

Improved Lower Bounds for all Odd-Query Locally Decodable Codes

Arpon Basu

ab5541@princeton.edu
Princeton University

Jun-Ting Hsieh*

juntingh@cs.cmu.edu
Carnegie Mellon University

Pravesh K. Kothari[†]

kothari@cs.princeton.edu
Princeton University

Andrew D. Lin

andrewlin@princeton.edu
Princeton University

November 21, 2024

Abstract

We prove that for every odd $q \geq 3$, any q -query binary, possibly non-linear locally decodable code (q -LDC) $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ must satisfy $k \leq \tilde{O}(n^{1-2/q})$. For even q , this bound was established in a sequence of works [KT00, GKST06, KW04]. For $q = 3$, the above bound was achieved in a recent work [AGKM23] using an argument that crucially exploits known exponential lower bounds for 2-LDCs. Their strategy hits an inherent bottleneck for $q \geq 5$.

Our key insight is identifying a general sufficient condition on the hypergraph of local decoding sets called *t*-approximate strong regularity. This condition demands that 1) the number of hyperedges containing any given subset of vertices of size t (i.e., its *co-degree*) be equal to the same but arbitrary value d_t up to a multiplicative constant slack, and 2) all other co-degrees be upper-bounded *relative* to d_t . This condition significantly generalizes related proposals in prior works [GKM22, HKM23, AGKM23, HKM⁺24] that demand *absolute* upper bounds on all co-degrees.

We give an argument based on spectral bounds on *Kikuchi Matrices* that lower bounds the blocklength of any LDC whose local decoding sets satisfy *t*-approximate strong regularity for any $t \leq q$. Crucially, unlike prior works, our argument works despite having no non-trivial absolute upper bound on the co-degrees of any set of vertices. To apply our argument to arbitrary q -LDCs, we give a new, greedy, approximate strong regularity decomposition that shows that arbitrary, dense enough hypergraphs can be partitioned (up to a small error) into approximately strongly regular pieces satisfying the required relative bounds on the co-degrees.

*Supported by NSF CAREER Award #2047933.

[†]Supported by NSF CAREER Award #2047933, Alfred P. Sloan Fellowship and a Google Research Scholar Award.

Contents

1	Introduction	1
1.1	Brief Overview	2
2	Warmup: A Combinatorial Proof for the Linear Case	4
2.1	LDCs and Even Covers	4
2.2	The Kikuchi Graph	6
2.3	Approximate Strong Regularity	7
2.4	Avoiding clustering of edges in the Kikuchi graph	9
2.5	Proof of Lemma 2.8	10
2.6	Regularity Decomposition	12
2.7	Non-linear LDC Lower Bounds	13
3	Preliminaries	15
3.1	Locally Decodable Codes	15
3.2	Concentration Inequalities	16
4	Proof of the Main Theorem	16
4.1	Proof of Lemma 4.6 : refuting regular q -XOR	19
4.2	Proof of Lemma 4.15 : bounding heavy vertices	24
5	Approximate Strong Regularity Decomposition	27
5.1	Proof of Lemma 5.1 : a good index exists	30
	Acknowledgments	31
	References	31

1 Introduction

A binary error-correcting code $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ is *locally decodable* if it admits a simple “local” decoding algorithm. Given a corrupted received word $y \in \{\pm 1\}^n$ and an index $1 \leq i \leq k$ of the message, the (randomized) algorithm reads at most q locations of y (where q is the *query complexity*) and is required to output the i -th bit of the message correctly with probability $> 1/2 + \Omega(1)$. The study of locally decodable codes (LDCs) goes back to the early days of algorithmic coding theory. Starting with the proof of the PCP theorem [AS98, ALM⁺98], their applications form a diverse list that includes worst-case to average-case reductions [Tre04], private information retrieval [Yek10], secure computation [IK04], derandomization [DS05], matrix rigidity [Dvi10], fault-tolerant computation [Rom06], and even data structure lower bounds [Wol09, CGW10].

The central question in the study of LDCs is the optimal trade-off between the *rate*, i.e., the smallest possible blocklength n as a function of k , and the query complexity q , assuming that the (normalized) distance of the code is a fixed constant $\delta \in (0, 1)$. For $q = 2$ (the smallest interesting query complexity), the binary Hadamard code achieves $n = 2^{O(k)}$. This exponential blow-up was eventually shown to be necessary even for non-linear codes over all finite alphabets [KT00, GKST06, KW04].¹

For $q \geq 3$, our knowledge is starkly limited. Over 15 years ago, a breakthrough result [Yek08, Efr09] showed that there are q -query LDCs (q -LDCs, henceforth) with a sub-exponential blocklength of $n \leq \exp(2^{\tilde{O}(\log^{1/(q-1)} k)}) \leq 2^{k^{o(1)}}$. On the other hand, a sequence of results [GKST06, KW04], beginning with the pioneering work of Katz and Trevisan [KT00], showed that $k \leq \tilde{O}(n^{1-2/q})$ for any *even* integer q . For odd q , since q -LDCs are $(q+1)$ -LDCs, we can invoke the above inequality to yield a weaker bound of $k \leq \tilde{O}(n^{1-2/(q+1)})$. In the following two decades, these bounds were state-of-the-art up to polylogarithmic factor improvements [Woo07, Woo08, Woo12, BCG20].

A recent work [AGKM23] introduced a new connection between q -LDCs and the problem of semirandom q -CSP refutation. By relying on spectral methods based on *Kikuchi Matrices* [WAM19] for semirandom CSP refutation [GKM22, HKM23], they obtained the first polynomial improvement on the above bounds for the case of $q = 3$. In particular, their bound of $k \leq \tilde{O}(n^{1/3})$ matches the expression $\tilde{O}(n^{1-2/q})$ proved in [KW04] for even q . The argument of [AGKM23] was later simplified in [HKM⁺24] to obtain a purely combinatorial proof for linear 3-LDCs that also removes all but one logarithmic factor. This framework eventually led to new progress giving exponential lower bounds on the blocklength of 3-query locally *correctable* codes [KM23, KM24, Yan24, AG24].

Beyond $q = 3$? Despite the quantitative improvements for $q = 3$, the ideas developed in [AGKM23, HKM⁺24] fail to give any improvement for any odd $q > 3$. Indeed, their arguments strongly exploit properties specific to $q = 3$, as we next explain. At a high level, the proof involves a decomposition step that starts from the query sets of a local decoder for a purported 3-LDC (with $k \gg n^{1/3}$) and outputs either (1) a 3-LDC where the query sets satisfy an additional pseudo-random property, or (2) an object that is essentially a 2-LDC with (a significantly reduced) message length of $\sim O(\log n)$. One must obtain a contradiction in either of the two cases to show that 3-LDCs with $k \gg n^{1/3}$ do

¹ For 2-query codes on alphabet sizes that grow with n , known lower bounds get progressively weaker [WW05]. For large enough alphabet sizes, this is inevitable as the problem is closely related to the 2-server private information retrieval (PIR) that does admit significant improvements on the parameters of Hadamard codes over large fields [DG16].

not exist. One can only get the 2-LDC-like object to have message length $\sim \log n$, but this suffices in their setting since 2-LDCs are known to require an exponential blocklength [GKST06, KW04]. The analogous plan for 5-LDCs would require an exponential lower bound on q -LDCs for all $q \leq 4$. But for $q = 4$, such a bound is not only far from the best-known quadratic lower bound but, given the Efremenko-Yekhanin codes, also *false*. This inherent bottleneck has so far prevented any improvement in the lower bounds for odd $q \geq 5$. Indeed, upgrading the lower bounds for $q \geq 5$ was explicitly identified as an outstanding open question in prior work [AGKM23].

In this work, we resolve this question and prove:

Theorem 1.1. *For any q -LDC $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ with a constant distance, $k \leq O(n^{1-2/q} \log^4 n)$. If E is linear, then, $k \leq O(n^{1-2/q} \log^2 n)$.*

While we omit a formal argument here, as in [AGKM23], our methods naturally extend to give a similar lower bound for LDCs over any finite (even $n^{o(1)}$) size alphabet. We note that in a concurrent and independent work, Janzer and Manohar [JM24] prove a result that obtains a similar bound as in Theorem 1.1 but with a single logarithmic factor.

1.1 Brief Overview

Let us explain our main ideas briefly (see Section 2 for a guided exposition with a complete proof for linear codes). After standard reductions (see Fact 3.4), the local decoder for a q -LDC can be described by a collection of q -uniform matchings $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$, one for every bit $1 \leq i \leq k$ of the message. To decode the i -th bit in a received corrupted codeword y , the decoder queries a uniformly random q -tuple in \mathcal{H}_i .

The idea from [AGKM23, HKM⁺24]: The key technical idea in [AGKM23] (and a combinatorial analog in [HKM⁺24]) is identifying a certain *well-spread* property of the hypergraph of the query sets $\bigcup_{i \in [k]} \mathcal{H}_i$ that allows spectral methods based on *Kikuchi Matrices* to succeed in proving that $k \lesssim n^{1-2/q}$. For any $q \geq 3$, their well-spread condition requires that no pair of $\{a, b\}$ of the codeword coordinates are queried in more than $O(\log n)$ different matchings \mathcal{H}_i . Briefly speaking, the well-spread property is used to argue that certain “Kikuchi graphs” obtained from the \mathcal{H}_i s are approximately regular, i.e., after dropping a negligible fraction of outlier vertices, all the remaining vertices have their degrees within a polylog n factor of the average.

The well-spread property (also called being a *design* in [BDYW11]) is a natural pseudo-random condition and is satisfied, e.g., by a collection of uniformly random hypergraph matchings with high probability. Its use in [AGKM23] was directly motivated by works on semirandom CSP refutation and the *hypergraph Moore bound* [GKM22, HKM23], where it leads to near-optimal results. Indeed, the first step in these works is to find a simple deterministic property of random hypergraphs that suffices for the success of spectral methods for “refuting” (i.e., certifying unsatisfiability) random CSP instances.

Given an arbitrary 3-LDC (i.e., arbitrary collection of matchings \mathcal{H}_i), one must now reduce the problem to the case where the local decoding sets satisfy the well-spread condition. This is easily ensured in their setting through a greedy decomposition: we iteratively peel off pairs that

occur $\gg O(\log n)$ times (call such pairs “heavy”). After the peeling-off process ends, the remaining elements in the matching must satisfy the well-spread condition by design. If the matchings retain a constant fraction of the constraints we began with, we can apply the argument for the well-spread case referenced above. If not, we look at all the query sets peeled off. By extending the original code to have coordinates for the peeled off “heavy pairs”, one can think of the query sets as describing an object that behaves essentially as a 2-LDC with message length $\lesssim \log n$ (this bound corresponds to the number of times a pair was queried in the original $\cup_i \mathcal{H}_i$). Given the known exponential lower bounds for 2-LDCs, this is enough for a contradiction.

The well-spread condition, when formalized for an arbitrary odd $q \geq 5$ requires control of co-degrees of all sets of size r for $r \leq 5$. Indeed, such conditions were formulated as notions of regularity in prior works [GKM22, HKM23]. Specifically, these works identified appropriate absolute thresholds $\tau_1, \tau_2, \dots, \tau_t$ in terms of the number of vertices n and related parameters, and posited that the co-degree of any size- r set Q (i.e., the number of hyperedges in \mathcal{H} containing Q) be upper bounded by τ_r . The precise thresholds are chosen to ensure the success of a certain spectral bound on Kikuchi matrices that forms the technical core of their argument.

For 5-LDCs, however, such an argument fails immediately since the threshold τ_3 for triples still turns out to be $\sim \log n$, and if we were to remove $\gg O(\log n)$ -heavy triples, we get an object that behaves essentially as a 4-LDC (instead of a 2-LDC) but still with message length $\sim \log n$. Unlike 2-LDCs, now there is no contradiction since there are 4-LDCs [Efr09, Yek08] with $\gg \log^{\omega(1)} n$ message length (and moreover, the best known lower bounds for 4-LDCs are just quadratic $k \lesssim n^{1/2}$).

Approximate strong regularity: Our key conceptual idea is to depart from the philosophy of reducing to well-spread $\mathcal{H} = \cup_i \mathcal{H}_i$ s. Instead, we prove that a significant relaxation that we call *approximate strong regularity* suffices for our spectral argument to succeed.

Approximate strong regularity does not demand any absolute upper bounds on the co-degrees of subsets of vertices in \mathcal{H} . Instead, given a $t \leq q$, we say that a \mathcal{H} is t -approximately strongly regular if 1) all sets of size t have roughly the same co-degree, say d_t , up to a fixed constant multiplicative slack, and 2) the co-degrees of every set of size $1 \leq r \leq q$ is appropriately bounded *relative to* d_t . Crucially, our relative bounds do not imply any non-trivial absolute upper bounds on the co-degrees. Thus, the co-degree of triples may be $\text{poly}(n)$ (as opposed to the $\sim \text{polylog } n$ tolerated in prior arguments).

We show that if $\mathcal{H} = \cup_i \mathcal{H}_i$ is t -approximately strongly regular, then certain Kikuchi graphs built from the \mathcal{H}_i s that arise in our spectral argument satisfy approximate regularity (i.e., dropping a few outlier vertices makes all degrees roughly equal). Our relative thresholds (see Definition 2.6) are chosen to make a key technical piece in the argument (where we prove that certain Kikuchi graphs obtained from the hypergraph \mathcal{H} are approximately regular) go through (and may not look motivated in first reading). As a result, we can argue that spectral methods succeed in establishing $k \lesssim n^{1-2/q}$ given approximate strong regularity of the underlying $\mathcal{H} = \cup_i \mathcal{H}_i$. We stress that this result holds regardless of how large d_r s themselves are and requires no pruning of heavy tuples, which is the key bottleneck in the prior approaches. We note that arguing the approximate regularity is the key technical component in all prior applications of the Kikuchi matrix method [GKM22, HKM⁺24, HKM23, AGKM23, KM23, Yan24, KM24].

Our second main technical component argues that any hypergraph can be decomposed (up to a negligible fraction of “error” hyperedges) into pieces that satisfy the required relative upper bounds on all co-degrees with respect to d_t for some t .

Outlook: Hypergraph decompositions to ensure approximate regularity of Kikuchi graphs have been a key tool in several recent results [GKM22, HKM23, AGKM23, HKM⁺24, KM23, KM24, Yan24]. In all these prior applications, regularity of Kikuchi graphs was ensured by forcing absolute upper bounds on the co-degrees of all subsets in the underlying hypergraphs. The key difficulty for improved lower bounds for q -LDCs for odd $q \geq 5$ was precisely the inapplicability of this natural idea. We thus expect that the notion of approximate strong regularity and the fact one can make this property hold without loss of generality by hypergraph decomposition will find more applications, especially in the context of the Kikuchi matrix method.

2 Warmup: A Combinatorial Proof for the Linear Case

In this section, we give a proof for the case of *linear* q -LDCs. The argument itself is short, but we provide additional commentary to explain our key ideas. Our proof exploits connections between LDCs and *even covers* in edge-colored hypergraphs. We will briefly discuss our approach for the non-linear case (i.e., [Theorem 1.1](#)) in [Section 2.7](#).

Theorem 2.1. *Let $q \geq 3$ be odd and let $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ be a linear q -LDC that corrects up to a δ fraction of errors. Then, $k \leq O(1/\delta) \cdot n^{1-2/q} \log^2 n$.*

In [Sections 2.1](#) and [2.2](#), we introduce standard notions of even covers and the Kikuchi graph. In [Section 2.3](#), we state our new notion of hypergraph regularity (called “approximate strong regularity”). This notion of regularity is a key technical idea in our work and departs from the usual *well-spread* conditions on the hypergraphs appearing in prior related works (see [Definition 2.7](#) for more details). In [Section 2.4](#), we give some explanation of the conditions, and in [Section 2.5](#) we prove a key result ([Lemma 2.8](#)) on the existence of “odd-colored” even covers when the hypergraph satisfies approximate strong regularity. In [Section 2.6](#), we discuss our decomposition algorithm that ensures approximate strong regularity for each piece except those containing a negligible fraction of the original hyperedges. At the end of [Section 2.6](#), we complete the proof of [Theorem 2.1](#).

2.1 LDCs and Even Covers

By standard reductions (see [Fact 3.4](#)), we can assume that the q -LDC is in the normal form. Thus, there are q -uniform matchings of size $\Omega(n)$, say $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$, that correspond to the local decoding sets for each of the message bits $1 \leq i \leq k$. Further, the decoder itself is linear. Thus, if $x = E(b)$ for a q -LDC $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$, then for every $C \in \mathcal{H}_i$, it must hold that $\prod_{u \in C} x_u = b_i$.

An *even cover* in a hypergraph is a collection of hyperedges that use every vertex an even number of times — a natural linear algebraic generalization of cycles in graphs. When viewing a hyperedge as the coefficient vector of a linear form over \mathbb{F}_2 , an even cover is simply a subset of equations that add up to 0 in \mathbb{F}_2 .

We begin with a simple connection between linear LDCs and even covers in edge-colored hypergraphs that has been standard in recent works on LDCs [AGKM23, HKM⁺24].

