

Discrete Gaussian Leftover Hash Lemma over Infinite Domains

Shweta Agrawal¹, Craig Gentry², Shai Halevi², and Amit Sahai¹

¹ UCLA

² IBM Research

Abstract. The classic Leftover Hash Lemma (LHL) is often used to argue that certain distributions arising from modular subset-sums are close to uniform over their finite domain. Though very powerful, the applicability of the leftover hash lemma to lattice based cryptography is limited for two reasons. First, typically the distributions we care about in lattice-based cryptography are *discrete Gaussians*, not uniform. Second, the elements chosen from these discrete Gaussian distributions lie in an infinite domain: a lattice rather than a finite field.

In this work we prove a “lattice world” analog of LHL over infinite domains, proving that certain “generalized subset sum” distributions are statistically close to well behaved discrete Gaussian distributions, even without any modular reduction. Specifically, given many vectors $\{\mathbf{x}_i\}_{i=1}^m$ from some lattice $L \subset \mathbb{R}^n$, we analyze the probability distribution $\sum_{i=1}^m z_i \mathbf{x}_i$ where the integer vector $\mathbf{z} \in \mathbb{Z}^m$ is chosen from a discrete Gaussian distribution. We show that when the \mathbf{x}_i ’s are “random enough” and the Gaussian from which the \mathbf{z} ’s are chosen is “wide enough”, then the resulting distribution is statistically close to a near-spherical discrete Gaussian over the lattice L . Beyond being interesting in its own right, this “lattice-world” analog of LHL has applications for the new construction of multilinear maps [5], where it is used to sample Discrete Gaussians obliviously. Specifically, given encoding of the \mathbf{x}_i ’s, it is used to produce an encoding of a near-spherical Gaussian distribution over the lattice. We believe that our new lemma will have other applications, and sketch some plausible ones in this work.

1 Introduction

The Leftover Hash Lemma (LHL) is a central tool in computer science, stating that universal hash functions are good randomness extractors. In a characteristic application, the universal hash function may often be instantiated by a simple inner product function, where it is used to argue that a random linear combination of some elements (that are chosen at random and then fixed “once and for all”) is statistically close to the uniform distribution over some finite domain. Though extremely useful and powerful in general, the applicability of the leftover hash lemma to lattice based cryptography is limited for two reasons. First, typically the distributions we care about in lattice-based cryptography are *discrete Gaussians*, not uniform. Second, the elements chosen from these discrete

Gaussian distributions lie in an infinite domain: a lattice rather than a finite field.

The study of discrete Gaussian distributions underlies much of the advances in lattice-based cryptography over the last decade. A discrete Gaussian distribution is a distribution over some fixed lattice, in which every lattice point is sampled with probability proportional to its probability mass under a standard (n -dimensional) Gaussian distribution. Micciancio and Regev have shown in [10] that these distributions share many of the nice properties of their continuous counterparts, and demonstrated their usefulness for lattice-based cryptography. Since then, discrete Gaussian distributions have been used extensively in all aspects of lattice-based cryptography (most notably in the famous “Learning with Errors” problem and its variants [14]). Despite their utility, we still do not understand discrete Gaussian distributions as well as we do their continuous counterparts.

A Gaussian Leftover Hash Lemma for Lattices?

The LHL has been applied often in lattice-based cryptography, but sometimes awkwardly. As an example, in the integer-based fully homomorphic encryption scheme of van Dijk et al. [18], ciphertexts live in the lattice \mathbb{Z} . Roughly speaking, the public key of that scheme contains many encryptions of zero, and encryption is done by adding the plaintext value to a subset-sum of these encryptions of zero. To prove security of this encryption method, van Dijk et al. apply the left-over hash lemma in this setting, but with the cost of complicating their encryption procedure by reducing the subset-sum of ciphertexts modulo a single large ciphertext, so as to bring the scheme back in to the realm of finite rings where the leftover hash lemma is naturally applied.³ It is natural to ask whether that scheme remains secure also without this artificial modular reduction, and more generally whether there is a more direct way to apply the LHL in settings with infinite rings.

As another example, in the recent construction of multilinear maps [5], Garg et. al. require a procedure to randomize “encodings” to break simple algebraic relations that exist between them. One natural way to achieve this randomization is by adding many random encodings of zero to the public parameters, and adding a random linear combination of these to re-randomize a given encoding (without changing the encoded value). However, in their setting, there is no way to “reduce” the encodings so that the LHL can be applied. Can they argue that the new randomized encoding yields an element from some well behaved distribution?

In this work we prove an analog of the leftover hash lemma over lattices, yielding a positive answers to the questions above. We use discrete Gaussian distributions as our notion of “well behaved” distributions. Then, for m vectors $\{\mathbf{x}_i\}_{i \in [m]}$ chosen “once and for all” from an n dimensional lattice $L \subset \mathbb{R}^n$,

³ Once in the realms of finite rings, one can alternatively use the generic proof of Rothblum [15], which also uses the LHL.

and a coefficient vector \mathbf{z} chosen from a discrete Gaussian distribution over the integers, we give sufficient conditions under which the distribution $\sum_{i=1}^m z_i \mathbf{x}_i$ is “well behaved.”

Oblivious Gaussian Sampler

Another application of our work is in the construction of an extremely simple *discrete Gaussian sampler* [6, 13]. Such samplers, that sample from a spherical discrete Gaussian distribution over a lattice have been constructed by [6] (using an algorithm by Klein [7]) as well as Peikert [13]. Here we consider a much simpler discrete Gaussian sampler (albeit a somewhat imperfect one). Specifically, consider the following sampler. In an offline phase, for $m > n$, the sampler samples a set of short vectors $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_m$ from L – e.g., using GPV or Peikert’s algorithm. Then, in the online phase, the sampler generates $\mathbf{z} \in \mathbb{Z}^m$ according to a discrete Gaussian and simply outputs $\sum_{i=1}^m z_i \mathbf{x}_i$. But does this simpler sampler work – i.e., can we say anything about its output distribution? Also, how small can we make the dimension m of \mathbf{z} and how small can we make the entries of \mathbf{z} ? Ideally m would be not much larger than the dimension of the lattice and the entries of \mathbf{z} have small variance – e.g., $\tilde{O}(\sqrt{n})$.

A very useful property of such a sampler is that it can be made *oblivious* to an explicit representation of the underlying lattice, which makes it applicable easily within an additively homomorphic scheme. Namely, if you are given lattice points encrypted under an additively homomorphic encryption scheme, you can use them to generate an encrypted well behaved Gaussian on the underlying lattice. Previous samplers [6, 13] are too complicated to use within an additively homomorphic encryption scheme ⁴.

Our Results

In this work, we obtain a discrete Gaussian version of the LHL over infinite rings. Formally, consider an n dimensional lattice L and (column) vectors $X = [\mathbf{x}_1 | \mathbf{x}_2 | \dots | \mathbf{x}_m] \in L$. We choose \mathbf{x}_i according to a discrete Gaussian distribution $\mathcal{D}_{L,S}$, where $\mathcal{D}_{L,S}$ is defined as $\mathcal{D}_{L,S,c}(\mathbf{x}) = \frac{\rho_{S,c}(\mathbf{x})}{\rho_{S,c}(L)}$ with $\rho_{S,c}(\mathbf{x}) \stackrel{\text{def}}{=} \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / s^2)$ and $\rho_{S,c}(A)$ for set A denotes $\sum_{\mathbf{x} \in A} \rho_{S,c}(\mathbf{x})$.

Let $\mathbf{z} \leftarrow \mathcal{D}_{\mathbb{Z}^m, s'}$, we analyze the conditions under which the vector $X \cdot \mathbf{z}$ is statistically close to a “near-spherical” discrete Gaussian. Formally, consider:

$$\mathcal{E}_{X, s'} \stackrel{\text{def}}{=} \{X \cdot \mathbf{z} : \mathbf{z} \leftarrow \mathcal{D}_{\mathbb{Z}^m, s'}\}$$

Then, we prove that $\mathcal{E}_{X, s'}$ is close to a discrete Gaussian over L of moderate “width”. Specifically, we show that for large enough s' , with overwhelming probability over the choice of X :

⁴ As noted by Peikert [13], one can generate an ellipsoidal Gaussian distribution over the lattice given a basis B by just outputting $\mathbf{y} \leftarrow B \cdot \mathbf{z}$ where \mathbf{z} is a discrete Gaussian, but this ellipsoidal Gaussian distribution would typically be very skewed.

1. $\mathcal{E}_{X,s'}$ is statistically close to the ellipsoid Gaussian $\mathcal{D}_{L,s'X^\top}$, over L .
2. The singular values of the matrix X are of size roughly $s\sqrt{m}$, hence the shape of $\mathcal{D}_{L,s'X^\top}$ is “roughly spherical”. Moreover, the “width” of $\mathcal{D}_{L,s'X^\top}$ is roughly $s's\sqrt{m} = \text{poly}(n)$.

We emphasize that it is straightforward to show that the covariance matrix of $\mathcal{E}_{X,s'}$ is exactly s'^2XX^\top . However, the technical challenge lies in showing that $\mathcal{E}_{X,s'}$ is close to a discrete Gaussian for a non-square X . Also note that for a square X , the shape of the covariance matrix XX^\top will typically be very “skewed” (i.e., the least singular value of X^\top is typically much smaller than the largest singular value). We note that the “approximately spherical” nature of the output distribution is important for performance reasons in applications such as GGH: These applications must choose parameters so that the least singular value of X “drowns out” vectors of a certain size, and the resulting vectors that they draw from $\mathcal{E}_{X,s'}$ grow in size with the largest singular value of X , hence it is important that these two values be as close as possible.

