Electronic Colloquium on Computational Complexity, Report No. 167 (2013)

Super-polylogarithmic hypergraph coloring hardness via low-degree long codes

Venkatesan Guruswami* Prahladh Harsha† Johan Håstad‡
Srikanth Srinivasan§ Girish Varma¶

November 28, 2013

Abstract

We prove improved inapproximability results for hypergraph coloring using the low-degree polynomial code (aka, the "short code" of Barak *et. al.* [FOCS 2012]) and the techniques proposed by Dinur and Guruswami [FOCS 2013] 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.
- Coloring a 3-colorable 3-uniform hypergraph with $(\log n)^{\Omega(1/\log\log\log n)}$ colors.

In each of these cases, the hardness results obtained are (at least) exponentially stronger than what was previously known for the respective cases. 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.

The fundamental bottleneck in obtaining coloring inapproximability results using the low-degree long code was a multipartite structural restriction in the PCP construction of Dinur-Guruswami. We are able to get around this restriction by simulating the multipartite structure implicitly by querying just one partition (albeit requiring 8 queries), which yields our result for 2-colorable 8-uniform hypergraphs. The result for 4-colorable 4-uniform hypergraphs is obtained via a "query doubling" method exploiting additional properties of the 8-query test. For 3-colorable 3-uniform hypergraphs, we exploit the ternary domain to design a test with an *additive* (as opposed to multiplicative) noise function, and analyze its efficacy in killing high weight Fourier coefficients via the pseudorandom properties of an associated quadratic form. The latter step involves extending the key algebraic ingredient of Dinur-Guruswami concerning testing binary Reed-Muller codes to the ternary alphabet.

*Computer Science Department, Carnegie Mellon University, USA. Research supported in part by a Packard Fellowship, US-Israel BSF grant number 2008293, and the US National Science Foundation Grant No. CCF-1115525. Email: guruswami@cmu.edu.

†Tata Institute of Fundamental Research, India. Part of the work was done while the author was visiting the Simons Institute for Theory of Computing. Email: prahladh@tifr.res.in.

‡KTH Royal Institute of Technology, Sweden. Part of the work was done while the author was visiting the Simons Institute for Theory of Computing. Partly supported by ERC grant 226203. Email: johanh@kth.se.

§Department of Mathematics, IIT Bombay, India. Email: srikanth@math.iitb.ac.in.

¶Tata Institute of Fundamental Research, India. Supported by Google India under the Google India PhD Fellowship Award. Email: girishrv@tifr.res.in.

ISSN 1433-8092

1 Introduction

The last two decades have seen tremendous progress in understanding the hardness of approximating constraint satisfaction problems. Despite this progress, the status of approximate coloring of constant colorable (hyper)graphs is not resolved and in fact, there is an exponential (if not doubly exponential) gap between the best known approximation algorithms and inapproximability results. The current best known approximation algorithms require at least $n^{\Omega(1)}$ colors to color a constant colorable (hyper)graph on n vertices while the best inapproximability results only rule out at best $(\log n)^{O(1)}$ (and in fact, in most cases, only $o(\log n)$) colors.

Given this disparity between the positive and negative results, it is natural to ask why current inapproximability techniques get stuck at the $\operatorname{poly}\log n$ color barrier. The primary bottleneck in going past polylogarithmic colors is the use of the *long code*, a quintessential ingredient in almost all tight inapproximability results, since it was first introduced by Bellare, Goldreich and Sudan [2]. The long code, as the name suggests, is the most redundant encoding, wherein a n-bit Boolean string x is encoded by a 2^{2^n}-bit string which 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 which prevents the coloring inapproximability to go past the $\operatorname{poly}\log n$ barrier. Recently, Barak *et. al.* [1], while trying to understanding the tightness of the Arora-Barak-Steurer algorithm for unique games, introduced the *short code*, also called the *low-degree long code* [4]. The low-degree long code is a puncturing of the long code in the sense, that it contains only the evaluations of low-degree functions (opposed to all functions). Barak *et. al.* [1] introduced the low-degree long code to prove exponentially stronger integrality gaps for Unique Games, and construct small set expanders whose Laplacians have many small eigenvalues,

Being a derandomization 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.* [1] used it obtain a more efficient version of the KKMO alphabet reduction [12] for Unique Games. However, using the low-degree long code towards improved reductions from Label Cover posed some challenges related to folding, and incorporating noise without giving up perfect completeness (which is crucial for results on coloring). Recently, Dinur and Guruswami [4] introduced a very elegant set of techniques to adapt the long code based inapproximability results to low-degree long codes. Using these techniques, they proved (1) improved inapproximability results for gap-$(1, \frac{15}{16} + \varepsilon)$-4SAT for $\varepsilon = \exp(-2^{\Omega(\sqrt{\log\log N})})$ (long code based reductions show for $\varepsilon = 1/\operatorname{poly}\log N$) and (2) hardness for a variant of approximate hypergraph coloring, with a gap of 2 and $\exp(2^{\Omega(\sqrt{\log\log N})})$ number of colors (where N is the number of vertices). It is to be noted that the latter is the first result to go beyond the logarithmic barrier for a coloring-type problem. However, the Dinur-Guruswami [4] results do not extend to standard (hyper)graph coloring hardness due to a multipartite structural bottleneck in the PCP construction, which we elaborate below.

As mentioned earlier, the two main contributions of Dinur-Guruswami [4] are (1) folding mechanism over the low-degree long code and (2) noise in the low-degree polynomials. The results of Bhattacharyya *et. al.* [3] and Barak *et. al.* [1] suggest that the product of d linearly independent affine functions suffices to work as noise for the low-degree long code setting (with degree = d) in the sense that it attenuates the contribution of large weight Fourier coefficients. However, this works only for PCP tests with imperfect completeness. Since approximate coloring results require perfect completeness, Dinur and Guruswami [4] inspired by the above result, develop a noise function which is the product of two random low-degree polynomials such that

the sum of the degrees is at most d. This necessitates restricting certain functions in the PCP test to be of smaller degree which in turn requires the PCP tests to query two types of tables – one a low-degree long code of degree d and another a low-degree long code of smaller degree. Though the latter table is a part of the former, a separate table is needed since otherwise the queries will be biased to the small degree portion of the low-degree long code. This multipartite structure is what precludes them from extending their result for standard coloring results. (Clearly, if the query of the PCP tests straddles two tables, then the associated hypergraph is trivially 2-colorable.)

