

Hardness of Bounded Distance Decoding on Lattices in ℓ_p Norms

Huck Bennett*

Chris Peikert†

March 17, 2020

Abstract

Bounded Distance Decoding $\text{BDD}_{p,\alpha}$ is the problem of decoding a lattice when the target point is promised to be within an α factor of the minimum distance of the lattice, in the ℓ_p norm. We prove that $\text{BDD}_{p,\alpha}$ is NP-hard under randomized reductions where $\alpha \rightarrow 1/2$ as $p \rightarrow \infty$ (and for $\alpha = 1/2$ when $p = \infty$), thereby showing the hardness of decoding for distances approaching the unique-decoding radius for large p . We also show *fine-grained* hardness for $\text{BDD}_{p,\alpha}$. For example, we prove that for all $p \in [1, \infty) \setminus 2\mathbb{Z}$ and constants $C > 1, \varepsilon > 0$, there is no $2^{(1-\varepsilon)n/C}$ -time algorithm for $\text{BDD}_{p,\alpha}$ for some constant α (which approaches $1/2$ as $p \rightarrow \infty$), assuming the randomized Strong Exponential Time Hypothesis (SETH). Moreover, essentially all of our results also hold (under analogous non-uniform assumptions) for BDD with *preprocessing*, in which unbounded precomputation can be applied to the lattice before the target is available.

Compared to prior work on the hardness of $\text{BDD}_{p,\alpha}$ by Liu, Lyubashevsky, and Micciancio (APPROX-RANDOM 2008), our results improve the values of α for which the problem is known to be NP-hard for all $p > p_1 \approx 4.2773$, and give the very first fine-grained hardness for BDD (in any norm). Our reductions rely on a special family of “locally dense” lattices in ℓ_p norms, which we construct by modifying the integer-lattice sparsification technique of Aggarwal and Stephens-Davidowitz (STOC 2018).

1 Introduction

Lattices in \mathbb{R}^n are a rich source of computational problems with applications across computer science, and especially in cryptography and cryptanalysis. (A lattice is a discrete additive subgroup of \mathbb{R}^n , or equivalently, the set of integer linear combinations of a set of linearly independent vectors.) Many important lattice problems appear intractable, and there is a wealth of research showing that central problems like the Shortest Vector Problem (SVP) and Closest Vector Problem (CVP) are NP-hard, even to approximate to within various factors and in various ℓ_p norms [vEB81, ABSS97, Ajt98, Mic98, Mic01, Kho03, Kho04, HR07, Mic12]. (For the sake of concision, throughout this introduction the term “NP-hard” allows for *randomized* reductions, which are needed in some important cases.)

Bounded Distance Decoding. In recent years, the emergence of lattices as a powerful foundation for cryptography, including for security against quantum attacks, has increased the importance of other lattice problems. In particular, many modern lattice-based encryption schemes rely on some form of the *Bounded Distance Decoding* (BDD) problem, which is like the Closest Vector Problem with a promise. An instance

*University of Michigan, hdbco@umich.edu.

†University of Michigan, cpeikert@umich.edu

of BDD_α for *relative distance* $\alpha > 0$ is a lattice \mathcal{L} and a target point t whose distance from the lattice is guaranteed to be within an α factor of the lattice’s minimum distance $\lambda_1(\mathcal{L}) = \min_{v \in \mathcal{L} \setminus \{0\}} \|v\|$, and the goal is to find a lattice vector within that distance of t ; when distances are measured in the ℓ_p norm we denote the problem $\text{BDD}_{p,\alpha}$. Note that when $\alpha < 1/2$ there is a unique solution, but the problem is interesting and well-defined for larger relative distances as well. We also consider *preprocessing* variants of CVP and BDD (respectively denoted CVPP and BDDP), in which unbounded precomputation can be applied to the lattice before the target is available. For example, this can model cryptographic contexts where a fixed long-term lattice may be shared among many users.

The importance of BDD(P) to cryptography is especially highlighted by the Learning With Errors (LWE) problem of Regev [Reg05], which is an average-case form of BDD that has been used (with inverse-polynomial α) in countless cryptosystems, including several that share a lattice among many users (see, e.g., [GPV08]). Moreover, Regev gave a worst-case to average-case reduction from BDD to LWE, so the security of cryptosystems is intimately related to the worst-case complexity of BDD.

Compared to problems like SVP and CVP, the BDD(P) problem has received much less attention from a complexity-theoretic perspective. We are aware of essentially only one work showing its NP-hardness: Liu, Lyubashevsky, and Micciancio [LLM06] proved that $\text{BDD}_{p,\alpha}$ and even $\text{BDDP}_{p,\alpha}$ are NP-hard for relative distances approaching $\min\{1/\sqrt{2}, 1/\sqrt[p]{2}\}$, which is $1/\sqrt{2}$ for $p \geq 2$. A few other works relate BDD(P) to other lattice problems (in both directions) in regimes where the problems are not believed to be NP-hard, e.g., [Mic08, DRS14, BSW16]. (Dadush, Regev, and Stephens-Davidowitz [DRS14] also gave a reduction that implies NP-hardness of $\text{BDD}_{2,\alpha}$ for any $\alpha > 1$, which is larger than the relative distance of $\alpha = 1/\sqrt{2} + \varepsilon$ achieved by [LLM06].)

Fine-grained hardness. An important aspect of hard lattice problems, especially for cryptography, is their *quantitative* hardness. That is, we want not only that a problem cannot be solved in polynomial time, but that it cannot be solved in, say, $2^{o(n)}$ time or even $2^{n/C}$ time for a certain constant C . Statements of this kind can be proven under generic complexity assumptions like the Exponential Time Hypothesis (ETH) of Impagliazzo and Paturi [IP01] or its variants like Strong ETH (SETH), via *fine-grained* reductions that are particularly efficient in the relevant parameters.

Recently, Bennett, Golovnev, and Stephens-Davidowitz [BGS17] initiated a study of the fine-grained hardness of lattice problems, focusing on CVP; follow-up work extended to SVP and showed more for CVP(P) [AS18a, ABGS19]. The technical goal of these works is a reduction having good *rank efficiency*, i.e., a reduction from k -SAT on n' variables to a lattice problem in rank $n = (C + o(1))n'$ for some constant $C \geq 1$, which we call the reduction’s “rank inefficiency.” (All of the lattice problems in question can be solved in $2^{n+o(n)}$ time [ADRS15, ADS15, AS18b], so $C = 1$ corresponds to optimal rank efficiency.) We mention that Regev’s BDD-to-LWE reduction [Reg05] has optimal rank efficiency, in that it reduces rank- n BDD to rank- n LWE. However, to date there are no fine-grained NP-hardness results for BDD itself; the prior NP-hardness proof for BDD [LLM06] incurs a large polynomial blowup in rank.

1.1 Our Results

We show improved NP-hardness, and entirely new fine-grained hardness, for Bounded Distance Decoding (and BDD with preprocessing) in arbitrary ℓ_p norms. Our work improves upon the known hardness of BDD in two respects: the relative distance α , and the rank inefficiency C (i.e., fine-grainedness) of the reductions. As p grows, both quantities improve, simultaneously approaching the unique-decoding threshold $\alpha = 1/2$ and optimal rank efficiency of $C = 1$ as $p \rightarrow \infty$, and achieving those quantities for $p = \infty$. We emphasize that these are the first fine-grained hardness results of any kind for BDD, for any ℓ_p norm.

Our main theorem summarizing the NP- and fine-grained hardness of BDD (with and without preprocessing) appears below in Theorem 1.1. For $p \in [1, \infty)$ and $C > 1$, the quantities α_p^* and $\alpha_{p,C}^*$ appearing in the theorem statement are certain positive real numbers that are decreasing in p and C , and approaching $1/2$ as $p \rightarrow \infty$ (for any C). See Figure 1 for a plot of their behavior, Equations (3.4) and (3.5) for their formal definitions, and Lemma 3.11 for quite tight closed-form upper bounds.

Theorem 1.1. *The following hold for $\text{BDD}_{p,\alpha}$ and $\text{BDDP}_{p,\alpha}$ in rank n :*

1. *For every $p \in [1, \infty)$ and constant $\alpha > \alpha_p^*$ (where $\alpha_p^* \leq \frac{1}{2} \cdot 4.6723^{1/p}$), and for $(p, \alpha) = (\infty, 1/2)$, there is no polynomial-time algorithm for $\text{BDD}_{p,\alpha}$ (respectively, $\text{BDDP}_{p,\alpha}$) unless $\text{NP} \subseteq \text{RP}$ (resp., $\text{NP} \subseteq \text{P/Poly}$).*
2. *For every $p \in [1, \infty)$ and constant $\alpha > \min\{\alpha_p^*, \alpha_2^*\}$, and for $(p, \alpha) = (\infty, 1/2)$, there is no $2^{o(n)}$ -time algorithm for $\text{BDD}_{p,\alpha}$ unless randomized ETH fails.*
3. *For every $p \in [1, \infty) \setminus \{2\}$ and constant $\alpha > \alpha_p^*$, and for $(p, \alpha) = (\infty, 1/2)$, there is no $2^{o(n)}$ -time algorithm for $\text{BDDP}_{p,\alpha}$ unless non-uniform ETH fails.*

Moreover, for every $p \in [1, \infty]$ and $\alpha > \alpha_2^$ there is no $2^{o(\sqrt{n})}$ -time algorithm for $\text{BDDP}_{p,\alpha}$ unless non-uniform ETH fails.*

4. *For every $p \in [1, \infty) \setminus 2\mathbb{Z}$ and constants $C > 1$, $\alpha > \alpha_{p,C}^*$, and $\epsilon > 0$, and for $(p, C, \alpha) = (\infty, 1, 1/2)$, there is no $2^{n(1-\epsilon)/C}$ -time algorithm for $\text{BDD}_{p,\alpha}$ (respectively, $\text{BDDP}_{p,\alpha}$) unless randomized SETH (resp., non-uniform SETH) fails.*

Although we do not have closed-form expressions for α_p^* and $\alpha_{p,C}^*$, we do get quite tight closed-form upper bounds (see Lemma 3.11). Moreover, it is easy to numerically compute close approximations to them, and to the values of p at which they cross certain thresholds. For example, $\alpha_p^* < 1/\sqrt{2}$ for all $p > p_1 \approx 4.2773$, so Item 1 of Theorem 1.1 improves on the prior best relative distance of any $\alpha > 1/\sqrt{2}$ for the NP-hardness of $\text{BDD}_{p,\alpha}$ in such ℓ_p norms [LLM06].

