

Exact SVP on Haar-Random Lattices at the Heuristic Sieving Exponents

Yang Zhou^{1,2,3} Zhedong Wang⁴ Mingsheng Wang²

¹State Key Laboratory of Information Security

²Institute of Information Engineering, Chinese Academy of Sciences

³School of Cyber Security, University of Chinese Academy of Sciences

⁴Shanghai Jiao Tong University

zhouyang2000@iie.ac.cn

wzdstill@sjtu.edu.cn

wangmingsheng@iie.ac.cn

Abstract

We give a Monte Carlo algorithm for exact Euclidean SVP on Haar-random lattices with arithmetic time $(3/2)^{n/2+o(n)}$ and space $(4/3)^{n/2+o(n)}$, up to polynomial factors in the logarithmic condition number of the input basis. The leading exponents match those of heuristic lattice sieving. For a set of lattices of Haar measure $1 - o(1)$, the algorithm succeeds with probability $1 - o(1)$ on every input basis.

Our algorithm regenerates each sieve list by running a Markov chain on the nonzero lattice points in a smaller ball. The parent list supplies signed increments. Each step reports all admissible proposals and assigns each the same probability, with the remaining probability assigned to staying put. The resulting chain is symmetric and has uniform stationary distribution. Trajectories using independent transition randomness produce endpoints jointly close to independent uniform samples once the chain mixes. Matrix Chernoff bounds transfer a Laplacian gap from the full angular graph to the graph defined by the parent list. The main proof task is to establish the required graph estimates on Haar-random lattices. We do so by adapting Rogers's moment method to centered trace moments and lattice-point counts. Spherical filtering reports the complete proposal sets and the final short differences within the claimed bounds.

Contents

1	Introduction	1
1.1	Our result	1
1.2	Sample costs in provable sieving and DGS	2
1.3	Proof overview	3
2	Preliminaries	5
2.1	Lattices and probability notation	6
2.2	Markov chains, mixing, and coupling	6
2.3	Matrix concentration	7
I	Reducing exact SVP to a spectral gap	8
3	A Markov chain for sampling from a smaller ball	8
3.1	Symmetry and uniform stationarity	8
3.2	Transferring the Laplacian gap to a random parent list	10
4	From uniform lists to a shortest vector	11
4.1	Joint sampling guarantees from independent trajectories	12
4.2	The algorithm: sampling and difference search	12
4.3	Representation counts and the final difference search	14
4.4	Proof of the main theorem	18
II	A uniform Laplacian gap for the full angular graph	19
5	Spectral contraction from centered trace moments	20
5.1	Rogers’s formula and centered lattice sums	20
5.2	Contraction of the continuous spherical-cap walk	22
5.3	Reducing lattice contraction to a trace moment	22
5.4	Bounding a fixed trace moment using integer relations	25
5.5	Summing the contraction error over all radii	28
6	Uniform degree bounds	28
6.1	Small and intermediate radii: conditional moments	29
6.2	Large radii: the covering radius	31
	Appendices	32
A	Spherical filtering for transitions and difference search	32
A.1	Reporting all admissible increments	33
A.2	Reporting short differences between two lists	33
A.3	Time and space from membership and candidate counts	34
A.4	Bounding filtering cost by moments of lattice-point counts	35
B	Coverage and streaming for product filters	36
B.1	Covering a near pair with a product filter	36
B.2	Enumerating filter buckets within the space bound	38
C	Uniform initialization and resource budgets	39
C.1	Exact uniform sampling in the initial ball	39
C.2	Choosing list sizes, capacities, and mixing times	39
C.3	Choosing filter parameters and bounding their cost	41
C.4	Converting expected work to time and space budgets	41
	References	43

1 Introduction

A full-rank lattice in $\mathbb{R}^n$ is the set of integer combinations of linearly independent vectors $b_1, \dots, b_n$:

$$L = \left\{ \sum_{i=1}^n z_i b_i : z_1, \dots, z_n \in \mathbb{Z} \right\}.$$

Given these vectors as an input basis, the exact Euclidean shortest vector problem (SVP) asks for a nonzero vector in L of minimum Euclidean length, denoted by $\lambda_1(L)$.

Algorithms for finding short lattice vectors have applications in integer programming, algorithmic number theory, and cryptanalysis [ADRS15, Section 1]. In the worst case, exact SVP is NP-hard under deterministic polynomial-time reductions [Wan26].

Ajtai, Kumar, and Sivakumar introduced lattice sieving to solve exact SVP in $2^{O(n)}$ time and space [AKS01]. Sieving maintains lists of lattice vectors and takes differences of nearby pairs to obtain shorter lattice vectors. Subsequent work improved the list sizes and the searches for nearby pairs. With spherical filtering, heuristic sieving attains time $(\sqrt{3/2})^{n+o(n)} = 2^{0.292481\dots n+o(n)}$ and space $(\sqrt{4/3})^{n+o(n)} = 2^{0.207518\dots n+o(n)}$ [NV08, BDGL16]. These analyses model successive lists as uniform samples from an annulus [NV08, Section 4.1] [BDGL16, Section 7], but selecting nearby pairs changes the distribution of their differences.

Provable sieve analyses obtain success guarantees through perturbations and worst-case packing bounds, with sample budgets above the heuristic list scale [NV08, MV10, PS09]. Discrete Gaussian sampling (DGS) assigns lattice points Gaussian weights, giving distributions that can be tracked rigorously. It yields exact SVP in $2^{n+o(n)}$ time [ADRS15]. Recent DGS-based constructions achieve $2^{0.5596n+o(n)}$ time [GFH26a, GFH26b] and $2^{n/2+o(n)}$ time and space [Hha26a, Hha26b].

For Haar-random determinant-one lattices, Siegel’s mean value theorem and Rogers’s moment formulas estimate lattice-point counts from volumes [Sie45, Rog55]. Pouly and Shen used these tools with DGS to obtain improved exact-SVP bounds on almost all lattices [PS26]. We attain the heuristic sieving bounds by using Markov chains to regenerate lists jointly close to independent uniform samples from successively smaller balls. The geometric estimates establish rapid mixing, while the construction controls the distributions of successive lists.

1.1 Our result

The input is an arbitrary basis B of the lattice. Our resource bounds retain a polynomial dependence on its logarithmic condition number

$$\ell_B = 1 + \log(\|B\|_F \|B^{-1}\|_F),$$

where $\|\cdot\|_F$ denotes the Frobenius norm.

Main result (informal). There is a Monte Carlo algorithm and a sequence of sets of determinant-one lattices of Haar measure $1 - o(1)$ with the following property. On every lattice in these sets, given any input basis, the algorithm returns a shortest nonzero vector with probability $1 - o(1)$. If $\ell_B = n^{O(1)}$, its arithmetic time and space have respective leading exponents in the dimension n

$$\frac{1}{2} \log_2 \frac{3}{2} = 0.292481\dots, \quad \frac{1}{2} \log_2 \frac{4}{3} = 0.207518\dots$$

Theorem 4.3 gives the formal statement, and Remark 4.4 bounds the input-dependent factors by $\text{poly}(n, \ell_B)$.

Reference	Input model	Analysis	Time	Space
BDGL (batched) [BDGL16]	Arbitrary	Heuristic	0.29248	0.20752
ADRS [ADRS15]	Arbitrary	Rigorous	1.00000	1.00000
Gao–Feng–Hu [GFH26b]	Arbitrary	Rigorous	0.55957	0.50000
Hhan [Hha26b]	Arbitrary	Rigorous	0.50000	0.50000
Pouly–Shen [PS26]	Haar-random	Rigorous	0.63269	0.50000
Doulgerakis–Laarhoven [DL26]	Haar-random	Rigorous	0.29248	0.20752
This work	Haar-random	Rigorous	0.29248	0.20752

The time and space columns give rounded leading exponents α in bounds of the form $2^{\alpha n + o(n)}$, suppressing input-dependent polynomial factors under each paper’s conventions. For this work, the displayed exponents assume $\ell_B = n^{O(1)}$. The BDGL row uses the batched Nguyen–Vidick variant described in [BDGL16, Section 7], with its heuristic assumptions on the sieve lists.

Concurrent work and k -tuple sieving. In independent and concurrent work, Doulgerakis and Laarhoven [DL26] obtain the same leading time and space exponents for exact SVP on Haar-random lattices. Their construction uses a binary tree of sieve lists and rejection sampling to maintain independent uniform samples in successive spherical shells. Our approach regenerates lists by mixing symmetric Markov chains with uniform stationary distributions. This formulation provides a natural route to k -tuple sieving: a transition uses a signed sum of $k - 1$ parent vectors as its increment, with a symmetric admissibility rule preserving the uniform stationary distribution. We have obtained preliminary results on this extension and are preparing the details for inclusion in the next update of this preprint.

1.2 Sample costs in provable sieving and DGS

Provable sieving. The provable analyses of AKS and its list and birthday variants use perturbations to ensure that two short outputs differ by a shortest vector [NV08, MV10, PS09]. Their sample bounds depend on both the number of possible outputs and the probability of useful perturbations. Worst-case packing bounds on output counts and on the number of vectors retained for reduction are exponentially larger than the volume estimates underlying heuristic sieving. Sharper counts still leave the perturbation cost: increasing the perturbation radius improves the overlap of two balls whose centers differ by a shortest vector, but enlarges the region containing useful outputs.

The Pujol–Stehlé analysis quantifies this tradeoff [PS09, Lemmas 3–5 and Section 3.2]. Let ξ and r be the perturbation and candidate radii in units of the shortest-vector length, and let p_ξ be the overlap probability. The analysis uses

$$r > 2\xi > 1, \quad p_\xi = \left(1 - \frac{1}{4\xi^2}\right)^{n/2+o(n)}.$$

Replacing only the candidate count by its Haar volume scale $r^{n+o(n)}$, while retaining these constraints, gives the birthday sample budget

$$N_{\text{PS}} = \frac{r^{n/2+o(n)}}{p_\xi} = \left(\frac{r}{1 - 1/(4\xi^2)}\right)^{n/2+o(n)} \geq \left(\frac{3\sqrt{3}}{2}\right)^{n/2-o(n)}.$$

This exceeds our target list size $(4/3)^{n/2+o(n)}$, before accounting for reduction or list searches.

Discrete Gaussian sampling. The DGS-based constructions of Gao, Feng, and Hu and of Hhan use the sampler of Aggarwal et al. [ADRS15] above the smoothing parameter. It combines samples along lattices $L_0 \supset L_1 \supset \dots$ satisfying $[L_i : L_{i+1}] = 2^a$ and $L_{i+2} \subseteq 2L_i$. Hence

$$2^{2a} = [L_i : L_{i+2}] \geq [L_i : 2L_i] = 2^n, \quad a \geq n/2.$$

Maintaining at least 2^a samples already costs $2^{n/2}$ time and space, before any final search [ADRS15, Definition 5.8 and Corollary 5.10]. Kim’s single-sample construction retains this batch cost [JK26]. Hhan’s coset-difference tree also requires coset tables of size $2^{n/2-o(n)}$ to keep its Gaussian signal above the estimation error [Hha26b, Section 1.1 and Remark 1.2].

1.3 Proof overview

We first construct a Markov chain that generates each list with an approximately uniform product distribution, assuming geometric estimates for the lattice. We then prove those estimates for Haar-random lattices using Rogers’s method. Let r_n be the radius of a volume-one ball. The algorithm starts with independent uniform lattice points in a large ball, decreases the radius by a factor $1 + 1/n$ at each stage, and stops at cr_n , where c must allow both rapid mixing and a successful final difference search.

A Markov chain with uniform stationary distribution. Freezing a list of reducers before generating independent outputs already appears in ListSieve-Birthday [PS09, Section 2]. We use a fixed parent list as a set of possible increments for a Markov chain. Let $D = (y_1, \dots, y_M)$ be a list of nonzero lattice points in the radius- $(1 + 1/n)R$ ball. The states are the nonzero lattice points in the radius- R ball. For each signed index (i, s) with $s \in \{-1, 1\}$, form $w = v + sy_i$. The proposal is admissible when w is a state and the angle between v and w is at most 60 degrees plus $O(1/n)$. Two boundary states of length R have a difference of length R at 60 degrees. The slack in the parent radius permits the $O(1/n)$ adjustment. The test concerns the current and next states, rather than the current state and a parent entry. A move may increase or decrease the length.

At each step, report the complete set of admissible signed indices, retaining distinct indices even when their vectors coincide. Let $d_D(v)$ be its size, and fix a common positive integer $K \geq \max_v d_D(v)$. Stay at v with probability $1 - d_D(v)/(2K)$; otherwise choose uniformly from the reported set. Each signed proposal then has probability $1/(2K)$. Its reverse uses the opposite sign and has the same probability, so the chain is symmetric with uniform stationary distribution. Writing L_D for the Laplacian of the graph defined by D and $g(P_D)$ for the spectral gap, we have

$$P_D = I - \frac{L_D}{2K}, \quad g(P_D) = \frac{\lambda_2(L_D)}{2K}.$$

Thus rapid mixing reduces to a Laplacian lower bound and a degree upper bound.

Regenerating the lists by rapid mixing. For a ball of volume V , form the full angular graph on its nonzero lattice points using this angular threshold. A cap of half-angle 60 degrees, up to $O(1/n)$, has spherical measure

$$p = (\sin 60^\circ)^{n+o(n)} = (3/4)^{n/2+o(n)}.$$

The geometric input is that point counts are close to V , all degrees are close to Vp , and the Laplacian gap is $\Omega(Vp)$.

For an independent uniform parent list of size M , the expected degree is $\Theta(Mp)$. The angular threshold ensures that every edge difference lies in the parent ball. Each sampled parent vector contributes a bounded positive semidefinite Laplacian. The expected sum is proportional to the Laplacian of the full angular graph. The matrix Chernoff bound transfers its gap to the sampled graph. For Mp at least a sufficiently large constant times the sum of the logarithms of the state-space size and the inverse failure probability, we obtain $K = \Theta(Mp)$ and $\lambda_2(L_D) = \Omega(Mp)$. The chain therefore has a constant spectral gap. Thus a list of size $p^{-1} = (4/3)^{n/2+o(n)}$ times the stated polynomial factors suffices (Lemma 3.5).

Run trajectories from a common starting state using independent transition randomness. Conditional on D and that state, their endpoints are jointly close to independent uniform samples. The estimate holds for every starting state, including one selected from D . Coupling successive stages adds the errors once per stage, so each list remains close to the uniform product distribution at its radius.

Finding a shortest vector among sample differences. At radius cr_n , split the output into two lists of size M and first analyze independent uniform samples. Let N be the number of nonzero lattice points in the ball. For a shortest vector v , let K_v count its representations $v = x - y$ by nonzero lattice points in the ball. These representations correspond to lattice points in the intersection of two radius- cr_n balls whose centers are separated by $\|v\| = (1 + o(1))r_n$. The lattice-point estimates give $N = (1 + o(1))c^n$ and $K_v \geq (c^2 - 1/4)^{n/2} / \text{poly}(n)$ (Lemma 4.6). For $M \geq (4/3)^{n/2}$, the number Z of sampled pairs with difference v therefore satisfies

$$\mathbb{E}Z = \frac{M^2 K_v}{N^2} \geq \frac{1}{\text{poly}(n)} \left(\frac{16(c^2 - 1/4)}{9c^4} \right)^{n/2}.$$

A second-moment argument gives success with high probability when this expectation grows exponentially. The joint total-variation bound transfers this guarantee to the algorithm.

Our spectral estimate also requires $Vp = (3c^2/4)^{n/2+o(n)}$ to grow exponentially at the terminal radius. Thus we choose c to satisfy

$$\frac{3}{4}c^2 > 1 \quad (\text{spectral estimate}), \quad \frac{9}{16}c^4 < c^2 - \frac{1}{4} \quad (\text{final difference search}).$$

They hold for c slightly above $\sqrt{4/3}$: at $c^2 = 4/3$, the second inequality is already strict, since $1 < 13/12$.

Implementing the transitions and difference search. Each transition requires the complete set of admissible signed proposals. Finding one useful list vector is insufficient to implement the stated transition probabilities. We use the BDGL product filters to report candidates, then test each candidate against the ball and angular conditions [BDGL16, Laa16]. Fresh repetitions make the probability of missing any admissible proposal small. The same filters report short differences between the final lists. Appendix A gives both reporting routines and their cost bounds. Gao, Feng, and Hu also use BDGL filtering with a rigorously controlled input distribution in their DGS-based algorithm [GFH26b].

To bound the filtering cost, we need control of every query along the trajectories. We first mix a common starting state to near uniformity, then compare with trajectories started at one common uniformly random state, independent of the parent list. By stationarity, each query in this reference execution is marginally uniform conditional on the parent list. The first two moments

of lattice-point counts then bound the expected numbers of filter memberships and candidate pairs. Coupling transfers the resulting cost bound to the actual execution.

With lists of size $(4/3)^{n/2+o(n)}$ and $(3/2)^{n/2+o(n)}$ implicit filter directions, these moment bounds give total time $(3/2)^{n/2+o(n)}$, up to the stated polynomial factors. Processing one bucket at a time keeps space at $(4/3)^{n/2+o(n)}$, with the same convention. The algorithm stops when a prescribed work budget is exceeded, and Markov’s inequality bounds the probability of this event. Combining this bound with the sampling and reporting errors proves the algorithmic guarantee under the stated geometric estimates.

Verifying the geometric estimates by Rogers’s method. To establish the geometric estimates, we expand centered lattice sums according to their integer relations and bound the resulting geometric integrals, following Rogers’s moment method [Rog55, Kim22, Kim24, SS19]. Existing limiting distributions for short-vector lengths and angles, including quantitative extensions, concern the first few or a growing number of short vectors [Sod11a, Sod11b, Kim20, Hol22]. Here the graph has exponentially many vertices and exponentially small cap neighborhoods, and the estimates must hold uniformly over its vertices and the prescribed radii. We therefore prove the required bounds directly from the moment formulas.

For the spectral bound, the continuous spherical-cap walk has a known contraction estimate [LMSY23]. Comparing its two-step averages with those of the full angular graph reduces the discrete estimate to a trace moment of the comparison error. In the centered Rogers expansion, terms with a free intermediate position cancel. Gram determinant estimates bound the remaining integer-relation terms. At a sufficiently large fixed moment order, these bounds give exponential decay for every fixed $c > \sqrt{4/3}$ (Lemma 5.9).

The same volume estimates bound centered moments of lattice-point counts in balls. Conditioning on a primitive vector gives bounds on lattice-point counts in caps centered at each lattice-point direction, and a union bound controls all degrees at small and intermediate radii. Covering-radius estimates handle larger radii. Together with contraction, these bounds verify the full angular graph estimates on a set of lattices of Haar probability $1 - o(1)$ independent of the input basis (Proposition II.1).

Organization. Section 2 records the standard mixing and concentration tools. Part I gives the sampling algorithm and final difference search assuming the full angular graph estimates. Part II proves those estimates through centered trace moments and degree concentration. The appendices give the filtering implementation, initialization, and explicit parameter choices and budgets.

