

A Note on Discrete Gaussian Combinations of Lattice Vectors

Divesh Aggarwal Oded Regev*

Received January 10, 2014; Revised April 4, 2016; Published June 8, 2016

Abstract: We prove a local central limit theorem for the sum of one-dimensional discrete Gaussians in n -dimensional space. In more detail, we analyze the distribution of $\sum_{i=1}^m v_i \mathbf{x}_i$ where $\mathbf{x}_1, \dots, \mathbf{x}_m$ are fixed vectors from some lattice $\mathcal{L} \subset \mathbb{R}^n$ and v_1, \dots, v_m are chosen independently from a discrete Gaussian distribution over \mathbb{Z} . We show that under a natural constraint on $\mathbf{x}_1, \dots, \mathbf{x}_m$, if the v_i are chosen from a wide enough Gaussian, the sum is statistically close to a discrete Gaussian over \mathcal{L} . We also analyze the case of $\mathbf{x}_1, \dots, \mathbf{x}_m$ that are themselves chosen from a discrete Gaussian distribution (and fixed).

Our results simplify and qualitatively improve upon a recent result by Agrawal, Gentry, Halevi, and Sahai [1].

Key words and phrases: Local central limit theorem, Discrete Gaussians, Lattices

1 Introduction

The study of discrete Gaussian distributions has been a crucial step in most of the work in lattice-based cryptography in the last few years. A discrete Gaussian distribution over some fixed lattice \mathcal{L} , denoted as $\mathcal{D}_{\mathcal{L},s}$ for some parameter $s > 0$, is a distribution in which each lattice point is sampled with probability proportional to the probability density function of a continuous Gaussian distribution of width s evaluated at that point. There have been several results, e.g. [11, 6, 12, 4, 10], showing that these functions have many useful properties similar to the corresponding continuous functions. However, some aspects of the discrete Gaussian distribution are not as well understood as those of its continuous counterpart.

*Courant Institute of Mathematical Sciences, New York University. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1320188. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Our Result. We analyze the following distribution. Let $\mathbf{x}_1, \dots, \mathbf{x}_m$ be m fixed points in \mathbb{Z}^n for some $m > n$. Let X be the $n \times m$ matrix formed by the vectors $\mathbf{x}_1, \dots, \mathbf{x}_m$ as the column vectors. For $r > 0$, define the distribution

$$\mathcal{E}_{X,r} := \{X \cdot \mathbf{v} : \mathbf{v} \leftarrow \mathcal{D}_{\mathbb{Z}^m, r}\}.$$

We show in Theorem 3.2 in Section 3 that if X satisfies a certain constraint and if r is large enough, then the distribution $\mathcal{E}_{X,r}$ is statistically close to the discrete Gaussian distribution with appropriate covariance (denoted as $\mathcal{D}_{\mathbb{Z}^n, rX^T}$). The constraint on X that we need is that X has small entries and additionally there are short vectors $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{Z}^m$ such that $X \cdot \mathbf{u}_i = \mathbf{e}_i$ for $1 \leq i \leq n$, where \mathbf{e}_i are unit vectors in n dimensions.

We then show that if the vectors $\mathbf{x}_1, \dots, \mathbf{x}_m$ are chosen according to the discrete Gaussian distribution, then X satisfies the aforementioned constraint, and hence the following result holds:

Theorem 1.1 (Informal). *For integers $n \geq 1$, $m = \text{poly}(n)$, $s = \Omega(\sqrt{\log(n/\varepsilon)})$, and $r = \tilde{\Omega}(ns\sqrt{\log(1/\varepsilon)})$, and $X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, s})^m$, we have with probability $1 - 2^{-n}$ that the statistical distance between $\mathcal{E}_{X,r}$ and $\mathcal{D}_{\mathbb{Z}^n, rX^T}$ is at most ε .*

The formal statement and proof appear as Theorem 5.1 in Section 5. We also mention there how to obtain a more general version for the case in which the columns of X are chosen from a discrete Gaussian distribution over an arbitrary lattice $\mathcal{L} \subset \mathbb{R}^n$; this follows easily from the case of \mathbb{Z}^n . The statement we prove is a bit stronger than what we claim: instead of statistical distance, we actually get point-wise closeness in probability mass (multiplicatively to within $1 \pm \varepsilon$). Such point-wise statements are referred to as *local* central limit theorems, and thus we can describe our result as a local central limit theorem for the sum of one-dimensional discrete Gaussians in n -dimensional space.

The motivation for studying the distribution $\mathcal{E}_{X,r}$ for this choice of X comes from a recent breakthrough construction of multilinear maps from ideal lattices by Garg, Gentry, and Halevi [5] (see [1] for details). Also, more recently, our result has been used by Ling et al. [9] to give a lattice-based traitor tracing scheme.

Our result improves on the main result of Agrawal, Gentry, Halevi, and Sahai [1] who showed a statement similar to Theorem 1.1 with the bound on r depending linearly on m . Improving this dependence on m was left by them as an open question, which we answer here.

Our Techniques. As was done in [1], the first part of our proof is to reduce the question to a question about bounding the successive minima of a so-called orthogonal lattice. Namely, let $A \subset \mathbb{Z}^m$ be the *orthogonal lattice* of X , i.e., the lattice formed by all integer vectors in \mathbb{Z}^m that are orthogonal to all n rows of X , which is (typically) a lattice of rank $m - n$. Our goal then is to find $m - n$ linearly independent vectors in it whose maximum length is as small as possible, i.e., we need to bound the last successive minimum λ_{m-n} of the lattice A . We feel that the techniques used in this part are rather standard, and instead of repeating the proofs, we cite the relevant lemmas from [1].

