

On the Lattice Smoothing Parameter Problem

Kai-Min Chung
Cornell University

Daniel Dadush
New York University

Feng-Hao Liu
Brown University

Chris Peikert
Georgia Institute of Technology

December 5, 2012

Abstract

The *smoothing parameter* $\eta_\varepsilon(\mathcal{L})$ of a Euclidean lattice \mathcal{L} , introduced by Micciancio and Regev (FOCS'04; SICOMP'07), is (informally) the smallest amount of Gaussian noise that “smooths out” the discrete structure of \mathcal{L} (up to error ε). It plays a central role in the best known worst-case/average-case reductions for lattice problems, a wealth of lattice-based cryptographic constructions, and (implicitly) the tightest known transference theorems for fundamental lattice quantities.

In this work we initiate a study of the complexity of approximating the smoothing parameter to within a factor γ , denoted γ -GapSPP. We show that (for $\varepsilon = 1/\text{poly}(n)$):

- $(2 + o(1))$ -GapSPP \in AM, via a Gaussian analogue of the classic Goldreich-Goldwasser protocol (STOC'98);
- $(1 + o(1))$ -GapSPP \in coAM, via a careful application of the Goldwasser-Sipser (STOC'86) set size lower bound protocol to thin shells in \mathbb{R}^n ;
- $(2 + o(1))$ -GapSPP \in SZK \subseteq AM \cap coAM (where SZK is the class of problems having statistical zero-knowledge proofs), by constructing a suitable instance-dependent commitment scheme (for a slightly worse $o(1)$ -term);
- $(1 + o(1))$ -GapSPP can be solved in deterministic $2^{O(n)} \text{polylog}(1/\varepsilon)$ time and $2^{O(n)}$ space.

As an application, we demonstrate a tighter worst-case to average-case reduction for basing cryptography on the worst-case hardness of the GapSPP problem, with $\tilde{O}(\sqrt{n})$ smaller approximation factor than the GapSVP problem. Central to our results are two novel, and nearly tight, characterizations of the magnitude of discrete gaussian sums over \mathcal{L} : the first relates these directly to the gaussian measure of the voronoi cell of \mathcal{L} , and the second to the fraction of overlap between Euclidean balls centered around points of \mathcal{L} .

1 Introduction

A (full-rank) n -dimensional lattice $\mathcal{L} = \mathcal{L}(B) = \{\sum_{i=1}^n c_i \mathbf{b}_i : c_i \in \mathbb{Z}\}$ is the set of all integer linear combinations of a set $B = \{\mathbf{b}_1, \dots, \mathbf{b}_n\} \subset \mathbb{R}^n$ of linearly independent vectors, called a basis of the lattice. It may also be seen as a discrete additive subgroup of \mathbb{R}^n . Lattices have been studied in mathematics for hundreds of years, and more recently have been at the center of many important developments in computer science, such as the LLL algorithm [LLL82] and its applications to cryptanalysis [Cop97] and error-correcting codes [CH11], and lattice-based cryptography [Ajt96] (including the first fully homomorphic encryption scheme [Gen09]).

Much recent progress in the computational study of lattices, especially in the realms of worst-case/average-case reductions and cryptography (as initiated by Ajtai [Ajt96]), has been made possible by the machinery of Gaussian measures and harmonic analysis. These tools were first employed for such purposes by Regev [Reg03] and Micciancio and Regev [MR04] (see also, e.g., [AR04, Reg05, Pei07, GPV08, Gen10, Pei10]), following their development by Banaszczyk [Ban93, Ban95, Ban96] to prove asymptotically tight (or nearly tight) transference theorems.

In particular, the notion from [MR04] of the *smoothing parameter* $\eta_\varepsilon(\mathcal{L})$ of a lattice \mathcal{L} plays a central role (sometimes implicitly) in the above-cited works, and so it is a key concept in the study of lattices from several perspectives. Informally, $\eta_\varepsilon(\mathcal{L})$ is the smallest amount s of Gaussian noise that completely “smooths out” the discrete structure of \mathcal{L} , up to statistical error ε . Formally, it is the smallest $s > 0$ such that the total Gaussian mass $\rho_{1/s}(\mathbf{w}) := \exp(-\pi s^2 \|\mathbf{w}\|^2)$, summed over all nonzero *dual* lattice vectors $\mathbf{w} \in \mathcal{L}^* \setminus \{\mathbf{0}\}$, is at most ε .¹ This condition is equivalent to the following “smoothing” condition: the distribution of a continuous Gaussian of width s , reduced modulo \mathcal{L} , has point-wise probability density within a $(1 \pm \varepsilon)$ factor of that of the uniform distribution over \mathbb{R}^n/\mathcal{L} .

Given the smoothing parameter’s central role in many mathematical and computational aspects of lattices, we believe it to be of comparable importance to other fundamental and well-studied geometric lattice quantities like the minimum distance, successive minima, covering radius, etc. While the smoothing parameter can be estimated by relating it to these other quantities [MR04, Pei07, GPV08], the bounds are quite coarse, typically yielding only $\tilde{\Omega}(\sqrt{n})$ -factor approximations.

We therefore initiate a study of the complexity of computing the smoothing parameter, with a focus on approximations. More formally, for an approximation factor $\gamma \geq 1$ and some $0 < \varepsilon < 1$ (which may both be functions of the lattice dimension n), we define γ -GapSPP $_\varepsilon$ to be the promise problem in which YES instances are lattices \mathcal{L} for which $\eta_\varepsilon(\mathcal{L}) \leq 1$, and NO instances are those for which $\eta_\varepsilon(\mathcal{L}) > \gamma$.

The dependence on ε . To understand the nature of GapSPP, it is important to notice that the value of ε has a large impact on the complexity of the problem. In particular, by known relations between the smoothing parameter and the shortest nonzero dual vector (see [MR04]), we have that

$$\sqrt{\log(1/\varepsilon)/\pi}/\lambda_1(\mathcal{L}^*) \leq \eta_\varepsilon(\mathcal{L}) \leq \sqrt{n}/\lambda_1(\mathcal{L}^*),$$

and hence for exponentially small error $\varepsilon = 2^{-\Omega(n)}$ the quantities $\eta_\varepsilon(\mathcal{L})$ and $\sqrt{n}/\lambda_1(\mathcal{L}^*)$ are within a constant factor of each other. Therefore, the (decision) Shortest Vector Problem γ -GapSVP is equivalent to γ -GapSPP $_{2^{-\Omega(n)}}$, up to a constant factor loss in the approximation. However, most uses of the smoothing parameter in the literature (e.g., worst-case to average-case reductions and transference theorems) work with either inverse polynomial $\varepsilon = n^{-O(1)}$ or “just barely” negligible $\varepsilon = n^{-\omega(1)}$ (e.g., $\varepsilon = n^{-\log n}$). For such

¹The dual lattice \mathcal{L}^* of \mathcal{L} is the set of all $\mathbf{y} \in \mathbb{R}^n$ for which $\langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}$ for every $\mathbf{x} \in \mathcal{L}$.

values of ε , the loss in approximation factor between GapSPP and GapSVP or other standard lattice problems can be as large as $\tilde{\Omega}(\sqrt{n})$, and as we will see, in this regime GapSPP behaves qualitatively differently from these other problems.

1.1 Results and Techniques

In this work, we prove several (possibly surprising) upper bounds on the complexity of γ -GapSPP $_\varepsilon$. Unless otherwise specified, the stated results hold for the setting $\varepsilon = n^{-O(1)}$. (We obtain results for smaller ε as well, but with slowly degrading approximation factors.) Similar results hold for a generalization of GapSPP which uses different values of ε for YES and NO instances (see Definition 2.3 and Corollary 2.5 for further details).

At a high level, we obtain several of our main results by noticing that the classic Goldreich–Goldwasser protocol [GG98], which was originally designed for approximating (the complement of) the GapSVP problem, can in fact be seen as more directly and tightly approximating the smoothing parameter (of the dual lattice). When viewed from this perspective, we show that slight variants of the GG protocol obtain an $2 + o(1)$ approximation for GapSPP, improving on the approximation for GapSVP by a $\tilde{O}(\sqrt{n})$ factor. Furthermore, using the known relations between GapSVP and GapSPP, one recover the original approximation factor for GapSVP. To obtain these tight approximation factors, as part of the main technical contributions of this paper, we develop two novel and nearly tight (up to a $2 + o(1)$ factor) geometric characterizations of the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ that elucidate the geometric content of the parameter ε .

Arthur-Merlin Protocols. We show that $(2 + o(1))$ -GapSPP \in AM \cap coAM, and moreover, that $(1 + o(1))$ -GapSPP \in coAM. That is, we give constant-round interactive proof systems which allow an unbounded prover to convince a randomized polynomial-time verifier that the smoothing parameter is “small,” and that it is “large.” In contrast with these positive results, we note that since the smoothing parameter is effectively determined by a sum over exponentially many lattice points, it is unclear whether γ -GapSPP is in NP or coNP for $\gamma = o(\sqrt{n})$. (For $\gamma = \Omega(\sqrt{n})$, known connections to other lattice quantities imply that γ -GapSPP $_\varepsilon \in$ NP \cap coNP.)

One important consequence of $(2 + o(1))$ -GapSPP \in AM \cap coAM is that the problem is not NP-hard (under Karp reductions, or “smart” Cook reductions [GS88]), unless coNP \subseteq AM [BHZ87] and the polynomial-time hierarchy collapses. Our result should also be contrasted with analogous results for approximating the Shortest and Closest Vector Problems, which are only known to be in NP \cap coAM for factors $\gamma \geq c\sqrt{n}/\log n$ [GG98], and in NP \cap coNP for factors $\gamma \geq c\sqrt{n}$ [AR04], as well as the results for approximating the Covering Radius Problem, whose 2-approximation is in AM but is in coAM only for $\gamma \geq c\sqrt{n}/\log n$, and in NP \cap coNP for $\gamma \geq \sqrt{n}$ [GMR04].

To prove that $(2 + o(1))$ -GapSPP \in AM, we use a Gaussian analogue of the Goldreich–Goldwasser protocol on the dual lattice \mathcal{L}^* , where the verifier samples from a Gaussian instead of a ball. (Interestingly, this leads to imperfect completeness, which turns out to be important for the tightness of the analysis.) More precisely, the verifier samples $\mathbf{x} \in \mathbb{R}^n$ from a Gaussian, reduces \mathbf{x} modulo (a basis of) the lattice \mathcal{L}^* , and sends the result to the prover. The prover’s task is to guess \mathbf{x} , and the verifier accepts or rejects accordingly. To prove that the protocol is complete and sound, we crucially rely on the following novel characterization of the smoothing parameter:

Voronoi Cell Characterization. For any $\varepsilon \in (0, 1)$, a scaling of the Voronoi cell² $\mathcal{V}(\mathcal{L}^*)$ by a factor $2\eta_\varepsilon(\mathcal{L})$ has Gaussian measure at least $1 - \varepsilon$, and an $\eta_\varepsilon(\mathcal{L})$ -scaling has Gaussian measure at most $1/(1 + \varepsilon)$.

²The Voronoi cell $\mathcal{V}(\mathcal{L}^*)$ is the set of points in \mathbb{R}^n that are closer to $\mathbf{0}$ than any other lattice point of \mathcal{L}^* , under ℓ_2 norm.

With this tool in hand, the analysis of the protocol is very simple. By the maximum likelihood principle, the optimal prover guesses correctly if and only if the verifier’s original sample lands inside the Voronoi cell, and hence the verifier’s acceptance probability is exactly the Gaussian measure of $\mathcal{V}(\mathcal{L}^*)$. See Section 3 for further details.

For proving $(1 + o(1))$ -GapSPP \in coAM, we rely on the classic set-size lower bound protocol of Goldwasser and Sipser [GS86]. In order to prove that the discrete Gaussian mass on $\mathcal{L}^* \setminus \{0\}$ is large, we apply the protocol to thin shells in \mathbb{R}^n , and rely on a discrete Gaussian concentration inequality of Banaszczyk [Ban93]. See Section 6 for an overview and full details.

Statistical Zero Knowledge Protocol. We prove that $(2 + o(1))$ -GapSPP \in SZK, the class of problems having statistical zero-knowledge proofs. We note that this result does not subsume the inclusion in $\text{AM} \cap \text{coAM}$ described above (as one might suspect, given that $\text{SZK} \subseteq \text{AM} \cap \text{coAM}$), due to a slightly worse dependence ε in the $o(1)$ term. To prove the theorem, we construct a new instance-dependent commitment scheme³ based on GapSPP, which is sufficiently binding (for an honest committer) and hiding (to a dishonest receiver). Constructing such a commitment scheme (with some additional observations in our case) is known to be sufficient for obtaining an SZK protocol [IOS97].

Our construction can be viewed as a generalization of an instance-dependent commitment scheme for $O(\sqrt{n/\log n})$ -GapSVP implicit in [MV03], which was also based on the Goldreich-Goldwasser protocol and is perfectly binding. At a very high level, the commitment scheme is based on revealing a “random” perturbed lattice point in \mathcal{L} , where the perturbation is taken uniformly from a ball of radius r . Roughly speaking, we get the binding property when there is only one lattice within distance r of the revealed perturbation, and get the hiding property when there are multiple such lattice points (which allow for equivocation). It turns out that the main measure of quality for the binding and hiding property corresponds to the fraction of overlap between the balls of radius r placed around lattice points of \mathcal{L} : less overlap means better binding, and more overlap yields better hiding. In [MV03], this overlap is analyzed in terms of the length λ_1 of the shortest nonzero vector of \mathcal{L} . In particular, if $r \leq \lambda_1/2$, then the balls are completely disjoint (perfect binding), and if $r \geq \Omega(\sqrt{n/\log n}) \cdot \lambda_1$, then a $1/\text{poly}(n)$ fraction of the ball around any lattice point overlaps with that of its nearest neighbor in the lattice, which gives non-negligible hiding.

The main insight which allows us obtain improved approximation factors when basing the commitment scheme on GapSPP is a new characterization of the smoothing parameter, which allows to get very fine control on the overlap.

