On the Lattice Distortion Problem*†

Huck Bennett[1], Daniel Dadush‡[2], and Noah Stephens-Davidowitz[3]

1 **Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, New York, USA**
hbennett@cs.nyu.edu

2 **Centrum Wiskunde & Informatica, Amsterdam, The Netherlands**
dadush@cwi.nl

3 **Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, New York, USA**
noahsd@gmail.com

Abstract

We introduce and study the *Lattice Distortion Problem* (LDP). LDP asks how "similar" two lattices are. I.e., what is the minimal distortion of a linear bijection between the two lattices? LDP generalizes the Lattice Isomorphism Problem (the lattice analogue of Graph Isomorphism), which simply asks whether the minimal distortion is one.

As our first contribution, we show that the distortion between any two lattices is approximated up to a $n^{O(\log n)}$ factor by a simple function of their successive minima. Our methods are constructive, allowing us to compute low-distortion mappings that are within a $2^{O(n \log\log n/\log n)}$ factor of optimal in polynomial time and within a $n^{O(\log n)}$ factor of optimal in singly exponential time. Our algorithms rely on a notion of basis reduction introduced by Seysen (Combinatorica 1993), which we show is intimately related to lattice distortion. Lastly, we show that LDP is NP-hard to approximate to within any constant factor (under randomized reductions), by a reduction from the Shortest Vector Problem.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases lattices, lattice distortion, lattice isomoprhism, geometry of numbers, basis reduction

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.9

1 Introduction

An n-dimensional *lattice* $\mathcal{L} \subset \mathbb{R}^n$ is the set of all integer linear combinations of linearly independent vectors $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ with $\mathbf{b}_i \in \mathbb{R}^n$. We write the lattice generated by basis B as $\mathcal{L}(B) = \{\sum_{i=1}^n a_i \mathbf{b}_i : a_i \in \mathbb{Z}\}$.

Lattices are very well-studied classical mathematical objects (e.g., [25, 9]), and over the past few decades, computational problems on lattices have found a remarkably large number of applications in computer science. Algorithms for lattice problems have proven to be quite useful, and they have therefore been studied extensively (e.g., [20, 16, 3, 24]). And, over the past twenty years, many strong cryptographic primitives have been constructed with

* The full version of this paper is available at http://arxiv.org/abs/1605.03613.

† This material is based upon work partially supported by the National Science Foundation under Grant No. CCF-1320188 and Grant No. CCF-1423228. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

‡ Daniel Dadush was supported by the NWO Veni grant 639.071.510.

their security based on the (worst-case) hardness of various computational lattice problems (e.g., [1, 23, 12, 11, 28, 8]).

In this paper, we address a natural question: how "similar" are two lattices? I.e., given lattices $\mathcal{L}_1, \mathcal{L}_2$, does there exist a linear bijective mapping $T : \mathcal{L}_1 \to \mathcal{L}_2$ that does not change the distances between points by much? If we insist that T exactly preserves distances, then this is the *Lattice Isomorphism Problem* (LIP), which was studied in [26, 32, 15, 21]. We extend this to the *Lattice Distortion Problem* (LDP), which asks how well such a mapping T can *approximately* preserve distances between points.

Given two lattices $\mathcal{L}_1, \mathcal{L}_2$, we define the *distortion* between them as

$$\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) = \min\{\|T\|\|T^{-1}\| \,:\, T(\mathcal{L}_1) = \mathcal{L}_2\} \ ,$$

where $\|T\| = \sup_{\|\mathbf{x}\|=1} \|T\mathbf{x}\|$ is the *operator norm*. The quantity $\kappa(T) = \|T\| \cdot \|T^{-1}\|$ is the *condition number* of T, which measures how much T "distorts distances" (up to a fixed scaling). It is easy to check that $\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2)$ bounds the ratio between most natural geometric parameters of $\mathcal{L}_1$ and $\mathcal{L}_2$ (up to scaling), and hence $\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2)$ is a strong measure of "similarity" between lattices. In particular, $\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) = 1$ if and only if $\mathcal{L}_1, \mathcal{L}_2$ are isomorphic (i.e., if and only if they are related by a scaled orthogonal transformation).

The *Lattice Distortion Problem* (LDP) is then defined in the natural way as follows. The input is two n-dimensional lattices $\mathcal{L}_1, \mathcal{L}_2$ (each represented by a basis), and the goal is to compute a bijective linear transformation T mapping $\mathcal{L}_1$ to $\mathcal{L}_2$ such that $\kappa(T) = \mathcal{D}(\mathcal{L}_1, \mathcal{L}_2)$. In this work, we study the approximate search and decisional versions of this problem, defined in the usual way. We refer to them as γ-LDP and γ-GapLDP respectively, where $\gamma = \gamma(n) \geq 1$ is the approximation factor. (See Section 2.4 for precise definitions.)

1.1 Our Contribution

As our first main contribution, we show that the distortion between any two lattices can be approximated by a natural function of geometric lattice parameters. Indeed, our proof techniques are constructive, leading to our second main contribution: an algorithm that computes low-distortion mappings, with a trade-off between the running time and the approximation factor. Finally, we show hardness of approximating lattice distortion.

To derive useful bounds on the distortion between two lattices, it is intuitively clear that one should study the "different scales over which the two lattices live." A natural notion of this is given by the successive minima, which are defined as follows. The i^{th} successive minimum, $\lambda_i(\mathcal{L})$, of $\mathcal{L}$ is the minimum radius $r > 0$ such that $\mathcal{L}$ contains i linearly independent vectors of norm at most r. For example, a lattice generated by a basis of orthogonal vectors of lengths $0 < a_1 \leq \cdots \leq a_n$ has successive minima $\lambda_i(\mathcal{L}) = a_i$. Since low-distortion mappings approximately preserve distances, it is intuitively clear that two lattices can only be related by a low-distortion mapping if their successive minima are close to each other (up to a fixed scaling).

Concretely, for two n-dimensional lattices $\mathcal{L}_1, \mathcal{L}_2$, we define

$$M(\mathcal{L}_1, \mathcal{L}_2) = \max_{i \in [n]} \frac{\lambda_i(\mathcal{L}_2)}{\lambda_i(\mathcal{L}_1)} \ , \tag{1}$$

which measures how much we need to scale up $\mathcal{L}_1$ so that its successive minima are at least as large as those of $\mathcal{L}_2$. For any linear map T from $\mathcal{L}_1$ to $\mathcal{L}_2$, it is easy to see that $\lambda_i(\mathcal{L}_2) \leq \|T\|\lambda_i(\mathcal{L}_1)$. Thus, by definition $M(\mathcal{L}_1, \mathcal{L}_2) \leq \|T\|$. Applying the same reasoning for T^{-1}, we derive the following simple lower bound on distortion.

$$\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) \geq M(\mathcal{L}_1, \mathcal{L}_2) \cdot M(\mathcal{L}_2, \mathcal{L}_1) \ . \tag{2}$$

We note that this lower bound is tight when $\mathcal{L}_1, \mathcal{L}_2$ are each generated by bases of orthogonal vectors. But, it is a priori unclear if any comparable upper bound should hold for general lattices, since the successive minima are a very "coarse" characterization of the geometry of the lattice. Nevertheless, we show a corresponding upper bound.

▶ **Theorem 1.** *Let $\mathcal{L}_1, \mathcal{L}_2$ be n-dimensional lattices. Then,*

$$M(\mathcal{L}_1, \mathcal{L}_2) \cdot M(\mathcal{L}_2, \mathcal{L}_1) \leq \mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) \leq n^{O(\log n)} \cdot M(\mathcal{L}_1, \mathcal{L}_2) \cdot M(\mathcal{L}_2, \mathcal{L}_1) \ .$$

In particular, Theorem 1, together with standard transference theorems (e.g., [7]), implies that $n^{O(\log n)}$-GapLDP is in $\mathsf{NP} \cap \mathsf{coNP}$. While the factor on the right-hand side of the theorem might be far from optimal, we show in Section 5.1 that it cannot be improved below $\Omega(\sqrt{n})$. Intuitively, this is because there exist lattices that are much more dense than $\mathbb{Z}^n$ over large scales but still have $\lambda_i(\mathcal{L}) = \Theta(1)$ for all i. I.e., there exist very dense lattice sphere packings (see, e.g., [31]).

To prove the above theorem, we make use of the intuition that a low-distortion mapping T from $\mathcal{L}_1$ to $\mathcal{L}_2$ should map a "short" basis B_1 of $\mathcal{L}_1$ to a "short" basis B_2 of $\mathcal{L}_2$. (Note that the condition $TB_1 = B_2$ completely determines $T = B_2 B_1^{-1}$.) The difficulty here is that standard notions of "short" fail for the purpose of capturing low-distortion mappings. In particular, in Section 5.2, we show that Hermite-Korkine-Zolotarev (HKZ) reduced bases, one of the strongest notions of "shortest possible" lattice bases, do not suffice by themselves for building low-distortion mappings. (See Section 2.6 for the definition of HKZ-reduced bases.) In particular, we give a simple example of a lattice $\mathcal{L}$ where an HKZ-reduced basis of $\mathcal{L}$ misses the optimal distortion $\mathcal{D}(\mathbb{Z}^n, \mathcal{L})$ by an exponential factor.