As a few other example values and their consequences under Theorem 1.1, we have $\alpha_2^* \approx 1.05006$, $\alpha_{3,2}^* \approx 1.1418$, and $\alpha_{3,5}^* \approx 0.917803$. So by Item 2, BDD in the Euclidean norm for any relative distance $\alpha > 1.05006$ requires $2^{\Omega(n)}$ time assuming randomized ETH. And by Item 4, for every $\epsilon > 0$ there is no $2^{(1-\epsilon)n/2}$ -time algorithm for $\text{BDD}_{3,1.1418}$, and no $2^{(1-\epsilon)n/5}$ -time algorithm for $\text{BDD}_{3,0.917803}$, assuming randomized SETH.

1.2 Technical Overview

As in prior NP-hardness reductions for SVP and BDD (and fine-grained hardness proofs for the former) [Ajt98, Mic98, Kho04, LLM06, HR07, Mic12, AS18a], the central component of our reductions is a family of rank- n lattices $\mathcal{L} \subset \mathbb{R}^d$ and target points $\mathbf{t} \in \mathbb{R}^d$ having a certain “local density” property in a desired ℓ_p norm. Informally, this means that \mathcal{L} has “large” minimum distance $\lambda_1^{(p)}(\mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L} \setminus \{0\}} \|\mathbf{v}\|_p$, i.e., there are no “short” nonzero vectors, but has many vectors “close” to the target \mathbf{t} . More precisely, we want $\lambda_1^{(p)}(\mathcal{L}) \geq r$ and $N_p(\mathcal{L}, \alpha r, \mathbf{t}) = \exp(n^{\Omega(1)})$ for some relative distance α , where

$$N_p(\mathcal{L}, s, \mathbf{t}) := |\{\mathbf{v} \in \mathcal{L} : \|\mathbf{v} - \mathbf{t}\|_p \leq s\}|$$

denotes the number of lattice points within distance s of \mathbf{t} .




Figure 1: Left: bounds on the relative distances $\alpha = \alpha(p)$ for which $\text{BDD}_{\alpha,p}$ was proved to be NP-hard in the ℓ_p norm, in this work and in [LLM06]; the crossover point is $p_1 \approx 4.2773$. (The plots include results obtained by norm embeddings [RR06], hence they are maximized at $p = 2$.) Right: our bounds $\alpha_{p,C}^*$ on the relative distances $\alpha > \alpha_{p,C}^*$ for which there is no $2^{(1-\varepsilon)n/C}$ -time algorithm for $\text{BDD}_{p,\alpha}$ for any $\varepsilon > 0$, assuming randomized SETH.

Micciancio [Mic98] constructed locally dense lattices with relative distance approaching $2^{-1/p}$ in the ℓ_p norm (for every finite $p \geq 1$), and used them to prove the NP-hardness of γ -approximate SVP in ℓ_p for any $\gamma < 2^{1/p}$. Subsequently, Liu, Lyubashevsky, and Micciancio [LLM06] used these lattices to prove the NP-hardness of BDD in ℓ_p for any relative distance $\alpha > 2^{-1/p}$. However, these works observed that the relative distance depends on p in the opposite way from what one might expect: as p grows, so does α , hence the associated NP-hard SVP approximation factors and BDD relative distances *worsen*. Yet using norm embeddings, it can be shown that ℓ_2 is essentially the “easiest” ℓ_p norm for lattice problems [RR06], so hardness in ℓ_2 implies hardness in ℓ_p (up to an arbitrarily small loss in approximation factor). Therefore, the locally dense lattices from [Mic98] do not seem to provide any benefits for $p > 2$ over $p = 2$, where the relative distance approaches $1/\sqrt{2}$. In addition, the rank of these lattices is a large polynomial in the relevant parameter, so they are not suitable for proving fine-grained hardness.¹

Local density via sparsification. More recently, Aggarwal and Stephens-Davidowitz [AS18a] (building on [BGS17]) proved fine-grained hardness for *exact* SVP in ℓ_p norms, via locally dense lattices obtained in a different way. Because they target exact SVP, it suffices to have local density for relative distance $\alpha = 1$, but for fine-grained hardness they need $N_p(\mathcal{L}, r, \mathbf{t}) = 2^{\Omega(n)}$, preferably with a large hidden constant (which determines the rank efficiency of the reduction). Following [MO90, EOR91], they start with the integer lattice \mathbb{Z}^n and all- $\frac{1}{2}$ s target vector $\mathbf{t} = \frac{1}{2}\mathbf{1} \in \mathbb{R}^n$. Clearly, there are 2^n lattice vectors all at distance $r = \frac{1}{2}n^{1/p}$ from \mathbf{t} in the ℓ_p norm, but the minimum distance of the lattice is only 1, so the relative distance of the “close” vectors is $\alpha = r$, which is far too large.

To improve the relative distance, they increase the minimum distance to at least $r = \frac{1}{2}n^{1/p}$ using the elegant technique of *random sparsification*, which is implicit in [EOR91] and was first used for proving NP-hardness of approximate SVP in [Kho03, Kho04]. The idea is to upper-bound the number $N_p(\mathbb{Z}^n, r, \mathbf{0})$ of “short” lattice points of length at most r , by some Q . Then, by taking a random sublattice $\mathcal{L} \subset \mathbb{Z}^n$ of determinant (index) slightly larger than Q , with noticeable probability none of the “short” nonzero vectors will be included in \mathcal{L} , whereas roughly $2^n/Q$ of the vectors “close” to \mathbf{t} will be in \mathcal{L} . So, as long as $Q = 2^{(1-\Omega(1))n}$, there are sufficiently many lattice vectors at the desired relative distance from \mathbf{t} .

¹We mention that Khot [Kho04] gave a different construction of locally dense lattices with other useful properties, but their relative distance is no smaller than that of Micciancio’s construction in any ℓ_p norm, and their rank is also a large polynomial in the relevant parameter.

Bounds for $N_p(\mathbb{Z}^n, r, \mathbf{0})$ were given by Mazo and Odlyzko [MO90], by a simple but powerful technique using the theta function $\Theta_p(\tau) := \sum_{z \in \mathbb{Z}^n} \exp(-\tau|z|^p)$. They showed (see Proposition 2.12) that

$$N_p(\mathbb{Z}^n, r, \mathbf{0}) \leq \min_{\tau > 0} \exp(\tau \cdot r^p) \cdot \Theta_p(\tau)^n = \left(\min_{\tau > 0} \exp(\tau/2^p) \cdot \Theta_p(\tau) \right)^n, \quad (1.1)$$

where the equality is by $r = \frac{1}{2}n^{1/p}$. So, Aggarwal and Stephens-Davidowitz need $\min_{\tau > 0} \exp(\tau/2^p) \cdot \Theta_p(\tau) < 2$, and it turns out that this is the case for every $p > p_0 \approx 2.1397$. (They also deal with smaller p by using a different target point \mathbf{t} .)

This work: local density for small relative distance. For the NP- and fine-grained hardness of BDD we use the same basic approach as in [AS18a], but with the different goal of getting local density for as small of a relative distance $\alpha < 1$ as we can manage. That is, we still have 2^n integral vectors all at distance $r = \frac{1}{2}n^{1/p}$ from the target $\mathbf{t} = \frac{1}{2}\mathbf{1} \in \mathbb{R}^n$, but we want to “sparsify away” all the nonzero integral vectors of length less than r/α . So, we want the right-hand side of the Mazo-Odlyzko bound (Equation (1.1)) to be at most $2^{(1-\Omega(1))n}$ for as large of a positive hidden constant as we can manage. More specifically, for any $p \geq 1$ and $C > 1$ (which ultimately corresponds to the reduction’s rank inefficiency) we can obtain local density of at least $2^{n/C}$ close vectors at any relative distance greater than

$$\alpha_{p,C}^* := \inf\{\alpha^* > 0 : \min_{\tau > 0} \exp(\tau/(2\alpha^*)^p) \cdot \Theta_p(\tau) \leq 2^{1-1/C}\}.$$

The value of $\alpha_{p,C}^*$ is strictly decreasing in both p and C , and for large C and $p > p_1 \approx 4.2773$ it drops below the relative distance of $1/\sqrt{2}$ approached by the local-density construction of [Mic98] for ℓ_2 (and also ℓ_p by norm embeddings.) This is the source of our improved relative distance for the NP-hardness of BDD in high ℓ_p norms.

We also show that obtaining local density by sparsifying the integer lattice happens to yield a very simple reduction to BDD from the *exact* version of CVP, which is how we obtain fine-grained hardness. Given a CVP instance consisting of a lattice and a target point, we essentially just take their direct sum with the integer lattice and the $\frac{1}{2}\mathbf{1}$ target (respectively), then sparsify. (See Lemma 3.5 and Theorem 3.3 for details.) Because this results in the (sparsified) locally dense lattice having $2^{\Omega(n)}$ close vectors all *exactly* at the threshold of the BDD promise, concatenating the CVP instance either keeps the target within the (slightly weaker) BDD promise, or puts it just outside. This is in contrast to the prior reduction of [LLM06], where the close vectors in the locally dense lattices of [Mic98] are at various distances from the target, hence a reduction from *approximate*-CVP with a large constant factor is needed to put the target outside the BDD promise. While approximating CVP to within any constant factor is known to be NP-hard [ABSS97], no fine-grained hardness is known for approximate CVP, except for factors just slightly larger than one [ABGS19].

1.3 Discussion and Future Work

Our work raises a number of interesting issues and directions for future research. First, it highlights that there are now two incomparable approaches for obtaining local density in the ℓ_p norm—Micciancio’s construction [Mic98], and sparsifying the integer lattice [EOR91, AS18a]—with each delivering a better relative distance for certain ranges of p . For $p \in [1, p_1 \approx 4.2773]$, Micciancio’s construction (with norm embeddings from ℓ_2 , where applicable) delivers the better relative distance, which approaches $\min\{1/\sqrt[p]{2}, 1/\sqrt{2}\}$. Moreover, this is essentially optimal in ℓ_2 , where $1/\sqrt{2}$ is unachievable due to the Rankin bound, which says that in \mathbb{R}^n we can have at most $2n$ subunit vectors with pairwise distances of $\sqrt{2}$ or more.

A first question, therefore, is whether relative distance less than $1/\sqrt{2}$ can be obtained for all $p > 2$. We conjecture that this is true, but can only manage to prove it via sparsification for all $p > p_1 \approx 4.2773$. More generally, an important open problem is to give a unified local-density construction that subsumes both of the above-mentioned approaches in terms of relative distance, and ideally in rank efficiency as well. In the other direction, another important goal is to give lower bounds on the relative distance in general ℓ_p norms. Apart from the Rankin bound, the only bound we are aware of is the trivial one of $\alpha \geq 1/2$ implied by the triangle inequality, which is essentially tight for ℓ_1 and tight for ℓ_∞ (as shown by [Mic98] and our work, respectively).

