New (and Old) Proof Systems for Lattice Problems

Navid Alamati* Chris Peikert† Noah Stephens-Davidowitz‡

December 19, 2017

Abstract

We continue the study of *statistical zero-knowledge* (SZK) proofs, both interactive and noninteractive, for computational problems on point lattices. We are particularly interested in the problem GapSPP of approximating the ε*-smoothing parameter* (for some $\varepsilon < 1/2$) of an n-dimensional lattice. The smoothing parameter is a key quantity in the study of lattices, and GapSPP has been emerging as a core problem in lattice-based cryptography, e.g., in worst-case to average-case reductions.

We show that GapSPP admits SZK proofs for *remarkably low* approximation factors, improving on prior work by up to roughly $\sqrt{n}$. Specifically:

- There is a *noninteractive* SZK proof for $O(\log(n)\sqrt{\log(1/\varepsilon)})$-approximate GapSPP. Moreover, for any negligible ε and a larger approximation factor $\widetilde{O}(\sqrt{n\log(1/\varepsilon)})$, there is such a proof with an *efficient prover*.
- There is an (interactive) SZK proof with an efficient prover for $O(\log n + \sqrt{\log(1/\varepsilon)/\log n})$-approximate coGapSPP. We show this by proving that $O(\log n)$-approximate GapSPP is in coNP.

In addition, we give an (interactive) SZK proof with an efficient prover for approximating the lattice *covering radius* to within an $O(\sqrt{n})$ factor, improving upon the prior best factor of $\omega(\sqrt{n\log n})$.

*Computer Science and Engineering, University of Michigan. Email: alamati@umich.edu.

†Computer Science and Engineering, University of Michigan. Email: cpeikert@umich.edu. This material is based upon work supported by the National Science Foundation under CAREER Award CCF-1054495 and CNS-1606362, the Alfred P. Sloan Foundation, and by a Google Research Award. The views expressed are those of the authors and do not necessarily reflect the official policy or position of the National Science Foundation, the Sloan Foundation, or Google.

‡Courant Institute of Mathematical Sciences, New York University. Email: noahsd@gmail.com. Supported by the National Science Foundation (NSF) under Grant No. CCF-1320188, and the Defense Advanced Research Projects Agency (DARPA) and Army Research Office (ARO) under Contract No. W911NF-15-C-0236. Part of this work was done while visiting the second author at the University of Michigan.

1 Introduction

Informally, a *proof system* [GMR85, Bab85a] is a protocol that allows a (possibly unbounded and malicious) prover to convince a skeptical verifier of the truth of some statement. A proof system is *zero knowledge* if the verifier "learns nothing more" from the interaction, other than the statement's veracity. The system is said to be *statistical* zero knowledge if the revealed information is negligible, even to an unbounded verifier; the class of problems having such proof systems is called SZK. Since their introduction, proof systems and zero-knowledge have found innumerable applications in cryptography and complexity theory. As a few examples, they have been used in constructions of secure multiparty computation [GMW87], digital signatures [BG89], actively secure public-key encryption [NY90], and "ZAPs" [DN00]. And if a problem has an SZK (or even coAM) proof, it is not NP-hard unless the polynomial-time hierarchy collapses [BHZ87], so interactive proofs have been used as evidence against NP-hardness; see, e.g., [GMR85, GMW91, GG98, HR14].

A proof system is *noninteractive* [BDMP88, GSV99] if it consists of just one message from the prover, assuming both it and the verifier have access to a truly random string. Noninteractive statistical zero-knowledge (NISZK) proof systems are especially powerful cryptographic primitives: they have minimal message complexity; they are concurrently and even "universally" composable [Can01]; and their security holds against unbounded malicious provers *and* verifiers, without any computational assumptions. However, we do not understand the class NISZK of problems that have noninteractive statistical zero-knowledge proof systems nearly as well as SZK. In particular, while NISZK is known to have complete problems, it is not known whether it is closed under complement or disjunction [GSV99], unlike SZK [SV97, Oka96].

Lattices and proofs. An n-dimensional lattice is a (full-rank) discrete additive subgroup of $\mathbb{R}^n$, and consists of all integer linear combinations of some linearly independent vectors $\mathbf{B} = \{\mathbf{b}_1, \ldots, \mathbf{b}_n\}$, called a *basis* of the lattice. Lattices have been extensively studied in computer science, and lend themselves to many natural computational problems. Perhaps the most well-known of these are the *Shortest Vector Problem* (SVP), which is to find a shortest nonzero vector in a given lattice, and the *Closest Vector Problem* (CVP), which is to find a lattice point that is closest to a given vector in $\mathbb{R}^n$. Algorithms for these problems and their approximation versions have many applications in computer science; see, e.g., [LLL82, Len83, Kan83, Odl90, JS98, NS01, DPV11]. In addition, many cryptographic primitives, ranging from public-key encryption and signatures to fully homomorphic encryption, are known to be secure assuming the (worst-case) hardness of certain lattice problems (see, e.g., [MR04, Reg05, GPV08, Pei09, BV11, BGV12]).

Due to the importance of lattices in cryptography, proof systems and zero-knowledge protocols for lattice problems have received a good deal of attention. Early on, Goldreich and Goldwasser [GG98] showed that for $\gamma = O(\sqrt{n/\log n})$, the γ-approximate Shortest and Closest Vector Problems, respectively denoted γ-GapSVP and γ-GapCVP, have SZK proof systems; this was later improved to coNP for $\gamma = O(\sqrt{n})$ factors [AR04].[1] Subsequently, Micciancio and Vadhan [MV03] gave different SZK proofs for the same problems, where the provers are *efficient* when given appropriate witnesses; this is obviously an important property if the proof systems are to be used by real entities as components of other protocols. Peikert and Vaikuntanathan [PV08] gave the first *noninteractive* statistical zero-knowledge proof systems for certain lattice problems, showing that, for example, $O(\sqrt{n})$-coGapSVP has an NISZK proof. The proof systems from [PV08] also have efficient provers, although for larger $\widetilde{O}(n)$ approximation factors.

[1] As described, the proofs from [GG98] are statistical zero knowledge against only *honest* verifiers, but any such proof can unconditionally be transformed to one that is statistical zero knowledge against *malicious* verifiers [GSV98]). We therefore ignore the distinction for the remainder of the paper.

Gaussians and the smoothing parameter. *Gaussian measures* have become an increasingly important tool in the study of lattices. For $s > 0$, the Gaussian measure of parameter (or width) s on $\mathbb{R}^n$ is defined as $\rho_s(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/s^2)$; for a lattice $\mathcal{L} \subset \mathbb{R}^n$, the Gaussian measure of the lattice is then

$$\rho_s(\mathcal{L}) := \sum_{\mathbf{v}\in\mathcal{L}} \rho_s(\mathbf{v}).$$

Gaussian measures on lattices have innumerable applications, including in worst-case to average-case reductions for lattice problems [MR04, Reg05], the construction of cryptographic primitives [GPV08], the design of algorithms for SVP and CVP [ADRS15, ADS15], and the study of the geometry of lattices [Ban93, Ban95, DR16, RS17].

In all of the above applications, a key quantity is the lattice *smoothing parameter* [MR04]. Informally, for a parameter $\varepsilon > 0$ and a lattice $\mathcal{L}$, the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ is the minimal Gaussian parameter that "smooths out" the discrete structure of $\mathcal{L}$, up to error ε. Formally, for $\varepsilon > 0$ we define

$$\eta_\varepsilon(\mathcal{L}) := \min\{s > 0 \ : \ \rho_{1/s}(\mathcal{L}^*) \le 1 + \varepsilon\} \ ,$$

where $\mathcal{L}^* := \{\mathbf{w} \in \mathbb{R}^n \ : \ \forall \mathbf{y} \in \mathcal{L}, \ \langle \mathbf{w}, \mathbf{y}\rangle \in \mathbb{Z}\}$ is the dual lattice of $\mathcal{L}$. All of the computational applications from the previous paragraph rely in some way on the "smoothness" of the Gaussian with parameter $s \ge \eta_\varepsilon(\mathcal{L})$ where $2^{-n} \ll \varepsilon < 1/2$.[2] For example, several of the proof systems from [PV08] start with deterministic reductions to an intermediate problem, which asks whether a lattice is "smooth" or well-separated.

The GapSPP problem. Given the prominence of the smoothing parameter in the theory of lattices, it is natural to ask about the complexity of computing it. Chung *et al.* [CDLP13] formally defined the problem γ-$\mathsf{GapSPP}_\varepsilon$ of approximating the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ to within a factor of $\gamma \ge 1$ and gave upper bounds on its complexity in the form of proof systems for *remarkably low* values of γ. For example, they showed that γ-$\mathsf{GapSPP}_\varepsilon \in \mathsf{SZK}$ for $\gamma = O(1 + \sqrt{\log(1/\varepsilon)/\log n})$. This in fact subsumes the prior result that $O(\sqrt{n/\log n})$-$\mathsf{GapSVP} \in \mathsf{SZK}$ of [GG98], via known relationships between the minimum distance and the smoothing parameter.

Chung *et al.* also showed a worst-case to average-case (quantum) reduction from $\widetilde{O}(\sqrt{n}/\alpha)$-GapSPP to a very important average-case problem in lattice-based cryptography, Regev's Learning With Errors (LWE), which asks us to decode from a random "q-ary" lattice under error proportional to α [Reg05]. Again, this subsumes the prior best reduction for GapSVP due to Regev. Most recently, Dadush and Regev [DR16] showed a similar worst-case to average-case reduction from GapSPP to the Short Integer Solution problem [Ajt96, MR04], another widely used average-case problem in lattice-based cryptography.

In hindsight, the proof systems and reductions of [GG98, Reg05, MR04] can most naturally be viewed as applying to GapSPP all along. This suggests that GapSPP may be a better problem than GapSVP on which to base the security of lattice-based cryptography. However, both [CDLP13] and [DR16] left open several questions and asked for a better understanding of the complexity of GapSPP. In particular, while interactive proof systems for this problem seem to be relatively well understood, nothing nontrivial was previously known about *noninteractive* proof systems (whether zero knowledge or not) for this problem.

[2] For $\varepsilon = 2^{-\Omega(n)}$ the smoothing parameter is determined (up to a constant factor) by the dual minimum distance, so it is much less interesting to consider as a separate quantity. The upper bound of $1/2$ could be replaced by any constant less than one. For $\varepsilon \ge 1$, $\eta_\varepsilon(\mathcal{L})$ is still formally defined, but its interpretation in terms of the "smoothness" of the corresponding Gaussian measure over $\mathcal{L}$ is much less clear.

1.1 Our Results

In this work we give new proof systems for lattice problems, and extend the reach of prior proof systems to new problems. Our new results, and how they compare to the previous state of the art, are as follows.

Our first main result is a NISZK proof system for γ-GapSPP$_\varepsilon$ with $\gamma = O(\log(n)\sqrt{\log(1/\varepsilon)})$. This improves, by a $\Theta(\sqrt{n}/\log n)$ factor, upon the previous best approximation factor of $\gamma = O(\sqrt{n\log(1/\varepsilon)})$, which follows from [PV08].

