

Adversary Lower Bounds for Lattice Problems

Pushkar S Joglekar[‡] Sandip R. Shinde^{‡*} Aarti Amod Agarkar[§]

September 8, 2026

Abstract

The Shortest Vector Problem (SVP) and the Closest Vector Problem (CVP) are the fundamental algorithmic questions in the geometry of numbers. In the past two decades, their algorithmic complexity has been studied quite extensively due to their connection with lattice based cryptosystems. In this paper, we study these problems in the setting of generic black-box metric query model, which captures algorithms that rely on distance evaluations without exploiting specific geometry of the underlying metric. We prove that any randomized algorithm for solving SVP or CVP within any constant factor $\gamma < 2/\sqrt{3} \approx 1.154$ which accesses the metric as a black-box oracle needs to make at least $2^{\Omega(n)}$ queries to the metric oracle.

Our results rely on fundamental tools from the geometry of numbers and convex geometry. For our SVP lower bound, we use Siegel's Mean Value Theorem and Rogers' second moment bound to probabilistically construct an adversarial metric by introducing $2^{\Omega(n)}$ disjoint local convex perturbations to the underlying Euclidean norm. The lower bound for CVP follows from the SVP lower bound by observing that the known Turing reduction from SVP to CVP works in the black-box metric query model.

1 Introduction

Integer lattices are discrete additive subgroups of $\mathbb{R}^n$. The Shortest Vector Problem (SVP) and Closest Vector Problem (CVP) are most important algorithmic questions over integer lattices. Given an integer lattice by a basis and a target vector v , CVP asks to find a lattice vector u that minimizes the distance to v with respect to a given norm.

*This research was funded by Science and Engineering Research Board (SERB), Government of India grant number CRG/2020/001456. The author would like to thank SERB for the funding.

[†]Vishwakarma Institute of Technology, Pune, India, email: joglekar.pushkar@gmail.com

[‡]Vishwakarma Institute of Technology, Pune, India, email: sandeep.shinde@vit.edu

[§]Vishwakarma Institute of Technology, Pune, India, email: aarti.agarkar@vit.edu

Similarly, SVP asks to find the shortest non-zero vector in the lattice with respect to given norm.

The security of modern lattice based cryptography rely on the presumed intractability of the SVP, CVP and related lattice problems. So recently these problems have been studied extensively from computational complexity perspective. Theoretical computer science often relies on black-box models to separate inherent combinatorial difficulty from specific structural aspects of the problem. A classic parallel can be found in the generic group model: computing discrete logarithms requires exponential time if the algorithm accesses group operation through black-box oracle [30], demonstrating that the problem is fundamentally hard unless specific algebraic structure of the underlying group is exploited. In this paper, we apply an analogous philosophy to lattice problems: we investigate the fundamental hardness of SVP and CVP when the algorithm is allowed to access the underlying metric only through black-box oracle; it can evaluate distances but cannot exploit the specific geometry of the underlying metric.

The algorithmic complexity of SVP, CVP and related lattice problems has been studied for their exact as well as approximation versions. The exact versions of both the problems have $2^{O(n)}$ algorithms [3], [24]. We also know efficient polynomial time algorithms to solve the problems within exponential factor of approximation [22], [8].

The exact versions of CVP and SVP are NP-hard, in fact both the problems are even hard to approximate under some believable complexity theory assumptions [23], [26]. We also know strong conditional hardness results for these problem. For instance, assuming the Exponential Time Hypothesis (ETH), computing exact SVP and CVP is known to require $2^{\Omega(n)}$ time [11], [1]. However, proving unconditional computational lower bounds for the approximation variants, the regime in which practical lattice cryptography operates, remains completely out of reach of existing techniques.

This naturally motivates the study of lattice problems within the generic black-box metric query model. We say an algorithm accesses the norm through a generic black-box oracle if the explicit geometric or algebraic definition of the norm is unknown to the algorithm; instead, whenever it wants to find the norm of a vector $w \in \mathbb{Q}^n$, it queries the oracle, which returns the norm of w in $O(1)$ time.

This model captures algorithms that do not exploit the specific geometry of underlying metric. A notable example of such algorithms include Ajtai-Kumar-Sivakumar (AKS) randomized sieving [3], which solves the exact shortest lattice vector problem in $2^{O(n)}$ time. It turns out that the AKS sieving technique works for general symmetric norms given by gauge functions which subsumes usual ℓ_p norms (see e.g. [14], [7]). In fact, it even adapts to asymmetric norms [17]. Moreover, AKS sieving, as well as some of the subsequent algorithms based upon it, can be seen to need only black-box oracle access for the underlying norm [7]. By restricting the algorithm to access the underlying metric through queries to a black-box, we capture the capabilities of generic AKS-style sieving and other related combinatorial algorithms.

1.1 Our Contributions

In this work, we establish adversary lower bounds on the query complexity of randomized algorithms operating in the black-box metric framework. Our main results are:

1. **Exponential Query Lower Bound for SVP:** We prove that any randomized algorithm requires at least $2^{\Omega(n)}$ queries to the metric oracle to approximate SVP within any constant factor $\gamma < 2/\sqrt{3} \approx 1.154$.
2. **Transferability to CVP:** We prove that this $2^{\Omega(n)}$ lower bound extends directly to approximate CVP within the same constant factor $\gamma < 2/\sqrt{3}$.

1.2 Technical Overview

The main idea of our lower bound for SVP relies on designing an adversarial metric oracle that isolates a specific target vector. To force a randomized algorithm to make $2^{\Omega(n)}$ queries, the oracle partitions the space into $2^{\Omega(n)}$ strictly disjoint conical regions. It then constructs a new metric by applying local convex perturbations to the underlying Euclidean ball exclusively within these cones. By selecting one of these conical regions uniformly at random and radially stretching the metric inside it, the oracle embeds the target vector while leaving the rest of the space unperturbed.

The Ideal Setting (Exact Shortest Vectors): Consider an idealized setting where the input lattice $\mathcal{L}$ possesses $2^{\Omega(n)}$ exact shortest vectors (i.e., a lattice exhibiting an exponential kissing number). By the definition of the first successive minimum $\lambda_1(\mathcal{L})$, any two non-collinear shortest vectors $u, v \in \Lambda$ must satisfy $\|u \pm v\|_2 \geq \lambda_1$. This implies that the angle θ subtended by u and v at the origin is bounded from below by $\pi/3$. This strict angular separation allows the adversary to construct $2^{\Omega(n)}$ strictly disjoint conical regions of half-angle $\alpha < \pi/6$, each radiating from the origin in the direction of a distinct shortest vectors. The adversary selects a target vector v^* uniformly at random. To hide this vector, the oracle defines a new metric by radially dilating the unperturbed Euclidean unit ball strictly within the chosen conical region to achieve a stretch factor γ . Crucially, to ensure that this local perturbation maintains the strict mathematical convexity of the resulting unit ball, the maximum achievable factor γ depends directly on the lower bound of the angle between the target vectors. Using basic trigonometry we can see that the stretch factor is bounded by $\gamma \leq 1/\cos(\alpha)$. Because the angular separation of shortest vectors restricts the cone's half-angle to $\alpha < \pi/6$, the approximation factor is fundamentally capped at $\gamma < 1/\cos(\pi/6) = 2/\sqrt{3}$. Because the conical regions are perfectly disjoint and the local convex perturbations are exclusively within the conical regions, the norms of all other shortest vectors remains unchanged, forcing any algorithm to query $2^{\Omega(n)}$ disjoint spatial regions to locate v^* .

However, relying on this idealized scenario presents a formidable mathematical hurdle. Explicitly constructing lattice families with exponential large kissing numbers remains a notoriously difficult open problem in the geometry of numbers [12]. The

current best known explicit construction based on Barnes-Wall lattices [9] achieves only quasi-polynomial bound ($n^{\Theta(\log n)}$) on the kissing number [16]. In 2019 Vlăduț published a paper claiming a construction of lattices with exponential kissing number [34]. However, Bennett, Golovnev and Stephens-Davidowitz [12] showed formally that the arguments there (and in a follow-up paper) are invalid, after which Vlăduț withdrew the arXiv version of [34] in 2024; a subsequent preprint attempting to repair the construction [35] was also withdrawn by its author in 2025.

While probabilistic arguments guarantee exponential kissing numbers for general non-lattice sphere packings, they fail for lattices: a random lattice sampled under the Haar-Siegel measure has an exact kissing number of just two almost surely, as accidental length equalities occur with probability zero. So even showing an *existence* of lattice families with exponential kissing number appears to be a formidable open problem.

Our Resolution (The dense ϵ -shells): We bypass this theoretical barrier by demonstrating that the adversarial construction does not strictly require exact shortest vectors. It is mathematically sufficient to hide the target among a exponentially many "almost" shortest lattice vectors. We relax the exactness requirement to an ϵ -shell; annular subset of $\mathbb{R}^n$ containing lattice vectors with lengths in the interval $[\lambda_1, (1 + \epsilon)\lambda_1]$. By simple application of Siegel's Mean Value Theorem and Rogers' variance bound we probabilistically guarantee the existence of a lattice Λ containing $2^{\Omega(n)}$ vectors strictly bounded within this narrow shell. We note that Siegel's mean value theorem, Rogers second moment bound are quite well-known tools in geometry of numbers and in fact their use to prove existence of lattice families with certain properties is also quite well-studied, see e.g. [36] for a latest related work relying on these theorems.

Next, we analyze the angular separation of the vectors in the ϵ -shell and prove that we can choose sufficiently small ϵ (depending on the constant δ) such that the angle θ between any two vectors within this specific shell remains strictly bounded below by $\pi/3 - \delta$ where $\delta > 0$ is an absolute constant. This preserves the necessary angular separation to construct $2^{\Omega(n)}$ disjoint conical caps on the unperturbed Euclidean unit ball, finalizing the adversarial metric construction without requiring explicit exponential kissing numbers.

Because lattices with exact exponential kissing numbers are not known to exist, our geometric relaxation to a probabilistically guaranteed ϵ -shell is quite crucial for our proof.

Transferring the Bound to CVP: To extend our lower bound to the Closest Vector Problem, we utilize the standard Turing reduction from SVP to CVP [19]. We note that this reduction works with respect to any convex norm and can be seen to operate entirely in a black-box manner, it simply queries the CVP subroutine using the exact same metric oracle provided for SVP. Consequently, any randomized algorithm attempting to solve CVP in the generic metric model is forced to interact with our unaltered adversarial metric, directly transferring the $2^{\Omega(n)}$ query lower bound to CVP without requiring any structural changes to the oracle.

Remark 1.1. *While concepts like "angles" and "tangential cones" provide useful intuition in lower dimensions, visual geometry in n -dimensional space can often be highly misleading. So in our proofs we rely on concrete algebraic definitions (dot products, gauge functions, convex hulls, etc) rather than spatial analogies to ensure the rigor.*

van Emde Boas [33] proved NP-hardness of CVP with respect to any ℓ_p norm and for SVP with respect to ℓ_∞ norm. In fact, based on some natural complexity theory assumptions we know that CVP is a hard problem for certain approximation factors as large as $2^{O(\frac{\log n}{\log \log n})}$ (see e.g. [5], [18]). Proving that CVP is hard to approximate within polynomial factor is an important open problem in the area. van Emde Boas [33] conjectured hardness for SVP in ℓ_2 norm. and it remained probably the biggest open problem in the area for almost two decades. In a seminal paper Ajtai [2] proved that SVP is NP-hard under randomized reduction. NP-hardness of approximation version of SVP has also been studied relying on some complexity theory assumptions. For example, for any $\epsilon > 0$ there is no polynomial time algorithm approximating SVP on n -dimensional lattice in ℓ_p norm to within a factor of $2^{O((\log n)^{1-\epsilon})}$ under some complexity theory assumptions [20]. [23] is for an excellent survey on the complexity of lattice problems.

