Algorithms for Sparse LWE and LPN with Small Secrets

Shashwat Agrawal
Indian Institute of Technology Delhi
csz248012@iitd.ac.in

Amitabha Bagchi
Indian Institute of Technology Delhi
bagchi@cse.iitd.ac.in

Rajendra Kumar
Indian Institute of Technology Delhi
rajendra@cse.iitd.ac.in

Abstract

We present new sample-runtime tradeoffs for the decisional sparse Learning With Errors (LWE) and sparse Learning Parity with Noise (LPN) problems over $\mathbb{Z}_q$, specifically in regimes where the secret vector is constrained by a small l_∞ norm. While small-secret constraints are useful for the practical efficiency of lattice-based cryptography—such as homomorphic encryption and zero-knowledge proofs—the extent to which an adversary can exploit these bounds when the coefficient matrix is sparse is an open question. We address this by reducing the distinguishing task to a relaxed variant of the Short Integer Solution (SIS) problem, where the strict $A^{\mathrm{T}}\boldsymbol{c} = 0$ requirement is replaced with an l_1-norm bound on $A^{\mathrm{T}}\boldsymbol{c}$. To solve this relaxed SIS problem, we design an algorithm that samples distinct, non-trivial walks on a Kikuchi graph having close end points. For LWE, this approach directly separates planted from random instances. For LPN, where the noise is uniformly distributed over non-zero elements, the proof is more involved. We first derive a different reduction from LPN to (relaxed) SIS and then extend the anti-concentration framework given by Gupta, He, O'Donnell, and Singer (SODA 2026).

1 Introduction

The Learning With Errors (LWE) problem, introduced by Regev [Reg05], and its close cousin, the Learning Parity with Noise (LPN) problem [Ale03], form central hardness assumptions in modern cryptography. In LWE, samples from the distribution $\mathcal{A}_{s,\chi}$ on $\mathbb{Z}_q^n \times \mathbb{Z}_q$ are generated by selecting $\boldsymbol{a} \in \mathbb{Z}_q^n$ uniformly at random, an error $e \in \mathbb{Z}_q$ from a distribution Ψ (typically a discrete Gaussian with width $r < q$), and computing $(\boldsymbol{a}, \langle \boldsymbol{a}, \boldsymbol{s} \rangle + e)$ modulo q; the decision version requires distinguishing these from uniform samples on $\mathbb{Z}_q^n \times \mathbb{Z}_q$. LPN, most commonly studied over $\mathbb{Z}_2$ with Bernoulli noise or more generally over modulus q with small hamming weight error, shares similar structure. Most lattice-based cryptosystems rely on LWE's conjectured intractability against classical and quantum algorithms [Reg05; Pei09; Gen09; BV14; Pei16], while LPN's hardness rests on decades of cryptanalysis [Gol00; CM01; Fei02; MST03; AOW15; AL16; KMOW17], offering competitive efficiency despite fewer constructions [Pie12; ADINZ17a]. Variants of these problems are considered to balance security with efficiency of cryptographic primitives. Different choices of modulus, noise distribution, assumption on secret vector $\boldsymbol{s}$, or structural constraints (such as sparsity) can lead to cryptographic primitives with faster key generation, smaller ciphertexts, or reduced computational overhead, while still relying on problems which are believed to have same hardness [LPR13; ABW10].

Small Secret and Sparsity. In this work, we focus on LWE and LPN with both small secret and sparsity constraints. A small secret version restricts the secret vector $\boldsymbol{s}$ to have a bounded ℓ_∞ norm. A sparse version restricts the coefficient vectors $\boldsymbol{a} \in \mathbb{Z}_q^n$ have only k non-zero entries for a parameter $k < n$.

Small secrets in the Learning With Errors (LWE) problem are essential for the practical efficiency of lattice-based cryptography, achieving optimization without sacrificing asymptotic security. It is well-established that the secret can follow the noise distribution [ACPS09], be securely constrained via dimension-modulus switching [BLPRS13], or be assumed strictly binary by scaling the dimension from n to $n \log q$ [Mic18]. For Fully Homomorphic Encryption (FHE), the small secrets constraint controls the rapid noise expansion during homomorphic evaluations, which directly enables efficient modulus and key switching [BGV14; DM15; CGGI20]. This same mechanism facilitates a wide range of advanced crypto primitives, including Circular-Secure Encryption [ACPS09], leakage-resilient cryptosystems [GKPV10], Attribute-Based Encryption (ABE) [CW23; HLL23], and lattice-based Zero-Knowledge Proofs (ZKPs) [Lyu09; ESLL19].

Similarly, the motivation for sparsity is that the samples can be computed and stored with a roughly $O(n/k)$ factor improvement in efficiency. Sparse LWE was introduced by Jain, Lin, and Saha [JLS24]. Assuming hardness of sparse LWE, the authors construct secret key Linearly Homomorphic Encryption (LHE) schemes with slightly super-constant/constant overheads, and secret key homomorphic encryption for arbitrary constant-degree polynomials with slightly super-constant/constant overheads. They also perform some cryptanalysis by suggesting two attacks for sparse LWE, namely the sparse vector in kernel attack and the dense minor attack. On analyzing the feasibility of these attacks for $k \geq \Omega(\log n)$ and $m = \text{poly}(n)$, it turns out that these attacks will take exponential time in n with high probability. Hence, the authors conjecture that sparse LWE with sparsity $k \geq \Omega(\log n)$ and polynomially many samples $m = \text{poly}(n)$ is as hard as the standard LWE with dimension $n' = \Theta(n)$ and polynomially many samples. Subsequently, Bangachev, Bresler, Tiegel and Vaikuntanathan [BBTV25] give reductions from dense LWE to sparse LWE, and similarly from dense LPN to sparse LPN. Their reductions aim to establish $n^{\Omega(k)}$-hardness of sparse LWE under the assumption that dense LWE requires exponential time, but still incur an additional $1/\log k$ multiplicative factor in the exponent. Recently, Agrawal, Bagchi and Kumar [ABK26] gave sample-runtime tradeoff algorithms for these problems using a spectral

method. For a tradeoff parameter $k/2 \leq l \leq n$, their algorithms takes roughly $(qn)^l$ time when there are roughly $(qn/l)^{k/2}$ samples.

Looking at the impact of small secret and sparsity constraints, it is natural to consider the LWE problem with both constraints. All the aforementioned hardness results for the small secret LWE rely on the assumption that the coefficient matrix A is chosen uniformly at random. We do not currently know how to obtain similar theoretical guaranties when sparsity constraints are applied to the matrix. In the absence of these reductions, a natural question arises:

Is small-secret sparse LWE fundamentally easier than sparse LWE? Specifically, to what extent can an algorithmic adversary exploit a bounded secret when the coefficient matrix is sparse?

The analysis of the hardness of small secret sparse LWE was also explicitly left as an open question in [JLS24]. We can also ask the same question for a small secret sparse LPN.

1.1 Our Contributions

We denote k-LIN $(\mathbb{Z}_q)$ for the space of collections of $\mathbb{Z}_q$-linear equations where the left-hand sides have exactly k non-zero coefficients. A k-LIN $(\mathbb{Z}_q)$ instance can be viewed as $(A, \boldsymbol{b}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$ where A is the coefficient matrix having k-sparse rows, and $\boldsymbol{b}$ is the vector of right-hand sides. For the distinguishing problem, we are supposed to determine whether $\boldsymbol{b}$ is random in $\mathbb{Z}_q^m$ or is planted with a secret solution $\boldsymbol{s} \in \mathbb{Z}_q^n$ satisfying $\|\boldsymbol{s}\|_\infty \leq B$ and some noise $\boldsymbol{e} \in \mathbb{Z}_q^m$, i.e. $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$. Typically, we assume that the noise in each equation is independently drawn from some noise distribution Ψ on $\mathbb{Z}_q$, i.e. $\boldsymbol{e} \leftarrow \Psi^m$. In this work, we consider two types of noise distributions:

- $\|\boldsymbol{e}\|_\infty \ll q$ with high probability; the discrete Gaussian distribution for LWE satisfies this.
- Ψ is uniform on $\mathbb{Z}_q \setminus \{0\}$, and has a different (higher) mass at 0. This is the distribution for LPN over $\mathbb{Z}_q$.

We obtain trade-offs between the sample complexity and runtime in our distinguishing algorithms.

The following theorem talks about a distinguisher for bounded planted noise:

Theorem 1.1. *(Informal version of Theorem 5.1) Suppose $\mathcal{I}$ is a k-LIN$(\mathbb{Z}_q)$ instance on n dimensions which is either*

- *random, or*
- *planted with a secret vector $\boldsymbol{s} \in \mathbb{Z}_q^n$ satisfying $\|\boldsymbol{s}\|_\infty \leq B$, noise $\boldsymbol{e}$ satisfying $\|\boldsymbol{e}\|_\infty \leq D$.*

Let $\lceil k/2 \rceil \leq l \leq n$ be a parameter, and $q \geq \Omega(Dl \log(Bn))$. There exists an algorithm to distinguish whether $\mathcal{I}$ is random or planted in $\widetilde{O}\left((Bn)^{l/2+k}\right)$ time, provided the number of samples is $m \geq (n/l)^k \cdot l \ln(Bn)$.

The LWE problem typically uses a discrete Gaussian with width $r < q/2$, denoted $\Psi_{\mathbb{Z}_q,r}$ as the noise distribution. It is known that the noise of dimension m ($\boldsymbol{e} \sim \Psi_{\mathbb{Z}_q,r}^m$) has a bound of $r\sqrt{T}$. Consequently, we will be able to use the above result for distinguishing sparse LWE instances.

The following theorem talks about a distinguisher when the noise is uniform over non-zero points in $\mathbb{Z}_q$, i.e. the LPN distribution:

Theorem 1.2. *(Informal version of Theorem 5.7) Suppose $\mathcal{I}$ is a k-LIN$(\mathbb{Z}_q)$ instance on n dimensions which is either*

- *random, or*
- *planted with a secret vector* $\boldsymbol{s} \in \mathbb{Z}_q^n$ *satisfying* $\|\boldsymbol{s}\|_\infty \leq B$*, noise* $\boldsymbol{e} \leftarrow \Psi^m$ *with* $\Psi(0) = 1 - \mu$*,* $\Psi(a) = \mu/(q-1)$ *for all* $a \in \mathbb{Z}_q \setminus \{0\}$*.*

Let $\lceil k/2 \rceil \leq l \leq n$ *be a parameter. There exists an algorithm to distinguish whether* $\mathcal{I}$ *is random or planted in* $\widetilde{O}\left(\left(1 - \frac{q\mu}{q-1}\right)^{-8/\varepsilon} (Bn/\varepsilon)^{l(1/2+\varepsilon)+k}\right)$ *time for some parameter* $\varepsilon \in (0, 0.5)$*, provided the number of samples is* $m \geq \left(1 - \frac{q\mu}{q-1}\right)^{-8/\varepsilon} \left(\frac{n}{l}\right)^k \cdot (l \log (Bn/\varepsilon))$*.*

We provide a high-level overview of our techniques in Section 2.

1.2 Related Work

1.2.1 Sparse k-XOR (k-LIN $(\mathbb{Z}_2)$)

Complexity of the k-LIN $(\mathbb{Z}_2)$ problem (popularly known as k-XOR problem) has been studied for over forty years in average-case complexity and has intimate connections to constraint satisfaction problems [Hås84; Gol11; AOW15; AL16; KMOW17; GKM22; HKM23; SOKB25; GHOS26]. Again, the objective has been to determine the computational complexity as a function of the sample complexity m. For $m \leq n$, the random and the planted noisy instances are statistically indistinguishable. Raghavendra, Rao and Schramm [RRS17] showed that sparse LPN over $\mathbb{Z}_2$ can be solved in $n^{O(l)}$ time when m is a quantity of the form $l \left(\frac{n}{l}\right)^{k/2} \log^{c(k)}(n)$ using a degree l "Sum of squares" method for $l \geq k/2$. An improved (and the best known) bound of the form $l \left(\frac{n}{l}\right)^{k/2} \log n$ for m with the same time complexity can be obtained using the "Kikuchi method" by Wein, Alaoui, and Moore [WAM19] (independently by Hastings [Has20]).[1]. More recently, Gupta, He, O'Donnell, and Singer [GHOS26] used the Kikuchi graph defined by [WAM19] and gave a new classical algorithm with a nearly quadratic speedup over the Kikuchi Method for the same sample complexity.

1.2.2 k-LIN (R)

Until recently, no optimal sample-runtime tradeoff results were known for k-LIN $(\mathbb{Z}_q)$ problem for $q > 2$.

Kocurek and Manohar [KM25] gave sample-runtime tradeoff algorithms for *refuting* semirandom[2] instances of k-LIN (R) when R is a field or an abelian group. In the task of refutation, the goal is to certify that a random instance does not have any assignment satisfying many constraints. They also establish a nearly matching lower bound for the required sample complexity for any algorithm falling within a "sum-of-squares" hierarchy. In addition, they showed existence of a *short linear dependent* set of rows in the LHS matrix A, which translates to existence of a short integer solution in A^{T}. We actually give algorithms to *find* such a short integer solution, in fact find a relaxed variant of the short integer solution, and then use it for the distinguishing task.

Recently, Agrawal, Bagchi and Kumar [ABK26] gave sample-runtime tradeoffs for the distinguishing task in k-LIN (R) where R is a finite commutative ring. They only solve the distinguishing task when the coefficient matrix A is random, whereas we solve it for semirandom A, i.e., arbitrary A with k-sparse rows. Moreover, they do not consider these problems with any assumption on the secret vector. Even if we compare our algorithm with theirs without the small secrets bound,

[1] The Kikuchi method is inspired from Kikuchi free energy, a marginal approximation to Gibbs free energy [Kik51]

[2] when the left-hand sides are arbitrary k-sparse

ours has a nearly quadratic speedup ($(qn)^{k+l(1/2+\varepsilon)}$ vs $(qn)^l$) than theirs when $q = O(n/l)$. For $q = \omega(n/l)$ too, we have a substantial speedup ($(qn)^{k+l(1/2+\varepsilon)}$ vs q^l).

1.3 Organization

In Section 2 we give a high level overview of our techniques. Section 3 introduces basic notations, the sparse LWE and LPN problems over higher modulus, standard results and the Kikuchi construction from [ABK26]. In Section 4 we solve a relaxed variant of the SIS problem for sparse matrices, which is the key ingredient for our attacks on k-LIN $(\mathbb{Z}_q)$. In Section 5 we utilize solutions of the relaxed SIS problem for distinguishing k-LIN $(\mathbb{Z}_q)$ instances.

2 Technical Overview

In this work, we are interested in solving the following problem: Given a k-LIN $(\mathbb{Z}_q)$ instance $(A, \boldsymbol{b}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$ where A has k-sparse rows, determine which of the two is the case:

- Random Instance: $\boldsymbol{b}$ is uniformly random in $\mathbb{Z}_q^m$
- Planted Instance: $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$ for unknown secret vector $\boldsymbol{s} \in \mathbb{Z}_q^n$, noise $\boldsymbol{e} \in \mathbb{Z}_q^m$ such that
  - $\boldsymbol{s} \in \mathbb{Z}_q^n$ satisfies $\|\boldsymbol{s}\|_\infty \le B$
  - $\boldsymbol{e} = (e_1, \ldots, e_m)^{\mathrm{T}}$ where $e_i \overset{\text{iid}}{\sim} \Psi$ for some noise distribution Ψ on $\mathbb{Z}_q$.

We will solve this for two kinds of noise distributions Ψ, one where $e_i \sim \Psi$ is bounded with high probability (like LWE), and the other which has uniform mass at $\mathbb{Z}_q \setminus \{0\}$ i.e., the LPN distribution noise. In both cases, we use reductions of the the LWE and LPN problems to the short integer solutions (SIS) problem. To the best of our knowledge, the reduction from LPN (over higher modulus) to SIS was not known before, which may be of independent interest. Then in Subsection 2.3 we give a high-level overview of how to attack the SIS problem when the instance matrix A has k-sparse rows.

2.1 Using Relaxed SIS for Bounded Noise

Given a LWE instance $(A, \boldsymbol{b}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$, a well-known observation [MR09] is the following: if we can find a non-trivial and *short* $\boldsymbol{c} \in \mathbb{Z}_q^m$ such that $\boldsymbol{c}^{\mathrm{T}} A = \boldsymbol{0}^{\mathrm{T}}$, we can hope to determine whether the instance is planted or random by examining $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} \in \mathbb{Z}_q$. This follows by observing that if $\boldsymbol{b}$ is random in $\mathbb{Z}_q^m$, then $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$ is also random in $\mathbb{Z}_q$ (as long as for some $i \in [m]$, c_i has a multiplicative inverse in $\mathbb{Z}_q$). In particular, $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right|$ [3] will have an expected value of $q/4$. On the other hand, if $\boldsymbol{b}$ is planted, say $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$, then $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} = \boldsymbol{c}^{\mathrm{T}}\boldsymbol{e}$. When $\boldsymbol{e}$ comes from a distribution that ensures that $\|\boldsymbol{e}\|_\infty$ is not too large with high probability (say, discrete Gaussian of sufficiently small width), we can get a bound on $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{e}\right| \le \|\boldsymbol{c}\|_1 \|\boldsymbol{e}\|_\infty$. Here the shortness assumption on $\boldsymbol{c}$ helps us: If $\|\boldsymbol{c}\|_1 \|\boldsymbol{e}\|_\infty$ is sufficiently smaller than $q/4$, we can distinguish the random case from the planted case with a decent success probability based on how small $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right|$ is. Again note that $\boldsymbol{c}$ should have some c_i entry to be a unit in $\mathbb{Z}_q$ to ensure that $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$ is uniform over $\mathbb{Z}_q$ in the random case. Let us call such a $\boldsymbol{c}$ non-trivial.[4]

[3] for $x \in \mathbb{Z}_q$, we assume $x \in \{-\lceil q/2\rceil + 1, \ldots, \lfloor q/2\rfloor\}$ while taking $|x|$

[4] In literature, we often take q to be prime, in which case this notion of non-trivial and non-zero are equivalent.

Notation 2.1. *We call $\boldsymbol{c} \in \mathbb{Z}_q^m$ non-trivial if there exists $i \in [m]$ such that c_i has a multiplicative inverse in $\mathbb{Z}_q$.*

Given $A \in \mathbb{Z}_q^{m \times n}$, the problem of finding a non-trivial and short $\boldsymbol{c} \in \mathbb{Z}_q^m$ such that $A^{\mathrm{T}}\boldsymbol{c} = \boldsymbol{0}$ (or equivalently $\boldsymbol{c}^{\mathrm{T}}A = \boldsymbol{0}^{\mathrm{T}}$) is known as the Short Integer Solution (SIS) problem. This problem itself is assumed to be computationally hard. Let us see how the small secret property for LWE (or any bounded noise variant) enables us to consider a relaxed version of the SIS problem[5] which will still suffice for our strategy to distinguish the random case from the planted case. In Subsection 2.3, we will discuss a high-level idea of how to attack this relaxed SIS problem for A^{T} when A has k-sparse rows.

Let us again examine $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right|$ in the planted case, this time without the restriction of $A^{\mathrm{T}}\boldsymbol{c} = \boldsymbol{0}$:

$$\begin{aligned}\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| &= \left|\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s} + \boldsymbol{c}^{\mathrm{T}}\boldsymbol{e}\right| \\ &\leq \left|\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right| + \left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{e}\right| \\ &\leq \left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \|\boldsymbol{s}\|_\infty + \|\boldsymbol{c}\|_1 \|\boldsymbol{e}\|_\infty\end{aligned}$$

If we have an appropriate bound on $\|\boldsymbol{s}\|_\infty$ then we don't need to find a $\boldsymbol{c}$ such that $A^{\mathrm{T}}\boldsymbol{c} = \boldsymbol{0}$, all we need is to ensure $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1$ is bounded and we will still have a low value of $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right|$ in the planted case. To realise this argument, we define a *relaxed* variant of SIS:

Definition 2.2. *Let $A \in \mathbb{Z}_q^{m \times n}$, let $R \geq 0$ be a relaxation parameter and let $T > 0$ be a shortness parameter. The R-relaxed T-short integer solution problem is to find a non-trivial $\boldsymbol{c} \in \mathbb{Z}_q^{m \times n}$ such that $\left\|A^T\boldsymbol{c}\right\|_1 \leq R$ and $\|\boldsymbol{c}\|_1 \leq T$.*

Now we address the question: What are acceptable values of R and T for an attack on this problem? Note that the relaxed problem has two parameters: a bound B on the secret, i.e., $\|\boldsymbol{s}\|_\infty \leq B$, and a bound D on the noise, i.e., $\|\boldsymbol{e}\|_\infty \leq D$. Putting the solution from Definition 2.2 into the calculation above we get that

$$\begin{aligned}\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| &\leq R \cdot \|\boldsymbol{s}\|_\infty + T \cdot \|\boldsymbol{e}\|_\infty \\ &\leq RB + TD\end{aligned}$$

If we are able to solve the relaxed SIS problem for $R \leq \frac{\varepsilon q}{8B}$ and $T \leq \frac{\varepsilon q}{8D}$, we will have $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| \leq \varepsilon q/4$ in the planted case, and thus we will be able to use the distinguishing argument described above.

