On Module Unique-SVP and NTRU

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

[1] Inria Lyon, Lyon, France
[2] ENS de Lyon, Lyon, France
[3] Univ. Bordeaux, CNRS, Inria, Bordeaux INP, IMB, Talence, France
[4] Institut Universitaire de France, Paris, France

Abstract. The NTRU problem can be viewed as an instance of finding a short non-zero vector in a lattice, under the promise that it contains an exceptionally short vector. Further, the lattice under scope has the structure of a rank-2 module over the ring of integers of a number field. Let us refer to this problem as the module unique Shortest Vector Problem, or mod-uSVP for short. We exhibit two reductions that together provide evidence the NTRU problem is not just a particular case of mod-uSVP, but representative of it from a computational perspective.

First, we reduce worst-case mod-uSVP to worst-case NTRU. For this, we rely on an oracle for id-SVP, the problem of finding short non-zero vectors in ideal lattices. Using the worst-case id-SVP to worst-case NTRU reduction from Pellet-Mary and Stehlé [ASIACRYPT'21], this shows that worst-case NTRU is equivalent to worst-case mod-uSVP.

Second, we give a random self-reduction for mod-uSVP. We put forward a distribution D^{uSVP} over mod-uSVP instances such that solving mod-uSVP with a non-negligible probability for samples from D^{uSVP} allows to solve mod-uSVP in the worst-case. With the first result, this gives a reduction from worst-case mod-uSVP to an average-case version of NTRU where the NTRU instance distribution is inherited from D^{uSVP}. This worst-case to average-case reduction requires an oracle for id-SVP.

1 Introduction

Let K be a number field, $\mathcal{O}_K$ its ring of integers and $\|\cdot\|$ the ℓ_2-norm in the complex embedding vector space. A notable example is $K = \mathbb{Q}[x]/(x^d+1)$ with d a power of 2: in this case, we have $\mathcal{O}_K = \mathbb{Z}[X]/\Phi(X)$ and $\|a\| = (d\sum_i |a_i|^2)^{1/2}$ for all $a = \sum_{0\le i<d} a_i x^i \in K$. In the (search) NTRU problem, one is given $h \in R_q := \mathcal{O}_K/q\mathcal{O}_K$ with the promise that there exists a pair $(f,g) \in \mathcal{O}_K^2$ such that $gh = f \bmod q\mathcal{O}_K$ and $\|f\|, \|g\|$ are significantly smaller than $\sqrt{q}$ (by a factor γ called the gap of the NTRU instance, see Definition 2.15 for a formal definition). The goal is to find a short multiple of the pair (f,g). An efficient algorithm for the NTRU problem for appropriate parameters would lead to a cryptanalysis of the seminal NTRU encryption scheme [HPS98], a variant of which appears among the finalists of the NIST post-quantum cryptography standardization process [CDH+20].

It was noticed very early that the NTRU problem can be interpreted in terms of Euclidean lattices [HPS98,CS97]. Indeed, the set $L_h := \{(a,b)^T \in K^2 : bh = a \bmod q\mathcal{O}_K\}$ forms a $(2d)$-dimensional lattice, when viewing $\mathcal{O}_K$ as a d-dimensional lattice via the embedding map (or, more elementarily for the running example, using the polynomial expressions). The lattice is described by h, from which a basis can be computed. This lattice has two peculiar properties. First, it contains an unusually short non-zero vector (f, g). Indeed, for most h's, we have $\det L_h = \Delta_K \cdot q^d$, where Δ_K refers to the field discriminant; our running example satisfies $\Delta_K = d^d$. As a result, one would expect the shortest non-zero vectors to have ℓ_2-norm around $q^{1/2}$, up to limited factors depending on Δ_K and d; but $(f, g)^T$ is much shorter, by assumption. However, this is not quite an instance of the unique Shortest Vector Problem (uSVP), as L_h does not contain just one exceptionally short non-zero vector (up to sign), but d linearly independent short vectors: in our running example, the $(x^i \cdot f, x^i \cdot g)^T$'s for $i \in [d]$ are linearly independent and belong to L_h and; in the general case, a short $\mathbb{Z}$-basis of $\mathcal{O}_K$ can be used in place of the x^i's. This leads us to the second peculiarity of the L_h lattice: as it is invariant under multiplication by elements of $\mathcal{O}_K$, it is a rank-2 $\mathcal{O}_K$-module. We hence have a rank-2 $\mathcal{O}_K$-module with the promise that it contains an unusually short non-zero vector, i.e., an unusually dense rank-1 submodule. We call mod-uSVP the problem of finding a short non-zero vector in rank-2 module containing an unusually short vector. In this introduction, we call gap of the mod-uSVP instance the ratio between the root determinant of the lattice (which predicts what would be expected for the euclidean norm of the shortest vector) and the actual euclidean norm of a shortest non-zero vector (see Definition 2.12 for a formal definition).

Search NTRU and mod-uSVP actually come with two flavors. The most natural one, described above, asks to recover a short vector of the corresponding rank-2 module. This is the variant we implicitly consider in this introduction when we discuss NTRU and mod-uSVP. As mentioned above, the NTRU and mod-uSVP lattices not only contain an unexpectedly short vector, but also an unexpectedly dense rank-1 sublattice. The second variant, which we refer to as $\mathrm{NTRU}_{\mathrm{mod}}$ or mod-$\mathrm{uSVP}_{\mathrm{mod}}$, asks to recover a basis of this dense submodule.

As seen above, the NTRU problem can be viewed as a special case of a lattice problem. It is however unclear if its instances are representative instances of some standard lattice problem, or, more precisely, if they are computationally equivalent to general instances of such a problem. In [Pei16, Section 4.4.4], Peikert sketched a reduction from a decision version of the NTRU problem to the Ring Learning With Errors (RLWE) problem [SSTX09,LPR10]; this reduction can be adapted to the search NTRU problem we consider here. Note that under some parameter constraints, RLWE is computationally equivalent to the Shortest Independent Vectors Problem for rank-2 modules [LS15,AD17] (mod-SIVP), which consists in finding $2d$ linearly independent vectors whose longest one is not much longer than optimal. Oppositely, in a recent work, Pellet-Mary and Stehlé [PS21] exhibited a reduction from the Shortest Vector Problem for lattices corresponding to ideals of $\mathcal{O}_K$ (id-SVP) to NTRU. Enhanced by the id-

SVP self-reducibility from [dBDPW20], this leads to a reduction from worst-case id-SVP to an average-case version of the NTRU problem.

Overall, we see that NTRU sits between id-SVP and mod-SIVP. Interestingly, id-SVP admits algorithms that outperform generic lattice reduction algorithms [LLL82,Sch87] for some parameter ranges [CDW21,PHS19]. As such a phenomenon is unknown in the case for mod-SIVP, there is potentially quite some room between id-SVP and mod-SIVP. With this state of affairs, it is unclear which of these problems captures the true hardness of NTRU, or if NTRU lies somewhere strictly in between.

Contributions. We give evidence that the NTRU problem is not just a particular case of mod-uSVP, but actually representative of it. More precisely, we show that worst-case NTRU is computationally equivalent to worst-case mod-uSVP, and that worst-case and an appropriately defined average-case mod-uSVP are also computationally equivalent, provided we have an oracle for id-SVP in both cases (and up to reduction losses). Together, these results imply that worst-case mod-uSVP reduces to average-case NTRU, provided we have an oracle for id-SVP. Combining this result with the reduction from worst-case id-SVP to worst-case NTRU from [PS21], this also implies that worst-case NTRU is computationally equivalent to worst-case mod-uSVP, without an id-SVP oracle.

Our first result is a collection of four reductions from the four variants of mod-uSVP (average case vs worst-case and vector vs module) to the corresponding four variants of NTRU, relying on an approximate id-SVP oracle. We give below a simplified version of one of these reductions, in the special case of power-of-two cyclotomic fields. More details and the other reductions can be found in Theorem 4.1.

Theorem 1.1 (Simplified version of Theorem 4.1). *Let K be a power-of-two cyclotomic field of degree d. Let $\gamma_{\mathrm{SVP}}, \gamma^+, \gamma_{\mathrm{NTRU}} > 1$. For all $q \geq 2^d \cdot \mathrm{poly}(\gamma^+)$ and $\gamma^- \geq \mathrm{poly}(d) \cdot \gamma_{\mathrm{NTRU}} \cdot \sqrt{\gamma_{\mathrm{HSVP}}}$, (worst-case)* mod-uSVP$_{\mathrm{mod}}$ *with gap in $[\gamma^-, \gamma^+]$ reduces in polynomial time to (worst-case)* NTRU$_{\mathrm{mod}}$ *with modulus q and gap $\geq \gamma_{\mathrm{NTRU}}$ and (worst-case)* id-SVP *with approximation factor γ_{SVP}.*

More concretely, when starting from a mod-uSVP instance for which the shortest non-zero vectors are $\approx \gamma$ times smaller than the root determinant, the reduction produces an NTRU instance satisfying $\sqrt{q}/(\|f\| + \|g\|) \approx \gamma^{O(1)}$, up to factors depending on field invariants. This transformation can be used to derive a reduction from average-case mod-uSVP to average-case NTRU (where the NTRU distribution is induced by the mod-uSVP distribution) and a reduction from worst-case mod-uSVP to worst-case NTRU (and similarly for the variants searching a dense rank-1 submodule). To achieve this transformation, an id-SVP oracle is required to find non-zero vectors in ideals within a factor $\gamma^{O(1)}$ from optimal. Note that for cyclotomic fields, the algorithm from [CDW21] allows to implement the oracle in quantum polynomial time when $\gamma \approx 2^{\sqrt{d}}$. Note also that [PS21] showed a reduction from worst-case id-SVP to worst-case NTRU, which is compatible with the reduction from worst-case mod-uSVP to worst-case

NTRU (relying on an id-SVP oracle). Combining both, we then obtain a reduction from worst-case mod-uSVP to worst-case search NTRU which *does not* rely on an id-SVP oracle. A drawback of the reduction is that it results in an NTRU modulus q of the order of $\approx 2^d$, even for small gap parameters γ. The modulus can be decreased by allowing the reduction to be more costly. Using lattice reduction algorithms [Sch87], one can reach $q \approx \gamma^{O(1)} \cdot \beta^{O(d/\beta)}$ if allowing for a reduction that runs in time polynomial in d, 2^β, $\log \Delta_K$ and $\zeta_K(2)$ (where ζ_K refers to the Dedekind zeta function). The quantities $\log \Delta_K$ and $\zeta_K(2)$ depend on the number field, and may not be polynomially bounded in the field degree d. In our running example, we have $\log \Delta_K = O(d)$ and $\zeta_K(2) = O(1)$ (see [SS13]).

Second, we exhibit a random self-reducibility property for mod-uSVP$_{\text{mod}}$. More explicitly, we give a reduction from worst-case mod-uSVP$_{\text{mod}}$ for rank-2 modules to an average-case version of itself, whose instances can be sampled from efficiently. The reduction preserves the gap parameter γ, up to factors depending on field invariants, and runs in time polynomial in $\log \Delta_K$.

Theorem 1.2 (Simplified version of Theorem 6.1, under ERH). *Let K be a power-of-two cyclotomic field of degree d. For any gap* $\text{poly}(d) < \gamma \leq 2^{O(d)}$, *there exists an efficiently samplable distribution D_γ^{uSVP} over* uSVP *instances with gap $\geq \gamma$ such that worst-case* mod-uSVP$_{\text{mod}}$ *with gap $\geq \gamma' = \gamma \cdot \text{poly}(d)$ reduces in polynomial time to average-case* mod-uSVP$_{\text{mod}}$ *for instance distribution D_γ^{uSVP}.*

Combined with the first reduction, the above allows to map a worst-case instance of mod-uSVP$_{\text{mod}}$ to an average-case instance of NTRU$_{\text{mod}}$, where the NTRU$_{\text{mod}}$ instance distribution is inherited from the average-case mod-uSVP distribution. This reduction relies on an id-SVP oracle. Since mod-uSVP$_{\text{mod}}$ and mod-uSVP are computationally equivalent (up to polynomial losses) when we have an id-SVP oracle, this also provides a reduction from worst-case uSVP to average-case NTRU. Contrary to the reduction from worst-case uSVP to worst-case NTRU, we cannot use the result of [PS21] to get rid of the id-SVP oracle. This is because the average-case distribution of NTRU instances that is produced by our reduction may not be compatible with the one used in [PS21].

We summarize the known reductions between variants of mod-uSVP and NTRU in Figure 1. Note that the reductions may not be composable due to incompatible parameter restrictions or instance distributions.

Technical overview. The NTRU problem is a restriction of mod-uSVP modules with a basis of a specific shape. In general, a rank-2 module M is represented by a pseudo-basis, i.e., two vectors $(\mathbf{b}_1, \mathbf{b}_2)$ in K^2 and two ideals I_1, I_2 of $\mathcal{O}_K$ such that $M = \mathbf{b}_1 I_1 + \mathbf{b}_2 I_2$. When the two ideals I_1 and I_2 are both equal to $\mathcal{O}_K$, the pseudo-basis is a basis, and the module is said to be free (note that a free module is a module that has at least one basis, but not all of its pseudo-bases will satisfy $I_1 = I_2 = \mathcal{O}_K$). In the NTRU problem, the instance is a basis $(\mathbf{b}_1, \mathbf{b}_2)$ of a free module contained in $\mathcal{O}_K^2$, with $\mathbf{b}_1 = (1, h)^T$ for some $h \in \mathcal{O}_K$ and $\mathbf{b}_2 = (0, q)^T$ for some integer q which is a parameter of the NTRU problem. Hence, the only degree of freedom in this basis comes from the choice of h. The NTRU problem then asks to solve mod-uSVP in this very specific module.

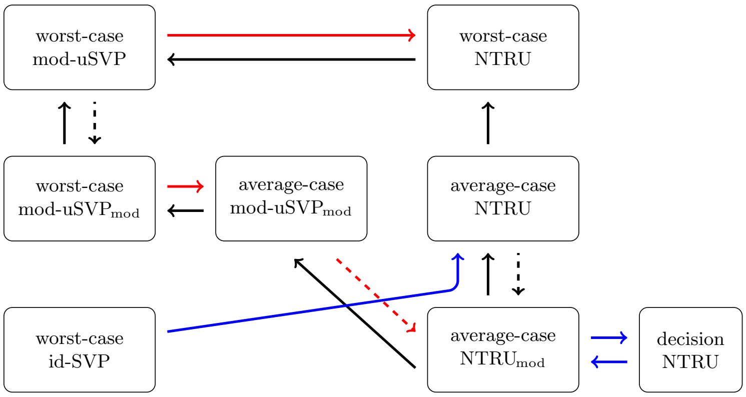

Fig. 1. Known reductions between NTRU and mod-uSVP variants. Dashed arrows require an id-SVP oracle. Blue arrows are proven in [PS21] and red arrows are proven in this article. The black arrows are folklore.

In the reduction from mod-uSVP to NTRU, we start with an arbitrary pseudo-basis of an arbitrary module M, and transform it into an NTRU basis. We then call the NTRU solver on this NTRU instance and lift the solution back to the original mod-uSVP module. In order to meaningfully lift a short vector (or a dense rank-1 submodule) back, we require our transformation to preserve the geometry of the rank-2 module M as much as possible. Our transformation proceeds in four main steps.

First of all, we transform the input module $M \subset K^2$ into an integral module whose volume is bounded from below and above by quantities depending only on the parameters of the reduction (NTRU modules are in $\mathcal{O}_K^2$ and have volume q^d). This is done by scaling M to the desired volume, and then rounding it to an integral module with a very close geometry. This rounding is performed by sampling two quasi-orthogonal vectors in the dual of M, and multiplying M on the left by the matrix whose rows are these two vectors. Multiplication on the left corresponds to a distortion of the ambient space, but since the two vectors are quasi orthogonal, this does not change the geometry too much. Also, as the row vectors of the sampled matrix belong to the dual of M, the resulting module is integral.

Our second step aims at obtaining the triangular shape of the NTRU basis. To do so, we compute the Hermite Normal Form of the pseudo-basis. With some probability, the two coefficients on the first row of the pseudo-basis will be coprime, leading to an HNF basis with a 1 as a top-left coefficient, exactly what we need for an NTRU instance. This is where $\zeta_K(2)$ comes into play, as it closely relates to the probability that two random elements of $\mathcal{O}_K$ are coprime.

At this point, our pseudo-basis still has coefficient ideals. We remove them with an id-SVP solver: we compute short x_1 and x_2 in the ideals I_1 and I_2, respectively, and then replace the pseudo-basis $((\mathbf{b}_1, \mathbf{b}_2), (I_1, I_2))$ by the basis $(x_1\mathbf{b}_1, x_2\mathbf{b}_2)$. This step has the effect of slightly sparsifying the module, i.e., it leads to a rank-2 submodule whose determinant is not much larger. If our gap is sufficiently large compared to the approximation factor of the id-SVP solver, our sparsified module will still contain an unexpectedly short non-zero vector.

We now have a basis of a free module with vectors of the form $(1, h')^T$ and $(0, b)^T$, with h' and b in $\mathcal{O}_K$. Our last step consists in replacing b by the NTRU parameter q. This is done by multiplying the second coordinates of both our basis vectors by q/b. If $q/b \approx 1$ (which we can ensure thanks to the id-SVP solver), then this does not change the geometry of the module too much.

To conclude, the transformation we have described allows us to transform any module of rank-2 with an unexpectedly short vector into an NTRU module with roughly the same geometry. The transformation is reversible, hence, we can lift any short vector or dense module found in the NTRU module back to the original rank-2 module. Since this transformation is a Karp reduction, it can be used to reduce average-case variants of mod-uSVP to average-case variants of NTRU where the NTRU distribution is inherited from the one on the uSVP instances.

For the random self-reducibility of mod-uSVP$_{\text{mod}}$, we start with an arbitrary rank-2 module M and want to randomize it so that the distribution of the output module M' does not depend on M. Once again, we design the transformation so that it preserves the geometry of the module, to be able to meaningfully lift any dense rank-1 submodule of M' back to a dense rank-1 submodule of M. For this reduction, we assume that all our worst-case modules live in $K_{\mathbb{R}}^2 = (K \otimes_{\mathbb{Q}} \mathbb{R})^2$ and have fixed volume (which we can always achieve by scaling the module). We also assume that the ℓ_2-norm of their shortest non-zero vectors is exactly $1/\gamma < 1$. This restriction to modules with a known gap can be waived, by guessing the gap and sparsifying the module (see Section 6).

Let us explain the main ideas behind the randomization in the simpler case of $K = \mathbb{Q}$. We have a lattice $M \subset \mathbb{R}^2$ with volume 1 and shortest non-zero vector $\mathbf{s}$ with $\|\mathbf{s}\| = 1/\gamma$. Up to rotation of the ambient space, we can assume that $\mathbf{s} = (1/\gamma, 0)^T$. Let us take $\mathbf{t} \in \mathbb{R}^2$ such that $(\mathbf{s}, \mathbf{t})$ forms a basis of M. Since the volume of M is 1, we know that $\mathbf{t} = (t_0, \gamma)^T$ for some $t_0 \in \mathbb{R}$. Up to the rotation of the ambient space, the quantity t_0 is the only degree of freedom. Note also that the lattice only depends on $t_0 \bmod 1/\gamma$. Let $\pi_{\mathbf{s}}(\mathbf{t})$ denote the quantity t_0, i.e., the norm of the orthogonal projection of $\mathbf{t}$ onto span($\mathbf{s}$). This discussion shows that the lattice M is uniquely determined by the span of its shortest non-zero vector and the quantity $\gamma \cdot \pi_{\mathbf{s}}(\mathbf{t}) \bmod 1$. Hence, to "hide" the lattice M, it suffices to "hide" these two quantities. Note that we use the vectors $\mathbf{s}$ and $\mathbf{t}$ for our reasoning, but we usually do not have access to them: we randomize our module by performing only operations that can be done on any of the bases of M (for $K_{\mathbb{R}}^2$ instead of $\mathbb{R}^2$, we expect that finding the analogue of $(\mathbf{s}, \mathbf{t})$ is difficult).

In order to hide the span of $\mathbf{s}$, one can apply a uniform orthonormal transformation to the ambient space. To hide the quantity $\gamma \cdot \pi_{\mathbf{s}}(\mathbf{t}) \bmod 1$, we "blur" the ambient space, by applying to it a transformation that is close to orthogonal, but not fully so. By appropriately choosing the transformation, one can obliviously transform the quantity $\gamma \cdot \pi_{\mathbf{s}}(\mathbf{t})$ into $x \cdot \gamma \cdot \pi_{\mathbf{s}}(\mathbf{t}) + y$, where x and y are some random variables. Recall that this quantity only matters modulo 1. Hence, if the standard deviation of y is sufficiently large compared to 1, then $y \bmod 1$ will be uniformly distributed and will hide the original value of $\pi_{\mathbf{s}}(\mathbf{t})$. The existence of a gap ensures that a close-to-orthogonal transformation suffices for this purpose.

This intuition over $\mathbb{R}^2$ explains one component of our randomization procedure, which we call the geometric randomization (see Section 5.2). Another important part of our randomization, which we call the coefficient randomization (Section 5.1), focuses on the coefficient ideals of the pseudo-basis (which are just $\mathbb{Z}$ for lattices). The transformation described above will have the effect of randomizing the vectors $\mathbf{b}_1$ and $\mathbf{b}_2$ of a pseudo-basis of our module M, but will have no impact on the coefficients ideals I_1 and I_2.

In order to hide those ideals, the first step is to multiply the module M by some uniformly distributed ideal I, using [dBDPW20]. Our new coefficient ideals $I \cdot I_1$ and $I \cdot I_2$ will then be uniformly distributed too. This is however not sufficient to fully hide the ideals, since the quotient $(I \cdot I_1)/(I \cdot I_2)$ is constant. In order to hide this last quantity, or decouple the ideals, we sparsify the module with respect to some prime ideal $\mathfrak{p}$: concretely, we take a uniformly random rank-2 submodule of M among those of index $\mathfrak{p}$.[5] This process generalizes lattice sparsification as introduced in [Kho06]. Lattice sparsification is a classic tool to remove one (or several) annoying vectors in a lattice. Here, the purpose is different: it has the effect of obliviously multiplying I_1 by $\mathfrak{p}$ while leaving I_2 unchanged (with probability close to 1). By [dBDPW20], the uniform distribution over bounded-norm prime ideals is close to the uniform distribution over norm-1 ideals (after renormalization of their norm), in the sense that little remains to be done to obtain the latter distribution. As a result, this sparsification enables us to (almost) randomize both I_1 and I_2, independently of one another. The gap to perfect randomization is handled by carefully studying the distribution resulting from the geometric and coefficient randomization (Section 5.3).

Summing up, our randomization consists in two main steps: a distortion of the ambient space, which randomizes the vectors $(\mathbf{b}_1, \mathbf{b}_2)$ and a sparsification, which hides the coefficient ideals I_1 and I_2 (together with the multiplication of the module by a random ideal I). Interestingly, we note that these two operations are similar (though adapted to rank-2 modules) to the ones that were used in [dBDPW20] to randomize ideal lattices.

The transformation described above allows us to transform an arbitrary module M of $K_{\mathbb{R}}^2$ into a random module M' of $K_{\mathbb{R}}^2$ whose distribution is independent of the input module. One last subtlety to handle in order to have a full worst-case to average-case reduction is to compute a canonical representation of the module M'. Indeed, the pseudo-basis of the properly distributed module M' that we have at the end of the randomization procedure might leak information about the input module M. Unfortunately, one cannot compute HNF bases in $K_{\mathbb{R}}^2$ (the HNF gives a canonical representation of rational lattices). In order to obtain a

[5] For two rank-2 modules $M' \subseteq M$ with pseudo-bases $((\mathbf{b}'_1, I'_1), (\mathbf{b}'_2, I'_2))$ and $((\mathbf{b}_1, I_1), (\mathbf{b}_2, I_2))$ respectively, we say that M' has index $\mathfrak{p}$ in M if $\det_K(\mathbf{b}'_1, \mathbf{b}'_2) \cdot I'_1 I'_2 = \mathfrak{p} \cdot \det_K(\mathbf{b}_1, \mathbf{b}_2) \cdot I_1 I_2$.

canonical representation of M', we then round it to a close module in $\mathcal{O}_K^2$ for which we will be able to compute an HNF pseudo-basis. The rounding procedure is the same as the one described in the reduction from uSVP to NTRU, and the distribution of the output pseudo-basis only depends on the input module and not on the specific pseudo-basis that is provided to represent it.

Discussion. A question arising from our reduction concerns the possibility to sample an NTRU instance from the distribution obtained at the end of the reduction, together with a short secret vector of the corresponding NTRU module. The difficulty stems from the fact that the output NTRU distribution we obtain after the reduction is not easy to describe, except as "the distribution obtained by running the reduction". The same difficulty also appeared in [PS21], where it was tackled by running the reduction to sample from the average-case NTRU distribution (and keeping in mind some quantities generated during the reduction in order to create a short vector of the output NTRU module). In our case, we face two additional difficulties when trying to apply the same strategy. First, we note that even sampling from the NTRU distribution, without asking for a short vector of the corresponding module, does not seem straightfoward. Since our mod-uSVP to NTRU reduction requires an id-SVP solver and takes subexponential time if one wants to reach small NTRU modulus q, it does not provide an efficient sampling algorithm for our final NTRU distribution. Secondly, our reduction allows us to lift a short vector from the NTRU module back to the uSVP module, but it is not so clear whether the converse is also possible (i.e., starting with a known vector of the uSVP module and obtaining a short vector of the final NTRU module). This is because of the sparsification step: when we sparsify a lattice, we can lift a vector from the sparser lattice back to the denser lattice (this is actually the same vector), but the converse seems more difficult.

Another question we leave open is about the compatibility of our reduction with those from [PS21]. Our worst-case mod-uSVP$_{\mathrm{mod}}$ to average-case NTRU$_{\mathrm{mod}}$ reduction produces a new distribution over NTRU instances. It is unclear whether this distribution can be used in the search to decision reduction from [PS21]. It is also unclear how it compares to the one produced by the worst-case id-SVP to average-case NTRU reduction from [PS21].

It should be noted that the regime where NTRU is provably secure (see [SS13]) is completely distinct from the regime required by our reductions. Indeed, the regime of [SS13] requires that f and g are slightly larger than $\sqrt{q}$, whereas our reduction requires f and g to be significantly smaller than $\sqrt{q}$. In other words, we are in a regime where NTRU is a uSVP instance (and we are trying to show that in this regime, it is representative of all uSVP instances), whereas [SS13] works in a regime where an NTRU instance is statistically close to uniform; in particular, in that regime, the underlying lattice is not a uSVP instance. The regime of the overstretch-NTRU attacks (including [KF17]) is also distinct from ours, but in the opposite direction. In these attacks, it is assumed that $\|f\|$ and $\|g\|$ are $\mathrm{poly}(d)$ and q grows; whereas in our case, we have $\|f\|$ and $\|g\|$ of the form $\sqrt{q}/\mathrm{poly}(d)$. Said differently, in those attacks, the short vector is short in absolute terms, whereas in our case it is short relative to what it would be for a random lattice of the same volume. We leave as an open problem to check whether these two regimes can be made to intersect.

2 Preliminaries

We use standard Landau notations, with underlying constants that are absolute (e.g., they do not depend on the specfic choice of number field). We consider column vectors (unless they are explicitly transposed). Vectors and matrices are respectively denoted in bold lowercase and uppercase fonts. For a vector $\mathbf{x} \in \mathbb{C}^k$, we let $\|\mathbf{x}\|$ denote its Hermitian norm.

We let $\mathcal{D}(c, s)$ refer to the normal distribution over $\mathbb{R}$ of center c and standard deviation $s > 0$. For X a set that is finite or has finite Lebesgue measure, we let $\mathcal{U}(X)$ denote the uniform distribution over X. For two distributions D_1, D_2 with compatible supports, we let $\mathrm{SD}(D_1, D_2) = \int |D_1(t) - D_2(t)| \, \mathrm{d}t/2$ refer to their statistical distance. For D_1, D_2 with $\mathrm{Supp}(D_1) \subseteq \mathrm{Supp}(D_2)$, we let $\mathrm{RD}(D_1 \parallel D_2) = \int D_1(t)^2/D_2(t) \, \mathrm{d}t$ refer to their Rényi divergence of order 2. The probability preservation property states that for any event E, the inequality $D_1(E) \geq D_2(E)^2/\mathrm{RD}(D_1 \parallel D_2)$ holds.

For a lattice L, we let $D_{L,s,\mathbf{c}}$ denote the Gaussian distribution of support L, standard deviation parameter s and center parameter $\mathbf{c} \in \operatorname{span} L$. We will use the following lemma, to sample discrete (tail-cut) Gaussian distributions. This lemma is adapted from [GPV08, Theorem 4.1]. A proof of this precise formulation can be found in [PS21, Lemma 2.2].

Lemma 2.1. *There exists a polynomial time algorithm that takes as input a basis* $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ *of an* n*-dimensional lattice* L*, a parameter* $s \geq \sqrt{n} \cdot \max_i \|\mathbf{b}_i\|$ *and a center* $\mathbf{c} \in \operatorname{span} L$ *and outputs a sample from a distribution* $\hat{D}_{\mathbf{B},s,\mathbf{c}}$ *such that*

- $\mathrm{SD}(D_{L,s,\mathbf{c}}, \hat{D}_{\mathbf{B},s,\mathbf{c}}) \leq 2^{-\Omega(n)}$;
- *for all* $\mathbf{v} \leftarrow \hat{D}_{\mathbf{B},s,\mathbf{c}}$*, it holds that* $\|\mathbf{v} - \mathbf{c}\| \leq \sqrt{n} \cdot s$.

Some results are obtained under the Extended Riemann Hypothesis (ERH).

