SIAM J. COMPUT.
Vol. 46, No. 1, pp. 132–159

SUPER-POLYLOGARITHMIC HYPERGRAPH COLORING HARDNESS VIA LOW-DEGREE LONG CODES*

VENKATESAN GURUSWAMI†, PRAHLADH HARSHA‡, JOHAN HÅSTAD§, SRIKANTH SRINIVASAN¶, AND GIRISH VARMA‡

Abstract. We prove improved inapproximability results for hypergraph coloring using the low-degree polynomial code (aka the "short code" of Barak et al. [*SIAM J. Comput.*, 44 (2015), pp. 1287–1324]) and the techniques proposed by Dinur and Guruswami [*Israel J. Math.*, 209 (2015), pp. 611–649] to incorporate this code for inapproximability results. In particular, we prove quasi NP-hardness of the following problems on n-vertex hypergraphs: coloring a 2-colorable 8-uniform hypergraph with $2^{2^{\Omega(\sqrt{\log\log n})}}$ colors; coloring a 4-colorable 4-uniform hypergraph with $2^{2^{\Omega(\sqrt{\log\log n})}}$ colors; and coloring a 3-colorable 3-uniform hypergraph with $(\log n)^{\Omega(1/\log\log\log n)}$ colors. For the first two cases, the hardness results obtained are superpolynomial in what was previously known, and in the last case it is an exponential improvement. In fact, prior to this result, $(\log n)^{O(1)}$ colors was the strongest quantitative bound on the number of colors ruled out by inapproximability results for $O(1)$-colorable hypergraphs, and $(\log\log n)^{O(1)}$ for $O(1)$-colorable, 3-uniform hypergraphs.

Key words. hardness of approximation, hypergraph coloring, short code

AMS subject classifications. 68QXX, 68Q17, 05C15

DOI. 10.1137/140995520

1. Introduction.

A *k-uniform hypergraph* G is a pair (V, E), where V is a set of vertices and E is a collection of k-element subsets of V (i.e., $E \subseteq \binom{V}{k}$). These k-element subsets are called the *hyperedges* of the hypergraph G. An *independent set* in a hypergraph $G = (V, E)$ is a subset I of vertices such that no hyperedge is completely contained inside I. A *q-coloring* of a hypergraph is a map from V to the set $\{1, \ldots, q\}$ such that no hyperedge is monochromatic (i.e., every hyperedge has at least two distinct colors among its vertices). A hypergraph is said to be *q-colorable* if such a q-coloring exists or equivalently if the set of vertices can be partitioned into q independent sets. The *hypergraph coloring problem* is that of finding, given G, the smallest q such that G is q-colorable. When $k = 2$, the hypergraph is just a graph and the hypergraph coloring problem is the standard graph coloring problem.

Graph and hypergraph coloring problems have been studied extensively. The first nontrivial case of this problem is when $k = q = 2$, i.e., checking whether a graph is 2-colorable or, equivalently, whether the graph is a bipartite graph, and this turns

*Received by the editors November 13, 2014; accepted for publication (in revised form) September 9, 2016; published electronically February 22, 2017. A preliminary version of this paper appeared in *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, 2014 [GHH+14].
http://www.siam.org/journals/sicomp/46-1/99552.html

Funding: The first author was supported in part by a Packard Fellowship, U.S.–Israel BSF grant 2008293, and the U.S. National Science Foundation grant CCF-1115525. The third author was partly supported by ERC grant 226203. The fifth author was supported by Google India under the Google India PhD Fellowship Award.

†Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213 (guruswami@cmu.edu).

‡Tata Institute of Fundamental Research, Mumbai 400005, India (prahladh@tifr.res.in, girishrv@tifr.res.in). Part of the work was done while the second author was visiting the Simons Institute for Theory of Computing.

§KTH Royal Institute of Technology, Stockholm S-100 44, Sweden (johanh@kth.se). Part of the work was done while this author was visiting the Simons Institute for Theory of Computing.

¶Department of Mathematics, IIT Bombay, Mumbai 400076, India (srikanth@math.iitb.ac.in).

out to be easy. Every other case (for larger values of k or q) happens to be NP-hard; determining whether a graph is 3-colorable is a classical NP-hard problem, while for $k \geq 3$, even determining whether a given k-uniform hypergraph is 2-colorable is known to be NP-hard. This latter property (2-colorability of hypergraphs for $k \geq 3$), also referred to as *Property B*, has received a lot of attention in the extremal combinatorics literature.

Given that even checking whether a graph is 3-colorable is NP-hard, it is natural to ask whether there are approximately optimal coloring algorithms in the following sense: For a parameter $M > 3$, is there an algorithm that on input a 3-colorable graph colors it with at most M colors? It is of course trivial to color any graph with n colors where n is the number of vertices in the graph: Assign each vertex a different color. A long sequence of works [Wig83, Blu94, KMS98, BK97, AC06, Chl07, KT12, KT14], using both combinatorial and semidefinite programming techniques, give efficient polynomial-time algorithms to color 3-colorable graphs with n^{δ} colors, where the current best value of δ is approximately $0.199\ldots$. On the other hand, the best-known NP-hardness results for approximately coloring 3-colorable graphs are only able to prove that it is NP-hard to 4-color a 3-colorable graph! (See [KLS00, GK04, BG16].) Under variants of the Unique Games Conjecture, one can show hardness of $O(1)$-coloring 3-colorable graphs [DMR09]. Better hardness results are known for larger values of q: Huang [Hua13] showed that it is NP-hard to color a q-colorable graph with $2^{\Omega(q^{1/3})}$ colors, improving on the previous bound of $2^{\Omega(\log^2 q)}$ due to Khot [Kho01]. Observe the huge disparity between the known upper and lower bounds.

The situation for hypergraph coloring ($k \geq 3$) is slightly better. Moving from graphs ($k = 2$) to hypergraphs ($k \geq 3$) makes the algorithmic problem even harder, and thus the best-known algorithms still require $n^{\Omega(1)}$ colors to color a 2-colorable hypergraph [KNS01, CF96, AKMR96, KT14]. From the inapproximability perspective, Guruswami, Håstad, and Sudan [GHS02] proved the first superconstant lower bounds, showing the quasi NP-hardness of coloring 2-colorable 4-uniform hypergraphs with $\Omega(\frac{\log\log n}{\log\log\log n})$ colors. Since that work, there have been many others showing quasi NP-hardness for different values of k and q. A significant result in this line of work is that of Khot [Kho02a], wherein he obtained the first polylogarithmic lower bound, showing quasi NP-hardness of coloring q-colorable 4-uniform hypergraphs with $(\log n)^{\Omega(q)}$ colors for $q \geq 7$. More recently, Dinur and Guruswami [DG15, Appendix B] obtained a similar (but incomparable) polylogarithmic hardness result for 2-colorable 8-uniform hypergraphs. Our first results yield a "superpolynomial" improvement over these results.

THEOREM 1.1 (2-colorable 8-uniform hypergraphs). *Under the assumption that* $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{2^{O(\sqrt{\log\log n})}})$, *there is no polynomial-time algorithm that, when given as input an 8-uniform hypergraph H on N vertices, can distinguish between the following:*

- *H is 2-colorable.*
- *H has no independent set of size $N/2^{2^{O(\sqrt{\log\log N})}}$.*

THEOREM 1.2 (4-colorable 4-uniform hypergraphs). *Under the assumption that* $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{2^{O(\sqrt{\log\log n})}})$, *there is no polynomial-time algorithm that, when given as input a 4-uniform hypergraph H on N vertices, can distinguish between the following:*

- *H is 4-colorable.*
- *H has no independent set of size $N/2^{2^{O(\sqrt{\log\log N})}}$.*

We remark that all of the above-mentioned hardness results (including ours) prove stronger lower bounds than stated in the discussion: They show hardness of finding an

independent size of $N/M(N)$, which in turn implies hardness of coloring with $M(N)$ colors. Thus, our two results imply that there do not exist polynomial-time algorithms that color 2-colorable 8-uniform (similarly, 4-colorable 4-uniform) hypergraphs with $2^{2^{O(\sqrt{\log\log N})}}$ colors unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{2^{O(\sqrt{\log\log n})}})$.

We then ask whether we can prove coloring inapproximability for even smaller uniformity, i.e., $k = 3$, the case of 3-uniform hypergraphs. The best known inapproximability results in this context are as follows. Khot [Kho02b] showed that it is quasi-NP-hard to find an independent set of size $N/(\log\log N)^{1/9}$ in a given N-vertex 3-colorable 3-uniform hypergraph. Dinur, Regev, and Smyth [DRS05] showed that it is quasi-NP-hard to color a 2-colorable 3-uniform hypergraph with $(\log\log N)^{1/3}$ colors. Our third result yields an "exponential" improvement.

THEOREM 1.3 (3-colorable 3-uniform hypergraphs). *Under the assumption that* $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{2^{O(\log\log n/\log\log\log n)}})$, *there is no polynomial-time algorithm that, when given as input a 3-uniform hypergraph H on N vertices, can distinguish between the following:*

- *H is 3-colorable.*
- *H has no independent set of size $N/2^{O(\log\log N/\log\log\log N)}$.*

As with the previous two results, Theorem 1.3 implies a result on the hardness of approximate coloring. More precisely, unless $\mathsf{NP} \subseteq \mathsf{DTIME}(n^{2^{O(\log\log n/\log\log\log n)}})$, there does not exist a polynomial-time algorithm that colors 3-colorable 3-uniform hypergraphs with $2^{\log\log N/\log\log\log N}$ colors.

1.1. Proof approach. All known hypergraph coloring inapproximability results are obtained using the machinery of probabilistically checkable proofs (PCPs). A PCP construction in which the verifier queries k locations of the proof and accepts if the symbols in these locations are not all equal (this is referred to as the not-all-equal (NAE) predicate) naturally yields a hardness result for approximate coloring of k-uniform hypergraphs via the following correspondence. The vertices of the hypergraph correspond to the locations in the PCPs, while the hyperedges correspond to the k-sized queries of the verifier.

Before we proceed to explain the ideas in our proofs, let us first try to understand why all previous hypergraph coloring inapproximability techniques got stuck at the $\operatorname{poly}\log n$ color barrier. Constructions of PCPs with a specific predicate (NAE in our case) typically proceed along the following lines. An outer PCP verifier is first constructed using the hardness of the so-called label cover problem. This is then composed with an inner verifier that makes tests restricted to the predicate corresponding to the desired hardness result (NAE in the case of coloring). One of the quintessential ingredients in the inner verifier construction (in almost all known inapproximability results) is the *long code*, first introduced by Bellare, Goldreich, and Sudan [BGS98]. The long code, as the name suggests, is a highly redundant encoding of its input (in fact, it is *the most* redundant encoding that doesn't repeat symbols). Under the long code, an n-bit Boolean string x is encoded by a 2^{2^n}-bit string that consists of the evaluation of *all* Boolean functions on n bits at the point x. It is this doubly exponential blowup of the long code that prevents the coloring inapproximability from going past the $\operatorname{poly}\log n$ barrier.

Recently, Barak et al. [BGH+15], while trying to understand the tightness of the Arora–Barak–Steurer algorithm for unique games, introduced the *short code*, also called the *low-degree long code* [DG15]. The low-degree long code is a puncturing of the long code that contains only the evaluations of low-degree functions (as opposed to

all functions). Because this is a "shortening" of the long code, one might hope to use the low-degree long code as a more size-efficient surrogate for the long code in inapproximability results. In fact, Barak et al. [BGH+15] used it to obtain a more efficient version of the MAXCUT inapproximability result of Khot et al. [KKMO07], assuming the Unique Games Conjecture, as well as integrality gap instances for Unique Games against exponentially more rounds of the Sherali–Adams + SDP hierarchy than known previously. The short code was also used by Kane and Meka [KM13] to construct instances of Uniform Sparsest Cut with an exponentially larger integrality gap ($\exp(\Omega(\sqrt{\log \log n}))$ compared to earlier $\Omega(\log \log n)$) against powerful semidefinite programs.

One of the challenges in PCP constructions that imply hardness for coloring problems is the requirement of *perfect completeness*, i.e., for Yes instances, there must exist a proof that is accepted by the verifier with probability 1 (this corresponds to the graph or hypergraph obtained by the reduction being properly colorable with every (hyper)-edge legally colored). In contrast, reductions based on the Unique Games Conjecture inherently lack perfect completeness; this offers flexibility to add noise to the queries made by the inner verifier, which in turns aids in the soundness analysis of the PCP. In the context of PCPs with perfect completeness, Dinur and Guruswami [DG15] introduced some elegant techniques to adapt the long-code-based inapproximability results to low-degree long codes. Barak et al. [BGH+15] obtained their results by discovering an intimate connection between Reed–Muller testing of Bhattacharyya et al. [BKS+10] and analysis of the low-degree long code. Exploring this connection further, Dinur and Guruswami [DG15] proved a new result for testing Reed–Muller codes over $\mathbb{F}_2$ (i.e., testing whether a given function is close to a low-degree polynomial over $\mathbb{F}_2$), which we describe below.

Let P_d^n be the set of degree d polynomials on n variables over $\mathbb{F}_2$, and let $\chi_f(g) = \chi_g(f) := (-1)^{\sum_{x\in\mathbb{F}_2^n} f(x)g(x)}$ denote the correlation between two functions $f, g : \mathbb{F}_2^n \to \mathbb{F}_2$. Thus, $\chi_f(g) = \chi_g(f) = 1$ if f and g are orthogonal over $\mathbb{F}_2$ (i.e., $\sum_{x\in\mathbb{F}_2^n} f(x)g(x) = 0$) and -1 otherwise. It is well known that P_{n-d-1}^n is exactly the set of functions that are orthogonal to *all* functions in P_d^n. In particular, $\chi_\beta(gh) = 1$ if $\beta \in \mathsf{P}_{n-d-1}^n$, $g \in \mathsf{P}_{d/4}^n$, and $h \in \mathsf{P}_{3d/4}^n$. On the other hand, if $\beta \notin \mathsf{P}_{n-d-1}^n$, we have $\mathbb{E}_{f\in\mathsf{P}_d^n}[\chi_\beta(f)] = 0$. The Dinur–Guruswami testing result states that if $\beta : \mathbb{F}_2^n \to \mathbb{F}_2$ is far from any degree $n-d-1$ polynomial, then for most degree $d/4$ polynomials, g, β will only be orthogonal to roughly *half* of the polynomials gh as h varies over degree $3d/4$ polynomials. More precisely, if $\beta : \mathbb{F}_2^n \to \mathbb{F}_2$ is $2^{d/2}$-far[1] from P_{n-d-1}^n, then $\mathbb{E}_{g\in\mathsf{P}_{d/4}^n} |\mathbb{E}_{h\in\mathsf{P}_{3d/4}^n}[\chi_\beta(gh)]|$ is doubly exponentially small in d (see Theorem 2.12 for the exact statement). This Reed–Muller testing result let them analyze the low-degree long code and construct inner verifiers with perfect completeness. Our first two hypergraph coloring results are obtained by constructing an appropriate NAE-predicate inner verifier using these techniques of Dinur and Guruswami.