Fortunately, we show that a suitable notion of shortness does exist for building low-distortion mappings by making a novel connection between low-distortion mappings and a notion of basis reduction introduced by Seysen [30]. In particular, for a basis $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ and dual basis $B^* = B^{-T} = [\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*]$, Seysen's condition number is defined as

$$S(B) = \max_{i \in [n]} \|\mathbf{b}_i\| \|\mathbf{b}_i^*\| \ .$$

Note that we always have $\langle \mathbf{b}_i, \mathbf{b}_i^* \rangle = 1$, so this parameter measures how tight the Cauchy-Schwarz inequality is over all primal-dual basis-vector pairs. We extend this notion and define $S(\mathcal{L})$ as the minimum of $S(B)$ over all bases B of $\mathcal{L}$. Using this notion, we give an effective version of Theorem 1 as follows.

▶ **Theorem 2.** *Let $\mathcal{L}_1, \mathcal{L}_2$ be n-dimensional lattices. Let $B_1, B_2 \in \mathbb{R}^{n \times n}$ be bases of $\mathcal{L}_1, \mathcal{L}_2$ whose columns are sorted in non-decreasing order of length. Then, we have that*

$$M(\mathcal{L}_1, \mathcal{L}_2) M(\mathcal{L}_2, \mathcal{L}_1) \leq \kappa(B_2 B_1^{-1}) \leq n^4 S(B_1)^2 S(B_2)^2 \cdot M(\mathcal{L}_1, \mathcal{L}_2) M(\mathcal{L}_2, \mathcal{L}_1) \ .$$

In particular, we have that

$$M(\mathcal{L}_1, \mathcal{L}_2) M(\mathcal{L}_2, \mathcal{L}_1) \leq \mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) \leq n^4 S(\mathcal{L}_1)^2 S(\mathcal{L}_2)^2 \cdot M(\mathcal{L}_1, \mathcal{L}_2) M(\mathcal{L}_2, \mathcal{L}_1) \ .$$

From here, the bound in Theorem 1 follows directly from the following (surprising) theorem of Seysen.

▶ **Theorem 3** (Seysen [30])**.** *For any $\mathcal{L} \subset \mathbb{R}^n$, $S(\mathcal{L}) \leq n^{O(\log n)}$.*

This immediately yields an algorithm for approximating the distortion between two lattices, by using standard lattice algorithms to approximate $M(\mathcal{L}_1, \mathcal{L}_2)$ and $M(\mathcal{L}_2, \mathcal{L}_1)$. But,

Seysen's proof of the above theorem is actually constructive! In particular, he shows how to efficiently convert any suitably reduced lattice basis into a basis with a low Seysen condition number. (See Section 2.6.2 for details.) Using this methodology, combined with standard basis reduction techniques, we derive the following time-approximation trade-off for γ-LDP.

▶ **Theorem 4** (Algorithm for LDP). *For any $\log n \leq k \leq n$, there is an algorithm solving $k^{O(n/k+\log n)}$-LDP in time $2^{O(k)}$.*

In other words, using the bounds in Theorem 1 together with known algorithms, we are able to approximate the distortion between two lattices. But, with a bit more work, we are able to solve *search* LDP by explicitly computing a low-distortion mapping between the input lattices.

We also prove the following lower bound for LDP.

▶ **Theorem 5** (Hardness of LDP). *γ-GapLDP is NP-hard under randomized polynomial-time reductions for any constant $\gamma \geq 1$.*

In particular, we show a reduction from approximating the (decisional) Shortest Vector Problem (GapSVP) over lattices to γ-GapLDP, where the approximation factor that we obtain for GapSVP is $O(\gamma)$. Since hardness of GapSVP is quite well-studied [2, 22, 17, 14], we are immediately able to import many hardness results to GapLDP. (See Corollary 30 and Theorem 31 for the precise statements.)

1.2 Comparison to related work

The main related work of which we are aware is that of Haviv and Regev [15] on the Lattice Isomorphism Problem (LIP). In their paper, they give an $n^{O(n)}$-time algorithm for solving LIP exactly, which proceeds by cleverly identifying a small candidate set of bases of $\mathcal{L}_1$ and $\mathcal{L}_2$ that must be mapped to each other by any isomorphism. One might expect that such an approach should also work for the purpose of solving LDP either exactly or for approximation factors below $n^{O(\log n)}$. However, the crucial assumption in LIP, that vectors in one lattice must be mapped to vectors of the same length in the other, completely breaks down in the current context. We thus do not know how to extend their techniques to LDP.

Much more generally, we note that LIP is closely related to the Graph Isomorphism Problem (GI). For example, both problems are in SZK but not known to be in P (although recent work on algorithms for GI has been quite exciting [6]!), and GI reduces to LIP [32]. Therefore, LDP is qualitatively similar to the Approximate Graph Isomorphism Problem, which was studied by Arora, Frieze, and Kaplan [4], who showed an upper bound, and Arvind, Köbler, Kuhnert, and Vasudev [5], who proved both upper and lower bounds. In particular, [5] showed that various versions of this problem are NP-hard to approximate to within a constant factor. Qualitatively, these hardness results are similar to our Theorem 5.

1.3 Conclusions and Open Problems

In conclusion, we introduce the Lattice Distortion Problem and show a connection between LDP and the notion of Seysen-reduced bases. We use this connection to derive time-approximation trade-offs for LDP. We also prove approximation hardness for GapLDP, showing a qualitative difference with LIP (which is unlikely to be NP-hard under reasonable complexity theoretic assumptions).

One major open question is what the correct bound in Theorem 3 is. In particular, there are no known families of lattices for which the Seysen condition number is provably

superpolynomial, and hence it is possible that $S(\mathcal{L}) = \text{poly}(n)$ for any n-dimensional lattice $\mathcal{L}$. A better bound would immediately improve our Theorem 2 and give a better approximation factor for GapLDP.

We also note that all of our algorithms solve LDP for arguably very large approximation factors $n^{\Omega(\log n)}$. We currently do not even know whether there exists a fixed-dimension polynomial-time algorithm for γ-LDP for any $\gamma = n^{o(\log n)}$. The main problem here is that we do not have any good characterization of nearly optimal distortion mappings between lattices.

Organization. In Section 2, we present necessary background material. In Section 3, we give our approximations for lattice distortion, proving Theorems 2 and 4. In Section 4, we give the hardness for lattice distortion, proving Theorem 5. In Section 5, we give some illustrative example instances of lattice distortion.

2 Preliminaries

For $\mathbf{x} \in \mathbb{R}^n$, we write $\|\mathbf{x}\|$ for the Euclidean norm of $\mathbf{x}$. We omit any mention of the bit length in the running time of our algorithms. In particular, all of our algorithms take as input vectors in $\mathbb{Q}^n$ and run in time $f(n) \cdot \text{poly}(m)$ for some f, where m is the maximal bit length of an input vector. We therefore suppress the factor of $\text{poly}(m)$.

2.1 Lattices

The i^{th} *successive minimum* of a lattice $\mathcal{L}$ is defined as $\lambda_i(\mathcal{L}) = \inf\{r > 0 : \dim(\text{span}(rB_2^n \cap \mathcal{L})) \geq i\}$. That is, the first successive minimum is the length of the shortest non-zero lattice vector, the second successive minimum is the length of the shortest lattice vector which is linearly independent of a vector achieving the first, and so on. When $\mathcal{L}$ is clear from context, we simply write λ_i.

The *dual lattice* of $\mathcal{L}$ is defined as $\mathcal{L}^* = \{\mathbf{x} \in \mathbb{R}^n : \forall \mathbf{y} \in \mathcal{L}\ \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$. If $\mathcal{L} = \mathcal{L}(B)$ then $\mathcal{L}^* = \mathcal{L}(B^*)$ where $B^* = B^{-T}$, the inverse transpose of B. We call $B^* = [\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*]$ the *dual basis* of B, and write $\lambda_i^* = \lambda_i(\mathcal{L}^*)$. We will repeatedly use Banaszczyk's Transference Theorem, which relates the successive minima of a lattice to those of its dual.

▶ **Theorem 6** (Banaszczyk's Transference Theorem [7]). *For every rank n lattice $\mathcal{L}$ and every $i \in [n]$, $1 \leq \lambda_i(\mathcal{L})\lambda_{n-i+1}(\mathcal{L}^*) \leq n$.*

Given a lattice $\mathcal{L}$, we define the *determinant* of $\mathcal{L}$ as $\det(\mathcal{L}) := |\det(B)|$, where B is a basis with $\mathcal{L}(B) = \mathcal{L}$. Since two bases B, B' of $\mathcal{L}$ differ by a unimodular transformation, we have that $|\det(B)| = |\det(B')|$ so that $\det(\mathcal{L})$ is well-defined.

