On Approximating the Covering Radius and Finding Dense Lattice Subspaces

Daniel Dadush*
Centrum Wiskunde & Informatica

Abstract

In this work, we give a novel algorithm for computing dense lattice subspaces, a conjecturally tight characterization of the lattice covering radius, and provide a bound on the slicing constant of lattice Voronoi cells. Our work is motivated by the pursuit of faster algorithms for integer programming, for which we give a conditional speedup based on the recent resolution of the ℓ_2 Kannan-Lovász conjecture. Through these results, we hope to motivate further study of the interplay between the recently developed reverse Minkowski theory, lattice algorithms and convex geometry.

On the algorithmic side, our main contribution is a $2^{O(n)}$-time algorithm for computing a $O(C_\eta(n))$-approximate sublattice of minimum normalized determinant on any n-dimensional lattice, where $C_\eta(n) = O(\log n)$ is the reverse Minkowski constant in dimension n. Our method for finding dense lattice subspaces is surprisingly simple: we iteratively descend to a random co-dimension 1 subspace chosen to be the orthogonal space to a discrete Gaussian sample from the dual lattice. Applying this algorithm within a "filtration reduction" scheme, we further show how to compute a $O(C_\eta(n)$-approximate canonical filtration of any lattice, which corresponds to a canonical way of decomposing a lattice into dense blocks. As a primary application, we get the first $2^{O(n)}$-time algorithm for computing a sparse lattice projection whose "volume radius" provides a lower bound on the lattice covering radius that is tight within a $O(\log^{2.5} n)$-factor. This provides an efficient algorithmic version of the ℓ_2 Kannan-Lovász conjecture, which was recently resolved by Regev and Stephens-Davidowitz (STOC '2017).

On the structural side, we prove a new lower bound on the covering radius which combines volumetric lower bounds across a chain of lattice projections. Assuming Bourgain's slicing conjecture restricted to Voronoi cells of stable lattices, our lower bound implies (somewhat surprisingly) that the problem of approximating the lattice covering radius to within a constant factor is in coNP. Complementing this result, we show that the slicing constant of any n-dimensional Voronoi cell is bounded by $O(C_{KL,2}(n)) = O(\log^{1.5} n)$, the ℓ_2 Kannan-Lovász constant, which complements the $O(\log n)$ bound of Regev and Stephens-Davidowitz for stable Voronoi cells.

*Supported by the NWO Veni grant 639.071.510. dadush@cwi.nl

1 Introduction

A k-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$ is the integer span of k linearly independent vectors $\mathbf{b}_1, \ldots, \mathbf{b}_k \in \mathbb{R}^n$. Many of the applications of lattices, be it to cryptography, optimization or coding, revolve around a fine-grained understanding of lattice geometry. The most fundamental geometric lattice parameter is the lattice determinant $\det(\mathcal{L}) := \mathrm{vol}_k(\mathbf{B}[0,1)^k) = \sqrt{\det(\mathbf{B}^\mathsf{T}\mathbf{B})}$, for any basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_k)$ of $\mathcal{L}$, where $1/\det(\mathcal{L})$ asymptotically corresponds the number of lattice points per unit volume for any "large enough" set in $\mathrm{span}(\mathcal{L})$. Relations between basic lattice parameters, such as the length of the shortest-zero vector $\lambda_1(\mathcal{L}) = \min\{\|\mathbf{y}\| : \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}$ or the covering radius $\mu(\mathcal{L}) = \min\{s \geq 0 : \mathrm{span}(\mathcal{L}) = \mathcal{L} + s(\mathbb{B}_2^n \cap \mathrm{span}(\mathcal{L}))\}$, where $\mathbb{B}_2^n$ is the unit Euclidean ball in $\mathbb{R}^n$, have been the object of intense study since the beginning of the study of lattices. For example, finding the densest lattice packing of $\mathbb{R}^n$ by Euclidean balls of radius $1/2$ corresponds to the n-dimensional lattice $\mathcal{L}$ whose shortest non-zero vector has length $\lambda_1(\mathcal{L}) = 1$ and whose determinant is minimum. As another example, Minkowski's first fundamental theorem implies the basic bound $\lambda_1(\mathcal{L}) \leq \sqrt{n}\det(\mathcal{L})^{1/n}$ for any n-dimensional lattice $\mathcal{L}$.

The Flatness Theorem. In the pursuit of a better understanding of lattice geometry, an important research direction has been the search for tight approximate min-max relations between different lattice parameters, known as transference theorems. These duality relations are often most naturally expressed as inequalities between geometric parameters of a lattice and that of its dual. The dual lattice $\mathcal{L}^* := \{\mathbf{y} \in \mathrm{span}(\mathcal{L}) : \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z} \ \ \forall \mathbf{x} \in \mathcal{L}\}$ is the set of vectors having integral inner products with all the vectors in $\mathcal{L}$, which can be thought of the normal vectors of lattice hyperplanes in $\mathcal{L}$. Perhaps the most fundamental facts about duality are that $\mathcal{L}^{**} = \mathcal{L}$ and that $\det(\mathcal{L})\det(\mathcal{L}^*) = 1$ (i.e. volumetric information perfectly "dualizes"). The most famous transference theorem relates the covering radius of $\mathcal{L}$ to the length of the shortest vector of $\mathcal{L}^*$, known as Khinchine's flatness theorem, where the bounds stated below are due to Banaszczyk [Ban93]:

$$\frac{1}{2} \leq \mu(\mathcal{L})\lambda_1(\mathcal{L}^*) \leq n \tag{1}$$

The "algorithmification" of a generalization of the above inequality will be of principal interest to this work, as we elaborate on below. To understand this inequality, we note that the LHS is the "easy side", which corresponds to the fact that $1/\lambda_1(\mathcal{L}^*)$ is precisely the largest possible spacing between any two consecutive and parallel lattice hyperplanes in $\mathcal{L}$. The harder RHS in essence says that any ball of radius $\mu(\mathcal{L})$ intersects at most $O(n)$ of these "maximally spaced" hyperplanes, i.e. from the collection $\{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{z} \rangle = i\}$, $i \in \mathbb{Z}$, where $\mathbf{z}$ is any shortest-nonzero vector of $\mathcal{L}^*$. For the RHS, even proving finiteness of the is highly non-trivial, and indeed the original proof of Khinchine [Khi48] only achieved a bound of $n!$.

From Flatness to Integer Programming. The classical integer programming problem (IP) is given an appropriately represented convex body K (e.g. K is a polytope given in inequality representation) to find an integer point in K or decide $K \cap \mathbb{Z}^n = \emptyset$. A major push to improve the flatness bound was given by the breakthrough work of Lenstra [Len83], who gave $2^{O(n^2)}$-time algorithm for IP based on a (weaker) algorithmic version of the above transference theorem. The main dichotomy that Lenstra exploited is that either K is very "fat", where one can find an integer point by an appropriate rounding procedure (e.g. a closest vector computation), or K is "flat", where then

there exists an integer direction along which K intersects a small number of integer hyperplanes that one recurses on. To derive this dichotomy from (1), one may apply John's theorem to K to get an affine transformation satisfying $\mathbb{B}_2^n \subseteq \mathbf{T}K - \mathbf{c} \subseteq n\mathbb{B}_2^n$ and the flatness theorem to $\mathcal{L} = \mathbf{T}\mathbb{Z}^n$. Being "fat" corresponds to $\mu(\mathcal{L}) \leq 1$, in which case K is guaranteed to contain a point in $\mathbb{Z}^n$, which is found by computing a closest vector in $\mathcal{L}$ to $\mathbf{c}$. In the opposite "flat" case, one easily derives from (1) that K intersects at most $O(n^2)$ integer hyperplanes with normal vector $T^\mathsf{T}\mathbf{z} \in \mathbb{Z}^n$, where $\mathbf{z}$ a shortest non-zero vector of $\mathcal{L}^*$.

The reduction above used as subroutines the classical shortest and closest vector problems (SVP & CVP), the problems of computing the shortest non-zero vector of a lattice given by its basis and the closest vector of a lattice to a given target respectively. Assuming access to SVP & CVP oracles as well as basic subroutines for K, Lenstra's insight was that up to a $\text{poly}(n)$-factor the right hand side of (1) bounds the branching factor of the search tree needed to solve an IP. This discovery lead to intense activity on achieving the tight ℓ_2 transference bounds above [Bab86, KL88, Hs88, LLS90, Ban93], yielding $n^{O(n)}$ estimates on the IP search tree size, as well as the development fast basis reduction algorithms for SVP & CVP [LLL82, Sch87, Kan87], for which [Kan87] gave $n^{O(n)}$-time and $\text{poly}(n)$-space algorithms.

Since that time there has been tremendous progress in our algorithmic knowledge for SVP and CVP with respect to the Euclidean norm, with the development lattice algorithms based on randomized sieving [AKS01, AKS02], Voronoi cell computations [SFS09, MV13], and mostly recently discrete Gaussian sampling [ADRS15, ADS15], where now both problems can be solved exactly in randomized $2^{n+o(n)}$-time and -space [ADRS15, ADS15]. However, similar progress for the IP problem has proved elusive. Despite the aforementioned developments and the many new tools in algorithmic convex geometry, the only improvement to the $n^{O(n)}$ complexity has been to the constant c in the exponent, where the best known bound is $c = 1$ [Dad12].

The Kannan-Lovász Conjecture. A significant stumbling block for most IP algorithms has been the reliance on hyperplane based branching, which morally only takes advantage of flatness one "direction at a time". Indeed, there are simple examples of convex bodies containing no integer points in their interior which intersect $\Omega(n)$ hyperplanes from any family of consecutive parallel integer hyperplanes (e.g. take an n-scaling of the standard simplex), and thus induce a branching factor of at least $\Omega(n)$. While one may hope that such worst-case nodes in the search tree occur only infrequently, it is unclear whether even a sophisticated amortized analysis could yield an averaging branching factor that is subpolynomial in n. As a way to address this issue, Kannan and Lovász [KL88] introduced a "higher dimensional" version of the flatness theorem which provides a way to leverage the power of many flatness directions at once. Since for subpolynomial branching factors, one cannot afford the loss due to ellipsoidal approximation (i.e. John's theorem), we state it in its general convex body version. For a convex set K and lattice $\mathcal{L}$, with $K \subseteq \text{span}(\mathcal{L}) \subseteq \mathbb{R}^n$, we define the covering radius of K with respect to $\mathcal{L}$ by $\mu(K, \mathcal{L}) = \inf\{r > 0 : \mathcal{L} + rK = \text{span}(\mathcal{L})\}$. For any n-dimensional convex body K and lattice $\mathcal{L}$ in $\mathbb{R}^n$, [KL88] proved the following:

$$1 \leq \mu(K, \mathcal{L}) \min_{\substack{\text{subspace } W \\ k:=\dim(W)\geq 1}} \frac{\text{vol}_k(\pi_W(K))^{1/k}}{\det(\pi_W(\mathcal{L}))^{1/k}} \leq C_{KL}(n), \tag{2}$$

where π_W denotes the orthogonal projection onto W [1]. While not entirely obvious, a first important remark is that the standard flatness theorem corresponds to restricting to one dimensional projections W, and hence the above is its multi-dimensional analogue. To understand the "easy" LHS, one can verify that it directly follows from two basic facts: (1) if sK covers space with respect $\mathcal{L}$ (i.e. $\mathcal{L} + sK = \mathbb{R}^n$), then $\mathrm{vol}_n(sK) \geq \det(\mathcal{L})$ (i.e. sK must be at least as big as "empty space" around $\mathcal{L}$), and (2) the covering property is preserved under projections, i.e. $sK + \mathcal{L} = \mathbb{R}^n \Rightarrow \pi_W(sK) + \pi_W(\mathcal{L}) = W$. For the constant $C_{KL}(n)$, [KL88] proved the upper bound $C_{KL}(n) \leq n$ but could only give a lower bound of $C_{KL}(n) = \Omega(\log n)$, leading to the following question:

Conjecture 1.1 (Kannan-Lováz (KL) Conjecture)**.** *Is $C_{KL}(n) = O(\log n)$?*

A tight connection between the KL conjecture and IP algorithms was given by the author in [Dad12], where it was shown that assuming an oracle for computing KL projections achieving (2), one can give an $O(C_{KL}(n))^n$ time algorithm, i.e. achieving an average branching factor of $O(C_{KL}(n))$. Hence, assuming the KL conjecture and an efficient KL projection algorithm, one would get $O(\log n)^n$-time algorithm. We note that the $O(n)^n$ algorithm from [Dad12] is in fact derived from this result, using a $2^{O(n)}$-time general norm SVP solver to instantiate an oracle achieving the $C_{KL}(n) \leq n$ bound. We note that the best k-dimensional projection above was in fact shown to be computable in $n^{O(kn)}$-time in [DM13], yielding an $n^{O(n^2)}$-time algorithm for finding the global minimum, which is unfortunately far too slow to be useful for IP.

To understand the IP connection, we sketch the dichotomy implied by the above flatness theorem when our goal is to compute a point in $K \cap \mathcal{L}$ (i.e. IP with $\mathcal{L}$ instead of $\mathbb{Z}^n$). In the "fat" case, we will ask for $\mu(K, \mathcal{L}) \leq 1/2$, so not only does K contain a point inside $\mathcal{L}$, but it contains one "deep" inside K, i.e. for any $\mathbf{c} \in K$, $(\frac{1}{2}K + \frac{1}{2}\mathbf{c}) \cap \mathcal{L} \neq \emptyset$. In this case, one can now algorithmically find such a lattice point in $2^{O(n)}$-time by performing a general norm approximate closest vector computation starting from the "center" $\mathbf{c}$ of K (e.g. barycenter), for which there are now various algorithms [Dad14, DK16]. In the opposite "flat" case, assuming for simplicity the worst-case $\mu(K, \mathcal{L}) = 1/2$, (2) now yields the existence of a subspace W (provided by the oracle), $\dim(W) = k \geq 1$, such that $\mathrm{vol}_k(\pi_W(K))/\det(\pi_W(\mathcal{L})) \leq 2^k C_{KL}(n)^k$. Using the fact that $\pi_W(K)$ covers space w.r.t. $\pi_W(\mathcal{L})$ (since $\mu(K, \mathcal{L}) = 1/2$), a standard packing argument allows us to conclude that the projection $\pi_W(K)$ is "uniformly sparse" w.r.t. $\pi_W(\mathcal{L})$, meaning that

$$|(\pi_W(K) + \mathbf{t}) \cap \pi_W(\mathcal{L})| \leq 2^{O(k)} \mathrm{vol}_k(\pi_W(K))/\det(\pi_W(\mathcal{L})) \leq 2^{O(k)} C_{KL}(n)^k, \forall \mathbf{t} \in W. \quad (3)$$

The uniform sparsity condition was shown in [DPV11, Dad12] (combining the use of M-ellipsoid and Voronoi cell computations) to make the sets $(\pi_W(K) + \mathbf{t}) \cap \pi_W(\mathcal{L})$ "easy to enumerate". More precisely, they can be enumerated in $2^{O(k)} C_{KL}(n)^k$-time using 2^n space. Thus, one can use this subroutine – which provides the key extra leverage over basic hyperplane branching – to recurse on all the subproblems indexed by $\pi_W(K) \cap \pi_W(\mathcal{L})$. A straightforward computation reveals that applying the above recursively yields an IP search tree of size $2^{O(n)} C_{KL}(n)^n$.

ℓ_2 Kannan-Lovász conjecture: While the general KL conjecture remains open, the ℓ_2 version of the conjecture (i.e. where K is the Euclidean ball) has recently been resolved up to polyogarithmic

[1] While $\pi_W(\mathcal{L})$ might not be discrete (i.e. not a lattice) for an arbitrary subspace W, one can simply define $\det(\pi_W(\mathcal{L})) = 0$ in this case. As is well-known however, $\pi_W(\mathcal{L})$ is a lattice iff $W^\perp$ admits a basis in $\mathcal{L}$ iff W admits a basis in $\mathcal{L}$.

factors in breakthrough work by Regev and Stephens-Davidowitz [RSD17]. Using the standard estimate $\mathrm{vol}_k(\mathbb{B}_2)^{1/k} = (1+o(1))\sqrt{2\pi e/k}$, the ℓ_2 version can be stated in simplified form as follows. For any n-dimensional lattice $\mathcal{L}$,

$$\frac{1}{\sqrt{2\pi e}} \leq \mu(\mathcal{L}) \min_{\substack{\text{subspace W} \\ k:=\dim(W)\geq 1}} (\sqrt{k}\det(\pi_W(\mathcal{L}))^{1/k})^{-1} \leq C_{KL,2}(n),$$

where [RSD17] proved that the ℓ_2 KL constant satisfies $C_{KL,2}(n) = O(\log^{3/2}(n))$. We note that the only known lower bound is $C_{KL,2}(n) = \Omega(\sqrt{\log n})$ (slightly weaker than the lower bound $C_{KL}(n) = \Omega(\log n)$) achieved by the lattice generated by $\mathbf{e}_1, \mathbf{e}_2/\sqrt{2}, \ldots, \mathbf{e}_n/\sqrt{n}$, $i \in [n]$, i.e. a scaling of the standard basis.