2.1 Number Fields

Let K be a number field of degree $d \geq 2$ and ring of integers $\mathcal{O}_K$. Let $K_{\mathbb{R}} = K \otimes_{\mathbb{Q}} \mathbb{R}$. We identify any element of K with its canonical embedding vector $\sigma : x \mapsto (\sigma_1(x), \cdots, \sigma_d(x))^T \in \mathbb{C}^d$. This leads to an identification of $K_{\mathbb{R}}$ with $\{\mathbf{y} \in \mathbb{C}^d : \forall i \in [r_1], y_i \in \mathbb{R} \text{ and } \forall i \in [r_2], \overline{y_{r_1+r_2+i}} = y_{r_1+i}\}$, where r_1 and r_2 respectively denote the number of real and pairs of complex embeddings. Note that the set $K_{\mathbb{R}}$ is a real vector subspace of dimension d embedded (via σ) in $\mathbb{C}^d$ and that $\sigma(\mathcal{O}_K)$ is a full rank lattice in $K_{\mathbb{R}}$. The (absolute) discriminant Δ_K is defined as $\Delta_K = |\det(\sigma(\mathcal{O}_K))^2|$. We have $d = O(\log \Delta_K)$, for Δ_K growing to infinity.

For $x \in K_{\mathbb{R}}$, we define $\overline{x} \in K_{\mathbb{R}}$ as the element obtained by componentwise complex conjugation of the canonical embedding vector of x. We extend this notation to vectors and matrices over $K_{\mathbb{R}}$, and let $\mathbf{x}^\dagger$ denote $\overline{\mathbf{x}}^T$ for any $\mathbf{x} \in K_{\mathbb{R}}^n$. We define $\overline{K}$ and $\overline{\mathcal{O}_K}$ as the subsets of $K_{\mathbb{R}}$ obtained by applying complex conjugation to elements of K and $\mathcal{O}_K$, respectively. For $\mathbf{x}, \mathbf{y} \in K_{\mathbb{R}}^n$, we define $\langle \mathbf{x}, \mathbf{y} \rangle_{K_{\mathbb{R}}} = \mathbf{x}^\dagger \cdot \mathbf{y} \in K_{\mathbb{R}}$ and $\|\mathbf{x}\| = \|\sigma(\langle \mathbf{x}, \mathbf{x} \rangle_{K_{\mathbb{R}}})\|^{1/2}$. The (absolute value of the) algebraic norm of $x \in K_{\mathbb{R}}$ is defined as $\mathcal{N}(x) = \prod_i |\sigma_i(x)|$. The algebraic norm of $\mathbf{x} \in K_{\mathbb{R}}^n$ is defined as $\mathcal{N}(\mathbf{x}) = \mathcal{N}(\langle \mathbf{x}, \mathbf{x} \rangle_{K_{\mathbb{R}}})^{1/2}$.

We define $K_{\mathbb{R}}^+$ as the subset of $K_{\mathbb{R}}$ corresponding to having all y_i's being positive real numbers. For $x \in K_{\mathbb{R}}^+$, we define $x^{1/2}$ as the element of $K_{\mathbb{R}}^+$ obtained by taking the square-roots of the embeddings.

We let $\mathcal{O}_K^\times = \{x \in \mathcal{O}_K : \mathcal{N}(x) = 1\}$ denote the set of units of $\mathcal{O}_K$ and $\operatorname{Log} \mathcal{O}_K^\times = \{(\log |\sigma_i(x)|)_i : x \in \mathcal{O}_K^\times\} \subset \mathbb{R}^d$ denote the log-unit lattice. Note that $\operatorname{span}_{\mathbb{R}}(\operatorname{Log} \mathcal{O}_K^\times) = E := \{\mathbf{y} \in \mathbb{R}^d : \sum y_i = 0 \wedge \forall i \in [r_2], y_{r_1+r_2+i} = y_{r_1+i}\}$, by Dirichlet's unit theorem. For $\zeta \in E$, we define $\exp(\zeta)$ as the element of $K_{\mathbb{R}}^+$ whose i-th embedding is $\exp(\zeta_i)$, for all i.

In this work, we assume that we know a LLL-reduced [LLL82] $\mathbb{Z}$-basis $(r_i)_{i \leq d}$ of $\mathcal{O}_K$. We define $\delta_K = \max_i \|r_i\|_\infty$. We have $1 \leq \delta_K \leq \Delta_K^{O(1)}$: the left inequality follows from the fact that $\|r\|_\infty \geq 1$ for all $r \in \mathcal{O}_K \setminus \{0\}$, whereas the right inequality derives from Minkowski's second theorem and the LLL-reducedness of the r_i's. In the case of cyclotomic number fields, taking the power basis gives $\delta_K = 1$. For $x = \sum_i x_i r_i \in K_{\mathbb{R}}$, we define $\lfloor x \rceil = \sum_i \lfloor x_i \rceil r_i$. We will use the notation $\{x\} = x - \lfloor x \rceil$. We have $\|\{x\}\|_\infty \leq d \cdot \delta_K$, and hence $\|\{x\}\| \leq d^{3/2} \cdot \delta_K$.

We will consider the following distributions over $K_{\mathbb{R}}$. Note that for $r \in K_{\mathbb{R}}^+$, the distribution of $r \cdot x$ for $x \sim \mathcal{D}_{K_{\mathbb{R}}}(c, \mathbf{s})$ is $\mathcal{D}_{K_{\mathbb{R}}}(r \cdot c, (\sigma_i(r) \cdot \mathbf{s}_i)_i)$.

Definition 2.2. *Let $\mathbf{s} \in \mathbb{R}_{>0}^{r_1+r_2}$. We define the normal distribution $\mathcal{D}_{K_{\mathbb{R}}}(c, \mathbf{s})$ of center $c \in K_{\mathbb{R}}$ and standard deviation vector $\mathbf{s}$ as the distribution obtained by independently sampling real numbers $(y)_{i \in [d]}$ with*

$$\begin{cases} y_j \sim \mathcal{D}(0, s_j) & \text{for } j \in [r_1] \\ y_{r_1+j}, y_{r_1+r_2+j} \sim \mathcal{D}(0, s_{r_1+j}) & \text{for } j \in [r_2] \end{cases}$$

and then returning $c + y$ where $y \in K_{\mathbb{R}}$ is such that $\sigma_j(y) = y_j$ for $j \in [r_1]$ and $\sigma_{r_1+j}(y) = y_{r_1+j} + i y_{r_1+j}$ for $j \in [r_2]$.

We define $\chi_{K_{\mathbb{R}}}$ as the distribution of $(\langle \mathbf{x}, \mathbf{x} \rangle_{K_{\mathbb{R}}})^{1/2}$ for $\mathbf{x} \in K_{\mathbb{R}}^2$ sampled according to $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)^2$.

For a matrix $\mathbf{B} \in K_{\mathbb{R}}^{n \times n}$, we define $\det(\mathbf{B}) = \mathcal{N}(\det_{K_{\mathbb{R}}}(\mathbf{B}))$. We say that $\mathbf{B}$ is orthogonal if $\mathbf{B}^\dagger \cdot \mathbf{B} = \mathbf{I}$, which implies that $\det(\mathbf{B}) = 1$. We let $\mathcal{O}_n(K_{\mathbb{R}})$ denote the set of orthogonal matrices. If a matrix $\mathbf{B} \in K_{\mathbb{R}}^{n \times n}$ has $K_{\mathbb{R}}$-linearly independent columns (i.e., no non-trivial linear combination is zero), then it admits a QR-factorization $\mathbf{B} = \mathbf{QR}$ with $\mathbf{Q} \in \mathcal{O}_n(K_{\mathbb{R}})$ and $\mathbf{R} \in K_{\mathbb{R}}^{n \times n}$ upper triangular with diagonal elements in $K_{\mathbb{R}}^+$ (see, e.g., [LPSW19, Section 2.3]).

2.2 Ideals

A fractional ideal (resp. oriented replete ideal) is a subset of K of the form $x \cdot I$ for some $x \in K^\times$ (resp. $x \in K_\mathbb{R}^\times$) and $I \subseteq \mathcal{O}_K$ an integral ideal. Unless specified otherwise, by default, an ideal will refer to an oriented replete ideal. For I ideal of K, we define the ideal $\overline{I} = \{\overline{x} : x \in I\}$ of $\overline{K}$. Using the canonical embedding, any non-zero ideal is identified to a d-dimensional lattice, called ideal lattice. The algebraic norm of an integral ideal I is $\mathcal{N}(I) := |\mathcal{O}_K/I|$ if it is non-zero and zero otherwise. This is extended to oriented replete ideals xI with $x \in K_\mathbb{R}^\times$ and I an integral ideal by setting $\mathcal{N}(xI) = \mathcal{N}(x) \cdot \mathcal{N}(I)$.

For I_1 and I_2 integral, the product ideal $I_1 I_2$ is the ideal spanned by all $x_1 \cdot x_2$ with $x_1 \in I_1$ and $x_2 \in I_2$. An integral ideal I is said prime if it cannot be written as $I = I_1 \cdot I_2$ with I_1, I_2 integral and both distinct from $\mathcal{O}_K$. For any $B \geq 0$, we let $\pi_K(B)$ denote the number of prime ideals with algebraic norm $\leq B$. Under the ERH, there exists an absolute constant c such that for any $B \geq (\log \Delta_K)^c$, we have $\pi_K(B) \in (B/\log B) \cdot [0.9, 1.1]$ (see [BS96, Theorem 8.7.4]). If $x_1 I_1$ and $x_2 I_2$ are two ideals with I_1 and I_2 integral, we define their product as $(x_1 I_1) \cdot (x_2 I_2) = (x_1 x_2)(I_1 I_2)$. The inverse of an ideal I is $I^{-1} = \{x \in K_\mathbb{R}^\times : xI \subseteq \mathcal{O}_K\}$.

We will use algorithms from [dBDPW20] to sample among different classes of ideals.

Lemma 2.3 (Adapted from [dBDPW20, Lemma 2.2], ERH). *There exists an algorithm $\mathcal{A}$ and an absolute constant c such that for any $B \geq (\log \Delta_K)^c$, algorithm $\mathcal{A}$ on input B runs in time $\mathrm{poly}(\log B, d)$ and returns a prime ideal uniformly among prime ideals of norm $\leq B$.*

We will also rely on Algorithm 2.1, which is adapted from [dBDPW20, Theorem 3.3], to sample (essentially) uniformly in the set $\mathcal{I}_1$ of norm-1 ideals, in time polynomial in $\log B$. Note that [dBDPW20] considers norm-1 ideals xI with I integral and all $\sigma_i(x)$'s being positive integers. This discrepancy is handled by introducing u at Step 3. The standard deviation in Step 2 and tailcut may seem arbitrary at first sight: these choices simplify the analysis of the module randomization (in Section 5.3). A proof of the following lemma is given in Appendix B.

Algorithm 2.1 `Ideal-Sample`$_B$

1: Sample $\mathfrak{p}$ uniformly among prime ideals of norms $\leq B$, using Lemma 2.3;
2: Sample $\zeta \in E$ from the centered normal law with standard deviation $d^{-3/2}$, conditioned on $\|\zeta\| \leq 1/d$;
3: Sample u uniform in $\{x \in K_\mathbb{R}, \forall i \in [d] : |\sigma_i(x)| = 1\}$;
4: Return $u \cdot \exp(\zeta) \cdot \mathfrak{p}/\mathcal{N}^{1/d}(\mathfrak{p})$.

Lemma 2.4 (Adapted from [dBDPW20, Theorem 3.3], ERH). *There exists an absolute constant c such that for any $B \geq (d^d \Delta_k)^c$,* `Ideal-Sample`$_B$ *runs in time polynomial in $\log B$ and its output distribution is within $2^{-\Omega(d)}$ statistical distance from $\mathcal{U}(\mathcal{I}_1)$.*

2.3 Modules

A module is a subset of some $K_{\mathbb{R}}^m$ of the form $M = \sum_{i \leq k} \mathbf{b}_i I_i$ where the I_i's are non-zero ideals and the $\mathbf{b}_i$'s are $K_{\mathbb{R}}$-linearly independent. This is written compactly as $M = \mathbf{B} \cdot \mathbb{I}$ (where $\mathbf{B}$ is the matrix whose columns are the $\mathbf{b}_i$ and $\mathbb{I} = (I_1, \dots, I_k)$). The tuple $((I_1, \mathbf{b}_1), \dots, (I_k, \mathbf{b}_k))$ is called a pseudo-basis of M and is written compactly as $(\mathbf{B}, \mathbb{I})$. The integer k is the rank of M. We define $\mathcal{N}(M) = \det(\mathbf{B}) \cdot \prod_{i \leq k} \mathcal{N}(I_i)$. Note that for $d = m = 1$, this matches the norm of an ideal. Using the canonical embedding, any rank-k module is identified to a (kd)-dimensional lattice, called module lattice. In particular, we define $\det(M)$ as the determinant of the module lattice. Note that $\det(M) = \mathcal{N}(M) \cdot \Delta_K^{k/2}$. The module successive minima $\lambda_i(M)$ for $i \in [kd]$ are defined similarly. We will also be interested in the module norm-minimum $\lambda_1^{\mathcal{N}}(M) = \inf\{\mathcal{N}(N) : N \text{ rank-1 submodule of } M\}$. A rank-1 submodule of M is said *densest* if it reaches $\lambda_1^{\mathcal{N}}(M)$.

The dual of a module M is defined as $M^\vee = \{\mathbf{b}^\vee \in \mathrm{span}_{K_{\mathbb{R}}}(M) : \forall \mathbf{b} \in M, \langle \mathbf{b}^\vee, \mathbf{b} \rangle_{K_{\mathbb{R}}} \in \mathcal{O}_K\}$: note that $M^\vee$ is an $\overline{\mathcal{O}_K}$-module, $\sigma(M^\vee)$ is the dual lattice of $\sigma(M)$ and $(\mathbf{B} \cdot \mathbb{I})^\vee = (\mathbf{B}^{-\dagger} \cdot \mathbb{J})$, where $J_i = (\overline{I_i})^{-1}$ for all $i \leq k$.

For any full-rank module $M \subseteq K^m$, there exists a pseudo-basis $(\mathbf{B}, \mathbb{I})$ such that $\mathbf{B} \in K^{m \times m}$ is lower-triangular with ones on the diagonal. It is called a Hermite Normal Form of M and can be computed in polynomial time from any finite set of pairs $\{(I_i, \mathbf{b}_i)\}_i$ such that $M = \sum_i \mathbf{b}_i I_i$ and the $\mathbf{b}_i$'s are not necessarily independent [BP91,Coh96,BFH17].

Definition 2.5. *Let M be a module. A submodule $N \subseteq M$ is said to be primitive if it satisfies any of the three equivalent conditions:*

- *the module N is maximal for the inclusion in the set of submodules of M of rank at most* $\mathrm{rank}(N)$*;*
- *there is a module N' with $M = N + N'$ and* $\mathrm{rank}(M) = \mathrm{rank}(N) + \mathrm{rank}(N')$*;*
- *we have $N = M \cap \mathrm{span}_K(N)$.*

In particular, any densest rank-1 submodule of M is primitive.

A proof that the three conditions are equivalent is provided in Appendix B. The last statement follows from Condition 1.

The latter lemma allows us to conclude that the module norm-minimum is reached (see Appendix B for a proof).

Lemma 2.6. *For any module M, there exists a rank-1 submodule N of M such that $\mathcal{N}(N) = \lambda_1^{\mathcal{N}}(M)$.*

The following result provides a lower bound on the probability that a rank-1 module $\mathbf{v} \cdot \mathcal{O}_K$ is primitive in a rank-k module M, when $\mathbf{v} \in M$ is sampled from a sufficiently wide Gaussian distribution. Taking $M = \mathcal{O}_K^k$, this provides in particular a lower bound on the probability that k elements sampled independently of a Gaussian distribution in $\mathcal{O}_K$ are relatively coprime. This result generalizes [SS13, Lemma 4.4], which proved the result for $k = 2$ and $M = \mathcal{O}_K^2$

(with a proof inspired from [Sit10]). The proof for the general case with rank-k modules is very similar to the special case $M = \mathcal{O}_K^2$, hence we postpone it to Appendix B.4. In this work, we will only use Lemma 2.7 for modules of rank-2, however, for the sake of re-usability, we state and prove it for modules of arbitrary ranks.

Lemma 2.7. *There exists an absolute polynomial P such that the following holds. For any $\delta \geq 0$, degree-d number field K, integer $k \geq 2$, rank-k module $M \subset K_{\mathbb{R}}^k$, if $\mathbf{c} \in \operatorname{span}_{K_{\mathbb{R}}}(M)$ and $\varsigma > 0$ are such that $\|\mathbf{c}\| \leq \delta \cdot \varsigma$ and $\varsigma \geq \lambda_{kd}(M) \cdot P(\Delta_K^{1/d}, k, d, \delta, \lambda_{kd}(M)/\lambda_1(M))$, then it holds that*

$$\Pr_{\mathbf{v} \leftarrow D_{M,\varsigma,\mathbf{c}}} \big(\mathbf{v} \cdot \mathcal{O}_K \text{ is primitive in } M\big) \geq \frac{1}{4\zeta_K(k)},$$

where $\zeta_K(\cdot)$ is the Dedekind zeta function of K and the λ_i's refer to the minima of the lattice $\sigma(M)$.

2.4 Rank-2 Modules with a Gap

In this work, we are interested in rank-2 modules that contain an unexpectedly dense rank-1 submodule, i.e., in modules M with $\lambda_1^{\mathcal{N}}(M)$ significantly smaller than $\sqrt{\mathcal{N}(M)}$. We define the gap of M by

$$\gamma(M) = \left(\frac{\mathcal{N}(M)^{\frac{1}{2}}}{\lambda_1^{\mathcal{N}}(M)}\right)^{\frac{1}{d}}.$$

The following lemma shows that if the gap is sufficiently large, then the densest rank-1 submodule is unique. A proof may be found in Appendix B.

Lemma 2.8. *Let M be a rank-2 module with gap $\gamma > 0$ and N a densest rank-1 submodule of M. If N' is a rank-1 submodule of M with $\mathcal{N}(N') < \gamma^d \sqrt{\mathcal{N}(M)}$, then $N' \subseteq N$.*

In particular, for $\gamma > 1$, the densest rank-1 submodule is unique and any vector $\mathbf{b} \in M$ with $\|\mathbf{b}\| < \gamma \cdot \mathcal{N}(M)^{1/(2d)}$ belongs to it.

In the following, when a rank-2 module M has a gap larger than 1, we will implicitly use Lemma 2.8 when referring to the densest rank-1 submodule of M. Most rank-2 modules we will consider will have gap larger than 1.

This can be used to show that we can use the QR-factorization to precisely describe rank-2 modules (see Appendix B for a proof).

Lemma 2.9. *Let M be a rank-2 module with gap $\gamma > 0$. Then M can be written as*

$$\frac{\mathcal{N}^{\frac{1}{2d}}(M)}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right),$$

where $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}})$, $r \in K_{\mathbb{R}}$, J_1 and J_2 are norm-1 ideals. We call this a QR-standard-form for M.

We note that there are multiple QR-standard forms for any module M, as units of $\mathbb{C}$ can be transferred from the ideal coefficients to the matrix $\mathbf{Q}$. In the following section, we will be interested in modules with specific distributions expressed in terms of QR-standard forms. It will then be convenient to define a module by a (well-distributed) QR-standard form. Note that the modules we define this way have norm 1.

Definition 2.10. *For any $\mathbf{Q} \in \mathcal{O}_2(K_\mathbb{R}), \gamma > 0, r \in K_\mathbb{R}$ and norm-1 ideals J_1, J_2, we define*

$$\texttt{QRSF-2-Mod}(\mathbf{Q}, \gamma, J_1, J_2, r) = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right).$$

We will use the following result on the first and last minimum of the dual of a rank-2 module with a gap. The proof is provided in Appendix B.

Lemma 2.11. *Let M be a rank-2 module in $K_\mathbb{R}^2$ with gap $\gamma(M) \geq 1$. Then*

$$\lambda_{2d}(M^\vee) \leq 2\sqrt{d} \cdot \gamma(M) \cdot \mathcal{N}(M)^{-\frac{1}{2d}}$$
$$\lambda_1(M^\vee)^{-1} \leq 2d \cdot \gamma(M) \cdot \mathcal{N}(M)^{1/(2d)} \cdot \delta_K \cdot \Delta_K^{\frac{1}{2d}}.$$

2.5 Algorithmic Problems

In this section, we define different variants of the unique-SVP problem for rank-2 modules, as well as variants of the NTRU problem. The definitions of the different NTRU problems differ slightly from the ones defined in [PS21]: this is to emphasize the resemblance between uSVP and NTRU. The difference between the NTRU definitions in this work and the ones in [PS21] are sufficiently minor that they can be reduced to one another without difficulty, and we hence opted to keep the same names.

Definition 2.12 (γ-uSVP instance). *Let $\gamma > 0$. A γ-uSVP instance consists in a pseudo-basis $(\mathbf{B}, \mathbb{I})$ of a rank-2 module $M \subset K^2$ such that M contains a non-zero vector $\mathbf{s}$ with $\|\mathbf{s}\| \leq \gamma^{-1} \cdot \mathcal{N}(M)^{1/(2d)}$.*

Note that any module M associated to a γ-uSVP instance contains the rank-1 submodule $\mathbf{s}\mathcal{O}_K$ whose norm is $\leq \sqrt{\mathcal{N}(M)}/\gamma^d$. By Lemma 2.8, this implies that if $\gamma > 1$, then the module M has a unique densest rank-1 submodule.

Definition 2.13 ($(\mathcal{D}, \gamma, \gamma')$-uSVP$_{\text{vec}}$ and (γ, γ')-wc-uSVP$_{\text{vec}}$). *Let $\gamma \geq \gamma' > 0$ and $\mathcal{D}$ a distribution over γ-uSVP instances. The $(\mathcal{D}, \gamma, \gamma')$ average-case unique SVP problem for rank-2 modules ($(\mathcal{D}, \gamma, \gamma')$-uSVP$_{\text{vec}}$ for short) asks, given as input a pseudo-basis of some rank-2 module M sampled from $\mathcal{D}$, to compute a vector $\mathbf{s} \in M \setminus \{\mathbf{0}\}$ such that $\|\mathbf{s}\| \leq \mathcal{N}(M)^{1/(2d)}/\gamma'$. The advantage of an algorithm $\mathcal{A}$ against the $(\mathcal{D}, \gamma, \gamma')$-uSVP$_{\text{vec}}$ problem is defined as*

$$\text{Adv}(\mathcal{A}) = \Pr_{(\mathbf{B},\mathbb{I}) \leftarrow \mathcal{D}} \left(\mathcal{A}((\mathbf{B}, \mathbb{I})) = \mathbf{s} \text{ with } \left| \begin{array}{l} \mathbf{s} \in M \setminus \{\mathbf{0}\} \\ \|\mathbf{s}\| \leq \mathcal{N}(M)^{1/(2d)}/\gamma' \end{array} \right. \right),$$

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

The worst-case variant ((γ,γ')-wc-uSVP$_{\text{vec}}$) asks to solve this problem for any γ-uSVP instance $(\mathbf{B},\mathbb{I})$.

Definition 2.14 ($(\mathcal{D},\gamma)$-uSVP$_{\text{mod}}$ and γ-wc-uSVP$_{\text{mod}}$). *Let $\gamma > 0$ and $\mathcal{D}$ a distribution over γ-uSVP instances. The $(\mathcal{D},\gamma)$ unique SVP problem for rank-2 modules ($(\mathcal{D},\gamma)$-uSVP$_{\text{mod}}$ for short) asks, given as input a γ-uSVP module M sampled from $\mathcal{D}$, to recover a densest rank-1 submodule $N \subset M$. The advantage of an algorithm $\mathcal{A}$ against the $(\mathcal{D},\gamma)$-uSVP$_{\text{mod}}$ problem is defined as*

$$\mathrm{Adv}(\mathcal{A}) = \Pr_{(\mathbf{B},\mathbb{I})\leftarrow\mathcal{D}}\left(\mathcal{A}((\mathbf{B},\mathbb{I})) = N \text{ with } \middle| \begin{array}{l} N \subset M \text{ with } \mathrm{rk}(N) = 1 \\ \mathcal{N}(N) = \lambda_1^{\mathcal{N}}(M) \end{array}\right),$$

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

The worst-case variant (γ-wc-uSVP$_{\text{mod}}$) asks to solve this problem for any γ-uSVP instance $(\mathbf{B},\mathbb{I})$.

We can now define the NTRU problems, as special cases of the uSVP variants above.

Definition 2.15 (NTRU instance). *Let $q \geq 2$ be an integer, and $\gamma > 0$ a real number. A (γ,q)-NTRU instance is a γ-uSVP instance whose pseudo-basis is required to be of the form $((\mathbf{b}_1,\mathcal{O}_K),(\mathbf{b}_2,\mathcal{O}_K))$ with $\mathbf{b}_1 = (1,h)^T$ for some $h \in \mathcal{O}_K$ and $\mathbf{b}_2 = (0,q)^T$.*

Comparison with [PS21]. In [PS21], an NTRU instance consists in the single element $h \in R_q$, whereas we consider it as a basis of a rank-2 module in this work. Both formalisms are equivalent, since one can reconstruct the basis of the rank-2 module from h (and also q, which is a parameter of the problem). A second difference comes from the fact that [PS21] requires the short vector $\mathbf{s} = (s_1,s_2)^T$ to satisfy $\|s_1\|,\|s_2\| \leq \sqrt{q}/\gamma$, whereas we require that $\|\mathbf{s}\| \leq \sqrt{q}/\gamma$. This means that a (γ,q)-NTRU instance for us is a (γ,q)-NTRU instance for [PS21], but the converse does not hold: a (γ,q)-NTRU instance for [PS21] is only guaranteed to be a $(\sqrt{2}\cdot\gamma,q)$-NTRU instance for us.

Definition 2.16 (NTRU problems). *Let $q \geq 2$, $\gamma \geq \gamma' > 0$ and $\mathcal{D}$ a distribution over (γ,q)-NTRU instances. The $(\mathcal{D},\gamma,\gamma',q)$-NTRU$_{\text{vec}}$ problem, (γ,γ',q)-wc-NTRU$_{\text{vec}}$ problem, $(\mathcal{D},\gamma,q)$-NTRU$_{\text{mod}}$ problem and (γ,q)-wc-NTRU$_{\text{mod}}$ problem are the restrictions of the uSVP problems to (γ,q)-NTRU instances.*

From the definitions of the NTRU and uSVP problems, one can see that the average case NTRU$_{\text{vec}}$ and NTRU$_{\text{mod}}$ problems reduce to the worst-case uSVP$_{\text{vec}}$ and uSVP$_{\text{mod}}$ problems. In the next sections, we will show that the converse also holds, provided we have an oracle solving ideal-SVP.

Finally, we also recall the definition of the Hermite shortest vector problem in ideal lattices (id-HSVP).

Definition 2.17 (γ-id-HSVP)**.** *Let $\gamma \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$. Given as input a fractional ideal $I \subset K$, the γ-id-HSVP problem asks to find an element $x \in I \setminus \{0\}$ such that $\|x\| \leq \gamma \cdot \mathcal{N}(I)^{1/d}$.*

By Minkowski's theorem, this problem is well-defined for any $\gamma \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$.

3 New Tools on Module Lattices

In this section, we present new tools to manipulate module lattices. For the sake of re-usability, we describe them for modules of arbitrary ranks, but we will use them only in rank 2 in the reductions of the present work. The missing proofs of this section are available in Appendix C.

3.1 Module Sparsification

An essential ingredient in the module randomization of Section 5 is sparsification. In this subsection, we extend to modules the definition and some properties of sparsification over lattices [Kho06].

Definition 3.1. *Let M a module, $\mathfrak{p}$ a prime ideal, $\overline{\mathbf{b}^\vee} \in (M^\vee / \mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ and $\mathbf{b}^\vee$ a lift of $\overline{\mathbf{b}^\vee}$ in $M^\vee$. The sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$ is the submodule*

$$M' = \left\{ \mathbf{m} \in M,\ \langle \mathbf{b}^\vee, \mathbf{m} \rangle_{K_\mathbb{R}} \in \mathfrak{p} \right\}.$$

The submodule M' does not depend on the choice of the vector $\mathbf{b}^\vee$ lifting $\overline{\mathbf{b}^\vee}$.

Note that $M \subseteq M' \subseteq \mathfrak{p}M$, implying that M' has the same rank as M. As showed by the following two lemmas, sparsification increases the module norm by a factor $\mathcal{N}(\mathfrak{p})$ and an arbitrary rank-1 submodule of M is not contained in M' (except with probability $\leq 1/\mathcal{N}(\mathfrak{p})$).

Lemma 3.2. *Let M a module, $\mathfrak{p}$ a prime ideal and $\overline{\mathbf{b}^\vee} \in (M^\vee / \mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$. Let M' be the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$. Then $\mathcal{N}(M') = \mathcal{N}(\mathfrak{p}) \cdot \mathcal{N}(M)$.*

Lemma 3.3. *Let M a rank-k module, $\mathfrak{p}$ a prime ideal and $\mathbf{b}I$ a primitive rank-1 submodule of M. Let $\overline{\mathbf{b}^\vee}$ be uniformly distributed in $(M^\vee / \mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ and M' be the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$. Then $\mathbf{b}\mathfrak{p}I \subseteq M'$ and, except with probability $1/\mathcal{N}(\mathfrak{p}) - 1/\mathcal{N}(\mathfrak{p})^k$, we have $\mathbf{b}I \not\subset M'$.*

The following lemma states that a module sparsification can be efficiently computed. The algorithm generalizes the one for lattice sparsification, detailed, e.g., in [BSW16].

Lemma 3.4. *There exists a polynomial-time algorithm taking as inputs an arbitrary pseudo-basis of $M \subset K_\mathbb{R}^k$, a prime ideal $\mathfrak{p}$ and $\overline{\mathbf{b}^\vee} \in (M^\vee / \mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ and computing a pseudo-basis of the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$.*

3.2 Module Rounding

In this section, we describe the `DualRound` algorithm that rounds a rank-k module contained in $K_{\mathbb{R}}^k$ into a module contained in $\mathcal{O}_K^k$ (with a close geometry), in a way that does not depend on how the module in $K_{\mathbb{R}}^k$ was represented. We do that by sampling almost orthogonal vectors in the dual lattice, in a similar fashion to what was done in [dBDPW20] in the context of ideal lattices. We believe that this technique of rounding via the dual might have other applications, especially in situations where one would like to have the analogue of an HNF basis for lattices with real coefficients.