1.1 Hypergraph coloring results

In this work, we show how this multipartite structural restriction can be overcome, thus yielding (standard) coloring inapproximability results. The first of our results extends the result of Dinur-Guruswami [4]: variant of 6-uniform hypergraph coloring result to a standard hypergraph coloring result, albeit of larger uniformity, namely 8.

Theorem 1.1 (2-colorable 8-uniform hypergraphs)**.** *Assuming* $\mathsf{NP} \not\subseteq \mathsf{DTIME}(n^{2^{O(\sqrt{\log \log n})}})$*, there is no polynomial time algorithm which, 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* $N/2^{2^{O(\sqrt{\log \log N})}}$.

This result is obtained using the framework of Dinur-Guruswami [4] by showing that the two additional queries can be used to simulate queries into the smaller table via queries into the larger table.

We note that prior to this result, $(\log N)^{\Omega(1)}$ colors was the strongest quantitative bound on hardness for hypergraph coloring: Khot obtained such a result for coloring 7-colorable 4-uniform hypergraphs [10] while Dinur and Guruswami [4] obtained a similar (but incomparable) result for 2-colorable 6-uniform hypergraphs both using the long code.

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 $e_1, e_2, e_1', e_2', e_3, e_4, e_3', e_4'$ in the completeness case satisfy

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

which is stronger than the required

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

Furthermore, for each i, the queries e_i and e_i' 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_1') and is expected to return 2 bits which 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,

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

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.

Theorem 1.2 (4-colorable 4-uniform hypergraphs)**.** *Assuming* NP $\not\subseteq$ DTIME$(n^{2^{O(\sqrt{\log\log n})}})$*, there is no polynomial time algorithm which, 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 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 [10] 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[1], (the proof of) Theorem 1.2 demonstrates a Boolean 4CSP for which it is quasi-NP-hard to distinguish between covering number of 2 vs. $\exp(\sqrt{\log\log N})$. The previous best result for a Boolean 4CSP was 2 vs. $\log\log N$, due to Dinur and Kol [6].

We then ask if we can prove coloring inapproximability for even smaller uniformity, i.e., 2 and 3 (graphs and 3-uniform hypergraphs respectively). We show that we can use a different noise function over $\mathbb{F}_3$ to obtain the following inapproximability result for 3-colorable 3-uniform hypergraphs.

Theorem 1.3 (3-colorable 3-uniform hypergraphs)**.** *Assuming* NP $\notin$ DTIME$(n^{2^{O(\log\log n/\log\log\log n)}})$*, there is no polynomial time algorithm which, 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* $N/2^{O(\log\log N/\log\log\log N)}$.

Prior to this result, the best inapproximability result for O(1)-colorable 3-uniform hypergraphs were as follows: Khot [11] showed that it is quasi-NP-hard to color a 3-colorable 3-uniform hypergraphs with $(\log\log N)^{1/9}$ colors and Dinur, Regev and Smyth [7] showed that it is quasi-NP-hard to color a 2-colorable 3-uniform hypergraphs with $(\log\log N)^{1/3}$ colors (observe that $2^{O(\log\log N/\log\log\log N)}$ is exponentially larger than $(\log\log N)^{\Omega(1)}$). For 2-colorable 3-uniform hypergraphs, the result of Dinur *et. al.* [7] only rules out colorability by $(\log\log N)^{\Omega(1)}$, while a recent result due to Khot and Saket [13] shows that it is hard to find a δN-sized independent set in a given N-vertex 2-colorable 3-uniform hypergraph assuming the d-to-1 games conjecture. Our improved inapproximability result is obtained by adapting Khot's proof to the low-degree long code using the new noise function over $\mathbb{F}_3$. We remark that this result is not as strong as the previous two ($2^{O(\log\log N/\log\log\log N)}$ instead of $2^{2^{O(\sqrt{\log\log N})}}$) as for 3-uniform hypergraphs, the starting point is a multilayered smooth label cover instance instead of just label cover, which causes a blowup in size and a corresponding deterioration in the parameters.

[1]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

1.2 Low-degree long code analysis via Reed-Muller testing

One of the key contributions of Barak *et. al.* [1] was the discovery of a connection between Reed-Muller testing and the analysis of the low-degree long code. In particular, they showed the following. Let P^n_d set of degree d polynomials on n variables over $\mathbb{F}_2$. For functions $\beta, g : \mathbb{F}_2^n \to \mathbb{F}_2$, let $\chi_\beta(g) = (-1)^{\sum_{x\in\mathbb{F}_2^n} \beta(x)g(x)}$. Barak *et. al.* oberved that if β is far from the set P^n_{n-d-1} of degree $n-d-1$ polynomials, then one can bound the expectation $|\mathbb{E}_\mu[\chi_\beta(\eta)]|$ for a random low-weight η using a powerful result on Reed-Muller testing over $\mathbb{F}_2$ due to Bhattacharyya *et. al.* [3]. This demonstrates that the noise function η attenuates the contribution of high-order Fourier coefficients and is thus useful in the low-degree long code analysis. However, this noise η has imperfect completeness and Dinur-Guruswami had to prove a new result on Reed-Muller testing over $\mathbb{F}_2$ to construct a noise function that allows for perfect completeness. They showed that if β is $2^{d/2}$-far from P^n_{n-d-1}, then $\mathbb{E}_{g\in\mathsf{P}^n_{d/4}}\left|\mathbb{E}_{h\in\mathsf{P}^n_{3d/4}}[\chi_\gamma(gh)]\right|$ was doubly exponentially small in d (see Theorem 2.12 for a fomal statement). This allowed them to extend some of the long code based inapproximability with perfect completeness to the low-degree long code setting. Tests based on the above property need to access functions of different degree (e.g., g, gh in the above discussion) and this results in a multipartite structure in the low-degree long code tables of [4]. The results for 2-colorable 8-uniform hypergraphs and 4-uniform 4-colorable hypergraphs are obtained using the above result of [4].