2.2 Using Relaxed SIS for Uniform noise at non-zero points

The other noise distribution of interest to us is that of the LPN problem, which gives uniform mass at all non-zero points in $\mathbb{Z}_q$, say $\Psi(a) = \mu/(q-1)$ for $a \in \mathbb{Z}_q \setminus \{0\}$ and $\Psi(0) = 1 - \mu$. We can no longer use the above technique, since $\|\boldsymbol{e}\|_\infty$ can now be $\geq q/4$ with considerable probability. In fact, the distributions of $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$ for the planted and random cases can be statistically close when $\boldsymbol{c}$ has even moderately many non-zero entries. To mitigate these issues, we do two modifications from our previous approach:

[5] This problem is also known as the *Approximate Trapdoor* in the cryptographic literature [CGM19]. In this work, we prefer to use the term relaxed SIS.

- Instead of working with $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right|$, we work with a different transformation over $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$. It is still desirable to transform this random variable to $\mathbb{R}$ or $\mathbb{C}$, so that analysing random variables is easier. A natural choice is the canonical character function $\chi : \mathbb{Z}_q \to \mathbb{C}$ given by $\chi(a) := e^{2\pi\iota a/q}$. This is an injective group homomorphism, so it will preserve the statistical distance between the two cases of $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$. Moreover, due to the special nature of noise here (uniform over non-zero points), it will be amenable to analysis.

  When $\boldsymbol{b}$ is random over $\mathbb{Z}_q^m$, $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$ is random over $\mathbb{Z}_q$ and then $\mathbb{E}\left[\chi\left(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right)\right] = 0$. Let us now examine the case when $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$.

$$\chi\left(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right) = \chi\left(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right)\chi\left(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{e}\right) \;= \chi\left(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right)\prod_i \chi\left(c_i e_i\right)$$
$$\Rightarrow \mathbb{E}\left[\chi\left(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right)\right] = \chi\left(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right)\prod_i \mathbb{E}\left[\chi\left(c_i e_i\right)\right] \quad \text{(since } e_i\text{'s are independent)}$$

  Since e_i has uniform mass over non-zero points, $c_i e_i$ has uniform mass at all non-zero points in the subgroup $c_i\mathbb{Z}_q$ [6]. Figure 1 shows distributions of $\chi(e)$ and $\chi(3e)$ when the ambient ring is $\mathbb{Z}_9$.



(a) Distribution of $\chi(e)$ (b) Distribution of $\chi(3e)$

Figure 1: e is the LPN noise over $\mathbb{Z}_9$ with $\mu = 0.7$

  Moreover, it can be shown that for $c_i \neq 0$, $\mathbb{E}[\chi(c_i e_i)] = \mathbb{E}[\chi(e_i)] = \rho := 1 - \frac{q\mu}{q-1}$. Hence

$$\mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right] = \chi\left(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right)\rho^{\|\boldsymbol{c}\|_0} \;\geq \chi\left(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right)\rho^{\|\boldsymbol{c}\|_1}$$
$$\Rightarrow \mathrm{Re}\left(\mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right]\right) \geq \cos\left(2\pi\left(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right)/q\right)\rho^{\|\boldsymbol{c}\|_1} \;\geq \cos\left(2\pi RB/q\right)\rho^{T},$$

  i.e., the real part of $\mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right]$ is bounded away from 0 in the planted case, as opposed to the random case where it is 0. We can see this visually in Figure 1. In both cases the complex parts are symmetrically distributed around 0 and so cancel each other out. In the case of $\chi(3e)$ it is clear that the real part is non-zero.

[6] if c_i is a unit, then $c_i\mathbb{Z}_q = \mathbb{Z}_q$, else it is a proper subgroup

- For the noise parameter μ we are interested in, ρ^T may be too small for us to distinguish the two cases. So, instead of obtaining one solution to the relaxed SIS problem, we obtain a sufficiently large collection of solutions $\mathcal{C} = \{\boldsymbol{c} : \left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \leq R, \|\boldsymbol{c}\|_1 \leq T\}$ so that every two distinct $\boldsymbol{c}, \boldsymbol{c}' \in \mathcal{C}$ have supports not contained in each other.

We then hope to distinguish on the basis of the quantity: $\psi_{\mathcal{C}} = \sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})$. When $\boldsymbol{b}$ is random, $\mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right] = 0$ and hence $\mathbb{E}[\psi_{\mathcal{C}}] = 0$. When $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$ we have

$$\mathrm{Re}\left(\psi_{\mathcal{C}}\right) = \mathrm{Re}\left(\sum_{\boldsymbol{c}\in\mathcal{C}} \mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right]\right) \geq |\mathcal{C}| \cos\left(2\pi RB/q\right) \rho^T$$

So if we have a sufficiently large collection $\mathcal{C}$, we have a separation in the expected value of $\mathrm{Re}\left(\psi_{\mathcal{C}}\right)$ (and hence $|\psi_{\mathcal{C}}|$).

However we face a challenge for proving tail bounds for $|\psi_{\mathcal{C}}|$ in the two cases: the summands $\{\chi\left(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right)\}_{\boldsymbol{c}\in\mathcal{C}}$ need not be independent for the collection $\mathcal{C}$ (though they will be pairwise independent). Indeed, we only manage to obtain $\mathcal{C}$ so that any two $\boldsymbol{c}, \boldsymbol{c}'$ have non-overlaps in their supports; but the support of $\boldsymbol{c}$ can be contained in that of $\mathcal{C} \setminus \{\boldsymbol{c}\}$. This prevents us from using standard concentration inequalities. For the random case, we are able to use Chebyshev's inequality to show an upper bound on $\left|\sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right|$. For the planted case, doing the same is not helpful since the variance of $\psi_{\mathcal{C}}$ can be much larger than its expected value. Instead we show *anti-concentration* of $\left|\sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right|$ around 0, for which we use an analogous idea as [GHOS26] of developing anti-concentration results for a certain class of multi-linear polynomials.

2.3 Finding Relaxed SIS

We are left to describe how to actually find R-relaxed short integer solutions for A^{T}. For this, we utilize the Kikuchi Graph construction for A. The Kikuchi construction was introduced for the k-XOR (k-LIN $(\mathbb{Z}_2)$) problem by Wein, Alaoui, and Moore [WAM19] (and implicitly by Hastings [Has20]). More recently, it was extended to k-LIN $(\mathbb{Z}_q)$ for a finite commutative ring R by [KM25; ABK26]. These works used the spectral norm of adjacency matrix of the Kikuchi Graph to solve the problem. However, computing the spectral norm takes $\widetilde{\Omega}(N)$ time for N being the vertex size of the graph. With our approach, we will not only obtain a nearly quadratic speedup in runtime for the standard k-LIN $(\mathbb{Z}_q)$ distinguishing problem, but will also obtain a much higher speedup when the secret vector in the planted case is small ($\|\boldsymbol{s}\|_\infty \leq B$).

The vertices in the Kikuchi Graph $\mathcal{K}_l(A)$ defined by [ABK26] are l-sparse vectors in $\mathbb{Z}_q^n$ (for some parameter $k/2 \leq l \leq n$), and $\boldsymbol{u} \xrightarrow{\boldsymbol{a}_i} \boldsymbol{v}$ (resp. $\boldsymbol{u} \xrightarrow{-\boldsymbol{a}_i} \boldsymbol{v}$) is an edge if $\boldsymbol{v} - \boldsymbol{u} = \boldsymbol{a}_i$ (resp. $\boldsymbol{v} - \boldsymbol{u} = -\boldsymbol{a}_i$) for some row $\boldsymbol{a}_i \in A$. Notice that $\boldsymbol{c}^{\mathrm{T}}A$ is related to a (directed) walk in the graph; indeed $\boldsymbol{c}^{\mathrm{T}}A$ is a linear combination of rows of A, and so is the addition of all edge labels of a walk in $\mathcal{K}_l(A)$. If the walk is $W = (\boldsymbol{w}_S, \ldots, \boldsymbol{w}_T)$, then the latter is $\boldsymbol{w}_T - \boldsymbol{w}_S$. Moreover, $\|\boldsymbol{c}\|_1$ is upper bounded by the length of the walk. This reduces the problem of R-relaxed SIS for A^{T} to the problem of finding a walk W in $\mathcal{K}_l(A)$ such that $\|\boldsymbol{w}_T - \boldsymbol{w}_S\|_1 \leq R$. Our strategy of finding such a walk is straightforward: we sample sufficiently many (say L) random walks of length T from a vertex $\boldsymbol{w}_0$ so that there are two walks $W_i = \left(\boldsymbol{w}_0, \boldsymbol{w}_1^{(i)}, \ldots, \boldsymbol{w}_T^{(i)}\right)$, $W_j = \left(\boldsymbol{w}_0, \boldsymbol{w}_1^{(j)}, \ldots, \boldsymbol{w}_T^{(j)}\right)$ satisfy $\left\|\boldsymbol{w}_T^{(i)} - \boldsymbol{w}_T^{(j)}\right\| \leq R$. We will then output the $2T$-length walk $W = W_i \circ W_j^{\text{reversed}}$ as the desired walk. We will show that we need to sample roughly $L = \widetilde{O}\left(\sqrt{\frac{N}{R^l}}\right) = \widetilde{O}\left(\left(\frac{qn}{R}\right)^{l/2}\right)$ many walks to achieve the above with a good probability of success. This is the dominating factor in our runtime.

We still face a couple of more challenges. The first is to ensure that $\boldsymbol{c}$ obtained from such a walk is non-trivial. The second is to produce a collection of multiple distinct solutions $\mathcal{C}$, as mentioned in above Subsection for the sparse LPN problem. We ensure both by arguing the following: With high probability, the set of edge labels which occur in a 'good' walk exactly once is not contained in any fixed label set of size at most $2T$. This ensures the first challenge: that a $\boldsymbol{c}$ obtained from a good walk is non-trivial since the set of edge labels which occur once is non-empty. To produce multiple solutions, we repeatedly produce good walks whose set of edge labels are not contained in the edge labels of any previous walk.

3 Preliminaries

3.1 Notations

Let $\mathbb{N}, \mathbb{R}, \mathbb{C}$ denote the set of naturals, reals, complex numbers respectively. Let $\iota = \sqrt{-1}$. For $n \in \mathbb{N}$, $[n] := \{1, 2, \ldots, n\}$. For a set S, $\mathrm{Unif}(S)$ denotes the uniform distribution on S. For a probability distribution $\mathcal{D}$ over S, $\mathrm{Supp}(\mathcal{D}) \subseteq S$ denotes its support, i.e. the subset where $\mathcal{D}$ has non-zero probability mass. For $a \in \mathbb{C}$, $\mathrm{Re}(a)$ denotes its real part, a^* denotes its complex conjugate.

We denote vectors by emboldened small letters (e.g. $\boldsymbol{x}$) and matrices by capital letters (e.g. A). For a vector $\boldsymbol{x}$, unless stated otherwise, x_j denotes the j^{th} coordinate of $\boldsymbol{x}$. $\mathbf{0}$ denotes a null vector. For $t \leq n$, $(\mathbb{Z}_q^n)_t$ denotes the subset of $\mathbb{Z}_q^n$ consisting of all vectors with exactly t coordinates in $\mathbb{Z}_q \setminus \{0\}$.

For a function $f : \mathbb{N} \to \mathbb{N}$, the notation $\widetilde{O}(f(n))$ denotes $O(f(n) \cdot \mathtt{poly}(\log f(n)))$ for some polynomial $\mathtt{poly}(\cdot)$.

For a directed graph $\mathcal{G}$, let $\text{in-deg}_{\mathcal{G}}(T)$, $\text{out-deg}_{\mathcal{G}}(T)$ denote the number of incoming, outgoing edges of a vertex T of $\mathcal{G}$ respectively. We may skip writing $\mathcal{G}$ in the subscript when the graph is clear in the context.

3.2 'Norm' operation over $\mathbb{Z}_q^n$

For $\boldsymbol{x} \in \mathbb{Z}_q^n$, $1 \leq p \leq \infty$, we consider $\|\boldsymbol{x}\|_p$ which we define as the standard ℓ_p norm on $\boldsymbol{x}$ while viewing the entries of $\boldsymbol{x}$ in $\{-\lceil q/2 \rceil + 1, \ldots, \lfloor q/2 \rfloor\}$. In addition, we use the notation $\|\boldsymbol{x}\|_0$ to denote the number of non-zero entries (the Hamming weight) in $\boldsymbol{x}$.

Definition 3.1. *For $\boldsymbol{x} \in \mathbb{Z}_q^n$, let $\widetilde{\boldsymbol{x}} \in \mathbb{Z}^n$ be such that $\widetilde{x}_i \in \{-\lceil q/2 \rceil + 1, \ldots, \lfloor q/2 \rfloor\}$ and $\widetilde{\boldsymbol{x}}$ (mod q) = $\boldsymbol{x}$. Then $\|\boldsymbol{x}\|_p := \|\widetilde{\boldsymbol{x}}\|_p$.*

This is not a well-defined norm operation since it does not satisfy the absolute homogeneity property; $\|\lambda \boldsymbol{x}\|_p \neq |\lambda| \, \|\boldsymbol{x}\|_p$ in general. However it satisfies the other two norm properties, i.e. it has positive definiteness ($\|\boldsymbol{x}\|_p = 0 \iff \boldsymbol{x} = \mathbf{0}$) and triangle inequality. Let us see a proof of the latter.

Observation 3.2. *For $1 \leq p \leq \infty$, $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{Z}_q^n$, $\|\boldsymbol{u} + \boldsymbol{v}\|_p \leq \|\boldsymbol{u}\|_p + \|\boldsymbol{v}\|_p$.*

Proof. It suffices to show that for $u_j, v_j \in \mathbb{Z}_q$, $|u_j + v_j| \leq |u_j| + |v_j|$. Without loss of generality, let us assume $u_j, v_j \in \{-\lceil q/2 \rceil + 1, \ldots, \lfloor q/2 \rfloor\}$.

- Suppose $-\lceil q/2 \rceil + 1 \leq u_j + v_j \leq \lfloor q/2 \rfloor$. Then the claim holds by the usual triangle inequality.
- Suppose $u_j + v_j < -\lceil q/2 \rceil + 1$. Then $|u_j + v_j| = u_j + v_j + q \leq \lfloor q/2 \rfloor$ and $|u_j| + |v_j| = -(u_j + v_j) \geq \lfloor q/2 \rfloor$.

- Suppose $u_j + v_j > \lfloor q/2 \rfloor$. Then $|u_j + v_j| = q - u_j - v_j \leq \lfloor q/2 \rfloor$ and $|u_j| + |v_j| = u_j + v_j > \lfloor q/2 \rfloor$.

□

As a consequence, the operation also enjoys Hölder's inequality, since the proof of this inequality only requires positive definiteness and triangle inequality conditions of the norm (see, e.g., Chapter 6 in [Fol99]).

Observation 3.3. *Let $p, p' \in [1, \infty]$ satisfy $\frac{1}{p} + \frac{1}{p'} = 1$. For $\boldsymbol{u}, \boldsymbol{v} \in \mathbb{Z}_q^n$, $\sum_{j=1}^n |u_j v_j| \leq \|\boldsymbol{u}\|_p \|\boldsymbol{v}\|_{p'}$.*

3.3 Probability and Concentration Inequalities

The following fact follows from McLaurin inequality (see eg., §2.22 of [HLP52]) and says that the probability of collision is minimized when points are picked uniformly at random.

Fact 3.1. *For every finite set S, distribution $\mathcal{D}$ on S, and $k \in \mathbb{N}$,*

$$\Pr_{X_1,\ldots,X_L \overset{iid}{\sim} \mathcal{D}} [\exists \text{ collision among } X_1, \ldots, X_L] \geq$$

$$\Pr_{X_1,\ldots,X_L \overset{iid}{\sim} Unif(S)} [\exists \text{ collision among } X_1, \ldots, X_L] = 1 - \prod_{j=1}^{L-1} \left(1 - \frac{j}{|S|}\right)$$

Let us state tail bound results for random variables (both scalar and matrix) which we will use in this work.

Theorem 3.2. *(Chernoff Bound) Let $X = \sum_{i=1}^n X_i$ be the sum of independent Bernoulli random variables with $E[X] = \mu$. Then*

$$\forall\ \delta \geq 0,\ \Pr[X \geq (1+\delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2+\delta}}$$
$$\forall\ 0 \leq \delta \leq 1,\ \Pr[X \leq (1-\delta)\mu] \leq e^{-\frac{\delta^2 \mu}{2}}$$

Theorem 3.3. *(Hoeffding Inequality) Let $X_1, \ldots, X_n$ be independent and identically distributed random variables on $[a, b]$, let $\overline{X} = \frac{1}{n} \sum_{i \in [n]} X_i$, and let $\mu := \mathbb{E}[\overline{X}]$. Then*

$$\Pr\left[\overline{X} \leq \mu - t\right] < \exp\left(-\frac{2nt^2}{(b-a)^2}\right)$$
$$\Pr\left[\overline{X} \geq \mu + t\right] < \exp\left(-\frac{2nt^2}{(b-a)^2}\right)$$

Theorem 3.4. *(Azuma-Hoeffding Inequality, see, e.g., Chapter 1 [BLM13]) Let $\{X_k\}_{k=0}^n$ be a martingale (or supermartingale) with respect to a filtration $\{\mathcal{F}_k\}$, meaning $\mathbb{E}[X_k \mid \mathcal{F}_{k-1}] \leq X_{k-1}$.*

Suppose there exist constants $c_k > 0$ such that the differences are bounded with probability 1:

$$|X_k - X_{k-1}| \leq c_k \quad \text{for all } k = 1, \ldots, n \tag{1}$$

Then, for any $t > 0$, the following upper tail inequality holds:

$$P(X_n - X_0 \leq -t) \leq \exp\left(-\frac{t^2}{2\sum_{k=1}^n c_k^2}\right) \tag{2}$$

3.4 k-LIN $(\mathbb{Z}_q)$

We now state definitions and problems related to sparse system of linear equations over $\mathbb{Z}_q$. In this subsection, we often borrow notations and results from [ABK26].

Definition 3.4. *A k-LIN$(\mathbb{Z}_q)$ instance $\mathcal{I} = \{(\boldsymbol{a}_i, b_i)_{i\in[m]}\}$ consists of vectors $\boldsymbol{a}_i \in (\mathbb{Z}_q^n)_k$ and $b_i \in \mathbb{Z}_q$.*

For a k-LIN $(\mathbb{Z}_q)$ instance $\mathcal{I}$, we view $(\boldsymbol{a}, b) \in \mathcal{I}$ as a $\mathbb{Z}_q$-linear homogeneous equation in n variables $x_1, \ldots, x_n$: $\sum_{j\in[n]} a_j x_j = b$.

Now we will define random and planted versions of k-LIN $(\mathbb{Z}_q)$. Note that (unlike [ABK26]) in these definitions, we do not assume any distribution on the left-hand sides; they are arbitrary k-sparse vectors. In this case, the instance is said to be *semirandom*. However, in the rest of the discussion, we will drop the semirandom term for brevity.