We sometimes work with lattices that do not have full rank—i.e., lattices generated by d linearly independent vectors $\mathcal{L} = \mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_d)$ with $d < n$. In this case, we simply identify $\text{span}(\mathbf{b}_1, \ldots, \mathbf{b}_d)$ with $\mathbb{R}^d$ and consider the lattice to be embedded in this space.

2.2 Linear mappings between lattices

We next characterize linear mappings between lattices in terms of bases.

▶ **Lemma 7.** *Let $\mathcal{L}_1, \mathcal{L}_2$ be full-rank lattices. Then a mapping $T : \mathcal{L}_1 \to \mathcal{L}_2$ is bijective and linear if and only if $T = BA^{-1}$ for some bases A, B of $\mathcal{L}_1, \mathcal{L}_2$ respectively. In particular, for any basis A of $\mathcal{L}_1$, $T(A)$ is a basis of $\mathcal{L}_2$.*

Proof. We first show that such a mapping is a bijection from $\mathcal{L}_1$ to $\mathcal{L}_2$. Let $T = BA^{-1}$ where $A = [\mathbf{a}_1, \ldots, \mathbf{a}_n]$ and $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ are bases of $\mathcal{L}_1, \mathcal{L}_2$ respectively. Because T has full rank, it is injective as a mapping from $\mathbb{R}^n$ to $\mathbb{R}^n$, and it is therefore injective as a mapping from $\mathcal{L}_1$ to $\mathcal{L}_2$. We have that for every $\mathbf{w} \in \mathcal{L}_2$, $\mathbf{w} = \sum_{i=1}^n c_i \mathbf{b}_i$ with $c_i \in \mathbb{Z}$. Let $\mathbf{v} = \sum_{i=1}^n c_i \mathbf{a}_i \in \mathcal{L}_1$. Then, $T(\mathbf{v}) = T(\sum_{i=1}^n c_i \mathbf{a}_i) = \sum_{i=1}^n c_i \mathbf{b}_i = \mathbf{w}$. Therefore, T is a bijection from $\mathcal{L}_1$ to $\mathcal{L}_2$.

We next show that any linear map T with $T(\mathcal{L}_1) = \mathcal{L}_2$ must have this form. Let $A = [\mathbf{a}_1, \ldots, \mathbf{a}_n]$ be a basis of $\mathcal{L}_1$, and let $B = T(A)$. We claim that $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ is a basis of $\mathcal{L}_2$.

Let $\mathbf{w} \in \mathcal{L}_2$. Because T is a bijection between $\mathcal{L}_1$ and $\mathcal{L}_2$, there exists $\mathbf{v} \in \mathcal{L}_1$ such that $T\mathbf{v} = w$. Using the definition of a basis and the linearity of T,

$$\mathbf{w} = T\mathbf{v} = T\big(\sum_{i=1}^n c_i \mathbf{a}_i\big) = \sum_{i=1}^n c_i \mathbf{b}_i,$$

for some $c_1, \ldots, c_n \in \mathbb{Z}$. Because $\mathbf{w}$ was picked arbitrarily, it follows that B is a basis of $\mathcal{L}_2$. ◀

2.3 Seysen's condition number $S(B)$

Seysen shows how to take any basis with relatively low multiplicative drop in its Gram-Schmidt vectors and convert it into a basis with relatively low $S(B) = \max_i \|\mathbf{b}_i\|\|\mathbf{b}_i^*\|$ [30]. By combining this with Gama and Nguyen's slide reduction technique [10], we obtain the following result.

▶ **Theorem 8.** *For every $\log n \le k \le n$ there exists an algorithm that takes a lattice $\mathcal{L}$ as input and computes a basis B of $\mathcal{L}$ with $S(B) \le k^{O(n/k+\log k)}$ in time $2^{O(k)}$.*

In particular, applying Seysen's procedure to slide-reduced bases suffices. We include a proof of Theorem 8 and a high-level description of Seysen's procedure in Section 2.6.

2.4 The Lattice Distortion Problem

▶ **Definition 9.** For any $\gamma = \gamma(n) \ge 1$, the γ-Lattice Distortion Problem (γ-LDP) is the search problem defined as follows. The input consists of two lattices $\mathcal{L}_1, \mathcal{L}_2$ (represented by bases $B_1, B_2 \in \mathbb{Q}^{n \times n}$). The goal is to output a matrix $T \in \mathbb{R}^{n \times n}$ such that $T(\mathcal{L}_1) = \mathcal{L}_2$ and $\kappa(T) \le \gamma \cdot \mathcal{D}(\mathcal{L}_1, \mathcal{L}_2)$.

▶ **Definition 10.** For any $\gamma = \gamma(n) \ge 1$, the γ-GapLDP is the promise problem defined as follows. The input consists of two lattices $\mathcal{L}_1, \mathcal{L}_2$ (represented by bases $B_1, B_2 \in \mathbb{Q}^{n \times n}$) and a number $c \ge 1$. The goal is to decide between a 'YES' instance where $\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) \le c$ and a 'NO' instance where $\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2) > \gamma \cdot c$.

2.5 Complexity of LDP

We show some basic facts about the complexity of GapLDP. First, we show that the Lattice Isomorphism Problem (LIP) corresponds to the special case of GapLDP where $c = 1$. LIP takes bases of $\mathcal{L}_1, \mathcal{L}_2$ as input and asks if there exists an orthogonal linear transformation O such that $O(\mathcal{L}_1) = \mathcal{L}_2$. Haviv and Regev [15] show that there exists an $n^{O(n)}$-time algorithm for LIP, and that LIP is in the complexity class SZK.

▶ **Lemma 11.** *There is a polynomial-time reduction from LIP to 1-GapLDP.*

Proof. Let $\mathcal{L}_1, \mathcal{L}_2$ be an LIP instance. First check that $\det(\mathcal{L}_1) = \det(\mathcal{L}_2)$. If not, then output a trivial 'NO' instance of 1-GapLDP. Otherwise, map the LIP instance to the 1-GapLDP instance with the same input bases and $c = 1$. For any $T : \mathcal{L}_1 \to \mathcal{L}_2$, we must have $\det(T) = 1$, and therefore $\kappa(T) = 1$ if and only if $\|T\| = \left\|T^{-1}\right\| = 1$. So, this is a 'YES' instance of GapLDP if and only if $\mathcal{L}_1, \mathcal{L}_2$ are isomorphic. ◀

▶ **Lemma 12.** 1-*GapLDP is in* NP.

Proof. Let $I = (\mathcal{L}_1, \mathcal{L}_2, c)$ be an instance of GapLDP, and let s be the length of I. We will show that for a 'YES' instance, there are bases A, B of $\mathcal{L}_1, \mathcal{L}_2$ respectively such that $T = BA^{-1}$ requires at most poly(s) bits to specify and $\kappa(T) \le c$. Assume without loss of generality that $\mathcal{L}_1, \mathcal{L}_2 \subseteq \mathbb{Z}^n$. Otherwise, scale the input lattices to achieve this at the expense of a factor s blow-up in input size.

To satisfy $\|T\| \left\|T^{-1}\right\| \le c$, we must have that $|t_{ij}| \le \|T\| \le c \cdot \det(\mathcal{L}_2)/\det(\mathcal{L}_1) \le c \cdot \det(\mathcal{L}_2)$ for each entry t_{ij} of T. By Cramer's rule, each entry of A^{-1} and hence T will be an integer multiple of $\frac{1}{\det \mathcal{L}_1}$, so we can assume without loss of generality that the denominator of each entry of T is $\det \mathcal{L}_1$.

Combining these bounds and applying Hadamard's inequality, we get that $|t_{ij}|$ takes at most

$$\log(c \cdot \det(\mathcal{L}_1) \det(\mathcal{L}_2)) \le \log\left(c \cdot \prod_{i=1}^{n} \|\mathbf{a}_i\| \prod_{i=1}^{n} \|\mathbf{b}_i\|\right)$$

bits to specify. Accounting for the sign of each t_{ij}, it follows that T takes at most $n^2 \cdot \log(2c \cdot \prod_{i=1}^{n} \|\mathbf{a}_i\| \|\mathbf{b}_i\|) \le n^2 \cdot (s+1)$ bits to specify. ◀

We remark that we can replace c with the quantity $n^{O(\log n)} M(\mathcal{L}_1, \mathcal{L}_2) M(\mathcal{L}_2, \mathcal{L}_1)$ (as given by the upper bound in Theorem 1) in the preceding argument to obtain an upper bound on the distortion of an *optimal* mapping T that does not depend on c.

2.6 Basis reduction

In this section, we define various notions of basis reductions and show how to use them to prove Theorem 8.