2 Preliminaries

Computational model. We use an exact real arithmetic model in which arithmetic operations, comparisons, square roots, logarithms, exponentials, and integer rounding have unit cost. Drawing an independent uniform sample from $[0, 1)$ or a standard Gaussian sample also has unit cost. We charge the standard polynomial arithmetic costs for sampling from balls and spheres and for generating random orthogonal matrices. Space is the number of real or integer storage cells used. The algorithm requires no basis reduction. Following [PS26, Section 3.4], we state the dimension exponents separately from polynomial factors depending on the input basis. Theorem 4.3 and Remark 4.4 specify this dependence.

2.1 Lattices and probability notation

All geometric vectors are columns and norms are Euclidean. The input basis $B = (b_1, \dots, b_n)$ generates $L = B\mathbb{Z}^n$, with $\det L = |\det B|$. Exact SVP asks for a nonzero vector of length

$$\lambda_1(L) = \min_{x \in L \setminus \{0\}} \|x\|.$$

Write $\mathcal{B}_n(t, R) = \{x : \|x - t\| \leq R\}$, omitting a zero center, and let vol denote volume. We scale B by $(\det L)^{-1/n}$ to make $\det L = 1$ and restore the original scale on output. Let r_n satisfy $\text{vol } \mathcal{B}_n(r_n) = 1$.

A Haar-random lattice is sampled from the invariant probability measure on the space of determinant-one lattices, identified with $\text{SL}_n(\mathbb{R})/\text{SL}_n(\mathbb{Z})$. The input basis is arbitrary, and the algorithm uses independent randomness.

Write $U(A)$ for the uniform distribution on A , using counting measure for finite sets, volume for regions, and surface measure for spheres, each normalized to total mass one. Let $\mathcal{L}(X)$ denote the distribution of a random variable X . Subscripts and conditioning in $\mathbb{E}$ and $\mathbb{P}$ specify the source of randomness.

For probability measures α, β on the same space, their total variation distance is

$$\|\alpha - \beta\|_{\text{TV}} = \sup_A |\alpha(A) - \beta(A)|,$$

where A ranges over measurable events. On a finite set this is $\frac{1}{2} \sum_x |\alpha(x) - \beta(x)|$.

2.2 Markov chains, mixing, and coupling

A Markov chain on a finite state space S is specified by a transition matrix P with nonnegative entries and $\sum_{w \in S} P(v, w) = 1$ for every v . Conditional on the current state v , the next state is drawn according to $P(v, \cdot)$, independently of the earlier states. We represent distributions as row vectors: an initial distribution α becomes αP^t after t steps. In particular, $P^t(z, \cdot)$ is the distribution at time t when the chain starts at z .

A probability distribution π is *stationary* if $\pi P = \pi$. The chain is *reversible* with respect to π if it satisfies detailed balance,

$$\pi(v)P(v, w) = \pi(w)P(w, v) \quad (v, w \in S).$$

Summing over v shows that detailed balance implies stationarity. All finite chains used in our algorithm have symmetric transition matrices, so they are reversible with respect to $U(S)$. When the graph of positive transitions is connected, this stationary distribution is unique. The spectral gap quantifies convergence to it.

Definition 2.1 (Absolute spectral gap). For a symmetric transition matrix P on a finite state space S of size at least two, define

$$g(P) = 1 - \sup_{\substack{f: S \rightarrow \mathbb{R}, f \neq 0 \\ \sum_v f(v) = 0}} \frac{\|Pf\|}{\|f\|}.$$

Here f is a real column vector indexed by S .

The constant vector is an eigenvector of P with eigenvalue one. The absolute gap measures contraction on the orthogonal subspace of zero-sum vectors. In our chains, every state has holding probability $P(v, v) \geq 1/2$, which makes all eigenvalues nonnegative. Thus $g(P) = 1 - \lambda_2(P)$, where $\lambda_2(P)$ is the second-largest eigenvalue. The following bound translates the spectral gap into convergence in total variation.

Theorem 2.2 (Spectral mixing, Levin–Peres [LPW17, proof of Theorem 12.4, p. 164]). *Let P be a symmetric transition matrix on a state space S of size $N \geq 2$, with $P(v, v) \geq 1/2$ for every v and absolute spectral gap at least $g > 0$. For every state z and integer $t \geq 0$,*

$$\left\| P^t(z, \cdot) - U(S) \right\|_{\text{TV}} \leq \frac{1}{2} \sqrt{N} e^{-gt}. \quad (2.1)$$

For $0 < \varepsilon < 1$, $O(g^{-1}(\log N + \log(1/\varepsilon)))$ steps suffice for error at most ε , uniformly over the starting state. Since total variation bounds the difference in probability of every event, this estimate transfers success guarantees for uniform samples to the chain’s output.

To pass from a single endpoint to a list and propagate errors across stages, we use couplings. A coupling of distributions α, β is a pair of random variables with these marginal distributions.

Proposition 2.3 (Maximal coupling, Levin–Peres [LPW17, Proposition 4.7]). *For probability distributions α, β on the same finite set, there is a pair (X, Y) with these marginals and*

$$\mathbb{P}\{X \neq Y\} = \|\alpha - \beta\|_{\text{TV}}.$$

Every coupling has disagreement probability at least this distance.

For independent samples, the product bound adds the errors in the individual coordinates. If α_i, β_i are distributions on the same finite space for each i , then

$$\left\| \bigotimes_{i=1}^Q \alpha_i - \bigotimes_{i=1}^Q \beta_i \right\|_{\text{TV}} \leq \sum_{i=1}^Q \|\alpha_i - \beta_i\|_{\text{TV}}. \quad (2.2)$$

Take independent maximal couplings of the coordinate pairs and apply a union bound.

Errors inherited from earlier stages do not increase when the same randomized procedure is applied to both input distributions. Let α, β be distributions on a finite input space, and let T be the transition matrix of a randomized procedure from this space to a finite output space, which may include a failure symbol. Then

$$\|\alpha T - \beta T\|_{\text{TV}} \leq \|\alpha - \beta\|_{\text{TV}}. \quad (2.3)$$

Couple the two inputs maximally and use the same randomness in both executions. Whenever the inputs agree, the outputs agree as well.

2.3 Matrix concentration

We use matrix Chernoff for sampled Laplacians and its scalar case for degrees. Write $\lambda_{\min}$ and $\lambda_{\max}$ for the extreme eigenvalues.

Theorem 2.4 (Matrix Chernoff bound, Tropp [Tro12, Corollary 5.2 and Remark 5.3]). *Let $Y_1, \dots, Y_M$ be independent real symmetric $d \times d$ random matrices with $0 \preceq Y_i \preceq B_* I$ almost surely, where $B_* > 0$ and $\preceq$ is the positive-semidefinite order. Put*

$$\mu_- = \lambda_{\min} \left(\sum_i \mathbb{E} Y_i \right), \quad \mu_+ = \lambda_{\max} \left(\sum_i \mathbb{E} Y_i \right).$$

Then

$$\mathbb{P} \left\{ \lambda_{\min} \left(\sum_i Y_i \right) \leq \mu_- / 2 \right\} \leq d \exp \left(-\frac{\mu_-}{8B_*} \right).$$

If $\mu_+ > 0$, then, for every $u \geq \mu_+$,

$$\mathbb{P} \left\{ \lambda_{\max} \left(\sum_i Y_i \right) \geq u \right\} \leq d \left(\frac{e\mu_+}{u} \right)^{u/B_*}.$$

Part I

Reducing exact SVP to a spectral gap

3 A Markov chain for sampling from a smaller ball

We maintain a list whose joint distribution is close to that of independent uniform samples from the current ball. Each stage regenerates the list by independent trajectories in a smaller ball. Since the starting state is selected from the parent list, the chain must mix uniformly over its starting state. We reduce this requirement to the Laplacian gap of the full angular graph in Proposition II.1.

At a stage with target radius R , fix a parent list of M nonzero lattice points in the radius- $(1 + 1/n)R$ ball. Its signed entries supply the increments of the chain. The state stays within the target ball, while its norm may increase or decrease from one step to the next. Each stage generates a full new list, retaining repeated vectors with their multiplicities. A parent point in the target ball supplies a starting state. Section 4.2 specifies the initialization and radius schedule.

3.1 Symmetry and uniform stationarity

Fix a stage whose smaller ball has radius R , and set

$$\tau = 1 + \frac{1}{n}, \quad b = 1 - \frac{\tau^2}{2}. \tag{3.1}$$

Definition 3.1 (Graph defined by a parent list). The state space is $S = (L \cap \mathcal{B}_n(R)) \setminus \{0\}$, and the parent list $D = (y_1, \dots, y_M)$ consists of nonzero lattice points of norm at most τR . A move from v to $w = v \pm y_i$ is admissible if $w \in S$ and the angle between v and w is at most $\arccos b$. With $[M] = \{1, \dots, M\}$, the signed indices of admissible proposals are

$$\mathcal{I}_D(v) = \{(i, s) : i \in [M], s \in \{-1, 1\}, w = v + sy_i \in S, \langle v, w \rangle \geq b \|v\| \|w\|\}. \tag{3.2}$$

Distinct indices count separately even when they represent the same vector. Let $A_D(v, w)$ count the proposals from v to w , let $d_D(v) = |\mathcal{I}_D(v)|$, and let $L_D = \text{diag}(d_D) - A_D$ be the graph Laplacian.

The angle in (3.2) is between the states v and w , whose difference is the signed parent entry $sy_i = w - v$. For fixed v , the allowed directions of w form a spherical cap, whose continuous analogue is the spherical-cap walk. Every pair of states satisfying the angular condition has its difference in the parent ball:

$$\|v - w\|^2 \leq \|v\|^2 + \|w\|^2 - 2b\|v\|\|w\| \leq 2(1 - b)R^2 = \tau^2 R^2.$$

Indeed, the maximum of the middle expression over both lengths in $[0, R]$ occurs when both equal R , since $b < 1/2$.

Choosing an admissible proposal uniformly at every step would generally give stationary probabilities proportional to the degrees. To obtain uniform stationarity, we give each signed proposal the same probability using a common denominator at every state, and assign the remaining probability to staying put.

Definition 3.2 (Markov chain defined by a parent list). Suppose $S \neq \emptyset$, and let K be a positive integer with $K \geq \max_{v \in S} d_D(v)$. Define $P_D : S \times S \rightarrow [0, 1]$ by

$$P_D(v, w) = \begin{cases} A_D(v, w)/(2K), & w \neq v, \\ 1 - d_D(v)/(2K), & w = v. \end{cases} \quad (3.3)$$

The algorithm fixes K in advance and writes $\perp$ for failure. With the radius and capacity fixed, one step is as follows. Each call uses independent transition randomness.

```

1: procedure STEP( $D, v$ )
2:    $\mathcal{I} \leftarrow$  all admissible signed indices at  $v$ , as in (3.2).
3:   if  $|\mathcal{I}| > K$  then
4:     return  $\perp$ 
5:   end if
6:   With probability  $1 - |\mathcal{I}|/(2K)$ , return  $v$ .
7:   Choose  $(i, s)$  uniformly from  $\mathcal{I}$ .
8:   return  $v + sy_i$ 
9: end procedure

```

If $\mathcal{I}$ is empty, the step stays with probability one. Distinct parent indices count separately, as in Definition 3.1. Reporting the full set $\mathcal{I}_D(v)$ gives exactly (3.3). Stopping at the first admissible proposal would change this transition rule.

Lemma 3.3 (Uniform stationarity). *The matrix P_D is symmetric and stochastic, satisfies $P_D(v, v) \geq 1/2$, and is reversible with stationary distribution $\pi = U(S)$.*

Proof. Every proposal changes the state, so $\sum_{w \neq v} A_D(v, w) = d_D(v)$. Thus every row of P_D sums to one, and $K \geq d_D(v)$ gives nonnegative entries and $P_D(v, v) \geq 1/2$. Reversing the sign, $(i, s) \mapsto (i, -s)$, is a multiplicity-preserving bijection between proposals from v to w and those from w to v . Hence A_D , and therefore P_D , is symmetric. The uniform measure then satisfies detailed balance and is stationary. $\square$

This holds for every fixed list satisfying the capacity condition. If the graph is connected, its stationary distribution is unique, and $P_D(v, v) \geq 1/2$ ensures convergence from every state. The filtering implementation is coupled to the chain with complete proposal enumeration. Its first missed proposal is charged to coupling failure.

To apply the mixing bound of Theorem 2.2, it remains to lower-bound $g_D = g(P_D)$. Since P_D is symmetric and $P_D(v, v) \geq 1/2$, Definition 2.1 gives

$$P_D = I - \frac{L_D}{2K}, \quad g_D = \frac{\lambda_2(L_D)}{2K}, \quad (3.4)$$

where I is the identity and $\lambda_2(L_D)$ is the second-smallest Laplacian eigenvalue. Our target is therefore $\lambda_2(L_D) = \Omega(K)$.

3.2 Transferring the Laplacian gap to a random parent list

We compare the random parent list with the graph that uses every admissible lattice increment.

Definition 3.4 (Full angular graph). The full angular graph has vertex set S , with an edge between distinct v, w exactly when $\langle v, w \rangle \geq b \|v\| \|w\|$. Its Laplacian is L_0 . Let S^{n-1} be the unit sphere and e_1 the first coordinate vector. Set

$$V = \text{vol } \mathcal{B}_n(R), \quad p = \Pr_{u \sim U(S^{n-1})} \{\langle u, e_1 \rangle \geq b\}.$$

The spherical-cap probability predicts degree Vp . Since $b = 1/2 + O(1/n)$,

$$p = n^{-1/2} (3/4)^{n/2} e^{O(1)}. \quad (3.5)$$

Computable bounds on p appear in Appendix C.

Fix $1/2 < \rho < 1$ and put $\kappa = (1 - \rho)/2$. The spectral analysis proves the following geometric input, with Haar probability $1 - o(1)$ at the required radii: point counts are close to V , all degrees are close to Vp , and

$$\lambda_2(L_0) \geq \kappa Vp. \quad (3.6)$$

The matrix Chernoff bound (Theorem 2.4) transfers these estimates to the chain P_D .

Lemma 3.5 (A random parent list preserves the Laplacian gap). *Fix a lattice and radius satisfying the point-count, degree, and spectral estimates above. Sample M parent points independently and uniformly from $(L \cap \mathcal{B}_n(\tau R)) \setminus \{0\}$. If Mp exceeds a sufficiently large constant multiple of $\log(|S|/\varepsilon)$, a capacity $K = \Theta(Mp)$ can be fixed in advance so that, with probability at least $1 - \varepsilon$,*

$$\max_v d_D(v) \leq K, \quad \lambda_2(L_D) = \Omega(Mp), \quad g_D \geq g_0 > 0. \quad (3.7)$$

The bound g_0 is computable from the sampling parameters, independent of the list and starting state, and bounded below by a positive constant depending only on the fixed ρ . Appendix C derives these parameters.

Proof. Lower-bound the Laplacian gap. Let N_+ count the nonzero points in the parent ball. Each sample y_i , used with both signs, contributes a random Laplacian X_i . Every vertex meets at most two edges contributed by this sample. For every real vector z , the inequality $(z_v - z_w)^2 \leq 2z_v^2 + 2z_w^2$ gives

$$0 \leq z^t X_i z = \sum_{\{v,w\} \text{ contributed by } y_i} (z_v - z_w)^2 \leq 4 \sum_v z_v^2.$$

Here t denotes transpose. Thus $0 \preceq X_i \preceq 4I$, and

$$L_D = \sum_{i=1}^M X_i, \quad \mathbb{E} X_i = \frac{2}{N_+} L_0. \quad (3.8)$$

For the second identity, fix an edge $\{v, w\}$ of the full angular graph. The geometric bound above puts both $w - v$ and $v - w$ in the parent ball. They are distinct nonzero points. Either choice for y_i contributes the undirected edge once, so

$$(X_i)_{vw} = -\mathbf{1}\{y_i = w - v\} - \mathbf{1}\{y_i = v - w\}, \quad \mathbb{E}(X_i)_{vw} = -\frac{2}{N_+},$$

where $\mathbf{1}\{\cdot\}$ denotes the indicator. Nonedges contribute zero. The diagonal is the degree, hence

$$\mathbb{E}(X_i)_{vv} = -\sum_{w \neq v} \mathbb{E}(X_i)_{vw} = \frac{2}{N_+} d_0(v),$$

with $d_0(v)$ the degree of v in the full angular graph. Edges contributed by one sample need not be independent.

Every X_i maps the constant vector to zero. Let the columns of W be an orthonormal basis of its orthogonal complement, the subspace of coordinate sum zero. Apply Theorem 2.4 to the independent matrices $W^t X_i W$, of dimension $N - 1$, where $N = |S|$. Their norms are at most four, and the smallest eigenvalue of their expected sum is

$$\mu_{\min} = \frac{2M}{N_+} \lambda_2(L_0).$$

Taking $d = N - 1$ and $B_* = 4$ gives

$$\Pr\{\lambda_2(L_D) \leq \mu_{\min}/2\} \leq (N - 1)e^{-\mu_{\min}/32}.$$

Since $N_+ = \Theta(\tau^n V) = \Theta(V)$ and $\lambda_2(L_0) \geq \kappa V p$, we obtain $\mu_{\min} = \Omega(Mp)$ and constants $a, a' > 0$ depending only on κ such that

$$\Pr\{\lambda_2(L_D) < aMp\} \leq |S|e^{-a'Mp}.$$

Upper-bound the degrees. At a fixed vertex, the degree is a sum of M independent variables in $[0, 2]$ with total mean $O(Mp)$. Choose $K = \Theta(Mp)$ at least u times every such mean, for any fixed $u > e$. The upper-tail bound in Theorem 2.4, with $d = 1$ and $B_* = 2$, bounds the probability of degree at least K by $(e/u)^{K/2} = e^{-\Omega(Mp)}$. A zero mean gives degree zero almost surely. A union bound over vertices gives the required capacity bound. Combining the two estimates with (3.4) proves the claim. $\square$

Multiplicity is essential here: the edge weights in $L_D = \sum_i X_i$ record all parent indices. Storing equal vectors with their frequencies preserves both the sampling model and (3.8). The spectral gap of P_D is now reduced to the three geometric estimates for the full angular graph.

4 From uniform lists to a shortest vector

We first turn the spectral gap of P_D into a guarantee for an entire output list. We then give the algorithm and prove its success and resource bounds.

4.1 Joint sampling guarantees from independent trajectories

Lemma 4.1 (Joint sampling from a fixed parent list). *Fix a parent list D satisfying (3.7), a state $z \in S$, and an integer $Q \geq 1$. Run Q copies of P_D from z with independent transition randomness. Their endpoints after t steps satisfy*

$$\left\| \mathcal{L}(X_1, \dots, X_Q \mid D, z) - U(S)^{\otimes Q} \right\|_{\text{TV}} \leq \frac{Q}{2} \sqrt{|S|} e^{-g_0 t}. \quad (4.1)$$

The bound is uniform over D and z subject to these conditions.

