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

Tamalika Mukherjee[1⋆] and Noah Stephens-Davidowitz[2⋆⋆]

[1] Purdue University, tmukherj@purdue.edu
[2] Cornell University, noahsd@gmail.com

Abstract. This is the extended abstract of [MS20]. See the full version at eprint:2019/1142.
We show how to generalize lattice reduction algorithms to module lattices. Specifically, we 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 high-dimensional generalization of the ($\mathbb{Z}$-)basis of a lattice.
The particular value of γ that we achieve depends on the underlying number field K, the order $R \subseteq \mathcal{O}_K$, and the embedding (as well as, of course, k and β). However, for reasonable choices of these parameters, the resulting value of γ 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 higher-rank instances of approximate ModuleSVP.
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, while our reduction works . Indeed, at a high level our reduction can be thought of as a generalization of theirs in roughly the same way that block reduction generalizes LLL reduction.
In this extended abstract, we present a special case of the more general result to appear in the full version [MS20].

1 Introduction

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

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

⋆ This work was done while being supported by The Center for Science of Information, an NSF Science and Technology Center, Cooperative Agreement # CCF 0939370.
⋆⋆ Part of this work was supported by NSF-BSF grant number 1718161 and NSF CAREER Award number 1350619 via Vinod Vaikuntanathan. Part of this work was done while the author was at the Centre for Quantum Technologies at the National University of Singapore, Massachusetts Institute of Technology, and the Simons Institute in Berkeley.

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 of this very large approximation factor, the LLL algorithm has found innumerable applications [LLL82, Bab86, SE94, 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]—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 (the decision version of) γ'-SVP is hard in the worst case [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 generic lattice-based constructions is their inefficiency (though, see [ABD+19]). Loosely speaking, this inefficiency arises from the fact that a lattice in dimension d typically requires about d^2 numbers to specify—at least 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], since such lattices can be described succinctly.

In particular, cryptographers typically 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 an order $R \subseteq \mathcal{O}_K$ (i.e., a discrete full-rank subring, 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 $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in K^\ell$,

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

By embedding the number field K into $\mathbb{Q}^n$ (or by equipping K with an inner product, which is what we do in the sequel), we can view module lattices as (ℓn)-dimensional "plain" lattices. In particular, it makes sense to talk about the length of module elements. A key parameter is the *rank*

k of the module lattice, which is the dimension of its K-span. We typically think of n as large (i.e., $n \to \infty$) and k as a relatively small constant.[3]

We can then define (γ, k)-ModuleSVP over R as the restriction of γ-SVP to rank-k module lattices $\mathcal{M} \subset K^\ell$ over R (under some inner product). 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 still under consideration by NIST would be broken in practice if significantly faster algorithms were found for ModuleSVP [NIS18]. (Just one relies on "plain" lattices [ABD+19].) 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 (growing) 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 by an adversary 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 cryptographic schemes based on structured lattices. (We are intentionally ignoring many important details here for simplicity. See [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 (surprising) line of work in the $k = 1$ case suggests that we need a better understanding of (γ, k)-ModuleSVP for all γ and k.

To that end, we observe that 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 approx-

[3] Notice that module lattices correspond exactly to lattices that are closed under a certain set of linear transformations—the linear transformations corresponding to multiplication by elements of R.

imation factor γ depending on d, m, and γ'. Indeed, the LLL algorithm can be viewed as 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 more-or-less our only non-trivial algorithms for superconstant approximation factors. (See [ALNS19].)

In other words, to solve γ-SVP (or, for that matter, (γ, k)-ModuleSVP) for superconstant γ, the best known strategy works by reducing 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 $d := kn$. (We have chosen this rather strange parameterization to more easily compare with our results for ModuleSVP.) For cryptanalysis, we typically must take $\beta = \Omega(k)$ and $\gamma' \leq \text{poly}(d)$ in order to achieve a final approximation factor γ that is polynomial in the dimension $d = kn$.

1.1 Our results

Lattice reduction for Modules. Our primary contribution is the following reduction.

Theorem 1 (Informal, see the discussion below and the full version [MS20]). *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 this generalizes to arbitrary orders and a more general class of embeddings that we call "semicanonical." They also showed a very interesting algorithm for $(\gamma, 2)$-ModuleSVP, which requires a CVP oracle over a lattice depending only on R. 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. (In this extended abstract, we only present a special case of the $\beta = 2$ reduction. See [MS20] for the general reduction.)

In the general case $\beta \geq 2$, we note the obvious resemblance between the approximation factor achieved by Theorem 1 and the approximation factor shown in 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]).[4] Therefore, we can interpret Theorem 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 does not mention the number field K, the associated embedding, or the order $R \subseteq \mathcal{O}_K$. In fact, the reduction works for any nice enough number field K, *any* order $R \subseteq \mathcal{O}_K$, and a reasonably large class of embeddings that we call semicanonical. These are generalizations of the canonical embedding that might prove useful in other settings. (Formally, we consider semicanonical *inner products* on K. See Sections 1.2 and 2.5.) Furthermore, the approximation factor that we achieve depends on certain geometric properties of the order and the embedding. (See the full version [MS20] for the precise statement.) The approximation factor shown in Theorem 1 is (a loose upper bound on) what we achieve for the canonical embedding of the ring of integers of a cyclotomic number field.

Two variants. As additional contributions, we note that our reduction can also be used to solve two variants of ModuleSVP.

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 satisfying this inequality. (The minimal value of γ for which γ-HSVP is a total problem is called *Hermite's constant*, which explains the name.) In particular, $(\gamma\sqrt{kn}, k)$-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. (The same is true of many "plain" basis reduction algorithms [GN08, ALNS19].) This variant of SVP is enough for most cryptanalytic applications, so that this

[4] Indeed, if we take $n = 1$ and $\gamma' = 1$, then we recover the original slide reduction algorithm from [GN08]. Specializing further to $\beta = 2$ recovers LLL.

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 2 (Informal, see the full version [MS20]). *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 2 is (a loose upper bound on) what we achieve for the canonical embedding of the ring of integers of a cyclotomic number field.

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 is perhaps best viewed as a consequence of Theorem 3. We again note the obvious similarity between Theorem 3 and Eq. (1). (There is an analogous result for what we might call "RankinDIP," in honor of Rankin's constants, 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 $\det(\mathcal{M})^{1/(kn)}$. For simplicity, we do not bother to make this formal.)

Theorem 3 (Informal, see the full version [MS20]). *For $2 \leq \beta < k$ with β dividing k, there is an efficient reduction from (γ, k)-DIP to (γ', β)-DIP, where*

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

Again, the resulting approximation factor depends on the geometry of the order R, and the above result corresponds to the case when $R = \mathcal{O}_K$ is the ring of integers of a number field K under the canonical embedding.

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 "blocks"—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 modules over more general orders R. In other words, 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 principal. More generally, all rank-k module lattices have an R-basis consisting of k vectors if and only if R is a principal ideal domain. Typically, the rings that interest us are *not* principal ideal domains.) This means that basis-reduction techniques do not really make sense over an R-basis.

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 $j - i + 1$. (We are being deliberately vague about what we mean by "projection" here. See Sections 1.2 and 2.5.) They are well-behaved in other ways as well. For example, (for nice enough embeddings) 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.[5]

From vectors to ideals (or sublattices). By working with filtrations, our reduction is most naturally viewed as a variant of basis reduction with the Gram-Schmidt vectors $\pi_{\mathcal{L}_{i-1}^\perp}(\boldsymbol{b}_i)$ replaced by ideals $\pi_{\mathcal{M}_{i-1}^\perp}(\mathcal{M}_i)$, and lengths replaced by the determinant. This naturally gives rise to Theorem 3—a reduction from DIP to DIP.

