

A simple and sharper proof of the hypergraph Moore bound

Jun-Ting Hsieh* Pravesh K. Kothari† Sidhanth Mohanty‡

July 21, 2022

Abstract

The hypergraph Moore bound is an elegant statement that characterizes the extremal trade-off between the girth — the number of hyperedges in the smallest cycle or *even cover* (a subhypergraph with all degrees even) and size — the number of hyperedges in a hypergraph. For graphs (i.e., 2-uniform hypergraphs), a bound tight up to the leading constant was proven in a classical work of Alon, Hoory and Linial [AHL02]. For hypergraphs of uniformity $k > 2$, an appropriate generalization was conjectured by Feige [Fei08]. The conjecture was settled up to an additional $\log^{4k+1} n$ factor in the size in a recent work of Guruswami, Kothari and Manohar [GKM21]. Their argument relies on a connection between the existence of short even covers and the spectrum of a certain randomly signed *Kikuchi* matrix. Their analysis, especially for the case of odd k , is significantly complicated.

In this work, we present a substantially simpler and shorter proof of the hypergraph Moore bound. Our key idea is the use of a new *reweighted* Kikuchi matrix and an *edge deletion* step that allows us to drop several involved steps in [GKM21]’s analysis such as combinatorial bucketing of rows of the Kikuchi matrix and the use of the Schudy–Sviridenko polynomial concentration. Our simpler proof also obtains tighter parameters: in particular, the argument gives a new proof of the classical Moore bound of [AHL02] with no loss (the proof in [GKM21] loses a $\log^3 n$ factor), and loses only a single logarithmic factor for all $k > 2$ -uniform hypergraphs.

As in [GKM21], our ideas naturally extend to yield a simpler proof of the full trade-off for strongly refuting smoothed instances of constraint satisfaction problems with similarly improved parameters.

*Carnegie Mellon University. juntingh@cs.cmu.edu. Supported by NSF CAREER Award #2047933.

†Carnegie Mellon University. praveshk@cs.cmu.edu. Supported by NSF CAREER Award #2047933, Alfred P. Sloan Fellowship and a Google Research Scholar Award.

‡UC Berkeley. sidhanthm@cs.berkeley.edu. Supported by a Google PhD Fellowship.

1 Introduction

What is the maximum girth of a graph on n vertices and average degree d ? For d -regular graphs, a simple “ball growing” argument shows that the graph must have a cycle of length at most $2\log_{d-1} n + 2$. This threshold is called the *Moore bound* [Wik22] (see Page 180 of [Big93]) and graphs achieving it are called Moore graphs. In a classical paper that resolved a question of Bollobás [Bol78], Alon, Hoory and Linial [AHL02] proved that the same upper bound holds even for irregular graphs. Later on, Hoory [Hoo02] obtained a better bound for bipartite graphs and Babu and Radhakrishnan [BR14] found an elegant proof based on the entropy of random walks.

Girth-density trade-offs for hypergraphs. This work is about a natural and well-studied generalization of the Moore bound to $k > 2$ -uniform hypergraphs. A cycle¹ in a hypergraph, more descriptively called an *even cover*, is a collection of hyperedges such that every vertex participates in an even number of them. The girth of a hypergraph is the smallest size of an even cover in it. When specialized to graphs, an even cover is simply a union of cycles and thus, this formulation naturally generalizes the standard notion of girth in graphs.

Analogously to the Moore bound, understanding the maximum number of hyperedges that one can pack in a hypergraph while avoiding an even cover of a given length is a basic *hypergraph Turán* problem. Hypergraph Turán problems are typically significantly more difficult than their counterparts in graphs. Indeed, even the original hypergraph Turán conjecture from the 1940s that studies an appropriate analog of triangle free hypergraphs is still open. We direct the reader to the recent survey of Keevash [Kee11] for an overview of hypergraph Turán theory.

Applications of girth-density trade-offs. Like the graph Moore bound, girth-density trade-offs for hypergraphs have foundational connections to several research directions in theoretical computer science. One source of such applications is the observation that for a collection of linear equations in \mathbb{F}_2 on n variables, if we associate each equation to the set indicated by its coefficient vector, then the girth of the resulting hypergraph on $[n]$ is the same as the size of the smallest linearly dependent subset of equations. As a consequence, rate vs distance trade-offs for *low density parity check* (LDPC) codes are equivalent to the girth vs size trade-offs for k -uniform hypergraphs with hyperedges corresponding to the columns of the parity check matrix. As a result, there is an extensive line of work that studies the girth-density trade-offs for hypergraphs (see e.g. [BKHL99, BMS08, AF09]).

Naor and Verstrate [NV08] started a systematic study of hypergraph girth density trade-offs. They were explicitly motivated by mapping the rate-distance trade-offs for LDPC codes and computing product representations of square integers arises as a step in sub-exponential time algorithms for integer factoring. In particular, they showed that every k -uniform hypergraph on n vertices and $O(n^{k/2} \log n)$ hyperedges must have an even cover of size $O(\log n)$. Improving the bounds of [NV08] for $k = 3$, Feige [Fei08] proved that every 3-uniform hypergraph on n vertices and $O(n^{3/2}) \log \log n$ hyperedges has an even cover of length $O(\log n)$. Feige’s motivation was a connection, via the connection to linear equations modulo 2 discussed above, to refuting random 3SAT formulas and generalizations. In particular, by exploiting this improved bound, Feige derived a weak refutation algorithm for smoothed 3SAT formulas with $O(n^{1.5} \log \log n)$ constraints.

¹There are several well-studied combinatorial notions of [cycles](#) in contrast to the more linear algebraic notion of even covers.

Hypergraph Moore bound. Feige’s result leaves open the uncharted territory of hypergraph sizes between $m \sim n$ and $m \sim n^{k/2}$ — a polynomially large multiplicative interval when $k > 2$. The work of Feige, Kim and Ofek [FKO06] found an intriguing connection between the girth bounds in this interesting regime and the foundational average-case problem of refuting random 3SAT formulas [Fei02]. They observed that *random* hypergraphs with $m \gtrsim n^{1.4}$ hyperedges² must have an even cover of length $O(n^{0.2})$ and used a tour de force argument based on the second moment method to establish that at the same density, random hypergraphs should contain $n^{1.4}$ different almost disjoint even covers of size $n^{0.2}$. As a consequence, they obtained their celebrated result on the existence of polynomial size witnesses of unsatisfiability for random 3SAT formulas with $O(n^{1.4})$ constraints — a threshold that is $n^{0.1}$ factor smaller than the $m \gtrsim n^{1.5}$ bound for the best known efficient refutation algorithms. Motivated both by whether FKO witnesses could be efficiently constructed (and potentially refute a strong form of Feige’s Random 3SAT hypothesis [Fei02]) and investigating whether such certificates exist in semirandom and smoothed 3SAT formulas, Feige [Fei08] conjectured the following *hypergraph Moore bound*.

Conjecture 1.1 (Hypergraph Moore Bound (Feige’s conjecture), Conjecture 1.2 of [Fei08]). *For every $k \in \mathbb{N}$ and $1 \leq r \leq n$, every hypergraph with n vertices and $m \gtrsim n \left(\frac{n}{r}\right)^{\frac{k}{2}-1}$ hyperedges has an even cover of size $O(r \log n)$.*

In addition to a complete rate-distance profile for LDPC codes, Feige’s conjecture implies (see Section 9 in [GKM21] for an exposition) a significantly simpler and 2nd-moment-method-free proof of the existence of the FKO [FKO06] refutation witnesses below the spectral threshold for random 3SAT (and other CSPs) that also generalizes to semirandom and smoothed instances³.

Feige’s conjecture was recently settled by Guruswami, Kothari and Manohar [GKM21] up to an additional $\log^{4k+1} n$ multiplicative factor in the density m . Their proof goes via a new connection between the existence of small even covers in k -uniform hypergraphs and sub-exponential size spectral refutations of semirandom k -XOR formulas via a certain *Kikuchi* matrix.

While [GKM21] begins with an elegant and simple observation, their technical analysis especially for odd k (the “hard” case in all algorithms and certificates for refutation) is quite complicated and involves manipulating the Kikuchi matrix via “row bucketing” and “row pruning” in various steps and invoking the Schudy–Sviridenko concentration inequality [SS12] (that extends the breakthrough work of Kim and Vu [KV00]) for polynomials with combinatorial structure in the monomials. As a consequence, even for the simplest case of $k = 2$ (i.e., recovering the classical Moore bound), their proof incurs an additional $\log^3 n$ factor.

1.1 Our results

The main result of this work is a simple and short proof of the hypergraph Moore bound that is *almost* tight up to a single logarithmic factor.

Theorem 1.2. *For every $k \in \mathbb{N}$ and $1 \leq r \leq n$, every hypergraph on n vertices and $m \gtrsim n \log n \cdot \left(\frac{n}{r}\right)^{\frac{k}{2}-1}$ hyperedges has an even cover of size $O(r \log n)$.*

²Throughout this work, we will use the notation $f \gtrsim g$ to stand for “there exists a constant $C > 0$ such that $f \geq Cg$ ”.

³A smoothed Boolean CSP instance is obtained by starting from a worst-case instance and perturbing the literal patterns by independently flipping each with some small constant probability (with probability $1/2$ in the special case of the semirandom model). In particular, in contrast to random CSPs where the variables in every clause are generated uniformly at random, smoothed and semirandom CSP instances have a worst-case clause structure.

In [Section 2.1](#) and [Appendix A](#), as evidence of the power of our proof strategy, we obtain yet another proof of the classical Moore bound [[AHL02](#)] with the same leading constant. Independently of our work, David Munhá Correia and Benny Sudakov [[CS22](#)] informed us that they have found a simple, combinatorial argument for analyzing the Kikuchi matrix to prove a hypergraph Moore bound for even arity k that also loses only a single logarithmic factor.

Our techniques extend to give a simple and tighter proof of a sub-exponential time strong refutation algorithm for semirandom k -XOR formulas when the number of constraints is below the “spectral threshold” $n^{k/2}$, which is spelled out in [Section 4](#). Via the standard XOR trick (see, for example, [[AOW15](#)]), this recovers a tighter trade-off for refuting smoothed Boolean constraint satisfaction problems as in [[GKM21](#)]. Prior to our work, a bound tight up to $\log n$ factors was not known even for the (significantly) easier setting of *fully random* k -XOR refutation for odd k (the argument of [[WAM19](#)] obtains such a result for even k) where the best known bound due to [[RRS17](#)] loses a $\log^{2k} n$ factor.

Theorem 1.3 (Informal). *Fix $k \in \mathbb{N}$ and $r \leq n$, there is an $n^{O(r)}$ -time algorithm such that given a semirandom k -XOR instance ψ with n variables and $m \gtrsim n \log n \cdot (\frac{n}{r})^{\frac{k}{2}-1}$ constraints, it certifies that ψ is not $(1/2 + 0.01)$ -satisfiable.*

We believe that the last remaining logarithmic factor in the theorems above is also unnecessary. However, removing it seems related to certain technical difficulties that arise in beating the logarithmic factor incurred in spectral norm bounds for the matrix Rademacher series [[Tro15](#)]. In particular, the tightest known proof for the closely related problem of refuting fully random k -XOR formulas below the spectral threshold also loses a $\log n$ factor for even k [[WAM19](#)] and a $\log^{2k} n$ -factor for odd k [[RRS17](#)]. For some easier settings such as refutation in the polynomial time regime [[dT22](#)] and understanding the SDP value of random NAE-3SAT and generalizations [[FM17](#), [DMO⁺19](#), [MOP20](#)], recent works manage to circumvent this difficulty by application of powerful tools such as the Ihara–Bass formula and the largest eigenvalue of non-backtracking walk matrices. Our proof suggests a natural and more elementary route to removing this final logarithmic factor but requires resolving the count of certain “walks” that arise in our analysis.

Key ideas. The abstract strategy employed by [[GKM21](#)] is to construct a so called *Kikuchi* matrix $A_{\mathcal{H}}$ (first introduced in the work of [[WAM19](#)] on Gaussian tensor PCA) associated with our hypergraph \mathcal{H} where:

- (i) $\mathbf{1}^\top A_{\mathcal{H}} \mathbf{1}$ is a surrogate for the number of hyperedges in \mathcal{H} .
- (ii) The lack of short even covers in \mathcal{H} can be turned into a certificate that $\mathbf{1}^\top A_{\mathcal{H}} \mathbf{1}$ is small, which then translates to a bound on $|\mathcal{H}|$.

The certificate used by [[GKM21](#)] is $\|A_{\mathcal{H}}\|_{\infty \rightarrow 1}$, which they control by bucketing the rows by weight, bounding the spectral norm of each submatrix, and stitching these norms together.

Our key insight is in the style of certificate we provide — we give a matrix Q such that $Q \succeq A_{\mathcal{H}}$. Such a certificate implies a bound of $\text{tr}(Q)$ on our surrogate for $|\mathcal{H}|$. The inequality $Q \succeq A_{\mathcal{H}}$ is equivalent to proving $\|Q^{-1/2} A_{\mathcal{H}} Q^{-1/2}\|_2 \leq 1$, which can be done via the trace moment method with relative ease. Our reweighting strategy is akin to constructing a “diagonal weighted” *dual solution* for certifying upper bounds on the value of the basic SDP relaxation for quadratic

optimization problems on the hypercube such as Max-Cut and the Grothendieck problems. Our *reweighting* strategy simplifies the analysis, removes the need for the “row bucketing” step in [GKM21], and lets us obtain a sharper result.

