Algorithms for the Densest Sub-Lattice Problem

Daniel Dadush* Daniele Micciancio†

December 24, 2012

Abstract

We give algorithms for computing the densest k-dimensional sublattice of an arbitrary lattice, and related problems. This is an important problem in the algorithmic geometry of numbers that includes as special cases Rankin's problem (which corresponds to the densest sublattice problem with respect to the Euclidean norm, and has applications to the design of lattice reduction algorithms), and the shortest vector problem for arbitrary norms (which corresponds to setting $k = 1$) and its dual ($k = n - 1$). Our algorithm works for any norm and has running time $k^{O(k \cdot n)}$ and uses $2^n \operatorname{poly}(n)$ space. In particular, the algorithm runs in single exponential time $2^{O(n)}$ for any constant $k = O(1)$.

*Georgia Tech. Atlanta, GA, USA. Email: dndadush@gmail.com

†University of California, San Diego. 9500 Gilman Dr., Mail Code 0404, La Jolla, CA 92093, USA. Email: daniele@cs.ucsd.edu. This work was supported in part by NSF grant CNS-1117936. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

1 Introduction

A lattice is a discrete subgroup of $\mathbb{R}^n$. Computational problems on lattices play an important role in many areas of computer science, mathematics and engineering, including cryptography, cryptanalysis, combinatorial optimization, communication theory and algebraic number theory. The most famous problem on lattices is the shortest vector problem (SVP), which asks to find the shortest nonzero vector in an input lattice. Several generalizations and variants of SVP naturally arise both in the theoretical study of lattices and in their applications. To start with, lattice problems can be defined with respect to any norm, but, for simplicity, in this introductory discussion we focus on the standard Euclidean norm $\|\mathbf{x}\| = \sqrt{\sum_i |x_i|^2}$. The most common generalization of SVP encountered in computer science is the k-shortest independent vectors problem (k-SIVP): given a lattice, find k linearly independent lattice vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ such that the maximum length $\max_i \|\mathbf{v}_i\|$ is minimized. SVP is recovered as a special case of k-SIVP by setting $k = 1$. At the other end of the spectrum, for $k = n$ equal to the dimension of the lattice, we get the classic SIVP problem arising in the construction of cryptographic functions with worst-case/average-case connection [Ajt96, MR04, Reg05]. Finally, for arbitrary k, k-SIVP is the computational problem naturally associated to the lattice successive minima $\lambda_1, \ldots, \lambda_n$, some of the most fundamental mathematical parameters describing the geometry of a lattice.[1] In this paper we consider a different, still very natural and equally motivated generalization of SVP: the k-dimensional Densest Sublattice Problem (k-DSP), i.e., given a lattice, find k linearly independent lattice vectors $\mathbf{v}_1, \ldots, \mathbf{v}_k$ that generate a sublattice achieving the smallest possible determinant $\det(\mathbf{v}_1, \ldots, \mathbf{v}_k)$.[2] The k-DSP was first explicitly formulated as a computational problem in [GHGKN06] under the name[3] of "smallest volume problem", but it was already implicit in several previous works both in mathematics and theoretical computer science. (See discussion below.) As for SIVP, when $k = 1$, the 1-DSP is equivalent to SVP, and therefore also to 1-SIVP. However, for other values of k, k-DSP and k-SIVP are different problems: they both require finding an optimal set of k linearly independent lattice vectors, but with respect to different quality measures, maximum length for SIVP and lattice determinant for DSP. The difference between the two problems is well illustrated by the fact that for $k = n$ the k-DSP is trivial, while k-SIVP is NP-hard [BS99]. For arbitrary values of k, both problems appear hard, but there is no clear relation between them.

Motivating applications The k-DSP (and its generalization to arbitrary norms, see Section 4) arises naturally in many important problems in mathematics, combinatorial optimization, and the computational study of lattices.

For example, in mathematics, the k-DSP problem with respect to the Euclidean norm is tightly related to the study of Rankin constants $\gamma_{n,k}$. These are fundamental parameters in the geometry of numbers introduced by Rankin [Ran55] as a generalization of Hermite's constants $\gamma_n = \gamma_{n,1}$, and defined as

$$\gamma_{n,k} = \sup_{\Lambda} \left(\min \frac{\det([\mathbf{b}_1, \ldots, \mathbf{b}_k])}{\det(\Lambda)^{k/n}} \right)^2$$

where the minimum is computed over all linearly independent lattice vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k \in \Lambda$, and the supremum is over all n-dimensional lattices. It is immediate to see that the quantity in the expression of the

[1] Related to k-SIVP is the *successive minima problem*, which asks to find n linearly independent lattice vectors $\mathbf{v}_1, \ldots, \mathbf{v}_n$ that simultaneously achieve all successive minima $\|\mathbf{v}_i\| = \lambda_i$ for $i = 1, \ldots, n$.

[2] We recall that the derminant is the volume of the parallelepiped spanned by $\mathbf{v}_1, \ldots, \mathbf{v}_k$, and equals the inverse density of the lattice points within their k-dimensional linear span.

[3] In this paper, we introduce a different name because "densest sublattice problem" more properly describes the generalization of this problem to arbitrary norm.

supremum is precisely (up to scaling and squaring) the objective function of the k-DSP problem. Determining the value of Rankin constants is a classical hard problem in mathematics, and to date, the value of $\gamma_{n,k}$ is known only for a handful of cases. (See Section 5.) An efficient algorithm to solve k-DSP immediately gives a powerful computational tool to determine lower bounds on $\gamma_{n,k}$.

In computer science and cryptography, Rankin constants and the associated k-DSP problem in the Euclidean norm has been suggested as a building block for novel basis reduction algorithms [GHGKN06], as well as an analytical method to understand the limitation of more classical block reduction methods [Sch87]. However, little progress was made along this direction since [GHGKN06], largely because of the lack of efficient algorithms to solve k-DSP. More specifically, [GHGKN06] gave an approximation algorithm for k-DSP (called the *transference reduction* algorithm), which results in a suboptimal basis block-reduction method, and is provably inferior to other techniques based on SVP and Hermite's constant γ_n [GN08]. We remark that part of the difficulty of evaluating the potential of the Rankin reduction algorithm of [GHGKN06] is due to the fact that the value of Rankin constants $\gamma_{n,k}$ is not known except for a handful of values of n, k. An algorithm to solve k-DSP would serve both as a tool to study the value of Rankin constants $\gamma_{n,k}$, and also as a method to instantiate the Rankin reduction framework of [GHGKN06] and better assess its potential.

The generalization of k-DSP to arbitrary norms, while nontrivial even to define (see Section 3), arises naturally in the context of applications, and in particular, may be useful in the development of faster algorithms for integer programming. The asymptotically fastest known algorithm to solve integer programming is the one of Kannan [Kan87], running in time $O(n^{O(n)})$, with some recent work [HK10, DPV11] improving the constant in the exponent. Kannan's algorithm works by reducing an integer programming instance in n variables, to $n^{O(r)}$ instances in $(n-r)$ variables, for some unspecified $r = 1, \ldots, n$. Recursing, this yields an algorithm with running time $n^{O(n)}$. The problem of finding an optimal decomposition (of the kind used by [Kan87]) into the smallest possible number of $(n-r)$-dimensional subproblems can be formulated as an r-DSP instance in an appropriate norm. Based on the best known upper and lower bounds in asymptotic convex geometry, this could lead to integer programming algorithms with running time as low as $(\log n)^{O(n)}$, much smaller than the current $n^{O(n)}$.

Similar ideas may also lead to better polynomial space algorithms for the closest vector problem with preprocessing. This and other possible potential applications are described in more detail in Section 5.

State of the art It is easy to see that k-DSP (for any fixed k) is at least as hard as SVP. (For example, one can map any SVP instance in dimension n to a corresponding k-DSP instance in dimension $n+k-1$ simply by adding $k-1$ very short vectors orthogonal to the original lattice.) In particular, just like SVP [Ajt98, Mic98, Kho03, HR07, Mic12], k-DSP is NP-hard (at least under randomized reductions) for any k, and it cannot be solved in subexponential time under standard complexity assumptions. A simple lattice duality argument (see Section 3) also shows that k-DSP is equivalent to $(n-k)$-DSP, where n is the dimension of the lattice. But beside that, not much is known about the computational complexity of k-DSP. In particular, while the algorithmic study of SVP, CVP and SIVP has received much attention, leading to practical heuristics [SE94, SH95, NV08, MV10a, GNR10, WLTB11] and asymptotically efficient algorithms with single exponential running time $2^{O(n)}$ [AKS01, BN07, AJ08, MV10a, PS09, MV10b], the only known algorithm for DSP in the literature is the one of [GHGKN06] for the special case of 4-dimensional lattices. We remark that [GHGKN06] also mentions that the general problem can be solved by a "gigantic" exhaustive search over all LLL reduced bases of the input lattice, resulting in $2^{O(n^3)}$ running time. (See discussion at the beginning of Section 3.1 for details.) The algorithms presented in this paper are also based on a form of exhaustive search, and this is unavoidable because of NP-hardness, but the search is over a much smaller space.

Our contribution The main result of this paper is the first nontrivial algorithm for the solution of the DSP problem in arbitrary dimension. Specifically, our algorithm runs in $k^{O(kn)}$ time and $2^n \operatorname{poly}(n)$ space, which for constant $k = O(1)$, is a single exponential function $2^{O(n)}$ of the dimension, similar to the complexity of the best known algorithm for SVP.

We remark that, as already noted in [GHGKN06], Hermite-Korkine-Zolotarev (HKZ) reduction does not solve k-DSP, and there are lattices such that no solution to k-DSP contains a shortest lattice vector. (See Section 3.1.) So, it is not clear how known algorithms or techniques for SVP, SIVP, etc. [SE94, SH95, NV08, MV10a, GNR10, WLTB11, AKS01, BN07, AJ08, MV10a, PS09, MV10b] can be used to approach the k-DSP. The main technical contribution of this paper is the realization (and simple proof) that either a k-DSP solution contains the shortest lattice vectors (in which case one can efficiently find them, and then recursively solve a lower dimensional problem), or one can efficiently generate a relatively short list of at most $O(k)^n$ lattice vectors that contains as a subset the solution to k-DSP. In the latter case, the problem can be solved in time $k^{O(kn)}$ simply by enumerating all possible subsets of the list of size k. This last enumeration step is the current bottleneck of our algorithm, as all other operations can be performed in at most $O(k)^n$ time. We describe several heuristics that can be used to substantially improve the running time of the enumeration step, making our algorithm reasonably practical, but without, at present, leading to substantial provable improvements in the running time. It is an interesting open problem if the complexity of the enumeration step (and therefore, the entire k-DSP algorithm) can be brought down from $k^{O(kn)}$ to $k^{O(n)}$, or even $2^{O(n)}$ for arbitrary k.

Our algorithm also generalizes to versions of the k-DSP in arbitrary norms (see Section 3), which has applications to Integer Programming (see Section 5 for details). At the technical level, the generalization is highly nontrivial, using much machinery from high dimensional convex geometry, but the resulting running time is similar to the Euclidean case: k-DSP with respect to arbitrary norms can be solved in time $k^{O(kn)}$ where n is the dimension of the lattice.

Organization The rest of the paper is organized as follows. In Section 2 we give some general background on lattices. In Section 3 we present a solution to the smallest volume problem, i.e., the densest sublattice problem with respect to the Euclidean norm. In Section 4 we treat the problem for arbitrary norm. Section 5 concludes the paper with a more detailed description of potential applications and directions for further research.

2 Preliminaries

Convex Bodies: For subsets $A, B \subseteq \mathbb{R}^n$, and scalars $s, t \in \mathbb{R}$ we define the *Minkowski Sum* as $sA + tB = \{s\mathbf{a} + t\mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B\}$. $K \subseteq \mathbb{R}^n$ is a *convex body* if it is full-dimensional, compact and convex. K is *symmetric* if $K = -K$. Let $\mathcal{B}_2^n(\mathbf{c}, r) = \{\mathbf{x} \in \mathbb{R}^n \colon \|\mathbf{x} - \mathbf{c}\|_2 \leq r\}$ be the n-dimensional Euclidean ball of radius r centered at $\mathbf{c}$. When $\mathbf{c}$ and/or r are omitted, they are assumed to be $\mathbf{c} = \mathbf{0}$ and $r = 1$. The superscript n and subscript 2 are also often omitted when clear from the context. K is said to be (r, R)-*centered* if $r\mathcal{B}_2^n \subseteq K \subseteq R\mathcal{B}_2^n$. When interacting algorithmically with K, K will be presented by a (weak-)membership oracle, i.e. given a point $\mathbf{x} \in \mathbb{R}^n$, the oracle returns 1 if $\mathbf{x} \in K$ and 0 otherwise. The complexity of our algorithms will be computed in terms of the number of oracle queries and arithmetic operations. In what follows, K will denote a symmetric (r, R)-centered convex body - presented by a membership oracle in our algorithms - unless otherwise specified.

We define the *norm induced by* K (or the gauge function of K) as $\|\mathbf{x}\|_K = \inf\{s \geq 0 : \mathbf{x} \in sK\}$ for $\mathbf{x} \in \mathbb{R}^n$. It is straightforward to verify that $\|\cdot\|_K$ satisfies all the norm axioms.

Lattices: $\Lambda \subseteq \mathbb{R}^n$ is a k-dimensional lattice if Λ can be represented as $\oplus_{i=1}^{k} \mathbb{Z}\mathbf{b}_i$ for linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k \in \mathbb{R}^n$. A lattice is *full-dimensional* if $k = n$. We denote the determinant of Λ as $\det(\Lambda) = \sqrt{\det(\mathrm{Gram}(\mathbf{b}_1, \ldots, \mathbf{b}_k))}$ where $\mathrm{Gram}(\mathbf{b}_1, \ldots, \mathbf{b}_k)_{ij} = \langle \mathbf{b}_i, \mathbf{b}_j \rangle$. The determinant is invariant under a change of basis of Λ (since all bases are related by unimodular transformations) and hence is a lattice invariant.

A lattice Λ is a *K-packing* if the (interiors of the) bodies $K + \mathbf{x}$ (with $\mathbf{x} \in \Lambda$) are mutually disjoint. A natural quantity associated to any packing $P = \{K + \mathbf{x} \colon \mathbf{x} \in \Lambda\}$ is the *density* $\delta(K, \Lambda) = \mathrm{vol}(K \cap \mathrm{span}(\Lambda))/\det(\Lambda)$, i.e., the fraction of space occupied by P. Finding a (full-dimensional) lattice packing Λ which achieves the highest possible density for a convex body K is a classical problem in mathematics. A lattice Λ is a *K-covering* if the (closure of the) bodies $K + \mathbf{x}$ (with $\mathbf{x} \in \Lambda$) cover the entire linear space spanned by Λ, i.e. $\mathrm{span}(\Lambda) \subseteq \Lambda + K$. In this case, the quantity $\delta(K, \Lambda) = \mathrm{vol}(K \cap \mathrm{span}(\Lambda))/\det(\Lambda)$ is called the *thickness* of the lattice covering, and equals the average number of copies of the body $K + \mathbf{x}$ (with $\mathbf{x} \in \Lambda$) that cover a random point in $\mathrm{span}(\Lambda)$. Finding a (full-dimensional) lattice covering Λ that is as thin as possible (i.e., achieves the lowest possible thickness) for a convex body K is also a classical problem in mathematics. Notice that if Λ is a K-packing, then for any linear subspace M, $\Lambda \cap M$ is a $(K \cap M)$-packing. Similarly, if Λ is a K-covering, then for any linear subspace W, $\pi_W(\Lambda)$ is a $\pi_W(K)$-covering, where π_W is the orthogonal projection onto W.

We define the i^{th} minimum of Λ with respect to K as $\lambda_i(K, \Lambda) = \inf\{s \geq 0 : \dim(sK \cap \Lambda) \geq i\}$ for $1 \leq i \leq \dim(\Lambda)$. The first minimum λ_1 is often denoted simply λ. When K is the Euclidean ball $\mathcal{B}(1) = \{\mathbf{x} \colon \|\mathbf{x}\|_2 \leq 1\}$, it is often omitted, and we write $\lambda_i(\Lambda)$.

Notice that Λ is a (cK)-packing if and only if $c \leq \lambda(K, \Lambda)/2$. Another important quantity associated to a lattice Λ is the covering radius $\mu(K, \Lambda)$ which is defined as the smallest real $c > 0$ such that Λ is a cK-covering. Again, for the Euclidean norm we simply write $\mu(\Lambda)$.

The dual of a lattice Λ is the set of vectors $\Lambda^* = \{\mathbf{x} \in \mathrm{span}(\Lambda) \colon \forall \mathbf{y} \in \Lambda . \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. The dual of a lattice is a lattice, and has determinant $\det(\Lambda^*) = 1/\det(\Lambda)$. The quantities λ_i and μ (with respect to the Euclidean norm) of a lattice Λ and its dual Λ^* are related by the inequalities $\lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_n \leq 2\mu \leq \sqrt{n}\lambda_n$ (e.g., see [MG02]) and transference bound $1 \leq 2\mu \cdot \lambda_1^* \leq n$ [Ban93a]. Similar (but quantitatively weaker) bounds hold for arbitrary norms.

We will use Minkowski's classical first and second theorems [Min10]:

Theorem 2.1. *Let $K, \Lambda \subseteq \mathbb{R}^n$ be a symmetric convex body and n-dimensional lattice. Then*

$$\lambda_1(K, \Lambda) \leq 2\frac{\det(\Lambda)^{\frac{1}{n}}}{\mathrm{vol}_n(K)^{\frac{1}{n}}}.$$

More generally, we have that

$$\prod_{i=1}^{n} \lambda_i(K, \Lambda)^{\frac{1}{n}} \leq 2\frac{\det(\Lambda)^{\frac{1}{n}}}{\mathrm{vol}_n(K)^{\frac{1}{n}}}.$$

In what follows, Λ will always denote an n dimensional lattice, presented by a basis $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^n$.