Theorem 1.1. *For any $\varepsilon \in (0, 1/2)$, $O(\log(n)\sqrt{\log(1/\varepsilon)})$-GapSPP$_\varepsilon \in$ NISZK.*

In fact, we demonstrate two different proof systems to establish this theorem (see Section 3). The first is identical to a proof system from [PV08], but with a very different analysis that relies on a recent geometric theorem of [RS17]. However, this proof system only works for negligible $\varepsilon < n^{-\omega(1)}$, so we also show an alternative that works for any $\varepsilon \in (0, 1/2)$ via reduction to the NISZK-complete *Entropy Approximation* problem [GSV99].

The prover in the proof system from [PV08] can be made efficient at the expense of a factor of $O(\sqrt{n\log n})$ in the approximation factor. From this we obtain the following.

Theorem 1.2. *For any negligible $0 < \varepsilon < n^{-\omega(1)}$, there is a* NISZK *proof system with an* efficient prover *for $O\big(\sqrt{n\log^3(n)\log(1/\varepsilon)}\big)$-GapSPP$_\varepsilon$.*

Next, we show that $O(\log n)$-GapSPP$_\varepsilon \in$ coNP for any $\varepsilon \in (0, 1)$. This improves, again by up to a $\Theta(\sqrt{n}/\log n)$ factor, the previous best known result of $O(1 + \sqrt{n/\log(1/\varepsilon)})$-GapSPP$_\varepsilon \in$ coNP, which follows from [Ban93].

Theorem 1.3. *For any $\varepsilon \in (0, 1/2)$, $O(\log n)$-GapSPP$_\varepsilon \in$ coNP.*

From this, together with the SZK protocol of [CDLP13] and the result of Nguyen and Vadhan [NV06] that any problem in SZK $\cap$ NP has an SZK proof system with an efficient prover, we obtain the following corollary. (The proof systems in [CDLP13] do not have efficient provers.)

Corollary 1.4. *For any $\varepsilon \in (0, 1/2)$, there is an* SZK *proof system with an* efficient prover *for $O(\log n + \sqrt{\log(1/\varepsilon)/\log n})$-coGapSPP$_\varepsilon$.*

Finally, we observe that $O(\sqrt{n})$-GapCRP $\in$ SZK, where GapCRP is the problem of approximating the *covering radius*, i.e., the maximum possible distance from a given lattice. For comparison, the previous best approximation factor was from [PV08], who showed that γ-GapCRP $\in$ NISZK $\subseteq$ SZK for any $\gamma = \omega(\sqrt{n\log n})$. We obtain this result via a straightforward reduction to $O(1)$-GapSPP$_\varepsilon$ for constant $\varepsilon < 1/2$, which, to recall, is in SZK [CDLP13]. Furthermore, since Guruswami, Micciancio, and Regev showed that $O(\sqrt{n})$-GapCRP $\in$ NP $\cap$ coNP [GMR04], it follows that the protocol can be made efficient.

Theorem 1.5. *We have $O(\sqrt{n})$-GapCRP $\in$ SZK. Furthermore, $O(\sqrt{n})$-GapCRP and $O(\sqrt{n})$-coGapCRP each have an* SZK *proof system with an efficient prover.*

1.2 Techniques

Sparse projections. Our main technical tool will be *sparse lattice projections*. In particular, we use the *determinant* of a lattice, defined as $\det(\mathcal{L}) := |\det(\mathbf{B})|$ for any basis $\mathbf{B}$ of $\mathcal{L}$, as our measure of sparsity.[3] It

[3]This is indeed a measure of sparsity because $1/\det(\mathcal{L})$ is the average number of lattice points inside a random shift of any unit-volume body, or equivalently, the limit as r goes to infinity of the number of lattice points per unit volume in a ball of radius r.

is an immediate consequence of the Poisson Summation Formula (Lemma 2.5) that $\det(\mathcal{L})^{1/n} \leq 2\eta_{1/2}(\mathcal{L})$. Notice that this inequality formalizes the intuitive notion that "a lattice cannot be smooth and sparse simultaneously."

Dadush and Regev made the simple observation that the same statement is true when we consider *projections* of the lattice [DR16]. I.e., for any projection π such that $\pi(\mathcal{L})$ is still a lattice, we have $\det(\pi(\mathcal{L}))^{1/\operatorname{rank}(\pi(\mathcal{L}))} \leq 2\eta_{1/2}(\mathcal{L})$, where $\operatorname{rank}(\pi(\mathcal{L}))$ is the dimension of the span of $\pi(\mathcal{L})$. (Indeed, this fact is immediate from the above together with the identity $(\pi(\mathcal{L}))^* = \mathcal{L}^* \cap \operatorname{span}(\pi(\mathcal{L}))$.) Therefore, if we define

$$\eta_{\det}(\mathcal{L}) := \max_{\pi} \det(\pi(\mathcal{L}))^{1/\operatorname{rank}(\pi(\mathcal{L}))} ,$$

where the maximum is taken over all projections π such that $\pi(\mathcal{L})$ is a lattice, then we have

$$\eta_{\det}(\mathcal{L}) \leq 2\eta_{1/2}(\mathcal{L}) . \tag{1.1}$$

Dadush and Regev conjectured that Equation (1.1) is tight up to a factor of $\operatorname{polylog}(n)$. I.e., up to polylog factors, a lattice is not smooth if and only if it has a sparse projection. Regev and Stephens-Davidowitz proved this conjecture [RS17], and the resulting theorem, presented below, will be our main technical tool.

Theorem 1.6 ([RS17]). *For any lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\eta_{1/2}(\mathcal{L}) \leq 10(\log n + 2)\eta_{\det}(\mathcal{L}) .$$

I.e., if $\eta_{1/2}(\mathcal{L}) \geq 10(\log n + 2)$, then there exists a lattice projection π such that $\det(\pi(\mathcal{L})) \geq 1$.

coNP proof system. Notice that Theorem 1.6 (together with Equation (1.1)) immediately implies that $O(\log n)$-$\mathsf{GapSPP}_{\varepsilon}$ is in coNP for $\varepsilon = 1/2$. Indeed, a projection π such that $\det(\pi(\mathcal{L}))^{1/\operatorname{rank}(\pi(\mathcal{L}))} \geq \eta_{1/2}(\mathcal{L})/O(\log n)$ can be used as a witness of "non-smoothness." Theorem 1.6 shows that such a witness always exists, and Equation (1.1) shows that no such witness exists with $\det(\pi(\mathcal{L}))^{1/\operatorname{rank}(\pi(\mathcal{L}))} > 2\eta_{1/2}(\mathcal{L})$. In order to extend this result to all $\varepsilon \in (0, 1)$, we use basic results about how $\eta_{\varepsilon}(\mathcal{L})$ varies with ε. (See Section 4.)

NISZK proof systems. We give two different NISZK proof systems for $O(\log(n)\sqrt{\log(1/\varepsilon)})$-$\mathsf{GapSPP}_{\varepsilon}$, both of which rely on Theorem 1.6.

Our first proof system (shown in Figure 1, Section 3.1) uses many vectors $\mathbf{t}_1, \ldots, \mathbf{t}_m$ sampled uniformly at random from a fundamental region of the lattice $\mathcal{L}$ as the common random string. The prover samples short vectors $\mathbf{e}_i$ (for $i = 1, \ldots, m$) from the *discrete Gaussian* distributions over the lattice cosets $\mathbf{e}_i + \mathcal{L}$. The verifier accepts if and only if the matrix $\mathbf{E} = \sum \mathbf{e}_i \mathbf{e}_i^T$ has small enough spectral norm. (I.e., the verifier accepts if the $\mathbf{e}_i$ are "short in all directions.") In fact, Peikert and Vaikuntanathan used the exact same proof system for the different lattice problem $O(\sqrt{n})$-$\mathsf{coGapSVP}$, and their proofs of correctness and zero knowledge also apply to our setting. However, the proof of soundness is quite different: we show that, if the lattice has a sparse projection π, then $\operatorname{dist}(\pi(\mathbf{t}_i), \pi(\mathcal{L}))$ will tend to be fairly large. It follows that $\sum \|\pi(\mathbf{e}_i)\|^2 = \operatorname{Tr}\left(\sum \pi(\mathbf{e}_i)\pi(\mathbf{e}_i)^T\right)$ will be fairly large with high probability, and therefore $\sum \mathbf{e}_i \mathbf{e}_i^T$ must have large spectral norm.

Our second proof system follows from a reduction to the Entropy Approximation problem, which asks to estimate the entropy of the output distribution of a circuit on random input. Goldreich, Sahai, and Vadhan [GSV99] showed that Entropy Approximation is NISZK-complete, so that a problem is in NISZK if and only if it can be (Karp-)reduced to approximating the entropy of a circuit. If $\eta_{\varepsilon}(\mathcal{L})$ is small, then we

know that a continuous Gaussian modulo the lattice will be very close to the uniform distribution, and so (a suitable discretization of) this distribution will have high entropy. On the other hand, if $\eta_\varepsilon(\mathcal{L})$ is large, then Theorem 1.6 says that most of the measure of a continuous Gaussian modulo the lattice lies in a low-volume subset of $\mathbb{R}^n/\mathcal{L}$, and so (a discretization of) this distribution must have low entropy.

This second proof system works for a wider range of ε. In particular, the first proof system is only statistical zero knowledge when ε is negligible in the input size, whereas the second proof system works for any $\varepsilon \in (0, 1/2)$.

1.3 Organization

The remainder of the paper is organized as follows.

- In Section 2 we recall the necessary background on lattices, proof systems, and probability.
- In Section 3 we give two different NISZK proof systems for $O(\log(n)\sqrt{\log(1/\varepsilon)})$-GapSPP$_\varepsilon$.
- In Section 4 we give a coNP proof system for $O(\log n)$-GapSPP$_\varepsilon$.
- In Section 5 we show that $O(\sqrt{n})$-GapCRP $\in$ SZK, via a simple reduction to $O(1)$-GapSPP$_{1/4}$.

2 Preliminaries

2.1 Notation

For any positive integer d, $[d]$ denotes the set $\{1, \ldots, d\}$. We use bold lower-case letters to denote vectors. We write matrices in capital letters. The ith component (column) of a vector $\mathbf{x}$ (matrix $\mathbf{X}$) is written as $\mathbf{x}_i$ ($\mathbf{X}_i$). The function $\log$ denotes the natural logarithm unless otherwise specified. For $\mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\| := \sqrt{x_1^2 + x_2^2 + \cdots + x_n^2}$ is the Euclidean norm. For a matrix $\mathbf{A} \in \mathbb{R}^{n \times m}$, $\|\mathbf{A}\| := \max_{\|\mathbf{x}\|=1} \|\mathbf{A}\mathbf{x}\|$ is the operator norm.

