

The (Generalized) Orthogonality Dimension of (Generalized) Kneser Graphs: Bounds and Applications

Alexander Golovnev* Ishay Haviv†

Abstract

The *orthogonality dimension* of a graph $G = (V, E)$ over a field \mathbb{F} is the smallest integer t for which there exists 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 $\langle u_v, u_{v'} \rangle = 0$ whenever v and v' are adjacent vertices in G . The study of the orthogonality dimension of graphs is motivated by various application in information theory and in theoretical computer science. The contribution of the present work is two-folded.

First, we prove that there exists a constant c such that for every sufficiently large integer t , it is NP-hard to decide whether the orthogonality dimension of an input graph over \mathbb{R} is at most t or at least $3t/2 - c$. At the heart of the proof lies a geometric result, which might be of independent interest, on a *generalization* of the orthogonality dimension parameter for the family of *Kneser graphs*, analogously to a long-standing conjecture of Stahl (J. Comb. Theo. Ser. B, 1976).

Second, we study the smallest possible orthogonality dimension over finite fields of the complement of graphs that do not contain certain fixed subgraphs. In particular, we provide an explicit construction of triangle-free n -vertex graphs whose complement has orthogonality dimension over the binary field at most $n^{1-\delta}$ for some constant $\delta > 0$. Our results involve constructions from the family of *generalized Kneser graphs* and they are motivated by the rigidity approach to circuit lower bounds. We use them to answer a couple of questions raised by Codenotti, Pudlák, and Resta (Theor. Comput. Sci., 2000), and in particular, to disprove their Odd Alternating Cycle Conjecture over every finite field.

1 Introduction

A t -dimensional *orthogonal representation* of a graph $G = (V, E)$ 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 $\langle u_v, u_{v'} \rangle = 0$ whenever v and v' are adjacent vertices in G . The *orthogonality dimension* of a graph G over \mathbb{F} , denoted by $\bar{\zeta}(G, \mathbb{F})$, is the smallest integer t for which there exists a t -dimensional orthogonal representation of G over \mathbb{F} .¹ The orthogonality dimension parameter is closely related to several other well-studied graph parameters. In particular, for every graph G and every field \mathbb{F} , $\bar{\zeta}(G, \mathbb{F})$ is sandwiched between the clique number and the chromatic number of G , that is, $\omega(G) \leq \bar{\zeta}(G, \mathbb{F}) \leq \chi(G)$.

*School of Engineering and Applied Sciences, Harvard University, Cambridge, Massachusetts, USA. Research supported by a Rabin Postdoctoral Fellowship.

†School of Computer Science, The Academic College of Tel Aviv-Yaffo, Tel Aviv 61083, Israel.

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. We have decided to use here the other definition, but one may view the notation $\bar{\zeta}(G, \mathbb{F})$ as standing for $\zeta(\bar{G}, \mathbb{F})$.

Orthogonal representations of graphs have been found useful over the years for various applications in information theory and in theoretical computer science. They were originally introduced over the real field in a seminal work of Lovász [31], where it was used to define the influential Lovász ϑ -function. The latter was used in [31] to determine the Shannon capacity, a notoriously difficult information-theoretic graph parameter, of the cycle on five vertices, and in the last decades it was successfully applied in algorithmic and combinatorial results (see, e.g., [27, 17, 3]). The orthogonality dimension of graphs plays an important role in several areas of computational complexity. Over finite fields, the orthogonality dimension and its extension due to Haemers [20] to a graph parameter called *minrank* have attracted a significant attention in circuit complexity, and more specifically, in the study of Valiant’s rigidity approach to circuit lower bounds [41] (see, e.g., [11, 35, 19]). Over the complex field, the orthogonality dimension was used in a characterization of the quantum communication complexity of promise equality problems [12, 4, 5] and in the study of the quantum chromatic number [8, 37]. The orthogonality dimension parameter was also investigated in the contexts of hardness of approximation [34, 28], integrality gaps for linear programming [25, 24], and algorithms based on semi-definite programming [9, 22].

The present work studies two aspects of the orthogonality dimension of graphs. First, we prove an NP-hardness result for approximating the orthogonality dimension of graphs over the real field \mathbb{R} . At the heart of the proof lies a geometric result, which might be of independent interest, on a *generalization* of the orthogonality dimension parameter for the family of *Kneser graphs*, analogously to a long-standing graph-theoretic conjecture due to Stahl [38]. The second aspect of the orthogonality dimension parameter considered in this work, motivated by the area of circuit complexity, is that of determining the smallest possible orthogonality dimension over finite fields of the complement of graphs that do not contain certain fixed subgraphs. In this context, we prove a new bound on the minrank parameter over finite fields for the family of *generalized Kneser graphs*. The bound is used to settle a couple of questions raised by Codenotti, Pudlák, and Resta in [11] and to disprove their Odd Alternating Cycle Conjecture over every finite field.

1.1 Our Contribution

1.1.1 The Generalized Orthogonality Dimension of Kneser Graphs

We start by considering the computational hardness of determining the orthogonality dimension of graphs over the real field \mathbb{R} . The challenge of understanding the hardness of this parameter was posed already in the late eighties by Lovász, Saks, and Schrijver [33] (see also [32]), and yet, the problem is far from being well-understood. It is easy to see that deciding whether an input graph G satisfies $\bar{\zeta}(G, \mathbb{R}) \leq t$ can be solved in polynomial running-time for $t \in \{1, 2\}$, and Peeters [34] has shown that it is NP-hard for $t \geq 3$. His result is known to imply that for every $t \geq 6$ it is NP-hard to decide whether an input graph G satisfies $\bar{\zeta}(G, \mathbb{R}) \leq t$ or $\bar{\zeta}(G, \mathbb{R}) \geq \lceil 4t/3 \rceil$ (see [22]). In the current work, we improve on the $4/3$ multiplicative gap and prove the following.

Theorem 1.1. *There exists a constant c such that for every sufficiently large integer t , it is NP-hard to decide whether an input graph G satisfies $\bar{\zeta}(G, \mathbb{R}) \leq t$ or $\bar{\zeta}(G, \mathbb{R}) \geq 3t/2 - c$.*

It is worth noting that in order to obtain hardness results for the orthogonality dimension parameter, it is natural to employ known hardness results regarding the closely related chromatic

number of graphs. Indeed, it is easy to verify (see, e.g., [22]) that every graph G satisfies

$$\log_3 \chi(G) \leq \bar{\xi}(G, \mathbb{R}) \leq \chi(G),$$

hence hardness of deciding whether an input graph G satisfies $\chi(G) \leq t_1$ or $\chi(G) \geq t_2$ immediately implies the hardness of deciding whether it satisfies $\bar{\xi}(G, \mathbb{R}) \leq t_1$ or $\bar{\xi}(G, \mathbb{R}) \geq \log_3 t_2$. In particular, a result of Dinur, Mossel, and Regev [14] on the hardness of the chromatic number implies that assuming a certain variant of the unique games conjecture, deciding whether a given graph G satisfies $\bar{\xi}(G, \mathbb{R}) \leq 3$ or $\bar{\xi}(G, \mathbb{R}) \geq t$ is NP-hard for every $t \geq 4$. However, if one is interested in standard NP-hardness for the orthogonality dimension, the state of the art for the hardness of the chromatic number does not seem to imply any hardness results, despite some remarkable recent progress [7, 42]. Moreover, most hardness proofs for the chromatic number crucially use the fact that an upper bound on the independence number of a graph implies a strong lower bound on its chromatic number (namely, $\chi(G) \geq \frac{|V(G)|}{\alpha(G)}$), whereas an analogue of such a statement for the orthogonality dimension does not hold in general (see, e.g., [22, Proposition 2.2]).

One technique for proving hardness results for the chromatic number that can be applied for the orthogonality dimension is that of Garey and Johnson [18], who have related hardness of graph coloring to the *multichromatic numbers of Kneser graphs*. The k th multichromatic number of a graph G , denoted by $\chi_k(G)$, is the smallest number of colors needed in order to assign to every vertex of G a set of k colors so that adjacent vertices are assigned to disjoint sets. Notice that $\chi_1(G)$ is simply the standard chromatic number $\chi(G)$. The family of Kneser graphs is defined as follows.

