2016 IEEE 57th Annual Symposium on Foundations of Computer Science

Towards Strong Reverse Minkowski-type Inequalities for Lattices

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

Oded Regev
Courant Institute of Mathematical Sciences
New York University
New York, NY

***Abstract*—We present a natural reverse Minkowski-type inequality for lattices, which gives *upper bounds* on the number of lattice points in a Euclidean ball in terms of sublattice determinants, and conjecture its optimal form. The conjecture exhibits a surprising wealth of connections to various areas in mathematics and computer science, including a conjecture motivated by integer programming by Kannan and Lovász (Annals of Math. 1988), a question from additive combinatorics asked by Green, a question on Brownian motions asked by Saloff-Coste (Colloq. Math. 2010), a theorem by Milman and Pisier from convex geometry (Ann. Probab. 1987), worst-case to average-case reductions in lattice-based cryptography, and more. We present these connections, provide evidence for the conjecture, and discuss possible approaches towards a proof. Our main technical contribution is in proving that our conjecture implies the ℓ_2 case of the Kannan and Lovász conjecture. The proof relies on a novel convex relaxation for the covering radius, and a rounding procedure based on "uncrossing" lattice subspaces.**

***Keywords*-lattices; geometry; Minkowski's first theorem**

I. Introduction

A lattice $\mathcal{L} \subset \mathbb{R}^n$ is defined as the set of all integer linear combinations of n linearly independent vectors $B = (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ in $\mathbb{R}^n$, where we call B a basis for $\mathcal{L}$. The determinant of $\mathcal{L}$ is defined as $\det(\mathcal{L}) = |\det(B)|$, which is invariant to the choice of basis for $\mathcal{L}$. The determinant measures the "global density" of the lattice. More precisely, $\det(\mathcal{L})^{-1}$ is the asymptotic number of lattice points per unit volume.

One of the earliest and most important results in this area is Minkowski's (First) Theorem from 1889, which guarantees the existence of lattice points in any large enough symmetric convex body. In particular, it implies that any lattice $\mathcal{L}$ with $\det(\mathcal{L}) \leq 1$ must contain exponentially in n many points of Euclidean norm at most $\sqrt{n}$. Informally, Minkowski's theorem can be described as saying that "global density implies local density".

The starting point of our investigation is an attempt to *reverse* this implication. Assume $\mathcal{L}$ has local density, i.e., exponentially many points of norm at most $\sqrt{n}$. Can we conclude that it must also have global density, i.e., $\det(\mathcal{L}) \leq 1$? A moment's thought reveals that we cannot hope for such a strong conclusion. Indeed, consider the lattice $\mathcal{L}$ generated by the basis $\{\mathbf{e}_1, \ldots, \mathbf{e}_{n-1}, M\mathbf{e}_n\}$ where M is arbitrarily large. Then $\det(\mathcal{L}) = M$ yet $\mathcal{L}$ has exponentially many points of norm at most $\sqrt{n}$.

A natural conjecture would therefore say that local density implies global density *in a subspace*, or in other words, that there is a (not necessarily full-rank) sublattice with low determinant. This is the essence of our main conjecture.

To make this more precise, consider the lattice $\mathbb{Z}^n$. Obviously, all its sublattices have determinant at least 1 (as the determinant of any integer matrix is integer). Moreover, it is not difficult to see that the number of integer points in a ball of radius r is approximately n^{r^2} for $r \ll \sqrt{n}$. Our main conjecture basically says that $\mathbb{Z}^n$ *has the highest local density among all lattices whose sublattices all have determinant at least* 1. See Conjecture III.1 for the formal statement, which also makes the connection to the so-called smoothing parameter explicit.

Conjecture I.1. *(Main conjecture, informal) Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice whose sublattices all have determinant at least* 1. *Then, for any $r > 0$, the number of lattice points of Euclidean norm at most r is at most* $\exp(\operatorname{poly}\log n \cdot r^2)$.

While one can consider other variants of the conjecture (see the stronger variant in Section IX and the weaker variants in Section V), the main conjecture is the most appealing, as described next.

Connections: The conjecture exhibits a surprising wealth of connections to various areas in mathematics, ranging from additive combinatorics to convex geometry and to heat diffusion on manifolds. We mention some of those next.

- The Kannan-Lovász (KL) conjecture [1] is a nearly tight characterization of the covering radius of a lattice with respect to a given convex body in terms of projection volumes and determinants. The main technical contribution of our paper, appearing in Section IV, is to show that the ℓ_2 version of the KL conjecture [1] is implied up to poly-logarithmic factors by the main conjecture. This implication is quite nontrivial, and a proof overview will be given below. Very briefly, we rely on a novel convex relaxation for the covering radius and a rounding strategy for the corresponding dual program to extract the relevant subspace. We remark

0272-5428/16 $31.00 © 2016 IEEE
DOI 10.1109/FOCS.2016.55

that it is possible that our main conjecture in fact implies the *general* KL conjecture.

- In 2007, it was suggested by Green [2] that a certain strong variant of the polynomial Freiman-Ruzsa conjecture over the integers is false. In work in progress of the second-named author with Lovett [3], we prove that this is indeed the case assuming the main conjecture.[1] No unconditional proof of this is currently known.
- The following question was considered by Saloff-Coste (see, e.g., [4], [5], [6] and especially [7, Problem 11]). For a lattice $\mathcal{L} \subset \mathbb{R}^n$, consider the Brownian motion on the flat torus $\mathbb{R}^n/\mathcal{L}$ starting from the origin. Its density at any time $t > 0$ is a Gaussian of standard deviation $\sqrt{t}$ reduced modulo $\mathcal{L}$. Saloff-Coste asked whether its mixing time in the total variation (or equivalently, L_1) sense is approximately the same as that in the (easier to analyze) L_2 sense. This is equivalent to asking whether the smoothing parameter of a lattice is approximately the same as the "L_1 smoothing parameter." In Section VII we provide more background and show that such a result follows from our main conjecture. We remark that this seems to be the weakest implication of the main conjecture that is already an interesting question in its own right.
- In Section VIII we describe some connections to computational complexity. We show that our main conjecture implies a very tight reduction from the problem of approximation the smoothing parameter to the problem of sampling points from a given coset of a lattice according to a Gaussian (or even subgaussian) distribution. Combined with known reductions, this shows that the hardness of SIS, a central average-case cryptographic problem, can be based on the worst-case hardness of approximation the smoothing parameter to within $\tilde{O}(\sqrt{n})$. We also mention an easy implication to the complexity of the problem of approximating the covering radius.

Evidence for the conjecture: First, as we show in an omitted section, the conjecture passes some basic sanity checks. For instance, it is true for natural families of lattices including "rectangular" lattices (or more generally, direct sum of lattices for which the conjecture holds) as well as random lattices. We also show there that a weaker bound on the number of lattice points is true (namely, with $\mathrm{poly}\log n$ replaced by $\sqrt{n}$).

Second, as we describe in an omitted section, the conjecture has a certain "continuous relaxation" which talks about volumes of sections instead of determinants of lattice subspaces. Somewhat surprisingly, that relaxation is known to be true and follows from a celebrated theorem by Milman and Pisier. This is the most non-trivial implication of the main conjecture that we know is true.

Finally, we believe that there are several approaches for proving the main conjecture itself that are worth exploring in more detail, including one based on additive combinatorics. We also tried exploring whether there are natural weaker variants of the conjecture that might be easier to prove. Those variants are described in Section V, and some are quite natural in their own right. The variant described in Section VI is particularly appealing, and we believe that it can be attacked using a Fourier analytic approach.

Related work: There has been considerable work on the problem of bounding or counting the number of lattice points in a ball, perhaps the earliest reference being Gauss's circle problem. Most of the work in this area, however, considers "large" convex bodies and shows that for such bodies the number of lattice points is very close to what one would expect by a volume heuristic. As such, this seems too coarse to capture the subspace structure highlighted by the main conjecture. See, e.g., [8] and references therein for the kind of results proved in this area since the early 20th century.