Our Techniques

Our main result can be argued along the following broad outline. Our first theorem (Theorem 2) says that the distribution of $X \cdot z \leftarrow \mathcal{E}_{X,s'}$ is indeed statistically close to a discrete Gaussian over L , as long as s' exceeds the smoothing parameter of a certain “orthogonal lattice” related to X (denoted A). Next, Theorem 3 clarifies that A will have a small smoothing parameter as long as X^\top is “regularly shaped” in a certain sense. Finally, we argue in Lemma 8 that when the columns of X are chosen from a discrete Gaussian, $\mathbf{x}_i \leftarrow \mathcal{D}_{L,S}$, then X^\top is “regularly shaped,” i.e. has singular values all close to $\sigma_n(S)\sqrt{m}$.

The analysis of the smoothing parameter of the “orthogonal lattice” A is particularly challenging and requires careful analysis of a certain “dual lattice” related to A . Specifically, we proceed by first embedding A into a full rank lattice A_q and then move to study M_q – the (scaled) dual of A_q . Here we obtain a lower bound on $\lambda_{n+1}(M_q)$, i.e. the $n+1^{\text{th}}$ minima of M_q . Next, we use a theorem by Banaszczyk to convert the lower bound on $\lambda_{n+1}(M_q)$ to an upper bound on $\lambda_{m-n}(A_q)$, obtaining $m-n$ linearly independent, bounded vectors in A_q . We argue that these vectors belong to A , thus obtaining an upper bound on $\lambda_{m-n}(A)$. Relating $\lambda_{m-n}(A)$ to $\eta_\epsilon(A)$ using a lemma by Micciancio and Regev completes the analysis. (We note that probabilistic bounds on the minima and smoothing parameter A_q, M_q are well known in the case when the entries of matrix X are uniformly random mod q (e.g. [6]), but here we obtain bounds in the case when X has Gaussian entries significantly smaller than q .)

To argue that X^\top is regularly shaped, we begin with the literature of random matrices which establishes that for a matrix $H \in \mathbb{R}^{m \times n}$, where each entry of H is distributed as $\mathcal{N}(0, s^2)$ and m is sufficiently greater than n , the singular values of H are all of size roughly $s\sqrt{m}$. We extend this result to discrete Gaussians – showing that as long as each vector $\mathbf{x}_i \leftarrow \mathcal{D}_{L,S}$ where S is “not too small” and “not too skewed”, then with high probability the singular values of X^\top are all of size roughly $s\sqrt{m}$.

Related Work

Properties of linear combinations of discrete Gaussians have been studied before in some cases by Peikert [13] as well as more recently by Boneh and Freeman [3]. Peikert’s “convolution lemma” (Theorem 3.1 in [13]) analyzes certain cases in which a linear combination of discrete Gaussians yields a discrete Gaussian, in the one dimensional case. More recently, Boneh and Freeman [3] observed that under certain conditions, a linear combination of discrete Gaussians over a lattice is also a discrete Gaussian. However, the deviation of the Gaussian needed to achieve this are quite large. Related questions were considered by Lyubashevsky [9] where he computes the expectation of the inner product of discrete Gaussians.

Discrete Gaussian samplers have been studied by [6] (who use an algorithm by [7]) and [13]. These works describe a discrete Gaussian sampling algorithm that takes as input a ‘high quality’ basis B for an n dimensional lattice L and output a sample from $\mathcal{D}_{L,s,c}$. In [6], $s \geq \|\tilde{B}\| \cdot \omega(\sqrt{\log n})$, and $\tilde{B} = \max_i \|\tilde{b}_i\|$ is the Gram Schmidt orthogonalization of B . In contrast, the algorithm of [13] requires $s \geq \sigma_1(B)$, i.e. the largest singular value of B , but is fully parallelizable. Both these samplers take as input an explicit description of a “high quality basis” of the relevant lattice, and the quality of their output distribution is related to the quality of the input basis.

Peikert’s sampler [13] is elegant and its complexity is difficult to beat: the only online computation is to compute $\mathbf{c} - B_1 \lfloor B_1^{-1}(\mathbf{c} - \mathbf{x}_2) \rfloor$, where \mathbf{c} is the center of the Gaussian, B_1 is the sampler’s basis for its lattice L , and \mathbf{x}_2 is a vector that is generated in an offline phase (freshly for each sampling) in a way designed to “cancel” the covariance of B_1 so as to induce a purely spherical Gaussian. However, since our sampler just directly takes an integer linear combination of lattice vectors, and does not require extra precision for handling the inverse B_1^{-1} , it might outperform Peikert’s in some situations, at least when $\mathbf{c} = 0$.

2 Preliminaries

We say that a function $f : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is negligible (and write $f(\lambda) < \text{negl}(\lambda)$) if for every d we have $f(\lambda) < 1/\lambda^d$ for sufficiently large λ . For two distributions \mathcal{D}_1 and \mathcal{D}_2 over some set Ω the statistical distance $\text{SD}(\mathcal{D}_1, \mathcal{D}_2)$ is

$$\text{SD}(\mathcal{D}_1, \mathcal{D}_2) \stackrel{\text{def}}{=} \frac{1}{2} \sum_{x \in \Omega} \left| \Pr_{\mathcal{D}_1}[x] - \Pr_{\mathcal{D}_2}[x] \right|$$

Two distribution ensembles $\mathcal{D}_1(\lambda)$ and $\mathcal{D}_2(\lambda)$ are statistically close or statistically indistinguishable if $\text{SD}(\mathcal{D}_1(\lambda), \mathcal{D}_2(\lambda))$ is a negligible function of λ .

2.1 Gaussian Distributions

For any real $s > 0$ and vector $\mathbf{c} \in \mathbb{R}^n$, define the (spherical) Gaussian function on \mathbb{R}^n centered at \mathbf{c} with parameter s as $\rho_{s,c}(\mathbf{x}) = \exp(-\pi \|\mathbf{x} - \mathbf{c}\|^2 / s^2)$

for all $\mathbf{x} \in \mathbb{R}^n$. The *normal distribution* with mean μ and deviation σ , denoted $\mathcal{N}(\mu, \sigma^2)$, assigns to each real number $x \in \mathbb{R}$ the probability density $f(x) = \frac{1}{\sigma\sqrt{2\pi}} \cdot \rho_{\sigma\sqrt{2\pi}, \mu}(x)$. The n -dimensional (spherical) continuous Gaussian distribution with center \mathbf{c} and uniform deviation σ^2 , denoted $\mathcal{N}^n(\mathbf{c}, \sigma^2)$, just chooses each entry of a dimension- n vector independently from $\mathcal{N}(c_i, \sigma^2)$.

The n -dimensional spherical Gaussian function generalizes naturally to ellipsoid Gaussians, where the different coordinates are jointly Gaussian but are neither identical nor independent. In this case we replace the single variance parameter $s^2 \in \mathbb{R}$ by the covariance matrix $\Sigma \in \mathbb{R}^{n \times n}$ (which must be positive-definite and symmetric). To maintain consistency of notations between the spherical and ellipsoid cases, below we let S be a matrix such that $S^\top \times S = \Sigma$. Such a matrix S always exists for a symmetric Σ , but it is not unique. (In fact there exist such S 'es that are not even n -by- n matrices, below we often work with such rectangular S 'es.)

For a rank- n matrix $S \in \mathbb{R}^{m \times n}$ and a vector $\mathbf{c} \in \mathbb{R}^n$, the ellipsoid Gaussian function on \mathbb{R}^n centered at \mathbf{c} with parameter S is defined by

$$\rho_{S, \mathbf{c}}(\mathbf{x}) = \exp(-\pi(\mathbf{x} - \mathbf{c})^\top (S^\top S)^{-1}(\mathbf{x} - \mathbf{c})) \quad \forall \mathbf{x} \in \mathbb{R}^n.$$

Obviously this function only depends on $\Sigma = S^\top S$ and not on the particular choice of S . It is also clear that the spherical case can be obtained by setting $S = sI_n$, with I_n the n -by- n identity matrix. Below we use the shorthand $\rho_s(\cdot)$ (or $\rho_S(\cdot)$) when the center of the distribution is $\mathbf{0}$.

2.2 Matrices and Singular Values

In this note we often use properties of rectangular (non-square) matrices. For $m \geq n$ and a rank- n matrix⁵ $X' \in \mathbb{R}^{m \times n}$, the pseudoinverse of X' is the (unique) m -by- n matrix Y' such that $X'^\top Y' = Y'^\top X' = I_n$ and the columns of Y' span the same linear space as those of X' . It is easy to see that Y' can be expressed as $Y' = X'(X'^\top X')^{-1}$ (note that $X'^\top X'$ is invertible since X' has rank n).

For a rank- n matrix $X' \in \mathbb{R}^{m \times n}$, denote $U_{X'} = \{\|X'\mathbf{u}\| : \mathbf{u} \in \mathbb{R}^n, \|\mathbf{u}\| = 1\}$. The *least singular value* of X' is then defined as $\sigma_n(X') = \inf(U'_{X'})$ and similarly the *largest singular value* of X' is $\sigma_1(X') = \sup(U'_{X'})$. Some properties of singular values that we use later in the text are stated in Fact 1.

Fact 1 For rank- n matrices $X', Y' \in \mathbb{R}^{m \times n}$ with $m \geq n$, the following holds:

1. If $X'^\top X' = Y'^\top Y'$ then X', Y' have the same singular values.
2. If Y' is the (pseudo)inverse of X' then the singular values of X', Y' are reciprocals.
3. If X' is a square matrix (i.e., $m = n$) then X', X'^\top have the same singular values.

⁵ We use the notation X' instead of X to avoid confusion later in the text where we will instantiate $X' = X^\top$

4. If $\sigma_1(Y') \leq \delta \sigma_n(X')$ for some constant $\delta < 1$, then $\sigma_1(X' + Y') \in [1 - \delta, 1 + \delta] \sigma_1(X')$ and $\sigma_n(X' + Y') \in [1 - \delta, 1 + \delta] \sigma_n(X')$. \square

It is well known that when m is sufficiently larger than n , then the singular values of a “random matrix” $X' \in \mathbb{R}^{m \times n}$ are all of size roughly \sqrt{m} . For example, Lemma 1 below is a special case of [8, Thm 3.1], and Lemma 2 can be proved along the same lines of (but much simpler than) the proof of [17, Corollary 2.3.5].