More broadly, for the BDD relative distance parameter α there are three regimes of interest: the local-density regime, where we know how to prove NP-hardness; the unique-decoding regime $\alpha < 1/2$; and (at least in some ℓ_p norms, including ℓ_2) the intermediate regime between them. It would be very interesting, and would seem to require new techniques, to show NP-hardness outside the local-density regime. One potential route would be to devise a *gap amplification* technique for BDD, analogous to how SVP has been proved to be NP-hard to approximate to within any constant factor [Kho04, HR07, Mic12]. Gap amplification may also be interesting in the absence of NP-hardness, e.g., for the inverse-polynomial relative distances used in cryptography. Currently, the only efficient gap amplification we are aware of is a modest one that decreases the relative distance by any $(1 - 1/n)^{O(1)}$ factor [LM09].

A final interesting research direction is related to the *unique* Shortest Vector Problem (uSVP), where the goal is to find a shortest nonzero vector v in a given lattice, under the promise that it is unique (up to sign). More generally, approximate uSVP has the promise that all lattice vectors not parallel to v are a certain factor γ as long. It is known that *exact* uSVP is NP-hard in ℓ_2 [KS01], and by known reductions it is straightforward to show the NP-hardness of *2-approximate* uSVP in ℓ_∞ . Can recent techniques help to prove NP-hardness of γ -approximate uSVP, for some constant $\gamma > 1$, in ℓ_p for some finite p , or specifically for ℓ_2 ? Do NP-hard approximation factors for uSVP grow smoothly with p ?

Acknowledgments. We thank Noah Stephens-Davidowitz for sharing his plot-generating code from [AS18a] with us.

2 Preliminaries

For any positive integer q , we identify the quotient group $\mathbb{Z}_q = \mathbb{Z}/q\mathbb{Z}$ with some set of distinguished representatives, e.g., $\{0, 1, \dots, q-1\}$. Let $B^+ := (B^t B)^{-1} B^t$ denote the Moore-Penrose pseudoinverse of a real-valued matrix B with full column rank. Observe that $B^+ v$ is the unique coefficient vector c with respect to B of any $v = Bc$ in the column span of B .

2.1 Problems with Preprocessing

In addition to ordinary computational problems, we are also interested in (promise) problems with *preprocessing*. In such a problem, an instance (x_P, x_Q) is comprised of a “preprocessing” part x_P and a “query” part x_Q , and an algorithm is allowed to perform unbounded computation on the preprocessing part before receiving the query part.

Formally, a preprocessing problem is a relation $\Pi = \{(x_P, x_Q), y\}$ of instance-solution pairs, where $\Pi_{\text{inst}} := \{(x_P, x_Q) : \exists y \text{ s.t. } ((x_P, x_Q), y) \in \Pi\}$ is the set of problem instances, and $\Pi_{(x_P, x_Q)} := \{y : ((x_P, x_Q), y) \in \Pi\}$ is the set of solutions for any particular instance (x_P, x_Q) . If every instance $(x_P, x_Q) \in \Pi_{\text{inst}}$ has exactly one solution that is either YES or NO, then Π is called a decision problem.

Definition 2.1. A *preprocessing algorithm* is a pair (P, Q) where P is a (possibly randomized) function representing potentially unbounded computation, and Q is an algorithm. The execution of (P, Q) on an input (x_P, x_Q) proceeds in two phases:

- first, in the preprocessing phase, P takes x_P as input and produces some preprocessed output σ ;
- then, in the query phase, Q takes both σ and x_Q as input and produces some ultimate output.

The running time T of the algorithm is defined to be the time used in the query phase alone, and is considered as a function of the total input length $|x_P| + |x_Q|$. The length of the preprocessed output is defined as $A = |\sigma|$, and is also considered as a function of the total input length. Note that without loss of generality, $A \leq T$.

If (P, Q) is deterministic, we say that it solves preprocessing problem Π if $Q(P(x_P), x_Q) \in \Pi_{(x_P, x_Q)}$ for all $(x_P, x_Q) \in \Pi_{\text{inst}}$. If (P, Q) is potentially randomized, we say that it solves Π if

$$\Pr[Q(P(x_P), x_Q) \in \Pi_{(x_P, x_Q)}] \geq \frac{2}{3}$$

for all $(x_P, x_Q) \in \Pi_{\text{inst}}$, where the probability is taken over the random coins of both P and Q .²

As shown below using a routine quantifier-swapping argument (as in Adleman's Theorem [Adl78]), it turns out that for NP relations and decision problems, any randomized preprocessing algorithm can be derandomized if the length of the query input x_Q is polynomial in the length of the preprocessing input x_P . So for convenience, in this work we allow for randomized algorithms, only switching to deterministic ones for our ultimate hardness theorems.

Lemma 2.2. *Let preprocessing problem Π be an NP relation or a decision problem for which $|x_Q| = \text{poly}(|x_P|)$ for all $(x_P, x_Q) \in \Pi_{\text{inst}}$. If Π has a randomized T -time algorithm, then it has a deterministic $T \cdot \text{poly}(|x_P| + |x_Q|)$ -time algorithm with $T \cdot \text{poly}(|x_P| + |x_Q|)$ -length preprocessed output.*

Proof. Let $q(\cdot)$ be a polynomial for which $|x_Q| \leq q(|x_P|)$ for all $(x_P, x_Q) \in \Pi_{\text{inst}}$. Let (P, Q) be a randomized T -time algorithm for Π , which by standard repetition techniques we can assume has probability strictly less than $\exp(-q(|x_P|))$ of being incorrect on any $(x_P, x_Q) \in \Pi_{\text{inst}}$, with only a $\text{poly}(|x_P| + |x_Q|)$ -factor overhead in the running time and preprocessed output length. Fix some arbitrary x_P . Then by the union bound over all $(x_P, x_Q) \in \Pi_{\text{inst}}$ and the hypothesis, we have

$$\Pr[\exists (x_P, x_Q) \in \Pi_{\text{inst}} : Q(P(x_P), x_Q) \notin \Pi_{(x_P, x_Q)}] < 1.$$

So, there exist coins for P and Q for which $Q(P(x_P), x_Q) \in \Pi_{(x_P, x_Q)}$ for all $(x_P, x_Q) \in \Pi_{\text{inst}}$. By fixing these coins we make P a deterministic function of x_P , and we include the coins for Q along with the preprocessed output $P(x_P)$, thus making Q deterministic as well. The resulting deterministic algorithm solves Π with the claimed resources, as needed. \square

Reductions for preprocessing problems. We need the following notions of reductions for preprocessing problems. The following generalizes Turing reductions and Cook reductions (i.e., polynomial-time Turing reductions).

²Note that it could be the case that some preprocessed outputs fail to make the query algorithm output a correct answer on some, or even all, query inputs.

Definition 2.3. A *Turing reduction* from one preprocessing problem X to another one Y is a pair of algorithms (R_P, R_Q) satisfying the following properties: R_P is a (potentially randomized) function with access to an oracle P , whose output length is polynomial in its input length; R_Q is an algorithm with access to an oracle Q ; and if (P, Q) solves problem Y , then (R_P, R_Q) solves problem X . Additionally, it is a *Cook reduction* if R_Q runs in time polynomial in the total input length of R_P and R_Q .

Similarly, the following generalizes mapping reductions and Karp reductions (i.e., polynomial-time mapping reductions) for decision problems.

Definition 2.4. A *mapping reduction* from one preprocessing decision problem X to another one Y is a pair (R_P, R_Q) satisfying the following properties: R_P is a deterministic function whose output length is polynomial in its input length; R_Q is a deterministic algorithm; and for any YES (respectively, NO) instance (x_P, x_Q) of X , the output pair (y_P, y_Q) is a YES (resp., NO) instance of Y , where (y_P, y_Q) are defined as follows:

- first, R_P takes x_P as input and outputs some (σ', y_P) , where σ' is some “internal” preprocessed output;
- then, R_Q takes (σ', x_Q) as input and outputs some y_Q .

Additionally, it is a *Karp reduction* if R_Q runs in time polynomial in the total input length of R_P and R_Q .

It is straightforward to see that if X mapping reduces to Y , and there is a deterministic polynomial-time preprocessing algorithm (P_Y, Q_Y) that solves Y , then there is also one (P_X, Q_X) that solves X , which works as follows:

1. the preprocessing algorithm P_X , given a preprocessing input x_p , first computes $(\sigma', y_P) = R_P(x_p)$, then computes $\sigma_Y = P_Y(y_P)$ and outputs $\sigma_X = (\sigma', \sigma_Y)$;
2. the query algorithm Q_X , given $\sigma_X = (\sigma', \sigma_Y)$ and a query input x_Q , computes $y_Q = R_Q(\sigma', x_Q)$ and finally outputs $Q_Y(\sigma_Y, y_Q)$.

2.2 Lattices

A *lattice* is the set of all integer linear combinations of some linearly independent vectors $\mathbf{b}_1, \dots, \mathbf{b}_n$. It is convenient to arrange these vectors as the columns of a matrix. Accordingly, we define a *basis* $B = (\mathbf{b}_1, \dots, \mathbf{b}_n) \in \mathbb{R}^{d \times n}$ to be a matrix with linearly independent columns, and the lattice generated by basis B as

$$\mathcal{L}(B) := \left\{ \sum_{i=1}^n a_i \mathbf{b}_i : a_1, \dots, a_n \in \mathbb{Z} \right\}.$$

Let \mathcal{B}_p^d denote the centered unit ℓ_p ball in d dimensions. Given a lattice $\mathcal{L} \subset \mathbb{R}^d$ of rank n , for $1 \leq i \leq n$ let

$$\lambda_i^{(p)}(\mathcal{L}) := \inf\{r > 0 : \dim(\text{span}(r \cdot \mathcal{B}_p^d \cap \mathcal{L})) \geq i\}$$

denote the i th successive minimum of \mathcal{L} with respect to the ℓ_p norm.

We denote the ℓ_p distance of a vector \mathbf{t} to a lattice \mathcal{L} as

$$\text{dist}_p(\mathbf{t}, \mathcal{L}) := \min_{\mathbf{v} \in \mathcal{L}} \|\mathbf{v} - \mathbf{t}\|_p.$$

2.3 Bounded Distance Decoding (with Preprocessing)

The primary computational problem that we study in this work is the Bounded Distance Decoding Problem (BDD), which is a version of the Closest Vector Problem (CVP) in which the target vector is promised to be relatively close to the lattice.

