

Approximating the Orthogonality Dimension of Graphs and Hypergraphs*

Ishay Haviv[†]

Abstract

A t -dimensional orthogonal representation of a hypergraph is an assignment of nonzero vectors in \mathbb{R}^t to its vertices, such that every hyperedge contains two vertices whose vectors are orthogonal. The orthogonality dimension of a hypergraph H , denoted by $\bar{\zeta}(H)$, is the smallest integer t for which there exists a t -dimensional orthogonal representation of H . In this paper we study computational aspects of the orthogonality dimension of graphs and hypergraphs. We prove that for every $k \geq 4$, it is NP-hard (resp. quasi-NP-hard) to distinguish n -vertex k -uniform hypergraphs H with $\bar{\zeta}(H) \leq 2$ from those satisfying $\bar{\zeta}(H) \geq \Omega(\log^\delta n)$ for some constant $\delta > 0$ (resp. $\bar{\zeta}(H) \geq \Omega(\log^{1-o(1)} n)$). For graphs, we relate the NP-hardness of approximating the orthogonality dimension to a variant of a long-standing conjecture of Stahl. We also consider the algorithmic problem in which given a graph G with $\bar{\zeta}(G) \leq 3$ the goal is to find an orthogonal representation of G of as low dimension as possible, and provide a polynomial time approximation algorithm based on semidefinite programming.

1 Introduction

A t -dimensional *orthogonal representation* of a hypergraph $H = (V, E)$ is an assignment of a nonzero real vector $u_v \in \mathbb{R}^t$ to every vertex $v \in V$, such that every hyperedge $e \in E$ contains two vertices $v, v' \in e$ whose vectors u_v and $u_{v'}$ are orthogonal. The *orthogonality dimension* of H , denoted by $\bar{\zeta}(H)$, is the smallest integer t for which there exists a t -dimensional orthogonal representation of H .¹ The notion of orthogonal representations was introduced for graphs by Lovász [43] in the study of the Shannon capacity and was later involved in a geometric characterization of connectivity properties of graphs by Lovász, Saks, and Schrijver [45]. The orthogonality dimension over the complex field was used by de Wolf [18] in a characterization of the quantum one-round communication complexity of promise equality problems and by Cameron et al. [14] in the study of the quantum chromatic number of graphs (see also [49, 10, 11]). An extension of orthogonal representations, called orthogonal bi-representations, was introduced by Haemers [31] and has found several further applications to information theory and to theoretical computer science.

*A preliminary version appeared in Proceedings of the 44th International Symposium on Mathematical Foundations of Computer Science (MFCS), pages 39:1–39:15, 2019.

[†]School of Computer Science, The Academic College of Tel Aviv-Yaffo, Tel Aviv 61083, Israel. Research supported in part by the Israel Science Foundation (grant No. 1218/20).

¹Orthogonal representations of *graphs* are sometimes defined in the literature as orthogonal representations of the complement, namely, the definition requires vectors associated with *non-adjacent* vertices to be orthogonal. In this paper we have decided to use the other definition because it is extended more naturally to hypergraphs. For a graph G , one can view the notation $\bar{\zeta}(G)$ as standing for $\zeta(\bar{G})$.

Orthogonal representations can be viewed as a generalization of hypergraph vertex colorings, one of the most fundamental and extensively studied topics in graph theory. Recall that a hypergraph H is said to be c -colorable if one can assign one of c colors to every vertex of H such that no hyperedge is monochromatic. The chromatic number of H , denoted by $\chi(H)$, is the smallest integer c for which H is c -colorable. Obviously, every c -coloring of H induces a c -dimensional orthogonal representation of H by assigning the i th unit vector $e_i \in \mathbb{R}^c$ to every vertex colored by the i th color. On the other hand, given a t -dimensional orthogonal representation $(u_v)_{v \in V}$ of H one can assign to every vertex v the vector in $\{-1, 0, +1\}^t$ that consists of the signs of the entries of u_v , and since nonzero orthogonal vectors have distinct sign vectors it follows that H is 3^t -colorable. We conclude that every hypergraph H satisfies

$$\log_3 \chi(H) \leq \bar{\xi}(H) \leq \chi(H). \quad (1)$$

The upper bound in (1) can clearly be tight (take, e.g., a complete graph), and it turns out that there exist graphs whose orthogonality dimension is exponentially smaller than their chromatic number (see Proposition 2.2).

The current work studies the problem of approximating the orthogonality dimension of graphs and hypergraphs. This research direction was already suggested in the late eighties by Lovász et al. [45], who remarked that computing the orthogonality dimension of graphs seems to be a difficult task (see also [44, Chapter 10]). Nevertheless, the only hardness result we are aware of for this graph parameter is the one of Peeters [47], who proved that for every $t \geq 3$ it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G) \leq t$ (see the recent follow-up work [27]). Before stating our hardness and algorithmic results, we overview related previous work on graph and hypergraph coloring.

1.1 Graph and Hypergraph Coloring

It is well known that the problem of deciding whether an input graph is c -colorable can be easily solved in polynomial time for $c \in \{1, 2\}$ and is NP-hard for every $c \geq 3$ [35].

In 1976, Garey and Johnson [25] have discovered an interesting connection between hardness of graph coloring and the multichromatic numbers of Kneser graphs. For integers $d \geq 2s$, the *Kneser graph* $K(d, s)$ is the graph whose vertices are all the s -subsets of $[d]$, where two sets are adjacent if they are disjoint. A k -tuple coloring f of a graph $G = (V, E)$ is an assignment $f(v)$ of a set of k colors to every vertex $v \in V$ such that $f(v) \cap f(v') = \emptyset$ whenever v and v' are adjacent in G . The k th *multichromatic number* of G , denoted by $\chi_k(G)$, is the smallest integer c for which G has a k -tuple coloring with c colors. Equivalently, $\chi_k(G)$ is the smallest integer c for which there exists a homomorphism from G to the Kneser graph $K(c, k)$. Note that $\chi_1(G)$ is simply the standard chromatic number $\chi(G)$. In the seventies, Stahl [51] has made the following conjecture regarding the multichromatic numbers of Kneser graphs (see also [23]).

Conjecture 1.1 ([51]). *For all integers k and $d \geq 2s$,*

$$\chi_k(K(d, s)) = \left\lceil \frac{k}{s} \right\rceil \cdot (d - 2s) + 2k.$$

More than forty years later, Conjecture 1.1 is still widely open. While the right-hand side in the conjecture is known to form an upper bound on $\chi_k(K(d, s))$ for all values of k , d and s , the

conjecture was confirmed only for a few special cases. For $k = 1$, the conjecture was proved by Lovász [42] in a breakthrough application of algebraic topology confirming a conjecture by Kneser [38]. Stahl [51] proved that the conjecture holds whenever $1 \leq k \leq s$, $d = 2s + 1$, or k is divisible by s . Garey and Johnson [25] proved the case of $s = 3$ and $k = 4$, namely, that $\chi_4(K(d, 3)) = 2d - 4$, and used it in the analysis of a reduction from 3-colorability to prove that for every $c \geq 6$, it is NP-hard to distinguish between graphs G that satisfy $\chi(G) \leq c$ and those that satisfy $\chi(G) \geq 2c - 4$.

In 1993, Khanna, Linial, and Safra [37] proved that it is NP-hard to decide whether an input graph G satisfies $\chi(G) \leq 3$ or $\chi(G) \geq 5$ (see also [30]). As observed in [9], combining this result with the proof technique of [25] and the case of $s = 3$ and $k = 5$ in Conjecture 1.1 proved by Stahl [52] (who confirmed there the conjecture for $s \leq 3$ and all integers k and d), it follows that for every $c \geq 6$ it is NP-hard to distinguish between the cases $\chi(G) \leq c$ and $\chi(G) \geq 2c - 2$. Brakensiek and Guruswami [9] improved this result using different techniques and proved the NP-hardness of deciding whether a given graph G satisfies $\chi(G) \leq c$ or $\chi(G) \geq 2c - 1$ for all $c \geq 3$. The latter condition was improved to $\chi(G) \geq 2c$ by Bulín, Krokhin, and Opršal [13], and recently it was significantly improved to $\chi(G) \geq \lfloor \frac{c}{2} \rfloor$ for all $c \geq 6$ by Wrochna and Živný [54]. We note that Dinur, Mossel, and Regev [19] proved that assuming a certain variant of the unique games conjecture, deciding whether a given graph G satisfies $\chi(G) \leq 3$ or $\chi(G) \geq c$ is NP-hard for every constant $c \geq 4$.

We next consider, for any constant $k \geq 3$, the problem of deciding whether an input k -uniform hypergraph (i.e., a hypergraph each of its hyperedges contains exactly k vertices) is c -colorable. While the problem can clearly be solved in polynomial time for $c = 1$, it was shown to be NP-hard for $c = 2$ and $k = 3$ by Lovász [41], motivating the study of the following gap problem: Given an n -vertex k -uniform hypergraph H , decide whether $\chi(H) \leq 2$ or $\chi(H) \geq c$. Guruswami, Håstad, and Sudan [29] proved that the problem is NP-hard for $k \geq 4$ and every constant $c \geq 3$. By combining their proof with the later PCP theorem of Moshkovitz and Raz [46], this result also follows for the super-constant $c = \Omega(\frac{\log \log n}{\log \log \log n})$. For $k = 3$, the NP-hardness was proved for every constant $c \geq 3$ by Dinur, Regev, and Smyth [20]. Their proof approach was extended in a recent work of Bhangale [6], who obtained NP-hardness for every $k \geq 4$ with the improved super-constant $c = \Omega(\log^\delta n)$ where $\delta > 0$ is some constant. Under the complexity assumption $\text{NP} \not\subseteq \text{DTIME}(2^{\text{poly}(\log n)})$, several stronger hardness results are known. This includes the case of $k = 3$ and $c = \Omega(\sqrt[3]{\log \log n})$ proved in [20] and the case of $k \geq 4$ and $c = \Omega(\frac{\log n}{\log \log n})$ proved in [6]. For additional related results see, e.g., [6, Table 1] and the references therein.

On the algorithmic side, significant efforts have been made in the literature to obtain polynomial time algorithms for coloring n -vertex 3-colorable graphs using as few colors as possible. This line of research was initiated by a simple algorithm of Wigderson [53] that used $O(\sqrt{n})$ colors. In a series of increasingly sophisticated combinatorial algorithms, Blum [7] improved the number of colors to $\tilde{O}(n^{3/8})$. Then, Karger, Motwani, and Sudan [34] introduced an algorithm based on a semidefinite relaxation and improved the number of colors to $\tilde{O}(n^{1/4})$. Combining the combinatorial approach of [7] and the semidefinite relaxation of [34], Blum and Karger [7, 8] improved it to $\tilde{O}(n^{3/14})$, which was later improved by Arora, Chlamtac, and Charikar [3] and by Chlamtac [16] to $\tilde{O}(n^{0.2111})$ and $\tilde{O}(n^{0.2072})$ respectively. The combinatorial component of these algorithms was recently improved by Kawarabayashi and Thorup [36], who reduced the number of colors to $\tilde{O}(n^{0.19996})$. Halperin et al. [32] have obtained analogue results for coloring n -vertex

c -colorable graphs for all constants $c \geq 4$, e.g., for $c = 4$ there exists an efficient algorithm that uses $\tilde{O}(n^{7/19})$ colors.

For hypergraphs, there exists a simple combinatorial algorithm that given an n -vertex k -uniform 2-colorable hypergraph finds in polynomial time a coloring with $\tilde{O}(n^{1-1/k})$ colors, as was shown independently by Alon et al. [2] and by Chen and Frieze [15]. For $k = 3$, this algorithm was combined in [2, 15] with the semidefinite programming approach of [34] to obtain a better bound of $\tilde{O}(n^{2/9})$, which was later improved to $\tilde{O}(n^{1/5})$ by Krivelevich, Nathaniel, and Sudakov [40]. We note, however, that Alon et al. [2] have provided evidence that the powerful semidefinite approach cannot be applied to coloring k -uniform hypergraphs for $k \geq 4$.