[5] 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 quite similar to filtrations. E.g., a pseudobasis can be converted into the filtration given by $\mathcal{M}_i := \mathcal{I}_1\boldsymbol{b}_1 + \cdots + \mathcal{I}_i\boldsymbol{b}_i$.

Indeed, this DIP-to-DIP reduction actually "never looks at the length of a vector." It only considers determinants of submodules. 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 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 (which 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 is a partial converse of Minkowski's theorem for ideals. The quantitative result depends on the geometry of the order R, which is the main reason that 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 freely between DIP and ModuleSVP (at the cost of a higher approximation factor), which yields our main result.

Projections In order for our reduction to make sense, we need some kind of notion of "projection." In particular, we need to make sense of the "projection of a module lattice $\mathcal{M} \subset K^\ell$ orthogonal to some submodule lattice $\mathcal{M}' \subseteq \mathcal{M}$" (since this is necessary to define, e.g., $\mathcal{M}_{[i,j]}$). In what follows, we use the word *projection* to mean any $\mathbb{Q}$-linear map that equals its own square.

One way to define projection is by noting that our notion of length in K^ℓ comes from viewing $K^\ell = K \oplus \cdots \oplus K$ as an $n\ell$-dimensional $\mathbb{Q}$-vector space, and fixing some inner product $\langle \cdot, \cdot \rangle_\rho$ on K (which immediately yields an inner product on K^ℓ). Indeed, it does not make sense to talk about ModuleSVP without first fixing some notion of length in K^ℓ, and the most natural notion is given by $\|\boldsymbol{x}\|_\rho^2 := \langle \boldsymbol{x}, \boldsymbol{x} \rangle_\rho$. We can then define our projection as simply the standard orthogonal projection over any $\mathbb{Q}$ vector space. The projection map $\Pi_{\rho,W}$ onto a subspace $W \subseteq K^\ell$ is

the unique $\mathbb{Q}$-linear map that leaves W unchanged and maps to zero all elements that are $\mathbb{Q}$-orthogonal to W.

This is of course the most natural notion of projection, and the projection Π_ρ has many nice properties (since it is just the standard notion of $\mathbb{Q}$-linear orthogonal projection). For example, Π_ρ is contracting (i.e., it cannot increase the length of a vector), and $\det(\mathcal{M}) = \det(V^\perp \cap \mathcal{M}) \cdot \det(\Pi_{\rho,V}(\mathcal{M}))$ (where length and the determinant are defined in terms of the inner product $\langle \cdot, \cdot \rangle_\rho$). However, the lattice $\Pi_{\rho,V}(\mathcal{M})$ might *not* be a module lattice. This is a serious issue because we wish to call our ModuleSVP oracle on this projection.

Another idea is to define a K-linear "inner product" $\langle \cdot, \cdot \rangle_K$ over K^ℓ, given by $\langle \boldsymbol{x}, \boldsymbol{y} \rangle_K := \sum_{i=1}^{\ell} x_i \overline{y_i}$, where $\overline{y_i}$ is the complex conjugate of y_i.[6] We can then define $(\mathcal{M}')^\perp := \{\boldsymbol{x} \in K^\ell \ : \ \forall \boldsymbol{y} \in \mathcal{M}', \ \langle \boldsymbol{y}, \boldsymbol{x} \rangle_K = 0\}$ and define the projection mapping $\Pi_K : K^\ell \to K^\ell$ to be the unique K-linear map that leaves $(M')^\perp$ fixed and sends all elements in $\mathcal{M}'$ to $\mathbf{0}$.

Since the map Π_K is K-linear (by definition), it maps the module lattice $\mathcal{M}$ to another module lattice $\Pi_K(\mathcal{M})$. So, it does not have the problem that Π_ρ had. However, Π_K might not interact nicely with $\langle \cdot, \cdot \rangle_\rho$. E.g., Π_K might increase the length of a vector (under the norm induced by Π_ρ), and we might *not* have $\det(\mathcal{M}) = \det(\mathcal{M}') \cdot \det(\Pi_K(\mathcal{M}))$. This is a big problem, since it means that, e.g., non-zero projections of short vectors in $\mathcal{M}$ "might not be found by a ModuleSVP oracle called on $\Pi_\rho(\mathcal{M})$." More generally, basis reduction algorithms rely heavily on both the contracting nature of projection and the identity $\det(\mathcal{M}) = \det(\mathcal{M}') \det(\Pi_\rho(\mathcal{M}))$.

In summary, Π_ρ is the "right" notion of orthogonal projection from a *geometric* perspective, since it behaves nicely in terms of geometric quantities like lengths and determinants. On the other hand, Π_K is the "right" notion of orthogonal projection from a *algebraic* perspective, since it preserves the module structure of lattices. Indeed, there is a sense in which Π_ρ is the *only* projection map that is "nice" geometrically, and Π_K algebraically.

We therefore restrict our attention to number fields K and inner products $\langle \cdot, \cdot \rangle_\rho$ for which $\Pi_\rho = \Pi_K$, so that a single projection has both the algebraic and geometric properties that we need. In particular, we take number fields K that are closed under complex conjugate (these are the totally real fields and CM fields) and inner products $\langle \cdot, \cdot \rangle_\rho$ that "respect

[6] Taking the complex conjugate is necessary to guarantee that $\langle \boldsymbol{x}, \boldsymbol{x} \rangle_K$ is non-zero (and totally positive) for $\boldsymbol{x} \neq \mathbf{0}$. Formally, this is not quite an inner product because the base field is neither $\mathbb{R}$ nor $\mathbb{C}$. But, it *is* a non-degenerate conjugate symmetric sesquilinear form, which makes the analogy useful.

field multiplication" in the sense that $\langle \alpha \boldsymbol{x}, \boldsymbol{y} \rangle_\rho = \langle \boldsymbol{x}, \overline{\alpha} \boldsymbol{y} \rangle_\rho$. Such *semi-canonical* inner products have a simple characterization in terms of (full-rank) linear maps $T : K \to \mathbb{Q}$:

$$\langle \boldsymbol{x}, \boldsymbol{y} \rangle_\rho := \sum_i T(x_i \overline{y_i}) \ .$$

(The *canonical* inner product is the important special case when $T := \mathrm{Tr}_{K/\mathbb{Q}}$ is the trace map.)

These same restrictions are also exactly what is needed to guarantee that the dual $\mathcal{M}^*$ of a module lattice is also a module lattice (which we also need for our reduction, for $k > 2$). See Section 2.5 for more details and other equivalent definitions.

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 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.[7] They also show a surprising algorithm for $(\gamma, 2)$-ModuleSVP (formally, a quantum polynomial-time reduction from this problem to the Closest Vector Problem over a lattice that depends only on K), 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

[7] 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.

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 order R is typically not Euclidean—or even a principal 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 nice filtration.

Acknowledgements

The authors thank Léo Ducas, Chris Peikert, and Alice Silverberg for very helpful discussions. We are also indebted to the anonymous Eurocrypt 2020 reviewers, who identified errors in an earlier version of this work, and the anonymous Crypto 2020 reviewers for their helpful comments.

