working draft

Minkowski's conjecture in dimensions at most 42

Or Ordentlich
The Hebrew University of Jerusalem

Barak Weiss
Tel Aviv University

Boaz Klartag
Tel Aviv University
Weizmann Institute of Science

September 6, 2026

Abstract

We prove Minkowski's conjecture on products of inhomogeneous linear forms in dimensions at most 42. The proof is computer-assisted. Its main geometric ingredient is a codimension-one inequality for the second moment of a lattice Voronoi cell. This inequality retains the saving obtained by choosing a lattice point in a different parallel layer. A constrained first-variation argument makes it possible to use the inequality at a maximum subject to a lower bound on the dual minimum. Combining it with a decomposition into scaled stable lattices gives an induction on the dimension that also retains a quantitative second-moment deficit. We reduce the induction to explicit inequalities in two real variables and one integer rank. These inequalities, and all Hermite bounds used in them, are verified with integer and rational arithmetic by the accompanying Python program. The verifier uses only the Python standard library.

Contents

1 Introduction

Let $L \subset \mathbb{R}^n$ be a lattice of full rank. Its determinant $\det L$ is the volume of a fundamental domain. Write

$$\mathcal{N}(L) = \sup_{x\in\mathbb{R}^n} \inf_{v\in L} \prod_{i=1}^{n} |x_i - v_i|. \tag{1.1}$$

Minkowski's conjecture asserts that $\mathcal{N}(L) \le 2^{-n} \det L$. We prove the following.

Theorem 1.1. *For every full-rank lattice $L \subset \mathbb{R}^n$ with $1 \le n \le 42$,*

$$\mathcal{N}(L) \le \frac{\det L}{2^n}. \tag{1.2}$$

Equivalently, if $B \in \mathrm{GL}_n(\mathbb{R})$ and $b \in \mathbb{R}^n$, then

$$\inf_{m\in\mathbb{Z}^n} \prod_{i=1}^{n} |(Bm+b)_i| \le \frac{|\det B|}{2^n}.$$

Throughout the paper, the formulation uses an infimum; no assertion that this infimum is attained is needed. The constant is sharp, as is seen from $L = \mathbb{Z}^n$ and $x = (1/2, \dots, 1/2)$.

The geometric statement that we prove is stronger than what is needed for the final reduction. A lattice is always considered in its real linear span, with the induced Euclidean metric. It is *unimodular* if its determinant is one. A unimodular lattice is *stable* if every nonzero sublattice has determinant at least one in its own span. Its Voronoi cell is

$$V(L) = \{x \in \operatorname{span} L : |x| \le |x - v| \text{ for every } v \in L\}.$$

We use the mean squared distance

$$M(L) = \frac{1}{\det L} \int_{V(L)} |x|^2 \, dx. \tag{1.3}$$

For a unimodular rank-d lattice, its normalized second moment is $M(L)/d$. In particular,

$$M(\mathbb{Z}^d) = \frac{d}{12}. \tag{1.4}$$

The normalization in (1.3) will also be used for lattices whose determinant is not one. With this convention,

$$M(aL) = a^2 M(L) \qquad (a > 0). \tag{1.5}$$

Theorem 1.2. *Every stable unimodular lattice L of rank $d \leq 42$ satisfies*

$$M(L) \leq \frac{d}{12}. \tag{1.6}$$

Shapira and Weiss [5] proved that the closure of every positive diagonal orbit in the space of unimodular lattices contains a stable lattice. They used this to reduce Minkowski's conjecture to a covering-radius bound for stable lattices. The inequality proved by Autissier [1] and later independently obtained by Magazinov [3],

$$R(L)^2 \leq 3M(L),$$

where $R(L)$ is the Euclidean covering radius, makes (1.6) sufficient for their reduction. We recall the complete deduction in Section 2.

Regev and Stephens-Davidowitz [4] developed a method based on stable lattices and their canonical filtration. They also studied the second moment of a Voronoi cell and its first variation. We use these ingredients, giving the lattice and variational details needed below. The additional step is to retain more information when a lattice is split into parallel layers.

1.1 The main inequality

Suppose that u is a unit vector and that

$$L = \Gamma + \mathbb{Z}(v + hu), \qquad \Gamma \subset u^\perp, \quad v \in u^\perp, \quad h > 0,$$

where Γ has full rank in $u^\perp$. Besides $M(L)$, write

$$C_L = \frac{1}{\det L} \int_{V(L)} x x^T \, dx, \qquad M(L) = \operatorname{tr} C_L. \tag{1.7}$$

The cell is centrally symmetric, so C_L is also the covariance matrix of its normalized volume measure. We prove

$$M(L) + \langle C_L u, u \rangle \leq M(\Gamma) + \frac{h^2}{6}. \tag{1.8}$$

The usual comparison with a product fundamental domain gives only

$$M(L) \leq M(\Gamma) + \frac{h^2}{12}.$$

It approximates a point using its nearest vertical layer. Allowing a different layer may reduce the horizontal error at the cost of increasing the vertical error. Inequality (1.8) says that the total improvement pays for this additional vertical error. Its proof is an elementary calculation for a periodic sequence of weighted points on the real line. We give the calculation, including the treatment of skipped layers, in Section 4.

1.2 Why an ordinary dimension induction is not enough

At an interior maximum of M among unimodular lattices, the matrix C_L is scalar. If $y \in L^*$ is a primitive dual vector, the hyperplane lattice $\Gamma = L \cap y^\perp$ has layer height $1/|y|$. This gives a useful choice of u and h in (1.8).

However, Γ is generally not unimodular. Its canonical filtration has factors of the form $\sqrt{t_i} S_i$, where S_i is stable and unimodular and $t_i \geq 1$. Applying only $M(S_i) \leq \operatorname{rank}(S_i)/12$ loses too much. The last factor has a quantitative lower bound on its dual minimum. We therefore prove, at the same time as (1.6), lower bounds for

$$D_d(a) = \inf\{d - 12M(L) : L \text{ is stable and unimodular of rank } d,\ \lambda_1(L^*)^2 \geq a\}. \tag{1.9}$$

Here λ_1 denotes the length of a shortest nonzero vector, and the infimum of the empty set is $+\infty$.

A maximum subject to the dual-minimum constraint in (1.9) need not have scalar covariance. Section 5 proves the substitute that is needed: some shortest dual vector y has

$$\left\langle C_L \frac{y}{|y|}, \frac{y}{|y|} \right\rangle \geq \frac{M(L)}{d}.$$

This lets us continue to use the layer inequality at a constrained maximum.

The resulting induction involves the squared dual minimum c, the largest canonical scale t, and the rank r of its factor. The geometric reduction to these parameters is in Section 6. Sections 7 and 8 describe the two finite checks: exact polynomial certificates for the Hermite bounds, and exact interval bounds for the dimension induction. The latter section includes a proof that successful checks imply Theorem 1.2. No search over lattices, numerical optimization, or unproved numerical approximation is part of the verification.

2 From second moments to Minkowski's conjecture

This section proves Theorem 1.1 from Theorem 1.2. We first recall the sharp conversion from second moment to covering radius. All integrals on a torus are with respect to Haar probability measure.

Lemma 2.1 (Autissier; Magazinov)**.** *For every lattice L in its span,*

$$R(L)^2 \leq 3M(L), \qquad R(L) = \max_{x \in \operatorname{span} L} \operatorname{dist}(x, L). \tag{2.1}$$

Proof. Put $q(x) = \operatorname{dist}(x, L)^2$. This is a continuous periodic function, and its torus mean is $M(L)$. Choose y with $q(y) = R(L)^2$. For a given x, let $v, w \in L$ be nearest points to x and $x - y$, respectively. The parallelogram identity gives

$$\begin{aligned} 2q(x) + 2q(x-y) &= 2|x - v|^2 + 2|x - y - w|^2 \\ &= |y - v + w|^2 + |2x - y - v - w|^2 \\ &\geq q(y) + q(2x - y). \end{aligned}$$

Both translation and the doubling map preserve Haar probability measure on the torus. Integration therefore gives

$$4M(L) \geq R(L)^2 + M(L),$$

which is (2.1). This is the proof in Autissier [1, Lemma 3.2]; see also Magazinov [3]. $\square$

Let

$$A_n = \{\mathrm{diag}(a_1, \dots, a_n) : a_i > 0,\ \prod_i a_i = 1\}.$$

We use the following published theorem, without reproving its dynamical part.