Our proof for odd k requires combining our reweighted Kikuchi matrix with a new “edge deletion” operation that controls the “heavy rows” in the Kikuchi matrix. At a high level, our strategy involves deleting an appropriately chosen set of entries of the Kikuchi matrix in comparison to the “row pruning” strategy of [GKM21] which involves deleting entire rows (vertices in the Kikuchi graph). This seemingly technical change leads to a great deal of simplification and in particular allows replacing the use of the Schudy–Sviridenko inequality [SS12] and the carefully introduced logarithmic factors in the hypergraph regularity decomposition in [GKM21].

Organization. The rest of this paper is organized as follows. In Section 2, we give a complete (and sharper) proof of the hypergraph Moore bound for the even arity case, starting with a proof of a weak version (that loses additional constant factors) of the classical Moore bound using our ideas. We will include detailed commentary for the sake of exposition and a short overview of the additional ideas (including our new edge deletion trick) to handle the odd arity case. In Section 3, we will give a proof of the hypergraph Moore bound for odd arity. Finally, in Section 4, we will extend our techniques to obtain strong refutation algorithms for semirandom and smoothed Boolean CSPs.

2 Warm-up: hypergraph Moore bound in the even arity case

In this section, we will give a proof of the Moore bound for hypergraphs of even arity with the goal of providing an exposition of our main ideas. As an illustration of the power of our reweighting idea, in Section 2.1 we will give a simple proof of the classical Moore bound [AHL02] that is tight up to an absolute constant factor (as opposed to the $\log^3 n$ loss incurred by the strategy of [GKM21]).⁴ In Section 2.2, we will generalize the reweighting idea to prove hypergraph Moore bound for all even arities. Finally in Section 2.3, we will discuss the key new idea of *edge deletions* that is crucial for our simpler and tighter proof for the case of odd k .

2.1 Weak Moore bound for graphs

In this section, we prove a weak Moore bound for graphs to illustrate our reweighting strategy in a simple setting. The resulting bound is weak in the sense that it incurs a constant factor loss when compared to [AHL02]. In Appendix A, we implement this strategy (in a way that is less generalizable to hypergraphs) to recover the tight $2 \log_{d-1} n$ bound.

We note that [GKM21] also proved a weaker Moore bound (Proposition 2.3 of [GKM21]) to illustrate their “row bucketing” strategy that partitions the vertices into $O(\log n)$ buckets, each of which has vertices with degrees within a multiplicative constant factor of each other. This strategy splits the adjacency matrix A into $O(\log^2 n)$ pieces and ends up requiring an average degree $d \gtrsim \log^3 n$ in order to contain a cycle of length $O(\log n)$.

⁴In Appendix A, we present a proof that uses one additional tool to recover the classical Moore bound for irregular graphs with the same leading constant.

This simple exercise will show how our reweighting handles different degrees automatically, avoiding the lossy row bucketing step completely.

Proposition 2.1 (Weak Moore bound for irregular graphs). *Every graph with n vertices and average degree $d > 16$ has a cycle of length at most $2\lceil \log_{(d/16)} n \rceil$.*

The core of the proof of [Proposition 2.1](#) is the following spectral norm bound on the reweighted adjacency matrix.

Claim 2.2. *Let G be a graph with n vertices and average degree $d > 1$ that has no cycle of length $\leq \ell$ for some even $\ell \in \mathbb{N}$. Let A be the $\{0, 1\}$ adjacency matrix of G , and let $\Gamma = D + d\mathbb{I}$ be the diagonal matrix such that $D_{uu} = d_u$ where d_u is the degree of vertex u . Then, $\|\Gamma^{-1/2} A \Gamma^{-1/2}\|_2 < \frac{2n^{1/\ell}}{\sqrt{d}}$.*

We now complete the proof of [Proposition 2.1](#).

Proof of [Proposition 2.1](#) by [Claim 2.2](#). Suppose G has no cycle of length $\leq \ell$, then [Claim 2.2](#) implies that $A \prec \frac{2n^{1/\ell}}{\sqrt{d}} \Gamma$. Then, the quadratic form $\mathbf{1}^\top A \mathbf{1} < \frac{2n^{1/\ell}}{\sqrt{d}} \text{tr}(\Gamma)$ since $\mathbf{1}^\top \Gamma \mathbf{1} = \text{tr}(\Gamma)$. By definition, $\mathbf{1}^\top A \mathbf{1} = nd$ and $\text{tr}(\Gamma) = \sum_{u=1}^n (d_u + d) = 2nd$. Thus, $n^{1/\ell} > \sqrt{d}/4$, and taking logs, we get

$$\frac{1}{\ell} \log n > \frac{1}{2} \log(d/16) \Rightarrow \frac{\ell}{2} < \log_{d/16} n.$$

ℓ is even, so we have $\ell < 2\lceil \log_{d/16} n \rceil$. Thus, by the contrapositive, G must contain a cycle of length $2\lceil \log_{d/16} n \rceil$. This completes the proof. \square

We now prove [Claim 2.2](#) using the well-known trace moment method, which reduces to counting weighted closed walks in the graph. In the analysis, we will see exactly how the choice of the reweighting matrix Γ accounts for different vertex degrees.

Proof of [Claim 2.2](#). Let $\tilde{A} = \Gamma^{-1/2} A \Gamma^{-1/2}$. For even $\ell \in \mathbb{N}$, the trace moment method states that $\|\tilde{A}\|_2^\ell \leq \text{tr}(\tilde{A}^\ell) = \text{tr}((\Gamma^{-1} A)^\ell)$, which is a summation of all (weighted) closed walks of length ℓ in G . Since there is no cycle of length $\leq \ell$, the only closed walks are the ones that *backtrack* to the original vertex, meaning that there can be at most $\ell/2$ “new” edges and at least $\ell/2$ “old” edges in the walk. We encode each closed walk $u_1 \rightarrow u_2 \rightarrow \dots \rightarrow u_\ell \rightarrow u_1$ as follows,

- Choose a starting vertex $u_1 \in [n]$.
- One bit $b_i \in \{0, 1\}$ at each step i to encode whether this step uses a new edge or an old one.
 - If $b_i = 0$ (new edge), select one of u_i 's neighbors as u_{i+1} .
 - If $b_i = 1$ (old edge), we must backtrack to the previous vertex u_{i-1} .

For $b \in \{0, 1\}$ and $u \in [n]$, let $N_b(u) \subseteq [n]$ be the possible next steps in the walk from u . Then, simply expanding $\text{tr}((\Gamma^{-1} A)^\ell)$, we get

$$\text{tr}((\Gamma^{-1} A)^\ell) = \sum_{b \in \{0, 1\}^\ell} \sum_{u_1 \in [n]} \sum_{u_2 \in N_{b_1}(u_1)} \Gamma_{u_1 u_1}^{-1} \sum_{u_3 \in N_{b_2}(u_2)} \Gamma_{u_2 u_2}^{-1} \dots \sum_{u_{\ell+1} \in N_{b_\ell}(u_\ell)} \Gamma_{u_\ell u_\ell}^{-1} \cdot \mathbf{1}(u_{\ell+1} = u_1).$$

As we can see, each step $u_i \rightarrow u_{i+1}$ gets a factor $\Gamma_{u_i u_i}^{-1} = \frac{1}{d_{u_i} + d}$. We can now bound the above by observing that if $b_i = 0$ (new edge), then $|N_0(u_i)| \leq d_{u_i}$ and

$$\sum_{u_{i+1} \in N_0(u_i)} \Gamma_{u_i u_i}^{-1} \leq \frac{d_{u_i}}{d_{u_i} + d} < 1,$$

and if $b_i = 0$ (old edge), then $|N_1(u_i)| = 1$ (the previous step) and

$$\sum_{u_{i+1} \in N_1(u_i)} \Gamma_{u_i u_i}^{-1} \leq \frac{1}{d_{u_i} + d} < \frac{1}{d}.$$

Finally, considering $b \in \{0, 1\}^\ell$, $u_1 \in [n]$, and there are at least $\ell/2$ old edges, we have

$$\text{tr}((\Gamma^{-1}A)^\ell) < 2^\ell n \left(\frac{1}{d}\right)^{\ell/2},$$

and taking the ℓ -th root completes the proof. \square

2.2 The case of even arity hypergraphs

In this section, we prove the existence of small even covers in even arity hypergraphs.

Theorem 2.3 (Theorem 1.2, even k). *For even $k \in \mathbb{N}$ and any $r \in \mathbb{N}$ with $k \leq r \leq n/8$, any k -uniform hypergraph \mathcal{H} with n vertices and $m \geq 128n \log n \cdot \left(\frac{n}{r}\right)^{k/2-1}$ hyperedges has an even cover of size at most $\lceil r \log_2 n \rceil + 1$.*

The proof is simple and almost identical to the proof of the weak Moore bound (Proposition 2.1) but with A being the adjacency matrix of the *Kikuchi graph* which we define below.

Definition 2.4 (Kikuchi graph). Let \mathcal{H} be a k -uniform hypergraph on vertex set $[n]$ for even k . For an integer parameter r , define the *Kikuchi graph* K_r associated to \mathcal{H} is a graph on vertex set $\binom{[n]}{r}$ such that a pair of vertices $S, T \in \binom{[n]}{r}$ have an edge between them if the symmetric difference $S \oplus T \in \mathcal{H}$. For such an edge, we write $S \overset{C}{\leftrightarrow} T$ and think of the edge as “colored” by $C \in \mathcal{H}$ where $C = S \oplus T$. We call the adjacency matrix A of K_r the *Kikuchi matrix*.

The key insight of [GKM21] (and also our starting point) is relating even covers in \mathcal{H} to cycles in the associated Kikuchi graph. For sets $R_1, R_2, \dots, R_\ell \subseteq [n]$ let $\oplus_{i \leq \ell} R_i$ denote the set of elements of $[n]$ that appear in an odd number of R_i s (i.e., the sum modulo 2 of the indicator vectors of R_i s).

Observation 2.5 (Closed walks in the Kikuchi graph). Let \mathcal{H} be a k -uniform hypergraph on $[n]$ for even k and let $S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_\ell \rightarrow S_1$ be a closed walk on vertices in K_r such that for every $i \leq \ell$, $S_i \overset{C_i}{\leftrightarrow} S_{i+1}$ for $C_1, C_2, \dots, C_\ell \in \mathcal{H}$ (denoting $S_{\ell+1} = S_1$). Then, $\oplus_{i \leq \ell} C_i = 0$. Further, if \mathcal{H} has no even cover of length ℓ , then every hyperedge in \mathcal{H} appears an even number of times in the multiset $\{C_1, C_2, \dots, C_\ell\}$. We will call such walks in K_r *trivial*.

Proof. Note that $S_i \oplus S_{i+1} = C_i$ for every $i \leq \ell$. If we add both sides of all ℓ such equalities then each S_i occurs in exactly two of the equations so the LHS must be 0. Thus, $\oplus_{i \leq \ell} C_i = 0$.

Next, we repeatedly remove hyperedges that occur an even number of times in the multiset $\{C_1, C_2, \dots, C_\ell\}$ to obtain a collection of $\ell' \leq \ell$ distinct hyperedges of \mathcal{H} . The sum (modulo 2) of the remaining hyperedge should still be 0 as we removed hyperedges in pairs. The resulting ℓ' must be 0 as otherwise the remaining hyperedges form an even cover of length $\ell' \leq \ell$. \square

Consider a hypergraph \mathcal{H} with n vertices and m hyperedges, and its associated Kikuchi graph (V, E) with parameter r . Each $C \in \mathcal{H}$ introduces $\frac{1}{2} \binom{k}{k/2} \binom{n-k}{r-k/2}$ edges in the Kikuchi graph (select $k/2$ vertices from C and select $r - k/2$ vertices from $[n] \setminus C$ to complete S), thus the total edges $|E| = \frac{1}{2} \binom{k}{k/2} \binom{n-k}{r-k/2} \cdot m$. Let d_S be the degree of $S \in V$, and let d denote the average degree, then simple calculations show that

$$d = \frac{\binom{k}{k/2} \binom{n-k}{r-k/2} m}{\binom{n}{r}} \geq \left(\frac{r}{n}\right)^{k/2} m \cdot \binom{k}{k/2} \left(1 - \frac{2r}{n}\right)^{k/2} \left(1 - \frac{k}{2r}\right)^{k/2} \geq \frac{1}{2} \left(\frac{r}{n}\right)^{k/2} m \quad (1)$$

when $k \leq r \leq n/8$.

We will follow the reweighting strategy with $\Gamma = D + d\mathbb{I}$ to bound the spectral norm of the reweighted Kikuchi matrix. The following lemma is analogous to [Claim 2.2](#).

Lemma 2.6. *Let $k, r, n \in \mathbb{N}$ such that $k \leq r \leq n$, and let $\ell \in \mathbb{N}$ be even. Let A be the Kikuchi matrix with parameter r of a k -uniform hypergraph \mathcal{H} on n vertices, and let $\Gamma = D + d\mathbb{I}$ where D is the degree matrix and d is the average degree of the Kikuchi graph. Suppose there is no even cover of size at most ℓ in \mathcal{H} , then*

$$\left\| \Gamma^{-1/2} A \Gamma^{-1/2} \right\|_2 < 2n^{r/\ell} \sqrt{\frac{\ell}{d}}.$$

We can immediately complete the proof of [Theorem 2.3](#).