For a basis $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$, we write $\pi_i^{(B)} := \pi_{\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\}^\perp}$ to represent projection onto the subspace $\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\}^\perp$. We then define the *Gram-Schmidt* orthogonalization of B, $(\tilde{b}_1, \ldots, \tilde{b}_n)$ as $\tilde{b}_i = \pi_i^{(B)}(\mathbf{b}_i)$. By construction the vectors $\tilde{b}_1, \ldots, \tilde{b}_n$ are orthogonal, and each $\mathbf{b}_i$ is a linear combination of $\tilde{b}_1, \ldots, \tilde{b}_i$. We define $\mu_{ij} = \frac{\langle \mathbf{b}_i, \tilde{b}_j \rangle}{\langle \tilde{b}_j, \tilde{b}_j \rangle}$.

We define the QR-decomposition of a full-rank matrix B as $B = QR$ where Q has orthonormal columns, and R is upper triangular. The QR-decomposition of a matrix is unique, and can be computed efficiently by applying Gram-Schmidt orthogonalization to the columns of B.

Unimodular matrices, denoted $GL(n, \mathbb{Z})$, form the multiplicative group of $n \times n$ matrices with integer entries and determinant ± 1.

▶ **Fact 13.** $\mathcal{L}(B) = \mathcal{L}(B')$ *if and only if there exists* $U \in GL(n, \mathbb{Z})$ *such that* $B' = B \cdot U$.

Based on this, a useful way to view basis reduction is as right-multiplication by unimodular matrices.

2.6.1 Slide-reduced bases

A very strong notion of basis reduction introduced by Korkine and Zolotareff [18] gives one way of formalizing what it means to be a "shortest-possible" lattice basis.

▶ **Definition 14** ([18], Definition 1 in [30]). Let B be a basis of $\mathcal{L}$. $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ is HKZ (Hermite-Korkine-Zolotareff) reduced if

1. $\forall j < i,\ |\mu_{ij}| \leq \frac{1}{2}$;
2. $\|b_1\| = \lambda_1(\mathcal{L}(B))$; and
3. if $n > 1$, then $[\pi_2^{(B)}(\mathbf{b}_2), \ldots, \pi_2^{(B)}(\mathbf{b}_n)]$ is an HKZ basis of $\pi_2^{(B)}(\mathcal{L})$.

By definition, the first vector $\mathbf{b}_1$ in an HKZ basis is a shortest vector in the lattice. Furthermore, computing an HKZ basis can be achieved by making n calls to an SVP oracle. So, the two problems have the same time complexity up to a factor of n. In particular, computing HKZ bases is NP-hard.

Gama and Nguyen (building on the work of Schnorr [29]) introduced the notion of slide-reduced bases [10], which can be thought of as a relaxed notion of HKZ bases that can be computed more efficiently.

▶ **Definition 15** ([10, Definition 1]). Let B be a basis of $\mathcal{L} \subset \mathbb{Q}^n$ and $\varepsilon > 0$. We say that B is ε-DSVP (dual SVP) reduced if its corresponding dual basis $[\mathbf{b}_1^*, \ldots, \mathbf{b}_n^*]$ satisfies $\|\mathbf{b}_n^*\| \leq (1+\varepsilon) \cdot \lambda_1(\mathcal{L}^*)$.

Then, for $k \geq 2$ an integer dividing n, we say that $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ is (ε, k)-slide reduced if

1. $\forall j < i,\ |\mu_{ij}| \leq \frac{1}{2}$;
2. $\forall 0 \leq i \leq n/k - 1$, the "projected truncated basis" $[\pi_{ik+1}^{(B)}(\mathbf{b}_{ik+1}), \ldots, \pi_{ik+1}^{(B)}(\mathbf{b}_{ik+k})]$ is HKZ reduced; and
3. $\forall 0 \leq i \leq n/k-2$, the "shifted projected truncated basis" $[\pi_{ik+2}^{(B)}(\mathbf{b}_{ik+2}), \ldots, \pi_{ik+2}^{(B)}(\mathbf{b}_{ik+k+1})]$ is ε-DSVP reduced.

▶ **Theorem 16** ([10]). *There is an algorithm that takes as input a lattice $\mathcal{L} \subset \mathbb{Q}^n$, $\varepsilon > 0$, and integer $k \geq \log n$ dividing n and outputs a (k, ε)-slide-reduced basis of $\mathcal{L}$ in time* $\mathrm{poly}(1/\varepsilon) \cdot 2^{O(k)}$.

We will be particularly concerned with the ratios between the length of the Gram-Schmidt vectors of a given basis. We prefer bases whose Gram-Schmidt vectors do not "decay too quickly," and we measure this decay by

$$\eta(B) = \max_{i \leq j} \frac{\left\|\tilde{b}_i\right\|}{\left\|\tilde{b}_j\right\|} .$$

Previous work bounded $\eta(B)$ for HKZ-reduced bases as follows.

▶ **Theorem 17** ([19, Proposition 4.2]). *For any HKZ-reduced basis B over $\mathbb{Q}^n$, $\eta(B) \leq n^{O(\log n)}$.*

Using Theorem 17 and some of the results in [10] we get a bound on $\eta(B)$ for slide-reduced bases.

▶ Proposition 18. For any $\log n \leq k \leq n$, there is an algorithm that takes as input a lattice $\mathcal{L} \subset \mathbb{Q}^n$ and outputs a basis B of $\mathcal{L}$ such that $\eta(B) \leq k^{O(n/k + \log k)}$. Furthermore, the algorithm runs in time $2^{O(k)}$.

See the full version of this paper for a proof of Proposition 18.

2.6.2 Seysen bases

Although slide-reduced bases B consist of short vectors and have bounded $\eta(B)$, they make only weak guarantees about the length of vectors in the dual basis B^*. Of course, one way to compute a basis whose dual will contain short dual basis is short is to simply compute B such that B^* is a suitably reduced basis of $\mathcal{L}^*$. Such a basis B is called a dual-reduced basis, and sees use in applications such as [15].

However, we would like to compute a basis such that the vectors in B and B^* are both short, which Seysen addressed in his work [30]. Seysen's main result finds a basis B such that both B and B^* are short by dividing this problem into two subproblems. The first involves finding a basis with small $\eta(B)$, as in Section 2.6.1. The second subproblem, discussed in [30, Section 3], involves conditioning unipotent matrices. Let $N(n,\mathbb{R})$ be the multiplicative group of unipotent $n \times n$-matrices. That is, a matrix $A \in N(n,\mathbb{R})$ if $a_{ii} = 1$ and $a_{ij} = 0$ for $i > j$ (i.e., A is upper triangular and has ones on the main diagonal). Let $N(n,\mathbb{Z})$ be the subgroup of $N(n,\mathbb{R})$ with integer entries. Because $N(n,\mathbb{Z})$ is a subset of $GL(n,\mathbb{Z})$, we trivially have that $\mathcal{L}(B) = \mathcal{L}(B \cdot U)$ for every $U \in N(n,\mathbb{Z})$.

Let $\|B\|_\infty := \max_{i,j\in[n]} |b_{ij}|$ denote the largest magnitude of an entry in B. We follow Seysen [30] and define $S'(B) = \max\{\|B\|_\infty, \|B^{-1}\|_\infty\}$. We also let

$$\zeta(n) = \sup_{A \in N(n,\mathbb{R})} \left\{ \inf_{U \in N(n,\mathbb{Z})} \{S'(A \cdot U)\} \right\}.$$

▶ **Theorem 19** ([30, Prop. 5 and Thm. 6]). *There exists an algorithm* Seysen *that takes as input $A \in N(n,\mathbb{R})$ and outputs $A \cdot U$ where $U \in N(n,\mathbb{Z})$ and $S'(A \cdot U) \le n^{O(\log n)}$ in time $O(n^3)$. In particular, $\zeta(n) \le n^{O(\log n)}$.*

Let $B = QR$ be a QR-decomposition of B. We may further decompose R as $R = DR'$, where $d_{ii} = \|\tilde{b}_i\|$ and

$$r'_{ij} = \begin{cases} 0 & \text{if } j < i, \\ 1 & \text{if } j = i, \\ \mu_{ji} & \text{if } j > i. \end{cases}$$

In particular, note that $R' \in N(n,\mathbb{R})$. It is easy to see that $\eta(B)$ controls $\|D\|\|D^{-1}\|$. On the other hand, using the bound on $\zeta(n)$, we can always multiply B on the right by $U \in N(n,\mathbb{Z})$ to control the size of $\|R'\|\|R'^{-1}\|$. Roughly speaking, these two facts imply Theorem 20.

