

On Solving Ideal-SVP and Ideal-BDD with Group Representations

HongyuanQu^{1,2}[0000–0002–1684–0819] and Guangwu
Xu^{1,2}[0000–0001–6200–3264] (✉)

¹ Key Laboratory of Cryptologic Technology and Information Security of
Ministry of Education, Qingdao, Shandong, 266237, China

qhy1qhy@mail.sdu.edu.cn, gxu4sdq@sdu.edu.cn

² School of Cyber Science and Technology, Shandong University,
Qingdao, Shandong, 266237, China.

Abstract. Ideal lattices over number fields play a central role in post-quantum cryptography, as evidenced by the NIST-standardized schemes Kyber and Dilithium. Understanding the hardness of ideal-lattice problems such as the Unique Shortest Vector Problem (Id-uSVP) and Bounded Distance Decoding (Id-BDD) is therefore essential. Recent works have shown that certain algebraic symmetries can significantly reduce the dimension of the Ideal-SVP problem. In particular, subfield attacks exploit the decomposition group to solve SVP in a smaller subfield. However, existing reductions are limited to SVP and rely on multiplicatively closed substructures.

In this work, we introduce a representation-theoretic framework for solving Id-uSVP and Id-BDD. Let $K/\mathbb{Q}$ be an abelian extension whose Galois group G is known, and let I be an integral ideal of K stabilized by a nontrivial subgroup $H \subseteq G$. We use the irreducible rational representations of H to construct scaled projection operators q_i that decompose K into a direct sum of linear subspaces V_i . Applying these operators to I yields low-dimensional lattices $I_i = q_i(I)$. We show that an Id-BDD instance (t, I) can be solved by solving independent BDD instances on the I_i 's and recombining the solutions, provided the target is within a distance bounded by $\lambda_1(I)/(2|H|^2)$. For Id-uSVP, we prove that under mild uniqueness conditions a shortest vector of I lies in at least one of the nonzero I_i . In cyclotomic fields, the rotation property gives two stronger results: (i) every nonzero I_i contains a shortest vector of I , so solving SVP on the lowest-rank nonzero I_i suffices; and (ii) BDD is further accelerated by reducing a single instance to n low-rank instances on the same lowest-rank nonzero I_i .

To the best of our knowledge, our approach is the first to employ group representation theory directly in lattice cryptanalysis. It addresses the limitation of previous subfield methods that they cannot solve BDD, removes the need for multiplicative closedness, and remains at least as efficient as prior SVP reductions. These results refine the understanding of the hardness of structured lattice problems and have implications for the security analysis of cryptosystems with additional algebraic structure.

1 Introduction

Lattice-based cryptography has become one of the most important components of post-quantum cryptography. The area was initiated by Ajtai’s work on the Short Integer Solution problem (SIS) [1], showing that solving SIS on average is as hard as solving worst-case approximate SVP. Regev later introduced the Learning with Errors problem (LWE) and proved a quantum reduction from worst-case approximate SVP/SIVP to the average-case LWE [22]; subsequent work replaced the quantum reduction by classical reductions under related assumptions [20, 7]. These results provide strong evidence that lattice-based schemes can resist quantum attacks, in contrast to schemes based on factoring or discrete logarithms, which are vulnerable to Shor’s algorithm. However, early lattice-based constructions were rather inefficient: they required large public keys, large ciphertexts, and expensive matrix-vector arithmetic over large moduli. This efficiency gap has motivated the use of lattices with additional algebraic structure.

The most widely adopted approach is to work with ideal lattices over number fields, especially those over power-of-two cyclotomic fields. In these fields the ring of integers can be represented as $R = \mathbb{Z}[X]/(X^n + 1)$, where n is a power of two. Multiplication in R can be computed in quasi-linear time using the Number Theoretic Transform (NTT), and ring elements require only linear space. The Ring-LWE and Ring-SIS problems, as well as their module variants, inherit these efficiency benefits while maintaining worst-case hardness reductions to structured lattice problems [17, 16, 15]. The practical importance of these structured lattices is reflected in the NIST post-quantum standardization process: the selected key-encapsulation mechanism Kyber (ML-KEM) [5] and the signature scheme Dilithium (ML-DSA) [10] are built on module-lattice problems over power-of-two cyclotomic rings. Other prominent schemes, such as Falcon [11] and NTRU-based constructions [13, 3], rely on closely related structured lattices. Consequently, understanding the hardness of ideal-lattice problems is a central task for post-quantum security.

However, the algebraic structure of ideal lattices is a double-edged sword: it enhances efficiency but also provides attackers with a richer toolkit, including automorphisms, subfields, ideal factorization, and norms. Hence, it is therefore important to understand the hardness of the ideal-lattice analogues of fundamental problems. In this work, we focus on two such variants: the Unique Shortest Vector Problem (uSVP) and Bounded Distance Decoding (BDD). In uSVP, one seeks the unique shortest nonzero lattice vector under a guarantee of significant separation from non-parallel vectors; in BDD, given a target point within a bounded distance from the lattice, the goal is to recover its nearest lattice vector. When instantiated over ideal lattices, these become Id-uSVP and Id-BDD. Their significance is multifold: Ring-LWE and Module-LWE rely on the worst-case hardness of ideal/module lattice problems like Id-SVP; NTRU-type schemes naturally reduce to uSVP or BDD in structured polynomial rings; and search LWE itself is fundamentally a BDD instance on q -ary lattices, whose ideal variants inherit the same interpretation. Consequently, progress in developing

attack methods for Id-uSVP and Id-BDD directly impacts both the theoretical foundations and the concrete security of structured lattice cryptography.

The relevance of these problems goes beyond direct attacks. As already pointed out in the literature, even if an efficient algorithm for Id-uSVP would not have a direct impact on the security of Ring-LWE, it would make the reduction to Id-uSVP vacuous, and hence let some interesting regime of Ring-LWE without lower bound security guarantees [6]. In addition, several cryptographic problems are themselves instances of Id-BDD; a notable example is the partial Vandermonde knapsack problem, which can be formulated as a BDD problem in an ideal lattice [9]. These observations motivate the search for new and more powerful algorithms for Id-uSVP and Id-BDD.

1.1 Previous Works

A major line of work has investigated conditions under which Ideal-SVP becomes significantly easier due to algebraic structure. Inspired by Bernstein’s logarithm-subfield attack [4], Albrecht, Bai and Ducas [2] and Cheon, Jeong and Lee [8] independently proposed subfield attacks in 2016 against overstretched NTRU, where the modulus is much larger than in standard NTRUEncrypt. Kirchner and Fouque [14] later refined these attacks and improved their practical performance. In 2021, Pan, Xu, Wadleigh and Cheng introduced a fundamentally new technique for solving SVP in prime ideals and general ideals over power-of-two cyclotomic fields [19]. Their method exploits the decomposition group of prime ideals to project the problem into a subfield, thereby dramatically reducing the lattice dimension and yielding a polynomial approximation factor. This approach was extended by Porter, Mendelsohn and Ling [21] to a broader class of ideals, whose characterization is more involved and depends on the factorization properties of the ideal and its algebraic norm. Boudgoust, Gachon and Pellet-Mary [6] further improved these results by obtaining a very simple sufficient condition for an ideal to be a “vulnerable ideal”; the resulting class strictly contains those of [19] and [21]. They also applied their algorithm to solve the decisional Partial Vandermonde Knapsack problem.

In 2024, Das and Joux [9] exploited the cyclic symmetry of $\mathbb{Z}[X]/(X^n \pm 1)$ to solve the search version of the Partial Vandermonde Knapsack problem, which is essentially an Ideal-BDD instance. Their method decomposes the original BDD instance into several low-dimensional BDD instances via the cyclic structure, and shows that the Partial Vandermonde Knapsack problem fails to provide the intended security for certain parameter sets.

1.2 Our Contributions

The approach of [6] reduces approximate SVP on an H -stabilized ideal I to approximate SVP on $I \cap K^H$, where K^H is the fixed field of H . Since $[K^H : \mathbb{Q}] = n/|H|$, this lowers the dimension from n to $n/|H|$. Under certain non-ramification conditions, they prove that $(I \cap K^H) \cdot \mathcal{O}_K = I$, so no information is lost. This is an elegant result. However, this technique has two limitations: the lift

of a short vector from $I \cap K^H$ to K increases its length by a factor at most $\sqrt{|H|}$, and the resulting vector is not necessarily a shortest vector of I ; moreover, this technique cannot be applied to BDD on I . Consequently, this approach only handles the decisional variant of the Partial Vandermonde Knapsack problem indirectly via a dual perspective, and it does not directly solve the search variant, which is essentially a BDD instance. [9] insightfully proposed a method that directly solves the search version of the Partial Vandermonde Knapsack problem, essentially reducing the BDD problem on I to several BDD problems on $I \cap K^H$. However, their method relies not only on the fact that I is stabilized by H , but also on the cyclic symmetry of $\mathbb{Z}[X]/(X^N + 1)$, where N is a power of two.

The following table summarizes the main limitations of prior approaches and highlights the advantages of our framework.

Table 1. Comparison of our contributions with prior work

Method	Solves search Id-BDD?	Id-SVP type	Required condition	Rank reduction
[6]	No	approx	non-ramification conditions	$n/ H $
[9]	Yes	—	prime/power-of-two cyclotomic fields	$n/ H $
Ours (general)	Yes	exact	abelian extensions	$< n$
Ours (cyclotomic)	Yes	exact	general cyclotomic extensions	$\leq n/ H $

In this paper, we focus on the setting that is most relevant to cryptographic applications: abelian extensions (which include all cyclotomic fields). Under this standard assumption, we present a general algorithm for solving BDD on ideal lattices, requiring only that the ideal I is stabilized by H and that $K/\mathbb{Q}$ is abelian. Unlike prior work, our method imposes no non-ramification conditions whatsoever. To the best of our knowledge, this is the first result that directly solves search BDD on ideals in general abelian extensions without relying on additional algebraic constraints such as power-of-two cyclotomic structure or ramification assumptions. Consequently, our method greatly broadens the scope of previous approaches within the abelian setting.

The key observation underlying this generalization is as follows: all previous reductions ultimately invoke BKZ on the reduced lattice, and BKZ exploits only the linear structure of the lattice, making no use of its multiplicative (ideal) structure. Therefore, the multiplicative closedness of the reduced substructure is not essential for the final lattice reduction step. This motivates us to adopt a more general framework based on group representation theory. Concretely, if we temporarily ignore the multiplicative structure, K is simply an n -dimensional $\mathbb{Q}$ -vector space, on which each element of H acts as an invertible $\mathbb{Q}$ -linear transformation, thereby yielding a rational representation $\psi : H \rightarrow \text{GL}_{\mathbb{Q}}(K)$. In this work, we focus on the case where $K/\mathbb{Q}$ is abelian (so both G and H are abelian),

since the representation theory of finite abelian groups is particularly elegant and well-behaved.

We begin by establishing that all irreducible rational representations $\phi_1, \dots, \phi_k$ of H , along with their characters $\chi_1, \dots, \chi_k$, are computable in polynomial time. With these characters, we define the scaled projection operators

$$q_i = \sum_{\sigma \in H} \chi_i(\sigma) \sigma, \quad i = 1, \dots, k.$$

Let $V_i = \text{Im}(q_i) = q_i(K)$. These subspaces yield the canonical decomposition $K = V_1 \oplus \dots \oplus V_k$ under the action of H . Among these components, exactly one is the fixed field K^H ; without loss of generality, we identify it as $V_1 = K^H$. The operators satisfy the fundamental identity

$$\sum_{i=1}^k q_i = |H| \cdot \text{id}_K,$$

which implies that for every $\alpha \in K$, $\sum_{i=1}^k q_i(\alpha) = |H| \cdot \alpha$.