The main part of our proof (and that of [1]) is therefore bounding $\lambda_{m-n}(A)$. The approach we take is using a simple, yet powerful, idea from a recent result by Kuperberg, Lovett, and Peled [7] (who were interested in an entirely different application, namely, showing the existence of some combinatorial structures like t -wise independent permutations). Using this idea, we get that in order to prove a bound on the successive minima of the orthogonal lattice, it suffices to satisfy two constraints: (1) that X has small entries, and (2) that there are short vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$ such that $X \cdot \mathbf{u}_i = \mathbf{e}_i$ for $1 \leq i \leq n$, where \mathbf{e}_i are unit vectors in n dimensions.

This already gives us sufficient constraints on X and r under which $\mathcal{E}_{X,r}$ is guaranteed to be close to a discrete Gaussian. In order to complete the proof of Theorem 1.1, it suffices to show that the X chosen there satisfies these two constraints with high probability. The first constraint, namely, that X has small entries

follows easily by definition. Proving the second constraint is trickier and we show this in Section 4 using a careful application of the pigeon-hole principle.

Comparison with [1]. As mentioned above, our Theorem 1.1 is similar to the main result of [1], the difference being the quantitative improvement in the bound on r . Namely, the bound on r obtained by [1] is $\Omega(mn \log(1/\varepsilon))$ which could be much worse than our bound for large m . We note that our result gives a worse bound if $s \gg m$. However, for applications like [5, 9], one chooses $s = \Theta(\sqrt{\log(n/\varepsilon)})$ in which case our result gives a better bound for all m, n, ε .

Another advantage of our approach is that our Theorem 3.2 (which we use to prove Theorem 1.1) applies to a *general* choice of X , and not just to a random X . This might be useful in future applications.

In terms of techniques, as mentioned above, the most challenging step in both our proof and the proof in [1] is to bound the last successive minimum λ_{m-n} of the orthogonal lattice A , which is a lattice of rank $m-n$ in dimension m . The way this is done in [1] is by first defining a superlattice A_q of A which is of *full rank*, and then considering the dual lattice M_q of A_q . They then obtain a lower bound on $\lambda_{n+1}(M_q)$, which using Banaszczyk’s transference theorem [3], implies an upper bound on $\lambda_{m-n}(A_q)$. Finally, they argue that this is also an upper bound on $\lambda_{m-n}(A)$. This somewhat indirect proof results in the bound on r depending linearly on m . Improving this dependence was left as an open question in [1]. In comparison, our method to bound $\lambda_{m-n}(A)$ is more direct.

Other Related Work. Micciancio and Peikert [10] recently showed hardness results for the main lattice-based cryptographic problems. One of the key ingredients in the proof was a new convolution theorem which was a strengthening of a previous similar result by Peikert [12]. This theorem, like our result, also looks at sums of discrete Gaussian samples. The difference is that in their statement the combination vector \mathbf{v} is fixed and the matrix X comprises of (spherical) discrete Gaussian with parameter s significantly bigger than the smoothing parameter of the underlying lattice. In our case, the matrix is fixed “once and for all,” and only \mathbf{v} varies. In our setting we essentially analyze the sum of 1-dimensional discrete Gaussians in n -dimensional space.

We note that by combining the result of [10] with that of [1], it might be possible to derive a result similar to our Theorem 1.1, i.e., to improve the dependence of r on m . The idea would be to partition the sum $\sum_{i=1}^m v_i \mathbf{x}_i$ into small blocks, argue that each block is close to a discrete Gaussian (using [1]) and then argue that the overall sum must therefore also be close to a discrete Gaussian (using [10]). As this approach would almost certainly lead to worse parameters and a far more complicated proof, we did not attempt to pursue it.

Finally, in a recent work, Langlois, Stehlé, and Steinfeld [8] showed another Gaussian local central limit theorem specialized to the case of the ring of integers of an algebraic number field.

2 Preliminaries

All logarithms, unless otherwise stated, are to the base 2. Natural logarithms, i.e., to the base e , are denoted by \ln . The norm $\|\cdot\|$ considered in this paper is the ℓ_2 norm, unless otherwise stated. The transpose and inverse of a matrix A are denoted as A^T and A^{-1} , respectively. The inverse transpose of a matrix A is denoted as A^{-T} .

Let D be a discrete distribution. We denote by $D[x]$ the probability it assigns to x , and by $X \sim D$ a random variable distributed according to D . For two distributions D, D' their statistical distance is $\Delta(D; D') = \frac{1}{2} \sum_x |D[x] - D'[x]|$.

Gaussian function. For any $s > 0$, the (spherical) Gaussian function on \mathbb{R}^n with parameter s is defined as $\rho_s(\mathbf{x}) = \exp(-\pi\|\mathbf{x}\|^2/s^2)$ for all $\mathbf{x} \in \mathbb{R}^n$. If $s = 1$, then it is omitted. For a rank- n matrix $S \in \mathbb{R}^{m \times n}$, the ellipsoidal Gaussian function on \mathbb{R}^n with covariance matrix $\Sigma = S^T S$ is defined as:

$$\forall \mathbf{x} \in \mathbb{R}^n, \rho_S(\mathbf{x}) = \exp(-\pi \mathbf{x}^T (S^T S)^{-1} \mathbf{x}) .$$

When $S = sI_n$, ρ_S is the same as ρ_s . Also, if $m = n$, then $\rho_S(\mathbf{x}) = \rho(S^{-T} \mathbf{x})$.

Lattices. A lattice is a discrete additive subgroup of \mathbb{R}^m . A set of linearly independent vectors that generates a lattice is called a basis and is denoted by $\mathbf{B} = \{\mathbf{b}_1, \dots, \mathbf{b}_n\} \subset \mathbb{R}^m$ for integers $m \geq n \geq 1$. The lattice generated by the basis \mathbf{B} is

$$\mathcal{L} = \mathcal{L}(\mathbf{B}) = \left\{ \mathbf{B}\mathbf{z} = \sum_{i=1}^n z_i \mathbf{b}_i : \mathbf{z} \in \mathbb{Z}^n \right\} .$$

We say that the rank of this lattice is n and its dimension is m . For $i = 1, \dots, n$, the *successive minimum* $\lambda_i(\mathcal{L})$ is defined as the smallest value such that a ball of radius $\lambda_i(\mathcal{L})$ centered around the origin contains at least i linearly independent lattice vectors.