Proof. Conditional on (D, z) , the endpoint distribution is $(P_D^t(z, \cdot))^{\otimes Q}$. By Theorem 2.2, each factor is within $\frac{1}{2} \sqrt{|S|} e^{-g_0 t}$ of $U(S)$ in total variation. Apply the product bound (2.2). $\square$

The target product distribution does not depend on D or z , so averaging over pairs satisfying the hypotheses preserves the bound. In particular, the starting state may depend on the entire parent list. Taking the single-trajectory error to be ε/Q gives joint error ε at a cost of $O(\log Q)$ additional steps.

The actual parent list is itself only approximately independent and uniform. Let η_{in} and η_{out} denote the total-variation errors of the input and output lists from their ideal product distributions, counting failure as an additional outcome. By (2.3), applying the same stage to the actual and ideal input lists preserves the bound η_{in} on their output distance. If $\varepsilon_{\text{stage}}$ is the sum of the error bounds for failure of the list conditions, failure to find a starting state, mixing, filtering, and budget exhaustion with ideal input, then

$$\eta_{\text{out}} \leq \eta_{\text{in}} + \varepsilon_{\text{stage}}.$$

The inherited error is charged once per stage. Appendix C assigns allowances whose sum, including initialization, is $o(1)$.

4.2 The algorithm: sampling and difference search

Set $R_j = c\tau^j r_n$, $S_j = (L \cap \mathcal{B}_n(R_j)) \setminus \{0\}$, and $\pi_j = U(S_j)$. The stopping constant c is fixed by the analysis of the final difference search below. To choose the initial radius, let $\rho_B = \frac{1}{2} \sum_i \|b_i\|$ bound the radius of the fundamental cell $B[-1/2, 1/2]^n$. Take the least integer $J \geq 1$ such that

$$R_J \geq n^2 \max(\rho_B, r_n). \quad (4.2)$$

To specify the trajectory lengths without counting the states, use the computable bounds

$$U_j = (2 \lceil R_j \|B^{-1}\|_F \rceil + 1)^n, \quad \Lambda = 1 + n + \log U_J + \log(J + 1). \quad (4.3)$$

Here $\|B^{-1}\|_F$ is the Frobenius norm. Every $x = Bz \in S_j$ has $\|z\|_\infty \leq R_j \|B^{-1}\|_F$, so $|S_j| \leq U_j$.

Set $\xi = 1/[n(J + 1)]$. Compute M, K, g_0 by the explicit rules (C.2)–(C.3) in Appendix C. These rules use only the input basis, the dimension, and the fixed constants, and give

$$M \geq (4/3)^{n/2}, \quad M = p^{-1} \text{poly}(n, J, \Lambda), \quad K = \text{poly}(n, J, \Lambda), \quad g_0 = \Omega(1).$$

On a lattice satisfying the geometric estimates, an independent uniform parent list fails (3.7) with probability at most ξ and contains no starting state in the next ball with probability at most ξ . These are the only properties of the parameter choices used below. Define

$$t_j(\delta) = \left\lceil g_0^{-1} \max\{0, \tfrac{1}{2} \log U_j + \log(1/(2\delta))\} \right\rceil, \quad 0 < \delta < 1. \quad (4.4)$$

By (2.1), this many steps give endpoint error at most δ . Let $Q_j = M$ for $j > 0$ and $Q_0 = 2M$. A shared burn-in of length $t_j(\xi)$ permits comparison of query costs with the reference execution defined in Remark 4.5. The subsequent Q_j trajectories use independent transition randomness for $s_j = t_j(\xi/Q_j)$ steps. Lemma 4.1 gives joint output error at most ξ uniformly over their common starting state. Appendix C verifies these bounds and supplies the resource budgets.

Algorithm 4.2 (Sampling and final difference search).

Input: Normalized basis B and the original scale.

Use the radius schedule and trajectory lengths above, and the parameters and budgets in Appendix C.

```

1:  $D \leftarrow ()$ 
2: while  $|D| < M$  do
3:   Draw  $u \sim U(\mathcal{B}_n(R_J + \rho_B))$  independently.
4:    $x \leftarrow B \lfloor B^{-1}u + \frac{1}{2}\mathbf{1} \rfloor$ 
5:   if  $0 < \|x\| \leq R_J$  then
6:     Append  $x$  to  $D$ .
7:   end if
8: end while
9: for  $j = J - 1, \dots, 0$  do
10:  Freeze  $D$  and use radius  $R_j$  for this stage.
11:   $Z \leftarrow$  the first element of  $D$  in  $S_j$ .
12:  if no such element exists then
13:    return  $\perp$ 
14:  end if
15:  for  $t = 1, \dots, t_j(\xi)$  do ▷ Burn-in by full scans
16:     $Z \leftarrow \text{STEP}(D, Z)$ 
17:  end for
18:   $(X_1, \dots, X_{Q_j}) \leftarrow (Z, \dots, Z)$ 
19:  for  $t = 1, \dots, s_j$  do ▷ One batch of filtered queries per step
20:    for  $a = 1, \dots, Q_j$  do
21:       $X_a \leftarrow \text{STEP}(D, X_a)$  with independent transition coins.
22:    end for
23:  end for
24:   $D \leftarrow (X_1, \dots, X_{Q_j})$ 
25: end for
26: Split  $D$  by index into two lists of  $M$  points.
27: Report cross-list pairs  $(x, y)$  with  $0 < \|x - y\| \leq R_0$ .
28: if no pair is reported then
29:   return  $\perp$ 
30: end if
31: return the shortest reported difference, restoring its scale.

```

Burn-in scans D . At each output step, spherical filtering reports all admissible proposals for the whole batch before any transitions are made. The same routine reports pairs for the final difference search. Each batch uses fresh filter randomness independent of its fixed inputs and transition coins. Appendix A implements these reporting operations. Any detected subroutine failure or exhausted resource budget returns $\perp$. A missed report or failure to find a shortest vector may instead produce a nonshortest output.

Theorem 4.3 (Exact shortest vectors in Haar-random lattices). *For all sufficiently large n , there is a measurable set of determinant-one lattices of Haar probability $1 - o(1)$ such that, for every lattice in the set and every input basis of that lattice, the algorithm returns an exact shortest nonzero*

vector with probability $1 - o(1)$ over its own randomness. Its time and space are bounded by

$$T \leq (3/2)^{n/2} 2^{o(n)} \text{poly}(n, J, \Lambda), \quad (4.5)$$

$$S \leq (4/3)^{n/2} 2^{o(n)} \text{poly}(n, J, \Lambda). \quad (4.6)$$

The algorithm returns failure if a budget is exhausted. This set of lattices is independent of the input basis, and the conditional success bound is uniform over this set and its bases. The bounds include every step of the algorithm. On arbitrary lattices, only the resource bounds are asserted.

Remark 4.4 (Dependence on the input basis). The parameters satisfy $J, \Lambda = O(n\ell_B)$, so the polynomial factors in the theorem are bounded by $\text{poly}(n, \ell_B)$. To see this, work with the normalized basis and put $A = \|B\|_F$ and $C = \|B^{-1}\|_F$. The determinant-one condition gives $A, C \geq \sqrt{n}$, while $\rho_B \leq \sqrt{n}A/2$ and $r_n = \Theta(\sqrt{n})$. Since $\log \tau = \Theta(1/n)$, the choice of J yields

$$J = O(n [\log n + \log(1 + \rho_B/r_n)]) = O(n\ell_B).$$

Minimality of J also gives $R_J C \leq \text{poly}(n)AC$, hence $\log U_J = O(n\ell_B)$ and $\Lambda = O(n\ell_B)$. In particular, $\ell_B = n^{O(1)}$ suffices for the stated dimension exponents. The success guarantee still holds for every basis, with the corresponding resource bounds.

Remark 4.5 (Reference execution and budgets). Endpoint mixing controls the output distribution. To bound the work along the trajectories, we compare their queries with those of a stationary reference execution.

Fix a lattice and a stage with nonempty state space S , and sample D from the ideal product distribution of the parent list. Let G be the event that D satisfies (3.7) and contains a starting state in S . For every D in G , suppose burn-in has total-variation error at most δ from $U(S)$, and let ε_{mix} be the joint endpoint error from Lemma 4.1.

For every D , define a reference execution with capacity $\max\{K, \max_v d_D(v)\}$, starting all output trajectories at one common $U(S)$ state independent of D . Use independent transition coins and complete proposal sets. Each query is then uniform conditional on D , though the queries need not be independent. Let W be the total work of the fresh filtering routines along these trajectories, without capacity or budget cutoffs. These reports do not determine the reference transitions. If $\varepsilon_{\text{filt}}$ bounds the probability of any missed proposal along the reference trajectories, an output-query budget $B_{\text{out}} > 0$ gives

$$\eta_{\text{out}} \leq \eta_{\text{in}} + \Pr(G^c) + \delta + \varepsilon_{\text{filt}} + \varepsilon_{\text{mix}} + \frac{\mathbb{E}W}{B_{\text{out}}}.$$

Here $\eta_{\text{in}}, \eta_{\text{out}}$ include failure as in Section 4.1. Indeed, couple the input lists, then, on G , couple the burn-in endpoint to the reference start. With shared coins and filters, the executions agree until an omission or cutoff. Markov's inequality controls W , and Lemma 4.1 controls the untruncated endpoints. Both $\Pr(G^c)$ and $\mathbb{E}W$ use the unconditioned ideal distribution of the parent list. In particular, no bound on $\mathbb{E}[W \mid G]$ is needed or inferred.

4.3 Representation counts and the final difference search

The geometric estimates require

$$\frac{3}{4}c^2 > 1. \quad (4.7)$$

To determine the constraint on c for the final difference search, we need two quantities: the ball count $N_0 = |S_0|$ and the number $K_v = |S_0 \cap (S_0 + v)|$ of representations of a specified short vector v as a difference. The first fixes the sampling probability of each point, and the second counts successful pairs.

Lemma 4.6 (Counts in balls and representations of short vectors). *Fix $c > 1$ with $c^2 - 1/4 > 1$. With Haar probability $1 - o(1)$, the shortest length lies in $[r_n n^{-1/n}, r_n n^{1/n}]$, and simultaneously for every primitive lattice vector v in this window,*

$$N_0 = (1 + o(1))c^n,$$

$$K_v \geq (c^2 - 1/4)^{n/2} / \text{poly}(n).$$

In particular, the representation bound applies to every shortest vector.

We prove this lemma using a conditional count in the intersection of two balls. The next two results give unconditional moments and the distribution conditional on a primitive vector.

Theorem 4.7 (Siegel–Rogers moment formulas [Sie45, Rog55]). *Let $n \geq 3$ and let L be a Haar-random determinant-one lattice in $\mathbb{R}^n$. For a nonnegative bounded compactly supported Borel function φ , put $\Sigma_\varphi = \sum_{x \in L \setminus \{0\}} \varphi(x)$ and $\zeta(n) = \sum_{a \geq 1} a^{-n}$. Then*

$$\mathbb{E}_L \Sigma_\varphi = \int_{\mathbb{R}^n} \varphi(x) dx, \quad \text{Var}_L(\Sigma_\varphi) = \frac{1}{\zeta(n)} \sum_{a \geq 1} \sum_{b \in \mathbb{Z} \setminus 0} \int_{\mathbb{R}^n} \varphi(ax) \varphi(bx) dx.$$

In particular, for a bounded Borel set A , if $N_A = \#((L \setminus \{0\}) \cap A)$ and $V_A = \text{vol } A$, then $\mathbb{E} N_A = V_A$ and $\text{Var } N_A \leq C V_A$ for an absolute constant C .

The displayed second-moment form is recorded in [PS26, Corollary 1]. A modern treatment of the real-lattice mean formulas is [Kim22].

The count-variance bound. Take $\varphi = \mathbf{1}_A$. Each integral counts continuous configurations in which two collinear integer multiples lie in A . Changing variables separately by $y = ax$ and $y = bx$ bounds it by $\max(a, |b|)^{-n} V_A$. There are $2(2h - 1)$ pairs with $\max(a, |b|) = h$, giving $\text{Var } N_A \leq 4V_A \sum_{h \geq 1} h^{1-n} \leq 4\zeta(2)V_A$. $\square$

This formula handles the ball count directly. The region defining K_v depends on a lattice vector, so we condition on a fixed primitive vector. A nonzero vector is *primitive* if division by an integer greater than one does not leave it in the lattice.

Theorem 4.8 (Conditioning on a primitive vector, Strömbergsson [Str11, §5.1, (5.1)–(5.3), and Proposition 5.1]). *Let $n \geq 2$, let L be a Haar-random determinant-one lattice, and let L_{prim} denote its primitive vectors. For every $v \neq 0$, there is a conditional probability measure on lattices containing v primitively, with expectation $\mathbb{E}_v^{\text{prim}}$, such that every nonnegative measurable H satisfies*

$$\mathbb{E}_L \sum_{v \in L_{\text{prim}}} H(L, v) = \frac{1}{\zeta(n)} \int_{\mathbb{R}^n} \mathbb{E}_v^{\text{prim}} H(L, v) dv.$$

Write $h = \|v\|$ and $m = n - 1$, rotate so that $v = h e_1$, and identify $e_1^\perp$ with $\mathbb{R}^m$. Under this conditional measure,

$$L = \{h(j + \omega(z))e_1 + h^{-1/m}z : z \in \Lambda', j \in \mathbb{Z}\}, \quad (4.8)$$

where Λ' is a Haar-random determinant-one lattice in $\mathbb{R}^m$. Conditional on Λ' , the map $\omega : \Lambda' \rightarrow \mathbb{T}$, with $\mathbb{T} = \mathbb{R}/\mathbb{Z}$, is a uniformly random homomorphism: its values on a lattice basis are independent uniform points of $\mathbb{T}$.

The variance case of the following lemma suffices for the representation count. Its higher moments will also be used for the uniform degree bounds in Section 6.

Lemma 4.9 (Moments of lattice sums with random shifts). *Let $m \geq 3$, let Λ' be a Haar-random determinant-one lattice in $\mathbb{R}^m$, and, conditional on Λ' , let $\omega : \Lambda' \rightarrow \mathbb{T}$ be a uniformly random homomorphism. Fix a ball $B \subset \mathbb{R}^m$ centered at the origin, of volume U , a constant $J_1 > 0$, and a measurable function $0 \leq F(z, t) \leq J_1 \mathbf{1}_B(z)$, all independently of (Λ', ω) . Define*

$$Z = \sum_{z \in \Lambda' \setminus \{0\}} F(z, \omega(z)), \quad \nu = \int_{\mathbb{R}^m} \int_{\mathbb{T}} F(z, t) dt dz,$$

where dt is Haar probability on $\mathbb{T}$. Then $\mathbb{E}Z = \nu$, and for every even integer $2 \leq k < m$,

$$\mathbb{E}(Z - \nu)^k \leq J_1^k \mathbb{E}(N_B - U)^k, \quad N_B = \#((\Lambda' \setminus \{0\}) \cap B). \quad (4.9)$$

In particular, $\text{Var } Z \leq C J_1^2 U$ for an absolute constant C .

Proof. Averaging each shift and applying the first lattice mean gives $\mathbb{E}Z = \nu$. For higher moments, use the integer-relation formula and its centering rule (Theorem 5.2 and Corollary 5.3), proved independently in Section 5.1, retaining the shifts. For linearly independent lattice vectors $y_1, \dots, y_r$, the vector of shifts $t = (\omega(y_1), \dots, \omega(y_r))^t$ is uniform on $\mathbb{T}^r$: their integer coordinate matrix induces a surjective homomorphism $\mathbb{T}^m \rightarrow \mathbb{T}^r$, which preserves Haar probability. The relation term indexed by $A \in \mathcal{A}_{r,k}$ is therefore

$$\int_{\mathbb{R}^{m \times r}} \mathbb{E}_{t \sim U(\mathbb{T}^r)} \prod_{j=1}^k F(Y A e_j, (A^t t)_j) dY,$$

where e_j is the j -th standard column vector.

If position j participates in no relation, an integer change of basis separates its Euclidean coordinate and shift from all other positions. Integrating them contributes exactly ν . This term cancels the corresponding choice of $-\nu$ in the expansion of $(Z - \nu)^k$, as in Corollary 5.3. Only nonnegative terms with no free position remain. Bounding each factor by $J_1 \mathbf{1}_B$ removes the shifts and leaves the same centered relation sum for N_B , proving (4.9). All rearrangements are justified by $Z^a \leq J_1^a N_B^a$ for $a \leq k < m$ and the moment integrability in Theorem 5.2. Taking $k = 2$ and applying the count-variance bound of Theorem 4.7 gives the last assertion. $\square$

Proof of Lemma 4.6. The ball at the stopping radius has volume c^n . Theorem 4.7 gives

$$\mathbb{E}N_0 = c^n, \quad \frac{\text{Var } N_0}{(\mathbb{E}N_0)^2} = O(c^{-n}) = o(1).$$

Chebyshev's inequality proves the ball estimate. The same mean and variance give the shortest-length window: the ball of radius $r_n n^{-1/n}$ has expected count $1/n$, so Markov's inequality bounds its probability of being nonempty by $1/n$. The ball of radius $r_n n^{1/n}$ has mean n and variance at most Cn , so its probability of being empty is at most C/n . Hence

$$\mathbb{P}\{\lambda_1(L) \notin [r_n n^{-1/n}, r_n n^{1/n}]\} \leq (1 + C)/n.$$

Count along lines parallel to the short vector. Fix a primitive v in this window. Write $h = \|v\|$, $m = n - 1$, rotate to $v = h e_1$, and use the conditional representation (4.8). Set

$$E_v = \mathcal{B}_n(R_0) \cap \mathcal{B}_n(v, R_0), \quad \mu_v = \text{vol } E_v.$$

On the line $h^{-1/m}z + \mathbb{R}v$, the number of points in E_v with shift t is $F_v(z, t) = \sum_{j \in \mathbb{Z}} \mathbf{1}_{E_v}(h(j + t)e_1 + h^{-1/m}z)$. Let $Z_v = \sum_{z \in \Lambda' \setminus 0} F_v(z, \omega(z))$ count the points off the axis $\mathbb{R}v$. Such a point and its translate by $-v$ are both nonzero, so $K_v \geq Z_v$. Each line contributes at most $J_1 = 1 + 2R_0/h = O(1)$ points. With $\rho_v = \sqrt{R_0^2 - h^2/4}$, the projected region lies in $B = h^{1/m}\mathcal{B}_m(\rho_v)$. For independent uniform $Z_0 \in B$ and $\Theta \in \mathbb{T}$,

$$0 \leq F_v \leq J_1 \mathbf{1}_B, \quad U = \text{vol } B = h \text{vol } \mathcal{B}_m(\rho_v), \quad U \mathbb{E} F_v(Z_0, \Theta) = \mu_v.$$

The last identity follows by averaging the shift t : the axial scaling has Jacobian h and the transverse scaling has Jacobian h^{-1} . Thus $\mathbb{E}_v^{\text{prim}} Z_v = \mu_v$.