Theorem 2.2 (Shapira–Weiss [5], Theorem 1.1)**.** *For every unimodular lattice $L \subset \mathbb{R}^n$, the closure of $A_n L$ contains a stable lattice.*

We spell out the continuity issue in applying this theorem to (1.1). The function $\mathcal{N}$ is A_n-invariant. It is also upper semicontinuous in the space of lattices. Indeed, write $L = B\mathbb{Z}^n$ and represent a translate by $x = Bz$ with $z \in [0,1]^n$. The function

$$(B, z) \longmapsto \inf_{m \in \mathbb{Z}^n} \prod_{i=1}^{n} |(B(z-m))_i|$$

is upper semicontinuous, being an infimum of continuous functions. Its maximum over the fixed compact cube $[0,1]^n$ is upper semicontinuous in B. To verify the last assertion directly, take a sequence $B_j \to B$ and maximizing points z_j. A subsequence along the upper limit has $z_j \to z$, and upper semicontinuity gives the required upper bound by the value at (B, z).

Now let L_0 be a stable lattice in $\overline{A_n L}$. The preceding observations imply

$$\mathcal{N}(L) \le \mathcal{N}(L_0). \tag{2.2}$$

For each x, choose a Euclidean nearest point $v \in L_0$. The arithmetic–geometric mean inequality gives

$$\prod_{i=1}^{n} |x_i - v_i| \le \left(\frac{|x-v|^2}{n}\right)^{n/2} \le \left(\frac{R(L_0)^2}{n}\right)^{n/2}. \tag{2.3}$$

If $n \le 42$, Theorem 1.2 and Lemma 2.1 bound the last expression by 2^{-n}. This proves (1.2) when $\det L = 1$. For general L, apply the unimodular result to $(\det L)^{-1/n} L$ and use the homogeneity of the product. This completes the reduction.

3 Lattice facts used in the proof

We collect the lattice facts needed for the induction. Apart from compactness of lattices with a fixed lower bound on their minimum, the arguments in this section are elementary. They also explain precisely what is meant by a factor of the canonical filtration. The stable-lattice framework and these compactness properties are discussed in [4, 5].

3.1 Primitive sublattices and orthogonal quotients

A sublattice $K \subset L$ is *primitive* if $K = L \cap \operatorname{span} K$. Equivalently, a basis of K can be extended to a basis of L. The saturation of any sublattice is primitive and has no larger determinant. Thus stability need only be tested on primitive sublattices.

When K is primitive, we write L/K for the orthogonal projection of L onto $(\operatorname{span} K)^\perp$, within $\operatorname{span} L$. This is a lattice of rank $\operatorname{rank} L - \operatorname{rank} K$. To see discreteness, extend a basis of K to a basis of L. The projections of the added basis vectors are linearly independent and generate L/K.

The dual lattice is

$$L^* = \{y \in \operatorname{span} L : \langle y, v \rangle \in \mathbb{Z} \text{ for all } v \in L\}.$$

We adopt the conventions $\det\{0\} = 1$ and $M(\{0\}) = 0$.

Lemma 3.1. *If $K \subset L$ is primitive, then*

$$\det L = \det K \det(L/K), \qquad (L/K)^* = L^* \cap (\operatorname{span} K)^\perp. \tag{3.1}$$

The orthogonal projection of L^* onto $\operatorname{span} K$ *is K^*. In particular,*

$$\det\bigl(L^* \cap (\operatorname{span} K)^\perp\bigr) = \frac{\det K}{\det L}. \tag{3.2}$$

Proof. Choose a basis of K and extend it to a basis of L. Computing the volume of its parallelepiped by first using the vectors in K and then projecting the remaining vectors gives the determinant identity.

For $y \in (\operatorname{span} K)^\perp$, taking scalar products commutes with orthogonal projection. Thus y is integer-valued on L/K exactly when it is integer-valued on L. This proves the second identity in (3.1).

The projection of a vector of L^* is certainly integer-valued on K. Conversely, an integer-valued linear functional on K extends to one on L, because a basis of K extends to a basis of L. Its representing vector in L^* projects to the given vector in K^*. Finally, $\det Q^* = 1/\det Q$ for every full-rank lattice in its span. Apply this to $Q = L/K$ and use the first identity. □

Corollary 3.2. *The dual of a stable unimodular lattice is stable.*

Proof. Let $B \subset L^*$ be primitive. Apply (3.2) with L^* in place of L. Since $\det L = 1$, it gives

$$\det B = \det\bigl(L \cap (\operatorname{span} B)^\perp\bigr) \geq 1.$$

The assertion also covers $B = L^*$ by the convention for the zero lattice. Passing to saturations handles nonprimitive sublattices. □

A primitive dual vector selects parallel lattice layers with a particularly simple height.

Lemma 3.3. *Let L be unimodular and let $y \in L^*$ be primitive and nonzero. Put*

$$c = |y|^2, \qquad u = y/\sqrt{c}, \qquad \Gamma = L \cap y^\perp.$$

Then

$$\det \Gamma = \sqrt{c}, \qquad L = \Gamma + \mathbb{Z}(v + c^{-1/2}u) \quad \textit{for some } v \in y^\perp. \tag{3.3}$$

Proof. The homomorphism $x \mapsto \langle y, x\rangle$ from L to $\mathbb{Z}$ is onto. Indeed, if its image were $m\mathbb{Z}$ with $m > 1$, then y/m would belong to L^*, contradicting primitivity. Choose $z \in L$ with $\langle y, z\rangle = 1$. Its component parallel to u is $c^{-1/2}u$, and every lattice vector differs from an integer multiple of z by a vector of Γ. The quotient L/Γ is therefore the one-dimensional lattice of spacing $c^{-1/2}$. The determinant identity in Lemma 3.1 gives $\det \Gamma = \sqrt{c}$. □

3.2 Fundamental domains and second moments

Let μ_L be Haar probability measure on $\operatorname{span} L/L$. The nearest-point error map sends a point of the torus to its representative in $V(L)$, except on the set where a nearest point is not unique. This exceptional set is contained in a countable union of affine hyperplanes and has measure zero. On any bounded fundamental domain the map is a piecewise translation, and it sends μ_L to normalized Lebesgue measure on $V(L)$. Consequently,

$$M(L) = \int_{\operatorname{span} L/L} \operatorname{dist}(x, L)^2 \, d\mu_L(x). \tag{3.4}$$

The corresponding mean of the outer product of the nearest-point error is C_L. This observation permits us to compute these moments on fundamental domains other than the Voronoi cell.

Lemma 3.4 (Splitting the second moment)**.** *For a primitive sublattice $K \subset L$,*

$$M(L) \leq M(K) + M(L/K). \tag{3.5}$$

Proof. Identify $V(K)$ and $V(L/K)$ with subsets of orthogonal subspaces. Their sum

$$F = V(K) + V(L/K)$$

is a fundamental domain for L, up to boundaries. To reduce a point modulo L, first reduce its projection modulo L/K, using a lift in L of the chosen quotient vector. Then reduce its remaining component modulo K. The representatives are unique away from the boundaries of the two Voronoi cells.

The mean of the periodic function $\mathrm{dist}(x, L)^2$ on F equals $M(L)$, by (3.4). Since $0 \in L$, this function is bounded above by $|x|^2$. Orthogonality gives $|x_1 + x_2|^2 = |x_1|^2 + |x_2|^2$, and normalized integration on the product gives

$$\frac{1}{\det L} \int_F |x|^2 \, dx = M(K) + M(L/K).$$

This proves (3.5). □

3.3 Continuity, compactness, and the boundary of stability

We use the usual topology on lattices: locally a convergent sequence can be written as $L_j = B_j \mathbb{Z}^d$ with $B_j \to B$ invertible. In these coordinates both $M(L)$ and C_L are continuous. Here are the details, which will also justify the limiting argument in the layer inequality.

Represent torus points by Bz, $z \in [0, 1]^d$. When B varies in a compact neighborhood of an invertible matrix, its smallest singular value is bounded below and its operator norm is bounded above. A nearest integer translate m must satisfy

$$|B(z - m)| \leq |Bz|.$$

These two uniform bounds imply that only finitely many $m \in \mathbb{Z}^d$ can be relevant, uniformly in z and in this neighborhood. The minimum of the associated finitely many squared lengths is continuous. At almost every z, the nearest translate for the limiting matrix is unique, so the minimizing error vector also converges. The errors are uniformly bounded. Dominated convergence proves continuity of M and C_L. The same finite-candidate argument proves continuity of $\lambda_1(L)$ and $\lambda_1(L^*)$.

Write $\mathcal{S}_d$ for the space of stable unimodular rank-d lattices. Mahler's compactness criterion states that unimodular lattices with a uniform positive lower bound on λ_1 form a relatively compact set. Since stability implies $\lambda_1(L) \geq 1$, and sublattice determinant inequalities pass to limits, $\mathcal{S}_d$ is compact. In particular,

$$\{L \in \mathcal{S}_d : \lambda_1(L^*)^2 \geq a\} \tag{3.6}$$

is compact for every a.

