.

Lemma C.2. *Algorithm C.1 runs in polynomial time. On input I, x, M with $x \in I \setminus \{0\}$, it outputs $y \in I \setminus \{0\}$ that is $(1 + 2/M)$-balanced and satisfies*

$$\|y\| \leq (1 + M/2) \cdot d^{3/2} \cdot \delta_K \cdot \|x\|_\infty.$$

Proof. The running time follows directly from the description of the algorithm. Let y be the output of Algorithm C.1 on input I, x and M. We have, by property of the nearest plane algorithm (see [Bab86, Th. 3.1]), that there exist $\mu_1, \ldots, \mu_d$ in $[-1/2, 1/2]$ such that

$$\|y - \mathbf{t}\|_\infty \leq \|y - \mathbf{t}\| \leq \Big(\sum_i \mu_i^2 \cdot \|\mathbf{b}_i^*\|^2 \Big)^{1/2}.$$

By Lemma 2.8, we have $\|\mathbf{B}_I^*\| \leq \delta_K \cdot \|x\| \leq s$. We hence obtain that $\|y - \mathbf{t}\|_\infty \leq \sqrt{d}s/2$. As a result, we have $|y_i| \in [s\sqrt{d}M/2, s\sqrt{d}(M+2)/2]$ for all i.

We then have $\mathcal{N}^{1/d}(y) \geq s\sqrt{d}M/2 \geq M/(M+2)\|y\|_\infty$. The same holds for y^{-1}, which gives that y is $(1 + 2/M)$-balanced. Finally, the inequality $\|y\| \leq \sqrt{d} \cdot \|y\|_\infty$ gives the desired bound on the norm of y. □

Corollary C.3. *For any $\gamma \geq 1$ and $\eta > 1$, there is a Karp polynomial time reduction from* id-BHSVP$^{\eta}_{\gamma'}$ *to* id-HSVP$_{\gamma}$*, where $\gamma' = \gamma \cdot \delta_K \cdot d^{3/2} \cdot \eta/(\eta - 1)$.*

Note that the converse reduction holds without any parameter loss, by definition of id-BHSVP.

Proof. Let I be an instance of id-BHSVP$^{\eta}_{\gamma'}$, and assume we have an oracle $\mathcal{O}$ for id-HSVP$_{\gamma}$. Let x be the output of $\mathcal{O}$ on I. We let $M = 2/(\eta - 1)$ and return $y = \texttt{BalanceElement}(I, x, M)$. The fact that y is a valid id-BHSVP$^{\eta}_{\gamma'}$ solution follows from the definition of M and Lemma C.2. □

Corollary C.4. *For any constant $\eta > 1$,* id-BHSVP$^{\eta}_{\gamma_{easy}(\eta)}$ *can be solved in polynomial time for $\gamma_{easy}(\eta) = \delta_K \cdot d^{3/2} \cdot (\eta/(\eta - 1)) \cdot 2^d$.*

Proof. The result follows from Corollary C.3, by using the LLL algorithm to solve id-HSVP$_{\gamma}$. □

C.2 Finding a non-trivial solution to inv-HSVP using a $\mathcal{P}^{-1}_{A,B}$-avg-HSVP oracle

The reduction from Theorem 2.10 is an iterative reduction, which proceeds by iteratively improving an existing solution with the usage of an oracle solving $\mathcal{P}^{-1}_{A,B}$-avg-id-HSVP$_{\gamma_{\text{avg}}}$. In this subsection, we focus on the main ingredient of one iteration of the reduction, the `SampleSmall` algorithm, presented in Algorithm C.2. The objective of this algorithm is, given as input $\mathfrak{b}^{-1}$ the inverse of a prime ideal, to find a non-trivial short non-zero vector in $\mathfrak{b}^{-1}$. Indeed, since $\mathfrak{b}$ is integral, we know that $\mathcal{O}_K \subseteq \mathfrak{b}^{-1}$, so the short non-zero vectors of $\mathcal{O}_K$ give trivial solutions to short non-zero vectors in $\mathfrak{b}^{-1}$. The objective of the `SampleSmall` algorithm is to find slightly shorter vectors than those trivial short vectors lying in $\mathcal{O}_K$ (which will exist if the norm of $\mathfrak{b}$ is large enough). In particular, we would like to obtain $x \in \mathfrak{b}^{-1} \setminus \{0\}$ with $\|x\| < 1$, so that multiplying by x decrease the euclidean norm. This will be used in the reduction to iteratively decreases the norm of a short non-zero vector found in our input ideal I.