Another direction of research was to find an efficient algorithm to solve CVP and SVP exactly. [21], [13] gave $2^{O(n \log n)}$ time *deterministic* algorithm to solve SVP or CVP exactly with respect to ℓ_p norm. In a breakthrough paper [3] Ajtai, Kumar and Sivakumar gave a $2^{O(n)}$ time *randomized* exact algorithm for SVP for ℓ_2 norm. Subsequently, using sieving technique they gave a $2^{O(n)}$ time approximation algorithm for CVP [4]. We also know a $2^{O(n)}$ exact algorithm for CVP in ℓ_2 norm due to Micciancio and Volgarious [24]. There are several results which give exponential time algorithms for CVP/ SVP by improving the constants in the exponent of the running time. See e.g. [28] for a recent $2^{0.802n}$ algorithm for a constant factor approximation of CVP.

The AKS sieving technique is generic; it works not only for standard ℓ_p norms but can also be adapted for norms given by gauge functions [14], [7] and even asymmetric norms [17]. Following the original AKS result, there are some follow-up papers which provides heuristics to significantly improved the exact time and space complexities for both SVP and CVP [25], [10] (they have exponential running time but improve the constant in the exponent as compared to original AKS).

We also know strong conditional hardness results for these problem. For example, assuming the Exponential Time Hypothesis or strong ETH we know solving SVP and CVP approximately within certain factors requires exponential time [11], [1].

When discussing the theoretical limits on solving lattice problems, it is critical to distinguish between computational time lower bounds and query complexity lower bounds. The lower bounds based on the ETH measure the computational time required by a Turing machine that has full, explicit access to the lattice basis and the exact norm definition. These bounds work for any algorithm for the problems but they are conditional; they depend on the assumption that the ETH holds true. By contrast, the adversary bounds explored in this work measure query complexity under a black-box

oracle model. They establish an unconditional lower bound of $2^{\Omega(n)}$ on the query complexity of the algorithm but they work only for specific class of algorithms: namely the algorithms which access the metric through a black-box oracle. This holds regardless of the algorithm's internal computational power. Even if $P=NP$, the algorithm mathematically lacks the geometric information required to output the correct answer.

The concept of utilizing tangential conical regions to analyze this geometric information was previously explored by Arvind and one of the current authors [6] for the exact Shortest Vector Problem. However, their geometric analysis was inherently limited to lattice families possessing at most a polynomial ($\text{poly}(n)$) kissing number. Furthermore, their lower bound applies specifically to algorithms required to output all shortest vectors *exactly*, rather than finding just one. Their lower bound doesn't work for approximation algorithms for the problem. In contrast, our adversarial framework bounds the query complexity of finding a single approximate shortest vector (up to a factor of $\gamma < 2/\sqrt{3}$). By relaxing the requirement of exact shortest vectors and utilizing probabilistically guaranteed dense ϵ -shells, we establish a strictly exponential $2^{\Omega(n)}$ lower bound for approximate SVP and CVP.

2 Preliminaries

2.1 Convex bodies and gauge functions

A *convex body* $K \subset \mathbb{R}^n$ is a compact and convex subset of $\mathbb{R}^n$ with a non-empty interior. K is called *O-symmetric* if for all $x \in K$, $-x \in K$.

Next we formally define functions which are useful for defining general notion of norms w. r. t. O-symmetric convex bodies. They are sometimes referred as Minkowski functionals in the literature. We call them gauge functions following Siegel [31].

Definition 2.1 (gauge function). [31] *A function $f : \mathbb{R}^n \rightarrow \mathbb{R}^+$ is called a gauge function if it satisfies following properties:*

1. $f(x) > 0$ for all $x \in \mathbb{R}^n \setminus \{0\}$ and $f(x) = 0$ if $x = 0$.
2. $f(\lambda x) = \lambda f(x)$ for all $x \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}^+$.
3. $f(x + y) \leq f(x) + f(y)$ for all $x, y \in \mathbb{R}^n$.
4. $f(x) = f(-x)$ for all $x \in \mathbb{R}^n$.

For $u \in \mathbb{R}^n$, we denote $f(u)$ by $\|u\|_f$ and call it norm of u with respect to gauge function f . We know that usual l_p norms satisfy all the above properties. A gauge function f naturally defines a metric on $\mathbb{R}^n$. For points $x, y \in \mathbb{R}^n$ the distance between them is simply given by $f(x - y)$. For $x \in \mathbb{R}^n$ and $r > 0$, let $B_f(x, r)$ denote the f -ball of radius r with center x :

$$B_f(x, r) = \{y \in \mathbb{R}^n | f(x - y) \leq r\}$$

Proposition 2.2. [31] *Let $f : \mathbb{R}^n \rightarrow \mathbb{R}^+$ be any gauge function then $B_f(0, 1)$ (a unit radius ball around origin) is a n dimensional bounded O-symmetric convex body. Conversely for any n dimensional bounded O-symmetric convex body C , there is a gauge function $f : \mathbb{R}^n \rightarrow \mathbb{R}^+$ such that $B_f(0, 1) = C$.*

So given any bounded O-symmetric convex body K there is a natural gauge function f_K associated with it and for any $x \in \mathbb{R}^n$ we can define norm of x with respect to K as

$$\|x\|_K = f_K(x) = \inf\{s > 0 \mid x \in sK\}$$




Figure 1: Norm w.r.t. O-symmetric convex body

2.2 Integer lattices and lattice problems

Let $B = \{b_1, \dots, b_n\} \in \mathbb{R}^n$ be a collection of n linearly independent vectors. The lattice $\mathcal{L}$ generated by B is the set of all integer linear combinations of $b_1, b_2, \dots, b_n$:

$$\mathcal{L} = \{x_1 b_1 + x_2 b_2 + \dots + x_n b_n \mid x_i \in \mathbb{Z} \text{ for } i \in [n]\}$$

. The set B is called as a basis for the lattice $\mathcal{L}$.

The length of shortest non-zero vector in a lattice $\mathcal{L}$ with respect to metric specified by a bounded O-symmetric convex body K is called as the first successive minima of $\mathcal{L}$ and denoted by $\lambda_1(\mathcal{L}, K)$:

$$\lambda_1(\mathcal{L}, K) = \inf\{\|v\|_K \mid v \in \mathcal{L} \setminus \{0\}\}$$

If the lattice and underlying metric is clear from the context we simply use λ_1 to denote the first successive minima.

Definition 2.3 (Shortest Vector Problem). *Given an integer lattice $\mathcal{L}$ by a basis, the shortest vector problem (SVP) is to find a shortest non zero vector in $\mathcal{L}$ w.r.t. a given*

metric specified by O -symmetric convex body K . In other words, the goal is to find $u \in \mathcal{L}$ such that

$$\|u\|_K = \lambda_1(\mathcal{L}, K)$$

In the approximation version SVP_γ of the problem, the goal is to find $u \in \mathcal{L}$ such that

$$\|u\|_K \leq \gamma \cdot \lambda_1(\mathcal{L}, K)$$

, where $\gamma \geq 1$ is called as factor of approximation.

Definition 2.4 (Closest Vector Problem). *Given an integer lattice $\mathcal{L}$ by a basis and a vector $v \in \mathbb{R}^n$, the closest vector problem (CVP) is to find $u \in \mathcal{L}$ closest to v w.r.t. a given metric specified by O -symmetric convex body K . In other words, the goal is to find $u \in \mathcal{L}$ such that*

$$\|v - u\|_K = \min_{w \in \mathcal{L}} \|v - w\|_K$$

In the approximation version CVP_γ of the problem, the goal is to find $u \in \mathcal{L}$ such that

$$\|v - u\|_K \leq \gamma \cdot \min_{w \in \mathcal{L}} \|v - w\|_K$$

Definition 2.5 (Kissing Number). *The kissing number of an integer lattice $\mathcal{L}$ is number of shortest non-zero vectors in $\mathcal{L}$ (w. r. t. given metric). In other words, kissing number is equal to number of lattice vectors of length $\lambda_1(\mathcal{L}, K)$.*

The kissing number of lattices is mainly studied w. r. t. ℓ_2 norm. From simple packing argument we can see that kissing number of any n dimensional lattice is upper bounded by $2^{O(n)}$. We refer to [16] for interesting history of kissing numbers.

Table 1 gives kissing numbers for some well-known lattice families in Euclidean space.

As noted in the introduction, showing existence of lattice families with $2^{\Omega(n)}$ kissing number remains a big open problem in geometry of numbers [12]. In the context of our adversarial metric construction, the kissing number dictates the number of strictly disjoint conical regions the oracle can embed. As we do not have explicit (or even existential) lattice families of dimension n with $2^{\Omega(n)}$ kissing number, it makes the relaxed ϵ -shell construction presented in Section 3 quite crucial to get $2^{\Omega(n)}$ lower bound on query complexity.

2.3 The Space of Lattices and Probabilistic Methods

For construction of our adversarial metric, we need to show existence of lattices which have exponentially many "almost" shortest lattice vectors. To prove the existence of such lattices without relying on explicit constructions, we need to use probabilistic method over the space of all lattices. Let $X_n = SL(n, \mathbb{R})/SL(n, \mathbb{Z})$ denote the continuous space of all full-rank lattices in $\mathbb{R}^n$ with a fundamental parallelepiped of volume 1 (unit-covolume lattices). This space admits a unique, translation-invariant probability

Table 1: Well-known lattice families and their kissing numbers.

Lattice Family	Dim. (n)	Kissing No. $\tau(\Lambda)$	Geometric Significance
Integer Lattice ($\mathbb{Z}^n$)	n	$2n$	The standard hypercubic lattice. Its kissing number grows only linearly.
Root Lattice (A_n)	$n \geq 1$	$n(n+1)$	Formed by points in $\mathbb{Z}^{n+1}$ whose coordinates sum to zero. Exhibits quadratic growth.
Root Lattice (D_n)	$n \geq 3$	$2n(n-1)$	Consists of vectors in $\mathbb{Z}^n$ with an even coordinate sum.
E_8 Lattice	8	240	Achieves the theoretically maximum possible kissing number in 8 dimensions.
Leech Lattice (Λ_{24})	24	196,560	Achieves the strictly optimal kissing number and sphere packing density in 24 dimensions.
Barnes-Wall (BW_n)	$n = 2^d$	$n^{\Theta(\log n)}$	The currently best-known explicit lattice construction for large n , achieving quasi-polynomial growth.

measure known as the normalized Haar-Siegel measure, denoted by μ . By equipping X_n with this measure, we can mathematically treat a lattice Λ as a random variable drawn from X_n , allowing us to compute the expected number of lattice points falling into arbitrary geometric regions. To compute these expectations, we rely on a foundational result by Siegel [32], [15] which bridges the discrete counting of lattice points with continuous Lebesgue integration.