As a consequence, every n-dimensional lattice $\mathcal{L}$ admits a "sparse projection" $\pi_W(\mathcal{L})$, $\dim(W) = k \geq 1$, whose volume radius $\approx \sqrt{k}\det(\pi_W(\mathcal{L}))^{1/k}$ (the radius of the k-dimensional ball of volume $\det(\pi_W(\mathcal{L}))$) certifies a lower bound on the covering radius $\mu(\mathcal{L})$ that is tight up to a $O(\log^{3/2} n)$ factor.

To compare to the basic flatness theorem 1, we first note that one can restrict the minimization to subspaces W which are lattice subspaces of $\mathcal{L}^*$, i.e. which admit a basis of vectors in $\mathcal{L}^*$, since otherwise the projections are not discrete and their determinants are 0. For such a subspace W, the standard duality relations yields $\pi_W(\mathcal{L})^* = \mathcal{L}^* \cap W$ and $\det(\mathcal{L}^* \cap W) = \det(\pi_W(\mathcal{L}))^{-1}$. Thus, one derives the equivalence

$$\min_{\substack{\text{subspace W} \\ k:=\dim(W)\geq 1}} (\sqrt{k}\det(\pi_W(\mathcal{L}))^{1/k})^{-1} = \min_{\substack{W \text{ subspace of } \mathcal{L}^* \\ k:=\dim(W)\geq 1}} \det(\mathcal{L}^* \cap W)^{1/k}/\sqrt{k}.$$

Restricting to a one dimensional W lattice subspace of $\mathcal{L}^*$, we have that $\det(\mathcal{L}^* \cap W) = \det(\mathbb{Z}\mathbf{y}) = \|\mathbf{y}\|$, for some non-zero $\mathbf{y} \in \mathcal{L}^*$. Hence, as claimed previously, the minimum over 1 dimensional subspaces equals $\lambda_1(\mathcal{L}^*)$, corresponding to the standard flatness theorem. From an algorithmic point of view, the one dimensional setting corresponds to solving the shortest vector problem in $\mathcal{L}^*$. Restricting to dimension k, one must find the k-dimensional sublattice of minimum determinant in $\mathcal{L}^*$, which as alluded to previously, was shown in [DM13] to be solvable in $k^{O(kn)}$-time via a massive enumeration of possible bases.

Questions. Motivated by the pursuit of faster IP algorithms a natural question is whether we can actually compute the above ℓ_2 KL projections efficiently (i.e. avoiding massive enumeration), giving a satisfactory algorithmic version of the ℓ_2 KL theorem. While it is unclear whether this would help achieve $n^{o(n)}$-time IP algorithms, as we will explain later, under the assumption that "symmetrization" does not shrink the covering radius (implied by the KL conjecture) it does in fact imply a $n^{n/2+o(n)}$-time IP algorithm.

Another natural question is to understand to what degree of approximation can we certify lower bounds on the lattice covering radius? Thus far, the KL theory has relied on the on the use of a *single projection* to lower bound the covering radius, and thus it is tempting to ask whether it is possible to combine information from multiple projections to achieve a better approximation factor.

2 Our contributions

We answer both of the above questions positively. Our first main result is as follows:

Theorem 2.1 (ℓ_2 KL projection)**.** *Given an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$, there is a $2^{1.6n+o(n)}$-time and $2^{n+o(n)}$-space algorithm randomized algorithm which with high probability computes a lattice subspace W of $\mathcal{L}^*$, where $k := \dim(W) \geq 1$, satisfying $\mu(\mathcal{L}) \leq O(\log^{2.5} n\sqrt{k}\det(\pi_W(\mathcal{L}))^{1/k})$.*

The above algorithmically recovers the [RSD17] ℓ_2-KL bound up to an $O(\log n)$ factor. To achieve the above result, our main technical primitive is an algorithm for computing approximately densest lattice subspaces in any lattice, which may be of independent interest. To state our result, we define some convenient notation.

Definition 2.2 (Normalized Determinant and Determinantal Minima)**.** *For a lattice $\mathcal{L}$, $\dim(\mathcal{L}) \geq 1$, we define its normalized determinant*

$$\mathrm{nd}(\mathcal{L}) = \det(\mathcal{L})^{1/\dim(\mathcal{L})},$$

and its minimum normalized determinant as

$$\tau(\mathcal{L}) = \min_{\text{sublattice } M \subseteq \mathcal{L}, \dim(M) \geq 1} \mathrm{nd}(M).$$

Our main algorithm for computing dense sublattices is given below.

Theorem 2.3 (Densest Subspace)**.** *Given an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$, there is a $2^{1.6n+o(n)}$-time and $2^{n+o(n)}$-space algorithm randomized algorithm which with high probability computes a sublattice $M \subseteq \mathcal{L}$, $\dim(M) \geq 1$, satisfying $\mathrm{nd}(M) \leq O(\log n)\tau(\mathcal{L})$.*

As opposed to the [DM13] algorithm, the above crucially uses the flexibility to adaptively choose the lattice dimension to avoid exhaustive enumeration. The algorithm, which is surprisingly simple, is based on discrete Gaussian sampling, and its analysis crucially relies the reverse Minkowski theorem of [RSD17] which we discuss in the techniques section. To use the dense sublattice algorithm to find ℓ_2-KL projections, we iteratively apply it to find an approximate version of the so-called canonical filtration of the lattice. This filtration corresponds to a canonical way of decomposing a lattice into "dense blocks", which is crucially used in [RSD17] in both the proof of the reverse Minkowski theorem and the ℓ_2-KL conjecture. An appropriate subblock of this decomposition will yield the desired sparse projection.

On the structural side, we give a new lower bound on the covering radius of any lattice, which combines volumetric lower bounds across a chain of lattice subspaces.

Lemma 2.4 (Chain Lower Bound)**.** *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be an n-dimensional lattice. Let $\{0\} = W_0 \subset W_1 \subset \cdots \subset W_k = \mathbb{R}^n$ be lattice subspaces of $\mathcal{L}$, where $d_i := \dim(W_i)$, $i \in [n]$. Then,*

$$\mu(\mathcal{L})^2 \geq \Omega(1)\sum_{i=1}^{k}(d_i - d_{i-1})\mathrm{nd}(\mathcal{L}/W_{i-1})^2,$$

where $\mathcal{L}/W_{i-1} := \pi_{W_{i-1}^{\perp}}(\mathcal{L})$, $i \in [k]$.

Note that the above lower bound is easy to compute given bases of the lattices $\mathcal{L}/W_0, \ldots, \mathcal{L}/W_{k-1}$, and yield NP certificate for lower bounds on the covering radius. To prove the above, we rely on a semidefinite programming lower bound for $\mu(\mathcal{L})^2$, implicit in the work of [DR16], and show how to "compile" any chain of subspaces into a solution of this program. Using the proof the ℓ_2-KL conjecture of [RSD17], which shows that one should instantiate the bound using the so-called canonical filtration of $\mathcal{L}$, we get that these lower bounds are suprisingly tight. More precisely, we get the following:

Theorem 2.5. *For $n \in \mathbb{N}$, define*

$$L^{\mu}(n) = \max\{\mu(\mathcal{L})/\sqrt{n} : n \text{ dimensional lattice } \mathcal{L}, \mathrm{nd}(\mathcal{L}) = \tau(\mathcal{L}) = 1\}.$$

Then, for any n-dimensional lattice $\mathcal{L}$, the best possible lower bound on $\mu(\mathcal{L})$ taken from Lemma 2.4 is tight within an $O(L^{\mu}(n))$ factor.

The maximum on $L^{\mu}(n)$ is taken over what are known as *stable* lattices in dimension n, which are lattice having determinant 1 but whose sublattices all have determinant at least 1. It is conjectured in [RSD17] that this maximum is in fact attained at $\mathcal{L} = \mathbb{Z}^n$, resulting in the bound $L^{\mu}(n) = 1/2$. Under this assumption, we conclude the existence of NP-certificates lower bounding the covering radius that are tight up to a constant factor.

Furthermore, [RSD17] prove the bound $L^{\mu}(n) = O(\log n)$ and show that $L^{\mu}(n)$ is constant factor equivalent to the worst-case slicing constant of any n-dimensional Voronoi cell of a stable lattice. We provide the relevant definitions below.

Definition 2.6 (Voronoi Cell)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$, we define the Voronoi cell of $\mathcal{L}$ by*

$$\mathcal{V}(\mathcal{L}) = \{\mathbf{x} \in \mathrm{span}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle \leq \frac{1}{2}\|\mathbf{y}\|^2 \ \forall \mathbf{y} \in \mathcal{L}\}\,.$$

In words, the Voronoi cell is the set of all points in $\mathrm{span}(\mathcal{L})$ *that are closer to the origin than any other lattice point.*

Definition 2.7 (Slicing Constant)**.** *For a symmetric convex body $K \subseteq \mathbb{R}^n$. The slicing constant L_K of K is defined by*

$$L_K^2 := \min_{\mathbf{T}:\ \mathrm{vol}_n(\mathbf{T}K)=1} \mathop{\mathbb{E}}_{\mathbf{X} \sim \mathrm{unif}(K)} [\|\mathbf{T}\mathbf{X}\|^2/n],$$

where the minimum is take over all invertible linear transformations $\mathbf{T}$ such that $\mathrm{vol}_n(\mathbf{T}K) = 1$. Among n-dimensional symmetric convex bodies, it is well-known that the Euclidean ball $\mathbb{B}_2^n$ has the smallest slicing constant, which satisfies $L_{\mathbb{B}_2^n} = (1 + o(1))/\sqrt{2\pi e}$.

A major conjecture in convex geometry is Bourgain's slicing conjecture, which asks if the slicing constant of any symmetric convex body is $O(1)$. The best known bound of $O(n^{1/4})$ is due to Klartag [Kla06]. As [RSD17] prove a bound of $O(\log n)$ on the slicing constant of Voronoi cells of stable lattices, it is natural to ask whether such a bound can be extended to arbitrary Voronoi cells. We answer this in the affirmative below.

Theorem 2.8 (Slicing Constant of Voronoi Cells)**.** *For any n-dimensional lattice $\mathcal{L}$, its Voronoi cell $\mathcal{V} := \mathcal{V}(\mathcal{L})$ satisfies $L_{\mathcal{V}} = O(C_{KL,2}(n)) = O(\log^{3/2} n)$, where $C_{KL,2}(n)$ is the ℓ_2 Kannan-Lovász constant in dimension n.*

To prove the above theorem, we in fact show that every Voronoi cell $\mathcal{V} \subseteq \mathbb{R}^n$, one can find an ellipsoid E, such that $\mathcal{V} \subseteq E$ and $\mathrm{vol}(E)^{1/n} \leq O(C_{KL,2}(n))\,\mathrm{vol}_n(\mathcal{V})$. This bounds what is known as the outer volume ratio of Voronoi cells, from which a bound on the slicing constant easily follows from known techniques.

Application to IP. We show that using Theorem 2.1 we can in fact derive a conditionally faster algorithm for IP. The improvement is based on the conjecture that the following symmetrization parameter is small:

Definition 2.9. *For $n \in \mathbb{N}$, define*

$$C_{\mathrm{sym}}(n) = \sup\{\mu(K, \mathbb{Z}^k)/\mu(K-K, \mathbb{Z}^k) : 1 \leq k \leq n, \text{ convex body } K \subseteq \mathbb{R}^k\}$$

In the above, $C_{\mathrm{sym}}(n)$ measures by how much the covering radius can drop when moving from a convex body K to its difference body $K - K$. Note that since we maximize over all bodies above, by linear transformation the bound also holds for all lattices as well. It is in fact easy to show that $C_{\mathrm{sym}}(n) \leq 4C_{KL}(n)$. This follows directly from the so-called Rogers-Shephard inequality [RS57], which says that $\mathrm{vol}_n(K-K)^{1/n} \leq 4\,\mathrm{vol}_n(K)^{1/n}$. Note that this implies that the volumetric lower bounds in 2 can only drop by a factor of 4 when moving K to $K - K$ and hence the bound $\mu(K, \mathcal{L}) \leq 4C_{KL}(n)\mu(K-K, \mathcal{L})$. While the KL conjecture implies the bound of $O(\log n)$ on $C_{\mathrm{sym}}(n)$, it is entirely possible that $C_{\mathrm{sym}}(n)$ is in fact $O(1)$.

Using the above parameter, we state the guarantees for our IP algorithm, which makes use of the KL projection algorithm 2.1.

Theorem 2.10. *Given an appropriately represented convex body $K \subseteq \mathbb{R}^n$, there exists a randomized $O(C_{\mathrm{sym}}(n)\sqrt{n}\log^{2.5}(n))^n$-time and $2^{n+o(n)}$-space algorithm which with high probability, either correctly decides that $K \cap \mathbb{Z}^n = \emptyset$ or outputs a point in $K \cap \mathbb{Z}^n$.*

In particular, assuming that $C_{\mathrm{sym}}(n) = n^{o(1)}$, the above algorithm runs in $n^{n/2+o(n)}$-time, reducing the constant in the exponent to $1/2$ for IP. While this would be only a modest improvement, perhaps more interestingly, it yields good motivation for exploring the basic consequences of the KL theory. Furthermore, one would expect that the task of bounding $C_{\mathrm{sym}}(n)$ is substantially easier than proving the full KL conjecture.

We now sketch the idea of the algorithm below and leave a formal proof to the full version. The main idea is in fact quite simple. Firstly, using the framework of [Dad12], as described in the introduction, we need only provide an oracle to compute a sparse projection. In particular, given a lattice $\mathcal{L}$ and convex body K in $\mathbb{R}^n$, we need only find a projection π_W, $k := \dim(W) \geq 1$, satisfying

$$\mathrm{vol}_k(\pi_W(K))^{1/k}/\det(\pi_W(\mathcal{L}))^{1/k} \leq O(C_{\mathrm{sym}}(n)\sqrt{n}\log^{2.5}(n))$$

under the assumption that $\mu(K, \mathcal{L}) = 1/2$.

To achieve this, we will use the fact for the symmetric convex body $K - K$, John's theorem yields an ellipsoid E which satisfies $E/\sqrt{n} \subseteq K - K \subseteq E$, i.e. better than the n-factor sandwiching one gets for general convex bodies. Furthermore, one can in fact compute such an ellipsoid E, with a sandwiching factor $O(\sqrt{n})$ instead of $\sqrt{n}$, in $2^{O(n)}$-time given a membership oracle for $K - K$ (see for example [HK13]). From here, we have the inequalities

$$\mu(E, \mathcal{L}) \geq \frac{\mu(K-K, \mathcal{L})}{O(\sqrt{n})} \geq \frac{\mu(K, \mathcal{L})}{O(C_{\mathrm{sym}}(n)\sqrt{n})} = \frac{1}{O(C_{\mathrm{sym}}(n)\sqrt{n})}\,.$$

We now apply Theorem 2.1 to find a sparse projection π_W for E w.r.t. $\mathcal{L}$ (by first applying the linear transformation to $\mathcal{L}$ that sends E to $\mathbb{B}_2^n$). By the guarantees of the algorithm,

$$\frac{\mathrm{vol}_k(\pi_W(K))^{1/k}}{\det(\pi_W(\mathcal{L}))^{1/k}} \leq \frac{\mathrm{vol}_k(\pi_W(E))^{1/k}}{\det(\pi_W(\mathcal{L}))^{1/k}} \leq O(\log^{2.5}(n)/\mu(E,\mathcal{L})) = O(C_{\mathrm{sym}}(n)\sqrt{n}\log^{2.5}(n)),$$

as needed. As a final remark, the above proof is simply an "algorithmification" of the bound $C_{KL}(n) \leq C_{\mathrm{sym}}(n) \cdot C_{KL,2}(n) \cdot \sqrt{n} = O(C_{\mathrm{sym}}(n)\sqrt{n}\log^{3/2} n)$.

2.1 Techniques: Finding Dense Lattice Subspaces

We describe here the main ideas behind our algorithm for finding dense lattice subspaces, deferring discussion of the remaining results to the relevant sections of the paper. To begin, for any set $A \subset \mathbb{R}^n$, we define the discrete Gaussian sum

$$\rho(A) = \sum_{\mathbf{y} \in A} e^{-\pi\|\mathbf{y}\|^2}.$$

We define the discrete Gaussian distribution $D_{\mathcal{L}}$ to be the distribution satisfying

$$\Pr_{\mathbf{X} \sim D_{\mathcal{L}}}[\mathbf{X} = \mathbf{y}] = \rho(\mathbf{y})/\rho(\mathcal{L})$$

for $\mathbf{y} \in \mathcal{L}$ and 0 otherwise. Our algorithms will rely on the ability to sample from the discrete Gaussian distribution. For this purpose, we rely on the results of [ADRS15], which give a $2^{n+o(n)}$-time and space algorithm discrete Gaussian sampling (DGS).

