Faster Deterministic Volume Estimation in the Oracle Model via Thin Lattice Coverings

Daniel Dadush

Centrum Wiskunde & Informatica, The Netherlands
`dadush@cwi.nl`

--- **Abstract** ---

We give a $2^{O(n)}(1+1/\varepsilon)^n$ time and poly(n)-space deterministic algorithm for computing a $(1+\varepsilon)^n$ approximation to the volume of a general convex body K, which comes close to matching the $(1+c/\varepsilon)^{n/2}$ lower bound for volume estimation in the oracle model by Bárány and Füredi (STOC 1986, Proc. Amer. Math. Soc. 1988). This improves on the previous results of Dadush and Vempala (Proc. Nat'l Acad. Sci. 2013), which gave the above result only for *symmetric bodies* and achieved a dependence of $2^{O(n)}(1 + \log^{5/2}(1/\varepsilon)/\varepsilon^3)^n$.

For our methods, we reduce the problem of volume estimation in K to counting lattice points in $K \subseteq \mathbb{R}^n$ (via enumeration) for a specially constructed lattice $\mathcal{L}$: a so-called *thin covering of space* with respect to K (more precisely, for which $\mathcal{L} + K = \mathbb{R}^n$ and $\mathrm{vol}_n(K)/\det(\mathcal{L}) = 2^{O(n)}$). The trade off between time and approximation ratio is achieved by scaling down the lattice.

As our main technical contribution, we give the first deterministic $2^{O(n)}$-time and poly(n)-space construction of thin covering lattices for general convex bodies. This improves on a recent construction of Alon et al. (STOC 2013) which requires exponential space and only works for symmetric bodies. For our construction, we combine the use of the M-ellipsoid from convex geometry (Milman, C. R. Math. Acad. Sci. Paris 1986) together with lattice sparsification and densification techniques (Dadush and Kun, SODA 2013; Rogers, J. London Math. Soc. 1950).

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases Deterministic Volume Estimation, Convex Geometry, Lattice Coverings of Space, Lattice Point Enumeration

Digital Object Identifier 10.4230/LIPIcs.SOCG.2015.704

1 Introduction

The problem of estimating the volume of a convex body is one of the most fundamental and well studied problems in high dimensional geometry. It is also one of the most striking examples of the *power of randomization*. In [11, 12], Bárány and Füredi showed that any deterministic volume algorithm for n-dimensional convex bodies having access only to a membership oracle (which returns whether a point is in the convex body or not), requires at least $(1+c/\varepsilon)^{n/2}$ membership queries to estimate volume to within a $(1+\varepsilon)^n$ factor, for $c > 0$ an absolute constant any ε small enough. In particular, an $O(1)$-approximation requires $n^{\Omega(n)}$ queries. In a breakthrough result however, Dyer, Frieze and Kannan [9] showed that if the algorithm is allowed to err with small probability, then even a $(1 + \varepsilon)$ approximation can be obtained in poly$(n, 1/\varepsilon)$-time. Their algorithm relied on novel Monte Carlo Markov Chain techniques that spurred much further research. These works left a major open question: can the volume algorithm be made deterministic when the description of the convex body is given explicitly (e.g. a polytope given by its inequalities)?

A related (and more modest) question, which has only recently received attention, is whether one can come close to matching the lower bounds of Bárány and Füredi for

deterministic volume computation in the oracle model. We note it was open to achieve such bounds deterministically even for explicitly presented polytopes. This was recently answered in the affirmative by Vempala and the author in [8], which gave a deterministic $2^{O(n)}(1 + \log^{5/2}(1/\varepsilon)/\varepsilon^3)^n$-time and polynomial space algorithm for estimating the volume of a *symmetric* convex body K (K is symmetric if $K = -K$) to within $(1 + \varepsilon)^n$. The main tool developed there was an algorithmic version of (variants of) Milman's construction for the *M-ellipsoid* in convex geometry [18]. An M-ellipsoid of an n-dimensional convex body K is an ellipsoid E (an ellipsoid is a linear transformation of the Euclidean ball) satisfying that $2^{O(n)}$ translates of E suffice to cover K and vice versa. Note that the volume of an M-ellipsoid of K immediately provides a $2^{O(n)}$ factor approximation to the volume of K.

From the above, two natural avenues of improvement were to reduce the dependence on ε and to generalize the result to asymmetric convex bodies.

2 Main Contribution

We make improvements on both of the last two fronts. Our main result is stated below.

▶ **Theorem 1** (Volume Estimation). *For a convex body $K \subseteq \mathbb{R}^n$ given by a membership oracle, and any $\varepsilon > 0$, one can compute $V \geq 0$ satisfying $\mathrm{vol}_n(K) \leq V \leq (1 + \varepsilon)^n \mathrm{vol}_n(K)$ in deterministic $2^{O(n)}(1 + 1/\varepsilon)^n$-time and $\mathrm{poly}(n)$-space.*

Both the algorithm and that of [8] share the same high level approach, namely, reducing volume estimation to counting lattice points within a carefully chosen convex body and lattice.

We note that if we are satisfied with a c^n approximation of volume for some large enough $c > 0$, then the volume of an M-ellipsoid is already a good enough volume approximation for K and hence lattice point counting is not needed. This extends to asymmetric convex bodies as well, by replacing K with the symmetric body $K - K$ (an oracle for which can be efficiently computed, see [13]) and using the standard inequalities

$$2^n \mathrm{vol}_n(K) \leq \mathrm{vol}_n(K - K) \leq \binom{2n}{n} \mathrm{vol}_n(K) \qquad \text{(see [23])}.$$

Hence the above result is truly interesting for the case of small constant ε.

Our runtime improvement over the algorithm of [8] comes from a much more efficient reduction from volume estimation to lattice point counting. In particular, the crucial ingredient in our improved reduction is the use of so-called *thin lattice coverings of space* with respect to K (and related convex bodies). The heart of our volume algorithm, and our main technical contribution, is a deterministic construction of thin-covering lattices for general convex bodies with *good enumeration properties*, that is, where lattice point enumeration can be performed efficiently using only polynomial space. This improves on a recent thin-lattice construction of [1] which requires exponential space and only works for symmetric bodies.

Organization. The remainder of this paper is organized as follows. First, we shall explain the reduction between volume estimation and lattice point counting, which will motivate the need for thin covering lattices and other related concepts. Second, we will present the polynomial space lattice point enumeration technique we use – Schorr-Euchner enumeration – and briefly discuss its implementation and associated challenges. Third, we give the formal statements of our main thin lattice construction and related algorithms, and their relations to prior work. Finally, in the remainder, we shall detail the main ideas behind the thin covering lattice construction.

3   Preliminaries

Basic concepts.   For two sets $A, B \subseteq \mathbb{R}^n$, we define their Minkowski sum $A + B = \{\mathbf{a} + \mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B\}$. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we write $\langle \mathbf{x}, \mathbf{y} \rangle = \sum_{i=1}^{n} x_i y_i$ to denote the standard inner product and $\|\mathbf{x}\|_2 = \sqrt{\langle \mathbf{x}, \mathbf{x} \rangle}$ for the Euclidean norm. We let $B_2^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq 1\}$ denote the unit Euclidean ball in $\mathbb{R}^n$. For a set $A \subseteq \mathbb{R}^n$, we denote its interior by A°. A convex body $K \subseteq \mathbb{R}^n$ is a compact convex set with non-empty interior. A function $f : \mathbb{R}^n \to \mathbb{R}$ is L-Lipshitz if $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, $|f(\mathbf{x}) - f(\mathbf{y})| \leq L \|\mathbf{x} - \mathbf{y}\|_2$.

Lattices.   We give some basic definitions of lattice concepts.

▶ **Definition 2** (Lattices and Bases). A full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is defined as all integer combinations of some basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n) \in \mathbb{R}^{n \times n}$. In particular, $\mathcal{L} = B\mathbb{Z}^n$. The determinant of $\mathcal{L}$ is defined as $\det(\mathcal{L}) = |\det(B)|$, which is invariant to the choice of lattice basis. We define $\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*$, the associated dual basis, to be the unique vectors satisfying $\langle \mathbf{b}_i, \mathbf{b}_j^* \rangle = 1$ if $i = j$ and 0 otherwise (corresponding to the columns of B^{-T}).