For the case of 3-uniform 3-colorable hypergraphs, we adapt Khot's hardness result [Kho02b] to the low-degree long code setting. To analyze the low-degree long code in this setting, we prove the following testing result for Reed–Muller codes over $\mathbb{F}_3$ (i.e., moving to a ternary alphabet instead of the binary alphabet $\mathbb{F}_2$). Let P_d^n now denote the set of degree d polynomials on n variables over $\mathbb{F}_3$, and let $\chi_f(g) = \chi_g(f) := \omega^{\sum_{x\in\mathbb{F}_3^n} f(x)g(x)}$ denote the correlation between two functions $f, g : \mathbb{F}_3^n \to \mathbb{F}_3$,

[1] We say that g is Δ-far from a class of functions $\mathcal{F}$ if for all $f \in \mathcal{F}$, we have $|\{x \in \mathbb{F}_2^n | f(x) \neq g(x)\}| \geq \Delta$ (note that we are using the nonnormalized Hamming distance).

where $\omega = e^{2\pi i/3}$. Thus, $\chi_f(g) = \chi_g(f) = 1$ iff f and g are orthogonal over $\mathbb{F}_3$ (i.e., $\sum_{x\in\mathbb{F}_3^n} f(x)g(x) = 0$). It is known that $\mathsf{P}^n_{2n-\ell-1}$ is exactly the set of functions that are orthogonal to *all* functions in P^n_ℓ. In particular, $\chi_\beta(p^2) = 1$ if $\beta \in \mathsf{P}^n_{2n-2d-1}$ and $p \in \mathsf{P}^n_d$. Similar in spirit to the Dinur–Guruswami testing result, we prove the following result: If $\beta : \mathbb{F}_3^n \to \mathbb{F}_3$ is far from any degree $2n-2d-1$ polynomial, then the correlation of β with the square of a random degree d polynomial is very small. More precisely, if $\beta : \mathbb{F}_3^n \to \mathbb{F}_3$ is $3^{d/2}$-far from $\mathsf{P}^n_{2n-2d-1}$, then $|\mathbb{E}_{p\in\mathsf{P}^n_d}[\chi_\beta(p^2)]|$ is doubly exponentially small in d. This is proved by considering the associated quadratic form Q^β defined as

$$Q(\beta) := \sum_{x\in\mathbb{F}_3^n} \beta(x)\cdot \mathrm{eval}(x)\,\mathrm{eval}(x)^T,$$

where $\mathrm{eval}(x)$ is the column-vector of evaluation of all degree d monomials at the point x. Observe that this quadratic form $Q(\beta)$ satisfies $p^TQ(\beta)p = \sum_{x\in\mathbb{F}_3^n}\beta(x)p^2(x)$. It is well known that the distance of the random variable p^TAp (for random p and fixed symmetric A) is inverse exponential in the rank of the quadratic form A. The testing result is thus proved by showing the following result on the rank of the quadratic form $Q(\beta)$: If the distance of β from polynomials of degree $2n-2d-1$ is at least $3^{d/2}$, then the rank of the matrix $Q(\beta)$ is $3^{\Omega(d)}$. This rank bound is proved along the lines of Dinur and Guruswami's result [DG15] using the Reed–Muller tester analysis of Haramaty, Shpilka, and Sudan [HSS13] over general fields instead of the Bhattacharyya et al. [BKS+10] analysis over $\mathbb{F}_2$.

1.2. Subsequent work. Saket [Sak14] recently obtained the following improved inapproximability result for 2-colorable 4-uniform hypergraphs: It is quasi-NP-hard to color a 2-colorable 4-uniform hypergraph with $(\log n)^c$ colors for some constant c. He obtained this improvement by giving an improved analysis using reverse hypercontractivity of the long-code-based test of Guruswami, Håstad, and Sudan [GHS02].

Subsequent to our result, Khot and Saket [KS14], in a significant improvement, showed that it is quasi-NP-hard to color a 2-colorable 12-uniform hypergraph with $2^{(\log n)^c}$ colors for some constant $c \in (0,1)$. This result is obtained by constructing a powerful outer verifier with a strong soundness property, referred to as the *superposition complexity*, which is then composed with an inner verifier based on the quadratic code (equivalently, the low-degree long code of degree two). Huang [Hua15] then gave a slightly simpler construction of PCPs with superposition complexity. Surprisingly, while previous results (including the results in the current paper) employed the Reed–Muller testing results in the analysis of the inner verifier, Khot and Saket (and then Huang) used these testing results in the analysis of their outer verifier. Varma [Var15] then showed how to improve the uniformity of the Khot–Saket 12-query verifier to an 8-query inner verifier based on the reductions in the current paper to yield the following result: It is quasi-NP-hard to color a 2-colorable 8-uniform hypergraph (similarly, a 4-colorable 4-uniform hypergraph) with $2^{(\log n)^c}$ colors for some constant $c \in (0,1)$.

Organization. We start with some preliminaries in section 2. Theorems 1.1, 1.2, and 1.3 are proved in sections 4, 5, and 6, respectively. The proof of the latter theorem requires a technical claim about low-degree polynomials over $\mathbb{F}_3$, which we prove in section 3.

2. Preliminaries.

2.1. Label cover. All of our reductions start from an appropriate instance of the label cover problem, bipartite or multipartite. A bipartite label cover instance

consists of a bipartite graph $G = (U, V, E)$, label sets Σ_U, Σ_V, and a set of projection constraints $\Pi = \{\pi_{uv} : \Sigma_U \to \Sigma_V | (u, v) \in E\}$. We consider label cover instances obtained from 3SAT instances in the following natural manner.

DEFINITION 2.1 (r-repeated label cover). *Let φ be a* 3SAT *instance with X as the set of variables and C the set of clauses. The r-repeated bipartite label cover instance $I(\varphi)$ is specified by the following:*

- *A graph $G := (U, V, E)$, where $U := C^r, V := X^r$.*
- *$\Sigma_U := \{0,1\}^{3r}, \Sigma_V := \{0,1\}^r$.*
- *There is an edge $(u, v) \in E$ if the tuple of variables v can be obtained from the tuple of clauses u by replacing each clause by a variable in it.*
- *The constraint $\pi_{uv} : \{0,1\}^{3r} \to \{0,1\}^r$ is simply the projection of the assignments on $3r$ variables in all the clauses in u to the assignments on the r variables in v.*
- *For each u there is a set of r functions $\{f_i^u : \{0,1\}^{3r} \to \{0,1\}\}_{i=1}^r$ such that $f_i^u(a) = 0$ iff the assignment a satisfies the ith clause in u. Note that f_i^u depends only on the 3 variables in the ith clause.*

A labeling $L_U : U \to \Sigma_U, L_V : V \to \Sigma_V$ satisfies an edge (u, v) iff $\pi_{uv}(L_U(u)) = L_V(v)$ and $L_U(u)$ satisfies all the clauses in u. Let $OPT(I(\varphi))$ be the maximal fraction of constraints that can be satisfied by any labeling.

The following theorem is obtained by applying Raz's parallel repetition theorem [Raz98] with r repetitions on hard instances of MAX-3SAT where each variable occurs the same number of times [Fei98].

THEOREM 2.2. *There is an algorithm that on input a* 3SAT *instance φ and $r \in \mathbb{N}$ outputs an r-repeated label cover instance $I(\varphi)$ in time $n^{O(r)}$ with the following properties:*

- *If $\varphi \in$ 3SAT, then $OPT(I(\varphi)) = 1$.*
- *If $\varphi \notin$ 3SAT, then $OPT(I(\varphi)) \le 2^{-\varepsilon_0 r}$ for some universal constant $\varepsilon_0 \in (0, 1)$.*

Moreover, the underlying graph G is both left and right regular.

2.2. Multilayered smooth label cover. For our hardness results for 3-uniform 3-colorable hypergraphs, we need a multipartite version of label cover, satisfying a smoothness condition, which was introduced by Khot [Kho02b].

DEFINITION 2.3 (smoothness). *Let I be a bipartite label cover instance specified by $((U, V, E), \Sigma_U, \Sigma_V, \Pi)$. Then I is η-*smooth *iff for every $u \in U$ and two distinct labels $a, b \in \Sigma_U$,*

$$\Pr_v[\pi_{uv}(a) = \pi_{uv}(b)] \le \eta,$$

where v is a random neighbor of u.

DEFINITION 2.4 (r-repeated ℓ-layered η-smooth label cover). *Let $T := \lceil \ell/\eta \rceil$ and let φ be a* 3SAT *instance with X as the set of variables and C the set of clauses. The r-repeated ℓ-layered η-smooth label cover instance $I(\varphi)$ is specified by the following:*

- *An ℓ-partite graph with vertex sets $V_0, \ldots, V_{\ell-1}$. Elements of V_i are tuples of the form (C', X'), where C' is a set of $(T + \ell - i)r$ clauses and X' is a set of ir variables.*
- *$\Sigma_{V_i} := \{0,1\}^{m_i}$, where $m_i := 3(T + \ell - i)r + ir$, which corresponds to all Boolean assignments to the clauses and variables corresponding to a vertex in layer V_i.*
- *For $0 \le i < j < \ell$, $E_{ij} \subseteq V_i \times V_j$ denotes the set of edges between layers V_i and V_j. For $v_i \in V_i$, $v_j \in V_j$, there is an edge $(v_i, v_j) \in E_{ij}$ iff v_j can*

be obtained from v_i by replacing some $(j-i)r$ clauses in v_i with variables occurring in the clauses, respectively.

- *The constraint $\pi_{v_iv_j}$ is the projection of assignments for clauses and variables in v_i to those of v_j.*
- *For each $i < \ell$, $v_i \in V_i$, there are $(T+\ell-i)r$ functions $f_j^{v_i} : \{0,1\}^{3(T+\ell-i)r+ir} \to \{0,1\}$, one for each clause j in v_i such that $f_j^{v_i}(a) = 0$ iff a satisfies the clause j. This function depends only on the 3 coordinates in j.*

Given a labeling $L_i : V_i \to \Sigma_{V_i}$ for all the vertices, an edge $(v_i, v_j) \in E_{ij}$ is satisfied iff $L_i(v_i)$ satisfies all the clauses in v_i, $L_j(v_j)$ satisfies all the clauses in v_j, and $\pi_{v_iv_j}(L_i(v_i)) = L_j(v_j)$. Let $OPT_{ij}(I(\varphi))$ be the maximum fraction of edges in E_{ij} that can be satisfied by any labeling.

The following theorem was proved by Dinur et al. [DGKR05] in the context of hypergraph vertex cover inapproximability (also see [DRS05]).

THEOREM 2.5. *There is an algorithm that on input a* 3SAT *instance φ and $\ell, r \in \mathbb{N}, \eta \in [0,1)$ outputs an r-repeated ℓ-layered η-smooth label cover instance $I(\varphi)$ in time $n^{O((1+1/\eta)\ell r)}$ with the following properties:*

1. *For all $0 \le i < j < \ell$, the bipartite label cover instance on $I_{ij} = ((V_i, V_j, E_{ij}), \Sigma_{V_i}, \Sigma_{V_j}, \Pi_{ij})$ is η-smooth.*
2. *For $1 < m < \ell$, any m layers $0 \le i_1 < \cdots < i_m \le \ell - 1$, and any $S_{i_j} \subseteq V_{i_j}$ such that $|S_{i_j}| \ge \frac{2}{m}|V_{i_j}|$, there exist distinct i_j and $i_{j'}$ such that the fraction of edges between S_{i_j} and $S_{i_{j'}}$ relative to $E_{i_ji_{j'}}$ is at least $1/m^2$.*
3. *If $\varphi \in$ 3SAT, then there is a labeling for $I(\varphi)$ that satisfies all the constraints.*
4. *If $\varphi \notin$ 3SAT, then*

$$OPT_{i,j}(I(\varphi)) \le 2^{-\Omega(r)} \quad \forall 0 \le i < j \le \ell.$$

2.3. Low-degree long code. Let $\mathbb{F}_p$ be the finite field of size p where p is a prime. The results in this section apply when $p = 2, 3$. The choice of p will be clear from the context, and hence the dependence of p on the quantities defined will be omitted. Let P_d^n be the set of polynomials of degree at most d on n variables over $\mathbb{F}_p$. Let $\mathfrak{F}_n := \mathsf{P}_{(p-1)n}^n$. Note that $\mathfrak{F}_n$ is the set of all functions from $\mathbb{F}_p^n$ to $\mathbb{F}_p$. $\mathfrak{F}_n$ is an $\mathbb{F}_p$-vector space of dimension p^n, and P_d^n is its subspace of dimension $n^{O(d)}$. The Hamming distance between f and $g \in \mathfrak{F}_n$, denoted by $\Delta(f,g)$, is the number of inputs on which f and g differ. When $S \subseteq \mathfrak{F}_n$, $\Delta(f,S) := \min_{g\in S} \Delta(f,g)$. We say f is Δ-far from S if $\Delta(f,S) \ge \Delta$ and f is Δ-close to S otherwise. Given $f, g, \in \mathfrak{F}_n$, the dot product between them is defined as

$$\langle f, g\rangle := \sum_{x\in\mathbb{F}_p^n} f(x)g(x).$$

For a subspace $S \subseteq \mathfrak{F}_n$, the dual subspace is defined as

$$S^\perp := \{g \in \mathfrak{F}_n : \forall f \in S, \langle g, f\rangle = 0\}.$$

The following theorem relating dual spaces is well known and is used to index the characters of P_d^n (see Lemma 2.10).

LEMMA 2.6. $(\mathsf{P}_d^n)^\perp = \mathsf{P}_{(p-1)n-d-1}^n$.

Proof. First note that the dimensions of the two subspaces are equal by a counting argument. Next we show that $(\mathsf{P}_d^n)^\perp \supseteq \mathsf{P}_{(p-1)n-d-1}^n$. We just need to show that for

any monomial of degree $(p-1)n-d-1$ with individual degrees $< p$, the dot product with any monomial of degree d with individual degrees $< p$ is 0. The product of any such pair of monomials is a monomial with total degree at most $(p-1)n-1$, and hence has a variable with degree $< p-1$. Without loss of generality, let this variable be x_1 with degree $t < p-1$. Notice that $\sum_{x_1 \in \mathbb{F}_p} x_1^t = 0$, and hence the dot product is 0. ∎

We need the following Schwartz–Zippel-like lemma for degree d polynomials. It is used in the soundness analysis of the low-degree long code tests, to lower bound the rejection probabilities.

LEMMA 2.7 (Schwartz–Zippel lemma [HSS13, Lemma 3.2]). *Let $f \in \mathbb{F}_p[x_1, \dots, x_n]$ be a nonzero polynomial of degree at most d with individual degrees at most $p-1$. Then $\Pr_{a \in \mathbb{F}_p^n}[f(a) \neq 0] \geq p^{-d/(p-1)}$.*

We now define the low-degree long code (introduced as the short code by Barak et al. [BGH+15] in the $\mathbb{F}_2$ case).

DEFINITION 2.8 (low-degree long code). *For $a \in \mathbb{F}_p^n$, the degree d long code for a is a function $\mathsf{LC}_d(a) : \mathsf{P}_d^n \to \mathbb{F}_p$ defined as*

$$\mathsf{LC}_d(a)(f) := f(a).$$

Note that for $d = (p-1)n$, this matches with the definition of the original long code over the alphabet $\mathbb{F}_p$.

DEFINITION 2.9 (characters). *A character of P_d^n is a function $\chi : \mathsf{P}_d^n \to \mathbb{C}$ such that*

$$\chi(0) = 1 \quad \text{and} \quad \forall f, g \in \mathsf{P}_d^n,\ \chi(f+g) = \chi(f)\chi(g).$$

The following lemma lists the basic properties of characters. They are used in the soundness analysis of the low-degree long code tests, analogous to Fourier analysis for long code tests.

