Lattice Reduction for Modules, or How to Reduce ModuleSVP to ModuleSVP

Tamalika Mukherjee *
Purdue University
tmukherj@purdue.edu

Noah Stephens-Davidowitz†
Massachusetts Institute of Technology
noahsd@gmail.com

October 3, 2019

Abstract

We show how to generalize lattice reduction algorithms to module lattices in order to reduce γ-approximate ModuleSVP over module lattices with rank $k \geq 2$ to γ'-approximate ModuleSVP over module lattices with rank $2 \leq \beta \leq k$. To do so, we modify the celebrated slide-reduction algorithm of Gama and Nguyen to work with module filtrations, a higher-dimensional analogue of the ($\mathbb{Z}$-)basis of a lattice.

The particular value of γ that we achieve depends on the underlying number field K, the ring $R \subset K$, and the embedding (as well as, of course, k and β). However, for reasonable choices of these parameters, the approximation factor that we achieve is surprisingly close to the one achieved by "plain" lattice reduction algorithms, which require an arbitrary SVP oracle in the same dimension. In other words, we show that ModuleSVP oracles are nearly as useful as SVP oracles for solving approximate ModuleSVP in higher dimensions.

Our result generalizes the recent independent result of Lee, Pellet-Mary, Stehlé, and Wallet, which works in the important special case when $\beta = 2$ and $R = \mathcal{O}_K$ is the ring of integers of K under the canonical embedding. Indeed, at a high level our reduction can be thought of as a generalization of theirs in roughly the same way that slide reduction generalizes LLL reduction.

1 Introduction

A (rational) lattice $\mathcal{L} \subset \mathbb{Q}^d$ is the set of all integer linear combinations of finitely many generating vectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in \mathbb{Q}^d$,

$$\mathcal{L} := \{z_1\boldsymbol{y}_1 + \cdots + z_m\boldsymbol{y}_m \ : \ z_i \in \mathbb{Z}\} \ .$$

For an approximation factor $\gamma \geq 1$, the γ-approximate Shortest Vector Problem (γ-SVP) asks us to find a non-zero vector $\boldsymbol{y} \in \mathcal{L}$ whose length is within a factor γ of the minimum possible.

Lattices have played a key role in computer science since Lenstra, Lenstra, and Lovász published their celebrated LLL algorithm, which solves γ-SVP for $\gamma = 2^{O(d)}$ in polynomial time [LLL82], essentially by reducing the problem to many instances of exact SVP in two dimensions. In spite

*This work was done while being supported by The Center for Science of Information, an NSF Science and Technology Center, Cooperative Agreement # CCF 0939370.

†Supported by NSF-BSF grant number 1718161 and NSF CAREER Award number 1350619 via Vinod Vaikuntanathan.

of this very large approximation factor, this algorithm has found innumerable applications [LLL82, Sha84, Bab86, SE94, NS01, NV10, FS10].

Lattices have taken on an even larger role in recent years because of the growing importance of lattice-based cryptography [Ajt96, HPS98, GPV08, Reg09, Pei09, SSTX09, LPR10, Pei16, Mah18]—that is, cryptography whose security relies on the hardness of γ-SVP or a closely related problem for some γ (typically, $\gamma = \text{poly}(d)$). These schemes have several advantages, such as worst-case to average-case reductions, which show that some of these schemes are actually provably secure under the assumption that γ'-SVP is hard [Ajt96, MR07, Reg09, LPR10, LS15, PRS17]. They are also thought to be secure against quantum attackers, and for this reason, they are likely to be standardized by NIST (the United States' National Institute for Standards and Technology) for widespread use in the near future [NIS18].

However, one drawback of these schemes is their inefficiency (though, see [BCD+16, ABD+19] for some improvements). Loosely speaking, this inefficiency arises from the fact that a lattice in dimension d "is typically specified with d^2 numbers"—d generating vectors, each with d coordinates. To get around this, cryptographers often use lattices with certain additional symmetries [HPS98, PR06, SSTX09, LPR10, SS11, LS12, DD12, LS15, PRS17].

In particular, they use (variants of) *module lattices*. For a number field K of degree n (i.e., $K := \mathbb{Q}[x]/p(x)$ for an irreducible polynomial $p(x)$ of degree n) with a full-rank discrete subring $R \subset K$ (such as $\mathbb{Z}[x]/p(x)$ when $p \in \mathbb{Z}[x]$ is monic, or the ring of integers $\mathcal{O}_K$ of K), a module lattice over R is the set of all R-linear combinations of finitely many generating vectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in K^k$,

$$\mathcal{M} := \{r_1\boldsymbol{y}_1 + \cdots + r_m\boldsymbol{y}_m \ : \ r_i \in R\} \ .$$

By embedding the number field K into $\mathbb{R}^n$ (or, equivalently, by equipping K with an inner product), we can view module lattices as $(k \cdot n)$-dimensional "plain" lattices. We typically think of n as large (i.e., $n \to \infty$) and k as a relatively small constant.

We can then define (γ, k)-ModuleSVP over R as the restriction of γ-SVP to module lattices $\mathcal{M} \subseteq K^k$ over R. Clearly, (γ, k)-ModuleSVP is no harder than γ-SVP over lattices with rank kn. A key question is whether we can do (significantly) better. In other words, are there (significantly) faster algorithms for ModuleSVP than there are for SVP? Does the specialization to module lattices (which yields large efficiency benefits for cryptography) impact security?

Many cryptographic schemes rely on the assumption that no such algorithms exist. E.g., about half of the candidate encryption schemes under consideration by NIST would be broken in practice if significantly faster algorithms were found for ModuleSVP [NIS18]. (Just one relies on "plain" lattices.) We would therefore like to understand the hardness of ModuleSVP as soon as possible.

Until recently, one might have conjectured that (γ, k)-ModuleSVP is essentially as hard as γ-SVP on rank kn lattices for all γ and k. However, a recent and still active line of work has shown much faster algorithms for the $k = 1$ case [CGS14, CDPR16, CDW17, Duc17, DPW19, PHS19], in which case the problem is called IdealSVP. Most cryptographic schemes are not known to be broken by these algorithms, or even with access to an oracle for exact IdealSVP. However, similar improvement for the case $k = 2$ would yield faster algorithms for both the Ring-LWE problem [SSTX09, LPR10] and the NTRU problem [HPS98], which would break most of the cryptographic schemes based on structured lattices. (We are intentionally ignoring many important details here for simplicity. We refer the reader to [Pei15, Duc17, DPW19, PHS19] for a more careful discussion.)

Therefore, (ignoring a number of important details) the security of many cryptographic schemes essentially relies on the assumption that (γ, k)-ModuleSVP for $k \geq 2$ is qualitatively different than

γ-IdealSVP $= (\gamma, 1)$-ModuleSVP. More generally, this recent line of work in the $k = 1$ cases suggests that we need a better understanding of (γ, k)-ModuleSVP for all γ and k.

Much of our understanding of γ-SVP comes from *basis reduction algorithms* [LLL82, SE94, GN08, MW16, ALNS19]. These algorithms allow us to reduce γ-SVP in a high dimension d to γ'-SVP in a lower dimension m (known as the block size) for some approximation factor γ depending on d, m, and γ'. Indeed, the LLL algorithm is an example of such a reduction for the case $m = 2$. For the approximation factors relevant to cryptography, our fastest algorithms rely on basis reduction. In fact, these are essentially our only non-trivial provably correct algorithms for not-too-small approximation factors (formally, for any $\gamma \gg \sqrt{d}$ [ALNS19]).

I.e., to solve γ-SVP (or, for that matter, (γ, k)-ModuleSVP) for not-too-small γ, the best strategy we have is to reduce the problem to many instances of SVP with a smaller approximation factor over lower-dimensional "blocks." The current state of the art, due to [ALNS19] and building heavily on the work of Gama and Nguyen [GN08], achieves an approximation factor of

$$\gamma = \gamma' \cdot (\gamma' \sqrt{\beta n})^{\frac{2(k-\beta)}{\beta - 1/n}} \tag{1}$$

for block size $m := \beta n$ and dimension $d := kn$. For cryptanalysis, we typically must take $\beta = \Omega(k)$ and $\gamma' \leq \text{poly}(d)$ in order to achieve an approximation factor γ that is polynomial in the dimension $d = kn$. (We have of course chosen this rather strange parameterization to more easily compare with our results for ModuleSVP.)

1.1 Our results

1.1.1 Lattice reduction for Modules.

Our primary contribution is the following reduction.

Theorem 1.1 (Informal, see the discussion below and Theorem 5.10)**.** *For $2 \leq \beta < k$ with β dividing k, there is an efficient reduction from (γ, k)-ModuleSVP to (γ', β)-ModuleSVP, where*

$$\gamma = (\gamma')^2 n \cdot (\gamma' \sqrt{\beta} n)^{\frac{2(k-\beta)}{\beta - 1}} .$$

The case $\beta = 2$ is of particular interest because of its relevance to cryptography. We note that, before this work was finished, Lee, Pellet-Mary, Stehlé, and Wallet published essentially the same reduction for this important special case [LPSW19]. (Formally, they only showed this for the canonical embedding for the ring of integers of a number field, but it is easy to see that it generalizes. They also showed a very interesting algorithm for $(\gamma, 2)$-ModuleSVP, which requires preprocessing. We refer the reader to [LPSW19] for the details.) For this $\beta = 2$ case, the reduction can be viewed as a generalization of the LLL algorithm. (We present the $\beta = 2$ case separately in Section 4.)

In the general case $\beta \geq 2$, we note the obvious resemblance between the approximation factor achieved by Theorem 1.1 and Eq. (1). Indeed, our reduction can be viewed as a generalization of Gama and Nguyen's celebrated slide reduction [GN08] to the module case (see also [ALNS19]). Therefore, we can interpret Theorem 1.1 as saying that "a ModuleSVP oracle is almost as good as a generic SVP oracle for basis reduction over module lattices."

Finally, notice that this informal version of Theorem 1.1 does not mention the number field K, the associated embedding, or the ring R. In fact, the reduction works for *any* number field K,

(discrete, full-rank) ring $R \subset K$, and embedding of K into $\mathbb{R}^n$, with two caveats. (See Theorem 5.10 for the precise statement.)

First, the approximation factor that we achieve depends on certain geometric properties of the ring and the embedding. The approximation factor shown in Theorem 1.1 is (a loose upper bound on) what we achieve for the canonical embedding of the ring of integers of a cyclotomic number field.

Second, for $\beta > 2$ and ring embeddings that do not satisfy a certain duality condition (which is, for example, satisfied by any ring in the canonical embedding that is closed under complex conjugation), our reduction requires a ModuleSVP oracle over dual module lattices as well.

1.1.2 Two variants.

As additional contributions, we note that our reduction can also be used to solve two variants of ModuleSVP. (Like Theorem 1.1, these two variants also require an oracle over dual module lattices for $\beta > 2$ and insufficiently nice ring embeddings.)

The first variant is known as ModuleHSVP (where the H is in honor of Hermite). This problem asks us to find a non-zero vector that is short relative to the determinant of the module lattice $\mathcal{M}$, rather than relative to the shortest non-zero vector. I.e., (γ, k)-ModuleHSVP asks us to find a non-zero vector $\boldsymbol{x}$ in a rank-k module lattice $\mathcal{M}$ with $\|\boldsymbol{x}\| \leq \gamma \cdot \det(\mathcal{M})^{1/(kn)}$. For $\gamma \gtrsim \sqrt{kn}$, there is always a non-zero vector of this weight. In particular, $(\sqrt{n}\gamma, k)$-ModuleHSVP trivially reduces to (γ, k)-ModuleSVP, but our reduction achieves a better approximation factor than what one would obtain by combining this trivial reduction with Theorem 1.1. (The same is true of most "plain" basis reduction algorithms [GN08, ALNS19], though some only work for Hermite SVP [MW16].) This variant of SVP is enough for most cryptanalytic applications, so that this better approximation factor could prove to be quite useful in practice. (In particular, the analogous result for plain basis reduction algorithms is often used in cryptanalysis.)

Theorem 1.2 (Informal, see Theorem 5.10)**.** *For $2 \leq \beta < k$ with β dividing k, there is an efficient reduction from (γ_H, k)-ModuleHSVP to (γ', β)-ModuleSVP, where*

$$\gamma_H := \gamma' \sqrt{n} \cdot (\gamma' \sqrt{\beta n})^{\frac{k-1}{(\beta-1)}}$$

Again, the approximation factor shown in Theorem 1.2 is what we achieve for the canonical embedding of the ring of integers of a cyclotomic number field. See Theorem 5.10 for the general result.