Lemma 3.5. *A lattice $L \in \mathcal{S}_d$ is an interior point of $\mathcal{S}_d$ if and only if*

$$\det K > 1$$

for every nonzero proper primitive sublattice $K \subset L$. If L is a boundary point, it has such a sublattice with $\det K = 1$*, and both K and L/K are stable and unimodular.*

Proof. For each rank k, determinants of rank-k sublattices are lengths of nonzero decomposable vectors in the exterior-power lattice $\bigwedge^k L$. There are only finitely many such vectors in a bounded ball. Thus, if all proper primitive determinants exceed one, those which are at most two have a positive gap above one. For a linear map sufficiently close to the identity, their determinants remain above one. All larger determinants remain above one as well, since the action on each exterior power changes lengths by a factor uniformly close to one. A unimodular deformation of L is therefore still stable. This proves the interior assertion in one direction.

If K has rank k with $0 < k < d$ and $\det K = 1$, consider the trace-zero symmetric map which equals $-\mathrm{Id}$ on $\operatorname{span} K$ and $(k/(d-k))\mathrm{Id}$ on its orthogonal complement. Its exponential preserves the determinant of L, but decreases $\det K$. Arbitrarily small positive deformations thus leave $\mathcal{S}_d$, so L is a boundary point.

Finally, K is stable because it is a sublattice of L. The quotient has determinant one by Lemma 3.1. If $B \subset L/K$ is a primitive sublattice, its inverse image $\widetilde{B}$ in L contains K and satisfies

$$\det \widetilde{B} = \det K \det B = \det B.$$

Stability of L gives $\det B \geq 1$. This proves stability of the quotient. $\square$

3.4 A filtration with stable quotients

The next lemma is the form of the canonical filtration used in the proof. We include its construction, rather than assuming that a quotient of a stable lattice is automatically stable after normalization. That assertion would not be valid without further information.

Lemma 3.6 (Stable factors)**.** *Suppose that every nonzero sublattice of a rank-k lattice Γ has determinant at least one. There is a primitive filtration*

$$\{0\} = K_0 \subset K_1 \subset \cdots \subset K_s = \Gamma$$

whose successive orthogonal quotient lattices are

$$K_i/K_{i-1} = \sqrt{t_i}\, S_i.$$

Here S_i is stable and unimodular of rank r_i, and

$$1 \leq t_1 < \cdots < t_s, \qquad \sum_{i=1}^{s} r_i = k, \qquad \prod_{i=1}^{s} t_i^{r_i} = (\det \Gamma)^2. \tag{3.7}$$

Moreover,

$$M(\Gamma) \leq \sum_{i=1}^{s} t_i M(S_i). \tag{3.8}$$

Proof. Among nonzero primitive sublattices $K \subset \Gamma$, minimize

$$(\det K)^{1/\operatorname{rank} K}.$$

This minimum is attained. Indeed, Γ itself gives an upper bound, and for each possible rank, bounded determinant leaves only finitely many possibilities for the exterior product of a basis. Among minimizing sublattices choose one of largest rank. Denote it by K_1, its rank by r_1, and the minimum by α_1.

For every rank-j sublattice $B \subset K_1$, the minimizing property implies $\det B \geq \alpha_1^j$. Also $\det K_1 = \alpha_1^{r_1}$. Thus $S_1 = \alpha_1^{-1} K_1$ is stable and unimodular.

Consider a nonzero primitive rank-j sublattice B of the orthogonal quotient Γ/K_1. Its inverse image $\widetilde{B} \subset \Gamma$ is primitive, contains K_1, and has rank $r_1 + j$. By minimality of α_1,

$$\det K_1 \det B = \det \widetilde{B} \geq \alpha_1^{r_1+j}.$$

Hence $\det B \geq \alpha_1^j$. Equality is impossible: it would make $\widetilde{B}$ a minimizing sublattice of rank larger than r_1. The minimum scale in Γ/K_1 is therefore strictly larger than α_1.

Repeat the construction in this quotient, and lift the resulting primitive sublattices back to Γ. The ranks decrease at each quotient step, so the procedure terminates. Its scales α_i are strictly increasing, and $\alpha_1 \geq 1$ by the hypothesis on Γ. Set $t_i = \alpha_i^2$. The rank and determinant identities give (3.7). Repeated use of Lemma 3.4, followed by (1.5), gives (3.8). □

We also fix the normalization of the Hermite constant:

$$\gamma_k = \sup_{\operatorname{rank} L = k} \frac{\lambda_1(L)^2}{(\det L)^{2/k}}. \tag{3.9}$$

Only upper bounds for these constants will be used. Every upper bound needed in this paper is obtained from an exact polynomial certificate in Section 7.

4 The layer inequality

We now prove (1.8). The one-dimensional calculation is stated separately so that its constants and the treatment of inactive layers can be checked directly.

4.1 A periodic weighted grid

Lemma 4.1. *Let $a > 0$, let $N \geq 1$ be an integer, and let $(w_j)_{j\in\mathbb{Z}}$ be an N-periodic sequence of real numbers. For almost every $s \in \mathbb{R}$, let $j(s)$ minimize*

$$w_j + a(s-j)^2 \qquad (j \in \mathbb{Z}).$$

Then

$$\frac{1}{N}\int_0^N \left[w_{j(s)} + 2a(s-j(s))^2\right] ds \leq \frac{1}{N}\sum_{j=0}^{N-1} w_j + \frac{a}{6}. \tag{4.1}$$

Proof. The weights are bounded, so a minimizer exists and only finitely many indices can occur on a bounded interval of s. Two distinct competing parabolas have different linear terms after the common term as^2 is removed. They are equal at at most one point. Thus the minimizer is unique except on a locally finite set.

An index is called active if it minimizes on an interval of positive length. For a fixed index, the set where it minimizes is an interval, since comparison with each other index is a linear inequality in s. The active indices occur in increasing order as s increases. Indeed, the function with larger index has the more negative coefficient of s after subtracting as^2. Periodicity implies

$$j(s+N) = j(s) + N$$

away from ties. In particular, the active indices and their cell boundaries repeat after translation by N.

Let $j < k$ be consecutive active indices. Put

$$\ell = k - j, \qquad \delta = w_k - w_j.$$

Their cells meet at the point

$$b = \frac{j+k}{2} + \frac{\delta}{2a\ell}. \tag{4.2}$$

We first compute the following contribution associated with this gap:

$$\begin{aligned}\int_j^b [w_j + 2a(s-j)^2]\,ds + \int_b^k [w_k + 2a(s-k)^2]\,ds \\ = \frac{\ell}{2}(w_j + w_k) + \frac{a\ell^3}{6}.\end{aligned} \tag{4.3}$$

Both integrals are understood as signed integrals. This convention is needed because b need not lie between j and k.

Here is the algebra in (4.3). Set $z = \delta/(2a\ell)$, so that $b - j = \ell/2 + z$ and $k - b = \ell/2 - z$. The contribution from the weights is

$$w_j(\ell/2 + z) + w_k(\ell/2 - z) = \frac{\ell}{2}(w_j + w_k) - \frac{\delta^2}{2a\ell}.$$

The contribution from the doubled quadratic terms is

$$\begin{aligned}\frac{2a}{3}\left[(\ell/2 + z)^3 + (\ell/2 - z)^3\right] &= \frac{a\ell^3}{6} + 2a\ell z^2 \\ &= \frac{a\ell^3}{6} + \frac{\delta^2}{2a\ell}.\end{aligned}$$

The two terms involving δ^2 cancel.

The intermediate indices $j + r$, $1 \le r < \ell$, cannot improve the value at b. Comparing their values with the common value of the indices j and k, and substituting (4.2), gives

$$w_{j+r} \ge \left(1 - \frac{r}{\ell}\right) w_j + \frac{r}{\ell} w_k + ar(\ell - r). \tag{4.4}$$

Geometrically, the points $(j, w_j + aj^2)$ form a discrete lifted diagram. Consecutive active indices are consecutive vertices of its lower convex hull, and (4.4) says that intermediate lifted points lie above the joining segment.

Summing (4.4) and using

$$\sum_{r=1}^{\ell-1} r(\ell - r) = \frac{\ell(\ell^2 - 1)}{6}$$

yields

$$\frac{w_j + w_k}{2} + \sum_{r=1}^{\ell-1} w_{j+r} + \frac{a\ell}{6} \ge \frac{\ell}{2}(w_j + w_k) + \frac{a\ell^3}{6}. \tag{4.5}$$

It remains to sum the gap contributions correctly. Enumerate the active indices in one period as

$$j_0 < j_1 < \cdots < j_m = j_0 + N,$$

and let b_i be the boundary between j_i and j_{i+1}. Put $b_{-1} = b_{m-1} - N$. The actual cell of j_i is $[b_{i-1}, b_i]$. In the sum of (4.3), the two signed integrals involving this index combine as

$$\int_{b_{i-1}}^{j_i} + \int_{j_i}^{b_i} = \int_{b_{i-1}}^{b_i}.$$

For the index at the end of the period, translate by N and use periodicity of the weights. Thus the sum of (4.3) is the actual selected integral over a period of length N.