Our algorithms will require enumerating lattice points inside an arbitrary scaling of a norm ball. This can be done efficiently with the algorithm of [MV10b] for the case of the Euclidean norm, and can be extended to arbitrary norms via an ellipsoid covering technique [DPV11]. The following is a refinement of the previous algorithms from [MV10b, DV12, Dad12] (see Theorem 5.2.6 of [Dad12]).

Theorem 2.2 (Lattice Point Enumerator)**.** *For a scalar $t \geq 1$ and symmetric convex body $K \subseteq \mathbb{R}^n$, the set $S = \{\mathbf{y} \in \Lambda : \|\mathbf{y}\|_K \leq t\lambda_1(K, \Lambda)\}$ of all lattice points of norm less than $t\lambda_1(K, \Lambda)$ has size at most*

$|S| \leq (1+2t)^n$ *and can be enumerated (one element at a time) in deterministic* $2^{O(n)}t^n$ *time and* $2^n \operatorname{poly}(n)$ *space.*

3 Densest Sublattice Problem

In this section, we define the Densest Sublattice Problem (and its "dual", the Sparsest Projection Problem), and discuss our algorithm for the simpler setting of the Euclidean norm. We remark that the algorithm for arbitrary norms presented in Section 4 is conceptually similar, but mathematically more involved, since it requires analyzing the behavior of sectional volumes of general convex bodies.

Let Λ be an arbitrary n-dimensional lattice, let $K \subset \mathbb{R}^n$ be a symmetric convex body, and fix a $k \leq n$. In the Euclidean setting K is simply $\mathcal{B}_2$. By scaling the body K, we may assume that the lattice Λ is a K-packing, i.e. $\Lambda + K$ is comprised of disjoint copies of K centered around points in Λ. In the Euclidean setting, this corresponds to a classical sphere packing. The *Densest k-dimensional Sublattice Problem* (k-DSP) with respect to K and Λ, is the problem of finding a sublattice $M \subseteq \Lambda$, $\dim(M) = k$ over which the packing $\Lambda + K$ is as *dense* as possible, i.e., the relative density

$$\delta(K \cap \operatorname{span}(M), M) = \frac{\operatorname{vol}_k(K \cap \operatorname{span}(M))}{\det(M)}$$

is maximized, where vol_k is the k-dimensional volume. Notice that in the Euclidean case the volume $\operatorname{vol}_k(\mathcal{B}_2 \cap \operatorname{span}(M))$ is just the volume of a k-dimensional ball, and hence does not depend on M. Therefore, k-DSP in the ℓ_2 norm asks to find a k-dimensional sublattice that maximizes the quantity $1/\det(M)$, or equivalently minimizes the determinant $\det(M)$. Hence k-DSP in the Euclidean norm is the same as the smallest volume (determinant) problem of Gama et al. [GHGKN06], and described in the introduction.

We also define a dual problem, the *Sparsest* (or *Thinnest*) *k-dimensional Projection Problem* (k-SPP). This time, we assume (after a suitable scaling of K) that Λ is a K-covering. The k-SPP is the problem of finding a dual sublattice $W \subseteq \Lambda^*$, where $\dim(W) = k$, such that the projection of the covering $\Lambda + K$ onto (the linear span of) W is as *sparse* (or *thin*) as possible, i.e. the relative density

$$\delta(\pi_W(K), \pi_W(\Lambda)) = \frac{\operatorname{vol}_k(\pi_W(K))}{\det(\pi_W(\Lambda))}$$

is minimized, where π_W denotes the orthogonal projection onto $\operatorname{span}(W)$. As for k-DSP, in the Euclidean setting the volume $\operatorname{vol}_k(\pi_W(\mathcal{B}_2(\mu)))$ does not depend on the choice of W. (It only depends on the dimension $\dim(W) = k$ which is fixed.) Therefore k-SPP in the ℓ_2 norm asks to minimize the quantity $1/\det(\pi_W(\Lambda)) = \det(\operatorname{span}(W) \cap \Lambda^*)$. So, k-SPP is equivalent to solving k-DSP in the dual lattice.

In Section 4 we will see that the k-DSP and k-SPP in general norms are also equivalent, but only in an approximate sense. In the Euclidean case, k-DSP and k-SPP can also be related as follows. Let $M = \{\mathbf{x} \in \Lambda \mid \forall \mathbf{y} \in W. \langle \mathbf{x}, \mathbf{y} \rangle = 0\}$ be the orthogonal complement of W. Clearly, $\dim(M) = n - k$, and the product $\det(M) \cdot \det(\pi_W(\Lambda)) = \det(\Lambda)$ is a constant independent of the choice of W and M. So, W maximizes $\det(\pi_W(\Lambda))$ (and solves k-SPP) if and only if M minimizes $\det(M)$ (and solves $(n-k)$-DSP). Therefore solving k-SPP on a lattice Λ is equivalent to solving $(n-k)$-DSP on the same lattice.

In summary, for the Euclidean case, both k-DSP and k-SPP can be equivalently formulated as the problem of finding a k-dimensional sublattice of minimal determinant, which is the problem we solve in this section.

3.1 Geometry of Smallest Determinant Sub-lattices

Before we describe the main idea behind our solution, we recall the only other known method to solve k-DSP in arbitrary dimension. In [GHGKN06] it is observed that if B is a basis such that its first k vectors solve k-DSP, then applying the LLL reduction algorithm to it does not modify the lattice generated by the first k vectors. Therefore, one may assume without loss of generality, that the basis B is LLL reduced. The number of LLL reduced bases for any n-dimensional lattice is bounded by a constant that only depends on n. So, [GHGKN06] suggests that k-DSP can be solved by enumerating all LLL reduced bases of the lattice, and for each basis, compute the determinant of the first k basis vectors. While [GHGKN06] gives neither a concrete bound on the number of LLL reduced bases, nor an algorithm to enumerate them, the number of LLL reduced bases can be shown to be at most $2^{O(n^3)}$ [Mic11], and the proof of the bound can be easily translated into an algorithm. This yields an algorithm to solve k-DSP for arbitrary k with running time $2^{O(n^3)}$. This is a very large time bound: quoting [GHGKN06], this method consists of a "gigantic exhaustive search". Some optimizations are possible, e.g., one may further assume that the first k vectors of the LLL reduced basis are HKZ reduced. This and other optimizations apply to our algorithm as well, and they are discussed in Section 3.3. Here we observe that, for small k, the running time of the LLL exhaustive search algorithm can be improved to $2^{O(kn^2)}$ by noticing that we only need to enumerate the first k vectors of all LLL reduced bases: it is easy to adapt the argument of [Mic11] to show that these first k vectors can only take $2^{O(kn^2)}$ possible values, and the proof is again algorithmic. But can we go below $2^{O(kn^2)}$, and ideally solve k-DSP in single exponential time $2^{O(n)}$ at least for constant $k = O(1)$?

As we are working with exponential time algorithms, it may be tempting to use stronger notions of reduced bases than LLL. For example, one may consider enumerating all Hermite-Korkine-Zolotarev reduced bases of the input lattice, which results in a much smaller search space than LLL. Unfortunately, as shown in [GHGKN06], HKZ reduction does not solve k-DSP, and in fact there are lattices such that any sublattice that contains the shortest lattice vector (as the first k vectors of any HKZ reduced basis) cannot be a solution to k-DSP. For example, consider the 2-DSP on the 3-dimensional lattice generated by the basis $\mathbf{b}_1 = (1, -1, 0)$, $\mathbf{b}_2 = (0, -1, 1)$, $\mathbf{b}_3 = (a, a, a)$ where $\sqrt{1/2} < a < \sqrt{2/3}$. The vectors $\mathbf{b}_1, \mathbf{b}_2$ generate the "exagonal" lattice, with minimum distance $\lambda_1(\mathbf{b}_1, \mathbf{b}_2) = \sqrt{2}$ and determinant $\det(\mathbf{b}_1, \mathbf{b}_2) = \sqrt{3}$. The vector $\mathbf{b}_3$ is orthogonal to $\mathbf{b}_1, \mathbf{b}_2$, and has length $\|\mathbf{b}_3\| = a\sqrt{3} < \sqrt{2}$. It follows that $\|\mathbf{b}_3\| = \lambda_1(\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3)$ and $\mathbf{b}_3$ is the shortest lattice vector. However, any two dimensional sub-lattice containing $\mathbf{b}_3$ has determinant at least $\|\mathbf{b}_3\| \cdot \lambda(\mathbf{b}_1, \mathbf{b}_2) = a\sqrt{3}\sqrt{2} > \sqrt{3} = \det(\mathbf{b}_1, \mathbf{b}_2)$, and it is not a 2-DSP solution.

The next lemma gives the crucial insight behind our solution to the smallest volume problem. Simply stated, it says that, while the solution to k-DSP may not contain a shortest lattice vector, the shortest vector in any k-DSP solution cannot be much longer than the shortest vector of the whole lattice. In fact, either the k-DSP solution to lattice Λ contains all lattice vectors of length $\lambda_1(\Lambda)$, or one can generate a list of at most $(2k+1)^n$ lattice vectors of length at most $k\lambda_1(\Lambda)$ that contains the k-DSP solution as a subset.

Lemma 3.1. *Let Λ be an n-dimensional lattice and $M \subseteq \Lambda$ be a k-dimensional sub-lattice of minimum determinant. Let $\mathbf{v} \in \Lambda$ be any lattice vector. Then, either $\mathbf{v} \in M$ or $\|\mathbf{v}\| \geq 1/\lambda_1(M^*)$. In particular, using the standard inequality $\lambda_k \leq 2\mu$ and the transference bound $\lambda_1^* \cdot \mu \leq k/2$, we get that either M contains all shortest lattice vectors $\Lambda \cap \mathcal{B}_2^n(\lambda_1)$ or*

$$\lambda_k(M) \leq 2\mu(M) \leq k/\lambda_1(M^*) \leq k\lambda_1(\Lambda).$$

Proof. Let $\mathbf{w}_k \in M^*$ a dual vector of minimal length $\|\mathbf{w}_k\| = \lambda_1(M^*)$, and extend $\mathbf{w}_k$ to a basis $[\mathbf{w}_1, \ldots, \mathbf{w}_{k-1}, \mathbf{w}_k]$ for M^*. Let $[\mathbf{b}_1, \ldots, \mathbf{b}_k]$ be the dual basis. This is a basis for lattice M such that $\langle \mathbf{b}_i, \mathbf{w}_j \rangle$ equals 1 when $i = j$, and 0 when $i \neq j$. The component of $\mathbf{b}_k$ orthogonal to $\mathbf{b}_1, \ldots, \mathbf{b}_{k-1}$ has

length $\|\mathbf{b}_k^*\| = 1/\|\mathbf{w}_k\| = 1/\lambda_1(M^*)$. Now, consider the sublattice $M' = [\mathbf{b}_1, \ldots, \mathbf{b}_{k-1}, \mathbf{v}]$, obtained replacing the last vector $\mathbf{b}_k$ with $\mathbf{v}$. If $\mathbf{v} \notin M$, then the dimension of M' is still k. Moreover, the determinant of M' is

$$\det(M') \leq \frac{\|\mathbf{v}\|}{\|\mathbf{b}_k^*\|} \cdot \det(M) = \|\mathbf{v}\|\lambda_1(M^*) \cdot \det(M).$$

It follows from the minimality of $\det(M)$ that $\|\mathbf{v}\|\lambda_1(M^*) \geq 1$. □

3.2 Algorithm

Let Λ be an n-dimensional lattice, and λ be the length of its shortest nonzero vector. From Section 3.1, we know that a minimum determinant sub-lattice either contains all lattice vectors of length λ, or it contains a set of k linearly independent vectors of length at most $k \cdot \lambda$. This easily yields a simple enumerative algorithm. The pseudocode of the algorithm is given in Algorithm 1.

Algorithm 1 Algorithm MinDet(Λ, k)

Input: Lattice $\Lambda \subseteq \mathbb{R}^n$, a target dimension $k \geq 0$, $k \leq \dim(\Lambda)$.
Output: $W \subseteq \mathrm{span}(\Lambda)$ a linear subspace, $\dim(\Lambda \cap W) = k$, such that $\det(\Lambda \cap W)$ is minimized.

```
if k = 0 then return {0}
W ← span({y ∈ Λ : ‖y‖_2 ≤ λ_1(Λ)})
if dim(W) ≤ k then
    Λ^⊥ ← π_{W^⊥}(Λ) where W^⊥ is the orthogonal complement of W
    W ← MinDet(Λ^⊥, k − dim(W)) + W
for all T ⊆ {y ∈ Λ : ‖y‖_2 ≤ k λ_1(Λ)} of size |T| = k do
    W' ← span(T)
    if dim(W') = k ∧ (dim(W) ≠ k ∨ det(W' ∩ Λ) < det(W ∩ Λ)) then
        W ← W'
return W.
```

Theorem 3.2. *Algorithm MinDet correctly solves k-DSP in $k^{O(kn)}$ time and $2^n \,\mathrm{poly}(n)$ space.*

Proof. The correctness of the algorithm easily follows from Lemma 3.1. Let W_1 be the linear space spanned by all vectors in Λ of minimal length $\lambda_1(\Lambda)$. The algorithm first finds a densest k-dimensional sublattice that contains all lattice vectors in W_1. (If $\dim(W_1) > k$, then no solution exist.) Notice that $\det(M \cap \Lambda) = \det(W_1 \cap \Lambda) \cdot \det(\pi_{W_1^\perp}(M \cap \Lambda))$ and $\dim(\pi_{W_1^\perp}(M)) = \dim(M) - \dim(W_1) = k - \dim(W_1)$. Since $\det(W_1 \cap \Lambda)$ is fixed, the determinant $\det(M \cap \Lambda)$ is minimized when $\det(\pi_{W_1^\perp}(M \cap \Lambda))$ is minimal. But $\pi_{W_1^\perp}(M \cap \Lambda)$ is a $(k - \dim(W_1))$-dimensional sublattice of $\pi_{W_1^\perp}(\Lambda)$ and therefore its linear span $\pi_{W_1^\perp}(M)$ solves the $(k - \dim(W_1))$-DSP on $\pi_{W_1^\perp}(\Lambda)$, and it is found by the recursive call.

Then, the algorithm considers all sets of k vectors of length at most $k\lambda_1(\Lambda)$, and keeps track of the best solution found. By Lemma 3.1, either the solution to the k-DSP contains all vectors in W_1 (and it is found before entering the loop), or it contains k-linearly independent vectors of length at most $k\lambda_1(\Lambda)$, and it is found during one of the iterations.

For the complexity analysis, we first note that each invocation of Algorithm 1 makes at most one recursive invocation to itself. Moreover, the recursive call is for a strictly smaller value of k. So, the total number of recursive calls is bounded by k, and it is enough to bound the time consumed by a single invocation in the

main body of the algorithm (without counting the cost of the recursive call). Next, observe that each of the basic lattice operations, such as computing a basis for a projected lattice, and computing lattice determinants, can be done in polynomial time.

The algorithm requires to enumerate all lattice vectors S of length bounded by $k\lambda_1(\Lambda)$. This is done using the lattice point enumerator of Theorem 2.2, which requires $2^{O(n)}$ time and $2^n \operatorname{poly}(n)$ space. Computing the linear span of all vectors of length $\lambda_1(\Lambda)$ is easily accomplished using linear algebra. Next to enumerate over all subsets of size k in S during the for loop, we use nested calls to the lattice point enumerator. Since $|S| \leq (1+2k)^n$, this requires at most $|S|^{k-1} = (1+2k)^{n(k-1)}$ calls to the enumerator, where each enumeration over S takes $2^{O(n)}k^n$ time and uses $2^n \operatorname{poly}(n)$ space. Since the other operations execute in $\operatorname{poly}(n)$ time, the for loop requires at most $2^{O(n)}k^{nk} = k^{O(nk)}$ time and $2^n \operatorname{poly}(n)$ space to execute. So, the total running time is $k^{O(nk)}$ and the space usage is $2^n \operatorname{poly}(n)$ as claimed. □

We remark that Algorithm 1 admits several optimizations that can be used to reduce the search space of the enumeration stage, and substantially improve the running time, leading to a reasonably practical algorithms for small values of k and moderate values of n. Some possible optimizations are described in Section 3.3.

3.3 Optimizations

We have proved that the running time of Algorithm 1 for k-DSP is at most $k^{O(kn)}$, with the running time dominated by the iteration over all subsets $T \subseteq S$ of size k. We remark that this exhaustive search over the sets T can be pruned is several ways, possibly leading to a reasonably efficient algorithm for small values of k and n. In this section we briefly examine some simple ways to prune the search space.

To start with, notice that one can assume without loss of generality that the DSP solution M is a Hermite-Korkine-Zolotarev (HKZ) basis, because one can always apply HKZ reduction to M without changing the sublattice it spans. (Notice this is different from saying we can assume M is part of an HKZ basis for the whole lattice Λ, which as already remarked, is not generally true.) We recall that a basis $\mathbf{b}_1, \ldots, \mathbf{b}_k$ is HKZ reduced if $\mathbf{b}_1$ is a shortest lattice vector, and the projected basis $\pi_1(\mathbf{b}_2), \ldots, \pi_1(\mathbf{b}_k)$ is in turn HKZ reduced. So, HKZ reducedness can be easily checked in time $k \cdot 2^{O(n)}$ by means of k SVP computations. Also, notice that if $\mathbf{b}_1, \ldots, \mathbf{b}_k$ is HKZ reduced, then all of its prefixes $\mathbf{b}_1, \ldots, \mathbf{b}_j$ ($j \leq k$) are also HKZ reduced. So, the DSP algorithm can be modified by always checking that the partial basis $\mathbf{b}_1, \ldots, \mathbf{b}_j$ is HKZ reduced. While this may not yield major asymptotic improvements in the running time, it is likely to reduce the search space by quite a bit, as HKZ reduced bases are only a small fraction of all bases of a given lattice with vectors of bounded norm.