Lemma 1. *There exists a universal constant $C > 1$ such that for any $m > 2n$, if the entries of $X' \in \mathbb{R}^{m \times n}$ are drawn independently from $\mathcal{N}(0, 1)$ then $\Pr[\sigma_n(X') < \sqrt{m}/C] < \exp(-O(m))$.* \square

Lemma 2. *There exists a universal constant $C > 1$ such that for any $m > 2n$, if the entries of $X' \in \mathbb{R}^{m \times n}$ are drawn independently from $\mathcal{N}(0, 1)$ then $\Pr[\sigma_1(X') > C\sqrt{m}] < \exp(-O(m))$.* \square

Corollary 1. *There exists a universal constant $C > 1$ such that for any $m > 2n$ and $s > 0$, if the entries of $X' \in \mathbb{R}^{m \times n}$ are drawn independently from $\mathcal{N}(0, s^2)$ then*

$$\Pr[s\sqrt{m}/C < \sigma_n(X') \leq \sigma_1(X') < sC\sqrt{m}] > 1 - \exp(-O(m)). \quad \square$$

Remark. The literature on random matrices is mostly focused on analyzing the “hard cases” of more general distributions and m which is very close to n (e.g., $m = (1 + o(1))n$ or even $m = n$). For our purposes, however, we only need the “easy case” where all the distributions are Gaussian and $m \gg n$ (e.g., $m = n^2$), in which case all the proofs are much easier (and the universal constant from Corollary 1 gets closer to one).

2.3 Lattices and their Dual

A lattice $L \subset \mathbb{R}^n$ is an additive discrete sub-group of \mathbb{R}^n . We denote by $\text{span}(L)$ the linear subspace of \mathbb{R}^n , spanned by the points in L . The rank of $L \subset \mathbb{R}^n$ is the dimension of $\text{span}(L)$, and we say that L has full rank if its rank is n . In this work we often consider lattices of less than full rank.

Every (nontrivial) lattice has bases: a basis for a rank- k lattice L is a set of k linearly independent points $\mathbf{b}_1, \dots, \mathbf{b}_k \in L$ such that $L = \{\sum_{i=1}^k z_i \mathbf{b}_i : z_i \in \mathbb{Z} \forall i\}$. If we arrange the vectors \mathbf{b}_i as the columns of a matrix $B \in \mathbb{R}^{n \times k}$ then we can write $L = \{B\mathbf{z} : \mathbf{z} \in \mathbb{Z}^k\}$. If B is a basis for L then we say that B spans L .

Definition 1 (Dual of a Lattice). *For a lattice $L \subset \mathbb{R}^n$, its dual lattice consists of all the points in $\text{span}(L)$ that are orthogonal to L modulo one, namely:*

$$L^* = \{\mathbf{y} \in \text{span}(L) : \forall \mathbf{x} \in L, \langle \mathbf{x}, \mathbf{y} \rangle \in \mathbb{Z}\}$$

Clearly, if L is spanned by the columns of some rank- k matrix $X \in \mathbb{R}^{n \times k}$ then L^* is spanned by the columns of the pseudoinverse of X . It follows from the definition that for two lattices $L \subseteq M$ we have $M^* \cap \text{span}(L) \subseteq L^*$.

Banasczyk provided strong transference theorems that relate the size of short vectors in L to the size of short vectors in L^* . Recall that $\lambda_i(L)$ denotes the i -th minimum of L (i.e., the smallest s such that L contains i linearly independent vectors of size at most s).

Theorem 1 (Banasczyk [2]). *For any rank- n lattice $L \subset \mathbb{R}^m$, and for all $i \in [n]$,*

$$1 \leq \lambda_i(L) \cdot \lambda_{n-i+1}(L^*) \leq n.$$

2.4 Gaussian Distributions over Lattices

The *ellipsoid discrete Gaussian distribution* over lattice L with parameter S , centered around \mathbf{c} , is

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

where $\rho_{S,\mathbf{c}}(A)$ for set A denotes $\sum_{\mathbf{x} \in A} \rho_{S,\mathbf{c}}(\mathbf{x})$. In other words, the probability $\mathcal{D}_{L,S,\mathbf{c}}(\mathbf{x})$ is simply proportional to $\rho_{S,\mathbf{c}}(\mathbf{x})$, the denominator being a normalization factor. The same definitions apply to the spherical case, which is denoted by $\mathcal{D}_{L,s,\mathbf{c}}(\cdot)$ (with lowercase s). As before, when $\mathbf{c} = \mathbf{0}$ we use the shorthand $\mathcal{D}_{L,S}$ (or $\mathcal{D}_{L,s}$). The following useful fact that follows directly from the definition, relates the ellipsoid Gaussian distributions over different lattices:

Fact 2 *Let $L \subset \mathbb{R}^n$ be a full-rank lattice, $\mathbf{c} \in \mathbb{R}^n$ a vector, and $S \in \mathbb{R}^{m \times n}$, $B \in \mathbb{R}^{n \times n}$ two rank- n matrices, and denote $L' = \{B^{-1}\mathbf{v} : \mathbf{v} \in L\}$, $\mathbf{c}' = B^{-1}\mathbf{c}$, and $S' = S \times (B^\top)^{-1}$. Then the distribution $\mathcal{D}_{L,S,\mathbf{c}}$ is identical to the distribution induced by drawing a vector $\mathbf{v} \leftarrow \mathcal{D}_{L',S',\mathbf{c}'}$ and outputting $\mathbf{u} = B\mathbf{v}$. \square*

A useful special case of Fact 2 is when L' is the integer lattice, $L' = \mathbb{Z}^n$, in which case L is just the lattice spanned by the basis B . In other words, the ellipsoid Gaussian distribution on $L(B)$, $\mathbf{v} \leftarrow \mathcal{D}_{L(B),S,\mathbf{c}}$, is induced by drawing an integer vector according to $\mathbf{z} \leftarrow \mathcal{D}_{\mathbb{Z}^n,S',\mathbf{c}'}$ and outputting $\mathbf{v} = B\mathbf{z}$, where $S' = S(B^{-1})^\top$ and $\mathbf{c}' = B^{-1}\mathbf{c}$.

Another useful special case is where $S = sB^\top$, so S is a square matrix and $S' = sI_n$. In this case the ellipsoid Gaussian distribution $\mathbf{v} \leftarrow \mathcal{D}_{L,S,\mathbf{c}}$ is induced by drawing a vector according to the *spherical Gaussian* $\mathbf{u} \leftarrow \mathcal{D}_{L',s,\mathbf{c}'}$ and outputting $\mathbf{v} = \frac{1}{s}S^\top\mathbf{u}$, where $\mathbf{c}' = s(S^\top)^{-1}\mathbf{c}$ and $L' = \{s(S^\top)^{-1}\mathbf{v} : \mathbf{v} \in L\}$.

Smoothing parameter. As in [10], for lattice L and real $\epsilon > 0$, the *smoothing parameter* of L , denoted $\eta_\epsilon(L)$, is defined as the smallest s such that $\rho_{1/s}(L^* \setminus \{\mathbf{0}\}) \leq \epsilon$. Intuitively, for a small enough ϵ , the number $\eta_\epsilon(L)$ is sufficiently larger than L 's fundamental parallelepiped so that sampling from the corresponding Gaussian “wipes out the internal structure” of L . Thus, the sparser the lattice, the larger its smoothing parameter.

It is well known that for a spherical Gaussian with parameter $s > \eta_\epsilon(L)$, the size of vectors drawn from $\mathcal{D}_{L,s}$ is bounded by $s\sqrt{n}$ whp (cf. [10, Lemma 4.4], [12, Corollary 5.3]). The following lemma (that follows easily from the spherical case and Fact 2) is a generalization to ellipsoid Gaussians.

Lemma 3. *For a rank- n lattice L , vector $\mathbf{c} \in \mathbb{R}^n$, constant $0 < \epsilon < 1$ and matrix S s.t. $\sigma_n(S) \geq \eta_\epsilon(L)$, we have that for $\mathbf{v} \leftarrow \mathcal{D}_{L,s,c}$,*

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

Moreover, for every $\mathbf{z} \in \mathbb{R}^n$ $r > 0$ it holds that

$$\Pr_{\mathbf{v} \leftarrow \mathcal{D}_{L,s,c}} (|\langle \mathbf{v} - \mathbf{c}, \mathbf{z} \rangle| \geq r\sigma_1(S)\|\mathbf{z}\|) \leq 2en \cdot \exp(-\pi r^2).$$

The proof can be found in the long version [1].

The next lemma says that the Gaussian distribution with parameter $s \geq \eta_\epsilon(L)$ is so smooth and “spread out” that it covers the approximately the same number of L -points regardless of where the Gaussian is centered. This is again well known for spherical distributions (cf. [6, Lemma 2.7]) and the generalization to ellipsoid distributions is immediate using Fact 2.

Lemma 4. *For any rank- n lattice L , real $\epsilon \in (0, 1)$, vector $\mathbf{c} \in \mathbb{R}^n$, and rank- n matrix $S \in \mathbb{R}^{m \times n}$ such that $\sigma_n(S) \geq \eta_\epsilon(L)$, we have $\rho_{S,c}(L) \in [\frac{1-\epsilon}{1+\epsilon}, 1] \cdot \rho_S(L)$.* \square

Regev also proved that drawing a point from L according to a spherical discrete Gaussian and adding to it a spherical continuous Gaussian, yields a probability distribution close to a continuous Gaussian (independent of the lattice), provided that both distributions have parameters sufficiently larger than the smoothing parameter of L .