The above identity suggests a natural reduction for solving BDD on the ideal I . For each component i , we consider the projected lattice $I_i := q_i(I)$. Given a target $t \in K$, we solve the BDD instances $(q_i(t), I_i)$ for $i = 1, \dots, k$, obtaining solutions $v_1, \dots, v_k$, and then reconstruct the global solution as

$$v = \frac{1}{|H|} \sum_{i=1}^k v_i.$$

We prove that whenever $\text{dist}(t, I) \leq \lambda_1(I)/(2|H|^2)$ and each sub-instance is solved correctly, the reconstructed v is indeed the closest vector in I to t . This yields a valid reduction for every $\frac{1}{\gamma}$ -Id-BDD instance with $\gamma \geq 2|H|^2$. Thus, whenever $|H|$ is constant, this gives a reduction for constant γ . Since smaller γ means harder instances, our algorithm can solve even those with small constant approximation factors.

Furthermore, under the uSVP assumption with $\gamma \geq |H|^2$, we prove that at least one of the components $I_1, \dots, I_t$ contains a shortest nonzero vector of I . To the best of our knowledge, this observation has not been reported in previous work. In particular, the trivial component I_1 is contained in $K^H \cap I$ and has rank $n/|H|$, which coincides with the rank of the lattice considered in [6]. Thus, for the approximate Ideal-SVP problem studied there, solving a single SVP instance on I_1 already matches the efficiency of [6]. Moreover, by solving SVP on all nonzero components and taking the shortest returned vector, our method recovers an exact shortest vector, whereas [6] only solves the approximate version. Additionally, since every component has rank strictly less than n , the worst-case complexity is exponentially smaller than that of the original rank- n problem.

When K is a cyclotomic field, the techniques of [9] can be adapted to further accelerate our Id-BDD algorithm. The reason is that, in this case, for every

$\alpha \in \mathcal{O}_K$ and every $i \in \mathbb{Z}$, we have $\|\alpha\| = \|\zeta^i \alpha\|$, where ζ is a primitive root of unity ($K = \mathbb{Q}(\zeta)$). This property allows us to reduce an Ideal-BDD instance (t, I) to the following n BDD subinstances:

$$(q_j(t), I_j), (q_j(\zeta t), I_j), \dots, (q_j(\zeta^{n-1}t), I_j),$$

where I_j has rank at most $n/|H|$. Finally, using the solutions $v_0, \dots, v_{n-1}$ obtained from these instances, we construct a linear system with a unique solution, from which the closest lattice vector v to t can be recovered via Gaussian elimination.

The same framework also applies to the γ -uSVP problem on ideals in cyclotomic fields. Specifically, when $\gamma \geq |H|^2$, we prove that every nonzero component among $I_1, \dots, I_t$ contains a shortest nonzero vector of I . To the best of our knowledge, this observation has not appeared in previous work. Consequently, it suffices to solve SVP only on the nonzero component I_j of minimal rank; the resulting vector is guaranteed to be a shortest nonzero vector of I , and moreover $\text{rank}(I_j) \leq n/|H|$. Detailed proofs are deferred to Section 4.

In summary, compared with prior subfield-based methods, our representation-theoretic framework offers the following key advantages:

- **First use of group representations.** To the best of our knowledge, this is the first time irreducible rational representations of the stabilizer group H are directly used to decompose and solve ideal-lattice problems.
- **No multiplicative closedness required.** The subspaces V_i need not be subrings or ideals; only the $\mathbb{Q}$ -linear action of H is used.
- **General reduction for BDD.** For general abelian extensions, a BDD instance (t, I) is reduced to several lower-dimensional BDD instances $(q_i(t), I_i)$ on the components of the canonical decomposition; solving these independently and recombining the solutions recovers the closest vector. This reduction applies to arbitrary abelian extensions and any subgroup H , and reduces the rank from n to strictly less than n . To the best of our knowledge, this is the first such reduction for BDD on ideals in general abelian extensions.
- **New observations for uSVP.** For uSVP on ideals stabilized by H , under the γ -uSVP promise with $\gamma \geq |H|^2$, a shortest vector of I appears in at least one of the nonzero projections I_i ; hence, by solving SVP on the nonzero components and taking the shortest returned vector, one recovers a shortest vector of I . Since each component has rank $< n$, this yields an exponential speedup over the original rank- n problem. To the best of our knowledge, this observation has not been reported in previous work. Our framework also covers the setting of [6] as a special case: in their approximate Ideal-SVP setting, it suffices to solve the trivial component I_1 of rank $n/|H|$, achieving at least the same dimension reduction while producing a sparser lattice.
- **Further speedups for cyclotomic rings.** When $K = \mathbb{Q}(\zeta)$ is a cyclotomic field, the additional rotation property $\|\alpha\| = \|\zeta^i \alpha\|$ leads to even stronger reductions. For Id-BDD, a single instance (t, I) is transformed into n instances $(q_j(\zeta^i t), I_j)$, all on the same low-rank lattice I_j with rank at most $n/|H|$; their solutions are combined via a uniquely solvable linear system to

recover the original closest vector. For Id-uSVP, under the γ -uSVP promise with $\gamma \geq |H|^2$, every nonzero component I_i contains a shortest vector of I , so it suffices to solve SVP on the component of minimal positive rank, which is at most $n/|H|$. These refinements provide an additional speedup beyond the general case.

2 Preliminaries

For every positive integer n , let S_n denote the symmetric group on n elements. Matrices are denoted by uppercase letters, and vectors by lowercase boldface letters; all vectors are column vectors. For a square matrix M , $\det(M)$ denotes its determinant. For $\mathbf{v} \in \mathbb{C}^n$, $\|\mathbf{v}\|$ denotes its Hermitian norm. We sometimes write $(a_1, \dots, a_n)$ for a column vector. Given vectors $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathbb{C}^n$, $\text{Span}_{\mathbb{R}}(\mathbf{x}_1, \dots, \mathbf{x}_m)$ denotes the subspace spanned by all real linear combinations of $\mathbf{x}_1, \dots, \mathbf{x}_m$; $\text{Span}_{\mathbb{R}}(S)$ is defined similarly when S is a set. For $r > 0$, let $\mathcal{B}_n(r)$ denote the n -dimensional ball of radius r . For a vector space V , let id_V denote the identity map on V . A linear transformation p on V is called a projection if $p^2 = p$.

Let V be a finite-dimensional vector space over a field K , and let L be an extension field of K . The scalar extension of V is the tensor product $L \otimes_K V$, whose elements are finite sums of the form $\sum_i \alpha_i \otimes \mathbf{v}_i$, where $\alpha_i \in L$ and $\mathbf{v}_i \in V$. The space $L \otimes_K V$ has a natural L -vector space structure, with scalar multiplication defined by $\alpha \cdot (\alpha' \otimes \mathbf{v}) := (\alpha\alpha') \otimes \mathbf{v}$. If $\mathbf{v}_1, \dots, \mathbf{v}_n$ is a basis of V , then $1 \otimes \mathbf{v}_1, \dots, 1 \otimes \mathbf{v}_n$ is a basis of $L \otimes_K V$. Thus scalar extension provides a natural way to enlarge the field of scalars of V from K to L .

2.1 Number Fields

In this subsection we introduce some necessary facts from algebraic number theory; for further details we refer to [18].

Every number field K is a finite simple extension of $\mathbb{Q}$; that is, there exists $\alpha \in \mathbb{C}$ such that $K = \mathbb{Q}(\alpha)$. Here α is an algebraic number, meaning that there exists an irreducible polynomial $f(X) \in \mathbb{Z}[X]$ with $f(\alpha) = 0$. In this case K can be equivalently represented as $K = \mathbb{Q}[X]/(f(X))$. This extension is denoted by $K/\mathbb{Q}$. For elements of K , we sometimes identify elements of K with their polynomial representatives when no confusion arises. The degree of the extension $K/\mathbb{Q}$ is defined as the dimension of K as a $\mathbb{Q}$ -vector space, denoted by $[K : \mathbb{Q}]$. If $L \subseteq K$ is a subfield of K , the degree $[K : L]$ of the extension K/L is defined similarly as the dimension of K as an L -vector space. Moreover, one has $[K : \mathbb{Q}] = [K : L] \cdot [L : \mathbb{Q}]$.

2.2 Galois Extensions and Abelian Extensions

Let $K/\mathbb{Q}$ be a number field of degree n . Then there exist exactly n distinct field embeddings $\sigma_1, \dots, \sigma_n : K \rightarrow \mathbb{C}$. The extension $K/\mathbb{Q}$ is called Galois if the

images of all embeddings $\sigma_1, \dots, \sigma_n$ are contained in K , equivalently, if each σ_i is an automorphism of K . In this case we write $\text{Gal}(K/\mathbb{Q}) = \{\sigma_1, \dots, \sigma_n\}$ and call it the Galois group of $K/\mathbb{Q}$.

Let H be a subgroup of $\text{Gal}(K/\mathbb{Q})$. The fixed field of H is defined by

$$K^H = \{x \in K : \sigma(x) = x \text{ for all } \sigma \in H\}.$$

The fundamental theorem of Galois theory asserts that the maps

$$H \mapsto K^H, \quad L \mapsto \text{Gal}(K/L)$$

are inverse to each other, thereby establishing a one-to-one correspondence between the subgroups of $\text{Gal}(K/\mathbb{Q})$ and the intermediate fields of $K/\mathbb{Q}$. If, in addition, $\text{Gal}(K/\mathbb{Q})$ is abelian, then $K/\mathbb{Q}$ is called an abelian extension. In this paper we mainly consider the case where $K/\mathbb{Q}$ is an abelian extension. The cyclotomic fields used in lattice-based cryptography are all abelian extensions.

2.3 Canonical embedding

For a number field K , write $n = [K : \mathbb{Q}]$, and let $\sigma_1, \dots, \sigma_n$ denote all embeddings of K into $\mathbb{C}$. Define the canonical embedding $\pi : K \rightarrow \mathbb{C}^n$ by

$$x \mapsto (\sigma_1(x), \dots, \sigma_n(x)).$$

For any $\alpha \in K$, let $\|\alpha\|$ denote the Hermitian norm of the vector $\pi(\alpha)$, i.e., $\|\alpha\| := \|\pi(\alpha)\|$.

If $K/\mathbb{Q}$ is a Galois extension, then each σ_i is an automorphism of K . For any $\alpha \in K$ and any σ_i , we have

$$\|\sigma_i(\alpha)\| = \|\alpha\|,$$

since composition with σ_i on the right permutes the set $\{\sigma_1, \dots, \sigma_n\}$, so the coordinates of $\pi(\sigma_i(\alpha))$ are obtained by permuting those of $\pi(\alpha)$, and hence their Hermitian norms coincide.

2.4 Ideal Lattices

An element $\alpha \in K$ is called an algebraic integer if there exists a monic polynomial $f(X) \in \mathbb{Z}[X]$ with integer coefficients such that $f(\alpha) = 0$. The set of all algebraic integers in K forms a ring, called the ring of integers of K , denoted by $\mathcal{O}_K$. Write $[K : \mathbb{Q}] = n$. Then $\mathcal{O}_K$ is a free $\mathbb{Z}$ -module of rank n , and any $\mathbb{Z}$ -basis of $\mathcal{O}_K$ is called an integral basis of K . Let $\alpha_1, \dots, \alpha_n$ be any integral basis of K , and let $\sigma_1, \dots, \sigma_n$ be all embeddings of K into $\mathbb{C}$. The discriminant of K is defined by

$$\Delta_K = (\det(\sigma_i(\alpha_j))_{i,j})^2.$$

Let $I \subseteq \mathcal{O}_K$ be an ideal of $\mathcal{O}_K$. Then I is also a free $\mathbb{Z}$ -module of rank n . The norm of I is defined as the index of I as a submodule of $\mathcal{O}_K$, that is, $N(I) = [\mathcal{O}_K : I]$. The image of I under the canonical embedding is denoted by

$\pi(I)$; it forms a rank- n lattice. Lattices arising in this way are called ideal lattices. A lattice basis is given by the image under the canonical embedding of any basis of I . The determinant of $\pi(I)$ is equal to $N(I) \cdot \sqrt{|\Delta_K|}$. Note that $\sqrt{|\Delta_K|}$ is the determinant of $\mathcal{O}_K$ regarded as an ideal lattice. Let $\lambda_1(I), \dots, \lambda_n(I)$ denote the successive minima of $\pi(I)$ with respect to the Hermitian norm.