2 Preliminaries

For $x \in \mathbb{C}$, we write $\overline{x}$ for the complex conjugate of x. For a $\mathbb{Q}$-subspace $V \subseteq \mathbb{Q}^d$ and a rational-valued inner product $\langle \cdot, \cdot \rangle_\rho$, we define the ρ-orthogonal projection onto V as the unique $\mathbb{Q}$-linear map $\Pi_{\rho,V} : \mathbb{Q}^d \to \mathbb{Q}^d$ that satisfies $\Pi_{\rho,V}(\boldsymbol{x}) = \boldsymbol{x}$ for $\boldsymbol{x} \in V$ and $\Pi_{\rho,V}(\boldsymbol{x}) = \boldsymbol{0}$ if $\langle \boldsymbol{y}, \boldsymbol{x} \rangle_\rho = 0$ for all $\boldsymbol{y} \in V$. We write $\langle \cdot, \cdot \rangle_{\mathbb{Q}}$ for the standard inner product over $\mathbb{Q}^d$.

2.1 Lattices

A lattice $\mathcal{L} \subset \mathbb{R}^d$ is the $\mathbb{Z}$-span of finitely many vectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in \mathbb{Q}^d$ such that

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

If $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m$ are $\mathbb{Q}$-linearly independent vectors, then we sometimes call this a $\mathbb{Z}$-basis, and we write $m := \mathrm{rank}_{\mathbb{Q}}(\mathcal{L})$. For any lattice $\mathcal{L} \subset \mathbb{Q}^d$ and sublattice $\mathcal{L}' \subseteq \mathcal{L}$, we say that $\mathcal{L}'$ is *primitive* if $\mathcal{L}' = \mathcal{L} \cap \mathrm{span}_{\mathbb{Q}}(\mathcal{L}')$. If $\mathcal{L}'$ is primitive and $W \subseteq \mathrm{span}_{\mathbb{Q}}(\mathcal{L}')$ is a $\mathbb{Q}$-subspace, then $W \cap \mathcal{L}'$ is also a primitive sublattice with $\mathrm{rank}_{\mathbb{Q}}(W \cap \mathcal{L}') = \dim_{\mathbb{Q}}(W)$.

The lattice determinant is $\det(\mathcal{L}) := \sqrt{\det(\mathbf{G})}$, where $\mathbf{G} \in \mathbb{Q}^{m \times m}$ is the Gram matrix $G_{i,j} := \langle \boldsymbol{b}_i, \boldsymbol{b}_j \rangle_{\mathbb{Q}}$ of $\mathbf{B} = (\boldsymbol{b}_1, \ldots, \boldsymbol{b}_m) \in \mathbb{Q}^{d \times m}$ for any $\mathbb{Z}$-basis $\mathbf{B}$ of $\mathcal{L}$ (the choice of basis does not matter). If $\mathcal{L}' \subset \mathcal{L}$ is primitive and $W \subset \mathbb{Q}^d$ is the subspace of all vectors that are $\mathbb{Q}$-orthogonal to $\mathcal{L}'$, then $\det(\mathcal{L}) = \det(\mathcal{L}') \det(\Pi_{Q,W}(\mathcal{L}))$.

We write $\lambda_1(\mathcal{L}) := \min_{\boldsymbol{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \langle \boldsymbol{y}, \boldsymbol{y} \rangle_{\mathbb{Q}}^{1/2}$ for the length of a 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,

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

The dual has as a basis $\mathbf{B}\mathbf{G}^{-1}$ for any basis $\mathbf{B}$ of $\mathcal{L}$ with Gram matrix $\mathbf{G}$, and in particular, $(\mathcal{L}^*)^* = \mathcal{L}$ and $\det(\mathcal{L}^*) = 1/\det(\mathcal{L})$. We also have the identity $\Pi_{\mathbb{Q},W}(\mathcal{L})^* = W \cap \mathcal{L}^*$ for any subspace $W \subset \mathbb{Q}^n$, provided that $\Pi_{\mathbb{Q},W}(\mathcal{L})$ is a lattice. (Equivalently, this holds for any subspace W that is spanned by dual lattice vectors, or equivalently, a subspace W such that the subspace of vectors $\mathbb{Q}$-orthogonal to W is spanned by lattice vectors.)

For a positive integer k, Hermite's constant is

$$\delta_k := \sup \lambda_1(\mathcal{L}) / \det(\mathcal{L})^{1/k} \ ,$$

where the supremum is over all lattices with rank k. Minkowski's celebrated theorem shows us that $\delta_k \leq \sqrt{2k/(\pi e)}$, and this is known to be tight up to a small constant factor.

2.2 Number fields

A number field K is a finite degree algebraic field extension of the rational numbers $\mathbb{Q}$, i.e., $K \cong \mathbb{Q}[x]/p(x)$ for some irreducible polynomial $p(x) \in \mathbb{Q}[x]$. The degree $n = [K : \mathbb{Q}]$ 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.)

We associate a rational-valued inner product $\langle \cdot, \cdot \rangle_\rho : K \times K \to \mathbb{Q}$ with our number field K, which satisfies the usual three properties of symmetry, linearity in the first argument, and positive definiteness.

2.3 Orders, ideals, and module lattices

For a number field K, the set of all algebraic integers in K, denoted by $\mathcal{O}_K \subset K$, forms a ring (under the usual addition and multiplication

operations in K), called the ring of integers of K. (An algebraic integer is a root of a monic polynomial with coefficients in $\mathbb{Z}$.) The ring of integers $\mathcal{O}_K$ is a free $\mathbb{Z}$-module of rank $n = [K : \mathbb{Q}]$, i.e., it is the set of all $\mathbb{Z}$-linear combinations of some basis $B = \{b_1, \ldots, b_n\} \subset \mathcal{O}_K$.

An *order* of K is a subring $R \subseteq \mathcal{O}_K$ which is also a free $\mathbb{Z}$-module of rank n.

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-span of finitely many vectors $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in K^\ell$

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

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 rank-one module lattices are isomorphic to ideals (under an appropriate scaling of the inner product). We say that such an ideal is *principal* if it is the R-span of a single element $\boldsymbol{x} \in K^\ell$, and we say that $\boldsymbol{x}$ *generates* the ideal.

As the name suggests, module lattices are themselves lattices (when viewed as subsets of $\mathbb{Q}^{kn}$). 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 $r_i \boldsymbol{y}_j$. In particular, if we fix some inner product $\langle \cdot, \cdot \rangle_\rho$ on K^ℓ, where this inner product is defined more rigorously in Subsection 2.5, then we can define, e.g., $\det(\mathcal{M})$, $\lambda_1(\mathcal{M})$, $\mathcal{M}^*$, $\mathrm{rank}_{\mathbb{Q}}(\mathcal{M})$, primitive submodules, etc., in the natural way (see Subsection 2.5).

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

2.4 The canonical embedding and CM/TR fields

The *canonical embedding* of a number field $K := \mathbb{Q}[x]/p(x)$ is an invertible $\mathbb{Q}$-linear map $\sigma : K \to \mathbb{C}^n$. Up to a reordering of the coordinates, it is the unique such map such that field multiplication between two elements $y = (y_1, \ldots, y_n) \in \sigma(K) \subset \mathbb{C}^n$ and $y' = (y'_1, \ldots, y'_n) \in \sigma(K) \subset \mathbb{C}^n$ is coordinate-wise, i.e., $\sigma(yy') = (y_1 y'_1, y_2 y'_2, \ldots, y_n y'_n)$. Equivalently, the embedding $\sigma(y)$ of y is $\sigma(y) = (\sigma_1(y), \ldots, \sigma_n(y)) \in \mathbb{C}^n$, 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.