Fact 2.2 (Lemma 2.7 of [HKM⁺24]). *Let $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ be a binary, linear q -query LDC associated with q -uniform matchings $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$. Suppose we color each hyperedge in \mathcal{H}_i with color i . Then, any even cover in $\mathcal{H} = \bigcup_{i \in [k]} \mathcal{H}_i$ must use color i an even number of times for each $i \in [k]$.*

This fact itself is simple to prove. Suppose an even cover uses a subset $S_i \subseteq \mathcal{H}_i$ of hyperedges from the i -th matching, and suppose $|S_i|$ is odd. Then, for any message vector $b \in \{\pm 1\}^k$ and the corresponding codeword $x = E(b)$, it must hold that $x_C = \prod_{u \in C} x_u = b_i$ for every $C \in S_i$. If we take product of the equations corresponding to edges in the even cover, then the LHS must square out to 1 while the RHS is the product of b_i s for i such that S_i has odd size. This is a contradiction if we set such message bit to -1 .

LDC Lower Bounds and Weakly Rainbow Even Covers: Thus, to prove that $k \ll n^{1-2/q}$ in any linear q -LDC, it suffices to show that if $k \gg n^{1-2/q}$ (i.e., $\mathcal{H} = \bigcup_i \mathcal{H}_i$ is a sufficiently dense hypergraph) then there is an even cover in \mathcal{H} that uses some color an *odd* number of times.

This goal has a nice combinatorial interpretation. Notice that \mathcal{H} is a *properly hyperedge-colored* hypergraph — since the hyperedges incident on any single vertex are all of distinct colors as the \mathcal{H}_i s are matchings. Thus, our goal is to show that a properly edge-colored hypergraph that is dense enough must have an even cover that uses some color an odd number of times. In fact, we will prove the stronger statement that such a hypergraph must contain an even cover that uses exactly one hyperedge of some color. Let us make this into a definition before we move on:

Definition 2.3 (Weak Rainbow Even Covers). *A **weakly rainbow** even cover in a properly hyperedge-colored hypergraph \mathcal{H} is an even cover that uses exactly one hyperedge of some color.*

Showing the existence of weakly rainbow even covers in dense enough hypergraphs is a variant of the well-studied rainbow cycle problem in graphs. Such problems were introduced in the pioneering work of [KMSV07]. Their work included, among many results, a *weak rainbow cycle* theorem for graphs: every properly edge-colored graph on n vertices with average degree $\geq O(\log n)$ has a cycle that uses some color exactly once². There is a long line of work building on their work including the ones leading up to the recent (almost) resolution of the related *rainbow cycle conjecture* [DLS13, Jan20, Tom24, KLLT24, JS24, ABS⁺23]. Here, let us state a generalization that we will apply as a black-box:

Fact 2.4 (Lemma 2.5 of [HKM⁺24]). *Let G be an n -vertex graph where each edge in G is assigned a set of s colors, and suppose G has average degree $d \geq 40s \log n$. Suppose that for every vertex $v \in G$ and color c , the number of edges incident on v whose assigned set of colors contains c is at most $d/20s \log n$. Then, G contains a closed walk, of size at most $2 \log n$, such that some color appears exactly once.*

Our main result for linear LDCs is a natural hypergraph analog of the weak rainbow cycle theorem for graphs (see [Lemma 2.8](#) for the key technical component). We will prove this by a

² We call this version “weak” to distinguish from the better-known rainbow cycle conjecture asks for a cycle where every color is used at most once.

reduction to edge-colored graphs and applying [Fact 2.4](#). This reduction is based on Kikuchi graphs that we next introduce.

2.2 The Kikuchi Graph

As in the prior works, a key technical idea is the use of Kikuchi graphs. In the even q case, for $\ell \leq n$, the level- ℓ Kikuchi graph associated with an edge-colored hypergraph \mathcal{H} is an edge-colored graph defined as follows:

- The vertex set is $\binom{[n]}{\ell}$, i.e., all ℓ -sized subsets of $[n]$.
- For each hyperedge $C \in \mathcal{H}$ with color i , we add edges (S, T) , denoted by $S \xleftrightarrow{C} T$, such that $S \oplus T = C$ and $|S \cap C| = |T \cap C| = q/2$, and color them with i .

The key to the utility of these graphs is an elementary but important connection [[GKM22](#)] between cycles/walks in the Kikuchi graph and even covers in \mathcal{H} . As a result, finding a weak rainbow even cover in \mathcal{H} reduces to finding a weak rainbow cycle in the edge-colored Kikuchi graph. For even q , this approach easily recovers the best known lower bounds for q -LDCs [[GKST06](#), [KW04](#)].

For odd q , the challenge with this approach is that the Kikuchi graph above does not make sense (for $S \oplus T = C$ and $|S| = |T| = \ell$, C must be of even size). To get around this issue, all prior works [[GKM22](#), [HKM23](#), [AGKM23](#), [HKM⁺24](#)] use variants of Kikuchi graphs where each transition corresponds to a pair (C, C') of hyperedges (instead of a single C for the even case above).

Let us informally describe the Kikuchi graph we use (see [Definition 4.11](#) and the illustration [Figure 2](#)). For a parameter $t \in \{1, 2, \dots, q-1\}$ (that will eventually be chosen carefully), our Kikuchi graph has transitions corresponding to pairs (C, C') that intersect in t vertices. In this case, note that $C \oplus C'$ is a set of size $2(q-t)$, which is even. Our Kikuchi graph takes all such pairs in \mathcal{H} and includes edges $S \xleftrightarrow{C, C'} T$.³

Since \mathcal{H} is properly hyperedge-colored, C and C' must have different colors $i \neq j$. Thus, each edge of the Kikuchi graph is associated with 2 distinct colors. We note that in [[AGKM23](#)], the t above is hardcoded to 1; in our case we must carefully choose t as a function of the hypergraph.

Weak rainbow even covers by applying [Fact 2.4](#) to Kikuchi graphs: Our Kikuchi graph naturally has edges that have a pair of colors. We say that a graph (where each edge can have multiple colors) is Δ -properly edge-colored if, for each color i , every vertex is incident to at most Δ edges whose color set includes the color i .

We want to apply [Fact 2.4](#) to the Kikuchi graph with $s = 2$. For this, we need to argue that the Kikuchi graph contains a large subgraph that

- (1) has average degree $\bar{d}(K) \geq O(\log N)$,
- (2) is $\bar{d}(K)/O(\log N)$ -properly edge-colored — that is, given any color i , the number of edges with a pair of colors including i incident on any vertex is at most $\bar{d}(K)/O(\log N)$.

³ For the arguments that follow, we need to setup this graph so that S and T each intersect both C, C' in equal, up to a slack of at most 1 element, but we omit a discussion of this issue here.

The average degree of the Kikuchi graph is easily described as a function of ℓ and $|\mathcal{H}|$ (see [Fact 4.12](#)), and setting $\ell = n^{1-2/q}$ will guarantee $\bar{d}(K) \geq O(\log N)$. The sticking point is ensuring that the Kikuchi graph is $\bar{d}(K)/O(\log N)$ -properly edge-colored.

Remark 2.5 (Prior works). Ensuring the proper edge-coloring property (which is essentially equivalent to approximate regularity of the Kikuchi graphs) is in fact the key technical step in all applications of the Kikuchi matrix method so far [[GKM22](#), [HKM23](#), [HKM⁺24](#), [AGKM23](#), [KM23](#), [KM24](#)]. In all these prior works, such a property was ensured when the underlying hypergraph satisfies a natural well-spread condition. This well-spread condition in the work on 3-LDCs [[AGKM23](#)] asks that in \mathcal{H} , no pair of vertices appear in more than $O(\log n)$ hyperedges. The remaining part of the argument is then giving a different argument when \mathcal{H} does not satisfy this well-spread property. In all prior applications, this is done by constructing a hypergraph of lower uniformity. For the case of 3-LDCs, this gives a graph and the $O(\log n)$ threshold for pairs being “heavy” above translates into its average degree being $\gg O(\log n)$. This average degree is sufficient to get cycles in such a graph that use some color exactly once. For the case of 5-LDCs, however, the same reduction produces a hypergraph of uniformity 4 with average degree $O(\log n)$. Such a hypergraph can manifestly only have even covers that use every color even number of times (for e.g., those coming from the Efremenko-Yekhanin matching vector codes [[Yek08](#), [Efr09](#)]!).

Our main idea to overcome this obstacle is finding a significantly more general condition — *approximate strong regularity* — on hypergraphs such that the associated Kikuchi graphs still possess the Δ -properly edge-colored property. Our more general condition can violate the well-spread property of hypergraphs drastically, and pairs, triples, quadruples, etc. can appear in an arbitrarily large number of hyperedges.

2.3 Approximate Strong Regularity

Given a hypergraph \mathcal{H} , the co-degree of a set Q of vertices, which we denote as $d_{\mathcal{H},Q}$, is the number of hyperedges in \mathcal{H} that contain Q . Moreover, we denote $d_{\mathcal{H},t} := \max_{|Q|=t} d_{\mathcal{H},Q}$ for all $t \in [q]$. We will omit the dependence on \mathcal{H} and write d_t for simplicity. Note that we have $d_1 \geq d_2 \geq \dots \geq d_q$ and $d_1 \geq |\mathcal{H}|/n$.

Our key technical piece shows the existence of a weakly rainbow even cover if \mathcal{H} satisfies a property called approximate strong regularity. This property demands that there be some $t \leq q$ such that 1) \mathcal{H} be partitioned into groups of roughly the same size d_t such that hyperedges in each group all intersect in a fixed set of size t , and, 2) the co-degrees of any subset of r vertices (for $1 \leq r \leq q$ is small *relative to* d_t). This demand for relative as opposed to absolute upper bounds is a key departure from similar notions of regularity appearing in prior works (and their insufficiency is a key reason for their inapplicability in proving improved lower bounds for odd $q \geq 5$).

The precise relative bounds below can look daunting. The choice is dictated by a key technical piece in our proof, where we show the approximate regularity of Kikuchi graphs that arise in our analysis. We note that arguing the approximate regularity of Kikuchi graphs is the centerpiece in all applications of the Kikuchi matrix method so far [[GKM22](#), [HKM23](#), [HKM⁺24](#), [AGKM23](#), [KM23](#),

Yan24, KM24]). We invite the reader to ignore the quantitative requirements in the first reading of this overview — we explain their origin in the next subsection (Section 2.4).

Definition 2.6 (Good index t). Let $d_1 \geq d_2 \geq \dots \geq d_q$. We say that an index t is *good* with respect to the tuple (d_1, d_2, \dots, d_q) if the following conditions hold:

- (1) $d_r/d_t \leq n^{1-\frac{2r}{q}}$ for every $1 \leq r \leq \lceil \frac{q-t}{2} \rceil$,
- (2) $d_r/d_t \leq n^{-\frac{2}{q}(r-t)+\frac{1}{q}(t-1(t \text{ even}))}$ for every $t \leq r \leq \lfloor \frac{q+t}{2} \rfloor$,
- (3) If $t < q/2$, then, $d_t \geq d_1 n^{-\frac{2}{q}(t-1)}$, and, if $t > q/2$, then, $d_t \geq d_1 n^{-1+2/q}$.

We now define t -approximate strong regularity of a q -uniform hypergraph \mathcal{H} that posits the above conditions are satisfied for some $1 \leq t \leq q$.

Definition 2.7 (Approximate Strong Regularity). For a given $t \in [q]$, a q -uniform hypergraph \mathcal{H} on $[n]$ is called t -approximately strongly regular if there exists a partitioning $\mathcal{H} = \mathcal{H}_1^{(t)} \sqcup \dots \sqcup \mathcal{H}_{p_t}^{(t)}$ such that:

- (1) For every $\theta \in [p_t]$, there is some $Q_\theta \in \binom{[n]}{t}$ such that all $C \in \mathcal{H}_\theta^{(t)}$ contain Q_θ .
- (2) $d_t/2 \leq |\mathcal{H}_\theta^{(t)}| \leq d_t$ for all $\theta \in [p_t]$.
- (3) t is a good index with respect to the sequence $d_1 \geq \dots \geq d_q$ (recall that $d_r := \max_{|Q|=r} d_{\mathcal{H},Q}$).

Let us interpret this definition. First, the numbers d_r give an upper bound on the co-degree of every size- r subset of vertices in \mathcal{H} . Thus, the first 2 conditions imply that the hypergraph can be partitioned into pieces, each containing a $(d_t/2)$ -heavy subset of t vertices. Condition (3) states that all other co-degrees are bounded relative to d_t , and moreover, item (3) of Definition 2.6 gives a lower bound on d_t in terms of d_1 , which we know is at least $|\mathcal{H}|/n$.

We can now state our weak hypergraph rainbow lemma for approximately strongly regular hypergraphs.

Lemma 2.8 (Weak Rainbow Bound for Approximately Strongly Regular Hypergraphs). *Let $q \in \mathbb{N}$ be odd and $t \in [q]$. There is a universal constant A depending only on q such that the following holds: Let \mathcal{H} be a properly hyperedge-colored, q -uniform, t -approximately strongly regular hypergraph on $[n]$ with average degree $k \geq An^{1-\frac{2}{q}} \log n$.*

Then, \mathcal{H} contains a weak rainbow even cover.

Before we prove Lemma 2.8, let us first explain the conditions of a good index in Definition 2.6. One can also see these conditions arise directly in the proof of Lemma 2.8 in Section 2.5. Later in Section 2.6, we will show that approximate strong regularity can be assumed without loss of generality by decomposing the given hypergraph (though losing a log factor).

2.4 Avoiding clustering of edges in the Kikuchi graph

In this section, we provide a high-level explanation (with a concrete example) showing why the conditions for a good index in [Definition 2.6](#) imply that the Kikuchi graph satisfies the proper coloring condition required in [Fact 2.4](#).

The Kikuchi graph K at level ℓ we described for \mathcal{H} is obtained by picking an integer $t \leq q$ and then constructing edges corresponding to pairs (C, C') such that C, C' intersect in a set of size t . The choice of t needs to be carefully done to ensure the $\bar{d}(K)/O(\log N)$ -proper coloring property of the resulting Kikuchi graph. Recall that to apply [Fact 2.4](#), we need 1) the average degree of the Kikuchi graph to be $\bar{d}(K) \geq O(\log N)$, and 2) the graph is $\sim \bar{d}(K)/\log N$ -properly edge-colored. Let's set $t = 1$ for the sake of illustration and first compute the average degree $\bar{d}(K)$ of the Kikuchi graph. Since $d_1 \sim k \sim n^{1-2/q}$ (each matching is near perfect), the number of pairs (C, C') in \mathcal{H} that intersect in one vertex is $\sim nk^2$. Each such pair then has $2(q-1)$ non-overlapping vertices. A random vertex S of the Kikuchi graph behaves essentially like a ℓ/n -biased random set, and thus, a given (C, C') generates an edge incident on it with probability $(\ell/n)^{q-1}$. Thus, $\bar{d}(K) \sim (\ell/n)^{q-1}nk^2 \geq \log N \sim \ell \log n$ if $\ell = n^{1-2/q}$ and $k \geq \tilde{O}(n^{1-2/q})$.

Now fix a color i . We must argue that the number of pairs (C, C') that induce an edge on a vertex S of the Kikuchi graph K is⁴ at most $\bar{d}(K)/O(\log N) \sim \text{polylog } n$. It suffices to argue that this holds for almost all vertices since we can then delete a negligible fraction of edges to ensure it for all without noticeably hurting the average degree.

We next show a concrete scenario where the above property fails (here, due to the co-degree d_3 of triples being too large). Let us set $q = 5$ for this illustration since the issue already pops up for this first interesting case. For every $C \in \mathcal{H}_i$, we can have $\sim k$ different C'_1, C'_2, \dots , hyperedges in \mathcal{H} that intersect C in exactly one vertex, say u . All such pairs (C, C'_j) contribute edges in our Kikuchi graph with a color pair containing i . Suppose that there are d_3 different C'_j 's all containing a fixed pair $\{v, w\}$. In this case, d_3 different C'_j 's contain the same triple $\{u, v, w\}$. See [Figure 1](#) for an illustration.

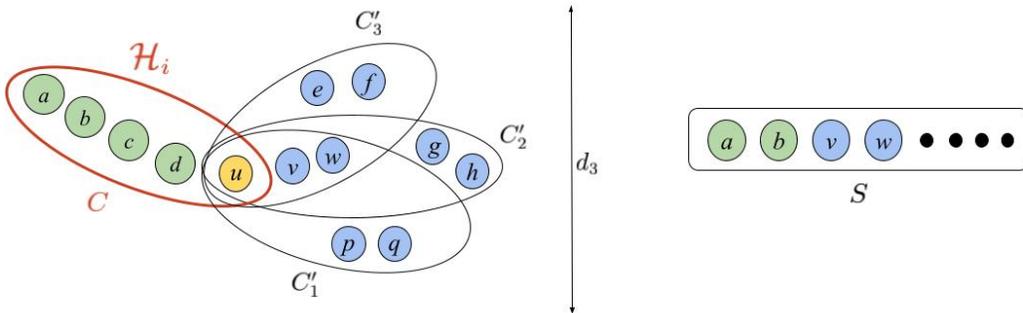


Figure 1: This example shows, for $q = 5$, a scenario where sets S as shown have $\geq d_3$ edges where one of the colors is i .