2.5 Hard Problems on Ideal Lattices

In this paper we mainly consider the following problems on ideal lattices: uSVP and BDD. For the shortest vector problem, in some cases the shortest nonzero vector of an ideal lattice is not unique, and the shortest vectors are closely related to each other. For example, let I be an ideal lattice in a number field K , and let $\alpha \in I$ be a shortest nonzero vector. If σ is an automorphism of K with $\sigma(\alpha) \in I$, then $\sigma(\alpha)$ is also a shortest nonzero vector of I , since $\|\alpha\| = \|\sigma(\alpha)\|$. Another special case is that of cyclotomic fields. Let n be a positive integer. In an ideal lattice I of the cyclotomic field $\mathbb{Q}[\zeta_n]$, the shortest nonzero vector is not unique, because for every element α in the ring, we have

$$\|\alpha\| = \|\zeta_n \alpha\| = \dots = \|\zeta_n^{\varphi(n)-1} \alpha\|.$$

Recall that the norm here is the Hermitian norm of the vector under the canonical embedding. Therefore, in this case we have

$$\lambda_1(I) = \lambda_2(I) = \dots = \lambda_{\varphi(n)}(I).$$

To handle the above situations, we generalize the original definition of uSVP as follows.

Definition 1 (γ -Id-gap-uSVP $_K$). Let $\gamma \geq 1$, and let K be a number field of degree n with ring of integers $\mathcal{O}_K$. The γ -Id-gap-uSVP $_K$ problem is: given an ideal I of $\mathcal{O}_K$ such that for every $\alpha \in I \setminus \{0\}$, either $\|\alpha\| = \lambda_1(I)$ or $\|\alpha\| \geq \gamma \lambda_1(I)$, find any shortest nonzero vector of I .

The following is the definition of BDD on ideal lattices.

Definition 2 ($\frac{1}{\gamma}$ -Id-BDD $_K$). Let $\gamma > 2$ be a real number, and let K be a number field of degree n with ring of integers $\mathcal{O}_K$. Let I be an ideal of $\mathcal{O}_K$. The $\frac{1}{\gamma}$ -Id-BDD $_K$ problem is: given $\mathbf{t} \in \text{Span}_{\mathbb{R}}(\pi(I))$ such that $\mathbf{t} = \mathbf{v} + \mathbf{e}$, where $\mathbf{v} \in \pi(I)$ and $0 < \|\mathbf{e}\| \leq \frac{1}{\gamma} \lambda_1(I)$, find $\mathbf{v}$.

Note that this definition guarantees the uniqueness of the solution to $\frac{1}{\gamma}$ -Id-BDD $_K$.

We also need the definition of approximate SVP and BDD problems for general lattices.

Definition 3 (γ -SVP). Let $\gamma \geq 1$. The γ -SVP problem is: given a basis $\mathbf{B} \in \mathbb{C}^{n \times m}$ of a rank- m lattice L in n -dimensional space, find a nonzero $\mathbf{x} \in L$ such that $\|\mathbf{x}\| \leq \gamma \cdot \lambda_1(L)$.

Definition 4 ($\frac{1}{\gamma}$ -BDD). *Let $\gamma > 2$ be a real number. The $\frac{1}{\gamma}$ -BDD problem is: given a basis $\mathbf{B} \in \mathbb{C}^{n \times m}$ of a rank- m lattice L in n -dimensional space and a target vector $\mathbf{t} \in \text{Span}_{\mathbb{R}}(L)$ such that $\mathbf{t} = \mathbf{v} + \mathbf{e}$, where $\mathbf{v} \in L$ and $0 < \|\mathbf{e}\| \leq \frac{1}{\gamma} \lambda_1(L)$, find $\mathbf{v}$.*

Observe that as γ increases, all of the above problems, namely γ -Id-gap-uSVP $_K$ and $\frac{1}{\gamma}$ -Id-BDD $_K$ on ideal lattices, and γ -SVP and $\frac{1}{\gamma}$ -BDD on general lattices, become easier. Currently, there is a consensus in the cryptography community that for $\gamma \leq n^2$, all the above problems are hard in the worst case, in the sense that the best known attacks have exponential or subexponential asymptotic complexity.

The BKZ algorithm is a standard algorithm for solving γ -uSVP and $\frac{1}{\gamma}$ -BDD. The running time of this algorithm has been formally studied in [12], with the following conclusion.

Lemma 1 ([12, Theorem 1]). *There is a classical probabilistic algorithm that takes as input a basis $\mathbf{B} \in \mathbb{C}^{n \times m}$ of a lattice L of rank m and a parameter $\gamma = \text{poly}(m)$, and solves γ -uSVP and $\frac{1}{\gamma}$ -BDD in L in time $\text{poly}(n) \cdot 2^{O(\frac{m \log m}{\log \gamma})}$.*

It is easy to observe that when the dimension n of the ambient space and the rank m of the lattice are close, for instance $n = \text{poly}(m)$, the complexity of using BKZ to solve γ -uSVP and $\frac{1}{\gamma}$ -BDD on L is mainly determined by $\frac{m \log m}{\log \gamma}$.

2.6 Complex Representations and Canonical Decomposition of Finite Abelian Groups

This paper uses the theory of rational representations of finite abelian groups to solve hard problems over ideal lattices. In this section we introduce some basic facts about complex representations of finite abelian groups. These facts will be used later to develop the corresponding theory of rational representations. For a detailed treatment of complex representations, see [23].

Let G be a finite abelian group, let K be a field, and let V be an n -dimensional vector space over K . Denote by $\text{GL}(V)$ the set of all invertible linear transformations of V . Under composition of maps, $\text{GL}(V)$ forms a group. A group homomorphism $\phi: G \rightarrow \text{GL}(V)$ is called a K -representation of G . The vector space V is called the representation space of ϕ , and its dimension n is called the degree of ϕ , denoted by $\deg \phi$. Clearly, the map sending every element of G to the identity transformation of V is a K -representation of G , called the trivial representation.

Let ϕ and ψ be two K -representations of G , with representation spaces V and W , respectively. If there exists an invertible linear map $\sigma: V \rightarrow W$ such that

$$\psi(g) \circ \sigma = \sigma \circ \phi(g)$$

for every $g \in G$, then ϕ and ψ are said to be isomorphic representations.

After fixing a basis of V , we identify $\text{GL}(V)$ with $\text{GL}_n(K)$; accordingly, for each $g \in G$ we may regard $\phi(g)$ as an element of $\text{GL}_n(K)$. Let $\text{Tr}(\phi(g))$ denote the trace of $\phi(g)$. The function

$$\chi_\phi: G \rightarrow K, \quad g \mapsto \text{Tr}(\phi(g))$$

is called the character of the representation ϕ . In this paper we mainly consider the cases where K is the field of complex numbers $\mathbb{C}$ or the field of rational numbers $\mathbb{Q}$; the corresponding characters are then called complex characters and rational characters, respectively.

A subspace $U \subseteq V$ is called an invariant subspace of V if for every $g \in G$ and every $u \in U$, we have $\phi(g)(u) \in U$. In this case, the restriction of ϕ to U is a representation of G on U , denoted by ϕ_U , and is called a subrepresentation of ϕ . If the only invariant subspaces of V are $\{0\}$ and V , then the representation ϕ is said to be irreducible. The intersection of invariant subspaces is again an invariant subspace.

The representation ϕ is said to be completely reducible if for every invariant subspace U of V , there exists an invariant subspace W of V such that $V = U \oplus W$. In this case, for every $v \in V$ there exist unique elements $u \in U$ and $w \in W$ such that $v = u + w$, and for every $g \in G$,

$$\phi(g)(v) = \phi_U(g)(u) + \phi_W(g)(w).$$

With respect to a suitable basis, the matrix of $\phi(g)$ is then block diagonal:

$$\begin{bmatrix} \phi_U(g) & 0 \\ 0 & \phi_W(g) \end{bmatrix}.$$

We say that ϕ is the direct sum of ϕ_U and ϕ_W , and write $\phi = \phi_U \oplus \phi_W$. It is easy to verify that

$$\chi_\phi = \chi_{\phi_U} + \chi_{\phi_W}.$$

Direct sums of finitely many representations are defined similarly.

We have the following well-known results.

Theorem 1 (Maschke's Theorem). *Let K be a field of characteristic 0, and let G be a finite group. Then every K -representation of G is completely reducible.*

The following is a corollary of Maschke's theorem.

Corollary 1. *Let K be a subfield of $\mathbb{C}$, let G be a finite group, and let V be a finite-dimensional vector space over K . If $\phi: G \rightarrow \text{GL}(V)$ is a K -representation of G , then V can be decomposed into a direct sum of invariant subspaces.*

We also need the following form of Schur's lemma.

Lemma 2 (Schur's lemma [23, Chapter II, Proposition 4]). *Let $\phi_1: G \rightarrow \text{GL}(V_1)$ and $\phi_2: G \rightarrow \text{GL}(V_2)$ be two non-zero irreducible K -representations of G , and let $\sigma: V_1 \rightarrow V_2$ be a linear map satisfying*

$$\sigma \circ \phi_1(g) = \phi_2(g) \circ \sigma, \quad \forall g \in G.$$

Then

1. if ϕ_1 is not isomorphic to ϕ_2 , then $\sigma = 0$;
2. if $V_1 = V_2$ and $\phi_1 = \phi_2$, then either $\sigma = 0$ or σ is an invertible linear map. In particular, if K is algebraically closed, then σ is a homothety, i.e., a scalar multiple of the identity.

Let G be a finite multiplicative abelian group. By the structure theorem for finitely generated abelian groups, G has the following elementary divisor decomposition:

$$G = \langle g_1 \rangle \times \cdots \times \langle g_n \rangle \cong \bigoplus_{j=1}^n \mathbb{Z}_{p_j^{a_j}}, \quad (1)$$

where the orders of $g_1, \dots, g_n \in G$ are $p_1^{a_1}, \dots, p_n^{a_n}$ respectively, with $p_1, \dots, p_n$ primes and $a_1, \dots, a_n$ positive integers.

Every irreducible complex representation of G is one-dimensional. In particular, for each $\mathbf{k} \in \mathbb{Z}_{p_1^{a_1}} \times \cdots \times \mathbb{Z}_{p_n^{a_n}}$, define

$$\phi_{\mathbf{k}}: G \rightarrow \mathbb{C}^\times, \quad \prod_{j=1}^n g_j^{e_j} \mapsto \prod_{j=1}^n e^{2\pi i k_j e_j / p_j^{a_j}}.$$

Then $\phi_{\mathbf{k}}$ is an irreducible complex representation of G . Let

$$\tilde{G} = \{\phi_{\mathbf{k}} : \mathbf{k} \in \mathbb{Z}_{p_1^{a_1}} \times \cdots \times \mathbb{Z}_{p_n^{a_n}}\}.$$

Then $\tilde{G}$ is the set of all irreducible complex representations of G . Under multiplication, $\tilde{G}$ forms a group, and $G \cong \tilde{G}$. The trivial one-dimensional complex representation of G is the identity element of $\tilde{G}$, and its character is identically 1.

If the elementary divisor form of G is known, then all irreducible complex representations of G and their complex characters can be computed in time polynomial in $|G|$.