Lower-bound the mean. Write $x = (h/2 + s)e_1 + z$, where $z \perp e_1$. The cylinder

$$|s| \leq h/n, \quad \|z\|^2 \leq \rho_v^2 - 2h^2/n$$

lies in E_v , since both squared lengths are at most $R_0^2 - 2h^2/n + h|s| + s^2 \leq R_0^2$. As $h/\rho_v = \Theta(1)$, its volume gives

$$\mu_v \geq \frac{2U}{n} \left(1 - \frac{2h^2}{n\rho_v^2}\right)^{(n-1)/2} \geq CU/n.$$

Using $h/r_n = \exp(O(\log n/n))$ and the formula for ball volume gives $U = (c^2 - 1/4)^{n/2} n^{O(1)}$.

Concentrate the count and remove the conditioning. Lemma 4.9 gives

$$\mathbb{E}_v^{\text{prim}}(Z_v - \mu_v)^2 \leq C J_1^2 U.$$

Thus the conditional probability that $K_v < \mu_v/2$ is at most $4C J_1^2 U / \mu_v^2 \leq C' n^2 / U = e^{-\Omega(n)}$, uniformly in the window. The expected number of primitive vectors in that window is at most n . Theorem 4.8 therefore makes the expected number violating this bound exponentially small. Markov's inequality gives the bound simultaneously for every primitive vector in the window. Combining this with the ball count and shortest-length estimate proves the lemma. $\square$

Lemma 4.10 (A shortest vector among differences of uniform samples). *Fix $c > 1$ such that $c^2 - 1/4 > 1$ and*

$$\frac{9}{16}c^4 < c^2 - 1/4. \tag{4.10}$$

There is a set of lattices of Haar probability $1 - o(1)$ such that, on every lattice in the set, two independent lists of $M \geq (4/3)^{n/2}$ independent uniform points in S_0 contain a shortest cross-list difference with probability $1 - o(1)$.

Proof. Fix a lattice satisfying Lemma 4.6 and a shortest vector v . Let Z count cross-list pairs whose difference is v . Each pair is uniform on the N_0^2 possible ordered pairs, of which K_v succeed. Therefore

$$\mathbb{E} Z = \frac{M^2 K_v}{N_0^2} \geq \frac{1}{\text{poly}(n)} \left(\frac{16(c^2 - 1/4)}{9c^4} \right)^{n/2}.$$

The mean grows exponentially by (4.10).

To bound the variance, only pairs sharing a sample need attention. Write $N = N_0$ and retain K_v for the number of representations. Among ordered pairs of candidate pairs, M^2 coincide and $2M^2(M - 1)$ share exactly one sample index. Their joint success probabilities are K_v/N^2 and K_v/N^3 , respectively: the equation $x - y = v$ determines either endpoint from the other. Candidate

pairs sharing no index are independent. Bounding each variance or covariance by the corresponding joint success probability gives

$$\text{Var } Z \leq M^2 \frac{K_v}{N^2} + 2M^2(M-1) \frac{K_v}{N^3} \leq \mathbb{E}Z \left(1 + \frac{2M}{N}\right).$$

Chebyshev's inequality and $\mathbb{E}Z = M^2 K_v / N^2$ therefore imply

$$\Pr\{Z = 0\} \leq \frac{1}{\mathbb{E}Z} + \frac{2N}{MK_v} = \frac{1}{\mathbb{E}Z} + \frac{2}{\sqrt{K_v \mathbb{E}Z}} \leq \frac{1}{\mathbb{E}Z} + \frac{2}{\sqrt{\mathbb{E}Z}} = o(1),$$

where the last step uses $K_v \geq 1$. □

The geometric condition (4.7) and the condition for the final difference search in (4.10) hold together in the nonempty interval

$$\sqrt{4/3} < c < \frac{\sqrt{8+2\sqrt{7}}}{3}. \quad (4.11)$$

The interval is nonempty because at $c^2 = 4/3$ the inequality for the final difference search is already strict: its two sides are $1 < 13/12$. Every c in the interval also satisfies $c^2 - 1/4 > 1$. Fix any such constant c for the algorithm. Its choice does not affect the leading time and space exponents.

The actual $2M$ endpoints are jointly close to the independent uniform samples just analyzed, and splitting them by index adds no distribution error. On a lattice satisfying the preceding estimates, $\lambda_1(L) = (1 + o(1))r_n < R_0$, so the explicit threshold for the final difference search includes every shortest vector. Once the reported differences include a shortest vector, exact comparison returns a shortest nonzero lattice vector.

4.4 Proof of the main theorem

Proof of Theorem 4.3 assuming Proposition II.1. Intersect the event in Proposition II.1 with the events for short-vector representations in Lemma 4.6 and filtering costs in Appendix A.4. Each holds with Haar probability $1 - o(1)$ independently of the input basis, so their intersection does as well.

Fix such a lattice and any input basis. Initialization gives an exactly uniform parent list. At each stage, Lemma 3.5 gives a constant spectral gap, and Lemma 4.1 regenerates a list close to the uniform product distribution. The error bounds for the list conditions, finding a starting state, mixing, filtering, and budget exhaustion sum to $o(1)$. The stagewise coupling adds each error once. Lemma 4.10 adds another $o(1)$ for failure of the final difference search.

Each trajectory takes $\text{poly}(n, J, \Lambda)$ steps. The shared burn-in scans the parent list at cost M times a polynomial. For the remaining steps, let $F = (3/2)^{n/2} 2^{o(n)}$ be the number of filter directions. The first- and second-moment estimates in Appendix A, equation (A.7), bound the expected batch cost of the stationary chain by $2^{o(n)} \text{poly}(n, J, \Lambda)$ times

$$M + F[1 + (3c^2/4)^{-n/2}].$$

Indeed, substitute $M = (4/3)^{n/2} \text{poly}(n, J, \Lambda)$, $Q \leq 2M$, and $V \geq c^n$: direct comparisons cost M times a polynomial, filter memberships and linearly independent candidate pairs cost F times a polynomial, and the collinear contribution has the remaining factor $(3c^2/4)^{-n/2}$. The same bound

applies to the final difference search with two radius- R_0 lists and $Q = M$. In particular, the term due to collinear pairs satisfies

$$\frac{FMQ(3/4)^{n/2}}{V} \leq F \text{poly}(n, J, \Lambda) (3c^2/4)^{-n/2}.$$

The condition $3c^2/4 > 1$ makes this factor decay. Markov's inequality, with the polynomial error allowances of Appendix C, turns expected costs into budgets. Summing over stages proves (4.5).

Streaming requires $2^{o(n)} \text{poly}(n, J, \Lambda)(M+Q+QK)$ storage. Here $Q \leq 2M$ and K is polynomial, so this is $(4/3)^{n/2} 2^{o(n)} \text{poly}(n, J, \Lambda)$, proving (4.6). The final difference search stores only its shortest current difference. All failure exits respect the budgets, while the success claim holds on the stated set of lattices. $\square$

It remains to prove Proposition II.1, which supplies the point counts, degree bounds, and Laplacian gap at all prescribed radii.

Part II

A uniform Laplacian gap for the full angular graph

The Markov chain construction reduces list generation to the geometric estimates in Proposition II.1. We prove these estimates using centered trace moments and lattice-point counts. The contraction and degree bounds must hold uniformly over the vertices and radii required by the algorithm.

Fix a lattice L and radius R . Write $S_R = (L \cap \mathcal{B}_n(R)) \setminus \{0\}$, $N_R = |S_R|$, $V = (R/r_n)^n$, and $\nu = Vp$. Distinct vertices x, y are adjacent when $\langle x, y \rangle \geq b \|x\| \|y\|$. Their adjacency matrix, degrees, and Laplacian are A_0, d_0 , and $L_0 = \text{diag}(d_0) - A_0$. Recall that $\tau = 1 + 1/n$, $b = 1 - \tau^2/2$, and $p = n^{-1/2}(3/4)^{n/2} e^{O(1)}$. Subscripts L on expectations and probabilities refer to Haar-random lattices.

Recall that $1/2 < \rho < 1$ and $\kappa = (1 - \rho)/2$.

Proposition II.1 (Geometry of the full angular graph). *Fix $c > \sqrt{4/3}$. With Haar probability $1 - o(1)$, simultaneously at all radii $R_j = c\tau^j r_n$, $j \geq 0$,*

$$\lambda_2(L_0) \geq \kappa\nu. \tag{II.1}$$

The same lattices satisfy $N_R = (1 + o(1))V$ and $d_0(x) = (1 + o(1))\nu$ for every vertex. All estimates are uniform over this infinite grid.

The matrix $M_0 = (A_0 + I)/\nu$ has row sums $(d_0(x) + 1)/\nu$. The next lemma reduces the Laplacian gap to degree concentration and contraction of M_0 on zero-sum vectors.

Lemma II.2 (From contraction and degrees to a Laplacian gap). *Suppose $N_R \geq 2$ and*

$$\max_{x \in S_R} \left| \frac{d_0(x) + 1}{\nu} - 1 \right| \leq \epsilon_n, \quad \|M_0 f\| \leq \rho \|f\| \quad \text{whenever} \quad \sum_x f(x) = 0.$$

Then $\lambda_2(L_0) \geq \nu(1 - \epsilon_n - \rho)$.

Proof. The identity

$$L_0 = \nu(I - M_0) + \text{diag}(d_0(x) + 1 - \nu)$$

and Cauchy–Schwarz give, for every nonzero zero-sum f ,

$$\frac{f^t L_0 f}{\nu \|f\|^2} \geq 1 - \epsilon_n - \frac{f^t M_0 f}{\|f\|^2} \geq 1 - \epsilon_n - \frac{\|M_0 f\|}{\|f\|} \geq 1 - \epsilon_n - \rho.$$

Minimize this Rayleigh quotient on the orthogonal complement of the constant vector. $\square$

The limiting continuous contraction is $1/2$, so contraction by ρ leaves a positive gap once the degree error tends to zero. We prove these two estimates on the radius grid.

5 Spectral contraction from centered trace moments

We first record the identity for centered lattice sums used throughout the proof. We then compare the continuous spherical-cap walk with the full angular graph.

5.1 Rogers’s formula and centered lattice sums

In the trace expansion below, intermediate points contribute centered lattice sums, while the end-points remain lattice sums. We therefore need Rogers’s formula with only a specified subset of coordinates centered. Integer coordinates express both the resulting cancellations and the later determinant bounds through the same relation matrices.

Definition 5.1 (Relation representatives). For $1 \leq r \leq k < n$, let $\mathcal{A}_{r,k}$ contain one representative of each class of matrices $A \in \mathbb{Z}^{r \times k}$ with no zero column and $A\mathbb{Z}^k = \mathbb{Z}^r$, modulo left multiplication by $\text{GL}_r(\mathbb{Z})$. A position j *participates in a relation* if some $\xi \in \mathbb{Z}^k$ satisfies $A\xi = 0$ and $\xi_j \neq 0$. Write $\mathcal{F}(A)$ for the set of remaining positions, which we call *free*.

Theorem 5.2 (Rogers’s formula in integer coordinates [SS22, Theorem 1.5, (1.12)]). *Let $1 \leq k < n$ and let $F : (\mathbb{R}^n)^k \rightarrow [0, \infty)$ be bounded, Borel measurable, and compactly supported. Then*

$$\mathbb{E}_L \sum_{x_1, \dots, x_k \in L \setminus \{0\}} F(x_1, \dots, x_k) = \sum_{r=1}^k \sum_{A \in \mathcal{A}_{r,k}} \int_{\mathbb{R}^{n \times r}} F(YA) dY.$$

Both sides are finite, and the identity also holds for bounded compactly supported real-valued F .

The source writes $F(YB^t)$, where the columns of B form a basis of a saturated subgroup of $\mathbb{Z}^k$. Taking $A = B^t$, Smith normal form identifies saturation with $A\mathbb{Z}^k = \mathbb{Z}^r$. Setting F to zero whenever an argument vanishes removes the origin and zero-column terms. For the equivalence with the classical formula [Rog55] and absolute integrability, see [SS22, Section 2 and Remark 2.2]. The integral is independent of the representative: replacing A by GA , $G \in \text{GL}_r(\mathbb{Z})$, amounts to the substitution $Y' = YG$, whose Jacobian is one.

Write S_i for summation over $x_i \in L \setminus \{0\}$ and $\int_i$ for integration over $x_i \in \mathbb{R}^n$, leaving the other coordinate labels unchanged. Centering selected coordinates gives the following form of the usual Rogers cancellation [SS19, proof of Theorem 2.3].

Corollary 5.3 (Partially centered relation sum). *For F as above and nonempty $B \subseteq \{1, \dots, k\}$,*

$$\mathbb{E}_L \left[\prod_{i \notin B} S_i \prod_{j \in B} (S_j - \mathbf{l}_j) \right] F = \sum_{r=1}^{k-1} \sum_{\substack{A \in \mathcal{A}_{r,k} \\ B \cap \mathcal{F}(A) = \emptyset}} \int F(YA) dY.$$

All expansions are absolutely integrable.

Proof. Choose a compact ball C containing every coordinate projection of $\text{supp } F$, and put $N_C = \#((L \setminus \{0\}) \cap C)$. Theorem 5.2 applied to $\mathbf{1}_C^{\otimes k}$ gives $\mathbb{E}_L N_C^k < \infty$. The sum of the absolute values of the expanded terms is bounded by $\|F\|_\infty \mathbb{E}_L (N_C + \text{vol } C)^k$, which justifies the rearrangements below.

Fix $J \subseteq \{1, \dots, k\}$ and integrate its positions first. Apply Theorem 5.2 to the remaining positions and insert an identity block for the integrated ones. In the column order (J^c, J) , this sends A' to $\text{diag}(A', I_{|J|})$, preserving the product Lebesgue measure. The insertion is a bijection onto relation classes whose positions in J are free. Indeed, for any such A , if H is generated by its columns outside J , freeness and $A\mathbb{Z}^k = \mathbb{Z}^r$ give

$$\mathbb{Z}^r = H \oplus \bigoplus_{j \in J} \mathbb{Z}a_j.$$

A basis of H followed by these a_j recovers the block form, uniquely up to the unimodular basis change defining the class of A' . When J is all positions, the block is simply I_k . Consequently,

$$\mathbb{E}_L \left[\prod_{i \notin J} S_i \prod_{j \in J} \mathbf{l}_j \right] F = \sum_{r=1}^k \sum_{\substack{A \in \mathcal{A}_{r,k} \\ J \subseteq \mathcal{F}(A)}} \int F(YA) dY.$$

Expanding the centered coordinates, the coefficient of each relation class is therefore

$$\sum_{J \subseteq B \cap \mathcal{F}(A)} (-1)^{|J|} = \mathbf{1}\{B \cap \mathcal{F}(A) = \emptyset\}.$$

Full-rank classes have every position free and thus cancel because $B \neq \emptyset$. □

Corollary 5.4 (Positivity of centered counts). *For $1 \leq k < n$ and bounded, nonnegative, compactly supported Borel $F : (\mathbb{R}^n)^k \rightarrow \mathbb{R}$,*

$$\mathbb{E}_L \left[\prod_{i=1}^k (S_i - \mathbf{l}_i) \right] F \geq 0.$$

All expanded terms are absolutely integrable. Moreover, $\mathbb{E}_L N_C^k < \infty$ for every bounded Borel set C , where $N_C = \#((L \setminus \{0\}) \cap C)$.

Proof. With $B = \{1, \dots, k\}$, every relation integral remaining in Corollary 5.3 is nonnegative. The moment assertion follows from Theorem 5.2 applied to $\mathbf{1}_C^{\otimes k}$. □

5.2 Contraction of the continuous spherical-cap walk

Definition 5.5 (Spherical-cap transition). Let σ be uniform probability measure on S^{n-1} and put

$$a(u, v) = \mathbf{1}\{\langle u, v \rangle \geq b\}, \quad T(u, v) = a(u, v)/p.$$

Here T is the transition density with respect to σ : $(Tg)(u) = \int T(u, v)g(v) d\sigma(v)$ is the expected value of g after one step from u . For $x \neq 0$, write $\hat{x} = x/\|x\|$. For real square-integrable g , write $\|g\|_{L^2(\sigma)}^2 = \mathbb{E}_{U \sim \sigma} g(U)^2$.

Theorem 5.6 (Mixing of the spherical-cap walk [LMSY23, Corollary 4.7]). *For the spherical-cap walk on S^{n-1} with threshold $b \in (0, 1)$ and cap measure p , when $nb^2 \rightarrow \infty$ there is $\beta_n = (1 + o_{nb^2}(1))b$ such that, for every initial probability distribution α and every integer $j \geq 1$,*

$$\|\alpha T^j - \sigma\|_{\text{TV}} \leq \sqrt{\frac{1}{2} \log \frac{1}{p}} \beta_n^{j-1}. \quad (5.1)$$

The error in β_n is independent of α and j .

Here αT^j is the distribution after j steps of the spherical-cap walk from α . The source's parameters d, τ, ρ, P_p are denoted here by n, b, σ, T . To compare spectra, we convert this total variation rate into an L^2 contraction rate.

Lemma 5.7 (Continuous contraction). *For every mean-zero $h \in L^2(\sigma)$,*

$$\|Th\|_{L^2(\sigma)} \leq \beta_n \|h\|_{L^2(\sigma)}, \quad \beta_n = \frac{1}{2} + o(1). \quad (5.2)$$

Proof. Let $\Pi h = \int h d\sigma$ be the projection onto constants, and set $C_n = \sqrt{\frac{1}{2} \log(1/p)}$. Applying Theorem 5.6 to a point mass at u gives

$$\sup_u \int |T^j(u, v) - 1| d\sigma(v) \leq 2C_n \beta_n^{j-1}.$$

The kernel is symmetric, so the same bound holds for column integrals. The Schur bound therefore yields

$$\|T^j - \Pi\|_{L^2 \rightarrow L^2} \leq 2C_n \beta_n^{j-1}.$$

Since T is self-adjoint and $T\Pi = \Pi T = \Pi$, we have $\|T - \Pi\|_{L^2 \rightarrow L^2}^j = \|T^j - \Pi\|_{L^2 \rightarrow L^2}$. For each fixed sufficiently large n , take j th roots and let $j \rightarrow \infty$ to obtain $\|T - \Pi\|_{L^2 \rightarrow L^2} \leq \beta_n$. Finally, $b = 1/2 + O(1/n)$ gives $\beta_n = 1/2 + o(1)$. $\square$

5.3 Reducing lattice contraction to a trace moment

For $U \sim \sigma$, define

$$T^2(u, v) = \mathbb{E}_U [T(u, U)T(U, v)], \quad \mathbf{R}_L(x, z) = M_0^2(x, z) - \frac{T^2(\hat{x}, \hat{z})}{V} \quad (x, z \in S_R).$$

The density T^2 is the measure of the intersection of two caps divided by p^2 . The real symmetric matrix $\mathbf{R}_L$ compares this intersection measure with the number of lattice intermediates.