Our second variant has no analogue for plain lattices. We consider the (γ, k)-Dense Ideal Problem ((γ, k)-DIP), in which the goal is to find a rank-one submodule $\mathcal{M}'$ (i.e., an ideal) such that $\det(\mathcal{M}')^{1/n}$ is within a factor γ of the minimum possible. This problem is in a sense more natural in our context. Indeed, Theorem 1.1 is perhaps best viewed as a consequence of Theorem 1.3. We again note the obvious similarity between Theorem 1.3 and Eq. (1). (There is an analogous result for what we might call "RankinDIP," which asks us to find an ideal whose determinant is small relative to $\det(\mathcal{M})^{1/(nk)}$, just like ModuleHSVP asks for a vector that is short relative to $\det(\mathcal{M})^{1/(nk)}$. For simplicity, we do not bother to make this formal.)

Theorem 1.3 (Informal, see Corollary 5.8)**.** *For $2 \leq \beta < k$ with β dividing k, there is an efficient reduction from (γ, k)-DIP to (γ', β)-DIP, where*

$$\gamma := \gamma' \cdot (\gamma' \sqrt{\beta n})^{\frac{2(k-\beta)}{\beta-1}} \ .$$

In fact, Theorem 1.3 holds as stated for any number field, ring, and embedding. I.e., the approximation factor does not depend on the geometry of the ring.

1.2 Our techniques

From bases to filtrations. Lattice basis reduction algorithms take as input a ($\mathbb{Z}$-)basis $(\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d)$ of a lattice $\mathcal{L} \subset \mathbb{Q}^d$ and they iteratively "shorten" the basis vectors using an oracle for SVP in $m < d$ dimensions. More specifically, let $\mathcal{L}_i$ be the lattice spanned by $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_i$. Basis reduction algorithms work by finding short vectors in lattices of the form $\mathcal{L}_{[i,j]} := \pi_{\mathcal{L}_{i-1}^{\perp}}(\mathcal{L}_j)$, where $\pi_{\mathcal{L}_i^{\perp}}$ represents projection onto the subspace orthogonal to $\mathcal{L}_i$. In the basis reduction literature, the $\mathcal{L}_i$ and $\mathcal{L}_{[i,j]}$ are typically not defined explicitly. Instead, corresponding bases for these lattices are defined.

To generalize this idea to module lattices, our first challenge is to find the appropriate analogue of a basis. Indeed, while lattices with rank d over $\mathbb{Z}$ have a $\mathbb{Z}$-basis consisting of d (linearly independent) lattice vectors, the analogous statement is typically not true for more general rings R. I.e., our module lattice $\mathcal{M}$ of rank k will not always have an R-basis consisting of only k elements. (E.g., rank one module lattices are ideals, and they have an R-basis consisting of a single element if and only if they are principle. More generally, all rank k module lattices have an R-basis consisting of k vectors if and only if R is a principle ideal domain.) This means that basis-reduction techniques do not really make sense over an R-basis, since, for example, the projection orthogonal to the first i elements in an R-basis is not necessarily a module lattice with rank $k - i$.

So, instead of generalizing $\mathbb{Z}$-bases themselves, we work directly with the sublattices $\mathcal{L}_i$ and blocks $\mathcal{L}_{[i,j]}$. To that end, we define a *module filtration* $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$ as a sequence of k (primitive) submodules with strictly increasing ranks (over K). Filtrations have the nice property that the projection $\mathcal{M}_{[i,j]} := \pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}_j)$ of $\mathcal{M}_j$ orthogonal to $\mathcal{M}_i$ is itself a module lattice with rank $j - i + 1$. They are well-behaved in other ways as well. For example, the determinant of $\mathcal{M}$ is given by the product of the determinants of the rank-one projections $\widetilde{\mathcal{M}}_i := \pi_{\mathcal{M}_{i-1}}(\mathcal{M}_i)$, which is analogous to the fact that the determinant of a lattice is given by the product of the lengths of the Gram-Schmidt vectors $\widetilde{\boldsymbol{b}}_i$ of any basis. These are the key properties that allow us to perform basis reduction using SVP oracle calls only on module lattices.[1]

From vectors to ideals (or sublattices). By working with filtrations, our reduction is most naturally viewed as a "basis reduction algorithm" with Gram-Schmidt vectors $\pi_{\mathcal{L}_{i-1}}(\boldsymbol{b}_i)$ replaced by ideals $\pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}_i)$, and lengths replaced by the determinant. This naturally gives rise to Theorem 1.3—a reduction from DIP to DIP. Indeed, this DIP-to-DIP reduction actually "never looks at a vector," and it can be viewed as a specialization to module lattices of a more general reduction from the problem of finding dense rank-n sublattices of a kn-dimensional lattice to the problem of finding dense rank-n sublattices in a βn-dimensional lattice (though we do not bother to show this formally).

From ideals back to vectors. In order to obtain our main result, we must convert this DIP-to-DIP reduction into a reduction from ModuleSVP to ModuleSVP. To do so, we use well-known

[1] In [FS10, LPSW19], the authors work with *pseudobases*, which consist of vectors $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_k \in K^k$ *and* ideals $\mathcal{I}_1, \ldots, \mathcal{I}_k \subset K$ such that $\mathcal{M} = \mathcal{I}_1\boldsymbol{b}_1 + \cdots + \mathcal{I}_k\boldsymbol{b}_k$. These are essentially equivalent to filtrations. E.g., they can be converted into filtrations by setting $\mathcal{M}_i := \mathcal{I}_1\boldsymbol{b}_1 + \cdots + \mathcal{I}_i\boldsymbol{b}_i$.

relationships between the length of short non-zero vectors and the determinants of dense rank-one submodules. Specifically, we use (1) Minkowski's theorem, which states that any dense submodule must contain a short vector (and holds for all lattices, not just module lattices); and (2) the fact that the R-span of a short vector must be a relatively dense ideal, which has no analogue for lattices in general. (The latter property depends on the geometry of the ring, which is why our approximation factors also depend on this geometry.)

Therefore, a ModuleSVP oracle can be used to find a short vector, which must generate a dense ideal. And, we may use a DIP oracle to find a low-rank submodule that contains a short vector. This allows us to move between DIP and ModuleSVP, which yields our main result.

1.3 Related work

The most closely related work to this paper is the recent independent work of Lee, Pellet-Mary, Stehlé, and Wallet [LPSW19], which was published before this work was finished. [LPSW19] proved Theorem 1.1 in the important special case when $\beta = 2$ and $R = \mathcal{O}_K$ is the ring of integers of the number field K under the canonical embedding. Their reduction is essentially identical to ours, though they use a formally different notion of a reduced basis that seems not to generalize quite as nicely for larger β.[2] They also show a surprising algorithm for $(\gamma, 2)$-ModuleSVP (formally, a reduction from this problem to the Closest Vector Problem over a fixed lattice), which can be used to instantiate the $(\gamma, 2)$-ModuleSVP oracle.

For $\beta > 2$, our reductions are generalizations of the slide-reduction algorithm of Gama and Nguyen [GN08], and our work is largely inspired by theirs. Indeed, both our notion of a reduced filtration and our algorithm for constructing one are direct generalizations of the corresponding ideas in [GN08] from bases of $\mathbb{Z}$-lattices to filtrations of module lattices.

There are also other, rather different notions of basis reduction for module lattices from prior work. For example, for certain Euclidean domains, Napias showed that the LLL algorithm (and Gauss's algorithm for rank-two lattices) generalizes quite nicely, with no need for an oracle [Nap96]. Follow-up work showed how to extend this to more Euclidean domains [GLM09, KL17]. However, it seems that algorithms of this type can only work in the Euclidean case [LPL18], and for the cryptographic applications that interest us most, the ring R is typically not Euclidean—or even a principle ideal domain. (The algorithm of [LPSW19] for $(\gamma, 2)$-ModuleSVP is particularly surprising precisely *because* it seems to mimic Gauss's algorithm even though it works for non-Euclidean rings.) In another direction, Fieker and Stehlé showed how to efficiently convert an LLL-reduced $\mathbb{Z}$-basis for a module lattice into an LLL-reduced pseudobasis, which in our language is essentially a filtration that is reduced in a certain sense [FS10]. I.e., they show how to efficiently convert a relatively short $\mathbb{Z}$-basis into a relatively short filtration.

Finally, there are many works showing relationships between variants of LWE (an average-case problem related to SVP) over different rings and other algebraic structures [LS12, DD12, RSSS17, RSW18, BBPS18, PP19], and other works showing reductions from various worst-case problems on structured lattices to such variants of LWE and SIS [LM06, PR06, Mic07, SSTX09, LPR10, LS12, LS15, RSSS17, PRS17, RSW18].

[2] Specifically, in the notation introduced above, they work with the ratio of $\det(\pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}_i))$ to $\det(\pi_{\mathcal{M}_i^{\perp}}(\mathcal{M}_{i+1}))$, while we work with the ratio of $\det(\pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}_i))$ relative to the minimum possible for a rank-one submodule of $\pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}_{i+1})$. The distinction is not particularly important for $\beta = 2$, but the analogous conditions for $\beta > 2$ are quite different. In particular, the most natural generalization of the first notion seems to only yield a solution to ModuleHSVP.

2 Preliminaries

For $\boldsymbol{x} \in K^k$ and a subspace $W \subseteq K^k$, we write $\pi_W(\boldsymbol{x})$ for the projection of $\boldsymbol{x}$ onto W. I.e., if $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_\ell$ is an orthonormal basis of W, then $\pi_W(\boldsymbol{x}) = \langle \boldsymbol{b}_1, \boldsymbol{x} \rangle \cdot \boldsymbol{b}_1 + \cdots + \langle \boldsymbol{b}_\ell, \boldsymbol{x} \rangle \cdot \boldsymbol{b}_\ell$. We similarly define $\pi_W(S) := \{\pi_W(\boldsymbol{x}) \ : \ \boldsymbol{x} \in S\}$ for subsets $S \subseteq K^k$.

2.1 Lattices and Rankin's constants

A lattice $\mathcal{L} \subset \mathbb{R}^d$ with rank m is the $\mathbb{Z}$-span of some basis $\mathbf{B} := (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m)$ of $\mathbb{R}$-linearly independent vectors,

$$\mathcal{L} := \{z_1 \boldsymbol{b}_1 + \cdots + z_m \boldsymbol{b}_m \ : \ z_i \in \mathbb{Z}\} \ .$$

We sometimes call this a $\mathbb{Z}$-basis, and we write $m := \mathrm{rank}_{\mathbb{R}}(\mathcal{L})$. Equivalently, lattices are additive subgroups of $\mathbb{R}^d$ with at most finitely many points inside any bounded set. Equivalently again, a lattice $\mathcal{L}$ is the $\mathbb{Z}$-span of any (not necessarily linearly independent) vectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_{m'} \in \mathbb{R}^d$ if any $\mathbb{R}$-linear dependence $a_1 \boldsymbol{x}_1 + \cdots + a_\ell \boldsymbol{x}_\ell = 0$ for $a_i \in \mathbb{R}$ and $\boldsymbol{x}_i \in \mathcal{L}$ is also a $\mathbb{Q}$-linear dependence, i.e., $a_i \in \mathbb{Q}$. (E.g., the $\mathbb{Z}$-span of $(1,0) \in \mathbb{R}^2$ and $(\sqrt{2}, 0) \in \mathbb{R}^2$ is not a lattice, but the $\mathbb{Z}$-span of $(1,0) \in \mathbb{R}^2$ and $(\sqrt{2}, 1) \in \mathbb{R}^2$ is a lattice.)

The lattice determinant $\det(\mathcal{L}) := \sqrt{\det(\mathbf{B}^T \mathbf{B})}$ for any basis $\mathbf{B}$ of $\mathcal{L}$ (the choice of basis does not matter). A lattice is full rank if $m = d$, and we sometimes assume that lattices are full rank, which we may do without loss of generality by identifying $\mathrm{span}_{\mathbb{R}}(\mathcal{L})$ with $\mathbb{R}^m$.

We write $\lambda_1(\mathcal{L}) := \min_{\boldsymbol{y} \in \mathcal{L} \setminus \{\boldsymbol{0}\}} \|\boldsymbol{y}\|$ for a length of the shortest non-zero vector in $\mathcal{L}$.