Also related is the work on *stable lattices* (sometimes known as "semistable"). These are lattices of determinant 1 whose sublattices all have determinant at least 1. Properties of stable lattices have been studied since the 1970s in connection with algebra, topology, and geometry. See [9] and references therein. In particular, Shapira and Weiss [9] consider a strong quantitative variant of the ℓ_2 version of the KL conjecture, and show that it implies the so-called Minkowski conjecture. It remains to be seen how exactly our main conjecture connects to the work on stable lattices.

Proof overview of the main technical theorem: Our main technical theorem shows how to derive the ℓ_2 case of the Kannan-Lovász (KL) conjecture from our main conjecture. To explain the KL conjecture, recall that the *covering radius* μ of a lattice is the maximum distance a point in space can be from the lattice. So for instance $\mu(\mathbb{Z}^n) = \sqrt{n}/2$. An easy lower bound on $\mu(\mathcal{L})$ is given in terms of the determinant of $\mathcal{L}$. Namely, since balls of radius $\mu(\mathcal{L})$ centered at all points of $\mathcal{L}$ cover $\mathbb{R}^n$, we obtain by volume considerations that $\mu(\mathcal{L}) \gtrsim \sqrt{n}(\det(\mathcal{L}))^{1/n}$. This bound can be far from tight, e.g., for the lattice generated by the basis $\{\mathbf{e}_1, \ldots, \mathbf{e}_{n-1}, M\mathbf{e}_n\}$ where M is large. As was the case for the main conjecture, we can tighten the bound by maximizing over subspaces, namely, we consider the maximum of $\sqrt{\dim(W)}(\det(\pi_W(\mathcal{L})))^{1/\dim(W)}$ over all subspaces W where π_W denotes the projection on W. Since $\mu(\mathcal{L}) \geq \mu(\pi_W(\mathcal{L}))$ this is clearly a lower bound on $\mu(\mathcal{L})$. The ℓ_2 KL conjecture says that this is nearly tight, i.e., that we can also *upper bound* $\mu(\mathcal{L})$ by the same maximum up to polylogarithmic terms.

As should be obvious by now, the ℓ_2 KL conjecture has a similar flavor to our main conjecture (and even more so when comparing the formal definitions, see Conjecture III.1 and Conjecture IV.1). The difference is that our conjecture

[1] We remark that for this it would suffice to prove the conclusion in Conjecture I.1 for $r = \sqrt{n}$.

tries to capture the number of points in a ball (or to be precise, the smoothing parameter), whereas ℓ_2 KL tries to capture the covering radius.

The first step in our proof is to bound from above the covering radius by a convex program we call μ_{sm}. Intuitively and informally, the convex program tries to find the "smallest" covariance matrix A such that the distribution obtained by picking a random lattice point and adding a Gaussian random variable with covariance A is close to uniform over $\mathbb{R}^n$. Here by smallest we mean the expected squared norm of a Gaussian with that covariance, which is simply $\mathrm{tr}(A)$. A Gaussian that satisfies this property is said to "smooth" $\mathcal{L}$. It is intuitively clear (and not difficult to prove formally) that this program bounds $\mu(\mathcal{L})^2$ from above, since in order for the Gaussian to smooth $\mathcal{L}$, it must "reach" all points in space, and hence its norm must be at least $\mu(\mathcal{L})$. To state this program formally, we recall that smoothness has an elegant equivalent definition (which follows from the Poisson summation formula) in terms of a Gaussian sum on dual lattice points, namely

$$\mu_{\mathrm{sm}}(\mathcal{L})^2 = \min\Big\{\mathrm{tr}(A) : A \succeq 0, \sum_{\mathbf{y}\in\mathcal{L}^*\setminus\{\mathbf{0}\}} e^{-\pi \mathbf{y}^{\mathsf{T}} A \mathbf{y}} \leq 1/2\Big\}. \tag{1}$$

At this point the natural thing to do would be to consider the dual of μ_{sm}, which is a maximization problem, and use its optimal solution to find a subspace the projection on which has large determinant, as needed for the ℓ_2 KL conjecture. Unfortunately, the dual program seems difficult to deal with since the single constraint in (1) "entangles" all the information about lattice points in a complicated way, and it is not clear how it would help in identifying such a subspace.

This is were our main conjecture comes in. We formulate another convex program μ_{det} which, assuming the main conjecture, bounds μ_{sm} from above up to polylogarithmic factors. In more detail, that program also tries to find the covariance matrix with smallest trace, but its constraints basically say that A should be such that in the dual lattice, all sublattice determinants are large (relative to A). By the main conjecture, large sublattice determinants imply small number of points in balls, which in turn, implies smoothing. (In fact, the formal statement of the main conjecture in Conjecture III.1 is already stated in terms of a Gaussian sum over lattice points as in (1), so there is no need for this detour through number of points in balls.) Since under the main conjecture, the constraints in μ_{det} imply those in μ_{sm} (i.e., are weaker), we obtain that μ_{det} bounds μ_{sm} from above up to polylogarithmic factors, as desired.

Since μ_{det} directly puts in a constraint for every subspace, the subspace structure comes out explicitly, and we are finally in position take the dual. It turns out that the dual of μ_{det} has a reasonably nice form, and a solution to it can be seen as some kind of mixture of various lattice subspaces. The last and most technically demanding part of the proof is to "round" that dual solution, i.e., we show how to take an arbitrary mixture of lattice subspaces and extract from it just one lattice subspace that is nearly as good. There are several steps to this proof, the most interesting one being a sort of "uncrossing inequality," showing that if the solution includes two subspaces V and W, we can replace them with the subspaces $V+W$ and $V\cap W$ in a way that does not decrease the goal function. We now repeat this uncrossing step over and over again. Notice that we make progress as long as there are two subspaces such that neither is contained in the other. Therefore, after sufficiently many iterations, we arrive at a chain, i.e., a sequence of subspaces $W_1 \subseteq W_2 \subseteq \cdots \subseteq W_m$. Using careful bucketing, we show that one of these subspaces must be nearly as good as the mixture, and this completes the proof.

Outline: Due to lack of space, almost all preliminaries, proofs, and some statements, are omitted. We refer the reader to the arXiv version for a more complete treatment http://arxiv.org/abs/1606.06913.

Acknowledgements: We thank Mark Rudelson for pointing us to the Milman-Pisier theorem and Chris Peikert for suggesting an early version of the weak conjecture and its relation to L_1 smoothing. We also thank them and Kai-Min Chung, Shachar Lovett, and Noah Stephens-Davidowitz for useful discussions. DD was supported by the NWO Veni grant 639.071.510. OR was supported by the Simons Collaboration on Algorithms and Geometry and by the National Science Foundation (NSF) under Grant No. CCF-1320188. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the NSF.

II. Preliminaries

We write $X \lesssim Y$ to mean that there exists a universal constant $C > 0$ such that $X \leq CY$, and similarly for $X \gtrsim Y$.

Definition II.1 (Matrix Slice)**.** *We denote $X^{\cap W}$, the slice of X on W, to be the unique PSD matrix satisfying*

$$\mathbf{y}^{\mathsf{T}}(X^{\cap W})\mathbf{y} = \min_{\mathbf{w}\in W^{\perp}} (\mathbf{y}+\mathbf{w})^{\mathsf{T}} X (\mathbf{y}+\mathbf{w}), \quad \forall \mathbf{y} \in \mathbb{R}^n .$$

If $X = \begin{pmatrix} A & C \\ C^{\mathsf{T}} & B \end{pmatrix} \in \mathbb{S}^n_+$, $A \in \mathbb{R}^{k\times k}$, $C \in \mathbb{R}^{k\times(n-k)}$, $B \in \mathbb{R}^{(n-k)\times(n-k)}$, then the slice X on $W = \mathbb{R}^k \times 0^{n-k}$ is the Schur complement of X with respect B (lifted to live in the full space), that is

$$X^{\cap W} = \begin{pmatrix} A - CB^{+}C^{\mathsf{T}} & 0^{k\times(n-k)} \\ 0^{(n-k)\times k} & 0^{(n-k)\times(n-k)} \end{pmatrix} .$$

Definition II.2 (Projected Determinant)**.** *Define the projected determinant of X on W by*

$$\det_W(X) \overset{\mathrm{def}}{=} \det(O_W^{\mathsf{T}} X O_W) \tag{2}$$

where O_W is any matrix whose columns form an orthonormal basis of W.