For each $\mathbf{k} \in \mathbb{Z}_{p_1^{a_1}} \times \cdots \times \mathbb{Z}_{p_n^{a_n}}$, let $\chi_{\mathbf{k}}$ denote the complex character of $\phi_{\mathbf{k}}$. The distinct characters satisfy the following orthogonality relation:

$$\frac{1}{|G|} \sum_{g \in G} \overline{\chi_{\mathbf{k}}(g)} \chi_{\mathbf{\ell}}(g) = \begin{cases} 1, & \text{if } \mathbf{k} = \mathbf{\ell}, \\ 0, & \text{otherwise.} \end{cases}$$

We denote these irreducible characters by $\chi_1, \dots, \chi_{|G|}$. For any complex representation ψ of G , there exist unique nonnegative integers $m_1, \dots, m_{|G|}$ such that

$$\chi_{\psi} = m_1 \chi_1 + \cdots + m_{|G|} \chi_{|G|}.$$

Using the orthogonality relations, we obtain

$$m_j = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_{\psi}(g)} \chi_j(g), \quad j = 1, \dots, |G|.$$

We shall need the following standard result.

Theorem 2. *Two complex representations ϕ and ψ of a finite abelian group G are isomorphic if and only if their characters coincide, i.e., $\chi_\phi(g) = \chi_\psi(g)$ for all $g \in G$.*

By Maschke's theorem and its corollary, the representation space V of a complex representation ψ can be decomposed into a finite direct sum of irreducible invariant subspaces:

$$V = U_1 \oplus \cdots \oplus U_s.$$

For each $i = 1, \dots, |G|$, let V_i denote the direct sum of all subspaces among $U_1, \dots, U_s$ that are isomorphic to ϕ_i . Then clearly

$$V = V_1 \oplus \cdots \oplus V_{|G|}.$$

This decomposition is called the canonical decomposition. Its main advantage is the following uniqueness property.

Theorem 3 ([23, Theorem 8]).

1. *The decomposition $V = V_1 \oplus \cdots \oplus V_{|G|}$ does not depend on the initial choice of decomposition of V into irreducible subrepresentations.*
2. *In this decomposition, the projection p_i of V onto V_i is given by*

$$p_i = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_i(g)} \psi(g). \quad (2)$$

This decomposition has the following key property: for each $i \in \{1, \dots, |G|\}$, the restriction of p_i to V_i is the identity map on V_i , and its restriction to $\bigoplus_{j \neq i} V_j$ is the zero map. Consequently, $\sum_{i=1}^{|G|} p_i$ is the identity map on V , i.e.,

$$\sum_{i=1}^{|G|} p_i = \frac{1}{|G|} \sum_{g \in G} \left(\sum_{i=1}^{|G|} \overline{\chi_i(g)} \right) \psi(g) = \text{id}_V. \quad (3)$$

3 Canonical Decomposition of Rational Representations of Finite Abelian Groups and Its Application to Solving Id-BDD and Id-uSVP Problems

In this section, we use the results on complex representations of finite abelian groups in Section 2.6 to establish the necessary facts about rational representations and their canonical decomposition. We then show that when an ideal I satisfies certain symmetry conditions, the canonical decomposition of rational representations can substantially reduce the complexity of solving the $\frac{1}{\gamma}$ -Id-BDD $_K$ and γ -Id-gap-uSVP $_K$ problems on I , where $\gamma \leq n^2$.

3.1 Canonical Decomposition of Rational Representations

Let G be a finite abelian group. We first use the complex representations and complex characters of G to determine all irreducible rational representations of G and their characters, and then construct the canonical decomposition of rational representations. Suppose that the elementary divisor form of G is known, as in (1). We keep the notation of the previous subsection: let ϕ_j and χ_j ($j = 1, \dots, |G|$) be all irreducible complex representations of G and their characters. Let n denote the order of G , i.e., $n = |G|$. Then for every $g \in G$, we have $g^n = 1$; hence for every $j = 1, \dots, |G|$ and every $g \in G$, we have $\chi_j(g)^n = 1$. Therefore $\chi_j(g) \in \mathbb{Q}[\zeta_n]$, where $\zeta_n = e^{2\pi i/n}$.

For any $\sigma \in \text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$, any $j \in \{1, \dots, |G|\}$, and any $g \in G$, since $\phi_j(g) \in \mathbb{Q}[\zeta_n]$, we may define a new representation $\sigma(\phi_j)$ by setting

$$(\sigma(\phi_j))(g) := \sigma(\phi_j(g)) \in \mathbb{Q}[\zeta_n].$$

It is straightforward to verify that $\sigma(\phi_j)$ is again a one-dimensional irreducible complex representation of G , and its character is given by

$$\chi_{\sigma(\phi_j)} = \sigma(\chi_j).$$

Since σ is an automorphism of $\mathbb{Q}[\zeta_n]$ and hence a bijection, it induces a permutation on the set $\tilde{G}$ of all irreducible complex representations of G . Thus there exists a permutation $\tau_\sigma \in S_{|G|}$ such that

$$\sigma(\chi_j) = \chi_{\tau_\sigma(j)}, \quad j = 1, \dots, |G|.$$

This defines an action of $\text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$ on $\tilde{G}$. Consequently, $\tilde{G}$ decomposes into a disjoint union of orbits. Let t be the number of orbits, and denote these orbits by $O_1, \dots, O_t$. In other words,

$$\tilde{G} = O_1 \sqcup \dots \sqcup O_t,$$

where each orbit O_i has the form

$$O_i = \{\sigma(\phi_j) : \sigma \in \text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})\}$$

for some representative ϕ_j .

Let $\phi: G \rightarrow \text{GL}(V)$ be an irreducible rational representation of G , and let $W = \mathbb{C} \otimes_{\mathbb{Q}} V$. Then ϕ extends naturally to a complex representation

$$\phi \otimes_{\mathbb{Q}} \mathbb{C}: G \rightarrow \text{GL}(W),$$

defined on pure tensors by

$$(\phi \otimes_{\mathbb{Q}} \mathbb{C})(g)(z \otimes v) = z \otimes \phi(g)(v)$$

for every $g \in G$ and $z \otimes v \in W$.

Let $v_1, \dots, v_d$ be a basis of V . Then $1 \otimes v_1, \dots, 1 \otimes v_d$ is a basis of W , and with respect to these two bases, the matrices of $\phi(g)$ and $(\phi \otimes_{\mathbb{Q}} \mathbb{C})(g)$ coincide for every $g \in G$. Hence

$$\chi_\phi = \chi_{\phi \otimes_{\mathbb{Q}} \mathbb{C}}.$$

Therefore χ_ϕ can be uniquely expressed as a nonnegative integer linear combination of the irreducible complex characters of G :

$$\chi_\phi = \sum_{j=1}^{|G|} m_j \chi_j, \quad m_j = \frac{1}{|G|} \sum_{g \in G} \overline{\chi_\phi(g)} \chi_j(g) \in \mathbb{Z}_{\geq 0}, \quad j = 1, \dots, |G|.$$

Since ϕ is a rational representation, for every $\sigma \in \text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$ we have $\sigma(\phi) = \phi$ and therefore $\sigma(\chi_\phi) = \chi_\phi$. On the other hand,

$$\sigma(\chi_\phi) = \sum_{j=1}^{|G|} m_j \sigma(\chi_j) = \sum_{j=1}^{|G|} m_j \chi_{\tau_\sigma(j)} = \sum_{j=1}^{|G|} m_j \chi_j = \sum_{j=1}^{|G|} m_{\tau_\sigma(j)} \chi_{\tau_\sigma(j)}.$$

Comparing coefficients, we obtain

$$m_j = m_{\tau_\sigma(j)} \quad (4)$$

for every $j \in \{1, \dots, |G|\}$ and every $\sigma \in \text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$. This shows that the multiplicities of irreducible complex characters lying in the same Galois orbit are equal.

We then have the following conclusion.

Theorem 4. *Let $\psi_j = \bigoplus_{\phi \in O_j} \phi$ and let $\theta_j = \sum_{\phi \in O_j} \chi_\phi$ for $j = 1, \dots, t$. Then $\psi_1, \dots, \psi_t$ are all non-isomorphic irreducible rational representations of G , and $\theta_1, \dots, \theta_t$ are their characters. Moreover,*

$$\frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \theta_j(g) = \delta_{ij} |O_i|, \quad i, j \in \{1, \dots, t\}. \quad (5)$$

Proof. We first prove that for each $j \in \{1, \dots, t\}$, ψ_j is a rational representation. This is equivalent to showing that there exists a rational representation $\rho_j: G \rightarrow \text{GL}(W_j)$ such that $\rho_j \otimes_{\mathbb{Q}} \mathbb{C} \cong \psi_j$.

Fix an arbitrary $\phi_0 \in O_j$, and set

$$S = \{\phi_0(g) : g \in G\}, \quad K = \mathbb{Q}(S).$$

Then $K \subseteq \mathbb{Q}[\zeta_n]$, where n is the order of G . Observe that every $x \in S$ is of the form $x = e^{2\pi i k/d}$ for some integers k, d with $\gcd(k, d) = 1$, so x is a primitive d -th root of unity. Consequently all primitive d -th roots of unity are powers of x , and therefore the minimal polynomial of x splits completely in K . Since this holds for every $x \in S$, we see that K is the splitting field of S over $\mathbb{Q}$.

Since $K \subseteq \mathbb{Q}[\zeta_n]$ and $\mathbb{Q}[\zeta_n]/\mathbb{Q}$ is a Galois extension, every element of $\text{Gal}(K/\mathbb{Q})$ is the restriction to K of some element of $\text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$. Let

$$O_j = \{\sigma_1(\phi_0), \dots, \sigma_{|O_j|}(\phi_0)\},$$

where $\sigma_1, \dots, \sigma_{|O_j|} \in \text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$. Then the restrictions $\sigma_1|_K, \dots, \sigma_{|O_j|}|_K$ belong to $\text{Gal}(K/\mathbb{Q})$, and they are all distinct because the representations $\sigma_\ell(\phi_0)$ are distinct elements of O_j . Therefore, for every $\ell \in \{1, \dots, |O_j|\}$ and every $g \in G$, we have $\sigma_\ell(\phi_0)(g) \in K$.

Conversely, let $\sigma \in \text{Gal}(K/\mathbb{Q})$. Then σ extends to an automorphism of $\mathbb{Q}[\zeta_n]$, and $\sigma(\phi_0)$ is still an irreducible complex representation of G . Since O_j is an orbit under $\text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$, we have $\sigma(\phi_0) \in O_j$, so $\sigma(\phi_0) = \sigma_\ell(\phi_0)$ for some ℓ . By the definition of K , an element of $\text{Gal}(K/\mathbb{Q})$ is uniquely determined by its action on the set S , so we must have $\sigma = \sigma_\ell|_K$. This shows that

$$\text{Gal}(K/\mathbb{Q}) = \{\sigma_1|_K, \dots, \sigma_{|O_j|}|_K\},$$

and hence $[K : \mathbb{Q}] = |O_j|$.

For each $g \in G$, multiplication by $\phi_0(g)$ defines a $\mathbb{Q}$ -linear transformation of K :

$$L_g : K \rightarrow K, \quad x \mapsto \phi_0(g)x.$$

Thus $\rho_j : G \rightarrow \text{GL}(K)$, $g \mapsto L_g$, is a rational representation of G . Moreover, for every $g \in G$,

$$\chi_{\rho_j}(g) = \text{Tr}(L_g) = \text{Tr}_{K/\mathbb{Q}}(\phi_0(g)) = \sum_{\ell=1}^{|O_j|} \sigma_\ell(\phi_0)(g) = \sum_{\phi \in O_j} \chi_\phi(g) = \theta_j(g).$$

By Theorem 2, we conclude that $\rho_j \otimes_{\mathbb{Q}} \mathbb{C} \cong \psi_j$.