The *dual lattice* $\mathcal{L}^*$ is the set of vectors in the span of $\mathcal{L}$ whose inner product with all lattice vectors is integral,

$$\mathcal{L}^* := \{\boldsymbol{w} \in \mathrm{span}_{\mathbb{R}}(\mathcal{L}) \ : \ \forall \boldsymbol{y} \in \mathcal{L}, \ \langle \boldsymbol{w}, \boldsymbol{y} \rangle \in \mathbb{Z}\} \ .$$

The dual has as a basis $\mathbf{B}(\mathbf{B}^T \mathbf{B})^{-1}$ for any basis $\mathbf{B}$ of $\mathcal{L}$, and in particular, $(\mathcal{L}^*)^* = \mathcal{L}$ and $\det(\mathcal{L}^*) = 1/\det(\mathcal{L})$. We also have the identity $\pi_W(\mathcal{L})^* = W \cap \mathcal{L}^*$ for any subspace $W \subset \mathbb{R}^n$, provided that $\pi_W(\mathcal{L})$ is a lattice. (Equivalently, this holds for any subspace W that is spanned by dual lattice vectors.)

For a full-rank lattice $\mathcal{L} \subset \mathbb{R}^d$ and $1 \le m \le d$, we write

$$\delta_m(\mathcal{L}) := \frac{1}{\det(\mathcal{L})^{1/d}} \min_{\mathcal{L}' \subseteq \mathcal{L}} \det(\mathcal{L}')^{1/m} \ ,$$

where the minimum is over sublattices with rank m. (The minimum is achieved because the lattice is discrete.) Then, Rankin's constants are defined as

$$\delta_{d,m} := \sup_{\mathcal{L} \subset \mathbb{R}^d} \delta_m(\mathcal{L}) \ ,$$

where the supremum is over full-rank lattices.

For example, in the case $m = 1$, Rankin's constant $\delta_{d,1}$ is just Hermite's constant, i.e., the maximal value of $\lambda_1(\mathcal{L})$ for determinant-one lattices. For convenience, we write $\delta_d := \delta_{d,1}$ for Hermite's constant.

Minkowski's celebrated theorem shows us that $\delta_d \le \sqrt{2d/(\pi e)}$, and this is known to be tight up to a small constant factor. For larger m, Rankin's constant is not known as precisely, but we do know the following.

Theorem 2.1. *For any $n, \beta \geq 1$*

$$C(\beta n)^{(1-1/\beta)/2} \leq \delta_{\beta n,n} \leq (\beta n)^{(\beta-1)/2 \cdot \log((\beta+1/n)/(\beta-1+1/n))} \leq \sqrt{\beta n}$$

for some universal constant $C > 0$.

The lower bound is from [SW14], and the upper bound is from [HS12, Lemmas 4.2 and 4.3]. We typically think of β as a small constant.

2.2 Number fields and (linear) embeddings

A number field K is an algebraic extension of the rational numbers $\mathbb{Q}$. I.e., $K = \mathbb{Q}[x]/p(x)$ for some irreducible polynomial $p(x) \in \mathbb{Q}[x]$. The degree n of the number field is simply the degree of the polynomial p. In particular, a degree-n number field is isomorphic as a $\mathbb{Q}$-vector space to $\mathbb{Q}^n$. (To see this, notice that the elements $1, x, x^2, \ldots, x^{n-1} \in K$ form a $\mathbb{Q}$-basis for K.)

For our purposes, an embedding of a number field K with degree n into $\mathbb{R}^n$ is simply an injective linear map $f : K \to \mathbb{R}^n$ such that $\mathrm{span}_{\mathbb{R}}(f(K)) = \mathbb{R}^n$. (These are *linear* embeddings, not to be confused with *field* embeddings, which must map a field to a field.) E.g., the coefficient embedding simply maps the polynomial $y := y_0 + y_1 x + \cdots + y_{n-1}x^{n-1} \in K = \mathbb{Q}[x]/p(x)$ to the vector $(y_0, y_1, \ldots, y_{n-1})$.

We always assume that a number field comes equipped with an embedding, and since the embedding is injective by assumption, we do not bother to write the map f explicitly. Instead, we simply make no distinction between a field element $y \in K$ and its unique image $f(y)$ under the embedding. E.g., we write $\mathrm{span}_{\mathbb{R}}(K) := \mathrm{span}_{\mathbb{R}}(f(K))$; $y = (y_1, \ldots, y_n) := (f(y)_1, \ldots, f(y)_n) \in R^n$; $\|y\| := \|f(y)\|$, where $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^n$; etc. Notice in particular that this notion of length extends to vectors $\boldsymbol{y} := (y_1, \ldots, y_k) \in K^k$ as

$$\|\boldsymbol{y}\| := (\|y_1\| + \cdots + \|y_k\|)^{1/2} .$$

2.3 Rings, ideals, and module lattices

A discrete, full-rank subring $R \subset K$ is a subring of K such that (1) R is finitely generated; (2) any element $x \in K$ can be represented as $x = r/q$ for some $q \in \mathbb{Z} \setminus \{0\}$ and $r \in R$. Throughout this paper, we always assume that our ring has these properties.

A (fractional) *ideal* $\mathcal{I}$ of R is the R-span of finitely many elements $y_1, \ldots, y_m \in K$,

$$\mathcal{I} := \{r_1 y_1 + \cdots + r_m y_m \ : \ r_i \in R\} .$$

More generally, a module lattice $\mathcal{M}$ over R is the R-spans of finitely many vectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in K^k$

$$\mathcal{M} := \{r_1 \boldsymbol{y}_1 + \cdots + r_m \boldsymbol{y}_m \ : \ r_i \in R\} .$$

The *rank* (over K) of a module lattice is the dimension (over K) of its span (over K), $\mathrm{rank}_K(\mathcal{M}) := \dim_K(\mathrm{span}_K(\mathcal{M}))$. We abuse language a bit and sometimes refer to rank-one module lattices as ideals since (as we will see below), rank-one module lattices are isomorphic to ideals. We say that such an ideal is *principal* if it is the R-span of a single element $\boldsymbol{x} \in K^k$, and we say that $\boldsymbol{x}$ *generates* the ideal.

As the name suggests, module lattices are themselves lattices (when viewed as subsets of $\mathbb{R}^{kn}$). To see this, it suffices to take a $\mathbb{Z}$-basis $r_1, \ldots, r_n$ of R and to observe that $\mathcal{M}$ is the $\mathbb{Z}$-span of $r_i \boldsymbol{y}_j$.[3] This in particular means that it makes sense to talk about, e.g., $\det(\mathcal{M})$, $\lambda_1(\mathcal{M})$, $\mathcal{M}^*$, $\mathrm{rank}_{\mathbb{R}}(\mathcal{M})$, etc.

Furthermore, we have $\mathrm{rank}_{\mathbb{R}}(\mathcal{M}) = n \cdot \mathrm{rank}_K(\mathcal{M})$. To see this, it suffices to notice that for any $S \subseteq K^k$, $\dim_{\mathbb{R}} \mathrm{span}_{\mathbb{R}}(\{r\boldsymbol{y} \ : \ r \in R, \boldsymbol{y} \in S\}) = n \cdot \dim_K \mathrm{span}_K(S)$.

We say that a ring $R \subset K$ (with some embedding) is *dual-closed* if for every module lattice $\mathcal{M} \subset K^k$, the dual lattice $\mathcal{M}^*$ is also a module lattice.

Notice that for any K-linear transformation $T : K^k \to K^{k'}$ and any module lattice $\mathcal{M} \subset K^k$, $T(\mathcal{M})$ is also a module lattice. (This is not true for $\mathbb{R}$-linear transformations.)

For example, if $T : K^k \to K^{k'}$ is an isometry on $\mathrm{span}_K(\mathcal{M})$ with $k' := \mathrm{rank}_K(\mathcal{M})$, then T allows us to identify $\mathrm{span}_K(\mathcal{M})$ with $K^{\mathrm{rank}_K(\mathcal{M})}$, and therefore to assume without loss of generality that $\mathcal{M}$ is full rank. For example, if $\mathcal{M}$ has rank one, then we may assume without loss of generality that it lies in K, i.e., that it is a (fractional) ideal. (This justifies our use of the term ideal to refer to all rank-one module lattices.)

More generally, for any K-subspace $V \subseteq K^k$ and any module lattice $\mathcal{M} \subset K^k$, $\mathcal{M} \cap V$ is a module lattice (possibly $\{\mathbf{0}\}$).

Another important example is when $T := \pi_{(\mathcal{M}')^\perp}$ is the projection map onto the space orthogonal to some submodule lattice $\mathcal{M}' \subseteq \mathcal{M}$, in which case we have $\det(\mathcal{M}) = \det(\mathcal{M}') \det(\pi_{(\mathcal{M}')^\perp}(\mathcal{M}))$.

2.4 Some geometric quantities of rings and modules

For a ring R in a particular embedding, we define

$$\alpha_R := \min_{\mathcal{I} \subset R} \frac{\lambda_1(\mathcal{I})}{\det(\mathcal{I})^{1/n}} \ ,$$

where the minimum is over all ideals $\mathcal{I} \subset R$. (Notice that α_R depends heavily on the choice of embedding, so perhaps formally we should write $\alpha_{R,f}$, where f is the embedding. We write α_R instead for simplicity.) The minimum is in fact achieved. Notice that α_R also bounds $\lambda_1(\mathcal{M}) / \det(\mathcal{M})^{1/n}$ for any rank-one module lattice $\mathcal{M} \subset K^k$ (since these are just ideals in disguise).

For a module lattice $\mathcal{M}$, we define

$$\tau_1(\mathcal{M}) := \min_{\mathcal{I} \subset \mathcal{M}} \det(\mathcal{I})^{1/n} \ ,$$

where the minimum is over the rank-one submodules $\mathcal{I} \subset \mathcal{M}$ (i.e., ideals). (Again, the minimum is achieved because $\mathcal{M}$ is discrete.) This quantity can be viewed as a different way to generalize $\lambda_1(\mathcal{L})$ to module lattices over arbitrary rings. I.e., the rank-one "submodules" of a "module" $\mathcal{L}$ over $\mathbb{Z}$ are lattices spanned by a single vector, and the determinant of such a "submodule" is just the length of this vector (which is unique up to sign). So, over $\mathbb{Z}$, $\tau_1 = \lambda_1$. For higher-dimensional rings R, the rank-one module lattices are n-dimensional lattices, which do not naturally correspond to a single vector. So, τ_1 and λ_1 are distinct quantities.

We do, however, have the simple inequality $\tau_1(\mathcal{M}) \leq \delta_{kn,n} \det(\mathcal{M})^{1/(kn)}$, and the following relationship between τ_1 and λ_1, which is governed by α_R.

[3] Since $\mathrm{span}_{\mathbb{R}}(K) = \mathbb{R}^n$ and K is isomorphic as a vector space to $\mathbb{Q}^n$, any $\mathbb{R}$-linear independence in K^k is also a $\mathbb{Q}$-linear dependence. It follows that the $\mathbb{Z}$-span of any vectors in K^k is a lattice, and in particular, $\mathcal{M}$ is a lattice.

Lemma 2.2. *For a module lattice $\mathcal{M}$,*

$$\frac{\lambda_1(\mathcal{M})}{\delta_n} \leq \tau_1(\mathcal{M}) \leq \frac{\lambda_1(\mathcal{M})}{\alpha_R} .$$

Proof. Let $\mathcal{I} \subset \mathcal{M}$ be the ideal generated by a non-zero shortest vector in $\mathcal{M}$. I.e., $\lambda_1(\mathcal{I}) = \lambda_1(\mathcal{M})$. Then from the definition of α_R, we know

$$\det(\mathcal{I})^{1/n} \leq \frac{\lambda_1(\mathcal{I})}{\alpha_R} . \tag{2}$$

Since $\mathcal{I} \subset \mathcal{M}$, we also have that

$$\tau_1(\mathcal{M}) \leq \det(\mathcal{I})^{1/n} . \tag{3}$$

Combining Eqs. (2) and (3) yields the upper bound.

For the lower bound, let $\mathcal{I}' \subset \mathcal{M}$ be an ideal satisfying $\det(\mathcal{I}')^{1/n} = \tau_1(\mathcal{M})$. Then by the definition of Hermite's constant, we have

$$\lambda_1(\mathcal{I}') \leq \delta_n \det(\mathcal{I}')^{1/n} = \delta_n \tau_1(\mathcal{M}) .$$

The lower bound follows by noting that $\lambda_1(\mathcal{M}) \leq \lambda_1(\mathcal{I}')$. □

2.5 The canonical embedding

The *canonical embedding* of a number field $K := \mathbb{Q}[x]/p(x)$ is most naturally viewed as an embedding into $\mathbb{C}^n$. (Since the resulting vector space is only n-dimensional over $\mathbb{R}$, this embedding is isometric to a real embedding that is slightly more cumbersome to define, known as the Minkowski embedding. In the sequel, we make no distinction between the two, since they are isometric.) Up to a reordering of the coordinates, it is the unique embedding such that field multiplication between two elements $y = (y_1, \ldots, y_n) \in K \subset \mathbb{C}^n$ and $y' = (y'_1, \ldots, y'_n) \in K \subset \mathbb{C}^n$ is coordinate-wise, i.e., $y \cdot y' = (y_1 y'_1, y_2 y'_2, \ldots, y_n y'_n)$. Equivalently, the embedding $f(y)$ of y is $f(y) = (\sigma_1(y), \ldots, \sigma_n(y)) \in \mathbb{C}^n$, where the σ_i are the n distinct field embeddings of K into $\mathbb{C}$. Alternatively, if we view $y := y(x) \in \mathbb{Q}[x]/p(x)$ as a polynomial, then the σ_i correspond to polynomial evaluation at the n distinct roots in $\mathbb{C}$ of the defining polynomial p of K.

