Solving the Shortest Vector Problem in $2^{0.7314n+o(n)}$ Time via Discrete Gaussian Sampling on Superlattices

Yiming Gao[1] Yansong Feng[2] Honggang Hu[1,3]

[1]School of Cyber Science and Technology, University of Science and Technology of China
qw1234567@mail.ustc.edu.cn

[2]Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China
fengyansong@amss.ac.cn

[3]Hefei National Laboratory, Hefei, China
hghu2005@ustc.edu.cn

Abstract

We give a classical randomized algorithm for the Euclidean Shortest Vector Problem (SVP) on arbitrary full-rank lattices. It runs in $2^{0.7314n+o(n)}$ time and $2^{n/2+o(n)}$ space. More precisely, it returns a shortest vector with high probability and runs in

$$2^{E_0 n+o(n)} \quad \text{time and} \quad 2^{n/2+o(n)} \quad \text{space}, \qquad E_0 = 0.73133754\ldots.$$

This improves the $2^{n+o(n)}$ classical bound of Aggarwal, Dadush, Regev, and Stephens-Davidowitz (ADRS, STOC 2015). It also beats the previous best known worst-case quantum bounds of Aggarwal, Chen, Kumar, and Shen (ACKS, SIAM J.Comp.2025), namely $2^{0.9497n+o(n)}$ without QRAM and $2^{0.8345n+o(n)}$ with QRAM.

The algorithm constructs a random prime-index superlattice $\Gamma \supset \mathcal{L}$ and applies the honest discrete Gaussian sampler of ADRS to Γ with a parameter above the smoothing threshold. Then it simply scans the resulting samples and retains the shortest nonzero one that lies in $\mathcal{L}$.

The analysis passes to the dual lattice $M = \Gamma^*$. A random linear constraint reduces the expected contribution of vectors outside $p\mathcal{L}^*$ by a factor smaller than $1/p$. A dual minimum bound, obtained by applying Poisson summation to a shortest dual line, controls the forced Gaussian mass on $p\mathcal{L}^*$. Consequently, Γ is smooth at the sampling scale and a fixed shortest vector of $\mathcal{L}$ is hit with probability at least $2^{-E_0 n-o(n)}$ per ideal sample.

1 Introduction

Let $\mathcal{L} \subset \mathbb{R}^n$ be a full-rank lattice. The Euclidean Shortest Vector Problem asks for a nonzero vector $v \in \mathcal{L}$ satisfying

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

Ajtai, Kumar, and Sivakumar initiated the provable sieving approach with a randomized $2^{O(n)}$-time algorithm [AKS01]. Subsequent work developed both practical and rigorous sieve variants [NV08, PS09], while Micciancio and Voulgaris obtained a deterministic $2^{2n+o(n)}$-time algorithm based on Voronoi cells [MV13]. Aggarwal, Dadush, Regev, and Stephens-Davidowitz [ADRSD15] later gave a randomized $2^{n+o(n)}$-time and space algorithm by reducing SVP to centered discrete Gaussian sampling. This remains the previous best worst-case classical running time bound without a space restriction.

A second result of ADRS is a substantially faster sampler at parameters above the smoothing parameter. Given an n-dimensional lattice $\mathcal{L}$ and a parameter

$$u > \sqrt{2}\,\eta_{1/2}(\mathcal{L}),$$

their honest discrete Gaussian sampler produces a distribution statistically close to $m = \lceil 2^{n/2} \rceil$ independent samples from $D_{\mathcal{L},u}$ in $2^{n/2+o(n)}$ time and space. This fast sampler has led to improved space-efficient and quantum algorithms [CCL18, ACKS25]. With $2^{n/2+o(n)}$ classical space, the previous classical running time is $2^{1.669n+o(n)}$, while the best known worst-case quantum bounds are $2^{0.9497n+o(n)}$ without QRAM and $2^{0.8345n+o(n)}$ with QRAM [ACKS25].

The difficulty is that SVP requires useful Gaussian mass on a shortest vector, and the corresponding parameter can lie below $\eta_{1/2}(\mathcal{L})$. Pouly and Shen exploit the fact that almost all Haar–Siegel random lattices have unusually favorable smoothing behavior, obtaining a $2^{0.63269n+o(n)}$ algorithm in that average case setting [PS26]. Here we take a different route: instead of assuming that the input lattice is typical, we replace an arbitrary input lattice by a random superlattice whose smoothing behavior we can control.

With $\beta = 2^{0.401}$, the leading exponents admit a direct comparison:

$$E_{\mathrm{PS26}} = \frac{1}{2} + \frac{1}{4e \ln 2} = 0.63269\ldots, \qquad E_0 = \frac{1}{2} + \frac{\beta^2}{4e \ln 2}.$$

Thus the two analyses share the same direct hit calculation at the ADRS sampling threshold, while the factor β^2 records the worst-case geometric loss in controlling Gaussian mass. The random superlattice construction is what makes this calculation available for an arbitrary input lattice.

The parameters must balance sampling and smoothing. A smaller sampling parameter gives a shortest vector more Gaussian weight, but honest DGS requires the parameter to remain above smoothing. Increasing the superlattice index p improves smoothing, but it also enlarges the Gaussian normalizer and lowers the hit probability. We therefore choose p just large enough to make the superlattice smooth at the useful scale.

More precisely, we construct a prime-index superlattice $\Gamma \supset \mathcal{L}$. The construction starts from a random subspace of dimension $n-1$ in the dual quotient $\mathcal{L}^*/p\mathcal{L}^*$. A fixed, structured superlattice would generally preserve too much of the short vector structure of $\mathcal{L}^*$; the random constraint is what suppresses this structure. We also avoid rejection sampling an entire discrete Gaussian table on $\mathcal{L}$: because the algorithm only needs one shortest vector, it is enough to scan the raw samples from Γ and test membership in $\mathcal{L}$.

This construction has two useful features. First, every vector of $\mathcal{L}$, and in particular every shortest vector, belongs to every possible Γ. Thus the random choice never removes the target vector. Second, the dual $M = \Gamma^*$ behaves like a random index-p sublattice of $\mathcal{L}^*$, so most nonzero dual vectors survive with probability approximately $1/p$. The only vectors that survive every choice are those in $p\mathcal{L}^*$. The main geometric step is therefore to control these forced vectors. We do so by applying one-dimensional Poisson summation to a shortest dual line, which converts the primal Gaussian mass bound into the needed lower bound on $\lambda_1(\mathcal{L}^*)$.

As in [ADRSD15, Section 2.4], all displayed resource bounds record their dependence on the dimension n; factors polynomial in the input encoding length are suppressed.

1.1 Main result

Write

$$\beta := 2^{0.401}, \qquad D := \frac{\beta^2}{e \ln 2} = 0.92535018\ldots, \qquad E_0 := \frac{1}{2} + \frac{D}{4} = 0.73133754\ldots.$$

Theorem 1.1 (Main theorem)**.** *There exist a classical randomized algorithm and an absolute constant n_0 such that, for every $n \geq n_0$, given a rational basis of a full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$, the algorithm returns a shortest nonzero vector of $\mathcal{L}$ with probability at least $1 - 2^{-\Omega(n \log n)}$. Under the dimension-only arithmetic convention above, its running time is*

$$2^{E_0 n + o(n)},$$

and its space usage is

$$2^{n/2 + o(n)}.$$

For the quantitative proof, let

$$a_n := \left\lceil \frac{n}{\sqrt{\log_2 n}} \right\rceil.$$

We prove the more precise running time bound

$$2^{E_0 n + O(n/\sqrt{\log n})}.$$

Since $a_n = o(n)$ and $E_0 < 0.7314$, this implies the stated $2^{0.7314n + o(n)}$ bound.

