Some Easy Instances of Ideal-SVP and Implications on the Partial Vandermonde Knapsack Problem

Katharina Boudgoust[1], Erell Gachon[2], and Alice Pellet-Mary[2]

katharina.boudgoust@cs.dk.au, erell.gachon@u-bordeaux.fr, alice.pellet-mary@math.u-bordeaux.fr

[1] Aarhus University, Denmark
[2] Univ. Bordeaux, CNRS, INRIA, Bordeaux INP, IMB, UMR 5251, F-33400 Talence, France

Abstract. In this article, we generalize the works of Pan et al. (Eurocrypt'21) and Porter et al. (ArXiv'21) and provide a *simple* condition under which an ideal lattice defines an easy instance of the shortest vector problem. Namely, we show that the more automorphisms stabilize the ideal, the easier it is to find a short vector in it. This observation was already made for prime ideals in Galois fields, and we generalize it to *any* ideal (whose prime factors are not ramified) of *any* number field. We then provide a cryptographic application of this result by showing that particular instances of the partial Vandermonde knapsack problem, also known as partial Fourier recovery problem, can be solved classically in polynomial time. As a proof of concept, we implemented our attack and managed to solve those particular instances for concrete parameter settings proposed in the literature. For random instances, we can halve the lattice dimension with non-negligible probability.

1 Introduction

Euclidean lattices are mathematical objects that play an important role in many areas of mathematics and computer science. There are several computational problems related to lattices that are proven to be NP-hard, for instance, the problem of finding a shortest vector (SVP) or a set of shortest independent vectors (SIVP) in a given lattice. A standard relaxation consists in solving them only up to some approximation factor $\gamma \geq 1$, denoted γ-S(I)VP. It is commonly conjectured that the problems remain hard to solve for approximation factors that are polynomial in the lattice rank. Their presumed intractability provides a fundamental starting point for the construction of provably secure cryptographic schemes, shown in the seminal works of Ajtai [Ajt96] and Regev [Reg05].

Unfortunately, all cryptographic schemes relying on the hardness of those lattice problems inherently suffer from large keys and slow computation times, being quadratic in the security parameter. In order to improve efficiency, problems on *structured* lattices have been introduced, e.g., [Mic02, LM06, PR06, SSTX09, LPR10, LS15]. The most popular setting is to consider O_K-modules of rank r, where O_K is the ring of integers of some number field K of degree d. By applying the d different field embeddings from K to $\mathbb{C}$, any O_K-module of rank r is mapped to a lattice of rank $d \cdot r$. Those lattices inherit the module structure (i.e., closed with respect to scalar multiplication by ring elements) and are called *module lattices*. If the module rank equals 1, they are called *ideal lattices*.

Many structured lattice assumptions, such as Ring-LWE [SSTX09, LPR10], NTRU [HPS98] or Module-LWE [LS15] can be solved with an SVP solver in module lattices of small rank (≥ 2). This motivates the study of the hardness of SVP in module lattices. To start tackling this problem, many algorithms have focused on the special case of solving SVP in rank-1 modules, that is in ideals. This restricted problem is denoted by Id-SVP. While solving Id-SVP is not known to break any of the three lattice assumptions mentioned above, studying this (potentially easier) problem can be seen as a first step to better understand the hardness behind algebraically structured lattices. Another motivation for studying Id-SVP comes from the fact that the first hardness result for Ring-LWE was a reduction from worst-case Id-SVP [SSTX09, LPR10]. This reduction only provides a lower bound on the hardness of Ring-LWE, and we have today a stronger reduction, from worst-case SVP in modules of rank ≥ 2, for some more restricted regime of parameters of Ring-LWE [AD17]. Still, even if an efficient algorithm for Id-SVP would not have a direct impact on the security of Ring-LWE, it would make the reduction from Id-SVP vacuous, and hence let some interesting regime of Ring-LWE without lower bound security guarantees.

Even though most of the lattice-based cryptographic schemes are not known to reduce to SVP in ideal lattices (but in module lattices of rank ≥ 2), there are a few counter-examples. They can be found among the first constructions of FHE schemes by Gentry [Gen09] or, as we will see below, in the constructions based on the partial Vandermonde knapsack problem [HPS+14] (also known as the partial Fourier recovery problem).

Hardness of Id-SVP. The hardness of Id-SVP has attracted a lot of work in recent years. On the one hand, some works have proven worst-case to average-case reductions for problems in ideal lattices [Gen09, dBDPW20]. They proved that there exist distributions over the set of ideal lattices such that an ideal chosen from this distribution is "as hard as possible". More formally, if one can solve Id-SVP for such random lattices with non-negligible probability, then one can solve Id-SVP in any ideal lattice.

On the other hand, several works have shown weaknesses of Id-SVP for specific choices of ideals or parameters. Cramer et al. [CDPR16] showed that Id-SVP can be solved in quantum polynomial time for *principal* ideals (i.e., ideals generated by a single ring element) of cyclotomic fields, when the generator is sampled

from a Gaussian distribution. It is also known that the relaxed variant of Id-SVP with a large approximation factor $\approx 2^{\sqrt{d}}$ can be solved in quantum polynomial time in cyclotomic fields of degree d [CDW21]. In 2021, Pan et al. [PXWC21] showed that, for some prime ideals with a lot of symmetries (in Galois number fields), the Id-SVP problem can be solved classically in polynomial time, with a polynomial approximation factor. This was extended by Porter et al. [PML21, Theorem 3] to a larger class of ideals, whose characterization is harder to state and relies on factoring properties of the ideal, as well as its algebraic norm.

Finally, there is a line of work targeting Id-SVP for all ideals of all number fields, for various approximation factors [PHS19, BR20, BLNR21]. However, the algorithms require an exponential-time pre-processing, and are at the moment no better than lattice reduction algorithms that work on unstructured lattices (e.g. BKZ).

Partial Vandermonde Knapsack. In the late 90's, Hoffman et al. [HKJL+00] patented a method for user identification and digital signatures based on the difficulty of recovering a constrained polynomial from partial information. Afterwards, the partial information was specified as a partial list of the polynomial's Fourier transform resulting in a signature scheme called PASS Sign [HPS+14]. The constraint regarding the polynomial was to choose its coefficients uniformly at random over a bounded set. Lu et al. [LZA18] moved from the Fourier transform (evaluation at all roots of unity) over cyclic rings to the Vandermonde transform (evaluation only at the *primitive* roots of unity) over cyclotomic rings.

The hardness assumption that underlies PASS Sign, as given in [LZA18], is the following. Let q be a prime and let m be an integer such that there exists a primitive m-th root of unity in the quotient ring $\mathbb{Z}_q := \mathbb{Z}/q\mathbb{Z}$. In this case, there exist exactly $d = \phi(m)$ such primitive roots $\{\omega_j\}_{1 \le j \le d}$, where ϕ is Euler's totient function. Further, let $g(X)$ be a polynomial of degree less than d having small integer coefficients. Its Vandermonde transform $\mathbf{V}(g) \in \mathbb{Z}_q^d$ is defined as $g(\omega_j)_{1 \le j \le d} \bmod q$. For a subset $\Omega \subseteq \{1, \ldots, d\}$ of size t, its partial Vandermonde transform $\mathbf{V}_\Omega(g) \in \mathbb{Z}_q^t$ is given by $g(\omega_j)_{j \in \Omega}$. The partial Vandermonde knapsack problem (PV-Knap) asks, given $\mathbf{V}_\Omega(g)$, to recover $g(X)$.[3]

As observed by Boudgoust et al. [BSS22, Bou21], recovering a short polynomial while having access only to a partial list of its Vandermonde transform can be seen as a problem over an ideal lattice. More precisely, in the mathematical setting above, we know that the ideal generated by q in the m-th cyclotomic ring $O_K = \mathbb{Z}[X]/\Phi_m(X)$ completely splits into d prime ideals, where $\Phi_m(X)$ denotes the m-th cyclotomic polynomial.[4] More precisely, it yields $qO_K = \prod_{j=1}^{d} \mathfrak{p}_j$, where $\mathfrak{p}_j = qO_K + (X - \omega_j)O_K$. Providing the evaluations $g(\omega_j)_{j \in \Omega}$ corresponds to specifying the coset $h := g \bmod I_\Omega$ with respect to the ideal $I_\Omega := \prod_{j \in \Omega} \mathfrak{p}_j$. Hence, PV-Knap essentially requires to recover g (with small coefficients) given h, which yields a problem over ideal lattices.

[3] In this paper, we only consider regimes where the solution to this problem is unique.

[4] For the sake of simplicity, we focus on cyclotomic fields in the introduction but stress that PV-Knap can be defined over any number field.

Contributions. The results of this work can be divided into three different parts. First, we show in Section 3 that Id-SVP can be solved efficiently for ideal lattices with a lot of symmetries, generalizing the results of [PXWC21,PML21]. We then show in Section 4 that there exist *bad* instances of PV-Knap, that are easy to solve using the algorithm above. Last, we present the results of our implemented attacks against different parameter sets and design choices for PV-Knap proposed in the literature (Section 5).

Contribution 1. In [PXWC21], the authors identified a class of "bad ideal lattices", i.e., ideal lattices in which Id-SVP can be solved efficiently with a polynomial approximation factor: *prime* ideals in Galois number fields that are above a prime of $\mathbb{Q}$ splitting into many prime factors. This result was later extended to a larger class of ideals (not necessarily prime) in [PML21, Theorem 3]. However, the characterization of the bad ideal lattices of [PML21] is significantly more complex than the one in [PXWC21], and depends on the algebraic norm of the ideal, as well as some hard to compute quantities, related to the ideal's prime decomposition. In this work, we improve upon those results in two ways:

1. we obtain a very simple sufficient condition for an ideal to be a bad ideal;
2. the class of bad ideals that we obtain from this simple condition contains the ones of [PXWC21] and [PML21], while being strictly larger.

We observe that the condition "a prime ideal is above a prime of $\mathbb{Q}$ splitting into many prime factors" from [PXWC21] can be rephrased more simply as a condition on the prime ideal having many symmetries (this observation was also made in [PML21]). By symmetry we mean here that the prime ideal is fixed (as a set) when applying an automorphism of the number field K in which the ideal lives. With this, we are able to generalize the result of [PXWC21] to *any* ideal (modulo a small condition on their algebraic norm) in *any* number field (not necessarily Galois).

Overall, we obtain the following result (informally stated here, see Theorem 3.1 for a formal statement): one can solve Id-SVP in an ideal lattice I in time roughly $\exp(d/n_I)$, where d is the degree of the number field K and n_I is the number of automorphisms of K that fix I as a set (this is an integer between 1 and d). If I has no symmetries, then $n_I = 1$ (I is always fixed by the identity), and we recover the run times of standard lattice reduction algorithms. This result can also be extended to approximation variants of Id-SVP, leading to an algorithm with approximation factor $\gamma \geq 1$ and time roughly $\exp(d/(n_I \cdot \log(\gamma)))$.

Testing whether an ideal I is fixed by an automorphism τ of K can be done efficiently if we have a description of τ and a basis of I. Contrary to previous works, this does not require any knowledge about the factorization of the ideal I. Hence, our characterization of bad ideals can be easily checked and may be useful to cryptographers introducing new assumptions related to ideal lattices.

We note that [PML21] also provides at the bottom of p.14 a simplified condition for their result, which does not require the knowledge of the factorization of I, but still depends on its algebraic norm. This simplified condition however is quite loose, and our simple condition above captures more ideals. This is for

instance the case for ideals I of norm $\geq 2^d$ which have many symmetries but whose prime factors have individually very few symmetries: our condition shows that these ideals are bad, whereas the condition of [PML21] does not capture them. Looking ahead, this special family of ideals is exactly the one arising when we transform a PV-Knap instance into an Id-SVP instance.

The fact that Id-SVP is easier to solve in lattices fixed by automorphisms of K is not very surprising. Indeed, we know that an element of K fixed by some automorphisms is actually an element from a subfield of K of smaller dimension. The same holds for ideals: an ideal I fixed by n_I automorphisms can be seen as an ideal in a subfield L of K (this formulation requires some care, it is made formal in Lemma 3.3, which is the main new technical material of this contribution), whose degree is exactly d/n_I (the more automorphisms, the smaller the degree of L). When looking for a short vector in I, one can consider I as an ideal of L instead of K, i.e., a lattice of smaller dimension d/n_I.

Finally, we remark that the results of [PXWC21, PML21] in all Galois fields are only mathematical results characterizing bad ideals and not algorithms. Both works then used this mathematical result to provide an Id-SVP algorithm, but they did so only in cyclotomic number fields. Generalizing the algorithm to other number fields was left as an open problem in [PXWC21, Remark 1]. In this work, we provide both the mathematical result (Theorem 3.2) and the algorithm (Theorem 3.1) for all number fields.

We would like to stress again that our algorithm only solves specific instances of Id-SVP. Hence, it does not have any implications to the hardness of structured problems such as Ring-SIS or Ring-LWE, as their hardness is based on the *worst-case* hardness of Id-SVP, and the reductions are only one-way.

Contribution 2. We now explain how the algorithm above can be used to solve some particular instances of PV-Knap in polynomial time. Recall that PV-Knap asks to recover $g \in O_K$ of small coefficients given $g \bmod I_\Omega$ for the ideal I_Ω. Note that it is easy to find a $g' \in O_K$ of unbounded coefficients such that $g' = g \bmod I_\Omega$. Thus, solving PV-Knap essentially requires to find the (unique) element $h' \in I_\Omega$ that is "close" to g', that is $g' - h' = g$. When interpreting the ideal I_Ω as an ideal lattice, this yields an instance of the bounded distance decoding problem (BDD), as we show in Section 4.1. We then argue in Section 4.2 that BDD in any ideal I reduces to SVP in its inverse ideal I^{-1}. To do so, we first use Babai's rounding algorithm to reduce BDD in I to SIVP in its dual $I^\vee$. Then, we use that for ideal lattices SIVP reduces to SVP and that we can go from the dual $I^\vee$ to the inverse I^{-1}. All lattice problems are considered with respect to an approximation factor that we specify for general number fields in Lemma 4.1. We provide simplified parameter conditions for power-of-two and prime cyclotomics (Corollary 4.2 and 4.3). We conclude this part by providing in Section 4.3 concrete choices of Ω for which we obtain a polynomial time algorithm that solves PV-Knap in I_Ω (using the results from Section 3).