A number field K is totally real if all of its embeddings are real. It is totally imaginary if none of its embeddings lie in $\mathbb{R}$. In the sequel, we exclusively work over totally real fields or CM-fields. A number field K is a CM-field if it is a quadratic extension of K/F where the base field F is totally real but K is totally imaginary. We write "CM/TR field" to represent number fields that are either CM or totally real. One of the useful properties of CM/TR fields (and the one that we require) is that complex conjugation on $\mathbb{C}$ induces an automorphism on K which is independent of its embedding into $\mathbb{C}$. In other words, for every element $x \in K$, there exists an element $\overline{x} \in K$ such that every field embedding maps $\overline{x}$ to the complex conjugate of the embedding of x. This property allows us to define an "inner product" over K (see next Section).

The inner product induced by the canonical embedding is given by $\langle x, y \rangle_\sigma := \sum \sigma_i(x)\overline{\sigma_i(y)} \in \mathbb{Q}$ for any $x, y \in K$. We will not discuss the canonical embedding much explicitly in the sequel, but it is a very useful and important example.

2.5 Semicanonical inner products

For $\boldsymbol{w}, \boldsymbol{y} \in K^\ell$ for a CM/TR number field K, we define the "inner product" (conjugate symmetric form with $\langle \boldsymbol{w}, \boldsymbol{w} \rangle_K \neq 0$ for $\boldsymbol{w} \neq \boldsymbol{0}$) over K as $\langle \boldsymbol{w}, \boldsymbol{y} \rangle_K := \sum_{i=1}^{k} w_i \overline{y_i}$. We say that $\boldsymbol{w}$ and $\boldsymbol{y}$ are "K-orthogonal" if $\langle \boldsymbol{w}, \boldsymbol{y} \rangle_K = 0$. For a module lattice $\mathcal{M} \subset K^\ell$, we write

$$\mathcal{M}^\perp := \{\boldsymbol{x} \in K^\ell \ : \ \forall \boldsymbol{y} \in \mathcal{M}, \ \langle \boldsymbol{y}, \boldsymbol{x} \rangle_K = 0\}$$

for the set of vectors that are K-orthogonal to $\mathcal{M}$. This is a K-subspace of K^ℓ with dimension equal to $\ell - \mathrm{rank}_K(\mathcal{M})$.

In analogy with ρ-orthogonal projection, for a K-subspace $V \subseteq K^\ell$ we define the "K-orthogonal projection map onto V" $\mathit{\Pi}_{K,V} : K^\ell \to K^\ell$ as the unique K-linear map satisfying $\mathit{\Pi}_{K,V}(\boldsymbol{x}) = \boldsymbol{x}$ for $\boldsymbol{x} \in V$ and $\mathit{\Pi}_{K,V}(\boldsymbol{x}) = \boldsymbol{0}$ if $\langle \boldsymbol{y}, \boldsymbol{x} \rangle_K = 0$ for all $\boldsymbol{y} \in V$.

We now introduce the related notion of a semicanonical inner product, which is a generalization of the inner product induced by the canonical embedding, $\langle x, y \rangle_\sigma = \sum_i \sigma_i(x)\overline{\sigma_i(y)}$ described in the previous section. Semicanonical inner products share many of the nice geometric properties of $\langle \cdot, \cdot \rangle_\sigma$, as we will see below.

Definition 4 (Semicanonical inner product). *Given a rational-valued inner product $\langle \cdot, \cdot \rangle_\rho$ over a CM/TR number field K, we say that ρ is semicanonical if $\langle yz, w \rangle_\rho = \langle y, \overline{z}w \rangle_\rho$ for $w, y, z \in K$.*

It is easy to see that the inner product $\langle \cdot, \cdot \rangle_\sigma$ is semicanonical since for any $w, y, z \in K$, $\langle yz, w \rangle_\sigma = \sum_i \sigma_i(yz)\overline{\sigma_i(w)} = \sum_i \sigma_i(yz)\sigma_i(\overline{w}) = \sum_i \sigma_i(y)\sigma_i(z\overline{w}) = \sum_i \sigma_i(y)\overline{\sigma_i(\overline{z}w)} = \langle y, \overline{z}w \rangle_\sigma$.

For $\boldsymbol{w}, \boldsymbol{y} \in K^\ell$, we define $\langle \boldsymbol{w}, \boldsymbol{y} \rangle_\rho := \sum_{i=1}^k \langle w_i, y_i \rangle_\rho$. We also write $\|\boldsymbol{w}\|_\rho^2 := \langle \boldsymbol{w}, \boldsymbol{w} \rangle_\rho$.

Lemma 5. *Given a CM/TR number field K and inner product $\langle \cdot, \cdot \rangle_\rho$, the following statements are equivalent.*

1. *For $w, y \in K$, there exists a $\mathbb{Q}$-linear transformation $T : K \to \mathbb{Q}$ such that*[8]
$$\langle w, y \rangle_\rho = T(w\overline{y}) \ .$$
2. *ρ is semicanonical.*
3. *For $\boldsymbol{w}, \boldsymbol{y} \in K^\ell$, $\langle \boldsymbol{w}, \boldsymbol{y} \rangle_K = 0$ if and only if $\langle \alpha \boldsymbol{w}, \boldsymbol{y} \rangle_\rho = 0$ for all $\alpha \in K$.*
4. *For any $\boldsymbol{y} \in K^\ell$ and K-subspace $V \subseteq K^\ell$, we have*
$$\Pi_{K,V}(\boldsymbol{y}) = \Pi_{\rho,V}(\boldsymbol{y}) \ .$$

Proof. **(1 $\Leftrightarrow$ 2).**

Assume that Condition 2 holds. Define the transformation $T : K \to \mathbb{Q}$ as,
$$T(z) := \langle z, 1 \rangle_\rho \ .$$
Since $\langle \cdot, \cdot \rangle_\rho$ is $\mathbb{Q}$-linear, we have that T is $\mathbb{Q}$-linear. For any $w, y \in K$, $\langle w, y \rangle_\rho = \langle w\overline{y}, 1 \rangle_\rho = T(w\overline{y})$.

Now, assume that Condition 1 holds, i.e. there exists a $\mathbb{Q}$-linear transformation $T : K \to \mathbb{Q}$ such that $\langle w, y \rangle_\rho = T(w\overline{y})$. For $w, y, z \in K$, we have
$$\langle yz, w \rangle_\rho = T(yz\overline{w}) = \langle y, \overline{z}w \rangle_\rho \ .$$
Therefore ρ is semicanonical.

(2 $\Leftrightarrow$ 3).

We will first assume Condition 2 and show that Condition 3 holds. Note that Condition 3 is a biconditional statement. We prove the forward direction first.

[8] For the special case of the canonical embedding $\langle \cdot, \cdot \rangle_\sigma$, T is the trace.