The dual lattice \mathcal{L}^* of \mathcal{L} is defined as $\mathcal{L}^* = \{\mathbf{x} \in \mathbb{R}^n : \langle \mathcal{L}, \mathbf{x} \rangle \subseteq \mathbb{Z}\}$. For a lattice \mathcal{L} and positive real $\varepsilon > 0$, the smoothing parameter $\eta_\varepsilon(\mathcal{L})$ is the smallest real $s > 0$ such that $\rho_{1/s}(\mathcal{L}^* \setminus \{\mathbf{0}\}) \leq \varepsilon$, where $\rho_{1/s}(A)$ for a set A denotes $\sum_{\mathbf{x} \in A} \rho_{1/s}(\mathbf{x})$.

For a rank- n lattice \mathcal{L} , $S \in \mathbb{R}^{m \times n}$, and $\mathbf{c} \in \mathbb{R}^n$, the ellipsoidal Gaussian distribution with parameter S and support $\mathcal{L} + \mathbf{c}$ is defined as:

$$\forall \mathbf{x} \in \mathcal{L} + \mathbf{c}, \mathcal{D}_{\mathcal{L}+\mathbf{c},S}(\mathbf{x}) = \frac{\rho_S(\mathbf{x})}{\rho_S(\mathcal{L} + \mathbf{c})} ,$$

where $\rho_S(A)$ for a set A denotes $\sum_{\mathbf{x} \in A} \rho_S(\mathbf{x})$.

Matrices and Singular Values. Let m, n be integers with $m \geq n$. For a rank- n matrix $S \in \mathbb{R}^{m \times n}$, there exist orthogonal matrices $U \in \mathbb{R}^{m \times m}$, $V \in \mathbb{R}^{n \times n}$ (i.e., $U^T U = I$, and $V^T V = I$) and a diagonal matrix¹

$$\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n) \in \mathbb{R}^{m \times n}, \text{ with } \sigma_1 \geq \dots \geq \sigma_n \geq 0 ,$$

such that $S = U \Sigma V^T$. The values $\sigma_1, \dots, \sigma_n$ are called the singular values of S . The largest singular value of S is denoted as $\sigma_1(S)$, and the least singular value is $\sigma_n(S)$.

Some Known Results.

Lemma 2.1 ([1, Lemma 3]). For any $c \geq 1/\sqrt{2\pi}$, rank- n lattice \mathcal{L} , $0 < \varepsilon < 1$, and matrix S s.t. $\sigma_n(S) \geq \eta_\varepsilon(\mathcal{L})$, we have

$$\Pr_{\mathbf{v} \leftarrow \mathcal{D}_{\mathcal{L},S}} (\|\mathbf{v}\| \geq \sigma_1(S) c \sqrt{n}) \leq \frac{1+\varepsilon}{1-\varepsilon} (c\sqrt{2\pi e} \cdot e^{-\pi c^2})^n .$$

Lemma 2.2 ([1, Lemma 4]). For any rank- m lattice \mathcal{L} , $0 < \varepsilon < 1$, vector $\mathbf{c} \in \mathbb{R}^m$, and full-rank matrix $S \in \mathbb{R}^{m \times m}$, such that $\sigma_m(S) \geq \eta_\varepsilon(\mathcal{L})$, we have

$$\rho_S(\mathcal{L} + \mathbf{c}) \in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1 \right] \cdot \rho_S(\mathcal{L}) .$$

¹The notation $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ means that $\Sigma_{ij} = \sigma_i$ if $i = j$, and 0 otherwise.

Lemma 2.3 ([1, Corollary 2]). For any full-rank lattice $\mathcal{L} \in \mathbb{R}^n$, $\varepsilon > 0$, $c > 2$, and a rank- n matrix S such that $\sigma_n(S) \geq (1+c)\eta_\varepsilon(\mathcal{L})$, the following holds. For any $T \subset \mathcal{L}$, and any $\mathbf{v} \in \mathcal{L}$,

$$\mathcal{D}_{\mathcal{L},S}(T) - \mathcal{D}_{\mathcal{L},S}(T - \mathbf{v}) \leq \frac{\operatorname{erf}(q/2 + 2q/c)}{\operatorname{erf}(2q)} \cdot \frac{1 + \varepsilon}{1 - \varepsilon},$$

where $q = \frac{\|\mathbf{v}\|\sqrt{\pi}}{\sigma_n(S)}$, and erf is the Gaussian error function defined as $\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$.

Lemma 2.4 ([2, Azuma's inequality, Chapter 7]). Let X_0, X_1, \dots be a possibly infinite sequence of random variables that form a discrete-time sub-martingale, i.e., for all $n \geq 0$,

$$E[X_{n+1} | X_1, \dots, X_n] \geq X_n.$$

If for all $n \geq 0$, $\Pr(|X_n - X_{n-1}| \leq c) = 1$, then for any integer N and positive real t ,

$$\Pr(X_N - X_0 \leq -t) \leq \exp\left(\frac{-t^2}{2Nc^2}\right).$$

3 Sufficient Condition for X

Let $X \in \mathbb{Z}^{n \times m}$ be a full row-rank matrix with $m \geq n$. For a full-rank matrix $R \in \mathbb{R}^{m \times m}$, and $\mathbf{c} \in \mathbb{R}^m$, we consider the distribution $\mathcal{E}_{X,R,\mathbf{c}}$ which is a generalization of $\mathcal{E}_{X,r}$ defined in Section 1.

$$\mathcal{E}_{X,R,\mathbf{c}} := \{X \cdot \mathbf{v} : \mathbf{v} \leftarrow \mathcal{D}_{\mathbb{Z}^m + \mathbf{c}, R}\}.$$

Setting $R = rI_m$, and $\mathbf{c} = \mathbf{0}$, we obtain the distribution $\mathcal{E}_{X,r}$. In this section, we establish a sufficient condition on X such that $\mathcal{E}_{X,R,\mathbf{c}}$ is statistically close to a discrete Gaussian distribution, namely, to $\mathcal{D}_{\mathbb{Z}^n + X\mathbf{c}, RX^T}$.