Contribution 3. As a third contribution, we implemented the algorithm of Section 3 in SageMath and used the observations of Section 4 to solve PV-Knap over

cyclotomic fields for different choices of Ω. Globally, we tested our attack for two different strategies on how to select Ω. In the first scenario, the set Ω is chosen in an advantageous way (for the attacker) to make the related PV-Knap problem easy. More concretely, we choose Ω so that I_Ω is stable by many automorphisms of the underlying number field K. Our experimental results confirm our asymptotic results from Section 4. Applied to different parameter sets that were proposed in the literature [HPS+14, LZA18], we can solve PV-Knap in few minutes or even in few seconds. In the second scenario, we study the case where Ω is chosen at random. For random Ω, the ideal I_Ω is with overwhelming probability not stable by any non-trivial subgroup of the Galois group of K. Thus, one might think that our algorithm won't improve the cryptanalysis of PV-Knap in this case. Perhaps surprisingly, we can still use our algorithm to distinguish PV-Knap instances from random instances with non-negligible probability. The main idea is to forget some of the i's in the set Ω. In general, reducing the size of Ω makes the problem harder, since our target BDD instance lies now in a denser lattice. However, by carefully discarding some elements of Ω, we may hope to obtain a subset Ω' such that $I_{\Omega'}$ is stable by some non-trivial automorphism, hence reducing the dimension of the ideals by some (small) factor. Overall, we observe that for all sets of parameters that we considered, there is a non-negligible probability to sample a random Ω for which one can reduce the dimension of the lattice problem by a factor 2. Finally, we run a full distinguishing attack on the smaller parameter set of [LZA18], which was supposed to provide 128 bits of security. Using the model of [MW18] for bit-security, we show that this set of parameters actually provides at most 87 bits of security against distinguishing attackers. We describe all results of our experiments in more details in Section 5.

Implications to cryptography.

Id-SVP *algorithm.* As explained above, our Id-SVP algorithm only provides improvement compared to standard lattice reduction algorithms if the ideal I is fixed by at least one non-trivial automorphism of K. This is a strong requirement on the ideal, and we expect that random ideals do not usually satisfy this condition (for most of the natural distributions on ideals, such as uniform ideals of norm bounded by some bound B). We note however that choosing ideal lattices with a lot of symmetries may be tempting for cryptographic constructions, as this may lead to faster algorithms. We see our results as a warning to cryptographers: one should not use ideal lattices with symmetries. The exhibition of bad instances of PV-Knap is an illustration of such misuse of ideal lattices.

Summing up, we believe that cryptographers willing to introduce new assumptions based on Id-SVP should follow the following guidelines:

1. check if the scheme can be modified such that the underlying rank increases from 1 to 2 in order to rely on Mod-SVP instead of Id-SVP;
2. if not possible, use random ideals sampled from one of the distributions for which we have a worst-case to average-case reduction [Gen09, dBDPW20];

3. if also not possible, then avoid the known bad ideals: ideals generated by an element sampled from a Gaussian distribution in a cyclotomic number field [CDPR16, CDW21] or ideals fixed by some non-trivial automorphism of the number field (this work);
4. in both cases, do not rely on the hardness of Id-SVP for approximation factors larger than $2^{\sqrt{d}}$ in cyclotomic fields, with d the degree of the number field [CDW21].

PV-Knap *attacks.* As described above, PV-Knap was first studied in the context of the signature scheme PASS Sign [HPS+14, LZA18]. Its key generation algorithm constructs an instance of PV-Knap (over either cyclic or cyclotomic rings), where the secret key is a ternary polynomial and the public key is given by a partial list of its Fourier/Vandermonde coefficients. Hence, solving the search variant of such PV-Knap instances translates to secret key recovery attacks against PASS Sign. In 2015, Hoffstein and Silverman [HS15] designed a public key encryption scheme called PASS Encrypt whose mathematical building blocks resemble those of PASS Sign. Later, the scheme was slightly modified in order to provide a proof of security with respect to concretely defined hardness assumptions by Boudgoust et al. [BSS22], accessible via one of the author's thesis manuscript [Bou21, Ch. 5+7]. In both variants, the key generation algorithms are the same as for PASS Sign, and thus, solving PV-Knap similarly leads to a secret key recovery attack against PASS Encrypt. Doröz et al. [DHSS20] used PASS Sign to design a signature scheme offering public aggregation of signatures independently issued from different users on different messages, called MMSA(TK). An attacker who is able to recover the secret key of a given "challenge" public key clearly violates the security notion used for aggregate signatures.

We would like to highlight again that our attacks on PV-Knap only impact some specific choices of the set Ω, or decrease the lattice dimension by a factor 2 when Ω is randomly chosen. Hence, they can be prevented by choosing Ω carefully (for instance randomly) and possibly increasing the dimension slightly.

2 Preliminaries

Vectors and matrices are written respectively in bold small letters and bold capital letters. Given a vector $\mathbf{v}$ in $\mathbb{R}^n$ or in $\mathbb{C}^n$, we denote $\|\mathbf{v}\|$ its Euclidean norm (or Hermitian norm if $\mathbf{v}$ has complex coordinates) and $\|\mathbf{v}\|_\infty$ its infinity norm. For a matrix $\mathbf{M}$, we write $\mathbf{M}^T$ for its transposed matrix. By default, we consider matrices with column vectors.

2.1 Number fields

In this section we recall some definitions and properties about number fields and Galois theory that are used in the article. More information can be found in [Mar77, Chapters 2-4 and Appendix B].

A number field K is a field of the form $K = \mathbb{Q}[X]/f(X)$, where $f(X)$ is irreducible over $\mathbb{Q}$. The degree of K is its dimension as a $\mathbb{Q}$-vector space, which is equal to the degree of f (hence, it is always finite). In this article, K and L always refer to number fields, with K of degree d. When $L \subseteq K$, we say that K is a field extension of L and write K/L. We let $[K : L]$ denote the degree of the extension, that is the dimension of K as an L-vector space. The degree of a tower of extensions $K/L/M$ is multiplicative, i.e., $[K : M] = [K : L] \cdot [L : M]$.

Canonical embedding. For a number field K of degree d, we let $\sigma_1, \cdots, \sigma_d$ denote the embeddings of K in $\mathbb{C}$. Using those, we define the canonical embedding of K as $\Sigma_K : K \to \mathbb{C}^d$, where $x \mapsto (\sigma_1(x), \cdots, \sigma_d(x))^T$. The trace $\mathrm{Tr}_K \colon K \to \mathbb{Q}$ is defined as the sum of the embeddings, i.e., for any $x \in K$, we have $\mathrm{Tr}_K(x) = \sum_{j=1}^{d} \sigma_j(x)$. Note that if $K/L/\mathbb{Q}$ is a tower of number fields, then any element x of L is also an element of K, and we can consider both $\Sigma_K(x)$ and $\Sigma_L(x)$. These two vectors are related, since we know (see for instance [Mar77, Theorem 50]) that every complex embedding of L extends to exactly $[K : L]$ complex embeddings of K. Hence, the coordinates of $\Sigma_K(x)$ are the same as the ones of $\Sigma_L(x)$, repeated $[K : L]$ times each. From this, we see that

$$\|\Sigma_K(x)\| = \sqrt{[K : L]} \cdot \|\Sigma_L(v)\|. \tag{2.1}$$

Galois theory. The automorphism group of a field extension K/L, denoted by $\mathrm{Aut}_L(K)$, is the set of all K-automorphisms τ such that $\tau(x) = x$ for all $x \in L$. The number of such automorphisms is always at most the degree of the field extension, that is $|\mathrm{Aut}_L(K)| \leq [K : L]$.

Definition 2.1 (Fixed fields). *Given a field extension K/L and a subgroup H of $\mathrm{Aut}_L(K)$, the fixed field of H is the subfield K_H of K defined by $K_H = \{x \in K \,|\, \tau(x) = x\,, \forall \tau \in H\}$. This fields contains L (i.e., we have $K/K_H/L$).*

The extension K/L is said to be Galois if and only if $|\mathrm{Aut}_L(K)| = [K : L]$.[5] In this case, we can also use the notation $\mathrm{Gal}(K/L)$ to refer to the automorphism group $\mathrm{Aut}_L(K)$, and we call it the Galois group of the extension. When the extension K/L is Galois, Galois theory tells us that there is a one to one correspondence between subgroups of the Galois group $\mathrm{Gal}(K/L)$ and subfields of K containing L (see [Mar77, Theorem 55]). This correspondence is given by the maps $H \subseteq \mathrm{Gal}(K/L) \mapsto K_H$ and $L \subset K' \subseteq K \mapsto \mathrm{Aut}_{K'}(K)$.

Lemma 2.2 ([Lan02, Theorem 1.8, Chapter 6]). *Let K/L be an extension (not necessarily Galois). Then, for any subgroup H of $\mathrm{Aut}_L(K)$, the extension K/K_H is Galois and $\mathrm{Gal}(K/K_H) = H$.*

[5] This is not the standard definition, see for instance [Mar77, Theorem 52] for a proof that this is an equivalent definition.

Ring of integers and discriminant. For a number field $K = \mathbb{Q}[X]/f(X)$, we write O_K its ring of integer, that is the subset of elements of K that are roots of a monic integer polynomial. It can be shown that O_K is a free $\mathbb{Z}$-module of rank d, where d is the degree of the number field. In other words, there exists a basis $r_1, \ldots, r_d \in O_K$ such that every element in O_K can be uniquely represented as an integer linear combination of those vectors. Often, we assume the knowledge of a short basis $r_1, \ldots, r_d$ of O_K, where the shortness is measured with respect to the canonical embedding Σ_K. To ease notations, we define the constant $C_K^\infty = \max_j \|\Sigma_K(r_j)\|_\infty$, which is used in Section 4. It always holds $\mathbb{Z}[X]/f(X) \subseteq O_K$ and for some number fields it also holds $O_K \subseteq \mathbb{Z}[X]/f(X)$ (e.g. for cyclotomic fields, see below). Note that being an integer is a property of the element, that does not depends on the number field. Hence, if K and L are two number fields with $L \subseteq K$, then we have that $O_L = O_K \cap L$.

The (absolute value of the) discriminant of a number field K is defined as $\Delta_K = |\det(\sigma_i(r_j))_{i,j}|^2$, where (r_j) is any basis of O_K. Given a tower of number fields $K/L/\mathbb{Q}$, it holds that $\Delta_K \geq \Delta_L^{[K:L]}$ (cf. [Mar77, Exercise 23]).

Product of sets. Let X and Y be two subsets of the same field K (so that we can add and multiply their elements). We define the product of X and Y by

$$X \cdot Y = \{\sum_{i=1}^{r} x_i y_i \,|\, r \geq 0,\, x_i \in X,\, y_i \in Y\}.$$

Note that this product is well defined for any sets X and Y, and not only ideals. This is useful when we consider ideals of subfields, which are not necessarily ideals in the larger field. The product of two sets enjoys commutative and associative properties: $X \cdot Y = Y \cdot X$ and $(X \cdot Y) \cdot Z = X \cdot (Y \cdot Z)$.

Ideals. An integral ideal I of a number field K is a subgroup of O_K such that $I \cdot O_K = I$. A fractional ideal $J \subset O_K$ is a set of the form $J = 1/D \cdot I$, where $D \in \mathbb{Z}_{>0}$ and I is an integral ideal. By default, we use the word "ideal" to refer to fractional ideals, and we specify "integral ideal" when we restrict ourselves to ideals contained in O_K. For $\alpha \in K$, we denote by $\alpha \cdot O_K = \{\alpha \cdot x \,|\, x \in O_K\}$ the ideal generated by α.

The product of two ideals (using the product of sets defined above) is also an ideal. The set of all non zero ideals forms a group with this product, i.e., for any non-zero ideal I, there exists an ideal I^{-1} such that $I \cdot I^{-1} = O_K$. The norm over K of an integral ideal I is defined as $\mathcal{N}_K(I) = |O_K/I|$, and the norm of a fractional ideal $J = 1/D \cdot I$ (with I integral) is defined as $\mathcal{N}_K(J) = 1/D^d \cdot \mathcal{N}(I)$. The norm function is multiplicative, that is $\mathcal{N}_K(I \cdot J) = \mathcal{N}_K(I) \cdot \mathcal{N}_K(J)$ for every integral ideals I and J of K.

We say that an integral ideal I divides another integral ideal J, denoted by $I|J$, if there exists some integral ideal I' such that $J = I \cdot I'$. This is equivalent to $J \subseteq I$ (see [Mar77, Corollary 3, Theorem 15]).

Proposition 2.3. *Given a tower of number fields $K/L/\mathbb{Q}$, the following holds:*

(1) If I is an integral ideal of K, then $I \cap O_L = I \cap L$ is an integral ideal of L.
(2) If J is an integral ideal of L, then $J \cdot O_K$ is an integral ideal of K.
(3) If J_1, J_2 are integral ideals of L, then $(J_1 \cdot O_K) \cdot (J_2 \cdot O_K) = (J_1 \cdot J_2) \cdot O_K$.
(4) If J is an integral ideal of L, then $\mathcal{N}_K(J \cdot O_K) = \mathcal{N}_L(J)^{[K:L]}$.

Proof. The first two items of the proposition immediately follow from the definition of an integral ideal. The third point is implied by the properties of set multiplication stated above (and the fact that $O_K \cdot O_K = O_K$). Finally, the fourth point is proven for instance in [Mar77, Theorem 22, point (b)]. □

We define the dual of an ideal I by $I^\vee = \{x \in K : \mathrm{Tr}_K(xy) \in \mathbb{Z}, \forall y \in I\}$. For any ideal I, its dual and inverse ideal are related to each other via the dual of the corresponding ring of integers, i.e., $I^\vee = I^{-1} O_K^\vee$ (see for instance [Con]). In the case where the ring of integers O_K is of the form $O_K = \mathbb{Z}[X]/f(X)$ for some irreducible polynomial f, we have $O_K^\vee = f'(X)^{-1} \cdot O_K$.

The following definition introduces the notion of decomposition group and decomposition field of an ideal. These notions are usually only defined for prime ideals (see for instance [Mar77, Chapter 4]), but we generalize the terminology to any ideal, since this is needed for the rest of the article.

Definition 2.4. *Let K be a number field and I be an ideal of K. The decomposition group of I is the subgroup H_I of $\mathrm{Aut}_\mathbb{Q}(K)$ defined by $H_I = \{\tau \in \mathrm{Aut}_\mathbb{Q}(K) \mid \tau(I) = I\}$.[6] The decomposition field of I, denoted by K_I, if the fixed field of H_I (cf. Definition 2.1).*

Prime ideals. A non-zero integral ideal $\mathfrak{p}$ of a number field K is said to be prime if it is maximal, i.e., it is different from O_K and the only ideals that contain it are itself and O_K. Any non-zero integral ideal I in a number field K admits a unique decomposition into prime ideals $I = \prod_{\mathfrak{p} \text{ prime}} \mathfrak{p}^{\alpha_\mathfrak{p}}$, where $\alpha_\mathfrak{p} \geq 0$.