Lemma 5 (Claim 3.9 of [14]). *Fix any n -dimensional lattice $L \subset \mathbb{R}^n$, real $\epsilon \in (0, 1/2)$, and two reals s, r such that $\frac{rs}{\sqrt{r^2+s^2}} \geq \eta_\epsilon(L)$, and denote $t = \sqrt{r^2+s^2}$.*

Let $\mathcal{R}_{L,r,s}$ be a distribution induced by choosing $\mathbf{x} \leftarrow \mathcal{D}_{L,s}$ from the spherical discrete Gaussian on L and $\mathbf{y} \leftarrow \mathcal{N}^n(0, r^2/2\pi)$ from a continuous Gaussian, and outputting $\mathbf{z} = \mathbf{x} + \mathbf{y}$. Then for any point $\mathbf{u} \in \mathbb{R}^n$, the probability density $\mathcal{R}_{L,r,s}(\mathbf{u})$ is close to the probability density under the spherical continuous Gaussian $\mathcal{N}^n(0, t^2/2\pi)$ upto a factor of $\frac{1-\epsilon}{1+\epsilon}$:

$$\frac{1-\epsilon}{1+\epsilon} \mathcal{N}^n(0, t^2/2\pi)(\mathbf{u}) \leq \mathcal{R}_{L,r,s}(\mathbf{u}) \leq \frac{1+\epsilon}{1-\epsilon} \mathcal{N}^n(0, t^2/2\pi)(\mathbf{u})$$

In particular, the statistical distance between $\mathcal{R}_{L,r,s}$ and $\mathcal{N}^n(0, t^2/2\pi)$ is at most 4ϵ .

More broadly, Lemma 5 implies that for any event $E(\mathbf{u})$, we have

$$\Pr_{\mathbf{u} \leftarrow \mathcal{N}^n(0, t^2/2\pi)} [E(\mathbf{u})] \cdot \frac{1-\epsilon}{1+\epsilon} \leq \Pr_{\mathbf{u} \leftarrow \mathcal{R}_{L,r,s}} [E(\mathbf{u})] \leq \Pr_{\mathbf{u} \leftarrow \mathcal{N}^n(0, t^2/2\pi)} [E(\mathbf{u})] \cdot \frac{1+\epsilon}{1-\epsilon}$$

Another useful property of “wide” discrete Gaussian distributions is that they do not change much by short shifts. Specifically, if we have an arbitrary subset of the lattice, $T \subseteq L$, and an arbitrary *short vector* $\mathbf{v} \in L$, then the probability mass of T is not very different than the probability mass of $T - \mathbf{v} = \{\mathbf{u} - \mathbf{v} : \mathbf{u} \in T\}$. Below let $\text{erf}(\cdot)$ denote the Gauss error function.

Lemma 6. *Fix a lattice $L \subset \mathbb{R}^n$, a positive real $\epsilon > 0$, and two parameters s, c such that $c > 2$ and $s \geq (1 + c)\eta_\epsilon(L)$. Then for any subset $T \subset L$ and any additional vector $\mathbf{v} \in L$, it holds that $\mathcal{D}_{L,s}(T) - \mathcal{D}_{L,s}(T - \mathbf{v}) \leq \frac{\text{erf}(q(1+4/c)/2)}{\text{erf}(2q)} \cdot \frac{1+\epsilon}{1-\epsilon}$, where $q = \|\mathbf{v}\| \sqrt{\pi}/s$.*

We provide the proof in A.1.

One useful special case of Lemma 6 is when $c = 100$ (say) and $\|\mathbf{v}\| \approx s$, where we get a bound $\mathcal{D}_{L,s}(T) - \mathcal{D}_{L,s}(T - \mathbf{v}) \leq \frac{\text{erf}(0.52\sqrt{\pi})}{\text{erf}(2\sqrt{\pi})} \cdot \frac{1+\epsilon}{1-\epsilon} \approx 0.81$. We note that when $\frac{\|\mathbf{v}\|}{s} \rightarrow 0$, the bound from Lemma 6 tends to (just over) $1/4$, but we note that we can make it tend to zero with a different choice of parameters in the proof (namely making $H'_\mathbf{v}$ and $H''_\mathbf{v}$ thicker, e.g. $H''_\mathbf{v} = H_\mathbf{v}$ and $H'_\mathbf{v} = 2H_\mathbf{v}$). Lemma 6 extends easily also to the ellipsoid Gaussian case, using Fact 2:

Corollary 2. *Fix a lattice $L \subset \mathbb{R}^n$, a positive real $\epsilon > 0$, a parameter $c > 2$ and a rank- n matrix S such that $s \stackrel{\text{def}}{=} \sigma_n(S) \geq (1 + c)\eta_\epsilon(L)$. Then for any subset $T \subset L$ and any additional vector $\mathbf{v} \in L$, it holds that $\mathcal{D}_{L,S}(T) - \mathcal{D}_{L,S}(T - \mathbf{v}) \leq \frac{\text{erf}(q(1+4/c)/2)}{\text{erf}(2q)} \cdot \frac{1+\epsilon}{1-\epsilon}$, where $q = \|\mathbf{v}\| \sqrt{\pi}/s$.*

Micciancio and Regev give the following bound on the smoothing parameter in terms of the primal lattice.

Lemma 7. *[Lemma 3.3 of [10]] For any n -dimensional lattice L and positive real $\epsilon > 0$,*

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

In particular, for any superlogarithmic function $\omega(\log n)$, there exists a negligible function $\epsilon(n)$ such that $\eta_\epsilon(L) \leq \sqrt{\omega(\log n)} \cdot \lambda_n(L)$.

3 Our Discrete Gaussian LHL

Consider a full rank lattice $L \subseteq \mathbb{Z}^n$, some negligible $\epsilon = \epsilon(n)$, the corresponding smoothing parameter $\eta = \eta_\epsilon(L)$ and parameters $s > \Omega(\eta)$, $m > \Omega(n \log n)$, and $s' > \Omega(\text{poly}(n) \log(1/\epsilon))$. The process that we analyze begins by choosing “once and for all” m points in L , drawn independently from a discrete Gaussian with parameter s , $\mathbf{x}_i \leftarrow \mathcal{D}_{L,s}$.⁶

⁶ More generally, we can consider drawing the vectors \mathbf{x}_i from an ellipsoid discrete Gaussian, $\mathbf{x}_i \leftarrow \mathcal{D}_{L,S}$, so long as the least singular value of S is at least s .

Once the \mathbf{x}_i 's are fixed, we arrange them as the columns of an n -by- m matrix $X = (\mathbf{x}_1 | \mathbf{x}_2 | \dots | \mathbf{x}_m)$, and consider the distribution $\mathcal{E}_{X,s'}$, induced by choosing an integer vector \mathbf{v} from a discrete spherical Gaussian with parameter s' and outputting $\mathbf{y} = X \cdot \mathbf{v}$:

$$\mathcal{E}_{X,s'} \stackrel{\text{def}}{=} \{X \cdot \mathbf{v} : \mathbf{v} \leftarrow \mathcal{D}_{\mathbb{Z}^m, s'}\}. \quad (1)$$

Our goal is to prove that $\mathcal{E}_{X,s'}$ is close to the ellipsoid Gaussian $\mathcal{D}_{L, s' X^\top}$, over L . We begin by proving that the singular values of X^\top are all roughly of the size $s\sqrt{m}^\top$.

Lemma 8. *There exists a universal constant $K > 1$ such that for all $m \geq 2n$, $\epsilon > 0$ and every n -dimensional real lattice $L \subset \mathbb{R}^n$, the following holds: choosing the rows of an m -by- n matrix X^\top independently at random from a spherical discrete Gaussian on L with parameter $s > 2K\eta_\epsilon(L)$, $X^\top \leftarrow (\mathcal{D}_{L,s})^m$, we have*

$$\Pr \left[s\sqrt{2\pi m}/K < \sigma_n(X^\top) \leq \sigma_1(X^\top) < sK\sqrt{2\pi m} \right] > 1 - (4m\epsilon + O(\exp(-m/K))).$$

The proof can be found in the long version [1].

3.1 The Distribution $\mathcal{E}_{X,s'}$ Over \mathbb{Z}^n

We next move to show that with high probability over the choice of X , the distribution $\mathcal{E}_{X,s'}$ is statistically close to the ellipsoid discrete Gaussian $\mathcal{D}_{L, s' X^\top}$. We first prove this for the special case of the integer lattice, $L = \mathbb{Z}^n$, and then use that special case to prove the same statement for general lattices. In either case, we analyze the setting where the columns of X are chosen from an ellipsoid Gaussian which is “not too small” and “not too skewed.”

Parameters. Below n is the security parameters and $\epsilon = \text{negligible}(n)$. Let S be an n -by- n matrix such that $\sigma_n(S) \geq 2K\eta_\epsilon(\mathbb{Z}^n)$, and denote $s_1 = \sigma_1(S)$, $s_n = \sigma_n(S)$, and $w = s_1/s_n$. (We consider w to be a measure for the “skewness” of S .) Also let m, q, s' be parameters satisfying $m \geq 10n \log q$, $q > 8m^{5/2}n^{1/2}s_1w$, and $s' \geq 4wm^{3/2}n^{1/2} \ln(1/\epsilon)$. An example setting of parameters to keep in mind is $m = n^2$, $s_n = \sqrt{n}$ (which implies $\epsilon \approx 2^{-\sqrt{n}}$), $s_1 = n$ (so $w = \sqrt{n}$), $q = 8n^7$, and $s' = n^5$.

Theorem 2. *For ϵ negligible in n , let $S \in \mathbb{R}^{n \times n}$ be a matrix such that $s_n = \sigma_n(S) \geq 18K\eta_\epsilon(\mathbb{Z}^n)$, and denote $s_1 = \sigma_1(S)$ and $w = s_1/s_n$. Also let m, s' be parameters such that $m \geq 10n \log(8m^{5/2}n^{1/2}s_1w)$ and $s' \geq 4wm^{3/2}n^{1/2} \ln(1/\epsilon)$.*

Then, when choosing the columns of an n -by- m matrix X from the ellipsoid Gaussian over \mathbb{Z}^n , $X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, S})^m$, we have with all but probability $2^{-O(m)}$ over the choice of X , that the statistical distance between $\mathcal{E}_{X,s'}$ and the ellipsoid Gaussian $\mathcal{D}_{\mathbb{Z}^n, s' X^\top}$ is bounded by 2ϵ .