Given access to DGS sampler, our algorithm for computing an $O(\log n)$-densest sublattice in $2^{O(n)}$-time is remarkably simple. The idea will be to iteratively reduce dimension by 1 starting from the full lattice $\mathcal{L}$ until we get to $\{0\}$, keeping track of the densest lattice we've seen through the process.

Starting at $M \leftarrow \mathcal{L} \subset \mathbb{R}^n$, the full lattice, we proceed as follows. We sample a DGS sample $\mathbf{X}$ on D_{sM^*}, where $s = \tau(M)/2$, and replace $M \leftarrow M \cap \mathbf{X}^{\perp}$. We repeat the process until $M = \{0\}$ and return the densest sublattice found.

The fundamental claim underlying the algorithm is the following.

Claim 2.11. *Let W be any lattice subspace of M satisfying $\mathrm{nd}(M \cap W) = \tau(M)$. Assume that $\mathrm{nd}(M) \geq \Omega(\log n)\tau(M)$. Then, with probability $\Omega(1)$, $\mathbf{X}$ sampled as above is non-zero and orthogonal to Λ.*

Assuming the above, it is clear that the above procedure has probability at least $2^{-O(n)}$ of finding an $O(\log n)$-approximately densest sublattice. This is because at every step of the procedure, we are either already done, or we reduce dimension by 1 and maintain that M continues to contain a densest sublattice $M \cap W$ in the next iteration with constant probability. Furthermore, we can only decrease dimension n times. Thus, running the above procedure $2^{O(n)}$ times yields the desired algorithm.

It remains to prove the claim. For this purpose, we first note s above is chosen so that $\tau(sM^*) = s/\tau(M) = 1/2$. Thus, sM^* has determinant 2^{-k}, $k := \dim(M)$, i.e. sM^* is somewhat "dense". From here, standard estimates reveal that Gaussian mass is relatively large, namely $\rho(sM^*) \geq 1/\det(sM^*) \geq 2^k$, which directly imples that the probability that $\mathbf{X} \sim D_{sM^*}$ equals 0 is at most 2^{-k}.

Given the above, it suffices to show that $\mathbf{X}$ is orthogonal to any densest lattice subspace W of M with good probability. From here, using determinantal arithmetic, our assumption that $\mathrm{nd}(M) = \Omega(\log n)\tau(M)$ and $\mathrm{nd}(M \cap W) = \tau(M)$ implies that $\mathrm{nd}(\pi_W(M^*)) = \mathrm{nd}(M \cap W)^{-1} = 1/\tau(M)$. Slightly less obvious is that in fact $\tau(\pi_W(M^*)) = \mathrm{nd}(\pi_W(M^*))$, i.e. $\pi_W(M^*)$ must be its own densest sublattice, though this is again a consequence of determinantal arithmetic. By our choise of s, we know have that $\tau(\pi_W(sM^*)) = \mathrm{nd}(M)/(2\tau(M)) = \Omega(\log n)$, i.e. $\pi_W(sM^*)$ is "uniformly sparse", i.e. contains no dense lattice subspaces.

Recall, that we wish to show that $\mathbf{X} \perp W \Leftrightarrow \pi_W(\mathbf{X}) = \mathbf{0}$. From here, a standard corellation inequality reveals that $\Pr[\pi_W(\mathbf{X}) = \mathbf{0}] \geq 1/\rho(\pi_W(sM^*))$, which is the probability of hitting 0 if $\mathbf{X}$ were sampled directly from $D_{\pi_W(sM^*)}$. Hence, we are now left with the task of showing that $\rho(\pi_W(sM^*)) = O(1)$, i.e. of showing small Gaussian mass under the assumption of "uniform sparsity".

Enter Reverse Minkowski. The reverse Minkowski theorem of [RSD17], first conjectured in [DR16], gives quantitative bounds on how many short lattice vectors are required to ensure the existence of a sublattice of small normalized determinant. More specifically, they show that if $\mathcal{L} \subset \mathbb{R}^n$ contains at least $2^{\Omega(\log^2 nk)}$ points of of length at most $r > 0$, then $\mathcal{L}$ must contain a sublattice of normalized determinant at most $r/\sqrt{k}$, i.e. $\tau(\mathcal{L}) \leq r/\sqrt{k}$. One can conveniently formalize the above in terms of discrete Gaussian sums, which will be more directly useful to us.

Definition 2.12 (Reverse Minkowski Constant)**.** *For $n \in \mathbb{N}$, define $C_\eta(n)$ to the smallest number such that for any lattice $\mathcal{L}$ of dimension at most n with $\tau(\mathcal{L}) \geq C_\eta(n)$ satisfies $\rho(\mathcal{L}) \leq 3/2$.*

Their reverse Minkowski theorem can now be stated as follows:

Theorem 2.13. *[RSD17] $C_\eta(n) = O(\log n)$.*

The above theorem now directly implies the desired claim, by choosing constants appropriately to ensure that $\tau(\pi_W(sM^*)) \geq C_\eta(n)$.

2.2 Organization

In section 3, we introduce the concepts needed throughout the paper. In section 4, we give our algorithm for computing dense lattice subspaces. In section 5, we give a novel characterization of the covering radius which enables the tighter complexity classification for GapCRP. In section 6, we show how to use our algorithm for finding dense lattice subspaces to compute approximately stable filtrations and good KL projections. Lastly, in section 7, we give the $O(\log^{3/2} n)$ bound on the slicing constant of Voronoi cells.

Acknowledgment. We are grateful to Noah Stephens-Davidowitz and Oded Regev for useful conversations. In particular, Oded Regev suggested factoring the slicing bound through the volume ratio as in 7.1.

3 Preliminaries

We write $X \lesssim Y$ to mean that there exists universal constant $C > 0$ such that $X \leq CY$, and similarly for $X \gtrsim Y$. We write ln to denote the natural logarithm and log to denote the logarithm

base 2. For a natural number $n \in \mathbb{N}$, we define $[n] = \{1, \ldots, n\}$ and $[n]_0 = [n] \cup \{0\}$. We denote the d-dimensional Lebesgue measure by $\mathrm{vol}_d(\cdot)$. For a subset $A \subseteq \mathbb{R}^n$, we define $\mathrm{span}(A)$ to be the linear span of A.

For a matrix $M \in \mathbb{R}^{n \times m}$, we define the matrix transpose $(M^\mathsf{T})_{ij} = M_{ji}$, $j \in [n]$,$i \in [m]$. We define the standard inner product between vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ by $\langle \mathbf{x}, \mathbf{y} \rangle := \mathbf{x}^\mathsf{T}\mathbf{y} = \sum_{i=1}^n x_i y_i$.

Definition 3.1 (Euclidean Ball)**.** *The Euclidean norm of a vector $\mathbf{x} \in \mathbb{R}^n$ is define as $\|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$. Define $\mathbb{B}_2^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq 1\}$ to the unit Euclidean ball. Letting $\omega_n = \mathrm{vol}_n(\mathbb{B}_2^n)$, we have the standard inequalities*

$$\frac{1}{2\sqrt{n\pi}} \left(\frac{2\pi e}{n}\right)^{n/2} \leq \omega_n \leq \frac{1}{\sqrt{n\pi}} \left(\frac{2\pi e}{n}\right)^{n/2} .$$

Definition 3.2 (Orthogonal Projection and Complement)**.** *For a linear subspace $W \subseteq \mathbb{R}^n$, we define π_W to be the orthogonal projection onto W. We define $W^\perp = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle = 0, \forall \mathbf{y} \in W\}$ to be the orthogonal complement of W.*

Definition 3.3 (Positive Semidefinite Matrices)**.** *A matrix $X \in \mathbb{R}^{n \times n}$ is positive semidefinite (PSD) if $X = X^\mathsf{T}$ and $\mathbf{y}^\mathsf{T} X \mathbf{y} \geq 0$, $\forall \mathbf{y} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$. X is positive definite (PD) if the previous inequality holds strictly for all non-zero $\mathbf{y} \in \mathbb{R}^n$.*

We denote the set of $n \times n$ PSD matrices by $\mathbb{S}_+^n$. We use the notation $A \succeq B \Leftrightarrow A - B \in \mathbb{S}_+^n$ to denote the Löwner ordering on PSD matrices. In particular, $A \succeq 0$ iff A is PSD. Similarly, we use the notation $A \succ B$ if $A - B$ is positive definite.

Definition 3.4 (Ellipsoidal Norm)**.** *Given an $n \times n$ positive definite matrix $Q \succ 0$, we define the norm induced by Q as $\|\mathbf{x}\|_Q = \sqrt{\mathbf{x}^\mathsf{T} Q \mathbf{x}}$, for any $\mathbf{x} \in \mathbb{R}^n$.*

Definition 3.5 (Projected Determinant)**.** *Given $Q \in \mathbb{S}_+^n$ and a subspace $W \subseteq \mathbb{R}^n$, we define*

$$\det_W(Q) = \det(U^\mathsf{T} Q U),$$

where the columns of U form an orthonormal basis of W. Note this definition is invariant to the choice of orthonormal basis for W.

3.1 Lattice Basics

Definition 3.6 (Lattices and their Bases)**.** *A d-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$ corresponds to all integer combinations of some d linear independent vectors $B = (b_1, \ldots, b_d) \in \mathbb{R}^n$, which is known as a* basis *of $\mathcal{L}$. We shall use the notation $\mathcal{L}(B)$ to denote the lattice generated by the basis B. We define the dimension $\dim(\mathcal{L})$ of $\mathcal{L}$ to be the dimension of its linear span, or equivalently, the minimum number of elements in any basis of $\mathcal{L}$.*

The dual lattice *of $\mathcal{L}$ is defined to be*

$$\mathcal{L}^* = \{\mathbf{x} \in \mathrm{span}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z} \forall \mathbf{y} \in \mathcal{L}\} .$$

In particular, if B is a basis for $\mathcal{L}$ then $B(B^\mathsf{T}B)^{-1}$ is a basis for $\mathcal{L}^*$.

Definition 3.7 (Lattice Determinant)**.** *For a d-dimensional lattice $\mathcal{L} := \mathcal{L}(B) \subset \mathbb{R}^n$, we define the determinant of $\mathcal{L}$ to be $\det(\mathcal{L}) := \det(B^\mathsf{T}B)^{1/2}$. Note that the determinant is invariant to the choice of*

basis. By convention the determinant of the trivial lattice $\{\mathbf{0}\}$ *is* 1*, that is* $\det(\{\mathbf{0}\}) = 1$*. It is easy to verify that the determinant of the dual lattice* $\mathcal{L}^*$ *satisfies* $\det(\mathcal{L}^*) = 1/\det(\mathcal{L})$*.*

We define the normalized determinant *of* $\mathcal{L}$ *by* $\mathrm{nd}(\mathcal{L}) = \det(\mathcal{L})^{1/\dim(\mathcal{L})}$*. Note that the normalized determinant is homogeneous, namely for* $t > 0$*,* $\mathrm{nd}(t\mathcal{L}) = t \cdot \mathrm{nd}(\mathcal{L})$*.*

Definition 3.8 (Lattice Subspaces and Projections)**.** *Let* $\mathcal{L} \subset \mathbb{R}^n$ *be a lattice. A linear subspace* $W \subseteq \mathbb{R}^n$ *is a lattice subspace of* $\mathcal{L}$ *if* W *admits a basis in* $\mathcal{L}$*, i.e.* $W = \mathrm{span}(\mathcal{L} \cap W)$*. We define* $\mathcal{L}/W := \pi_{W^\perp}(\mathcal{L})$ *to be the projection of* $\mathcal{L}$ *orthogonal to* W*.*

Lemma 3.9. *Let* $\mathcal{L} \subset \mathbb{R}^n$ *be a lattice and let* $W \subseteq \mathbb{R}^n$ *be a linear subspace. Then the following holds:*

1. $W \subseteq \mathbb{R}^n$ *is a lattice subspace of* $\mathcal{L}$ *iff* $W^\perp \cap \mathrm{span}(\mathcal{L})$ *is a lattice subspace of* $\mathcal{L}^*$*.*
2. $\mathcal{L}/W := \pi_{W^\perp}(\mathcal{L})$ *is a lattice* $\Leftrightarrow$ W *is a lattice subspace of* $\mathcal{L}$*.*

The following lemma explains how to dualize sections and projections of a lattice.

Lemma 3.10 (Section and Projection Duality)**.** *Let* $W_1 \subseteq W_2$ *be lattice subspaces of a lattice* $\mathcal{L} \subset \mathbb{R}^n$*. Then* $((\mathcal{L} \cap W_2)/W_1)^* = (\mathcal{L} \cap W_1^\perp)/W_2^\perp$*.*

The following lemma gives the formula for how to decompose the determinant of a lattice with respect to a lattice subspace.

Lemma 3.11 (Lattice Determinant Lemma)**.** *Let* $W \subseteq \mathbb{R}^n$ *be a lattice subspace of* $\mathcal{L}$*, then*

$$\det(\mathcal{L}) = \det(\mathcal{L} \cap W)\det(\mathcal{L}/W)$$

3.2 Lattice Parameters and Geometry

Definition 3.12 (Successive Minima)**.** *For a lattice* $\mathcal{L} \subset \mathbb{R}^n$*, define* $\lambda_i(\mathcal{L}) = \inf\{r \geq 0 : \dim(\mathcal{L} \cap r\mathbb{B}_2^n) \geq i\}$ *for* $i \in \dim(\mathcal{L})$*.*

Definition 3.13 (Determinantal Minima)**.** *For a lattice* $\mathcal{L} \subset \mathbb{R}^n$ *and* $k \in [\dim(\mathcal{L})]$*, define*

$$\tau_k(\mathcal{L}) = \min_{\substack{M \subseteq \mathcal{L} \text{ sublattice} \\ \dim(M) = k}} \det(M),$$

to be the k^{th} *minimum determinant (or determinantal minimum) of* $\mathcal{L}$*. By convention, we define* $\tau_0(\mathcal{L}) = 1$*. Define*

$$\tau(\mathcal{L}) = \min_{k \in [\dim(\mathcal{L})]} \tau_d(\mathcal{L})^{1/d}$$

to be the minimum normalized determinant of $\mathcal{L}$*.*

Note that for any minimizing sublattice M *for* $\tau_k(\mathcal{L})$*, we must have* $M = \mathrm{span}(M) \cap \mathcal{L}$*, since otherwise replacing* M *by* $\mathrm{span}(M) \cap \mathcal{L}$ *would decrease the determinant. Therefore, we can equivalently above minimize over k-dimensional lattice subspaces instead of sublattices.*

The following lemma shows the simple duality relation satisfied by the determinantal minima.

Lemma 3.14 (Determinantal Duality). *For a lattice $\mathcal{L} \subset \mathbb{R}^n$ and $d \in [\dim(\mathcal{L})]_0$, then*

$$\tau_d(\mathcal{L}^*) = \tau_{\dim(\mathcal{L})-d}(\mathcal{L}) / \det(\mathcal{L}) \, .$$

Proof. Let W denote a $\dim(\mathcal{L}) - d$-dimensional lattice subspace. By the lattice determinant Lemma 3.11, we have that

$$\det(\mathcal{L} \cap W) \det(\mathcal{L}/W) = \det(\mathcal{L}) \Rightarrow \frac{\det(\mathcal{L} \cap W)}{\det(\mathcal{L})} = \frac{1}{\det(\mathcal{L}/W)} \, .$$

By duality, $\frac{1}{\det(\mathcal{L}/W)} = \det((\mathcal{L}/W)^*) = \det(\mathcal{L}^* \cap W^\perp)$, and hence

$$\frac{\det(\mathcal{L} \cap W)}{\det(\mathcal{L})} = \det(\mathcal{L}^* \cap W^\perp) \, .$$

Since $W^\perp \cap \text{span}(\mathcal{L})$ is a lattice subspace of $\mathcal{L}^*$ and $\dim(W^\perp \cap \text{span}(\mathcal{L})) = \dim(\mathcal{L}) - \dim(W) = d$, by Lemma 3.9 we get the equality

$$\begin{aligned} \frac{\tau_{\dim(\mathcal{L})-d}(\mathcal{L})}{\det(\mathcal{L})} &= \min\{\frac{\det(\mathcal{L} \cap W)}{\det(\mathcal{L})} : W \text{ lattice subspace of } \mathcal{L}, \dim(W) = \dim(\mathcal{L}) - d\} \\ &= \min\{\det(\mathcal{L}^* \cap W) : W \text{ lattice subspace of } \mathcal{L}^*, \dim(W) = d\} = \tau_d(\mathcal{L}^*) \, , \end{aligned}$$

as needed. □

We now present a restatement of (a slight generalization of) Minkowski's Second Theorem as a quantitive relationship between determinantal and successive minima.

