Constant-Factor Hardness of the Euclidean Lattice Covering Radius

Zhijie Zheng[1] Yucong He[2] Chengzhi Xie[3]

[1]Peking University
[2]Beijing Institute of Technology
[3]National University of Singapore

Abstract

We prove that approximating the Euclidean covering radius of a rational lattice within an absolute constant greater than one is NP-hard. A deterministic polynomial-time many-one reduction distinguishes full-rank rational lattices of covering radius at most a rational threshold r from those of covering radius greater than $(1+10^{-6})r$. This gives a dimension-independent approximation gap, complementing Vallentin's recent exact-hardness result, whose reduction has an inverse-polynomial relative separation. In particular, Euclidean covering radius admits no deterministic polynomial-time approximation scheme unless P = NP.

The hardness persists for sparse bases arbitrarily close to the identity. For every fixed condition-number cap $\rho > 1$, approximation within some constant factor greater than one remains NP-hard for rational symmetric positive-definite bases with condition number at most ρ and at most ten nonzero entries in each row and column. The gap depends on the prescribed condition number or closeness to the identity.

The reduction encodes a parity optimum in the distance from the half-integer center to the lattice. A weighted localization theorem bounds how much the squared distance can increase at any other target. For quadratic forms close to a positive diagonal matrix, the error is quadratic in the relative row-sum perturbation. Combined with rational degree normalization, this transfers a signed parity gap of order tm to the squared radius while keeping the geometric error at $O(t^2m)$. A fixed perturbation therefore preserves a constant gap in every dimension. An explicit parity gadget and rational threshold rounding make the reduction deterministic with polynomial bit complexity.

1 Introduction

The covering radius measures how far a point can be from its nearest lattice point. For an invertible matrix $B \in \mathbb{R}^{n\times n}$ and the lattice $\Lambda = B\mathbb{Z}^n$, it is

$$\mu(\Lambda) = \sup_{y\in\mathbb{R}^n} \min_{z\in\mathbb{Z}^n} \|y - Bz\|_2. \tag{1.1}$$

Equivalently, it is the smallest radius for which Euclidean balls centered at the lattice points cover all of $\mathbb{R}^n$. The outer optimization ranges over target points y, and the inner optimization finds the nearest lattice point by ranging over integer vectors z. Thus the covering radius is the supremal radius of an open Euclidean ball containing no lattice point.

The computational study of this parameter connects lattice algorithms with quantitative bounds in the geometry of numbers [GMR05, Dad19]. A central question is the accuracy achievable in polynomial time. For a fixed relative error, the relevant hardness statement must separate the two cases by a factor bounded away from one as the dimension grows.

Recent results sharpen the distinction between exact computation and approximation. Vallentin [Val26] proved Π_2^{P}-completeness of exact Euclidean covering radius using a relative separation that decreases as an inverse polynomial in the dimension. Constant-factor hardness was established for sufficiently large ℓ_p norms [HR12, BL26]. We prove a constant-factor gap for the Euclidean norm under deterministic polynomial-time many-one reductions from an NP-complete language.

Definition 1.1. For a fixed constant $\gamma > 1$, the promise problem $\mathrm{GapCRP}_2(\gamma)$ has as input an invertible matrix $B \in \mathbb{Q}^{n\times n}$ and a positive rational number r. Its YES instances satisfy $\mu(B\mathbb{Z}^n) \le r$, and its NO instances satisfy $\mu(B\mathbb{Z}^n) > \gamma r$.

All rational numbers are represented in binary. A deterministic many-one reduction to this promise problem sends every YES input of its source language to a YES instance and every NO input to a NO instance.

Theorem 1.2. *The problem* $\mathrm{GapCRP}_2(1 + 10^{-6})$ *is NP-hard under deterministic polynomial-time many-one reductions. The reduction outputs an invertible square rational basis and a positive rational radius.*

The theorem establishes hardness for the fixed factor $1 + 10^{-6}$ and every smaller factor greater than one, independently of the dimension. It therefore excludes a deterministic polynomial-time approximation scheme unless P = NP. Section 6 gives the precise conclusions for one-sided and two-sided numerical estimates. With small-radius instances designated YES, the promise is NP-hard; reversing the YES and NO cases gives a coNP-hard promise.

The identity basis has covering radius $\sqrt{n}/2$, making small perturbations a natural restricted setting for approximation. Our second result shows that constant-factor hardness persists within this setting. For every fixed $\rho > 1$, approximation within some constant factor $\gamma(\rho) > 1$ remains NP-hard for rational symmetric positive-definite bases of Euclidean condition number at most ρ, with at most ten nonzero entries in each row and column. For any fixed $\eta > 0$, the construction can also satisfy $\|B - I\|_2 \le \eta$. The gap depends on the prescribed condition number or distance from the identity. Thus even sparse perturbations of the orthogonal lattice retain approximation hardness for a sufficiently small fixed relative error; see Theorem 7.1.

Preserving a combinatorial gap in the geometry. A reduction can encode a parity optimum in the distance from one selected point to the lattice. The covering radius also maximizes over the

target point, so a deepest hole elsewhere could increase the radius enough to obscure the encoded gap. The geometric task is to bound this increase at a scale that preserves a constant multiplicative separation.

For a diagonal quadratic form, the half-integer point $h = \frac{1}{2}\mathbf{1}$ is a deepest hole. Theorem 2.1 controls the movement of deepest holes under a perturbation in two steps. Local comparisons with adjacent integer points bound every closest-point residual. At a deepest hole, the closest residuals must have both signs in each coordinate, which forces the hole to remain close to h. Averaging the distances to two antipodal integer points then cancels the first-order dependence on the displacement. The remaining error is quadratic in the perturbation size.

To encode a signed parity instance with m edges, we use a rational basis whose Gram matrix has the form

$$Q = D + tA + \frac{t^2}{4} AD^{-1}A,$$

where A records the constraints and D scales each coordinate according to its incidence degree. At h, the term tA records the optimum satisfaction score. The degree normalization ensures both that the total diagonal weight is $O(m)$ and that the relative row perturbation is $O(t)$. The squared covering radius consequently has an $O(t^2 m)$ error around a quantity with a first-order $\Omega(tm)$ satisfaction gap. Choosing a sufficiently small fixed t makes the latter dominate. The resulting gap is independent of dimension.

