The isodual constant of the LLL potential: exact values and structure

Nick Howgrave-Graham
nickhg@gmail.com

Preprint, 9 November 2026

Abstract

The LLL potential $\Psi^2(B) = \prod_k \det \mathrm{Gram}(b_1, \ldots, b_k)$, minimized over the bases of a lattice L and normalized by its determinant, defines an isodual lattice invariant $\eta(L)$. Its supremum η_n over n-dimensional lattices, introduced in [17], is the staircase case of the generalized Hermite constants of Watanabe and Meyer, for which explicit values were known only for $n \le 4$. We prove that $\eta_5 = 6$, attained only by D_5 and its dual: the flag analogue of Korkine and Zolotarev's value of γ_5. The natural reduction to four-dimensional lattices, through a function Φ_5, does not suffice: the global maximum of Φ_5 is $2^{16}/(3^4 \cdot 5^3) \approx 6.4727$, attained not at D_4 but at an even 135-modular lattice M_*, which we have not found in the literature. A second stage, over the lifts of the remaining region, completes the proof; both stages are computer-assisted, with about 9×10^9 leaves, each re-proved in interval arithmetic by a separately written checker; code and data are public. We also prove $\eta_8 = 2304$, attained by E_8, and show that E_6, E_7, E_8 and their duals are strict local maxima of η (for $n = 6, 7$ these are the conjectured extremal lattices). Asymptotically, $\log \eta_n = (\frac{1}{4} + o(1))\, n^2 \log n$.

For lattice reduction we introduce the Ψ-Hermite constant μ_m, the largest first Gram–Schmidt norm in a potential-minimal basis of an m-dimensional lattice, and bound it through η_m. This gives $\mu_m = \gamma_m$, Hermite's constant, for $m \le 5$, and $\mu_m \le \exp(O(\log^2 m))$ in general.

1 Introduction

The potential $\Psi^2(B) = \prod_k G_k$, the product of the prefix Gram determinants of a basis, is the quantity whose decrease proves that LLL terminates. (Only the square enters; we keep it, as in [17].) Minimized over all bases of a lattice L, and normalized to be invariant under scaling, it becomes a lattice invariant $\eta(L)$. This invariant is *isodual*: $\eta(L) = \eta(L^*)$. Its supremum η_n over n-dimensional lattices is a Hermite-type constant for the whole flag of a basis, rather than for its first vector. Hermite's constant γ_n is known only for $n \le 8$ and $n = 24$ [6]; in dimension 5 it is due to Korkine and Zolotarev [24], attained only by D_5. Our main result is the flag analogue in the same dimension: $\eta_5 = 6$, attained only by D_5 and its dual D_5^*, the two ends of a flag playing symmetric roles.

The constant was defined, and $\eta_2 = \gamma_2$ and $\eta_3 = 3/2$ computed, in ePrint 2007/105 (March 2007) [17]. It is the staircase case $\gamma_{n,(n-1,\ldots,1)}$ of the generalized Hermite constants of Watanabe [46] and Meyer [30]; Meyer's preprint (April 2008) computes the same value 3/2 independently, and Meyer's thesis [29, Prop. 4.16] also obtains $\eta_4 = 3$, attained by D_4, by the same Mordell-type argument as ours. Meyer observes that very little is known about explicit values of these constants [31]; we are not aware of any explicit value for $n \ge 5$.

The two main theorems are the value of η_5 and the lattice that makes it hard to prove.

Theorem A (Theorem 14). $\eta_5 = 6$. A five-dimensional lattice L has $\eta(L) = 6$ if and only if L is similar to D_5 or to D_5^*.

For a lattice M of rank $n-1$ and determinant 1, let $\kappa(M) = \sup_\varphi \min_{s \in M \setminus 0} |s|^2/(1 - \|\varphi(s)\|^2)$, where φ ranges over linear forms on M modulo integer-valued ones (the ways a lattice can be lifted over M) and $\|x\|$ is the distance from x to the nearest integer (§5.1). Let L have rank n, and let v be a shortest vector of L (any choice). Proposition 10 shows $\eta(L) \le \Phi_n(\pi_v L)$ with $\Phi_n(M) = \kappa(M)^{(n-1)/2}\eta(M)$, where $\pi_v L$ is rescaled to determinant 1; so $\eta_5 \le \sup_M \Phi_5(M)$ over four-dimensional M.

Theorem B (Theorem 20, Proposition 17). The global maximum of Φ_5 is $2^{16}/(3^4 \cdot 5^3) \approx 6.4727$, attained only at lattices similar to the lattice M_* of Definition 16. M_* is an even 135-modular lattice, the unique densest in its genus and a Hermitian lattice of rank 2 over $\mathbb{Z}[\omega]$, with $|\operatorname{Aut}(M_*)| = 72$. In particular, no bound on η_5 through Φ_5 alone can reach 6.

The obvious candidate for the maximum of Φ_5 is D_4, the projection of D_5^* along a shortest vector, with $\Phi_5(D_4) \approx 6.348$, already above 6. But the landscape rises past D_4: along the isodual geodesic from D_4 to $A_2 \otimes A_2$, on which M_* lies (§6), Φ_5 climbs to M_* before falling (Figure 1). So the proof of Theorem A needs a second stage, which bounds the lifts over the region where $\Phi_5 \ge 6$ directly.

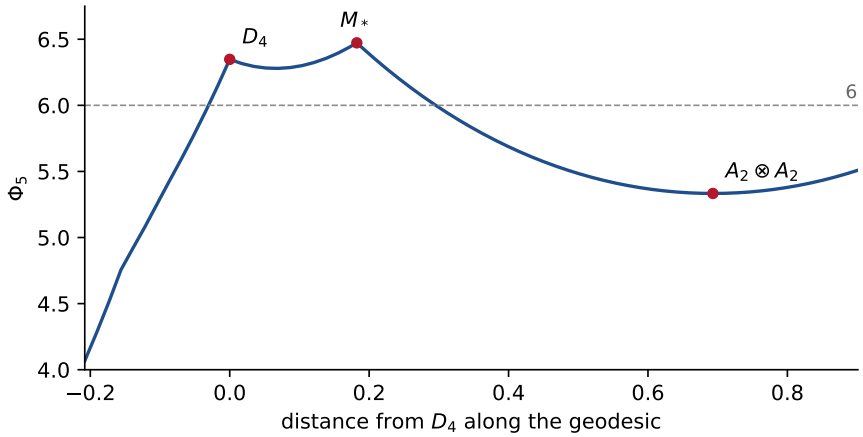


Figure 1: Φ_5 along the geodesic of isodual lattices from D_4 through M_* to $A_2 \otimes A_2$, against the distance from D_4 in the affine-invariant metric on Gram matrices (the metric d of Theorem 13): M_* is at $\ln \frac{6}{5}$ and $A_2 \otimes A_2$ at $\ln 2$; negative values continue the geodesic beyond D_4. The dashed line is $\Phi_5 = 6$, the value of η_5; a proof of $\eta_5 = 6$ through Φ_5 alone would need the whole curve, and every other lattice, below it. The curve is numerical (computed as described in §6); the three marked values, $3(36\sqrt{2}/35)^2$, $2^{16}/(3^4 \cdot 5^3)$ and $\frac{16}{3}$, are exact. Φ_5 is symmetric about $A_2 \otimes A_2$ along the geodesic, where it has a local minimum.

This paper makes two contributions, linked by one inequality: η_m controls the quality constant μ_m of potential-minimal bases (Theorem 35).

1. **The constants** (§3–8). Besides Theorems A and B, we prove $\eta_8 = 2304$ (and that D_4 is the only lattice attaining Meyer's value $\eta_4 = 3$), compute η for all irreducible root lattices, and observe that for $n = 3, \dots, 8$ the critical lattice has a flag that is extremal in every rank simultaneously. In the proof of Theorem A, a strict local maximality theorem for D_5 and

D_5^* with an explicit radius supplies the equality case. A Coxeter-number lemma bounds lifts over root lattices and exposes an inductive structure: for $n = 6, 7, 8$, lifts over the previous critical lattice cannot be extremal, and dimension 5 is the exceptional case (Proposition 28). Asymptotically, $\log \eta_n = (\frac{1}{4} + o(1))\, n^2 \log n$ (§8).

2. **Consequences for reduction** (§9). Potential-based algorithms (PotLLL, DeepBKZ, PotBKZ) come with termination guarantees. For PotLLL, whose output is LLL-reduced, the known quality bound is LLL's, and we know of no bound for these algorithms that reflects the potential condition itself. We introduce the Ψ-Hermite constant μ_m, the largest first Gram–Schmidt norm of a potential-minimal basis, and a key inequality relating it to η_m, which gives $\mu_m = \gamma_m$ for $m \leq 5$. For large m the bound is only quasi-polynomial.

What is proved. Every theorem, proposition, lemma and corollary is proved. Those whose proofs are marked "computer-assisted" (Theorems 13, 19, 20 and 14, and Propositions 21 and 22) rest on exact rational computations or on logged searches in which every leaf is re-proved by a separately written checker in interval arithmetic; the exact values computed in Proposition 17 and Remark 18, and the identities used for Proposition 26, rest on exact rational computations. Numerical observations are stated as such and are not used in any proof. For $n = 6, 7$ we prove $H_n \leq \eta_n \leq U_n$, where $H_6 = 128\sqrt{3}/9$ and $H_7 = 144$ are attained by E_n and E_n^*, and $U_6 = 16\sqrt{3} \approx 27.71$ and $U_7 = 16\sqrt{3} \cdot 3^{7/4} \approx 189.5$ (Table 3). We also prove that E_n and E_n^* are strict local maxima of η with an explicit radius (Proposition 26): every lattice within distance $1/10$ of E_6 or E_6^*, or within $3/50$ of E_7 or E_7^*, has $\eta \leq H_n$, with equality only at those lattices (distance in the metric d of §5.2, after scaling). The same proposition shows that E_8 is a strict local maximum, within distance $1/20$. What remains open is the global statement $\eta_n = H_n$ (Conjecture 25), that no lattice outside these neighbourhoods does better; for $n = 6$ there is numerical evidence (§7). Also open are Conjecture 30, with numerical support, and the global uniqueness of E_8 as the lattice attaining η_8.

Reading guide. Readers interested in the constants can skip §9, except Theorem 35 and Corollary 36 for $m \leq 4$, which the computer-assisted proofs use. Readers interested in reduction need §2, Lemma 3, the exact values of §4, and Theorem 14 with Corollary 36; the rest of §5, §§6–7 and §8 can be skipped.

Related work.

- *Generalized Hermite constants.* Watanabe [46] and Meyer [30] define $\gamma_{n,\lambda}$ for each partition λ. For the staircase $\lambda = (n-1, \ldots, 1)$ Meyer's height is $\prod_{k<n} \det L_k$, so $\eta_n = \gamma_{n,(n-1,\ldots,1)}$. Equivalently, Ψ^2 is the height of the complete flag of a basis in the sense of Thunder and Kim [43, 20]. Meyer [31, Thm. 5.1] shows, via designs, that D_4, E_6, E_7, E_8, Leech, Barnes–Wall and Thompson–Smith are extreme (local maxima) for partitions satisfying $\lfloor \lambda \rfloor < n/2$, where $\lfloor \lambda \rfloor$ denotes the number of parts of λ. The staircase partition has $n - 1 \geq n/2$ parts, so this does not cover η_n. Meyer also develops the Voronoi theory of these constants [30, §3]: a form is extreme for λ if and only if it is perfect and eutactic for λ. His thesis [29] computes several of them by the Mordell inequality, among them $\eta_4 = 3$. Our local results at D_5 (Theorem 13) are a quantitative version of this theory, with an explicit radius. Rankin's constants [37] entered lattice reduction in [11].

- *Isoduality.* Symplectic lattices, including NTRU lattices, are isodual, and [12] exploits this for reduction. Since $\eta(L) = \eta(L^*)$, isodual lattices are natural candidates for maximizers of

η. Isodual lattices were studied by Conway and Sloane [7], following Buser and Sarnak [5] and Bergé and Martinet [2]. The lattices around M_* turn out to be symplectic as well (§6).

- *Potential-based reduction.* PotLLL [10] later used the same potential, from the original LLL analysis [27], to control deep insertions in LLL. DeepBKZ [47], self-dual DeepBKZ [48] and PotBKZ with a self-dual variant [39] extend this to blocks, with provable termination. Self-dual BKZ [32] alternates primal and dual SVP tours, and slide reduction [13] combines primal and dual SVP oracles with a provable Hermite factor.

- *Random lattices and enumeration.* Laarhoven [26] studies the minimal enumeration-tree exponent over bases of Haar-random lattices. His upper bounds split along a dense half-rank sublattice and reduce the two halves by dual HKZ and HKZ, which is the construction of Lemma 3; the idea goes back to [11]. His lower bounds count complete flags using Kim's formula [20], and he notes that products of covolumes along flags occur in the theory of generalized Hermite constants. Theorem 31 shows that the random-lattice value of this product and its extremal value η_n agree to leading order.

- Our contribution to reduction is a structural one: the potential condition loses nothing against Hermite's constant for blocks of size at most 5 (Corollary 36). We do not propose a new algorithm.

2 Preliminaries

Lattices are described by Gram matrices and all lengths are *squared*. For a basis $b_1, \dots, b_n$ of L, let $L_k = \operatorname{span}_{\mathbb{Z}}(b_1, \dots, b_k)$, let $G_k = \det \operatorname{Gram}(L_k)$, let $\Delta = G_n$ (the Gram determinant of L), and let $r_i = \|b_i^*\|^2$. The LLL potential is

$$\Psi^2(B) = \prod_{k=1}^{n} G_k = \prod_{i=1}^{n} r_i^{n+1-i}.$$

Definition 1. $\eta(L) = \min_B \Psi^2(B)/\Delta^{(n+1)/2}$, and $\eta_n = \sup_L \eta(L)$. A basis attaining the minimum is Ψ-*minimal*.

- *Isoduality.* For the reversed dual basis, $G_k^* = G_{n-k}/\Delta$, so $\Psi^2(B^*) = \Psi^2(B)/\Delta^{n+1}$ and $\eta(L) = \eta(L^*)$. The reversed dual of a Ψ-minimal basis is Ψ-minimal.

- *Heredity.* Every contiguous projected block $\pi_{i-1}(b_i, \dots, b_j)$ of a Ψ-minimal basis is Ψ-minimal: changing the basis of that block alters only $G_i, \dots, G_{j-1}$, and it alters them by exactly the block's own potential.

- *Small cases.* For $n = 2$, $\eta(L) = \lambda_1^2/\sqrt{\Delta}$, so $\eta_2 = \gamma_2 = 2/\sqrt{3}$.

Let γ_n be Hermite's constant and $\gamma(L) = \lambda_1^2/\Delta^{1/n}$. Let $\gamma_{n,k}$ be Rankin's constant (squared normalization, as in [40]), the maximum over L of $\min_M \det(M)/\Delta^{k/n}$ over rank-k sublattices M, and let $\gamma'_{n,k}$ be the Bergé–Martinet constant [2], the supremum over L of $\sqrt{\gamma_{n,k}(L)\gamma_{n,k}(L^*)}$ (for $k = 1$ this is their γ'_n, tabulated with lower bounds in [7, Table I]). Every rank-k sublattice of a lattice with minimum λ satisfies the *Hermite bound* $\det M \geq \lambda^k/\gamma_k^k$.

Table 1 lists the constants we import, in the squared normalization used throughout (lengths are squared, $\gamma_n = \max \lambda_1^2/\Delta^{2/n}$). Meyer's normalization of $\gamma_{n,\lambda}$ agrees with ours: his value for the staircase in dimension 4 is $3 = \eta_4$ [29, Prop. 4.16].

constant	value	source	used in
γ_2, γ_3	$2/\sqrt{3}$, $2^{1/3}$	Lagrange, Gauss; see [8, Ch. 1]	$\eta_2 = \gamma_2$, Cor. 36
γ_4	$\sqrt{2}$, only D_4	[23], see [8, Ch. 1]	Thm. 7, App. C
γ_5	$8^{1/5}$, only D_5	[24]	Thm. 14 (range of r)
$\gamma_6, \gamma_7, \gamma_8$	$(64/3)^{1/6}$, $64^{1/7}$, 2	Blichfeldt; see [8, Ch. 1]; only E_6, E_7, E_8 [45]	H_n, Cor. 36, Prop. 26
γ_m, all m	$\le (4/\pi)\Gamma(1+\frac{m}{2})^{2/m} \le 1+\frac{m}{4}$	Minkowski; see [8, Ch. 1]	Thm. 31
density Δ_m, large m	$\le 2^{-0.599m(1+o(1))}$	[19]	Cor. 32
$\gamma_{8,4}$	4	[40, Thm. 3.1]	Thm. 8
γ'_5, γ'_7	$\sqrt{2}$, $\sqrt{3}$	[36]	Table 3
$\gamma'_{6,2}$	2	[40]	Table 3
η_3	$\frac{3}{2}$	Prop. 6 ([17], [30])	Thm. 7, App. E
η_4	3	Thm. 7 ([29])	Thm. 35, Apps. C–E
(K1), (K2)	flag counts	[20, Thms. 3, 5]	Thm. 31

Table 1: Imported constants (squared normalization), each proved in the cited source.

Lemma 2 ([16, Lemma 3.2.5]; [35, Lemma 1])**.** *Let $M \subseteq L$ have rank k. Then there are indices $t_1 < \cdots < t_k$ with $\det(M) \ge \prod_i r_{t_i}$ and $M \subseteq \operatorname{span}(b_1, \ldots, b_{t_k})$.*

Proof. Write a basis $d_1, \ldots, d_k$ of M in the coordinates of B. Integer row operations triangularize it, so the last nonzero coordinate of d_i is at index t_i, with $t_1 < \cdots < t_k$ and nonzero integer pivots.

The projection of d_i orthogonal to $d_1, \ldots, d_{i-1}$ has squared length at least that of its projection orthogonal to $b_1, \ldots, b_{t_i-1}$. That second projection is a nonzero integer multiple of $b^*_{t_i}$, so its squared length is at least r_{t_i}. Multiplying over i gives the result. □

This appears in the author's thesis [16, p. 47] and was found independently by Pataki and Tural [35]. Under their name it underlies the overstretched-NTRU analyses of [21] and [9].

Exact Ψ-minimization. For an m-dimensional lattice we use the recursion $\Psi^2_{\min}(L) = \min_v |v|^{2m}\Psi^2_{\min}(\pi_v L)$ over primitive v with $|v|^2 \le \mu_m \Delta^{1/m}$ (μ_m as in §9; the proved values for $m \le 5$ are Theorem 35 and its corollary, and for $m \ge 6$, where μ_m is not known, we use the bound $\mu_m \le (4/3)^{(m-1)/2}$ of §9). The search is seeded with an LLL basis. A branch is pruned when $|v|^{2m}$ times the lower bound $\prod_k \big(\prod \text{of the } k \text{ smallest } r_i(\pi_v L)\big)$ exceeds the best value found; that lower bound is Lemma 2. On random lattices of dimension 6–8 this searches 8–100 times fewer nodes. The exact values of η computed for particular lattices in this paper, such as $\eta(M_*)$ in Proposition 17, use this recursion.

3 A splitting inequality

Lemma 3. *For $1 \le k < n$ (with $\eta_1 := 1$), $\eta_n \le \eta_k\, \eta_{n-k}\, (\gamma'_{n,k})^{n/2} \le \eta_k\, \eta_{n-k}\, \gamma_{n,k}^{n/2}$. More precisely, $\eta(L) \le \eta_k \eta_{n-k} \big(\det M/\Delta^{k/n}\big)^{n/2}$ for every rank-k sublattice $M \subseteq L$.*

Proof. We may take M primitive, of determinant D. Let π be the projection orthogonal to M, so that πL has rank $n-k$ and determinant Δ/D.

Take a Ψ-minimal basis of M and append lifts of a Ψ-minimal basis of πL. The result is a basis of L, and its first k prefix determinants are those of M's basis. For $j > k$ the prefix determinants

are D times those of πL's basis. Hence

$$\Psi^2 \le \eta_k D^{(k+1)/2} \cdot D^{n-k} \cdot \eta_{n-k} (\Delta/D)^{(n-k+1)/2} = \eta_k \eta_{n-k}\, D^{n/2} \Delta^{(n-k+1)/2}.$$

Dividing by $\Delta^{(n+1)/2}$ gives the second statement, and minimizing over M gives $\eta(L) \le \eta_k \eta_{n-k} \gamma_{n,k}(L)^{n/2}$. The same bound holds for L^*, and $\eta(L) = \eta(L^*)$. Taking the geometric mean gives the γ' form. □

For $k = 1$ (equivalently $k = n-1$, by duality) the lemma is a Mordell-type recursion, and for the staircase it is exactly [30, Prop. 40] with $m = n-1$, with the same constant; that proposition bounds $\gamma_{n,\lambda} \le \gamma_{m,\lambda} \gamma_{n,m}^{|\lambda|/m}$. We did not find the general-k product form or the Bergé–Martinet refinement in the literature.

4 Exact values

4.1 Root lattices

For E_6, E_7, E_8 the root flag meets the Hermite bound in every rank (e.g. $G_5(E_6) = 4 = 32/8$). For A_n and D_n it does not in general (e.g. $G_4(A_4) = 5 > 4$, $G_6(D_7) = 4 > 3$), and the following lemma handles these families for all n.

Lemma 4. *Let M be a sublattice of rank k. If $M \subseteq A_n$ then $\det M \ge k+1$. If $M \subseteq D_n$ then $\det M \ge \min(k+1, 4)$.*

Proof. Write a basis of M as the rows of a $k \times N$ integer matrix with columns $c_1, \dots, c_N$ (with $N = n+1$ for $A_n \subset \mathbb{Z}^{n+1}$ and $N = n$ for $D_n \subset \mathbb{Z}^n$). Let p_S denote its $k \times k$ minors. By Cauchy–Binet, $\det M = \sum_S p_S^2$.