A d-dimensional Euclidean lattice $\mathcal{L} \subset \mathbb{R}^n$ is defined as the set of all integer linear combinations of d linearly independent vectors $B = (\mathbf{b}_1, \ldots, \mathbf{b}_d)$ in $\mathbb{R}^n$, where we call B a basis for $\mathcal{L}$. If $d = n$ we say that the lattice is full rank. The *determinant* of $\mathcal{L}$ is defined as $\det(\mathcal{L}) = \sqrt{\det(B^{\mathsf{T}}B)}$, which is invariant to the choice of basis for $\mathcal{L}$.

The dual lattice of $\mathcal{L}$ is $\mathcal{L}^* = \{\mathbf{y} \in \mathrm{span}(\mathcal{L}) : \langle \mathbf{y}, \mathbf{x} \rangle \in \mathbb{Z}, \forall \mathbf{x} \in \mathcal{L}\}$. It is easy to verify that $B(B^{\mathsf{T}}B)^{-1}$ yields a basis for $\mathcal{L}^*$ and that $\det(\mathcal{L}^*) = 1/\det(\mathcal{L})$. We say that a subspace $W \subseteq \mathbb{R}^n$ is a *lattice subspace* of $\mathcal{L} \subset \mathbb{R}^n$ if W admits a basis of vectors in $\mathcal{L}$. We will need a few important facts about lattice subspaces and projections. Firstly, the projection $\pi_W(\mathcal{L})$ onto a subspace $W \subseteq \mathbb{R}^n$ is a lattice (i.e., discrete) if and only if W is a lattice subspace of $\mathcal{L}^*$. Furthermore, $\pi_W(\mathcal{L})^* = \mathcal{L}^* \cap W$. Secondly, for two lattice subspaces V, W of $\mathcal{L}$, both the intersection $V \cap W$ and the sum $V + W$ is a lattice subspace of $\mathcal{L}$.

Definition II.3. *The* covering radius *of $\mathcal{L}$ is*

$$\mu(\mathcal{L}) = \inf\{r \geq 0 : \mathrm{span}(\mathcal{L}) \subseteq \mathcal{L} + rB_2^n\} \ .$$

For a positive definite $X \succ 0$ and a countable set $T \subseteq \mathbb{R}^n$, define

$$\rho_X(T) = \sum_{\mathbf{y} \in T} e^{-\pi \mathbf{y}^{\mathsf{T}} X^{-1} \mathbf{y}}.$$

We extend this to positive semidefinite $X \succeq 0$ by

$$\rho_X(T) = \sum_{\mathbf{y} \in T \cap \mathrm{im}(X)} e^{-\pi \mathbf{y}^{\mathsf{T}} X^{+} \mathbf{y}}.$$

Note that with the above definition, $\rho_X(T)$ is a continuous function over all of $\mathbb{S}_+^n$.

We remark that the above notation is slightly non-standard, in that we parametrize the ρ with respect to X and not $X^{1/2}$. This notation will however be more convenient for us. For $s > 0$, we will often denote $\rho_{s^2 I}$ by ρ_{s^2}.

We define $\eta_\varepsilon(\mathcal{L})$ the ε-smoothing parameter of $\mathcal{L}$ as the unique $s > 0$ satisfying

$$\rho_{1/s^2}(\mathcal{L}^*) = \sum_{\mathbf{y} \in \mathcal{L}^*} e^{-\pi \|s\mathbf{y}\|^2} = 1 + \varepsilon \ .$$

We will simply say the smoothing parameter of $\mathcal{L}$ to denote $\eta(\mathcal{L}) \overset{\mathrm{def}}{=} \eta_{1/2}(\mathcal{L})$.

III. The Main Conjecture

We now formally state the main conjecture.

Conjecture III.1. *(Main conjecture) Let $C_\eta(n) > 0$ be the smallest number such that for any $\mathcal{L} \subset \mathbb{R}^n$,*

$$\eta(\mathcal{L}) \leq C_\eta(n) \max_{W \neq \{\mathbf{0}\} \text{ lattice subspace of } \mathcal{L}^*} (\det(\mathcal{L}^* \cap W))^{-1/\dim(W)}. \tag{3}$$

Then $C_\eta(n) \leq \mathrm{poly}\log n$.

Note that by homogeneity, to prove Conjecture III.1 it suffices to show that $\eta(\mathcal{L}) \leq \mathrm{poly}\log n$ whenever all sublattices of $\mathcal{L}^*$ have determinant at least 1. We now show an equivalence between the smoothing parameter and a bound on the number of dual lattice points at distance r of the form $e^{(sr)^2}$. This formalizes the equivalence between the above conjecture and Conjecture I.1.

Lemma III.2. *For an n-dimensional lattice $\mathcal{L}$, the following inequality holds:*

$$\eta(\mathcal{L})/\sqrt{3} \leq \max_{r > 0} \frac{\sqrt{\log(|\mathcal{L}^* \cap rB_2^n|)/\pi}}{r} \leq \eta(\mathcal{L}) \ .$$

IV. The Main Conjecture implies the Kannan-Lovász Conjecture

We start by stating the ℓ_2 KL conjecture.

Conjecture IV.1 (The ℓ_2 Kannan-Lovász Conjecture)**.** *Let $C_{KL}(n) > 0$ be the smallest number such that for any $\mathcal{L} \subset \mathbb{R}^n$,*

$$\mu(\mathcal{L}) \leq C_{KL}(n) \max_{\substack{W \text{ lattice subspace of } \mathcal{L}^* \\ 1 \leq d = \dim(W) \leq n}} \sqrt{d} \det(\mathcal{L}^* \cap W)^{-1/d} \ . \tag{4}$$

Then $C_{KL}(n) \leq \mathrm{poly}\log n$.

We note that the reverse direction is easy to prove, i.e., that the $\max$ in Eq. (4) is $\lesssim \mu(\mathcal{L})$. We also note that the best known lower bound on $C_{KL}(n)$ is $O(\sqrt{\log n})$, obtained by the lattice generated by the basis $B = (\mathbf{e}_1, \mathbf{e}_2/\sqrt{2}, \ldots, \mathbf{e}_n/\sqrt{n})$. With the above formulation, we can state the main result of this paper as follows.

Theorem IV.2. *The ℓ_2 Kannan-Lovász conjecture holds with bound $C_{KL}(n) = O(\log n) C_\eta(n)$.*

In an omitted section we provide some optional background on the KL conjecture. The rest of this section is dedicated to the proof of Theorem IV.2.

A. Proof outline

We now give a high level overview of the proof. The first step of the proof, appearing in Section IV-B, is to formulate a convex relaxation μ_{sm} of the covering radius μ. This convex relaxation is quite natural, and can be described as measuring the "most efficient" way to smooth a lattice using an ellipsoidal Gaussian. We will explore it further in Section VI. Next, we use our main conjecture to arrive at a further convex relaxation, which we dub μ_{det}. This is done in Section IV-C. The final and arguably most interesting part of the proof is to round the dual formulation of μ_{det}. This is done in Section IV-D. The proof of the theorem is then just a combination of the theorems appearing in the following subsections

B. Smooth μ bound

The first step in the proof is to bound the covering radius by a convex relaxation we call μ_{sm}. The fact that the covering radius is at most $\sqrt{n}$ times the smoothing parameter is standard by now (it is already implicit in [10]), and the theorem can be seen as a slight extension of this standard fact to non-spherical Gaussians. One interesting aspect of the statement, though, is that the resulting minimization problem is convex.