Ball Overlap Characterization. For $\varepsilon \geq 2^{-o(n)}$, Euclidean balls of radius $R = \sqrt{n/(2\pi)}/(2\eta_\varepsilon(\mathcal{L}^*))$ centered at all points of \mathcal{L} overlap in at most a 2ε fraction of their mass, and balls of radius $(2 + o(1))R$ overlap in at least an $\varepsilon/2$ fraction of their mass.

From the above we are able to determine, to within a factor $2 + o(1)$, whether balls placed at points of \mathcal{L} overlap in at most or at least an ε fraction of their mass, based solely on the smoothing parameter (of the dual lattice). Intuitively, this is because the smoothing parameter takes into account all the lattice points in \mathcal{L} , and hence is able to provide much better “global” information about the overlap. We refer the reader to Section 4 for further details and discussion.

Application to Worst-Case/Average-Case Reductions. As an application, we also obtain a worst-case to average-case reduction from GapSPP to the Learning With Errors problem (LWE) [Reg05], which has

³Roughly speaking, an instance-dependent commitment scheme for a language L is a commitment scheme that can depend on the instance x and such that only one of the (statistical) hiding and binding properties are required to hold, depending on whether $x \in L$.

a tighter connection factor than the known reductions from GapSVP [Reg05, Pei09]. Roughly speaking, the goal of LWE is to solve n -dimensional random noisy linear equations modulo some q , where Gaussian noise with standard deviation αq is added to each equation. The LWE problem is extremely versatile as a basis for numerous cryptographic constructions (e.g., [PW08, GPV08, CHKP10, BV11]). Regev’s celebrated result [Reg05] showed a quantum reduction from solving worst-case γ -GapSVP (among other problems) to solving LWE with $\gamma = \tilde{O}(n/\alpha)$. Furthermore, Peikert [Pei09] showed a corresponding classical reduction, when the modulus $q \geq 2^{n/2}$. Therefore, the security of LWE-based cryptographic constructions can be based on the worst-case hardness of the GapSVP problem.

We observe that the reductions of [Reg05, Pei09] in fact implicitly solve the GapSPP problem. Thus, by slightly modifying the last step of those reductions, we obtain corresponding quantum/classical reductions from γ -GapSPP $_\varepsilon$ (with $\varepsilon = \text{negl}(n)$) to LWE with $\gamma = O(\sqrt{n}/\alpha)$. As a consequence, the security of LWE-based cryptographic constructions can be based on the worst-case hardness of a potentially harder lattice problem.

The application to worst-case/average-case reduction follows by noting that the reduction of [Pei09] solves GapSVP by running the Goldreich-Goldwasser protocol, where the prover’s strategy is simulated by using a bounded distance decoding (BDD) oracle, which in turn is implemented using the LWE oracle. To obtain a tighter reduction from GapSPP to LWE, we observe that the quality of the BDD oracle depends directly on the smoothing parameter, as opposed to the length of the shortest vector. In light of this, we instead solve GapSPP using the Gaussian analogue of the Goldreich-Goldwasser protocol described above, while still using a bounded distance decoding (BDD) oracle to simulate the prover’s strategy. See Section 5 for further details.

Algorithm for GapSPP. We give a deterministic $2^{O(n)}$ polylog($1/\varepsilon$)-time and $2^{O(n)}$ -space algorithm for deciding $(1 + o(1))$ -GapSPP. For this we use recent algorithms of [MV10, DPV11] for enumerating lattice points in \mathcal{L}^* to estimate the Gaussian mass. The full details are in Section 7.2.

Perspectives and Open Questions. Our initial work on the complexity of the GapSPP problem opens up several directions for further study of the smoothing parameter from a computational perspective. Perhaps the most intriguing question is whether $(2 + o(1))$ -GapSPP is SZK-complete. A positive answer might lead to progress on the long-standing goal of basing cryptography on general complexity classes. Some reason for optimism comes from its rather unusual complexity: like SZK-complete problems, $(2 + o(1))$ -GapSPP is in SZK but is not known to be in NP or coNP. We are unaware of any other problems (aside from SZK-complete ones) having these characteristics.

In a related direction, in this work we focus on the standard “ L_∞ notion” of the smoothing parameter $\eta_\varepsilon(\mathcal{L})$, whereas the complexity of a related “ L_1 notion” of the smoothing parameter, denoted $\eta_\varepsilon^{(1)}(\mathcal{L})$, also seems quite interesting. More precisely, $\eta_\varepsilon(\mathcal{L})$ can be defined equivalently as the smallest parameter s such that the distribution of a continuous Gaussian of width s , reduced modulo \mathcal{L} , has point-wise probability density within a $(1 \pm \varepsilon)$ factor of that of the uniform distribution on \mathbb{R}^n/\mathcal{L} . The L_1 variant $\eta_\varepsilon^{(1)}(\mathcal{L})$ of the smoothing parameter instead is defined to be the smallest parameter s such that the statistical distance (i.e., half of the L_1 distance) between the above two distributions is at most ε . (Clearly, $\eta_\varepsilon^{(1)}(\mathcal{L}) \leq \eta_\varepsilon(\mathcal{L})$.) By definition, the problem of approximating the L_1 smoothing parameter, denoted γ -GapSPP $_\varepsilon^{(1)}$, appears to naturally reduce to a well-known SZK-complete problem called Statistical Difference (SD) problem [SV03], which is a promise problem asking whether two input distributions (specified by circuits) have statistical distance less than α or greater than β . Thus, the problem appears to be in SZK and is another candidate SZK-complete lattice problem. Unfortunately, the above argument relies on $\eta_\varepsilon^{(1)}(\mathcal{L})$ being a monotonic function

in ε , which is a basic property that we do not know how to prove (or disprove)! In fact, we know very little about the L_1 smoothing parameter. Given the potentially interesting complexity of γ -GapSPP $_\varepsilon^{(1)}$, it seems worthwhile to further investigate the L_1 smoothing parameter, from both the geometric and computational perspectives.

Finally, we note that our results generally apply only in the setting where $\varepsilon < 1$. It seems quite interesting to understand how the complexity of GapSPP changes for larger ε . We remark that our geometric characterizations only “half fail” for larger ε . More precisely, in the regime $\eta_\varepsilon(\mathcal{L}) \geq 1$, $\varepsilon \geq 1$, we still get upper bounds on the Gaussian measure of the Voronoi cell, as well as lower bounds on the fraction of overlap for balls centered at lattice points. For our AM protocol, this implies that the prover generally fails to convince the verifier, and for our instant-dependent commitment scheme, this implies that it is always hiding. Interestingly, our coAM protocol still applies for larger ε , almost without change. Here the main issue is that we do not know a “good” geometric interpretation of the statement $\rho(\mathcal{L} \setminus \{\mathbf{0}\}) \leq \varepsilon$ for any $\varepsilon \geq 1$.

Organization. The rest of the paper is organized as follows. In Section 2 we give the basic preliminaries. In Section 3, we give our Arthur-Merlin protocol for showing that $(2 + o(1))$ -GapSPP \in AM (Theorem 3.1). In Section 4 we construct a statistical zero-knowledge proof for GapSPP (Theorem 4.1). In Section 5, we describe the reduction from GapSPP to LWE (Theorem 5.5). In Section 6, we show that $(1 + o(1))$ -GapSPP \in coAM (Theorem 6.1). In Section 7.2 we give a deterministic algorithm for computing the smoothing parameter (Theorem 7.1).

2 Preliminaries

For sets $A, B \subseteq \mathbb{R}^n$, denote their Minkowski sum by $A + B = \{\mathbf{a} + \mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B\}$. We let $B_2^n = \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq 1\}$ denote the unit Euclidean ball in \mathbb{R}^n , and $S^{n-1} = \partial B_2^n$ the unit sphere in \mathbb{R}^n . Unless stated otherwise, $\|\cdot\|$ denotes the Euclidean norm.

Lattices. A lattice $\mathcal{L} \subset \mathbb{R}^n$ with basis B , and its dual \mathcal{L}^* , are defined as in the introduction. For a basis B and a vector $\mathbf{x} \in \mathbb{R}^n$, we let $\mathbf{x} \bmod B$ denote the unique $\bar{\mathbf{x}} \in \mathcal{L} + \mathbf{x}$ such that $\bar{\mathbf{x}} = \sum_{i=1}^n c_i \mathbf{b}_i$ for $c_i \in [-\frac{1}{2}, \frac{1}{2})$. It can be computed efficiently from \mathbf{x} and B (treated as matrix of column vectors) as $\bar{\mathbf{x}} = \mathbf{x} - B \lfloor B^{-1} \mathbf{x} \rfloor$. We sometimes instead write $\mathbf{x} \bmod \mathcal{L}$ when the basis is implicit.

The Voronoi cell $\mathcal{V}(\mathcal{L})$ is the set of points in \mathbb{R}^n that are at least as close to $\mathbf{0}$ (under the ℓ_2 norm) as to any other vector in \mathcal{L} :

$$\begin{aligned} \mathcal{V}(\mathcal{L}) &= \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\|_2 \leq \|\mathbf{x} - \mathbf{y}\|_2, \forall \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\} \\ &= \{\mathbf{x} \in \mathbb{R}^n : \langle \mathbf{x}, \mathbf{y} \rangle \leq \frac{1}{2} \langle \mathbf{y}, \mathbf{x} \rangle, \forall \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}. \end{aligned}$$

When the lattice in question is clear we shorten $\mathcal{V}(\mathcal{L})$ to \mathcal{V} . Note that \mathcal{V} is a symmetric polytope that tiles space with respect to \mathcal{L} , i.e., $\mathcal{L} + \mathcal{V} = \mathbb{R}^n$ and for all distinct $x, y \in \mathcal{L}$, the sets $x + \mathcal{V}$ and $y + \mathcal{V}$ are interior disjoint.

Gaussian measures. Define the Gaussian function $\rho: \mathbb{R}^n \rightarrow \mathbb{R}^+$ as $\rho(\mathbf{x}) = e^{-\pi\|\mathbf{x}\|^2}$, and for real $s > 0$, define $\rho_s(\mathbf{x}) = \rho(\mathbf{x}/s) = e^{-\pi\|\mathbf{x}\|^2/s^2}$. For a countable subset $A \subseteq \mathbb{R}^n$, we define $\rho_s(A) = \sum_{\mathbf{x} \in A} \rho_s(\mathbf{x})$.

For a measurable subset $A \subseteq \mathbb{R}^n$, we define the Gaussian measure of A (parameterized by $s > 0$) as $\gamma_s(A) = \frac{1}{s^n} \int_A \rho_s(\mathbf{x}) d\mathbf{x}$. Note that $\gamma_s(\mathbb{R}^n) = 1$, so γ_s is a probability measure. For parameter $s > 0$, we let

D_s be the corresponding continuous Gaussian distribution with parameter s centered around $\mathbf{0}$:

$$D_s(A) = \gamma_s(A) \quad \forall \text{ measurable } A \subseteq \mathbb{R}^n.$$

Similarly, for any countable subset $T \subseteq \mathbb{R}^n$ for which $\rho_s(T)$ converges, define the discrete Gaussian distribution $D_{T,s}$ over T by

$$D_{T,s}(\mathbf{x}) = \frac{\rho_s(\mathbf{x})}{\rho_s(T)} \quad \forall \mathbf{x} \in T.$$

We usually consider the discrete Gaussian over a lattice \mathcal{L} , i.e., where $T = \mathcal{L}$, though there will be situations where T corresponds a union of cosets of \mathcal{L} . In all these cases, $\rho_s(T)$ converges.

The following gives the standard concentration bounds for the continuous and discrete Gaussians.

Lemma 2.1 ([Ban93, Ban95]). *Let $X \in \mathbb{R}^n$ be distributed as D_s or $D_{\mathcal{L},s}$ for an n -dimensional lattice \mathcal{L} . For any $\mathbf{v} \in \mathbb{R}^n \setminus \{\mathbf{0}\}$ and $t > 0$, we have*

$$\Pr[\langle X, \mathbf{v} \rangle \geq t \|\mathbf{v}\|] \leq e^{-\pi(t/s)^2},$$

and for $\varepsilon > 0$ we have

$$\Pr[\|X\|^2 \geq (1 + \varepsilon)s^2 \frac{n}{2\pi}] \leq ((1 + \varepsilon)e^{-\varepsilon})^{n/2},$$

which for $0 < \varepsilon < \frac{1}{2}$ is bounded by $e^{-n\varepsilon^2/6}$.

The smoothing parameter. We recall the definition of the smoothing parameter from [MR04], and define the associated computational problem GapSPP.

Definition 2.2 (Smoothing Parameter). For a lattice \mathcal{L} and real $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ is the smallest $s > 0$ such that $\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon$.

Definition 2.3 (Smoothing Parameter Problem). For $\gamma = \gamma(n) \geq 1$ and positive $\varepsilon_Y = \varepsilon_Y(n), \varepsilon_N = \varepsilon_N(n)$ with $\varepsilon_Y \leq \varepsilon_N$, an instance of γ -GapSPP $_{\varepsilon_Y, \varepsilon_N}$ is a basis B of an n -dimensional lattice $\mathcal{L} = \mathcal{L}(B)$. It is a YES instance if $\eta_{\varepsilon_Y}(\mathcal{L}) \leq 1$, and is a NO instance if $\eta_{\varepsilon_N}(\mathcal{L}) > \gamma$. When $\varepsilon_Y = \varepsilon_N = \varepsilon$, we write γ -GapSPP $_\varepsilon$.

Notice that YES and NO instances are disjoint, since for a YES instance we have $\rho(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon_Y$, whereas for a NO instance we have $\rho(\mathcal{L}^* \setminus \{\mathbf{0}\}) \geq \rho_{1/\gamma}(\mathcal{L}^* \setminus \{\mathbf{0}\}) > \varepsilon_N \geq \varepsilon_Y$.

For the design and analysis of our interactive protocols, it is often convenient to use separate $\varepsilon_Y, \varepsilon_N$ parameters. The following lemma and its corollary then let us draw conclusions about GapSPP for a single ε parameter, for an (often slightly) larger approximation factor.

