

A Simpler Strong Refutation of Random k -XOR*

Kwangjun Ahn

Department of Electrical Engineering and Computer Science
Massachusetts Institute of Technology
Email: kjahn@mit.edu

August 11, 2020

Abstract

Strong refutation of random CSPs is a fundamental question in theoretical computer science that has received particular attention due to the long-standing gap between the information-theoretic limit and the computational limit. This gap is recently bridged by Raghavendra, Rao and Schramm where they study sub-exponential algorithms for the regime between the two limits. In this work, we take a simpler approach to their algorithm and analysis.

1 Introduction

Refutation of random instances of constraint satisfaction problems (random CSPs) is one of the central questions in theoretical computer science with numerous applications. Among many predicates (types of constraints), this paper considers the XOR predicate and studies the strong refutation of the corresponding random CSP. In fact, Allen, ODonnell and Witmer [2] demonstrate that one can use strong refutation algorithms for random XOR to refute random CSPs with other predicates¹. In particular, we consider:

Definition 1 (Random k -XOR). A random k -XOR with probability p (or equivalently, at density pn^{k-1}) refers to a set $\Phi = \{C_S\}$ of k -XOR constraints over n variables $x \in \{\pm 1\}^n$ obtained as per the following procedure:

1. First sample each of the n^k possible k -tuples with probability p independently.
2. For each sampled $S = (s_1, s_2, \dots, s_k) \in [n]^k$, include a k -XOR constraint $C_S : \prod_{i=1}^k x_{s_i} = \eta_S$, where η_S is i.i.d. Rademacher random variable.

For an assignment $x \in \{\pm 1\}^n$, let $P_\Phi(x)$ be the fraction of constraints satisfied by x .

*This work is presented at International Conference on Randomization and Computation (RANDOM) 2020.

¹For instance, it is demonstrated that one can refute random k -SAT by reducing it to strong refutations of random ℓ -XOR for $\ell = 1, 2, \dots, k$.

Remark 1. One can alternatively consider a model where we sample subsets of size k instead of k -tuples (there will be $\binom{n}{k}$ possible subsets in total). However, as noted in [2], the precise details of the random model are not relevant to the results to follow. For simplicity, we follow the prior works [2, 7] and consider the above k -tuple model throughout the paper.

Under this random k -XOR model, we study the strong refutation problem. To motivate the problem, it is a consequence of standard concentration inequalities that when the density is of $\omega(1)$ (i.e., $pn^{k-1} = \omega(1)$), with high probability, no assignment can satisfy more than a $1/2 + o(1)$ fraction of the constraints. Hence, it is a natural algorithmic question to ask whether one can certify such a fact. More specifically, we consider:

Definition 2 (Strong refutation). For a quantity $\alpha = \omega(1)$, an algorithm which takes a k -XOR instance and outputs a quantity \widehat{P}_Φ is said to strongly refute random k -XOR at density α if it satisfies:

1. For any k -XOR instance Φ , $P_\Phi(x) \leq \widehat{P}_\Phi$ for all assignments $x \in \{\pm 1\}^n$.
2. For a random k -XOR instance Φ with density α , $\widehat{P}_\Phi = 1/2 + o(1)$ with high probability.

However, the question of developing strong refutation algorithms for the density $\omega(1)$ turns out to be rather intractable. More specifically, the best known guarantees are obtained from spectral methods [2, 3] which require the density to be $\widetilde{\Omega}(n^{k/2-1})$. This computational limit of $\widetilde{\Omega}(n^{k/2-1})$ (also known as *spectral threshold*) is significantly larger than the information-theoretic threshold of $\omega(1)$, and this gap has been conjectured to be fundamental.

Recently, to bridge the gap, Raghavendra, Rao and Schramm investigate sub-exponential refutation algorithms below the spectral threshold [7]. Their results constitute a smooth trade-off between the density and the time complexity required for certifying unsatisfiability. More specifically, their algorithm parametrized by d achieves the following performance: For all $\delta \in [0, 1)$, their algorithm with $d = n^\delta$ finds a certificate at density $\widetilde{\Omega}(n^{(k/2-1)(1-\delta)})$ in time $\exp(\widetilde{O}(n^\delta))$. At $\delta \approx 0$, their result recovers the polynomial-time strong refutation result at the spectral threshold, while at $\delta \approx 1$, their result recovers an exponential-time strong refutation at the information-theoretic threshold.

This beautiful result, however, relies on an intricate analysis spanning over 20 pages as well as technical complications in algorithm steps, raising a question of whether one can simplify the analysis as well as the algorithm. This work addresses this question as follows:

1. This work simplifies the key technical component of the analysis in [7] (Section 4). More specifically, the spectral norm analysis [7, Theorem 4.4] is significantly simplified in this work relying on more elementary combinatorial arguments.
2. In addition, for even k , this work also simplifies their refutation algorithm by modifying the technical preprocessing step (Section 5). At a high level, the previous work requires $O(d)$ spectral norm computations of the matrix of size $n^{O(d) \times O(d)}$, whereas the approach in this paper only requires a *single* computation.

As a byproduct of our simpler approach, the theoretical guarantee in this paper comes with less technical conditions and enjoys better refutation performances.

2 Preliminary: spectral strong refutation algorithms

To set the stage for our main result, we first briefly review the spectral refutation algorithms in the prior works [5, 2, 3] that achieve the spectral threshold. For illustrative purpose, we focus throughout on the case when k is even; indeed, the odd k case follows similarly modulo some extra “tricks” to reduce it to the even case (see e.g. [2, Appendix A.2] for details).

We first represent the strong refutation problem as the problem of certifying an upper bound on a polynomial.

Definition 3 (Constraints tensor). Given a set of constraints Φ consisting of m constraints C_{S_1}, \dots, C_{S_m} , the *constraints tensor* of Φ is a n^k tensor \mathbf{T}^Φ defined as $\mathbf{T}_S^\Phi = \eta_{S_a}$ if $S = S_a$ for some $a = 1, \dots, m$ and $\mathbf{T}_S^\Phi = 0$ otherwise.

Definition 4 (Constraints polynomial). Given a set of constraints Φ , the *constraints polynomial* of Φ is a k -degree homogeneous polynomial f^Φ defined as $f^\Phi(x) := \langle \mathbf{T}^\Phi, x^{\otimes k} \rangle$.

Having the above definitions, it is straightforward to verify the following identity:

$$\begin{aligned} 2m \cdot \left(P_\Phi(x) - \frac{1}{2} \right) &= \sum_{\ell=1}^m \eta_{(i_\ell, j_\ell)} x_{i_\ell} x_{j_\ell} = \langle \mathbf{T}^\Phi, x^{\otimes k} \rangle = f^\Phi(x) \\ \iff P_\Phi(x) &= \frac{1}{2} + \frac{1}{2m} \cdot f^\Phi(x). \end{aligned} \quad (1)$$

Having established (1), the strong refutation problem turns into the problem of certifying a good upper bound on the constraints polynomial:

$$\max_{x \in \{\pm 1\}^n} P_\Phi(x) \leq \frac{1}{2} + o(1) \iff \max_{x \in \{\pm 1\}^n} f^\Phi(x) = o(m). \quad (2)$$

Now, the key idea of the spectral refutation is to certify an upper bound on the constraints polynomial by first computing a *matrix representation* of the polynomial and then computing the spectral norm of the matrix². We first formally define matrix representations:

Definition 5 (Matrix representation). We say an $n^{k/2} \times n^{k/2}$ matrix M is a matrix representation of a degree- k homogeneous polynomial f if we have $f(x) = (x^{\otimes k/2})^\top M x^{\otimes k/2}$. Here and below, we use $x^{\otimes k/2}$ to denote its vector flattening³.