⁷ Since we eventually apply the following lemmas to X^\top , we will use X^\top in the statement of the lemmas for consistency at the risk of notational clumsiness.

The rest of this subsection is devoted to proving Theorem 2. We begin by showing that with overwhelming probability, the columns of X span all of \mathbb{Z}^n , which means also that the support of $\mathcal{E}_{X,s'}$ includes all of \mathbb{Z}^n .

Lemma 9. *With parameters as above, when drawing the columns of an n -by- m matrix X independently at random from $\mathcal{D}_{\mathbb{Z}^n,S}$ we get $X \cdot \mathbb{Z}^m = \mathbb{Z}^n$ with all but probability $2^{-O(m)}$.*

The proof can be found in the long version [1].

From now on we assume that the columns of X indeed span all of \mathbb{Z}^n . Now let $A = A(X)$ be the $(m - n)$ -dimensional lattice in \mathbb{Z}^m orthogonal to all the rows of X , and for any $\mathbf{z} \in \mathbb{Z}^n$ we denote by $A_{\mathbf{z}} = A_{\mathbf{z}}(X)$ the \mathbf{z} coset of A :

$$A = A(X) \stackrel{\text{def}}{=} \{\mathbf{v} \in \mathbb{Z}^m : X \cdot \mathbf{v} = 0\} \text{ and } A_{\mathbf{z}} = A_{\mathbf{z}}(X) \stackrel{\text{def}}{=} \{\mathbf{v} \in \mathbb{Z}^m : X \cdot \mathbf{v} = \mathbf{z}\}.$$

Since the columns of X span all of \mathbb{Z}^n then $A_{\mathbf{z}}$ is nonempty for every $\mathbf{z} \in \mathbb{Z}^n$, and we have $A_{\mathbf{z}} = \mathbf{v}_{\mathbf{z}} + A$ for any arbitrary point $\mathbf{v}_{\mathbf{z}} \in A_{\mathbf{z}}$.

Below we prove that the smoothing parameter of A is small (whp), and use that to bound the distance between $\mathcal{E}_{X,s'}$ and $\mathcal{D}_{\mathbb{Z}^n,s'X^\top}$. First we show that if the smoothing parameter of A is indeed small (i.e., smaller than the parameter s' used to sample the coefficient vector \mathbf{v}), then $\mathcal{E}_{X,s'}$ and $\mathcal{D}_{\mathbb{Z}^n,s'X^\top}$ must be close.

Lemma 10. *Fix X and $A = A(X)$ as above. If $s' \geq \eta_\epsilon(A)$, then for any point $\mathbf{z} \in \mathbb{Z}^n$, the probability mass assigned to \mathbf{z} by $\mathcal{E}_{X,s'}$ differs from that assigned by $\mathcal{D}_{\mathbb{Z}^n,s'X^\top}$ by at most a factor of $(1 - \epsilon)/(1 + \epsilon)$, namely*

$$\mathcal{E}_{X,s'}(\mathbf{z}) \in \left[\frac{1-\epsilon}{1+\epsilon}, 1\right] \cdot \mathcal{D}_{\mathbb{Z}^n,s'X^\top}(\mathbf{z}).$$

In particular, if $\epsilon < 1/3$ then the statistical distance between $\mathcal{E}_{X,s'}$ and $\mathcal{D}_{\mathbb{Z}^n,s'X}$ is at most 2ϵ .

The proof can be found in Appendix A.2.

The smoothing parameter of A . We now turn our attention to proving that A is “smooth enough”. Specifically, for the parameters above we prove that with high probability over the choice of X , the smoothing parameter $\eta_\epsilon(A)$ is bounded below $s' = 4wm^{3/2}n^{1/2} \ln(1/\epsilon)$.

Recall again that $A = A(X)$ is the rank- $(m - n)$ lattice containing all the integer vectors in \mathbb{Z}^m orthogonal to the rows of X . We extend A to a full-rank lattice as follows: First we extend the rows space of X , by throwing in also the scaled standard unit vectors qe_i for the integer parameter q mentioned above ($q \geq 8m^{5/2}n^{1/2}s_1w$). That is, we let $M_q = M_q(X)$ be the full-rank m -dimensional lattice spanned by the rows of X and the vectors qe_i ,

$$M_q = \{X^\top \mathbf{z} + q\mathbf{y} : \mathbf{z} \in \mathbb{Z}^n, \mathbf{y} \in \mathbb{Z}^m\} = \{\mathbf{u} \in \mathbb{Z}^m : \exists \mathbf{z} \in \mathbb{Z}_q^n \text{ s.t. } \mathbf{u} \equiv X^\top \mathbf{z} \pmod{q}\}$$

(where we identify \mathbb{Z}_q above with the set $[-q/2, q/2) \cap \mathbb{Z}$). Next, let A_q be the dual of M_q , scaled up by a factor of q , i.e.,

$$\begin{aligned} A_q &= qM_q^* = \{\mathbf{v} \in \mathbb{R}^m : \forall \mathbf{u} \in M_q, \langle \mathbf{v}, \mathbf{u} \rangle \in q\mathbb{Z}\} \\ &= \{\mathbf{v} \in \mathbb{R}^m : \forall \mathbf{z} \in \mathbb{Z}_q^n, \mathbf{y} \in \mathbb{Z}^m, \mathbf{z}^\top X \cdot \mathbf{v} + q \langle \mathbf{v}, \mathbf{y} \rangle \in q\mathbb{Z}\} \end{aligned}$$

It is easy to see that $A \subset A_q$, since any $\mathbf{v} \in A$ is an integer vector (so $q \langle \mathbf{v}, \mathbf{y} \rangle \in q\mathbb{Z}$ for all $\mathbf{y} \in \mathbb{Z}^m$) and orthogonal to the rows of X (so $\mathbf{z}^\top X \cdot \mathbf{v} = 0$ for all $\mathbf{z} \in \mathbb{Z}_q^n$).

Obviously all the rows of X belong to M_q , and whp they are linearly independent and relatively short (i.e., of size roughly $s_1\sqrt{m}$). In Lemma 11 below we show, however, that whp over the choice of X 's, these are essentially the *only* short vectors in M_q .

Lemma 11. *Recall that we choose X as $X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, S})^m$, and let $w = \sigma_1(S)/\sigma_n(S)$ be a measure of the “skewness” of S . The $n+1$ 'st minima of the lattice $M_q = M_q(X)$ is at least $q/(4w\sqrt{mn})$, except with negligible probability over the choice of X . Namely, $\Pr_{X \leftarrow (\mathcal{D}_{\mathbb{Z}^n, S})^m}[\lambda_{n+1}(M_q) < q/(4w\sqrt{mn})] < 2^{-O(m)}$.*

Proof. We prove that with high probability over the choice of X , every vector in M_q which is *not* in the linear span of the rows of X is of size at least $q/4nw$.

Recall that every vector in M_q is of the form $X^\top \mathbf{z} + q\mathbf{y}$ for some $\mathbf{z} \in \mathbb{Z}_q^n$ and $\mathbf{y} \in \mathbb{Z}^m$. Let us denote by $[\mathbf{v}]_q$ the modular reduction of all the entries in \mathbf{v} into the interval $[-q/2, q/2)$, then clearly for every $\mathbf{z} \in \mathbb{Z}_q^n$

$$\|[X^\top \mathbf{z}]_q\| = \inf\{\|X^\top \mathbf{z} + q\mathbf{y}\| : \mathbf{y} \in \mathbb{Z}^m\}.$$

Moreover, for every $\mathbf{z} \in \mathbb{Z}_q^n, \mathbf{y} \in \mathbb{Z}^m$, if $X^\top \mathbf{z} + q\mathbf{y} \neq [X^\top \mathbf{z}]_q$ then $\|X^\top \mathbf{z} + q\mathbf{y}\| \geq q/2$. Thus it suffices to show that every vector of the form $[X^\top \mathbf{z}]_q$ which is not in the linear span of the rows of X has size at least $q/4nw$ (whp over the choice of X).

Fix a particular vector $\mathbf{z} \in \mathbb{Z}_q^n$ (i.e. an integer vector with entries in $[-q/2, q/2)$). For this fixed vector \mathbf{z} , let i_{\max} be the index of the largest entry in \mathbf{z} (in absolute value), and let z_{\max} be the value of that entry. Considering the vector $\mathbf{v} = [X^\top \mathbf{z}]_q$ for a random matrix X whose columns are drawn independently from the distribution $\mathcal{D}_{\mathbb{Z}^n, S}$, each entry of \mathbf{v} is the inner product of the fixed vector \mathbf{z} with a random vector $\mathbf{x}_i \leftarrow \mathcal{D}_{\mathbb{Z}^n, S}$, reduced modulo q into the interval $[-q/2, +q/2)$.

Denoting $s_1 = \sigma_1(S)$ and $s_n = \sigma_n(S)$, we now have two cases, either \mathbf{z} is “small”, i.e., $|z_{\max}| < q/(2s_1\sqrt{mn})$ or it is “large”, $|z_{\max}| \geq q/(2s_1\sqrt{mn})$. By the “moreover” part in Lemma 3 (with $r = \sqrt{m}$), for each \mathbf{x}_i we have $|\langle \mathbf{x}_i, \mathbf{z} \rangle| \leq s_1\sqrt{m}\|\mathbf{z}\|$ except with probability bounded below 2^{-m} . If \mathbf{z} is “small” then $\|\mathbf{z}\| \leq q/(2s_1\sqrt{m})$ and so we get

$$|\langle \mathbf{x}_i, \mathbf{z} \rangle| \leq \|\mathbf{z}\| \cdot s_1\sqrt{m} < q/2$$

except with probability $< 2^{-m}$. Hence except with probability $m2^{-m}$ all the entries of $X^\top \mathbf{z}$ are smaller than $q/2$ in magnitude, which means that $[X^\top \mathbf{z}]_q =$

$X^\top \mathbf{z}$, and so $[X^\top \mathbf{z}]_q$ belongs to the row space of X . Using the union bound again, we get that with all but probability $q^n \cdot m2^{-m} < m2^{-9m/10}$, the vectors $[X^\top \mathbf{z}]_q$ for all the “small” \mathbf{z} ’s belong to the row space of X .