1.2 Our Contribution

The present paper offers hardness and algorithmic results on the orthogonality dimension of graphs and hypergraphs. We first mention that known hardness results on the chromatic number can be used to derive hardness results for the orthogonality dimension. Indeed, the inequalities given in (1) yield that for every integers t_1 and t_2 , NP-hardness of deciding whether an input k -uniform hypergraph H satisfies $\chi(H) \leq t_1$ or $\chi(H) \geq t_2$ immediately implies the NP-hardness of deciding whether it satisfies $\bar{\xi}(H) \leq t_1$ or $\bar{\xi}(H) \geq \log_3 t_2$. In particular, the hardness results of [20, 30] imply that for all constants $k \geq 3$ and $t \geq 3$, it is NP-hard to decide whether an input k -uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq t$. For $k = 2$, such a result follows from [19] under a variant of the unique games conjecture. However, for super-constant hardness gaps this implication leads to an exponential loss. In particular, it follows from [6] that for every $k \geq 4$ it is NP-hard to distinguish n -vertex k -uniform hypergraphs H with $\bar{\xi}(H) \leq 2$ from those satisfying $\bar{\xi}(H) \geq \Omega(\log \log n)$. We prove that this exponential loss can be avoided.

Theorem 1.2. *Let $k \geq 4$ be a fixed integer.*

1. *There exists a constant $\delta > 0$ for which it is NP-hard to decide whether an input n -vertex k -uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq \log^\delta n$.*
2. *Assuming $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)})$, for every constant $c > 0$ there is no polynomial time algorithm that decides whether an input n -vertex k -uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq c \cdot \frac{\log n}{\log \log n}$.*

We next consider the hardness of approximating the orthogonality dimension of *graphs*. Our result involves a generalization of orthogonal representations of graphs defined as follows. A t -dimensional *orthogonal k -subspace representation* of a graph $G = (V, E)$ is an assignment of a subspace $U_v \subseteq \mathbb{R}^t$ with $\dim(U_v) = k$ to every vertex $v \in V$, such that the subspaces U_v and $U_{v'}$ are orthogonal whenever v and v' are adjacent in G . For a graph G , let $\bar{\xi}_k(G)$ denote the smallest integer t for which there exists a t -dimensional orthogonal k -subspace representation of G . Note that $\bar{\xi}_1(G) = \bar{\xi}(G)$ for every graph G . We prove the following result.

Theorem 1.3. *For every graph F , it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G) \leq \bar{\xi}_3(F)$ or $\bar{\xi}(G) \geq \bar{\xi}_4(F)$.*

With Theorem 1.3 in hand, it is of interest to find graphs F for which $\bar{\xi}_4(F)$ is large compared to $\bar{\xi}_3(F)$. We consider here, in light of Conjecture 1.1, the behavior of the $\bar{\xi}_k$ parameters on the

Kneser graphs $K(d, s)$. For $k = 1$, it was recently shown that the orthogonality dimension and the chromatic number coincide on Kneser graphs [33, 1]. We further observe, as an application of a result of Bukh and Cox [12], that the values of χ_k and $\bar{\xi}_k$ coincide on the Kneser graphs $K(d, s)$ for every k divisible by s (that is, for all integers $\ell \geq 1$ and $d \geq 2s$, $\bar{\xi}_{\ell \cdot s}(K(d, s)) = \chi_{\ell \cdot s}(K(d, s)) = \ell \cdot d$, and in particular $\bar{\xi}_3(K(d, 3)) = d$; see Corollary 5.8). It would be natural to ask whether this is also the case for $k = 4$ and $s = 3$.

Question 1.4. *Is it true that for every $d \geq 6$, $\bar{\xi}_4(K(d, 3)) = 2d - 4$?*

A positive answer to Question 1.4 would imply that for every $t \geq 6$, it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G) \leq t$ or $\bar{\xi}(G) \geq 2t - 4$, analogously to the hardness result of [25] for the chromatic number. Note that a result of [12] implies that every $d \geq 6$ satisfies $\bar{\xi}_4(K(d, 3)) \geq \lceil 4d/3 \rceil$ (see Lemma 5.7). Combining this bound with Theorem 1.3, it follows that for every $t \geq 6$ it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G) \leq t$ or $\bar{\xi}(G) \geq \lceil 4t/3 \rceil$. In the recent follow-up work [27], it was further shown that there exists a constant c such that for every sufficiently large d , it holds that $\bar{\xi}_4(K(d, 3)) \geq 3d/2 - c$. By Theorem 1.3, this implies that for every sufficiently large t , it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G) \leq t$ or $\bar{\xi}(G) \geq 3t/2 - c$.

We finally consider the algorithmic problem in which given an n -vertex k -uniform hypergraph H with constant orthogonality dimension, the goal is to find an orthogonal representation of H of as low dimension as possible. It is not difficult to show that a hypergraph H satisfies $\bar{\xi}(H) \leq 2$ if and only if it is 2-colorable. Hence, by the algorithm of Krivelevich et al. [40], given an n -vertex 3-uniform hypergraph H with $\bar{\xi}(H) \leq 2$ it is possible to efficiently find a coloring of H that uses $\tilde{O}(n^{1/5})$ colors, and, in particular, to obtain an orthogonal representation of H of this dimension. For graphs, the first nontrivial case is where we are given as input an n -vertex graph G with $\bar{\xi}(G) \leq 3$, for which we prove the following result.

Theorem 1.5. *There exists a randomized polynomial time algorithm that given an n -vertex graph G satisfying $\bar{\xi}(G) \leq 3$, finds a coloring of G that uses at most $\tilde{O}(n^{0.2413})$ colors. In particular, the algorithm finds an orthogonal representation of G of dimension $\tilde{O}(n^{0.2413})$.*

In fact, we prove a stronger statement than that of Theorem 1.5, allowing the input graph G to satisfy $\bar{\xi}_k(G) \leq 3k$ for some integer k (rather than the special case $k = 1$; see Theorem 5.14).

1.3 Overview of Proofs

1.3.1 Hardness of Approximating the Orthogonality Dimension of Hypergraphs

Theorem 1.2 is proved in two steps. In the first, we show that approximating the orthogonality dimension of k -uniform hypergraphs becomes harder as the uniformity parameter k grows, and in the second we prove the hardness result for 4-uniform hypergraphs. By combining the two, Theorem 1.2 follows. We elaborate below on each of these two steps.

The uniformity reduction. Our goal is to show that for every $k_1 \leq k_2$, one can efficiently transform a given k_1 -uniform hypergraph H_1 to a k_2 -uniform hypergraph H_2 so that $\bar{\xi}(H_1) = \bar{\xi}(H_2)$. We borrow a reduction used in [29] for hypergraph coloring and prove that it preserves the orthogonality dimension. For simplicity of presentation, let us consider here the case of $k_1 = 2$ and

$k_2 = 4$. Given an n -vertex graph $G = (V, E)$ we construct a 4-uniform hypergraph H whose vertex set consists of ℓ copies V_1, \dots, V_ℓ of V . For $i \in [\ell]$, let E_i denote the collection of 2-subsets of V_i that correspond to the edges of G . The hyperedges of H are defined as all possible unions of pairs of sets picked from distinct collections E_i and E_j .

As a warm-up, we observe that for a sufficiently large ℓ , say $\ell = n + 1$, we have $\chi(G) = \chi(H)$. Indeed, if G is c -colorable then the c -coloring of G applied to each of the copies of V in H implies that H is c -colorable. On the other hand, if G is not c -colorable then for every coloring of H by c colors, every graph (V_i, E_i) contains a monochromatic edge. By $\ell > c$, there are $i \neq j$ and sets $e_1 \in E_i$ and $e_2 \in E_j$ such that all vertices of $e_1 \cup e_2$ share the same color. This implies the existence of a monochromatic hyperedge in H , hence H is not c -colorable.

We next show that for a sufficiently large ℓ we have $\bar{\chi}(G) = \bar{\chi}(H)$. The first direction is equally easy, namely, if G has a t -dimensional orthogonal representation then by assigning its vectors to every copy of V in H we get a t -dimensional orthogonal representation of H . For the other direction, assume that G satisfies $\bar{\chi}(G) > t$, and suppose for the sake of contradiction that $\bar{\chi}(H) \leq t$, i.e., there exists a t -dimensional orthogonal representation $(u_v)_{v \in V(H)}$ of H . By $\bar{\chi}(G) > t$, for every $i \in [\ell]$ there are two vertices $a_i, b_i \in V_i$ adjacent in G whose vectors are not orthogonal, that is, $\langle u_{a_i}, u_{b_i} \rangle \neq 0$. Now, it suffices to show that for some $i \neq j$ the four vectors $u_{a_i}, u_{b_i}, u_{a_j}, u_{b_j}$ are pairwise non-orthogonal, as this would imply a contradiction to the fact that $\{a_i, b_i, a_j, b_j\}$ is a hyperedge of H . It is not difficult to see that for some $i \neq j$ the vectors u_{a_i} and u_{b_j} are not orthogonal. Indeed, let $M_1 \in \mathbb{R}^{\ell \times t}$ be the matrix whose rows are the vectors u_{a_i} for $i \in [\ell]$, and let $M_2 \in \mathbb{R}^{\ell \times t}$ be the matrix whose rows are the vectors u_{b_i} for $i \in [\ell]$. Consider the matrix $M \in \mathbb{R}^{\ell \times \ell}$ defined by $M = M_1 \cdot M_2^T$, and notice that its diagonal entries are all nonzero (because $\langle u_{a_i}, u_{b_i} \rangle \neq 0$ for every i) and that its rank is at most t . Assuming that $\ell > t$, the matrix M must have some nonzero non-diagonal entry (otherwise its rank is ℓ), implying that $\langle u_{a_i}, u_{b_j} \rangle \neq 0$ for some $i \neq j$. This, however, still does not complete the argument, since it might be the case that for these indices i and j , one of the inner products $\langle u_{a_i}, u_{a_j} \rangle$, $\langle u_{b_i}, u_{a_j} \rangle$, and $\langle u_{b_i}, u_{b_j} \rangle$ is zero, avoiding the contradiction.

To overcome this difficulty, we use a symmetrization argument showing that the assumption $\bar{\chi}(H) \leq t$ implies that H has some t' -dimensional orthogonal representation $(w_v)_{v \in V(H)}$, where t' is not too large, with the following symmetry property: For every i and j , if $\langle w_{a_i}, w_{b_j} \rangle \neq 0$ then the inner products $\langle w_{a_i}, w_{a_j} \rangle$, $\langle w_{b_i}, w_{a_j} \rangle$, and $\langle w_{b_i}, w_{b_j} \rangle$ are all nonzero. With such an orthogonal representation, applying the above argument with $\ell > t'$ would certainly imply a contradiction and complete the proof. We achieve the symmetry property for $t' = t^4$ using vector tensor products. Namely, we assign to every vertex a_i the vector $w_{a_i} = u_{a_i} \otimes u_{b_i} \otimes u_{a_i} \otimes u_{b_i}$, to every vertex b_i the vector $w_{b_i} = u_{a_i}^{\otimes 2} \otimes u_{b_i}^{\otimes 2}$, and to every other vertex v the vector $w_v = u_v^{\otimes 4}$. It is straightforward to verify that $(w_v)_{v \in V(H)}$ forms a t' -dimensional orthogonal representation of H . Moreover, by standard properties of tensor products we have

$$\langle w_{a_i}, w_{b_j} \rangle = \langle u_{a_i}, u_{a_j} \rangle \cdot \langle u_{b_i}, u_{a_j} \rangle \cdot \langle u_{a_i}, u_{b_j} \rangle \cdot \langle u_{b_i}, u_{b_j} \rangle,$$

which can be used to obtain that if $\langle w_{a_i}, w_{b_j} \rangle \neq 0$ then the other three inner products $\langle w_{a_i}, w_{a_j} \rangle$, $\langle w_{b_i}, w_{a_j} \rangle$, and $\langle w_{b_i}, w_{b_j} \rangle$ are nonzero as well. This completes the proof sketch for $k_1 = 2$ and $k_2 = 4$. For the proof of the general case, we generalize the tensor-based argument to k -tuples of vertices (see Lemma 3.3) and extend the matrix reasoning applied above using bounds on off-diagonal Ramsey numbers (see Lemma 3.2 and Remark 3.5).