If we have a matrix representation M of the constraints polynomial f^Φ , one can certify an upper bound by computing the spectral norm of the matrix representation:

$$\max_{x \in \{\pm 1\}^n} f^\Phi(x) = \max_{x \in \{\pm 1\}^n} (x^{\otimes k/2})^\top M x^{\otimes k/2} \leq n^{k/2} \|M\|, \quad (3)$$

where the inequality follows since $\|x^{\otimes k/2}\| = \sqrt{n^{k/2}}$.

Having (3), it is now crucial to find a matrix representation that results in a small spectral norm. It turns out that to achieve the spectral threshold, a simple matrix representation suffices. Let us denote by \mathbf{M}^Φ the natural $n^{k/2} \times n^{k/2}$ flattening of the constraints tensor \mathbf{T}^Φ . Certainly \mathbf{M}^Φ is a matrix representation, and hence, its symmetrization is also a matrix representation:

²We remark that many spectral methods in the literature can be understood as following this principle of computing the spectral norm of a matrix representation. We refer readers to [1, Section 9] for more examples of matrix representations arising in statistical problems.

³More formally, we regard $x^{\otimes k/2}$ as a vector of dimension $nk/2$ rather than as a $n^{k/2}$ tensor.

Definition 6 (Symmetric matrix representation). $S^\Phi := \frac{1}{2}[\mathbf{M}^\Phi + (\mathbf{M}^\Phi)^\top]$.

Indeed, it follows from a standard result in random matrix theory that the symmetric matrix representation S^Φ constructed from random k -XOR has the spectral norm $o(m)$ with high probability as soon as the density is above the spectral threshold, i.e., $\alpha = \tilde{\Omega}(n^{k/2-1})$ (see e.g. [2, Appendix A.1] for precise details).

Thus far, we present the spectral refutation algorithms in the prior arts that achieve the spectral threshold. Now, we move on to the result due to Raghavendra, Rao and Schramm [7]. It turns out that for strong refutation below the spectral threshold, one needs to rely on a higher-order symmetry. This will be the subject of the next section.

3 Higher-order symmetry for refutation below spectral threshold

In this section, we discuss the approach based on a higher-order symmetry due to Raghavendra, Rao and Schramm [7]. We remark that a similar technique was independently developed by Bhattachiprolu, Guruswami and Lee [4] under the context of finding an upper bound certificate of the tensor injective norm.

3.1 Higher-order type-symmetric matrix representation

To illustrate the main idea, we first define the types of the entries:

Definition 7 (Histogram tuples). Let $\text{hist}(I)$ be the n -tuple $(\alpha_1, \dots, \alpha_n)$ such that α_i is the number of times i appears in I , i.e., the histogram of the tuple I . Let $\text{hist}(I)! := \prod_{i=1}^n (\alpha_i)!$, where $0! = 1$ by convention.

Definition 8 (Types of entries). Given a matrix representation M of a constraints polynomial f^Φ , we say two entries $M_{I,J}$ and $M_{I',J'}$ have the same *type* if $\text{hist}(I) = \text{hist}(I')$ and $\text{hist}(J) = \text{hist}(J')$, i.e., for all $i \in [n]$, the number of i 's appearing in I (resp. J) is equal to that in I' (resp. J').

With this definition, one can easily notice that the entries of the same type corresponds to the coefficient of the same monomial in f^Φ . Now the key idea of [7] is to consider a matrix representation which distributes the coefficient of a monomial in f^Φ *equally* across the corresponding type of entries. It turns out that such a matrix representation has small spectral norm, resulting in a better refutation certificate.

Definition 9 (Type-symmetric matrix representation). We say a matrix representation is type-symmetric if the entries of the same type have the same value.

To maximize the gain from a type-symmetric matrix representation, [7] indeed considers a higher order matrix representation, which amounts to working with $(f^\Phi)^d$ instead of f^Φ for some $d > 1$ at the cost of increased computational complexity. Given a type-symmetric matrix representation $\mathbf{R}^{\Phi,d}$ of $(f^\Phi)^d$ (we defer the formal definition to Definition 11), we have $f^\Phi(x)^d = (x^{\otimes kd/2})^\top \mathbf{R}^{\Phi,d} x^{\otimes kd/2}$ from which one can conclude

$$\max_{x \in \{\pm 1\}^n} f^\Phi(x) = n^{k/2} \cdot \max_{x \in \frac{1}{\sqrt{n}} \cdot \{\pm 1\}^n} \left[(x^{\otimes kd/2})^\top \mathbf{R}^{\Phi,d} x^{\otimes kd/2} \right]^{1/d} \leq n^{k/2} \cdot \left\| \mathbf{R}^{\Phi,d} \right\|^{1/d}. \quad (4)$$

However, as mentioned in [7, Section 4], it turns out that the inequality in (4) is not tight enough for the desired result. To overcome this issue, [7] suggested the technique of removing high multiplicity rows/columns. This will be the subject of the next subsection.

3.2 Overcoming challenge with trimming rows/columns

Before getting into the technique in [7], let us first discuss why the inequality in (4) is not tight. The main reason for the looseness is the fact that the left hand side of the inequality is the maximum over the specific unit vectors of the form $\frac{1}{\sqrt{n}} \cdot \{\pm 1\}^n$, while the spectral norm certificate finds the maximum over all unit vectors. In particular, if the maximum of the spectral norm is achieved by a sparse vector, this certificate would no longer provide a good upper bound.

To cope with this issue, [7] employs the *trimming* step, in which they remove rows and columns of $\mathbf{R}^{\Phi, d}$ corresponding to index tuples with high multiplicities, i.e., I 's such that coordinate values of $\text{hist}(I)$ are large. This technical step indeed results in a better spectral norm bound as we shall see in Section 4.2.

3.3 Technical challenge of the approach in Raghavendra-Rao-Schramm

However, it turns out that analyzing this higher-order method with the trimming step is rather technical:

1. Note that the construction of symmetric matrix representation results in a rather complicated dependency structure across entries, making it hard to analyze its spectral norm. Indeed, the spectral norm analysis [7, Theorem 4.4] constitutes the main technical component of the analysis in [7].
2. Moreover, it turns out that justifying the validity of the trimming step also requires some technical modification of the algorithm together with an additional careful analysis. At a high level, these complications arise due to the fact that the trimmed matrix is no longer a matrix representation of the constraints polynomial. In particular, their approach requires computations of $O(d)$ spectral norms of matrices of size $n^{O(d) \times O(d)}$.

We will address the above challenges in order in the subsequent sections.