We write rB_2^n for the n-dimensional Euclidean ball of radius r. A set $S \subseteq \mathbb{R}^n$ is said to be *symmetric* if $-S = S$. The distance from a point $\mathbf{x} \in \mathbb{R}^n$ to a set $S \subseteq \mathbb{R}^n$ is defined to be $\mathrm{dist}(\mathbf{x}, S) = \inf_{\mathbf{s} \in S} \mathrm{dist}(\mathbf{x}, \mathbf{s})$. We write $S^\perp$ to denote the subspace of vectors orthogonal to S. For a set $S \subseteq \mathbb{R}^n$ and a point $\mathbf{x} \in \mathbb{R}^n$, $\pi_S(\mathbf{x})$ denotes the orthogonal projection of $\mathbf{x}$ onto $\mathrm{span}(S)$. For sets $A, B \subseteq \mathbb{R}^n$, we denote their Minkowski sum by $A + B = \{\mathbf{a} + \mathbf{b} : \mathbf{a} \in A, \mathbf{b} \in B\}$. We extend a function f to a countable set in the natural way by defining $f(A) := \sum_{a \in A} f(a)$.

Throughout the paper, we write C for an arbitrary universal constant $C > 0$, whose value might change from one use to the next.

2.2 Lattices

Here we provide some backgrounds on lattices. An n-dimensional *lattice* $\mathcal{L} \subset \mathbb{R}^n$ of rank d is the set of integer linear combinations of d linearly independent vectors $\mathbf{B} := (\mathbf{b}_1, \ldots, \mathbf{b}_d)$,

$$\mathcal{L} = \mathcal{L}(\mathbf{B}) = \left\{ \mathbf{B}\mathbf{z} = \sum_{i \in [d]} z_i \cdot \mathbf{b}_i : \mathbf{z} \in \mathbb{Z}^d \right\}.$$

We usually work with *full-rank* lattices, where $d = n$. A *sublattice* $\mathcal{L}' \subseteq \mathcal{L}$ is an additive subgroup of $\mathcal{L}$. The *dual lattice* of $\mathcal{L}$, denoted by $\mathcal{L}^*$, is defined as the set

$$\mathcal{L}^* = \left\{ \mathbf{y} \in \mathbb{R}^n : \forall \mathbf{v} \in \mathcal{L}, \langle \mathbf{v}, \mathbf{y} \rangle \in \mathbb{Z} \right\}$$

of all integer vectors having integer inner products with all vectors in $\mathcal{L}$. It is easy to check that $(\mathcal{L}^*)^* = \mathcal{L}$ and that, if $\mathbf{B}$ is a basis for $\mathcal{L}$, then $\mathbf{B}^* = \mathbf{B}(\mathbf{B}^T\mathbf{B})^{-1}$ is a basis for $\mathcal{L}^*$. The *fundamental parallelepiped* of a lattice $\mathcal{L}$ with respect to basis $\mathbf{B}$ is the set

$$\mathcal{P}(\mathbf{B}) = \Big\{ \sum_{i\in[d]} c_i \mathbf{b}_i : 0 \le c_i < 1 \Big\} .$$

It is easy to see that $\mathcal{P}(\mathbf{B})$ is a *fundamental domain* of $\mathcal{L}$. I.e., it tiles $\mathbb{R}^n$ with respect to $\mathcal{L}$. For any lattice $\mathcal{L}(\mathbf{B})$ and point $\mathbf{x} \in \mathbb{R}^n$, there exists a unique point $\mathbf{y} \in \mathcal{P}(\mathbf{B})$ such that $\mathbf{y} - \mathbf{x} \in \mathcal{L}(\mathbf{B})$. We denote this vector by $\mathbf{y} = \mathbf{x} \bmod \mathbf{B}$. Notice that $\mathbf{y}$ can be computed in polynomial time given $\mathbf{B}$ and $\mathbf{x}$. We sometimes write $\mathbf{x} \bmod \mathcal{L}$ when the specific fundamental domain is not important, and we write $\mathbb{R}^n/\mathcal{L}$ for an arbitrary fundamental domain.

The *determinant* of a lattice $\mathcal{L}$, is defined to be $\det(\mathcal{L}) = \sqrt{\det(\mathbf{B}^T\mathbf{B})}$. It is easy to verify that the determinant does not depend on the choice of basis and that $\det(\mathcal{L})$ is the volume of any fundamental domain of $\mathcal{L}$.

The *minimum distance* of a lattice $\mathcal{L}$, is the length of the shortest non-zero lattice vector,

$$\lambda_1(\mathcal{L}) := \min_{\mathbf{y}\in\mathcal{L}\setminus\{\mathbf{0}\}} \|\mathbf{y}\| .$$

Similarly, we define

$$\lambda_n(\mathcal{L}) := \min \max_i \|\mathbf{y}_i\| ,$$

where the minimum is taken over linearly independent lattice vectors $\mathbf{y}_1, \ldots, \mathbf{y}_n \in \mathcal{L}$. The *covering radius* of a lattice $\mathcal{L}$ is

$$\mu(\mathcal{L}) := \max_{\mathbf{t}\in\mathbb{R}^n} \operatorname{dist}(\mathbf{t}, \mathcal{L}) .$$

The *Voronoi cell* of a lattice $\mathcal{L}$ is the set

$$\mathcal{V}(\mathcal{L}) := \{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{t}\| \le \|\mathbf{y} - \mathbf{t}\|, \forall \mathbf{y} \in \mathcal{L} \setminus \{\mathbf{0}\}\}$$

of vectors in $\mathbb{R}^n$ that are closer to $\mathbf{0}$ than any other point of $\mathcal{L}$. It is easy to check that $\mathcal{V}(\mathcal{L})$ is a symmetric polytope and that it tiles $\mathbb{R}^n$ with respect to $\mathcal{L}$. The following claim is an immediate consequence of the fact that an n-dimensional unit ball has volume at most $(2\pi e/n)^{n/2}$.

Claim 2.1. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\mu(\mathcal{L}) \ge \sqrt{n/(2\pi e)} \cdot \det(\mathcal{L})^{1/n} .$$

Lemma 2.2. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and $r \ge 0$,*

$$|\mathcal{L} \cap rB_2^n| \le (5/\sqrt{n})^n \cdot \frac{(r + \mu(\mathcal{L}))^n}{\det(\mathcal{L})} .$$

Proof. For each vector $\mathbf{y} \in \mathcal{L} \cap rB_2^n$, notice that $\mathcal{V}(\mathcal{L}) + \mathbf{y} \subseteq (r + \mu(\mathcal{L}))B_2^n$. And, for distinct vectors $\mathbf{y}, \mathbf{y}' \in \mathcal{L}$, $\mathcal{V}(\mathcal{L}) + \mathbf{y}$ and $\mathcal{V}(\mathcal{L}) + \mathbf{y}'$ are disjoint (up to a set of measure zero). Therefore,

$$\operatorname{vol}((r + \mu(\mathcal{L}))B_2^n) \ge \operatorname{vol}\Big(\bigcup_{\mathbf{y}\in\mathcal{L}\cap rB_2^n} \mathcal{V}(\mathcal{L}) + \mathbf{y} \Big) = |\mathcal{L} \cap rB_2^n| \operatorname{vol}(\mathcal{V}(\mathcal{L})) = |\mathcal{L} \cap rB_2^n| \cdot \det(\mathcal{L}) .$$

The result follows by recalling that for any $r' > 0$, $\operatorname{vol}(r'B_2^n) \le (5r'/\sqrt{n})^n$. □

Lemma 2.3 ([GMR04]). *For any lattice $\mathcal{L} \subset \mathbb{R}^n$,*

$$\mathop{\mathbb{E}}_{\mathbf{t} \sim \mathbb{R}^n/\mathcal{L}}[\mathrm{dist}(\mathbf{t}, \mathcal{L})^2] \geq \mu(\mathcal{L})^2/4 \ ,$$

where $\mathbf{t} \in \mathbb{R}^n/\mathcal{L}$ is sampled uniformly at random.

Proof. Let $\mathbf{v} \in \mathbb{R}^n$ such that $\mathrm{dist}(\mathbf{v}, \mathcal{L}) = \mu(\mathcal{L})$. Notice that $\mathbf{v} - \mathbf{t} \bmod \mathcal{L}$ is uniformly distributed. And, by the triangle inequality, $\mathrm{dist}(\mathbf{v} - \mathbf{t}, \mathcal{L}) + \mathrm{dist}(\mathbf{t}, \mathcal{L}) \geq \mathrm{dist}(\mathbf{v}, \mathcal{L}) = \mu(\mathcal{L})$. So,

$$\mathop{\mathbb{E}}_{\mathbf{t} \sim \mathbb{R}^n/\mathcal{L}}[\mathrm{dist}(\mathbf{t}, \mathcal{L})] = \frac{1}{2} \cdot \mathop{\mathbb{E}}_{\mathbf{t} \sim \mathbb{R}^n/\mathcal{L}}[\mathrm{dist}(\mathbf{v} - \mathbf{t}, \mathcal{L}) + \mathrm{dist}(\mathbf{t}, \mathcal{L})] \geq \mu(\mathcal{L})/2 \ .$$

The result then follows by Markov's inequality. □

A *lattice projection* for a lattice $\mathcal{L} \subset \mathbb{R}^n$ is an orthogonal projection $\pi : \mathbb{R}^n \to \mathbb{R}^n$ defined by $\pi(\mathbf{x}) := \pi_{S^\perp}(\mathbf{x})$ for lattice vectors $S \subset \mathcal{L}$.

Claim 2.4. *For any $\mathcal{L} \subset \mathbb{R}^n$ and any lattice projection π, $\pi(\mathcal{L})$ is a lattice. Furthermore, if $\mathbf{t} \in \mathbb{R}^n/\mathcal{L}$ is sampled uniformly at random, then $\pi(\mathbf{t})$ is uniform mod $\pi(\mathcal{L})$.*

Proof. The first statement follows from the well known fact that, if $W = \mathrm{span}\, S$ for some set of lattice vectors $S \subset \mathcal{L}$, then there exists a basis $\mathbf{B} := (\mathbf{b}_1, \ldots, \mathbf{b}_n)$ of $\mathcal{L}$ such that $\mathrm{span}(\mathbf{b}_1, \ldots, \mathbf{b}_k) = W$, where $k := \dim W$. (See, e.g., [MG02].) From this, it follows immediately that $\pi(\mathbf{b}_{k+1}), \ldots \pi(\mathbf{b}_n)$ are linearly independent and $\pi(\mathcal{L})$ is the lattice spanned by these vectors, where $\pi := \pi_{S^\perp}$.