For A_n we have $\sum_j c_j = 0$. Fix S with $p_S \ne 0$ and $s \in S$. Substituting $c_s = -\sum_{j \ne s} c_j$ into column s of the minor gives $p_S = \mp \sum_{j \notin S} \pm p_{(S \setminus s) \cup j}$, since the terms with $j \in S$ vanish. So for each $s \in S$ some $p_{(S\setminus s)\cup j}$ is nonzero. These k sets are distinct (each determines s) and differ from S. Hence at least $k+1$ minors are nonzero integers, and $\det M \ge k+1$.

For D_n we have $\sum_j c_j \equiv 0 \pmod 2$, so the identity holds modulo 2. If some p_S is odd, there are at least $k+1$ odd minors and $\det M \ge k+1$. Otherwise all minors are even and at least one is nonzero, so $\det M \ge 4$. □

Theorem 5. *The values in Table 2 hold. In each case the root flag is Ψ-minimal, because each of its prefix determinants is individually minimal (Lemma 4 for A_n, D_n; the Hermite bound for E_6, E_7, E_8).*

L	$\eta(L)$	Ψ-minimal squared profile r_i
A_n	$(n+1)!/(n+1)^{(n+1)/2}$	$2, \frac{3}{2}, \frac{4}{3}, \dots, \frac{n+1}{n}$
D_n $(n \ge 4)$	$3 \cdot 2^{n-4}$	$2, \frac{3}{2}, \frac{4}{3}, 1, \dots, 1$
E_6	$128\sqrt{3}/9 \approx 24.63$	$2, \frac{3}{2}, \frac{4}{3}, 1, 1, \frac{3}{4}$
E_7	144	$2, \frac{3}{2}, \frac{4}{3}, 1, 1, \frac{3}{4}, \frac{2}{3}$
E_8	2304	$2, \frac{3}{2}, \frac{4}{3}, 1, 1, \frac{3}{4}, \frac{2}{3}, \frac{1}{2}$

Table 2: η of root lattices. For $n \le 7$ the values were also confirmed by exact enumeration of Ψ-minimal bases (for $n = 6, 7$ with the bound $\mu_n \le (4/3)^{(n-1)/2}$).

4.2 Dimensions 3, 4 and 8

Proposition 6 ([17, Thm. 1], [30, Prop. 47])**.** $\eta_3 = \frac{3}{2}$*, attained only by* A_3 *and* A_3^* *(see also [29, Prop. 4.14]).*

Theorem 7. $\eta_4 = 3$ *[29, Prop. 4.16], attained only by* D_4*.*

Proof. Lemma 3 with $k = 1$ and Proposition 6 give $\eta_4 \le \eta_3\gamma_4^2 = \frac{3}{2} \cdot 2 = 3$, and $\eta(D_4) = 3$ by Theorem 5. Equality forces $\gamma(L) = \gamma_4$, so L is critical, and D_4 is the unique critical lattice in dimension 4 [24]. □

The value and the attainment are due to Meyer [29, Prop. 4.16] ("atteinte sur D_4"), with the same proof. Uniqueness, which Meyer does not state, follows from the same argument, since equality forces $\gamma(L) = \gamma_4$ and D_4 is the unique critical lattice.

Theorem 8. $\eta_8 = 2304$*, attained by* E_8*.*

Proof. Lemma 3 with $k = 4$, together with $\gamma_{8,4} = 4$ [40, Thm. 3.1] (in the squared normalization; it is 2 in the covolume normalization), gives $\eta_8 \le \eta_4^2 \cdot 4^4 = 2304$, and $\eta(E_8) = 2304$ by Theorem 5. □

Uniqueness of E_8 is open (§10).

4.3 Flags that are extremal in every rank

Let $H_m = \gamma_m^{m(m-1)/2} / \prod_{k<m} \gamma_k^k$. For a critical lattice, normalized so that $\lambda_1^2 = \gamma_m$ and $\Delta = 1$, the Hermite bound gives $G_k \ge \gamma_m^k/\gamma_k^k$ for every flag, so $\eta \ge H_m$, and $\eta_m \ge H_m$ always. One checks that

$$H_3, \dots, H_8 = \tfrac{3}{2},\ 3,\ 6,\ \tfrac{128\sqrt{3}}{9},\ 144,\ 2304,$$

exactly η of $A_3, D_4, D_5, E_6, E_7, E_8$. So in each of these dimensions the critical lattice has a flag meeting the Hermite bound in every rank simultaneously (cf. [29, §4.2.1.2], where the same nested minimal sections are used for D_4, E_6 and E_8).

The dual-side split is tight. Lemma 3 with $k = n - 1$, applied to a single lattice, gives $\eta(L) \le \eta_{n-1}\,\gamma(L^*)^{n/2}$. At the critical lattices it is tight in the following sense: $H_{n-1}\gamma(L_n^*)^{n/2} = \eta(L_n)$, that is, it equals 3, 6, $128\sqrt{3}/9$, 144, 2304 for $L_n = D_4, D_5, E_6, E_7, E_8$. For $n \le 6$, where $\eta_{n-1} = H_{n-1}$ is known, this is the bound itself; for $n = 7, 8$ it is the bound with η_{n-1} replaced by the conjectured value H_{n-1}. Equivalently $\gamma(L_n^*) = (\gamma_n/\gamma_{n-1})^{n-1}$: the root flag meets the Hermite bound in rank $n-1$. Consequently, if $\eta_{n-1} = H_{n-1}$, a lattice with $\eta(L) > H_n$ must have both $\gamma(L)$ and $\gamma(L^*)$ above $\gamma(L_n^*) = (H_n/H_{n-1})^{2/n}$. Similarly, the $k = 1$ split applied along a shortest vector v gives $\eta(L) \le \gamma(L)^{n/2}\eta(\pi_v L)$, with equality at the root lattices, so a counterexample also needs $\eta(\pi_v L) > H_n/\gamma_n^{n/2}$ (e.g. $3/\sqrt{2}$ for $n = 5$, $16/3$ for $n = 6$).

5 Dimension 5

This section proves $\eta_5 = 6$. Lifts along a shortest vector reduce $\eta_{\ell+1}$ to a function of ℓ-dimensional lattices (§5.1); D_5 and D_5^* are strict local maxima of η with an explicit radius (§5.2); and a two-stage computer-assisted search gives the theorem (§5.3). The landscape of the auxiliary function Φ_5, which explains why a second stage is needed, is in §6, and dimensions 6 and 7 are in §7. Numerical observations are stated as such and are not used in any proof.

5.1 Lifts over root lattices

By the dual-side split paragraph of §4, a lattice with $\eta(L) > c$ needs $\eta(\pi_v L) > c/\gamma_5^{5/2}$ for a shortest vector v. For c close to $\eta_4\gamma_5^{5/2}$ this forces $\pi_v L$ to be close to D_4. The following shows that $\pi_v L \cong D_4$ itself is incompatible with v being long.

For $x \in \mathbb{R}$ write $\|x\|$ for the distance from x to the nearest integer.

Lemma 9. *Let R be an irreducible simply-laced root system with Coxeter number h. For every linear form φ on $\mathbb{R}R$ there is a root α with $\|\varphi(\alpha)\| \le 1/h$. The bound is attained: for $\varphi = \rho^\vee/h$, which takes the value $\mathrm{ht}(\alpha)/h$ on each positive root α, one has $\min_\alpha \|\varphi(\alpha)\| = 1/h$.*

Proof. Replacing φ by $\varphi\circ w$ (w in the Weyl group W) or by $\varphi+\lambda$ (with λ integer-valued on R) does not change the multiset $\{\|\varphi(\alpha)\|\}$. Every form is equivalent under these operations to one in the closed fundamental alcove [18, §4.3], that is, with $a_i := \varphi(\alpha_i) \ge 0$ for the simple roots $\alpha_1,\dots,\alpha_\ell$ and $a_0 := 1-\varphi(\theta) \ge 0$, where $\theta = \sum_i n_i\alpha_i$ is the highest root. Then $\|\varphi(\alpha_i)\| \le a_i$ and $\|\varphi(\theta)\| \le a_0$. With $n_0 = 1$ we have $\sum_{i=0}^{\ell} n_i a_i = 1$ and $\sum_{i=0}^{\ell} n_i = h$ [3, Ch. VI, §1.11, Prop. 31], so $\min_i a_i \le 1/h$. For $\varphi = \rho^\vee/h$, every a_i equals $1/h$, and the heights of positive roots lie in $[1, h-1]$. □

The bound is implicit in the theory of the principal element [25]; we have not found it stated in this form. For D_4, with roots $\pm e_i \pm e_j$, the form is $\varphi(x) = w\cdot x$ with $w = (0, \frac16, \frac13, \frac12)$, and the lemma reduces to a pigeonhole: the five alcove values $a_0,\dots,a_4$ satisfy $a_0 + a_1 + 2a_2 + a_3 + a_4 = 1$, so the smallest of them, t, has $6t \le 1$, that is $5t \le 1-t$. The lemma and its proof hold for every irreducible root system, since $\sum_{i\ge0} n_i = h$ in general and $\rho^\vee$ takes the value 1 on each simple root; only Corollary 11, which uses that the roots are the minimal vectors, needs R simply laced.

For a lattice M of rank ℓ with $\det M = 1$ define

$$\kappa(M) = \sup_{\varphi} \min_{s\in M\setminus 0} \frac{|s|^2}{1-\|\varphi(s)\|^2},$$

the supremum over linear forms φ on $M\otimes\mathbb{R}$ (modulo integer-valued forms, so κ is $GL_\ell(\mathbb{Z})$-invariant). Since $0 \le \|\varphi(s)\| \le \frac12$, taking s minimal gives $\kappa(M) \le \frac43\lambda_1(M)^2$, and any φ gives $\kappa(M) \ge \lambda_1(M)^2$.

Proposition 10. *Let L have rank $\ell+1$ and $\det L = 1$, let v be any shortest vector, and let M be $\pi_v L$ rescaled to determinant 1. Then $\gamma(L)^{(\ell+1)/\ell} \le \kappa(M)$ and $\eta(L) \le \kappa(M)^{\ell/2}\,\eta(M)$. Consequently*

$$\eta_{\ell+1} \le \sup_M\ \kappa(M)^{\ell/2}\,\eta(M),$$

the supremum over lattices M of rank ℓ and determinant 1.

Proof. Put $r = |v|^2 = \gamma(L)$. Then $\det \pi_v L = 1/r$, so $|s|_M^2 = r^{1/\ell}|s|_{\pi_v L}^2$. Each $s \in \pi_v L$ lifts to $s + \varphi(s)v \in L$ for a linear form φ, well defined modulo integer-valued forms, and the lattice vectors over s have norms $|s|_{\pi_v L}^2 + r(j+\varphi(s))^2$, $j\in\mathbb{Z}$. Since v is shortest, $|s|_{\pi_v L}^2 \ge r(1-\|\varphi(s)\|^2)$, that is $r^{(\ell+1)/\ell} \le |s|_M^2/(1-\|\varphi(s)\|^2)$, for every $s \ne 0$; so $r^{(\ell+1)/\ell} \le \kappa(M)$. The $k=1$ split along v gives $\eta(L) \le r^{(\ell+1)/2}\eta(\pi_v L) = \big(r^{(\ell+1)/\ell}\big)^{\ell/2}\eta(M)$. □

Different shortest vectors may give non-isometric projections; the bound holds for each. For $\ell = 4$ this reads $\eta_5 \le \sup_M \kappa(M)^2\eta(M)$: the five-dimensional problem becomes a single $GL_4(\mathbb{Z})$-invariant inequality on four-dimensional lattices. Since $\kappa \le \frac43\lambda_1^2 \le \frac43\gamma_4$, the trivial bound is $\eta_5 \le \frac{32}{9}\cdot 3 \approx 10.7$; the content is in bounding κ near D_4.

Corollary 11. *Let L have rank $\ell+1$ and let v be a shortest vector. Suppose $\pi_v L$ is similar to an irreducible simply-laced root lattice R of rank ℓ with Coxeter number h. Then*

$$\gamma(L)^{(\ell+1)/\ell} \le \frac{h^2}{h^2-1}\,\gamma(R), \qquad \eta(L) \le \gamma(L)^{(\ell+1)/2}\,\eta(R).$$

In particular:

- *for $R = D_4$: $\gamma(L) \le (36\sqrt{2}/35)^{4/5} \approx 1.3496$ and $\eta(L) \le 3\,(36\sqrt{2}/35)^2 \approx 6.35$;*
- *for $R = A_4$: $\gamma(L) \le \left(\frac{25}{24}\cdot 2\cdot 5^{-1/4}\right)^{4/5} \approx 1.304$ and $\eta(L) \le 4.17$.*

Proof. At determinant 1 the minimum of R is $\gamma(R)$, and its minimal vectors are the roots. By Lemma 9, for every φ some root has $\|\varphi(\alpha)\| \le 1/h$, so $\kappa(R) \le \gamma(R)/(1-h^{-2})$. Apply Proposition 10. □

Proposition 12 (κ-maximizers of root lattices)**.** *Let R be an irreducible simply-laced root lattice of rank ℓ with Weyl group W and Coxeter number h, let $P^\vee$ be the lattice of integer-valued forms on R and $Q^\vee \subseteq P^\vee$ the coroot lattice. Then $\kappa(R) = \gamma(R)\,h^2/(h^2-1)$, and the forms attaining it are exactly the $w\rho^\vee/h$, $w \in W$, modulo $P^\vee$. They are torsion points of order exactly h, and form $|W|/|P^\vee/Q^\vee|$ classes: $192/4 = 48$ for D_4 and $120/5 = 24$ for A_4, or 24 and 12 up to sign.*

Proof. By Corollary 11 and its proof, $\kappa(R) \le \gamma(R)\,h^2/(h^2-1)$. At $\varphi = \rho^\vee/h$ every root has $\|\varphi(\alpha)\| \ge 1/h$, and non-roots have $|s|^2 \ge 2\gamma(R) > \gamma(R)h^2/(h^2-1)$, so equality holds. For the maximizers, in the proof of Lemma 9 a form in the closed alcove with every $\|\varphi(\alpha)\| \ge 1/h$ has every $a_i \ge 1/h$, hence every $a_i = 1/h$ since $\sum_i n_i a_i = 1$; so the maximizers are the forms equivalent to $\rho^\vee/h$ under W and $P^\vee$. They have order exactly h, since $k\rho^\vee/h$ is integer-valued on R only if $h \mid k$. The classes are the W-orbit of $\rho^\vee/h$ in the torus of forms modulo $P^\vee$, and its stabilizer corresponds to the elements of $W \ltimes P^\vee$ fixing $\rho^\vee/h$. The affine Weyl group $W \ltimes Q^\vee$ acts simply transitively on alcoves, and $W \ltimes P^\vee$ is its extension by the group $\Omega \cong P^\vee/Q^\vee$ of elements mapping the alcove to itself [3, Ch. VI, §2]. Since $\rho^\vee/h$ lies inside the alcove, an element fixing it maps the alcove to itself, so lies in Ω; conversely Ω permutes the coordinates $a_0, \dots, a_\ell$, which are all equal at $\rho^\vee/h$, so it fixes $\rho^\vee/h$. Hence the stabilizer is Ω, and there are $|W|/|P^\vee/Q^\vee|$ classes. □

An exact search over all torsion points of order dividing 24 (for D_4) and 20 (for A_4) finds exactly these classes (`certificates/d4_kappa_classes.py`). So the case $\pi_v L \cong D_4$ in Corollary 11 gives at most $\Phi_5(D_4) = \kappa(D_4)^2\eta(D_4) = 3(36\sqrt{2}/35)^2 \approx 6.348$. The equality configuration ($\pi_v L \cong D_4$, $\varphi = \rho^\vee/h$, $\gamma(L) = 1.3496$) has a tie for the minimum among v and the lifts of the five roots $\alpha_1, \dots, \alpha_4, \theta$. Likewise, lifts over A_4 itself (numerically a second local maximum of η_4, §7.2) cannot reach $\eta = 6$.

Write $\Phi_n(M) = \kappa(M)^{(n-1)/2}\eta(M)$ for lattices M of rank $n-1$, so that $\eta_n \le \sup_M \Phi_n(M)$ by Proposition 10. Some values for $n = 5$:

M	$\eta(M)$	$\kappa(M)$	$\Phi_5(M)$
D_4	3	$36\sqrt{2}/35 \approx 1.4546$	6.348
A_4	2.1466	$\frac{25}{24}\cdot 2\cdot 5^{-1/4} \approx 1.3932$	4.167
$\mathbb{Z}^4$	1	$4/3$	1.778

(For non-root lattices numerical values of κ from random searches are only lower bounds, since the maximum sits at a non-smooth point with many ties.)

5.2 The local maximum at D_5

Throughout, $d(G, G') = \|\log(G^{-1/2}G'G^{-1/2})\|_F$ is the affine-invariant metric on Gram matrices of equal determinant. It is invariant under $G \mapsto UGU^\top$ ($U \in GL_n(\mathbb{Z})$) and under $G \mapsto G^{-1}$.

Theorem 13. *D_5 is a strict local maximum of η. More precisely, if G_0 is a Gram matrix of D_5 and $d(G, G_0) \le 1/10$, then $\eta(G) \le 6$, with equality only at G_0 (after scaling G to $\det G = \det G_0$; η is scale-invariant). The same holds about every Gram matrix of D_5^*.*

The radius $1/10$ is the one the (B) leaves of Theorem 14 use. With the exact first-order radius of Appendix A, the same argument gives $\eta \le 6$ for $d(G, G_0) \le 6/25$.

Proof (computer-assisted). Write $G = G_0^{1/2} e^{\Lambda} G_0^{1/2}$ with Λ traceless, so $d(G, G_0) = \|\Lambda\|_F =: t$.

- D_5 has 720 Ψ-minimal flags, enumerated with the LLL($\delta = 1$) bound on $|b_1|^2$ rather than μ_5, so there is no circularity with Corollary 36. Their first-order terms $L_f(\Lambda)$ span the 14-dimensional space of traceless Λ, and exact rational LPs show that their convex hull contains a ball of radius $r \ge 0.2392$ in the dual norm. So for every Λ some minimal flag has $L_f(\Lambda) \le -rt$.
- Since $\log\det$ is concave, each flag satisfies $\log h_f(G) \le \log h_f(G_0) + L_f(\Lambda) + 4\,\mathrm{tr}\,E$ with $E = e^{\Lambda} - I - \Lambda$, and $\mathrm{tr}\,E \le \frac{1}{2}t^2 e^{\sqrt{4/5}\,t}$.
- Hence $\log \eta(G) \le \log 6 - rt + 2t^2 e^{\sqrt{4/5}\,t} < \log 6$ for $0 < t \le 1/10$.
- $G \mapsto JG^{-1}J$ (the reversed dual basis) preserves η and d, which gives the statement for D_5^*.

The details are in Appendix A. □

Qualitatively, the local maximality of η at D_5 follows from Meyer's Voronoi-type theorem [30, Thm. 34]: a form is extreme (a local maximum) for λ if and only if it is perfect and eutactic for λ, that is, the gradients of its minimal length functions span the tangent space affinely and 0 lies in the relative interior of their convex hull. For the staircase λ the length functions are the $\log h_f$, and the first step of the proof verifies exactly these two properties (a fortiori for all minimal flags, since 720 of them already suffice). What Theorem 13 adds is strictness with an explicit radius, which the proof of $\eta_5 = 6$ requires.

5.3 The value $\eta_5 = 6$

Theorem 14. $\eta_5 = 6$. *A five-dimensional lattice L has $\eta(L) = 6$ if and only if L is similar to D_5 or to D_5^*.*

The proof has two computer-assisted stages, and uses only the following:

inputs		used for
Proposition 10, Lemma 3	$\to$	the reduction to $M = \pi_v L$, and the range of r
η_2, η_3, η_4 (Proposition 6, Theorem 7)	$\to$	flag bounds in both stages
Theorem 35 with $m \le 4$ (proved in §9, independently of §5)	$\to$	the root box of stage 1
$\gamma_5 = 8^{1/5}$ [24]	$\to$	the range of r in stage 2
Theorem 13 and its dual	$\to$	the (B) leaves of stage 2, and the equality case
stages 1 and 2 (App. E)	$\Rightarrow$	Theorem 14 $\Rightarrow$ Corollary 36: $\mu_5 = \gamma_5$

Corollary 36 for $m = 5$ comes after the theorem and is not used in it, so there is no circularity.

Proof (computer-assisted). By Theorem 5 and isoduality, $\eta(D_5) = \eta(D_5^*) = 6$. Let L have determinant 1, take any shortest vector v, let $r = |v|^2$, and let M be $\pi_v L$ rescaled to determinant 1, with a basis $b_1, \dots, b_4$ of Gram matrix S. In the orthogonal model, L has the basis $v, b_1', \dots, b_4'$ with $b_i' = r^{-1/8} b_i + \varphi_i v$, where $\varphi \in \mathbb{R}^4$ is the lifting form, and Gram matrix

$$G = \begin{pmatrix} r & r\varphi^\top \\ r\varphi & r^{-1/4} S + r\varphi\varphi^\top \end{pmatrix}, \qquad \det G = 1.$$

Changing b_i' by integer multiples of v, and v by $-v$, reduces φ to $[0, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]^3$. Moreover $r \le \gamma_5 = 8^{1/5}$ [24], and $r \ge 2^{2/5}$, since otherwise the split along v (Lemma 3) and Theorem 7 give $\eta(L) \le r^{5/2} \eta(M) \le 3r^{5/2} < 6$.