Definition 3.1. Let X be an $n \times m$ matrix and let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{Z}^m$ be the rows of X . For any positive $q_1 = q_1(n), q_2 = q_2(n)$, the matrix X is said to have quality (q_1, q_2) if all column-vectors of X have ℓ_2 norm at most q_1 and there exist pairwise orthogonal $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{Z}^m$, such that for all $i, j \in [n]$, $\mathbf{u}_i \cdot \mathbf{x}_j = \delta_{ij}$, and $\|\mathbf{u}_i\|_2 \leq q_2$.

Note that the fact that X has quality (q_1, q_2) for $q_1, q_2 \in \mathbb{R}^+$ implies that X is a full row-rank matrix, and that $X\mathbb{Z}^m = \mathbb{Z}^n$. In this section, we will prove the following theorem.

Theorem 3.2. For any $m > n \geq 1$, and $0 < \varepsilon < 1/3$, $\mathbf{c} \in \mathbb{R}^m$, if X is an $n \times m$ matrix of quality (q_1, q_2) , and $R \in \mathbb{R}^{m \times m}$ is full-rank such that

$$\sigma_m(R) \geq (1 + q_1 q_2) \cdot \sqrt{\frac{\ln(2(m-n)(1+1/\varepsilon))}{\pi}},$$

then

$$\Delta(\mathcal{E}_{X,R,\mathbf{c}}, \mathcal{D}_{\mathbb{Z}^n + X\mathbf{c}, RX^T}) \leq 2\varepsilon.$$

The proof of Theorem 3.2 follows from the sequence of lemmas below. Define a lattice $A = A(X)$ as the lattice containing integer vectors in \mathbb{Z}^m orthogonal to all the row vectors of X ,

$$A = A(X) := \{\mathbf{v} \in \mathbb{Z}^m : X \cdot \mathbf{v} = \mathbf{0}\}. \quad (3.1)$$

Note that the rank of A is $m - n$ since X has rank n and the vectors in A span the vector space orthogonal to the n row vectors of X . The following result is a slight extension of [1, Lemma 10] (mentioned already in a previous version of [9]). It shows (in the case of centered spherical distribution) that if r is bigger than the smoothing parameter of A , then $\mathcal{E}_{X,r}$ is statistically indistinguishable from $\mathcal{D}_{\mathbb{Z}^n, rX^T}$.

Lemma 3.3. *Let X be an $n \times m$ full row-rank matrix and let $A = A(X)$ be as defined in (3.1). For any $\varepsilon < 1/3$, for any full-rank matrix $R \in \mathbb{R}^{m \times m}$ such that $\sigma_m(R) > \eta_\varepsilon(A)$, and for any $\mathbf{c} \in \mathbb{R}^m$,*

$$\Delta(\mathcal{E}_{X,R,\mathbf{c}}, \mathcal{D}_{X\mathbb{Z}^m + X\mathbf{c}, RX^T}) \leq 2\varepsilon .$$

Proof. The support of $\mathcal{E}_{X,R,\mathbf{c}}$ is $X\mathbb{Z}^m + X\mathbf{c}$. Fix some $\mathbf{z} \in X\mathbb{Z}^m + X\mathbf{c}$. The probability mass assigned to \mathbf{z} by $\mathcal{E}_{X,R,\mathbf{c}}$ is proportional to $\rho_R(A_{\mathbf{z}}) = \rho(R^{-T}A_{\mathbf{z}})$, where

$$A_{\mathbf{z}} = A_{\mathbf{z}}(X) := \{\mathbf{v} \in \mathbb{Z}^m + \mathbf{c} : X \cdot \mathbf{v} = \mathbf{z}\} .$$

Note that $A_{\mathbf{z}} = A + \mathbf{w}_{\mathbf{z}}$, for any arbitrary element $\mathbf{w}_{\mathbf{z}} \in A_{\mathbf{z}}$ (since $A_{\mathbf{z}}$ is non-empty).

Consider now the quantity $\rho(R^{-T}A_{\mathbf{z}})$. Let $P = X \cdot R^T \in \mathbb{R}^{n \times m}$, and note that $\ker P = R^{-T} \ker X$. Let $Q = (PP^T)^{-1}P$ be the pseudo-inverse of P (i.e., $PQ^T = I_n$ and the rows of Q span the same linear subspace of \mathbb{R}^m as the rows of P).

Define $\mathbf{u}_{\mathbf{z}} := Q^T \mathbf{z}$. Note that $\mathbf{u}_{\mathbf{z}}$ is the point in the affine subspace $\ker P + R^{-T} \mathbf{w}_{\mathbf{z}}$ closest to the origin: it is clearly in this affine subspace, and moreover, being in the row space of Q , it must be orthogonal to $\ker Q = \ker P$. Therefore, for any $\mathbf{r} \in R^{-T}A_{\mathbf{z}}$, we have that $\mathbf{r} - \mathbf{u}_{\mathbf{z}}$ is orthogonal to $\mathbf{u}_{\mathbf{z}}$, and thus, $\rho(\mathbf{r}) = \rho(\mathbf{u}_{\mathbf{z}}) \cdot \rho(\mathbf{r} - \mathbf{u}_{\mathbf{z}})$. Hence,

$$\rho(R^{-T}A_{\mathbf{z}}) = \rho(\mathbf{u}_{\mathbf{z}}) \cdot \rho(R^{-T}A_{\mathbf{z}} - \mathbf{u}_{\mathbf{z}}) .$$