Relation to earlier work. Guruswami, Micciancio, and Regev [GMR05] established foundational complexity results for covering radius. Their arbitrary-constant-factor hardness concerns linear codes. For lattices, Haviv and Regev [HR12] proved constant-factor Π_2^{P}-hardness in sufficiently large ℓ_p norms, and Bennett and Ly [BL26] obtained NP-hardness for an explicit range of finite p above approximately 35.31. Theorem 1.2 treats the Euclidean norm, $p = 2$.

For an explicit comparison of the approximation gaps, Vallentin's reduction [Val26] uses $R = d/4$ and $\delta = 1/(100d^6)$, with squared-radius thresholds $R - 2\delta$ and $R + 4\delta$ and a rational radius satisfying $R \le r^2 < R + \delta$. These thresholds yield a relative separation $1 + \Theta(d^{-7})$. Our result establishes an absolute constant separation under a reduction proving NP-hardness.

The perturbative approach has a direct antecedent in Vallentin's analysis of deep holes and antipodal candidate nearest points [Val26, Lemmas 3.1–3.2]. We retain that antipodal averaging step. The new quantitative ingredient is a localization estimate measured relative to each row's diagonal scale. It bounds the geometric error in terms of the sum of those scales and permits a perturbation of fixed size in every dimension. Combining this estimate with rational degree normalization converts a constant satisfaction gap into a constant covering-radius gap.

Approximation algorithms and geometric certificates give further context. Guruswami, Micciancio, and Regev [GMR05] give randomized $2^{O(n)}$-time approximations for every fixed factor greater than one. The analysis of Haviv, Lyubashevsky, and Regev [HLR09], together with Magazinov's proof of their Voronoi-cell second-moment conjecture [Mag20], places $\mathrm{GapCRP}_2(\gamma)$ in AM for every fixed $\gamma > \sqrt{3}$. Dadush [Dad19] develops volumetric lower bounds from lattice projections and gives a randomized $2^{O(n)}$-time algorithm that computes such a bound within an $O(\log^{5/2} n)$ factor of the covering radius. These results concern general lattices; our localization estimate gives a quantitative bound for forms close to a positive diagonal matrix.

Organization. Section 2 proves the localization theorem. Section 3 derives a deterministic signed two-variable gap from Håstad's parity theorem [Hås01], including the exact sixteen-edge gadget. Section 4 gives the rational lattice encoding, and Section 5 completes the reduction with exact

radius rounding and bit-complexity bounds. Section 6 gives numerical approximation consequences. Section 7 uses an explicit expander to prove hardness for sparse bases near the identity.

2 A weighted localization estimate

For a symmetric positive-definite matrix $Q \in \mathbb{R}^{n\times n}$, define

$$d_Q(x) = \min_{z\in\mathbb{Z}^n} (x-z)^\top Q(x-z), \qquad \mu_Q^2 = \max_{x\in\mathbb{R}^n} d_Q(x). \tag{2.1}$$

Positive definiteness guarantees that the minimum exists. The function d_Q is continuous and $\mathbb{Z}^n$-periodic, so it attains a maximum on $[0,1]^n$. To see continuity, restrict the minimum locally to finitely many integer points: positive definiteness bounds all possible minimizers in a bounded Euclidean ball. When $Q = B^\top B$ and B is invertible, the change of variables $y = Bx$ gives

$$\mu(B\mathbb{Z}^n)^2 = \mu_Q^2. \tag{2.2}$$

Write $\|M\|_\infty = \max_i \sum_j |M_{ij}|$ for the induced maximum row-sum norm, and set

$$h = \frac{1}{2}\mathbf{1}, \qquad a_Q = \frac{1}{4}\min_{s\in\{-1,1\}^n} s^\top Q s. \tag{2.3}$$

The following theorem identifies the center value and bounds how much the value at a deepest hole can exceed it.

Theorem 2.1 (Weighted localization)**.** *Let $D = \operatorname{diag}(D_1,\dots,D_n)$ with $D_i > 0$, and let Q be symmetric positive definite. Suppose*

$$\kappa = \|D^{-1}Q - I\|_\infty < \frac{1}{3}, \qquad S = \sum_{i=1}^n D_i.$$

Then $d_Q(h) = a_Q$. Every maximizer of d_Q, after translation into $[0,1)^n$, satisfies

$$\left|x_i - \frac{1}{2}\right| \le \frac{\kappa}{1-\kappa} \qquad (1 \le i \le n), \tag{2.4}$$

and

$$0 \le \mu_Q^2 - a_Q \le S(1+\kappa)\left(\frac{\kappa}{1-\kappa}\right)^2. \tag{2.5}$$

The hypothesis compares each row with its own diagonal scale. Although $D^{-1}Q$ may be asymmetric, its induced row-sum norm controls closest-point residuals. Symmetry of Q enters later, when we compare the quadratic forms Q and D.

Lemma 2.2 (Closest residuals)**.** *Under the hypotheses of Theorem 2.1, let z be any closest integer point to x in the Q metric, and write $u = x - z$. Then*

$$\|u\|_\infty \le M_\kappa := \frac{1+\kappa}{2(1-\kappa)} < 1. \tag{2.6}$$

Proof. Comparing z with $z + e_i$ and $z - e_i$ gives

$$u^\top Q u \le (u \mp e_i)^\top Q(u \mp e_i) = u^\top Q u \mp 2(Qu)_i + Q_{ii}.$$

Thus $|(Qu)_i| \le Q_{ii}/2$. Put $H = D^{-1}Q = I + E$. Dividing the ith inequality by D_i and using $Q_{ii}/D_i \le 1+\kappa$, we obtain $\|Hu\|_\infty \le (1+\kappa)/2$. The Neumann series $H^{-1} = \sum_{k\ge 0}(-E)^k$ converges in the induced row-sum norm and satisfies $\|H^{-1}\|_\infty \le (1-\kappa)^{-1}$. This proves the first inequality in (2.6). The second follows from $\kappa < 1/3$. □