- *Stage 1.* The search of Appendices C and D (which proves Theorem 20 in §6) is run with the threshold 6. A box of Ψ-minimal bases of M closes when $\Phi_5 < 6$ on it, or when it contains no Ψ-minimal basis; a box that has not closed by the time its width falls below 0.12 is *collected*. By Proposition 10, $\eta(L) \le \Phi_5(M) < 6$ unless a Ψ-minimal basis of M lies in a collected box X.

- *Stage 2.* For each collected box X, a second branch-and-bound covers the 14-dimensional set of (S, φ, r) with S in X, φ as above and $r \in [2^{2/5}, 8^{1/5}]$. A cell closes when one of the following holds on all of it:

  (N) an explicit lattice vector is strictly shorter than v, so this parametrization has v not shortest;

  (F) a single flag of L, or the split of Lemma 3 along one of its sublattices with the known values $\eta_1, \dots, \eta_4$, gives $\eta(L) < 6$;

  (P) a convex combination of such bounds gives $\eta(L) < 6$;

  (B) $d(G, G_0) \le 1/10$ for a Gram matrix G_0 of D_5 or of D_5^*, so $\eta(L) \le 6$, with equality only at G_0, by Theorem 13.

Each stage writes a log of its search tree, every leaf is re-proved by a separately written checker in interval arithmetic, and a coverage script verifies that the leaves tile the search space and that the collected boxes are exactly the roots of stage 2 (Appendices D and E).

Now let $\eta(L) \ge 6$, and take v shortest and $b_1, \dots, b_4$ Ψ-minimal. By stage 1 the basis lies in a collected box, and by stage 2 the point (S, φ, r) lies in a leaf. It is not of type (N), since v is shortest, nor of types (F) or (P), which are strict. So it lies in a leaf of type (B), and $\eta(L) = 6$ with $G = G_0$, that is, L is similar to D_5 or D_5^*. In particular $\eta_5 = 6$. □

The proof does not use Propositions 21 or 22: no ball about $\pi_v D_5$ or M_* (§6) was used. Stage 1 has 3.2×10^8 leaves, 1.5×10^7 of them collected; merged into 3.6×10^6 boxes, these are the roots of stage 2, which has 8.7×10^9 leaves, only 47,437 of them of type (B). No leaf of stage 1 needed the ball of Proposition 21: every neighbourhood of $\pi_v D_5$ is collected, and the equality case is decided entirely in stage 2. With Corollary 36, Theorem 14 gives $\mu_5 = \gamma_5$.

Remark 15 (Vectors and flags). Korkine and Zolotarev showed that $\gamma_5 = 8^{1/5}$, attained only by D_5 [24]. For complete flags the extremal role is shared with the dual lattice: $\eta_5 = 6$ is attained by D_5 and D_5^*, and only by them. This is forced by isoduality, since $\eta(L) = \eta(L^*)$, but the two lattices play very different roles for vectors. D_5^* is far from critical, with $\gamma(D_5^*) = 2^{2/5} \approx 1.32$ against

$\gamma_5 \approx 1.52$; it is extremal at the other end of the flag. A rank-4 sublattice of L has determinant $|w|^2 \det L$ for a primitive dual vector w, so $\gamma_{5,4}(L) = \gamma(L^*)$, and the densest hyperplane sections occur for $L = D_5^*$. So D_5 is extremal for the first step of a flag and D_5^* for the last, and since Ψ^2 treats a flag and its reversed dual symmetrically, both maximize the product of all the steps. The value $\gamma(D_5^*) = 2^{2/5}$ is also exactly the threshold of the dual-side split, and D_5^* projects along a shortest vector onto D_4 (Remark 23). The same pattern appears in the other dimensions where η_n is known or conjectured. In dimension 3 the value $\frac{3}{2}$ is attained by A_3 and A_3^* [17, Thm. 1], [30, Prop. 47], [29, Prop. 4.14]. D_4 and E_8 are similar to their duals, so they are unique. The conjectured extremal lattices in dimensions 6 and 7 come in the pairs E_6, E_6^* and E_7, E_7^*.

6 The landscape of Φ_5

Proposition 10 alone cannot prove $\eta_5 = 6$: the maximum of Φ_5 is $2^{16}/(3^4 \cdot 5^3) \approx 6.4727 > 6$, attained at a lattice M_*. This section identifies M_*, proves that it is the global maximum (Theorem 20), and describes the other peaks of Φ_5. Theorem 14 uses none of it, only the search method of Appendices C and D.

Definition 16. Let M_* be the lattice with Gram matrix

$$\sqrt{\tfrac{12}{5}} \begin{pmatrix} 1 & -\frac{1}{2} & 0 & 0 \\ -\frac{1}{2} & 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & -\frac{1}{2} & \frac{8}{9} & \frac{1}{18} \\ 0 & -\frac{1}{2} & \frac{1}{18} & \frac{8}{9} \end{pmatrix} \qquad (\det = 1),$$

and let $\varphi_0 = (0, 0, \frac{1}{3}, \frac{1}{3})$. Up to scaling, M_* has the primitive integral Gram matrix

$$P_0 = \begin{pmatrix} 18 & -9 & 0 & 0 \\ -9 & 18 & -9 & -9 \\ 0 & -9 & 16 & 1 \\ 0 & -9 & 1 & 16 \end{pmatrix} \qquad (18 \text{ times the matrix above}),$$

with $\det P_0 = 135^2$. Equivalently, up to scaling, M_* is the Hermitian lattice of rank 2 over the Eisenstein integers $\mathbb{Z}[\omega]$ with Gram matrix $\left(\begin{smallmatrix} 8 & 3-2\omega \\ 3-2\bar{\omega} & 8 \end{smallmatrix}\right)$, viewed as a $\mathbb{Z}$-lattice through its trace form (Proposition 17(8)).

Proposition 17 (The lattice M_*)**.** *1. $\kappa(M_*) = \sqrt{12/5}$, attained only at $\varphi = \pm\varphi_0$ modulo integer-valued forms.*

2. $\eta(M_*) = 2^{14}/(3^5 \cdot 5^2) \approx 2.6970$, attained by the flag $(-1,-1,0,-1)$, $(-1,-1,1,-1)$, $(-1,0,-1,-1)$.

3. Hence $\Phi_5(M_*) = 2^{16}/(3^4 \cdot 5^3) \approx 6.4727 > \Phi_5(D_4) \approx 6.348$.

4. M_* is isodual, and even 135-modular: P_0 and $135P_0^{-1}$ are even and integral, and $U(135P_0^{-1})U^\top = P_0$ for the unimodular U with rows $(0,0,-1,0)$, $(0,0,0,1)$, $(-1,1,0,-1)$, $(0,-1,1,0)$.

5. In P_0 the norms begin 16 (18 vectors), 18 (6), 30 (6), 34 (18). So M_* has 9 pairs of minimal vectors and is not perfect (in dimension 4 that needs at least 10 pairs), $\gamma(M_*) = 16/\sqrt{135} \approx 1.3771 < \sqrt{2} = \gamma(D_4)$, and M_* is not similar to $A_2 \otimes A_2$ (which has $\gamma = \frac{4}{3}$).

6. $|\operatorname{Aut}(M_*)| = 72$. *$M_*$ is not strongly 135-modular: for the exact divisors $d = 5$ and $d = 27$ of 135, the lattice $\sqrt{d}\,(M_*^{\#} \cap \frac{1}{d}M_*)$ is not isometric to M_*.*

7. *The genus of P_0 has class number 6, and M_* is the unique lattice of largest minimum, that is the unique densest lattice, in it.*

8. *M_* has exactly four automorphisms g with $g^2 + g + 1 = 0$. Letting ω act as one of them makes M_* a Hermitian lattice of rank 2 over $\mathbb{Z}[\omega]$, with Gram matrix $H = \left(\begin{smallmatrix} 8 & 3-2\omega \\ 3-2\bar{\omega} & 8 \end{smallmatrix}\right)$ and $\det H = 45$, whose trace form $\operatorname{Tr}_{\mathbb{Q}(\omega)/\mathbb{Q}} h$ is isometric to P_0.*

Proof. (1) In the Gram matrix of Definition 16 the 9 pairs of minimal vectors have norm $\frac{8}{9}\sqrt{12/5}$, and every other nonzero vector has norm at least $\sqrt{12/5}$. So $\min_s k_s(\varphi) \geq \sqrt{12/5}$ if and only if $\|\varphi(s)\| \geq \frac{1}{3}$ for each minimal pair. At φ_0, $\|\varphi_0(s)\|$ is $\frac{1}{3}$ unless $3 \mid s_3 + s_4$, in which case it is 0, and no minimal vector has $3 \mid s_3 + s_4$; so $\min_s k_s(\pm\varphi_0) = \sqrt{12/5}$. Conversely, the minimal vectors include $s_1 = e_4$, $s_2 = e_3$, $s_3 = e_2 + e_4$ and $s_4 = e_1 + e_2 + e_4$, a basis of $\mathbb{Z}^4$, and also $s_2 + s_3$, $s_2 + s_4$ and $s_2 + s_3 - s_1$. Write $a_i \in [\frac{1}{3}, \frac{2}{3}]$ for $\varphi(s_i)$ modulo 1. Then $a_2 + a_3 \in [\frac{2}{3}, \frac{4}{3}]$ must lie in $[\frac{1}{3}, \frac{2}{3}]$ modulo 1, which forces $a_2 = a_3 = t$ with $t \in \{\frac{1}{3}, \frac{2}{3}\}$. In the same way $s_2 + s_4$ gives $a_4 = t$, and $s_2 + s_3 - s_1$, with $2t - a_1 \in [2t - \frac{2}{3}, 2t - \frac{1}{3}]$, gives $2t - a_1 = t$, that is $a_1 = t$. Since the s_i form a basis, φ is determined modulo integer forms by the a_i, and $a_i = \frac{1}{3}$ for all i is φ_0. So $\kappa(M_*) = \sqrt{12/5}$, attained only at $\pm\varphi_0$. As an independent machine check, an exact enumeration over rational polyhedra (all 9 pairs, `certificates/mstar_unique.sage`) confirms this.

(2) The flag gives the upper bound. The lower bound comes from the exact Ψ-minimization $\eta(L) = \min_v |v|^n \eta(\pi_v L)/\sqrt{\det L}$, recursing to dimension 1. A Ψ-minimal basis has $|b_1|^{2n} \leq \mu_n^n \det L$, with $\mu_n^n = \frac{4}{3}, 2, 4$ for $n = 2, 3, 4$ (Theorem 35), so the minimum runs over finitely many primitive v. The computation is done in exact rational arithmetic, since η^2 and the Gram matrices stay rational.

(3) follows from (1) and (2). Items (4)–(8) are exact computations with P_0, in `certificates/`. For (7), the six representatives of the genus have automorphism groups of orders $8, 72, 16, 24, 12, 144$ and minima $4, 16, 10, 4, 12, 10$ (in the scale of P_0), and $\sum 1/|\operatorname{Aut}| = 1/3$ equals the mass of the genus, so the list is complete. For (8), $h(x, y) = \frac{1}{3}\sum_{k=0}^{2} \langle g^k x, y\rangle \omega^k$ is Hermitian for the $\mathbb{Z}[\omega]$-action through g, its trace is $\langle\,,\rangle$, and in a $\mathbb{Z}[\omega]$-basis e, f (with e, ge, f, gf a $\mathbb{Z}$-basis) its Gram matrix is H. □

We did not find M_*, or its extremal lift L_* (Remark 18), in the Nebe–Sloane catalogue [33], whose four-dimensional modular lattices have levels up to 15 and whose table of quaternary forms stops at discriminant 1732, or in [7]. Conway and Sloane [7] classify the isodual lattices of dimension 3 and, in dimensions 5–7, construct paths from L_n to L_n^* through an isodual lattice; in dimension 4 they record only the densest isodual lattice, D_4 (their Table I). For identification: M_* has theta series $1 + 18q^{16} + 6q^{18} + 6q^{30} + 18q^{34} + \cdots$ in the coordinates of P_0; it is the unique densest lattice in its genus, of class number 6, and Hermitian over $\mathbb{Z}[\omega]$ (Proposition 17(7), (8)). Its Hermitian genus has class number 4,[1] so M_* is not among the Hermitian genera of class number at most two classified by Kirschmer [22].

M_* is not isolated among isodual lattices. With U as in (4), the Gram matrices G with $U(cG^{-1})U^{\top} = G$ for some $c > 0$ form, up to scaling, a three-dimensional family through M_*

[1] Four of the six classes of the integral genus carry a $\mathbb{Z}[\omega]$-structure, one each up to automorphisms and isometric to its conjugate, all of determinant 45; one of them is $\langle 5\rangle \perp \langle 9\rangle$. They agree at 2 and 5, which are inert, and at 3 each splits as $\langle u\rangle \perp \langle 45/u\rangle$ with $u \equiv 2 \bmod 3$, a non-norm, so they form one Hermitian genus (`certificates/mstar_herm_genus.sage`).

that is totally geodesic for the metric d of Theorem 13 (it is the fixed set of an isometry). In the coordinates of P_0 it contains the Gram matrices

$$H_D = \begin{pmatrix} 2 & -1 & 0 & 0 \\ -1 & 2 & -1 & -1 \\ 0 & -1 & 2 & 0 \\ 0 & -1 & 0 & 2 \end{pmatrix}, \qquad H_A = \begin{pmatrix} 6 & -3 & 0 & 0 \\ -3 & 6 & -3 & -3 \\ 0 & -3 & 4 & 1 \\ 0 & -3 & 1 & 4 \end{pmatrix}$$

of D_4 and $A_2 \otimes A_2$, and M_* lies exactly on the geodesic joining them: $H_D^{-1}P_0$ and $H_D^{-1}H_A$ commute, and at determinant 1 their eigenvalues are $\sqrt{6/5}^{\pm 1}$ and $\sqrt{2}^{\pm 1}$ (each twice). So $d(D_4, M_*) = \ln \frac{6}{5}$, $d(M_*, A_2 \otimes A_2) = \ln \frac{5}{3}$ and $d(D_4, A_2 \otimes A_2) = \ln 2$; M_* is the point at parameter $\log_2 \frac{6}{5}$ on the geodesic from D_4 to $A_2 \otimes A_2$ (`certificates/geodesic_exact.py`). Of the three, D_4 is the shadow of D_5^* and $A_2 \otimes A_2$ that of $\pi_v E_6$ (Remark 24); between them, M_* is the global maximum of Φ_5 (Theorem 20, Figure 1).

The curve in Figure 1 is numerical. At each of 66 points, κ is computed in floating point by branch and bound over the torus of lifting forms φ, which brackets the global maximum (not a local one) to relative accuracy 10^{-8}. η is computed by exact Ψ-minimization of a rational Gram matrix whose entries are within 10^{-13} of the floating-point ones (`numerics/geodesic_phi5.sage`).

The geodesic from D_4 to $A_2 \otimes A_2$ lies in two further totally geodesic families. Let ω act as an automorphism g of M_* with $g^2 + g + 1 = 0$ (Proposition 17(8)); g also fixes D_4 and $A_2 \otimes A_2$, and the trace forms of rank-2 Hermitian forms over $\mathbb{Z}[\omega]$ for this action form a copy of hyperbolic 3-space through all three lattices. The two such actions, up to inversion, give two such families. All their lattices are proportional to symplectic lattices in the sense of [5, 1]: $z \mapsto H^{-1}S\bar{z}/\sqrt{-3}$, with $S = \left(\begin{smallmatrix} 0 & 1 \\ -1 & 0 \end{smallmatrix}\right)$, is a similarity of L onto L^* whose square is $-1/(3 \det H)$. The isodual family of Proposition 17(4) is of a different kind: its isoduality, normalized to an isometry, has order 4, and its square is not ± 1. At M_* the tangent spaces of the three families meet pairwise only in the directions of the geodesic and of scaling (exact computations in `certificates/herm_family.sage`).

Remark 18 (Limits of Proposition 10 in dimension 5). By Proposition 17, D_4 does not maximize Φ_5, and Proposition 10 alone cannot give better than $\eta_5 \le 2^{16}/(3^4 \cdot 5^3) \approx 6.4727$. M_* is a strict local maximum of Φ_5 (Proposition 22) and the global one (Theorem 20). The bound is far from tight at M_*: the extremal lift L_* (with $\pi_v L_* \cong M_*$, $\gamma(L_*) = (12/5)^{2/5}$ and lifting form φ_0) has $\eta(L_*) = 486/125 = 3.888$. Explicitly, take $r = 18$, so that the minimal vectors of P_0 (norm $16 = r(1 - \frac{1}{9})$) tie with v; in the basis $(v, b_1', \dots, b_4')$ the Gram matrix is $\left(\begin{smallmatrix} r & r\varphi_0^\top \\ r\varphi_0 & P_0 + r\varphi_0\varphi_0^\top \end{smallmatrix}\right)$, and dividing by 3 gives the primitive integral form

$$G(L_*) \sim \begin{pmatrix} 6 & 0 & 0 & 2 & 2 \\ 0 & 6 & -3 & 0 & 0 \\ 0 & -3 & 6 & -3 & -3 \\ 2 & 0 & -3 & 6 & 1 \\ 2 & 0 & -3 & 1 & 6 \end{pmatrix}, \qquad \det = 1350 = 2 \cdot 3^3 \cdot 5^2,$$

with norms beginning 6 (26 vectors), 8 (18), 10 (6), 12 (30), 14 (18); so L_* has 13 pairs of minimal vectors and is not perfect. The flag $e_5, e_3, e_2, e_4, (1, -1, -2, -1, -1)$ has prefix determinants $6, 27, 108, 405, 1350$, so $\eta(L_*) \le 6 \cdot 27 \cdot 108 \cdot 405/1350^2 = 486/125$, and the exact Ψ-minimization (over LLL($\delta = 1$)-reduced bases, as in Appendix A) shows equality (`certificates/lstar_check.sage`). Before Theorem 20, a numerical search found no larger value: 144 ascents of Φ_5 in floating point (κ by branch-and-bound over φ, exhaustive η), of which all 21 ending above 6.4 converged to M_*, from both D_4 and random starts.

Theorem 19. $\Phi_5(M) < 7$ *for every 4-dimensional lattice M, so $\eta_5 < 7$ (computer-assisted; Appendix C).*

This first, smaller run improves the bound $3 \cdot 2^{5/4} \approx 7.135$ of Table 3 and is superseded by Theorem 20; we keep it as an independent check of the machinery, whose method Appendix C describes. By Remark 18, Proposition 10 alone cannot go below 6.4727; a second run reaches that limit.

Theorem 20. $\Phi_5(M) \le 2^{16}/(3^4 \cdot 5^3)$ *for every 4-dimensional lattice M, with equality only for M similar to M_*. Hence $\eta_5 < 2^{16}/(3^4 \cdot 5^3) \approx 6.4727$ (superseded by Theorem 14; the content here is the first statement).*

Proof (computer-assisted). The search of Theorem 19 is run with the threshold $2^{16}/(3^4 \cdot 5^3)$ in place of 7. A box in which every Φ_5 is below the threshold closes as before. Near M_* that is impossible, since $\Phi_5(M_*)$ equals the threshold, and a box there closes when it lies inside the ball of Proposition 22 about one of the copies of M_* in the search space, where $\Phi_5 \le 2^{16}/(3^4 \cdot 5^3)$ with equality only at the center. The details are in Appendix D.

For the second statement let L have $\det L = 1$ and a shortest vector v, and let M be $\pi_v L$ rescaled to determinant 1. By Proposition 10, $\eta(L) \le \kappa(M)^2 \eta(M) = \Phi_5(M) \le 2^{16}/(3^4 \cdot 5^3)$. Equality needs M similar to M_* and $\gamma(L)^{5/4} = \kappa(M)$; by the proof of Proposition 10 the lifting form of L then attains $\kappa(M_*)$, so by Proposition 17 it is $\pm\varphi_0$, and L is similar to the lift L_* of Remark 18. But $\eta(L_*) \le 3.888$. So $\eta(L) < 2^{16}/(3^4 \cdot 5^3)$ for every L, and since η_5 is attained (by Mahler's compactness theorem), $\eta_5 < 2^{16}/(3^4 \cdot 5^3)$. □

The next two propositions use the metric d of Theorem 13. In both, the inradius is that of the first-order generalized gradient of $\log \Phi_5$ at the center, in the dual of the Frobenius norm on traceless Λ (where $G = G_0^{1/2} e^{\Lambda} G_0^{1/2}$); it is the sharp first-order rate of decrease.

For the next proposition we need $\pi_v D_5$ explicitly. Projecting $D_5 = \{x \in \mathbb{Z}^5 : \sum x_i \text{ even}\}$ along the root $v = e_4 + e_5$ and writing $y = (x_1, x_2, x_3, x_4 - x_5)$ gives

$$\pi_v D_5 \cong \Big\{ y \in \mathbb{Z}^4 : \sum y_i \text{ even} \Big\}, \qquad |y|^2 = y_1^2 + y_2^2 + y_3^2 + \tfrac{1}{2} y_4^2,$$

so $\pi_v D_5$ is D_4 with one coordinate axis scaled by $1/\sqrt{2}$. It has determinant 2 and minimum $\frac{3}{2}$, attained by the 6 pairs $\pm e_i \pm e_4$ ($i \le 3$), and the lifting form of D_5 over it is $\varphi(y) = y_4/2$, a 2-torsion point. (All roots of D_5 are equivalent under $\mathrm{Aut}(D_5)$, so the choice of v does not matter.) By contrast, projecting $D_5^* = \mathbb{Z}^5 \cup (\mathbb{Z}^5 + \frac{1}{2}\mathbf{1})$ along its shortest vector e_1 gives $\mathbb{Z}^4 \cup (\mathbb{Z}^4 + \frac{1}{2}\mathbf{1}) = D_4^*$, which is similar to D_4, again with a 2-torsion lifting form.