Definition 2.5. For $1 \leq p \leq \infty$ and $\alpha = \alpha(n) > 0$, the α -Bounded Distance Decoding problem in the ℓ_p norm ($\text{BDD}_{p,\alpha}$) is the (search) promise problem defined as follows. The input is (a basis of) a rank- n lattice \mathcal{L} and a target vector \mathbf{t} satisfying $\text{dist}_p(\mathbf{t}, \mathcal{L}) \leq \alpha(n) \cdot \lambda_1^{(p)}(\mathcal{L})$. The goal is to output a lattice vector $\mathbf{v} \in \mathcal{L}$ that satisfies $\|\mathbf{v} - \mathbf{t}\|_p \leq \alpha(n) \cdot \lambda_1^{(p)}(\mathcal{L})$.

The preprocessing (search) promise problem $\text{BDDP}_{p,\alpha}$ is defined analogously, where the preprocessing input is (a basis of) the lattice, and the query input is the target \mathbf{t} .

We note that in some works, BDD is defined to have the goal of finding a $\mathbf{v} \in \mathcal{L}$ such that $\|\mathbf{v} - \mathbf{t}\|_p = \text{dist}_p(\mathbf{t}, \mathcal{L})$. This formulation is clearly no easier than the one defined above. So, our hardness theorems, which are proved for the definition above, immediately apply to the alternative formulation as well.

We also remark that for $\alpha < 1/2$, the promise ensures that there is a *unique* vector \mathbf{v} satisfying $\|\mathbf{v} - \mathbf{t}\|_p \leq \alpha \cdot \lambda_1^{(p)}(\mathcal{L})$. However, BDD is still well defined for $\alpha \geq 1/2$, i.e., above the unique-decoding radius. As in prior work, our hardness results for $\text{BDD}_{p,\alpha}$ are limited to this regime.

To the best of our knowledge, essentially the only previous study of the NP-hardness of BDD is due to [LLM06], which showed the following result.³

Theorem 2.6 ([LLM06, Corollaries 1 and 2]). *For any $p \in [1, \infty)$ and $\alpha > 1/2^{1/p}$, there is no polynomial-time algorithm for $\text{BDD}_{p,\alpha}$ (respectively, with preprocessing) unless $\text{NP} \subseteq \text{RP}$ (resp., unless $\text{NP} \subseteq \text{P/Poly}$).*

Regev and Rosen [RR06] used norm embeddings to show that almost any lattice problem is at least as hard in the ℓ_p norm, for any $p \in [1, \infty]$, as it is in the ℓ_2 norm, up to an arbitrarily small constant-factor loss in the approximation factor. In other words, they essentially showed that ℓ_2 is the “easiest” ℓ_p norm for lattice problems. (In addition, their reduction preserves the rank of the lattice.) Based on this, [LLM06] observed the following corollary, which is an improvement on the factor α from Theorem 2.6 for all $p > 2$.

Theorem 2.7 ([LLM06, Corollary 3]). *For any $p \in [1, \infty)$ and $\alpha > 1/\sqrt{2}$, there is no polynomial-time algorithm for $\text{BDD}_{p,\alpha}$ (respectively, with preprocessing) unless $\text{NP} \subseteq \text{RP}$ (resp., unless $\text{NP} \subseteq \text{P/Poly}$).*

Figure 1 shows the bounds from Theorems 2.6 and 2.7 together with the new bounds achieved in this work as a function of p .

2.4 Sparsification

A powerful idea, first used in the context of hardness proofs for lattice problems in [Kho03], is that of random lattice sparsification. Given a lattice \mathcal{L} with basis B , we can construct a random sublattice $\mathcal{L}' \subseteq \mathcal{L}$ as

$$\mathcal{L}' = \{\mathbf{v} \in \mathcal{L} : \langle \mathbf{z}, B^+ \mathbf{v} \rangle = 0 \pmod{q}\}$$

for uniformly random $\mathbf{z} \in \mathbb{Z}_q^n$, where q is a suitably chosen prime.

³Additionally, [DRS14] gave a reduction from CVP to $\text{BDD}_{2,\alpha}$ but only for some $\alpha > 1$. Also, [Pei09, LM09] gave a reduction from GapSVP_γ to BDD, but only for large $\gamma = \gamma(n)$ for which GapSVP is not known to be NP-hard.

Lemma 2.8. *Let q be a prime and let $\mathbf{x}_1, \dots, \mathbf{x}_N \in \mathbb{Z}_q^n \setminus \{\mathbf{0}\}$ be arbitrary. Then*

$$\Pr_{\mathbf{z} \leftarrow \mathbb{Z}_q^n} [\exists i \in [N] \text{ such that } \langle \mathbf{z}, \mathbf{x}_i \rangle = 0 \pmod{q}] \leq \frac{N}{q}.$$

Proof. We have $\Pr[\langle \mathbf{z}, \mathbf{x}_i \rangle = 0] = 1/q$ for each \mathbf{x}_i , and the claim follows by the union bound. \square

The following corollary is immediate.

Corollary 2.9. *Let q be a prime and \mathcal{L} be a lattice of rank n with basis B . Then for all $r > 0$ and all $p \in [1, \infty]$,*

$$\Pr_{\mathbf{z} \leftarrow \mathbb{Z}_q^n} [\lambda_1^{(p)}(\mathcal{L}') < r] \leq \frac{N_p^o(\mathcal{L} \setminus \{\mathbf{0}\}, r, \mathbf{0})}{q},$$

where $\mathcal{L}' = \{\mathbf{v} \in \mathcal{L} : \langle \mathbf{z}, B^+ \mathbf{v} \rangle = 0 \pmod{q}\}$.

Theorem 2.10 ([Ste16, Theorem 3.1]). *For any lattice \mathcal{L} of rank n with basis B , prime q , and lattice vectors $\mathbf{x}, \mathbf{y}_1, \dots, \mathbf{y}_N \in \mathcal{L}$ such that $B^+ \mathbf{x} \not\equiv B^+ \mathbf{y}_i \pmod{q}$ for all $i \in [N]$, we have*

$$\frac{1}{q} - \frac{N}{q^2} - \frac{N}{q^{n-1}} \leq \Pr_{\mathbf{z}, \mathbf{c} \leftarrow \mathbb{Z}_q^n} [\langle \mathbf{z}, B^+ \mathbf{x} + \mathbf{c} \rangle = 0 \pmod{q} \wedge \langle \mathbf{z}, B^+ \mathbf{y}_i + \mathbf{c} \rangle \neq 0 \pmod{q} \forall i \in [N]] \leq \frac{1}{q} + \frac{1}{q^n}.$$

We will use only the lower bound from Theorem 2.10, but we note that the upper bound is relatively tight for $q \gg N$.

Corollary 2.11. *For any $p \in [1, \infty]$ and $r \geq 0$, lattice \mathcal{L} of rank n with basis B , vector \mathbf{t} , prime q , and lattice vectors $\mathbf{v}_1, \dots, \mathbf{v}_N \in \mathcal{L}$ such that $\|\mathbf{v}_i - \mathbf{t}\|_p \leq r$ for all $i \in [N]$ and such that all the $B^+ \mathbf{v}_i \pmod{q}$ are distinct, we have*

$$\Pr_{\mathbf{z}, \mathbf{c} \leftarrow \mathbb{Z}_q^n} [\text{dist}_p(\mathbf{t} + B\mathbf{c}, \mathcal{L}') \leq r] \geq \frac{N}{q} - \frac{N(N-1)}{q^2} - \frac{N(N-1)}{q^{n-1}},$$

where $\mathcal{L}' = \{\mathbf{v} \in \mathcal{L} : \langle \mathbf{z}, B^+ \mathbf{v} \rangle = 0 \pmod{q}\}$.

Proof. Observe that for each $i \in [N]$, the events

$$E_i := [\langle \mathbf{z}, B^+ \mathbf{v}_i \rangle = 0 \pmod{q} \text{ and } \langle \mathbf{z}, B^+ \mathbf{v}_j \rangle \neq 0 \pmod{q} \text{ for all } j \neq i]$$

are disjoint, and by invoking Theorem 2.10 with $\mathbf{x} = \mathbf{v}_i$ and the \mathbf{y}_j being the remaining \mathbf{v}_k for $k \neq i$, we have

$$\Pr_{\mathbf{z}, \mathbf{c}} [E_i] \geq \frac{1}{q} - \frac{N-1}{q^2} - \frac{N-1}{q^{n-1}}.$$

Also observe that if E_i occurs, then $\mathbf{v}_i + B\mathbf{c} \in \mathcal{L}'$ (also $\mathbf{v}_j + B\mathbf{c} \notin \mathcal{L}'$ for all $j \neq i$, but we will not need this). Therefore,

$$\text{dist}_p(\mathbf{t} + B\mathbf{c}, \mathcal{L}') \leq \|\mathbf{t} + B\mathbf{c} - (\mathbf{v}_i + B\mathbf{c})\| = \|\mathbf{t} - \mathbf{v}_i\| \leq r.$$

So, the probability in the left-hand side of the claim is at least

$$\Pr_{\mathbf{z}, \mathbf{c}} \left[\bigcup_{i \in [N]} E_i \right] = \sum_{i \in [N]} \Pr_{\mathbf{z}, \mathbf{c}} [E_i] \geq \frac{N}{q} - \frac{N(N-1)}{q^2} - \frac{N(N-1)}{q^{n-1}}. \quad \square$$

2.5 Counting Lattice Points in a Ball

Following [AS18a], for any discrete set A of points (e.g., a lattice, or a subset thereof), we denote the number of points in A contained in the closed and open (respectively) ℓ_p ball of radius r centered at a point \mathbf{t} as

$$N_p(A, r, \mathbf{t}) := |\{\mathbf{y} \in A : \|\mathbf{y} - \mathbf{t}\|_p \leq r\}|, \quad (2.1)$$

$$N_p^o(A, r, \mathbf{t}) := |\{\mathbf{y} \in A : \|\mathbf{y} - \mathbf{t}\|_p < r\}|. \quad (2.2)$$

Clearly, $N_p^o(A, r, \mathbf{t}) \leq N_p(A, r, \mathbf{t})$.

For $1 \leq p < \infty$ and $\tau > 0$ define

$$\Theta_p(\tau) := \sum_{z \in \mathbb{Z}} \exp(-\tau|z|^p).$$

We use the following upper bound due to Mazo and Odlyzko [MO90] on the number of short vectors in the integer lattice. We include its short proof for completeness.