Theorem IV.3 (Smooth μ bound)**.** *For an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$, let*

$$\mu_{\mathrm{sm}}(\mathcal{L})^2 = \min\Big\{\mathrm{tr}(A) : A \succeq 0, \sum_{\mathbf{y}\in\mathcal{L}^*\setminus\{\mathbf{0}\}} e^{-\pi \mathbf{y}^{\mathsf{T}} A \mathbf{y}} \le 1/2\Big\}. \tag{5}$$

Then $\mu(\mathcal{L}) \le 4\pi^{-1/2}\mu_{\mathrm{sm}}(\mathcal{L})$. Furthermore, the program defining $\mu_{\mathrm{sm}}(\mathcal{L})$ is convex.

C. Determinantal μ bound and its dual

Theorem IV.4 (Determinantal μ bound)**.** *For an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$, let*

$$\mu_{\det}(\mathcal{L})^2 = \min\Big\{\mathrm{tr}(A) : A \succeq 0, \det_W(A) \ge \frac{1}{\det(\mathcal{L}^* \cap W)^2}, \\ \forall\, W \neq \{\mathbf{0}\} \text{ lattice subspace of } \mathcal{L}^*\Big\}. \tag{6}$$

Then, the program defining $\mu_{\det}(\mathcal{L})$ is convex, and

$$\mu_{\det}(\mathcal{L})/2 \le \mu_{\mathrm{sm}}(\mathcal{L}) \le C_\eta(n)\mu_{\det}(\mathcal{L}) \ .$$

We now write down the dual of $\mu_{\det}$ and prove that strong duality holds (following a mostly-standard argument).

Theorem IV.5 (Dual of determinantal μ bound)**.** *For an n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$, let*

$$\begin{aligned} &\mu_{\det}^{\mathrm{dual}}(\mathcal{L})^2 = \textit{maximize } \sum_{i=1}^{m} \frac{d_i \det_{W_i}(X_i)^{1/d_i}}{\det(\mathcal{L}^* \cap W_i)^{2/d_i}} \\ &\textit{subject to} \\ &\sum_{i=1}^{m} X_i \preceq I \\ &W_i \neq \{\mathbf{0}\} \text{ lattice subspace of } \mathcal{L}^*, d_i = \dim(W_i), i \in [m] \\ &X_i \succeq 0,\ \mathrm{im}(X_i) = W_i, i \in [m] \\ &m \in \mathbb{N}. \end{aligned} \tag{7}$$

Then $\mu_{\det}(\mathcal{L}) = \mu_{\det}^{\mathrm{dual}}(\mathcal{L})$.

D. Subspace rounding

Here we prove Theorem IV.6, the most involved part of the reduction. The proof uses a certain "uncrossing" inequality, Lemma IV.7, proven in Section IV-E below.

Theorem IV.6 (Subspace Rounding)**.** *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice. Then*

$$\mu_{\det}^{\mathrm{dual}}(\mathcal{L})^2 \le \\ 24(\log_2 n + 1)^2 \max_{\substack{W \text{ lattice subspace of } \mathcal{L}^* \\ d=\dim(W)\in[n]}} \frac{d}{\det(\mathcal{L}^* \cap W)^{2/d}}.$$

E. The uncrossing inequality

The following is the uncrossing inequality.

Lemma IV.7. *For any lattice $\mathcal{L}$, lattice subspaces V, W, and positive semidefinite $X, Y \succeq 0$ with images V, W respectively,*

$$\frac{\det_V(X)}{\det(\mathcal{L}\cap V)^2} \cdot \frac{\det_W(Y)}{\det(\mathcal{L}\cap W)^2} \le \\ \frac{\det_{V\cap W}\big(\big(\frac{X+Y}{2}\big)^{\cap(V\cap W)}\big)}{\det(\mathcal{L}\cap(V\cap W))^2} \cdot \\ \frac{\det_{V+W}\big(X+Y-\big(\frac{X+Y}{2}\big)^{\cap(V\cap W)}\big)}{\det(\mathcal{L}\cap(V+W))^2} \ .$$

V. Weaker variants of the main conjecture

In an effort to make progress on the main conjecture, it is natural to consider weaker forms of it and hope that they would be easier to prove. In this section we describe four such forms, some quite natural in their own right, and describe their relationship. Those weaker forms are obtained by relaxing the quantity appearing in the right-hand side of Eq. (3) in the main conjecture, which we call here $\eta_{\det}$. The first two involve the quantities η_ρ and η_μ where instead of asking the sublattice $\mathcal{L}^* \cap W$ to have small determinant, we ask it to have large Gaussian mass (or equivalently, many lattice points in a ball) in the case of η_ρ, or its dual to have large covering radius in the case of η_μ. The remaining two quantities, η_ρ° and η_μ°, are obtained from the previous two by replacing the lattice subspace by an arbitrary ellipsoid. This avoids the discreteness inherent in the set of lattice subspaces, and might be more amenable to a proof.

Definition V.1. *For a lattice $\mathcal{L}$ we define the following quantities, where W always ranges over all lattice subspaces of $\mathcal{L}^*$ of positive dimension.*

$$\begin{aligned} \eta_{\det}(\mathcal{L}) &= \max_W (\det(\mathcal{L}^* \cap W))^{-1/\dim(W)} \\ \eta_\rho(\mathcal{L}) &= \max_{s>0, W} s \cdot (\log \rho_{1/s^2}(\mathcal{L}^* \cap W)/\dim W)^{1/2} \\ \eta_\mu(\mathcal{L}) &= \max_W \mu(\pi_W(\mathcal{L}))/\sqrt{\dim W} \\ \eta_\rho^\circ(\mathcal{L}) &= \sup_{X \succ 0} (\log \rho_X(\mathcal{L}^*)/\operatorname{tr} X)^{1/2} \\ \eta_\mu^\circ(\mathcal{L}) &= \sup_{R \textit{ non-singular}} \mu(R\mathcal{L})/(\mathrm{tr}(R^{\mathsf{T}}R))^{1/2} \\ &= \sup_{R \textit{ non-singular}, \|R\|_F \le 1} \mu(R\mathcal{L}) \end{aligned}$$

Remark. *It is easy to check from the definitions that all six "η-type" parameters are positively homogeneous, that is for $\tilde{\eta} \in \{\eta_{\mathrm{det}}, \eta_\rho, \eta_\mu, \eta_\rho^\circ, \eta_\mu^\circ, \eta\}$ as above, we have $\tilde{\eta}(\lambda\mathcal{L}) = \lambda\tilde{\eta}(\mathcal{L})$ for $\lambda > 0$. We will show a stronger property in Lemma V.6 below. Also, in the definitions of η_ρ° and η_μ° we can equivalently take the supremum over all (nonzero) matrices X, R, including singular ones. This holds due to standard limit arguments.*

Theorem V.2. *For any lattice $\mathcal{L}$ we have the inequalities*

$$\begin{array}{ccccccc} \frac{1}{C_\eta(n)}\eta(\mathcal{L}) \leq \eta_{\mathrm{det}}(\mathcal{L}) & \lesssim & \eta_\rho(\mathcal{L}) & \lesssim & \eta_\mu(\mathcal{L}) & & \\ & & \rotatebox{-90}{$\leq$} & & \rotatebox{-90}{$\leq$} & & \\ & & \eta_\rho^\circ(\mathcal{L}) & \lesssim & \eta_\mu^\circ(\mathcal{L}) & \lesssim & \eta(\mathcal{L}) \ . \end{array}$$

Lemma V.3. *For any lattice $\mathcal{L}$ and $s > 0$,*

$$\mu(\mathcal{L}) \geq \max_{s>0} s \cdot \sqrt{\frac{\log \rho(s\mathcal{L}^*)}{\pi}} \ . \tag{8}$$

One might wonder if the inequality in Lemma V.3 also holds in reverse, say up to polylogarithmic factors. It turns out that this is an easy consequence of the ℓ_2 KL conjecture, and we may therefore refer to it as *the weak-KL conjecture*. Intuitively, it shows that the covering radius can be characterized up to polylogarithmic factors by the norm distribution of points in the dual lattice. To make this more precise, we show in Lemma V.5 below that the maximum in (8) has a simple equivalent point counting formulation. In particular, we show that it is in essence the minimum $s > 0$ for which the number of points at any radius r is bounded by a function of the form e^{sr}. We note the interesting similarity with the point counting formulation of the smoothing parameter in Lemma III.2, which gives a bound of the form $e^{(sr)^2}$.