On the left side of (4.5), each intermediate weight is counted once, each active weight is counted as two halves, and the gap lengths add to N. Summing (4.5), using (4.3), and dividing by N proves (4.1). $\square$

4.2 Applying the weighted calculation to a lattice

Proposition 4.2 (Layer inequality)**.** *Let u be a unit vector, let Γ be a full-rank lattice in $u^\perp$, and let*

$$L = \Gamma + \mathbb{Z}(v + hu), \qquad v \in u^\perp, \quad h > 0.$$

Then

$$M(L) + \langle C_L u, u \rangle \le M(\Gamma) + \frac{h^2}{6}. \tag{4.6}$$

Proof. Put $a = h^2$ and let

$$q(x) = \operatorname{dist}(x, \Gamma)^2$$

be the squared-distance function on the horizontal torus $u^\perp/\Gamma$. At a point $x + hsu$, the squared distance to L is

$$\min_{j \in \mathbb{Z}} \{ q(x - jv) + a(s - j)^2 \}. \tag{4.7}$$

For a minimizing index j, the term $a(s-j)^2$ is exactly the squared component parallel to u of the nearest-point error.

First suppose that $Nv \in \Gamma$ for an integer $N \ge 1$. For fixed x, the weights

$$w_j = q(x - jv)$$

are N-periodic. Apply Lemma 4.1 and integrate x with respect to Haar probability measure on the horizontal torus. The average of each w_j is $M(\Gamma)$.

We identify the average of the selected expression on the other side. If $J(x, s)$ denotes either the minimum in (4.7) or the selected squared vertical error, then

$$J(x, s + 1) = J(x - v, s)$$

outside its measure-zero set of ties. It follows by translation invariance in x that its average over x and $s \in [0, N)$, divided by N, is its average over x and $s \in [0, 1)$.

The product of a fundamental domain for Γ and the vertical interval $[0, h)$ is a fundamental domain for L. The discussion preceding Lemma 3.4 shows that normalized nearest-point errors on this domain have the same distribution as normalized Lebesgue measure on $V(L)$. Therefore the average of (4.7) is $M(L)$ and the average of its selected vertical term is $\langle C_L u, u \rangle$. Lemma 4.1 now gives (4.6).

For general v, choose a sequence v_j converging to v such that each v_j is torsion modulo Γ. These points are dense in the horizontal torus: use rational coordinates in a basis of Γ. The corresponding lattices converge to L. By continuity of M and C_L, established in Section 3, the inequalities for these lattices pass to the limit. $\square$

For comparison with the ordinary splitting estimate, set

$$T = \langle C_L u, u \rangle, \qquad S = M(\Gamma) + \frac{h^2}{12} - M(L).$$

Every vertical coordinate of a lattice vector belongs to $h\mathbb{Z}$. Thus the selected vertical error is at least the distance to the nearest point of $h\mathbb{Z}$, and averaging gives $T \ge h^2/12$. Proposition 4.2 is exactly the assertion

$$S \ge T - \frac{h^2}{12} \ge 0.$$

This explains the additional term in the layer inequality. For an orthogonal product lattice, $S = 0$ and $T = h^2/12$, so equality holds.

5 First variation under a dual-minimum constraint

Fix a full-rank lattice L and a positive definite symmetric matrix Q. Define

$$\mathcal{F}_L(Q) = \int_{\mathbb{R}^d/L} \min_{v \in L} (x-v)^T Q (x-v)\, d\mu_L(x). \tag{5.1}$$

For every invertible matrix B,

$$\mathcal{F}_L(B^T B) = M(BL). \tag{5.2}$$

Indeed, the map induced by B sends Haar probability measure on $\mathbb{R}^d/L$ to Haar probability measure on $\mathbb{R}^d/BL$, and it sends the minimum in (5.1) to the Euclidean squared distance to BL.

Lemma 5.1 (First variation)**.** *For every symmetric matrix H,*

$$D\mathcal{F}_L(\mathrm{Id})[H] = \mathrm{tr}(HC_L). \tag{5.3}$$

Consequently,

$$\frac{d}{ds} M(e^{sH}L)\bigg|_{s=0} = 2\,\mathrm{tr}(HC_L). \tag{5.4}$$

Proof. Choose a bounded fundamental domain for L. For Q in a sufficiently small neighborhood of the identity, a uniform positive lower bound on its eigenvalues and comparison with the candidate $v = 0$ show that only finitely many lattice vectors can minimize in (5.1), uniformly on this domain. For almost every x, the Euclidean nearest vector $v(x)$ is unique. At each such point, the derivative of the minimum at $Q = \mathrm{Id}$ is

$$(x - v(x))^T H (x - v(x)).$$

The difference quotients are uniformly bounded, because the relevant error vectors are bounded. Differentiation under the integral is therefore justified by dominated convergence. The nearest-point error distribution gives

$$D\mathcal{F}_L(\mathrm{Id})[H] = \frac{1}{\det L} \int_{V(L)} x^T H x\, dx = \mathrm{tr}(HC_L).$$

For the second identity, use (5.2) with $B = e^{sH}$, so that $B^T B = e^{2sH}$ has derivative $2H$ at zero. $\square$

This is the Voronoi-moment first variation used in [4]. Notice that the normalized definition of M in (1.3) makes (5.2) valid even when $\det B \neq 1$.

Corollary 5.2. *At an interior local maximum of M on the space of unimodular lattices,*

$$C_L = \frac{M(L)}{d}\mathrm{Id}.$$

Proof. For every symmetric trace-zero H, the deformation $e^{sH}L$ is unimodular. Its derivative in (5.4) must vanish. The orthogonal complement, under the trace inner product, of all symmetric trace-zero matrices consists of the scalar matrices. $\square$

For the induction, the next statement is more useful than scalar covariance.

Lemma 5.3 (A shortest dual direction)**.** *Let L maximize M on the set in* (3.6)*. If L is an interior point of the stable locus, some shortest nonzero vector $y \in L^*$ satisfies*

$$\left\langle C_L \frac{y}{|y|}, \frac{y}{|y|} \right\rangle \geq \frac{M(L)}{d}. \tag{5.5}$$

Proof. Put $m = M(L)/d$ and $H = C_L - m\mathrm{Id}$. If $H = 0$, every nonzero direction satisfies the assertion. Otherwise $\operatorname{tr} H = 0$, and Lemma 5.1 gives

$$\frac{d}{ds} M(e^{sH}L)\bigg|_{s=0} = 2\operatorname{tr}(H^2) > 0. \tag{5.6}$$

Suppose that every shortest dual vector fails (5.5). Then $\langle Hy, y\rangle < 0$ for all such vectors. Since the dual of $e^{sH}L$ is $e^{-sH}L^*$,

$$\frac{d}{ds}|e^{-sH}y|^2\bigg|_{s=0} = -2\langle Hy, y\rangle > 0$$

for every shortest y.

There are finitely many shortest vectors and a positive gap to the next squared length. A uniform lower bound on the smallest singular value of e^{-sH} prevents all longer vectors from entering this gap for sufficiently small s. The preceding positive derivatives then show that the dual minimum strictly increases for sufficiently small positive s.

The deformation stays unimodular because $\operatorname{tr} H = 0$, and stays stable because L is an interior point. It satisfies the dual-minimum constraint and increases M by (5.6). This contradicts maximality. $\square$

Choose y as in Lemma 5.3, and put $c = |y|^2$. A shortest vector is primitive, since a nontrivial integer divisor would be shorter. Lemma 3.3 and Proposition 4.2 therefore give

$$\frac{d+1}{d} M(L) \le M(L \cap y^{\perp}) + \frac{1}{6c}. \tag{5.7}$$

The dual-minimum constraint may be active; Lemma 5.3 does not assert scalar covariance in that case. Inequality (5.7) is the precise substitute needed below.

6 The quantitative dimension induction

The purpose of this section is to replace a constrained maximum over lattices by a minimum over two real parameters and one integer. We keep track of the assumptions at each step so that the order of the induction is explicit.

6.1 Deficit bounds in lower dimensions

For $d \ge 1$ and $a \ge 0$, define $D_d(a)$ by (1.9). The function D_d is nondecreasing: increasing a restricts the set over which the infimum is taken. When the constrained set is nonempty, the infimum is attained, by (3.6) and continuity of M.

Suppose that the bound $M(L) \le q/12$ has already been proved for all stable rank-q lattices with $q < d$. Suppose also that we have nonnegative, nondecreasing lower bounds

$$b_q(a) \le D_q(a) \qquad (q < d). \tag{6.1}$$

A value $b_q(a) = +\infty$ is used only when the corresponding constrained set is empty. Values below one may be set to zero. The induction starts with

$$b_1(a) = \begin{cases} 0, & a \le 1, \\ +\infty, & a > 1. \end{cases} \tag{6.2}$$

Indeed, every unimodular one-dimensional lattice is isometric to $\mathbb{Z}$.