Lemma 2.3 (Signs of closest residuals at a deepest hole)**.** *Let Q be symmetric positive definite, and let x maximize d_Q. For every nonzero $v \in \mathbb{R}^n$, there is a closest integer point z with $v^\top Q(x-z) < 0$. In particular, for each coordinate i, some closest residual has positive ith coordinate and some closest residual has negative ith coordinate.*

Proof. Suppose that $v^\top Q(x-z) \geq 0$ for every closest point z. We first reduce to finitely many candidate integer points near x. If $\lambda > 0$ is the smallest eigenvalue of Q, then

$$(y-z)^\top Q(y-z) \geq \lambda \|y-z\|_2^2.$$

For y in a sufficiently small closed ball around x, all sufficiently distant integer points therefore have squared distance greater than $d_Q(x)+1$. Among the remaining finite set, every point that is not closest at x has a positive gap above $d_Q(x)$. By continuity, all these distances remain greater than $d_Q(x)$ in a small enough neighborhood.

For each closest point z and each $\varepsilon > 0$,

$$(x+\varepsilon v - z)^\top Q(x+\varepsilon v - z) = d_Q(x) + 2\varepsilon v^\top Q(x-z) + \varepsilon^2 v^\top Q v > d_Q(x).$$

Positive definiteness makes the increase strict even if the linear term vanishes. Taking ε small enough to remain in this neighborhood gives $d_Q(x+\varepsilon v) > d_Q(x)$, contradicting maximality. Finally, applying the result with $v = Q^{-1}e_i$ and with $v = -Q^{-1}e_i$ gives the two coordinate signs. □

Proof of Theorem 2.1. The argument has three parts. Bounds on closest residuals locate every deepest hole near h. Half-integrality then identifies the value at h, and averaging two antipodal lattice points bounds the change in value by a quadratic displacement term.

We first localize a deepest hole. Translate a maximizing x into $[0,1)^n$. Lemma 2.2 bounds every coordinate of every closest residual in absolute value by $M_\kappa < 1$. By Lemma 2.3, both signs occur in each coordinate. Consequently, a positive residual in coordinate i must equal $x_i - 0$, and a negative residual must equal $x_i - 1$. The existence of a negative residual also excludes $x_i = 0$. Hence

$$1 - M_\kappa \leq x_i \leq M_\kappa,$$

which is (2.4). The closest points that give the two signs may differ from one coordinate to another.

Next consider the center h. Its residuals have half-integer coordinates. Lemma 2.2 forces every closest residual to belong to $\{-1/2, 1/2\}^n$, and every vector in this set is a residual from an integer point. Thus $d_Q(h) = a_Q$.

To bound the effect of moving away from h, choose $s \in \{-1,1\}^n$ minimizing $s^\top Q s$ and write $x = h + \delta$. Both $h + s/2$ and $h - s/2$ are integral. The smaller of their squared distances from x is at most their average, so

$$\begin{aligned} d_Q(x) &\leq \frac{1}{2}(\delta - s/2)^\top Q(\delta - s/2) + \frac{1}{2}(\delta + s/2)^\top Q(\delta + s/2) \\ &= a_Q + \delta^\top Q \delta. \end{aligned} \tag{2.7}$$

The averaging cancels the linear dependence on δ.

Finally put $F = Q - D$. The row hypothesis gives $\sum_j |F_{ij}| \leq \kappa D_i$. Since F is symmetric, every $w \in \mathbb{R}^n$ satisfies

$$|w^\top F w| \leq \sum_{i,j} |F_{ij}| \frac{w_i^2 + w_j^2}{2} = \sum_i w_i^2 \sum_j |F_{ij}| \leq \kappa \sum_i D_i w_i^2. \tag{2.8}$$

Thus $\delta^\top Q\delta \le (1+\kappa)\sum_i D_i\delta_i^2$. At a maximizer, (2.4) gives $|\delta_i| \le \kappa/(1-\kappa)$ for every coordinate. Consequently,

$$\delta^\top Q\delta \le (1+\kappa)\sum_i D_i\left(\frac{\kappa}{1-\kappa}\right)^2 = S(1+\kappa)\left(\frac{\kappa}{1-\kappa}\right)^2.$$

Substituting into (2.7) proves the upper bound in (2.5). The displacement is of order κ, and antipodal averaging makes its contribution to the value quadratic in κ. Evaluating d_Q at h gives the lower bound. □

3 A deterministic signed parity gap

A signed two-variable instance is a loopless undirected multigraph in which each edge object prefers either equality or inequality of its endpoint signs. Parallel edge objects, including ones with opposite preferences, are allowed. For an assignment $s \in \{-1,1\}^n$, let $\mathrm{Score}(s)$ be the number of satisfied edge objects and let $C = \max_s \mathrm{Score}(s)$. All edges in this section have unit weight.

Proposition 3.1. *There is a deterministic polynomial-time many-one reduction from an NP-complete language to signed two-variable instances with $m > 0$ edges such that*

$$\textit{YES: } C \ge \alpha m, \qquad \textit{NO: } C \le \beta m, \qquad \alpha = \frac{95}{128}, \quad \beta = \frac{89}{128}. \tag{3.1}$$

The instances have no loops. Discarding isolated vertices leaves at most $2m$ vertices.

We derive the gap in two explicit steps: a rounding procedure converts a weighted parity instance into an unweighted one, and a fixed gadget replaces each three-variable equation by pairwise constraints. The complexity-theoretic input is the following consequence of Håstad's PCP theorem.

Theorem 3.2 (Håstad's weighted parity gap)**.** *There is a deterministic polynomial-time reduction from an NP-complete language to a list of $K \ge 1$ equations $x_{a_i}x_{b_i}x_{d_i} = c_i$ over signs, with $c_i \in \{-1,1\}$ and nonnegative rational weights w_i summing to one, such that the maximum satisfied total weight is at least $31/32$ on YES inputs and at most $33/64$ on NO inputs. The list and its rational weights have polynomial encoding length.*

This follows from the weighted construction in the proof of Håstad [Hås01, Theorem 5.4], with $\delta = 1/32$. The construction, on printed pages 28–29 of the author version, enumerates the polynomially many verifier choices and gives each resulting equation its rational probability. This enumeration makes the reduction deterministic. Repeated variables in an equation are allowed by the constructions below.