We will not discuss the canonical embedding explicitly in the sequel, but it is a very useful and important example.

Lemma 2.3. *For any number field K, any (discrete, full-rank) subring $R \subset K$ that is closed under complex conjugation is dual-closed in the canonical embedding.*

Proof. Since ring elements act diagonally in the canonical embedding, we have that for any $\boldsymbol{x} \in \mathcal{M}$, $\boldsymbol{w} \in \mathcal{M}^*$, and $r \in R$, $\langle \boldsymbol{x}, r\boldsymbol{w} \rangle = \langle \overline{r}\boldsymbol{x}, \boldsymbol{w} \rangle \in \mathbb{Z}$, where $\overline{r}$ represents the complex conjugate. So, if $\overline{r} \in R$, then $\overline{r}\boldsymbol{x} \in \mathcal{M}$ and therefore $\langle \boldsymbol{x}, r\boldsymbol{w} \rangle \in \mathbb{Z}$. □

We also have the following well-known property of the canonical embedding. (A lower bound on α_R follows by considering the algebraic norm of the shortest non-zero vector in the ideal and applying the inequality between arithmetic and geometric means. The matching upper bound is witnessed by R itself, which must contain the element $1 \in K$, which has $\|1\| = \sqrt{n}$ in the canonical embedding.)

Lemma 2.4. *For any (discrete, full-rank) subring $R \subset K$ of any number field K under the canonical embedding, we have*

$$\alpha_R = \frac{\sqrt{n}}{\det(R)^{1/n}} .$$

In particular, if $R := \mathcal{O}_K$ is the ring of integers of a cyclotomic number field K, then $\det(R)^{1/n} \leq \sqrt{n}$, so that $\alpha_R \geq 1$.

2.6 ModuleSVP, the Dense Ideal Problem, and their duals

We now provide the formal definition of ModuleSVP, and its variant the Dense Ideal Problem. We also define certain dual versions, which we will need in the sequel.

Definition 2.5 (ModuleSVP and its dual)**.** *For a number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 1$, and approximation factor $\gamma = \gamma(n,k) \geq 1$, (γ,k)-*ModuleSVP *is defined as follows. The input is (a generating set for) a module lattice $\mathcal{M} \subset K^k$ of rank k. The goal is to output a module element $\boldsymbol{x} \in \mathcal{M}$ such that $0 < \|\boldsymbol{x}\| \leq \gamma\lambda_1(\mathcal{M})$.*

Dual (γ,k)-ModuleSVP *has the same input, and the goal is to output a dual module element $\boldsymbol{x} \in \mathcal{M}^*$ such that $0 < \|\boldsymbol{x}\| \leq \gamma\lambda_1(\mathcal{M}^*)$.*

(γ,k)-ModuleSVP' is the union of these two problems. I.e., the input is a module and a string $s \in \{PRIMAL,\ DUAL\}$. If $s = PRIMAL$, the goal is to output a valid solution to (γ,k)-ModuleSVP. If $s = DUAL$, the goal is to output a valid solution to dual (γ,k)-ModuleSVP.

Definition 2.6 ((γ,k)-DIP and its dual)**.** *For a number field K, ring $R \subset K$, rank $k \geq 2$, and approximation factor $\gamma \geq 1$, the* (γ,k)-Dense Ideal Problem*, or (γ,k)-DIP, is the search problem defined as follows. The input is a module lattice $\mathcal{M} \subset K^k$ with rank k, and the goal is to find a submodule $\mathcal{M}' \subset \mathcal{M}$ with rank one (i.e., an ideal lattice) such that $\det(\mathcal{M}')^{1/n} \leq \gamma\tau_1(\mathcal{M})$.*

Dual (γ,k)-DIP *has the same input, and the goal is to output a submodule $\mathcal{M}' \subset \mathcal{M}^*$ of the dual with rank one (i.e., an ideal lattice) such that $\det(\mathcal{M}')^{1/n} \leq \gamma\tau_1(\mathcal{M}^*)$.*

(γ,k)-DIP' is the union of these two problems as above.

Notice that, if the ring R is dual-closed, then ModuleSVP, dual ModuleSVP, and ModuleSVP' are equivalent problems, as are DIP, dual DIP, and DIP'.

Furthermore, $(\gamma,2)$-DIP, dual $(\gamma,2)$-DIP, and $(\gamma,2)$-DIP' are equivalent. To see this, let $\mathcal{M}_1 \subset \mathcal{M}$ be a primitive rank-one submodule. Of course $\mathcal{M}_1$ defines a filtration, $\mathcal{M}_1 \subset \mathcal{M}_2 = \mathcal{M}$. Then, $\det(\mathcal{M}) = \det(\mathcal{M}_1)\det(\widetilde{\mathcal{M}}_2)$. The corresponding dual filtration is $\widetilde{\mathcal{M}}_2^* \subset \mathcal{M}^*$, and we have $\det(\widetilde{\mathcal{M}}_2^*) = 1/\det(\widetilde{\mathcal{M}}_2) = \det(\mathcal{M}_1)/\det(\mathcal{M})$. Therefore, (1) $\tau_1(\mathcal{M}^*) = \tau_1(\mathcal{M})/\det(\mathcal{M})^{1/n}$; and (2) $\mathcal{M}_1$ is a solution to $(\gamma,2)$-DIP on $\mathcal{M}$ if and only if $\widetilde{\mathcal{M}}_2^*$ is a solution to dual $(\gamma,2)$-DIP on $\mathcal{M}^*$.

Definition 2.7 ((γ,k)-ModuleHSVP)**.** *For a number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 1$, and approximation factor $\gamma = \gamma(n,k) \geq 1$, (γ,k)-*ModuleHSVP *is defined as follows. The input is (a generating set for) a module lattice $\mathcal{M} \subset K^k$ of rank k. The goal is to output a module element $\boldsymbol{x} \in \mathcal{M}$ such that $0 < \|\boldsymbol{x}\| \leq \gamma\det(\mathcal{M})^{1/(kn)}$.*

Theorem 2.8. *For any number field K, (discrete, full-rank) subring $R \subset K$, rank $\beta \geq 2$, and approximation factor $\gamma' \geq 1$, there exists a reduction from (γ,β)-DIP to (γ',β)-ModuleSVP where $\gamma := \frac{\gamma'\delta_n}{\alpha_R}$.*

Proof. The reduction takes as input a module $\mathcal{M}$ of rank β, and uses the output from the (γ', β)-ModuleSVP oracle which is a non-zero short vector $\boldsymbol{x}$ such that $0 < \|\boldsymbol{x}\| \leq \gamma' \lambda_1(\mathcal{M})$, to output a submodule $\mathcal{M}' \subset \mathcal{M}$ such that $\det(\mathcal{M}')^{1/n} \leq \gamma \tau_1(\mathcal{M})$.

Let $\mathcal{M}' := (\boldsymbol{x})$, i.e. $\mathcal{M}'$ is a principal ideal generated by $\boldsymbol{x}$. Note that $\lambda_1(\mathcal{M}') \leq \|\boldsymbol{x}\| \leq \gamma' \lambda_1(\mathcal{M})$. Then using Lemma 2.2, we have

$$\det(\mathcal{M}')^{1/n} \leq \frac{\lambda_1(\mathcal{M}')}{\alpha_R} \leq \frac{\gamma' \lambda_1(\mathcal{M})}{\alpha_R} \leq \frac{\gamma'}{\alpha_R} \cdot \delta_n \cdot \tau_1(\mathcal{M}) \ ,$$

as needed. □

2.7 On bit representations

Throughout this work, we follow the convention (common in the literature on lattices) of avoiding discussion of the particular bit representation of elements in K and their embeddings as much as possible. In practice, one can represent elements in K as polynomials with rational coefficients, and the coordinates in an embedding can be taken to be, say, algebraic numbers represented as roots of integer polynomials or sufficiently good rational approximations. The embedding itself is represented by specifying the embeddings of some basis of K. Since arithmetic operations may be performed efficiently with these representations, we are largely justified in ignoring such bit-level details.

There are two issues that arise, however, and we address them briefly here.

First, there is the question of whether the bit lengths of the numbers that we work with can become superpolynomial after polynomially many operations. This issue is well-studied in the context of basis reduction, and it may be avoided by, for example, ensuring that the basis remains LLL reduced at all times. In particular, in our reductions, we say without further explanation that we update a filtration such that a certain projection $\pi_i(\mathcal{M}_j)$ in the new filtration is equal to a certain module lattice. This operation can always be performed in such a way to keep the bit lengths bounded (under the assumption, valid in our case, that the target module lattice has determinant smaller than the corresponding projection $\pi_i(\mathcal{M}_j)$ in the filtration before the change)—e.g., by ensuring that the underlying $\mathbb{Z}$-basis is LLL reduced. We refer the reader to [GN08] for a more careful analysis in the context of slide reduction and [LPSW19] for discussion of similar issues in the context of module lattices. With this carefully swept under the rug, we content ourselves in the sequel with simply bounding the number of such operations performed by our reductions.

Second, we will actually need a minor relationship between the bit length of the representation of the embedding and the geometry of the ring R. To see why this is necessary, imagine that we could have a ring R such that $\lambda_1(R) < 2^{-m^{\omega(1)}}$, where m is the bit length of the description of R. Then, we could not even write down $\lambda_1(R)$ in polynomial time. Of course, this cannot happen for reasonable representations.

Fact 2.9. *If the number field K, its embedding, and the ring R are represented as described above, then for any integer $k \geq 1$ and any module $\mathcal{M} \subset K^k$*

$$2^{-\operatorname{poly}(m,k)} \leq \det(\mathcal{M}) \leq 2^{\operatorname{poly}(m,k)} \ ,$$

where m is the bit length of this description together with the description of a generating set for $\mathcal{M}$.

3 Filtrations

For a module lattice $\mathcal{M} \subset K^\ell$ over R with rank k, a *filtration* of $\mathcal{M}$ is a nested sequence $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of module lattices over R such that

1. **Primitivity:** $\mathcal{M}_i = \mathcal{M} \cap \mathrm{span}_K(\mathcal{M}_i)$;
2. **Increasing ranks:** $\mathrm{rank}_K(\mathcal{M}_i) = i$; and
3. **Rank-one projections:** $\widetilde{\mathcal{M}}_i := \pi_{\mathcal{M}_{i-1}^\perp}(\mathcal{M}_i)$ is a rank one module lattice over R (i.e., an ideal).

(In fact, primitivity together with the fact that $\mathcal{M}_i \subset \mathcal{M}_{i+1}$ is a strict containment already implies the other two conditions. E.g., this implies that $\mathrm{rank}_K(\mathcal{M}_i) < \mathrm{rank}_K(\mathcal{M}_{i+1})$, and since the ranks are positive integers with $\mathrm{rank}_K(\mathcal{M}_k) = k$, we must have $\mathrm{rank}_K(\mathcal{M}_i) = i$. Nevertheless, we find it helpful to state the other two conditions explicitly.) We also write $\mathcal{M}_{[i,j]} := \pi_{\mathcal{M}_{i-1}^\perp}(\mathcal{M}_j)$.

Filtrations for module lattices over R are analogues of bases for lattices over $\mathbb{Z}$. Specifically, the basis $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d \in \mathbb{R}^d$ of a lattice corresponds to the filtration given by $\mathcal{L}_i := \{z_1\boldsymbol{b}_1 + \cdots + z_i\boldsymbol{b}_i \,:\, z_j \in \mathbb{Z}\}$. The $\widetilde{\mathcal{M}}_i$ defined above are the analogues of the Gram-Schmidt orthogonalization $\widetilde{\boldsymbol{b}}_1, \ldots, \widetilde{\boldsymbol{b}}_d$ of a lattice over $\mathbb{Q}$. We therefore call $\widetilde{\mathcal{M}}_i$ an *R-Gram-Schmidt orthogonalization*.