Definition 1.2 (Kneser Graphs). *For integers $d \geq 2s$, the Kneser graph $K(d, s)$ is the graph whose vertices are all the s -subsets of $[d] = \{1, \dots, d\}$, where two sets are adjacent if they are disjoint.*

Note that the multichromatic numbers can be defined in terms of Kneser graphs, namely, $\chi_k(G)$ is the smallest integer d for which there exists a homomorphism from G to $K(d, k)$.

In the seventies, Stahl [38] has made the following conjecture.

Conjecture 1.3 (Stahl's Conjecture [38]). *For all integers k and $d \geq 2s$,*

$$\chi_k(K(d, s)) = \binom{k}{s} \cdot (d - 2s) + 2k.$$

Stahl's conjecture has received a significant attention in the literature over the years. Even very recently, it was related to the well-known recently disproved Hedetniemi's conjecture [40]. Nevertheless, more than forty years since it was proposed, Stahl's conjecture is still open. It is known that the right-hand side in Conjecture 1.3 forms an upper bound on $\chi_k(K(d, s))$, and that this bound is tight up to an additive constant that depends solely on s [10, 39]. The precise statement of the conjecture was confirmed only for a few special cases. This includes the case of $k = 1$ proved by Lovász [30], the cases of $s \leq 2, k \leq s, d = 2s + 1$, and k divisible by s proved by Stahl [38, 39], and the case of $s = 3$ and $k = 4$ proved by Garey and Johnson [18] (extended to $s = 3$ with any k in [39]). The result of [18] was combined there with a simple reduction to show that for every $t \geq 6$, it is NP-hard to decide whether a given graph G satisfies $\chi(G) \leq t$ or $\chi(G) \geq 2t - 4$.

The recent work [22] has suggested to borrow the reduction of [18] to prove hardness results for the orthogonality dimension parameter. This approach requires the following generalization of orthogonal representations of graphs over the reals.

Definition 1.4 (Orthogonal Subspace Representation). *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 $\bar{\zeta}_k(G, \mathbb{R})$ denote the smallest integer t for which there exists a t -dimensional orthogonal k -subspace representation of G .*

Note that for $k = 1$, Definition 1.4 coincides with the orthogonality dimension over the reals, and that for every graph G and every k it holds that $\bar{\zeta}_k(G, \mathbb{R}) \leq \chi_k(G)$.

Combining the hardness result of Peeters [34] with the reduction of [18] implies the following.

Proposition 1.5 ([22]). *For every graph F , it is NP-hard to decide whether an input graph G satisfies $\bar{\zeta}(G, \mathbb{R}) \leq \bar{\zeta}_3(F, \mathbb{R})$ or $\bar{\zeta}(G, \mathbb{R}) \geq \bar{\zeta}_4(F, \mathbb{R})$.*

With Proposition 1.5 in hand, it is of interest to find graphs F with a large gap between $\bar{\zeta}_3(F, \mathbb{R})$ and $\bar{\zeta}_4(F, \mathbb{R})$. In light of Conjecture 1.3, it is natural to consider the generalized orthogonality dimension parameters for the family of Kneser graphs. For $k = 1$, it was shown in [24] that the standard chromatic number and the standard orthogonality dimension over \mathbb{R} coincide on all Kneser graphs. In addition, a result of Bukh and Cox [6] implies that for every $d \geq 2s$ and every k , $\bar{\zeta}_k(K(d, s), \mathbb{R}) \geq kd/s$. This implies that the k th chromatic number and the k th orthogonality dimension over \mathbb{R} coincide on $K(d, s)$ whenever k is divisible by s .

In this work we initiate a systematic study of the generalized orthogonality dimension parameters of Kneser graphs, analogously to Conjecture 1.3. Let us already mention that the arguments applied in the study of Stahl's conjecture do not seem to extend to our question. The main reason is that the proofs in [38, 18, 10, 39] use Hilton-Milner-type theorems to characterize the possible structures of the independent sets induced by generalized colorings of Kneser graphs, whereas in our setting, orthogonal subspace representations do not naturally induce large independent sets and the problem seems to require a more geometric approach.

The first non-trivial case is that of Kneser graphs $K(d, s)$ with $s = 2$, for which we show that the generalized orthogonality dimension parameters are equal to the multichromatic numbers.

Theorem 1.6. *For all integers $k \geq 1$ and $d \geq 4$,*

$$\bar{\zeta}_k(K(d, 2), \mathbb{R}) = \left\lceil \frac{k}{2} \right\rceil \cdot (d - 4) + 2k.$$

We proceed by considering a general $s \geq 3$ and prove the following lower bound.

Theorem 1.7. *For every integers $k \geq s \geq 3$ there exists $c = c(s, k)$ such that for all integers $d \geq 2s$,*

$$\bar{\zeta}_k(K(d, s), \mathbb{R}) \geq \frac{k - \left\lceil \frac{k+1}{s} \right\rceil + 1}{s - 1} \cdot d - c.$$

Note that for $k = \ell \cdot s - 1$ where ℓ is an integer, the bound provided by Theorem 1.7 is tight up to the additive constant c . Indeed, in this case we get that there exists a constant c such that for all integers $d \geq 2s$ it holds that

$$\ell \cdot d - c \leq \bar{\zeta}_{\ell \cdot s - 1}(K(d, s), \mathbb{R}) \leq \chi_{\ell \cdot s - 1}(K(d, s)) \leq \ell \cdot d - 2.$$

Note further that for the special case of $k = 4$ and $s = 3$, Theorem 1.7 implies that there exists a constant c such that $\bar{\zeta}_4(K(d, 3), \mathbb{R}) \geq 3d/2 - c$ for every sufficiently large integer d . This, combined with Proposition 1.5 and the fact that $\bar{\zeta}_3(K(d, 3), \mathbb{R}) = d$, yields our hardness result Theorem 1.1.

It will be interesting to figure out if the bounds given in Theorem 1.7 can be tightened to the quantity given in the right-hand side of Conjecture 1.3, at least up to an additive term independent of d . In particular, it will be nice to decide whether for all integers $d \geq 6$ it holds that $\bar{\zeta}_4(K(d, 3), \mathbb{R}) = 2d - 4$. A positive answer would imply that for every $t \geq 6$, it is NP-hard to decide whether an input graph G satisfies $\bar{\zeta}(G) \leq t$ or $\bar{\zeta}(G) \geq 2t - 4$. We remark, however, that the approach suggested by Proposition 1.5 for the hardness of the orthogonality dimension cannot yield a multiplicative hardness gap larger than 2, as it is easy to see that every graph F satisfies $\bar{\zeta}_4(F, \mathbb{R}) \leq \bar{\zeta}(F, \mathbb{R}) + \bar{\zeta}_3(F, \mathbb{R}) \leq 2 \cdot \bar{\zeta}_3(F, \mathbb{R})$.

1.1.2 The Orthogonality Dimension of Generalized Kneser Graphs

We next consider the orthogonality dimension over finite fields of the complement of graphs that do not contain some fixed subgraphs. In fact, in this context we consider an extension of the orthogonality dimension parameter, called minrank, that was introduced by Haemers in [20] and is defined as follows.

Definition 1.8 (Minrank). *Let $G = (V, E)$ be a directed graph on the vertex set $V = [n]$ and let \mathbb{F} be a field. We say that an n by n matrix M over \mathbb{F} represents G if $M_{i,i} \neq 0$ for every $i \in V$, and $M_{i,j} = 0$ for every distinct $i, j \in V$ such that $(i, j) \notin E$. The minrank of G over \mathbb{F} is defined as*

$$\text{minrk}_{\mathbb{F}}(G) = \min\{\text{rank}_{\mathbb{F}}(M) \mid M \text{ represents } G \text{ over } \mathbb{F}\}.$$

The definition is naturally extended to (undirected) graphs by replacing every undirected edge with two oppositely directed edges.