Since $\mathbf{u}_{\mathbf{z}} \in R^{-T} \ker X + R^{-T} \mathbf{w}_{\mathbf{z}}$, it follows that

$$R^{-T}A_{\mathbf{z}} - \mathbf{u}_{\mathbf{z}} = R^{-T}(A - \mathbf{c}') ,$$

for some \mathbf{c}' in the span of the lattice A . Thus, using Lemma 2.2, we have that

$$\begin{aligned} \rho(R^{-T}A_{\mathbf{z}}) &= \rho(\mathbf{u}_{\mathbf{z}}) \cdot \rho_R(A - \mathbf{c}') \\ &\in \left[\frac{1-\varepsilon}{1+\varepsilon}, 1 \right] \cdot \rho_R(A) \cdot \rho(\mathbf{u}_{\mathbf{z}}) \\ &= \left[\frac{1-\varepsilon}{1+\varepsilon}, 1 \right] \cdot \rho_R(A) \cdot \rho(P^T(PP^T)^{-1}\mathbf{z}) = \left[\frac{1-\varepsilon}{1+\varepsilon}, 1 \right] \cdot \rho_R(A) \cdot \rho_{P^T}(\mathbf{z}) . \end{aligned}$$

This implies that the statistical distance between $\mathcal{E}_{X,R,\mathbf{c}}$ and $\mathcal{D}_{X\mathbb{Z}^m + X\mathbf{c}, RX^T}$ is at most $1 - \frac{1-\varepsilon}{1+\varepsilon} \leq 2\varepsilon$. \square

Thus, in order to bound the required statistical distance, we need to bound the smoothing parameter $\eta_\varepsilon(A)$, for which we use the following result.

Lemma 3.4 ([11, Lemma 3.3]). *For any rank- n lattice \mathcal{L} and any $\varepsilon > 0$,*

$$\eta_\varepsilon(\mathcal{L}) \leq \lambda_n(\mathcal{L}) \cdot \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}} .$$

In particular, $\eta_\varepsilon(\mathbb{Z}^n) \leq \sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}}$.

Therefore, we now only need to bound $\lambda_{m-n}(A)$, and then using Lemma 3.4, we get a bound on $\eta_\varepsilon(A)$. For this, we use the following result from [7], which we state and prove in our notation.

Lemma 3.5. *Let $X \in \mathbb{Z}^{n \times m}$ and let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{Z}^m$ be the rows of X . If X has quality (q_1, q_2) , then there exist linearly independent $\mathbf{v}_1, \dots, \mathbf{v}_{m-n} \in \mathbb{Z}^m$ such that $\mathbf{v}_k \cdot \mathbf{x}_j = 0$ and $\|\mathbf{v}_k\|_2 \leq 1 + q_1 q_2$ for all j, k . In particular $\lambda_{m-n}(A) \leq 1 + q_1 q_2$.*

Proof. Let $\mathbf{u}_1, \dots, \mathbf{u}_n$ be as in Definition 3.1. Define vectors $\mathbf{v}_1, \dots, \mathbf{v}_m$ as:

$$\mathbf{v}_k = \mathbf{e}_k - \sum_{i=1}^n x_{ik} \mathbf{u}_i,$$

where x_{ik} is the k -th coordinate of \mathbf{x}_i , and \mathbf{e}_k 's are standard unit vectors in m dimensions. Thus,

$$\mathbf{v}_k \cdot \mathbf{x}_j = \mathbf{e}_k \cdot \mathbf{x}_j - \sum_{i=1}^n x_{ik} (\mathbf{u}_i \cdot \mathbf{x}_j) = x_{jk} - x_{jk} = 0,$$

using the fact that $\mathbf{u}_i \cdot \mathbf{x}_j = \delta_{ij}$.

Using the fact that \mathbf{u}_i 's are orthogonal and the triangle inequality, we have that

$$\|\mathbf{v}_k\| \leq 1 + \sqrt{\sum_{i=1}^n |x_{ik}|^2 \|\mathbf{u}_i\|^2} \leq 1 + q_2 \sqrt{\sum_{i=1}^n |x_{ik}|^2} \leq 1 + q_1 q_2.$$

Also, clearly there are at least $m - n$ linearly independent vectors in $\mathbf{v}_1, \dots, \mathbf{v}_m$, since, by definition, the set of vectors $\{\mathbf{v}_1, \dots, \mathbf{v}_m, \mathbf{u}_1, \dots, \mathbf{u}_n\}$ together generate all the unit vectors \mathbf{e}_k , and hence span \mathbb{R}^m . \square

Combining Lemmata 3.3, 3.4, and 3.5 immediately leads to a proof of Theorem 3.2.

We note that one can relax Definition 3.1 by not requiring the pairwise orthogonality of the vectors $\mathbf{u}_1, \dots, \mathbf{u}_n$. In this case, everything works just as before except the last inequality which is replaced by

$$\|\mathbf{v}_k\| \leq 1 + \sum_{i=1}^n |x_{ik}| \|\mathbf{u}_i\| \leq 1 + q_2 \sum_{i=1}^n |x_{ik}|.$$

Thus, Theorem 3.2 still holds if we take q_1 to be a bound on the ℓ_1 norm of the columns of X (instead of the ℓ_2 norm as it is now).

4 A random matrix X is of good quality

In this section, we show that $X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, S})^m$ has quality (q_1, q_2) for “small” q_1, q_2 with overwhelming probability. For proving this result, we need the following claim that easily follows from the pigeon-hole principle.