For $m \in 4\mathbb{N}$, set $z_{m+1} = z_1$ and define the cyclic indicator

$$F_R(z_1, \dots, z_m) = \prod_{i=1}^m [\mathbf{1}\{0 < \|z_i\| \leq R\} \mathbf{1}\{\langle z_i, z_{i+1} \rangle \geq b \|z_i\| \|z_{i+1}\|\}]. \quad (5.3)$$

The number of arguments specifies m . Odd positions will be trace endpoints, and even positions will be intermediate points.

Lemma 5.8 (Contraction from trace moments). *If $r \geq 2$ is even, $2r < n$, and $t > 0$, then*

$$\mathbb{P}_L \left\{ \sup_{\substack{f \neq 0 \\ \sum_x f(x)=0}} \frac{\|M_0 f\|^2}{\|f\|^2} > \beta_n^2 + t \right\} \leq \frac{2^r}{t^r} \mathbb{E}_L \text{Tr}(\mathbf{R}_L^r).$$

The event is empty when $N_R < 2$, and the trace is zero when $N_R = 0$.

The comparison replaces two lattice averages by spherical averages. The first replacement produces the matrix $\mathbf{R}_L$. The second produces an integral operator with the same discrepancy kernel. We control their operator norms by even trace moments, and Corollary 5.4 bounds the expected continuous trace moment by the lattice trace moment.

Proof of Lemma 5.8. Fix $N_R \geq 2$ and a nonzero zero-sum f .

Replace the intermediate lattice point by a spherical point. Expanding the discrete square gives

$$\begin{aligned} \|M_0 f\|^2 &= \frac{1}{V} \sum_{x,z} f(x)f(z) \left[\frac{1}{V} \sum_y T(\hat{x}, \hat{y}) T(\hat{y}, \hat{z}) \right], \\ \frac{1}{V} \sum_{x,z} f(x)f(z) T^2(\hat{x}, \hat{z}) &= \frac{1}{V} \mathbb{E}_U \left(\sum_x f(x) T(\hat{x}, U) \right)^2. \end{aligned}$$

All sums range over S_R . The expression inside the continuous square naturally defines

$$h(u) = \frac{1}{N_R} \sum_x f(x) T(\hat{x}, u) = \mathbb{E}_X [f(X) T(\hat{X}, u)], \quad X \sim U(S_R).$$

Since $|T| \leq 1/p$, this function is square-integrable. Moreover, $\int h d\sigma = N_R^{-1} \sum_x f(x) = 0$. The continuous square is $(N_R^2/V) \mathbb{E} h(U)^2$.

Define the discrepancy between the two intermediate-point averages by

$$\Delta(u, v) = \frac{1}{V} \sum_{y \in S_R} T(u, \hat{y}) T(\hat{y}, v) - T^2(u, v). \quad (5.4)$$

Geometrically, $p^2 \Delta$ is the lattice count in a two-cap intersection, divided by V , minus its spherical probability. Subtracting the two squares yields

$$\|M_0 f\|^2 - \frac{N_R^2}{V} \mathbb{E} h(U)^2 = \frac{1}{V} \sum_{x,z} f(x) \Delta(\hat{x}, \hat{z}) f(z). \quad (5.5)$$

Since $\mathbf{R}_L(x, z) = \Delta(\hat{x}, \hat{z})/V$, this is

$$\|M_0 f\|^2 = \frac{N_R^2}{V} \mathbb{E} h(U)^2 + f^t \mathbf{R}_L f. \quad (5.6)$$

Replace the remaining lattice average by a spherical average. The definition of h and Cauchy–Schwarz, applied to $f(X)$ and $(Th)(\hat{X})$, give

$$\mathbb{E}h(U)^2 = \mathbb{E}_X[f(X)(Th)(\hat{X})], \quad (\mathbb{E}h(U)^2)^2 \leq \mathbb{E}f(X)^2 \mathbb{E}_X(Th)(\hat{X})^2.$$

To apply continuous contraction, expand and subtract the two averages:

$$\frac{N_R}{V} \mathbb{E}_X(Th)(\hat{X})^2 - \mathbb{E}_W(Th)(W)^2 = \mathbb{E}_{U,U'}[h(U)\Delta(U,U')h(U')],$$

where U, U', W are independent samples from σ . The kernels are $V^{-1} \sum_x T(U, \hat{x})T(\hat{x}, U')$ and $T^2(U, U')$, so their difference is the same Δ .

Let D_L be the integral operator with kernel Δ , used only in this proof. Lemma 5.7 gives

$$\frac{N_R}{V} \mathbb{E}_X(Th)(\hat{X})^2 \leq (\beta_n^2 + \|D_L\|_{\text{op}}) \mathbb{E}h(U)^2.$$

Combine this with Cauchy–Schwarz and cancel $\mathbb{E}h(U)^2$ when it is positive:

$$\frac{N_R}{V} \mathbb{E}h(U)^2 \leq (\beta_n^2 + \|D_L\|_{\text{op}}) \mathbb{E}f(X)^2.$$

The bound also holds when $h = 0$. Multiply by N_R , use $\|f\|^2 = N_R \mathbb{E}f(X)^2$, and substitute in (5.6). We obtain, for this fixed lattice,

$$\|M_0 f\|^2 \leq (\beta_n^2 + \|R_L\|_{\text{op}} + \|D_L\|_{\text{op}}) \|f\|^2. \quad (5.7)$$

Compare the expected trace moments. By (5.7), contraction can fail only if $\|R_L\|_{\text{op}} + \|D_L\|_{\text{op}} > t$. The matrix R_L and the kernel of D_L are real and symmetric. For even r , the largest r th power of an eigenvalue is at most their sum:

$$\|R_L\|_{\text{op}}^r \leq \text{Tr}(R_L^r), \quad \|D_L\|_{\text{op}}^r \leq \text{Tr}(D_L^r). \quad (5.8)$$

For D_L , boundedness and symmetry of the kernel ensure square-summability of the real eigenvalues, so this argument applies for every even $r \geq 2$.

Use F_R from (5.3) with $m = 2r$. Recall that every S_i sums over the same lattice $L \setminus \{0\}$, while $\mathbf{l}_i$ is unnormalized Lebesgue integration in coordinate i . Expanding the cyclic products gives, for each fixed lattice,

$$\begin{aligned} \text{Tr}(R_L^r) &= \frac{1}{(Vp)^{2r}} \left[\prod_{j=1}^r S_{2j-1} \prod_{j=1}^r (S_{2j} - \mathbf{l}_{2j}) \right] F_R, \\ \text{Tr}(D_L^r) &= \frac{1}{(Vp)^{2r}} \left[\prod_{j=1}^r \mathbf{l}_{2j-1} \prod_{j=1}^r (S_{2j} - \mathbf{l}_{2j}) \right] F_R. \end{aligned} \quad (5.9)$$

Indeed, each copy of Δ contributes one centered intermediate position with factor $1/(Vp^2)$. Each endpoint contributes another factor $1/V$: from $R_L(x, z) = \Delta(\hat{x}, \hat{z})/V$ in the first trace, and from replacing a spherical average by a normalized ball integral in the second. The angular factors are the same in both traces.

In the first line, expand each endpoint operator as $S_{2j-1} = (S_{2j-1} - \mathbf{l}_{2j-1}) + \mathbf{l}_{2j-1}$. Choosing the integral at every endpoint gives the second line. In every other term, first integrate the endpoints at which $\mathbf{l}_{2j-1}$ was chosen. The resulting function is nonnegative, bounded, and compactly supported,

and every remaining coordinate is centered. Corollary 5.4 therefore makes the Haar expectation of that term nonnegative. Since $2r < n$, all expansions are absolutely integrable. Together with the nonnegativity of even trace powers, this proves

$$0 \leq \mathbb{E}_L \operatorname{Tr}(D_L^r) \leq \mathbb{E}_L \operatorname{Tr}(\mathbf{R}_L^r). \quad (5.10)$$

Finally, Markov's inequality and $(a+b)^r \leq 2^{r-1}(a^r + b^r)$ give

$$\begin{aligned} \mathbb{P}_L\{\|\mathbf{R}_L\|_{\text{op}} + \|D_L\|_{\text{op}} > t\} &\leq \frac{2^{r-1}}{t^r} \mathbb{E}_L[\|\mathbf{R}_L\|_{\text{op}}^r + \|D_L\|_{\text{op}}^r] \\ &\leq \frac{2^r}{t^r} \mathbb{E}_L \operatorname{Tr}(\mathbf{R}_L^r). \end{aligned}$$

□

5.4 Bounding a fixed trace moment using integer relations

For $m \in 4\mathbb{N}$ and $n > m$, put

$$Q_m(R) = \mathbb{E}_L \operatorname{Tr}(\mathbf{R}_L^{m/2}).$$

Lemma 5.9 (Fixed trace moment bound). *For every fixed $c > \sqrt{4/3}$, there are $m \in 4\mathbb{N}$ and $\gamma > 0$, depending only on c , such that, for all sufficiently large n and every $R \geq cr_n$,*

$$Q_m(R) \leq e^{-\gamma n}.$$

The moment order is fixed independently of n . Lemma 5.8, applied with order $m/2$ and $t = \rho^2 - \beta_n^2$, then gives

$$\mathbb{P}_L\{\exists f : \sum_x f(x) = 0, \|M_0 f\| > \rho \|f\|\} \leq C_{\rho, m} Q_m(R), \quad (5.11)$$

since $\beta_n \rightarrow 1/2 < \rho$. We first bound each integer relation and then choose m so that their sum decays exponentially.

Expand the centered lattice sums. Temporarily fix $m \in 4\mathbb{N}$ and use F_R from (5.3). The first identity in (5.9), with trace order $m/2$, and Corollary 5.3 give

$$Q_m(R) = \frac{1}{V^m p^m} \sum_{r=1}^{m-1} \sum_{A \in \mathcal{A}_{r, m}^{\text{cen}}} J_A, \quad J_A = \int_{\mathbb{R}^{n \times r}} F_R(YA) dY, \quad (5.12)$$

where every even position participates in each retained relation. All terms are nonnegative, and the expansion is absolutely integrable for $n > m$.

Write a_i for the columns of A , put $a_{m+1} = a_1$, and define

$$h(A) = \max_{|J|=r} |\det A_J|, \quad D_A = \max \left\{ \det G : \begin{array}{l} G \succeq 0, \quad s_i = a_i^t G a_i \leq 1, \\ a_i^t G a_{i+1} \geq b \sqrt{s_i s_{i+1}} \quad (1 \leq i \leq m) \end{array} \right\}.$$

An invertible A_J bounds the feasible matrices, so the maximum exists. Both quantities are invariant under the representative of the relation class.

The next lemma bounds each relation integral using D_A and counts the relation classes with a given $h(A)$. After normalization by $V^m p^m$, we seek bounds $D_A \leq e^{O_m(1/n)} \delta_r$ with

$$c^{r-m} (4/3)^{m/2} \sqrt{\delta_r} < 1.$$

The factors c^{r-m} and $(4/3)^{m/2}$ come from V^{r-m} and p^{-m} , respectively. Decay in $h(A)$ will make the sum over relation classes convergent. The lemma also applies to ball counts, which have no angular constraints.

Lemma 5.10 (Volume and counting bounds). *Let F_R be a Borel indicator invariant under a common orthogonal transformation of its m arguments and supported on $\|z_i\| \leq R$ for every i . Put $J_A = \int F_R(YA) dY$, and suppose $\det(Y^t Y) \leq R^{2r} D_A$ whenever $F_R(YA) = 1$. For fixed m and $n > m$, uniformly in r, A, R ,*

$$J_A \leq \text{poly}(n) V^r h(A)^{-r} D_A^{(n-r)/2}, \quad \#\{A \in \mathcal{A}_{r,m} : h(A) = h\} \leq C_m h^{m^2/4}. \quad (5.13)$$

Here $V = \text{vol } \mathcal{B}_n(R)$, and the polynomial and C_m depend only on m . The counting bound also holds for every subfamily of relation classes.

Proof. Let $\iota : \mathbb{R}^r \rightarrow \mathbb{R}^n$ be the standard isometric embedding. Scaling and the rotation-invariant integration formula [SS19, Lemma 5.2] give

$$J_A = R^{nr} \frac{\prod_{j=n-r+1}^n \omega_j}{\prod_{j=1}^r \omega_j} \int_{\mathbb{R}^{r \times r}} F_R(R\iota(ZA)) \det(Z^t Z)^{(n-r)/2} dZ, \quad \omega_j = j \text{ vol } \mathcal{B}_j(1).$$

Choose J with $|\det A_J| = h(A)$. On the support, $\det(Z^t Z) \leq D_A$ and every column of ZA_J lies in the unit ball. Substituting $W = ZA_J$ bounds the integral by $D_A^{(n-r)/2} h(A)^{-r} \text{vol}(\mathcal{B}_r(1))^r$. Stirling's formula gives the volume bound.

For the count, choose the first column set J attaining h . Cramer's rule makes the J -columns of $A_J^{-1}A$ the identity and every other entry an integer multiple of $1/h$ in $[-1, 1]$. There are at most $\binom{m}{r} (2h+1)^{r(m-r)} \leq C_m h^{m^2/4}$ possibilities. This normalized matrix determines the class: if $A' = UA$ has the same row space, an integer right inverse C of A gives $U = A'C$ integer, and the reverse argument makes U^{-1} integer. $\square$

Bound the Gram determinants. First, Hadamard's inequality at a largest minor gives

$$D_A \leq h(A)^{-2} \leq 1. \quad (5.14)$$

A second elementary bound is

$$D_A \leq (1 - b^2)^{r-1} = e^{O_m(1/n)} (3/4)^{r-1}. \quad (5.15)$$

Indeed, for positive definite G , scan the vectors $G^{1/2}a_i$ in order and retain those that increase rank. Every retained vector after the first has a nonzero predecessor in the previous span, so its squared Gram-Schmidt length is at most $1 - b^2$. The selected integer minor has absolute value at least one. Singular G has determinant zero.

Writing $d = m - r$, the sequential bound gives the exponential base $\sqrt{4/3}(3c^2/4)^{-d/2}$. This is less than one for large d , but does not improve when m grows with d fixed. For the remaining ranks, we use all edges of the cycle and the fact that every even position participates in a relation. The next lemma turns this participation into a determinant bound.

Lemma 5.11 (Determinants of integer kernels). *Let $A \in \mathbb{Z}^{r \times m}$ satisfy $A\mathbb{Z}^m = \mathbb{Z}^r$, let $d = m - r \geq 1$, and let the columns of $B \in \mathbb{Z}^{m \times d}$ form an integer basis of $\ker(A) \cap \mathbb{Z}^m$. If B has s nonzero rows and H is positive definite, then*

$$\det(AHA^t) = \det H \det(B^t H^{-1} B), \quad \det(B^t B) \geq s - d + 1.$$

In particular, if $H \preceq \lambda I$, then

$$\det(AHA^t) \geq \det H \lambda^{-d}(s-d+1).$$

Proof. Take an integer right inverse C of A . The matrix $W = (C, B)$ is unimodular, since every integer vector x decomposes uniquely into CAx and an integer combination of the columns of B . The upper-left block of $W^{-1}HW^{-t}$ is AHA^t , and the complementary lower-right block of its inverse is $B^tH^{-1}B$. The complementary-minor identity gives the determinant formula.

Choose d independent rows of B . Each additional nonzero row has a nonzero coefficient in this row basis and can replace a corresponding basis row to give a nonzero minor. These $s-d$ row sets are distinct from one another and from the original basis set. All nonzero minors are integers, so Cauchy–Binet gives $\det(B^tB) \geq 1 + s - d$. Finally, $H^{-1} \succeq \lambda^{-1}I$ gives the stated consequence. $\square$

Let S_m be the cyclic shift matrix. The coefficient $1/2$ matches the limiting cap threshold. Normalizing by $1 - (1/2)^2 = 3/4$, set

$$H = \frac{4}{3}(I - \frac{1}{2}S_m)(I - \frac{1}{2}S_m)^t = \frac{1}{3}I + \frac{2}{3}L(C_m), \quad Q_A = \frac{r}{m}AHA^t.$$

For every feasible G , the vectors $w_i = G^{1/2}a_i$ satisfy $\|w_i\| \leq 1$ and $\|w_i - w_{i+1}\| \leq \tau$. Hence

$$\text{Tr}(GQ_A) = \frac{r}{m} \left(\frac{1}{3} \sum_i \|w_i\|^2 + \frac{2}{3} \sum_i \|w_i - w_{i+1}\|^2 \right) \leq r \left(\frac{1}{3} + \frac{2}{3}\tau^2 \right).$$

Arithmetic–geometric mean yields

$$D_A \leq e^{O_m(1/n)} / \det Q_A. \quad (5.16)$$

The factorization gives

$$\det H = (4/3)^m (1 - 2^{-m})^2, \quad H \preceq \frac{4}{3}(1 + 1/2)^2 I = 3I.$$

Every even position participates in a relation, so at least $m/2$ rows of its integer kernel basis are nonzero. Lemma 5.11 therefore gives, whenever $m > 2d$,

$$\det Q_A \geq (4/3)^m K_d(m), \quad K_d(m) = e^{-d}(1 - 2^{-m})^2 3^{-d}(m/2 - d + 1), \quad (5.17)$$

where we used $(1 - d/m)^{m-d} > e^{-d}$. The factor $(4/3)^m$ cancels the angular normalization in the exponential base. For each fixed d , the remaining gain $K_d(m)$ tends to infinity as m grows through multiples of four.

Proof of Lemma 5.9. Since $3c^2/4 > 1$, choose a fixed integer $D \geq 2$ with

$$\sqrt{4/3}(3c^2/4)^{-D/2} < 1.$$

Then choose a fixed $m \in 4\mathbb{N}$, with $m > 2D$, large enough that $c^{-d}/\sqrt{K_d(m)} < 1$ for every $1 \leq d < D$. Set

$$\delta_r = \begin{cases} (3/4)^{r-1}, & m-r \geq D, \\ (3/4)^m / K_{m-r}(m), & 1 \leq m-r < D. \end{cases}$$

The determinant estimates give $D_A \leq e^{O_m(1/n)} \delta_r$. Moreover,

$$b_r := c^{r-m} (4/3)^{m/2} \sqrt{\delta_r} < 1.$$

For $d = m - r \geq D$, this is $\sqrt{4/3} (3c^2/4)^{-d/2}$; for $d < D$, it is $c^{-d}/\sqrt{K_d(m)}$.