⁴ We note that when $t > 1$, the degree of the Kikuchi graph will in general be $\gg O(\log N)$ and depend on d_t . In that case, the number of edges incident on a typical S with a fixed color i that we can tolerate will scale relative to d_t .

A pair (C, C') contributes an edge on S if S intersects C and C' in $(q-1)/2 = 2$ elements each for $q = 5$. S already contains 2 vertices from C and by containing the pair $\{v, w\}$ two vertices from every one of the d_3 different C'_j 's. Thus, S has an edge corresponding to every one of the (C, C'_j) where C'_j contains $\{v, w\}$ and thus at least d_3 edges with color i . If $d_3 \gg \text{polylog } n$ (it can be $n^{\Omega(1)}$ in general!), we have failed to satisfy the proper coloring requirement.

One can generalize the above bad example into a setting where setting $t = 1$ in our Kikuchi graph simply fails to satisfy the proper coloring condition required in [Fact 2.4](#) for a constant fraction of vertices. More generally, for any given t , the above analysis necessitates that the ratio d_r/d_t be appropriately upper bounded as a function of r and t . Such restrictions give rise to the specific conditions in [Definition 2.6](#), where each condition corresponds to a bad scenario analogous to the example above. It turns out that we can argue (done in full in the next subsection!) that the list of bad scenarios the above conditions rule out is exhaustive. In other words, if every one of the conditions above is met, then, we can establish a $\bar{d}/O(\log N)$ upper bound on the number of edges of a fixed color i (out of a pair of colors) incident on almost all vertices of the Kikuchi graph.

2.5 Proof of [Lemma 2.8](#)

Let us now give a full proof of [Lemma 2.8](#).

Proof of [Lemma 2.8](#). Recall that $d_1 \geq d_2 \geq \dots \geq d_q$ denote the co-degrees of \mathcal{H} , and t is a good index as in [Definition 2.6](#). We first claim that

$$d_t \geq \max\left\{n^{1-\frac{2t}{q}}, 1\right\} \cdot \Omega(\log n). \quad (1)$$

To see this, note that by item (3) of [Definition 2.6](#), we have $d_t \geq d_1 \cdot \max\{n^{-\frac{2}{q}(t-1)}, n^{-1+\frac{2}{q}}\}$. Moreover, d_1 is at least the average degree, which is assumed to be $\geq \Omega(n^{1-2/q} \log n)$. This establishes [Eq. \(1\)](#).

Next, let $\ell := n^{1-\frac{2}{q}}$, and consider the level- ℓ Kikuchi graph K as described in [Section 2.2](#) by canceling t -tuples. By t -approximate strong regularity ([Definition 2.7](#)), \mathcal{H} has a partition $\{\mathcal{H}_\theta^{(t)}\}_{\theta \in [p_t]}$, each of size at least $d_t/2$. Thus, we can form $\sum_{\theta=1}^{p_t} \binom{|\mathcal{H}_\theta^{(t)}|}{2} \geq \Omega(d_t) \sum_{\theta=1}^{p_t} |\mathcal{H}_\theta^{(t)}| \geq \Omega(knd_t)$ pairs of hyperedges. Thus, the average degree of the Kikuchi graph K is

$$\bar{d}(K) \geq \Omega\left(\frac{\ell}{n}\right)^{q-t} \cdot knd_t \geq \Omega(d_t) \cdot n^{\frac{2}{q}(t-1)} \log n.$$

In particular, by [Eq. \(1\)](#), we have $\bar{d}(K) \geq \Omega(n^{1-\frac{2}{q}} \log^2 n) \geq \omega(1) \cdot \log |V(K)|$, since $|V(K)| = \binom{2n}{\ell} \leq (2n)^\ell$. This establishes the average degree lower bound required in [Fact 2.4](#).

Next, define

$$\Delta := \frac{\bar{d}(K)}{100\ell \log n} \geq A' d_t \cdot n^{\frac{2t}{q}-1}, \quad (2)$$

where A' is some large constant depending on q . It suffices to show that there is a subgraph of K which is Δ -properly edge-colored and has average degree $\bar{d}(K)/2$. This would complete the proof

since [Fact 2.4](#) guarantees the existence of a cycle, of length at most $2 \log |V(K)| = O(\ell \log n)$, that uses a color exactly once.

Fix any $C_1 \neq C_2 \in \mathcal{H}_\theta^{(t)}$ of colors i, j respectively. We know that $i \neq j$ since \mathcal{H} is properly edge-colored and $|C_1 \cap C_2| \geq t > 0$. Recall that the pair (C_1, C_2) forms a matching in K , denoted by $\{(S, T) : S \xleftrightarrow{C_1, C_2} T\}$. We will remove an edge (S, T) in this matching if S or T is incident to more than Δ edges containing color i or j . Our goal is to prove that this process removes at most $1/2$ fraction of the edges.

Consider a random edge $S \xleftrightarrow{C_1, C_2} T$ among the matching edges, which can be sampled by first choosing the symmetric difference $S \oplus T$ from $C_1 \setminus Q_\theta$ and $C_2 \setminus Q_\theta$, and then sampling $\ell - \lfloor \frac{q-t}{2} \rfloor$ vertices for $S \cap T$. By Markov's inequality,

$$\Pr_{(S,T)} [(S, T) \text{ is removed}] \leq \frac{1}{\Delta} \cdot \mathbb{E}_{(S,T)} [\text{number of edges incident to } (S, T) \text{ with color } i \text{ or } j].$$

We will upper bound the right-hand side via the co-degree guarantees from [Definition 2.6](#). There are two possible ways that an incident edge can have color i (same analysis for j).

- (S, T) is incident to an edge formed by C_1, C_3 . In this case, C_3 must also be in $\mathcal{H}_\theta^{(t)}$ and thus must have $|C_2 \cap C_3| \geq |Q_\theta| = t$. Suppose $|C_2 \cap C_3| = t + a$ for some $a \geq 0$. For a random (S, T) to be incident to an edge formed by (C_1, C_3) , one must sample at least $\max\{\lfloor \frac{q-t}{2} \rfloor - a, 0\}$ elements from $C_3 \setminus C_2$ to complete $S \cap T$. This probability is at most $O\left(\frac{\ell}{n}\right)^{\max\{\lfloor \frac{q-t}{2} \rfloor - a, 0\}}$. On the other hand, there are at most $O(d_{t+a})$ such C_3 with $|C_2 \cap C_3| = t + a$. Thus, using [Eq. \(2\)](#), the contribution from this case is at most

$$\begin{aligned} \frac{1}{\Delta} \sum_{a=0}^{q-t} O(d_{t+a}) \cdot O\left(\frac{\ell}{n}\right)^{\max\{\lfloor \frac{q-t}{2} \rfloor - a, 0\}} &\leq \frac{O(1)}{A'} \sum_{a=0}^{\lfloor \frac{q-t}{2} \rfloor} \frac{d_{t+a}}{d_t} \cdot n^{1-\frac{2t}{q}} \cdot n^{-\frac{2}{q}(\lfloor \frac{q-t}{2} \rfloor - a)} \\ &\leq \frac{O(1)}{A'} \sum_{a=0}^{\lfloor \frac{q-t}{2} \rfloor} \frac{d_{t+a}}{d_t} \cdot n^{\frac{2a}{q} - \frac{1}{q}(t-1)(t \text{ even})} \\ &\leq O(1/A'). \end{aligned}$$

Here, we use that d_s is decreasing with s to handle the terms with $a > \lfloor \frac{q-t}{2} \rfloor$, and the assumption $d_{t+a} \leq d_t \cdot n^{-\frac{2a}{q} + \frac{1}{q}(t-1)(t \text{ even})}$.

- (S, T) is incident to an edge formed by (C_3, C_4) , where $C_3 \neq C_1$ and C_3 has color i . Since \mathcal{H} is properly edge-colored, we have $C_1 \cap C_3 = \emptyset$ and $C_3, C_4 \in \mathcal{H}_{\theta'}^{(t)}$ for some $\theta' \neq \theta$. Suppose $|C_4 \cap C_2| = s$. Then, to form $S \cap T$, one must sample $\lfloor \frac{q-t}{2} \rfloor$ (resp. $\lceil \frac{q-t}{2} \rceil$) elements from C_3 and $\max\{\lceil \frac{q-t}{2} \rceil - s, 0\}$ (resp. $\max\{\lfloor \frac{q-t}{2} \rfloor - s, 0\}$) elements from $C_4 \setminus C_2$, hence at least $\max\{q-t-s, \lfloor \frac{q-t}{2} \rfloor\}$. The number of such pairs (C_3, C_4) can be upper bounded as follows: (1) if $s = 0$, then there are n choices for C_3 with color i , which belongs to some $\mathcal{H}_{\theta'}^{(t)}$, and $|\mathcal{H}_{\theta'}^{(t)}| \leq d_t$ choices for C_4 , (2) if $s \geq 1$, then there are $O(d_s)$ choices for C_4 , which belongs to

some $\mathcal{H}_{\theta'}^{(t)}$, and at most 1 choice for $C_3 \in \mathcal{H}_{\theta'}^{(t)}$ of color i . Thus, using [Eq. \(2\)](#), the contribution from this case is at most

$$\begin{aligned} & \frac{1}{\Delta} \left(n d_t \cdot O\left(\frac{\ell}{n}\right)^{q-t} + \sum_{s=1}^q O(d_s) \cdot O\left(\frac{\ell}{n}\right)^{\max\{q-t-s, \lfloor \frac{q-t}{2} \rfloor\}} \right) \\ & \leq \frac{O(1)}{A'} \cdot n^{1-\frac{2t}{q}} \left(n \cdot n^{-\frac{2}{q}(q-t)} + \sum_{s=1}^{\lceil \frac{q-t}{2} \rceil} \frac{d_s}{d_t} \cdot n^{-\frac{2}{q}(q-t-s)} \right) \\ & \leq O(1/A'). \end{aligned}$$

Again, we use that d_s is decreasing with s to handle the terms with $s > \lceil \frac{q-t}{2} \rceil$, and the assumption $d_s \leq d_t \cdot n^{1-\frac{2s}{q}}$.

Therefore, for large enough A' (depending only on q), we have that $\Pr_{(S,T)}[(S,T) \text{ is removed}] < 1/2$. Thus, the edge deletion process will remove at most $1/2$ of the edges. This completes the proof. \square

2.6 Regularity Decomposition

Our notion of approximate strong regularity ([Definition 2.7](#)) requires a good index t as in [Definition 2.6](#). For all we know, such an index may not even exist.

There is always a good t : One of our main technical insights is that given any hypergraph, *for every possible sequence of co-degrees* $d_1 \geq d_2 \geq \dots \geq d_q$, there is always an index $t \in [q]$ that satisfies all the conditions in [Definition 2.6](#)! The proof is elementary and short. We argue that either a $t_0 < q/2$ that extremizes a natural measure is good, or, if not, then an extremal d_t for $t > q/2$ that violates the requirements for t_0 being a good index works. See [Lemma 5.1](#) for the formal statement and [Section 5.1](#) for the proof. We remark that such a proof is possible precisely because all inequalities in [Definition 2.6](#) are relative.

Given that a good index always exists, we next describe how to decompose the hypergraph into $\sim \log n$ pieces, where each piece is t -approximately strongly regular for some $t \in [q]$. We refer readers to [Section 5](#) for more details.

Decomposition in prior works: All recent works on even covers and LDC lower bounds [[GKM22](#), [HKM23](#), [AGKM23](#), [HKM⁺24](#)] decompose the hypergraph into $\mathcal{H}^{(1)}, \dots, \mathcal{H}^{(q)}$ as follows: Carefully choose thresholds τ_s for all $s \in [q]$. Then, starting with $t = q$, (1) if there exists a $Q \in \binom{[n]}{t}$ such that $|\{C \in \mathcal{H} : Q \subseteq C\}| \geq \tau_t$, then add these hyperedges to $\mathcal{H}^{(t)}$ and remove them from \mathcal{H} , (2) if no such Q exists, then decrement t and repeat unless $t = 1$.

By this decomposition, we have the guarantee that for each $\mathcal{H}^{(t)}$, the co-degrees satisfy $d_s < \tau_s$ for all $s > t$, otherwise these hyperedges would have been put in $\mathcal{H}^{(s)}$. Unfortunately, this does not suffice for us, as [Definition 2.6](#) requires bounds on co-degrees for all s .

Approximate Strong Regularity decomposition: Rather than decomposing to satisfy some fixed thresholds on the co-degrees, we do the following:

- (1) Set the “current” hypergraph to be \mathcal{H} .
- (2) Apply [Lemma 5.1](#) on its co-degree sequence $d_1 \geq \dots \geq d_q$ to obtain a good index $t \in [q]$.
- (3) Greedily extract hyperedges containing “heavy” t -tuples from the current hypergraph, i.e. take $Q \in \binom{[n]}{t}$ that maximizes $|\mathcal{H}_{|Q}^{\text{curr}}| = |\{C \in \mathcal{H}' : Q \subseteq C\}|$, and remove $\mathcal{H}_{|Q}^{\text{curr}}$ from $\mathcal{H}^{\text{curr}}$.
- (4) Repeat step (3) until the d_t -value in $\mathcal{H}^{\text{curr}}$ falls below half of its original value (in the parent hypergraph).
- (5) At this point, all the $\mathcal{H}_{|Q}^{\text{curr}}$ s which were removed form a t -approximately strongly regular hypergraph. We designate the hypergraph remaining as the next “current” hypergraph, and repeat the above process.

Note that in the above process, the value of d_t in the “current” graph goes down by a factor of 2 for some $t \in [q]$ for every iteration, and thus the above process can only run for $\sim q \log_2 |\mathcal{H}|$ steps (if the process goes on for longer, then some d_t gets halved $\geq q \log_2 |\mathcal{H}|/q \geq \log_2 |\mathcal{H}|$ times).

Therefore, this decomposition produces $O(q \log n)$ subhypergraphs, each of which is approximately strongly regular, and one of them contains at least $\Omega(\frac{1}{q \log n})$ fraction of \mathcal{H} .

Claim 2.9. *There exists a subhypergraph $\mathcal{H}^\pi \subseteq \mathcal{H}$ with at least $\Omega\left(\frac{|\mathcal{H}|}{q \log n}\right)$ hyperedges and is t -approximately strongly regular for some $t \in [q]$.*

See [Lemma 4.7](#) and its proof in [Section 5](#) for a more general result that we need for the non-linear case. For linear LDCs, the largest piece in the decomposition suffices for our analysis.

Now, [Theorem 2.1](#) is a simple corollary of [Lemma 2.8](#) and [Claim 2.9](#).

Proof of [Theorem 2.1](#). Towards a contradiction, let the message length satisfy $k \geq \Omega(\delta^{-1} n^{1-\frac{2}{q}} \log^2 n)$. In the underlying hypergraph \mathcal{H} , each of the k matchings has size $\Omega(\delta n)$. Let \mathcal{H}^π be the subhypergraph of \mathcal{H} from [Claim 2.9](#). Then, \mathcal{H}^π has average degree $\Omega(n^{1-\frac{2}{q}} \log n)$, and is t -approximately strongly regular for some $t \in [q]$, thus we can apply [Lemma 2.8](#) and conclude that \mathcal{H}^π (hence also \mathcal{H}) must contain a weak rainbow even cover. This is a contradiction via [Fact 2.2](#). \square

2.7 Non-linear LDC Lower Bounds

To prove lower bounds for non-linear q -LDCs, we depart from the combinatorial lens, since merely showing the existence of a weak rainbow even cover does not suffice. Instead, we use the connection to CSP refutation discovered by [\[AGKM23\]](#): Given a q -query normally decodable code ([Definition 3.3](#)) $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ associated with q -uniform matchings $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$, for each message $b \in \{\pm 1\}^k$, we can construct a q -XOR instance Ψ_b over $x \in \{\pm 1\}^n$ with clauses $x_C = b_i$ for each $i \in [k]$ and $C \in \mathcal{H}_i$. Here, we denote $x_C = \prod_{u \in C} x_u$.

[Definition 3.3](#) can be interpreted as that every clause C in Ψ_b is satisfied by $x = E(b)$ for “most” messages $b \in \{\pm 1\}^k$. This implies that the value of Ψ_b — the maximum number of satisfied constraints among x — must be large for “most” $b \in \{\pm 1\}^k$. Thus, to rule out the existence of a

q -LDC (with k too large), it suffices to prove that for a random $b \in \{\pm 1\}^k$, Ψ_b is *highly unsatisfiable* — no assignment satisfies $\frac{1}{2} + \Omega(1)$ fraction of the clauses. This is closely related to semirandom CSP refutation since both problems involve proving unsatisfiability of instances where the underlying hypergraph structure is arbitrary but the signs are random.