Proof of [Theorem 2.3](#) by [Lemma 2.6](#). Suppose that there is no even cover of size $\leq \ell := \lceil r \log_2 n \rceil$ (assume this is even, otherwise add 1). Then, $n^{r/\ell} \leq 2$ and [Lemma 2.6](#) states that the Kikuchi graph (V, E) satisfies $A \prec 4\sqrt{\ell/d} \cdot \Gamma$ where $\Gamma = D + d\mathbb{I}$. Then,

$$\mathbf{1}^\top A \mathbf{1} < 4\sqrt{\frac{\ell}{d}} \cdot \text{tr}(\Gamma) = 4\sqrt{\frac{\ell}{d}} \cdot \sum_{S \in V} (d_S + d) = 8\sqrt{\frac{\ell}{d}} \cdot |V|d.$$

On the other hand, $\mathbf{1}^\top A \mathbf{1} = 2|E| = |V|d$. Thus, we have $d < 64\ell$. By [Eq. \(1\)](#) we have $d \geq \frac{1}{2} \left(\frac{r}{n}\right)^{k/2} m$ when $k \leq r \leq n/8$. Thus, if there is no even cover of size $\leq \ell$, then $m < 128n \log n \cdot \left(\frac{n}{r}\right)^{k/2-1}$, completing the proof. \square

Now, we prove [Lemma 2.6](#) by counting weighted closed walks in the Kikuchi graph, essentially the same way we prove [Claim 2.2](#).

Proof of [Lemma 2.6](#). Let $\tilde{A} = \Gamma^{-1/2} A \Gamma^{-1/2}$. We use the trace power method:

$$\|\tilde{A}\|_2^\ell \leq \text{tr}(\tilde{A}^\ell) = \text{tr}((\Gamma^{-1} A)^\ell).$$

We upper bound $\text{tr}((\Gamma^{-1} A)^\ell)$ by counting (weighted) closed walks of length ℓ in the Kikuchi graph. Note that each edge (S, T) of the Kikuchi graph corresponds to a hyperedge $S \oplus T \in \mathcal{H}$. Since there is no even covers of size at most ℓ , any closed walk must contain an even number of each hyperedge in \mathcal{H} .

We can encode a closed walk $S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_\ell \rightarrow S_1$ as follows:

- Choose a starting vertex $S_1 \in V$.
- One bit $b_i \in \{0, 1\}$ at each step i to encode whether this step uses a new hyperedge or an old one.

- If $b_i = 0$ (new hyperedge), select one of S_i 's neighbors as S_{i+1} .
- If $b_i = 1$ (old hyperedge), select an old hyperedge C from the previous steps, and set $S_{i+1} = S_i \oplus C$.

Note that there are at most $\ell/2$ new hyperedges and at least $\ell/2$ old hyperedges since each hyperedge must occur an even number of times. For $b \in \{0, 1\}$ and $S \in V$, let $N_b(S) \subseteq V$ be the possible next steps in the walk from S (according to b). Each step $S_i \rightarrow S_{i+1}$ gets a factor $(\Gamma^{-1}A)_{S_i, S_{i+1}} = \Gamma_{S_i, S_i}^{-1} = \frac{1}{d_{S_i} + d}$. Thus,

$$\text{tr}((\Gamma^{-1}A)^\ell) = \sum_{b \in \{0, 1\}^\ell} \sum_{S_1 \in V} \sum_{S_2 \in N_{b_1}(S_1)} \frac{1}{d_{S_1} + d} \sum_{S_3 \in N_{b_2}(S_2)} \frac{1}{d_{S_2} + d} \cdots \sum_{S_{\ell+1} \in N_{b_\ell}(S_\ell)} \frac{\mathbf{1}(S_{\ell+1} = S_1)}{d_{S_\ell} + d}.$$

We can upper bound the above as follows. If $b = 0$, then $|N_0(S_i)| \leq d_{S_i}$ and $\sum_{S_{i+1} \in N_0(S_i)} \Gamma_{S_i, S_i}^{-1} \leq \frac{d_{S_i}}{d_{S_i} + d} < 1$. If $b = 1$, then $|N_1(S_i)| \leq \ell$ as there are only ℓ options to choose one of the previous steps, and $\sum_{S_{i+1} \in N_1(S_i)} \Gamma_{S_i, S_i}^{-1} \leq \frac{\ell}{d_{S_i} + d} < \frac{\ell}{d}$. Furthermore, we can assume that $\ell \leq d$, otherwise we can simply treat all steps as new hyperedges.

Finally, $b \in \{0, 1\}^\ell$, there are $|V| = \binom{n}{r}$ choices for the starting vertex S_1 , and there are at least $\ell/2$ old hyperedges. Thus, we have

$$\text{tr}((\Gamma^{-1}A)^\ell) < 2^\ell \binom{n}{r} \left(\frac{\ell}{d}\right)^{\ell/2} \leq 2^\ell n^r \left(\frac{\ell}{d}\right)^{\ell/2}.$$

Taking the ℓ -th root completes the proof. \square

Summary. As the above short and simple arguments illustrate, reweighted Kikuchi matrix appears to be a clean and simple way to handle irregularities in the degree of graphs (Kikuchi or otherwise) in spectral double counting. For $k > 2$, we suspect that the extra $\log n$ factor incurred in our analysis can likely be removed by a better counting of the weighted closed walks.

2.3 Overview of the odd arity case

As in many previous works on refuting constraint satisfaction problems, the odd arity case requires significantly more work. Indeed, even the definition of the Kikuchi graph (Definition 2.4) only makes sense when k is even. We present the proof of the odd arity case in Section 3, and here we outline some of our key ideas.

Bipartite hypergraph. The main insight is to transform the hypergraph \mathcal{H} to a “bipartite” hypergraph (this abstraction is closely related to the Cauchy-Schwarz trick in the context of odd-arity CSP refutation). First, we partition the hyperedges of \mathcal{H} into $\mathcal{H}_1, \dots, \mathcal{H}_p$ such that for each \mathcal{H}_i , all hyperedges in \mathcal{H}_i contains a “center” vertex $u_i \in [n]$. We denote $\tilde{\mathcal{H}}_i$ to be $\{C \setminus \{u_i\} : C \in \mathcal{H}_i\}$, i.e. removing the center vertex, and denote $\tilde{C} := C \setminus \{u_i\}$. Then, we construct a $2(k-1)$ -uniform hypergraph as follows: for each $i \in [p]$ and each distinct pair $C, C' \in \mathcal{H}_i$, we add a hyperedge $\tilde{C} \oplus \tilde{C}'$ (let's assume \tilde{C}, \tilde{C}' are disjoint for now).

Let's make some quick calculations. Suppose that \mathcal{H} has m hyperedges, and suppose that there are roughly $p \approx n$ partitions and each partition size is $\approx m/n$. Then, the new $2(k-1)$ -uniform hypergraphs will contain roughly $n \cdot (m/n)^2 = \frac{m^2}{n}$ hyperedges. Now, since $2(k-1)$ is even, we can apply our bound for the even case: there is an even cover of size $r \log n$ when $\frac{m^2}{n} \geq \tilde{O}(n) \left(\frac{n}{r}\right)^{(k-1)-1}$, meaning $m \geq \tilde{O}(n) \left(\frac{n}{r}\right)^{\frac{k}{2}-1}$, the correct bound!

The issue is that this new hypergraph has small even covers for trivial reasons: any 3 pairs (C_1, C_2) , (C_2, C_3) and (C_3, C_1) from \mathcal{H}_i form an even cover of size 3. Nevertheless, we can proceed to analyze the Kikuchi matrix of the new hypergraph (Definition 3.5), assuming that there is no small even cover in the *original* hypergraph. Note that now an edge (S, T) is associated with 2 hyperedges C, C' from the same \mathcal{H}_i , which we denote as $S \xleftrightarrow{C, C'} T$. Assuming that there is no even cover of size $\leq 2\ell$, we bound the number of "trivial" closed walks where each hyperedge is used an even number of times.

Encoding a closed walk. The standard technique of bounding counts of closed walks in the trace moment method is to give a small *encoding* of a walk. In our case, in a length- ℓ closed walk of the Kikuchi graph, each step is associated with two hyperedges, and we have two types of steps:

1. a step using 2 *new* hyperedges, and
2. a step using at least 1 *old* hyperedge.

The first type is bounded exactly the same way as the even arity case by our weight matrix Γ , the trouble is the second type: while we can easily encode one edge in the step, we need too many bits to encode the other edge.

Deleting bad edges of the Kikuchi graph. The main insight is that in the end, we only care about bounding $\mathbf{1}^\top A \mathbf{1}$. Again, let d be the average degree of the Kikuchi graph, $A \in \mathbb{R}^{N \times N}$ be the Kikuchi matrix, and $\Gamma = D + d\mathbb{I}$ be our diagonal weight matrix. If we delete (say) half of the edges of the Kikuchi graph such that we have $\|\Gamma^{-1/2} A' \Gamma^{-1/2}\|_2 \leq \lambda(d)$, where A' is the modified Kikuchi matrix, for some "good enough" $\lambda(d)$, then we will have $\frac{Nd}{2} \leq \mathbf{1}^\top A' \mathbf{1} \leq \lambda \operatorname{tr}(\Gamma) = \lambda \cdot 2Nd$, essentially only losing a constant factor in the density.

We define an appropriate edge deletion process, prove that the fraction of edges removed is small (Claim 3.9), and show that the resulting subgraph has combinatorial properties that let us encode steps of the second type efficiently (Lemma 3.8).

Improving the row pruning step of [GKM21]. The analysis of [GKM21] also requires reducing the Kikuchi graph to obtain certain combinatorial properties. However, instead of deleting "bad" edges, they delete "bad" vertices, which they defined as vertices that are bad for *some* i in the bipartite hypergraph (they call this row pruning as each row of the Kikuchi matrix corresponds to a vertex). Crucially, doing so requires a union bound over i , hence they need a strong bound on fraction of bad vertices for each i . Furthermore, they proved their bound using tail inequalities for low-degree polynomials by Schudy and Sviridenko [SS12], which is a powerful black-box concentration inequality but loses log factors and requires involved analysis. All this combined with the row bucketing step introduces several log factors.

Hyperedges with large intersections. It turns out that the fraction of bad edges highly depends on *large intersections* of hyperedges in \mathcal{H} . To bound the fraction of edges deleted, we require our hypergraph to be somewhat “regular” – that is, no small subset appears in more than an appropriately chosen threshold of hyperedges in \mathcal{H} . To this end, we invoke the hypergraph regularity decomposition of [GKM21] (with more transparently chosen thresholds that do not involve carefully chosen logarithmic factors) to decompose the hypergraph into at most k subhypergraphs such that each piece satisfies the required regularity conditions (see [Algorithm 1](#) and [Observation 3.3 & 3.2](#)). Then, there must be one subhypergraph $\mathcal{H}^{(i)}$ with at least m/k hyperedges, and we will show that there exists an even cover within $\mathcal{H}^{(i)}$.

3 Hypergraph Moore bound for odd arity hypergraphs

In this section we prove the hypergraph Moore bound for k -uniform hypergraphs when k is odd.

Theorem 3.1 ([Theorem 1.2](#), odd k). *There is a universal constant B such that for any odd $k \in \mathbb{N}$, and any $r \in \mathbb{N}$ satisfying $2k \leq r \leq \frac{n}{Bk}$, any k -uniform hypergraph \mathcal{H} with n vertices and $m \geq B^k n \log n \cdot \left(\frac{n}{r}\right)^{k/2-1}$ hyperedges has an even cover of size at most $r \log_2 n$.*

Our proof strategy broadly involves the following steps.

- **Hypergraph decomposition.** We partition \mathcal{H} into subhypergraphs $\mathcal{H}^{(0)}, \mathcal{H}^{(1)}, \dots, \mathcal{H}^{(k-1)}$ with the property that every size- $(i+1)$ set in $\mathcal{H}^{(i)}$ is contained in only a small number of clauses, and every clause in $\mathcal{H}^{(i)}$ intersects many other clauses at a size- i set. One of the $\mathcal{H}^{(i)}$ must contain at least m/k clauses, and we find an even cover in that $\mathcal{H}^{(i)}$.
- **Large i .** When $i \geq \frac{k+1}{2}$, we give a direct reduction to the hypergraph Moore bound for even arity hypergraphs and apply [Theorem 2.3](#).
- **Kikuchi graph.** To handle the remaining values of i , we show the existence of an even cover by proving the contrapositive — a hypergraph with no small even covers has a bounded number of hyperedges. To achieve this, we appropriately define the Kikuchi graph for odd arity hypergraphs, and show that the adjacency matrix \hat{A} of some suitably chosen subgraph (via the “edge deletion process” described below) satisfies $\hat{A} \preceq Q$ for some diagonal matrix Q . Then the resulting inequality $\mathbf{1}^\top \hat{A} \mathbf{1} \leq \text{tr}(Q)$ can be rearranged to bound the number of hyperedges.
- **Trace method.** The way we prove $\hat{A} \preceq Q$ is by using the trace moment method to show $\left\| Q^{-1/2} \hat{A} Q^{-1/2} \right\|_2 \leq 1$. Bounding a high trace power of $Q^{-1/2} \hat{A} Q^{-1/2}$ corresponds to bounding the total weight of closed walks that use every hyperedge an even number of times in the Kikuchi graph.
- **Edge deletion process.** We delete a small fraction of the edges in K_r with the guarantee that in the resulting subgraph any clause participates in only a small number of incident edges to every vertex.

Hypergraph decomposition. We describe our algorithm to partition our hypergraph.