Lemma 2.4. *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be an n dimensional lattice. If $\rho_s(\mathcal{L} \setminus \{\mathbf{0}\}) \leq \varepsilon < 1$, then letting $t = \sqrt{1 + \log(r)/\log(\varepsilon^{-1})}$ for any $r \geq 1$, we have*

$$\rho_{s/t}(\mathcal{L} \setminus \{\mathbf{0}\}) \leq \frac{1}{r} \rho_s(\mathcal{L} \setminus \{\mathbf{0}\}) \leq \varepsilon/r.$$

Proof. By scaling \mathcal{L} , it suffices to prove the claim for $s = 1$. Since $t \geq 1$, we have

$$\begin{aligned} \rho_{1/t}(\mathcal{L} \setminus \{\mathbf{0}\}) &= \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} e^{-\pi\|\mathbf{y}\|^2} = \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} (e^{-\pi\|\mathbf{y}\|^2})^{t^2} \\ &\leq \left(\sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} e^{-\pi\|\mathbf{y}\|^2} \right)^{t^2} = \rho(\mathcal{L} \setminus \{\mathbf{0}\})^{t^2} \leq \rho(\mathcal{L} \setminus \{\mathbf{0}\}) \cdot \varepsilon^{t^2-1}. \end{aligned}$$

To finish the proof, note that $\varepsilon^{t^2-1} = 1/r$, as needed. □

Corollary 2.5. *For any $\varepsilon_N < 1$, there is a trivial reduction from γ' -GapSPP $_{\varepsilon_Y}$ to γ -GapSPP $_{\varepsilon_Y, \varepsilon_N}$, where $\gamma' = \gamma \cdot \sqrt{\log(\varepsilon_Y^{-1}) / \log(\varepsilon_N^{-1})}$.*

The proof follows by a routine calculation, letting $\varepsilon = \varepsilon_N$ and $r = \varepsilon_N / \varepsilon_Y$ in the above lemma. As a few notable examples, if ε_Y and ε_N are both fixed constants, then the loss γ' / γ in approximation factor from the reduction is a constant strictly greater than 1. But if ε_Y is constant and $\varepsilon_N = (1 + o(1)) \cdot \varepsilon_Y$, or if $\varepsilon_Y = o(1)$ and $\varepsilon_N \leq C \cdot \varepsilon_Y$ for a constant $C \geq 1$, then the loss in approximation factor is only $1 + o(1)$.

3 AM Protocol for GapSPP

Here we give an Arthur-Merlin protocol for 2-GapSPP $_{\varepsilon_Y, \varepsilon_N}$, defined formally in Protocol 1. It is simply a Gaussian variant of the classic Goldreich-Goldwasser protocol [GG98], which was originally developed to prove that approximating the Closest and Shortest Vector Problems to within a $c\sqrt{n/\log n}$ factor is in coAM. In our protocol, instead of choosing an error vector \mathbf{x} from the ball of radius $c\sqrt{n/\log n}$, the verifier chooses \mathbf{x} from a continuous Gaussian distribution of parameter 1. It then reduces \mathbf{x} modulo the lattice (actually, the dual lattice \mathcal{L}^* in our setting) and challenges the prover to find the original vector \mathbf{x} .

For intuition on why this protocol is complete and sound, first observe that the optimal prover strategy is maximum likelihood decoding of the verifier's challenge $\bar{\mathbf{x}} = \mathbf{x} \bmod \mathcal{L}^*$, i.e., to return a most-likely element in the coset $\mathbf{x}' \in \mathcal{L}^* + \bar{\mathbf{x}}$. Because the Gaussian function is decreasing in $\|\mathbf{x}'\|$, the prover should therefore return the shortest element $\mathbf{x}' \in \mathcal{L}^* + \bar{\mathbf{x}}$, i.e., the unique $\mathbf{x}' \in \mathcal{V}(\mathcal{L}^*) \cap (\mathcal{L}^* + \bar{\mathbf{x}})$. (We can ignore the measure-zero event that $\bar{\mathbf{x}}$ is equidistant from two or more points in \mathcal{L}^*). The verifier can therefore be made to accept with probability $\gamma(\mathcal{V}(\mathcal{L}^*))$, and no more. Note that unlike the original Goldreich-Goldwasser protocol, ours does not have perfect completeness, and in fact this is essential for establishing such a small approximation factor for GapSPP.

For completeness, consider a YES instance where $\eta_{\varepsilon_Y}(\mathcal{L}) \leq 1/2$, i.e., $\rho_2(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon_Y$. (For convenience, here we scale the 2-GapSPP problem so that NO instances have $\eta_{\varepsilon_N}(\mathcal{L}) > 1$.) Intuitively, because the measure on $\mathcal{L}^* \setminus \{\mathbf{0}\}$ is small, these lattice points are all far from the origin and so $\mathcal{V}(\mathcal{L}^*)$ captures most of the Gaussian measure γ ; Lemma 3.4 makes this formal. Finally, for soundness we consider the case where the discrete measure on nonzero lattice points is relatively large, i.e., $\rho_1(\mathcal{L}^* \setminus \{\mathbf{0}\}) > \varepsilon_N$. Conversely to the above, this intuitively means that the continuous Gaussian measure $\gamma(\mathcal{V}(\mathcal{L}^*))$ cannot be too large, and Lemma 3.4 again makes this precise.

Algorithm 1 Gaussian Goldreich-Goldwasser (GGG) Protocol

Input: Basis $B \subset \mathbb{R}^n$ of a lattice $\mathcal{L} = \mathcal{L}(B)$.

- 1: Verifier chooses Gaussian $\mathbf{x} \leftarrow D_1$ and sends $\bar{\mathbf{x}} = \mathbf{x} \bmod \mathcal{L}^*$ to prover.
 - 2: Prover returns an $\mathbf{x}' \in \mathbb{R}^n$.
 - 3: Verifier accepts if $\mathbf{x}' = \mathbf{x}$.
-

Theorem 3.1. *For $0 < \varepsilon \leq \delta < \frac{1}{2}$, Protocol 1 on lattice $\mathcal{L} = \mathcal{L}(B)$ satisfies:*

1. *Completeness: If $\eta_\varepsilon(\mathcal{L}) \leq \frac{1}{2}$, then there exists a prover that makes the verifier accept with probability at least $1 - \varepsilon$.*
2. *Soundness: If $\eta_{\frac{\delta}{1-\delta}}(\mathcal{L}) \geq 1$, then the verifier rejects with probability at least δ .*

In particular, $2\text{-GapSPP}_{\varepsilon, \delta/(1-\delta)} \in \text{AM}$ when $\delta - \varepsilon \geq 1/\text{poly}(n)$. Moreover, when $\varepsilon = \text{negl}(n)$ the protocol is honest-verifier statistical zero knowledge, i.e., $2\text{-GapSPP}_{\varepsilon, \delta/(1-\delta)} \in \text{HVSZK} = \text{SZK}$.

By applying Corollary 2.5, we obtain the following upper bounds on the complexity of $\gamma\text{-GapSPP}_\varepsilon$ for different ranges of ε .

Corollary 3.2. *For the following $\varepsilon(n) < 1$, we have $\gamma\text{-GapSPP}_\varepsilon \in \text{AM}$ for the following $\gamma(n)$:*

- If $\varepsilon(n) \leq \text{negl}(n)$, then $\gamma = O(\sqrt{\log(\varepsilon^{-1})/\log n})$.
- If $1/\text{poly}(n) \leq \varepsilon(n) \leq o(1)$, then $\gamma = (2 + o(1))$.
- If $\varepsilon(n) \geq \Omega(1)$, then $\gamma = O(1)$.

The next two lemmas provide the crux of the proof of Theorem 3.1.

Lemma 3.3. *Let $S \subseteq \mathbb{R}^n$ be symmetric (i.e., $S = -S$) measurable set. Then for any $\mathbf{y} \in \mathbb{R}^n$,*

$$\gamma_s(S + \mathbf{y}) \geq \gamma_s(S) \cdot \rho_s(\mathbf{y}).$$

Proof. By scaling S and \mathbf{y} , it suffices to prove the claim for $s = 1$. For any $t \in \mathbb{R}$, note that $\cosh(t) = \frac{1}{2}(e^t + e^{-t}) \geq 1$. We have

$$\begin{aligned} \gamma(S + \mathbf{y}) &= \int_S e^{-\pi\|\mathbf{y}-\mathbf{x}\|^2} dx = \int_S \frac{1}{2}(e^{-\pi\|\mathbf{y}-\mathbf{x}\|^2} + e^{-\pi\|\mathbf{y}+\mathbf{x}\|^2}) dx && \text{(symmetry of } S\text{)} \\ &= e^{-\pi\|\mathbf{y}\|^2} \int_S e^{-\pi\|\mathbf{x}\|^2} \cdot \frac{1}{2} \left(e^{2\pi\langle \mathbf{x}, \mathbf{y} \rangle} + e^{-2\pi\langle \mathbf{x}, \mathbf{y} \rangle} \right) dx && \text{(expanding the squares)} \\ &\geq \rho(\mathbf{y}) \int_S \rho(\mathbf{x}) dx = \rho(\mathbf{y}) \cdot \gamma(S). && \square \end{aligned}$$

The following crucial lemma establishes a tight relationship between discrete Gaussian sums on \mathcal{L} and the Gaussian mass of the Voronoi cell.

Lemma 3.4 (Voronoi Cell Characterization). *Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a lattice with Voronoi cell $\mathcal{V} = \mathcal{V}(\mathcal{L})$, and let $s > 0$. Then*

$$\frac{\rho_s(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho_s(\mathcal{L})} \leq 1 - \gamma_s(\mathcal{V}) \leq \rho_{2s}(\mathcal{L} \setminus \{\mathbf{0}\}).$$

In particular, letting $s_\varepsilon = \eta_\varepsilon(\mathcal{L}^)$ for some $\varepsilon \in (0, 1)$, we have that $\gamma(2s_\varepsilon\mathcal{V}) \geq 1 - \varepsilon$ and $\gamma(s_\varepsilon\mathcal{V}) \leq \frac{1}{1+\varepsilon}$.*

Proof. By scaling \mathcal{L} , it suffices to prove the claim for $s = 1$. We first show the upper bound. Let $X \in \mathbb{R}^n$ be distributed as D_1 , and note that $1 - \gamma(\mathcal{V}) = \Pr[X \notin \mathcal{V}]$. By the union bound and Lemma 2.1,

$$\begin{aligned} \Pr[X \notin \mathcal{V}] &= \Pr\left[\bigcup_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \{\langle X, \mathbf{y} \rangle > \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle\} \right] \leq \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \Pr[\langle X, \mathbf{y} \rangle > \frac{1}{2} \langle \mathbf{y}, \mathbf{y} \rangle] \\ &\leq \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} e^{-\pi\|\mathbf{y}/2\|^2} = \rho_2(\mathcal{L} \setminus \{\mathbf{0}\}). \end{aligned}$$

We now prove the lower bound. Since \mathcal{V} tiles space with respect to \mathcal{L} , by applying Lemma 3.3 with $S = \mathcal{V}$, we have

$$1 - \gamma(\mathcal{V}) = \gamma(\mathbb{R}^n \setminus \mathcal{V}) = \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \gamma(\mathcal{V} + \mathbf{y}) \geq \gamma(\mathcal{V}) \cdot \rho(\mathcal{L} \setminus \{\mathbf{0}\}),$$

Rearranging terms and using $\rho(\{\mathbf{0}\}) = 1$, we have $1 - \gamma(\mathcal{V}) \geq 1 - 1/\rho(\mathcal{L}) = \rho(\mathcal{L} \setminus \{\mathbf{0}\})/\rho(\mathcal{L})$, as desired. Finally, the ‘‘in particular’’ claim follows from $\gamma(r\mathcal{V}) = \gamma_{1/r}(\mathcal{V})$ and an easy calculation. \square

Proof of Theorem 3.1. As already argued above, the optimal prover strategy given $\bar{x} \in \mathbb{R}^n$ is maximum likelihood decoding, and the optimal prover can make the verifier accept with probability $\gamma(\mathcal{V}(\mathcal{L}^*))$. Completeness and soundness now follow immediately from Lemma 3.4, as already outlined in the overview.

For honest-verifier statistical zero-knowledge when $\varepsilon = \text{negl}(n)$, the simulator just chooses $\mathbf{x} \leftarrow D_1$ as the verifier's randomness, and outputs \mathbf{x} as the message from the prover. Because the prover also returns \mathbf{x} with probability at least $1 - \varepsilon$ in the real protocol, the simulated transcript is within negligible statistical distance of the real transcript. \square

4 SZK Protocol for GapSPP

This section is devoted to showing that $(2 + o(1))\text{-GapSPP}_{1/\text{poly}(n)}$ is in SZK.

Theorem 4.1. *For every $\varepsilon : \mathbb{N} \rightarrow [0, 1]$ such that $\frac{1}{\text{poly}(n)} \leq \varepsilon(n) \leq \frac{1}{36}$, we have $2 \cdot (1 + \delta)\text{-GapSPP}_{\varepsilon, 12\varepsilon} \in \text{SZK}$, where $\delta = \sqrt{\frac{3}{2n} \ln \frac{4}{\varepsilon}}$.*

As before, the following corollary gives the implied upper bound on the complexity of $\gamma\text{-GapSPP}_\varepsilon$ (by applying Corollary 2.5).