We need to show that for vectors $\boldsymbol{w}, \boldsymbol{y} \in K^\ell$ and $\alpha \in K$ satisfying $\langle \boldsymbol{w}, \boldsymbol{y} \rangle_K = \sum_{i=1}^k w_i \overline{y_i} = 0$, we have $\langle \alpha \boldsymbol{w}, \boldsymbol{y} \rangle_\rho = 0$. This follows directly from Condition 2,

$$\langle \alpha \boldsymbol{w}, \boldsymbol{y} \rangle_\rho = \sum_{i=1}^k \langle \alpha w_i, y_i \rangle_\rho = \sum_{i=1}^k \langle \alpha, \overline{w_i} y_i \rangle_\rho = \langle \alpha, \sum_{i=1}^k \overline{w_i} y_i \rangle_\rho = \langle \alpha, 0 \rangle_\rho = 0 \, .$$

Now we prove the backward direction for Condition 3, i.e. for vectors $\boldsymbol{w}, \boldsymbol{y} \in K^\ell$ such that for all $\alpha \in K$, $\langle \alpha \boldsymbol{w}, \boldsymbol{y} \rangle_\rho = 0$, we need to show that $\langle \boldsymbol{w}, \boldsymbol{y} \rangle_K = 0$. Based on our assumption and following the calculations above, we get

$$0 = \langle \alpha \boldsymbol{w}, \boldsymbol{y} \rangle_\rho = \langle \alpha, \sum_{i=1}^k \overline{w_i} y_i \rangle_\rho \, .$$

Since the above expression holds for all $\alpha \in K$, suppose that $\alpha = \sum_{i=1}^k \overline{w_i} y_i$, in which case, the above expression becomes $\langle \alpha, \alpha \rangle_\rho = 0$ which implies $\alpha = 0$, or in other words $\sum_{i=1}^k \overline{w_i} y_i = 0$.

Finally, we assume that Condition 3 holds and prove Condition 2. For $\alpha, w', y' \in K$, let $\boldsymbol{w} := (\alpha w', w', 0, \ldots, 0), \boldsymbol{y} := (y', -\overline{\alpha} y', 0, \ldots, 0)$. Observe that

$$\langle \boldsymbol{w}, \boldsymbol{y} \rangle_K = (\alpha w') \overline{y'} + (w')(-\alpha \overline{y'}) = 0 \, .$$

By Condition 3, this implies that $\langle \alpha \boldsymbol{w}, \boldsymbol{y} \rangle_\rho = \langle \alpha w', y' \rangle_\rho + \langle w', -\overline{\alpha} y' \rangle_\rho = 0$. In other words, $\langle \alpha w', y' \rangle_\rho = \langle w', \overline{\alpha} y' \rangle_\rho$. Therefore, ρ must be semicanonical.

(3 $\Leftrightarrow$ 4).

This follows immediately from the definitions of $\Pi_{K,V}$ and $\Pi_{\rho,V}$. In particular, both maps are $\mathbb{Q}$-linear (though $\Pi_{K,V}$ is also K-linear), which means that it suffices to show that they behave identically on some $\mathbb{Q}$-basis of K^ℓ. Indeed, by definition, both of them act as the identity map on V, and their kernels are respectively the subspace of K-orthogonal vectors to V and ρ-orthogonal vectors to V. Therefore, the two maps are the same if and only if the subspace of K-orthogonal vectors equals the subspace of ρ-orthogonal vectors. □

Corollary 6. *For a CM/TR field K, associated semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, order $R \subseteq \mathcal{O}_K$, module lattice $\mathcal{M} \subset K^\ell$ over R, and K-subspace $W \subseteq K^\ell$,*

1. *If R is closed under conjugation then the dual denoted by $\mathcal{M}^*$ is also a module lattice, which satisfies $\det(\mathcal{M}^*) = 1/\det(\mathcal{M})$.*

2. *For any primitive submodule* $\mathcal{M}' \subset \mathcal{M}$ *we have that*

$$\det(\mathcal{M}) = \det(\mathcal{M}') \det(\Pi_{K,(\mathcal{M}')^\perp}(\mathcal{M})) \ .$$

3. *For any* $\boldsymbol{y} \in K^\ell$, $\|\Pi_{\rho,W}(\boldsymbol{y})\|_\rho \le \|\boldsymbol{y}\|_\rho$.
4. *If* $\mathcal{M}$ *has rank* k*, and* $\mathcal{M}' := \Pi_{\rho,W}(\mathcal{M})$ *is also a module lattice* with rank k*, then* $\det(\mathcal{M}') \le \det(\mathcal{M})$.

Proof. To show Item 1, we need to show that for any $\boldsymbol{y} \in \mathcal{M}^*$ and $r \in R$, $r\boldsymbol{y} \in \mathcal{M}^*$. For any $\boldsymbol{w} \in \mathcal{M}$, by the semicanonical property, $\langle \boldsymbol{w}, r\boldsymbol{y} \rangle_\rho = \langle \overline{r}\boldsymbol{w}, \boldsymbol{y} \rangle_\rho$. Since R is closed under conjugation, $\overline{r} \in R$, and $\overline{r}\boldsymbol{w} \in \mathcal{M}$. Since $\boldsymbol{y} \in \mathcal{M}^*$ is a dual vector, $\langle \overline{r}\boldsymbol{w}, \boldsymbol{y} \rangle_\rho \in \mathbb{Z}$, which implies $\langle \boldsymbol{w}, r\boldsymbol{y} \rangle_\rho \in \mathbb{Z}$, i.e., $r\boldsymbol{y} \in \mathcal{M}^*$.

To show Item 2, recall from Section 2.1 that for a lattice $\mathcal{L} \subset \mathbb{Q}^d$ with primitive sublattice $\mathcal{L}' \subset \mathcal{L}$, we have the analogous fact: $\det(\mathcal{L}) = \det(\mathcal{L}') \det(\Pi_{\mathbb{Q},V}(\mathcal{L}))$, where V is the $\mathbb{Q}$-subspace of vectors that are $\mathbb{Q}$-orthogonal to $\mathcal{L}$. Since module lattices $\mathcal{M}$ under the inner product $\langle \cdot, \cdot \rangle_\rho$ are in fact lattices, it follows that $\det(\mathcal{M}) = \det(\mathcal{M}') \det(\Pi_{\rho,W}(\mathcal{M}))$. Finally, by Lemma 5, $\Pi_{\rho,W} = \Pi_{K,W}$, so that the identity holds for $\Pi_{K,W}$ as well.

Similarly, Items 3 and 4 follow from the corresponding facts about projections over $\mathbb{Q}$. □

2.6 Some geometric quantities of orders and module lattices

For an order R of a number field K of degree n with an inner product $\langle \cdot, \cdot \rangle_\rho$, we define

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

where the infimum is over all rank-one modules $\mathcal{I} \subset K^\ell$. (Notice that α_R depends heavily on the choice of inner product $\langle \cdot, \cdot \rangle_\rho$, so perhaps formally we should write $\alpha_{R,\rho}$. We write α_R instead for simplicity.)

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 infimum is over the rank-one submodules $\mathcal{I} \subset \mathcal{M}$ (i.e., ideals). This quantity can be viewed as a different way to generalize $\lambda_1(\mathcal{L})$ to module lattices over arbitrary orders. 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. So,

over $\mathbb{Z}$, $\tau_1 = \lambda_1$. For higher-dimensional orders 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 define

$$\mu_{R,k} := \sup_{\mathcal{M}} \frac{\tau_1(\mathcal{M})}{\det(\mathcal{M})^{1/(kn)}} ,$$

where the supremum is over all rank-k module lattices $\mathcal{M} \subset K^{k'}$ (for any integer $k' \geq k$). (This can be thought of as the module analogue of either Rankin's constant or Hermite's constant.)

For a module lattice $\mathcal{M}$ of rank k, we have the simple inequality $\tau_1(\mathcal{M}) \leq \mu_{R,k} \det(\mathcal{M})^{1/(kn)}$, and the following relationship between τ_1 and λ_1, which is governed by α_R.