LEMMA 2.10. *Let $\{1, \omega, \dots, \omega^{p-1}\}$ be the pth roots of unity, and for $\beta \in \mathfrak{F}_n, f \in \mathsf{P}_d^n$, let $\chi_\beta(f) := \omega^{\langle \beta, f \rangle}$.*

- *The characters of P_d^n are $\{\chi_\beta : \beta \in \mathfrak{F}_n\}$.*
- *For any $\beta, \beta' \in \mathfrak{F}_n$, $\chi_\beta = \chi'_\beta$ iff $\beta - \beta' \in (\mathsf{P}_d^n)^\perp$.*
- *For $\beta \in (\mathsf{P}_d^n)^\perp$, χ_β is the constant 1 function.*
- *For all β, there exists β' such that $\beta - \beta' \in (\mathsf{P}_d^n)^\perp$ and $|\operatorname{support}(\beta')| = \Delta(\beta, (\mathsf{P}_d^n)^\perp)$ (i.e., the constant 0 function is (one of) the closest function to β' in $(\mathsf{P}_d^n)^\perp$). We call such a β' a minimum support function for the coset $\beta + (\mathsf{P}_d^n)^\perp$.*
- *Characters form an orthonormal basis for the vector space of functions from P_d^n to $\mathbb{C}$, under the inner product $\langle A, B \rangle := \mathbb{E}_{f \in \mathsf{P}_d^n}[A(f)\overline{B(f)}]$.*
- *Any function $A : \mathsf{P}_d^n \to \mathbb{C}$ can be uniquely decomposed as*

$$A(f) = \sum_{\beta \in \Lambda_d^n} \widehat{A}(\beta)\chi_\beta(f),$$

where $\widehat{A}(\beta) := \mathbb{E}_{g \in \mathsf{P}_d^n}[A(g)\overline{\chi_\beta(g)}]$ and Λ_d^n is the set of minimum support functions, one for each of the cosets in $\mathfrak{F}_n/(\mathsf{P}_d^n)^\perp$, with ties broken arbitrarily.

- *Parseval's identity: For any $A : \mathsf{P}_d^n \to \mathbb{C}$, $\sum_{\beta \in \Lambda_d^n} |\widehat{A}(\beta)|^2 = \mathbb{E}_{f \in \mathsf{P}_d^n}[|A(f)|^2]$. In particular, if $A : \mathsf{P}_d^n \to \{1, \omega, \dots, \omega^{p-1}\}$, $\sum_{\beta \in \Lambda_d^n} |\widehat{A}(\beta)|^2 = 1$.*

The following lemma relates characters over different domains related by coordinate projections.

LEMMA 2.11. *Let $m \le n$ and $\pi : \mathbb{F}_p^n \to \mathbb{F}_p^m$ be a (coordinate) projection; i.e., there exist indices $1 \le i_i < \cdots < i_m \le n$ such that $\pi(x_1, \ldots, x_n) = (x_{i_1}, \ldots, x_{i_m})$. Then for $f \in \mathsf{P}_d^m$, $\beta \in \mathsf{P}_d^n$,*

$$\chi_\beta(f \circ \pi) = \chi_{\pi_p(\beta)}(f),$$

where $\pi_p(\beta)(y) := \sum_{x \in \pi^{-1}(y)} \beta(x)$.

Proof.

$$\begin{aligned}\chi_\beta(f \circ \pi) &= \omega^{\sum_{x \in \mathbb{F}_3^n} f(\pi(x))\beta(x)} = \omega^{\sum_{y \in \mathbb{F}_3^m} f(y)\left(\sum_{x \in \pi^{-1}(y)} \beta(x)\right)} \\ &= \omega^{\sum_{y \in \mathbb{F}_3^m} f(y)\pi_p(\beta)(y)} = \chi_{\pi_p(\beta)}(f). \qquad \square\end{aligned}$$

Dinur and Guruswami [DG15] proved the following theorem about Reed–Muller codes over $\mathbb{F}_2$ using the testing result of Bhattacharyya et al. [BKS+10].

THEOREM 2.12 (see [DG15, Theorem 1]). *Let d be a multiple of 4 and $p = 2$. If $\gamma \in \mathfrak{F}_n$ is $2^{d/2}$-far from $(\mathsf{P}_d^n)^\perp = \mathsf{P}_{n-d-1}^n$, then*

$$\mathop{\mathbb{E}}_{g \in \mathsf{P}_{d/4}^n} \left[\left| \mathop{\mathbb{E}}_{h \in \mathsf{P}_{3d/4}^n} [\chi_\gamma(gh)] \right| \right] \le 2^{-2^{(d/4-2)}}. \tag{2.1}$$

(Theorem 1 in [DG15] states an upper bound of $2^{-4 \cdot 2^{d/4}}$ on the expectation in (2.1); however, the proof in fact shows the bound $2^{-2^{(d/4-2)}}$. This small change is inconsequential, and the key is the doubly exponential decay in d.)

2.4. Folding over satisfying assignments.

LEMMA 2.13. *Let $d > 1$, let X be a set of $p^d - 1$ points in $\mathbb{F}_p^n$, and let $f : X \to \mathbb{F}_p$ be an arbitrary function. Then there exists a polynomial q of degree at most $(p-1)d$ such that q agrees with f on all points in X.*

Proof. By Lemmas 2.6 and 2.7, any nonzero polynomial in $(\mathsf{P}_{(p-1)d}^n)^\perp$ has support size at least p^d. In other words, the evaluations of $(\mathsf{P}_{(p-1)d}^n)^\perp$ at $\mathbb{F}_p^n$ form a code of distance at least p^d. Therefore its dual code, namely the evaluations of $\mathsf{P}_{(p-1)d}^n$ at $\mathbb{F}_p^n$, induces a $(p^d - 1)$-wise independent distribution. Hence, it is possible to interpolate a degree $(p-1)d$ polynomial to take on any desired values at an arbitrary subset of $p^d - 1$ points in $\mathbb{F}_p^n$. $\square$

For any set S, a function $A : \mathsf{P}_{(p-1)d}^n \to S$ is said to be *folded* over a subspace $J \subseteq \mathsf{P}_{(p-1)d}^n$ if A is constant over cosets of J in $\mathsf{P}_{(p-1)d}^n$.

FACT 2.14. *Given a function $A : \mathsf{P}_{(p-1)d}^n / J \to S$, there is a unique function $A' : \mathsf{P}_{(p-1)d}^n \to S$ that is folded over J such that for $g \in \mathsf{P}_{(p-1)d}^n$, $A'(g) = A(g + J)$. We call A' the lift of A.*

Given $q_1, \ldots, q_k \in \mathsf{P}_{3(p-1)}^n$, let

$$J(q_1, \ldots, q_k) := \left\{ \sum_i r_i q_i : r_i \in \mathsf{P}_{(p-1)(d-3)}^n \right\}.$$

The following lemma shows that if a function is folded over $J = J(q_1, \ldots, q_k)$, then it cannot have weight on small support characters that are nonzero on J (this is a generalization of the corresponding lemma in [DG15] to other fields).

LEMMA 2.15. *Let $\beta \in \mathfrak{F}_n$ be such that $|\operatorname{support}(\beta)| < p^{d-3}$, and there exists $x \in \operatorname{support}(\beta)$ with $q_i(x) \neq 0$ for some i. Then if $A : \mathsf{P}_d^n \to \mathbb{C}$ is folded over $J = J(q_1, \ldots, q_k)$, then $\widehat{A}(\beta) = 0$.*

Proof. Construct a polynomial r that is zero at all points in support of β except at x. From Lemma 2.13, its possible to construct such a polynomial of degree at most $(p-1)(d-3)$. Then we have that $rq_i \in J$ and $\langle \beta, rq_i \rangle \neq 0$. Now

$$\begin{aligned}
\mathop{\mathbb{E}}_h [A(h)\chi_\beta(h)] &= \frac{1}{p} \mathop{\mathbb{E}}_h [A(h)\chi_\beta(h) + A(h+rq_i)\chi_\beta(h+rq_i) + \cdots \\
&\qquad + A(h+(p-1)rq_i)\chi_\beta(h+(p-1)rq_i)] \\
&= \frac{1}{p} \mathop{\mathbb{E}}_h [A(h)\chi_\beta(h) + A(h)\chi_\beta(h+rq_i) + \cdots + A(h)\chi_\beta(h+(p-1)rq_i)] \\
&= \frac{1}{p} \mathop{\mathbb{E}}_h [A(h)\chi_\beta(h)(1 + \chi_\beta(rq_i) + \cdots + \chi_\beta((p-1)rq_i))] \\
&= 0 \qquad (\text{since } \chi_\beta(rq_i) \neq 1). \qquad \square
\end{aligned}$$

3. Correlation with a random square. In this section, we analyze the quantity $\langle \beta, p^2 \rangle$, where $p \in \mathsf{P}_d^n$ is chosen uniformly at random and $\beta : \mathbb{F}_3^n \to \mathbb{F}_3$ is a fixed function having distance exactly Δ from $(\mathsf{P}_{2d}^n)^\perp = \mathsf{P}_{2n-2d-1}^n$.

Throughout this section, we work over the field $\mathbb{F}_3$. For $a \in \mathbb{N}^n$, let $|a| := \sum_i a_i$ and let x^a denote the monomial $\prod_i x_i^{a_i}$. Over $\mathbb{F}_3$, the individual degrees are at most 2 (since $x^3 \equiv x$). Hence, we assume without loss of generality that the coefficient vector $a \in \{0,1,2\}^n$. In this notation, $p(x) = \sum_{|a| \leq d} p_a x^a$, where p_a are chosen independently and uniformly at random from $\mathbb{F}_3$. For $x \in \mathbb{F}_3^n$, let e_x be the column vector of evaluation of all degree d monomials at x, i.e., $e_x := (x^a)_{|a| \leq d}$. Then $p(x) = p^T e_x$, where p is now thought of as the column vector $(p_a)_{|a| \leq d}$ and hence, $p^2(x) = (p^T e_x)^2 = p^T (e_x e_x^T) p$. Thus we have

$$\langle \beta, p^2 \rangle = \sum_x \beta(x) \left(p^T e_x e_x^T p \right) = p^T \left(\sum_x \beta(x) e_x e_x^T \right) p.$$

We are thus interested in the quadratic form represented by the matrix $Q^\beta := \sum_x \beta(x) e_x^T e_x$. Observe that all β belonging to the same coset in $\mathsf{P}_{2n}^n / \mathsf{P}_{2n-2d-1}^n$ have the same value for $\langle \beta, p^2 \rangle$ and the matrix Q^β. Hence, by Lemma 2.10, we might without loss of generality assume that β satisfies $|\operatorname{support}(\beta)| = \Delta$. The following lemma (an easy consequence of [LN97, Theorem 6.21]) shows that it suffices to understand the rank of Q^β.

LEMMA 3.1. *Let A be an $n \times n$, symmetric matrix with entries from $\mathbb{F}_3$. The statistical distance of the random variable $p^T A p$ from uniform is $\exp(-\Omega(\operatorname{rank}(A)))$.*

In the next sequence of lemmas, we relate $\operatorname{rank}(Q^\beta)$ to Δ. In particular, we show that $\operatorname{rank}(Q^\beta)$ is equal to Δ if $\Delta \leq 3^{d/2}$, and is exponential in d otherwise. Recall that over $\mathbb{F}_3$, P_{2n}^n is the set of all functions from $\mathbb{F}_3^n$ to $\mathbb{F}_3$ and $(\mathsf{P}_{2d}^n)^\perp = \mathsf{P}_{2n-2d-1}^n$.

LEMMA 3.2. $\operatorname{rank}(Q^\beta) \leq \Delta$.

Proof. By assumption, β satisfies $\Delta = |\operatorname{support}(\beta)|$. The lemma follows from the fact that $e_x e_x^T$ are rank one matrices and $Q^\beta = \sum_x \beta(x) e_x e_x^T$. $\square$

LEMMA 3.3. *If $\Delta < 3^{d/2}$, then $\operatorname{rank}(Q^\beta) = \Delta$.*

Proof. By assumption, β satisfies $\Delta = |\operatorname{support}(\beta)|$ and $Q^\beta = \sum_x \beta(x) e_x e_x^T$. Since $(\mathsf{P}^n_d)^\perp = \mathsf{P}^n_{2n-d-1}$ and any nonzero polynomial with degree $2n-d-1$ has support at least $3^{d/2}$ (Lemma 2.7), arguing as in Lemma 2.13, any $\lceil 3^{d/2} \rceil - 1$ vectors e_x are linearly independent. In particular, the Δ vectors e_x for x in support(β) are linearly independent. Consider any nonzero v in the kernel of the matrix Q^β. The linear independence of e_x gives that $e_x^T v = 0$ for all $x \in \operatorname{support}(\beta)$. Hence, the kernel of Q^β resides in a Δ-codimensional space, which implies that $\operatorname{rank}(Q^\beta) = \Delta$. ∎

We conjecture that Lemma 3.3 holds for larger values of Δ, but for our purposes we only need a lower bound on the rank when $\Delta \geq 3^{d/2}$.

LEMMA 3.4. *There exists a constant d_0 such that if $d > d_0$ and $\Delta > 3^{d/2}$, then* $\operatorname{rank}(Q^\beta) \geq 3^{d/9}$.

Proof. The proof of this theorem is similar to the proof of [DG15, Theorem 17] for the $\mathbb{F}_2$ case, and we follow it step by step. Define $B^n_{d,k}(\beta) := \{q \in \mathsf{P}^n_k : q\beta \in \mathsf{P}^n_{2n-2d-1+k}\}$.

CLAIM 3.5. $\ker(Q^\beta) = B^n_{d,d}(\beta)$.

Proof. The matrix Q^β satisfies that $Q^\beta(a,b) = \langle \beta, x^a x^b \rangle$ for all $a, b \in \{0,1,2\}^n$, $|a|, |b| \leq d$. Using this description of Q^β, we obtain the following description of $\ker(Q^\beta)$:

$$\begin{aligned}
(h_a)_{|a|\leq d} \in \ker(Q^\beta) &\Leftrightarrow \forall a : |a| \leq d, \sum_{b:|b|\leq d} \langle \beta, x^a x^b \rangle h_b = 0 \\
&\Leftrightarrow \forall a : |a| \leq d, \left\langle \beta, x^a \sum_{b:|b|\leq d} h_b x^b \right\rangle = 0 \\
&\Leftrightarrow \forall a : |a| \leq d, \langle \beta x^a, h \rangle = 0 \\
&\Leftrightarrow \forall q \in \mathsf{P}^n_d, \langle \beta q, h \rangle = 0 \\
&\Leftrightarrow \forall q \in \mathsf{P}^n_d, \langle \beta h, q \rangle = 0 \\
&\Leftrightarrow \beta h \in \mathsf{P}^n_{2n-d-1}.
\end{aligned}$$ ∎

Thus to prove Lemma 3.4, it suffices to show that $\operatorname{rank}(Q^\beta) = \dim(\mathsf{P}^n_d / B^n_{d,d}(\beta)) \geq 3^{d/9}$. Towards this end, we define

$$\Phi_{d,k}(D) := \min_{n > d/2, \beta \in \mathsf{P}^n_{2n} : \Delta(\beta, \mathsf{P}^n_{2n-2d-1}) > D} \dim(\mathsf{P}^n_k / B^n_{d,k}(\beta)). \tag{3.1}$$

In terms of $\Phi_{d,k}$, Lemma 3.4 now reduces to showing that $\Phi_{d,d}(3^{d/2}) \geq 3^{d/9}$. We obtain this lower bound by recursively bounding this quantity. The following serves as the base case of the recursion.