Proposition 21 ($\pi_v D_5$ is a sharp local maximum of Φ_5)**.** *Let G_0 be a Gram matrix of $\pi_v D_5$, the projection of D_5 along a shortest vector. Then $\Phi_5(G) \le 6$ whenever $d(G, G_0) \le 1/20$, with equality only at G_0. The generalized gradient of $\log \Phi_5$ at G_0 has inradius exactly $1/(3\sqrt{3})$.*

Proof (computer-assisted). Appendix B. □

At G_0, κ is attained at a single class φ_0 modulo integer forms, a 2-torsion point (the glue of D_5 over $\pi_v D_5$), with 13 active pairs: 6 with $\|\varphi_0(s)\| = \frac{1}{2}$ and 7 with $\varphi_0(s) \in \mathbb{Z}$. Uniqueness is elementary. With $\det G_0 = 2$ the six half vectors h have norm $\frac{3}{2}$, so $\min_s k_s(\varphi) \ge 2$ forces $\|\varphi(h)\| = \frac{1}{2}$ for each of them. The half vectors generate the lattice (in the coordinates above they are $\pm e_i \pm e_4$, $i \le 3$, and $e_i + e_j = (e_i + e_4) + (e_j - e_4)$; see also Appendix B), so this determines φ modulo integer forms, as φ_0; since $\min_s k_s(\varphi_0) = 2$, $\kappa_{\mathrm{raw}}(G_0) = 2$, attained only at φ_0; here $\kappa_{\mathrm{raw}}(G)$ denotes $\sup_\varphi \min_{s \ne 0} s^\top G s / (1 - \|\varphi(s)\|^2)$ for the Gram matrix G as given, so that $\kappa = \kappa_{\mathrm{raw}} / \det(G)^{1/4}$. The minimum defining η is attained by 36 flags.

Proposition 22 (M_* is a sharp local maximum of Φ_5)**.** *Let G_0 be a Gram matrix of M_* (Definition 16). Then $\Phi_5(G) \le 2^{16}/(3^4 \cdot 5^3)$ whenever $d(G, G_0) \le 1/40$, with equality only at G_0. The generalized gradient of $\log \Phi_5$ at G_0 has inradius exactly $11/(16\sqrt{2})$.*

Proof (computer-assisted). Appendix B. □

Here κ is attained at the single 3-torsion class $\varphi_0 = (0, 0, \frac{1}{3}, \frac{1}{3})$ (Proposition 17, up to sign), with 12 active pairs (9 with $\|\varphi_0(s)\| = \frac{1}{3}$, the minimal vectors, and 3 with $\varphi_0(s) \in \mathbb{Z}$), and η by 54 flags.

Remark 23 (The landscape of Φ_5). $\pi_v D_5$ (6) and M_* (6.4727) are sharp local maxima of Φ_5 (Propositions 21 and 22), and each has a unique κ-maximizer class (up to sign), at a torsion point of order 2 and 3 respectively (after Proposition 21, and Proposition 17). Also $\Phi_5(D_4) = 3(36\sqrt{2}/35)^2 \approx 6.348$ exactly, and the κ-maximizers of D_4 are the 24 classes (up to sign) of principal elements $w\rho^\vee/6$, $w \in W(D_4)$, all of order 6 (Proposition 12). The rest of this remark is numerical (floating point) and is not used in any proof. D_4 appears to be a third local maximum: Φ_5 decreases at D_4 in every direction tested (2,000 random directions in all nine dimensions, radius 0.002 to 0.03, κ by branch and bound). Along the slowest first-order descent direction from $\pi_v D_5$, Φ_5 falls to about 5.86 and rises again to a copy of D_4 (at distance 0.600), which is joined to M_* by a pass at height about 6.28. So $\{\Phi_5 > 6\}$ appears to be a neighborhood of D_4 and M_*, separated from $\pi_v D_5$. In floating-point searches $\pi_v E_6$ also has a unique maximizer class for Φ_6 (order 2). Figure 2 shows a two-dimensional section through all three lattices.

The three peaks are tied to the extremal five-dimensional lattices by the two inequalities of Proposition 10:

$$\begin{aligned}
&\eta(D_5) = \gamma(D_5)^{5/2}\,\eta(\pi_v D_5) = \Phi_5(\pi_v D_5) = 6,\\
&\eta(D_5^*) = \gamma(D_5^*)^{5/2}\,\eta(D_4) = 2 \cdot 3 = 6 < \Phi_5(D_4) = 3\big(\tfrac{36}{35}\sqrt{2}\big)^2 \approx 6.348,\\
&\eta(L_*) \le \tfrac{486}{125} < \gamma(L_*)^{5/2}\,\eta(M_*) = \Phi_5(M_*) = \tfrac{2^{16}}{3^4 \cdot 5^3} \approx 6.4727.
\end{aligned}$$

The bound $\eta(L) \le \gamma(L)^{5/2}\eta(\pi_v L)$ is lost at L_*, the bound $\gamma(L)^{5/4} \le \kappa(\pi_v L)$ is lost at D_5^* (by the factor $h^2/(h^2-1) = 36/35$ in κ), and neither is lost at D_5. Numerically (floating point), over M_* itself the lifts of largest η have $\eta \approx 5.16$, and near M_* they stay below about 5.2.

Remark 24 (Peaks of Φ_n for $n = 4, 5, 6$). Lower local maxima of η cast shadows in the same sense. The following is exact.

- $\pi_v A_4 \cong (\mathbb{Z}^3, I + \frac{1}{2}J)$ (J the all-ones matrix), and $\Phi_4(\pi_v A_4) = \eta(A_4) = 24/(5\sqrt{5})$. Indeed $\eta(\pi_v A_4) = \frac{6}{5}$ by exact Ψ-minimization, and $\kappa_{\mathrm{raw}}(\pi_v A_4) = 2$, attained only at $\varphi = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$: the minimal vectors $\pm e_i$ have norm $\frac{3}{2}$ and every other vector norm at least 2, so $k_{e_i}(\varphi) = \frac{3}{2}/(1 - \|\varphi(e_i)\|^2)$ is at most 2, with equality only when $\|\varphi(e_i)\| = \frac{1}{2}$; at $\varphi = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$ the vectors $e_i - e_j$ give exactly 2.

- $\pi_u(\pi_v E_6) \cong A_2 \otimes A_2$ for every shortest vector u of $\pi_v E_6$, and $\Phi_5(A_2 \otimes A_2) = \frac{16}{3} = \eta(\pi_v E_6)$. Here $\eta(A_2 \otimes A_2) = \frac{64}{27}$ by exact Ψ-minimization, and $\kappa(A_2 \otimes A_2) = \frac{3}{2}$, attained only at $\pm\varphi_0$ with $\varphi_0 = \frac{1}{3}$ on each $\alpha_i \otimes \beta_j$. In $A_2 \otimes A_2$ (determinant 81) the 9 minimal pairs are the $\alpha \otimes \beta$ for roots α, β, of norm 4, and all other vectors have norm at least 6. So $\min_s k_s(\varphi) \ge \frac{9}{2}$ forces $\|\varphi(\alpha \otimes \beta)\| \ge \frac{1}{3}$ for all of them. Then $(\alpha_1 + \alpha_2) \otimes \beta_j$ forces $\varphi(\alpha_1 \otimes \beta_j) = \varphi(\alpha_2 \otimes \beta_j) \in \{\frac{1}{3}, \frac{2}{3}\}$ modulo 1, and $\alpha_i \otimes (\beta_1 + \beta_2)$ does the same in the other factor, so $\varphi = \pm\varphi_0$, where the minimum is exactly $\frac{9}{2}$ ($\frac{3}{2}$ at determinant 1).

- $\pi_w(A_4^*) \cong A_3^*$ (bcc), the shadow of the top; and $\Phi_4(A_3) = \kappa(A_3)^{3/2}\eta(A_3) \approx 2.337$ with $\kappa(A_3) = \frac{16}{15}\gamma(A_3)$, the Coxeter value over the previous critical lattice.

The following is numerical (floating point) and is not used in any proof. In 200 ascents of Φ_4 with exact κ and η at the end points, every end point converges to bcc (188, three of them only from randomly rotated simplices), to fcc (4, one of them along a ridge on which the first-order slope vanishes) or to $\pi_v A_4$ (5); three torsion-seeded starts degenerated. At fcc and $\pi_v A_4$ the first-order test of Propositions 21 and 22, solved in floating point at the exact lattices, gives sharp local maxima (inradius at least 0.031 and 0.105). $A_2 \otimes A_2$, however, is not a local maximum of Φ_5: along the isodual geodesic through M_* (Figure 1), Φ_5 has a local minimum at $A_2 \otimes A_2$ and rises on both sides, at second order; Nelder–Mead re-ascents had stalled there. For Φ_6, 464 ascents found no peak above 22 other than $\pi_v E_6$ and D_5^* (after Conjecture 30). The peaks and shadows found are:

kind	$n = 4$	$n = 5$	$n = 6$ (above 22)
shadows of the extremal pair, $\pi_v L_n$ and $\pi_w(L_n^*)$	bcc (3)	$\pi_v D_5$ (6), D_4 (6.348)	$\pi_v E_6$, D_5^* (H_6)
shadows of lower local maxima of η_n	$\pi_v A_4$ (2.147)	$A_2 \otimes A_2$ (16/3, not a peak)	none found
Coxeter peaks over L_{n-1}	fcc (2.337)	$= D_4$	— ($\Phi_6(D_5) \approx 17.6$)
none of these	—	M_* (6.4727)	none found

So a shadow of a local maximum of η_n need not be a peak of Φ_n. M_* is the only peak found that exceeds the conjectured value H_n of η_n, and it is the one that stops Proposition 10 at 6.4727: its lift L_* has $\eta(L_*) = 486/125$, far below $\Phi_5(M_*)$. Like $A_2 \otimes A_2$, M_* is isodual with 9 pairs of minimal vectors and a unique 3-torsion κ-maximizer class, but only $A_2 \otimes A_2$ is a shadow. The unidentified local maximum of η_5 near 5.03 (§7.2) would cast a further shadow, which we have not computed. In dimension 5, $D_4 = \pi_w(D_5^*)$ is both a dual shadow and the Coxeter peak over $L_4 = D_4$; its value 6.348 exceeds $\eta(D_5^*) = 6$ by the factor $(36/35)^2$ in κ^2.

7 Dimensions 6 and 7

7.1 Bounds and the conjecture

n	lower bound	upper bound	from
5	6 (D_5)	6	Thm. 14
6	$128\sqrt{3}/9 \approx 24.63$ (E_6)	$16\sqrt{3} \approx 27.71$	$k = 2$, $\gamma'_{6,2} = 2$
7	144 (E_7)	$16\sqrt{3} \cdot 3^{7/4} \approx 189.5$	$k = 1$, $\gamma'_7 = \sqrt{3}$

Table 3: Bounds for $n = 5, 6, 7$; for $n = 6, 7$ from Lemma 3. The values γ'_5, γ'_7 are due to [36] and $\gamma'_{6,2}$ to [40]. Given $\eta_5 = 6$, the $n = 6$ entry is also the plain $k = 1$ bound $\eta_5\gamma_6^3 = 16\sqrt{3}$.

For $n = 5$, Lemma 3 gives $3 \cdot 2^{5/4} \approx 7.135$ ($k = 1$, $\gamma'_5 = \sqrt{2}$), Theorem 19 gives 7, Theorem 20 gives $2^{16}/(3^4 \cdot 5^3) \approx 6.4727$, the limit of the reduction to Φ_5, and Theorem 14 gives the exact value.

Conjecture 25. *$\eta_6 = 128\sqrt{3}/9$, attained only by E_6 and E_6^*, and $\eta_7 = 144$, attained only by E_7 and E_7^*.*

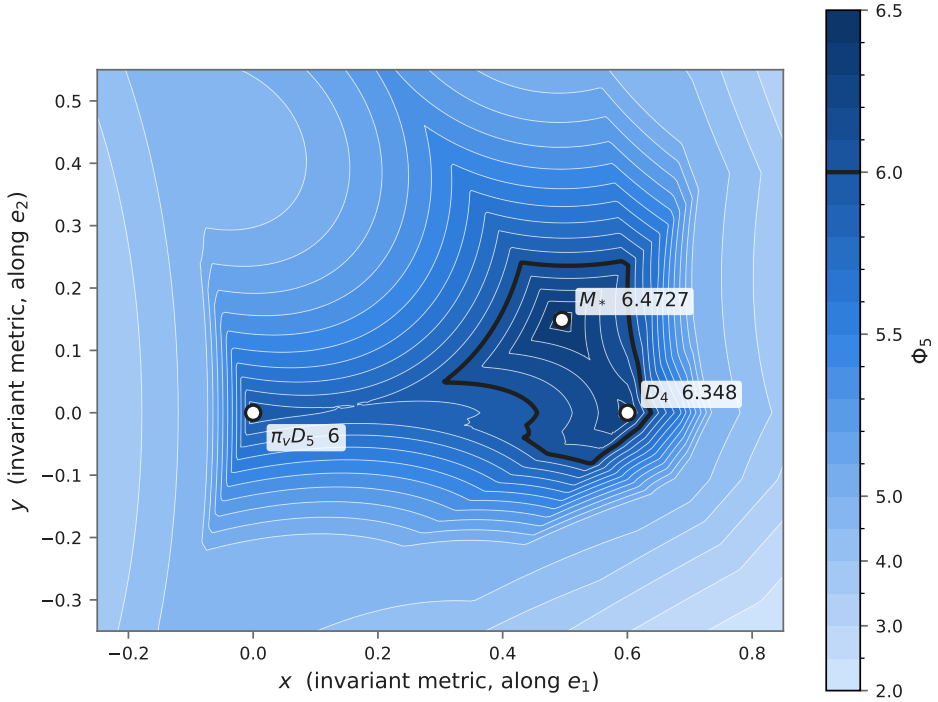


Figure 2: Φ_5 on the 2-plane through $\pi_v D_5$, the nearest copy of D_4 and the nearest copy of M_*, in exponential coordinates $G = G_0^{1/2} \exp(xe_1 + ye_2) G_0^{1/2}$ about $G_0 = \pi_v D_5$ (e_1, e_2 orthonormal for the invariant metric at G_0, so distances from $\pi_v D_5$ in the picture are true distances; D_4 at distance 0.600, M_* at 0.517). The bold curve is $\Phi_5 = 6$. $\pi_v D_5$ and D_4 are the projections of D_5 and D_5^* along a shortest vector. $\pi_v D_5$ and M_* are sharp local maxima (Propositions 21 and 22); D_4 is, numerically, a lower sharp peak, joined to M_* by a pass at about 6.28. A two-dimensional section of a nine-dimensional landscape; values are numerical (floating point: κ by branch and bound over the torus, η by exhaustive flag search), not certified.

Thus $\eta_n = H_n$ for $n \leq 5$ and $n = 8$ ([17] and Theorems 7, 14 and 8), and the conjecture is that this holds for all $n \leq 8$. For $n = 5$ no split reaches 6, so a different method is needed (§5.3). Equality holds only at D_5 and D_5^*. These are not similar (D_5 is not isodual), and both attain 6 because $\eta(L) = \eta(L^*)$.

The conjectured maximizers are at least strict local maxima. This rules out competitors near E_6 and E_7 only, and is not evidence of global maximality; the evidence for that is the Φ_6 landscape below. The same argument applies to E_8, where $\eta_8 = 2304$ is proved (Theorem 8): there it gives local uniqueness, that no lattice near E_8 other than E_8 attains η_8, while global uniqueness is open (§10).

Proposition 26. *E_6, E_7 and E_8 are strict local maxima of η: if G_0 is a Gram matrix of E_n ($n = 6, 7, 8$) and $d(G, G_0) \leq \rho_n$, with $\rho_6 = 3/40$, $\rho_7 = 3/50$ and $\rho_8 = 1/20$, then $\eta(G) \leq H_n$, with equality only at G_0 (after scaling to $\det G = \det G_0$). For E_6 the radius can be taken to be $1/10$. The same holds about every Gram matrix of E_n^*.*

Proof. Write $G = G_0^{1/2} e^{\Lambda} G_0^{1/2}$ with Λ traceless and $t = \|\Lambda\|_F = d(G, G_0)$, as in Theorem 13. Let

$f = (L_1 \subset \cdots \subset L_{n-1})$ be a Ψ-minimal flag of G_0, P_k the orthogonal projection onto the span of L_k (in coordinates in which G_0 is the identity), and $C_f = \sum_{k<n} P_k - \frac{n-1}{2} I$, which is traceless. Since $\det G = \det G_0$, Lemma 38 applied to each L_k gives

$$\log \eta(G) \le \log H_n + \langle C_f, \Lambda \rangle + \tfrac{n-1}{2}\, t^2 e^{\sqrt{(n-1)/n}\, t}.$$

The eigenvalues of C_f are $j - \frac{n-1}{2}$ $(0 \le j < n)$, so $\|C_f\|_F^2 = n(n^2-1)/12$, which is $35/2$, 28 and 42 for $n = 6, 7, 8$.

The flags. Let F be the orbit of the root flag of G_0 under the Weyl group $W = W(E_n)$, and $N = |F|$. Every flag in F is Ψ-minimal, since the root flag is (Theorem 5) and W acts by automorphisms. Nothing has to be enumerated: the argument uses only that F is a W-orbit of Ψ-minimal flags.

The design identities. Let $D = n(n+1)/2 - 1$ be the dimension of the space of traceless Λ (20, 27 and 35). Then

$$\sum_{f \in F} C_f = 0, \qquad \sum_{f \in F} \langle C_f, \Lambda \rangle^2 = \frac{N \|C_f\|_F^2}{D}\, \|\Lambda\|_F^2 \quad \text{for all traceless } \Lambda.$$

Indeed W acts irreducibly on traceless symmetric matrices for E_6, E_7 and E_8 (the characters $\phi_{20,2}$, $\phi_{27,2}$ and $\phi_{35,2}$; see [14]). So $\sum_f C_f$, which is fixed by W, is 0. The operator $\Lambda \mapsto \sum_f \langle C_f, \Lambda \rangle C_f$ is self-adjoint and commutes with W, so it is a scalar by Schur's lemma (a self-adjoint element of the commutant, a real division algebra, has a real eigenvalue), and taking the trace gives the scalar $N\|C_f\|_F^2/D$. At D_5 this fails: the corresponding space splits as $4 \oplus 10$. An analogous design argument for Hermite's constant is classical [28].

The radius. Fix Λ with $\|\Lambda\|_F = t$ and put $x_f = -\langle C_f, \Lambda \rangle$, $m = \max_f x_f$ and $M = \|C_f\|_F t$. Then $-M \le x_f \le m$, so $(m - x_f)(x_f + M) \ge 0$. Averaging over F, the x_f have mean 0 and mean square $\|C_f\|_F^2 t^2/D$, so $mM \ge \|C_f\|_F^2 t^2/D$, that is, some flag has $\langle C_f, \Lambda \rangle \le -rt$ with $r = \|C_f\|_F/D$: $r = \sqrt{70}/40$, $2\sqrt{7}/27$ and $\sqrt{42}/35$ for E_6, E_7 and E_8. Hence $\log \eta(G) \le \log H_n - rt + \frac{n-1}{2} t^2 e^{\sqrt{(n-1)/n}\, t}$, and in interval arithmetic the right-hand side is below $\log H_n$ for $0 < t \le \rho_n$. For E_6, an exact LP certificate as in Appendix A (a cross-polytope in $\Gamma = \mathrm{conv}\{C_f\}$, each vertex an explicit convex combination of flags) gives $r \ge 3/10$, and with it the radius $1/10$. As in Theorem 13, $G \mapsto JG^{-1}J$ gives the statement for E_n^*. □

The irreducibility was also checked exactly, by computing the commutant of the reflection action over $\mathbb{Q}$ (dimension 1 for E_6, E_7, E_8). For any lattice whose automorphism group acts irreducibly on traceless symmetric matrices, the same argument gives a first-order radius at least $\|C_f\|_F/D = \sqrt{n(n+1)/(3(n-1))}/(n+2)$, so such a lattice is a strict local maximum of η.

As an independent check of the identities for E_6 and E_7, we enumerated all Ψ-minimal flags and verified both identities in exact arithmetic. By the Hermite bound, $\det L_k \ge (2/\gamma_k)^k$, which is $2, 3, 4, 4, 4, 3$ for $k = 1, \ldots, 6$, and the root flag attains every bound (Theorem 5). So the Ψ-minimal flags are the chains with these determinants, and L_k is a critical lattice of minimum 2, hence by the uniqueness of critical lattices ([23, 24] for $k \le 5$, [45] for $k = 6$) one of $A_1, A_2, A_3, D_4, D_5, E_6$. Each step is generated by a root: if $L_{k+1} = L_k + \mathbb{Z}w$ with w reduced modulo L_k, then $|w|^2 \le \det L_{k+1}/\det L_k + \mathrm{cov}(L_k)^2$, with cov the covering radius, and this is $\frac{3}{2} + \frac{1}{2}$, $\frac{4}{3} + \frac{2}{3}$, $1 + 1$, $1 + 1$, $\frac{3}{4} + \frac{5}{4}$, that is 2 at every step [8]. Building the chains from the 72 (resp. 126) roots gives 19,440 Ψ-minimal flags of E_6 and 544,320 of E_7; for E_7 the sums are taken by dynamic programming over the distinct sublattices, 63, 336, 1,260, 315 and 378 in ranks 1 to 5. (Artefacts: `e6_flags.sage`, `e7_design.sage`, `e6_cross.sage` and `schur_check.sage`, each under five minutes, in `certificates/` of the code release.)