Definition 3.5. *(Semirandom Random instance) A k-LIN$(\mathbb{Z}_q)$ instance $\mathcal{I} = \{(\boldsymbol{a}_i, b_i)_{i\in[m]}\}$ is said to be a random instance if $b_i \xleftarrow{iid} Unif(\mathbb{Z}_q)$.*

Definition 3.6. *(Semirandom Planted instance) A k-LIN$(\mathbb{Z}_q)$ instance $\mathcal{I} = \{(\boldsymbol{a}_i, b_i)_{i\in[m]}\}$ is said to be a planted instance with respect to secret vector $\boldsymbol{s} \in \mathbb{Z}_q^n$ and noise distribution Ψ on $\mathbb{Z}_q$ if for each $i \in [m]$, $b_i = \langle \boldsymbol{a_i}, \boldsymbol{s}\rangle + e_i$ where $e_i \xleftarrow{iid} \Psi$.*

3.4.1 Sparse LWE and Sparse LPN

We first recall the discrete Gaussian distribution on the integers and $\mathbb{Z}_q$.

Definition 3.7. *The discrete Gaussian distribution on the integers with width r, denoted $\Psi_{\mathbb{Z},r}$, is the distribution having probability mass at $x \in \mathbb{Z}$ proportional to $\exp\left(-\pi x^2/r^2\right)$.*

Definition 3.8. *The discrete Gaussian distribution on $\mathbb{Z}_q$ with width r, denoted $\Psi_{\mathbb{Z}_q,r}$ is the distribution having probability mass at $j \in \mathbb{Z}_q$ proportional to $\sum_{x\in q\mathbb{Z}+j} e^{-\pi x^2/r^2}$.*

The following observations state that the ℓ_2 norm of a vector sampled from the discrete Gaussian distribution is bounded with high probability.

Observation 3.9. *(Lemma 1.5(i) in [Ban93]) Let $n \in \mathbb{N}$ and $r > 0$ be the width parameter. If a noise vector $\mathbf{e} \in \mathbb{Z}^n$ is sampled from $(\Psi_{\mathbb{Z},r})^n$ (each coordinate is independently sampled from $\Psi_{\mathbb{Z},r}$), then its Euclidean norm is bounded by $r\sqrt{n}$ with overwhelming probability:*

$$\Pr_{\mathbf{e}\sim\left(\Psi_{\mathbb{Z},r}\right)^n}\left[\|\mathbf{e}\|_2 > r\sqrt{n}\right] \le 2^{-n+1} \tag{3}$$

If $q > r\sqrt{n}$, we will get a similar result for $\mathbf{e} \sim \left(\Psi_{\mathbb{Z}_q,r}\right)^n$.

Observation 3.10. *Let $n \in \mathbb{N}$ and $r > 0$ be the width parameter. Let $q > r\sqrt{n}$ If a noise vector $\mathbf{e} \in \mathbb{Z}^n$ is sampled from $\left(\Psi_{\mathbb{Z}_q,r}\right)^n$ (each coordinate is independently sampled from $\Psi_{\mathbb{Z}_q,r}$), then its Euclidean norm is bounded by $r\sqrt{n}$ with overwhelming probability:*

$$\Pr_{\mathbf{e}\sim\left(\Psi_{\mathbb{Z}_q,r}\right)^n}\left[\|\mathbf{e}\|_2 > r\sqrt{n}\right] \le 2^{-n+2} \tag{4}$$

The k-sparse LWE problem was first introduced in [JLS24]. We add a parameter B which is the ℓ_∞ bound for the secret vector $\boldsymbol{s}$ in the planted case.

Definition 3.11. *For an integer dimension n, sparsity parameter k, sample complexity m, modulus q, noise parameter r and secret bound parameter B, we define the k-LWE$_{n,m,q,r,B}$ problem as a distinguishing problem that is given m samples from either of two k-LIN$(\mathbb{Z}_q)$ distributions, and is required to guess the distribution. The first distribution is a planted k-LIN$(\mathbb{Z}_q)$ distribution with secret vector $\boldsymbol{s}$ satisfying $\|\boldsymbol{s}\|_\infty \leq B$ and noise distribution $\Psi_{\mathbb{Z}_q,r}$ and the second distribution is a random k-LIN$(\mathbb{Z}_q)$ distribution.*

We also define a sparse and small secret version of the Learning Parity with Noise (LPN) problem for higher moduli, the standard version of which was defined in [ADINZ17b].

Definition 3.12. *For an integer dimension n, sparsity parameter k, sample complexity m, modulus q, noise parameter μ and secret bound parameter B, we define the k-LPN$_{n,m,q,\mu,B}$ problem as a distinguishing problem that is given m samples from either of two k-LIN$(\mathbb{Z}_q)$ distributions, and is required to guess the distribution. The first distribution is a planted k-LIN$(\mathbb{Z}_q)$ distribution with secret vector $\boldsymbol{s}$ satisfying $\|\boldsymbol{s}\|_\infty \leq B$ and noise distribution $\xi_{\mathbb{Z}_q,\mu}$ on $\mathbb{Z}_q$ where $\xi_{\mathbb{Z}_q,\mu}(0) = 1-\mu$, $\forall j \in \mathbb{Z}_q \setminus \{0\}$ $\xi_{\mathbb{Z}_q,\mu}(j) = \mu/(q-1)$. The second distribution is a random k-LIN$(\mathbb{Z}_q)$ distribution.*

3.4.2 The Kikuchi Construction for k-LIN $(\mathbb{Z}_q)$

In [ABK26], for a ring R, a graph corresponding to a k-LIN (R) instance was defined, where the signs of the edges came from the right-hand sides of the instance. For this work, we can skip talking about the signs, so we present their Kikuchi construction only using the left-hand sides of the instance.

Construction 3.13. *Let $A \in \mathbb{Z}_q^{m\times n}$ consist of rows in $\left(\mathbb{Z}_q^n\right)_k$. Let $\lceil k/2 \rceil \leq l \leq n$ be a parameter. $\mathcal{K}_l(A)$ is a directed graph over vertices indexed by $\left(\mathbb{Z}_q^n\right)_l$. For every row $\boldsymbol{a}$ of A, $\mathcal{K}_l(A)$ has edges $\boldsymbol{u} \xrightarrow{\boldsymbol{a}} \boldsymbol{v}$ and $\boldsymbol{v} \xrightarrow{-\boldsymbol{a}} \boldsymbol{u}$ for all $\boldsymbol{u}, \boldsymbol{v} \in \left(\mathbb{Z}_q^n\right)_l$ which satisfy $\boldsymbol{v} - \boldsymbol{u} = \boldsymbol{a}$.*

We wish to compute the number of edges in $\mathcal{K}_l(A)$ yielded by one row in A. Let us rewrite Construction 3.13 in a way that is more amenable to computing this.

Construction 3.14. *(Equivalent to Construction 3.13) For every row $\boldsymbol{a}$ in A, $\mathcal{K}_l(A)$ has all edges $\boldsymbol{u} \xrightarrow{\boldsymbol{a}} \boldsymbol{v}$, $\boldsymbol{v} \xrightarrow{-\boldsymbol{a}} \boldsymbol{u}$ which satisfy the following:*

1. *For every $j \in [n]$ such that $a_j = 0$, exactly one of the following holds:*
   - *(a)* $u_j = v_j = 0$.
   - *(b)* $u_j = v_j \neq 0$.
2. *For every $j \in [n]$ such that $a_j \in \mathbb{Z}_q \setminus \{0\}$, exactly one of the following holds:*
   - *(a)* $u_j = -a_j, v_j = 0$.
   - *(b)* $u_j = 0, v_j = a_j$.
   - *(c)* $u_j, v_j \in \mathbb{Z}_q \setminus \{0\}$ *such that* $v_j - u_j = a_j$.

Let $r_{1a}, r_{1b}, r_{2a}, r_{2b}, r_{2c}$ denote the number of indices on which conditions 1(a), 1(b), 2(a), 2(b), 2(c) are applied respectively. It can be seen that the following must be satisfied:

$$r_{1a} + r_{1b} + r_{2a} + r_{2b} + r_{2c} = n \quad \text{(total indices)} \tag{5}$$
$$r_{1b} + r_{2a} + r_{2c} = l \quad \text{(non-zero entries of } \boldsymbol{u}\text{)} \tag{6}$$
$$r_{1b} + r_{2b} + r_{2c} = l \quad \text{(non-zero entries of } \boldsymbol{v}\text{)} \tag{7}$$
$$r_{2a} + r_{2b} + r_{2c} = k \quad \text{(non-zero entries of } \boldsymbol{a}\text{)} \tag{8}$$

If $r_{2c} = r \in \{0, 1, \ldots, k\}$, we obtain $r_{2a} = r_{2b} = (k-r)/2$ from (6), (7), (8), $r_{1b} = l - (k+r)/2$ from (6), $r_{1a} = n - l - (k-r)/2$ from (5). For indices j of condition 1(b), we have $q-1$ choices of (u_j, v_j), and for indices j of condition 2(c), we have $q-2$ choices of (u_j, v_j).

Thus we get the number of constraints for each original constraint as: [7]

$$\theta^{(n,l,k)} = 2\sum_{r=0}^{t} \binom{k}{r} (q-2)^r \binom{k-r}{(k-r)/2} \binom{n-k}{l-(k+r)/2} (q-1)^{l-(k+r)/2} \tag{9}$$

Note that in (9), only the terms where r has same parity as k are positive.

We can obtain two lower bounds for $\theta^{(n,l,k)}$. One of them will come from the summand of $r = 0$ or $r = 1$, depending on whether k is even or odd, respectively.

$$\text{Assuming } \lceil k/2 \rceil \le l \le n - \lfloor k/2 \rfloor, \quad \theta^{(n,l,k)} \ge \frac{2^{k-1}}{(k+1)} \binom{n-k}{l-\lceil k/2 \rceil} (q-1)^{l-\lfloor k/2 \rfloor} \tag{10}$$

The other bound will come from the $r = k$ summand term.

$$\text{Assuming } k \le l \le n, \quad \theta^{(n,l,k)} \ge 2^{1-k} \binom{n-k}{l-k} (q-1)^l \tag{11}$$

Depending on whether q is small or large, we will use these bounds for our analysis.

4 Solving a relaxed SIS problem using Kikuchi Graphs

In this section, we will attack the R-relaxed T-short integer solution problem (Definition 2.2) for the instance A^{T} when A has k-sparse rows. For this, we will utilize the construction of Kikuchi Graph of A. In particular, we will relate the required solution with a certain kind of walk in the Kikuchi graph of A. In the context of this work, a walk in a directed graph $\mathcal{G} = (V, E)$ is only in the direction of the edges; i.e. $W = (w_0, w_1, \ldots, w_T)$ is a walk only if $(w_{t-1}, w_t) \in E \ \forall t \in [T]$.

Construction 4.1. *Suppose $A \in \mathbb{Z}_q^{m \times n}$ has k-sparse rows, and $\boldsymbol{a}_i$ denotes its i^{th} row. Let $W = (\boldsymbol{u}_0, \boldsymbol{u}_1, \ldots, \boldsymbol{u}_T)$ be a walk in $\mathcal{K}_l(A)$. Construct $\boldsymbol{c} = (c_i)_{i \in [m]}$ as follows:*

- $\boldsymbol{c} = \boldsymbol{0}$ *initially.*
- *For $r \in [T]$:*
  - *if $\boldsymbol{u}_r - \boldsymbol{u}_{r-1} = \boldsymbol{a}_i$ for some $i \in [m]$ then add 1 to c_i*
  - *else if $\boldsymbol{u}_r - \boldsymbol{u}_{r-1} = -\boldsymbol{a}_i$ then add -1 to c_i.*

Observation 4.2. *After applying Construction 4.1, $\|\boldsymbol{c}\|_1 \le T$*

Proof. For each edge $\boldsymbol{u}_{r-1} \to \boldsymbol{u}_r$, we add or subtract 1 in a coordinate of $\boldsymbol{c}$, which increases its ℓ_1 norm by at most 1. □

Observation 4.3. *After applying Construction 4.1, $A^T \boldsymbol{c} = \boldsymbol{u}_T - \boldsymbol{u}_0$.*

[7] in this equation we use the convention that if s is not a non-negative integer or $s > t$, then $\binom{t}{s} = 0$

Proof. c_i is the count of number of times $\boldsymbol{a}_i = \boldsymbol{u}_r - \boldsymbol{u}_{r-1}$ occurs minus the number of times $-\boldsymbol{a}_i = \boldsymbol{u}_r - \boldsymbol{u}_{r-1}$ occurs in the walk $W = (\boldsymbol{u}_0, \boldsymbol{u}_1, \ldots, \boldsymbol{u}_T)$. Thus

$$\begin{aligned} A^{\mathrm{T}}\boldsymbol{c} &= \sum_{i\in[m]} c_i \boldsymbol{a}_i \\ &= (\boldsymbol{u}_1 - \boldsymbol{u}_0) + \ldots + (\boldsymbol{u}_T - \boldsymbol{u}_{T-1}) \quad = \boldsymbol{u}_T - \boldsymbol{u}_0 \end{aligned}$$

□

By Observation 4.2 and Observation 4.3, if we could obtain a walk of length T so that its end points are R apart in the ℓ_1 norm, we solve the R-relaxed T-short integer solution problem.

We devote the rest of this section to finding such a walk. There is one more subtle but 'non-trivial' aspect, ensuring that $\boldsymbol{c}$ obtained from a walk is *non-trivial*. To ensure this, we will define a goodness notion on walks and obtain good walks.

In Subsection 4.1 we set up some required tools and results on graphs. In Subsection 4.2 we describe the algorithms to sample such walks. We defer analyzing their correctness and complexity to Subsection 4.3.

4.1 Useful Definitions and Results on Graphs

Before venturing into solving the relaxed SIS problem (on A^{T} with A having k-sparse rows), let us set up some useful properties on certain directed graphs that we will utilize in this section.

Definition 4.4. *For a directed graph $\mathcal{G} = (V, E)$, let ν denote the* out-degree distribution *over V given by $\nu(v) \propto$ out-deg(v) for all $v \in V$.*

Definition 4.5. *The random walk distribution of length T from w_0, denoted walk$^T_{w_0}$, is a distribution over walks of length T with initial vertex w_0 defined as follows: To sample $W = (w_0, w_1, \ldots, w_T) \sim$ walk$^T_{w_0}$, iteratively sample w_{t+1} to be a random neighbor of w_t for $1 \le t < T$.*

It is well known that in undirected graphs, the marginal distribution of w_t in above definition is simply ν, i.e. ν is the stationary distribution with respect to the Markov process of random walks. For directed graphs, we can only establish this when in-deg$(v) =$ out-deg(v) $\forall v \in V$.

Observation 4.6. *Let $\mathcal{G} = (V, E)$ be a directed graph such that for every vertex $v \in V$, in-deg$(v) =$ out-deg(v). If $w_0 \sim \nu$ and $W = (w_0, w_1, \ldots, w_T) \sim$ walk$^T_{w_0}$, then the marginal distribution of w_t is ν (for every $1 \le t \le T$).*

Proof. Let us assume by inductive hypothesis that $w_{t-1} \sim \nu$ for some $t > 0$.

$$\begin{aligned} \Pr[w_t] &= \sum_{(w_{t-1}, w_t)\in E} \Pr[w_t | w_{t-1}] \Pr[w_{t-1}] \\ &= \sum_{(w_{t-1}, w_t)\in E} \frac{1}{\text{out-deg}(w_{t-1})} \frac{\text{out-deg}(w_{t-1})}{|E|} \quad = \frac{\text{in-deg}(w_t)}{|E|} = \frac{\text{out-deg}(w_t)}{|E|} \end{aligned}$$

□

As mentioned in Definition 2.2, the walk we derive $\boldsymbol{c}$ from should yield a non-trivial $\boldsymbol{c}$. One simple way of ensuring this is to show that the walk contains some particular $\boldsymbol{a}_i$ only once. This will directly mean that $c_i = \pm 1$ and hence non-trivial. We will show that we can select such walks with good probability. To do this we will execute the following recipe:

- We will first show that for directed graphs with the in-degree equal to out-degree at each vertex, a random walk begun from the stationary distribution will always visit a good proportion of vertices whose degree is comparable to the average degree. (Observation 4.9 and Lemma 4.1.)

- We will then show that the Kikuchi graph $\mathcal{K}_l(A)$ has the property that there are at most two edges corresponding to each $\boldsymbol{a}_i \in \mathbb{Z}_q^m$ of A incoming or outgoing at each vertex. (Observation 4.11.)
- Putting these two facts together we can see intuitively that whenever a random walk in $\mathcal{K}_l(A)$ visits a vertex of high degree it has the chance of moving through an edge label $\boldsymbol{a}_i \in \mathbb{Z}_q^m$ of A that it has not used before, since each such edge label is represented very few times at each vertex. In fact we can take this intuition forward and show that a random walk will have a good number of edges corresponding to constraints that are used only a single time through the walk, which is what we require. (Lemma 4.2)

We now begin this recipe by defining a 'goodness' property of a walk in a directed graph $\mathcal{G}$. Intuitively, a walk is good if it does not have too many vertices of low degree.

Definition 4.7. *Let $\mathcal{G} = (V, E)$ be a directed graph. Let $\Delta = \frac{|E|}{|V|}$ be the average out-degree of the graph. We say that a vertex $v \in V$ is bad if out-deg$(v) < 3\Delta/16$. For a directed walk $W = (w_0, \ldots, w_T)$ in $\mathcal{G}$, let Bad$(W) := |\{0 \le i < T : w_i \text{ is bad}\}|$. We say that W is good if Bad$(W) < T/4$.*

Remark 4.8. *The constants* $3/16$ *and* $1/4$ *in Definition 4.7 used to define* bad *and* good *respectively can be replaced by any suitable constants $\beta \in (0, 1], \gamma \in [\beta, 0.5)$ respectively.*

As indicated, the machinery for good walks is built to obtain non-trivial relaxed short integer solutions. Intuitively, good walks will comprise a 'diverse' set of edges so that the labels don't all 'cancel out' while applying Construction 4.1.

The following results show that random walks are good with a decent probability.

Observation 4.9. *If $\mathcal{G}$ has average out-degree Δ, then*

$$\Pr_{v \sim \nu}[\textit{out-deg}(v) < \beta\Delta] < \beta.$$

Proof. Let $S := \{v : \text{out-deg}(v) < \beta\Delta\}$. Then

$$\Pr_{v \sim \nu}[\text{out-deg}(v) < \beta\Delta] = \frac{\sum_{v \in S} \text{out-deg}(v)}{|E|} < \frac{\sum_{v \in S} \beta\Delta}{|E|} = \frac{\beta\Delta|S|}{|E|} \le \frac{\beta\Delta|V|}{|E|} = \beta \qquad \square$$

Lemma 4.1. *Let $\mathcal{G} = (V, E)$ be a directed graph such that for every vertex $v \in V$, in-deg(v) = out-deg(v). Suppose $w_0 \sim \nu$ and $W \sim \textit{walk}_{w_0}^T$, then*

$$\Pr[W \textit{ is good}] > 1/4$$

Proof. By Observation 4.6 and Observation 4.9, we get that $\Pr[w_t \text{ is bad}] < 3/16$ for each $w_t \in W$, where the probability is over both $w_0 \sim \nu$ and $W \sim \text{walk}_{w_0}^T$. Thus

$$\mathop{\mathbb{E}}_{w_0 \sim \nu, W \sim \text{walk}_{w_0}^T}[\text{Bad}(W)] < 3T/16$$

By Markov's inequality, $\Pr[W \text{ is good}] = \Pr[\text{Bad}(W) < T/4] > 1/4$. $\square$

The Kikuchi graph has labels associated to each edge; the label for an edge $(\boldsymbol{u}, \boldsymbol{v})$ is $\boldsymbol{v} - \boldsymbol{u}$. To capture this, let us introduce edge labels/colors in our directed graph model.

Definition 4.10. *Let $\mathcal{G} = (V, E)$ be a directed graph and $\Gamma : E \to [m]$ be a function. For $t \in \mathbb{N}$, Γ is a t-deficient edge coloring if every vertex has at most t incoming and at most t outgoing edges of every color, i.e.*

$$\forall v \in V, \forall a \in [m] : |\{w : e = (v, w) \in E, \Gamma(e) = a\}| \leq t, \ |\{w : e = (w, v) \in E, \Gamma(e) = a\}| \leq t$$

We then say that $\mathcal{G} = (V, E, \Gamma)$ is a t-deficient edge colored directed graph.