CLAIM 3.6. *For $k > 2d$, for all D, $\Phi_{d,k}(D) = 0$, and for $k \leq 2d$, $\Phi_{d,k}(1) \geq 1$.*

Proof. Let β be the polynomial that attains the minimum in (3.1). The first part of the claim follows from the fact that if $k > 2d$, then $B^n_{d,k}(\beta) = \mathsf{P}^n_k$.

Now for the second part. Since $\beta \notin \mathsf{P}^n_{2n-2d-1}$, there is a monomial x^a with $|a| \leq 2d$ such that

$$\langle \beta, x^a \rangle \neq 0 \Longleftrightarrow \langle \beta x^a, 1 \rangle \neq 0 \Longleftrightarrow \beta x^a \notin \mathsf{P}^n_{2n-1}.$$

If $|a| \leq k$, then $x^a \notin B^n_{d,k}(\beta)$ and we are done. Otherwise, consider b such that $b \leq a$ coordinatewise and $|b| = k$. Suppose $x^b \beta \in \mathsf{P}^n_{2n-2d-1+k}$; then $x^a \beta \in \mathsf{P}^n_{2n-1}$,

which is a contradiction. Hence, $x^b\beta \notin \mathsf{P}^n_{2n-2d-1+k}$, and the second part of the claim follows. ∎

For the induction step, we need the following claim.

CLAIM 3.7. *There exists a constant d_0 such that if $3^5 < \Delta < 3^d$, $d > d_0$ where β is Δ-far from $\mathsf{P}^n_{2n-2d-1}$, then there exists nonzero $\ell \in \mathsf{P}^n_1$ such that for all $c \in \mathbb{F}_3$, $\beta|_{\ell=c}$ are $\Delta/27$ far from the restriction of $\mathsf{P}^n_{2n-2d-1}$ to affine hyperplanes.*

See Appendix A for a proof of Claim 3.7 from Theorems 4.16 and 1.7 of [HSS13].

CLAIM 3.8. *If $3^5 \le D \le 3^d$ and $d > d_0$, then*

$$\Phi_{d,k}(D) \ge \Phi_{d-1,k}(D/27) + \Phi_{d-1,k-1}(D/27) + \Phi_{d-1,k-2}(D/27).$$

Proof. From Lemma 3.7, we get that there exists nonzero $\ell \in \mathsf{P}^n_1$ such that for all $c \in \mathbb{F}_3$, $\beta|_{\ell=c}$ is $D/27$ far from $\mathsf{P}^{n-1}_{2n-2d-1}$. By applying a change of basis, we can assume that $\ell = x_n$.

Let $\beta = (x_n^2-1)\gamma + x_n\eta + \theta$ and $q = (x_n^2-1)r + (x_n-1)s + t$, where $\gamma, \eta, \theta, r, s, t$ do not depend on x_n. Note that $\theta-\gamma, \theta+\eta, \theta-\eta$ are $D/27$ far from $\mathsf{P}^{n-1}_{2n-2d-1}$. Expanding the product βq, we have

$$\beta q = (x_n^2-1)\left((\theta-\gamma)r + \gamma t + \eta s - \gamma s\right) + (x_n-1)\left((\theta-\eta)s + \eta t\right) + (\theta+\eta)t.$$

Comparing terms, we observe that $\beta q \in \mathsf{P}^n_{2n-2d-1+k}$ iff the following three items are true:

1. $(\theta-\gamma)r + \gamma t + \eta s - \gamma s \in \mathsf{P}^{n-1}_{2n-2d-1+k-2}$,
2. $(\theta-\eta)s + \eta t \in \mathsf{P}^{n-1}_{2n-2d-1+k-1}$,
3. $(\theta+\eta)t \in \mathsf{P}^{n-1}_{2n-2d-1+k}$.

Since $r \in \mathsf{P}^n_{k-2}, s \in \mathsf{P}^n_{k-1}, t \in \mathsf{P}^n_k$, this is equivalent to the following (written in reverse order):

1. $t \in B^{n-1}_{d-1,k}(\theta+\eta)$,
2. $s \in -\eta t + B^{n-1}_{d-1,k-1}(\theta-\eta)$,
3. $r \in \gamma s - \eta s - \gamma t + B^{n-1}_{d-1,k-2}(\theta-\gamma)$.

Since t, s, r belong to sets with the same size as $B^{n-1}_{d-1,k}(\theta+\eta)$, $B^{n-1}_{d-1,k-1}(\theta-\eta)$, $B^{n-1}_{d-1,k-2}(\theta-\gamma)$, respectively, and each choice gives a distinct element of $B^n_{d,k}(\beta)$, we get the following equality:

$$\dim(B^n_{d,k}(\beta)) = \dim(B^{n-1}_{d-1,k}(\theta+\eta)) + \dim(B^{n-1}_{d-1,k-1}(\theta-\eta)) + \dim(B^{n-1}_{d-1,k-2}(\theta-\gamma)).$$

Combining this with $\dim(\mathsf{P}^n_k) = \dim(\mathsf{P}^{n-1}_k) + \dim(\mathsf{P}^{n-1}_{k-1}) + \dim(\mathsf{P}^{n-1}_{k-2})$, we obtain

$$\begin{aligned}\dim(\mathsf{P}^n_k/B^n_{d,k}(\beta)) &= \dim(\mathsf{P}^{n-1}_k/B^{n-1}_{d-1,k}(\theta+\eta)) + \dim(\mathsf{P}^{n-1}_{k-1}/B^{n-1}_{d-1,k-1}(\theta-\eta)) \\ &\quad + \dim(\mathsf{P}^{n-1}_{k-2}/B^{n-1}_{d-1,k-2}(\theta-\gamma)) \\ &\ge \Phi_{d-1,k}(D/27) + \Phi_{d-1,k-1}(D/27) + \Phi_{d-1,k-2}(D/27).\end{aligned}$$

The last inequality follows from the fact that $\theta-\gamma, \theta+\eta, \theta-\eta$ are $D/27$ far from $\mathsf{P}^{n-1}_{2n-2d-1} = \mathsf{P}^{n-1}_{2(n-1)-2(d-1)-1}$. Thus, the claim is proved. ∎

To prove Lemma 3.4, we start with $\Phi_{d,d}(3^{d/2})$ and apply Claim 3.8 recursively $d/6-2$ times and finally use the base case from Claim 3.6 (this can be done as long as $d/6-2 \le d/2$). This gives $\operatorname{rank}(Q^\beta) \ge \Phi_{d,d}(3^{d/2}) \ge 3^{d/6-2} \ge 3^{d/9}$ as long as d_0 is large enough. ∎

4. Hardness of coloring 2-colorable 8-uniform hypergraphs. We prove the theorem by a reduction from 3SAT via the instances of the repeated label cover problem obtained in Theorem 2.2. Let $r \in \mathbb{N}$ be a parameter, which we will fix later, and let $I(\varphi)$ be an instance of the r-repeated label cover obtained in Theorem 2.2 starting from a 3SAT instance φ.

We denote by $G = (U, V, E)$ the underlying left and right regular bipartite graph. For $u \in U$ and $i \in [3r]$, fix functions $f_i^u : \{0,1\}^{3r} \to \{0,1\}$ as in Definition 2.1. Throughout this section, we work over $\mathbb{F}_2$. For a degree parameter d, which we will determine later, and a vertex $u \in U$, we define the subspace $J_u := \{\sum_{i=1}^{3r} r_i f_i^u : r_i \in \mathsf{P}^{3r}_{(d-3)}\}$. Note that since each f_i^u depends only on 3 variables, it is a polynomial of degree at most 3 and hence, J_u is indeed a subspace of P^{3r}_d. Let N_u denote the cardinality of the quotient space P^{3r}_d / J_u.

We now define the hypergraph H produced by the reduction. The vertices of H—denoted $V(H)$—are obtained by replacing each $u \in U$ by a block $\mathcal{B}_u$ of N_u vertices, which we identify with elements of P^{3r}_d / J_u. Let N denote $|V(H)| = \sum_{u \in U} N_u$.

We think of a 2-coloring of $V(H)$ as a map from $V(H)$ to $\mathbb{F}_2$. Given a coloring $A : V(H) \to \mathbb{F}_2$, we denote by $A_u : \mathsf{P}^{3r}_d / J_u \to \mathbb{F}_2$ the restriction of A to the block $\mathcal{B}_u$ (under our identification of $\mathcal{B}_u$ with P^{3r}_d / J_u). Let $A'_u : \mathsf{P}^{3r}_d \to \mathbb{F}_2$ denote the lift of A_u as defined in Fact 2.14.

The (weighted) edge set $E(H)$ of H is specified implicitly by the following PCP verifier for the label cover instance $I(\varphi)$, which expects as its input a 2-coloring $A : V(H) \to \mathbb{F}_2$.

2-Color 8-Uniform Test(d).

1. Choose a uniformly random $v \in V$ and then choose $u, w \in U$ uniformly random neighbors of v (by the right regularity of G, both (u, v) and (u, w) are uniform random edges in E). Let π denote $\pi_{uv} : \mathbb{F}_2^{3r} \to \mathbb{F}_2^r$ and similarly, let π' be π_{wv}.
2. Choose $f \in \mathsf{P}^r_d$, $e_1, e_2, e_3, e_4 \in \mathsf{P}^{3r}_d$, $g_1, g_2 \in \mathsf{P}^{3r}_{d/4}$, and $h_1, h_2, h_3, h_4 \in \mathsf{P}^{3r}_{3d/4}$ independently and uniformly at random. Define functions $\eta_1, \eta_2, \eta_3, \eta_4 \in \mathsf{P}^{3r}_d$ as follows:
$$\eta_1 := 1 + f \circ \pi + g_1 h_1, \qquad \eta_3 := f \circ \pi' + g_2 h_3,$$
$$\eta_2 := 1 + f \circ \pi + (1 + g_1) h_2, \qquad \eta_4 := f \circ \pi' + (1 + g_2) h_4.$$
3. Accept iff $A'_u(e_1), A'_u(e_1 + \eta_1), A'_u(e_2), A'_u(e_2 + \eta_2), A'_w(e_3), A'_w(e_3 + \eta_3), A'_w(e_4), A'_w(e_4 + \eta_4)$ are not all equal.

We now analyze the above test.

LEMMA 4.1 (completeness). *If φ is satisfiable, then there exists a 2-coloring $A : V(H) \to \mathbb{F}_2$ such that the verifier accepts with probability 1. In other words, the hypergraph H is 2-colorable.*

Proof. Since φ is satisfiable, Theorem 2.2 tells us that there are labelings $L_U : U \to \mathbb{F}_2^{3r}$ and $L_V : V \to \mathbb{F}_2^r$ such that for all $u \in U$, $L_U(u)$ satisfies all the clauses in U and moreover, for every edge $(u, v) \in E$, we have $\pi_{uv}(L_U(u)) = L_V(v)$. Fix such L_U, L_V. Let a_u denote $L_U(u)$ for any $u \in U$ and let b_v denote $L_V(v)$ for any $v \in V$.

Now, the coloring $A : V(H) \to \mathbb{F}_2$ is defined to ensure that for each $u \in U$, its restriction A_u is such that its lift $A'_u = \mathsf{LC}_d(a_u)$. Note that this makes sense since $\mathsf{LC}_d(a_u)$ is folded over J_u: Indeed, given any $g \in \mathsf{P}^{3r}_d$ and $h = \sum_i r_i f_i^u \in J_u$, we have $\mathsf{LC}_d(a_u)(g + h) = g(a_u) + h(a_u) = g(a_u)$ as $h(a_u) = \sum_i r_i(a_u) f_i^u(a_u) = 0$ for any satisfying assignment a_u of the clauses corresponding to u.

We now show that the verifier accepts A with probability 1. Fix any choices of $v \in V$, $u, w \in U$, f, e_i, h_i $(i \in [4])$, and g_i $(i \in [2])$ as in the test. By the definitions of L_U and L_V, we must have $\pi(a_u) = \pi'(a_w) = b_v$. This implies that the 8 positions in A viewed by the verifier, respectively, contain the following values:

$$\begin{aligned}
&e_1(a_u), \quad e_1(a_u) + 1 + f(b_v) + g_1(a_u)h_1(a_u),\\
&e_2(a_u), \quad e_2(a_u) + 1 + f(b_v) + (1 + g_1(a_u))h_2(a_u),\\
&e_3(a_w), \quad e_3(a_w) + f(b_v) + g_2(a_w)h_3(a_w),\\
&e_4(a_w), \quad e_4(a_w) + f(b_v) + (1 + g_2(a_w))h_4(a_w).
\end{aligned}$$

If $f(b_v) = 0$, then

$$\begin{aligned}
&e_1(a_u) \neq e_1(a_u) + 1 + f(b_v) + g_1(a_u)h_1(a_u) \quad \text{or}\\
&e_2(a_u) \neq e_2(a_u) + 1 + f(b_v) + (1 + g_1(a_u))h_2(a_u).
\end{aligned} \tag{4.1}$$

Else if $f(b_v) = 1$, then

$$\begin{aligned}
&e_3(a_w) \neq e_3(a_w) + f(b_v) + g_2(a_w)h_3(a_w) \quad \text{or}\\
&e_4(a_w) \neq e_4(a_w) + f(b_v) + (1 + g_2(a_w))h_4(a_w).
\end{aligned} \tag{4.2}$$

Thus, the verifier always accepts. ∎

Remark 4.2. Lemma 4.1 actually yields a stronger statement. Let us group the probes of the verifier as $(e_i, e_i + \eta_i)$ for $i \in [4]$. Then, for the given coloring A in Lemma 4.1 and any random choices of the verifier, there is some $i \in [4]$ such that A is not constant on inputs in the ith group. We use this in section 5 to devise a 4-query verifier over an alphabet of size 4.