Note that for every graph G and every field \mathbb{F} , $\text{minrk}_{\mathbb{F}}(\bar{G}) \leq \bar{\zeta}(G, \mathbb{F})$.²

We consider here the question of whether there are graphs with no short odd cycles and yet low minrank over finite fields. This question is motivated by the area of circuit complexity, and more specifically, by Valiant's approach to circuit lower bounds [41], as described next. The *rigidity* of an n by n matrix M over a field \mathbb{F} with respect to a given parameter r is the smallest number of entries that one has to change in M in order to reduce its rank over \mathbb{F} to below r . Roughly speaking, it was shown in [41] that n by n matrices with large rigidity for $r = \varepsilon \cdot n$ where $\varepsilon > 0$ is a constant can be used to obtain superlinear lower bounds on the size of logarithmic depth arithmetic circuits computing linear transformations. In 2000, Codenotti, Pudlák, and Resta [11] have proposed the *Odd Alternating Cycle Conjecture*, stated below. By an *alternating odd cycle* we refer to a directed graph which forms an odd cycle when the orientation of the edges is ignored, and such that the orientation of the edges alternates with one exception.

Conjecture 1.9 (The Odd Alternating Cycle Conjecture [11]). *For every field \mathbb{F} there exist $\varepsilon > 0$ and an odd integer ℓ such that every n -vertex directed graph G with $\text{minrk}_{\mathbb{F}}(G) \leq \varepsilon \cdot n$ contains an alternating cycle of length ℓ .*

²Indeed, given a t -dimensional orthogonal representation of an n -vertex graph G over a field \mathbb{F} , consider the matrix $B \in \mathbb{F}^{n \times t}$ whose rows contain the vectors associated with the vertices of G . Then, the n by n matrix $B \cdot B^T$ represents \bar{G} and has rank at most t over \mathbb{F} , hence $\text{minrk}_{\mathbb{F}}(\bar{G}) \leq t$.

It was proved in [11] that Conjecture 1.9 implies, if true, that certain explicit circulant matrices have superlinear rigidity. In contrast, for $\ell = 3$ it was shown in [11] that there are n -vertex (undirected) triangle-free graphs G satisfying $\text{minrk}_{\mathbb{F}}(G) \leq O(n^{3/4})$ for every field \mathbb{F} , and it was left open whether the statement of Conjecture 1.9 may hold for larger values of ℓ . In the recent work [23] the conjecture was disproved over the real field, but remained open for finite fields which are of special interest in circuit complexity. For the orthogonality dimension over the binary field \mathbb{F}_2 , it was shown in [11] that there exist triangle-free n -vertex graphs G satisfying $\bar{\xi}(\overline{G}, \mathbb{F}_2) = n/4 + 2$. It was asked there whether every n -vertex graph G satisfying $\bar{\xi}(\overline{G}, \mathbb{F}_2) \leq n/4 + 1$ must contain a triangle.

In the current work we prove a new upper bound on the minrank parameter over finite fields of *generalized Kneser graphs*. In these graphs the vertices are all the s -subsets of a universe $[d]$, where two sets are adjacent if their intersection size is smaller than some integer m . Note that for $m = 1$ we get the standard family of Kneser graphs (see Definition 1.2). In the proof we modify and extend an argument of [23], which is based on linear spaces of multivariate polynomials, building on a previous work of Alon [2]. For the precise statement, see Theorem 3.2. We turn to describe several applications of our bound.

As a first application, we establish an explicit construction of graphs that do not contain short odd cycles and yet have low minrank over every finite field.

Theorem 1.10. *For every odd integer $\ell \geq 3$ there exists $\delta = \delta(\ell) > 0$ such that for every sufficiently large integer n , there exists an n -vertex graph G with no odd cycle of length at most ℓ such that for every finite field \mathbb{F} ,*

$$\text{minrk}_{\mathbb{F}}(G) \leq n^{1-\delta}.$$

Theorem 1.10 immediately implies that the Odd Alternating Cycle Conjecture is false over every finite field, even for undirected graphs. This can be viewed as falling into the recent line of non-rigidity results based on the polynomial method (see, e.g., [1, 15, 16]).

We next consider the behavior of the orthogonality dimension over the binary field of the complement of triangle-free graphs. It is relevant to mention here that in the proof of Theorem 1.10, the matrices that imply the stated bound on the minrank are symmetric (see Remark 3.3). For the binary field, this can be combined with a matrix decomposition result due to Lempel [29] to obtain the following theorem, which answers a question of [11] negatively.

Theorem 1.11. *There exists a constant $\delta > 0$ such that for every sufficiently large integer n there exists a triangle-free n -vertex graph G such that $\bar{\xi}(\overline{G}, \mathbb{F}_2) \leq n^{1-\delta}$.*

The above result can also be stated in terms of nearly orthogonal systems. For a field \mathbb{F} , a system of vectors in \mathbb{F}^m is said to be *nearly orthogonal* if every vector of the system is not self-orthogonal and any set of three of them contains an orthogonal pair. For the real field, it was proved by Rosenfeld [36] that every nearly orthogonal system in \mathbb{R}^m has size at most $2m$. Theorem 1.11 shows that the situation is quite different over the binary field. Namely, it implies that there exists a constant $\delta > 0$ such that for infinitely many integers m there exists a nearly orthogonal system in \mathbb{F}_2^m of size at least $m^{1+\delta}$.

We finally mention that our bound on the minrank parameter of generalized Kneser graphs can be used to obtain graphs with a constant *vector chromatic number* χ_v (see Definition 3.9) whose complement has a polynomially large minrank over every finite field.

Theorem 1.12. *There exists a constant $\delta > 0$ such that for infinitely many integers n there exists an n -vertex graph G such that $\chi_v(G) \leq 3$ and yet $\text{minrk}_{\mathbb{F}}(\overline{G}) \geq n^\delta$ for every finite field \mathbb{F} .*

The interest in such graphs comes from the semidefinite programming algorithmic approach applied in [9] for approximating the minrank parameter. As explained in [21], such graphs imply a limitation on this approach, which is based on the constant vector chromatic number of the complement of the instances. Theorem 1.12 improves on [21, Theorem 1.3] where the bound on the minrank is shown only for sufficiently large finite fields.

1.2 Outline

The rest of the paper is organized as follows. In Section 2, we prove our bounds on the generalized orthogonality dimension parameters of Kneser graphs (Theorems 1.6 and 1.7) and derive our hardness result (Theorem 1.1). In Section 3, we prove our bound on the minrank parameter over finite fields of generalized Kneser graphs and deduce Theorems 1.10, 1.11, and 1.12.

2 The Generalized Orthogonality Dimension of Kneser Graphs

In this section we study the generalized orthogonality dimension parameters of Kneser graphs, namely, the quantities $\overline{\xi}_k(K(d, s))$ (recall Definitions 1.2 and 1.4), and prove Theorems 1.6 and 1.7. We start with a linear algebra lemma that will be useful in our proofs.

2.1 Linear Algebra Lemma

Lemma 2.1. *Let U be a subspace of \mathbb{R}^t with $\dim(U) = \ell$, let \mathcal{W} be a finite collection of subspaces of \mathbb{R}^t , and let $\ell' \leq \ell$ be an integer satisfying $\dim(U \cap W) \leq \ell'$ for every $W \in \mathcal{W}$. Then, there exists a subspace U' of U with $\dim(U') = \ell - \ell'$ such that $\dim(U' \cap W) = 0$ for every $W \in \mathcal{W}$.*

Intuitively, given a subspace U and a collection \mathcal{W} as in the lemma, a ‘random-like’ subspace U' of U with dimension $\ell - \ell'$ is expected to have a trivial intersection with each of the subspaces of \mathcal{W} , and thus to satisfy the assertion of the lemma. To prove it formally, we use the following well-known fact.

