On Approximating the Covering Radius and Finding Dense Lattice Subspaces[*]

Daniel Dadush
dadush@cwi.nl
Centrum Wiskunde & Informatica
Amsterdam, The Netherlands

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.

CCS CONCEPTS

• **Theory of computation → Approximation algorithms analysis**; **Complexity classes**; **Randomness, geometry and discrete structures**.

KEYWORDS

Lattice Problems, Kannan-Lovász Conjecture, Integer Programming

ACM Reference Format:
Daniel Dadush. 2019. On Approximating the Covering Radius and Finding Dense Lattice Subspaces. In *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC '19), June 23–26, 2019, Phoenix, AZ, USA.* ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/3313276.3316397

1 INTRODUCTION

A k-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$ is the integer span of k linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_k \in \mathbb{R}^n$. 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 $\mathbf{B} = (\mathbf{b}_1, \dots, \mathbf{b}_k)$ of $\mathcal{L}$, where $1/\det(\mathcal{L})$ corresponds to 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 non-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 as the normal vectors of lattice hyperplanes

[*] This is an extended abstract. All proofs are deferred to the full version of the paper.

STOC '19, June 23–26, 2019, Phoenix, AZ, USA

ACM ISBN 978-1-4503-6705-9/19/06...$15.00
https://doi.org/10.1145/3313276.3316397

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 [7]:

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

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 is highly non-trivial, and indeed the original proof of Khinchine [20] 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 [24], 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 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 [6, 7, 17, 19, 22], yielding $n^{O(n)}$ estimates on the IP search tree size, as well as the development fast basis reduction algorithms for SVP & CVP [18, 23, 28], for which [18] gave $n^{O(n)}$-time and 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 [4, 5], Voronoi cell computations [25, 30], and mostly recently discrete Gaussian sampling [1–3], where now both problems can be solved exactly in randomized $2^{n+o(n)}$-time and -space [1, 2]. 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 $c = 1$ [8].

The Kannan-Lovász Conjecture. A 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 [19] 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 $\mathbb{R}^n$, [19] 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 $\text{vol}_n(sK) \geq \det(\mathcal{L})$ (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 $C_{KL}(n)$, [19] 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 connection between the KL conjecture and IP algorithms was given by the author in [8], 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

[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 $\mathcal{L}$.

efficient KL projection algorithm, one would get $O(\log n)^n$-time algorithm. We note that the $O(n)^n$ algorithm from [8] is in fact derived from this result, using a $2^{O(n)}$-time general norm SVP solver to instantiate an oracle achieving the $C_{KL}(n) \le n$ bound. We note that the best k-dimensional projection above was in fact shown to be computable in $k^{O(kn)}$-time in [11], 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 $\mu(K, \mathcal{L}) \le 1/2$, 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 [9, 10]. 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), $\dim(W) = k \ge 1$, such that $\mathrm{vol}_k(\pi_W(K))/\det(\pi_W(\mathcal{L})) \le 2^k C_{KL}(n)^k$. 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

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

The uniform sparsity condition was shown in [8, 12] (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 factors in breakthrough work by Regev and Stephens-Davidowitz [26]. 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 $\mathcal{L}$,

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

where [26] 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" $\pi_W(\mathcal{L})$, $\dim(W) = k \ge 1$, 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.

From the complexity perspective, the above has strong complexity implications for the covering radius problem. Given as input a lattice $\mathcal{L}$ and a number $t > 0$, the α-approximate covering radius problem (α-GapCRP), is to decide whether $\mu(\mathcal{L}) \le t$ (YES instance) or $\mu(\mathcal{L}) > \alpha t$ (NO instance). Since the determinant of a lattice projection can be computed in polynomial time, one derives that α-GapCRP $\in$ coNP for $\alpha = O(C_{KL,2}(n)) = O(\log^{3/2}(n))$. Prior to resolution of the ℓ_2-KL conjecture, it was shown in [14] that 2-GapCRP $\in$ AM, $\sqrt{n/\log n}$-GapCRP $\in$ coAM and $\sqrt{n}$-GapCRP $\in$ NP $\cap$ coNP. Thus, the ℓ_2-KL bound gives an exponential improvement to the coNP classification for GapCRP. In terms of hardness, it was shown in [15] that the c_p-GapCRP is Π_2-hard for any large enough ℓ_p norms (here we measure the covering radius under the ℓ_p instead of ℓ_2 norm).