Hardness for 4-uniform hypergraphs. We next consider the hardness of approximating the orthogonality dimension of 4-uniform hypergraphs. A significant difficulty in proving such a result lies at the challenge of proving strong lower bounds on the orthogonality dimension. For the sake of comparison, in most hardness proofs for hypergraph coloring the lower bound on the chromatic number of the hypergraph H constructed by the reduction is shown by an upper bound on its independence number $\alpha(H)$ and the standard inequality $\chi(H) \geq \frac{|V(H)|}{\alpha(H)}$. This approach cannot be used for the orthogonality dimension, which in certain cases can be exponentially smaller than this ratio (see Proposition 2.2). Exceptions of this approach, where the lower bound on the chromatic number is not proved via the independence number, are the works of Dinur et al. [20] and Bhangale [6] on which we elaborate next.

A standard technique in proving hardness of approximation results is to reduce from the Label Cover problem, in which given a collection of constraints over a set of variables the goal is to decide whether there exists an assignment that satisfies all the constraints or any assignment satisfies only a small fraction of them (see Section 4.1). In such reductions, every variable over a domain $[R]$ is encoded via an error-correcting code known as the long code, and the constraints are replaced by “inner” constraints designed for the specific studied problem. One way to view the long code is as the graph whose vertices are all subsets of $[R]$ where two sets are adjacent if they are disjoint [21]. In the hardness proof of [20] for the chromatic number of 3-uniform hypergraphs, this graph was replaced by the induced subgraph that consists only of subsets of a given size (i.e., a Kneser graph), where a label $\alpha \in [R]$ is encoded by the 2-coloring of the vertices according to whether the sets contain α or not. The analysis of [20] is crucially based on the large chromatic number of Kneser graphs [42] and on the property that every coloring of Kneser graphs with number of colors smaller than their chromatic number enforces a large color class that includes a monochromatic edge. The latter property was proved in [20] using the chromatic number of the Schrijver graph [50], a vertex-critical subgraph of the Kneser graph. The approach of [20] was recently extended by Bhangale [6], who used in his long code construction only the vertices of the Schrijver graph. The fact that this subgraph has much fewer vertices and yet large chromatic number has led to improved hardness factors. However, for the analysis to work the “inner” constraints had to include four vertices, and this is the reason that the result was obtained for 4-uniform hypergraphs (and not for 3-uniform hypergraphs as in [20]).

In the current work we prove the hardness of approximating the orthogonality dimension of 4-uniform hypergraphs using the reduction applied in [6] for hypergraph coloring. While we achieve the same hardness factors as in [6], the analysis relies on several different ideas and tools. This includes the aforementioned symmetrization argument for orthogonal representations (see Lemma 3.3), a lower bound of Golovnev et al. [28] on the sparsity of low rank matrices with nonzero entries on the diagonal (see Lemma 4.5), and the orthogonality dimension of Schrijver graphs determined in [33] (see Theorem 2.3).

1.3.2 Hardness of Approximating the Orthogonality Dimension of Graphs

Theorem 1.3 relates the hardness of approximating the orthogonality dimension of graphs to orthogonal subspace representations. Our starting point is the NP-hardness of deciding whether an input graph G satisfies $\bar{\xi}(G) \leq 3$ [47]. Following an approach of Garey and Johnson [25], our reduction constructs a graph G' defined as the lexicographic product of some fixed graph F

and the input graph G . Namely, we replace every vertex of F by a copy of G and replace every edge of F by a complete bipartite graph between the vertex sets associated with its endpoints (see Definition 5.2). We then show that if $\bar{\xi}(G) \leq 3$ then G' has a $\bar{\xi}_3(F)$ -dimensional orthogonal representation, whereas if $\bar{\xi}(G) \geq 4$ the orthogonality dimension of G' is at least $\bar{\xi}_4(F)$. It would be interesting to figure out the best hardness factors that Theorem 1.3 can yield (see Question 1.4). We note, though, that our approach is limited to multiplicative hardness gaps bounded by 2, as it is easy to see that every graph F satisfies $\bar{\xi}_4(F) \leq \bar{\xi}_3(F) + \bar{\xi}_1(F) \leq 2 \cdot \bar{\xi}_3(F)$.

1.3.3 Coloring Graphs with Orthogonality Dimension Three

Consider the problem in which given an n -vertex graph G satisfying $\bar{\xi}(G) \leq 3$, the goal is to find an orthogonal representation of G of as low dimension as possible. Employing an approach of [17], we attempt to find a coloring of G with a small number of colors, as this in particular gives an orthogonal representation of the same dimension. As mentioned before, for every $c \geq 3$ there are known efficient algorithms for coloring n -vertex c -colorable graphs, however, our only guarantee on G is that its orthogonality dimension is at most 3. Interestingly, it follows from a theorem of Kochen and Specker [39] (see also [26]) that the largest possible chromatic number of such a graph is 4. It follows that given an n -vertex graph G with $\bar{\xi}(G) \leq 3$, one can simply apply the efficient algorithm of [32] for coloring 4-colorable graphs to obtain a coloring of G by $\tilde{O}(n^\gamma)$ colors where $\gamma = 7/19 \approx 0.368$.

To improve on this bound, we show an efficient algorithm that finds a large independent set in a given graph G satisfying $\bar{\xi}(G) \leq 3$. We consider two cases according to the maximum degree in the graph. If G has a vertex of large degree then the algorithm finds a large independent set in its neighborhood. This can be done since the assumption $\bar{\xi}(G) \leq 3$ implies that the neighborhood of every vertex of G is 2-colorable (just as for 3-colorable graphs). Otherwise, in case that all the degrees in G are small, we use an algorithm of Karger et al. [34] based on a semidefinite relaxation of the chromatic number, called the *vector chromatic number*. Our analysis relies on a result by Lovász [43] relating the (strict) vector chromatic number of graphs to their orthogonality dimension. Now, by repeatedly omitting independent sets in G , we obtain a coloring that uses $\tilde{O}(n^{1/4})$ colors. This can be slightly improved to $\tilde{O}(n^{0.2413})$ by applying the refined analysis of Arora et al. [3] for the rounding algorithm of [34].

As already mentioned, our algorithm can handle any graph G that satisfies $\bar{\xi}_k(G) \leq 3k$ for some integer k , rather than for $k = 1$ (see Theorem 5.14). The generalized analysis involves a connection, recently proved by Bukh and Cox [12], between the (strict) vector chromatic number and the graph parameters $\bar{\xi}_k$.

1.4 Outline

The rest of the paper is organized as follows. In Section 2, we provide some background on the orthogonality dimension and on the Kneser and Schrijver graphs. In Section 3, we present and analyze the uniformity reduction. In Section 4, we prove our hardness result for 4-uniform hypergraphs and derive Theorem 1.2. In the final Section 5, we prove our hardness and algorithmic results on the orthogonality dimension of graphs, confirming Theorems 1.3 and 1.5.

2 Preliminaries

2.1 Orthogonality Dimension

We define the orthogonality dimension of hypergraphs over a general field.

Definition 2.1. A t -dimensional orthogonal representation of a hypergraph $H = (V, E)$ over a field \mathbb{F} is an assignment of a vector $u_v \in \mathbb{F}^t$ with $\langle u_v, u_v \rangle \neq 0$ to every vertex $v \in V$, such that for every hyperedge $e \in E$ there are two vertices $v, v' \in e$ satisfying $\langle u_v, u_{v'} \rangle = 0$. The orthogonality dimension of a hypergraph $H = (V, E)$ over \mathbb{F} , denoted by $\bar{\xi}(H, \mathbb{F})$, is the smallest integer t for which there exists a t -dimensional orthogonal representation of H over \mathbb{F} . For the real field \mathbb{R} , we let $\bar{\xi}(H)$ stand for $\bar{\xi}(H, \mathbb{R})$.

The following proposition shows that there are graphs whose orthogonality dimension is exponentially smaller than their chromatic number.

Proposition 2.2. There exists an explicit family of graphs G_t such that $\bar{\xi}(G_t) \leq t$ and $\chi(G_t) \geq 2^{\Omega(t)}$.

Proof: Let t be an integer divisible by 4. Consider the graph $G_t = (V, E)$ whose vertices are all the $\frac{t}{2}$ -subsets of $[t]$ where two sets $A, B \in V$ are adjacent if their intersection size satisfies $|A \cap B| = \frac{t}{4}$. Notice that by $|A| = |B| = \frac{t}{2}$, this condition is equivalent to $|A \triangle B| = \frac{t}{2}$. Assign to every vertex A the vector $u_A \in \{\pm 1\}^t$ where $(u_A)_i = +1$ if $i \in A$ and $(u_A)_i = -1$ otherwise. We claim that $(u_A)_{A \in V}$ is an orthogonal representation of G_t . Indeed, for every adjacent vertices $A, B \in V$ we have

$$\langle u_A, u_B \rangle = (-1) \cdot |A \triangle B| + (t - |A \triangle B|) = t - 2 \cdot |A \triangle B| = 0.$$

This implies that $\bar{\xi}(G_t) \leq t$. On the other hand, by a celebrated result of Frankl and Rödl [24], $\alpha(G_t) \leq 2^{c \cdot t}$ for some $c < 1$, implying that

$$\chi(G_t) \geq \frac{|V|}{\alpha(G_t)} \geq \frac{\binom{t}{t/2}}{2^{c \cdot t}} \geq 2^{(1-c-o(1)) \cdot t},$$

completing the proof. ■

2.2 Kneser and Schrijver Graphs

For integers $d \geq 2s$, the *Kneser graph* $K(d, s)$ is the graph whose vertices are all the s -subsets of $[d]$, where two sets are adjacent if they are disjoint.

A set $A \subseteq [d]$ is said to be *stable* if it does not contain two consecutive elements modulo d (that is, if $i \in A$ then $i + 1 \notin A$, and if $d \in A$ then $1 \notin A$). In other words, a stable subset of $[d]$ is an independent set in the cycle C_d with the numbering from 1 to d along the cycle. For integers $d \geq 2s$, the *Schrijver graph* $S(d, s)$ is the graph whose vertices are all the stable s -subsets of $[d]$, where two sets are adjacent if they are disjoint.

The orthogonality dimension of the Kneser and Schrijver graphs was determined in [33] using topological methods.

Theorem 2.3 ([33]). For every $d \geq 2s$, $\bar{\xi}(K(d, s)) = \bar{\xi}(S(d, s)) = d - 2s + 2$.

The number of vertices in $K(d, s)$ is clearly $\binom{d}{s}$. We need the following simple bound, given in [20], on the number of vertices in $S(d, s)$.

Lemma 2.4 ([20]). For every $d \geq 2s$, the number of vertices in $S(d, s)$ is at most $\binom{d}{d-2s}$.

3 The Uniformity Reduction

In this section we reduce the problem of approximating the orthogonality dimension of k_1 -uniform hypergraphs over a field \mathbb{F} to approximating it on k_2 -uniform hypergraphs, where $k_1 \leq k_2$.

3.1 Ramsey Numbers

For integers s and t , the *Ramsey number* $R(s, t)$ is defined as the smallest integer n such that every n -vertex graph contains a clique of size s or an independent set of size t (or both). We need the following well-known upper bound on Ramsey numbers due to Erdős and Szekeres [22].

Theorem 3.1 ([22]). *For all integers s and t , $R(s, t) \leq \binom{t+s-2}{s-1}$.*

The following lemma relates Ramsey numbers to nonzero patterns of matrices with nonzero entries on the diagonal.