LEMMA 4.3 (soundness). *Let $d \geq 16$ be a multiple of 4, let $\delta > 0$, and let ε_0 be the constant from Theorem 2.2. If φ is unsatisfiable and H contains an independent set of size δN, then $\delta^8 \leq 2^{d/2+1} \cdot 2^{-\varepsilon_0 r} + 2^{-2^{d/8}}$.*

Proof. Fix any independent set $\mathcal{I} \subseteq V(H)$ of size δN. Let $A : V(H) \to \{0,1\}$ be the indicator function of $\mathcal{I}$. For $u \in U$, let $A_u : \mathsf{P}_d^{3r}/J_u \to \{0,1\}$ denote the restriction of A to the block of vertices corresponding to u, and let $A'_u : \mathsf{P}_d^{3r} \to \{0,1\}$ be the lift of A_u. Note that we have $\mathbb{E}_{(g+J_u)\in \mathsf{P}_d^{3r}/J_u}[A_u(g)] = \mathbb{E}_{g\in\mathsf{P}_d^{3r}}[A'_u(g)]$ for any $u \in U$. In particular,

$$\mathop{\mathbb{E}}_{u\in U}\mathop{\mathbb{E}}_{g\in\mathsf{P}_d^{3r}}[A'_u(g)] = \mathop{\mathbb{E}}_{u\in U}\mathop{\mathbb{E}}_{(g+J_u)\in\mathsf{P}_d^{3r}/J_u}[A_u(g)] \geq \delta. \tag{4.3}$$

Since $\mathcal{I}$ is an independent set, in particular it must be the case that the probability that a random edge (chosen according to the probability distribution defined on $E(H)$ by the PCP verifier) completely lies inside $\mathcal{I}$ is 0. We note that another expression for this probability is given by the quantity $\mathbb{E}_{v\in V,u,w\in U}[Q(v,u,w)]$, where $v \in V$ and $u, w \in U$ are as chosen by the PCP verifier described above and $Q(v,u,w)$ is defined as follows:

$$Q(v,u,w) := \mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\mathop{\mathbb{E}}_{\substack{e_1,e_2\\ e_3,e_4}}\left[\prod_{i\in[2]} A'_u(e_i)A'_u(e_i+\eta_i)A'_w(e_{i+2})A'_w(e_{i+2}+\eta_{i+2})\right]\right].$$

We analyze the right-hand side of the above using its Fourier expansion (see Lemma 2.10). As defined in section 2.3, let Λ_d^{3r} be a set of minimum weight coset

representatives of the cosets of $(\mathsf{P}_d^{3r})^\perp$ in $\mathfrak{F}_{3^r}$. Standard computations yield the following:

$$Q(v,u,w) = \sum_{\substack{\alpha_1,\alpha_2 \\ \beta_1,\beta_2 \in \Lambda_d^{3r}}} \underbrace{\left(\prod_{i\in[2]} \widehat{A'_u}(\alpha_i)^2 \widehat{A'_w}(\beta_i)^2\right) \mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2 \\ \eta_3,\eta_4}} \left[\prod_{i\in[2]} \chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]}_{\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)}. \tag{4.4}$$

When v,u,w are clear from context, we use $\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)$ instead of $\xi_{v,u,w}(\alpha_1, \alpha_2,\beta_1,\beta_2)$.

We analyze the above expression by breaking it up as follows. Let

$$\begin{aligned} \text{FAR} &:= \{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in (\Lambda_d^{3r})^4 : \max\{\Delta(\alpha_i,\mathsf{P}_d^{3r}),\Delta(\beta_i,\mathsf{P}_d^{3r})\} \geq 2^{d/2}\}, \\ \text{NEAR} &:= (\Lambda_d^{3r})^4 \setminus \text{FAR}. \end{aligned}$$

We now make the following claim for every v,u,w, the proof of which is deferred to the end of the section.

Claim 4.4. *For $d \geq 16$, we have $\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{FAR}} |\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)| \leq 2^{-2^{d/8}}$.*

Substituting in (4.4), we have, for any $v \in V$ and $u,w \in U$,

$$\begin{aligned} Q(v,u,w) &\geq \sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}} \xi(\alpha_1,\alpha_2,\beta_1,\beta_2) - \sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{FAR}} |\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)| \\ &\geq \sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}} \xi(\alpha_1,\alpha_2,\beta_1,\beta_2) - 2^{-2^{-d/8}}. \end{aligned} \tag{4.5}$$

Now fix any $(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \text{NEAR}$. We analyze the expectation term in $\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)$ further as follows:

$$\begin{aligned} &\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2 \\ \eta_3,\eta_4}} \left[\prod_{i\in[2]} \chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right] \\ &= \mathop{\mathbb{E}}_{\substack{g_1,g_2,f \\ h_1,\ldots,h_4}} [\chi_{\alpha_1}(1+f\circ\pi+g_1h_1)\chi_{\alpha_2}(1+f\circ\pi+(1+g_1)h_2)\chi_{\beta_1}(f\circ\pi'+g_2h_3) \\ &\qquad \cdot \chi_{\beta_2}(f\circ\pi'+(1+g_2)h_4)] \\ &= \mathop{\mathbb{E}}_{g_i,h_j} \left[\prod_{i\in[2]} \chi_{\alpha_i}(1+(1+i+g_1)h_i)\chi_{\beta_i}((1+i+g_2)h_{i+2}) \cdot \mathop{\mathbb{E}}_f \left[\chi_{\pi_2(\alpha_1+\alpha_2)+\pi'_2(\beta_1+\beta_2)}(f)\right]\right], \end{aligned} \tag{4.6}$$

where π_2 and π'_2 are as defined in Lemma 2.11. The innermost expectation is 0 unless $\chi_{\pi_2(\alpha_1+\alpha_2)+\pi'_2(\beta_1+\beta_2)}$ is the trivial character on P_d^r or, equivalently, $\gamma := \pi_2(\alpha_1+\alpha_2)+\pi'_2(\beta_1+\beta_2) \in (\mathsf{P}_d^r)^\perp$.

We claim that this implies that $\gamma = 0$. To see this, we observe from the definition of π_2 and π'_2 that $|\operatorname{support}(\gamma)| \leq \sum_{i\in[2]} |\operatorname{support}(\alpha_i)| + |\operatorname{support}(\beta_i)| \leq 4\cdot 2^{d/2}$, since $(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \text{NEAR}$ and $|\operatorname{support}(\alpha)| = \Delta(\alpha,(\mathsf{P}_d^{3r})^\perp)$ for $\alpha \in \Lambda_d^{3r}$. However, if $\gamma \neq 0$ and $\gamma \in (\mathsf{P}_d^r)^\perp$, by Lemma 2.7, we must have $|\operatorname{support}(\gamma)| \geq 2^d > 4\cdot 2^{d/4}$ since

$d \geq 8$. This implies that $\gamma = 0$. Substituting in (4.6), we get

$$\begin{aligned}(4.7)\quad &\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]} \chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]\\ &= \begin{cases} 0 & \text{if } \pi_2(\alpha_1+\alpha_2)+\pi_2'(\beta_1+\beta_2)\neq 0,\\ \mathbb{E}_{g_j,h_i}\left[\prod_{i\in[2]} \chi_{\alpha_i}(1+(1+i+g_1)h_i)\chi_{\beta_i}((1+i+g_2)h_{i+2})\right] & \text{otherwise.}\end{cases}\end{aligned}$$

Substituting back in (4.5), we have

$$(4.8)\qquad Q(v,u,w) = \sum_{\substack{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR}:\\ \pi_2(\alpha_1+\alpha_2)+\pi_2'(\beta_1+\beta_2)=0}} \xi(\alpha_1,\alpha_2,\beta_1,\beta_2) - 2^{-2^{d/8}}.$$

We partition the terms in the above sum further into

$$\begin{aligned}\text{NEAR}_0 &:= \{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR} : \pi_2(\alpha_1+\alpha_2)=\pi_2'(\beta_1+\beta_2)=0\},\\ \text{NEAR}_1 &:= \{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR} : \pi_2(\alpha_1+\alpha_2)=\pi_2'(\beta_1+\beta_2)\neq 0\}\end{aligned}$$

and make the following claims about the contributions of these subsets.

CLAIM 4.5. $\mathbb{E}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR}_1} |\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right] \leq 2^{d/2+1}\cdot 2^{-\varepsilon_0 r}$.

CLAIM 4.6. *Let δ be the fractional size of the independent set.*

$$\mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR}_0} \xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right] \geq \delta^8.$$

Assuming these claims for now, we can finish the proof of Lemma 4.3 as follows. By (4.8),

$$\begin{aligned}0 &= \mathop{\mathbb{E}}_{v,u,w}[Q(v,u,w)]\\ &\geq \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR}_0} \xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right]\\ &\quad - \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR}_1} |\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right] - 2^{-2^{d/8}}\\ &\geq \delta^8 - 2^{d/2+1}\cdot 2^{-\varepsilon_0 r} - 2^{-2^{d/8}}.\end{aligned}$$

∎

We now turn to the proofs of Claims 4.4–4.6.

Proof of Claim 4.4. Fix any $(\alpha_1,\alpha_2,\beta_1,\beta_2)\in$ FAR. Conditioned on any choice of

f, the expectation term in $|\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)|$ may be bounded as follows:

(4.9)
$$\begin{aligned}
&\left|\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]}\chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]\right| \\
&= \left|\mathop{\mathbb{E}}_{\substack{g_1,g_2\\ h_1,\ldots,h_4}}[\chi_{\alpha_1}(1+f\circ\pi+g_1h_1)\chi_{\alpha_2}(1+f\circ\pi+(1+g_1)h_2)\chi_{\beta_1}(f\circ\pi'+g_2h_3)\right. \\
&\qquad\left.\cdot\,\chi_{\beta_2}(f\circ\pi'+(1+g_2)h_4)]\right| \\
&\le \mathop{\mathbb{E}}_{g_1,g_2}\left[\prod_{i\in[2]}\left|\mathop{\mathbb{E}}_{h_i}[\chi_{\alpha_i}(1+f\circ\pi+(1+i+g_1)h_i)]\right|\cdot\left|\mathop{\mathbb{E}}_{h_{i+2}}[\chi_{\beta_i}(f\circ\pi'+(1+i+g_2)h_{i+2})]\right|\right] \\
&= \mathop{\mathbb{E}}_{g_1,g_2}\left[\prod_{i\in[2]}\left|\mathop{\mathbb{E}}_{h_i}[\chi_{\alpha_i}((1+i+g_1)h_i)]\right|\cdot\left|\mathop{\mathbb{E}}_{h_{i+2}}[\chi_{\beta_i}((1+i+g_2)h_{i+2})]\right|\right] \\
&\le \mathop{\mathbb{E}}_{g_1,g_2}\left[\min\left\{\left|\mathop{\mathbb{E}}_{h_i}[\chi_{\alpha_i}((1+i+g_1)h_i)]\right|,\left|\mathop{\mathbb{E}}_{h_{i+2}}[\chi_{\beta_i}((1+i+g_2)h_{i+2})]\right| : i\in[2]\right\}\right] \\
&\le \min\left\{\mathop{\mathbb{E}}_{g_1}\left[\left|\mathop{\mathbb{E}}_{h_i}[\chi_{\alpha_i}((1+i+g_1)h_i)]\right|\right],\mathop{\mathbb{E}}_{g_2}\left[\left|\mathop{\mathbb{E}}_{h_{i+2}}[\chi_{\beta_i}((1+i+g_2)h_{i+2})]\right|\right] : i\in[2]\right\}.
\end{aligned}$$

Note that for any $i\in[2]$, $(1+i+g_1)$ and $(1+i+g_2)$ are uniformly random elements of $\mathsf{P}^{3r}_{d/4}$ that are independent of $h_1,\ldots,h_4$. Moreover, since $(\alpha_1,\alpha_2,\beta_1,\beta_2)\in$ FAR, we know that there is a $\gamma\in\{\alpha_1,\alpha_2,\beta_1,\beta_2\}$ such that $\Delta(\gamma,(\mathsf{P}^{3r}_d)^\perp)\ge 2^{d/2}$. Therefore, by Theorem 2.12, we have

$$\mathop{\mathbb{E}}_{g\in\mathsf{P}^{3r}_{d/4}}\left[\left|\mathop{\mathbb{E}}_{h\in\mathsf{P}^{3r}_{3d/4}}[\chi_\gamma(gh)]\right|\right]\le 2^{-2^{(d/4-2)}}\le 2^{-2^{d/8}},$$

where the second inequality follows because $d\ge 16$. Substituting the above in (4.9), we obtain

$$\left|\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]}\chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]\right|\le 2^{-2^{d/8}}.$$

Thus, we obtain

$$\begin{aligned}
&\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\mathrm{FAR}}|\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)| \\
&\qquad\le 2^{-2^{d/8}}\cdot\sum_{\alpha_1,\alpha_2,\beta_1,\beta_2\in\Lambda^{3r}_d}\left(\prod_{i\in[2]}\widehat{A'_u}(\alpha_i)^2\widehat{A'_w}(\beta_i)^2\right)\le 2^{-2^{d/8}},
\end{aligned}$$

where the last inequality follows from Parseval's identity and the fact that $|A(x)|\le 1$ for all $x\in V(H)$. ∎

Proof of Claim 4.5. We use a Fourier decoding argument. Formally, we sample random labelings $L_U: U\to\mathbb{F}_2^{3r}$ and $L_V: V\to\mathbb{F}_3^r$ such that $L_U(u)$ satisfies all clauses

in u and such that

$$\Pr_{(u,v)\in E, L_U, L_V}[\pi_{uv}(L_U(u)) = L_V(v)] \geq \frac{1}{2^{d+2}} \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_1} |\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right]. \tag{4.10}$$

Since $OPT(I(\varphi)) \leq 2^{-\varepsilon_0 r}$, the left-hand side of the above inequality is at most $2^{-\varepsilon_0 r}$. This implies the claim.

Define $L_U : U \to \mathbb{F}_2^{3r}$ as follows: Given $u \in U$, we sample a random pair $\alpha_1, \alpha_2 \in \Lambda_d^{3r}$ such that $|\alpha_1|, |\alpha_2| < 2^{d/2}$ with probability proportional to $\widehat{A'_u}(\alpha_1)^2 \widehat{A'_u}(\alpha_2)^2$, and set $L_U(u)$ to be a_u for a uniformly random a_u chosen from $\mathrm{support}(\alpha_1) \cup \mathrm{support}(\alpha_2)$. Since $|\alpha_1|, |\alpha_2| < 2^{d/2} < 2^{d-4}$, by Lemma 2.15, any α_1, α_2 sampled as above is supported only on satisfying assignments of all the clauses in u.

We also define $L_V : V \to \mathbb{F}_2^r$ similarly: Given $v \in V$, we sample a random neighbor $w \in U$ of v and choose at random a pair $\beta_1, \beta_2 \in \Lambda_d^{3r}$ such that $|\beta_1|, |\beta_2| < 2^{d/2}$ with probability proportional to $\widehat{A'_w}(\beta_1)^2 \widehat{A'_w}(\beta_2)^2$, and set $L_V(v)$ to be $\pi_{wv}(a_w)$ for a uniformly random a_w chosen from $\mathrm{support}(\beta_1) \cup \mathrm{support}(\beta_2)$.