Lemma 3.3 (Deterministic unweighting)**.** *From the instance in Theorem 3.2, one can construct in deterministic polynomial time $M \ge 1$ unweighted three-variable sign equations whose maximum number $S_{\max}$ of simultaneously satisfied equations obeys*

$$\textit{YES: } S_{\max} \ge \frac{15}{16}M, \qquad \textit{NO: } S_{\max} \le \frac{9}{16}M. \tag{3.2}$$

Proof. Let $L = 128K$ and replace equation i by $\lfloor Lw_i \rfloor$ copies. The number $M = \sum_i \lfloor Lw_i \rfloor$ satisfies $127K \le M \le 128K$. If an assignment has weighted value v and satisfies S of the new equations, then

$$Lv - K \le S \le Lv.$$

On a YES input, take an assignment with $v \geq 31/32$. Since $M \leq L$ and $Lv - K \geq 0$,

$$\frac{S}{M} \geq \frac{Lv - K}{L} \geq \frac{31}{32} - \frac{1}{128} = \frac{123}{128} > \frac{15}{16}.$$

On a NO input, every assignment satisfies

$$\frac{S}{M} \leq \frac{L(33/64)}{127K} = \frac{66}{127} < \frac{9}{16}.$$

Exact rational arithmetic computes each multiplicity, and the total output size is $O(K)$. Thus the number of copies is polynomially bounded even when a common denominator is numerically large. □

Lemma 3.4 (The sixteen-edge cube gadget)**.** *An equation $x_1x_2x_3 = c$ can be replaced by sixteen signed two-variable constraints with four private auxiliary variables and a shared reference sign z. With $z = 1$ and the original variables fixed, the maximum number of satisfied gadget constraints is twelve if the equation is satisfied and ten otherwise.*

Proof. Label the vertices of a cube by $\{0,1\}^3$. Its four even-parity vertices are terminals with signs

$$q_{000} = z, \qquad q_{011} = cx_1, \qquad q_{101} = x_2, \qquad q_{110} = x_3.$$

The four odd-parity vertices are private auxiliary variables. Place an inequality constraint on each of the twelve cube edges and an equality constraint on each of the four main diagonals. The terminal sign cx_1 is represented by a literal: when $c = -1$, switch the preferences of its four incident edges in this gadget. This switching is local to that occurrence of x_1.

Each auxiliary vertex is adjacent to all four terminals. It prefers equality with the opposite terminal and inequality with the other three. Denote the terminal signs by $q_1, \ldots, q_4$, and set $T = \sum_i q_i$. If the auxiliary opposite terminal j receives sign y, its four satisfaction indicators sum to

$$\frac{1 + yq_j}{2} + \sum_{i \neq j} \frac{1 - yq_i}{2} = 2 + \frac{y}{2}(2q_j - T).$$

Choosing the better of $y = 1$ and $y = -1$ therefore satisfies

$$2 + \frac{1}{2}\left| q_j - \sum_{i \neq j} q_i \right| = 2 + \frac{1}{2}|2q_j - T|$$

constraints. The four auxiliary choices are independent, so the total optimum is

$$8 + \frac{1}{2}\sum_{j=1}^{4} |2q_j - T|. \tag{3.3}$$

If $\prod_j q_j = 1$, all terminal signs agree or two have each sign. In either case the absolute-value sum is eight and the optimum is twelve. If $\prod_j q_j = -1$, one or three signs are negative, the sum is four, and the optimum is ten. For $z = 1$, the product is $cx_1x_2x_3$, so the first case is exactly satisfaction of the source equation.

The identity holds for every terminal assignment, including those obtained from repeated original variables. Every edge joins an even-parity terminal to a fresh odd-parity auxiliary vertex. Identifying repeated terminals therefore creates only parallel edges and preserves looplessness. □

This is the cube construction in Håstad [Hås01, Theorem 5.14], attributed there to Sorkin. The calculation above fixes the normalization and exact scores used in the reduction.

Proof of Proposition 3.1. Apply Lemma 3.3 and replace each of the M equations by the gadget in Lemma 3.4. Use a single shared new variable z and otherwise disjoint auxiliary sets. Since all constraints are pairwise, negating all signs preserves every score, so an assignment of maximum score may be chosen with $z = 1$. The private auxiliaries can then be optimized independently, giving

$$m = 16M, \qquad C = 10M + 2S_{\max}.$$

Substituting (3.2) proves (3.1), because

$$\frac{10 + 2(15/16)}{16} = \frac{95}{128}, \qquad \frac{10 + 2(9/16)}{16} = \frac{89}{128}.$$

The constructions are deterministic and polynomial in size. Lemma 3.4 gives looplessness. Finally, the sum of the nonzero incidence degrees is $2m$, so at most $2m$ vertices remain after isolated vertices are discarded. □

4 Encoding the optimum in a rational lattice

Let a loopless signed instance have n nonisolated vertices and $m > 0$ unit edge objects. Define its symmetric signed adjacency matrix by

$$A_{ij} = \#\{\text{inequality edges } ij\} - \#\{\text{equality edges } ij\}, \qquad A_{ii} = 0.$$

Let d_i be the incidence degree counting all edge objects, before cancellation in A. Then

$$\sum_j |A_{ij}| \le d_i, \qquad \sum_i d_i = 2m, \qquad s^\top A s = 2m - 4\,\mathrm{Score}(s). \tag{4.1}$$

Indeed, an inequality edge contributes $2s_is_j$ to the quadratic form and has satisfaction indicator $(1 - s_is_j)/2$; an equality edge contributes $-2s_is_j$ and has indicator $(1 + s_is_j)/2$. Summing gives the identity, including when opposite-signed parallel edges cancel in A.

A diagonal matrix of rational scales accounts for the variation in vertex degrees. For each i, let $T_i = 2^{k_i}$, where k_i is the smallest nonnegative integer satisfying $T_i^2 \ge d_i$, and put

$$T = \mathrm{diag}(T_1, \dots, T_n), \qquad D = T^2, \qquad S = \sum_i T_i^2. \tag{4.2}$$

Since $d_i \ge 1$, minimality gives

$$d_i \le T_i^2 < 4d_i, \qquad 2m \le S < 8m. \tag{4.3}$$

These powers can be found by integer comparisons.