Theorem 3.15 (Minkowski's Second Theorem). *For an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$ and $d \in [n]$, we have that*

$$d^{-d/2} \prod_{i=1}^{d} \lambda_i(\mathcal{L}) \leq \tau_d(\mathcal{L}) \leq \prod_{i=1}^{d} \lambda_i(\mathcal{L}) \, .$$

Definition 3.16 (Voronoi Cell). *For a lattice $\mathcal{L} \subset \mathbb{R}^n$, we define the Voronoi cell of $\mathcal{L}$ by*

$$\mathcal{V}(\mathcal{L}) = \{\mathbf{x} \in \text{span}(\mathcal{L}) : \langle \mathbf{x}, \mathbf{y} \rangle \leq \frac{1}{2}\|\mathbf{y}\|^2 \; \forall \mathbf{y} \in \mathcal{L}\} \, .$$

In words, the Voronoi cell is the set of all points in $\text{span}(\mathcal{L})$ *that are closer to the origin than any other lattice point.*

Definition 3.17 (Covering Radius). *For a lattice $\mathcal{L} \subset \mathbb{R}^n$ and $x \in \mathbb{R}^n$, define $d(\mathbf{x}, \mathcal{L}) = \min\{\|\mathbf{y} - \mathbf{x}\|_2 : \mathbf{y} \in \mathcal{L}\}$ to be the distance of $\mathbf{x}$ to $\mathcal{L}$. From here, we define the covering radius of $\mathcal{L}$ as*

$$\mu(\mathcal{L}) = \max\{d(\mathbf{x}, \mathcal{L}) : \mathbf{x} \in \text{span}(\mathcal{L})\} = \max\{\|\mathbf{x}\| : \mathbf{x} \in \mathcal{V}\} \, ,$$

to be the largest ℓ_2 distance of any point in the linear span of $\mathcal{L}$ to itself. We also define the smoother average version of the covering radius

$$\bar{\mu}(\mathcal{L})^2 = \mathop{\mathbb{E}}_{\mathbf{x} \sim \text{span}(\mathcal{L})/\mathcal{L}}[d(\mathbf{x}, \mathcal{L})^2] := \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathcal{V}}[\|\mathbf{x}\|^2] \, .$$

The following lemma yields the tight relationship between the averaged and normal version of the covering radius.

Lemma 3.18. *[HLR09] For any lattice $\mathcal{L}$, $\bar{\mu}(\mathcal{L}) \leq \mu(\mathcal{L}) \leq 2\bar{\mu}(\mathcal{L})$.*

Definition 3.19 (ℓ_2 Kannan Lovász Constant)**.** *Define $C_{KL,2}(n)$ to be the least positive number such that for any lattice $\mathcal{L}$ of dimension at most n, we have*

$$\mu(\mathcal{L}) \leq C_{KL,2}(n) \max_{\substack{W \text{ lattice subspace of } \mathcal{L}^* \\ \dim(W) \in [\dim(\mathcal{L})]}} \sqrt{\dim(W)} \cdot \mathrm{nd}(\pi_W(\mathcal{L})) \,.$$

3.3 Hermite Korkine Zolotareff (HKZ) Bases

For a basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ and $I \subseteq [n]$, we define B_I to the restriction of $\mathbf{B}$ to the columns in I. We define the Gram Schmidt Orthogonalization of $\mathbf{B}$ by $\widetilde{\mathbf{B}} = (\widetilde{\mathbf{b}}_1, \ldots, \widetilde{\mathbf{b}}_n)$ where

$$\widetilde{\mathbf{b}}_i = \pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}(\mathbf{b}_i),\ \forall i \in [n] \,,$$

where $\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\}^\perp$ denotes the orthogonal complement of $\mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_{i-1})$.

For a lattice $\mathcal{L}$, a basis $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ is an HKZ-basis if $\widetilde{\mathbf{b}}_i = \lambda_1(\pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}(\mathcal{L}))$.

The following is derived from the algorithm of [ADRS15] for the Shortest Vector Problem.

Theorem 3.20. *Given an n-dimensional lattice $\mathcal{L} := \mathcal{L}(\mathbf{B})$, $\mathbf{B} \in \mathbb{Q}^{n\times n}$, $\kappa \geq 1$, there is a $2^{n+\log(\kappa)+o(n)}$-time and -space algorithm which computes an HKZ-basis $\mathbf{B}'$ of $\mathcal{L}$ with probability at least $1 - 2^{-\kappa n}$.*

The following tells us how well an HKZ basis can approximate the minimum subdeterminants of a lattice.

Lemma 3.21. *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice and let $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ denote an HKZ-basis of $\mathcal{L}$. Then*

$$d^{-d/2} \det(\mathcal{L}(\mathbf{B}_{[d]})) \leq \tau_d(\mathcal{L}) \leq \det(\mathcal{L}(\mathbf{B}_{[d]})) \,.$$

In particular,

$$\frac{1}{\sqrt{n}} \min_{d\in[n]} \mathrm{nd}(\mathcal{L}(\mathbf{B}_{[d]})) \leq \tau(\mathcal{L}) \leq \min_{d\in[n]} \mathrm{nd}(\mathcal{L}(\mathbf{B}_{[d]})) \,.$$

Proof. The upper bound is trivial, so we focus on the lower bound. For this purpose, we claim that

$$\det(\mathcal{L}(\mathbf{B}_{[d]})) = \prod_{i=1}^{d} \|\widetilde{\mathbf{b}}_i\| \leq \prod_{i=1}^{d} \lambda_i(\mathcal{L}) \,.$$

Given the claim, the Lemma follows directly from Minkowski's second theorem. To prove the claim, it suffices to show that $\|\widetilde{\mathbf{b}}_i\| \leq \lambda_i(\mathcal{L})$, for $i \in [n]$. Let $\mathbf{v}_1, \ldots, \mathbf{v}_i \in \mathcal{L}$, denote linear independent vectors satisfying $\|\mathbf{v}_j\| = \lambda_j(\mathcal{L})$, $j \in [i]$. Since they span an i-dimensional space, one of them must be mapped to something non-zero under the projection $\pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}$. Letting $j^* \in [i]$, denote an index such that $\pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}(\mathbf{v}_j) \neq 0$, we get that

$$\|\widetilde{\mathbf{b}}_i\| = \lambda_1(\pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}(\mathcal{L})) \leq \|\pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}(v_j)\| \leq \|\mathbf{v}_j\| \leq \lambda_j(\mathcal{L}) \leq \lambda_i(\mathcal{L}) \,,$$

as needed.

For the last statement, it follows directly from the first part and the definition of $\tau(\mathcal{L})$. □

3.4 Stable Lattices and the Canonical Filtration

In this section, we will explain how one can decompose any lattice into *well-conditioned* pieces, which are known as stable lattices. The theory of stable lattices was developed by [HN75, Stu76, Gra84], see [Gra84] for a thorough treatment.

Definition 3.22 (Stable Lattice)**.** *A lattice $\mathcal{L}$ is stable if all sublattices of $\mathcal{L}$ have normalized determinant at least $\mathrm{nd}(\mathcal{L})$. Note that if $\mathcal{L}$ has determinant 1, then it is stable if all its sublattices have determinant at least 1. We note that in most treatments stability requires $\det(\mathcal{L}) = 1$, but it will be convenient for us to relax this condition.*

Definition 3.23 (Lattice Filtration)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$, we define an ordered sequence $\mathcal{F} = (W_0, \ldots, W_k)$ of lattice subspaces of $\mathcal{L}$ to be a filtration of size k if:*

1. **Boundary condition:** $W_0 = \{\mathbf{0}\}$, $W_k = \mathrm{span}(\mathcal{L})$.
2. **Chain condition:** $W_{i-1} \subset W_i$, $i \in [k]$.

We define the dual filtration $\mathcal{F}^* = (W_k^\perp \cap \mathrm{span}(\mathcal{L}), W_{k-1}^\perp \cap \mathrm{span}(\mathcal{L}), \ldots, W_0^\perp \cap \mathrm{span}(\mathcal{L}))$, namely the complementary filtration for the dual lattice.

Definition 3.24 (Canonical Polytope)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$, we define the canonical polytope of $\mathcal{L}$*

$$\mathcal{P}(\mathcal{L}) = \mathrm{conv}((\dim(W), \ln(\det(\mathcal{L} \cap W))) : W \text{ lattice subspace of } \mathcal{L}) + \{(0, y) : y \geq 0\},$$

i.e. all points above the convex hull $\mathrm{conv}((\dim(W), \ln(\det(\mathcal{L} \cap W))) : W$ lattice subspace of $\mathcal{L})$ in $\mathbb{R}^2$. Since only the "lowest" points matter in the definition of the polytope, it is easy to check that

$$\mathcal{P}(\mathcal{L}) = \mathrm{conv}((i, \ln(\tau_d(\mathcal{L}))) : i \in [\dim(\mathcal{L})]_0) + \{(0, y) : y \geq 0\}.$$

We recall that by determinantal duality (Lemma 3.14), $\det_d(\mathcal{L}^*) = \tau_{\dim(\mathcal{L})-d}(\mathcal{L})/\det(\mathcal{L})$. Given the above, it is easy to check that the canonical polytope $\mathcal{P}(\mathcal{L}^*)$ of the dual lattice can be obtained by reflecting $\mathcal{P}(\mathcal{L})$ about the axis $\{(x, y) : x = \dim(\mathcal{L})/2\}$ and shifting by down by $\ln(\det(\mathcal{L}))$.

From the canonical polytope, we define the canonical filtration:

Definition 3.25 (Canonical Filtration)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$, we define a canonical filtration $\mathcal{F}_\mathcal{L} = (W_0, \ldots, W_k)$ of $\mathcal{L}$ to be a sequence of subspaces ordered in increasing order of dimension for which the point set $\{(W_i, \ln\det(\mathcal{L} \cap W_i)) : i \in [k]_0\}$ is exactly the set of vertices of the lower hull of $\mathcal{P}(\mathcal{L})$.*

While it may seem from the above definition that there may be many canonical filtrations, and that the canonical filtration need not form a chain (and hence be a filtration in the sense of 3.23), rather remarkably, all the properties one could hope for in fact do hold. Furthermore, the canonical filtration will give us a powerful way to decompose a lattice into stable (i.e. "well-conditioned") lattices as is summarized by the following lemma:

Lemma 3.26. *[RSD17, Proposition 2.5] For a lattice $\mathcal{L} \subset \mathbb{R}^n$, let $\mathcal{F}_\mathcal{L} = (W_0, \ldots, W_k)$ be a canonical filtration of $\mathcal{L}$. Then the following holds:*

1. *$\mathcal{F}$ is the unique canonical filtration of $\mathcal{L}$ and defines a filtration in the sense of 3.23.*
2. *The dual filtration $\mathcal{F}^*$ is the canonical filtration of $\mathcal{L}^*$.*
3. *$\mathrm{nd}((\mathcal{L}_i \cap W_i)/W_{i-1}) < \mathrm{nd}((\mathcal{L} \cap W_{i+1})/W_i)$, $i \in [k-1]$.*
4. *The lattice $(\mathcal{L}_i \cap W_i)/W_{i-1}$, $i \in [k]$, is stable.*

3.5 Gaussian Measures

Definition 3.27 (Gaussian Mass)**.** *Define $\rho(\mathbf{x}) = e^{-\pi\|\mathbf{x}\|^2}$ for $\mathbf{x} \in \mathbb{R}^n$. We extend this definition to any countable set $T \subseteq \mathbb{R}^n$ by $\rho(T) = \sum_{\mathbf{x}\in T} \rho(\mathbf{x})$.*

Definition 3.28 (Poisson Summation Formula)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$ and shift $\mathbf{t} \in \mathrm{span}(\mathcal{L})$, we have that*

$$\rho(\mathcal{L} + \mathbf{t}) = \frac{1}{\det(\mathcal{L})} \sum_{\mathbf{y}\in\mathcal{L}^*} e^{2\pi i\langle \mathbf{y},\mathbf{t}\rangle}\rho(\mathbf{y}) \,.$$

We will need the following standard lemmas. We include proofs for completeness.

Lemma 3.29. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice. Then for $s \geq 1$,*

$$\rho(s\mathcal{L} \setminus \{\mathbf{0}\}) \leq \rho(\mathcal{L} \setminus \{\mathbf{0}\})^{s^2} \,.$$

Proof.

$$\begin{aligned}\rho(s\mathcal{L} \setminus \{\mathbf{0}\}) &= \sum_{\mathbf{y}\in\mathcal{L}\setminus\{\mathbf{0}\}} \rho(\mathbf{y})^{s^2} \leq (\sum_{\mathbf{y}\in\mathcal{L}\setminus\{\mathbf{0}\}} \rho(\mathbf{y}))^{s^2} \quad (\text{ since } s \geq 1)\\ &= \rho(\mathcal{L} \setminus \{\mathbf{0}\})^{s^2} \,.\end{aligned}$$

□

Lemma 3.30. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice. Then*

$$\max_{\mathbf{t}\in\mathrm{span}(\mathcal{L})} \rho(\mathcal{L} + \mathbf{t}) = \rho(\mathcal{L}) \,.$$

Furthermore, $\rho(\mathcal{L}) = \frac{\rho(\mathcal{L}^*)}{\det(\mathcal{L})} \geq \frac{1}{\det(\mathcal{L})}$.

Proof. For $\mathbf{t} \in \mathrm{span}(\mathcal{L})$, by the Poisson summation formula

$$\rho(\mathcal{L} + \mathbf{t}) = \frac{1}{\det(\mathcal{L})} \sum_{\mathbf{y}\in\mathcal{L}^*} e^{2\pi i\langle \mathbf{y},\mathbf{t}\rangle}\rho(\mathbf{y}) \leq \frac{1}{\det(\mathcal{L})} \sum_{\mathbf{y}\in\mathcal{L}^*} \rho(\mathbf{y}) = \frac{\rho(\mathcal{L}^*)}{\det(\mathcal{L})} = \rho(\mathcal{L}) \,,$$

as needed. The furthermore now follows from the above and the fact that $\rho(\mathcal{L}^*) \geq \rho(0) = 1$. □

Lemma 3.31. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice and W be a lattice subspace of $\mathcal{L}^*$. Then*

$$\rho(\mathcal{L}) \leq \rho(\pi_W(\mathcal{L}))\rho(\mathcal{L} \cap W^\perp) \,.$$

Proof. For $\mathbf{a} \in \pi_W(\mathcal{L})$, we have that $\pi_W^{-1}(\mathbf{a}) \cap \mathcal{L} = (\mathbf{a} + W^\perp) \cap \mathcal{L}$. Choose a representative $\hat{\mathbf{a}} \in (\mathbf{a} + W^\perp) \cap \mathcal{L}$, and note now that $(\mathbf{a} + W^\perp) \cap \mathcal{L} = \hat{\mathbf{a}} + (\mathcal{L} \cap W^\perp)$ and that $\hat{\mathbf{a}} - \mathbf{a} \in W^\perp$. From here, we see that

$$\begin{aligned}\rho(\mathcal{L}) &= \rho(\bigcup_{\mathbf{a}\in A} (\hat{\mathbf{a}} + (\mathcal{L} \cap W^\perp))) = \sum_{\mathbf{a}\in\pi_W(\mathcal{L})} \rho(\hat{\mathbf{a}} + (\mathcal{L} \cap W^\perp))\\ &= \sum_{\mathbf{a}\in\pi_W(\mathcal{L})} \rho(\mathbf{a})\rho((\hat{\mathbf{a}} - \mathbf{a}) + (\mathcal{L} \cap W^\perp)) \quad \left(\text{ since } \mathbf{a} \perp W^\perp \right)\\ &\leq \sum_{\mathbf{a}\in\pi_W(\mathcal{L})} \rho(\mathbf{a})\rho(\mathcal{L} \cap W^\perp) \quad (\text{ by Lemma 3.30 })\\ &= \rho(\pi_W(\mathcal{L}))\rho(\mathcal{L} \cap W^\perp) \,,\end{aligned}$$

as needed. □

Definition 3.32 (Discrete Gaussian)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$, we define the discrete Gaussian $D_{\mathcal{L}}$ to be the distribution satisfying $\Pr_{\mathbf{x}\sim D_T}[\mathbf{x} = \mathbf{y}] = \rho(\mathbf{y})/\rho(\mathcal{L}), \forall \mathbf{y} \in \mathcal{L}$.*

3.6 Lattice Problems and Algorithms

Definition 3.33 (Gap Covering Radius Problem)**.** *For $\gamma := \gamma(n) \geq 1$, the γ-GapCRP is given as input a lattice $\mathcal{L} := \mathcal{L}(B)$, $B \in \mathbb{Q}^{n\times n}$, and a number t, to reject if $\mu(\mathcal{L}) \leq t$ and accept if $\mu(\mathcal{L}) > \gamma t$.*

Theorem 3.34 (DGS Sampler)**.** *For an n-dimensional lattice $\mathcal{L}$ and $\kappa \geq 1$, there is a $2^{n+\mathrm{polylog}(\kappa)+o(n)}$-time and -space algorithm that generates $2^{n/2}$ samples within statistical distance $2^{-\kappa n}$ of i.i.d. $D_{\mathcal{L}}$.*

3.7 Convex Geometry

A major conjecture in convex geometry is the Bourgain's slicing conjecture:

Conjecture 3.35. *Does every symmetric convex body of volume 1 admit a hyperplane section whose volume is lower bounded by a constant? Equivalently, for any convex body $K \subseteq \mathbb{R}^n$ is $L_K = O(1)$.*

The slicing conjecture is known to hold for many classes of convex bodies. In particular, it is known for bodies of bounded outer volume ratio.