Lemma 7. *Given a number field K, order $R \subseteq \mathcal{O}_K$, a module lattice $\mathcal{M}$, and an inner product $\langle \cdot, \cdot \rangle_\rho$*

$$\frac{\lambda_1(\mathcal{M})}{\delta_n} \leq \tau_1(\mathcal{M}) \leq \frac{\lambda_1(\mathcal{M})}{\alpha_R} , \tag{2}$$

$$1 \leq \mu_{R,k} \leq \frac{\sqrt{kn}}{\alpha_R} . \tag{3}$$

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

$$\det(\mathcal{I})^{1/n} \leq \frac{\lambda_1(\mathcal{I})}{\alpha_R} . \tag{4}$$

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

$$\tau_1(\mathcal{M}) \leq \det(\mathcal{I})^{1/n} . \tag{5}$$

Combining Eqs. (4) and (5) yields the upper bound in (2).

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 in (2) follows by noting that $\lambda_1(\mathcal{M}) \leq \lambda_1(\mathcal{I}')$.

Observe that for rank k module lattices, Minkowski's theorem gives us $\lambda_1(\mathcal{M}) \leq \sqrt{kn} \det(\mathcal{M})^{1/(kn)}$. Combining this relation with the upper bound from (2) yields the upper bound in (3). The lower bound is witnessed by, e.g., $\mathcal{M} = R^k$, which satisfies $\tau_1(\mathcal{M}) = \det(R)^{1/n} = \det(M)^{1/(kn)}$. □

We also have the following well-known property of the canonical embedding.

Lemma 8. *For any order $R \subseteq \mathcal{O}_K$ of any number field K under the inner product $\langle \cdot, \cdot \rangle_\sigma$ induced by 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 $\alpha_R \geq 1$.

2.7 ModuleSVP and the Dense Ideal Problem

We now provide the formal definition of ModuleSVP, and its variant the Dense Ideal Problem.

Definition 9 (ModuleSVP). *For a number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 1$, approximation factor $\gamma = \gamma(R, k) \geq 1$, and inner product $\langle \cdot, \cdot \rangle_\rho$, (γ, k)-*ModuleSVP *is defined as follows. The input is (a generating set for) a module lattice $\mathcal{M} \subset K^\ell$ with rank k. The goal is to output a module element $\boldsymbol{x} \in \mathcal{M}$ such that $0 < \|\boldsymbol{x}\|_\rho \leq \gamma \lambda_1(\mathcal{M})$.*

Definition 10 (The Dense Ideal Problem). *For a number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 2$, approximation factor $\gamma = \gamma(R, k) \geq 1$, and inner product $\langle \cdot, \cdot \rangle_\rho$, the* (γ, k)-Dense Ideal Problem*, or (γ, k)-DIP, is the search problem defined as follows. The input is a (generating set for) module lattice $\mathcal{M} \subset K^\ell$ 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 $\det(\mathcal{M}')^{1/n} \leq \gamma \tau_1(\mathcal{M})$.*

Definition 11 (ModuleHSVP). *For a number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 2$, approximation factor $\gamma = \gamma(R, k) \geq 1$, and inner product $\langle \cdot, \cdot \rangle_\rho$, (γ, k)-*ModuleHSVP *is defined as follows. The input is (a generating set for) a module lattice $\mathcal{M} \subset K^\ell$ with rank k. The goal is to output a module element $\boldsymbol{x} \in \mathcal{M}$ such that $0 < \|\boldsymbol{x}\|_\rho \leq \gamma \det(\mathcal{M})^{1/(kn)}$.*

Notice that a solution to the above problem is guaranteed to exist if $\gamma \geq \delta_{kn}$.

Theorem 12. *For a number field K, order $R \subseteq \mathcal{O}_K$, rank $\beta \geq 2$, approximation factor $\gamma' = \gamma'(R, \beta) \geq 1$, and an inner product $\langle \cdot, \cdot \rangle_\rho$, 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 lattice $\mathcal{M}$ of rank β, and uses the output from the (γ', β)-ModuleSVP oracle which is a non-zero vector $\boldsymbol{x} \in \mathcal{M}$ such that $0 < \|\boldsymbol{x}\|_\rho \leq \gamma' \lambda_1(\mathcal{M})$, to output a submodule $\mathcal{M}' \subset \mathcal{M}$ such that $\det(\mathcal{M}')^{1/n} \leq \gamma \tau_1(\mathcal{M})$.

Let $\mathcal{M}' := R\boldsymbol{x}$, i.e. $\mathcal{M}'$ is a principal ideal generated by $\boldsymbol{x}$. Note that $\lambda_1(\mathcal{M}') \leq \|\boldsymbol{x}\|_\rho \leq \gamma' \lambda_1(\mathcal{M})$. Then using Lemma 7, 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. □

2.8 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. In practice, one can represent elements in K as polynomials with rational coefficients, and the inner product can be represented by specifying the pairwise inner products of basis elements (i.e., as a quadratic form). Since arithmetic operations may be performed efficiently with these representations, we are largely justified in ignoring such bit-level details.

In the full version [MS20], we discuss this a bit more, but see [GN08] for a more detailed discussion about the bit-level complexity of basis reduction, and [LPSW19] for a similar discussion in the context of module lattices specifically. We will need one fact that makes use of the bit-level representation.

Fact 13. *If the number field K, its inner product $\langle \cdot, \cdot \rangle_\rho$, and the order $R \subseteq \mathcal{O}_K$ are represented as described above, then for any integer $\ell \geq 1$ and any module lattice $\mathcal{M} \subset K^\ell$*

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

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 an order $R \subseteq \mathcal{O}_K$ with rank k over a CM/TR field 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:** $\widetilde{\mathcal{M}}_i := \Pi_{K,\mathcal{M}_{i-1}^\perp}(\mathcal{M}_i)$ is a rank-one module lattice over R.

(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 $\mathcal{M}_{[i,j]} := \Pi_{K,\mathcal{M}_{i-1}^\perp}(\mathcal{M}_j)$, which we call a *block* of the filtration. We also adopt the convention that $\mathcal{M}_0 = \{\mathbf{0}\}$ is the zero module.

Filtrations for module lattices over R are analogues of bases for lattices over $\mathbb{Z}$. Specifically, the basis $\boldsymbol{b}_1, \ldots, \boldsymbol{b}_d \in \mathbb{Q}^d$ of a lattice naturally 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*.

It is perhaps not immediately obvious that filtrations are nice to work with, or even that they always exist. So, we first note that they exist and can be found efficiently.

Fact 14. *For a CM/TR number field K, order $R \subseteq \mathcal{O}_K$, an inner product $\langle \cdot, \cdot \rangle_\rho$, and a module lattice $\mathcal{M} \subset K^\ell$ with rank k, there exists a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$.*

Furthermore, R-generating sets for the $\mathcal{M}_i$ can be computed efficiently (given an R-generating set for $\mathcal{M}$), and if ρ is semicanonical, $\det(\mathcal{M}) = \det(\widetilde{\mathcal{M}}_1) \cdots \det(\widetilde{\mathcal{M}}_k)$.

Proof. Let $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_m \in K^\ell$ be an R-generating set for $\mathcal{M}$, and suppose without loss of generality that $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_k$ are linearly independent over K. We take $\mathcal{M}_i := \mathcal{M} \cap \mathrm{span}_K(\boldsymbol{y}_1, \ldots, \boldsymbol{y}_i)$. An R-generating set for $\mathcal{M}_i$ can be computed by finding a $\mathbb{Z}$-basis for $\mathcal{M}_i$ (as a lattice) and then noting that a $\mathbb{Z}$-basis is also an R-generating set.