Weaker forms of the main conjecture: By replacing the quantity η_{det} appearing in the main conjecture with one of other four quantities appearing in Definition V.1, we get weaker forms of the main conjecture. Some of those are quite natural. For instance, by using η_μ, we obtain a conjecture that can be interpreted as saying that if $\mathcal{L}$ is not smooth, then there is a certificate for that in the form of a projection where the Gaussian "does not reach" the covering radius of the projected lattice. Also, the conjecture obtained from η_ρ° is quite appealing as both sides of the inequality only involve the Gaussian mass. Finally, by using $\eta_\mu^\circ(\mathcal{L})$ we obtain the weakest form of the main conjecture. It turns out that this weakest form is sufficient for the applications in Section VII and Section VIII-A, and we therefore define it explicitly.

Conjecture V.4. *(Weak conjecture) Let $C_\eta^{(\mu,\circ)}(n) > 0$ be the smallest number such that for any n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\eta(\mathcal{L}) \leq C_\eta^{(\mu,\circ)}(n)\eta_\mu^\circ(\mathcal{L}) \ . \tag{9}$$

Then $C_\eta^{(\mu,\circ)}(n) \leq \mathrm{poly}\log n$.

In Section V-A we will show that the weak conjecture combined with the ℓ_2-KL conjecture imply the main conjecture.

Point counting formulation: Here we formulate the claim made after Lemma V.3.

Lemma V.5. *Let $\mathcal{L} \subset \mathbb{R}^n$ be an n-dimensional lattice. Then*

$$\begin{aligned} 2\sqrt{\pi}e^{3/2} \max_{r>0} \frac{\log(|\mathcal{L} \cap rB_2^n|)}{2\pi r} &\geq \max_{s>0} s\sqrt{\frac{\log \rho(s\mathcal{L})}{\pi}} \\ &\geq \max_{r>0} \frac{\log(|\mathcal{L} \cap rB_2^n|)}{2\pi r}. \end{aligned}$$

Continuity of the η parameters: We end this section by stating basic continuity properties of the η-type parameters.

Lemma V.6. *Let $\tilde{\eta} \in \{\eta_{\mathrm{det}}, \eta_\rho, \eta_\mu, \eta_\rho^\circ, \eta_\mu^\circ, \eta\}$. For $\mathcal{L} \subset \mathbb{R}^n$ an n-dimensional lattice, for an invertible transformation $T \in \mathbb{R}^{n \times n}$,*

$$\|T^{-1}\|^{-1}\tilde{\eta}(\mathcal{L}) \leq \tilde{\eta}(T\mathcal{L}) \leq \|T\|\tilde{\eta}(\mathcal{L}) \ .$$

In particular, $\tilde{\eta}(\mathcal{L}) = \tilde{\eta}(T\mathcal{L})$ if T is an orthogonal transformation. Furthermore, the map $T \mapsto \tilde{\eta}(T\mathcal{L})$, with domain equal to the space of $n \times n$ invertible matrices, is continuous.

A. The Kannan-Lovász conjecture and the weak conjecture

The goal of this section is to show that the ℓ_2 Kannan-Lovász conjecture is equivalent to a polylogarithmic bound on the worst case ratio between η_μ° and η_{det}.

Definition V.7. *Let $C_{(\mu,\circ)}^{\mathrm{det}}(n)$ denote the smallest number such that for any n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\eta_\mu^\circ(\mathcal{L}) \leq C_{(\mu,\circ)}^{\mathrm{det}}(n) \cdot \eta_{\mathrm{det}}(\mathcal{L}) \ .$$

The main equivalence is given below.

Theorem V.8. *For $n \geq 1$, $C_{(\mu,\circ)}^{\mathrm{det}}(n) \leq C_{KL}(n) \lesssim \log n \cdot C_{(\mu,\circ)}^{\mathrm{det}}(n)$.*

The first inequality immediately implies that the weak conjecture together with the ℓ_2-KL conjecture imply the main conjecture. The second inequality is a sharpened version of our main result, Theorem IV.2, since $C_{(\mu,\circ)}^{\mathrm{det}}(n) \lesssim C_\eta(n)$ by Theorem V.2.

VI. Smooth μ tightness and the weak conjecture

Recall from Theorem IV.3 the convex program μ_{sm} that provides an upper bound on μ.

Definition VI.1. *(Smooth μ tightness) Let $C^{(s\mu)}(n) > 0$ be the smallest number such that for any n-dimensional lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\mu_{\mathrm{sm}}(\mathcal{L}) \leq C^{(s\mu)}(n)\mu(\mathcal{L}) \ .$$

Intuitively, $C^{(s\mu)}(n)$ being small means that it is always possible to smooth a lattice with a (not necessarily spherical) Gaussian whose expected norm is nearly as small as possible, namely, not much more than the covering radius. Obviously, one cannot smooth a lattice with a Gaussian of expected norm less than the covering radius – this is precisely the content of Theorem IV.3.

The main result of this section (omitted) is that $C^{(s\mu)}(n) \approx C_\eta^{(\mu,\circ)}(n)$. This implies that the weak conjecture (Conjecture V.4) is equivalent to the statement that $C^{(s\mu)}(n) \leq \operatorname{poly}\log n$. We mention in passing that the proof of Theorem IV.2 combined with the easy reverse direction of the KL conjecture already implies that for any lattice $\mathcal{L}$, $\mu_{\rm sm}(\mathcal{L}) \lesssim C_\eta(n)\mu(\mathcal{L})$, i.e., that $C^{(s\mu)}(n) \lesssim C_\eta(n)$.

At a high level, we show that $C^{(s\mu)}(n)$ is in essence the "dual" form of $C_\eta^{(\mu,\circ)}(n)$.

VII. Mixing Time of Brownian Motion on the Torus

Given a full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$, consider the Brownian motion on $\mathbb{R}^n/\mathcal{L}$ starting from the origin. The probability density function at time $t > 0$ is given by $f_t : \mathbb{R}^n/\mathcal{L} \to \mathbb{R}^+$,

$$f_t(\mathbf{x}) = t^{-n/2}\rho_t(\mathcal{L} + \mathbf{x}) \ .$$

As t goes to infinity, the Brownian motion converges to the uniform distribution. For $1 \leq p < \infty$, the L_p *mixing time* $\tau_p(\mathcal{L})$ is defined as

$$\inf\Big\{t > 0 : \Big(\det(\mathcal{L})^{-1}\int_{\mathbb{R}^n/\mathcal{L}} |\det(\mathcal{L})f_t(\mathbf{x}) - 1|^p d\mathbf{x}\Big)^{1/p} < \frac{1}{4}\Big\}$$

and extended to $p = \infty$ by

$$\tau_\infty(\mathcal{L}) = \inf\big\{t > 0 \ : \ \forall \mathbf{x}, \ |\det(\mathcal{L})f_t(\mathbf{x}) - 1| < 1/4\big\} \ .$$

We clearly have that $\tau_p(\mathcal{L}) \leq \tau_q(\mathcal{L})$ for any $1 \leq p \leq q \leq \infty$. Also, the constant $1/4$ is arbitrary (see, e.g., [11, Section 4.5]).