For the case of 3-uniform 3-colorable hypergraphs, we observe that if we extend the alphabet to ternary (i.e., $\mathbb{F}_3$ instead of $\mathbb{F}_2$), we can design a noise function that has both perfect completeness and does not result in a multipartite structural restriction. Let P^n_d now denote the set of degree d polynomials on n variables over $\mathbb{F}_3$. We show that if $\beta : \mathbb{F}_3^n \to \mathbb{F}_3$ is $3^{d/2}$-far from $\mathsf{P}^n_{2n-2d-1}$, then $\left|\mathbb{E}_{p\in\mathsf{P}^n_d}[\chi_\beta(p^2)]\right|$ is doubly exponentially small in d. This is proved by showing the following pseudorandom property of 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. If the distance of β from polynomials of degree $2n-2d-1$, denoted by $\Delta_d(\beta)$ is at least $3^{d/2}$, then the rank of the matrix $Q(\beta)$ is exponential in d and is otherwise equal to the distance $\Delta_d(\beta)$. This rank bound is proved along the lines of [4] using the Reed-Muller tester analysis of Haramaty, Shpilka and Sudan [9] over general fields instead of the Bhattacharyya *et. al.* [3] analysis over $\mathbb{F}_2$.

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 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:*

- *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 [15] with r repetitions on hard instances of MAX-3SAT where each variable occurs the same number of times [8].

Theorem 2.2. *There is an algorithm which 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$(I(\varphi)) \leq 2^{-\varepsilon_0 r}$ for some universal constant $\varepsilon_0 \in (0, 1)$.*

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

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.

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 $a, b \in \Sigma_U$*

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

where v is a random neighbour of u.

Definition 2.4 (r-repeated ℓ-layered η-smooth label cover)**.** *Let $T := \lceil \ell/\eta \rceil$ and 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:*

- *An ℓ-partite graph with vertex sets $V_0, \cdots V_{\ell-1}$. 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 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* $\pi_{v_i v_j}$ *is the projection of assignments for clauses and variables in* v_i *to that 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 only depends 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* $\pi_{v_i v_j}(L_i(v_i)) = L_j(v_j)$*. 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.* [5] in the context of hypergraph vertex cover inapproximability (also see [7]).

Theorem 2.5. *There is an algorithm which on input a* 3SAT *instance* φ *and* $\ell, r \in \mathbb{N}, \eta \in [0,1)$ *outputs a* 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. $\forall\, 0 \le i < j < \ell$*, the bipartite label cover instance on* $I_{ij} = \big((V_i, V_j, E_{ij}), \Sigma_{V_i}, \Sigma_{V_j}, \Pi_{ij}\big)$ *is* η*-smooth.*
2. *For* $1 < m < \ell$*, any* m *layers* $0 \le i_1 < \cdots < i_m \le \ell - 1$*, any* $S_{i_j} \subseteq V_{i_j}$ *such that* $|S_{i_j}| \ge \frac{2}{m}|V_{i_j}|$*, there exists distinct* i_j *and* $i_{j'}$ *such that the fraction of edges between* S_{i_j} *and* $S_{i_{j'}}$ *relative to* $E_{i_j i_{j'}}$ *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*
$$OPT_{i,j}(I(\varphi)) \le 2^{-\Omega(r)}, \quad \forall 0 \le i < j \le \ell.$$

2.2 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 context and hence the dependence of p on the quantities defined will be omitted. Let P_d^n be the set of degree d polynomials 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 a $\mathbb{F}_p$-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.

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

We need the following Schwartz-Zippel-like Lemma for degree d polynomials.

Lemma 2.7 (Schwartz-Zippel lemma [9, Lemma 3.2]). *Let $f \in \mathbb{F}_p[x_1, \cdots, x_n]$ be a non-zero 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.* [1] 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*

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

The following lemma lists the basic properties of characters.

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

- *The characters of P_d^n are $\{\chi_\beta : \beta \in \mathfrak{F}_n\}$.*
- *For any $\beta, \beta' \in \mathfrak{F}_n$, $\chi_\beta = \chi'_\beta$ if and only if $\beta - \beta' \in (\mathsf{P}_d^n)^\perp$.*
- *For $\beta \in (\mathsf{P}_d^n)^\perp$, χ_β is the constant 1 function.*
- *$\forall \beta, \exists \beta'$ 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 forms an orthonormal basis for the vector space of functions from P_d^n to $\mathbb{C}$, under the inner product $\langle A, B \rangle := \mathbb{E}_{f \in \mathsf{P}_d^n}\left[A(f)\overline{B(f)}\right]$*
- *Any function $A : \mathsf{P}_d^n \to \mathbb{C}$ can be uniquely decomposed as*

$$A(f) = \sum_{\beta \in \Lambda_d^n} \widehat{A}(\beta)\chi_\beta(f) \text{ where } \widehat{A}(\beta) := \mathop{\mathbb{E}}_{g \in \mathsf{P}_d^n}\left[A(g)\overline{\chi_\beta(g)}\right],$$

  where 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 function $A : \mathsf{P}_d^n \to \mathbb{C}$, $\sum_{\beta \in \Lambda_d^n} |\widehat{A}(\beta)|^2 = \mathbb{E}_{f \in \mathsf{P}_d^n}[|A(f)|^2]$. In particular, if $A : \mathsf{P}_d^n \to \{1, \omega, \cdots, \omega^{p-1}\}$, $\sum_{\beta \in \Lambda_d^n} |\widehat{A}(\beta)|^2 = 1$.*

The following lemma relates characters over different domains related by co-ordinate projections.

Lemma 2.11. *Let $m \leq n$ and $\pi : \mathbb{F}_p^n \to \mathbb{F}_p^m$ be a (co-ordinate) projection i.e., there exist indices $1 \leq i_i < \cdots < i_m \leq n$ such that $\pi(x_1, \ldots, x_n) = (x_{i_1}, \cdots, x_{i_m})$. 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.

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

Dinur and Guruswami [4] proved the following theorem about Reed-Muller codes over $\mathbb{F}_2$ using Bhattacharyya *et. al.* [3] testing result.