The fact about the determinants follows from Item 2 in Corollary 6. □

Finally, given a semi-canonical inner product $\langle \cdot, \cdot \rangle_\rho$ and R that is closed under conjugation, 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 $\Pi_{K,\mathcal{M}_{k-1}^\perp}(\mathcal{M})^* \subset \Pi_{K,\mathcal{M}_{k-2}^\perp}(\mathcal{M})^* \subset \cdots \subset \Pi_{K,\mathcal{M}_1^\perp}(\mathcal{M})^* \subset \mathcal{M}^*$, where $\Pi_{K,\mathcal{M}_i^\perp}(\mathcal{M})^*$ is a module lattice with

rank $k-i$. Equivalently, the dual filtration is given by $\mathcal{M}^* \cap \mathcal{M}_{k-1}^\perp \subset (\mathcal{M}^* \cap \mathcal{M}_{k-2}^\perp) \subset \cdots \subset (\mathcal{M}^* \cap \mathcal{M}_1^\perp) \subset \mathcal{M}^*$. In particular, the R-Gram-Schmidt orthogonalization of the dual filtration is the dual of 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 $\beta = 2$ and when the number field is sufficiently nice. The results here are strictly generalized by and subsumed by those presented in the full version [MS20], and the proofs have many common features. (Our proofs are also essentially the same as those in [LPSW19].)

Recall that we denote blocks of the filtration $\mathcal{M}_1 \subset \ldots \subset \mathcal{M}_k = \mathcal{M}$ as $\mathcal{M}_{[i,j]} = \Pi_{K,\mathcal{M}_{i-1}^\perp}(\mathcal{M}_j)$, and rank-one projections as $\widetilde{\mathcal{M}}_i = \Pi_{K,\mathcal{M}_{i-1}^\perp}(\mathcal{M}_i)$.

Definition 15 (DIP reduction). *For a CM/TR number field K, an order $R \subseteq \mathcal{O}_K$, an inner product $\langle \cdot, \cdot \rangle_\rho$, and approximation factor $\gamma \geq 1$, a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of a module $\mathcal{M}$ over R is γ-DIP-reduced if* $\det(\mathcal{M}_1)^{1/n} \leq \gamma \cdot \tau_1(\mathcal{M})$.

Definition 16 (γ-reduced filtration). *For a CM/TR number field K, an order $R \subseteq \mathcal{O}_K$, an inner product $\langle \cdot, \cdot \rangle_\rho$, and approximation factor $\gamma \geq 1$, a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ of a module $\mathcal{M}$ over R 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 17. *For a CM/TR number field K, an order $R \subseteq \mathcal{O}_K$, approximation factor $\gamma = \gamma(R,k) \geq 1$, a semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, and a γ-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$, we have*

$$\det(\mathcal{M}_1)^{1/n} \leq (\gamma \mu_{R,2})^{2(i-1)} \det(\widetilde{\mathcal{M}}_i)^{1/n} ,$$

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,

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

where the last equality follows from Fact 14 (since ρ is semicanonical). Rearranging, we see that $\det(\widetilde{\mathcal{M}}_i)^{1/n} \le (\gamma\mu_{R,2})^2 \det(\widetilde{\mathcal{M}}_{i+1})^{1/n}$. By a simple induction argument, we see that $\det(\mathcal{M}_1)^{1/n} \le (\gamma\mu_{R,2})^{2(i-1)} \det(\widetilde{\mathcal{M}}_i)^{1/n}$. □

Lemma 18. *For a CM/TR number field K, an order $R \subseteq \mathcal{O}_K$, an approximation factor $\gamma \ge 1$, and a semicanonical inner product $\langle\cdot,\cdot\rangle_\rho$, if a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is γ-reduced, then*

$$\det(\mathcal{M}_1)^{1/n} \le \gamma \cdot (\gamma\mu_{R,2})^{2(k-2)} \cdot \tau_1(\mathcal{M}) \ , \ \textit{and} \tag{6}$$

$$\det(\mathcal{M}_1)^{1/n} \le (\gamma\mu_{R,2})^{k-1} \cdot \det(\mathcal{M})^{1/(kn)} \ . \tag{7}$$

Proof. First, suppose that $\tau_1(\mathcal{M}_2) = \tau_1(\mathcal{M})$. Then, the result is immediate, from the fact that the filtration is γ-reduced, i.e., $\det(\mathcal{M}_1)^{1/n} \le \tau_1(\mathcal{M}_2) = \tau_1(\mathcal{M})$.

Otherwise, let $i \in [2, k-1]$ be such that $\tau_1(\mathcal{M}_{i+1}) = \tau_1(\mathcal{M})$ but $\tau_1(\mathcal{M}_{i-1}) \ne \tau_1(\mathcal{M})$. 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 $\det(\mathcal{M}')^{1/n} = \tau_1(\mathcal{M})$. Since $\mathcal{M}_{i-1}$ is primitive, $\mathcal{M}' \not\subset \operatorname{span}_K \mathcal{M}_{i-1}$. Therefore, $\Pi_{K,\mathcal{M}_{i-1}^\perp}(\mathcal{M}') \subset \mathcal{M}_{[i,i+1]}$ is a non-zero rank-one module lattice. It follows that

$$\tau_1(\mathcal{M}_{[i,i+1]}) \le \det(\Pi_{K,\mathcal{M}_{i-1}^\perp}(\mathcal{M}'))^{1/n} \le \det(\mathcal{M}')^{1/n} = \tau_1(\mathcal{M}) \ ,$$

where the second inequality is Item 4 of Corollary 6. Then, since the filtration is γ-reduced,

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

By combining the expression above with Lemma 17, we have

$$\det(\mathcal{M}_1)^{1/n} \le \gamma \cdot (\gamma\mu_{R,2})^{2(i-1)} \cdot \tau_1(\mathcal{M}) \ , \tag{8}$$

and recalling that $i \le k-1$, we obtain Eq. (6).

Again, recall from Lemma 17 that $\det(\mathcal{M}_1)^{1/n} \le (\gamma\mu_{R,2})^{2(i-1)} \det(\widetilde{\mathcal{M}}_i)^{1/n}$. Taking the product of these inequalities for $1 \le i \le k$, we see that

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

Raising both sides to the power $1/k$ yields Eq. (7). □

Corollary 19. *For a CM/TR number field K, an order $R \subseteq \mathcal{O}_K$, an approximation factor $\gamma = \gamma(R,k) \geq 1$, and a semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, if a filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ is γ-reduced, then*

$$\lambda_1(\mathcal{M}_1) \leq \frac{\gamma \delta_n}{\alpha_R} \cdot (\gamma \mu_{R,2})^{2(k-2)} \cdot \lambda_1(\mathcal{M}) \ , \ \textit{and} \tag{9}$$

$$\lambda_1(\mathcal{M}_1) \leq \delta_n (\gamma \mu_{R,2})^{(k-1)} \cdot \det(\mathcal{M})^{1/(kn)} \ . \tag{10}$$

Proof. By combining Eq. (6) from Lemma 18 with Lemma 7, we have

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

Using the definition of Hermite's constant δ_n with the above relation, we obtain Eq. (9):