Claim 4.1. *Let B, n be positive integers and let $\ell = \ell(B, n)$ be an integer such that $(2B\ell + 1)^n < 2^\ell$. Also let $\mathbf{x}_1, \dots, \mathbf{x}_\ell \in \mathbb{Z}^n$ be such that $\|\mathbf{x}_j\|_\infty \leq B$ for all j . Then there exist $\alpha_1, \dots, \alpha_\ell \in \{-1, 0, 1\}$, not all zero, such that $\sum_{j=1}^\ell \alpha_j \mathbf{x}_j = \mathbf{0}$. In particular, if $Bn \geq 100$ then this holds for $\ell = \lfloor 2n \log(Bn) \rfloor$.*

Proof. Any 0/1 combination of $\mathbf{x}_1, \dots, \mathbf{x}_\ell$ is in $\{-B\ell, \dots, -1, 0, 1, \dots, B\ell\}^n$, and thus the total number of distinct resulting vectors is at most $(2B\ell + 1)^n < 2^\ell$. By the pigeon-hole principle there exist two distinct 0/1 combinations that result in the same vector, which implies the result by taking their difference.

If $\ell = \lfloor 2n \log(Bn) \rfloor$, then

$$2^\ell > \frac{(Bn)^{2n}}{2} > ((Bn)^2/2)^n,$$

and

$$(2B\ell + 1)^n \leq (4Bn \log(Bn) + 1)^n < (5Bn \log(Bn))^n.$$

Since $Bn > 10 \log(Bn)$ holds for $Bn \geq 100$, we have that $(2B\ell + 1)^n < 2^\ell$. \square

Lemma 4.2. *Let $n \geq 100$ be an integer, and let $\varepsilon = \varepsilon(n) \in (0, 1/1000)$. Let $S \in \mathbb{R}^{n \times n}$ be such that $\sigma_n(S) \geq 9\sqrt{\frac{\ln(2n(1+\varepsilon))}{\pi}}$. Denote $\sigma_1 = \sigma_1(S)$ and $\sigma_n = \sigma_n(S)$. Let m be an integer such that $m > 30n \log(\sigma_1 n)$ and let $X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, S})^m$ be an $n \times m$ matrix. Then, with probability $1 - 2^{-n}$, X has quality $(\sigma_1 \sqrt{n \log m}, 2\sqrt{30n \log(\sigma_1 n)})$.*

Proof. By Lemma 2.1, with $c = \sqrt{\log m}$, we have using the union bound that with probability $1 - 2^{-n-1}$, all column-vectors of X have norm at most $\sigma_1 \sqrt{n \log m}$. It is sufficient to show that with probability $1 - 2^{-n-1}$, there exist pairwise orthogonal $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{Z}^m$ such that $\|\mathbf{u}_i\|_2 \leq 2\sqrt{30n \log(\sigma_1 n)}$, and $X \cdot \mathbf{u}_i = \mathbf{e}_i$ for $1 \leq i \leq n$.

Denote $t := 3n \log(\sigma_1 n)$. For each $i = 1, \dots, n$, we will show that with probability at least $1 - 2^{-t/10}$, there exists a \mathbf{u}_i with the desired properties; the result will then follow by the union bound. Let $\mathbf{x}_1, \dots, \mathbf{x}_m$ be the columns of X . Define sets $\mathcal{S}_j \subset \mathbb{Z}^n$ for $1 \leq j \leq m$ as follows:

$$\forall j \in [m], \mathcal{S}_j = \left\{ \sum_{k=1}^j \alpha_k \mathbf{x}_k : \alpha_k \in \{-1, 0, 1\} \right\}.$$

From now on we show the existence of \mathbf{u}_1 and at the end of the proof we explain how to show the existence of the other \mathbf{u}_i . We will show that with probability $1 - 2^{-t/10}$, $\mathcal{S}_{\lfloor 10t \rfloor} \cap (\mathcal{S}_{\lfloor 10t \rfloor} + \mathbf{e}_1)$ is non-empty, which implies that \mathbf{e}_1 can be obtained as a $\{-2, -1, 0, 1, 2\}$ combination of the first $\lfloor 10t \rfloor$ vectors. Then we are done since the vector \mathbf{u}_1 can be defined as this coefficient vector with first $\lfloor 10t \rfloor$ coordinates in $\{-2, -1, 0, 1, 2\}$ and the remaining coordinates zero. It is easy to see that $\|\mathbf{u}_1\|$ is at most $2\sqrt{10t}$.

The idea used to prove this is that if this does not happen, then \mathcal{S}_j and $\mathcal{S}_j + \mathbf{e}_1$ must be disjoint for all $j \leq \lfloor 10t \rfloor$, and so at most one of \mathbf{x}_{j+1} and $\mathbf{x}_{j+1} - \mathbf{e}_1$ is in \mathcal{S}_j for any $j \leq \lfloor 10t \rfloor$. Using the randomness of \mathbf{x}_j , we conclude that with constant probability, $\mathbf{x}_{j+1} \notin \mathcal{S}_j$ and this cannot happen for many values of j by Claim 4.1.

Define binary random variables Y_1, \dots, Y_m as follows: $Y_j = 1$ if and only if either

- $\mathcal{S}_j \cap (\mathcal{S}_j + \mathbf{e}_1)$ is non-empty, or
- $\mathbf{x}_{j+1} \notin \mathcal{S}_j$ and $\|\mathbf{x}_{j+1}\|_\infty \leq \sigma_1 \sqrt{n}$.

Consider the following claims about the distribution of Y_1, \dots, Y_m .

Claim 4.3. *For any values of y_1, \dots, y_{j-1} , and for all j , $\Pr(Y_j = 1 \mid Y_1 = y_1, \dots, Y_{j-1} = y_{j-1}) \geq 0.3$.*

Proof. Throughout this proof, we condition on arbitrary values of $\mathbf{x}_1, \dots, \mathbf{x}_j$, and our goal is to show that the probability of $Y_j = 1$ is at least 0.3.