Let $(u, v) \in E$ be a uniformly random edge of G and consider the probability that $\pi_{uv}(L_U(u)) = L_V(v)$. This probability can clearly be lower bounded as follows:

$$\Pr_{(u,v)\in E, L_U, L_V}[\pi(L_U(u)) = L_V(v)] \geq \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{\substack{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}: \\ \pi(\mathrm{support}(\alpha_1)\cup\mathrm{support}(\alpha_2))\cap \\ \pi'(\mathrm{support}(\beta_1)\cup\mathrm{support}(\beta_2))\neq\emptyset}} \prod_{i\in[2]} \widehat{A'_u}(\alpha_i)^2 \widehat{A'_w}(\beta_i)^2\right] \cdot \frac{1}{2^{d+2}},$$

where π denotes π_{uv} and π' denotes π_{wv}. Observe that if $(\alpha_1, \alpha_2, \beta_1, \beta_2) \in \mathrm{NEAR}_1$, then $\pi_2(\alpha_1 + \alpha_2) = \pi'_2(\beta_1 + \beta_2) \neq 0$ and, in particular,

$$\pi(\mathrm{support}(\alpha_1) \cup \mathrm{support}(\alpha_2)) \cap \pi'(\mathrm{support}(\beta_1) \cup \mathrm{support}(\beta_2)) \neq \emptyset.$$

Therefore, we get the following, which implies (4.10) and hence proves the claim:

$$\Pr_{(u,v)\in E, L_U, L_V}[\pi(L_U(u)) = L_V(v)] \geq \frac{1}{2^{d+2}} \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_1} \prod_{i\in[2]} \widehat{A'_u}(\alpha_i)^2 \widehat{A'_w}(\beta_i)^2\right]. \qquad \square$$

Proof of Claim 4.6. We argue below that for any $v \in V$ and its neighbors $u, w \in U$ and any $(\alpha_1, \alpha_2, \beta_1, \beta_2) \in \mathrm{NEAR}_0$,

$$\xi(\alpha_1, \alpha_2, \beta_1, \beta_2) \geq 0. \tag{4.11}$$

Given (4.11), we have

$$\mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_0} \xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right] \geq \mathop{\mathbb{E}}_{v,u,w}[\xi_{v,u,w}(0,0,0,0)] = \mathop{\mathbb{E}}_{v,u,w}\left[\widehat{A'_u}(0)^4 \widehat{A'_w}(0)^4\right].$$

Conditioned on $v \in V$, u and w are independent and randomly chosen neighbors of v. Thus, the above may be further lower bounded as follows:

$$\mathop{\mathbb{E}}_{v,u,w}\left[\widehat{A'_u}(0)^4\widehat{A'_w}(0)^4\right] = \mathop{\mathbb{E}}_{v}\left[\left(\mathop{\mathbb{E}}_{u:(u,v)\in E}\left[\widehat{A'_u}(0)^4\right]\right)^2\right]$$
$$\geq \left(\mathop{\mathbb{E}}_{(u,v)\in E}\left[\widehat{A'_u}(0)\right]\right)^8 = \left(\mathop{\mathbb{E}}_{u\in U, g\in \mathbb{P}^{3r}_d}[A'_u(g)]\right)^8 \geq \delta^8,$$

where the first inequality follows from repeated applications of the Cauchy–Schwarz inequality and the last from (4.3).

For any v, u, w, and $(\alpha_1, \alpha_2, \beta_1, \beta_2) \in \text{NEAR}_0$, it remains to prove (4.11) (i.e., nonnegativity of $\xi_{v,u,w}(\alpha_1, \alpha_2, \beta_1, \beta_2)$). From (4.4), it suffices to argue the nonnegativity of

(4.12)
$$\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]}\chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]$$
$$= \mathop{\mathbb{E}}_{g_1,g_2}\left[\prod_{i\in[2]}\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i}(1+(1+i+g_1)h_i)\right]\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i}((1+i+g_2)h_{i+2})\right]\right]$$
$$= \mathop{\mathbb{E}}_{g_1,g_2}\left[(-1)^{\sum_x \alpha_1(x)+\alpha_2(x)}\cdot\prod_{i\in[2]}\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i(1+i+g_1)}(h_i)\right]\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i(1+i+g_2)}(h_{i+2})\right]\right],$$

where we have used (4.7) for the first equality and the fact that $\chi_\alpha(gh) = \chi_{\alpha g}(h)$ for the second. We claim that all the terms inside the final expectation are nonnegative.

First, since $(\alpha_1, \alpha_2, \beta_1, \beta_2) \in \text{NEAR}_0$, we have that $\pi_2(\alpha_1 + \alpha_2) = 0$ and hence $(-1)^{\sum_x \alpha_1(x)+\alpha_2(x)} = (-1)^{\sum_y \pi_2(\alpha_1+\alpha_2)(y)} = 1$. Second, the orthonormality of characters implies that for any $\alpha \in \mathfrak{F}_{3r}$, we have that $\mathbb{E}_{h\in\mathbb{P}^r_{3d/4}}[\chi_\alpha(h)] \in \{0, 1\}$ and hence nonnegative.

This shows that the right-hand side of (4.12) is nonnegative and hence proves (4.11). ∎

Proof of Theorem 1.1. Given the completeness (see Lemma 4.1) and soundness (see Lemma 4.3), we only need to fix parameters. Let $d = C \log r$ for a large enough constant $C \geq 16$ to be determined shortly. By Lemma 4.3, if H has an independent set of size δN, then $\delta^8 \leq 2^{d/2}\cdot 2^{-\varepsilon_0 r} + 2^{-2^{d/8}} < 2^{-\varepsilon_0 r/2}$ for large enough $C > 0$ and $r \in \mathbb{N}$. Hence, H has no independent sets of $\delta' N$, where $\delta' = 2^{-\varepsilon_0 r/16}$.

The hypergraph H can be produced in time polynomial in $N = n^{O(r)}2^{r^{O(d)}} = n^{O(r)}2^{r^{O(\log r)}}$. Setting $r = 2^{\Theta(\sqrt{\log\log n})}$, we get $N = n^{2^{O(\sqrt{\log\log n})}}$ and $\delta' = 2^{-\Omega(r)} = 2^{-2^{\Theta(\sqrt{\log\log n})}} = 2^{-2^{\Theta(\sqrt{\log\log N})}}$, proving Theorem 1.1. ∎

5. Hardness of coloring 4-colorable 4-uniform hypergraphs. This construction is motivated by Remark 4.2 above. We observe that the 8-query PCP test used in the above inapproximability result has a stronger completeness guarantee than required to prove the above result: The 8 queries of the not-all-equal (NAE) PCP test, say $\{e_i, e'_i\}_{i=1}^4$ in the completeness case, satisfy

$$\bigvee_{i=1}^{4}\mathsf{NAE}(A(e_i), A(e'_i)),$$

which is stronger than the required

$$\mathsf{NAE}(A(e_1), A(e_1'), A(e_2), A(e_2'), A(e_3), A(e_3'), A(e_4), A(e_4')).$$

Furthermore, for each $i \in \{1, 4\}$, the queries $e_i, e_i', e_{i+1}, e_{i+1}'$ appear in the same table. This lets us perform the following "doubling of queries": Each location is now indexed by a pair of queries, e.g., (e_1, e_2), and is expected to return 2 bits that are the answers to the two queries, respectively. The stronger completeness property yields a 4-query NAE PCP test over an alphabet of size 4 with the completeness property,

$$\mathsf{NAE}(B(e_1, e_2), B(e_1', e_2')) \vee \mathsf{NAE}(B(e_3, e_4), B(e_3', e_4')),$$

which suffices for the completeness for proving inapproximability results for 4-colorable 4-uniform hypergraphs. We show that the soundness analysis also carries over to yield the following hardness for 4-colorable 4-uniform hypergraphs.

We remark that the doubling method, mentioned above, when used in the vanilla long code setting (as opposed to low-degree long code setting) already yields the following inapproximability: It is quasi-NP-hard to color a 4-colorable 4-uniform hypergraph with $(\log N)^{\Omega(1)}$ colors. This result already improves upon the above mentioned result of Khot [Kho02a] for 7-colorable 4-uniform hypergraphs. Another feature of the doubling method is that although the underlying alphabet is of size 4, namely $\{0, 1\}^2$, it suffices for the soundness analysis to perform standard Fourier analysis over $\mathbb{F}_2$.

In the language of covering complexity,[2] (the proof of) Theorem 1.2 demonstrates a Boolean 4CSP for which it is quasi-NP-hard to distinguish between a covering number of 2 versus $\exp(\sqrt{\log \log N})$. The previous best result for a Boolean 4CSP was 2 versus $\log \log N$, due to Dinur and Kol [DK13].

We now turn to the formal construction of the verifier each of whose queries correspond to 2 queries of the verifier described above. Let $I(\varphi)$, $G = (U, V, E)$, and J_u $(u \in U)$ be defined as in section 4.

Now the vertices of the hypergraph H produced by the reduction denoted by $V(H)$ are obtained by replacing each $u \in U$ by a block $\mathcal{B}_u$ of N_u^2 vertices, which we identify with elements of $\mathsf{P}_d^{3r}/J_u \times \mathsf{P}_d^{3r}/J_u$. Let N denote $|V(H)| = \sum_{u \in U} N_u^2$.

We think of a 4-coloring of $V(H)$ as a map from $V(H)$ to the 4-element set $\mathbb{F}_2 \times \mathbb{F}_2$. Given a coloring $A : V(H) \to \mathbb{F}_2 \times \mathbb{F}_2$, we denote by $A_u : \mathsf{P}_d^{3r}/J_u \times \mathsf{P}_d^{3r}/J_u \to \mathbb{F}_2 \times \mathbb{F}_2$ the restriction of A to the block $\mathcal{B}_u$. Let $A_u' : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \mathbb{F}_2 \times \mathbb{F}_2$ denote the lift of A_u, as defined by $A_u'(g_1, g_2) := A_u(g_1 + J_u, g_2 + J_u)$.

The verifier is defined as follows. The verifier is identical to the verifier in section 4 but for the doubling of queries.

4-Color 4-Uniform Test(d).

1. Choose a uniformly random $v \in V$ and then choose $u, w \in U$ uniformly random neighbors of v. Let π denote $\pi_{uv} : \mathbb{F}_2^{3r} \to \mathbb{F}_2^r$ and similarly, let π' be π_{wv}.
2. Choose $f \in \mathsf{P}_d^r$, $e_1, e_2, e_3, e_4 \in \mathsf{P}_d^{3r}$, $g_1, g_2 \in \mathsf{P}_{d/4}^{3r}$, and $h_1, h_2, h_3, h_4 \in \mathsf{P}_{3d/4}^{3r}$ independently and uniformly at random. Define functions $\eta_1, \eta_2, \eta_3, \eta_4 \in \mathsf{P}_d^{3r}$ as follows:

$$\eta_1 := 1 + f \circ \pi + g_1 h_1, \qquad \eta_3 := f \circ \pi' + g_2 h_3,$$
$$\eta_2 := 1 + f \circ \pi + (1 + g_1) h_2, \qquad \eta_4 := f \circ \pi' + (1 + g_2) h_4.$$

[2] The covering number of a CSP is the minimal number of assignments to the vertices so that each hyperedge is covered by at least one assignment.

3. Accept iff $A'_u(e_1,e_2)$, $A'_u(e_1+\eta_1,e_2+\eta_2)$, $A'_w(e_3,e_4)$, $A'_w(e_3+\eta_3,e_4+\eta_4)$ are not all equal.

The analysis of the above test closely follows that of the 2-color 8-uniform test.

LEMMA 5.1 (completeness). *If φ is satisfiable, then there exists a 4-coloring $A : V(H) \to \mathbb{F}_2 \times \mathbb{F}_2$ such that the verifier accepts with probability 1. In other words, the hypergraph H is 4-colorable.*

Proof. The proof follows directly from Remark 4.2. ∎

The soundness lemma requires us to perform Fourier analysis on functions $A : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \{0,1\}$, for which we need the following easily verifiable facts.

FACT 5.2. *Let $A : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \mathbb{C}$ be any function. A nonzero function $\chi : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \mathbb{C}$ is a character if $\chi(g_1+h_1, g_2+h_2) = \chi(g_1,g_2)\chi(h_1,h_2)$.*

- *$\chi : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \mathbb{C}$ is a character iff there exist $(\alpha_1,\alpha_2) \in \mathfrak{F}_{3r} \times \mathfrak{F}_{3r}$ such that $\chi(g_1,g_2) = \chi_{\alpha_1}(g_1)\chi_{\alpha_2}(g_2)$ for any $g_1, g_2 \in \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r}$, where χ_{α_1} and χ_{α_2} are characters of P_d^{3r}.*
- *(α_1,α_2) and (β_1,β_2) yield the same character iff $(\alpha_1-\beta_1), (\alpha_2-\beta_2) \in (\mathsf{P}_d^{3r})^\perp$.*
- *Folding: Let $A : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \mathbb{C}$ be any function folded over the subgroup $J \times J$, where $J := \{\sum_{i=1}^k r_i q_i : r_i \in \mathsf{P}_{d-3}^{3r}\}$ and $q_1, \ldots, q_k \in \mathsf{P}_3^{3r}$. Then, for any $(\alpha_1,\alpha_2) \in \mathfrak{F}_{3r} \times \mathfrak{F}_{3r}$ such that $|\alpha_j| := \Delta(\alpha_j, (\mathsf{P}_d^{3r})^\perp) < 2^{d-3}$ for $j \in \{1,2\}$ and $\widehat{A}(\alpha_1,\alpha_2) \neq 0$, it must be the case that $\mathrm{support}(\alpha_1) \cup \mathrm{support}(\alpha_2)$ only contains x such that $q_i(x) = 0$ for each $i \in [k]$.*

LEMMA 5.3 (soundness). *Let $d \geq 8$ be a multiple of 4, let $\delta > 0$, and let ε_0 be the constant from Theorem 2.2. If φ is unsatisfiable and H contains an independent set of size δN, then $\delta^4 \leq 2^{d/2+1} \cdot 2^{-\varepsilon_0 r} + 2^{-4 \cdot 2^{-d/4}}$.*

The proof of Lemma 5.3 is similar to the proof of Lemma 4.3. The parameters are set exactly as in Theorem 1.1 to yield Theorem 1.2.

Proof of Lemma 5.3. As the proof is similar to that of Lemma 4.3, we only give a proof sketch, highlighting the salient differences.

As before, fix any independent set $\mathcal{I} \subseteq V(H)$ of size δN. Let $A : V(H) \to \{0,1\}$ be the indicator function of $\mathcal{I}$. We have $\mathbb{E}_{u \in U} \mathbb{E}_{g_1,g_2 \in \mathsf{P}_d^{3r}} [A'_u(g_1,g_2)] \geq \delta$.

Again, we analyze $\mathbb{E}_{v \in V, u,w \in U}[Q(v,u,w)]$, which gives the probability that a random edge (chosen according to the probability distribution defined on $E(H)$ by the PCP verifier) completely lies inside the independent set $\mathcal{I}$ and is hence 0. Here, $Q(v,u,w)$ is defined as follows:

$$Q(v,u,w) := \mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2 \\ \eta_3,\eta_4}} \left[\mathop{\mathbb{E}}_{\substack{e_1,e_2 \\ e_3,e_4}} [A'_u(e_1,e_2) A'_u(e_1+\eta_1,e_2+\eta_2) A'_w(e_3,e_4) A'_w(e_3+\eta_3,e_4+\eta_4)] \right].$$

The Fourier expansion of this expression (see Fact 5.2) yields the following. From Fact 5.2, we have that $\mathcal{C}'_d := \Lambda_d^{3r} \times \Lambda_d^{3r}$ gives us all the distinct characters of $\mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r}$. Standard computations give us

$$Q(v,u,w) = \sum_{\substack{\alpha_1,\alpha_2 \\ \beta_1,\beta_2 \in \Lambda_d^{3r}}} \underbrace{\widehat{A'_u}(\alpha_1,\alpha_2)^2 \widehat{A'_w}(\beta_1,\beta_2)^2 \mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2 \\ \eta_3,\eta_4}} \left[\prod_{i \in [2]} \chi_{\alpha_i}(\eta_i) \chi_{\beta_i}(\eta_{i+2}) \right]}_{\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)}.$$

As in Lemma 4.3, let $\mathrm{FAR} := \{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in (\Lambda_d^{3r})^4 : \max\{\Delta(\alpha_i, \mathsf{P}_d^{3r}), \Delta(\beta_i, \mathsf{P}_d^{3r})\} \geq 2^{d/2}\}$, $\mathrm{NEAR} := (\Lambda_d^{3r})^4 \setminus \mathrm{FAR}$, $\mathrm{NEAR}_0 := \{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \mathrm{NEAR} : \pi_2(\alpha_1+\alpha_2) =$

$\pi_2'(\beta_1+\beta_2)=0\}$, and $\text{NEAR}_1 := \{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \text{NEAR} : \pi_2(\alpha_1+\alpha_2)=\pi_2'(\beta_1+\beta_2)\neq 0\}$.