We next turn to analyzing “large” \mathbf{z} ’s. Fix one “large” vector \mathbf{z} , and for that vector define the set of “bad” vectors $\mathbf{x} \in \mathbb{Z}^n$, i.e. the ones for which $|\langle \mathbf{z}, \mathbf{x} \rangle|_q < q/4nw$ (and the other vectors $\mathbf{x} \in \mathbb{Z}^n$ are “good”). Observe that if \mathbf{x} is “bad”, then we can get a “good” vector by adding to it the i_{\max} ’th standard unit vector, scaled up by a factor of $\mu = \min(\lceil s_n \rceil, \lfloor q/2z_{\max} \rfloor)$, since

$$|\langle \mathbf{z}, \mathbf{x} + \mu \mathbf{e}_{i_{\max}} \rangle|_q = |\langle \mathbf{z}, \mathbf{x} \rangle + \mu z_{\max}|_q \geq \mu |z_{\max}| - |\langle \mathbf{z}, \mathbf{x} \rangle|_q \geq q/4nw.$$

(The last two inequalities follow from $q/2nw < \mu |z_{\max}| \leq q/2$ and $|\langle \mathbf{z}, \mathbf{x} \rangle|_q < q/(4w\sqrt{mn})$.) Hence the injunction $\mathbf{x} \mapsto \mathbf{x} + \mu \mathbf{e}_{i_{\max}}$ maps “bad” \mathbf{x} ’es to “good” \mathbf{x} ’es. Moreover, since the \mathbf{x} ’es are chosen according to the wide ellipsoid Gaussian $\mathcal{D}_{\mathbb{Z}^n, S}$ with $\sigma_n(S) = s_n \geq \eta_\epsilon(\mathbb{Z}^n)$, and since the scaled standard unit vectors are short, $\mu < s_n + 1$, then by Lemma 6 the total probability mass of the “bad” vectors \mathbf{x} differs from the total mass of the “good” vectors $\mathbf{x} + \mu \mathbf{e}_{i_{\max}}$ by at most 0.81. It follows that when choosing $\mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}^n, S}$, we have $\Pr_{\mathbf{x}} [|\langle \mathbf{z}, \mathbf{x} \rangle|_q < q/(4w\sqrt{mn})] \leq (1 + 0.81)/2 < 0.91$. Thus the probability that all the entries of $[X^\top \mathbf{z}]_q$ are smaller than $q/(4w\sqrt{nm})$ in magnitude is bounded by $(0.91)^m = 2^{-0.14m}$. Since $m > 10n \log q$, we can use the union bound to conclude that the probability that there exists some “large” vector for which $\|[X^\top \mathbf{z}]_q\| < q/(4w\sqrt{mn})$ is no more than $q^n \cdot 2^{-0.14m} < 2^{-O(m)}$.

Summing up the two cases, with all but probability $2^{-O(m)}$ over the choice of X , there does not exist any vector $\mathbf{z} \in \mathbb{Z}_q^n$ for which $[X^\top \mathbf{z}]_q$ is linearly independent of the rows of X and yet $\|[X^\top \mathbf{z}]_q\| < q/(4w\sqrt{mn})$.

Corollary 3. *With the parameters as above, the smoothing parameter of $A = A(X)$ satisfies $\eta_\epsilon(A) \leq s' = 4wm^{3/2}n^{1/2} \ln(1/\epsilon)$, except with probability $2^{-O(m)}$.*

The proof can be found in the long version [1].

Putting together Lemma 10 and Corollary 3 completes the proof of Theorem 2. \square

3.2 The Distribution $\mathcal{E}_{X, s'}$ Over General Lattices

Armed with Theorem 2, we turn to prove the same theorem also for general lattices.

Theorem 3. *Let L be a full-rank lattice $L \subset \mathbb{R}^n$ and B a matrix whose columns form a basis of L . Also let $M \in \mathbb{R}^{n \times n}$ be a full rank matrix, and denote $S = M(B^\top)^{-1}$, $s_1 = \sigma_1(S)$, $s_n = \sigma_n(S)$, and $w = s_1/s_n$. Finally let ϵ be negligible in n and m, s' be parameters such that $m \geq 10n \log(8m^{5/2}n^{1/2}s_1w)$ and $s' \geq 4wm^{3/2}n^{1/2} \ln(1/\epsilon)$.*

If $s_n \geq \eta_\epsilon(\mathbb{Z}^n)$, then, when choosing the columns of an n -by- m matrix X from the ellipsoid Gaussian over L , $X \leftarrow (\mathcal{D}_{L, M})^m$, we have with all but probability $2^{-O(m)}$ over the choice of X , that the statistical distance between $\mathcal{E}_{X, s'}$ and the ellipsoid Gaussian $\mathcal{D}_{L, s' X^\top}$ is bounded by 2ϵ .

This theorem is an immediate corollary of Theorem 2 and Fact 2. The proof can be found in the long version [1].

4 Applications

In this section, we discuss the application of our discrete Gaussian LHL in the construction of multilinear maps from lattices [5]. This construction is illustrative of a “canonical setting” where our lemma should be useful.

Brief overview of the GGH Construction. To begin, we provide a very high level overview of the GGH construction, skipping most details. We refer the reader to [5] for a complete description. In [5], the mapping $a \rightarrow g^a$ from bilinear maps is viewed as a form of “encoding” $a \mapsto Enc(a)$ that satisfies some properties:

1. Encoding is easy to compute in the forward direction and hard to invert.
2. Encoding is additively homomorphic and also one-time multiplicatively homomorphic (via the pairing).
3. Given $Enc(a), Enc(b)$ it is easy to test whether $a = b$.
4. Given encodings, it is hard to test more complicated relations between the underlying scalars. For example, BDDH roughly means that given $Enc(a), Enc(b), Enc(c), Enc(d)$ it is hard to test if $d = abc$.

In [5], the authors construct encodings from ideal lattices that approximately satisfy (and generalize) the above properties. Skipping most of the details, [5] roughly used a specific (NTRU-like) lattice-based homomorphic encryption scheme, where $Enc(a)$ is just an encryption of a . The ability to add and multiply then just follows from the homomorphism of the underlying cryptosystem, and GGH described how to add to this cryptosystem a “broken secret key” that cannot be used for decryption but is good enough for testing if two ciphertexts encrypt the same element. (In the terminology from [5], this broken key is called the *zero-test parameter*.)

In the specific cryptosystem used in the GGH construction, ciphertexts are elements in some polynomial ring (represented as vectors in \mathbb{Z}^n), and additive/multiplicative homomorphism is implemented simply by addition and multiplication in the ring. A natural way to enable encoding is to publish a single ciphertext that encrypts/encodes 1, $\mathbf{y}_1 = Enc(1)$. To encode any other plaintext element a , we can use the multiplicative homomorphism by setting $Enc(a) = a \cdot \mathbf{y}_1$ in the ring. However this simple encoding is certainly not hard to decode: just dividing by \mathbf{y}_1 in the ring suffices! For the same reason, it is also not hard to determine “complex relations” between encoding.

Randomizing the encodings. To break these simple algebraic relations, the authors include in the public parameters also “randomizers” \mathbf{x}_i ($i = 1, \dots, m$), which are just random encryptions/encodings of zero, namely $\mathbf{x}_i \leftarrow Enc(0)$. Then to re-randomize the encoding $\mathbf{u}_a = a \cdot \mathbf{y}_1$, they add to it a “random linear combination” of the \mathbf{x}_i ’s, and (by additive homomorphism) this is another

encoding of the same element. This approach seems to be thwart the simple algebraic decoding from above, but what can be said about the resulting encodings? Here is where GGH use our results to analyze the probability distribution of these re-randomized encodings.

In a little more detail, an instance of the GGH encoding includes an ideal lattice L and a secret ring element z , and an encoding of an element a has the form e_a/z where e_a is a short element that belongs to the same coset of L as the “plaintext” a . The x_i ’s are therefore ring elements of the form b_i/z where the b_i ’s are short vectors in L . Denoting by X the matrix with the x_i as columns and by B the matrix with the numerators b_i as columns, i.e., $X = (x_1 | \dots | x_m)$ and $B = (b_1 | \dots | b_m)$. Re-randomizing the encoding $u_a = e_a/z$ is obtained by choosing a random coefficient vector $r \leftarrow D_{\mathbb{Z}^m, \sigma^*}$ (for large enough σ^*), and setting

$$u' := u_a + Xr = \frac{e_a + Br}{z}.$$

Since all the b_i ’s are in the lattice L , then obviously $e_a + Br$ is in the same coset of L as e_a itself. Moreover since the b_i ’s are short and so are the coefficients of r , then also so is $e_a + Br$. Hence u' is a valid encoding of the same plaintext a that was encoded in u_a .

Finally, using our Theorem 3 from this work, GGH can claim that the distribution of u is nearly independent of the original u_a (conditioned on its coset). If the b_i ’s are chosen from a wide enough spherical distribution, then our Gaussian LHL allows them to conclude that Br is close to a wide ellipsoid Gaussian. With appropriate choice of σ^* the “width” of that distribution is much larger than the original e_a , hence the distribution of $e_a + Br$ is nearly independent of e_a , conditioned on the coset it belongs to.

5 Discussion

Unlike the classic LHL, our lattice version of LHL is less than perfect – instead of yielding a perfectly spherical Gaussian, it only gives us an approximately spherical one, i.e. $\mathcal{D}_{L, s', X^\top}$. Here approximately spherical means that all the singular values of the matrix X^\top are within a small, constant sized interval. It is therefore natural to ask: 1) Can we do better and obtain a perfectly spherical Gaussian? 2) Is an approximately spherical Gaussian sufficient for cryptographic applications?