Corollary 4.2. *For every $\varepsilon : \mathbb{N} \rightarrow (0, 1)$, if $1/\text{poly}(n) \leq \varepsilon(n) \leq o(1)$, then $(2 + o(1))\text{-GapSPP}_\varepsilon \in \text{SZK}$. If $\varepsilon(n) \leq \text{negl}(n)$, then $O\left(\sqrt{\frac{\log(1/\varepsilon)}{\log n}}\right)\text{-GapSPP}_\varepsilon \in \text{SZK}$. Finally, if $\Omega(1) \leq \varepsilon(n) \leq 1/3$, then $O(1)\text{-GapSPP}_\varepsilon \in \text{SZK}$.*

Our construction follows a classic approach of constructing an *instance-dependent (ID) commitment scheme* for GapSPP, which is known to be sufficient for obtaining a SZK protocol [IOS97]. With an additional observation, we show that a *significantly weaker* notion of ID commitment schemes is sufficient to obtain SZK protocols; roughly speaking, we only need an ID bit-commitment scheme that is sufficiently binding for an *honest* sender, and hiding (from a dishonest receiver). Specifically, it is sufficient to construct a “non-trivial” ID commitment scheme defined as follows.

Definition 4.3. Let Π be a promise problem. A (non-interactive) instance-dependent bit-commitment scheme Com for Π is a PPT algorithm that on input an instance $x \in \{0, 1\}^n$ and a bit $b \in \{0, 1\}$, outputs a commitment $\text{Com}_x(b) \in \{0, 1\}^*$. Let $p = p(n), q = q(n) \in (0, 1)$. We define (weak) binding and hiding properties of Com as follows.

- *Statistical honest-sender q -binding for YES instances:* For every $x \in \Pi_Y$ and $b \in \{0, 1\}$,

$$\Pr[\text{Com}_x(b) \in \text{supp}(\text{Com}_x(\bar{b}))] \leq q(|x|).$$

(Note that when $\text{Com}_x(b) \notin \text{supp}(\text{Com}_x(\bar{b}))$, the commitment $\text{Com}_x(b)$ cannot be opened to \bar{b} . Thus, the above condition implies that the binding property can be broken with probability at most q .)

- *Statistical p -hiding for NO instances:* For every $x \in \Pi_N$,

$$\Delta(\text{Com}_x(0), \text{Com}_x(1)) \leq p(|x|).$$

(The above condition implies that given $\text{Com}_x(b)$ for a random b , one can only predict b correctly with probability at most $(1 + p)/2$, which means that the hiding property can be broken with advantage at most p .)

Com is *non-trivial* if Com is statistical p -hiding and statistical honest-sender q -binding with $p + q \leq 1 - 1/\text{poly}(n)$.⁴ Com is *secure* if Com is statistical p -hiding and statistical honest-sender q -binding with negligible p and q .

In the next subsection, we focus on constructing a non-trivial ID commitment schemes for $(2 + o(1))$ -GapSPP $_{1/\text{poly}(n)}$. We present more detailed background for ID commitment schemes and discuss why it is sufficient to construct SZK protocols in Section 4.3.

4.1 A Non-Trivial ID Commitment Scheme for GapSPP

In this section, we construct a non-trivial ID bit-commitment scheme SPCom for $(2 + o(1))$ -GapSPP $_{1/\text{poly}(n)}$. Our construction can be viewed as a generalization of an instance-dependent commitment scheme implicit in [MV03] for $O(\sqrt{n/\log n})$ -GapSVP.⁵ To explain the intuition behind our construction, it is instructive to first consider the construction of ID commitment scheme for GapSVP (for simplicity, below we describe commitment to a random b): To commit, a sender first selects a “random” lattice point $\mathbf{v} \in \mathcal{L}$ (see Figure 4.1 for the precise distribution) and adds a random noise vector \mathbf{e} drawn from a ball of certain radius (say, $r = 1/2$) to \mathbf{v} ; let $\mathbf{w} = \mathbf{v} + \mathbf{e}$. Intuitively, the vector \mathbf{w} is binding to \mathbf{v} if the noise is sufficiently short. To actually commit to a bit, the sender also samples a random hash function h , and commits to the hashed bit $b = h(\mathbf{v})$. Namely, (\mathbf{w}, h) is a commitment $\text{Com}_{\mathcal{L}}(b)$ to $b = h(\mathbf{v})$.

Intuitively, if the length of the shortest vector $\lambda_1(\mathcal{L}) \geq 1$, then all balls centered at lattice points $\mathbf{v} \in \mathcal{L}$ of radius $r = 1/2$ are disjoint, and thus $\text{Com}_{\mathcal{L}}(b) = (\mathbf{w}, h)$ is perfect binding. On the other hand, if the shortest vector is too short, say, $\lambda_1(\mathcal{L}) \leq O(\sqrt{(\log n)/n})$, then \mathbf{w} may fall in the intersection region of two (or more) balls with non-negligible probability, using the symmetry of the lattice and the fact that the balls centered around the origin and a shortest non-zero vector have non-negligible overlap. When \mathbf{w} lies in the balls centered at \mathbf{v}_1 and \mathbf{v}_2 and $h(\mathbf{v}_1) \neq h(\mathbf{v}_2)$, the commitment $\text{Com}_{\mathcal{L}}(b) = (\mathbf{w}, h)$ does not reveal the committed value b , which intuitively achieves hiding. Indeed, the above argument can be formalized readily, yielding an ID bit-commitment scheme for $O(\sqrt{n/\log n})$ -GapSVP with perfect binding and weak hiding properties.

Note that in the above commitment scheme, the quality of the hiding property depends on how much the ball $\mathbf{v} + rB_2^n$ overlaps with the balls around surrounding lattice points. However, in the above analysis, we only exploit the overlap contributed by a nearest lattice point to \mathbf{v} , ignoring the overlap contributed by all other balls. In general, such an approach can only give a very coarse approximation of the overlap, which one can see from the example of extremal lattices where there are exponentially many lattice points of length roughly equal to that of the shortest vector. As a result, using this approach one can only obtain a non-trivial ID bit-commitment scheme for γ -GapSVP with $\gamma \geq \Omega(\sqrt{n/\log n})$.

Our key observation is that, when we switch from GapSVP to GapSPP, the above construction gives a non-trivial ID bit-commitment scheme for γ -GapSPP $_{1/\text{poly}(n)}$ with $\gamma = 2 + o(1)$. This stems from our new ball overlap characterization of the smoothing parameter, which gives us much finer control on the amount overlap we obtain in the above protocol. We formalize this characterization as follows:

⁴This is in contrast to the fact that one can construct a (trivial) p -hiding and q -binding commitment scheme for every $p + q \geq 1$. For example, defining $\text{Com}_x(b) = b$ gives $p = 1$ and $q = 0$, and defining $\text{Com}_x(b) = 0$ gives $p = 0$ and $q = 1$. More generally, defining $\text{Com}_x(b)$ to be b with probability α and 00 with probability $1 - \alpha$ gives $p = \alpha$ and $q = 1 - \alpha$.

⁵While [MV03] constructed their protocol by combining the reduction from GapSVP to GapCVP with Goldreich-Levin hardcore predicate, their construction can be interpreted as implicitly constructing an ID bit-commitment scheme for GapSVP by first constructing one with *perfect* binding but weak hiding, and then amplifying the hiding property.

Lemma 4.4 (Ball Overlap Characterization). Let \mathcal{L} be an n dimensional lattice. For $r > 0$, define

$$\text{Overlap}(\mathcal{L}, r) \stackrel{\text{def}}{=} \frac{\text{vol}_n \left(\bigcup_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} (rB_2^n \cap (rB_2^n + \mathbf{y})) \right)}{\text{vol}_n(rB_2^n)},$$

which denotes the fraction of overlap of a ball of radius r centered at a point in \mathcal{L} with balls of equal radius centered at all other lattice points. Then for $\varepsilon \in (2^{o(-n)}, 1/3)$, setting $r_\varepsilon = \sqrt{\frac{n}{2\pi}} / (2\eta_\varepsilon(\mathcal{L}^*))$, the following holds:

1. For $0 \leq r \leq r_\varepsilon$, we have $\text{Overlap}(\mathcal{L}, r) \leq 2\varepsilon$.
2. For any $r \geq 2(1 + \delta) \cdot r_\varepsilon$ where $\delta = \sqrt{\frac{3}{2n} \ln \frac{4}{\varepsilon}}$, we have $\text{Overlap}(\mathcal{L}, r) \geq \varepsilon/2$.

The above lemma says that up to a factor of $2 + o(1)$, the smoothing parameter $\eta_\varepsilon(\mathcal{L}^*)$ characterizes the required radius for balls on \mathcal{L} to have roughly ε fraction of overlap. As we shall see shortly, the amount of overlap tightly characterizes the binding and hiding property of the commitment scheme described above. As such, by choosing ε_Y and ε_N with a small constant factor gap, the above construction yields a non-trivial ID bit-commitment scheme for γ -GapSPP $_{\varepsilon_Y, \varepsilon_N}$ with $\gamma = 2 + o(1)$.

We proceed to formally define our ID bit-commitment scheme SPCom for GapSPP in Fig 1, and establish its binding and hiding properties. We prove the binding and hiding properties in Lemma 4.5 and 4.6, respectively, and summarize the properties of SPCom in Lemma 4.10. We defer the proofs of all geometric lemmas (in particular, the Ball Overlap Characterization) to subsection 4.2.

We remark that since we are approximating $\eta_\varepsilon(\mathcal{L})$, the following protocol operates directly on \mathcal{L}^* . For simplicity of notation, for a basis B of \mathcal{L} , we write $B^* = (B^{-1})^t$ to denote the corresponding dual basis of \mathcal{L}^* .

Let $\mathcal{H} = \{h : \{0, 1\}^n \rightarrow \{0, 1\}\}$ be a pairwise-independent hash family.
On input a lattice basis B and a bit $b \in \{0, 1\}$,

- Sample uniformly random $\mathbf{z} \leftarrow \{0, 1\}^n$ and $h \leftarrow \mathcal{H}$ jointly subject to $h(\mathbf{z}) = b$. (This can be done by rejection sampling, or equivalently by sampling uniform $\mathbf{z} \leftarrow \{0, 1\}^n$ first, and then sampling $h \leftarrow \mathcal{H}$ conditioned on $h(\mathbf{z}) = b$.)
- Sample $\mathbf{e} \leftarrow rB_2^n$ with $r = \frac{1}{2} \sqrt{\frac{n}{2\pi}}$.
- Let $\mathbf{v} = B^* \mathbf{z}$ and $\mathbf{w} = (\mathbf{v} + \mathbf{e} \bmod 2B^*)$.
- Output $\text{SPCom}_B(b) = (\mathbf{w}, h)$.

Figure 1: SPCom: a non-trivial ID commitment scheme for GapSPP.

The following two technical lemmas establish the (weak) binding and hiding properties of SPCom.

Lemma 4.5. For every $b \in \{0, 1\}$,

$$\Pr[\text{SPCom}_B(b) \in \text{supp}(\text{SPCom}_B(\bar{b}))] \leq \text{Overlap}(\mathcal{L}^*, r).$$

Proof. Let $S = \bigcup_{\mathbf{y} \in \mathcal{L}^* \setminus \{0\}} (rB_2^n \cap (rB_2^n + \mathbf{y}))$. By definition, $\text{SPCom}_B(b)$ is generated by sampling $\mathbf{e} \leftarrow rB_2^n$ and $h \leftarrow \mathcal{H}, \mathbf{z} \in \{0, 1\}^n$ such that $h(\mathbf{z}) = b$, and outputting $(\mathbf{w}, h) = (\mathbf{v} + \mathbf{e} \bmod 2B^*, h)$, where $\mathbf{v} = B^*\mathbf{z}$. Thus, we can write $\mathbf{w} = \mathbf{v} + \mathbf{u} + \mathbf{e}$ for some $\mathbf{u} \in 2\mathcal{L}^*$.

The event $\text{SPCom}_B(b) \in \text{supp}(\text{SPCom}_B(\bar{b}))$ means that there exists some $\mathbf{z}' \in \{0, 1\}^n, \mathbf{e}' \in rB_2^n$ such that $h(\mathbf{z}') = \bar{b}$ and $w = (\mathbf{v}' + \mathbf{e}' \bmod 2B^*)$, where $\mathbf{v}' = B^*\mathbf{z}'$. Similarly, we can write $\mathbf{w} = \mathbf{v}' + \mathbf{u}' + \mathbf{e}'$ for some $\mathbf{u}' \in 2\mathcal{L}^*$.

