Kernelization for Orthogonality Dimension

Ishay Haviv
School of Computer Science, The Academic College of Tel Aviv-Yaffo, Israel

Dror Rabinovich
School of Computer Science, The Academic College of Tel Aviv-Yaffo, Israel

Abstract

The orthogonality dimension of a graph over $\mathbb{R}$ is the smallest integer d for which one can assign to every vertex a nonzero vector in $\mathbb{R}^d$ such that every two adjacent vertices receive orthogonal vectors. For an integer d, the d-ORTHO-DIM$_\mathbb{R}$ problem asks to decide whether the orthogonality dimension of a given graph over $\mathbb{R}$ is at most d. We prove that for every integer $d \geq 3$, the d-ORTHO-DIM$_\mathbb{R}$ problem parameterized by the vertex cover number k admits a kernel with $O(k^{d-1})$ vertices and bit-size $O(k^{d-1} \cdot \log k)$. We complement this result by a nearly matching lower bound, showing that for any $\varepsilon > 0$, the problem admits no kernel of bit-size $O(k^{d-1-\varepsilon})$ unless $\mathsf{NP} \subseteq \mathsf{coNP/poly}$. We further study the kernelizability of orthogonality dimension problems in additional settings, including over general fields and under various structural parameterizations.

2012 ACM Subject Classification Theory of computation → Parameterized complexity and exact algorithms; Mathematics of computing → Graph coloring; Mathematics of computing → Combinatorial algorithms

Keywords and phrases Orthogonality dimension, Fixed-parameter tractability, Kernelization, Graph coloring

Digital Object Identifier 10.4230/LIPIcs.IPEC.2024.8

Related Version *Full Version*: https://arxiv.org/abs/2408.08380

Funding *Ishay Haviv*: Research supported by the Israel Science Foundation (grant No. 1218/20).
Dror Rabinovich: Research supported by the Israel Science Foundation (grant No. 1218/20).

Acknowledgements We are grateful to the anonymous reviewers for their constructive and helpful feedback.

1 Introduction

For a field $\mathbb{F}$ and an integer d, a d-dimensional orthogonal representation of a graph $G = (V, E)$ over $\mathbb{F}$ is an assignment of a vector $u_v \in \mathbb{F}^d$ with $\langle u_v, u_v \rangle \neq 0$ to each vertex $v \in V$, such that for every two adjacent vertices v and v' in G, it holds that $\langle u_v, u_{v'} \rangle = 0$. We consider here the standard inner product, defined for any two vectors $x, y \in \mathbb{F}^d$ by $\langle x, y \rangle = \sum_{i=1}^{d} x_i \cdot y_i$ with operations performed over the field $\mathbb{F}$. The orthogonality dimension of a graph G over $\mathbb{F}$, denoted by $\overline{\xi}_\mathbb{F}(G)$, is the smallest integer d for which G admits a d-dimensional orthogonal representation over $\mathbb{F}$ (see Definition 10 and Remark 11).

The notion of orthogonal representations over the real field $\mathbb{R}$ was introduced in 1979 by Lovász [26], who used them to define the celebrated ϑ-function that was motivated by questions in information theory on the Shannon capacity of graphs. Over the years, orthogonal representations and the orthogonality dimension have found a variety of applications in several areas of research. In graph theory, orthogonal representations over the reals were used by Lovász, Saks, and Schrijver [28] to characterize connectivity properties of graphs (see also [27, Chapter 10]). In computational complexity, the orthogonality dimension over finite fields was related to lower bounds in circuit complexity by Codenotti, Pudlák, and Resta [13] (see also [19, 18]). Over the complex field, it was employed by de Wolf [15] to determine the quantum one-round communication complexity of promise equality problems (see also [9, 6, 7]). Additional notable applications from the area of information theory are related to index coding [3, 1], distributed storage [2], and hat-guessing games [30].

The question of the complexity of determining the orthogonality dimension of a given graph over a specified field was proposed in 1989 by Lovász et al. [28]. For a field $\mathbb{F}$ and an integer d, consider the decision problem that given a graph G asks to decide whether $\overline{\xi}_{\mathbb{F}}(G) \leq d$. It is easy to see that the problem can be solved efficiently for $d \in \{1, 2\}$, because a graph G satisfies $\overline{\xi}_{\mathbb{F}}(G) \leq 1$ if and only if it is edgeless, and it satisfies $\overline{\xi}_{\mathbb{F}}(G) \leq 2$ if and only if it is bipartite. For every $d \geq 3$, however, it was shown by Peeters [29] in 1996 that the problem is NP-hard for every field $\mathbb{F}$. More recently, it was shown in [12] that for every sufficiently large integer d, it is NP-hard to distinguish graphs G that satisfy $\overline{\xi}_{\mathbb{F}}(G) \leq d$ from those satisfying $\overline{\xi}_{\mathbb{F}}(G) \geq 2^{(1-o(1))\cdot d/2}$, provided that $\mathbb{F}$ is either a finite field or $\mathbb{R}$.