First let us consider whether we can make the Gaussian perfectly spherical. Indeed, as the number of lattice vectors m grows larger, we expect the greatest and least singular value of the discrete Gaussian matrix X to converge – this would imply that as $m \rightarrow \infty$, the linear combination $\sum_{i=1}^m z_i x_i$ does indeed behave like a spherical Gaussian. While we do not prove this, we refer the reader to [16] for intuitive evidence. However, the focus of this work is small m (e.g., $m = \tilde{O}(n)$) suitable for applications, in which case we do not know how to prove the same.

This leads to the second question: is approximately spherical good enough? This depends on the application. We have already seen that it is sufficient for GGH encodings [5], where a canonical, wide-enough, but non-spherical Gaussian is used to “drown out” an initial encoding, and send it to a canonical distribution of encodings that encode the same value. Our LHL shows that one can sample from such a canonical approximate Gaussian distribution without using the initial Gaussian samples “wastefully”.

On the other hand, we caution the reader that if the application requires the basis vectors $\mathbf{x}_1, \dots, \mathbf{x}_m$ to be kept secret (such as when the basis is a trapdoor), then one must carefully consider whether our Gaussian sampler can be used safely. This is because, as demonstrated by [11] and [4], lattice applications where the basis is desired to be secret can be broken completely even if partial information about the basis is leaked. In an application where the trapdoor is available explicitly and oblivious sampling is not needed, it is safer to use the samplers of [6] or [13] to sample a perfectly spherical Gaussian that is statistically independent of the trapdoor.

Acknowledgments. The first and fourth authors were supported in part from a DARPA/ONR PROCEED award, NSF grants 1228984, 1136174, 1118096, and 1065276, a Xerox Faculty Research Award, a Google Faculty Research Award, an equipment grant from Intel, and an Okawa Foundation Research Grant. This material is based upon work supported by the Defense Advanced Research Projects Agency through the U.S. Office of Naval Research under Contract N00014-11-1-0389. The views expressed are those of the author and do not reflect the official policy or position of the Department of Defense, the National Science Foundation, or the U.S. Government.

The second and third authors were supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center (DoI/NBC) contract number D11PC20202. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

References

1. Shweta Agrawal, Craig Gentry, Shai Halevi, and Amit Sahai. Discrete gaussian leftover hash lemma over infinite domains. <http://eprint.iacr.org/2012/714>, 2012.
2. Wojciech Banaszczyk. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen*, 296(4):625–635, 1993.
3. Dan Boneh and David Mandall Freeman. Homomorphic signatures for polynomial functions. In *Eurocrypt*, volume 6632 of *Lecture Notes in Computer Science*, pages 149–168. Springer, 2011.

4. Léo Ducas and Phong Q. Nguyen. Learning a zonotope and more: Cryptanalysis of ntrusign countermeasures. In *ASIACRYPT*, volume 7658 of *Lecture Notes in Computer Science*, pages 433–450, 2012.
5. Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices and applications. In *Eurocrypt*, volume 7881 of *Lecture Notes in Computer Science*, pages 1–17. Springer, 2013. Full version in <http://eprint.iacr.org/2013/610>.
6. Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. Trapdoors for hard lattices and new cryptographic constructions. In Cynthia Dwork, editor, *STOC*, pages 197–206. ACM, 2008.
7. Philip Klein. Finding the closest lattice vector when it’s unusually close. In *Proceedings of the eleventh annual ACM-SIAM symposium on Discrete algorithms*, SODA’00, pages 937–941, 2000.
8. A. E. Litvak, A. Pajor, M. Rudelson, and N. Tomczak-Jaegermann. Smallest singular value of random matrices and geometry of random polytopes. *Advances in Mathematics*, 195(2), 2005.
9. Vadim Lyubashevsky. Lattice signatures without trapdoors. In *Eurocrypt*, volume 7237 of *Lecture Notes in Computer Science*, pages 738–755. Springer, 2012.
10. Daniele Micciancio and Oded Regev. Worst-case to average-case reductions based on gaussian measures. *SIAM J. Computing*, 37(1):267–302, 2007.
11. Phong Q. Nguyen and Oded Regev. Learning a parallelepiped: Cryptanalysis of GGH and NTRU signatures. *J. Cryptol.*, 22(2):139–160, April 2009.
12. Chris Peikert. Limits on the hardness of lattice problems in l_p norms. *Computational Complexity*, 17(2):300–351, 2008.
13. Chris Peikert. An efficient and parallel gaussian sampler for lattices. In *Crypto*, volume 6223 of *Lecture Notes in Computer Science*, pages 80–97. Springer, 2010.
14. Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *JACM*, 56(6), 2009.
15. Ron Rothblum. Homomorphic encryption: From private-key to public-key. In *TCC*, volume 6597 of *Lecture Notes in Computer Science*, pages 219–234. Springer, 2011.
16. Mark Rudelson and Roman Vershynin. Non-asymptotic theory of random matrices: extreme singular values. In *International Congress of Mathematicians*, 2010.
17. Terence Tao. *Topics in random matrix theory*, volume 132 of *Graduate Studies in Mathematics*. American Mathematical Society, 2012.
18. Marten van Dijk, Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. Fully homomorphic encryption over the integers. In Henri Gilbert, editor, *Advances in Cryptology - EUROCRYPT’10*, volume 6110 of *Lecture Notes in Computer Science*, pages 24–43. Springer, 2010.

A More Proofs

A.1 Proof of Lemma 6

Proof. Clearly for any fixed \mathbf{v} , the set that maximizes $\mathcal{D}_{L,s}(T) - \mathcal{D}_{L,s}(T - \mathbf{v})$ is the set of all vectors $\mathbf{u} \in L$ for which $\mathcal{D}_{L,s}(\mathbf{u}) > \mathcal{D}_{L,s}(\mathbf{u} - \mathbf{v})$, which we denote by $T_{\mathbf{v}} \stackrel{\text{def}}{=} \{\mathbf{u} \in L : \mathcal{D}_{L,s}(\mathbf{u}) > \mathcal{D}_{L,s}(\mathbf{u} - \mathbf{v})\}$. Observe that for any $\mathbf{u} \in L$ we have $\mathcal{D}_{L,s}(\mathbf{u}) > \mathcal{D}_{L,s}(\mathbf{u} - \mathbf{v})$ iff $\rho_s(\mathbf{u}) > \rho_s(\mathbf{u} - \mathbf{v})$, which is equivalent to

$\|\mathbf{u}\| < \|\mathbf{u} - \mathbf{v}\|$. That is, \mathbf{u} must lie in the half-space whose projection on \mathbf{v} is less than half of \mathbf{v} , namely $\langle \mathbf{u}, \mathbf{v} \rangle < \|\mathbf{v}\|^2/2$. In other words we have

$$T_{\mathbf{v}} = \{\mathbf{u} \in L : \langle \mathbf{u}, \mathbf{v} \rangle < \|\mathbf{v}\|^2/2\},$$

which also means that $T_{\mathbf{v}} - \mathbf{v} = \{\mathbf{u} \in L : \langle \mathbf{u}, \mathbf{v} \rangle < -\|\mathbf{v}\|^2/2\} \subseteq T_{\mathbf{v}}$. We can therefore express the difference in probability mass as $\mathcal{D}_{L,s}(T_{\mathbf{v}}) - \mathcal{D}_{L,s}(T_{\mathbf{v}} - \mathbf{v}) = \mathcal{D}_{L,s}(T_{\mathbf{v}} \setminus (T_{\mathbf{v}} - \mathbf{v}))$. Below we denote this set-difference by

$$H_{\mathbf{v}} \stackrel{\text{def}}{=} T_{\mathbf{v}} \setminus (T_{\mathbf{v}} - \mathbf{v}) = \left\{ \mathbf{u} \in L : \langle \mathbf{u}, \mathbf{v} \rangle \in \left(-\frac{\|\mathbf{v}\|^2}{2}, \frac{\|\mathbf{v}\|^2}{2}\right] \right\}.$$

That is, $H_{\mathbf{v}}$ is the “slice” in space of width $\|\mathbf{v}\|$ in the direction of \mathbf{v} , which is symmetric around the origin. The arguments above imply that for any set T we have $\mathcal{D}_{L,s}(T) - \mathcal{D}_{L,s}(T - \mathbf{v}) \leq \mathcal{D}_{L,s}(H_{\mathbf{v}})$. The rest of the proof is devoted to upper-bounding the probability mass of that slice, i.e., $\mathcal{D}_{L,s}(H_{\mathbf{v}}) = \Pr_{\mathbf{u} \leftarrow \mathcal{D}_{L,s}}[\mathbf{u} \in H_{\mathbf{v}}]$.

To this end we consider the slightly thicker slice, say $H'_{\mathbf{v}} = (1 + \frac{\epsilon}{c})H_{\mathbf{v}}$, and the random variable \mathbf{w} , which is obtained by drawing $\mathbf{u} \leftarrow \mathcal{D}_{L,s}$ and adding to it a continuous Gaussian variable of “width” s/c . We argue that \mathbf{w} is somewhat likely to fall outside of the thick slice $H'_{\mathbf{v}}$, but conditioning on $\mathbf{u} \in H_{\mathbf{v}}$ we have that \mathbf{w} is very unlikely to fall outside of $H'_{\mathbf{v}}$. Putting these two arguments together, we get that \mathbf{u} must have significant probability of falling outside $H_{\mathbf{v}}$, thereby getting our upper bound.

In more detail, denoting $r = s/c$ we consider drawing $\mathbf{u} \leftarrow \mathcal{D}_{L,s}$ and $\mathbf{z} \leftarrow \mathcal{N}^n(0, r^2/2\pi)$, and setting $\mathbf{w} = \mathbf{u} + \mathbf{z}$. Denoting $t = \sqrt{r^2 + s^2}$, we have that $s \leq t \leq s(1 + \frac{1}{c})$ and $rs/t \geq s/(c+1) \geq \eta_{\epsilon}(L)$. Thus the conditions of Lemma 5 are met, and we get that \mathbf{w} is distributed close to a normal random variable $\mathcal{N}^n(0, t^2/2\pi)$, upto a factor of at most $\frac{1+\epsilon}{1-\epsilon}$.