▶ **Theorem 20** ([30, Theorem 7]). *Let $B =$ Seysen(B') where B' is a matrix. Then $S(B) \le n \cdot \eta(B') \cdot \zeta(n)^2$.*

Proof of Theorem 8. Let $B =$ Seysen(B'), where B' is a basis as computed in Proposition 18. We then have that

$$\begin{aligned} S(B) &\le n \cdot \eta(B') \cdot \zeta(n)^2 && \text{(by Theorem 20)} \\ &\le n \cdot k^{O(n/k+\log k)} \cdot \zeta(n)^2 && \text{(by Proposition 18)} \\ &\le n \cdot k^{O(n/k+\log k)} \cdot (n^{O(\log n)})^2 && \text{(by Theorem 19)} \\ &\le k^{O(n/k+\log k)}. \end{aligned}$$

We can compute B' in $2^{O(k)}$ time using Proposition 18. Moreover, by Theorem 19, Seysen runs in $O(n^3)$ time. Therefore the algorithm runs in $2^{O(k)}$ time. ◀

3 Approximating lattice distortion

In this section, we show how to compute low-distortion mappings between lattices by using bases with low $S(B)$.

3.1 Basis length bounds in terms of $S(B)$

Call a basis $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ *sorted* if $\|\mathbf{b}_1\| \leq \cdots \leq \|\mathbf{b}_n\|$. Clearly, $\|\mathbf{b}_i\| / \lambda_i \geq 1$ for a sorted basis B. Note that sorting B does not change $S(B)$, since $S(\cdot)$ is invariant under permutations of the basis vectors.

A natural way to quantify the "shortness" of a lattice basis is to upper bound $\|\mathbf{b}_k\| / \lambda_k$ for all $k \in [n]$. For example, [19] shows that $\|\mathbf{b}_k\| / \lambda_k \leq \sqrt{n}$ when B is an HKZ basis. We prove a characterization of Seysen-reduced bases, showing that *both* the primal basis vectors and the dual basis vectors are not much longer than the successive minima. Namely, we show that $S(B)$ is an upper bound on both $\|\mathbf{b}_k\| / \lambda_k$ and $\|\mathbf{b}_k^*\| / \lambda_{n-k+1}^*$ for sorted bases B. This characterization is key to bounding the distortion between two lattices, and it might be of independent interest.

▶ **Lemma 21.** *Let B be a sorted basis of $\mathcal{L}$. Then $\max_{k \in [n]} \|\mathbf{b}_k\| / \lambda_k \leq S(B)$.*

Proof. For every $k \in [n]$, we have

$$\begin{aligned}
\|\mathbf{b}_k\| / \lambda_k &\leq \|\mathbf{b}_k\| \lambda_{n-k+1}^* && \text{(by the lower bound in Theorem 6)} \\
&\leq \|\mathbf{b}_k\| \max_{i \in \{k, \ldots, n\}} \|\mathbf{b}_i^*\| && \text{(the } \mathbf{b}_i^* \text{ are linearly independent)} \\
&\leq \max_{i \in \{k, \ldots, n\}} \|\mathbf{b}_i\| \, \|\mathbf{b}_i^*\| && (B \text{ is sorted}) \\
&\leq S(B).
\end{aligned}$$

◀

▶ **Lemma 22.** *Let B be a sorted basis of $\mathcal{L}$. Then $\max_{k \in [n]} \|\mathbf{b}_k^*\| / \lambda_{n-k+1}^* \leq S(B)$.*

Proof. For every $k \in [n]$, we have

$$\frac{\|\mathbf{b}_k^*\|}{\lambda_{n-k+1}^*} \leq \frac{\|\mathbf{b}_k\| \, \|\mathbf{b}_k^*\|}{\lambda_k \lambda_{n-k+1}^*} \leq \max_{i \in [n]} \|\mathbf{b}_i\| \, \|\mathbf{b}_i^*\| = S(B).$$

The first inequality follows from the assumption that B is sorted, and the second follows from the lower bound in Theorem 6. ◀

3.2 Approximating LDP using Seysen bases

In this section, we bound the distortion $\mathcal{D}(\mathcal{L}_1, \mathcal{L}_2)$ between lattices $\mathcal{L}_1, \mathcal{L}_2$. The upper bound is constructive and depends on $S(B_1), S(B_2)$, which naturally leads to Theorem 4.

▶ **Lemma 23.** *Let $A = [\mathbf{a}_1, \ldots, \mathbf{a}_n]$ and $B = [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ be sorted bases of $\mathcal{L}_1, \mathcal{L}_2$ respectively. Then,*

$$\left\| BA^{-1} \right\| \leq n^2 S(A) S(B) M(\mathcal{L}_1, \mathcal{L}_2).$$

Proof.

$$\begin{aligned}
\left\|BA^{-1}\right\| &= \left\|\sum_{i=1}^{n} \mathbf{b}_i(\mathbf{a}_i^*)^T\right\| \\
&\leq \sum_{i=1}^{n} \left\|\mathbf{b}_i(\mathbf{a}_i^*)^T\right\| && \text{(by triangle inequality)} \\
&= \sum_{i=1}^{n} \|\mathbf{b}_i\| \, \|\mathbf{a}_i^*\| \\
&\leq n \max_{i\in[n]} \|\mathbf{b}_i\| \, \|\mathbf{a}_i^*\| \\
&\leq nS(B) \max_{i\in[n]} \lambda_i(\mathcal{L}_2) \, \|\mathbf{a}_i^*\| && \text{(by Lemma 21)} \\
&\leq nS(A)S(B) \max_{i\in[n]} \lambda_i(\mathcal{L}_2)\lambda_{n-i+1}^*(\mathcal{L}_1) && \text{(by Lemma 22)} \\
&\leq n^2 S(A)S(B)M(\mathcal{L}_1,\mathcal{L}_2). && \text{(by Theorem 6)}
\end{aligned}$$

◀

Proof of Theorem 2. Note that by definition there always exist bases B_1, B_2 of $\mathcal{L}_1, \mathcal{L}_2$ respectively achieving $S(B_i) = S(\mathcal{L}_i)$. Therefore, applying Lemma 23 twice to bound both $\left\|B_2B_1^{-1}\right\|$ and $\left\|B_1B_2^{-1}\right\|$, we get the upper bound.

For the lower bound, let $\mathbf{v}_1, \ldots, \mathbf{v}_n \in \mathcal{L}_1$ be linearly independent vectors such that $\|\mathbf{v}_i\| = \lambda_i(\mathcal{L}_1)$ for every i. Then, for every i,

$$\lambda_i(\mathcal{L}_2) \leq \max_{j\in[i]} \|T\mathbf{v}_j\| \leq \|T\| \max_{j\in[i]} \|\mathbf{v}_j\| = \|T\| \, \lambda_i(\mathcal{L}_1).$$

Rearranging, we get that $\lambda_i(\mathcal{L}_2)/\lambda_i(\mathcal{L}_1) \leq \|T\|$. This holds for arbitrary i, so in particular $\max_{i\in[n]} \lambda_i(\mathcal{L}_2)/\lambda_i(\mathcal{L}_1) = M(\mathcal{L}_1,\mathcal{L}_2) \leq \|T\|$. The same computation with $\mathcal{L}_1, \mathcal{L}_2$ reversed shows that $M(\mathcal{L}_2,\mathcal{L}_1) \leq \left\|T^{-1}\right\|$. Multiplying these bounds together implies the lower bound in the theorem statement. ◀

We can now prove Theorem 4.

Proof of Theorem 4. Let $(\mathcal{L}_1, \mathcal{L}_2)$ be an instance of LDP. For $i = 1, 2$, compute a basis B_i of $\mathcal{L}_i$ using the algorithm described in Theorem 8 with parameter k. We have that $S(B_i) \leq k^{O(n/k+\log k)}$. This computation takes $2^{O(k)}$ time. The algorithm then simply outputs $T = B_2B_1^{-1}$.

By Lemma 23 and the upper bounds on $S(B_i)$, we get that $\kappa(T) \leq k^{O(n/k+\log k)} \cdot M(\mathcal{L}_1,\mathcal{L}_2) \cdot M(\mathcal{L}_2,\mathcal{L}_1)$. This is within a factor of $k^{O(n/k+\log k)} \cdot n^{O(\log n)} = k^{O(n/k+\log k)}$ of $\mathcal{D}(\mathcal{L}_1,\mathcal{L}_2)$ by Theorem 1. So, the algorithm is correct. ◀

4 Hardness of LDP

In this section, we prove the hardness of γ-GapLDP. (See Theorem 31.) Our reduction works in two steps. First, we show how to use an oracle for GapLDP to solve a variant of GapCVP that we call γ-GapCVP$^\alpha$. (See Definition 24 and Theorem 26.) Given a CVP instance consisting of a lattice $\mathcal{L}$ and a target vector $\mathbf{t}$, our idea is to compare "$\mathcal{L}$ with $\mathbf{t}$ appended to it" to "$\mathcal{L}$ with an extra orthogonal vector appended to it." (See Eq. (3).) We show that, if $\text{dist}(\mathbf{t}, \mathcal{L})$ is small, then these lattices will be similar. On the other hand, if (1) $\text{dist}(k\mathbf{t}, \mathcal{L})$ is

large for all non-zero integers k, and (2) $\lambda_1(\mathcal{L})$ is not too small; then the two lattices must be quite dissimilar.