Theorem 2.12 ([4, 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*

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

2.3 Folding over satisfying assignments

Lemma 2.13. *Let $d > 1$, X be a set of $p^d - 1$ points in $\mathbb{F}_p^n$ and $f : X \to \mathbb{F}_p$ 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 polynomial in $(\mathsf{P}_{(p-1)d}^n)^\perp$ has suppport size at least p^d. Hence, it is possible to interpolate a degree $(p-1)d$ polynomial through $p^d - 1$ points. $\square$

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 $q_1, \cdots, q_k \in \mathsf{P}_{3(p-1)}^n$, 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 non-zero on J (this is a generalization of the corresponding lemma in [4] to arbitrary fields).

Lemma 2.15. *Let $\beta \in \mathfrak{F}_n$ is such that $|\mathrm{support}(\beta)| < p^{d-3}$, and there exists $x \in \mathrm{support}(\beta)$ 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 which 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

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

$\square$

3 Correlation with a random square

In this section, we analyze the quantity

$$\langle \beta, p^2 \rangle,$$

where $p \in \mathsf{P}^n_d$ is chosen uniformly at random and $\beta : \mathbb{F}_3^n \to \mathbb{F}_3$ is a fixed function having distance exactly Δ from $(\mathsf{P}^n_{2d})^\perp = \mathsf{P}^n_{2n-2d-1}$.

Throughout this section, we work over the field $\mathbb{F}_3$. For $a \in \mathbb{N}^n$, let $|a| := \sum_i a_i$ and 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 wlog. that the coefficient vector $a \in \{0,1,2\}^n$. In this notation, $p(x) = \sum_{|a| \le d} p_a x^a$ 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., $e_x := (x^a)_{|a| \le d}$. Then $p(x) = p^T e_x$ where p is now thought of as the column vector $(p_a)_{|a| \le d}$ and hence, $p^2(x) = (p^T e_x)^2 = p^T (e_x e_x^T) p$.

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

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 $\mathsf{P}^n_{2n}/\mathsf{P}^n_{2n-2d-1}$ have the same value for $\langle \beta, p^2 \rangle$ and the matrix Q^β. Hence, by Lemma 2.10, we might wlog. assume that β satisifies $\mathrm{support}(\beta) = \Delta$. The following lemma (an easy consequence of [14, Theorem 6.21]), shows that it suffices to understand the rank of Q^β.

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

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

Lemma 3.2. $\mathrm{rank}(Q^\beta) \le \Delta$.

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

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

Proof. By assumption, β satisifies $\Delta = \mathrm{support}(\beta)$ 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 non-zero polynomial with degree $2n-d-1$ has support at least $3^{d/2}$ (Lemma 2.7), any $\lceil 3^{d/2} \rceil - 1$ vectors e_x are linearly independent. In particular, the Δ vectors e_x for x in $\mathrm{support}(\beta)$ are linearly independent. Consider any non-zero v in the kernel of the matrix Q^β. The linear independence of e_x's gives that $e_x^T v = 0$ for all $x \in \mathrm{support}(\beta)$. Hence, the kernel of Q^β resides in a Δ-codimensional space which implies that $\mathrm{rank}(Q^\beta) = \Delta$. □

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 $\Delta \ge 3^{d/2}$.

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

Proof. The proof of this theorem is similar to the proof of [4, Theorems 15,17] for the $\mathbb{F}_2$ case and we follow it step by step. Define $B^n_{d,k}(\beta) := \left\{q \in \mathsf{P}^n_k : q\beta \in \mathsf{P}^n_{2n-2d-1+k}\right\}$.

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 $\ker(Q^\beta)$.

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

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

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

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

Proof. Let β be the polynomial which 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

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

If $|a| \leq k$, $x^a \notin B^n_{d,k}(\beta)$ and we are done. Otherwise, consider b such that $b \leq a$ coordinate-wise 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. $\square$

For the induction step, we need the following result from Haramaty, Shpilka and Sudan [9].

Claim 3.7 ([9, Theorems 4.16, 1.7])**.** *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 $\forall c \in \mathbb{F}_3$, $\beta|_{\ell=c}$ 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 [9].

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 $\ell \in \mathsf{P}_1^n$ such that for all $c \in \mathbb{F}_3$, $\beta|_{\ell=c}$ is $\Delta/27$ 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 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 belongs 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

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

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, 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 $\mathrm{rank}(Q^\beta) \ge \Phi_{d,d}(3^{d/2}) \ge 3^{d/6-2} \ge 3^{d/9}$ 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 that we will fix later and let $I(\varphi)$ be an instance of 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 that we will determine later and a vertex $u \in U$, we define the subspace J_u of P_d^{3r} as follows:

$$J_u := \left\{ \sum_{i=1}^{3r} r_i f_i^u : r_i \in \mathsf{P}_{(d-3)}^{3r} \right\}.$$

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_d^{3r}. Let N_u denote the cardinality of the quotient space P_d^{3r}/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_d^{3r}/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 $A_u : \mathsf{P}_d^{3r}/J_u \to \mathbb{F}_2$ the restriction of A to the block $\mathcal{B}_u$ (under our identification of $\mathcal{B}_u$ with P_d^{3r}/J_u). Let $A'_u : \mathsf{P}_d^{3r} \to \mathbb{F}_2$ 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 unifom 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 $f \in \mathsf{P}_d^r$, $e_1, e_2, e_3, e_4 \in \mathsf{P}_d^{3r}$, and $g_1, g_2 \in \mathsf{P}_{d/4}^{3r}$ 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.

$$\begin{aligned} \eta_1 &:= 1 + f \circ \pi + g_1 h_1, & \eta_3 &:= f \circ \pi' + g_2 h_3, \\ \eta_2 &:= 1 + f \circ \pi + (1 + g_1) h_2, & \eta_4 &:= f \circ \pi' + (1 + g_2) h_4. \end{aligned}$$

3. Accept if and only if $A'_u(e_1), A'_u(e_1+\eta_1), A'_u(e_2), A'_u(e_2+\eta_2), A'_w(e_3), A'_w(e_3+\eta_3), A'_w(e_4), A'_w(e_4+\eta_4)$ 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 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 $g \in \mathsf{P}_d^{3r}$ 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 $v \in V$ and $u, w \in U$, 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:

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

If $f(b_v) = 0$, then either the first two values or the third and fourth values are unequal, whereas if $f(b_v) = 1$, then one of the last two pairs must be unequal. 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 $d \geq 8$ be a multiple of 4 , $\delta > 0$ and ε_0 be the constant from Theorem 2.2. If φ is unsatisfiable and H contains an independent set of size δN, then $\delta^8 \leq 2^{d/2} \cdot 2^{-\varepsilon_0 r} + 2^{-4 \cdot 2^{-d/4}}$.*