Since the continuous Gaussian distribution is spherical, we can consider expressing it in an orthonormal basis with one vector in the direction of \mathbf{v} . When expressed in this basis, we get the event $\mathbf{z} \in H'_{\mathbf{v}}$ exactly when the coefficient in the direction of \mathbf{v} (which is distributed close to the 1-dimensional Gaussian $\mathcal{N}(0, t^2/2\pi)$) exceeds $\|\mathbf{v}(1 + \frac{\epsilon}{c})/2\|$ in magnitude. Hence we have

$$\begin{aligned} \Pr[\mathbf{w} \in H'_{\mathbf{v}}] &\leq \Pr_{\alpha \leftarrow \mathcal{N}(0, t^2/2\pi)}[|\alpha| \leq \|\mathbf{v}\|] \cdot \frac{1+\epsilon}{1-\epsilon} \\ &= \text{erf}\left(\frac{\|\mathbf{v}\|\sqrt{\pi}(1 + \frac{\epsilon}{c})}{2t}\right) \cdot \frac{1+\epsilon}{1-\epsilon} \leq \text{erf}\left(\frac{\|\mathbf{v}\|\sqrt{\pi}(1 + \frac{\epsilon}{c})}{2s}\right) \cdot \frac{1+\epsilon}{1-\epsilon} \end{aligned}$$

On the other hand, consider the conditional probability $\Pr[\mathbf{w} \in H'_{\mathbf{v}} | \mathbf{u} \in H_{\mathbf{v}}]$: Let $H''_{\mathbf{v}} = \frac{\epsilon}{c}H_{\mathbf{v}}$, then if $\mathbf{u} \in H_{\mathbf{v}}$ and $\mathbf{z} \in H''_{\mathbf{v}}$, then it must be the case that $\mathbf{w} = \mathbf{u} + \mathbf{z} \in H'_{\mathbf{v}}$. As before, we can consider the continuous Gaussian on \mathbf{z} in an orthonormal basis with one vector in the direction of \mathbf{v} , and we get

$$\begin{aligned} \Pr[\mathbf{w} \in H'_{\mathbf{v}} | \mathbf{u} \in H_{\mathbf{v}}] &\geq \Pr[\mathbf{z} \in H''_{\mathbf{v}} | \mathbf{u} \in H_{\mathbf{v}}] = \Pr[\mathbf{z} \in H''_{\mathbf{v}}] \\ &= \Pr_{\beta \leftarrow \mathcal{N}(0, r^2/2\pi)}[|\beta| \leq 2\|\mathbf{v}\|/c] = \text{erf}(\|\mathbf{v}\|2\sqrt{\pi}/cr) = \text{erf}(2\|\mathbf{v}\|\sqrt{\pi}/s) \end{aligned}$$

Putting the last two bounds together, we get

$$\begin{aligned} \operatorname{erf}\left(\frac{\|\mathbf{v}\|\sqrt{\pi}(1+\frac{4}{c})}{2s}\right) \cdot \frac{1+\epsilon}{1-\epsilon} &\geq \Pr[\mathbf{w} \in H'_v] \geq \Pr[\mathbf{u} \in H_v] \cdot \Pr[\mathbf{w} \notin H'_v | \mathbf{u} \in H_v] \\ &\geq \Pr[\mathbf{u} \in H_v] \cdot \operatorname{erf}\left(\frac{\|\mathbf{v}\|2\sqrt{\pi}}{s}\right) \end{aligned}$$

from which we conclude that $\Pr[\mathbf{u} \in H_v] \leq \frac{\operatorname{erf}(\|\mathbf{v}\|\sqrt{\pi}(1+4/c)/2s)}{\operatorname{erf}(\|\mathbf{v}\|2\sqrt{\pi}/s)} \cdot \frac{1+\epsilon}{1-\epsilon}$, as needed.

A.2 Proof of Lemma 10

Proof. Fix some $\mathbf{z} \in \mathbb{Z}^n$. The probability mass assigned to \mathbf{z} by $\mathcal{E}_{X,s'}$ is the probability of drawing a random vector according to the discrete Gaussian $\mathcal{D}_{\mathbb{Z}^m,s'}$ and hitting some $\mathbf{v} \in \mathbb{Z}^m$ for which $X \cdot \mathbf{v} = \mathbf{z}$. In other words, this is exactly the probability mass assigned by $\mathcal{D}_{\mathbb{Z}^m,s'}$ to the coset $A_{\mathbf{z}}$. Below let $T = T(X) \subseteq \mathbb{R}^m$ be the linear subspace containing the lattice A , and $T_{\mathbf{z}} = T_{\mathbf{z}}(X) \subseteq \mathbb{R}^m$ be the affine subspace containing the coset $A_{\mathbf{z}}$:

$$T = T(X) = \{\mathbf{v} \in \mathbb{R}^m : X \cdot \mathbf{v} = 0\}, \quad \text{and} \quad T_{\mathbf{z}} = T_{\mathbf{z}}(X) = \{\mathbf{v} \in \mathbb{R}^m : X \cdot \mathbf{v} = \mathbf{z}\}.$$

Let Y be the pseudoinverse of X (i.e. $XY^\top = I_n$ and the rows of Y span the same linear sub-space as the rows of X). Let $\mathbf{u}_{\mathbf{z}} = Y^\top \mathbf{z}$, and we note that $\mathbf{u}_{\mathbf{z}}$ is the point in the affine space $T_{\mathbf{z}}$ closest to the origin: To see this, note that $\mathbf{u}_{\mathbf{z}} \in T_{\mathbf{z}}$ since $X \cdot \mathbf{u}_{\mathbf{z}} = X \times Y^\top \mathbf{z} = \mathbf{z}$. In addition, $\mathbf{u}_{\mathbf{z}}$ belongs to the row space of Y , so also to the row space of X , and hence it is orthogonal to T .

Since $\mathbf{u}_{\mathbf{z}}$ is the point in the affine space $T_{\mathbf{z}}$ closest to the origin, it follows that for every point in the coset $\mathbf{v} \in A_{\mathbf{z}}$ we have $\|\mathbf{v}\|^2 = \|\mathbf{u}_{\mathbf{z}}\|^2 + \|\mathbf{v} - \mathbf{u}_{\mathbf{z}}\|^2$, and therefore

$$\rho_{s'}(\mathbf{v}) = e^{-\pi(\|\mathbf{v}\|/s')^2} = e^{-\pi(\|\mathbf{u}_{\mathbf{z}}\|/s')^2} \cdot e^{-\pi(\|\mathbf{v}-\mathbf{u}_{\mathbf{z}}\|/s')^2} = \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \rho_{s'}(\mathbf{v} - \mathbf{u}_{\mathbf{z}}).$$

This, in turn, implies that the total mass assigned to $A_{\mathbf{z}}$ by $\rho_{s'}$ is

$$\rho_{s'}(A_{\mathbf{z}}) = \sum_{\mathbf{v} \in A_{\mathbf{z}}} \rho_{s'}(\mathbf{v}) = \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \sum_{\mathbf{v} \in A_{\mathbf{z}}} \rho_{s'}(\mathbf{v} - \mathbf{u}_{\mathbf{z}}) = \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \rho_{s'}(A_{\mathbf{z}} - \mathbf{u}_{\mathbf{z}})$$

Fix one arbitrary point $\mathbf{w}_{\mathbf{z}} \in A_{\mathbf{z}}$, and let $\delta_{\mathbf{z}}$ be the distance from $\mathbf{u}_{\mathbf{z}}$ to that point, $\delta_{\mathbf{z}} = \mathbf{u}_{\mathbf{z}} - \mathbf{w}_{\mathbf{z}}$. Since $A_{\mathbf{z}} = \mathbf{w}_{\mathbf{z}} + A$, we get $A_{\mathbf{z}} - \mathbf{u}_{\mathbf{z}} = A - \delta_{\mathbf{z}}$, and together with the equation above we have:

$$\begin{aligned} \rho_{s'}(A_{\mathbf{z}}) &= \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \rho_{s'}(A_{\mathbf{z}} - \mathbf{u}_{\mathbf{z}}) = \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \rho_{s'}(A - \delta_{\mathbf{z}}) \\ &= \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \rho_{s',\delta_{\mathbf{z}}}(A) \stackrel{\text{Lemma 4}}{=} \rho_{s'}(\mathbf{u}_{\mathbf{z}}) \cdot \rho_{s'}(A) \cdot \left[\frac{1-\epsilon}{1+\epsilon}, 1\right]. \end{aligned} \quad (3)$$

As a last step, recall that $\mathbf{u}_{\mathbf{z}} = Y^\top \mathbf{z}$ where $YY^\top = (XX^\top)^{-1}$. Thus $\rho_{s'}(\mathbf{u}_{\mathbf{z}}) =$

$$\rho_{s'}(Y^\top \mathbf{z}) = \exp(-\pi|\mathbf{z}^\top YY^\top \mathbf{z}|/s'^2) = \exp\left(-\pi|\mathbf{z}^\top ((s'X)(s'X)^\top)^{-1} \mathbf{z}|\right) = \rho_{(s'X)^\top}(\mathbf{z}).$$

Putting everything together we get

$$\mathcal{E}_{X,s'}(\mathbf{z}) = \mathcal{D}_{\mathbb{Z}^m,s'}(A_{\mathbf{z}}) = \frac{\rho_{s'}(A_{\mathbf{z}})}{\rho_{s'}(\mathbb{Z}^m)} \in \rho_{(s'X^\top)}(\mathbf{z}) \cdot \frac{\rho_{s'}(A)}{\rho_{s'}(\mathbb{Z}^m)} \cdot \left[\frac{1-\epsilon}{1+\epsilon}, 1\right]$$

The term $\frac{\rho_{s'}(A)}{\rho_{s'}(\mathbb{Z}^m)}$ is a normalization factor independent of \mathbf{z} , hence the probability mass $\mathcal{E}_{X,s'}(\mathbf{z})$ is proportional to $\rho_{(s'X^\top)}(\mathbf{z})$, upto some “deviation factor” in $[\frac{1-\epsilon}{1+\epsilon}, 1]$.