Remark 27 (Eutaxy radii of E_6 and E_7). For the true first-order radii we have exact upper bounds. Along $\Lambda = I - 3P$, with P the projection onto the plane of the two minuscule weights ω_1, ω_6 of E_6, $\min_f \langle C_f, \Lambda\rangle = -3$ (attained by 5,184 flags) and $\|\Lambda\|_F^2 = 12$, so $r(E_6) \le \sqrt{3}/2$; this is the analogue of $r(D_4) = 2/3$ (Appendix A). Along $\Lambda = I - 7P$, with P the projection onto the minuscule weight ω_7 of E_7, the minimum is $-35/6$ and $\|\Lambda\|_F^2 = 42$, so $r(E_7) \le \frac{5}{6}\sqrt{7/6}$. These directions have a uniform explanation. A minuscule weight w is a deep hole of L, and since $L_k \subseteq L$, its projection $P_k w$ onto the span of L_k lies in the Voronoi cell of L_k (for $x \in L_k$, Pythagoras turns $|w| \le |w - x|$ into $|P_k w| \le |P_k w - x|$), so it has norm at most $\mathrm{cov}(L_k)^2$. Along $\Lambda = I - nww^\top/|w|^2$ this gives $\min_f \langle C_f, \Lambda\rangle \ge -n\big(\sum_{k<n} \mathrm{cov}(L_k)^2/|w|^2 - \frac{n-1}{2}\big)$, with equality for the flags attaining every $\mathrm{cov}(L_k)^2$; for E_7 this is $-7 \cdot \frac{5}{6} = -\frac{35}{6}$. The same single-weight bound gives $r(D_5) \le 4/(3\sqrt{5})$, and a version with the plane of two minuscule weights gives the bounds for E_6 and $r(D_4) \le 2/3$; at D_4 and D_5 these are the exact values. The following is numerical (floating point) and is not used in any proof. 300 polar-vertex ascents for E_6 find no direction below $\sqrt{3}/2$ and return to it in 204 runs, so numerically $r(E_6) = \sqrt{3}/2$; if so, the radius in Proposition 26 for E_6 can be taken to be $27/100$. For E_7, over the 544,320 Ψ-minimal flags (their C_f are distinct), 300 ascents find no direction below $\frac{5}{6}\sqrt{7/6}$ and return to it in 281 runs, with 51,840 flags active there; the next local maximum found has $r \approx 1.478$. Among the directions $I - cP$ with P the projection onto the span of a set of fundamental weights, ω_7 alone is the best. So numerically $r(E_7) = \frac{5}{6}\sqrt{7/6}$; if so, the radius for E_7 can be taken to be $6/25$.

Proposition 28 (Lifts over the previous critical lattice)**.** *Let $5 \le n \le 8$, let L_{n-1} be the critical lattice in dimension $n-1$ (D_4, D_5, E_6, E_7), with Coxeter number h, and let L have rank n and a shortest vector v such that $\pi_v L$ is similar to L_{n-1}. Then*

$$\gamma(L) \le c_n := \Big(\frac{h^2}{h^2-1}\,\gamma_{n-1}\Big)^{(n-1)/n}, \qquad \eta(L) \le c_n^{n/2} H_{n-1} = \Phi_n(L_{n-1}),$$

and $\eta(L) > H_n$ requires $\gamma(L) > (H_n/H_{n-1})^{2/n}$. The values are (upper bounds rounded up, required values rounded down):

n	L_{n-1}	h	*cap* c_n *on* $\gamma(L)$	*required* $\gamma(L)$	$\Phi_n(L_{n-1})$	H_n
5	D_4	*6*	*1.3496*	*1.3195*	*6.348*	*6*
6	D_5	*8*	*1.4329*	*1.6012*	*17.66*	*24.63*
7	E_6	*12*	*1.5577*	*1.6561*	*116.2*	*144*
8	E_7	*18*	*1.6864*	*2*	*1164.6*	*2304*

So for $n = 6, 7, 8$, $\eta(L) < H_n$. For $n = 5$, $\eta(L) > 6$ forces $\gamma(L) \in (2^{2/5}, c_5]$.

Proof. Apply Corollary 11 with $R = L_{n-1}$ and $\ell = n - 1$. Since L_{n-1} is critical, $\gamma(L_{n-1}) = \gamma_{n-1}$, which gives the cap on $\gamma(L)$, and $\eta(L) \le \gamma(L)^{n/2}\eta(L_{n-1}) = \gamma(L)^{n/2} H_{n-1}$ (Theorem 5). This gives the bound on $\eta(L)$, which equals $\Phi_n(L_{n-1})$ by Proposition 12, and the condition $\gamma(L) > (H_n/H_{n-1})^{2/n}$. The table evaluates these with $\gamma_4, \dots, \gamma_7$ from Table 1 and $H_4, \dots, H_8$ from Theorem 5; for $n = 5$, $(H_5/H_4)^{2/5} = 2^{2/5}$. □

The proposition does not use η_{n-1} or the conjecture, and the $n = 8$ row is consistent with Theorem 8. The lower end of the band for $n = 5$ is attained, by D_5^*: its projection along a shortest vector is D_4, its lifting form is the 2-torsion class (not near $\rho^\vee/6$), and both the split bound and the condition that v be shortest are equalities there, so $\eta(D_5^*) = 2^{5/5} \cdot 3 = 6$.

Remark 29 (Inductive structure). Dimension 5 is exceptional because the Coxeter number of D_4 is small, and this is one reason why η_5 is the hardest case: for the threshold 6, lifts over D_4 survive in the thin band of Proposition 28. For $5 \le n \le 8$, a lattice with $\eta(L) > H_n$ and a shortest vector v needs $\eta(\pi_v L) > H_n/\gamma_n^{n/2}$ by the split along v, and this threshold is exactly $\eta(\pi_v L_n)$ ($3/\sqrt{2}$, 16/3, 18, 144 for $n = 5, \dots, 8$). The following is heuristic only and is not used in any proof: for $n \ge 6$ the problem should localize to lifts over $\pi_v L_n$, that is, to a neighborhood of L_n (excluding $\pi_v L \cong L_{n-1}$ exactly does not exclude nearby lattices).

In terms of Φ_n: since L_n is a lift of $\pi_v L_n$ along a shortest vector, $\kappa(\pi_v L_n) \ge \gamma_n^{n/(n-1)}$, and the tightness of the split at root lattices gives $\Phi_n(\pi_v L_n) \ge H_n$. Numerically, $\kappa(\pi_v L_n) = \gamma_n^{n/(n-1)}$ locally for $L_n = E_6, E_7$. By contrast, Proposition 28 gives $\Phi_6(D_5) \approx 17.66 < H_6$, $\Phi_7(E_6) \approx 116.2 < H_7$ and $\Phi_8(E_7) \approx 1164.6 < H_8$.

Conjecture 30. *For $n = 6, 7, 8$, $\sup_M \Phi_n(M) = H_n$, attained exactly at $\pi_v L_n$ and $\pi_w(L_n^*)$ (v and w shortest vectors), which coincide for $n = 8$.*

Proposition 10 does not use η_{n-1}, so a certificate of Conjecture 30 for $n = 6$ or 7 would prove $\eta_6 = 128\sqrt{3}/9$ or $\eta_7 = 144$ directly, with no second stage. In dimension 5 this route stops at exactly $2^{16}/(3^4 \cdot 5^3) \approx 6.4727 > H_5$ (Remark 18, Theorem 20), which is why the proof of Theorem 14 needs one. Since Φ_5 has a local maximum above $\Phi_5(D_4)$ (at M_*), the first check for Conjecture 30 is a numerical search. The evidence for $n = 6$ is numerical (floating point) and is not used in any proof. 464 Nelder–Mead ascents of Φ_6[2] found no value above H_6: every end point above 22 re-ascends to a copy of $\pi_v E_6$ or of $D_5^* \cong \pi_w(E_6^*)$, and the best come within 10^{-5} of H_6. Both points appear to be sharp local maxima, with first-order inradius at least 0.314 and 0.344 (maximizers searched among the 4-torsion points). Along the geodesic from $\pi_v E_6$ to the nearest copy of D_5^* (distance 1.0224), Φ_6 falls monotonically to about 18.24 and rises back to H_6 (`numerics/scan.sage`), so the two peaks are separated by a deep valley, down to 74% of H_6. In dimension 5, by contrast, the geodesic from D_4 towards $A_2 \otimes A_2$ first rises, to $\Phi_5(M_*) \approx 6.4727$ (Figure 1), and the same ascent started near D_4 exceeds $\Phi_5(D_4)$ in half of the runs. For $n = 7$, numerically $\Phi_7(E_6^*) = 144 = H_7$. A narrow region with $\Phi_6 > H_6$ could still have been missed.

7.2 Numerical and computational evidence

Numerical evidence. The following observations are numerical and are not used in any proof. A Voronoi-style ascent of η in floating point (active set of near-minimal flags plus an LP step) gives:

- $n = 4$: 24 random starts all end at ≤ 3. Most converge to D_4; a second local maximum sits near 2.145, consistent with $\eta(A_4) = 5!/5^{5/2} \approx 2.1466$.
- $n = 5$: no start exceeds 6. Starts near D_5 return to ≈ 6, further local maxima sit near 5.33 (consistent with $\eta(\pi_v E_6) = 16/3$), and one converged run ends near 5.03 at a lattice we have not identified (4 pairs of minimal vectors, not isodual). (For Φ_5 see Remark 23.)

Computer-assisted global bounds. The proofs of Theorems 19, 20 and 14 cover a space of Ψ-minimal bases of determinant 1 by boxes in interval arithmetic, each certified by an integer flag or sublattice or discarded by the window constraints of Appendix C; the method and the leaf counts are in Appendices C, D and E. The window constraints use $\mu_2 = \gamma_2$, $\mu_3 = \gamma_3$ and $\mu_4 = \gamma_4$

[2] 128 random starts; 64 seeded at 5-torsion points; 48 near $\pi_v E_6$; 32 each near D_5 and seeded at torsion points of orders $2, 3, 4, 6$; 16 each near A_5, A_5^*, D_5^* and $\mathbb{Z}^5$. Here κ is computed by branch and bound over φ, and η by exact recursion at the end points.

(Theorem 35, which for $m = 4$ uses Theorem 7), so all three computations depend on those two theorems. An earlier version of the method also reproduced $\eta_4 = 3$.

8 Asymptotics

Theorem 31. $\log \eta_n = (\frac{1}{4} + o(1))\, n^2 \log n$.

The upper bound is elementary. The lower bound is essentially Laarhoven's area rule [26, Prop. 7.1], which rests on Kim's mean-value formula for flags [20, Thm. 5 and Corollary]; we add only the uniformity over all profiles (Steps 1–2 below).

Proof. Upper bound. Iterating Lemma 3 with $k = 1$ gives $\eta_n \le \prod_{m \le n} \gamma_m^{m/2}$, and $\gamma_m \le 1 + m/4$, so $\sum_{m \le n} \frac{m}{2} \log \gamma_m = (\frac{1}{4} + o(1))\, n^2 \log n$.

Lower bound. Let L be Haar-random with $\det L = 1$ (covolume 1). Kim and Laarhoven use covolumes, so write $\operatorname{vol} M = \sqrt{\det \operatorname{Gram}(M)}$ and, for a basis B, $P_k(B) = \log \operatorname{vol} L_k = \frac{1}{2} \log G_k$. Then $\log \Psi^2(B) = 2 \sum_{k=1}^{n-1} P_k(B)$, since $P_n = 0$. It suffices to show:

> $(\star)$ for each $\varepsilon > 0$, with probability $1 - o(1)$ every basis of L has $\sum_k P_k(B) \ge (\frac{1}{8} - \varepsilon)\, n^2 \log n$,

because then $\eta_n \ge \eta(L) \ge n^{(1/4 - 2\varepsilon) n^2}$. We use two inputs from [20], with $V(m)$ the volume of the unit m-ball and the convention $\zeta(1) = 1$:

(K1) [20, Thm. 3, with $k = 1$] (also [41, 44]): for $1 \le d \le n - 1$, the average, over Haar-random lattices L of determinant 1, of the number of $A \in \operatorname{Gr}(L, d)$ (the primitive rank-d sublattices) with $\det A \le H$ is $a(n, d) H^n$, where $a(n, d) = \frac{1}{n} \binom{n}{d} \prod_{i=1}^{d} \frac{V(n-i+1)\zeta(i)}{V(i)\zeta(n-i+1)}$. Here Kim's $\det A = \|v_1 \wedge \cdots \wedge v_d\|$ is the covolume $\operatorname{vol} A$, the square root of the Gram determinant.

(K2) [20, Thm. 5 and Corollary]: for $0 = d_0 < d_1 < \cdots < d_k < d_{k+1} = n$, the average number, over Haar-random lattices of determinant 1, of flags $A_1 \subseteq \cdots \subseteq A_k$ of primitive sublattices with $\operatorname{rk} A_i = d_i$ and $\det A_i \le H_i$ is $a(n, \mathbf{d}) \prod_{i=1}^{k} H_i^{d_{i+1} - d_{i-1}}$, where $a(n, \mathbf{d}) = a(n, d_1) \prod_{i=1}^{k-1} \frac{n - d_{i-1}}{d_{i+1} - d_{i-1}}\, a(n - d_i, d_{i+1} - d_i)$. For complete flags ($d_i = i$, $k = n - 1$) every exponent $d_{i+1} - d_{i-1}$ is 2, and $a(n, \mathbf{d}) = C_n := \frac{n!}{2^{n-1}} \prod_{m=2}^{n} \frac{V(m)}{2\zeta(m)}$. Indeed $a(m, 1) = V(m)/(2\zeta(m))$, since $V(1) = 2$ and $\zeta(1) = 1$, and the recursion gives $a(n, \mathbf{d}) = a(n, 1) \prod_{i=1}^{n-2} \frac{n-i+1}{2}\, a(n - i, 1)$, where $\prod_{i=1}^{n-2} (n - i + 1) = n!/2$. So the expected number of complete flags with $\operatorname{vol} M_k \le X_k$ for all k is $C_n \prod_k X_k^2$. Since $\log V(m) = -\frac{m}{2} \log m + O(m)$, $\log C_n = -\frac{1}{4} n^2 \log n + O(n^2)$; cf. [26, (7.3)].

Step 1: no tiny sublattices. The product in $a(n, r)$ equals $\prod_{j=1}^{n} V(j) \big/ \big(\prod_{i \le r} V(i) \prod_{i \le n-r} V(i) \big)$ up to the ζ factors, so $a(n, r) = a(n, n - r)$. For $r \le n/2$ each ratio $V(n - i + 1)/V(i)$ is at most $V(5)/V(1) < 3$, the ζ factors are at most $\zeta(2)^n$ and $\binom{n}{r} \le 2^n$, so $a(n, r) \le e^{O(n)}$ for all r. Taking $H = n^{-n}$ in (K1) and summing over r, with probability $1 - o(1)$ every primitive sublattice of L has $\log \operatorname{vol} \ge -n \log n$, so every basis has $P_k \ge -n \log n$ for all k.

Step 2: discretize. Put $\delta = \varepsilon/2$. Suppose Step 1 holds and some basis has $\sum_k P_k \le (\frac{1}{8} - \varepsilon) n^2 \log n$. Then each $P_k \le (\frac{1}{8} - \varepsilon) n^2 \log n + n \cdot n \log n \le \frac{9}{8} n^2 \log n$. Set $h_k = \delta \lceil P_k / (\delta n \log n) \rceil$. Then h lies in the fixed set $\mathcal{H}_n = (\delta \mathbb{Z} \cap [-1, 2n])^{n-1}$ of size $n^{O(n)}$, the prefix flag of the basis has $\operatorname{vol} L_k \le n^{n h_k}$ for all k, and $\sum_k h_k \le (\frac{1}{8} - \varepsilon + \delta) n$.

Step 3: union bound. For each $h \in \mathcal{H}_n$ with $\sum_k h_k \le (\frac{1}{8} - \varepsilon + \delta) n$, (K2) bounds the expected number of flags with $\operatorname{vol} M_k \le n^{n h_k}$ by $C_n \exp\big(2n \log n \sum_k h_k\big) \le \exp\big(-2(\varepsilon - \delta) n^2 \log n + O(n^2) \big)$.

Summing over the $n^{O(n)}$ vectors h and applying Markov's inequality, $(\star)$ fails with probability at most $o(1)+n^{-\varepsilon n^2+O(n^2/\log n)}=o(1)$. In particular, for large n such a lattice exists, so $\eta_n \ge \eta(L) \ge n^{(1/4-2\varepsilon)n^2}$. $\square$

Theorem 31 replaces the $O(n^3)$ bound of the 2007 paper and determines the leading constant. Only $\log C_n$ up to $O(n^2)$ is used in the proof; the exact C_n enters Corollary 32.

Corollary 32 (Second-order term)**.** *Let* $c_{\mathrm{KL}} = 2^{2-2\cdot 0.599} \approx 1.7435$. *Then*

$$\tfrac14 n^2 \log n - (a+o(1))\, n^2 \;\le\; \log \eta_n \;\le\; \tfrac14 n^2 \log n - (b+o(1))\, n^2,$$

with $a = \frac18 + \frac14(1+\log 2\pi) \approx 0.8345$ *and* $b = \frac18 + \frac14 \log(2\pi e/c_{\mathrm{KL}}) = a - 0.2005\log 2 \approx 0.6955$.

Proof. Two sums. By Euler–Maclaurin, $\sum_{m\le n} \frac{m}{2}\log m = \frac14 n^2 \log n - \frac18 n^2 + O(n\log n)$. By Stirling's formula, $\log V(m) = \frac{m}{2}\log \pi - \log\Gamma(\frac{m}{2}+1) = -\frac{m}{2}\log m + \frac{m}{2}(1+\log 2\pi) + O(\log m)$.

The constant C_n. In $\log C_n = \log n! - (n-1)\log 2 + \sum_{m=2}^{n} \big(\log V(m) - \log 2 - \log \zeta(m)\big)$ the terms other than $\log V(m)$ contribute $O(n\log n)$, since $1 \le \zeta(m) \le \zeta(2)$. With the two sums,

$$\log C_n = -\tfrac14 n^2 \log n + a\, n^2 + O(n\log n).$$

Lower bound. We show $\log \eta_n \ge -\log C_n - 2\varepsilon n^2$ for each $\varepsilon > 0$ and all large n, which with the formula for $\log C_n$ gives the lower bound. Let L be Haar-random with $\det L = 1$, and use the notation of the proof of Theorem 31. Put $S = -\frac12 \log C_n - \varepsilon n^2$, so that $S \le (\frac18 + o(1))\, n^2 \log n$, and $\delta = (\log n)^{-2}$. By Step 1 of that proof, with probability $1-o(1)$ every primitive sublattice of L has $\log \mathrm{vol} \ge -n\log n$. Suppose this holds and some basis B has $\sum_k P_k(B) \le S$. Then $-n\log n \le P_k \le S + n^2 \log n \le 2n^2 \log n$ for every k and large n. Put $h_k = \delta\lceil P_k/(\delta n \log n)\rceil$. Then $h = (h_k)$ lies in $\mathcal{H}_n = (\delta\mathbb{Z} \cap [-1, 2n+1])^{n-1}$, which has at most $((2n+2)/\delta + 1)^{n-1} = e^{O(n\log n)}$ elements, since $\log(1/\delta) = 2\log\log n$. The prefix flag of B has $\mathrm{vol}\, L_k \le n^{nh_k}$ for all k, and

$$n\log n \sum_k h_k \le \sum_k P_k + (n-1)\,\delta\, n \log n \le S + \frac{n^2}{\log n}.$$

For each such h, (K2) bounds the expected number of complete flags with $\mathrm{vol}\, M_k \le n^{nh_k}$ for all k by

$$C_n \exp\Big(2n\log n \sum_k h_k\Big) \le C_n \exp\Big(2S + \frac{2n^2}{\log n}\Big) = \exp\Big(-2\varepsilon n^2 + \frac{2n^2}{\log n}\Big).$$

Summing over the $e^{O(n\log n)}$ vectors h gives $e^{-2\varepsilon n^2 + o(n^2)} = o(1)$. By Markov's inequality, with probability $1-o(1)$ every basis of L has $\sum_k P_k(B) > S$, so $\eta(L) = \min_B \Psi^2(B) > e^{2S}$. Hence $\log \eta_n \ge 2S = -\log C_n - 2\varepsilon n^2$ for large n.

Upper bound. By the proof of Theorem 31, $\log \eta_n \le \sum_{m=2}^{n} \frac{m}{2}\log \gamma_m$. The Kabatiansky–Levenshtein bound $\Delta_m \le 2^{-0.599 m(1+o(1))}$ on the packing density [19] (see [8, Ch. 1]) and $\gamma_m = 4(\Delta_m/V(m))^{2/m}$, with $V(m)^{-2/m} = (1+o(1))\, m/(2\pi e)$ by Stirling, give $\log \gamma_m \le \log m - \log(2\pi e/c_{\mathrm{KL}}) + \varepsilon_m$ with $\varepsilon_m \to 0$. Given $\varepsilon > 0$, choose m_0 with $\varepsilon_m \le \varepsilon$ for $m > m_0$; then $\sum_{m\le n} \frac{m}{2}\varepsilon_m \le O_{m_0}(1) + \frac14 \varepsilon n^2$, so $\sum_{m\le n} \frac{m}{2}\varepsilon_m = o(n^2)$. With the first sum,

$$\log \eta_n \le \tfrac14 n^2 \log n - \tfrac18 n^2 - \tfrac14 \log(2\pi e/c_{\mathrm{KL}})\, n^2 + o(n^2) = \tfrac14 n^2 \log n - (b+o(1))\, n^2.$$

Finally $\frac14 \log c_{\mathrm{KL}} = \frac14(2 - 1.198)\log 2 = 0.2005 \log 2$, so $b = a - 0.2005\log 2$. $\square$

The lower constant a is exactly what $\sum_{m\le n}\frac{m}{2}\log\gamma_m$ gives with the Minkowski–Hlawka value $\gamma_m \sim m/(2\pi e)$ [15]. So the second-order gap for η_n is the classical gap in the asymptotics of Hermite's constant.