Theorem 2.6 (Siegel’s Mean Value Theorem). *For any Riemann-integrable function $f : \mathbb{R}^n \rightarrow \mathbb{R}$ of compact support, the expected sum of f over the non-zero vectors of a random unit-covolume lattice $\Lambda \sim \mu$ is exactly equal to the Lebesgue integral of f over $\mathbb{R}^n$:*

$$\mathbb{E}_{\Lambda \sim \mu} \left[\sum_{x \in \Lambda \setminus \{0\}} f(x) \right] = \int_{\mathbb{R}^n} f(x) dx$$

We need to analyze the variance of these lattice point distributions, for this purpose we utilize Rogers’ integration formula [27], which bounds the second moment of the Haar-Siegel measure. In the original paper Rogers gave explicit series for the variance, the upper bound on the variance as stated below is from [29].

Theorem 2.7 (Rogers-Schmidt variance bound). *Let $n \geq 3$. For any bounded, measurable subset $S \subset \mathbb{R}^n$, let $N(S, \Lambda) = |\Lambda \cap S|$. Under the Haar-Siegel measure on X_n , the variance of $N(S, \Lambda)$ is bounded by a constant multiple of its volume:*

$$\text{Var}(N(S, \Lambda)) \leq C \cdot \text{Vol}(S)$$

where $C > 0$ is an absolute constant independent of the volume of S .

2.4 The Black-Box Oracle Model

In this model, the algorithm does not know the gauge function explicitly. It must learn the geometry by querying the oracle for the norm of arbitrary vectors $w \in \mathbb{Q}^n$. Our lower bounds hold against general queries; the algorithm is not restricted to querying only displacement vectors between lattice points and targets.

An algorithm $\mathcal{A}$ operates in this model if it does not possess an explicit geometric or algebraic description of the convex body K or the corresponding gauge function. Instead, $\mathcal{A}$ is provided with a black-box oracle $\mathcal{O}_K$. For any query vector $x \in \mathbb{Q}^n$, the oracle returns the exact norm of x w. r. t. the metric given by the convex body K in $O(1)$ time:

$$\mathcal{O}_K(x) = \|x\|_K$$

. The query complexity of an algorithm $\mathcal{A}$ is defined as the total number of calls made to the oracle $\mathcal{O}_K$. The algorithm may perform unbounded internal computational steps, which are not charged. This model perfectly captures generic distance-based combinatorial algorithms, such as AKS randomized sieving.

3 Construction of Adversarial Convex Body

First we prove existence of lattices with dense ϵ -shell which play a crucial role in the construction of adversarial convex body.

3.1 Lattices with dense ϵ -shells

Our construction of adversarial metric oracle crucially relies on existence of a lattice $\mathcal{L}$ which has exponentially many lattice points inside a ball of radius $(1 + \epsilon) \cdot \lambda_1(\mathcal{L})$.

We prove existence of such a lattice in the Theorem below.

Theorem 3.1. *For any constant $\epsilon > 0$ and sufficiently large n , there exists a full-rank lattice $\Lambda \subset \mathbb{R}^n$ with first successive minimum $\lambda_1 > 0$ such that the number of lattice vectors in the ϵ -shell S_ϵ is exponential in n :*

$$|\Lambda \cap S_\epsilon| \geq 2^{\Omega(n)}$$

where

$$S_\epsilon = \{x \in \mathbb{R}^n \mid \lambda_1 \leq \|x\|_2 \leq \lambda_1 \cdot (1 + \epsilon)\}$$




Figure 2: Dense ϵ -shell

Proof. Let $X_n = SL(n, \mathbb{R})/SL(n, \mathbb{Z})$ denote the space of unit co-volume lattices in $\mathbb{R}^n$, with the unique normalized Haar-Siegel probability measure μ . Let $\mathcal{L}$ is sampled from X_n with respect to measure μ . Let $B(R) \subset \mathbb{R}^n$ be the ball of radius R around origin w. r. t. ℓ_2 norm. We carefully choose R such that the volume of $B(R)$ is exactly a small constant, say $\text{Vol}(B(R)) = 1/3$. Let $A_\epsilon(R)$ denote the closed ϵ -shell surrounding this ball, defined as:

$$A_\epsilon(R) = \{x \in \mathbb{R}^n \mid R \leq \|x\|_2 \leq R \cdot (1 + \epsilon)\}$$

. Consider following two events:

- E_1 be the event that the shortest nonzero vector of $\mathcal{L}$ has length at least R .
- E_2 be the event that $A_\epsilon(R)$ contains $2^{\Omega(n)}$ vectors of lattice $\mathcal{L}$.

If we can show that both the events E_1, E_2 happen with sufficiently large probabilities then using union bound we can argue that event $E_1 \cap E_2$ happens with a non-zero probability which would suffice to show existence of lattice with required property. Let us analyze probabilities of the events E_1 and E_2 .

Let f be the indicator function of $B(R)$. From Siegel's Mean Value Theorem (Theorem 2.6) it follows that the expected number of non-zero lattice vectors of $\mathcal{L}$ strictly inside $B(R)$ is precisely volume of $B(R)$ which is $1/3$. By Markov's inequality, the probability that $B(R)$ contains at least one lattice point is at most $1/3$. Therefore, with probability at least $2/3$, the shortest non-zero vector of $\mathcal{L}$ has length $\geq R$. Hence $\mathbb{P}[E_1] \geq 2/3$.

Next, we analyze the probability of event E_2 . Let us denote the volume of the shell $A_\epsilon(R)$ by V_A . Clearly V_A grows exponentially with the dimension n :

$$V_A = \text{Vol}(B(R)) ((1 + \epsilon)^n - 1) = \frac{1}{3} ((1 + \epsilon)^n - 1) = 2^{\Omega(n)}$$

. By Theorem 2.6, the expected number of lattice points in $A_\epsilon(R)$ is exactly V_A :

$$\mathbb{E}_{\mathcal{L} \sim \mu} |\mathcal{L} \cap A_\epsilon(R)| = V_A$$

. By Theorem 2.7 we know that the variance of the random variable $|\mathcal{L} \cap A_\epsilon(R)|$ is upper bounded by CV_A :

$$\text{Var} |\mathcal{L} \cap A_\epsilon(R)| \leq CV_A$$

where C is an absolute constant.

Applying Chebyshev's inequality to the random variable $N_A = |\mathcal{L} \cap A_\epsilon(R)|$, the probability that the number of lattice points deviates from the expectation by a large factor is bounded by:

$$\mathbb{P} \left(|N_A - V_A| \geq \frac{V_A}{2} \right) \leq \frac{4 \cdot \text{Var}(N_A)}{V_A^2} \leq \frac{4C}{V_A}$$

Since $V_A = 2^{\Omega(n)}$, this probability goes to 0 exponentially fast as n approaches ∞ . Thus, with probability $1 - o(1)$, a random lattice contains $\Theta(V_A) = 2^{\Omega(n)}$ points in $A_\epsilon(R)$. So $\mathbb{P}[E_2] \geq 1 - o(1)$.

By the union bound, the event $E_1 \cap E_2$ happens with a strictly positive probability (at least $2/3 - o(1) > 0$). So there exist lattice Λ with no non-zero lattice points strictly inside $B(R)$ and contains $2^{\Omega(n)}$ lattice points in the shell $A_\epsilon(R)$. For this specific lattice, the first successive minimum satisfies $\lambda_1 \geq R$. Consequently, every lattice vector x residing in the shell $A_\epsilon(R)$ has length bounded by:

$$\lambda_1 \leq \|x\|_2 \leq R \cdot (1 + \epsilon) \leq \lambda_1 \cdot (1 + \epsilon)$$

This implies that all $2^{\Omega(n)}$ vectors in $A_\epsilon(R)$ lie strictly within the desired ϵ -shell S_ϵ . This completes the proof. $\square$

3.2 Angular separation of ϵ -shell vectors

Let $\Lambda \subset \mathbb{R}^n$ be a full rank lattice. Let

$$S_\epsilon = \{x \in \mathbb{R}^n \mid \lambda_1(\Lambda) \leq \|x\|_2 \leq \lambda_1(\Lambda) \cdot (1 + \epsilon)\}$$

To ensure that the adversarial conical perturbations are strictly disjoint, we must prove that the angle between any two target vectors in the ϵ -shell S_ϵ is sufficiently large. We use simple geometry to prove the lower bound on the angular separation. A simple useful fact is: norm of difference of any two non-collinear lattice vectors is bounded from below by the length of shortest non-zero vector in Λ , $\lambda_1(\Lambda)$.

Lemma 3.2. *For any constant $\delta > 0$, there is a constant $\epsilon > 0$ (ϵ depends on δ) such that for any two distinct non-collinear lattice vectors $x, y \in \Lambda \cap S_\epsilon$ the angle θ between them is bounded from below by $\pi/3 - \delta$.*

Proof. Let x and y be two distinct, non-collinear lattice vectors in the ϵ -shell S_ϵ . By the definition of the ϵ -shell, the ℓ_2 norm of both vectors satisfy:

$$\lambda_1 \leq \|x\|_2, \|y\|_2 \leq (1 + \epsilon)\lambda_1$$

As Λ is a lattice, the difference $x - y$ is also a non-zero lattice vector. By the definition of the first successive minimum, the length of this difference vector cannot be strictly smaller than λ_1 , $\|x - y\|_2 \geq \lambda_1$. Let θ be angle between x and y . Consider the geometric triangle formed by the origin, x , and y . Using cosine rule, we get:

$$\cos \theta = \frac{\|x\|_2^2 + \|y\|_2^2 - \|x - y\|_2^2}{2\|x\|_2\|y\|_2}$$

To find the minimum possible angle θ , we must find the maximum possible value for $\cos \theta$. We may assume that the numerator is positive: otherwise $\cos \theta \leq 0 < \frac{1}{2} + 2\epsilon + \epsilon^2$ and the bound derived below holds trivially. Since the upper bound obtained for the numerator below is itself positive, we may bound the numerator from above and the denominator from below independently.

To maximize $\cos \theta$ we replace $\|x\|_2$ and $\|y\|_2$ with their maximum possible value $(1 + \epsilon)\lambda_1$ in the numerator, and we replace the subtracted term $\|x - y\|_2$ with its minimum possible value λ_1 . We replace both $\|x\|_2$ and $\|y\|_2$ in the denominator with their minimum possible value λ_1 . Applying these bounds yields:

$$\cos \theta \leq \frac{(1 + \epsilon)^2\lambda_1^2 + (1 + \epsilon)^2\lambda_1^2 - \lambda_1^2}{2\lambda_1 \cdot \lambda_1}$$

Which implies

$$\cos \theta \leq \frac{2(1 + \epsilon)^2 - 1}{2} = (1 + \epsilon)^2 - \frac{1}{2}$$

Expanding the squared term gives:

$$\cos \theta \leq 1 + 2\epsilon + \epsilon^2 - 1/2 = \frac{1}{2} + 2\epsilon + \epsilon^2$$

We have $\cos(\pi/3) = 1/2$. By the continuity of the cosine function, for any arbitrarily small $\delta > 0$, we can choose a sufficiently small $\epsilon > 0$ such that the small additive error $(2\epsilon + \epsilon^2)$ satisfies:

$$\cos \theta \leq \frac{1}{2} + 2\epsilon + \epsilon^2 \leq \cos\left(\frac{\pi}{3} - \delta\right)$$

Because the cosine function is monotonically decreasing on the interval $[0, \pi]$, this implies:

$$\theta \geq \frac{\pi}{3} - \delta$$

□

3.3 Integer lattice with dense ϵ -shell

The Haar–Siegel dense-shell argument in the Section 3.1 produces a lattice Λ with a *real* basis. One possible formulation is therefore an exact-real Blum–Shub–Smale (BSS) model, allowing real inputs and queries. If the intended model uses rational/integer lattice and rational query points (subject to a separate convention for exact oracle answers), we need to obtain an integer/rational lattice from the real lattice Λ which has similar exponential bound on number of vectors in ϵ -shell and bound on angular separation of these vectors (similar to Lemma 3.2).