Remark 1.2. Suppose β can be replaced by some $\widetilde{\beta} < \beta$ in the uniform bound on the number of lattice points given by Lemma A.1. The same shell argument then gives the corresponding Gaussian mass bound, and the rest of the analysis gives

$$\widetilde{E} = \frac{1}{2} + \frac{\widetilde{\beta}^2}{4e \ln 2} < E_0.$$

A very recent manuscript of OpenAI reports asymptotic bounds for spherical codes that strictly improve the optimized Kabatiansky–Levenshtein exponent at every fixed angle [Ope26, Chapter 2, Corollary 1.3]. These bounds suggest that the uniform estimate for the number of lattice points, and hence the final exponent, could be improved. Replacing β by the resulting smaller value $\widetilde{\beta}$ would lower the leading exponent from 0.7313 to about 0.7296. Because the manuscript is very recent and currently under peer review, we use the classical KL bound in our stated result and do not include this numerical improvement in the theorem.

1.2 Proof overview

The proof has three main ingredients: choosing a scale at which the input lattice has small Gaussian mass, constructing a random superlattice that is smooth at that scale, and showing that a sample from the superlattice directly hits a shortest vector of $\mathcal{L}$ with the claimed probability.

The useful scale. Suppose first that $\lambda_1(\mathcal{L})$ is known, and define

$$s_0^2 := \frac{2\pi e}{\beta^2 n} \lambda_1(\mathcal{L})^2.$$

We choose $s = (1 + O(1/n))\, s_0$ and

$$t = (1 + O(1/n))\sqrt{2}\, s.$$

The value s_0 is the largest scale at which the quantitative Kabatiansky–Levenshtein Gaussian mass bound still guarantees

$$\rho_s(\mathcal{L}) \leq 2^{O(\log n)}.$$

The factor $\sqrt{2}$ in t is dictated by the threshold of the fast honest DGS algorithm.

Smoothing a random superlattice. Set $r = 1/s$. Our goal is to prove

$$\rho_r(M \setminus \{0\}) < \frac{1}{2},$$

because this implies $\eta_{1/2}(\Gamma) < s$. Let $\Lambda = \mathcal{L}^*$ and choose a prime

$$p = \Theta\left(2^{a_n} \frac{\det \mathcal{L}}{s^n}\right).$$

Select a uniformly random nonzero functional $\varphi : \Lambda/p\Lambda \to \mathbb{F}_p$ and set

$$M := \ker \varphi, \qquad \Gamma := M^*.$$

Then $\mathcal{L} \subset \Gamma$, $[\Gamma : \mathcal{L}] = p$, and $\det \Gamma = \det \mathcal{L}/p$. In particular, $\det \Gamma = \Theta(s^n/2^{a_n})$ and $s^n/\det \Gamma = \Theta(2^{a_n})$. This places Γ near the critical density at scale s: a Gaussian of parameter s has volume s^n, so $\det \Gamma \asymp s^n$ is the natural determinant scale. We make Γ denser by the subexponential factor 2^{a_n}, which supplies enough slack for smoothing while contributing only $o(n)$ to the final exponent.

The expected nonzero Gaussian mass of M is

$$\mathbb{E}_\varphi[\rho_r(M \setminus \{0\})] = \rho_r(p\Lambda) - 1 + \frac{p^{n-1} - 1}{p^n - 1}(\rho_r(\Lambda) - \rho_r(p\Lambda)).$$

The second term is the nonforced contribution. Poisson summation gives

$$\rho_r(\Lambda) = \frac{\det \mathcal{L}}{s^n}\rho_s(\mathcal{L}),$$

so the choice of p bounds this contribution by $2^{-a_n + O(\log n)}$.

The first term is forced because $p\Lambda$ lies in every kernel. The dual minimum bound in Lemma 4.2 gives

$$\frac{\det \mathcal{L}\, \lambda_1(\mathcal{L}^*)}{s^{n-1}} \geq \frac{1}{\rho_s(\mathcal{L})} \geq 2^{-O(\log n)}.$$

Consequently,

$$ps\lambda_1(\mathcal{L}^*) \geq 2^{a_n - O(\log n)},$$

which makes $\rho_r(p\mathcal{L}^* \setminus \{0\})$ superexponentially small. Markov's inequality therefore shows that, with probability $1 - 2^{-a_n + O(\log n)}$,

$$\rho_r(M \setminus \{0\}) < \frac{1}{2}.$$

Since $\Gamma^* = M$, this implies $\eta_{1/2}(\Gamma) < s$ and hence

$$t > \sqrt{2}\, \eta_{1/2}(\Gamma).$$

Directly hitting a shortest vector. On such a good kernel, Poisson summation gives

$$\rho_t(\Gamma) = \frac{pt^n}{\det \mathcal{L}}\rho_{1/t}(M) \leq 2^{n/2 + a_n + O(1)}.$$

Fix any shortest vector $v \in \mathcal{L}$. Since $\mathcal{L} \subset \Gamma$, the vector v occurs in $D_{\Gamma,t}$ with probability

$$\Pr_{X \sim D_{\Gamma,t}}[X = v] = \frac{\exp(-\pi\lambda_1(\mathcal{L})^2/t^2)}{\rho_t(\Gamma)}.$$

The scale choice gives

$$\exp(-\pi\lambda_1(\mathcal{L})^2/t^2) \geq 2^{-(D/4)n},$$

and therefore

$$\Pr[X = v] \geq 2^{-(1/2+D/4)n-a_n-O(1)} = 2^{-E_0 n-a_n-O(1)}.$$

The exponent has the transparent decomposition

$$E_0 = \underbrace{\frac{1}{2}}_{(t/s)^n} + \underbrace{\frac{D}{4}}_{\text{Gaussian weight of } v}.$$

Finally, a polynomial-size grid based on LLL supplies a branch with $s = (1 + O(1/n))s_0$ without knowing $\lambda_1(\mathcal{L})$. On that branch, $\lceil n^2 2^{E_0 n+a_n} \rceil$ ideal samples hit v except with probability $\exp(-\Omega(n^2))$. Independent repetitions of the random kernel amplify the smoothing success probability to $1 - 2^{-\Omega(n \log n)}$. The actual sampler outputs are statistically close to the ideal independent samples, and the algorithm keeps only the shortest nonzero sample that passes the exact membership test for $\mathcal{L}$.

1.3 Concurrent work

Hhan [Hha26] gives randomized algorithms for SVP using $2^{n/2+o(n)}$ space. Their expected running times are $2^{0.603867n+o(n)}$ and $2^{0.541051n+o(n)}$ in the classical and quantum settings, respectively. The method uses random affine lattice cosets and importance sampling to recover a shortest vector from the Hessian of the periodic Gaussian evaluated at a midpoint.

Kim [Kim26] gives an algorithm that draws one discrete Gaussian sample at an arbitrary parameter in expected $2^{n/2+o(n)}$ time. This yields an algorithm for SVP in $2^{0.7315n+o(n)}$ time and $2^{n/2+o(n)}$ space. Both methods sample from a random superlattice above its smoothing parameter and retain samples in the input lattice. Kim uses a quotient whose order is a power of two, while our method uses an index-p kernel chosen for a candidate value of $\lambda_1(\mathcal{L})$ and controls the mass on $p\mathcal{L}^*$ separately.

2 Preliminaries

All logarithms are natural unless a base is shown. We write $B_2^d(R)$ for the closed Euclidean ball of radius R in $\mathbb{R}^d$ and d_{TV} for total variation distance.