There are only finitely many ranks. Choose a fixed $\epsilon \in (0, 1)$ sufficiently small that

$$b'_r := c^{r-m} (4/3)^{m/2} \delta_r^{(1-\epsilon)/2} < 1 \quad (1 \leq r < m).$$

Interpolating with (5.14) gives $D_A \leq e^{O_m(1/n)} \delta_r^{1-\epsilon} h(A)^{-2\epsilon}$. Since $V \geq c^n$ and $p = n^{-1/2} (3/4)^{n/2} e^{O(1)}$, Lemma 5.10 implies

$$\frac{J_A}{V^m p^m} \leq \text{poly}(n) (b'_r)^n h(A)^{-\epsilon n - (1-\epsilon)r} \leq \text{poly}(n) (b'_r)^n h(A)^{-\epsilon n}.$$

Thus the nonnegative relation expansion satisfies

$$Q_m(R) \leq \text{poly}(n) \sum_{r=1}^{m-1} (b'_r)^n \sum_{h \geq 1} h^{m^2/4 - \epsilon n} \leq e^{-\gamma n}$$

for some $\gamma > 0$ and all sufficiently large n . The sum over h is bounded for all sufficiently large n . Every choice and constant depends only on c and is independent of R . $\square$

5.5 Summing the contraction error over all radii

Fix the m supplied by Lemma 5.9. It remains to make the contraction bound simultaneous over all prescribed radii. Put $R_0 = cr_n$ and $V_0 = c^n$. Lemma 5.9 gives $Q_m(R_0) = e^{-\Omega(n)}$.

In a rank- r relation integral, the substitution $Y = (R/R_0)Y'$ preserves angles and multiplies volume by $(V/V_0)^r$. Normalization therefore multiplies the relation by $(V/V_0)^{r-m}$. Positivity and $r \leq m-1$ imply

$$Q_m(R) \leq \frac{V_0}{V} Q_m(R_0) \quad (R \geq R_0). \quad (5.18)$$

Since $V_j/V_0 = \tau^{nj}$, the expected number of failing radii bounds the probability of any failure:

$$\mathbb{P}_L\{\text{contraction fails at some } R_j\} \leq C_{\rho,m} \sum_{j \geq 0} Q_m(R_j) \leq \frac{C_{\rho,m} Q_m(R_0)}{1 - \tau^{-n}} = e^{-\Omega(n)}.$$

The denominator tends to $1 - e^{-1} > 0$.

6 Uniform degree bounds

It remains to bound the row sums in Lemma II.2. We treat the grid radii $cr_n \leq R \leq n^2 r_n$ by conditioning on a primitive vector, and all larger radii by the covering radius. The dividing radius ensures that the covering error $n\mu(L)/R$ tends to zero, where $\mu(L)$ is the covering radius.

6.1 Small and intermediate radii: conditional moments

Proposition 6.1 (Uniform degrees). *Fix $c > \sqrt{4/3}$. With Haar probability $1 - o(1)$, simultaneously for all grid radii $R = c\tau^j r_n$ in $[cr_n, n^2 r_n]$,*

$$N_R = (1 + o(1))V, \quad d_0(x) = (1 + o(1))\nu \quad (x \in S_R). \quad (6.1)$$

More precisely, there is a constant $\eta > 0$ such that N_R/V and every $(d_0(x) + 1)/\nu$ differ from one by at most $e^{-\eta n}$, except on an event of probability $e^{-\Omega(n)}$.

The total count follows from $\mathbb{E}_L N_R = V$ and $\text{Var}_L(N_R) = O(V)$, by Theorem 4.7. For degrees, we first bound the conditional cap count of one fixed primitive vector, then average over these vectors.

Condition on a primitive vector. Vertices on the same ray have the same cap count $d_0(x) + 1$, so it suffices to consider primitive vectors. Apply Theorem 4.8 to the indicator that a primitive vector in the radius- R ball has a degree deviation larger than the stated tolerance. It gives

$$\mathbb{E}_L \#\{\text{bad primitive vectors in } S_R\} = \frac{1}{\zeta(n)} \int_{\mathcal{B}_n(R)} \mathbb{P}_v^{\text{prim}}\{\text{the cap count at } v \text{ is bad}\} dv.$$

Here $\mathbb{P}_v^{\text{prim}}$ is the conditional lattice distribution from that theorem. Thus a uniform bound on the integrand, multiplied by V , bounds the probability that some direction violates the degree bound. Fix a primitive vector v , put $h = \|v\|$ and $m = n - 1$, and rotate so that $v = he_1$.

Count cap points along lines parallel to the fixed vector. Using the conditional representation (4.8), set

$$C = \{x : 0 < \|x\| \leq R, \ x_1 \geq b\|x\|\}, \quad F(z, t) = \sum_{j \in \mathbb{Z}} \mathbf{1}_C(h(j+t)e_1 + h^{-1/m}z).$$

The value $F(z, t)$ counts cap points on the line $h^{-1/m}z + \mathbb{R}v$ with shift t . The count off the axis $\mathbb{R}v$ is

$$Z = \sum_{z \in \Lambda' \setminus \{0\}} F(z, \omega(z)).$$

The shifts $\omega(z)$ come from the same random homomorphism. Each line contributes at most $J_1 = 1 + 2R/h$ points, and only z in $B = h^{1/m}\mathcal{B}_m(R\sqrt{1-b^2})$ can contribute. Hence

$$\begin{aligned} 0 \leq F(z, t) &\leq J_1 \mathbf{1}_B(z), & U &:= \text{vol } B = h \text{vol } \mathcal{B}_m(R\sqrt{1-b^2}), \\ U \mathbb{E}F(Z_0, \Theta) &= \nu, & Z_0 &\sim U(B), \quad \Theta \sim U(\mathbb{T}) \text{ independently.} \end{aligned} \quad (6.2)$$

Indeed, integrating over t and summing the axial translates covers the whole axis. The axial Jacobian h and transverse Jacobian h^{-1} cancel, so the last mean equals $\text{vol } C = \nu$. The first-moment formula then gives $\mathbb{E}Z = \nu$. Nonzero multiples of v contribute an additional $O(J_1)$ axial points. Deleting the vertex itself changes the count by one.

Equation (6.2) verifies the hypotheses of Lemma 4.9. We may therefore bound the centered moment of the count in a cap by a moment of the count in a ball, keeping all dependencies among the shifts.

Bound the centered moments of ball counts. The conditional estimate must survive the outer integral over primitive vectors, whose volume is V . We therefore seek a relative moment bound of

order e^{-Tm}/U : the factor $1/U$ offsets most of this volume, and choosing T sufficiently large absorbs the remaining factor V/U .

For a Haar-random determinant-one lattice $\Lambda' \subset \mathbb{R}^m$ and an origin-centered ball $B \subset \mathbb{R}^m$ of volume U , write

$$\mathcal{M}_k(U) = \mathbb{E}(N_B - U)^k, \quad N_B = \#((\Lambda' \setminus \{0\}) \cap B).$$

Lemma 6.2 (Moments of centered ball counts). *For every fixed $a > 1$ and $T > 0$, there is a fixed even integer k such that, for all sufficiently large dimensions m and every $U \geq a^m$,*

$$\mathcal{M}_k(U) \leq U^{k-1} e^{-Tm}. \quad (6.3)$$

Proof. The fully centered relation expansion writes $\mathcal{M}_k(U)$ as a nonnegative sum of integrals J_A of ranks $1 \leq r < k$. The proof of Lemma 5.10, with the cap indicators omitted, gives

$$J_A \leq \text{poly}(m) U^r h(A)^{-r} D_A^{(m-r)/2}, \quad \#\{A : h(A) = h\} \leq C_k h^{k^2/4},$$

where D_A maximizes $\det G$ subject to $G \succeq 0$ and $a_i^t G a_i \leq 1$ for every column. Hadamard's inequality gives $D_A \leq h(A)^{-2} \leq 1$.

Put $d = k - r$. Full centering makes every row of an integer kernel basis nonzero. Lemma 5.11, with $H = I$, therefore gives $\det(AA^t) \geq k - d + 1$. Since $\text{Tr}(G(r/k)AA^t) \leq r$, arithmetic-geometric mean yields

$$D_A \leq \frac{(k/r)^r}{\det(AA^t)} < \frac{e^d}{k - d + 1}.$$

Fix $T_* > T$ and choose an integer D such that $a^{1-D} < e^{-T_*}$. Then choose a sufficiently large even $k > D$ so that

$$a^{1-d} \sqrt{\frac{e^d}{k - d + 1}} < e^{-T_*} \quad (1 \leq d < D).$$

Use $\delta_r = e^d/(k - d + 1)$ for $d < D$ and $\delta_r = 1$ otherwise. Thus $D_A \leq \delta_r$ and every $a^{1-d} \sqrt{\delta_r} < e^{-T_*}$. Since these inequalities are strict, a sufficiently small fixed $\epsilon \in (0, 1)$ preserves

$$a^{1-d} \delta_r^{(1-\epsilon)/2} < e^{-T_*}.$$

Interpolate with $D_A \leq h(A)^{-2}$. Since $U^{1-d} \leq a^{m(1-d)}$, the volume bound gives

$$\frac{J_A}{U^{k-1}} \leq \text{poly}(m) (a^{1-d} \delta_r^{(1-\epsilon)/2})^m h(A)^{-\epsilon m - (1-\epsilon)r}.$$

Summing the counting bound over r and h proves

$$\frac{\mathcal{M}_k(U)}{U^{k-1}} \leq \text{poly}(m) e^{-T_* m} \sum_{h \geq 1} h^{k^2/4 - \epsilon m} \leq e^{-Tm}.$$

The sum over h is bounded for all sufficiently large m , uniformly in U . $\square$

Proof of Proposition 6.1. To keep the contribution from the axis polynomial, first exclude lattices with $\lambda_1(L) < r_n/2$. The first-moment formula bounds their probability by 2^{-n} . Outside this event, every primitive vector satisfies $h \geq r_n/2$. We may therefore restrict the outer integral to $r_n/2 \leq h \leq R$, keeping the original conditional distribution for each fixed primitive vector.

Fix $c > \sqrt{4/3}$, choose constants $1 < a < \sqrt{3c^2/4}$ and $T > \frac{1}{2} \log(4/3)$, and take the fixed even k from Lemma 6.2. The ball-volume formula gives, uniformly over the retained vectors and radii,

$$U \geq (3c^2/4)^{n/2} n^{-O(1)} \geq a^{n-1}, \quad \frac{J_1 U}{\nu} = O(n), \quad \frac{V}{U} \leq \text{poly}(n)(4/3)^{n/2}.$$

Applying Lemmas 4.9 and 6.2 in dimension $m = n - 1$ yields

$$\frac{\mathbb{E}(Z - \nu)^k}{\nu^k} \leq \left(\frac{J_1 U}{\nu} \right)^k \frac{\mathcal{M}_k(U)}{U^k} \leq \frac{\text{poly}(n) e^{-T(n-1)}}{U}.$$

For $r_n/2 \leq h \leq R \leq n^2 r_n$, the additional axial count is polynomial in n , whereas $\nu \geq (3c^2/4)^{n/2} n^{-O(1)}$. Its relative contribution is therefore exponentially small.

Markov's inequality at relative tolerance $e^{-\eta n}$, followed by the outer integral over v , bounds the probability that some direction violates the degree bound by

$$2^{-n} + \text{poly}(n) e^{k\eta n - T(n-1)} (4/3)^{n/2} \leq 2^{-n} + \text{poly}(n) e^{-[T - \frac{1}{2} \log(4/3) - k\eta]n}.$$

Choose $\eta > 0$ sufficiently small that

$$k\eta < \min\{T - \frac{1}{2} \log(4/3), \frac{1}{2} \log(3c^2/4), \log c\}.$$

This also absorbs the axial contribution and makes the total-count failure probability $O(e^{2\eta n}/V)$ exponentially small. A union bound over the $O(n \log n)$ grid radii proves Proposition 6.1 for every fixed $c > \sqrt{4/3}$. $\square$

6.2 Large radii: the covering radius

For $R \geq n^2 r_n$, Voronoi cells compare point counts with volumes uniformly over every direction. Write $\mu(L)$ for the covering radius, the supremum over points in $\mathbb{R}^n$ of their distance to L . Equivalently, this is the circumradius of a Voronoi cell.

Theorem 6.3 (Random-lattice covering, Ordentlich–Regev–Weiss [ORW22, Corollary 1.4]). *There are absolute constants $C, c_0 > 0$ such that, for all sufficiently large n , a Haar-random determinant-one lattice satisfies*

$$\mathbb{P}_L \left\{ \text{vol } \mathcal{B}_n(\mu(L)) > Cn^3 \right\} \leq e^{-c_0 n}.$$

Since $\text{vol } \mathcal{B}_n(r_n) = 1$, this theorem gives

$$\mu(L) \leq (Cn^3)^{1/n} r_n = O(r_n)$$

outside an exponentially small exceptional set.

Proposition 6.4 (Uniform counts at large radii). *Suppose $\det L = 1$ and $\mu(L) \leq Cr_n$. For all sufficiently large n , the estimates (6.1) hold uniformly over all real $R \geq n^2 r_n$. Moreover, for every ball,*

$$\#(L \cap \mathcal{B}_n(t, s)) \leq \text{vol } \mathcal{B}_n(s + \mu(L)).$$

Proof. Put $\mu = \mu(L)$, and let C_x be the volume-one Voronoi cell centered at x . We first give one count estimate for any convex body K containing a ball $\mathcal{B}_n(a, r)$, with $\mu < r$. Write $K \ominus \mathcal{B}_n(\mu) = \{y : y + \mathcal{B}_n(\mu) \subseteq K\}$ for its inner parallel body. Rounding moves points by at most μ , so

$$K \ominus \mathcal{B}_n(\mu) \subseteq \bigcup_{x \in L \cap K} C_x \subseteq K + \mathcal{B}_n(\mu).$$

Convexity and $\mathcal{B}_n(a, r) \subseteq K$ give

$$\begin{aligned} a + (1 - \mu/r)(K - a) &\subseteq K \ominus \mathcal{B}_n(\mu), \\ K + \mathcal{B}_n(\mu) &\subseteq a + (1 + \mu/r)(K - a). \end{aligned}$$

Both inclusions follow by taking convex combinations with the contained ball. Taking volumes yields

$$(1 - \mu/r)^n \text{vol } K \leq \#(L \cap K) \leq (1 + \mu/r)^n \text{vol } K.$$

Apply this first to $K = \mathcal{B}_n(R)$, whose inradius is R . For a unit direction u , the closed cap cone

$$K = \{x : \|x\| \leq R, \langle x, u \rangle \geq b \|x\|\}$$

is convex and contains $\mathcal{B}_n(3Ru/4, R/4)$: in this ball, $\langle x, u \rangle \geq R/2$, $\|x\| \leq R$, and $0 < b \leq 1/2$ for sufficiently large n . Thus the same estimate applies with $r = R/4$. Since $n\mu/R = O(1/n)$ in the stated range, both counts differ from their volumes by a relative $O(n\mu/R)$. Removing the origin gives

$$\frac{N_R}{V} = 1 + O(n\mu/R) + O(V^{-1}), \quad \frac{d_0(x) + 1}{\nu} = 1 + O(n\mu/R) + O(\nu^{-1}),$$

uniformly in R and the direction. These errors are $o(1)$. Finally, the cells centered in any ball $\mathcal{B}_n(t, s)$ lie in $\mathcal{B}_n(t, s + \mu)$. Taking volumes proves the remaining assertion. $\square$

Complete the spectral estimate for the full angular graph.

Proof of Proposition II.1. Intersect the covering event, the small- and intermediate-radius degree event, and the contraction event on the full grid. This intersection still has Haar probability $1 - o(1)$. On it, all grid radii satisfy the degree condition with $\epsilon_n = o(1)$ and the contraction condition with the fixed $\rho < 1$. Lemma II.2 therefore gives, uniformly over these radii,

$$\lambda_2(L_0) \geq \nu(1 - \epsilon_n - \rho) \geq \kappa\nu$$

for all sufficiently large n . The accompanying point counts and degree estimates follow from the two radius ranges above. $\square$

Proposition II.1 supplies the geometric input to Lemma 3.5, completing the proof of Theorem 4.3.

Appendices

A Spherical filtering for transitions and difference search

The Markov chain requires all admissible signed proposals at each query state. Its transition rule depends on the complete proposal set and its size. We use the product filters of Becker, Ducas, Gama, and Laarhoven to report this set and the short differences between the final lists [BDGL16, Laa16].

For a fixed parent list and a fixed batch of queries, the reporting routine tests every candidate exactly. The probability that it completes a batch while omitting any admissible proposal is at most the assigned error allowance. It returns failure if a capacity or time budget is exceeded. This coverage guarantee holds for every fixed batch. We bound the expected cost using the reference execution of Remark 4.5.

A.1 Reporting all admissible increments

Fix the parent list $D = (y_1, \dots, y_M)$ and all Q queries for the current step. For a query v and a signed parent entry (i, s) , put $x = -v/R$ and $y = sy_i/(\tau R)$. An admissible move satisfies $\|x - \tau y\| \leq 1$, hence $\|x - y\| \leq 1 + 1/n$. We use this distance bound to generate candidates, then test each candidate against the state-space and angular conditions in (3.2). We compare every data point of norm below $\sqrt{3/4}$ with all queries, and every query of norm below $\sqrt{3/4}$ with the entire parent list. The remaining points use the following filters.

Definition A.1 (Filter region). For a unit direction u , let

$$E_u = \{z : \sqrt{3/4} \leq \|z\| \leq 1, \quad \|z - u/2\|^2 \leq 3/4\}. \quad (\text{A.1})$$

For $r = \|z\|$, membership is equivalent to $\langle \hat{z}, u \rangle \geq A(r)$, where $A(r) = r - 1/(2r)$. Points sharing a filter are candidate pairs.

Following the product-code construction of [BDGL16, Definition 5.1], split the space into $m = \lceil \log n \rceil$ orthogonal blocks whose dimensions differ by at most one. Sample directions independently in each block, concatenate them with weights proportional to the square roots of the block dimensions, and apply an independent Haar rotation. The original construction uses equal blocks. Appendix B proves the coverage property needed for the present block sizes. The parameters there give

$$F = (3/2)^{n/2} 2^{o(n)} \quad (\text{A.2})$$

implicit directions using $2^{o(n)}$ representation space. Every global direction has a uniform marginal distribution.

Algorithm A.2 (Batch queries with independent product codes).

Input: Frozen parent list D , Q current queries, and $0 < \zeta < 1$.

During enumeration, return $\perp$ immediately if the budget is exhausted or a query has more than K admissible proposals.

- 1: $\ell \leftarrow \lceil \log_2(2MQ/\zeta) \rceil$
 - 2: Generate ℓ fresh rotations and block codes, independently of the history and transition coins.
 - 3: Perform the direct comparisons for points of small norm.
 - 4: Process each code by prefixes and buckets, testing every candidate exactly against (3.2).
 - 5: Deduplicate reports of the same (i, s) , preserving distinct parent indices even when their vectors coincide.
 - 6: Order each proposal list lexicographically by (i, s) , with -1 before $+1$.
 - 7: Execute (3.3) with independent transition coins.
-

Each output step uses fresh codes. Burn-in still scans the parent list directly.