Next we prove that ψ_j is irreducible over $\mathbb{Q}$. Suppose that ψ_j is reducible. By Maschke's theorem, there exists a nonempty proper subset $O'_j \subsetneq O_j$ such that

$$\psi_j = \left(\bigoplus_{\phi \in O'_j} \phi \right) \oplus \left(\bigoplus_{\psi \in O_j \setminus O'_j} \psi \right),$$

where both direct sums are rational representations. Without loss of generality, assume $\phi_0 \in O'_j$. For every $\ell \in \{1, \dots, |O_j|\}$,

$$\bigoplus_{\phi \in O'_j} \phi = \bigoplus_{\phi \in O'_j} \sigma_\ell(\phi).$$

Since the representations in O_j are pairwise non-isomorphic irreducible complex representations, this equality implies that O'_j is invariant under the action of $\text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$, which contradicts the fact that O_j is an orbit. Hence every ψ_j is an irreducible rational representation.

By (4), the $\psi_1, \dots, \psi_t$ account for all irreducible rational representations of G , because any rational representation decomposes into a direct sum of irreducible complex representations and its character must have equal multiplicities on each Galois orbit. The fact that they are pairwise non-isomorphic follows from the

orthogonality relation (5), which we now prove. Since the orbits $O_1, \dots, O_t$ are pairwise disjoint, the orthogonality of irreducible complex characters yields

$$\begin{aligned} \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \theta_j(g) &= \frac{1}{|G|} \sum_{g \in G} \left(\sum_{\phi \in O_i} \overline{\chi_\phi(g)} \right) \left(\sum_{\psi \in O_j} \chi_\psi(g) \right) \\ &= \sum_{\phi \in O_i} \sum_{\psi \in O_j} \left(\frac{1}{|G|} \sum_{g \in G} \overline{\chi_\phi(g)} \chi_\psi(g) \right) \\ &= \delta_{ij} |O_i|. \end{aligned}$$

This completes the proof. $\square$

Note that the orbit of the trivial one-dimensional complex representation consists only of itself, so G admits a trivial irreducible rational representation whose character is identically 1.

We continue to use the above notation. Let $\psi_1, \dots, \psi_t$ denote all irreducible rational representations of G , and let $\theta_1, \dots, \theta_t$ denote their characters.

Let $\psi: G \rightarrow \text{GL}(V)$ be an arbitrary finite-dimensional rational representation of G . By Maschke's theorem and its corollary, V can be decomposed into a direct sum of irreducible invariant subspaces: $V = U_1 \oplus \dots \oplus U_s$. For each $i = 1, \dots, t$, let V_i be the direct sum of all subspaces among $U_1, \dots, U_s$ that are isomorphic to ψ_i . Then

$$V = V_1 \oplus \dots \oplus V_t.$$

This decomposition is also called the canonical decomposition of V . The properties of the canonical decomposition for complex representations carry over naturally to rational representations.

Theorem 5. 1. *The decomposition $V = V_1 \oplus \dots \oplus V_t$ does not depend on the initial decomposition of V into irreducible subrepresentations.*
 2. *For $i = 1, \dots, t$, define*

$$p_i = \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \psi(g).$$

Then $V_i = \text{Im } p_i$ for each $i = 1, \dots, t$, and

$$\sum_{i=1}^t p_i = \text{id}_V. \quad (6)$$

Proof. The first assertion follows from the second, since the projections p_i do not depend on the choice of the decomposition. We now prove the second assertion. Fix $j \in \{1, \dots, t\}$ and assume $V_j \neq \{0\}$. Take any $U_k \subseteq V_j$ such that U_k is isomorphic to ψ_j . Then the restriction of $\psi(g)$ to U_k is equivalent to ψ_j . For

every $h \in G$,

$$\begin{aligned}
\psi_j(h) \circ (p_i|_{U_k}) \circ \psi_j(h)^{-1} &= \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \psi_j(h) \psi_j(g) \psi_j(h^{-1}) \\
&= \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \psi_j(hgh^{-1}) \\
&= \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(h^{-1}gh)} \psi_j(g) \\
&= \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \psi_j(g) \\
&= p_i|_{U_k}.
\end{aligned}$$

By Schur's lemma, if $i \neq j$, then $p_i|_{U_k} = 0$, and hence $p_i|_{V_j} = 0$. If $i = j$, then by Theorem 4,

$$\text{Tr}(p_i|_{U_k}) = \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \theta_i(g) = |O_i| \neq 0.$$

Thus $p_i|_{U_k}$ is invertible, and hence $p_i|_{V_i}$ is invertible.

Now let $W = \mathbb{C} \otimes_{\mathbb{Q}} V$, and let $\psi \otimes_{\mathbb{Q}} \mathbb{C}: G \rightarrow \text{GL}(W)$ be the scalar extension of ψ . For $i = 1, \dots, t$, set

$$q_i = \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} (\psi \otimes_{\mathbb{Q}} \mathbb{C})(g).$$

Then

$$\sum_{j=1}^t q_j = \frac{1}{|G|} \sum_{g \in G} \left(\sum_{j=1}^t \overline{\theta_j(g)} \right) (\psi \otimes_{\mathbb{Q}} \mathbb{C})(g).$$

By the definition of θ_j , the sum $\sum_{j=1}^t \overline{\theta_j(g)}$ is the sum of all irreducible complex characters of G . Therefore, by (3), $\sum_{j=1}^t q_j$ is the identity map on W . Since V is embedded in W and the projections p_i are the restrictions of q_i to V , it follows that $\sum_{j=1}^t p_j$ is the identity map on V . $\square$

We next show that, given the elementary divisor form of G , the canonical decomposition of any rational representation of G can be computed in polynomial time.

Theorem 6. *Let G be a finite abelian group of order n whose elementary divisor form is known, and let V be an m -dimensional vector space over $\mathbb{Q}$. For any rational representation $\psi: G \rightarrow \text{GL}(V)$, the canonical decomposition of V with respect to ψ can be computed in time $\text{poly}(m, n)$.*

Algorithm 1 Algorithm for computing the canonical decomposition

Require: A rational representation $\psi: G \rightarrow \text{GL}(V)$ and the elementary divisor form of G as in (1).

Ensure: The canonical decomposition $V = V_1 \oplus \cdots \oplus V_t$.

- 1: Using the elementary divisor form of G , compute all irreducible complex representations of G and their characters.
- 2: Compute $\text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})$, where $n = |G|$.
- 3: Use Theorem 4 to obtain all irreducible rational characters $\theta_1, \dots, \theta_t$ of G .
- 4: For $i = 1, \dots, t$, construct the projection map $p_i = \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \psi(g)$.
- 5: For $i = 1, \dots, t$, compute $V_i = \text{Im } p_i$.
- 6: **return** $V_1, \dots, V_t$.

Proof. Algorithm 1 computes the desired decomposition. Its first three steps determine all irreducible rational characters $\theta_1, \dots, \theta_t$ of G . Since $|G| = n$ and $|\text{Gal}(\mathbb{Q}[\zeta_n]/\mathbb{Q})| \leq n$, these steps require only $\text{poly}(n)$ arithmetic operations on values in cyclotomic fields. The fourth step constructs the projections

$$p_i = \frac{1}{|G|} \sum_{g \in G} \overline{\theta_i(g)} \psi(g), \quad i = 1, \dots, t.$$

Each $\psi(g)$ is a $m \times m$ matrix over $\mathbb{Q}$, and the sum involves n terms; hence each p_i is obtained in $\text{poly}(m, n)$ time, and there are at most n such projections. Finally, the fifth step computes $V_i = \text{Im } p_i$ using standard linear algebra, again in $\text{poly}(m, n)$ time. By Theorem 5, the resulting subspaces form the canonical decomposition of V . Thus the total running time is $\text{poly}(m, n)$. $\square$

3.2 Solving $\frac{1}{\gamma}$ -Id-BDD $_K$ via canonical decomposition of rational representations

In this subsection we use the theory developed in the previous subsection to show that when an ideal lattice satisfies certain symmetry conditions, the BDD problem on it can be solved much more efficiently by exploiting this symmetry and the canonical decomposition of rational representations.

Let $K/\mathbb{Q}$ be an abelian extension with known Galois group G , and set $n := [K : \mathbb{Q}]$, so that $|G| = n$. Let π denote the canonical embedding of K into $\mathbb{C}^n$. Let $I \subseteq \mathcal{O}_K$ be an ideal, and let $\mathbf{t} \in \text{Span}_{\mathbb{R}}(\pi(I))$ be a target vector such that $\mathbf{t} = \mathbf{v} + \mathbf{e}$ for some $\mathbf{v} \in \pi(I)$ and $0 < \|\mathbf{e}\| \leq \frac{1}{\gamma} \lambda_1(I)$. The goal of $\frac{1}{\gamma}$ -Id-BDD $_K$ is to recover $\mathbf{v}$.

As in [6], we consider the case where I is stabilized by a nontrivial subgroup H of G ; that is, for every $\sigma \in H$ and every $\alpha \in I$, we have $\sigma(\alpha) \in I$. Write $|H| = d$, so that $d \mid n$. Let K^H denote the fixed field of H , namely

$$K^H = \{\alpha \in K : \sigma(\alpha) = \alpha \text{ for all } \sigma \in H\}.$$

The field K is a $\mathbb{Q}$ -vector space, and every element of G acts on K as an invertible $\mathbb{Q}$ -linear transformation. Since H is a subgroup of G , we thus obtain the following

rational representation of H :

$$\psi: H \rightarrow \text{GL}(K), \quad \sigma \mapsto \sigma.$$

Using Algorithm 1, we can compute in $\text{poly}(n)$ time all irreducible rational representations of H and their rational characters; denote them by $\psi_1, \dots, \psi_t$ and $\chi_1, \dots, \chi_t$. We also obtain the projection operators associated with the canonical decomposition of K as a $\mathbb{Q}$ -vector space:

$$p_i = \frac{1}{|H|} \sum_{\sigma \in H} \overline{\chi_i(\sigma)} \sigma, \quad i = 1, \dots, t.$$

Set $V_i = \text{Im } p_i$ for $i = 1, \dots, t$. Each V_i is a $\mathbb{Q}$ -linear subspace of K , but not necessarily a subfield of K . Without loss of generality, assume that ψ_1 is the trivial one-dimensional rational representation of H , so that $\chi_1(\sigma) = 1$ for all $\sigma \in H$, and hence

$$p_1 = \frac{1}{|H|} \sum_{\sigma \in H} \sigma.$$

Proposition 1. $V_1 = K^H$, i.e., the fixed field of H .

Proof. For any $\alpha \in V_1$, there exists $\beta \in K$ such that $\alpha = p_1(\beta)$, i.e.,

$$\alpha = \frac{1}{|H|} \sum_{\sigma \in H} \sigma(\beta).$$

For any $\tau \in H$, we have

$$\tau(\alpha) = \tau \left(\frac{1}{|H|} \sum_{\sigma \in H} \sigma(\beta) \right) = \frac{1}{|H|} \sum_{\sigma \in H} \tau \circ \sigma(\beta) = \frac{1}{|H|} \sum_{\sigma \in H} \sigma(\beta) = \alpha.$$

Thus $\alpha \in K^H$. Conversely, for any $\alpha \in K^H$, we have

$$\alpha = \frac{1}{|H|} \sum_{\sigma \in H} \alpha = \frac{1}{|H|} \sum_{\sigma \in H} \sigma(\alpha) = p_1(\alpha).$$

Therefore $\alpha \in \text{Im } p_1 = V_1$. Hence $V_1 = K^H$. $\square$

Observe that $\dim K^H = [G : H] < \dim K$ because H is a nontrivial subgroup of G , and since $\sum_{i=1}^t \dim V_i = \dim K$, it follows that $\max_{1 \leq i \leq t} \dim V_i < \dim K$. Define $m := \max_{1 \leq i \leq t} \dim V_i$.