Using the Poisson summation formula and Parseval's identity, one can show that $\tau_\infty(\mathcal{L}) \leq 2\tau_2(\mathcal{L})$ (and we can even take $1/2$ instead of $1/4$ in the definition of τ_2). This property is not unique to Brownian motions on the torus – the L_∞ mixing time of any reversible Markov chain is always at most twice the L_2 mixing time (see, e.g., the appendix of [12]).

What about the L_1 mixing time? How much smaller can it be than the L_2 (or L_∞) mixing time? Analyzing the L_1 mixing time of Markov chains is generally quite hard. We note that there are examples of random walks on finite transitive graphs where the L_1 and L_2 mixing times differ greatly [13]. Still, one can hope that L_1 and L_2 mixing times are close when considering Brownian motion on manifolds such as the torus. This and related questions were considered by Saloff-Coste (see, e.g., [4], [5], [6]) who also asks it explicitly in a recent survey [7, Problem 11].

We next prove that under the weak conjecture, L_1 mixing time is approximately the same as the L_∞ mixing time. Intuitively, the proof proceeds as follows. If we are below the L_∞ mixing time, then by the weak conjecture, there is a certificate for that in the form of a Euclidean structure (or equivalently, a linear transformation) under which the Brownian motion does not reach the covering radius. But if this is the case, then surely it cannot even mix in the L_1 sense.

Theorem VII.1. *For any n-dimensional lattice $\mathcal{L}$, $\tau_\infty(\mathcal{L}) \lesssim C_\eta^{(\mu,\circ)}(n)^2\,\tau_1(\mathcal{L})$.*

VIII. Computational complexity and cryptography

In this section, we present the complexity implications of the weak conjecture (Conjecture V.4) for approximating the smoothing parameter (Section VIII-A), and of the ℓ_2 KL conjecture (Conjecture IV.1) for approximating the covering radius (Section VIII-B).

A. The Weak Conjecture and GapSPP

We consider the following two computational problems. The first, GapSPP, is the problem of approximating the smoothing parameter $\eta(\mathcal{L})$ of an input lattice $\mathcal{L}$. The second, Discrete Gaussian Sampling (DGS), is the task of generating samples distributed according to $D_{\mathcal{L}+\mathbf{t},s}$ for parameters $s \geq \eta(\mathcal{L})$. Here, the *discrete Gaussian distribution* $D_{\mathcal{L}+\mathbf{t},s}$ is the discrete distribution with support $\mathcal{L} + \mathbf{t}$ whose probability mass function is proportional to the restriction of the Gaussian density of standard deviation s to $\mathcal{L} + \mathbf{t}$.

The smoothing parameter and the discrete Gaussian distribution are closely related. For instance, one of the main important properties of the smoothing parameter is that for $s = \tilde{\Omega}(\eta(\mathcal{L}))$, the discrete Gaussian distribution $D_{\mathcal{L}+\mathbf{t},s}$ "behaves like" a continuous Gaussian of standard deviation s in terms of its global statistics such as moments. Also, both play a fundamental role in lattice-based cryptography, in particular in the best known worst-case to average-case reductions for lattice problems (e.g., [14], [15], [16], [17], [18]). Finally, they recently featured in the fastest known provable algorithms for lattice problems [19].

Given the tight relationship between the smoothing parameter and the discrete Gaussian distribution, a natural question is whether one can compute a good approximation of the smoothing parameter (that is, solve GapSPP) using only oracle access to a discrete Gaussian sampler. The best known reduction is from $\tilde{O}(\sqrt{n})$-GapSPP to DGS sampling, and is implicit in [14]. The main goal of this section is to show that conditioned on the weak conjecture, one obtains an exponential improvement in the approximation factor, namely a reduction from $\operatorname{poly}\log n$-GapSPP to DGS. The formal statement appears in Theorem VIII.4.

Connection to lattice-based cryptography: Together with prior work [14], [16], [18], this reduction directly implies a worst case to average case reduction from $\tilde{O}(\sqrt{n})$-GapSPP to the Shortest Integer Solution problem (SIS), one of the base hard problems in lattice-based cryptography (see Definition VIII.5). The best unconditional approximation factor is $O(n)$. We will describe this in more detail in Section VIII-A2.

For the Learning with Errors (LWE) problem (see [15]), the other and perhaps most versatile base problem in lattice-based cryptography, it was shown in [17] that $\tilde{O}(\sqrt{n}/\alpha)$-GapSPP reduces to LWE, where α is the LWE error parameter. Interestingly, they also show that the reduction of $\tilde{O}(n/\alpha)$-GapSVP to LWE in [15] can be recovered by running the GapSPP reduction and using the known relations between the smoothing parameter and the shortest vector in the dual lattice (i.e., one can factor the reduction through GapSPP).

Given the above results, it seems that GapSPP might be a good alternative to the standard worst-case problems such as GapSVP or SIVP for worst case to average case reductions. Indeed, the obtained approximation factor is an $\tilde{O}(\sqrt{n})$ factor better when reducing to SIS (conditionally) and LWE (unconditionally), though one may argue that it is perhaps somewhat dubious to compare approximation factors with respect to different lattice problems.

On a concluding note, we remark that other than the results presented here and those in [17], very little work has been done to understand the fine-grained complexity of GapSPP. We hope here to have helped motivate its further study.

Overview of the reduction: We first explain the known reduction from $\tilde{O}(\sqrt{n})$-GapSPP to DGS sampling implicit in [14]. Given as input a lattice $\mathcal{L}$ and $s > 0$, the reduction simply calls the DGS oracle with $\mathcal{L}$, s, and $\mathbf{t} = 0$, and then runs a certain statistical test to check if the output "looks like" a discrete Gaussian distribution. Specifically, the test checks that (1) all vectors are in $\mathcal{L}$ and are of length $O(s\sqrt{n})$ and (2) they span $\mathbb{R}^n$. It is obvious that this reduction can be implemented in polynomial time. Correctness follows by showing that (1) for any parameter $s = \Omega(\eta(\mathcal{L}))$, $\tilde{O}(n)$ DGS samples on an n-dimensional lattice $\mathcal{L}$ are likely to be of length $O(s\sqrt{n})$ and span $\mathbb{R}^n$, and that (2) for $s = \tilde{O}(\eta(\mathcal{L})/\sqrt{n})$, any set of lattice points of length $O(s\sqrt{n})$ will be contained in a proper subspace of $\mathcal{L}$. We note that in the second case, it is crucial that the test is guaranteed to fail regardless of the distribution of samples, since the oracle can behave arbitrarily for s below the smoothing parameter.

Our improved reduction is from $\operatorname{poly}\log n$-GapSPP to DGS sampling, and is conditioned on the weak conjecture. Let us assume that we need to distinguish between $\eta(\mathcal{L}) \leq 1$ and $\eta(\mathcal{L}) \geq \operatorname{poly}\log n$. We first pick a coset $\mathbf{t}$ of $\mathcal{L}$ uniformly at random. We then ask the oracle to produce $O(n)$ DGS samples at parameter $O(1)$ over $\mathcal{L} + \mathbf{t}$. We then compute the empirical second moment matrix C over these samples, and accept if the largest eigenvalue of C is $O(1)$ and reject otherwise. The fact that this test succeeds in a yes instance follows from standard concentration arguments for subgaussian random variables. However, for no instances, we require the weak conjecture to prove soundness (against any distribution, not just DGS). At a technical level, we will use the weak conjecture to deduce that with constant probability over $\mathbf{t}$, for some linear map R, $\|R\|_F \leq 1$, the set $R(\mathcal{L} + \mathbf{t})$ consists only vectors of length $\Omega(1)$. The mere existence of this matrix R will turn out to be enough to force the covariance matrix of *any distribution* on $\mathcal{L} + \mathbf{t}$ to have largest eigenvalue of size $\Omega(1)$. This concludes the reduction.