Proposition 4.1 (Quantitative rational encoding)**.** *For a fixed rational $0 < t \le 1/16$, set*

$$B = T + \frac{t}{2} T^{-1} A. \tag{4.4}$$

The matrix B is rational and invertible. If $C = \max_s \mathrm{Score}(s)$, then

$$\frac{S}{4} + t\left(\frac{m}{2} - C\right) \le \mu(B\mathbb{Z}^n)^2 \le \frac{S}{4} + t\left(\frac{m}{2} - C\right) + 14t^2 m. \tag{4.5}$$

Every entry of B has $O(\log(m+1))$ bits for fixed t, and the basis is computable in deterministic polynomial time.

Proof. We verify invertibility, compute the center value, and then apply the localization theorem.

First, symmetry and (4.1) give, for every real v,

$$|v^\top Av| \le \sum_{i,j} |A_{ij}| \frac{v_i^2 + v_j^2}{2} \le \sum_i d_i v_i^2 \le v^\top Dv.$$

Hence $D + (t/2)A$ is positive definite for $t < 2$. The factorization $B = T^{-1}[D + (t/2)A]$ proves invertibility. Its Gram matrix is exactly

$$Q = B^\top B = D + tA + \frac{t^2}{4} AD^{-1}A. \tag{4.6}$$

Let $K = D^{-1}A$. By (4.3), $\|K\|_\infty \le 1$, and therefore

$$D^{-1}Q = I + tK + \frac{t^2}{4}K^2, \qquad \kappa := \|D^{-1}Q - I\|_\infty \le t + \frac{t^2}{4} \le \frac{17}{16}t \le \frac{1}{8}. \tag{4.7}$$

Theorem 2.1 thus applies directly to the Gram matrix of the rational basis (4.4).

For a sign vector s, the diagonal term is $s^\top Ds = S$, and (4.1) determines the linear term. The quadratic perturbation is nonnegative and satisfies

$$s^\top AD^{-1}As = \sum_i \frac{(As)_i^2}{T_i^2} \le \sum_i \frac{d_i^2}{T_i^2} \le \sum_i d_i = 2m. \tag{4.8}$$

In the center value $a_Q = \frac{1}{4}\min_s s^\top Qs$, the last term of (4.6) has coefficient $t^2/16$. Equation (4.8) therefore bounds its contribution between zero and $t^2 m/8$. Bounding the value from below for every sign vector and from above at an assignment of score C gives

$$\frac{S}{4} + t\left(\frac{m}{2} - C\right) \le a_Q \le \frac{S}{4} + t\left(\frac{m}{2} - C\right) + \frac{t^2 m}{8}. \tag{4.9}$$

Displacement of a deepest hole from the center contributes the additional error in (2.5). By (4.3) and (4.7), it is at most

$$S(1+\kappa)\left(\frac{\kappa}{1-\kappa}\right)^2 \le 8m \cdot \frac{9}{8}\left(\frac{17}{16}t \cdot \frac{8}{7}\right)^2 = \frac{2601}{196}t^2 m.$$

Since $2601/196 + 1/8 = 5251/392 < 14$, equations (2.2), (2.5), and (4.9) prove (4.5).

Finally, $|A_{ij}| \le m$ and $T_i^2 < 4d_i \le 4m$, so every T_i and every entry of A has $O(\log(m+1))$ bits. For fixed rational t, the entries in (4.4) have the same asymptotic bit length. Computing the degrees, powers of two, and matrix entries requires polynomially many operations on integers of polynomial bit length. □

5 The constant gap and exact radius rounding

The encoding separates the squared radii by a constant multiple of m. We now convert that separation into the promised multiplicative gap in the radius, using a rational threshold rounded upward so that the YES condition is preserved.

Proof of Theorem 1.2. Take the instance from Proposition 3.1, discard isolated vertices, and apply Proposition 4.1 with

$$t = \frac{1}{1000}, \qquad \gamma = 1 + \frac{1}{10^6}. \tag{5.1}$$

These constants are independent of the input and dimension. Define rational squared-radius thresholds

$$U = \frac{S}{4} + t\left(\frac{1}{2} - \alpha\right)m + 14t^2 m, \qquad L = \frac{S}{4} + t\left(\frac{1}{2} - \beta\right)m. \tag{5.2}$$

A YES input has $C \geq \alpha m$, and (4.5) gives $\mu(B\mathbb{Z}^n)^2 \leq U$. A NO input has $C \leq \beta m$, and the same estimate gives $\mu(B\mathbb{Z}^n)^2 \geq L$. The exact separation is

$$L - U = \left(\frac{3}{64000} - \frac{14}{10^6}\right)m = \frac{263}{8\,000\,000}m. \tag{5.3}$$

Also, using (4.3),

$$\frac{m}{4} < U < 2m - \left(\frac{31}{128000} - \frac{14}{10^6}\right)m < 2m - 10^{-4}m. \tag{5.4}$$

For the lower bound, $S/4 \geq m/2$ and $31/128000 < 1/4$ suffice.

We will choose a rational $r \geq \sqrt{U}$ so that $\gamma^2 r^2 < L$. The separation $L - U$ must cover two increases: $r^2 - U$ from rounding the threshold and $(\gamma^2 - 1)r^2$ from imposing the multiplicative gap. For an exact upward rounding, let

$$H = 10^{12}m, \qquad k = \min\{j \in \mathbb{Z}_{\geq 0} : j^2 \geq H^2 U\}, \qquad r = k/H. \tag{5.5}$$

Since $U > 0$, the radius r is positive. Minimality of k implies $\sqrt{U} \leq r < \sqrt{U} + 1/H$, and hence

$$0 \leq r^2 - U < \frac{2\sqrt{2m}}{H} + \frac{1}{H^2} < \frac{m}{10^8}. \tag{5.6}$$

The last inequality follows from $m \geq 1$ and $\sqrt{2m} \leq 2m$. Together with (5.4), it also gives $r^2 < 2m$. Thus

$$\begin{aligned} \gamma^2 r^2 - U &= (r^2 - U) + (\gamma^2 - 1)r^2 \\ &< [10^{-8} + 2(2 \cdot 10^{-6} + 10^{-12})]m \\ &= \frac{2\,005\,001}{500\,000\,000\,000}m < \frac{263}{8\,000\,000}m = L - U. \end{aligned} \tag{5.7}$$

Consequently $\gamma^2 r^2 < L$. The YES case satisfies $\mu(B\mathbb{Z}^n) \leq r$, whereas the NO case satisfies $\mu(B\mathbb{Z}^n) \geq \sqrt{L} > \gamma r$.