We next analyze an interior maximum of M in the rank-d set defining $D_d(a)$. Choose a shortest dual vector as in Lemma 5.3, and write

$$c = |y|^2, \qquad \Gamma = L \cap y^{\perp}.$$

Then $\det \Gamma = \sqrt{c}$, and every sublattice of Γ has determinant at least one. Apply Lemma 3.6. Its factors are $\sqrt{t_i} S_i$, of ranks r_i, with

$$\sum_i r_i = d-1, \qquad \prod_i t_i^{r_i} = c, \qquad 1 \le t_1 < \cdots < t_s. \tag{6.3}$$

6.2 The last factor and its dual minimum

Let $\sqrt{t}\, S = \Gamma/K$ be the last factor, where $t = t_s$, $r = r_s$, and $K = K_{s-1}$ is the penultimate member of the filtration. The case $K = \{0\}$ is allowed. From (6.3),

$$t^r \le c \le t^{d-1}. \tag{6.4}$$

The first inequality uses that all other scales are at least one; the second uses that all scales are at most t.

Consider the dual sublattice

$$\Delta = L^* \cap (\operatorname{span} K)^{\perp}.$$

It has rank $r+1$, contains y, and has no shorter nonzero vector than y. Its determinant is determined by Lemma 3.1 and by the last quotient:

$$\det(\Delta)^2 = \det(K)^2 = \frac{c}{t^r}, \qquad \lambda_1(\Delta)^2 = c. \tag{6.5}$$

The Hermite inequality in dimension $r+1$ gives

$$c \le \gamma_{r+1} \left(\frac{c}{t^r}\right)^{1/(r+1)}.$$

Raising to the power $r+1$ and cancelling c yields

$$(ct)^r \le \gamma_{r+1}^{r+1}. \tag{6.6}$$

There is also a lower bound for the dual minimum of S. To identify the relevant projection precisely, work in the quotient space $(\operatorname{span} K)^{\perp}$. There, L/K is a lattice with dual Δ, and Γ/K is a primitive hyperplane sublattice. The projection assertion in Lemma 3.1, applied in this quotient space, says that the projection of Δ onto the span of Γ/K is

$$(\Gamma/K)^* = t^{-1/2} S^*.$$

The complementary direction is the line spanned by y.

Take any nonzero vector w in this projection and lift it to a vector of Δ. Subtracting an integer multiple of y makes the component parallel to y have length at most $|y|/2$. The resulting vector is still nonzero, because its projection is w, and its squared length is at least c by (6.5). Orthogonality therefore gives

$$|w|^2 + \frac{c}{4} \ge c.$$

After rescaling by $\sqrt{t}$, this proves

$$\lambda_1(S^*)^2 \ge \frac{3}{4} ct, \qquad \frac{3}{4} ct \le \gamma_r. \tag{6.7}$$

The second assertion follows because S^* is unimodular.

Now (6.1) applies to S:

$$12M(S) \le r - b_r(3ct/4).$$

Use this bound for the last factor and the basic lower-dimensional theorem for every other factor. Equation (3.8) gives

$$12M(\Gamma) \le \sum_i r_i t_i - t\, b_r(3ct/4). \tag{6.8}$$

The positive quantity $t\, b_r(3ct/4)$ is the information that would be lost in a dimension induction retaining only $M(S) \le r/12$.

6.3 A two-parameter lower bound

The logarithmic mean

$$\ell(t) = \int_0^1 t^s\, ds = \frac{t-1}{\log t} \quad (t > 1), \qquad \ell(1) = 1$$

is nondecreasing on $[1, \infty)$. Convexity of $s \mapsto t^s$ gives $t^s \le 1 - s + st$, and hence

$$\ell(t) \le \frac{t+1}{2}. \tag{6.9}$$

Since each $t_i \le t$, equations (6.3) imply

$$\begin{aligned} \sum_i r_i(t_i - 1) &= \sum_i r_i \ell(t_i) \log t_i \\ &\le \ell(t) \sum_i r_i \log t_i = \ell(t) \log c \le \frac{t+1}{2} \log c. \end{aligned} \tag{6.10}$$

The terms with $t_i = 1$ are zero, so no limiting convention is needed in this sum.

Combining (5.7), (6.8), and (6.10), we find

$$\frac{12(d+1)}{d} M(L) \le d - 1 + \frac{t+1}{2} \log c - t\, b_r(3ct/4) + \frac{2}{c}.$$

Rearranging gives the basic recursive estimate

$$d - 12M(L) \ge \Phi_d(r, c, t), \tag{6.11}$$

where

$$\Phi_d(r, c, t) = \frac{d}{d+1} \left[2 - \frac{2}{c} - \frac{t+1}{2} \log c + t\, b_r(3ct/4) \right]. \tag{6.12}$$

Let $G_k \ge \gamma_k$ be any certified upper bounds, with $G_1 = 1$. The parameters produced by the geometry satisfy

$$\begin{gathered} 1 \le r < d, \qquad \max\{1, a\} \le c \le G_d, \qquad t \ge 1, \\ t^r \le c \le t^{d-1}, \qquad (ct)^r \le G_{r+1}^{r+1}, \qquad 3ct/4 \le G_r. \end{gathered} \tag{6.13}$$

Here $c \ge 1$ follows from stability of the dual, and $c \ge a$ is the imposed constraint. Replacing the Hermite constants by upper bounds only enlarges the parameter domain.

6.4 The boundary case

Suppose that a maximum in the constrained rank-d set lies on the boundary of $\mathcal{S}_d$. Lemma 3.5 supplies a proper nonzero primitive sublattice K with determinant one. Both K and $Q = L/K$ are stable and unimodular. Moreover,

$$Q^* = L^* \cap (\operatorname{span} K)^\perp \subset L^*,$$

so $\lambda_1(Q^*)^2 \geq a$. Put $q = \operatorname{rank} Q$. Splitting the second moment gives

$$\begin{aligned} d - 12M(L) &\geq (\operatorname{rank} K - 12M(K)) + (q - 12M(Q)) \\ &\geq b_q(a). \end{aligned} \tag{6.14}$$

We used the basic lower-dimensional theorem for K and the constrained deficit bound for Q.

It follows that lower bounds for $D_d(a)$ can be obtained by taking the minimum of the boundary bounds $b_q(a)$, $1 \leq q < d$, and lower bounds for (6.12) on (6.13). Compactness ensures that these two cases account for every nonempty constrained class: its maximum of M is either an interior point or a boundary point of the stable locus.

6.5 The first-counterexample test and the order of induction

Before constructing nonnegative deficit bounds in dimension d, we must prove the basic theorem in that dimension. The following elementary exclusion makes this possible without assuming anything about b_d.

Lemma 6.1. *Assume* (1.6) *in every dimension below d. If* (1.6) *fails in dimension d, and L maximizes M on $\mathcal{S}_d$, then L is an interior point and*

$$\lambda_1(L^*)^2 > 2. \tag{6.15}$$

Proof. A boundary maximum splits into two lower-dimensional stable unimodular lattices. Lemma 3.4 would give $M(L) \leq d/12$, so the maximum is interior. Corollary 5.2 gives $C_L = m\mathrm{Id}$, where $m = M(L)/d > 1/12$.

Choose a shortest dual vector, and put $c = |y|^2$. For the factors of $\Gamma = L \cap y^\perp$, repeat each scale t_i exactly r_i times. This gives $d-1$ numbers at least one, with product c. The elementary inequality

$$x_1 + \cdots + x_k \leq k - 1 + \prod_{i=1}^{k} x_i \qquad (x_i \geq 1) \tag{6.16}$$

follows by induction from $(x-1)(y-1) \geq 0$. Consequently, the lower-dimensional basic theorem and splitting give

$$12M(\Gamma) \leq \sum_i r_i t_i \leq d - 2 + c.$$

Using scalar covariance in the layer inequality yields

$$12(d+1)m \leq d - 2 + c + \frac{2}{c}.$$

For $1 \leq c \leq 2$, the identity

$$c + \frac{2}{c} - 3 = \frac{(c-1)(c-2)}{c} \leq 0$$

would imply $m \leq 1/12$, a contradiction. Thus $c > 2$. $\square$

Proposition 6.2 (Scalar criterion for the next dimension)**.** *Assume the basic theorem and nonnegative deficit bounds* $b_q \le D_q$ *for all* $q < d$*. If*

$$\Phi_d(r, c, t) \ge 0 \tag{6.17}$$

for every triple satisfying (6.13) *with* $c \ge 2$*, then the basic theorem holds in dimension* d*.*

Proof. Otherwise choose a maximum as in Lemma 6.1. It is interior, has $c > 2$, and its parameters satisfy (6.13). The recursive estimate (6.11) and (6.17) imply $d - 12M(L) \ge 0$, contradicting its choice. □

The induction proceeds in two stages at each dimension. First, using only lower-dimensional deficit bounds, verify (6.17) and deduce the basic theorem from Proposition 6.2. Second, use the interior and boundary estimates to obtain lower bounds for $D_d(a)$. At this second stage, any lower bound may be replaced by its maximum with zero, because the basic theorem in dimension d has just been proved. Section 8 implements these two stages using a finite table of thresholds and exact lower bounds on rectangular parameter boxes.