Another optimization is obtained by observing that the shortest vector in M is always much shorter than $n\lambda_1(\Lambda)$, as shown in the next lemma.

Lemma 3.3. *Let $M \subseteq \Lambda$ be a k-dimensional sub-lattice of minimum determinant. Then*

$$\lambda_1(M) \leq \gamma_k \lambda_1(\Lambda)$$

Proof. Let $\mathbf{b}_1, \ldots, \mathbf{b}_n$ be a lattice basis of Λ such that $M = (\mathbf{b}_1, \ldots, \mathbf{b}_k)$. We want to bound from above the quantity $\lambda_1(M)/\lambda_1(\Lambda)$. We first compute a lower bound on $\lambda_1(\Lambda)$.

Let π_k the orthogonal projection onto the complement of $\mathbf{b}_1, \ldots, \mathbf{b}_{k-1}$, and $\mathbf{b}_k^* = \pi_k(\mathbf{b}_k)$ the component of $\mathbf{b}_k$ orthogonal to $\mathbf{b}_1, \ldots, \mathbf{b}_{k-1}$. Notice that $\mathbf{b}_k^*$ must be a shortest vector in $\pi_k(\Lambda)$, because otherwise $\det(M) = \det([\mathbf{b}_1, \ldots, \mathbf{b}_{k-1}]) \cdot \|\mathbf{b}_k^*\|$ can be further reduced. Now, let $\mathbf{v} = \sum_i x_i \mathbf{b}_i$ the shortest vector of Λ. If $x_i \neq 0$ for some $i \geq k$, then $\pi_k(\mathbf{v}) \neq \mathbf{0}$ and $\|\mathbf{v}\| \geq \|\pi_k(\mathbf{v})\| \geq \lambda_1(\pi_k(\Lambda)) = \|\mathbf{b}_k^*\|$. Otherwise,

if $x_i = 0$ for all $i \geq k$, the vector $\mathbf{v}$ belongs to the lattice generated by M (in fact, it belongs to the sublattice generated by the first $k-1$ basis vectors of M), and $\|\mathbf{v}\| \geq \lambda_1(M)$. Either way, we have $\lambda_1(\Lambda) = \|\mathbf{v}\| \geq \min(\lambda_1(M), \|\mathbf{b}_k^*\|)$.

This gives an upper bound

$$\lambda_1(M)/\lambda_1(\Lambda) \leq \lambda_1(M)/\min(\lambda_1(M), \|\mathbf{b}_k^*\|) = \max(1, \lambda_1(M)/\|\mathbf{b}_k^*\|).$$

By Minkowski's theorem we know that $\lambda_1(M) \leq \sqrt{\gamma_k}\det(M)^{1/k}$. Now, assume without loss of generality that $\mathbf{b}_k^*/\|\mathbf{b}_k^*\|^2$ is a shortest vector in the dual lattice M^*, as this can be achieved without changing the lattice generated by M. Applying Minkowski's theorem to the dual of lattice M we get $1/\|\mathbf{b}_k^*\| \leq \sqrt{\gamma_k}\det(M)^{-1/k}$. Combining the two inequalities, we get

$$\lambda_1(M)/\lambda_1(\Lambda) \leq \sqrt{\gamma_k}\det(M)^{1/k}\sqrt{\gamma_k}\det(M)^{-1/k} = \gamma_k.$$

□

It follows from Lemma 3.3 that not all vectors in S need to be examined as potential first elements of M, but only those of length bounded by $\gamma_k\lambda_1(\Lambda)$. We remark that γ_k is much smaller than k. For example, Table 1 shows that the value of $\gamma_k = \gamma_{k,1}$ for $k \leq 8$ is bounded by 2. For arbitrary k, it is known that $\gamma_k \leq 1 + k/4$. Asymptotically, $\gamma_k \leq k \cdot (1.744\ldots/(2\pi e) + o(1)) \approx k/10$. So, the first vector of M can be chosen from a much smaller portion of S. Of course, the bound from Lemma 3.3 can also be applied recursively to the selection of the other basis vectors of M. Namely, after the first i vectors $\mathbf{b}_1, \ldots, \mathbf{b}_i$ have been selected, one can project Λ orthogonally to $\mathbf{b}_1, \ldots, \mathbf{b}_i$ and require that the projection of next vector $\mathbf{b}_{i+1}$ has length at most γ_{k-i} times the shortest vector in the projected lattice.

Another possible improvement is offered by employing better asymptotic estimates in the transference bound $\lambda_k \cdot \lambda_1^* \leq k/(2\pi)(1 + O(1/\sqrt{k}))$ used in Lemma 3.1.

Finally, the following simple lemma shows that also shortest dual lattice vectors offer useful information about the k-DSP solution.

Lemma 3.4. *Let Λ be an n-dimensional lattice and $M \subseteq \Lambda$ be a k-dimensional sub-lattice of minimum determinant. Let $\mathbf{u} \in \Lambda^*$ be any dual lattice vector. Then, either $\mathbf{u}$ is orthogonal to M or $\|\mathbf{u}\| \geq \lambda_1(M^*)$. In particular, either M is orthogonal to all the shortest nonzero vectors in Λ^*, or $\lambda_1(M^*) \leq \lambda_1(\Lambda^*)$.*

Proof. Notice that $M^* = \pi_M(\Lambda^*)$ where π_M is the orthogonal projection onto the linear space spanned by M. Assuming that $\mathbf{u}$ is not orthogonal to M, we get $\pi_M(\mathbf{u}) \neq 0$, and $\lambda_1(M^*) \leq \|\pi_M(\mathbf{u})\| \leq \|\mathbf{u}\|$. □

Using Lemma 3.4, one can consider as a separate case solutions M that are orthogonal to all the shortest dual vectors in Λ^*. This allows to reduce the dimension n of the lattice, leading to an easier to solve problem, similarly to what was already done by considering solutions that contain all vectors in Λ of length $\lambda_1(\Lambda)$. If the optimal solution is not found, then one can assume that the lattice Λ and the optimal solution M satisfy $1/\lambda_1(\Lambda^*) \leq 1/\lambda_1(M^*) \leq \lambda_1(\Lambda)$, and try to use this relation to further prune the enumeration.

4 Densest Sub-lattice Problem in General Norms

In this section we generalize our results for k-DSP and k-SPP to arbitrary norms. The definition of these problems for general norm was already given in Section 3. Here we first give some additional background on convex geometry in Section 4.1. Then, in Section 4.2 we establish an approximate duality relationship

between k-SPP and k-DSP, and describe how our algorithm for k-DSP can be used to solve k-SPP as well. The presentation of the algorithm for k-DSP in arbitrary norm follows the same structure used in Section 3 for the Euclidean case, first establishing some useful geometric properties in Section 4.3, and then using these properties to derive an algorithm in Section 4.4.

For an n-dimensional lattice Λ and convex body $K \subseteq \mathbb{R}^n$, for each $k \in [n]$, we define the lattice parameter

$$\tau_k(K, \Lambda) = \min_{\substack{\text{sublattice } M \subseteq \Lambda \\ \dim(M)=k}} \frac{\det(M)}{\text{vol}_k(K \cap \text{span}(M))} \tag{4.1}$$

The quantity $\tau_k(K, \Lambda)$ corresponds to the value of an optimal k-DSP solution with respect to K and Λ.

4.1 Preliminaries

Convex Bodies: For a subset $T \subseteq \mathbb{R}^n$, we write $W_T = \text{span}(T)$ for the linear span of T. For $K \subseteq \mathbb{R}^n$, a compact convex set containing the origin in its relative interior, we define

$$K^* = \{\mathbf{x} \in W_K : \langle \mathbf{x}, \mathbf{y} \rangle \leq 1 \ \ \forall \mathbf{y} \in K\}.$$

By classical duality we have that $(K^*)^* = K$.

A matrix $A \in \mathbb{R}^{n \times n}$ is positive semi-definite if A is symmetric and $\mathbf{x}^t A \mathbf{x} \geq 0$ for all $\mathbf{x} \in \mathbb{R}^n$. For a positive semi-definite matrix $A \in \mathbb{R}^{n \times n}$, define $\|\mathbf{x}\|_A = \sqrt{\mathbf{x}^t A \mathbf{x}}$ for $\mathbf{x} \in \mathbb{R}^n$. For a linear subspace $W \subseteq \mathbb{R}^n$, we let $W^\perp$ denote its orthogonal complement. We define $E_W(A) = \{\mathbf{x} \in W : \|\mathbf{x}\|_A \leq 1\}$. By convention, we write $E(A)$ for $E_{\mathbb{R}^n}(A)$. We note that $\|\mathbf{x}\|_A = \|\mathbf{x}\|_{E(A)}$, for $\mathbf{x} \in \mathbb{R}^n$. If A is non-singular when restricted to W, we define A_W^{-1}, the inverse of A restricted to W, as the unique positive semi-definite matrix satisfying $A_W^{-1} \pi_W = A_W^{-1}$ and $A A_W^{-1} = \pi_W$. From here, we note that $E_W(A)^* = E_W(A_W^{-1})$. We define the determinant of A restricted to W, to be $\det_W(A) = \det(O_W^t A O_W)$, where O_W denotes any orthonormal basis matrix for the subspace W (note that the quantity is invariant under change of basis). From here we may express the volume formula:

$$\text{vol}_r(E_W(A)) = \det_W(A)^{-\frac{1}{2}} \text{vol}_r(\mathcal{B}_2^r)$$

where $r = \dim(W)$.

For a non-empty symmetric compact convex set $K \subseteq \mathbb{R}^n$, $r = \dim(K)$, we define the covariance matrix of K as

$$\text{cov}(K)_{ij} = \frac{1}{\text{vol}_r(K)} \int_K \mathbf{x}_i \mathbf{x}_j d\mathbf{x} \quad \forall i, j \in [n]$$

where we integrate using the r-dimensional Lebesgue measure restricted to W_K. We note that $\text{cov}(K)$ is a positive semi-definite matrix. We define the Binet ellipsoid of K as $B_K = E_{W_K}(\text{cov}(K))$, and the inertial ellipsoid of K as $E_K = B_K^* = E_{W_K}(\text{cov}(K)_{W_K}^{-1})$.

Define the isotropic constant K, as

$$L_K = \frac{\det_{W_K}(\text{cov}(K))^{\frac{1}{2r}}}{\text{vol}_r(K)^{\frac{1}{r}}}$$

where $r = \dim(K)$.

Let $\mathcal{K}_n$ denote the set of symmetric convex bodies in $\mathbb{R}^n$, and let $L_n = \sup_{K \in \mathcal{K}_n} L_K$. Clearly, for a non-empty symmetric convex set $K \subseteq \mathbb{R}^n$, we have that $L_K \leq L_{\dim(K)}$. The following theorem of Klartag, which improves on the estimate of Bourgain [Bou86], gives the best known bound on the isotropic constant:

Theorem 4.1 (K06). *The isotropic constant L_n satisfies $\Omega(1) \leq L_n \leq O(n^{\frac{1}{4}})$*

We will need the following theorem of Hensley [Hen80], which relates the volume of hyperplane sections to a body's covariance:

Theorem 4.2. *Let $K \subseteq \mathbb{R}^n$ be a k-dimensional symmetric compact convex set. Then for $\mathbf{v} \in W_K$, $\|\mathbf{v}\|_2 = 1$, we have that*

$$\frac{1}{2\sqrt{3}} \frac{\mathrm{vol}_k(K)}{\mathrm{vol}_{k-1}(K \cap \mathbf{v}^\perp)} \leq \|\mathbf{v}\|_{\mathrm{cov}(K)} \leq \frac{1}{2} \frac{\mathrm{vol}_k(K)}{\mathrm{vol}_{k-1}(K \cap \mathbf{v}^\perp)}$$

The following gives bounds on how well the inertial ellipsoid approximates a convex body. The estimates below are from [KLS95]:

Theorem 4.3. *For a non-empty symmetric compact convex set $K \subseteq \mathbb{R}^n$, where $r = \dim(K)$, the inertial ellipsoid E_K satisfies*

$$\sqrt{\frac{r+2}{r}} \cdot E_K \subseteq K \subseteq \sqrt{r(r+2)} \cdot E_K \tag{4.2}$$

The following theorem gives upper and lower estimates on the volume product, a fundamental quantity in convex geometry. The upper bound follows from the work of Blashke [Bla18] and Santaló [San49]. The lower bound was first established by Bourgain and Milman [BM87], and was recently refined by Kuperberg [Kup08], as well as by Nazarov [Naz09], where Kuperberg achieves the best constants.

Theorem 4.4. *Let K be a symmetric convex body in $\mathbb{R}^n$. Then we have*

$$\left(\frac{\pi e(1+o(1))}{n}\right)^n \leq \mathrm{vol}_n(K)\mathrm{vol}_n(K^*) \leq \mathrm{vol}_n(\mathcal{B}_2^n)^2. \tag{4.3}$$

where the upper bound holds at equality if and only if K is an ellipsoid.

We remark that the upper and lower bounds match within a 2^n factor since $\mathrm{vol}(\mathcal{B}_2^n)^2 = \left(\frac{2\pi e(1+o(1))}{n}\right)^n$.

Lastly, we will need the classical Rogers-Shephard [RS57] inequality:

Theorem 4.5. *Let $K \subseteq \mathbb{R}^n$ be a convex body. Then*

$$\mathrm{vol}_n(K) \leq \mathrm{vol}_n(K - K) \leq \binom{2n}{n} \mathrm{vol}_n(K) < 4^n \mathrm{vol}_n(K)$$

where the second inequality holds at equality iff K is a simplex.

Logconcave functions: For a function $f : \mathbb{R}^n \to \mathbb{R}_+$ we denote $\mathrm{support}(f) = \overline{\{\mathbf{x} : f(\mathbf{x}) > 0\}}$, where $\overline{A}$ is the topological closure of the set A. f is *logconcave* if

$$f(\alpha\mathbf{x} + (1-\alpha)\mathbf{y}) \geq f^\alpha(\mathbf{x}) f^{1-\alpha}(\mathbf{y})$$

for all $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ and $\alpha \in [0,1]$. f is $\frac{1}{k}$-concave, $k \geq 1$, if forall $\mathbf{x}, \mathbf{y} \in \mathrm{support}(f)$ and $\alpha \in [0,1]$ we have that

$$f(\alpha\mathbf{x} + (1-\alpha)\mathbf{y})^{\frac{1}{k}} \geq \alpha f(\mathbf{x})^{\frac{1}{\mathbf{k}}} + (1-\alpha) f(\mathbf{y})^{\frac{1}{k}}$$

Note that $\frac{1}{k}$-concave functions are logconcave (also logconcave functions can be thought of as $\frac{1}{k}$-concave for $k = \infty$). Also, note that logconcavity of f implies that $\mathrm{support}(f)$ is a closed convex set.

Logconcave functions arise naturally in the study of convex bodies. In particular, the indicator function of a convex body is logconcave. More generally, a classic result is that any marginal of a logconcave function is again logconcave.

Theorem 4.6. *Let $f\colon \mathbb{R}^n \to \mathbb{R}_+$ be a logconcave function. Let $W \subseteq \mathbb{R}^n$ be a linear subspace. Then the marginal function $f^W : W \to \mathbb{R}_+$ defined by*

$$f^W(\mathbf{x}) = \int_{\mathbf{y}\in W^\perp} f(\mathbf{y}+\mathbf{x})\mathrm{d}\mathbf{y}, \quad \mathbf{x} \in W,$$

is logconcave. In particular, for any convex body $K \subseteq \mathbb{R}^n$, the marginal function $\mathrm{vol}_K^W : W \to \mathbb{R}_+$ defined by

$$\mathrm{vol}_K^W(\mathbf{x}) = \mathrm{vol}(K \cap (W^\perp + \mathbf{x})), \quad \mathbf{x} \in W,$$

is $\frac{1}{k}$-concave where $k = \dim(W^\perp)$.

From logconcave functions we can also build convex bodies. The following theorem of Ball [Bal88] describes this:

Theorem 4.7. *Let f be an integrable logconcave function (i.e. $\int_{\mathbb{R}^n} f(\mathbf{x})\mathrm{d}\mathbf{x} < \infty$) satisfying $f(0) \in \mathrm{interior}(\{\mathbf{x} \in \mathbb{R}^n : f(\mathbf{x}) > 0\})$. Then for $p \geq 0$, the following set*

$$K_p(f) = \{\mathbf{x} \in \mathbb{R}^n : \int_{r=0}^{\infty} f(r\mathbf{x}) r^{p-1}\mathrm{d}r \geq \frac{f(0)}{p}\}$$

is a convex body, corresponding to the norm

$$\|\mathbf{x}\|_{f,p} = \frac{1}{\sqrt[p]{p\int_{r\in\mathbb{R}_+} \frac{f(r\mathbf{x})}{f(0)} r^{p-1}\mathrm{d}r}}.$$

Furthermore, for $W \subseteq \mathbb{R}^n$ a linear subspace, $\dim(W) = p$, we have that

$$\int_W \frac{f(\mathbf{x})}{f(0)}\mathrm{d}\mathbf{x} = \mathrm{vol}(K_p(f) \cap W)$$

where we integrate with respect to the p-dimensional Lebesgue measure restricted to W.

Notice that if I_K is the indicator function of a convex body K, then $K_p(I_K) = K$ is the original convex body for any p, and $\|\mathbf{x}\|_{I_K,p} = \|\mathbf{x}\|_K$. One can also define $K_\infty(f)$ and the corresponding norm $\|\mathbf{x}\|_{f,\infty}$ as the limit of $K_p(f)$ and $\|\mathbf{x}\|_{f,p}$ for $p \to \infty$. The result $K_\infty(f)$ is a convex body, and equals precisely the support $\overline{\{\mathbf{x} \mid f(\mathbf{x}) > 0\}}$ of the function f.