For every fixed near pair in the outer annulus, one code contains a common filter with probability at least $1/2$, uniformly over unequal block sizes and nearly parallel pairs (Lemma B.1). Repeating ℓ times and taking a union bound over at most $2MQ$ pairs bounds the batch's missed-pair probability by ζ . Conditional on the history at the start of the step, the queries are fixed and the fresh codes remain independent, so the same bound applies to adaptive trajectories. Couple the implementation to complete neighbor enumeration by sharing the transition coins. They agree until the first missed pair or budget failure. The fixed ordering removes any dependence on the filters' reporting order.

A.2 Reporting short differences between two lists

We now implement the final difference search in Algorithm 4.2. Let $R_0 = cr_n$, and let $(a_1, \dots, a_M)$ and $(b_1, \dots, b_M)$ be the two final lists of samples from the ball. Use the first list as queries and the

second as data, report all differences satisfying

$$0 < \|a_i - b_j\| \leq R_0,$$

and return the shortest. Since the typical shortest length is $(1 + o(1))r_n < R_0$, this threshold includes the shortest differences used in the success argument in Lemma 4.10.

Set $x = a_i/R_0$ and $y = b_j/R_0$. Both lie in the unit ball, and the target condition is $\|x - y\| \leq 1$. As in a step of the Markov chain, directly compare pairs with at least one norm below $\sqrt{3/4}$ and use the regions E_u for the remaining pairs. The final difference search uses the directions of the two samples themselves. The normalization $x = -v/R$ for a Markov chain step uses a minus sign because the proposal is $v + sy_i$.

After fixing the two lists, independently generate $\ell_{\text{end}} = \lceil \log_2(M^2/\zeta_{\text{end}}) \rceil$ fresh codes, where $0 < \zeta_{\text{end}} < 1$ is the allowed failure probability for this search. The fixed-pair coverage bound and a union bound over M^2 pairs ensure that every target difference is reported with probability at least $1 - \zeta_{\text{end}}$. In the batch algorithm, replace the proposal lists by a single shortest nonzero difference found so far. Repeated reports only repeat a length comparison. The chain's degree cutoff is not used. Return the shortest difference after processing all buckets, or fail if none exists or the budget is exhausted.

For ideal independent lists, a fixed pair is uniform on $S_0 \times S_0$. Thus the filter membership and comparison bounds below apply with both sampling radii equal to R_0 and $Q = M$. Direct comparisons contribute the inner-ball term $M^2(3/4)^{n/2}$ in expectation. Candidate counts include equal vectors, whose zero differences are discarded. Coupling the entire final list to the ideal product distribution transfers coverage, budget, and success events with a single total-variation error.

A.3 Time and space from membership and candidate counts

We first record the deterministic enumeration bound underlying both implementations.

Lemma A.3 (Enumeration with blocks of different sizes [BDGL16, Algorithm 1 and proof of Lemma 5.1]). *Let $C_i \subset \mathbb{R}^{d_i}$ be fixed block codes of sizes $B_i \geq 1$, where $\sum_i d_i = n$, and let Q_{rot} be an orthogonal matrix. For a fixed target t and real threshold a , all $z \in Q_{\text{rot}}(C_1 \times \cdots \times C_m)$ satisfying $\langle t, z \rangle \geq a$ can be enumerated in time*

$$O\left(n^2 + \sum_i d_i B_i + \sum_i B_i \log B_i + m(N_{\text{out}} + 1)\right).$$

Here N_{out} is the output count. Each output gives the original sampling indices of its m block codewords, without computing the rotated vector. The enumeration may be restricted to a specified prefix or stopped after a specified number of outputs.

Proof. Use Algorithm 1 of [BDGL16], allowing each level its own size B_i . Rotation, block inner products, and sorting cost $O(n^2)$, $O(\sum_i d_i B_i)$, and $O(\sum_i B_i \log B_i)$, respectively. Keep a prefix only if its value plus the largest possible suffix reaches a . Every retained prefix therefore has an output descendant, independently of block sizes. The enumeration visits $O(m(N_{\text{out}} + 1))$ nodes, where the added one covers empty output. $\square$

This uses the deterministic enumeration procedure underlying the published average-time bound.

Count one filter membership whenever a point z belongs to a filter E_u . Group points by complete codeword, splitting prefixes when their capacity is exceeded and releasing each bucket after processing it. Let J_D and J_X count memberships for data points and queries, and let C_{pair} count exact comparisons, including those for points of small norm. For fixed inputs these counts depend on the fresh codes. Every choice of codes satisfies

$$T_{\text{batch}} \leq 2^{o(n)} \text{poly}(n, \ell)(M + Q + J_D + J_X + C_{\text{pair}}), \quad (\text{A.3})$$

$$S_{\text{batch}} \leq 2^{o(n)} \text{poly}(n, \ell)(M + Q + QK). \quad (\text{A.4})$$

Appendix B proves the streaming bounds. The final difference search stores only the shortest nonzero difference found so far.

A.4 Bounding filtering cost by moments of lattice-point counts

By (A.3), it remains to bound the membership and comparison counts. Use the reference execution of Remark 4.5. Conditional on L, D , every query X is marginally uniform on S , including when D fails the gap or capacity conditions.

Write $N_R = |S|$, let N_+ be the number of nonzero parent-ball points, and let $N_A(L)$ count nonzero lattice points in a region A . For fixed regions A, B , the random count $\sum_i \mathbf{1}_A(X) \mathbf{1}_B(y_i)$ counts candidate pairs of a query and a parent entry. Conditional on the parent list, the probability that $X \in A$ is N_A/N_R . Averaging the number of parent entries in B therefore gives

$$\begin{aligned} \mathbb{E}_{D, X|L} \sum_{i=1}^M \mathbf{1}_A(X) \mathbf{1}_B(y_i) &= \frac{N_A(L)}{N_R} \mathbb{E}_{D|L} \sum_{i=1}^M \mathbf{1}_B(y_i) \\ &= \frac{M}{N_R N_+} N_A(L) N_B(L). \end{aligned} \quad (\text{A.5})$$

The expectation averages over the full product distribution of D , without conditioning on the list satisfying its spectral or capacity bounds. Taking Haar expectation on the right gives the mean candidate count for random lattices.

Apply Theorem 4.7 and Cauchy–Schwarz to the centered counts: $\text{Cov}(N_A, N_B) \leq C\sqrt{\text{vol } A \text{ vol } B}$. For the filter regions $A = -RE_u$ and $B = \tau RE_u$, the volumes are at most $V(3/4)^{n/2}$ and $\tau^n V(3/4)^{n/2}$. Thus

$$\mathbb{E}_L N_A(L) N_B(L) \leq \text{vol } A \text{ vol } B + C\tau^n V(3/4)^{n/2}. \quad (\text{A.6})$$

The covariance term accounts for collinear lattice relations. Uniform marginals of the fresh code-words and linearity of expectation now allow us to sum the mean work over buckets.

On the counting event $N_R \asymp V$ and $N_+ \asymp \tau^n V$, the first moment bounds the direct comparisons by $O(MQ(3/4)^{n/2})$ and filter memberships by $O(F(M + Q)(3/4)^{n/2})$ in expectation. The second moment bounds candidate pairs by $O(FMQ[(3/4)^n + (3/4)^{n/2}/V])$. Substitution in (A.3) bounds the mean batch cost by $2^{o(n)} \text{poly}(n, \ell)$ times

$$(M + Q) + MQ(3/4)^{n/2} + F[(M + Q)(3/4)^{n/2} + MQ(3/4)^n + MQ(3/4)^{n/2}/V]. \quad (\text{A.7})$$

Here and below, a bound restricted to the counting event means that the corresponding nonnegative cost is multiplied by that event’s indicator before expectation. This permits the deterministic denominator bounds without conditioning the Haar moment formulas.

These are the cost bounds used in the algorithmic reduction of Section 4.

The lattice estimates must hold on a set independent of the input basis. This follows by bounding a fixed collection of geometric quantities. For $a \in \{1, \tau\}$, take

$$\frac{N_{\mathcal{B}_n(\sqrt{3/4}R)}(L)}{V(3/4)^{n/2}}, \quad \frac{\mathbb{E}_u N_{RE_u}(L)}{V(3/4)^{n/2}}, \quad \frac{\mathbb{E}_u [N_{RE_u}(L) N_{aRE_u}(L)]}{a^n (V^2(3/4)^n + V(3/4)^{n/2})}. \quad (\text{A.8})$$

The expectations here are over a uniform direction u . Since $L = -L$, the sign in the query region does not affect these counts. The three quantities control direct comparisons, memberships, and candidate pairs, respectively. Their Haar means are bounded by an absolute constant, by the preceding first- and second-moment estimates. The case $a = 1$ includes the final difference search.

Take these quantities at every grid radius $R \leq \tau n^2 r_n$, including the parent radius for each smaller target ball, and call the resulting collection $\mathcal{C}_n$. It has size $O(n \log n)$ and depends only on n and the fixed c . Thus

$$\mathbb{P}_L \{ \max_{G \in \mathcal{C}_n} G(L) > A_n \} \leq \frac{C|\mathcal{C}_n|}{A_n} = o(1) \quad \text{if } n \log n = o(A_n). \quad (\text{A.9})$$

Intersect this event's complement with the uniform point-count event. The bounds on N_R, N_+ now control the denominators in (A.5). For any basis and any M, Q, ℓ , these parameters are only deterministic coefficients multiplying the same geometric quantities. At large radii, Proposition 6.4 gives bounds uniform over all $R \geq n^2 r_n$ and directions u . For example, with $\mu = \mu(L)$,

$$N_{-RE_u}(L) \leq \text{vol } \mathcal{B}_n(\sqrt{3/4}R + \mu) = V(3/4)^{n/2} \exp(O(n\mu/R)),$$

since $-RE_u$ lies in the ball of radius $\sqrt{3/4}R$ centered at $-Ru/2$. The parent filter has the same bound after scaling. These events cover every required radius. Finally, on each lattice in this set, apply Markov's inequality to the reference execution to choose a budget. The coupling after burn-in transfers the bound on budget exhaustion to the actual execution.

B Coverage and streaming for product filters

This appendix proves the coverage and enumeration guarantees used in Algorithm A.2. The explicit block parameters and budgets appear in the initialization appendix.

B.1 Covering a near pair with a product filter

We construct a vector of squared norm at most $1/3 + o(1)$ satisfying both filter inequalities. Its projections will give success probability at least $(2/3)^{d_i/2 - o(d_i)}$ in a block of dimension d_i , which determines the total number of codewords.

Lemma B.1 (Coverage for a fixed pair). *Let $x = ru, y = sv$, where u, v are unit vectors, $r, s \in [\sqrt{3/4}, 1]$, and $\|x - y\| \leq 1 + 1/n$. There is a random product code with $F = (3/2)^{n/2} 2^{o(n)}$ codewords and $2^{o(n)}$ representation space such that x, y share a filter with probability $1 - o(1)$, uniformly over all such pairs. In particular, the probability is at least $1/2$ for all sufficiently large n .*

We use three standard geometric facts: the first bounds the norm of the constructed vector, the second preserves the filter inequalities under a random projection, and the third bounds the probability of sampling a suitable direction.

Theorem B.2 (Planar Jung theorem [CKW20, Appendix A, Theorem 9]). *Every nonempty compact set in $\mathbb{R}^2$ of diameter D is contained in a closed disk of radius at most $D/\sqrt{3}$.*

Lemma B.3 (Random orthogonal projection [DG03, Lemma 2.2]). *Let Π project onto a uniformly random k -dimensional subspace of $\mathbb{R}^n$, where $1 \leq k < n$. For every fixed $z \neq 0$ and $0 < \epsilon < 1$,*

$$\mathbb{P} \left\{ \left| \|\Pi z\|^2 - \frac{k}{n} \|z\|^2 \right| > \epsilon \frac{k}{n} \|z\|^2 \right\} \leq 2e^{-k\epsilon^2/12}.$$

This is the relative-deviation consequence of the cited lemma at thresholds $1 - \epsilon$ and $1 + \epsilon$.

Lemma B.4 (Two-dimensional projection of a uniform sphere point [BDGL16, Appendix A, proof of Lemma A.1]). *Let $k \geq 3$ and let Z be uniform on S^{k-1} . In orthonormal coordinates on any fixed two-dimensional plane, its projection has density*

$$\frac{k-2}{2\pi} (1 - \|z\|^2)^{(k-4)/2} \mathbf{1}_{\{\|z\| < 1\}}, \quad z \in \mathbb{R}^2,$$

with respect to two-dimensional Lebesgue measure.

Proof of Lemma B.1. Construct a vector satisfying both filter inequalities. Put $a = A(r)$ and $d = A(s)$. In a two-dimensional plane containing $0, x, y$, take their smallest enclosing disk, with center w_0 . Orthogonal projection implies $w_0 \in \text{span}\{u, v\}$. The three distances are at most $1 + 1/n$, so Theorem B.2 bounds the radius by $(1 + 1/n)/\sqrt{3}$.

For large n ,

$$\langle x, y \rangle = \frac{r^2 + s^2 - \|x - y\|^2}{2} \geq \frac{3/2 - (1 + 1/n)^2}{2} > 0.$$

The origin lies on the boundary of the smallest disk. Otherwise only x, y constrain the radius, and the disk has diameter xy . But this disk contains the origin exactly when $\langle x, y \rangle \leq 0$, a contradiction. Its radius is consequently $\|w_0\|$. From $\|x - w_0\| \leq \|w_0\|$,

$$\|w_0\|^2 \leq \frac{(1 + 1/n)^2}{3} = \frac{1}{3} + O(1/n), \quad \langle w_0, u \rangle \geq \frac{r}{2} \geq r - \frac{1}{2r} = a.$$

The last inequality uses $r \leq 1$. Similarly $\langle w_0, v \rangle \geq d$.

Preserve a margin in every block. Take $m = \lceil \log n \rceil$ blocks with dimensions d_i differing by at most one, and put $\lambda_i = d_i/n$. Choose a margin $\epsilon_f \rightarrow 0$ slowly enough that $\epsilon_f^2 \min_i d_i \gg \log m$. For a sufficiently large fixed C , put $w = w_0 + C\epsilon_f(u + v)$. Since $\langle u, v \rangle > 0$, each inner product increases by at least $C\epsilon_f$. The two inequalities therefore have margin $\Omega(\epsilon_f)$, while $\|w\|^2 \leq 1/3 + O(\epsilon_f)$.

After a Haar rotation, every block is a uniform subspace of its prescribed dimension. Apply Lemma B.3 to u, v, w and their nonzero pairwise sums and differences, then take a union bound over blocks. Except with probability $O(m)e^{-\Omega(\epsilon_f^2 \min_i d_i)} = o(1)$, all these projected squared norms have relative error at most ϵ_f . For $z, z' \in \{u, v, w\}$, polarization gives

$$|\langle z_i, z'_i \rangle - \lambda_i \langle z, z' \rangle| \leq \frac{\epsilon_f \lambda_i}{4} (\|z + z'\|^2 + \|z - z'\|^2) = O(\epsilon_f \lambda_i),$$

where z_i, z'_i denote the block projections. A zero sum or difference has zero error. This argument uses only the individual block marginals and does not require independence between blocks.

Write u_i, v_i, w_i for their projections to block i . A successful block direction $z_i \in S^{d_i-1}$ must satisfy

$$\langle z_i, u_i/\sqrt{\lambda_i} \rangle \geq a, \quad \langle z_i, v_i/\sqrt{\lambda_i} \rangle \geq d.$$

The vector $w_i/\sqrt{\lambda_i}$ satisfies both inequalities with margin $\Omega(\epsilon_f)$ and has squared norm at most $1/3 + O(\epsilon_f)$.

Find a suitable codeword in each block. Fix a rotation satisfying the preceding projection bounds. Every block codeword is independently uniform on S^{d_i-1} . Project to a plane containing u_i, v_i , fixing an additional direction before sampling if they are parallel. The vector w_i lies in their span. For large n , $d_i \geq 4$ and $1/d_i = o(\epsilon_f)$. A projection in the disk of radius $1/d_i$ centered at $w_i/\sqrt{\lambda_i}$ satisfies both filter constraints, and every point of this disk has squared norm at most $1/3 + O(\epsilon_f)$. The density lower bound in Lemma B.4, multiplied by the disk area, gives the conditional success probability

$$p_i \geq d_i^{-C} (2/3)^{d_i/2} \exp(-O(\epsilon_f d_i)).$$

Choose $h_n = o(n)$ with $h_n \gg n\epsilon_f + m \log n$, and take

$$t_i = \left\lceil ((3/2)^{n/2} e^{h_n})^{d_i/n} \right\rceil$$

codewords in block i . Since $d_i \asymp n/m$,

$$\log(t_i p_i) \geq (d_i/n) h_n - O(\epsilon_f d_i + \log d_i) \gg \log m.$$

Conditional on the rotation, the t_i codewords are independent, so block failure has probability at most $(1 - p_i)^{t_i} \leq e^{-t_i p_i}$. A union bound over blocks and the rotation failure event gives total failure probability $o(1)$. Concatenating one successful direction per block with weights $\sqrt{\lambda_i}$ gives global inner products at least a, d . Finally, $\prod_i t_i = (3/2)^{n/2} 2^{o(n)}$ and $\sum_i t_i = 2^{o(n)}$, as required. $\square$

B.2 Enumerating filter buckets within the space bound

Having proved coverage, we establish the cost bounds (A.3) and (A.4) for every choice of codes. Lemma A.3 and the $2^{o(n)}$ total block-code size give $2^{o(n)}$ preparation work per enumeration, followed by a polynomial multiple of the output count.

Label every block codeword by its original sampling index. A complete codeword is keyed by the tuple of these labels, and a prefix fixes the first labels in this tuple. Sorting inner products for a particular point changes only the order in which labels are visited. Enumeration restricted to a prefix fixes those labelled block codewords and enumerates the remaining blocks. Set the bucket capacity to $B_{\text{cap}} = 2M + Q$. Enumerate the data memberships for a prefix, stopping at the $(B_{\text{cap}} + 1)$ st membership and splitting the prefix if necessary. When the capacity is respected, group the data by complete codeword and process queries one at a time against the resulting dictionary. At depth m , a prefix specifies one complete codeword and has at most $2M$ data memberships, so splitting terminates.

Let b_{max} be the largest block-code size and I the number of nodes whose membership count exceeds B_{cap} . At each depth, the memberships under distinct visited prefixes are disjoint, and each such node contains more than B_{cap} of them. Consequently,

$$I \leq \frac{m J_D}{B_{\text{cap}}}, \quad \#\{\text{visited nodes}\} \leq 1 + b_{\text{max}} I.$$

Rescanning the $2M + Q = B_{\text{cap}}$ inputs at each visited node thus costs at most $2^{o(n)}(M + Q + mb_{\text{max}}J_D)$. At a fixed depth, the truncated outputs from distinct prefixes are disjoint subsets of the J_D data memberships, so their total is at most J_D . Summing over depths gives $O(mJ_D)$ outputs. Query enumeration is performed only at accepted prefixes. These prefixes are pairwise disjoint, so it reports at most J_X memberships in total. Since $b_{\text{max}} = 2^{o(n)}$, these bounds give the claimed enumeration overhead.