At a high level, we choose a reduced basis of Λ , multiply the basis by a sufficiently large integer, and round its entries. The perturbation is uniform over all lattice coefficient vectors, so it controls the new first minimum as well as the selected near-short vectors; Minkowski’s second theorem together with a reduced-basis estimate bounds the precision needed for rounding. The result is an integral lattice with a polynomial-size basis, exponentially many near-short vectors, and the same angular slack. The above high level plan is executed in Theorem 3.3. We give the proof in the Appendix.

Theorem 3.3 (Polynomially encoded integral dense shell). *For every fixed $\epsilon > 0$ and $0 < \delta < \pi/3$, there are constants $c = c(\epsilon, \delta) > 0$, $C = C(\epsilon, \delta) > 0$, and n_0 such that for every $n \geq n_0$ there exist a nonsingular integer matrix $A \in \mathbb{Z}^{n \times n}$ and a symmetric set $U \subset S_\epsilon(\mathcal{M}) \setminus \{0\}$ where $\mathcal{M}$ is a lattice generated by columns of A satisfying*

1. *the binary encoding length of A is at most Cn^4 ;*
2. *$|V| \geq 2^{cn}$ where $V = U/\{\pm 1\}$ be a collection of vectors obtained by identifying $u, -u \in U$. That is, V is collection of pairwise non-antipodal points in U .*
3. *$\text{angle}(u, v) > \pi/3 - \delta$ whenever $u \neq v$ in V , where*

$$\text{angle}(u, v) := \arccos \left(\frac{|\langle u, v \rangle|}{\|u\|_2 \|v\|_2} \right).$$

Remark 3.4 (Model boundary). *The integral basis makes the SVP input and the bases, targets, returned lattice vectors, and residuals in the parity reduction rational. The result is existential, as is the original random-lattice witness. It does not encode arbitrary exact real answers from the unknown norm oracle; an exact-real atomic-answer convention, or a specified approximate-value model, is still needed for that separate issue. Also, the exponential count is for fixed ϵ and δ . If either varies with n , both the exponent and the displayed precision bound must retain their dependence on the internal η .*

3.4 Conical Regions and Adversarial Convex Body

Next we define exponentially many conical regions and our adversarial convex body and prove their useful properties.

Definition 3.5 (Conical Region). For a given vector $u \in \mathbb{R}^n \setminus \{0\}$ and an angle parameter $\theta \in [0, \pi]$, the conical region $\mathcal{F}_{u,\theta}$ is defined as:

$$\mathcal{F}_{u,\theta} = \left\{ x \in \mathbb{R}^n \setminus \{0\} \mid \frac{\langle x, u \rangle}{\|x\|_2 \|u\|_2} \geq \cos \theta \right\}$$

Let $\delta > 0$ be a constant. By Lemma 3.2 we know that there is a constant ϵ (which depends on δ) such that angular separation between any two non-collinear lattice points in ϵ -shell S_ϵ is at least $\pi/3 - \delta$ for any lattice Λ . Let $\Lambda \subset \mathbb{R}^n$ be a full rank lattice such that there are at least $2^{\Omega(n)}$ lattice vectors in the ϵ -shell S_ϵ . We know that such lattice exists by Theorem 3.1. Using simple geometry and trigonometry we prove that conical regions corresponding to any two non-collinear lattice vectors $u, v \in S_\epsilon$ are disjoint for $\theta = \pi/6 - \delta/2$.

Lemma 3.6. Let $u, v \in \Lambda \cap S_\epsilon$ be two distinct, non-collinear lattice vectors from the ϵ -shell, and let $\delta > 0$ be the gap parameter from Lemma 3.2. If $\theta < \pi/6 - \delta/2$, then $\mathcal{F}_{u,\theta} \cap \mathcal{F}_{v,\theta} = \emptyset$.




Figure 3: Disjoint conical regions

Proof. We will give a proof by contradiction. Assume there exists a non-zero vector $x \in \mathcal{F}_{u,\theta} \cap \mathcal{F}_{v,\theta}$. Let $\tilde{x} = x/\|x\|_2$, $\tilde{u} = u/\|u\|_2$, and $\tilde{v} = v/\|v\|_2$ denote the unit vectors along the direction of x , u and v respectively. By the definition of the conical regions, we have:

$$\langle \tilde{x}, \tilde{u} \rangle = c_1 \geq \cos \theta \quad \text{and} \quad \langle \tilde{x}, \tilde{v} \rangle = c_2 \geq \cos \theta$$

We split $\tilde{u}$ and $\tilde{v}$ into components parallel and orthogonal to $\tilde{x}$:

$$\tilde{u} = c_1 \tilde{x} + w_u \quad \text{where} \quad \langle w_u, \tilde{x} \rangle = 0$$

$$\tilde{v} = c_2 \tilde{x} + w_v \quad \text{where} \quad \langle w_v, \tilde{x} \rangle = 0$$

Because $\tilde{u}$ and $\tilde{v}$ are unit vectors, by the Pythagorean theorem, the squared norms of their orthogonal components are exactly $\|w_u\|_2^2 = 1 - c_1^2$ and $\|w_v\|_2^2 = 1 - c_2^2$. Let us consider the normalized inner product between the two target vectors u and v :

$$\frac{\langle u, v \rangle}{\|u\|_2 \|v\|_2} = \langle \tilde{u}, \tilde{v} \rangle = \langle c_1 \tilde{x} + w_u, c_2 \tilde{x} + w_v \rangle$$

As $\langle \tilde{x}, \tilde{x} \rangle = 1$ and the cross terms vanish, this simplifies to:

$$\langle \tilde{u}, \tilde{v} \rangle = c_1 c_2 + \langle w_u, w_v \rangle$$

Clearly, $\langle w_u, w_v \rangle \geq -\|w_u\|_2 \|w_v\|_2$. Substituting the norms gives:

$$\langle \tilde{u}, \tilde{v} \rangle \geq c_1 c_2 - \sqrt{1 - c_1^2} \sqrt{1 - c_2^2}$$

Let us define the function $f(c_1, c_2) = c_1 c_2 - \sqrt{1 - c_1^2} \sqrt{1 - c_2^2}$. For $c_1, c_2 \in [0, 1]$, standard calculus shows that f is monotonically increasing with respect to both c_1 and c_2 . Therefore, since $c_1 \geq \cos \theta$ and $c_2 \geq \cos \theta$, the absolute minimum of this expression occurs when $c_1 = \cos \theta$ and $c_2 = \cos \theta$. Substituting these minimum values for c_1, c_2 yields:

$$\begin{aligned} \langle \tilde{u}, \tilde{v} \rangle &\geq \cos^2 \theta - \sqrt{1 - \cos^2 \theta} \sqrt{1 - \cos^2 \theta} \\ \langle \tilde{u}, \tilde{v} \rangle &\geq \cos^2 \theta - \sin^2 \theta = \cos(2\theta) \end{aligned}$$

We are given that $\theta < \pi/6 - \delta/2$, which implies $2\theta < \pi/3 - \delta$. Because the cosine function is monotonically decreasing on $[0, \pi]$, we have:

$$\cos(2\theta) > \cos\left(\frac{\pi}{3} - \delta\right)$$

Consequently:

$$\frac{\langle u, v \rangle}{\|u\|_2 \|v\|_2} > \cos\left(\frac{\pi}{3} - \delta\right)$$

But by Lemma 3.2, for any two distinct vectors $u, v \in \Lambda \cap S_\epsilon$, the normalized inner product (which is equal to cosine of the angle between them) is bounded above:

$$\frac{\langle u, v \rangle}{\|u\|_2 \|v\|_2} \leq \cos\left(\frac{\pi}{3} - \delta\right)$$

This is a contradiction. Therefore, no such non-zero vector x exists, and $\mathcal{F}_{u,\theta} \cap \mathcal{F}_{v,\theta} = \emptyset$. $\square$

Figure 3 shows the disjoint conical regions centered around points $u, v \in S_\epsilon$.




Figure 4: Adversarial Convex Body

Definition 3.7 (Adversarial Convex Body). Given $u \in S_\epsilon$ and the stretching factor $\gamma > 1$ we define the adversarial body $\mathcal{B}_{u,\gamma}$ as the convex hull of $B \cup \{v, -v\}$:

$$\mathcal{B}_{u,\gamma} = \text{conv}(B \cup \{v, -v\}) \quad \text{where} \quad v = \gamma \lambda_1(\Lambda) \tilde{u}$$

where $B = \{x \in \mathbb{R}^n \mid \|x\|_2 \leq \lambda_1(\Lambda)\}$ denote the standard ℓ_2 ball of radius $\lambda_1(\Lambda)$ around origin and $\tilde{u} = u/\|u\|_2$ is a unit vector in the direction of u . Figure 4 gives visual representation of the adversarial convex body centered around point u .

We can see that the adversarial body is a constant-factor well-rounded perturbations of a Euclidean ball.

Lemma 3.8. Let $u \in S_\epsilon$, $\theta \in [0, \pi]$ and the stretching factor $\gamma < 1/\cos \theta$. Let B be a ball of radius $\lambda_1(\Lambda)$ around origin. Then any point $x \in \mathcal{B}_{u,\gamma}$ but not in B is contained in the union of symmetric conical regions defined by the angle θ :

$$\mathcal{B}_{u,\gamma} \setminus B \subset \mathcal{F}_{u,\theta} \cup \mathcal{F}_{-u,\theta}$$

Proof. Let $x \in \mathcal{B}_{u,\gamma} \setminus B$. The shaded region in Figure 3 indicates conical region around u , there will be such region around $-u$ symmetrically.