7 Exact Hermite bounds from polynomial witnesses

We now specify the numerical inputs G_d and explain how they are verified. They come from the Fourier method of Cohn and Elkies [2]. For clarity, we give the lattice version of the argument and the polynomial identities used by the verifier. No table of numerical Hermite constants is taken as an assumption.

Our Fourier-transform convention is

$$\widehat{f}(\xi) = \int_{\mathbb{R}^d} f(x) e^{-2\pi i \langle x, \xi \rangle}\, dx.$$

Lemma 7.1 (Fourier certificate for a Hermite bound)**.** *Suppose that a real Schwartz function* f *on* $\mathbb{R}^d$ *satisfies*

$$f(0) = \widehat{f}(0) > 0, \qquad f(x) \le 0 \text{ if } |x| \ge R, \qquad \widehat{f}(\xi) \ge 0 \text{ for all } \xi. \tag{7.1}$$

Then $\gamma_d \le R^2$*.*

Proof. Let L be any full-rank lattice with $\lambda_1(L) \ge R$. The sign of f gives

$$\sum_{v \in L} f(v) \le f(0).$$

Poisson summation and the sign of $\widehat{f}$ give

$$\sum_{v \in L} f(v) = \frac{1}{\det L} \sum_{w \in L^*} \widehat{f}(w) \ge \frac{\widehat{f}(0)}{\det L}.$$

The Schwartz assumption ensures absolute convergence of these sums. Since $f(0) = \widehat{f}(0) > 0$, we obtain $\det L \ge 1$.

For an arbitrary lattice K, apply this to $L = (R/\lambda_1(K))K$. It gives

$$\left(\frac{R}{\lambda_1(K)} \right)^d \det K \ge 1,$$

or $\lambda_1(K)^2 / (\det K)^{2/d} \le R^2$. Take the supremum over K. □

7.1 The polynomial Fourier identity

Put $X = 2\pi|x|^2$. For $k \geq 0$, let

$$P_k(X) = 2^k k!\, L_k^{d/2-1}(X),$$

where the generalized Laguerre polynomials may be defined by

$$\sum_{k=0}^{\infty} L_k^{\alpha}(X) z^k = (1-z)^{-\alpha-1} \exp\left(-\frac{Xz}{1-z}\right), \qquad |z| < 1. \tag{7.2}$$

They satisfy the integer-coefficient recurrence

$$\begin{aligned} P_0(X) = 1, \qquad P_1(X) = d - 2X, \\ P_{k+1}(X) = (4k + d - 2X)P_k(X) - 2k(2k + d - 2)P_{k-1}(X). \end{aligned} \tag{7.3}$$

This follows by differentiating (7.2) in z and comparing coefficients. It is also the recurrence used to construct the polynomials in the code.

The relevant Fourier identity is

$$(\widehat{P_k(2\pi|x|^2)e^{-\pi|x|^2}})(\xi) = (-1)^k P_k(2\pi|\xi|^2) e^{-\pi|\xi|^2}. \tag{7.4}$$

Here is a direct verification of the normalization. Multiply (7.2), with $\alpha = d/2 - 1$ and $X = 2\pi|x|^2$, by $e^{-\pi|x|^2}$. The result is

$$(1-z)^{-d/2} \exp\left(-\pi \frac{1+z}{1-z} |x|^2\right).$$

Using the Fourier transform of a Gaussian, its Fourier transform is

$$(1+z)^{-d/2} \exp\left(-\pi \frac{1-z}{1+z} |\xi|^2\right).$$

This is the original expression with z replaced by $-z$. Comparing coefficients proves the identity for the Laguerre polynomials and hence for P_k. Coefficient comparison can be justified by differentiating under the integral for real z in a neighborhood of zero, where a fixed integrable Gaussian dominates the derivatives. This is the polynomial Fourier basis used in [2, Section 7].

For each dimension, the data file supplies integers $a_0, \ldots, a_m$ and a positive rational number X_0. Define

$$F(X) = \sum_{k=0}^{m} a_k P_k(X), \qquad \widetilde{F}(X) = \sum_{k=0}^{m} (-1)^k a_k P_k(X). \tag{7.5}$$

The associated functions $f(x) = F(2\pi|x|^2)e^{-\pi|x|^2}$ and $\widehat{f}(x) = \widetilde{F}(2\pi|x|^2)e^{-\pi|x|^2}$ form a Fourier pair by (7.4). The exact checks are

$$F(0) = \widetilde{F}(0) > 0, \qquad F(X) \leq 0 \quad (X \geq X_0), \qquad \widetilde{F}(X) \geq 0 \quad (X \geq 0). \tag{7.6}$$

Lemma 7.1 then gives

$$\gamma_d \leq \frac{X_0}{2\pi}. \tag{7.7}$$

7.2 Checking the polynomial signs with integers

The file `anc/hermite_witnesses.json` supplies, for each $2 \le d \le 42$, the integer coefficients in (7.5), the rational number X_0, and a list of common rational double roots. The degree is at most 47. The verifier first reconstructs F and $\widetilde{F}$ using (7.3) and checks their values at zero. Only after this common normalization has been checked may the two polynomials be rescaled separately by positive constants.

Let the supplied common roots be $\rho_1, \dots, \rho_s$. The verifier divides both polynomials by

$$P(X) = \prod_{j=1}^{s} (X - \rho_j)^2$$

and divides $\widetilde{F}$ by the further factor $(X - X_0)^2$. Every division is exact: the remainder must be zero as a rational number. The quotient is multiplied by a positive common denominator and divided by the positive greatest common divisor of its integer coefficients. These operations preserve its sign. Thus it remains to prove nonnegativity of a positive multiple of $-F/P$ on $[X_0, \infty)$ and of a positive multiple of $\widetilde{F}/(P(X - X_0)^2)$ on $[0, \infty)$.

We use the following elementary sign test. Let $p(X) = \sum_{j=0}^{m} p_j X^j$ be a polynomial with rational coefficients. To prove $p(X) \ge 0$ for $X \ge a$, first replace it by $p(X + a)$, clearing denominators with a positive multiplier. Choose $B > 0$. On $[0, B]$, its Bernstein representation is

$$p(X) = \sum_{k=0}^{m} \beta_k \binom{m}{k} (X/B)^k (1 - X/B)^{m-k}, \qquad \beta_k = \sum_{j=0}^{k} \frac{\binom{k}{j}}{\binom{m}{j}} p_j B^j. \tag{7.8}$$

The coefficient formula follows from

$$u^j = \sum_{k=j}^{m} \frac{\binom{k}{j}}{\binom{m}{j}} \binom{m}{k} u^k (1 - u)^{m-k}.$$

If every β_k is nonnegative, then $p \ge 0$ on $[0, B]$. If every ordinary coefficient of $p(X + B)$ is nonnegative, then $p \ge 0$ on $[B, \infty)$. These two finite tests prove the required sign on the whole ray. All denominators in (7.8) can be cleared by a positive integer.

For the supplied data, both tests succeed for each reduced polynomial. The verifier tries $B = 1, 2, 4, \dots$ until the tail test passes, and then checks the Bernstein coefficients on the finite interval. No root isolation, approximate sign evaluation, or interval subdivision is required. The recorded value of B and the result of each test appear in the output file.

The rational root locations and coefficient lists are witnesses, not assumptions about polynomial signs. Their validity is established by these exact divisions and sign checks. Consequently, the method by which the coefficients were first selected plays no role in the verification.

7.3 Converting the witnesses to rational bounds

The dimension induction uses rational upper bounds $G_d = g_d/1000$. To compute them without a floating-point value of π, the verifier checks

$$\pi > \pi_0 := \frac{314159}{100000}. \tag{7.9}$$

One exact check uses Machin's identity

$$\pi = 16 \arctan(1/5) - 4 \arctan(1/239).$$

The six-term alternating sum for arctan(1/5) is a lower bound, and $\arctan(1/239) < 1/239$. Thus

$$16\sum_{j=0}^{5}\frac{(-1)^j}{(2j+1)5^{2j+1}}-\frac{4}{239}>\frac{314159}{100000}$$

is a sufficient check; it is an integer comparison after clearing denominators. Machin's identity follows from the tangent addition formula, with the angle in the first quadrant.

For a witness with parameter X_0, set

$$g_d=\left\lceil\frac{1000X_0}{2\pi_0}\right\rceil,\qquad G_d=\frac{g_d}{1000}. \tag{7.10}$$

Equations (7.7) and (7.9) show that $\gamma_d \le G_d$. We use $G_1 = 1$ separately. For example, in dimension 42 the certificate has

$$X_0=\frac{19199}{500},\qquad G_{42}=\frac{6112}{1000}.$$

All the rational bounds used in the proof are listed in Section 9. The listing is an output of the exact witness verification, not an additional input that must be trusted.