Fact 2.2. *Let U be a subspace of \mathbb{R}^t , and let \mathcal{W} be a finite collection of proper subspaces of U . Then, \mathcal{W} does not cover U , that is, there exists a vector $u \in U$ such that $u \notin W$ for every $W \in \mathcal{W}$.*

Proof: We may assume without loss of generality that $U = \mathbb{R}^t$ (otherwise apply to U and to the subspaces of \mathcal{W} an invertible linear transformation from U to $\mathbb{R}^{\dim(U)}$) and that $\dim(W) = t - 1$ for every $W \in \mathcal{W}$ (otherwise replace W with an arbitrary subspace of dimension $t - 1$ that contains it). Consider the set $S = \{(1, \alpha, \alpha^2, \dots, \alpha^{t-1}) \mid \alpha \in \mathbb{R}\} \subseteq \mathbb{R}^t$. Every subspace $W \in \mathcal{W}$ consists of all the points $x \in \mathbb{R}^t$ satisfying a certain linear equation $\sum_{i=1}^t a_i x_i = 0$, so the intersection $S \cap W$ is the set of all points $(1, \alpha, \alpha^2, \dots, \alpha^{t-1})$ where α satisfies $\sum_{i=1}^t a_i \alpha^{i-1} = 0$. Since a polynomial of degree $t - 1$ has at most $t - 1$ zeros, it follows that $S \cap W$ is finite for every W , hence the union of the subspaces of the finite collection \mathcal{W} intersects S at a finite number of points. This implies that \mathcal{W} does not cover S , and in particular, it does not cover the entire space \mathbb{R}^t . \blacksquare

We use Fact 2.2 to prove the following lemma.

Lemma 2.3. *Let U be a subspace of \mathbb{R}^t with $\dim(U) = \ell$, and let \mathcal{W} be a finite collection of subspaces of U . Then for every integer $\ell' \leq \ell$ there exists a subspace U' of U with $\dim(U') = \ell'$ such that for every $W \in \mathcal{W}$ it holds that*

$$\dim(U' \cap W) = \max(0, \dim(W) + \ell' - \ell).$$

Proof: Let U and \mathcal{W} be as in the statement of the lemma. It suffices to show that for every $\ell' \leq \ell$ there exists a subspace U' of U with $\dim(U') = \ell'$ such that for every $W \in \mathcal{W}$ it holds that

$$\dim(U' \cap W) \leq \max(0, \dim(W) + \ell' - \ell). \quad (1)$$

Indeed, such a subspace U' also satisfies for every subspace $W \in \mathcal{W}$ that $\dim(U' \cap W) \geq 0$ and that

$$\begin{aligned} \dim(U' \cap W) &= \dim(U') + \dim(W) - \dim(U' + W) \geq \dim(U') + \dim(W) - \dim(U + W) \\ &= \dim(U') + \dim(W) - \dim(U) = \dim(W) + \ell' - \ell, \end{aligned}$$

so the inequality in (1) is in fact an equality.

We prove the existence of a subspace U' satisfying (1) by induction on ℓ' . For $\ell' = 0$ the statement trivially holds for the choice $U' = \{0\}$. Assume that for $1 \leq \ell' \leq \ell$ there exists a subspace U'' of U with $\dim(U'') = \ell' - 1$ such that for every $W \in \mathcal{W}$ it holds that

$$\dim(U'' \cap W) \leq \max(0, \dim(W) + \ell' - \ell - 1). \quad (2)$$

Consider the collection \mathcal{W}' of all the proper subspaces of U in $\{U'' + W \mid W \in \mathcal{W}\}$. By Fact 2.2, there exists a vector $u \in U$ that does not lie in any of the subspaces of \mathcal{W}' . Define

$$U' = U'' + \text{span}(u).$$

Since $u \notin U''$ we have $\dim(U') = \dim(U'') + 1 = \ell'$. Fix a subspace $W \in \mathcal{W}$. If $U'' + W = U$ then

$$\dim(U'' \cap W) = \dim(U'') + \dim(W) - \dim(U'' + W) = \dim(W) + \ell' - \ell - 1,$$

and thus

$$\dim(U' \cap W) \leq \dim(U'' \cap W) + 1 = \dim(W) + \ell' - \ell \leq \max(0, \dim(W) + \ell' - \ell),$$

as required. Otherwise, $U'' + W$ is a proper subspace of U , hence $\dim(U'' + W) \leq \ell - 1$. It follows that

$$\dim(U'' \cap W) = \dim(U'') + \dim(W) - \dim(U'' + W) \geq \dim(W) + \ell' - \ell,$$

which using the induction hypothesis given in (2) implies that $\dim(U'' \cap W) = 0$. Our choice of u guarantees that $u \notin U'' + W$, that is, $(u + U'') \cap W = \emptyset$, hence $\dim(U' \cap W) = \dim(U'' \cap W) = 0$. In particular, $\dim(U' \cap W) \leq \max(0, \dim(W) + \ell' - \ell)$, and we are done. \blacksquare

Equipped with Lemma 2.3, we are ready to derive Lemma 2.1.

Proof of Lemma 2.1: Apply Lemma 2.3 to the subspace U with the collection of its subspaces $\{U \cap W \mid W \in \mathcal{W}\}$. We obtain a subspace U' of U with $\dim(U') = \ell - \ell'$ such that for every $W \in \mathcal{W}$ it holds that $\dim(U' \cap W) = \max(0, \dim(U \cap W) - \ell') = 0$, as desired. \blacksquare

2.2 The case $s = 2$

We turn to prove Theorem 1.6 which determines the generalized orthogonality dimension parameters of Kneser graphs $K(d, s)$ for $s = 2$.

Proof of Theorem 1.6: Fix an integer $k \geq 1$. For the upper bound, recall that for all integers $d \geq 4$ we have

$$\bar{\zeta}_k(K(d, 2), \mathbb{R}) \leq \chi_k(K(d, 2)) = \left\lceil \frac{k}{2} \right\rceil \cdot (d - 4) + 2k.$$

For the lower bound, we consider the induced subgraph of $K(d, 2)$, denoted by $K^-(d, 2)$, obtained from $K(d, 2)$ by removing one of its vertices, say, the vertex $\{1, 2\}$. We turn to prove that for all integers $d \geq 4$ it holds that

$$\bar{\zeta}_k(K^-(d, 2), \mathbb{R}) \geq \left\lceil \frac{k}{2} \right\rceil \cdot (d - 4) + 2k, \quad (3)$$

which immediately implies the required lower bound on $\bar{\zeta}_k(K(d, 2), \mathbb{R})$ as well. To this end, we apply an induction on d . For $d = 4$, the graph $K(d, 2)$ is a perfect matching on 6 vertices, hence its subgraph $K^-(d, 2)$ clearly contains an edge. Since every orthogonal k -subspace representation of this graph assigns to the vertices of this edge orthogonal k -subspaces it follows that

$$\bar{\zeta}_k(K^-(4, 2), \mathbb{R}) \geq 2k,$$

as desired. Now, fix some $d > 4$. Assuming that (3) holds for $d - 1$, we turn to prove it for d .

Recall that the vertex set V of $K^-(d, 2)$ consists of all the 2-subsets of $[d]$ except $\{1, 2\}$. Let $(U_A)_{A \in V}$ be a t -dimensional orthogonal k -subspace representation of $K^-(d, 2)$. We proceed by considering the following two cases.

Assume first that there exists some $i \geq 4$ for which

$$\dim(U_{\{1,3\}} \cap U_{\{1,i\}}) \geq \left\lceil \frac{k}{2} \right\rceil. \quad (4)$$

In this case, consider the induced subgraph of $K^-(d, 2)$ on the vertex set V' obtained from V by removing the vertex $\{3, i\}$ and all the vertices that include the element 1. Notice that this subgraph is isomorphic to $K^-(d - 1, 2)$ and that every vertex of V' is disjoint from either $\{1, 3\}$ or from $\{1, i\}$ (or both). This implies that the restriction $(U_A)_{A \in V'}$ of the given assignment to the vertices of V' forms an orthogonal k -subspace representation of $K^-(d - 1, 2)$, all of whose subspaces lie in the subspace of \mathbb{R}^t that is orthogonal to $U = U_{\{1,3\}} \cap U_{\{1,i\}}$. By applying an orthogonal linear transformation from this subspace to $\mathbb{R}^{t - \dim(U)}$, we obtain that

$$\bar{\zeta}_k(K^-(d - 1, 2), \mathbb{R}) \leq t - \dim(U) \leq t - \left\lceil \frac{k}{2} \right\rceil,$$

where in the second inequality we have used (4). Using the induction hypothesis, this implies that

$$t \geq \bar{\zeta}_k(K^-(d - 1, 2), \mathbb{R}) + \left\lceil \frac{k}{2} \right\rceil \geq \left\lceil \frac{k}{2} \right\rceil \cdot (d - 5) + 2k + \left\lceil \frac{k}{2} \right\rceil = \left\lceil \frac{k}{2} \right\rceil \cdot (d - 4) + 2k,$$

and we are done.