Lemma 3.2. *Let \mathbb{F} be a field, and let M be an $n \times n$ matrix over \mathbb{F} with nonzero entries on the diagonal. Denote $t = \text{rank}_{\mathbb{F}}(M)$ and suppose that $n \geq R(s, t + 1)$. Then there exists a set $C \subseteq [n]$ of size $|C| = s$ such that for every $i, i' \in C$ it holds that $M_{i,i'} \neq 0$ or $M_{i',i} \neq 0$ (or both).*

Proof: Let $M \in \mathbb{F}^{n \times n}$ be a matrix with nonzero entries on the diagonal such that $t = \text{rank}_{\mathbb{F}}(M)$ and $n \geq R(s, t + 1)$. Define a graph G on the vertex set $[n]$ where two distinct vertices i and i' are adjacent if $M_{i,i'} \neq 0$ or $M_{i',i} \neq 0$ (or both). Observe that any principal sub-matrix of M that corresponds to an independent set in G is diagonal and has nonzero entries on the diagonal. This implies that the size of any independent set in G does not exceed the rank t of M , that is, $\alpha(G) \leq t$. Since G has at least $R(s, t + 1)$ vertices, it follows that G contains a clique C of size s . This set C satisfies the requirement of the lemma. ■

3.2 Symmetrization Lemma

We turn to prove our symmetrization lemma for orthogonal representations, which is used in the analysis of the uniformity reduction (Section 3.3) and in our hardness proof for 4-uniform hypergraphs (Section 4). The lemma says, roughly speaking, that given an orthogonal representation of a hypergraph H and a collection A of pairwise disjoint k -tuples of vertices of H , one can construct another orthogonal representation of H with a polynomially related dimension such that the following property holds: For every two k -tuples a and b in A such that the vectors associated with the first vertex of a and the second vertex of b are not orthogonal, it holds that all the vectors associated with the vertices of a are not orthogonal to all the vectors associated with the vertices of b . The lemma is used to prove lower bounds on the orthogonality dimension of hypergraphs. Indeed, in order to show that an assignment of vectors to the vertices does not form an orthogonal representation, one has to show that the vectors associated with the vertices of some hyperedge are pairwise non-orthogonal. For a hyperedge that consists of vectors from some k -tuples of A , the lemma essentially allows us to consider only the first two vertices of every such k -tuple.

In the proof of the lemma we use the notion of tensor product of vectors. Recall that for a field \mathbb{F} , the tensor product of the vectors $u \in \mathbb{F}^n$ and $u' \in \mathbb{F}^{n'}$, denoted by $u \otimes u'$, is a vector in $\mathbb{F}^{n \cdot n'}$ with coordinates corresponding to all products $u_i \cdot u'_{i'}$ for $i \in [n]$ and $i' \in [n']$. We let $u^{\otimes k}$ denote

the tensor product of the vector u with itself k times. It is well known and easy to verify that for every $u, w \in \mathbb{F}^n$ and $u', w' \in \mathbb{F}^{n'}$ it holds that $\langle u \otimes u', w \otimes w' \rangle = \langle u, w \rangle \cdot \langle u', w' \rangle$. We also need the following notation. For a set V and a k -tuple $a \in V^k$ we let $a[i]$ stand for the i th component of a , that is, $a = (a[1], a[2], \dots, a[k])$.

Lemma 3.3. *Let $k \geq 2$ be an integer, let $H = (V, E)$ be a hypergraph with a t -dimensional orthogonal representation $(u_v)_{v \in V}$ over a field \mathbb{F} , and let $A \subseteq V^k$ be a collection of pairwise disjoint k -tuples of vertices of H . Suppose that for every $a \in A$ and $i, i' \in [k]$, $\langle u_{a[i]}, u_{a[i']} \rangle \neq 0$. Then there exists a t^{k^2} -dimensional orthogonal representation $(w_v)_{v \in V}$ of H over \mathbb{F} such that*

1. *for every $a \in A$ and $i, i' \in [k]$, $\langle w_{a[i]}, w_{a[i']} \rangle \neq 0$, and*
2. *for every (distinct) $a, b \in A$ such that $\langle w_{a[1]}, w_{b[2]} \rangle \neq 0$ it holds that $\langle w_{a[i]}, w_{b[i']} \rangle \neq 0$ for all $i, i' \in [k]$.*

Proof: For a hypergraph $H = (V, E)$ and a collection $A \subseteq V^k$ of pairwise disjoint k -tuples of vertices of H , let $(u_v)_{v \in V}$ be a t -dimensional orthogonal representation of H over \mathbb{F} as in the lemma. We assign to the vertices of H t^{k^2} -dimensional vectors $(w_v)_{v \in V}$ over \mathbb{F} as follows. For every $a \in A$, define

$$\begin{aligned} w_{a[1]} &= (u_{a[1]} \otimes u_{a[2]} \otimes \dots \otimes u_{a[k]})^{\otimes k}, \\ w_{a[2]} &= u_{a[1]}^{\otimes k} \otimes u_{a[2]}^{\otimes k} \otimes \dots \otimes u_{a[k]}^{\otimes k}. \end{aligned}$$

For every other vertex $v \in V$, i.e., a vertex that does not appear in the first two coordinates of the k -tuples of A , we define $w_v = u_v^{\otimes k^2}$. Note that the assignment is well defined since the k -tuples of A are pairwise disjoint. (In fact, we could weaken the condition of pairwise disjointness in the lemma and allow $a[i] = b[i']$ for $a, b \in A$ and $i, i' \geq 3$.)

We first observe that the assignment $(w_v)_{v \in V}$ is an orthogonal representation of H over \mathbb{F} . For every vertex $v \in V$, the vector w_v is a tensor product of k^2 vectors of the form $u_{v'}$ with $v' \in V$. This implies that $\langle w_v, w_v \rangle$ is a product of inner products of the form $\langle u_{v'}, u_{v'} \rangle$ with $v' \in V$, and since they are all nonzero it follows that $\langle w_v, w_v \rangle \neq 0$. Now, since $(u_v)_{v \in V}$ is an orthogonal representation of H , it suffices to show that for every two vertices $v, v' \in V$, it holds that $\langle w_v, w_{v'} \rangle = 0$ whenever $\langle u_v, u_{v'} \rangle = 0$. Consider two vertices $v, v' \in V$ such that $\langle u_v, u_{v'} \rangle = 0$. Notice that w_v and $w_{v'}$ are tensor products of k^2 vectors, the $(k+1)$ th of which is u_v and $u_{v'}$ respectively, hence $\langle w_v, w_{v'} \rangle = 0$.

We next verify that the orthogonal representation $(w_v)_{v \in V}$ of H over \mathbb{F} satisfies the properties required by the lemma. For the first item, observe that for every $a \in A$ each vector $w_{a[i]}$ is a tensor product of k^2 vectors of the form $u_{a[j]}$ with $j \in [k]$. Hence, the inner product $\langle w_{a[i]}, w_{a[i']} \rangle$ for $i, i' \in [k]$ is a product of inner products of the form $\langle u_{a[j]}, u_{a[j']} \rangle$ with $j, j' \in [k]$, which are all nonzero by assumption, hence $\langle w_{a[i]}, w_{a[i']} \rangle \neq 0$. For the second item, consider distinct $a, b \in A$ such that $\langle w_{a[1]}, w_{b[2]} \rangle \neq 0$. Observe that

$$\langle w_{a[1]}, w_{b[2]} \rangle = \prod_{(j, j') \in [k] \times [k]} \langle u_{a[j]}, u_{b[j']} \rangle,$$

which implies that $\langle u_{a[j]}, u_{b[j']} \rangle \neq 0$ for all $j, j' \in [k]$. Since every inner product $\langle w_{a[i]}, w_{b[i']} \rangle$ for $i, i' \in [k]$ is a product of inner products of the form $\langle u_{a[j]}, u_{b[j']} \rangle$ with $j, j' \in [k]$, we derive that $\langle w_{a[i]}, w_{b[i']} \rangle \neq 0$, and we are done. \blacksquare

3.3 The Reduction

We are ready to prove the main result of this section.

Theorem 3.4. *Let $k_2 \geq k_1 \geq 2$ be constants. For every parameter $m = m(n) \leq n$ there exists a polynomial time reduction that given a k_1 -uniform hypergraph H_1 on n vertices outputs a k_2 -uniform hypergraph H_2 on $n \cdot m^{O(1)}$ vertices such that for every field \mathbb{F} ,*

1. $\bar{\xi}(H_2, \mathbb{F}) \leq \bar{\xi}(H_1, \mathbb{F})$, and
2. if $\bar{\xi}(H_2, \mathbb{F}) \leq m$ then $\bar{\xi}(H_1, \mathbb{F}) \leq \bar{\xi}(H_2, \mathbb{F})$.

Proof: Let $k_2 \geq k_1 \geq 2$ be constants and put $s = \lceil k_2/k_1 \rceil$. For a given parameter $m = m(n) \leq n$ consider the reduction that given a k_1 -uniform hypergraph $H_1 = (V, E)$ on n vertices outputs the k_2 -uniform hypergraph $H_2 = (V', E')$ defined as follows. Its vertex set is defined as $V' = V_1 \cup \dots \cup V_\ell$ where each V_i is a copy of the vertex set V of H_1 and

$$\ell = \binom{m^{k_1^2} + s - 1}{s - 1}. \quad (2)$$

Let E_i denote the collection of k_1 -subsets of V_i that correspond to the hyperedges in the hypergraph H_1 . A hyperedge of H_2 is given by the union of $s - 1$ sets picked from different E_i 's together with additional $k_2 - (s - 1) \cdot k_1 \leq k_1$ vertices chosen from a set picked from another E_i . More precisely, for every distinct indices i_1, \dots, i_s , choices of $e_{i_j} \in E_{i_j}$ for $j \in [s]$, and a set $e'_{i_s} \subseteq e_{i_s}$ of size $|e'_{i_s}| = k_2 - (s - 1) \cdot k_1$ we include in H_2 the hyperedge $e_{i_1} \cup \dots \cup e_{i_{s-1}} \cup e'_{i_s}$. Note that H_2 is a k_2 -uniform hypergraph on $n \cdot \ell = n \cdot O(m^{k_1^2 \cdot (s-1)})$ vertices and that it can be constructed in polynomial running time.

To prove Item 1 of the theorem, assume that there exists a t -dimensional orthogonal representation of H_1 over \mathbb{F} . For the hypergraph H_2 , assign the same vectors to each of the ℓ copies of the vertex set V . Since every hyperedge of H_2 contains a k_1 -subset from some E_i , two of the vectors associated with its vertices are orthogonal. It follows that this assignment is a t -dimensional orthogonal representation of H_2 over \mathbb{F} , hence $\bar{\xi}(H_2, \mathbb{F}) \leq \bar{\xi}(H_1, \mathbb{F})$.

To prove Item 2, denote $t = \bar{\xi}(H_2, \mathbb{F}) \leq m$ and let $(u_v)_{v \in V'}$ be a t -dimensional orthogonal representation of H_2 over \mathbb{F} . Assume by contradiction that $\bar{\xi}(H_1, \mathbb{F}) > t$. Then, the restriction $(u_v)_{v \in V_i}$ of the given orthogonal representation to any V_i does not form an orthogonal representation of the hypergraph (V_i, E_i) . This implies that for every $i \in [\ell]$ there exists a k_1 -tuple $a_i \in V_i^{k_1}$ whose vertices form a hyperedge in (V_i, E_i) such that $\langle u_{a_i[j]}, u_{a_i[j']} \rangle \neq 0$ for all $j, j' \in [k_1]$.

Applying Lemma 3.3 to the orthogonal representation $(u_v)_{v \in V'}$ of H_2 over \mathbb{F} with the pairwise disjoint k_1 -tuples $\{a_i \mid i \in [\ell]\}$, we obtain a $t^{k_1^2}$ -dimensional orthogonal representation $(w_v)_{v \in V'}$ of H_2 over \mathbb{F} that satisfies the following properties:

1. $\langle w_{a_i[j]}, w_{a_i[j']} \rangle \neq 0$ for every $i \in [\ell]$ and $j, j' \in [k_1]$, and
2. for every (distinct) $i, i' \in [\ell]$ such that $\langle w_{a_i[1]}, w_{a_{i'}[2]} \rangle \neq 0$ it holds that $\langle w_{a_i[j]}, w_{a_{i'}[j']} \rangle \neq 0$ for all $j, j' \in [k_1]$.