Algorithm 1. We partition \mathcal{H} into hypergraphs $\mathcal{H}^{(0)}, \dots, \mathcal{H}^{(k-1)}$ via the following algorithm.

1. Set $t = k - 1$ and $\mathcal{H}_{\text{current}} := \mathcal{H}$.
 2. Set counter $s = 1$. While there is $U \subseteq [n]$ such that $|U| = t$ and $|\{C \in \mathcal{H}_{\text{current}} : U \subseteq C\}| \geq \max\left\{2, \left(\frac{n}{r}\right)^{\frac{k}{2}-t}\right\}$:
 - (a) Choose U satisfying the condition and let $\mathcal{H}_s^{(t)}$ be a subset of $\{C \in \mathcal{H}_{\text{current}} : U \subseteq C\}$ of size $\max\left\{2, \left(\frac{n}{r}\right)^{\frac{k}{2}-t}\right\}$.
 - (b) Add all clauses in $\mathcal{H}_s^{(t)}$ to $\mathcal{H}^{(t)}$.
 - (c) Delete all clauses in $\mathcal{H}_s^{(t)}$ to $\mathcal{H}_{\text{current}}$.
 - (d) Increment s by 1.
 3. Decrement t by 1. If $t > 0$, go back to step 2; otherwise take the remaining clauses in $\mathcal{H}_{\text{current}}$ and add them to $\mathcal{H}^{(0)}$.
-

First, observe that the largest subhypergraph $\mathcal{H}^{(i)}$ in the partition produced by our algorithm must have at least $\frac{m}{k}$ hyperedges. Next, observe that $i \neq 0$ because if $|\mathcal{H}^{(0)}| \geq m/k$, then there must be a $j \in [n]$ such that $|\{C \in \mathcal{H}^{(0)} : j \in C\}| \geq \frac{m}{nk} \gg \left(\frac{n}{r}\right)^{k/2-1}$, which would have been added to $\mathcal{H}^{(1)}$. Our goal in the rest of the proof is to find a small even cover in $\mathcal{H}^{(i)}$. The following observations articulate the properties of $\mathcal{H}^{(i)}$ we need that are guaranteed by the algorithm.

Observation 3.2. $\mathcal{H}^{(i)}$ can be partitioned into $\mathcal{H}_1^{(i)}, \dots, \mathcal{H}_p^{(i)}$ where for each $j \in [p]$, there is a set U_j of size i such that every $C \in \mathcal{H}_j^{(i)}$ contains U_j , and $|\mathcal{H}_j^{(i)}| \geq \left(\frac{n}{r}\right)^{\frac{k}{2}-i}$ and $p \leq m \cdot \left(\frac{r}{n}\right)^{\frac{k}{2}-i}$.

Observation 3.3. For $s \geq 1$ and any $U \subseteq [n]$ such that $|U| = i + s$, the number of hyperedges in $\mathcal{H}^{(i)}$ containing U is at most $\max\left\{1, \left(\frac{n}{r}\right)^{\frac{k}{2}-s-i}\right\}$, otherwise they would have been added to $\mathcal{H}^{(i+s)}$.

Reduction to even arity case when $i \geq \frac{k+1}{2}$. In this case, by [Observation 3.3](#), each pair $C \neq C'$ in any $\mathcal{H}_j^{(i)}$ must satisfy $C \cap C' = U_j$. The following makes the reduction from finding even covers in $\mathcal{H}^{(i)}$ if $i \geq \frac{k+1}{2}$ to the even arity case concrete.

Lemma 3.4. Let \mathcal{H} be a k -uniform hypergraph on n vertices with no even cover of size $r \log_2 n$. Fix $1 \leq i \leq k - 1$. Suppose $\mathcal{H}_1, \dots, \mathcal{H}_p$ are disjoint subsets of \mathcal{H} such that for each $j \in [p]$, $|\mathcal{H}_j| \geq 2$ and all pairs of hyperedges $C \neq C' \in \mathcal{H}_j$ satisfy $C \cap C' = U_j$ for some $U_j \subseteq [n]$ of size i . Then,

$$\sum_{j=1}^p |\mathcal{H}_j| \leq O(n \log n) \left(\frac{2n}{r}\right)^{k-i-1}.$$

In particular, when $i \geq \frac{k+1}{2}$ the above is at most $O(n \log n) \cdot \left(\frac{n}{r}\right)^{k/2-1}$.

Proof. Given such disjoint subsets $\mathcal{H}_1, \dots, \mathcal{H}_p$, we can construct a $2(k-i)$ -uniform hypergraph $\widehat{\mathcal{H}}$ by the following: for each $j \in [p]$, arbitrarily order the edges: $\mathcal{H}_j = (C_1, \dots, C_{|\mathcal{H}_j|})$. Then, add the hyperedge $C_s \oplus C_{s+1}$ to $\widehat{\mathcal{H}}$ for $s = 1, \dots, |\mathcal{H}_j| - 1$. By assumption $|C_s \cap C_{s+1}| = |U_j| = i$, thus $|C_s \oplus C_{s+1}| = 2(k-i)$. The resulting $\widehat{\mathcal{H}}$ has

$$|\widehat{\mathcal{H}}| = \sum_{j=1}^p |\mathcal{H}_j| - 1 \geq \frac{1}{2} \sum_{j=1}^p |\mathcal{H}_j|$$

hyperedges, since $|\mathcal{H}_j| \geq 2$ for all $j \in [p]$.

We claim that $\widehat{\mathcal{H}}$ cannot have an even cover of size at most $\frac{t}{2} \log_2 n$. First, if $\widehat{\mathcal{H}}$ has repeated hyperedges, then there must exist $j \neq j' \in [p]$ and $C_1, C_2 \in \mathcal{H}_j, C'_1, C'_2 \in \mathcal{H}_{j'}$ such that $C_1 \oplus C_2 = C'_1 \oplus C'_2$, but then $\{C_1, C_2, C'_1, C'_2\}$ would be an even cover of size 4 in \mathcal{H} . Now, suppose $\widehat{\mathcal{H}}$ has no repeated edges but has an even cover of size ℓ . Then, for any \widehat{C} in the even cover, we can uniquely identify $j \in [p]$ and $s \leq |\mathcal{H}_j| - 1$ such that $C_s, C_{s+1} \in \mathcal{H}_j$ and $\widehat{C} = C_s \oplus C_{s+1}$. Furthermore, by construction there must be at least two $C_s, C_{s'} \in \mathcal{H}_j$ that each occurs only once. Therefore, these edges must form an even cover of size at most 2ℓ in \mathcal{H} .

Since $2(k-i)$ is even and $\widehat{\mathcal{H}}$ has no even cover of size $\frac{t}{2} \log_2 n$, we can apply [Theorem 2.3](#) to show that

$$|\widehat{\mathcal{H}}| \leq O(n \log n) \left(\frac{2n}{r} \right)^{k-i-1}.$$

This completes the proof. \square

Henceforth, we assume $i \leq \frac{k-1}{2}$, which is the case we need an appropriate Kikuchi graph for odd arity hypergraphs.

Kikuchi matrix for odd arity hypergraphs. The following is the same Kikuchi graph defined in [\[GKM21, Definition 6.2\]](#).

Definition 3.5 (Colored Kikuchi graphs and subgraphs). Fix $r \in \mathbb{N}$ and $t \in \{1, \dots, k-1\}$ such that $2k \leq r \leq n$. Let $\mathcal{H}_1, \dots, \mathcal{H}_p$ be p disjoint sets of hyperedges such that for each $i \in [p]$, all hyperedges in \mathcal{H}_i have a common subset $U_i \subset [n]$ where $|U_i| = t$. For each $C \in \mathcal{H}_i$, denote $\widetilde{C} := C \setminus U_i$, and denote $\widetilde{\mathcal{H}}_i := \{\widetilde{C} : C \in \mathcal{H}_i\}$ which can be viewed as a $(k-t)$ -uniform hypergraph. We define the *colored* Kikuchi graph K_r as follows.

The vertex set $V(K_r)$ consists of subsets of $[n] \times [2]$ of size r , where $S \in V$ is viewed as $(S^{(1)}, S^{(2)})$ where $S^{(1)}, S^{(2)} \subseteq [n]$ are colored *green* and *blue* respectively. For each $i \in [p]$ and each $C \neq C' \in \mathcal{H}_i$, let $\widetilde{C}^{(1)}$ be \widetilde{C} colored green and $\widetilde{C}'^{(2)}$ be \widetilde{C}' colored blue, and we add an edge between $S, T \in V$, denoted $S \xleftrightarrow{C, C'} T$, if $S \oplus T = \widetilde{C}^{(1)} \oplus \widetilde{C}'^{(2)}$ and if one of the following holds,

- $|\widetilde{C} \cap S^{(1)}| = |\widetilde{C}' \cap T^{(2)}| = \left\lceil \frac{k-t}{2} \right\rceil$ and $|\widetilde{C}' \cap S^{(2)}| = |\widetilde{C} \cap T^{(1)}| = \left\lfloor \frac{k-t}{2} \right\rfloor$, or
- $|\widetilde{C} \cap S^{(1)}| = |\widetilde{C}' \cap T^{(2)}| = \left\lfloor \frac{k-t}{2} \right\rfloor$ and $|\widetilde{C}' \cap S^{(2)}| = |\widetilde{C} \cap T^{(1)}| = \left\lceil \frac{k-t}{2} \right\rceil$, or

[Figure 1](#) shows an example of two edges $C, C' \in \mathcal{H}_i$ forming an edge (S, T) in the Kikuchi graph.

We say that the edge (S, T) is *type- i* , and for $S \in V$, we define the *type- i* degree as

$$d_{S,i} := \left| \left\{ C \in \mathcal{H}_i : |\widetilde{C} \cap S^{(1)}| \text{ or } |\widetilde{C} \cap S^{(2)}| \in \left\{ \left\lceil \frac{k-t}{2} \right\rceil, \left\lfloor \frac{k-t}{2} \right\rfloor \right\} \right\} \right|.$$

We call any subgraph of the colored Kikuchi graph as a *colored Kikuchi subgraph*.




Figure 1: An example of Definition 3.5 with $k = 5$ and $t = 2$. On the left are two 5-uniform hyperedges in \mathcal{H}_i with common intersection $U_i = \{1, 2\}$ and $\tilde{C} = \{3, 4, 5\}$, $\tilde{C}' = \{6, 7, 8\}$. On the right, S and T are vertices in the Kikuchi graph where $S^{(1)} = \{3, 4, 6\}$, $T^{(1)} = \{5, 6\}$ are colored green, and $S^{(2)} = \{6, 9\}$, $T^{(2)} = \{7, 8, 9\}$ are colored blue. C and C' form an edge between S, T because $|\tilde{C} \cap S^{(1)}| = 2$, $|\tilde{C} \cap T^{(1)}| = 1$, $|\tilde{C}' \cap S^{(2)}| = 1$, and $|\tilde{C}' \cap T^{(2)}| = 2$.

Remark 3.6 (Purpose of coloring). The coloring in Definition 3.5 is needed because $C \neq C' \in \mathcal{H}_i$ may have intersection larger than t , meaning $|C \oplus C'| = |\tilde{C} \oplus \tilde{C}'| < 2(k - t)$, making the analysis complicated. Coloring \tilde{C}, \tilde{C}' with different colors automatically makes $\tilde{C}^{(1)}, \tilde{C}'^{(2)}$ disjoint, i.e. $|S \oplus T| = |\tilde{C}^{(1)} \oplus \tilde{C}'^{(2)}| = 2(k - t)$. Note also that a vertex $S \subseteq [n] \times [2]$ may contain two copies of some element in $[n]$ with different colors, as shown in Figure 1.

Observation 3.7 (Parameters of the Kikuchi graph). The Kikuchi graph (V, E) defined in Definition 3.5 has $|V| = \binom{2n}{r}$, and each distinct pair $C, C' \in \mathcal{H}_i$ contributes a collection of edges $E_{C, C'}$ in E , where

$$|E_{C, C'}| = \alpha_t := \binom{k-t}{\lfloor \frac{k-t}{2} \rfloor} \binom{k-t}{\lceil \frac{k-t}{2} \rceil} \binom{2n-2(k-t)}{r-(k-t)} \cdot 2^{1(k-t \text{ is odd})}$$

by first choosing $\tilde{C} \cap S^{(1)}, \tilde{C}' \cap S^{(2)}$ (or $\tilde{C} \cap S^{(2)}, \tilde{C}' \cap S^{(1)}$) and completing S 's remaining $r - (k - t)$ elements. Thus, $|E| = \sum_{i=1}^p \binom{|\mathcal{H}_i|}{2} \cdot \alpha_t$, and standard calculations show that when $2k \leq r \leq n/8$, the average degree $d = \frac{2|E|}{|V|}$ satisfies

$$\left(\frac{r}{2n}\right)^{k-t} \sum_{i=1}^p \binom{|\mathcal{H}_i|}{2} \leq d \leq 2^{2k} \left(\frac{r}{2n}\right)^{k-t} \sum_{i=1}^p \binom{|\mathcal{H}_i|}{2}.$$

Our ideal hope is that the adjacency matrix A of the Kikuchi graph, constructed from $\mathcal{H}^{(i)} = (\mathcal{H}_1^{(i)}, \dots, \mathcal{H}_p^{(i)})$, is bounded in the PSD order by some low-trace diagonal matrix Q . To achieve this, we prove the following lemma analogous to Lemma 2.6, but with the additional requirement that $d_{S,i}$ is small for all $S \in V(K_r)$ and $i \in [p]$. The proof is almost identical to the proof of Lemma 2.6 but the encoding for an “old hyperedge” step is different.