The second statement follows from the following similarly well known fact. Let $\widetilde{\mathbf{b}}_i := \pi_{\{\mathbf{b}_1, \ldots, \mathbf{b}_{i-1}\}^\perp}(\mathbf{b}_i)$ be the Gram-Schmidt vectors of the basis $\mathbf{B}$ described above. Then, the hyperrectangle

$$\widetilde{R} := \Big\{ \sum_i a_i \widetilde{\mathbf{b}}_i \ : \ -1/2 < a_i \leq 1/2 \Big\}$$

is a fundamental domain of the lattice. (See, e.g., [Bab85b]) I.e., for each $\mathbf{t} \in \mathbb{R}^n/\mathcal{L}$, there is a unique representative $\widetilde{\mathbf{t}} \in \widetilde{R}$ with $\widetilde{\mathbf{t}} \equiv \mathbf{t} \bmod \mathcal{L}$. The result then follows by noting that, if $\widetilde{\mathbf{t}} \in \widetilde{R}$ is chosen uniformly at random, then clearly $\pi(\widetilde{\mathbf{t}}) \in \pi(\widetilde{R})$ is uniform in $\pi(\widetilde{R})$, which is a fundamental region of $\pi(\mathcal{L})$. □

2.3 Gaussian Measure

Here we review some useful background on Gaussians over lattices. For a positive parameter $s > 0$ and vector $\mathbf{x} \in \mathbb{R}^n$, we define the Gaussian mass of $\mathbf{x}$ as $\rho_s(\mathbf{x}) = e^{-\pi \|\mathbf{x}\|^2/s^2}$. For a measurable set $A \subseteq \mathbb{R}^n$, we define $\gamma_s(A) = s^{-n} \int_A \rho_s(\mathbf{x}) \, \mathrm{d}\mathbf{x}$. It is easy to see that $\gamma_s(\mathbb{R}^n) = 1$ and hence γ_s is a probability measure. We define the *discrete Gaussian distribution* over a countable set A as

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

In all cases, the parameter s is taken to be one when omitted. The following lemma is the Poisson Summation Formula for the Gaussian mass of a lattice.

Lemma 2.5. *For any (full-rank) lattice $\mathcal{L}$ and $s > 0$,*

$$\rho_s(\mathcal{L}) = \frac{1}{\det(\mathcal{L})} \cdot \rho_{1/s}(\mathcal{L}^*) \ .$$

We will also need Banaszczyk's celebrated lemma [Ban93, Lemma 1.5].

Lemma 2.6 ([Ban93]). *For any lattice $\mathcal{L} \subset \mathbb{R}^n$, shift vector $\mathbf{t} \in \mathbb{R}^n$, and $r \geq 1/\sqrt{2\pi}$,*

$$\rho((\mathcal{L}+\mathbf{t}) \setminus \sqrt{n}B_2^n) \leq \left(\sqrt{2\pi e r^2}e^{-\pi r^2}\right)^n \cdot \rho(\mathcal{L}) \ .$$

Micciancio and Regev introduced a lattice parameter called the smoothing parameter. For an n-dimensional lattice $\mathcal{L}$ and $\varepsilon > 0$, the *smoothing parameter* $\eta_\varepsilon(\mathcal{L})$ is defined as the smallest s such that $\rho_{1/s}(\mathcal{L}^*) \leq 1+\varepsilon$. The motivation for defining smoothing parameter comes from the following two facts [MR04].

Claim 2.7. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$, shift vector $\mathbf{t} \in \mathbb{R}^n$, $\varepsilon \in (0,1)$, and parameter $s \geq \eta_\varepsilon(\mathcal{L})$,*

$$\frac{1-\varepsilon}{1+\varepsilon} \cdot \rho_s(\mathcal{L}) \leq \rho_s(\mathcal{L}-\mathbf{t}) \leq \rho_s(\mathcal{L}) \ .$$

Lemma 2.8. *For any lattice $\mathcal{L}, \mathbf{c} \in \mathbb{R}^n$ and $s \geq \eta_\varepsilon(\mathcal{L})$,*

$$\Delta((D_s \bmod \mathbf{B}), U(\mathbb{R}^n/\mathcal{L})) \leq \varepsilon/2 \ ,$$

where D_s is the continuous Gaussian distribution with parameter s and $U(\mathbb{R}^n/\mathcal{L})$ denotes the uniform distribution over $\mathbb{R}^n/\mathcal{L}$.

We use the following epsilon-decreasing tool which has been introduced in [CDLP13].

Lemma 2.9 ([CDLP13], Lemma 2.4). *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and any $0 < \varepsilon' \leq \varepsilon < 1$,*

$$\eta_{\varepsilon'}(\mathcal{L}) \leq \sqrt{\log(1/\varepsilon')/\log(1/\varepsilon)} \cdot \eta_\varepsilon(\mathcal{L}) \ .$$

Proof. We may assume without loss of generality that $\eta_\varepsilon(\mathcal{L}) = 1$. Notice that this implies that $\lambda_1(\mathcal{L}^*) \geq \sqrt{\log(1/\varepsilon)/\pi}$. Then, for any $s \geq 1$,

$$\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) = \sum \exp(-\pi(s^2-1)\|\mathbf{w}\|^2) \cdot \rho(\mathbf{w}) \leq \exp(-\pi(s^2-1)\lambda_1(\mathcal{L})^2)\rho(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon^{s^2} \ .$$

Setting $s := \sqrt{\log(1/\varepsilon')/\log(1/\varepsilon)}$ gives the result. □

Lemma 2.10. *For any lattice $\mathcal{L} \subset \mathbb{Q}^n$ with basis $\mathbf{B}$ whose bit length is β and any $\varepsilon \in (0,1/2)$, we have $\eta_\varepsilon(\mathcal{L}(\mathbf{B})) \leq 2^{\mathrm{poly}(\beta)}\sqrt{\log(1/\varepsilon)}$, and $\lambda_n(\mathcal{L}) \leq 2\mu(\mathcal{L}) \leq 2^{\mathrm{poly}(\beta)}$.*

2.4 Sampling from the Discrete Gaussian

For any $\mathbf{B} = (\mathbf{b}_1, \ldots, \mathbf{b}_n) \in \mathbb{R}^{n \times n}$, let

$$\|\widetilde{\mathbf{B}}\| := \max_i \|\pi_{\{\mathbf{b}_1,\ldots,\mathbf{b}_{i-1}\}^\perp}(\mathbf{b}_i)\| \ ,$$

i.e., $\|\widetilde{\mathbf{B}}\|$ is the length of the longest Gram-Schmidt vector of $\mathbf{B}$.

We recall the following result from a sequence of works due to Klein [Kle00]; Gentry, Peikert, and Vaikuntanathan [GPV08]; and Brakerski et al. [BLP+13].

Theorem 2.11. *There is an efficient algorithm that takes as input a basis $\mathbf{B} \in \mathbb{Q}^{n \times n}$ and any parameter $s \geq \|\widetilde{\mathbf{B}}\|\sqrt{\log n}$ and outputs a sample from $D_{\mathcal{L},s}$, where $\mathcal{L} \subset \mathbb{R}^n$ is the lattice generated by $\mathbf{B}$.*

Corollary 2.12. *There is an efficient algorithm that takes as input a (basis for a) lattice $\mathcal{L} \subset \mathbb{Q}^n$ and parameter $s \geq 2^n \eta_\varepsilon(\mathcal{L})$ and outputs a sample from $D_{\mathcal{L},s}$.*

Proof. Combine the above with the celebrated LLL algorithm [LLL82], which in particular allows us to find a basis for $\mathcal{L}$ with $\|\widetilde{\mathbf{B}}\| \leq 2^{n/2}\eta_\varepsilon(\mathcal{L})$. □

We also need the following result, which is implicit in [Ban93]. See, e.g., [DR16] for a proof.

Lemma 2.13. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and $\varepsilon \in (0, 1/2)$,*

$$\lambda_n(\mathcal{L}) \leq 2\mu(\mathcal{L}) \leq \sqrt{n} \cdot \eta_\varepsilon(\mathcal{L}) \ .$$

In particular, there exists a basis $\mathbf{B}$ of $\mathcal{L}$ with $\|\widetilde{\mathbf{B}}\| \leq \lambda_n(\mathcal{L}) \leq \sqrt{n} \cdot \eta_{1/2}(\mathcal{L})$.

Corollary 2.14. *For any lattice $\mathcal{L} \subset \mathbb{Q}^n$ with basis $\mathbf{B}$, there exists preprocessing P whose size is polynomial in the bit length of $\mathbf{B}$ and an efficient algorithm that, on input P and $s \geq \sqrt{n \log n} \cdot \eta_{1/2}(\mathcal{L})$ outputs a sample from $D_{\mathcal{L},s}$.*

Proof. By Lemma 2.13, there exists a basis $\mathbf{B}'$ with $\|\widetilde{\mathbf{B}'}\| \leq \sqrt{n} \cdot \eta_{1/2}(\mathcal{L})$. By Lemma 2.10, the bit length of $\mathbf{B}'$ is polynomial in the bit length of $\mathbf{B}$. We use this as our preprocessing P. The result then follows by Theorem 2.11. □

2.5 Computational Problems

Here we define two promise problems that will be considered in this paper.

Definition 2.15 (Covering Radius Problem). For any approximation factor $\gamma = \gamma(n) \geq 1$, an instance of γ-GapCRP is a (basis for a) lattice $\mathcal{L} \subset \mathbb{Q}^n$. It is a YES instance if $\mu(\mathcal{L}) \leq 1$ and a NO instance if $\mu(\mathcal{L}) > \gamma$.

Definition 2.16 (Smoothing Parameter Problem). For any approximation factor $\gamma = \gamma(n) \geq 1$ and $\varepsilon = \varepsilon(n) > 0$, an instance of γ-GapSPP$_\varepsilon$ is a (basis for a) lattice $\mathcal{L} \subset \mathbb{Q}^n$. It is a YES instance if $\eta_\varepsilon(\mathcal{L}) \leq 1$ and a NO instance if $\eta_\varepsilon(\mathcal{L}) > \gamma$.

We will need the following result from [CDLP13].

Theorem 2.17. *For any $\varepsilon \in (0, 1/2)$, γ-GapSPP$_\varepsilon$ is in SZK for $\gamma = O(1 + \sqrt{\log(1/\varepsilon)/\log(n)})$.*[4]

2.6 Noninteractive Proof Systems

Definition 2.18 (Noninteractive Proof System). A pair (P, V) is a *noninteractive proof system* for a promise problem $\Pi = (\Pi^{\text{YES}}, \Pi^{\text{NO}})$ if P is a (possibly unbounded) algorithm and V is a polynomial-time algorithm such that

- *Completeness:* for every $x \in \Pi_n^{\text{YES}}$, $\Pr[V(x, r, P(x, r)) \text{ accepts}] \geq 1 - \varepsilon$; and

[4] In [CDLP13], this result is proven only for $\varepsilon < 1/3$. However, it is immediate from, e.g., Lemma 2.9 that the result can be extended to any $\varepsilon < 1/2$.

- *Soundness:* for every $x \in \Pi_n^{\text{NO}}$, $\Pr[\exists\ \pi\ :\ V(x, r, \pi) \text{ accepts}] \leq \varepsilon$,

where n is the input length, $\varepsilon = \varepsilon(n) \leq \text{negl}(n)$, and the probabilities are taken over r, which is sampled uniformly at random from $\{0,1\}^{\text{poly}(n)}$.

A noninteractive proof system (P, V) for a promise problem $\Pi = (\Pi^{\text{YES}}, \Pi^{\text{NO}})$ is *statistical zero knowledge* if there exists a probabilistic polynomial-time algorithm S (called a *simulator*) such that for all $x \in \Pi^{\text{YES}}$, the statistical distance between $S(x)$ and $(r, P(x, r))$ is negligible in n. The class of promise problems having noninteractive statistical zero-knowledge proof systems is denoted NISZK.

2.7 Probability

The *entropy* of a random variable X over a countable set S is given by

$$H(X) := \sum_{a \in S} \Pr[X = a] \cdot \log_2(1/\Pr[X = a]) \ .$$

We will also need the Chernoff-Hoeffding bound [Hoe63].