We are left with the case where for every $i \geq 4$ it holds that

$$\dim(U_{\{1,3\}} \cap U_{\{1,i\}}) \leq \left\lceil \frac{k}{2} \right\rceil - 1.$$

Apply Lemma 2.1 to the k -subspace $U_{\{1,3\}}$ and the collection $\{U_{\{1,i\}} \mid 4 \leq i \leq d\}$. It follows that there exists a subspace U of $U_{\{1,3\}}$ with $\dim(U) = k - (\lceil \frac{k}{2} \rceil - 1) \geq \lceil \frac{k}{2} \rceil$ such that for every $i \geq 4$ it holds that $\dim(U \cap U_{\{1,i\}}) = 0$. Consider the induced subgraph of $K^-(d, 2)$ on the vertex set V' obtained from V by removing the vertex $\{2, 3\}$ and all the vertices that include the element 1. As before, this subgraph is isomorphic to the graph $K^-(d-1, 2)$.

We define an orthogonal k -subspace representation of this graph as follows. Let B be a set in V' . If $3 \notin B$ we define $\tilde{U}_B = U_B$. Otherwise we have $B = \{3, i\}$ for some $i \geq 4$, and we let $\tilde{U}_{\{3,i\}}$ be the projection of $U_{\{1,i\}}$ to the subspace of \mathbb{R}^t that is orthogonal to U . Note that the fact that $\dim(U \cap U_{\{1,i\}}) = 0$ guarantees that $\dim(\tilde{U}_{\{3,i\}}) = \dim(U_{\{1,i\}}) = k$.

To prove that the assignment $(\tilde{U}_B)_{B \in V'}$ forms an orthogonal k -subspace representation of the graph, let B_1 and B_2 be disjoint sets in V' . If $3 \notin B_1 \cup B_2$ then we have $\tilde{U}_{B_1} = U_{B_1}$ and $\tilde{U}_{B_2} = U_{B_2}$, so it is clear that \tilde{U}_{B_1} and \tilde{U}_{B_2} are orthogonal. Otherwise, assume without loss of generality that $B_1 = \{3, i\}$ for some $i \geq 4$ and that $3 \notin B_2$. In this case we have $\tilde{U}_{B_2} = U_{B_2}$, and since B_2 is disjoint from B_1 it is also disjoint from $\{1, i\}$ and from $\{1, 3\}$, hence \tilde{U}_{B_2} is orthogonal to both $U_{\{1,i\}}$ and $U_{\{1,3\}}$ as well as to the projection \tilde{U}_{B_1} of $U_{\{1,i\}}$ to the subspace orthogonal to $U \subseteq U_{\{1,3\}}$. We get that \tilde{U}_{B_1} and \tilde{U}_{B_2} are orthogonal, as required.

Finally, observe that all the subspaces \tilde{U}_B lie in the subspace of \mathbb{R}^t that is orthogonal to U . Indeed, for sets B with $3 \in B$ this follows from the definition of \tilde{U}_B , and for the other sets this holds because they are disjoint from $\{1, 3\}$. By applying an orthogonal linear transformation from this subspace to $\mathbb{R}^{t-\dim(U)}$, we obtain that

$$\bar{\xi}_k(K^-(d-1, 2), \mathbb{R}) \leq t - \dim(U) \leq t - \left\lceil \frac{k}{2} \right\rceil,$$

and as in the previous case, by the induction hypothesis it follows that $t \geq \lceil \frac{k}{2} \rceil \cdot (d-4) + 2k$, completing the proof. \blacksquare

2.3 General s

We now prove Theorem 1.7 which provides a lower bound on the generalized orthogonality dimension parameters of Kneser graphs $K(d, s)$ for $s \geq 3$.

Proof of Theorem 1.7: Fix integers $k \geq s \geq 3$ and denote $m = \lceil \frac{k+1}{s} \rceil$. Let $d_0 = d_0(s, k)$ be a sufficiently large integer to be determined later. We apply an induction on d . To do so, we define $c = c(s, k)$ to be sufficiently large, say, $c = \frac{k-m+1}{s-1} \cdot (d_0 + s - 2)$, so that the statement of the theorem trivially holds for all integers $d \leq d_0 + s - 2$, and turn to prove the statement for $d \geq d_0$ assuming that it holds for $d - (s - 1)$.

Let $(U_A)_{A \in V}$ be a t -dimensional orthogonal k -subspace representation of $K(d, s)$. We start with some notation. For an s -subset A of $[d]$, an element $i \in A$, and an s -subset B of $[d]$ satisfying $A \cap B = \{i\}$, we let $\mathcal{G}_{A,i}(B)$ denote the collection that consists of the set B and all the sets obtained from B by replacing i with some element from $A \setminus \{i\}$. Note that $|\mathcal{G}_{A,i}(B)| = s$. We say that a vertex

A of $K(d, s)$ is *good* (with respect to the given orthogonal subspace representation) if there exists an $i \in A$ such that for every vertex B satisfying $A \cap B = \{i\}$ it holds that $\dim(U_A \cap U_C) \leq m - 1$ for some $C \in \mathcal{G}_{A,i}(B)$.

Assume first that there exists a good vertex A in $K(d, s)$ associated with an element $i \in A$. Applying Lemma 2.1, we get that there exists a $(k - m + 1)$ -subspace U of U_A such that for every vertex B satisfying $A \cap B = \{i\}$ it holds that $\dim(U \cap U_C) = 0$ for some $C \in \mathcal{G}_{A,i}(B)$. We define an orthogonal k -subspace representation of the graph $K(d - (s - 1), s)$ on the ground set $[d] \setminus (A \setminus \{i\})$ as follows. Let B be an s -subset of $[d] \setminus (A \setminus \{i\})$. If $i \notin B$ we define $\tilde{U}_B = U_B$. Otherwise, we have $A \cap B = \{i\}$, and we let \tilde{U}_B be the projection of U_C to the subspace of \mathbb{R}^t orthogonal to U , where $C \in \mathcal{G}_{A,i}(B)$ is a set satisfying $\dim(U \cap U_C) = 0$. Note that this condition guarantees that $\dim(\tilde{U}_B) = \dim(U_C) = k$.

We claim that the subspaces \tilde{U}_B form an orthogonal k -subspace representation of the graph $K(d - (s - 1), s)$. To see this, let B_1 and B_2 be disjoint s -subsets of $[d] \setminus (A \setminus \{i\})$. If $i \notin B_1 \cup B_2$ then we have $\tilde{U}_{B_1} = U_{B_1}$ and $\tilde{U}_{B_2} = U_{B_2}$, so it is clear that \tilde{U}_{B_1} and \tilde{U}_{B_2} are orthogonal. Otherwise, assume without loss of generality that $i \in B_1$ and $i \notin B_2$. In this case, $\tilde{U}_{B_2} = U_{B_2}$, and \tilde{U}_{B_1} is the projection of U_C to the subspace of \mathbb{R}^t orthogonal to U for some $C \in \mathcal{G}_{A,i}(B_1)$. Since B_2 is disjoint from A , it follows that the subspace \tilde{U}_{B_2} is orthogonal to U_A as well as to its subspace U . It also follows that B_2 is disjoint from every set in $\mathcal{G}_{A,i}(B_1)$, hence the subspace \tilde{U}_{B_2} is orthogonal to U_C . We get that \tilde{U}_{B_2} is orthogonal to \tilde{U}_{B_1} , as required.

Now, observe that the above orthogonal k -subspace representation of $K(d - (s - 1), s)$ lies in the subspace of \mathbb{R}^t that is orthogonal to the $(k - m + 1)$ -subspace U . Indeed, for sets B with $i \in B$ this follows from the definition of \tilde{U}_B , and for the other sets this holds because they are disjoint from A . By applying an orthogonal linear transformation from this subspace to $\mathbb{R}^{t - \dim(U)}$, it follows that

$$\bar{\xi}_k(K(d - (s - 1), s), \mathbb{R}) \leq t - \dim(U) = t - (k - m + 1).$$

Using the induction hypothesis, this implies that

$$t \geq \frac{k - m + 1}{s - 1} \cdot (d - (s - 1)) - c + (k - m + 1) = \frac{k - m + 1}{s - 1} \cdot d - c,$$

and we are done.