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

Eq. (10) follows by directly applying the definition of Hermite's constant to Eq. (7) from Lemma 18. □

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 20 ((γ,k)-RFP). *For a CM/TR number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 1$, approximation factor $\gamma = \gamma(R,k) \geq 1$, and inner product $\langle \cdot, \cdot \rangle_\rho$, the (γ,k)-Reduced Filtration Problem, or (γ,k)-RFP, is the search problem defined as follows. The input is (a generating set for) a module lattice $\mathcal{M} \subset K^\ell$ 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 21. *For any CM/TR number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 2$, approximation factor $\gamma = \gamma(R,k) \geq 1$, semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, 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 $\mathcal{M} \subset K^\ell$ with

rank k, the reduction first computes a filtration $\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$ (as in Fact 14). It then repeatedly updates this filtration in place 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 submodule 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 $\widetilde{\mathcal{M}}_i = \widetilde{\mathcal{M}}'_i$ and leaves $\mathcal{M}_j$ unchanged for $j \neq i$. (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 $\Pi_{K,\mathcal{M}_{i-1}^{\perp}}(W) = \mathrm{span}_K(\widetilde{\mathcal{M}}'_i)$ and $\mathcal{M}_{i-1} \subset W$ and setting $\mathcal{M}_i := W \cap \mathcal{M}$. As we noted in Section 2.1, $\mathcal{M}_i$ will then be a primitive submodule with rank i, and it follows from the conditions on W that $\mathcal{M}_{i-1} \subset \mathcal{M}_i \subset \mathcal{M}_{i+1}$ and $\widetilde{\mathcal{M}}_i = \widetilde{\mathcal{M}}'_i$.)

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, $(1+\varepsilon)^n \det(\widetilde{\mathcal{M}}'_i) \geq \det(\widetilde{\mathcal{M}}_i)$.

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

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

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 13, $-\log(\Phi(\mathcal{M}_1, \ldots, \mathcal{M}_k))$ 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, since ρ

is semicanonical, by Item 2 of Corollary 6, we have

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

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 24.

Corollary 22. *For any CM/TR number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 2$, approximation factor $\gamma' = \gamma'(R,k) \geq 1$, semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, and constant $\varepsilon > 0$, there exists an efficient reduction from (γ, k)-DIP to $(\gamma', 2)$-DIP where*

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

Proof. The reduction takes as input a (generating set of a) module lattice $\mathcal{M}$ of rank k and runs the $((1+\varepsilon)\gamma', k)$-RFP procedure from Theorem 21, using the $(\gamma', 2)$-DIP oracle, receiving as output some $((1+\varepsilon)\gamma')$-reduced filtration $\mathcal{M}_1 \subset \cdots \subset \mathcal{M}_k = \mathcal{M}$ of $\mathcal{M}$. Finally, the reduction outputs $\mathcal{M}_1$.

Clearly, the reduction runs in polynomial time. By Eq. (6) from Lemma 18, we must have

$$\det(\mathcal{M}_1)^{1/n} \leq (1+\varepsilon)\gamma' \cdot ((1+\varepsilon)\gamma' \cdot \mu_{R,2})^{2(k-2)} \tau_1(\mathcal{M}) = \gamma \tau_1(\mathcal{M}) ,$$

as needed. □

Corollary 23. *For any CM/TR number field K, order $R \subseteq \mathcal{O}_K$ closed under conjugation, rank $k \geq 2$, approximation factor $\gamma' = \gamma'(R,k) \geq 1$, semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, 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 reduction takes as input a (generating set of a) module lattice $\mathcal{M}$ of rank k. It then runs the procedure from Theorem 21 with $\gamma := \gamma'\delta_n/\alpha_R$. Each time that this procedure requires a call to its $(\gamma, 2)$-DIP procedure, it uses the procedure from Theorem 12 and its $(\gamma', 2)$-ModuleSVP oracle to solve the $(\gamma, 2)$-DIP instance.

Clearly, the reduction runs in polynomial time and outputs a γ_R-reduced filtration of $\mathcal{M}$, where $\gamma_R = (1+\varepsilon)\gamma = (1+\varepsilon)\frac{\gamma'\delta_n}{\alpha_R}$. □

Theorem 24 (Main Theorem). *For any CM/TR number field K, order $R \subseteq \mathcal{O}_K$, rank $k \geq 2$, approximation factor $\gamma = \gamma(R,k) \geq 1$, semicanonical inner product $\langle \cdot, \cdot \rangle_\rho$, and constant $\varepsilon > 0$, there is an efficient reduction from (γ, k)-ModuleSVP to $(\gamma', 2)$-ModuleSVP where*

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

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

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

Proof. In fact, the reduction is the same for both ModuleSVP and ModuleHSVP. On input (a generating set for) a module lattice $\mathcal{M} \subset K^\ell$ with rank k, the reduction proceeds as follows. It obtains a γ_R-reduced filtration $\mathcal{M}_1 \subset \mathcal{M}_2 \subset \cdots \subset \mathcal{M}_k$ using its $(\gamma', 2)$-ModuleSVP oracle, where $\gamma_R := (1+\varepsilon)\frac{\gamma' \delta_n}{\alpha_R}$ (by Corollary 23). It then calls its $(\gamma', 2)$-ModuleSVP on $\mathcal{M}_2$ which outputs a vector $\boldsymbol{x}$ such that $0 < \|\boldsymbol{x}\|_\rho \leq \gamma' \lambda_1(\mathcal{M}_2)$. It then simply outputs this vector.

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

$$0 < \|\boldsymbol{x}\|_\rho \leq \gamma' \lambda_1(\mathcal{M}_2) \leq \gamma' \lambda_1(\mathcal{M}_1) .$$

By Eq. (9) of Corollary 19,

$$\begin{aligned} \lambda_1(\mathcal{M}_1) &\leq \frac{\gamma_R \delta_n}{\alpha_R} \cdot (\gamma_R \mu_{R,2})^{2(k-2)} \cdot \lambda_1(\mathcal{M}) \\ &= \frac{(1+\varepsilon)\gamma' \delta_n^2}{\alpha_R^2} \cdot \left((1+\varepsilon)\frac{\gamma' \delta_n}{\alpha_R}\mu_{R,2}\right)^{2(k-2)} \cdot \lambda_1(\mathcal{M}) \end{aligned}$$

Combining the above two expressions, we get

$$0 < \|\boldsymbol{x}\|_\rho \leq \frac{(1+\varepsilon)\gamma'^2 \delta_n^2}{\alpha_R^2} \cdot \left((1+\varepsilon)\frac{\gamma' \delta_n}{\alpha_R}\mu_{R,2}\right)^{2(k-2)} \cdot \lambda_1(\mathcal{M}) .$$

Therefore,

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

as needed.

Similarly, by Eq. (10) of Corollary 19,

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

which gives the reduction from ModuleHSVP. □

Bibliography

[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.

[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 *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 *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.

[KL17] Taechan Kim and Changmin Lee. Lattice reductions over Euclidean rings with applications to cryptanalysis. In *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.

[LPL18] S. Lyu, C. Porter, and C. Ling. Performance limits of lattice reduction over imaginary quadratic fields with applications to compute-and-forward. In *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.

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

[MS20] Tamalika Mukherjee and Noah Stephens-Davidowitz. Lattice reduction for modules, or how to reduce ModuleSVP to ModuleSVP. In *CRYPTO*, 2020. https://eprint.iacr.org/2019/1142.

[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.

[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 *Eurocrypt*, 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.

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

[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.