For probability distributions P, Q on a common countable space Ω, define their statistical (equivalently, total-variation) distance by

$$d_{\mathrm{TV}}(P, Q) := \frac{1}{2} \sum_{\omega \in \Omega} |P(\omega) - Q(\omega)|.$$

2.1 Lattices

A full-rank lattice has the form $\mathcal{L} = B\mathbb{Z}^n$ for a nonsingular $B \in \mathbb{R}^{n\times n}$, where B is called a basis of $\mathcal{L}$. Its determinant is $\det \mathcal{L} := |\det B|$. More generally, for any lattice $\mathcal{L}$, $\det \mathcal{L}$ denotes its covolume in $\mathrm{span}(\mathcal{L})$. We denote by $\lambda_1(\mathcal{L})$ the first minimum of $\mathcal{L}$, defined as the length of a shortest nonzero vector of $\mathcal{L}$:

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

More generally, if $\mathcal{L} \subseteq \mathbb{R}^m$ is a rank-n lattice, then its dual lattice is

$$\mathcal{L}^* := \{y \in \operatorname{span}(\mathcal{L}) : \langle x, y\rangle \in \mathbb{Z} \text{ for every } x \in \mathcal{L}\}.$$

For a full-rank lattice $\mathcal{L} = B\mathbb{Z}^n$, this agrees with $\mathcal{L}^* = B^{-T}\mathbb{Z}^n$.

Similarly, for a lattice basis $B = (b_1, \ldots, b_n)$, the dual basis

$$B^* = (b_1^*, \ldots, b_n^*)$$

is the unique basis satisfying

$$\langle b_i^*, b_j\rangle = \begin{cases} 1, & i = j, \\ 0, & i \neq j. \end{cases}$$

For a full-rank lattice, we have

$$B^* = B^{-T}, \qquad \det \mathcal{L}^* = \frac{1}{\det \mathcal{L}}.$$

2.2 Discrete Gaussian distribution

We use the Gaussian convention

$$\rho_s(x) := \exp\left(-\pi \frac{\|x\|^2}{s^2}\right), \qquad \rho_s(S) := \sum_{x \in S} \rho_s(x),$$

and define the discrete Gaussian distribution by

$$D_{\mathcal{L},s}(x) := \frac{\rho_s(x)}{\rho_s(\mathcal{L})} \qquad (x \in \mathcal{L}).$$

Poisson summation gives the identity

$$\rho_s(\mathcal{L}) = \frac{s^n}{\det \mathcal{L}} \rho_{1/s}(\mathcal{L}^*). \tag{1}$$

Definition 2.1 (Smoothing parameter)**.** For $\epsilon > 0$, the smoothing parameter $\eta_\epsilon(\mathcal{L})$ is the unique positive value satisfying

$$\rho_{1/\eta_\epsilon(\mathcal{L})}(\mathcal{L}^* \setminus \{0\}) = \epsilon.$$

The map $u \mapsto \rho_u(\mathcal{L} \setminus \{0\})$ is continuous and strictly increasing. We will use this elementary observation to preserve a strict inequality in the sampling threshold.

We state quantitative forms of the published estimates used below. Their constants and dimension thresholds are absolute. The Gaussian mass estimate is uniform in both the lattice and its parameter. We prove this uniformity in Appendix A.

Theorem 2.2 (Uniform quantitative Gaussian mass bound)**.** *There is an absolute constant $C_{\mathrm{G}} > 0$ with the following property. For every rank-d lattice $\mathcal{L}$ in a Euclidean space and every $u > 0$,*

$$\rho_u(\mathcal{L}) \le 1 + \left(\frac{\beta^2 u^2 d}{2\pi e\, \lambda_1(\mathcal{L})^2}\right)^{d/2+1} (d+1)^{C_{\mathrm{G}}}. \tag{2}$$

Moreover, if $u \le \sqrt{2\pi/d}\, \lambda_1(\mathcal{L})$, then

$$\rho_u(\mathcal{L}) \le 1 + \exp\left(-\pi \frac{\lambda_1(\mathcal{L})^2}{u^2}\right) \beta^d (d+1)^{C_{\mathrm{G}}}. \tag{3}$$

This is a quantitative form of Aggarwal et al. [ADRSD15, Lemma 4.2]. The cited proof has a small quantifier gap: the estimate for the number of lattice points from [PS09, Lemma 3] is stated for a fixed normalized radius, whereas the shell radii in the application vary with the dimension. Appendix A proves this uniformity by applying the Kabatiansky–Levenshtein estimate only on a fixed compact interval and using an elementary packing bound for the remaining Gaussian tail; Gaussian scaling then handles large parameters.

Theorem 2.3 (Honest DGS)**.** *There is an algorithm that, on input an n-dimensional lattice $\mathcal{L}$, a parameter $u > 0$, and a confidence parameter $\kappa = \Omega(n)$, runs in*

$$2^{n/2+\mathrm{polylog}(\kappa)+o(n)}$$

time and space. The joint statistical distance between its output distribution and an ideal sequence of independent samples from $D_{\mathcal{L},u}$ is at most $\exp(-\Omega(\kappa))$. If

$$u > \sqrt{2}\,\eta_{1/2}(\mathcal{L}),$$

the ideal sequence has length $m = \lceil 2^{n/2} \rceil$. Moreover, for the choice $\kappa = n^2$ used below, the running time and space are

$$2^{n/2+O(n/\log n+\log^5 n)}.$$

This is the honest DGS guarantee of Aggarwal et al. [ADRSD15, Theorem 5.11]. For the stated parameters, the proof uses $a = n/2 + O(n/\log n)$, $\ell = O(\log^4 n)$, and $r = O(n/\log n)$, with $M = (C\kappa^4)^{\ell+1}2^a$. Substituting $\kappa = n^2$ gives the stated bound. Below the threshold, the sampler may output fewer samples or no samples.

3 Random Prime-Index Superlattices

We first describe the random lattice construction used by the algorithm. For every choice of the random functional, it produces a prime-index superlattice of $\mathcal{L}$. An exact first moment identity describes the effect of the randomness on its dual Gaussian mass.

3.1 The construction and its geometry

Let $\Lambda := \mathcal{L}^*$. Fix a prime p and identify

$$V := \Lambda/p\Lambda \cong \mathbb{F}_p^n.$$

Choose a nonzero functional $\varphi \in V^* \setminus \{0\}$ uniformly at random, and define

$$M := \{x \in \Lambda : \varphi(x \bmod p\Lambda) = 0\}, \qquad \Gamma := M^*. \tag{4}$$

Lemma 3.1 (Index and determinant)**.** *The construction in Equation (4) satisfies*

$$p\Lambda \subset M \subset \Lambda, \qquad [\Lambda : M] = p,$$

and

$$\mathcal{L} \subset \Gamma, \qquad [\Gamma : \mathcal{L}] = p, \qquad \det\Gamma = \frac{\det\mathcal{L}}{p}.$$

Proof. The kernel of a nonzero functional on $\mathbb{F}_p^n$ has dimension $n-1$, so $[V : \ker\varphi] = p$. Pulling the kernel back to Λ proves the first line. Duality reverses inclusion, so $\Gamma = M^* \supset \Lambda^* = \mathcal{L}$. Finally,

$$\det M = p\det\Lambda = \frac{p}{\det\mathcal{L}}, \qquad \det\Gamma = (\det M)^{-1} = \frac{\det\mathcal{L}}{p}.$$

The index formula follows by taking the ratio of determinants. □

3.2 Gaussian mass of a random kernel

Lemma 3.2 (Expected mass of a random kernel)**.** *For every $r > 0$,*

$$\mathbb{E}_\varphi[\rho_r(M \setminus \{0\})] = \rho_r(p\Lambda) - 1 + \frac{p^{n-1}-1}{p^n-1}(\rho_r(\Lambda) - \rho_r(p\Lambda)). \tag{5}$$

Proof. Every vector in $p\Lambda$ belongs to M for every φ. These vectors contribute the forced mass $\rho_r(p\Lambda) - 1$, where the subtraction removes the origin.

Now fix $x \in \Lambda \setminus p\Lambda$. Its residue $\bar{x} \in V$ is nonzero. There are $p^n - 1$ nonzero functionals on V. Exactly $p^{n-1} - 1$ of them are both nonzero and vanish on $\bar{x}$: all functionals on the quotient $V/\operatorname{span}(\bar{x})$ except the zero functional. Hence

$$\Pr_\varphi[x \in M] = \frac{p^{n-1}-1}{p^n-1}.$$

Multiplying by $\rho_r(x)$ and summing over $x \in \Lambda \setminus p\Lambda$ proves Equation (5). □

Remark 3.3. The coefficient of the nonforced mass is strictly smaller than $1/p$. The forced term cannot be reduced by choosing a different random functional; it must be controlled geometrically through the choice of p.