We next show that γ-GapCVP$^\alpha$ is as hard as GapSVP. (See Theorem 29.) This reduction is a variant of the celebrated reduction of [13]. It differs from the original in that it "works in base p" instead of in base two, and it "adds an extra coordinate to $\mathbf{t}$." We show that this is sufficient to satisfy the promises required by γ-GapCVP$^\alpha$.

Both reductions are relatively straightforward.

4.1 Reduction from a variant of CVP

▶ **Definition 24.** For any $\gamma = \gamma(n) \geq 1$ and $\alpha = \alpha(n) > 0$, γ-GapCVP$^\alpha$ is the promise problem defined as follows. The input is a lattice $\mathcal{L} \subset \mathbb{Q}^n$, a target $\mathbf{t} \in \mathbb{Q}^n$, and a distance $d > 0$. It is a 'YES' instance if $\text{dist}(\mathbf{t}, \mathcal{L}) \leq d$ and a 'NO' instance if $\text{dist}(k\mathbf{t}, \mathcal{L}) > \gamma d$ for all non-zero integers k *and* $d < \alpha \cdot \lambda_1(\mathcal{L})$.

We will need the following characterization of the operator norm of a matrix in terms of its behavior over a lattice. Intuitively, this says that "a lattice has a point in every direction."

▶ **Fact 25.** *For any matrix $A \in \mathbb{R}^{n \times n}$ and (full-rank) lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\|A\| = \sup_{\mathbf{y} \in \mathcal{L} \setminus \{0\}} \frac{\|A\mathbf{y}\|}{\|\mathbf{y}\|} .$$

Proof. It suffices to note that, for any $\mathbf{x} \in \mathbb{R}^n$ with $\|\mathbf{x}\| = 1$ and any full-rank lattice $\mathcal{L} \subset \mathbb{R}^n$, there is a sequence $\mathbf{y}_1, \mathbf{y}_2, \ldots$ of vectors $\mathbf{y}_i \in \mathcal{L}$ such that

$$\lim_{m \to \infty} \frac{\mathbf{y}_m}{\|\mathbf{y}_m\|} = \mathbf{x} .$$

Indeed, this follows immediately from the fact that the rationals are dense in the reals. ◀

▶ **Theorem 26.** *For any $\gamma = \gamma(n) \geq 1$, there is an efficient reduction from γ'-GapCVP$^{1/\gamma'}$ to γ-GapLDP, where $\gamma' = O(\gamma)$.*

Proof. On input $\mathcal{L} \subset \mathbb{Q}^n$ with basis $(\mathbf{b}_1, \ldots, \mathbf{b}_n)$, $\mathbf{t} \in \mathbb{Q}^n$, and $d > 0$, the reduction behaves as follows. Let $\mathcal{L}_1 := \mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_n, r \cdot \mathbf{e}_{n+1})$ with $r > 0$ to be set in the analysis. Let $\mathcal{L}_2 := \mathcal{L}(\mathbf{b}_1, \ldots, \mathbf{b}_n, \mathbf{t} + r \cdot \mathbf{e}_{n+1})$. I.e.,

$$\mathcal{L}_1 = \mathcal{L}\begin{pmatrix} B & \mathbf{0} \\ 0 & r \end{pmatrix} \qquad \mathcal{L}_2 = \mathcal{L}\begin{pmatrix} B & \mathbf{t} \\ 0 & r \end{pmatrix} . \tag{3}$$

(Formally, we must embed the $\mathbf{b}_i$ and $\mathbf{t}$ in $\mathbb{Q}^{n+1}$ under the natural embedding, but we ignore this for simplicity.) The reduction then calls its γ-GapLDP oracle with input $\mathcal{L}_1$, $\mathcal{L}_2$, and $c > 0$ to be set in the analysis and outputs its response.

It is clear that the reduction runs in polynomial time. Suppose that $\text{dist}(\mathbf{t}, \mathcal{L}) \leq d$. We note that $\mathcal{L}_2$ does not change if we shift $\mathbf{t}$ by a lattice vector. So, we may assume without loss of generality that $\mathbf{0}$ is a closest lattice vector to $\mathbf{t}$ and therefore $\|\mathbf{t}\| \leq d$.

Let $B_1 := [\mathbf{b}_1, \ldots, \mathbf{b}_n, r \cdot \mathbf{e}_{n+1}]$ and $B_2 := [\mathbf{b}_1, \ldots, \mathbf{b}_n, \mathbf{t} + r \cdot \mathbf{e}_{n+1}]$ be the bases from the reduction. It suffices to show that $\kappa(B_2 B_1^{-1})$ is small. Indeed, for any $\mathbf{y} \in \mathcal{L}_1$, we can write $\mathbf{y} = (\mathbf{y}', kr)$ for some $k \in \mathbb{Z}$ and $\mathbf{y}' \in \mathcal{L}$. Then, we have

$$\|B_2 B_1^{-1} \mathbf{y}\| = \|(\mathbf{y}' + k\mathbf{t}, kr)\| \leq \|(\mathbf{y}', kr)\| + |k|\|\mathbf{t}\| \leq (1 + d/r)\|\mathbf{y}\| .$$

Similarly, $\|B_2B_1^{-1}\mathbf{y}\| \geq \|\mathbf{y}\| - |k|\|\mathbf{t}\| \geq (1-d/r)\|\mathbf{y}\|$. Therefore, by Fact 25, $\kappa(B_2B_1^{-1}) \leq (1+d/r)/(1-d/r)$. So, we take $c := (1+d/r)/(1-d/r)$, and the oracle will therefore output 'YES'.

Now, suppose $\text{dist}(z\mathbf{t}, \mathcal{L}) > 10\gamma d$ for all non-zero integers z, and $\lambda_1(\mathcal{L}) > 10\gamma d$. (I.e., we take $\gamma' = 10\gamma = O(\gamma)$.) Let A be a linear map with $A\mathcal{L}_1 = \mathcal{L}_2$. Note that A has determinant one, so that $\kappa(A) \geq \frac{\|A\mathbf{x}\|}{\|\mathbf{x}\|}$ for any $\mathbf{x} \in \mathbb{Q}^{n+1} \setminus \{\mathbf{0}\}$. We have that $A(\mathbf{0}, r) = (\mathbf{y}', kr)$ for some $\mathbf{y}' \in \mathcal{L} + k\mathbf{t}$ and $k \in \mathbb{Z}$. If $k \neq 0$, then $\|A(\mathbf{0}, r)\| \geq \text{dist}(k\mathbf{t}, \mathcal{L}) > 10\gamma d$. So, $\kappa(A) \geq \|A(\mathbf{0}, r)\|/r > 10\gamma d/r$.

If, on the other hand, $k = 0$, then $\mathbf{y}' \in \mathcal{L} \setminus \{\mathbf{0}\}$ and $\|A(\mathbf{0}, r)\| = \|(\mathbf{y}', 0)\| \geq \lambda_1(\mathcal{L}) > 10\gamma d$, so that we again have $\kappa(A) \geq \|A(\mathbf{0}, r)\|/r > 10\gamma d/r$. Taking $r = 2\gamma d$ gives $\kappa(A) > \gamma \cdot c$, so that the oracle will output 'NO', as needed. ◀

4.2 Hardness of This Variant of GapCVP

We recall the definition of (the decision version of) γ-GapSVP.

▶ **Definition 27.** For any $\gamma = \gamma(n) \geq 1$, γ-GapSVP is the promise problem defined as follows: The input is a lattice $\mathcal{L} \subset \mathbb{Q}^n$, and a distance $d > 0$. It is a 'YES' instance if $\lambda_1(\mathcal{L}) \leq d$ and a 'NO' instance if $\lambda_1(\mathcal{L}) > \gamma d$.

Haviv and Regev (building on work of Ajtai, Micciancio, and Khot [2, 22, 17]) proved the following strong hardness result for γ-GapSVP [14].

▶ **Theorem 28** ([14, Theorem 1.1])**.**

1. *γ-GapSVP is* NP*-hard under randomized polynomial-time reductions for any constant $\gamma \geq 1$. I.e., there is no randomized polynomial-time algorithm for γ-GapSVP unless* $\mathsf{NP} \subseteq \mathsf{RP}$.
2. *$2^{\log^{1-\varepsilon} n}$-GapSVP is* NP*-hard under randomized quasipolynomial-time reductions for any constant $\varepsilon > 0$. I.e., there is no randomized polynomial-time algorithm for $2^{\log^{1-\varepsilon} n}$-GapSVP unless* $\mathsf{NP} \subseteq \mathsf{RTIME}(2^{\text{polylog}(n)})$.
3. *$n^{c/\log\log n}$-GapSVP is* NP*-hard under randomized subexponential-time reductions for some universal constant $c > 0$. I.e., there is no randomized polynomial-time algorithm for $n^{c/\log\log n}$-GapSVP unless* $\mathsf{NP} \subseteq \mathsf{RSUBEXP} := \bigcap_{\delta>0} \mathsf{RTIME}(2^{n^\delta})$.