Let $M_1 \in \mathbb{F}^{\ell \times t^{k_1^2}}$ be the matrix whose rows are the vectors $w_{a_i[1]}$ for $i \in [\ell]$, and let $M_2 \in \mathbb{F}^{\ell \times t^{k_1^2}}$ be the matrix whose rows are the vectors $w_{a_i[2]}$ for $i \in [\ell]$. Consider the matrix $M \in \mathbb{F}^{\ell \times \ell}$ defined

by $M = M_1 \cdot M_2^T$, and notice that $M_{i,i'} = \langle w_{a_i[1]}, w_{a_{i'}[2]} \rangle$ for every $i, i' \in [\ell]$ and that its rank is at most $t^{k_1^2}$. Property (1) of the orthogonal representation $(w_v)_{v \in V'}$ implies that the diagonal entries of M are all nonzero. By the Erdős-Szekeres bound on Ramsey numbers (Theorem 3.1) and the definition of ℓ in (2), we have $\ell \geq R(s, m^{k_1^2} + 1) \geq R(s, t^{k_1^2} + 1)$. Hence, we can apply Lemma 3.2 to obtain that there exists a set $C \subseteq [\ell]$ of size $|C| = s$ such that for every $i, i' \in C$ it holds that $M_{i,i'} \neq 0$ or $M_{i',i} \neq 0$.

To complete the proof recall that the hypergraph H_2 includes a hyperedge whose vertices all appear in the k_1 -tuples a_i with $i \in C$. We will get a contradiction by proving that the vectors assigned to these vertices by the orthogonal representation $(w_v)_{v \in V'}$ of H_2 are pairwise non-orthogonal. Indeed, by Property (1) of $(w_v)_{v \in V'}$, we have $\langle w_{a_i[j]}, w_{a_{i'}[j']} \rangle \neq 0$ for every $i \in C$ and $j, j' \in [k_1]$. In addition, for distinct $i, i' \in C$ it holds that $M_{i,i'} \neq 0$ or $M_{i',i} \neq 0$, that is, $\langle w_{a_i[1]}, w_{a_{i'}[2]} \rangle \neq 0$ or $\langle w_{a_{i'}[1]}, w_{a_i[2]} \rangle \neq 0$. By Property (2) of $(w_v)_{v \in V'}$ we have $\langle w_{a_i[j]}, w_{a_{i'}[j']} \rangle \neq 0$ for all $i, i' \in C$ and $j, j' \in [k_1]$, and we are done. ■

Remark 3.5. We note that the use of Ramsey numbers in the proof of Theorem 3.4 is not essential. As pointed out to us by an anonymous reviewer, to prove the assertion of the theorem it suffices to analyze the reduction for the case $k_2 = k_1 + 1$ which can be repeatedly applied to imply the general result. For this special case, the Ramsey numbers are not needed as demonstrated in the discussion in Section 1.3. Yet, we have decided to present here the direct analysis for general k_1 and k_2 since it shows that the reduction used in [29] for hypergraph coloring perfectly preserves the orthogonality dimension, and because the running time of this reduction is slightly better than that of repeatedly applying the reduction for $k_2 = k_1 + 1$.

We derive the following corollary.

Corollary 3.6. Let $k_2 \geq k_1 \geq 2$ be constants. There exists a polynomial time reduction that given a k_1 -uniform hypergraph H_1 on n vertices outputs a k_2 -uniform hypergraph H_2 on $n^{O(1)}$ vertices such that for every field \mathbb{F} , $\bar{\xi}(H_1, \mathbb{F}) = \bar{\xi}(H_2, \mathbb{F})$.

Proof: Apply Theorem 3.4 with $m(n) = n$. By Item 1, we have $\bar{\xi}(H_2, \mathbb{F}) \leq \bar{\xi}(H_1, \mathbb{F}) \leq n$. Combining this with Item 2, we get that $\bar{\xi}(H_1, \mathbb{F}) \leq \bar{\xi}(H_2, \mathbb{F})$, and we are done. ■

4 The Orthogonality Dimension of Hypergraphs

We prove the following hardness result for the orthogonality dimension of 4-uniform hypergraphs over the real field (see Remark 4.11).

Theorem 4.1.

1. There exists a constant $\delta > 0$ for which it is NP-hard to decide whether an input n -vertex 4-uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq \log^\delta n$.
2. Assuming $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)})$, for every constant $c > 0$ there is no polynomial time algorithm that decides whether an input n -vertex 4-uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq c \cdot \frac{\log n}{\log \log n}$.

Theorem 1.2 follows by combining Theorem 4.1 with Corollary 3.6.

4.1 Label Cover

Theorem 4.1 is proved by a reduction from the Label Cover problem, defined as follows.

Definition 4.2. In the Label Cover problem the input $\mathcal{L} = (U, V, E, R, L, \phi)$ consists of a (bi-regular) graph (U, V, E) where every vertex of U is associated with a variable with range $[R]$ and every vertex of V is associated with a variable with range $[L]$. Every edge $(x, z) \in E$ is associated with a projection constraint $\phi_{x \rightarrow z} : [R] \rightarrow [L]$. Given an instance \mathcal{L} the goal is to find an assignment ρ to the variables of $U \cup V$ that maximizes the number of edges $(x, z) \in E$ such that $\phi_{x \rightarrow z}(\rho(x)) = \rho(z)$.

The following theorem follows from the PCP theorem [4, 5] combined with the parallel repetition theorem [48].

Theorem 4.3. For every $\ell = \ell(n)$ there exists a reduction from 3-SAT to Label Cover that given a 3-SAT instance φ of size n outputs in running time $n^{O(\ell)}$ a Label Cover instance \mathcal{L} of size $n^{O(\ell)}$ with range size $2^{O(\ell)}$ such that

- if φ is satisfiable then there exists an assignment to \mathcal{L} that satisfies all the constraints, and
- if φ is not satisfiable then every assignment to \mathcal{L} satisfies at most $2^{-\Omega(\ell)}$ fraction of the constraints.

We also need the following result of [46] which gives better parameters for sub-constant soundness error.

Theorem 4.4 ([46]). For every $\varepsilon = \varepsilon(n)$ there exists a reduction from 3-SAT to Label Cover that given a 3-SAT instance φ of size n outputs in running time $\text{poly}(n, \frac{1}{\varepsilon})$ a Label Cover instance \mathcal{L} of size $n^{1+o(1)} \cdot \frac{1}{\varepsilon^{O(1)}}$ with range size $\exp(\frac{1}{\varepsilon^{O(1)}})$ such that

- if φ is satisfiable then there exists an assignment to \mathcal{L} that satisfies all the constraints, and
- if φ is not satisfiable then every assignment to \mathcal{L} satisfies at most ε fraction of the constraints.

4.2 Sparsity of Low Rank Matrices

The following lemma, proved by Golovnev, Regev, and Weinstein [28], relates the sparsity of a matrix with nonzero entries on the diagonal to its rank.

Lemma 4.5 ([28]). For every field \mathbb{F} and an $n \times n$ matrix M over \mathbb{F} with nonzero entries on the diagonal, the number $s(M)$ of nonzero entries in M satisfies $s(M) \geq \frac{n^2}{4 \cdot \text{rank}_{\mathbb{F}}(M)}$.

4.3 Proof of Theorem 4.1

The proof of Theorem 4.1 uses a reduction of Bhangale [6], described below.

The reduction. The reduction gets as input an instance $\mathcal{L} = (U, V, E, R, L, \phi)$ of the Label Cover problem (see Definition 4.2) and produces a 4-uniform hypergraph H . Let t be an integer parameter to be determined later, and set $s = \lceil (R - t)/2 \rceil$.

The vertices of H include for every variable $x \in U$ a copy $C[x]$ of the vertex set of the Schrijver graph $S(R, s)$ (see Section 2.2). Formally, for every variable $x \in U$ we define $C[x]$ by

$$C[x] = \{x\} \times V(S(R, s)),$$

so that every vertex of $C[x]$ is referred to as a pair (x, A) , where A is a stable s -subset of $[R]$. The vertex set of H is $V(H) = \cup_{x \in U} C[x]$.

We next define the hyperedges of H . We identify here sets with their characteristic vectors, that is, for a set A we let A_α be 1 if $\alpha \in A$ and 0 otherwise. For every two variables $x, y \in U$ with a common neighbor $z \in V$ in \mathcal{L} , the set $\{(x, A), (x, B), (y, C), (y, D)\}$ forms a hyperedge in H if for every $\alpha, \beta \in [R]$ such that $\phi_{x \rightarrow z}(\alpha) = \phi_{y \rightarrow z}(\beta)$ it holds that $\{A_\alpha, B_\alpha, C_\beta, D_\beta\} = \{0, 1\}$. This completes the description of the reduction.

To provide some intuition, we note that the role of the component $C[x]$ in H is to encode the value of the variable x , where an assignment of $\alpha \in [R]$ to x is encoded by the 2-coloring of $C[x]$ in which every vertex (x, A) is colored by $A_\alpha \in \{0, 1\}$. This encoding is used in the following simple proof of the completeness of the reduction.

Proposition 4.6 (Completeness). *If the Label Cover instance \mathcal{L} is satisfiable then $\bar{\xi}(H) \leq 2$.*

Proof: Assume that there exists a satisfying assignment ρ to the variables of $U \cup V$ such that $\phi_{x \rightarrow z}(\rho(x)) = \rho(z)$ for every edge (x, z) of \mathcal{L} . Consider the 2-coloring of H that assigns to every vertex (x, A) the color $A_{\rho(x)}$. We claim that this coloring is proper. To see this, consider a hyperedge $\{(x, A), (x, B), (y, C), (y, D)\}$ of H defined with respect to a common neighbor $z \in V$ of x and y . Since ρ is a satisfying assignment, we have $\phi_{x \rightarrow z}(\rho(x)) = \rho(z)$ and $\phi_{y \rightarrow z}(\rho(y)) = \rho(z)$, and thus $\phi_{x \rightarrow z}(\rho(x)) = \phi_{y \rightarrow z}(\rho(y))$. By the definition of the hyperedges of H we conclude that $\{A_{\rho(x)}, B_{\rho(x)}, C_{\rho(y)}, D_{\rho(y)}\} = \{0, 1\}$, hence the hyperedge is not monochromatic. This implies that $\chi(H) \leq 2$ and thus $\bar{\xi}(H) \leq 2$, as required. ■

We turn to prove the soundness of the reduction.

Proposition 4.7 (Soundness). *If $\bar{\xi}(H) \leq t$ then there exists an assignment to the variables of $U \cup V$ that satisfies at least $\frac{1}{4 \cdot t^6}$ fraction of the edges of \mathcal{L} .*

Proof: Assume that $\bar{\xi}(H) \leq t$. Then there exists an assignment of a nonzero vector $u_v \in \mathbb{R}^t$ to every vertex $v \in V(H)$ such that every hyperedge of H contains two vertices assigned to orthogonal vectors. For every variable $x \in U$, the vertices of $C[x]$ can be viewed as the vertices of the Schrijver graph $S(R, s)$, which by Theorem 2.3 satisfies

$$\bar{\xi}(S(R, s)) = R - 2s + 2 = R - 2 \cdot \lceil (R - t)/2 \rceil + 2 \geq R - (R - t + 1) + 2 = t + 1.$$

This implies that for every variable $x \in U$ the assignment $(u_v)_{v \in C[x]}$ does not form an orthogonal representation of the graph $S(R, s)$, hence there exist two vertices $a_x = (x, A^{(x)})$ and $b_x = (x, B^{(x)})$ in $C[x]$ such that $A^{(x)} \cap B^{(x)} = \emptyset$ and $\langle u_{a_x}, u_{b_x} \rangle \neq 0$.

Applying Lemma 3.3 to the orthogonal representation $(u_v)_{v \in V(H)}$ of H with the pairwise disjoint pairs $\{(a_x, b_x) \mid x \in U\}$, we obtain a t^4 -dimensional orthogonal representation $(w_v)_{v \in V(H)}$ of H that satisfies the following properties:

1. for every $x \in U$, $\langle w_{a_x}, w_{b_x} \rangle \neq 0$, and
2. for every $x, y \in U$ such that $\langle w_{a_x}, w_{b_y} \rangle \neq 0$, the inner products $\langle w_{a_x}, w_{a_y} \rangle$, $\langle w_{b_x}, w_{a_y} \rangle$, and $\langle w_{b_x}, w_{b_y} \rangle$ are all nonzero.