Proposition 2.12 ([MO90]). *For any $p \in [1, \infty)$, $r > 0$, and $n \in \mathbb{N}$,*

$$N_p(\mathbb{Z}^n, r, \mathbf{0}) \leq \min_{\tau > 0} \exp(\tau r^p) \cdot \Theta_p(\tau)^n.$$

Proof. For $\tau > 0$ we have

$$\Theta_p(\tau)^n = \sum_{z \in \mathbb{Z}^n} \exp(-\tau\|z\|_p^p) \geq \sum_{z \in \mathbb{Z}^n \cap rB_p^n} \exp(-\tau\|z\|_p^p) \geq \exp(-\tau r^p) \cdot N_p(\mathbb{Z}^n, r, \mathbf{0}).$$

The result follows by rearranging and taking the minimum over all $\tau > 0$. □

2.6 Hardness Assumptions

We recall the Exponential Time Hypothesis (ETH) of Impagliazzo and Paturi [IP01], and several of its variants. These hypotheses make stronger assumptions about the complexity of the k -SAT problem than the assumption $P \neq NP$, and serve as highly useful tools for studying the fine-grained complexity of hard computational problems. Indeed, we will show that strong fine-grained hardness for BDD follows from these hypotheses.

Definition 2.13. The *(randomized) Exponential Time Hypothesis* ((randomized) ETH) asserts that there is no (randomized) $2^{o(n)}$ -time algorithm for 3-SAT on n variables.

Definition 2.14. The *(randomized) Strong Exponential Time Hypothesis* ((randomized) SETH) asserts that for every $\varepsilon > 0$ there exists $k \in \mathbb{Z}^+$ such that there is no (randomized) $2^{(1-\varepsilon)n}$ -time algorithm for k -SAT on n variables.

For proving hardness of lattice problem with preprocessing, we define (Max-) k -SAT with preprocessing as follows. The preprocessing input is a size parameter n , encoded in unary. The query input is a k -SAT formula ϕ with n variables and m (distinct) clauses, together with a threshold $W \in \{0, \dots, m\}$ in the case of Max- k -SAT. For k -SAT, it is a YES instance if ϕ is satisfiable, and is a NO instance otherwise. For Max- k -SAT, it is a YES instance if there exists an assignment to the variables of ϕ that satisfies at least W of its clauses, and is a NO instance otherwise.

Observe that because the preprocessing input is just n , a preprocessing algorithm for (Max-) k -SAT with preprocessing is equivalent to a (non-uniform) family of circuits for the problem *without* preprocessing. Also, for any fixed k , because there are only $O(n^k)$ possible clauses on n variables, the length of the query input for (Max-) k -SAT instances having preprocessing input n is $\text{poly}(n)$, so we get the following corollary of Lemma 2.2.

Corollary 2.15. *If (Max-) k -SAT with preprocessing has a randomized $T(n)$ -time algorithm, then it has a deterministic $T(n) \cdot \text{poly}(n)$ -time algorithm using $T(n) \cdot \text{poly}(n)$ -length preprocessed output.*

Following, e.g., [SV19, ABGS19], we also define *non-uniform* variants of ETH and SETH, which deal with the complexity of k -SAT with preprocessing. More precisely, non-uniform ETH asserts that no family of size- $2^{o(n)}$ circuits solves 3-SAT on n variables (equivalently, 3-SAT with preprocessing does not have a $2^{o(n)}$ -time algorithm), and non-uniform SETH asserts that for every $\varepsilon > 0$ there exists $k \in \mathbb{Z}^+$ such that no family of circuits of size $2^{(1-\varepsilon)n}$ solves k -SAT on n variables (equivalently, k -SAT with preprocessing does not have a $2^{(1-\varepsilon)n}$ -time algorithm). These hypotheses are useful for analyzing the fine-grained complexity of preprocessing problems.

One might additionally consider “randomized non-uniform” versions of (S)ETH. However, Corollary 2.15 says that a randomized algorithm for (Max-) k -SAT with preprocessing can be derandomized with only polynomial overhead, so randomized non-uniform (S)ETH is equivalent to (deterministic) non-uniform (S)ETH, so we only consider the latter.

Finally, we remark that one can define weaker versions of randomized or non-uniform (S)ETH with Max-3-SAT (respectively, Max- k -SAT) in place of 3-SAT (resp., k -SAT). Many of our results hold even under these weaker hypotheses. In particular, the derandomization result in Corollary 2.15 applies to both k -SAT and Max- k -SAT.

3 Hardness of $\text{BDD}_{p,\alpha}$

In this section, we present our main result by giving a reduction from a known-hard variant GapCVP'_p of the Closest Vector Problem (CVP) to BDD. We perform this reduction in two main steps.

1. First, in Section 3.1 we define a variant of $\text{BDD}_{p,\alpha}$, which we call (S, T) - $\text{BDD}_{p,\alpha}$. Essentially, an instance of this problem is a lattice that may have up to S “short” nonzero vectors of ℓ_p norm bounded by some r , and a target vector that is “close” to—i.e., within distance αr of—at least T lattice vectors. (The presence of short vectors prevents this from being a true $\text{BDD}_{p,\alpha}$ instance.) We then give a reduction, for $S \ll T$, from (S, T) - $\text{BDD}_{p,\alpha}$ to $\text{BDD}_{p,\alpha}$ itself, using sparsification.
2. Then, in Section 3.2 we reduce from GapCVP'_p to (S, T) - $\text{BDD}_{p,\alpha}$ for suitable $S \ll T$ whenever α is sufficiently large as a function of p (and the desired rank efficiency), based on analysis given in Section 3.3 and Lemma 3.11.

3.1 (S, T) -BDD to BDD

We start by defining a special decision variant of BDD. Essentially, the input is a lattice and a target vector, and the problem is to distinguish between the case where there are few “short” lattice vectors but many lattice vectors “close” to the target, and the case where the target is not close to the lattice. There is a gap factor between the “close” and “short” distances, and for technical reasons we count only those “close” vectors having binary coefficients with respect to the given input basis.

Definition 3.1. Let $S = S(n), T = T(n) \geq 0, p \in [1, \infty]$, and $\alpha = \alpha(n) > 0$. An instance of the decision promise problem (S, T) -BDD $_{p,\alpha}$ is a lattice basis $B \in \mathbb{R}^{d \times n}$, a distance $r > 0$, and a target $\mathbf{t} \in \mathbb{R}^d$.

- It is a YES instance if $N_p^o(\mathcal{L}(B) \setminus \{\mathbf{0}\}, r, \mathbf{0}) \leq S(n)$ and $N_p(B \cdot \{0, 1\}^n, \alpha r, \mathbf{t}) \geq T(n)$.
- It is a NO instance if $\text{dist}_p(\mathbf{t}, \mathcal{L}(B)) > \alpha r$.

The search version is: given a YES instance (B, r, \mathbf{t}) , find a $\mathbf{v} \in \mathcal{L}(B)$ such that $\|\mathbf{v} - \mathbf{t}\|_p \leq \alpha r$.

The preprocessing search and decision problems (S, T) -BDDP $_{p,\alpha}$ are defined analogously, where the preprocessing input is B and r , and the query input is \mathbf{t} .

We stress that in the preprocessing problems BDDP, the distance r is part of the preprocessing input; this makes the problem no harder than a variant where r is part of the query input. So, our hardness results for the above definition immediately apply to that variant as well. However, our reduction *from* (S, T) -BDDP (given in Lemma 3.2) critically relies on the fact that r is part of the preprocessing input.

Clearly, there is a trivial reduction from the decision version of (S, T) -BDD $_{p,\alpha}$ to its search version (and similarly for the preprocessing problems): just call the oracle for the search problem and test whether it returns a lattice vector within distance αr of the target. So, to obtain more general results, our reductions involving (S, T) -BDD will be *from* the search version, and *to* the decision version.

Reducing to BDD. We next observe that for $S(n) = 0$ and any $T(n) > 0$, there is *almost* a trivial reduction from (S, T) -BDD $_{p,\alpha}$ to ordinary BDD $_{p,\alpha}$, because YES instances of the former satisfy the BDD $_{p,\alpha}$ promise. (See below for the easy proof.) The only subtlety is that we want the BDD $_{p,\alpha}$ oracle to return a lattice vector that is within distance αr of the target; recall that the definition of BDD $_{p,\alpha}$ only guarantees distance $\alpha \cdot \lambda_1^{(p)}(\mathcal{L}(B))$. This issue is easily resolved by modifying the lattice to *upper bound* its minimum distance by r , which increases the lattice's rank by one. (For the alternative definition of BDD described after Definition 2.5, the trivial reduction works, and no increase in the rank is needed.)

Lemma 3.2. *For any $T(n) > 0, p \in [1, \infty]$, and $\alpha = \alpha(n) > 0$, there is a deterministic Cook reduction from the search version of $(0, T(n))$ -BDD $_{p,\alpha}$ (resp., with preprocessing) in rank n to BDD $_{p,\alpha}$ (resp., with preprocessing) in rank $n + 1$.*

Proof. The reduction works as follows. On input (B, r, \mathbf{t}) , call the BDD $_{p,\alpha}$ oracle on

$$B' := \begin{pmatrix} B & 0 \\ 0 & r \end{pmatrix}, \quad \mathbf{t}' := \begin{pmatrix} \mathbf{t} \\ 0 \end{pmatrix},$$

and (without loss of generality) receive from the oracle a vector $\mathbf{v}' = (\mathbf{v}, zr)$ for some $\mathbf{v} \in \mathcal{L}$ and $z \in \mathbb{Z}$. Output \mathbf{v} .

We analyze the reduction. Let $\mathcal{L} = \mathcal{L}(B)$ and $\mathcal{L}' = \mathcal{L}(B')$. Because the input is a YES instance, we have $N_p^o(\mathcal{L} \setminus \{\mathbf{0}\}, r, \mathbf{0}) = 0$ and hence $\lambda_1^{(p)}(\mathcal{L}) \geq r$, so $\lambda_1^{(p)}(\mathcal{L}') = r$. Moreover, $N_p(B \cdot \{0, 1\}^n, \alpha r, \mathbf{t}) > 0$ implies that $\text{dist}_p(\mathbf{t}, \mathcal{L}) \leq \alpha r = \alpha \cdot \lambda_1^{(p)}(\mathcal{L}')$. So, (B', \mathbf{t}') satisfies the BDD $_{p,\alpha}$ promise, hence the oracle is obligated to return some $\mathbf{v}' = (\mathbf{v}, zr) \in \mathcal{L}'$ where $\mathbf{v} \in \mathcal{L}$ and $\alpha r = \alpha \lambda_1^{(p)}(\mathcal{L}') \geq \|\mathbf{v}' - \mathbf{t}'\|_p \geq \|\mathbf{v} - \mathbf{t}\|_p$. Therefore, the output \mathbf{v} of the reduction is a valid solution.