Fact 3.1. *For every module lattice $\mathcal{M} \subset K^k$ over $R \subset K$ with rank k, there exists a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$.*

Furthermore, the $\mathcal{M}_i$ can be computed efficiently (given an R-basis for $\mathcal{M}$), and $\det(\mathcal{M}) = \det(\widetilde{\mathcal{M}}_1) \cdots \det(\widetilde{\mathcal{M}}_k)$.

Proof. Let $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_\ell$ be an R-basis for $\mathcal{M}$, and suppose without loss of generality that $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_k$ are linearly independent over K. We take $\mathcal{M}_i := \mathcal{M} \cap \mathrm{span}_K(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i)$.

The fact about the determinants follows from the analogous fact for general lattices. Specifically, for any lattice $\mathcal{L} \subset \mathbb{R}^d$ and any sequence of primitive sublattices $\mathcal{L}_1 \subset \mathcal{L}_2 \subset \cdots \subset \mathcal{L}_k = \mathcal{L}$, $\det(\mathcal{L}) = \prod_i \det(\mathcal{L}_i)/\det(\mathcal{L}_{i-1}) = \prod_i \det(\pi_{\mathcal{L}_{i-1}^\perp}(\mathcal{L}))$ (where we have used the convention $\mathcal{L}_0 := \{\mathbf{0}\}$ and $\det(\mathcal{L}_0) := 1$). □

Finally, each filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$ induces a *dual filtration* given by $\pi_{\mathcal{M}_{k-1}^\perp}(\mathcal{M})^* \subset \pi_{\mathcal{M}_{k-2}^\perp}(\mathcal{M})^* \subset \cdots \subset \pi_{\mathcal{M}_1^\perp}(\mathcal{M})^* \subset \mathcal{M}^*$, where the $\mathcal{M}_i$ are module lattices over the dual ring R'. Equivalently, the dual filtration is given by $(\mathcal{M}^* \cap \mathcal{M}_{k-1}^\perp) \subset (\mathcal{M}^* \cap \mathcal{M}_{k-2}^\perp) \subset \cdots \subset (\mathcal{M}^* \cap \mathcal{M}_1^\perp) \subset \mathcal{M}^*$. While we will not use the dual filtration explicitly, it is helpful to keep it in mind. In particular, the R'-Gram-Schmidt orthogonalization of the dual filtration is the reverse of the R-Gram-Schmidt orthogonalization of the original filtration, in analogy to the reversed dual basis $\mathbf{B}^{-s}$ that is commonly used in basis reduction. (See, e.g., [GN08, MW16].)

4 An LLL-style algorithm for the special case of $\beta = 2$

Here, we present our reductions in the special case when $\beta = 2$. The results here are strictly generalized by and subsumed by those in Section 5, and the proofs have many common features. (Our proofs are also essentially the same as those in [LPSW19].) However, the case $\beta = 2$ is considerably simpler, and we therefore include a separate section for this case. (To make comparison

easier, we have given this section and Section 5 identical structures. E.g., plugging $\beta = 2$ into Lemma 5.4 yields Lemma 4.4, and the same is true for, e.g., Theorems 5.10 and 4.10.)

In particular, we do not need to mention the dual module lattice at all in this section. We can instead use the following simple notion of a reduced filtration.

Definition 4.1 (DIP reduction)**.** *For $\gamma \geq 1$, a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ is γ-DIP-reduced if* $\det(\mathcal{M}_1)^{1/n} \leq \gamma \cdot \tau_1(\mathcal{M})$.

Definition 4.2 (γ-reduced filtration)**.** *For $\gamma \geq 1$, we say that a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is* γ-reduced *if $\mathcal{M}_{[i,i+1]}$ is γ-DIP-reduced for all $i \in [1, k-1]$.*

We now show a number of properties of γ-reduced filtrations that make them useful for solving ModuleSVP and its variants.

Lemma 4.3. *For any γ-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$, we have* $\det(\mathcal{M}_1)^{1/n} \leq (\gamma\delta_{2n,n})^{2(i-1)} \det(\widetilde{\mathcal{M}}_i)^{1/n}$ *for all $1 \leq i \leq k$.*

Proof. Since $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is γ-reduced,

$$\begin{aligned}\det(\widetilde{\mathcal{M}}_i)^{1/n} &\leq \gamma \cdot \tau_1(\mathcal{M}_{[i,i+1]}) \\ &\leq \gamma \cdot \delta_{2n,n} \cdot \det(\mathcal{M}_{[i,i+1]})^{1/(2n)} \\ &= \gamma \cdot \delta_{2n,n} \cdot \left(\det(\widetilde{\mathcal{M}}_i)\det(\widetilde{\mathcal{M}}_{i-1})\right)^{1/(2n)} .\end{aligned}$$

Rearranging, we see that $\det(\widetilde{\mathcal{M}}_i)^{1/n} \leq (\gamma\delta_{2n,n})^2 \det(\widetilde{\mathcal{M}}_{i+1})^{1/n}$. By a simple induction argument, we see that $\det(\mathcal{M}_1)^{1/n} \leq (\gamma\delta_{2n,n})^{2(i-1)} \det(\widetilde{\mathcal{M}}_i)^{1/n}$. □

Lemma 4.4. *If a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is γ-reduced, then*

$$\det(\mathcal{M}_1)^{1/n} \leq \gamma \cdot (\gamma\delta_{2n,n})^{2(k-2)} \cdot \tau_1(\mathcal{M}) , \text{ and} \tag{4}$$

$$\det(\mathcal{M}_1)^{1/n} \leq (\gamma\delta_{2n,n})^{(k-1)} \cdot \det(\mathcal{M})^{1/kn} . \tag{5}$$

Proof. Let $i \in [1, k-1]$ be such that $\tau_1(\mathcal{M}_{i+1}) = \tau_1(\mathcal{M})$ but $\tau_1(\mathcal{M}_{i-1}) \neq \tau_1(\mathcal{M})$ (where we use the convention that $\tau_1(\mathcal{M}_0) \neq \tau_1(\mathcal{M})$). Since $\mathcal{M}_k = \mathcal{M}$, there must exist such an i. In particular, there exists some rank-one module lattice $\mathcal{M}' \subset \mathcal{M}_{i+1}$ with $\mathcal{M}' \not\subset \mathcal{M}_{i-1}$ such that $\det(\mathcal{M}')^{1/n} = \tau_1(\mathcal{M})$. Therefore (since the $\mathcal{M}_i$ are primitive), $\pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}') \subset \mathcal{M}_{[i,i+1]}$ is a non-zero rank-one module lattice. It follows that $\tau_1(\mathcal{M}_{[i,i+1]}) \leq \det(\pi_{\mathcal{M}_{i-1}^{\perp}}(\mathcal{M}'))^{1/n} \leq \det(\mathcal{M}')^{1/n} = \tau_1(\mathcal{M})$. Then, by the γ-reduced property of the filtration,

$$\det(\widetilde{\mathcal{M}}_i)^{1/n} \leq \gamma\tau_1(\mathcal{M}_{[i,i+1]}) \leq \gamma\tau_1(\mathcal{M}) .$$

By combining the expression above with Lemma 4.3, we have

$$\det(\mathcal{M}_1)^{1/n} \leq \gamma \cdot (\gamma\delta_{2n,n})^{2(i-1)} \cdot \tau_1(\mathcal{M}) , \tag{6}$$

and recalling that $i \leq k-1$, we obtain Eq. (4).

By taking the product of the expression obtained from Lemma 4.3 over $1 \leq i \leq k$, we see that

$$\det(\mathcal{M}_1)^{k/n} \leq (\gamma\delta_{2n,n})^{k(k-1)} \det(\mathcal{M})^{1/n} .$$

Raising both sides to the power $1/k$ yields Eq. (5). □

Corollary 4.5. *If a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is γ-reduced, then*

$$\lambda_1(\mathcal{M}_1) \le \frac{\gamma\delta_n}{\alpha_R} \cdot (\gamma\delta_{2n,n})^{2(k-2)} \cdot \lambda_1(\mathcal{M}) \ , \ \text{and} \tag{7}$$

$$\lambda_1(\mathcal{M}_1) \le \delta_n(\gamma\delta_{2n,n})^{(k-1)} \cdot \det(\mathcal{M})^{1/kn} \ . \tag{8}$$

Proof. By combining Eq (4) from Lemma 4.4 with Lemma 2.2, we have

$$\det(\mathcal{M}_1)^{1/n} \le \gamma \cdot (\gamma\delta_{2n,n})^{2(k-2)} \cdot \tau_1(\mathcal{M}) \le \gamma \cdot (\gamma\delta_{2n,n})^{2(k-2)} \cdot \frac{\lambda_1(\mathcal{M})}{\alpha_R} \ .$$

Using the definition of Hermite's constant δ_n, with the above relation, we obtain Eq. (7):

$$\lambda_1(\mathcal{M}_1) \le \delta_n \det(\mathcal{M}_1)^{1/n} \le \delta_n \cdot \gamma(\gamma\delta_{2n,n})^{2(k-2)} \cdot \frac{\lambda_1(\mathcal{M})}{\alpha_R} \ .$$

Eq. (8) follows by directly applying the definition of Hermite's constant to Eq. (5) from Lemma 4.4. □

4.1 Finding γ-reduced filtrations

We are now ready to show how to find a γ-reduced filtration with access to a $(\gamma, 2)$-ModuleSVP oracle. The reduction is a natural analogue of the LLL algorithm, and essentially identical to the reduction in [LPSW19].

Definition 4.6 ((γ, k)-RFP)**.** *For a number field K, (discrete, full-rank) ring $R \subset K$, rank $k \ge 2$, and approximation factor $\gamma \ge 1$, the (γ, k)-Reduced Filtration Problem, or (γ, k)-RFP, is the search problem defined as follows. The input is a module lattice $\mathcal{M} \subset K^k$ with rank k, and the goal is to find a γ-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$.*

Theorem 4.7. *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \ge 2$, approximation factor $\gamma \ge 1$, and constant $\varepsilon > 0$, there is an efficient reduction from $((1+\varepsilon)\gamma, k)$-RFP to $(\gamma, 2)$-DIP.*

Proof. The idea is to use our $(\gamma, 2)$-DIP oracle to compute a $(1+\varepsilon)\gamma$-reduced filtration just like the LLL algorithm computes a reduced basis. In particular, on input (a generating set for) a module lattice $\mathcal{M} \subset K^k$ with rank k, the reduction first computes a filtration $\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$. It then repeatedly updates this filtration in places as follows.

For each $\mathcal{M}_{[i,i+1]}$, the reduction calls the $(\gamma, 2)$-DIP oracle with $\mathcal{M}_{[i,i+1]}$ as input and receives as output some rank-one ideal $\widetilde{\mathcal{M}}'_i \subset \mathcal{M}_{[i,i+1]}$. We may assume without loss of generality that $\widetilde{\mathcal{M}}'_i$ is a primitive sublattice of $\mathcal{M}_{[i,i+1]}$, i.e., that $\widetilde{\mathcal{M}}'_i = \mathcal{M}_{[i,i+1]} \cap \mathrm{span}_K(\widetilde{\mathcal{M}}'_i)$. If $(1+\varepsilon)^n \det(\widetilde{\mathcal{M}}'_i) < \det(\widetilde{\mathcal{M}}_i)$ then the reduction sets $\mathcal{M}_i$ so that $\widetilde{\mathcal{M}}_i = \widetilde{\mathcal{M}}'_i$. (Formally, the reduction can do this by, e.g., picking any i-dimensional K-subspace W of $\mathrm{span}_K(\mathcal{M}_{i+1})$ such that $\pi_{\mathcal{M}_{i-1}^{\perp}}(W) = \mathrm{span}_K(\widetilde{\mathcal{M}}'_i)$ and setting $\mathcal{M}_i := W \cap \mathcal{M}$.)

The reduction terminates and outputs the current filtration when none of these checks results in an update to the filtration, i.e., when for all i, $(1+\varepsilon)^n \det(\widetilde{\mathcal{M}}'_i) \ge \det(\widetilde{\mathcal{M}}_i)$.

We first observe that the output filtration is indeed $(1+\varepsilon)\gamma$-reduced. To see this, notice that the reduction only terminates if the filtration satisfies

$$\det(\widetilde{\mathcal{M}}_i)^{1/n} \le (1+\varepsilon)\det(\widetilde{\mathcal{M}}'_i)^{1/n} \le (1+\varepsilon)\gamma \cdot \tau_1(\mathcal{M}_{[i,i+1]}) \ ,$$

as needed.