In this article, we are interested in moving prime ideals from a field to a subfield and vice versa. This relates to the terminology of primes lying above or below another prime, as defined in the following lemma.

Lemma 2.5 ([Mar77, Theorem 19]). *Let $K/L/\mathbb{Q}$ be a tower of number fields. Let $\mathfrak{p}$ be a prime ideal of K and $\mathfrak{q}$ be a prime ideal of L. The following conditions are equivalent:*

$$(1)\ \ \mathfrak{p} \cap L = \mathfrak{q} \qquad \text{and} \qquad (2)\ \ \mathfrak{p} | (\mathfrak{q} \cdot O_K).$$

When these conditions hold, we say that $\mathfrak{p}$ lies above $\mathfrak{q}$, or that $\mathfrak{q}$ lies below $\mathfrak{p}$.

Lemma 2.6 ([Mar77, Theorem 20]). *Every prime ideal $\mathfrak{p}$ of K lies above exactly one prime ideal of L. Every prime ideal $\mathfrak{q}$ of L lies below at least one prime ideal of K.*

[6] Note that the equality $\tau(I) = I$ means that the two sets are equal, but it does not mean that all the elements of I are fixed by τ.

If $L = \mathbb{Q}$, this lemma implies that any prime ideal $\mathfrak{p}$ of K lies over exactly one rational prime $q \in \mathbb{Z}$. It then holds that $\mathcal{N}_K(\mathfrak{p}) = q^r$ for some $r \in \{1, \cdots, d\}$.

Let $\mathfrak{p}$ be a prime of K and let $\mathfrak{q}$ be the unique prime of L below $\mathfrak{p}$. We say that $\mathfrak{p}$ is ramified in K/L if $\mathfrak{p}^\alpha | (\mathfrak{q} \cdot O_K)$ for some exponent $\alpha \geq 2$ (by Lemma 2.5, we know that $\alpha \geq 1$). The largest integer α such that $\mathfrak{p}^\alpha | (\mathfrak{q} \cdot O_K)$ is called the ramification index of $\mathfrak{p}$ in K/L. In this article, we are mostly interested in prime ideals that are not ramified. This is the most frequent case, since only a finite number of prime ideals are ramified in K/L (cf. [Mar77, Cor. 3 after Thm. 24]).

Lemma 2.7. *If a prime ideal $\mathfrak{p}$ of K is unramified in $K/\mathbb{Q}$, then it is also unramified in K/L for all subfields L of K containing $\mathbb{Q}$.*

Proof. Since $\mathfrak{p}$ is unramified in $K/\mathbb{Q}$, by definition, it only appears once in the prime decomposition of $(\mathfrak{p} \cap \mathbb{Q}) \cdot O_K$. Moreover, since L contains $\mathbb{Q}$, we know that $(\mathfrak{p} \cap \mathbb{Q}) \cdot O_K \subset (\mathfrak{p} \cap L) \cdot O_K$, i.e., $(\mathfrak{p} \cap L) \cdot O_K$ divides $(\mathfrak{p} \cap \mathbb{Q}) \cdot O_K$. This means that $\mathfrak{p}$ may appear at most once in the prime decomposition of $(\mathfrak{p} \cap L) \cdot O_K$, i.e., $\mathfrak{p}$ is unramified in K/L. □

This observation enables us to discard all possible ramified ideals in any subfield of K, by discarding the ones that are ramified in $K/\mathbb{Q}$. Moreover, we know that if a prime $\mathfrak{p}$ is ramified in $K/\mathbb{Q}$, then it is above some $q \in \mathbb{Q}$ that divides Δ_K.

Lemma 2.8 ([Mar77, Thm. 23]). *Let K/L be Galois. If $\mathfrak{p}$ is a prime ideal of K over a prime ideal $\mathfrak{q}$ of L, then for any $\tau \in \mathrm{Gal}(K/L)$, the ideal $\tau(\mathfrak{p})$ is also a prime ideal of K over $\mathfrak{q}$. Conversely, for any two prime ideals $\mathfrak{p}$ and $\mathfrak{p}'$ of K over the same prime $\mathfrak{q}$ of L, there exists a $\tau \in \mathrm{Gal}(K/L)$ such that $\tau(\mathfrak{p}) = \mathfrak{p}'$.*

Cyclotomic fields. Cyclotomic fields form a special class of number fields. For some integer $m \geq 2$, the m-th cyclotomic field can be described as $K = \mathbb{Q}[X]/\Phi_m(X)$, where its defining polynomial $\Phi_m(X)$ is the m-th cyclotomic polynomial. Its degree equals $\deg(\Phi_m(X)) = \phi(m)$, where $\phi(\cdot)$ is Euler's totient function. If $K = \mathbb{Q}[X]/\Phi_m(X)$ is a cyclotomic field, then $(1, X, \ldots, X^{\phi(m)-1})$ forms a basis of O_K, also called the power basis (cf. [Was82, Theorem 2.6]). In other words, $O_K = \mathbb{Z}[X]/\Phi_m(X)$ and we can set the constant C_K^∞ from above as 1.

All cyclotomic fields are Galois and their Galois group is abelian (cf. [Mar77, Corollary 2, Theorem 3]). The following lemma holds for any finite abelian group. We instantiate it directly with $\mathrm{Gal}(K/\mathbb{Q})$.

Lemma 2.9. *Let K be the m-th cyclotomic number field. For every $r | \phi(m)$, there is a subgroup H of $\mathrm{Aut}_\mathbb{Q}(K)$ of cardinality r.*

Proof. Since $\mathrm{Aut}_\mathbb{Q}(K)$ is a finite abelian group of cardinality $\phi(m)$, it can be decomposed into $\mathrm{Aut}_\mathbb{Q}(K) \simeq \prod_{p|\phi(m)} G_p$, where G_p is a group of cardinality p^{α_p} and $\phi(m) = \prod_p p^{\alpha_p}$. In each p-group G_p, there is a subgroup of cardinality p^{β_p} for any $\beta_p \leq \alpha_p$. Taking the product of these subgroups, one can create a subgroup of $\mathrm{Aut}_\mathbb{Q}(K)$ of any cardinality dividing $\phi(m)$. □

The discriminant of the m-th cyclotomic field K is $\Delta_K = \frac{m^{\phi(m)}}{\prod_{p|m} p^{\phi(m)/(p-1)}} \leq m^{\phi(m)}$ (cf. [Was82, Prop. 2.7]). For m a power of 2, it simplifies to $\Delta_K = \phi(m)^{\phi(m)}$.

2.2 Lattices

For a lattice L, we denote $\lambda_1(L)$ its first minimum, i.e., $\lambda_1(L) = \min_{\mathbf{v} \in L \setminus \{0\}} \|\mathbf{v}\|$. The determinant of L is given by $\det(L) = \sqrt{|\det(\mathbf{B}^T \cdot \mathbf{B})|}$ where $\mathbf{B}$ is any basis of L. Minkowski's theorem states that for any lattice L of rank n, it holds that $\lambda_1(L) \leq \sqrt{n} \cdot \det(L)^{1/n}$. We use the notation $\mathrm{Span}_{\mathbb{R}}(L)$ to refer to the real vector space spanned by the vectors of L. Further, we define the dual lattice of L as $L^\vee = \{x \in \mathrm{Span}_{\mathbb{R}}(L) \colon \langle x, y \rangle \in \mathbb{Z} \ \forall \ y \in L\}$. If $\mathbf{B}$ is a basis of L, then $\mathbf{B}^\vee = (\mathbf{B}^T)^{-1}$ is a basis of $L^\vee$. This implies that $\det(L^\vee) = 1/\det(L)$.

Ideal lattices. When we embed an ideal I of K into $\mathbb{C}^d$ using the canonical embedding, the resulting set $\Sigma_K(I)$ is a lattice of rank d, called an ideal lattice. The determinant of the ideal lattice $\Sigma_K(I)$ is $\det(\Sigma_K(I)) = \mathcal{N}_K(I) \cdot \sqrt{\Delta_K}$. The duality notions of ideals and lattices are closely related. Indeed, it holds that $\Sigma_K(I)^\vee = \overline{\Sigma_K(I^\vee)}$, where $\bar{\cdot}$ denotes the complex conjugation and $\overline{L} := \{\overline{x} \,|\, x \in L\}$ for any lattice $L \subset \mathbb{C}^d$. From this, we see that

$$\det(\Sigma_K(I^{-1})) = \det(\Sigma_K(I)^\vee) \cdot \Delta_K = \det(\Sigma_K(I^\vee)) \cdot \Delta_K. \tag{2.2}$$

In the case of ideal lattices, the minimum of a lattice is closely related to the normalized algebraic norm of the ideal

$$\sqrt{d} \cdot \mathcal{N}_K(I)^{1/d} \leq \lambda_1(\Sigma_K(I)) \leq \Delta_K^{1/(2d)} \cdot \sqrt{d} \cdot \mathcal{N}_K(I)^{1/d}, \tag{2.3}$$

where the first inequality comes from the arithmetic-geometric means inequality applied to a shortest vector of $\Sigma_K(I)$ and the second is Minkowski's theorem.

Algorithmic problems over ideal lattices. In this work, we are interested in three algorithmic problems that we state over ideal lattices: the shortest vector problem (SVP), the shortest independent vector problem (SIVP) and the bounded distance decoding (BDD) problem, all three in their so-called Hermite variant. Whereas in the original formulation those problems are defined with respect to the minimum λ_1 of a lattice L, their Hermite variant phrases them with respect to the determinant $\det(L)$ of the lattice. As we explained above, for ideal lattices both quantities are closely related and only differ by a factor $\Delta_K^{1/(2d)}$ (Equation 2.3). One of the advantages when working with the Hermite variant is that the quantity $\det(L)$ is easier to compute than the quantity $\lambda_1(L)$. The three problems are defined as follows.

Definition 2.10 (γ-Id-HSVP$_K$). *Let $\gamma \geq 1$ and K be a number field of degree d with ring of integers O_K. The γ-Id-HSVP$_K$ problem asks, given as input an ideal I of O_K, to find a non-zero element $v \in I$ such that*

$$\|\Sigma_K(v)\| \leq \gamma \cdot \det(\Sigma_K(I))^{1/d}.$$

This problem always has a solution as long as $\gamma \geq \sqrt{d}$. There exist in the literature different algorithms for solving Id-HSVP$_K$. One is the BKZ algorithm [SE94], which works for all lattices. The run time of (a variant of) this algorithm was formally studied in [HPS11], achieving the following complexity.

Lemma 2.11 ([HPS11, Theorem 1]). *There is a classical probabilistic algorithm that takes as input a basis $\mathbf{B}_L \in \mathbb{Q}^n$ of a lattice L of rank n, a parameter $\gamma \in [\sqrt{n}, 2^n]$, and solves γ-HSVP in L in time* $\mathrm{poly}(n, \mathrm{size}(\mathbf{B}_L)) \cdot 2^{O(n \log(n)/\log(\gamma))}$.

There exist also special algorithms for Id-HSVP, relying on the algebraic properties of the ideals to find short vectors more efficiently. More details about these algorithms may be found in Appendix A (we don't use them in this article).

Definition 2.12 (γ-Id-HSIVP$_K$). *Let $\gamma \geq 1$ and K be a number field of degree d with ring of integers O_K. The γ-Id-HSIVP$_K$ problem asks, given as input an ideal I of O_K, to output d linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_d \in \Sigma_K(I)$ such that $\max_j \|\mathbf{b}_j\| \leq \gamma \cdot \det(\Sigma_K(I))^{1/d}$.*

The Hermite variant of the BDD problem over ideal lattices has no official name in the literature yet, we simply call it Hermite Ideal BDD (or Id-HBDD).

Definition 2.13 (worst-case γ-Id-HBDD$_K$). *For $\gamma > 2\Delta_K^{1/2d}/\sqrt{d}$, K a number field of degree d with ring of integers O_K and I an ideal of O_K, the worst-case γ-Id-HBDD$_K$ is the following. Given as input any $\mathbf{t} \in \mathrm{Span}_{\mathbb{R}}(\Sigma_K(I))$ with the promise that $\mathbf{t} = \mathbf{v} + \mathbf{e}$ with $\mathbf{v} \in \Sigma_K(I)$ and $\|\mathbf{e}\| \leq 1/\gamma \cdot \det(\Sigma_K(I))^{1/d}$, the problem asks to output $\mathbf{v}$.*

Note that the constraint $\gamma > 2\Delta_K^{1/2d}/\sqrt{d}$ ensures that there is a unique $\mathbf{v} \in \Sigma_K(I)$ with $\|\mathbf{v}-\mathbf{t}\| \leq 1/\gamma \cdot \det(\Sigma_K(I))^{1/d} \leq 1/2 \cdot \lambda_1(\Sigma_K(I))$, using Equation 2.3. Hence, the Id-HBDD problem is well defined.[7] The terminology "worst-case" means that we ask an algorithm to be able to solve the problem for all choices of input $\mathbf{t}$ that satisfy the promise.

2.3 Representation and size of algebraic objects

Given a rational number $z = x/y \in \mathbb{Q}$ with x and y coprime integers, we denote by $\mathrm{size}(z)$ the quantity $\log_2 |x| + \log_2 |y|$. Up to a bit of sign, this corresponds to the bit-length needed to represent z. For a matrix $\mathbf{M} = (z_{ij})_{i,j}$ over $\mathbb{Q}$, its $\mathrm{size}(\mathbf{M})$ corresponds to the sum of $\mathrm{size}(z_{ij})$ over all its entries z_{ij}.

Given a number field K of degree d, we often need to assume the knowledge of a basis matrix $\mathbf{B}_K$ of its ring of integers O_K. This basis consists of all the (floating points approximations of the) complex vectors $\Sigma_K(r_i)$, where $(r_i)_i \in$

[7] For arbitrary Euclidean lattices, it is much harder to give concrete conditions which ensure a unique solution for HBDD. This is why we think the definition of this problem only makes sense in the ideal setting.

O_K^d forms a $\mathbb{Z}$-basis of O_K. We use the notation $\text{size}(\mathbf{B}_K) = \max_{i,j}(\log|(\mathbf{B}_K)_{i,j}|)$, where $(\mathbf{B}_K)_{i,j}$ are the coefficients of $\mathbf{B}_K$.