Lemma 2.19 (Chernoff-Hoeffding bound). *Let $X_1, \ldots, X_m \in [0,1]$ be independent and identically distributed random variables with $\overline{X} := \mathbb{E}[X_i]$. Then, for any $s > 0$,*

$$\Pr\left[m\overline{X} - \sum X_i \geq s\right] \leq \exp(-s^2/(2m)) \ .$$

Finally, we will need a minor variant of the above inequality.

Lemma 2.20. *Let $X_1, \ldots, X_m \in \mathbb{R}$ be independent (but not necessarily identically distributed) random variables. Suppose that there exists an $\alpha \geq 0$ and $s > 0$ such that for any $r > 0$,*

$$\Pr[|X_i| \geq r] \leq \alpha \exp(-r^2/s^2) \ .$$

Then, for any $r > 0$,

$$\Pr\left[\sum X_i^2 \geq r\right] \leq (1+\alpha)^m \exp(-r/(2s^2)) \ .$$

Proof. For any index i, we have

$$\begin{aligned}
\mathbb{E}[\exp(X_i^2/(2s^2))] &= 1 + \frac{1}{s^2} \cdot \int_0^\infty r \exp(r^2/(2s^2)) \Pr[|X_i| \geq r] \, \mathrm{d}r \\
&\leq 1 + \frac{\alpha}{s^2} \cdot \int_0^\infty r \exp(-r^2/(2s^2)) \, \mathrm{d}r \\
&= 1 + \alpha \ .
\end{aligned}$$

Since the X_i are independent, it follows that

$$\mathbb{E}\left[\exp\left(\sum X_i^2/(2s^2)\right)\right] = \mathbb{E}\left[\prod_i \exp(X_i^2/(2s^2))\right] \leq (1+\alpha)^m \ .$$

The result then follows by Markov's inequality. □

3 Two NISZK Proofs for GapSPP

Recall the definition

$$\eta_{\det}(\mathcal{L}) := \max_{\pi} \det(\pi(\mathcal{L}))^{1/\operatorname{rank}(\pi(\mathcal{L}))} .$$

We will also need the following definition from [DR16],

$$C_\eta(n) := \sup_{\mathcal{L}} \frac{\eta_{1/2}(\mathcal{L})}{\eta_{\det}(\mathcal{L})} ,$$

where the supremum is taken over all lattices $\mathcal{L} \subset \mathbb{R}^n$. In this notation, Theorem 1.6 is equivalent to the inequality

$$C_\eta(n) \le 10(\log n + 2) .$$

We note that the true value of $C_\eta(n)$ is still not known. (In particular, the best lower bound is $C_\eta(n) \ge \sqrt{\log(n)/\pi} + o(1)$, which follows from the fact that $\eta_{1/2}(\mathbb{Z}^n) = \sqrt{\log(n)/\pi} + o(1)$.) We therefore state our results in terms of $C_\eta(n)$.

3.1 An Explicit Proof System

We first consider the NISZK proof system for $\sqrt{n}$-coGapSVP due to [PV08], shown in Figure 1. We show that this is actually also a NISZK proof system for $O(\sqrt{\log(1/\varepsilon)} \cdot \log n)$-GapSPP$_\varepsilon$ for negligible ε. (In Section 3.2, we show a different proof system that works for all $\varepsilon \in (0, 1/2)$, also with an approximation factor of $O(\log(n)\sqrt{\log(1/\varepsilon)})$.)

NISZK proof system for GapSPP.

Common Input: A basis $\mathbf{B}$ for a lattice $\mathcal{L} \subset \mathbb{Q}^n$.

Random Input : m vectors $\mathbf{t}_1, \ldots, \mathbf{t}_m \in \mathcal{P}(\mathbf{B})$, sampled uniformly at random.

Prover P: Sample m vectors $\mathbf{e}_1, \ldots, \mathbf{e}_m \in \mathbb{R}^n$ independently from $D_{\mathcal{L}+\mathbf{t}_i}$, and output them as the proof.

Verifier V: Accept if and only if $\mathbf{e}_i \equiv \mathbf{t}_i \bmod \mathcal{L}$ for all i *and* $\left\| \sum \mathbf{e}_i \mathbf{e}_i^T \right\| \le 3m$.

Figure 1: The non-interactive zero-knowledge proof system for GapSPP, where $m := 100n$.

Theorem 3.1. *For any $\varepsilon \le \operatorname{negl}(n)$, γ-GapSPP$_\varepsilon$ is in NISZK for*

$$\gamma := O(C_\eta(n)\sqrt{\log(1/\varepsilon)}) \le O(\log(n)\sqrt{\log(1/\varepsilon)})$$

via the proof system shown in Figure 1.

We will prove in turn that the proof system is statistical zero knowledge, complete, and sound. In fact, the proofs of statistical zero knowledge and completeness are nearly identical to the corresponding proofs in [PV08].

To prove the zero-knowledge property of the proof system, we consider the simulator that behaves as follows. Let $\mathbf{e}_1, \ldots, \mathbf{e}_m \in \mathbb{R}^m$ be sampled independently from the continuous Gaussian centered at $\mathbf{0}$. Let $\mathbf{t}_1, \ldots, \mathbf{t}_m \in \mathcal{P}(\mathbf{B})$ such that $\mathbf{e}_i \equiv \mathbf{t}_i \bmod \mathcal{L}$. The simulator then outputs $\mathbf{t}_1, \ldots, \mathbf{t}_m$ as the random input and $\mathbf{e}_1, \ldots, \mathbf{e}_m$ as the proof.

Lemma 3.2 (Statistical zero knowledge). *For any $\varepsilon \in (0, 1)$ and lattice $\mathcal{L} \subset \mathbb{Q}^n$ with $\eta_\varepsilon(\mathcal{L}) \leq 1$, the output of the simulator described above is within statistical distance εm of honestly generated random input and an honestly generated proof as in Figure 1. In particular, the proof system in Figure 1 is statistical zero knowledge for negligible ε.*

Proof. Notice that, conditioned on the random input $\mathbf{t}_i$, the distribution of $\mathbf{e}_i$ is exactly $D_{\mathcal{L}+\mathbf{t}_i,s}$. So, we only need to show that the random input $\mathbf{t}_1, \ldots, \mathbf{t}_m \in \mathcal{P}(\mathbf{B})$ chosen by the simulator is within statistical distance εm of uniform. Indeed, this follows from Lemma 2.8 and the union bound. □

The proof of completeness is a bit tedious and nearly identical to proofs of similar statements in [AR04, PV08, DR16]. We include a proof in Appendix A.

Lemma 3.3 (Completeness). *For any lattice $\mathcal{L} \subset \mathbb{Q}^n$ with $\eta_{1/2}(\mathcal{L}) \leq 1$, the proof given in Figure 1 will be accepted except with negligible probability. I.e., the proof system is complete.*

Soundness. We now show the soundness of the proof system shown in Figure 1, using Theorem 1.6. We note that [DR16] contains an implicit proof of a very similar result in a different context. (Dadush and Regev conjectured a form of Theorem 1.6 and showed a number of implications [DR16]. In particular, they showed that with non-negligible probability over a *single* uniformly random shift $\mathbf{t} \in \mathbb{R}^n/\mathcal{L}$, there is no list of vectors $\mathbf{e}_1, \ldots, \mathbf{e}_m \in \mathcal{L} + \mathbf{t}$ with small covariance.)

Theorem 3.4. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ with basis $\mathbf{B}$ satisfying $\eta_{1/2}(\mathcal{L}) \geq 100 C_\eta(n)$ (and in particular any lattice with $\eta_{1/2}(\mathcal{L}) \geq 1000(\log(n) + 2)$), if $\mathbf{t}_1, \ldots, \mathbf{t}_m$ are sampled uniformly from $\mathbb{R}^n/\mathcal{L}$, then the probability that there exists any proof $\mathbf{e}_1, \ldots, \mathbf{e}_m$ with $\mathbf{e}_i \equiv \mathbf{t} \bmod \mathcal{L}$ and*

$$\left\| \sum \mathbf{e}_i \mathbf{e}_i^T \right\| \leq 3m$$

is at most $\exp(-\Omega(m^2))$.

Proof. By the definition of $C_\eta(n)$ there is a lattice projection π such that $\det(\pi(\mathcal{L})) \geq 100^k$, where $k := \operatorname{rank}(\pi(\mathcal{L}))$. For any $\mathbf{e}_1, \ldots, \mathbf{e}_m$ with $\mathbf{e}_i \equiv \mathbf{t}_i \bmod \mathcal{L}$, we have

$$\begin{aligned}
\left\| \sum \mathbf{e}_i \mathbf{e}_i^T \right\| &\geq \left\| \sum \pi(\mathbf{e}_i) \pi(\mathbf{e}_i)^T \right\| \\
&\geq \frac{1}{k} \operatorname{Tr} \left(\sum \pi(\mathbf{e}_i) \pi(\mathbf{e}_i)^T \right) \\
&= \frac{1}{k} \sum \|\pi(\mathbf{e}_i)\|^2 \\
&\geq \frac{1}{k} \sum \operatorname{dist}(\pi(\mathbf{t}_i), \pi(\mathcal{L}))^2 ,
\end{aligned}$$

where the first inequality on the spectral norms follows from the fact that $\langle \mathbf{u}, \pi(\mathbf{e}_i) \rangle = \langle \pi(\mathbf{u}), \pi(\mathbf{e}_i) \rangle$ and $\|\pi(\mathbf{u})\| \leq \|\mathbf{u}\|$; the second inequality follows from the fact that the spectral norm is the largest eigenvalue and the trace is the sum of the k eigenvalues; and the equality is by definition of trace.

Now by Claim 2.4, $\pi(\mathbf{t}_i)$ is uniformly distributed mod $\pi(\mathcal{L})$, and therefore by Lemma 2.3,

$$\mathbb{E}[\operatorname{dist}(\pi(\mathbf{t}_i), \pi(\mathcal{L}))^2] \geq \mu(\pi(\mathcal{L}))^2/4.$$

Furthermore, since the $\mathbf{t}_i$ are independent and identically distributed with $\operatorname{dist}(\pi(\mathbf{t}_i), \pi(\mathcal{L})) \leq \mu(\pi(\mathcal{L}))$, we can apply the Chernoff-Hoeffding bound (Lemma 2.19) to get

$$\Pr\left[\sum \operatorname{dist}(\pi(\mathbf{t}_i), \pi(\mathcal{L}))^2 \leq m\mu(\pi(\mathcal{L}))^2/5\right] \leq \exp(-Cm^2)\ .$$

The result follows by noting that $\mu(\pi(\mathcal{L}))^2/(5k) \geq 3$ by Claim 2.1, together with the fact that $\det(\pi(\mathcal{L})) \geq 100^k$. □

Corollary 3.5 (Soundness). *For any $\varepsilon \in (0, 1/2)$ and lattice $\mathcal{L} \subset \mathbb{R}^n$ with basis $\mathbf{B}$ satisfying $n \geq 2$ and $\eta_\varepsilon(\mathcal{L}) \geq 100C_\eta(n)\sqrt{\log(1/\varepsilon)}$ (and in particular any lattice with $\eta_\varepsilon(\mathcal{L}) \geq 1000(\log(n)+2)\sqrt{\log(1/\varepsilon)}$), if $\mathbf{t}_1, \ldots, \mathbf{t}_m$ are sampled uniformly from $\mathcal{P}(\mathbf{B})$, then the probability that there exists a proof $\mathbf{e}_1, \ldots, \mathbf{e}_m$ with $\mathbf{e}_i \equiv \mathbf{t} \bmod \mathcal{L}$ and*

$$\left\| \sum \mathbf{e}_i \mathbf{e}_i^T \right\| \leq 3m$$

is at most $\exp(-\Omega(m^2))$. In other words, the proof system in Figure 1 is $\exp(-\Omega(m^2))$-statistically sound.