The reduction and its output have polynomial bit complexity. The source reduction, unweighting, and cube construction are deterministic and have polynomial size. The dimension is $n \leq 2m$. Proposition 4.1 gives basis entries with $O(\log(m+1))$ bits. Because t, α, β are fixed rationals, S, U, H also have $O(\log(m+1))$ bits. Writing $U = p/q$ with positive integers, the comparison defining k is the integer comparison $qj^2 \geq H^2 p$. Binary search on $[0, 2Hm + 1]$ finds the least such j; its upper endpoint exceeds $H\sqrt{U}$ because $U < 2m$ and $m \geq 1$. There are $O(\log(m+1))$ comparisons on integers of $O(\log(m+1))$ bits. Thus k and r have $O(\log(m+1))$ bits, and exact integer arithmetic computes the entire threshold. A dense output basis uses $O(n^2 \log(m+1))$ bits. This proves that the mapping is a deterministic polynomial-time many-one reduction. □

6 Consequences for numerical approximation

The decision gap yields the following bounds for numerical approximation under one-sided and two-sided output guarantees.

Corollary 6.1. *Unless* P = NP*, no deterministic polynomial-time algorithm, on every rational full-rank basis B, outputs a rational number $\widehat{\mu}$ satisfying*

$$\mu(B\mathbb{Z}^n) \le \widehat{\mu} \le (1 + 10^{-6})\mu(B\mathbb{Z}^n).$$

There is also no deterministic polynomial-time algorithm guaranteeing $\mu(B\mathbb{Z}^n)/a \le \widehat{\mu} \le a\mu(B\mathbb{Z}^n)$ for any fixed rational $a > 1$ with $a^2 \le 1 + 10^{-6}$. In particular, either convention admits no polynomial-time approximation scheme unless P = NP*.*

Proof. Set $\gamma = 1 + 10^{-6}$. In the first convention, a YES instance of Theorem 1.2 has $\widehat{\mu} \le \gamma r$, while a NO instance has $\widehat{\mu} \ge \mu(B\mathbb{Z}^n) > \gamma r$. Comparison with the rational threshold γr therefore decides the promise problem.

In the second convention, a YES instance has $\widehat{\mu} \le ar$, while a NO instance has $\widehat{\mu} \ge \mu(B\mathbb{Z}^n)/a > \gamma r/a \ge ar$. Comparison with ar again decides the problem. A polynomial-time approximation scheme at a fixed small accuracy, for example 10^{-7}, would give one of these guarantees in polynomial time. □

Remark 6.2. For a one-sided lower estimate, the guarantee $\mu/(1 + 10^{-6}) \le \widehat{\mu} \le \mu$ is excluded by comparing $\widehat{\mu}$ with r. The complexity assumption P ≠ NP gives these exclusions for deterministic polynomial-time algorithms.

7 Hardness near the identity

The normalization in Section 4 accommodates arbitrary degrees through varying diagonal scales. To obtain well-conditioned bases, we first convert the signed instance to bounded weighted degree, then apply Theorem 2.1 with $D = I$. Write $\mathrm{cond}_2(B) = \sigma_{\max}(B)/\sigma_{\min}(B)$ for the Euclidean condition number.

Theorem 7.1. *For every fixed real $\rho > 1$, there is a fixed rational $\gamma(\rho) > 1$ such that* $\mathrm{GapCRP}_2(\gamma(\rho))$ *is NP-hard under deterministic polynomial-time many-one reductions, even for rational symmetric positive-definite bases B with $\mathrm{cond}_2(B) \le \rho$. Every row and column has at most ten nonzero entries, and, for fixed ρ, all entries belong to a fixed finite set of rationals.*

More explicitly, any fixed positive rational

$$t \le \min\left\{10^{-7}, \frac{2(\rho - 1)}{25(\rho + 1)}\right\} \qquad \textit{permits} \qquad \gamma(\rho) = 1 + \frac{t}{1000}. \tag{7.1}$$

The same construction satisfies $\|B - I\|_2 \le \|B - I\|_\infty \le 25t/2$.

In particular, $t = 10^{-7}$ gives gap $1 + 10^{-10}$ with condition number at most $(1 + 1.25 \cdot 10^{-6})/(1 - 1.25 \cdot 10^{-6})$. For any fixed $\eta > 0$, choosing also $t \le 2\eta/25$ bounds both displayed norms by η. If ρ or η is irrational, choose a smaller fixed rational t; the reduction then uses only that rational constant.

7.1 An explicit expander and the cloud reduction

The degree reduction replaces each vertex by a cloud of copies connected by an expander. Internal equality constraints ensure that assigning the majority sign to an entire cloud preserves or increases the total score. The expansion constant below is sufficient for an exact relation between the original and converted optima.

Lemma 7.2 (Expander input)**.** *For an integer $q \geq 2$, take vertex set $\mathbb{Z}_q^2$ and the four permutations, with arithmetic modulo q,*

$$F_1(x, y) = (x + 2y, y), \quad F_2(x, y) = (x + 2y + 1, y),$$
$$F_3(x, y) = (x, y + 2x), \quad F_4(x, y) = (x, y + 2x + 1).$$

For every vertex v and every $j \in \{1, 2, 3, 4\}$, create one undirected edge object $\{v, F_j(v)\}$. Retain coincident endpoint pairs and reverse-direction duplicates as separate objects, and count each loop twice toward degree. The resulting multigraph is eight-regular. Its adjacency matrix P, with each loop counted twice on the diagonal, satisfies

$$f^\top P f \leq \frac{73}{10}\|f\|_2^2 \qquad \text{whenever } f^\top \mathbf{1} = 0. \tag{7.2}$$

This is Construction 8.1 of Hoory, Linial, and Wigderson [HLW06]. Their Fourier proof in Section 8, printed pages 503–508, supplies the estimate used here: the coefficient 3.65 on page 507 bounds the four-forward-map quadratic form in their equation (12), on page 504. The adjacency quadratic form is twice that form, giving 73/10. Their Theorem 8.2 gives the stronger bound $5\sqrt{2}$. We use (7.2) to derive the boundary bound needed for the degree reduction.