Algorithm C.2 `SampleSmall`$_{A,B}$

Input: A basis of an integral ideal $\mathfrak{b}$.
Oracles: $\mathcal{O}$ for $\mathcal{P}^{-1}_{A,B}$-avg-id-HSVP$_{\gamma_{\text{avg}}}$, $\mathcal{F}$ for factoring integral ideals.
Output: $x \in \mathfrak{b}^{-1}$ or $x = \perp$.

1: Compute a basis $\mathbf{B}$ of $\mathfrak{b}^{-1}$ with $\|\mathbf{B}^*\| \leq \delta_K$ (using `InvertIdeal`).
2: Set $(\mathfrak{p}, w)$ be the output of `SampleWithTrap`$_{A,B}$ on input $(A, B, 2^{-(d+1)}, \mathbf{B})$ (this relies on $\mathcal{F}$).
3: Set $v = \mathcal{O}(\mathfrak{p}^{-1})$.
4: If $v \neq \perp$, then return $v \cdot w$.
5: Else, return $\perp$.

Theorem C.5. *Let $\gamma_{avg} \geq 1$ and $3 \leq A < B$ satisfying $B/|\mathcal{P}_{A,B}|, B/A \leq \mathrm{poly}(\log \Delta_K)$. Let $\mathcal{O}$ be an oracle solving $\mathcal{P}_{A,B}^{-1}$-avg-id-HSVP$_{\gamma_{avg}}$ with success probability $\delta \geq 2^{-d}$ and let $\mathcal{F}$ be an ideal-factoring oracle.*

On input a non-zero integral ideal $\mathfrak{b}$ and given access to $\mathcal{O}$ and $\mathcal{F}$, Algorithm `SampleSmall`$_{A,B}$ *runs in expected time* $\mathrm{poly}(\log \Delta_K, \log B, \log \mathcal{N}(\mathfrak{b}))$*, and performs only one call to $\mathcal{O}$ and possibly multiple calls to $\mathcal{F}$ for integral ideals of norm* $\mathrm{poly}(\log \Delta_K, \log B, \log \mathcal{N}(\mathfrak{b}))$ *bits. It outputs $x \neq \perp$ with probability $\geq \delta/2$ and, when this is the case, it holds that $x \in \mathfrak{b}^{-1} \setminus \{0\}$ and*

$$\|x\| \leq \frac{10\gamma_{avg}(d + \ln B) \cdot \Delta_K^{1/(2d)}}{A^{1/d}} \cdot \max\left(\delta_K, \left(B \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1})\right)^{1/d}\right).$$

Proof. We first focus on the running time of the algorithm. Every step can be performed in polynomial time. For Step 1, we use Lemma 2.8 and the fact that $\mathfrak{b}$ is integral. For Step 2, we use Lemma 2.9. Note that Step 3 is not inside a loop, hence the call to $\mathcal{O}$ is performed only once.

Let us now prove that the algorithm returns an element $x \neq \perp$ with probability at least $\delta/2$. Note that by Lemma 2.9, the distribution $\mathcal{D}$ of the ideal $\mathfrak{p}$ given as input to $\mathcal{O}$ is within statistical distance $\leq 2^{-(d+1)} \leq \delta/2$ from uniform over $\mathcal{P}_{A,B}$ (here we used the lower bound $\delta \geq 2^{-d}$). Since we know that $\mathcal{O}$ has success probability δ when its input $\mathfrak{p}^{-1}$ is distributed uniformly in $\mathcal{P}_{A,B}^{-1}$, this proves that the probability that $\mathcal{O}$ succeeds in solving id-HSVP$_{\gamma_{\mathrm{avg}}}$ in Step 3 of the algorithm is at least $\delta - \delta/2 \geq \delta/2$, as desired.