Once the r_i's are fixed, every element x of K can be represented as a rational vector $(x_1, \cdots, x_d)$, such that $x = \sum_i x_i r_i$. This gives us an exact representation for the elements of K, in the basis $(r_i)_i$. Note that an element x is in O_K if and only if the vector $(x_1, \cdots, x_d)$ is in $\mathbb{Z}^d$. For $x \in K$, we let $\text{size}(x)$ denote the size of the vector $(x_1, \cdots, x_d) \in \mathbb{Q}^d$, as defined above (note that this depends on the choice of the r_i's, which are assumed to be fixed once and for all).

An ideal I of K is represent by a $\mathbb{Z}$-basis $(b_1, \cdots, b_d) \in K^d$ (i.e., $I = \{\sum_i x_i b_i \,|\, x_i \in \mathbb{Z}\}$). Every element $b_i \in K$ in the basis is represented by a vector in $\mathbb{Q}^d$, as explained in the previous paragraph. We call basis of I the matrix $\mathbf{B}_I$ whose columns are the vectors corresponding to the b_i's. This is a matrix in $\mathbb{Q}^{d\times d}$ (and in $\mathbb{Z}^{d\times d}$ if I is integral), and we use the notation $\text{size}(\mathbf{B}_I)$ as defined above.

An automorphism $\tau \in \text{Aut}_{\mathbb{Q}}(K)$ is represented by a $d \times d$ matrix $\mathbf{M}_\tau$ whose coefficients are such that $\tau(r_j) = \sum_i (\mathbf{M}_\tau)_{i,j} r_i$ (i.e., the j-th column of $\mathbf{M}_\tau$ corresponds to the coordinates of $\tau(r_j)$ in the basis $(r_i)_i$). Since $\tau(x)$ is an algebraic integer if x is, then $\mathbf{M}_\tau$ has integer coefficients. We let $\text{size}(\tau)$ denote the size of the integral $d \times d$ matrix $\mathbf{M}_\tau$, as defined above.

2.4 The Partial Vandermonde Knapsack Problem

The partial Vandermonde knapsack problem (PV-Knap) was first introduced by Hoffstein et al. [HPS+14][8] and later reformulated over number fields by Lu et al. [LZA18]. As observed by Boudgoust [Bou21, Sec. 5.2], the problem can be phrased as a problem over ideal lattices. We use this formulation in the following. For completeness, we provide an explanation why both, the original and the ideal formulation, are equivalent in Appendix B.

Let K be a number field of degree d with ring of integers O_K. Further, let q be a prime integer such that the ideal generated by q splits in exactly d different prime ideals, i.e., $qO_K = \prod_{j=1}^d \mathfrak{p}_j$, where $\mathfrak{p}_j$ is a prime ideal of norm q. For $t \le d$, we define $\mathcal{P}_t = \{\Omega \subseteq \{1, \ldots, d\} \colon |\Omega| = t\}$. For any $\Omega \in \mathcal{P}_t$, we set $I_\Omega := \prod_{j\in\Omega} \mathfrak{p}_j$, yielding an ideal of norm q^t.

Definition 2.14 (PV-Knap). *Let K, O_K, d, q and t be as above. Fix $\Omega \in \mathcal{P}_t$ and let ψ be a distribution over O_K such that $\max_{a\leftarrow\psi} ||\Sigma_K(a)|| \le B$ for some positive real B fulfilling $2B < \sqrt{d} \cdot q^{t/d}$. Sample $a \leftarrow \psi$. Given $b = a \bmod I_\Omega$, the partial Vandermonde knapsack problem $\text{PV-Knap}_{\Omega,\psi}$ asks to find a.*

The constraint $2B < \sqrt{d} \cdot q^{t/d}$ ensures that there is a unique a in the support of ψ such that $b = a \bmod I_\Omega$. By Eq. 2.3, we know that the minimum of $\Sigma_K(I_\Omega)$ with respect to the Euclidean norm is bounded from below by $\sqrt{d} \cdot \mathcal{N}(I_\Omega)^{1/d} = \sqrt{d} \cdot q^{t/d}$. If there were two solutions $a \neq a' \in O_K$ such that $a = a' \bmod I_\Omega$, then the element $a - a'$ would lie in I_Ω and its Euclidean norm with respect to the canonical embedding would be bounded above by $2B < \sqrt{d} \cdot q^{t/d} \le \lambda_1(\Sigma_K(I_\Omega))$, leading to a contradiction. Hence, the PV-Knap problem is well defined.

[8] Even though they originally called it the partial Fourier recovery problem.

We can also define a decision variant of PV-Knap in the natural way. Given Ω and $b+I_\Omega$, one has to decide whether b was defined as in the problem's definition above or if it was sampled uniformly at random.

Whereas in the above definition PV-Knap is defined over O_K, the problem is in some works (e.g. [HPS+14, HS15, DHSS20]) defined over the cyclic ring $\mathbb{Z}[X]/(X^N-1)$ for some prime integer N. We recall its concrete formulation in Appendix B. Note that those rings are closely connected to prime cyclotomic number fields as the polynomial X^N-1 factors into two particular irreducible polynomials. More precisely, we have $X^N-1=(X-1)\cdot\Phi_N(X)$ and thus there exists an injective ring morphism from $\mathbb{Z}[X]/(X^N-1)$ to $\mathbb{Z}[X]/(X-1)\times\mathbb{Z}[X]/\Phi_N(X)$. Using this morphism, one can transform a PV-Knap instance over the cyclic ring into a PV-Knap instance over $\mathbb{Z}[X]/\Phi_N(X)$, i.e., the ring of integers of a cyclotomic fields. A solution to PV-Knap in $\mathbb{Z}[X]/\Phi_N(X)$ can then be lifted back to $\mathbb{Z}[X]/(X^N-1)$ by guessing the last coordinate in $\mathbb{Z}[X]/(X-1)$. Hence, even though the results of Section 4 are formulated for number fields, they also apply to the original parameter setting of [HPS+14].

In our definition, the bound B is with respect to the canonical embedding Σ_K, whereas in the former works, it was with respect to the coefficient embedding. In most of the number fields used in lattice-based cryptography, we know how to go from one embedding to another. For instance, for the m-th cyclotomic field we obtain a bound B in the canonical embedding by multiplying a bound B' in the coefficient embedding by the factor $\sqrt{m}$. A prominent choice in [LZA18, HPS+14] is $B'=1$, which yields $B=\sqrt{m}$. In the case of power-of-two cyclotomics this bound can be tightened to $B=\sqrt{d}$ where $d=m/2$.

The definition we present doesn't specify how to choose Ω, which we exploit in Section 4 when finding *bad* choices of Ω. This follows the same design choice as [HPS+14, HS15, DHSS20]. Other works [LZA18, BSS22, Bou21] decided to sample Ω uniformly at random over the set $\mathcal{P}_t$, which has an important effect on the performance of our attacks as we elaborate later in Section 5.

3 Easy Instances of Ideal-SVP

The objective of this section is to prove the following theorem, which gives a simple and sufficient condition under which the Id-HSVP$_K$ problem is easy in an ideal lattice. The condition requires the ideal I to have no ramified prime factors. By Lemma 2.7, this is the case if the ideal's algebraic norm is coprime with the discriminant of K. Hence, the condition can be verified easily, without computing the prime factorization of the ideal.

Theorem 3.1. *Let K be a number field of degree d and I be an integral ideal of K whose prime factors are not ramified in $K/\mathbb{Q}$. There is an algorithm that takes as input a basis $\mathbf{B}_K$ of O_K, a representation G of $\mathrm{Aut}_\mathbb{Q}(K)$, a basis $\mathbf{B}_I$ of I and a parameter $\gamma\geq 2\sqrt{d}$ and solves γ-Id-HSVP$_K$ in I in classical time*

$$\exp\Big(O\Big(\frac{d\cdot\log(d)}{n_I\cdot\log(\gamma/\sqrt{n_I})}\Big)\Big)\cdot\mathrm{poly}(\mathrm{size}(\mathbf{B}_I),\mathrm{size}(\mathbf{B}_K),\mathrm{size}(G)),$$

where $n_I := |H_I|$ is the number of K-automorphisms that fix I as a set.

3.1 Reducing the ideal in a subfield

In this section, we ignore the representation of the mathematical objects, and concentrate on the following mathematical result. It states that if an ideal is fixed by a sufficiently large group of automorphisms, then one can find a short vector of it by looking for short vectors of its intersection with a subfield of smaller dimension. Hence, we can reduce the dimension of the problem.

Theorem 3.2. *Let K be a number field and I be an integral ideal of K whose prime factors are not ramified in $K/\mathbb{Q}$. Let $K_I \subseteq K$ be the decomposition field of I (see Def. 2.4). We write $d = [K : \mathbb{Q}]$ and $d_I = [K_I : \mathbb{Q}]$, and we let $\gamma \geq 1$.*

Then, any $v \in I \cap K_I$ which is a solution to γ-Id-HSVP$_{K_I}$ in $I \cap K_I$ is also a solution to γ'-Id-HSVP$_K$ in I, where

$$\gamma' = \gamma \cdot \sqrt{d/d_I}.$$

This generalizes Theorem 4 of [PXWC21] to non-prime ideals I, and to number fields that are not necessarily Galois. The latter is easily obtained from the observation that the extension K/K_I is always Galois, even if $K/\mathbb{Q}$ is not (Lemma 2.2). The generalization to non-prime ideals requires more work. The main difficulty of this generalization lies in proving the following lemma.

Lemma 3.3. *Let K/L be a Galois extension of number fields. Let I be an integral ideal of K whose prime factors are not ramified in K/L. If $\sigma(I) = I$ for all $\sigma \in \mathrm{Gal}(K/L)$, then it holds that*

$$I = (I \cap O_L) \cdot O_K.$$

Intuitively, this lemma means that when intersecting the ideal I with the subfield L, one loses no information on I, since it can be recovered simply by multiplying by O_K again. This conveys the intuition that the short vectors of I should also be contained into the intersection $I \cap O_L$.

Proof. Note that the inclusion $I \supseteq (I \cap O_L) \cdot O_K$ always holds, even if I is divisible by ramified primes, or if $\sigma(I) \neq I$ for some $\sigma \in \mathrm{Gal}(K/L)$. However, in the general case, this inclusion is usually not an equality: the set $(I \cap O_L) \cdot O_K$ can be much sparser than I, hence losing information about I. In the rest of this proof, we focus on proving the reverse inclusion $I \subseteq (I \cap O_L) \cdot O_K$.

First, we group the prime factors of I into groups of primes that are all above the same prime in O_L. In other words, we write $I = \prod_{\mathfrak{q} \text{ prime of } O_L} I_{\mathfrak{q}}$, where $I_{\mathfrak{q}} = \prod_{\mathfrak{p}_i \text{ prime of } O_K \text{ above } \mathfrak{q}} \mathfrak{p}_i^{\alpha_i}$.

Let us fix a prime ideal $\mathfrak{q}$ in O_L, which does not ramify in O_K (recall that we required that the prime factors of I are not ramified in K/L). Since $\mathfrak{q}$ does not ramify, we know that $\mathfrak{q} \cdot O_K = \mathfrak{p}_1 \cdots \mathfrak{p}_r$ for some distinct prime ideals $\mathfrak{p}_i$ of O_K.

Next, since K/L is Galois, we know that $\mathrm{Gal}(K/L)$ acts transitively on the $\mathfrak{p}_i$, i.e., for every indices i, j, there is some $\sigma \in \mathrm{Gal}(K/L)$ such that $\sigma(\mathfrak{p}_i) = \mathfrak{p}_j$. Using

that $\sigma(I) = I$ for all $\sigma \in \mathrm{Gal}(K/L)$ and that the prime decomposition of an ideal is unique, we conclude that all the $\mathfrak{p}_i$ appear with the same exponent in the prime decomposition of I. Hence, $I_\mathfrak{q} = \prod_{i=1}^{r} \mathfrak{p}_i^{\alpha_\mathfrak{q}} = (\mathfrak{q}O_K)^{\alpha_\mathfrak{q}}$, for some $\alpha_\mathfrak{q} \geq 0$.

Summing up, we can write I as a product $I = \prod_i (\mathfrak{q}_i \cdot O_K)^{\alpha_i}$, for some prime ideals $\mathfrak{q}_i$ of O_L and $\alpha_i \geq 1$. We see here that the condition $\sigma(I) = I$ for all $\sigma \in \mathrm{Gal}(K/L)$ (and the fact that I is not divisible by any ramified prime) implies the natural intuition that I is an ideal of O_L, lifted in O_K.

Using this equation, let us now prove that $I \cap O_L = \prod_i \mathfrak{q}_i^{\alpha_i}$. The inclusion $\prod_i \mathfrak{q}_i^{\alpha_i} \subseteq I \cap O_L$ follows from

$$\prod_i \mathfrak{q}_i^{\alpha_i} = \prod_i (\mathfrak{q}_i \cdot O_L)^{\alpha_i} \subseteq \prod_i (\mathfrak{q}_i \cdot O_K)^{\alpha_i} = I.$$

Since $I \cap O_L$ is an ideal of O_L and we have seen that $\prod_i \mathfrak{q}_i^{\alpha_i} \subseteq I \cap O_L$, i.e., $(I \cap O_L) | \prod_i \mathfrak{q}_i^{\alpha_i}$, we know that $(I \cap O_L) = \prod_i \mathfrak{q}_i^{\beta_i}$ for some $\beta_i \leq \alpha_i$. Multiplying this equation by O_K we obtain

$$(I \cap O_L) \cdot O_K = (\prod_i \mathfrak{q}_i^{\beta_i}) \cdot O_K = \prod_i (\mathfrak{q}_i \cdot O_K)^{\beta_i}.$$

We have already seen that $(I \cap O_L) \cdot O_K \subseteq I$. Hence we obtain $\prod_i (\mathfrak{q}_i \cdot O_K)^{\beta_i} \subseteq \prod_i (\mathfrak{q}_i \cdot O_K)^{\alpha_i}$, which holds only if $\beta_i = \alpha_i$ (since $\beta_i \leq \alpha_i$). We then conclude that $(I \cap O_L) = \prod_i \mathfrak{q}_i^{\alpha_i}$ as desired.

Finally, multiplying this equation by O_K, we obtain

$$(I \cap O_L) \cdot O_K = (\prod_i \mathfrak{q}_i^{\alpha_i}) \cdot O_K = \prod_i (\mathfrak{q}_i \cdot O_K)^{\alpha_i} = I,$$

as desired. □

With this lemma at hand, we are now ready to prove Theorem 3.2. This proof follows almost directly the one of Theorem 4 of [PXWC21].