Now let us relate the Kikuchi Graph $\mathcal{K}_l(A)$ with a t-deficient edge colored directed graph.

Observation 4.11. *Let $A \in \mathbb{Z}_q^{m \times n}$ have k-sparse rows and $\mathcal{K}_l(A)$ be its Kikuchi graph for a parameter $l \geq k/2$. $\mathcal{K}_l(A)$ satisfies the property that every edge $e = (\boldsymbol{u}, \boldsymbol{v})$ has a corresponding reverse edge $e^{\text{rev}} = (\boldsymbol{v}, \boldsymbol{u})$. Further, for every edge $(\boldsymbol{u}, \boldsymbol{v})$ satisfying $\boldsymbol{v} - \boldsymbol{u} = \boldsymbol{a}_i$ for some row $\boldsymbol{a}_i \in \mathbb{Z}_q^m$ of A, let Γ map the edges $(\boldsymbol{u}, \boldsymbol{v}), (\boldsymbol{v}, \boldsymbol{u})$ to color i. Then $\mathcal{K}_l(A)$ is a 2-deficient edge colored directed graph.*

Proof. The first property follows directly from Construction 3.13. For any vectors $\boldsymbol{u} \in (\mathbb{Z}_q^n)_l$, $\boldsymbol{a}_i \in (\mathbb{Z}_q^n)_k$ we can have at most one $\boldsymbol{v} \in (\mathbb{Z}_q^n)_l$ such that $\boldsymbol{v} - \boldsymbol{u} = \boldsymbol{a}_i$ and at most one $\boldsymbol{v} \in (\mathbb{Z}_q^n)_l$ such that $\boldsymbol{v} - \boldsymbol{u} = -\boldsymbol{a}_i$. This means that any vertex $\boldsymbol{u}$ has at most two outgoing edges with color i, and similarly at most two incoming edges with color i, which establishes that $\mathcal{K}_l(A)$ is 2-deficient edge colored. □

As per our recipe, our plan now is to show that we can select walks with single frequency edge labels with good probability. In fact, we will show that we can find multiple walks such that the set of singleton edges of any walk is not contained in the label set of any other walk. The latter will be helpful for us for the LPN problem where we will require multiple relaxed integer solutions.

Definition 4.12. *Let $\mathcal{G} = (V, E, \Gamma)$ be a t-deficient edge colored directed graph. Let S be a multiset of edges. We denote Single(S) to be the set of singleton colors in S, i.e., the set of colors which occur only once in $(\Gamma(e))_{e \in S}$.*

As indicated above, a random walk in a t-deficient edge colored graph, which is good, will hit a diverse set of colors, since there are many vertices adjacent to a diverse set of colors. This will ensure the presence of colors that occur only once in the walk.

We now state a result which says that two walks that are both good will together comprise singleton colors. In fact, the result says that the set of singleton colors is unlikely to be contained in a fixed small set of colors. This will not only help retrieve a non-trivial relaxed SIS, but also multiple such solutions with pairwise distinct supports. We prove this in Appendix A analogously to Lemma 22 of [GHOS26].

Lemma 4.2. *Let $\mathcal{G} = (V, E, \Gamma)$ be a t-deficient edge colored directed graph with average degree Δ. Fix $w_0 \in V$ and an 'avoidance' set of colors $B \subseteq [m]$ such that $|B| \leq 2T$. Suppose we independently sample $W_1, W_2 \sim \text{walk}_{w_0}^T$ and define the event:*

$$\mathcal{E} : W_1, W_2 \text{ are good and } \text{Single}(W_1 \sqcup W_2) \subseteq B.$$

Then we have $\Pr[\mathcal{E}] \leq (350tT/\Delta)^{T/2}$.

4.2 Algorithms for relaxed SIS

We now give algorithms to find solution(s) to the relaxed SIS problem for a sparse matrix, i.e., given $A \in \mathbb{Z}_q^{m \times n}$ with k-sparse rows, we find non-zero $\boldsymbol{c} \in \mathbb{Z}_q^m$ such that $\|\boldsymbol{c}\|_1 \leq T$, $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \leq R$. As discussed at the beginning of this section, this reduces to finding a walk of length T such that the difference between its starting and ending vertices has a small ℓ_1 norm.

Our strategy is simple: we sample many random walks from a starting vertex $\boldsymbol{w}_0$ (which is sampled according to $\nu(\mathcal{K}_l(A))$) and hope that the end points of two walks have a difference of small ℓ_1 norm.

Algorithm 1 FIND-GOOD-CLOSE-WALK(A, l, T, R, δ)

Require: A matrix $A \in \mathbb{Z}_q^{m \times n}$ with k-sparse rows, Kikuchi Graph parameter $\lceil k/2 \rceil \leq l \leq n$, walk length $T \in \mathbb{N}$, relaxation parameter $R \geq 0$, error threshold $\delta \in (0, 1)$

1: Sample $\boldsymbol{w}_0 \sim \nu\left(\mathcal{K}_l(A)\right)$.

2: Draw $W_1, \ldots, W_L \sim \mathrm{walk}_{w_0}^T$ independently, where $L := 8\sqrt{\log(2/\delta)}\sqrt{\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}}$

3: **if** there exist $i \neq j \in [L]$ such that $W_i = \left(\boldsymbol{w}_0, \ldots, \boldsymbol{w}_T^{(i)}\right), W_j = \left(\boldsymbol{w}_0, \ldots, \boldsymbol{w}_T^{(j)}\right)$ are both good and satisfy $\left\|\boldsymbol{w}_T^{(i)} - \boldsymbol{w}_T^{(j)}\right\|_1 \leq r$ **then**

4: **return** $W_i \circ W_j^{\mathrm{rev}}$

5: **end if**

6: **return** `FAIL`

We defer analysing the correctness and complexity of this algorithm to Subsection 4.3.

For our purpose, we will also require multiple distinct solutions to the relaxed SIS problem[8]. The 'distinctness' property we will realize in our collection of solutions $\mathcal{C}$ is somewhat strong: For any two $\boldsymbol{c}, \boldsymbol{c}' \in \mathcal{C}$, there exists $i \in [m]$ such that $|c_i| = 1$ and $c'_i = 0$.

We now present the algorithm FIND-DISTINCT-NONTRIVIAL-CLOSE-WALKS (Algorithm 2) for finding walks corresponding to such solutions. The strategy is straightforward: we repeatedly sample close walks via Algorithm 1 and add them to our collection if they satisfy a distinctness property with respect to all walks in the current collection. In particular, we require that the set of singleton edge labels (see Definition 4.12) of the current walk must not be contained in the edge labels of a previous walk.

4.3 Analysis of algorithms

In this subsection, $\mathcal{G}$ refers to the Kikuchi graph $\mathcal{K}_l(A)$.

4.3.1 Analysing Find-Good-Close-Walk

By Lemma 4.1 for a (fixed) $i \in [L]$, $\Pr[W_i \text{ is good}] > 1/4$. By Chernoff bound (Theorem 3.2), there are at least $L/8$ walks which are good with high probability. Of such good walks, we need to argue (for the chosen value of L) that a couple of them have terminal vertices with ℓ_1 distance at most R. Note that we are effectively independently sampling $L/8$ terminal vertices from some distribution $\mathcal{D}$ on the vertex set, and we henceforth analyze the probability of sampling two vertices from $\mathcal{D}$ which are at most R distance apart.

[8] specifically for solving the sparse LPN problem for higher modulus

Algorithm 2 FIND-DISTINCT-NONTRIVIAL-CLOSE-WALKS(A, l, T, R, ζ)

Require: A matrix $A \in \mathbb{Z}_q^{m \times n}$ with k-sparse rows, Kikuchi Graph parameter $\lceil k/2 \rceil \le l \le n$, walk length T, relaxation parameter $R \ge 0$, parameter $\zeta \in (0, 0.5)$.

$\mathcal{W} := \emptyset$

for $r = 1, \ldots, \left\lceil \left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}} \right)^{\zeta} \right\rceil$ **do**

  Fix $\delta = 0.5 \left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}} \right)^{-2\zeta}$

  $W_r \leftarrow$ FIND-GOOD-CLOSE-WALK(A, l, T, R, δ)

  if Single$(W_r) \neq \emptyset$, and Single$(W_r) \neq \Gamma(W_{r'})$ for all $r' < r$ **then**

    add W_r to $\mathcal{W}$

  end if

end for

return $\mathcal{W}$

Let us call $S \subseteq (\mathbb{Z}_q^n)_l$ a set of *R-faraway* points if for all $\boldsymbol{u}, \boldsymbol{v} \in S$ we have $\|\boldsymbol{u} - \boldsymbol{v}\|_1 > R$. S is a maximal set of R-faraway points if for every $\boldsymbol{v} \in (\mathbb{Z}_q^n)_l \setminus S$, there exists $\boldsymbol{u} \in S$ such that $\|\boldsymbol{u} - \boldsymbol{v}\|_1 \le R$. We first show that the uniform distribution over a maximal set of faraway points is a worst-case distribution for the event of sampling nearby points.

Proposition 4.3. *Let $\mathcal{D}$ be a distribution on vertices of $\mathcal{G}$. Suppose $S \subseteq$ Supp$(\mathcal{D})$ is a maximal set of $R/2$-faraway points. Then*

$$\Pr_{\boldsymbol{w}_1, \ldots, \boldsymbol{w}_L \overset{iid}{\sim} \mathcal{D}} \left[\exists i, i' \in [L], i \neq i' : \ \|\boldsymbol{w}_i - \boldsymbol{w}_{i'}\|_1 \le R \right] \ge 1 - \prod_{j=1}^{L-1} \left(1 - \frac{j}{|S|} \right) \ge 1 - \exp\left(-\frac{L(L-1)}{2\,|S|} \right)$$

Proof. Let $\mathcal{E}$ denote the event: $\exists i, i' \in [L], i \neq i' : \ \|\boldsymbol{w}_i - \boldsymbol{w}_i'\|_1 \le R$.

Let us define $\varphi : \text{Supp}(\mathcal{D}) \to S$ which maps every point to a nearest point in S (w.r.t. the ℓ_1 norm). There may exist more than one nearest point in S for some point $\boldsymbol{v} \in \text{Supp}(\mathcal{D})$ in which case $\varphi(\boldsymbol{v})$ is chosen arbitrarily amongst the nearest points. The map φ partitions Supp$(\mathcal{D})$ into $|S|$ *clusters*;

$$\text{Supp}(\mathcal{D}) = \bigsqcup_{\boldsymbol{u} \in S} \varphi^{-1}(\boldsymbol{u})$$

By the maximality condition on S, each point $\boldsymbol{v} \in \varphi^{-1}(\boldsymbol{u})$ satisfies $\|\boldsymbol{v} - \boldsymbol{u}\|_1 \le R/2$. By Triangle inequality (Observation 3.2), all points in $\varphi^{-1}(\boldsymbol{u})$ are within distance R from each other.

Now $\mathcal{E}^c$ is the event that the points $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_L \overset{iid}{\sim} \mathcal{D}$ are at distance $> R$ from each other. This is contained in the event that each point $\boldsymbol{w}_1, \ldots, \boldsymbol{w}_L$ is chosen from a distinct partition $\varphi^{-1}(\boldsymbol{u})$.

Let $p'_{\boldsymbol{u}}$ be the probability of a point being chosen from $\varphi^{-1}(\boldsymbol{u})$. We have

$$\Pr[\mathcal{E}^c] \le L! \sum_{\{\boldsymbol{u}_1, \ldots, \boldsymbol{u}_L\} \subseteq S} p'_{\boldsymbol{u}_1} \ldots p'_{\boldsymbol{u}_L}$$

The RHS in above equation represents the probability of collision when sampling L points from S with distribution p'. We can apply Fact 3.1 on this equation to get

$$\Pr[\mathcal{E}^c] \le \prod_{j=1}^{L-1} \left(1 - \frac{j}{|S|} \right)$$

□

Proposition 4.3 gives us a lower bound for the probability of getting two nearby vertices while sampling L independent vertices from any distribution; in particular this lower bound for any distribution is achieved when S is a maximal $R/2$-faraway set of maximum size. Let us now obtain an upper bound on $|S|$.

Observation 4.13. *Let $\boldsymbol{u} \in (\mathbb{Z}_q)_l^n$, let $R < q/2$. There are at least $\binom{R+l}{l}$ vectors in $(\mathbb{Z}_q)_l^n$ within an ℓ_1 distance of R from $\boldsymbol{u}$.*

Proof. We only compute the number of vectors $\boldsymbol{x} \in (\mathbb{Z}_q)_l^n$ within distance R from $\boldsymbol{u}$ which are non-zero at the same l positions as those in $\boldsymbol{u}$.

$$\|\boldsymbol{x} - \boldsymbol{u}\|_1 \leq R \;\Rightarrow |x_{j_1} - u_{j_1}| + \ldots + |x_{j_l} - u_{j_l}| \leq R$$

Each of $|x_{j_p} - u_{j_p}|$ can possibly take values in $\{0, \ldots, \lfloor q/2 \rfloor\}$. Hence, the above equation is satisfied for at least $\binom{R+l}{l}$ values of $\boldsymbol{x}$. □

Lemma 4.4. *Let $S \subseteq (\mathbb{Z}_q^n)_l$ be a R-faraway set. Then $|S| \leq \frac{\binom{n}{l}(q-1)^l}{\binom{R/2+l}{l}}$*

Proof. If S is a set of R-faraway points, and we define

$$B_{R/2}(\boldsymbol{u}) := \{\boldsymbol{v} \in (\mathbb{Z}_q^n)_l : \|\boldsymbol{u} - \boldsymbol{v}\|_1 \leq R/2\}$$

then $B_{R/2}(\boldsymbol{u})$ are pairwise disjoint sets for $\boldsymbol{u} \in S$. So we have

$$\sum_{\boldsymbol{u} \in S} \left|B_{R/2}(\boldsymbol{u})\right| \leq \binom{n}{l}(q-1)^l$$

and the result follows by using the lower bound for $\left|B_{R/2}(\boldsymbol{u})\right|$ from Observation 4.13. □

We are now ready to prove the correctness of Algorithm 1.

Lemma 4.5. *Let $A \in \mathbb{Z}_q^{m \times n}$ have k-sparse rows and $\lceil k/2 \rceil \leq l \leq n$ be a parameter. Let $N := \binom{n}{l}(q-1)^l \geq 64 \log(2/\delta)$. FIND-GOOD-CLOSE-WALK$(A, l, T, R, \delta)$ successfully returns a walk $W = W_i \circ W_j^{rev} = (\boldsymbol{u}_0, \ldots, \boldsymbol{u}_{2T})$ in $\mathcal{K}_l(A)$ where $\|\boldsymbol{u}_{2T} - \boldsymbol{u}_0\|_1 \leq R$, both W_i, W_j are of length T and are good, except with probability at most δ.*

Proof. Firstly, by Lemma 4.1 for a (fixed) $i \in [L]$, $\Pr[W_i \text{ is good}] > 1/4$. By Chernoff bound (Theorem 3.2), we get that there are at least $L/8$ walks which are good except with probability at most $\exp(-L/64) \leq \delta/2$ for the lower bound on N. Conditioned on this event, let $\mathcal{D}$ be the distribution of final vertices of the good walks. By Proposition 4.3, we get that the probability of all terminal vertices of the walks being R-faraway is at most

$$\prod_{j=1}^{L/8-1} \left(1 - \frac{j}{|S|}\right) \leq \exp\left(-\frac{L^2}{64\,|S|}\right) \leq \delta/2$$

where $S \subseteq \text{Supp}(\mathcal{D})$ is a maximal set of R-faraway points, and we substituted the upper bound for $|S|$ (Lemma 4.4) and the algorithm's value for L to obtain the last inequality. Adding both the failure probabilities gives us a failure probability of δ. □

The second part of our task is to ensure that $\boldsymbol{c}$ is non-trivial. We will achieve the stronger condition that $|c_i| = 1$ for some $i \in [m]$. In fact, we will show an even stronger result that such an i is outside any fixed subset $B \subseteq [m]$, $|B| \leq 2T$. This will be particularly helpful in producing multiple short integer solutions $\boldsymbol{c}$ which have pairwise distinct supports.

Corollary 4.6. *Let $\alpha \in (0,1)$ be a fixed parameter*[9] *and let $\delta \in (0, 0.5)$ be a failure probability parameter (for Algorithm 1). Let $A \in \mathbb{Z}_q^{m \times n}$ have k-sparse rows. Let $B \subseteq [m]$, $|B| \leq 2T$ be a subset of colors.*

Consider the following assumptions about the parameters:

- $N := \binom{n}{l}(q-1)^l \geq 64 \log(2/\delta)$
- $T \geq \frac{2}{\log(1/\alpha)} \left(\log \left(N / \binom{R/4+l}{l} \right) + \log(1/\delta) + \log\log(1/\delta) + 4 \right)$
- $\Delta := \frac{m\theta^{(n,l,k)}}{N} \geq \frac{700}{\alpha} T$
- $|B| \leq 2T$

Then Find-Good-Close-Walk(A, l, T, R, δ) *successfully returns a walk $W = (\boldsymbol{u}_0, \ldots, \boldsymbol{u}_{2T})$ in the Kikuchi graph $\mathcal{K}_l(A)$ such that $\|\boldsymbol{u}_{2T} - \boldsymbol{u}_0\|_1 \leq R$ and Single$(W) \not\subseteq B$ except with probability at most 2δ.*

Proof. By Lemma 4.5, we get that Find-Good-Close-Walk(A, l, T, R, δ) returns a walk $W = (\boldsymbol{u}_0, \ldots, \boldsymbol{u}_{2T})$ in $\mathcal{K}_l(A)$ where $\|\boldsymbol{u}_{2T} - \boldsymbol{u}_0\|_1 \leq R$ with both halves good except with probability δ. The algorithm also samples $\binom{L}{2} \leq 32 \log(2/\delta) \frac{N}{\binom{R/4+l}{l}}$ pairs of walks sampled independently from $\text{walk}_{w_0}^T$. Lemma 4.2 says that any such good pair W_1, W_2 satisfies Single$(W_1 \sqcup W_2) \subseteq B$ with probability at most $\left(\frac{700T}{\Delta}\right)^{T/2} \leq \frac{\delta \binom{R/4+l}{l}}{32N \log(2/\delta)}$. By a union bound, any of the pairs can have Single$(W_1 \sqcup W_2) \subseteq B$ with probability at most δ. Thus the total failure probability is bounded by 2δ. □

Thus Algorithm 1 returns a walk which is 'non-trivial' in the sense that its set of singleton colors is not contained in any small fixed set B. This in turn means that the vector $\boldsymbol{c} \in \mathbb{Z}_q^m$ obtained from such a walk (via Construction 4.1) chooses non-zero rows in A with multiplicity 1 ($|c_i| = 1$ for $i \notin B$ for some fixed B). Let us show this along with arguing the complexity of the algorithm.