4 A simpler spectral norm analysis

In this section, we provide a simpler spectral norm analysis of the symmetric matrix representation. As we mentioned in the previous section, the symmetric matrix representation has an intricate dependency structure between entries and hence the standard tools such as matrix Chernoff bound [8] does not apply. Hence, we need to rely on more direct analysis based on the trace power method:

Proposition 1 (Trace power method). *Let $n, \ell \in \mathbb{N}$, let $c \in \mathbb{R}$, and let M be a symmetric $n \times n$ random matrix. Then,*

$$\mathbb{E} \text{Tr}(M^{2\ell}) \leq \beta \implies \Pr\left(\|M\| \leq c \cdot \beta^{1/2\ell}\right) \geq 1 - c^{-2\ell}.$$

Proof. The proof follows from the fact that $\|M\|^{2\ell} = \|M^{2\ell}\| \leq \text{Tr}(M^{2\ell})$ together with Markov's inequality. \square

Hence, to come up with a probabilistic upper bound on the spectral norm, one need to bound the trace power term. However, in contrast to well-known settings in random matrix theory, our matrix

of interest M has $kd/2$ -tuples for its row/column indices, which renders computing the trace power term more complicated. In particular, for an integer ℓ , the trace power term can be represented as

$$\sum_{I^{(1)}, \dots, I^{(2\ell)} \in [n]^{kd/2}} \mathbb{E} \left[\prod_{j=1}^{2\ell} M_{I^{(j)}, I^{(j+1)}} \right], \quad (5)$$

where indices are read modulo- 2ℓ , i.e., $I_{2\ell+1}$ denotes I_1 . As a warm-up, we first analyze the symmetric matrix representation, i.e. $M = \mathbf{R}^{\Phi, d}$.

4.1 Warm-up: analysis for higher-order type-symmetric matrix

In this section we apply the trace power method to $M = \mathbf{R}^{\Phi, d}$ as a warm-up. Let us first formally define $\mathbf{R}^{\Phi, d}$. To that end, we first recall the symmetric matrix representation S^Φ . By its definition (Definition 6), S^Φ is a $n^{k/2} \times n^{k/2}$ symmetric random matrix with independent mean-zero entries taking values in $[-1, 1]$. Now, let $S^{\Phi, d}$ be the d -th Kronecker power of S^Φ , i.e., for $k/2$ -tuples U_1, \dots, U_d and V_1, \dots, V_d ,

$$S_{(U_1, \dots, U_d), (V_1, \dots, V_d)}^{\Phi, d} = S_{U_1, V_1}^\Phi \times S_{U_2, V_2}^\Phi \times \dots \times S_{U_d, V_d}^\Phi. \quad (6)$$

Now, the symmetric matrix representation is obtained from $S^{\Phi, d}$ by replacing each entry with the average of the entries of the same type as the corresponding entry. To formally define, we begin with some notations:

Definition 10 (Permutations). For each positive integers n, q and $I = (i_1, \dots, i_q) \in [n]^q$, let \mathbb{S}_q be the set of permutations on $[q]$. For a permutation $\pi \in \mathbb{S}_q$ and a subtuple $U = (i_{j_1}, \dots, i_{j_\ell})$ of I , let $\pi(U) := (i_{\pi(j_1)}, \dots, i_{\pi(j_\ell)})$.

Now based on these notations, we formally define $\mathbf{R}^{\Phi, d}$ as follows:

Definition 11 (Higher-order symmetric matrix representation). For an even integer k and $d \geq 1$, $\mathbf{R}^{\Phi, d}$ is an $n^{kd/2 \times kd/2}$ matrix representation of $(f^\Phi)^d$ defined as

$$\mathbf{R}_{I, J}^{\Phi, d} = \frac{1}{|\mathbb{S}_{kd/2}|^2} \sum_{\pi, \sigma \in \mathbb{S}_{kd/2}} S_{\pi(I), \sigma(J)}^{\Phi, d}. \quad (7)$$

Now having the formal definition of $\mathbf{R}^{\Phi, d}$, one can write the trace power term (5) as follows (where we write each $kd/2$ -tuple as $I^{(\cdot)} = (U_1^{(\cdot)}, U_2^{(\cdot)}, \dots, U_d^{(\cdot)})$):

$$\begin{aligned} & \frac{1}{|\mathbb{S}_{kd/2}|^{4\ell}} \cdot \sum_{\substack{I^{(j)} \in [n]^{kd/2} \\ j=1, \dots, 2\ell}} \sum_{\substack{\pi_j, \sigma_j \in \mathbb{S}_{kd/2} \\ j=1, \dots, 2\ell}} \mathbb{E} \left[\prod_{j=1}^{2\ell} S_{\pi_j(I^{(j)}), \sigma_j(I^{(j+1)})}^{\Phi, d} \right] \\ &= \frac{1}{|\mathbb{S}_{kd/2}|^{4\ell}} \cdot \sum_{\substack{I^{(j)} \in [n]^{kd/2} \\ j=1, \dots, 2\ell}} \sum_{\substack{\pi_j, \sigma_j \in \mathbb{S}_{kd/2} \\ j=1, \dots, 2\ell}} \mathbb{E} \left[\prod_{j=1}^{2\ell} \prod_{s=1}^d S_{\pi_j(U_s^{(j)}), \sigma_j(U_s^{(j+1)})}^\Phi \right], \end{aligned} \quad (8)$$

Although (8) looks quite complicated, note that one can actually simplify it further.

Definition 12 (Partition of the index set according equality). Given $\{I^{(j)}\}$, $\{\pi_j\}$ and $\{\sigma_j\}$ ($j = 1, \dots, 2\ell$), we define $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})$ to be the partition of the index set $\mathcal{I} := \{(j, s) : j = 1, \dots, 2\ell, s = 1, \dots, d\}$ according to the equivalence relation

$$(j, s) \sim (j', s') \Leftrightarrow \begin{cases} (\pi_j(U_s^{(j)}), \sigma_j(U_s^{(j+1)})) = (\pi_{j'}(U_{s'}^{(j')}), \sigma_{j'}(U_{s'}^{(j'+1)})) \text{ or} \\ (\pi_j(U_s^{(j)}), \sigma_j(U_s^{(j+1)})) = (\sigma_{j'}(U_{s'}^{(j'+1)}), \pi_{j'}(U_{s'}^{(j')})). \end{cases}$$

We denote by $|\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})|$ the number of equivalence classes in the partition.

Since S^Φ is a symmetric random matrix with mean zero entries, it follows that the summand in (8) corresponding to $\{I^{(j)}\}$, $\{\pi_j\}$ and $\{\sigma_j\}$ is equal to zero if the partition $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})$ contains an equivalence class of odd size.

Hence, in order to have a nonzero summand, every equivalence class of $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})$ must have even size.

Definition 13. Given $\{I^{(j)}\}$, $\{\pi_j\}$ and $\{\sigma_j\}$ ($j = 1, \dots, 2\ell$), we say the partition of the index set $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})$ is called an *even* partition if all equivalence classes have even size.

When $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})$ is even, since each entry of S^Φ is in $[-1, 1]$, one can easily upper bound the summand of (8) explicitly in terms of the number of equivalence classes:

$$\mathbb{E} \left[\prod_{j=1}^{2\ell} \prod_{s=1}^d S_{\pi_j(U_s^{(j)}), \sigma_j(U_s^{(j+1)})}^\Phi \right] \leq p^{|\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})|}. \quad (9)$$

Using the upper bound (9), and grouping the trace power term so that each group contains the summand corresponding to the same partition, we obtain

$$(8) \leq \frac{1}{|\mathbb{S}_{kd/2}|^{4\ell}} \sum_{\mathcal{Q}: \text{even}} \left[p^{|\mathcal{Q}|} \cdot \text{Num}(\mathcal{Q}) \right], \quad (10)$$

where $\text{Num}(\mathcal{Q}) := |\{(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\}) : \text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\}) = \mathcal{Q}\}|$. Therefore, to upper bound the trace power term, one needs to upper-estimate $\text{Num}(\mathcal{Q})$ for each \mathcal{Q} . Although the counting $\text{Num}(\mathcal{Q})$ looks complicated, the symmetry saves the day.