Proof. By Lemma 2.9, we have $\eta_{1/2} \geq 100C_\eta(n)$, and the result follows from Theorem 3.4. □

Making the prover efficient. Finally, following [PV08] we observe that the prover in the proof system shown in Figure 1 can be made efficient if we relax the approximation factor. In particular, if $\eta_\varepsilon(\mathcal{L}) \leq 1/\sqrt{n \log n}$, then by Corollary 2.14, there is in fact an efficient prover. Theorem 1.2 then follows immediately from the above analysis.

3.2 A Proof via Entropy Approximation

We recall from Goldreich, Sahai, and Vadhan [GSV99] the Entropy Approximation problem, which asks us to approximate the entropy of the distribution obtained by calling some input circuit $\mathcal{C}$ on the uniform distribution over its input space. In particular, we recall that [GSV99] proved that this problem is NISZK-complete. (Formally, we only need the fact that Entropy Approximation is in NISZK.)

Definition 3.6. An instance of the Entropy Approximation problem is a circuit $\mathcal{C}$ and an integer k. It is a YES instance if $H(\mathcal{C}(U)) > k+1$ and a NO instance if $H(\mathcal{C}(U)) < k-1$, where U is the uniform distribution on the input space of $\mathcal{C}$.

Theorem 3.7 ([GSV99]). *Entropy Approximation is* NISZK*-complete.*

In the rest of this section, we show a Karp reduction from $O(\log(n)\sqrt{\log(1/\varepsilon)})$-$\mathsf{GapSPP}_\varepsilon$ to Entropy Approximation. I.e., we give an efficient algorithm that takes as input a basis for a lattice $\mathcal{L}$ and outputs a circuit $\mathcal{C}_\mathcal{L}$ such that (1) if $\eta_\varepsilon(\mathcal{L}) \leq 1$, then $H(\mathcal{C}_\mathcal{L}(U))$ is large; but (2) if $\eta_\varepsilon(\mathcal{L}) \geq C\log(n)\sqrt{\log(1/\varepsilon)}$, then $H(\mathcal{C}_\mathcal{L}(U))$ is small.

Intuitively, we want to use a circuit that samples from the *continuous* Gaussian with parameter one modulo the lattice $\mathcal{L}$. Then, by Claim 2.7, if $\eta_\varepsilon(\mathcal{L}) \leq 1$, the resulting distribution will be nearly uniform over $\mathbb{R}^n/\mathcal{L}$. On the other hand, we know that, with high probability, the continuous Gaussian lies in a set of

volume roughly one. And, by definition, if $\eta_\varepsilon(\mathcal{L}) \geq \Omega(C_\eta(n)\sqrt{\log(1/\varepsilon)})$, then there exists a projection π such that, say, $\mathrm{vol}(\pi(\mathbb{R}^n/\mathcal{L})) = \det(\pi(\mathcal{L})) \geq 100$. Therefore, the projected Gaussian lies in a small fraction of $\pi(\mathbb{R}^n/\mathcal{L})$ with high probability.

To make this precise, we must discretize $\mathbb{R}^n/\mathcal{L}$ appropriately to, say, $(\mathcal{L}/q)/\mathcal{L}$ for some large integer $q > 1$ and sample from a discretized version of the continuous Gaussian. Naturally, we choose $D_{\mathcal{L}/q}$. The following theorem shows that $D_{\mathcal{L}/q} \bmod \mathcal{L}$ lies in a small subset of $(\mathcal{L}/q)/\mathcal{L}$ when $\eta_{1/2}(\mathcal{L})$ is large.

Theorem 3.8. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ with sufficiently large n and integer $q \geq 2^n(\eta_{2^{-n}}(\mathcal{L}) + \mu(\mathcal{L}))$, if $\eta_{1/2}(\mathcal{L}) \geq 1000C_\eta(n)$ (and in particular if $\eta_{1/2}(\mathcal{L}) \geq 10^4(\log(n)+2)$), then there is a subset $S \subset (\mathcal{L}/q)/\mathcal{L}$ with $|S| \leq q^n/200$ such that*

$$\Pr_{\mathbf{X} \sim D_{\mathcal{L}/q} \bmod \mathcal{L}}[\mathbf{X} \in S] \geq \frac{9}{10} .$$

Proof. It is easy to see that $D_{\mathcal{L}/q}$ is statistically close to the distribution obtained by sampling from a *continuous* Gaussian with parameter one and rounding to the closest vector in $\mathcal{L}/q$. (One must simply recall from Lemma 2.6 that nearly all of the mass of $D_{\mathcal{L}/q}$ lies in a ball of radius $\sqrt{n}$ and notice that for such short points, shifts of size $\mu(\mathcal{L}/q) < 2^{-n}$ have little effect on the Gaussian mass.) It therefore suffices to show that the above probability is at least $19/20$ when $\mathbf{X}$ is sampled from this new distribution. We write $\mathsf{CVP}(\mathbf{t})$ for the closest vector in $\mathcal{L}/q$ to $\mathbf{t}$.

By assumption, there is a lattice projection π onto a k-dimensional subspace such that $\det(\pi(\mathcal{L})) \geq 1000^k$. Notice that $\|\pi(\mathsf{CVP}(\mathbf{t}))\| \leq \|\pi(\mathbf{t})\| + \mu(\mathcal{L})/q \leq \|\mathbf{t}\| + 2^{-n}$ for any $\mathbf{t} \in \mathbb{R}^n$. In particular, if $\mathbf{X}$ is sampled from a continuous Gaussian with parameter one,

$$\Pr\left[\|\pi(\mathsf{CVP}(\mathbf{X}))\| \geq \sqrt{k}\right] \leq \Pr\left[\|\pi(\mathbf{X})\| \geq \sqrt{k} - 2^{-n}\right] \leq \frac{1}{20} ,$$

where we have applied Lemma 2.20. But, by Lemma 2.2, there are at most $(q/200)^k$ points $\mathbf{y} \in (\pi(\mathcal{L})/q)/\pi(\mathcal{L})) \cap \sqrt{k}B_2^k$. Therefore, there are at most $q^n/200^k \leq q^n/200$ points $\mathbf{y} \in (\mathcal{L}/q)/\mathcal{L}$ with $19/20$ of the mass, as needed. □

Corollary 3.9. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ with $n \geq 2$, $\varepsilon \in (0, 1/2)$, and integer $q \geq 2$, let $\mathbf{X} \sim D_{\mathcal{L}/q} \bmod \mathcal{L}$. Then,*

1. *if $\eta_\varepsilon(\mathcal{L}) \leq 1$, then $H(\mathbf{X}) > n \log_2 q - 2$; but*
2. *if $\eta_\varepsilon(\mathcal{L}) \geq 1000C_\eta(n) \cdot \sqrt{\log(1/\varepsilon)}$ (and in particular if $\eta_\varepsilon(\mathcal{L}) \geq 10^4 \log(n)\sqrt{\log(1/\varepsilon)}$) and $q \geq 2^n(\eta_{2^{-n}}(\mathcal{L}) + \mu(\mathcal{L}))$, then $H(\mathbf{X}) < n \log_2 q - 6$.*

Proof. Suppose that $\eta_\varepsilon(\mathcal{L}) \leq 1$. Then, by Claim 2.7, for any $\mathbf{y} \in (\mathcal{L}/q)/\mathcal{L}$,

$$\Pr_{\mathbf{X} \sim D_{\mathcal{L}/q} \bmod \mathcal{L}}[\mathbf{X} = \mathbf{y}] = \frac{\rho(\mathcal{L} + \mathbf{y})}{\rho(\mathcal{L}/q)} \leq \frac{1+\varepsilon}{1-\varepsilon} \cdot \frac{1}{q^n} .$$

It follows that

$$H(D_{\mathcal{L}/q} \bmod \mathcal{L}) \geq n \log_2 q + \log_2(1-\varepsilon) - \log_2(1+\varepsilon) > n \log_2 q - 2 ,$$

as needed.

Suppose, on the other hand, that $\eta_\varepsilon(\mathcal{L}) \geq 1000 C_\eta(n) \cdot \sqrt{\log(1/\varepsilon)}$ and $q \geq 2^n(\eta_{2^{-n}}(\mathcal{L}) + \mu(\mathcal{L}))$. By Lemma 2.9, $\eta_{1/2}(\mathcal{L}) \geq 1000 C_\eta(n)$, so that by Theorem 3.8, there is a set S of size $|S| = q^n/200$ with at least $9/10$ of the mass of $D_{\mathcal{L}/q} \bmod \mathcal{L}$. Therefore,

$$H(D_{\mathcal{L}/q} \bmod \mathcal{L}) \leq \frac{9}{10} \cdot \log_2 |S| + \frac{1}{10} \cdot n \log_2 q < n \log_2 q - 6 \ ,$$

as needed. □

Corollary 3.9 shows that, in order to reduce $O(\log(n)\sqrt{\log(1/\varepsilon)})$-$\mathsf{GapSPP}_\varepsilon$ to Entropy Approximation, it suffices to construct a circuit that samples from $D_{\mathcal{L}/q} \bmod \mathcal{L}$. The main result of this section follows immediately from Corollary 2.12.

Theorem 3.10. *There is an efficient Karp reduction from γ-$\mathsf{GapSPP}_\varepsilon$ to Entropy Approximation for*

$$\gamma := O(C_\eta(n)\sqrt{\log(1/\varepsilon)}) \leq O(\log(n)\sqrt{\log(1/\varepsilon)}) \ .$$

and any $\varepsilon \in (0, 1/2)$. I.e., γ-$\mathsf{GapSPP}_\varepsilon$ is in NISZK.

Proof. The reduction behaves as follows on input $\mathcal{L} \subset \mathbb{Q}^n$. By Lemma 2.10, we can find an integer $q \geq 2$ with polynomial bit length that satisfies $q \geq 2^n(\eta_{2^{-n}}(\mathcal{L}) + \mu(\mathcal{L}))$. The reduction constructs the circuit $\mathcal{C}_{\mathcal{L}/q}$ from Corollary 2.12 and outputs the modified circuit $\mathcal{C}_{(\mathcal{L}/q)/\mathcal{L}}$ that takes the output from $\mathcal{C}_{\mathcal{L}/q}$ and reduces it modulo $\mathcal{L}$. It then outputs the Entropy Approximation instance $(\mathcal{C}_{(\mathcal{L}/q)/\mathcal{L}}, k := n \log_2 q - 4)$.