Lemma 4.8. *Let $f : \mathbb{R}_+ \to \mathbb{R}_+$ be an integrable function such $\max_x f(x) = f(0) = 1$ and f is $\frac{1}{k}$-concave. Then for any $p \geq 1$ we have that*

$$\left(p\int_{\mathbb{R}_+} x^{p-1}f(x)\mathrm{d}x\right)^{\frac{1}{p}} \leq \mathrm{vol}_1(\mathrm{support}(f)) \leq \binom{k+p}{p}^{\frac{1}{p}} \left(p\int_{\mathbb{R}_+} x^{p-1}f(x)\mathrm{d}x\right)^{\frac{1}{p}}$$

Proof. By $\frac{1}{k}$-concavity of f, we have that $\mathrm{support}(f)$ is a closed interval, i.e. $\mathrm{support}(f) = [0, l]$ for some $l > 0$. We first claim that $l < \infty$. Since $f^{\frac{1}{k}}$ is concave on $[0, l]$ and $f(0) = 1$ and $f(l) \geq 0$, we see that $f(x) \geq \left(1 - \frac{x}{l}\right)^k$ for $x \in [0, l]$. Therefore

$$\int_{\mathbb{R}_+} f(x)\mathrm{d}x \geq \int_0^l \left(1-\frac{x}{l}\right)^k \mathrm{d}x = -\frac{l}{k+1}\left(1-\frac{x}{l}\right)^{k+1}\Big|_0^l = \frac{l}{k+1}$$

Since f is assumed to be integrable, we have that $\frac{l}{k+1} < \infty \Leftrightarrow l < \infty$ as needed.

For the first inequality, since $f(x) \leq 1$, we have that

$$\left(p\int_{\mathbb{R}_+} x^{p-1}f(x)\mathrm{d}x\right)^{\frac{1}{p}} \leq \left(p\int_0^l x^{p-1}\mathrm{d}x\right)^{\frac{1}{p}} = l = \mathrm{vol}_1(\mathrm{support}(f))$$

For the second inequality, since $f(x) \geq \left(1-\frac{x}{l}\right)^k$, a straightforward computation gives us that

$$\left(p\int_{\mathbb{R}_+} x^{p-1}f(x)\mathrm{d}x\right)^{\frac{1}{p}} \geq \left(p\int_0^l x^{p-1}\left(1-\frac{x}{l}\right)^k\mathrm{d}x\right)^{\frac{1}{p}} = \binom{k+p}{p}^{-\frac{1}{p}} l$$

as needed. □

Lemma 4.9. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body. Then for $\mathbf{x} \in \mathbb{R}^n$, we have that*

$$\|\mathbf{x}\|_K = \frac{2\|\mathbf{x}\|_2}{\mathrm{vol}_1(K \cap \mathbb{R}x)}$$

Furthermore, for a k-dimensional linear subspace W, for $\mathbf{x} \in W$ we have that

$$\|\mathbf{x}\|_{K_1^W} = \frac{2\mathrm{vol}_k(K \cap W)\|\mathbf{x}\|_2}{\mathrm{vol}_{k+1}(K \cap (W + \mathbb{R}\mathbf{x}))}$$

Proof. We prove the first equality. Note that when $\mathbf{x} = \mathbf{0}$, both expressions are 0, hence we may assume that $\mathbf{x} \neq \mathbf{0}$ and that $\|\mathbf{x}\|_K > 0$. Since $\|\mathbf{x}\|_K = \inf\{s \geq 0 : \mathbf{x} \in sK\} = \inf\{s \geq 0 : \frac{\mathbf{x}}{s} \in K\}$, we have that $K \cap \mathbb{R}_+\mathbf{x} = [0, \frac{\mathbf{x}}{\|\mathbf{x}\|_K}]$, i.e. the line segment from 0 to $\frac{\mathbf{x}}{\|\mathbf{x}\|_K}$. By symmetry of K, we get that $K \cap \mathbb{R}\mathbf{x} = [-\frac{\mathbf{x}}{\|\mathbf{x}\|_K}, \frac{\mathbf{x}}{\|\mathbf{x}\|_K}]$. Therefore we have that $\mathrm{vol}_1(K \cap \mathbb{R}\mathbf{x}) = 2\|\frac{\mathbf{x}}{\|\mathbf{x}\|_K}\|_2 = 2\frac{\|\mathbf{x}\|_2}{\|\mathbf{x}\|_K}$ and hence $\frac{2\|\mathbf{x}\|_2}{\mathrm{vol}_1(K \cap \mathbb{R}\mathbf{x})} = \|\mathbf{x}\|_K$ as needed.

For $\mathbf{x} \in W$, by the first part, we have that

$$\|\mathbf{x}\|_{K_1^W} = \frac{2\|\mathbf{x}\|_2}{\mathrm{vol}_1(K_1^W \cap \mathbb{R}\mathbf{x})}$$

By Theorem 4.7, we note that since $\mathbb{R}\mathbf{x}$ is 1 dimensional, we have that

$$\mathrm{vol}_1(K_1^W \cap \mathbb{R}\mathbf{x}) = \int_{\mathbb{R}\mathbf{x}} \frac{\mathrm{vol}_K^W(\mathbf{y})}{\mathrm{vol}_K^W(\mathbf{0})}\mathrm{d}\mathbf{y} = \int_{\mathbb{R}\mathbf{x}} \frac{\mathrm{vol}_k(K \cap (W + \mathbf{y}))}{\mathrm{vol}_k(K \cap W)}\mathrm{d}\mathbf{y} = \frac{\mathrm{vol}_{k+1}(K \cap (W + \mathbb{R}\mathbf{x}))}{\mathrm{vol}_k(K \cap W)}$$

Combining the last two equalities now yields the claim. □

We will need the following lemma which provides sandwiching estimates between a particular class of bodies introduced by Ball [Bal88] and an associated projected body:

Lemma 4.10. *Let $K \subseteq \mathbb{R}^n$ be a symmetric convex body. Take $W \subseteq \mathbb{R}^n$, $\dim(W) = n - k$, and $p \in \mathbb{N}$. Let $K_p^W = K_p(\mathrm{vol}_K^W)$, where vol_K^W denotes the marginal function as defined in 4.6. Then for $\mathbf{x} \in W$*

$$\binom{k+p}{p}^{-\frac{1}{p}} \|\mathbf{x}\|_{K_p^W} \leq \|\mathbf{x}\|_{K_\infty^W} = \|\mathbf{x}\|_{\pi_W(K)} \leq \|\mathbf{x}\|_{K_p^W}$$

Proof. By convexity, it is easy to see that $\text{support}(\text{vol}_K^W) = \overline{\{\mathbf{x} \in W : \text{vol}(\mathbf{x} + W^\perp \cap K) > 0\}} = \pi_W(K)$. We claim that $\text{vol}_K^W(\mathbf{y}) \leq \text{vol}_K^W(0)$ for $\mathbf{y} \in \pi_W(K)$. By central symmetry of K, we know that

$$\text{vol}_K^W(\mathbf{y}) = \text{vol}_k((\mathbf{y} + W^\perp) \cap K) = \text{vol}((-\mathbf{y} + W^\perp) \cap K) = \text{vol}_K^W(-\mathbf{y})$$

Hence by $\frac{1}{k}$-concavity of vol_K^W, we have that

$$\text{vol}_K^W(\mathbf{y})^{\frac{1}{k}} = \frac{1}{2}(\text{vol}_K^W(\mathbf{y}))^{\frac{1}{k}} + \frac{1}{2}(\text{vol}_K^W(-\mathbf{y}))^{\frac{1}{k}} \leq (\text{vol}_K^W(\frac{1}{2}\mathbf{y} + \frac{1}{2}(-\mathbf{y})))^{\frac{1}{k}} = (\text{vol}_K^W(\mathbf{0}))^{\frac{1}{k}}$$

as needed. Define $h : \mathbb{R}_+ \to \mathbb{R}_+$ as $h(r) = \frac{\text{vol}_K^W(r\mathbf{x})}{\text{vol}_K^W(\mathbf{0})}$. Note that $\max_{x \geq 0} h(x) = h(0) = 1$, h is $\frac{1}{k}$-concave and

$$\begin{aligned}\text{vol}_1(\text{support}(h)) &= \text{vol}_1(\{r \geq 0 : \text{vol}_K^W(r\mathbf{x}) > 0\}) = \text{vol}_1(\{r \geq 0 : r\mathbf{x} \in \pi_W(K)\}) \\ &= \text{vol}_1(\{r : 0 \leq r \leq \frac{1}{\|\mathbf{x}\|_{\pi_W(K)}}\}) = \frac{1}{\|\mathbf{x}\|_{\pi_W(K)}}\end{aligned}$$

Now remembering that

$$\|\mathbf{x}\|_{K_p^W} = \left(p\int_0^\infty \frac{\text{vol}_K^W(r\mathbf{x})}{\text{vol}_K^W(\mathbf{0})} r^{p-1}\text{d}r\right)^{-\frac{1}{p}} = \left(p\int_0^\infty h(r)r^{p-1}\text{d}r\right)^{-\frac{1}{p}}$$

Applying Lemma 4.8 to h, we have that

$$\|\mathbf{x}\|_{\pi_W(K)} \leq \|\mathbf{x}\|_{K_p^W} \leq \binom{k+p}{p}^{\frac{1}{p}} \|\mathbf{x}\|_{\pi_W(K)}$$

as needed. □

4.2 Duality Relations

In this section, we exhibit an approximate duality relation between k-SPP and k-DSP in general norms.

Lemma 4.11. *Let $K, \Lambda \subseteq \mathbb{R}^n$ denote a convex body and lattice in $\mathbb{R}^n$. Let W denote a lattice subspace of Λ^*, where $\dim(W) = k$, and let $M = \Lambda^* \cap W$. Then*

$$\left(\frac{\pi e(1+o(1))}{4k}\right)^k \frac{\det(M)}{\text{vol}_k((K-K)^* \cap W)} \leq \frac{\text{vol}_k(\pi_W(K))}{\det(\pi_W(\Lambda))} \leq \left(\frac{2\pi e(1+o(1))}{k}\right)^k \frac{\det(M)}{\text{vol}_k((K-K)^* \cap W)}$$

Proof. By a standard duality argument, we note that $\pi_W(\Lambda)^* = \Lambda^* \cap W = M$ and hence $\det(\pi_W(\Lambda)) = \det(M)^{-1}$. Next we know that $\pi_W(K-K) = \pi_W(K) - \pi_W(K)$, where we note that $\pi_W(K-K)$ is symmetric. Therefore, by the Blashke-Santaló inequality (upper bound in Theorem 4.4), we have that

$$\text{vol}_k(\pi_W(K)) \leq \text{vol}_n(\pi_W(K-K)) \leq \left(\frac{2\pi e(1+o(1))}{k}\right)^k \frac{1}{\text{vol}_k(\pi_W(K-K)^*)}$$

Using the a standard duality relation, we see that $\pi_W(K-K)^* = (K-K)^* \cap W$ and hence $\text{vol}_k(\pi_W(K-K)^*) = \text{vol}_k((K-K)^* \cap W)$. Next, by the Rogers-Shepard and Bourgain-Milman inequality (Theorem 4.5 and lower bound in 4.4 respectively), we see that

$$\text{vol}_k(\pi_W(K)) \geq \frac{1}{4^k}\text{vol}_k(\pi_W(K-K)) \geq \left(\frac{\pi e(1+o(1))}{4k}\right)^k \frac{1}{\text{vol}_k((K-K)^* \cap W)}$$

Putting together these inequalities, yields the statement. □

Using the above estimates, we get the following immediate corollary:

Corollary 4.12. *The k-SPP problem with respect to K and Λ can be approximated up to an 8^k factor (ignoring lower order terms) by the k-DSP problem with respect to $(K-K)^*$ and Λ^*.*

From the perspective of intended applications, for example with respect to Integer Programming (see Section 5), we are mostly interested in the $1/k$-power of the k-SPP and k-DSP value, and hence the above reduction yields an essentially "constant" factor approximation, which is more than enough for our purposes. Given this reduction, we will focus solely on giving algorithms for k-DSP problem. Furthermore, we note that our methods can easily be adapted to output all approximate solutions to the k-DSP (in essentially the same running time), and hence modulo technicalities associated with volume computation, we can essentially get an exact algorithm for k-SPP with the above reduction.

4.3 Geometry of Densest Sublattices

The main goal of this section is to establish the general norm analogue of Lemma 3.1. This will enable us to employ the same algorithmic approach for solving k-DSP under general norms as was used in the ℓ_2 setting. We obtain this result in Lemma 4.14, where we show that a k-DSP minimizing sublattice either contains all shortest vectors or has covering radius by $O(k^2\lambda_1)$. In Section 4.5, we study further geometric properties of densest sublattices.

The following Lemma provides us with the main relation needed for us to reason about densest sublattices.

Lemma 4.13. *Let $K, \Lambda \subseteq \mathbb{R}^n$ be a symmetric convex body and n-dimensional lattice respectively. For k, $2 \leq k \leq n$, let $N \subseteq M \subseteq \Lambda$, be $k-1$ and k dimensional sublattices respectively satisfying*

$$\tau_{k-1}(K, M) = \frac{\det(N)}{\mathrm{vol}_{k-1}(K \cap W_N)} \quad \textit{and} \quad \tau_k(K, \Lambda) = \frac{\det(M)}{\mathrm{vol}_k(K \cap W_M)}.$$

Then letting E_M denote the inertial ellipsoid of $K_M = K \cap W_M$, we have that

$$\frac{1}{\sqrt{3}\lambda_1(E_M^*, M^*)} \leq 2\frac{\tau_k(K, \Lambda)}{\tau_{k-1}(K, M)} = \lambda_1(K_1^{W_N^\perp}, \pi_{W_N^\perp}(\Lambda)) \leq \frac{1}{\lambda_1(E_M^*, M^*)}$$

Proof. We begin by establishing the two equalities above. Take $\mathbf{x} \in \Lambda \setminus N$, and let $\mathbf{x}^* = \pi_{W_N^\perp}(\mathbf{x})$. By Lemma 4.9, we note that

$$\begin{aligned} \frac{\det(N + \mathbb{Z}\mathbf{x})}{\mathrm{vol}_k(K \cap (W_N + \mathbb{R}\mathbf{x}))} &= \frac{\det(N)}{\mathrm{vol}_{k-1}(K \cap W_N)} \cdot \frac{\mathrm{vol}_{k-1}(K \cap W_N)\|\mathbf{x}^*\|_2}{\mathrm{vol}_k(K \cap (W_N + \mathbb{R}\mathbf{x}^*))} \\ &= \tau_{k-1}(K, M) \cdot \frac{1}{2}\|\mathbf{x}^*\|_{K_1^{W_N^\perp}} \end{aligned} \tag{4.4}$$

Take $\mathbf{y} \in \Lambda \setminus N$, such that $N + \mathbb{Z}\mathbf{y} = M$, and let $\mathbf{y}^* = \pi_{W_N^\perp}(\mathbf{y})$. Since M achieves the minimum for $\tau_k(K, \Lambda)$, by Equation (4.4) we see that $\frac{\tau_k(K,\Lambda)}{\tau_{k-1}(K,M)} = \frac{1}{2}\|\mathbf{y}^*\|_{K_1^{W_N^\perp}}$. Furthermore, for $\mathbf{x} \in \Lambda \setminus N$ we have that

$$\begin{gathered} \frac{\det(M)}{\mathrm{vol}_k(K \cap W_M)} \leq \frac{\det(N + \mathbb{Z}\mathbf{x})}{\mathrm{vol}_k(K \cap (W_N + \mathbb{R}\mathbf{x}))} \Leftrightarrow \\ \|\mathbf{y}^*\|_{K_1^{W_N^\perp}} \leq \|\mathbf{x}^*\|_{K_1^{W_N^\perp}} \end{gathered}$$

Since $\pi_{W_N^\perp}(\Lambda \setminus N) = \pi_{W_N^\perp}(\Lambda) \setminus \{\mathbf{0}\}$, we have that

$$2\frac{\tau_k(K,\Lambda)}{\tau_{k-1}(K,M)} = \|\mathbf{y}^*\|_{K_1^{W_N^\perp}} = \min_{\mathbf{x}\in\pi_{W_N^\perp}(\Lambda)\setminus\{\mathbf{0}\}} \|\mathbf{x}\|_{K_1^{W_N^\perp}} = \lambda_1(K_1^{W_N^\perp}, \pi_{W_N^\perp}(\Lambda))$$

as needed.

We now prove the two inequalities. Take $\mathbf{y} \in M^*$. Since $\mathbf{y}$ is a nonzero dual element, we have that $M \cap (\mathbb{R}\mathbf{y})^\perp$ is a $k-1$ dimensional sublattice of M, and $\pi_{\mathbb{R}\mathbf{y}}(M)$ is a lattice projection. From here we have that

$$\begin{aligned}\det(M)\|\mathbf{y}\|_2 &= \det(M\cap\mathbf{y}^\perp)\det(\pi_{\mathbf{y}}(M))\|\mathbf{y}\|_2 = \det(M\cap\mathbf{y}^\perp)\frac{\|\mathbf{y}\|_2}{\det(\pi_{\mathbf{y}}(M)^*)}\\ &= \det(M\cap\mathbf{y}^\perp)\frac{\|\mathbf{y}\|_2}{\det(M^*\cap\mathbb{R}\mathbf{y})} \geq \det(M\cap\mathbf{y}^\perp)\end{aligned} \tag{4.5}$$

where the last inequality holds with equality iff $\mathbb{Z}\mathbf{y} = M^* \cap \mathbb{R}\mathbf{y}$, i.e. iff $\mathbf{y}$ is primitive. Furthermore, it is a basic fact that any sublattice of $T \subseteq M$ of codimension 1, satisfying $T \cap M = \mathrm{span}(T) \cap M$, can be obtained as $T = M \cap \mathbb{R}\mathbf{y}^\perp$ for a primitive vector $\mathbf{y} \in M^* \setminus \{\mathbf{0}\}$.