Motivated by the diverse applications of orthogonality dimension, the present paper delves into the computational complexity of this graph quantity from the perspective of parameterized complexity. We study the decision problems associated with orthogonality dimension with respect to various structural parameterizations, with a particular attention dedicated to the vertex cover number parameterization. We exhibit fixed-parameter tractability results for such problems, along with upper and lower bounds on their kernelizability. Our approach draws its inspiration from prior work on the parameterized complexity of coloring problems, most notably by Jansen and Kratsch [23] (see also [22, Chapter 7]) and by Jansen and Pieterse [24]. In what follows, we provide an overview on relevant research on coloring problems and then turn to a description of our contribution. We use here standard notions from the area of parameterized complexity, whose definitions can be found in Section 2.4 (see also, e.g., [14, 17]).

Graph coloring is a cornerstone concept in graph theory that has been extensively studied from a computational point of view. For an integer q, a q-coloring of a graph G is an assignment of a color to each vertex of G from a set of q colors. The coloring is said to be proper if it assigns distinct colors to every two adjacent vertices in the graph. A graph G is called q-colorable if it admits a proper q-coloring, and the smallest integer q for which G is q-colorable is called the chromatic number of G and is denoted by $\chi(G)$. For an integer q, let q-Coloring denote the decision problem that given a graph G asks to decide whether $\chi(G) \leq q$. The Coloring problem is defined similarly, with the only, yet crucial, difference that the number of colors q is not fixed but forms part of the input. It is well known that the q-Coloring problem can be solved in polynomial time for $q \in \{1, 2\}$ and is NP-complete for every $q \geq 3$. This implies that the Coloring problem, parameterized by the number of colors q, is not fixed-parameter tractable unless $\mathsf{P} = \mathsf{NP}$.

The study of the parameterized complexity of coloring problems was initiated in 2003 by Cai [8], who proposed the following terminology. For a family of graphs $\mathcal{G}$ and for an integer k, let $\mathcal{G} + k\mathrm{v}$ denote the family of all graphs that can be obtained from a graph of $\mathcal{G}$ by adding at most k vertices (with arbitrary neighborhoods). Equivalently, a graph $G = (V, E)$ lies in $\mathcal{G} + k\mathrm{v}$ if there exists a set $X \subseteq V$ of size $|X| \leq k$, referred to as a modulator, such that the graph $G \setminus X$ obtained from G by removing the vertices of X lies in $\mathcal{G}$. For example, letting Empty denote the family of all edgeless graphs, the family Empty $+ k\mathrm{v}$ consists of all the graphs that admit a vertex cover of size at most k. For an integer q, the q-Coloring problem on $\mathcal{G} + k\mathrm{v}$ graphs is the parameterized problem defined as follows.

Input: A graph $G = (V, E)$ and a set $X \subseteq V$ such that $G \setminus X \in \mathcal{G}$.

Question: Is $\chi(G) \leq q$?

Parameter: The size $k = |X|$ of the modulator X.

As before, the Coloring problem on $\mathcal{G} + k\mathrm{v}$ graphs is defined similarly, except that the number of colors q is not fixed but forms part of the input.

A common parameterization of graph problems that received a considerable amount of attention in the literature is that of the vertex cover number, corresponding to the family $\mathcal{G} = \textsc{Empty}$ (see, e.g., [16]). It is well known that the Coloring problem on $\textsc{Empty} + k\mathrm{v}$ graphs is fixed-parameter tractable. Nevertheless, Bodlaender, Jansen, and Kratsch [5] proved that the problem does not admit a kernel of polynomial size under the assumption $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$, whose refutation is known to imply the collapse of the polynomial-time hierarchy [32]. Yet, for any fixed integer $q \geq 3$, Jansen and Kratsch [23] showed that the q-Coloring problem on $\textsc{Empty} + k\mathrm{v}$ graphs admits a kernel with $O(k^q)$ vertices which can be encoded in $O(k^q)$ bits (see also [16]). This result was improved by Jansen and Pieterse [24] as an application of an algebraic sparsification technique they introduced in [25]. It was shown in [24] that for every $q \geq 3$, the q-Coloring problem on $\textsc{Empty}+k\mathrm{v}$ graphs admits a kernel with $O(k^{q-1})$ vertices and bit-size $O(k^{q-1} \cdot \log k)$ (see [24] for various generalizations). On the contrary, it was shown in [23, 24] that for every $q \geq 3$ and any $\varepsilon > 0$, the problem does not admit a kernel that can be encoded in $O(k^{q-1-\varepsilon})$ bits unless $\mathsf{NP} \subseteq \mathsf{coNP/poly}$, thereby settling its kernelization complexity up to a multiplicative $k^{o(1)}$ term.