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

$$\begin{aligned} &\min_{\substack{\text{subspace W} \\ k:=\dim(W) \ge 1}} (\sqrt{k} \det(\pi_W(\mathcal{L}))^{1/k})^{-1} = \\ &\qquad \min_{\substack{W \text{ subspace of } \mathcal{L}^* \\ k:=\dim(W) \ge 1}} \det(\mathcal{L}^* \cap W)^{1/k}/\sqrt{k}. \end{aligned} \tag{4}$$

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 [11] 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 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 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 [26] 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 result on 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 [11] 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 on the reverse Minkowski theorem of [26] 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 [26] 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 yields an NP certificate for a lower bound 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 [13], and show how to "compile" any chain of subspaces into a solution of this program. Using the proof of the ℓ_2-KL conjecture of [26], 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 $\mathcal{L}$, we have that

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

where the maximum is taken over chains of lattice subspaces $W_0 \subset W_1 \subset \cdots \subset W_k$ as in Lemma 2.4. In particular, $O(L^\mu(n))$-GapCRP $\in$ coNP.

The maximum in the definition of $L^\mu(n)$ is taken over what are known as *stable* lattices of dimension n, which are lattices having determinant 1 and whose sublattices all have determinant at least 1. Using the tools developed to prove the reverse Minkowski theorem (described further below), [26] established the upper bound $L^\mu(n) = O(\log n)$. Furthermore, it was conjectured by Shapira and Weiss [29] that $L^\mu(n)$ is maximized at the integer lattice $\mathbb{Z}^n$, resulting in the bound $L^\mu(n) = 1/2$. Under this assumption, we get the existence of NP-certificates lower bounding the covering radius that are tight up to a constant factor, that is $O(1)$-GapCRP $\in$ coNP. We believe this conditional result to be rather surprising. To the author's knowledge, it was not even conjectured that GapCRP would be in coNP for an approximation factor not depending on the lattice dimension n. Unconditionally, the above theorem shows that $O(\log n)$-GapCRP $\in$ coNP, improving on the $O(\log^{3/2} n)$-approximation factor of [26].

As part of their proof, [26] also show that $L^\mu(n)$ is in fact constant factor equivalent to the worst-case slicing constant of any n-dimensional Voronoi cell of a stable lattice. This shows that the above conjecture reduces to a special case of the so-called slicing conjecture in convex geometry. 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 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 $L_{\mathbb{B}_2^n} = (1 + o(1))/\sqrt{2\pi e}$.

Bourgain's slicing conjecture 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 [21]. As [26] 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 $\mathrm{vol}_n(E)^{1/n} \leq O(C_{KL,2}(n)) \,\mathrm{vol}_n(\mathcal{V})^{1/n}$. 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

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

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 [27], 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 [8], 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 [16]). 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,

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

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 the main ideas behind our algorithm for finding dense lattice subspaces, deferring discussion of the remaining results to the full version 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 [1], 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 $s = \mathrm{nd}(M)/2$, 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.

Lemma 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 W.*

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 with constant probability, we reduce dimension by 1 and maintain the property that M contains a densest sublattice $M \cap W$ in the next iteration. 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 $\mathrm{nd}(sM^*) = s/\mathrm{nd}(M) = 1/2$. 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 implies 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 $\mathrm{nd}(\pi_W(M^*)) = \mathrm{nd}(M \cap W)^{-1} = 1/\tau(M)$. 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 choice of s, we have that $\tau(\pi_W(sM^*)) = \mathrm{nd}(M)/(2\tau(M)) = \Omega(\log n)$. That is, $\pi_W(sM^*)$ is "uniformly sparse" in that it 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 [26], first conjectured in [13], 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 be the smallest number such that for any lattice $\mathcal{L}$ of dimension at most n with $\tau(\mathcal{L}) \geq C_\eta(n)$ satisfies $\rho(\mathcal{L}) \leq 3/2$.

The reverse Minkowski theorem can now be stated as follows:

THEOREM 2.13. *[26]* $C_\eta(n) = O(\log n)$.

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

ACKNOWLEDGMENTS. 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. The research was supported by the NWO Veni grant 639.071.510.

REFERENCES

[1] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. 2015. Solving the Shortest Vector Problem in 2^n Time via Discrete Gaussian Sampling. In *STOC*. Available at http://arxiv.org/abs/1412.7994.

[2] Divesh Aggarwal, Daniel Dadush, and Noah Stephens-Davidowitz. 2015. 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)*. IEEE, 563–582.