Finally, observe that all of the above also constitutes a valid reduction for the preprocessing problems, because B' depends only on the preprocessing part B, r of the input. \square

We now present a more general randomized reduction from (S, T) - $\text{BDD}_{p,\alpha}$ to $\text{BDD}_{p,\alpha}$, which works whenever $T(n) \geq 10S(n)$. The essential idea is to sparsify the input lattice, so that with some noticeable probability no short vectors remain, but at least one vector close to the target does remain. In this case, the result will be an instance of $(0, 1)$ - $\text{BDD}_{p,\alpha}$, which reduces to $\text{BDD}_{p,\alpha}$ as shown above.

We note that the triangle inequality precludes the existence of (S, T) - $\text{BDD}_{p,\alpha}$ instances with $T > S + 1$ and $\alpha \leq 1/2$, so with this approach we can only hope to show hardness of $\text{BDD}_{p,\alpha}$ for $\alpha > 1/2$, i.e., the unique-decoding regime remains out of reach.

Theorem 3.3. *For any $S = S(n) \geq 1$ and $T = T(n) \geq 10S$ that is efficiently computable (for unary n), $p \in [1, \infty]$, and $\alpha = \alpha(n) > 0$, there is a randomized Cook reduction with no false positives from the search version of (S, T) - $\text{BDD}_{p,\alpha}$ (resp., with preprocessing) in rank n to $\text{BDD}_{p,\alpha}$ (resp., with preprocessing) in rank $n + 1$.*

Proof. By Lemma 3.2, it suffices to give such a reduction to $(0, 1)$ - $\text{BDD}_{p,\alpha}$ in rank n , which works as follows. On input (B, r, \mathbf{t}) , let $\mathcal{L} = \mathcal{L}(B)$. First, randomly choose a prime q where $10T \leq q \leq 20T$. Then sample $\mathbf{z}, \mathbf{c} \in \mathbb{Z}_q^n$ independently and uniformly at random, and define

$$\mathcal{L}' := \{\mathbf{v} \in \mathcal{L} : \langle \mathbf{z}, B^+ \mathbf{v} \rangle = 0 \pmod{q}\} \text{ and } \mathbf{t}' := \mathbf{t} + B\mathbf{c}.$$

Let B' be a basis of \mathcal{L}' . (Such a basis is efficiently computable from B , \mathbf{z} , and q . See, e.g., [Ste16, Claim 2.15].) Invoke the $(0, 1)$ - $\text{BDD}_{p,\alpha}$ oracle on (B', r, \mathbf{t}') , and output whatever the oracle outputs.

We now analyze the reduction. We are promised that (B, r, \mathbf{t}) is a YES instance of (S, T) - $\text{BDD}_{p,\alpha}$, and it suffices to show that (B', r, \mathbf{t}') is a YES instance of $(0, 1)$ - $\text{BDD}_{p,\alpha}$, i.e., $\lambda_1^{(p)}(\mathcal{L}') \geq r$ and $\text{dist}_p(\mathbf{t}', \mathcal{L}') \leq \alpha r$, with some positive constant probability. By Corollary 2.9 we have

$$\Pr[\lambda_1^{(p)}(\mathcal{L}') < r] \leq \frac{N_p^o(\mathcal{L} \setminus \{\mathbf{0}\}, r, \mathbf{0})}{q} \leq \frac{S}{q} \leq \frac{1}{100}.$$

Furthermore, because there are T vectors $\mathbf{v}_i \in \mathcal{L}$ for which $\|\mathbf{v}_i - \mathbf{t}\|_p \leq \alpha r$, and their coefficient vectors $B^+ \mathbf{v}_i \in \{0, 1\}^n$ are distinct (as integer vectors, and hence also modulo q), by Corollary 2.11 we have

$$\Pr[\text{dist}_p(\mathbf{t}', \mathcal{L}') \leq \alpha r] \geq \frac{T}{q} - \frac{T^2}{q^2} - \frac{T^2}{q^{n-1}} \geq \frac{1}{20} - \frac{1}{400} - \frac{1}{400q^{n-3}}.$$

Therefore, by the union bound we have

$$\Pr[\lambda_1^{(p)}(\mathcal{L}') \geq r \text{ and } \text{dist}_p(\mathbf{t}', \mathcal{L}') \leq \alpha r] \geq \frac{1}{20} - \frac{1}{400} - \frac{1}{400q^{n-3}} - \frac{1}{100} \geq \frac{1}{40}$$

for all $n \geq 3$, as desired.

Finally, the above also constitutes a valid reduction for the preprocessing problems (in the sense of Definition 2.3), because B' depends only on B from the preprocessing part of the input and the reduction's own random choices (and r remains unchanged). \square

3.2 GapCVP' to (S, T) -BDD

Here we show that a known-hard variant of the (exact) Closest Vector Problem reduces to (S, T) -BDD (in its decision version).

Definition 3.4. For $p \in [1, \infty]$, the (decision) promise problem GapCVP'_p is defined as follows: an instance consists of a basis $B \in \mathbb{R}^{d \times n}$ and a target vector $\mathbf{t} \in \mathbb{R}^d$.

- It is a YES instance if there exists $\mathbf{x} \in \{0, 1\}^n$ such that $\|B\mathbf{x} - \mathbf{t}\|_p \leq 1$.
- It is a NO instance if $\text{dist}_p(\mathbf{t}, \mathcal{L}(B)) > 1$.

The preprocessing (decision) promise problem $\text{GapCVPP}'_p$ is defined analogously, where the preprocessing input is B and the query input is \mathbf{t} .

Observe that for GapCVP'_p the distance threshold is 1 (and not some instance-dependent value) without loss of generality, because we can scale the lattice and target vector. The same goes for $\text{GapCVPP}'_p$, with the caveat that any instance-dependent distance threshold would need to be included in the *preprocessing* part of the input, not the query part. (See Remark 3.10 below for why this is essentially without loss of generality, under a mild assumption on the $\text{GapCVPP}'_p$ instances.) We remark that some works define these problems with a stronger requirement that in the NO case, $\text{dist}_p(z\mathbf{t}, \mathcal{L}(B)) > r$ for all $z \in \mathbb{Z} \setminus \{0\}$. We will not need this stronger requirement, and some of the hardness results for GapCVP'_p that we rely on are not known to hold with it, so we use the weaker requirement.

We next describe a simple transformation on lattices and target vectors: we essentially take a direct sum of the input lattice with the integer lattice of any desired dimension n and append an all- $\frac{1}{2}$ s vector to the target vector.

Lemma 3.5. For any $n' \leq n$, define the following transformations that map a basis B' of a rank- n' lattice \mathcal{L}' to a basis B of a rank- n lattice \mathcal{L} , and a target vector \mathbf{t}' to a target vector \mathbf{t} :

$$B := \begin{pmatrix} \frac{1}{2}B' & 0 \\ I_{n'} & 0 \\ 0 & I_{n-n'} \end{pmatrix}, \quad \mathbf{t} := \frac{1}{2} \begin{pmatrix} \mathbf{t}' \\ \mathbf{1}_{n'} \\ \mathbf{1}_{n-n'} \end{pmatrix}, \quad (3.1)$$

and define

$$s_p = s_p(n) := \frac{1}{2}(n+1)^{1/p} \text{ for } p \in [1, \infty), \text{ and } s_\infty := 1/2. \quad (3.2)$$

Then:

1. $N_p^o(\mathcal{L}, r, \mathbf{0}) \leq N_p^o(\mathbb{Z}^n, r, \mathbf{0})$ for all $r \geq 0$;
2. if there exists an $\mathbf{x} \in \{0, 1\}^{n'}$ such that $\|B'\mathbf{x} - \mathbf{t}'\|_p \leq 1$, then $N_p(B \cdot \{0, 1\}^n, s_p, \mathbf{t}) \geq 2^{n-n'}$;
3. if $\text{dist}_p(\mathbf{t}', \mathcal{L}') > 1$ then $\text{dist}_p(\mathbf{t}, \mathcal{L}) > s_p$.

Proof. Item 1 follows immediately by construction of B , because vectors $\mathbf{v}' = (\frac{1}{2}B'\mathbf{x}, \mathbf{x}, \mathbf{y}) \in \mathcal{L}$ for $\mathbf{x}, \mathbf{y} \in \mathbb{Z}^n$ correspond bijectively to vectors $\mathbf{v} = (\mathbf{x}, \mathbf{y}) \in \mathbb{Z}^n$, and $\|\mathbf{v}\|_p \leq \|\mathbf{v}'\|_p$.

For Item 2, for every $\mathbf{y} \in \{0, 1\}^{n-n'}$, the vector $\mathbf{v} := (\frac{1}{2}B'\mathbf{x}, \mathbf{x}, \mathbf{y}) \in \mathcal{L}$ satisfies

$$\|\mathbf{v} - \mathbf{t}\|_p^p = \frac{\|B'\mathbf{x} - \mathbf{t}'\|_p^p}{2^p} + \frac{n}{2^p} \leq s_p^p$$

for finite p , and $\|\mathbf{v} - \mathbf{t}\|_\infty = \max(\frac{1}{2}\|B'\mathbf{x} - \mathbf{t}'\|_\infty, \frac{1}{2}) = \frac{1}{2} = s_\infty$. The claim follows.