Let $\mathbf{y} = \mathbf{v}' + \mathbf{u}' - \mathbf{v} - \mathbf{u}$, and note that $\mathbf{y} \in \mathcal{L}^*$. The facts that $\mathbf{w} = \mathbf{v} + \mathbf{u} + \mathbf{e} = \mathbf{v}' + \mathbf{u}' + \mathbf{e}'$ and $\mathbf{e}' \in rB_2^n$ imply that $\mathbf{e} \in rB_2^n + \mathbf{y}$, which implies $\mathbf{e} \in S$. As the event in the LHS implies $\mathbf{e} \in S$, it follows that

$$\Pr[\text{SPCom}_B(b) \in \text{supp}(\text{SPCom}_B(\bar{b}))] \leq \Pr[\mathbf{e} \in S] = \frac{\text{vol}_n(S)}{\text{vol}_n(rB_2^n)} = \text{Overlap}(\mathcal{L}^*, r). \quad \square$$

Lemma 4.6.

$$\Delta(\text{SPCom}_B(0), \text{SPCom}_B(1)) \leq 1 - (\text{Overlap}(\mathcal{L}^*, r)/2).$$

Proof. Let $S = \bigcup_{\mathbf{y} \in \mathcal{L}^* \setminus \{0\}} (rB_2^n \cap (rB_2^n + \mathbf{y}))$. Define random variables $(W_0, H_0) = \text{SPCom}_B(0)$ and $(W_1, H_1) = \text{SPCom}_B(1)$. Observe that the marginal distributions of W_0 and W_1 are identical (following by the fact that $h \leftarrow \mathcal{H}$ maps every $\mathbf{z} \in \{0, 1\}^n$ to a uniformly random bit), we have

$$\begin{aligned} \Delta(\text{SPCom}_B(0), \text{SPCom}_B(1)) &= (1/2) \cdot \sum_{\mathbf{w}, h} \left| \Pr_{W_0, H_0}[(\mathbf{w}, h)] - \Pr_{W_1, H_1}[(\mathbf{w}, h)] \right| \\ &= (1/2) \cdot \sum_{\mathbf{w}} \Pr_{W_0}[\mathbf{w}] \cdot \sum_h \left| \Pr_{H_0}[h|W_0 = \mathbf{w}] - \Pr_{H_1}[h|W_1 = \mathbf{w}] \right| \\ &= \mathbb{E}_{\mathbf{w} \leftarrow W_0} [\Delta(H_0|W_0=\mathbf{w}, H_1|W_1=\mathbf{w})] \end{aligned}$$

For every $\mathbf{w} \in (\mathbb{R}^n \bmod 2B^*)$, define $T_{\mathbf{w}} = (rB_2^n + \mathbf{w}) \cap \mathcal{L}^*$ and $T'_{\mathbf{w}} = (T_{\mathbf{w}} \bmod 2B^*)$. we rely on the following two technical claims to upper bound the statistical distance. Note that the event $|T_{\mathbf{w}}| \geq 2$ is equivalent to the event $\mathbf{e} \in S$, where \mathbf{e} is the error vector used to generate \mathbf{w} , and hence

$$\Pr[|T_{\mathbf{w}}| \geq 2] = \Pr[\mathbf{e} \in S] = \text{Overlap}(\mathcal{L}^*, r).$$

Claim 4.7.

$$\Pr_{\mathbf{w} \leftarrow W_0}[|T'_{\mathbf{w}}| \geq 2] = \Pr_{\mathbf{w} \leftarrow W_0}[|T_{\mathbf{w}}| \geq 2].$$

Claim 4.8. For every $\mathbf{w} \in (\mathbb{R}^n \bmod 2B^*)$ with $|T'_{\mathbf{w}}| \geq 2$,

$$\Delta(H_0|W_0=\mathbf{w}, H_1|W_1=\mathbf{w}) \leq 1/2.$$

The above two claims imply that

$$\begin{aligned} \Delta(\text{SPCom}_B(0), \text{SPCom}_B(1)) &\leq \Pr_{\mathbf{w} \leftarrow W_0}[|T'_{\mathbf{w}}| \geq 2] \cdot (1/2) + \Pr_{\mathbf{w} \leftarrow W_0}[|T'_{\mathbf{w}}| = 1] \cdot 1 \\ &= \text{Overlap}(\mathcal{L}^*, r) \cdot (1/2) + (1 - \text{Overlap}(\mathcal{L}^*, r)) \\ &= 1 - \text{Overlap}(\mathcal{L}^*, r)/2, \end{aligned}$$

as desired. It remains to prove the claims.

Proof. (of Claim 4.7) Let $\{\mathbf{y}_1, \dots, \mathbf{y}_t\} = T_{\mathbf{w}} = \mathcal{L}^* \cap (rB2^n + \mathbf{w})$, where the \mathbf{y}_i s are ordered such that $\|\mathbf{y}_i - \mathbf{w}\|_2 \leq \|\mathbf{y}_{i+1} - \mathbf{w}\|_2$. By assumption, we have that $t \geq 2$. To prove that $|T'_{\mathbf{w}}| = |T_{\mathbf{w}} \pmod{2B^*}| \geq 2$, it suffices to show that $\mathbf{y}_1 \neq \mathbf{y}_2 \pmod{2B^*}$. Assume not, then note that $\bar{\mathbf{y}} = \frac{1}{2}(\mathbf{y}_1 + \mathbf{y}_2) \in \mathcal{L}^*$, $\bar{\mathbf{y}} \neq \mathbf{y}_1$, $\bar{\mathbf{y}} \neq \mathbf{y}_2$. Furthermore, by the triangle inequality

$$\|\bar{\mathbf{y}} - \mathbf{w}\|_2 = \left\| \frac{1}{2}(\mathbf{y}_1 + \mathbf{y}_2) - \mathbf{w} \right\|_2 \leq \frac{1}{2} \|\mathbf{y}_1 - \mathbf{w}\|_2 + \frac{1}{2} \|\mathbf{y}_2 - \mathbf{w}\|_2 \leq \|\mathbf{y}_2 - \mathbf{w}\|_2 \leq r$$

Hence $\bar{\mathbf{y}} \in T_{\mathbf{w}}$. We now examine two cases. If $\|\mathbf{y}_1 - \mathbf{w}\|_2 = \|\mathbf{y}_2 - \mathbf{w}\|_2$, then since $\mathbf{y}_1 \neq \mathbf{y}_2$, the above inequality must hold strictly. But then $\|\bar{\mathbf{y}} - \mathbf{w}\|_2 < \|\mathbf{y}_1 - \mathbf{w}\|_2$, which contradicts the fact that \mathbf{y}_1 is a closest lattice vector to \mathbf{w} . If $\|\mathbf{y}_1 - \mathbf{w}\|_2 < \|\mathbf{y}_2 - \mathbf{w}\|_2$, then $\|\bar{\mathbf{y}} - \mathbf{w}\|_2 < \|\mathbf{y}_2 - \mathbf{w}\|_2$, which contradicts that \mathbf{y}_2 is a closest lattice vector to \mathbf{w} after \mathbf{y}_1 . The claim thus follows. \square

Proof. (of Claim 4.8) Let $T'_{\mathbf{w}} = \{\mathbf{v}_1, \dots, \mathbf{v}_t\}$ and let \mathbf{z}_i be the coordinates of \mathbf{v}_i with respect to the basis B^* for every $i \in [t]$. Note that by construction, conditioned on \mathbf{w} , the random variable $\mathbf{z} \in \{0, 1\}^n$ becomes uniform over the $\{\mathbf{z}_1, \dots, \mathbf{z}_t\}$.

Now, consider a probability space \mathcal{P} defined by independent random variables (I, H) , where I is a uniformly random index in $[t]$ and H is a random hash function in \mathcal{H} . Define a random variable $B = H(\mathbf{z}_I)$. Note that by the construction, for $b \in \{0, 1\}$, the random variable $H_b|_{W_b=\mathbf{w}}$ has identical distribution to the random variable $H|_{B=b}$ in \mathcal{P} . Thus, our goal can be rephrased as to upper bound $\Delta(H|_{B=0}, H|_{B=1})$.

By Bayes' rule,

$$\Pr[H|_{B=b} = h] = \frac{\Pr[H = h] \Pr[B = b|H = h]}{\Pr[B = b]} = 2 \Pr[H = h] \cdot \frac{\#\{i : h(\mathbf{z}_i) = b\}}{t},$$

thus

$$\Delta(H|_{B=0}, H|_{B=1}) = \sum_h \Pr[H = h] \frac{|\#\{i : h(\mathbf{z}_i) = 0\} - \#\{i : h(\mathbf{z}_i) = 1\}|}{t}.$$

Intuitively, since \mathcal{H} is a pairwise-independent hash-family, the discrepancy $\#\{i : h(\mathbf{z}_i) = 0\} - \#\{i : h(\mathbf{z}_i) = 1\}$ should be small on expectation. We prove this presently.

Claim 4.9. For $t \geq 2$, $\mathbb{E}_{h \leftarrow \mathcal{H}}[|\#\{i : h(\mathbf{z}_i) = 0\} - \#\{i : h(\mathbf{z}_i) = 1\}|] \leq \frac{t}{2}$

Proof. For $i \in [t]$, let $X_i = (-1)^{h(\mathbf{z}_i)} \in \{-1, 1\}$. Since $\Pr[h(\mathbf{z}_i) = 0] = \Pr[h(\mathbf{z}_i) = 1] = 1/2$, we have that $\mathbb{E}[X_i] = 0$ for all $i \in [t]$. Furthermore, by pairwise independence we have that $\mathbb{E}[X_i X_j] = \mathbb{E}[X_i] \mathbb{E}[X_j] = 0$ for distinct $i, j \in [t]$. By definition of X_i , it is easy to verify that

$$|\#\{i : h(\mathbf{z}_i) = 0\} - \#\{i : h(\mathbf{z}_i) = 1\}| = \left| \sum_{i=1}^t X_i \right|.$$

By Jensen's inequality and pairwise independence, we obtain the inequality

$$\mathbb{E} \left[\left[\sum_{i=1}^t X_i \right]^2 \right] \leq \mathbb{E} \left[\left(\sum_{i=1}^t X_i \right)^2 \right] = \sum_{1 \leq i, j \leq t} \mathbb{E}[X_i X_j] = \sum_{i=1}^t \mathbb{E}[X_i^2] = t$$

Taking a square root, the above inequality gives us $\mathbb{E}[|\#\{i : h(\mathbf{z}_i) = 0\} - \#\{i : h(\mathbf{z}_i) = 1\}|] \leq \sqrt{t}$. Since $\sqrt{t} \leq t/2$ for $t \geq 4$, the claim holds for all $t \geq 4$.

It remains to prove the claim for $t = 2, 3$. For $t = 2$, h acts like a truly random hash function, and hence a direct computation yields

$$\mathbb{E}[|X_1 + X_2|] = 2 \Pr[X_1 = X_2] + 0 \Pr[X_1 \neq X_2] = 2(1/2) + 0 = 1, \quad (4.1)$$

as needed. For the case $t = 3$, we have that

$$\begin{aligned} \mathbb{E}[|X_1 + X_2 + X_3|] &= 3 \Pr[X_1 = X_2 = X_3] + 1 \Pr[X_1, X_2, X_3 \text{ not all equal}] \\ &= 3 \Pr[X_1 = X_2 = X_3] + 1(1 - \Pr[X_1 = X_2 = X_3]) = 1 + 2 \Pr[X_1 = X_2 = X_3]. \end{aligned}$$

By inclusion exclusion we get that

$$\begin{aligned} 1 &= \Pr[\exists X_i = 1] + \Pr[X_1 = X_2 = X_3 = -1] \\ &= \sum_{i=1}^3 \Pr[X_i = 1] - \sum_{1 \leq i < j \leq 3} \Pr[X_i = X_j = 1] + \Pr[X_1 = X_2 = X_3 = 1] + \Pr[X_1 = X_2 = X_3 = -1] \\ &= \sum_{i=1}^3 \Pr[X_i = 1] - \sum_{1 \leq i < j \leq 3} \Pr[X_i = X_j = 1] + \Pr[X_1 = X_2 = X_3] \end{aligned}$$

By rearranging the above equality and using pairwise independence, we get

$$\Pr[X_1 = X_2 = X_3] = 1 - \sum_{i=1}^3 \Pr[X_i = 1] + \sum_{1 \leq i < j \leq 3} \Pr[X_i = X_j = 1] = 1 - 3(1/2) + 3(1/4) = \frac{1}{4} \quad (4.2) \quad \square$$

Combining Equations (4.1) and (4.2), we get that $\mathbb{E}[|X_1 + X_2 + X_3|] = 1 + 2(1/4) = 3/2$, as needed. \square

From the above claim, we observe that $\Delta(H|_{B=0}, H|_{B=1}) \leq \frac{t/2}{t} \leq 1/2$ for every $t \geq 2$ as needed. \square

Finally, we prove the ID binding and hiding properties of SPCom by Lemma 4.4, 4.5, and 4.6.

Lemma 4.10. *For every $\varepsilon : \mathbb{N} \rightarrow [0, 1]$ such that $1/\text{poly}(n) \leq \varepsilon(n) \leq 1/36$, SPCom is a non-trivial ID commitment scheme for $2 \cdot (1 + \delta)$ -GapSPP $_{\varepsilon, 12\varepsilon}$ with $\delta = \sqrt{\frac{3}{2n} \ln \frac{4}{\varepsilon}}$. Specifically, SPCom is (2ε) -binding for the YES-instances and $(1 - 3\varepsilon)$ -hiding for the NO-instances of $2 \cdot (1 + \delta)$ -GapSPP $_{\varepsilon, 12\varepsilon}$, respectively.*

Proof. For YES-instances where $\eta_\varepsilon(\mathcal{L}) \leq 1$, by Part 1. of Lemma 4.4 and noting that $r \geq r_\varepsilon$,

$$\text{Overlap}(\mathcal{L}, r) \leq 2\varepsilon.$$

Thus, by Lemma 4.5, SPCom is (2ε) -binding for the YES-instances. On the other hand, for NO-instances where $\eta_\varepsilon(\mathcal{L}) \geq 2 \cdot (1 + \delta)$, by Part 2. of Lemma 4.4 and noting that $r \geq 2 \cdot (1 + \delta) \cdot r_\varepsilon$,

$$\text{Overlap}(\mathcal{L}, r) \geq 12\varepsilon/2 = 6\varepsilon.$$

Thus, by Lemma 4.6, Com is $(1 - 3\varepsilon)$ -hiding for the NO-instances. \square

Theorem 4.1 then follows by combining Lemma 4.10 and Theorem 4.14 stated in the next section. We remark that our SZK protocol for $(2 + o(1))$ -GapSPP $_{1/\text{poly}(n)}$ does not have efficient prover strategy, since we do not know if the problem is in NP or coNP.

4.2 Geometric Lemmas

Here we prove several geometric lemmas with the goal of establishing the ball overlap characterization (Lemma 4.4). The first lemma gives a standard upper bound on the volume of the intersection of two euclidean balls (see [Bal97, Lemma 2.2], noting that the ball intersection volume is at most twice that of the spherical cap).

Lemma 4.11. *For $r > 0$, let $s = r\sqrt{2\pi/n}$. Then for any $\mathbf{y} \in \mathbb{R}^n$,*

$$\frac{\text{vol}_n(rB_2^n \cap (\mathbf{y} + rB_2^n))}{\text{vol}_n(rB_2^n)} \leq 2e^{-\pi\|\frac{\mathbf{y}}{2s}\|^2}.$$

The following lemma will allow us to transfer lower bounds on the gaussian measure of the overlap region to lower bounds on uniform measure, and will be important in the proof of the ball overlap characterization.