Charge direct comparisons of points with small norm to C_{pair} . This gives (A.3). Traverse the prefix tree depth first, generating child prefixes as needed. Store only one bounded bucket, the prefix stack, and at most K admissible signed indices for each query. This gives (A.4). Codes and queries are processed sequentially. For each query, maintain a balanced search tree of signed parent indices. It removes duplicate reports and gives the required lexicographic order by in-order traversal, preserving the multiplicity of equal vectors at different parent indices.

C Uniform initialization and resource budgets

This appendix specifies the sampling parameters and resource budgets used in Algorithm 4.2, and verifies their probability bounds.

C.1 Exact uniform sampling in the initial ball

The cell $\mathcal{P}_B = B[-1/2, 1/2]^n$ has volume one and lies in $\mathcal{B}_n(\rho_B)$. For every $x \in S_J \cup \{0\}$, the entire cell $x + \mathcal{P}_B$ lies in the proposal ball and rounds to x . Put $N_J = |S_J|$ and $V_{\text{prop}} = \text{vol } \mathcal{B}_n(R_J + \rho_B)$. Every target point has proposal probability $1/V_{\text{prop}}$, so acceptance has probability N_J/V_{prop} and the accepted point is uniform on S_J .

The cells of these target points cover $\mathcal{B}_n(R_J - \rho_B)$. Consequently,

$$\Pr\{\text{accept}\} \geq \left(\frac{R_J - \rho_B}{R_J + \rho_B}\right)^n - V_{\text{prop}}^{-1} \geq \left(\frac{1 - n^{-2}}{1 + n^{-2}}\right)^n - n^{-2n} =: \alpha_n \rightarrow 1.$$

Here $\rho_B/R_J \leq n^{-2}$ and $V_{\text{prop}}^{-1} \leq n^{-2n}$, so the bound holds for every input basis. Independent proposals give independent uniform accepted values, independently of their waiting times. Conditional on collecting M values within the fixed proposal budget, the accepted list therefore has distribution $U(S_J)^{\otimes M}$. The algorithm reports a shortage as failure. Including this outcome, the total-variation distance from the ideal product distribution is at most the shortage probability.

C.2 Choosing list sizes, capacities, and mixing times

We first bound the cap probability by computable quantities. Let $a_n = \Gamma(n/2)/(\sqrt{\pi}\Gamma((n-1)/2))$ and put

$$p_- = \frac{a_n}{n-1}(1-b^2)^{(n-1)/2}, \quad p_+ = p_-/b, \quad \theta_{\pm} = p_{\pm}/\tau^n. \quad (\text{C.1})$$

The first coordinate Z of a uniform sphere point has density $a_n(1-t^2)^{(n-3)/2}$, and $p = \mathbb{P}\{Z \geq b\}$. Insert a factor t in this integral, use $b \leq t \leq 1$, and integrate the derivative of $(1-t^2)^{(n-1)/2}$ to obtain $p_- \leq p \leq p_+$.

Recall $\xi = 1/[n(J+1)]$ and $\kappa = (1-\rho)/2$. The cap bounds above give $\theta_- \leq p/\tau^n \leq \theta_+$. Define

$$\mu_-(M) = \frac{2\kappa M \theta_-}{1+\kappa}, \quad \mu_+(M) = \frac{2(1+\kappa)M \theta_+}{1-\kappa}, \quad \sigma_n = \frac{1-\kappa}{(1+\kappa)\tau^n}.$$

These bound the second eigenvalue of $\mathbb{E}L_D$, the expected degree, and the probability that a parent sample lies in the next ball. Choose the least integer $M \geq (4/3)^{n/2}$ satisfying

$$U_J e^{-\mu_-(M)/32} \leq \xi/2, \quad e^{-M\sigma_n} \leq \xi. \quad (\text{C.2})$$

Next choose the least integer $K \geq e\mu_+(M)$ satisfying

$$U_J \left(\frac{e\mu_+(M)}{K} \right)^{K/2} \leq \xi/2, \quad g_0 = \frac{\mu_-(M)}{4K}. \quad (\text{C.3})$$

The two tail bounds come from Theorem 2.4: a sample contributes a Laplacian of norm at most four and a degree at most two. Consider a lattice whose point-count and degree relative errors are at most κ and whose Laplacian gap satisfies (3.6). The proof of Lemma 3.5 gives

$$\mu_{\min} \geq \mu_-(M), \quad \max_v \mathbb{E}d_D(v) \leq \mu_+(M).$$

The two Chernoff bounds and (C.2)–(C.3) give failure probability at most $\xi/2$ for each estimate. On success, $\lambda_2(L_D) \geq \mu_-(M)/2$ and the chain has spectral gap at least $g_0 = \mu_-(M)/(4K)$. A parent sample lies in the target ball with probability at least σ_n , so the probability of having no starting state is at most ξ .

The parameter conditions are monotone in the ranges searched and can be solved by doubling and binary search using $O(\log M + \log K)$ tests. Since $\theta_+/\theta_- = 1/b = O(1)$ and $\sigma_n = \Omega(1)$, they give

$$M = \Theta(\theta_-^{-1} \log(U_J/\xi)), \quad K = \Theta(M\theta_+), \quad g_0 = \Omega(1).$$

For the capacity upper bound, a sufficiently large fixed multiple of $\mu_+(M)$ satisfies its tail condition. All constants are independent of the lattice, basis, and stage. Allocate the remaining error bounds as follows.

- Allow at most

$$N_I = \left\lceil \frac{\max\{2M, 8 \log(1/\xi)\}}{\alpha_n} \right\rceil$$

initialization proposals. Their number of acceptances has mean $\mu \geq \alpha_n N_I \geq 2M$. The scalar Chernoff bound gives shortage probability at most $e^{-\mu/8} \leq \xi$. Thus both numerical factors in this budget come from the stated tail bound.

- At stage j , run the shared burn-in for $t_j(\xi)$ steps, then each of the Q_j output trajectories for $s_j = t_j(\xi/Q_j)$ steps. Burn-in error and the joint output mixing error are each at most ξ .
- Use fresh filter codes at each output step. Allocate $\xi/\max(1, s_j)$ to each batch and use

$$\ell_j = \lceil \log_2(2MQ_j \max(1, s_j)/\xi) \rceil$$

independent repetitions.

Split the last list by fixed indices into two batches of M points. For the final difference search, the probability of missing a pair and the probability of exhausting the time budget are each at most ξ .

C.3 Choosing filter parameters and bounding their cost

Take $m = \lceil \log n \rceil$ and partition the coordinates into m blocks whose dimensions differ by at most one, placing larger blocks first. Choose polynomial-time computable sequences ϵ_f, h_n satisfying

$$\epsilon_f \rightarrow 0, \quad \epsilon_f^2 \min_i d_i \gg \log m, \quad n\epsilon_f + m \log n \ll h_n \ll n.$$

The first two conditions control projection errors. The last provides enough codewords in each block while keeping the overhead subexponential. For example, $\epsilon_f = 1/m$ and $h_n = n/\sqrt{m}$ satisfy them. Fix any such choice, and set

$$t_i = \left\lceil ((3/2)^{n/2} e^{h_n})^{d_i/n} \right\rceil, \quad F = \prod_i t_i.$$

Lemma B.1 applies, with $F = (3/2)^{n/2} 2^{o(n)}$ and $\sum_i t_i = 2^{o(n)}$.

We record the operation counts to specify the budget. Put $b_{\max} = \max_i t_i$ and

$$\begin{aligned} L_* &= 1 + \lceil \log_2((2M + Q + 1)(F + 1)) \rceil, \\ A_{\text{prep}} &= n^2 + \sum_i d_i t_i + \sum_i t_i \lceil \log_2(t_i + 1) \rceil, \\ P_T &= n^3 + (1 + mb_{\max})A_{\text{prep}} + m^2 + mL_* + n + L_*. \end{aligned}$$

The n^3 term pays for generating a Haar rotation by Gaussian QR decomposition. The three terms of A_{prep} pay for applying it, block inner products, and sorting. The streaming proof bounds the number of scanned prefixes by $1 + mb_{\max}J_D/(2M + Q)$. Truncated enumeration costs $O(m^2 J_D)$; sorting and binary search on codeword keys use $O(mL_*)$ operations per membership, exact vector tests use $O(n)$ per candidate, and a balanced search tree deduplicates signed parent indices in $O(L_*)$ operations per report. Thus one code costs at most a fixed constant times $P_T(M + Q + J_D + J_X + C_{\text{pair}})$.

Choose a polynomial-time computable sequence $A_n = n^{O(1)}$ with $n \log n = o(A_n)$, for example $A_n = n^2$. Equations (A.8)–(A.9) bound all required geometric coefficients by A_n outside a set of Haar probability $o(1)$. The covering-radius bound handles all larger radii.

Replace V^{-1} in (A.7) by c^{-n} , and define

$$\begin{aligned} \mathcal{W}(M, Q) &= M + Q + MQ(3/4)^{n/2} \\ &\quad + F[(M + Q)(3/4)^{n/2} + MQ(3/4)^n + MQ(3/4)^{n/2}/c^n], \\ \mathcal{H}(M, Q, \ell) &= C_B \ell P_T A_n \mathcal{W}(M, Q). \end{aligned}$$

Here C_B is an absolute constant in the operation and geometric mean bounds. On a set of lattices of Haar probability $1 - o(1)$ independent of the input basis, $\mathcal{H}$ bounds the expected batch cost of the stationary chain.

C.4 Converting expected work to time and space budgets

The total time budget for the output queries of stage j is

$$B_j = \lceil s_j \mathcal{H}(M, Q_j, \ell_j) / \xi \rceil. \quad (\text{C.4})$$

Charge burn-in separately by its fixed number of full scans. For the final difference search, take $Q = M$, $\ell = \lceil \log_2(2M^2/\xi) \rceil$, and time budget $\lceil \mathcal{H}(M, M, \ell)/\xi \rceil$. Generate and process the codes one at a time. A batch uses at most

$$C_S \left[(n + m + 1)(2M + Q) + QK + n^2 + \sum_i (d_i + 1)t_i + m^2 \right]$$

arithmetic storage cells, for an absolute constant C_S . The terms store the input vectors and one bucket with its codeword keys, proposal tables, a rotation, the block codes and their scores, and the prefix stack.

Remark 4.5 transfers the reference cost bound to the actual output trajectories by Markov's inequality and the burn-in coupling. For the final difference search, use the coupling of the endpoint lists to independent uniform samples instead. All budgets are computed in advance, and exceeding one causes failure. At every stage, failure of the parent-list conditions, absence of a starting state, burn-in error, joint output mixing error, filter omission, and budget exhaustion each contribute at most a constant multiple of ξ . Including initialization and filtering for the final difference search, the total error is $O((J + 1)\xi) = O(1/n)$. The coupling charges the actual parent list's joint error only once per stage. Failure to find a shortest vector is bounded separately by the second-moment argument in the main text.

All extra budget factors are $2^{o(n)} \text{poly}(n, J, \Lambda)$, giving the bounds in Theorem 4.3. Remark 4.4 states when these yield the claimed dimension exponents.

References

- [ADRS15] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling: Extended abstract. In *Proceedings of STOC 2015*, pp. 733–742, 2015. doi:[10.1145/2746539.2746606](https://doi.org/10.1145/2746539.2746606).
- [AKS01] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of STOC 2001*, pp. 601–610, 2001. doi:[10.1145/380752.380857](https://doi.org/10.1145/380752.380857).
- [BDGL16] Anja Becker, Léo Ducas, Nicolas Gama, and Thijs Laarhoven. New directions in nearest neighbor searching with applications to lattice sieving. In *Proceedings of SODA 2016*, pp. 10–24, 2016. doi:[10.1137/1.9781611974331.ch2](https://doi.org/10.1137/1.9781611974331.ch2). Published full text.
- [CKW20] Aruni Choudhary, Siargey Kachanovich, and Mathijs Wintraecken. Coxeter triangulations have good quality. *Mathematics in Computer Science*, 14:141–176, 2020. doi:[10.1007/s11786-020-00461-5](https://doi.org/10.1007/s11786-020-00461-5).
- [DG03] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003. doi:[10.1002/rsa.10073](https://doi.org/10.1002/rsa.10073). Author version.
- [DL26] Emmanouil Doulgerakis and Thijs Laarhoven. Smooth Sailing through Spherical Shells: Provable Random-Lattice Sieving in Time $2^{0.292n}$. Preprint, 2026. [arXiv:2609.37482v1](https://arxiv.org/abs/2609.37482).
- [GFH26a] Yiming Gao, Yansong Feng, and Honggang Hu. Solving the Shortest Vector Problem in $2^{0.7314n+o(n)}$ Time via Discrete Gaussian Sampling on Superlattices. Cryptology ePrint Archive, Paper 2026/1587, 2026. [ePrint 2026/1587](https://eprint.iacr.org/2026/1587).
- [GFH26b] Yiming Gao, Yansong Feng, and Honggang Hu. Discrete Gaussian Sampling Meets BDGL Decoding: Solving the Shortest Vector Problem in $2^{0.5596n+o(n)}$ Time. Cryptology ePrint Archive, Paper 2026/1844, 2026. [ePrint 2026/1844](https://eprint.iacr.org/2026/1844).
- [Hha26a] Minki Hhan. Solving the Shortest Vector Problem in $2^{0.6039n}$ Time via Mid-point Hessian. Preprint, 2026. [arXiv:2608.02478](https://arxiv.org/abs/2608.02478).
- [Hha26b] Minki Hhan. Finding a Shortest Vector and More in $2^{n/2+o(n)}$ Time using q -ary Coset Difference Tree. Preprint, 2026. [arXiv:2609.02764](https://arxiv.org/abs/2609.02764).
- [Hol22] Kristian Holm. On the distribution of angles between increasingly many short lattice vectors. *Journal of Number Theory*, 240:357–403, 2022. doi:[10.1016/j.jnt.2022.02.001](https://doi.org/10.1016/j.jnt.2022.02.001).
- [JK26] Jiseung Kim. One Discrete Gaussian Sample in $2^{n/2+o(n)}$ Time. Cryptology ePrint Archive, Paper 2026/1599, 2026. [ePrint 2026/1599](https://eprint.iacr.org/2026/1599).
- [Kim20] Seungki Kim. Random lattice vectors in a set of size $O(n)$. *International Mathematics Research Notices*, 2020(5):1385–1416, 2020. doi:[10.1093/imrn/rny060](https://doi.org/10.1093/imrn/rny060).
- [Kim22] Seungki Kim. Mean value formulas on sublattices and flags of the random lattice. *Journal of Number Theory*, 241:330–351, 2022. doi:[10.1016/j.jnt.2022.03.013](https://doi.org/10.1016/j.jnt.2022.03.013). Author version.
- [Kim24] Seungki Kim. Adelic Rogers integral formula. *Journal of the London Mathematical Society*, 109(1):e12830, 2024. doi:[10.1112/jlms.12830](https://doi.org/10.1112/jlms.12830).

- [Laa16] Thijs Laarhoven. *Search problems in cryptography: from fingerprinting to lattice sieving*. PhD thesis, Technische Universiteit Eindhoven, 2016. [Full text](#).
- [LMSY23] Siqi Liu, Sidhanth Mohanty, Tselil Schramm, and Elizabeth Yang. Local and global expansion in random geometric graphs. In *Proceedings of STOC 2023*, pp. 817–825, 2023. [doi:10.1145/3564246.3585106](#). [Full version](#).
- [LPW17] David A. Levin and Yuval Peres, with contributions by Elizabeth L. Wilmer. *Markov Chains and Mixing Times*, second edition. American Mathematical Society, 2017. [Full text](#).
- [MV10] Daniele Micciancio and Panagiotis Voulgaris. Faster exponential time algorithms for the shortest vector problem. In *Proceedings of SODA 2010*, pp. 1468–1480, 2010. [doi:10.1137/1.9781611973075.119](#).
- [NV08] Phong Q. Nguyen and Thomas Vidick. Sieve algorithms for the shortest vector problem are practical. *Journal of Mathematical Cryptology*, 2(2):181–207, 2008. [doi:10.1515/JMC.2008.009](#).
- [ORW22] Or Ordentlich, Oded Regev, and Barak Weiss. New bounds on the density of lattice coverings. *Journal of the American Mathematical Society*, 35(1):295–308, 2022. [doi:10.1090/jams/984](#). [Author version](#).
- [PS26] Amaury Pouly and Yixin Shen. Solving the Shortest Vector Problem in $2^{0.63269n+o(n)}$ Time on Random Lattices. In *EUROCRYPT 2026*, LNCS 16544, pp. 92–123, 2026. [doi:10.1007/978-3-032-25327-9_4](#). [Author full version](#).
- [PS09] Xavier Pujol and Damien Stehlé. Solving the shortest lattice vector problem in time $2^{2.465n}$. Cryptology ePrint Archive, Paper 2009/605, 2009. [ePrint 2009/605](#).
- [Rog55] C. A. Rogers. Mean values over the space of lattices. *Acta Mathematica*, 94:249–287, 1955. [doi:10.1007/BF02392493](#).
- [Sie45] Carl Ludwig Siegel. A mean value theorem in geometry of numbers. *Annals of Mathematics*, 46(2):340–347, 1945. [doi:10.2307/1969027](#).
- [Sod11a] Anders Södergren. On the Poisson distribution of lengths of lattice vectors in a random lattice. *Mathematische Zeitschrift*, 269(3–4):945–954, 2011. [doi:10.1007/s00209-010-0772-8](#).
- [Sod11b] Anders Södergren. On the distribution of angles between the N shortest vectors in a random lattice. *Journal of the London Mathematical Society*, 84(3):749–764, 2011. [doi:10.1112/jlms/jdr032](#).
- [Str11] Andreas Strömbergsson. On the probability of a random lattice avoiding a large convex set. *Proceedings of the London Mathematical Society*, 103(6):950–1006, 2011. [doi:10.1112/plms/pdr021](#). [Author version](#).
- [SS19] Andreas Strömbergsson and Anders Södergren. On the generalized circle problem for a random lattice in large dimension. *Advances in Mathematics*, 345:1042–1074, 2019. [doi:10.1016/j.aim.2019.01.034](#). [Author version](#).

- [SS22] Andreas Strömbergsson and Anders Södergren. On a mean value formula for multiple sums over a lattice and its dual. Preprint, 2022. [arXiv:2211.05454](#). We use its equivalent integer-coordinate form of the classical Rogers formula.
- [Tro12] Joel A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12:389–434, 2012. [doi:10.1007/s10208-011-9099-z](#). [Author version](#).
- [Wan26] Daqing Wan. NP-hardness of SVP in Euclidean Space. arXiv:2603.27398, version 3, 2026. [arXiv:2603.27398v3](#).