Definition 3.36 (Outer Volume Ratio)**.** *For a symmetric convex body $K \subseteq \mathbb{R}^n$, define its outer volume ratio as*

$$\mathrm{vr}^\circ(K) = \min\{ \left(\frac{\mathrm{vol}_n(E)}{\mathrm{vol}_n(K)}\right)^{\frac{1}{n}} : K \subseteq E, E \text{ ellipsoid } \}$$

We show the easy connection between the slicing constant the volume ratio below for completeness.

Lemma 3.37. *For a symmetric convex body $K \subseteq \mathbb{R}^n$, $L_K \leq \frac{(1+o(1))}{\sqrt{2\pi e}}\mathrm{vr}^\circ(K)$.*

Proof. Let $E \supseteq K$ be the ellipsoid attaining the outer volume ratio. Let T be the linear transformation such that TK has volume 1 and for which $TE = s\mathbb{B}_2^n$ is a scaling of the Euclidean ball. We now that have that

$$L_K^2 \leq \mathop{\mathbb{E}}_{X\sim TK}[\|X\|^2]/n \leq s^2/n\,,$$

since $TK \subseteq s\mathbb{B}_2^n$. Since TK has volume one and by affine invariance $TE = s\mathbb{B}_2^n$ attains the volume outer ratio for TK, we have that $\mathrm{vr}^\circ(K) = \mathrm{vol}_n(s\mathbb{B}_2^n)^{1/n} = s\omega_n^{1/n}$. Therefore

$$s^2/n = (\omega_n^{-2/n}/n)\mathrm{vr}^\circ(K)^2 \leq \frac{1+o(1)}{2\pi e}\cdot \mathrm{vr}^\circ(K)^2\,,$$

as needed. □

The best known general upper bound on the slicing constant is $O(n^{1/4})$ due to Klartag [Kla06], which improved upon the bound of $O(n^{1/4}\log n)$ due to Bourgain [Bou91].

4 Finding Dense Sublattices

A natural question is whether one can efficiently find a subspace that whose normalized determinant to within a small factor of $\tau(\mathcal{L})$? In this section, we show that up to an approximation factor of $O(C_\eta(n))$, the reverse Minkowski constant, the answer is affirmative. To formalize this, we define the following computational problem:

Definition 4.1 (Densest Sublattice Problem)**.** *For $\gamma := \gamma(n)$, the γ-DSP is given an n-dimensional lattice $\mathcal{L} := \mathcal{L}(\mathbf{B})$, $\mathbf{B} \in \mathbb{Q}^{n\times n}$, to compute a basis of a sublattice $M \subseteq \mathcal{L}$, $M \neq \{\mathbf{0}\}$, satisfying $\mathrm{nd}(M) \leq \gamma\tau(\mathcal{L})$.*

We note that the problem of exactly computing the k-dimensional sublattice of minimum determinant was previously considered in [DM13]. The corresponding algorithm required $k^{O(kn)}$ time, which used exhaustive enumeration to find a basis generating this densest sublattice. In contrast, here we show that if we are looking for the densest sublattice across all dimensions and are willing to tolerate a $O(\log n)$ approximation factor, then there is a very natural procedure for finding an approximate densest sublattice in single exponential time.

Theorem 4.2. *For $\alpha \geq 1$, $\kappa \geq 1$, there is a $2^{1.6n+\log(\kappa)+o(n)}$-time and $2^{n+o(n)}$-space algorithm which solves $2C_\eta(n)$-DSP on n-dimensional lattices with probability at least $1 - 2^{-\kappa n}$.*

The algorithm will consist of calling the subroutine Algorithm 1 repeatedly. The main idea will be iteratively induct on a random co-dimension 1 sublattice chosen from a discrete Gaussian distribution.

Function `DenseSub`$(\mathcal{L}, \varepsilon)$
  Input: Lattice $\mathcal{L}$, $n := \dim(\mathcal{L}) \geq 1$, error parameter $\varepsilon > 0$
  Output: Approximately densest sublattice of $\mathcal{L}$
1 **if** $n \geq 2$ **then**
2 Sample $\mathbf{X}$ within statistical distance $\varepsilon/2^n$ from $D_{\mathrm{nd}(\mathcal{L})\mathcal{L}^*/2}$.
3 **if** $\mathbf{X} \neq \mathbf{0}$ **then**
4 $M \leftarrow$ `DenseSub`$(\mathcal{L} \cap \mathbf{X}^\perp, \varepsilon)$.
5 **if** $\mathrm{nd}(M) < \mathrm{nd}(\mathcal{L})$ **then return** M.
6 **return** $\mathcal{L}$.

Algorithm 1: Finding a Dense Lattice Subspace

The following lemma yields the core of the analysis of the algorithm. In particular, it says that if the current lattice $\mathcal{L}$ is not an approximate densest sublattice, then a sample from the appropriate discrete Gaussian on the dual lattice will be orthogonal to any densest sublattice with lower bounded probability. In the lemma below, recall that the first non-trivial subspace of the canonical filtration is the span of the inclusion-wise largest densest sublattice of $\mathcal{L}$.

Lemma 4.3. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a k-dimensional lattice and let $\mathcal{F}_\mathcal{L} = (W_0, \ldots, W_l)$ denote its canonical filtration. Assume that $stC_\eta(k)\tau(\mathcal{L}) \leq \mathrm{nd}(\mathcal{L})$, for $s, t \geq 1$. Then,*

$$\Pr_{\mathbf{X}\sim D_{\mathrm{nd}(\mathcal{L})\mathcal{L}^*/s}}[\mathbf{X} \neq 0, \mathbf{X} \perp W_1] \geq \frac{1}{1+2^{-t^2}} - s^{-k}.$$

Proof. Noting that the statement is invariant under scaling $\mathcal{L}$, we may assume that $\det(\mathcal{L}) = \det(\mathcal{L}^*) = 1$. Given our assumption that $\mathcal{L}$ is not a densest sublattice, we have that the canonical filtration is non-trivial and hence $\{\mathbf{0}\} \subset W_1 \subset \mathcal{L}$. Recall that by definition $\mathrm{nd}(\mathcal{L} \cap W_1) = \tau(\mathcal{L})$.

We begin by bounding the probability of the bad event $\mathbf{X} = 0$. For this event, we have that

$$\Pr_{\mathbf{x} \sim D_{\mathcal{L}^*/s}}[\mathbf{X} = \mathbf{0}] = \frac{1}{\rho(\mathcal{L}^*/s)} \underbrace{\leq}_{\text{Lemma 3.30}} \det(\mathcal{L}^*/s) = s^{-k}, \tag{4}$$

where the last equality follows by assumption that $\det(\mathcal{L}^*) = 1$ and $\dim(\mathcal{L}^*) = k$.

We now lower bound the probability that $\mathbf{X} \perp W_1$. We first lower bound this probability by the reciprocal of the Gaussian mass of the projection of $\mathcal{L}^*$ onto W_1:

$$\Pr_{\mathbf{x} \sim D_{\mathcal{L}^*/s}}[\mathbf{X} \perp W_1] = \frac{\rho((\mathcal{L}^* \cap W_1^\perp)/s)}{\rho(\mathcal{L}^*/s)} \underbrace{\geq}_{\text{Lemma 3.31}} \frac{\rho((\mathcal{L}^* \cap W^\perp)/s)}{\rho((\mathcal{L}^* \cap W_1^\perp)/s)\rho(\pi_{W_1}(\mathcal{L}^*)/s)} = \frac{1}{\rho(\pi_{W_1}(\mathcal{L}^*)/s)}. \tag{5}$$

It now suffices to derive an appropriate upper bound on the Gaussian mass of $\pi_{W_1}(\mathcal{L}^*)/s$, for which we will use the reverse Minkowski theorem. Recall that the dual filtration

$$\mathcal{F}^*_{\mathcal{L}} = (\{\mathbf{0}\}, \ldots, W_1^\perp \cap \mathrm{span}(\mathcal{L}^*), \mathrm{span}(\mathcal{L}^*)).$$

In particular, the "last block" of the dual filtration corresponds to the lattice $\mathcal{L}^*/W_1^\perp = \pi_{W_1}(\mathcal{L}^*)$ which is non-trivial by assumption as well as stable by Lemma 3.26. Furthermore, $\mathrm{nd}(\pi_{W_1}(\mathcal{L}^*)) = 1/\mathrm{nd}(\mathcal{L} \cap W_1) \geq stC_\eta(k)$ since $\det(\mathcal{L}) = 1$. Since $\mathcal{L}^*$ is k-dimensional, by the reverse Minkowski theorem we deduce that

$$\rho(\pi_{W_1}(\mathcal{L}^*)/st) \leq 3/2 \Rightarrow \rho(\pi_{W_1}(\mathcal{L}^* \setminus \{\mathbf{0}\})/st) \leq 1/2.$$

Given the above, we can deduce the bound

$$\rho(\pi_{W_1}(\mathcal{L}^*)/s) = 1 + \rho(\pi_{W_1}(\mathcal{L}^* \setminus \{\mathbf{0}\})/s) \underbrace{\leq}_{\text{Lemma 3.29}} 1 + \rho(\pi_{W_1}(\mathcal{L}^* \setminus \{\mathbf{0}\})/st)^{t^2} \leq 1 + 2^{-t^2}. \tag{6}$$

Combining (4),(5) and (6), we derive the desired bound

$$\begin{aligned}\Pr[\mathbf{X} \neq 0, \mathbf{X} \perp W_1] &\geq \Pr[\mathbf{X} \perp W_1] - \Pr[\mathbf{X} = 0] \\ &\geq \frac{1}{\rho(\pi_{W_1}(\mathcal{L}^*)/s)} - s^{-k} \geq \frac{1}{1 + 2^{-t^2}} - s^{-k}.\end{aligned}$$

□

We can now prove Theorem 4.2. We note that the constant 2 in the approximation factor $2C_\eta(n)$ was chosen to simplify the analysis and has not been optimized.

Proof of Theorem 4.2. Given an n-dimensional lattice $\mathcal{L}$ and $\kappa \geq 1$, the full algorithm makes $T := 2^{6n+o(n)}\kappa n = 2^{6n+\log\kappa+o(n)}$ calls to $\texttt{DenseSub}(\mathcal{L}, 2^{-n})$ and returns the densest sublattice found over all calls. Each call to $\texttt{DenseSub}(\mathcal{L}, 2^{-n})$ induces at most n recursive calls, one for each dimension in $[n]$. The work done at each level is dominated by the call to the DGS sampler (Theorem 3.34), with statistical distance requirement $2^{-n-d} \geq 4^{-n}$ on a lattice of dimension $d \in [n]$.

The total work per call to `DenseSub` is thus at most $n2^{n+o(n)} = 2^{n+o(n)}$ time using $2^{n+o(n)}$-space and hence the total runtime of the algorithm is at most $T2^{n+o(n)} = 2^{1.6n+o(n)+\log\kappa}$ with space usage $2^{n+o(n)}$, as needed.

To prove correctness, we will show that each run of Algorithm 1 has probability at least $2^{-.6n-o(n)}$ of finding a $2C_\eta(n)$-approximate densest sublattice. Given this, the probability that we fail to find the desired sublattice over the course of the algorithm is at most $(1-2^{-.6n-o(n)})^T \leq e^{-\kappa n}$ as needed.

To prove the desired probability bound, we first assume that Algorithm 1 has access to perfect DGS samples. Let $P(n)$, $n \geq 1$, denote the worst-case probability that $\texttt{DenseSub}(\Lambda, 0)$ (i.e. perfect DGS samples) outputs a $2C_\eta(n)$-approximate densest sublattice on a n-dimensional lattice Λ. First note that $P(1) = 1$, since a one dimensional lattice is always trivially its own densest sublattice. We now claim that $P(n) \geq (\frac{2}{3} - 1/2^n)P(n-1)$, for $n \geq 2$. Let Λ be a n dimensional lattice. If Λ is already a $2C_\eta(n)$-approximate densest sublattice, $\texttt{DenseSub}(\Lambda, 0)$ clearly outputs correctly with probability 1, so assume not. Let W denote the first non-trivial subspace in the canonical filtration of Λ, recalling that this is the span of the inclusion-wise largest densest sublattice of Λ. From here, applying Lemma 4.3 with $t = 1$ and $s = 2$, we see that in the first recursion level of $\texttt{DenseSub}(\Lambda, 0)$, the computed DGS Sample $\mathbf{X} \sim D_{\mathrm{nd}(\Lambda)\Lambda^*/2}$ satisfies (a) $\mathbf{X} \neq 0, \mathbf{X} \perp W$ with probability at least $\frac{1}{1+2^{-t^2}} - s^{-n} = \frac{2}{3} - 2^{-n}$. Conditioning on this event, the next recursive call is to $\texttt{DenseSub}(\Lambda \cap \mathbf{X}^\perp, 0)$ where $\dim(\Lambda \cap \mathbf{X}^\perp) = n-1$ and $W \subseteq \mathrm{span}(\Lambda \cap \mathbf{X}^\perp)$. In particular, $\tau(\Lambda) = \tau(\Lambda \cap \mathbf{X}^\perp)$. Thus, conditioning on the event (a), if the recursive call to $\Lambda \cap \mathbf{X}^\perp$ succeeds then so does the call to $\texttt{DenseSub}(\Lambda, 0)$ (noting that $2C_\eta(n) \geq 2C_\eta(n-1)$). Since the probability of success of the recursive call is at least $P(n-1)$, the total probability of success is at least $(\frac{2}{3} - 2^{-n})P(n-1)$, as needed. Expanding the recurrence relation, a straightforward calculation reveals that $P(n) \geq \prod_{i=2}^{n}(\frac{2}{3} - 2^{-i}) \geq \frac{2}{3}^{n-1}/4 \geq 2^{-.6n-o(n)}$. To prove the success probability lower bound with imperfect DGS samples, note that the statistical distance with the perfect sampling case is at most $\sum_{i=2}^{n} 2^{-n}/2^i \leq 2^{-n}$ and hence the probability of success remains at least $2^{-.6n-o(n)} - 2^{-n} = 2^{-.6n-o(n)}$, as needed. □

5 Characterizing the Covering Radius via Slicing

In this section, we show that we can certify lower bounds on the covering radius that are tight up to the slicing constant of stable Voronoi cells. We begin with the relevant definitions.

Definition 5.1. *For $n \in \mathbb{N}$, define*

1. $L_n^{\mathrm{stable}} = \max\{L_{\mathcal{V}} : \mathcal{V}$ *is the Voronoi cell of a stable lattice of dimension at most* $n\}$.
2. $C_{\overline{KL},2}^{\mathrm{stable}}(n) = \max\{\bar{\mu}(\mathcal{L})/\sqrt{d} : \mathcal{L}$ *is a* $d \leq n$ *dimensional stable lattice of determinant* $1\}$.

An important result of [RSD17] is equivalent characterization of the worst-case $\bar{\mu}$ bound for stable lattices as the worst-case slicing constant of their Voronoi cells.

Theorem 5.2. *[RSD17]* $L_n^{\mathrm{stable}} = C_{\overline{KL},2}^{\mathrm{stable}}(n)$.

We will need the following useful extension of the average $\bar{\mu}$ which takes in an additional inner product.

Definition 5.3 (Generalized $\bar{\mu}$)**.** *For a lattice $\mathcal{L} \subset \mathbb{R}^n$ and $Q \in \mathbb{S}^n_+$, $\mathrm{im}(Q)$ a lattice subspace of $\mathcal{L}^*$, define $d(\mathbf{x}, \mathcal{L}, Q) = \min\{\|\mathbf{x} - \mathbf{y}\|_Q : \mathbf{y} \in \mathcal{L}\}$. Define*

$$\bar{\mu}(\mathcal{L}, Q)^2 = \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{span}(\mathcal{L})/\mathcal{L}}[d(\mathbf{x}, \mathcal{L}, Q)^2] \,.$$

The first lemma we will need is slightly generalized volumetric lower bound for covering radius under a general ellipsoidal (semi-)norm.