Corollary 4.7. *Let $\alpha \in (0,1)$ be a fixed parameter. Let k, n, l, q be positive integers, $R \geq 0$ be a relaxation parameter, $\delta \in (0, 0.5)$ be a failure parameter. Consider a matrix $A \in \mathbb{Z}_q^{m \times n}$ having k-sparse rows.*

1. *Suppose $\lceil k/2 \rceil \leq l \leq n - \lfloor k/2 \rfloor$, and $m \geq \frac{1}{\alpha} \cdot l \log\left(\frac{qn}{R/4+l}\right) \cdot (q-1)^{\lfloor k/2 \rfloor} \left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil}$ where $C > 0$ is an absolute constant. There is an algorithm that runs in time $\widetilde{O}\left((q-1)^{\lfloor k/2 \rfloor} \left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil} \sqrt{\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}}\right)$*[10] *and finds $\boldsymbol{c} \in \mathbb{Z}_q^m$ having $\|\boldsymbol{c}\|_1 \leq \frac{4}{\log(1/\alpha)} \left(l \log\left(\frac{qn}{R/4+l}\right) + \log(1/\delta) + \log\log(1/\delta) + 4 \right)$, $c_i \in \{\pm 1\}$ for some $i \in [m]$, $\left\| A^T \boldsymbol{c} \right\|_1 \leq R$, except with probability 2δ.*

[9] we can choose α to be any constant now; in Subsection 5.2 we will implicitly choose an appropriate value

[10] here, $\widetilde{O}$ also hides $\log(1/\delta)$ and lower terms

2. *Suppose* $k \le l \le n$, *and* $m \ge \frac{1}{\alpha} \cdot l \log\left(\frac{qn}{R/4+l}\right) \cdot \left(\frac{Cn}{l}\right)^k$ *where* $C > 0$ *is an absolute constant. There is an algorithm that runs in time* $\widetilde{O}\left((Cn/l)^k \sqrt{\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}}\right)$ *and finds* $\boldsymbol{c} \in \mathbb{Z}_q^m$ *having* $\|\boldsymbol{c}\|_1 \le \frac{4}{\log(1/\alpha)}\left(l \log\left(\frac{qn}{R/4+l}\right) + \log(1/\delta) + \log\log(1/\delta) + 4\right)$, $c_i \in \{\pm 1\}$ *for some* $i \in [m]$, $\left\|A^T \boldsymbol{c}\right\|_1 \le R$, *except with probability* 2δ.

Proof. The corollary has two items because of two lower bounds for $\theta^{(n,l,k)}$ in (10),(11) out of which one is better depending on q. Both have the same proof idea; below we prove the second item.

Recall that the Kikuchi graph of A, $\mathcal{K}_l(A)$ is 2-deficient edge colored (Observation 4.11), has $N = \binom{n}{l}(q-1)^l$ vertices and $m\theta^{(n,l,k)}$ edges ($\theta^{(n,l,k)}$ is defined in (9)). We use the FIND-GOOD-CLOSE-WALK algorithm (Algorithm 1).

To apply Corollary 4.6, we also require $\Delta = \frac{m\theta^{(n,l,k)}}{N} \ge \frac{700}{\alpha}T$ which gives us $m \ge l\log(qn) \cdot \left(\frac{Cn}{l}\right)^k$ (here we used (11) for a lower bound on $\theta^{(n,l,k)}$). Now by Corollary 4.6, the algorithm returns a closed walk W in the Kikuchi Graph of length $2T$, which has at least one edge color occuring exactly once ($\text{Single}(W) \notin \emptyset$). We can use Construction 4.1 on W to get $\boldsymbol{c} \in \mathbb{Z}_q^m$ which satisfies $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \le R$ (from Observation 4.3). The singleton color condition ensures that there is some entry c_i in $\boldsymbol{c}$ which is ± 1, and the length of W is an upper bound on $\|\boldsymbol{c}\|_1$, i.e. $\|\boldsymbol{c}\|_1 \le 2T$.

Let us now argue the time complexity. In Algorithm 1, step 2 is the dominating step which takes $L = \widetilde{O}\left(\sqrt{\frac{N}{\binom{R/4+l}{l}}}\right)$ iterations. It involves finding neighbors of vertices of the Kikuchi Graph $\mathcal{K}_l(\mathcal{I})$. Each vertex of the graph can be represented in $\widetilde{O}(l)$ bits, and there are $m = l\log\left(\frac{qn}{R/4+l}\right) \cdot \left(\frac{Cn}{l}\right)^k$ possible neighbors [11] (corr. to m constraints) for any vertex. Since we sample a walk of length T, step 1 takes $\widetilde{O}\left(\sqrt{\frac{N}{\binom{R/4+l}{l}}} mlT\right) = \widetilde{O}\left(m\sqrt{\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}}\right)$ time. □

4.3.2 Analysing Find-Distinct-Nontrivial-Close-Walks

The correctness of FIND-DISTINCT-NONTRIVIAL-CLOSE-WALKS (Algorithm 2) largely follows from correctness of Algorithm 1 since it uses the latter as a subroutine. We are mainly left to argue that repeatedly invoking the latter will yield enough number of distinct non-trivial good walks. We will again take the help of Lemma 4.2 for this.

Corollary 4.8. *Let* $\alpha \in (0,1)$ *be a fixed parameter and let* $\zeta \in (0, 0.5)$ *be a parameter for number of iterations (for Algorithm 2). Let* $A \in \mathbb{Z}_q^{m \times n}$ *have* k*-sparse rows and* $\lceil k/2 \rceil \le l \le n$ *be a parameter.*

Consider the following assumptions about the parameters:

- $N := \binom{n}{l}(q-1)^l \ge 450$
- $T \ge \frac{2}{\log(1/\alpha)}\left((1+2\zeta)\log\left(N/\binom{R/4+l}{l}\right) + \log\log\left(N/\binom{R/4+l}{l}\right) + 5\right)$
- $\Delta := \frac{m\theta^{(n,l,k)}}{N} \ge \frac{700}{\alpha}T$

[11] if there are more than the required number of samples, we may discard them

Then with probability at least $1-\exp\left(-0.4\left(\frac{N}{\binom{R/4+l}{l}}\right)^{\zeta}\right)$, FIND-DISTINCT-NONTRIVIAL-CLOSE-WALKS(A,l,T,R,ζ) *returns a list* $\mathcal{W}$ *of* $0.1\left\lceil\left(\frac{N}{\binom{R/4+l}{l}}\right)^{\zeta}\right\rceil$ *nontrivial walks* $W=(\boldsymbol{u}_0,\ldots,\boldsymbol{u}_{2T})$ *such that* $\|\boldsymbol{u}_{2T}-\boldsymbol{u}_0\|_1\le R$, *with* Single$(W)\not\subseteq\Gamma(W')$ *for any* $W,W'\in\mathcal{W}$.

Proof. As fixed in the algorithm, $\delta=0.5\left(N/\binom{R/4+l}{l}\right)^{-2\zeta}$. Let $Q=\left\lceil\left(\frac{N}{\binom{R/4+l}{l}}\right)^{\zeta}\right\rceil$. From Corollary 4.6, the r^{th} iteration produces a closed walk W_r which satisfies Single$(W_r)\not\subseteq\Gamma(W_{r'})$ for some previous walk $W_{r'}$, except with probability 2δ. By a union bound, Single$(W_r)\not\subseteq\Gamma(W_{r'})$ for every $r'<r$ except with probability $2\delta\cdot r$. Let X_r be the indicator variable for the success of this event, i.e. $\Pr[X_r=1]\ge 1-2\delta\cdot r$. We construct $Y_r=\mathbb{E}\left[\sum_{i=1}^{Q}X_i\mid X_1,\ldots,X_r\right]$ which is the Doob Martingale associated with the sum $\sum_{i=1}^{Q}X_i$. Note that $Y_Q=\sum_{i=1}^{Q}X_i$ and $Y_0=\mathbb{E}\left[\sum_{i=1}^{Q}X_i\right]$. Since $\{Y_r\}_{r=0}^{Q}$ is a martingale, and satisfies $|Y_r-Y_{r-1}|\le 1$, we can apply Azuma-Hoeffding's inequality (Theorem 3.4) to it, getting

$$\Pr\left[Y_Q\le Y_0-t\right]\le\exp\left(-\frac{t^2}{2Q}\right)\tag{12}$$

Noting that $Y_0=\sum_{r=1}^{Q}\mathbb{E}[X_r]\ge Q-Q(Q+1)\delta>Q-1$, we get

$$\Pr\left[Y_Q\le 0.1Q\right]\le\exp\left(-0.4Q\right)\quad(\text{using } t=0.9Q)$$

□

Corollary 4.9. *Let* $\alpha\in(0,1)$ *be a fixed parameter. Let* k,n,l,q *be positive integers,* $R\ge 0$ *be a relaxation parameter. Consider* $A\in\mathbb{Z}_q^{m\times n}$ *having* k*-sparse rows.*

1. *Suppose* $\lceil k/2\rceil\le l\le n-\lfloor k/2\rfloor$, *and* $m\ge\frac{1}{\alpha}\cdot l\log\left(\frac{qn}{R/4+l}\right)\cdot(q-1)^{\lfloor k/2\rfloor}\left(\frac{Cn}{l}\right)^{\lceil k/2\rceil}$ *where* $C>0$ *is an absolute constant. There is an algorithm that runs in time* $\widetilde{O}\left((q-1)^{\lfloor k/2\rfloor}\left(\frac{Cn}{l}\right)^{\lceil k/2\rceil}\left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}\right)^{1/2+\zeta}\right)$ *and finds a set* $\mathcal{C}\subseteq\mathbb{Z}_q^m$ *of size* $0.1\left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}\right)^{\zeta}$ *such that*

   - $\forall\boldsymbol{c}\in\mathcal{C},\left\|A^T\boldsymbol{c}\right\|_1\le R$, $\|\boldsymbol{c}\|_1\le\frac{4}{\log(1/\alpha)}\left((1+2\zeta)l\log\left(\frac{qn}{R/4+l}\right)+\log l+\log\left(\log\left(\frac{qn}{R/4+l}\right)\right)+5\right)$
   - *For all distinct* $\boldsymbol{c},\boldsymbol{c}'\in\mathcal{C}$, $\exists i\in[m],|c_i|=1,c_i'=0$.

   The failure probability of the algorithm is $(qn)^{-\Omega(l)}$.

2. *Suppose* $k\le l\le n$, *and* $m\ge\frac{1}{\alpha}\cdot l\log\left(\frac{qn}{R/4+l}\right)\cdot\left(\frac{Cn}{l}\right)^{k}$ *where* $C>0$ *is an absolute constant. There is an algorithm that runs in time* $\widetilde{O}\left((Cn/l)^k\left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}\right)^{1/2+\zeta}\right)$ *and finds a set* $\mathcal{C}\subseteq\mathbb{Z}_q^m$ *of size* $0.1\left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}}\right)^{\zeta}$ *such that*

   - $\forall\boldsymbol{c}\in\mathcal{C},\left\|A^T\boldsymbol{c}\right\|_1\le R$, $\|\boldsymbol{c}\|_1\le\frac{4}{\log(1/\alpha)}\left((1+2\zeta)l\log\left(\frac{qn}{R/4+l}\right)+\log l+\log\log\left(\frac{qn}{R/4+l}\right)+5\right)$
   - *For all distinct* $\boldsymbol{c},\boldsymbol{c}'\in\mathcal{C}$, $\exists i\in[m],|c_i|=1,c_i'=0$.

   The failure probability of the algorithm is $(qn)^{-\Omega(l)}$.

Proof. The corollary has two items because of two lower bounds for $\theta^{(n,l,k)}$ in (10),(11) out of which one is better depending on q. Both have the same proof idea; below we prove the second item.

Recall that the Kikuchi graph of A, $\mathcal{K}_l(A)$ has $N = \binom{n}{l}(q-1)^l$ vertices and $m\theta^{(n,l,k)}$ edges ($\theta^{(n,l,k)}$ is defined in (9)). We use the FIND-DISTINCT-NONTRIVIAL-CLOSE-WALKS algorithm (Algorithm 2), and setting the minimum T required in Corollary 4.8.

To apply Corollary 4.8, we also require $\Delta = \frac{m\theta^{(n,l,k)}}{N} \geq \frac{700}{\alpha}T$ which gives us $m \geq \frac{1}{\alpha} \cdot l\ln(qn) \cdot \left(\frac{Cn}{l}\right)^k$ (here we used (11) for a lower bound on $\theta^{(n,l,k)}$). The correctness of the algorithm follows from Corollary 4.8.

Let us now argue the time complexity. Algorithm 2 calls Algorithm 1 $\left\lceil \left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}} \right)^\zeta \right\rceil$ times. This contributes to the dominating time for each such call, so the total time so far (using Corollary 4.7) is $\widetilde{O}\left(\left(\frac{N}{\binom{R/4+l}{l}} \right)^\zeta (Cn/l)^k \sqrt{\frac{N}{\binom{R/4+l}{l}}} \right) = \widetilde{O}\left((Cn/l)^k \left(\frac{\binom{n}{l}(q-1)^l}{\binom{R/4+l}{l}} \right)^{1/2+\zeta} \right)$. □

In Section 5 we will see how to utilize the solution(s) to the relaxed SIS problem to solve the k-LIN $(\mathbb{Z}_q)$ problem for two different kinds of noise distributions.

5 Distinguishing k-LIN $(\mathbb{Z}_q)$ instances

Consider a k-LIN $(\mathbb{Z}_q)$ instance $\mathcal{I} = \{(\boldsymbol{a}_i, b_i)_{i\in[m]}\}$ which we view as $(A, \boldsymbol{b}) \in \mathbb{Z}_q^{m\times n} \times \mathbb{Z}_q^m$ where $\boldsymbol{a}_i$ is the i^{th} row of A. We are promised that

- either $\boldsymbol{b}$ is uniformly random in $\mathbb{Z}_q^m$, or
- $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$ for unknown $\boldsymbol{s} \in \mathbb{Z}_q^n$ satisfying $\|\boldsymbol{s}\|_\infty \leq B$, and $e_i \xleftarrow{\text{iid}} \Psi$ for some distribution Ψ on $\mathbb{Z}_q$.

We will now see how to utilize a relaxed SIS for the distinguishing task. We will consider two kinds of noise distributions: one where the noise is sufficiently bounded with high probability (Subsection 5.1) and the other noise being the LPN noise distribution (Subsection 5.2).

5.1 Distinguishing k-LIN $(\mathbb{Z}_q)$ with bounded noise

In this subsection, we assume that the noise distribution Ψ on $\mathbb{Z}_q$ is *bounded* (with high probability), i.e. for $e \sim \Psi$, $|e| \leq D'$ with high probability. This will give us a bound of the form: for $\widetilde{\boldsymbol{e}} \sim \Psi^T$ for some $T \in \mathbb{N}$, $\|\widetilde{\boldsymbol{e}}\|_\infty \leq D$ with high probability.

From Corollary 4.7 we can obtain a solution to the relaxed SIS instance A, i.e. we can obtain $\boldsymbol{c} \in \mathbb{Z}_q^m$ such that $\|\boldsymbol{c}\|_1$, $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1$ are bounded, $c_i \in \{\pm 1\}$ for some $i \in [m]$. Let us first examine the planted case. Let us denote $\widetilde{\boldsymbol{c}}$, $\widetilde{\boldsymbol{e}}$ as versions of $\boldsymbol{c}$,$\boldsymbol{e}$ truncated on the support of $\boldsymbol{c}$. This support size ($\|\boldsymbol{c}\|_0$) is bounded by $\|\boldsymbol{c}\|_1$ for which the bound by Corollary 4.7 prevails, same as the value of $2T$ in the theorem statement. Then $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} = \boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s} + \boldsymbol{c}^{\mathrm{T}}\boldsymbol{e} = \boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s} + \widetilde{\boldsymbol{c}}^{\mathrm{T}}\widetilde{\boldsymbol{e}}$. Then

$$\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| = \left|\boldsymbol{c}^{\mathrm{T}}(A\boldsymbol{s}+\boldsymbol{e})\right| \leq \left|\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right| + \left|\widetilde{\boldsymbol{c}}^{\mathrm{T}}\widetilde{\boldsymbol{e}}\right| \quad \text{(using Observation 3.2)} \tag{13}$$

$$\leq \left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \|\boldsymbol{s}\|_\infty + \|\widetilde{\boldsymbol{c}}\|_1 \|\widetilde{\boldsymbol{e}}\|_\infty \quad \text{(using Observation 3.3)} \tag{14}$$

$$\leq \left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \cdot B + \|\boldsymbol{c}\|_1 \cdot D \tag{15}$$

In the random case, $\boldsymbol{b}$ is uniformly random in $\mathbb{Z}_q^m$. Further, Corollary 4.7 ensures that $\boldsymbol{c}$ has $c_i \in \{\pm 1\}$ for at least one coordinate $i \in [m]$. As a consequence $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} = c_i b_i + \sum_{i' \neq i} c_{i'} b_{i'}$ is uniformly random in $\mathbb{Z}_q$. This implies, for $\varepsilon \in (0,1)$, $\Pr\left[\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| > \varepsilon q/4\right] = 1 - \varepsilon/2$.

To distinguish on the basis of $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right|$, we need to ensure that $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| \leq \varepsilon q/4$. Recall that the relaxation paramter R was an upper bound for $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1$, which we can set so that $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \cdot B \leq \varepsilon q/8$. Further, we will derive another condition on q to ensure $\|\boldsymbol{c}\|_1 \cdot D \leq \varepsilon q/8$.

Theorem 5.1. *Let n, m, q, k be positive integers, $D > 0$ be a error bound parameter, $0 < B \leq q/2$ be a secret vector bound, $\varepsilon \in (0, 0.5)$ be a failure parameter. Consider a k-LIN$(\mathbb{Z}_q)$ instance $\mathcal{I}$ of dimension n and m constraints which is either random or is planted with some secret vector $\boldsymbol{s} \in \mathbb{Z}_q^n$ satisfying $\|\boldsymbol{s}\|_\infty \leq B$ and noise distribution Ψ on $\mathbb{Z}_q$, where for $T = l \log(32Bn/\varepsilon)$ we have $\Pr_{\widetilde{e} \sim \Psi^T}\left[\|\widetilde{\boldsymbol{e}}\|_\infty > D\right] < \varepsilon'$. There exists an algorithm to distinguish the distribution of $\mathcal{I}$ with failure probability at most $2.5\varepsilon + \varepsilon'$ provided the following:*

1. *Suppose the following on the parameters*
   - $\lceil k/2 \rceil \leq l \leq n - \lfloor k/2 \rfloor$
   - $q \geq \frac{16lD}{\varepsilon} \log(32Bn/\varepsilon)$
   - $m \geq l \log(32Bn/\varepsilon) \cdot (q-1)^{\lfloor k/2 \rfloor} \left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil}$ *where $C > 0$ is an absolute constant.*

   The algorithm takes time $\widetilde{O}\left((q-1)^{\lfloor k/2 \rfloor} \left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil} (32Bn/\varepsilon)^{l/2}\right)$.

2. *Suppose the following on the parameters*
   - $k \leq l \leq n$
   - $q \geq \frac{16lD}{\varepsilon} \log(32Bn/\varepsilon)$
   - $m \geq l \log(32Bn/\varepsilon) \cdot \left(\frac{Cn}{l}\right)^k$ *where $C > 0$ is an absolute constant.*

   The algorithm takes time $\widetilde{O}\left((Cn/l)^k (32Bn/\varepsilon)^{l/2}\right)$.

Proof. The theorem has two items because of two lower bounds for $\theta^{(n,l,k)}$ in (10),(11) out of which one is better depending on q. Both have the same proof idea; below we prove the second item.

Let us invoke the algorithm in Corollary 4.7 with relaxation parameter $R = \frac{\varepsilon q}{8B}$ and failure parameter $\delta := \varepsilon$. We choose the fixed parameter α in the Corollary to be an appropriate constant so that T and m of the Corollary are as given in this Theorem. Let $\boldsymbol{c} \in \mathbb{Z}_q^m$ be the vector returned by the algorithm. The algorithm fails with probability at most 2ε and its time complexity matches the claimed one by using the value of R and the bound on q in the expression of Corollary 4.7.

If the instance is random, $\boldsymbol{b}$ is uniformly random in $\mathbb{Z}_q^m$. Further, Corollary 4.7 ensures that $\boldsymbol{c}$ has $c_i \in \{\pm 1\}$ for at least one coordinate $i \in [m]$. As a consequence $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} = c_i b_i + \sum_{i' \neq i} c_{i'} b_{i'}$ is uniformly random in $\mathbb{Z}_q$. As a consequence, for $\varepsilon \in (0,1)$ $\Pr\left[\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| \leq \varepsilon q/4\right] = \varepsilon/2$.

Suppose the instance $(A, \boldsymbol{b})$ is planted, $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$. Let us denote $\widetilde{\boldsymbol{c}}$, $\widetilde{\boldsymbol{e}}$ as versions of $\boldsymbol{c}$,$\boldsymbol{e}$ truncated on the support of $\boldsymbol{c}$. This support size ($\|\boldsymbol{c}\|_0$) is bounded by $\|\boldsymbol{c}\|_1$ for which the bound by Corollary 4.7 prevails, same as the value of T in the theorem statement. Then $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} = \boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s} + \boldsymbol{c}^{\mathrm{T}}\boldsymbol{e} = \boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s} + \widetilde{\boldsymbol{c}}^{\mathrm{T}}\widetilde{\boldsymbol{e}}$. By (15), we have (except with probability ε')

$$\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| \leq \left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \cdot B + \|\boldsymbol{c}\|_1 \cdot D \ \leq \varepsilon q/4$$

by substituting the relaxation parameter R and the lower bound for q.