▶ **Definition 3** (Gram Schmid Projections). For a basis $\mathbf{b}_1, \ldots, \mathbf{b}_n \in \mathbb{R}^n$, we define the i^{th} Gram-Schmidt projection π_i, $i \in [n+1]$, to be the orthogonal projection onto the orthogonal complement of the linear span of $\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}$. Note that π_1 is the identity on $\mathbb{R}^n$ and π_{n+1} is the identically 0 map.

▶ **Definition 4** (Basis Parallelipiped). For a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ with basis B, we define $\mathcal{P}(B) = B[-1/2, 1/2)^n$ to be the half-open symmetric parallelepiped. Note that $\text{vol}_n(\mathcal{P}(B)) = \det(\mathcal{L})$.

▶ **Definition 5** (Sublattice Index). For a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and full rank sublattice $\mathcal{L}' \subseteq \mathcal{L}$, we define the index of $\mathcal{L}'$ in $\mathcal{L}$, denoted $[\mathcal{L} : \mathcal{L}']$, as $|\{\mathbf{y} + \mathcal{L}' : \mathbf{y} \in \mathcal{L}\}| < \infty$ (i.e. number of shifts of $\mathcal{L}'$ in $\mathcal{L}$). Here, we have the fundamental identity $[\mathcal{L} : \mathcal{L}'] = \det(\mathcal{L}')/\det(\mathcal{L})$.

▶ **Definition 6** (Lattice Tiling). A measurable set $A \subseteq \mathbb{R}^n$ tiles with respect to a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ (and vice versa) if for every $\mathbf{x} \in \mathbb{R}^n$ there is a *unique* $\mathbf{y} \in \mathcal{L}$ such that $\mathbf{x} \in \mathbf{y} + A$. Here, A is said to be a *fundamental domain* of $\mathcal{L}$.

   A basic fact is that every fundamental domain of $\mathcal{L}$ has the same volume. In particular, since $\mathcal{P}(B)$ is a fundamental domain, every fundamental domain of $\mathcal{L}$ has volume $\det(\mathcal{L})$.

Computational model.   For a convex body $K \subseteq \mathbb{R}^n$, a membership oracle O_K for K takes as input $\mathbf{x} \in \mathbb{R}^n$ and returns 1 if $\mathbf{x} \in K$ and 0 otherwise. K is $(\mathbf{a}_0, r, R)$-centered, for $r, R > 0$ and $\mathbf{a}_0 \in \mathbb{R}^n$, if $r B_2^n \subseteq K - \mathbf{a}_0 \subseteq R B_2^n$. When we refer to K being centered, we shall mean that the *centering guarantees* $(\mathbf{a}_0, r, R)$ exist and are implicitly passed to any algorithm operating on K and that the complexity of this algorithm may depend on these guarantees. For $\varepsilon > 0$, we define $K^\varepsilon = K + \varepsilon B_2^n$ and $K^{-\varepsilon} = \{\mathbf{x} \in K : \mathbf{x} + \varepsilon B_2^n \subseteq K\}$. A weak membership oracle O_K for K, takes an additional parameter $\varepsilon > 0$, and only guarantees that $O_K(\mathbf{x}, \varepsilon) = 1$ if $\mathbf{x} \in K^{-\varepsilon}$ and 0 if $\mathbf{x} \notin K^\varepsilon$. All our algorithms will operate on centered convex bodies equipped with (weak) membership oracles, and the complexity of our algorithms will be measured by the number of arithmetic operations and oracle calls they perform.

   One of the main algorithmic tools we will use is the following classical result in convex optimization:

▶ **Theorem 7** (Convex Optimization [25, 13]). *Let $K \subseteq \mathbb{R}^n$ be a centered convex body given by a weak membership oracle O_K. Let $f : \mathbb{R}^n \to \mathbb{R}$ denote an L-Lipshitz convex function.*

Then, for $\varepsilon > 0$, a vector $\mathbf{y} \in K$ satisfying

$$f(\mathbf{y}) - \varepsilon \leq \min_{\mathbf{x} \in K} f(\mathbf{x}) \leq f(\mathbf{y})$$

can be computed using a polynomial number of arithmetic operations, oracle calls and evaluations of f.

4   From Volume Estimation to Counting Lattice Points

In this section, we will show how to reduce the volume estimation problem to counting lattice points inside a well-chosen convex body. We will primarily concern ourselves with the task of minimizing the number of lattice points we need to enumerate to achieve a desired approximation factor. The important details regarding how to efficiently enumerate these lattice points is left to later sections.

To build intuition, we shall first try to estimate the volume of a convex body $K \subseteq \mathbb{R}^n$ by counting the number of points it contains in the standard integer lattice $\mathbb{Z}^n$. Through this attempt, we will expose some of the main ingredients needed to make volume approximation efficient.

For the integer lattice, the canonical relation between lattice point counting and volume is simply derived by associating every point in $\mathbf{y} \in \mathbb{Z}^n$ with the half open cube around it, i.e. $C = [-1/2, 1/2)^n + \mathbf{y}$. Since these shifted cubes have volume 1 and are all disjoint, the count $|K \cap \mathbb{Z}^n|$ is the same as the volume of the set $S = (K \cap \mathbb{Z}^n) + C$. Now as is, the set S may both miss parts of and "stick out" of K, so it is difficult to deduce any relationship between their volumes. To fix one of these problems, note that the cubes around each integer point form a *tiling of space*, that is every point in $\mathbb{R}^n$ is in exactly one such cube. Hence if we enlarge S to contain all the cubes centered around $\mathbb{Z}^n$ that *touch* K – formally, we redefine $S = ((K - C) \cap \mathbb{Z}^n) + C$ – then we are guaranteed that S covers K. In particular,

$$|(K - C) \cap \mathbb{Z}^n| = \mathrm{vol}_n(S) \geq \mathrm{vol}_n(K).$$

Note then that the volume of S can be computed if we can enumerate the integer points in $K - C$ (we defer for now the discussion of how to do this efficiently). So now, from the perspective of approximation, we are left with the problem that S may stick out very far from K, and hence may have very large volume compared to K. Indeed, this may easily happen (say if K is a ball of tiny radius), since we have made no assumptions on K.

Regardless, if we scale down $\mathbb{Z}^n$ and $C = [-1/2, 1/2)^n$ by ε, then as we let $\varepsilon \to 0$, the volume of S (defined on the scaled down lattice and cube) will clearly converge to the volume of K since S will converge to K. Given this, we are lead to two basic questions. Firstly, how small do we need to make ε to get a $(1 + \varepsilon)^n$ approximation of volume? Secondly, how many lattice points do we need to enumerate to compute this approximation? Crucially, the answer to this last question will essentially determine the complexity of the algorithm.

To get a quantitative estimate, let us normalize the geometry by assuming that $\pm C \subseteq K/2$. (while requiring the condition for both C and $-C$ is essentially redundant here, it will be very important when we generalize the forthcoming analysis.) Note that this can always be achieved by an appropriate shift and scaling of K. Letting $S_\varepsilon = ((K - \varepsilon C) \cap \varepsilon \mathbb{Z}^n) + \varepsilon C$, for $\varepsilon > 0$, by the same reasoning as before we have that

$$\mathrm{vol}_n(K) \leq \mathrm{vol}_n(S_\varepsilon) = \mathrm{vol}_n(\varepsilon C)|(K - \varepsilon C) \cap \mathbb{Z}^n| = \varepsilon^n |(K - \varepsilon C) \cap \mathbb{Z}^n|. \tag{1}$$

Furthermore, since $\pm C \subseteq K/2$, we have that

$$\text{vol}_n(S_\varepsilon) = \text{vol}_n(((K - \varepsilon C) \cap \varepsilon \mathbb{Z}^n) + \varepsilon C) \leq \text{vol}_n(K + \varepsilon(C - C))$$
$$\leq \text{vol}_n(K + \varepsilon(K/2 + K/2)) = \text{vol}_n((1 + \varepsilon)K) = (1 + \varepsilon)^n \text{vol}_n(K), \tag{2}$$

where the last two equalities hold by convexity of K and the homogeneity of volume respectively. Hence, from the above computing a $(1 + \varepsilon)^n$ approximation to $\text{vol}_n(K)$ reduces to enumerating the points in $(K - \varepsilon C) \cap \varepsilon \mathbb{Z}^n$. Combining (1),(2) and rearranging, we see that the number of points we must enumerate is bounded by

$$|(K - \varepsilon C) \cap \mathbb{Z}^n| \leq (1 + 1/\varepsilon)^n \text{vol}_n(K) = 2^n (1 + 1/\varepsilon)^n \left(\text{vol}_n(K/2)/\text{vol}_n(C)\right).$$

Now, if we believe that the correct measure of complexity is simply the number of lattice points we must enumerate (ignoring the actual complexity of enumeration for now), then we would achieve the complexity estimate in Theorem 1 if $\text{vol}_n(K/2)/\text{vol}_n(C) = 2^{O(n)}$. However, it is clear that not every convex body K can be scaled and shifted such that $\pm C \subseteq K/2$ and $\text{vol}_n(K/2)/\text{vol}_n(C) = 2^{O(n)}$.