Lemma 5.4. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice and $Q \in \mathbb{S}^n_+$, $W := \mathrm{im}(Q)$ a lattice subspace of $\mathcal{L}^*$, $d := \dim(W)$. Then*

$$\bar{\mu}(\mathcal{L}, Q)^2 \geq d L_{\mathbb{B}_2^d} \cdot \det(Q)_W^{1/d} \det(\mathcal{L}/W^\perp)^{2/d} \,.$$

Proof. Let $Q = \mathbf{B}\mathbf{B}^\mathsf{T}$ denote the Cholesky factorization of Q, where $\mathbf{B}$ is nonsingular. Note that $\mathrm{im}(Q) = \mathrm{span}(\mathbf{B})$ and hence $\ker(\mathbf{B}^\mathsf{T}) = \mathrm{im}(Q)^\perp = \ker(Q)$. In particular, $\ker(\mathbf{B}^\mathsf{T}) \cap \mathrm{span}(\mathcal{L})$ is a lattice subspace of $\mathcal{L}$. Thus $\Lambda = \mathbf{B}^\mathsf{T}\mathcal{L}$ is discrete and hence a lattice. We now claim that $\bar{\mu}(\mathcal{L}, Q)^2 = \bar{\mu}(\mathbf{B}^\mathsf{T}\mathcal{L})$. Firstly, by construction $\|\mathbf{x} - \mathbf{y}\|_Q = \|\mathbf{B}^\mathsf{T}(\mathbf{x} - \mathbf{y})\|_2$. Next, it is easy to check that for $\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})$, that $\mathbf{B}^\mathsf{T}\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\Lambda)/\Lambda)$. Therefore

$$\mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})}[d(\mathbf{x}, \mathcal{L}, Q)^2] = \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})}[d(\mathbf{B}^\mathsf{T}\mathbf{x}, \mathbf{B}^\mathsf{T}\mathcal{L})^2] = \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\Lambda)/\Lambda)}[d(\mathbf{x}, \Lambda)^2] \,.$$

From here, letting $d = \dim(\Lambda)$, we have that

$$\mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\Lambda)/\Lambda)}[d(\mathbf{x}, \Lambda)^2] = \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathcal{V}(\Lambda))}[\|\mathbf{x}\|^2] \geq d L^2_{\mathbb{B}_2^d} \mathrm{vol}_d(\mathcal{V}(\Lambda))^{2/d} = d L^2_{\mathbb{B}_2^d} \det(\Lambda)^{2/d} \,.$$

We now conclude the proof by the identity:

$$\det(\Lambda)^2 = \det(\mathbf{B}^\mathsf{T}\mathbf{B}) \det(\mathcal{L}/W^\perp)^2 = \det_W(Q) \det(\mathcal{L}/W^\perp)^2 \,.$$

□

The next lemma establishes a crucial *concavity* property for $\bar{\mu}$, establishing very useful lower bounds on the covering radius. This bound is in fact already implicit in [DR16], but somewhat hidden, so we give a direct and simple proof here.

Lemma 5.5. *Let $\mathcal{L} \subset \mathbb{R}^n$ be a lattice. Let $Q_1, \ldots, Q_k \in \mathbb{S}^n_+$ such that $W_i := \mathrm{im}(Q_i)$, $i \in [k]$, is a lattice subspace of $\mathcal{L}^*$, where $d_i := \dim(W_i)$, satisfying $\sum_{i=1}^k Q_i \preceq I_n$. Then*

$$\bar{\mu}(\mathcal{L})^2 \geq \sum_{i=1}^k \bar{\mu}(\mathcal{L}, Q_i)^2 \gtrsim \sum_{i=1}^k d_i \det_{W_i}(Q_i)^{1/d_i} \det(\mathcal{L}/W_i^\perp)^{2/d_i} \,.$$

Proof. We prove the first inequality. The second inequality is an immediate consequence of lemma 5.4 combined with the fact that $L_{\mathbb{B}_2^n} \geq \frac{1-o(1)}{\sqrt{2\pi e}}$. For the first inequality, we have that

$$\begin{aligned}
\bar{\mu}(\mathcal{L})^2 &= \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})}[d(\mathbf{x}, \mathcal{L})^2] = \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})}[\min_{\mathbf{y} \in \mathcal{L}} \|\mathbf{x} - \mathbf{y}\|_2^2] \geq \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})}[\min_{\mathbf{y} \in \mathcal{L}} \sum_{i=1}^k \|\mathbf{x} - \mathbf{y}\|_{Q_i}^2] \\
&\geq \mathop{\mathbb{E}}_{\mathbf{x} \sim \mathrm{unif}(\mathrm{span}(\mathcal{L})/\mathcal{L})}[\sum_{i=1}^k \min_{\mathbf{y} \in \mathcal{L}} \|\mathbf{x} - \mathbf{y}\|_{Q_i}^2] = \sum_{i=1}^k \bar{\mu}(\mathcal{L}, Q_i)^2 \,,
\end{aligned}$$

as needed. □

The next technical lemma is the crucial ingredient to our improved analysis of the approximability of the covering radius. Here, we use the power of the above "dual program", to derive useful lower bounds from any lattice filtration. We note that the upper bound give below comes from a direct "Babai style" analysis over the orthogonal blocks $(\mathcal{L} \cap W_i)/W_{i-1}, i \in [k]$ and is standard.

Lemma 5.6. *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice, and let $\{\mathbf{0}\} = W_0 \subset W_1 \subset W_2 \subset W_k = \mathbb{R}^n$ be a filtration of $\mathcal{L}$ with $d_i := \dim(W_i)$, $i \in [k]$. Then*

$$\mu(\mathcal{L})^2 \leq \sum_{i=1}^{k} (d_i - d_{i-1}) \mu((\mathcal{L} \cap W_i)/W_{i-1})^2 \,, \tag{7}$$

and

$$\mu(\mathcal{L})^2 \gtrsim \sum_{i=1}^{k} (d_i - d_{i-1}) \det((\mathcal{L}/W_{i-1}))^{2/(n-d_{i-1})} \,. \tag{8}$$

Furthermore, if

$$\det((\mathcal{L} \cap W_i)/W_{i-1})^{1/(d_i - d_{i-1})} \leq \det((\mathcal{L} \cap W_{i+1})/W_i)^{1/(d_{i+1} - d_i)} \quad i \in \{1, \ldots, k-1\} \,, \tag{9}$$

then

$$\mu(\mathcal{L})^2 \gtrsim \sum_{i=1}^{k} (d_i - d_{i-1}) \det((\mathcal{L} \cap W_i)/W_{i-1})^{2/(d_i - d_{i-1})} \,. \tag{10}$$

To prove the lemma we will need the following fact about numbers.

Lemma 5.7. *Let $0 = l_0 < l_1 < \cdots < l_k = n$. Then there exists non-negative numbers $\lambda_{ij} \geq 0$, $i, j \in [k]$, $\lambda_{ij} = 0$ if $j > l_i$, satisfying*

$$\begin{array}{lll} 1) & l_i \cdot (\prod_{j=1}^{l_i} \lambda_{ij})^{1/l_i} & \gtrsim l_i - l_{i-1} \quad \forall i \in [k] \\ 2) & \sum_{i=1}^{k} \lambda_{ij} & \leq 1 \quad \forall j \in [n] \,. \end{array}$$

Proof. To begin, we consider the case where $k = n$ and $l_i = i, \forall i \in [n]$. We will deduce the solution for general case from this solution. For the special case, we use the solution

$$\hat{\lambda}_{ij} = j/(2i^2), \quad j \leq i \,.$$

To check feasibility, for property 1 we have that

$$i \cdot (\prod_{j=1}^{i} \hat{\lambda}_{ij})^{1/i} = i \cdot \frac{(i!)^{1/i}}{2i^2} \geq 1/(2e),$$

where the last inequality follows by Stirling's approximation. For property 2, we have that

$$\sum_{i=1}^{n} \hat{\lambda}_{ij} = \sum_{i=j}^{n} j/(2i^2) \leq j \int_{j}^{\infty} 1/x^2 dx = 1,$$

as needed.

For the general case, the strategy is simply to "collapse" the global solution down. Precisely, we will use the solution

$$\lambda_{ij} = \sum_{r=l_{i-1}+1}^{l_i} \hat{\lambda}_{rj} \,.$$

Since by construction $\sum_{i=1}^{k} \lambda_{ij} = \sum_{i=1}^{n} \hat{\lambda}_{ij} \leq 1$, we clearly satisfy constraint 2. Therefore it suffices to check constraint 1. For this, we have that

$$\begin{aligned}
l_i \cdot (\prod_{j=1}^{l_i} \lambda_{ij})^{1/l_i} &= l_i \cdot (\prod_{j=1}^{l_i} (\sum_{r=l_{i-1}+1}^{l_i} \hat{\lambda}_{rj}))^{1/l_i} = l_i \cdot (\prod_{j=1}^{l_i} (\frac{j}{2} \cdot \sum_{r=\max\{l_{i-1}+1,j\}}^{l_i} \frac{1}{r^2}))^{1/l_i} \\
&\geq \frac{l_i (l_i!)^{1/l_i}}{2} \cdot (\prod_{j=1}^{l_i} \int_{\max\{l_{i-1}+1,j\}}^{l_i+1} \frac{1}{x^2} dx)^{1/l_i} \\
&\geq \frac{l_i^2}{2e} \cdot (\prod_{j=1}^{l_{i-1}} \frac{l_i - l_{i+1}}{(l_{i-1}+1)(l_i+1)} \cdot \prod_{j=l_{i-1}+1}^{l_i} \frac{l_i+1-j}{(l_i+1)j})^{1/l_i} \\
&= \frac{l_i^2}{2e} \cdot (\frac{(l_i - l_{i+1})^{l_{i-1}}}{(l_{i-1}+1)^{l_{i-1}}(l_i+1)^{l_{i-1}}} \cdot \frac{(l_i - l_{i-1})!}{(l_i+1)^{l_i - l_{i-1}}} \cdot \frac{l_{i-1}!}{l_i!})^{1/l_i} \\
&= \frac{l_i - l_{i-1}}{2e} \cdot \frac{l_i^2}{(l_i+1)l_i} (\frac{l_i^{l_i}}{l_i!} \cdot \frac{(l_i - l_{i-1})!}{(l_i - l_{i-1})^{l_i - l_{i-1}}} \cdot \frac{l_{i-1}!}{(l_{i-1}+1)^{l_{i-1}}})^{1/l_i} \\
&\gtrsim l_i - l_{i-1} \quad (\text{ by Sterling's approximation }) \,,
\end{aligned}$$

as needed.

□

Proof of Theorem 5.6. The proof of the upper bound is standard and can be found for example in [RSD17, Proposition 6.4]. We thus focus on the lower bound.

For the lower bound the strategy will be to construct matrices $Q_1, \ldots, Q_k$ satisfying the conditions of 5.5 achieving a value that is within a constant factor of the desired bound in 8.

Let $V_i := W_{k-i}^{\perp}$, $l_i := \dim(V_i) = n - d_i$, $i \in \{0, \ldots, k\}$ where $\{\mathbf{0}\} = V_0 \subset V_1 \subset \cdots \subset V_k = \mathbb{R}^n$ are lattice subspaces of $\mathcal{L}^*$. Let us now apply an orthogonal transformation to such that the subspaces $V_0, V_1, \ldots, V_k$ align with the coordinate subspaces. This allows to without loss of generality assume that $V_i = \text{span}(e_1, \ldots, e_{l_i})$, where $e_1, \ldots, e_n$ are the standard basis. Now let $\lambda_{ij}, i, j \in [n]$ be defined as in Lemma 5.7 with respect to $l_1, \ldots, l_k$. Let $D_i = \text{diag}(\lambda_{i1}, \ldots, \lambda_{in})$, $i \in [k]$, denote the corresponding diagonal matrices. By the properties guaranteed by Lemma 5.7, we directly get that $\sum_{i=1}^{k} D_i \preceq I_n$ and that $\text{im}(D_i) = V_i$, $i \in [k]$. Therefore by Lemma 5.5, we have

that

$$\begin{aligned}
\mu(\mathcal{L})^2 &\geq \bar{\mu}(\mathcal{L})^2 \\
&\gtrsim \sum_{i=1}^{k} l_i \det_{V_i}(D_i)^{1/l_i} \det(\mathcal{L}/V_i^{\perp})^{2/l_i} \\
&= \sum_{i=1}^{k} l_i (\prod_{j=1}^{l_i} \lambda_{ij})^{1/l_i} \det(\mathcal{L}/V_i^{\perp})^{2/l_i} \\
&\gtrsim \sum_{i=1}^{k} (l_i - l_{i-1}) \det(\mathcal{L}/V_i^{\perp})^{2/l_i} \quad \text{(by property 1)} \\
&= \sum_{i=1}^{k} (d_{k-i+1} - d_{k-i}) \det(\mathcal{L}/W_{k-i})^{2/(n-d_{k-i})} \\
&= \sum_{i=1}^{k} (d_i - d_{i-1}) \det(\mathcal{L}/W_{i-1})^{2/(n-d_{i-1})} ,
\end{aligned}$$

as needed.

For the furthermore, we have that

$$\begin{aligned}
\mu(\mathcal{L})^2 &\gtrsim \sum_{i=1}^{k} (d_i - d_{i-1}) \det(\mathcal{L}/W_{i-1})^{2/(n-d_{i-1})} \\
&= \sum_{i=1}^{k} (d_i - d_{i-1}) \left(\prod_{j=i}^{k} \det((\mathcal{L} \cap W_j)/W_{j-1})^2 \right)^{1/(n-d_{i-1})} \\
&= \sum_{i=1}^{k} (d_i - d_{i-1}) \prod_{j=i}^{k} \left(\det((\mathcal{L} \cap W_j)/W_{j-1})^{2/(d_j - d_{j-1})} \right)^{(d_j - d_{j-1})/(n-d_{i-1})} \\
&\geq \sum_{i=1}^{k} (d_i - d_{i-1}) \det((\mathcal{L} \cap W_j)/W_{j-1})^{2/(d_i - d_{i-1})} \quad \text{(by inequality (9))} ,
\end{aligned}$$

as needed. □

The following is the main theorem of this section. The upper bound was obtained in [RSD17] while the nearly matching lower bound is novel.

Theorem 5.8. *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice. Let $\mathcal{F}_{\mathcal{L}} = (\{\mathbf{0}\} = W_0, \ldots, W_k = \mathbb{R}^n)$ be the canonical filtration of $\mathcal{L}$. Then*

$$\mu(\mathcal{L})^2 \leq 4(L_n^{\text{stable}})^2 \sum_{i=1}^{k} (d_i - d_{i-1}) \text{nd}((\mathcal{L} \cap W_i)/W_{i-1})^2 \tag{11}$$

$$\mu(\mathcal{L})^2 \gtrsim \sum_{i=1}^{k} (d_i - d_{i-1}) \text{nd}((\mathcal{L} \cap W_i)/W_{i-1})^2 . \tag{12}$$

Proof. From Lemma 3.26 the canonical filtration satisfies the inequalities in (9) and that each block $(\mathcal{L} \cap W_i)/W_{i-1}$, $i \in [k]$, is stable.

Clearly, the lower bound follows directly from Lemma 5.6. For the upper bound, by Lemma 5.6, Theorem 5.2 and stability of the blocks, we have that

$$\begin{aligned}\mu(\mathcal{L})^2 &\leq \sum_{i=1}^{k} \mu((\mathcal{L} \cap W_i)/W_{i-1})^2 \leq 4 \sum_{i=1}^{k} \bar{\mu}((\mathcal{L} \cap W_i)/W_{i-1})^2 \\ &\leq 4(L_n^{\text{stable}})^2 \sum_{i=1}^{k} (d_i - d_{i-1}) \text{nd}((\mathcal{L} \cap W_i)/W_{i-1})^2 \,,\end{aligned}$$

as needed. □

We note that since the dual of the canonical filtration is the canonical filtration of the dual, by simply "inverting" the normalized determinants in (11), we again get a nearly tight bound for $\mu(\mathcal{L}^*)$. This gives a surprisingly tight duality relation relating bounds for $\mu(\mathcal{L})$ and $\mu(\mathcal{L}^*)$.

Using the above, we derive the improved complexity classification for GapCRP as a simple corollary.

Corollary 5.9. *$O(L_n^{\text{stable}})$-GapCRP $\in$ coNP.*

Proof. Let $c_1, c_2 \geq 1$ be constants to be chosen later. Let $\gamma = c_1 L_n^{\text{stable}}$. Given an instance $(\mathcal{L}, t)$ of γ-coGapCRP for an n-dimensional lattice $\mathcal{L}$, the verifier accepts on any filtration $\{\mathbf{0}\} = W_0 \subset \cdots \subset W_k$ of $\mathcal{L}$, $d_i := \dim(W_i)$, $i \in [k]$, such that

$$c_2^2 t^2 < \sum_{i=1}^{k} (d_i - d_{i-1}) \det(\mathcal{L}/W_{i-1})^{2/(n-d_{i-1})} \,. \tag{13}$$

Clearly, this check can be performed in polynomial time, so we need only check the completeness and soundness.

Firstly, let us choose c_2 large enough, so that if the test (13) passes we are guaranteed that $\mu(\mathcal{L}) > t$. The existence of a universal constant c_2 is given by Lemma 5.6. Second, let us choose c_1 large enough, so that if $\mu(\mathcal{L}) > \gamma t$, there exists choice of filtration such that test (13) passes. The existence of a universal constant c_1 is now guaranteed by Lemma 5.6 and Theorem 5.8 by our choice of $\gamma = \Omega(L_n^{\text{stable}})$.