Remark 33 (Flag heights). Following Thunder [43], Kim [20] defines the height of a flag $A_1 \subseteq \cdots \subseteq A_k$ as $\prod_i (\operatorname{vol} A_i)^{d_{i+1}-d_{i-1}}$. For a complete flag this is $\prod_k (\operatorname{vol} L_k)^2 = \prod_k G_k$, so $\Psi^2(B)$ is exactly the height of the flag of B, and $\eta(L)$ is the minimal normalized height of a complete flag of L, consistent with Meyer's description of η_n as a flag-height constant. Kim also shows that the *average* number of complete flags of height $\le H$ is infinite [20, Corollary to Thm. 5]; this is why Step 1, which excludes very small sublattices, is needed before the union bound over ceilings.

9 Consequences for reduction: the Ψ-Hermite constant

Definition 34. The *Ψ-Hermite constant* is $\mu_m = \sup r_1(B)/\Delta^{1/m}$, the supremum over m-dimensional lattices and their Ψ-minimal bases B.

Trivially $\gamma_m \le \mu_m$. Size reduction does not change Ψ, and in a Ψ-minimal basis no swap of neighbours decreases it, so $r_{i+1} \ge \frac{3}{4} r_i$; hence $\mu_m \le (4/3)^{(m-1)/2}$. For $m = 2$, $\Psi^2 = r_1$, so $\mu_2 = \gamma_2$. In a Ψ-minimal basis b_1 need not be a shortest vector.

Theorem 35 (Key inequality). $\mu_m^{m(m-1)/2} \le \eta_m \prod_{k<m} \mu_k^k$.

Proof. Normalize $\Delta = 1$. Each prefix L_k of a Ψ-minimal basis is Ψ-minimal, so $G_k \ge (r_1/\mu_k)^k$. Hence $\eta_m \ge \Psi^2 = \prod_{k<m} G_k \ge r_1^{m(m-1)/2} / \prod_{k<m} \mu_k^k$. □

If $\eta_m = H_m$ and $\mu_k = \gamma_k$ for $k < m$, the right-hand side is $\gamma_m^{m(m-1)/2}$. This is where the exact values of §§4–5 pay off.

Corollary 36.

1. *If $\eta_m = H_m$ and $\mu_k = \gamma_k$ for all $k < m$, then $\mu_m = \gamma_m$. In particular $\mu_3 = \gamma_3$, $\mu_4 = \gamma_4$ and, by Theorem 14, $\mu_5 = \gamma_5$.*
2. *If η_6 and η_7 are as conjectured, then $\mu_m = \gamma_m$ for all $m \le 8$.*

Corollary 37 (Quasi-polynomial bound). *For all m, $\mu_m \le \gamma_2 \prod_{j=3}^{m} \gamma_j^{1/(j-1)} = \exp\big((\frac{1}{2} + o(1))\log^2 m\big)$.*

Proof.

- Write $a_m = \log \mu_m$ and $e_m = \sum_{j\le m} \frac{j}{2}\log\gamma_j$, which is the iterated $k = 1$ bound on $\log \eta_m$.
- The right-hand side of Theorem 35 is increasing in each a_k, so $a_m \le A_m$, where A_m is defined by equality.
- Subtracting the equality for $m-1$ from that for m and rearranging gives $A_m = A_{m-1} + \log\gamma_m/(m-1)$.

□

This compares with the exponential bound $(4/3)^{(m-1)/2}$ from LLL. Large m needs care: $\log H_m \approx m^2/4$ while $\log \eta_m \sim \frac{1}{4} m^2 \log m$ (Theorem 31). So $\eta_m = H_m$ fails eventually, and from then on Theorem 35 no longer gives $\mu_m = \gamma_m$. Since it is only an upper bound, this says nothing about whether $\mu_m = \gamma_m$ itself fails (§10).

Blockwise reduction. Theorem 35 is what a quality bound for blockwise potential reduction needs: it is the ingredient that a primal–dual window argument, as in slide reduction [13], uses to control the Hermite factor of bases whose windows of m consecutive projected vectors are all Ψ-minimal. We treat blockwise Ψ_m reduction, its δ-relaxation and the cost of the Ψ-minimization oracle in a separate note.

Numerically (floating-point choice of lattices, exact Ψ-minimization), among perturbations of A_m, D_m, E_m and their duals at four scales, with hill climbs on r_1^m/Δ (200, 192 and 128 lattices for $m = 6, 7, 8$), no Ψ-minimal basis with $r_1 > \gamma_m \Delta^{1/m}$ was found; the largest values of $r_1^m/(\gamma_m^m \Delta)$ were 0.89, 0.84 and 0.61. In about 7–8% of the lattices the first vector of some Ψ-minimal basis is not a shortest vector, though it satisfies the bound (`numerics/mu_check.sage`).

10 Open problems

1. **A shorter proof of $\eta_5 = 6$.** Theorem 14 rests on about 9×10^9 certified leaves, almost all of them in the second stage, over the lifts of the region where $\Phi_5 \geq 6$. Numerically that region is a neighborhood of D_4 and M_* (Figure 2), the lifts over it reach $\eta = 6$ only at D_5^* (over D_4), and over M_* they stay below about 5.2 (Remark 23). A sharper form of Proposition 10, one that sees that the lifts over M_* are far from tight, or a local argument at D_5^* of the kind used for $\pi_v D_5$, would replace most of the computation.

2. **η_6, η_7.** Certify Conjecture 30: a search over 5- (resp. 6-) dimensional lattices bounding Φ_n above, with a local certificate for the combined min-type (η) and max-min-type (κ) function at $\pi_v L_n$ and $\pi_w(L_n^*)$. Alternatively, with $\beta_n = \sup_L \min(\gamma(L), \gamma(L^*))$: since $\eta_5 = 6$, $\beta_6 = \gamma(E_6^*)$ would give $\eta_6 = 128\sqrt{3}/9$ (for $n = 5$ the analogue fails numerically, $\beta_5 \geq 1.371 > \gamma(D_5^*)$).

3. **Maximizers of Φ_n.** At the known local maxima of Φ_5 and Φ_6 the κ-maximizers are torsion points: a single class of order 2 at $\pi_v D_5$, of order 3 at M_* and of order 2 at $\pi_v E_6$, and 24 classes up to sign, all of order 6 and forming one orbit, at D_4 (Remark 23). Are they always torsion points forming a single orbit under the automorphism group? (Proved at $\pi_v D_5$, M_* and D_4, numerical at $\pi_v E_6$; for D_4 see Proposition 12.)

4. **Hermitian and symplectic families.** Do the maximizers of Φ_n, or the lattices that block the reduction to Φ_n, lie on Hermitian or symplectic families in other dimensions? (E_6 is a rank-3 $\mathbb{Z}[\omega]$-lattice, but rank-3 Hermitian lattices need not be isodual.)

5. **Uniqueness of E_8.** Locally this is proved: by Proposition 26, every lattice within distance $1/20$ of E_8 other than E_8 has $\eta < 2304$. Globally it is open. Take a flag starting at a shortest vector. This gives $\eta(L) \leq \lambda_1^8 \eta_7$, so $\eta_7 = 144$ forces $\lambda_1^2 = 2$, i.e. L critical, i.e. E_8. Alternatively, equality in the $k = 4$ split forces $\gamma_{8,4}(L) = 4$ and forces both M and πL (and their duals) to be similar to D_4; this structure may suffice without η_7.

6. **Eutaxy radii.** The first-order radius $r(L)$, the largest r with $\min_f \langle C_f, \Lambda \rangle \leq -r\|\Lambda\|_F$ for every traceless Λ, is $2/3$ for D_4 and $4/(3\sqrt{5})$ for D_5 (Appendix A). For E_6 and E_7 we know only $3/10 \leq r(E_6) \leq \sqrt{3}/2$ and $2\sqrt{7}/27 \leq r(E_7) \leq \frac{5}{6}\sqrt{7/6}$ (Proposition 26 and Remark 27); numerically both upper bounds are the true values. Determine them. In each known case the minimizing direction is $I - cP$ with P the projection onto a span of minuscule weights. Is there a structural proof, for example a higher-moment version of the design argument?

7. **Asymptotics.** Determine the second-order term: if $\log \eta_n = \frac{1}{4} n^2 \log n - cn^2 + o(n^2)$, then $c \in [b, a] \approx [0.6955, 0.8345]$ by Corollary 32. The lower end of this interval is tied to the

Kabatiansky–Levenshtein bound for Hermite's constant, so narrowing it needs either a better upper bound for γ_m or an argument that does not go through $\prod \gamma_m^{m/2}$.

8. **The Ψ-Hermite constant.** Is $\mu_m = \gamma_m$ for all m? It holds for $m \le 5$ (Corollary 36), and Theorem 35 stops giving it once $\eta_m > H_m$. Failing that, is μ_m bounded by a polynomial in m? Corollary 37 gives only $\exp\big((\frac{1}{2} + o(1)) \log^2 m\big)$, against $\gamma_m = \Theta(m)$. Is there a direct proof of $\mu_5 \le \gamma_5$? Corollary 36 obtains it from Theorem 14, so it rests on the whole computer-assisted proof, although the obstacle there, the peak at M_*, has nothing to do with first vectors.

9. **Symplectic lattices.** For symplectic lattices, such as NTRU lattices, one may restrict the minimum to symplectic bases ($BJB^T = J$ for the standard alternating J, up to scaling). Their Gram–Schmidt profiles are palindromic, so at determinant 1 the prefix determinants satisfy $G_{2m-k} = G_k$ and $\Psi^2 = G_m \prod_{k<m} G_k^2$. This defines $\eta^{sp}(L) \ge \eta(L)$, the constant relevant to symplectic reduction [12]. Is $\eta^{sp} = \eta$ on symplectic lattices, i.e. does every symplectic lattice have a symplectic Ψ-minimal basis?

Data and code availability

- *Code:* the checkers, the coverage scripts, the exact certificates, the hash manifests, a worked example and instructions for verifying the results at four levels of effort are at `https://github.com/nickhg/isodual_constant`; file names in the text refer to it. Release v1.1 matches this version of the paper, and all releases are archived on Zenodo, doi:10.5281/zenodo.23650216.

- *Data:* the certificate logs (about 263 GB in about 70,000 files) are at `https://huggingface.co/datasets/nickhg1971/isodual-constant-certificates`, doi:10.57967/hf/11260 (revision `96e9321`).

The SHA-256 of every file of the dataset, and those of the logs as recorded when they were produced, are part of the archived code release, so any copy of the data can be checked against them.

Acknowledgements

AI tools (Anthropic's Claude) were used for literature search, checking algebra, writing and running the numerical and computer-assisted experiments, and drafting text. All results and text were checked by the author, who takes full responsibility for the content. We thank Brian Moulton for independently checking many of the computations, for a careful reading that found several errors, and for several results: the exact value $r(D_4) = 2/3$, the first (computational) proof that E_6 is a strict local maximum of η, the extension of Proposition 26 to E_8, and the explanation of the upper bounds on the eutaxy radii via covering radii.

References

[1] A.-M. Bergé, Symplectic lattices, in *Quadratic Forms and Their Applications* (Dublin, 1999), E. Bayer-Fluckiger, D. Lewis and A. Ranicki (eds.), Contemp. Math. 272, Amer. Math. Soc., 2000, pp. 9–26.

[2] A.-M. Bergé and J. Martinet, Sur un problème de dualité lié aux sphères en géométrie des nombres, *J. Number Theory* 32 (1989), 14–42.

[3] N. Bourbaki, *Lie Groups and Lie Algebras, Chapters 4–6*, Springer, 2002.

[4] W. Bruns, B. Ichim, C. Söger and U. von der Ohe, *Normaliz*: algorithms for rational cones and affine monoids, `https://www.normaliz.uni-osnabrueck.de`.

[5] P. Buser and P. Sarnak, On the period matrix of a Riemann surface of large genus (with an appendix by J. H. Conway and N. J. A. Sloane), *Invent. Math.* 117 (1994), 27–56.

[6] H. Cohn and A. Kumar, Optimality and uniqueness of the Leech lattice among lattices, *Ann. of Math.* 170 (2009), 1003–1050.

[7] J. H. Conway and N. J. A. Sloane, On lattices equivalent to their duals, *J. Number Theory* 48 (1994), 373–382.

[8] J. H. Conway and N. J. A. Sloane, *Sphere Packings, Lattices and Groups*, 3rd ed., Springer, 1999.

[9] L. Ducas and W. van Woerden, NTRU fatigue: how stretched is overstretched?, *ASIACRYPT 2021*, LNCS 13093, 3–32.

[10] F. Fontein, M. Schneider and U. Wagner, PotLLL: a polynomial time version of LLL with deep insertions, *Des. Codes Cryptogr.* 73 (2014), 355–368.

[11] N. Gama, N. Howgrave-Graham, H. Koy and P. Q. Nguyen, Rankin's constant and blockwise lattice reduction, *CRYPTO 2006*, LNCS 4117, 112–130.

[12] N. Gama, N. Howgrave-Graham and P. Q. Nguyen, Symplectic lattice reduction and NTRU, *EUROCRYPT 2006*, LNCS 4004, 233–253.

[13] N. Gama and P. Q. Nguyen, Finding short lattice vectors within Mordell's inequality, *STOC 2008*, 207–216.

[14] M. Geck and G. Pfeiffer, *Characters of Finite Coxeter Groups and Iwahori–Hecke Algebras*, London Math. Soc. Monographs 21, Oxford University Press, 2000.

[15] E. Hlawka, Zur Geometrie der Zahlen, *Math. Z.* 49 (1943), 285–312.

[16] N. Howgrave-Graham, *Computational mathematics inspired by RSA*, PhD thesis, University of Bath, 1998.

[17] N. Howgrave-Graham, Isodual reduction of lattices, Cryptology ePrint Archive, Report 2007/105.

[18] J. E. Humphreys, *Reflection Groups and Coxeter Groups*, Cambridge University Press, 1990.

[19] G. A. Kabatiansky and V. I. Levenshtein, Bounds for packings on a sphere and in space, *Problems Inform. Transmission* 14 (1978), 1–17.

[20] S. Kim, Mean value formulas on sublattices and flags of the random lattice, *J. Number Theory* 241 (2022), 330–351, doi:10.1016/j.jnt.2022.03.013.

[21] P. Kirchner and P.-A. Fouque, Revisiting lattice attacks on overstretched NTRU parameters, *EUROCRYPT 2017*, LNCS 10210, 3–26.

[22] M. Kirschmer, *Definite quadratic and hermitian forms with small class number*, Habilitationsschrift, RWTH Aachen University, 2016; tables at `https://www.math.rwth-aachen.de/homes/Markus.Kirschmer/forms/`.

[23] A. Korkine and G. Zolotareff, Sur les formes quadratiques, *Math. Ann.* 6 (1873), 366–389.

[24] A. Korkine and G. Zolotareff, Sur les formes quadratiques positives, *Math. Ann.* 11 (1877), 242–292.

[25] B. Kostant, The principal three-dimensional subgroup and the Betti numbers of a complex simple Lie group, *Amer. J. Math.* 81 (1959), 973–1032.

[26] T. Laarhoven, Asymptotics of random-lattice enumeration: beyond HKZ via dense sublattices, arXiv:2610.15265 (2026).

[27] A. K. Lenstra, H. W. Lenstra and L. Lovász, Factoring polynomials with rational coefficients, *Math. Ann.* 261 (1982), 515–534.

[28] J. Martinet, *Perfect Lattices in Euclidean Spaces*, Grundlehren der mathematischen Wissenschaften 327, Springer, 2003.

[29] B. Meyer, *Constantes d'Hermite et théorie de Voronoï*, Thèse de doctorat, Université Bordeaux 1, 2008.

[30] B. Meyer, Generalised Hermite constants, Voronoi theory and heights on flag varieties, *Bull. Soc. Math. France* 137 (2009), 127–158.

[31] B. Meyer, Extreme lattices and vexillar designs, *J. Algebra* 322 (2009), 4368–4381.

[32] D. Micciancio and M. Walter, Practical, predictable lattice basis reduction, *EUROCRYPT 2016*, LNCS 9665, 820–849.

[33] G. Nebe and N. J. A. Sloane, *A Catalogue of Lattices*, `https://www.math.rwth-aachen.de/~Gabriele.Nebe/LATTICES/` (accessed October 2026).

[34] A. Neumaier, *Interval Methods for Systems of Equations*, Cambridge University Press, 1990.

[35] G. Pataki and M. Tural, On sublattice determinants in reduced bases, arXiv:0804.4014 (2008).

[36] C. Poor and D. S. Yuen, The Bergé–Martinet constant and slopes of Siegel cusp forms, *Bull. London Math. Soc.* 38 (2006), 913–924.

[37] R. A. Rankin, On positive definite quadratic forms, *J. London Math. Soc.* 28 (1953), 309–314.

[38] S. M. Rump, Verification methods: rigorous results using floating-point arithmetic, *Acta Numer.* 19 (2010), 287–449.

[39] A. Sato and M. Yasuda, A new BKZ-type reduction with provable termination and development of its self-dual variant, *Japan J. Indust. Appl. Math.* 43 (2026), doi:10.1007/s13160-026-00799-6.

[40] K. Sawatani, T. Watanabe and K. Okuda, A note on the Hermite–Rankin constant, *J. Théor. Nombres Bordeaux* 22 (2010), 209–217.

[41] W. M. Schmidt, Asymptotic formulae for point lattices of bounded determinant and subspaces of bounded height, *Duke Math. J.* 35 (1968), 327–339.

[42] A. Sibidanov, P. Zimmermann and S. Glondu, The CORE-MATH project, *29th IEEE Symposium on Computer Arithmetic (ARITH 2022)*, 26–34.

[43] J. L. Thunder, Asymptotic estimates for rational points of bounded height on flag varieties, *Compositio Math.* 88 (1993), 155–186.

[44] J. L. Thunder, Higher-dimensional analogs of Hermite's constant, *Michigan Math. J.* 45 (1998), 301–314.

[45] N. M. Vetchinkin, Uniqueness of classes of positive quadratic forms on which values of the Hermite constant are attained for $6 \leq n \leq 8$, *Trudy Mat. Inst. Steklov.* 152 (1980), 34–86.

[46] T. Watanabe, On an analog of Hermite's constant, *J. Lie Theory* 10 (2000), 33–52.

[47] J. Yamaguchi and M. Yasuda, Explicit formula for Gram–Schmidt vectors in LLL with deep insertions and its applications, *NuTMiC 2017*, LNCS 10737 (2018), 142–160.

[48] M. Yasuda, Self-dual DeepBKZ for finding short lattice vectors, *J. Math. Cryptol.* 14 (2020), 84–94.

A The local certificate at D_5

The three local certificates (this appendix and Appendix B) bound the flag terms with one lemma.

Lemma 38. *Let G_0 be positive definite, $G = G_0^{1/2} e^{\Lambda} G_0^{1/2}$ with Λ symmetric, and U a $k \times n$ matrix of rank k. Let P be the orthogonal projection onto the span of $G_0^{1/2} U^\top$. Then*

$$\log\det(UGU^\top) - \log\det(UG_0U^\top) \le \operatorname{tr}(P\Lambda) + \tfrac{1}{2}\|\Lambda\|_F^2 e^{\|\Lambda\|_2},$$

and if Λ is $n \times n$ and traceless, $\|\Lambda\|_2 \le \sqrt{(n-1)/n}\,\|\Lambda\|_F$.

Proof. $\log\det$ is concave, so $\log\det(UGU^\top) \le \log\det A + \operatorname{tr}(A^{-1}U(G-G_0)U^\top)$ with $A = UG_0U^\top$. Here $G - G_0 = G_0^{1/2}(\Lambda + E)G_0^{1/2}$ with $E = e^{\Lambda} - I - \Lambda$, and $\operatorname{tr}(A^{-1}UG_0^{1/2}XG_0^{1/2}U^\top) = \operatorname{tr}(PX)$. Since $e^{\lambda} - 1 - \lambda \ge 0$, $E \succeq 0$ and $\operatorname{tr}(PE) \le \operatorname{tr} E = \sum_i (e^{\lambda_i} - 1 - \lambda_i) \le \sum_i \frac{1}{2}\lambda_i^2 e^{|\lambda_i|} \le \frac{1}{2}\|\Lambda\|_F^2 e^{\|\Lambda\|_2}$. For the last claim, a traceless Λ with largest eigenvalue modulus λ has the other eigenvalues summing to $\mp\lambda$, so $\|\Lambda\|_F^2 \ge \lambda^2 + \lambda^2/(n-1)$. □

The inequality is global (not a Taylor remainder), and the factor $e^{\|\Lambda\|_2}$ appears once.

Center and coordinates. Since d is invariant under $G \mapsto UGU^\top$, the center may be any Gram matrix of D_5; we take its Cartan matrix G_0 ($\det G_0 = 4$). Write $G = G_0^{1/2} e^{\Lambda} G_0^{1/2}$ with Λ traceless and $t = \|\Lambda\|_F = d(G, G_0)$, and $Y = G_0^{1/2}\Lambda G_0^{1/2}$, so that linear functionals of Λ are rational functionals of Y. For a flag $f = (U_1, \dots, U_4)$ let $h_f(G) = \prod_{k \le 4} \det(U_k G U_k^\top)$; then $\eta(G) = \min_f h_f(G)/\det(G)^2$, and $\det G = \det G_0$.