We remember that the E_M^* (polar of the inertial ellipsoid) is the binet ellipsoid of K_M, and hence $\|\mathbf{y}\|_{E_M^*} = \|\mathbf{y}\|_{\mathrm{cov}(K_M)}$ for $\mathbf{y} \in W_M$. From Theorem 4.2, we have that

$$\begin{aligned}&\frac{1}{2\sqrt{3}}\frac{\mathrm{vol}_k(K_M)}{\mathrm{vol}_{k-1}(K\cap\mathbf{y}^\perp)} \leq \|\mathbf{y}\|_{E_M^*} \leq \frac{1}{2}\frac{\mathrm{vol}_k(K_M)}{\mathrm{vol}_{k-1}(K\cap\mathbf{y}^\perp)} \quad \forall \mathbf{y}\in W_K, \|\mathbf{y}\|_2 = 1 \quad \Leftrightarrow\\ &\frac{1}{2\sqrt{3}}\frac{\mathrm{vol}_k(K_M)}{\mathrm{vol}_{k-1}(K\cap\mathbf{y}^\perp)}\|\mathbf{y}\|_2 \leq \|\mathbf{y}\|_{E_M^*} \leq \frac{1}{2}\frac{\mathrm{vol}_k(K_M)}{\mathrm{vol}_{k-1}(K\cap\mathbf{y}^\perp)}\|\mathbf{y}\|_2 \quad \forall \mathbf{y}\in W_K\end{aligned} \tag{4.6}$$

By Equations (4.5) and (4.6), for $\mathbf{y} \in M^* \setminus \{\mathbf{0}\}$, we have that

$$\begin{aligned}\frac{\det(M)}{\mathrm{vol}_k(K_M)}\|\mathbf{y}\|_{E_M^*} &\geq \frac{1}{2\sqrt{3}}\frac{\det(M)}{\mathrm{vol}_k(K_M)}\frac{\mathrm{vol}_k(K_M)\|\mathbf{y}\|_2}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)}\\ &= \frac{1}{2\sqrt{3}}\frac{\det(M)\|\mathbf{y}\|_2}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)} \geq \frac{1}{2\sqrt{3}}\frac{\det(M\cap\mathbf{y}^\perp)}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)}\end{aligned} \tag{4.7}$$

Furthermore, if $\mathbf{y}$ is primitive, we have that

$$\begin{aligned}\frac{\det(M)}{\mathrm{vol}_k(K_M)}\|\mathbf{y}\|_{E_M^*} &\leq \frac{1}{2}\frac{\det(M)}{\mathrm{vol}_k(K_M)}\frac{\mathrm{vol}_k(K_M)\|\mathbf{y}\|_2}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)}\\ &= \frac{1}{2}\frac{\det(M)\|\mathbf{y}\|_2}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)} = \frac{1}{2}\frac{\det(M\cap\mathbf{y}^\perp)}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)}\end{aligned} \tag{4.8}$$

From Equation (4.7), we now have that

$$\begin{aligned}\tau_k(K,\Lambda)\lambda_1(E_M^*, M^*) &= \min_{\mathbf{y}\in M^*\setminus\{\mathbf{0}\}}\frac{\det(M)}{\mathrm{vol}_k(K_M)}\|\mathbf{y}\|_{E_M^*} \geq \frac{1}{2\sqrt{3}}\min_{\mathbf{y}\in M^*\setminus\{\mathbf{0}\}}\frac{\det(M\cap\mathbf{y}^\perp)}{\mathrm{vol}_{k-1}(K_M\cap\mathbf{y}^\perp)}\\ &\geq \frac{1}{2\sqrt{3}}\tau_{k-1}(K,M).\end{aligned}$$

Therefore, $2\frac{\tau_k(K,\Lambda)}{\tau_{k-1}(K,M)} \geq \frac{1}{\sqrt{3}}\frac{1}{\lambda_1(E_M^*,M^*)}$ as needed.

For the reverse inequality, let $\mathbf{y} \in \Lambda^*$ be the primitive nonzero vector satisfying $N = M \cap \mathbf{y}^\perp$. Note that such a vector must exist, since otherwise $N \subsetneq M \cap \mathrm{span}(N)$ and hence

$$\frac{\det(M \cap \mathrm{span}(N))}{\mathrm{vol}_{k-1}(K \cap \mathrm{span}(N))} < \frac{\det(N)}{\mathrm{vol}_{k-1}(K \cap \mathrm{span}(N))} = \tau_{k-1}(K, M),$$

a clear contradiction. Using Equation (4.8), we now see that

$$\tau_k(K,\Lambda)\lambda_1(E_M^*, M^*) \leq \frac{\det(M)}{\mathrm{vol}_k(K_M)}\|\mathbf{y}\|_{E_M^*} \leq \frac{1}{2}\frac{\det(N)}{\mathrm{vol}_{k-1}(K \cap W_N)} = \frac{1}{2}\tau_{k-1}(K,\Lambda).$$

Therefore, $2\frac{\tau_k(K,\Lambda)}{\tau_{k-1}(K,M)} \leq \frac{1}{\lambda_1(E_M^*, M^*)}$ as need. □

Lemma 4.14. *Let $M \subseteq \Lambda$ be a minimizer for k-DSP with respect to the norm induced by $K \subseteq \mathbb{R}^n$. Then for any $\mathbf{y} \in \Lambda$, either $\mathbf{y} \in M$ or $\mu(K, M) \leq \frac{\sqrt{3}}{2}k^2\|\mathbf{y}\|_K$. In particular, either M contains all vectors of length $\lambda_1(K,\Lambda)$, or $\mu(K, M) \leq \frac{\sqrt{3}}{2}k^2\lambda_1(K,\Lambda)$.*

Proof. Assume that $2 \leq k \leq n-1$. Take $\mathbf{y} \in \Lambda$ and assume that $\mathbf{y} \notin M$. Let $N \subseteq M$ be as in Theorem 4.13. By Lemma 4.10, we have that

$$\|\mathbf{y}\|_K \geq \|\mathbf{y}\|_{\pi_{W_N^\perp}(K)} \geq \binom{k-1+1}{1}^{-1}\|\mathbf{y}\|_{K_1^{W_N^\perp}} = \frac{1}{k}\|\mathbf{y}\|_{K_1^{W_N^\perp}} \geq \frac{1}{k}\lambda_1(K_1^{W_N^\perp}, \pi_{W_N^\perp}(\Lambda))$$

By Theorem 4.13, we have that

$$\lambda_1(K_1^{W_N^\perp}, \pi_{W_N^\perp}(\Lambda)) \geq \frac{1}{\sqrt{3}\lambda_1(E_M^*, M^*)}$$

where E_M is the inertial ellipsoid of $K_M = K \cap W_M$. By the transference theorem for ℓ_2 [Ban93a], we get

$$\frac{\sqrt{3}}{2}k^2\|\mathbf{y}\|_K \geq \frac{k}{2\lambda_1(E_M^*, M^*)} \geq \mu(E_M, M)$$

By Theorem 4.3 we have that $E_M \subseteq K_M$, and hence $\mu(E_M, M) \geq \mu(K, M)$. Combining this with the above inequality yields the result. □

4.4 Algorithms

In this section, we give an algorithm to solve the k-DSP in general norms (Algorithm 2). The structure of the algorithm is essentially identical to the one for the Euclidean norm (Algorithm 1). The main difference is that the analysis is more mathematically involved, due to the need to bound many convex geometric quantities related to the sectional volumes of convex bodies (the bulk of this analysis appears in the previous section).

One technical detail is that our k-DSP algorithm requires the ability to compute volumes of k-dimensional sections of a symmetric convex body. In general, computing such volumes exactly in the oracle model is impossible; however, estimating these volumes within $(1+\epsilon)$ is achievable. For our purposes, this means that we cannot hope to compute anything more than a $(1+\epsilon)$-approximate minimizer to the k-DSP problem (which is sufficient for all the intended applications). The main method to estimate volumes in the oracle model uses random sampling techniques based on geometric random walks. These techniques yield polynomial time Monte Carlo algorithms for $(1+\epsilon)$ volume estimation (see [DFK91]). Even though these algorithms

are randomized, the probability of error for $(1+\epsilon)$ volume estimation can easily be reduced to $2^{-\Omega(\mathrm{poly}(n))}$ in polynomial time, which is much smaller than the inverse of the running time of our k-DSP algorithm. Therefore, we can assume that during the course of the algorithm, all calls to the volume oracle return correct $(1+\epsilon)$-approximations (where this assumption holds with overwhelming probability).

Given the previous remarks, we will for simplicity, assume that our k-DSP algorithm has access to an exact deterministic oracle for computing volumes. We will therefore refer to the runtime of the algorithm as the number of arithmetic operations and calls to the volume and norm oracles. We note that dealing with an approximate oracle makes no essential difference in the analysis: it is easy to check that the same algorithm, given an $(1+\epsilon)$-approximate volume oracle, outputs a $(1+O(\epsilon))$ approximate solution to the k-DSP.

Algorithm 2 Algorithm DSP(K, Λ, k, H)

Input: An (r, R)-centered symmetric convex body $K \subseteq \mathbb{R}^n$ presented by a membership oracle, a basis $(\mathbf{b}_1, \ldots, \mathbf{b}_n)$ for a lattice $\Lambda \subseteq \mathbb{R}^n$, a target dimension $0 \leq k \leq n$, and a lattice subspace $H \subseteq \mathrm{span}(\Lambda)$, $\dim(H) \leq k$.

Output: W a lattice subspace of Λ, $\dim(W) = k$, $H \subseteq W$, W minimizes $\frac{\det(\Lambda \cap W)}{\mathrm{vol}_k(K \cap W)}$ over all k-dimensional lattice subspaces containing H.

$d \leftarrow k - \dim(H)$
if $d = 0$ **then**
  return H
$K^\pi = \pi_{H^\perp}(K), \quad \Lambda^\pi = \pi_{H^\perp}(\Lambda)$
$W \leftarrow \mathrm{NULL}$
$V \leftarrow \mathrm{span}(\{\mathbf{y} \in \Lambda^\pi : \|\mathbf{y}\|_{K^\pi} \leq \lambda_1(K^\pi, \Lambda^\pi)\})$
if $\dim(V) \leq d$ **then**
  $W \leftarrow \mathrm{DSP}(K, \Lambda, k, H + V)$
for all $T \subseteq \{\mathbf{y} \in \Lambda^\pi : \|\mathbf{y}\|_{K^\pi} \leq \sqrt{3}ekd\, \lambda_1(K^\pi, \Lambda^\pi)\}$ of size $|T| = d$ **do**
  $W' \leftarrow H + \mathrm{span}(T)$
  if $\dim(W') = k \wedge \left(W = \mathrm{NULL} \vee \frac{\det(\Lambda \cap W')}{\mathrm{vol}_k(K \cap W')} < \frac{\det(\Lambda \cap W)}{\mathrm{vol}_k(K \cap W)}\right)$ **then**
    $W \leftarrow W'$
return W.

Theorem 4.15 (k-DSP Correctness)**.** *Algorithm 2 is correct. A call to the algorithm on input* $\mathrm{DSP}(K, \Lambda, k, \emptyset)$ *runs in* $2^{O(n)} k^{2kn}\, \mathrm{poly}(n)$ *time and uses* $2^n\, \mathrm{poly}(n)$ *space.*

Proof.

Correctness: We proceed by induction on k. For $k = 0$, the algorithm is trivially correct, so let us assume that $k \geq 1$. We show that the recursive step in the algorithm is correct.

Let K^π, Λ^π and d be defined as in the algorithm, and let $p = \dim(H)$. Let M denote a k-dimensional lattice subspace of Λ containing H which minimizes $\frac{\det(\Lambda \cap M)}{\mathrm{vol}_k(K \cap M)}$ subject to containing H. Let $M^\pi = \pi_{H^\perp}(M)$, where we note that $\dim(M^\pi) = \dim(M) - \dim(H) = k - p = d$ (where d is as defined as in the algorithm). Then we have the identity

$$\frac{\det(\Lambda \cap M)}{\mathrm{vol}_k(K \cap M)} = \frac{\det(\Lambda \cap H)}{\mathrm{vol}_p(K \cap H)} \frac{\mathrm{vol}_p(K \cap H)\det(\Lambda^\pi \cap M^\pi)}{\mathrm{vol}_k(K \cap (H + M^\pi))} \tag{4.9}$$

Let $\mathrm{vol}_W^{H^\perp}$ denote the marginal function of K on $H^\perp$, that is $\mathrm{vol}_K^{H^\perp}(\mathbf{x}) = \mathrm{vol}_p((\mathbf{x}+H)\cap K)$ for $\mathbf{x}\in H^\perp$. Then since $\dim(M^\pi) = d$, by Theorem 4.7 we have that

$$\frac{\mathrm{vol}_k(K\cap(H+M^\pi))}{\mathrm{vol}_p(K\cap H)} = \int_{M^\pi} \frac{\mathrm{vol}_K^{H^\perp}(\mathbf{x})}{\mathrm{vol}_K^{H^\perp}(\mathbf{0})}\mathrm{d}\mathbf{x} = \mathrm{vol}_d(K_d^{H^\perp}\cap M^\pi). \tag{4.10}$$

Combining Equations (4.9) and (4.10), we see that

$$\frac{\det(\Lambda\cap M)}{\mathrm{vol}_k(K\cap M)} = \frac{\det(\Lambda\cap H)}{\mathrm{vol}_p(K\cap H)}\frac{\det(\Lambda^\pi\cap M^\pi)}{\mathrm{vol}_d(K_d^{H^\perp}\cap M^\pi)} \tag{4.11}$$

Since Equation (4.11) holds for any k-dimensional lattice subspace containing H, we conclude that finding the k-DSP minimizer for K and Λ containing H is equivalent to solving d-DSP for $K_d^{H^\perp}$ and Λ^π.

We now show that at the end of the current node of the recursion, the algorithm has computed a k-dimensional lattice subspace W, $H\subseteq W$, whose density is at least as high as M's, i.e. for which

$$\frac{\det(\Lambda\cap W)}{\mathrm{vol}_k(K\cap W)} \le \frac{\det(\Lambda\cap M)}{\mathrm{vol}_k(K\cap M)}$$

This will suffice to prove correctness.

In the first step, the algorithm computes the subspace V, corresponding to the span of the shortest vectors of Λ^π under the norm induced by K^π. If $V\subseteq M^\pi$, then clearly $\dim(V)\le\dim(M^\pi) = \dim(M) - \dim(H) = d$. The algorithm therefore enters the if statement at line 6. By the induction hypothesis, after the execution of line 7, we have that W is a k-DSP minimizer for K and Λ containing $H+V$. Since $H+V\subseteq H+M^\pi = M$, we have that the value of W is at least as good as M, and hence is also optimal. Note that by optimality, the subspace W remains unchanged after the execution of the for loop on lines 8-11, since no other valid lattice subspace can increase the density of W. Hence when $V\subseteq M^\pi$, the recursive step correctly returns a k-DSP minimizer for K and Λ containing H.

Now assume that V is not contained in M^π. We will now show that one of the iterations of the for loop on lines 8-11 corresponds to the subspace M^π, and hence the subspace W returned by the algorithm at this recursion node will have density as least as large as M's (since $H+M^\pi = M$). This will suffice to prove correctness.

To show that M^π is found by the for loop, it suffices to show that M^π is spanned by lattice vectors of Λ^π whose norm with respect to K^π is bounded by $\sqrt{3}ekd\lambda_1(K^\pi,\Lambda^\pi)$ (since $\dim(M^\pi)=d$). In particular, it suffices to show that $\mu(K^\pi,M^\pi)\le\frac{\sqrt{3}ekd}{2}\lambda_1(K^\pi,\Lambda^\pi)$ since $\lambda_d(K^\pi,M^\pi)\le 2\mu(K^\pi,M^\pi)$.

Since V is not contained in M^π, there exists $\mathbf{y}\in\Lambda^\pi$, where $\|\mathbf{y}\|_{K^\pi} = \lambda_1(K^\pi,\Lambda^\pi)$ such that $\mathbf{y}\notin M^\pi$. Since M^π is a d-DSP minimizer with respect to $K_d^{H^\perp}$ and Λ^π, by Lemma 4.14 we have that

$$\mu(K_d^{H^\perp},M^\pi)\le\frac{\sqrt{3}}{2}d^2\|\mathbf{y}\|_{K_d^{H^\perp}} \tag{4.12}$$

Since $K^\pi = K_\infty^{H^\perp}$ by Lemma 4.10, we have that

$$\|\mathbf{y}\|_{K_d^{H^\perp}} \le \binom{p+d}{d}^{\frac{1}{d}}\|\mathbf{y}\|_{K^\pi} = \binom{k}{d}^{\frac{1}{d}}\lambda_1(K^\pi,\Lambda^\pi)\le\frac{ek}{d}\lambda_1(K^\pi,\Lambda^\pi) \tag{4.13}$$

Again by Lemma 4.10, we have that $\|\mathbf{x}\|_{K^\pi}\le\|\mathbf{x}\|_{K_d^{H^\perp}}\ \forall\mathbf{x}\in H^\perp \Leftrightarrow K_d^{H^\perp}\subseteq K^\pi$. Hence we get that $\mu(K^\pi,\Lambda^\pi)\le\mu(K_d^{H^\perp},\Lambda^\pi)$. Combining this with Equations (4.12) and (4.13), we obtain

$$\mu(K^\pi,\Lambda^\pi)\le\mu(K_d^{H^\perp},\Lambda^\pi)\le\frac{\sqrt{3}}{2}d^2\|\mathbf{y}\|_{K_d^{H^\perp}}\le\frac{\sqrt{3}ekd}{2}\lambda_1(K^\pi,\Lambda^\pi)$$

as needed. This completes the proof of correctness.