Given the choice of c_1, c_2 completeness and soundness are now direct. Firstly, if $(\mathcal{L}, t)$ is a NO instance, i.e. $\mu(\mathcal{L}) \leq t$, then the verifier will never accept by our choice of c_2. Second, if $(\mathcal{L}, t)$ is a YES instance, i.e. $\mu(\mathcal{L}) > \gamma t$, then by our choice of c_1 there exists a filtration on which the verifier accepts. □

6 Computing Approximately Stable Filtrations

From the previous section, a natural question is whether we can efficiently build a filtration $\mathcal{F}$ which "approximates" the canonical filtration of $\mathcal{L}$. In this section, we show that we can indeed get a meaningful approximation using the DSP algorithm from section 4. We state our notion of approximation below:

Definition 6.1 (Approximately Stable Filtrations)**.** *For an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$ and $\gamma \geq 1$, we define a filtration $\mathcal{F} = (\{0\} = W_0, W_1, \ldots, W_k = \mathbb{R}^n)$ for $\mathcal{L}$ to be γ-stable if*

1. $\mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1}) < \mathrm{nd}((\mathcal{L} \cap W_{i+1})/W_i)$ *for* $i \in [k]$.
2. $\mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1}) \leq \gamma\tau((\mathcal{L} \cap W_i)/W_{i-1})$, $i \in [k]$.
3. $\mathrm{nd}((\mathcal{L}^* \cap W_{i-1}^{\perp})/W_i^{\perp}) \leq \gamma\tau((\mathcal{L}^* \cap W_{i-1}^{\perp})/W_i^{\perp})$, $i \in [k]$.

We note stability is closed until duality, namely a filtration $\mathcal{F}$ of $\mathcal{L}$ is γ-stable $\mathcal{L}$ iff the dual filtration $\mathcal{F}^* = (W_k^{\perp}, W_{k-1}^{\perp}, \ldots, W_0^{\perp})$ of $\mathcal{L}^*$ is γ-stable. This derives from the duality relation $((\mathcal{L} \cap W_i)/W_{i-1})^* = (\mathcal{L}^* \cap W_{i-1}^{\perp})/W_i^{\perp}$.

Furthermore, it is not hard using the techniques behind Lemma 3.26 to show that if $\mathcal{F}$ is 1-stable then $\mathcal{F}$ is the canonical filtration.

Definition 6.2 (Filtration Polytope)**.** *Given a filtration $\mathcal{F} = (W_0 = \{0\}, W_1, \ldots, W_k = \mathbb{R}^n)$ of an n-dimensional lattice $\mathcal{L}$, we define the filtration polytope*

$$\mathcal{P}(\mathcal{F}) = \mathrm{conv}((\dim(W_i), \ln(\det(\mathcal{L} \cap W_i))) : i \in [k]_0) + \{(0, y) : y \geq 0\},$$

i.e. the upwards closure of the convex hull $\mathrm{conv}((\dim(W_i), \ln(\det(W_i))) : i \in [k]_0)$.

For the sake of intuition, we give a pictorial interpretation of conditions 1-3 on $\mathcal{F}$ in terms of the filtration polytope $\mathcal{P}(\mathcal{F})$:

1. The points $(\dim(W_i), \ln(\det(W_i)))$, $i \in [k]$, are the vertices of the lower hull of $P(\mathcal{F})$.
2. For $i \in [k]$, the slope of any edge of the canonical polytope $\mathcal{P}((\mathcal{L} \cap W_i)/W_{i-1})$ is at least $\ln(\mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1})) - \ln\gamma$.
3. For $i \in [k]$, the slope of any edge of the canonical polytope $\mathcal{P}((\mathcal{L} \cap W_i)/W_{i-1})$ is at most $\ln(\mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1})) + \ln\gamma$.

For conditions 2-3 above, it is useful to image overlaying the drawings of $\mathcal{P}((\mathcal{L} \cap W_i)/W_{i-1})$, $i \in [k]$, with that of $\mathcal{P}(\mathcal{F})$. Namely, for each $i \in [k]$, translate $\mathcal{P}((\mathcal{L} \cap W_i)/W_{i-1})$ by $(\dim(W_{i-1}), \ln(\det(\mathcal{L} \cap W_{i-1}))$ so that it aligns with the edge from $(\dim(W_{i-1}), \ln(\det(\mathcal{L} \cap W_{i-1}))$ to $(\dim(W_i), \ln(\det(\mathcal{L} \cap W_i)))$ on $\mathcal{P}(\mathcal{F})$.

Before providing an algorithm for computing γ-stable filtrations, we show that γ-stable filtrations also provide nearly tight upper and lower bounds for the covering radius, by adapting the proof of Theorem 5.8.

Lemma 6.3. *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice. Let $\mathcal{F} = (\{\mathbf{0}\} = W_0, \ldots, W_k = \mathbb{R}^n)$ be a γ-stable filtration of $\mathcal{L}$, where $d_i := \dim(W_i)$, $i \in [k]_0$. Then*

$$\mu(\mathcal{L})^2 \leq 4\gamma^2 (L_n^{\mathrm{stable}})^2 \sum_{i=1}^{k} (d_i - d_{i-1}) \mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1})^2 \tag{14}$$

$$\mu(\mathcal{L})^2 \gtrsim \sum_{i=1}^{k} (d_i - d_{i-1}) \mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1})^2 . \tag{15}$$

Furthermore,

$$\mu(\mathcal{L})^2 \leq 4 \log n \gamma^2 (L_n^{\mathrm{stable}})^2 \max_{i \in [k]} \dim(\mathcal{L}/W_{i-1}) \cdot \mathrm{nd}(\mathcal{L}/W_{i-1})^2. \tag{16}$$

Proof. As before, by the non-decreasing property of the normalized block determinants, the lower bound follows directly from Lemma 5.6. For the upper bound, as before by Lemma 5.6, we have that

$$\mu(\mathcal{L})^2 \leq \sum_{i=1}^{k} \mu((\mathcal{L} \cap W_i)/W_{i-1})^2 \,. \tag{17}$$

We now use γ-stability to analyze the covering radius of the individual blocks. For $i \in [k]$, let $\mathcal{F}_i = (V_{i,0}, \ldots, V_{i,k_i})$ denote the canonical filtration of $\mathcal{L}_i := (\mathcal{L} \cap W_i)/W_{i-1}$. By property 3 and γ-stability, we have that $\mathrm{nd}((\mathcal{L}_i \cap V_{i,j})/V_{i,j-1}) \leq \gamma \mathrm{nd}(\mathcal{L}_i)$ (all slopes are upper bounded), and hence by Theorem 5.8, we have that

$$\begin{aligned}
\mu(\mathcal{L}_i)^2 &\leq 4(L_n^{\text{stable}})^2 \sum_{j=1}^{k_i} (\dim(V_{i,j}) - \dim(V_{i,0})) \mathrm{nd}((\mathcal{L}_i \cap V_{i,j})/V_{i,j-1})^2 \\
&\leq 4(L_n^{\text{stable}})^2 \sum_{j=1}^{k_i} (\dim(V_{i,j}) - \dim(V_{i,0})) \gamma^2 \mathrm{nd}(\mathcal{L}_i)^2 = 4\gamma^2 (L_n^{\text{stable}})^2 \dim(\mathcal{L}_i) \mathrm{nd}(\mathcal{L}_i)^2 \,.
\end{aligned} \tag{18}$$

Inequality 14 now follows by combining (17), (18).

For the furthermore, by the stability of the filtration, recall that $\mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1})$ is increasing for $i \in [k]$. For $i \in [k]$, let $l_i := \dim(\mathcal{L}/W_{i-1})$, where have the identity

$$\dim(\mathcal{L}/W_{i-1}) = n - d_{i-1} = \sum_{j=i}^{k} d_j - d_{j-1}.$$

Given the above,

$$\begin{aligned}
\sum_{i=1}^{k} (d_i - d_{i-1}) \mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1})^2 &\leq \sum_{i=1}^{k} (d_i - d_{i-1}) \prod_{j=i}^{k} \mathrm{nd}((\mathcal{L} \cap W_j)/W_{j-1})^{2(d_j - d_{j-1})/l_i} \\
&= \sum_{i=1}^{k} \frac{d_i - d_{i-1}}{l_i} \cdot l_i \, \mathrm{nd}(\mathcal{L}/W_{i-1})^2 \\
&\leq (\sum_{i=1}^{k} \frac{d_i - d_{i-1}}{l_i}) \cdot \max_{i \in [k]} l_i \, \mathrm{nd}(\mathcal{L}/W_{i-1})^2 \\
&\leq (\sum_{i=1}^{n} \frac{1}{i}) \cdot \max_{i \in [k]} l_i \, \mathrm{nd}(\mathcal{L}/W_{i-1})^2 \\
&\leq \log n \cdot \max_{i \in [k]} l_i \, \mathrm{nd}(\mathcal{L}/W_{i-1})^2.
\end{aligned}$$

Inequality 16 now follows by combining 14 with the above. □

We now give the following simple algorithm for computing $O(C_\eta(n))$-stable filtrations. The algorithm initializes itself with a good basis, and iteratively refines the current partition in a process we call *filtration reduction*. At a high level, the algorithm alternates between refining the current filtration using a DSP algorithm and coarsening it by restricting to the vertices of the filtration polytope. The main idea is to track progress using the potential $\mathrm{vol}_2(\mathcal{P}(\mathcal{L}) \setminus \mathcal{P}(\mathcal{F}))$, i.e. the volume of the space under the current filtration polytope and above the canonical polytope.

Function `StableFiltration(`$\mathcal{L}$`,`κ`)`

Input: $\mathcal{L} := \mathcal{L}(\mathbf{B})$, $\mathbf{B} \in \mathbb{Q}^{n\times n}$, $\kappa \geq 1$

Output: $3C_\eta(n)$-stable filtration $\mathcal{F}$

$T \leftarrow \lfloor (n^2/2)\ln n/\ln(3/2)\rfloor + 1$; updates $\leftarrow 0$.

$\mathbf{B} \leftarrow (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ an HKZ basis of $\mathcal{L}$ w.p. at least $1 - 2^{-\kappa n}/(T+1)$.

$\mathcal{F} \leftarrow (\mathrm{span}(B_{[i]}) : i \in [n]_0)$.

repeat

  stable $\leftarrow$ true.

  $(W_0, \ldots, W_k) \leftarrow$ elements of $\mathcal{F}$.

  if $\exists\, i$ *s.t.* $(\dim(W_i), \ln(\det(\mathcal{L}\cap W_i)))$ *not a vertex of* $\mathcal{P}(\mathcal{F})$ **then**

    Remove all the non-vertices from $\mathcal{F}$.

    stable $\leftarrow$ false.

  end

  else

    for $i \in [k]$ **do**

      $\mathcal{L}_i \leftarrow (\mathcal{L}\cap W_i)/W_{i-1}$.

      $M_{i,l} \leftarrow 2C_\eta(n)$-approximate densest sublattice of $\mathcal{L}_i$

      w.p. at least $1 - 2^{-\kappa n}/(T+1)$.

      if $\mathrm{nd}(M_{i,l}) < (2/3)\mathrm{nd}(\mathcal{L}_i)$ **then**

        `//` $\mathcal{P}(\mathcal{L}_i)$ `has edge of slope` $< -\ln(3/2) + \ln \mathrm{nd}(\mathcal{L}_i)$

        Add $W_{i-1} + \mathrm{span}(M_{i,l})$ to $\mathcal{F}$ in between W_{i-1} and W_i.

        stable $\leftarrow$ false; updates $\leftarrow$ updates $+1$.

        break from for loop.

      end

      $M_{i,h} \leftarrow 2C_\eta(n)$-approximate densest sublattice of $\mathcal{L}_i^*$

      w.p. at least $1 - 2^{-\kappa n}/(T+1)$.

      if $\mathrm{nd}(M_{i,h}) < (2/3)\mathrm{nd}(\mathcal{L}_i^*)$ **then**

        `//` $\mathcal{P}(\mathcal{L}_i)$ `has edge of slope` $> \ln(3/2) + \ln \mathrm{nd}(\mathcal{L}_i)$

        Add $\mathrm{span}(M_{i,h})^\perp \cap W_i$ to $\mathcal{F}$ in between W_{i-1} and W_i.

        stable $\leftarrow$ false; updates $\leftarrow$ updates $+1$.

        break from for loop.

      end

    end

  end

until *stable = true or updates* $\geq T$

return $\mathcal{F}$

Algorithm 2: Computing an Approximately Stable Filtration

Theorem 6.4. *Given as input $\mathcal{L} := \mathcal{L}(B) \subset \mathbb{R}^n$, an n-dimensional lattice and parameter $\kappa \geq 1$, using $2^{1.6n+\log(\kappa)+o(n)}$ time and $2^{n+o(n)}$ space algorithm* (2) *computes a $3C_\eta(n)$-stable filtration of $\mathcal{L}$ with probability at least $1 - 2^{-\kappa n}$.*

Before proving the above, we show how to use it directly obtain an algorithm to compute KL projections that are tight up to a $O(\log^{2.5} n)$ factor.

Proof of Theorem 2.1. Applying Theorem 6.4, we compute a $\gamma(n) := 3C_\eta(n)$ stable filtration $\mathcal{F} = \{W_0, \ldots, W_k\}$ for $\mathcal{L}$ and return the maximizer of $\max_{i\in[k]} \sqrt{\dim(\mathcal{L}/W_{i-1})}\mathrm{nd}(\mathcal{L}/W_{i-1})$. The running time is immediate. For correctness, by Lemma 6.3 part (16), we have that

$$\mu(\mathcal{L}) \leq 2\gamma(n) L_n^{\text{stable}} \sqrt{\log n} \max_{i\in[k]} \sqrt{\dim(\mathcal{L}/W_{i-1})}\mathrm{nd}(\mathcal{L}/W_{i-1})$$

where $2\gamma(n)L_n^{\text{stable}}\sqrt{\log n} = 6C_\eta(n)L_n^{\text{stable}}\sqrt{\log n} = O(\log^{2.5} n)$, as needed. □

Proof of Theorem 6.4. By the design of the algorithm, the number of calls to either the densest sublattice algorithm or the HKZ basis algorithm is upper bounded by $T + 1 = O(n^2 \log n)$. Since the probability of error for each call is at most $2^{-\kappa n}/(T+1)$, the probability that any call fails during the course of the algorithm is bounded by $2^{-\kappa n}$.

We first bound the running time of the algorithm. Since we never remove non-vertex elements from $\mathcal{F}$ twice in a row, the total number of iterations of the repeat loop is bounded by $2T$. Since each call densest sublattice algorithms takes $2^{1.6n+\log(\kappa)+o(n)}$ time and the HKZ algorithm takes $2^{n+o(n)}$ time, the total running time is clearly bounded by $(2T+1)2^{1.6n+\log(\kappa)+o(n)} = 2^{1.6n+\log(\kappa)+o(n)}$ and total space usage is $2^{n+o(n)}$.

Now let us assume all the calls to the DSP algorithm and the HKZ basis algorithm return correctly, which as argued above occurs with probability at least $1 - 2^{-\kappa n}$. Under this assumption, we claim that the algorithm computes a $3C_\eta(n)$-stable filtration. Firstly, we remark that the if statement on line 7 checks condition 1, the if statement on line 16 checks condition 2, and the if statement on line 23 checks condition 3 for $3C_\eta(n)$-stability. Thus, if all the tests pass the returned $\mathcal{F}$ is indeed $3C_\eta(n)$ stable.

It remains to show that all the tests pass before the update variable gets incremented T times. For this purpose, we will rely on the potential $\mathrm{vol}_2(\mathcal{P}(\mathcal{L}) \setminus \mathcal{P}(F))$. Note that this quantity is always non-negative. Since the canonical polytope $\mathcal{P}(\mathcal{L})$ always lies below $\mathcal{P}(F)$ for any filtration, i.e. $\mathcal{P}(F) \subseteq \mathcal{P}(\mathcal{L})$, the potential is 0 iff $\mathcal{F}$ is the canonical filtration. Note that removing non-vertices from $\mathcal{F}$ doesn't change the filtration polytope, and hence the potential can only change when $\mathcal{F}$ is updated inside the for loop on line 12.

The following two claims will yield the result:

1. The initial potential is upper bounded $(n^2/2)\ln n$.
2. The potential drops by at least $\ln(3/2)$ every time an update occurs.

Assuming the above, the number of updates performed by the algorithm is upper bounded by $\lfloor (n^2/2)\ln n / \ln(3/2) \rfloor < T$, as needed. We now prove the claims.