We turn to show that there exists an assignment to the variables of $U \cup V$ that satisfies at least $\frac{1}{4 \cdot t^6}$ fraction of the edges of \mathcal{L} . To this end, we define a *randomized assignment* to the variables of $U \cup V$ as follows. For every $x \in U$ consider the set $E(x) = [R] \setminus (A^{(x)} \cup B^{(x)})$. Notice that the disjointness of the sets $A^{(x)}$ and $B^{(x)}$ implies that

$$|E(x)| = R - 2 \cdot s = R - 2 \cdot \lceil (R - t)/2 \rceil \leq t.$$

We assign to every variable $x \in U$ an assignment $\rho(x)$ chosen uniformly at random from $E(x)$. We further assign to every variable $z \in V$ an assignment $\rho(z)$ that maximizes the number of constraints involving z that can be satisfied, given that the assignment of every $x \in U$ is chosen from $E(x)$. More formally, for a variable $z \in V$ let $U_z \subseteq U$ be the set of neighbors of z in U and define $\rho(z)$ to be some $\beta \in [L]$ with largest number of variables $x \in U_z$ such that $\beta \in \phi_{x \rightarrow z}(E(x))$.

We need the following three claims.

Claim 4.8. *For every variable $z \in V$, there exists a variable $x \in U_z$ for which at least $\frac{1}{4 \cdot t^4}$ fraction of the variables y of U_z satisfies $\langle w_{a_x}, w_{b_y} \rangle \neq 0$.*

Claim 4.9. *For every variable $z \in V$ and two variables $x, y \in U_z$ such that $\langle w_{a_x}, w_{b_y} \rangle \neq 0$, we have $\phi_{x \rightarrow z}(E(x)) \cap \phi_{y \rightarrow z}(E(y)) \neq \emptyset$.*

Claim 4.10. *Let \mathcal{F} be a collection of ℓ -subsets of $[R]$ that includes a set that intersects every set of \mathcal{F} . Then there exists an element of $[R]$ that belongs to at least $\frac{1}{\ell}$ fraction of the sets of \mathcal{F} .*

Let us first show that these three claims complete the proof of the proposition. We claim that for every variable $z \in V$, the randomized assignment ρ satisfies in expectation at least $\frac{1}{4 \cdot t^6}$ fraction of the constraints involving z . Indeed, fix a variable $z \in V$. By Claim 4.8, there exists a variable $x \in U_z$ for which at least $\frac{1}{4 \cdot t^4}$ fraction of the variables y of U_z satisfies $\langle w_{a_x}, w_{b_y} \rangle \neq 0$. By Claim 4.9, this x satisfies $\phi_{x \rightarrow z}(E(x)) \cap \phi_{y \rightarrow z}(E(y)) \neq \emptyset$ for at least $\frac{1}{4 \cdot t^4}$ fraction of the variables y of U_z . Applying Claim 4.10 to these sets $\phi_{y \rightarrow z}(E(y))$, we obtain that there exists an element of $[R]$ that belongs to at least $\frac{1}{|E(x)|} \cdot \frac{1}{4 \cdot t^4} \geq \frac{1}{4 \cdot t^5}$ fraction of the sets $\phi_{y \rightarrow z}(E(y))$ with $y \in U_z$. By the definition of $\rho(z)$, we get that $\rho(z) \in \phi_{y \rightarrow z}(E(y))$ for at least $\frac{1}{4 \cdot t^5}$ fraction of the variables $y \in U_z$. Since $\rho(y)$ is chosen uniformly at random from $E(y)$, ρ satisfies in expectation at least $\frac{1}{t} \cdot \frac{1}{4 \cdot t^5} = \frac{1}{4 \cdot t^6}$ fraction of the constraints involving z , as required. Now, by linearity of expectation, ρ satisfies in expectation at least $\frac{1}{4 \cdot t^6}$ fraction of the constraints of \mathcal{L} , so in particular, there exists an assignment satisfying at least $\frac{1}{4 \cdot t^6}$ fraction of the constraints of \mathcal{L} , and we are done.

It remains to prove Claims 4.8, 4.9, and 4.10.

Proof of Claim 4.8: Fix a variable $z \in V$ and denote $n = |U_z|$. Let $M_1 \in \mathbb{R}^{n \times t^4}$ be the matrix whose rows are the vectors w_{a_x} for $x \in U_z$, and let $M_2 \in \mathbb{R}^{n \times t^4}$ be the matrix whose rows are the vectors w_{b_x} for $x \in U_z$. Consider the matrix $M \in \mathbb{R}^{n \times n}$ defined by $M = M_1 \cdot M_2^T$, and notice that $M_{x,y} = \langle w_{a_x}, w_{b_y} \rangle$ for every $x, y \in U_z$ and that its rank is at most t^4 . Property (1) of the orthogonal representation $(w_v)_{v \in V(H)}$ implies that the diagonal entries of M are all nonzero, so we can apply Lemma 4.5 to obtain that $s(M) \geq \frac{n^2}{4 \cdot t^4}$. In particular, there exists a row of M with at least $\frac{n}{4 \cdot t^4}$ nonzero entries. This implies that there exists a variable $x \in U_z$ for which at least $\frac{1}{4 \cdot t^4}$ fraction of the variables y of U_z satisfies $\langle w_{a_x}, w_{b_y} \rangle \neq 0$, as required. ■

Proof of Claim 4.9: Fix a variable $z \in V$ and two variables $x, y \in U_z$ such that $\langle w_{a_x}, w_{b_y} \rangle \neq 0$. Assume by contradiction that $\phi_{x \rightarrow z}(E(x)) \cap \phi_{y \rightarrow z}(E(y)) = \emptyset$.

We first show that the set $e = \{(x, A^{(x)}), (x, B^{(x)}), (y, A^{(y)}), (y, B^{(y)})\}$ is a hyperedge of H . To see this, suppose that some $\alpha, \beta \in [R]$ satisfy $\phi_{x \rightarrow z}(\alpha) = \phi_{y \rightarrow z}(\beta)$. By the assumption that $\phi_{x \rightarrow z}(E(x))$ and $\phi_{y \rightarrow z}(E(y))$ are disjoint, it follows that $\alpha \notin E(x)$ or $\beta \notin E(y)$, that is, $\alpha \in A^{(x)} \cup B^{(x)}$ or $\beta \in A^{(y)} \cup B^{(y)}$. In the former case, using the fact that $A^{(x)} \cap B^{(x)} = \emptyset$, we have $\{A_\alpha^{(x)}, B_\alpha^{(x)}\} = \{0, 1\}$ implying that e is a hyperedge. The latter case is handled similarly.

We now show that the vectors in $\{w_{a_x}, w_{b_x}, w_{a_y}, w_{b_y}\}$ are pairwise non-orthogonal. By Property (1) of the vectors in $(w_v)_{v \in V(H)}$ we have $\langle w_{a_x}, w_{b_x} \rangle \neq 0$ and $\langle w_{a_y}, w_{b_y} \rangle \neq 0$. By assumption we have $\langle w_{a_x}, w_{b_y} \rangle \neq 0$, and using Property (2) we get that $\langle w_{a_x}, w_{a_y} \rangle, \langle w_{b_x}, w_{a_y} \rangle$, and $\langle w_{b_x}, w_{b_y} \rangle$ are all nonzero as well. Since e is a hyperedge of H , this yields a contradiction to the fact that $(w_v)_{v \in V(H)}$ is an orthogonal representation of H , and we are done. ■

Proof of Claim 4.10: Assume that the set $A \in \mathcal{F}$ intersects every set of \mathcal{F} . Denote $A = \{a_1, \dots, a_\ell\}$ and put $\mathcal{A}_i = \{B \in \mathcal{F} \mid a_i \in B\}$ for every $i \in [\ell]$. Since $\mathcal{F} = \cup_{i \in [\ell]} \mathcal{A}_i$, it follows that

$$|\mathcal{F}| = |\cup_{i \in [\ell]} \mathcal{A}_i| \leq \sum_{i \in [\ell]} |\mathcal{A}_i|,$$

implying that there exists an $i \in [\ell]$ such that $|\mathcal{A}_i| \geq \frac{|\mathcal{F}|}{\ell}$. Therefore, some a_i belongs to at least $\frac{1}{\ell}$ fraction of the sets of \mathcal{F} , as required. ■

The proof of the proposition is completed. ■

To derive Theorem 4.1 we just have to set the parameters appropriately, as is done below.

Proof of Theorem 4.1, Item 1: Let ε and t be two parameters. By Theorem 4.4, there exists a reduction that maps a 3-SAT instance φ of size n in running time $\text{poly}(n, \frac{1}{\varepsilon})$ to a Label Cover instance \mathcal{L} of size $n^{1+o(1)} \cdot \frac{1}{\varepsilon^{O(1)}}$ with range size $R \leq \exp(\frac{1}{\varepsilon^{O(1)}})$ such that if φ is satisfiable then so is \mathcal{L} and if φ is not satisfiable then every assignment to \mathcal{L} satisfies at most ε fraction of the constraints. We proceed by mapping the instance \mathcal{L} to a 4-uniform hypergraph H using the reduction described above. By Lemma 2.4 and the fact that $R - 2s \leq t$, the number of vertices in $S(R, s)$ is at most $\binom{R}{R-2s} \leq R^t$. Hence, the number N of vertices in H satisfies

$$N \leq n^{1+o(1)} \cdot \frac{1}{\varepsilon^{O(1)}} \cdot R^t \leq n^{1+o(1)} \cdot \exp\left(\frac{1}{\varepsilon^{O(1)}}\right)^t.$$

By Proposition 4.6, if φ is satisfiable then $\bar{\xi}(H) \leq 2$. By Proposition 4.7, if $\bar{\xi}(H) \leq t$ then there exists an assignment to the variables of $U \cup V$ that satisfies at least $\frac{1}{4 \cdot t^6}$ fraction of the constraints of \mathcal{L} . In particular, for $\varepsilon < \frac{1}{4 \cdot t^6}$, if φ is not satisfiable then $\bar{\xi}(H) > t$.

Now, for a sufficiently small constant $\delta > 0$ set $t = \log^\delta n$ and, say, $\varepsilon = \frac{1}{\log^{7\delta} n}$, so that $\varepsilon < \frac{1}{4 \cdot t^6}$. For these parameters the running time of the reduction is polynomial in n and the hypergraph H has $N = n^{1+o(1)}$ vertices. The reduction implies that for some $\delta' > 0$ it is NP-hard to decide whether an input N -vertex 4-uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq \log^{\delta'} N$, completing the proof. ■

Proof of Theorem 4.1, Item 2: Let ℓ and t be two parameters. By Theorem 4.3, there exists a reduction that maps a 3-SAT instance φ of size n in running time $n^{O(\ell)}$ to a Label Cover instance \mathcal{L} of size $n^{O(\ell)}$ with range size $R \leq 2^{O(\ell)}$ such that if φ is satisfiable then so is \mathcal{L} and if φ is not satisfiable then every assignment to \mathcal{L} satisfies at most $2^{-\Omega(\ell)}$ fraction of the constraints. We proceed by mapping the instance \mathcal{L} to a 4-uniform hypergraph H using the reduction described above. Using Lemma 2.4, the number N of vertices in H satisfies

$$N \leq n^{O(\ell)} \cdot R^t \leq n^{O(\ell)} \cdot 2^{O(\ell \cdot t)}.$$

By Proposition 4.6, if φ is satisfiable then $\bar{\xi}(H) \leq 2$. By Proposition 4.7, if $\bar{\xi}(H) \leq t$ then there exists an assignment to the variables of $U \cup V$ that satisfies at least $\frac{1}{4 \cdot t^6}$ fraction of the constraints of \mathcal{L} . In particular, for $\ell = \Omega(\log t)$, if φ is not satisfiable then $\bar{\xi}(H) > t$.