Runtime: We first note that the recursion tree for Algorithm 2 corresponds to a simple path, since we create at most one subproblem per recursion node. Furthermore, the length of this path is clearly at most k, since we start with $H = \{\mathbf{0}\}$, and add at least 1 dimension to H at each step.

Let us examine the total amount of work done at a recursion node indexed by the partial k-DSP solution H, where $\dim(H) = p \leq k$. Note that K^π and Λ^π are both $n - p$ dimensional. First, the time to compute the subspace V is bounded by the time to enumerate all the shortest non-zero vectors of Λ^π under the norm induced by K^π. This requires $\mathrm{poly}(n)2^{O(n-p)} = 2^{O(n)}$ time and $2^{n-p}\,\mathrm{poly}(n) = 2^n\,\mathrm{poly}(n)$ space using the Lattice Point Enumerator (Theorem 2.2). Next, in the for loop, we iterate over all subsets of size at most $d = k - p$ in $\{\mathbf{y} \in \Lambda^\pi : \|\mathbf{y}\|_{K^\pi} \leq \sqrt{3}ekd\lambda_1(K^\pi, \Lambda^\pi)\}$. Iterating over this set of vectors once requires $\mathrm{poly}(n)2^{O(n-p)}(kd)^{n-p} = 2^{O(n)}k^{2n}$ time and $2^n\,\mathrm{poly}(n)$ space using the Lattice Point Enumerator. Hence iterating over subsets of size d can be achieved in time $2^{O(n)}k^{2nd} = 2^{O(n)}k^{2kn}$ using $2^n\,\mathrm{poly}(n)$ space.

Therefore summing over at all most k levels of the recursion, the entire algorithm requires at most $k2^{O(n)}k^{2kn} = 2^{O(n)}k^{2kn}$ time and $\mathrm{poly}(n)2^n$ space. □

4.5 Geometry of Densest Sublattices (Continued)

In this section, we explore further geometric properties of densest sublattices and the τ_i parameters for general norms. In Lemma 4.16, we give absolute bounds on the τ_i parameters in terms of the successive minima. In Lemma 4.17, we show that the τ_i can be bounded as a function of τ_{i+1}, i.e. that the τ_i parameters do not grow too rapidly as i decreases. In Lemma 4.18, we generalize Lemma 3.3 to arbitrary norms, showing that the length of shortest nonzero vector of a k-DSP minimizer is at most a $O(k^{\frac{7}{4}})$ factor larger than the shortest nonzero vector of the lattice.

Lemma 4.16. *Let M be a minimizer for the Densest k-dimensional Sublattice Problem with respect to K and Λ. Then*

$$\frac{1}{2^k}\prod_{i=1}^{k}\lambda_i(K,\Lambda) \leq \tau_k(K,\Lambda) \leq \frac{k!}{2^k}\prod_{i=1}^{k}\lambda_k(K,\Lambda)$$

Furthermore, when $K = \mathcal{B}_2^n$, the factor $k!/2^k$ can be decreased to $\mathrm{vol}_k(\mathcal{B}_2^k)^{-1} = (\sqrt{\frac{k}{2\pi e}}(1 + o(1)))^k$.

Proof. Since $M \subseteq \Lambda$, $\dim(M) = k$, we immediately have that $\lambda_i(K,\Lambda) \leq \lambda_i(K,M)$ for $i \in [k]$. Next by Minkowski's Second Theorem we have that

$$\prod_{i=1}^{k}\lambda_i(K,M) \leq 2^k\frac{\det(M)}{\mathrm{vol}_k(K \cap \mathrm{span}(M))}. \tag{4.14}$$

This yields the first inequality. Pick linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k \in \Lambda$ such that $\|\mathbf{b}_i\|_K = \lambda_i(K,\Lambda)$. Let $S = \sum_{i=1}^{k}\mathbb{Z}\mathbf{b}_i$ denote the lattice generated by $\mathbf{b}_1, \ldots, \mathbf{b}_k$. Since M is a Densest k-dimensional Sublattice, by definition we have that

$$\frac{\det(M)}{\mathrm{vol}_k(K \cap \mathrm{span}(M))} \leq \frac{\det(S)}{\mathrm{vol}_k(K \cap \mathrm{span}(S))} \tag{4.15}$$

Let $\mathrm{Gram}(\mathbf{b}_1, \ldots, \mathbf{b}_k)$ denote the Gramm matrix of the vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k$, i.e. $\mathrm{Gram}(\mathbf{b}_1, \ldots, \mathbf{b}_k)_{ij} = \langle \mathbf{b}_i, \mathbf{b}_j \rangle$, $1 \leq i, j \leq k$. Next we see that

$$\begin{aligned}\det(S) = \det\left(\mathrm{Gram}(\mathbf{b}_1, \ldots, \mathbf{b}_k)\right)^{1/2} &= \det\left(\mathrm{Gram}\left(\frac{\mathbf{b}_1}{\|\mathbf{b}_1\|_K}, \ldots, \frac{\mathbf{b}_k}{\|\mathbf{b}_k\|_K}\right)\right)^{1/2} \prod_{i=1}^{k} \|\mathbf{b}_i\|_K \\ &= \det\left(\mathrm{Gram}\left(\frac{\mathbf{b}_1}{\lambda_1(K,\Lambda)}, \ldots, \frac{\mathbf{b}_k}{\lambda_k(K,\Lambda)}\right)\right)^{1/2} \prod_{i=1}^{k} \lambda_i(K,\Lambda)\end{aligned}$$

Let $\mathbf{b}_i' = \frac{\mathbf{b}_i}{\lambda_i(K,\Lambda)}$ for $i \in [k]$. Note that $\|\mathbf{b}_i'\|_K = \frac{\|\mathbf{b}_i\|_K}{\lambda_i(K,\Lambda)} = 1$ and hence $\mathbf{b}_i' \in K$. Since K is symmetric and convex, we have that $\mathrm{conv}\{\pm\mathbf{b}_1', \ldots, \pm\mathbf{b}_k'\} \subseteq K \cap \mathrm{span}(S)$ and hence

$$\mathrm{vol}_k(\mathrm{conv}\{0, \mathbf{b}_1', \ldots, \mathbf{b}_k'\}) = 2^{-k}\mathrm{vol}_k(\mathrm{conv}(\{\pm\mathbf{b}_1', \ldots, \pm\mathbf{b}_k'\})) \leq 2^{-k} \cdot \mathrm{vol}_k(K \cap \mathrm{span}(S))$$

where conv denotes the convex hull. Next, a standard computation yields that

$$\det(\mathrm{Gram}(\mathbf{b}_1', \ldots, \mathbf{b}_k'))^{1/2} = \mathrm{vol}_k(\{\sum_{i=1}^{k} a_i\mathbf{b}_i' : a_i \in [0,1], i \in [k]\}) = k!\,\mathrm{vol}_k(\mathrm{conv}\{0, \mathbf{b}_1', \ldots, \mathbf{b}_k'\})$$

Therefore

$$\begin{aligned}\frac{\det(S)}{\mathrm{vol}_k(K \cap \mathrm{span}(S))} &= \frac{\mathrm{vol}_k(\mathrm{conv}\{0, \mathbf{b}_1', \ldots, \mathbf{b}_k'\})}{\mathrm{vol}_k(K \cap \mathrm{span}(S))} k! \prod_{i=1}^{k} \lambda_i(K,\Lambda) \\ &\leq \frac{k!}{2^k} \prod_{i=1}^{k} \lambda_i(K,\Lambda)\end{aligned} \tag{4.16}$$

as needed.

When $K = \mathcal{B}_2^n$, this can be improved using Hadamard's inequality to get the bound

$$\det(\mathrm{Gram}(\mathbf{b}_1, \ldots, \mathbf{b}_k))^{1/2} \leq \prod_{i=1}^{k} \|\mathbf{b}_i\|_2 = \prod_{i=1}^{k} \lambda_i(\mathcal{B}_2^n, \Lambda).$$

Lastly we note that $\mathrm{vol}_k(\mathcal{B}_2^n \cap \mathrm{span}(S)) = \mathrm{vol}_k(\mathcal{B}_2^k)$. Plugging in these estimates yields the result. □

Lemma 4.17. *Let $K, \Lambda \subseteq \mathbb{R}^n$ be a symmetric convex body and n-dimensional lattice respectively. Then for $2 \leq k \leq n$, we have that*

$$\tau_{k-1}(K,\Lambda) \leq 2\sqrt{3\gamma_k} L_k \tau_k(K,\Lambda)^{\frac{k-1}{k}} = O(k^{\frac{3}{4}})\tau_k(K,\Lambda)^{\frac{k-1}{k}}$$

Proof. Let M_k denote a k dimensional sublattice attaining $\tau_k(K,\Lambda)$. Note that by inclusion, we have that $\tau_{k-1}(K,\Lambda) \leq \tau_{k-1}(K,M_k)$, hence it suffices to prove the claim for

$$\tau_{k-1}(K,M_k) \text{ and } \tau_k(K,M_k) = \frac{\det(M_k)}{\mathrm{vol}_k(K \cap \mathrm{span}(M_k))}.$$

Hence by making $\mathrm{span}(M_k)$ the ambient subspace, we may assume that $k = n$, and that $\mathrm{span}(M_k) = \mathbb{R}^n$.

Let E_K denote the inertial ellipsoid of K. From Theorem 4.13, we have that

$$\frac{1}{\sqrt{3}\lambda_1(E_K^*, \Lambda^*)} \leq 2\frac{\tau_n(K,\Lambda)}{\tau_{n-1}(K,\Lambda)} \Leftrightarrow \tau_{n-1}(K,\Lambda) \leq 2\sqrt{3}\lambda_1(E_K^*, \Lambda^*)\frac{\det(\Lambda)}{\mathrm{vol}_n(K)} \tag{4.17}$$

Noting that $E_K^* = \mathrm{cov}(K)^{-\frac{1}{2}}\mathcal{B}_2^n$, we have that

$$\begin{aligned}\lambda_1(E_K^*, \Lambda^*) &= \lambda_1(\mathcal{B}_2^n, \mathrm{cov}(K)^{\frac{1}{2}}\Lambda^*) \leq \sqrt{\gamma_n}\det(\mathrm{cov}(K)^{\frac{1}{2}}\Lambda^*)^{\frac{1}{n}} = \sqrt{\gamma_n}\frac{\det(\mathrm{cov}(K))^{\frac{1}{2n}}}{\det(\Lambda)^{\frac{1}{n}}} \\ &= \sqrt{\gamma_n}L_K\frac{\mathrm{vol}_n(K)^{\frac{1}{n}}}{\det(\Lambda)^{\frac{1}{n}}} \leq \sqrt{\gamma_n}L_n\frac{\mathrm{vol}_n(K)^{\frac{1}{n}}}{\det(\Lambda)^{\frac{1}{n}}}\end{aligned} \tag{4.18}$$

Substituting the bound from Equation (4.18) into Equation (4.17) directly yields the first inequality. To derive the asymptotic bound, we use the bounds $\sqrt{\gamma_k} = O(k^{\frac{1}{2}})$ and that $L_k = O(k^{\frac{1}{4}})$. □

Lemma 4.18. *Let $M \subseteq \Lambda$ be a minimizer for the k-DSP with respect to the norm induced by $K \subseteq \mathbb{R}^n$. Then*

$$\lambda_1(K,\Lambda) \leq 2\sqrt{3}L_k k\sqrt{\gamma_k}\lambda_1(K,\Lambda) = O(k^{\frac{7}{4}})\lambda_1(K,\Lambda)$$

Proof. We first note that $2\sqrt{3}L_k k\sqrt{\gamma_k} \geq 1$ for all $k \geq 1$. This follows from the fact that $L_k \geq L_{[-1,1]^k} = \frac{1}{2\sqrt{3}}$ (direct computation), and that $\sqrt{\gamma_k} \geq \frac{\lambda_1(\mathcal{B}_2^k, \mathbb{Z}^k)}{\det(\mathbb{Z}^k)^{\frac{1}{n}}} = 1$. For $k = 1$, we note that $M = \mathbb{Z}\mathbf{v}$ for some $\mathbf{y} \in \Lambda \setminus \{\mathbf{0}\}$ where

$$\tau_1(K,\Lambda) = \frac{\det(\mathbb{Z}\mathbf{y})}{\mathrm{vol}_1(K \cap \mathbb{R}\mathbf{y})} = \frac{\|\mathbf{y}\|_2}{\mathrm{vol}_1(K \cap \mathbb{R}\mathbf{y})} = \frac{1}{2}\|\mathbf{y}\|_K.$$

Since M achieves the minimum value $\tau_1(K,\Lambda)$, we see that M is necessarily spanned by a shortest nonzero vector of the lattice, and hence $\lambda_1(K,M) = \lambda_1(K,\Lambda)$ as needed. For $k = n$, we clearly have that $M = \Lambda$, and hence trivially $\lambda_1(K,M) = \lambda_1(K,\Lambda)$.

Now assume that $2 \leq k \leq n-1$. Let $N \subseteq M$ be a $k-1$ dimensional sublattice satisfying $\frac{\det(N)}{\mathrm{vol}_{k-1}(K \cap W_N)} = \tau_{k-1}(M)$. Take $\mathbf{y} \in \Lambda \setminus \{\mathbf{0}\}$ satisfying $\|\mathbf{y}\|_K = \lambda_1(K,\Lambda)$. By the same reasoning as described in Lemma 4.14, we have that

$$\|\mathbf{y}\|_K \geq \|\mathbf{y}\|_{\pi_{W_N^\perp}(K)} \geq \frac{1}{k}\|\mathbf{y}\|_{K_1^{W_N^\perp}} \geq \frac{1}{k}\lambda_1(K_1^{W_N^\perp}, \pi_{W_N^\perp}(\Lambda)) \geq \frac{1}{\sqrt{3}k\lambda_1(E_M^*, M^*)}$$

From Equation 4.18 in Lemma 4.17, we have that

$$\frac{1}{\lambda_1(E_M^*, M^*)} \geq \frac{1}{L_k\sqrt{\gamma_k}}\frac{\det(M)^{\frac{1}{k}}}{\mathrm{vol}_k(K_M)^{\frac{1}{k}}}$$

By Minkowski's first theorem, we have that

$$\frac{\det(M)^{\frac{1}{k}}}{\mathrm{vol}_k(K_M)^{\frac{1}{k}}} \geq \frac{1}{2}\lambda_1(K,M)$$

Putting everything together we get that

$$\lambda_1(K,\Lambda) \geq \frac{1}{k}\lambda_1(K_1^{W_N^\perp}, \pi_{W_N^\perp}(\Lambda)) \geq \frac{1}{\sqrt{3}k\lambda_1(E_M^*, M^*)} \geq \frac{1}{\sqrt{3}L_k k\sqrt{\gamma_k}} \frac{\det(M)^{\frac{1}{k}}}{\mathrm{vol}_k(K_M)^{\frac{1}{k}}}$$
$$\geq \frac{1}{2\sqrt{3}L_k k\sqrt{\gamma_k}}\lambda_1(K,M)$$

as needed. For the asymptotic bound, as in Lemma 4.17, we use the fact that $\sqrt{\gamma_k}L_k = O(k^{\frac{3}{4}})$. □

5 Applications, Research Directions and Open Problems

In this section we give a more detailed description of possible applications and motivating problems for the algorithms presented in this paper. We remark that the examples presented in the following subsections are more directions for further research than fully developed applications. In fact, k-DSP and k-SPP are only part of the solution to these problems, and there are several open problems yet to be solved to obtain results along the lines suggested here.

5.1 Rankin constant

Rankin constants are fundamental parameters in the geometry of numbers introduced in [Ran55] as a generalization of Hermite's constants γ_n. We recall that for any $k \leq n$ and n-dimensional lattice Λ, the constant $\gamma_{n,k}(\Lambda)$ is defined as

$$\gamma_{n,k}(\Lambda) = \left(\min \frac{\det([\mathbf{b}_1,\ldots,\mathbf{b}_k])}{\det(\Lambda)^{k/n}}\right)^2$$

where the minimum is computed over all linearly independent lattice vectors $\mathbf{b}_1,\ldots,\mathbf{b}_k \in \Lambda$. The Rankin constant $\gamma_{n,k}$ is defined as the maximum of $\gamma_{n,k}(\Lambda)$ over all n-dimensional lattices Λ.

As a special case, for $k = 1$, we obtain Hermite's constants $\gamma_{n,1}(\Lambda) = \gamma_n(\Lambda) = (\lambda_1(\Lambda)/\det(\Lambda)^{1/n})^2$ and $\gamma_{n,1} = \gamma_n$. It follows by duality that $\gamma_{n,k} = \gamma_{n,n-k}$, and therefore we also have $\gamma_{n,n-1} = \gamma_n$.

Till recently, the only known value of Rankin's constant (beside those implied by the known values of Hermite's constant γ_n, and the trivial $\gamma_{n,n} = 1$) was $\gamma_{4,2}$. A few other values $\gamma_{6,2} = \gamma_{6,4}, \gamma_{8,2} = \gamma_{8,6}, \gamma_{8,3} = \gamma_{8,5}, \gamma_{8,4}$ were recently determined in [SWO10]. Table 1 shows all the currently known values of $\gamma_{n,k}$. As the table shows, Rankin's constants $\gamma_{n,k}$ are known only for a handful of values of the parameters n, k. A natural application of our algorithm is the study Rankin constant. In particular, for any $k \leq n$ and rank n lattice Λ, our algorithm allows to compute $\gamma_{n,k}(\Lambda)$, and provide a lower bound on $\gamma_{n,k}$.