Lemma 3.8. *Let $r \geq 2k$. Given disjoint hyperedges $\mathcal{H}_1, \dots, \mathcal{H}_p$, let \hat{A} be the adjacency matrix of any colored Kikuchi subgraph \hat{K}_r as defined in Definition 3.5, and let $\Gamma = D + d\mathbf{I}$ where D is the degree matrix*

and d is the average degree of G . Fix $\eta \in \mathbb{R}$ and let $\ell \in \mathbb{N}$ be even. Suppose there is no even cover of size at most ℓ , and suppose $d_{S,i} \leq \eta$ for all $S \in V$ and $i \in [p]$. Then,

$$\left\| \Gamma^{-1/2} \widehat{A} \Gamma^{-1/2} \right\|_2 \leq 2n^{r/\ell} \sqrt{\frac{2\eta\ell}{d}}.$$

Proof. Let $\widetilde{A} = \Gamma^{-1/2} \widehat{A} \Gamma^{-1/2}$. We again use the trace power method:

$$\|\widetilde{A}\|_2^\ell \leq \text{tr}(\widetilde{A}^\ell) = \text{tr}((\Gamma^{-1}A)^\ell).$$

Note that each edge (S, T) in \widehat{A} corresponds to two hyperedges of the same type (both from some \mathcal{H}_i), one green and one blue, and since there is no even covers of size at most ℓ , any closed walk must contain an even number of each hyperedge.

We encode a closed walk $S_1 \rightarrow S_2 \rightarrow \dots \rightarrow S_\ell \rightarrow S_1$ as follows:

- Starting vertex $S_1 \in V$.
- One bit $b_i \in \{0, 1\}$ at step i to encode whether this step uses two new hyperedges or one (or more) old hyperedge.
 - If $b_i = 0$ (two new hyperedges), select one of S_i 's neighbors as S_{i+1} .
 - If $b_i = 1$ (old hyperedge), select an old green (or blue) hyperedge C from the previous steps, and select a blue (or green) hyperedge C' incident to S_i .

Recall that for $b \in \{0, 1\}$, we write $N_b(S)$ as the possible next steps in the walk from S . Using the same analysis as the proof of [Lemma 2.6](#), for $b = 0$,

$$\sum_{S_{i+1} \in N_0(S_i)} \frac{1}{d_{S_i} + d} \leq 1,$$

and for $b = 1$, suppose the old edge is of type $j \in [p]$, then $|N_1(S_i)| \leq 2\ell d_{S_i, j}$ (one previous step, 2 colors), thus

$$\sum_{S_{i+1} \in N_b(S_i)} \frac{1}{d_{S_i} + d} \leq \frac{2\ell d_{S_i, j}}{d_{S_i} + d} \leq \frac{2\eta\ell}{d}.$$

We can assume that $2\eta\ell \leq d$, otherwise we can simply treat all steps as new hyperedges.

There are $\binom{2n}{r} \leq \left(\frac{2en}{r}\right)^r \leq n^r$ (since $r \geq 2k$ and $k \geq 3$) choices to pick the starting vertex S_1 . Furthermore, there can be at most $\ell/2$ steps that use two new hyperedges, i.e. $|b| \geq \ell/2$, thus

$$\text{tr}((\Gamma^{-1}\widehat{A})^\ell) \leq n^r \sum_{b \in \{0,1\}^\ell} \left(\frac{2\eta\ell}{d}\right)^{|b|} \leq 2^\ell n^r \left(\frac{2\eta\ell}{d}\right)^{\ell/2}.$$

Taking the ℓ -th root completes the proof. □

Construction of colored Kikuchi subgraph. Unfortunately, the requirement for all $d_{S,i}$ to be bounded by a small η prohibits us from obtaining a good bound on the adjacency matrix of the full colored Kikuchi graph K_r using [Lemma 3.8](#). This motivates dropping a small number of edges from K_r , and bounding the adjacency matrix \widehat{A} of the resulting subgraph \widehat{K}_r instead. Thus, we proceed with identifying a suitable colored Kikuchi subgraph \widehat{K}_r of $\mathcal{H}^{(i)}$ with adjacency matrix \widehat{A} via the following *edge deletion process*:

Start with the colored Kikuchi graph K_r , and delete every edge $\{S, T\}$ caused by a pair of clauses C, C' such that S or T has strictly more than 1 edge that C or C' participates in.

To obtain a handle on the average degree of \widehat{K}_r , we first show that the number of edges of K_r we delete to obtain \widehat{K}_r is only a small fraction of the total number of edges, and then the desired lower bound follows from a lower bound on $|E(K_r)|$.

Analyzing the edge deletion process. We find it convenient to think of the fraction of deleted edges as the *probability that a uniformly random edge in K_r is absent in \widehat{K}_r* . With this probabilistic interpretation in hand, observe that a uniformly random edge in K_r is the same as choosing a uniformly random pair of clauses (C, C') such that C and C' both belong to the same $\mathcal{H}_j^{(i)}$ and then choosing a random edge $\{S, T\}$ in $E_{C, C'}$, the collection of edges adorned by (C, C') . We will use the notation $C'' \rightarrow_C S$ to mean $|\widetilde{C}'' \cap S| = |\widetilde{C} \cap S|$, where we recall from [Definition 3.5](#) that $\widetilde{C} := C \setminus U_j$ with U_j being the size- i common intersection of $\mathcal{H}_j^{(i)}$. We then show the following.

Claim 3.9 (Deletion probability). *For every pair of clauses (C, C') such that C and C' belong to the same $\mathcal{H}_j^{(i)}$ for some $j \in [p]$,*

$$\Pr_{\{S, T\} \sim E_{C, C'}} [\{S, T\} \text{ deleted}] \leq k \cdot 4^{k+1} \sqrt{\frac{r}{n}}.$$

Proof. Recall that we defined $\widetilde{C} = C \setminus U_j$ and $\widetilde{C}' = C' \setminus U_j$. The distribution of $S = (S^{(1)}, S^{(2)})$ (the green and blue vertices) is uniform on all sets such that:

- $|\widetilde{C} \cap S^{(1)}| = \left\lceil \frac{k-i}{2} \right\rceil, |\widetilde{C}' \cap S^{(2)}| = \left\lfloor \frac{k-i}{2} \right\rfloor$, or
- $|\widetilde{C} \cap S^{(1)}| = \left\lfloor \frac{k-i}{2} \right\rfloor, |\widetilde{C}' \cap S^{(2)}| = \left\lceil \frac{k-i}{2} \right\rceil$.

Then, by union bound,

$$\begin{aligned} \Pr_{\{S, T\} \sim E_{C, C'}} [\{S, T\} \text{ deleted}] &\leq \Pr_{\{S, T\} \sim E_{C, C'}} [\exists C'' \rightarrow_C S^{(1)} : C'' \in \mathcal{H}_j^{(i)}, C'' \neq C] + \\ &\quad \Pr_{\{S, T\} \sim E_{C, C'}} [\exists C'' \rightarrow_{C'} S^{(2)} : C'' \in \mathcal{H}_j^{(i)}, C'' \neq C'] + \\ &\quad \Pr_{\{S, T\} \sim E_{C, C'}} [\exists C'' \rightarrow_C T^{(1)} : C'' \in \mathcal{H}_j^{(i)}, C'' \neq C] + \\ &\quad \Pr_{\{S, T\} \sim E_{C, C'}} [\exists C'' \rightarrow_{C'} T^{(2)} : C'' \in \mathcal{H}_j^{(i)}, C'' \neq C'] \\ &= 4 \Pr_{\{S, T\} \sim E_{C, C'}} [\exists C'' \rightarrow_C S^{(1)} : C'' \in \mathcal{H}_j^{(i)}, C'' \neq C] \end{aligned}$$

then by Markov's inequality,

$$\begin{aligned} &\leq 4 \mathbb{E}_{\{S, T\} \sim E_{C, C'}} \left[\mathbf{1}_{C'' : C'' \rightarrow_C S^{(1)}, C'' \in \mathcal{H}_j^{(i)}, C'' \neq C} \right] \\ &= 4 \sum_{\substack{C'' : C'' \in \mathcal{H}_j^{(i)} \\ C'' \neq C}} \Pr_{\{S, T\} \sim E_{C, C'}} [C'' \rightarrow_C S^{(1)}] \end{aligned} \tag{2}$$

Once the intersection of S with \tilde{C} and \tilde{C}' is chosen, the remaining elements are selected uniformly at random without replacement. For fixed $C'' \neq C \in \mathcal{H}_j^{(i)}$, since they contain U_j of size i , $|\tilde{C}'' \cap \tilde{C}| = |C'' \cap C| - i$, and S must include $\lfloor \frac{k-i}{2} \rfloor - (|C'' \cap C| - i)$ additional elements from $\tilde{C}'' \setminus \tilde{C}$ for $C'' \rightarrow_C S^{(1)}$ to hold. Thus,

$$\Pr_{\{S,T\} \sim E_{C,C'}} [C'' \rightarrow_C S^{(1)}] \leq 2^k \left(\frac{r}{n}\right)^{\lfloor \frac{k-i}{2} \rfloor - |C'' \cap C| + i}.$$

Thus, we can prove:

$$(2) \leq 4 \cdot 2^k \sum_{s=i}^{k-1} \sum_{\substack{U \subseteq C \\ |U|=s}} \sum_{\substack{C'' : C'' \in \mathcal{H}_j^{(i)} \\ C'' \neq C \\ C'' \cap C = U}} \left(\frac{r}{n}\right)^{\lfloor \frac{k-i}{2} \rfloor - s + i} \quad (3)$$

By [Observation 3.3](#), we can bound the above as

$$\begin{aligned} &\leq 4 \cdot 2^k \sum_{s=i}^{k-1} \sum_{\substack{U \subseteq C \\ |U|=s}} \binom{n}{r}^{\frac{k}{2} - s} \left(\frac{r}{n}\right)^{\frac{k-i}{2} - \frac{\mathbf{1}[k-i \text{ odd}]}{2} - s + i} \\ &\leq k \cdot 4^{k+1} \sqrt{\frac{r}{n}}, \end{aligned}$$

as $\frac{i}{2} - \frac{\mathbf{1}[k-i \text{ odd}]}{2} \geq \frac{1}{2}$ for all $i \geq 1$ when k is odd. \square

Lower bound on average degree in A . By choosing B large enough, the upper bound on r , and [Claim 3.9](#), the fraction of edges we delete from the original colored Kikuchi graph K_r to obtain \hat{K}_r is at most .5 and hence $d(\hat{K}_r) \geq .5d(K_r)$ where $d(K_r)$ and $d(\hat{K}_r)$ are the average degrees in K_r and \hat{K}_r respectively. Thus, we know:

$$d(K_r) \geq \left(\frac{r}{2n}\right)^{k-i} \sum_{j=1}^p \binom{|\mathcal{H}_j^{(i)}|}{2} \geq \left(\frac{r}{2n}\right)^{k-i} \cdot p \cdot \binom{m/kp}{2} \geq \left(\frac{r}{2n}\right)^{k-i} \cdot \frac{m^2}{4k^2p}$$

where the first inequality uses [Observation 3.7](#), and the second inequality is due to Jensen's inequality.

By the upper bound $p \leq m \cdot \left(\frac{r}{n}\right)^{\frac{k}{2} - i}$ as noted in [Observation 3.2](#):

$$d(K_r) \geq \frac{1}{4k^2 2^k} \cdot \left(\frac{r}{n}\right)^{k-i} \cdot \left(\frac{n}{r}\right)^{\frac{k}{2} - i} \cdot m = \frac{1}{4k^2 2^k} \cdot \left(\frac{r}{n}\right)^{\frac{k}{2}} \cdot m.$$

As an upshot, we know:

Claim 3.10. $d(\hat{K}_r) \geq \frac{1}{8k^2 2^k} \cdot \left(\frac{r}{n}\right)^{k/2} \cdot m.$

Spectral double counting. With a lower bound on $d(\widehat{K}_r)$ in hand, we are now ready to perform our weighted spectral double counting argument to complete the proof of [Theorem 3.1](#).

Proof of [Theorem 3.1](#). Recall that our goal is to prove that there is a small even cover in $\mathcal{H}^{(i)}$, the largest piece obtained from the decomposition, and also recall that if $i \geq \frac{k+1}{2}$, then we are done by [Lemma 3.4](#). Hence, we assume $i \leq \frac{k-1}{2}$ for the rest of the proof.

Suppose there are no even covers in \mathcal{H} of size $\ell = r \log n$, then there are also none in $\mathcal{H}^{(i)}$ from [Lemma 3.8](#) we get:

$$\left\| \Gamma^{-1/2} \widehat{A} \Gamma^{-1/2} \right\|_2 \leq 4 \sqrt{\frac{2\ell}{d(\widehat{K}_r)}}.$$

Thus, $\widehat{A} \preceq 4 \sqrt{\frac{2\ell}{d(\widehat{K}_r)}} \Gamma$, and by taking the quadratic form with the all-ones vector, we get:

$$2|E(\widehat{K}_r)| = \mathbf{1}^\top \widehat{A} \mathbf{1} \leq 4 \sqrt{\frac{2\ell}{d(\widehat{K}_r)}} \cdot \text{tr}(\Gamma) = 16 \sqrt{\frac{2\ell}{d(\widehat{K}_r)}} \cdot |E(\widehat{K}_r)|,$$

which implies

$$d(\widehat{K}_r) \leq 128\ell,$$

and by our lower bound on $d(\widehat{K}_r)$ from [Claim 3.10](#), we get

$$\frac{1}{8k^2 2^k} \cdot \left(\frac{r}{n}\right)^{\frac{k}{2}} \cdot m \leq 128r \log n,$$

which we can rearrange as

$$m \leq B^k n \log n \cdot \left(\frac{n}{r}\right)^{k/2-1}.$$

for some large enough constant B . Thus, if m is lower bounded as in the theorem statement, there must be an even cover of size $\ell \log n$. \square

4 Strong refutation of semirandom k -XOR

In this section, we show that our reweighted Kikuchi matrix and edge deletion process yield a significantly simpler analysis of strong refutation algorithms for semirandom k -XOR formulas and lose only a single $\log n$ factor in the density. Combined with Feige's "XOR principle" [[Fei02](#), [AOW15](#)], we also obtain refutation algorithms for all *smoothed* Boolean CSPs. We will omit such reduction in this work and direct the reader to [[GKM21](#)] for a detailed exposition.