Thus we can output PLANTED based on whether $\left|\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}\right| \leq \varepsilon q/4$ and be right except with probability at most $2.5\varepsilon + \varepsilon'$. □

For sparse LWE, the noise distribution is typically a discrete Gaussian $\Psi_{\mathbb{Z}_q,r}$ with width parameter $r < q$. For $\widetilde{\boldsymbol{e}} \sim \Psi^T_{\mathbb{Z}_q,r}$, have a bound of $\|\widetilde{\boldsymbol{e}}\|_\infty \leq r\sqrt{T}$ (Observation 3.10) with probability $2^{-\Omega(T)} = (Bn/\varepsilon)^{-\Omega(l)}$.

Corollary 5.2. *Let n, m, q, k be positive integers, $0 < B \leq q/2$ be a parameter for secret vector bound, $\varepsilon \in (0, 0.5)$ be a failure parameter. There exists an algorithm that solves the decisional k-LWE$_{n,m,q,r,B}$ problem with failure probability at most $2.5\varepsilon + (Bn/\varepsilon)^{-\Omega(k)}$ provided the following:*

1. *Suppose the following on the parameters*
   - $\lceil k/2 \rceil \leq l \leq n - \lfloor k/2 \rfloor$
   - $q \geq \frac{16r}{\varepsilon}\left(l \log(32Bn/\varepsilon)\right)^{3/2}$
   - $m \geq l \log\left(32Bn/\varepsilon\right) \cdot (q-1)^{\lfloor k/2 \rfloor} \left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil}$ *where $C > 0$ is an absolute constant.*

   The algorithm takes time $\widetilde{O}\left((q-1)^{\lfloor k/2 \rfloor} \left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil} (32Bn/\varepsilon)^{l/2}\right)$.

2. *Suppose the following on the parameters*
   - $k \leq l \leq n$
   - $q \geq \frac{16r}{\varepsilon}\left(l \log(32Bn/\varepsilon)\right)^{3/2}$
   - $m \geq l \log\left(32Bn/\varepsilon\right) \cdot \left(\frac{Cn}{l}\right)^{k}$ *where $C > 0$ is an absolute constant.*

   The algorithm takes time $\widetilde{O}\left((Cn/l)^k (32Bn/\varepsilon)^{l/2}\right)$.

Remark 5.1. *Note that the hardness of standard LWE from lattice problems was established by Regev [Reg05] for any $r > 2\sqrt{n}$ and $r < q$. The small secrets LWE (in the non-sparse regime) is also equivalent to standard LWE ([BLPRS13]). On the other hand, Corollary 5.2 solves k-sparse LWE in this parameter regime as long as, say $q \gg r(k \log n)^{3/2}$.*

5.2 Solving sparse LPN

Recall that the noise distribution $\xi_{\mathbb{Z}_q,\mu}$ in the planted case of an LPN instance (Definition 3.12) has a uniform mass over all non-zero points in $\mathbb{Z}_q$. With this noise distribution, there are a couple of issues due to which we need to alter our distinguishing strategy from Subsection 5.1.

Firstly, $e \leftarrow \xi_{\mathbb{Z}_q,\mu}$ is not sufficiently bounded as all non-zero points in $\mathbb{Z}_q$ have the same probability mass. This means that the distinguishing technique in Subsection 5.1 will not be useful. Instead of analyzing the absolute value of $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$; we will work with a different function of $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$. Consider the canonical character of $\mathbb{Z}_q$

$$\begin{aligned} \chi :& \mathbb{Z}_q \to \mathbb{C} \\ & x \mapsto e^{2\pi\iota x/q} \end{aligned}$$

which is an (additive) injective group homomorphism.

Observation 5.2. *Let $c \in \mathbb{Z}_q \setminus \{0\}$. Then $\sum_{a\in\mathbb{Z}_q} \chi(ca) = \sum_{a=0}^{q-1} e^{2\pi\iota ca/q} = 0$.*

Secondly, the distributions of $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{e}$ and $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b}$ for $\boldsymbol{b} \leftarrow \mathrm{Unif}(\mathbb{Z}_q^m)$ can be statistically very close if $\boldsymbol{c}$ has even moderately many non-zero entries. The latter distribution is uniform over $\mathbb{Z}_q$, as discussed in Subsection 5.1. To analyze the former distribution, firstly, if c_i is a unit in $\mathbb{Z}_q$, then c_ie_i has the same distribution as e_i, since $c_ie_i = 0$ iff $e_i = 0$, and the probability mass of e_i at any point in $\mathbb{Z}_q \setminus \{0\}$ is the same. Suppose the non-zero c_i's are all units, then $\sum_{i\in[m]} c_ie_i$ has the same distribution as $\sum_{i:c_i\neq 0} e_i$. This distribution has roughly $(1-\mu)^{\|\boldsymbol{c}\|_0}$ mass on 0, and the remaining mass is uniform over non-zero points. For $\boldsymbol{c}$ obtained by algorithms in Section 4 $\|\boldsymbol{c}\|_0$ can be as high as $\Omega(l\log(qn))$ which means this quantity is not significantly larger than $1/q$, and the distribution can be very close to uniform. To mitigate this issue, we will obtain multiple distinct instances $\mathcal{C}$ of $\boldsymbol{c} \in \mathbb{Z}_q^m$.

Overall, we will distinguish based on the quantity: $\sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})$.

When $\boldsymbol{b}$ is random in $\mathbb{Z}_q^m$, $\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b} = \sum_{i\in[m]} c_ib_i$ is random in $\mathbb{Z}_q$ because Corollary 4.7 ensures that $\boldsymbol{c}$ has $c_i \in \{\pm 1\}$ for at least one coordinate $i \in [m]$. Using this Observation 5.2 we get

$$\sum_{\boldsymbol{c}\in\mathcal{C}} \mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right] = 0 \tag{16}$$

To analyze the expectation when $\boldsymbol{b}$ is planted, let us first make a couple of observations.

Observation 5.3. *Let $e \sim \xi_{\mathbb{Z}_q,\mu}$, $c \in \mathbb{Z}_q \setminus \{0\}$. Then $\mathbb{E}[\chi(ce)] = 1 - \frac{q\mu}{q-1}$*

Proof.

$$\begin{aligned}\mathbb{E}[\chi(ce)] = \sum_{a=0}^{q-1} \chi(ca)\xi_{Z_q,\mu}(a) &= 1-\mu+(\mu/(q-1))\sum_{a=1}^{q-1}\chi(ca)\\ &= 1-\mu-\mu/(q-1) \text{ (from Observation 5.2)}\end{aligned}$$

□

Observation 5.4. *Suppose $c \in \mathbb{Z}_q$ such that $|c| \le B' \le q/4$. Then $Re\,(\chi(c)) \ge \cos(2\pi B'/q)$*

Proof. When $|c| \le q/4$, $\mathrm{Re}\,(\chi(c)) = \cos(2\pi c/q)$ is decreasing (resp. increasing) for $c \ge 0$ (resp. $c < 0$). □

For the planted case, we have $\boldsymbol{b} = A\boldsymbol{s} + \boldsymbol{e}$. Additionally we assume that the secret vector $\boldsymbol{s}$ satisfies $\|\boldsymbol{s}\|_\infty \le B$. We will utilize multiple short integer solutions for A from Section 4; in particular there is an algorithm (Corollary 4.9) that yields a collection $\mathcal{C} = \{\boldsymbol{c} \in \mathbb{Z}_q^m\}$ of distinct non-trivial $\boldsymbol{c} \in \mathbb{Z}_q^m$ such that $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1 \le R$ and $\|\boldsymbol{c}\|_1 \le 2T$ for appropriate T as per the Corollary bound for $\|\boldsymbol{c}\|_1$.

Let $\rho := 1 - \frac{q\mu}{q-1}$. We have

$$\begin{aligned}\sum_{\boldsymbol{c}\in\mathcal{C}} \mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right] &= \sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s})\,\mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{e})\right] &(17)\\ &= \sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}) \prod_i \mathbb{E}\left[\chi(c_ie_i)\right] \quad (\text{since } e_i\text{'s are independent}) &(18)\\ &= \sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s})\rho^{\|\boldsymbol{c}\|_0} \quad (\text{from Observation 5.3}) &(19)\\ \Rightarrow \mathrm{Re}\left(\sum_{\boldsymbol{c}\in\mathcal{C}} \mathbb{E}\left[\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right]\right) &\ge \cos(2\pi RB/q)\sum_{\boldsymbol{c}\in\mathcal{C}} \rho^{\|\boldsymbol{c}\|_0} \quad \left(\text{Observations 3.3, 5.4 on } \left|\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s}\right|\right) &(20)\\ &\ge \cos(2\pi RB/q)\sum_{\boldsymbol{c}\in\mathcal{C}} \rho^{\|\boldsymbol{c}\|_1} \ge \cos(2\pi RB/q)\,|\mathcal{C}|\,\rho^{2T} &(21)\end{aligned}$$

So we managed to obtain a separation in the quantity $\sum_{\boldsymbol{c}\in\mathcal{C}} \mathrm{Re}\left(\chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right)$ in expectation for the random and planted cases, and we hope to realize a separation with high probability. However we face a challenge for this; the summands, though pairwise independent need not be independent as distinct $\boldsymbol{c} \in \mathcal{C}$ may share non-zero coordinates. This prevents us from using standard concentration inequalities. For the random case, we are able to use the second moment method to show an upper bound on $\left|\sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right|$. For the planted case, we show *anti-concentration* of $\left|\sum_{\boldsymbol{c}\in\mathcal{C}} \chi(\boldsymbol{c}^{\mathrm{T}}\boldsymbol{b})\right|$ around 0. We use an analogous idea as [GHOS26] of developing anti-concentration results for a certain class of polynomials.

5.2.1 Polynomial Anticoncentration

Let us define a 'block-bivariate' m-variate polynomial with respect to a partition of $[m]$. This is a generalization of block-linear polynomials used in [GHOS26]. These are polynomials over two sets of indeterminates $\{x_1,\ldots,x_m\}$,$\{y_1,\ldots,y_m\}$. Every monomial in such a polynomial consists of only one set of x_i's and y_i's from each part of the partition.

Definition 5.5. *Let m, T be integers. Let*

$$P(x_1,\ldots,x_m,y_1,\ldots,y_m) = \sum_{S\subseteq[m]} a_S \prod_{i\in S} x_i y_i^{d_i}$$

be a polynomial, where $a_S \in \mathbb{C}$. Let $\boldsymbol{\pi} = (\pi_1,\ldots,\pi_T)$ be a partition of $[m]$. For $S \subseteq [m]$, we say $\boldsymbol{\pi}$ shatters *S if $|S \cap \pi_j| \leq 1$ for every $j \in [T]$. Define the block-bivariate restriction of P to $\boldsymbol{\pi}$ as*

$$P_{\boldsymbol{\pi}}(x_1,\ldots,x_m) := \sum_{\substack{S\subseteq[m]\\ \pi \text{ shatters } S}} a_S \prod_{i\in S} x_i y_i^{d_i}$$

We say that P is $\boldsymbol{\pi}$-block-bivariate if $P_{\boldsymbol{\pi}} = P$.

The following result says that for a random equi-partition $\boldsymbol{\pi}$ of $[m]$, the block-bivariate restriction of a polynomial to $\boldsymbol{\pi}$ does not lose too many terms.

Lemma 5.3. *(Lemma 26 in [GHOS26]) Let $m, T \in \mathbb{N}$, and let $P(x_1,\ldots,x_m)$ be a sum of μ monomials in the x_i's of degree at most T. Then for a uniformly random equipartition $\boldsymbol{\pi} = (\pi_1,\ldots,\pi_T)$ of $[m]$ ($|\pi_1| = \ldots = |\pi_{T-1}| \geq |\pi_T|$),*

$$\Pr[\textit{number of monomials in } P_{\boldsymbol{\pi}} \geq 0.5 \cdot \mu \cdot \exp(-T)] \geq \exp(-2T).$$

For our purpose, we will give an anti-concentration result for block-bivariate polynomials. We will reduce this task for block-bivariate polynomials to a similar task for linear polynomials. We will need to impose a special condition on such a block-bivariate polynomial: $\mathbb{E}[Y^d] = \mathbb{E}[Y]$ for all $d \in [q-1]$. We will be able to show this is the case when $Y_i = \chi(b_i)$ by virtue of Observation 5.3.

The polynomial we consider also has X_i's as a separate set of variables. Let us justify why we use them. Consider a linear polynomial where each indeterminate is a discrete random variable. Even if a polynomial has mean far from, say 1, it can still have non-negligible probability of being close to 1 if the variance is large. This issue can be circumvented if the indeterminates were *continuous* random variables with bounded density functions. In this case, even if the variance of the polynomial is large, the individual random variables are not concentrated in any small interval and it can be shown that the polynomial's value can also not be concentrated in any small interval. For this reason, we will add continuous noise to our random variables of interest $\{Y_i = \chi(b_i)\}$, which is done by the random variables $\{X_i\}_{i\in[m]}$.

We make this concrete in Appendix B with similar flavor as Section 5 of [GHOS26]. We will use the following anti-concentration result for our planted case.

Lemma 5.4. *Let $\mathcal{D}$ be a distribution on $[0,1]$ with probability density $f_{\mathcal{D}} \le 1$. Let $\mathcal{D}'$ be another distribution on $\{x \in \mathbb{C} : |x| \le 1\}$. Furthermore, for $Y \sim \mathcal{D}'$ we have $\mathbb{E}\left[Y^d\right] = \mathbb{E}[Y] \in \mathbb{R}_{>0}$ for all $d \in [q-1]$. Fix integers $m, T \ge 1$ and a partition $\boldsymbol{\pi} = (\pi_1, \dots, \pi_T)$ of $[m]$. Let*

$$P(x_1, \dots, x_m, y_1, \dots, y_m) = \sum_{S \subseteq [m]} a_S \prod_{i \in S} x_i y_i^{d_i}$$

be a $\boldsymbol{\pi}$-block-bivariate polynomial of total degree $\le T$, such that the exponents $d_i \in [q-1]$. Define

$$\mu := \mathop{\mathbb{E}}_{X_i \overset{iid}{\sim} \mathcal{D}, Y_i \overset{iid}{\sim} \mathcal{D}'} [P(X_1, \dots, X_m, Y_1, \dots, Y_m)]$$

Then for an absolute constant C_{anti}:

$$\Pr_{X_i \overset{iid}{\sim} \mathcal{D}, Y_i \overset{iid}{\sim} \mathcal{D}'} \left[|P(X_1, \dots, X_m, Y_1, \dots, Y_m)| < C_{anti}^{-10T} |\mu| \right] \le 0.9^T.$$

5.2.2 Distinguishing using relaxed SIS

Our technique will be to use the collection $\mathcal{C}$ of distinct solutions $\boldsymbol{c} \in \mathbb{Z}_q^m$ of the relaxed SIS problem for A obtained from Corollary 4.9. We first discard those $\boldsymbol{c} \in \mathcal{C}$ which are not shattered by a random equipartition of $[m]$. If there are still sufficiently many solutions left, we first add a uniform noise to $\chi(b_i)$ (for reasons discussed in above Subsubsection 5.2.1) and then use the distinguishing quantity as discussed above. We present algorithm DISTINGUISH-LPN-INSTANCE (Algorithm 3) for this purpose.

Let us analyze the correctness and complexity of Algorithm 3.

Firstly, we observe that the algorithm does not output `FAIL` with high probability. Indeed, Corollary 4.8 and Lemma 5.3 guarantee that for a random equi-partition $\boldsymbol{\pi}$, there are at least $0.05 N^{\zeta} e^{-2T}$ shattered terms of $\mathcal{C}$ with probability at least e^{-4T}, i.e. the loop breaks in any iteration with probability at least e^{-4T}. Thus the loop breaks in e^{5T} iterations with probability at least $1 - (1 - e^{-4T})^{e^{5T}} \ge 1 - e^{-e^T}$. Hence, in the rest of the argument, we condition on the event that the algorithm does not output `FAIL`.

Now, notice that our output decision is based on the following quantity:

$$\psi_{\mathcal{C}'} := \sum_{\boldsymbol{c} \in \mathcal{C}'} \prod_{\substack{i \in [m] \\ c_i \neq 0}} \eta_i \left(\chi(c_i b_i) \right) \tag{22}$$

where η_i are i.i.d. random variables from Unif$([0,1])$.

Now $\psi_{\mathcal{C}'}$ is a polynomial in $\{\eta_1, \dots, \eta_m, \chi(b_1), \dots, \chi(b_m)\}$, as for $c_i \in \mathbb{Z}_q \setminus \{0\}$ we can write $\chi(c_i b_i) = \chi(b_i)^{c_i'}$ where $c_i' \in [q-1]$, $c_i' \cong c_i \mod q$. This will help us analyze the expectation and use the anti-concentration results.

The expectation follows similarly to our calculations in (16), (21):

$$\mathbb{E}[\psi_{\mathcal{C}'}] = 0 \text{ when } \mathcal{I} \text{ is random} \tag{23}$$

$$\mathrm{Re}\left(\mathbb{E}[\psi_{\mathcal{C}'}]\right) = \mathrm{Re}\left(\sum_{\boldsymbol{c} \in \mathcal{C}'} \chi(\boldsymbol{c}^{\mathrm{T}} A \boldsymbol{s}) \prod_{\substack{i \in [m] \\ c_i \neq 0}} \mathbb{E}[\eta_i] \, \mathbb{E}[\chi(c_i e_i)] \right) = \mathrm{Re}\left(\sum_{\boldsymbol{c} \in \mathcal{C}'} \chi(\boldsymbol{c}^{\mathrm{T}} A \boldsymbol{s}) \prod_{\substack{i \in [m] \\ c_i \neq 0}} \frac{\rho}{2} \right) \tag{24}$$

$$\ge \cos(2\pi RB/q) \left|\mathcal{C}'\right| (\rho/2)^{2T} \quad = \cos(\varepsilon) \left|\mathcal{C}'\right| (\rho/2)^{2T} \text{ when } \mathcal{I} \text{ is planted} \tag{25}$$

Algorithm 3 DISTINGUISH-LPN-INSTANCE$(\mathcal{I}, l, \varepsilon)$

Require: A k-LPN$_{n,m,q,\mu,B}$ instance $\mathcal{I} = \{(\boldsymbol{a}_i, b_i)\}$, Kikuchi graph parameter $k/2 \le l \le n$, failure parameter $\varepsilon \in (0, 0.5)$

$R = \frac{\varepsilon q}{2\pi B}$, $\zeta = \varepsilon$, $A = [\{\boldsymbol{a}_i\}]^{\mathrm{T}}$, $N := \binom{n}{l}(q-1)^l$

$T = \frac{\zeta}{4\log\left(2eC_{\mathrm{anti}}^{10}/\left(1-\frac{q\mu}{q-1}\right)\right)}\left((1+2\zeta)\log\left(N/\binom{R/4+l}{l}\right) + \log\log\left(N/\binom{R/4+l}{l}\right) + 5\right)$

$\mathcal{W} \leftarrow$ FIND-DISTINCT-NONTRIVIAL-CLOSE-WALKS(A, l, T, R, ζ)

Apply Construction 4.1 on the walks in $\mathcal{W}$ to obtain the collection $\mathcal{C}$

for $s = 1, \ldots, e^{5T}$ **do**

  Sample a random partition $\boldsymbol{\pi}$ of $[m]$ into $2T$ equally-sized parts ($|\pi_1| = \ldots = |\pi_{2T-1}| \ge |\pi_{2T}|$)

  Let $\mathcal{C}' \subseteq \mathcal{C}$ be the set of $\boldsymbol{c} \in \mathbb{Z}_q^m$ whose support is shattered by $\boldsymbol{\pi}$

  if $|\mathcal{C}'| \ge 0.05\left(N/\binom{R/4+l}{l}\right)^{\zeta} \cdot e^{-2T}$ **then**

    restrict $\mathcal{C}'$ to $0.05\left(N/\binom{R/4+l}{l}\right)^{\zeta} \cdot e^{-2T}$ arbitrary elements

    break the loop

  end if

end for

if the loop never broke **then**

  output `FAIL`

end if

Sample $\eta_1, \ldots, \eta_m \overset{iid}{\sim} \mathrm{Unif}([0,1])$

Define $\psi_{\mathcal{C}'} := \sum_{\boldsymbol{c}\in\mathcal{C}'} \prod_{\substack{i\in[m]\\ c_i\neq 0}} \eta_i\left(\chi(c_i b_i)\right)$

if $|\psi_{\mathcal{C}'}| \ge \left(N/\binom{R/4+l}{l}\right)^{0.6\zeta}$ **then**

  output `PLANTED`

else

  output `RANDOM`

end if

Hence we will establish that $|\Psi_{\mathcal{C}'}|$ will be low in the random case and high in the planted case.