On the other hand, it is easy to see that the above analysis can be substantially generalized. More precisely, instead of relying on the integer lattice, we may use an arbitrary lattice $\mathcal{L} = B\mathbb{Z}^n$, for some basis B. Instead of cubes (or parallelepipeds), we may use any measurable set $F \subseteq \mathbb{R}^n$ which tiles with respect to $\mathcal{L}$. From here, if there exists $\mathbf{c} \in K$, such that $\pm F \subseteq (K - \mathbf{c})/2$ (note that F need no longer be symmetric), then by the same analysis as above we have that

$$\text{vol}_n(K) \leq \varepsilon^n \cdot \text{vol}_n(F) \cdot |(K - \varepsilon F) \cap \varepsilon \mathcal{L}| \leq (1 + \varepsilon)^n \text{vol}_n(K). \tag{3}$$

When trying to use the above formula to approximate volume, one may rightly worry that the set F above maybe quite complicated, and hence of limited algorithmic use. Fortunately, it turns out that we won't actually need to know F at all – we will only need to rely on its *existence* – and, in fact, only knowledge of the point $\mathbf{c}$ will be required. To justify this, we first remark that F is a fundamental domain, and hence $\text{vol}_n(F) = \det(\mathcal{L})$, which is easily computable given B.

Let $K[\mathbf{c}] = (K - \mathbf{c}) \cap (\mathbf{c} - K)$ denote the symmetrization of K about $\mathbf{c}$ (note that $K[\mathbf{c}]$ is indeed symmetric). By construction, we see that

$$\pm F \subseteq \pm K[\mathbf{c}]/2 = K[\mathbf{c}]/2 \subseteq (K - \mathbf{c})/2.$$

From here, it is not hard to check that replacing $K - \varepsilon F$ by $K + \varepsilon K[\mathbf{c}]/2$ in (3) yields

$$\text{vol}_n(K) \leq \varepsilon^n \cdot \det(\mathcal{L}) \cdot |(K + \varepsilon K[\mathbf{c}]/2) \cap \varepsilon \mathcal{L}| \leq (1 + \varepsilon)^n \text{vol}_n(K). \tag{4}$$

The above formula will indeed form the basis of our algorithmic approach, where we note that a membership oracle for $K + \varepsilon K[\mathbf{c}]/2$ (under mild assumptions on $\mathbf{c}$) can be efficiently constructed from a membership oracle for K (see [13]). Rearranging as before, we get that the number of lattice points we need to enumerate to compute the desired approximation is bounded by

$$|(K + \varepsilon K[\mathbf{c}]/2) \cap \varepsilon \mathcal{L}| \leq 2^n (1 + 1/\varepsilon)^n \underbrace{\frac{\text{vol}_n(K)}{\text{vol}_n(K[\mathbf{c}])}}_{(a)} \underbrace{\frac{\text{vol}_n(K[\mathbf{c}]/2)}{\det(\mathcal{L})}}_{(b)} \tag{5}$$

Hence, to achieve the desired complexity bound, we will need both the expressions (a) and (b) to be bounded by $2^{O(n)}$. More precisely, we will need to compute a point $\mathbf{c} \in K$ and a lattice $\mathcal{L} \subseteq \mathbb{R}^n$ such that

1. $\mathrm{vol}_n(K) \leq 2^{O(n)} \mathrm{vol}_n(K[\mathbf{c}])$.
2. $\exists\ F \subseteq K[\mathbf{c}]$ a fundamental domain for $\mathcal{L}$, and $\mathrm{vol}_n(K[\mathbf{c}]) \leq 2^{O(n)} \det(\mathcal{L})$.

We note that condition (1) becomes trivial if K is already symmetric, since we can simply choose $\mathbf{c} = \mathbf{0}$. In condition (2), note that we have, for convenience of notation, multiplied the required conditions for $\mathcal{L}$ in (5) by 2.

We now relate some initial details of how to find $\mathbf{c}$ and $\mathcal{L}$ satisfying these conditions, deferring the full discussion of our methods to later sections. The plan here is to treat each condition separately. In particular, we will first choose $\mathbf{c}$ to satisfy (1) and then pick $\mathcal{L}$ satisfying (2).

Choosing the lattice $\mathcal{L}$. Once we have chosen $\mathbf{c}$, we wish to choose a lattice satisfying condition (2). For this purpose, we will only use the fact that $K[\mathbf{c}]$ is a symmetric convex body (which is why we can treat both conditions separately). As a first remark, we note that the existence of fundamental domain $F \subseteq K[\mathbf{c}]$ is equivalent to asking that $K[\mathbf{c}]$ *cover space* with respect to $\mathcal{L}$.

▶ **Definition 8** (Lattice Covering). A measurable $A \subseteq \mathbb{R}^n$ is covering with respect to a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ (and vice versa) if $\mathcal{L} + A = \mathbb{R}^n$. The covering induced by A and $\mathcal{L}$ is said to be α-thin, $\alpha \geq 1$, if $\mathrm{vol}_n(A)/\det(\mathcal{L}) \leq \alpha$.

Indeed, assuming that $\mathcal{L} + K[\mathbf{c}] = \mathbb{R}^n$, we can recover a suitable fundamental domain F by picking (in a measurable way) a unique representative of in $(\mathcal{L} + \mathbf{x}) \cap K[\mathbf{c}]$, for each distinct coset $\mathcal{L} + \mathbf{x}$, $\mathbf{x} \in \mathbb{R}^n$. We note that this simply corresponds to throwing away the "overrepresented" parts of $K[\mathbf{c}]$. From this discussion, we see that every covering of space must have thinness at least 1. Note that at a high level, the covering induced by $K[\mathbf{c}]$ and $\mathcal{L}$ being α-thin means that on average points in $\mathbb{R}^n$ are covered by at most α lattice shifts of $K[\mathbf{c}]$ (and clearly at least 1).

We can now restate our goal as that of constructing a lattice $\mathcal{L}$ forming a $2^{O(n)}$-thin covering with respect to $K[\mathbf{c}]$. We give a detailed accounting of how to build such lattices in section 8.1.