If $\mathcal{S}_j \cap (\mathcal{S}_j + \mathbf{e}_1)$ is non-empty (an event that depends only on $\mathbf{x}_1, \dots, \mathbf{x}_j$), then we are done. Assume therefore that our choice of $\mathbf{x}_1, \dots, \mathbf{x}_j$ is such that \mathcal{S}_j and $\mathcal{S}_j + \mathbf{e}_1$ are disjoint. By Lemma 2.1, with all but exponentially small probability in n , $\|\mathbf{x}_{j+1}\|_\infty \leq \|\mathbf{x}_{j+1}\|_2 \leq \sigma_1 \sqrt{n}$. It therefore suffices to show that $\mathbf{x}_{j+1} \in \mathcal{S}_j$ with probability at most, say, 0.695. This follows since from Lemma 3.4, we get that $\sigma_n \geq 9\eta_\varepsilon(\mathbb{Z}^n)$, and thus we have by Lemma 2.3 that

$$\Pr(\mathbf{x}_{j+1} \in \mathcal{S}_j) - \Pr(\mathbf{x}_{j+1} \in \mathcal{S}_j + \mathbf{e}_1) \leq \frac{\operatorname{erf}(3\sqrt{\pi}/(4\sigma_n))}{\operatorname{erf}(2\sqrt{\pi}/\sigma_n)} \cdot \frac{1 + \varepsilon}{1 - \varepsilon} < 0.39,$$

where the last inequality follows because erf is nearly linear for small arguments. Finally, the claim follows by noticing that the sum of the two probabilities in the left hand side is at most 1. \square

Claim 4.4.

$$\Pr(Y_1 + \dots + Y_{\lfloor 10t \rfloor} \geq 2t) \geq 1 - 2^{-t/10}.$$

Proof. Define random variables Z_0, Z_1, \dots recursively as $Z_0 = 0$ and for $i \geq 1$, $Z_i = Z_{i-1} + Y_i - \frac{3}{10}$, so for all j ,

$$Y_1 + \dots + Y_j = Z_j - Z_0 + \frac{3j}{10}.$$

Clearly Z_0, Z_1, \dots is a sub-martingale, and for all $j \geq 1$, $|Z_j - Z_{j-1}| \leq \frac{7}{10}$. Thus, using Lemma 2.4, we get that

$$\Pr(Y_1 + \dots + Y_{\lfloor 10t \rfloor} \leq 2t) = \Pr(Z_{\lfloor 10t \rfloor} - Z_0 \leq -t) \leq \exp\left(\frac{-t^2}{2\lfloor 10t \rfloor(0.7)^2}\right) \leq 2^{-t/10},$$

which implies the result. \square

Claim 4.5. *If $Y_1 + \dots + Y_{\lfloor 10t \rfloor} \geq 2t$, then $\mathcal{S}_{\lfloor 10t \rfloor} \cap (\mathcal{S}_{\lfloor 10t \rfloor} + \mathbf{e}_1)$ is non-empty.*

Proof. Assume $Y_{j_1} = \dots = Y_{j_{\lfloor 2t \rfloor}} = 1$ and $\mathcal{S}_{\lfloor 10t \rfloor} \cap (\mathcal{S}_{\lfloor 10t \rfloor} + \mathbf{e}_1)$ is empty. Consider the vectors $\mathbf{x}_{j_1+1}, \dots, \mathbf{x}_{j_{\lfloor 2t \rfloor}+1}$. We have that $\|\mathbf{x}_{j_k+1}\|_\infty \leq \sigma_1 \sqrt{n}$ for $1 \leq k \leq \lfloor 2t \rfloor$. Since $2t = 6n \log(\sigma_1 n) \geq 2n \log(\sigma_1 n^{3/2})$, by Claim 4.1, there exist $\alpha_1, \dots, \alpha_{\lfloor 2t \rfloor} \in \{-1, 0, 1\}$, not all zero, such that

$$\sum_{k=1}^{\lfloor 2t \rfloor} \alpha_k \mathbf{x}_{j_k+1} = \mathbf{0}.$$

Let k be the largest index such that $\alpha_k \neq 0$. This implies $\mathbf{x}_{j_k+1} \in \mathcal{S}_{j_k}$, contradicting the fact that $Y_{j_k} = 1$. \square

Claim 4.3, 4.4, and 4.5 together prove the existence of \mathbf{u}_1 . To obtain a vector \mathbf{u}_2 that is orthogonal to \mathbf{u}_1 , we do the following. Let X' be the $(n+1) \times m$ matrix formed by adding another row, \mathbf{u}_1 , to X . We follow the same argument as above with X replaced by X' and \mathbf{e}_1 replaced by \mathbf{e}_2 . Let $\mathbf{x}'_j \in \mathbb{Z}^{n+1}$ be \mathbf{x}_j concatenated with u_{1j} , the j -th coordinate of \mathbf{u}_1 . By construction $\|\mathbf{u}_1\|_\infty \leq 2$, and hence $\|\mathbf{x}'_j\|_\infty \leq \sigma_1 \sqrt{n}$ if and only if $\|\mathbf{x}_j\|_\infty \leq \sigma_1 \sqrt{n}$. For completing the proof, we need to bound $\Pr(\mathbf{x}'_{j+1} \in \mathcal{S}') - \Pr(\mathbf{x}'_{j+1} \in \mathcal{S}' + \mathbf{e}_2)$ for an arbitrary set $\mathcal{S}' \subset \mathbb{Z}^{n+1}$. Since the $(n+1)$ -th coordinate \mathbf{x}'_{j+1} is fixed to be $u_{1(j+1)}$, we have that

$$\Pr(\mathbf{x}'_{j+1} \in \mathcal{S}') - \Pr(\mathbf{x}'_{j+1} \in \mathcal{S}' + \mathbf{e}_2) = \Pr(\mathbf{x}_{j+1} \in \mathcal{S}) - \Pr(\mathbf{x}_{j+1} \in \mathcal{S} + \mathbf{e}_2),$$

where $\mathcal{S} \subset \mathbb{Z}^n$ is the set of vectors formed by projecting the vectors of $\mathcal{S}'' = \{\mathbf{y} \in \mathcal{S}' : y_{n+1} = u_{1(j+1)}\}$ onto the first n coordinates. The remainder of the proof is exactly the same. We similarly get $\mathbf{u}_3, \dots, \mathbf{u}_n$, such that for all $i \leq n$, \mathbf{u}_i is orthogonal to $\mathbf{u}_1, \dots, \mathbf{u}_{i-1}$. \square

5 Sum of Discrete Gaussian Distribution

In this section, we state the main result of this paper, the proof of which follows easily from the results proved in Section 3 and Section 4.