Theorem 4.1 (Semirandom k -XOR refutation). *Fix $k \in \mathbb{N}$. There is an algorithm with parameter $r \in \mathbb{N}$, $2k \leq r \leq n/8$ that takes as input a semirandom k -XOR instance*

$$\psi(x) = \frac{1}{m} \sum_{C \in \mathcal{H}} b_C x_C$$

where \mathcal{H} is a k -uniform hypergraph with n vertices and m hyperedges, and each $b_C \in \{\pm 1\}$ is chosen uniformly at random. The algorithm has the following guarantee: there is a universal constant C such that if $m \geq C^k n \log n \cdot \left(\frac{n}{r}\right)^{\frac{k}{2}-1} \varepsilon^{-4}$ for $\varepsilon \in (0, 1/2)$, then with probability over $1 - \frac{1}{\text{poly}(n)}$ over $\{b_C\}_{C \in \mathcal{H}}$, the algorithm runs in time $n^{O(r)}$ and certifies that $\psi(x) \leq \varepsilon$.

Remark 4.2 (Refutation strength: dependence on ε). For the even arity case, we actually obtain a stronger guarantee (weaker requirement) of $m \geq O(n \log n) \cdot \left(\frac{n}{r}\right)^{\frac{k}{2}-1} \varepsilon^{-2}$. For the odd arity case however, our analysis incurs a (likely suboptimal) dependence of $1/\varepsilon^4$ on the refutation strength (i.e., the upper bound on the value of the input k -XOR instance), though improving the $1/\varepsilon^5$ dependence of [GKM21, Theorem 5.1]. In contrast, a $1/\varepsilon^2$ dependence is known to hold for fully random k -XOR instances [RRS17]. Apart from a somewhat unsatisfying deficiency, this suboptimality turns out to be consequential – in particular, it changes the threshold at which efficient FKO refutation witnesses exist for semirandom k -SAT (and other CSPs) by a polynomial factor in n . Finding the “right” dependence of $1/\varepsilon^2$ (for the odd case) is an interesting open problem.

Our refutation algorithm will utilize the same Kikuchi graphs from Definition 2.4 and Definition 3.5 but with signs added to the edges in the natural way.

Definition 4.3 (Signed Kikuchi graph). Let \mathcal{H} be a k -uniform hypergraph associated with $\{\pm 1\}$ signs $\{b_C\}_{C \in \mathcal{H}}$. For the even arity case, let A_b be the signed adjacency matrix of the Kikuchi graph from Definition 2.4 where each edge $S \xleftrightarrow{C} T$ has a sign b_C . For the odd arity case, let A_b be the signed adjacency matrix of the Kikuchi graph from Definition 3.5 where each edge $S \xleftrightarrow{C, C'} T$ has a sign $b_C b_{C'}$.

4.1 Refuting semirandom even arity XOR

In this section, we prove Theorem 4.1 when k is even. As we will see in the short proof, our idea of the reweighted Kikuchi matrix from the hypergraph Moore bound naturally applies here, and in fact, we obtain the “right” $1/\varepsilon^2$ dependence in this case, i.e., we can certify that $\psi(x) \leq \varepsilon$ when $m \geq O(n \log n) \cdot \left(\frac{n}{r}\right)^{\frac{k}{2}-1} \varepsilon^{-2}$.

Recall that in the Kikuchi graph (V, E) , each $C \in \mathcal{H}$ contributes $\alpha := \frac{1}{2} \binom{k}{k/2} \binom{n-k}{r-k/2}$ edges in E , hence $|E| = \frac{1}{2} |V| d = m\alpha$. Thus, it is clear that

$$\psi(x) = \frac{1}{m} \cdot \frac{1}{\alpha} \sum_{(S, T) \in E} b_{S \oplus T} x_{S \oplus T} = \frac{1}{\binom{n}{r} d} (x^{\odot r})^\top A_b x^{\odot r} \quad (4)$$

where $x^{\odot r} \in \{\pm 1\}^{\binom{n}{r}}$ and the S -entry of $x^{\odot r}$ is x_S for $S \subseteq [n]$, $|S| = r$.

We now follow the same reweighting strategy: with $\Gamma = D + d\mathbb{I}$, we bound the spectral norm of the reweighted Kikuchi matrix $\|\Gamma^{-1/2} A_b \Gamma^{-1/2}\|_2$ with an almost identical proof as Lemma 2.6.

Lemma 4.4. *Let k be even and $r \in \mathbb{N}$. Let A_b be the signed Kikuchi graph with random $\{\pm 1\}$ coefficients $\{b_C\}_{C \in \mathcal{H}}$, and let $\Gamma = D + d\mathbb{I}$ where D is the degree matrix and d is the average degree of the Kikuchi graph. Then, with probability at least $1 - \frac{1}{\text{poly}(n)}$ over the randomness of $\{b_C\}_{C \in \mathcal{H}}$,*

$$\left\| \Gamma^{-1/2} A_b \Gamma^{-1/2} \right\|_2 \leq O\left(\sqrt{\frac{r \log n}{d}} \right).$$

Proof. Let $\tilde{A}_b = \Gamma^{-1/2} A_b \Gamma^{-1/2}$. We again use the trace power method $\|\tilde{A}_b\|_2^\ell \leq \text{tr}((\Gamma^{-1} A_b)^\ell)$ where we choose an even $\ell = 2 \lceil r \log_2 n \rceil$. Observe that in expectation, $\mathbb{E}_b \text{tr}((\Gamma^{-1} A_b)^\ell)$ counts the closed walks that use each hyperedge an even number of times. This is exactly the same as Lemma 2.6

where we count closed walks in an unsigned Kikuchi graph assuming there is no even cover of size $\leq \ell$. Thus, [Lemma 2.6](#) shows that

$$\mathbb{E}_b \operatorname{tr}((\Gamma^{-1} A_b)^\ell) \leq 2^\ell n^r \left(\frac{\ell}{d}\right)^{\ell/2} \leq O\left(\frac{\ell}{d}\right)^{\ell/2}$$

when $\ell \geq r \log_2 n$. Then, by Markov's inequality, for any $\lambda > 0$,

$$\Pr_b \left[\|\tilde{A}_b\|_2 \geq \lambda \right] = \Pr_b \left[\|\tilde{A}_b\|_2^\ell \geq \lambda^\ell \right] \leq \lambda^{-\ell} \cdot \mathbb{E}_b \operatorname{tr}((\Gamma^{-1} A_b)^\ell) \leq O\left(\frac{\ell}{\lambda^2 d}\right)^{\ell/2}$$

Choosing $\lambda = O(\sqrt{\ell/d})$ completes the proof. \square

We can complete the proof of [Theorem 4.1](#) for even k .

Proof of [Theorem 4.1](#) for even k . Let A_b be the signed Kikuchi graph with signs $\{b_C\}_{C \in \mathcal{H}}$, let $\Gamma = D + d\mathbb{I}$ where D is the degree matrix and d is the average degree of the Kikuchi graph, and let $\tilde{A}_b = \Gamma^{-1/2} A_b \Gamma^{-1/2}$. The certification algorithm is simply to compute $\|\tilde{A}_b\|_2$. Since $A_b \preceq \|\tilde{A}_b\|_2 \cdot \Gamma$, and $\operatorname{tr}(\Gamma) = 2 \binom{n}{r} d$, by [Lemma 4.4](#),

$$\psi(x) = (4) \leq \frac{1}{\binom{n}{r} d} \|\tilde{A}_b\|_2 \cdot \operatorname{tr}(\Gamma) \leq O\left(\sqrt{\frac{r \log n}{d}}\right)$$

using the fact that $x^{\odot r} \in \{\pm 1\}^{\binom{n}{r}}$ and $(x^{\odot r})^\top \Gamma x^{\odot r} = \operatorname{tr}(\Gamma)$. There is some constant C such that when $m \geq Cn \log n \cdot \left(\frac{n}{r}\right)^{\frac{k}{2}-1} \varepsilon^{-2}$, by [Eq. \(1\)](#) the average degree $d \geq \frac{1}{2} \left(\frac{r}{n}\right)^{k/2} m = \frac{C}{2} r \log n \cdot \varepsilon^{-2}$, thus giving us $\psi(x) \leq \varepsilon$. This completes the proof. \square

4.2 Refuting semirandom odd arity XOR

Our proof of [Theorem 4.1](#) for the odd arity case closely mimics the steps taken in proving the hypergraph Moore bound for odd arity hypergraphs ([Theorem 3.1](#)). Given a semirandom k -XOR instance ψ on hypergraph \mathcal{H} with random signs $\{b_C\}_{C \in \mathcal{H}}$, we first apply the following hypergraph decomposition algorithm (a variant of [Algorithm 1](#)) to decompose the hypergraph into subhypergraphs $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(k-1)}$. The main difference compared to [Algorithm 1](#) is that in the final step, we add the "leftover" hyperedges to $\mathcal{H}^{(1)}$ instead of an extra $\mathcal{H}^{(0)}$.

Algorithm 2 (Hypergraph decomposition). Given a k -uniform hypergraph \mathcal{H} on n vertices and m hyperedges, and thresholds $\tau_1, \dots, \tau_{k-1} \geq 2$, we partition \mathcal{H} into hypergraphs $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(k-1)}$ via the following algorithm.

1. Set $t = k - 1$ and $\mathcal{H}_{\text{current}} := \mathcal{H}$.
2. Set counter $s = 1$. While there is $T \subseteq [n]$ such that $|T| = t$ and $|\{C \in \mathcal{H}_{\text{current}} : T \subseteq C\}| \geq \tau_t$:
 - (a) Choose T satisfying the condition and let $\mathcal{H}_s^{(t)}$ be a subset of $\{C \in \mathcal{H}_{\text{current}} : T \subseteq C\}$ of size τ_t .
 - (b) Add all clauses in $\mathcal{H}_s^{(t)}$ to $\mathcal{H}^{(t)}$.

- (c) Delete all clauses in $\mathcal{H}_s^{(t)}$ to $\mathcal{H}_{\text{current}}$.
 - (d) Increment s by 1.
3. Decrement t by 1. If $t > 0$, go back to step 2; otherwise take the remaining clauses in $\mathcal{H}_{\text{current}}$ and partition them into n parts F_1, \dots, F_n where each clause C goes to some F_i such that $i \in C$. Add F_1, \dots, F_n to $\mathcal{H}^{(1)}$ and terminate.
-

Notations and parameters. Throughout this section we will use the following notations.

- In [Algorithm 2](#), we set thresholds $\tau_t = \max\left\{1, \left(\frac{n}{r}\right)^{\frac{k}{2}-t}\right\} \cdot 4k\varepsilon^{-2}$.
- In the decomposition, each $\mathcal{H}^{(t)}$ contains p_t groups $\mathcal{H}_1^{(t)}, \dots, \mathcal{H}_{p_t}^{(t)}$ where group $\mathcal{H}_i^{(t)}$ has a center $T_i^{(t)}$ of size t , and for each $C \in \mathcal{H}_i^{(t)}$, we write $\tilde{C} = C \setminus T_i^{(t)}$.
- Each $|\mathcal{H}_i^{(t)}| = \tau_t$, with the exception that $|\mathcal{H}_i^{(1)}| \leq \tau_1$ may have different sizes (the leftover hyperedges in [Algorithm 2](#)). Let $m_t := \sum_{i=1}^{p_t} |\mathcal{H}_i^{(t)}|$ be the total number of hyperedges in $\mathcal{H}^{(t)}$.
- When $t = 1$ and $m \geq C^k n \log n \cdot \left(\frac{n}{r}\right)^{\frac{k}{2}-1} \varepsilon^{-4}$ for a large enough constant C , we have $m \geq n\tau_1$, hence $p_1 \leq \frac{m}{\tau_1} + n \leq \frac{2m}{\tau_1}$. Thus, we will use $p_t \tau_t \leq 2m$ for all $t \in [k-1]$.
- For each $t \in [k-1]$, the colored Kikuchi graph (V, E) obtained from $\mathcal{H}^{(t)} = (\mathcal{H}_1^{(t)}, \dots, \mathcal{H}_{p_t}^{(t)})$ (from [Definition 3.5](#)) has edges $|E| = \alpha_t \sum_{i=1}^{p_t} \binom{|\mathcal{H}_i^{(t)}|}{2} \leq \frac{1}{2} \alpha_t m_t \tau_t$, where $\alpha_t \approx \left(\frac{2n}{r}\right)^{r-(k-t)}$ is the number of edges contributed by each distinct pair $C, C' \in \mathcal{H}_i$ (see [Observation 3.7](#)).