First order. The 720 Ψ-minimal bases of D_5 are enumerated exactly (with the recursion of §2); they give 720 flags f with $h_f(G_0)/\det(G_0)^2 = 6$, with 720 distinct gradients. (The enumeration runs over all size-reduced bases satisfying the Lovász condition with $\delta = 1$: a branch-and-bound that visits 1,461 projected lattices and returns exactly these 720 bases. By heredity (§2) every two-dimensional projected block of a Ψ-minimal basis is Ψ-minimal, so its first vector is a shortest vector of the block and the Lovász condition holds with $\delta = 1$; size reduction changes neither the Gram–Schmidt norms nor the flag. Hence every Ψ-minimal flag is represented, and such bases satisfy $|b_1|^2 \le (4/3)^{(n-1)/2} \det^{1/n}$, so there are finitely many, and μ_5 is not needed. Only the stated properties of these 720 flags are used below; that they are all the minimal flags is not needed for the bound.) By Lemma 38, with P_k the orthogonal projection onto the span of $G_0^{1/2} U_k^\top$,

$$\log h_f(G) \le \log h_f(G_0) + L_f(\Lambda) + \sum_{k \le 4} \operatorname{tr}(P_k E), \qquad L_f(\Lambda) = \sum_{k \le 4} \operatorname{tr}(P_k \Lambda),$$

where $E = e^{\Lambda} - I - \Lambda \succeq 0$, so each $\operatorname{tr}(P_k E) \le \operatorname{tr} E$. In the coordinates Y, $L_f(\Lambda) = \operatorname{tr}(Q_f Y)$ with $Q_f = \sum_{k \le 4} U_k^\top (U_k G_0 U_k^\top)^{-1} U_k$, which is rational. For a basis of the space of functionals, take the traceless projections $X - \frac{1}{5}\operatorname{tr}(G_0^{-1}X)\,G_0$ of $X = E_{ii}$ and $X = E_{ij} + E_{ji}$ ($i < j$), in lexicographic order of (i, j), keep the first 14 that are linearly independent, take the dual basis of functionals,

and orthogonalize it by Gram–Schmidt in the dual norm; call the result $e_1, \dots, e_{14}$. For each signed axis an exact LP over the rationals finds the largest $t_j^{\pm}$ with $\pm t_j^{\pm} e_j \in \Gamma = \operatorname{conv}\{L_f\}$. So Γ contains the cross-polytope with half-axes $a_j = \min(t_j^+, t_j^-)\|e_j\|_*$, and with it the dual-norm ball of radius $r = (\sum_j a_j^{-2})^{-1/2} \geq 0.2392$. Equivalently, $\min_f L_f(\Lambda) \leq -rt$ for every Λ.

The proof needs only this lower bound, but the inradius of Γ is at most $4/(3\sqrt{5})$. Take $\Lambda_0 = I - 5P$, with P the orthogonal projection onto the spinor fundamental weight ω_4 (Bourbaki numbering), so that $Y_0 = G_0 - 4E_{44}$. Then $\min_f L_f(\Lambda_0) = -8/3$, attained by 120 of the 720 flags, and $\|\Lambda_0\|_F^2 = 20$, so no ball of radius larger than $(8/3)/\sqrt{20} = 4/(3\sqrt{5})$ fits in Γ.

The exact radius. In fact the inradius of Γ equals $4/(3\sqrt{5})$. This is not used elsewhere. In orthonormal coordinates of $\mathbb{R}^5 \supset D_5$, an orthogonal change of frame turns $L_f(\Lambda)$ into $\langle C_f, \Lambda\rangle$ with $C_f = \sum_{k\leq 4} P_{V_k} - 2I$, where V_k is spanned by the first k vectors of the flag. Replacing Λ by $-\Lambda$, the claim is that $\max_f \langle C_f, \Lambda\rangle \geq r_0\|\Lambda\|_F$ for every traceless symmetric Λ, with $r_0 = 4/(3\sqrt{5})$. The proof is computer-assisted.

- Signed permutations are automorphisms of D_5 and permute the flags, so it suffices to take Λ with $\Lambda_{11} \geq \cdots \geq \Lambda_{55}$ and $\Lambda_{1j} \geq 0$. In this region the extremal direction is $J - I$.
- Near $J - I$: the 120 flags with $\langle C_f, J - I\rangle = 8/3$ have tangential parts whose convex hull contains a ball of radius s, by an inscribed cross-polytope (26 exact LPs), with $s/r_0 > 27/100$. Writing a unit Λ at angle θ from $J - I$ as $\cos\theta$ times the unit vector of $J - I$ plus $\sin\theta\, w$, some active flag gives $\langle C_f, \Lambda\rangle \geq r_0\cos\theta + s\sin\theta \geq r_0$ whenever $\tan(\theta/2) \leq 27/100$.
- Elsewhere, write $\Lambda = \sum_i x_i E_i$ in an orthogonal rational basis and scale to $\max_i |x_i| = 1$. The 28 faces of $[-1, 1]^{14}$ are covered by 199,619 dyadic boxes. Each box either misses the region, or lies in the cap above, or has a rational $\sigma > 0$ and a rational convex combination z of the C_f with $\langle z, \Lambda\rangle - \|\Lambda\|_F^2/(2\sigma) \geq 8\sigma/45$ on the box. By the AM–GM inequality, the last condition gives $\langle z, \Lambda\rangle \geq r_0\|\Lambda\|_F$. Its left side is a separable concave quadratic in x, so its minimum over a box is the sum of coordinate-wise minima at the endpoints, computed exactly.

A separately written checker verifies, in exact rational arithmetic, that the flags are Ψ-minimal and invariant under signed permutations, the cap witnesses, and every box. The convex combinations come from a floating-point LP and are only candidates. With $r = 4/(3\sqrt{5})$ the radius argument below gives $\eta \leq 6$ for $d(G, G_0) \leq 6/25$ (checked in interval arithmetic). Theorem 13 is stated with $1/10$ because that is all the proof of Theorem 14 uses, and its certificate is much smaller.

In Meyer's terms [30, Def. 30], the spanning says that D_5 is perfect and the ball inside Γ that it is (strictly) eutactic for the staircase partition; the radius r quantifies the eutaxy.

Second order. By Lemma 38, $\operatorname{tr} E \leq \frac{1}{2}t^2 e^{\sqrt{4/5}\,t}$, so the four prefixes contribute at most $2t^2 e^{\sqrt{4/5}\,t}$. Hence

$$\log\eta(G) \leq \log 6 - rt + 2t^2 e^{\sqrt{4/5}\,t}.$$

Radius. The right-hand side is below $\log 6$ for $0 < t \leq \rho$ whenever $r > 2\rho e^{\sqrt{4/5}\,\rho}$. This holds for $\rho = 1/10$, checked in interval arithmetic (margin 0.0205); the largest admissible radius is about 0.1086. By isoduality, $G \mapsto JG^{-1}J$ maps the ball about G_0 isometrically onto a ball about a Gram matrix of D_5^* and preserves η.

For D_4 the same computation (144 minimal flags, 9 directions, remainder $3\operatorname{tr} E$ with $|\lambda| \leq \sqrt{3/4}\,t$) can be done with the exact inradius. Here Γ is 9-dimensional with 144 vertices and 105,184

facets, and its inradius is $r = 2/3$, attained along $\Lambda_0 = I - 2P$ with P the orthogonal projection onto the plane of two of the three minuscule weights $\omega_1, \omega_3, \omega_4$: for these, $\min_f L_f(\Lambda_0) = -4/3$, attained by 54 flags, and $\|\Lambda_0\|_F^2 = 4$. Hence $\eta \le 3$ for $d(G, G_0) \le 33/100$, with equality only at G_0. The facets were enumerated with Normaliz [4], and every facet was checked against all 144 vertices.

Artefacts. `d5_lemma.sage` (about three minutes in Sage); `radius_d4d5.sage` (about a minute; it also needs Normaliz) for the upper bound at D_5 and the exact radius at D_4; and, for the exact radius at D_5, the producer `rad_produce.py` (about ten minutes on eight cores; it writes the 448 logs of the cover) and the checker `rad_check.py` (about a minute, exact rational arithmetic), with its inputs (the 720 bases and the cap witnesses). All are in `certificates/` of the code release.

B Local certificates for Φ_5

This appendix proves Propositions 21 and 22. The two proofs share one scheme; the differences are noted at the end.

Setting. Fix the center G_0 (a Gram matrix of $\pi_v D_5$ or of M_*) and write $G = G_0^{1/2} e^{\Lambda} G_0^{1/2}$ with Λ symmetric and traceless (equal determinant), $t = \|\Lambda\|_F = d(G, G_0)$. Then $\Phi_5(G) = \kappa_{\mathrm{raw}}(G)^2 \det(G)^{-1/2} \tilde{\eta}(G)$, where $\kappa_{\mathrm{raw}}(G) = \sup_{\varphi} \min_{s \neq 0} k_s$, $k_s = q_s/(1 - \|\varphi(s)\|^2)$, $q_s = s^{\top} G s$, and $\tilde{\eta}(G) = \min_f h_f(G)/\det(G)^{3/2}$, $h_f = \prod_{k \le 3} \det(U_k G U_k^{\top})$ over flags $f = (U_k)$. Put $\psi(\tau) = -\log(1 - \tau^2)$, convex and increasing on $[0, 1)$. At G_0 the maximum over φ is attained at a class φ_0, with a finite set of active s (those with $k_s = \kappa_{\mathrm{raw}}$) and of minimal flags.

Remainder bounds (by hand).

- *Vectors.* $\log q_s(G) \le \log q_s(G_0) + \ell_s(\Lambda) + t^2/4$, with ℓ_s linear. Indeed $\log(u^{\top} e^{\Lambda} u)$ for a unit vector u is the log moment generating function of the spectrum of Λ under the weights u, and the spectrum has spread at most $\sqrt{2}\, t$; Hoeffding's lemma bounds the excess over the mean by $(\sqrt{2}t)^2/8$.

- *Flags.* $\log h_f(G) \le \log h_f(G_0) + L_f(\Lambda) + 3 \operatorname{tr} E$, with L_f linear and $\operatorname{tr} E \le \frac{1}{2} t^2 e^{\sqrt{3/4}\, t}$: Lemma 38 for each of the three prefixes, in rank 4.

So $\log \Phi_5(G) - \log \Phi_5(G_0) \le$ (a first-order term in Λ, φ) $+ \operatorname{rem}(t)$ with $\operatorname{rem}(t) = \frac{1}{2} t^2 + \frac{3}{2} t^2 e^{\sqrt{3/4}\, t}$ (the first part from $2 \times t^2/4$ for κ^2).

Exterior and interior. Put $U = ct$ for a constant c. A φ is *exterior* if some active s has $\|\varphi(s)\|$ at least U away from its value at φ_0, on the side that lowers k_s. There, convexity of ψ bounds $\psi(\|\varphi(s)\|)$ by its value at φ_0 minus a multiple of U. With $B \ge \max_{\|\Lambda\|_F = 1} [2\ell_s(\Lambda) + \min_f L_f(\Lambda)]$, computed as a dual-norm bound over a convex combination of the L_f, the exterior condition reads $B\, t + \operatorname{rem}(t) <$ (gain from U). Otherwise φ is *interior*: $\varphi = \pm\varphi_0 + \Delta$ with no active s moved U or more to the lowering side. One then has to show that there is a single such component, i.e. that this forces Δ to be close to 0 modulo integer forms, with every active $|s \cdot \Delta|$ of order U.

- At $\pi_v D_5$ this follows from the Smith form: the 6 half vectors generate $\mathbb{Z}^4$, so Δ is determined modulo $\mathbb{Z}^4$ by the values $s \cdot \Delta$ on them.

- At M_* it is a rigorous cover of the torus $[0,1]^4$ by exact rational cells. Each cell is exterior; or *far* (some short s has $q_s(G_0)/(1-\max_{\text{cell}}\|\varphi(s)\|^2)$ below the threshold the ball can reach); or *near* $\pm\varphi_0$, with $|s\cdot\Delta|<\frac{1}{6}$ for every active s and a fixed nearest integer shift. In a near cell the exterior case fails, so $\sigma_s\, s\cdot\Delta > -U$ for every third vector s (where $\|\varphi(s)\| = \frac{1}{3}+\sigma_s\, s\cdot\Delta$). The third vectors are positively dependent, so the polytope $\{\Delta : \sigma_s\, s\cdot\Delta \geq -1 \text{ for every third } s\}$ is bounded, and therefore $|s\cdot\Delta| \leq C_s U$ for every active s, with C_s the maximum of $|s\cdot\Delta|$ over its vertices. This is what makes the second-order remainder below $O(t)$.

The polytope. On the interior, ψ is replaced by its tangent at φ_0 plus a second-order remainder (θ,ε at $\pi_v D_5$; ψ'' on $[0,\frac{1}{3}+C_sU]$ at M_*). The first-order term is then $\max_\Delta \min_s[\ell_s(\Lambda)+a_s\, s\cdot\Delta]+\min_f L_f(\Lambda)$. By LP duality it equals $\min_{g\in\Gamma} g(\Lambda)$ with

$$\Gamma = \Gamma_\kappa + \operatorname{conv}\{L_f\}, \qquad \Gamma_\kappa = \Big\{ \sum_s \lambda_s \ell_s : \lambda \geq 0,\ \sum_s \lambda_s = 2,\ \sum_s \lambda_s a_s s = 0 \Big\}.$$

If Γ contains the dual-norm ball of radius r about 0, then $\min_\Gamma g(\Lambda) \leq -rt$. The facets of Γ are computed exactly over $\mathbb{Q}$, every facet is checked to have positive right-hand side, and r is the smallest facet distance. The final scalar inequalities (interior: $r > \operatorname{rem}(t)/t+$ second-order terms; exterior: gain $> B + \operatorname{rem}(t)/t$; all increasing in t, so checked at $t=\rho$) are verified in interval arithmetic.

Data.

	$\pi_v D_5$	M_*
torsion of φ_0	2	3
active pairs	6 half + 7 integer	9 third + 3 integer
minimal flags	36	54
Γ as used in the proof	432 vertices, 14199 facets*	468 vertices, 11162 facets
unperturbed inradius	$1/(3\sqrt{3})$	$11/(16\sqrt{2})$
(ρ, c)	(1/20, 11/20)	(1/40, 9/10)
B	≤ 1.1686	≤ 1.1963
margins (interior, exterior)	0.0277, 0.1276	0.0691, 0.0460

*The perturbed polytope $\Gamma(\theta,\varepsilon)$ at $\rho = 1/20$; the unperturbed Γ at $\pi_v D_5$, whose inradius is $1/(3\sqrt{3})$, has 6483 facets. At $\pi_v D_5$ the radius used is that of the perturbed polytope $\Gamma(\theta,\varepsilon)$, $r = 0.1310$; at M_* the second-order terms are kept outside the polytope, and $r = 11/(16\sqrt{2})$ is used directly. At M_* the cover has 3560 cells (3046 exterior, 514 near, 0 far). At $\rho = 1/20$ the M_* interior inequality fails, because of the second-order term of the integer vectors; $\rho = 1/40$ suffices.

Artefacts. `pivd5_lemma.sage` and `mstar_lemma.sage`, with their input data (`pivd5_data.json`, `mstar_data.json`, written and independently re-verified) and outputs; about one hour each in Sage. They are in `certificates/` of the code release.

C The certificate for $\eta_5 < 7$

Result and inputs. The computation certifies $\sup_M \Phi_5(M) < 7$, the supremum over 4-dimensional lattices M with $\det M = 1$. Hence $\eta_5 < 7$ by Proposition 10. The only proved facts it uses are:

- $\mu_2 = \gamma_2$, $\mu_3 = \gamma_3$, $\mu_4 = \gamma_4$ (Theorem 35 and its corollary; $\mu_4 = \gamma_4$ rests on $\eta_4 = 3$);
- $\eta_2 = 2/\sqrt{3}$ and $\eta_3 = 3/2$, for the window potentials;
- $\eta_4 = 3$ (Theorem 7), used only as a cap $\eta(M) \le 3$ in the η bound;
- $\gamma_4^2 = 2$.

Coordinates and the root box. Write G_k for the Gram determinant of $b_1, \dots, b_k$ and $r_k = G_{k+1}/G_k$ (indexed from 0). A basis is described by $x = (y_1, y_2, y_3, \mu_{10}, \mu_{20}, \mu_{21}, \mu_{30}, \mu_{31}, \mu_{32})$, where $y_k = \log G_k$ ($y_0 = y_4 = 0$) and the μ_{ij} are the Gram–Schmidt coefficients. So $G(x) = \mu R \mu^T$, with μ unit lower triangular and $R = \mathrm{diag}(r_k)$, $r_k = e^{y_{k+1}-y_k}$, and $\eta(M) = \prod_k G_k = e^{y_1+y_2+y_3}$ when the basis is Ψ-minimal.

Every M has a basis in the root box B below.

1. Ψ-minimal bases exist, by discreteness.
2. Size reduction ($b_i \mapsto b_i - \sum_{j<i} m_j b_j$, $m_j \in \mathbb{Z}$) leaves every G_k unchanged. So a Ψ-minimal basis can be size-reduced, giving $|\mu_{ij}| \le \frac{1}{2}$.
3. Replacing b_i by $-b_i$ changes no G_k, so the signs can be chosen with $\mu_{i,i-1} \ge 0$.
4. The y-ranges:
   - Every 2-window of a Ψ-minimal basis is Ψ-minimal, so $r_{i+1} \ge \frac{3}{4} r_i$. The usual LLL argument then gives $G_k \le (4/3)^{k(4-k)/2}$.
   - For $k = 1$ and $k = 3$, the window bounds give the sharper $G_1 \le \mu_4$ and (dually) $G_3 \le \mu_4$.
5. Only lattices with $\Phi_5(M) \ge 7$ need to be considered. Now $\kappa(M) \le \frac{4}{3}\lambda_1(M)^2 \le \frac{4}{3}\gamma_4$, since $\|\varphi(s)\| \le \frac{1}{2}$; hence $\kappa(M)^2 \le 32/9$. So $\Phi_5(M) \ge 7$ forces $\eta(M) \ge c := \frac{9 \cdot 7}{32}$. With the upper bounds $\bar{y}_j$ above, this gives $y_k \ge \log c - \sum_{j \ne k} \bar{y}_j$.

The program uses $c = 1.96874999 < 63/32$ (outward-rounded), so B is slightly larger than necessary.

The certificate. The program `certify` splits B along y_1 into 480 root slices. Slice 469 is further split into 96 pieces. It then subdivides each slice by bisection and writes the tree to a log. The claim for each leaf is one of two:

(D) the leaf contains no Ψ-minimal basis; or

(A) $\Phi_5(M) < 7$ for every M with a basis in the leaf.

The checker. The program `check_phi` is implemented separately from `certify`; both were written in the same project. They share only the coordinates and the log format. The checker:

- reads each log;
- replays the bisections exactly (the midpoints are recomputed in round-to-nearest, as `certify` computed them);
- re-proves every leaf with its own bounds, subdividing a leaf further where those bounds need it.

This extra subdivision (up to a fixed number of bisections, 24 by default) is part of the trusted checker. It is sound because a claim proved on every part of a subdivision of a leaf holds on the leaf. A separate script checks that the root boxes of all logs tile B exactly, with no gaps or overlaps along y_1 and equal ranges in the other eight coordinates.

Arithmetic. All arithmetic is interval arithmetic in the round-upward mode, with lower ends computed as $-((-a) \circ b)$. The functions exp and log are those of CORE-MATH [42], which are correctly rounded in every rounding mode, widened by one ulp; $\sqrt{}$ is correctly rounded by IEEE 754. The midpoints of bisections are computed in round-to-nearest, as `certify` computed them, by switching the rounding mode around that one operation. The checkers are compiled for baseline x86-64 (SSE2, no fused multiply-add) with `gcc -O2 -frounding-math -ffp-contract=off -static`, and CORE-MATH, which uses fused multiply-adds explicitly, with `-mfma`; so the checkers' own arithmetic contains no fused operations, and only CORE-MATH's explicit fused multiply-adds use the instruction. This applies to every quantity below, including $G(x^*)$, $f(G^*)$ and P.

(D): no Ψ-minimal basis. Each test below is a strict inequality that holds on the whole leaf. So whenever a test discards a leaf, every basis in it is beaten strictly by another basis; two tied Ψ-minimal bases can never eliminate each other.

- Window bounds: a window $b_i, \dots, b_j$ of length L is Ψ-minimal, so $r_i \le \mu_L (G_j/G_{i-1})^{1/L}$ and $r_j \ge (G_j/G_{i-1})^{1/L}/\mu_L$.
- Window potentials: a window of length $L < 4$, normalized to determinant 1, has $\Psi^2 \le \eta_L$.
- The two-dimensional reduction bound $r_{i+1} \ge (1 - \mu_{i+1,i}^2)\, r_i$.
- Deep insertions: some insertion of b_k at position $i < k$ lowers Ψ^2 on the whole leaf. The same test is applied to the reversed dual basis.
- Elementary changes: some basis $b_i \pm b_j$, transposition $b_i \leftrightarrow b_j$, or move of b_k to position i has Ψ^2 smaller than $e^{y_1+y_2+y_3}$ on the whole leaf. The test bounds the difference $\log \Psi_U^2(G(x)) - (y_1 + y_2 + y_3)$ as a single function by the method of the η bound below. Subtracting a linear function keeps it concave, and the mean-value gradient enclosure picks up -1 in each y-component. Near points where the two bases tie, the margin is first order in the distance, and bounding the two sides separately would lose exactly that much.

(A): the η bound. At the center x^* of the leaf the checker searches, in floating point, for an integer flag U that is good at x^*. The search only proposes candidates; every bound below is rigorous for any U. Then $\eta(M) \le \Psi_U^2(G)$.

Let $f(G) = \sum_k w_k \log\det(U_k G U_k^T)$, where U_k is the first k rows of U and $w_k \ge 0$. Two choices are used:

- $w = (1, 1, 1)$ gives $\log \Psi_U^2$;
- $w_k = 2$ for a single k (with the others 0) gives the split bound $\eta \le \eta_k \eta_{4-k} G_k(U)^2$.