For Item 3, for finite p we have

$$\text{dist}_p(\mathbf{t}, \mathcal{L})^p \geq \frac{\text{dist}_p(\mathbf{t}', \mathcal{L}')^p}{2^p} + \frac{n}{2^p} > \frac{n+1}{2^p} = s_p^p,$$

and for $p = \infty$ we immediately have $\text{dist}_\infty(\mathbf{t}, \mathcal{L}) \geq \frac{1}{2} \text{dist}_\infty(\mathbf{t}', \mathcal{L}') > \frac{1}{2} = s_\infty$, as needed. \square

Corollary 3.6. *For any $p \in [1, \infty]$, $\alpha > 0$, and $\text{poly}(n')$ -bounded $n \geq n'$, there is a deterministic Karp reduction from GapCVP'_p (resp., with preprocessing) in rank n' to the decision version of (S, T) - $\text{BDD}_{p, \alpha}$ (resp., with preprocessing) in rank n , where $S(n) = N_p^o(\mathbb{Z}^n \setminus \{\mathbf{0}\}, s_p/\alpha, \mathbf{0})$ for s_p as defined in Equation (3.2), and $T(n) = 2^{n-n'}$.*

Proof. Given an input GapCVP'_p instance (B', \mathbf{t}') , the reduction simply outputs $(B, r = s_p/\alpha, \mathbf{t})$, where B, \mathbf{t} are as in Equation (3.1). Observe that this is also valid for the preprocessing problems because B and r depend only on B' . Correctness follows immediately by Lemma 3.5. \square

3.3 Setting Parameters

We now investigate the relationship among the choice of ℓ_p norm (for finite p), the BDD relative distance α , and the rank ratio $C := n/n'$, subject to the constraint

$$N_p^o(\mathbb{Z}^n, s_p/\alpha, \mathbf{0}) \leq 2^{n-n'}/10 = T(n)/10, \quad (3.3)$$

so that the reductions in Corollary 3.6 and Theorem 3.3 can be composed. For $p \in [1, \infty)$ and $C > 1$, define

$$\alpha_{p,C}^* := \inf\{\alpha^* > 0 : \min_{\tau > 0} \exp(\tau/(2\alpha^*)^p) \cdot \Theta_p(\tau) \leq 2^{1-1/C}\}, \quad (3.4)$$

$$\alpha_p^* := \lim_{C \rightarrow \infty} \alpha_{p,C}^* = \inf\{\alpha^* > 0 : \min_{\tau > 0} \exp(\tau/(2\alpha^*)^p) \cdot \Theta_p(\tau) \leq 2\}. \quad (3.5)$$

These quantities are well defined because for any $C > 1$ we have $2^{1-1/C} > 1$, so the inequality in Equation (3.4) is satisfied for sufficiently large τ and α^* . Moreover, it is straightforward to verify that $\alpha_{p,C}^*$ is strictly decreasing in both p and C , and α_p^* is strictly decreasing in p . Although it is not clear how to solve for these quantities in closed form, it is possible to approximate them numerically to good accuracy (see Figure 1), and to get quite tight closed-form upper bounds (see Lemma 3.11). We now show that to satisfy Equation (3.3) it suffices to take any constant $\alpha > \alpha_{p,C}^*$.

Corollary 3.7. *For any $p \in [1, \infty)$, $C \geq 1$, and constant $\alpha > \alpha_{p,C}^*$ (Equation (3.4)), there is a deterministic Karp reduction from GapCVP'_p (resp., with preprocessing) in rank n' to the decision version of (S, T) - $\text{BDD}_{p, \alpha}$ (resp., with preprocessing) in rank $n = Cn'$, where $S(n) = T(n)/10$ and $T(n) = 2^{(1-1/C)n}$.*

Proof. Recalling that $s_p = \frac{1}{2}(n+1)^{1/p}$, by Proposition 2.12, $N_p^o(\mathbb{Z}^n, s_p/\alpha, \mathbf{0})$ is at most

$$\begin{aligned} N_p^o(\mathbb{Z}^n, s_p/\alpha, \mathbf{0}) &\leq \min_{\tau > 0} \exp(\tau \cdot (s_p/\alpha)^p) \cdot \Theta_p(\tau)^n \\ &= \min_{\tau > 0} \exp(\tau \cdot (n+1)/(2\alpha)^p) \cdot \Theta_p(\tau)^n \\ &= \left(\min_{\tau > 0} \exp(\tau/(n(2\alpha)^p)) \cdot \exp(\tau/(2\alpha)^p) \cdot \Theta_p(\tau) \right)^n. \end{aligned}$$

Because $\alpha > \alpha_{p,C}^*$, we have that $\min_{\tau > 0} \exp(\tau/(2\alpha)^p) \cdot \Theta_p(\tau)$ is a constant strictly less than $2^{1-1/C}$. So, $N_p^o(\mathbb{Z}^n, s_p/\alpha, \mathbf{0}) \leq 2^{(1-1/C)n}/10 = T(n)/10$ for all large enough n . The claim follows from Corollary 3.6. \square

Theorem 3.8. *For any $p \in [1, \infty)$, $C \geq 1$, and constant $\alpha > \alpha_{p,C}^*$, there is a randomized Cook reduction with no false positives from GapCVP'_p (resp., with preprocessing) in rank n' to $\text{BDD}_{p, \alpha}$ (resp., with preprocessing) in rank $n = Cn' + 1$. Furthermore, the same holds for $p = \infty$, $C = 1$, $\alpha = 1/2$, and the reduction is deterministic.*

Proof. For finite p , we simply compose the reductions from Corollary 3.7 and Theorem 3.3, with the trivial decision-to-search reduction for (S, T) -BDD $_{p, \alpha}$ in between.

For $p = \infty$, we first invoke the deterministic reduction from Corollary 3.6, from GapCVP' $_{\infty}$ in rank n' to (S, T) -BDD $_{\infty, 1/2}$ in rank $Cn' = n'$, where $S = N_{\infty}^o(\mathbb{Z}^n \setminus \{\mathbf{0}\}, 1, \mathbf{0}) = 0$ and $T = 2^0 > 0$. By Lemma 3.2, the latter problem reduces deterministically to BDD $_{\infty, 1/2}$ in rank $n' + 1$.

Lastly, all of these reductions work for the preprocessing problems as well, because their component reductions do. \square

3.4 Putting it all Together

We now combine our reductions from GapSVP' to BDD with prior hardness results for GapCVP' (stated below in Theorem 3.9) to obtain our ultimate hardness theorems for BDD. We first recall relevant known hardness results for GapCVP' $_p$ and GapCVPP' $_p$.

Theorem 3.9 ([Mic01, BGS17, ABGS19]). *The following hold for GapCVP' $_p$ and GapCVPP' $_p$ in rank n :*

1. *For every $p \in [1, \infty]$, GapCVP' $_p$ is NP-hard, and GapCVPP' $_p$ has no polynomial-time (preprocessing) algorithm unless $\text{NP} \subseteq \text{P/Poly}$.*
2. *For every $p \in [1, \infty]$, there is no $2^{o(n)}$ -time randomized algorithm for GapCVP' $_p$ unless randomized ETH fails.*
3. *For every $p \in [1, \infty] \setminus \{2\}$, there is no $2^{o(n)}$ -time algorithm for GapCVPP' $_p$, and there is no $2^{o(\sqrt{n})}$ -time algorithm for GapCVPP' $_2$, unless non-uniform ETH fails.*
4. *For every $p \in [1, \infty] \setminus 2\mathbb{Z}$ and every $\varepsilon > 0$, there is no $2^{(1-\varepsilon)n}$ -time randomized algorithm for GapCVP' $_p$ (respectively, GapCVPP' $_p$) unless randomized SETH (resp., non-uniform SETH) fails.*

Remark 3.10. Several of the above results are stated slightly differently from what appears in [Mic01, BGS17, ABGS19]. First, all of the above results for GapCVP' $_p$ (respectively, GapCVPP' $_p$) are instead stated for GapCVP $_p$ (resp., GapCVPP $_p$). However, inspection shows that the reductions are indeed to GapCVP' $_p$ or GapCVPP' $_p$, so this difference is immaterial.

Second, the above statements ruling out *randomized* algorithms for GapCVP' $_p$ assuming randomized (S)ETH are instead phrased in [BGS17, ABGS19] as ruling out *deterministic* algorithms for GapCVP' $_p$ assuming deterministic (S)ETH. However, because these results are proved via deterministic reductions, randomized algorithms for GapCVP' $_p$ have the consequences claimed above.

Third, the above results for GapCVPP' $_p$ follow from the reductions given in (the proofs of) [Mic01], [ABGS19, Theorem 4.3], [BGS17, Theorem 1.4 and Lemma 6.1], and [ABGS19, Theorem 4.6]. However, those reductions all prove hardness for the variant of GapCVPP' $_p$ where the distance threshold r is part of the *query* input, rather than the preprocessing input. Inspection of [ABGS19, Theorem 4.6] shows that r is fixed in the output instance, so this difference is immaterial in that case. We next describe how to handle this difference for the remaining cases. Below we give, for any $p \in [1, \infty)$, a straightforward rank-preserving mapping reduction (in the sense of Definition 2.4) from the variant of GapCVPP' $_p$ where the distance threshold r is part of the query input to the variant where it is part of the preprocessing input, assuming that r is always at most some r^* that depends only on B , and whose length $\log r^*$ is polynomial in the length of B . Inspection shows that such an r^* does indeed exist for the reductions given in [Mic01], [ABGS19, Theorem 4.3], and [BGS17, Lemma 6.1], which handles the second difference for those cases.

The mapping reduction (R_P, R_Q) in question maps $(B, (t, r)) \mapsto ((B', r^*), t')$ as follows. First, R_P takes B as input, and sets $B' := \begin{pmatrix} B \\ \mathbf{0}^t \end{pmatrix}$; it also outputs $\sigma' = r^*$ as side information for R_Q . Then, R_Q

takes (t, r) and r^* as input, and outputs $t' := (t, ((r^*)^p - r^p)^{1/p})$. Using the guarantee that $r^* \geq r$, it is straightforward to check that the output instance $((B', r^*), t')$ is a YES instance (respectively, NO instance) if the input instance $(B, (t, r))$ is a YES instance resp., NO instance, as required.

Finally, we again remark that several of the hardness results in Theorem 3.9 in fact hold under weaker versions of randomized or non-uniform (S)ETH that relate to Max-3-SAT (respectively, Max- k -SAT), instead of 3-SAT (resp. k -SAT). Therefore, it is straightforward to obtain corresponding hardness results for BDD(P) under these weaker assumptions as well.