Thus the smoothing analysis reduces to two separate tasks: reducing the nonforced mass by the factor $1/p$, and making the forced mass on $p\mathcal{L}^*$ negligible. We next choose the scale and prime to meet both requirements.

4 Choosing the Scale and Prime

The algorithm does not know $\lambda_1(\mathcal{L})$, so it searches a grid of polynomial size. We show that one branch has the required primal Gaussian mass and then choose a prime that controls both terms in the first moment identity of Lemma 3.2 at the same time.

4.1 The scale grid

An LLL reduced basis gives a value $\widehat{\lambda}$ satisfying

$$\lambda_1(\mathcal{L}) \le \widehat{\lambda} \le 2^{(n-1)/2}\lambda_1(\mathcal{L})$$

in polynomial time [LLL82]. Define

$$\widehat{\lambda}_j := \frac{\widehat{\lambda}}{(1+1/n)^j}, \qquad 0 \le j \le J := \left\lceil \log_{1+1/n}(2^{(n-1)/2}) \right\rceil + 1.$$

Since $J = O(n^2)$, all branches contribute only a polynomial overhead. At least one index j_* satisfies

$$\lambda_1(\mathcal{L}) \le \widehat{\lambda}_{j_*} < (1+1/n)\lambda_1(\mathcal{L}). \tag{6}$$

For every branch, let

$$s_j^\circ := \sqrt{\frac{2\pi e}{\beta^2 n}}\,\widehat{\lambda}_j$$

and use directed rational approximations satisfying

$$s_j^\circ \le s_j \le (1+n^{-3})s_j^\circ, \qquad (1+1/n)\sqrt{2}\, s_j \le t_j \le (1+n^{-3})(1+1/n)\sqrt{2}\, s_j. \tag{7}$$

On the correct branch, write $s := s_{j_*}$, $t := t_{j_*}$, and $r := 1/s$. If

$$s_0 := \sqrt{\frac{2\pi e}{\beta^2 n}}\,\lambda_1(\mathcal{L}),$$

then

$$s_0 \le s < (1+1/n)(1+n^{-3})s_0 = (1+O(1/n))s_0. \tag{8}$$

Lemma 4.1 (Primal Gaussian mass)**.** *On the correct branch,*

$$\rho_s(\mathcal{L}) \le 2^{O(\log n)}.$$

Proof. Apply Equation (2) to the lattice $\mathcal{L}$ with $d = n$, and use Equation (8). The expression inside parentheses is

$$\frac{\beta^2 s^2 n}{2\pi e\,\lambda_1(\mathcal{L})^2} \le (1+O(1/n))^2.$$

Raising this to $n/2+1$ contributes only $e^{O(1)}$. The factor $(n+1)^{C_{\mathrm{G}}}$ and the additive 1 give the claimed bound. □

4.2 A dual minimum bound

The forced mass depends on $\lambda_1(\mathcal{L}^*)$. The next bound applies Poisson summation to the line generated by a shortest dual vector. It converts the primal Gaussian mass into the dual estimate needed to choose the prime.

Lemma 4.2 (Dual minimum from primal Gaussian mass)**.** *For every full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$ and every $s > 0$,*

$$\frac{\det \mathcal{L}\,\lambda_1(\mathcal{L}^*)}{s^{n-1}} \ge \frac{1}{\rho_s(\mathcal{L})}. \tag{9}$$

Proof. Let $w \in \mathcal{L}^*$ be a shortest nonzero vector. Since $\mathbb{Z}w \subseteq \mathcal{L}^*$, the one-dimensional Poisson identity gives

$$\rho_{1/s}(\mathcal{L}^*) \ge \rho_{1/s}(\mathbb{Z}w) = \rho_{1/(s\|w\|)}(\mathbb{Z}) \ge \frac{1}{s\|w\|}.$$

Poisson summation for $\mathcal{L}$ also gives

$$\rho_{1/s}(\mathcal{L}^*) = \frac{\det \mathcal{L}}{s^n}\rho_s(\mathcal{L}).$$

Combining the two displays and using $\|w\| = \lambda_1(\mathcal{L}^*)$ proves (9). □

Thus the primal mass bound in Lemma 4.1 gives the dual estimate needed to make the forced term negligible.

4.3 Prime selection and the valid branch

The prime has two roles. It must be large enough to reduce the nonforced mass by a factor of about $1/p$ and to move the forced lattice $p\mathcal{L}^*$ beyond the relevant Gaussian scale. On the other hand, the normalizer $\rho_t(\Gamma)$ grows with p, which lowers the hit probability. We therefore choose the smallest useful scale for p, with the subexponential factor 2^{a_n} providing smoothing slack. Put

$$P := 2^{a_n} \frac{\det \mathcal{L}}{s^n}. \tag{10}$$

Poisson summation and the zero vector in $\mathcal{L}^*$ show

$$\frac{\det \mathcal{L}}{s^n} = \frac{\rho_{1/s}(\mathcal{L}^*)}{\rho_s(\mathcal{L})} \geq \frac{1}{\rho_s(\mathcal{L})} \geq 2^{-O(\log n)}.$$

Thus $P \geq 2^{a_n - O(\log n)} > 2$ for all sufficiently large n. Compute a certified rational upper bound $\overline{P}$ with $P \leq \overline{P} \leq 2P$, and put $X = \lceil \overline{P} \rceil$ and $T = \lceil Cn \log n \log(2X) \rceil$. Independently sample uniform integers from $[X, 2X]$, testing each deterministically for primality, and stop at the first prime or after T trials. If none is found, declare the branch failed. By the prime number theorem, for all sufficiently large X, a uniformly random integer in $[X, 2X]$ is prime with probability $\Omega(1/\log X)$. For a sufficiently large absolute constant C, the failure probability is therefore $2^{-\Omega(n \log n)}$. On success, $P \leq p \leq 5P$, and therefore

$$2^{a_n} \frac{\det \mathcal{L}}{s^n} \leq p \leq 5 \cdot 2^{a_n} \frac{\det \mathcal{L}}{s^n}. \tag{11}$$

Proposition 4.3 (Valid parameter branch)**.** *For all sufficiently large n, the scale grid contains an index j_* whose bounded prime search succeeds except with probability $2^{-\Omega(n \log n)}$. On success, the resulting parameters $s = s_{j_*}$, $t = t_{j_*}$, and p satisfy*

$$s_0 \leq s < (1 + 1/n)(1 + n^{-3}) s_0,$$
$$(1 + 1/n)\sqrt{2}\, s \leq t \leq (1 + n^{-3})(1 + 1/n)\sqrt{2}\, s,$$

as well as

$$\rho_s(\mathcal{L}) \leq 2^{O(\log n)}, \qquad 2^{a_n} \frac{\det \mathcal{L}}{s^n} \leq p \leq 5 \cdot 2^{a_n} \frac{\det \mathcal{L}}{s^n},$$

and

$$ps\lambda_1(\mathcal{L}^*) \geq 2^{a_n - O(\log n)}. \tag{12}$$

Proof. The existence of j_* and the polynomial bound on the number of branches follow from Equation (6) and the definition of J. The scale and sampling parameter bounds follow from Equations (7) and (8), and the Gaussian mass bound is Lemma 4.1. On this branch, $P > 2$ for all sufficiently large n, so the bounded prime search has the stated success probability and Equation (11) gives the prime window. Finally,

$$\begin{aligned} ps\lambda_1(\mathcal{L}^*) &\geq 2^{a_n} \frac{\det \mathcal{L}\, \lambda_1(\mathcal{L}^*)}{s^{n-1}} \\ &\geq 2^{a_n - O(\log n)} \end{aligned}$$

by Lemma 4.2 and Lemma 4.1. □

The finitely many dimensions before this asymptotic regime are handled by SVP in fixed dimensions.