Proof (Proof of Theorem 3.2). Let $v \in I \cap K_I$ be a solution to γ-Id-HSVP$_{K_I}$, that is, v is non-zero and satisfies

$$\|\Sigma_{K_I}(v)\| \leq \gamma \cdot \Delta_{K_I}^{1/(2d_I)} \cdot \mathcal{N}_{K_I}(I \cap K_I)^{1/d_I}.$$

We want to show that $\Sigma_K(v)$ is also a short non-zero vector of the ideal lattice $\Sigma_K(I)$. Clearly, v is non-zero and lies in I, hence we focus on its euclidean norm. Since $v \in K_I$, we know that the coordinates of $\Sigma_K(v)$ are the same as the ones of $\Sigma_{K_I}(v)$, repeated d/d_I times each. Hence we have $\|\Sigma_K(v)\| = \sqrt{d/d_I} \cdot \|\Sigma_{K_I}(v)\|$ (cf. Equation (2.1)). We know from preliminaries that $\Delta_{K_I}^{1/(2d_I)} \leq \Delta_K^{1/(2d)}$. It remains to upper bound $\mathcal{N}_{K_I}(I \cap K_I)^{1/d_I}$.

Let us apply Lemma 3.3 to $L = K_I$. To see that we can indeed apply the lemma, note that K/K_I is Galois thanks to Lemma 2.2. Note also that the prime factors of I are not ramified in $K/\mathbb{Q}$ by assumption, hence, thanks to Lemma 2.7,

they are also not ramified in K/K_I. Finally, note that by definition of K_I, it holds that I is fixed by all automorphisms of $\mathrm{Aut}_{K_I}(K)$. We can then apply Lemma 3.3 and use item (4) of Proposition 2.3 to obtain

$$\mathcal{N}_K(I) = \mathcal{N}_K((I \cap O_{K_I}) \cdot O_K) = \mathcal{N}_{K_I}(I \cap O_{K_I})^{d/d_I} = \mathcal{N}_{K_I}(I \cap K_I)^{d/d_I},$$

where we used the fact that I is integral and so $I \cap K_I = I \cap O_{K_I}$.

Combining everything, we finally obtain

$$\begin{aligned}\|\Sigma_K(v)\| &= \sqrt{d/d_I} \cdot \|\Sigma_{K_I}(v)\| \leq \gamma \cdot \sqrt{d/d_I} \cdot \Delta_{K_I}^{1/(2d_I)} \cdot \mathcal{N}_{K_I}(I \cap K_I)^{1/d_I} \\ &\leq \gamma \cdot \sqrt{d/d_I} \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}_K(I)^{1/d} \\ &= \gamma' \cdot \Delta_K^{1/(2d)} \cdot \mathcal{N}_K(I)^{1/d}.\end{aligned}$$

□

3.2 Proof of Theorem 3.1

We are now ready to prove Theorem 3.1. The algorithm to solve Id-HSVP$_K$ in I is described in Algorithm 3.1. It computes the intersection of I with K_I, in order to reduce the dimension of the lattice, solves Id-HSVP$_{K_I}$ in this lattice of smaller dimension, and then uses Theorem 3.2 to claim that the vector it founds is indeed a solution to Id-HSVP$_K$ in I. The proof below shows its correctness, specifies its run time and the size of the objects that are manipulated.

Algorithm 3.1 Solving Id-HSVP$_K$ in an easy ideal I

Input: A basis of O_K, the group of endomorphisms $\mathrm{Aut}_{\mathbb{Q}}(K)$, an ideal I without ramified prime factors, a parameter $\gamma \geq \sqrt{d}$
Output: A solution to γ-Id-HSVP$_K$ in I

$H_I = \{\}$
for $\tau \in \mathrm{Aut}_{\mathbb{Q}}(K)$ **do**
  Compute a basis of $\tau(I)$.
  if $\tau(I) = I$ **then**
    Add τ to H_I.
  end if
end for
Compute a basis of K_I, the subfield of K fixed by H_I.
Compute a basis of $J = I \cap K_I$
Solve γ'-Id-HSVP$_{K_I}$ in J with $\gamma' = \gamma/\sqrt{|H_I|}$, to obtain an element $x \in J$
return x

Proof (Proof of Theorem 3.1).

Correctness. Since I has no ramified prime factors, we know from Theorem 3.2 that the element x obtained by solving γ'-Id-HSVP$_{K_I}$ in $I \cap K_I$ is also a solution to $(\gamma' \cdot \sqrt{d/d_I})$-Id-HSVP$_K$ in I. Using the fact that K/K_I is a Galois extension, we know that $|H_I| = |\operatorname{Gal}(K/K_I)| = [K : K_I] = d/d_I$. Hence, by choice of γ', we obtain that $\gamma' \cdot \sqrt{d/d_I} = \gamma$. We conclude that x is indeed a solution to γ-Id-HSVP$_K$ in I as desired.

Run time. Observe that the for loop of the algorithm runs at most $|\operatorname{Aut}_{\mathbb{Q}}(K)| \leq [K : \mathbb{Q}] = d$ times. At each iteration of the loop, we need to compute a basis of $\tau(I)$. Recall that we know a basis $(x_1, \cdots, x_d)$ of I, where the elements x_i are represented by their integral vector in the known basis $\mathbf{B}_K = (r_1, \cdots, r_d)$ of O_K. Recall that the automorphisms τ of $\operatorname{Aut}_{\mathbb{Q}}(K)$ are linear transforms that are represented by a integral matrices. Hence, to compute a basis of $\tau(I)$, it is sufficient to multiply the matrix corresponding to τ with the basis matrix of I. This is a multiplication of two integral matrices of dimension d, which can be performed in time polynomial in d and in the bit-size of the entry of the two bases. Testing the equality $\tau(I) = I$ can be done by testing whether each vector of the basis of $\tau(I)$ is in the integer span of the basis of I and conversely. This is again polynomial in d and the bit-size of the entries of the two bases.

Let us now consider the computation of a basis of K_I. This is a real subspace of K of dimension $d_I = d/|H_I|$. This subspace is defined by a collection of linear equations $\tau(x) = x$ for all $\tau \in H_I$. Hence, one can compute a basis of this subspace by computing the kernel (over $\mathbb{Q}$) of a matrix with dimension $|H_I| \times d$ and whose coefficients are integers of bit-size polynomial in the input bit size. This can be done in time polynomial in d and in the bit-size of the coefficients of the matrices corresponding to the automorphisms τ.

Finally, the intersection of a lattice with a rational vector space can be performed in polynomial time (cf. Lemma A.1), and so the basis of $J = I \cap K_I$ can be computed efficiently.

Once J is computed, we run an Id-HSVP$_{K_I}$ solver on it. To do so, we use the BKZ algorithm for which we have concrete run time bounds (cf. Lemma 2.11). This algorithm forgets about the ideal structure of the lattice and simply requires as input a basis of the lattice $\Sigma_{K_I}(J)$. In order to obtain such a basis, we can multiply the basis of J (over $(r_1, \cdots, r_d)$) by the matrix $\mathbf{B}_K$ formed by the (known) embeddings $\Sigma_K(r_1), \cdots, \Sigma_K(r_d)$. This gives us a basis of $\Sigma_K(J)$. In order to obtain a basis of $\Sigma_{K_I}(J)$, we then simply remove the multiple coordinates that appear in $\Sigma_K(J)$. These operations can be performed in polynomial time. The BKZ algorithm with parameter γ' then runs in time poly(input size) $\cdot$ $2^{O(d_I \log(d_I)/\log(\gamma'))}$, since the lattice $\Sigma_{K_I}(J)$ has rank d_I. Note that we used the fact that $\gamma \geq \sqrt{d}$, so that $\gamma' \geq \sqrt{d_I}$ and hence we can indeed apply Lemma 2.11 with parameter γ'. Note also that since $\gamma \geq 2\sqrt{d}$, we have $\gamma' = \gamma/\sqrt{|H_I|} \geq 2$, hence $\log(\gamma')$ is not zero and we can indeed divide by it. □

In the proof of Theorem 3.1, we decided to use the BKZ algorithm to solve the Id-HSVP instance in J, which enables us to have concrete values for the run time of the algorithm in the theorem statement. One might want to use here

another Id-HSVP algorithm instead of BKZ. However, among the three classes of algorithms mentioned in Section 2.3 and in Appendix A that solve Id-HSVP, it seems that only the BKZ algorithm can be used easily here. Indeed, the solver from [CDW17, CDW21] only works for cyclotomic fields. In our case, even if we restrict to K being cyclotomic, its subfield K_I may not be cyclotomic, so we cannot always use this solver. In the case of the solver from [PHS19], it requires an exponential pre-processing on the number field. This might be a reasonable assumption when the number field is fixed, but in our case the number field K_I depends on the input ideal I. Hence we would need to perform the pre-processing every time, which would be worse than just running the BKZ algorithm.

4 Easy Instances of Partial Vandermonde Knapsack

In this section, we explain how one can reduce the problem of recovering the secret element of a PV-Knap instance to the problem of finding a short vector in the ideal lattice I_Ω^{-1}, depending on Ω. We conclude the section by remarking that, for some choices of the set Ω, the ideal I_Ω^{-1} is stabilized by a large subgroup of the automorphism group of K, leading to an efficient SVP solver in I_Ω^{-1}, and hence to an efficient attack against PV-Knap (for these specific choices of Ω).

4.1 PV-Knap as an Instance of Ideal Hermite BDD

Recall the definition of the partial Vandermonde knapsack problem (PV-Knap) as introduced in Section 2.4 and the definitions of some algorithmic problems over ideal lattices (Id-HSVP, Id-HSIVP, Id-HBDD) as introduced in Section 2.2. Let ψ denote a B-bounded distribution over the ring of integers O_K with respect to the canonical embedding and the Euclidean norm, i.e., $\max_{a\leftarrow\psi} ||\Sigma_K(a)|| \leq B$. Further, let $b = a \bmod I_\Omega$ be an instance of PV-Knap$_{\Omega,\psi}$. Recall that $I_\Omega = \prod_{j\in\Omega} \mathfrak{p}_j$ where the $\mathfrak{p}_j$ come from the prime ideal factorization of the ideal qO_K and $\Omega \subseteq \{1,\dots,d\}$ with $|\Omega| = t$. It follows from the definition that this is exactly an instance of γ_1-Id-HBDD$_K$ for the ideal lattice I_Ω, with

$$\gamma_1 = \frac{\det(\Sigma_K(I_\Omega))^{1/d}}{B} = \frac{q^{t/d}\cdot \Delta_K^{1/(2d)}}{B}.$$

4.2 Reduction from Ideal Hermite BDD to Ideal Hermite SVP in the Inverse Ideal

We now show a sequence of reductions that overall reduce Id-HBDD for an ideal I to Id-HSVP in its inverse ideal I^{-1}.

Lemma 4.1. *Let $K = \mathbb{Q}[x]/f(X)$ be a number field of degree d and discriminant Δ_K with $f(X)$ its defining polynomial and let $(r_1,\dots,r_d)$ be a known basis of O_K. Let $\gamma_1, \gamma_4 > 0$ be such that*

$$\gamma_1 > \gamma_4 \cdot 2\Delta_K^{1/d} \cdot C_K^\infty,$$

where $C_K^\infty = \max_j \|\Sigma_K(r_j)\|_\infty$. *For any fractional ideal* I *in* K, *there is a (deterministic) polynomial-time reduction from* γ_1-Id-HBDD$_K$ *in* I *to* γ_4-Id-HSVP$_K$ *in* I^{-1}.

If in addition $O_K = \mathbb{Z}[X]/f(X)$, *then the* γ_1 *can even be as small as*

$$\gamma_1 > \gamma_4 \cdot 2\Delta_K^{1/d} \cdot \|\Sigma_K(1/f'(X))\|_\infty \cdot C_K^\infty.$$

(Note that this improves upon the previous bound only if $\|\Sigma_K(1/f'(X))\|_\infty < 1$.*)*

In the case of power-of-two and prime cyclotomics, the parameter conditions simplify to the following.

Corollary 4.2. *Let* K *be the* m*-th cyclotomic number field, where* m *is a power of two, of degree* $d = m/2$. *There is an efficient reduction from* γ_1-Id-HBDD$_K$ *in* I *to* γ_4-Id-HSVP$_K$ *in* I^{-1}, *as long as*

$$\gamma_1 > 2\gamma_4.$$

Proof. Using the power basis implies $C_K^\infty = 1$ for all cyclotomic fields and for power-of-two cyclotomics it yields $\Delta_K^{1/d} = d$. Further, $O_K = \mathbb{Z}[X]/f(X)$ with $f(X) = X^d - 1$ and thus $f'(X) = d \cdot X^{d-1}$, completing the proof. □

Corollary 4.3. *Let* K *be the* m*-th cyclotomic number field, where* $m \geq 2$ *is a prime, of degree* $d = m - 1$. *There is an efficient reduction from* γ_1-Id-HBDD$_K$ *in* I *to* γ_4-Id-HSVP$_K$ *in* I^{-1}, *as long as*

$$\gamma_1 > 4\gamma_4.$$

Proof. Again, the power basis leads to $C_K^\infty = 1$ and for prime cyclotomics it yields $\sqrt{m} \leq \Delta_K^{1/d} \leq m$. Furthermore, $O_K = \mathbb{Z}[X]/f(X)$ with $f(X) = \frac{X^m-1}{X-1}$ and thus $f'(X) = \frac{mX^{m-1}\cdot(X-1)-(X^m-1)}{(X-1)^2} = \frac{mX^{m-1}}{X-1}$ (since $X^m = 1$ by definition), leading to $\|\Sigma_K(1/f'(X))\|_\infty \leq \frac{2}{m}$ and thus completing the proof. □

We prove Lemma 4.1 in the following three steps.

Step 1: From Id-HBDD in I to Id-HSIVP in $I^\vee$. This reduction is well-known for BDD and SIVP in their standard formulation and works for any lattice, not only for ideal lattices. It corresponds to solving BDD in a lattice L by using the so-called Babai's rounding algorithm [Bab86], whose performance can be assessed by looking at the size of the vectors of the dual basis of $L^\vee$ (see for instance [CDPR16, Claim 2.1]). For completeness, we detail out how to proceed for the Hermite variant, and quantify the loss in the approximation factor for this variant in the following.

Lemma 4.4 (Id-HBDD to Id-HSIVP). *Let* I *be a fractional ideal of a number field* K *of degree* d. *There is a (deterministic) polynomial-time reduction from* γ_1-Id-HBDD$_K$ *in* I *to* γ_2-Id-HSIVP$_K$ *in* $I^\vee$, *for any* $2\gamma_2 < \gamma_1$.