Choosing the center $\mathbf{c}$. To compute $\mathbf{c}$, we will require the following measure of symmetry:

▶ **Definition 9** (Kovner-Besicovitch Symmetry Measure). For a convex body $K \subseteq \mathbb{R}^n$, we define its *Kovner-Besicovitch* measure of symmetry (see [15]) as

$$\mathrm{Sym}_{kb}(K) = \max_{\mathbf{c} \in K} \mathrm{vol}_n(K[\mathbf{c}])/\mathrm{vol}_n(K), \quad \text{where } K[\mathbf{c}] = (K - \mathbf{c}) \cap (\mathbf{c} - K). \tag{6}$$

Note that K is symmetric (about some center) iff $\mathrm{Sym}_{kb}(K) = 1$. For $\mathbf{c} \in K$, we define its *KB value* to be $\mathrm{vol}_n(K[\mathbf{c}])/\mathrm{vol}_n(K)$. Clearly, to satisfy condition (1), the best center we can choose is simply that of maximum KB value. For such a maximizer to be useful, we must at least convince ourselves that best center has KB value at least $2^{-O(n)}$. For this purpose, let X denote a uniform random variable over K. By a classical computation, we have that

$$\mathbb{E}_X \left[\frac{\mathrm{vol}_n(K[X])}{\mathrm{vol}_n(K)} \right] = \int_K \frac{\mathrm{vol}_n(K[\mathbf{x}])}{\mathrm{vol}_n(K)^2} \mathrm{d}\mathbf{x} = \int_K \int_K \frac{\mathbf{1}[2\mathbf{x} - \mathbf{y} \in K]}{\mathrm{vol}_n(K)^2} \mathrm{d}\mathbf{y}\mathrm{d}\mathbf{x}$$
$$= \int_K \frac{\mathrm{vol}_n((K + \mathbf{y})/2)}{\mathrm{vol}_n(K)^2} \mathrm{d}\mathbf{y} = 2^{-n}.$$

By the probabilistic method, we therefore have that $\mathrm{Sym}_{kb}(K) \geq 2^{-n}$, which is more than good enough for us. Furthermore, it was actually shown in [19] that the centroid $\mu = \mathbb{E}[X]$ of K has KB value at least 2^{-n}. Hence, with the aid of random sampling techniques over convex bodies [9], computing a point with good KB value is rather straightforward.

Since our goal is to get a deterministic algorithm however, we cannot rely on random sampling methods. Perhaps surprisingly, our approach for computing a high KB value point will be to approximately solve the optimization problem in (6). Indeed, by the Brunn-Minkowski inequality (which states that $\mathrm{vol}(A)^{1/n} + \mathrm{vol}(B)^{1/n} \leq \mathrm{vol}(A + B)^{1/n}$ for $A, B, A + B$ measurable), the function $f(\mathbf{c}) = \mathrm{vol}_n(K[\mathbf{c}])^{1/n}$ is in fact *concave* over K. Hence, maximizing f is a concave optimization problem.

We define a point $\mathbf{c} \in K$ to be an α-approximate KB point for K, $0 < \alpha \leq 1$, if its KB value $\mathrm{vol}_n(K[\mathbf{c}])/\mathrm{vol}_n(K)$ is at least an α-factor of $\mathrm{Sym}_{kb}(K)$. For our purposes, it will suffice to be able to compute a $2^{-O(n)}$ approximate KB point, which we note corresponds to computing a constant factor approximation to $\max_{\mathbf{c} \in K} f(\mathbf{c})$. We will actually be able to compute $(1 + \varepsilon)^{-n}$-approximation KB points for any desired $\varepsilon > 0$ (see Theorem 22). Our approximation algorithm will be somewhat non-trivial, requiring many calls to our volume algorithm over symmetric bodies (noting that each $K[\mathbf{c}]$ is symmetric). We defer the full discussion to section 8.2.

5 Schnorr-Euchner Enumeration

The currently most powerful polynomial space lattice point enumeration strategy is *Schnorr-Euchner* enumeration. It is the primary enumeration method for all polynomial space solvers for the Closest Vector Problem (CVP) under the Euclidean norm (given a target $\mathbf{t}$ and lattice $\mathcal{L}$, find the closest vector in $\mathcal{L}$ to $\mathbf{t}$), and will form the core of our enumeration algorithm. We now explain how to adapt it to enumerate lattice points in general convex bodies (it was originally specified only for Euclidean balls, see for example [16]), and present some of its important properties.

High level algorithm. Given a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $\mathcal{L}$ and a convex body K, Schnorr-Euchner builds all feasible solutions to $\{\mathbf{z} \in \mathbb{Z}^n : \sum_{i=1}^n z_i \mathbf{b}_i \in K\}$, corresponding to $\mathcal{L} \cap K$, using a search tree over the coefficients. The nodes at level i of the tree, $i \in \{0, \ldots, n\}$, correspond to integral assignments of the last i coefficients that are "feasible" for K. Precisely, a partial assignment $z_{n-i+1}, \ldots, z_n \in \mathbb{Z}$ is feasible for K if $\exists\, r_1, \ldots, r_{n-i} \in \mathbb{R}$ such that

$$\sum_{j=1}^{n-i} r_j \mathbf{b}_j + \sum_{j=n-i+1}^{n} z_j \mathbf{b}_j \in K. \tag{7}$$

By convention, we consider the root (level 0) to have an empty assignment, which is feasible iff $K \neq \emptyset$. From a level i node, with partial assignment $z_{n-i+1}, \ldots, z_n \in \mathbb{Z}$, we recurse on all feasible extensions $z_{n-i}, \ldots, z_n$ with $z_{n-i} \in \mathbb{Z}$. By convexity of K, the set of integer assignments for z_{n-i} inducing a feasible extension form a consecutive interval, which will allow us to enumerate them efficiently.

Implementation. Since the nature of computations in the oracle model are always approximate, we will have to relax the notion of feasible partial assignment when implementing the above algorithm. In particular, we will only be able to determine where a partial assignment is either not feasible for K or feasible for K^ε, for any desired error tolerance $\varepsilon > 0$. The

exact guarantees for our enumeration algorithm, which will be sufficient for all intended applications, are stated below.

▶ **Lemma 10** (Enumeration Complexity). *Let $K \subseteq \mathbb{R}^n$ be a $(\mathbf{a}_0, r, R)$-centered convex body given a weak membership oracle, and let $\mathcal{L} \subseteq \mathbb{R}^n$ a full rank lattice with basis $B = (\mathbf{b}_1, \dots, \mathbf{b}_n)$. Then for $0 < \varepsilon < 1$, a set S satisfying $K \cap \mathcal{L} \subseteq S \subseteq K^\varepsilon \cap \mathcal{L}$ can be enumerated, where every point is outputted exactly once, using polynomial space and time polynomial times*

$$\sum_{i=0}^{n} |\pi_{n-i+1}(K^\varepsilon) \cap \pi_{n-i+1}(\mathcal{L})|,$$

where $\pi_1, \dots, \pi_n$ are the Gram-Schmidt projections of B.

Proof. Given the high level description above, to fully describe the algorithm, it remains to describe how we compute all feasible extensions of a giving partial assignment. In the algorithm, we will guarantee that we enumerate over all partial assignments feasible for K, while enumerating at most over all partial assignments feasible for K^ε.

Extending a partial assignment. Let $\mathbf{b}_1^*, \dots, \mathbf{b}_n^*$ be the associated dual basis for B. Assume that we are at a level i recursion node, $0 \leq i \leq n$, with an associated partial assignment $z_{n-i+1}, \dots, z_n \in \mathbb{Z}$. To begin processing this node, we first check that the partial assignment is feasible. Letting $\mathbf{t} = \sum_{j=n-i+1}^{n} z_j \mathbf{b}_j$, and

$$d = \min_{\mathbf{x} \in K} \|\pi_{n-i+1}(\mathbf{x} - \mathbf{t})\|_2,$$

we use Theorem 7 to compute $d' \in \mathbb{R}$ satisfying $d' \leq d \leq d' + \varepsilon$. If $d' > 0$, we conclude that the partial assignment is infeasible for K and terminate the node, and if $d' \leq 0$, we conclude that it is feasible for K^ε and continue.