8 A finite problem that implies the theorem

We give a complete finite specification of the dimension induction. It is useful to separate its mathematical soundness from the implementation details. Every quantity defined in this section is either an integer, a rational number, or the symbol $+\infty$. The program evaluates precisely the stated lower bounds.

8.1 The grids and the meaning of the table

Fix the denominators

$$q=200,\qquad p=1000,\qquad Q=10^6,\qquad Q_{\log}=10^{12}. \tag{8.1}$$

The c-mesh has width $1/q$, the t-mesh has width $1/p$, deficit bounds are stored in units of $1/Q$, and logarithm upper bounds in units of $1/Q_{\log}$.

Put $G_{\max}=\max_{1\le d\le 42} G_d$ and

$$I=\lfloor qG_{\max}\rfloor-q,\qquad a_i=1+\frac{i}{q}\quad(0\le i\le I+1).$$

We construct table entries $B_{d,i}$ for $1\le d\le 42$ and $0\le i\le I$. Their intended meaning is

$$\begin{aligned}&L\in\mathcal{S}_d,\quad \lambda_1(L^*)^2\ge a_i \quad\Longrightarrow\quad d-12M(L)\ge B_{d,i}/Q;\\ &B_{d,i}=+\infty \quad\Longrightarrow\quad \{L\in\mathcal{S}_d:\lambda_1(L^*)^2\ge a_i\}=\varnothing.\end{aligned} \tag{8.2}$$

For thresholds below one, we use the lower bound zero. At dimension one the entries are

$$B_{1,0}=0,\qquad B_{1,i}=+\infty\quad(i>0). \tag{8.3}$$

We first give a common bounded range for t. If a triple is admissible, then $t^r\le c$ and $(ct)^r\le G_{r+1}^{r+1}$ imply

$$t^{r(r+1)}\le(ct)^r\le G_{r+1}^{r+1},\quad\text{hence}\quad t^r\le G_{r+1}.$$

The exact checks on the rational bounds include

$$G_{r+1} \le (231/200)^r \qquad (1 \le r \le 41). \tag{8.4}$$

Thus $1 \le t \le 231/200 = 1.155$. The t-intervals are

$$[t_j, t_{j+1}], \qquad t_j = 1 + j/p, \quad 0 \le j < J, \qquad J = 155.$$

Their union contains the full permitted range, including its upper endpoint.

A parameter box is

$$[c_-, c_+] \times [t_-, t_+] = [a_i, a_{i+1}] \times [t_j, t_{j+1}]. \tag{8.5}$$

At dimension d it is enough to consider

$$0 \le i \le s_d := \lfloor qG_d \rfloor - q.$$

The last box is kept even if its right endpoint exceeds G_d. This enlarges the domain and cannot invalidate a lower bound.

8.2 Exact rules for excluding a box

For a proposed rank r, the box can contain an admissible triple only if

$$\begin{aligned} t_-^r &\le c_+, \\ (c_- t_-)^r &\le G_{r+1}^{r+1}, \\ \frac{3}{4} c_- t_- &\le G_r, \\ c_- &\le t_+^{d-1}. \end{aligned} \tag{8.6}$$

We retain the box only when all four inequalities hold. Each is a necessary condition, obtained by using the appropriate endpoints in (6.13). Therefore a discarded box contains no admissible triple. Retaining a box does not assert that it contains a lattice or even an admissible triple; including extra boxes merely weakens the test.

There is one additional valid exclusion, supplied by the lower-dimensional induction. Put

$$k = \left\lfloor q \frac{3}{4} c_- t_- \right\rfloor - q. \tag{8.7}$$

If $k < 0$, set $b = 0$. If $k \ge 0$, set $b = B_{r,k}$. The third inequality in (8.6) ensures $k \le I$. If $b = +\infty$, discard the box. In fact, an actual factor arising from this box would have

$$\lambda_1(S^*)^2 \ge \frac{3}{4} ct \ge \frac{3}{4} c_- t_- \ge a_k,$$

contradicting the empty-set assertion in (8.2). If b is finite, the same reasoning gives

$$t\, b_r(3ct/4) \ge t_- \frac{b}{Q} \tag{8.8}$$

when the deficit bounds in (6.12) are taken from the table. We used $b \ge 0$ in multiplying by the lower endpoint t_-.

All comparisons in (8.6) are made by integer arithmetic. For example, writing $c_- = C/q$, $c_+ = (C+1)/q$, $t_- = T/p$, and $G_{r+1} = g_{r+1}/1000$, the first two comparisons are

$$qT^r \le (C+1)p^r, \qquad 1000^{r+1}(CT)^r \le g_{r+1}^{r+1}(qp)^r.$$

Thus there is no ambiguity from evaluating a real power or comparing rounded decimal values.

8.3 Rational upper bounds for logarithms

We require upper bounds for $\log a_i$. Set $L_0 = 0$ and, for $i \geq 1$, define the integers

$$L_i = L_{i-1} + \left\lceil Q_{\log} \frac{2(q+i)-1}{2(q+i)(q+i-1)} \right\rceil . \tag{8.9}$$

Then

$$U_i := L_i / Q_{\log} \geq \log a_i . \tag{8.10}$$

Indeed,

$$\log a_i = \int_q^{q+i} \frac{dx}{x}.$$

The function $1/x$ is convex on $(0, \infty)$, so its integral over $[C-1, C]$ is at most the trapezoidal average

$$\frac{1}{2}\left(\frac{1}{C-1} + \frac{1}{C}\right) = \frac{2C-1}{2C(C-1)}.$$

Summing these bounds and rounding each increment upward proves (8.10). No estimate for an asymptotic error term is needed: convexity gives the correct direction on every interval.

8.4 The lower bound attached to a retained box

For a retained box, let

$$A = \frac{t_+ + 1}{2}.$$

Because $c \geq 1$ and $t \leq t_+$, the part of the bracket in (6.12) not involving the deficit is at least

$$\phi_A(c) = 2 - 2/c - A \log c.$$

Its derivative is

$$\phi_A'(c) = \frac{2 - Ac}{c^2}.$$

This derivative changes sign at most once, from positive to negative. Hence ϕ_A has no interior minimum on $[c_-, c_+]$, and

$$\phi_A(c) \geq \min\{\phi_A(c_-), \phi_A(c_+)\}. \tag{8.11}$$

This endpoint observation avoids estimating $2/c$ and $\log c$ at different endpoints.

Define the integers

$$\begin{aligned} E_- &= 2Q - \left\lceil \frac{2Q}{c_-} \right\rceil - \lceil QAU_i \rceil, \\ E_+ &= 2Q - \left\lceil \frac{2Q}{c_+} \right\rceil - \lceil QAU_{i+1} \rceil, \\ E &= \min\{E_-, E_+\}. \end{aligned} \tag{8.12}$$

Equations (8.10) and (8.11) imply $E/Q \leq \phi_A(c)$ throughout the box. With b defined after (8.7), assign to the box the integer

$$J_{d,r,i,j} = \left\lfloor \frac{d}{d+1} \left(E + \lfloor t_- b \rfloor \right) \right\rfloor . \tag{8.13}$$

Every operation in (8.12)–(8.13) is on rationals and integers. Floors are genuine floors, also for negative numbers.

Lemma 8.1. *Assume* (8.2) *in rank r, with nonnegative finite table entries. On every retained box, any geometric parameters from Section 6 satisfy*

$$d - 12M(L) \geq J_{d,r,i,j}/Q. \tag{8.14}$$

Proof. The negative terms in (6.12) are bounded below by E/Q. The last term is bounded below by $t_{-}b/Q$ by (8.8). Replacing it by $\lfloor t_{-}b \rfloor/Q$ weakens the bound. Multiply by the positive factor $d/(d+1)$ and round downward once more. This gives (8.13) and (8.14). □

8.5 The finite test and the update rule

For each $i \leq s_d$, let

$$I_{d,i} = \min J_{d,r,i,j}, \tag{8.15}$$

where the minimum is over all retained boxes at this i and all $1 \leq r < d$. A minimum over an empty set is $+\infty$. For $i > s_d$, set $I_{d,i} = +\infty$.

Since $a_q = 2$, the first-counterexample test in dimension d is the finite inequality

$$m_d := \min_{q \leq i \leq s_d} I_{d,i} \geq 0. \tag{8.16}$$

If $s_d < q$, the minimum is over an empty set and the test passes. In that case the Hermite bound already prevents $c > 2$.