A general point of $\mathcal{B}_{u,\gamma} = \text{conv}(B \cup \{v, -v\})$ has the form $x = \alpha b + \beta v + \beta'(-v)$ with $b \in B$, $\alpha, \beta, \beta' \geq 0$ and $\alpha + \beta + \beta' = 1$. As $\mathcal{B}_{u,\gamma}$ and B are centrally symmetric, we may assume $\beta \geq \beta'$: otherwise we replace x by $-x$ and prove $-x \in \mathcal{F}_{u,\theta}$, which is equivalent to $x \in \mathcal{F}_{-u,\theta}$. Let $s = \beta - \beta' \geq 0$, so that $x = \alpha b + sv$ with $\alpha = 1 - \beta - \beta' \leq 1 - s$. If $s = 1$ then $x = v$, which lies in $\mathcal{F}_{u,\theta}$ trivially. If $s = 0$ then $x = \alpha b \in B$ (as $\alpha \leq 1$, $0 \in B$ and B is convex), contradicting $x \notin B$. So $s \in (0, 1)$, and as $\alpha/(1-s) \leq 1$, the point $b' = \frac{\alpha}{1-s}b$ lies in B , again by convexity of B with $0 \in B$. Renaming b' as b , we can assume that $x = (1-s)b + sv$ for some point $b \in B$ and scalar $s \in (0, 1)$. On substituting $v = \gamma \lambda_1 \tilde{u}$, we get

$$b = \frac{x - s\gamma \lambda_1 \tilde{u}}{1-s}$$

. As $b \in B$, we have $\|b\|_2^2 \leq \lambda_1^2$. Taking the squared norm and multiplying by $(1-s)^2$ gives:

$$\|x\|_2^2 - 2s\gamma\lambda_1\langle x \cdot \tilde{u} \rangle + s^2\gamma^2\lambda_1^2 \leq (1-s)^2\lambda_1^2$$

By rearranging the terms we get:

$$2s\gamma\lambda_1\langle x \cdot \tilde{u} \rangle \geq \|x\|_2^2 + s^2\gamma^2\lambda_1^2 - (1-s)^2\lambda_1^2 \quad (1)$$

Now we will show that the right hand side of the above equation is at least $2\lambda_1\|x\|_2$. To see this, let E is the difference:

$$E = \|x\|^2 - \lambda_1^2(1-s)^2 + s^2\gamma^2\lambda_1^2 - 2s\lambda_1\|x\|$$

By expanding and rearranging the terms we get:

$$E = \|x\|^2 - 2s\lambda_1\|x\| - \lambda_1^2 + 2s\lambda_1^2 - s^2\lambda_1^2 + s^2\gamma^2\lambda_1^2$$

Which implies

$$E = (\|x\|^2 - \lambda_1^2) - 2s\lambda_1(\|x\| - \lambda_1) + s^2\lambda_1^2(\gamma^2 - 1)$$

By taking out the common factor $(\|x\| - \lambda_1)$ from the first two terms:

$$E = (\|x\| - \lambda_1)(\|x\| + \lambda_1 - 2s\lambda_1) + s^2\lambda_1^2(\gamma^2 - 1)$$

Now we can see that all three terms in the above expression are greater than zero. Firstly, as $x \notin B$ we get $(\|x\| - \lambda_1) > 0$. Since $s < 1$ and $\|x\| > \lambda_1$ we get

$$\|x\| + \lambda_1 - 2s\lambda_1 > \lambda_1 + \lambda_1(1-2s) = 2\lambda_1(1-s) > 0$$

. Since $\gamma > 1$, the term $s^2\lambda_1^2(\gamma^2 - 1)$ is strictly positive. Which implies $E > 0$. So we have proved:

$$\|x\|^2 + s^2\gamma^2\lambda_1^2 - \lambda_1^2(1-s)^2 > 2s\lambda_1\|x\|$$

Hence from Equation 1 we get

$$2s\gamma\lambda_1\langle x \cdot \tilde{u} \rangle > 2s\lambda_1\|x\|$$

Since $\gamma < 1/\cos\theta$, we get

$$\frac{\langle x \cdot \tilde{u} \rangle}{\|x\|} > 1/\gamma > \cos\theta$$

Thus proving $x \in \mathcal{F}_{u,\theta}$. □

Lemma 3.9. *Let Λ be an integer lattice and B be a ball of radius $\lambda_1(\Lambda)$ around origin. For an angle $\theta \in (0, \pi)$ let $\gamma < 1/\cos\theta$ be the stretching factor. Then for any vector $x \notin \mathcal{F}_{u,\theta} \cup \mathcal{F}_{-u,\theta}$, its norm with respect to the adversarial body $\mathcal{B}_{u,\gamma}$ is exactly equal to its norm with respect to the original body B :*

$$\|x\|_{\mathcal{B}_{u,\gamma}} = \|x\|_B$$

Proof. By definition, the norm of vector x with respect to the two bodies are:

$$\|x\|_B = \inf\{t > 0 \mid x/t \in B\}$$

$$\|x\|_{\mathcal{B}_{u,\gamma}} = \inf\{t > 0 \mid x/t \in \mathcal{B}_{u,\gamma}\}$$

As the ball B is a subset of the adversarial body ($B \subseteq \mathcal{B}_{u,\gamma}$), for any scalar t if $x/t \in B$ then $x/t \in \mathcal{B}_{u,\gamma}$. So the infimum for $\mathcal{B}_{u,\gamma}$ is less than or equal to the infimum for B , which gives us:

$$\|x\|_{\mathcal{B}_{u,\gamma}} \leq \|x\|_B$$

Now, we will rule out a possibility that $\|x\|_{\mathcal{B}_{u,\gamma}} < \|x\|_B$, hence proving the equality. Suppose :

$$\|x\|_{\mathcal{B}_{u,\gamma}} < \|x\|_B$$

If this strict inequality holds, then by the definition of the infimum, there must exist some scalar t^* that satisfies:

$$\|x\|_{\mathcal{B}_{u,\gamma}} \leq t^* < \|x\|_B$$

such that the scaled point $y = x/t^*$ is contained in the adversarial body ($y \in \mathcal{B}_{u,\gamma}$) but not in the original body B :

$$y \in \mathcal{B}_{u,\gamma} \setminus B$$

As $\gamma < 1/\cos \theta$ by Lemma 3.8, we get:

$$y \in \mathcal{F}_{u,\theta} \cup \mathcal{F}_{-u,\theta}$$

By the definition of conical regions, this means the normalized inner product of y satisfies the cone inequality. Because substituting $y = x/t^*$ yields:

$$\frac{\langle (x/t^*) \cdot \tilde{u} \rangle}{\|x/t^*\|_2} = \frac{\frac{1}{t^*} \langle x \cdot \tilde{u} \rangle}{\frac{1}{t^*} \|x\|_2} = \frac{\langle x \cdot \tilde{u} \rangle}{\|x\|_2} \geq \cos \theta$$

This implies $x \in \mathcal{F}_{u,\theta} \cup \mathcal{F}_{-u,\theta}$, a contradiction. As the strict inequality leads to a contradiction, we conclude that $\|x\|_{\mathcal{B}_{u,\gamma}} = \|x\|_B$ $\square$

4 Adversary Lower bounds for SVP and CVP

4.1 Lower bound for deterministic algorithms for SVP

Theorem 4.1 (Deterministic Lower Bound for SVP). *There exists a polynomially encoded lattice family $\{\Lambda\}_n$ of rank n (i.e. description of Λ has $\text{poly}(n)$ bit size) such that, any deterministic algorithm $\mathcal{A}$ that solves the Shortest Vector Problem (SVP) on Λ within an approximation factor of $\gamma' < 2/\sqrt{3}$ using a black-box metric oracle must make at least $2^{\Omega(n)}$ queries to the oracle.*

Proof. We choose stretching factor γ as a constant such that $\gamma' < \gamma < 2/\sqrt{3}$. So $1/\gamma > \sqrt{3}/2 = \cos(\pi/6)$. Due to continuity of cosine function and the fact that it is monotonically decreasing in the interval $[0, \pi/2]$, there exists a non-zero constant δ such that $1/\gamma > \cos(\pi/6 - \delta/2)$. We choose angle θ such that $\cos \theta$ is in the open interval $(\cos(\pi/6 - \delta/2), 1/\gamma)$. So we have $\arccos(1/\gamma) < \theta < \pi/6 - \delta/2$ or equivalently

$$1/\gamma > \cos \theta > \cos(\pi/6 - \delta/2)$$

We choose $\epsilon < \min\{\epsilon(\delta), \frac{\gamma - \gamma'}{\gamma'}, 1\}$, where $\epsilon(\delta)$ is the constant given by Lemma 3.2.

Let Λ be the *real* lattice obtained by Theorem 3.1. Using Theorem 3.3 we can without loss of generality assume that Λ is a *polynomial sized integer lattice* with $2^{\Omega(n)}$ vectors in the ϵ -shell

$$S_\epsilon = \{x \in \mathbb{R}^n \mid \lambda_1(\Lambda) \leq \|x\|_2 \leq \lambda_1(\Lambda) \cdot (1 + \epsilon)\}$$

. Let $U = \Lambda \cap S_\epsilon$ be the set of target lattice vectors located in the ϵ -shell S_ϵ . We have, $|U| = 2^{\Omega(n)}$. For any pair of non-collinear points in U , the angle between them is at least $\pi/3 - \delta$. Note also that, as $\epsilon < 1$, any two distinct non-antipodal points of U are automatically non-collinear: for any $w \in \Lambda \setminus \{0\}$, the set $\Lambda \cap \mathbb{R}w$ equals $\mathbb{Z}w_0$ for some $w_0 \in \Lambda$ with $\|w_0\|_2 \geq \lambda_1$, so two collinear points of the shell are either equal, antipodal, or one of them has norm at least $2\lambda_1 > (1 + \epsilon)\lambda_1$, which is impossible in S_ϵ . Let B be ball of radius $\lambda_1(\Lambda)$ around origin.

We will give a proof by contradiction. Suppose there is a deterministic algorithm $\mathcal{A}$ that solves SVP within factor γ' by making $N = 2^{o(n)}$ queries to the metric oracle. During the execution suppose $\mathcal{A}$ makes a fixed set of queries $Q = \{q_1, q_2, \dots, q_N\}$ and outputs a non-zero lattice vector y . Oracle answers all queries w. r. t. convex body B as a metric, that is on query q_i oracle returns $\|q_i\|_B$ for $i \in [N]$. For $x \in U$, let $\mathcal{C}_{[x],\theta} = \mathcal{F}_{x,\theta} \cup \mathcal{F}_{-x,\theta}$. Let $x, y \in U$ be non antipodal points ($y \neq -x$). As $\theta < \pi/6 - \delta/2$, by applying Lemma 3.6 to x, y and $x, -y$ we get:

$$\mathcal{F}_{x,\theta} \cap \mathcal{F}_{y,\theta} = \emptyset \text{ and } \mathcal{F}_{x,\theta} \cap \mathcal{F}_{-y,\theta} = \emptyset$$

As any point $u \in \mathcal{F}_{x,\theta}$ iff $-u \in \mathcal{F}_{-x,\theta}$, we symmetrically get:

$$\mathcal{F}_{-x,\theta} \cap \mathcal{F}_{-y,\theta} = \emptyset \text{ and } \mathcal{F}_{-x,\theta} \cap \mathcal{F}_{y,\theta} = \emptyset$$

Consequently we get that intersection of double cones centered at x and y is empty:

$$\mathcal{C}_{[x],\theta} \cap \mathcal{C}_{[y],\theta} = \emptyset$$

Now $N = 2^{o(n)}$ which implies $N + 1 = |Q \cup \{y\}| = 2^{o(n)}$. We have $|U| = 2^{\Omega(n)}$, which implies $|U|/2 = 2^{\Omega(n)}$. This implies there is $u^* \in U$ such that:

$$(Q \cup \{y\}) \cap (\mathcal{F}_{u^*,\theta} \cup \mathcal{F}_{-u^*,\theta}) = \emptyset$$

. The adversary now sets the true metric to be the metric given by the adversarial body $\mathcal{B}_{u^*, \gamma}$. Since $1/\gamma > \cos(\pi/6 - \delta/2)$ and none of the queried vectors are in $\mathcal{F}_{u^*, \theta} \cup \mathcal{F}_{-u^*, \theta}$ by Lemma 3.9 we have

$$\|q_i\|_{\mathcal{B}_{u^*, \gamma}} = \|q_i\|_B \text{ for } i \in [N] \quad \text{and} \quad \|y\|_{\mathcal{B}_{u^*, \gamma}} = \|y\|_B$$

So the algorithm $\mathcal{A}$ can not distinguish whether the queries are answered with respect to B or $\mathcal{B}_{u^*, \gamma}$. With respect to the adversarial metric $\mathcal{B}_{u^*, \gamma}$, the boundary in the direction of u^* is expanded by the stretching factor γ :

$$\|u^*\|_{\mathcal{B}_{u^*, \gamma}} = \frac{\|u^*\|_B}{\gamma}$$

Indeed, the linear functional $x \mapsto \langle x, \tilde{u}^* \rangle$, where $\tilde{u}^* = u^*/\|u^*\|_2$, attains its maximum over $\mathcal{B}_{u^*, \gamma} = \text{conv}(B \cup \{v, -v\})$ only at the point $v = \gamma\lambda_1\tilde{u}^*$, since its maximum over B is $\lambda_1 < \gamma\lambda_1$. Hence the ray from the origin through u^* exits $\mathcal{B}_{u^*, \gamma}$ exactly at v , and by positive homogeneity of the gauge every point x on this ray satisfies $\|x\|_{\mathcal{B}_{u^*, \gamma}} = \|x\|_2/(\gamma\lambda_1) = \|x\|_B/\gamma$.