For the first claim, recall that the initial filtration $(\mathrm{span}(B_{[i]}) : i \in [n]_0)$ are the subspaces of an HKZ basis for $\mathcal{L}$. By lemma 3.21, we have that

$$i^{-i/2} \det(B_{[i]}) \leq \tau_i(\mathcal{L}) \Rightarrow \ln(\det(B_{[i]})) - (n/2)\ln n \leq \ln \tau_i(\mathcal{L}) \,,$$

for $i \in [n]_0$. In particular, the points $\{(i, \ln\det(\tau_i(\mathcal{L})) : i \in [n]_0)\}$ are contained inside $\mathcal{P}(\mathcal{F}) - (0, (n/2)\ln n)$, the initial filtration polytope shifted down by $(n/2)\ln n$. Since the canonical polytope is the upwards closure of the convex hull of these points, by convexity of $\mathcal{P}(\mathcal{F})$, we get that $\mathcal{P}(\mathcal{L}) \subseteq \mathcal{P}(\mathcal{F}) - (0, (n/2)\ln n)$. In particular, the initial potential satisfies

$$\mathrm{vol}_2(\mathcal{P}(\mathcal{L}) \setminus \mathcal{P}(\mathcal{F})) \leq \mathrm{vol}_2((\mathcal{P}(\mathcal{F}) - (0, (n/2)\ln n)) \setminus \mathcal{P}(\mathcal{F})) = (n^2/2)\ln n\,,$$

as needed.

For a lattice subspace W of $\mathcal{L}$, we shall use the notation $p_W := (\dim(W), \ln(\det(\mathcal{L} \cap W)))$ to denote the associated point in the canonical polytope. For the second claim, let V be a subspace we wish add to $\mathcal{F}$ in between adjacent subspace W_{i-1}, W_i during an update step. Note that by construction, in either line 17 or 24, V satisfies $W_{i-1} \subset V \subset W_i$. Letting $\mathcal{F}'$ denote the updated filtration with V inserted, we claim that

$$\mathrm{vol}_2(\mathcal{P}(\mathcal{F}')) \geq \mathrm{vol}_2(\mathcal{P}(\mathcal{F})) + \mathrm{vol}_2(\Delta) \tag{19}$$

where

$$\Delta = \mathrm{conv}(p_{W_{i-1}}, p_V, p_{W_i})\,,$$

denotes the triangle induced by W_{i-1}, V, W_i. Note that since the canonical polytope contains the polytope of any filtration, (19) implies that the potential of $\mathcal{F}'$ drops by at least $\mathrm{vol}_2(\Delta)$.

We now prove (19). Since by construction $p_{W-1}, p_V, p_{W_i} \in \mathcal{P}(\mathcal{F}')$, by convexity of $\mathcal{P}(\mathcal{F}')$, it suffices to show that the triangle Δ is interior disjoint from $\mathcal{P}(\mathcal{F})$. Since the segment $[p_{W_{i-1}}, p_{W_i}]$ is an edge of $\mathcal{P}(\mathcal{F})$ (ensured by the test on line 7) and $\mathcal{P}$ is upwards closed, it suffices to show that p_V is below the segment $[p_{W_{i-1}}, p_{W_i}]$, since then the line through $p_{W_{i-1}}, p_{W_i}$ separates the interior of Δ from $\mathcal{F}$. Let s_l, s_u, s_m denote the slopes of the segments $[p_{W_{i-1}}, p_V]$,$[p_V, p_{W_i}]$ and $[p_{W_{i-1}}, p_{W_i}]$ respectively. Since the x-coordinates of $p_{W_{i-1}}, p_V, p_{W_i}$ satisfy $\dim(W_{i-1}) < \dim(V) < \dim(W_i)$, it suffices to show that either $s_l < s_m$ or that $s_u > s_m$. The expressions for the slopes are easily derived as follows:

$$\begin{aligned}
s_l &= \frac{\ln(\det(\mathcal{L} \cap V)) - \ln(\det(\mathcal{L} \cap W_{i-1}))}{\dim(V) - \dim(W_{i-1})} = \ln(\mathrm{nd}((\mathcal{L} \cap V)/W_{i-1}))\,,\\
s_u &= \frac{\ln(\det(\mathcal{L} \cap W_i)) - \ln(\det(\mathcal{L} \cap V))}{\dim(W_i) - \dim(V)} = \ln(\mathrm{nd}((\mathcal{L} \cap W_i)/V))\,,\\
s_m &= \frac{\ln(\det(\mathcal{L} \cap W_i)) - \ln(\det(\mathcal{L} \cap W_{i-1}))}{\dim(W_i) - \dim(W_{i-1})} = \ln(\mathrm{nd}((\mathcal{L} \cap W_i)/W_{i-1}))\,.
\end{aligned} \tag{20}$$

Let $\mathcal{L}_i := (\mathcal{L} \cap W_i)/W_{i-1}$. We now examine cases (a) and (b). For case (a), assume that $V = W_{i-1} + \mathrm{span}(M_{i,l})$ as in line 17, where $M_{i,l}$ is a primitive sublattice of $\mathcal{L}_i$. Then, since $\mathcal{L} \cap W/W_{i-1} = M_{i,l}$, by the check on line 16,

$$s_u = \ln(\mathrm{nd}(M_{i,l})) < \ln(2/3\mathrm{nd}(\mathcal{L}_i)) = -\ln(3/2) + s_m < s_m\,, \tag{21}$$

as needed. For case (b), assume that $V = \mathrm{span}(M_{i,h})^\perp \cap W_i$ as in line 24, where $M_{i,h}$ is a primitive sublattice of $\mathcal{L}_i^*$. By the check on line 23, we have that

$$\mathrm{nd}(M_{i,h}) \leq \frac{2}{3}\mathrm{nd}(\mathcal{L}_i^*) \Leftrightarrow \mathrm{nd}(M_{i,h}^*) \geq \frac{3}{2}\mathrm{nd}(\mathcal{L}_i) \Leftrightarrow \ln(\mathrm{nd}(M_{i,h}^*)) \geq \ln(3/2) + s_m$$

From here, it is not hard to verify that $(\mathcal{L} \cap W_i)/V = M^*_{i,h}$, and hence $s_u = \ln((\det \cap W_i)/V) \geq \ln(3/2) + s_m > s_m$, as needed.

To finish the proof of claim 2, we must show that $\mathrm{vol}_2(\Delta) \geq \ln(3/2)$. From here, a direct computation reveals (i.e. base $\times$ height /2) that

$$\begin{aligned} \mathrm{vol}_2(\Delta) &= (\dim(W_i) - \dim(W_{i-1}))(\dim(V) - \dim(W_{i-1}))(s_m - s_l)/2 \\ &= (\dim(W_i) - \dim(W_{i-1}))(\dim(W_i) - \dim(V))(s_u - s_m)/2\,. \end{aligned} \tag{22}$$

Since V is strictly in between W_{i-1} and W_i, we see that $\dim(W_i) - \dim(W_{i-1}) \geq 2$, and both $\dim(V) - \dim(W_{i-1}), \dim(W_i) - \dim(V) \geq 1$. From the above arguments, we also have that in case (a) $s_m - s_l \geq \ln(3/2)$ and (b) $s_u - s_m \geq \ln(3/2)$. Combining these estimates together, we conclude that $\mathrm{vol}_2(\Delta) \geq \ln(3/2)$, as needed. □

7 Bounding The Slicing Constant of Voronoi Cells

In this section, we show that Voronoi cells admit a rather good unconditional bound on the slicing constant, using the recent resolution of the ℓ_2 Kannan-Lovász conjecture [RSD17]. More precisely, we show that the outer volume ratio is any Voronoi cell is bounded by the ℓ_2 Kannan-Lovász constant.

Theorem 7.1. *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice and let $\mathcal{V} := \mathcal{V}(\mathcal{L})$. Then*

$$\mathrm{vr}^\circ(\mathcal{V}(\mathcal{L})) = O(C_{KL,2}(n))\,.$$

Proof. To begin we shall first recursively construct orthogonal subspaces $W_1, \ldots, W_k \subseteq \mathbb{R}^n$ for some $k \in [n]$, corresponding to "good" KL projections, which we will use to define the axes of the containing ellipsoid for $\mathcal{V}$.

Precisely, we start by defining W_1 as the lattice subspace of $\mathcal{L}^*$ maximizing

$$\max_{1 \leq d = \dim(W) \leq n} \frac{\det(\pi_W(\mathcal{L}))^{\frac{1}{d}}}{\mathrm{vol}_d(\mathbb{B}_2^d)^{\frac{1}{d}}}\,.$$

If $W_1 = \mathbb{R}^n$ stop, otherwise define W_2 recursively in the same way with respect $\mathcal{L}_2 = \mathcal{L} \cap W_1^\perp$ and the subspace $\mathrm{span}(\mathcal{L}_2) = W_1^\perp$, and so forth.

Let $d_1, \ldots, d_k$ denote the dimensions of $W_1, \ldots, W_k$, and let $\mathcal{L}_i = \mathcal{L} \cap \sum_{j=i}^k W_j$ for $i \in [n]$, noting that $\mathcal{L}_1 = \mathcal{L}$. By construction these subspaces are all orthogonal, $\sum_{i=1}^k W_i = \mathbb{R}^n$, and for $i \in [k]$, the subspace W_i is a maximizing KL-subspace for the sublattice $\mathcal{L}_i$ of $\mathcal{L}$. In particular, defining

$$r_i := \frac{\det(\pi_{W_i}(\mathcal{L}_i))^{1/d_i}}{\mathrm{vol}_{d_i}(\mathbb{B}_2^{d_i})^{\frac{1}{d_i}}}\,,$$

we have that $\mu(\mathcal{L}_i) \leq O(C_{KL,2}(n))r_i$ (the $O(1)$ comes from the slightly different normalization we use here to define KL subspaces).

Now examine the convex body

$$K = \{\mathbf{x} \in \mathbb{R}^n : \|\pi_{W_i}(\mathbf{x})\|_2 \leq r_i, i \in [k]\}\,.$$

We claim that

1. $\mathrm{vol}_n(K) = \det(\mathcal{L})$.
2. $\mathcal{V}(\mathcal{L}) \subseteq O(C_{KL,2}(n))K$.

We prove (1). Since the subspaces $W_1, \dots, W_k$ are orthogonal, we see that

$$\mathrm{vol}_n(K) = \prod_{i=1}^{k} \mathrm{vol}_{d_i}(r_i \mathbb{B}_2^{d_i}) = \prod_{i=1}^{k} \det(\pi_{W_i}(\mathcal{L}_i)) = \det(\mathcal{L}) \,,$$

where in the last equality we inductively use the identity

$$\det(\Lambda) = \det(\Lambda \cap W^{\perp}) \det(\pi_W(\Lambda))$$

for any lattice Λ and lattice subspace of W of Λ^*.

We prove (2). Take $x \in \mathcal{V}$. Then for $i \in [k]$, we have that

$$\|\pi_{W_i}(x)\| \le \|\pi_{\mathrm{span}(\mathcal{L}_i)}(x)\| \,.$$

Since $\mathcal{L}_i$ is a sublattice of $\mathcal{L}$, we know that $\pi_{\mathrm{span}(\mathcal{L}_i)}(\mathcal{V}) \subseteq \mathcal{V}(\mathcal{L}_i)$, and hence $\pi_{\mathrm{span}(\mathcal{L}_i)}(\mathbf{x}) \in \mathcal{V}(\mathcal{L}_i)$. In particular,

$$\|\pi_{\mathrm{span}(\mathcal{L}_i)}(x)\| \le \mu(\mathcal{L}_i) \le O(C_{KL,2}(n)) r_i \,,$$

as needed.

Note that if K were an ellipsoid then $O(C_{KL}(n))K$ would be a witness showing a $O(C_{KL,2}(n))$ bound on the outer volume ratio of $\mathcal{V}$. Fortunately, K is a direct product of ellipsoids, which themselves has constant volume ratio.

To see this, let $E = \{\mathbf{x} \in \mathbb{R}^n : \sum_{i=1}^{k} \frac{d_i}{n} \|\pi_{W_i}(x)\|^2 / r_i^2 \le 1\}$. Clearly, $K \subseteq E$ and E is an ellipsoid. Thus, to complete the theorem, it suffices to show that $\mathrm{vol}_n(E) \le c^n \mathrm{vol}_n(K)$ for an absolute constant $c > 1$.

A direct computation reveals that

$$\mathrm{vol}_n(E) = \prod_{i=1}^{k} r_i^{d_i} \left(\frac{n}{d_i}\right)^{d_i/2} \mathrm{vol}_n(\mathbb{B}_2^n) \,.$$

Recalling that

$$\mathrm{vol}_n(K) = \prod_{i=1}^{k} r_i^{d_i} \mathrm{vol}_{d_i}(\mathbb{B}_2^{d_i}) \,,$$

we have that

$$\begin{aligned} \frac{\mathrm{vol}_n(E)}{\mathrm{vol}_n(K)} &= \left(\prod_{i=1}^{k} \frac{1}{\mathrm{vol}_{d_i}(\mathbb{B}_2^{d_i}) d_i^{d_i/2}}\right) n^{n/2} \mathrm{vol}_n(\mathbb{B}_2^n) \\ &\le n^{n/2} \mathrm{vol}_n(\mathbb{B}_2^n) \le \sqrt{2\pi e}^n \,, \end{aligned}$$

as needed. □

By combining Lemma 3.37 with Theorem 7.1, we derive the following immediate corollary:

Corollary 7.2. *The Voronoi cell $\mathcal{V} \subseteq \mathbb{R}^n$ of an n-dimensional lattice satisfies $L_{\mathcal{V}} = O(C_{KL,2}(n)) = O(\log^{3/2}(n))$.*

References

[ADRS15] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time via discrete Gaussian sampling. In *STOC*, 2015. Available at `http://arxiv.org/abs/1412.7994`.

[ADS15] Divesh Aggarwal, Daniel Dadush, and Noah Stephens-Davidowitz. Solving the closest vector problem in 2^n time–the discrete Gaussian strikes again! In *IEEE 56th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 563–582. IEEE, 2015.

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

[AKS02] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *CCC*, pages 41–45, 2002.

[Bab86] László Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.

[Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.

[Bou91] J. Bourgain. On the distribution of polynomials on high-dimensional convex sets. In *Geometric aspects of functional analysis (1989–90)*, volume 1469 of *Lecture Notes in Math.*, pages 127–137. Springer, Berlin, 1991.

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

[Dad14] Daniel Dadush. A randomized sieving algorithm for approximate integer programming. *Algorithmica*, 70(2):208–244, 2014. Preliminary version in LATIN 2012.

[DK16] Daniel Dadush and Gábor Kun. Lattice sparsification and the approximate closest vector problem. *Theory of Computing*, 12:Paper No. 2, 34, 2016. Preliminary version in SODA 2012.

[DM13] Daniel Dadush and Daniele Micciancio. Algorithms for the densest sub-lattice problem. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1103–1122. SIAM, Philadelphia, PA, 2013.

[DPV11] Daniel Dadush, Chris Peikert, and Santosh Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science—FOCS 2011*, pages 580–589. IEEE Computer Soc., Los Alamitos, CA, 2011.

[DR16] Daniel Dadush and Oded Regev. Towards strong reverse minkowski-type inequalities for lattices. In *IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 2016.

[Gra84] Daniel R. Grayson. Reduction theory using semistability. *Comment. Math. Helv.*, 59(4):600–634, 1984.

[HK13] Robert Hildebrand and Matthias Kppe. A new lenstra-type algorithm for quasiconvex polynomial integer minimization with complexity $2^{O(n \log n)}$. *Discrete Optimization*, 10(1):69 – 84, 2013.

[HLR09] Ishay Haviv, Vadim Lyubashevsky, and Oded Regev. A note on the distribution of the distance from a lattice. *Discrete and Computational Geometry*, 41(1):162–176, 2009.

[HN75] G. Harder and M. S. Narasimhan. On the cohomology groups of moduli spaces of vector bundles on curves. *Math. Ann.*, 212:215–248, 1974/75.

[Hs88] J. Hå stad. Dual vectors and lower bounds for the nearest lattice point problem. *Combinatorica*, 8(1):75–81, 1988.

[Kan87] Ravi Kannan. Minkowski's convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, 1987.

[Khi48] A. Khinchine. A quantitative formulation of the approximation theory of Kronecker. *Izvestiya Akad. Nauk SSSR. Ser. Mat.*, 12:113–122, 1948.

[KL88] Ravi Kannan and László Lovász. Covering minima and lattice-point-free convex bodies. *Ann. of Math. (2)*, 128(3):577–602, 1988.

[Kla06] B. Klartag. On convex perturbations with a bounded isotropic constant. *Geom. Funct. Anal.*, 16(6):1274–1290, 2006.

[Len83] H. W. Lenstra, Jr. Integer programming with a fixed number of variables. *Math. Oper. Res.*, 8(4):538–548, 1983.

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

[LLS90] J. C. Lagarias, H. W. Lenstra, Jr., and C.-P. Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990.

[MV13] Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. *SIAM Journal on Computing*, 42(3):1364–1391, 2013.

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

[RSD17] Oded Regev and Noah Stephens-Davidowitz. A Reverse Minkowksi Theorem. In *STOC*, 2017.

[Sch87] C.-P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoret. Comput. Sci.*, 53(2-3):201–224, 1987.

[SFS09] Naftali Sommer, Meir Feder, and Ofir Shalvi. Finding the closest lattice point by iterative slicing. *SIAM J. Discrete Math.*, 23(2):715–731, 2009.

[Stu76] Ulrich Stuhler. Eine Bemerkung zur Reduktionstheorie quadratischer Formen. *Arch. Math. (Basel)*, 27(6):604–610, 1976.