We write Ψ_b as a degree- q polynomial in $x = (x_1, \dots, x_n)$: $\Psi_b(x) = \sum_{i=1}^k \sum_{C \in \mathcal{H}_i} b_i x_C$. Denoting $m = |\mathcal{H}| = \sum_{i=1}^k |\mathcal{H}_i|$, we have $\Psi_b(x) \in [-m, m]$ for all $x \in \{\pm 1\}^n$. Denote $\text{val}(\Psi_b) = \max_{x \in \{\pm 1\}^n} \Psi_b(x)$. Note that $\text{val}(\Psi_b) \geq 0$ since $\mathbb{E}_x[\Psi_b(x)] = 0$ under a uniformly random x . A q -LDC must satisfy $\mathbb{E}_b[\text{val}(\Psi_b)] \geq \Omega(m)$ (see [Observation 4.3](#)); thus we would like to prove that $\mathbb{E}_b[\text{val}(\Psi_b)] \leq o(m)$ for any arbitrary \mathcal{H} when $k \gg n^{1-2/q}$.

The key connection to Kikuchi graphs is as follows. For even q , $\Psi_b(x)$ is equivalent to the quadratic form of the *signed* Kikuchi matrix $K_b := \sum_{i=1}^k b_i K_i$, where K_i is the unsigned Kikuchi matrix of \mathcal{H}_i . More specifically, $\Psi_b(x) = c \cdot z^\top K_b z$, where $z \in \{\pm 1\}^N$ is some vector of monomials of x and $c > 0$ is some scaling. Then, since $\|z\|_2^2 = N$ for any x , it suffices to prove an upper bound on $\|K_b\|_2$. In particular, we can upper bound $\mathbb{E}_b[\|K_b\|_2] = \mathbb{E}_b[\|\sum_{i=1}^k b_i K_i\|_2]$ using the well-known Matrix Khintchine inequality ([Fact 3.5](#)).

For odd q , however, we cannot write $\Psi_b(x)$ directly as a quadratic form, so we need the *Cauchy-Schwarz trick* which we describe next.

Cauchy-Schwarz trick: Let us first consider a subhypergraph \mathcal{H}' such that every $C \in \mathcal{H}'$ contains some tuple $Q \in \binom{[n]}{t}$, where $t \in [q]$. Let $\Psi_b^{\mathcal{H}'}$ be the sub-instance of Ψ_b corresponding to \mathcal{H}' . With slight abuse of notation, for $C \in \mathcal{H}_i$, we denote $b_C = b_i$. Then,

$$\Psi_b^{\mathcal{H}'}(x)^2 = \left(\sum_{C \in \mathcal{H}'} b_C x_C \right)^2 = |\mathcal{H}'| + \sum_{C \neq C' \in \mathcal{H}'} b_C b_{C'} x_{C \oplus C'}.$$

Notice that $|C \oplus C'| = 2(q - t)$ when $|C \cap C'| = t$. Recall that the Kikuchi graph for odd q ([Definition 4.11](#)) is formed precisely by taking pairs of hyperedges C, C' with intersection t and canceling them. It is not hard to show that the second term in the above can be expressed as a quadratic form of the signed Kikuchi matrix (see [Lemma 4.14](#)), though each edge $S \xleftrightarrow{C, C'} T$ now gets a correlated sign $b_C b_{C'}$. Note also that since $C \cap C' \neq \emptyset$ and $\mathcal{H}_1, \dots, \mathcal{H}_k$ are hypergraph matchings, C and C' must belong to different \mathcal{H}_i 's and thus $b_C, b_{C'}$ are different random variables.

Pruning heavy vertices: In the discussions above, we ignored an important detail: $\mathbb{E}_b[\|K_b\|_2]$ can be large due to some large degree vertices in the Kikuchi graph, i.e., rows/columns in the matrix with large ℓ_1 norm. On the other hand, to upper bound $z^\top K_b z$ for $z \in \{\pm 1\}^N$, it suffices to upper bound the *infinity-to-1 norm* defined as $\|K_b\|_{\infty \rightarrow 1} := \max_{u, v \in \{\pm 1\}^N} u^\top K_b v$. Then, letting \tilde{K}_b be the matrix obtained by removing heavy rows and columns, we can bound $\|K_b\|_{\infty \rightarrow 1} \leq \|K_b - \tilde{K}_b\|_{\infty \rightarrow 1} + \|\tilde{K}_b\|_{\infty \rightarrow 1}$.⁵ We can thus bound $\mathbb{E}_b[\|\tilde{K}_b\|_{\infty \rightarrow 1}]$ by its spectral norm using the Matrix Khintchine Inequality ([Fact 3.5](#)).

On the other hand, to bound $\|K_b - \tilde{K}_b\|_{\infty \rightarrow 1}$, we need to upper bound the contribution from those pruned rows and columns. Consider a random $S \in \binom{[n]}{\ell}$ (i.e., a random vertex in the Kikuchi

⁵ This is the row pruning step in [[GKM22](#)].

graph), and let $s \in \{0, 1\}^n$ be the indicator vector of S . For the purpose of this overview, we think of s as a vector with i.i.d. Bernoulli variables with bias ℓ/n . Then, the random variable $\deg(S)$ can be expressed as a polynomial in s of degree at most q . To bound the tail probabilities of $\deg(S)$, we now invoke a standard polynomial concentration inequality (Fact 3.6) by Schudy and Sviridenko [SS12], which improves upon the classical work of [KV00].

Crucially, Fact 3.6 requires upper bounds on the *partial derivatives* of the polynomial. It turns out that these derivatives can be upper bounded by the co-degrees of the hypergraph \mathcal{H} (see Lemma 4.20)! Moreover, the co-degree bounds required coincide with those stated in Lemma 2.8.

Therefore, our strategy is to use the decomposition described in Section 2.6 on the q -XOR instance Ψ_b , which gives us sub-instances whose underlying hypergraphs are t -approximately strongly regular. Then, we obtain upper bounds on each sub-instance by using the Cauchy-Schwarz trick (with the good index t) and the infinity-to-1 norm of the corresponding Kikuchi matrix, thus finishing the proof. This is articulated in Section 4.2.

3 Preliminaries

For a $N \times N$ matrix A , we define $\|A\|_2 := \max_{u \neq 0} \|Au\|_2 / \|u\|_2$ to be its spectral norm and $\|A\|_{\infty \rightarrow 1} := \max_{u, v \in \{\pm 1\}^N} u^T A v$ to be the $\infty \rightarrow 1$ norm.

Hypergraphs. In this work, we will work with q -uniform hypergraphs \mathcal{H} over $[n]$. Unless otherwise stated, our hypergraphs are allowed to have parallel or “multi”-edges. We write $|\mathcal{H}|$ to denote the number of hyperedges in \mathcal{H} .

We first recall the standard notion of *co-degrees* in a hypergraph.

Definition 3.1 (Co-degrees). Let \mathcal{H} be a hypergraph on $[n]$. For any $Q \subseteq [n]$, let $\mathcal{H}_{|Q} := \{C \in \mathcal{H} : Q \subseteq C\}$ be the subset of hyperedges containing Q . We define the co-degree of Q in \mathcal{H} to be $d_{\mathcal{H}, Q} = |\mathcal{H}_{|Q}|$. We let $d_{\mathcal{H}, t} := \max_{Q \in \binom{[n]}{t}} d_{\mathcal{H}, Q}$ be the maximum possible co-degree of any subset of size t . We’ll refer to $d_{\mathcal{H}, Q}, d_{\mathcal{H}, t}$ simply as d_Q, d_t respectively when \mathcal{H} is clear from the context.

Observe that $d_1 \geq d_2 \geq \dots \geq d_q$ and $d_1 \geq |\mathcal{H}|/n$.

A hypergraph \mathcal{H} is a matching if all its hyperedges are pairwise disjoint. A proper edge-coloring of a hypergraph \mathcal{H} is an assignment of colors to the hyperedges of \mathcal{H} such that all hyperedges incident on any vertex get distinct colors.

3.1 Locally Decodable Codes

Definition 3.2 (Locally Decodable Code (LDC)). A code $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ is (q, δ, ε) -locally decodable if there exists a randomized decoding algorithm $\text{Dec}(\cdot)$ such that $\text{Dec}(\cdot)$ is given oracle access to $y \in \{0, 1\}^n$, takes an $i \in [k]$ as an input, and satisfies the following properties:

1. $\text{Dec}(\cdot)$ makes $\leq q$ queries to the string y .
2. For all $b \in \{\pm 1\}^k$, and for all $y \in \{0, 1\}^n$ such that $\text{dist}(y, C(b)) \leq \delta n$, $\Pr[\text{Dec}^y(i) = b_i] \geq \frac{1}{2} + \varepsilon$.

Here $\text{dist}(x, y) := \#\{v \in [n] : x_v \neq y_v\}$ denotes the Hamming distance between x and y .

Definition 3.3 (Normally Decodable Codes). A code $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ is called (q, δ, ε) -normally decodable if there exist q -uniform hypergraph matchings $\mathcal{H}_1, \dots, \mathcal{H}_k$, each with at least δn hyperedges, such that for every $C \in \mathcal{H}_i$, we have $\Pr_{b \leftarrow \{\pm 1\}^k} [b_i = \prod_{v \in C} E(b)_v] \geq \frac{1}{2} + \varepsilon$.

The following reduction, which is a variant of a classical result of [KT00], allows us to work with normally decodable codes without loss of generality.

Fact 3.4 (Reduction to Normally Decodable Codes, Lemma 6.2 in [Yek12]). Let $C : \{0, 1\}^k \rightarrow \{0, 1\}^n$ be a (q, δ, ε) -LDC. Then there is a code $C' : \{\pm 1\}^k \rightarrow \{\pm 1\}^{O(n)}$ that is $(q, \delta', \varepsilon')$ -normally decodable, with $\delta' \geq \varepsilon \delta / 3q^2 2^{q-1}$ and $\varepsilon' \geq \varepsilon / 2^{2q}$.

3.2 Concentration Inequalities

The following is a standard concentration result of random matrices.

Fact 3.5 (Matrix Khintchine Inequality [TJ74, LPP91]; Theorem 4.1.1 of [Tro15]). Let A_1, A_2, \dots , be arbitrary $M \times N$ matrices. Then,

$$\mathbb{E}_{b_1, b_2, \dots \in \{\pm 1\}} \left\| \sum_i b_i A_i \right\|_2 \leq \sqrt{2 \log(M+N)} \cdot \max \left\{ \left\| \sum_i A_i A_i^\top \right\|_2^{1/2}, \left\| \sum_i A_i^\top A_i \right\|_2^{1/2} \right\},$$

where the expectation is taken over independent and uniform $b_i \in \{\pm 1\}$.

We also utilize the following polynomial concentration inequality.

Fact 3.6 (Polynomial Concentration; Theorem 1.1 in [SS12]). Let $f(x) = f(x_1, \dots, x_n)$ be a multilinear polynomial of degree q with non-negative coefficients. Let Y_1, \dots, Y_n be i.i.d Bernoulli random variables. Write $f(Y) := f(Y_1, \dots, Y_n)$, and let $\mu_0 := \mathbb{E}f(Y)$. For $r \in [q]$, define

$$\mu_r := \max_{Z \in \binom{[n]}{r}} \mathbb{E}_Y (\partial_Z f)(Y),$$

where $\partial_Z f := (\prod_{i \in Z} \partial_{x_i}) f$. Then there exists an absolute constant $W \geq 1$ such that for any $\lambda > 0$, we have:

$$\Pr(|f(Y) - \mu| \geq \lambda) \leq e^2 \cdot \max \left\{ \max_{r \in [q]} \exp \left(-\frac{\lambda^2}{\mu_0 \mu_r W^q} \right), \max_{r \in [q]} \exp \left(-\left(\frac{\lambda}{\mu_r W^q} \right)^{1/r} \right) \right\}.$$

4 Proof of the Main Theorem

In this section, we introduce the core components of our argument and use them to prove:

Theorem 4.1 (Main Theorem). Let $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ be a (q, δ, ε) -LDC for any $q \geq 3$. Then,

$$k \leq n^{1-\frac{2}{q}} \cdot \frac{2^{O(q)} \log^4 n}{\varepsilon^6 \delta^2}.$$

Given [Fact 3.4](#), we can assume that the LDC $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ is in the normal form. Let $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$ be the q -uniform matchings of size $\geq \delta n$ for the associated normal decoder. By a slight abuse of notation, we will use \mathcal{H} to denote both the tuple $(\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k)$ and the hypergraph formed by their union $\bigcup_{i=1}^k \mathcal{H}_i$. As in [\[AGKM23\]](#), we will associate a system of q -XOR constraints that encode the success of the decoder for E . Let us first define this q -XOR system.

Definition 4.2 (q -XOR Polynomial for a Normal LDC). Given a collection of q -uniform matchings $\mathcal{H} = (\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k)$ on $[n]$ and any $b \in \{\pm 1\}^k$, let

$$\Psi_b^{\mathcal{H}}(x) = \sum_{i=1}^k b_i \sum_{C \in \mathcal{H}_i} x_C,$$

be a degree- q polynomial in variables x_1, x_2, \dots, x_n corresponding to the codeword bits. We drop the superscript and write Ψ_b when the collection of q -uniform matchings \mathcal{H} is clear from the context. Let $\text{val}(\Psi_b) = \max_{x \in \{\pm 1\}^n} \Psi_b(x)$ be the *value* of Ψ_b .

Informally speaking, on average over $b \in \{\pm 1\}^k$, each monomial in Ψ_b corresponds to a local decoding constraint that is satisfied with a non-trivial advantage by the codeword $E(b)$. This is immediate given [Definition 3.3](#), and we record it as the following observation:

Observation 4.3. For a (q, δ, ε) -normally decodable code E , we have $\mathbb{E}_{b \in \{\pm 1\}^k} [\Psi_b(E(b))] \geq 2\varepsilon |\mathcal{H}|$.

Remark 4.4. Note that for a normally decodable code, \mathcal{H} is a union of k hypergraph matchings and since $k \leq n$, the hypergraph thus has size $\leq nk \leq n^2$.

Our goal in the rest of this section is to prove an upper bound on $\mathbb{E}_b[\text{val}(\Psi_b)]$, which, when combined with [Observation 4.3](#), will give us [Theorem 4.1](#). To do this, we will first identify a sufficient condition on \mathcal{H} called *approximate strong regularity* ([Definition 2.7](#)), and argue that this condition can be ensured without the loss of generality by an appropriate decomposition applied to an arbitrary hypergraph \mathcal{H} .

Approximate Strong Regularity Our notion of approximate strong regularity can be seen as a generalization of similar definitions used in in [\[GKM22, HKM23\]](#). Their notions of regularity demand absolute upper bounds on all co-degrees in the hypergraph. In contrast, our definition below identifies a special index t and demands 1) that all size t sets have their co-degrees close up to a factor 2 to a fixed value d_t , and 2) the co-degree of sets of size $r \neq t$ be appropriately small *relative to d_t* . Crucially, our definition neither demands nor guarantees any non-trivial absolute upper bound on any co-degree. The following definition (discussed earlier in [Section 2](#)) notes the relative upper bounds we need so as to make a crucial technical component of our argument go through.

Definition (Restatement of [Definition 2.6](#)). Let $d_1 \geq d_2 \geq \dots \geq d_q$. We say that an index t is *good* with respect to the tuple (d_1, d_2, \dots, d_q) if the following conditions hold:

(1) $d_r/d_t \leq n^{1-\frac{2r}{q}}$ for every $1 \leq r \leq \lceil \frac{q-t}{2} \rceil$,

$$(2) \quad d_r/d_t \leq n^{-\frac{2}{q}(r-t)+\frac{1}{q}(t-1(t \text{ even}))} \text{ for every } t \leq r \leq \lfloor \frac{q+t}{2} \rfloor,$$

$$(3) \quad \text{If } t < q/2, \text{ then, } d_t \geq d_1 n^{-\frac{2}{q}(t-1)}, \text{ and, if } t > q/2, \text{ then, } d_t \geq d_1 n^{-1+2/q}.$$

We can now define approximate strong regularity. Informally speaking, \mathcal{H} is approximately regular if one can partition the hyperedges of \mathcal{H} into groups so that 1) the hyperedges in a group all contain some set of size t , 2) each group has the same size up to a factor 2, and, for the co-degree tuple formed by d_r being the maximum possible co-degree of any set of size r , t is a good index.

Definition (Restatement of [Definition 2.7](#)). For a given $t \in [q]$, a q -uniform hypergraph \mathcal{H} on $[n]$ is called t -approximately strongly regular if there exists a partitioning $\mathcal{H} = \mathcal{H}_1^{(t)} \sqcup \dots \sqcup \mathcal{H}_{p_t}^{(t)}$ such that:

$$(1) \quad \text{For every } \theta \in [p_t], \text{ there is some } Q_\theta \in \binom{[n]}{t} \text{ such that all } C \in \mathcal{H}_\theta^{(t)} \text{ contain } Q_\theta.$$

$$(2) \quad d_t/2 \leq |\mathcal{H}_\theta^{(t)}| \leq d_t \text{ for all } \theta \in [p_t].$$

$$(3) \quad t \text{ is a good index with respect to the sequence } d_1 \geq \dots \geq d_q \text{ (recall that } d_r := \max_{|Q|=r} d_{\mathcal{H},Q}).$$

Remark 4.5. It is easy to see that we have $p_t \leq \frac{2|\mathcal{H}|}{d_t}$. Also, let $P = \sqcup_{t \in [q]} P_t$, where P_t is the set of all indices π such that \mathcal{H}^π is t -approximately strongly regular.