If $i = n$, we output the lattice point $\sum_{i=1}^{n} z_i \mathbf{b}_i \in K^\varepsilon \cap \mathcal{L}$ and terminate the node. If $i < n$, we now compute the possible feasible extensions with $z_{n-i} \in \mathbb{Z}$, where we shall guarantee that all integral extensions feasible for K are found and that all examined extensions are feasible for K^ε. Let $\bar{\mathbf{b}}_{n-i}^* = \pi_{n-i+1}(\mathbf{b}_{n-i}^*)$ and $\hat{\mathbf{b}}_{n-i}^* = \mathbf{b}_{n-i}^* - \bar{\mathbf{b}}_{n-i}^*$. Set $M = 4\|\hat{\mathbf{b}}_{n-i}^*\|R/\varepsilon$, and let

$$u = \max_{\mathbf{x} \in K} \left\langle \bar{\mathbf{b}}_{n-i}^*, \mathbf{t} \right\rangle + \left\langle \hat{\mathbf{b}}_{n-i}^*, \mathbf{x} \right\rangle - M\|\pi_{n-i+1}(\mathbf{x} - \mathbf{t})\|_2$$

$$l = \min_{\mathbf{x} \in K} \left\langle \bar{\mathbf{b}}_{n-i}^*, \mathbf{t} \right\rangle + \left\langle \hat{\mathbf{b}}_{n-i}^*, \mathbf{x} \right\rangle + M\|\pi_{n-i+1}(\mathbf{x} - \mathbf{t})\|_2 .$$

Using Theorem 7, we compute $u' \in \mathbb{R}$ satisfying $u \leq u' \leq u + \varepsilon\|\hat{\mathbf{b}}_{n-i}^*\|_2/2$ and $l' \in \mathbb{R}$ satisfying $l \geq l' \geq l - \varepsilon\|\hat{\mathbf{b}}_{n-i}^*\|_2/2$. We now recurse on the integral extensions $z_{n-i} \in \{z \in \mathbb{Z} : l' \leq z \leq u'\}$.

Correctness. We must guarantee that the above algorithm correctly returns a set of points between $K \cap \mathcal{L}$ and $K^\varepsilon \cap \mathcal{L}$. Due to lack of space, we defer this analysis to the full version of the paper.

Complexity analysis. To bound the runtime of the above algorithm, we remark that the work done at each node in the recursion tree is polynomial (noting that the work enumerating $\{z \in \mathbb{Z} : l' \leq z \leq u'\}$ can be charged to a node's children), hence it suffices to bound the number of nodes in the tree. Given the above analysis, for each i, $0 \leq i \leq n$, the nodes

are level i are each associated with a distinct point in $\pi_{n-i+1}(K^\varepsilon) \cap \pi_{n-i+1}(\mathcal{L})$. Hence, the complexity of the algorithm is indeed polynomial times $\sum_{i=0}^{n} |\pi_{n-i+1}(K^\varepsilon) \cap \pi_{n-i+1}(\mathcal{L})|$, as needed. ◀

Motivated by the above lemma, we define the following measure of enumeration complexity.

▶ **Definition 11** (Schnorr-Euchner Enumerable). A convex body $K \subseteq \mathbb{R}^n$ is α-*Schnorr-Euchner enumerable*, or α-SE, with respect to a basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ for $\mathcal{L}$ (or vice versa) if for every shift $\mathbf{t}$, $\mathbf{t} \in \mathbb{R}^n$, and level i, $i \in \{1, \ldots, n\}$, we have that $|\pi_{n-i+1}(K + \mathbf{t}) \cap \pi_{n-i+1}(\mathcal{L})| \leq \alpha$, i.e. the number of distinct feasible partial assignments for $K + \mathbf{t}$ with respect to B at level i is bounded by α.

As explained previously, the total number of feasible partial assignment controls the essential complexity of Schnorr-Euchner enumeration. The usefulness of the α-SE property for K is that it will enable us to bound the complexity of Schnorr-Euchner enumeration for general convex sets via their *covering numbers* with respect to K.

▶ **Definition 12** (Covering Numbers). For two sets $C, D \subseteq \mathbb{R}^n$, we denote the *covering number* of C with respect to D

$$N(C, D) = \min\{|T| : T \subseteq \mathbb{R}^n, C \subseteq T + D\} \ .$$

C, D have covering numbers bounded by (c_1, c_2) if $N(C, D) \leq c_1$ and $N(D, C) \leq c_2$.

The following corollary, which will be crucial to making our volume algorithm efficient, is immediate:

▶ **Corollary 13.** *Let $K \subseteq \mathbb{R}^n$ be a convex body and $\mathcal{L} \subseteq \mathbb{R}^n$ be a full rank lattice with basis B. Assume that K is α-SE with respect to B. Then for any convex body $C \subseteq \mathbb{R}^n$, C is $\alpha N(C, K)$-SE with respect to B. In particular, if C is centered and equipped with a weak membership oracle, then for any $\varepsilon' > 0$ and $\mathbf{t} \in \mathbb{R}^n$, a set S satisfying $(C + t) \cap \mathcal{L} \subseteq S \subseteq (C^{\varepsilon'} + \mathbf{t}) \cap \mathcal{L}$ can be enumerated using polynomial space in time polynomial times $\alpha \cdot N(C, K)$.*

To help make the above bounds effective, we will use the fact that covering numbers for convex bodies are tightly controlled by volumes. We note that we will generally be use these estimates with respect to different scalings of the same convex body (or one of its symmetrizations).

▶ **Theorem 14** (Covering Bounds [24]). *For convex bodies $C, D \subseteq \mathbb{R}^n$, we have that*

$$\frac{\mathrm{vol}_n(C - D)}{\mathrm{vol}_n(D - D)} \leq N(C, D) \leq n(\log n + \log\log n + 5) \frac{\mathrm{vol}_n(C - D)}{\mathrm{vol}_n(D)}.$$

The next lemma two lemmas will enable us to get the main estimates we will use to bound SE-complexity.

▶ **Lemma 15.** *Let $K \subseteq \mathbb{R}^n$ be a convex body, and let $\mathcal{L} \subseteq \mathbb{R}^n$ be a full rank lattice with basis $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$. Then K is $N(K, \mathcal{P}(B))$-SE with respect to B.*

Proof. Let $T \subseteq \mathbb{R}^n$ satisfy $K \subseteq T + \mathcal{P}(B)$ and $|T| = N(K, \mathcal{P}(B))$. Letting $\pi_1, \ldots, \pi_n$ denote the Gram-Schmidt projections of B, it is easy to check that $\pi_i(\mathcal{P}(B))$, $i \in [n]$, is the parallelepiped of the basis $\pi_i(\mathbf{b}_i), \ldots, \pi_i(\mathbf{b}_n)$ for $\pi_i(\Lambda)$, and hence is a fundamental domain of $\pi_i(\Lambda)$. Given this, for each $\mathbf{x} \in T$, $|\pi_i(\mathbf{x} + \mathcal{P}(B)) \cap \pi_i(\Lambda)| = 1$. Since $\pi_i(T + \mathcal{P}(B))$ covers $\pi_i(K) \cap \pi_i(\Lambda)$, we deduce that $|\pi_i(K) \cap \pi_i(\Lambda)| \leq |T|$. Hence, K is $|T|$-SE as needed. ◀

▶ **Lemma 16** (Robustness of SE-complexity). *Let $K \subseteq \mathbb{R}^n$ be a convex body, $\mathcal{L} \subseteq \mathbb{R}^n$ be a full rank lattice with basis B. If K is α-SE with respect to B, then given a basis $\widetilde{B}$ of*
1. *$\mathcal{L}' \subseteq \mathcal{L}$, a full rank sublattice, a basis B' of $\mathcal{L}'$ for which K is α-SE*
2. *$\mathcal{L} \subseteq \mathcal{L}'$, a full rank superlattice, a basis B' of $\mathcal{L}'$ for which K is $\alpha \cdot [\mathcal{L}' : \mathcal{L}]$-SE*
can be computed in polynomial time.

6 Lattice Packing and Covering

We now present some additional relevant lattice concepts. We refer the reader to book [14] for a comprehensive reference.

For a symmetric convex body K, we define $\|\mathbf{x}\|_K = \inf\{s \geq 0 : \mathbf{x} \in sK\}$ as the norm induced by K, which satisfies all norm properties.

▶ **Definition 17** (Lattice Packing). A measurable set $A \subseteq \mathbb{R}^n$ packs with respect to a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ (and vice versa) if the translates $\mathbf{y} + A$, $\mathbf{y} \in \mathcal{L}$, are mutually disjoint.

The packing induced by A and $\mathcal{L}$ is α-dense if $\mathrm{vol}_n(A)/\det(\mathcal{L}) \geq \alpha$. We note that packing density is always less than 1.