Note that every irreducible complex character of H is a root of unity and hence an algebraic integer. For any irreducible rational character χ_i of H , it is a sum of several irreducible complex characters, so it is also an algebraic integer. Since $\chi_i(\sigma)$ is both a rational number and an algebraic integer for every $\sigma \in H$, it follows that $\chi_i(\sigma) \in \mathbb{Z}$.

For each $i \in \{1, \dots, t\}$, define

$$q_i = \sum_{\sigma \in H} \overline{\chi_i(\sigma)} \sigma = |H| \cdot p_i.$$

Let $I_i = q_i(I) = \{q_i(\alpha) : \alpha \in I\}$, $i = 1, \dots, t$. We then have the following important result.

Lemma 3. *For every $i \in \{1, \dots, t\}$, we have $I_i \subseteq I \cap V_i$, and I_i is a free $\mathbb{Z}$ -module of rank $\dim V_i$. Consequently, $\pi(I_i)$ is a lattice of rank $\dim V_i$, where π is the canonical embedding of K into $\mathbb{C}^n$.*

Proof. Clearly I_i is a $\mathbb{Z}$ -module, since I is a $\mathbb{Z}$ -module and q_i is a linear map. Since $\chi_i(\sigma) \in \mathbb{Z}$ for every $\sigma \in H$ and I is stabilized by H , it follows that $q_i(\alpha) \in I$ for every $\alpha \in I$; hence $I_i \subseteq I$. Moreover, $q_i(K) = |H|p_i(K) = V_i$ since V_i is a $\mathbb{Q}$ -vector space, which implies $I_i \subseteq I \cap V_i$. As $I_i \subseteq \mathcal{O}_K$ and $\mathcal{O}_K$ is a free $\mathbb{Z}$ -module of rank n , the submodule I_i is also free over $\mathbb{Z}$; denote its rank by r_i . Note that $\sum_{i=1}^t I_i \subseteq I$. On the other hand, for any $\alpha \in I$, by Theorem 5 and (6) we have

$$\sum_{i=1}^t q_i(\alpha) = |H| \sum_{i=1}^t p_i(\alpha) = |H|\alpha.$$

Therefore,

$$|H| \cdot I \subseteq \sum_{i=1}^t I_i \subseteq I.$$

It follows that the rank of $\sum_{i=1}^t I_i$ equals the rank of I , which is n . Since the rank of $\sum_{i=1}^t I_i$ is at most $\sum_{i=1}^t r_i$, we obtain $n \leq \sum_{i=1}^t r_i$. On the other hand, $I_i \subseteq V_i$ implies $r_i \leq \dim V_i$, and $\sum_{i=1}^t \dim V_i = n$. Combining these inequalities yields $r_i = \dim V_i$ for each i . $\square$

We now prove that the rank- n BDD problem on I can be reduced to at most t smaller BDD instances of rank at most m . Here the rank refers to the rank of the lattice.

Theorem 7. *Let $K/\mathbb{Q}$ be an abelian extension of degree n with Galois group G , and let $I \subseteq \mathcal{O}_K$ be an integral ideal. Suppose that I is stabilized by a nontrivial subgroup H of G . Then for any $\gamma \geq 2|H|^2$, the rank- $n \frac{1}{\gamma}$ -Id-BDD $_K$ problem on I can be reduced in polynomial time to at most $|H|$ sub-problems of $\frac{|H|^2}{\gamma}$ -BDD, each of rank at most m . Here $m = \max_{1 \leq i \leq t} \dim V_i$, where $V_1, \dots, V_t$ are the components of the canonical decomposition of K under the action of H .*

Proof. The following algorithm fulfills the requirements of the theorem.

Algorithm 2 Reduction algorithm for Id-BDD**Require:** K, G, I, H , target vector $\mathbf{t}$ **Ensure:** several $\frac{|H|^2}{\gamma}$ -BDD instances $(\mathbf{t}_1, I_1), \dots, (\mathbf{t}_t, I_t)$

- 1: Using the elementary divisor form of H , compute all irreducible complex representations of H and their characters.
- 2: Compute $\text{Gal}(\mathbb{Q}[\zeta_d]/\mathbb{Q})$, where $d = |H|$.
- 3: Use Theorem 4 to obtain all irreducible rational characters $\chi_1, \dots, \chi_t$ of H .
- 4: For $i = 1, \dots, t$, construct the maps $q_i = \sum_{\sigma \in H} \chi_i(\sigma) \sigma$.
- 5: For $i = 1, \dots, t$, compute $I_i = q_i(I) = \{q_i(\alpha) : \alpha \in I\}$.
- 6: For $i = 1, \dots, t$, compute $\mathbf{t}_i = q_i(\mathbf{t})$.
- 7: **return** $(\mathbf{t}_1, I_1), \dots, (\mathbf{t}_t, I_t)$.

The above algorithm runs in time $\text{poly}(n)$. By Theorem 4, the number t of irreducible rational representations of H is at most $|H|$.

From the preceding analysis and Lemma 3, we have $\dim \pi(I_i) = \dim V_i \leq m$, where π is the canonical embedding of K into $\mathbb{C}^n$. Next we prove that for each $i \in \{1, \dots, t\}$, the distance from $\mathbf{t}_i$ to $\pi(I_i)$ is at most $\frac{|H|^2}{\gamma} \lambda_1(\pi(I_i))$.

By assumption, the distance from $\mathbf{t}$ to $\pi(I)$ is at most $\frac{1}{\gamma} \lambda_1(\pi(I))$; that is, $\mathbf{t} = \mathbf{v} + \mathbf{e}$, where $\mathbf{v} \in \pi(I)$, $\mathbf{e} \in \text{Span}_{\mathbb{R}}(\pi(I))$, and $\|\mathbf{e}\| \leq \frac{1}{\gamma} \lambda_1(\pi(I))$. For any $\sigma \in H$, we have

$$\|\sigma(\mathbf{t})\| = \|\mathbf{t}\|, \quad \|\sigma(\mathbf{e})\| = \|\mathbf{e}\|.$$

Furthermore, since every irreducible complex character of H is a root of unity, its absolute value is 1. Therefore, for each $i \in \{1, \dots, t\}$ and every $\sigma \in H$, we have $|\chi_i(\sigma)| \leq |O_i| \leq |H|$, where O_i is the corresponding orbit under the action of $\text{Gal}(\mathbb{Q}[\zeta_d]/\mathbb{Q})$ with $d = |H|$. Hence, by the triangle inequality,

$$\|q_i(\mathbf{e})\| \leq \sum_{\sigma \in H} |\overline{\chi_i(\sigma)}| \cdot \|\sigma(\mathbf{e})\| \leq |H| \cdot |O_i| \cdot \|\mathbf{e}\| \leq |H|^2 \|\mathbf{e}\| \leq \frac{|H|^2}{\gamma} \lambda_1(\pi(I)) \leq \frac{|H|^2}{\gamma} \lambda_1(\pi(I_i)),$$

where the last inequality holds because $I_i \subseteq I$ implies $\lambda_1(\pi(I_i)) \geq \lambda_1(\pi(I))$.

Since q_i is linear, for every $i \in \{1, \dots, t\}$ we have

$$\mathbf{t}_i = q_i(\mathbf{t}) = q_i(\mathbf{v}) + q_i(\mathbf{e}),$$

with $q_i(\mathbf{v}) \in I_i$. The above estimate shows that the distance from $\mathbf{t}_i$ to I_i is at most $\frac{|H|^2}{\gamma} \lambda_1(\pi(I_i))$. Thus, by solving the at most t sub-BDD problems $(\mathbf{t}_1, I_1), \dots, (\mathbf{t}_t, I_t)$ to obtain $q_1(\mathbf{v}), \dots, q_t(\mathbf{v})$, and using the identity $\sum_{i=1}^t q_i = |H| \cdot \text{id}_K$, we can recover the correct solution $\mathbf{v}$ to the original BDD instance $(\mathbf{t}, I)$. $\square$

We now analyze when Theorem 7 speeds up the solution of $\frac{1}{\gamma}$ -Id-BDD $_K$. According to Lemma 1, directly solving $\frac{1}{\gamma}$ -Id-BDD $_K$ has complexity $\text{poly}(n) \cdot 2^{O(n \log n / \log \gamma)}$, while a rank- m $\frac{|H|^2}{\gamma}$ -BDD problem has complexity $\text{poly}(n) \cdot$

$2^{O(m \log m / \log(\gamma/|H|^2))}$. Therefore, the reduction of Theorem 7 accelerates the solution of $\frac{1}{\gamma}$ -Id-BDD $_K$ whenever

$$\frac{n \log n}{\log \gamma} > \frac{m \log m}{\log(\gamma/|H|^2)}. \quad (7)$$

This condition is readily satisfied in typical cryptographic settings. Indeed, in lattice-based schemes one usually has $\gamma = \text{poly}(n)$, so that $\log \gamma = \Theta(\log n)$. Consequently, the exponent governing the direct attack is essentially linear in n , whereas the exponent for the reduced rank- m instances is essentially linear in m . Since $m < n$, this already yields an exponential improvement. Moreover, in practically relevant cases the difference between m and n can be substantial: for example, when I is stabilized by a subgroup H whose action decomposes K into several components of comparable size, m may be as small as $O(n/|H|)$. In addition, because $|H|$ is typically a small constant while γ is polynomial in n , the denominator $\log(\gamma/|H|^2)$ differs from $\log \gamma$ only by an additive constant. Thus the reduction decreases the main exponential term from roughly $2^{O(n \log n / \log \gamma)}$ to $2^{O(m \log m / \log \gamma)}$, which represents an exponential speedup whenever $m \ll n$. The inequality (7) therefore holds for essentially all parameters of interest, and the improvement becomes more pronounced as the symmetry of I increases.

3.3 Solving γ -Id-gap-uSVP $_K$ via canonical decomposition of rational representations

The canonical decomposition of rational representations can also be used to solve the γ -Id-gap-uSVP $_K$ problem. We keep the notation of the previous subsection.

Theorem 8. *Let $K/\mathbb{Q}$ be an abelian extension of degree n with Galois group G , and let $I \subseteq \mathcal{O}_K$ be an integral ideal. Suppose that I is stabilized by a nontrivial subgroup H of G . Then for any $\gamma > |H|^2$, the rank- n γ -Id-gap-uSVP $_K$ problem on I can be reduced in polynomial time to at most t instances of the γ -SVP problem, each of rank at most m , where $m = \max_{1 \leq i \leq t} \dim V_i$, and $V_1, \dots, V_t$ are the components of the canonical decomposition of K under the action of H .*

Proof. The following algorithm realizes the reduction.

Algorithm 3 Reduction algorithm for Id-uSVP

Require: K, G, I, H

Ensure: A list of γ -SVP instances $I_1, \dots, I_t$.

- 1: Using the elementary divisor form of H , compute all irreducible complex representations of H and their characters.
 - 2: Compute $\text{Gal}(\mathbb{Q}[\zeta_d]/\mathbb{Q})$, where $d = |H|$.
 - 3: Use Theorem 4 to obtain all irreducible rational characters $\chi_1, \dots, \chi_t$ of H .
 - 4: For $i = 1, \dots, t$, construct the maps $q_i = \sum_{\sigma \in H} \overline{\chi_i(\sigma)} \sigma$.
 - 5: For $i = 1, \dots, t$, compute $I_i = q_i(I) = \{q_i(\alpha) : \alpha \in I\}$.
 - 6: **return** $I_1, \dots, I_t$.
-

The above algorithm runs in time $\text{poly}(n)$, which is polynomial in the input size.