`DualRound` is parameterized by a standard deviation parameter $\varsigma > 0$, a BKZ block-size $\beta \in \{2, \ldots, kd\}$ and an error bound $\varepsilon > 0$. It starts by computing a short $\mathbb{Z}$-basis of $\mathbf{C}^\vee$, by using a provable variant of the BKZ algorithm [Sch87,HPS11,GN08,ALNS20]. This offers different runtime-quality trade-offs. It then uses the discrete Gaussian sampler from Lemma 2.1 with orthogonal center parameters $\mathbf{t}_i$.

Algorithm 3.1 Algorithm $\texttt{DualRound}_{\varsigma,\beta,\varepsilon}$

Input: A pseudo-basis $(\mathbf{B}, \mathbb{I})$ of a rank-k module $M \subset K_{\mathbb{R}}^k$.
1: Compute a $\mathbb{Z}$-basis of $M^\vee$;
2: Run BKZ with block-size β on it to obtain a new $\mathbb{Z}$-basis $\mathbf{C}^\vee$ of $M^\vee$;
3: Set $R = \varepsilon^{-1}\sqrt{kd}\varsigma$;
4: For $i \in [k]$, set $\mathbf{t}_i = R \cdot \mathbf{e}_i$, where $\mathbf{e}_i$ is the i-th canonical unit vector of $K_{\mathbb{R}}^k$;
5: For $i \in [k]$, sample $\mathbf{y}_i \leftarrow \hat{D}_{\mathbf{C}^\vee,\varsigma,\mathbf{t}_i}$;
6: Return $\mathbf{Y} = (\mathbf{y}_1 | \ldots | \mathbf{y}_k)^\dagger$.

Lemma 3.5. *Let $(\mathbf{B}, \mathbb{I})$ be a pseudo-basis of a rank-k module $M \subset K_{\mathbb{R}}^k$. Let $\beta \in \{2, \cdots, kd\}$, $\varepsilon > 0$, and ς be such that $\varsigma \geq (kd)^{kd/\beta+3/2} \cdot \lambda_{kd}(M^\vee)$. Algorithm* `DualRound` *runs in time polynomial in $2^\beta, \log(\varsigma/\varepsilon)$ and the bitsize of its input. Further, on input $(\mathbf{B}, \mathbb{I})$, $\texttt{DualRound}_{\varsigma,\beta,\varepsilon}$ outputs a matrix $\mathbf{Y} \in \mathrm{M}_k(K_{\mathbb{R}})$ such that*

- *$(\mathbf{Y} \cdot \mathbf{B}) \cdot \mathbb{I}$ is contained in $\mathcal{O}_K^k$;*
- *$\mathbf{Y} = R \cdot \mathbf{I}_k + \mathbf{E}$ for $R = \varepsilon^{-1}\sqrt{kd}\varsigma > 0$ and $\|e_{ij}\| \leq \varepsilon R$ for all $i, j \in [k]$.*

Moreover, if $(\mathbf{B}', \mathbb{I}')$ is another pseudo-basis of M and if $\mathbf{Y}'$ is the output of `DualRound` *given this pseudo-basis as input, then*

$$\mathrm{SD}(\mathbf{Y}, \mathbf{Y}') \leq 2^{-\Omega(kd)}.$$

Note that Lemma 3.5 does not necessarily ensure that the matrix $\mathbf{Y}$ is invertible, hence the module $\mathbf{Y} \cdot \mathbf{B} \cdot \mathbb{I}$ might not be of rank k. However, by choosing ε sufficiently small and using the second condition on $\mathbf{Y}$, one can make sure that $\mathbf{Y}$ is indeed invertible. This is the purpose of Lemma 3.6.

Lemma 3.6. *Let $\mathbf{Y} \in K_{\mathbb{R}}^{k\times k}$ be such that $\mathbf{Y} = R \cdot \mathbf{I}_k + \mathbf{E}$ for some $R > 0$ and $\|e_{ij}\| \leq \varepsilon \cdot R$ for all $i, j \in [k]$. Assume that $\varepsilon \leq 1/(2k)$. Then $\mathbf{Y}$ is invertible and we have $\mathbf{Y}^{-1} = R^{-1} \cdot \mathbf{I}_k + \mathbf{E}'$, with $\|e'_{ij}\| \leq (k+1) \cdot \varepsilon \cdot R^{-1}$ for all $i, j \in [k]$. Further, it holds that $\det(\mathbf{Y}) \in [(1 + (k+1)(k+2)\varepsilon)^{-d/2}, (1+3\varepsilon)^{d/2}] \cdot R^{kd}$.*

4 From uSVP to NTRU

In this section, we prove the following result

Theorem 4.1. *Let K be a number field of degree d with $\zeta_K(2) = 2^{o(d)}$ and let $\gamma^+ > 0$ (recall that $\zeta_K(\cdot)$ denotes the Dedekind zeta function of K). There exists $q_0 = \mathrm{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+) \in \mathbb{R}_{\geq 0}$ and an algorithm* `uSVP-to-NTRU` *such that the following holds. For any $q \geq q_0$, $\gamma_{\mathrm{NTRU}} \geq \gamma'_{\mathrm{NTRU}} > 1$, $\gamma_{\mathrm{HSVP}} \geq \sqrt{d}\Delta_K^{1/(2d)}$, let*

$$\gamma_{\mathrm{uSVP}} = \gamma_{\mathrm{NTRU}} \cdot \sqrt{\gamma_{\mathrm{HSVP}}} \cdot 16\sqrt{2} \cdot d^{3/2} \cdot \delta_K$$
$$\gamma'_{\mathrm{uSVP}} = \frac{\gamma'_{\mathrm{NTRU}}}{\gamma_{\mathrm{HSVP}}^{3/2} \cdot 4 \cdot d^{9/2} \cdot \delta_K^2}.$$

For any distribution $\mathcal{D}_{\mathrm{uSVP}}$ over γ_{uSVP}-uSVP instances with gap $\leq \gamma^+$, let $\mathcal{D}_{\mathrm{NTRU}}$ be the distribution `uSVP-to-NTRU`$(\mathcal{D}_{\mathrm{uSVP}}, q, \gamma_{\mathrm{HSVP}})$. *We have four reductions*

- *from $(\mathcal{D}_{\mathrm{uSVP}}, \gamma_{\mathrm{uSVP}})$-uSVP$_{\mathrm{mod}}$ to $(\mathcal{D}_{\mathrm{NTRU}}, \gamma_{\mathrm{NTRU}}, q)$-NTRU$_{\mathrm{mod}}$;*
- *from γ_{uSVP}-wc-uSVP$_{\mathrm{mod}}$ restricted to modules with gap $\leq \gamma^+$ to $(\gamma_{\mathrm{NTRU}}, q)$-wc-NTRU$_{\mathrm{mod}}$;*
- *from $(\mathcal{D}_{\mathrm{uSVP}}, \gamma_{\mathrm{uSVP}}, \gamma'_{\mathrm{uSVP}})$-uSVP$_{\mathrm{vec}}$ to $(\mathcal{D}_{\mathrm{NTRU}}, \gamma_{\mathrm{NTRU}}, \gamma'_{\mathrm{NTRU}}, q)$-NTRU$_{\mathrm{vec}}$;*
- *from $(\gamma_{\mathrm{uSVP}}, \gamma'_{\mathrm{uSVP}})$-wc-uSVP$_{\mathrm{vec}}$ restricted to modules with gap $\leq \gamma^+$ to $(\gamma_{\mathrm{NTRU}}, \gamma'_{\mathrm{NTRU}}, q)$-wc-NTRU$_{\mathrm{vec}}$.*

Given access to an oracle solving γ_{HSVP}-id-HSVP, the four reductions run in time polynomial in their input size, in $\exp(\frac{d\log(d)}{\log(2q/q_0)})$ and in $\zeta_K(2)$.

The outline of the reduction is given in Figure 2. Note that the quantity $\zeta_K(2)$ may be exponential in d for some number fields (which may impact on the run-time of the reduction, or even on the applicability of the reduction since we require $\zeta_K(2) = 2^{o(d)}$). In the case of power-of-two cyclotomic fields, it was proven in [SS13, Lemma 4.2] that $\zeta_K(2) = O(1)$. The missing proofs of this section are available in Appendix D.

4.1 Pre-conditioning the uSVP Instance

In this section, we use algorithm `DualRound` to pre-process the input module and control its volume. In order to have the Hermite Normal Form of our integral module look like an NTRU instance, we slightly modify the geometry of our input module to make it have what we call the coprime property (see Definition 4.2). Hence, we describe an algorithm, called `PreCond` (see Algorithm D.1), which combines all this and transform any uSVP instance (with a lower bounded gap) into a new uSVP instance with roughly the same geometry and with all the properties we will require in Section 4.2.

Definition 4.2 (Coprime property). *We say that a rank-2 module $M \subseteq \mathcal{O}_K^2$ has the* coprime property *if it holds that*

$$\{x \in \mathcal{O}_K \mid \exists y \in \mathcal{O}_K,\ (x, y)^T \in M\} = \mathcal{O}_K.$$

In other words, the module M has the coprime property if the ideal spanned by the first coordinate of all the vectors of M is equal to $\mathcal{O}_K$.

We note that having the coprime property is not very constraining. In fact, any module can be applied a small distorsion in order to ensure the coprime property. This is formalized in Lemma 4.3 below.

Lemma 4.3. *Let $(\mathbf{B}, \mathbb{I})$ be a pseudo-basis of a rank-2 module $M \subset K^2$ with gap $\gamma(M) \geq 1$. There exists some $V_0 > 0$ with $V_0^{1/(2d)} = \mathrm{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma(M))$ and an algorithm `PreCond` such that the following holds. Let $\beta \in \{2, \cdots, 2d\}$ and $V > 0$ be such that $V^{1/(2d)} \geq (2d)^{2d/\beta} \cdot V_0^{1/(2d)}$. Then, on input $(\mathbf{B}, \mathbb{I})$, V and β, algorithm `PreCond` outputs a matrix $\mathbf{Y} \in \mathrm{GL}_2(K)$ such that*

- *if $(\mathbf{B}, \mathbb{I})$ is a γ_{uSVP}-uSVP instance, then $(\mathbf{YB}, \mathbb{I})$ is a γ'_{uSVP}-uSVP instance for $\gamma'_{\mathrm{uSVP}} = \gamma_{\mathrm{uSVP}}/(2\sqrt{2})$;*
- *the rank-2 module $M' := \mathbf{YB} \cdot \mathbb{I}$ is contained in $\mathcal{O}_K^2$;*
- *$\mathcal{N}(M') \in [1/2^d, 2^d] \cdot V$;*
- *M' has the coprime property;*
- *$\mathbf{Y} = R \cdot \mathbf{I}_2 + \mathbf{E}$ for some $R = V^{1/(2d)} \cdot \mathcal{N}(M)^{-1/(2d)} > 0$ and $\|e_{ij}\| \leq R/5$ for all $1 \leq i, j \leq 2$.*

Assume that $\zeta_K(2) \leq 2^{o(d)}$. Then Algorithm `PreCond` runs in expected time polynomial in its input bitsize, in 2^β and in $\zeta_K(2)$.

4.2 Transforming a uSVP Instance into an NTRU Instance

As the NTRU modules are free, the second step of our reduction finds a free module containing our uSVP instance and transforms it into an NTRU instance. For this purpose, we use the `BalanceIdeal` algorithm (cf Algorithm D.2) that takes as input any fractional ideal I and uses a γ_{HSVP}-id-HSVP oracle to output a balanced element x such that $\langle x \rangle$ contains I but is not much larger than it.

Lemma 4.4. *There exists an algorithm `BalanceIdeal` that takes as input a fractional ideal $I \subset K$ and a parameter $\gamma_{\mathrm{HSVP}} \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$, and outputs an element $x \in K$ such that $I \subseteq \langle x \rangle$ and $|\sigma_i(x)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$ for all $i \leq d$, where $\sigma = \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(I)^{-1/d}$.*

Moreover, given access to a γ_{HSVP}-id-HSVP oracle, it runs in polynomial time and makes one call to the γ_{HSVP}-id-HSVP oracle.

We can now describe our algorithm transforming a uSVP instance into an NTRU instance: Algorithm 4.1. The operations done by this algorithm are summarised in Figure 2 and proven in Lemma 4.6.

Algorithm 4.1 Algorithm `Conditioned-to-NTRU`

Input: A pseudo-basis $\mathbf{B}_1 \cdot \mathbb{I}$ of a rank-2 module in $\mathcal{O}_K^2$ and some parameters q and γ_{HSVP}

Output: A basis $\mathbf{B}_4$ of a free rank-2 module and some auxiliary information `aux`

Compute the HNF pseudo-basis $\mathbf{B}_2 \cdot \mathbb{J}$ of the rank-2 module spanned by $\mathbf{B}_1 \cdot \mathbb{I}$
Let $a = \mathbf{B}_2[1,0]$ # $\mathbf{B}_2 = \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix}$
Sample $b \leftarrow$ `BalanceIdeal`$(J_2, \gamma_{\text{HSVP}})$
Compute $h = \lfloor a \cdot q / b \rceil$
Return $\mathbf{B}_4 = \begin{pmatrix} 1 & 0 \\ h & q \end{pmatrix}$ and `aux` $= (a, b, J_1, J_2)$

Lemma 4.5. *Let $\gamma_{\text{HSVP}} \geq \sqrt{d}\Delta_K^{1/(2d)}$, $q \in \mathbb{Z}_{>0}$ and $(\mathbf{B}, \mathbb{I})$ be a pseudo-basis of a rank-2 module $M \subseteq \mathcal{O}_K^2$. Assume that we have access to a γ_{HSVP}-id-HSVP oracle. On input γ_{HSVP}, q and $(\mathbf{B}, \mathbb{I})$, algorithm* `Conditioned-to-NTRU` *runs in polynomial time in the bitsize of its input and makes one call to the γ_{HSVP}-id-HSVP oracle.*

Lemma 4.6. *Let $\gamma_{\text{HSVP}} \geq \sqrt{d} \cdot \Delta_K^{1/(2d)}$, $\gamma_{\text{NTRU}} > 1$ and $q \in \mathbb{Z}_{>0}$ be some parameters. Define*

$$V = \gamma_{\text{HSVP}}^d \cdot q^d \cdot d^d$$
$$\textit{and } \gamma_{\text{uSVP}} = \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 8 \cdot d^{3/2} \cdot \delta_K.$$

Let $(\mathbf{B}, \mathbb{I})$ be any γ_{uSVP}-uSVP instance in $\mathcal{O}_K^2$, with the coprime property and with norm in $[1/2^{2d} \cdot V, 2^{2d} \cdot V]$. Then on input $(\mathbf{B}, \mathbb{I}), \gamma_{\text{HSVP}}, q$, the algorithm `Conditioned-to-NTRU` *outputs $(\mathbf{B}_4, \textit{aux})$ such that $\mathbf{B}_4$ is a $(\gamma_{\text{NTRU}}, q)$-NTRU instance.*

The `aux` information output by algorithm `Conditioned-to-NTRU` will be used in Algorithms D.4 and D.3 to lift any short vector / dense submodule from the NTRU instance back to the uSVP instance. The proofs of Lemmas 4.5 and 4.6 are available in Appendices D.4 and D.5 respectively.

4.3 Lifting back Short Vectors and Dense Submodules

In this section, we prove that using the auxiliary information `aux` produced by Algorithm `Conditioned-to-NTRU`, one can lift a short vector or a densest submodule from the output NTRU instance back to the input uSVP instance. The proofs of Lemmas 4.7 and 4.8 are available in Appendices D.6 and D.7 respectively.

Lemma 4.7. *There exists an algorithm* `LiftMod` *such that the following holds. Let q, γ_{HSVP} and $(\mathbf{B}, \mathbb{I})$ be as in Lemma 4.6. Let M_1 denote the rank-2 module generated by $(\mathbf{B}, \mathbb{I})$, $[\mathbf{C}, (a, b, J_1, J_2)] \leftarrow$* `Conditioned-to-NTRU`$((\mathbf{B}, \mathbb{I}), q, \gamma_{\text{HSVP}})$ *and let M_4 denote the rank-2 free module generated by $\mathbf{C}$.*

Module	Pseudo-basis	Short vector
M_1	$\left[\begin{matrix} I_1 & I_2 \\ \begin{pmatrix} b_{11} & b_{12} \\ b_{21} & b_{22} \end{pmatrix} \end{matrix}\right]$	$\mathbf{s}_1 = \begin{pmatrix} u \\ v \end{pmatrix}$
	$\downarrow$ Step 1 (HNF)	
$M_2 = M_1$	$\left[\begin{matrix} J_1 & J_2 \\ \begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix} \end{matrix}\right]$	$\mathbf{s}_2 = \mathbf{s}_1$
	$\downarrow$ Step 2 (Principalization)	
$M_3 \supseteq M_2$	$\left[\begin{matrix} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ a & b \end{pmatrix} \end{matrix}\right]$	$\mathbf{s}_3 = \mathbf{s}_2$
	$\downarrow$ Step 3 (distortion + rounding)	
M_4	$\left[\begin{matrix} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ \lfloor a \cdot q/b \rceil & q \end{pmatrix} \end{matrix}\right]$	$\mathbf{s}_4 = \begin{pmatrix} u \\ v \cdot q/b - u \cdot \{a \cdot q/b\} \end{pmatrix}$

Fig. 2. Outline of algorithm `Conditioned-to-NTRU`.

Let $(\mathbf{v}, J)$ be a pseudo-basis of the densest rank-1 submodule of M_4. Then, on input $a, b, (\mathbf{C}, \mathcal{O}_K^2)$ and $(\mathbf{v}, J)$, algorithm `LiftMod` *outputs $\mathbf{w} \in K$ such that $\mathrm{span}_K(\mathbf{w}) \cap M_1$ is the densest rank-1 submodule of M_1.*

Moreover, algorithm `LiftMod` *runs in polynomial time.*

Lemma 4.8. *There exists an algorithm* `LiftVec` *such that the following holds. Let $q, \gamma_{\mathrm{HSVP}}$ and $(\mathbf{B}, \mathbb{I})$ be as in Lemma 4.6. Let M_1 denote the rank-2 module generated by $(\mathbf{B}, \mathbb{I})$, $[\mathbf{C}, \boldsymbol{aux}] \leftarrow$* `Conditioned-to-NTRU`*$((\mathbf{B}, \mathbb{I}), q, \gamma_{\mathrm{HSVP}})$ and let M_4 denote the rank-2 free module generated by $\mathbf{C}$.*

Let $\mathbf{s} \in M_4$. Then, on input $\boldsymbol{aux}, \gamma_{\mathrm{HSVP}}, (\mathbf{C}, \mathcal{O}_K^2)$ and $\mathbf{s}$, algorithm `LiftVec` *outputs a vector $\mathbf{t} \in M$ such that $\|\mathbf{t}\| \leq \|\mathbf{s}\| \cdot 68 \cdot \gamma_{\mathrm{HSVP}}^2 \cdot d^4 \cdot \delta_K^2$.*

If given access to a γ_{HSVP}-id-HSVP oracle, algorithm `LiftVec` *runs in polynomial time and makes 1 call to the oracle.*

Combining all the results of this section, one can prove Theorem 4.1.

5 Randomization of Rank-2 Modules with Gaps

A rank-2 module with a gap can, by Lemma 2.9 and the fact that densest submodules are primitive, be written as $M = \mathbf{u} \cdot J_1 + \mathbf{v} \cdot J_2$ where $\mathbf{u} \cdot J_1$ is a densest rank-1 submodule of M. Informally, the goal of this section is to randomize $\mathbf{u}, \mathbf{v}, J_1, J_2$ without changing the gap too much. The missing proofs of this section are available in Appendix E.

We first describe the average-case distribution we are considering. Note that the gap parameter γ' is itself a random variable.

Definition 5.1. *Let $\gamma > 0$ and $B > 2$. We define the distribution $D^{\text{module}}_{B,\gamma}$ over rank-2 and norm-1 modules by:*

$$D^{\text{module}}_{B,\gamma} = \texttt{QRSF-2-Mod}(\mathbf{Q}, \gamma', I_1, I_2, r),$$

where

- *the matrix $\mathbf{Q}$ is uniform in $\mathcal{O}_2(K_{\mathbb{R}})$;*
- *the gap parameter γ' is set as $\gamma' = \gamma \cdot \mathcal{N}(c/a)^{1/(2d)}/B^{1/(2d)}$ with $(a, c) \in K^2_{\mathbb{R}}$ distributed as $\chi_{K_{\mathbb{R}}} \times \mathcal{D}(0, 1)$ conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$;*
- *the ideals I_1, I_2 are uniform in $\mathcal{I}_1$ (the set of norm-1 ideals);*
- *the element r is uniform in $K_{\mathbb{R}} \bmod \gamma'^{-2} \cdot I_1 I_2^{-1}$.*

We now state the main theorem of this section, which can be viewed as a worst-case to average-case reduction for rank-2 modules with a gap.

Theorem 5.2 (ERH). *For all $B \geq (d^d \Delta_k)^{\Omega(1)}$ and $\gamma \geq B^{1/(2d)}$ there exists a procedure $\texttt{Randomize}_B$ that runs in time polynomial in $\log B$ and the bitsize of its input, and such that on input a pseudo-basis $(\mathbf{B}, \mathbb{I})$ of a rank-2 and norm-1 module M of gap γ outputs a pair $((\mathbf{B}', \mathbb{I}'), \boldsymbol{aux})$ such that*

- *the pseudo-basis $(\mathbf{B}', \mathbb{I}')$ spans a rank-2 and norm-1 module M';*
- *any event that holds for $D^{\text{module}}_{B,\gamma}$ with probability $\varepsilon \geq 2^{-o(d)}$ also holds for M' with probability $\Omega(\varepsilon^4)$ over the internal randomness of $\texttt{Randomize}_B$.*

Further, there exists a deterministic algorithm $\texttt{Recover}$ that runs in time polynomial in the bitsize of its input such that for M' as above, if U' is a densest rank-1 submodule of M', then $\texttt{Recover}(U', \boldsymbol{aux})$ returns the densest rank-1 submodule of M, with probability $1-2^{-\Omega(d)}$ over the randomness of $\texttt{Randomize}_B$.

We note that the theorem does not state that the output distribution of $\texttt{Randomize}_B$ is $D^{\text{module}}_{B,\gamma}$, but only that they are close in the sense that any event that holds with sufficient probability for $D^{\text{module}}_{B,\gamma}$ also holds for the output distribution of $\texttt{Randomize}_B$ with a polynomially related probability.

$\texttt{Randomize}_B$ is described in Algorithm 5.6. It consists of two main steps: a coefficient randomization (described in Section 5.1), whose purpose is to randomize the ideal coefficients; and a geometric randomization (described in Section 5.2), whose purpose is to randomize the pseudo-basis matrix. Section 5.3 compares the distribution that would ideally be returned by the composition of the coefficient and geometric randomizations, with the distribution of the pseudo-basis in Definition 5.1. Finally, we complete the proof of Theorem 5.2 in Section 5.4.

5.1 Coefficient Randomization

In the coefficient randomization step, our aim is to randomize the ideal coefficients of a good pseudo-basis (i.e., whose first pair corresponds to the densest rank-1 submodule), given an arbitrary pseudo-basis of a rank-2 module. One may multiply the whole pseudo-basis by a random ideal, but this only randomizes the pair of ideal coefficients. More precisely, this leaves the ratio of the ideal coefficients unchanged. To decouple the ideal coefficients, we use module sparsification, as described in Section 3. This first step towards coefficient randomization is formally described in Algorithm 5.1. Steps 1 and 3 are respectively performed using Lemmas 2.3 and 3.4.

Algorithm 5.1 Partial Coefficient Randomization: $\texttt{Partial-CR}_B$

Input: A pseudo-basis of a rank-2 module M.

1: Sample $\mathfrak{p}$ uniformly among prime ideals of norms $\leq B$;
2: Sample $\overline{\mathbf{b}^\vee}$ uniformly in $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$;
3: Return a pseudo-basis of the sparsification of M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$ along with $\mathfrak{p}$.

Theorem 5.3 (ERH). *Let $B \geq (\log \Delta_K)^{\Omega(1)}$. The runtime of $\texttt{Partial-CR}_B$ is polynomial in $\log B$ and the bitsize of its input. Let $(\mathbf{B}, \mathbb{I})$ be a pseudo-basis of a rank-2 module M, and let $(J_1, \mathbf{u}), (J_2, \mathbf{v})$ be an arbitrary pseudo-basis of M. Let M' be the rank-2 module spanned by the pseudo-basis output by $\texttt{Partial-CR}_B$ when given $(\mathbf{B}, \mathbb{I})$ as input, let $\overline{\mathbf{b}^\vee}$ be the element of $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$ sampled in $\texttt{Partial-CR}_B$ and let $\mathbf{b}^\vee$ be a lift of $\overline{\mathbf{b}^\vee}$ in $M^\vee$.*
Then, with probability $1 - (1/B)^{\Omega(1)}$, we have $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \notin \mathfrak{p}J_1^{-1}$. In that case, there exists $x \in J_1 J_2^{-1}$ such that

$$M' = \mathbf{u} \cdot \mathfrak{p}J_1 + (\mathbf{v} + x\mathbf{u}) \cdot J_2.$$

Assume further that $\gamma(M) \geq B^{1/(2d)}$ and that $\mathbf{u} \cdot J_1$ is the densest rank-1 submodule of M. Then, still when $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \notin \mathfrak{p}J_1^{-1}$, we have that $\gamma(M') = \gamma(M)/\mathcal{N}(\mathfrak{p})^{1/(2d)} > 1$ and $\mathbf{u} \cdot \mathfrak{p}J_1$ is the densest rank-1 submodule of M'.

The result follows from Lemmas 5.4 and 5.5, whose proofs are postponed to Appendix E.

Lemma 5.4. *Borrowing the notations of Theorem 5.3, we have*

$$\mathbf{u} \cdot \mathfrak{p}J_1 \subset M' \quad \text{and} \quad \mathbf{u} \cdot J_1 \not\subset M',$$

with probability $1 - (1/B)^{\Omega(1)}$ over the choices of $\mathfrak{p}$ and $\overline{\mathbf{b}^\vee}$.

Lemma 5.5. *Borrowing the notations of Theorem 5.3 and assuming that $\mathbf{u} \cdot J_1 \not\subset M'$, there exists $x \in J_1 J_2^{-1}$ such that $(\mathbf{v} + x\mathbf{u}) \cdot J_2 \subset M'$.*

We now describe the coefficient randomization. Ideally, we would have access to a pseudo-basis $((J_1, \mathbf{u}), (J_2, \mathbf{v}))$ of the module M under scope, for which the densest rank-1 submodule is $\mathbf{u} \cdot J_1$. We would multiply J_1 by a random ideal and J_2 by another random ideal. Unfortunately, given only access to an arbitrary pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of M, this seems difficult to achieve obliviously. Instead, we use algorithm `Ideal-Sample` (Algorithm 2.1) to obtain a uniform norm-1 ideal I, and multiply M by it. This will obliviously multiply J_1 and J_2 by I. As this distribution is invariant by ideal multiplication, the ideal $J_2 I / \mathcal{N}(J_2)^{1/d}$ will be uniform among norm-1 ideals. It remains to obliviously randomize the first ideal independently of the second one. For this purpose, we use `Partial-CR` (Algorithm 5.1), which has the effect of obliviously multiplying the first ideal with a random prime ideal $\mathfrak{p}$ while leaving the second one unchanged (with overwhelming probability). Note that multiplying by a random prime ideal is the main component of the ideal randomization algorithm `Ideal-Sample`. In a sense, this "almost" randomizes J_1.

Algorithm 5.2 describes the process on the input basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$. The corresponding randomization performed on the hidden pseudo-basis $((J_1, \mathbf{u}), (J_2, \mathbf{v}))$ is described in Algorithm 5.3. Note that there is no need for Algorithm 5.3 to be efficient as its sole purpose is to describe the behavior of Algorithm 5.2 on the hidden pseudo-basis.

In Theorem 5.6, we show that the resulting distributions on the output modules are statistically close, and describe the evolution of the densest rank-1 submodule.

Algorithm 5.2 Real Coefficient Randomization: $\texttt{Real-CR}_{B,B'}$

Input: A pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of a module $M \subset K_{\mathbb{R}}^2$.
Let $((I'_1, \mathbf{b}'_1), (I'_2, \mathbf{b}'_2)), \mathfrak{p}$ be the output of $\texttt{Partial-CR}_B$ on input $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$;
Sample $\mathfrak{q}$ using $\texttt{Ideal-Sample}_{B'}$;
Let $\mathbf{b}''_i = \mathbf{b}'_i / \mathcal{N}(\mathfrak{p})^{1/(2d)}$ for $i \in [2]$;
Return $((\mathfrak{q} I'_1, \mathbf{b}''_1), (\mathfrak{q} I'_2, \mathbf{b}''_2)), \mathfrak{p}, \mathfrak{q}$.

Algorithm 5.3 Ideal Coefficient Randomization: $\texttt{Ideal-CR}_B$

Input: $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}}), \gamma > 1, J_1, J_2$ ideals of norm 1, $r \in K_{\mathbb{R}}$;
Let $M = \texttt{QRSF-2-Mod}(\mathbf{Q}, \gamma, J_1, J_2, r)$;
Let $\mathbf{u} = 1/\gamma \cdot \mathbf{Q} \cdot (1, 0)^T$ and $\mathbf{v} = \gamma \cdot \mathbf{Q} \cdot (r, 1)^T$;
Sample $\mathfrak{p}$ uniformly among prime ideals of norms $\leq B$;
Sample $\mathbf{b}^\vee$ in $M^\vee$, uniform in $M^\vee / \mathfrak{p} M^\vee$ conditioned on $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_{\mathbb{R}}} \notin \mathfrak{p} J_1^{-1}$;
Find $x \in J_1 J_2{}^{-1}$ such that $\langle \mathbf{b}^\vee, \mathbf{v} + x \cdot \mathbf{u} \rangle_{K_{\mathbb{R}}} \in \mathfrak{p} J_2^{-1}$;
Sample J uniformly among norm-1 ideals;
Return $(\mathbf{Q}, \gamma / \mathcal{N}(\mathfrak{p})^{1/(2d)}, J_1 J_2{}^{-1} J \mathfrak{p} / \mathcal{N}^{1/d}(\mathfrak{p}), J, r + x)$.