As can be seen from the above description, our reduction (as well as that in [14]) actually do not require true discrete Gaussian samples (i.e., samples from $D_{\mathcal{L}+\mathbf{t},s}$) – any subgaussian distribution on $\mathcal{L} + \mathbf{t}$ would be equally good. We make this formal in Definition VIII.3 below, where we define the computational problem of discrete subgaussian sampling (DSGS). This mild strengthening of the reduction turns out to be quite useful for the connection to SIS, since some of the known reductions to SIS only produce subgaussian samples but not true discrete Gaussian samples.

1) Reduction of GapSPP to Discrete Subgaussian Sampling: We begin with the necessary preliminaries on subgaussian random variables.

Definition VIII.1 (Subgaussian Random Variable)**.** *We say that a random variable $X \in \mathbb{R}$ is s-subgaussian or subgaussian with parameter s, for $s > 0$, if for all $t \geq 0$,*

$$\Pr[|X| \geq t] \leq 2e^{-(t/s)^2/2} \ .$$

We note that the canonical example of a 1*-subgaussian distribution is* $N(0,1)$ *itself. For a vector-valued random variable $X \in \mathbb{R}^n$, we say that X is s-subgaussian if all its one-dimensional marginals are, i.e., if $\forall \theta \in S^{n-1}$, the random variable $\langle X, \theta \rangle$ is s-subgaussian.*

We now define the main lattice problems of interest in this section.

Definition VIII.2 (Gap Smoothing Parameter Problem)**.** *For $\alpha = \alpha(n) \geq 1$, $\varepsilon = \varepsilon(n) \geq 0$, α-GapSPP$_\varepsilon$ (the Smoothing Parameter Problem) is defined as follows: given a basis B for a lattice $\mathcal{L} \subset \mathbb{R}^n$ and $s > 0$, decide whether $\eta_\varepsilon(\mathcal{L}) < s$ (YES instance) or $\eta_\varepsilon(\mathcal{L}) \geq \alpha s$ (NO instance).*

Definition VIII.3 (Discrete Subgaussian Sampling Problem)**.** *For σ a function that maps lattices to non-negative real numbers, and $m = m(n) \in \mathbb{N}$, m-DSGS$^\sigma$ (the Discrete Subgaussian Sampling Problem) is defined as follows: given a basis $\mathbf{B}$ for a lattice $\mathcal{L} \subset \mathbb{R}^n$, a shift $\mathbf{t} \in \mathbb{R}^n$, and a parameter $s \geq \sigma(\mathcal{L})$, output a sequence of m independent and identically distributed s-subgaussian random vectors supported on $\mathcal{L} + \mathbf{t}$. We note that the function σ need not be efficiently computable, and that the output of the sampler is only guaranteed when $s \geq \sigma(\mathcal{L})$.*

The main result of this section is the following.

Theorem VIII.4. *For any $\alpha = \alpha(n) \geq 1$, there exists a polynomial time reduction from $O(\alpha C_\eta^{(\mu,\circ)}(n))$-GapSPP$_{1/2}$ to $O(n)$-DSGS$^{\alpha\eta}$.*

2) The implications to lattice-based cryptography: Reductions from worst-case lattice problems to the average-case problem SIS are often stated as reductions from standard lattice problems such as SIVP or GapSVP. Here we observe that they are in fact implicitly reductions from DSGS. As such, they can be combined with the reduction from Theorem VIII.4 to obtain a reduction from the worst case problem GapSPP to SIS. For concreteness, we will follow here the reduction from [16], which is a simplification of the reduction in [14]. A similar conclusion should apply to more recent and refined reductions, such as the one in [18].

We first define the average-case SIS problem. A possible setting of parameters is $q = \mathrm{poly}(n)$, $m = \lceil n \log q \rceil$, and $\beta = \sqrt{m}$. We note that with this setting solutions are guaranteed to exist, and so the problem is not vacuous.

Definition VIII.5 (Small Integer Solution Problem)**.** *For integers q, m, and n, and a real number $\beta \geq 1$, the $\mathrm{SIS}_{q,m,\beta}$ problem asks to find with non-negligible probability, given a matrix $\mathbf{A}$ chosen uniformly from $\mathbb{Z}_q^{n\times m}$, a nonzero $\mathbf{e} \in \mathbb{Z}^m$ such that $\mathbf{A}\mathbf{e} = 0 \bmod q$ and $\|\mathbf{e}\|_2 \leq \beta$.*

Theorem VIII.6. *For any $m(n), \beta(n) = \mathrm{poly}(n)$, and prime $q(n) = n \cdot \beta(n) \cdot \omega(\sqrt{\log n})$, there is a poly-time reduction from DSGS^{σ} for $\sigma(\mathcal{L}) = \tilde{O}(\beta(n)\eta(\mathcal{L}))$ (with any desired polynomial number of samples) to solving $\mathrm{SIS}_{q,m,\beta}$ on the average with non-negligible probability.*

Combined with Theorem VIII.4, this yields a reduction from $\tilde{O}(\beta(n)C_\eta^{(\mu,\circ)}(n))$-$\mathrm{GapSPP}_{1/2}$ to $\mathrm{SIS}_{q,m,\beta}$. In particular, with the setting of parameters mentioned above Definition VIII.5, and assuming the weak conjecture, we obtain a reduction from $\tilde{O}(\sqrt{n})$-$\mathrm{GapSPP}_{1/2}$ to $\mathrm{SIS}_{q,m,\beta}$, an improvement upon the best known reduction by an $\tilde{O}(\sqrt{n})$ factor.

B. KL and the Covering Radius Problem

The ℓ_2 KL conjecture has an easy application to the computational complexity of lattice problems. Although this application is not directly related to the main topic of this paper, we record it here for future reference and for further motivation.

Consider the decision version of the γ-approximate covering radius problem, denoted GapCRP_γ, where $\gamma = \gamma(n) > 0$ is an approximation factor. Here, we are given a lattice $\mathcal{L}$ and a number $r > 0$ and the goal is to decide if the covering radius $\mu(\mathcal{L})$ of $\mathcal{L}$ is at most r (YES instances) or more than $\gamma \cdot r$ (NO instances). This problem was considered by Guruswami et al. [20] who showed that $\mathrm{GapCRP}_{\sqrt{n}} \in \mathsf{NP}$ and that $\mathrm{GapCRP}_2 \in \mathsf{AM}$. They also showed that $\mathrm{GapCRP}_{\sqrt{n}} \in \mathsf{coNP}$ and that $\mathrm{GapCRP}_{\sqrt{n/\log n}} \in \mathsf{coAM}$. Here we observe that the ℓ_2 KL conjecture implies an exponential improvement on the two latter results. Namely, we have that $\mathrm{GapCRP}_{C_{KL}(n)} \in \mathsf{coNP}$. Indeed, this is easy to prove. Construct a verifier that is given as proof a lattice subspace W of $\mathcal{L}^*$ and verifies that $\det(\mathcal{L}^* \cap W) \leq d^{d/2}$, where $d = \dim(W)$. If $\mu(\mathcal{L}) \geq C_{KL}(n)$ then such a subspace exists by definition. If, on the other hand, $\mu(\mathcal{L}) \leq 1$, then the easy reverse direction of Eq. (4) shows that no such subspace can exist.

IX. Limits for Strong Reverse Minkowski Inequalities

In this somewhat more speculative section, we discuss a possible stronger form of the main conjecture that, if true, would truly deserve the name "reverse Minkowski." We recall that the main conjecture bounds from above the number of lattice points at any radius $r > 0$ by a function of the form $e^{(\mathrm{poly}\log n)r^2}$ when all sublattice determinants are at least 1. It is a priori unclear why assuming such a uniform upper bound on all radii is the "natural" thing to ask for (hopefully, we have at least demonstrated its usefulness), and one may be tempted to ask: given a radius $r > 0$, what is the tightest bound on the number of lattice points at this radius in terms of sublattice determinants?