Our algorithm also allows to study the average value of $\gamma_{n,k}(\Lambda)$ when Λ is chosen at random according to some specified distribution over rank n lattices. The significance of determining the average value of Rankin's constant for randomly chosen lattices will be explained in the next subsection.

5.2 Block reduction algorithms

In [GHGKN06] it is shown that Rankin constants play a role in the analysis of Schnorr's block Korkine-Zolotarev reduction [Sch87], similar to the role played by Hermite constant γ_2 in the analysis of the LLL algorithm [LLL82]. Specifically, the LLL algorithm approximates the shortest vector problem in n-dimensional lattices within a worst case factor $\gamma(n) \approx \gamma_{2,1}^{n/2}$. In [GHGKN06] it is observed that Schnorr's

agorithm [Sch87] is based on the (approximate) solution of the "smallest volume problem" (which is the same as our DSP) for parameters $(n, k) = (2k, k)$. However, no efficient algorithm is given in [GHGKN06] to solve this problem, and that paper resorts to an approximation algorithm, called "transference reduction", which results in poorer SVP approximation factors than the optimal $\gamma(n) \approx \gamma_{2k,k}^{n/(2k)}$. Our algorithm allows to provably achieve SVP approximation $\gamma(n) \approx \gamma_{2k,k}^{n/(2k)}$, though at a high computational price due to the $k^{O(k^2)}$ running time of our algorithm for $n = 2k$.

Clearly, any comparison between block reduction algorithms should take into account time-approximation trade-offs. So, evaluating the practicality of a block reduction algorithm based on our DSP solution would require optimized implementations and a careful study of the parameter space which is beyond the scope of this paper. Still, developing algorithms based on Rankin constants seems an interesting research direction, which had seen no progress since [GHGKN06] due to the lack of algorithms to solve the underlying Rankin problem. We remark that the LLL algorithm typically achieves much better approximation factors than the worst-case $\gamma(n) \approx \gamma_{2,1}^{n/2}$ [NS06]. The exceptional performance of LLL in practice is due to the fact that the 2-dimensional lattices that arise when applying LLL on an n-dimensional lattice are somehow random, and typically achieve a much smaller value of $\gamma_2(\Lambda_2) \leq \gamma_2 = \gamma_{2,1}$. Similarly, one can expect that block reduction algorithms based on the exact solution of k-DSP are likely to perform much better in practice than the worst-case bound based on Rankin constant $\gamma_{n,k}$.

We suggest a further generalization of block reduction algorithms that encompasses as special cases both the Rankin reduction algorithm of [GHGKN06] and the newer sliding reduction algorithm of [GN08] based on Mordell's inequality. The idea is to partition the basis vectors of a lattice $\mathbf{b}_1, \ldots, \mathbf{b}_n$ into blocks of alternating size $k, k', k, k', \ldots$. Then, for each pair of consecutive blocks, solve the corresponding k-DSP or k'-DSP on the $k + k'$ dimensional (projected) sublattice defined by the blocks. Notice that by duality, k-DSP and k'-DSP are equivalent problems (under the ℓ_2-norm), so we only need to be able to solve $\min(k, k')$-DSP. The algorithm may operate on all pairs of consecutive blocks iteratively, just like in LLL and other block reduction algorithms. Many previous reduction algorithms can be seen as special cases of this general framework:

- The LLL algorithm corresponds to setting $k = k' = 1$
- The Rankin algorithm of [GHGKN06] corresponds to larger values of $k = k'$
- Interestingly, using duality, also the sliding algorithm of [GN08] (which is the algorithm currently providing the best asymptotic worst-case results) can be shown to be a special case of our general algorithm where $k = 1$ and k' is arbitrary.

The DSP algorithm proposed in this paper allows to implement a general block reduction algorithm for arbitrary k, k'. We ask: what values of k, k' lead to the best approximation guarantees (possibly for a fixed value of the block size $m = k + k'$? At what running time cost? Previous work [GN08] suggests that using $(k, k') = (1, m - 1)$ is better than the parameters $(k, k') = (m/2, m/2)$ underlying the algorithm of [GHGKN06]. But this may be due to the fact that [GHGKN06] didn't have a method to solve the half volume problem. Also, we see no reason why $(1, m - 1)$-reduction should give better results than, say, $(2, m - 2)$-reduction. Our algorithm allows to explore the entire range of parameters and perform a comprehensive study of generalized (k, k')-block reduction.

5.3 Integer Programming

The Integer Linear Programming (ILP) problem is: given a linear system $A\mathbf{x} \leq \mathbf{b}$ in $\mathbb{R}^n$, decide whether there exists an integer solution $\mathbf{x} \in \mathbb{Z}^n$ to the system. More generally, we examine the convex Integer Non Linear Programming (INLP) problem, where given a convex body K and lattice Λ in $\mathbb{R}^n$ we must decide whether $K \cap \Lambda \neq \emptyset$.

In [Len83], Lenstra gave a $2^{O(n^3)}$-time[4] algorithm for ILP (which easily generalizes to INLP), and provided a general framework for solving integer programs using techniques from the geometry of numbers. Using an algorithmic version Kinchine's Flatness Theorem [Kin48], Lenstra showed that any n-dimensional IP can be reduced $2^{O(n^2)}$ IPs of dimension $n-1$, yielding a recursive solution procedure. To perform the reduction, Lenstra chooses a "flatness direction" $\mathbf{y} \in \Lambda^* \setminus \{\mathbf{0}\}$ of K, with which he slices the feasible region K into consecutive parallel hyperplanes (containing all the lattice points in K) of the form $\{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{y}, \mathbf{x} \rangle = t\} \cap K$ for $t \in \mathbb{Z}$. The main challenge in this method is finding a dual vector $\mathbf{y}$ resulting in the smallest number of intersecting hyperplanes, which Lenstra noticed exactly corresponds to solving the SVP on Λ^* with respect to the width norm of K and which he solves approximately. In [Kan87], Kannan generalized and improved Lenstra's approach giving a $O(n^{2.5})^n$-time IP algorithm. Here, Kannan proposes a higher dimensional reduction procedure, where he uses an exact SVP solver to build an entire "short" basis (HKZ basis) for the lattice to reduce an n-dimensional IP to $O(n^{2.5})^r$ IPs of dimension $n-r$, for some $r \in [n]$ (Lenstra's approach fixes $r = 1$).

Though later work on IP focused almost exclusively on improving the flatness bounds achievable via Lenstra's framework [Ban93b, BLPS99, HK10, DPV11], Kannan's work suggests that higher dimensional reduction may in fact be more efficient. Moreover, in [KL88], Kannan and Lovasz implicitly study the inherent efficiency of higher dimensional reduction, where they prove the following:

Theorem 5.1 (Higher Dimensional Flatness)**.** *Let $K, \Lambda \subseteq \mathbb{R}^n$ be a convex body and lattice. Then*

$$\mu(K,\Lambda) \inf_{W \subseteq \Lambda^*} \left(\frac{\mathrm{vol}(\pi_W(K))}{\det(\pi_W(\Lambda))} \right)^{\frac{1}{\dim(\mathrm{span}(W))}} \leq n \tag{5.1}$$

where π_W denotes orthogonal projection onto $\mathrm{span}(W)$.

We claim that the best subspace W to use for higher dimensional reduction exactly correponds to the minimizing subspace in the above expression. Note that this subspace can be found by n computations of k-SPP for $k \in [n]$. Given W, the reduction procedure computes the set $S = \pi_W(K) \cap \pi_W(\Lambda)$ (which is easier than $K \cap \Lambda$), and for each $s \in S$, recursively solves the $n - \dim(W)$ dimensional IP with respect to $(s + W^{\perp}) \cap K$ and Λ. As we will show in the next paragraph, the size of $|S|^{\frac{1}{\dim(W)}}$ (which essentially determines the reduction complexity) can be bounded in terms of the volumetric estimate above. From the perspective of Kannan and Lovasz [KL88], the above theorem is an inhomogeneous generalization of Minkowski's first theorem, which demonstrates the existence of a non-zero lattice point based on purely volumetric assumptions. Importantly, they conjecture that the upper bound of n in (5.1) is not tight, and show only that a valid upper bound must be $\Omega(\log n)$. For the standard flatness theorem (which is expressed by forcing $\dim(W) = 1$ in (5.1), see [KL88]) on the other hand, it is known that the the upper bound is $\Omega(n)$. While the latter statement makes it unlikely that one can speed up IP beyond $n^{O(n)}$-time using only Lenstra's approach, the former statement leaves open the possibility for an $O(\log n)^n$-time IP algorithms based solely on higher dimensional reduction. We justify this assertion in the next paragraph.

[4]Here we restrict our attention to the dependence on n, where we note that the remaining dependencies are polynomial.

$n=1$	1								
$n=2$	$\frac{2}{\sqrt{3}}$	1							
$n=3$	$\sqrt[3]{2}$	$\sqrt[3]{2}$	1						
$n=4$	$\sqrt{2}$	$3/2$	$\sqrt{2}$	1					
$n=5$	$2^{3/5}$			$2^{3/5}$	1				
$n=6$	$\sqrt[6]{64/3}$	$3^{2/3}$		$3^{2/3}$	$\sqrt[6]{64/3}$	1			
$n=7$	$2^{6/7}$					$2^{6/7}$	1		
$n=8$	2	3	4	4	4	3	2	1	
$n=24$	4							4	1
$\gamma_{n,k}$	$k=1$	$k=2$	$k=3$	$k=4$	$k=5$	$k=6$	$k=7$	$k=23$	$k=24$

Table 1: Known values of Rankin constants $\gamma_{n,k}$.

Volume Bound on the number of Lattice Points We will first show that the number of lattices points in K acts like volume once $\mu(K,\Lambda) \leq 1$. Let $G(K,\Lambda) = \max_{\mathbf{x}\in\mathbb{R}^n} |(K+\mathbf{x})\cap\Lambda| = \max_{\mathbf{x}\in\mathbb{R}^n} |K\cap(\Lambda+\mathbf{x})|$ be the maximum number of lattice points in a shifted copy of K.

Lemma 5.2. *Let $K \subseteq \mathbb{R}^n$ be a convex body and let Λ be an n-dimensional lattice. Then*

$$\frac{\mathrm{vol}(K)}{\det(\Lambda)} \leq G(K,\Lambda) \leq \max\{1,\mu(K,\Lambda)^n\}2^n\frac{\mathrm{vol}(K)}{\det(\Lambda)}$$

Proof. Let F denote any fundamental region of Λ. Then note that

$$\mathrm{vol}(K) = \int_F |(\Lambda+\mathbf{x})\cap K| d\mathbf{x}$$

Since $\mathrm{vol}(F) = \det(\Lambda)$, there must exist $\mathbf{x}\in F$ such that $|(\Lambda+\mathbf{x})\cap K| \geq \frac{\mathrm{vol}(K)}{\det(\Lambda)}$ as needed.

By shifting K, we may assume that $\mathbf{0}\in K$. Now let $\leq_{lex}$ denote the standard lexicographic ordering on $\mathbb{R}^n$. For K,Λ and $\mathbf{x}\in\mathbb{R}^n$ define $\mathrm{CVP}(K,\Lambda,\mathbf{x}) = \arg\min_{\mathbf{y}\in\Lambda}\|\mathbf{y}-\mathbf{x}\|_K$. Define the Voronoi cell of Λ with respect to K as the set $\mathcal{V}(\Lambda,K)$ of all points $\mathbf{x}\in\mathbb{R}^n$ such that $\mathbf{0}$ is the lexicographically first element of $\mathrm{CVP}(-K,\Lambda,\mathbf{x})$, i.e., $\mathbf{0}\in\mathrm{CVP}(-K,\Lambda,\mathbf{x})$ and $\forall\,\mathbf{y}\in\mathrm{CVP}(-K,\Lambda,\mathbf{x})$. $\mathbf{0}\leq_{lex}\mathbf{y}$. We claim that $\mathcal{V}(\Lambda,K)$ is a fundamental region of Λ, i.e. $\mathcal{V}(\Lambda,K)$ induces an Λ-periodic tiling of $\mathbb{R}^n$. Take $\mathbf{x},\mathbf{y}\in\Lambda$, $\mathbf{x}\neq\mathbf{y}$. We must show that $\mathbf{x}+\mathcal{V}(\Lambda,K)\cap\mathbf{y}+\mathcal{V}(\Lambda,K)=\emptyset$. Assume not and pick $\mathbf{z}\in\mathbf{x}+\mathcal{V}(\Lambda,K)\cap\mathbf{y}+\mathcal{V}(\Lambda,K)$. Then we see that $\mathbf{x},\mathbf{y}\in\mathrm{CVP}(-K,\Lambda,\mathbf{z})$, but then $\mathbf{x}\leq_{lex}\mathbf{y}$ and $\mathbf{y}\leq_{lex}\mathbf{x}$, a contradiction. Now take $\mathbf{z}\in\mathbb{R}^n$. Let $\mathbf{x}$ denote the lex least element of $\mathrm{CVP}(-K,\Lambda,\mathbf{z})$. Then clearly $\mathbf{z}\in\mathbf{x}+\mathcal{V}(\Lambda,K)$ as needed.

By the definition of $\mathcal{V}(\Lambda,K)$ we see that $\mu(K,\Lambda) = \sup_{\mathbf{x}\in\mathcal{V}(\Lambda,K)}\|\mathbf{x}\|_{-K}$. Therefore for $\mathbf{x}\in\mathcal{V}(\Lambda,K)$, we have that $0\in\mathbf{x}-\mu(K,\Lambda)K \Rightarrow \mathbf{x}\in\mu(K,\Lambda)K$, hence $\mathcal{V}(\Lambda,K)\subseteq\mu(K,\Lambda)K$.

Now we note that it suffices to prove the theorem when $\mu(K,\Lambda)\leq 1$, i.e. when $\mathcal{V}(K,\Lambda)\subseteq K$. Take $\mathbf{x}\in\mathbb{R}^n$. We wish to bound $|K\cap(\Lambda+\mathbf{x})|$. Since $\mathcal{V}(K,\Lambda)\subseteq K$, we note that for $\mathbf{y}\in K\cap(\Lambda+\mathbf{x})$, we have that

$$\mathbf{y}+\mathcal{V}(K,\Lambda)\subseteq K+\mathcal{V}(K,\Lambda)\subseteq K+K=2K$$

Furthermore, since $\mathcal{V}(K,\Lambda)$ is a fundamental region of Λ, we have that $\mathrm{vol}(\mathcal{V}(K,\Lambda)) = \det(\Lambda)$, and so

$$\begin{aligned}\mathrm{vol}(2K) &\geq \mathrm{vol}\left(\bigcup_{\mathbf{y}\in(\Lambda+\mathbf{x})\cap K} \mathbf{y}+\mathcal{V}(K,\Lambda)\right)\\ &= |(\Lambda+\mathbf{x})\cap K|\,\mathrm{vol}(\mathcal{V}(K,\Lambda)) = |(\Lambda+\mathbf{x})\cap K|\det(\Lambda)\end{aligned}$$

Hence $|(\Lambda+\mathbf{x})\cap K| \leq \frac{\mathrm{vol}(2K)}{\det(\Lambda)} = 2^n\frac{\mathrm{vol}(K)}{\det(\Lambda)}$. □

The above lemma helps justify the assertion that Theorem 5.1 is relevant to higher dimensional reduction techniques for IP. To recap, when solving IP with respect to K and Λ, higher dimensional reduction seeks to find a subspace W for which $|\pi_W(K)\cap\pi_W(\Lambda)|^{\frac{1}{\dim(W)}}$ is as small as possible, where each point in $\pi_W(K)\cap\pi_W(\Lambda)$ represents a lower dimensional subproblem to be solved recursively. Since $\mu(\pi_W(K),\pi_W(\Lambda)) \leq \mu(K,\Lambda)$, the above Lemma combined with Theorem 5.1 shows that for minimizing projection W, assuming $\mu(K,\Lambda) \geq \frac{1}{2}$ (which will be the case in IP after some initial preprocessing), that $G(\pi_W(K),\pi_W(\Lambda)) \leq (2n)^{\dim(W)}$. Therefore, using the DPV lattice point enumerator (see Theorem 5.2.6 of [Dad12]) we can enumerate $\pi_W(K)\cap\pi_W(\Lambda)$ in $O(n)^{\dim(W)}$-time. Furthermore, any improvement to the bound of Theorem 5.1 would lead to a corresponding improvement.

The main challenges associated with using higher dimensional reduction for IP here are two fold. Firstly, we need an algorithm to compute the minimizing subspace W. As noted in Section 5, the problem of finding W easily reduces to a sequence of k-SPPs (and hence by duality to k-DSP) for $k \in [n]$. However, the $k^{O(kn)}$ algorithm for $k-SDP$ is currently too expensive too run for large k without surpassing the complexity of IP. Second, we need to develop a stronger bound for the right hand side of Theorem 5.1 than the current $O(n)$ bound to make provable progress in the theory of IP.

5.4 Closest Vector Problem with Preprocessing

The closest vector problem (CVP) asks, given a lattice Λ and a target point $\mathbf{t}$, to find the lattice point closest to $\mathbf{t}$. As usual, the problem can be defined with respect to any norm, but the Euclidean norm is the most common, and the one we focus on here. The CVP with preprocessing (CVPP) is a variant of CVP where the lattice Λ can be arbitrarily preprocessed, and it is particularly relevant in applications where the lattice is known long in advance, or it is used with several target vectors $\mathbf{t}$. In this section we describe a potential application of our DSP algorithm to the solution of CVPP.