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,

$$\mathop{\mathbb{E}}_{u \in U} \mathop{\mathbb{E}}_{g \in \mathsf{P}_d^{3r}} \left[A'_u(g)\right] = \mathop{\mathbb{E}}_{u \in U} \mathop{\mathbb{E}}_{(g+J_u) \in \mathsf{P}_d^{3r}/J_u} [A_u(g)] \geq \delta. \tag{4.1}$$

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 2.2, let Λ_d^{3r} be a set of minimum weight coset representatives of the cosets of $(\mathsf{P}_d^{3r})^\perp$ in $\mathfrak{F}_{3r}$. Standard computations yield the following:

$$Q(v,u,w) = \sum_{\substack{\alpha_1,\alpha_2\\ \beta_1,\beta_2 \in \Lambda_d^{3r}}} \underbrace{\left(\prod_{i \in [2]} \widehat{A'_u}(\alpha_i)^2 \widehat{A'_w}(\beta_i)^2 \right) \mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}} \left[\prod_{i \in [2]} \chi_{\alpha_i}(\eta_i) \chi_{\beta_i}(\eta_{i+2}) \right]}_{\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)}. \tag{4.2}$$

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

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

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. $\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{FAR}} |\xi(\alpha_1, \alpha_2, \beta_1, \beta_2)| \leq 2^{-4\cdot 2^{d/4}}$.

Substituting in (4.2), we have for any $v \in V$ and $u, w \in U$,

$$\begin{aligned} Q(v, u, w) &\geq \sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}} \xi(\alpha_1, \alpha_2, \beta_1, \beta_2) - \sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{FAR}} |\xi(\alpha_1, \alpha_2, \beta_1, \beta_2)| \\ &\geq \sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR}} \xi(\alpha_1, \alpha_2, \beta_1, \beta_2) - 2^{-4\cdot 2^{-d/4}}. \end{aligned} \tag{4.3}$$

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.

$$\begin{aligned} &\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}} \left[\prod_{i\in[2]} \chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right] \\ &= \mathop{\mathbb{E}}_{\substack{g_1,g_2,f\\ h_1,\ldots,h_4}} \left[\chi_{\alpha_1}(1 + f\circ\pi + g_1h_1)\chi_{\alpha_2}(1 + f\circ\pi + (1+g_1)h_2)\chi_{\beta_1}(f\circ\pi' + g_2h_3)\chi_{\beta_2}(f\circ\pi' + (1+g_2)h_4)\right] \\ &= \mathop{\mathbb{E}}_{g_i,h_j} \left[\prod_{i\in[2]} \chi_{\alpha_i}(1 + (1+i+g_1)h_i)\chi_{\beta_i}((1+i+g_2)h_{i+2}) \cdot \mathop{\mathbb{E}}_{f}\left[\chi_{\pi_2(\alpha_1+\alpha_2)+\pi_2'(\beta_1+\beta_2)}(f)\right]\right]. \end{aligned} \tag{4.4}$$

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 $|\operatorname{support}(\gamma)| \geq 2^d > 4 \cdot 2^{d/2}$ since $d \geq 8$. This implies that $\gamma = 0$. Substituting in (4.4), we get

$$\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}} \left[\prod_{i\in[2]} \chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right] = \begin{cases} 0, \text{ if } \pi_2(\alpha_1 + \alpha_2) + \pi_2'(\beta_1 + \beta_2) \neq 0, \\ \mathbb{E}_{g_j,h_i}\left[\prod_{i\in[2]} \chi_{\alpha_i}(1 + (1+i+g_1)h_i)\chi_{\beta_i}((1+i+g_2)h_{i+2})\right], \text{otherwise.} \end{cases} \tag{4.5}$$

Substituting back in (4.3), we have

$$Q(v, u, w) = \sum_{\substack{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\text{NEAR:}\\ \pi_2(\alpha_1+\alpha_2)+\pi_2'(\beta_1+\beta_2)=0}} \xi(\alpha_1, \alpha_2, \beta_1, \beta_2) - 2^{-4\cdot 2^{-d/4}}. \tag{4.6}$$

We partition the terms in the above sum further into $\text{NEAR}_0 := \{(\alpha_1, \alpha_2, \beta_1, \beta_2) \in \text{NEAR} : \pi_2(\alpha_1+\alpha_2) = \pi_2'(\beta_1+\beta_2) = 0\}$ and $\text{NEAR}_1 := \{(\alpha_1, \alpha_2, \beta_1, \beta_2) \in \text{NEAR} : \pi_2(\alpha_1+\alpha_2) = \pi_2'(\beta_1+\beta_2) \neq 0\}$.

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

Claim 4.6. $\mathbb{E}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_0} \xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right] \ge \delta^8$.

Assuming these claims for now, we can finish the proof of Lemma 4.3 as follows. By (4.6),

$$\begin{aligned}
0 &= \mathop{\mathbb{E}}_{v,u,w}[Q(v,u,w)] \\
&\ge \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_0} \xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)\right] - \mathop{\mathbb{E}}_{v,u,w}\left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_1} |\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right] - 2^{-4\cdot 2^{-d/4}} \\
&\ge \delta^8 - 2^{d/2}\cdot 2^{-\varepsilon_0 r} - 2^{-4\cdot 2^{-d/4}}.
\end{aligned}$$

□

We now fill in the proofs of Claims 4.4–4.6.