▶ **Definition 18** (Minimum Distance). For a symmetric convex body $K \subseteq \mathbb{R}^n$ and full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$, we denote $\lambda_1(K, \mathcal{L}) = \min_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \|\mathbf{y}\|_K$, the minimum distance of $\mathcal{L}$ under $\| \cdot \|_K$ (length of shortest non-zero vector).

▶ **Definition 19** (Packing and Covering Radius). Let $K \subseteq \mathbb{R}^n$ be a convex body and $\mathcal{L} \subseteq \mathbb{R}^n$ be a full rank lattice.

Let $\varrho(K, \mathcal{L}) = \lambda_1(K - K, \mathcal{L})$ denote the packing radius of K with respect to $\mathcal{L}$. K° packs with respect to $\mathcal{L}$ iff $\varrho(K, \mathcal{L}) \geq 1$. If K is symmetric $\varrho(K, \mathcal{L}) = \lambda_1(K, \mathcal{L})/2$.

Let $\mu(K, \mathcal{L}) = \inf\{s \geq 0 : \mathcal{L} + sK = \mathbb{R}^n\}$ denote the covering radius of K with respect to $\mathcal{L}$. K covers with respect to $\mathcal{L}$ iff $\mu(K, \mathcal{L}) \leq 1$.

▶ **Lemma 20.** *Let $K \subseteq \mathbb{R}^n$ be a convex body and let $\mathcal{L} \subseteq \mathbb{R}^n$ be a full rank lattice. Then, if K covers with respect $\mathcal{L}$ and $\varrho(K, \mathcal{L}) \geq 1/\beta$, $\beta > 0$, then the covering induced by K and $\mathcal{L}$ is β^n-thin.*

Proof. By assumption $\varrho(K, \mathcal{L}) \geq 1/\beta$, and hence $(K/\beta)^\circ$ packs with respect to $\mathcal{L}$. In particular, $\mathrm{vol}_n(K/\beta) \leq \det(\mathcal{L})$. Therefore, the thinness is covering induced by K and $\mathcal{L}$ is bounded by $\mathrm{vol}_n(K)/\det(\mathcal{L}) \leq \mathrm{vol}_n(K)/\mathrm{vol}_n(K/\beta) = \beta^n$, as needed. ◀

7 Thin Covering Lattices

Our main technical contribution is a deterministic construction for thin covering lattices with good Schnorr-Euchner enumeration properties. We state its guarantees below.

▶ **Theorem 21** (Thin Lattice). *Let $K \subseteq \mathbb{R}^n$ be $(\mathbf{a}_0, r, R)$-centered convex body given by a weak membership oracle. Then, there is a deterministic $2^{O(n)}$-time and $\mathrm{poly}(n)$-space algorithm that constructs a basis B for a full rank lattice $\mathcal{L} \subseteq \mathbb{R}^n$ and a point $\mathbf{c} \in K$, satisfying*
1. *$\mathbf{c}$ is a $(6/7)^n$-approximate KB point for K and $K[\mathbf{c}]$ is $(\mathbf{c}, r/(30n), 2R)$-centered.*
2. *$K[\mathbf{c}]$ covers with respect to $\mathcal{L}$ and has packing radius $\varrho(K[\mathbf{c}], \mathcal{L}) \geq 1/3$.*
3. *$K[\mathbf{c}]$ is $2^{O(n)}$-SE with respect to B.*

▶ **Remarks**. If K is symmetric, we can specialize the above theorem by setting $\mathbf{c} = \mathbf{0}$, in which case $K[\mathbf{c}] = K$. By Lemma 20, in the above theorem, we have that $\mathcal{L}$ forms a 3^n-thin

covering with respect to $K[\mathbf{c}]$. Next, since $K[\mathbf{c}] \subseteq K - \mathbf{c}$, $\mathcal{L}$ also covers with respect to K. In particular, the thinness of the covering induced by K and $\mathcal{L}$ is bounded by

$$\mathrm{vol}_n(K)/\det(\mathcal{L}) = \mathrm{vol}_n(K)/\mathrm{vol}(K[\mathbf{c}]) \cdot \mathrm{vol}_n(K[\mathbf{c}])/\det(\mathcal{L}) \leq 2^n (7/6)^n 3^n = 7^n.$$

Hence $\mathcal{L}$ is $2^{O(n)}$-thin covering lattice for both $K[\mathbf{c}]$ and K.

Volume estimation.   We now use the above construction to prove our main volume estimation result.

Proof of Theorem 1 (Volume Estimation). We wish to compute V such that $\mathrm{vol}_n(K) \leq V \leq (1+\varepsilon)^n \mathrm{vol}_n(K)$ for $0 < \varepsilon < 1$, where $K \subseteq \mathbb{R}^n$ is a $(\mathbf{a}_0, r, R)$-centered convex body given by a weak membership oracle.

To begin, we construct the lattice $\mathcal{L}$ with basis B, and point $\mathbf{c} \in K$ as guaranteed by Theorem 21. From here, we construct a weak membership oracle O_C for $C = K + (\varepsilon/3)K[\mathbf{c}]$ from the weak membership oracle for K (see [13] for details). Note that $K[\mathbf{c}]$ is $(\mathbf{c}, r/(30n), 2R)$-centered and C is $(\mathbf{a}_0, r, 2R)$-centered. From here, letting $\varepsilon' = \varepsilon r/(180n)$, we use Corollary 13 on inputs C, $(\varepsilon/3)\mathcal{L}$, $(\varepsilon/3)B$ and ε' to enumerate S, satisfying

$$(K + (\varepsilon/3)K[\mathbf{c}]) \cap (\varepsilon/3)\mathcal{L} \subseteq C^{\varepsilon'} \cap (\varepsilon/3)\mathcal{L} \subseteq (K + (\varepsilon/2)K[\mathbf{c}]) \cap (\varepsilon/3)\mathcal{L}$$

in time

$$2^{O(n)} N(K + (\varepsilon/2)K[\mathbf{c}], (\varepsilon/3)K[\mathbf{c}]) = 2^{O(n)} (1 + 1/\varepsilon)^n,$$

where the last inequality follows by Theorem 14. From here, we return $V = |S| \det(\mathcal{L})(\varepsilon/3)^n$ (note that we need only count each element of S as it is outputted, which requires only polynomial space). The fact that V satisfies the required bounds follows directly from the discussions in section 4 (see Equation (4)).   ◀

Comparison with prior constructions.   Much work has been dedicated to proving the existence of extremely thin-lattice coverings [21, 20, 22, 4, 10] – much of instigated by C.A. Rogers – where the best construction [22] provides $n^{\log n + O(1)}$-thin coverings for any convex body K.

All of these constructions rely on sampling from a probabilistic ensembles of lattices, occasionally with some additional post processing, and are intrinsically difficult to derandomize. More problematically however, these ensembles produce lattices that are as "hard as possible" (see for example, section 2 in [3]) to enumerate from with known polynomial space methods, severely complicating their use in our context (and in many others in fact).

Given the above discussion, the construction in Theorem 21 gives the *first existential construction* of "easy to enumerate" thin-covering lattices for general convex bodies. As an added bonus of our construction, when the convex body K is symmetric, the covering lattice we construct has packing radius at least $1/3$ and has the property that CVP under the norm $\| \cdot \|_K$ can be solved in $2^{O(n)}$ time and poly(n) space (since this reduces to enumeration inside shifts of K). While building thin covering lattices for ℓ_p norms is trivial – $2n^{-1/p}\mathbb{Z}_n$ is a $2^{O(n)}$-thin covering lattice for the ℓ_p norm – building ones with packing radius $\Omega(1)$. In fact, even for the ℓ_2 norm, there is no known explicit construction of such a lattice. While the packing radius property is not necessary in our main application, we believe it might be useful elsewhere, such as in lattice based schemes for Locality Sensitive Hashing (see [2] for an application using the 24-dimensional Leech lattice).