Theorem 5.6 (ERH). *Assume that $B' \geq (d^d \Delta_K)^{\Omega(1)}$ and $B \geq (\log \Delta_K)^{\Omega(1)}$. The runtime of $\texttt{Real-CR}_{B,B'}$ is polynomial in $\log(BB')$ and the bitsize of its input.*

Let $M = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right) \subset K_{\mathbb{R}}^2$ a module with norm 1, in QR-standard form. Then the distribution of the module output by $\texttt{Real-CR}_{B,B'}$ on *input an arbitrary pseudo-basis of M is within statistical distance $(1/B)^{\Omega(1)}+2^{-d}$ of* $\texttt{QRSF-2-Mod}(\texttt{Ideal-CR}_B(\mathbf{Q},\gamma,J_1,J_2,r))$.

Assume further that $\gamma \geq B^{1/(2d)}$ and let U denote the densest rank-1 submodule of M. Let $(M',\mathfrak{p},\mathfrak{q})$ be the output of $\texttt{Real-CR}_{B,B'}$ *on input M. Then, with probability $1-(1/B)^{\Omega(1)}$, we have that $\gamma(M')=\gamma(M)/\mathcal{N}(\mathfrak{p})^{1/(2d)}>1$ and the densest rank-1 submodule of M' is*

$$\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}\cdot U\cdot \mathfrak{q}\frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}.$$

5.2 Geometric Randomization

In the geometric module randomization, we will use a distribution D_{distort} over $K_{\mathbb{R}}^{2\times 2}$ whose purpose is to distort the geometric relationship between the densest rank-1 submodule and the complementing rank-1 submodule of the rank-2 module under scope. We define D_{distort} as $\mathcal{D}_{K_{\mathbb{R}}}(0,1)^{2\times 2}$ conditioned on the event that $|\det(\sigma_i(\mathbf{D}))|>1/d$ holds for all $i\in[d]$.

The following lemmas describe useful properties of the distribution D_{distort}.

Lemma 5.7. *The following properties hold.*

- *The distribution D_{distort} can be sampled from in time polynomial in d.*
- *The distribution D_{distort} is invariant by multiplication on the left and the right by matrices in $\mathcal{O}_2(K_{\mathbb{R}})$.*

Lemma 5.8. *Let D be the distribution over $K_{\mathbb{R}}^{2\times 2}$ of*

$$\mathbf{Q}\cdot\begin{pmatrix} a & b\\ 0 & c\end{pmatrix}$$

where $\mathbf{Q}\leftarrow\mathcal{U}(\mathcal{O}_2(K_{\mathbb{R}}))$, $a\leftarrow\chi_{K_{\mathbb{R}}}$ and $b,c\leftarrow\mathcal{D}_{K_{\mathbb{R}}}(0,1)$, conditioned on the event that for all $i\in[d]$ we have $|\sigma_i(a\cdot c)|\geq 1/d$. Then $D=D_{\text{distort}}$.

Let $((J_1,\mathbf{u}),(J_2,\mathbf{v}))$ be a pseudo-basis of a rank-2 module M. Assume that $\mathbf{u}\cdot J_1$ is the densest rank-1 submodule, but that we have access to this pseudo-basis only indirectly, via an arbitrary pseudo-basis of M. Write

$$(\mathbf{u}|\mathbf{v})=\mathbf{Q}\cdot\begin{pmatrix} 1 & r\\ 0 & 1\end{pmatrix},$$

for some $r\in K_{\mathbb{R}}$. The purpose of the geometric randomization is to map r to some r' that is uniform modulo $J_1J_2^{-1}$, while at the same time not distorting the module M too much, so that the randomized M still has a gap and its rank-1 densest submodule is related to $\mathbf{u}\cdot J_1$. For this purpose, we multiply M on the left by a matrix sampled from D_{distort}. For the analysis, it is convenient to take it Gaussian, and to avoid a potentially large distortion, we avoid matrix samples with small determinant. This corresponds to algorithm `Real-GR` (Algorithm 5.4). The effect on the hidden pseudo-basis $((J_1,\mathbf{u}),(J_2,\mathbf{v}))$ is described in algorithm `Ideal-GR` (Algorithm 5.5). In Theorem 5.9, we show that the resulting module distributions are identical, and describe the evolution of the densest rank-1 sublattice.

Algorithm 5.4 Real Geometric Randomization: `Real-GR`

Input: A pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of a norm-1 module $M \subset K_{\mathbb{R}}^2$.

Sample $\mathbf{D} \leftarrow D_{\text{distort}}$ (using Lemma 5.7);
$(\mathbf{b}'_1 | \mathbf{b}'_2) \leftarrow \det(\mathbf{D})^{-1/(2d)} \cdot \mathbf{D} \cdot (\mathbf{b}_1 | \mathbf{b}_2)$;
Return $((I_1, \mathbf{b}'_1), (I_2, \mathbf{b}'_2)), \mathbf{D}$.

Algorithm 5.5 Ideal Geometric Randomization: `Ideal-GR`

Input: $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}}), \gamma > 1, J_1, J_2$ ideals of norm 1, $r \in K_{\mathbb{R}}$;

Sample $a \leftarrow \chi_{K_{\mathbb{R}}}$ and $c \leftarrow \mathcal{D}(0,1)$ conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$;
Sample $b \leftarrow \mathcal{D}(0,1)$;
Sample $\mathbf{Q}' \leftarrow \mathcal{U}(\mathcal{O}_2(K_{\mathbb{R}}))$;
Set $J'_1 = a/\mathcal{N}^{1/d}(a) \cdot J_1$ and $J'_2 = c/\mathcal{N}^{1/d}(c) \cdot J_2$;
Set $\gamma' = \gamma \cdot \mathcal{N}(c/a)^{1/(2d)}$;
Set $r' = (b + ar)/c$;
Return $(\mathbf{Q}', \gamma', J'_1, J'_2, r')$.

Theorem 5.9. *Algorithm* `Real-GR` *runs in polynomial time. Let* $M = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right) \subset K_{\mathbb{R}}^2$ *a module with norm 1, in QR-standard-form. Let M' be the module spanned by the output of* `Real-GR` *on input an arbitrary pseudo-basis of M. Then the distribution of M' is identical to the distribution* `QRSF-2-Mod(Ideal-GR(`$\mathbf{Q}, \gamma, J_1, J_2, r$`))`.

Further, if $\gamma > d$ and U is the densest rank-1 submodule of M, then, with probability $1 - 2^{-\Omega(d)}$, we have $\gamma(M') > 1$ and the densest rank-1 submodule of M' is $\det(\mathbf{D})^{-1/(2d)} \cdot \mathbf{D} \cdot U$, where $\mathbf{D}$ is the Gaussian matrix sampled during the execution of `Real-GR`.

5.3 On the `Ideal-GR` ∘ `Ideal-CR` Distribution

We define a few probability distributions over the inputs of `QRSF-2-Mod`, which we will use to show that the operations performed on the available arbitrary pseudo-basis randomize the rank-2 module, so that the input module is "forgotten" in the output module distribution while at the same time controlling the evolution of the densest rank-1 submodule.

Definition 5.10. *Let $B \geq 2$ and $\gamma > 0$. We consider the following random variables, which are assumed independent (unless stated otherwise).*

- $\mathbf{Q}$ *uniform in $\mathcal{O}_2(K_{\mathbb{R}})$;*
- *$b \in K_{\mathbb{R}}$ distributed as $\mathcal{D}_{K_{\mathbb{R}}}(0,1)$;*
- *$(a, c) \in K_{\mathbb{R}}^2$ distributed as $\chi_{K_{\mathbb{R}}} \times \mathcal{D}_{K_{\mathbb{R}}}(0,1)$ conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a \cdot c)| \geq 1/d$; we define $\gamma' = \gamma \cdot \mathcal{N}(c/a)^{1/(2d)}/B^{1/(2d)}$;*
- *$\mathfrak{p}$ uniform among prime ideals of norms $\leq B$;*
- *I_1, I_2, J uniform in $\mathcal{I}_1$ (the set of norm-1 ideals);*
- *$\zeta \in E$ sampled from the centered normal law of standard deviation $d^{-3/2}$, conditioned on $\|\zeta\| \leq 1/d$;*

- *u uniform in $\{x \in K_{\mathbb{R}}, \forall i \in [d] : |\sigma_i(x)| = 1\}$;*
- *r' uniform in $K_{\mathbb{R}} \bmod \gamma'^{-2} \cdot I_1 I_2^{-1}$.*

Let $J_1, J_2 \in \mathcal{I}_1$ and $r \in K_{\mathbb{R}}$ arbitrary. Let x be as in Step 5 of `Ideal-CR`$_B$*, when given $(\mathbf{Q}, \gamma, J_1, J_2, r)$ as input and with the variable $\mathfrak{p}$ of* `Ideal-CR`$_B$ *being the random variable above. In order to simplify the notations, we define the random variable:*

$$I(J_1, J_2) = \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{au}{c\exp(\zeta)} \cdot J_1 {J_2}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})} \ \in \mathcal{I}_1.$$

Let $r''(J_1, J_2)$ be uniformly distributed in $K_{\mathbb{R}} \bmod \gamma'^{-2} \cdot I(J_1, J_2) \cdot J^{-1}$.

We define the following distributions of the form $(\tilde{\mathbf{Q}}, \tilde{\gamma}, \tilde{I}_1, \tilde{I}_2, \tilde{r})$, where the random variables $\tilde{r}$ is defined modulo $\tilde{\gamma}^{-2} \cdot \tilde{I}_1 \cdot \tilde{I}_2^{-1}$:

$$D_{B,\gamma}^{\mathrm{rand}} : \left(\mathbf{Q},\ \gamma\frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \frac{a}{\mathcal{N}^{\frac{1}{d}}(a)} J_1 {J_2}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ \frac{c}{\mathcal{N}^{\frac{1}{d}}(c)} \cdot J,\ \frac{b + a(r+x)}{c}\right),$$

$$D_{B,\gamma}^{(1)} : \left(\mathbf{Q},\ \gamma\frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{au}{c} \cdot J_1 {J_2}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ J,\ u\frac{b + a(r+x)}{c}\right),$$

$$D_{B,\gamma}^{(2)} : \left(\mathbf{Q},\ \gamma \cdot \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ I(J_1, J_2),\ J,\ u\frac{b + a(r+x)}{c\exp(\zeta)}\right),$$

$$D_{B,\gamma}^{(3)} : \left(\mathbf{Q},\ \gamma',\ I(J_1, J_2),\ J,\ \frac{B^{\frac{1}{d}}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})} \cdot u\frac{b + a(r+x)}{c\exp(\zeta)}\right),$$

$$D_{B,\gamma}^{(4)} : (\mathbf{Q},\ \gamma',\ I(J_1, J_2),\ J,\ r''(J_1, J_2)),$$

$$D_{B,\gamma}^{\mathrm{target}} : (\mathbf{Q},\ \gamma',\ I_1,\ I_2,\ r').$$

Note that $D_{B,\gamma}^{\mathrm{rand}}$ is the distribution obtained by composing `Ideal-CR`$_B$ (Algorithm 5.3) and `Ideal-GR` (Algorithm 5.5), on an input of the form $(\mathbf{Q}_0, \gamma, J_1, J_2, r)$ with (γ, J_1, J_2, r) as above and $\mathbf{Q}_0 \in \mathcal{O}_2(K_{\mathbb{R}})$ arbitrary. These algorithms significantly randomize the QR-standard form, but it still depends on (J_1, J_2, r). On the other hand, the distribution $D_{B,\gamma}^{\mathrm{target}}$ is independent of (J_1, J_2, r). Our goal is to show that these two distributions are similar, in the sense that any event that holds with some probability $\varepsilon \geq 2^{-o(d)}$ for one holds with probability $\varepsilon^{O(1)}$ for the other one.

For this purpose, we consider the intermediate (hybrid) distributions of Definition 5.10. To help the reader, we use two colours in the definition of the successive distributions. The entries of the tuples that are in red are those that change compared to the previous distribution. The variables with blue background are those that depend on (J_1, J_2, r). The relations between the distributions of Definition 5.10 are pictorially summarized in Figure 3. The lemmas formally stating these relations and their proofs are provided in Appendix E. Some of the relations require $B \geq (d^d \Delta_K)^{\Omega(1)}$ or $\gamma \geq d^{1/4} \cdot \Delta_K^{1/(2d)}$.

$$D^{\mathrm{rand}}_{B,\gamma} = D^{(1)}_{B,\gamma} \xrightarrow{\mathrm{RD}_2=O(1)} D^{(2)}_{B,\gamma} \xrightarrow{\mathrm{RD}_2=O(1)} D^{(3)}_{B,\gamma} \xleftrightarrow{\mathrm{SD}=2^{-\Omega(d)}} D^{(4)}_{B,\gamma} \xleftrightarrow{\mathrm{SD}=2^{-\Omega(d)}} D^{\mathrm{target}}_{B,\gamma}$$

Fig. 3. The relations between the distributions of Definition 5.10, proved in Lemmas E.1, E.2, E.3, E.4 and E.6. Here $D \xrightarrow{\mathrm{RD}_2=O(1)} D'$ means $\mathrm{RD}(D' \parallel D) = O(1)$ and $D \xrightarrow{\mathrm{SD}=2^{-\Omega(d)}} D'$ means $\mathrm{SD}(D, D') = 2^{-\Omega(d)}$.

5.4 Full Module Randomization

The full randomization algorithm $\texttt{Randomize}_B$ (Algorithm 5.6) is the composition of algorithms `Real-CR` and `Real-GR`.

Algorithm 5.6 (Real) Full Randomization: $\texttt{Randomize}_B$

Input: A pseudo-basis $(\mathbf{B}, \mathbb{I})$ of a norm-1 module $M \subset K^2_{\mathbb{R}}$.
1: Apply $\texttt{Real-CR}_{B,(d^d \Delta_K)^{\Omega(1)}}$ to $(\mathbf{B}, \mathbb{I})$ and let $((\mathbf{B}^\circ, \mathbb{I}^\circ), \mathfrak{p}, \mathfrak{q})$ be the output;
2: Apply `Real-GR` to $(\mathbf{B}^\circ, \mathbb{I}^\circ)$ and let $((\mathbf{B}', \mathbb{I}'), \mathbf{D})$ be the output;
3: Return $((\mathbf{B}', \mathbb{I}'), \texttt{aux})$ with $\texttt{aux} = (\mathfrak{p}, \mathfrak{q}, \mathbf{D})$.

Let $((\mathbf{B}', \mathbb{I}'), \texttt{aux})$ be an output of $\texttt{Randomize}_B$, and U' be a rank-1 submodule of the module spanned by $(\mathbf{B}', \mathbb{I}')$. We define:

$$\texttt{Recover}(U', \texttt{aux} = (\mathfrak{p}, \mathfrak{q}, D)) \;=\; (\mathcal{N}(\mathfrak{p}) \cdot \det(\mathbf{D}))^{\frac{1}{2d}} \cdot \mathbf{D}^{-1} \cdot U' \cdot \mathfrak{q}^{-1} \mathfrak{p}^{-1}.$$

With these choices of algorithms $\texttt{Randomize}_B$ and `Recover`, we can finally prove Theorem 5.2. For this purpose, we show that the module distribution that is output from the randomization algorithm (on an arbitrary input) and the distribution $D^{\mathrm{module}}_{B,\gamma}$ from Definition 5.1 are close in the mixed "SD plus RD" sense of Figure 3. The full proof is available in Appendix E.1.

6 Random Self-Reducibility of Module uSVP

The main result of this section is the following worst-case to average-case reduction for $\mathrm{uSVP}_{\mathrm{mod}}$.

Theorem 6.1 (ERH). *There exist $\gamma_0 = (d\Delta_K^{1/d})^{O(1)}$ and a family of distributions $(D^{\mathrm{uSVP}}_\gamma)_{\gamma \ge \gamma_0}$ such that the following properties hold for any $\gamma \ge \gamma_0$:*

- *if $\gamma \le (2^d \Delta_K^{1/d})^{O(1)}$, then D^{uSVP}_γ can be sampled from in time polynomial in $\log \Delta_K$;*
- *with probability $1 - 2^{-\Omega(d)}$, a sample from D^{uSVP}_γ is a pseudo-basis of a rank-2 module $M \subseteq \mathcal{O}_K^2$ with gap $\gamma(M) \ge \gamma \cdot \sqrt{d}{\Delta_K}^{1/(2d)}$; in particular, these are γ-uSVP instances;*
- *there exists a Karp reduction from γ'-wc-$\mathrm{uSVP}_{\mathrm{mod}}$ to $(D^{\mathrm{uSVP}}_\gamma, \gamma)$-$\mathrm{uSVP}_{\mathrm{mod}}$, with $\gamma' = \gamma \cdot (d \cdot \Delta_K^{1/d})^{O(1)}$; the reduction runs in time polynomial in $\log \Delta_K$ and the input bitsize.*

Note that the restriction on γ for the first condition is very mild, as in this parameter range, uSVP_{mod} can be solved in polynomial time using the LLL algorithm [LLL82]. We now proceed in two steps. We first define and study the distribution D^{uSVP}, and then prove Theorem 6.1.

6.1 A Distribution over uSVP Instances

Let $\gamma > 1$. The distribution D_γ^{uSVP} is defined as follows:

- sample a module from $D_{B,\gamma'}^{\text{module}}$ along with a pseudo-basis $(\mathbf{B}, \mathbb{I})$, with $B = (d^d \Delta_K)^{O(1)}$ and $\gamma' = 2\gamma \cdot \sqrt{d}{\Delta_K}^{1/(2d)} \cdot \sqrt{d} B^{1/d}$ (see Definition 5.1) and using `Ideal-Sample` to sample from $\mathcal{I}_1$;
- call $\texttt{DualRound}_{\varsigma,\beta,\varepsilon}(\mathbf{B}, \mathbb{I})$ with $\varsigma = (2^d \Delta_K^{1/d})^{O(1)}$, $\beta = 2$ and $\varepsilon = 1/(2d)^{3/2}$, and let $\mathbf{Y}$ denote the output;
- return $\texttt{HNF}(\mathbf{Y} \cdot \mathbf{B}, \mathbb{I})$.

The first two statements of Theorem 6.1 are implied by the following lemmas, whose proofs can be found in Appendix F.

Lemma 6.2. *A sample M from $D_{B,\gamma'}^{\text{module}}$ has gap $\gamma(M) \geq \gamma'/(\sqrt{d}B^{1/d})$, with probability $1 - 2^{-\Omega(d)}$.*

Using the latter result and Lemma 2.11, we obtain that the assumptions of Lemma 3.5 are satisfied. This implies that the above sampling algorithm runs in time polynomial in $\log \Delta_K$. By Lemmas 3.5 and 3.6, the output is a pseudo-basis of a rank-2 module in $\mathcal{O}_K^2$.

Lemma 6.3. *Let $\gamma > 2$. Let $(\mathbf{B}, \mathbb{I})$ be a pseudo-basis of a rank-2 module M with gap γ. Let $\mathbf{Y}$ denote the output of $\texttt{DualRound}_{\varsigma,\beta,\varepsilon}(\mathbf{B}, \mathbb{I})$ with $\varsigma = \gamma \cdot (2d)^{2d+3}$, $\beta = 2$ and $\varepsilon = 1/(2d)^{3/2}$. Then the module spanned by $(\mathbf{Y} \cdot \mathbf{B}, \mathbb{I})$ has gap $\geq \gamma/2$.*

The definition of D_γ^{uSVP} and Lemmas 6.2 and 6.3 implies that the modules whose pseudo-basis are sampled from D_γ^{uSVP} have gap $\geq \gamma \cdot \sqrt{d}{\Delta_K}^{1/(2d)}$, and hence are γ-uSVP instances with overwhelming probability.

6.2 Reducing Worst-Case Instances to D^{uSVP} Instances

We first introduce intermediate problems, that will allow us to split the reduction into several steps.

Definition 6.4. *Let $\gamma > 1$. A γ-$\text{uSVP}^{\mathcal{N}}$ instance consists in a pseudo-basis $(\mathbf{B}, \mathbb{I})$ of a rank-2 module $M \subset K^2$ such that $\gamma(M) \geq \gamma$.*

Let $\mathcal{D}$ a distribution over γ-$\text{uSVP}^{\mathcal{N}}$ instances. The (D, γ)-$\text{uSVP}_{\text{mod}}^{\mathcal{N}}$ problem asks, given as input a sample $(\mathbf{B}, \mathbb{I})$ from $\mathcal{D}$, to recover a densest rank-1 submodule of the module spanned by $(\mathbf{B}, \mathbb{I})$.

The worst-case variant γ-wc-$\text{uSVP}_{\text{mod}}^{\mathcal{N}}$ asks to solve this problem for any γ-$\text{uSVP}^{\mathcal{N}}$ instance.

The $\gamma^{\approx}$-wc-$\text{uSVP}_{\text{mod}}^{\mathcal{N}}$ variant is the restriction of γ-wc-$\text{uSVP}_{\text{mod}}^{\mathcal{N}}$ to the γ-$\text{uSVP}^{\mathcal{N}}$ instances whose spanned modules M satisfy $\gamma(M) \in [\gamma, \gamma \cdot (1 + 1/d)]$.

Note that worst-case wc-uSVP$_{\mathrm{mod}}$ reduces to wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ as the existence of a short non-zero vector implies the one of a dense rank-1 module. Similarly, uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ reduces to uSVP$_{\mathrm{mod}}$ with a loss of a $(\sqrt{d}\Delta_K^{1/d})$ factor in the parameters, thanks to Minkoswki's theorem. To prove the third statement of Theorem 6.1, it hence suffices to reduce wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ to uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ for distribution D_γ^{uSVP}. The result follows from Lemmas 6.5 and 6.7.

The first lemma states that to solve γ-wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ (in which the gap is only bounded from below), then it suffices to solve $\gamma^{\approx}$-wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ (in which the gap is almost known). It relies on sparsification.

Lemma 6.5 (ERH). *Let $\gamma, \gamma' > 1$ satisfying $\gamma' \geq 2\log(\Delta_K)^{O(1/d)} \cdot \gamma$. Then γ'-wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ reduces to $\gamma^{\approx}$-wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$. The reduction runs in time polynomial in $(\log \Delta_K)^{O(1)}$ and its input bitsize and succeeds with probability $\Omega(1/(d^2 + \log \Delta_K))$.*

Using the Rényi divergence, it is possible to relate the success probability of an algorithm towards solving uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ for samples from D_γ^{uSVP} with the same probability for $D_{\gamma'}^{\mathrm{uSVP}}$, when γ and γ' are sufficiently close.

Lemma 6.6. *Let $\gamma, \gamma', \gamma'' > 1$ with $\gamma' \in \gamma \cdot [1, 1 + 1/d]$ and $\gamma'' = \gamma/(d\Delta_K^{1/d})^{O(1)}$. Then any algorithm that solves $(D_\gamma^{\mathrm{uSVP}}, \gamma'')$-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ with probability ε also solves $(D_{\gamma'}^{\mathrm{uSVP}}, \gamma'')$-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ with probability $\Omega(\varepsilon^2)$.*

Equipped with the latter result, we are now able to state the worst-case to average case component of the reduction.

Lemma 6.7 (ERH). *Let $\gamma, \gamma', \gamma'' > 1$ with $\gamma' = \gamma \cdot (d\Delta_K^{1/d})^{O(1)}$ and $\gamma'' = \gamma/(d\Delta_K^{1/d})^{O(1)}$. Then there is a reduction from $\gamma^{\approx}$-wc-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ to $(D_{\gamma'}^{\mathrm{uSVP}}, \gamma'')$-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$. The reduction runs in time polynomial in $\log \Delta_K$ and the input bitsize, and if the $(D_{\gamma'}^{\mathrm{uSVP}}, \gamma'')$-uSVP$_{\mathrm{mod}}^{\mathcal{N}}$ oracle succeeds with probability $\varepsilon \geq 2^{-o(d)}$, then the reduction succeeds with probability $\varepsilon^{O(1)}$.*

Acknowledgments. The authors thank Koen de Boer, Guillaume Hanrot and Aurel Page for insightful 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). The last author was supported in part by the European Union Horizon 2020 Research and Innovation Program Grant 780701.

References

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

ALNS20. D. Aggarwal, J. Li, P. Q. Nguyen, and N. Stephens-Davidowitz. Slide reduction, revisited - filling the gaps in SVP approximation. In *CRYPTO*, 2020.

Apo98. T. M. Apostol. *Introduction to analytic number theory.* 1998.

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

BFH17. J.-F. Biasse, C. Fieker, and T. Hofmann. On the computation of the HNF of a module over the ring of integers of a number field. *J Symb Comput*, 2017.

BP91. W. Bosma and M. Pohst. Computations with finitely generated modules over Dedekind domains. In *ISSAC*, 1991.

BS96. E. Bach and J. O. Shallit. *Algorithmic Number Theory: Efficient Algorithms.* 1996.

BSW16. S. Bai, D. Stehlé, and W. Wen. Improved reduction from the bounded distance decoding problem to the unique shortest vector problem in lattices. In *ICALP*, 2016.

CDH$^+$20. C. Chen, O. Danba, J. Hoffstein, A. Hülsing, J. Rijneveld, T. Saito, J. M. Schanck, P. Schwabe, W. Whyte, K. Xagawa, T. Yamakawa, and Z. Zhang. NTRU: A submission to the NIST post-quantum standardization effort. Available at https://www.ntru.org/, 2020.

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

Coh96. H. Cohen. Hermite and Smith normal form algorithms over Dedekind domains. *Math of Comput*, 1996.

CS97. D. Coppersmith and A. Shamir. Lattice attacks on NTRU. In *EUROCRYPT*, 1997.

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

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

GAL13. M. Gil, F. Alajaji, and T. Linder. Rényi divergence measures for commonly used univariate continuous distributions. *Inform Sciences*, 2013.

GN08. N. Gama and P. Q. Nguyen. Finding short lattice vectors within Mordell's inequality. In *STOC*, 2008.

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

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

HPS11. G. Hanrot, X. Pujol, and D. Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In *CRYPTO*, 2011.

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

Kho06. S. Khot. Hardness of approximating the shortest vector problem in high ℓ_p norms. *J Comput Sys Sci*, 2006.

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

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

LPSW19. C. Lee, A. Pellet-Mary, D. Stehlé, and A. Wallet. An LLL algorithm for module lattices. In *ASIACRYPT*, 2019.

LS15. A. Langlois and D. Stehlé. Worst-case to average-case reductions for module lattices. *Design Code Cryptogr*, 2015.
MR07. D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J Comput*, 2007.
Neu99. J. Neukirch. *Algebraic number theory*. 1999.
Pei16. C. Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 2016.
PHS19. A. Pellet-Mary, G. Hanrot, and D. Stehlé. Approx-SVP in ideal lattices with pre-processing. In *EUROCRYPT*, 2019.
PRS17. C. Peikert, O. Regev, and N. Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In *STOC*, 2017.
PS21. A. Pellet-Mary and D. Stehlé. On the hardness of the NTRU problem. In *ASIACRYPT*, 2021.
Sch87. C.-P. Schnorr. A hierarchy of polynomial lattice basis reduction algorithms. *Theor Comput Sci*, 1987.
Sit10. B. D. Sittinger. The probability that random algebraic integers are relatively r-prime. *J Number Theory*, 2010.
SS13. D. Stehlé and R. Steinfeld. Making NTRUEncrypt and NTRUSign as secure as standard worst-case problems over ideal lattices. Available at https://eprint.iacr.org/2013/004, 2013.
SSTX09. D. Stehlé, R. Steinfeld, K. Tanaka, and K. Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, 2009.
Ste15. N. Stephens-Davidowitz. Dimension-preserving reductions between lattice problems. Available at http://noahsd.com/latticeproblems.pdf, 2015.

A Properties of the Rényi Divergence

We will use the following result that bounds the Rényi divergence between two zero-centered normal distributions over $K_{\mathbb{R}}$. It follows from standard divergence bounds on Gaussians such as in [GAL13, Table 2] (note that in this work the Rényi divergence is the logarithm of ours).

Lemma A.1. *Let $a, b \in K_{\mathbb{R}}^{+}$. Let $\mathbf{a} = (\sigma_i(a))_{i \in [r_1+r_2]}$ and $\mathbf{b} = (\sigma_i(b))_{i \in [r_1+r_2]}$. If $2b_i - a_i > 0$ for all $i \in [r_1 + r_2]$, then we have*

$$\mathrm{RD}(\mathcal{D}_{K_{\mathbb{R}}}(0, \mathbf{a}) \parallel \mathcal{D}_{K_{\mathbb{R}}}(0, \mathbf{b})) \leq \mathcal{N}\left(\frac{b^2}{a(2b-a)}\right)^{\frac{1}{2}}.$$

We will also use the following technical lemma on the Rényi divergence of a product of random variables.