We can now prove our main theorem, restated from the introduction:

Theorem 1.1. *The following hold for $\text{BDD}_{p,\alpha}$ and $\text{BDDP}_{p,\alpha}$ in rank n :*

1. *For every $p \in [1, \infty)$ and constant $\alpha > \alpha_p^*$ (where $\alpha_p^* \leq \frac{1}{2} \cdot 4.6723^{1/p}$), and for $(p, \alpha) = (\infty, 1/2)$, there is no polynomial-time algorithm for $\text{BDD}_{p,\alpha}$ (respectively, $\text{BDDP}_{p,\alpha}$) unless $\text{NP} \subseteq \text{RP}$ (resp., $\text{NP} \subseteq \text{P/Poly}$).*
2. *For every $p \in [1, \infty)$ and constant $\alpha > \min\{\alpha_p^*, \alpha_2^*\}$, and for $(p, \alpha) = (\infty, 1/2)$, there is no $2^{o(n)}$ -time algorithm for $\text{BDD}_{p,\alpha}$ unless randomized ETH fails.*
3. *For every $p \in [1, \infty) \setminus \{2\}$ and constant $\alpha > \alpha_p^*$, and for $(p, \alpha) = (\infty, 1/2)$, there is no $2^{o(n)}$ -time algorithm for $\text{BDDP}_{p,\alpha}$ unless non-uniform ETH fails.*

Moreover, for every $p \in [1, \infty]$ and $\alpha > \alpha_2^$ there is no $2^{o(\sqrt{n})}$ -time algorithm for $\text{BDDP}_{p,\alpha}$ unless non-uniform ETH fails.*

4. *For every $p \in [1, \infty) \setminus 2\mathbb{Z}$ and constants $C > 1$, $\alpha > \alpha_{p,C}^*$, and $\epsilon > 0$, and for $(p, C, \alpha) = (\infty, 1, 1/2)$, there is no $2^{n(1-\epsilon)/C}$ -time algorithm for $\text{BDD}_{p,\alpha}$ (respectively, $\text{BDDP}_{p,\alpha}$) unless randomized SETH (resp., non-uniform SETH) fails.*

Proof. For BDD, each item of the theorem follows from the corresponding item of Theorem 3.9, followed by Theorem 3.8 and then (where needed) rank-preserving norm embeddings from ℓ_2 to ℓ_p [RR06]. (Also, Lemma 3.11 below provides the upper bound on α_p^* .) The claims for BDDP follow similarly, combined with the well-known fact that $\text{P/Poly} = \text{BPP/Poly}$ and Corollary 2.15.⁴

3.5 An Upper Bound on $\alpha_{p,C}^*$ and α_p^*

We conclude with closed-form upper bounds on $\alpha_{p,C}^*$ and α_p^* . The main idea is to replace $\Theta_p(\tau)$ with an upper bound of $\Theta_1(\tau)$ (which has a closed-form expression) in Equations (3.4) and (3.5), then directly analyze the value of $\tau > 0$ that minimizes the resulting expressions. This leads to quite tight bounds (and also yields tighter bounds than the techniques used in the proof of [AS18a, Claim 4.4], which bounds a related quantity). For example, $\alpha_2^* \approx 1.05006$, and the upper bound in Lemma 3.11 gives $\alpha_2^* \leq 1.08078$; similarly, $\alpha_5^* \approx 0.672558$ and the upper bound in Lemma 3.11 gives $\alpha_5^* \leq 0.680575$.

Lemma 3.11. *Define*

$$g(\sigma, \tau) := \exp(\tau/\sigma) \cdot \left(\frac{2}{1 - \exp(-\tau)} - 1 \right)$$

⁴In fact, $\text{P/Poly} = \text{BPP/Poly}$ also follows as a corollary of the more general derandomization result in Lemma 2.2.

and $\tau^*(\sigma) := \operatorname{arcsinh}(\sigma) = \ln(\sigma + \sqrt{1 + \sigma^2})$. Let σ^* and σ_C^* for $C > 1$ be the (unique) constants for which $g(\sigma^*, \tau^*(\sigma^*)) = 2$ and $g(\sigma_C^*, \tau^*(\sigma_C^*)) = 2^{1-1/C}$. Then for any $p \in [1, \infty)$, we have

$$\alpha_{p,C}^* \leq \frac{1}{2} \cdot (\sigma_C^*)^{1/p} \quad \text{and} \quad \alpha_p^* \leq \frac{1}{2} \cdot (\sigma^*)^{1/p} \leq \frac{1}{2} \cdot 4.6723^{1/p}.$$

In particular, $\alpha_{p,C}^* \rightarrow 1/2$ as $p \rightarrow \infty$ for any fixed $C > 1$, and therefore $\alpha_p^* \rightarrow 1/2$ as $p \rightarrow \infty$.

Proof. For any $\tau > 0$, by the definition of $\Theta_p(\tau)$ and the formula for summing geometric series we have

$$\Theta_p(\tau) \leq \Theta_1(\tau) = 1 + 2 \sum_{i=1}^{\infty} \exp(-\tau)^i = \frac{2}{1 - \exp(-\tau)} - 1. \quad (3.6)$$

Define the objective function

$$f(p, \alpha) := \min_{\tau > 0} \exp(\tau / (2\alpha)^p) \cdot \Theta_p(\tau)$$

to be the expression that is upper-bounded in Equations (3.4) and (3.5). For any fixed $\alpha > 0$, set $\sigma := (2\alpha)^p$. Applying Equation (3.6), it follows that $f(p, \alpha) \leq g(\sigma, \tau)$ for any $\tau > 0$. This implies that if there exists some $\tau > 0$ satisfying $g(\sigma, \tau) \leq 2$ then $\alpha_p^* \leq \frac{1}{2}\sigma^{1/p}$, and similarly, if $g(\sigma, \tau) \leq 2^{1-1/C}$ then $\alpha_{p,C}^* \leq \frac{1}{2}\sigma^{1/p}$.

By standard calculus,

$$\frac{\partial g}{\partial \tau} = \frac{e^{\tau/\sigma}}{1 - e^{-\tau}} \cdot \left((1 + e^{-\tau})/\sigma - 2e^{-\tau}/(1 - e^{-\tau}) \right).$$

Setting the right-hand side of the above expression equal to 0 and solving for τ yields the single real solution

$$\tau = \tau^*(\sigma) = \operatorname{arcsinh}(\sigma) = \ln(\sigma + \sqrt{1 + \sigma^2}),$$

which is a local minimum, and therefore a global minimum of $g(\sigma, \tau)$ for any fixed $\sigma > 0$.

Define the univariate function $g^*(\sigma) := g(\sigma, \tau^*(\sigma))$. The fact that σ^* and σ_C^* exist and are unique follows by noting that $\lim_{\sigma \rightarrow 0^+} g^*(\sigma) = \infty$, that $\lim_{\sigma \rightarrow \infty} g^*(\sigma) = 1$, and that $g^*(\sigma)$ is strictly decreasing in $\sigma > 0$. By definition of σ^* (respectively, σ_C^*), it follows that $g^*(\sigma^*) = 2$ for $\alpha = \frac{1}{2}(\sigma^*)^{1/p}$, and $g^*(\sigma_C^*) = 2^{1-1/C}$ for $\alpha = \frac{1}{2}(\sigma_C^*)^{1/p}$, as desired. Moreover, one can check numerically that $\sigma^* \leq 4.6723$. \square

References

- [ABGS19] D. Aggarwal, H. Bennett, A. Golovnev, and N. Stephens-Davidowitz. Fine-grained hardness of CVP(P)— Everything that we can prove (and nothing else). <https://arxiv.org/abs/1911.02440>, 2019.
- [ABSS97] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes, and systems of linear equations. *J. Comput. Syst. Sci.*, 54(2):317–331, 1997.
- [Adl78] L. M. Adleman. Two theorems on random polynomial time. In *FOCS*, pages 75–83. 1978.
- [ADRS15] D. Aggarwal, D. Dadush, O. Regev, and N. Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling: Extended abstract. In *STOC*, pages 733–742. 2015.

- [ADS15] D. Aggarwal, D. Dadush, and N. Stephens-Davidowitz. Solving the closest vector problem in 2^n time – the discrete Gaussian strikes again! In *FOCS*, pages 563–582. 2015.
- [Ajt98] M. Ajtai. The shortest vector problem in L_2 is NP-hard for randomized reductions (extended abstract). In *STOC*, pages 10–19. 1998.
- [AS18a] D. Aggarwal and N. Stephens-Davidowitz. (Gap/S)ETH hardness of SVP. In *STOC*, pages 228–238. 2018.
- [AS18b] D. Aggarwal and N. Stephens-Davidowitz. Just take the average! An embarrassingly simple 2^n -time algorithm for SVP (and CVP). In *Symposium on Simplicity in Algorithms*, volume 61, pages 12:1–12:19. 2018.
- [BGS17] H. Bennett, A. Golovnev, and N. Stephens-Davidowitz. On the quantitative hardness of CVP. In *FOCS*. 2017.
- [BSW16] S. Bai, D. Stehlé, and W. Wen. Improved reduction from the bounded distance decoding problem to the unique shortest vector problem in lattices. In *ICALP*, pages 76:1–76:12. 2016.
- [DRS14] D. Dadush, O. Regev, and N. Stephens-Davidowitz. On the closest vector problem with a distance guarantee. In *IEEE Conference on Computational Complexity*, pages 98–109. 2014.
- [EOR91] N. D. Elkies, A. M. Odlyzko, and J. A. Rush. On the packing densities of superballs and other bodies. *Inventiones mathematicae*, 105:613–639, December 1991.
- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008.
- [HR07] I. Haviv and O. Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory of Computing*, 8(1):513–531, 2012. Preliminary version in *STOC* 2007.
- [IP01] R. Impagliazzo and R. Paturi. On the complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001.
- [Kho03] S. Khot. Hardness of approximating the shortest vector problem in high ℓ_p norms. *J. Comput. Syst. Sci.*, 72(2):206–219, 2006.
- [Kho04] S. Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, 2005. Preliminary version in *FOCS* 2004.
- [KS01] R. Kumar and D. Sivakumar. On the unique shortest lattice vector problem. *Theor. Comput. Sci.*, 255(1-2):641–648, 2001.
- [LLM06] Y. Liu, V. Lyubashevsky, and D. Micciancio. On bounded distance decoding for general lattices. In *APPROX-RANDOM*, pages 450–461. 2006.
- [LM09] V. Lyubashevsky and D. Micciancio. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In *CRYPTO*, pages 577–594. 2009.
- [Mic98] D. Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM J. Comput.*, 30(6):2008–2035, 2000. Preliminary version in *FOCS* 1998.

- [Mic01] D. Micciancio. The hardness of the closest vector problem with preprocessing. *IEEE Trans. Information Theory*, 47(3):1212–1215, 2001. doi:10.1109/18.915688.
- [Mic08] D. Micciancio. Efficient reductions among lattice problems. In *SODA*, pages 84–93. 2008.
- [Mic12] D. Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory of Computing*, 8(1):487–512, 2012.
- [MO90] J. E. Mazo and A. M. Odlyzko. Lattice points in high-dimensional spheres. *Monatshefte für Mathematik*, 110:47–61, March 1990.
- [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009.
- [Reg05] O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6):1–40, 2009. Preliminary version in *STOC* 2005.
- [RR06] O. Regev and R. Rosen. Lattice problems and norm embeddings. In *STOC*, pages 447–456. 2006.
- [Ste16] N. Stephens-Davidowitz. Discrete Gaussian sampling reduces to CVP and SVP. In *SODA*, pages 1748–1764. 2016.
- [SV19] N. Stephens-Davidowitz and V. Vaikuntanathan. SETH-hardness of coding problems. In *FOCS*, pages 287–301. 2019.
- [vEB81] P. van Emde Boas. Another NP-complete problem and the complexity of computing short vectors in a lattice. Technical Report 81-04, University of Amsterdam, 1981.