In particular, to prove Theorem 5, it suffices to reduce γ'-GapSVP to γ-CVP$^{1/\gamma}$ for $\gamma' = O(\gamma)$.

▶ **Theorem 29.** *For any $1 \leq \gamma = \gamma(n) \leq \text{poly}(n)$, there is an efficient reduction from γ'-GapSVP to γ-GapCVP$^{1/\gamma}$, where $\gamma' = \gamma \cdot (1 + o(1))$.*

Proof. Let p be a prime with $10\gamma n \leq p \leq 20\gamma n \leq \text{poly}(n)$. We take $\gamma' = \gamma \cdot (1 + o(1))$ so that

$$\gamma = \frac{\gamma'}{\sqrt{1 - \gamma'^2/(p-1)^2}}.$$

On input a basis $B := [\mathbf{b}_1, \ldots, \mathbf{b}_n]$ for a lattice $\mathcal{L} \subset \mathbb{Q}^n$, and $d > 0$, the reduction behaves as follows. For $i = 1, \ldots, n$, let $\mathcal{L}_i := \mathcal{L}(\mathbf{b}_1, \ldots, p\mathbf{b}_i, \ldots, \mathbf{b}_n)$ be "$\mathcal{L}$ with its ith basis vector multiplied by p." And, for all i and $1 \leq j < p$, let $\mathbf{t}_{i,j} := j\mathbf{b}_i + r\mathbf{e}_{n+1}$, with $r > 0$ to be set in the analysis. For each i, j, the reduction calls its γ-GapCVP$^{1/\gamma}$ oracle on input $\mathcal{L}_i$, $\mathbf{t}_{i,j}$, and $d' := \sqrt{d^2 + r^2}$. Finally, it outputs 'YES' if the oracle answered 'YES' for any query. Otherwise, it outputs 'NO'.

It is clear that the algorithm is efficient. Note that

$$\operatorname{dist}(j\mathbf{b}_i, \mathcal{L}_i) = \min\left\{ \left\| \sum_{\ell=1}^{n} a_\ell \mathbf{b}_\ell \right\| \ : \ a_\ell \in \mathbb{Z},\ a_i \equiv j \bmod p \right\} .$$

In particular, $\lambda_1(\mathcal{L}) = \min_{i,j} \operatorname{dist}(j\mathbf{b}_i, \mathcal{L}_i)$.

So, suppose $\lambda_1(\mathcal{L}) \le d$. Then, there must be some i, j such that $\operatorname{dist}(\mathbf{t}_{i,j}, \mathcal{L}_i)^2 \le r^2 + \lambda_1(\mathcal{L})^2 \le r^2 + d^2 = d'^2$. So, the oracle answers 'YES' at least once.

Now, suppose $\lambda_1(\mathcal{L}) > \gamma' d$. Since $\mathcal{L}_i \subset \mathcal{L}$, we have $\lambda_1(\mathcal{L}_i) \ge \lambda_1(\mathcal{L}) > \gamma' d$, and therefore $d < \lambda_1(\mathcal{L}_i)/\gamma' < \lambda_1(\mathcal{L}_i)/\gamma$, as needed. And, by the above observation, we have $\operatorname{dist}(j\mathbf{b}_i, \mathcal{L}_i) \ge \lambda_1(\mathcal{L}) > \gamma' d$ for all $1 \le i \le n$ and $1 \le j < p$. Furthermore, for any integer $1 \le z < p$, we have $\operatorname{dist}(zj\mathbf{b}_i, \mathcal{L}_i) = \operatorname{dist}((zj \bmod p) \cdot \mathbf{b}_i, \mathcal{L}_i) > \gamma' d$, where we have used the fact that p is prime so that $zj \not\equiv 0 \bmod p$. It follows that $\operatorname{dist}(z\mathbf{t}_{i,j}, \mathcal{L}_i) > \operatorname{dist}(zj\mathbf{b}_i, \mathcal{L}_i) > \gamma' d$. And, for $z \ge p$, it is trivially the case that $\operatorname{dist}(z\mathbf{t}_{i,j}, \mathcal{L}_i) \ge zr \ge pr$. Taking $r := \gamma' d/(p-1)$, we have that in both cases

$$\operatorname{dist}(z\mathbf{t}_{i,j}, \mathcal{L}_i) > \gamma' d = \frac{\gamma' d'}{\sqrt{1 - r^2}} = \frac{\gamma' d'}{\sqrt{1 - \gamma'^2/(p-1)^2}} = \gamma d .$$

So, the oracle will always answer 'NO'. ◀

▶ **Corollary 30.** *For any $1 \le \gamma = \gamma(n) \le \operatorname{poly}(n)$, there is an efficient reduction from γ'-GapSVP to γ-GapLDP, where $\gamma' = O(\gamma)$.*

Proof. Combine Theorems 26 and 29. ◀

With this, the proof of our main hardness result is immediate.

▶ **Theorem 31.** *The three hardness results in Theorem 28 hold with GapLDP in place of GapSVP.*

Proof. Combine Theorem 28 with Corollary 30. ◀

5 Some illustrative examples

5.1 Separating distortion from the successive minima

We now show that, for every n, there exists a $\mathcal{L}$ such that $\mathcal{D}(\mathcal{L}, \mathbb{Z}^n) \ge \Omega(\sqrt{n}) \cdot M(\mathcal{L}, \mathbb{Z}^n) \cdot M(\mathbb{Z}^n, \mathcal{L})$. Indeed, it suffices to take any lattice with $\det(\mathcal{L})^{1/n} \le O(n^{-1/2})$ but $\lambda_i(\mathcal{L}) = \Theta(1)$. (This is true for almost all lattices in a certain precise sense. See, e.g., [31].)

▶ **Lemma 32.** *For any $n \ge 1$, there is a lattice $\mathcal{L} \subset \mathbb{Q}^n$ such that $\det(\mathcal{L})^{1/n} \le O(n^{-1/2})$ and $\lambda_i(\mathcal{L}) = \Theta(1)$ for all i.*

▶ Proposition 33. For any $n \ge 1$, there exists a lattice $\mathcal{L} \subset \mathbb{Q}^n$ such that

$$\mathcal{D}(\mathcal{L}, \mathbb{Z}^n) \ge \Omega(\sqrt{n}) \cdot M(\mathcal{L}, \mathbb{Z}^n) \cdot M(\mathbb{Z}^n, \mathcal{L}) .$$

Proof. Let $\mathcal{L} \subset \mathbb{Q}^n$ be any lattice as in Lemma 32. In particular, $M(\mathcal{L}, \mathbb{Z}^n) \cdot M(\mathbb{Z}^n, \mathcal{L}) = O(1)$. However, for any linear map T with $T(\mathcal{L}) = \mathbb{Z}^n$, we of course have

$$\|T\| \ge |\det(T)|^{1/n} = \det(\mathbb{Z}^n)^{1/n} / \det(\mathcal{L})^{1/n} \ge \Omega(\sqrt{n}) .$$

(To see the first inequality, it suffices to recall that $|\det(T)| = \prod \sigma_i$ and $\|T\| = \max \sigma_i$, where the σ_i are the singular values of T.) And, $T^{-1}\mathbf{e}_1$ must be a non-zero lattice vector, so $\|T^{-1}\| \ge \|T^{-1}\mathbf{e}_1\| \ge \lambda_1(\mathcal{L}) \ge \Omega(1)$. Therefore, $\kappa(T) = \|T\|\|T^{-1}\| \ge \Omega(\sqrt{n})$, as needed. ◀

5.2 Non-optimality of HKZ bases for distortion

We show an example demonstrating that mappings between lattices built using HKZ bases are non-optimal in terms of their distortion. Namely, we give a family of $n \times n$ HKZ bases $\{B_n\}$ such that $\mathcal{D}(\mathbb{Z}^n, \mathcal{L}(B_n)) \leq n^{O(\log n)}$, but where the mapping $T = B_n$ from $\mathbb{Z}^n$ to $\mathcal{L}(B_n)$ has exponential distortion. This shows the necessity of using Seysen reduction in addition to HKZ reduction.

▶ **Theorem 34.** *For every $n \geq 1$, there exists an $n \times n$ HKZ basis B such that $\mathcal{D}(\mathbb{Z}^n, \mathcal{L}(B)) \leq n^{O(\log n)}$, but $\kappa(B) \geq \Omega(1.5^n)$.*

Recall that $\|B\|_\infty$ denotes the largest magnitude of an entry in B. It holds that $\|B\|_\infty \leq \|B\| \leq n \|B\|_\infty$.

▶ **Lemma 35.** *Let $B = B_n$ denote the $n \times n$ basis defined as*

$$b_{ij} = \begin{cases} 0 & \text{if } j < i, \\ 1 & \text{if } j = i, \\ -\frac{1}{2} & \text{if } j > i. \end{cases}$$

Then B is an HKZ basis and $\kappa(B) = \Omega(1.5^n)$.