Note that the expectation term in $\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)$ is *exactly* the same as that in $\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)$ in Lemma 4.3. This means that the remaining computations can be carried out almost exactly as in Lemma 4.3.

The following can be proved in the same way as Claims 4.4–4.6.

CLAIM 5.4. *For any fixed v,u,w, we have* $\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{FAR}} |\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)| \le 2^{-4\cdot 2^{-d/4}}$.

CLAIM 5.5. $\mathbb{E}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}_1} |\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right] \le 2^{d/2+1}\cdot 2^{-\varepsilon_0 r}$.

(There is a small difference here from the proof of Claim 4.5 owing to the fact that the Fourier coefficients appearing in $\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)$ have a slightly different form. The only change that needs to be made is to sample $\alpha_1,\alpha_2\in\Lambda_d^{3r}$ and $\beta_1,\beta_2\in\Lambda_d^{3r}$ with probability proportional to $\widehat{A'}_u(\alpha_1,\alpha_2)^2$ and $\widehat{A'}_w(\beta_1,\beta_2)^2$, respectively.)

CLAIM 5.6. $\mathbb{E}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}_0} \xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right] \ge \delta^4$.

As in Lemma 4.3, the above can be used to show

$$\begin{aligned}
0 &\ge \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}_0} \xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right.\\
&\qquad\left.+\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}_1} \xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right] - 2^{-4\cdot 2^{-d/4}}\\
&\ge \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}_0} \xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right]\\
&\quad - \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}_1} |\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right] - 2^{-4\cdot 2^{-d/4}}\\
&\ge \delta^4 - 2^{d/2+1}\cdot 2^{-\varepsilon_0 r} - 2^{-4\cdot 2^{-d/4}}.
\end{aligned}$$

This completes the proof of Lemma 5.3. ∎

6. Hardness of coloring 3-colorable 3-uniform hypergraphs. This construction is an adaptation of Khot's construction [Kho02b] to the low-degree long code setting. We prove the theorem by a reduction from 3SAT via the instances of the multilayered label cover problem obtained in Theorem 2.5. Let r,ℓ,η be parameters, which will be determined later, and let $I(\varphi)$ be an instance of the r-repeated ℓ-layered η-smooth label cover instance with constraint graph $G=(V_0,\dots,V_{\ell-1},\{E_{ij}\}_{0\le i<j<\ell})$ obtained from the 3SAT instance φ. We use the results from the preliminaries with the field set to $\mathbb{F}_3=\{0,1,2\}$. For every layer i and every vertex $v\in V_i$, let $\{c_1,\dots,c_{(T+\ell-i)r}\}$ be the clauses corresponding to v, where $T=\lceil l/\eta\rceil$ as in Definition 2.4. We construct polynomials $\{p_1,\dots,p_{(T+\ell-i)r}\}$ of degree at most 6 over $\mathbb{F}_3$ such that p_j depends only on variables in c_j with the following properties. Let $a\in\mathbb{F}_3^3$. If $a\notin\{0,1\}^3$, then $p_j(a)\neq 0$. Otherwise $p_j(a)=0$ iff $c_j(a)=1$. For a degree parameter d, which we will determine later, for each vertex v define the subspace $J_v := \left\{\sum_i q_i p_i : q_i\in \mathsf{P}_{2d-6}^{m_v}\right\}$, where $m_v := m_i = 3(T+\ell-i)r+ir$.

We now define the hypergraph H produced by the reduction. The vertices of H—denoted $V(H)$—are obtained by replacing each $v \in G$ by a block $\mathcal{B}_v$ of $N_v := |\mathsf{P}_{2d}^{m_v}/J_v|$ vertices, which we identify with elements of $\mathsf{P}_{2d}^{m_v}/J_v$. Let N denote $|V(H)| = \sum_v N_v$.

We think of a 3-coloring of $V(H)$ as a map from $V(H)$ to $\mathbb{F}_3$. Given a coloring $A : V(H) \to \mathbb{F}_3$, we denote by $A_v : \mathsf{P}_{2d}^{m_v}/J_v \to \mathbb{F}_3$ the restriction of A to the block $\mathcal{B}_v$. Let $A'_v : \mathsf{P}_{2d}^{m_v} \to \mathbb{F}_3$ denote the lift of A_v as defined in Fact 2.14.

The (weighted) edge set $E(H)$ of H is specified implicitly by the following PCP verifier.

3-Color 3-Uniform Test(d).

1. Choose two layers $0 \le i < j < \ell$ uniformly at random and then choose a uniformly random edge $(u, v) \in E_{ij}$. Let π denote $\pi_{uv} : \mathbb{F}_3^{m_u} \to \mathbb{F}_3^{m_v}$.
2. Choose $p \in \mathsf{P}_d^{m_u}, g \in \mathsf{P}_{2d}^{m_u}$, and $f \in \mathsf{P}_{2d}^{m_v}$ independently and uniformly at random and let $g' := p^2 + 1 - g - f \circ \pi$.
3. Accept iff $A'_v(f), A'_u(g), A'_u(g')$ are not all equal.

The above hypergraph construction explains the reasons (as in [DRS05, Kho02b]) for using the multilayered label cover. Unlike the constructions in the previous two sections, the hyperedges in the 3-uniform case straddle both sides of the corresponding edge (u, v) in the label cover instance. Hence, if constructed from the bipartite label cover, the corresponding 3-uniform hypergraph will also be bipartite and hence always 2-colorable irrespective of the label cover instance. Using the multilayered construction gets around this problem.

LEMMA 6.1 (completeness). *If $\varphi \in \mathsf{3SAT}$, then there is proof $A : V(H) \to \mathbb{F}_3$ that the verifier accepts with probability 1. In other words, the hypergraph H is 3-colorable.*

Proof. Since $\varphi \in \mathsf{3SAT}$, Theorem 2.5 tells us that there are labelings $L_i : V_i \to \{0,1\}^{m_i}$ for $0 \le i < \ell$ that satisfy all the constraints in $I(\varphi)$. For all $i, v \in V_i$, we set $A_v : \mathsf{P}_{2d}^{m_v}/J_v \to \mathbb{F}_3$ such that its lift $A'_v = \mathsf{LC}_{2d}(L_i(v))$. This is possible since A'_v is folded over J_v. For any edge (u, v) between layers i, j, with labels $L_i(u) = a, L_j(v) = b$ such that $\pi(a) = b$, $(A'_v(f), A'_u(g), A'_u(g')) = (f(b), g(a), g'(a))$. The lemma follows by observing that $g'(a) + g(a) + f(b) \neq 0$ always (since $p^2(a) + 1 \neq 0$). ∎

LEMMA 6.2 (soundness). *Let $\ell = 32/\delta^2$. If $\varphi \notin \mathsf{3SAT}$ and H contains an independent set of size $\delta|V(H)|$, then*

$$\delta^5/2^9 \le 2^{-\Omega(r)} \cdot 3^d + \eta \cdot 3^d + \exp(-3^{\Omega(d)}).$$

Proof. Let $A : V(H) \to \{0,1\}$ be the characteristic function of the independent set of fractional size exactly δ. We have that for all v, $\mathbb{E}_{g \in P_{2d}^{m_v}/J_v}[A_v(g)] = \mathbb{E}_{g \in P_{2d}^{m_v}}[A'_v(g)]$, where A'_v is the lift of A_v. Define

$$Q(u, v) := \mathop{\mathbb{E}}_{f,g,p}\left[A'_v(f) A'_u(g) A'_u(p^2 + 1 - f \circ \pi - g)\right].$$

Observe that $\mathbb{E}_{i,j,u,v}[Q(u, v)] = 0$, as A corresponds to an independent set. Using Lemma 2.10, we have the following Fourier expansion of Q:

$$Q(u, v) = \sum_{\alpha,\beta,\gamma} \widehat{A}'_v(\alpha) \widehat{A}'_u(\beta) \widehat{A}'_u(\gamma) \mathop{\mathbb{E}}_{f,g,p}[\chi_\alpha(f) \chi_\beta(g) \chi_\gamma(g')], \tag{6.1}$$

where the summation is over $\alpha \in \Lambda_{2d}^{m_v}$, $\beta, \gamma \in \Lambda_{2d}^{m_u}$ and Λ is as defined in Lemma 2.10. From the orthonormality of characters, the nonzero terms satisfy $\beta = \gamma$ and $\alpha =$

$\pi_3(\beta)$. Substituting in (6.1), we get

$$Q(u,v) = \sum_{\beta} \underbrace{\widehat{A'_u}(\beta)^2 \widehat{A'_v}(\pi_3(\beta)) \mathop{\mathbb{E}}_{p} \left[\chi_\beta(p^2+1)\right]}_{\xi_{u,v}(\beta)}. \tag{6.2}$$

CLAIM 6.3. *If we set $\ell = 32/\delta^2$, then there exist layers $0 \le i < j < \ell$ such that $\mathbb{E}_{(u,v)\in E_{ij}}[\xi_{u,v}(0)] \ge \delta^5/2^9$.*

Proof. Since A' has fractional size δ, there exists a set S of vertices of fractional size $\delta/2$ such that for all $v \in S$, $\widehat{A'_v}(0) = \mathbb{E}_f[A'_v(f)] \ge \delta/2$. Furthermore, there exist $\delta\ell/4$ layers, in which the fractional size of $S_i := S \cap V_i$ in layer V_i is at least $\delta/4$. Since $\ell = 32/\delta^2$, we obtain from Theorem 2.5 that there exist layers i, j such that the fraction of edges in E_{ij} between S_i and S_j is at least $\delta' = \delta^2/64$. From above, we have that

$$\mathop{\mathbb{E}}_{(u,v)\in E_{ij}}[\xi_{u,v}(0)] \ge \delta' \cdot (\delta/2)^3 \ge \delta^5/2^9.$$

For the rest of the proof, layers i, j will be fixed as given by Claim 6.3. To analyze the expression in (6.2), we consider the following breakup of $\Lambda_{2d}^{m_i} \setminus \{0\}$ for every $(u,v) \in E_{ij}$:

$$\begin{aligned}
\text{FAR} &:= \{\beta \in \Lambda_{2d}^{m_i} : \Delta(\beta, (\mathsf{P}_{2d}^{m_i})^\perp) \ge 3^{d/2}\}, \\
\text{NEAR}_1 &:= \{\beta \in \Lambda_{2d}^{m_i} \setminus \text{FAR} : \beta \ne 0 \text{ and } \pi_3(\beta) \notin (\mathsf{P}_{2d}^{m_i})^\perp\}, \\
\text{NEAR}_0 &:= \{\beta \in \Lambda_{2d}^{m_i} \setminus \text{FAR} : \beta \ne 0 \text{ and } \pi_3(\beta) \in (\mathsf{P}_{2d}^{m_i})^\perp\}.
\end{aligned}$$

In Claims 6.4, 6.5, and 6.6, we bound the absolute values of the sum of $\mathbb{E}_{u,v}[\xi_{u,v}(\beta)]$ for β in FAR, NEAR_0, and NEAR_1, respectively.

CLAIM 6.4. $|\mathbb{E}_{(u,v)\in E_{ij}}[\sum_{\beta\in\text{FAR}} \xi_{u,v}(\beta)]| \le \exp(-3^{\Omega(d)})$.

CLAIM 6.5. $|\mathbb{E}_{(u,v)\in E_{ij}}[\sum_{\beta\in\text{NEAR}_1} \xi_{u,v}(\beta)]| \le 2^{-\Omega(r)} \cdot 3^d$.

CLAIM 6.6. $|\mathbb{E}_{(u,v)\in E_{ij}}[\sum_{\beta\in\text{NEAR}_0} \xi_{u,v}(\beta)]| \le \eta \cdot 3^d$.

Combined with Claim 6.3, this exhausts all of the terms in the expansion (6.2). Lemma 6.2 now follows from Claims 6.3–6.6. ∎

We now proceed to the proofs of Claims 6.4–6.6.

Proof of Claim 6.4.

$$\left|\mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in\text{FAR}} \xi_{u,v}(\beta)\right]\right| \le \mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in\text{FAR}} |\widehat{A'_u}(\beta)|^2 \cdot |\widehat{A'_v}(\pi_3(\beta))| \cdot \left|\mathop{\mathbb{E}}_{p}\left[\omega^{\langle\beta,p^2+1\rangle}\right]\right|\right].$$

The quantity $\langle\beta, p^2\rangle$ is analyzed in section 3. Let z be a uniformly random $\mathbb{F}_3$ element. By Lemmas 3.1 and 3.4, we get that the statistical distance between the distributions of $\langle\beta, p^2+1\rangle$ and z is $\exp(-3^{\Omega(d)})$. Since the $\mathbb{E}_z[\omega^z] = 0$, we have that $|\mathbb{E}_p[\omega^{\langle\beta,p^2+1\rangle}]| \le \exp(-3^{\Omega(d)})$. The claim follows since $|\widehat{A'_v}(\alpha)| \le 1$ for any α and $\sum_\beta |\widehat{A'_u}(\beta)|^2 \le 1$. ∎

Proof of Claim 6.5. It suffices to bound the following for proving the claim:

$$\mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in \mathrm{NEAR}_1} |\widehat{A'_u}(\beta)|^2 \cdot |\widehat{A'_v}(\pi_3(\beta))|\right]$$

$$\leq \mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sqrt{\sum_{\beta\in \mathrm{NEAR}_1} |\widehat{A'_u}(\beta)|^2 \cdot |\widehat{A'_v}(\pi_3(\beta))|^2}\sqrt{\sum_{\beta\in \mathrm{NEAR}_1} |\widehat{A'_u}(\beta)|^2}\right] \text{ (Cauchy–Schwarz)}$$

$$\leq \sqrt{\mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in \mathrm{NEAR}_1} |\widehat{A'_u}(\beta)|^2 \cdot |\widehat{A'_v}(\pi_3(\beta))|^2\right]} \text{ (Jensen's inequality).}$$ ∎

We bound the above using a Fourier decoding argument as in the proof of Claim 4.5. For every vertex $v \in V_i \cup V_j$, pick a random β according to $|\widehat{A'_v}(\beta)|^2$ (note $\sum_\beta |\widehat{A'_v}(\beta)|^2 \leq 1$) and assign a random labeling to v from the support of β. By an argument identical to the proof of Claim 4.5, we get (using the soundness of the multilayered label cover from Theorem 2.5)

$$\frac{1}{3^d} \mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in \mathrm{NEAR}_1} |\widehat{A'_v}(\pi_3(\beta))|^2 |\widehat{A'_u}(\beta)|^2\right] \leq 2^{-\Omega(r)}.$$

Proof of Claim 6.6. We bound this sum using the smoothness property of the label cover instance:

$$\mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in \mathrm{NEAR}_0} |\widehat{A'_u}(\beta)|^2 \cdot |\widehat{A'_v}(\pi_3(\beta))|\right]$$

$$\leq \mathop{\mathbb{E}}_{u\in V_i}\left[\sum_{\beta\notin \mathrm{FAR}\cup\{0\}} \Pr_{v:(u,v)\in E_{ij}}\left[\pi_3(\beta) \in (\mathsf{P}_{2d}^{m_v})^\perp\right] \cdot |\widehat{A'_u}(\beta)|^2\right].$$

We now argue that for every u and $\beta \notin \mathrm{FAR} \cup \{0\}$,

$$\Pr_{(u,v)\in E_{ij}}\left[\pi_3(\beta) \notin (\mathsf{P}_{2d}^{m_v})^\perp\right] \leq 3^d \cdot \eta.$$

This combined with the fact that $\sum_\beta |\widehat{A'_u}(\beta)|^2 \leq 1$ yields the claim. For every $u \in V_i$ and β such that $0 \neq |\operatorname{support}(\beta)| = \Delta(\beta, (\mathsf{P}_{2d}^{m_u})^\perp) \leq 3^{d/2}$, by the smoothness property (Theorem 2.5), we have that with probability at least $1 - 3^d\eta$, we have

$$\forall a \neq a' \in \operatorname{support}(\beta), \quad \pi(a) \neq \pi(a'). \tag{6.3}$$

When (6.3) holds, we have $\pi_3(\beta) \neq 0$. Now since

$$|\operatorname{support}(\pi_3(\beta))| \leq |\operatorname{support}(\beta)| \leq 3^{d/2}$$

and nonzero polynomials in $(\mathsf{P}_{2d}^{m_v})^\perp$ have support at least 3^d, we can further conclude that $\pi_3(\beta) \notin (\mathsf{P}_{2d}^{m_v})^\perp$ whenever (6.3) holds. ∎

Proof of Theorem 1.3. Given the completeness (see Lemma 6.1) and soundness (see Lemma 6.2), we only need to fix parameters. Let n be the size of the 3SAT instance and N the size of the hypergraph produced by the reduction.