We are left with the case where no vertex of $K(d, s)$ is good, for which we need the following lemma.

Lemma 2.4. *If a vertex A of $K(d, s)$ is not good then there exists a nonzero vector $u_A \in U_A$ such that the number of vertices D of $K(d, s)$ for which U_D is not orthogonal to u_A is at most*

$$\binom{2s - 1}{2} \cdot \binom{d - 2}{s - 2}.$$

We first show how Lemma 2.4 completes the proof of the theorem. Assume that no vertex of $K(d, s)$ is good, and consider the following process: We start with the entire vertex set of $K(d, s)$, and in every iteration we choose an arbitrary vertex A associated with its nonzero vector $u_A \in U_A$ from Lemma 2.4 and eliminate all vertices whose subspaces are not orthogonal to u_A . The nonzero vectors associated with the chosen vertices are clearly pairwise orthogonal, and their number, just

like the number of iterations in the process, is at least

$$\frac{\binom{d}{s}}{\binom{2s-1}{2} \cdot \binom{d-2}{s-2}} \geq \frac{k-m+1}{s-1} \cdot d - c,$$

where the inequality holds for every $d \geq d_0$ assuming that $d_0 = d_0(s, k)$ is sufficiently large (because the left-hand side of the inequality is quadratic in d whereas the right-hand side is linear in d). However, the size of the obtained orthogonal set cannot exceed the dimension t , hence

$$t \geq \frac{k-m+1}{s-1} \cdot d - c,$$

and we are done. It remains to prove Lemma 2.4.

Proof of Lemma 2.4: Assume that A is not a good vertex of $K(d, s)$ and fix an arbitrary $i \in A$. Then there exists a vertex B satisfying $A \cap B = \{i\}$ such that $\dim(U_A \cap U_C) \geq m$ for every vertex $C \in \mathcal{G}_{A,i}(B)$. Denote $\mathcal{G}_{A,i}(B) = \{C_1, \dots, C_s\}$, and recall that every set C_j intersects A at one distinct element. For every $j \in [s]$ define $V_j = U_A \cap U_{C_j}$ and $W_j = V_1 + \dots + V_j$. Note that

$$W_1 \subseteq W_2 \subseteq \dots \subseteq W_s \subseteq U_A.$$

Since $\dim(W_1) = \dim(V_1) \geq m$ and $\dim(U_A) = k < m \cdot s$, there must exist some $j \in [s-1]$ for which

$$\dim(W_{j+1}) - \dim(W_j) < m. \quad (5)$$

For this j , we have

$$\begin{aligned} \dim(V_{j+1} \cap W_j) &= \dim(V_{j+1}) + \dim(W_j) - \dim(V_{j+1} + W_j) \\ &= \dim(V_{j+1}) + \dim(W_j) - \dim(W_{j+1}) > m - m = 0, \end{aligned}$$

where the inequality follows by combining (5) with the fact that $\dim(V_{j+1}) \geq m$. This implies that there exists a nonzero vector u_A in $V_{j+1} \cap W_j$. Observe that

$$u_A \in U_A \cap U_{C_{j+1}} \cap (U_{C_1} + \dots + U_{C_j}).$$

Now, consider a vertex D of $K(d, s)$ whose subspace U_D is not orthogonal to u_A . It follows that U_D is not orthogonal to U_A , to $U_{C_{j+1}}$, and to $U_{C_1} + \dots + U_{C_j}$, hence D intersects the sets A and C_{j+1} as well as at least one of the sets C_1, \dots, C_j . We claim that D must include at least two elements from $A \cup B$. Indeed, D intersects A but if D includes from $A \cup B$ only one element and this element belongs to A then D either does not intersect C_{j+1} or does not intersect any of C_1, \dots, C_j . It follows that the number of vertices D for which U_D is not orthogonal to u_A is bounded from above by the number of s -subsets of $[d]$ that include at least two elements from the $2s-1$ elements of $A \cup B$. The latter is at most $\binom{2s-1}{2} \cdot \binom{d-2}{s-2}$, as required. \blacksquare

The proof is completed. \blacksquare

As immediate corollaries of Theorem 1.7, we obtain the following.

Corollary 2.5. For every integers $s \geq 3$ and $\ell \geq 2$ there exists $c = c(s, \ell)$ such that for all integers $d \geq 2s$,

$$\bar{\xi}_{\ell, s-1}(K(d, s), \mathbb{R}) \geq \ell \cdot d - c.$$

As mentioned before, the bound given in Corollary 2.5 is tight up to the additive constant c .

Corollary 2.6. There exists an absolute constant c such that for all integers $d \geq 6$,

$$\bar{\xi}_4(K(d, 3), \mathbb{R}) \geq 3d/2 - c.$$

Equipped with Corollary 2.6, we are ready to deduce Theorem 1.1.

Proof of Theorem 1.1: Let t be a sufficiently large integer. Recall that a result of [6] implies that $\bar{\xi}_3(K(t, 3), \mathbb{R}) = t$, whereas Corollary 2.6 implies that $\bar{\xi}_4(K(t, 3), \mathbb{R}) \geq 3t/2 - c$ for an absolute constant c . Applying Proposition 1.5 with $F = K(t, 3)$, it follows that it is NP-hard to decide whether an input graph G satisfies $\bar{\xi}(G, \mathbb{R}) \leq t$ or $\bar{\xi}(G, \mathbb{R}) \geq 3t/2 - c$, as desired. ■

3 The Minrank of Generalized Kneser Graphs

In this section we consider a generalization of the family of Kneser graphs, defined as follows.

Definition 3.1 (Generalized Kneser Graphs). For integers $m \leq s \leq d$, the generalized Kneser graph $K^<(d, s, m)$ is the graph whose vertices are all the s -subsets of $[d]$, where two sets A, B are adjacent if $|A \cap B| < m$.

For this family of graphs, we prove the following upper bound on the minrank parameter over finite fields (recall Definition 1.8).

Theorem 3.2. For all integers $m \leq s \leq d$ and for every finite field \mathbb{F} ,

$$\text{minrk}_{\mathbb{F}}(K^<(d, s, m)) \leq \sum_{i=0}^{s-m} \binom{d}{i}.$$

Moreover, the bound on the minrank can be achieved by a symmetric matrix.

Remark 3.3. Theorem 3.2 guarantees that the bound on the minrank can be achieved by a symmetric matrix. This will be crucial for one of our applications, namely, for a construction of triangle-free graphs whose complement has low orthogonality dimension over the binary field \mathbb{F}_2 (see Section 3.3.2). We remark, however, that for undirected graphs and for fields of characteristic different from 2, attaining the bound on the minrank by a symmetric matrix can be achieved easily with a factor of 2 worse bound on the minrank. Indeed, if a matrix M represents a graph G over a field \mathbb{F} of characteristic different from 2 and satisfies $\text{rank}_{\mathbb{F}}(M) = r$ then the matrix $M + M^T$ also represents G and has rank at most $2r$ over \mathbb{F} . This argument does not hold over fields of characteristic 2, since in this case the diagonal entries of $M + M^T$ are all zeros.

As in the previous section, we start with a simple linear algebra lemma.

3.1 Linear Algebra Lemma

Lemma 3.4. For a graph G on the vertex set $[n]$, let $M \in \mathbb{Z}^{n \times n}$ be an integer matrix such that $M_{i,i} = 1$ for every $i \in [n]$, and $M_{i,j} = 0$ for every distinct non-adjacent vertices i and j in G . Then, for every finite field \mathbb{F} , $\text{minrk}_{\mathbb{F}}(G) \leq \text{rank}_{\mathbb{R}}(M)$.

We need the following fact.

Fact 3.5. Let p be a prime and let M be an integer matrix. Then, the matrix $M' = M \pmod{p}$ satisfies $\text{rank}_{\mathbb{F}_p}(M') \leq \text{rank}_{\mathbb{R}}(M)$.