Proof. For every basis A, it holds that $\min_{i \in [n]} \|\widetilde{\mathbf{a}}_i\| \leq \lambda_1(\mathcal{L}(A))$ (see, e.g., [27]). Note that for $i \geq 0$ the i^{th} Gram-Schmidt vector of $(\pi_k^{(B)}(\mathbf{b}_k), \ldots, \pi_k^{(B)}(\mathbf{b}_n))$ is simply $\tilde{b}_{i+k}$. Let $k \in [n]$. We then have that $1 = \min_{i \in [n]} \left\|\tilde{b}_i\right\| \leq \lambda_1(\pi_k(\mathcal{L}))$. On the other hand, $\lambda_1(\pi_k(\mathcal{L})) \leq \left\|\tilde{b}_k\right\| = \left\|\pi_k^{(B)}(\mathbf{b}_k)\right\| = 1$, implying that $\lambda_1(\pi_k(\mathcal{L}(B))) = 1$. It follows that B is an HKZ basis.

Because $\|B\|_\infty = 1$, it suffices to show that $\left\|B^{-1}\right\|_\infty \geq \Omega(1.5^n)$. Let $\mathbf{x}$ denote the nth column of B^{-1}. We must then have that $B\mathbf{x} = e_n$. Because B is upper triangular, we get the following formula by back substitution (see, e.g., [33]):

$$x_j = \begin{cases} 1 & \text{if } j = n, \\ \frac{1}{2} \sum_{k=j+1}^{n} x_k & \text{otherwise.} \end{cases} \tag{4}$$

We therefore have that $x_n = 1, x_{n-1} = \frac{1}{2}$. Using Eq. (4), we get that for $1 \leq m \leq n-2$,

$$x_m = \frac{1}{2} \sum_{k=m+1}^{n} x_k = \frac{1}{2} \cdot x_{m+1} + \frac{1}{2} \cdot \sum_{k=m+2}^{n} x_k = 1.5 \cdot x_{m+1} \, .$$

Applying this formula recursively, we get that $x_m = \frac{1}{2} \cdot 1.5^{n-m-1}$ for $1 \leq m \leq n-1$. ◀

The proof of Theorem 34 follows.

Proof of Theorem 34. Let $B' = B_n$ be an HKZ basis as specified in Lemma 35, and take I_n as the basis of $\mathbb{Z}_n$. Then $\kappa(B' \cdot I_n) = \Omega(1.5^n)$.

On the other hand, let $B = \textsc{Seysen}(B')$. Then, because $\eta(B') = 1$, $S(B) = n^{O(\log n)}$ by Theorem 20. Clearly, $\lambda_i(\mathbb{Z}_n) = 1$ for all $i \in [n]$. On the other hand, $1 \leq \lambda_i(\mathcal{L}(B)) \leq \sqrt{n}$ for all $i \in [n]$. The lower bound holds because $\min \left\|\tilde{b}_i\right\| = 1$, and the upper bound comes from the fact that $\|\mathbf{b}'_i\| \leq \sqrt{n}$ for all $i \in [n]$ and the linear independence of the $\mathbf{b}'_i$.[1] It follows that $M(\mathbb{Z}^n, \mathcal{L}(B)) \leq \sqrt{n}$ and $M(\mathcal{L}(B), \mathbb{Z}^n) \leq 1$. Applying Lemma 23 to B and B^{-1}, we then get that $\kappa(B \cdot I_n) \leq n^{O(\log n)}$. ◀

[1] In fact, $\lambda_n(\mathcal{L}(B)) = O(1)$.

Acknowledgements. We thank Oded Regev for pointing us to Seysen's paper and for many helpful conversations. The concise proof of Lemma 22 using the Transference Theorem is due to Michael Walter. We thank Paul Kirchner for suggesting the use of slide-reduced bases, and for identifying a minor bug in an earlier version of this paper.

References

1 Miklós Ajtai. Generating hard instances of lattice problems. In *STOC*, 1996.
2 Miklós Ajtai. The Shortest Vector Problem in L2 is NP-hard for randomized reductions. In *STOC*, 1998. doi:10.1145/276698.276705.
3 Miklós Ajtai, Ravi Kumar, and D. Sivakumar. A sieve algorithm for the Shortest Lattice Vector Problem. In *STOC*, pages 601–610, 2001. doi:10.1145/380752.380857.
4 Sanjeev Arora, Alan Frieze, and Haim Kaplan. A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. *Mathematical Programming*, 2002. doi:10.1007/s101070100271.
5 Vikraman Arvind, Johannes Köbler, Sebastian Kuhnert, and Yadu Vasudev. Approximate Graph Isomorphism. In *Mathematical Foundations of Computer Science*, 2012.
6 L. Babai. Graph Isomorphism in quasipolynomial time, 2016. http://arxiv.org/abs/1512.03547.
7 W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(1):625–635, 1993. doi:10.1007/BF01445125.
8 Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In *ITCS*, 2014. doi:10.1145/2554797.2554799.
9 J. Conway and N.J.A. Sloane. *Sphere Packings, Lattices and Groups*. Springer New York, 1998.
10 Nicolas Gama and Phong Q. Nguyen. Finding short lattice vectors within Mordell's inequality. In *STOC*, 2008. doi:10.1145/1374376.1374408.
11 Craig Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, 2009.
12 Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, 2008.
13 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. doi:10.1016/S0020-0190(99)00083-6.
14 Ishay Haviv and Oded Regev. Tensor-based hardness of the Shortest Vector Problem to within almost polynomial factors. *Theory of Computing*, 8(23):513–531, 2012. Preliminary version in STOC'07.
15 Ishay Haviv and Oded Regev. On the Lattice Isomorphism Problem. In *SODA*, 2014. doi:10.1137/1.9781611973402.29.
16 Ravi Kannan. Minkowski's convex body theorem and Integer Programming. *Mathematics of Operations Research*, 12(3):pp. 415–440, 1987. URL: http://www.jstor.org/stable/3689974.
17 Subhash Khot. Hardness of approximating the Shortest Vector Problem in lattices. *Journal of the ACM*, 52(5):789–808, September 2005. Preliminary version in FOCS'04.
18 A. Korkine and G. Zolotareff. Sur les formes quadratiques. *Mathematische Annalen*, 6(3):366–389, 1873. doi:10.1007/BF01442795.
19 J. C. Lagarias, Hendrik W. Lenstra Jr., and Claus-Peter Schnorr. Korkin-Zolotarev bases and successive minima of a lattice and its reciprocal lattice. *Combinatorica*, 10(4):333–348, 1990. doi:10.1007/BF02128669.
20 A.K. Lenstra, H.W. Lenstra Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, 1982. doi:10.1007/BF01457454.

21 Hendrik W. Lenstra Jr. and Alice Silverberg. Lattices with symmetry, 2014. http://arxiv.org/abs/1501.00178.

22 Daniele Micciancio. The Shortest Vector Problem is NP-hard to approximate to within some constant. *SIAM Journal on Computing*, 30(6):2008–2035, March 2001. Preliminary version in FOCS 1998.

23 Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal on Computing*, 37(1):267–302, 2007.

24 Daniele Micciancio and Panagiotis Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. *SIAM J. Comput.*, 42(3):1364–1391, 2013. doi:10.1137/100811970.

25 H. Minkowski. *Geometrie der Zahlen*. Number v. 1 in Geometrie der Zahlen. B.G. Teubner, 1910.

26 W. Plesken and B. Souvignier. Computing isometries of lattices. *J. Symbolic Comput.*, 24(3-4):327–334, 1997. Computational algebra and number theory (London, 1993).

27 Oded Regev. Lecture notes from course on "Lattices in Computer Science", 2009. URL: http://www.cims.nyu.edu/~regev/teaching/lattices_fall_2009/index.html.

28 Oded Regev. On lattices, Learning with Errors, random linear codes, and cryptography. *Journal of the ACM*, 56(6):Art. 34, 40, 2009. doi:10.1145/1568318.1568324.

29 C.P. Schnorr. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science*, 1987. doi:10.1016/0304-3975(87)90064-8.

30 Martin Seysen. Simultaneous reduction of a lattice basis and its reciprocal basis. *Combinatorica*, 13(3):363–376, 1993. doi:10.1007/BF01202355.

31 Carl Ludwig Siegel. A mean value theorem in geometry of numbers. *Annals of Mathematics*, 46(2):pp. 340–347, 1945. URL: http://www.jstor.org/stable/1969027.

32 Mathieu Dutour Sikiric, Achill Schürmann, and Frank Vallentin. Complexity and algorithms for computing Voronoi cells of lattices. *Math. Comput.*, 78(267):1713–1731, 2009. doi:10.1090/S0025-5718-09-02224-8.

33 Lloyd N. Trefethen and David Bau III. *Numerical Linear Algebra*. SIAM: Society for Industrial and Applied Mathematics, 1997. ISBN 0898713617.