Lemma A.2. *Let X, Y be independent random variables in $\mathbb{R}$ with probability distributions D_X, D_Y. Assume that D_X is non-zero over $\mathbb{R}$ (whereas Y can even be discrete). Then*

$$\mathrm{RD}(X \cdot Y \parallel X) \leq \left(\mathbb{E}_{y \sim D_Y}\left(\mathrm{RD}(X \cdot y \parallel X)\right)^{\frac{1}{2}}\right)^2.$$

Proof. Let D' be the distribution probability of $X \cdot Y$. We have, for all $t \in \mathbb{R}$:

$$D'(t) = \int_y D_Y(y) D_X\left(\frac{t}{y}\right) \mathrm{d}y\,.$$

This implies that:

$$\begin{aligned}\mathrm{RD}(X \cdot Y \parallel X) &= \int_t \frac{1}{D_X(t)} \left(\int_y D_Y(y) D_X\left(\frac{t}{y}\right) \mathrm{d}y\right)^2 \mathrm{d}t \\ &= \int_{y_1, y_2} D_Y(y_1) D_Y(y_2) \int_t \frac{D_X\left(\frac{t}{y_1}\right) D_X\left(\frac{t}{y_2}\right)}{D_X(t)} \mathrm{d}t\, \mathrm{d}y_1\, \mathrm{d}y_2\,.\end{aligned}$$

By the Cauchy-Schwartz inequality, we have

$$\begin{aligned}\left(\int_t \frac{D_X\left(\frac{t}{y_1}\right) D_X\left(\frac{t}{y_2}\right)}{D_X(t)} \mathrm{d}t\right)^2 &\leq \int_t \frac{\left(D_X\left(\frac{t}{y_1}\right)\right)^2}{D_X(t)} \mathrm{d}t \cdot \int_t \frac{\left(D_X\left(\frac{t}{y_2}\right)\right)^2}{D_X(t)} \mathrm{d}t \\ &= \mathrm{RD}(X \cdot y_1 \parallel X) \cdot \mathrm{RD}(X \cdot y_2 \parallel X).\end{aligned}$$

Overall, we obtain that

$$\mathrm{RD}(X \cdot Y \parallel X) \leq \left(\int_y D_Y(y) \Big(\mathrm{RD}(X \cdot y \parallel X)\Big)^{\frac{1}{2}} \mathrm{d}y\right)^2,$$

which completes the proof. ◻

B Missing proofs from Section 2

B.1 Proof of Lemma 2.4

Let $\zeta' \in E$ sampled from the centered normal law with standard deviation $d^{-3/2}$, $z' = \exp(\zeta')$. We use the notations from [dBDPW20] and instantiate [dBDPW20, Theorem 3.3] with $s = d^{-3/2}$, $\varepsilon = 2^{-d}$, $N = 1$ and

$$k = \frac{\Theta(d \log d) + \log\big(\mathrm{Vol}(\mathrm{Pic}_K^0)\big)}{\log d}.$$

By the log-unit lattice smoothing analysis from [dBDPW20, Appendix B.1], the condition on N in [dBDPW20, Theorem 3.3] is satisfied. Now, note that the bound on $\mathrm{Vol}(\mathrm{Pic}_K^0)$ in [dBDPW20, Lemma 2.3] implies that $k \leq O(d + \log \Delta_K / \log d)$. Therefore, our lower bound on B implies the one in [dBDPW20, Theorem 3.3]. By [dBDPW20, Theorem 3.3], we deduce that the distribution of the projection of $z' \cdot \mathfrak{p}/\mathcal{N}^{1/d}(\mathfrak{p})$ into Pic_K^0 is within 2^{-d} statistical distance from $\mathcal{U}(\mathrm{Pic}_K^0)$, implying by Gaussian tail-bounds that the distribution of $\exp(\zeta) \cdot \mathfrak{p}/\mathcal{N}^{1/d}(\mathfrak{p})$ is within $2^{-\Omega(d)}$ statistical distance from $\mathcal{U}(\mathrm{Pic}_K^0)$. The proof can be completed by using [dBDPW20, Lemma 2.7]. □

B.2 Equivalence of the Conditions in Definition 2.5

Assume that N is maximal for the inclusion. Let $m = \mathrm{rank}(M)$ and $k = \mathrm{rank}(N)$ and write $N = \sum_{i\in[k]} \mathbf{c}_i J_i$. By [FS10, Theorem 4], there exists a pseudo-basis $(\mathbf{b}_i, I_i)_{i\in[m]}$ of M such that $\mathrm{span}_{i\in[k]}(\mathbf{b}_i I_i) = \mathrm{span}_{i\in[k]}(\mathbf{c}_i J_i) = \mathrm{span}(N)$. By maximality of N this implies that $N = \sum_{i\in[k]} \mathbf{b}_i \cdot I_i$. Taking $N' = \sum_{i>k} \mathbf{b}_i \cdot I_i$ allows to conclude that $M = N + N'$ and $\mathrm{rank}(M) = \mathrm{rank}(N) + \mathrm{rank}(N')$.

Now, assume that there is a module N' with $M = N + N'$ and $\mathrm{rank}(M) = \mathrm{rank}(N) + \mathrm{rank}(N')$. As $N \subseteq M$, we have $N \subseteq M \cap \mathrm{span}_K(N)$. Further, by the rank equality, we must have $N' \cap \mathrm{span}_K(N) = \{\mathbf{0}\}$. Then we have

$$N \subseteq M \cap \mathrm{span}_K(N) = (N + N') \cap \mathrm{span}_K(N) = N \cap \mathrm{span}_K(N) \subseteq N.$$

Finally, assume that $N = M \cap \mathrm{span}_K(N)$. Let P with $\mathrm{rank}(P) = \mathrm{rank}(N)$ and $N \subseteq P$. We have that $\mathrm{span}_K(N) \subseteq \mathrm{span}_K(P)$, and hence $\mathrm{span}_K(N) = \mathrm{span}_K(P)$ by equality of the dimensions. Then we have

$$N \subseteq P \subseteq M \cap \mathrm{span}_K(P) = M \cap \mathrm{span}_K(N) = N.$$

This completes the equivalency proof. □

B.3 Proof of Lemma 2.6

Let k denote the rank of M. By Minkowski's theorem, there exists a non-zero vector in M of ℓ_2-norm $\leq \sqrt{kd} \cdot (\det M)^{1/(kd)}$. By considering the rank-1 module that it spans, we obtain that $\lambda_1^{\mathcal{N}}(M) \leq (kd)^{d/2} \cdot (\det M)^{1/k}$. Now, by using Minkowski's theorem again, we obtain that all rank-1 submodules of norm $\le (kd)^{d/2} \cdot (\det M)^{1/k}$ contain a non-zero vector of M of ℓ_2-norm $\le \sqrt{k}d \cdot \Delta_K^{1/(2d)} \cdot (\det M)^{1/(kd)}$. By discreteness of M, the non-zero vectors $\{\mathbf{s}_i\}_{i\ge1}$ of M with ℓ_2-norm $\le \sqrt{k}d \cdot \Delta_K^{1/(2d)} \cdot (\det M)^{1/(kd)}$ form a finite set. Now, we can consider all the maximal rank-1 submodules of M containing at least one of these vectors. By Condition 3 of Definition 2.5, two maximal rank-1 submodules of M containing the same vector $\mathbf{s}_i$ must be equal, hence there are only finitely many such submodules. This allows us to conclude that the infimum corresponding to $\lambda_1^{\mathcal{N}}(M)$ is over a finite set and must be reached. □

B.4 Proof of Lemma 2.7

We first recall the lemma statement.

Lemma B.1. *There exists an absolute polynomial P such that the following holds. For any number field K of degree d, integer $k \ge 2$, rank-k module $M \subset K_{\mathbb{R}}^k$ and real number $\delta \ge 0$, if $\mathbf{c} \in \mathrm{span}_{K_{\mathbb{R}}}(M)$ and $\varsigma \in \mathbb{R}_{>0}$ are such that $\|\mathbf{c}\| \le \delta \cdot \varsigma$ and $\varsigma \ge \lambda_{kd}(M) \cdot P(\Delta_K^{1/d}, k, d, \delta, \lambda_{kd}(M)/\lambda_1(M))$, then it holds that*

$$\Pr_{\mathbf{v} \leftarrow D_{M,\varsigma,\mathbf{c}}} \big(\mathbf{v} \cdot \mathcal{O}_K \text{ is primitive in } M\big) \ge \frac{1}{4\zeta_K(k)},$$

where $\zeta_K(\cdot)$ is the Dedekind zeta function of the number field K.

Before proving the lemma, we recall some facts regarding the Dedekind zeta function (see, e.g., [Neu99, Chapter 7] for more details). First, let us define the Möbius function of a field K. It is defined over integral ideals of $\mathcal{O}_K$ by

$$\mu_K\left(\prod_{i=1}^r \mathfrak{p}_i^{e_i}\right) := \begin{cases} 1 & \text{if } r = 0 \\ (-1)^r & \text{if } e_1 = \cdots = e_r = 1 \\ 0 & \text{otherwise} \end{cases}$$

where the $\mathfrak{p}_i$'s are distinct prime ideals. For any $s > 1$, the two following equations holds, where the sums are over integral ideals of $\mathcal{O}_K$:

$$\zeta_K(s) = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{1}{\mathcal{N}(\mathfrak{a})^s}$$
$$\zeta_K(s)^{-1} = \sum_{\mathfrak{a} \subseteq \mathcal{O}_K} \frac{\mu_K(\mathfrak{a})}{\mathcal{N}(\mathfrak{a})^s}.$$

The Dedekind zeta function is well-defined for any $s > 1$ (i.e., the sums above are absolutely converging for $s > 1$).

Lemma B.2. *Let $H(N) := |\{\mathfrak{a} \subseteq \mathcal{O}_K \text{ ideal} \mid \mathcal{N}(\mathfrak{a}) \le N\}|$, for $N \ge 0$. For any $s > 1$ and number field K, it holds that $H(N) \le \zeta_K(s) \cdot N^s$.*

Proof. This follows from $\zeta_K(s) \ge \sum_{\mathfrak{a}, \mathcal{N}(\mathfrak{a}) \le N} \frac{1}{\mathcal{N}(\mathfrak{a})^s} \ge H(N)/N^s$. □

Lemma B.3. *For any $s \geq 3/2$ and degree-d number field K, it holds that $\zeta_K(s) \leq 2^{2d}$.*

Proof. We use the Euler product form of the Dedekind zeta function

$$\begin{aligned}
\zeta_K(s) &= \prod_{\substack{\mathfrak{p}\subset\mathcal{O}_K \\ \mathfrak{p}\text{ prime}}} \frac{1}{1-\mathcal{N}(\mathfrak{p})^{-s}} \\
&= \prod_{\substack{p\in\mathbb{Z} \\ p\text{ prime}}} \prod_{\substack{\mathfrak{p}|p\mathcal{O}_K \\ \mathfrak{p}\text{ prime}}} \frac{1}{1-\mathcal{N}(\mathfrak{p})^{-s}} \\
&\leq \prod_{\substack{p\in\mathbb{Z} \\ p\text{ prime}}} \prod_{\substack{\mathfrak{p}|p\mathcal{O}_K \\ \mathfrak{p}\text{ prime}}} \frac{1}{1-p^{-s}} \\
&\leq \prod_{\substack{p\in\mathbb{Z} \\ p\text{ prime}}} \left(\frac{1}{1-p^{-s}}\right)^d \\
&\leq \zeta_{\mathbb{Q}}(s)^d \leq \zeta_{\mathbb{Q}}(3/2)^d \leq 4^d,
\end{aligned}$$

where we used the fact that $\zeta_{\mathbb{Q}}(3/2) \approx 2.6 \leq 4$. □

Finally, we will use the following notations and facts regarding Gaussian distributions. We let $\rho_{\varsigma,\mathbf{c}}(\mathbf{x}) = \exp\big(-\pi\|\mathbf{x}-\mathbf{c}\|^2/\varsigma^2\big)$. We also write $D_{M,\varsigma,\mathbf{c}}(\mathbf{x})$ the probability that the Gaussian distribution $D_{M,\varsigma,\mathbf{c}}$ outputs the vector $\mathbf{x}$, i.e., $D_{M,\varsigma,\mathbf{c}}(\mathbf{x}) = \rho_{\varsigma,\mathbf{c}}(\mathbf{x})/\rho_{\varsigma,\mathbf{c}}(M)$:

- From [Ban93, Lemma 1.5], we know that for any rank-n lattice L and $c > 1/\sqrt{2\pi}$, we have $\rho_{\varsigma,\mathbf{c}}(\{\mathbf{v} \in L \,|\, \|\mathbf{v}-\mathbf{c}\| > c\cdot\sqrt{n}\cdot\varsigma\}) \leq 2C^n \cdot \rho_\varsigma(L)$, where $C = c\cdot\sqrt{2\pi\mathrm{e}}\cdot\mathrm{e}^{-\pi\cdot c^2} < 1$.
- From [MR07, Lemma 3.3], we know that for any $\varepsilon > 0$ and rank-n lattice L, the smoothing parameter of L satisfies $\eta_\varepsilon(L) \leq \sqrt{\ln(2n(1+1/\varepsilon))/\pi}\cdot\lambda_n(L)$.
- Finally, from the proof of [MR07, Lemma 4.4], we know that if $\varsigma \geq \eta_\varepsilon(L)$, then it holds that $\rho_{\varsigma,\mathbf{c}}(L) \in [1-\varepsilon, 1+\varepsilon]\cdot\varsigma^n/\det(L)$.

Proof (Lemma 2.7). We follow the same proof structure as in [SS13, Lemma 4.4]. Let us fix some number field K of degree d, integer $k \geq 2$, rank-k module $M \subset K_{\mathbb{R}}^k$ and real number $\delta \geq 0$. Let

$$\varsigma_0 = 2^{45}\cdot\Delta_K^{7/(2d)}\cdot k^8\cdot d^4\cdot(k^2\cdot d^2+\delta^3)\cdot\lambda_1(M)^{-3}\cdot\lambda_{kd}(M)^4.$$

Observe that $\varsigma_0 = \lambda_{kd}(M)\cdot P(\Delta_K^{1/d}, k, d, \delta, \lambda_{kd}(M)/\lambda_1(M))$ for some absolute polynomial P. We will prove that the lemma holds for this polynomial P.

Let us then fix some ς and $\mathbf{c} \in \mathrm{span}_{K_{\mathbb{R}}}(M)$ such that $\|\mathbf{c}\| \leq \delta\cdot\varsigma$ and $\varsigma \geq \varsigma_0$. We define the following quantities.

$$\varepsilon = 2^{-2kd-5}$$
$$B_1 = \frac{\varsigma^d}{\sqrt{\Delta_K}\cdot\lambda_{kd}(M)^d\cdot\max\left(\varepsilon^{-1/k},(2\ln(2/\varepsilon))^d\right)}$$
$$B_2 = \left(2\cdot\sqrt{kd}\cdot\varsigma+\|\mathbf{c}\|\right)^d\cdot\sqrt{\Delta_K}\cdot\lambda_1(M)^{-d}.$$

With these notations, we are ready to bound from below the probability that a Gaussian element in M is primitive. To do so, observe that for $\mathbf{v}\in M$, the rank-1 module $\mathbf{v}\mathcal{O}_K$ is not primitive in M if an only if there exists some prime ideal $\mathfrak{p}$ such that $\mathbf{v}\in\mathfrak{p}\cdot M$. Indeed, let us define $I=\{x\in K\,|\,x\cdot\mathbf{v}\in\mathcal{O}_K\}$. One can check that I is a fractional ideal with $\mathcal{O}_K\subseteq I$. Moreover, by definition of a primitive submodule, we have $I=\mathcal{O}_K$ if and only if $\mathbf{v}\cdot\mathcal{O}_K$ is a primitive submodule of M. Let $\mathfrak{a}=I^{-1}$, which is an integral ideal. By definition of I and $\mathfrak{a}$, we have that $\mathbf{v}\in\mathfrak{a}\cdot M$. If $\mathbf{v}\cdot\mathcal{O}_K$ is not primitive, then $\mathfrak{a}\neq\mathcal{O}_K$ so there exists $\mathfrak{p}|\mathfrak{a}$, and it holds that $\mathbf{v}\in\mathfrak{p}\cdot M$. Reciprocally, if $\mathbf{v}\in\mathfrak{p}\cdot M$ for some prime ideal $\mathfrak{p}$, then $\mathbf{v}\cdot\mathfrak{p}^{-1}\subset M$, so $I\neq\mathcal{O}_K$ and $\mathbf{v}\cdot\mathcal{O}_K$ is not primitive in M.

From this observation, we can rewrite

$$\begin{aligned}\Pr_{\mathbf{v}\leftarrow D_{M,\varsigma,\mathbf{c}}}\Big(\mathbf{v}\cdot\mathcal{O}_K\text{ is primitive in }M\Big) &= D_{M,\varsigma,\mathbf{c}}\Big(M\setminus\bigcup_{\mathfrak{p}\text{ prime}}\mathfrak{p}M\Big)\\ &\geq D^T_{M,\varsigma,\mathbf{c}}\Big(M\setminus\bigcup_{\mathfrak{p}\text{ prime}}\mathfrak{p}M\Big),\end{aligned}$$

with $D^T_{M,\varsigma,\mathbf{c}}$ being the truncated Gaussian function, defined as $D^T_{M,\varsigma,\mathbf{c}}(\mathbf{v}) = D_{M,\varsigma,\mathbf{c}}(\mathbf{v})$ if $\mathbf{v}\neq\mathbf{0}$ and $\|\mathbf{v}-\mathbf{c}\|\leq 2\cdot\sqrt{kd}$, and $D^T_{M,\varsigma,\mathbf{c}}(\mathbf{v})=0$ otherwise (note that $D^T_{M,\varsigma,\mathbf{c}}$ does not sum to 1 and is not a probability distribution).

Let us then focus on $p := D^T_{M,\varsigma,\mathbf{c}}\big(M\setminus\bigcup_{\mathfrak{p}\text{ prime}}\mathfrak{p}M\big)$. Observe that for any distinct prime ideals $\mathfrak{p}_1,\ldots,\mathfrak{p}_t$, it holds that $\cap_{i\leq t}(\mathfrak{p}_i\cdot M)=(\prod_{i\leq t}\mathfrak{p}_i)\cdot M$. Hence, from the inclusion-exclusion principle, we obtain

$$p=\sum_{\mathfrak{a}\subseteq\mathcal{O}_K}\mu_K(\mathfrak{a})\cdot D^T_{M,\varsigma,\mathbf{c}}\Big(\mathfrak{a}\cdot M\Big)=\sum_{\substack{\mathfrak{a}\subseteq\mathcal{O}_K\\ \mathcal{N}(\mathfrak{a})\leq B_2}}\mu_K(\mathfrak{a})\cdot D^T_{M,\varsigma,\mathbf{c}}\Big(\mathfrak{a}\cdot M\Big),$$

where the sums are above the integral ideals $\mathfrak{a}\subseteq\mathcal{O}_K$ and B_2 was defined at the start of the proof. The second equality comes from the fact that if $\mathcal{N}(\mathfrak{a})>B_2$, then $\mathfrak{a}\cdot M$ does not contain any non-zero vector shorter than $2\sqrt{kd}\cdot\varsigma+\|\mathbf{c}\|$ (since otherwise $M=\mathfrak{a}^{-1}\cdot(\mathfrak{a}\cdot M)$ would contain a non-zero vector smaller than $(2\sqrt{kd}\cdot\varsigma+\|\mathbf{c}\|)\cdot\mathcal{N}(\mathfrak{a})^{-1/d}\cdot\Delta_K^{1/(2d)}<\lambda_1(M)$, contradicting the definition of $\lambda_1(M)$). This implies that $\mathfrak{a}\cdot M$ does not contain any non-zero vector in the ball $\{\mathbf{v}\,|\,\|\mathbf{v}-\mathbf{c}\|\leq 2\sqrt{kd}\cdot\varsigma\}$, hence $D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a}\cdot M)=0$.

Combining this with the equation recalled before the proof relating the Dedekind zeta function and the Möbius function, we obtain

$$|p - \zeta_K(k)^{-1}| \leq \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_1}} \left| D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) - \mathcal{N}(\mathfrak{a})^{-k} \right| + \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ B_1 < \mathcal{N}(\mathfrak{a}) \leq B_2}} D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) + \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) > B_1}} \mathcal{N}(\mathfrak{a})^{-k}$$

Note that by definition of B_1 and B_2, it holds that $B_1 \leq B_2$. We will bound each one of the three sums above by $\zeta_K(k)^{-1}/4$, which will prove the result.

Let us start with the first sum. Let $\mathfrak{a}$ be an integral ideal with $\mathcal{N}(\mathfrak{a}) \leq B_1$. We know that $\lambda_{kd}(\mathfrak{a} \cdot M) \leq \lambda_1^\infty(\mathfrak{a}) \cdot \lambda_{kd}(M)$ (since multiplying kd linearly independent vectors from M by a shortest vector of $\mathfrak{a}$ provides kd linearly independent vectors of $\mathfrak{a}M$). Moreover, since $\mathcal{N}(\mathfrak{a}) \leq B_1$, we know that $\lambda_1^\infty(\mathfrak{a}) \leq B_1^{1/d} \cdot \Delta_K^{1/(2d)}$. By definition of B_1 and ε, we have

$$\varsigma \geq B_1^{\frac{1}{d}} \cdot \Delta_K^{\frac{1}{2d}} \cdot \lambda_{kd}(M) \cdot 2\ln(2/\varepsilon) \geq \lambda_{kd}(\mathfrak{a}M) \cdot 2\ln(2/\varepsilon) \geq \eta_\varepsilon(\mathfrak{a}M).$$

Since $\varsigma \geq \eta_\varepsilon(\mathfrak{a}M)$, we know that $\rho_{\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) \in [1-\varepsilon, 1+\varepsilon] \cdot \varsigma^{kd}/\det(\mathfrak{a}M)$. From [Ban93, Lemma 1.5] (recalled above) with $c = 2$, we also know that $\rho_{\varsigma,\mathbf{c}}(\{\mathbf{v} \in \mathfrak{a}M \mid \|\mathbf{v} - \mathbf{c}\| > 2 \cdot \sqrt{kd} \cdot \varsigma\}) \leq \varepsilon \cdot \rho_\varsigma(\mathfrak{a}M) \leq 2\varepsilon \cdot \varsigma^{kd}/\det(\mathfrak{a}M)$.

Recall also that by definition of B_1, we have $\varsigma \geq \Delta_K^{1/(2d)} \cdot \mathcal{N}(\mathfrak{a})^{1/d} \cdot \lambda_{kd}(M) \cdot \varepsilon^{-1/(kd)}$. Hence, we obtain that $\rho_{\varsigma,\mathbf{c}}(\mathbf{0}) \leq \varepsilon \cdot \varsigma^{kd}/\det(\mathfrak{a}M)$. Combining everything, this implies that

$$\rho_{\varsigma,\mathbf{c}}\Big(\mathfrak{a} \cdot M \setminus \{\mathbf{v} \mid \mathbf{v} = \mathbf{0} \text{ or } \|\mathbf{v} - \mathbf{c}\| > 2 \cdot \sqrt{kd} \cdot \varsigma\}\Big) \in [1 - 4\varepsilon, 1 + \varepsilon] \cdot \frac{\varsigma^{kd}}{\det(\mathfrak{a}M)}.$$

By definition of $D^T_{M,\varsigma,\mathbf{c}}$, this implies that

$$D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) \in \left[\frac{1-4\varepsilon}{1+\varepsilon}, \frac{1+\varepsilon}{1-\varepsilon}\right] \cdot \frac{\det(M)}{\det(\mathfrak{a}M)} \subset [1 - 5\varepsilon, 1 + 4\varepsilon] \cdot \frac{1}{\mathcal{N}(\mathfrak{a})^k},$$

where we used the identities $\det(M) = \Delta_K^{k/2} \cdot \mathcal{N}(M)$ and $\det(\mathfrak{a}M) = \Delta_K^{k/2} \cdot \mathcal{N}(\mathfrak{a}M)$. This concludes the upper bound on the first sum

$$\sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_1}} \left| D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) - \mathcal{N}(\mathfrak{a})^{-k} \right| \leq 5\varepsilon \cdot \sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) \leq B_1}} \frac{1}{\mathcal{N}(\mathfrak{a})^k} \leq 5\varepsilon\zeta_K(k) \leq \frac{1}{4\zeta_K(k)},$$

by definition of ε, and using Lemma B.3 to assert that $\zeta_K(k)^2 \leq 2^{4d} \leq 2^{2kd}$.

Let us now consider the second sum. Let $\mathfrak{a}$ be an integral ideal with $B_1 < \mathcal{N}(\mathfrak{a}) \leq B_2$. Let us define $I = \lceil (\mathcal{N}(\mathfrak{a})/B_1)^{1/d} \rceil^{-1} \cdot \mathfrak{a}$. This is a fractional ideal with $\mathcal{N}(I) \in [1/2^d, 1] \cdot B_1$. Moreover, we have $\mathfrak{a} \subseteq I$, hence $D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) \leq$

$D^T_{M,\varsigma,\mathbf{c}}(I \cdot M)$, where we let $D^T_{M,\varsigma,\mathbf{c}}(\mathbf{v}) := \rho_{\varsigma,\mathbf{c}}(\mathbf{v})/\rho_{\varsigma,\mathbf{c}}(M)$ for any $\mathbf{v} \in K_{\mathbb{R}}$, even those not in M (note that since I is a fractional ideal, then $I \cdot M$ needs not be contained in M). Observe that everything we did above when $\mathcal{N}(\mathfrak{a}) \leq B_1$ can be adapted to a fractional ideal I of norm $\mathcal{N}(I) \leq B_1$. Hence, we have (using the analysis for the first sum):

$$D^T_{M,\varsigma,\mathbf{c}}(I \cdot M) \leq (1+4\varepsilon) \cdot \frac{1}{\mathcal{N}(I)^k} \leq \frac{2^{kd+1}}{B_1^k}.$$

We can hence bound the second sum from above by

$$\begin{aligned}\sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ B_1 < \mathcal{N}(\mathfrak{a}) \leq B_2}} D^T_{M,\varsigma,\mathbf{c}}(\mathfrak{a} \cdot M) &\leq \frac{|\{\mathfrak{a} \subseteq \mathcal{O}_K \,|\, \mathcal{N}(\mathfrak{a}) \leq B_2\}| \cdot 2^{kd+1}}{B_1^k} \\ &\leq \frac{\zeta_K(s_0) \cdot B_2^{s_0} \cdot 2^{kd+1}}{B_1^k} \\ &\leq B_1^{-(k-s_0)} \cdot \Big(\zeta_K(s_0) \cdot 2^{kd+1} \cdot (B_2/B_1)^{s_0}\Big),\end{aligned}$$

where we used Lemma B.2 for the second inequality, with some $s_0 \in (1,k)$. Let us choose $s_0 = \max(3/2, k/2)$. This choice of s_0 ensures that $s_0 \in [3/2, k)$ and that $s_0/(k-s_0) \leq 3$ and $k/(k-s_0) \leq 4$ for any $k \geq 2$. Using Lemma B.3, the definitions of s_0, B_1, B_2 and the lower bound on ς, one can check that this is $\leq 1/4 \cdot \zeta_K(k)^{-1}$.

We are finally left with the last sum. Recall that $H(N)$ denotes the number of integral ideals of norm $\leq N$. With this notation, we can rewrite

$$\sum_{\substack{\mathfrak{a} \subseteq \mathcal{O}_K \\ \mathcal{N}(\mathfrak{a}) > B_1}} \mathcal{N}(\mathfrak{a})^{-k} = \sum_{N > B_1} \frac{H(N) - H(N-1)}{N^k} \leq \sum_{N > B_1} H(N) \cdot \Big(\frac{1}{N^k} - \frac{1}{(N+1)^k}\Big).$$

Let us prove that the last sum above is absolutely converging (in order to prove that our transformation was valid). Using Lemma B.2 with $s = 1.5$, we know that $H_N \leq \zeta_K(1.5) \cdot N^{1.5}$, where the quantity $\zeta_K(1.5)$ depends on the number field but is fixed when N tends to infinity. Hence, the quantity inside the sum is bounded by $O(N^{1.5} \cdot N^{k-1}/N^{2k}) = O(1/N^{k-0.5}) = O(1/N^{1.5})$ since $k \geq 2$. We conclude that the sum is converging absolutely as desired.

Let us now compute an upper bound on this sum. Since $N \geq B_1 \geq k$, we know that $(N+1)^k - N^k \leq k^2 \cdot N^{k-1}$. Applying Lemma B.2 again with $s = s_0 \in (1,k)$, we obtain

$$\begin{aligned}\sum_{N > B_1} H(N) \cdot \Big(\frac{1}{N^k} - \frac{1}{(N+1)^k}\Big) &\leq \zeta_K(s_0) \cdot \sum_{N > B_1} \frac{k^2 \cdot N^{s_0+k-1}}{N^{2k}} \\ &\leq \zeta_K(s_0) \cdot k^2 \cdot \int_{\lfloor B_1 \rfloor}^{+\infty} x^{-(k+1-s_0)} \mathrm{d}x \\ &= \zeta_K(s_0) \cdot k^2 \cdot (k-s_0)^{-1} \cdot \lfloor B_1 \rfloor^{-(k-s_0)}\end{aligned}$$

Using $s_0 = \max(3/2, k/2)$ again, the definitions of B_1, the lower bound on ς and Lemma B.3, one can check that $\sum_{\substack{\mathfrak{a}\subseteq\mathcal{O}_K \\ \mathcal{N}(\mathfrak{a})>B_1}} \mathcal{N}(\mathfrak{a})^{-k} \leq 1/4 \cdot \zeta_K(k)^{-1}$ as desired.

□

B.5 Proof of Lemma 2.8

Let $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ be a pseudo-basis of M and write $N = \mathbf{s}_1 J_1$ and $N' = \mathbf{s}_2 J_2$. There exists $\mathbf{Z} \in K^{2\times 2}$ such that $\mathbf{S} = \mathbf{BZ}$. Assume by contradiction that $\mathrm{span}_K(N') \neq \mathrm{span}_K(N)$. In that case, the matrix $\mathbf{Z}$ has rank 2, the vectors $\mathbf{s}_1$ and $\mathbf{s}_2$ are $K_{\mathbb{R}}$-linearly independent and $M' = \mathbf{s}_1 J_1 + \mathbf{s}_2 J_2$ is a submodule of M. By using a QR-factorization $\mathbf{S} = \mathbf{QR}$, one sees that $\det(\mathbf{S}) = \mathcal{N}(r_{11})\mathcal{N}(r_{22})$ and $\mathcal{N}(M') \leq \mathcal{N}(N)\mathcal{N}(N')$. We hence obtain:

$$\mathcal{N}(M) \leq \mathcal{N}(M') \leq \mathcal{N}(N)\mathcal{N}(N') < \frac{\sqrt{\mathcal{N}(M)}}{\gamma^d}\left(\gamma^d\sqrt{\mathcal{N}(M)}\right) = \mathcal{N}(M),$$

which gives a contradiction. We thus have that $\mathrm{span}_K(N') = \mathrm{span}_K(N)$. Definition 2.5 allows us to conclude that $N' \subseteq N$.