5 Making the Superlattice Smooth

We now fix the valid branch given by Proposition 4.3 and apply the random construction of Equation (4). The first moment identity separates the expected dual mass into a nonforced contribution and a forced contribution, which we bound separately.

5.1 Bounding the expected dual mass

Lemma 5.1 (Nonforced mass)**.** *On the valid parameter branch,*

$$\frac{\rho_r(\mathcal{L}^*)}{p} \le 2^{-a_n+O(\log n)}.$$

Proof. Apply Poisson summation (1) to $\mathcal{L}$ at parameter s:

$$\rho_r(\mathcal{L}^*) = \rho_{1/s}(\mathcal{L}^*) = \frac{\det \mathcal{L}}{s^n}\rho_s(\mathcal{L}).$$

By Lemma 4.1, the last factor is $2^{O(\log n)}$. Dividing by Equation (11) proves the result. □

The remaining term is forced into every kernel and therefore requires the dual-product estimate in Equation (12).

Lemma 5.2 (Forced mass)**.** *On the valid parameter branch,*

$$\rho_r(p\mathcal{L}^* \setminus \{0\}) \le \exp(-2^{2a_n-O(\log n)}).$$

Proof. Since scaling the lattice and inversely scaling the Gaussian parameter are equivalent,

$$\rho_r(p\mathcal{L}^* \setminus \{0\}) = \rho_{r/p}(\mathcal{L}^* \setminus \{0\}).$$

By Equation (12),

$$ps\lambda_1(\mathcal{L}^*) \ge 2^{a_n-O(\log n)}.$$

Because $r = 1/s$, this is the ratio

$$\frac{\lambda_1(\mathcal{L}^*)}{r/p} = ps\lambda_1(\mathcal{L}^*).$$

This lower bound tends to infinity superpolynomially. In particular, the condition for Equation (3) holds for all sufficiently large n. Applying that inequality to $\mathcal{L}^*$ at parameter r/p gives

$$\rho_{r/p}(\mathcal{L}^* \setminus \{0\}) \le \exp\left(-\pi(ps\lambda_1(\mathcal{L}^*))^2\right)\beta^n(n+1)^{C_{\mathrm{G}}} \le \exp(-2^{2a_n-O(\log n)}).$$ □

Combining Lemmas 3.2, 5.1 and 5.2 and $(p^{n-1}-1)/(p^n-1) < 1/p$ gives the expected mass estimate

$$\mathbb{E}_\varphi[\rho_r(M \setminus \{0\})] \le \rho_r(p\mathcal{L}^* \setminus \{0\}) + \frac{\rho_r(\mathcal{L}^*)}{p} \le 2^{-a_n+O(\log n)}. \tag{13}$$

5.2 High-probability smoothing

Proposition 5.3 (Random kernel smoothing)**.** *On the valid parameter branch, with probability* $1 - 2^{-a_n+O(\log n)}$ *over the random nonzero functional* φ*,*

$$\rho_r(M \setminus \{0\}) < \frac{1}{2}. \tag{14}$$

Every M *satisfying Equation (14) also satisfies*

$$t > \sqrt{2}\,\eta_{1/2}(\Gamma), \qquad \Gamma = M^*. \tag{15}$$

Proof. By Equation (13), Markov's inequality gives

$$\Pr_{\varphi}[\rho_r(M \setminus \{0\}) \geq 1/2] \leq 2\mathbb{E}_{\varphi}[\rho_r(M \setminus \{0\})] \leq 2^{-a_n+O(\log n)}.$$

For the second statement, recall that $\Gamma^* = M$. By definition,

$$\rho_{1/\eta_{1/2}(\Gamma)}(M \setminus \{0\}) = \frac{1}{2}.$$

This Gaussian mass is continuous and strictly increasing in its parameter. Thus Equation (14) implies

$$r < \frac{1}{\eta_{1/2}(\Gamma)}.$$

Thus $\eta_{1/2}(\Gamma) < s$. Since $t \geq (1 + 1/n)\sqrt{2}s$, this gives $t > \sqrt{2}\eta_{1/2}(\Gamma)$. □

6 Hitting a Shortest Vector

Condition on a kernel satisfying Equation (14). The smoothing guarantee allows the honest DGS algorithm to produce a full block of samples from $D_{\Gamma,t}$. It remains to bound the normalizer and hence the probability that one such sample is a shortest vector of $\mathcal{L}$.

6.1 Bounding the normalizer

Lemma 6.1 (Normalizer of the superlattice)**.** *If* M *satisfies Equation (14), then*

$$\rho_t(\Gamma) \leq 2^{n/2+a_n+O(1)}.$$

Proof. Since $1/t < r/\sqrt{2} < r$, Equation (14) implies

$$1 \leq \rho_{1/t}(M) < \frac{3}{2}.$$

By Lemma 3.1, $\det \Gamma = \det \mathcal{L}/p$, and Poisson summation gives

$$\rho_t(\Gamma) = \frac{t^n}{\det \Gamma}\rho_{1/t}(\Gamma^*) = \frac{pt^n}{\det \mathcal{L}}\rho_{1/t}(M).$$

Using $t^n = 2^{n/2+O(1)}s^n$ and Equation (11),

$$\rho_t(\Gamma) \leq \frac{3}{2}\,\Theta\left(2^{a_n}\frac{\det \mathcal{L}}{s^n}\right)\frac{2^{n/2+O(1)}s^n}{\det \mathcal{L}} = 2^{n/2+a_n+O(1)}.$$ □

6.2 The direct hit bound

Fix any shortest vector $v \in \mathcal{L}$, independently of the algorithmic randomness.

Lemma 6.2 (Hit probability for one sample)**.** *If M satisfies Equation (14), then*

$$\Pr_{X \sim D_{\Gamma,t}}[X = v] \geq 2^{-E_0 n - a_n - O(1)}.$$

Proof. By Lemma 3.1, $v \in \mathcal{L} \subset \Gamma$. Therefore

$$\Pr[X = v] = \frac{\exp(-\pi \lambda_1(\mathcal{L})^2/t^2)}{\rho_t(\Gamma)}.$$

The scale choice gives

$$\begin{aligned} \frac{\pi \lambda_1(\mathcal{L})^2}{t^2} &\leq \frac{\pi \lambda_1(\mathcal{L})^2}{2s_0^2} \\ &= \frac{\beta^2}{4e}\, n = \frac{D}{4}(\ln 2) n. \end{aligned}$$

Hence the numerator is at least $2^{-(D/4)n}$. Combining this with Lemma 6.1 proves

$$\Pr[X = v] \geq 2^{-n/2 - a_n - (D/4)n - O(1)} = 2^{-E_0 n - a_n - O(1)}.$$ □

Remark 6.3. For every superlattice $\Gamma \supset \mathcal{L}$ and every $t > 0$,

$$D_{\Gamma,t} \mid \{X \in \mathcal{L}\} = D_{\mathcal{L},t}, \qquad \Pr_{X \sim D_{\Gamma,t}}[X \in \mathcal{L}] = \frac{\rho_t(\mathcal{L})}{\rho_t(\Gamma)}.$$

Indeed, for each $x \in \mathcal{L}$, conditioning changes the probability $\rho_t(x)/\rho_t(\Gamma)$ by the common normalizing factor $\rho_t(\Gamma)/\rho_t(\mathcal{L})$. Rejection sampling all the way down to a large $D_{\mathcal{L},t}$ table can be expensive. The algorithm only needs one shortest vector and therefore scans the raw $D_{\Gamma,t}$ samples directly.

7 The Algorithm and Its Analysis

7.1 Algorithm

The complete procedure is given in Algorithm 1. The scale grid supplies a branch close to $\lambda_1(\mathcal{L})$, the Q independent kernels amplify the smoothing probability, and the R sampler calls supply the required N samples. The membership test keeps exactly the samples that realize a direct hit. The constant $C_Q > 0$ is chosen sufficiently large in the correctness proof.