Lemma 4.12. *Let $K \subseteq \mathbb{R}^n$ be a convex body containing the origin. Then for any $r, s > 0$ we have that*

$$\frac{\gamma_s(rB_2^n \setminus K)}{\gamma_s(rB_2^n)} \leq \frac{\text{vol}_n(rB_2^n \setminus K)}{\text{vol}_n(rB_2^n)}$$

Proof. The plan of the proof is to show that since gaussian measure is more biased towards the origin than the uniform measure, switching from gaussian to uniform pushes measures outside of K . To make this precise, we note that

$$\frac{\gamma_s(rB_2^n \setminus K)}{\gamma_s(rB_2^n)} = \int_{S^{n-1}} \int_0^r I[t\theta \notin K] \frac{e^{-\pi\frac{t^2}{s^2}}}{\gamma_s(rB_2^n)s^n} t^{n-1} dt d\theta \quad (4.3)$$

where $d\theta$ is the Haar measure on the unit sphere $S^{n-1} \subseteq \mathbb{R}^n$. Furthermore, note that

$$\frac{\text{vol}_n(rB_2^n \setminus K)}{\text{vol}_n(rB_2^n)} = \int_{S^{n-1}} \int_0^r I[t\theta \notin K] \frac{1}{\text{vol}_n(rB_2^n)} t^{n-1} dt d\theta \quad (4.4)$$

Since both the uniform and gaussian measure are spherically symmetric, we must have that

$$\int_0^r \frac{1}{\text{vol}_n(rB_2^n)} t^{n-1} dt = \int_0^r \frac{e^{-\pi(t/s)^2}}{s^n \gamma_s(rB_2^n)} t^{n-1} dt \quad (4.5)$$

Let $f, g : [0, r] \rightarrow \mathbb{R}_+$ be defined by $f(t) = \frac{1}{\text{vol}_n(rB_2^n)}$ and $g(t) = \frac{e^{-\pi(t/s)^2}}{s^n \gamma_s(rB_2^n)}$. Since f is constant, g is decreasing, and $\int_0^r f(t)t^{n-1} dt = \int_0^r g(t)t^{n-1} dt$ (by equation (4.5)), there exists $b \in (0, r)$ such that $f(t) \leq g(t)$ on $[0, b]$ and $f(t) \geq g(t)$ on $(b, r]$. Given this geometry, we clearly have that for all $c \in [0, r]$

$$\int_c^r f(t)t^{n-1} dt \geq \int_c^r g(t)t^{n-1} dt \quad (4.6)$$

Since $K \subseteq \mathbb{R}^n$ is a convex body containing $\mathbf{0}$, for every line segment $\{t\theta : 0 \leq t \leq r\}$, we have that $\{t\theta : 0 \leq t \leq r\} \setminus K = \{t\theta : c \leq t \leq r\}$ for some $c > 0$. From the above inequality (equation (4.6)), we now see that

$$\int_0^r I[t\theta \notin K] f(t)t^{n-1} dt \geq \int_0^r I[t\theta \notin K] g(t)t^{n-1} dt \quad (4.7)$$

for all $\theta \in S^{n-1}$. Combining equations (4.3) and (4.4) with the above inequality (4.7) yields the result. \square

The following lemma establishes the necessary bounds for our ball overlap characterization of the smoothing parameter.

Lemma 4.13. *Let \mathcal{L} denote an n -dimensional lattice. For $r > 0$, $0 < \delta < \frac{1}{4}$ and $s = r\sqrt{\frac{2\pi}{n}}$, the following holds:*

$$\left(\frac{\rho_{\frac{s}{1+\delta}}(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho_{\frac{s}{1+\delta}}(\mathcal{L})} - e^{-\frac{2n}{3}\delta^2} \right) \leq \text{Overlap}(\mathcal{L}, r) \leq 2\rho_{2s}(\mathcal{L} \setminus \{\mathbf{0}\}).$$

Proof. For simplicity of notation, we define $S = \bigcup_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} (rB_2^n \cap (rB_2^n + \mathbf{y}))$, noting that $\text{Overlap}(\mathcal{L}, r) = \text{vol}_n(S) / \text{vol}_n(rB_2^n)$. For the upper bound, by the union bound and Lemma 4.11, we have

$$\frac{\text{vol}_n(S)}{\text{vol}_n(rB_2^n)} \leq \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} \frac{\text{vol}_n(rB_2^n \cap rB_2^n + \mathbf{y})}{\text{vol}_n(rB_2^n)} \leq \sum_{\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}} 2e^{-\pi\|\frac{\mathbf{y}}{2s}\|^2} = 2\rho_{2s}(\mathcal{L} \setminus \{\mathbf{0}\})$$

as needed.

For the lower bound, let $\mathcal{V} = \mathcal{V}(\mathcal{L})$. We claim that $rB_2^n \setminus \mathcal{V} \subseteq S$. To see this take $\mathbf{x} \in rB_2^n \setminus \mathcal{V}$. Since $\mathbf{x} \notin \mathcal{V}$, there exists $\mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}$ such that $\|\mathbf{y} - \mathbf{x}\| < \|\mathbf{x}\| \leq r$. Therefore $\mathbf{x} \in rB_2^n \cap \mathbf{y} + rB_2^n$ as needed.

Let $s' = \frac{s}{1+\delta}$. By Lemma 3.4, we have that $\gamma_{s'}(\mathcal{V}) \leq \frac{1}{\rho_{s'}(\mathcal{L})}$. Let $X \in \mathbb{R}^n$ denote the gaussian with distribution γ . By the standard gaussian tailbound (Lemma 2.1), we have that

$$\gamma_{s'}(\mathbb{R}^n \setminus rB_2^n) = \Pr[\|s'X\|^2 \geq r^2] = \Pr[\|X\|^2 \geq (1+\delta)^2 \frac{n}{2\pi}] \leq \Pr[\|X\|^2 \geq (1+2\delta) \frac{n}{2\pi}] \leq e^{-\frac{2n}{3}\delta^2}.$$

Therefore we have that

$$\gamma_{s'}(rB_2^n \setminus \mathcal{V}) \geq 1 - \gamma_{s'}(\mathbb{R}^n \setminus \mathcal{V}) - \gamma_{s'}(\mathbb{R}^n \setminus rB_2^n) \geq \frac{\rho_{s'}(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho_{s'}(\mathcal{L})} - e^{-\frac{2n}{3}\delta^2}$$

By Lemma 4.12, we have that

$$\frac{\rho_{s'}(\mathcal{L} \setminus \{\mathbf{0}\})}{\rho_{s'}(\mathcal{L})} - e^{-\frac{2n}{3}\delta^2} \leq \frac{\gamma_{s'}(rB_2^n \setminus \mathcal{V})}{\gamma_{s'}(rB_2^n)} \leq \frac{\text{vol}_n(rB_2^n \setminus \mathcal{V})}{\text{vol}_n(rB_2^n)} \leq \frac{\text{vol}_n(S)}{\text{vol}_n(rB_2^n)}$$

as needed. \square

Proof of Lemma 4.4 (Ball Overlap Characterization). Let $r_\varepsilon = \frac{1}{2\eta_\varepsilon(\mathcal{L}^*)} \sqrt{\frac{n}{2\pi}}$ and $s = \frac{1}{\eta_\varepsilon(\mathcal{L}^*)}$. Note that by definition of s , we have that $\rho_s(\mathcal{L}) \leq \varepsilon$. For Part 1, by Lemma 4.13, using the fact that $s = 2r_\varepsilon \sqrt{\frac{2\pi}{n}}$, we get

$$\text{Overlap}(\mathcal{L}, r_\varepsilon) \leq 2\rho_s(\mathcal{L}) \leq 2\varepsilon.$$

Furthermore, for every $0 < r \leq r_\varepsilon$, since $\eta_\varepsilon(\mathcal{L})$ is a monotonically decreasing function in ε , $r = r_{\varepsilon'}$ for some $\varepsilon' \leq \varepsilon$. Thus,

$$\text{Overlap}(\mathcal{L}, r) \leq 2\varepsilon' \leq 2\varepsilon.$$

For Part 2, by 4.13 and the fact that $2^{-o(n)} \leq \varepsilon \leq 1/3$ and $\delta = \sqrt{\frac{3}{2n} \ln \frac{4}{\varepsilon}} < 1/4$, we get

$$\text{Overlap}(\mathcal{L}, 2 \cdot (1+\delta) \cdot r_\varepsilon) \geq \frac{\rho_s(\mathcal{L}^* \setminus \{\mathbf{0}\})}{\rho_s(\mathcal{L}^*)} - e^{-\frac{2n}{3}\delta^2} \geq \frac{\varepsilon}{1+\varepsilon} - \frac{\varepsilon}{4} \geq \frac{\varepsilon}{2}.$$

Again, for every $r \geq 2 \cdot (1+\delta)r_\varepsilon$, by the monotonicity of $\eta_\varepsilon(\mathcal{L})$ in ε , $r = r_{\varepsilon'}$ for some $\varepsilon' \geq \varepsilon$, and thus,

$$\text{Overlap}(\mathcal{L}, r) \geq \varepsilon'/2 \geq \varepsilon/2. \quad \square$$

4.3 Background and From ID Commitment Schemes to SZK Protocols

An ID commitment scheme Com for a promise problem Π is a commitment scheme that can depend on the instance x and such that only one of the hiding and binding properties are required to hold, depending on whether x is an YES or NO instance. Since only one of the hiding and binding properties needs to hold at a time, it is possible to achieve both the statistical hiding and statistical binding properties, and thus useful for constructing SZK protocols.

Typically, one requires the hiding property to hold for the YES instances and the binding property to hold for the NO instances, and such an ID commitment scheme readily gives a SZK protocol with soundness error $1/2$. On the other hand, an ID commitment scheme with reverse guarantees, i.e., binding for YES instances and hiding for NO instances, also readily gives a honest verifier SZK protocol, where the verifier commits to a random bit b and the prover's task is to guess the bit b correctly. Furthermore, since the verifier (who is the sender of the ID commitment scheme) is honest, the binding property only needs to hold with respect to the honest sender (referred to as "honest-sender binding property"). Since $\text{HVSZK} = \text{SZK}$ [GSV98], an ID commitment scheme that is honest-sender binding for YES instances and hiding for NO instance is also sufficient for showing that the promise problem is in SZK. Note that since only honest-sender binding property is required, we can without loss of generality assume that a commitment scheme is non-interactive (by letting the sender emulate the receiver and send the emulated view to the receiver). Thus, such a commitment scheme is simply an algorithm.

We observe that, the existing security amplification techniques for regular commitment schemes can be applied to the instance-dependent setting. As a consequence, any ID commitment scheme with "non-trivial" honest-sender binding and hiding properties is sufficient to obtain SZK protocols. More precisely, as formally defined in Definition 4.3, we consider ID commitment schemes Com with weak p -hiding and q -binding properties, where the hiding and binding properties can be broken with "advantage" at most p and q , respectively, and we say Com is "non-trivial" if $p + q \leq 1 - 1/\text{poly}(n)$. Known security amplification results for commitment schemes (for the case of statistical security) [DKS99] state that any non-trivial commitment scheme can be amplified to one with full-fledge security (i.e., both p and q are negligible). The same conclusion holds for ID commitment schemes, and thus to construct a SZK protocol for a language L , it suffice to construct a non-trivial honest-sender binding ID commitment scheme for L .

Theorem 4.14. *Let Π be a promise problem. Suppose there exists a non-trivial ID commitment scheme for Π , then $\Pi \in \text{SZK}$.*

Proof. (sketch) The theorem can be proved by applying known technique/results for regular commitment schemes to the instance-dependent setting. Briefly, security amplification of commitment schemes can be done using the following two operations [DKS99].

- **Repetition.** Given Com and $k \in \mathbb{N}$, define $\text{Com}'_x(b) = (\text{Com}_x(b; r_1), \dots, \text{Com}_x(b; r_k))$, i.e., concatenation of k commitments of Com using independent randomness. This amplifies the binding property but degrades the hiding property. Specifically, if Com is p -hiding and q -binding, then Com' is $(1 - (1 - p)^k)$ -hiding and q^k -binding.
- **Sharing.** Given Com and $k \in \mathbb{N}$, define $\text{Com}'_x(b) = (\text{Com}_x(b_1; r_1), \dots, \text{Com}_x(b_k; r_k))$, where b_1, \dots, b_k are chosen randomly subject to $b_1 \oplus \dots \oplus b_k = b$, and r_1, \dots, r_k are independent randomness. This amplifies the hiding property but degrades the binding property. Specifically, if Com is p -hiding and q -binding, then Com' is p^k -hiding and $1 - (1 - q)^k$ -binding.

It can be shown (as in [DKS99]) that as long as $p + q \leq 1 - 1/\text{poly}(n)$, one can amplify a p -hiding and q -binding commitment scheme Com to a secure Com' by alternately applying repetition and sharing operations with carefully chosen parameters k 's, and the resulting Com' calls Com in a black-box way $\text{poly}(n)$ times.

Once we have a secure non-interactive instance-dependent bit-commitment scheme for Π , we can readily construct a two-message honest-verifier SZK protocol for L as follows: On input $x \in \{0, 1\}^n$,

- V samples random $b \leftarrow \{0, 1\}$, computes and sends $\text{Com}_x(b)$ to P .
- P sends b' to V as his guess of b .
- V accepts iff $b' = b$.

It is not hard to see that the binding and hiding properties translate to the completeness and $1/2$ -soundness for the protocol, and a simulator can generate the view by emulating V and outputting $(\text{Com}_x(b), b)$. Since $\text{HVSZK} = \text{SZK}$, we have $\Pi \in \text{SZK}$. \square

Remark 4.15. Interestingly, as a by-product, an SZK-complete problem called ‘‘Image Intersection Density’’ (IID) (defined by [BOG03] and proved to be SZK-complete by [CCKV08]) can naturally be interpreted as a weak ID bit-commitment scheme as defined in Definition 4.3, which allows us to (immediately) obtain an optimal ‘‘polarization’’ result to the problem.