Lemma 7.3 (Minority boundary)**.** *In the multigraph of Lemma 7.2, every set U of $k \leq q^2/2$ vertices has at least $7k/20$ crossing edge objects. The same holds after deleting all loops. For $q = 1$, the statement holds as well, with the empty set as its only admissible set.*

Proof. For $q \geq 2$, put $f = \mathbf{1}_U - (k/q^2)\mathbf{1}$. A crossing undirected edge contributes one to the Laplacian energy, and a loop contributes zero. Thus

$$|\partial U| = f^\top (8I - P) f \geq \left(8 - \frac{73}{10}\right)\|f\|_2^2 = \frac{7}{10}k\left(1 - \frac{k}{q^2}\right) \geq \frac{7}{20}k.$$

Deleting loops preserves every boundary. The assertion for $q = 1$ follows from $U = \varnothing$. □

Lemma 7.4 (Exact bounded-degree conversion)**.** *From a signed instance with $m > 0$ unit edge objects, no loops, and no isolated vertices, one can construct deterministically a signed instance with positive integer edge weights, $N \leq 8m$ vertices, and incident total weight at most 25 at each vertex. If e is the number of its internal cloud edges, counted with multiplicity, its total weight is $W = m + 3e$ and its maximum satisfied weight is exactly*

$$C' = 3e + C. \tag{7.3}$$

Each vertex has at most nine distinct neighbors, and the construction has polynomial size.

Proof. Replace a vertex v of incidence degree d_v by a cloud of $N_v = q_v^2$ vertices, where q_v is the least positive integer satisfying $q_v^2 \geq d_v$. Designate d_v distinct vertices as ports, one for each original incidence; the remaining vertices are dummies. Since $q_v < \sqrt{d_v} + 1 \leq 2\sqrt{d_v}$,

$$d_v \leq N_v \leq 4d_v, \qquad N = \sum_v N_v \leq 8m.$$

Within each cloud, insert the multigraph of Lemma 7.2, delete its loops, and give every remaining edge an equality preference of weight three. A one-vertex cloud has no internal edges. Replace each original edge by a unit-weight edge with the same preference between its two designated ports. Each port has one external edge, and each vertex has internal incidence degree at most eight. Thus the incident total weight is at most $1 + 3 \cdot 8 = 25$, with at most nine distinct neighbors.

Consider any sign assignment. Choose a majority sign in each cloud, breaking ties arbitrarily, and change all minority signs to it. Let U_v be the set of changed vertices in cloud v and put $k_v = |U_v| \le N_v/2$. An internal equality edge changes from unsatisfied to satisfied exactly when it crosses U_v; all other internal edges keep their status. The internal gain is therefore $3 \sum_v |\partial U_v| \ge (21/20) \sum_v k_v$ by Lemma 7.3.

For the external edges, every lost unit of score has a changed endpoint. Each port belongs to one external edge, so the number of such edges is at most the number of changed ports, which is at most $\sum_v k_v$. An edge whose two endpoints change keeps its status and may still be counted twice in this upper bound. Combining the internal gain and the external loss gives

$$\text{Score(after)} - \text{Score(before)} \ge \frac{21}{20} \sum_v k_v - \sum_v k_v = \frac{1}{20} \sum_v k_v \ge 0.$$

Thus making every cloud constant never decreases the total score.

An assignment that is constant on each cloud satisfies all internal constraints, of total weight $3e$. Its external score is precisely the score of the corresponding original assignment. The maximum over these assignments is consequently $3e + C$. The rounding argument proves that this is also the unrestricted optimum, establishing (7.3). The explicit maps, integer square roots, and port choices are deterministic. Before loop deletion there are at most four internal edge objects per vertex, so the construction uses $O(m)$ vertices and edges. □

Applying Lemma 7.4 to Proposition 3.1, define the computable thresholds

$$a = 3e + \alpha m, \qquad b = 3e + \beta m. \tag{7.4}$$

They satisfy $0 \le b < a \le W$. The resulting promise has $C' \ge a$ on YES inputs and $C' \le b$ on NO inputs, and

$$a - b = \frac{3}{64} m \ge \frac{3}{512} N. \tag{7.5}$$

Both thresholds include the same internal equality contribution $3e$.

7.2 The near-identity basis and its error

Let A be the signed weighted adjacency matrix of the cloud instance, with positive entries for inequality and negative entries for equality, and with zero diagonal. Aggregating parallel edges can only decrease the absolute row sum. Thus

$$\|A\|_\infty \le \Delta := 25, \qquad s^\top A s = 2W - 4\,\text{Score}(s), \qquad W \le \frac{\Delta N}{2}. \tag{7.6}$$

Every eigenvalue of the symmetric matrix A has absolute value at most Δ: apply the eigenvector equation at a coordinate of maximum absolute value. Hence $\|A\|_2 \le \Delta$ as well.

For fixed rational t as in (7.1), put

$$B = I + \frac{t}{2} A, \qquad Q = B^\top B = I + tA + \frac{t^2}{4} A^2. \tag{7.7}$$

The eigenvalues of B lie in $[1-\Delta t/2, 1+\Delta t/2]$, so B is symmetric positive definite and

$$\operatorname{cond}_2(B) \le \frac{1+\Delta t/2}{1-\Delta t/2} \le \rho, \qquad \|B-I\|_2 \le \|B-I\|_\infty \le \Delta t/2. \tag{7.8}$$

Each row has at most nine off-diagonal nonzero entries, and symmetry gives the same column bound. All entries of A are integers of absolute value at most 25. Thus the basis has the sparsity bound and the fixed finite set of rational entries stated in Theorem 7.1.