To arrive at a plausible candidate for such a bound, let us recall Minkowski's first theorem, which provides volumetric *lower* bounds on lattice point counts. The lower bound is stated in terms of the determinant of the full lattice, but obviously, having a sublattice of small determinant would also suffice, as we can invoke Minkowski's theorem inside the subspace spanned by that sublattice. More precisely, it follows from Minkowski's theorem that for any lattice $\mathcal{L}$,

$$|\mathcal{L} \cap rB_2^n| \geq M((r/2), \mathcal{L}) \ ,$$

where

$$M(r, \mathcal{L}) = \max_{\substack{W \text{ lattice subspace of } \mathcal{L} \\ 0 \leq d = \dim(W) \leq n}} \mathrm{vol}_d(rB_2^d)/\det(\mathcal{L} \cap W) \ . \tag{10}$$

By convention, the quotient is 1 for $d = 0$, and hence we note that $M(r, \mathcal{L}) \geq 1$ always.

It now becomes natural to ask whether the volumetric bound $M(r, \mathcal{L})$ is far from being tight. To this end, we introduce the following definition.

Definition IX.1. *Let $C_M(n)$ denote the smallest number such that for any $r > 0$ and any n-dimensional lattice $\mathcal{L}$,*

$$|rB_2^n \cap \mathcal{L}| \leq M(C_M(n)r, \mathcal{L}) \ .$$

We may speculate that $C_M(n) \leq \mathrm{poly}\log n$. We note that while this would be amazing if true, a counterexample would possibly be just as (or more) instructive and yield useful insights into the structure of lattice points. We also note that this would imply the main conjecture, Conjecture III.1, as shown in the following proposition.

Proposition IX.2. *For any n, $C_\eta(n) = O(C_M(n))$.*

The best upper bound we can prove is $C_M(n) = O(\sqrt{n})$, as shown in Theorem IX.4 below. We note that by Proposition IX.2, this recovers the best known bound $O(\sqrt{n})$ on $C_\eta(n)$.

General convex bodies: One can take the above discussion a step further and wonder why we should restrict ourselves to scalings of the Euclidean ball, since Minkowski's theorem holds more generally for any symmetric convex body. Indeed, slightly overloading notation, for any symmetric convex body $K \subseteq \mathbb{R}^n$, Minkowski's theorem implies that $|\mathcal{L} \cap K| \geq M((K/2), \mathcal{L})$, where $M(K, \mathcal{L})$ is defined as in (10) with rB_2^d replaced by $K \cap W$.

Definition IX.3. *Let $C'_M(n)$ denote the smallest number such that for any n-dimensional lattice $\mathcal{L}$, and any symmetric convex body $K \subseteq \mathbb{R}^n$,*

$$|\mathcal{L} \cap K| \leq M(C'_M(n)K, \mathcal{L}) \ .$$

Clearly $C'_M(n) \geq C_M(n)$. Theorem IX.4 below gives the best upper bound we can prove, $C'_M(n) \leq O(n)$. Given the foregoing discussion, one might be tempted to ask whether the upper bound on $C'_M(n)$ can be improved, say to $\mathrm{poly} \log n$. Unfortunately, we can show that this is false: in Theorem IX.5 we prove that $C'_M(n) \geq \Omega(n^{1/4-\varepsilon})$. The proof crucially uses a convex body K that is *very far* from a Euclidean ball, hence we may still hope that a better inequality is possible for the Euclidean ball, namely, that $C_M(n) \leq \mathrm{poly} \log n$.

A. Best known upper bounds

Theorem IX.4 (Weak Reverse Minkowski)**.** *For a symmetric convex body K and n-dimensional lattice $\mathcal{L}$ in $\mathbb{R}^n$, $|K \cap \mathcal{L}| \leq M(3nK, \mathcal{L})$. Furthermore, for any radius $r > 0$, $|rB_2^n \cap \mathcal{L}| \leq M(6\sqrt{n}r, \mathcal{L})$.*

The above inequality is in fact relatively simple to prove and so we find it somewhat surprising that it was not discovered earlier.

B. Lower bound for general convex bodies

Theorem IX.5. *For any $\varepsilon > 0$ there exists a constant $C > 0$ such that for any n large enough, there exists an n-dimensional symmetric convex body K and lattice $\mathcal{L}$ such that $M(Cn^{1/4-\varepsilon}K, \mathcal{L}) \leq |\mathcal{L} \cap K|$.*

References

[1] R. Kannan and L. Lovász, "Covering minima and lattice-point-free convex bodies," *Ann. of Math. (2)*, vol. 128, no. 3, pp. 577–602, 1988.

[2] B. Green, 2007, a guest post on Terence Tao's blog.

[3] S. Lovett and O. Regev, "A conditional counterexample to a polynomial Freiman-Ruzsa conjecture over the integers," 2016, in preparation.

[4] L. Saloff-Coste, "Precise estimates on the rate at which certain diffusions tend to equilibrium," *Math. Z.*, vol. 217, no. 4, pp. 641–677, 1994.

[5] ——, "On the convergence to equilibrium of Brownian motion on compact simple Lie groups," *J. Geom. Anal.*, vol. 14, no. 4, pp. 715–733, 2004.

[6] A. Bendikov and L. Saloff-Coste, "Central Gaussian convolution semigroups on compact groups: a survey," *Infin. Dimens. Anal. Quantum Probab. Relat. Top.*, vol. 6, no. 4, pp. 629–659, 2003.

[7] L. Saloff-Coste, "Analysis on compact Lie groups of large dimension and on connected compact groups," *Colloq. Math.*, vol. 118, no. 1, pp. 183–199, 2010.

[8] V. Bentkus and F. Götze, "On the lattice point problem for ellipsoids," *Acta Arith.*, vol. 80, no. 2, pp. 101–125, 1997.

[9] U. Shapira and B. Weiss, "Stable lattices and the diagonal group," 2016, to appear in J. of the European Math. Society.

[10] W. Banaszczyk, "New bounds in some transference theorems in the geometry of numbers," *Mathematische Annalen*, vol. 296, no. 4, pp. 625–635, 1993.

[11] D. A. Levin, Y. Peres, and E. L. Wilmer, *Markov chains and mixing times*. American Mathematical Society, Providence, RI, 2009, with a chapter by James G. Propp and David B. Wilson.

[12] R. Montenegro and P. Tetali, "Mathematical aspects of mixing times in Markov chains," *Found. Trends Theor. Comput. Sci.*, vol. 1, no. 3, pp. x+121, 2006.

[13] Y. Peres and D. Revelle, "Mixing times for random walks on finite lamplighter groups," *Electron. J. Probab.*, vol. 9, pp. no. 26, 825–845, 2004.

[14] D. Micciancio and O. Regev, "Worst-case to average-case reductions based on Gaussian measures," *SIAM J. Comput.*, vol. 37, no. 1, pp. 267–302 (electronic), 2007.

[15] O. Regev, "On lattices, learning with errors, random linear codes, and cryptography," *J. ACM*, vol. 56, no. 6, pp. Art. 34, 40, 2009, preliminary version in STOC 2005.

[16] C. Gentry, C. Peikert, and V. Vaikuntanathan, "Trapdoors for hard lattices and new cryptographic constructions [extended abstract]," in *STOC'08*. ACM, New York, 2008, pp. 197–206.

[17] K.-M. Chung, D. Dadush, F.-H. Liu, and C. Peikert, "On the lattice smoothing parameter problem," in *2013 IEEE Conference on Computational Complexity—CCC 2013*. IEEE Computer Soc., Los Alamitos, CA, 2013, pp. 230–241.

[18] D. Micciancio and C. Peikert, "Hardness of SIS and LWE with small parameters," in *CRYPTO 2013*, 2013, pp. 21–39.

[19] D. Aggarwal, D. Dadush, O. Regev, and N. Stephens-Davidowitz, "Solving the shortest vector problem in 2^n time via discrete Gaussian sampling," in *STOC*, 2015.

[20] V. Guruswami, D. Micciancio, and O. Regev, "The complexity of the covering radius problem," *Comput. Complexity*, vol. 14, no. 2, pp. 90–121, 2005.