Algorithm 1: Direct Hit SVP

Input: A rational nonsingular matrix B defining $\mathcal{L} = B\mathbb{Z}^n \subset \mathbb{R}^n$.

Output: A candidate $v_{\mathrm{best}} \in \mathcal{L} \setminus \{0\}$, or $\perp$.

Set $\kappa \leftarrow n^2$, $m \leftarrow \lceil 2^{n/2} \rceil$, and $N \leftarrow \lceil n^2\, 2^{E_0 n + a_n} \rceil$;
Set $R \leftarrow \lceil N/m \rceil$ and $Q \leftarrow \lceil C_Q (\log_2 n)^{3/2} \rceil$;
Set $v_{\mathrm{best}} \leftarrow \perp$;
foreach *scale grid point* $\widehat{\lambda}_j$ **do**
  Compute the directed rational values s_j, t_j from Equation (7);
  Set $P_j \leftarrow 2^{a_n} \det(\mathcal{L})/s_j^n$ and compute a certified $\overline{P}_j \in [\max\{2, P_j\}, 2\max\{2, P_j\}]$;
  Set $X_j \leftarrow \lceil \overline{P}_j \rceil$ and run the bounded prime search of Section 4.3;
  if *the search returns a prime* p_j **then**
    for $q = 1, \dots, Q$ **do**
      Choose a uniformly random nonzero $\varphi_{j,q} : \mathcal{L}^*/p_j\mathcal{L}^* \to \mathbb{F}_{p_j}$;
      Set $M_{j,q} \leftarrow \ker \varphi_{j,q}$ and $\Gamma_{j,q} \leftarrow M_{j,q}^*$;
      Set $c \leftarrow 0$;
      for $\ell = 1, \dots, R$ **do**
        Invoke the honest DGS algorithm of Theorem 2.3 on $(\Gamma_{j,q}, t_j, \kappa)$ and let $\mathcal{B}$ be its output block;
        foreach x *among the first* $\min\{N - c, |\mathcal{B}|\}$ *samples of* $\mathcal{B}$ **do**
          if $x \in \mathcal{L} \setminus \{0\}$ *and* ($v_{\mathrm{best}} = \perp$ *or* $\|x\| < \|v_{\mathrm{best}}\|$) **then**
            Set $v_{\mathrm{best}} \leftarrow x$;
          end
          Set $c \leftarrow c + 1$;
        end
        Discard $\mathcal{B}$;
      end
    end
  end
end
return v_{best};

7.2 Correctness

Theorem 7.1 (Correctness of the direct hit algorithm)**.** *For all sufficiently large n, Algorithm 1 returns a shortest nonzero vector of $\mathcal{L}$ with probability at least $1 - 2^{-\Omega(n \log n)}$.*

Proof. **The ideal experiment.** Fix a shortest vector $v \in \mathcal{L}$. There is a correct scale branch by Equation (6). On one repetition of this branch, Proposition 5.3 bounds the probability of a bad kernel by $2^{-a_n + O(\log n)}$.

Condition on the kernel being good and set $q := \Pr[D_{\Gamma,t} = v]$. By Lemma 6.2,

$$Nq \geq n^2 2^{E_0 n + a_n} 2^{-E_0 n - a_n - O(1)} = n^2 2^{-O(1)} = \Omega(n^2).$$

The first N samples produced by the R ideal blocks are independent, and hence

$$\Pr[\text{no sample equals } v] = (1 - q)^N \leq e^{-Nq} \leq \exp(-\Omega(n^2)).$$

Thus the ideal failure probability of one repetition is at most

$$2^{-a_n+O(\log n)} + \exp(-\Omega(n^2)) \le 2^{-a_n/2}$$

for all sufficiently large n. The Q repetitions use independent kernels and fresh ideal sampler randomness. Since

$$Qa_n = \Theta(n \log n),$$

choosing C_Q sufficiently large makes the probability that every repetition on the correct branch misses v at most $2^{-3n \log n}$.

From ideal to actual samples. Order all hDGS invocations across the scale branches and repetitions, and let H_i record the realized scales, primes, kernels, and all outputs preceding invocation i. For every history h that can arise in either the actual or the ideal experiment, let $P_i(h, \cdot)$ and $Q_i(h, \cdot)$ denote, respectively, the actual and ideal conditional output kernels of invocation i. Conditional on $H_i = h$, the lattice and Gaussian parameter supplied to this invocation are fixed. Since the invocation uses fresh internal randomness, the uniform guarantee of Theorem 2.3 yields a quantity $\delta_n = \exp(-\Omega(n^2))$, independent of i and h, such that

$$d_{\mathrm{TV}}\big(P_i(h, \cdot), Q_i(h, \cdot)\big) \le \delta_n.$$

On a good realized kernel, the ideal output block consists of exactly $m = \lceil 2^{n/2} \rceil$ independent samples. On all other histories, the ideal experiment retains the corresponding ideal output block, including its possibly random length.

Let $P^{(i)}$ and $Q^{(i)}$ denote the distributions of the actual and ideal transcript prefixes through invocation i, respectively. Thus, writing juxtaposition for the action of a transition kernel on a distribution,

$$P^{(i)} = P^{(i-1)}P_i, \qquad Q^{(i)} = Q^{(i-1)}Q_i.$$

Using $Q^{(i-1)}P_i$ as an intermediate transcript distribution, the triangle inequality, the contraction of total variation distance under a common Markov kernel, and the uniform conditional bound above give

$$\begin{aligned} d_{\mathrm{TV}}\big(P^{(i)}, Q^{(i)}\big) &\le d_{\mathrm{TV}}\big(P^{(i-1)}P_i, Q^{(i-1)}P_i\big) + d_{\mathrm{TV}}\big(Q^{(i-1)}P_i, Q^{(i-1)}Q_i\big) \\ &\le d_{\mathrm{TV}}\big(P^{(i-1)}, Q^{(i-1)}\big) + \delta_n. \end{aligned}$$

Since the two experiments have the same initial transcript distribution, iteration yields

$$d_{\mathrm{TV}}\big(P^{(K)}, Q^{(K)}\big) \le K\delta_n,$$

where the total number of hDGS invocations satisfies

$$K \le O(n^2)QR \le O(n^2)Q\left(1 + \frac{N}{m}\right) = 2^{(E_0-1/2)n+a_n+O(\log n)}.$$

Consequently, $K\delta_n = \exp(-\Omega(n^2))$. Finally, the membership filter and the selection of the minimum are deterministic postprocessing operations and therefore cannot increase total variation distance.

Conclusion. The bounded prime search on the correct branch fails with probability $2^{-\Omega(n \log n)}$. Passing from ideal to actual sampler outputs loses only $\exp(-\Omega(n^2))$. If the correct branch sees v, then the final candidate has length exactly $\lambda_1(\mathcal{L})$: every candidate is explicitly checked to be a nonzero member of $\mathcal{L}$, and no such vector can be shorter than v. Increasing the constants in the prime search and in Q if necessary, a union bound over the stated events proves the theorem. □

7.3 Complexity

Theorem 7.2 (Time and space complexity)**.** *In the dimension-only arithmetic model, Algorithm 1 runs in time*

$$2^{E_0 n + O(n/\sqrt{\log n})}$$

and uses space

$$2^{n/2 + O(n/\log n + \log^5 n)}.$$

Proof. One hDGS invocation costs

$$2^{n/2 + O(n/\log n + \log^5 n)}$$

time and space. On a good branch, an ideal block contains $m = \lceil 2^{n/2} \rceil$ samples. Thus the total sampling time per scale branch, including all Q repetitions, is

$$\begin{aligned} QR \cdot 2^{n/2 + O(n/\log n + \log^5 n)} &\le Q\left(1 + \frac{N}{m}\right) 2^{n/2 + O(n/\log n + \log^5 n)} \\ &= 2^{E_0 n + a_n + O(n/\log n + \log^5 n)} \\ &= 2^{E_0 n + O(n/\sqrt{\log n})}. \end{aligned}$$