Our key technical argument is captured in the following lemma that shows an upper bound on $\mathbb{E}_b [\text{val}(\Psi_b^{\mathcal{H}})^2]$ if \mathcal{H} is t -approximately strongly regular.

Lemma 4.6. Let $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$ be q -uniform matchings on $[n]$ such that $\cup_{i \in [k]} \mathcal{H}_i$ is t -approximately strongly regular with degree sequence $d_1 \geq d_2 \geq \dots \geq d_q$ such that $k \geq d_1 \geq \ell := n^{1-2/q} \log n$. Then we have

$$\mathbb{E}_b \left[\text{val} \left(\Psi_b^{\mathcal{H}} \right)^2 \right] \leq \frac{2|\mathcal{H}|^2}{d_t} + |\mathcal{H}| n \sqrt{k\ell} O(\log n)^{1/2}.$$

We will prove [Lemma 4.6](#) in [Section 4.1](#) modulo a few technical results proved in subsequent sections. To finish the proof, we will show in [Section 5](#) that every $\mathcal{H} = (\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k)$ can be decomposed into approximately strongly regular pieces up to a negligible error.

Lemma 4.7. Let \mathcal{H} be a q -uniform hypergraph on $[n]$. Let $\eta > 0$. Then, there is a collection $\{\mathcal{H}^\pi : \pi \in P\}$, with $|P| \leq q \lceil \log_2 |\mathcal{H}| \rceil + 1$, of hypergraphs such that:

$$1. \quad \text{For every } \pi \in P, \text{ there exists } t \in [q] \text{ such that } \mathcal{H}^\pi \text{ is } t\text{-approximately strongly regular.}$$

$$2. \quad \sqcup_{\pi \in P} \mathcal{H}^\pi \subseteq \mathcal{H},$$

$$3. \quad |\sqcup_{\pi \in P} \mathcal{H}^\pi| \geq (1 - \eta)|\mathcal{H}|,$$

$$4. \quad |\mathcal{H}^\pi| \geq \frac{\eta|\mathcal{H}|}{q \lceil \log_2 |\mathcal{H}| \rceil + 1} \text{ for all } \pi \in P.$$

Let us complete the proof of [Theorem 4.1](#) using [Lemmas 4.6](#) and [4.7](#).

Proof of Theorem 4.1. Let $E : \{\pm 1\}^k \rightarrow \{\pm 1\}^n$ be a (q, δ, ε) -LDC. Let $\varepsilon' = \varepsilon/2^{2q}$ and $\delta' = \varepsilon\delta/3q^22^{q-1}$ be the parameters obtained from [Fact 3.4](#) when we reduce to the normal form. From [Observation 4.3](#), we thus have that: $\mathbb{E}_{b \sim \{\pm 1\}^k} [\text{val}(\Psi_b)] \geq 2\varepsilon'|\mathcal{H}|$.

Set $\ell = n^{1-\frac{2}{q}} \log n$. Suppose towards a contradiction that $k \geq An^{1-\frac{2}{q}} \cdot \frac{q^2 \log^4 n}{\varepsilon^4 \delta^2}$ for a large enough constant $A > 0$. Note that this expression is at most $n^{1-\frac{2}{q}} \cdot \frac{2^{O(q)} \log^4 n}{\varepsilon^6 \delta^2}$.

Let us apply the results we developed so far to construct an upper bound on $\mathbb{E}_b[\text{val}(\Psi_b)]$. First, we apply the decomposition from [Lemma 4.7](#) for $\eta = \varepsilon'/2$ and get a collection of subhypergraphs $\{\mathcal{H}^\pi : \pi \in P\}$. Here, we have $|P| \leq O(q \log n)$ and $\sum_{\pi \in P} |\mathcal{H}^\pi| \geq (1 - \eta)|\mathcal{H}|$. Let P_t be the set of all indices π such that \mathcal{H}^π is t -approximately strongly regular (where t is the good index). For simplicity, we denote $d_r^\pi := d_{\mathcal{H}^\pi, r}$ to be the co-degree bounds of \mathcal{H}^π . Then we can write:

$$\begin{aligned} \mathbb{E}_b[\text{val}(\Psi_b)] &\leq \mathbb{E}_b \left[\text{val} \left(\sum_{\pi \in P} \Psi_b^{\mathcal{H}^\pi} \right) \right] + \eta|\mathcal{H}| \leq \sum_{\pi \in P} \mathbb{E}_b \left[\text{val} \left(\Psi_b^{\mathcal{H}^\pi} \right) \right] + \eta|\mathcal{H}| \\ &\leq \sqrt{O(q \log n) \cdot \sum_{\pi \in P} \mathbb{E}_b \left[\text{val} \left(\Psi_b^{\mathcal{H}^\pi} \right)^2 \right]} + \eta|\mathcal{H}|, \end{aligned} \quad (3)$$

where the last inequality follows from Cauchy-Schwarz.

For each $\pi \in P$, we have $|\mathcal{H}^\pi| \geq \frac{\eta|\mathcal{H}|}{q \log_2 |\mathcal{H}|+1} \geq \frac{\eta\delta'kn}{3q \log n}$ since $|\mathcal{H}| \leq n^2$. Thus, we have $d_1^\pi \geq |\mathcal{H}^\pi|/n \geq \frac{\eta\delta'k}{3q \log n} \gg \ell$ by the lower bound on k . This allows us to invoke [Lemma 4.6](#) (which requires $d_1^\pi \geq \ell$) for all \mathcal{H}^π for $\pi \in P_t$:

$$\mathbb{E}_b \left[\text{val} \left(\Psi_b^{\mathcal{H}^\pi} \right)^2 \right] \leq \frac{2|\mathcal{H}^\pi|^2}{d_t^\pi} + O(1) \cdot n\sqrt{k\ell}(\log n)^{1/2}|\mathcal{H}^\pi|.$$

Since t is a good index for \mathcal{H}^π , condition (3) of [Definition 2.6](#) gives that $d_t^\pi \geq d_1^\pi n^{-1+2/q} \geq |\mathcal{H}^\pi| n^{-2+2/q}$. Thus, the first term above is upper bounded by $2|\mathcal{H}^\pi| n^{2-2/q}$. Thus, summing over $\pi \in P_t$ and $t \in [q]$, we have

$$\sum_{\pi \in P} \mathbb{E}_b \left[\text{val} \left(\Psi_b^{\mathcal{H}^\pi} \right)^2 \right] \leq O(1) \cdot |\mathcal{H}| \left(n^{2-\frac{2}{q}} + n\sqrt{k\ell \log n} \right).$$

We next show that the above is upper bounded by $\frac{c\varepsilon^2}{q \log n} \cdot |\mathcal{H}|^2$ for some small enough constant c .

Note that $|\mathcal{H}| \geq \delta'kn$. Clearly, we have $n^{2-\frac{2}{q}} \leq o\left(\frac{\varepsilon^2}{q \log n}\right) \cdot |\mathcal{H}|$ by our choice of k . Next, we verify that $n\sqrt{k\ell \log n} \leq O\left(\frac{\varepsilon^2}{q \log n}\right) \cdot \delta'kn$. This is equivalent to $k \geq \Omega\left(\frac{q^2}{\varepsilon^4 \delta^2}\right) \ell \log^3 n$. Since $\ell = n^{1-\frac{2}{q}} \log n$, this is satisfied by our choice of k .

Therefore, by [Eq. \(3\)](#) we have $\mathbb{E}_b[\text{val}(\Psi_b)] \leq (\varepsilon'/2 + \eta)|\mathcal{H}| \leq \varepsilon'|\mathcal{H}|$. This contradicts that $\mathbb{E}_b[\text{val}(\Psi_b)] \geq 2\varepsilon'|\mathcal{H}|$. \square

4.1 Proof of [Lemma 4.6](#): refuting regular q -XOR

In this section, we prove [Lemma 4.6](#). To do this, we would like to write Ψ_b as a quadratic form of a ‘‘Kikuchi’’ matrix and then upper bound it using the spectral norm of the matrix. As in prior works,

the main issue is that $\Psi_b(x)$ is an odd-degree homogeneous polynomial while Kikuchi matrices naturally yield even-degree polynomials as their quadratic forms. Therefore, we will upper bound $\text{val}(\Psi_b)^2$ and use the Cauchy-Schwarz trick as described in [Section 2.7](#). This yields an even-degree polynomial, which we will then analyze a spectral upper bound for. We first define this even-degree polynomial and then show how it can be written as a quadratic form of an appropriate Kikuchi matrix. In order to ensure independence of some matrices that arise later in our analysis, we will define our polynomial with respect to a fixed partition $L \sqcup R$ of $[k]$.

Definition 4.8. Let \mathcal{H} be a t -approximately strongly regular hypergraph. Let $\{\mathcal{H}_\theta^{(t)}\}_{\theta \in [p_t]}$ be the partition given by approximate strong regularity, so that for $\theta \in [p_t]$, there is a $Q_\theta \in \binom{[n]}{t}$ such that every $C \in \mathcal{H}_\theta$ satisfies $Q_\theta \subseteq C$. For $b \in \{\pm 1\}^k$ and a partition $L \sqcup R = [k]$, we define

$$f_{L,R,t}^{\mathcal{H}}(x) := \sum_{\theta \in [p_t]} \sum_{\substack{i \in L \\ j \in R}} b_i b_j \sum_{\substack{C \in \mathcal{H}_\theta^{(t)} \cap \mathcal{H}_i, \\ C' \in \mathcal{H}_\theta^{(t)} \cap \mathcal{H}_j}} x_{C \setminus Q_\theta} x_{C' \setminus Q_\theta}.$$

Remark 4.9 (Unraveling the definition of $f_{L,R,t}^{\mathcal{H}}$). To define $f_{L,R,t}^{\mathcal{H}}$, we work with an arbitrary partition of the original matchings $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k$ into a “left” group L and a “right” group R (this partitioning later ensures independence when we apply matrix concentration).

Then, $f_{L,R,t}^{\mathcal{H}}$ can be viewed as merging pairs of hyperedges between the left and right groups within each $\mathcal{H}_\theta^{(t)}$. That is, $f_{L,R,t}^{\mathcal{H}}$ includes one monomial for each pair $C, C' \in \mathcal{H}_\theta^{(t)}$ such that $C \in \mathcal{H}_i$ for $i \in L$ and $C' \in \mathcal{H}_j$ for $j \in R$, obtained by multiplying $b_i x_C$ and $b_j x_{C'}$. Since $Q_\theta \subseteq C, C'$, the product has a squared term $x_{Q_\theta}^2$ which “cancels” out to 1.

The Cauchy-Schwarz trick: Let us relate $\mathbb{E}_{L,R}[\text{val}(f_{L,R,t}^{\mathcal{H}})]$ with $\Psi_b^{\mathcal{H}}$ and conclude that it is enough to bound $\mathbb{E}_{L,R}[\text{val}(f_{L,R,t}^{\mathcal{H}})]$. Here, the expectation is over a uniformly random partition $L \cup R$ of $[k]$. We emphasize again that t -approximate strong regularity ([Definition 2.7](#)) guarantees a partitioning $\{\mathcal{H}_\theta^{(t)}\}_{\theta \in [p_t]}$.

Lemma 4.10 (Cauchy-Schwarz Trick). *Let \mathcal{H} be a t -approximately strongly regular hypergraph. Then we have*

$$\text{val}(\Psi_b^{\mathcal{H}})^2 \leq \frac{2|\mathcal{H}|^2}{d_t} + \frac{8|\mathcal{H}|}{d_t} \mathbb{E}_{L,R} \left[\text{val}(f_{L,R,t}^{\mathcal{H}}) \right].$$

Proof. For convenience, we will slightly abuse notation and denote b_C to be b_i for $C \in \mathcal{H}_i$. We have:

$$\Psi_b^{\mathcal{H}}(x) = \sum_{C \in \mathcal{H}} b_C x_C = \sum_{\theta \in [p_t]} x_{Q_\theta} \sum_{C \in \mathcal{H}_\theta^{(t)}} b_C x_{C \setminus Q_\theta}.$$

Thus, by the Cauchy-Schwarz inequality:

$$\Psi_b^{\mathcal{H}}(x)^2 \leq \left(\sum_{\theta \in [p_t]} x_{Q_\theta}^2 \right) \sum_{\theta \in [p_t]} \left(\sum_{C \in \mathcal{H}_\theta^{(t)}} b_C x_{C \setminus Q_\theta} \right)^2 = p_t \sum_{\theta \in [p_t]} \left(|\mathcal{H}_\theta^{(t)}| + \sum_{C \neq C' \in \mathcal{H}_\theta^{(t)}} b_C b_{C'} x_{C \setminus Q_\theta} x_{C' \setminus Q_\theta} \right).$$

Note that all hyperedges in $\mathcal{H}_\theta^{(t)}$ contain Q_θ , so they must belong to distinct matchings. For any $C \neq C' \in \mathcal{H}_\theta$ where $C \in \mathcal{H}_i$ and $C' \in \mathcal{H}_j$ for some $i \neq j \in [k]$, when we take a random partition $L \sqcup R = [k]$, we have that $i \in L$ and $j \in R$ with probability exactly $1/4$. Thus, the above expression equals

$$p_t |\mathcal{H}| + p_t \cdot 4 \mathbb{E}_{L,R} \sum_{\theta \in [p_t]} \sum_{i \in L, j \in R} b_i b_j \sum_{\substack{C \in \mathcal{H}_i \cap \mathcal{H}_\theta^{(t)}, \\ C' \in \mathcal{H}_j \cap \mathcal{H}_\theta^{(t)}}} x_{C \setminus Q_\theta} x_{C' \setminus Q_\theta} = p_t \left(|\mathcal{H}| + 4 \mathbb{E}_{L,R} \left[f_{L,R,t}^{\mathcal{H}}(x) \right] \right).$$

Here, the last inequality uses our definition of $f_{L,R,t}^{\mathcal{H}}$ in [Definition 4.8](#). Thus, since we have $p_t \leq \frac{2|\mathcal{H}|}{d_t}$ by the remark following [Definition 2.7](#), we apply Jensen's Inequality to get

$$\text{val}(\Psi_b)^2 \leq \frac{2|\mathcal{H}|^2}{d_t} + \max_{x \in \{\pm 1\}^n} \frac{8|\mathcal{H}|}{d_t} \mathbb{E}_{L,R} \left[f_{L,R,t}^{\mathcal{H}}(x) \right] \leq \frac{2|\mathcal{H}|^2}{d_t} + \frac{8|\mathcal{H}|}{d_t} \mathbb{E}_{L,R} \left[f_{L,R,t}^{\mathcal{H}}(x) \right]. \quad \square$$

$f_{L,R,t}^{\mathcal{H}}$ as a quadratic form We will write $f_{L,R,t}^{\mathcal{H}}$ as a quadratic form of a natural (signed) *Kikuchi* matrix (same as the one appearing in [\[HKM23\]](#)). See [Figure 2](#) for an example. For any set S , and any $X, Y \subseteq S$, define $X \oplus Y := (X \setminus Y) \cup (Y \setminus X)$ to be the symmetric difference of X and Y . We define Kikuchi graphs associated with the monomials appearing in $f_{L,R,t}^{\mathcal{H}}$.

Definition 4.11 (Level- ℓ Kikuchi Graph). Given a t -approximately strongly regular hypergraph \mathcal{H} , hypergraph matchings $\mathcal{H} = (\mathcal{H}_1, \dots, \mathcal{H}_k)$, and a partition $\{\mathcal{H}_\theta^{(t)}\}_{\theta \in [p_t]}$ from approximate strong regularity, we define the Kikuchi graph as follows.

We create two copies $(x_u, 1)$ and $(x_u, 2)$ of the original n variables x_u for each $1 \leq u \leq n$. The Kikuchi graph has vertices indexed by subsets $S \in \binom{[n] \times [2]}{\ell}$ for a parameter $\ell \in \mathbb{N}$. For any $C \neq C' \in \mathcal{H}_\theta^{(t)}$, let $\tilde{C} := C \setminus Q_\theta$ and $\tilde{C}' := C' \setminus Q_\theta$. For any $S, T \in \binom{[n] \times [2]}{\ell}$, we say that $S \xleftrightarrow{C, C'} T$ if:

1. $S \oplus T = X \oplus Y$, for $X := \tilde{C} \times \{1\}, Y := \tilde{C}' \times \{2\}$. Clearly, $X \oplus Y = X \cup Y$ as X, Y are disjoint.
2. When $q - t$ is even, $|S \cap X| = |S \cap Y| = |T \cap X| = |T \cap Y| = (q - t)/2$.
3. When $q - t$ is odd, either

$$(a) |S \cap X| = |T \cap Y| = \left\lceil \frac{q-t}{2} \right\rceil, |S \cap Y| = |T \cap X| = \left\lfloor \frac{q-t}{2} \right\rfloor \text{ or}$$

$$(b) |S \cap X| = |T \cap Y| = \left\lfloor \frac{q-t}{2} \right\rfloor, |S \cap Y| = |T \cap X| = \left\lceil \frac{q-t}{2} \right\rceil.$$

Notice that each pair (C, C') contributes a matching (at most one edge per vertex) in the Kikuchi graph. The following result gives the size of this matching.