Definition 14. $\text{Num}(\mathcal{Q} \mid \{\pi_j\}) := |\{(\{I^{(j)}\}, \{\sigma_j\}) : \text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\}) = \mathcal{Q}\}|$.

First, one can easily verify the following based on a simple symmetry argument (here Id denotes the identity permutation in $\mathbb{S}_{kd/2}$):

Claim 1. $\text{Num}(\mathcal{Q} \mid \{\text{Id}\}) = \text{Num}(\mathcal{Q} \mid \{\pi_j\})$ for any $\{\pi_j\}$.

Proof. See Section A.1. □

Due to Claim 1, it follows that:

$$\text{Num}(\mathcal{Q}) = |\mathbb{S}_{kd/2}|^{2\ell} \cdot \text{Num}(\mathcal{Q} \mid \{\text{Id}\}). \quad (11)$$

Hence, with this argument, we reduce the problem of counting triples $(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\})$ into the problem of counting pairs $(\{I^{(j)}\}, \{\sigma_j\})$. Now let us further reduce the problem. To that end, we first define:

Definition 15. We say a collection of index tuples $\{I^{(j)}\}$ is \mathcal{Q} -valid if there exist $\{\sigma_j\}$ such that $\text{Par}(\{I^{(j)}\}, \{\text{Id}\}, \{\sigma_j\})$ is equal to \mathcal{Q} .

Claim 2. For any \mathcal{Q} -valid $\{I^{(j)}\}$, there are at most $(kd/2)^{k|\mathcal{Q}|/2} \cdot \prod_{j=1}^{2\ell} \text{hist}(I^{(j)})!$ different $\{\sigma_j\}$'s such that $\text{Par}(\{I^{(j)}\}, \{\text{Id}\}, \{\sigma_j\}) = \mathcal{Q}$.

Proof. The proof is based on an elementary counting argument. See Section A.2. \square

Due to Claim 2, now we have:

$$\text{Num}(\mathcal{Q}) \leq |\mathbb{S}_{kd/2}|^{2\ell} \cdot (kd/2)^{k|\mathcal{Q}|/2} \cdot \sum_{\{I^{(j)}\} : \mathcal{Q}\text{-valid}} \left[\prod_{j=1}^{2\ell} \text{hist}(I^{(j)})! \right]. \quad (12)$$

Putting this back to (10), we obtain the following result:

Theorem 1. For even k and $d \geq 1$, let $\mathbf{R}^{\Phi, d}$ be the $n^{kd/2 \times kd/2}$ higher-order symmetric matrix representation (Definition 11) of random k -XOR. Then, the following upper bound on the trace power term holds:

$$\mathbb{E} \text{Tr}((\mathbf{R}^{\Phi, d})^{2\ell}) \leq \frac{1}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{\mathcal{Q}: \text{even}} \left[\left(p(kd/2)^{k/2} \right)^{|\mathcal{Q}|} \cdot \sum_{\substack{\{I^{(j)}\} : \\ \mathcal{Q}\text{-valid}}} \left[\prod_{j=1}^{2\ell} \text{hist}(I^{(j)})! \right] \right].$$

Having established Theorem 1, one can slightly modify the proof to handle the trimmed matrix from Section 3.2. This will be the focus of the next subsection.

4.2 A simpler spectral norm analysis of the trimmed matrix

Having established Theorem 1, which explicitly characterizes the upper bound on the trace power term in terms of $\text{hist}(I^{(j)})!$'s, one can now quantitatively understand the trimming technique due to Raghavendra, Rao and Schramm [7]. In particular, we will shortly demonstrate that our Theorem 1 recovers the main technical result [7, Theorem 4.4]. This is remarkable as our proof is much simpler than the original proof in [7].

The problem with the upper bound in Theorem 1 is that the value $\text{hist}(I^{(j)})!$ could be in general large. For instance, if $I^{(j)}$ is the $kd/2$ -tuple consisting only of index 1, then $\text{hist}(I^{(j)})! = (kd/2)!$, which turns out to be too large for our desired result. Now having observed this, one can now see how the trimming preprocessing of [7] helps reduce the spectral norm: by removing rows/columns corresponding to the index tuples with high multiplicities, one can significantly reduce the upper bound. More formally, following [7], if we remove the rows/columns corresponding to the index tuples I 's such that $\text{hist}(I)$ has a coordinate value larger than $10 \log n$, we have the following:

Corollary 1. For even k and $d \geq 1$, let $\mathbf{R}^{\Phi, d, \text{trim}}$ be the $n^{kd/2 \times kd/2}$ matrix obtained from the $\mathbf{R}^{\Phi, d}$ (Definition 11) by removing all rows/columns I 's such that $\text{hist}(I)$ has a coordinate value larger than $10 \log n$. Assume that $d^{k/2-1} n^{k/2} p > 1$. Then, the following spectral norm bound holds with probability at least $1 - n^{-2}$:

$$\left\| \mathbf{R}^{\Phi, d, \text{trim}} \right\|^{1/d} \leq c \cdot \frac{e^{3k/4} \cdot 10^{5k/2}}{(k/2)^{k/4}} \cdot \frac{n^{k/4} p^{1/2}}{d^{(k-2)/4}} \cdot \log^{5k/2+1} n.$$

for some absolute constant $c > 0$.

Remark 2. Although we focus on the even k case throughout the proof for simplicity, we note that a similar argument applies to the case of odd k following the “tricks” [7, Section 4.2] based on Cauchy Schwartz inequality. Consequently, our proof technique provides a simpler proof of the main technical statement for the odd k case [7, Theorem 4.13]

Proof of Corollary 1. From Theorem 1, we have the following upper bound on the trace power term:

$$\mathbb{E} \operatorname{Tr}((\mathbf{R}^{\Phi, d, \text{trim}})^{2\ell}) \leq \frac{1}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{\mathcal{Q}: \text{ even}} \left[\left(p(kd/2)^{k/2} \right)^{|\mathcal{Q}|} \cdot \sum_{\substack{\{I^{(j)}\} : \\ \mathcal{Q}\text{-valid}}} \left[\prod_{j=1}^{2\ell} \operatorname{hist}(I^{(j)})! \right] \right].$$

On the other hand, due to the trimming procedure, each coordinate value of the tuple $\operatorname{hist}(I^{(j)})$ is upper bounded by $10 \log n$, from which we have the following upper bound on the $\operatorname{hist}(I^{(j)})!$:

$$\operatorname{hist}(I^{(j)})! \leq ((10 \log n)!)^{kd/2} \leq (10 \log n)^{5kd \log n} \leq n^{5kd \log(10 \log n)}.$$

The trimming step gives us an uniform upper bound on $\operatorname{hist}(I^{(j)})!$, and hence, it suffices to upper bound the number of \mathcal{Q} -valid $\{I^{(j)}\}$'s:

Claim 3. *For any even partition \mathcal{Q} , there are at most $n^{k(|\mathcal{Q}|+d)/2}$ \mathcal{Q} -valid $\{I^{(j)}\}$'s.*

Proof. The proof is elementary. See Section A.3. □

Due to Claim 3, the upper bound on the trace power term becomes:

$$\begin{aligned} \mathbb{E} \operatorname{Tr}((\mathbf{R}^{\Phi, d, \text{trim}})^{2\ell}) &\leq \frac{n^{5kd \log(10 \log n)}}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{\mathcal{Q}: \text{ even}} \left[\left(p(kd/2)^{k/2} \right)^{|\mathcal{Q}|} \cdot n^{k(|\mathcal{Q}|+d)/2} \right] \\ &= \frac{n^{5kd \log(10 \log n) + kd/2}}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{M=1}^{d\ell} \left[N_M \cdot \left(p(nkd/2)^{k/2} \right)^M \right], \end{aligned} \quad (13)$$

where N_M is the number of even partitions of size M and we have $M \leq d\ell$ in the range of summation since an even partition has size at most $d\ell$. Thus, the last ingredient is to bound the number of even partitions:

Claim 4. $N_M \leq \binom{2d\ell}{M} \cdot M^{2d\ell-M}$ for all $1 \leq M \leq d\ell$.

Proof. The first term in the upper bound accounts for the number of different ways of choosing M representative indices in \mathcal{I} , and the second term counts the number of different ways of assigning the other indices to the M representative elements. □

Due to Claim 4, the upper bound (13) becomes:

$$\frac{n^{5kd \log(10 \log n) + kd/2}}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{M=1}^{d\ell} \left[\binom{2d\ell}{M} \cdot M^{2d\ell-M} \cdot \left(p(nkd/2)^{k/2} \right)^M \right]. \quad (14)$$

Having established (14), the rest of the proof is straightforward calculations. We first upper bound each term in the above summand as follows: (i) $\binom{2d\ell}{M} \leq 2^{2d\ell}$, (ii) $M^{2d\ell-M} \leq (d\ell)^{2d\ell-M} \leq$

$d^{2d\ell-M} \ell^{2d\ell}$, and (iii) $((kdn/2)^{k/2} p)^M \leq ((dn)^{k/2} p)^M \cdot (k/2)^{kd\ell/2}$. Then, the summand in (14) is upper bounded by

$$2^{2d\ell} \cdot d^{2d\ell-M} \ell^{2d\ell} \cdot (k/2)^{kd\ell/2} \left((dn)^{k/2} p \right)^M = (2\ell)^{2d\ell} (k/2)^{kd\ell/2} \cdot d^{2d\ell} \left(d^{k/2-1} n^{k/2} p \right)^M.$$

Using this upper bound, it follows that

$$\begin{aligned} (14) &\leq \frac{(2\ell)^{2d\ell} (k/2)^{kd\ell/2} n^{5kd \log(10 \log n) + kd/2} d^{2d\ell}}{|\mathbb{S}_{kd/2}|^{2\ell}} \cdot \sum_{M=1}^{d\ell} \left(d^{k/2-1} n^{k/2} p \right)^M \\ &\leq \frac{(2\ell)^{2d\ell} (k/2)^{kd\ell/2} e^{kd\ell} n^{5kd \log(10 \log n) + kd/2} d^{2d\ell}}{(kd/2)^{kd\ell}} \cdot d\ell \cdot \left(d^{k/2-1} n^{k/2} p \right)^{d\ell}, \end{aligned} \quad (15)$$

where the inequality follows from the facts that $|\mathbb{S}_{kd/2}|^{2\ell} = ((kd/2)!)^{2\ell} \geq \left(\frac{kd/2}{e}\right)^{kd\ell}$ ($\because n! \geq (n/e)^n$) and $d^{k/2-1} n^{k/2} p > 1$. Reorganizing terms in (15), we obtain

$$(2\ell)^{2d\ell+1} e^{kd\ell} (k/2)^{-kd\ell/2} d^{-kd\ell/2+d\ell+1} n^{kd\ell/2+5kd \log(10 \log n) + kd/2} p^{d\ell}.$$

Invoking Proposition 1 and using the fact that $f(x) = x^{1/x}$ is bounded on $[1, \infty)$, $\|M\|^{1/d}$ is upper bounded by

$$c \cdot (2\ell) e^{k/2} (k/2)^{-k/4} d^{-k/4+1/2} n^{k/4+5k \log(10 \log n)/(2\ell)+k/(4\ell)} p^{1/2}$$

with probability at least $1 - e^{-2\ell}$ for some absolute constant $c > 0$. Choosing $\ell = \log n$, we complete the proof. \square

Thus far, we have addressed the first challenge in Section 3.3 by developing a simpler spectral norm analysis of the type-symmetric representation as well as the trimmed matrix. Now, we move on to the second challenge: as mentioned in Section 3.3, the trimmed matrix $\mathbf{R}^{\Phi, d, \text{trim}}$ is no longer a matrix representation of $(f^\Phi)^d$, it requires additional non-trivial modifications of the algorithm steps as well as analysis.

5 A simpler spectral refutation with re-scaling entries

In this section, we address the second challenge from Section 3.3 and develop a simpler spectral refutation algorithm. Our main idea is to re-scale the rows/columns of $\mathbf{R}^{\Phi, d}$. To describe our re-scaling step, we first revisit the upper bound from Theorem 1:

$$\mathbb{E} \text{Tr}((\mathbf{R}^{\Phi, d})^{2\ell}) \leq \frac{1}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{\mathcal{Q}: \text{even}} \left[\left(p(kd/2)^{k/2} \right)^{|\mathcal{Q}|} \cdot \sum_{\substack{\{I^{(j)}\} \\ \mathcal{Q}\text{-valid}}} \left[\prod_{j=1}^{2\ell} \text{hist}(I^{(j)})! \right] \right]. \quad (16)$$

As we have discussed in Section 4.2, we need to cancel out the $\text{hist}(I^{(j)})!$ terms in the bound to reduce the spectral norm. Our approach is to appropriately re-scale $\mathbf{R}^{\Phi, d}$ so that one can remove the $\prod_{j=1}^{2\ell} \text{hist}(I^{(j)})!$ terms in the upper bound (16). In particular, if we divide the (I, J) -th entry of $\mathbf{R}^{\Phi, d}$ by $\sqrt{\text{hist}(I)! \cdot \text{hist}(J)!}$, the $\prod_{j=1}^{2\ell} \text{hist}(I^{(j)})!$ term will be exactly canceled out by the re-scaling. More formally, we define the following vector and its corresponding diagonal matrix:

Definition 16 (Re-scaling factors). Let hist be an $n^{kd/2}$ -dimensional vector whose I -th coordinate is defined as $\text{hist}_I := \sqrt{\text{hist}(I)!}$ for each $I \in [n]^{kd/2}$. We define D_{hist} to be an $n^{kd/2} \times n^{kd/2}$ diagonal matrix whose (I, I) -th entry is defined as hist_I .

Using Definition 16, one can precisely achieve the re-scaling discussed above as follows:

Definition 17 (Re-scaled matrix representation). $\mathbf{R}^{\Phi, d, \text{rescale}} := D_{\text{hist}}^{-1} \cdot \mathbf{R}^{\Phi, d} \cdot D_{\text{hist}}$.