There are only $O(n^2)$ branches. Invocations and branches are run sequentially, no two hDGS blocks are stored at the same time, and only one accepted candidate is kept outside the active block. Hence the space is

$$2^{n/2 + O(n/\log n + \log^5 n)}.$$

Outside the hDGS calls, the algorithm performs at most

$$JQN = O(n^2 QN)$$

lattice membership tests and norm comparisons, as well as $JQ = O(n^2 Q)$ kernel constructions. In the dimension-only arithmetic model, these operations take

$$O(n^2 QN)\,\mathrm{poly}(n) = 2^{E_0 n + a_n + O(\log n)}$$

time and polynomial additional space. The scale computations and bounded prime searches contribute only factors polynomial in n and in the suppressed input encoding lengths. These contributions are dominated by the hDGS time bound above and do not change the stated time or space bounds. Since $n/\log n + \log^5 n = o(a_n)$, the displayed estimates prove the theorem. □

Remark 7.3 (Exact implementation)**.** The resource bounds above use the dimension-only arithmetic model specified in the introduction. For completeness, all algebraic steps outside the hDGS calls can be carried out exactly on rational input.

Suppose $\mathcal{L} = B\mathbb{Z}^n$ for a rational nonsingular matrix B. Then $\mathcal{L}^* = B^{-T}\mathbb{Z}^n$. A functional on $\mathcal{L}^*/p\mathcal{L}^*$ is represented by a nonzero vector $a \in \mathbb{F}_p^n$, acting as

$$\varphi(B^{-T}z) = a^T z \bmod p.$$

The coefficient lattice

$$Z_a := \{z \in \mathbb{Z}^n : a^T z = 0 \bmod p\}$$

has an integer basis computable by Hermite or Smith normal form. If U is such a basis matrix, then

$$C := B^{-T}U, \qquad M = C\mathbb{Z}^n, \qquad \Gamma = C^{-T}\mathbb{Z}^n.$$

Certified directed rational interval arithmetic computes the values in Equation (7) and the bound $\overline{P}$. It preserves $t > \sqrt{2}s$ and $p \geq P$ by construction, while the $1 + O(1/n)$ scale factors contribute only $2^{O(1)}$. The bounded prime search uses exact integer arithmetic and deterministic primality testing.

Finally, membership in $\mathcal{L}$ is exact:

$$x \in \mathcal{L} \quad \Longleftrightarrow \quad B^{-1}x \in \mathbb{Z}^n.$$

All candidate vectors output by hDGS are represented in the basis of Γ, so this test is exact rational linear algebra followed by an integrality check.

Proof of Theorem 1.1. Choose an absolute n_0 so that Theorem 7.1 holds for every $n \geq n_0$. That theorem gives the claimed success probability, while Theorem 7.2 gives the stated resource bounds in this model. Since $n/\sqrt{\log n} = o(n)$ and $n/\log n + \log^5 n = o(n)$, these bounds are respectively $2^{E_0 n + o(n)}$ and $2^{n/2+o(n)}$. Remark 7.3 gives exact construction and membership testing on rational input. Dimensions below n_0 use SVP in fixed dimensions. □

A Uniform Quantitative Gaussian Mass Bounds

The proof has two stages. First, the KL spherical code estimate gives a uniform lattice point count over every fixed compact range of radii. Second, a shell decomposition and Gaussian scaling turn this count into the uniform mass bound in Theorem 2.2. We begin by recording the numerical margin used in the first stage. For a spherical code, let $A(d, \theta)$ denote the maximum number of unit vectors in $\mathbb{R}^d$ having pairwise angle at least θ. Put

$$\theta_c := 62.99^\circ$$

and define the KL rate

$$G(\theta) := \frac{1+\sin\theta}{2\sin\theta}\log_2\frac{1+\sin\theta}{2\sin\theta} - \frac{1-\sin\theta}{2\sin\theta}\log_2\frac{1-\sin\theta}{2\sin\theta}.$$

At the cutoff angle,

$$\begin{aligned} G(\theta_c) &= 0.33765073\ldots, \\ F_c &:= -\frac{1}{2}\log_2(1-\cos\theta_c) - 0.099 \\ &= 0.33770649\ldots. \end{aligned}$$

Thus

$$\mu := F_c - G(\theta_c) = 5.57637 \cdot 10^{-5} > 0. \tag{16}$$

Lemma A.1 (Uniform KL point count)**.** *For every fixed $R_0 \geq 1$, there are constants C_{R_0} and d_{R_0} such that the following holds. If $d \geq d_{R_0}$ and $\mathcal{L}$ is a rank-d lattice normalized by $\lambda_1(\mathcal{L}) = 1$, then, for every $R \in [1, R_0]$,*

$$|\mathcal{L} \cap B_2^d(R)| \leq C_{R_0}(d+1)\beta^d R^d. \tag{17}$$

Proof. The KL spherical code estimate at the single fixed angle θ_c [KL78, Section 6, Theorem 4, (54)] gives

$$A(d, \theta_c) \le 2^{G(\theta_c)d + o(d)}.$$

For each fixed $0 < \theta \le \theta_c$, the KL angular comparison inequality [KL78, Section 6, inequality before Corollary 1] gives

$$\begin{aligned}\limsup_{d\to\infty} \frac{1}{d} \log_2 A(d, \theta) &\le G(\theta_c) + \log_2 \frac{\sin(\theta_c/2)}{\sin(\theta/2)} \\ &= F(\theta) - \mu,\end{aligned}$$

where

$$F(\theta) := -\frac{1}{2} \log_2(1 - \cos(\min\{\theta, \theta_c\})) - 0.099.$$

Now fix a compact interval $I \subset (0, \theta_c]$. By uniform continuity of F, choose a finite grid in I so fine that rounding an angle down to the grid increases F by at most $\mu/4$. At every grid point the previous limsup is at most $F - \mu$, so one dimension threshold makes the finite collection of bounds at most $2^{d(F-3\mu/4)}$. Monotonicity of $A(d, \theta)$ then gives the stronger bound $A(d, \theta) \le 2^{d(F(\theta)-\mu/2)}$ for all $\theta \in I$. Monotonicity also handles $\theta \ge \theta_c$. In particular,

$$A(d, \theta) \le 2^{dF(\theta)} \tag{18}$$

uniformly on every fixed compact subinterval of $(0, \pi]$.

For the point count, put $\alpha := 1 + 1/d$ and partition the nonzero points into shells $r \le \|x\| < \alpha r$. For distinct points x, y in one such shell, the lattice minimum and the cosine law give

$$\cos \angle(x, y) \le \frac{\alpha + \alpha^{-1}}{2} - \frac{1}{2\alpha^2 r^2}.$$

For $1 \le r \le R_0$, the resulting angles stay in a fixed compact subinterval of $(0, \pi]$ for all sufficiently large d. If θ_r denotes the angle on the right side, then

$$2r^2(1 - \cos\theta_r) = \alpha^{-2} - r^2(\alpha + \alpha^{-1} - 2) = 1 + O_{R_0}(1/d).$$

Equation (18) therefore bounds every shell by

$$|\mathcal{L} \cap \{x : r \le \|x\| < \alpha r\}| \le 2^{0.401d + C_{R_0} r^d}.$$

There are $O_{R_0}(d)$ shells below R_0. Summing them and including the origin proves (17). □

Write

$$N_{\mathcal{L}}(R) := |\mathcal{L} \cap B_2^d(R)|$$

for a rank-d lattice $\mathcal{L}$ normalized by $\lambda_1(\mathcal{L}) = 1$.

For comparison, the elementary disjoint-ball argument is uniform over all radii:

$$N_{\mathcal{L}}(R) \le (1 + 2R)^d. \tag{19}$$

Indeed, the radius-1/2 balls centered at $\mathcal{L} \cap B_2^d(R)$ are disjoint and lie in $B_2^d(R + 1/2)$.