Let $\alpha \in I$ be a shortest nonzero vector. By Theorem 5 we have

$$\sum_{i=1}^t q_i = |H| \cdot \text{id}_K.$$

Applying this identity to α gives

$$\sum_{i=1}^t q_i(\alpha) = |H| \cdot \alpha \neq 0.$$

Hence there exists at least one index $j \in \{1, \dots, t\}$ such that $q_j(\alpha) \neq 0$. For this j , the triangle inequality yields

$$\|q_j(\alpha)\| = \left\| \sum_{\sigma \in H} \overline{\chi_j(\sigma)} \sigma(\alpha) \right\| \leq \sum_{\sigma \in H} |\chi_j(\sigma)| \cdot \|\sigma(\alpha)\| \leq |H|^2 \cdot \|\alpha\|.$$

Since $\gamma > |H|^2$, it follows that $\|q_j(\alpha)\| < \gamma \cdot \|\alpha\|$. By the gap condition in Definition 1, every nonzero vector of I has length either $\|\alpha\|$ or at least $\gamma\|\alpha\|$. As $q_j(\alpha) \in I_j \subseteq I$ and $\|q_j(\alpha)\| < \gamma\|\alpha\|$, we must have $\|q_j(\alpha)\| = \|\alpha\|$; thus $q_j(\alpha)$ is also a shortest nonzero vector of I .

Now consider the γ -SVP instances returned by Algorithm 3. For each i , solve γ -SVP on $\pi(I_i)$ and let $\mathbf{v}_i$ be the returned vector. Then $\|\mathbf{v}_i\| \leq \gamma \lambda_1(\pi(I_i))$. For the index j constructed above, we have $\lambda_1(\pi(I_j)) = \|\alpha\|$, and hence $\|\mathbf{v}_j\| \leq \gamma\|\alpha\|$. By the same gap property, $\|\mathbf{v}_j\|$ must equal $\|\alpha\|$. Therefore, among the returned vectors $\mathbf{v}_1, \dots, \mathbf{v}_t$, at least one has length $\|\alpha\|$; since all other returned vectors have length at least $\|\alpha\|$, selecting the vector of minimal length yields a shortest nonzero vector of I . This solves the original γ -Id-gap-uSVP $_K$ instance. $\square$

According to Lemma 1, directly solving γ -Id-gap-uSVP $_K$ has complexity $\text{poly}(n) \cdot 2^{O(n \log n / \log \gamma)}$. After the reduction of Theorem 8, we solve at most t instances of γ -SVP, each of rank at most m . Since $t \leq |H|$ is polynomial in n , the total complexity is

$$\text{poly}(n) \cdot 2^{O(m \log m / \log \gamma)}.$$

Since $m < n$, this gives a significant exponential speedup over the original rank- n problem.

Our approach is related to the method of [6], which reduces the original approximate Ideal-SVP problem to an approximate Ideal-SVP problem on $K^H \cap I$. That reduction requires certain ramification conditions on I , whereas we instead require $K/\mathbb{Q}$ to be abelian. Our method further extends their approach to solve the γ -uSVP problem on I under the condition $\gamma > |H|^2$, at the cost of solving several lower-dimensional SVP instances. In the worst case, this still yields an

exponential speedup over solving the original n -dimensional uSVP instance directly. To the best of our knowledge, this observation has not been reported in previous work.

We note that when the goal is only to solve the approximate Ideal-SVP problem, our reduction can be simplified to match the strategy of [6]: it suffices to solve approximate SVP on the component I_1 and output the resulting vector. Since $\text{rank}(I_1) = [K^H : \mathbb{Q}] = n/|H|$, which equals the rank of $K^H \cap I$, and since I_1 is sparser than $K^H \cap I$ (in fact, $I_1 \subseteq K^H \cap I$), this special case achieves at least the same efficiency as the method of [6].

4 Further Reductions for Ideal Lattices in Cyclotomic Rings

In cryptographic constructions, for efficiency one often works with cyclotomic rings, i.e., the ring of integers $\mathcal{O}_K$ of a cyclotomic field $K = \mathbb{Q}(\zeta)$, where ζ is a primitive root of unity. These rings enjoy the following useful property: for every element $\alpha \in \mathcal{O}_K$ and every integer i , the norm is invariant under multiplication by ζ^i , namely,

$$\|\alpha\| = \|\zeta^i \alpha\|, \quad \forall i \in \mathbb{Z}. \quad (8)$$

Here the norm denotes the Hermitian norm of the image under the canonical embedding. In this section we prove that this property can be used to further reduce the complexity of solving the $\frac{1}{\gamma}$ -Id-BDD $_K$ and γ -Id-gap-uSVP $_K$ problems studied in the preceding sections. A similar idea has been used in [9] to solve the Partial Vandermonde Knapsack problem, which can be reduced to $\frac{1}{\gamma}$ -Id-BDD $_K$; however, their method is restricted to the special cases of power-of-two cyclotomic fields or prime-order cyclotomic fields. The method of this section can therefore be viewed as a further generalization of [9].

4.1 Solving $\frac{1}{\gamma}$ -Id-BDD $_K$ in Cyclotomic Rings

Theorem 9. *Let $K = \mathbb{Q}(\zeta)$ be a cyclotomic field of degree n , with ring of integers $\mathcal{O}_K$. Let $I \subseteq \mathcal{O}_K$ be a full-rank ideal, and suppose I is stabilized by a non-trivial subgroup H of the Galois group $G = \text{Gal}(K/\mathbb{Q})$. Then, for any $\gamma \geq 2|H|^2$, the rank- $n \frac{1}{\gamma}$ -Id-BDD $_K$ problem on I can be reduced in polynomial time to n instances of the $\frac{|H|^2}{\gamma}$ -BDD problem, each of rank m , where*

$$m = \min_{\substack{1 \leq i \leq t \\ \dim V_i > 0}} \dim V_i,$$

and $V_1, \dots, V_t$ are the nonzero components of the canonical decomposition of K under the action of H .

Proof. The following algorithm realizes the reduction.

Algorithm 4 Reduction algorithm for Id-BDD in cyclotomic rings**Require:** K, G, I, H , target vector $\mathbf{t}$ **Ensure:** Several $\frac{|H|^2}{\gamma}$ -BDD instances $(\mathbf{t}_0, I'), \dots, (\mathbf{t}_{n-1}, I')$

- 1: Using the elementary divisor form of H , compute all irreducible complex representations of H and their characters.
- 2: Compute $\text{Gal}(\mathbb{Q}[\zeta_d]/\mathbb{Q})$, where $d = |H|$.
- 3: Use Theorem 4 to obtain all irreducible rational characters $\chi_1, \dots, \chi_t$ of H .
- 4: For $i = 1, \dots, t$, construct the maps $q_i = \sum_{\sigma \in H} \overline{\chi_i(\sigma)} \sigma$.
- 5: For $i = 1, \dots, t$, compute $I_i = q_i(I) = \{q_i(\alpha) : \alpha \in I\}$.
- 6: $j \leftarrow \arg \min_{\substack{1 \leq i \leq t \\ \dim V_i > 0}} \dim V_i$.
- 7: $I' \leftarrow I_j$.
- 8: For $i = 0, \dots, n-1$, compute $\mathbf{t}_i = q_j(\zeta^i \mathbf{t})$.
- 9: **return** $(\mathbf{t}_0, I'), \dots, (\mathbf{t}_{n-1}, I')$.

Clearly the running time of the above algorithm is $\text{poly}(n)$.

For every $i \in \{0, 1, \dots, n-1\}$, we have $\zeta^i \mathbf{t} \neq \mathbf{0}$ since $\mathbf{t} \neq \mathbf{0}$. Write $\mathbf{t} = \mathbf{v} + \mathbf{e}$, where $\mathbf{v} \in \pi(I)$, $0 < \|\mathbf{e}\| \leq \frac{1}{\gamma} \lambda_1(\pi(I))$, and π denotes the canonical embedding. For every $i \in \{0, 1, \dots, n-1\}$, we have $\zeta^i \mathbf{t} = \zeta^i \mathbf{v} + \zeta^i \mathbf{e}$, where $\zeta^i \mathbf{v} \in \pi(I)$ because I is closed under multiplication, and by (8) we have $\|\zeta^i \mathbf{e}\| = \|\mathbf{e}\|$. By the triangle inequality,

$$\begin{aligned} \|q_j(\zeta^i \mathbf{e})\| &= \left\| \sum_{\sigma \in H} \overline{\chi_j(\sigma)} \sigma(\zeta^i \mathbf{e}) \right\| \\ &\leq |H|^2 \cdot \|\mathbf{e}\| \leq \frac{|H|^2}{\gamma} \lambda_1(\pi(I)) \leq \frac{1}{2} \lambda_1(\pi(I)) \leq \frac{1}{2} \lambda_1(\pi(I_j)), \end{aligned}$$

where the last inequality follows from $I_j \subseteq I$.

Since q_j is linear, we have $\mathbf{t}_i = q_j(\zeta^i \mathbf{t}) = q_j(\zeta^i \mathbf{v}) + q_j(\zeta^i \mathbf{e})$, where $q_j(\zeta^i \mathbf{v}) \in \pi(I_j)$. The above inequality shows that the distance from $\mathbf{t}_i$ to $\pi(I_j)$ is at most $\frac{|H|^2}{\gamma} \lambda_1(\pi(I_j))$.

Finally, we show that by solving these n instances of $\frac{|H|^2}{\gamma}$ -BDD and using the resulting solutions $\mathbf{v}_0, \dots, \mathbf{v}_{n-1}$, one can recover the original solution $\mathbf{v}$ in polynomial time. By the proof of Theorem 5, q_j is an invertible linear transformation on V_j . For each $i \in \mathbb{Z}$, let $\tau_i : \pi(K) \rightarrow \pi(K)$, $\mathbf{v} \mapsto \zeta^i \mathbf{v}$, be the invertible linear map induced by multiplication by $\zeta^i \in \mathcal{O}_K$. Then

$$\mathbf{v}_i = q_j \circ \tau_i(\mathbf{v}), \quad i = 0, 1, \dots, n-1. \quad (9)$$

Here $\mathbf{v}$ is unknown, and we obtain a linear system with n unknowns. This system clearly has a solution $\mathbf{v}$; we now prove that the solution is unique, so $\mathbf{v}$ can be recovered by Gaussian elimination in polynomial time.

Suppose there exist distinct $\mathbf{v}, \mathbf{v}' \in \pi(K)$ such that (9) holds for both. Then

$$\mathbf{0} = q_j \circ \tau_i(\mathbf{v} - \mathbf{v}'), \quad i = 0, 1, \dots, n-1,$$

which is equivalent to

$$\zeta^i(\mathbf{v} - \mathbf{v}') \in \ker q_j, \quad i = 0, 1, \dots, n-1.$$

Write $\mathbf{v} = \pi(\alpha)$, $\mathbf{v}' = \pi(\alpha')$, and set $\beta = \alpha - \alpha' \neq 0$. The equations give $q_j(\zeta^i \beta) = 0$ for all i . Since $\{\zeta^k \beta\}_{k=0}^{n-1}$ is a $\mathbb{Z}$ -basis of (β) , we get $(\beta) \subseteq \ker q_j$, hence $\dim_{\mathbb{Q}} \ker q_j \geq n$. But $\dim_{\mathbb{Q}} \ker q_j = n - \dim_{\mathbb{Q}} \operatorname{Im} q_j < n$ because $\dim_{\mathbb{Q}} \operatorname{Im} q_j > 0$. Contradiction. Thus $\mathbf{v} = \mathbf{v}'$. $\square$

Note that with this method, the complexity of solving $\frac{1}{\gamma}$ -Id-BDD $_K$ can be further reduced to $\operatorname{poly}(n) \cdot 2^{O(m \log m / \log(\gamma/|H|^2))}$, where m is now the smallest positive dimension among $V_1, \dots, V_t$, and we have the known upper bound $m \leq \frac{n}{|H|}$. In contrast, the complexity in Section 3.2 has the same form but with m equal to the largest dimension among $V_1, \dots, V_t$. Therefore, when the smallest positive dimension is strictly smaller than the largest dimension, the method of Theorem 9 yields an additional exponential improvement over the method of Section 3.2.