Finally, let us prove the upper bound on $\|x\|$ when $x \neq \perp$. In this case, we have $x = v \cdot w$ and use the upper bounds on v and on w (from Lemma 2.9) to obtain

$$\begin{aligned}
\|x\| &\leq \|v\| \cdot \|w\| \\
&\leq \gamma_{\mathrm{avg}} \cdot \mathrm{Vol}(\mathfrak{p}^{-1})^{1/d} \cdot 2(5d + \ln B + \ln(48)) \cdot \max\left(\delta_K, \left(\Delta_K \cdot B \cdot \mathcal{N}(\mathfrak{b}^{-1})\right)^{1/d}\right) \\
&\leq \frac{10\gamma_{\mathrm{avg}}(d + \ln B) \cdot \Delta_K^{1/(2d)}}{A^{1/d}} \cdot \max\left(\delta_K, \left(B \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1})\right)^{1/d}\right),
\end{aligned}$$

where we used the fact that $B \geq 3$. This completes the proof. □

For simplicity, we will use the following corollary, where we use the Extended Riemann Hypothesis in order to estimate the number of prime ideals in the set $\mathcal{P}_{A,B}$ and simplify the conditions.

Corollary C.6 (ERH). *Let $\gamma_{avg} \geq 1$ and $3 \leq A \leq (\Delta_K)^{d^{O(1)}}$. Let $\mathcal{O}$ be an oracle solving $\mathcal{P}_{A,4A}^{-1}$-avg-id-HSVP$_{\gamma_{avg}}$ with success probability $\delta \in (0, 1]$ and let $\mathcal{F}$ be an oracle factoring integral ideals. Let $\varepsilon \in (0, 1)$ and assume that*

$$A^{1/d} \geq 10 \cdot \gamma_{avg} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K \cdot \varepsilon^{-1}.$$

Then there exists an algorithm $\mathcal{A}$ that takes as input any integral ideal $\mathfrak{b}$ with $\mathcal{N}(\mathfrak{b}) \geq 4A$ and outputs $x \in \mathfrak{b}^{-1} \setminus \{0\}$ such that $\|x\| \leq \varepsilon$. If given access to $\mathcal{O}$ and $\mathcal{F}$, algorithm $\mathcal{A}$ runs in expected time $\mathrm{poly}(\log \Delta_K, \log(\mathcal{N}(\mathfrak{b})), 1/\delta)$ *and calls $\mathcal{F}$ on ideals of norm at most* $\mathrm{poly}(\log \Delta_K, \log \mathcal{N}(\mathfrak{b}))$ *bits.*

Proof. Algorithm $\mathcal{A}$ consists in repeatedly running $\texttt{SampleSmall}_{A,B}$ with $B = 4A$, on input $\mathfrak{b}$, until it outputs $x \neq \perp$. Let us prove that A and B satisfy the constraints required in Theorem C.5. If $A \leq \text{poly}(\log \Delta_K)$, then $4A/|\mathcal{P}_{A,4A}|$ is polynomial. Else, the ERH implies that

$$\frac{4A}{|\mathcal{P}_{A,B}|} \leq O(\ln A) \leq \text{poly}(\log \Delta_K).$$

The claim on the running time of algorithm $\mathcal{A}$ and the fact that $x \in \mathfrak{b}^{-1} \setminus 0$ follow from Theorem C.5. Note that Theorem C.5 is guaranteed to work only if the success probability δ of $\mathcal{O}$ is at least 2^{-d}. If the success probability is smaller than this quantity, algorithm $\mathcal{A}$ simply runs an SVP solver on ideal $\mathfrak{b}^{-1}$ and returns a shortest non-zero vector. This shortest non-zero vector will have Euclidean norm $\leq \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot A^{1/d} \leq \varepsilon$ by assumption on A, and the call to the SVP solver has a running time $2^{O(d)} = \text{poly}(1/\delta)$.

We now bound $\|x\|$. From Theorem C.5 and by choice of B, we know that

$$\|x\| \leq \frac{10\gamma_{\text{avg}}(d + \ln(4A)) \cdot \Delta_K^{1/(2d)}}{A^{1/d}} \cdot \max\left(\delta_K, \left(4A \cdot \Delta_K \cdot \mathcal{N}(\mathfrak{b}^{-1})\right)^{1/d}\right).$$

Since $\mathcal{N}(\mathfrak{b}) \geq 4A$ and $\delta_K \geq \lambda_d(\mathcal{O}_K) \geq \Delta_K^{1/(2d)}$, it holds that $(4A \cdot \Delta_K/\mathcal{N}(\mathfrak{b}))^{1/d} \leq \delta_K \cdot \Delta_K^{1/(2d)}$, and hence

$$\|x\| \leq \frac{10\gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K}{A^{1/d}} \leq \varepsilon.$$