Now, for an arbitrarily large constant $c' > 0$ put $t = c' \cdot \log n$ and define $\ell = c'' \cdot \log \log n$ for a sufficiently large $c'' > 0$. For these parameters the running time of the reduction is $n^{O(\ell)} = n^{O(\log \log n)}$ and the hypergraph H has $N = n^{O(\log \log n)}$ vertices. The reduction implies that assuming $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)})$, there is no polynomial time algorithm that decides whether an input N -vertex 4-uniform hypergraph H satisfies $\bar{\xi}(H) \leq 2$ or $\bar{\xi}(H) \geq c \cdot \frac{\log N}{\log \log N}$, where $c > 0$ can be arbitrarily large. This completes the proof. \blacksquare

Remark 4.11. A recent work of Alishahi and Meunier [1] extends Theorem 2.3 and shows that for every $d \geq 2s$ and for every field \mathbb{F} , it holds that $\bar{\xi}(S(d, s), \mathbb{F}) = d - 2s + 2$. This, combined with the above proof and with Corollary 3.6, implies that the hardness results given in Theorem 1.2 hold for the orthogonality dimension over every field \mathbb{F} . We omit the details.

5 The Orthogonality Dimension of Graphs

In this section we focus on the orthogonality dimension of graphs and prove Theorems 1.3 and 1.5.

5.1 Orthogonal Subspace Representations

We generalize the notion of orthogonal representations over the real field by assigning to every vertex a subspace of a given dimension, so that adjacent vertices are assigned to orthogonal subspaces.

Definition 5.1. A t -dimensional orthogonal k -subspace representation of a graph $G = (V, E)$ is an assignment of a subspace $U_v \subseteq \mathbb{R}^t$ with $\dim(U_v) = k$ to every vertex $v \in V$, such that the subspaces U_v and $U_{v'}$ are orthogonal whenever v and v' are adjacent in G . For a graph G , let $\bar{\xi}_k(G)$ denote the smallest integer t for which there exists a t -dimensional orthogonal k -subspace representation of G .

Clearly, $\bar{\xi}(G) = \bar{\xi}_1(G)$ for every graph G . It is also easy to see that the multichromatic numbers of graphs, defined in Section 1.1, bound the parameters $\bar{\xi}_k$ from above, namely, $\bar{\xi}_k(G) \leq \chi_k(G)$ for every G and k .

5.2 Hardness

In this section we prove Theorem 1.3, namely that for every graph F , it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G) \leq \bar{\xi}_3(F)$ or $\bar{\xi}(G) \geq \bar{\xi}_4(F)$. The proof employs the notion of lexicographic product of graphs, defined as follows.

Definition 5.2. The lexicographic product of the graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, denoted by $G_1 \bullet G_2$, is the graph whose vertex set is $V_1 \times V_2$ where two vertices (x_1, y_1) and (x_2, y_2) are adjacent if either $\{x_1, x_2\} \in E_1$ or $x_1 = x_2$ and $\{y_1, y_2\} \in E_2$.

One can view the graph $G_1 \bullet G_2$ as the graph obtained from G_1 by replacing every vertex by a copy of G_2 and replacing every edge by a complete bipartite graph between the vertex sets associated with its endpoints. We need the following property of the orthogonality dimension of lexicographic products of graphs.

Lemma 5.3. For every two graphs G_1 and G_2 , $\bar{\xi}(G_1 \bullet G_2) = \bar{\xi}_k(G_1)$ where $k = \bar{\xi}(G_2)$.

Proof: Let $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ be two graphs and denote $k = \bar{\xi}(G_2)$.

We first prove that $\bar{\xi}(G_1 \bullet G_2) \geq \bar{\xi}_k(G_1)$. Denote $t = \bar{\xi}(G_1 \bullet G_2)$, then there exists a t -dimensional orthogonal representation $(u_{(x,y)})_{(x,y) \in V_1 \times V_2}$ of $G_1 \bullet G_2$. For every $x \in V_1$, let U_x denote the subspace of \mathbb{R}^t spanned by all vectors $u_{(x,y)}$ with $y \in V_2$. By the definition of $G_1 \bullet G_2$, the subspaces U_x and $U_{x'}$ are orthogonal whenever x and x' are adjacent in G_1 . Further, for every $x \in V_1$, the restriction of the given orthogonal representation to the copy of G_2 associated with x forms an orthogonal representation of G_2 , so by $k = \bar{\xi}(G_2)$ it follows that $\dim(U_x) \geq k$. This implies that there exists a t -dimensional orthogonal k -subspace representation of G_1 , hence $\bar{\xi}_k(G_1) \leq t$, as required.

We next prove that $\bar{\xi}(G_1 \bullet G_2) \leq \bar{\xi}_k(G_1)$. Denote $t = \bar{\xi}_k(G_1)$, then there exists a t -dimensional orthogonal k -subspace representation $(U_x)_{x \in V_1}$ of G_1 . By $k = \bar{\xi}(G_2)$, there exists a k -dimensional orthogonal representation $(u_y)_{y \in V_2}$ of G_2 . For every $x \in V_1$, the fact that $\dim(U_x) = k$ implies that there exists an orthogonal linear transformation T_x from \mathbb{R}^k onto U_x . We assign to every vertex $(x, y) \in V_1 \times V_2$ of $G_1 \bullet G_2$ the nonzero vector $w_{(x,y)} = T_x(u_y) \in \mathbb{R}^t$. We claim that this is a t -dimensional orthogonal representation of $G_1 \bullet G_2$. To see this, let (x_1, y_1) and (x_2, y_2) be two adjacent vertices in $G_1 \bullet G_2$. If x_1 and x_2 are adjacent in G_1 then the subspaces U_{x_1} and U_{x_2} are orthogonal, hence the vectors $w_{(x_1,y_1)} \in U_{x_1}$ and $w_{(x_2,y_2)} \in U_{x_2}$ are orthogonal as well. Otherwise, $x_1 = x_2$ and the vertices y_1 and y_2 are adjacent in G_2 . This implies that the vectors u_{y_1} and u_{y_2} are orthogonal, and since T_{x_1} preserves inner products it follows that $T_{x_1}(u_{y_1})$ and $T_{x_1}(u_{y_2})$ are orthogonal, and we are done. ■

Equipped with Lemma 5.3, we are ready to prove Theorem 1.3.

Proof of Theorem 1.3: Fix a graph F . We reduce from the NP-hard problem of deciding whether an input graph G satisfies $\bar{\xi}(G) \leq 3$ [47]. The reduction maps an input graph G to the lexicographic product $G' = F \bullet G$ of F and G . The graph G' can clearly be constructed in polynomial time. The correctness of the reduction follows from Lemma 5.3. Indeed, we have $\bar{\xi}(G') = \bar{\xi}_k(F)$ for $k = \bar{\xi}(G)$, so if $\bar{\xi}(G) \leq 3$ then $\bar{\xi}(G') \leq \bar{\xi}_3(F)$ and if $\bar{\xi}(G) \geq 4$ then $\bar{\xi}(G') \geq \bar{\xi}_4(F)$. ■

5.3 Algorithm

Before presenting our algorithm, we provide some background on the vector chromatic number of graphs.

5.3.1 Vector Chromatic Number

Consider the following two relaxations of the chromatic number of graphs, due to Karger, Motwani, and Sudan [34].

Definition 5.4. For a graph $G = (V, E)$ the vector chromatic number of G , denoted by $\chi_v(G)$, is the minimal real value of $\kappa > 1$ for which there exists an assignment of a unit real vector u_v to every vertex $v \in V$ such that $\langle u_v, u_{v'} \rangle \leq -\frac{1}{\kappa-1}$ whenever v and v' are adjacent in G .

Definition 5.5. For a graph $G = (V, E)$ the strict vector chromatic number of G , denoted by $\chi_v^{(s)}(G)$, is the minimal real value of $\kappa > 1$ for which there exists an assignment of a unit real vector u_v to every vertex $v \in V$ such that $\langle u_v, u_{v'} \rangle = -\frac{1}{\kappa-1}$ whenever v and v' are adjacent in G .

It is well known and easy to verify that for every graph G , $\chi_v(G) \leq \chi_v^{(s)}(G) \leq \chi(G)$. Karger et al. [34] have obtained the following algorithmic result.

Theorem 5.6 ([34]). *There exists a randomized polynomial time algorithm that given an n -vertex graph G with maximum degree at most Δ and $\chi_v(G) \leq \kappa$ for some $\kappa \geq 2$, finds an independent set of size $\tilde{\Omega}(\frac{n}{\Delta^{1-2/\kappa}})$.*

Note that the well-known Lovász ϑ -function, introduced in [43], is known to satisfy $\vartheta(G) = \chi_v^{(s)}(\overline{G})$ for every graph G [34]. Combining this with a result of [43], it follows that the orthogonality dimension forms an upper bound on the strict vector chromatic number, that is, $\overline{\xi}(G) \geq \chi_v^{(s)}(G)$ for every graph G . This was recently generalized by Bukh and Cox as follows (see [12, Proposition 23]).

Lemma 5.7 ([12]). *For every graph G and an integer k , $\overline{\xi}_k(G) \geq k \cdot \chi_v^{(s)}(G)$.*

We derive that the graph parameters $\overline{\xi}_k$ and χ_k coincide on Kneser graphs $K(d, s)$ whenever k is divisible by s .

Corollary 5.8. *For all integers $\ell \geq 1$ and $d \geq 2s$, $\overline{\xi}_{\ell \cdot s}(K(d, s)) = \ell \cdot d$.*

Proof: For the upper bound on $\overline{\xi}_{\ell \cdot s}(K(d, s))$, recall that Conjecture 1.1 was confirmed for $k = \ell \cdot s$ in [51], hence $\overline{\xi}_{\ell \cdot s}(K(d, s)) \leq \chi_{\ell \cdot s}(K(d, s)) = \ell \cdot d$. For the lower bound, combine Lemma 5.7 with the fact that $\chi_v^{(s)}(K(d, s)) = \frac{d}{s}$ (see [43]), to get that $\overline{\xi}_{\ell \cdot s}(K(d, s)) \geq \ell \cdot d$. ■

5.3.2 The Algorithm

We present an efficient algorithm that given a graph G that satisfies $\overline{\xi}(G) \leq 3$ (or, more generally, $\overline{\xi}_k(G) \leq 3k$ for some k), finds a coloring of G with relatively few colors. We use the following simple claim of Blum [7] which reduces the algorithmic task of coloring a graph to the algorithmic task of finding a large independent set in it.

Claim 5.9 ([7]). *Let \mathcal{G} be a graph family which is closed under taking induced subgraphs, let $c_1, c_2 > 1$ be arbitrary fixed constants, and let $f : \mathbb{N} \rightarrow \mathbb{N}$ be any non-decreasing function satisfying $c_1 \cdot f(n) \leq f(2n) \leq c_2 \cdot f(n)$ for all sufficiently large n . Then if there exists a (randomized) polynomial time algorithm which finds an independent set of size $f(n)$ in any n -vertex graph $G \in \mathcal{G}$, then there exists a (randomized) polynomial time algorithm which finds an $O(\frac{n}{f(n)})$ -coloring of any n -vertex graph $G \in \mathcal{G}$.*

We need the following two simple lemmas.

Lemma 5.10. *Let $G = (V, E)$ be a graph such that $\bar{\xi}_k(G) \leq 2k$ for some integer k . Then G is 2-colorable.*

Proof: Let $G = (V, E)$ be a graph satisfying $\bar{\xi}_k(G) \leq 2k$, and let $(U_v)_{v \in V}$ be a $2k$ -dimensional orthogonal k -subspace representation of G . It suffices to prove that every connected component of G is 2-colorable. Fix a vertex v in some connected component of G , and observe that there exists a unique subspace of \mathbb{R}^{2k} of dimension k orthogonal to U_v . This implies that the orthogonal subspace representation of G provides a 2-coloring of the connected component of v , where the vertices of even distance from v are assigned to U_v and the vertices of odd distance from v are assigned to its orthogonal complement U_v^\perp , so we are done. ■

For a graph G and a vertex v , let $N(v)$ denote the neighborhood of v in G and let $G[N(v)]$ denote the subgraph of G induced by $N(v)$.