As $u^* \in S_\epsilon$, we have $\|u^*\|_B \leq 1 + \epsilon$ we get:

$$\|u^*\|_{\mathcal{B}_{u^*, \gamma}} \leq \frac{1 + \epsilon}{\gamma}$$

As the output vector y is a lattice point, we have $\|y\|_B \geq 1$. As $\|y\|_{\mathcal{B}_{u^*, \gamma}} = \|y\|_B$ we get $\|y\|_{\mathcal{B}_{u^*, \gamma}} \geq 1$.

The approximation factor achieved by the algorithm is the ratio between the norm of its output y and the norm of the shortest vector u^* with respect to $\mathcal{B}_{u^*, \gamma}$:

$$\text{Approximation ratio} = \frac{\|y\|_{\mathcal{B}_{u^*, \gamma}}}{\|u^*\|_{\mathcal{B}_{u^*, \gamma}}} \geq \frac{1}{\frac{1+\epsilon}{\gamma}} = \frac{\gamma}{1+\epsilon} > \gamma'$$

The last inequality follows from our choice $\epsilon < \frac{\gamma - \gamma'}{\gamma'}$. This is a contradiction as algorithm fails to output a vector with norm within γ' factor of the norm of the shortest vector. □

4.2 Lower bound for deterministic algorithms for CVP

Theorem 4.2 (Deterministic Lower Bound for CVP). *There exists a polynomially encoded family of CVP instances $\{(\Lambda_{n,i}, t_{n,i}) \mid i \in [n]\}_{n \geq 1}$, with each $\Lambda_{n,i}$ a lattice of rank n and each instance of $\text{poly}(n)$ bit size, such that any deterministic algorithm $\mathcal{A}$ which, using a black-box metric oracle, solves each of these CVP instances within an approximation factor of $\gamma' < 2/\sqrt{3}$ must make at least $2^{\Omega(n)}$ queries to the oracle on some instance.*

Proof. We give a proof by contradiction, using well known polynomial-time Turing reduction from SVP to CVP [19]. First we recall the reduction from SVP to CVP from [19]. Given a lattice Λ with basis $\{b_1, b_2, \dots, b_n\}$, the reduction constructs n sub-lattices. For each $i \in \{1, \dots, n\}$, it defines the lattice Λ_i as the lattice generated by $\{b_1, \dots, b_{i-1}, 2b_i, b_{i+1}, \dots, b_n\}$. The basis for Λ_i is obtained from the basis of Λ by replacing b_i by $2b_i$. The algorithm $\mathcal{A}_{\text{SVP}}$ then calls the CVP algorithm exactly n times, asking it to find the closest vector in Λ_i to the target point $t_i = b_i$. Note that

$$b_i - \Lambda_i = \{Bc \mid c \in \mathbb{Z}^n, c_i \text{ is odd}\} \subseteq \Lambda \setminus \{0\}$$

The shortest of the resulting difference vectors is guaranteed to solve the γ' -approximate SVP for the original lattice Λ if $\mathcal{A}_{\text{CVP}}$ solves γ' -approximate CVP.

The correctness of the reduction relies on parity of the lattice coefficients to isolate the shortest vector into a shifted sub-lattice. Crucially, it depends only on the translation-invariance of the distance function and does not invoke inner products, angles, or any specific geometric properties of the norm. Therefore, the reduction works for arbitrary norms given by gauge functions. Also the reduction doesn't need to see the metric explicitly and works perfectly in the black-box set up (accessing metric through a black-box oracle).

The instance family of the statement is $(\Lambda_{n,i}, t_{n,i}) = (\Lambda_i, b_i)$ for $i \in [n]$, where Λ is the adversarial lattice family of Theorem 4.1 with basis $\{b_1, \dots, b_n\}$. So clearly, if we have a deterministic algorithm to solve CVP on these instances within factor γ' which accesses metric through black-box oracle and makes at most $2^{o(n)}$ queries to the oracle per instance, then we have similar algorithm for SVP which makes at most $n \cdot 2^{o(n)} + n = 2^{o(n)}$ oracle queries (the additive n accounts for selecting the shortest of the n difference vectors), a contradiction by Theorem 4.1. Thus proving the desired. $\square$

4.3 Lower Bound against Randomized Algorithms (via Yao's Minimax Principle)

In this Section we use Yao's Minmax principle to extend our lower bounds to randomized algorithms. The principle states that the expected cost of any randomized algorithm on the worst-case input is lower-bounded by the expected cost of the best deterministic algorithm on a chosen probability distribution over the inputs. As a result, to prove that any randomized algorithm must fail with high probability using few queries, it suffices to define a distribution of inputs and prove that any deterministic algorithm fails with high probability over this distribution.

Theorem 4.3 (Randomized Lower Bound for SVP). *There exists a polynomially encoded lattice family $\{\Lambda\}_n$ of rank n (i.e. description of Λ has $\text{poly}(n)$ bit size) such that the following holds: if a randomized algorithm $\mathcal{R}$ which accesses the metric through*

a black-box oracle solves the Shortest Vector Problem (SVP) on Λ within an approximation factor of $\gamma' < 2/\sqrt{3}$ with constant success probability (say, with probability $\geq 2/3$), then the expected number of oracle queries $\mathcal{R}$ makes must be at least $2^{\Omega(n)}$.

Proof. We choose lattice Λ , target vector set U , constants γ, ϵ , convex bodies B and $\mathcal{B}_{u,\gamma}$ for $u \in U$ exactly as in the proof of Theorem 4.1. As $u \in S_\epsilon$ iff $-u \in S_\epsilon$ we define $V = U/\{1, -1\}$, a set of points obtained from U by identifying $u \in U$ with $-u$. That is V is a set of mutually non antipodal points from U .

We define a probability distribution $\mathcal{D}$ over the the collection of convex bodies $\{\mathcal{B}_{u,\gamma} | u \in V\}$. To construct an instance from $\mathcal{D}$, we draw a target vector u^* uniformly at random from the set V . We set the true metric of the oracle to be the corresponding adversarial body $\mathcal{B}_{u^*,\gamma}$. Let $\mathcal{A}$ be any deterministic algorithm that makes at most $N = 2^{o(n)}$ queries to the metric oracle. We analyze the probability that $\mathcal{A}$ successfully achieves a γ' -approximation when the metric is drawn from $\mathcal{D}$. As $\mathcal{A}$ is deterministic, its sequence of queries is entirely fixed by the oracle's responses. Run $\mathcal{A}$ against the oracle for the unperturbed ball B ; this yields a fixed query set $Q = \{q_1, q_2, \dots, q_N\}$ and a fixed output y .

Using arguments similar in the proof of Theorem 4.1 based on Lemma 3.6, we get the double conical regions $\mathcal{C}_{[u],\theta} = \mathcal{F}_{u,\theta} \cup \mathcal{F}_{-u,\theta}$ for all $u \in V$ are mutually disjoint. Therefore, any single query $q_i \in Q$ can fall into the perturbed double cone of at most one target vector in V . Consequently, the entire set of N queries can intersect the cones of at most N distinct target vectors. Because the true target u^* is chosen uniformly at random from V , the probability that any of the algorithm's queries hit the perturbed cones of u^* is bounded by:

$$\Pr[Q \cap \mathcal{C}_{[u^*],\theta} \neq \emptyset] \leq \frac{N}{|V|} = \frac{2^{o(n)}}{2^{\Omega(n)}} = 2^{-\Omega(n)}$$

So with probability $1 - 2^{-\Omega(n)}$, none of the queries in Q fall inside the cones of u^* . On this event, the run of $\mathcal{A}$ against the true oracle for $\mathcal{B}_{u^*,\gamma}$ produces exactly the same transcript, and hence the same output y : by induction on the query index, if the first $i-1$ answers agree then the i -th query is the same vector q_i in both runs, and by Lemma 3.9 (as $q_i \notin \mathcal{C}_{[u^*],\theta}$) the answers to it agree as well:

$$\|q_i\|_{\mathcal{B}_{u^*,\gamma}} = \|q_i\|_B$$

Moreover, if the output y (a non-zero lattice vector) also lies outside $\mathcal{C}_{[u^*],\theta}$, then the algorithm fails: $\|y\|_{\mathcal{B}_{u^*,\gamma}} = \|y\|_B \geq 1$, while $\gamma' \cdot \lambda_1(\Lambda, \mathcal{B}_{u^*,\gamma}) \leq \gamma' \cdot \|u^*\|_{\mathcal{B}_{u^*,\gamma}} \leq \gamma'(1+\epsilon)/\gamma < 1$ by the choice $\epsilon < \frac{\gamma-\gamma'}{\gamma'}$, so y is not a γ' -approximate shortest vector. As the double cones are mutually disjoint, the probability that the fixed output y falls into the double cone of the randomly chosen u^* is at most $1/|V| = 2^{-\Omega(n)}$. So, the total probability that the deterministic algorithm $\mathcal{A}$ succeeds either by successfully querying the cone or by blindly guessing a valid output vector is bounded by the union bound:

$$\Pr[\mathcal{A} \text{ succeeds}] \leq \frac{N}{|V|} + \frac{1}{|V|} = \frac{N+1}{|V|} = 2^{-\Omega(n)}$$