Specifically, the input to the IID problem is two distributions (X, Y) specified by circuits, where the YES instance satisfying $\Delta(X, Y) \leq a$ and the NO instance satisfying $\Pr[X \notin \text{supp}(Y)] \geq b$ and $\Pr[Y \notin \text{supp}(X)] \geq b$, where $a, b \in (0, 1)$ are parameters of the problem. By defining X and Y as commitment to 0 and 1 respectively, the condition to YES instance corresponds to statistical a -hiding and the condition to NO instance corresponds to statistical honest-sender $(1 - b)$ -binding.⁶ Interpreting the IID problem as a weak ID bit-commitment scheme makes it natural to apply the security amplification result of commitment schemes [DKS99], which gives an optimal polarization result of the problem, stating that the IID problem with parameters $a(n) - b(n) \geq 1/\text{poly}(n)$ is complete for SZK. This improves the previous known result in [CCKV08], which holds for constants $a > b$. In fact, the security amplification and polarization techniques exploit identical operations. The stronger result from the security amplification literature is obtained by applying the repetition and sharing operations more carefully.

5 Applications to Worst-case to Average-case Reductions

Our study of GapSPP has natural applications to the context of worst-case to average case reductions. In particular, we show that we can relate the hardness of average-case hard learning with error (LWE) problems and worst-case hard GapSPP problems with a tighter connection factor. Our result directly implies the worst-case to average-case result from GapSVP to LWE obtained by Regev [Reg05] and Peikert [Pei09]. First we review the LWE problem.

Definition 5.1 (Learning with Error Problem [Reg05]). Let $q = q(n) \in \mathbb{N}$, $\alpha = \alpha(n) \in (0, 1)$. Let Φ_α be the distribution on $[0, 1)$ obtained by drawing a sample from the Gaussian distribution with standard deviation α and reducing it modulo 1. Define A_{s, Φ_α} to be the distribution on $\mathbb{Z}_q^n \times [0, 1)$ obtained by choosing

⁶The binding and hiding properties hold for reverse instances, but one can instead consider the complement of the IID problem to obtain a consistent definition since SZK is close under complement.

a vector $\mathbf{a} \in \mathbb{Z}_q^n$ uniformly at random, choosing an error term $e \leftarrow \Phi_\alpha$, and outputting $(\mathbf{a}, \langle \mathbf{a}, \mathbf{s} \rangle / q + e)$ where the addition is performed in modulo 1.

The goal of the learning with errors problem $\text{LWE}_{q,\alpha}$ in n dimensions is, given access to any desired $\text{poly}(n)$ numbers of samples from $A_{\mathbf{s},\Phi_\alpha}$ for a random $\mathbf{s} \leftarrow \mathbb{Z}_q^n$, to find \mathbf{s} (with overwhelming probability).

Following [Reg05, Pei09], we use the bounded decoding BDD problem as an intermediate step in our reduction. Here we instead parameterize the α -BDD problem with α relative to the smoothing parameter (as opposed to the shortest vector used in literature); this is essential for us to obtain tighter reduction for GapSPP.

Definition 5.2 (Bounded Distance Decoding Problem (α -BDD $_\varepsilon$)). Given a lattice basis B and a vector \mathbf{t} such that $\text{dist}(\mathbf{t}, \mathcal{L}(B)) < \alpha/\eta_\varepsilon(\mathcal{L}(B)^*)$, find the lattice vector $\mathbf{v} \in \mathcal{L}(B)$ such that $\text{dist}(\mathbf{t}, \mathbf{v}) \leq \alpha/\eta_\varepsilon(\mathcal{L}(B)^*)$.

We recall the following Lemma from Regev [Reg05] and Peikert [Pei09] that reduce solving worst-case BDD problem to solving LWE through quantum and classic reductions, respectively.

Lemma 5.3 ([Reg05, Pei09]). *Let $q(n) \in \mathbb{N}$, $\alpha(n) \in (0, 1)$, $\varepsilon(n)$ be a negligible function such that $\alpha \cdot q > 2\sqrt{n}$. There exists a PPT quantum reduction from solving $\alpha/2$ -BDD $_\varepsilon$ in the worst case (with overwhelming probability) to solving $\text{LWE}_{q,\alpha}$ using $\text{poly}(n)$ samples.*

If in addition $q \geq 2^{n/2}$, then there exists a classical reduction from solving $\alpha/2$ -BDD $_\varepsilon$ in the worst case (with overwhelming probability) to solving $\text{LWE}_{q,\alpha}$ using $\text{poly}(n)$ samples.

We note that the reason $\varepsilon = \text{negl}(n)$ in the above Lemma is to guarantee that the LWE samples generated during the reduction are within negligible statistical distance from “true” LWE samples.

We now establish a new result that relates BDD and GapSPP. Our new observation is that the prover in the GGG protocol (Algorithm 1) can be implemented by a BDD oracle. Thus, if one has a BDD solver, one can solve the GapSPP problem. We note that we only need the BDD oracle to work for YES instances, and hence we require $\varepsilon_Y = \text{negl}(n)$ while leaving $\varepsilon_N = \frac{1}{\text{poly}(n)}$. More precisely, we have the following lemma.

Lemma 5.4. *Let $\alpha(n) \in (0, 1)$, $\varepsilon_Y(n) \in \text{negl}(n)$ and $\varepsilon_N \in 1/\text{poly}(n)$. There exists a PPT Turing reduction from solving \sqrt{n}/α -GapSPP $_{\varepsilon_Y, \varepsilon_N}$ to solving α -BDD $_{\varepsilon_Y}$.*

Proof. For convenience, we scale the \sqrt{n}/α -GapSPP problem so that YES instances have $\eta_{\varepsilon_Y}(\mathcal{L}) \leq \alpha/\sqrt{n}$, and NO instances have $\eta_{\varepsilon_N}(\mathcal{L}) > 1$. Let B be an input of the problem \sqrt{n}/α -GapSPP. We run the GGG protocol as Algorithm 1 on input B , where the prover’s strategy is implemented using the α -BDD $_{\varepsilon_Y}$ solver. Then we output the verifier’s decision.

Now we describe the analysis. For NO instances, by an identical analysis to Theorem 3.1, the above algorithm rejects with probability at least $\varepsilon_N/(1 + \varepsilon_N) > 1/\text{poly}(n)$. For YES instances, we observe that the optimal prover’s strategy can be emulated if $\|\mathbf{x}\|$ is less than the BDD decoding distance $\alpha/\eta_{\varepsilon_Y}(\mathcal{L}) \geq \sqrt{n}$. By the Gaussian tail bound as Lemma 2.1, we have $\Pr[\|\mathbf{x}\| \geq \sqrt{n}] < e^{-\Omega(n)}$. Recall that by Lemma 3.4, in GGG protocol the verifier rejects the optimal prover with probability $1 - \gamma_1(\mathcal{V}(\mathcal{L}^*)) \leq \rho_2(\mathcal{L}^* \setminus \mathbf{0}) \leq \varepsilon_Y$. Thus, by a union bound the algorithm rejects with probability at most $\varepsilon_Y + e^{-\Omega(n)} \leq \text{negl}(n)$.

Putting together the above lemmas, we obtain a tighter worst-case to average-case reduction from GapSPP to LWE.

Theorem 5.5. *Let $q(n) \in \mathbb{N}$, $\alpha(n) \in (0, 1)$, $\varepsilon_Y(n) \in \text{negl}(n)$ and $\varepsilon_N \in 1/\text{poly}(n)$ such that $\alpha \cdot q > 2\sqrt{n}$. There exists a PPT quantum reduction from solving $2\sqrt{n}/\alpha$ -GapSPP $_{\varepsilon_Y, \varepsilon_N}$ in the worst case (with overwhelming probability) to solving $\text{LWE}_{q, \alpha}$ using $\text{poly}(n)$ samples.*

If in addition $q \geq 2^{n/2}$, then there exists a classical reduction from solving $2\sqrt{n}/\alpha$ -GapSPP $_{\varepsilon_Y, \varepsilon_N}$ in the worst case (with overwhelming probability) to solving $\text{LWE}_{q, \alpha}$ using $\text{poly}(n)$ samples.

Remark 5.6. By using the following relation of shortest vectors and smoothing parameters by Micciancio and Regev [MR04]:

$$\frac{\sqrt{\log(1/\varepsilon)}}{\sqrt{\pi}\lambda_1(\mathcal{L}^*)} \leq \eta_\varepsilon(\mathcal{L}) \leq \frac{\sqrt{n}}{\lambda_1(\mathcal{L}^*)} \text{ for } \varepsilon \in [2^{-n}, 1],$$

the above theorem implies that there exists a corresponding PPT quantum/classical reduction from $(c \cdot \frac{n}{\alpha\sqrt{\log n}})$ -GapSVP to $\text{LWE}_{q, \alpha}$ for any constant $c > 0$.

6 Co-AM Protocol for GapSPP

In this section, we describe an co-AM protocol for GapSPP. Formally, we establish the following:

Theorem 6.1. *For any $\alpha \geq 1/\text{poly}(n)$ and $\varepsilon_Y, \varepsilon_N$ such that $\varepsilon_N \geq (1 + 1/\text{poly}(n)) \cdot \varepsilon_Y$, we have $(1 + \alpha)$ -GapSPP $_{\varepsilon_Y, \varepsilon_N} \in \text{coAM}$.*

By applying Corollary 2.5, we obtain the following upper bound on the complexity of γ -GapSPP $_\varepsilon$.

Corollary 6.2. *For every $\varepsilon : \mathbb{N} \rightarrow (0, 1)$ such that $\varepsilon(n) < 1 - 1/\text{poly}(n)$, we have $(1 + o(1))$ -GapSPP $_\varepsilon \in \text{coAM}$.*

Our main tool is the classic set size lower bound protocol by Goldwasser and Sipser [GS86]. We use this protocol to show that the smoothing parameter should be at least as large as some quantity. To show that $\eta(\mathcal{L})$ is large, equivalently we are showing that the discrete Gaussian weights are large for the points in \mathcal{L}^* inside the \sqrt{n} ball⁷. (The Gaussian weights outside the ball becomes exponentially small.)

The set size lower bound protocol gives a very accurate approximation of lattice points inside the \sqrt{n} ball, but its set size is not sufficient to approximate the Gaussian weights. The two points inside the ball could have lengths that differ a lot, and thus their Gaussian weights differ even more. Our new observation is that we can partition the \sqrt{n} ball into different shells (con-centered at $\mathbf{0}$), and then use the set size protocol to approximate the number of lattice points lying in each shell. Since every point in the same shell has roughly the same length and thus Gaussian weight, we can approximate the total Gaussian weights in a shell according to the size. Thus, summing up the Gaussian weight of each shell, we are able to approximate the Gaussian weights inside the \sqrt{n} ball. Thus, we are able to show that the Gaussian weights inside the ball are large, and thus η is large.

First we describe the set size lower bound protocol:

Definition 6.3 (Set size lower bound protocol [GS86]). Let V be a probabilistic polynomial time verifier, and P be a (computationally unbounded) prover. Let $S \subseteq \{0, 1\}^n$ be a set whose membership can be efficiently certified. The two parties hold common inputs 1^n and $K \in \mathbb{N}$.

We say $\langle P, V \rangle$ is a $(1 - \gamma)$ -approximation protocol of the set size $|S|$ if the following conditions hold:

⁷Actually the radius needs to depend on the parameter ε_Y . Here for simplicity we think ε_Y as a constant.

- (Completeness) If $|S| \geq K$, then V will always accept.
- (Soundness) If $|S| < (1 - \gamma) \cdot K$, then V will accept with probability at most $\text{negl}(n)$ for some negligible function $\text{negl}(\cdot)$.

Now we recall the classic construction of the set size lower bound protocol:

Theorem 6.4 ([GS86]). *For any set $S \in \{0, 1\}^n$ whose membership can be efficiently certified, and any $\gamma = 1/\text{poly}(n)$, there exists a public-coin, 2-round $(1 - \gamma)$ -approximation protocol of the set size $|S|$.*

Moreover, for any $k = \text{poly}(n)$, we can run the protocol k -times in parallel for k set-number pairs $\{(S_i, K_i)\}_{i \in [k]}$, and the resulting protocol has perfect completeness and negligible soundness error. Here soundness error means the probability that there exists some $i^ \in [k]$ such that $|S_{i^*}| \leq (1 - \gamma) \cdot K_{i^*}$ but V accepts.*

Proof of Theorem 6.1. To show the theorem, we first describe a coAM protocol $\langle P, V \rangle$ in the following. Note that the verifier in a coAM protocol must accept the NO instances and reject the YES instances of $(1 + \alpha)$ -GapSPP $_{\varepsilon_Y, \varepsilon_N}$. For convenience, the YES or NO instances here are with respect to the GapSPP problem, so the completeness means the verifier accepts any NO instance, and the soundness means he rejects any YES instance.

Let B be an n -dimensional basis of a lattice \mathcal{L} as input to the prover and verifier, satisfying either $\eta_{\varepsilon_N}(\mathcal{L}) \geq (1 + \alpha)$ (NO instance) or $\eta_{\varepsilon_Y}(\mathcal{L}) \leq 1$ (YES instance), where $\alpha \geq 1/\text{poly}(n)$, $\varepsilon_N \geq (1 + 1/\text{poly}(n)) \cdot \varepsilon_Y$. The prover and the verifier agree on the following parameters:

Parameters. Let $R = n \cdot (1 + \log(1/\varepsilon_Y))$, $1 - \beta = \frac{2\varepsilon_Y}{\varepsilon_Y + \varepsilon_N}$, and let $T = \lceil \frac{\log \sqrt{R}}{\log(1+\alpha)} \rceil$. We know for $\alpha \geq 1/\text{poly}(n)$ being noticeable, we have T bounded by some polynomial, i.e. $T \leq \text{poly}(n)$. Then we define spaces $S_0 \stackrel{\text{def}}{=} \{\mathbf{v} \in \mathcal{L}^* : 0 < \|\mathbf{v}\| \leq 1\}$, and $S_i \stackrel{\text{def}}{=} \{\mathbf{v} \in \mathcal{L}^* : (1 + \alpha)^{i-1} < \|\mathbf{v}\| \leq (1 + \alpha)^i\}$, for $i \in [T]$. Pictorially, these S_i 's form a partition of space inside the region of $\sqrt{R}B_2^n$. Each S_i is a shell that contains lattice points from length $(1 + \alpha)^{i-1}$ to $(1 + \alpha)^i$.