Proof: It suffices to show that if some rows v_1, \dots, v_k of M are linearly dependent over \mathbb{R} then, considered modulo p , they are also linearly dependent over \mathbb{F}_p . To see this, assume that there exist $a_1, \dots, a_k \in \mathbb{R}$, at least one of which is nonzero, for which $\sum_{i=1}^k a_i v_i = 0$. Since the v_i 's are integer vectors it can be assumed that $a_1, \dots, a_k \in \mathbb{Z}$ and that $\text{gcd}(a_1, \dots, a_k) = 1$. This implies that they are not all zeros modulo p . Therefore, the same coefficients, considered modulo p , provide a non-trivial combination of the corresponding rows of M' with sum zero, and we are done. ■

Proof of Lemma 3.4: Let \mathbb{F} be a finite field and denote its characteristic by p . For a graph G and an integer matrix M as in the statement of the lemma, consider the matrix $M' = M \pmod{p}$. Observe that M' represents G over \mathbb{F}_p , hence by Fact 3.5, $\text{minrk}_{\mathbb{F}_p}(G) \leq \text{rank}_{\mathbb{F}_p}(M') \leq \text{rank}_{\mathbb{R}}(M)$. Since \mathbb{F}_p is a subfield of \mathbb{F} , it holds that $\text{minrk}_{\mathbb{F}}(G) \leq \text{minrk}_{\mathbb{F}_p}(G)$, and we are done. ■

3.2 Proof of Theorem 3.2

We are ready to prove Theorem 3.2.

Proof of Theorem 3.2: Consider the polynomial $q \in \mathbb{R}[x]$ defined by

$$q(x) = \binom{x-m}{s-m} = \frac{1}{(s-m)!} \cdot (x-m)(x-(m+1)) \cdots (x-(s-1)).$$

Notice that q is an integer-valued polynomial of degree $s-m$. Let $f : \{0,1\}^d \times \{0,1\}^d \rightarrow \mathbb{R}$ be the function defined by

$$f(x,y) = q\left(\sum_{i=1}^d x_i y_i\right)$$

for every $x, y \in \{0,1\}^d$. Expanding f as a linear combination of monomials, the relation $z^2 = z$ for $z \in \{0,1\}$ implies that one can reduce to 1 the exponent of each variable occurring in a monomial. It follows that f can be represented as a multilinear polynomial in the $2d$ variables of x and y . By combining terms involving the same monomial in the variables of x , one can write f as

$$f(x,y) = \sum_{i=1}^R g_i(x) h_i(y)$$

for an integer R and functions $g_i, h_i : \{0,1\}^d \rightarrow \mathbb{R}$, $i \in [R]$, such that the g_i 's are distinct multilinear monomials of total degree at most $s-m$ in d variables. It follows that $R \leq \sum_{i=0}^{s-m} \binom{d}{i}$.

Now, let M_1 and M_2 be the $2^d \times R$ matrices whose rows are indexed by $\{0, 1\}^d$ and whose columns are indexed by $[R]$, defined by $(M_1)_{x,i} = g_i(x)$ and $(M_2)_{x,i} = h_i(x)$. Then, the rank over \mathbb{R} of the matrix $M = M_1 \cdot M_2^T$ is at most R and for every $x, y \in \{0, 1\}^d$ it holds that $M_{x,y} = f(x, y)$. By the definition of f the matrix M is symmetric, and since q is an integer-valued polynomial, all of its entries are integer.

Finally, let V be the vertex set of $K^{\leq}(d, s, m)$, that is, the collection of all s -subsets of $[d]$, and identify every vertex $A \in V$ with an indicator vector $c_A \in \{0, 1\}^d$ in the natural way. Observe that for every $A, B \in V$ we have

$$M_{c_A, c_B} = f(c_A, c_B) = q(|A \cap B|).$$

Hence, for every $A \in V$ we have $|A| = s$ and thus $M_{c_A, c_A} = q(s) = 1$, whereas for every distinct non-adjacent $A, B \in V$ we have $m \leq |A \cap B| \leq s - 1$ and thus $M_{c_A, c_B} = q(|A \cap B|) = 0$. Since the restriction of M to $V \times V$ is symmetric and has rank at most R over the reals, Lemma 3.4 implies that $\text{minrk}_{\mathbb{F}}(K^{\leq}(d, s, m)) \leq R$ for every finite field \mathbb{F} and that the bound can be achieved by a symmetric matrix, as desired. \blacksquare

3.3 Applications

We gather below several applications of Theorem 3.2.

3.3.1 The Odd Alternating Cycle Conjecture over Finite Fields

We turn to disprove Conjecture 1.9 over every finite field. We will use the simple fact that generalized Kneser graphs do not contain short odd cycles, as stated below (see, e.g., [13, 23]).

Lemma 3.6. *Let $\ell \geq 3$ be an odd integer. For every even integer d and an integer $m \leq \frac{d}{2\ell}$, the graph $K^{\leq}(d, \frac{d}{2}, m)$ contains no odd cycle of length at most ℓ .*

We prove the following theorem, confirming Theorem 1.10.

Theorem 3.7. *For every odd integer $\ell \geq 3$ there exists $\delta = \delta(\ell) > 0$ such that for every sufficiently large integer n , there exists an n -vertex graph G with no odd cycle of length at most ℓ such that for every finite field \mathbb{F} ,*

$$\text{minrk}_{\mathbb{F}}(G) \leq n^{1-\delta}.$$

Moreover, the bound on the minrank can be achieved by a symmetric matrix.

Proof: Fix an odd integer $\ell \geq 3$. For an integer d divisible by 2ℓ , consider the graph $G = K^{\leq}(d, \frac{d}{2}, m)$ where $m = \frac{d}{2\ell}$. By Lemma 3.6, G contains no odd cycle of length at most ℓ . As for the minrank parameter, Theorem 3.2 implies that for every finite field \mathbb{F} ,

$$\text{minrk}_{\mathbb{F}}(G) \leq \sum_{i=0}^{d/2-m} \binom{d}{i} \leq 2^{H(\frac{1}{2} - \frac{m}{d}) \cdot d} = 2^{H(\frac{1}{2} - \frac{1}{2\ell}) \cdot d},$$

where H stands for the binary entropy function. Since G has $|V| = \binom{d}{d/2} = 2^{(1-o(1)) \cdot d}$ vertices, for any $\delta > 0$ such that $H(\frac{1}{2} - \frac{1}{2\ell}) < 1 - \delta$ we have $\text{minrk}_{\mathbb{F}}(G) \leq |V|^{1-\delta}$ for every sufficiently large integer d . The proof is completed by considering, for every sufficiently large integer n , some n -vertex subgraph of the graph defined above, where d is the smallest integer divisible by 2ℓ such that $n \leq \binom{d}{d/2}$. \blacksquare

3.3.2 Triangle-free Graphs and the Orthogonality Dimension over the Binary Field

We turn to prove Theorem 1.11. Its proof adopts the following special case of a result due to Lempel [29].

Lemma 3.8 ([29]). *Let M be an n by n symmetric matrix over the binary field \mathbb{F}_2 with at least one nonzero diagonal entry and rank r . Then, there exists an n by r matrix B over \mathbb{F}_2 satisfying $M = B \cdot B^T$.*

Proof of Theorem 1.11: Apply Theorem 3.7 with $\ell = 3$ to obtain some $\delta > 0$ such that for every sufficiently large integer n , there exist a triangle-free n -vertex graph G and an n by n symmetric matrix M over \mathbb{F}_2 of rank $r = \text{rank}_{\mathbb{F}_2}(M) \leq n^{1-\delta}$ that represents G . By Lemma 3.8, there exists an n by r matrix B over \mathbb{F}_2 satisfying $M = B \cdot B^T$. By assigning the i th row of B to the i th vertex of G we get an r -dimensional orthogonal representation of \overline{G} over \mathbb{F}_2 , hence $\overline{\xi}(\overline{G}, \mathbb{F}_2) \leq r \leq n^{1-\delta}$. ■

3.3.3 The Vector Chromatic Number vs. Minrank

The vector chromatic number of graphs, introduced by Karger, Motwani, and Sudan in [26], is defined as follows.