It remains to show that the reduction terminates in polynomial time. Our proof is more-or-less identical to the celebrated proof in [LLL82] (and the proof in [LPSW19]). Consider the potential function

$$\Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k) := \prod_{i=1}^{k} \det(\mathcal{M}_i) \ .$$

At the beginning of the reduction, $\log \Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k)$ is bounded by a polynomial in the input size (since Φ is efficiently computable). And, by Fact 2.9, $-\log(\Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k))$ is bounded by a polynomial in the input size throughout the reduction. Therefore, it suffices to show that the potential decreases by at least, say, a constant factor every time that the reduction updates the filtration.

Consider a step in the reduction in which it updates $\mathcal{M}_i$. Denote $\widehat{\mathcal{M}}_0$ as $\mathcal{M}_i$ before the update and $\widehat{\mathcal{M}}_1$ as $\mathcal{M}_i$ after the update. Then

$$\det(\widehat{\mathcal{M}}_1) = \det(\mathcal{M}_{i-1}) \det(\widetilde{\mathcal{M}}_i') < \det(\mathcal{M}_{i-1}) \frac{\det(\widetilde{\mathcal{M}}_i)}{(1+\varepsilon)^n} = \frac{\det(\widehat{\mathcal{M}}_0)}{(1+\varepsilon)^n} \ .$$

The other terms $\det(\mathcal{M}_j)$ for $i \neq j$ in the definition of Φ remain unchanged. Thus, the potential function decreases by a factor of at least $(1+\varepsilon)^n$ after each update, as needed. □

Finally, we derive the main results of this section as corollaries of Theorem 4.10.

Corollary 4.8. *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, approximation factor $\gamma' \geq 1$, and constant $\varepsilon > 0$, there exists an efficient reduction from (γ, k)-DIP to $(\gamma', 2)$-DIP where*

$$\gamma := (1+\varepsilon)\gamma' \cdot ((1+\varepsilon)\gamma' \cdot \delta_{2n,n})^{2(k-2)} \ .$$

Proof. The reduction takes as input a module lattice $\mathcal{M}$ of rank k and calls the $((1+\varepsilon)\gamma', k)$-RFP oracle using the $(\gamma', 2)$-DIP oracle from Theorem 4.7. By Eq. (4) from Lemma 4.4, $\gamma := (1+\varepsilon)\gamma' \cdot ((1+\varepsilon)\gamma' \cdot \delta_{2n,n})^{2(k-2)}$. This yields as output a rank-one module $\mathcal{M}_1 \subset \mathcal{M}$, such that $\det(\mathcal{M}_1)^{1/n} \leq \gamma \tau_1(\mathcal{M})$. □

Corollary 4.9. *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, approximation factor $\gamma' \geq 1$, and constant $\varepsilon > 0$, there exists an efficient reduction from (γ_R, k)-RFP to $(\gamma', 2)$-ModuleSVP where $\gamma_R := (1+\varepsilon)\frac{\gamma'\delta_n}{\alpha_R}$.*

Proof. The $(\gamma, 2)$-DIP procedure in Theorem 4.7 calls the $(\gamma', 2)$-ModuleSVP oracle from Theorem 2.8, thereby getting a constant of $\gamma := \frac{\gamma'\delta_n}{\alpha(R)}$. From Theorem 4.7, we get a γ_R-reduced filtration, where $\gamma_R = (1+\varepsilon)\gamma$. In other words, $\gamma_R := (1+\varepsilon)\frac{\gamma'\delta_n}{\alpha_R}$. □

Theorem 4.10 (Main Theorem)**.** *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, approximation factor $\gamma' \geq 1$, and constant $\varepsilon > 0$, there is an efficient reduction from (γ, k)-ModuleSVP to $(\gamma', 2)$-ModuleSVP where*

$$\gamma := (1+\varepsilon) \cdot \left(\frac{\gamma'\delta_n}{\alpha_R}\right)^2 \cdot \left((1+\varepsilon)\gamma' \cdot \frac{\delta_n \delta_{2n,n}}{\alpha_R}\right)^{2(k-2)} \ .$$

There is also an efficient reduction from (γ_H, k)-ModuleHSVP to $(\gamma', 2)$-ModuleSVP, where

$$\gamma_H := \gamma'\delta_n \cdot \left((1+\varepsilon)\gamma' \cdot \frac{\delta_n \delta_{2n,n}}{\alpha_R}\right)^{(k-1)} \ .$$

Proof. In fact, the reduction is the same for both ModuleSVP and ModuleHSVP. On input (a generating set for) a module lattice $\mathcal{M} \subset K^k$ with rank k, the reduction proceeds as follows (for both ModuleHSVP and ModuleSVP). It obtains a γ_R-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ using $(\gamma', 2)$-ModuleSVP oracle, where $\gamma_R := (1+\varepsilon)\frac{\gamma'\delta_n}{\alpha_R}$ (by Corollary 4.9). It then calls its $(\gamma', 2)$-ModuleSVP on $\mathcal{M}_2$ which outputs a vector $\boldsymbol{x}$ such that $0 < \|\boldsymbol{x}\| \le \gamma'\lambda_1(\mathcal{M}_2)$. This vector $\boldsymbol{x}$ is the output of our (γ, k)-ModuleSVP.

Since $\mathcal{M}_1 \subset \mathcal{M}_2$, we have

$$0 < \|\boldsymbol{x}\| \le \gamma'\lambda_1(\mathcal{M}_2) \le \gamma'\lambda_1(\mathcal{M}_1) \,.$$

By Eq. (7) of Corollary 4.5,

$$\lambda_1(\mathcal{M}_1) \le \frac{\delta_n \delta_{2n,n}^{2(k-2)} \gamma_R^{(2k-3)}}{\alpha_R} \cdot \lambda_1(\mathcal{M}) = \frac{\delta_n \delta_{2n,n}^{2(k-2)}}{\alpha_R} \cdot \left((1+\varepsilon)\frac{\delta_n\gamma'}{\alpha_R}\right)^{(2k-3)} \cdot \lambda_1(\mathcal{M}) \,.$$

Combining the above two expressions, we get

$$0 < \|\boldsymbol{x}\| \le (1+\varepsilon)^{(2k-3)} \delta_{2n,n}^{2(k-2)} \cdot \left(\frac{\delta_n\gamma'}{\alpha_R}\right)^{2(k-1)} \cdot \lambda_1(\mathcal{M}) \,.$$

Therefore,

$$\gamma = (1+\varepsilon)^{(2k-3)} \delta_{2n,n}^{2(k-2)} \cdot \left(\frac{\delta_n\gamma'}{\alpha_R}\right)^{2(k-1)} \,,$$

as needed.

Similarly, by Eq. (8) of Corollary 4.5,

$$\begin{aligned} \|\boldsymbol{x}\| &\le \gamma'\delta_n \cdot (\gamma_R \delta_{2n,n})^{(k-1)} \cdot \det(\mathcal{M})^{1/kn} \\ &= \gamma'\delta_n \cdot ((1+\varepsilon)\gamma'\delta_n\delta_{2n,n}/\alpha_R)^{(k-1)} \cdot \det(\mathcal{M})^{1/kn} \,, \end{aligned}$$

which gives the reduction from ModuleHSVP. □

5 Slide-reduced filtrations for modules

Throughout this section, p will always denote the number of β-blocks in a filtration of a rank k module, i.e., $k = \beta p$. We also write $\delta := \delta_{\beta n,n} \le \sqrt{\beta n}$ for Rankin's constant.

As in [GN08], we will need a dual notion of DIP-reduced filtrations (in analogy with the notions of SVP-reduced and DSVP-reduced bases in [GN08]), which we will combine together with DIP-reduced filtrations to define our notion of slide reduction. While in [GN08], reduction is defined by comparing lengths of certain vectors to λ_1 of a particular lattice, we compare the *determinants* of certain *ideals* to τ_1 of the analogous module. I.e., our definitions are a high-dimensional analogue of those in [GN08], replacing lengths of vectors with determinants of high-dimensional (ideal) sublattices.

Definition 5.1 (DualDIP reduction)**.** *For $\gamma \ge 1$, a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ is γ-DualDIP-reduced if $\gamma \cdot \det(\widetilde{\mathcal{M}}_k)^{1/n} \ge 1/\tau_1(\mathcal{M}^*)$.*

This is in fact the dual notion of DIP reduction, which we can see by recalling the notion of the dual filtration, as defined in Section 3. We then see that a filtration is DualDIP-reduced if and only if its dual filtration is DIP-reduced, and in particular, a dual DIP oracle is sufficient to obtain a DualDIP-reduced filtration.

We can now "glue" DIP-reduced and DualDIP-reduced filtrations together to obtain a notion of slide-reduced filtration, which is of course a generalization of the notion of a slide-reduced basis from [GN08]. Indeed, once we have the right primitive notions of reduced filtrations, the right generalization of slide-reduced filtrations is clear.

Definition 5.2 ((γ, β)-slide-reduced filtration)**.** *For an approximation factor $\gamma \geq 1$ and an integer block size $\beta \geq 2$, a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ of a rank k module lattice $\mathcal{M}$ where $k = \beta p$ is (γ, β)-slide reduced if it satisfies the following two conditions:*

- ***Primal Conditions.*** *For all $i \in [0; p-1]$, the block $\mathcal{M}_{[i\beta+1, i\beta+\beta]}$ is γ-DIP-reduced.*
- ***Dual Conditions.*** *For all $i \in [0; p-2]$, the block $\mathcal{M}_{[i\beta+2, i\beta+\beta+1]}$ is γ-DualDIP-reduced.*

The following lemma shows how the primal and dual conditions combine to guarantee nice behavior of the R-Gram-Schmidt orthogonalization $\widetilde{\mathcal{M}}_i$.

Lemma 5.3. *If a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is (γ, β)-slide-reduced, then*

$$\det(\mathcal{M}_1)^{1/n} \leq (\gamma\delta)^{2i\beta/(\beta-1)} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} ,$$

for $i \in [1, p-1]$.

Proof. From the primal condition of Definition 5.2,

$$\begin{aligned}
\det(\widetilde{\mathcal{M}}_{i\beta+1})^{\beta/n} &\leq \left(\gamma \cdot \tau_1(\mathcal{M}_{[i\beta+1, i\beta+\beta]})\right)^{\beta} \\
&\leq \left(\gamma\delta \cdot \det(\mathcal{M}_{[i\beta+1, i\beta+\beta]})^{1/(\beta n)}\right)^{\beta} \\
&= (\gamma\delta)^{\beta} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} \det(\mathcal{M}_{[i\beta+2, i\beta+\beta]})^{1/n} .
\end{aligned}$$

Therefore, we have

$$\det(\widetilde{\mathcal{M}}_{i\beta+1})^{(\beta-1)/n} \leq (\gamma\delta)^{\beta} \cdot \det(\mathcal{M}_{[i\beta+2, i\beta+\beta]})^{1/n} . \tag{9}$$

From the dual condition of Definition 5.2,

$$\begin{aligned}
\gamma^{\beta} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+\beta+1})^{\beta/n} &\geq \frac{1}{\tau_1(\mathcal{M}^*_{[i\beta+2, i\beta+\beta+1]})^{\beta}} \\
&\geq \delta^{-\beta} \cdot \det(\mathcal{M}_{[i\beta+2, i\beta+\beta+1]})^{1/n} \\
&= \delta^{-\beta} \cdot \det(\mathcal{M}_{[i\beta+2, i\beta+\beta]})^{1/n} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+\beta+1})^{1/n} .
\end{aligned}$$

Therefore, we have

$$\det(\mathcal{M}_{[i\beta+2, i\beta+\beta]})^{1/n} \leq (\gamma\delta)^{\beta} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+\beta+1})^{(\beta-1)/n} . \tag{10}$$

By combining Eqs. (9) and (10), for $i \in [0, p-2]$,

$$\det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} \leq (\gamma\delta)^{2\beta/(\beta-1)} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+\beta+1})^{1/n} .$$

Then, by a simple induction argument, we see that

$$\det(\mathcal{M}_1)^{1/n} \leq (\gamma\delta)^{2i\beta/(\beta-1)} \cdot \det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} ,$$

for $i \in [0, p-1]$. □

This next lemma and its corollary show why slide-reduced filtrations are useful for solving ModuleSVP, ModuleHSVP, and DIP. In particular, the submodule $\mathcal{M}_1$ of a slide-reduced filtration is guaranteed to have small determinant (as we show in Lemma 5.4) and to contain a short non-zero vector (as we show in Corollary 5.5). Lemma 5.4 is a direct high-dimensional generalization of [GN08, Theorem 1]. Indeed, setting $R = \mathbb{Z}$ and therefore $n = 1$, which in particular implies that $\tau_1 = \lambda_1$, directly recovers [GN08, Theorem 1].