Fact 4.12 (Observation 3.7 of [\[HKM23\]](#)). *The Kikuchi graph defined in [Definition 4.11](#) has $N = \binom{2n}{\ell}$ vertices, and each pair (C, C') contributes a D -sized matching in the Kikuchi graph, where*

$$D := \binom{q-t}{\lfloor \frac{q-t}{2} \rfloor} \binom{2n - 2(q-t)}{\ell - (q-t)} \cdot 2^{\mathbf{1}_{(q-t \text{ is odd})}} \geq N \left(\frac{\ell}{2n} \right)^{q-t}.$$

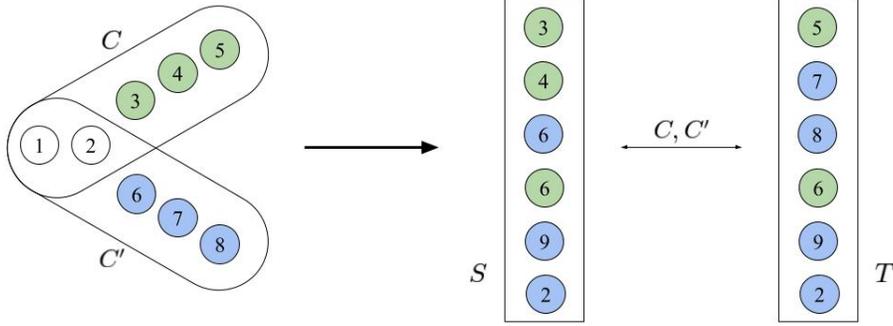


Figure 2: In this example, the hyperedges C, C' induce an edge between S, T . Here, $q = 5, t = 2$. Vertices of the form $(x, 1)$ are colored green, while vertices of the form $(x, 2)$ are colored blue. Note that $S \oplus T = (C \setminus Q_\theta) \times \{1\} \cup (C' \setminus Q_\theta) \times \{2\}$. Also note how S contains $2 = \lfloor \frac{q-t}{2} \rfloor$ green elements and $1 = \lfloor \frac{q-t}{2} \rfloor$ blue element from $S \oplus T$, while T contains $1 = \lfloor \frac{q-t}{2} \rfloor$ green element and $2 = \lfloor \frac{q-t}{2} \rfloor$ blue elements from $S \oplus T$.

We next define the *signed* Kikuchi matrix according to $b \in \{\pm 1\}^k$. We will use \bar{K} to denote unsigned matrices and K to denote signed matrices according to b .

Definition 4.13 (Signed Kikuchi Matrix). Assume the same setting as Definition 4.11. Let $N := \binom{2n}{\ell}$. Let $\bar{K}^{(C, C')} \in \{0, 1\}^{N \times N}$ be the matrix with rows/columns indexed by $\binom{[n] \times [2]}{\ell}$ and $\bar{K}^{(C, C')}(S, T) = \mathbf{1}\{S \xleftrightarrow{C, C'} T\}$, i.e., the adjacency matrix of the matching formed by (C, C') . For $i \neq j \in [k]$, define the unsigned matrix

$$\bar{K}_{i, j, t} := \sum_{\theta \in [p_t]} \sum_{\substack{C \in \mathcal{H}_i \cap \mathcal{H}_\theta^{(t)} \\ C' \in \mathcal{H}_j \cap \mathcal{H}_\theta^{(t)}}} \bar{K}^{(C, C')}.$$

Next, for $b \in \{\pm 1\}^k$ and a partition $L \sqcup R = [k]$, we define the signed matrix

$$K_{i, t} := \sum_{j \in R} b_j \bar{K}_{i, j, t}.$$

Here, we omit the dependence on b, L and R for simplicity.

We now observe that $f_{L, R, t}^{\mathcal{H}}$ is a quadratic form of $\sum_{i \in L} b_i K_{i, t}$.

Lemma 4.14. For any $x \in \{\pm 1\}^n$, let $z \in \{\pm 1\}^N$ for $N = \binom{2n}{\ell}$ be defined so that for any $S = S_1 \sqcup S_2$ where $|S| = \ell$, $S_1 \subseteq [n] \times \{1\}$, and $S_2 \subseteq [n] \times \{2\}$, we have $z_S := \prod_{u \in S_1} x_u \prod_{u \in S_2} x_u$. Then,

$$z^\top \left(\sum_{i \in L} b_i K_{i, t} \right) z = D \cdot f_{L, R, t}^{\mathcal{H}}(x).$$

Thus, $\text{val}(f_{L, R, t}^{\mathcal{H}}) \leq D^{-1} \|\sum_{i \in L} K_{i, t}\|_{\infty \rightarrow 1}$.

Proof. From [Fact 4.12](#), we have that $z^\top \bar{K}^{(C,C')} z = D \cdot x_{C \setminus Q_\theta} x_{C' \setminus Q_\theta}$. Thus,

$$D \cdot f_{L,R,t}^{\mathcal{H}}(x) = \sum_{i \in L, j \in R} b_i b_j \cdot z^\top \bar{K}_{i,j,t} z = z^\top \left(\sum_{i \in L} b_i K_{i,t} \right) z \leq \left\| \sum_{i \in L} b_i K_{i,t} \right\|_{\infty \rightarrow 1}. \quad \square$$

Spectral bound via Matrix Khintchine We will use the Matrix Khintchine inequality ([Fact 3.5](#)) to upper bound the spectral norm of $\sum_{i \in L} b_i K_{i,t}$. The left-right partitioning scheme ensures that the $K_{i,t}$'s are independent random matrices as i varies in L . Recall that $d_r = \max_{Q \in \binom{[n]}{r}} d_{\mathcal{H},Q}$ is the maximum co-degree of size- r subsets. The following key technical result shows that one can upper bound the maximum degree of $K_{i,t}$ according to d_t for all but a negligible fraction of vertices in $K_{i,t}$ for any $i \in L$.

Lemma 4.15. *Assume the same setting as [Definition 4.11](#) and suppose that $k \geq d_1 \geq \ell = n^{1-2/q} \log n$. Then $d_t \geq d_1 n^{-2(t-1)/q}$ if $t < q/2$ and $d_t \geq d_1 n^{-1+2/q}$ if $t > q/2$. Moreover, for every partition $L \sqcup R = [k]$ and every $K_{i,t}$ for $i \in L$, the following holds: there exists a constant W such that for all but a n^{-100q} fraction of vertices $S \in \binom{[n] \times [2]}{\ell} = V(K_{i,t})$, we have*

$$\deg_{K_{i,t}}(S) \leq \left(\frac{\ell}{n} \right)^{q-t} n d_t \cdot W^q,$$

where $\deg_{K_{i,t}}(S)$ denotes the maximum ℓ_1 norm of row S in $K_{i,t}$ over all b_j for $j \in R$.

Note that the lower bound on d_t in [Lemma 4.15](#) and $d_1 \geq n^{1-2/q}$ ensure that $\left(\frac{\ell}{n}\right)^{q-t} n d_t \geq 1$.

We will prove [Lemma 4.15](#) in [Section 4.2](#), using a key technical lemma ([Lemma 5.1](#)) that guarantees a “good” index t . Having the above ℓ_1 norm upper bound for most rows in $K_{i,t}$ allows us to easily obtain an upper bound on $\left\| \sum_{i \in L} b_i K_{i,t} \right\|_{\infty \rightarrow 1}$ using Matrix Khintchine (after removing the bad rows and columns).

Lemma 4.16. *We have $D^{-1} \left\| \sum_{i \in L} b_i K_{i,t} \right\|_{\infty \rightarrow 1} \leq \sqrt{k \ell} n d_t O(\log n)^{1/2}$.*

Proof. Fix any partition $L \sqcup R = [k]$. For every $i \in L$, let $\tilde{K}_{i,t}$ be obtained by zeroing out any row and column indexed by S such that $\deg_{K_{i,t}}(S) > \left(\frac{\ell}{n}\right)^{q-t} n d_t \cdot W^q$. Then, for any $i \in L$, $K_{i,t}$ and $\tilde{K}_{i,t}$ differ in at most $N \cdot k(n^{-100q} + e^{-\ell/4})$ rows and columns (by [Lemma 4.15](#)). Further, observe that the ℓ_1 norm of any row of $K_{i,t}$ is at most $|\mathcal{H}|^2 \ell^q \leq n^{3q}$. As a result,

$$\left\| \sum_{i \in L} b_i K_{i,t} \right\|_{\infty \rightarrow 1} \leq \left\| \sum_{i \in L} b_i \tilde{K}_{i,t} \right\|_{\infty \rightarrow 1} + 2N \left(n^{-100q} + e^{-\ell/4} \right) \cdot n^{3q} \leq \left\| \sum_{i \in L} b_i \tilde{K}_{i,t} \right\|_{\infty \rightarrow 1} + o(N), \quad (4)$$

where we use the naive ℓ_1 norm bound for the rows and columns removed from $K_{i,t}$.

Next, observe that $\left\| \sum_{i \in L} b_i \tilde{K}_{i,t} \right\|_{\infty \rightarrow 1} \leq N \left\| \sum_{i \in L} b_i \tilde{K}_{i,t} \right\|_2$. Applying Matrix Khintchine inequality ([Fact 3.5](#)) and the triangle inequality:

$$\mathbb{E}_{b \in \{\pm 1\}^k} \left[\left\| \sum_{i \in L} b_i \tilde{K}_{i,t} \right\|_2 \right] \leq \sigma \sqrt{2 \log N}, \text{ where } \sigma^2 = \left\| \sum_{i \in L} \tilde{K}_{i,t}^2 \right\|_2 \leq |L| \cdot \max_{i \in L} \left\| \tilde{K}_{i,t} \right\|_2^2.$$

Now, since $\widetilde{K}_{i,t}$ has small degrees, we can upper bound its spectral norm by the largest ℓ_1 norm of any row, i.e., $\max_S \deg_{\widetilde{K}_{i,t}}(S)$. So, we have $\sigma \leq \sqrt{k} \cdot (\ell/n)^{q-t} n d_t \cdot W^q$.

Plugging back in (4), we have:

$$\mathbb{E}_{b \in \{\pm 1\}^k} \left[\left\| \sum_{i \in L} b_i K_{i,t} \right\|_{\infty \rightarrow 1} \right] \leq O(N \sqrt{\log N}) \sqrt{k} \left(\frac{\ell}{n} \right)^{q-t} n d_t \cdot W^q + o(N). \quad (5)$$

We verify that the first term dominates the second above. By assumption, $k \geq d_1 \geq \ell = n^{1-2/q} \log n$, and by Lemma 4.15, we have $d_t \geq d_1 n^{-2(t-1)/q}$, thus $\sqrt{k} \ell \left(\frac{\ell}{n} \right)^{q-t} \cdot n d_t \geq d_t n^{2(t-1)/q} \geq 1$. Then, using Fact 4.12, we have $N/D \leq (2n/\ell)^{q-t}$, and $\log N \leq O(\ell \log n)$. Thus, multiplying (5) by D^{-1} gives the desired upper bound $\sqrt{k} \ell n d_t O(\log n)^{1/2}$. \square

With the Cauchy-Schwarz trick (Lemma 4.10) and the bound on $\|\sum_{i \in L} b_i K_{i,t}\|_{\infty \rightarrow 1}$ (Lemma 4.16), Lemma 4.6 is almost an immediate corollary.

Proof of Lemma 4.6. Using Lemma 4.10, we have that

$$\text{val}(\Psi_b)^2 \leq \frac{2|\mathcal{H}|^2}{d_t} + \frac{8|\mathcal{H}|}{d_t} \mathbb{E}_{L,R} \left[\text{val} \left(f_{L,R,t}^{\mathcal{H}} \right) \right]. \quad (6)$$

By Lemma 4.14, we can write $f_{L,R,t}^{\mathcal{H}}(x)$ as D^{-1} times a quadratic form of the signed Kikuchi matrix $\sum_{i \in L} b_i K_{i,t}$ from Definition 4.13. Then, by Lemma 4.16,

$$\mathbb{E}_{L,R} \mathbb{E}_b [\text{val}(f_{L,R,t})] \leq \frac{1}{D} \mathbb{E}_{L,R} \mathbb{E}_b \left[\left\| \sum_{i \in L} b_i K_{i,t} \right\|_{\infty \rightarrow 1} \right] \leq \sqrt{k} \ell \cdot n d_t O(\log n)^{1/2}.$$

Plugging back this bound in (6) completes the proof. \square

4.2 Proof of Lemma 4.15: bounding heavy vertices

In this section, we prove Lemma 4.15. First, we observe that the degree of any vertex S of $K_{i,t}$ can be upper bounded by a polynomial Deg in the 0-1 indicator of the vertex S . We then use a polynomial concentration inequality (Fact 3.6) to analyze the concentration of Deg as a function of a uniformly random s . The concentration inequality relies on expected partial derivatives of Deg that have natural combinatorial interpretations and can be described as a function of the co-degree sequence $d_1 \geq d_2 \geq \dots \geq d_q$ of the input hypergraph $\mathcal{H} = (\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_k)$. In particular, the required bounds are precisely those such that d_t is a good index as defined in the technical overview (Section 2.3), which we restate below.

Definition (Restatement of Definition 2.6). Let $d_1 \geq d_2 \geq \dots \geq d_q$. We say that an index t is *good* with respect to the tuple (d_1, d_2, \dots, d_q) if the following conditions hold:

- (1) $d_r/d_t \leq n^{1-\frac{2r}{q}}$ for every $1 \leq r \leq \lceil \frac{q-t}{2} \rceil$,
- (2) $d_r/d_t \leq n^{-\frac{2}{q}(r-t) + \frac{1}{q}(t-1 \text{ (t even)})}$ for every $t \leq r \leq \lfloor \frac{q+t}{2} \rfloor$,

(3) If $t < q/2$, then, $d_t \geq d_1 n^{-\frac{2}{q}(t-1)}$, and, if $t > q/2$, then, $d_t \geq d_1 n^{-1+2/q}$.

We now prove [Lemma 4.15](#).

Definition 4.17 (The Degree Polynomial). Assume the same setting as [Definition 4.11](#), and fix an $i \in [k]$. Let $s, s' \in \{0, 1\}^n$ be indicator vectors of the sets $S \cap ([n] \times \{1\})$ and $S \cap ([n] \times \{2\})$, respectively. Consider the polynomial $\text{Deg}(s, s')$ defined so that

$$\text{Deg}(s, s') = \sum_{\theta \in [p_i]} \sum_{\substack{C \in \mathcal{H}_i \cap \mathcal{H}_\theta^{(t)} \\ C' \neq C \in \mathcal{H}_\theta^{(t)}}} \sum_{\substack{R \in \binom{C \setminus Q_\theta}{\lfloor (q-t)/2 \rfloor} \\ R' \in \binom{C' \setminus Q_\theta}{\lfloor (q-t)/2 \rfloor}}} s_R s'_{R'}$$

if t is even, and,

$$\text{Deg}(s, s') = \sum_{\theta \in [p_i]} \sum_{\substack{C \in \mathcal{H}_i \cap \mathcal{H}_\theta^{(t)} \\ C' \neq C \in \mathcal{H}_\theta^{(t)}}} \left(\sum_{\substack{R \in \binom{C \setminus Q_\theta}{\lfloor (q-t)/2 \rfloor} \\ R' \in \binom{C' \setminus Q_\theta}{\lfloor (q-t)/2 \rfloor}}} s_R s'_{R'} + \sum_{\substack{R \in \binom{C \setminus Q_\theta}{\lfloor (q-t)/2 \rfloor} \\ R' \in \binom{C' \setminus Q_\theta}{\lfloor (q-t)/2 \rfloor}}} s_R s'_{R'} \right)$$

if t is odd.

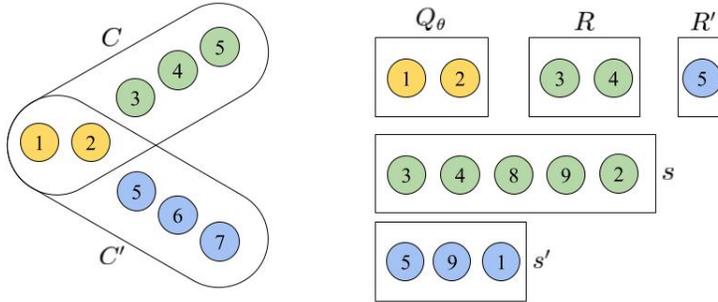


Figure 3: Example of a contributing monomial to $\text{Deg}(s, s')$. Here, $q = 5, t = 2$. As in [Fig. 2](#), vertices of the form $(x, 1)$ are colored green, while vertices of the form $(x, 2)$ are colored blue. The elements of Q_θ are colored yellow. Note that $q - t = 3$ is odd, and S indicated by (s, s') contains $2 = \lfloor \frac{q-t}{2} \rfloor$ green elements from $(C \setminus Q_\theta) \times \{1\}$, and $1 = \lfloor \frac{q-t}{2} \rfloor$ blue element from $(C' \setminus Q_\theta) \times \{2\}$.