The running time is clear. Suppose that $\eta_\varepsilon(\mathcal{L}) \leq 1$. Then, by Corollary 3.9,

$$H(D_{\mathcal{L}/q} \bmod \mathcal{L}) > n \log_2 q - 2 \ .$$

Since the output of $\mathcal{C}_{(\mathcal{L}/q)/\mathcal{L}}$ is statistically close to $D_{\mathcal{L}/q} \bmod \mathcal{L}$, it follows that $H(\mathcal{C}_{(\mathcal{L}/q)/\mathcal{L}}(U)) > n \log_2 q - 3$, as needed.

If, on the other hand, $\eta_\varepsilon(\mathcal{L}) \geq \Omega(C_\eta(n) \cdot \sqrt{\log(1/\varepsilon)})$, then by Corollary 3.9,

$$H(D_{\mathcal{L}/q} \bmod \mathcal{L}) < n \log_2 q - 6 \ .$$

Since the output of $\mathcal{C}_{(\mathcal{L}/q)/\mathcal{L}}$ is statistically close to $D_{\mathcal{L}/q} \bmod \mathcal{L}$, it follows that $H(\mathcal{C}_{(\mathcal{L}/q)/\mathcal{L}}(U)) < n \log_2 q - 5$. □

4 A coNP Proof for $O(\log n)$-GapSPP

We will need the following result from [RS17], which extends Theorem 1.6 to smaller ε by noting that $\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\})$ decays at least as quickly as $\rho_{1/s}(\lambda_1(\mathcal{L}^*))$.

Theorem 4.1. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and any $\varepsilon \in (0, 1/2)$,*

$$\eta_\varepsilon(\mathcal{L})^2 \leq C_\eta(n)^2 \eta_{\det}(\mathcal{L})^2 + \frac{\log(1/\varepsilon)}{\pi \lambda_1(\mathcal{L}^*)^2} \leq 100(\log n + 2)^2 \eta_{\det}(\mathcal{L})^2 + \frac{\log(1/\varepsilon)}{\pi \lambda_1(\mathcal{L}^*)^2} \ .$$

Proof. We may assume without loss of generality that $\eta_{\det}(\mathcal{L}) = 1$. Then, by definition, $\rho_{1/C_\eta(n)}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq 1/2$. Therefore, for any $s \geq C_\eta(n)$,

$$\begin{aligned}\rho_{1/s}(\mathcal{L}^*) &= 1 + \sum_{\mathbf{w} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \exp(-\pi(s^2 - C_\eta(n)^2)\|\mathbf{w}\|^2)\rho_{1/C_\eta(n)}(\mathbf{w}) \\ &\leq 1 + \sum_{\mathbf{w} \in \mathcal{L}^* \setminus \{\mathbf{0}\}} \exp(-\pi(s^2 - C_\eta(n)^2)\lambda_1(\mathcal{L}^*)^2)\rho_{1/C_\eta(n)}(\mathbf{w}) \\ &\leq 1 + \exp(-\pi(s^2 - C_\eta(n)^2)\lambda_1(\mathcal{L}^*)^2)/2 \ ,\end{aligned}$$

and the result follows. □

Next, we prove an easy lower bound with a similar form (by taking the average of two trivial lower bounds).

Lemma 4.2. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and any $\varepsilon \in (0, 1/2)$,*

$$\eta_\varepsilon(\mathcal{L})^2 \geq \eta_{\det}(\mathcal{L})^2/8 + \frac{\log(2/\varepsilon)}{2\pi\lambda_1(\mathcal{L}^*)^2} \ .$$

Proof. First, note that $\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \geq 2\rho_{1/s}(\lambda_1(\mathcal{L}^*))$. Rearranging, we see that

$$\eta_\varepsilon(\mathcal{L})^2 \geq \frac{\log(2/\varepsilon)}{\pi\lambda_1(\mathcal{L}^*)^2} \ .$$

On the other hand, recall that for any lattice projection π onto a subspace W, $\det(\mathcal{L}^* \cap W) = 1/\det(\pi(\mathcal{L}))$. I.e., $\eta_{\det}(\mathcal{L}) = \max_{\mathcal{L}' \subseteq \mathcal{L}^*} \det(\mathcal{L}')^{-1/\operatorname{rank}(\mathcal{L}')}$. So, suppose $s \leq \eta_{\det}(\mathcal{L})/2$. Then, by Lemma 2.5,

$$\rho_{1/s}(\mathcal{L}^*) = \max_{\mathcal{L}' \subseteq \mathcal{L}^*} \rho_{1/s}(\mathcal{L}') \geq \max_{\mathcal{L}' \subseteq \mathcal{L}^*} s^{-\operatorname{rank}(\mathcal{L}')}/\det(\mathcal{L}') \geq 2 \ .$$

So, $\eta_\varepsilon(\mathcal{L})^2 \geq \eta_1(\mathcal{L})^2 \geq \eta_{\det}(\mathcal{L})^2/4$. The result follows by taking the average of the two bounds. □

The main theorem of this section now follows immediately.

Theorem 4.3. *For any $\varepsilon \in (0, 1/2)$, γ-$\mathsf{GapSPP}_\varepsilon$ is in coNP for $\gamma = O(C_\eta(n)) \leq O(\log n)$.*

Proof. Let $\gamma := 2\sqrt{2}C_\eta(n)$. On input a lattice $\mathcal{L} \subset \mathbb{R}^n$, the prover simply sends a lattice projection π with $\det(\pi(\mathcal{L}))^{1/\operatorname{rank}(\pi(\mathcal{L}))} = \eta_{\det}(\mathcal{L})$ and a vector $\mathbf{w} \in \mathcal{L}^*$ with $\|\mathbf{w}\| = \lambda_1(\mathcal{L}^*)$. The verifier checks that π is indeed a lattice projection and that $\mathbf{w} \in \mathcal{L}^* \setminus \{\mathbf{0}\}$. It then answers NO if and only if

$$\gamma^2 \det(\pi(\mathcal{L}))^{2/\operatorname{rank}(\pi(\mathcal{L}))}/8 + \frac{\log(1/\varepsilon)}{\pi\|\mathbf{w}\|^2} > \gamma^2 \ . \tag{4.1}$$

To prove completeness, suppose that $\eta_\varepsilon(\mathcal{L}) > \gamma$. Then, by Theorem 4.1,

$$\gamma^2\eta_{\det}(\mathcal{L})^2/8 + \frac{\log(1/\varepsilon)}{\pi\lambda_1(\mathcal{L}^*)^2} \geq \eta_\varepsilon(\mathcal{L})^2 > \gamma^2 \ .$$

I.e., there exists a valid proof, as needed.

To prove soundness, suppose that $\eta_\varepsilon(\mathcal{L}) \le 1$. Then, by Lemma 4.2,

$$\eta_{\det}(\mathcal{L})^2/8 + \frac{\log(1/\varepsilon)}{2\pi\lambda_1(\mathcal{L}^*)^2} \le \eta_\varepsilon(\mathcal{L})^2 \le 1 \ .$$

Therefore,

$$\begin{aligned}\gamma^2\eta_{\det}(\mathcal{L})^2/8 + \frac{\log(1/\varepsilon)}{\pi\lambda_1(\mathcal{L}^*)^2} &\le \frac{\gamma^2\eta_{\det}(\mathcal{L})^2/8 + \frac{\log(1/\varepsilon)}{\pi\lambda_1(\mathcal{L}^*)^2}}{\eta_{\det}(\mathcal{L})^2/8 + \frac{\log(1/\varepsilon)}{2\pi\lambda_1(\mathcal{L}^*)^2}} \\ &\le \max\{\gamma^2, 2\} \\ &\le \gamma^2 \ .\end{aligned}$$

In other words, Equation (4.1) cannot hold for any pair $\mathbf{w} \in \mathcal{L}^* \setminus \{\mathbf{0}\}$ and lattice projection π. I.e., the verifier will always answer YES, as needed. □

Finally, we derive the following corollary.

Corollary 4.4. *For any $\varepsilon \in (0, 1/2)$, γ-coGapSPP$_\varepsilon$ has an* SZK *proof system with an efficient prover for*

$$\gamma := O(C_\eta(n) + \sqrt{\log(1/\varepsilon)/\log n}) \le O(\log n + \sqrt{\log(1/\varepsilon)/\log n}) \ .$$

Proof. By Theorem 2.17, γ-GapSPP$_\varepsilon$ is in SZK. Since SZK is closed under complements [SV97, Oka96], γ-coGapSPP$_\varepsilon$ is in SZK as well. By Theorem 4.3, γ-coGapSPP$_\varepsilon$ is in NP. The result then follows by the fact that any language in SZK ∩ NP has an SZK proof system with an efficient prover [NV06]. □

5 An SZK Proof for $O(\sqrt{n})$-GapCRP

In this section we prove that $O(\sqrt{n})$-GapCRP is in SZK, which improves the previous known result by a $\omega(\sqrt{\log n})$ factor [PV08]. First we need the following result from [CDLP13].

Lemma 5.1. *For any lattice $\mathcal{L}$ and parameter $s > 0$,*

$$\rho_s(\mathcal{L}) \cdot \gamma_s(\mathcal{V}(\mathcal{L})) \le 1 \ .$$

Here we prove an upper bound on the smoothing parameter of a lattice in terms of its covering radius. This bound is implicit in [DR16].

Lemma 5.2. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ and $\varepsilon > 0$, we have*

$$\eta_\varepsilon(\mathcal{L}) \le \sqrt{\frac{\pi}{\log(1+\varepsilon)}} \cdot \mu(\mathcal{L}) \ .$$

In particular, $\eta_\varepsilon(\mathcal{L}) \le O(\mu(\mathcal{L}))$ for any $\varepsilon \ge \Omega(1)$.

Proof.

$$\begin{aligned}
\rho_{1/s}(\mathcal{L}^*) &= s^{-n} \cdot \det(\mathcal{L}) \cdot \rho_s(\mathcal{L}) && \text{(Lemma 2.5)} \\
&\leq \frac{s^{-n} \cdot \det(\mathcal{L})}{\gamma_s(\mathcal{V}(\mathcal{L}))} && \text{(Lemma 5.1)} \\
&\leq \frac{s^{-n} \cdot \det(\mathcal{L})}{\int_{\mathcal{V}(\mathcal{L})} s^{-n} \cdot \exp(-\pi \mathbf{x}^2/s^2)\, \mathrm{d}\mathbf{x}} \\
&\leq \frac{\det(\mathcal{L})}{\int_{\mathcal{V}(\mathcal{L})} \exp(-\pi \mu(\mathcal{L})^2/s^2)\, \mathrm{d}\mathbf{x}} \\
&\leq \exp(\pi \mu(\mathcal{L})^2/s^2)\ ,
\end{aligned}$$

where we used the fact that $\|\mathbf{x}\| \leq \mu(\mathcal{L})$ for any $\mathbf{x} \in \mathcal{V}(\mathcal{L})$. By setting $s = \sqrt{\frac{\pi}{\log(1+\varepsilon)}} \cdot \mu(\mathcal{L})$ we have the desired result. □

Theorem 5.3. *The problem $O(\sqrt{n})$-GapCRP has an SZK proof system with an efficient prover, as does $O(\sqrt{n})$-coGapCRP .*

Proof. Fix some some constant $\varepsilon \in (0, 1/2)$. By Lemma 2.13 and Lemma 5.2, we know that there exist C_1 and C_2 such that

$$C_1 \eta_\varepsilon(\mathcal{L}) \leq \mu(\mathcal{L}) \leq C_2 \sqrt{n} \cdot \eta_\varepsilon(\mathcal{L}),$$

and hence there is a simple reduction from $O(\sqrt{n})$-GapCRP to $O(1)$-GapSPP$_\varepsilon$. It follows from Theorem 2.17 that $O(\sqrt{n})$-GapCRP is in SZK. To see that the prover can be made efficient, we recall from [GMR04] that $O(\sqrt{n})$-GapCRP is in NP ∩ coNP. The result then follows by the fact that any language in SZK ∩ NP has an SZK proof system with an efficient prover [NV06]. □

References

[ADRS15] D. Aggarwal, D. Dadush, O. Regev, and N. Stephens-Davidowitz. Solving the shortest vector problem in 2^n time using discrete Gaussian sampling. In *STOC*, pages 733–742. 2015.