Now let $\mathcal{R}$ be a randomized algorithm that, on every input metric $\mathcal{B}_{u,\gamma}$ with $u \in V$, outputs a γ' -approximate shortest vector with probability at least $2/3$, and let T denote the maximum, over these inputs, of its expected number of oracle queries. Truncate $\mathcal{R}$ after $18T$ queries: by Markov's inequality the truncation aborts any run with probability at most $1/18$, so the truncated algorithm $\mathcal{R}'$ always makes at most $18T$ queries and succeeds with probability at least $2/3 - 1/18 > 1/2$ on every input. Draw u^* uniformly at random from V ; averaging over u^* and over the internal coins of $\mathcal{R}'$, the success probability of $\mathcal{R}'$ exceeds $1/2$, so some fixing of the coins yields a deterministic algorithm making at most $N = 18T$ queries whose success probability over $u^* \sim \text{Uniform}(V)$ exceeds $1/2$. By the argument above (which is valid for any N), that success probability is at most $(N+1)/|V|$, so $1/2 < (18T+1)/|V|$ and hence $T \geq (|V| - 2)/36 = 2^{\Omega(n)}$: any randomized algorithm which solves γ' -approximate SVP with success probability at least $2/3$ must make at least $2^{\Omega(n)}$ expected queries. Thus proving the desired. $\square$

Corollary 4.4 (Randomized Lower Bound for CVP). *There exists a polynomially encoded family of CVP instances $\{(\Lambda_{n,i}, t_{n,i}) \mid i \in [n]\}_{n \geq 1}$, with each $\Lambda_{n,i}$ a lattice of rank n and each instance of $\text{poly}(n)$ bit size, such that the following holds: if a randomized algorithm $\mathcal{R}$ which accesses the metric through a black-box oracle solves each of these CVP instances within an approximation factor of $\gamma' < 2/\sqrt{3}$ with constant success probability (say, with probability $\geq 2/3$), then the expected number of oracle queries $\mathcal{R}$ makes on some instance must be at least $2^{\Omega(n)}$.*

Proof. We give a proof by contradiction, using the well-known Turing reduction from SVP to CVP [19]. The instance family of the statement is $(\Lambda_{n,i}, t_{n,i}) = (\Lambda_i, b_i)$ for $i \in [n]$, as in Theorem 4.2. Let $\mathcal{A}_{CVP}$ be a randomized algorithm that solves γ' -approximate CVP on these instances with a success probability of at least $2/3$, making an expected $Q = 2^{o(n)}$ queries to the metric oracle per instance. The reduction to SVP requires invoking the CVP algorithm exactly n times on different shifted sub-lattices. To ensure the overall SVP reduction succeeds with a high constant probability, we must formally bound the probability of failure across all n independent calls. First, we amplify the success probability of $\mathcal{A}_{CVP}$. By running $\mathcal{A}_{CVP}$ independently $k = \Theta(\log n)$ times for a single instance and selecting the valid output vector closest to the target (the one which has smallest distance to the target), we reduce the failure probability of a single CVP invocation from $1/3$ to at most $\frac{1}{3n}$. This amplification increases the expected queries per CVP instance to $\Theta(\log n) \cdot Q$ (we may assume $Q \geq 1$), including the $\Theta(\log n)$ oracle evaluations needed to select the closest valid output. Now we apply the union bound. The probability that any of the n amplified CVP calls fails is bounded by:

$$\Pr[\text{failure}] \leq n \cdot \left(\frac{1}{3n} \right) = \frac{1}{3}$$

Consequently, this yields a randomized algorithm for SVP that succeeds with probability at least $1 - 1/3 = 2/3$. The total expected number of queries N made by this resulting SVP algorithm to the metric oracle is:

$$N = n \cdot \Theta(\log n) \cdot Q = \Theta(n \log n) \cdot 2^{o(n)} = 2^{o(n)}$$

This is a direct contradiction with Theorem 4.3. As from Theorem 4.3 we know that any randomized algorithm solving γ' -approximate SVP with success probability $\geq 2/3$ must make at least $2^{\Omega(n)}$ expected queries. So, the expected query complexity Q of the original randomized CVP algorithm can not be $2^{o(n)}$. $\square$

4.4 Informal discussion: tightness of the Bound for the ℓ_2 Metric

If the adversary chooses the standard Euclidean ℓ_2 ball as the base metric C , the stretch factor cannot exceed $1/\cos(\theta)$, where θ is the half-angle of the assigned cone. This is not a limitation of our proof technique, but a strict geometric property of the ℓ_2 norm. Suppose the adversary attempts to construct some updated, O -symmetric convex body C' to embed the target vector $v = \gamma \lambda_1 u$. To preserve the spatial disjointness required for the lower bound, C' must identically match the base ball C everywhere outside the assigned cone $\mathcal{F}_{u,\theta}$. Because C' is a convex body containing both the unperturbed base ball C and the new stretched peaks $\{v, -v\}$, it must strictly contain their minimal convex hull: $\text{conv}(C \cup \{v, -v\}) \subseteq C'$. Because the boundary of the ℓ_2 ball C is uniformly curved, this minimal convex hull is defined by the exact tangent lines drawn from the peak v to the surface of the sphere. These tangent lines form a right triangle with the origin and the point of tangency. By standard trigonometry, the angle this tangential boundary makes with the central axis u is exactly $\arccos(1/\gamma)$. If the adversary attempts to increase the stretch factor strictly beyond $\gamma = 1/\cos(\theta)$, this tangent angle becomes greater than θ . This physically forces the outer boundary of the minimal convex hull to cross outside the allowed cone $\mathcal{F}_{u,\theta}$. Since C' must contain this hull, C' inevitably leaves the region as well. Once the perturbation leaves this region, it intersects with the cones assigned to other target vectors, destroying the spatial disjointness required for the lower bound. Therefore, for algorithms operating over the ℓ_2 metric, nothing better than the $1/\cos(\theta)$ limit is geometrically possible within this black-box framework. We will make these informal observations more precise in the full version of the paper.

5 Concluding Remarks and Open Problems

In this paper, we proved unconditional $2^{\Omega(n)}$ query complexity lower bounds on any randomized algorithm for SVP (respect. CVP) which accesses the metric through black-box oracle and solves SVP (respect. CVP) within factor $\gamma < 2/\sqrt{3}$. We observed that if the base metric is the ℓ_2 ball, its curvature possibly limits the stretch factor to $\gamma \leq 1/\cos(\theta)$. This raises a natural question: can the lower bound technique

achieve a proof for larger approximation factors by using a different base body? Bodies with flat faces, like the ℓ_∞ polytope, theoretically allow for a much larger local stretch without leaving a single assigned cone. However, the flat-faced bodies have extreme corners. When the convex hull wraps around the stretched peak and these corners, it inevitably intersects neighboring cones, destroying the disjoint packing required for the lower bound. So metric like ℓ_∞ might not give any better results. We will explore this further in the full version of the paper. We leave following open question for future work:

What is the maximum achievable stretch factor γ for an arbitrary o -symmetric convex body? Specifically, what is the exact mathematical trade-off between a body having flat faces (to maximize local stretch) and remaining globally round (to successfully pack $2^{\Omega(n)}$ disjoint target cones)?

Historically, it is believed that CVP is a harder computational problem than SVP. But our technique doesn't prove any better lower bound result for CVP compared to SVP. The obvious reason for this is we use Turing reduction from SVP to CVP to prove the lower bound for CVP using the lower bound for SVP. It raises a natural question: can we prove better lower bound for CVP (which works for larger approximation factors)? Of course such a lower-bound proof must directly construct adversarial body for CVP and the argument shouldn't be via SVP to CVP Turing reduction.

6 Acknowledgments

The authors acknowledge the use of Google Gemini Pro for assistance in refining the exposition and streamlining the structure of the manuscript. We thank Anthropic's Claude (Claude Fable 5, used through Claude Code) for discussions during preparation of the proof of Theorem 3.3. All core ideas, mathematical formulations, and theoretical contributions are the original work of the human authors. The authors take full responsibility for the correctness and final content of this paper.

References

- [1] Divesh Aggarwal and Noah Stephens-Davidowitz. (Gap/S)ETH hardness of SVP. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 228–238, 2018.
- [2] Miklós Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In Jeffrey Scott Vitter, editor, *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*, pages 10–19. ACM, 1998. doi:10.1145/276698.276705.

- [3] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the shortest lattice vector problem. In Jeffrey Scott Vitter, Paul G. Spirakis, and Mihalis Yannakakis, editors, *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, July 6-8, 2001, Heraklion, Crete, Greece*, pages 601–610. ACM, 2001. doi:10.1145/380752.380857.
- [4] Miklós Ajtai, Ravi Kumar, and D. Sivakumar. Sampling short lattice vectors and the closest lattice vector problem. In *Proceedings of the 17th Annual IEEE Conference on Computational Complexity, Montréal, Québec, Canada, May 21-24, 2002*, pages 53–57. IEEE Computer Society, 2002. doi:10.1109/CCC.2002.1004339.
- [5] Sanjeev Arora, László Babai, Jacques Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *Journal of Computer and System Sciences*, 54(2):317–331, 1997.
- [6] V. Arvind and Pushkar S. Joglekar. Lattice problems, gauge functions and parameterized algorithms, 2008. URL: <https://arxiv.org/abs/0804.4744>, arXiv:0804.4744.
- [7] Vikraman Arvind and Pushkar S. Joglekar. Some sieving algorithms for lattice problems. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2008, December 9-11, 2008, Bangalore, India*, volume 2 of *LIPIcs*, pages 25–36. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2008. doi:10.4230/LIPIcs.FSTTCS.2008.1738.
- [8] László Babai. On Lovász’ lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986.
- [9] Eric Stephen Barnes and Gordon Elliott Wall. Some extreme forms defined in terms of abelian groups. *Journal of the Australian Mathematical Society*, 1(1):47–63, 1959.
- [10] Anja Becker, Nicolas Gama, and Antoine Joux. A sieve algorithm based on overlattices. *LMS Journal of Computation and Mathematics*, 17(A):49–70, 2014. Special issue for the Proceedings of the 11th Algorithmic Number Theory Symposium (ANTS XI).
- [11] Huck Bennett, Alexander Golovnev, and Noah Stephens-Davidowitz. Fine-grained hardness of CVP(p)—everything that ETH implies. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 582–593. IEEE, 2017.

- [12] Huck Bennett, Alexander Golovnev, and Noah Stephens-Davidowitz. Difficulties constructing lattices with exponential kissing number from codes. *IEEE Transactions on Information Theory*, 71(10):7644–7648, 2025.
- [13] Johannes Blömer. Closest vectors, successive minima, and dual hkz-bases of lattices. In Ugo Montanari, José D. P. Rolim, and Emo Welzl, editors, *Automata, Languages and Programming, 27th International Colloquium, ICALP 2000, Geneva, Switzerland, July 9-15, 2000, Proceedings*, volume 1853 of *Lecture Notes in Computer Science*, pages 248–259. Springer, 2000. doi:10.1007/3-540-45022-X_22.
- [14] Johannes Blömer and Stefanie Naewe. Sampling methods for shortest vectors, closest vectors and successive minima. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *Automata, Languages and Programming, 34th International Colloquium, ICALP 2007, Wrocław, Poland, July 9-13, 2007, Proceedings*, volume 4596 of *Lecture Notes in Computer Science*, pages 65–77. Springer, 2007. doi:10.1007/978-3-540-73420-8_8.
- [15] John William Scott Cassels. *An Introduction to the Geometry of Numbers*. Classics in Mathematics. Springer Science & Business Media, 1997.
- [16] John Horton Conway and Neil J. A. Sloane. *Sphere Packings, Lattices and Groups*, volume 290. Springer-Verlag, New York, 3rd edition, 1999.
- [17] Daniel Dadush. A randomized sieving algorithm for approximate integer programming. *Algorithmica*, 70(2):208–244, 2014. URL: <https://doi.org/10.1007/s00453-013-9834-8>, doi:10.1007/s00453-013-9834-8.
- [18] Irit Dinur, Guy Kindler, and Shmuel Safra. Approximating-cvp to within almost-polynomial factors is np-hard. In *39th Annual Symposium on Foundations of Computer Science, FOCS '98, November 8-11, 1998, Palo Alto, California, USA*, pages 99–111. IEEE Computer Society, 1998. doi:10.1109/SFCS.1998.743433.
- [19] Oded Goldreich, Daniele Micciancio, Shmuel Safra, and Jean-Paul Seifert. Approximating shortest lattice vectors is not harder than approximating closest lattice vectors. *Information Processing Letters*, 71(2):55–61, 1999.
- [20] Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. In David S. Johnson and Uriel Feige, editors, *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 469–477. ACM, 2007. doi:10.1145/1250790.1250859.