Lemma 5.11. *Let $G = (V, E)$ be a graph such that $\bar{\xi}_k(G) \leq 3k$ for some integer k . Then for every vertex $v \in V$ the subgraph $G[N(v)]$ is 2-colorable.*

Proof: Let $G = (V, E)$ be a graph satisfying $\bar{\xi}_k(G) \leq 3k$, and let $(U_v)_{v \in V}$ be a $3k$ -dimensional orthogonal k -subspace representation of G . Let $v \in V$ be a vertex in G . The subspace U_v is orthogonal to all subspaces $U_{v'}$ with $v' \in N(v)$. By $\dim(U_v) = k$, the orthogonal complement of U_v in \mathbb{R}^{3k} has dimension $2k$. It follows that $G[N(v)]$ admits a $2k$ -dimensional orthogonal k -subspace representation, hence by Lemma 5.10 it is 2-colorable, as required. ■

We are ready to prove the following result.

Theorem 5.12. *There exists a randomized polynomial time algorithm that given an n -vertex graph G that satisfies $\bar{\xi}_k(G) \leq 3k$ for some k , finds a coloring of G that uses at most $\tilde{O}(n^{1/4})$ colors. In particular, the algorithm finds an orthogonal representation of G of dimension $\tilde{O}(n^{1/4})$.*

Proof: By Claim 5.9 it suffices to show that there exists a randomized polynomial time algorithm that given an n -vertex graph $G = (V, E)$ with $\bar{\xi}_k(G) \leq 3k$ for some k , finds in G an independent set of size $\tilde{\Omega}(n^{3/4})$. We consider two possible cases. Suppose first that there exists a vertex $v \in V$ in G whose degree is at least $\Delta = n^{3/4}$. Then by Lemma 5.11 the subgraph $G[N(v)]$ is 2-colorable, so we can find an independent set in G of size at least $\frac{\Delta}{2}$ by finding in polynomial time a 2-coloring of $G[N(v)]$ and taking a largest color class. Otherwise, the maximum degree of G is at most Δ . By Lemma 5.7 we have

$$\chi_v(G) \leq \chi_v^{(s)}(G) \leq \frac{\bar{\xi}_k(G)}{k} \leq 3,$$

so by Theorem 5.6 we can find in polynomial time an independent set of size $\tilde{\Omega}(\frac{n}{\Delta^{1/3}}) \geq \tilde{\Omega}(n^{3/4})$, and we are done. ■

To improve the number of used colors, we employ the following result that stems from the analysis by Arora et al. [3] of the semidefinite relaxation of [34]. (For an explicit statement, see [17, Lemma 4.12], applied with $\sigma = 0.5$, $c \approx 0.04843726$, and $\delta \approx 0.7587$.)

Theorem 5.13 ([3]). *There exists a randomized polynomial time algorithm that given an n -vertex graph G with maximum degree at most $\Delta = n^{0.7587}$ and $\chi_v(G) \leq 3$, finds an independent set in G of size at least $\tilde{\Omega}(n \cdot \Delta^{-0.3179}) \geq \tilde{\Omega}(n^{0.7587})$.*

By applying Theorem 5.13 instead of Theorem 5.6 in the proof of Theorem 5.12, we obtain the following slight improvement, confirming Theorem 1.5.

Theorem 5.14. *There exists a randomized polynomial time algorithm that given an n -vertex graph G that satisfies $\tilde{\zeta}_k(G) \leq 3k$ for some k , finds a coloring of G that uses at most $\tilde{O}(n^{0.2413})$ colors.*

Acknowledgements

We are grateful to Alexander Golovnev for useful discussions and to the anonymous reviewers for their valuable suggestions.

References

- [1] M. Alishahi and F. Meunier. Topological bounds for graph representations over any field. *arXiv*, 2019.
- [2] N. Alon, P. Kelsen, S. Mahajan, and R. Hariharan. Approximate hypergraph coloring. *Nord. J. Comput.*, 3(4):425–439, 1996. Preliminary version in SWAT’96.
- [3] S. Arora, E. Chlamtac, and M. Charikar. New approximation guarantee for chromatic number. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC’06)*, pages 215–224, 2006.
- [4] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998. Preliminary version in FOCS’92.
- [5] S. Arora and S. Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998. Preliminary version in FOCS’92.
- [6] A. Bhangale. NP-hardness of coloring 2-colorable hypergraph with poly-logarithmically many colors. In *Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP’18)*, pages 15:1–15:11, 2018.
- [7] A. Blum. New approximation algorithms for graph coloring. *J. ACM*, 41(3):470–516, 1994.
- [8] A. Blum and D. R. Karger. An $O(n^{3/14})$ -coloring algorithm for 3-colorable graphs. *Inf. Process. Lett.*, 61(1):49–53, 1997.
- [9] J. Brakensiek and V. Guruswami. New hardness results for graph and hypergraph colorings. In *Proceedings of the 31st Conference on Computational Complexity (CCC’16)*, pages 14:1–14:27, 2016.

- [10] J. Briët, H. Buhrman, D. Leung, T. Piovosan, and F. Speelman. Round elimination in exact communication complexity. In *Proceedings of the 10th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC'15)*, volume 44 of *LIPICs*, pages 206–225, 2015.
- [11] J. Briët and J. Zuiddam. On the orthogonal rank of Cayley graphs and impossibility of quantum round elimination. *Quantum Information & Computation*, 17(1&2):106–116, 2017.
- [12] B. Bukh and C. Cox. On a fractional version of Haemers’ bound. *IEEE Trans. Inform. Theory*, 65(6):3340–3348, 2019.
- [13] J. Bulín, A. A. Krokhin, and J. Opršal. Algebraic approach to promise constraint satisfaction. In *Proceedings of the 51st Annual ACM Symposium on Theory of Computing (STOC'19)*, pages 602–613, 2019.
- [14] P. J. Cameron, A. Montanaro, M. W. Newman, S. Severini, and A. J. Winter. On the quantum chromatic number of a graph. *Electr. J. Comb.*, 14(1), 2007.
- [15] H. Chen and A. M. Frieze. Coloring bipartite hypergraphs. In *Proceedings of the 5th International Conference on Integer Programming and Combinatorial Optimization (IPCO'96)*, pages 345–358, 1996.
- [16] E. Chlamtac. Approximation algorithms using hierarchies of semidefinite programming relaxations. In *Proceedings of the 48th Symposium on Foundations of Computer Science (FOCS'07)*, pages 691–701, 2007.
- [17] E. Chlamtáč and I. Haviv. Linear index coding via semidefinite programming. *Combinatorics, Probability & Computing*, 23(2):223–247, 2014. Preliminary version in SODA'12.
- [18] R. de Wolf. Quantum Computing and Communication Complexity. PhD thesis, Universiteit van Amsterdam, 2001.
- [19] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM Journal on Computing*, 39(3):843–873, 2009. Preliminary version in STOC'06.
- [20] I. Dinur, O. Regev, and C. D. Smyth. The hardness of 3-uniform hypergraph coloring. *Combinatorica*, 25(5):519–535, 2005. Preliminary version in FOCS'02.
- [21] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1):439–485, 2005. Preliminary version in STOC'02.
- [22] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935.
- [23] P. Frankl and Z. Füredi. Extremal problems concerning Kneser graphs. *J. Comb. Theory, Ser. B*, 40(3):270–284, 1986.
- [24] P. Frankl and V. Rödl. Forbidden intersections. *Trans. Amer. Math. Soc.*, 300(1):259–286, 1987.
- [25] M. R. Garey and D. S. Johnson. The complexity of near-optimal graph coloring. *J. ACM*, 23(1):43–49, 1976.

- [26] C. D. Godsil and J. Zaks. Colouring the sphere. *University of Waterloo research report*, CORR 88-12, 1988.
- [27] A. Golovnev and I. Haviv. The (generalized) orthogonality dimension of (generalized) Kneser graphs: Bounds and applications. *CoRR*, abs/2002.08580, 2020.
- [28] A. Golovnev, O. Regev, and O. Weinstein. The minrank of random graphs. *IEEE Trans. Inform. Theory*, 64(11):6990–6995, 2018. Preliminary version in RANDOM’17.
- [29] V. Guruswami, J. Håstad, and M. Sudan. Hardness of approximate hypergraph coloring. *SIAM J. Comput.*, 31(6):1663–1686, 2002. Preliminary version in FOCS’00.
- [30] V. Guruswami and S. Khanna. On the hardness of 4-coloring a 3-colorable graph. *SIAM J. Discrete Math.*, 18(1):30–40, 2004. Preliminary version in CCC’00.
- [31] W. H. Haemers. An upper bound for the Shannon capacity of a graph. In *Algebraic methods in graph theory, Vol. I, II (Szeged, 1978)*, volume 25 of *Colloq. Math. Soc. János Bolyai*, pages 267–272. North-Holland, Amsterdam, 1981.
- [32] E. Halperin, R. Nathaniel, and U. Zwick. Coloring k -colorable graphs using relatively small palettes. *J. Algorithms*, 45(1):72–90, 2002. Preliminary version in SODA’01.
- [33] I. Haviv. Topological bounds on the dimension of orthogonal representations of graphs. *Eur. J. Comb.*, 81:84–97, 2019.
- [34] D. R. Karger, R. Motwani, and M. Sudan. Approximate graph coloring by semidefinite programming. *J. ACM*, 45(2):246–265, 1998. Preliminary version in FOCS’94.
- [35] R. M. Karp. Reducibility among combinatorial problems. In *Proceedings of a Symposium on the Complexity of Computer Computations*, pages 85–103, 1972.
- [36] K. Kawarabayashi and M. Thorup. Coloring 3-colorable graphs with less than $n^{1/5}$ colors. *J. ACM*, 64(1):4:1–4:23, 2017. Preliminary versions in FOCS’12 and STACS’14.
- [37] S. Khanna, N. Linial, and S. Safra. On the hardness of approximating the chromatic number. *Combinatorica*, 20(3):393–415, 2000. Preliminary version in ISTCS’93.
- [38] M. Kneser. Aufgabe 360. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 58(2):27, 1955.
- [39] S. Kochen and E. P. Specker. The problem of hidden variables in quantum mechanics. *Journal of Mathematics and Mechanics*, 17(1):59–87, 1967.
- [40] M. Krivelevich, R. Nathaniel, and B. Sudakov. Approximating coloring and maximum independent sets in 3-uniform hypergraphs. *J. Algorithms*, 41(1):99–113, 2001. Preliminary version in SODA’01.
- [41] L. Lovász. Coverings and colorings of hypergraphs. In *Proceedings of the 4th Southeastern Conf. on Comb.*, pages 3–12. Utilitas Math., 1973.
- [42] L. Lovász. Kneser’s conjecture, chromatic number, and homotopy. *J. Comb. Theory, Ser. A*, 25(3):319–324, 1978.

- [43] L. Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(1):1–7, 1979.
- [44] L. Lovász. *Graphs and Geometry*, volume 65. Colloquium Publications, 2019.
- [45] L. Lovász, M. Saks, and A. Schrijver. Orthogonal representations and connectivity of graphs. *Linear Algebra and its Applications*, 114/115:439–454, 1989. Special Issue Dedicated to Alan J. Hoffman.
- [46] D. Moshkovitz and R. Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5):29:1–29:29, 2010. Preliminary version in FOCS’08.
- [47] R. Peeters. Orthogonal representations over finite fields and the chromatic number of graphs. *Combinatorica*, 16(3):417–431, 1996.
- [48] R. Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. Preliminary version in STOC’95.
- [49] G. Scarpa and S. Severini. Kochen-Specker sets and the rank-1 quantum chromatic number. *IEEE Trans. Inform. Theory*, 58(4):2524–2529, 2012.
- [50] A. Schrijver. Vertex-critical subgraphs of Kneser graphs. *Nieuw Arch. Wiskd.*, 26(3):454–461, 1978.
- [51] S. Stahl. n -tuple colorings and associated graphs. *J. Comb. Theory, Ser. B*, 20(2):185–203, 1976.
- [52] S. Stahl. The multichromatic numbers of some Kneser graphs. *Discrete Mathematics*, 185(1-3):287–291, 1998.
- [53] A. Wigderson. Improving the performance guarantee for approximate graph coloring. *J. ACM*, 30(4):729–735, 1983. Preliminary version in STOC’82.
- [54] M. Wrochna and S. Zivný. Improved hardness for H -colourings of G -colourable graphs. In *Proceedings of the 31st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’20)*, pages 1426–1435, 2020.