Then, following the same proof as that of Corollary 1, one can prove the following spectral norm bound:

Corollary 2. For even k and $d \geq 1$, let $\mathbf{R}^{\Phi, d, \text{rescale}}$ be the $n^{kd/2 \times kd/2}$ matrix obtained from the $\mathbf{R}^{\Phi, d}$ by re-scaling the rows/columns as per (17). Assume that $d^{k/2-1}n^{k/2}p > 1$. Then, the following spectral norm bound holds with probability at least $1 - n^{-2}$:

$$\left\| \mathbf{R}^{\Phi, d, \text{rescale}} \right\|^{1/d} \leq c \cdot \frac{e^{3k/4}}{(k/2)^{k/4}} \cdot \frac{n^{k/4}p^{1/2}}{d^{(k-2)/4}} \cdot \log n.$$

for some absolute constant $c > 0$.

Remark 3. Note that the spectral norm bound in Corollary 2 is better than the bound due to the trimming step (Corollary 1). This improvement actually leads to a better strong refutation guarantee as we shall see in Theorem 2. Also see Section 6 for an extensive comparison with [7].

Proof. Due to the re-scaling factor, following the proof of Theorem 1, we obtain the following bound on the trace bower term without the $\prod_{j=1}^{2\ell} \text{hist}(I^{(j)})!$ term:

$$\mathbb{E} \text{Tr}((\mathbf{R}^{\Phi, d, \text{rescale}})^{2\ell}) \leq \frac{1}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{\mathcal{Q}: \text{even}} \left[\left(p(kd/2)^{k/2} \right)^{|\mathcal{Q}|} \cdot \sum_{\substack{\{I^{(j)}\} : \\ \mathcal{Q}\text{-valid}}} [1] \right].$$

Now due to Claims 3 and 4, one can further upper bound the trace power term by

$$\frac{n^{kd/2}}{|\mathbb{S}_{kd/2}|^{2\ell}} \sum_{M=1}^{d\ell} \left[\binom{2d\ell}{M} \cdot M^{2d\ell-M} \cdot \left(p(nkd/2)^{k/2} \right)^M \right], \quad (17)$$

which is better than (14) by a multiplicative factor of $n^{5kd \log(10 \log n)}$. Now, following the exact same calculations as in the proof of Corollary 1 and choosing $\ell = \log n$, one can easily notice that the improvement by a multiplicative factor of $n^{5kd \log(10 \log n)}$ results in an improvement in the final bound by a multiplicative factor of $n^{5k \log(10 \log n)/(2\ell)} = n^{5k \log(10 \log n)/(2 \log n)} = (10 \log n)^{5k/2}$, which completes the proof. \square

With this re-scaled matrix $\mathbf{R}^{\Phi, d, \text{rescale}}$, one can also easily come up with a valid certificate for strong refutation (Definition 2):

Proposition 2. For any k -XOR instance Φ and assignment $x \in \{\pm 1\}^n$, we have

$$\left| P_{\Phi}(x) - \frac{1}{2} \right| \leq \frac{1}{2m} \left[\left\| \mathbf{R}^{\Phi, d, \text{rescale}} \right\| \cdot \left(\sum_{I \in [n]^{kd/2}} \text{hist}(I)! \right) \right]^{1/d}.$$

In other words, $\frac{1}{2} + \frac{1}{2m} \left[\left\| \mathbf{R}^{\Phi, d, \text{rescale}} \right\| \cdot \left(\sum_{I \in [n]^{kd/2}} \text{hist}(I)! \right) \right]^{1/d}$ is a valid certificate for strong refutation.

Proof. First, since $\mathbf{R}^{\Phi,d}$ is a matrix representation of $(f^\Phi)^d$, we have

$$f^\Phi(x)^d = (x^{\otimes kd/2})^\top \mathbf{R}^{\Phi,d} x^{\otimes kd/2}.$$

Hence, it follows that

$$\begin{aligned} f^\Phi(x)^d &= (D_{\text{hist}} x^{\otimes kd/2})^\top \cdot D_{\text{hist}} D_{\text{hist}}^{-1} \cdot \mathbf{R}^{\Phi,d} \cdot D_{\text{hist}}^{-1} D_{\text{hist}} \cdot x^{\otimes kd/2} \\ &= (D_{\text{hist}} x^{\otimes kd/2})^\top \cdot \mathbf{R}^{\Phi,d,\text{rescale}} \cdot D_{\text{hist}} x^{\otimes kd/2}. \end{aligned}$$

Consequently, we have

$$|f^\Phi(x)^d| \leq \left\| \mathbf{R}^{\Phi,d,\text{rescale}} \right\| \cdot \left\| D_{\text{hist}} x^{\otimes kd/2} \right\|^2 = \left\| \mathbf{R}^{\Phi,d,\text{rescale}} \right\| \cdot \left(\sum_{I \in [n]^{kd/2}} \text{hist}(I)! \right),$$

where the equality is due to the fact that $x^{\otimes kd/2}$ is an $n^{kd/2}$ -dimensional vector with coordinates equal to ± 1 . Therefore, the proposition follows thanks to the identity (1), which reads $P_\Phi(x) = \frac{1}{2} + \frac{1}{2m} \cdot f^\Phi(x)$. \square

Hence, in order to guarantee that the certificate from Proposition 2 works, our last ingredient is to show that the term $(\sum_{I \in [n]^{kd/2}} \text{hist}(I)!)$ is not too large compared to $\|x^{\otimes kd/2}\|^2 = n^{kd/2}$.

Proposition 3. *For even k and $d \geq 1$,*

$$\sum_{I \in [n]^{kd/2}} \text{hist}(I)! = \frac{(kd/2 + n - 1)!}{(n - 1)!} = (kd/2 + n - 1)(kd/2 + n - 2) \cdots n.$$

In particular, if $d \leq n$, we have $\sum_{I \in [n]^{kd/2}} \text{hist}(I)! \leq (k/2 + 1)^{kd/2} n^{kd/2}$.

Proof. We first group the terms in the summation according to the value of $\text{hist}(I)$:

$$\sum_{I \in [n]^{kd/2}} \text{hist}(I)! = \sum_{\substack{(s_1, s_2, \dots, s_n) \in (\mathbb{Z}_{\geq 0})^n: \\ \sum_i s_i = kd/2}} \sum_{\substack{I \in [n]^{kd/2}: \\ \text{hist}(I) = (s_1, s_2, \dots, s_n)}} \prod_{i=1}^n (s_i)!. \quad (18)$$

For each $(s_1, s_2, \dots, s_n) \in (\mathbb{Z}_{\geq 0})^n$, there are $\frac{(kd/2)!}{\prod_{i=1}^n (s_i)!}$ different I 's such that $\text{hist}(I) = (s_1, s_2, \dots, s_n)$. Hence, the right hand side of (18) becomes

$$\sum_{\substack{(s_1, s_2, \dots, s_n) \in (\mathbb{Z}_{\geq 0})^n: \\ \sum_i s_i = kd/2}} (kd/2)! = (kd/2)! \cdot \left| \left\{ (s_1, s_2, \dots, s_n) \in (\mathbb{Z}_{\geq 0})^n : \sum_i s_i = kd/2 \right\} \right|.$$

It is a simple enumerative combinatorics (c.f. stars and bars argument) to show that the number of feasible (s_1, \dots, s_n) 's is equal to $\binom{kd/2+n-1}{n-1} = \binom{kd/2+n-1}{kd/2}$. Therefore, the summation is equal to

$$\binom{kd/2 + n - 1}{kd/2} \cdot (kd/2)! = (kd/2 + n - 1)(kd/2 + n - 2) \cdots n,$$

which completes the proof. \square

Combining what we have obtained thus far, one can consider the following simpler refutation algorithm based on re-scaling entries:

Algorithm 1. *A simpler strong refutation algorithm with parameter d for even k .*

Input: A k -XOR instance Φ on n variables consisting of m clauses C_{S_1}, \dots, C_{S_m} for distinct tuples $S_1, \dots, S_m \in [n]^k$ and a parameter $d \in \mathbb{N}$.