Remark 4.18 (Understanding $\text{Deg}(s, s')$). In either case, note that $\text{Deg}(s, s')$ is a polynomial of degree $q - t$. For any (s, s') that indicates a subset of size ℓ in $[n] \times [2]$, $\text{Deg}(s, s')$ upper bounds the degree of the vertex S in $K_{i,t}$. This is because for each pair (C, C') that contributes a D -matching to $K_{i,t}$, there is a corresponding monomial that evaluates to at least 1 in $\text{Deg}(s, s')$. We note that in $\text{Deg}(s, s')$, we have contributions from all C, C' such that $C \in \mathcal{H}_i$ and C, C' are in some piece $\mathcal{H}_\theta^{(t)}$ of

the partition promised by the approximate strong regularity of \mathcal{H} . In particular, $\text{Deg}(s, s')$ contains contributions from pairs (C, C') such that $C' \in \mathcal{H}_j$ for some $j \in L$ even though such a pair does not contribute any edges to $K_{i,t}$. This, of course, can only hurt us in our upper bounds on the degree.

To understand the distribution of degrees of the vertices S of $K_{i,t}$, we must analyze the distribution of $\text{Deg}(s, s')$ as (s, s') vary over indicators of subsets of size exactly ℓ . As in [GKM22], we use a coupling argument to instead work with the ℓ/n -biased product distribution over $\{0, 1\}^{2n}$.

Fact 4.19 (Similar to Lemma 6.20 of [GKM22]). *Let U be the $\frac{\ell}{n}$ -biased product distribution on (s, s') . Let U' be the uniform distribution on $(s, s') \in \{0, 1\}^{2n}$ indicating sets of size exactly ℓ . Then for any λ , we have*

$$\Pr_{(s,s') \sim U'} [\text{Deg}(s, s') \leq \lambda] \geq \Pr_{(s,s') \sim U} [\text{Deg}(s, s') \leq \lambda] - e^{-\ell/4}.$$

Our plan now is to apply the polynomial concentration inequality to understand the behavior of the polynomial $\text{Deg}(s, s')$. The next lemma bounds the expectations of partial derivatives of $\text{Deg}(s, s')$ towards this application.

Lemma 4.20. *Suppose $(s, s') \in \{0, 1\}^{2n}$ has the $\frac{\ell}{n}$ -biased product distribution. Let Z_1 be any multiset of elements from s_1, s_2, \dots, s_n , and let Z_2 be any multiset of elements from s'_1, s'_2, \dots, s'_n . Then,*

$$\mathbb{E} \left[\left(\prod_{j \in Z_2} \partial_{s'_j} \right) \text{Deg}(s, s') \right] \leq 2^q \left(\frac{\ell}{n} \right)^{q-t-|Z_2|} d_{|Z_2|},$$

and,

$$\mathbb{E} \left[\left(\prod_{i \in Z_1} \partial_{s_i} \prod_{j \in Z_2} \partial_{s'_j} \right) \text{Deg}(s, s') \right] \leq 2^q \left(\frac{\ell}{n} \right)^{q-t-|Z_1|-|Z_2|} d_{|Z_2|+t}.$$

In particular, if $\gamma_i = \log_n d_i$ for all $i \in [q]$ and $\gamma_1 \geq 1 - 2/q$, and our hypergraph is t -approximately strongly regular,

$$\mathbb{E} \left[\left(\prod_{i \in Z_1} \partial_{s_i} \prod_{j \in Z_2} \partial_{s'_j} \right) \text{Deg}(s, s') \right] \leq 2^q \left(\frac{\ell}{n} \right)^{q-t} \frac{nd_t}{(\log n)^{|Z_1|+|Z_2|}}.$$

Proof. Without loss of generality, we will assume that Z_1, Z_2 are sets (since differentiating twice with respect to the same variable immediately makes the polynomial vanish).

When $Z_1 = Z_2 = \emptyset$, we observe that

$$\mathbb{E} [\text{Deg}(s, s')] \leq 2^{2q} \left(\frac{\ell}{n} \right)^{q-t} nd_t. \quad (7)$$

To see why, observe that the number of (C, C') that contribute monomials to $\text{Deg}(s, s')$ can be upper bounded by nd_t since there are $\leq n$ possible choices of $C \in \mathcal{H}_i$. Given $C \in \mathcal{H}_i$, there is at most one $Q_\theta \subseteq C$ such that $C \in \mathcal{H}_\theta^{(t)}$, and any Q_θ appears in at most d_t different C' in \mathcal{H} . Further, each pair (C, C') contributes at most 2^{2q} monomials corresponding R, R' that range over subsets of C, C'

respectively. Finally, each monomial has degree exactly $q - t$ and thus has expectation $(\ell/n)^{q-t}$ under the ℓ/n -biased product distribution.

Let us start with the case when $Z_1 = \emptyset$ and $Z_2 \neq \emptyset$. Clearly, $|Z_2| \leq \lceil (q-t)/2 \rceil$. Our overall strategy is similar to the above – we will count the number of pairs (C, C') that can contribute a non-zero monomial to $(\prod_{j \in Z_2} \partial_{s'_j}) \text{Deg}(s, s')$. First, each (C, C') contributes at most 2^{2q} different monomials, and each monomial has an expectation of $(\ell/n)^{q-t-|Z_2|}$. To count the number of pairs, the only monomials that survive after taking the partial derivatives are ones where $C' \supseteq Z_2$. Thus, the number of possible C' is at most $d_{|Z_2|}$. Given C' , we have a unique θ such that $C' \in \mathcal{H}_\theta^{(t)}$, and given Q_θ , there is at most one $C \in \mathcal{H}_i$ (since \mathcal{H}_i is a matching) such that $C \supseteq Q_\theta$. Thus, we obtain an upper bound of $d_{|Z_2|} 2^q (\ell/n)^{q-t-|Z_2|}$ in this case.

Now, $(\ell/n)^{q-t-|Z_2|} d_{|Z_2|} \leq (\ell/n)^{q-t} n d_t \cdot (\log n)^{-|Z_2|}$ if and only if $d_{|Z_2|} \leq (\ell/n)^{|Z_2|} n d_t \cdot (\log n)^{-|Z_2|}$. Since $\ell = n^{1-2/q} \log n$, this is true if and only if $d_{|Z_2|} \leq n^{1-2|Z_2|/q} d_t$ or $\gamma_{|Z_2|} \leq 1 - 2/q + \gamma_t$, a condition that is satisfied since t satisfies the first condition in [Definition 2.6](#).

Finally, consider the case when Z_1 is non-empty. Since the rest of the argument is similar, let's focus on counting the number of pairs (C, C') such that $C \supseteq Z_1$. Clearly, there is at most one $C \in \mathcal{H}_i$ that can contain Z_1 . Given C , we know the unique $\mathcal{H}_\theta^{(t)}$ that contains it and thus know Q_θ . So we are looking for C' that contain $Q_\theta \cup Z_2$ – a set of size $t + |Z_2|$. Thus, the number of such C' is at most $d_{t+|Z_2|}$.

Now, $(\ell/n)^{q-t-|Z_1|-|Z_2|} d_{t+|Z_2|} \leq (\ell/n)^{q-t} n d_t \cdot (\log n)^{-(|Z_1|+|Z_2|)}$ for $\ell = n^{1-2/q} \log n$ if and only if $d_{t+|Z_2|} \leq n^{1-2(|Z_1|+|Z_2|)/q} d_t$. If $|Z_2| \neq \lceil (q-t)/2 \rceil$, we set $|Z_1|$ to its maximum of $\lceil (q-t)/2 \rceil$ to see this is satisfied if $d_{t+|Z_2|} \leq n^{-2(|Z_2|-t)/q} n^{-(1/q)\mathbf{1}_{(q-t \text{ is odd})}} d_t$, which is true by the second condition in [Definition 2.6](#). Otherwise, we let $g = |Z_2| - 1$ and $f = |Z_1| + 1$, and since the degrees are decreasing, we get $d_{t+|Z_2|} \leq d_{t+g} \leq n^{1-2(f+g)/q} d_t = n^{1-2(|Z_1|+|Z_2|)/q} d_t$. \square

Proof of [Lemma 4.15](#). Given the degree sequence $d_1 \geq d_2 \geq \dots \geq d_q$, set $\gamma_i = \log_n d_i$ for every $1 \leq i \leq q$. Then, notice that the γ_i 's satisfy the hypothesis of [Lemma 5.1](#) since we set k large enough so that $d_1 \geq |\mathcal{H}|/n \geq n^{1-2/q}$. Since our hypergraph is t -approximately strongly regular, property (3) of [Definition 2.6](#) implies that $d_t \geq d_1 n^{-2(t-1)/q}$ if $t < q/2$ and $d_t \geq d_1 n^{-1+2/q}$ if $t > q/2$, i.e., the desired lower bound for d_t stated in [Lemma 4.15](#). Note that using the assumption $d_1 \geq \ell = n^{1-2/q}$, it follows that the upper bound on $\mathbb{E}[\text{Deg}(s, s')]$ in (7) is at least $n^{-\frac{2}{q}(q-t)} n d_t \geq 1$.

We apply [Fact 3.6](#) to $\text{Deg}(s, s')$ over the (ℓ/n) -biased product distribution over $\{0, 1\}^n$. Given [Lemma 4.20](#), we immediately obtain that for large enough absolute constant $W > 0$,

$$\Pr \left[\text{Deg}(s, s') \geq \left(\frac{\ell}{n} \right)^{q-t} n d_t W^q \right] \leq n^{-100q}.$$

Combining with [Fact 4.19](#) finishes the proof. \square

5 Approximate Strong Regularity Decomposition

In this section, we prove [Lemma 4.7](#). We use the following technical (but rather crucial!) result, which we prove in [Section 5.1](#), to show that a good index as defined in [Definition 2.6](#) always exists.

Lemma 5.1. Let $q \geq 3$ be an odd integer, and let $\gamma_1 \geq \gamma_2 \geq \dots \geq \gamma_{q-1} \geq \gamma_q \geq 0$ be real numbers. Then there exists $t \in [q]$ satisfying the following properties:

(1) $\gamma_r \leq \gamma_t + 1 - \frac{2r}{q}$ for all $1 \leq r \leq \lfloor \frac{q-t}{2} \rfloor$.

(2) $\gamma_r \leq \gamma_t - \frac{2(r-t)}{q} + \frac{1}{q}(t - \mathbf{1}(t \text{ even}))$ for all $t < r \leq \lfloor \frac{q+t}{2} \rfloor$.

(3) If $t < \frac{q}{2}$, then $\gamma_t - (1 - \frac{2t}{q}) \geq \gamma_1 - (1 - \frac{2}{q})$. If $t > \frac{q}{2}$, then $\gamma_t > \gamma_1 - (1 - \frac{2}{q})$.

Thus, by taking $\gamma_i = \log_n d_i$ for all $i \in [q]$, every q -uniform hypergraph contains a good index.

Now, we restate the approximate strong regularity decomposition lemma. We refer the reader to [Section 2.6](#) for an overview.

Lemma (Restatement of [Lemma 4.7](#)). Let \mathcal{H} be a q -uniform hypergraph on $[n]$. Let $\eta > 0$. Then, there is a collection $\{\mathcal{H}^\pi : \pi \in P\}$, with $|P| \leq q \lceil \log_2 |\mathcal{H}| \rceil + 1$, of hypergraphs such that:

1. For every $\pi \in P$, there exists $t \in [q]$ such that \mathcal{H}^π is t -approximately strongly regular.
2. $\sqcup_{\pi \in P} \mathcal{H}^\pi \subseteq \mathcal{H}$,
3. $|\sqcup_{\pi \in P} \mathcal{H}^\pi| \geq (1 - \eta)|\mathcal{H}|$,
4. $|\mathcal{H}^\pi| \geq \frac{\eta^{|\mathcal{H}|}}{q \lceil \log_2 |\mathcal{H}| \rceil + 1}$ for all $\pi \in P$.

We prove [Lemma 4.7](#) by describing an algorithm to produce the required partition, which uses as a subroutine the following [Algorithm 1](#).

Algorithm 1.

Input: A hypergraph \mathcal{H} on vertices $[n]$, and an integer $t \geq 1$.

Output: Two hypergraphs $(\mathcal{H}^{\text{regular}}, \mathcal{H}^{\text{residual}})$ where $\mathcal{H} = \mathcal{H}^{\text{regular}} \sqcup \mathcal{H}^{\text{residual}}$.

Operation:

1. Set $\mathcal{H}^{\text{regular}} := \emptyset, \mathcal{H}^{\text{residual}} := \mathcal{H}, D := d_{\mathcal{H}, t}$.
2. While $d_{\mathcal{H}^{\text{residual}}, t} \geq D/2$:
 - (a) Consider $Q \in \binom{V(\mathcal{H})}{t}$ such that $|\mathcal{H}_{|Q}^{\text{residual}}|$ is the maximum. Then perform $\mathcal{H}^{\text{regular}} \leftarrow \mathcal{H}^{\text{regular}} \cup \mathcal{H}_{|Q}^{\text{residual}}, \mathcal{H}^{\text{residual}} \leftarrow \mathcal{H}^{\text{residual}} \setminus \mathcal{H}_{|Q}^{\text{residual}}$.
3. Output $(\mathcal{H}^{\text{regular}}, \mathcal{H}^{\text{residual}})$.

Remark 5.2. Note that the output of this procedure depends on which Q we choose in [Item 2a](#) if there are multiple Q such that $|\mathcal{H}_{|Q}^{\text{residual}}|$ is maximal. Any arbitrary choice of Q is fine for our application.

In [Algorithm 1](#), observe that $d_{\mathcal{H}^{\text{regular}},r} \leq d_{\mathcal{H},r}$ for all $r \in [q]$, since $\mathcal{H}^{\text{regular}} \subseteq \mathcal{H}$. Also observe that if $d_{\mathcal{H},t} \geq 2$, then $d_{\mathcal{H}^{\text{residual}},t} < d_{\mathcal{H},t}/2 < d_{\mathcal{H},t} = d_{\mathcal{H}^{\text{regular}},t}$.

Our main decomposition algorithm iteratively applies [Algorithm 1](#).

Algorithm 2 (Approximate Strong Regularity Decomposition).

Input: \mathcal{H} , a q -uniform hypergraph on $[n]$.

Output: \mathcal{T} , a collection of hypergraphs partitioning \mathcal{H} .

Operation:

1. Define $\mathcal{H}^{\text{curr}} := \mathcal{H}, \mathcal{T} := ()$.
2. While $\mathcal{H}^{\text{curr}} \neq \emptyset$:
 - (a) Write $\gamma_i := \log_n d_{\mathcal{H}^{\text{curr}},i}$ for all $i \in [q]$, and let $t \in [q]$ be the index returned by invoking [Lemma 5.1](#) on $\{\gamma_i\}_{i \in [q]}$.
 - (b) Apply [Algorithm 1](#) on $\mathcal{H}^{\text{curr}}$ and t to get hypergraphs $(\mathcal{H}^{\text{curr}})^{\text{regular}}, (\mathcal{H}^{\text{curr}})^{\text{residual}}$. Write $\mathcal{H}^{\text{curr}} \leftarrow (\mathcal{H}^{\text{curr}})^{\text{residual}}, \mathcal{T} \leftarrow (\mathcal{T}, (\mathcal{H}^{\text{curr}})^{\text{regular}})$.
3. Return \mathcal{T} .

In the above algorithm, $(\mathcal{T}, \mathcal{H}')$ stands for appending the hypergraph \mathcal{H}' at the end of the tuple \mathcal{T} . We record some small propositions about the above algorithm.

Proposition 5.3. *Suppose \mathcal{T} is the output of [Algorithm 2](#) on a hypergraph \mathcal{H} . Let \mathcal{H}' be any but the last entry of \mathcal{T} . Let t be the index in the invocation of [Item 2b](#) which produced \mathcal{H}' . Then:*

1. \mathcal{H}' is t -approximately strongly regular.
2. Write $\gamma'_i := \log_n d_{\mathcal{H}',t}$. Then $\{\gamma'_i\}$ satisfies the conditions of [Lemma 5.1](#) with the “good index” being t .

Proof. Consider any $C \in \mathcal{H}'$. During the iteration of [Algorithm 1](#) on inputs $(\mathcal{H}^{\text{curr}}, t)$ where C was added to \mathcal{H}' , suppose that $\mathcal{H}^{\text{curr}} = \mathcal{H}''$. Note that $d_{\mathcal{H}',t} = d_{\mathcal{H}'',t}$. Then there exists $Q \in \binom{V(\mathcal{H})}{t}$ such that $Q \subseteq C \in \mathcal{H}''_{|Q}$ and $|\mathcal{H}''_{|Q}| \geq d_{\mathcal{H}',t}/2$, and $|\mathcal{H}''_{|Q}| \leq d_{\mathcal{H}',t}$ holds by definition. Thus \mathcal{H}^j can be written as the union of these $\mathcal{H}''_{|Q}$'s, where we enumerate arbitrarily by $[p_t]$.