Theorem 5.1. *Let $m > n \geq 100$ be integers, and let $\varepsilon = \varepsilon(n) \in (0, 1/1000)$. Let $R \in \mathbb{R}^{m \times m}$ be a full-rank matrix, and let $\mathbf{c} \in \mathbb{R}^m$. Let $S \in \mathbb{R}^{n \times n}$ be a full-rank matrix, and let $X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, S})^m$. If $m \geq 30n \log(\sigma_1(S)n)$, $\sigma_m(R) \geq 10n\sigma_1(S) \log m \sqrt{\log(1/\varepsilon) \log(n\sigma_1(S))}$, and $\sigma_n(S) \geq 9\sqrt{\frac{\ln(2n(1+1/\varepsilon))}{\pi}}$, then, with probability $1 - 2^{-n}$ over the choice of X , we have that*

$$\Delta(\mathcal{E}_{X, R, \mathbf{c}}, \mathcal{D}_{\mathbb{Z}^n + X\mathbf{c}, RX^T}) \leq 2\varepsilon.$$

Proof. Theorem 5.1 follows easily from Lemma 4.2 and Theorem 3.2 using the observation that

$$\begin{aligned}\sigma_m(R) &\geq 10n\sigma_1(S)\log m\sqrt{\log(1/\varepsilon)\log(n\sigma_1(S))} \\ &\geq \left(1+2\sigma_1(S)\sqrt{n\log m}\sqrt{30n\log(\sigma_1(S)n)}\right)\cdot\sqrt{\frac{\ln(2(m-n)(1+1/\varepsilon))}{\pi}}.\end{aligned}$$

□

It is easy to see that if we sample a vector x from $D_{\mathbb{Z}^n,S}$ then Bx is distributed like $D_{\mathcal{L},SB^T}$ where $\mathcal{L} = \mathcal{L}(B)$ is the lattice generated by the columns of B (see, e.g., [1, Fact 2]). As a result, we can extend Theorem 5.1 to the case in which X is sampled from $(D_{\mathcal{L},M})^m$ where $\mathcal{L} = \mathcal{L}(B)$ for some basis B , and $M = SB^T$. The only change is that the statistical closeness is to the distribution $D_{\mathcal{L}+X\mathbf{e},RX^T}$.

Acknowledgements: We thank Damien Stehlé for useful comments.

References

- [1] SHWETA AGRAWAL, CRAIG GENTRY, SHAI HALEVI, AND AMIT SAHAI: Discrete Gaussian leftover hash lemma over infinite domains. In *ASIACRYPT*, pp. 97–116, 2013. [1](#), [2](#), [3](#), [4](#), [5](#), [10](#)
- [2] NOGA ALON AND JOEL SPENCER: *The Probabilistic Method*. Wiley, 2004. [5](#)
- [3] WOJCIECH BANASZCZYK: New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(01):625–635, 1993. [3](#)
- [4] DAN BONEH AND DAVID FREEMAN: Homomorphic signatures for polynomial functions. In *Advances in Cryptology-EUROCRYPT*, pp. 149–168. Springer, 2011. [1](#)
- [5] SANJAM GARG, CRAIG GENTRY, AND SHAI HALEVI: Candidate multilinear maps from ideal lattices. In *Advances in Cryptology-EUROCRYPT*, pp. 1–17. Springer, 2013. [2](#), [3](#)
- [6] CRAIG GENTRY, CHRIS PEIKERT, AND VINOD VAIKUNTANATHAN: Trapdoors for hard lattices and new cryptographic constructions. In *40th Symposium on Theory of Computing-STOC*, pp. 197–206. ACM, 2008. [1](#)
- [7] GREG KUPERBERG, SHACHAR LOVETT, AND RON PELED: Probabilistic existence of rigid combinatorial structures. In *44th Symposium on Theory of Computing-STOC*, pp. 1091–1106. ACM, 2012. [2](#), [6](#)
- [8] ADELIN LANGLOIS, DAMIEN STEHLÉ, AND RON STEINFELD: GGHLite: More efficient multilinear maps from ideal lattices. In *Advances in Cryptology-EUROCRYPT 2014*, pp. 239–256. Springer, 2014. [3](#)
- [9] SAN LING, DUONG HIEU PHAN, DAMIEN STEHLÉ, AND RON STEINFELD: Hardness of k -LWE and applications in traitor tracing. In *Advances in Cryptology-CRYPTO 2014*, pp. 315–334. Springer, 2014. [2](#), [3](#), [5](#)
- [10] DANIELE MICCIANCIO AND CHRIS PEIKERT: Hardness of SIS and LWE with small parameters. In *Advances in Cryptology-CRYPTO*. Springer, 2013. [1](#), [3](#)

- [11] DANIELE MICCIANCIO AND ODED REGEV: Worst-case to average-case reductions based on Gaussian measures. *SIAM Journal on Computing*, 37(01):267–302, 2007. 1, 6
- [12] CHRIS PEIKERT: An efficient and parallel Gaussian sampler for lattices. In *Advances in Cryptology-CRYPTO*, pp. 80–97. Springer, 2010. 1, 3

AUTHORS

Divesh Aggarwal
Senior Researcher
Department of Computer and Communication Sciences
EPFL, Lausanne, Switzerland
dives.h.aggarwal@epfl.ch
<https://sites.google.com/site/diveshhomepage/>

Oded Regev
Professor
Courant Institute of Mathematical Sciences
New York University, New York, USA
<http://www.cims.nyu.edu/~regev/>