Let $\mathbf{B} = [\mathbf{b}_1,\ldots,\mathbf{b}_n]$ be a basis of an n-dimensional lattice. As usual, for any $i = 1,\ldots,n$, let π_i be the projection onto the orthogonal complement of $\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}$. The vectors $\mathbf{b}_i^* = \pi_i(\mathbf{b}_i)$ are the Gram-Schmidt orthogonalization of the basis vectors and can be used to measure the quality of a lattice basis in various ways, e.g., when used in conjunction with enumeration algorithms for the solution of the closest vector problem.

Enumeration based solutions to CVP work by first preprocessing the lattice to produce a reduced basis, and then use the resulting basis to find the lattice point closest to a given target using branch-and-bound methods. (See below for details.) The enumeration algorithm of Kannan [Kan87] solves CVP in time $n^{O(n)}$ and polynomial space $n^{O(1)}$ by first computing a Hermite-Korkine-Zolotarev (HKZ) reduced basis. Recently, an asymptotically faster CVP algorithm with single exponential running time $2^{O(n)}$ (for the Euclidean norm) has been discovered by Micciancio and Voulgaris [MV10b]. However, the algorithm of [MV10b] has the drawback of using also exponential space $2^{O(n)}$ to store the Voronoi cell of the lattice. Moreover, implementations of enumeration algorithms have been highly optimized, and run pretty well in practice. As a

result, the algorithm of [MV10b], while asymptotically faster than [Kan87], is not competitive for practical values of the dimension.[5]

For these reasons, it is an interesting open problem to determine if the running time of enumeration algorithms can be improved by using a different preprocessing of the basis. In this context, it is natural to consider the lattice as fixed, and disregard the time spent during the preprocessing (basis reduction) stage. This is justified by the fact that many applications require the solution of CVP for the same lattice and several target vectors $\mathbf{t}$, so the time spent during the preprocessing of the basis can be amortized over the solution of several CVP instances.

We show how our DSP algorithm can be used to find bases that are potentially much better than the HKZ reduced bases currently used in state-of-the-art lattice enumeration algorithms. First, let us recall how enumeration algorithms work. Let $\mathbf{B}$ be a (reduced) lattice basis and $\mathbf{t}$ a target vector. On input $\mathbf{B}$ and $\mathbf{t}$, enumeration algorithms produce a list containing all lattice vectors within a certain distance r from $\mathbf{t}$, and select, among these vectors, the one closest to the target. If the distance of $\mathbf{t}$ from the lattice is less than r (in particular, if $r = \mu(\Lambda)$ equals the covering radius of the lattice) then the algorithm solves the closest vector problem.

Enumeration algorithms work by enumerating not only lattice points, but also points in the projections $\pi_i(\mathcal{L}(\mathbf{B}))$ of the input lattice, and for any i, they enumerate the lattice points

$$L_i = \pi_i(\mathcal{L}(\mathbf{B})) \cap \pi_i(\mathbf{t} + \mathcal{B}_2 \cdot r) \supseteq \pi_i(\mathcal{L}(\mathbf{B}) \cap (\mathbf{t} + \mathcal{B}_2 \cdot r))$$

within distance r from the projected target $\pi_i(\mathbf{t})$. These projected lattice points can be naturally arranged into a tree, where L_i contains the nodes at level $n - i + 1$, and the parent of node $\pi_i(\mathbf{v})$ is $\pi_{i+1}(\mathbf{v})$. All the lattice points within distance r from $\mathbf{t}$ can then be enumerated using only polynomial space performing a depth first tree traversal, which has running time proportional to the number of nodes in the tree, and can be analyzed considering the number of nodes at each level L_i. When solving CVP (i.e. when $r = \mu(\Lambda)$), this number can be bounded volumetrically by the quantity $\mathrm{vol}(\mathcal{B} \cdot r)/\det(\Lambda)$ as shown in Lemma 5.2. Furthermore, this estimate is essentially tight, as justified by the Gaussian heuristic. Recall that the Gaussian heuristic states that the number of lattice points in a sufficiently large ball $\mathcal{B} \cdot r + \mathbf{t}$ around the target is roughly equal to $\mathrm{vol}(\mathcal{B} \cdot r)/\det(\Lambda)$. The heuristic is justified by the fact that for any given lattice Λ, the expected number of lattice points in such a ball $\mathcal{B} + \mathbf{t}$ when the target $\mathbf{t}$ is chosen uniformly at random (e.g., within a fundamental region of the lattice) is precisely $\mathrm{vol}(\mathcal{B})/\det(\Lambda)$. Using this estimate then, we get that the number of points within each level of the tree is

$$|L_{n-k+1}| \quad \approx \quad \mathrm{vol}(\pi_{n-k+1}(\mathcal{B}_2 \cdot \rho))/\det(\pi_{n-k+1}(\mathcal{L}(\mathbf{B}))).$$

We note that these ratios are exactly the quantites examined in Theorem 5.1. Specializing to the Euclidean norm, the best subspace for us to perform a projection is the minimizer of the following expression:

$$\Phi(\Lambda) = \min_{W \subseteq \Lambda^*} \left(\frac{\mathrm{vol}(\pi_W(\mathcal{B}_2 \cdot \rho))}{\det(\pi_W(\mathcal{L}(\mathbf{B})))} \right)^{1/\dim(W)}$$

where W ranges over all sublattices of Λ^* of arbitrary dimension $1 \leq \dim(W) \leq n$. The best currently known lower and upper bounds on this quantity [KL88] are

$$\sqrt{\log n} \leq \max_{\Lambda} \Phi(\Lambda) \leq \sqrt{n}.$$

[5] Clearly, due to the different asymptotics, for sufficiently large n, the algorithm of [MV10b] will be faster than [Kan87]. However, the crossover point is high enough that for those values of n, neither [MV10b] nor [Kan87] can be run in practice due to their enormous (exponential) running time.

We remark that the upper bound $\sqrt{n}$ is believed to be far from the true answer. We can make the heuristics assumption that the lower bound $\max_\Lambda \Phi(\Lambda) = \sqrt{\log n}$ is the correct value. The assumption is justified by the fact that finding lattices for which $\Phi(\Lambda) > \sqrt{\log n}$ would yield better lower bounds on $\max_\Lambda \Phi(\Lambda)$ than currently known. So, either such lattices do not exist, or at least they are hard to find, and arguably unlikely to occur in practice. Similar heuristic assumptions (concerning the so-called *kissing number*, another fundamental constant in the study of lattice packings) have been considered in the study of algorithms for the shortest vector problem [MV10a], and proved to yield very accurate predictions of the (empirical) performance of heuristic sieving algorithms for SVP even in moderately high dimension $n < 80$.

Assuming $\Phi(\Lambda) \leq \sqrt{\log n}$, we get that there is a level k of the tree and a basis $\mathbf{B}$ such that $|L_{n-k+1}| \leq (\log n)^{k/2}$. This allows to reduce a CVP instance in dimension n to $(\log n)^{k/2}$ instances in dimension $n-k$, by finding all points L_{n-k+1} in the projected lattice within distance ρ from the (projected) target, lift them to points in the original lattice, and then solving the corresponding CVP instance with lattice $\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_{n-k})$. This leads to an enumeration algorithm that runs in time $(\log n)^{n/2}$, while using polynomial space, provided we can solve the following problems:

- Find a k, and basis $\mathbf{B}$ such that $|L_{n-k+1}| \leq (\log n)^{k/2}$, and
- enumerate all the points in L_{n-k+1} in time $|L_{n-k+1}| \leq (\log n)^{k/2}$

Notice that the running time $(\log n)^{n/2}$ is not as small as the running time $2^{O(n)}$ achieved by [MV10b] using exponential space, but it is much smaller than the running time $n^{O(n)}$ of the best current (polynomial space) enumeration methods.

We remark that our DSP algorithm can be used to solve the first of the two subproblems described above. More specifically, for any k, we have that

$$\begin{aligned} |L_{n-k+1}| &\approx \frac{\mathrm{vol}_k(\mathcal{B}_2^k \rho)}{\det(\pi_{n-k+1}(\mathcal{L}(\mathbf{B})))} \\ &= \frac{\mathrm{vol}_k(\mathcal{B}_2^k) \cdot \rho^k}{\det(\mathcal{L}(\mathbf{B}))} \cdot \det(\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_{n-k})) \end{aligned}$$

is proportional to $\det(\mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_{n-k}))$. So, for any fixed k, minimizing $|L_{n-k+1}|$ is equivalent to solving $(n-k)$-DSP. We can solve $(n-k)$-DSP for all $k = 1, \ldots, n$, compute the corresponding quantity $|L_{n-k+1}|^{1/k}$, and select the value of k and associated basis $\mathbf{B}$ for which $|L_{n-k+1}|^{1/k} = \min_k |L_{n-k+1}|^{1/k} = \Phi(\Lambda) \leq \sqrt{\log n}$.

These leaves us with the problem of efficiently enumerating all the points in L_{n-k+1}. If $n-k+1$ is sufficiently small (e.g., logarithmic in n) this can be done using the methods of [MV10b, DPV11], which use space $\exp(k) = n^{O(1)}$. For larger k, the best current polynomial space algorithms are again enumeration methods, but they have running time $k^{O(k)}$. Here too our DSP algorithm can be useful. For simplicity, let us consider the extreme case where the quantity $|L_{n-k+1}|^{1/k}$ is minimized for $k = n$. The problem is that of enumerating all lattice points within distance $\rho(\Lambda)$ from a target point $\mathbf{t}$, for a lattice Λ such that the covering radius $\rho(\Lambda) \leq \sqrt{\log n} \cdot \det(\Lambda)^{1/n} / \mathrm{vol}(\mathcal{B}_2^n)^{1/n} = O(\sqrt{n \log n}) \det(\Lambda)^{1/n}$ is within a poly-logarithmic factor from Minkowski's bound. These are very dense lattices, where number of lattice points in a covering ball centered around an arbitrary target vector $\mathbf{t}$ can be bounded from above by a function of the form $(\log n)^{O(n)}$, rather than $n^{O(n)}$. When applying enumeration algorithms to this lattice, we obtain a tree such that the bottom layer has size $|L_1| = (\log n)^{O(n)}$. The problem is that the size of intermediate layers (e.g., $L_{n/2}$) can be much larger than $(\log n)^{O(n)}$. Using HKZ reduction, one can obtain a basis such that $|L_{n/2}| = n^{O(n)}$. A better basis, for which $|L_{n/2}|$ is as small as possible, can be found using our DSP algorithm. Finally, we

remark that the number of nodes at level $L_{n/2}$ for which the full enumeration tree contains a node at level L_1 (i.e., lead to lattice point within distance $\rho(\Lambda)$ from the target) is at most $(\log n)^{n/2}$, just because L_1 is small. So, if we could somehow determine, given a node, if the corresponding subtree leads to a leaf at leverl L_1, then we could prune the enumeration tree, and make it run in time $(\log n)^{n/2}$. No theoretical solution is known to this problem, but this is a task that may be addressed by pruning methods extensively used to speed up enumeration algorithms.

References

[AJ08] V. Arvind and P. S. Joglekar. Some sieving algorithms for lattice problems. In *FSTTCS*, pages 25–36. 2008.

[Ajt96] M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.

[Ajt98] M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19. 1998.

[AKS01] M. Ajtai, R. Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *STOC*, pages 601–610. 2001.

[Bal88] K. M. Ball. Logarithmically concave functions and sections of convex sets in rn. *Studia Mathematica*, 88:69–84, 1988.

[Ban93a] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.

[Ban93b] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296:625–635, 1993. ISSN 0025-5831.

[Bla18] W. Blaschke. Über affine geometry xiv: eine minimum aufgabe für legendres trägheits ellipsoid. *Ber. verh. sächs. Akad. d. Wiss.*, 70:72–75, 1918.

[BLPS99] W. Banaszczyk, A. Litvak, A. Pajor, and S. Szarek. The flatness theorem for nonsymmetric convex bodies via the local theory of Banach spaces. *Mathematics of Operations Research*, 24(3):728–750, 1999.

[BM87] J. Bourgain and V. D. Milman. New volume ratio properties for convex symmetric bodies in $\mathbb{R}^n$. *Inventiones Mathematicae*, 88:319–340, 1987.

[BN07] J. Blömer and S. Naewe. Sampling methods for shortest vectors, closest vectors and successive minima. *Theoretical Computer Science*, 410(18):1648–1665, April 2009. Preliminary version in ICALP 2007.

[Bou86] J. Bourgain. On high-dimensional maximal functions associated to convex bodies. *Amer. J. Math*, 108(6):1467–1476, 1986.

[BS99] J. Blömer and J.-P. Seifert. On the complexity of computing short linearly independent vectors and short bases in a lattice. In *STOC*, pages 711–720. 1999.

[Dad12] D. Dadush. *Integer Programming, Lattice Algorithms, and Deterministic Volume Estimation*. Ph.D. thesis, Georgia Institute of Technology, 2012.

[DFK91] M. Dyer, A. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991. ISSN 0004-5411. doi:http://doi.acm.org/10.1145/102782.102783.

[DPV11] D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via m-ellipsoid coverings. In *FOCS*, pages 580–589. 2011.

[DV12] D. Dadush and S. Vempala. Near optimal volume deterministic algorithms and lattice algorithms via m-ellipsoids. Arxiv, Report 1201.5972, 2012. `http://arxiv.org`.

[GHGKN06] N. Gama, N. Howgrave-Graham, H. Koy, and P. Nguyen. Rankin's constant and blockwise lattice reduction. In *Advances in Cryptology – Proceedings of CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 112–130. Springer, 2006.

[GN08] N. Gama and P. Q. Nguyen. Finding short lattice vectors within Mordell's inequality. In *Proc. of STOC*, pages 207–216. 2008.

[GNR10] N. Gama, P. Nguyen, and O. Regev. Lattice enumeration using extreme pruning. In *EUROCRYPT*. Springer, 2010.

[Hen80] D. Hensley. Slicing convex bodies - bounds for slice area in terms of bodies covariance. *Proceedings of the American Mathematical Society*, 79(4):619–625, 1980.

[HK10] R. Hildebrand and M. Köppe. A new Lenstra-type algorithm for quasiconvex polynomial integer minimization with complexity $2^{O(n \log n)}$. Arxiv, Report 1006.4661, 2010. `http://arxiv.org`.

[HR07] I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory of Computing*, 8(1):513–531, 2012. Preliminary version in STOC 2007.

[Kan87] R. Kannan. Minkowski's convex body theorem and integer programming. *Mathematics of operations research*, 12(3):415–440, August 1987.

[Kho03] S. Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005. Preliminary version in FOCS 2003.

[Kin48] A. Kinchine. A quantitative formulation of Kronecker's theory of approximation. *Izv. Acad. Nauk SSSR*, 12:113–122, 1948.

[KL88] R. Kannan and L. Lovász. Covering minima and lattice point free convex bodies. *Annals of Mathematics*, 128:577–602, 1988.

[KLS95] R. Kannan, L. Lovász, and M. Simonovits. Isoperimetric problems for convex bodies and a localization lemma. *Discrete & Computational Geometry*, 13:541–559, 1995.

[Kup08] G. Kuperberg. From the mahler conjecture to gauss linking integrals. *Geometric And Functional Analysis*, 18:870–892, 2008.

[Len83] H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, November 1983.

[LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.

[MG02] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, 2002.

[Mic98] D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035, 2000. Preliminary version in FOCS 1998.

[Mic11] D. Micciancio. How-many-lll-reduced-bases-are-there? answer 1. MathOverflow http://mathoverflow.net, June 2011. URL (accessed on 2011-10-23) http://mathoverflow.net/questions/57021/how-many-lll-reduced-bases-are-there.

[Mic12] D. Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory of Computing*, 8(1):487–512, 2012.

[Min10] H. Minkowski. *Geometrie Der Zahlen*. Leipzig and Berlin: R. G. Teubner, 1910.

[MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.

[MV10a] D. Micciancio and P. Voulgaris. Faster exponential time algorithms for the shortest vector problem. In *Proc. of SODA*, pages 1468–1480. 2010.

[MV10b] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. *SIAM J. on Computing*, 2012. To appear. Preliminary version in STOC 2010.

[Naz09] F. Nazarov. The Hörmander proof of the Bourgain-Milman theorem, 2009. Preprint.

[NS06] P. Nguyen and D. Stehlé. LLL on the average. In *Proc. of ANTS*, pages 238–256. 2006.

[NV08] P. Q. Nguyen and T. Vidick. Sieve algorithms for the shortest vector problem are practical. *Journal of Mathematical Cryptology*, 2(2):181–207, July 2008.

[PS09] X. Pujol and D. Stehlé. Solving the shortest lattice vector problem in time $2^{2.465n}$. Report 2009/605, IACR ePrint archive, December 2009.

[Ran55] R. A. Rankin. The closest packing of spherical caps in n dimensions. In *Proceedings of the Glasgow mathematical association*, volume 2, pages 139–144. Oliver and Boyd, 1955.

[Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in STOC 2005.

[RS57] C. Rogers and G. Shephard. The difference body of a convex body. *Arch. Math.*, 8:220–233, 1957.

[San49] L. A. Santaló. Un invariante afin para los cuerpos convexos del espacio de n dimensiones. *Portugalie Math.*, 8:155–161, 1949.

[Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theor. Comput. Sci.*, 53:201–224, 1987.

[SE94] C.-P. Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Mathematical programming*, 66(1-3):181–199, August 1994. Preliminary version in FCT 1991.

[SH95] C.-P. Schnorr and H. H. Hörner. Attacking the Chor-Rivest cryptosystem by improved lattice reduction. In *EUROCRYPT*, pages 1–12. 1995.

[SWO10] K. Sawatani, T. Watanabe, and K. Okuda. A note on the Hermite-Rankin constant. *J. Th. Nombres Bordeaux*, 22:209–217, 2010.

[WLTB11] X. Wang, M. Liu, C. Tian, and J. Bi. Improved Nguyen-Vidick heuristic sieve algorithm for shortest vector problem. In *Proceedings of ASIACCS '11*, pages 1–9. ACM, 2011.