The only previous algorithmic construction is due Alon et al [1], whose gave a deterministic $2^{O(n)}$-time and 2^n-space thin-lattice construction for *symmetric bodies* based on a greedy construction of Rogers [21] – which our construction is also based on – along with a 2^n space enumeration method. For their enumeration technique, they rely on the M-ellipsoid covering and Voronoi cell based enumeration algorithms of [17, 7, 5]. With these techniques, starting with a thin covering lattice $\mathcal{L}$ for a symmetric convex body K, they can enumerate $\mathcal{L} \cap C$, for any convex body C, in time $2^{O(n)} N(C, K)$ using 2^n space. Hence, the enumeration guarantees are similar to ours, though at the cost of exponential space.

Rogers' greedy construction. We now describe Roger's method and our related improvements. This construction starts with essentially any lattice $\mathcal{L}$ and symmetric convex body K such that $\varrho(K, \mathcal{L}) \geq 1$. The constructions proceeds by iteratively making $\mathcal{L}$ denser, by adding points in $\mathcal{L}/3$ to $\mathcal{L}$, while guaranteeing that the packing radius with respect to K stays at least 1. This will have the net effect of increasing the packing density by 3. Since the packing density cannot increase indefinitely (it can never go above 1), the densification process eventually stops, at which point one can conclude that the final lattice $\mathcal{L}$, after a factor 3 scaling, covers with respect to K and has packing radius at least 1/3.

A first main problem is that even if we initialize Rogers' construction with an "easy to enumerate" lattice $\mathcal{L}$, the final generated lattice maybe so far away from the initial lattice that it loses the easy enumeration property. To avoid this problem, we show that if we start the procedure with an easy to enumerate dense packing lattice for K, then the procedure converges fast enough for the final generated lattice to retain the easy enumeration property. To build the initial dense packing lattice, we begin with a lattice $\mathcal{L}$ with basis B derived from the axes of an M-ellipsoid of K, which satisfies that B is $2^{O(n)}$-SE with respect to K, that we then subsequently sparsify it, using techniques of [6], to make it induce a $2^{-O(n)}$-dense packing with respect to K.

A second problem with Roger's greedy construction is that it only directly works for symmetric bodies. In particular, if we start with an asymmetric convex body K, the final generated lattice will only be guaranteed to cover with respect to $K - K$ and not K (here, the only known relation is that $\mu(K, \mathcal{L}) \leq n\mu(K - K, \mathcal{L})$, which is far too weak). To circumvent this problem, we symmetrize K about an approximate KB point using an efficient algorithm to construct such points. Our algorithm to construct approximate KB points will in fact rely on many iterated calls of our volume algorithm and thin-lattice construction for symmetric convex bodies.

8 Techniques

We now detail the main ideas behind our thin lattice construction. We begin by describing our thin lattice construction for symmetric convex bodies, and continue with our algorithm computing approximation Kovner-Besicovitch points. We recover our full thin lattice construction (Theorem 21) by combining these two algorithms. Due to lack of space, we defer most proofs to the full version of the paper.

8.1 Thin Lattice Construction

We now describe our construction of thin covering lattices for symmetric convex bodies, corresponding to parts 2 and 3 of Theorem 21.

The construction will proceed in three stages. In the first stage, we build a base lattice Λ with a basis B derived from the axes of an M-ellipsoid E of K, for which K is $2^{O(n)}$-SE. In

the second stage, we sparsify the base lattice Λ so that it becomes a $2^{-O(n)}$-dense packing lattice $\mathcal{N}$ for K using techniques from [6]. In the last stage, we densify $\mathcal{N}$ using Rogers' procedure to derive the final $2^{O(n)}$-thin covering lattice $\mathcal{L}$.

Through these stages, our goal will be to guarantee that the "distance" of the base Λ to the final lattice $\mathcal{L}$, quantified by the product of indexes $[\Lambda : \mathcal{N}] \cdot [\mathcal{L} : \mathcal{N}]$, is bounded by $2^{O(n)}$. Having achieved this, the robustness of SE-complexity (see Lemma 16) will allow us to construct a basis for $\mathcal{L}$ with respect to which K is $2^{O(n)}$-SE. We now detail the main arguments underlying each stage.

M-Lattice. For the first stage, we define the basis B of Λ, so that $\mathcal{P}(B) \subseteq E$ is a maximum volume inscribed parallelepiped, where E an M-ellipsoid for K. Here it is not hard to check that $\text{vol}_n(E)/\mathcal{P}(B) = \text{vol}_n(B_2^n)/\text{vol}_n([-1/\sqrt{n}, 1/\sqrt{n})^n) = 2^{O(n)}$. Given this, $\mathcal{P}(B)$ inherits the covering properties of E with respect to K, in particular, $N(K, \mathcal{P}(B)), N(\mathcal{P}(B), K) = 2^{O(n)}$. In particular, $\det(\Lambda) = \text{vol}_n(\mathcal{P}(B)) = 2^{\Theta(n)}\text{vol}_n(K)$, and, by Lemma 15, K is $2^{O(n)}$-SE with respect to B.

Packing lattice. For the second stage, to make Λ a packing lattice, it suffices to "remove" all the lattice points in $\Lambda \cap 2K \setminus \{\mathbf{0}\}$ (by symmetry of K). By the covering properties, $|\Lambda \cap 2K| \leq N(2K, K) \cdot N(K, \mathcal{P}(B)) = 2^{O(n)}$, and hence, we may expect to find a sublattice $\mathcal{N}$ such that $[\Lambda : \mathcal{N}] = 2^{O(n)}$ and $\mathcal{N} \cap 2K = \{\mathbf{0}\}$. Indeed, a simple expectation argument shows that a "random" sublattice $\mathcal{N}$ of index $2^{O(n)}$ avoids all the non-zero points in $\Lambda \cap 2K$ with good probability (see [6]). Furthermore, one can find the sublattice $\mathcal{N}$ deterministically using the method of conditional expectations. Since $\mathcal{N}$ is a sublattice, note that by Lemma 16, a basis of $\mathcal{N}$ can be computed for which the SE-complexity of K does not increase compared to B. To see that $\mathcal{N}$ induces a $2^{-O(n)}$-dense packing for K, note that

$$\det(\mathcal{N}) = [\Lambda : \mathcal{N}] \det(\Lambda) = 2^{O(n)} \det(\Lambda)$$
$$= 2^{O(n)}\text{vol}_n(\mathcal{P}(B)) = 2^{O(n)}\text{vol}_n(K), \quad \text{as needed.}$$

Rogers' procedure. For the last stage, we initially set $\mathcal{L} \leftarrow \mathcal{N}$ and then iteratively densify $\mathcal{L}$ to get a $2^{O(n)}$-thin covering lattice. By assumption, $\mathcal{L}$ starts as a packing lattice for K, or equivalently, $\mathcal{L}$ has minimum distance $\lambda_1(K, \mathcal{L}) \geq 2$. To make $\mathcal{L}$ denser, we look for a point $\mathbf{x} \in \mathcal{L}/3$ at distance greater than 2 from $\mathcal{L}$ under $\|\cdot\|_K$. If such a point $\mathbf{x}$ is found, we set $\mathcal{L} \leftarrow \mathcal{L} + \{0, \pm\mathbf{x}\}$. By the distance assumption and symmetry of K, we maintain the invariant $\lambda_1(K, \mathcal{L}) \geq 2$, while decreasing the determinant by a factor 3.

Note that each successful iteration increases the packing density by 3. Since the packing density starts at $2^{-O(n)}$, this process must terminate in at most $O(n)$ steps. In particular, after termination, we have that $[\mathcal{L} : \mathcal{N}] = 3^{O(n)}$, and hence by Lemma 16 and our assumptions on $\mathcal{N}$, we can compute a B basis of $\mathcal{L}$ for which K is $2^{O(n)}$-SE (indeed, this can be done at every iteration). Next, at termination, we must have that every point in $\mathcal{L}/3$ is at distance less than 2 from $\mathcal{L}$. From here, it is not hard to show that every point in $\mathbb{R}^n$ is at distance at most $(3/2) \cdot 2 = 3$, i.e. $\mu(K, \mathcal{L}) \leq 3$. We can therefore return $\mathcal{L}/3$ as our covering lattice, which will have packing radius at least $1/3$ as desired.