[3] Divesh Aggarwal and Noah Stephens-Davidowitz. 2018. Just Take the Average! An Embarrassingly Simple 2^n-Time Algorithm for SVP (and CVP). In *1st Symposium on Simplicity in Algorithms (SOSA 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[4] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. 2001. A Sieve Algorithm for the Shortest Lattice Vector Problem. In *STOC*. 601–610. http://doi.acm.org/10.1145/380752.380857

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

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

[7] W. Banaszczyk. 1993. New bounds in some transference theorems in the geometry of numbers. *Math. Ann.* 296, 4 (1993), 625–635. https://doi.org/10.1007/BF01445125

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

[9] Daniel Dadush. 2014. A randomized sieving algorithm for approximate integer programming. *Algorithmica* 70, 2 (2014), 208–244. https://doi.org/10.1007/s00453-013-9834-8 Preliminary version in LATIN 2012.

[10] Daniel Dadush and Gábor Kun. 2016. Lattice sparsification and the approximate closest vector problem. *Theory of Computing* 12 (2016), Paper No. 2, 34. https://doi.org/10.4086/toc.2016.v012a880 Preliminary version in SODA 2012.

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

[12] Daniel Dadush, Chris Peikert, and Santosh Vempala. 2011. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *2011 IEEE 52nd Annual Symposium on Foundations of Computer Science—FOCS 2011*. IEEE Computer Soc., Los Alamitos, CA, 580–589. https://doi.org/10.1109/FOCS.2011.31

[13] Daniel Dadush and Oded Regev. 2016. Towards Strong Reverse Minkowski-Type Inequalities for Lattices. In *IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*.

[14] Venkatesan Guruswami, Daniele Micciancio, and Oded Regev. 2005. The complexity of the covering radius problem. *computational complexity* 14, 2 (Jun 2005), 90–121.

[15] Ishay Haviv and Oded Regev. 2006. Hardness of the covering radius problem on lattices. In *21st Annual IEEE Conference on Computational Complexity (CCC'06)*. IEEE, 14–pp.

[16] Robert Hildebrand and Matthias Köppe. 2013. A new Lenstra-type algorithm for quasiconvex polynomial integer minimization with complexity $2^{O(n \log n)}$. *Discrete Optimization* 10, 1 (2013), 69 – 84. https://doi.org/10.1016/j.disopt.2012.11.003

[17] J. Håstad. 1988. Dual vectors and lower bounds for the nearest lattice point problem. *Combinatorica* 8, 1 (1988), 75–81. https://doi.org/10.1007/BF02122554

[18] Ravi Kannan. 1987. Minkowski's convex body theorem and integer programming. *Math. Oper. Res.* 12, 3 (1987), 415–440. https://doi.org/10.1287/moor.12.3.415

[19] Ravi Kannan and László Lovász. 1988. Covering minima and lattice-point-free convex bodies. *Ann. of Math. (2)* 128, 3 (1988), 577–602. https://doi.org/10.2307/1971436

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

[21] B. Klartag. 2006. On convex perturbations with a bounded isotropic constant. *Geom. Funct. Anal.* 16, 6 (2006), 1274–1290. https://doi.org/10.1007/s00039-006-0588-1

[22] J. C. Lagarias, H. W. Lenstra, Jr., and C.-P. Schnorr. 1990. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica* 10, 4 (1990), 333–348. http://dx.doi.org/10.1007/BF02128669

[23] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. 1982. Factoring polynomials with rational coefficients. *Math. Ann.* 261, 4 (1982), 515–534. https://doi.org/10.1007/BF01457454

[24] H. W. Lenstra, Jr. 1983. Integer programming with a fixed number of variables. *Math. Oper. Res.* 8, 4 (1983), 538–548. https://doi.org/10.1287/moor.8.4.538

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

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

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

[28] C.-P. Schnorr. 1987. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoret. Comput. Sci.* 53, 2-3 (1987), 201–224. https://doi.org/10.1016/0304-3975(87)90064-8

[29] Uri Shapira and Barak Weiss. 2016. Stable lattices and the diagonal group. *Journal of the European Mathematical Society* 18, 8 (2016), 1753–1767.

[30] Naftali Sommer, Meir Feder, and Ofir Shalvi. 2009. Finding the closest lattice point by iterative slicing. *SIAM J. Discrete Math.* 23, 2 (2009), 715–731. https://doi.org/10.1137/060676362