1. Construct a higher-order symmetric matrix representation $\mathbf{R}^{\Phi, d}$ based on the k -XOR instance Φ as per Definition 11.
2. Compute $\mathbf{R}^{\Phi, d, \text{rescale}}$ as per Definition 17.

Output: $\widehat{P}_\Phi := \frac{1}{2} + \frac{1}{2m} \|\mathbf{R}^{\Phi, d, \text{rescale}}\|^{1/d} \cdot \left(\frac{(kd/2+n-1)!}{(n-1)!}\right)^{1/d}$.

Remark 4. A similar idea of re-scaling rows/columns with diagonal matrices to obtain a better certificate also appeared in the MAXCUT literature; see e.g. [6, Theorem 2.2].

Theorem 2. *Let $d \leq n$ be positive integers and k be an even integer. For any instance Φ of k -XOR, the output \widehat{P}_Φ of Algorithm 1 satisfies $|P_\Phi(x) - \frac{1}{2}| \leq \widehat{P}_\Phi - \frac{1}{2}$ for any $x \in \{\pm 1\}^n$. Assume further that Φ is an instance of random k -XOR with probability p (Definition 1). If $p \cdot d^{k/2-1} n^{k/2} > 1$, the following bound holds with probability at least $1 - O(n^{-1})$ for some absolute constant $c > 0$:*

$$\widehat{P}_\Phi - \frac{1}{2} \leq c \cdot \frac{\log n}{\sqrt{d^{k/2-1} n^{k/2} p}} \cdot \frac{e^{3k/4} \cdot (k/2 + 1)^{k/2}}{(k/2)^{k/4}}.$$

In particular, Algorithm 1 with parameter d certifies with high probability that $P_\Phi(x)$ is equal to $1/2 \pm o(1)$ for any $x \in \{\pm 1\}^n$ whenever $n^{k-1} p = \omega((n/d)^{k/2-1} \log^2 n)$.

Proof. First from Proposition 2, we have

$$\left| P_\Phi(x) - \frac{1}{2} \right| \leq \frac{1}{2m} \left[\|\mathbf{R}^{\Phi, d, \text{rescale}}\| \cdot \left(\sum_{I \in [n]^{kd/2}} \text{hist}(I)! \right) \right]^{1/d} = \widehat{P}_\Phi - \frac{1}{2}, \quad (19)$$

where the equality is due to Proposition 3. Hence the first part of the theorem is proved. As for the second part, it follows from Corollary 2 and Proposition 3 that with probability at least $1 - O(n^{-2})$:

$$\left[\|\mathbf{R}^{\Phi, d, \text{rescale}}\| \cdot \left(\sum_{I \in [n]^{kd/2}} \text{hist}(I)! \right) \right]^{1/d} \leq c \cdot \frac{n^{3k/4} p^{1/2}}{d^{(k-2)/4}} \log n \cdot \frac{e^{3k/4} \cdot (k/2 + 1)^{k/2}}{(k/2)^{k/4}} \quad (20)$$

for some absolute constant $c > 0$. Next, it follows from a standard concentration inequality (e.g. Chernoff bound) that with probability at least (say) $1 - n^{-10}$, $m \geq pn^k/2$. Putting these bounds back to (19), we obtain

$$\begin{aligned} \left| P_\Phi(x) - \frac{1}{2} \right| &\leq \frac{1}{pn^k} \cdot c \cdot \frac{n^{3k/4} p^{1/2}}{d^{(k-2)/4}} \log n \cdot \frac{e^{3k/4} \cdot (k/2 + 1)^{k/2}}{(k/2)^{k/4}} \\ &= c \cdot \frac{\log n}{\sqrt{d^{k/2-1} n^{k/2} p}} \cdot \frac{e^{3k/4} \cdot (k/2 + 1)^{k/2}}{(k/2)^{k/4}}, \end{aligned}$$

and hence, the second part of the theorem also follows. \square

6 Comparison with Raghavendra-Rao-Schramm

We compare Algorithm 1 with the refutation algorithm of Raghavendra, Rao and Schramm [7]. First, the algorithm steps in this paper is simpler than that of [7]. As we have discussed earlier, the trimming step in the algorithm of [7] causes some technical complications as the resulting matrix is no longer a matrix representation of $(f^\Phi)^d$. Indeed, their algorithm first constructs matrices of size $n^{kj/2} \times n^{kj/2}$ for $j \in [\delta d, d]$ and computes the spectral norms of those matrices to design a refutation certificate; see [7, Section 4.1.1] for details. This is in stark contrast with Algorithm 1 which only computes the spectral norm of a *single* matrix $\mathbf{R}^{\Phi, d, \text{rescale}}$ of size $n^{kd/2} \times n^{kd/2}$. In addition, while their certificate requires non-trivial analysis [7, Section 4.1.1] to guarantee its validity, the validity of our certificate \widehat{P}_Φ readily follows as we saw in Proposition 2.

As a result of the simpler approach in this paper, the theoretical guarantee in this paper comes with less technical conditions and enjoys a better refutation guarantee as well as density requirement. More specifically, unlike the guarantee in [7], our main theorem does not require a technical condition like $d \log n = O(n)$. Moreover, the density requirement for strong refutation reads $n^{k-1}p = \omega((n/d)^{k/2-1} \log^{2k} n)$ in [7], which is worse than that of this paper by a polylogarithmic factor (recall that the requirement in Theorem 2 reads $n^{k-1}p = \omega((n/d)^{k/2-1} \log^2 n)$). Lastly, even when the density requirement is fulfilled, their refutation guarantee reads $\frac{1}{2} + \gamma + o(1)$ for some constant $\gamma > 0$ that depends on a hyperparameter in the trimming step. On the other hand, this constant γ does not appear in the refutation guarantee of this paper.

7 Conclusion

In this paper, we establish a simpler approach to strong refutation of random k -XOR below the spectral threshold. Our simplification is two-fold. First, we provide a simpler spectral norm analysis of the certificate matrix of the previous work [7] (Section 4). Second, we develop a simple strong refutation algorithm for the even k case (Section 5). Thanks to our simpler approach, our main result (Theorem 2) enjoys a better theoretical guarantee under less assumptions. It is important to note that a recent work by Wein, El Alaoui and Moore also establishes a simpler strong refutation algorithm for random even k -XOR [9, Theorem F.1] with a different approach. Given the successful simplifications for the even k case, it would be interesting to see if one can come up with a simpler strong refutation algorithm for the odd k case.