Proof. Let I be an ideal lattice and let $\mathbf{t} \in \mathrm{Span}_{\mathbb{R}}(\Sigma_K(I))$ be an instance of γ_1-Id-HBDD$_K$, with the promise that $\mathbf{t} = \mathbf{v} + \mathbf{e}$ with $\mathbf{v} \in \Sigma_K(I)$ and $\|\mathbf{e}\| \leq 1/\gamma_1 \cdot \det(\Sigma_K(I))^{1/d}$. Further, assume that we are able to solve γ_2-Id-HSIVP$_K$ in the dual ideal $I^\vee$, i.e., we obtain d linearly independent vectors $\mathbf{b}_j^\vee \in \Sigma_K(I^\vee)$ for $j \in \{1, \ldots, d\}$ such that $\max_j \|\mathbf{b}_j^\vee\| \leq \gamma_2 \cdot \det(\Sigma_K(I^\vee))^{1/d}$.

Recall that $\Sigma_K(I)^\vee = \overline{\Sigma_K(I^\vee)}$, hence the vectors $\overline{\mathbf{b}_j^\vee}$ are in $\Sigma_K(I)^\vee$ (and have the same euclidean norm as the $\mathbf{b}_j^\vee$). For every j, by the definition of the dual lattice, the inner product $\langle \overline{\mathbf{b}_j^\vee}, \mathbf{t} \rangle$ belongs to $\mathbb{Z} + c_j$, where

$$|c_j| \leq \|\overline{\mathbf{b}_j^\vee}\| \cdot \|\mathbf{e}\| \leq \gamma_2/\gamma_1 \cdot \det(\Sigma_K(I^\vee))^{1/d} \cdot \det(\Sigma_K(I))^{1/d} \leq \gamma_2/\gamma_1 < 1/2.$$

Hence, one can round $\langle \overline{\mathbf{b}_j^\vee}, \mathbf{t} \rangle$ to the nearest integer to recover $\langle \overline{\mathbf{b}_j^\vee}, \mathbf{v} \rangle$. Doing this for all j's, we can then recover $\mathbf{v}$ by linear algebra.

We note again that this procedure, while expressed via the dual and in Hermite forms, is simply Babai's rounding algorithm performed with the basis $(\mathbf{b}_j)_j$ of $\Sigma_K(I)$,[9] defined as the dual basis of $(\overline{\mathbf{b}_j^\vee})_j$. □

Step 2: From Id-HSIVP in $I^\vee$ to Id-HSVP in $I^\vee$. This reduction step is special to ideal lattices, as it uses the fact that in the ideal case one short vector is enough to generate a set of linearly independent short vectors.

Lemma 4.5 (Id-HSIVP **to** Id-HSVP**).** *Let I be a fractional ideal of a number field K of degree d. Furthermore, let $r_1, \ldots, r_d \in O_K$ be a known basis of O_K. There is a (deterministic) polynomial-time reduction from γ_2-Id-HSIVP$_K$ in $I^\vee$ to γ_3-Id-HSVP$_K$ in $I^\vee$, where $\gamma_2 = C_K^\infty \cdot \gamma_3$ and $C_K^\infty = \max_j \|\Sigma_K(r_j)\|_\infty$.*

Proof. Assume that we are able to solve γ_3-Id-HSVP$_K$ for the ideal $I^\vee$, i.e., we obtain an element $x \in I^\vee$ of norm $\|\Sigma_K(x)\| \leq \gamma_3 \cdot \det(\Sigma_K(I^\vee))^{1/d}$. Since $I^\vee$ is an ideal and since we know a basis $(r_i)_i$ of O_K, we can transform this single short element into d linearly independent ones: $r_i \cdot x \in I^\vee$, for $i = 1$ to d. These elements satisfy

$$\|\Sigma_K(r_i \cdot x)\| \leq \|\Sigma_K(r_i)\|_\infty \cdot \|\Sigma_K(x)\| \leq C_K^\infty \cdot \gamma_3 \cdot \det(\Sigma_K(I^\vee))^{1/d}.$$

This solves γ_2-HSIVP in $I^\vee$. □

For a given number field K, the constant C_K^∞ is determined by the quality of a short basis for the ring of integers O_K with respect to the infinity norm that we are able to compute. Note that for cyclotomic fields, we know how to find a basis of infinity norm 1 (the power basis) and thus in this case $\gamma_2 = \gamma_3$.

[9] If $(\mathbf{b}_j^\vee)_j$ is not a basis of $\Sigma_k(I^\vee)$, it might be that $(\mathbf{b}_j)_j$ is only a basis of a slightly denser lattice containing $\Sigma_K(I)$. However, the constraints on the parameters imply that $\mathbf{t}$ is still a BDD instance in this denser lattice, with a solution in $\Sigma_K(I)$, hence this does not impact the reasoning.

Step 3: From Id-HSVP in $I^\vee$ to Id-HSVP in I^{-1}. In the last step, we go from the dual to the inverse ideal. This step is motivated from the fact that the shape of I_Ω coming from an instance of PV-Knap is very similar to the shape of its inverse $I_\Omega^{-1} = \frac{1}{q} I_{\Omega^c}$.

Lemma 4.6 (Id-HSVP in $I^\vee$ to Id-HSVP in I^{-1}). *Let I be a fractional ideal of a number field $K = \mathbb{Q}[x]/f(X)$ of degree d and discriminant Δ_K with $f(X)$ its defining polynomial. There is an efficient reduction from γ_3-Id-HSIVP$_K$ in $I^\vee$ to γ_4-Id-HSVP$_K$ in I^{-1}, for any $\gamma_3, \gamma_4 > 0$ such that $\gamma_3 = \gamma_4 \cdot \Delta_K^{1/d}$.*

Furthermore, if $O_K = \mathbb{Z}[X]/f(X)$, the reduction also holds for any $\gamma_3, \gamma_4 > 0$ such that $\gamma_3 = \gamma_4 \cdot \Delta_K^{1/d} \cdot \|\Sigma_K(1/f'(X))\|_\infty$.

Proof. Assume that we are able to solve γ_4-Id-HSVP$_K$ for the ideal I^{-1}, i.e., we obtain an element $x \in I^{-1}$ of norm $\|\Sigma_K(x)\| \leq \gamma_4 \cdot \det(\Sigma_K(I^{-1}))^{1/d}$. By the definition of the inverse and dual of I, it yields that $I^{-1} \subseteq I^\vee$ and thus the short vector $\Sigma_K(x)$ is already an element of the ideal lattice $\in \Sigma_K(I^\vee)$. As it yields that $\det(I^{-1}) = \Delta_K \cdot \det(I^\vee)$ (Equation 2.2), this vector solves $\gamma_4 \cdot \Delta_K$-Id-HSVP$_K$ in $I^\vee$, which proves the first part of the lemma.

Assume now that $O_K = \mathbb{Z}[X]/f(X)$ for some irreducible polynomial $f(X)$. In this specific case, it holds that $I^\vee = I^{-1} \cdot O_K^\vee$ with $O_K^\vee = 1/f'(X) \cdot O_K$. Thus, we can multiply x by the element $1/f'(X)$ and still obtain an element in $I^\vee$. Overall, we obtain a vector $\Sigma_K(x \cdot 1/f'(X))$ whose norm is bounded above by $\gamma_4 \cdot \Delta_K^{1/d} \cdot \|\Sigma_K(1/f'(\zeta))\|_\infty \cdot \det(\Sigma_K(I^\vee))^{1/d}$, which proves the second part of the lemma. □

4.3 Bad Choices of Ω

We now elaborate on how the above results lead to polynomial-time attacks against PV-Knap for some special choices of Ω. In the following, we restrict ourselves to number fields K that are cyclotomic with a conductor m which is either a power of two or a prime integer. These are the number fields used in the literature on PV-Knap, and restricting to these number fields simplifies our attack. Recall that we write $d = \phi(m)$ for the degree of K.

Let q, t, B and $\Omega \in \mathcal{P}_t$ be PV-Knap parameters satisfying $q^{t/d} \geq 8 \cdot B$ (note that this condition is slightly stronger than the condition required in Definition 2.14 of PV-Knap for the problem to be well defined).

Combining Theorem 3.1 and Corollaries 4.2 and 4.3, we obtain a solver for PV-Knap that runs in classical time

$$\exp\left(O\left(\frac{d\log(d)}{|H_{I_\Omega^{-1}}|\log(\gamma_4/|H_{I_\Omega^{-1}}|)}\right)\right) \cdot \mathrm{poly}\left(d, \log q\right), \tag{4.1}$$

where

$$\gamma_4 = \frac{q^{t/d}\sqrt{d}}{4B} \geq 2\sqrt{d}.$$

The last inequality comes from our lower bound on $q^{t/d}$ and is required to apply Thm. 3.1. Recall that for an ideal I the integer $|H_I|$ denotes the number of K-automorphisms that fix I as a set. By definition all ideals I_Ω are unramified.

In the rest of this section, we show that if $t \geq d/2$, then there are choices of Ω that make $|H_{I_\Omega^{-1}}|$ linear in the degree d of the number field, hence leading to polynomial-time attacks against PV-Knap for this choice of Ω. We also explain how the result degrades for smaller choices of t.

Special structure of I_Ω. First of all, we observe that an automorphism $\tau \in \mathrm{Aut}_\mathbb{Q}(K)$ fixes a fractional ideal I if and only if it fixes its inverse I^{-1}. Hence, we only focus here on the group of automorphisms H_{I_Ω} fixing I_Ω, instead of $H_{I_\Omega^{-1}}$.

Recall that I_Ω has a special structure, it is equal to $\prod_{i\in\Omega} \mathfrak{p}_i$, where the $\mathfrak{p}_i$'s are all distinct prime ideals above some fully splitting prime q. Recall also that cyclotomic fields are Galois, hence we can apply Lemma 2.8, which implies that

$$\{\mathfrak{p}_i \,|\, 1 \leq i \leq d\} = \{\tau(\mathfrak{p}_1) \,|\, \tau \in \mathrm{Aut}_\mathbb{Q}(K)\},$$

where $\mathfrak{p}_1$ is any of the prime ideals above q. Let us fix such a prime ideal $\mathfrak{p}_1$. From the equation above, we know that for any subgroup $H \subseteq \mathrm{Aut}_\mathbb{Q}(K)$, there exists a set $\Omega_H \subset \{1, \ldots, d\}$ with $|\Omega_H| = |H|$ such that

$$\{\tau(\mathfrak{p}_1) \,|\, \tau \in H\} = \{\mathfrak{p}_i \,|\, i \in \Omega_H\}.$$

Note that the set Ω_H also depends on the choice of $\mathfrak{p}_1$, but this choice has no impact on our attack, hence we do not mention it in the notation.

By definition of Ω_H, it holds that $I_{\Omega_H} = \prod_{i\in\Omega_H} \mathfrak{p}_i = \prod_{\tau\in H} \tau(\mathfrak{p}_1)$ is fixed by H. The same equation also shows that I_{Ω_H} is not fixed by any strictly larger group of automorphisms containing H.

To conclude, we have a way, given any subgroup H of $\mathrm{Aut}_\mathbb{Q}(K)$, to construct a subset $\Omega_H \in \{1, \cdots, d\}$ such that $|\Omega_H| = |H|$ and $H_{I_{\Omega_H}} = H$.

Subgroups of $\mathrm{Aut}_\mathbb{Q}(K)$ of the desired size. Recall that the set Ω of the PV-Knap instance has to have size t. If there exists a subgroup H of $\mathrm{Aut}_\mathbb{Q}(K)$ with size t, then the previous paragraph shows that one can find bad sets Ω of size t with $|H_{I_\Omega}| = t$. This leads to an attack against PV-Knap for those bad sets Ω whose run time is $\exp\left(O\left(\frac{d}{t}\right)\right) \cdot \mathrm{poly}\,(d, \log q)$. It is polynomial if $t = \Omega(d)$, as is usually the case in PV-Knap parameter sets (see for instance Section 5).

If there is no subgroup H of $\mathrm{Aut}_\mathbb{Q}(K)$ of size t, one can choose a subgroup H of maximal cardinality, subject to $|H| \leq t$. This provides a set $\Omega' = \Omega_H$ of cardinality $|H| \leq t$ such that $I_{\Omega'}$ is fixed by H. This set Ω' does not have the desired size. However, we observe that one can always transform a PV-Knap instance with respect to Ω into a PV-Knap instance with respect to Ω' for any $\Omega' \subseteq \Omega$. This is done by "forgetting" the value of $a \bmod \mathfrak{p}_i$ for the i's in $\Omega \setminus \Omega'$. Another way to phrase this is to observe that if $\Omega' \subset \Omega$, then I_Ω is a sublattice of $I_{\Omega'}$. Hence, we can view any BDD instance in I_Ω as a BDD instance in $I_{\Omega'}$, provided that the volume of ideal $I_{\Omega'}$ is not too small (so that the BDD instance is still close to a unique point of the ideal $I_{\Omega'}$).

This shows that, even when there are no subgroups H of $\mathrm{Aut}_{\mathbb{Q}}(K)$ with size t, one can find bad sets Ω of size t containing a subset Ω' fixed by some subgroup $H \subseteq \mathrm{Aut}_{\mathbb{Q}}(K)$ of cardinality

$$t_0 = \max(|H| \,:\, H \text{ subgroup of } \mathrm{Aut}_{\mathbb{Q}}(K) \text{ and } |H| \leq t).$$

If $q^{t_0/d} \geq 8 \cdot B$, we can solve PV-Knap for Ω in time $\exp\left(O\left(\frac{d}{t_0}\right)\right) \cdot \mathrm{poly}\,(d, \log q)$.

Finally, let us estimate the quantity t_0. We know from Lem. 2.9 that $\mathrm{Aut}_{\mathbb{Q}}(K)$ contains subgroups of any order dividing $\phi(m)$. Hence, one can take

$$t_0 = \max(r \,:\, r | \phi(m) \text{ and } r \leq t\}.$$

In the case of power-of-two cyclotomic fields, this means that we always have $t_0 \geq t/2$. Hence, if $t = \Omega(d)$, there always exist bad sets Ω for which the attack runs in polynomial time (provided that $q^{t/(2d)} \geq 8 \cdot B$).

In the case of prime conductors m, we know that $\phi(m) = d$ is odd, hence if $t \geq d/2$, then we have $t_0 \geq d/2$ and there also exist bad sets Ω for which the attack runs in polynomial time.