Let $d = C_1 \log\log(1/\delta')$, $\eta = (\delta')^5/C_2$, and $r = C_3 \log(1/\delta')$ for large enough constants C_1, C_2, C_3 and parameter $\delta' \in (0,1)$ to be determined shortly. By Lemma 6.2, if H has an independent set of size δN, then $\delta^5/2^9 \le 3^d \cdot 2^{-\Omega(r)} + 3^d \cdot \eta + \exp(-3^{\Omega(d)}) < (\delta')^5/2^9$ for large enough C_1, C_2, C_3. Hence, H has no independent sets of $\delta' N$.

The hypergraph H is of size $N = \ell n^{(1+1/\eta)\ell r} 3^{((1+1/\eta)\ell r)^{O(d)}}$. Setting $\ell = C_4/(\delta')^2$ gives $\log(1/\delta') = \Theta(\log\log n/\log\log\log n)$, and since $\log\log n = \Theta(\log\log N)$, we get that

$$N = n^{2^{O(\log\log n/\log\log\log n)}} \quad \text{and} \quad 1/\delta' = 2^{\Theta(\log\log N/\log\log\log N)}.$$

This completes the proof of Theorem 1.3. ∎

Appendix A. Proof of Claim 3.7. We need the following theorem due to Haramaty, Shpilka, and Sudan [HSS13].

THEOREM A.1 ([HSS13, Theorems 4.16 and 1.7] specialized to $\mathbb{F}_3$ and using absolute distances instead of fractional distances). *There exists a constant λ_3 such that the following holds. For $\beta : \mathbb{F}_3^n \to \mathbb{F}_3$, let $A_1, \ldots, A_K$ be hyperplanes such that $\beta|_{A_i}$ is Δ_1-close to some degree r polynomial on A_i. If $K > 3^{\lceil \frac{r+1}{2} \rceil + \lambda_3}$ and $\Delta_1 < 3^{n-r/2-2}/2$, then $\Delta(\beta, \mathsf{P}_r^n) \le 6\Delta_1 + 8 \cdot 3^n/K$.*

Setting the degree $r = 2n - 2d - 1$ in the above theorem implies that if there are $K > 3^{n-d+\lambda_3}$ hyperplanes $A_1, \ldots, A_K$ such that $\beta|_{A_i}$ is Δ_1-close to a degree $(2n-2d-1)$ polynomial on A_i, then $\Delta(\beta, \mathsf{P}_{2n-2d-1}^n) \le 6\Delta_1 + 8 \cdot 3^n/K$.

Suppose Claim 3.7 were false. Then, for every nonzero $l \in \mathsf{P}_1^n$, at least one of $\beta|_{\ell=0}$ or $\beta|_{\ell=1}$ or $\beta|_{\ell=2}$ is $\Delta/27$-close to a degree $(2n-2d-1)$ polynomial. We thus get $K = (3^n - 1)/2$ hyperplanes such that the restriction of β to these hyperplanes is $\Delta/27$-close to a degree $(2n-2d-1)$ polynomial. Observe that $K \ge 3^{n-d+\lambda_3}$ if $d \ge d_0 \ge \lambda_3 + 2$ and $\Delta/27 < 3^{n-(2n-2d-1)/2-2}/2 = 3^{d-1.5}/2$ if $\Delta < 3^d$. Hence, by Theorem A.1 we have $\Delta(\beta, \mathsf{P}_{2n-2d-1}^n) \le 6\Delta/27 + 2 \cdot 8 \cdot 3^n/(3^n - 1) < 6\Delta/27 + 32 < \Delta$ (since $\Delta \ge 3^4$). This contradicts the hypothesis that β is Δ-far from $\mathsf{P}_{2n-2d-1}^n$. ∎

REFERENCES

[AC06] S. ARORA AND E. CHLAMTAC, *New approximation guarantee for chromatic number*, in Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 215–224, https://doi.org/10.1145/1132516.1132548.

[AKMR96] N. ALON, P. KELSEN, S. MAHAJAN, AND H. RAMESH, *Approximate hypergraph coloring*, Nordic J. Comput., 3 (1996), pp. 425–439 (preliminary version in 5th SWAT, 1996); available online at https://www.cs.helsinki.fi/njc/References/alonkmr1996:425.html.

[BG16] J. BRAKENSIEK AND V. GURUSWAMI, *New hardness results for graph and hypergraph colorings*, in Proceedings of the 31st Conference on Computational Complexity, 2016, pp. 14:1–14:27, https://doi.org/10.4230/LIPIcs.CCC.2016.14.

[BGH+15] B. BARAK, P. GOPALAN, J. HÅSTAD, R. MEKA, P. RAGHAVENDRA, AND D. STEURER, *Making the long code shorter*, SIAM J. Comput., 44 (2015), pp. 1287–1324 (preliminary version in 53rd FOCS, 2012), https://doi.org/10.1137/130929394.

[BGS98] M. BELLARE, O. GOLDREICH, AND M. SUDAN, *Free bits, PCPs, and nonapproximability—towards tight results*, SIAM J. Comput., 27 (1998), pp. 804–915 (preliminary version in 36th FOCS, 1995), https://doi.org/10.1137/S0097539796302531.

[BK97] A. BLUM AND D. R. KARGER, *An $\tilde{O}(n^{3/14})$-coloring algorithm for 3-colorable graphs*, Inform. Process. Lett., 61 (1997), pp. 49–53, https://doi.org/10.1016/S0020-0190(96)00190-1.

[BKS+10] A. Bhattacharyya, S. Kopparty, G. Schoenebeck, M. Sudan, and D. Zuckerman, *Optimal testing of Reed-Muller codes*, in Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2010, pp. 488–497, https://doi.org/10.1109/FOCS.2010.54.

[Blu94] A. Blum, *New approximation algorithms for graph coloring*, J. ACM, 41 (1994), pp. 470–516 (preliminary versions in 31st FOCS, 1990, and 21st STOC, 1989), https://doi.org/10.1145/176584.176586.

[CF96] H. Chen and A. M. Frieze, *Coloring bipartite hypergraphs*, in Integer Programming and Combinatorial Optimization, Lecture Notes in Comput. Sci. 1084, W. H. Cunningham, S. T. McCormick, and M. Queyranne, eds., Springer, Berlin, Heidelberg, 1996, pp. 345–358, https://doi.org/10.1007/3-540-61310-2_26.

[Chl07] E. Chlamtac, *Approximation algorithms using hierarchies of semidefinite programming relaxations*, in Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2007, pp. 691–701, https://doi.org/10.1109/FOCS.2007.72.

[DG15] I. Dinur and V. Guruswami, *PCPs via the low-degree long code and hardness for constrained hypergraph coloring*, Israel J. Math., 209 (2015), pp. 611–649 (preliminary version in 54th FOCS, 2013), https://doi.org/10.1007/s11856-015-1231-3.

[DGKR05] I. Dinur, V. Guruswami, S. Khot, and O. Regev, *A new multilayered PCP and the hardness of hypergraph vertex cover*, SIAM J. Comput., 34 (2005), pp. 1129–1146 (preliminary version in 35th STOC, 2003), https://doi.org/10.1137/S0097539704443057.

[DK13] I. Dinur and G. Kol, *Covering CSPs*, in Proceedings of the 28th IEEE Conference on Computational Complexity, 2013, pp. 207–218, https://doi.org/10.1109/CCC.2013.29.

[DMR09] I. Dinur, E. Mossel, and O. Regev, *Conditional hardness for approximate coloring*, SIAM J. Comput., 39 (2009), pp. 843–873 (preliminary version in 38th STOC, 2006), https://doi.org/10.1137/07068062X.

[DRS05] I. Dinur, O. Regev, and C. D. Smyth, *The hardness of* 3*-uniform hypergraph coloring*, Combinatorica, 25 (2005), pp. 519–535 (preliminary version in 43rd FOCS, 2002), https://doi.org/10.1007/s00493-005-0032-4.

[Fei98] U. Feige, *A threshold of* $\ln n$ *for approximating set cover*, J. ACM, 45 (1998), pp. 634–652 (preliminary version in 28th STOC, 1996), https://doi.org/10.1145/285055.285059.

[GHH+14] V. Guruswami, P. Harsha, J. Håstad, S. Srinivasan, and G. Varma, *Super-polylogarithmic hypergraph coloring hardness via low-degree long codes*, in Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC), 2014, pp. 614–623, https://doi.org/10.1145/2591796.2591882.

[GHS02] V. Guruswami, J. Håstad, and M. Sudan, *Hardness of approximate hypergraph coloring*, SIAM J. Comput., 31 (2002), pp. 1663–1686 (preliminary version in 41st FOCS, 2000), https://doi.org/10.1137/S0097539700377165.

[GK04] V. Guruswami and S. Khanna, *On the hardness of* 4*-coloring a* 3*-colorable graph*, SIAM J. Discrete Math., 18 (2004), pp. 30–40 (preliminary version in 15th CCC, 2000), https://doi.org/10.1137/S0895480100376794.

[HSS13] E. Haramaty, A. Shpilka, and M. Sudan, *Optimal testing of multivariate polynomials over small prime fields*, SIAM J. Comput., 42 (2013), pp. 536–562 (preliminary version in 52nd FOCS, 2011), https://doi.org/10.1137/120879257.

[Hua13] S. Huang, *Improved hardness of approximating chromatic number*, in Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, Lecture Notes in Comput. Sci. 8096, P. Raghavendra, S. Raskhodnikova, K. Jansen, and J. D. P. Rolim, eds., Springer, Berlin, Heidelberg, 2013, pp. 233–243, https://doi.org/10.1007/978-3-642-40328-6_17.

[Hua15] S. Huang, $2^{(\log N)^{1/10-o(1)}}$ *Hardness for Hypergraph Coloring*, preprint, https://arxiv.org/abs/1504.03923, 2015.

[Kho01] S. Khot, *Improved inapproximability results for MaxClique, chromatic number and approximate graph coloring*, in Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS), 2001, pp. 600–609, https://doi.org/10.1109/SFCS.2001.959936.

[Kho02a] S. Khot, *Hardness results for approximate hypergraph coloring*, in Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing (STOC), 2002, pp. 351–359, https://doi.org/10.1145/509907.509962.

[Kho02b] S. Khot, *Hardness results for coloring* 3*-colorable* 3*-uniform hypergraphs*, in Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science (FOCS), 2002, pp. 23–32, https://doi.org/10.1109/SFCS.2002.1181879.

[KKMO07] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell, *Optimal inapproximability results for MAX-CUT and other* 2*-variable CSPs?*, SIAM J. Comput., 37 (2007), pp. 319–357 (preliminary version in 45th FOCS, 2004), https://doi.org/10.1137/S0097539705447372.

[KLS00] S. Khanna, N. Linial, and S. Safra, *On the hardness of approximating the chromatic number*, Combinatorica, 20 (2000), pp. 393–415 (preliminary version in 2nd Israel Symposium on Theory of Computing Systems, 1993), https://doi.org/10.1007/s004930070013.

[KM13] D. M. Kane and R. Meka, *A PRG for Lipschitz functions of polynomials with applications to sparsest cut*, in Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing (STOC), 2013, pp. 1–10, https://doi.org/10.1145/2488608.2488610.

[KMS98] D. R. Karger, R. Motwani, and M. Sudan, *Approximate graph coloring by semidefinite programming*, J. ACM, 45 (1998), 246–265 (preliminary version in 35th FOCS, 1994), https://doi.org/10.1145/274787.274791.

[KNS01] M. Krivelevich, R. Nathaniel, and B. Sudakov, *Approximating coloring and maximum independent sets in* 3*-uniform hypergraphs*, J. Algorithms, 41 (2001), pp. 99–113 (preliminary version in 12th SODA, 2001), https://doi.org/10.1006/jagm.2001.1173.

[KS14] S. Khot and R. Saket, *Hardness of coloring* 2*-colorable* 12*-uniform hypergraphs with* $\exp(\log^{\Omega(1)} n)$ *colors*, in Proceedings of the IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS), 2014, pp. 206–215, https://doi.org/10.1109/FOCS.2014.30.

[KT12] K. Kawarabayashi and M. Thorup, *Combinatorial coloring of* 3*-colorable graphs*, in Proceedings of the IEEE 53rd Annual Symposium on Foundations of Computer Science (FOCS), 2012, pp. 68–75, https://doi.org/10.1109/FOCS.2012.16.

[KT14] K. Kawarabayashi and M. Thorup, *Coloring* 3*-colorable graphs with* $o(n^{1/5})$ *colors*, in Proceedings of the 31st International Symposium on Theoretical Aspects of Computer Science (STACS), LIPIcs 25, E. W. Mayr and N. Portier, eds., Schloss Dagstuhl, Dagstuhl, Germany, 2014, pp. 458–469, https://doi.org/10.4230/LIPIcs.STACS.2014.458.

[LN97] R. Lidl and H. Niederreiter, *Finite Fields*, 2nd ed., Encyclopedia Math. Appl. 20, Cambridge University Press, Cambridge, UK, 1997, https://doi.org/10.1017/CBO9780511525926.

[Raz98] R. Raz, *A parallel repetition theorem*, SIAM J. Comput., 27 (1998), pp. 763–803 (preliminary version in 27th STOC, 1995), https://doi.org/10.1137/S0097539795280895.

[Sak14] R. Saket, *Hardness of finding independent sets in* 2*-colorable hypergraphs and of satisfiable CSPs*, in Proceedings of the 29th IEEE Conference on Computational Complexity, 2014, pp. 78–89, https://doi.org/10.1109/CCC.2014.16.

[Var15] G. Varma, *Reducing uniformity in Khot-Saket hypergraph coloring hardness reductions*, Chic. J. Theoret. Comput. Sci., 2015 (2015), 3, https://doi.org/10.4086/cjtcs.2015.003.

[Wig83] A. Wigderson, *Improving the performance guarantee for approximate graph coloring*, J. ACM, 30 (1983), pp. 729–735 (preliminary version in 14th STOC, 1982), https://doi.org/10.1145/2157.2158.