With these notations and parameters in mind, we can write $\psi(x)$ as

$$\psi(x) = \frac{1}{m} \sum_{t=1}^{k-1} \sum_{C \in \mathcal{H}^{(t)}} b_C x_C = \frac{1}{k} \sum_{t=1}^{k-1} \psi_t(x)$$

where $\psi_t(x) := \frac{k}{m} \sum_{i=1}^{p_t} \sum_{C \in \mathcal{H}_i^{(t)}} b_C x_C = \frac{k}{m} \sum_{i=1}^{p_t} x_{T_i} \sum_{C \in \mathcal{H}_i^{(t)}} b_C x_{\tilde{C}}.$ (5)

Essentially, each ψ_t is the sub-instance of ψ restricted to the partition $\mathcal{H}^{(t)}$. Recall that for the purpose of showing existence of even covers, we only need to focus on one $\mathcal{H}^{(t)}$. For refutation however, we need to certify a bound on $\psi_t(x)$ for all $t \in [k-1]$.

Lemma 4.5 (Refuting each ψ_t). *Fix an odd $k \in \mathbb{N}$, $t \in [k-1]$, and let $2k \leq r \leq n/8$. There is a constant C such that given a semirandom k -XOR instance ψ with n variables and $m \geq C^k n \log n \left(\frac{n}{r}\right)^{\frac{k}{2}-1} \varepsilon^{-4}$ clauses for $\varepsilon \in (0, 1/2)$, and suppose ψ_t is the subinstance from [Eq. \(5\)](#) obtained by the hypergraph decomposition algorithm ([Algorithm 2](#)), then with probability $1 - \frac{1}{\text{poly}(n)}$ over the random signs, we can certify that $\psi_t(x) \leq \varepsilon$ in $n^{O(r)}$ time.*

[Lemma 4.5](#) immediately completes the proof of [Theorem 4.1](#) for odd k .

Proof of Theorem 4.1 by Lemma 4.5. Given the hypergraph \mathcal{H} , we apply the hypergraph decomposition algorithm (Algorithm 2) with thresholds $\tau_1, \dots, \tau_{k-1}$ and obtain subinstances $\psi_1, \dots, \psi_{k-1}$ as in Eq. (5). For each $t \in [k-1]$, we can certify that $\psi_t(x) \leq \varepsilon$ by Lemma 4.5 with high probability, which immediately implies the desired bound $\psi(x) \leq \varepsilon$. \square

Edge deletion process. The proof of Lemma 4.5 requires deleting the “bad” edges from the signed Kikuchi matrix $A_b^{(t)}$ via a similar deletion process as the one used in the proof of Theorem 3.1, but with some parameter $\eta > 1$ instead of 1 and an additional *equalizing* step:

Start with the colored Kikuchi graph, and delete every edge $\{S, T\}$ caused by a pair of clauses $C, C' \in \mathcal{H}_i^{(t)}$ such that S or T has more than η edges that C or C' participates in.

Suppose $\rho < 1$ is the maximum fraction of edges deleted among all pairs of clauses. Then, for every $i \in [p_t]$ and every distinct pair $C, C' \in \mathcal{H}_i^{(t)}$, we delete (additional) edges caused by C, C' arbitrarily such that exactly ρ fraction of edges are deleted.

Observation 4.6 (Uniform deletion). The final step in the above edge deletion process ensures that every pair C, C' contributes the *same* number of edges $((1 - \rho)\alpha_t$ to be exact) in the Kikuchi graph.

Mirroring the proof of Claim 3.9 yields the following generalization.

Lemma 4.7 (Deletion rate). *Suppose a subhypergraph $\mathcal{H}^{(i)}$ satisfies that for any $s \geq i$ and any $T \subseteq [n]$ with $|T| = s$, the number of hyperedges in $\mathcal{H}^{(i)}$ containing T is at most τ_s , then the deletion process with parameter $\eta \geq 1$ satisfies*

$$\Pr_{\{S, T\} \sim E_{C, C'}} [\{S, T\} \text{ deleted}] \leq \frac{4^k}{\eta} \cdot \sum_{s=i}^{\lfloor \frac{k+i}{2} \rfloor} \tau_s \left(\frac{r}{n}\right)^{\lfloor \frac{k+i}{2} \rfloor - s}.$$

Proof. The proof is identical to the proof of Claim 3.9. Eq. (2) holds with an additional $1/\eta$ factor due to Markov’s inequality. The lemma statement then follows immediately from Eq. (3). \square

Proof of Lemma 4.5 via the Cauchy-Schwarz trick and the deletion process.

Proof of Lemma 4.5. We apply the Cauchy-Schwarz trick to ψ_t from Eq. (5):

$$\begin{aligned} \psi_t(x)^2 &\leq \frac{k}{m^2} \sum_{i=1}^{p_t} x_{T_i}^2 \cdot \sum_{i=1}^{p_t} \left(\sum_{C \in \mathcal{H}_i^{(t)}} b_C x_{\bar{C}} \right)^2 \leq \frac{kp_t}{m^2} \sum_{i=1}^{p_t} \sum_{C, C' \in \mathcal{H}_i^{(t)}} b_C b_{C'} x_{\bar{C}} x_{\bar{C}'} \\ &\leq \frac{kp_t m_t}{m^2} + \frac{kp_t}{m^2} \sum_{i=1}^{p_t} \sum_{C \neq C' \in \mathcal{H}_i^{(t)}} b_C b_{C'} x_{C \oplus C'} \end{aligned} \quad (6)$$

since $x \in \{\pm 1\}^n$, $b_C \in \{\pm 1\}$ and $\sum_{i=1}^{p_t} |\mathcal{H}_i^{(t)}| = m_t$. For the first term, since for all $t \in [k-1]$, we set $\tau_t \geq 4k\varepsilon^{-2}$ and $p_t \leq 2m/\tau_t \leq \frac{m\varepsilon^2}{2k}$, thus

$$\frac{kp_t m_t}{m^2} \leq \frac{\varepsilon^2}{2}. \quad (7)$$

We can now focus our attention on the second term in Eq. (6).

Given $\mathcal{H}^{(t)}$ and its partitions $\mathcal{H}_1^{(t)}, \dots, \mathcal{H}_{p_t}^{(t)}$ of size τ_t , and signs $\{b_C\}_{C \in \mathcal{H}^{(t)}}$, let $A_b^{(t)}$ be the signed Kikuchi matrix defined in [Definition 4.3](#), which is the signed version of the colored Kikuchi graph (V, E) from [Definition 3.5](#). Recall from [Observation 3.7](#) that each distinct pair $C, C' \in \mathcal{H}_i^{(t)}$ contributes $\alpha_t \approx \left(\frac{2n}{r}\right)^{r-(k-t)}$ edges in the graph. Thus, similar to [\(4\)](#) in the even case, we can write the second term of [Eq. \(6\)](#) as a quadratic form:

$$f_t(x) := \frac{kp_t}{m^2} \sum_{i=1}^{p_t} \sum_{C \neq C' \in \mathcal{H}_i^{(t)}} b_C b_{C'} x_{C \oplus C'} = \frac{kp_t}{2\alpha_t m^2} (x^{\odot r})^\top A_b^{(t)} x^{\odot r} \quad (8)$$

where $x^{\odot r} \in \{\pm 1\}^{\binom{2n}{r}}$ such that for $S \in [n] \times [2]$ with $S = (S^{(1)}, S^{(2)})$ (green and blue elements), the S -entry of $x^{\odot r}$ is $x_{S^{(1)} \oplus S^{(2)}}$.

We proceed to certify an upper bound on $f_t(x)$. Given the signed Kikuchi matrix $A_b^{(t)}$, we first apply the deletion process with parameter $\eta = B^k \varepsilon^{-2}$ for some large enough constant B . With the chosen thresholds τ_s , [Lemma 4.7](#) states that the deletion probability ρ is at most

$$\rho \leq \frac{4^k}{\eta} \cdot \sum_{s=t}^{\lfloor \frac{k+t}{2} \rfloor} 4k\varepsilon^{-2} \cdot \max\left\{1, \left(\frac{n}{r}\right)^{\frac{k}{2}-s}\right\} \cdot \left(\frac{r}{n}\right)^{\lfloor \frac{k+t}{2} \rfloor - s} \leq \frac{1}{2},$$

since $s \leq \lfloor \frac{k+t}{2} \rfloor$ in the summation and $\lfloor \frac{k+t}{2} \rfloor \geq \frac{k+1}{2}$ for all $t \geq 1$.

Let $\widehat{A}_b^{(t)}$ be the Kikuchi matrix after the deletion process. By [Observation 4.6](#), each distinct pair $C, C' \in \mathcal{H}_i^{(t)}$ contributes exactly $(1 - \rho)$ fraction of the original edges. Thus, we have

$$(x^{\odot r})^\top \widehat{A}_b^{(t)} x^{\odot r} = (1 - \rho) \cdot (x^{\odot r})^\top A_b^{(t)} x^{\odot r}. \quad (9)$$

Next, we follow the same argument as the proof of [Lemma 4.4](#) to analyze $\widehat{A}_b^{(t)}$, using the norm bound of [Lemma 3.8](#). Let $\Gamma = D + d\mathbb{I}$ where D is the degree matrix and d is the average degree, and let $\widetilde{A}_b = \Gamma^{-1/2} \widehat{A}_b^{(t)} \Gamma^{-1/2}$. To bound $\|\widetilde{A}_b\|_2$, we again use the trace power method $\|\widetilde{A}_b\|_2^\ell \leq \text{tr}((\Gamma^{-1} \widehat{A}_b^{(t)})^\ell)$ where we choose an even $\ell = 2\lceil r \log_2 n \rceil$. Observe that in expectation, $\mathbb{E}_b \text{tr}((\Gamma^{-1} A_b)^\ell)$ counts the closed walks that use each hyperedge an even number of times. This is exactly the same as [Lemma 3.8](#) where we count closed walks in an unsigned Kikuchi graph assuming there is no even cover of size $\leq \ell$. Furthermore, $d_{S,i} \leq \eta$ is automatically satisfied after the deletion process. Thus, we can directly apply [Lemma 3.8](#) and show that

$$\mathbb{E}_b \text{tr}((\Gamma^{-1} \widehat{A}_b^{(t)})^\ell) \leq 2^\ell n^r \left(\frac{2\eta\ell}{d}\right)^{\ell/2} \leq O\left(\frac{\eta\ell}{d}\right)^{\ell/2}$$

when $\ell \geq r \log_2 n$. Then, by Markov's inequality, we have that $\Pr_b \left[\|\widetilde{A}_b\|_2 \geq O\left(\sqrt{\frac{\eta\ell}{d}}\right) \right] \leq \frac{1}{\text{poly}(n)}$.

Thus, with high probability we have $\widehat{A}_b^{(t)} \preceq O\left(\sqrt{\frac{\eta\ell}{d}}\right) \cdot \Gamma$, then since $\text{tr}(\Gamma) = 4|E|$,

$$(x^{\odot r})^\top \widehat{A}_b^{(t)} x^{\odot r} \leq O\left(\sqrt{\frac{\eta\ell}{d}}\right) \cdot \text{tr}(\Gamma) = O\left(\sqrt{\frac{\eta\ell}{d}}\right) \cdot |E|.$$

Next, let $\widehat{f}_t(x) = \frac{kp_t}{2\alpha_t m^2} (x^{\odot r})^\top \widehat{A}_b^{(t)} x^{\odot r}$. By [Observation 3.7](#), we have $d \geq (\frac{r}{2n})^{k-t} \sum_{i=1}^{p_t} \binom{|\mathcal{H}_i^{(t)}|}{2}$ when $2k \leq r \leq n/8$. Plugging in parameters $|E| = \alpha_t \sum_{i=1}^{p_t} \binom{|\mathcal{H}_i^{(t)}|}{2}$, $p_t \tau_t \leq 2m$, $\eta = B^k \varepsilon^{-2}$, and $\ell = 2\lceil r \log_2 n \rceil$, standard calculations show that

$$\widehat{f}_t(x) \leq O(1) \frac{kp_t}{\alpha_t m^2} \sqrt{\frac{\eta \ell}{d}} |E| \leq O(1) \frac{kp_t}{m^2} \sqrt{\eta \ell \left(\frac{2n}{r}\right)^{k-t} \sum_{i=1}^{p_t} \binom{|\mathcal{H}_i^{(t)}|}{2}} \leq O(1) \sqrt{\frac{\eta r \log n}{m \tau_t} \left(\frac{2n}{r}\right)^{k-t}}.$$

Suppose $m \geq C^k n \log n \cdot (\frac{n}{r})^{\frac{k}{2}-1} \varepsilon^{-4}$ for some large enough constant C . We split into cases:

1. For $t \leq \frac{k-1}{2}$, we set $\tau_t = (\frac{n}{r})^{\frac{k}{2}-t} \cdot 4k\varepsilon^{-2}$, thus $\widehat{f}_t(x) \leq \frac{\varepsilon^2}{4}$.
2. For $t \geq \frac{k+1}{2}$, we set $\tau_t = 4k\varepsilon^{-2}$, thus $\widehat{f}_t(x) \leq \frac{\varepsilon^2}{4} (\frac{n}{r})^{\frac{k}{4}-\frac{t}{2}} < \frac{\varepsilon^2}{4}$.