Proof of Claim 4.4. Fix any $(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{FAR}$. Conditioned on any choice of f, the expectation term in $|\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)|$ may be bounded as follows:

$$\begin{aligned}
&\left|\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]}\chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]\right| \\
&= \left|\mathop{\mathbb{E}}_{\substack{g_1,g_2\\ h_1,\ldots,h_4}}\left[\chi_{\alpha_1}(1+f\circ\pi+g_1h_1)\chi_{\alpha_2}(1+f\circ\pi+(1+g_1)h_2)\chi_{\beta_1}(f\circ\pi'+g_2h_3)\chi_{\beta_2}(f\circ\pi'+(1+g_2)h_4)\right]\right| \\
&\le \mathop{\mathbb{E}}_{g_1,g_2}\left[\prod_{i\in[2]}\left|\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i}(1+f\circ\pi+(1+i+g_1)h_i)\right]\right|\cdot\left|\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i}(f\circ\pi'+(1+i+g_2)h_{i+2})\right]\right|\right] \\
&= \mathop{\mathbb{E}}_{g_1,g_2}\left[\prod_{i\in[2]}\left|\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i}((1+i+g_1)h_i)\right]\right|\cdot\left|\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i}((1+i+g_2)h_{i+2})\right]\right|\right] \\
&\le \mathop{\mathbb{E}}_{g_1,g_2}\left[\min\left\{\left|\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i}((1+i+g_1)h_i)\right]\right|,\left|\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i}((1+i+g_2)h_{i+2})\right]\right| : i\in[2]\right\}\right] \\
&\le \min\left\{\mathop{\mathbb{E}}_{g_1}\left[\left|\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i}((1+i+g_1)h_i)\right]\right|\right],\mathop{\mathbb{E}}_{g_2}\left[\left|\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i}((1+i+g_2)h_{i+2})\right]\right|\right] : i\in[2]\right\}. \qquad (4.7)
\end{aligned}$$

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 $(\alpha_1,\alpha_2,\beta_1,\beta_2)\in\mathrm{FAR}$, 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

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

Substituting the above in (4.7), we obtain

$$\left|\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]}\chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right]\right|\le 2^{-4\cdot 2^{d/2}}.$$

Thus, we obtain

$$\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{FAR}} |\xi(\alpha_1,\alpha_2,\beta_1,\beta_2)| \le 2^{-4\cdot 2^{d/2}} \cdot \sum_{\alpha_1,\alpha_2,\beta_1,\beta_2\in\Lambda_d^{3r}} \left(\prod_{i\in[2]} \widehat{A'_u}(\alpha_i)^2 \widehat{A'_w}(\beta_i)^2\right) \le 2^{-4\cdot 2^{d/2}},$$

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

$$\Pr_{(u,v)\in E, L_U, L_V} [\pi_{uv}(L_U(u)) = L_V(v)] \ge \frac{1}{2^{d/2}} \mathop{\mathbb{E}}_{v,u,w} \left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2)\in \mathrm{NEAR}_1} |\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)|\right]. \quad (4.8)$$

Since $OPT(I(\varphi)) \le 2^{-\varepsilon_0 r}$, 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.

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

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 (4.8) and hence proves the claim.

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

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

$$\xi(\alpha_1, \alpha_2, \beta_1, \beta_2) \ge 0. \quad (4.9)$$

Given (4.9), we have

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

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

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

where the first inequality follows from repeated applications of the Cauchy-Schwarz inequality and the last from (4.1).

For any v, u, w and $(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \mathrm{NEAR}_0$, it remains to prove (4.9) (i.e., non-negativity of $\xi_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)$). From (4.2), it suffices to argue the non-negativity of

$$\begin{aligned}\mathop{\mathbb{E}}_{\substack{\eta_1,\eta_2\\ \eta_3,\eta_4}}\left[\prod_{i\in[2]}\chi_{\alpha_i}(\eta_i)\chi_{\beta_i}(\eta_{i+2})\right] &= \mathop{\mathbb{E}}_{g_1,g_2}\left[\prod_{i\in[2]}\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i}(1+(1+i+g_1)h_i)\right]\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i}((1+i+g_2)h_{i+2})\right]\right]\\ &= \mathop{\mathbb{E}}_{g_1,g_2}\left[(-1)^{\sum_x \alpha_1(x)+\alpha_2(x)}\cdot\prod_{i\in[2]}\mathop{\mathbb{E}}_{h_i}\left[\chi_{\alpha_i(1+i+g_1)}(h_i)\right]\mathop{\mathbb{E}}_{h_{i+2}}\left[\chi_{\beta_i(1+i+g_2)}(h_{i+2})\right]\right],\end{aligned} \tag{4.10}$$

where we have used (4.5) 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 non-negative.

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

This shows that the right-hand side of (4.10) is non-negative. and hence proves (4.9). □

Proof of Theorem 1.1. Given the completeness (Lemma 4.1) and soundness (Lemma 4.3), we only need to fix parameters. Let $d = C\log r$ for a large enough constant $C \geq 8$ determined shortly. By Lemma 4.3, if H has an independent set of size δN, then $\delta^8 \leq 2^{d/2}\cdot 2^{-\varepsilon_0 r} + 2^{-4\cdot 2^{-d/4}} < 2^{-\varepsilon_0 r/2}$ 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 construct a new verifier each of whose queries correspond to two queries of the verifier described above. Let $I(\varphi)$, $G = (U, V, E)$, and J_u ($u \in U$) be as defined 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 $f \in \mathsf{P}_{d'}^r$, $e_1, e_2, e_3, e_4 \in \mathsf{P}_d^{3r}$, and $g_1, g_2 \in \mathsf{P}_{d/4}^{3r}$ 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.
$$\eta_1 := 1 + f \circ \pi + g_1 h_1, \qquad \eta_3 := f \circ \pi' + g_2 h_3,$$
$$\eta_2 := 1 + f \circ \pi + (1 + g_1) h_2, \qquad \eta_4 := f \circ \pi' + (1 + g_2) h_4.$$
3. Accept if and only if $A'_u(e_1, e_2), A'_u(e_1 + \eta_1, e_2 + \eta_2), A'_w(e_3, e_4), A'_w(e_3 + \eta_3, e_4 + \eta_4)$ 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. 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 non-zero 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 if and only if 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 if and only if $(\alpha_1 - \alpha_2), (\beta_1 - \beta_2) \in (\mathsf{P}_d^{3r})^\perp$.*
- *Folding: Fix $A : \mathsf{P}_d^{3r} \times \mathsf{P}_d^{3r} \to \mathbb{C}$ 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 , $\delta > 0$ and ε_0 be the constant from Theorem 2.2. If φ is unsatisfiable and H contains an independent set of size δN, then $\delta^4 \leq 2^{d/2} \cdot 2^{-\varepsilon_0 r} + 2^{-4 \cdot 2^{-d/4}}$.*

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 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:

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

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 $\textsc{far} := \{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in (\Lambda_d^{3r})^4 : \max\{\Delta(\alpha_i,\mathsf{P}_d^{3r}), \Delta(\beta_i,\mathsf{P}_d^{3r})\} \geq 2^{d/2}\}$, $\textsc{near} := (\Lambda_d^{3r})^4 \setminus \textsc{far}$, $\textsc{near}_0 := \{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \textsc{near} : \pi_2(\alpha_1+\alpha_2) = \pi'_2(\beta_1+\beta_2) = 0\}$, and $\textsc{near}_1 := \{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \textsc{near} : \pi_2(\alpha_1+\alpha_2) = \pi'_2(\beta_1+\beta_2) \neq 0\}$.