4.2 Solving γ -Id-gap-uSVP $_K$ in Cyclotomic Rings

We now show that the same rotation technique used for Id-BDD can also be applied to solve the γ -Id-gap-uSVP $_K$ problem on a component of minimal positive rank among $I_1, \dots, I_t$.

Theorem 10. *Let $K = \mathbb{Q}(\zeta)$ be a cyclotomic field of degree n , with ring of integers $\mathcal{O}_K$. Let $I \subseteq \mathcal{O}_K$ be a full-rank ideal, and suppose I is stabilized by a nontrivial subgroup H of the Galois group $G = \operatorname{Gal}(K/\mathbb{Q})$. Then for any $\gamma > |H|^2$, the rank- n γ -Id-gap-uSVP $_K$ problem on I can be reduced in polynomial time to a single γ -SVP problem of rank m , where*

$$m = \min_{\substack{1 \leq i \leq t \\ \dim V_i > 0}} \dim V_i,$$

and $V_1, \dots, V_t$ are the components of the canonical decomposition of K under the action of H .

Proof. The following algorithm realizes the reduction.

Algorithm 5 Reduction algorithm for Id-uSVP in cyclotomic rings**Require:** K, G, I, H **Ensure:** A γ -SVP instance I' .

- 1: Using the elementary divisor form of H , compute all irreducible complex representations of H and their characters.
- 2: Compute $\text{Gal}(\mathbb{Q}[\zeta_d]/\mathbb{Q})$, where $d = |H|$.
- 3: Use Theorem 4 to obtain all irreducible rational characters $\chi_1, \dots, \chi_t$ of H .
- 4: For $i = 1, \dots, t$, construct the maps $q_i = \sum_{\sigma \in H} \overline{\chi_i(\sigma)} \sigma$.
- 5: For $i = 1, \dots, t$, compute $I_i = q_i(I) = \{q_i(\alpha) : \alpha \in I\}$.
- 6: $j \leftarrow \arg \min_{\substack{1 \leq i \leq t \\ \dim V_i > 0}} \dim V_i$.
- 7: **return** I_j .

Clearly the above algorithm runs in time $\text{poly}(n)$. Let $\alpha \in I$ be a shortest nonzero vector. We first prove that for every nonzero component V_i of the canonical decomposition, the corresponding lattice $I_i = q_i(I)$ contains a shortest vector of I .

Fix such an index i . We claim that there exists $k \in \{0, \dots, n-1\}$ such that $q_i(\zeta^k \alpha) \neq 0$. Suppose, for contradiction, that $q_i(\zeta^k \alpha) = 0$ for all k . Then $\zeta^k \alpha \in \ker q_i$ for every k . Since the set $\{\zeta^k \alpha\}_{k=0}^{n-1}$ forms a $\mathbb{Z}$ -basis for the principal ideal (α) , and q_i is $\mathbb{Z}$ -linear, it follows that $(\alpha) \subseteq \ker q_i$. But (α) is a full-rank ideal of rank n , so $\dim_{\mathbb{Q}} \ker q_i \geq n$, contradicting $\dim \text{Im } q_i = \dim V_i > 0$. Hence such a k exists.

Set $\beta = \zeta^k \alpha$. By (8), β is again a shortest vector of I , and $q_i(\beta) \neq 0$. Moreover,

$$\|q_i(\beta)\| \leq |H|^2 \|\beta\| = |H|^2 \lambda_1(I) < \gamma \lambda_1(I),$$

since $\gamma > |H|^2$. Since $q_i(\beta) \in I_i \subseteq I$ and I satisfies the γ -gap condition of Definition 1, the vector $q_i(\beta)$ must itself be a shortest vector of I . Thus every nonzero component I_i contains a shortest vector of I .

Now let I_j be the nonzero component of minimal rank among $I_1, \dots, I_t$. Since I_j contains a shortest vector of I and inherits the gap condition, solving γ -SVP on I_j directly yields a shortest vector of I . $\square$

The complexity of this reduction is polynomial in n , and the resulting SVP instance has rank m . Using the BKZ complexity estimate from Lemma 1, the direct solution of the original rank- n problem requires time $\text{poly}(n) \cdot 2^{O(n \log n / \log \gamma)}$, whereas our reduced problem of rank m requires only $\text{poly}(n) \cdot 2^{O(m \log m / \log \gamma)}$. Since $m \leq n/|H|$ (as the trivial component has dimension $n/|H|$), this yields a substantial exponential speedup, especially when $|H|$ is large.

5 Conclusion and Open Problems

We introduced a representation-theoretic framework for solving Id-uSVP and Id-BDD in ideal lattices stabilized by a nontrivial abelian subgroup of the Galois group. By projecting onto rational components, we reduce both problems to

low-dimensional instances. In cyclotomic fields, the rotation property gives further improvements. To the best of our knowledge, our method is the first to use group representations directly in lattice cryptanalysis. It extends previous subfield techniques to BDD and to all cyclotomic fields, and avoids the requirement of multiplicative closure.

An interesting open problem is to extend our framework to non-abelian Galois extensions. In that setting, the representation theory is substantially more involved, requiring tools such as the Schur index, and the simple projection decomposition used here does not carry over directly. Due to space limitations, we did not investigate this direction, but we believe it is a natural and important avenue for future work.

References

1. Ajtai, M.: Generating hard instances of lattice problems (extended abstract). In: Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing. pp. 99–108. STOC '96, Association for Computing Machinery, New York, NY, USA (1996). <https://doi.org/10.1145/237814.237838>, <https://doi.org/10.1145/237814.237838>
2. Albrecht, M., Bai, S.,ucas, L.: A subfield lattice attack on overstretched ntru assumptions. In: Robshaw, M., Katz, J. (eds.) Advances in Cryptology – CRYPTO 2016. pp. 153–178. Springer Berlin Heidelberg, Berlin, Heidelberg (2016)
3. Bernstein, D.J., Chuengsatiansup, C., Lange, T., van Vredendaal, C.: Ntru prime: Reducing attack surface at low cost. In: Adams, C., Camenisch, J. (eds.) Selected Areas in Cryptography – SAC 2017. pp. 235–260. Springer International Publishing, Cham (2018)
4. Bernstein, Daniel J.: A subfield-logarithm attack against ideal lattices: computing algebraic number theory tackles lattice-based cryptography. (The crypto blog, 2014), <https://blog.cr.yp.to/20140213-ideal.html>
5. Bos, J., Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schanck, J.M., Schwabe, P., Seiler, G., Stehle, D.: Crystals - kyber: A cca-secure module-lattice-based kem. In: 2018 IEEE European Symposium on Security and Privacy (EuroS&P). pp. 353–367 (2018). <https://doi.org/10.1109/EuroSP.2018.00032>
6. Boudgoust, K., Gachon, E., Pellet-Mary, A.: Some easy instances of ideal-svp and implications on the partial vandermonde knapsack problem. In: Dodis, Y., Shrimpton, T. (eds.) Advances in Cryptology – CRYPTO 2022. pp. 480–509. Springer Nature Switzerland, Cham (2022)
7. Brakerski, Z., Langlois, A., Peikert, C., Regev, O., Stehlé, D.: Classical hardness of learning with errors. In: Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing. pp. 575–584. STOC '13, Association for Computing Machinery, New York, NY, USA (2013). <https://doi.org/10.1145/2488608.2488680>, <https://doi.org/10.1145/2488608.2488680>
8. Cheon, J.H., Jeong, J., Lee, C.: An algorithm for ntru problems and cryptanalysis of the ggh multilinear map without a low-level encoding of zero. LMS Journal of Computation and Mathematics **19**(A), 255–266 (2016). <https://doi.org/10.1112/S1461157016000371>
9. Das, D., Joux, A.: Key recovery attack on the partial vandermonde knapsack problem. In: Joye, M., Leander, G. (eds.) Advances in Cryptology – EUROCRYPT 2024. pp. 205–225. Springer Nature Switzerland, Cham (2024)

10. Ducas, L., Kiltz, E., Lepoint, T., Lyubashevsky, V., Schwabe, P., Seiler, G., Stehlé, D.: Crystals-dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems* **2018**(1), 238–268 (Feb 2018). <https://doi.org/10.13154/tches.v2018.i1.238-268>, <https://tches.iacr.org/index.php/TCHES/article/view/839>
11. Fouque, P.A., Hoffstein, J., Kirchner, P., Lyubashevsky, V., Pornin, T., Prest, T., Ricosset, T., Seiler, G., Whyte, W., Zhang, Z.: Falcon: Fast-fourier lattice-based compact signatures over ntru (2019), <https://api.semanticscholar.org/CorpusID:231637439>
12. Hanrot, G., Pujol, X., Stehlé, D.: Analyzing blockwise lattice algorithms using dynamical systems. In: Rogaway, P. (ed.) *Advances in Cryptology – CRYPTO 2011*. pp. 447–464. Springer Berlin Heidelberg, Berlin, Heidelberg (2011)
13. Hoffstein, J., Pipher, J., Silverman, J.H.: Ntru: A ring-based public key cryptosystem. In: Buhler, J.P. (ed.) *Algorithmic Number Theory*. pp. 267–288. Springer Berlin Heidelberg, Berlin, Heidelberg (1998)
14. Kirchner, P., Fouque, P.A.: Revisiting lattice attacks on overstretched ntru parameters. In: Coron, J.S., Nielsen, J.B. (eds.) *Advances in Cryptology – EUROCRYPT 2017*. pp. 3–26. Springer International Publishing, Cham (2017)
15. Langlois, A., Stehlé, D.: Worst-case to average-case reductions for module lattices. *Designs, Codes and Cryptography* **75**, 565 – 599 (2014), <https://api.semanticscholar.org/CorpusID:1853632>
16. Lyubashevsky, V., Micciancio, D.: Generalized compact knapsacks are collision resistant. In: Bugliesi, M., Preneel, B., Sassone, V., Wegener, I. (eds.) *Automata, Languages and Programming*. pp. 144–155. Springer Berlin Heidelberg, Berlin, Heidelberg (2006)
17. Lyubashevsky, V., Peikert, C., Regev, O.: On ideal lattices and learning with errors over rings. In: Gilbert, H. (ed.) *Advances in Cryptology – EUROCRYPT 2010*. pp. 1–23. Springer Berlin Heidelberg, Berlin, Heidelberg (2010)
18. Marcus, D.A.: *Number Fields*. Universitext, 2 edn. (2018). <https://doi.org/10.1007/978-3-319-90233-3>, originally published by Springer-Verlag New York, 1977
19. Pan, Y., Xu, J., Wadleigh, N., Cheng, Q.: On the ideal shortest vector problem over random rational primes. In: Canteaut, A., Standaert, F.X. (eds.) *Advances in Cryptology – EUROCRYPT 2021*. pp. 559–583. Springer International Publishing, Cham (2021)
20. Peikert, C.: Public-key cryptosystems from the worst-case shortest vector problem: extended abstract. In: *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*. pp. 333–342. STOC ’09, Association for Computing Machinery, New York, NY, USA (2009). <https://doi.org/10.1145/1536414.1536461>, <https://doi.org/10.1145/1536414.1536461>
21. Porter, C., Mendelsohn, A., Ling, C.: Subfield algorithms for ideal- and module-svp based on the decomposition group. *CoRR* **abs/2105.03219** (2021), <https://arxiv.org/abs/2105.03219>
22. Regev, O.: On lattices, learning with errors, random linear codes, and cryptography. *J. ACM* **56**(6) (Sep 2009). <https://doi.org/10.1145/1568318.1568324>, <https://doi.org/10.1145/1568318.1568324>
23. Serre, J.P.: *Linear representations of finite groups*. Springer-Verlag, New York (1977), translated from the second French edition by Leonard L. Scott, Graduate Texts in Mathematics, Vol. 42