- [21] Ravi Kannan. Minkowski's convex body theorem and integer programming. *Math. Oper. Res.*, 12(3):415–440, 1987. doi:10.1287/moor.12.3.415.
- [22] Lenstra H.W. Lenstra A. K. and Lovasz. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(1):515–534, 1982.
- [23] Daniele Micciancio and Shafi Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671. Springer Science & Business Media, 2002.
- [24] Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. *SIAM Journal on Computing*, 42(3):1364–1391, 2013.
- [25] Phong Q. Nguyen and Thomas Vidick. Sieve algorithms for the shortest vector problem are practical. *Journal of Mathematical Cryptology*, 2(2):181–207, 2008.
- [26] Chris Peikert. A decade of lattice cryptography. *Foundations and trends® in theoretical computer science*, 10(4):283–424, 2016.
- [27] Claude Ambrose Rogers. Mean values over the space of lattices. *Acta Mathematica*, 94(1):249–287, 1955.
- [28] Thomas Rothvoss and Moritz Venzin. Approximate CVP in time $2^{0.802n}$ – now in any norm! In *Integer Programming and Combinatorial Optimization (IPCO 2022)*, volume 13265 of *Lecture Notes in Computer Science*, pages 440–453. Springer, 2022. doi:10.1007/978-3-031-06901-7_32.
- [29] Wolfgang M. Schmidt. A metrical theorem in geometry of numbers. *Transactions of the American Mathematical Society*, 95(3):516–529, 1960.
- [30] Victor Shoup. Lower bounds for discrete logarithms and related problems. In *International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT)*, pages 256–266. Springer, Berlin, Heidelberg, 1997. doi:10.1007/3-540-69053-0_18.
- [31] C. L. Siegel. *Lectures on Geometry of Numbers*. Springer-Verlag publishing company. Springer, 1988.
- [32] Carl Ludwig Siegel. A mean value theorem in geometry of numbers. *Annals of Mathematics*, 46(2):340–347, 1945.
- [33] van Emde Boas. Another np-complete problem and the complexity of computing short vectors in a lattice. *Technical Report, University of Amsterdam.*, 81(04), 1981.

- [34] Serge Vlăduț. Lattices with exponentially large kissing numbers. *Moscow Journal of Combinatorics and Number Theory*, 8(2):163–177, 2019. doi: 10.2140/moscow.2019.8.163.
- [35] Serge Vlăduț. Lattices with exponentially large kissing numbers do exist, 2025. URL: <https://arxiv.org/abs/2411.07371>, arXiv:2411.07371.
- [36] Yutong Zhang and Yaoran Yang. A uniform random-lattice tail bound for the SVP kissing-profile parameter, 2026. arXiv:2605.21966.

7 Appendix

Proof of Theorem 3.3

Proof. All auxiliary precision is chosen inside the proof. Define

$$a_\delta := \sqrt{\frac{1}{2} + \cos\left(\frac{\pi}{3} - \delta\right)} - 1, \quad \eta := \frac{1}{2} \min\{1, \varepsilon, a_\delta\}. \quad (2)$$

Then $0 < \eta < \min\{1, \varepsilon, a_\delta\}$ and hence

$$\frac{1}{2} + 2\eta + \eta^2 < \cos\left(\frac{\pi}{3} - \delta\right). \quad (3)$$

The real witness. Let $\kappa_n := \text{vol}(B_2^n)$ and choose r_n so that $\kappa_n r_n^n = 1/3$. Now by choosing $R = r_n$ and shell width $\eta/2$ in the Proof of Theorem 2.6 WE GET A unit-covolume lattice Λ for which

$$\lambda_1(\Lambda) \geq r_n \quad (4)$$

and a symmetric set

$$U_0 \subseteq \Lambda \cap S_{\eta/2}(\Lambda), \quad |U_0| \geq \frac{(1 + \eta/2)^n - 1}{6}. \quad (5)$$

Indeed, the inner ball has volume $1/3$ and is empty with probability at least $2/3$ by Siegel's mean-value theorem and Markov's inequality. The surrounding annular region has volume $((1 + \eta/2)^n - 1)/3$, and the Rogers–Schmidt second-moment bound and Chebyshev's inequality give at least half that many points with probability $1 - o(1)$. The events intersect for large n . Retaining (4), rather than using only the relative-shell conclusion, is what makes the following conditioning bound possible.

A reduced basis with quantitative conditioning. Let $\lambda_1 \leq \dots \leq \lambda_n$ be the successive minima of Λ . Minkowski's second theorem gives

$$\kappa_n \prod_{i=1}^n \lambda_i \leq 2^n.$$

Together with (4) and $\kappa_n r_n^n = 1/3$, this yields

$$\lambda_n \leq \frac{2^n}{\kappa_n r_n^{n-1}} = 3 \cdot 2^n r_n. \quad (6)$$

The inclusion $[-1/\sqrt{n}, 1/\sqrt{n}]^n \subset B_2^n$ gives $\kappa_n \geq (2/\sqrt{n})^n$ and therefore $r_n < \sqrt{n}/2$.

Choose a size-reduced Korkine–Zolotarev basis $b_1, \dots, b_n$ of Λ . If b_i^* are its Gram–Schmidt vectors, then

$$\|b_i^*\|_2 \leq \lambda_i, \quad \|b_i\|_2^2 \leq \left(1 + \frac{i-1}{4}\right) \lambda_i^2. \quad (7)$$

For clarity, the first inequality follows because among i independent lattice vectors of length at most λ_i , one has nonzero projection orthogonal to $b_1, \dots, b_{i-1}$, whereas b_i^* is shortest in that projected lattice. The second follows from size reduction $|\mu_{ij}| \leq 1/2$ and $\|b_j^*\|_2 \leq \lambda_j \leq \lambda_i$. This also recalls the complete standard short-basis estimate being used here.

For the basis matrix $G = (b_1 | \dots | b_n)$, equations (6) and (7) imply

$$\|b_i\|_2 \leq \sqrt{n} \lambda_n < 3 \cdot 2^{n-1} n,$$

and hence

$$\|G\|_{2 \rightarrow 2} \leq \|G\|_F \leq H_n := 2^{n+1} n^{3/2}. \quad (8)$$

Since $|\det G| = 1$, the product of the singular values is one, whence

$$\|G^{-1}\|_{2 \rightarrow 2} \leq H_n^{n-1} = 2^{O(n^2)}. \quad (9)$$

Scale, round, and control every lattice vector. Choose the single integer scaling factor

$$Q := 1 + \left\lfloor \frac{4n \|G^{-1}\|_{2 \rightarrow 2}}{\eta} \right\rfloor, \quad A_{ij} := \lfloor Q G_{ij} \rfloor. \quad (10)$$

Put $\tilde{G} := A/Q$ and $T := \tilde{G}G^{-1}$. Nearest-integer rounding gives

$$\tau := \|T - I\|_{2 \rightarrow 2} \leq \frac{n}{2Q} \|G^{-1}\|_{2 \rightarrow 2} < \frac{\eta}{8} < 1. \quad (11)$$

Therefore, for every $z \in R^n$,

$$(1 - \tau) \|Gz\|_2 \leq \|\tilde{G}z\|_2 \leq (1 + \tau) \|Gz\|_2. \quad (12)$$

This uniform estimate makes A nonsingular and, after taking the minimum over $z \in \mathbb{Z}^n \setminus \{0\}$, gives

$$Q(1 - \tau) \lambda_1(\Lambda) \leq \lambda_1(A\mathbb{Z}^n) \leq Q(1 + \tau) \lambda_1(\Lambda). \quad (13)$$

Let $Z_0 := \{z \in \mathbb{Z}^n : Gz \in U_0\}$ and set $U := \{Az : z \in Z_0\}$. Nonsingularity preserves symmetry and cardinality. For $z \in Z_0$, equations (5), (12), and (13) imply

$$\frac{\|Az\|_2}{\lambda_1(A\mathbb{Z}^n)} \leq \frac{1 + \tau}{1 - \tau} \left(1 + \frac{\eta}{2}\right) < 1 + \eta < 1 + \varepsilon. \quad (14)$$

The strict middle inequality follows from $\tau < \eta/8$ and

$$(1 - \eta/8)(1 + \eta) - (1 + \eta/8)(1 + \eta/2) = \frac{\eta(4 - 3\eta)}{16} > 0.$$

Thus $U \subseteq S_\eta(A\mathbb{Z}^n) \subseteq S_\varepsilon(A\mathbb{Z}^n)$.

Count and angles. Write $\lambda_A := \lambda_1(A\mathbb{Z}^n)$. If $[u] \neq [v]$ are represented in U , choose $\sigma \in \{\pm 1\}$ with $\langle u, \sigma v \rangle = |\langle u, v \rangle|$. The vector $u - \sigma v$ is a nonzero lattice vector, so its norm is at least λ_A . Since the relevant cosine numerator is nonnegative,

$$\begin{aligned} \frac{|\langle u, v \rangle|}{\|u\|_2 \|v\|_2} &= \frac{\|u\|_2^2 + \|v\|_2^2 - \|u - \sigma v\|_2^2}{2 \|u\|_2 \|v\|_2} \\ &\leq \frac{2(1 + \eta)^2 - 1}{2} = \frac{1}{2} + 2\eta + \eta^2 < \cos\left(\frac{\pi}{3} - \delta\right), \end{aligned} \quad (15)$$

where the last step is (3). This proves the projective separation. Since U is symmetric and contains no zero vector, every antipodal orbit has size two. (The bound $\eta < 1$ also rules out distinct positive multiples on one lattice line.) Consequently,

$$|U/\{\pm 1\}| = \frac{|U_0|}{2} \geq \frac{(1 + \eta/2)^n - 1}{12} = 2^{\Omega_{\varepsilon, \delta}(n)}. \quad (16)$$

This supplies a constant $c(\varepsilon, \delta) > 0$. Antipodal vectors themselves have ordinary angle π , so the final sentence of the theorem also follows.

Polynomial encoding. Equations (8)–(10) give

$$Q \leq 1 + \frac{4nH_n^{n-1}}{\eta}, \quad |A_{ij}| \leq QH_n + \frac{1}{2}. \quad (17)$$

Hence every A_{ij} has

$$O\left(n^2 + n \log n + \log \frac{1}{\eta}\right) = O_{\varepsilon, \delta}(n^2)$$

bits. There are n^2 entries, so the total binary encoding length is $O_{\varepsilon, \delta}(n^4)$. Increasing C and n_0 if necessary proves the first item and completes the proof. $\square$