Each f is concave in G. So with $G^* = G(x^*)$ and $P = \nabla f(G^*) = \sum_k w_k U_k^T (U_k G^* U_k^T)^{-1} U_k$,

$$f(G(x)) \le f(G^*) + F(x) - F(x^*), \qquad F(x) = \langle P, G(x) \rangle = \operatorname{tr}(P \mu R \mu^T).$$

By the mean value theorem, $F(x) - F(x^*) = \nabla F(\xi) \cdot (x - x^*)$ for some ξ in the leaf, where

$$\frac{\partial F}{\partial \mu_{ab}} = 2(P\mu R)_{ab}, \qquad \frac{\partial F}{\partial y_s} = r_{s-1}(\mu^T P \mu)_{s-1,s-1} - r_s(\mu^T P \mu)_{ss}.$$

The checker encloses ∇F over the whole leaf. It also uses the cruder bound $\sum_{ij} P_{ij}(G_{ij}(x) - G^*_{ij})$, with each entry enclosed separately, and keeps the smaller of the two. The result is E, an upper bound for $\log \eta(M)$ on the leaf, capped at $\log 3$ (Theorem 7).

(A): the κ bound. Put $K = \sqrt{7/e^E}$; it is enough to show $\kappa(M) \le K$ on the leaf.

- Recall $\kappa(M) = \sup_\varphi \inf_{s \neq 0} k_s(\varphi)$, where $k_s(\varphi) = |s|^2/(1 - \|\varphi(s)\|^2)$. Restricting s to a finite set S can only raise the infimum, so it still gives an upper bound. The checker takes $S = \{s \in [-2,2]^4 \cap \mathbb{Z}^4 \setminus \{0\}$, first nonzero coordinate positive$\}$.
- For each $s \in S$, $|s|^2 = \sum_j r_j \big(\sum_{i \ge j} s_i \mu_{ij}\big)^2$ is bounded above on the leaf, say by n_s. Then $k_s(\varphi) \le K$ whenever $\|\varphi(s)\|^2 \le 1 - n_s/K =: \tau_s^2$.
- Shortcut: if some $n_s < \frac{3}{4}K$ (the checker tests $\tau_s > \frac{1}{2}$), then $\|\varphi(s)\| \le \frac{1}{2} < \tau_s$ for every φ, so $\kappa \le \frac{4}{3} n_s < K$.
- Otherwise φ runs over the torus, reduced by $\varphi \mapsto -\varphi$ and integer translation to $\varphi_1 \in [0, \frac{1}{2}]$, $\varphi_2, \varphi_3, \varphi_4 \in [0,1]$. This domain is covered by branch and bound: a φ-box is closed when some s has $\|\varphi(s)\| < \tau_s$ on the whole box. Here $\varphi(s)$ is enclosed by interval arithmetic and its distance to one integer is bounded.

Both tests are strict. In the branch and bound, every φ in the closed torus has some $s \in S$ with $k_s(\varphi) < K$. Since $\varphi \mapsto \min_{s \in S} k_s(\varphi)$ is continuous on the compact torus, its maximum is also $< K$. Hence $\kappa(M) < K$ and $\Phi_5(M) = \kappa^2 \eta < K^2 e^E \le 7$.

Because $E \le \log 3$, $K \ge \sqrt{7/3}$ always. So every leaf on which κ is bounded by $\sqrt{7/3}$ closes without any η bound.

The run.

- *Slices.* The box B was split into 480 root slices along y_1, and slice 469 into 96 further pieces, for 575 logs in all. The certificate tree has 2.39×10^7 leaves.
- *Time.* `certify` used about 180 vCPU-hours on a 32-vCPU cloud instance. `check_phi` used 4.7 vCPU-hours, with at most 24 extra bisections per leaf; the longest single slice took 21 minutes.
- *Result.* Every leaf of every log was re-proved by the final version of `check_phi` (the one described here and archived), and the logs are complete.
- *Coverage.* The root boxes tile a box containing B. This is checked in exact rational arithmetic (`scripts/root_box_check.py`): each double endpoint u of the union of the roots is a rational, the logarithmic endpoints of B have the form $\frac{1}{2}\log Q$ with Q rational, and $u \le \frac{1}{2}\log Q$ is decided by enclosing e^{2u} between rational bounds (a Taylor series with explicit remainder); the μ-endpoints are compared directly. The same check, with $c = 2048/1125$ and $27/16$, is run on the root boxes of Appendices D and E.
- *Software.* C99. The final re-check of all logs, with the build described above, is at the end of Appendix E. Code, scripts and hash manifests: `https://github.com/nickhg/isodual_constant`, which also points to the logs.

D The certificate for $\Phi_5 \le 2^{16}/(3^4 \cdot 5^3)$

The method is that of Appendix C, with the threshold 7 replaced by $P = 2^{16}/(3^4 \cdot 5^3)$, and with the changes below.

Threshold and root box. The η cut becomes $\eta(M) \ge 9P/32 = 2048/1125$; the program uses $c = 1.8204 < 2048/1125$. Leaves of type (A) prove $\Phi_5(M) < P_-$, where P_- is P rounded down to a double; the checker verifies $P_- \le P$ exactly, in integer arithmetic, before it accepts a log.

(B): the balls about M_*. Since $\Phi_5(M_*) = P$, no box containing a Gram matrix of M_* can be of type (A). A third kind of leaf is therefore allowed:

(B) every G in the leaf satisfies $d(G, G_0) \le 1/40$ for one of the Gram matrices G_0 of M_* that lie in the root box (there are 10), so $\Phi_5 \le P$ there, with equality only at G_0, by Proposition 22.

The checker reads each G_0 as $G_0 = UP_0U^\top/\sqrt{135}$, of determinant 1, for the integral matrix P_0 of Definition 16 and a unimodular U. It verifies $\det U = \pm 1$, $\det(UP_0U^\top) = 135^2$ and the adjugate exactly in integers, encloses $\sqrt{135}$ in an interval, and uses $G_0^{-1} = \sqrt{135}\,\mathrm{adj}(UP_0U^\top)/135^2$. On the leaf it encloses $\delta^2 = \mathrm{tr}((G_0^{-1}E)^2)$, $E = G - G_0$, and accepts the leaf if $\delta \le 1/41$. Since $d(G, G_0) = \|\log(I+WEW)\|_F$ with $W = G_0^{-1/2}$ and $\|WEW\|_F = \delta$, this gives $d \le \delta/(1-\delta) \le 1/40$. Leaves of types (A), (B) and (D) together cover B, so $\Phi_5 \le P$ for every M, with equality only at the 10 centers.

What the checker is told. For each leaf of type (A), `certify` records the flag that closed it. The checker tries that flag as well as its own floating-point search; as in Appendix C, a flag is only a candidate, and every bound is proved for whichever flag is used. `certify` can also close a box with a convex combination of the bounds of several flags, which the checker (one flag per box) does not attempt. Where the checker could not follow such a leaf, the piece was certified again with combinations disabled (a setting of the producer: the checker rejecting what it cannot prove is the system working as intended). This affected at least 14 pieces, all near M_*. The checker used at most 24 extra bisections per leaf, raised to 28 or 36 where needed; a single leaf, adjacent to a copy of M_*, needed 48.

Splitting the work. The root box is split along y_1 into 480 root slices, and the 90 hardest slices into 8 pieces each. 450 of these pieces are split further exactly as `certify` itself bisects: at each node the coordinate of largest width (the y widths counted half) is halved at $\frac{1}{2}(lo + hi)$, rounded to nearest. A *manifest* records each split: the piece's root box, the coordinate split at each internal node, and the leaves. 449 pieces were split three levels deep. The hardest piece, adjacent to a copy of M_*, was split adaptively: a sub-box whose certification did not finish within a minute was split three more levels, giving a tree of 505 leaves and depth up to 57. The coverage script checks that each manifest is a full binary tree, replays the bisections to recompute every leaf's box exactly, and requires each leaf to have exactly one log whose root is that box. The root boxes of the slices and of the unsplit pieces must then tile B as in Appendix C, and their union must equal the root box of an unsliced run. Splitting along `certify`'s own bisections repeats no work: for four pieces run both whole and split, the children together processed within 1% of the boxes of the whole run (for example 2.074M against 2.07M). Exact agreement is not expected, since `certify`'s search for flags depends on its internal state, and soundness does not depend on it.

The run.

- *Size.* 4,757 logs (660 slices and unsplit pieces, and 4,097 sub-boxes of the 450 split pieces) with 3.65×10^8 leaves in all, 106,690 of them of type (B).
- *Time.* `certify` used 824 vCPU-hours for the logs in the certificate. The whole run, on cloud instances, took about 4,700 vCPU-hours, including checking, work lost to preempted instances and duplicated runs of the hardest piece.
- *Result.* Every log was first checked on the machine that produced it, and its compressed log and outputs were copied and verified against their checksums before being accepted. All 4,757 logs passed the final re-check (end of Appendix E) with no failures: 4,716 at 24 extra bisections, 34 at 28, 6 at 36 and one at 48.
- *Coverage.* The coverage script accepts the 4,757 root boxes with the 450 manifests, and their union equals the root box, which contains B for $c = 2048/1125$ (checked exactly, as in Appendix C).

E The certificate for $\eta_5 = 6$

This appendix gives the details of the two stages in the proof of Theorem 14.

Stage 1. This is the search of Appendix D with three changes.

- The threshold is 6, which is exact in doubles, and the η cut becomes $\eta(M) \geq 9 \cdot 6/32 = 27/16$; the program uses $c = 1.68749999 < 27/16$.
- Balls of radius $1/20$ about the copies of $\pi_v D_5$ (Proposition 21) were available as leaves. None was used: no box around $\pi_v D_5$ closed in this way before it was collected.
- A box that has not closed when its width falls below $W = 0.12$ is a leaf of type (C), *collected*. It is written to the log exactly, in hexadecimal. The checker (`check_phi`) proves every other leaf as before, and outputs the (C) leaves of each log it accepts; these, and only these, are the input of stage 2. The coverage script counts a (C) leaf as covered, so stage 1 alone proves $\Phi_5 < 6$ outside the union of the (C) leaves.

The root box was split as in Appendix D, and each of the 742 hardest pieces into 32 sub-boxes along the same bisections: $24{,}133 = 389 + 742 \cdot 32$ logs in all, with $1{,}131 = 389 + 742$ roots tiling the root box and the 742 split pieces tiled by their sub-logs. They have 3.25×10^8 leaves, of which 15,419,650 are of type (C). All logs passed the checker, and the coverage script accepts them.

Merging. Two (C) leaves of the same piece that are equal in eight coordinates and abut in the ninth are replaced by their union, which is again a box, and this is repeated to a fixed point. This gives 3,605,989 boxes. Each merged box comes with its merge tree, and a separate script rebuilds every merged box from the original leaves, checking each merge exactly (hexadecimal comparison, no arithmetic) and that every leaf is used exactly once.

Stage 2: coordinates and root. The coordinates are $z = (r, \varphi_1, \ldots, \varphi_4, x)$, where x are the nine coordinates of stage 1 for the basis $b_1, \ldots, b_4$ of M, and $G(z)$ is the Gram matrix in the proof of Theorem 14. The root over a merged box X is $[r_{lo}, r_{hi}] \times [0, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]^3 \times X$, with $r_{lo} \leq 2^{2/5}$ and $r_{hi} \geq 8^{1/5}$ (rounded outward by two ulps; the checker verifies both). The merged boxes are cut into 15,841 jobs of consecutive boxes, and each box is the root of its own log.

Stage 2: leaves. A cell is bisected at the midpoint of one coordinate, rounded to nearest, and its leaves are of the four types of the proof of Theorem 14.

(N) An integer vector $w = (-k, c)$, $c \neq 0$, with $w^\top G w < G_{00} = r$ on the cell.

(F) A unimodular U and one of the bounds, with A_j the leading $j \times j$ block of $U^\top G U$,

$$\log \eta(L) \le \sum_{j=1}^{4} \log\det A_j, \qquad \log \eta(L) \le \log \eta_k + \log \eta_{5-k} + \tfrac{5}{2} \log\det A_k \quad (1 \le k \le 4),$$

the second by Lemma 3 (with $\det G = 1$, $\eta_1 = 1$, $\eta_2 = 2/\sqrt{3}$, $\eta_3 = \tfrac{3}{2}$, $\eta_4 = 3$), which is below $\log 6$ on the cell.

(P) Several such bounds with weights $\lambda_i \ge 0$ whose weighted mean is below $\log 6$ on the cell; since $\log \eta(L)$ is at most each of them, it is at most their mean.

(B) A unimodular W and $t \in \{0, 1\}$ with $d(W G W^\top, T_t) \le 1/10$ on the cell, where T_0 and T_1 are Gram matrices of D_5 and D_5^* of determinant 1.

Stage 2: the checker. `check_lift` is written separately from the producer and shares only the coordinates and the log format. It recomputes $G(z)$ and its derivatives from the formula, in interval arithmetic with outward rounding (the same primitives as `check_phi`), replays the bisections, and checks that each log covers its root. Each $\log\det A_j$ is bounded by the tangent inequality of the concave function $\log\det$, $\log\det A \le -\log\det Y - j + \operatorname{tr}(YA)$, valid for every positive definite Y; the checker takes Y to be the inverse of A_j at the centre of the cell, computed in floating point, and encloses $\log\det Y$ by an interval Cholesky factorization (if it succeeds, the matrix is positive definite and the product of the squared diagonal encloses the determinant [38]). Every bound is then a constant plus a linear function of the entries of G, enclosed over the cell by its mean-value form, intersected with the naive interval; both are standard enclosures of the range of a function over a box [34, Ch. 2]. For (B), with $T_t = LL^\top$, it encloses $\delta = \|L^{-1} W G W^\top L^{-\top} - I\|_F$ and requires $\delta \le 1/11$, so $d \le \delta/(1-\delta) \le 1/10$. A leaf that fails as given may be bisected up to 12 more times by the checker, each part being proved with the same certificate (part of the trusted checker, sound for the same reason as in Appendix C); over the whole run this happened 10,278,717 times. Unit tests include deliberately broken certificates of each type, all of which the checker rejects.

The producer. The producer `stage2` is floating point and proves nothing: its leaves are only candidates for the checker. By default it chooses flags by local descent from the permutations of the basis, and for each cell it tries, in order, (N), (F), (B) and an LP over the flags found so far, which gives (P). Every job was checked on the machine that produced it. A box whose log was rejected was produced again with a larger safety margin (10^{-4}, then 10^{-3}). A box that was still rejected, or that had not finished within 1,800 s, was bisected at the midpoint of its widest coordinate (ties to the lowest index), and each half was proved as a root of its own; the coverage script recomputes the same bisection.

Two incidents during the run concerned only this orchestration, not the checker or the coverage. First, the producer was changed during the run (a splitting hint and a depth guard: the first version could recurse without end near the boundary of the domain). Logs of both versions are accepted, since soundness rests on the checker alone. Second, the bisection fallback originally packed a job whose depth-6 bisection had given up. `check_lift` cannot see a missing log, but the coverage script, which recomputes the bisection tree from the chunk of merged boxes, does, and it found the

one such box. Its remaining parts, and the last two jobs, were proved box by box, with bisection allowed to depth 30 and with the producer's beam search over flags built from short vectors as a retry after a time-out. The cause was a local minimum of the producer's flag descent, at cells where the true η is about 3.9; this was a matter of completeness of the producer's heuristic, and no rejected or unsound leaf was involved.

Coverage. For each of the 742 pieces, the coverage script checks that the chunk files of its jobs concatenate exactly to the merged boxes of the piece, and that every merged box has either an accepted log whose root is exactly that box, or, recursively up to depth 30, accepted logs for the two halves of the bisection above. A log is accepted if its checker output reports the box exactly (hexadecimal) with the producer and checker versions of the run, and the job's checker output ends with no failures. With the coverage of stage 1 and the merge check, this covers every lift in the proof of Theorem 14.

The run.

- *Size.* 15,841 jobs over 3,605,989 merged boxes, of which 533 were bisected by the fallback. The checked leaves are

type	leaves
(N)	1,563,150,108
(F)	1,735,212,582
(P)	5,398,164,149
(B)	47,437
total	8,696,574,276

  The LP combinations (P) do most of the work, and the balls of Theorem 13 are needed only for 47,437 cells.

- *Time.* On cloud instances, stage 1 took about 1,400 vCPU-hours and stage 2 about 5,540, including checking on the producing machines and work lost to preempted instances. A first replay of stage 2 took 258, the box-by-box proofs about 40, and the final re-check below about 620.

- *Result.* Each job's logs were uploaded to cloud storage, with the storage's own SHA-256 checksum compared against the checksum recorded on the producing machine, before the job was counted.

The final re-check. Every log of the four certificates in this paper (Appendices C, D and this one) was re-checked at the end by a single build of `check_phi` and `check_lift`, with exp and log from CORE-MATH as in Appendix C: statically linked binaries built from the stamped version on each machine, with byte-identical results (the same SHA-256 on every machine, recorded with the outputs). For each log, the SHA-256 of the copy that was checked was compared with the one recorded when the log was produced.

- $\Phi_5 < 7$: all 575 logs passed, at 24 extra bisections.

- $\Phi_5 \le 2^{16}/(3^4 \cdot 5^3)$: all 4,757 logs passed, with the distribution above.

- Stage 1 of $\eta_5 = 6$: all 24,133 logs passed at 24 extra bisections, and for each the collected boxes output by the checker are byte-identical to those used as the roots of stage 2.
- Stage 2: all 15,839 archives and the 3 box-by-box proofs passed (of the 15,841 jobs, 15,839 have an archive; the other 2, and the one box missed by the bisection fallback, are the 3 box-by-box proofs), and the coverage script, run on these outputs, accepts all 742 pieces.

An independent check. A second stage-2 checker, `check_lift2`, was written from the log specification and the header comment of `check_lift`. The only thing it takes from `check_lift` is the table of the two integral targets of the (B) leaves, which the specification fixes only up to a choice of basis. It is written in Python with Arb ball arithmetic (including Arb's rigorous exp and log), with a different enclosure method (a mean-value form over each cell) and sharing no code with `check_lift`. It re-proved every (B) leaf of the certificate and a sample of the other leaves fixed in advance by seed: the 304 archives containing a (B) leaf and 300 others drawn at random, and in each archive each non-(B) leaf with probability 0.005. In all it checked 47,437 (B), 508,521 (F), 1,485,204 (P) and 362,365 (N) leaves in 604 archives, with no failures. It chooses how to subdivide a cell heuristically; soundness does not depend on this choice. Separately, the map from search coordinates to Gram matrices was recomputed from first principles in Sage, from explicit bases, and agrees with the checkers' to 10^{-57} at 10^5 random points.

Tests. Besides the broken certificates for `check_lift`, the code release has two more sets of tests. The first is a set of negative controls for `check_phi` at the two maximizers. A box of half-width 10^{-7} about a Gram matrix of M_* (resp. $\pi_v D_5$) is rejected without the balls, with the threshold P (resp. 6) and with a threshold 10^{-6} below it, although the checker may bisect it 24 more times. With a threshold 10^{-3} above, it is accepted, so the rejection comes from the value of Φ_5 there and not from the box. The second is a set of 30 damaged inputs to the coverage and merge scripts (a deleted, duplicated or truncated file or line, an endpoint shifted by one ulp, a wrong header, the output of another job), each of which they reject.

Independence of the checks. The checks are independent at the level of code, but not of authorship. `check_phi` and `check_lift` were written separately from the producers and share with them only the coordinates and the log format. `check_lift2` was written from the log specification and uses a different language, arithmetic library and enclosure method. The coordinate map was recomputed from first principles in Sage. All of these components were written with the same AI tools (see the acknowledgements), and the two C checkers implement the same leaf predicates as the producers. A shared misreading of the mathematical specification, such as the meaning of a leaf type, the root box or the coordinate map, would therefore not be caught by agreement between the programs. Against that risk the paper states each leaf predicate and the root box as mathematics (Appendices C–E), so that they can be checked against the text, and the checkers are short enough to read.

Bisection midpoints. The checkers recompute each midpoint as $m = \frac{1}{2}(\mathrm{lo} + \mathrm{hi})$, with the rounding mode switched to nearest, and give the same double m to both children $[\mathrm{lo}, m]$ and $[m, \mathrm{hi}]$. These cover the parent for any m with $\mathrm{lo} \le m \le \mathrm{hi}$, and that holds for $\frac{1}{2}(\mathrm{lo} + \mathrm{hi})$ in every rounding mode, since rounding is monotone and $2\,\mathrm{lo}$ and $2\,\mathrm{hi}$ are representable. So coverage does not depend on the rounding mode in force at the midpoint: a midpoint computed in another mode would change only the boxes on which the checker re-proves the leaves, not their union. Only the leaf predicates rely on directed rounding.

What has to be trusted. The producers (`certify`, `stage2`) and all orchestration scripts are untrusted: their output is only a candidate. For Theorems 19, 20 and 14, the trusted base consists of

- the two checkers, `check_phi` (696 lines of C) and `check_lift` (557), with CORE-MATH's exp and log, the C compiler and IEEE-754 arithmetic under directed rounding;
- for Theorem 13, which the (B) leaves invoke, SageMath's exact rational arithmetic and its interval arithmetic (`RealIntervalField`);
- the coverage and merge scripts (`cover_check.py`, `root_box_check.py`, `merge_check.py`, `cover_s2.py`; about 240 lines of Python, comparing hexadecimal floats exactly);
- the reductions on paper: the root box of Appendix C, the meaning of each leaf type, and the results the leaves invoke (Propositions 10 and 6, Lemma 3, Theorems 7, 35 and 13, and Proposition 22 for the balls of Theorem 20).

The exact certificates of the other computer-assisted results (Appendices A and B, Propositions 17 and 26) rest on SageMath and Python with exact rational arithmetic, and on SageMath's interval arithmetic where intervals are used. Every leaf claim is a finite inequality between interval expressions, so the checkers could in principle be replaced by a formalization in a proof assistant, which would remove the compiler and the C code from the trusted base.