5 Experimental Results

We implemented the attack described in Section 4 in SageMath [The20] to solve easy instances of PV-Knap over cyclotomic fields. The code is available at `https://github.com/apelletm/easy-PV-knap`.

We tested our attack in two significantly different scenarios. In the first one, the set Ω of the PV-Knap instance is fixed to make the problem easy (i.e., by choosing Ω such that I_Ω is stable by a lot of automorphisms of K, cf. Section 4.3). In the second scenario, we consider randomly chosen sets Ω.

Our results show that the easy cases are indeed easy: if Ω is badly chosen, one can solve PV-Knap (in both its search and decision versions) in a few seconds. Perhaps surprisingly, we observe that our attack can also be beneficial for randomly chosen sets Ω, for the decision variant of PV-Knap.

Generation of PV-Knap *instances.* We decided to generate PV-Knap instances whose parameters are as suggested in [HPS+14] and [LZA18]. These parameters are summarized in Table 1 below. All number fields are cyclotomic, m is the conductor of the cyclotomic field K, d is the degree of K, $t = |\Omega|$ is the size of Ω and q is a rational prime that fully splits in K. The last line of the table contains the security estimates provided in [HPS+14, LZA18] for these parameters.

As explained above, we consider two types of PV-Knap instances. The first type is what we call *worst-case instances*, where we choose the set Ω so that the ideal I_Ω is stable by many automorphisms of the number field K. For this case, the user can choose the size of the subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ fixing I_Ω.

The second type of instances we generate are what we call *random instances*. In this case, the set Ω is sampled uniformly at random among all the subsets of $\{1, \ldots, d\}$ of size t.

	LZA 1	LZA 2	HPSSW 1	HPSSW 2	HPSSW 3	HPSSW 4
m	1024	2048	433	577	769	1153
d	512	1024	432	576	768	1152
t	256	512	200	280	386	600
q	65537	65537	775937	743177	1047379	968521
estimated bit security	128	128	$\ll 62$	$\ll 80$	< 100	≤ 130

Table 1. Parameter sets used for the attack

Regarding historical choices, [LZA18] suggested taking the set Ω uniformly at random, while [HPS+14] seems to assume that Ω can be chosen arbitrarily (and fixed once and for all). Here, we consider all sets of parameters in both regimes where Ω is arbitrary or uniformly chosen.

In both cases, the secret element a and public element b were computed in the same way: we sample $a \in O_K$ uniformly with coefficients in $\{-1, 0, 1\}$ (note that we consider the coefficient embedding of a here, to be consistent with the way PV-Knap instances are described in [HPS+14, LZA18]). We then set $b = a \bmod I_\Omega$. [10]

Worst-case instances of Ω. In these experiments, we choose Ω so that I_Ω is stabilized by a large subset of $\mathrm{Aut}_{\mathbb{Q}}(K)$, as explained in Section 4.3. Note that for the HPSSW parameter sets, we do not have subgroups of $\mathrm{Aut}_{\mathbb{Q}}(K)$ of size exactly t. Hence, we use the technique described above: we take $t_0 = \max(r : r|\phi(m)$ and $r \leq t\}$, a bad set Ω' fixed by a subgroup of order t_0, and run the attack with this set Ω'.

In Table 2 below, we summarize some of the parameters related to the attack. Note that the quantity t_0 is always equal to either $d/2$ or $d/3$, hence we are in a regime where the lattice reduction step can be performed in dimension 2 or 3. Recall that the quantity B is an upper bound on the size of $\|\Sigma_K(a)\|$. In our case, since a has ternary coefficients, this is upper bounded by $\sqrt{d}$.

Recall that our attack from Section 4.3 was proven to work when $q^{t_0/d} \geq 8B$. This condition is not always satisfied for our parameter sets, however, we observed that in practice, the attack works for all parameter sets, even when the condition was not satisfied. This is not so surprising since the condition is a sufficient condition for the attack to provably work, but not a necessary one.

For each set of parameters described in Table 1, we performed 20 tests of our search and decision attacks, for an optimal set Ω (optimal for the attack, i.e., containing a subset Ω' fixed by a group of automorphisms of size t_0). The search and decision attacks both succeeded with probability 1 on all cases. They

[10] For the case of HPSSW parameters, the generation of a is slightly different, in order to be consistent with the specifications of [HPS+14]. They consider PV-Knap instances over the cyclic ring $\mathbb{Z}[X]/(X^m - 1)$ instead of O_K. For this specific case, we generate a with ternary coefficients in the ring $\mathbb{Z}[X]/(X^m - 1)$, and then reduce it modulo $\Phi_m(X)$ in order to map it to O_K and continue the attack in O_K, cf. Sec. 2.4.

	LZA 1	LZA 2	HPSSW 1	HPSSW 2	HPSSW 3	HPSSW 4
t	256	512	200	280	386	600
t_0	256	512	144	192	384	576
$q^{t_0/d}$	256.0	256.0	91.9	90.6	1023.4	984.1
$8 \cdot B$	181.0	256	166.3	192.2	221.8	271.6

Table 2. Some quantities related to the attack

took between 5 seconds for the smallest sets of parameters and 2 minutes for the largest ones, on a personal laptop (the timings are for performing the 20 tests, but the short vector in I_Ω^{-1} is computed only once).

For the large sets of parameters LZA 2 and HPSSW 4, we also tried the attack with not so optimal sets Ω: we chose Ω so that I_Ω was stable by a subset of $\mathrm{Aut}_\mathbb{Q}(K)$ of size 16, instead of the optimal subsets of size 512 and 576 respectively.[11] This means that the SVP instance we had to solve was in dimension 64 and 72 respectively (instead of dimension 2). Even in this less favorable scenario, the search attack succeeded with probability 1 over the 20 tests, and it ran in 2 minutes and 4 minutes respectively. Note that recovering the secret a already solves the decision variant of PV-Knap as well.

Our conclusion is that the easy instances of PV-Knap that we identified are really easy (solved in less than a few minutes on a personal laptop), even for number fields of large degree and concrete parameter sets, and even when the condition $q^{t_0/d} \geq 8B \cdot d^{3/2}$ is not satisfied. Hence, the choice of the set Ω should absolutely not be given to the attacker.

This worst-case attack can be considered to break (at least partially) the PV-Knap settings suggested in [HPS+14], since it wasn't specified how the set Ω should be chosen. For [LZA18], the authors require Ω to be uniformly sampled, hence the worst-case attack cannot be considered to break their settings.

Random choices of Ω: estimating the cost of lattice attacks. We now consider the cases where the set Ω is chosen uniformly at random among all sets of size t. In this situation, it is very unlikely that the ideal I_Ω is stable by any non trivial subgroup of the Galois group. Even for a subgroup of order 2, the probability that I_Ω is stable by this subgroup is roughly equal to $1/2^t$. Indeed, let $\tau \in \mathrm{Aut}_\mathbb{Q}(K)$ be an element of order 2 (i.e., $\tau(\tau(x)) = x$ for all x). The ideal I_Ω is stable by τ if and only if, for every $i \in \Omega$, we have $j \in \Omega$ where j is such that $\mathfrak{p}_j = \tau(\mathfrak{p}_i)$. Since Ω is chosen uniformly at random, the probability that $j \in \Omega$ is roughly $1/2$.

Even though I_Ω is very unlikely to be stabilized by a non trivial subgroup of $\mathrm{Aut}_\mathbb{Q}(K)$, we can still try to apply our attack here. The idea is always the same: we can forget some of the i's in Ω. As we have already seen, reducing the

[11] Note that here, we do not reduce the size of Ω below t_0: we take Ω as the union of multiple sets Ω', each one of size 16 such that $I_{\Omega'}$ is fixed by a subgroup H of $\mathrm{Aut}_\mathbb{Q}(K)$ of size 16 (the same H for all the Ω').

size of Ω by forgetting some of the i's makes the problem harder, since our target is a BDD instance in a denser lattice, and at some point the solution will not even be unique anymore. On the other hand, by discarding some of the elements of Ω in a carefully chosen way, we may hope to obtain a subset Ω' such that $I_{\Omega'}$ is stable by some non trivial subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$.

Our objective is then to reduce Ω to some subset Ω' sufficiently large such that b is still a BDD instance in $I_{\Omega'}$, but with $I_{\Omega'}$ stabilized by a subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ as large as possible. The objective of our experiments in this paragraph was to estimate by how much one can hope to reduce the lattice dimension by using this technique. In other words, what is the largest subset of $\mathrm{Aut}_{\mathbb{Q}}(K)$ that stabilizes a sufficiently large subset Ω' of Ω (so that the problem is still well defined with Ω')?

To estimate this quantity, we proceed in two steps. We first estimate the minimal size of Ω' that we can allow for the distinguishing attack to succeed. This is done experimentally, by estimating the size B of a shortest vector in $qI_{\Omega'}^{-1}$ for Ω' of a given length (note that the volume of $qI_{\Omega'}^{-1}$ is equal to $q^{d-|\Omega'|} \cdot \Delta_K^{1/2}$, which only depends on the size of Ω' and not the actual choice of Ω'). We then compute a short element v of length B and experimentally try to distinguish between $v \cdot a \bmod q$ with a uniformly distributed modulo q and $v \cdot a \bmod q$ with a randomly chosen with ternary coefficients (if v is sufficiently small, we expect that $v \cdot a \bmod q$ has more coefficients $< q/4$ when a is ternary than when a is uniform). This gives us an (experimental) lower bound on the size of Ω' we can take in order to distinguish PV-Knap instances from random elements, with a not too small advantage.

Once this lower bound t_0 on the size of Ω' is computed, we compute the largest subset of $\mathrm{Aut}_{\mathbb{Q}}(K)$ stabilizing a subset Ω' of Ω of size at least t_0. We do that for different random choices of Ω, and compute the probability (over the choice of Ω) that there exists a subset Ω' of Ω of size at least t_0 and such that $I_{\Omega'}$ is stabilized by a subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ of order $1, 2, 3, \ldots$.

We observe that, most of the time, there does not exist a subset Ω' with sufficiently large size and stabilized by a non-trivial subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$. In these cases, we cannot use our attack to lower the dimension of the lattices. However, in some cases, we were able to find a sufficiently large set Ω' stabilized by a subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ of order 2. In this case, one can reduce the dimension of the lattice in which to solve SVP by a factor 2. In Table 3 below, we show the empirical probability that Ω contains a large enough subset Ω' stabilized by a subgroup of order k of $\mathrm{Aut}_{\mathbb{Q}}(K)$, for $k = 1$ and $k = 2$ (we never observed a larger k experimentally).

We can see that for all parameter sets, there is a non-negligible probability to sample a random Ω that contains a good subset Ω' allowing to reduce the dimension of the lattice problem by a factor 2. Hence, by sampling many random PV-Knap instance, one can hope to obtain an easier than expected instance in a few trials (between 3 and 2500 trials depending on the parameter sets).

The fact that the probability to find a good subset Ω' increases when the dimension increases in the HPSSW parameter sets might seem surprising at first.

	LZA 1	LZA 2	HPSSW 1	HPSSW 2	HPSSW 3	HPSSW 4
Subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ of size 1	0.86	0.9996	0.98	0.94	0.55	0.65
Subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ of size 2	0.14	0.0004	0.02	0.06	0.45	0.35

Table 3. Probability to find a good subset Ω' in a random Ω

We believe that the explanation comes from the choice of t, which is $< d/2$ for HPSSW 1 and HPSSW 2 and is $> d/2$ for HPSSW 3 and HPSSW 4. The larger t, the easier it is to find a not too small subset that has some nice stabilizing properties. We also note that the probability to find a good set Ω seems to vary significantly with the choice of t, and with our estimate of t_0 (the minimal size of Ω' that we can allow). Running the same computation with a different random seed might produce significantly different probabilities. For this reason, the numbers in Table 3 are to be taken as order of magnitudes, and not precise estimates of the success probability.

We conclude that, even when the set Ω is chosen uniformly at random, there is some non-negligible probability that one can reduce the dimension of the lattice in which to solve SVP by a factor 2. This might significantly improve the run time of the attack, since the cost of SVP increases exponentially with the dimension of the lattice. Hence, one should be careful when choosing parameter sets for the PV-Knap problem.

Random choices of Ω: full distinguishing attack. Finally, we also ran the full distinguishing attack on the parameter set LZA 1, which was supposed to provide 128 bits of security.

We implemented the strategy described above: we sampled 3000 random PV-Knap instances, and kept the one whose set Ω contained the largest subset Ω' stabilized by a subgroup of $\mathrm{Aut}_{\mathbb{Q}}(K)$ of order 2. We then ran BKZ with block size ≤ 50 in the lattice $qI_{\Omega'}^{-1}$ to obtain a sufficiently short element v. This took time roughly 11 hours on a personal laptop. We then estimated empirically the probability success of our distinguishing attack given this short element v and random BDD targets b.

We concluded that our short vector v allows us to distinguish uniform targets b from PV-Knap ones with advantage at least 0.0005. We computed this advantage using 10^6 samples, to make sure that the advantage gap we computed was significant (Hoeffding's bound guarantees that our advantage is at least 0.0005, expect with probability at most 0.01). Overall, taking into account the fact that our attack chooses the best Ω among 3000 choices, this means that our distinguishing advantage is at least $3000^{-1} \cdot 0.0005 \geq 2^{-23}$, for a run time of less than 12 hours on a personal laptop with a 1.8 GHz processor, hence amounting to $\leq 2^{47}$ bit-operations. It was suggested in [MW18] to define the bit-security of a distinguishing problem as $\log_2(T/\varepsilon^2)$, where ε is the distinguishing advantage of the attacker and T is its time (or, in our case, its number

of bit-operations). Our attack hence shows that the parameter set LZA 1 enjoys at most $47 + 2 \cdot 23 = 93$ bits of security, which is significantly smaller than the expected 128 bits of security. We note however that this does not fully invalidate the claim made in [LZA18], since the 128 bit-security is claimed against search attackers, and not distinguishing attackers.

We could also increase the advantage of our attack a bit more, by spending more time on the lattice reduction phase, in order to obtain an even shorter element v. We did so with BKZ with block-size 55 and obtained an attack with advantage roughly $3000^{-1} \cdot 0.0044 \geq 2^{-20}$, for a total time ≤ 20 hours. This reduces the security of the parameter set LZA 1 even further to ≤ 87 bits of security (against distinguishing attackers).

This attack shows that the security estimate provided in [LZA18] for the first set of parameters is overestimated for distinguishing attackers, even when the set Ω is chosen uniformly at random. We expect that the other estimates provided in [LZA18] and [HPS+14] might also be overestimated, even though it might not be possible to actually run the full attack in a few hours on a laptop.