The last inequality follows from the assumption on A and ε. □

C.3 Iterating the reduction

In order to prove Theorem 2.10, we are going to use the id-BHSVP problem. Recall that the id-BHSVP problem is equivalent to the id-HSVP problem, up to some polynomial loss, so we can safely replace id-HSVP by id-BHSVP, which will make our reductions easier to prove. The lemma below states that if we have an oracle solving $\mathcal{P}_{A,4A}^{-1}$-avg-id-HSVP$_{\gamma_{\text{avg}}}$ and an algorithm solving inv-BHSVP$_{\gamma}^{\eta}$, then we can create an algorithm solving inv-BHSVP$_{\gamma'}^{\eta'}$ where γ' is slightly smaller than γ and η' is slightly larger than η (i.e., we can find smaller less balanced vectors in our ideals). This corresponds to one iteration of the full reduction.

For the whole subsection, we fix $\gamma_{\text{avg}} \geq 1$, $\varepsilon \in (0, 1)$ and $3 \leq A \leq (\Delta_K)^{d^{O(1)}}$ satisfying:

$$A^{1/d} \geq 10 \cdot \gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K \cdot \varepsilon^{-1}.$$

Lemma C.7 (ERH). *Let* $\gamma_{min} = (4A)^{1/d}/(\varepsilon \cdot \Delta_K^{1/(2d)})$, $\gamma > \gamma_{min}$ *and* $\eta \in (1, \gamma/\gamma_{min}]$.

$$\text{inv-BHSVP}_{\gamma'}^{\eta'} \textit{ reduces to } \text{inv-BHSVP}_{\gamma}^{\eta} \textit{ and } \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{avg}}$$

for $\eta' = \eta \cdot (1 + 1/d)$ *and* $\gamma' = 2 \cdot d^{5/2} \cdot \delta_K \cdot \varepsilon \cdot \gamma$. *If given access to an oracle* $\mathcal{F}$ *factoring integral ideals, the expected running time of the reduction is polynomial in* $\log \Delta_K$, $\log \gamma$, $1/\delta$ *and the size of its input, where* δ *is the success probability of the oracle for* $\mathcal{P}_{A,4A}^{-1}$-avg-id-HSVP$_{\gamma_{avg}}$. *Moreover, the oracle* $\mathcal{F}$ *is called on ideals whose norms have a bit-size* $\mathrm{poly}(\log \Delta_K, \log \gamma)$.

Proof. Assume that we are given $I = \mathfrak{b}^{-1}$ the inverse of an integral ideal. Let x be the output of the inv-BHSVP$_\gamma^\eta$ oracle on input I. As $\eta' \geq \eta$, the element x is η'-balanced. If $\|x\|_\infty \leq \varepsilon \cdot \gamma \cdot \mathrm{Vol}(I)^{1/d}$, then it is a solution for inv-BHSVP$_{\sqrt{d} \cdot \varepsilon \cdot \gamma}^{\eta'}$ and we can output it. Else, we have

$$|\mathcal{N}(x)| \geq \eta^{-d} \cdot \|x\|_\infty^d \geq \eta^{-d} \cdot \varepsilon^d \cdot \gamma^d \cdot \Delta_K^{1/2} \cdot \mathcal{N}(I).$$

Now we set $\mathfrak{b} = (x) \cdot I^{-1}$. This ideal is the inverse of an integral ideal, and by the previous inequality and the condition on η we have

$$\mathcal{N}(\mathfrak{b}) = \frac{\mathcal{N}(x)}{\mathcal{N}(I)} \geq \frac{\varepsilon^d \cdot \gamma^d \cdot \Delta_K^{1/2}}{\eta^d} \geq 4A.$$

This last inequality, and the definition of A meet the conditions of Corollary C.6, we then can make a call to $\texttt{SampleSmall}_{A,4A}(\mathfrak{b})$ and denote by y its output. The element y verifies $\|y\|_\infty \leq \varepsilon$ and $y \in \mathfrak{b}^{-1} \setminus \{0\}$.