Acknowledgements

The author thanks Vijay Bhattiprolu for suggesting the idea of re-scaling entries of the matrix representation and anonymous reviewers for valuable comments. The author acknowledges the Kwanjeong Educational Foundation scholarship and also partial support as a graduate research assistant from the NSF Grant (CAREER: 1846088).

References

- [1] Kwangjun Ahn, Dhruv Medarametla, and Aaron Potechin. Graph matrices: Norm bounds and applications. *arXiv*, 2020.

- [2] Sarah R Allen, Ryan O’Donnell, and David Witmer. How to refute a random CSP. In *Proceedings of the 56th FOCS*, pages 689–708. IEEE, 2015.
- [3] Boaz Barak and Ankur Moitra. Noisy tensor completion via the sum-of-squares hierarchy. In *COLT*, pages 417–445, 2016.
- [4] Vijay Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Sum-of-Squares Certificates for Maxima of Random Tensors on the Sphere. In *APPROX/RANDOM 2017*, volume 81, pages 31:1–31:20. LIPIcs, 2017.
- [5] Amin Coja-Oghlan, Andreas Goerdt, and André Lanka. Strong refutation heuristics for random k -SAT. *Combinatorics, Probability and Computing*, 16(1):5–28, 2007.
- [6] Charles Delorme and Svatopluk Poljak. The performance of an eigenvalue bound on the max-cut problem in some classes of graphs. *Discrete Mathematics*, 111(1-3):145–156, 1993.
- [7] Prasad Raghavendra, Satish Rao, and Tselil Schramm. Strongly refuting random CSPs below the spectral threshold. In *Proceedings of the 49th STOC*, pages 121–131. ACM, 2017.
- [8] Joel A Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12(4):389–434, 2012.
- [9] Alexander S Wein, Ahmed El Alaoui, and Cristopher Moore. The kikuchi hierarchy and tensor pca. In *Proceedings of the 60th FOCS*, pages 1446–1468. IEEE, 2019.

A Deferred proofs of claims

A.1 Proof of Claim 1

Recall that Claim 1 reads $\text{Num}(\mathcal{Q} \mid \{\text{Id}\}) = \text{Num}(\mathcal{Q} \mid \{\pi_j\})$ for any $\{\pi_j\}$. Let us arbitrarily fix a collection of permutations $\{\pi_j\}$. The main observation is that for any $\{I^{(j)}\}$ and $\{\sigma_j\}$, we have $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\}) = \text{Par}(\{\pi_j(I^{(j)})\}, \{\text{Id}\}, \{\sigma_j \circ \pi_j^{-1}\})$. This is a straightforward consequence of Definition 12. Hence, there is an one-to-one correspondence between the collection of pairs $(\{I^{(j)}\}, \{\sigma_j\})$ such that $\text{Par}(\{I^{(j)}\}, \{\pi_j\}, \{\sigma_j\}) = \mathcal{Q}$ and the collection such that $\text{Par}(\{I^{(j)}\}, \{\text{Id}\}, \{\sigma_j\}) = \mathcal{Q}$. This concluded the proof.

A.2 Proof of Claim 2

We first restate Claim 2: for any \mathcal{Q} -valid $\{I^{(j)}\}$, there are at most $(kd/2)^{k|\mathcal{Q}|/2} \cdot \prod_{j=1}^{2\ell} \text{hist}(I^{(j)})!$ different $\{\sigma_j\}$ ’s such that $\text{Par}(\{I^{(j)}\}, \{\text{Id}\}, \{\sigma_j\}) = \mathcal{Q}$.

We bound the number of feasible $\{\sigma_j\}$ ’s as we go through the index set $\mathcal{I} = \{(j, s) : j = 1, \dots, 2\ell, s = 1, \dots, d\}$ in the lexicographical order, i.e., $(1, 1), (1, 2), \dots, (1, d), (2, 1), \dots$ and so on. As we read the indices in such an order, we call an index (j, s) *new* if $(U_s^{(j)}, \sigma_j(U_s^{(j+1)}))$ is not equivalent to the previously appeared indices. Consider the indices $(j, 1), (j, 2), \dots, (j, d)$ for a fixed $j \in [2\ell]$. We consider two different scenarios:

1. First, suppose that all indices $(j, 1), (j, 2), \dots, (j, d)$ are old. Then it should be the case that for each $m = 1, \dots, kd/2$, $\sigma_j(m)$ is chosen so that the $\sigma_j(m)$ -th coordinate of $I^{(j+1)}$ respects

the previous appeared equivalent index. Having observed this, it readily follows that there are $\text{hist}(I^{(j+1)})!$ different choices for $\sigma_j(1), \dots, \sigma_j(kd/2)$ considering the permutation.

2. Now, suppose that there are $d_{\text{new}}^{(j)}$ new indices among $(j, 1), (j, 2), \dots, (j, d)$. For simplicity, assume that $U_{j,1}, \dots, U_{j,d_{\text{new}}^{(j)}}$ are new. Choosing the values $\sigma_j(1), \sigma_j(2), \dots, \sigma_j(kd_{\text{new}}^{(j)}/2)$ arbitrarily, there are at most

$$(kd/2)(kd/2 - 1) \cdots (kd/2 - kd_{\text{new}}^{(j)}/2 + 1) \leq (kd/2)^{kd_{\text{new}}^{(j)}/2}$$

different choices for $\sigma_j(1), \sigma_j(2), \dots, \sigma_j(kd_{\text{new}}^{(j)}/2)$. A similar counting to previous case yields that for the remaining values there are at most $\text{hist}(I^{(j+1)})!$ different choices.

Taking a product over all j 's, we complete the proof since $\sum_{m=1}^{2\ell} d_{\text{new}}^{(m)} = |\mathcal{Q}|$. □

A.3 Proof of Claim 3

Let \mathcal{Q} be an even partition. We count the number of possible \mathcal{Q} -valid $\{I^{(j)}\}$'s. First, let us choose $I^{(1)}$ arbitrarily. Note that there are $n^{kd/2}$ different ways of choosing I_1 . Now, consider $I_2, \dots, I_{2\ell}$. Similar to the proof of Claim 2, we will bound the number of feasible choices s we go through the index set $\mathcal{I} = \{(j, s) : j = 1, \dots, 2\ell, s = 1, \dots, d\}$ in the lexicographical order. Again, we call an index (j, s) *new* if $(\pi_j(U_s^{(j)}), \sigma_j(U_s^{(j+1)}))$ is not equivalent to the previously appeared indices.

Note that we only need to consider new indices because the tuples of old indices are fully determined by their previous appearance. We begin with the tuples $(1, 1), (1, 2), \dots, (1, d)$. Whenever we encounter a new tuple, say $(U_s^{(1)}, \sigma_1(U_s^{(2)}))$, we only need to specify $\sigma_1(U_s^{(2)})$ since $I^{(1)}$ is already fully specified. Hence, there are at most $n^{kd_{\text{new}}^{(1)}/2}$ different ways of choosing $I^{(2)}$, where $d_{\text{new}}^{(1)}$ is the number of new indices among $(1, 1), (1, 2), \dots, (1, d)$. By similar arguments, inductively for $j = 2, 3, \dots, 2\ell$, there are at most $n^{kd_{\text{new}}^{(j)}/2}$ different ways of choosing $I^{(j)}$. Taken collectively, we obtain the result since $\sum_{j=1}^{2\ell} d_{\text{new}}^{(j)} = |\mathcal{Q}|$.