Note that the expectation term in $\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)$ is *exactly* 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, 4.5 and 4.6.

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

Claim 5.5. $\mathbb{E}_{v,u,w} \left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \textsc{near}_1} |\xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2)| \right] \leq 2^{d/2} \cdot 2^{-\varepsilon_0 r}$.

(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. $\mathbb{E}_{v,u,w} \left[\sum_{(\alpha_1,\alpha_2,\beta_1,\beta_2) \in \textsc{near}_0} \xi'_{v,u,w}(\alpha_1,\alpha_2,\beta_1,\beta_2) \right] \geq \delta^4$.

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

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

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 [11] 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 that 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 $G = (V_0, \ldots, V_{\ell-1}, \{E_{ij}\}_{0\leq i<j<\ell})$ 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 $\{c_1, \cdots c_{(T+\ell-i)r}\}$ be the clauses corresponding to v where $T = \lceil l/\eta \rceil$ as in Definition 2.4. We construct polynomials $\{p_1, \cdots p_{(T+\ell-i)r}\}$ 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 that we will determine later, for each vertex v define the subspace J_v as follows:

$$
J_v := \left\{ \sum_i q_i p_i : q_i \in \mathsf{P}^{m_v}_{2d-6} \right\} \text{ where } m_v := m_i = 3(T+\ell-i)r + ir.
$$

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 $N_v := |\mathsf{P}^{m_v}_{2d}/J_v|$. vertices, which we identify with elements of $\mathsf{P}^{m_v}_{2d}/J_v$. 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 $A_v : \mathsf{P}^{m_v}_{2d}/J_v \to \mathbb{F}_3$ the restriction of A to the block $\mathcal{B}_v$. Let $A'_v : \mathsf{P}^{m_v}_{2d} \to \mathbb{F}_3$ 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 $0 \leq i < j < \ell$ 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 $p \in \mathsf{P}^{m_u}_d, g \in \mathsf{P}^{m_u}_{2d}$ and $f \in \mathsf{P}^{m_v}_{2d}$ independently and uniformly at random and let $g' := p^2 + 1 - g - f \circ \pi$.
3. Accept if and only if $A'_v(f), A'_u(g), A'_u(g')$ are not all equal.

The above hypergraph construction explains the reasons (as in [7, 11]) for using the multi-layered 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 $\varphi \in$ 3SAT, then there is proof $A : V(H) \to \mathbb{F}_3$ which the verifier accepts with probability 1. In other words, the hypergraph H is 3-colorable.*

Proof. Since $\varphi \in$ 3SAT, Theorem 2.5 tells us that there are labelings $L_i : V_i \to \{0,1\}^{m_i}$ for $0 \leq i < \ell$ which satisfy all the constraints in $I(\varphi)$. For $\forall i, v \in V_i$, 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 $\varphi \notin$ 3SAT and H contains a independent set of size $\delta|V(H)|$, then*

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

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

$$Q(u, v) = \sum_{\alpha,\beta,\gamma} \widehat{A'_v}(\alpha)\widehat{A'_u}(\beta)\widehat{A'_u}(\gamma) \mathop{\mathbb{E}}_{f,g,p}\left[\chi_\alpha(f)\chi_\beta(g)\chi_\gamma(p^2 + 1 - f \circ \pi - g)\right], \tag{6.1}$$

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 non-zero 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 $\ell = 32/\delta^2$, there exists layers $0 \leq i < j < \ell$ such that $\mathbb{E}_{(u,v)\in E_{ij}}[\xi_{u,v}(0)] \geq \delta^5/2^9$.*

Proof. Since A' has fractional size δ, there exists a set S of vertices of fractional size $\delta/2$ such that $\forall v \in S, \widehat{A'_v}(0) = \mathbb{E}_f[A'_v(f)] \geq \delta/2$. Furthermore, there exists $\delta\ell/4$ 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 exists 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

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

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 $(u,v) \in E_{ij}$: $\text{FAR} := \{\beta \in \Lambda_{2d}^{m_i} : \Delta(\beta, (\mathsf{P}_{2d}^{m_i})^\perp) \geq 3^{d/2}\}$, $\text{NEAR}_1 := \{\beta \in \Lambda_{2d}^{m_i} \setminus \text{FAR} : \beta \neq 0 \text{ and } \pi_3(\beta) \notin (\mathsf{P}_{2d}^{m_v})^\perp\}$ and $\text{NEAR}_0 := \{\beta \in \Lambda_{2d}^{m_i} \setminus \text{FAR} : \beta \neq 0 \text{ and } \pi_3(\beta) \in (\mathsf{P}_{2d}^{m_v})^\perp\}$. 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. $\left|\mathbb{E}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in\text{FAR}} \xi_{u,v}(\beta)\right]\right| \leq \exp(-3^{\Omega(d)})$.

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

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

Combined with Claim 6.3, this exhausts all 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, 6.5 and 6.6.

Proof of Claim 6.4.