On the other hand, Corollary 5.5 has no obvious analogue over $\mathbb{Z}$. In particular, Eq. (16) of Corollary 5.5 is identical to Eq. (12) of Lemma 5.4 over $\mathbb{Z}$, while the proof of Eq. (15) of Corollary 5.5 relies on the particular geometry of module lattices.

Lemma 5.4. *If a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is (γ, β)-slide-reduced, then*

$$\det(\mathcal{M}_1)^{1/n} \leq \gamma \cdot (\gamma\delta)^{\frac{2(k-\beta)}{\beta-1}} \cdot \tau_1(\mathcal{M}) , \text{ and} \tag{11}$$

$$\det(\mathcal{M}_1)^{1/n} \leq (\gamma\delta)^{\frac{k-1}{\beta-1}} \cdot \det(\mathcal{M})^{\frac{1}{kn}} . \tag{12}$$

Proof. Let us pick the first $i \in [0, p-1]$ such that $\tau_1(\mathcal{M}_{i\beta+\beta}) = \tau_1(\mathcal{M})$. In particular, there exists some rank-one module lattice $\mathcal{M}' \subset \mathcal{M}_{i\beta+\beta}$ with $\mathcal{M}' \not\subset \mathcal{M}_{i\beta}$ such that $\det(\mathcal{M}')^{1/n} = \tau_1(\mathcal{M})$. Therefore (since the $\mathcal{M}_i$ are primitive), $\pi_{\mathcal{M}_{i\beta}^{\perp}}(\mathcal{M}') \subset \mathcal{M}_{[i\beta+1, i\beta+\beta]}$ is a (non-zero) rank-one module lattice so that

$$\tau_1(\mathcal{M}_{[i\beta+1,i\beta+\beta]}) \leq \det(\pi_{\mathcal{M}_{i\beta}^{\perp}}(\mathcal{M}'))^{1/n} \leq \det(\mathcal{M}')^{1/n} = \tau_1(\mathcal{M}) .$$

Therefore, by the primal property,

$$\det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} \leq \gamma \cdot \tau_1(\mathcal{M}_{[i\beta+1,i\beta+\beta]}) \leq \gamma \cdot \tau_1(\mathcal{M}) .$$

Eq. (11) then follows by Lemma 5.3:

$$\det(\mathcal{M}_1)^{1/n} \leq \gamma^{\frac{2i\beta}{\beta-1}+1} \cdot \delta^{\frac{2i\beta}{\beta-1}} \cdot \tau_1(\mathcal{M}) \leq \gamma^{\frac{2(p-1)\beta}{\beta-1}+1} \cdot \delta^{\frac{2(p-1)\beta}{\beta-1}} \cdot \tau_1(\mathcal{M}) .$$

In order to derive Eq. (12), we take the product of Lemma 5.3 over $0 \leq i \leq p-1$,

$$\det(\mathcal{M}_1)^{p/n} \leq (\gamma\delta)^{\frac{p(p-1)\beta}{\beta-1}} \cdot \prod_{i=0}^{p-1} \det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} . \tag{13}$$

Using the primal condition, we also have,

$$\det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} \leq \gamma\delta \cdot \det(\mathcal{M}_{[i\beta+1,i\beta+\beta]})^{1/(\beta n)} . \tag{14}$$

By combining Eqs. (13) and (14), we see that

$$\begin{aligned}\det(\mathcal{M}_1)^{p/n} &\leq (\gamma\delta)^{\frac{p(p-1)\beta}{\beta-1}} \cdot (\gamma\delta)^p \cdot \prod_{i=0}^{p-1} \det(\mathcal{M}_{[i\beta+1,i\beta+\beta]})^{1/(\beta n)} \\ &= (\gamma\delta)^{\frac{p(k-1)}{\beta-1}} \cdot \det(\mathcal{M})^{1/(\beta n)},\end{aligned}$$

and the result follows. □

Corollary 5.5. *If a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is (γ, β)-slide-reduced, then*

$$\lambda_1(\mathcal{M}_1) \leq \frac{\gamma\delta_n}{\alpha_R} \cdot (\gamma\delta)^{\frac{2(k-\beta)}{\beta-1}} \cdot \lambda_1(\mathcal{M}) \ , \ \textit{and} \tag{15}$$

$$\lambda_1(\mathcal{M}_1) \leq \delta_n(\gamma\delta)^{\frac{k-1}{\beta-1}} \cdot \det(\mathcal{M})^{\frac{1}{kn}} \ . \tag{16}$$

Proof. Combining Lemma 2.2 and Eq. (11) from Lemma 5.4, we obtain

$$\det(\mathcal{M}_1)^{1/n} \leq \gamma^{\frac{2(k-\beta)}{\beta-1}+1} \cdot \delta^{\frac{2(k-\beta)}{\beta-1}} \cdot \frac{\lambda_1(\mathcal{M})}{\alpha_R} \ .$$

By the definition of Hermite's constant δ_n, we obtain Eq. (15).

Eq. (16) is follows directly from applying the definition of Hermite's constant to Eq. (12) from Lemma 5.4. □

5.1 Finding slide-reduced filtrations

We now show how to use a DIP oracle to build a slide-reduced filtration and then derive our main results.

Definition 5.6 ((γ, k, β)-RFP)**.** *For a number field K, ring $R \subset K$, rank $k \geq 2$, block size $\beta \geq 2$ dividing k, and approximation factor $\gamma \geq 1$, the (γ, k, β)-Reduced Filtration Problem, or (γ, k, β)-RFP, is the search problem defined as follows. The input is a module lattice $\mathcal{M} \subset K^k$ with rank k, and the goal is to find a (γ, β)-slide-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$.*

Theorem 5.7. *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, block size $\beta \geq 2$ dividing k, and approximation factor $\gamma \geq 1$, there is an efficient reduction from $((1+\varepsilon)\gamma, k, \beta)$-RFP to (γ, β)-DIP'.*

In particular, if R is dual-closed or $\beta = 2$, then there is a reduction from $((1+\varepsilon)\gamma, k, \beta)$-RFP to (γ, β)-DIP (as opposed to DIP').

Proof. On input (a generating set for) a module lattice $\mathcal{M} \subset K^k$ with rank k, the reduction first computes a filtration $\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$. It then repeatedly updates this filtration in place as follows.

1. **Primal reduction.** For each $\mathcal{M}_{[i\beta+1, i\beta+\beta]}$ where $i \in [0, p-1]$, the reduction calls its (γ, β)-DIP oracle (i.e., it calls its DIP' oracle with $s = \text{PRIMAL}$) with $\mathcal{M}_{[i\beta+1, i\beta+\beta]}$ as input, receiving as output $\widetilde{\mathcal{M}}'_{i\beta+1} \subset \mathcal{M}_{[i\beta+1, i\beta+\beta]}$. We may assume without loss of generality that $\widetilde{\mathcal{M}}'_{i\beta+1}$ is primitive, i.e., $\widetilde{\mathcal{M}}'_{i\beta+1} = \mathcal{M}_{[i\beta+1, i\beta+\beta]} \cap \text{span}_K(\widetilde{\mathcal{M}}'_{i\beta+1})$. If $(1+\varepsilon)^n \det(\widetilde{\mathcal{M}}'_{i\beta+1}) < \det(\widetilde{\mathcal{M}}_{i\beta+1})$, then the reduction updates the filtration so that $\widetilde{\mathcal{M}}_{i\beta+1} = \widetilde{\mathcal{M}}'_{i\beta+1}$, leaving the full block $\mathcal{M}_{[i\beta+1, i\beta+\beta]}$ unchanged. (Formally, to do this, the reduction can, e.g., pick any $(i\beta+1)$-dimensional K-subspace W_1 of $\mathcal{M}_{i\beta+\beta}$ such that $\pi_{\mathcal{M}_{i\beta}^{\perp}}(W_1) = \text{span}_K(\widetilde{\mathcal{M}}'_{i\beta+1})$ together with a nested sequence of subspaces $W_1 \subset W_2 \subset \cdots \subset W_\beta = \text{span}_K(\mathcal{M}_{i\beta+\beta})$ and set $\mathcal{M}_{i\beta+j} = \mathcal{M} \cap W_j$ for $j = 1, \ldots, \beta$.)

2. **Dual reduction.** For each $\mathcal{M}_{[i\beta+2, i\beta+\beta+1]}$ where $i \in [0, p-2]$, the reduction calls the dual (γ, β)-DIP oracle with $\mathcal{M}^*_{[i\beta+2, i\beta+\beta+1]}$ as input, receiving as output $(\widetilde{\mathcal{M}}')^*_{i\beta+\beta+1}$.

   If $\det(\widetilde{\mathcal{M}}'_{i\beta+\beta+1}) > (1+\varepsilon)^n \det(\widetilde{\mathcal{M}}_{i\beta+\beta+1})$, then it updates the filtration so that $\widetilde{\mathcal{M}}_{i\beta+\beta+1} = \widetilde{\mathcal{M}}'_{i\beta+\beta+1}$, leaving the full dual block $\mathcal{M}_{[i\beta+2, i\beta+\beta+1]}$ unchanged.

If no update is made in Step 2, then the algorithm terminates and outputs the filtration.

Our proof is more-or-less identical to the proof in [GN08]. We first observe that the output is in fact a $((1+\varepsilon)\gamma, \beta)$-reduced slide filtration of rank k. In order to see this, observe that the primal conditions are satisfied,

$$\det(\widetilde{\mathcal{M}}_{i\beta+1})^{1/n} \leq (1+\varepsilon)\det(\widetilde{\mathcal{M}}'_{i\beta+1})^{1/n} \leq (1+\varepsilon)\gamma\tau_1(\mathcal{M}_{[i\beta+1,i\beta+\beta]}) ,$$

after the end of Step 1. If no updates happen in Step 2, then clearly the primal conditions remain satisfied.

If no update happens in Step 2, this means that

$$\begin{aligned}\det(\widetilde{\mathcal{M}}_{i\beta+\beta+1})^{1/n} &\geq \frac{1}{(1+\varepsilon)}\det(\widetilde{\mathcal{M}}'_{i\beta+\beta+1})^{1/n} \\ &\geq \frac{1}{(1+\varepsilon)\gamma\tau_1(\mathcal{M}^*_{[i\beta+2,i\beta+\beta+1]})} ,\end{aligned}$$

i.e., the dual conditions are satisfied.

It remains to show that the reduction terminates efficiently. We will analyze the following potential function,

$$\Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k) = \prod_{i=1}^{p-1} \det(\mathcal{M}_{i\beta}) .$$

At the beginning of the reduction, $\log \Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k)$ is bounded by a polynomial in the input size (since Φ is efficiently computable). And, by Fact 2.9, $-\log(\Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k))$ is bounded by a polynomial in the input size throughout the reduction. Furthermore, the potential does not change at all in Step 1. Therefore, it suffices to show that the potential decreases by at least, say, a constant factor every time that the reduction updates the filtration in Step 2.

Observe that Φ strictly decreases after each dual step in which the filtration is updated. In order to see this, suppose such an update occurs on the dual block $\mathcal{M}_{[i\beta+2,i\beta+\beta+1]}$, and notice that all elements in Φ remain unchanged except $\det(\mathcal{M}_{i\beta+\beta})$, where

$$\det(\mathcal{M}_{i\beta+\beta}) = \det(\mathcal{M}_{i\beta+1})\det(\mathcal{M}_{[i\beta+2,i\beta+\beta]}) .$$

Let $\widehat{\mathcal{M}}_0$ be $\mathcal{M}_{[i\beta+2,i\beta+\beta]}$ before the update and let $\widehat{\mathcal{M}}_1$ be $\mathcal{M}_{[i\beta+2,i\beta+\beta]}$ after the update. Since $\det(\mathcal{M}_{[i\beta+2,i\beta+\beta+1]})$ remains unchanged after the dual reduction step, we have

$$\begin{aligned}\det(\widehat{\mathcal{M}}_1)\det(\widetilde{\mathcal{M}}'_{i\beta+\beta+1}) &= \det(\widehat{\mathcal{M}}_0)\det(\widetilde{\mathcal{M}}_{i\beta+\beta+1}) \\ &\leq \det(\widehat{\mathcal{M}}_0)\cdot\frac{\det(\widetilde{\mathcal{M}}'_{i\beta+\beta+1})}{(1+\varepsilon)^n} .\end{aligned}$$

Therefore,

$$\det(\widehat{\mathcal{M}}_1) \leq \frac{\det(\widehat{\mathcal{M}}_0)}{(1+\varepsilon)^n} .$$

It follows that $\det(\mathcal{M}_{i\beta+\beta})$ decreases by at least a factor of $(1+\varepsilon)^n$.