Proof of Theorem 2.2. It is enough first to treat $d \geq 2$; the finitely many omitted dimensions are handled at the end. By scaling, assume $\lambda_1(\mathcal{L}) = 1$. Put

$$u_0 := \sqrt{\frac{2\pi e}{\beta^2 d}}, \qquad \tau := 1 + \frac{1}{d}, \qquad R_i := \tau^i.$$

First suppose $u \leq u_0$. Partition $\mathcal{L} \setminus \{0\}$ into the shells $R_i \leq \|x\| < R_{i+1}$. Choose once and for all a sufficiently large constant R_0. For $R_i < R_0$, apply (17) on $[1, 2R_0]$. Since $\tau^d \leq e$ and there are only $O(d)$ such shells, their total mass is at most

$$\beta^d (d+1)^{O(1)} \max_{R \geq 1} \exp(-\pi R^2/u^2) R^d, \tag{20}$$

where the polynomial factor is independent of $\mathcal{L}$ and u.

For $R_i \geq R_0$, (19) and $1 + 2\tau R_i \leq 4R_i$ give the shell bound

$$(4R_i)^d \exp(-\pi R_i^2/u^2).$$

The choice $u \leq u_0$ implies

$$\frac{\pi}{u^2} \geq \frac{\beta^2}{2e} d.$$

Consequently, after increasing the fixed R_0 if necessary, these terms form a geometrically decreasing sequence. Explicitly, the ratio of successive displayed shell bounds is at most

$$\tau^d \exp\left(-\frac{\pi(\tau^2 - 1)R_i^2}{u^2}\right) \leq \exp\left(1 - \frac{\beta^2 R_i^2}{e}\right),$$

which is smaller than a fixed constant below one for $R_i \geq R_0$. If $u \leq \sqrt{2\pi/d}$, comparison with the value at $R = 1$ uses

$$\ln(4R/\beta) - \frac{R^2 - 1}{2} < 0 \qquad (R \geq R_0),$$

while for $\sqrt{2\pi/d} < u \leq u_0$ comparison with the maximum in (20) uses

$$\ln(4R) - \frac{\beta^2 R^2}{2e} < \frac{1}{2} \ln\left(\frac{\beta^2}{e}\right) \qquad (R \geq R_0).$$

Both left sides tend to $-\infty$ quadratically, so the same fixed R_0 works with a fixed strict margin. In the second range, put $z := \beta^2 u^2 d/(2\pi e)$. Then $z \geq \beta^2/e$, and hence the last display bounds the first tail term by $(\beta^2/e)^{d/2} \leq z^{d/2}$ times an exponentially small factor. The complete tail is therefore covered by the polynomial factor that depends only on the dimension in (20).

Elementary calculus now gives

$$\max_{R \geq 1} e^{-\pi R^2/u^2} R^d = \begin{cases} e^{-\pi/u^2}, & u \leq \sqrt{2\pi/d}, \\ \left(\dfrac{u^2 d}{2\pi e}\right)^{d/2}, & \sqrt{2\pi/d} < u \leq u_0. \end{cases} \tag{21}$$

The first case proves (3). In the second case, use $z = \beta^2 u^2 d/(2\pi e)$ from above. Here $\beta^2/e < z \leq 1$, so the constant factor between $z^{d/2}$ and $z^{d/2+1}$ can be included in $(d+1)^{O(1)}$. This proves (2) in the entire second range. In the first range, writing $x = u^2 d/(2\pi) \leq 1$ and using

$$\sup_{0 < x \leq 1} x^{-1} \exp\left(\frac{d}{2}\left(1 - \frac{1}{x} - \ln x\right)\right) \leq 2 \qquad (d \geq 2)$$

shows that the right side of (3) is at most a constant times the nontrivial term in (2). Thus the global bound also holds for this range.

It remains to treat $u \geq u_0$. The Gaussian scaling inequality [ADRSD15, Lemma 2.3] says that for every $a \geq 1$ and every rank-d lattice $\mathcal{L}$,

$$\rho_{av}(\mathcal{L}) \leq a^d \rho_v(\mathcal{L}).$$

The already proved bound at u_0 gives $\rho_{u_0}(\mathcal{L}) = (d+1)^{O(1)}$. Taking $a = u/u_0$ therefore yields

$$\rho_u(\mathcal{L}) \leq \left(\frac{u}{u_0}\right)^d (d+1)^{O(1)} = z^{d/2}(d+1)^{O(1)} \leq 1 + z^{d/2+1}(d+1)^{O(1)},$$

because now $z \geq 1$. Every polynomial factor introduced above is a function of d alone. For each of the finitely dimensions omitted by the quantitative estimate of lattice points, use (19) in every shell. In the range $u \leq u_0$, the same two comparisons with the first shell or the radial maximum give a uniformly convergent geometric tail; in the range $u \geq u_0$, Gaussian scaling reduces to the endpoint u_0. Thus there is a finite constant depending only on the dimension, uniform in $\mathcal{L}$ and u. By increasing the absolute exponent C_{G}, we can include all these finitely many dimensions in the factor $(d+1)^{C_{\mathrm{G}}}$. This proves the theorem uniformly over $\mathcal{L}$ and u. □

Use of Generative AI

During the early preparation of this manuscript, the authors used OpenAI Codex (GPT-5.6 sol) for drafting assistance and exploratory discussion of proof details. All mathematical claims and arguments in the final manuscript were independently verified by the authors, who take full responsibility for the content.

References

[ACKS25] Divesh Aggarwal, Yanlin Chen, Rajendra Kumar, and Yixin Shen. Improved classical and quantum algorithms for the shortest vector problem via bounded distance decoding. *SIAM Journal on Computing*, 54(2):233–278, 2025.

[ADRSD15] Divesh Aggarwal, Daniel Dadush, Oded Regev, and Noah Stephens-Davidowitz. Solving the shortest vector problem in 2^n time via discrete gaussian sampling. In *Proceedings of the 47th Annual ACM Symposium on Theory of Computing (STOC)*, pages 733–742, 2015.

[AKS01] Miklós Ajtai, Ravi Kumar, and Dandapani Sivakumar. A sieve algorithm for the shortest lattice vector problem. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 601–610, 2001.

[CCL18] Yanlin Chen, Kai-Min Chung, and Ching-Yi Lai. Space-efficient classical and quantum algorithms for the shortest vector problem. *Quantum Information and Computation*, 18(3–4):283–305, 2018.

[Hha26] Minki Hhan. Solving the shortest vector problem in time $2^{0.6039n+o(n)}$ via mid-point hessian, 2026. arXiv preprint arXiv:2608.02478.

[Kim26] Jiseung Kim. One discrete gaussian sample in $2^{n/2+o(n)}$ time, 2026. arXiv preprint arXiv:2608.03220.

[KL78] Grigorii A. Kabatiansky and Vladimir I. Levenshtein. On bounds for packings on a sphere and in space. *Problems of Information Transmission*, 14(1):1–17, 1978.

[LLL82] Arjen K. Lenstra, Hendrik W. Lenstra, Jr., and László Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982.

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

[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.

[Ope26] OpenAI. Ten advances in mathematics and theoretical computer science, 2026. Manuscript, Chapter 2: "Improved Bounds for Binary and Spherical Codes," released August 1, 2026.

[PS09] Xavier Pujol and Damien Stehlé. Solving the shortest lattice vector problem in time $2^{2.465n}$. Technical Report 2009/605, IACR Cryptology ePrint Archive, 2009.

[PS26] Amaury Pouly and Yixin Shen. Solving the shortest vector problem in $2^{0.63269n+o(n)}$ time on random lattices. In Joan Daemen and Emmanuel Thomé, editors, *Advances in Cryptology – EUROCRYPT 2026, Part IV*, volume 16544 of *Lecture Notes in Computer Science*, pages 92–123. Springer Nature Switzerland, 2026. Full version: IACR Cryptology ePrint Archive, Paper 2024/1805.