We now denote $y' = \texttt{BalanceElement}(\mathfrak{b}^{-1}, y, 2d)$. By Lemma C.2, we have that $y' \in \mathfrak{b}^{-1} \setminus 0$ is $(1 + 1/d)$-balanced and that

$$\|y'\| \leq (1 + d) \cdot d^{3/2} \cdot \delta_K \cdot \varepsilon \leq 2 \cdot d^{5/2} \cdot \delta_K \cdot \varepsilon$$

We then return $y' \cdot x$. We have that $y' \cdot x \in I$, and since x is η-balanced and y' is $(1 + 1/d)$-balanced, then xy' is η'-balanced and

$$\|x \cdot y'\| \leq \|y'\| \cdot \|x\| \leq 2 \cdot d^{5/2} \cdot \delta_K \cdot \varepsilon \cdot \gamma \cdot \mathrm{Vol}(I)^{1/d} = \gamma' \cdot \mathrm{Vol}(I)^{1/d}.$$

The running time of the algorithm comes from the running time of the call to $\texttt{SampleSmall}_{A,4A}(\mathfrak{b})$ and the running time of $\texttt{BalanceElement}(\mathfrak{b}^{-1}, y, 2d)$. The former is polynomial in $\log \Delta_K$, $\log \mathcal{N}(\mathfrak{b})$ and $1/\delta$ and requires factoring ideals of norm at most $\mathrm{poly}(\log \Delta_K, \log \mathcal{N}(\mathfrak{b}))$ bits. The latter has a running time polynomial in $\log \Delta_K$ and $\log \mathcal{N}(\mathfrak{b})$. Observe that $\mathcal{N}(\mathfrak{b}) = |\mathcal{N}(x)|/\mathcal{N}(I) \leq \|x\|^d / N(I) \leq \gamma^d \cdot \sqrt{\Delta_K}$. The result follows. □

We will now iterate Lemma C.7, instantiated with $\varepsilon = 1/2 \cdot (2 \cdot d^{5/2} \cdot \delta_K)^{-1}$. This choice of ε ensures that $\gamma' = \gamma/2$, i.e., the approximation factor is divided by 2 at every iteration of the reduction (at the cost of slightly less balanced elements). We will iterate this reduction step until we obtain a reduction from inv-BHSVP$_{\gamma'}^{\eta'}$ with an approximation factor γ' as small as possible, to inv-BHSVP$_\gamma^\eta$ with γ so large that it can be solved in polynomial time using the LLL algorithm. Hence, the only oracle that will remain for the reduction to work is the one solving $\mathcal{P}_{A,4A}^{-1}$-avg-id-HSVP$_{\gamma_{\mathrm{avg}}}$ (and the one factoring ideals, which can be quantumly efficiently instantiated).

Lemma C.8. *Let* $\gamma_{avg} \geq 1$, $3 \leq A \leq \Delta_K^{d^{O(1)}}$ *verifying*

$$A^{1/d} \geq \gamma_{avg} \cdot 40 \cdot d^{5/2} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K^2$$

and

$$\gamma_{min} = \frac{4 \cdot d^{5/2} \cdot \delta_K \cdot (4A)^{1/d}}{\Delta_K^{1/(2d)}}.$$

There exists a reduction

$$\textit{from} \quad \text{inv-BHSVP}^{2e}_{2e\gamma_{min}} \quad \textit{to} \quad \mathcal{P}^{-1}_{A,4A}\text{-avg-id-HSVP}_{\gamma_{avg}}.$$

Given access to an ideal-factoring oracle $\mathcal{F}$*, the expected running time of this reduction is polynomial in its input bit-size, in* $\log \Delta_K$ *and in* $1/\delta$*, where* $\delta \in (0, 1]$ *is the success probability of the* $\mathcal{P}^{-1}_{A,4A}$-avg-id-HSVP$_{\gamma_{avg}}$ *oracle. Moreover, the reduction calls* $\mathcal{F}$ *on integral ideals whose algebraic norms have bit-size* $\text{poly}(\log \Delta_K)$.

Proof. Let $\varepsilon = (4d^{5/2} \cdot \delta_K)^{-1}$. Define $\gamma_0 = \gamma_{\min} \cdot 2e \cdot 2^d$, $\eta_0 = 2$, and for any $k \in \{1, \ldots, d\}$ $\gamma_k = \gamma_0 \cdot 2^{-k}$ and $\eta_k = \eta_0 \cdot (1 + 1/d)^k$. Observe that, for any k, we have that $\gamma_k > \gamma_{\min}$ and $\eta_k \in (1, \gamma_k/\gamma_{\min}]$. Moreover, if we let $\varepsilon = (4d^{5/2} \cdot \delta_K)^{-1}$, then our choice of $\gamma_{\min}$ coincide with the definition of $\gamma_{\min}$ in Lemma C.7, and our choice of A satisfies the constraint $A^{1/d} \geq 10 \cdot \gamma_{\text{avg}} \cdot (d + \ln(4A)) \cdot \Delta_K^{1/d} \cdot \delta_K \cdot \varepsilon^{-1}$.