Assume now that $\gamma > 1$. Then the first statement implies that the densest rank-1 submodule N is unique. Let $\mathbf{b} \in M$ with $0 < \|\mathbf{b}\| < \gamma \cdot \mathcal{N}(M)^{1/(2d)}$. Then $\mathbf{b}\mathcal{O}_K$ is a rank-1 submodule of M and

$$\mathcal{N}(\mathbf{b}\mathcal{O}_K) \leq \|\mathbf{b}\|^d < \gamma^d \cdot \sqrt{\mathcal{N}(M)}.$$

By the above, we must have $\mathbf{b}\mathcal{O}_K \subseteq N$, which is equivalent to $\mathbf{b} \in N$. □

B.6 Proof of Lemma 2.9

Let N be a densest rank-1 submodule of M. By Definition 2.5, there exists a rank-1 submodule N' such that $M = N + N'$. Equivalently, we obtain a pseudo-basis $((I_1, \mathbf{b}_1), (I_2, \mathbf{b}_2))$ of M such that $N = \mathbf{b}_1 I_1$. Wlog, we may assume that $\mathcal{N}(I_1) = \mathcal{N}(I_2) = 1$, by multiplying $\mathbf{b}_i$ by $\mathcal{N}^{-1/d}(I_i)$ for $i \in [2]$. Let $\mathbf{Q} \in \mathcal{O}_2(K_{\mathbb{R}})$ and $\mathbf{R} \in K_{\mathbb{R}}^{2\times 2}$ upper triangular such that $\mathbf{B} = \mathbf{QR}$. Let $\mathbf{D}$ be the diagonal matrix with $d_1 = r_{11}/\mathcal{N}^{1/d}(r_{11})$ and $d_2 = r_{22}/\mathcal{N}^{1/d}(r_{22})$ as diagonal coefficients. Let $J_1 = d_1 I_1$, $J_2 = d_2 I_2$ and

$$\mathbf{B}' = \frac{\mathcal{N}^{\frac{1}{2d}}(M)}{\gamma} \cdot \mathbf{Q} \cdot \mathbf{R} \cdot \Big(\frac{\gamma}{\mathcal{N}^{\frac{1}{2d}}(M)}\mathbf{D}^{-1}\Big).$$

It now suffices to prove that $((J_1, \mathbf{b}'_1), (J_2, \mathbf{b}'_2))$ is a pseudo-basis of M of the desired form, i.e., to check that $\mathbf{R}' = \mathbf{R} \cdot (\gamma/\mathcal{N}^{1/(2d)}(M))\mathbf{D}^{-1}$ has diagonal coefficients equal to 1 and γ. We have $r'_{ii} = \mathcal{N}^{1/d}(r_{ii})(\gamma/\mathcal{N}^{1/(2d)}(M))$ for $i \in [2]$, by construction. The fact that $\mathcal{N}(r_{11}) = \lambda_1^{\mathcal{N}}(M)$ gives that $r'_{11} = 1$. The equality $\mathcal{N}(M) = \det(\mathbf{B}')$ provides the result. □

B.7 Proof of Lemma 2.11

From Banaszczyk's transference theorem [Ban93], we know that $1 \leq \lambda_{2d}(M^\vee) \cdot \lambda_1(M) \leq 2d$. We also know that $\lambda_1(M)^d \geq \sqrt{d} \cdot \lambda_1^{\mathcal{N}}(M)$, since for any vector $\mathbf{v} \in K_\mathbb{R}^2$ it holds that $\|\mathbf{v}\| \geq \sqrt{d} \cdot \mathcal{N}(\mathbf{v} \cdot \mathcal{O}_K)^{1/d}$ (by applying the inequality of arithmetic and geometric means to the squares of the coordinates of $\langle \mathbf{v}, \mathbf{v} \rangle_{K_\mathbb{R}}$). Further, from the definition of the gap of a module M, we know that $\lambda_1^{\mathcal{N}}(M) = \mathcal{N}(M)^{1/2}/\gamma(M)^d$. Combining these relations provides the upper bound on $\lambda_{2d}(M^\vee)$.

In order to get the upper bound on $\lambda_1(M^\vee)^{-1}$, we use the inequality $1 \leq \lambda_{2d}(M) \cdot \lambda_1(M^\vee)$. Hence, it suffices to bound $\lambda_{2d}(M)$ from above. To do so, we use Lemma 2.9. We know that there exist d $\mathbb{Z}$-linearly independent vectors in J_1 of norms $\leq \sqrt{d} \cdot \Delta_K^{1/(2d)} \cdot \delta_K$, and similarly in J_2. Hence, from the representation of M in Lemma 2.9, we obtain $2d$ linearly independent vectors in M of norms $\leq (\gamma(M) \cdot \sqrt{d} + d/\gamma(M)) \cdot \mathcal{N}(M)^{1/(2d)} \cdot \Delta_K^{1/(2d)} \cdot \delta_K$ (where we reduced the last d vectors using the first d ones). Since $\gamma(M) \geq 1$, this implies than $\lambda_{kd}(M) \leq 2d \cdot \gamma(M) \cdot \mathcal{N}(M)^{1/(2d)} \cdot \Delta_K^{1/(2d)} \cdot \delta_K$. □

C Missing proofs from Section 3

C.1 Proof of Lemma 3.2

The fact that $\overline{\mathbf{b}^\vee} \neq \mathbf{0}$ implies that the map $\mathbf{m} \mapsto \langle \mathbf{b}^\vee, \mathbf{m} \rangle_{K_\mathbb{R}}$ is a surjective homomorphism from M to $\mathcal{O}_K/\mathfrak{p}$ whose kernel is M'. This gives the following exact sequence of $\mathcal{O}_K$-modules:

$$0 \to M' \to M \to \mathcal{O}_K/\mathfrak{p} \to 0.$$

Now, note that $\mathcal{O}_K/\mathfrak{p}$ is isomorphic to the finite field of size $\mathcal{N}(\mathfrak{p})$. The exact sequence and the finiteness of $\mathcal{O}_K/\mathfrak{p}$ imply that $\mathcal{N}(M') = \mathcal{N}(M) \cdot |\mathcal{O}_K/\mathfrak{p}|$. The proof is completed by noting that $|\mathcal{O}_K/\mathfrak{p}| = \mathcal{N}(\mathfrak{p})$. □

C.2 Proof of Lemma 3.3

The fact that $\mathfrak{p} \cdot M \subset M'$ implies that $\mathbf{b}\mathfrak{p}I \subset M'$. We now prove the second property. As $\mathbf{b}I$ is primitive, there exists a pseudo-basis $(\mathbf{B}, \mathbb{I})$ of M such that $\mathbf{b}_1 = \mathbf{b}$ and $I_1 = I$ (see Definition 2.5). We start by noting that $\langle M^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \cdot I = \mathcal{O}_K$. Indeed, as $(\mathbf{B}, \mathbb{I})$ is a pseudo-basis of M, we have that $(\mathbf{B}^{-\dagger} \cdot \mathbb{J})$ is a pseudo-basis of $M^\vee$, with $J_i = (\overline{I_i})^{-1}$ for all i. Therefore:

$$\langle M^\vee, \mathbf{b} \rangle_{K_\mathbb{R}} \cdot I = \sum_i \left\langle \mathbf{b}_i^{-\dagger}, \mathbf{b}_1 \right\rangle_{K_\mathbb{R}} \cdot \mathcal{O}_K = \mathcal{O}_K.$$

The fact that $\langle M^\vee, \mathbf{b} \rangle_{K_\mathbb{R}} \cdot I = \mathcal{O}_K$ implies that the scalar product with $\mathbf{b}$ is a surjective homomorphism $M^\vee \to I^{-1}$. This induces a surjective homomorphism $M^\vee/\mathfrak{p}M^\vee \to I^{-1}/\mathfrak{p}I^{-1}$. Because of their respective ranks as $\mathcal{O}_K$-modules, the cardinality of $I^{-1}/\mathfrak{p}I^{-1}$ is $\mathcal{N}(\mathfrak{p})$ and the cardinality of $M^\vee/\mathfrak{p}M^\vee$ is $\mathcal{N}(\mathfrak{p})^k$. Lagrange's theorem (for groups) then implies that every element of $I^{-1}/\mathfrak{p}I^{-1}$ has exactly $\mathcal{N}(\mathfrak{p})^k/\mathcal{N}(\mathfrak{p}) = \mathcal{N}(\mathfrak{p})^{k-1}$ pre-images in $M^\vee/\mathfrak{p}M^\vee$ by this application. In particular, the zero element of $I^{-1}/\mathfrak{p}I^{-1}$ has $\mathcal{N}(\mathfrak{p})^{k-1}$ pre-images, including $\mathbf{0}$. Since $\overline{\mathbf{b}^\vee}$ is uniform in $(M^\vee/\mathfrak{p}M^\vee) \setminus \{\mathbf{0}\}$, this implies that the probability that $\langle \mathbf{b}^\vee, \mathbf{b}\rangle_{K_\mathbb{R}} \in \mathfrak{p}I^{-1}$ is $(\mathcal{N}(\mathfrak{p})^{k-1} - 1)/\mathcal{N}(\mathfrak{p})^k = 1/\mathcal{N}(\mathfrak{p}) - 1/\mathcal{N}(\mathfrak{p})^k$ over the choice of $\overline{\mathbf{b}^\vee}$.

To complete the proof, note that $\langle \mathbf{b}^\vee, \mathbf{b}\rangle_{K_\mathbb{R}} \notin \mathfrak{p}I^{-1}$ is equivalent to $\mathbf{b}I \not\subset M'$, by definition of M'. ☐

C.3 Proof of Lemma 3.4

Let $(\mathbf{B}, \mathbb{I})$ be a pseudo-basis of M with integral coefficient ideals I_i. As seen in Section 2.3, the pair $(\mathbf{B}^{-\dagger}, \mathbb{J})$ is a pseudo-basis of $M^\vee$, where $J_i = (\overline{I_i})^{-1}$ for all i. Take $\mathbf{u} \in \mathbb{J}$ such that $\mathbf{B}^{-\dagger} \cdot \mathbf{u}$ is a representative of $\overline{\mathbf{b}^\vee}$ in $M^\vee$. We have:

$$\begin{aligned} M' &= \left\{ \mathbf{B} \cdot \mathbf{v} : \mathbf{v} \in \mathbb{I} \text{ and } \left(\mathbf{B}^{-\dagger} \cdot \mathbf{u}\right)^\dagger \cdot \mathbf{B} \cdot \mathbf{v} \in \mathfrak{p} \right\} \\ &= \left\{ \mathbf{B} \cdot \mathbf{v} : \mathbf{v} \in \mathbb{I} \text{ and } \langle \mathbf{u}, \mathbf{v}\rangle_{K_\mathbb{R}} \in \mathfrak{p} \right\}. \end{aligned}$$

Let us define

$$N = \left\{ \mathbf{v} \in \mathbb{I} : \langle \mathbf{u}, \mathbf{v}\rangle_{K_\mathbb{R}} = 0 \right\} \quad \text{and} \quad N' = \left\{ \mathbf{v} \in \mathbb{I} : \langle \mathbf{u}, \mathbf{v}\rangle_{K_\mathbb{R}} \in \mathfrak{p} \right\}.$$

We use the $\mathbb{Z}$-basis $(r_i)_{i\in[d]}$ of $\mathcal{O}_K$ to identify N with a $\mathbb{Z}$-lattice corresponding to the orthogonal of an integer vector. A basis of this lattice can be computed in polynomial-time, and the basis vectors provide a set $(\mathbf{n}_i)_{i\in[kd]}$ of (non K-linearly independent) vectors in $\mathcal{O}_K^k$ such that $N = \sum_i \mathbf{n}_i \mathcal{O}_K$. The module N' is the rank-k module generated by the pseudo-basis

$$N' = \sum_{i=1}^{kd} \mathbf{n}_i \mathcal{O}_K + \sum_{i=1}^{k} \mathbf{e}_i \mathfrak{p},$$

where $\mathbf{e}_i$ is the i-th canonical unit vector. From the pairs $\{(\mathcal{O}_K, \mathbf{n}_i)\}_i$ and $\{(\mathfrak{p}, \mathbf{e}_i)\}_i$, we compute a Hermite Normal Form $(\mathbf{B}', \mathbb{I}')$ of the integral module N'. By definition of N', the pair $(\mathbf{B} \cdot \mathbf{B}', \mathbb{I}')$ is a pseudo-basis of M'. ☐

C.4 Proof of Lemma 3.5

In Step 2, we use one of the provable variants of the BKZ algorithm mentioned above, which allows us to obtain a basis $\mathbf{C}$ of $M^\vee$ such that $\max_i \|\mathbf{c}_i\| \le (kd)^{kd/\beta+1} \cdot \lambda_{kd}(M^\vee)$ in time polynomial in the bitsize of the input basis of $M^\vee$ and in 2^β. Note that these analyses of the algorithm under scope only prove that the algorithm solves $(kd)^{kd/\beta}$-SVP (i.e., outputs one short non-zero vector) and do not mention the approximation factor obtained for SIVP (the Shortest Independent Vector Problem). Hence, to obtain an upper bound on $\max_i \|\mathbf{c}_i\|$, we also use the polynomial-time reduction from $(\sqrt{n}\gamma)$-SIVP to γ-SVP for lattices of rank-n (see [Ste15, Page 1]), together with the fact that one can transform any set of n short linearly independent vectors of norm $\leq B$ in a rank-n lattice L into a basis of L with vectors of norms $\leq \sqrt{n} \cdot B$. Now, we observe that $\varsigma \geq \sqrt{kd} \cdot \max_i \|\mathbf{c}_i\|$, hence we can apply Lemma 2.1. This means in particular that the vectors $\mathbf{y}_i$ can be sampled in polynomial time, which completes the runtime analysis.

Let us now prove that the matrix $\mathbf{Y}$ satisfies the conditions of the theorem. First of all, note that since the vectors $\mathbf{y}_i$ are in $M^\vee$, then for all $\mathbf{v} \in M$ we have $\mathbf{Y} \cdot \mathbf{v} \in \mathcal{O}_K^k$, which proves the first point. For the second point, recall that we use a tail-cut distribution $\hat{D}_{\mathbf{C}^\vee,\varsigma,\mathbf{t}_i}$, hence, it holds that $\|\mathbf{y}_i - \mathbf{t}_i\| \leq \sqrt{kd} \cdot \varsigma = \varepsilon \cdot R$, as desired.

Finally, recall from Lemma 2.1 that the distribution $\hat{D}_{\mathbf{C}^\vee,\varsigma,\mathbf{t}_i}$ (which might depend on $\mathbf{C}^\vee$ and hence on $(\mathbf{B}, \mathbb{I})$) is within statistical distance at most $2^{-\Omega(kd)}$ from the Gaussian distribution $D_{M^\vee,\varsigma,\mathbf{t}_i}$, which is independent of the known basis of $M^\vee$. Hence, the distribution of $\mathbf{Y}$ is within statistical distance at most $k \cdot 2^{-\Omega(kd)} = 2^{-\Omega(kd)}$ from a distribution independent of the choice of the pseudo-basis $(\mathbf{B}, \mathbb{I})$. $\square$

C.5 Proof of Lemma 3.6

Wlog, we prove the result for $R = 1$. Note that the operator norm of $\mathbf{E}$ satisfies $\|\mathbf{E}\| \leq k\varepsilon < 1$. Therefore, the matrix $\sum_{i \geq 0}(-\mathbf{E})^i$ is well-defined, and satisfies $\mathbf{Y}^{-1} = \sum_{i \geq 0}(-\mathbf{E})^i$. We have $\mathbf{Y}^{-1} = \mathbf{I}_k + \mathbf{E}'$ with $\mathbf{E}' = -\mathbf{E} + \sum_{i \geq 2}(-\mathbf{E})^i$. Using the operator norm again, we obtain that $\|e'_{ij} + e_{ij}\| \leq (k\varepsilon)^2/(1-k\varepsilon) \leq k\varepsilon$ for all $i, j \in [k]$, by using assumption tha $k\varepsilon \leq 1/2$. This proves the first statement.

By Hadamard's inequality, we have

$$\det(\mathbf{Y}) \leq \left(\sqrt{(1+\varepsilon)^2 + (k-1)\varepsilon^2}\right)^d$$
$$\det\left(\mathbf{Y}^{-1}\right) \leq \left(\sqrt{(1+\varepsilon')^2 + (k-1)\varepsilon'^2}\right)^d,$$

with $\varepsilon' = (k+1)\varepsilon$. Simplifying the expressions using the facts that $\varepsilon' \leq 1$ and $k\varepsilon \leq 1/2$ leads to the second statement. $\square$

D Missing proofs from Section 4

D.1 Proof of Theorem 4.1

Theorem 4.1 is a direct corollary of the following more complete statement.

Theorem D.1. *Let K be a number field of degree d with $\zeta_K(2) = 2^{o(d)}$ and let $\gamma^+ > 0$. There exist three algorithms* `uSVP-to-NTRU`, `LiftVec'` *and* `LiftMod'` *and $q_0 = \mathrm{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+) \in \mathbb{R}_{\geq 0}$ such that the following holds.*

For any $q \geq q_0$, $\gamma_{\mathrm{NTRU}} > 1$, $\gamma_{\mathrm{HSVP}} \geq \sqrt{d}\Delta_K^{1/(2d)}$ and $(\mathbf{B}, \mathbb{I})$ pseudo-basis of a rank-2 module $M \subset K^2$ with $\gamma(M) \leq \gamma^+$, we have

- *Algorithm* `uSVP-to-NTRU` *takes as input* $(\mathbf{B}, \mathbb{I})$*,* q *and* γ_{HSVP} *and outputs a pseudo-basis* $(\mathbf{B}', \mathcal{O}_K^2)$ *of a rank* 2 *free module* $M' \subset \mathcal{O}_K^2$*, together with some auxiliary information* ***aux****. If* $(\mathbf{B}, \mathbb{I})$ *is a* γ_{uSVP}-uSVP *instance with*

$$\gamma_{\text{uSVP}} = \gamma_{\text{NTRU}} \cdot \sqrt{\gamma_{\text{HSVP}}} \cdot 16\sqrt{2} \cdot d^{3/2} \cdot \delta_K,$$

  then $(\mathbf{B}', \mathcal{O}_K^2)$ *is a* $(\gamma_{\text{NTRU}}, q)$-NTRU *instance. If given access to a* γ_{HSVP}-id-HSVP *oracle, it runs in time polynomial in its input bitsize, in* $\zeta_K(2)$ *and in* $\exp(\frac{d\log(d)}{\log(2q/q_0)})$ *and makes one call to the oracle.*
- *Algorithm* `LiftVec'` *takes as input a non-zero vector* $\mathbf{v}' \in M'$ *and the auxiliary information* ***aux****. It outputs a non-zero vector* $\mathbf{s} \in M$ *such that*

$$\|\mathbf{s}\| \leq 4 \cdot \gamma_{\text{HSVP}}^{3/2} \cdot d^{9/2} \cdot \delta_K^2 \cdot \frac{\|\mathbf{s}'\|}{\mathcal{N}(M')^{1/(2d)}} \cdot \mathcal{N}(M)^{1/(2d)}.$$

  If given access to a γ_{HSVP}-id-HSVP *oracle, it runs in polynomial time and makes one call to the oracle.*
- *Algorithm* `LiftMod'` *takes as input a pseudo-basis of a rank-*1 *densest submodule* N' *of* M' *and the auxiliary information* ***aux*** *and outputs a pseudo-basis of a rank-*1 *densest submodule* N *of* M*. It runs in polynomial time.*

Proof. Let $V_0 = \text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+)$ be as in Lemma 4.3 (defined using γ^+ instead of $\gamma(M)$). Define

$$q_0 = \frac{V_0^{1/d} \cdot 4d}{\gamma_{\text{HSVP}}}.$$

One can check that q_0 is indeed $\text{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma^+)$ as desired. We prove that the theorem holds for this choice of q_0.

Algorithm `uSVP-to-NTRU`*.* On input $(\mathbf{B}, \mathbb{I})$, q and γ_{HSVP}, `uSVP-to-NTRU` sets $V = \gamma_{\text{HSVP}}^d \cdot q^d \cdot d^d$ and $\beta = \lceil \frac{2d\log(2d)}{\log\left(\sqrt{q/q_0}\right)+\log(2d)} \rceil$. It then runs Algorithm `PreCond` on input $(\mathbf{B}, \mathbb{I})$, V and β, to obtain a matrix $\mathbf{Y} \in \text{GL}_2(K)$.

From the definition of q_0, V and β, one can check that $V^{1/(2d)} \geq (2d)^{2d/\beta} \cdot V_0^{1/(2d)}$. Moreover, we have $\gamma(M) \leq \gamma^+$ by assumption, hence we can apply Lemma 4.3. This implies in particular that the call to the `PreCond` algorithm runs in time polynomial in the input bitsize, in $2^\beta = 2^{O(d\log(d)/\log(2q/q_0))}$ and in $\zeta_K(2)$.

Algorithm `uSVP-to-NTRU` then runs the Algorithm `Conditioned-to-NTRU` on input $(\mathbf{YB}, \mathbb{I})$, q and γ_{HSVP}. It obtains a basis $\mathbf{B}'$ of a free module M' and some auxiliary information `aux`'. Algorithm `uSVP-to-NTRU` finally outputs $(\mathbf{B}', \mathcal{O}_K^2)$ and $\texttt{aux} = (\texttt{aux}', \mathbf{Y}, \gamma_{\text{HSVP}}, \mathbf{B}')$.

We know from Lemma 4.5 that the call to `Conditioned-to-NTRU` can be done in polynomial time, with one call to the γ_{HSVP}-id-HSVP oracle. This concludes the proof on the run time of Algorithm `uSVP-to-NTRU`.

Let us assume now that $(\mathbf{B}, \mathbb{I})$ was a γ_{uSVP}-uSVP instance, for γ_{uSVP} as in the theorem. We know from Lemma 4.3 that $(\mathbf{YB}, \mathbb{I})$ is a $\gamma_{\text{uSVP}}/(2\sqrt{2})$-uSVP instance. Moreover, still from Lemma 4.3, we know that the module spanned by $(\mathbf{YB}, \mathbb{I})$ is a rank-2 module in $\mathcal{O}_K^2$, with the coprime property and such that $\mathcal{N}(M') \in [1/2^d, 2^d] \cdot V$. Hence we can apply Lemma 4.6 and conclude that $(\mathbf{B}', \mathcal{O}_K^2)$ is a γ_{NTRU} instance as desired (note that V and $\gamma_{\mathrm{uSVP}}/(2\sqrt{2})$ have the desired shape for applying Lemma 4.6). This proves the first point of our theorem.

Algorithm `LiftVec'`. On input $\mathbf{s}' \in M'$ and `aux` $= ($`aux'`$, \mathbf{Y}, \gamma_{\mathrm{HSVP}}, \mathbf{B}')$, Algorithm `LiftVec'` runs `LiftVec`$($`aux'`$, \gamma_{\mathrm{HSVP}}, \mathbf{B}', \mathbf{s}')$ and gets a vector $\mathbf{t}$. It then outputs $\mathbf{Y}^{-1} \cdot t$. By Lemma 4.8, we know that the call to `LiftVec` can be performed in polynomial time, with one call to the id-HSVP oracle. This proves the run time of `LiftVec'`.

By Lemma 4.7 again, we know that $\|\mathbf{t}\| \leq \|\mathbf{s}'\| \cdot \gamma_{\mathrm{HSVP}}^2 \cdot d^4 \cdot \delta_K^2$. From the shape of $\mathbf{Y}$ and Lemma 3.6 instantiated with $\varepsilon = 1/5$, we obtain

$$\|\mathbf{Y}^{-1} \cdot \mathbf{t}\| \leq \frac{4\mathcal{N}(M)^{1/(2d)}}{V^{1/(2d)}} \cdot \|t\| \leq 4 \cdot \gamma_{\mathrm{HSVP}}^{3/2} \cdot d^{9/2} \cdot \delta_K^2 \cdot \frac{\|\mathbf{s}'\|}{\sqrt{q}} \cdot \mathcal{N}(M)^{1/(2d)}.$$

Using the fact that $\mathcal{N}(M') = q^{2d}$ provides the desired upper bound on the output size. Note also that by construction, $\mathbf{Y}^{-1} \cdot \mathbf{t}$ is indeed a non-zero vector in M.

Algorithm `LiftMod'`. Let us call $\tilde{M}$ the intermediate module $(\mathbf{Y} \cdot \mathbf{B}) \cdot \mathbb{I}$ computed by Algorithm `uSVP-to-NTRU`.

On input a pseudo-basis $(\mathbf{v}', J')$ of a densest rank-1 module of M' and `aux` $=$ $($`aux'`$, \mathbf{Y}, \gamma_{\mathrm{HSVP}}, \mathbf{B}')$, Algorithm `LiftMod'` runs `LiftMod`$($`aux'`$, \mathbf{B}', (\mathbf{v}', J'))$ and gets a vector $\mathbf{w}$. It then computes J such that $\mathrm{span}(\mathbf{w}) \cap \tilde{M} = \mathbf{w} \cdot J$, sets $\mathbf{v} = \mathbf{Y}^{-1} \cdot \mathbf{w}$ and outputs the pseudo basis $(\mathbf{v}, J)$.

From Lemma 4.7, we know that algorithm `LiftMod'` runs in polynomial time. Moreover, since $(\mathbf{v}', J')$ was a densest submodule of M', we know that $\mathbf{w} \cdot J$ is a densest submodule of the module $\tilde{M}$. Recall that we proved that $\tilde{M}$ is a $\gamma_{\mathrm{uSVP}}/(2\sqrt{2})$-uSVP instance, hence we have $\mathcal{N}(\mathbf{w} \cdot J)^{1/d} = \lambda_1^{\mathcal{N}}(\tilde{M}) \leq 2\sqrt{2}/\gamma_{\mathrm{uSVP}} \cdot \mathcal{N}(\tilde{M})^{1/(2d)}$. From the special shape of $\mathbf{Y}$, one can prove that $\mathcal{N}(\mathbf{Y}^{-1} \cdot \mathbf{w}) \leq 4^d \cdot R^{-d} \cdot \mathcal{N}(\mathbf{w})$, with $R = V^{1/(2d)} \cdot \mathcal{N}(M)^{-1/(2d)}$. Hence, we obtain

$$\mathcal{N}(\mathbf{v} \cdot J)^{1/d} \leq 4 \cdot \frac{2\sqrt{2}}{R \cdot \gamma_{\mathrm{uSVP}}} \cdot \mathcal{N}(\tilde{M})^{1/(2d)} \leq \frac{16}{\gamma_{\mathrm{uSVP}}} \cdot \mathcal{N}(M)^{1/(2d)},$$

where we used the definition of R and the fact that $\mathcal{N}(\tilde{M}) \leq 2^d \cdot V$ by Lemma 4.3. Since $\gamma_{\mathrm{uSVP}} > 16$, we conclude that $\mathbf{v} \cdot J$ is a rank-1 submodule of M with $\mathcal{N}(\mathbf{v} \cdot J) < \mathcal{N}(M)^{1/2}$ and so from Lemma 2.8, we conclude that $\mathbf{v} \cdot J$ is indeed the densest submodule of M.

D.2 Proof of Lemma 4.3

The algorithm `PreCond` is as follows.

Algorithm D.1 Algorithm `PreCond`

Input: A pseudo-basis $(\mathbf{B}, \mathbb{I})$ of a rank-2 module $M \subseteq K^2$, a parameter $V > 0$ and a block-size $\beta \in [2, 2d]$

Output: A matrix $\mathbf{Y} \in \mathrm{GL}_2(K)$

```
Set ς = V^{1/(2d)} · (5√(2d))^{-1} · N(M)^{-1/(2d)}
repeat
    Sample Y := (y_1, y_2)^T ← DualRound((B, I), ς, β, 1/5)
until y_1 · O_K is a primitive submodule of M^∨
Return Y
```

Proof. Let P be the polynomial from Lemma 2.7 and define

$$V_0^{\frac{1}{2d}} = 10\sqrt{2}\cdot d\cdot\gamma(M)\cdot\Big(P\big(\Delta_K^{1/d}, 2, d, 5\sqrt{2d}, 4d^{3/2}\cdot\gamma(M)^2\cdot\delta_K\cdot\Delta_K^{1/(2d)}\big)+(2d)^{3/2}\Big).$$

We will prove that the lemma holds for this choice of V_0. Note that $V_0^{1/(2d)}$ is indeed $\mathrm{poly}(\Delta_K^{1/d}, d, \delta_K, \gamma(M))$ as desired.

Let us first observe that, by using the lower bound on V, the definition of ς and V_0 and Lemma 2.11, one can prove that the lower bound $\varsigma \geq (2d)^{2d/\beta+3/2}\cdot\lambda_{2d}(\sigma(M^\vee))$ required in Lemma 3.5 is satisfied.

Applying Lemma 3.5, we know that the calls to Algorithm `DualRound` will take a time polynomial in the input bitsize and in 2^β. To estimate the number of such calls, let us use Lemma 2.7. The definition of V_0 ensures that $\varsigma \geq \lambda_{2d}(M^\vee)\cdot P(\Delta_K^{1/d}, 2, d, \|\mathbf{t}_1\|/\varsigma, \lambda_{kd}(M^\vee)/\lambda_1(M^\vee))$ as required by Lemma 2.7. Hence, we know that $\Pr_{\mathbf{y}\leftarrow D_{M^\vee,\varsigma,\mathbf{t}_1}}(\mathbf{y}\cdot\mathcal{O}_K \text{ is primitive in } M^\vee) \geq 1/(4\zeta_K(2))$. Using the fact that $\mathrm{SD}(D_{M^\vee,\varsigma,\mathbf{t}_1}, \hat{D}_{M^\vee,\varsigma,\mathbf{t}_1}) \leq 2^{-\Omega(d)}$ we conclude that the probability to exit the while loop is at least $1/(4\zeta_K(2)) - 2^{-\Omega(d)} \geq \zeta_K(2)^{-O(1)}$ at every iteration of the algorithm. This proves the expected run time of the algorithm.