For the random case, we are still able to prove an upper bound for $\Psi_{\mathcal{C}'}$ using the second moment method. For the planted case, we obtain a lower bound for $\Psi_{\mathcal{C}'}$ by proving anti-concentration around 0.

In the rest of the argument, when we write X^c for an indeterminate X and $c \in \mathbb{Z}_q$, we assume that $c \in \{0, \ldots, q-1\}$ while evaluating the exponent.

The following result is useful for the random case:

Proposition 5.5. *Let N, R, T, ζ be parameters fixed in Algorithm 3. Let*

$$P_{\mathcal{C}'}(x_1, \ldots, x_m, y_1, \ldots, y_m) := \sum_{\boldsymbol{c}\in\mathcal{C}'} \alpha_{\boldsymbol{c}} \prod_{\substack{i\in[m]\\ c_i\neq 0}} x_i y_i^{c_i}$$

where $|\alpha_{\boldsymbol{c}}| \le 1$. If $X_1, \ldots, X_m, Y_1, \ldots, Y_m$ are complex-valued independent random variables such that $|X_i|, |Y_i| \le 1$, and $\mathbb{E}[Y_i^c] = 0$ for all $c \in \mathbb{Z}_q \setminus \{0\}$, then:

$$\Pr\left[|P_{\mathcal{C}'}(X_1, \ldots, X_m, Y_1, \ldots, Y_m)| < \left(N/\binom{R/4+l}{l}\right)^{0.6\zeta}\right] > 1 - n^{-\Omega(\zeta l)}.$$

Proof. We show this using the second moment method. Firstly, $\mathbb{E}\left[P_{\mathcal{C}'}(X_1,\ldots,X_m,Y_1,\ldots,Y_m)\right]=0$ by the assumptions on X_i's and Y_i's. Further, Corollary 4.8 guarantees that (except with probability $\exp\left(-0.4\left(N/\binom{R/4+l}{l}\right)^{\zeta}\right)$) for any distinct $\boldsymbol{c},\boldsymbol{c}'\in\mathcal{C}'$, there exists $i\in[m]$ such that $c_i\neq 0$ and $c'_i=0$. This implies that the summands in $P_{\mathcal{C}'}(X_1,\ldots,X_m,Y_1,\ldots,Y_m)$ are pairwise independent. Thus

$$\begin{aligned}\mathbb{E}\left[\left|P_{\mathcal{C}'}(X_1,\ldots,X_m,Y_1,\ldots,Y_m)\right|^2\right]&=\mathbb{E}\left[\sum_{\boldsymbol{c}\in\mathcal{C}'}|\alpha_{\boldsymbol{c}}|^2\prod_{\substack{i\in[m]\\c_i\neq 0}}|X_i|^2\,|Y_i|^{2c_i}\right]\\&\leq|\mathcal{C}'|=0.05\left(N/\binom{R/4+l}{l}\right)^{\zeta}\cdot e^{-2T}\end{aligned}$$

where the second last inequality is by the bound on X_i's and Y_i's and the last inequality follows from line 9 in Algorithm 3. We complete the argument by applying Chebyshev's inequality:

$$\begin{aligned}&\Pr\left[\left|P_{\mathcal{C}'}(X_1,\ldots,X_m,Y_1,\ldots,Y_m)\right|\geq\left(N/\binom{R/4+l}{l}\right)^{0.6\zeta}\right]\\&\leq\left(N/\binom{R/4+l}{l}\right)^{-1.2\zeta}\mathrm{Var}(P_{\mathcal{C}'}(X_1,\ldots,X_m,Y_1,\ldots,Y_m))\\&\leq 0.05\left(N/\binom{R/4+l}{l}\right)^{-0.2\zeta}\cdot e^{-2T}\;=n^{-\Omega(\zeta l)}\end{aligned}$$

□

If $(A,\boldsymbol{b})$ is random, we can use $X_i=\eta_i$, $Y_i=\chi(b_i)$ and $\alpha_{\boldsymbol{c}}=\chi(\boldsymbol{c}^{\mathrm{T}}A\boldsymbol{s})$ in the polynomial of Proposition 5.5 to argue that $|\psi_{\mathcal{C}'}|$ is small with high probability.

Now we show a lower bound for the planted case:

Proposition 5.6. *Suppose $(A,\boldsymbol{b})$ is a planted k-LPN$_{n,m,q,\mu,B}$ instance where $\boldsymbol{b}=A\boldsymbol{s}+\boldsymbol{e}$. Let all symbols be as described in Algorithm 3. Let*

$$P_{\mathcal{C}',A,\boldsymbol{s}}(x_1,\ldots,x_m,y_1,\ldots,y_m):=\sum_{\boldsymbol{c}\in\mathcal{C}'}\chi(\boldsymbol{c}^TA\boldsymbol{s})\prod_{\substack{i\in[m]\\c_i\neq 0}}x_iy_i^{c_i}$$

Then:

$$\Pr\left[\left|P_{\mathcal{C}',A,\boldsymbol{s}}\left(\eta_1,\ldots,\eta_m,\chi(e_1),\ldots,\chi(e_m)\right)\right|\geq\left(N/\binom{R/4+l}{l}\right)^{0.6\zeta}\right]\geq 1-n^{-\Omega(\zeta l)}.$$

Proof. Note that $\|\boldsymbol{s}\|_\infty\leq B$ and $\mathcal{C}'$ consists of $0.05N^{\zeta}e^{-2T}$ vectors $\boldsymbol{c}$ with $\|\boldsymbol{c}\|_1\leq 2T$, $\left\|A^{\mathrm{T}}\boldsymbol{c}\right\|_1\leq R$. Let $\mu:=\mathbb{E}\left[P_{\mathcal{C}',A,\boldsymbol{s}}(\eta_1,\ldots,\eta_m,\chi(e_1),\ldots,\chi(e_m))\right]$. From (25) we have

$$\begin{aligned}\mathrm{Re}(\mu)=\mathrm{Re}\left(\mathbb{E}\left[P_{\mathcal{C}',A,\boldsymbol{s}}(\eta_1,\ldots,\eta_m,\chi(e_1),\ldots,\chi(e_m))\right]\right)&\geq\cos(\varepsilon)\left|\mathcal{C}'\right|(\rho/2)^{2T}\\&=0.05\cos(\varepsilon)N^{\zeta}e^{-2T}\,(\rho/2)^{2T}\end{aligned}$$

By Lemma 5.4, we have

$$\begin{aligned}
&\Pr\left[\left|P_{\mathcal{C}',A,\boldsymbol{s}}(\eta_1,\ldots,\eta_m,\chi(e_1),\ldots,\chi(e_m)\right| < 0.05C_{\text{anti}}^{-20T}N^{\zeta}e^{-2T}\left(\rho/2\right)^{2T}\right]\\
&\leq \Pr\left[\left|P_{\mathcal{C}',A,\boldsymbol{s}}(\eta_1,\ldots,\eta_m,\chi(e_1),\ldots,\chi(e_m)\right| < C_{\text{anti}}^{-20T}\text{Re}(\mu)\right]\\
&\leq \Pr\left[\left|P_{\mathcal{C}',A,\boldsymbol{s}}(\eta_1,\ldots,\eta_m,\chi(e_1),\ldots,\chi(e_m)\right| < C_{\text{anti}}^{-20T}\left|\mu\right|\right] \leq 0.9^{2T} \leq n^{-\Omega(\zeta l)}
\end{aligned}$$

To finish proving the result, observe that $0.05C_{\text{anti}}^{-20T}N^{\zeta}e^{-2T}\left(\rho/2\right)^{2T} \geq \left(N/\binom{R/4+l}{l}\right)^{0.6\zeta}$ for the chosen value of T in the algorithm. □

With these results, we obtain a distinguisher for higher modulus sparse LPN:

Theorem 5.7. *Let n, m, q, k be positive integers, $0 < B \leq q/2$ be a parameter for secret vector bound, $\varepsilon \in (0, 0.5)$ be a failure parameter. There exists an algorithm for solving the decisional k-LPN$_{n,m,q,\mu,B}$ problem with failure probability at most $n^{-\Omega(\varepsilon l)}$ provided the following:*

1. *Suppose the following on the parameters*
   - $\lceil k/2 \rceil \leq l \leq n - \lfloor k/2 \rfloor$
   - $m \geq \left(1 - \frac{q\mu}{q-1}\right)^{-8/\varepsilon}(l\log(8\pi Bn/\varepsilon))\cdot\left(\frac{C(q-1)n}{l}\right)^{k/2}$ *where $C > 0$ is an absolute constant.*

   The algorithm takes time $\widetilde{O}\left(\left(1 - \frac{q\mu}{q-1}\right)^{-8/\varepsilon}(q-1)^{\lfloor k/2 \rfloor}\left(\frac{Cn}{l}\right)^{\lceil k/2 \rceil}(8\pi Bn/\varepsilon)^{l(1/2+\varepsilon)}\right)$.

2. *Suppose the following on the parameters*
   - $k \leq l \leq n$
   - $m \geq \left(1 - \frac{q\mu}{q-1}\right)^{-8/\varepsilon}(l\log(8\pi Bn/\varepsilon))\cdot\left(\frac{Cn}{l}\right)^{k}$ *where $C > 0$ is an absolute constant.*

   The algorithm takes time $\widetilde{O}\left(\left(1 - \frac{q\mu}{q-1}\right)^{-8/\varepsilon}(Cn/l)^{k}(8\pi Bn/\varepsilon)^{l(1/2+\varepsilon)}\right)$.

Proof. The theorem has two items because of two lower bounds for $\theta^{(n,l,k)}$ in (10),(11) out of which one is better depending on q. Both have the same proof idea; below we prove the second item.

Let $(A, \boldsymbol{b})$ be a k-LPN$_{n,m,q,\mu,B}$ instance. Recall from Observation 5.3 $\rho := \mathbb{E}[\chi(\xi_{\mathbb{Z}_q,\mu})] = 1 - \frac{q\mu}{q-1}$. Propositions 5.5, 5.6 establish that Algorithm 3 when given as input a set $\mathcal{C}$ of $\boldsymbol{c} \in \mathbb{Z}_q^m$ of size $\|\boldsymbol{c}\|_1 \leq 2T$, with $|\mathcal{C}| \geq 0.1N^{\zeta}$, outputs a guess `RANDOM` or `PLANTED`, correct with probability at least $1 - N^{-\Omega(\zeta)}$. Using this after applying Algorithm 2 (whose correctness is given by Corollary 4.9) establishes the correctness of the algorithm.

Corollary 4.9 also establishes the sample complexity and time complexity of Algorithm 2 which matches the claimed sample and time complexity. We now show that the time of Algorithm 3 is dominated by this time complexity. The `for` loop in steps 1-7 makes $e^{5T} = N^{O(\zeta)}$ iterations. Within each iteration, step 5 takes $\widetilde{O}(m)$ time and steps 6-9 are doable in $O(|\mathcal{C}|\,T) = \widetilde{O}(N^{\zeta})$ time, with a total time of $N^{O(\zeta)}\widetilde{O}(m)$. The remaining steps are also doable in $\widetilde{O}(|\mathcal{C}'|\,T) = \widetilde{O}(N^{\zeta})$ time. Overall, the remaining steps take $N^{O(\zeta)}\widetilde{O}(m) = \widetilde{O}\left(m\left(\binom{n}{l}q^{l}\right)^{O(\zeta)}\right)$ time. □

References

[ABK26] Shashwat Agrawal, Amitabha Bagchi, and Rajendra Kumar. *Spectral Method attacks Sparse LWE, Sparse LPN and Beyond*. Cryptology ePrint Archive, Paper 2026/614. 2026. URL: https://eprint.iacr.org/2026/614.

[ABW10] Benny Applebaum, Boaz Barak, and Avi Wigderson. "Public-key cryptography from different assumptions". In: *Proceedings of the forty-second ACM symposium on Theory of computing*. 2010, 171–180.

[ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. "Fast cryptographic primitives and circular-secure encryption based on hard learning problems". In: *Advances in Cryptology–CRYPTO 2009*. Springer. 2009, pp. 595–618.

[ADINZ17a] Benny Applebaum, Ivan Damgård, Yuval Ishai, Michael Nielsen, and Lior Zichron. "Secure arithmetic computation with constant computational overhead". In: *Annual International Cryptology Conference*. Springer. 2017, pp. 223–254.

[ADINZ17b] Benny Applebaum, Ivan Damgård, Yuval Ishai, Michael Nielsen, and Lior Zichron. "Secure arithmetic computation with constant computational overhead". In: *Annual International Cryptology Conference*. Springer. 2017, pp. 223–254.

[AL16] Benny Applebaum and Shachar Lovett. "Algebraic attacks against random local functions and their countermeasures". In: *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*. 2016, pp. 1087–1100.

[Ale03] Michael Alekhnovich. "More on average case vs approximation complexity". In: *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.* IEEE. 2003, pp. 298–307.

[AOW15] Sarah R Allen, Ryan O'Donnell, and David Witmer. "How to refute a random CSP". In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*. IEEE. 2015, pp. 689–708.

[Ban93] Wojciech Banaszczyk. "New bounds in some transference theorems in the geometry of numbers". In: *Mathematische Annalen* 296.1 (1993), pp. 625–635.

[BBTV25] Kiril Bangachev, Guy Bresler, Stefan Tiegel, and Vinod Vaikuntanathan. "Near-optimal time-sparsity trade-offs for solving noisy linear equations". In: *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*. 2025, pp. 1910–1920.

[BGV14] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. "(Leveled) fully homomorphic encryption without bootstrapping". In: *ACM Transactions on Computation Theory (TOCT)* 6.3 (2014), pp. 1–36.

[BLM13] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities. A nonasymptotic theory of independence*, With a foreword by Michel Ledoux. Oxford University Press, Oxford, 2013, pp. x+481. ISBN: 978-0-19-953525-5. DOI: 10.1093/acprof:oso/9780199535255.001.0001. URL: https://doi.org/10.1093/acprof:oso/9780199535255.001.0001.

[BLPRS13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. "Classical hardness of learning with errors". In: *Proceedings of the forty-fifth annual ACM symposium on Theory of computing* (2013), pp. 575–584.

[BV14] Zvika Brakerski and Vinod Vaikuntanathan. "Lattice-based FHE as secure as PKE". In: *Proceedings of the 5th conference on Innovations in theoretical computer science*. 2014, pp. 1–12.

[CGGI20] Ilaria Chillotti, Nicolas Gama, Mariya Georgieva, and Malika Izabachène. "TFHE: Fast Fully Homomorphic Encryption Over the Torus: I. Chillotti et al." In: *Journal of Cryptology* 33.1 (2020), pp. 34–91.

[CGM19] Yilei Chen, Nicholas Genise, and Pratyay Mukherjee. "Approximate trapdoors for lattices and smaller hash-and-sign signatures". In: *International Conference on the Theory and Application of Cryptology and Information Security*. Springer. 2019, pp. 3–32.

[CM01] Mary Cryan and Peter Bro Miltersen. "On pseudorandom generators in NC0". In: *International Symposium on Mathematical Foundations of Computer Science*. Springer. 2001, pp. 272–284.

[CW23] Valerio Cini and Hoeteck Wee. "ABE for Circuits with $poly(\lambda)$-sized Keys from LWE". In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2023, pp. 435–446.

[DM15] Léo Ducas and Daniele Micciancio. "FHEW: bootstrapping homomorphic encryption in less than a second". In: *Annual international conference on the theory and applications of cryptographic techniques*. Springer. 2015, pp. 617–640.

[ESLL19] Muhammed F. Esgin, Ron Steinfeld, Joseph K. Liu, and Dongxi Liu. "Lattice-Based Zero-Knowledge Proofs: New Techniques for Shorter and Faster Constructions and Applications". In: *Advances in Cryptology – CRYPTO 2019*. Springer. 2019, pp. 115–146.

[Ess45] Carl-Gustav Esseen. "Fourier analysis of distribution functions. A mathematical study of the Laplace-Gaussian law". In: *Acta Math.* 77 (1945), pp. 1–125. ISSN: 0001-5962,1871-2509. DOI: 10.1007/BF02392223. URL: https://doi.org/10.1007/BF02392223.

[Fei02] Uriel Feige. "Relations between average case complexity and approximation complexity". In: *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*. 2002, pp. 534–543.

[Fol99] G. B. Folland. *Real Analysis: Modern Techniques and Their Applications*. 2nd. John Wiley & Sons, 1999.

[Gen09] Craig Gentry. "Fully homomorphic encryption using ideal lattices". In: *Proceedings of the forty-first annual ACM symposium on Theory of computing*. 2009, pp. 169–178.

[GHOS26] Meghal Gupta, William He, Ryan O'Donnell, and Noah G Singer. "A Classical Quadratic Speedup for Planted xor". In: *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2026, pp. 2455–2472. URL: https://arxiv.org/pdf/2508.09422.

[GKM22] Venkatesan Guruswami, Pravesh K. Kothari, and Peter Manohar. "Algorithms and certificates for Boolean CSP refutation: smoothed is no harder than random". In: *STOC '22—Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, New York, [2022] ©2022, pp. 678–689. ISBN: 978-1-4503-9264-8. DOI: 10.1145/3519935.3519955. URL: https://doi.org/10.1145/3519935.3519955.

[GKPV10] Shafi Goldwasser, Yael Tauman Kalai, Chris Peikert, and Vinod Vaikuntanathan. "Robustness of the learning with errors assumption". In: (2010).

[Gol00] Oded Goldreich. "Candidate One-Way Functions Based on Expander Graphs". In: (2000).

[Gol11] Oded Goldreich. "Candidate one-way functions based on expander graphs". In: *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation: In Collaboration with Lidor Avigad, Mihir Bellare, Zvika Brakerski, Shafi Goldwasser, Shai Halevi, Tali Kaufman, Leonid Levin, Noam Nisan, Dana Ron, Madhu Sudan, Luca Trevisan, Salil Vadhan, Avi Wigderson, David Zuckerman*. Springer, 2011, pp. 76–87.

[Hal77] G. Halász. "Estimates for the concentration function of combinatorial number theory and probability". In: *Period. Math. Hungar.* 8.3-4 (1977), pp. 197–211. ISSN: 0031-5303,1588-2829. DOI: 10.1007/BF02018403. URL: https://doi.org/10.1007/BF02018403.

[Has20] Matthew B Hastings. "Classical and quantum algorithms for tensor principal component analysis". In: *Quantum* 4 (2020), p. 237.

[Hås84] Johan Håstad. *An NP-complete problem: some aspects of its solution and some possible applications*. Institut Mittag-Leffler, 1984.

[HKM23] Jun-Ting Hsieh, Pravesh K. Kothari, and Sidhanth Mohanty. "A simple and sharper proof of the hypergraph Moore bound". In: *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, Philadelphia, PA, 2023, pp. 2324–2344. ISBN: 978-1-61197-755-4. DOI: 10.1137/1.9781611977554.ch89. URL: https://doi.org/10.1137/1.9781611977554.ch89.

[HLL23] Yin-Chun Hsieh, Huijia Lin, and Ji Luo. "Attribute-Based Encryption for Circuits of Unbounded Depth from Lattices". In: *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE. 2023, pp. 415–434.