Notice that no other terms in the potential change after such an update in Step 2. Therefore, the potential Φ decreases by a factor of at least $(1+\varepsilon)^n$ after the occurrence of each dual update, as needed. □

Corollary 5.8. *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, block size $\beta \geq 2$ dividing k, constant $\varepsilon > 0$, and approximation factor $\gamma' \geq 1$, there is an efficient reduction from (γ, k)-DIP to (γ', β)-DIP', where*

$$\gamma := (1+\varepsilon)\gamma' \cdot ((1+\varepsilon)\gamma'\delta)^{\frac{2(k-\beta)}{\beta-1}} .$$

If R is dual-closed or $\beta = 2$, then there is a reduction from (γ, k)-DIP to (γ', β)-DIP (as opposed to DIP').

Proof. The reduction takes as input a module lattice $\mathcal{M}$ of rank k and solves the corresponding $((1+\varepsilon)\gamma', k, \beta)$-RFP instance using its (γ', β)-DIP oracle as in Theorem 5.7. I.e., it finds a $((1+\varepsilon)\gamma', \beta)$-slide-reduced filtration $\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$. It then outputs $\mathcal{M}_1$. Then, by Eq. (11) from Lemma 5.4,

$$\det(\mathcal{M}_1)^{1/n} := ((1+\varepsilon)\gamma')^{\frac{2(k-\beta)}{\beta-1}+1} \cdot \delta^{\frac{2(k-\beta)}{\beta-1}} \cdot \tau_1(\mathcal{M}) ,$$

as needed. □

Corollary 5.9. *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, block size $\beta \geq 2$ dividing k, constant $\varepsilon > 0$, and approximation factor $\gamma' \geq 1$, there is an efficient reduction from (γ_R, β, k)-RFP to (γ', β)-ModuleSVP', where*

$$\gamma_R := (1+\varepsilon)\gamma'\delta_n/\alpha_R .$$

If R is dual-closed or $\beta = 2$, then there is a reduction from (γ_R, β, k)-RFP to (γ', β)-ModuleSVP (as opposed to ModuleSVP').

Proof. The reduction instantiates the (γ, β)-DIP' oracle needed in Theorem 5.7 using its (γ', β)-ModuleSVP oracle and the reduction from Theorem 2.8, where $\gamma := \gamma'\delta_n/\alpha_R$. This yields a γ_R-reduced filtration, where $\gamma_R = (1+\varepsilon)\gamma = (1+\varepsilon)\gamma'\delta_n/\alpha_R$. □

Theorem 5.10 (Main Theorem)**.** *For any number field K, (discrete, full-rank) ring $R \subset K$, rank $k \geq 2$, block size $\beta \geq 2$ dividing k, constant $\varepsilon > 0$, and approximation factor $\gamma' \geq 1$, there is an efficient reduction from (γ, k)-ModuleSVP to (γ', β)-ModuleSVP', where*

$$\gamma := (1+\varepsilon) \cdot \left(\gamma' \cdot \frac{\delta_n}{\alpha_R}\right)^2 \cdot \left((1+\varepsilon) \cdot \gamma' \cdot \frac{\delta\delta_n}{\alpha_R}\right)^{\frac{2(k-\beta)}{\beta-1}} .$$

There is also an efficient reduction from (γ_H, k)-ModuleHSVP to (γ', β)-ModuleSVP', where

$$\gamma_H := \gamma'\delta_n \cdot \left((1+\varepsilon)\gamma' \cdot \frac{\delta\delta_n}{\alpha_R}\right)^{\frac{k-1}{(\beta-1)}} .$$

If R is dual-closed or $\beta = 2$, then there is a reduction from (γ, k)-ModuleSVP and from (γ_H, k)-ModuleHSVP to (γ', β)-ModuleSVP (as opposed to ModuleSVP').

Proof. In fact, the reduction is the same for both ModuleSVP and ModuleHSVP. On input (a generating set for) a module lattice $\mathcal{M} \subset K^k$ with rank k, the reduction proceeds as follows. It first obtains a γ_R-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ using its (γ', β)-ModuleSVP oracle, where $\gamma_R := \frac{(1+\varepsilon)\gamma'\delta_n}{\alpha_R}$ (by Corollary 5.9). It then calls its (γ', β)-ModuleSVP oracle on $\mathcal{M}_\beta$, which returns a vector $\boldsymbol{x} \in \mathcal{M}_\beta \subseteq \mathcal{M}$ such that $0 < \|\boldsymbol{x}\| \leq \gamma'\lambda_1(\mathcal{M}_\beta)$. Finally, our reduction outputs $\boldsymbol{x}$.

Since $\mathcal{M}_1 \subset \mathcal{M}_\beta$, we have

$$0 < \|\boldsymbol{x}\| \le \gamma' \lambda_1(\mathcal{M}_\beta) \le \gamma' \lambda_1(\mathcal{M}_1) \ .$$

By Eq. (15) of Corollary 5.5,

$$\begin{aligned} \lambda_1(\mathcal{M}_1) &\le \frac{\gamma_R \delta_n}{\alpha_R} \cdot (\gamma_R \delta)^{\frac{2(k-\beta)}{\beta-1}} \cdot \lambda_1(\mathcal{M}) \\ &= (1+\varepsilon) \cdot \frac{\gamma' \delta_n^2}{\alpha_R^2} \cdot ((1+\varepsilon)\gamma' \delta_n \delta / \alpha_R)^{\frac{2(k-\beta)}{\beta-1}} \cdot \lambda_1(\mathcal{M}) \ . \end{aligned}$$

Combining the above two expressions, we get

$$0 < \|\boldsymbol{x}\| \le \gamma \lambda_1(\mathcal{M}) \ ,$$

as needed.

Similarly, by Eq. (16) of Corollary 5.5,

$$\begin{aligned} \|\boldsymbol{x}\| &\le \gamma' \delta_n (\gamma_R \delta)^{\frac{k-1}{(\beta-1)}} \cdot \det(\mathcal{M})^{\frac{1}{kn}} \\ &= \gamma' \delta_n ((1+\varepsilon)\gamma' \delta \delta_n / \alpha_R)^{\frac{k-1}{(\beta-1)}} \cdot \det(\mathcal{M})^{\frac{1}{kn}} \ , \end{aligned}$$

as needed. □

References

[ABD+19] Erdem Alkim, Joppe W. Bos, Léo Ducas, Patrick Longa, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Chris Peikert, Ananth Raghunathan, Douglas Stebila, Karen Easterbrook, and Brian A LaMacchia. FrodoKEM. https://frodokem.org/, 2019.

[Ajt96] Miklós Ajtai. Generating hard instances of lattice problems. In *STOC*, 1996.

[ALNS19] Divesh Aggarwal, Jianwei Li, Phong Q. Nguyen, and Noah Stephens-Davidowitz. Slide reduction, revisited—Filling the gaps in SVP approximation. https://arxiv.org/abs/1908.03724, 2019.

[Bab86] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1), 1986.

[BBPS18] Madalina Bolboceanu, Zvika Brakerski, Renen Perlman, and Devika Sharma. Order-LWE and the hardness of Ring-LWE with entropic secrets, 2018. https://eprint.iacr.org/2018/494.

[BCD+16] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: take off the ring! Practical, quantum-secure key exchange from LWE. In *CCS*, 2016.

[CDPR16] Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Eurocrypt*, 2016.

[CDW17] Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short Stickelberger class relations and application to Ideal-SVP. In *Eurocrypt*, 2017. https://eprint.iacr.org/2016/885.

[CGS14] Peter Campbell, Michael Groves, and Dan Shepherd. Soliloquy: a cautionary tale. ETSI 2nd Quantum-Safe Crypto Workshop, 2014.

[DD12] Léo Ducas and Alain Durmus. Ring-LWE in polynomial rings. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC*, 2012.

[DPW19] Léo Ducas, Maxime Plançon, and Benjamin Wesolowski. On the shortness of vectors to be found by the Ideal-SVP quantum algorithm. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO*, 2019.

[Duc17] Léo Ducas. Advances on quantum cryptanalysis of ideal lattices. *Nieuw Archief voor Wiskunde*, 18(5), 2017.

[FS10] Claus Fieker and Damien Stehlé. Short bases of lattices over number fields. In *ANTS*, 2010.

[GLM09] Y. H. Gan, C. Ling, and W. H. Mow. Complex lattice reduction algorithm for low-complexity full-diversity MIMO detection. *IEEE Transactions on Signal Processing*, 57(7), 2009.

[GN08] Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within Mordell's inequality. In *STOC*, 2008.

[GPV08] Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008. https://eprint.iacr.org/2007/432.

[HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: a ring-based public key cryptosystem. In *ANTS*, 1998.

[HS12] Guillaume Hanrot and Damien Stehle. A complete worst-case analysis of Kannans shortest lattice vector algorithm, 2012.

[KL17] Taechan Kim and Changmin Lee. Lattice reductions over Euclidean rings with applications to cryptanalysis. In Máire O'Neill, editor, *Cryptography and Coding*, 2017.

[LLL82] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4), 1982.

[LM06] Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP*, 2006.

[LPL18] S. Lyu, C. Porter, and C. Ling. Performance limits of lattice reduction over imaginary quadratic fields with applications to compute-and-forward. In *2018 IEEE Information Theory Workshop (ITW)*, 2018.

[LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and Learning with Errors over rings. In *Eurocrypt*, 2010.

[LPSW19] Changmin Lee, Alice Pellet-Mary, Damien Stehlé, and Alexandre Wallet. An LLL algorithm for module lattices. In *ASIACRYPT*, 2019. https://eprint.iacr.org/2019/1035.

[LS12] Adeline Langlois and Damien Stehlé. Hardness of decision (R)LWE for any modulus, 2012.

[LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography*, 75(3), 2015.

[Mah18] Urmila Mahadev. Classical Verification of Quantum Computations. In *arXiv:1804.01082 [quant-ph]*, 2018. http://arxiv.org/abs/1804.01082.

[Mic07] Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions. *Computational Complexity*, 16(4), 2007.

[MR07] Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal of Computing*, 37(1), 2007.

[MW16] Daniele Micciancio and Michael Walter. Practical, predictable lattice basis reduction. In *Eurocrypt*, 2016. http://eprint.iacr.org/2015/1123.

[Nap96] Huguette Napias. A generalization of the LLL-algorithm over Euclidean rings or orders. *Journal de Théorie des Nombres de Bordeaux*, 8(2), 1996.

[NIS18] Computer Security Division NIST. Post-quantum cryptography. https://csrc.nist.gov/Projects/Post-Quantum-Cryptography, 2018.

[NS01] Phong Q. Nguyen and Jacques Stern. The two faces of lattices in cryptology. In *CaLC*, 2001.

[NV10] Phong Q. Nguyen and Brigitte Vallée, editors. *The LLL algorithm: Survey and applications*. Springer-Verlag, 2010.

[Pei09] Chris Peikert. Public-key cryptosystems from the worst-case Shortest Vector Problem. In *STOC*, 2009.

[Pei15] Chris Peikert. What does GCHQ's "cautionary tale" mean for lattice cryptography? https://web.eecs.umich.edu/~cpeikert/soliloquy.html, 2015.

[Pei16] Chris Peikert. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science*, 10(4), 2016.

[PHS19] Alice Pellet-Mary, Guillaume Hanrot, and Damien Stehlé. Approx-SVP in ideal lattices with pre-processing. In Yuval Ishai and Vincent Rijmen, editors, *Eurocrypt*, 2019.

[PP19] Chris Peikert and Zachary Pepin. Algebraically structured LWE, revisited. In *TCC*, 2019.

[PR06] Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, 2006.

[PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of Ring-LWE for any ring and modulus. In *STOC*, 2017. https://eprint.iacr.org/2017/258.

[Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), 2009.

[RSSS17] Miruna Roşca, Amin Sakzad, Damien Stehlé, and Ron Steinfeld. Middle-Product Learning with Errors. In Jonathan Katz and Hovav Shacham, editors, *CRYPTO*, 2017.

[RSW18] Miruna Rosca, Damien Stehlé, and Alexandre Wallet. On the Ring-LWE and Polynomial-LWE problems. In Jesper Buus Nielsen and Vincent Rijmen, editors, *Eurocrypt*, 2018.

[SE94] Claus-Peter Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathmatical Programming*, 66, 1994.

[Sha84] Adi Shamir. A polynomial-time algorithm for breaking the basic Merkle-Hellman cryptosystem. *IEEE Trans. Inform. Theory*, 30(5), 1984.

[SS11] Damien Stehlé and Ron Steinfeld. Making NTRU as secure as worst-case problems over ideal lattices. In *EUROCRYPT*, 2011.

[SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, 2009.

[SW14] Uri Shapira and Barak Weiss. A volume estimate for the set of stable lattices. *Comptes Rendus Mathématique. Académie des Sciences. Paris*, 352(11), 2014.