The fact that $\mathbf{Y} = R\cdot\mathbf{I}_2 + \mathbf{E}$ with $\|e_{ij}\| \leq R/5$ and that $M' := \mathbf{YB}\cdot\mathbb{I}$ is included in $\mathcal{O}_K^2$ follows from Lemma 3.5 (instantiated with $\varepsilon = 1/5$). Since $\varepsilon = 1/5 \leq 1/4$, we can also use Lemma 3.6. This implies in particular that $\det(\mathbf{Y}) \in [1/2^d, 2^d]\cdot R^{2d}$, where $R = 5\cdot\sqrt{2d}\cdot\varsigma = V^{1/(2d)}\cdot\mathcal{N}(M)^{-1/(2d)}$ by definition of ς. This proves that $\mathbf{Y}$ is invertible, and so M' is indeed a rank-2 module. This also proves that $\mathcal{N}(M') = \det(\mathbf{Y})\cdot\mathcal{N}(M) \in [1/2^d, 2^d]\cdot V$.

Let us now show that if $(\mathbf{B}, \mathbb{I})$ was a γ_{uSVP}-uSVP instance, then $(\mathbf{YB}, \mathbb{I})$ is a γ'_{uSVP}-uSVP instance. Let $\mathbf{s} \in M$ be a short vector such that $\|\mathbf{s}\| \leq 1/\gamma_{\mathrm{uSVP}}\cdot\mathcal{N}(M)^{1/(2d)}$ (such a short vector exists if $(\mathbf{B}, \mathbb{I})$ is a γ_{uSVP}-uSVP instance). Define $\mathbf{s}' = \mathbf{Y}\cdot\mathbf{s}$, which is a vector of M'. We have

$$\|\mathbf{s}'\| \leq R\cdot\|\mathbf{s}\| + \|\mathbf{E}\cdot\mathbf{s}\| \leq 2R\cdot\|s\| \leq 2R\cdot\gamma_{\mathrm{uSVP}}^{-1}\cdot\mathcal{N}(M)^{1/(2d)}.$$

Recall that $\mathcal{N}(M') = \det(\mathbf{Y})\cdot\mathcal{N}(M) \geq 1/2^d\cdot R^{2d}\cdot\mathcal{N}(M)$. This finally implies that $\|\mathbf{s}'\| \leq 2\sqrt{2}\cdot\gamma_{\mathrm{uSVP}}^{-1}\cdot\mathcal{N}(M')^{1/(2d)}$, and so $(\mathbf{YB}, \mathbb{I})$ is indeed a γ'_{uSVP}-uSVP instance.

It finally remains to show that the module M' has the coprime property. This is implied by the fact that $\mathbf{y}_1\cdot\mathcal{O}_K$ is primitive in $M^\vee$. Indeed, by definition of M', we have that $\{x \in \mathcal{O}_K \,|\, \exists y \in \mathcal{O}_K \text{ s.t. } (x, y)^T \in M'\} = \{\langle\mathbf{y}_1, \mathbf{z}\rangle_{K_\mathbb{R}} \,|\, \mathbf{z} \in M\}$.

One can see from the definition that this set is an ideal of $\mathcal{O}_K$. Assume by contradiction that it is not equal to $\mathcal{O}_K$ and let $\mathfrak{p}$ be a prime ideal dividing it. Then it holds that $\mathbf{y}_1 \cdot \mathfrak{p}^{-1} \subset M^\vee$. But this is a rank-1 submodule of $M^\vee$ containing strictly the rank-1 module $\mathbf{y}_1 \cdot \mathcal{O}_K$, contradicting the assumption that $\mathbf{y}_1 \cdot \mathcal{O}_K$ is primitive in $M^\vee$. Hence, we conclude that M' has the coprime property. ☐

D.3 Proof of Lemma 4.3

The algorithm `BalanceIdeal` is as follows.

Algorithm D.2 Algorithm `BalanceIdeal`

Input: A $\mathbb{Z}$-basis of a fractional ideal $I \subset K$ and a parameter $\gamma_{\mathrm{HSVP}} \geq 1$
Output: An element $x \in K$

Using a γ_{HSVP}-id-HSVP oracle to get short linearly independent vectors of I^{-1}
Call a γ_{HSVP}-id-HSVP solver on I^{-1} to get $y \in I^{-1}$
Let $\mathbf{B} = (yr_1, \cdots, yr_d)$ (this is a $\mathbb{Z}$-basis of $\langle y \rangle$)
Using the short vectors to find a balanced element in I^{-1} by solving CVP
Let $\sigma = \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(I)^{-1/d}$ and $t = (\sigma, \cdots, \sigma)$
Write $t = \sum_i t_i \cdot yr_i$, with $t_i \in \mathbb{R}$
Define $s = \sum_i \lfloor t_i \rceil \cdot yr_i$
Return $x = s^{-1}$

Proof. One can check that all the steps of the algorithm, except for the one call to the γ_{HSVP}-HSVP oracle, can be performed in polynomial time.

Let us then prove correction, and start with $I \subseteq \langle x \rangle$. We know that $s \in \langle y \rangle$, by definition of s. Since $y \in I^{-1}$, it holds that $\langle s \rangle \subseteq I^{-1}$, which implies $I \subseteq \langle s \rangle^{-1} = \langle x \rangle$ as desired (provided that $s \neq 0$, which we will show below).

Let us know look at how balanced are the coordinates of s (and x). We have

$$\begin{aligned}
\|s - t\|_\infty &\leq \sum_i 1/2 \cdot \|yr_i\|_\infty \\
&\leq 1/2 \cdot \sum_i \|y\|_\infty \cdot \|r_i\|_\infty \\
&\leq d/2 \cdot \|y\| \cdot \delta_K \\
&\leq 1/(2d) \cdot \sigma,
\end{aligned}$$

where we used in the last inequality the fact that y is the output of the γ_{HSVP}-id-HSVP solver on I^{-1}, and hence $\|y\| \leq \gamma_{\mathrm{HSVP}} \cdot \mathcal{N}(I)^{-1/d}$. Since $\sigma_i(s)$ is the i-th coordinate of s and all the coordinates of t are equal to σ, this implies that $|\sigma_i(s)| \in [\sigma \cdot (1 - 1/(2d)), \sigma \cdot (1 + 1/(2d))]$ (and in particular $\sigma_i(s) \neq 0$ for all i's, so s is invertible). Using the facts that $\sigma_i(x) = \sigma_i(s)^{-1}$ and the convexity of the function $x \mapsto 1/x$ over $[1/2, 2]$ conclude the proof. ☐

D.4 Proof of Lemma 4.5

We know from preliminaries (cf Section 2.3) that the HNF basis of a module can be computed in polynomial time. From Lemma 4.4 we know that the algorithm `BalanceIdeal` runs in polynomial time and make one call to the γ_{HSVP} oracle. Note that the input ideal J_2 is indeed fractional (and even integral) since $M \subset \mathcal{O}_K^2$ and that $\gamma_{\mathrm{HSVP}} \geq \sqrt{d}\Delta_K^{1/(2d)}$ hence we can indeed run algorithm `BalanceIdeal`. Finally, the multiplications and rounding in the third step of the algorithm can be performed in polynomial time too. □

D.5 Proof of Lemma 4.6

Let us fix some δ, $\gamma_{\mathrm{HSVP}}, \gamma_{\mathrm{NTRU}}$ and q as in the theorem and define V and γ_{uSVP} accordingly.

Let $(\mathbf{B}, \mathbb{I})$ be the input pseudo-basis, spanning a rank-2 module $M_1 \subset \mathcal{O}_K^2$ with $\mathcal{N}(M_1) \in [1/2^{2d}, 2^{2d}] \cdot V$, with the coprime property, and which we know contains a non-zero vector $\mathbf{s}_1 = (u, v)^T \in M_1$ such that $\|\mathbf{s}_1\| \leq 1/\gamma_{\mathrm{uSVP}} \cdot \mathcal{N}(M_1)^{1/(2d)}$. We will see step by step how the module M_1 is modified by the algorithm, and what happens to its short non-zero vectors. This is summarized on Figure 2.

First step: HNF. After the HNF computation, we have a new pseudo-basis of the form

$$\left[\begin{array}{cc} J_1 & J_2 \\ \multicolumn{2}{c}{\begin{pmatrix} 1 & 0 \\ a & 1 \end{pmatrix}} \end{array} \right]$$

for some $a \in K$ and $J_1, J_2 \subset K$ (cf Section 2.3). This pseudo-basis generates a rank-2 module M_2 which is the same as the input module M_1. Hence, M_2 contains a short non-zero vector $\mathbf{s}_2 := \mathbf{s}_1$.

Since our module M_2 is integral, we know that both ideals J_1 and J_2 are integral. Also, since module $M_2 = M_1$ has the coprime property, we know that $J_1 = \mathcal{O}_K$. Finally, because of the shape of the pseudo-basis, it holds that $\mathcal{N}(J_1) \cdot \mathcal{N}(J_2) = \mathcal{N}(M_2) = \mathcal{N}(M_1)$, which yields $\mathcal{N}(J_2) \geq \mathcal{N}(M_1)$.

Second step: from pseudo-basis to basis. Let M_3 be the free module generated by the pseudo-basis

$$\left[\begin{array}{cc} \mathcal{O}_K & \mathcal{O}_K \\ \multicolumn{2}{c}{\begin{pmatrix} 1 & 0 \\ a & b \end{pmatrix}} \end{array} \right],$$

where $b \leftarrow$ `BalanceIdeal`$(J_2, \gamma_{\mathrm{HSVP}})$. Since $J_2 \subseteq \langle b \rangle$ by Lemma 4.4 and $J_1 = \mathcal{O}_K$, we conclude that $M_2 \subseteq M_3$. Hence, the short vector $\mathbf{s}_3 := \mathbf{s}_2$ is still in M_3.

Before moving to the next step, let us have a closer look at b. We know from Lemma 4.4 that $|\sigma_i(b)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$ for all $i \leq d$, with $\sigma = \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_2)^{-1/d}$. Using the lower bound on $\mathcal{N}(J_2)$ that we computed above, this shows that $\sigma \leq \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(M_1)^{-1/d}$.

Third step: transforming b into q. Let M_4 be the free module generated by the pseudo-basis

$$\left[\begin{array}{cc} \mathcal{O}_K & \mathcal{O}_K \\ \begin{pmatrix} 1 & 0 \\ h & q \end{pmatrix} \end{array} \right],$$

where $h = \lfloor a \cdot q/b \rceil$. This is the basis output by Algorithm `Conditioned-to-NTRU`. The new module M_4 does not contain M_1 anymore, however we will show that its geometry is close to the one of M_3, so that it has a short non-zero vector if M_3 does.

Recall that M_3 contains a short vector $\mathbf{s}_3 = (u, v)^T$ such that $\|\mathbf{s}_3\| \leq 1/\gamma_{\mathrm{uSVP}} \cdot \mathcal{N}(M_1)^{1/(2d)}$. Let $x \in \mathcal{O}_K$ be such that $\mathbf{s}_3 = u \cdot (1, a)^T + x \cdot (0, b)^T$. Define $\mathbf{s}_4 = u \cdot (1, h)^T + x \cdot (0, q)^T \in M_4 \setminus \{\mathbf{0}\}$. Unrolling the definition of h and using the equation $u \cdot a + x \cdot b = v$, one can rewrite $\mathbf{s}_4 = (u, v \cdot q/b - u \cdot \{a \cdot q/b\})^T$. We can upper bound the euclidean norm of $\mathbf{s}_4$ as follows

$$\begin{aligned}
\|\mathbf{s}_4\| &\leq \|u\| + \|v\| \cdot \|q/b\|_\infty + \|u\| \cdot \|\{a \cdot q/b\}\|_\infty \\
&\leq \gamma_{\mathrm{uSVP}}^{-1} \cdot \mathcal{N}(M_1)^{1/(2d)} \cdot \big(1 + q \cdot 2\sigma + d\delta_K\big) \\
&\leq \gamma_{\mathrm{uSVP}}^{-1} \cdot \mathcal{N}(M_1)^{1/(2d)} \cdot \big(q \cdot 2\sigma + 2 \cdot d\delta_K\big) \\
&\leq 2\gamma_{\mathrm{uSVP}}^{-1} \cdot d \cdot \delta_K \cdot \big(\mathcal{N}(M_1)^{1/(2d)} + q \cdot d \cdot \gamma_{\mathrm{HSVP}} \cdot \mathcal{N}(M_1)^{-1/(2d)}\big) \\
&\leq 1/\gamma_{\mathrm{NTRU}} \cdot \sqrt{q},
\end{aligned}$$

where in the last step we used the fact that $\mathcal{N}(M_1)^{1/(2d)} \in [1/2, 2] \cdot V^{1/(2d)}$ and the definitions of V and γ_{uSVP}. We conclude that the pseudo-basis output by Algorithm `Conditioned-to-NTRU` is indeed a γ_{NTRU}-NTRU instance, as desired. □

D.6 Proof of Lemma 4.7

Algorithm `LiftMod` is as follows.

Algorithm D.3 Algorithm `LiftMod`

Input: Two elements $a, b \in K$, an NTRU instance $((\mathbf{c}_1, \mathbf{c}_2), \mathcal{O}_K^2)$ and a pseudo-basis $(\mathbf{v}, J)$ of a rank-1 module in K^2.

Output: A vector $\mathbf{w} \in K^2$

1: Compute $x, y \in K$ such that $\mathbf{v} = x \cdot \mathbf{c}_1 + y \cdot \mathbf{c}_2$
2: Define $\mathbf{w} = x \cdot (1, a)^T + y \cdot (0, b)^T$
3: **Return w**

Proof. The run time follows from inspection of the algorithm.

Let $\mathbf{s}_1$ be a shortest vector of M_1. Since $\gamma_{\mathrm{uSVP}} > 1$, we know from Lemma 2.8 that $\mathbf{s}_1$ belongs to the densest rank-1 submodule of M_1, i.e., the densest submodule of M_1 is equal to $\mathrm{span}_K(\mathbf{s}_1) \cap M_1$.

Let us use the notations M_1, M_2, M_3 and M_4 as in Figure 2 (and in the proof of Lemma 4.6). Recall from the proof of Lemma 4.6 that $\mathbf{s}_1$ is still a vector of the rank-2 module M_3 spanned by $(1,a)^T, (0,b)^T$. Let $u, r \in \mathcal{O}_K$ be such that $\mathbf{s}_1 = u \cdot (1,a)^T + r \cdot (0,b)^T$. Recall again from the proof of Lemma 4.6 that $\mathbf{s}_4 = u \cdot \mathbf{c}_1 + r \cdot \mathbf{c}_2$ is an unexpectedly short vector of the output NTRU module M_4. More precisely, we proved that $\|\mathbf{s}_4\| \leq 1/\gamma_{\mathrm{NTRU}} \cdot \mathcal{N}(M_4)^{1/(2d)}$.

Using Lemma 2.8 again and the fact that $\gamma_{\mathrm{NTRU}} > 1$, we know that $\mathbf{s}_4$ belongs to the densest submodule of M_4. Since $(\mathbf{v}, J)$ is a pseudo-basis of this densest submodule, it should be that $\mathbf{v}$ and $\mathbf{s}_4$ are K-colinear, i.e., there exists $z \in K$ such that $\mathbf{v} = z \cdot \mathbf{s}_4 = zu \cdot \mathbf{c}_1 + zr \cdot \mathbf{c}_2$.

Hence, the elements x, y computed in the algorithms are equal to zu and zr respectively. This proves that $\mathbf{w} = x \cdot (1,a)^T + y \cdot (0,b)^T = z \cdot \mathbf{s}_1$. Hence, $\mathrm{span}_K(\mathbf{w}) = \mathrm{span}_K(\mathbf{s}_1)$ and the densest submodule of M is $\mathrm{span}_K(\mathbf{w}) \cap M_1$.

□

D.7 Proof of Lemma 4.8

Algorithm `LiftVec` is as follows.

Algorithm D.4 Algorithm `LiftVec`

Input: Some auxiliary information `aux` $= (a, b, J_1, J_2)$, a parameter γ_{HSVP}, an NTRU instance $((\mathbf{c}_1, \mathbf{c}_2), \mathcal{O}_K^2)$ and a vector $\mathbf{s} \in \mathbf{C} \cdot \mathcal{O}_K^2$

Output: A vector $\mathbf{w} \in K^2$

1: Compute $x, y \in \mathcal{O}_K$ such that $\mathbf{s} = x \cdot \mathbf{c}_1 + y \cdot \mathbf{c}_2$

2: Run $z \leftarrow$ `BalanceIdeal`$(\langle b \rangle \cdot J_1^{-1} \cdot J_2^{-1}, \gamma_{\mathrm{HSVP}})$

3: Compute $\mathbf{t} = z^{-1} \cdot \big(x \cdot (1,a)^T + y \cdot (0,b)^T\big)$

4: **Return t**

Proof. The running-time follows from inspection of the algorithm and from Lemma 4.4.

Let us show that the output $\mathbf{t}$ of the algorithm is indeed in the module M_1. Let us keep the notations M_1, M_2, M_3 and M_4 from Figure 2. In particular, $M_1 = M_2$ is the module generated by the pseudo-basis $((1,a)^T, (0,b)^T), (J_1, \langle b^{-1} \rangle \cdot J_2)$.

From Lemma 4.4, we know that $\langle b \rangle \cdot J_1^{-1} \cdot J_2^{-1} \subseteq \langle z \rangle$, i.e., $z^{-1} \in J_1 \cdot J_2 \cdot \langle b^{-1} \rangle$. Using the fact that J_1 and $J_2 \cdot \langle b^{-1} \rangle$ are both integral (recall that $J_2 \subseteq \langle b \rangle$ and that M_1 is in $\mathcal{O}_K^2$), this implies that $z^{-1} \in J_1 \cap J_2 \cdot \langle b^{-1} \rangle$. Since $x, y \in \mathcal{O}_K$, we conclude that $\mathbf{t} = z^{-1} \cdot x \cdot (1,a)^T + z^{-1} \cdot y \cdot (0,b)^T$ is in M_1 as desired.

Let us now upper bound the size of $\mathbf{t}$. Let us write $\mathbf{s} = (s_1, s_2)^T$ and express the coordinates of $\mathbf{t}$ in terms of s_1 and s_2. From the equation $\mathbf{s} = x \cdot (1, \lfloor a \cdot q/b \rceil)^T + y \cdot (0, q)^T$, we obtain $x = s_1$ and $s_2 = x \lfloor a \cdot q/b \rceil + yq$. This implies that $\mathbf{t} = z^{-1} \cdot \big(s_1, b/q \cdot (s_2 + s_1 \cdot \{a \cdot q/b\})\big)^T$. From this, we can upper bound

$$\|\mathbf{t}\| \leq \|z^{-1}\|_\infty \cdot \|\mathbf{s}\| \cdot \big(1 + \|b/q\|_\infty \cdot (1 + d\delta_K)\big).$$

Recall from the proof of Lemma 4.6 that for all $i \leq d$, we have $|\sigma_i(b)| \in [1 - 1/d, 1 + 1/d] \cdot \sigma^{-1}$, with $\sigma = \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_2)^{-1/d}$. Hence,

$$\|b/q\|_\infty \leq \frac{2 \cdot \mathcal{N}(J_2)^{1/d}}{\gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot q} \quad \text{and} \quad |\mathcal{N}(b)| \geq \frac{\mathcal{N}(J_2)}{(2\gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K)^d}.$$

From Lemma 4.4, we similarly know that $|\sigma_i(z)| \in [1-1/d, 1+1/d] \cdot \sigma_z^{-1}$ for all $i \leq d$, where $\sigma_z = \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_1 \cdot J_2 \cdot \langle b^{-1} \rangle)^{1/d}$. Hence we obtain

$$\|z^{-1}\|_\infty \leq 2 \cdot \gamma_{\mathrm{HSVP}} \cdot d^2 \cdot \delta_K \cdot \mathcal{N}(J_1 \cdot J_2 \cdot \langle b^{-1} \rangle)^{1/d} \leq 4 \cdot \gamma_{\mathrm{HSVP}}^2 \cdot d^4 \cdot \delta_K^2,$$

where we used the fact that $J_1 = \mathcal{O}_K$ thanks to the coprime property of M_1.

Finally, recall that $\mathcal{N}(J_2) = \mathcal{N}(M_1) \leq 2^{2d} \cdot V$, with $V^{1/d} = q \cdot \gamma_{\mathrm{HSVP}} \cdot d$. Combining everything provides the desired upper bound on $\|\mathbf{t}\|$. □

E Missing proofs from Section 5

E.1 Proof of Theorem 5.2

Note that the assumptions on B and γ in the theorem statement enable the use of all theorems and lemmas from Sections 5.1, 5.2 and 5.3. The runtime statement follows from the runtime statements in Theorems 5.6 and 5.9. By using Theorems 5.6 and 5.9, we also obtain that the pseudo-basis output by `Randomize` spans a rank-2 and norm-1 module.

Let M' be the module spanned by the output $(\mathbf{B}', \mathbb{I}')$ of `Randomize`, when given as input a module with gap γ. By Theorems 5.6 and 5.9, the distribution of M'' (over the internal randomness of `Randomize`) is within statistical distance $2^{-\Omega(d)}$ from $\texttt{QRSF-2-Mod}(D_{B,\gamma}^{\mathrm{rand}})$, where $D_{B,\gamma}^{\mathrm{rand}}$ is as defined in Definition 5.10. Now, we apply `QRSF-2-Mod` to all the distributions of Definition 5.10. By the probability preservation properties of the statistical distance and Rényi divergence, and by Lemmas E.1, E.2, E.3, E.4 and E.6, any event that occurs with probability $\varepsilon \geq 2^{-\Omega(d)}$ for $\texttt{QRSF-2-Mod}(D_{B,\gamma}^{\mathrm{target}})$ also holds with probability $\Omega(\varepsilon^4)$ for $\texttt{QRSF-2-Mod}(D_{B,\gamma}^{\mathrm{rand}})$. By observing that $\texttt{QRSF-2-Mod}(D_{B,\gamma}^{\mathrm{rand}})$ is exactly $D_{B,\gamma}^{\mathrm{module}}$, we obtain that any event that holds for $D_{B,\gamma}^{\mathrm{module}}$ with probability $\varepsilon \geq 2^{-o(d)}$ also holds for M' with probability $\Omega(\varepsilon^4)$ over the internal randomness of `Randomize`.

We now analyze `Recover`. Let M be the module spanned by $(\mathbf{B}, \mathbb{I})$. Let U be its densest rank-1 submodule. Let $((\mathbf{B}', \mathbb{I}'), \texttt{aux})$ be an output of `Randomize` when given $(\mathbf{B}, \mathbb{I})$ as input, and U' be a densest rank-1 submodule of M'. By Theorems 5.6 and 5.9, we have that with probability $1 - 2^{-\Omega(d)}$, the module M' has gap larger than 1 and its densest rank-1 submodule is

$$U' = (\mathcal{N}(\mathfrak{p}) \cdot \det(\mathbf{D}))^{-\frac{1}{2d}} \cdot \mathbf{D} \cdot U \cdot q\mathfrak{p}.$$

This completes the proof. □

E.2 Proof of Theorem 5.3

Assume that $\mathbf{u} \cdot J_1 \not\subset M'$, which holds with probability $1 - (1/B)^{\Omega(1)}$ by Lemma 5.4. We fix x as in Lemma 5.5. Let $M'' = \mathbf{u} \cdot \mathfrak{p}J_1 + (\mathbf{v} + x\mathbf{u}) \cdot J_2$. By Lemmas 5.4 and 5.5, we have that $M'' \subseteq M'$. By construction, the norm of M'' is $\mathcal{N}(\mathfrak{p}) \cdot \mathcal{N}(M)$, which is equal to $\mathcal{N}(M')$ by Lemma 3.2, leading to the equality $M' = M''$. This completes the proof of the first statement.

Assume that we have $M' = \mathbf{u} \cdot \mathfrak{p} J_1 + (\mathbf{v} + x\mathbf{u}) \cdot J_2$ and $\gamma(M) \geq B^{1/(2d)}$, and that $\mathbf{u} \cdot J_1$ is the densest rank-1 submodule of M. As $\mathbf{u} \cdot \mathfrak{p} J_1$ is a rank-1 submodule of M', we have:

$$\gamma(M') \geq \left(\frac{\sqrt{\mathcal{N}(M')}}{\mathcal{N}(\mathbf{u} \cdot \mathfrak{p} J_1)} \right)^{\frac{1}{d}} = \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \left(\frac{\sqrt{\mathcal{N}(M)}}{\mathcal{N}(\mathbf{u} \cdot J_1)} \right)^{\frac{1}{d}} = \frac{\gamma(M)}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}}.$$

As $\mathcal{N}(\mathfrak{p}) \leq B$, we obtain that $\gamma(M') \geq \gamma(M)/B^{1/(2d)} > 1$. By Lemma 2.8, we know that M' has a unique densest rank-1 submodule. Now, using the equalities above and the inequalities $\gamma(M) \geq B^{1/(2d)}$ and $\mathcal{N}(\mathfrak{p}) \leq B$, we have

$$\mathcal{N}(\mathbf{u} \cdot \mathfrak{p} J_1) = \mathcal{N}(\mathfrak{p})^{\frac{1}{2}} \cdot \frac{\mathcal{N}(M')^{\frac{1}{2}}}{\gamma(M)^d} \leq \mathcal{N}(M')^{\frac{1}{2}}.$$

Lemma 2.8 then implies that $\mathbf{u} \cdot \mathfrak{p} J_1$ is contained in the densest rank-1 submodule of M'. By primitivity (see Definition 2.5), we conclude that it is the densest rank-1 submodule of M'. ◻

E.3 Proof of Lemma 5.4

As $\mathbf{u}\mathfrak{p} J_1$ is a primitive rank-1 submodule of M, we can use Lemma 3.3. It implies that the result holds, except with probability $1/\mathcal{N}(\mathfrak{p}) - 1/\mathcal{N}(\mathfrak{p})^2$ over the choice of $\overline{\mathbf{b}^\vee}$.

The overall probability (including over the choice of $\mathfrak{p}$) that $\mathbf{u} \cdot J_1 \subset M'$ holds satisfies:

$$\begin{aligned}
\sum_{\mathcal{N}(\mathfrak{p}) \leq B} \Pr(\mathfrak{p}) \cdot \Pr\big(\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \in \mathfrak{p} J_1^{-1} \mid \mathfrak{p}\big) &= \frac{1}{\pi_K(B)} \sum_{\mathcal{N}(\mathfrak{p}) \leq B} \left(\frac{1}{\mathcal{N}(\mathfrak{p})} - \frac{1}{\mathcal{N}(\mathfrak{p})^2} \right) \\
&\leq \frac{1}{\pi_K(B)} \sum_{p \leq B} \sum_{\mathfrak{p} \mid p} \frac{1}{\mathcal{N}(\mathfrak{p})} \\
&\leq \frac{d}{\pi_K(B)} \sum_{p \leq B} \frac{1}{p},
\end{aligned}$$

where the sums indexed by $\mathfrak{p}$ are over the prime ideals of $\mathcal{O}_K$ and the sums indexed by p are over the prime integers. The last inequality comes from the facts that there are at most d ideals $\mathfrak{p}$ over p, and each of them has norm $\geq p$. As $\sum_{p \leq B} 1/p = \log \log B + O(1)$ (see, e.g., [Apo98, Theorem 4.2]) and $\pi_K(B) = \Theta(B/\log B)$, we obtain that the probability above is $\leq (1/B)^{\Omega(1)}$. ◻

E.4 Proof of Lemma 5.5

Let $j \in J_1$ with $j\mathbf{u} \notin M'$. Since $\langle \mathbf{b}^\vee, j\mathbf{u} \rangle_{K_\mathbb{R}}$ belongs to $\mathcal{O}_K \setminus \mathfrak{p}$ (by definition of j), we can take a representative $a \in \mathcal{O}_K$ of its inverse in $\mathcal{O}_K/\mathfrak{p}$. We define $y = -\langle \mathbf{b}^\vee, \mathbf{v} \rangle_{K_\mathbb{R}} \cdot a \in J_2^{-1}$. By construction, we have $\langle \mathbf{b}^\vee, \mathbf{v} + jy\mathbf{u} \rangle_{K_\mathbb{R}} \in \mathfrak{p} J_2^{-1}$. This implies that $(\mathbf{v} + jy\mathbf{u}) \cdot J_2 \subset M'$. Setting $x = jy$ provides the result. ◻

E.5 Proof of Theorem 5.6

The run-time bound follows from Theorem 5.3 and Lemma 2.4. Now, we write

$$M = \frac{1}{\gamma} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 + \begin{bmatrix} r \\ 1 \end{bmatrix} \cdot \gamma^2 \cdot J_2 \right) = \mathbf{u} \cdot J_1 + \mathbf{v} \cdot J_2.$$

Let $\mathfrak{p}, \overline{\mathbf{b}^\vee}$ and $\mathfrak{q}$ refer to the random variables sampled during the execution of `Real-CR` and let $\mathbf{b}^\vee$ be a representative of $\overline{\mathbf{b}^\vee}$ in $M^\vee$. By Theorem 5.3, we have $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \notin \mathfrak{p} J_1^{-1}$ with probability $1 - (1/B)^{\Omega(1)}$. In the following, we assume that this holds. We also replace the distribution of $\mathfrak{q}$ by the uniform distribution over norm-1 ideals. By Lemma 2.4, these two distributions are within 2^{-d} statistical distance from one another. These two assumptions account for the statistical distance upper bound in the theorem statement.