Then $\langle P, V \rangle$ does the following:

- P sends $K_0, K_1, K_2, \dots, K_T \in \mathbb{N}$ as claims of the sizes of $S_0, S_1, S_2, \dots, S_T$.
- Then for each pair (S_i, K_i) , P and V run the $(1 - \beta)$ -approximation protocol as its subroutine. These T approximation protocols are run in parallel. Note that $\frac{2\varepsilon_Y}{\varepsilon_Y + \varepsilon_N} \leq 1 - 1/\text{poly}(n)$ since $\varepsilon_N \geq (1 + 1/\text{poly}(n)) \cdot \varepsilon_Y$. Thus, $1 - \beta \leq 1 - 1/\text{poly}(n)$, which is within the range of parameters of the set size lower bound.
- In the end, V accepts if and only if all the approximation subprotocols are accepted, and $\sum_{0 \leq i \leq T} K_i \cdot e^{-\pi(1+\alpha)^{2i}} \geq (\varepsilon_Y + \varepsilon_N)/2$.

It is easy to see that the verifier can be implemented in probabilistic polynomial time. It remains to show the completeness and soundness. We show them by the following two claims:

Claim 6.5. *If B is a NO instance, V will always accept the honest prover's strategy.*

Proof. Let K_0, K_1, \dots, K_T be the values of the set sizes S_0, S_1, \dots, S_T , as the honest prover will always send the correct values. From the promise of NO instances, we know $\eta_{\varepsilon_N}(\mathcal{L}) \geq (1 + \alpha)$, which implies

$$q \stackrel{\text{def}}{=} \sum_{\mathbf{v} \in \mathcal{L}^* \setminus \{0\}} e^{-\pi(1+\alpha)^2 \|\mathbf{v}\|^2} \geq \varepsilon_N.$$

By rearranging the order of summation, we have

$$\begin{aligned} q &= \sum_{0 \leq i \leq T} \sum_{\mathbf{v} \in S_i} e^{-\pi(1+\alpha)^2 \|\mathbf{v}\|^2} + \sum_{\mathbf{v} \in \mathcal{L}^* \setminus (\sqrt{R} \cdot B_2^n)} e^{-\pi(1+\alpha)^2 \|\mathbf{v}\|^2} \\ &\leq \sum_{0 \leq i \leq T} \sum_{\mathbf{v} \in S_i} e^{-\pi(1+\alpha)^2 \|\mathbf{v}\|^2} + 2^{-n} \cdot \varepsilon_Y \\ &\leq \sum_{0 \leq i \leq T} K_i \cdot e^{-\pi(1+\alpha)^{2i}} + 2^{-n} \cdot \varepsilon_Y. \end{aligned}$$

The first equality comes from the rearrangement; the second line is a tail bound inequality by Lemma 2.1 by plugging suitable parameters; the last inequality is by the fact that $\mathbf{v} \in S_i$ implies $\|\mathbf{v}\| \geq (1 + \alpha)^{i-1}$ for $i \in [T]$.

Then we have

$$\sum_{0 \leq i \leq T} K_i \cdot e^{-\pi(1+\alpha)^{2i}} \geq \varepsilon_N - 2^{-n} \cdot \varepsilon_Y \geq (\varepsilon_Y + \varepsilon_N)/2,$$

for all sufficiently large n 's. This follows by the fact that $\varepsilon_N \geq (1 + 1/\text{poly}(n)) \cdot \varepsilon_Y$, and a straightforward examination. Thus, the verifier will always accept. \square

Claim 6.6. *If B is a YES instance, then no prover can convince the verifier with probability better than a negligible quantity.*

Proof. From the promise of YES instances, we know $\eta_{\varepsilon_Y}(\mathcal{L}) \leq 1$, which implies

$$q \stackrel{\text{def}}{=} \sum_{\mathbf{v} \in \mathcal{L}^* \setminus \{0\}} e^{-\pi \|\mathbf{v}\|^2} \leq \varepsilon_Y.$$

Similarly, we rearrange the order of summation and get

$$\begin{aligned} q &= \sum_{0 \leq i \leq T} \sum_{\mathbf{v} \in S_i} e^{-\pi \|\mathbf{v}\|^2} + \sum_{\mathbf{v} \in \mathcal{L}^* \setminus (\sqrt{R} \cdot B_2^n)} e^{-\pi \|\mathbf{v}\|^2} \\ &\geq \sum_{0 \leq i \leq T} \sum_{\mathbf{v} \in S_i} e^{-\pi \|\mathbf{v}\|^2} \\ &\geq \sum_{0 \leq i \leq T} |S_i| \cdot e^{-\pi(1+\alpha)^{2i}}. \end{aligned}$$

Suppose the prover sends some $K_0, K_1 \dots K_T$ such that $\sum_{0 \leq i \leq T} K_i \cdot e^{-\pi(1+\alpha)^{2i}} \geq (\varepsilon_Y + \varepsilon_N)/2$, it must be the case that $\exists i^* \in [T]$ such that $\frac{\varepsilon_Y}{(\varepsilon_Y + \varepsilon_N)/2} K_{i^*} = (1 - \beta) \cdot K_{i^*} \geq |S_{i^*}|$ from a simple counting argument. By the soundness of the $(1 - \beta)$ -approximation protocol, the verifier will catch this with probability $(1 - \text{negl}(n))$. Hence the verifier accepts a YES instance with only negligible probability. \square

Together with the two claims, the proof of the theorem is complete. \square

7 Deterministic Algorithm for Smoothing Parameter

In this section we show that $(1 + o(1))$ -GapSPP can be solved deterministically in time $2^{O(n)}$. In particular we are able to show the following theorem.

To show the theorem, we are going to establish the following lemma.

Theorem 7.1. *For any $\varepsilon_Y, \varepsilon_N : \mathbb{N} \rightarrow [0, 1]$ such that $\varepsilon_N(n) - \varepsilon_Y(n) \geq 1/2^{-2n}$, $1\text{-GapSPP}_{\varepsilon_Y, \varepsilon_N} \in \text{DTIME}(2^{O(n)})$.*

Together with Corollary 2.5, we are able to obtain the following corollary.

Corollary 7.2. *For any $\varepsilon : \mathbb{N} \rightarrow [0, 1]$ and $\varepsilon(n) \geq 2^{-n}$, the problem $(1+o(1))\text{-GapSPP}_\varepsilon \in \text{DTIME}(2^{O(n)})$.*

We will crucially use the following lattice point enumeration algorithm. The algorithm is a slight tweak of closest vector problem algorithm of Micciancio and Voulgaris [MV10], which was first used in [DPV11] to solve the shortest vector problem in general norms.

Proposition 7.3 ([MV10, DPV11], Algorithm Ball-Enum). *There is an algorithm Ball-Enum that given a radius $r > 0$, a basis B of an n -dimensional lattice \mathcal{L} , and $t \in \mathbb{R}^n$, lazily enumerates the set $\mathcal{L} \cap (rB_2^n + t)$ in deterministic time $2^{O(n)} \cdot (|\mathcal{L} \cap (t + rB_2^n)| + 1)$ using at most $2^{O(n)}$ space.*

Now we are ready to prove Theorem 7.1 using the above theorem.

Proof. Let B be an n -dimensional basis of a lattice \mathcal{L} satisfying either $\eta_{\varepsilon_Y} \leq 1$ or $\eta_{\varepsilon_N} \geq 1$, where $\varepsilon_Y, \varepsilon_N$ are parameters that the conditions in the theorem hold. Now we are going to describe an algorithm A on input B that distinguishes the two cases.

A runs the enumeration algorithm with the parameters $t = 0$, $r = \sqrt{n}$ to enumerate all points in $\mathcal{L}^* \cap \sqrt{n} \cdot B_2^n$. If A has already found $e^{\pi n} \cdot \varepsilon_N$ points from the enumeration algorithm, A terminates and rejects immediately. This is because

$$\sum_{\mathbf{v} \in \mathcal{L}^* \setminus \{0\}} e^{-\pi \|\mathbf{v}\|^2} \geq \sum_{\mathbf{v} \in (\mathcal{L}^* \setminus \{0\}) \cap \sqrt{n} \cdot B_2^n} e^{-\pi \|\mathbf{v}\|^2} \geq (e^{\pi n} \cdot \varepsilon_N) \cdot e^{-\pi n} = \varepsilon_N,$$

which already implies the case of no instances.

Otherwise A computes $u = \sum_{\mathbf{v} \in (\mathcal{L}^* \setminus \{0\}) \cap \sqrt{n} \cdot B_2^n} e^{-\pi \|\mathbf{v}\|^2}$. A accepts if $u \leq (\varepsilon_Y + \varepsilon_N)/2$, and otherwise rejects. The analysis of its completeness and soundness is very similar to that of Theorem 6.1, so we do not restate it here. It is not hard to see that this can be done in time $2^{O(n)} \cdot e^{\pi n} \cdot \varepsilon_N = 2^{O(n)}$.

Acknowledgments

We thank Oded Regev for helpful discussions. We also thank the anonymous reviewers for helpful suggestions.

References

- [Ajt96] M. Ajtai. Generating hard instances of lattice problems. *Quaderni di Matematica*, 13:1–32, 2004. Preliminary version in STOC 1996.

- [AR04] D. Aharonov and O. Regev. Lattice problems in $NP \cap coNP$. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004.
- [Bal97] K. M. Ball. An elementary introduction to modern convex geometry. In *S. Levy (Ed.), Flavors of Geometry, Number 31 in MSRI Publications*, pages 1–58, 1997.
- [Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.
- [Ban95] W. Banaszczyk. Inequalities for convex bodies and polar reciprocal lattices in R^n . *Discrete & Computational Geometry*, 13:217–231, 1995.
- [Ban96] W. Banaszczyk. Inequalities for convex bodies and polar reciprocal lattices in R^n II: Application of k -convexity. *Discrete and Computational Geometry*, 16:305–311, 1996. ISSN 0179-5376.
- [BHZ87] R. B. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? *Inf. Process. Lett.*, 25(2):127–132, 1987.
- [BOG03] M. Ben-Or and D. Gutfreund. Trading help for interaction in statistical zero-knowledge proofs. *J. Cryptology*, 16(2):95–116, 2003.
- [BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In *FOCS*, pages 97–106. 2011.
- [CCKV08] A. Chailloux, D. F. Ciocan, I. Kerenidis, and S. P. Vadhan. Interactive and noninteractive zero knowledge are equivalent in the help model. In *TCC*, pages 501–534. 2008.
- [CH11] H. Cohn and N. Heninger. Ideal forms of Coppersmith’s theorem and Guruswami-Sudan list decoding. In *ICS*, pages 298–308. 2011.
- [CHKP10] D. Cash, D. Hofheinz, E. Kiltz, and C. Peikert. Bonsai trees, or how to delegate a lattice basis. In *EUROCRYPT*, pages 523–552. 2010.
- [Cop97] D. Coppersmith. Small solutions to polynomial equations, and low exponent RSA vulnerabilities. *J. Cryptology*, 10(4):233–260, 1997.
- [DKS99] I. Damgård, J. Kilian, and L. Salvail. On the (im)possibility of basing oblivious transfer and bit commitment on weakened security assumptions. In *EUROCRYPT*, pages 56–73. 1999.
- [DPV11] D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, pages 580–589. 2011.
- [Gen09] C. Gentry. Fully homomorphic encryption using ideal lattices. In *STOC*, pages 169–178. 2009.
- [Gen10] C. Gentry. Toward basing fully homomorphic encryption on worst-case hardness. In *CRYPTO*, pages 116–137. 2010.
- [GG98] O. Goldreich and S. Goldwasser. On the limits of nonapproximability of lattice problems. *J. Comput. Syst. Sci.*, 60(3):540–563, 2000. Preliminary version in STOC 1998.
- [GMR04] V. Guruswami, D. Micciancio, and O. Regev. The complexity of the covering radius problem. *Computational Complexity*, 14(2):90–121, 2005. Preliminary version in CCC 2004.

- [GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008.
- [GS86] S. Goldwasser and M. Sipser. Private coins versus public coins in interactive proof systems. In *STOC*, pages 59–68. 1986.
- [GS88] J. Grollmann and A. L. Selman. Complexity measures for public-key cryptosystems. *SIAM J. Comput.*, 17(2):309–335, 1988.
- [GSV98] O. Goldreich, A. Sahai, and S. P. Vadhan. Honest-verifier statistical zero-knowledge equals general statistical zero-knowledge. In *STOC*, pages 399–408. 1998.
- [IOS97] T. Itoh, Y. Ohta, and H. Shizuya. A language-dependent cryptographic primitive. *J. Cryptology*, 10(1):37–50, 1997.
- [LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.
- [MR04] D. Micciancio and O. Regev. Worst-case to average-case reductions based on Gaussian measures. *SIAM J. Comput.*, 37(1):267–302, 2007. Preliminary version in FOCS 2004.
- [MV03] D. Micciancio and S. P. Vadhan. Statistical zero-knowledge proofs with efficient provers: Lattice problems and more. In *CRYPTO*, pages 282–298. 2003.
- [MV10] D. Micciancio and P. Voulgaris. A deterministic single exponential time algorithm for most lattice problems based on Voronoi cell computations. In *STOC*, pages 351–358. 2010.
- [Pei07] C. Peikert. Limits on the hardness of lattice problems in ℓ_p norms. *Computational Complexity*, 17(2):300–351, May 2008. Preliminary version in CCC 2007.
- [Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009.
- [Pei10] C. Peikert. An efficient and parallel Gaussian sampler for lattices. In *CRYPTO*, pages 80–97. 2010.
- [PW08] C. Peikert and B. Waters. Lossy trapdoor functions and their applications. In *STOC*, pages 187–196. 2008.
- [Reg03] O. Regev. New lattice-based cryptographic constructions. *J. ACM*, 51(6):899–942, 2004. Preliminary version in STOC 2003.
- [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.
- [SV03] A. Sahai and S. P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003.