For the radius estimate, $\Delta t \le 3/16$ implies

$$\kappa = \|Q-I\|_\infty \le \Delta t + \frac{(\Delta t)^2}{4} \le \frac{4}{3}\Delta t \le \frac{1}{4}.$$

Theorem 2.1 with $D = I$ bounds the geometric error by

$$N(1+\kappa)\left(\frac{\kappa}{1-\kappa}\right)^2 \le N \cdot \frac{5}{4}\left(\frac{4}{3}\Delta t \cdot \frac{4}{3}\right)^2 = \frac{320}{81}\Delta^2 t^2 N.$$

For every sign vector, $0 \le s^\top A^2 s = \|As\|_2^2 \le \Delta^2 N$. Together with (7.6) and (7.7), this bounds the center value between $N/4 + t(W/2 - C')$ and that quantity plus $\Delta^2 t^2 N/16$. Since

$$\frac{320}{81} + \frac{1}{16} = \frac{5201}{1296} < 5,$$

we obtain

$$\frac{N}{4} + t\left(\frac{W}{2} - C'\right) \le \mu(B\mathbb{Z}^N)^2 \le \frac{N}{4} + t\left(\frac{W}{2} - C'\right) + 3125t^2 N. \tag{7.9}$$

7.3 Thresholds and completion of the proof

Proof of Theorem 7.1. Use the deterministic source and cloud construction above, and form (7.7). Define

$$U = \frac{N}{4} + t(W/2 - a) + 3125t^2 N, \qquad L = \frac{N}{4} + t(W/2 - b).$$

By (7.9), the YES case has $\mu^2 \le U$ and the NO case has $\mu^2 \ge L$. Equations (7.5) and $t \le 10^{-7}$ give

$$L - U \ge \left(\frac{3}{512} - 3125t\right) tN \ge \frac{71}{12800} tN > \frac{tN}{200}. \tag{7.10}$$

Since $0 \le b < a \le W \le 25N/2$, we also have

$$U > \frac{N}{5}, \qquad U \le N\left(\frac{1}{4} + \frac{25t}{4} + 3125t^2\right) < \frac{N}{3} - 10^{-3}N. \tag{7.11}$$

For the lower bound, use $U \ge N(1/4 - 25t/4) > N/5$. The upper numerical inequality holds throughout the stated range up to $t = 10^{-7}$.

We again allocate the squared-radius separation to rounding and multiplicative expansion. We bound both costs by small constant multiples of tN, within the separation in (7.10). Set $h_0 = \lceil 10^6/t \rceil$ and $H = h_0 N$. Let k be the least nonnegative integer satisfying $k^2 \ge H^2 U$, and put $r = k/H$. As in (5.5), integer comparisons compute this threshold exactly. Since $\sqrt{U} \le r < \sqrt{U} + 1/H$,

$$0 \le r^2 - U < \frac{2\sqrt{U}}{H} + \frac{1}{H^2} \le \frac{2t}{10^6\sqrt{N}} + \frac{t^2}{10^{12}N^2} < \frac{tN}{10000}.$$

Here $U < N/3 < N$, $H \geq 10^6 N/t$, $N \geq 1$, and $t \leq 10^{-7}$ justify the successive inequalities. Equation (7.11) then gives $r^2 < N/3$. With $\gamma = 1 + t/1000$, we have $\gamma^2 - 1 \leq 3t/1000$, so

$$\gamma^2 r^2 - U < \frac{tN}{10000} + \frac{3t}{1000}\frac{N}{3} = \frac{11}{10000} tN < \frac{tN}{200} < L - U.$$

The YES case therefore satisfies $\mu \leq r$, and the NO case satisfies $\mu > \gamma r$.

The condition-number bound and the restrictions on the entries follow from (7.8) and the construction of the basis. All cloud edge weights are one or three, $N = O(m)$, and all entries of A are integers of absolute value at most 25. For fixed t, the entries of B have constant bit length, while U, H, k, r have $O(\log(N+1))$ bits with constants depending on t. Integer binary search computes k in polynomial time, as in the proof of Theorem 1.2. This completes the deterministic polynomial-time reduction. □

AI methodology

OpenAI Codex assisted with the introduction, English exposition, elaborating and checking proof details, LaTeX typesetting, reference verification, and consistency checks throughout the manuscript.

References

[BL26] Huck Bennett and Peter Ly. Hardness of the binary covering radius problem in large ℓ_p norms. In *APPROX/RANDOM 2026*, volume 392 of *LIPIcs*, pages 10:1–10:21, 2026. doi:10.4230/LIPIcs.APPROX/RANDOM.2026.10.

[Dad19] Daniel Dadush. On approximating the covering radius and finding dense lattice subspaces. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 1021–1026, 2019. doi:10.1145/3313276.3316397.

[GMR05] Venkatesan Guruswami, Daniele Micciancio, and Oded Regev. The complexity of the covering radius problem. *Computational Complexity*, 14(2):90–121, 2005. doi:10.1007/s00037-005-0193-y.

[Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001. doi:10.1145/502090.502098. Author version dated February 18, 2002: `https://people.kth.se/~johanh/optimalinap.pdf`.

[HR12] Ishay Haviv and Oded Regev. Hardness of the covering radius problem on lattices. *Chicago Journal of Theoretical Computer Science*, 2012, Article 4, pages 1–12. doi:10.4086/cjtcs.2012.004.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006. doi:10.1090/S0273-0979-06-01126-8.

[HLR09] Ishay Haviv, Vadim Lyubashevsky, and Oded Regev. A note on the distribution of the distance from a lattice. *Discrete & Computational Geometry*, 41(1):162–176, 2009. doi:10.1007/s00454-008-9123-5.

[Mag20] Alexander Magazinov. A proof of a conjecture by Haviv, Lyubashevsky and Regev on the second moment of a lattice Voronoi cell. *Advances in Geometry*, 20(1):117–120, 2020. doi:10.1515/advgeom-2018-0018.

[Val26] Frank Vallentin. The complexity of computing the covering radius of a Euclidean lattice. arXiv:2609.35027v1, 2026. https://arxiv.org/abs/2609.35027v1.