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

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 $\left|\mathbb{E}_p\left[\omega^{\langle\beta,p^2+1\rangle}\right]\right| \leq \exp(-3^{\Omega(d)})$. The claim follows since $\left|\widehat{A'_v}(\alpha)\right| \leq 1$ for any α and $\sum_\beta |\widehat{A'_u}(\beta)|^2 \leq 1$. □

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

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

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

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

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 labelcover from Theorem 2.5),

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

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

$$\mathop{\mathbb{E}}_{(u,v)\in E_{ij}}\left[\sum_{\beta\in \mathrm{NEAR}_0} |\widehat{A'_u}(\beta)|^2\cdot|\widehat{A'_v}(\pi_3(\beta))|\right] \le \mathop{\mathbb{E}}_{u\in V_i}\left[\sum_{\beta\notin \mathrm{FAR}\cup\{0\}} \Pr_{v:(u,v)\in E_{ij}}\left[\pi_3(\beta)\in (\mathsf{P}_{2d}^{m_v})^{\perp}\right]\cdot|\widehat{A'_u}(\beta)|^2\right].$$

We now argue that for every u and $\beta\notin \mathrm{FAR}\cup\{0\}$, $\Pr_{(u,v)\in E_{ij}}\left[\pi_3(\beta)\notin(\mathsf{P}_{2d}^{m_v})^{\perp}\right]$ is at most $3^d\cdot\eta$. This combined with the fact that $\sum_\beta|\widehat{A'_u}(\beta)|^2\le 1$ yields the claim. For every $u\in V_i$ and β such that $0\neq|\operatorname{support}(\beta)|=\Delta(\beta,(\mathsf{P}_{2d}^{m_u})^{\perp})\le 3^{d/2}$, by the smoothness property (Theorem 2.5), we have that with probability at least $1-3^d\eta$, we have

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

When (6.3) holds, we have $\pi_3(\beta)\neq 0$. Now since $|\operatorname{support}(\pi_3(\beta))|\le|\operatorname{support}(\beta)|\le 3^{d/2}$ and non-zero polynomials in $(\mathsf{P}_{2d}^{m_v})^{\perp}$ has 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 (Lemma 6.1) and soundness (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 produced by the reduction is of size $N=\ell n^{(1+1/\eta)\ell r}3^{((1+1/\eta)\ell r)^{O(d)}}$. Setting $\ell=C_4/(\delta')^2$ and $\log(1/\delta')=\Theta(\log\log n/\log\log\log n)$, we get that $N=n^{2^{O(\log\log n/\log\log\log n)}}$. Since $\log\log n=\Theta(\log\log N)$, we also get that $1/\delta'=2^{\Theta(\log\log N/\log\log\log N)}$. This completes the proof of Theorem 1.3. □

References

[1] Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer. Making the long code shorter. In *Proc. 53th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 370–379, 2012. arXiv:1111.0405, doi:10.1109/FOCS.2012.83.

[2] Mihir Bellare, Oded Goldreich, and Madhu Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM J. Computing*, 27(3):804–915, June 1998. (Preliminary version in *36th FOCS*, 1995). eccc:TR95-024, doi:10.1137/S0097539796302531.

[3] Arnab Bhattacharyya, Swastik Kopparty, Grant Schoenebeck, Madhu Sudan, and David Zuckerman. Optimal testing of Reed-Muller codes. In *Proc. 51st IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 488–497, 2010. arXiv:0910.0641, doi:10.1109/FOCS.2010.54.

[4] Irit Dinur and Venkatesan Guruswami. PCPs via low-degree long code and hardness for constrained hypergraph coloring. In *Proc. 54th IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 340–349, 2013. eccc:TR13-122, doi:10.1109/FOCS.2013.44.

[5] Irit Dinur, Venkatesan Guruswami, Subhash Khot, and Oded Regev. A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM J. Computing*, 34(5):1129–1146, 2005. (Preliminary version in *35th STOC*, 2003). arXiv:cs.CC/0304026, doi:10.1137/S0097539704443057.

[6] Irit Dinur and Gillat Kol. Covering CSPs. In *Proc. 28th IEEE Conference on Computational Complexity*, pages 207–218, 2013. eccc:TR12-088, doi:10.1109/CCC.2013.29.

[7] Irit Dinur, Oded Regev, and Clifford D. Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica*, 25(5):519–535, 2005. (Preliminary version in *43rd FOCS*, 2002). doi:10.1007/s00493-005-0032-4.

[8] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, July 1998. (Preliminary version in *28th STOC*, 1996). doi:10.1145/285055.285059.

[9] Elad Haramaty, Amir Shpilka, and Madhu Sudan. Optimal testing of multivariate polynomials over small prime fields. *SIAM J. Computing*, 42(2):536–562, 2013. (Preliminary version in *52nd FOCS*, 2011). eccc:TR11-059, doi:10.1137/120879257.

[10] Subhash Khot. Hardness results for approximate hypergraph coloring. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 351–359, 2002. doi:10.1145/509907.509962.

[11] Subhash Khot. Hardness results for coloring 3-colorable 3-uniform hypergraphs. In *Proc. 43rd IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 23–32, 2002. doi:10.1109/SFCS.2002.1181879.

[12] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Computing*, 37(1):319–357, 2007. (Preliminary version in *45th FOCS*, 2004). eccc:TR05-101, doi:10.1137/S0097539705447372.

[13] Subhash Khot and Rishi Saket. Hardness of finding independent sets in 2-colorable and almost 2-colorable hypergraphs. In *Proc. 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2014. (To appear). arXiv:1308.3247.

[14] Rudolf Lidl and Harald Niederreiter. *Finite Fields*, volume 2 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1997. doi:10.1017/CBO9780511525926.

[15] Ran Raz. A parallel repetition theorem. *SIAM J. Computing*, 27(3):763–803, June 1998. (Preliminary version in *27th STOC*, 1995). doi:10.1137/S0097539795280895.

A Proof of Claim 3.7

We need the following theorem due to Haramaty, Shpilka and Sudan [9].

Theorem A.1 ([9, Theorem 4.16, 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 $\Delta(\beta, \mathsf{P}_r^n) \leq 6\Delta_1 + 8 \cdot 3^n/K$.*

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 $(2n-2d-1)$ polynmial on A_i, then $\Delta(\beta, \mathsf{P}_{2n-2d-1}^n) \leq 6\Delta_1 + 8 \cdot 3^n/K$.

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 $K \geq 3^{n-d+\lambda_3}$ if $d \geq d_0 \geq \lambda_3 + 2$ 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 $\Delta(\beta, \mathsf{P}_{2n-2d-1}^n) \leq 6\Delta/27 + 2 \cdot 8 \cdot 3^n/(3^n-1) < 6\Delta/27 + 32 < \Delta$ (since $\Delta \geq 3^4$). This contradicts the hypothesis that β is Δ-far from $\mathsf{P}_{2n-2d-1}^n$. □

ECCC ISSN 1433-8092

http://eccc.hpi-web.de