[HLP52] Godfrey H Hardy, John E Littlewood, and George Pólya. *Inequalities*. 1952.

[JLS24] Aayush Jain, Huijia Lin, and Sagnik Saha. "A systematic study of sparse LWE". In: *Annual International Cryptology Conference*. Springer. 2024, pp. 210–245. URL: https://eprint.iacr.org/2024/1589.

[Kik51] Ryoichi Kikuchi. "A theory of cooperative phenomena". In: *Physical review* 81.6 (1951), p. 988.

[KM25] Nicholas Kocurek and Peter Manohar. "Spectral Refutations of Semirandom k-LIN over Larger Fields". In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2025)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. 2025, pp. 17–1. URL: https://arxiv.org/pdf/2508.18185.

[KMOW17] Pravesh K Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. "Sum of squares lower bounds for refuting any CSP". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 132–145.

[LPR13] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. "On ideal lattices and learning with errors over rings". In: *Journal of the ACM (JACM)* 60.6 (2013), pp. 1–35.

[Lyu09] Vadim Lyubashevsky. "Fiat-Shamir with Aborts: Applications to Lattice and Factoring-Based Signatures". In: *Advances in Cryptology – ASIACRYPT 2009*. Springer. 2009, pp. 598–616.

[Mic18] Daniele Micciancio. "On the hardness of learning with errors with binary secrets". In: *Theory of Computing* 14.1 (2018), pp. 1–17.

[MR09] Daniele Micciancio and Oded Regev. "Lattice-based Cryptography". In: *Post-Quantum Cryptography*. Ed. by Daniel J. Bernstein, Johannes Buchmann, and Erik Dahmen. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 147–191. ISBN: 978-3-540-88702-7. DOI: 10.1007/978-3-540-88702-7_5. URL: https://doi.org/10.1007/978-3-540-88702-7_5.

[MST03] Elchanan Mossel, Amir Shpilka, and Luca Trevisan. "On epsilon-Biased Generators in NC^ 0". In: *Annual Symposium on Foundations of Computer Science*. Vol. 44. IEEE Computer Society Press. 2003, pp. 136–145.

[Pei09] Chris Peikert. "Public-key cryptosystems from the worst-case shortest vector problem". In: *Proceedings of the forty-first annual ACM symposium on Theory of computing* (2009), pp. 333–342.

[Pei16] Chris Peikert. "A decade of lattice cryptography". In: *Foundations and trends® in theoretical computer science* 10.4 (2016), pp. 283–424.

[Pie12] Krzysztof Pietrzak. "Cryptography from learning parity with noise". In: *International Conference on Current Trends in Theory and Practice of Computer Science*. Springer. 2012, pp. 99–114.

[Reg05] Oded Regev. "On Lattices, Learning with Errors, Random Linear Codes, and Cryptography". In: *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing* (2005), pp. 84–93.

[RRS17] Prasad Raghavendra, Satish Rao, and Tselil Schramm. "Strongly refuting random CSPs below the spectral threshold". In: *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*. 2017, pp. 121–131.

[SOKB25] Alexander Schmidhuber, Ryan O'Donnell, Robin Kothari, and Ryan Babbush. "Quartic quantum speedups for planted inference". In: *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2025, pp. 905–913. URL: https://arxiv.org/abs/2406.19378.

[WAM19] Alexander Wein, Ahmed El Alaoui, and Cristopher Moore. "The Kikuchi hierarchy and tensor PCA". In: *Journal of the ACM* (2019).

A Proof of Lemma 4.2

We now prove the following lemma analogously as Lemma 22 of [GHOS26].

Lemma A.1. *Let $\mathcal{G} = (V, E, \Gamma)$ be a t-deficient edge colored directed graph with average degree Δ. Fix $w_0 \in V$ and an 'avoidance' set of colors $B \subseteq [m]$ such that $|B| \leq 2T$. Suppose we independently sample $W_1, W_2 \sim walk^T_{w_0}$ and define the event:*

$$\mathcal{E} : W_1, W_2 \text{ are good and } Single(W_1 \sqcup W_2) \subseteq B.$$

Then we have $\Pr[\mathcal{E}] \leq (350tT/\Delta)^{T/2}$.

Proof. For $j \in [2T]$, let e_j denote the j^{th} edge in W_1 if $j \leq T$, or the $(j-T)^{\text{th}}$ edge in W_2 if $j > T$. Similarly, let w_{j-1} denote the $(j-1)^{\text{th}}$ vertex in W_1 if $j \leq T$, or the $(j-1-T)^{\text{th}}$ vertex in W_2 if $j > T$.

For $j \in [2T]$, let us call the j^{th} time step *unusual* if:

1. w_{j-1} is not bad, yet
2. $\Gamma(e_j) \in \{\Gamma(e_1), \ldots, \Gamma(e_{j-1})\} \cup B$

We begin by arguing that if the event $\mathcal{E}$ occurs, then there is a large number of unusual steps.

First, let us identify the steps that violate condition (1). If $\mathcal{E}$ occurs then W_1, W_2 are good and then by Definition 4.7 there are less than $T/2$ vertices in W_1, W_2 combined which are bad. This means condition (1) is violated for less than $T/2$ steps.

Now let us turn to the violations of condition (2). If $\mathcal{E}$ occurs then $\text{Single}(W_1 \sqcup W_2) \subseteq B$. This means that if $\Gamma(e_j)$ is a color that appears only once, it cannot violate (2) since it must be present in B. Therefore $\Gamma(e_j)$ can violate (2) only if it appears more than once in the sequence. But since condition (2) is not violated at step j if $\Gamma(e_j)$ has already been encountered (i.e., it is in $\{\Gamma(e_1), \ldots, \Gamma(e_{j-1})\}$), every color that appears more than once can cause only one violation of (2). Hence, such violations can appear at no more than half the total number of steps, i.e., at no more than $2T/2 = T$ steps.

Thus both conditions are met for at least $2T - T - T/2 = T/2$ steps, i.e. there are at least $T/2$ unusual steps.

Let $U(W_1, W_2) \in \{0,1\}^{2T}$ denote the *unusualness* string, i.e. $U(W_1, W_2)_j = 1$ iff the j^{th} step is unusual. From above, we obtain

$$\Pr[\mathcal{E}] \leq \Pr[|U(W_1, W_2)| \geq T/2] = \sum_{u : |u| \geq T/2} \Pr[U(W_1, W_2) = u] \tag{26}$$

Let us now bound $\Pr[U(W_1, W_2) = u]$ for a fixed string $u \in \{0,1\}^{2T}$. For the same, we will bound the probability:

$$\Pr\left[U(W_1, W_2)_j = u_j = 1 \mid U(W_1, W_2)_{<j} = u_{<j}\right]$$

Note that for $U(W_1, W_2)_j = 1$, condition 1. says that w_{j-1} must not be bad, i.e. out-deg$(w_{j-1}) \geq 3\Delta/16$. Further, condition 2. says that

$$\Gamma(e_j) \in \{\Gamma(e_1), \ldots, \Gamma(e_{j-1})\} \cup B$$

There are at most $j - 1 + |B|$ choices for $\Gamma(e_j)$. Since the graph $\mathcal{G}$ is t-deficient edge colorable, we can have at most $t(j - 1 + |B|)$ outgoing edges from w_{j-1} that have colors in $\{\Gamma(e_1), \ldots, \Gamma(e_{j-1})\} \cup B$. The edge e_j is picked uniformly from the set of outgoing edges of w_{j-1}, thus we get

$$\Pr\left[U(W_1, W_2)_j = u_j = 1 \mid U(W_1, W_2)_{<j} = u_{<j}\right] \tag{27}$$

$$\leq \Pr\left[\Gamma(e_j) \in \{\Gamma(e_1), \ldots, \Gamma(e_{j-1})\} \cup B \mid w_{j-1} \text{ is not bad}\right] \tag{28}$$

$$\leq \frac{16t(j + |B|)}{3\Delta} \leq \frac{64tT}{3\Delta} \tag{29}$$

Thus,

$$\Pr[U(W_1, W_2) = u] = \prod_{j \in [2T]} \Pr\left[U(W_1, W_2)_j = u_j \mid U(W_1, W_2)_{<j} = u_{<j}\right] \tag{30}$$

$$\leq \prod_{\substack{j \in [2T] \\ u_j = 1}} \Pr\left[U(W_1, W_2)_j = u_j = 1 \mid U(W_1, W_2)_{<j} = u_{<j}\right] \tag{31}$$

$$\leq \left(\frac{64tT}{3\Delta}\right)^{|u|} \quad \text{(using (29))} \tag{32}$$

Finally, we use (32) in (26) to complete the proof:

$$\Pr[\mathcal{E}] \leq \sum_{u : |u| \geq T/2} \Pr[U(W_1, W_2) = u] \leq 2^{2T} \left(\frac{64tT}{3\Delta}\right)^{T/2} \leq \left(\frac{350tT}{\Delta}\right)^{T/2}$$

□

B Polynomial Anticoncentration

For our purpose, we will give an anti-concentration result for block-bivariate polynomials. We will reduce this task for block-bivariate polynomials to a similar task for linear polynomials. We will need to impose a special condition on such a block-bivariate polynomial: $\mathbb{E}[Y_i^d] = \mathbb{E}[Y_i]$ for all $d \in [q-1]$.

The definition for block-bivariate polynomial we consider also has X_i's as a separate set of variables. Let us justify why we use them. Consider a linear polynomial where each indeterminate is a discrete random variable. Even if a polynomial has mean far from, say 1, it can still have non-negligible probability of being close to 1 if the variance is large. This issue can be circumvented if the indeterminates were *continuous* random variables with bounded density functions. In this case, even if the variance of the polynomial is large, the individual random variables are not concentrated in any small interval and it can be shown that the polynomial's value can also not be concentrated in any small interval. So, in the definition of block-bivariate polynomial, the random variables $\{X_i\}_{i \in [m]}$ will serve the purpose of adding continuous noise to the (possibly discrete) random variables $\{Y_i\}_{i \in [m]}$.

Lemma B.1. *(generalizing Lemma 27 of [GHOS26]) Let $X_1, \ldots, X_m$ be independent random variables, where each X_i is complex valued, satisfies $|X_i| \leq 2$, has mean 0 and has p.d.f. bounded by 1. Let $L = c_1 X_1 + \cdots + c_m X_m$, where $c_1, \ldots, c_m \in \mathbb{C}$. Then for all $0 < \varepsilon < 1/2$, we have*

$$\Pr[|L - 1| \leq \varepsilon] \leq C_{anti} \cdot \varepsilon$$

where $C_{anti} < \infty$ is a universal constant.

Proof. Let $\sigma^2 = \sum_{i=1}^m |c_j|^2$ denote the variance scale of the linear combination.

- Suppose $\sigma^2 \leq \frac{1}{64 \log(1/\varepsilon)}$. Note that

$$|L - 1| \leq \varepsilon \implies |L| \geq 1 - \varepsilon > \frac{1}{2}$$

Because the variables are independent, complex-valued, and bounded by $|c_i X_i| \le 2|c_i|$, we invoke Hoeffding's Inequality (Thoerem 3.3) to bound the upper tail:

$$\Pr[|L-1| \le \varepsilon] \le \Pr\left[|L| \ge \frac{1}{2}\right] \le 4\exp\left(-\frac{(1/2)^2}{16\sum_{i=1}^m |c_i|^2}\right) = 4\exp\left(-\frac{1}{64\sigma^2}\right) \le 4\varepsilon$$

- Suppose $\sigma^2 > \frac{1}{64\log(1/\varepsilon)}$. For $w \in \mathbb{R}^2$, let $\phi_{c_i X_i}(w) := \mathbb{E}\left[e^{\iota\langle w, c_i X_i\rangle}\right]$ denote the characteristic function of $c_i X_i$, viewing the complex number $c_i X_i$ in $\mathbb{R}^2$. Because the density of X_i is bounded by 1, $\mathbb{E}\left[|X_i|^2\right] > 0$ and by Taylor Series expansion, there exists an absolute constant $\kappa > 0$ such that for sufficiently small w:

$$|\phi_{c_i X_i}(w)| \le \exp(-\kappa\,|c_i|^2\,|w|^2)$$

By independence, the characteristic function of the sum is $\phi_L(w) = \prod_{i=1}^m \phi_{c_i X_i}(w)$ and therefore:

$$|\phi_L(w)| \le \exp\left(-\kappa\sum_{j=1}^n |c_j|^2|w|^2\right) = \exp(-\kappa\sigma^2|w|^2)$$

We apply the two-dimensional Esséen Smoothing Lemma [Ess45; Hal77], which relates the small-ball probability to the integral of the characteristic function over a dual ball:

$$\Pr[|L-1| \le \varepsilon] \le C_0\varepsilon^2 \int_{\mathbb{R}^2} |\phi_L(w)|\,dw \le C_0\varepsilon^2 \int_{\mathbb{R}^2} \exp(-\kappa\sigma^2|w|^2)\,dw$$

Evaluating the Gaussian integral yields:

$$\Pr[|L-1| \le \varepsilon] \le \frac{C_0\pi}{\kappa}\cdot\frac{\varepsilon^2}{\sigma^2}$$

Substituting the lower bound $\sigma^2 > \frac{1}{64\log(1/\varepsilon)}$, we have:

$$\Pr[|L-1| \le \varepsilon] \le \left(\frac{64C_0\pi}{\kappa}\right)\varepsilon^2\log(1/\varepsilon) \le C_1\varepsilon$$ □

As a next step, we prove an anti-concentration result for affine polynomials.

Corollary B.2. *(similar to Corollary 29 in [GHOS26]) Let $X_1, \dots, X_m$ be independent random variables with mean $\rho \in \mathbb{R}_{>0}$, supported on $\{x \in \mathcal{C} : |x| \le 1\}$, and with pdf bounded by 1. Let L be an affine polynomial on n variables. Then for all $0 < \varepsilon < 1/2$, we have*

$$\Pr[|L(X_1, \dots, X_m)| \le \varepsilon L(\rho, \dots, \rho)] \le C_{\text{anti}} \cdot \varepsilon$$

(where C_{anti} is the constant from Lemma 27).

Proof. (same proof in [GHOS26]) Note that by affineness,

$$L(x_1, \dots, x_m) = L(x_1 - \rho, \dots, x_m - \rho) + L(\rho, \dots, \rho) - L(0, \dots, 0)$$
$$L(x_1 - \rho, \dots, x_m - \rho) - L(0, \dots, 0) \text{ is a homogeneous linear polynomial}$$

Hence we can apply Lemma B.1 to obtain

$$\Pr\left[\left|\frac{L(X_1 - \rho, \dots, X_n - \rho) - L(0, \dots, 0)}{-L(\rho, \dots, \rho)} - 1\right| \le \varepsilon\right] \le C_{\text{anti}} \cdot \varepsilon.$$

This event is equivalent to the desired event.

□

With Lemma B.1 and Corollary B.2, the following lemma can also be proved similar to Proposition 30 in [GHOS26]. This Proposition considered block linear polynomials, however we extend it to block bivariate polynomials in a special case for our purpose.

Lemma B.3. *(generalizing Proposition 30 of [GHOS26]) Let $\mathcal{D}$ be a distribution on $[0,1]$ with probability density $f_{\mathcal{D}} \leq 1$. Let $\mathcal{D}'$ be another distribution on $\{x \in \mathbb{C} : |x| \leq 1\}$. Furthermore, for $Y \sim \mathcal{D}'$ we have $\mathbb{E}[Y^d] = \mathbb{E}[Y] \in \mathbb{R}_{>0}$ for all $d \in [q-1]$. Fix integers $m, T \geq 1$ and a partition $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_T)$ of $[m]$. Let*

$$P(x_1, \ldots, x_m, y_1, \ldots, y_m) = \sum_{S \subseteq [m]} a_S \prod_{i \in S} x_i y_i^{d_i}$$

be a $\boldsymbol{\pi}$-block-bivariate polynomial of total degree $\leq T$, such that the exponents $d_i \in [q-1]$. Define

$$\mu := \mathop{\mathbb{E}}_{X_i \overset{iid}{\sim} \mathcal{D}, Y_i \overset{iid}{\sim} \mathcal{D}'} [P(X_1, \ldots, X_m, Y_1, \ldots, Y_m)]$$

Then for an absolute constant C_{anti}:

$$\Pr_{X_i \overset{iid}{\sim} \mathcal{D}, Y_i \overset{iid}{\sim} \mathcal{D}'} \left[|P(X_1, \ldots, X_m, Y_1, \ldots, Y_m)| < C_{anti}^{-10T} |\mu| \right] \leq 0.9^T.$$

Proof. Without loss of generality, we assume that for $t \in [T]$, $\pi_t = \{(t-1)\frac{m}{T} + 1, \ldots, t\frac{m}{T}\}$. Let $\kappa = \mathbb{E}[\mathcal{D}]$, $\rho = \mathbb{E}[\mathcal{D}']$. For $t \in [T] \cup \{0\}$ we define the random variable

$$R_t = P\left(X_1, \ldots, X_{tm/T}, \kappa, \ldots, \kappa, Y_1, \ldots, Y_{tm/T}, \rho, \ldots, \rho\right)$$

Consider $P_{t+1} := R_{t+1} | X_1, \ldots, X_{tm/T}, Y_1, \ldots, Y_{tm/T}$ which is a polynomial in variables $\{X_{tm/T+1}, \ldots, X_{(t+1)m/T}, Y_{tm/T+1}, \ldots, Y_{(t+1)m/T}\}$. Since P is $\boldsymbol{\pi}$-block bivariate, and for all $d \in [q-1], i \in [m]$, $\mathbb{E}[Y_i^d] = \mathbb{E}[Y_i] = \rho$, we can write P_{t+1} such that:

$$\begin{aligned}
&P_{t+1}(X_{tm/T+1}, \ldots, X_{(t+1)m/T}, Y_{tm/T+1}, \ldots, Y_{(t+1)m/T}) \\
&= c_0 + c_1 \cdot X_{tm/T+1} \cdot P_{t+1,1}(Y_{tm/T+1}) + \ldots + c_T \cdot X_{(t+1)m/T} \cdot P_{t+1,t}(Y_{(t+1)m/T}) \\
\mathbb{E}\left[X_{tm/T+1} \cdot P_{t+1,1}\right] &= \ldots = \mathbb{E}\left[X_{(t+1)m/T} \cdot P_{t+1,t}\right] = \kappa\rho
\end{aligned}$$

Also $\mathbb{E}[R_{t+1} | X_1, \ldots, X_{tm/T}, Y_1, \ldots, Y_{tm/T}] = R_t$. We can then apply Corollary B.2 to this polynomial to get:

$$\Pr\left[|R_{t+1}| \leq \varepsilon |R_t| \mid X_1, \ldots, X_{tm/T}, Y_1, \ldots, Y_{tm/T}\right] \leq C_{\text{anti}} \cdot \varepsilon \tag{33}$$

We can use (33) to get the following:

$$\mathbb{E}\left[\left|\frac{R_t}{R_{t+1}}\right|^{0.1} \mid X_1, \ldots, X_{tm/T}, Y_1, \ldots, Y_{tm/T}\right] \tag{34}$$

$$= \int_0^\infty \Pr\left[\left|\frac{R_t}{R_{t+1}}\right|^{0.1} > y \mid X_1, \ldots, X_{tm/T}, Y_1, \ldots, Y_{tm/T}\right] dy \tag{35}$$

$$= \int_0^\infty \Pr\left[|R_{t+1}| < y^{-10} |R_t| \mid X_1, \ldots, X_{tm/T}, Y_1, \ldots, Y_{tm/T}\right] dy \tag{36}$$

$$\leq 1 + \int_1^\infty C_{\text{anti}} \cdot y^{-10} \cdot dy \tag{37}$$

$$\leq 1 + 0.12 C_{\text{anti}} \tag{38}$$

Next, we will use (38) telescopically to get

$$\mathbb{E}\left[\left|\frac{R_0}{R_T}\right|^{0.1}\right] \leq (1 + 0.12C_{\text{anti}})^T \tag{39}$$

Note that $R_0 = \mu$ and $R_T = P(X_1, \ldots, X_m, Y_1, \ldots, Y_m)$. We get the desired result by applying Markov's inequality to (39) (and choosing a sufficiently large value of C_{anti}). □