Acknowledgments. We are grateful to Amin Sakzad, Damien Stehlé and Ron Steinfeld for helpful discussions. We also thank Pascal Autissier for spotting a mistake in a previous version of this article. This research was partly funded by the ANR CHARM project (ANR-21-CE94-0003) and further supported by the Danish Independent Research Council under project number 0165-00107B (C3PO).

References

AD17. Martin R. Albrecht and Amit Deo. Large modulus ring-LWE $\geq$ module-LWE. In *ASIACRYPT 2017*, volume 10624 of *Lecture Notes in Computer Science*, pages 267–296. Springer, 2017.

Ajt96. Miklós Ajtai. Generating hard instances of lattice problems (extended abstract). In *STOC*, pages 99–108. ACM, 1996.

Bab86. László Babai. On lovász'lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.

BLNR21. Olivier Bernard, Andrea Lesavourey, Tuong-Huy Nguyen, and Adeline Roux-Langlois. Log-s-unit lattices using explicit stickelberger generators to solve approx ideal-svp. *Cryptology ePrint Archive*, 2021.

Bou21. Katharina Boudgoust. *Theoretical Hardness of Algebraically Structured Learning With Errors*. PhD thesis, Universite Rennes 1, 2021. https://tel.archives-ouvertes.fr/tel-03534254/document.

BR20. Olivier Bernard and Adeline Roux-Langlois. Twisted-phs: Using the product formula to solve approx-svp in ideal lattices. In *Asiacrypt*, pages 349–380. Springer, 2020.

BSS22. Katharina Boudgoust, Amin Sakzad, and Ron Steinfeld. Vandermonde meets regev: Public key encryption schemes based on partial vandermonde problems. Cryptology ePrint Archive, Report 2022/679, 2022.

CDPR16. Ronald Cramer, Léo Ducas, Chris Peikert, and Oded Regev. Recovering short generators of principal ideals in cyclotomic rings. In *Eurocrypt*, pages 559–585. Springer, 2016.

CDW17. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Short Stickelberger class relations and application to ideal-SVP. In *Eurocrypt*, pages 324–348. Springer, 2017.

CDW21. Ronald Cramer, Léo Ducas, and Benjamin Wesolowski. Mildly short vectors in cyclotomic ideal lattices in quantum polynomial time. *Journal of the ACM (JACM)*, 68(2):1–26, 2021.

Coh13. Henri Cohen. *A course in computational algebraic number theory*, volume 138. Springer Science & Business Media, 2013.

Con. Keith Conrad. The different ideal. https://kconrad.math.uconn.edu/blurbs/gradnumthy/different.pdf, last accessed on 16.02.2022.

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

DHSS20. Yarkin Doröz, Jeffrey Hoffstein, Joseph H Silverman, and Berk Sunar. Mmsat: A scheme for multimessage multiuser signature aggregation. Cryptology ePrint Archive, Report 2020/520, 2020.

Gen09. Craig Gentry. *A fully homomorphic encryption scheme*. PhD thesis, Stanford University, 2009. http://crypto.stanford.edu/craig.

HKJL+00. Jeffrey Hoffstein, Burton S Kaliski Jr, Daniel Bennett Lieman, Matthew John Barton Robshaw, and Yiqun Lisa Yin. Secure user identification based on constrained polynomials, June 13 2000. US Patent 6,076,163. Filed October 20, 1997.

HPS98. Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In *ANTS*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.

HPS11. Guillaume Hanrot, Xavier Pujol, and Damien Stehlé. Analyzing blockwise lattice algorithms using dynamical systems. In *Crypto*, pages 447–464. Springer, 2011.

HPS+14. Jeff Hoffstein, Jill Pipher, John M Schanck, Joseph H Silverman, and William Whyte. Practical signatures from the partial fourier recovery problem. In *International Conference on Applied Cryptography and Network Security*, pages 476–493. Springer, 2014.

HS15. Jeffrey Hoffstein and Joseph H Silverman. Pass-encrypt: a public key cryptosystem based on partial evaluation of polynomials. *Designs, Codes and Cryptography*, 77(2):541–552, 2015.

Lan02. Serge Lang. *Algebra*. Springer, 2002.

LM06. Vadim Lyubashevsky and Daniele Micciancio. Generalized compact knapsacks are collision resistant. In *ICALP (2)*, volume 4052 of *Lecture Notes in Computer Science*, pages 144–155. Springer, 2006.

LPR10. Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In *EUROCRYPT*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23. Springer, 2010.

LS15. Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 75(3):565–599, 2015.

LZA18. Xingye Lu, Zhenfei Zhang, and Man Ho Au. Practical signatures from the partial fourier recovery problem revisited: A provably-secure and gaussian-distributed construction. In *ACISP*, volume 10946 of *Lecture Notes in Computer Science*, pages 813–820. Springer, 2018.

Mar77. Daniel A Marcus. *Number fields*, volume 2. Springer, 1977.

Mic02. Daniele Micciancio. Generalized compact knapsacks, cyclic lattices, and efficient one-way functions from worst-case complexity assumptions. In *FOCS*, pages 356–365. IEEE Computer Society, 2002.

MW18. Daniele Micciancio and Michael Walter. On the bit security of cryptographic primitives. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 3–28. Springer, 2018.

PHS19. Alice Pellet-Mary, Guillaume Hanrot, and Damien Stehlé. Approx-svp in ideal lattices with pre-processing. In *Eurocrypt*, pages 685–716. Springer, 2019.

PML21. Christian Porter, Andrew Mendelsohn, and Cong Ling. Subfield algorithms for ideal-and module-svp based on the decomposition group. *arXiv preprint arXiv:2105.03219*, 2021.

PR06. Chris Peikert and Alon Rosen. Efficient collision-resistant hashing from worst-case assumptions on cyclic lattices. In *TCC*, volume 3876 of *Lecture Notes in Computer Science*, pages 145–166. Springer, 2006.

PXWC21. Yanbin Pan, Jun Xu, Nick Wadleigh, and Qi Cheng. On the ideal shortest vector problem over random rational primes. In *Eurocrypt*, pages 559–583. Springer, 2021.

Reg05. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93. ACM, 2005.

SE94. Claus-Peter Schnorr and Martin Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66:181–199, 1994.

SSTX09. Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In *ASIACRYPT*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635. Springer, 2009.

The20. The Sage Developers. *SageMath, the Sage Mathematics Software System (Version 9.0)*, 2020. `https://www.sagemath.org`.

Was82. Lawrence C Washington. *Introduction to cyclotomic fields*, volume 83. Springer-Verlag, Berlin, 1982.

Appendix A Additional Preliminaries

A.1 Easy algorithmic problems on lattices

Lemma A.1. *Given a rational basis $\mathbf{B}_L$ of a lattice L which lives in a vector space $E \subseteq \mathbb{Q}^m$, and a basis $\mathbf{B}_F$ of a sub-space F of E, one can compute a basis of the lattice $L \cap F$ in time polynomial in $m, \text{size}(\mathbf{B}_L)$ and $\text{size}(\mathbf{B}_F)$.*

Proof. First, compute a basis $\mathbf{B}_{F^\perp}$ of the subspace $F^\perp$ orthogonal to F. By Gauss pivoting, this can be done in time polynomial in m and $\text{size}(\mathbf{B}_F)$. This basis is such that a vector $\mathbf{x} \in E$ is in F if and only if $\mathbf{B}_{F^\perp} \cdot \mathbf{x} = \mathbf{0}$. Recall that $\mathbf{x}$ is in the lattice L if and only if there exist some integer vector $\mathbf{y}$ such that $\mathbf{x} = \mathbf{B}_L \cdot \mathbf{y}$. Hence, to compute the vectors $\mathbf{x} \in L \cap F$, it is equivalent to compute the integer vectors $\mathbf{y}$ such that $\mathbf{B}_{F^\perp} \cdot \mathbf{B}_L \cdot \mathbf{y} = \mathbf{0}$. This means we want to find all integer solutions to a rational linear system. Since the system is homogeneous, we can multiply it by its common denominator to obtain an integral linear system. Then, computing the integral kernel of an integral matrix can be performed in polynomial time by computing its Hermite normal form, as described for instance in [Coh13, Algorithm 2.4.10]. □

A.2 Algorithms for Id-HSVP

We discuss here two algorithms that exploit the algebraic structure of ideal lattices in order to find short vectors in them (for the canonical embedding Σ_K).

In the special case of ideal lattices of a *cyclotomic* number fields K of degree d, one can use the algorithm of [CDW17, CDW21]. The algorithm was described only for prime power cyclotomic fields in [CDW17], and then generalized to any cyclotomic field in [CDW21]. This is a heuristic and quantum algorithm, that takes as input a basis $\mathbf{B}_K$ of O_K and a basis $\mathbf{B}_I$ of an ideal I. It solves γ-Id-HSVP$_K$ in I in quantum polynomial time for some approximation factor $\gamma = \exp(\sqrt{d} \cdot \log(d)^{O(1)})$.

In the case of ideal lattices in any number field K, one could also use the algorithm of [PHS19]. This algorithm is heuristic and consists in an offline phase, that only depends on the number field K, followed by an online phase that depends on both K and the ideal I. When the discriminant of the number field is not too large (for instance for cyclotomic fields), the online phase performs heuristically better than the BKZ algorithm. The offline phase however requires a time $2^{O(d)}$ which is asymptotically the same as the time required by the BKZ algorithm for the smallest approximation factor γ. Hence, when considering the run time of both the offline and the online phase, this algorithm is never better than the BKZ algorithm. Variants of this algorithm are known [BR20], which have better practical run times, but the same asymptotic complexity (no better than BKZ).

Appendix B Partial Vandermonde Knapsack Version 2

Let $K = \mathbb{Q}[X]/(f(X))$ be a number field of degree d with ring of integers O_K of the form $\mathbb{Z}[X]/f(X))$. Further, let q be a prime such that $f(X)$ completely splits into d linear factors modulo q, i.e., $f(X) = \prod_{j=1}^{d}(X - \omega_j)$, where $\omega_j \in \mathbb{Z}_q$.

For the case of K being the m-th cyclotomic field, this is the case for every prime q such that $q = 1 \bmod m$.

By the Kummer-Dedekind theorem we simultaneously obtain the factorization of the ideal generated by q into prime ideals over O_K. It yields $qO_K = \prod_{j=1}^{d} \mathfrak{p}_j$, where $\mathfrak{p}_j = qO_K + (X - \omega_j)O_K$.

The (discrete) Vandermonde transform of any field element $g(X) \in K$ is defined as the evaluation of $g(X)$ at the roots ω_j given by the factorization of $f(X)$, i.e., $\mathbf{V}(g(X)) = g(\omega_j)_{1 \le j \le d} \in \mathbb{Z}_q^d$. For $t \le d$, we set $\mathcal{P}_t = \{\Omega \subseteq \{1, \ldots, d\} \colon |\Omega| = t\}$ and $\Omega^c = \{1, \ldots, d\} \setminus \Omega$. We define the partial Vandermonde transform as $\mathbf{V}_\Omega(g(X)) = g(\omega_j)_{j \in \Omega} \in \mathbb{Z}_q^t$ and the complement partial Vandermonde transform as $\mathbf{V}_{\Omega^c}(g(X)) = g(\omega_j)_{j \in \Omega^c} \in \mathbb{Z}_q^{d-t}$.

Definition B.1 (PV-Knap Version 2). *Let K, O_K, d, q and t be as above. Fix $\Omega \in \mathcal{P}_t$ and let ψ be a distribution over O_K such that $\max_{g \leftarrow \psi} ||\Sigma_K(g)|| \le B$ for some positive real B fulfilling $2B < \sqrt{d} \cdot q^{t/d}$. Sample $g(X) \leftarrow \psi$. Given $h := \mathbf{V}_\Omega(g(X)) \in \mathbb{Z}_q^t$, the partial Vandermonde knapsack problem* PV-Knap$_{\Omega,\psi}$ *asks to find $g(X)$.*

Note that the partial Vandermonde transform of $g(X)$ with respect to the set Ω is zero modulo q if and only if $g(X)$ lies in the ideal $I_\Omega := \prod_{j \in \Omega} \mathfrak{p}_j$, as defined in Section 2.4. More precisely,

$$I_\Omega = \{g(X) \in O_K \colon \mathbf{V}_\Omega(g(X)) = 0 \bmod q\}.$$

Thus, for a given ring element $g(X)$ of small norm, the partial Vandermonde transform $\mathbf{V}_\Omega(g(X)) \bmod q$ is a way to specify the coset $g(X) + I_\Omega$, leading to the equivalence of Definition 2.14 and Definition B.1.

Partial Vandermonde Learning With Errors. Boudgoust et al. [BSS22, Bou21] introduce a dual problem to PV-Knap that they call partial Vandermonde learning with errors (PV-LWE). The duality connection is similar as the one between the standard Knapsack and LWE problems. They prove that both problems are equivalent to each other for power-of-two cyclotomic fields [Bou21, Sec. 5.3]. Hence, our attacks against PV-Knap also apply to PV-Knap in this case.

Partial Fourier Recovery Problem. As mentioned in the introduction and in the preliminaries, in some works the problem is defined via the Fourier transform instead of the Vandermonde transform and called the *Partial Fourier Recovery* problem [HPS+14, HS15, DHSS20]. For completeness, we provide its concrete formulation here as well.

Let N and q be primes, such that there exists a primitive N-th root of unity modulo q, denoted by ω. Further, consider the ring of polynomials $R :=$

$\mathbb{Z}[X]/(X^N-1)$. For any polynomial $g(X) \in R$, we define the Fourier transform as its evaluation at all the powers of ω (i.e., evaluation at all N-th roots of unity). More precisely, $\mathbf{F}(g(X)) = g(\omega^j)_{j\in 0,\ldots,N-1}$. Similarly as above, we can define the partial Fourier transform by restricting it to a subset $\Omega \subseteq \{0,\ldots,N-1\}$ of size $t \leq N$ which we denote by $\mathbf{F}_\Omega(g(X)) = g(\omega^j)_{j\in\Omega}$.

Definition B.2 (Partial Fourier Recovery). *Let R, N, q and t as above. Fix Ω and let ψ be a distribution over R providing polynomials with coefficients of small norm. Sample $g(X) \leftarrow \psi$. Given $\mathbf{F}_\Omega(g(X)) \in \mathbb{Z}_q^t$, the Partial Fourier Recovery problem asks to find $g(X)$.*

A current choice of ψ is the uniform distribution over ternary polynomials, i.e., polynomials with coefficients in $\{-1, 0, 1\}$.