Therefore, by calculating $\|\widetilde{A}_b\|_2$, which can be done in $n^{O(r)}$ time, we can certify that $\widehat{f}_t(x) \leq \frac{\varepsilon^2}{4}$. Combined with [Eq. \(9\)](#) and the bound of $\rho \leq 1/2$, we can certify that

$$f_t(x) \leq \frac{1}{1-\rho} \cdot \widehat{f}_t(x) \leq \frac{\varepsilon^2}{2},$$

and with [Eq. \(7\)](#), we can certify an upper bound on [Eq. \(6\)](#):

$$\psi_t(x)^2 \leq (7) + (8) \leq \frac{\varepsilon^2}{2} + f_t(x) \leq \varepsilon^2,$$

completing the proof. □

Acknowledgments

We would like to thank Omar Alrabiah for heroic feedback on an earlier draft of this paper and Benny Sudakov for an illuminating discussion and comments. J.H. would like to thank Peter Manohar for discussions about [\[GKM21\]](#). S.M. would like to thank Siqi Liu and Tselil Schramm for discussions on an independent problem that fueled some of the ideas. Finally, much of this work was conducted when the authors were visiting the program on “Computational Complexity of Statistical Inference” at Simons Institute for the Theory of Computing in Fall 2021. We would like to thank the institute for their hospitality and support! S.M. was also visiting Microsoft Research Redmond when part of this work was done.

References

- [AF09] Noga Alon and Uriel Feige. On the power of two, three and four probes. In *Proceedings of the twentieth annual ACM-SIAM symposium on Discrete algorithms*, pages 346–354. SIAM, 2009. 1
- [AHL02] Noga Alon, Shlomo Hoory, and Nathan Linial. The Moore bound for irregular graphs. *Graphs and Combinatorics*, 18(1):53–57, 2002. [0](#), [1](#), [3](#), [4](#), [25](#), [26](#)

- [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*, pages 689–708. IEEE, 2015. [3](#), [17](#)
- [Big93] Norman Biggs. *Algebraic graph theory*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, second edition, 1993. [1](#)
- [BKHL99] Claudia Bertram-Kretzberg, Thomas Hofmeister, and Hanno Lefmann. Sparse 0-1 matrices and forbidden hypergraphs. *Combinatorics, Probability and Computing*, 8(5):417–427, 1999. [1](#)
- [BMS08] Louay Bazzi, Mohammad Mahdian, and Daniel A Spielman. The minimum distance of turbo-like codes. *IEEE Transactions on Information Theory*, 55(1):6–15, 2008. [1](#)
- [Bol78] Béla Bollobás. *Extremal graph theory*, volume 11 of *London Mathematical Society Monographs*. Academic Press, Inc. [Harcourt Brace Jovanovich, Publishers], London-New York, 1978. [1](#)
- [BR14] S Ajesh Babu and Jaikumar Radhakrishnan. An entropy-based proof for the Moore bound for irregular graphs. In *Perspectives in Computational Complexity*, pages 173–181. Springer, 2014. [1](#), [25](#), [26](#)
- [CS22] David Munhá Correia and Benny Sudakov. Personal communication. 2022. [3](#)
- [DMO⁺19] Yash Deshpande, Andrea Montanari, Ryan O’Donnell, Tselil Schramm, and Subhabrata Sen. The threshold for SDP-refutation of random regular NAE-3SAT. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2305–2321. SIAM, 2019. [3](#)
- [dT22] Tommaso d’Orsi and Luca Trevisan. A Ihara-Bass Formula for Non-Boolean Matrices and Strong Refutations of Random CSPs. *arXiv preprint arXiv:2204.10881*, 2022. [3](#)
- [Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 534–543, 2002. [2](#), [17](#)
- [Fei08] Uriel Feige. Small linear dependencies for binary vectors of low weight. In *Building bridges*, pages 283–307. Springer, 2008. [0](#), [1](#), [2](#)
- [FKO06] Uriel Feige, Jeong Han Kim, and Eran Ofek. Witnesses for non-satisfiability of dense random 3CNF formulas. In *2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS’06)*, pages 497–508. IEEE, 2006. [2](#)
- [FM17] Zhou Fan and Andrea Montanari. How well do local algorithms solve semidefinite programs? In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 604–614, 2017. [3](#)
- [GKM21] Venkatesan Guruswami, Pravesh K Kothari, and Peter Manohar. Algorithms and Certificates for Boolean CSP Refutation: “Smoothed is no harder than Random”. *arXiv preprint arXiv:2109.04415*, 2021. [0](#), [2](#), [3](#), [4](#), [6](#), [9](#), [10](#), [12](#), [17](#), [18](#), [23](#)

- [Hoo02] Shlomo Hoory. The size of bipartite graphs with a given girth. *Journal of Combinatorial Theory, Series B*, 86(2):215–220, 2002. 1
- [Kee11] Peter Keevash. Hypergraph turan problems. *Surveys in combinatorics*, 392:83–140, 2011. 1
- [KV00] Jeong Han Kim and Van H Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000. 2
- [MOP20] Sidhanth Mohanty, Ryan ODonnell, and Pedro Paredes. The SDP value for random two-eigenvalue CSPs. In *37th International Symposium on Theoretical Aspects of Computer Science*, 2020. 3
- [NV08] Assaf Naor and Jacques Verstraëte. Parity check matrices and product representations of squares. *Combinatorica*, 28(2):163–185, 2008. 1
- [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*, pages 121–131, 2017. 3, 18
- [SS12] Warren Schudy and Maxim Sviridenko. Concentration and moment inequalities for polynomials of independent random variables. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 437–446. SIAM, 2012. 2, 4, 9
- [Tro15] Joel A Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015. 3
- [WAM19] Alexander S Wein, Ahmed El Alaoui, and Cristopher Moore. The Kikuchi hierarchy and tensor PCA. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1446–1468. IEEE, 2019. 3
- [Wik22] Wikipedia contributors. Moore graph — Wikipedia, the free encyclopedia, 2022. [Online; accessed 12-July-2022]. 1

A Alternative proof of the Moore bound for irregular graphs

We proved the weak Moore bound ([Proposition 2.1](#)) by showing that if there is no cycle of length $\leq \ell$, then $A \prec \frac{2n^{1/\ell}}{\sqrt{d}}(D + dI)$ ([Claim 2.2](#)) where D is the diagonal degree matrix and d is the average degree, which then gives us a bound of $2\lceil \log_{d/16} n \rceil$. In this section, we prove that using a more carefully chosen diagonal matrix Γ' , such a strategy can recover the exact Moore bound $2\log_{d-1} n$. This provides an alternative proof of the Moore bound in addition to the existing proofs by [\[AHL02\]](#) and [\[BR14\]](#)⁵.

Theorem A.1 (Moore bound for irregular graphs). *Suppose G is a graph on n vertices with average degree $d > 2$. Then G has a cycle of length $2(\lfloor \log_{d-1} n \rfloor + 1)$.*

⁵[\[AHL02\]](#) and [\[BR14\]](#) actually obtained a slightly more precise bound depending on whether the girth of the graph is odd or even.

The following lemma shows what the “correct” diagonal matrix should be to recover the exact Moore bound.

Lemma A.2. *Let G be a graph with n vertices and degree matrix D that has no cycle of length $\leq \ell$ for some even $\ell \in \mathbb{N}$. Then, the adjacency matrix A satisfies*

$$A \preceq n^{2/\ell} \mathbb{I} + n^{-2/\ell} (D - \mathbb{I}).$$

Proof of Theorem A.1 by Lemma A.2. Assuming there is no cycle of length $\leq \ell$, Lemma A.2 implies that

$$\vec{1}^\top A \vec{1} = nd \leq n \cdot (n^{2/\ell} + n^{-2/\ell} (d - 1)).$$

Let $x = n^{2/\ell}$, then we have $x^2 - dx + (d - 1) \geq 0$, which implies that $x \geq d - 1$ (as $x \leq 1$ is not valid). Taking logs, we get

$$\frac{2}{\ell} \log n \geq \log(d - 1) \implies \ell \leq 2 \log_{d-1} n.$$

ℓ is even, so $\ell < 2(\lfloor \log_{d-1} n \rfloor + 1)$. This completes the proof. \square

The proof of Lemma A.2 is based on *non-backtracking walks*, which are walks such that no edge is the inverse of its preceding edge. We note that both proofs of [AHL02] and [BR14] also analyze non-backtracking walks. For a graph G on n vertices with adjacency matrix A , we define $A^{(s)}$ to be the $n \times n$ matrix whose (u, v) entry counts the number of length- s non-backtracking walks between vertices u and v in G . The following is a standard fact.

Fact A.3 (Recurrence and generating function of $A^{(s)}$). *The non-backtracking matrices $A^{(s)}$ satisfy the following recurrence:*

$$\begin{aligned} A^{(0)} &= \mathbb{I}, \\ A^{(1)} &= A, \\ A^{(2)} &= A^2 - D, \\ A^{(s)} &= A^{(s-1)}A - A^{(s-2)}(D - \mathbb{I}), \quad s > 2. \end{aligned}$$

The recurrences imply that these matrices have a generating function:

$$J(t) := \sum_{s=0}^{\infty} A^{(s)} t^s = (1 - t^2) \cdot H(t)^{-1}, \text{ where } H(t) := \mathbb{I} - At + (D - \mathbb{I})t^2$$

for $t \in [0, 1)$ whenever the series converges.

We first prove the following lemma,

Lemma A.4. *Let $s, k \in \mathbb{N}$, $s \geq k$, and let q, r be the quotient and remainder of s divided by k , i.e. $s = qk + r$. Then,*

$$\text{tr}(A^{(s)}) \leq \sqrt{n} \cdot \|A^{(k)}\|_2^q \cdot \|A^{(r)}\|_F.$$

Proof. $\text{tr}(A^{(s)})$ counts the number of closed non-backtracking walks of length s in the graph. Now, consider the set of closed walks of length $s = qk + r$ such that after every k non-backtracking steps, we can “forget the previous step”, i.e. we are allowed to backtrack at step ik for every $i = 0, \dots, q$.

The number of such walks is $\text{tr}((A^{(k)})^q A^{(r)})$. The set of closed non-backtracking walk is clearly a subset of such walks, thus we have

$$\text{tr}(A^{(s)}) \leq \text{tr}((A^{(k)})^q A^{(r)}) \leq \left\| (A^{(k)})^q \right\|_F \cdot \left\| A^{(r)} \right\|_F.$$

Let $\lambda_1, \dots, \lambda_n$ be the eigenvalues of $A^{(k)}$ and $\lambda_{\max} = \|A^{(k)}\|_2$. Then,

$$\left\| (A^{(k)})^q \right\|_F = \sqrt{\sum_{i=1}^n \lambda_i^{2q}} \leq \sqrt{n} (\lambda_{\max})^q.$$

This completes the proof. \square

With [Fact A.3](#) and [Lemma A.4](#), we now prove [Lemma A.2](#) by analyzing the convergence of $J(t)$ as t increases from 0.

Proof of [Lemma A.2](#). Let A be the adjacency matrix of G with average degree $d > 2$, and let D be the diagonal degree matrix G . Recall the definitions $J(t) = \sum_{s=0}^{\infty} A^{(s)} t^s$ and $H(t) = \mathbb{I} - At + (D - \mathbb{I})t^2$ from [Fact A.3](#). We will analyze the convergence of $\text{tr}(J(t))$ as t increase from 0.

Observe that $J(0) = H(0) = \mathbb{I}$, and since $J(t)$ and $H(t)$ are both symmetric matrices, their eigenvalues move continuously on the real line as t increases from 0. Thus, suppose there is some $t^* \in (0, 1)$ such that $\text{tr}(J(t)) < \infty$ for all $t \in [0, t^*)$, then $H(t) \succ 0$ for all $t \in [0, t^*)$. This is easy to see because if not, then there must be some $t' \in [0, t^*)$ such that $H(t') \succeq 0$ but has a zero eigenvalue, and $\text{tr}(J(t'))$ will not converge.

We next show that we can take $t^* = n^{-2/\ell}$ assuming that G has no cycle of length $\leq \ell = 2k$. First, observe that every entry of $A^{(k)}$ must be either 0 or 1, otherwise if $A^{(k)}[i, j] > 1$ then there are two distinct length- k paths from i to j , meaning there is a cycle of length at most $2k = \ell$, a contradiction. Therefore, the L_1 norm of each row of $A^{(k)}$ is at most n , hence $\|A^{(k)}\|_2 \leq n$. Next, observe that for each $s \in \mathbb{N}$ we can write $s = qk + r$, and

$$J(t) = \sum_{s=0}^{\infty} A^{(s)} t^s \leq \sum_{r=0}^{k-1} \sum_{q=0}^{\infty} A^{(qk+r)} t^{qk+r}.$$

By [Lemma A.4](#), we have

$$\text{tr}(J(t)) \leq \sum_{r=0}^{k-1} t^r \sqrt{n} \|A^{(r)}\|_F \sum_{q=0}^{\infty} \|A^{(k)}\|_2^q \cdot t^{qk} \leq \sum_{r=0}^{k-1} t^r \sqrt{n} \|A^{(r)}\|_F \sum_{q=0}^{\infty} (nt^k)^q.$$

Thus, if $t < n^{-1/k} < 1$, then $\text{tr}(J(t)) < \infty$. Therefore, we have $H(t) \succ 0$ for all $t \in [0, n^{-1/k})$, and by continuity $H(n^{-1/k}) \succeq 0$, which means that

$$\mathbb{I} - n^{-1/k} A + n^{-2/k} (D - \mathbb{I}) \succeq 0 \implies A \preceq n^{2/\ell} \mathbb{I} + n^{-2/\ell} (D - \mathbb{I})$$

as $\ell = 2k$. This completes the proof. \square