Let $x \in J_1 {J_2}^{-1}$ as in Theorem 5.3. We have $\langle \mathbf{b}^\vee, \mathbf{v} + x\mathbf{u} \rangle_{K_\mathbb{R}} \in \mathfrak{p} J_2^{-1}$. For any choice of x such that the latter holds, the module M' corresponding to the output of `Real-CR` is, by Theorem 5.3:

$$M' = \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \cdot (\mathbf{u} \cdot J_1 \mathfrak{p}\mathfrak{q} + \mathbf{v}' \cdot J_2 \mathfrak{q}),$$

where $\mathbf{v}' = \mathbf{v} + x\mathbf{u}$. Note that the QR-factorization of the matrix $[\mathbf{u}|\mathbf{v}']$ is:

$$[\mathbf{u}|\mathbf{v}'] = \mathbf{Q} \cdot \begin{pmatrix} \frac{1}{\gamma} & \gamma \cdot (r + x) \\ 0 & \gamma \end{pmatrix}.$$

We define the norm-1 ideal $J = J_2 \mathfrak{q}$. We have:

$$\begin{aligned} M' &= \frac{1}{\gamma \cdot \mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 {J_2}^{-1} J \mathfrak{p} + \gamma^2 \cdot \begin{bmatrix} r + x \\ 1 \end{bmatrix} \cdot J \right) \\ &= \frac{1}{\sqrt{\gamma'}} \cdot \mathbf{Q} \cdot \left(\begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 {J_2}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})} + \gamma'^2 \cdot \begin{bmatrix} r + x \\ 1 \end{bmatrix} \cdot J \right), \end{aligned}$$

where $\gamma' = \gamma / \mathcal{N}(\mathfrak{p})^{1/(2d)}$. As the ideal $\mathfrak{q}$ is distributed uniformly over the set of norm-1 ideals, so is J. This implies that the distribution of M' matches that of `QRSF-2-Mod`(`Ideal-CR`$_B(\mathbf{Q}, \gamma, J_1, J_2, r)$).

Still assuming that we have $\langle \mathbf{b}^\vee, \mathbf{u} \rangle_{K_\mathbb{R}} \notin \mathfrak{p} J_1^{-1}$, Theorem 5.3 gives us that the densest rank-1 submodule of M' is:

$$\frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}} \cdot \frac{1}{\gamma} \mathbf{u} \cdot J_1 \mathfrak{p}\mathfrak{q} = \frac{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}}{\gamma} \cdot \mathbf{Q} \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1 \mathfrak{q} \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}.$$

This completes the proof of the theorem. □

E.6 Proof of Lemma 5.7

For the first statement, we prove that for $\mathbf{D}_0 \in \mathbb{R}^{2 \times 2}$ sampled from $\mathcal{D}(0,1)^{2 \times 2}$, we have $|\det \mathbf{D}_0| \geq 1/d$ with probability $1 - O((\log d)/d)$. As D_{distort} consists in $\leq d$ independent copies of the latter distribution, the probability of accepting a sample from $\mathcal{D}_{K_\mathbb{R}}(0,1)^{2\times 2}$ when rejecting to D_{distort} is at least $1/d^{O(1)}$.

Observe that $\mathbf{D}_0 = \|\mathbf{d}_1\| \cdot \|\mathbf{d}_2^*\|$, where $\mathbf{d}_1 \sim \mathcal{D}(0,1)^2$ is the first column of $\mathbf{D}_0$ and $\mathbf{d}_2^*$ is the projection of the second column orthogonally to the first. As $\mathbf{D}_0$ is invariant under rotations, conditioned on $\mathbf{d}_1$, the vector $\mathbf{d}_2^*$ is distributed as a sample from $\mathcal{D}(0,1)$ multiplied with a unit vector orthogonal to $\mathbf{d}_1$. For these reasons, it suffices to show that with probability $O((\log d)/d)$, the product of two iid samples x, y from $\mathcal{D}(0,1)$ has magnitude $\geq 1/d$. We have

$$\begin{aligned}\Pr_{x,y\leftarrow\mathcal{D}(0,1)}[|xy| < 1/d] &\leq O(1/d) \;+\; 4\cdot \Pr_{x,y\leftarrow\mathcal{D}(0,1)}[xy < 1/d \wedge x,y \in [1/d,1]] \\ &\leq O(1/d) \;+\; c\cdot \Pr_{x,y\leftarrow\mathcal{U}([1/d,1])}[xy < 1/d]\,,\end{aligned}$$

for some constant c. The latter is $O((\log d)/d)$, allowing to complete the proof of the first statement.

The second statement comes from the invariances of the determinant and vector Gaussian distribution under multiplication by an orthogonal matrix. $\square$

E.7 Proof of Lemma 5.8

We first show that without the conditioning, the matrix $\mathbf{D}$ from the lemma statement is distributed from $\mathcal{D}_{K_\mathbb{R}}(0,1)^{2\times 2}$. Let us write $\mathbf{D} = [\mathbf{d}_1|\mathbf{d}_2]$. Then $\mathbf{d}_1$ is the product of a uniform unit vector and an element sampled from $\chi_{K_\mathbb{R}}$. It is hence distributed as a Gaussian vector. Now, as the Gaussian vector distribution is invariant by multiplication by an orthogonal matrix, the distribution of $\mathbf{d}_2 = \mathbf{Q}\cdot(b,c)^T$ is $\mathcal{D}_{K_\mathbb{R}}(0,1)^2$, independently of $\mathbf{Q}$ and a.

To complete the proof, note that the conditioning is with respect to the event "$\forall i : |\det(\sigma_i(\mathbf{D}))| \geq 1/d$", for both D and D_{distort}. $\square$

E.8 Proof of Theorem 5.9

The runtime claim follows from Lemma 5.8. Now, let $\mathbf{D} \leftarrow D_{\text{distort}}$ be the matrix sampled in Step 1 of `Real-GR`. The matrix $\mathbf{D}\cdot\mathbf{Q}$ is also distributed from D_{distort}, by Lemma 5.7. By Lemma 5.8 we can write $\mathbf{DQ} = \mathbf{Q}'\cdot\begin{pmatrix} a & b \\ 0 & c\end{pmatrix}$ with $\mathbf{Q}' \leftarrow \mathcal{U}(\mathcal{O}_2(K_\mathbb{R}))$, $a \leftarrow \chi_{K_\mathbb{R}}$ and $b, c \leftarrow \mathcal{D}(0,1)$, conditioned on the event that for all $i \in [d]$ we have $|\sigma_i(a\cdot c)| \geq 1/d$. We can then write:

$$\mathbf{D}\cdot M = \mathbf{Q}'\cdot\begin{pmatrix} a & b \\ 0 & c\end{pmatrix}\cdot\begin{pmatrix} 1 & r \\ 0 & 1\end{pmatrix}\cdot\begin{bmatrix} 1/\gamma\cdot J_1 & \\ & \gamma\cdot J_2\end{bmatrix} = \mathbf{Q}'\cdot\begin{pmatrix} a & b+ar+b \\ 0 & c\end{pmatrix}\cdot\begin{bmatrix} 1/\gamma\cdot J_1 & \\ & \gamma\cdot J_2\end{bmatrix}.$$

Using the equality $\det\mathbf{D} = \mathcal{N}(ab)$, we obtain:

$$M' = |\det\mathbf{D}|^{-\frac{1}{2d}}\cdot\mathbf{D}\cdot M = \mathbf{Q}'\cdot\begin{pmatrix} 1 & r' \\ 0 & 1\end{pmatrix}\cdot\begin{bmatrix} 1/\gamma'\cdot J_1' & \\ & \gamma'\cdot J_2'\end{bmatrix},$$

where $r' = (b+ar)/c$, $\gamma' = \mathcal{N}(c/a)^{1/(2d)} \cdot \gamma$, $J_1' = (a/\mathcal{N}^{1/d}(a))J_1$ and $J_2' = (c/\mathcal{N}^{1/d}(c))J_2$. This proves the equality of distributions.

We now study $\gamma(M')$. For this, by the above, we can consider `Ideal-GR`. Thanks to the conditioning on the distribution of (a,c), we have:

$$\gamma' = \mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}} \cdot \gamma = \frac{\mathcal{N}(ac)^{\frac{1}{2d}}}{\mathcal{N}(a)^{\frac{1}{d}}}\gamma \geq \frac{1}{\sqrt{d}\cdot \mathcal{N}(a)^{\frac{1}{d}}}\gamma.$$

Now, note that without the conditioning, the coefficient a would be normally distributed, and the Gaussian tailbound would imply that $\mathcal{N}(a)^{1/d} \leq \sqrt{d}$ with probability $1 - 2^{-\Omega(d)}$. As the rejection occurs with probability at most $1 - 1/d^{O(1)}$ over the choice of (a,c) Gaussian, we still have that $\mathcal{N}(a)^{1/d} \leq \sqrt{d}$ with probability $1-2^{-\Omega(d)}$ for (a,c) distributed as in `Ideal-GR`. Overall, we obtain that $\gamma' \geq \gamma/d$ with probability $1-2^{-\Omega(d)}$. Using to the QR-standard form of M' with J_1' and J_2' of norm 1, we obtain that $\gamma(M') \geq \gamma' > 1$. By Lemma 2.8, the module M' has a unique rank-1 densest submodule. The QR-standard form leads us to consider the following rank-1 submodule of M':

$$U' = \frac{1}{\gamma'} \cdot \mathbf{Q}' \cdot \begin{bmatrix} 1 \\ 0 \end{bmatrix} \cdot J_1' = |\det \mathbf{D}|^{-1/(2d)} \cdot \mathbf{D} \cdot U.$$

It satisfies $\mathcal{N}(U') = 1/\gamma'^d \leq \gamma(M')^d$. By Lemma 2.8, it is contained in the unique densest rank-1 submodule of M. By primitivity, we have equality. ☐

E.9 Relations between the distributions of Definition 5.10

Let us first recall the definitions of the considered distributions.

$$D^{\text{rand}}_{B,\gamma} : \left(\mathbf{Q},\ \gamma\frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \frac{a}{\mathcal{N}^{\frac{1}{d}}(a)} J_1 {J_2}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ \frac{c}{\mathcal{N}^{\frac{1}{d}}(c)} \cdot J,\ \frac{b+a(r+x)}{c}\right),$$

$$D^{(1)}_{B,\gamma} : \left(\mathbf{Q},\ \gamma\frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right)\cdot\frac{au}{c}\cdot J_1 {J_2}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ J,\ u\frac{b+a(r+x)}{c}\right),$$

$$D^{(2)}_{B,\gamma} : \left(\mathbf{Q},\ \gamma\cdot\frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ I(J_1,J_2),\ J,\ u\frac{b+a(r+x)}{c\exp(\zeta)}\right),$$

$$D^{(3)}_{B,\gamma} : \left(\mathbf{Q},\ \gamma',\ I(J_1,J_2),\ J,\ \frac{B^{\frac{1}{d}}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}\cdot u\frac{b+a(r+x)}{c\exp(\zeta)}\right),$$

$$D^{(4)}_{B,\gamma} : (\mathbf{Q},\ \gamma',\ I(J_1,J_2),\ J,\ r''(J_1,J_2)),$$

$$D^{\text{target}}_{B,\gamma} : (\mathbf{Q},\ \gamma',\ I_1,\ I_2,\ r').$$

Where B, γ, J_1, J_1 and the random variables $\mathbf{Q}, a, b, c, x, \mathfrak{p}, I_1, I_2, I(\cdot,\cdot), J, \zeta, u, r', r''$ are defined in Definition 5.10.

Lemma E.1. *For any $B \geq 2, \gamma > 0$, $r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have*

$$D^{\text{rand}}_{B,\gamma}(J_1, J_2, r) = D^{(1)}_{B,\gamma}(J_1, J_2, r).$$

Proof. Let

$$A = \Big(\mathbf{Q},\ \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \frac{a}{\mathcal{N}^{\frac{1}{d}}(a)} J_1 J_2{}^{-1} J \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ \frac{c}{\mathcal{N}^{\frac{1}{d}}(c)} \cdot J,\ \frac{b + a(r+x)}{c}\Big)$$

be a sample from $D^{\text{rand}}_{B,\gamma}(J_1, J_2, r)$. As the distribution $\mathcal{U}(\mathcal{I}_1)$ is invariant by multiplication by a norm-1 ideal, the random variable $J' = c/\mathcal{N}^{1/d}(c) \cdot J$ is uniformly distributed in $\mathcal{I}_1$ (over the randomness of J, which is statistically independent of all other random variables). We have

$$A = \Big(\mathbf{Q},\ \gamma \frac{\mathcal{N}\left(\frac{c}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{a}{c} \cdot J_1 J_2{}^{-1} J' \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ J',\ \frac{b + a(r+x)}{c}\Big).$$

Now let u be uniform in $\{x \in K_{\mathbb{R}}, \forall i \in [d] : |\sigma_i(x)| = 1\}$, and $c' = cu$. As the distribution $\mathcal{D}_{K_{\mathbb{R}}}(0, 1)$ is invariant by multiplication by an element in this set, and the conditioning on (a, c) translates identically to (a, c'), the random variable (a, c') follows the same distribution as the random variable (a, c) (which is statistically independent of all other random variables). We have

$$A = \Big(\mathbf{Q},\ \gamma \frac{\mathcal{N}\left(\frac{c'}{a}\right)^{\frac{1}{2d}}}{\mathcal{N}(\mathfrak{p})^{\frac{1}{2d}}},\ \mathcal{N}^{\frac{1}{d}}\left(\frac{c}{a}\right) \cdot \frac{au}{c'} \cdot J_1 J_2{}^{-1} J' \frac{\mathfrak{p}}{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})},\ J',\ u\frac{b + a(r+x)}{c'}\Big).$$

We recognize the distribution $D^{(1)}_{B,\gamma}(J_1, J_2, r)$. □

Lemma E.2. *For any $B \geq 2, \gamma > 0$, $r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\mathrm{RD}\Big(D^{(2)}_{B,\gamma}(J_1, J_2, r)\ \Big\|\ D^{(1)}_{B,\gamma}(J_1, J_2, r)\Big) = O(1).$$

Proof. The result follows from the fact that $\mathcal{N}(c \cdot \exp(\zeta)) = \mathcal{N}(c)$, the data processing inequality and the bound:

$$\mathrm{RD}(c \cdot \exp(\zeta) \parallel c) = O(1).$$

The rest of the proof is devoted to establishing the latter.

Let $\zeta \in E$ fixed with $\|\zeta\|_\infty \leq 1/d$. When $d \geq 2$, we have that $2 - \exp(\zeta_i) > 0$ for all i. Therefore, by Lemma A.1 and the fact that $\mathcal{N}(\exp(\zeta)) = 1$, we have:

$$\mathrm{RD}(\mathcal{D}_{K_{\mathbb{R}}}(0, \exp(\zeta)) \parallel \mathcal{D}_{K_{\mathbb{R}}}(0, 1)) = \mathcal{N}(2 - \exp(\zeta))^{-\frac{1}{2}}.$$

As $|\zeta_i| \leq 1/d$ holds for all i, each embedding coefficient of $|(2 - \exp(\zeta))|$ is $\leq 1 - 1/d$. We hence obtain that

$$\mathcal{N}(2 - \exp(\zeta))^{-\frac{1}{2}} \leq (1 - 1/d)^{-\frac{d}{2}} = O(1).$$

To complete the proof, let us consider ζ as a random variable again. We use Lemma A.2 with $K_{\mathbb{R}}$ in place of $\mathbb{R}$ (which is fine, by the multiplicativity property of the Rényi divergence), to obtain:

$$\mathrm{RD}(c \cdot \exp(\zeta) \parallel c) \leq \mathbb{E}_{\zeta}\left(\mathrm{RD}(\mathcal{D}_{K_{\mathbb{R}}}(0, \exp(\zeta)) \parallel \mathcal{D}_{K_{\mathbb{R}}}(0,1))^{\frac{1}{2}}\right)^2.$$

By the analysis above, the latter upper bound is $O(1)$. □

Lemma E.3 (ERH). *For any $B \geq (\log \Delta_K)^{\Omega(1)}$, $\gamma > 0$, $r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\mathrm{RD}\left(D^{(3)}_{B,\gamma}(J_1, J_2, r) \,\middle\|\, D^{(2)}_{B,\gamma}(J_1, J_2, r)\right) = O(1).$$

Proof. Note that $D^{(3)}$ is obtained from $D^{(2)}$ by replacing all occurrences of c by $c \cdot \mathcal{N}^{1/d}(\mathfrak{p})/B^{1/d}$. The result then follows from the data processing inequality and the bound:

$$\mathrm{RD}\left(c \cdot \frac{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}{B^{\frac{1}{d}}} \,\middle\|\, c\right) = O(1)$$

The rest of the proof is devoted to proving the latter.

Let us fix a $\mathfrak{p}$ of norm $\leq B$, this implies that $2 - \mathcal{N}^{1/d}(\mathfrak{p})/B^{\frac{1}{d}} > 0$ so by Lemma A.1 we have

$$\mathrm{RD}\left(c \cdot \frac{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}{B^{\frac{1}{d}}} \,\middle\|\, c\right) \leq \mathcal{N}\left(\frac{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}{B^{\frac{1}{d}}} \cdot \left(2 - \frac{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}{B^{\frac{1}{d}}}\right)\right)^{-\frac{1}{2}} \leq \left(\frac{B}{\mathcal{N}(\mathfrak{p})}\right)^{\frac{1}{2}}.$$

Now, we consider $\mathfrak{p}$ as a random variable again. Thanks to the above, we have:

$$\mathbb{E}_{\mathfrak{p}}\left(\mathrm{RD}\left(c \cdot \frac{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}{B^{\frac{1}{d}}} \,\middle\|\, c\right)^{\frac{1}{2}}\right) \leq \frac{B^{\frac{1}{4}}}{\pi_K(B)} \sum_{\mathcal{N}(\mathfrak{p}) \leq B} \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{4}}}.$$

Abel's summation formula gives (see, e.g., [Apo98, Theorem 4.2]):

$$\begin{aligned}\sum_{\mathcal{N}(\mathfrak{p}) \leq B} \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{4}}} &= \frac{\pi_K(B)}{B^{\frac{1}{4}}} + \frac{1}{4}\int_2^B \pi_K(t) t^{-\frac{5}{4}}\,\mathrm{d}t \\ &\leq 1.1\frac{B^{\frac{3}{4}}}{\log B} + 1.1\int_{B_0}^B \frac{t^{-1/4}}{\log(t)}\,\mathrm{d}t + \int_2^{B_0} \pi_K(t) t^{-\frac{5}{4}}\,\mathrm{d}t\,,\end{aligned}$$

where $B_0 = (\log \Delta_K)^{\Omega(1)}$ is such that for $B \geq B_0$ we have $\pi_K(B) \leq 1.1B/\log B$ (see Section 2). The last term in the upper bound is $\leq 2^{-1/4}\pi_K(B_0) \leq B_0$. Assuming that $B \geq B_0^2$, the latter is $\leq B^{1/2}$. Overall, we obtain that

$$\sum_{\mathcal{N}(\mathfrak{p}) \leq B} \frac{1}{\mathcal{N}(\mathfrak{p})^{\frac{1}{4}}} \leq 5\frac{B^{\frac{3}{4}}}{\log(B)}.$$

Using the lower bound $\pi_K(B) \geq 0.9B/\log B$ from Section 2, we then obtain that

$$\mathbb{E}_{\mathfrak{p}} \left(\mathrm{RD}\left(c \cdot \frac{\mathcal{N}^{\frac{1}{d}}(\mathfrak{p})}{B^{\frac{1}{d}}} \,\middle\|\, c \right)^{\frac{1}{2}} \right) \leq O(1).$$

Finally, Lemma A.2 allows us to conclude. ◻

Lemma E.4. *For $B \geq 2$, $\gamma \geq d^{1/4} \Delta_K^{1/(2d)}$, $r \in K_{\mathbb{R}}$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\mathrm{SD}\left(D^{(3)}_{B,\gamma}(J_1, J_2, r), D^{4}_{B,\gamma}(J_1, J_2) \right) \leq 2^{-\Omega(d)}.$$

To prove Lemma E.4, we will use the following result on the closeness to uniformity of a Gaussian distribution over $K_{\mathbb{R}}$, when it is folded modulo an ideal lattice.

Lemma E.5 (Adapted from [PRS17, Lemma 6.9]). *Let I an ideal, $s \in K_{\mathbb{R}}^{+}$ and $\mathbf{s} = (\sigma_i(s))_{i \in [r_1+r_2]}$. If $\mathcal{N}(s) \geq \Delta_K \cdot \mathcal{N}(I)$, then we have:*

$$\mathrm{SD}\left(\mathcal{D}_{K_{\mathbb{R}}}(0, \mathbf{a}) \bmod I, \mathcal{U}(K_{\mathbb{R}} \bmod I) \right) \leq 2^{-\Omega(d)}.$$

Proof (Lemma E.4). We consider the following sample from $D^{(3)}_{B,\gamma}(J_1, J_2, r)$:

$$\left(\mathbf{Q}, \gamma', I(J_1, J_2), J, B^{\frac{1}{d}} \cdot u \frac{b + a(r+x)}{c \exp(\zeta) \mathcal{N}^{\frac{1}{d}}(\mathfrak{p})} \right).$$

Note that $b \sim \mathcal{D}_{K_{\mathbb{R}}}(0,1)$ is independent of all other variables and occurs only once in the sample above. Let $b' = B^{1/d} \cdot sb/(c \exp(\zeta) \mathcal{N}^{1/d}(\mathfrak{p}))$. Over the randomness of b (and assuming all other random variables are fixed), it is distributed as $\mathcal{D}_{K_{\mathbb{R}}}(0, B^{1/d}/(|c| \exp(\zeta) \mathcal{N}^{1/d}(\mathfrak{p})))$. We now consider the folding of b' modulo the ideal $I' := \gamma'^{-2} I(J_1, J_2) \cdot J^{-1}$. Lemma E.5 implies that if

$$\frac{B}{\mathcal{N}(c)\mathcal{N}(\mathfrak{p})} \geq \Delta_K \cdot \mathcal{N}(I'),$$

then $\mathrm{SD}(c' \bmod I', \mathcal{U}(K_{\mathbb{R}} \bmod I') \leq 2^{-\Omega(d)}$, leading to the result. It hence suffices to prove the premise.

As $I(J_1, J_2), J \in \mathcal{I}_1$ and $\mathcal{N}(\mathfrak{p}) \leq B$, using the definition of γ', it suffices that we have $\gamma^{2d} \geq \Delta_K \mathcal{N}(a)$. By the Gaussian tail bound, we have $\mathcal{N}(a) \leq d^{d/2}$ with probability $1 - 2^{-\Omega(d)}$, which suffices for our purposes. ◻

Lemma E.6. *For $B \geq (d^d \Delta_k)^{\Omega(1)}$, $\gamma > 0$ and $J_1, J_2 \in \mathcal{I}_1$, we have:*

$$\mathrm{SD}\left(D^{(4)}_{B,\gamma}(J_1, J_2), D^{\mathrm{target}}_{B,\gamma} \right) \leq 2^{-\Omega(d)}.$$

Proof. By Lemma 2.4, the distribution of $\frac{\mathfrak{p}}{\mathcal{N}^{1/d}(\mathfrak{p})} \cdot u \exp(-\zeta)$ is within statistical distance $2^{-\Omega(d)}$ from $\mathcal{U}(\mathcal{I}_1)$. The latter distribution being invariant by multiplication by norm-1 ideals, we obtain that the distribution of $I(J_1, J_2)$ is at statistical distance $2^{-\Omega(d)}$ from $\mathcal{U}(\mathcal{I}_1)$, over the random choices of $\mathfrak{p}$, s and ζ. As they are independent of $\mathbf{Q}$, γ' and J, and as the distribution of the last tuple entry is a function of the others, we obtain the result. ◻

F Missing proofs from Section 6

F.1 Proof of Lemma 6.2

Using the notations from Definition 5.1, the gap of the module M is equal to $\gamma(M) = \gamma'\mathcal{N}(c/a)^{1/(2d)}/B^{1/(2d)}$. Now, by the conditioning on the pair (a, c), we have $\mathcal{N}(c) \geq 1/(d^d\mathcal{N}(a))$. Also, by the Gaussian tail bound, we have $\|a\| \leq \sqrt{d}$ with probability $1 - 2^{-\Omega(d)}$. The inequality $\mathcal{N}(a) \leq \|a\|/\sqrt{d}$ then leads to the result. □

F.2 Proof of Lemma 6.3

Let us write $(\mathbf{B}, \mathbb{I}) = \mathbf{Q}\cdot((1,0)^T \cdot 1/\gamma J_1 + (r,1)^T \cdot \gamma J_2)$ and $\mathbf{Y} = R\cdot(\mathbf{I}+(2d)^{-3/2}\cdot\mathbf{E})$ with R as define in `DualRound` (Algorithm 3.1) and $\|e_{ij}\| \leq 1$ for all $i, j \in [2]$ (see Lemma 3.5). We consider the QR-factorization of $\mathbf{Q}^{-1} \cdot \mathbf{Y} \cdot \mathbf{Q}$:

$$\mathbf{Q}^{-1} \cdot \mathbf{Y} \cdot \mathbf{Q} = R \cdot \mathbf{Q}' \cdot \begin{bmatrix} x & y \\ 0 & z \end{bmatrix},$$

for some $x, y, z \in K_{\mathbb{R}}$. In particular, we have that $\mathcal{N}(x)$ is the algebraic norm of the first column of $\mathbf{I}_2 + \mathbf{E}'$, where $\mathbf{E}' = \mathbf{Q}^{-1} \cdot \mathbf{E} \cdot \mathbf{Q}$ satisfies $\|e'_{ij}\|_\infty \leq \sqrt{2d}$ for $i, j \in [2]$. This implies that $\mathcal{N}(x) \leq 1 + 1/(2d)$. In the same vein as in the proof of Theorem 5.9, this implies that

$$\mathcal{N}\big(\mathbf{Y} \cdot U\big) \leq R^d \cdot \big(1 + \frac{1}{2d}\big)^d \cdot \mathcal{N}(U) \leq R^d \cdot \sqrt{e} \cdot \mathcal{N}(U),$$

where $U = \mathbf{Q} \cdot (1,0)^T \cdot 1/\gamma J_1$. The result then follows from Lemma 3.6 and the fact that $\mathcal{N}(\mathbf{Y} \cdot M) = \det(\mathbf{Y}) \cdot \mathcal{N}(M)$. □

F.3 Proof of Lemma 6.5

Wlog, we may assume that the gap of the γ'-wc-uSVP$^{\mathcal{N}}_{\text{mod}}$ instance $(\mathbf{B}, \mathbb{I})$ satisfies $\gamma' \leq 2^d \Delta_K^{O(1/d)}$, as otherwise the problem can be solved in polynomial time using LLL [LLL82]. We cover the interval $[2\log(\Delta_K)^{O(1/d)} \cdot \gamma, 2^d\Delta_K^{O(1/d)}]$ by at most $O(d^2 + \log \Delta_K)$ intervals of the form $\gamma \cdot [(1 + 1/(3d))^i, (1 + 1/(3d))^{i+1}]$, and guess uniformly the i for which contains the gap of the module M spanned by $(\mathbf{B}, \mathbb{I})$. The guess is correct with probability $\Omega(1/(d^2 + \log \Delta_K))$ and, in the following, we only analyze what happens when this occurs.

The next step is to find a prime ideal $\mathfrak{p}$ such that $\mathcal{N}(\mathfrak{p})^{1/(2d)} \in \gamma' \cdot [(1 + 1/(3d))^{i-1}, (1 + 1/(3d))^i]$. As $\gamma' \geq 2\log(\Delta_K)^{O(1/d)}$, we can use Lemma 2.3 to sample $\mathfrak{p}$ uniformly among the prime ideals with norms $\leq (1 + 1/(3d))^{di}$. By the estimates on π_K stated in Section 2, the value $\mathcal{N}(\mathfrak{p})^{1/(2d)}$ belongs to the appropriate interval with probability $\Omega(1)$. We assume this is the case. Note that we then have that $\gamma'/\mathcal{N}(\mathfrak{p}) \in \gamma \cdot [1, 1 + 1/d]$.

We then sample $\overline{\mathbf{b}^\vee}$ uniformly in $(M^\vee/\mathfrak{p}M^\vee)\setminus\{\mathbf{0}\}$, and sparsify M by $(\overline{\mathbf{b}^\vee}, \mathfrak{p})$, using Lemma 3.4. By Lemmas 3.2 and 3.3, the gap of the sparsified module M'

is $\gamma'/\mathcal{N}(\mathfrak{p})$, with probability $\Omega(1)$, and the pseudo-basis of M' is a valid $\gamma^{\approx}$-wc-uSVP$^{\mathcal{N}}_{\text{mod}}$ instance. Finally, note that when the latter event occurs, if U is the densest rank-1 submodule of M, then $\mathfrak{p}U$ is the densest rank-1 submodule of M' (as in the proof of Theorem 5.6). This completes the description and the analysis of the reduction. □

F.4 Proof of Lemma 6.6

By Lemma 6.2, samples from D^{uSVP}_{γ} and $D^{\text{uSVP}}_{\gamma'}$ are indeed γ-uSVP$^{\mathcal{N}}$ instances.

Now, note that $D^{\text{uSVP}}_{\gamma'}$ is obtained from D^{uSVP}_{γ} by replacing all the occurences of c by $c \cdot (\gamma'/\gamma)^2$ in Definition 5.1. The result then follows from the data processing inequality and the bound:

$$\text{RD}\big(c \,\|\, c \cdot (\gamma'/\gamma)^2\big) = O(1)$$

The rest of the proof is devoted to proving the latter. We have $\gamma'/\gamma \geq 1$, implying that $2(\gamma'/\gamma)^2 - 1 \geq 1$. By Lemma A.1 this implies that

$$\text{RD}\big(c \,\|\, c \cdot (\gamma'/\gamma)^2\big) \leq \mathcal{N}\left(\frac{(\gamma'/\gamma)^4}{2(\gamma'/\gamma)^2 - 1}\right)^{1/2} \leq (1 + 1/d)^{2d} = O(1).$$

□

F.5 Proof of Lemma 6.7

The reduction first runs algorithm $\texttt{Randomize}_B$ from Theorem 5.2. It then calls $\texttt{DualRound}_{\varsigma,\beta,\varepsilon}$ and $\texttt{HNF}$. The parameters B, ς, β and ε are set exactly as in the sampling algorithm for D^{uSVP}. It then calls the $(D^{\text{uSVP}}_{\gamma'}, \gamma'')$-uSVP$^{\mathcal{N}}_{\text{mod}}$ oracle and pulls the returned rank-1 submodule back to a rank-1 submodule of the input module, using the $\mathbf{Y}$ matrix from `DualRound` and the `aux` output from `Randomize`.

The runtime bound comes from Theorem 5.2 and Lemma 3.5. Correctness follows from Theorem 5.2, Lemmas 6.3 and Lemma 6.6. □