[ADS15] D. Aggarwal, D. Dadush, and N. Stephens-Davidowitz. Solving the closest vector problem in 2^n time - the discrete Gaussian strikes again! In *FOCS*, pages 563–582. 2015.

[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 ∩ coNP. *J. ACM*, 52(5):749–765, 2005. Preliminary version in FOCS 2004.

[Bab85a] L. Babai. Trading group theory for randomness. In *STOC*, pages 421–429. 1985.

[Bab85b] L. Babai. On Lovász' lattice reduction and the nearest lattice point problem. *Combinatorica*, 6(1):1–13, 1986. Preliminary version in STACS 1985.

[Ban93] W. Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.

[Ban95] W. Banaszczyk. Inequalites for convex bodies and polar reciprocal lattices in $\mathbb{R}^n$. *Discrete & Computational Geometry*, 13:217–231, 1995.

[BDMP88] M. Blum, A. De Santis, S. Micali, and G. Persiano. Noninteractive zero-knowledge. *SIAM J. Comput.*, 20(6):1084–1118, 1991. Preliminary version in STOC 1988.

[BG89] M. Bellare and S. Goldwasser. New paradigms for digital signatures and message authentication based on non-interative zero knowledge proofs. In *CRYPTO*, pages 194–211. 1989.

[BGV12] Z. Brakerski, C. Gentry, and V. Vaikuntanathan. (Leveled) fully homomorphic encryption without bootstrapping. *TOCT*, 6(3):13, 2014. Preliminary version in ITCS 2012.

[BHZ87] R. B. Boppana, J. Håstad, and S. Zachos. Does co-NP have short interactive proofs? *Inf. Process. Lett.*, 25(2):127–132, 1987.

[BLP+13] Z. Brakerski, A. Langlois, C. Peikert, O. Regev, and D. Stehlé. Classical hardness of learning with errors. In *STOC*, pages 575–584. 2013.

[BV11] Z. Brakerski and V. Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. *SIAM J. Comput.*, 43(2):831–871, 2014. Preliminary version in FOCS 2011.

[Can01] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145. 2001.

[CDLP13] K. Chung, D. Dadush, F. Liu, and C. Peikert. On the lattice smoothing parameter problem. In *IEEE Conference on Computational Complexity*, pages 230–241. 2013.

[DN00] C. Dwork and M. Naor. Zaps and their applications. *SIAM J. Comput.*, 36(6):1513–1543, 2007.

[DPV11] D. Dadush, C. Peikert, and S. Vempala. Enumerative lattice algorithms in any norm via M-ellipsoid coverings. In *FOCS*, pages 580–589. 2011.

[DR16] D. Dadush and O. Regev. Towards strong reverse Minkowski-type inequalities for lattices. In *FOCS*, pages 447–456. 2016.

[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.

[GMR85] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989. Preliminary version in STOC 1985.

[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.

[GMW87] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or A completeness theorem for protocols with honest majority. In *STOC*, pages 218–229. 1987.

[GMW91] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity for all languages in NP have zero-knowledge proof systems. *J. ACM*, 38(3):691–729, 1991.

[GPV08] C. Gentry, C. Peikert, and V. Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In *STOC*, pages 197–206. 2008.

[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.

[GSV99] O. Goldreich, A. Sahai, and S. P. Vadhan. Can statistical zero knowledge be made non-interactive? or on the relationship of SZK and NISZK. In *CRYPTO*, pages 467–484. 1999.

[Hoe63] W. Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58:13–30, 1963.

[HR14] I. Haviv and O. Regev. On the lattice isomorphism problem. In *SODA*, pages 391–404. 2014.

[JS98] A. Joux and J. Stern. Lattice reduction: A toolbox for the cryptanalyst. *J. Cryptology*, 11(3):161–185, 1998.

[Kan83] R. Kannan. Improved algorithms for integer programming and related lattice problems. In *STOC*, pages 193–206. 1983.

[Kle00] P. N. Klein. Finding the closest lattice vector when it's unusually close. In *SODA*, pages 937–941. 2000.

[Len83] H. W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, November 1983.

[LLL82] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261(4):515–534, December 1982.

[MG02] D. Micciancio and S. Goldwasser. *Complexity of Lattice Problems: a cryptographic perspective*, volume 671 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, Massachusetts, 2002.

[MP12] D. Micciancio and C. Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In *EUROCRYPT*, pages 700–718. 2012.

[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.

[NS01] P. Q. Nguyen and J. Stern. The two faces of lattices in cryptology. In *CaLC*, pages 146–180. 2001.

[NV06] M. Nguyen and S. P. Vadhan. Zero knowledge with efficient provers. In *STOC*, pages 287–295. 2006.

[NY90] M. Naor and M. Yung. Public-key cryptosystems provably secure against chosen ciphertext attacks. In *STOC*, pages 427–437. 1990.

[Odl90] A. M. Odlyzko. The rise and fall of knapsack cryptosystems. In C. Pomerance, editor, *Cryptology and Computational Number Theory*, volume 42 of *Proceedings of Symposia in Applied Mathematics*, pages 75–88. 1990.

[Oka96] T. Okamoto. On relationships between statistical zero-knowledge proofs. *J. Comput. Syst. Sci.*, 60(1):47–108, 2000. Preliminary version in STOC 1996.

[Pei09] C. Peikert. Public-key cryptosystems from the worst-case shortest vector problem. In *STOC*, pages 333–342. 2009.

[PV08] C. Peikert and V. Vaikuntanathan. Noninteractive statistical zero-knowledge proofs for lattice problems. In *CRYPTO*, pages 536–553. 2008.

[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.

[RS17] O. Regev and N. Stephens-Davidowitz. A reverse Minkowski theorem. In *STOC*, pages 941–953. 2017.

[SV97] A. Sahai and S. P. Vadhan. A complete problem for statistical zero knowledge. *J. ACM*, 50(2):196–249, 2003. Preliminary version in FOCS 1997.

[Ver12] R. Vershynin. *Introduction to the non-asymptotic analysis of random matrices*, chapter 5, pages 210–268. Cambridge University Press, 2012. Available at `http://www-personal.umich.edu/~romanv/papers/non-asymptotic-rmt-plain.pdf`.

A Proof of Lemma 3.3

Definition A.1. For any $\delta > 0$, $S \subseteq \mathbb{R}^n$, we say that $A \subseteq S$ is a *δ-net of S* if for each $\mathbf{v} \in S$, there is some $\mathbf{u} \in A$ such that $\|\mathbf{u} - \mathbf{v}\| \leq \delta$.

Lemma A.2. *For any $\delta > 0$, there exists a δ-net of the unit sphere in $\mathbb{R}^n$ with at most $(1 + 2/\delta)^n$ points.*

Proof. Let N be maximal such that N points can be placed on the unit sphere in such a way that no pair of points is within distance δ of each other. Clearly, there exists a δ-net of size N.

So, it suffices to show that any collection of vectors A in the unit sphere with $|A| > (1 + 2/\delta)^n$ must contain two points within distance δ of each other. Let

$$\mathcal{B} := \bigcup_{\mathbf{u} \in A} ((\delta/2)B_2^n + \mathbf{u})$$

be the union of balls of radius $\delta/2$ centered at each point in A. Notice that $\mathcal{B} \subseteq (1 + \delta/2)B_2^n$. If all of these balls were disjoint, then we would have

$$\mathrm{vol}(\mathcal{B}_2^n) = |A| \cdot (\delta/2)^n \,\mathrm{vol}(B_2^n) > \mathrm{vol}((1 + \delta/2)B_2^n) \,,$$

a contradiction. Therefore, two such balls must overlap. I.e., there must be two points within distance δ of each other, as needed. □

We will need the following result from [Ver12, Lemma 5.4].

Lemma A.3. *For a symmetric matrix $M \in \mathbb{R}^{n \times n}$ and a δ-net of the unit sphere A with $\delta \in (0, 1/2)$,*

$$\|M\| \leq \frac{1}{1 - 2\delta} \cdot \max_{\mathbf{v} \in A} |\langle M\mathbf{v}, \mathbf{v}\rangle| \,.$$

We will also need the following result from [MP12, Lemma 2.8], which shows that the discrete Gaussian distribution is subgaussian.

Lemma A.4. *For any lattice $\mathcal{L} \subset \mathbb{R}^n$ with $\eta_{1/2}(\mathcal{L}) \leq 1$, shift vector $\mathbf{t} \in \mathbb{R}^n$, unit vector $\mathbf{v} \in \mathbb{R}^n$, and any $r > 0$,*

$$\Pr_{\mathbf{X} \sim D_{\mathcal{L}-\mathbf{t}}}[|\langle \mathbf{v}, \mathbf{X} \rangle| \geq r] \leq 10 \exp(-\pi r^2) \ .$$

Proof of Lemma 3.3. Let $\{\mathbf{v}_1, \ldots, \mathbf{v}_N\}$ be a $(1/10)$-net of the unit sphere with $N \leq 25^n$, as guaranteed by Lemma A.2. By Lemma A.4, we have that for any $\mathbf{e}_i$ in the proof, any $\mathbf{v}_j$, and any $r \geq 0$, $\Pr[|\langle \mathbf{v}_j, \mathbf{e}_i \rangle| \geq r] \leq 10 \exp(-\pi r^2)$. Therefore, by Lemma 2.20

$$\Pr\Big[\sum_i \langle \mathbf{v}_j, \mathbf{e}_i \rangle^2 \geq r\Big] \leq 2^m e^{-\pi r/2} \ .$$

Applying the union bound, we have

$$\Pr\Big[\exists j, \ \sum_i \langle \mathbf{v}_j, \mathbf{e}_i \rangle^2 \geq r\Big] \leq N 2^m e^{-\pi r/2} \ .$$

Taking $r := 2m$, we see that this probability is negligible. Applying Lemma A.3 shows that

$$\Big\| \sum_i \mathbf{e}_i \mathbf{e}_i^T \Big\| \leq 2m \cdot \frac{5}{4} < 3m \ ,$$

except with negligible probability, as needed. □