Only after (8.16) has passed, set, for $a_k \leq G_d$,

$$B_{d,k} = \max\left\{0, \min\left\{\min_{1 \leq r < d} B_{r,k},\ \min_{i \geq k} I_{d,i}\right\}\right\}. \tag{8.17}$$

For $a_k > G_d$, set $B_{d,k} = +\infty$. The second minimum in (8.17) is a suffix minimum. It accounts for every possible dual minimum $c \geq a_k$, rather than just the box starting at a_k.

The table update preserves nondecreasing entries in k: the boundary entries are nondecreasing by induction, suffix minima are nondecreasing, and taking a minimum and then a maximum with zero preserves this property. Thus the table also defines nondecreasing step-function lower bounds of the kind used in Section 6.

Theorem 8.2 (Soundness of the finite verification)**.** *Suppose that $G_d \geq \gamma_d$ for $1 \leq d \leq N$, that* (8.4) *holds for $r < N$, and that the procedure* (8.3)–(8.17) *passes* (8.16) *for each $2 \leq d \leq N$. Then every stable unimodular lattice of rank at most N satisfies $M(L) \leq \operatorname{rank}(L)/12$. Moreover, every table entry has the meaning* (8.2).

Proof. We prove both assertions by induction on d. They are immediate in rank one.

Assume they hold in all smaller ranks. Suppose first that the basic bound fails in rank d. Maximize M on $\mathcal{S}_d$. Lemma 6.1 gives an interior maximum with squared dual minimum $c > 2$. Section 6 supplies a rank $r < d$ and a scale t satisfying (6.13). Choose a box containing (c, t) with $i \geq q$ and $i \leq s_d$. Such a box exists because the grids cover the entire range, including endpoints.

None of the exclusions in (8.6) can discard this box. Nor can the child entry be $+\infty$, because the actual last factor is a stable lattice satisfying its threshold. Lemma 8.1 and (8.16) therefore give

$$d - 12M(L) \geq J_{d,r,i,j}/Q \geq I_{d,i}/Q \geq m_d/Q \geq 0.$$

This contradicts the choice of L. The basic bound is proved in dimension d.

We now justify (8.17). If $a_k > G_d$, the constrained set is empty by the Hermite bound. Otherwise assume it is nonempty and maximize M on it. If the maximum is on the boundary of stability, (6.14) and the lower-dimensional induction give a lower bound by one of the entries $B_{r,k}/Q$, with $r < d$. If it is interior, the constrained-direction lemma and the rest of Section 6 give actual parameters (r, c, t) with $c \geq a_k$. Their box has $i \geq k$, is not excluded, and has finite child entry. Lemma 8.1 gives a lower bound by the corresponding $I_{d,i}/Q$.

Thus, in either case, the inner minimum in (8.17) is a valid deficit lower bound. It cannot be $+\infty$ if the constrained set is nonempty, because one of the two cases supplies a finite bound. Finally, the basic theorem in rank d was proved in the preceding paragraph, so replacing the lower bound by its maximum with zero is legitimate. This proves (8.2) and completes the induction. □

The theorem isolates the computational problem from the geometric argument. To establish Theorem 1.2, it remains only to supply the Hermite witnesses and check the finite inequalities (8.16) through $N = 42$.

9 The verification through dimension 42

The two verifiers are `anc/verify.py` and `anc/verify_hermite.py`. Their input data are in `anc/hermite_witnesses.json`. The first program implements the induction in Section 8. It begins by running the second program on the polynomial witnesses. Thus the numerical bounds in the induction are recomputed from their exact certificates at every run.

The complete list of upper bounds and induction margins is given in Table 1. The entries g_d mean $G_d = g_d/1000$. The margin entries are the integers m_d in (8.16); the actual deficit lower bounds are $m_d/10^6$. A dash means that the minimum in the finite first-counterexample test is over no retained boxes, and hence has value $+\infty$.

With the mesh in (8.1), the program retains 901,245 candidate rank–box pairs after the first three exclusions in (8.6). For a fixed dimension it additionally applies the last exclusion and the empty-child test. The counts refer to intervals covering the real parameter space, not to sampled points. The final finite margin is

$$m_{42} = 45579 > 0. \tag{9.1}$$

All preceding dimensions pass as well. The positive values concern a hypothetical first-counterexample maximum, whose dual minimum exceeds 2. They do not assert a positive uniform deficit on the entire stable locus; the cubic lattice has deficit zero.

By Theorem 8.2, the successful checks prove Theorem 1.2. Section 2 then proves Theorem 1.1.

9.1 Reproducing and reading the computation

From the directory containing this paper's source, run

```
python anc/verify.py
```

Python 3.10 or later and the standard library suffice. The program requires assertions to be enabled and refuses to run under Python's `-O` option. It reads the rational witness file, checks all polynomial identities and signs, verifies (8.4), constructs the logarithm bounds and table entries, and tests (8.16) successively for $d = 2, \ldots, 42$. It stops with an error if any required check fails.

d	g_d	m_d	d	g_d	m_d
1	1000	base	22	3760	620285
2	1155	–	23	3881	621460
3	1305	–	24	4001	622512
4	1450	–	25	4121	619225
5	1591	–	26	4240	620143
6	1730	–	27	4359	620973
7	1866	–	28	4478	619382
8	2001	666625	29	4596	620119
9	2133	654916	30	4714	620794
10	2264	648000	31	4831	617600
11	2394	638129	32	4949	618204
12	2522	632985	33	5066	618772
13	2650	629074	34	5183	619308
14	2776	629378	35	5300	619814
15	2902	628005	36	5416	620292
16	3027	624872	37	5533	618301
17	3151	624164	38	5649	618729
18	3275	623013	39	5765	510954
19	3398	621176	40	5880	371266
20	3518	622680	41	5996	216387
21	3639	620254	42	6112	45579

Table 1: Exact Hermite bounds and the complete dimension induction. Every finite margin is nonnegative. The rank-one row is the initial condition.

The last lines of the output include

```
40 PASS 371266 /1000000
41 PASS 216387 /1000000
42 PASS 45579 /1000000
Max dimension verified: 42
```

The output file `anc/verification_result.json` records all rational Hermite bounds, polynomial sign checks, dimension margins, mesh denominators, and the candidate count. This output is not used as an input to the proof: it can be deleted and regenerated.

The implementation stores the table entries and margins as integers in units of $1/Q$. It uses the integer 10^{30} as a marker for $+\infty$, testing for the marker before arithmetic with a child entry. Every finite deficit-table entry is checked to lie between 0 and dQ, and the finite box computations are far below the marker. The code also checks that each row of deficit bounds is nondecreasing. Python's arbitrary-precision integers avoid overflow. The only floating-point quantity reported by the program is elapsed time; no proof decision depends on it.

For correspondence with the mathematical specification, the variable `cs` is q, `ts` is p, and `logs[i]` is L_i from (8.9). The variable `ia` is the child index in (8.7). The two endpoint integers in (8.12) are `base_left` and `base_right`; `lower` is the box score (8.13). The array `interior` is $I_{d,i}$, and the array `D` contains $B_{d,i}$. The proof of correctness is Theorem 8.2, together with the interval estimate in Lemma 8.1.

The supplied polynomial data and verifier files, identified by checksums, are part of the proof. A table of decimal Hermite bounds or positive margins alone would not establish the finite assertions.

Funding acknowledgments

The following research support is acknowledged: Israel Science Foundation grants 2019/19 (Barak Weiss), 2878/25 (Or Ordentlich), and 674/24 (Boaz Klartag), and the joint ISF–NSFC grant 3739/21 (Barak Weiss).

Statement on the use of artificial intelligence

This manuscript was written by ChatGPT, in a research dialogue directed by Boaz Klartag, following discussions with Or Ordentlich and Barak Weiss. ChatGPT also developed the additional mathematical arguments presented here and wrote the accompanying verification software and polynomial-certificate data. The Python program verifies the stated finite arithmetic assertions; the geometric and analytic arguments are written proofs, not proof-assistant formalizations. We have not yet checked this manuscript. No independent human verification is asserted by this statement.

References

[1] P. Autissier, *Un lemme matriciel effectif*, Math. Z. **273** (2013), no. 1–2, 355–361. doi:10.1007/s00209-012-1008-x.

[2] H. Cohn and N. Elkies, *New upper bounds on sphere packings I*, Ann. of Math. (2) **157** (2003), no. 2, 689–714. doi:10.4007/annals.2003.157.689.

[3] A. Magazinov, *A proof of a conjecture by Haviv, Lyubashevsky and Regev on the second moment of a lattice Voronoi cell*, Adv. Geom. **20** (2020), no. 1, 117–120. doi:10.1515/advgeom-2018-0018.

[4] O. Regev and N. Stephens-Davidowitz, *A reverse Minkowski theorem*, Ann. of Math. (2) **199** (2024), no. 1, 1–49. doi:10.4007/annals.2024.199.1.1.

[5] U. Shapira and B. Weiss, *Stable lattices and the diagonal group*, J. Eur. Math. Soc. **18** (2016), no. 8, 1753–1767. doi:10.4171/JEMS/628.