The last detail is to show that at each stage, we can find a "far away" point in $\mathcal{L}/3$ or decide that none exists in $2^{O(n)}$-time. By the above discussion, we can assume that at the current stage, we have a basis B for $\mathcal{L}$ for which K is $2^{O(n)}$-SE. From here, it is easy to see that there is a point in $\mathcal{L}/3$ at distance greater than 2 iff there exists $\mathbf{x} \in B\{0, \pm 1/3\}^n$ (yielding representatives for each coset in $(\mathcal{L}/3)/\mathcal{L}$) at distance greater than 2. Since a point

$\mathbf{x} \in \mathbb{R}^n$ is at distance greater than 2 from $\mathcal{L}$ iff $(\mathbf{x} + 2K) \cap \mathcal{L} = \emptyset$, one can test this property for any $\mathbf{x}$ in $2^{O(n)}$-time using Schnorr-Euchner enumeration. Repeating this test 3^n times for each point in $B\{0, \pm 1/3\}^n$ yields the result.

This completes our description thin covering lattice constructions for symmetric bodies.

8.2 Computing Approximate Kovner-Besicovitch Points

We state the guarantees for our algorithm computing approximate KB points below.

▶ **Theorem 22.** *Let $K \subseteq \mathbb{R}^n$ be a $(\mathbf{a}_0, r, R)$-centered convex body given by a weak membership oracle. Then, for $\varepsilon > 0$, one can compute a $(1 + \varepsilon)^{-n}$ approximate Kovner-Besicovitch point $\mathbf{c} \in K$, such that $K[\mathbf{c}]$ is $(\mathbf{c}, \varepsilon r/(5n), 2R)$-centered, in deterministic $2^{O(n)}(1 + 1/\varepsilon)^{2n+1}$ time and $\mathrm{poly}(n)$ space.*

▶ **Remark.** Part 1 of Theorem 21 follows by applying the Theorem 22 to K with $\varepsilon = 1/6$.

High level algorithm. First, by applying a suitable linear affine transformation to K (i.e. standard ellipsoidal rounding), we may assume that $B_2^n \subseteq K \subseteq (n+1)n^{1/2}B_2^n$. We now define the sequence of bodies $K_i = 2^i B_2^n \cap K$, for $i \in \{0, \dots, T\}$, $T = O(\log n)$, where $K_0 = B_2^n$ and $K_T = K$. For each K_i, $i \in [T-1]$, we will compute a 3^{-n} approximate KB point $\mathbf{c}_i$ for K_i from a 3^{-n}-approximate KB point $\mathbf{c}_{i-1}$ for K_{i-1}. Finally, in the last step, from K_{T-1} to K_T, we amplify this to $(1 + \varepsilon)^{-n}$ approximation. We note that we may start with $\mathbf{c}_0 = \mathbf{0}$, since this is the center of symmetry for $K_0 = B_2^n$. Furthermore, at each step, since the volume $\mathrm{vol}_n(K_i) \leq 2^n \mathrm{vol}_n(K_{i-1})$, the KB value of c_{i-1} with respect to K_i, $i \in [T]$, is at least $2^{-n} \cdot 3^{-n} \cdot 2^{-n} = 12^{-n}$.

To compute $\mathbf{c}_i$ starting from $\mathbf{c}_{i-1}$, we perform the following improvement steps: from our current solution for $\mathbf{c}_i$ (initialized at $\mathbf{c}_{i-1}$ during the first iteration), we begin by building a thin-covering lattice $\mathcal{L}$ with basis B for $K_i[\mathbf{c}_i]$ (note $K_i[\mathbf{c}_i]$ is symmetric). We then construct a covering of $(1/2)(K_i + \mathbf{c}_i)$ by $(\varepsilon/2)K_i[\mathbf{c}_i]$, whose centers are computed by enumerating $S = (1/2)((K_i + \varepsilon K_i[\mathbf{c}_i] + \mathbf{c}_i) \cap \varepsilon \mathcal{L})$ via Schnorr-Euchner enumeration using B. We then replace $\mathbf{c}_i$ by the element in S (noting that $S \subseteq K_i$) of largest approximate KB value, where for each $\mathbf{x} \in S$ we approximate $\mathrm{vol}_n(K_i[\mathbf{x}])$ to within $(1 + \varepsilon/10)^n$ using the volume algorithm for symmetric convex bodies. The concavity of the function $\mathrm{vol}_n(K[\mathbf{x}])^{1/n}$ will allow us to show that at each step, we improve the objective value by essentially a $(1 + c\varepsilon)^n$ factor. Hence $O(1/\varepsilon)$ iterations suffice to construct a near optimal solution.

References

1 N. Alon, A. Schraibman, T. Lee, and S. Vempala. The approximate rank of a matrix and its algorithmic applications. In *STOC*, 2013.

2 A. Andoni and P. Indyk. Near-optimal hashing algorithms for approximate nearest neighbor in high dimensions. In *FOCS*, pages 459–468, 2006.

3 A. Becker, N. Gama, and A. Joux. Solving shortest and closest vector problems: The decomposition approach. Cryptology Eprint. Report 2013/685, 2013.

4   G. J. Butler. Simultaneous packing and covering in euclidean space. *Proceedings of the London Mathematical Society*, 25(3):721–735, 1972.

5   D. Dadush. *Integer Programming, Lattice Algorithms, and Deterministic Volume Estimation*. PhD thesis, Georgia Institute of Technology, 2012.

6   D. Dadush and G. Kun. Lattice sparsification and the approximate closest vector problem. In *SODA*, 2013.

7   D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, 2011.

8   D. Dadush and S. Vempala. Near-optimal deterministic algorithms for volume computation via m-ellipsoids. *Proceedings of the National Academy of Sciences*, 2013.

9   M. E. Dyer, A. M. Frieze, and R. Kannan. A random polynomial-time algorithm for approximating the volume of convex bodies. *J. ACM*, 38(1):1–17, 1991. Preliminary version in STOC 1989.

10 U. Erez, S. Litsyn, and R. Zamir. Lattices which are good for (almost) everything. *IEEE Transactions on Information Theory*, 51(10):3401–3416, 2005.

11 Z. Füredi and I. Bárány. Computing the volume is difficult. In *STOC*, pages 442–447, New York, NY, USA, 1986. ACM.

12 Z. Füredi and I. Bárány. Approximation of the sphere by polytopes having few vertices. *Proceedings of the AMS*, 102(3), 1988.

13 M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer-Verlag, 1988.

14 P. M. Gruber. *Convex and discrete geometry*, volume 336. Springer Science & Business Media, 2007.

15 B. Grünbaum. Measures of symmetry for convex sets. In *Proceedings of the 7th Symposium in Pure Mathematics of the American Mathematical Society, Symposium on Convexity*, pages 233–270, 1961.

16 G. Hanrot and D. Stehlé. Improved analysis of Kannan's shortest lattice vector algorithm. In *CRYPTO*, pages 170–186, Berlin, Heidelberg, 2007. Springer-Verlag.

17 D. Micciancio and P. 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. Preliminary version in STOC 2010.

18 V. D. Milman. Inégalités de Brunn-Minkowski inverse et applications at la theorie locales des espaces normes. *C. R. Math. Acad. Sci. Paris*, 302(1):25–28, 1986.

19 V. D. Milman and A. Pajor. Entropy and asymptotic geometry of non-symmetric convex bodies. *Advances in Mathematics*, 152(2):314–335, 2000.

20 C. A. Rogers. Lattice coverings of space: The Minkowski-Hlawka theorem. *Proceedings of the London Mathematical Society*, s3-8(3):447–465, 1958.

21 C. A. Rogers. A note on coverings and packings. *Journal of the London Mathematical Society*, s1-25(4):327–331, 1950.

22 C. A. Rogers. Lattice coverings of space. *Mathematika*, 6:33–39, 6 1959.

23 C. A. Rogers and G. C. Shephard. The difference body of a convex body. *Archiv der Mathematik*, 8:220–233, 1957.

24 C. A. Rogers and C. Zong. Covering convex bodies by translates of convex bodies. *Mathematika*, 44:215–218, 6 1997.

25 D. B. Yudin and A. S. Nemirovski. Evaluation of the information complexity of mathematical programming problems (in russian). *Ekonomika i Matematicheskie Metody*, 13(2):3–45, 1976.