Definition 3.9 (Vector Chromatic Number). *For a graph $G = (V, E)$ the vector chromatic number of G , denoted by $\chi_v(G)$, is the minimal real value of $\kappa > 1$ such that there exists an assignment of a unit vector w_v to every vertex $v \in V$ satisfying the inequality $\langle w_v, w_{v'} \rangle \leq -\frac{1}{\kappa-1}$ whenever v and v' are adjacent in G .*

To prove Theorem 1.12, we need the following simple fact that relates the minrank of a graph to the minrank of its complement (see, e.g., [34, Remark 2.2]).

Fact 3.10. *For every field \mathbb{F} and an n -vertex graph G , $\text{minrk}_{\mathbb{F}}(G) \cdot \text{minrk}_{\mathbb{F}}(\overline{G}) \geq n$.*

Proof of Theorem 1.12: For an integer d divisible by 8, consider the graph $G = K^{<}(d, \frac{d}{2}, m)$ where $m = \frac{d}{8}$. We first claim that $\chi_v(G) \leq 3$. To see this, assign to every vertex A of G , representing a $\frac{d}{2}$ -subset of $[d]$, the unit vector $w_A \in \mathbb{R}^d$ defined by $(w_A)_i = \frac{1}{\sqrt{d}}$ if $i \in A$ and $(w_A)_i = -\frac{1}{\sqrt{d}}$ otherwise. Observe that every two distinct vertices A and B that are adjacent in G satisfy $|A \cap B| < \frac{d}{8}$ and thus $|A \triangle B| > \frac{3d}{4}$, implying that $\langle w_A, w_B \rangle = \frac{d-2|A \cap B|}{d} < -\frac{1}{2}$. This implies that $\chi_v(G) \leq 3$, as claimed. As for the minrank parameter, Theorem 3.2 implies that for every finite field \mathbb{F} ,

$$\text{minrk}_{\mathbb{F}}(G) \leq \sum_{i=0}^{d/2-m} \binom{d}{i} \leq 2^{H(\frac{1}{2}-\frac{m}{d}) \cdot d} = 2^{H(3/8) \cdot d},$$

where H stands for the binary entropy function. Since G has $n = \binom{d}{d/2} = 2^{(1-o(1)) \cdot d}$ vertices, for any $\delta < 1 - H(3/8)$ we have $\text{minrk}_{\mathbb{F}}(G) \leq n^{1-\delta}$ assuming that d is sufficiently large. By Fact 3.10, this implies that $\text{minrk}_{\mathbb{F}}(\overline{G}) \geq n^{\delta}$, and we are done. ■

References

- [1] J. Alman and R. R. Williams. Probabilistic rank and matrix rigidity. In *Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC'17)*, pages 641–652, 2017.

- [2] N. Alon. The Shannon capacity of a union. *Combinatorica*, 18(3):301–310, 1998.
- [3] N. Alon and N. Kahale. Approximating the independence number via the θ -function. *Math. Program.*, 80:253–264, 1998.
- [4] 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.
- [5] 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.
- [6] B. Bukh and C. Cox. On a fractional version of Haemers’ bound. *IEEE Trans. Inform. Theory*, 65(6):3340–3348, 2019.
- [7] 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.
- [8] 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.
- [9] E. Chlamtáč and I. Haviv. Linear index coding via semidefinite programming. *Combinatorics, Probability & Computing*, 23(2):223–247, 2014. Preliminary version in SODA’12.
- [10] V. Chvátal, M. R. Garey, and D. S. Johnson. Two results concerning multicoloring. *Annals of Discrete Math.*, 2:151–154, 1978.
- [11] B. Codenotti, P. Pudlák, and G. Resta. Some structural properties of low-rank matrices related to computational complexity. *Theor. Comput. Sci.*, 235(1):89–107, 2000. Preliminary version in ECCC’97.
- [12] R. de Wolf. Quantum Computing and Communication Complexity. PhD thesis, Universiteit van Amsterdam, 2001.
- [13] T. Denley. The odd girth of the generalised Kneser graph. *Eur. J. Comb.*, 18(6):607–611, 1997.
- [14] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. *SIAM J. Comput.*, 39(3):843–873, 2009. Preliminary version in STOC’06.
- [15] Z. Dvir and B. L. Edelman. Matrix rigidity and the Croot-Lev-Pach lemma. *Theory of Computing*, 15(1):1–7, 2019.
- [16] Z. Dvir and A. Liu. Fourier and circulant matrices are not rigid. In *34th Computational Complexity Conference (CCC’19)*, pages 17:1–17:23, 2019.
- [17] U. Feige. Randomized graph products, chromatic numbers, and the Lovász θ -function. *Combinatorica*, 17(1):79–90, 1997. Preliminary version in STOC’95.

- [18] M. R. Garey and D. S. Johnson. The complexity of near-optimal graph coloring. *J. ACM*, 23(1):43–49, 1976.
- [19] 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.
- [20] W. H. Haemers. On some problems of Lovász concerning the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(2):231–232, 1979.
- [21] I. Haviv. On minrank and the Lovász theta function. In *International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX’18)*, pages 13:1–13:15, 2018.
- [22] I. Haviv. Approximating the orthogonality dimension of graphs and hypergraphs. In *44th International Symposium on Mathematical Foundations of Computer Science (MFCS’19)*, pages 39:1–39:15, 2019.
- [23] I. Haviv. On minrank and forbidden subgraphs. *ACM Transactions on Computation Theory (TOCT)*, 11(4):20, 2019. Preliminary version in RANDOM’18.
- [24] I. Haviv. Topological bounds on the dimension of orthogonal representations of graphs. *Eur. J. Comb.*, 81:84–97, 2019.
- [25] S. Hu, I. Tamo, and O. Shayevitz. A bound on the Shannon capacity via a linear programming variation. *SIAM J. Discrete Math.*, 32(3):2229–2241, 2018. Preliminary version in ISIT’17.
- [26] 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.
- [27] D. E. Knuth. The sandwich theorem. *Electr. J. Comb.*, 1(A1):1–48, 1994.
- [28] M. Langberg and A. Sprintson. On the hardness of approximating the network coding capacity. *IEEE Trans. Inform. Theory*, 57(2):1008–1014, 2011. Preliminary version in ISIT’08.
- [29] A. Lempel. Matrix factorization over $GF(2)$ and trace-orthogonal bases of $GF(2^n)$. *SIAM J. Comput.*, 4(2):175–186, 1975.
- [30] L. Lovász. Kneser’s conjecture, chromatic number, and homotopy. *J. Comb. Theory, Ser. A*, 25(3):319–324, 1978.
- [31] L. Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(1):1–7, 1979.
- [32] L. Lovász. Chapter 5: Orthogonal representations and their dimension. *Lecture notes: Selected topics in graph theory*, 2016. Available at <http://web.cs.elte.hu/%7Elvasz/kurzusok/orth16-2.pdf>.
- [33] 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.
- [34] R. Peeters. Orthogonal representations over finite fields and the chromatic number of graphs. *Combinatorica*, 16(3):417–431, 1996.

- [35] S. Riis. Information flows, graphs and their guessing numbers. *Electr. J. Comb.*, 14(1), 2007.
- [36] M. Rosenfeld. Almost orthogonal lines in E^d . *DIMACS Series in Discrete Math.*, 4:489–492, 1991.
- [37] G. Scarpa and S. Severini. Kochen-Specker sets and the rank-1 quantum chromatic number. *IEEE Trans. Inform. Theory*, 58(4):2524–2529, 2012.
- [38] S. Stahl. n -tuple colorings and associated graphs. *J. Comb. Theory, Ser. B*, 20(2):185–203, 1976.
- [39] S. Stahl. The multichromatic numbers of some Kneser graphs. *Discrete Mathematics*, 185(1-3):287–291, 1998.
- [40] C. Tardif and X. Zhu. A note on Hedetniemi’s conjecture, Stahl’s conjecture and the Poljak-Rödl function. *Electr. J. Comb.*, 26(4):P4.32, 2019.
- [41] L. G. Valiant. Graph-theoretic arguments in low-level complexity. In *6th International Symposium on Mathematical Foundations of Computer Science (MFCS’77)*, pages 162–176, 1977.
- [42] M. Wrochna and S. Živný. 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.