We can then apply Lemma C.7 and we get, for any $0 \leq k < d$, that

$$\text{inv-BHSVP}^{\eta_{k+1}}_{\gamma_{k+1}} \leq \text{inv-BHSVP}^{\eta_k}_{\gamma_k} + \mathcal{P}^{-1}_{A,4A}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}.$$

By combining the reduction, we then have that

$$\text{inv-BHSVP}^{\eta_d}_{\gamma_d} \leq \text{inv-BHSVP}^{\eta_0}_{\gamma_0} + \mathcal{P}^{-1}_{A,4A}\text{-avg-id-HSVP}_{\gamma_{\text{avg}}}.$$

Now, from the definition of $\gamma_{\min}$ and the lower bound on $A^{1/d}$, one can check that $\gamma_0 \geq \delta_K \cdot d^{3/2} \cdot \left(\frac{\eta_0}{\eta_0 - 1}\right) \cdot 2^d$. Hence, by Corollary C.4 we have that $\text{inv-BHSVP}^{\eta_0}_{\gamma_0}$ can be solved in polynomial time.

Regarding the running time, our reduction consists in d consecutive reductions. From Lemma C.7, the k-th reduction has a running time polynomial in $\log \Delta_K$, $\log \gamma_k$ and $1/\delta$. Since for every k we have that $\log \gamma_k \leq \log \gamma_0 = \text{poly}(\log \Delta_K)$, we conclude that the total running time of the reduction is polynomial in $\log \Delta_K$ and $1/\delta$. The same argument also shows that the ideal-factoring oracle is only called on integral ideals whose norm have a bit-size $\text{poly}(\Delta_K)$. □

We are now ready to prove our main theorem of this section. To do so, we instantiate Lemma C.8 with an appropriate value of A, and combine the reduction with the ones from Appendix C.1 showing that inv-BHSVP is equivalent to id-HSVP (up to polynomial losses).

Proof (of Theorem 2.10). Let $C_{1,K}$ be minimal such that $C_{1,K} \geq 40 \cdot d^{5/2} \cdot (d + d^2 + \ln(4C_{1,K}^d)) \cdot \Delta_K^{1/d} \cdot \delta_K^2$. Then $C_{1,K} = \mathrm{poly}(\Delta_K^{1/d}, \log \Delta_K, \delta_K)$. Moreover, using the fact that $\gamma_{\mathrm{avg}} \leq 2^d$, one can check that

$$(\gamma_{\mathrm{avg}}^d \cdot C_{1,K}^d)^{1/d} \geq \gamma_{\mathrm{avg}} \cdot 40 \cdot d^{5/2} \cdot (d + \ln(4 \cdot \gamma_{\mathrm{avg}}^d \cdot C_{1,K}^d)) \cdot \Delta_K^{1/d} \cdot \delta_K^2.$$

This inequality also holds for any $A \geq \gamma_{\mathrm{avg}}^d \cdot C_{1,K}^d$. Hence, any such A with $A \leq \Delta_K^{d^{O(1)}}$ satisfies the conditions of Lemma C.8. Now let

$$C_{2,K} = 2e \cdot \frac{4 \cdot d^{5/2} \cdot \delta_K \cdot 4^{1/d}}{\Delta_K^{1/(2d)}} = \mathrm{poly}(\log \Delta_K, \delta_K).$$

We set $\gamma = A^{1/d} \cdot C_{2,K}$ and observe that $\gamma \geq 2e \cdot \gamma_{\min}$ for $\gamma_{\min}$ as in Lemma C.8. Then by the Lemmas C.8 and 2.6 we have:

$$\text{id-HSVP}_\gamma \;\leq\; \text{inv-HSVP}_\gamma \;\leq\; \text{inv-BHSVP}_\gamma^{2e} \;\leq\; \mathcal{P}_{A,4A}^{-1}\text{-avg-id-HSVP}_{\gamma_{\mathrm{avg}}},$$

where the second reduction comes from the definition of id-BHSVP (a solution to id-BHSVP$_\gamma^\eta$ in any fractional ideal I is by definition also a solution of id-HSVP$_\gamma$ in L). This completes the proof. □