The second statement follows by noting that $\gamma'_i = \gamma_t$, and $\gamma'_i \leq \gamma_i$. \square

Proposition 5.4. *For the above algorithm, $|\mathcal{T}| \leq q \lceil \log_2 |\mathcal{H}| \rceil + 1$.*

Proof. Write $\mathcal{H}' = \mathcal{H}^{\text{curr}}$ at any step of this algorithm, and let \mathcal{H}'' be $\mathcal{H}^{\text{curr}}$ after we apply [Item 2b](#). Note that for some $t \in [q]$, we had that $d_{\mathcal{H}'',t} < d_{\mathcal{H}',t}/2$. Now, assume for the sake of contradiction that $|\mathcal{T}| > q \lceil \log_2 |\mathcal{H}| \rceil + 1$. Then by the pigeonhole principle, there must be some $t \in [q]$ which was the index invoked in [Item 2b](#) more than $\lceil \log_2 |\mathcal{H}| \rceil$ times. Now, let \mathcal{H}''' be the last entry of \mathcal{T} . Then we have $d_{\mathcal{H}''',t} < d_{\mathcal{H},t}/2^{\lceil \log_2 |\mathcal{H}| \rceil} \leq 1$, where the last inequality follows since $d_{\mathcal{H},t} \leq |\mathcal{H}|$. But $d_{\mathcal{H}''',t} < 1$ is a contradiction, and the result follows. \square

We can finally prove [Lemma 4.7](#).

Proof of Lemma 4.7. Run [Algorithm 2](#) on \mathcal{H} to obtain \mathcal{T} . Write $T := |\mathcal{T}|$. By [Proposition 5.4](#), we know that $T \leq q \lceil \log_2 |\mathcal{H}| \rceil + 1$.

Write $\{\mathcal{H}^\pi : \pi \in P\}$ ⁶ to be the hypergraphs in \mathcal{T} with size $\geq \frac{\eta|\mathcal{H}|}{T}$. Note that the number of hyperedges of \mathcal{H} not contained in $\{\mathcal{H}^\pi : \pi \in P\}$ has size $< \frac{\eta|\mathcal{H}|}{T} \cdot T = \eta|\mathcal{H}|$. [Lemma 4.7](#) now follows from [Propositions 5.3](#) and [5.4](#). \square

5.1 Proof of [Lemma 5.1](#): a good index exists

In this section, we prove [Lemma 5.1](#). We first show that for inequality (1) in [Lemma 5.1](#), we may assume $r < t$.

Lemma 5.5. *In [Lemma 5.1](#), when $t < r \leq \left\lceil \frac{q-t}{2} \right\rceil$, we have (2) \implies (1).*

Proof. Since $t < \left\lceil \frac{q-t}{2} \right\rceil \leq \frac{q-t}{2}$, we have $3t \leq q$. Then, $-\frac{2(r-t)}{q} + \frac{1}{q}(t - \mathbf{1}(t \text{ even})) \leq \frac{1}{q}(3t - 2r) \leq 1 - \frac{2r}{q}$. This means that (2) is a tighter upper bound than (1). \square

We also show the following useful lemma.

Lemma 5.6. *If $t > q/2$, then inequality (2) is always satisfied.*

Proof. This follows from the descending properties of the γ_t 's. For any r such that $t < r \leq \left\lceil \frac{q+t}{2} \right\rceil$, we have $-\frac{2(r-t)}{q} + \frac{1}{q}(t - \mathbf{1}(t \text{ even})) \geq -\frac{2}{q} \cdot \frac{q-t-\mathbf{1}(t \text{ even})}{2} + \frac{1}{q}(t - \mathbf{1}(t \text{ even})) = \frac{1}{q}(2t - q) > 0$. Thus, since $r > t$, we have $\gamma_r \leq \gamma_t$, and inequality (2) holds. \square

Proof of Lemma 5.1. We use the following algorithm to find the desired index.

1. Set $t_0 < q/2$ to be the index that maximizes $\gamma_{t_0} + \frac{2t_0}{q}$.
2. If there is some $t > q/2$ such that that γ_t violates inequality (2) for γ_{t_0} , then return the t that maximizes $\gamma_t + \frac{2t}{q}$ over the violating indices.

First, we claim that for all $r < t_0$, γ_r satisfies inequality (1) for γ_{t_0} . This follows from the fact that t_0 maximizes $\gamma_{t_0} + \frac{2t_0}{q}$, thus we have

$$\gamma_r + \frac{2r}{q} \leq \gamma_{t_0} + \frac{2t_0}{q} \implies \gamma_r \leq \gamma_{t_0} + \frac{2t_0}{q} - \frac{2r}{q} < \gamma_{t_0} + 1 - \frac{2r}{q},$$

since $t_0 < q/2$. Moreover, we show that no index $r \in [t_0 + 1, \frac{q-1}{2}]$ violates inequality (2) for γ_{t_0} . Again, we use $\gamma_r \leq \gamma_{t_0} - \frac{2(r-t_0)}{q}$, which immediately implies that inequality (2) is satisfied. Therefore, if there is an index t that violates inequality (2) for γ_{t_0} , we must have $t > q/2$.

Next, [Lemma 5.6](#) states that inequality (2) is satisfied for any $r > t$ since $t > q/2$. Thus, we now verify that inequality (1) is satisfied for all $r \leq \left\lceil \frac{q-t}{2} \right\rceil$. Since t violates inequality (2) for γ_{t_0} , we must

⁶ P is just an indexing set, and π is the corresponding indexing variable.

have $t \leq \left\lfloor \frac{q+t_0}{2} \right\rfloor = \frac{q+t_0 - \mathbf{1}(t_0 \text{ even})}{2}$ (as q is odd) and $\gamma_t > \gamma_{t_0} - \frac{2(t-t_0)}{q} + \frac{1}{q}(t_0 - \mathbf{1}(t_0 \text{ even}))$, which imply that $\gamma_t > \gamma_{t_0} - 1 + \frac{2t_0}{q}$. Combined with $\gamma_r \leq \gamma_{t_0} + \frac{2(t_0-r)}{q}$, we get

$$\gamma_r \leq \gamma_t + 1 - \frac{2t_0}{q} + \frac{2(t_0-r)}{q} = \gamma_t + 1 - \frac{2r}{q},$$

thus satisfying inequality (1).

Finally, we show lower bounds on γ_{t_0} or γ_t . First, we have $\gamma_{t_0} + \frac{2t_0}{q} \geq \gamma_1 + \frac{2}{q}$ since $t_0 < q/2$. But $\gamma_{t_0} + \frac{2t_0}{q} \geq \gamma_1 + \frac{2}{q} \implies \gamma_{t_0} \geq 1 - \frac{2t_0}{q} + (\gamma_1 - (1 - \frac{2}{q}))$, as desired. For γ_t (if we reach step 2 of the algorithm), we have that $\gamma_t > \gamma_{t_0} - 1 + \frac{2t_0}{q} \geq \gamma_1 - (1 - \frac{2}{q})$. This completes the proof. \square

Acknowledgments

We thank Matija Bucić for helpful discussions on related topics. A.B. thanks Rohit Agarwal for helpful discussions and comments.

References

- [ABS⁺23] Noga Alon, Matija Bucić, Lisa Sauermann, Dmitrii Zakharov, and Or Zamir. Essentially tight bounds for rainbow cycles in proper edge-colourings. *arXiv preprint arXiv:2309.04460*, 2023. (pg. 5)
- [AG24] Omar Alrabiah and Venkatesan Guruswami. Near-tight bounds for 3-query locally correctable binary linear codes via rainbow cycles. *CoRR*, abs/2404.05864, 2024. (pg. 1)
- [AGKM23] Omar Alrabiah, Venkatesan Guruswami, Pravesh K. Kothari, and Peter Manohar. A near-cubic lower bound for 3-query locally decodable codes from semirandom CSP refutation. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 1438–1448. ACM, 2023. (pg. 0, 1, 2, 3, 4, 5, 6, 7, 8, 12, 13, 17)
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998. (pg. 1)
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of np. *Journal of the ACM (JACM)*, 45(1):70–122, 1998. (pg. 1)
- [BCG20] Arnab Bhattacharyya, L Sunil Chandran, and Suprovat Ghoshal. Combinatorial Lower Bounds for 3-Query LDCs. In *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151, page 85. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2020. (pg. 1)

- [BDYW11] Boaz Barak, Zeev Dvir, Amir Yehudayoff, and Avi Wigderson. Rank bounds for design matrices with applications to combinatorial geometry and locally correctable codes. In *Proceedings of the 43rd ACM Symposium on Theory of Computing, STOC 2011, San Jose, CA, USA, 6-8 June 2011*, pages 519–528. ACM, 2011. (pg. 2)
- [CGW10] Victor Chen, Elena Grigorescu, and Ronald de Wolf. Efficient and error-correcting data structures for membership and polynomial evaluation. In *27th International Symposium on Theoretical Aspects of Computer Science, STACS 2010, March 4–6, 2010, Nancy, France*, volume 5 of *LIPICs*, pages 203–214. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2010. (pg. 1)
- [DG16] Zeev Dvir and Sivakanth Gopi. 2-server PIR with subpolynomial communication. *Journal of the ACM (JACM)*, 63(4):1–15, 2016. (pg. 1)
- [DLS13] Shagnik Das, Choongbum Lee, and Benny Sudakov. Rainbow Turán problem for even cycles. *European Journal of Combinatorics*, 34(5):905–915, 2013. (pg. 5)
- [DS05] Zeev Dvir and Amir Shpilka. Locally decodable codes with 2 queries and polynomial identity testing for depth 3 circuits. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing, Baltimore, MD, USA, May 22–24, 2005*, pages 592–601. ACM, 2005. (pg. 1)
- [Dvi10] Zeev Dvir. On matrix rigidity and locally self-correctable codes. In *Proceedings of the 25th Annual IEEE Conference on Computational Complexity, CCC 2010, Cambridge, Massachusetts, USA, June 9–12, 2010*, pages 291–298. IEEE Computer Society, 2010. (pg. 1)
- [Efr09] Klim Efremenko. 3-query locally decodable codes of subexponential length. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 - June 2, 2009*, pages 39–44. ACM, 2009. (pg. 1, 3, 7)
- [GKM22] Venkatesan Guruswami, Pravesh K. Kothari, and Peter Manohar. Algorithms and certificates for Boolean CSP refutation: smoothed is no harder than random. In *STOC '22: 54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20–24, 2022*, pages 678–689. ACM, 2022. (pg. 0, 1, 2, 3, 4, 6, 7, 8, 12, 14, 17, 26)
- [GKST06] Oded Goldreich, Howard Karloff, Leonard J Schulman, and Luca Trevisan. Lower bounds for linear locally decodable codes and private information retrieval. *Computational Complexity*, 15(3):263–296, 2006. (pg. 0, 1, 2, 6)
- [HKM23] Jun-Ting Hsieh, Pravesh K. Kothari, and Sidhanth Mohanty. A simple and sharper proof of the hypergraph moore bound. *ACM-SIAM Symposium on Discrete Algorithms, SODA*, 2023. (pg. 0, 1, 2, 3, 4, 6, 7, 8, 12, 17, 21)
- [HKM⁺24] Jun-Ting Hsieh, Pravesh K Kothari, Sidhanth Mohanty, David Munhá Correia, and Benny Sudakov. Small Even Covers, Locally Decodable Codes and Restricted Subgraphs

- of Edge-Colored Kikuchi Graphs. *arXiv preprint arXiv:2401.11590*, 2024. (pg. [0](#), [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [12](#))
- [IK04] Yuval Ishai and Eyal Kushilevitz. On the hardness of information-theoretic multiparty computation. In *Advances in Cryptology - EUROCRYPT 2004, International Conference on the Theory and Applications of Cryptographic Techniques, Interlaken, Switzerland, May 2-6, 2004, Proceedings*, volume 3027 of *Lecture Notes in Computer Science*, pages 439–455. Springer, 2004. (pg. [1](#))
- [Jan20] Oliver Janzer. Rainbow Turán number of even cycles, repeated patterns and blow-ups of cycles. *Israel Journal of Mathematics*, 253:813–840, 2020. (pg. [5](#))
- [JM24] Oliver Janzer and Peter Manohar. A $k^{\frac{q}{q-2}}$ Lower Bound for Odd Query Locally Decodable Codes from Bipartite Kikuchi Graphs. In *Personal communication*, 2024. (pg. [2](#))
- [JS24] Oliver Janzer and Benny Sudakov. On the turán number of the hypercube. *Forum of Mathematics, Sigma*, 12:e38, 2024. (pg. [5](#))
- [KLLT24] Jaehoon Kim, Joonkyung Lee, Hong Liu, and Tuan Tran. Rainbow Cycles in Properly Edge-Colored Graphs. *Combinatorica*, 44(4):909–919, Aug 2024. (pg. [5](#))
- [KM23] Pravesh K. Kothari and Peter Manohar. An exponential lower bound for linear 3-query locally correctable codes. *CoRR*, abs/2311.00558, 2023. (pg. [1](#), [3](#), [4](#), [7](#), [8](#))
- [KM24] Pravesh K. Kothari and Peter Manohar. Superpolynomial lower bounds for smooth 3-lccs and sharp bounds for designs. *CoRR*, abs/2404.06513, 2024. (pg. [1](#), [3](#), [4](#), [7](#), [8](#))
- [KMSV07] Peter Keevash, Dhruv Mubayi, Benny Sudakov, and Jacques Verstraëte. Rainbow Turán Problems. *Combinatorics, Probability and Computing*, 16(1):109–126, 2007. (pg. [5](#))
- [KT00] Jonathan Katz and Luca Trevisan. On the efficiency of local decoding procedures for error-correcting codes. In *Proceedings of the thirty-second annual ACM symposium on Theory of computing*, pages 80–86, 2000. (pg. [0](#), [1](#), [16](#))
- [KV00] Jeong Han Kim and Van H Vu. Concentration of multivariate polynomials and its applications. *Combinatorica*, 20(3):417–434, 2000. (pg. [15](#))
- [KW04] Iordanis Kerenidis and Ronald de Wolf. Exponential lower bound for 2-query locally decodable codes via a quantum argument. *Journal of Computer and System Sciences*, 69(3):395–420, 2004. (pg. [0](#), [1](#), [2](#), [6](#))
- [LPP91] Françoise Lust-Piquard and Gilles Pisier. Noncommutative Khintchine and Paley inequalities. *Ark. Mat.*, 29(2):241–260, 1991. (pg. [16](#))
- [Rom06] Andrei E. Romashchenko. Reliable computations based on locally decodable codes. In *STACS 2006, 23rd Annual Symposium on Theoretical Aspects of Computer Science, Marseille, France, February 23-25, 2006, Proceedings*, volume 3884 of *Lecture Notes in Computer Science*, pages 537–548. Springer, 2006. (pg. [1](#))

- [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, SODA '12*, page 437–446, USA, 2012. Society for Industrial and Applied Mathematics. (pg. [15](#), [16](#))
- [TJ74] Nicole Tomczak-Jaegermann. The moduli of smoothness and convexity and the Rademacher averages of trace classes $S_p(1 \leq p < \infty)$. *Studia Math.*, 50:163–182, 1974. (pg. [16](#))
- [Tom24] István Tomon. Robust (rainbow) subdivisions and simplicial cycles. *Advances in Combinatorics*, Jan 8 2024. (pg. [5](#))
- [Tre04] Luca Trevisan. Some applications of coding theory in computational complexity. *arXiv preprint cs/0409044*, 2004. (pg. [1](#))
- [Tro15] Joel A. Tropp. An Introduction to Matrix Concentration Inequalities. *Found. Trends Mach. Learn.*, 8(1-2):1–230, 2015. (pg. [16](#))
- [WAM19] Alexander S. Wein, Ahmed El Alaoui, and Cristopher Moore. The Kikuchi Hierarchy and Tensor PCA. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1446–1468. IEEE Computer Society, 2019. (pg. [1](#))
- [Wol09] Ronald de Wolf. Error-correcting data structures. In *26th International Symposium on Theoretical Aspects of Computer Science, STACS 2009, February 26-28, 2009, Freiburg, Germany, Proceedings*, volume 3 of *LIPICs*, pages 313–324. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2009. (pg. [1](#))
- [Woo07] David Woodruff. New lower bounds for general locally decodable codes. In *Electronic Colloquium on Computational Complexity (ECCC)*, volume 14, 2007. (pg. [1](#))
- [Woo08] David Woodruff. Corruption and recovery-efficient locally decodable codes. In *Approximation, Randomization and Combinatorial Optimization. Algorithms and Techniques*, pages 584–595. Springer, 2008. (pg. [1](#))
- [Woo12] David P Woodruff. A quadratic lower bound for three-query linear locally decodable codes over any field. *Journal of Computer Science and Technology*, 27(4):678–686, 2012. (pg. [1](#))
- [WW05] Stephanie Wehner and Ronald de Wolf. Improved lower bounds for locally decodable codes and private information retrieval. In *International Colloquium on Automata, Languages, and Programming*, pages 1424–1436. Springer, 2005. (pg. [1](#))
- [Yan24] Tal Yankovitz. A stronger bound for linear 3-lcc. *Electron. Colloquium Comput. Complex.*, pages TR24–036, 2024. (pg. [1](#), [3](#), [4](#), [8](#))

- [Yek08] Sergey Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM (JACM)*, 55(1):1–16, 2008. (pg. [1](#), [3](#), [7](#))
- [Yek10] Sergey Yekhanin. *Locally Decodable Codes and Private Information Retrieval Schemes*. Information Security and Cryptography. Springer, 2010. (pg. [1](#))
- [Yek12] Sergey Yekhanin. Locally decodable codes. *Foundations and Trends® in Theoretical Computer Science*, 6(3):139–255, 2012. (pg. [16](#))