The paper [23] further studied the kernelization complexity of the q-Coloring problem on $\mathcal{G} + k\mathrm{v}$ graphs for general families $\mathcal{G}$. In particular, they considered graph families $\mathcal{G}$ that are hereditary (i.e., closed under removal of vertices) and that are, roughly speaking, local with respect to the q-List Coloring problem, in the sense that a NO instance of q-List Coloring that involves a graph from $\mathcal{G}$ must have a NO sub-instance whose size depends solely on q. For such families $\mathcal{G}$, it was shown in [23] that the q-Coloring problem on $\mathcal{G} + k\mathrm{v}$ graphs admits a kernel of polynomial size, and this result was complemented with a lower bound on the kernel size relying on the assumption $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$. These results apply, for example, for the families $\cup$Split and $\cup$Cochordal of the graphs whose connected components are split graphs and cochordal graphs respectively. On the other hand, strengthening a result of Bodlaender et al. [4], the authors of [23] proved that the 3-Coloring problem on $\textsc{Path}+k\mathrm{v}$ graphs does not admit a kernel of polynomial size unless $\mathsf{NP} \subseteq \mathsf{coNP/poly}$, where Path stands for the family of path graphs.

1.1 Our Contribution

This paper initiates a systematic study of the parameterized complexity of the orthogonality dimension of graphs. It is noteworthy that the orthogonality dimension of graphs is related to their chromatic number. Indeed, for every field $\mathbb{F}$ and for every graph G, it holds that $\overline{\xi}_{\mathbb{F}}(G) \leq \chi(G)$, because a proper q-coloring of G may be viewed as a q-dimensional orthogonal representation of G over $\mathbb{F}$ that uses only vectors from the standard basis of $\mathbb{F}^q$ (see Claim 12). Yet, it turns out that the two graph quantities can differ substantially, as there exist graphs where the orthogonality dimension is exponentially smaller than the chromatic number (see, e.g., [20, Proposition 2.2]). Our investigation of the parameterized complexity of orthogonality dimension problems aligns with the approach of [23, 24] for studying coloring problems within the parameterized complexity framework. While coloring problems are primarily combinatorial in nature, the attempt to prove analogous results for orthogonality dimension raises intriguing questions reflecting the algebraic aspects of this graph quantity.

We first introduce the decision problems associated with orthogonality dimension.

▶ **Definition 1.** *For a field $\mathbb{F}$, the Ortho-Dim$_{\mathbb{F}}$ problem is defined as follows.*

Input: A graph $G = (V, E)$ and an integer d.

Question: Is $\overline{\xi}_{\mathbb{F}}(G) \leq d$?

For a field $\mathbb{F}$ and a family of graphs $\mathcal{G}$, the (parameterized) Ortho-Dim$_{\mathbb{F}}$ problem on $\mathcal{G} + k\mathrm{v}$ graphs is defined as follows.

Input: A graph $G = (V, E)$, a set $X \subseteq V$ such that $G \setminus X \in \mathcal{G}$, and an integer d.
Question: Is $\overline{\xi}_{\mathbb{F}}(G) \leq d$?
Parameter: The size $k = |X|$ of the modulator X.

For a field $\mathbb{F}$, an integer d, and a family of graphs $\mathcal{G}$, the (parameterized) d-ORTHO-DIM$_{\mathbb{F}}$ problem on $\mathcal{G} + k\mathrm{v}$ graphs is defined as follows.

Input: A graph $G = (V, E)$ and a set $X \subseteq V$ such that $G \setminus X \in \mathcal{G}$.
Question: Is $\overline{\xi}_{\mathbb{F}}(G) \leq d$?
Parameter: The size $k = |X|$ of the modulator X.

Let us stress that the integer d forms part of the input in the ORTHO-DIM$_{\mathbb{F}}$ problem, whereas it is a fixed constant in the d-ORTHO-DIM$_{\mathbb{F}}$ problem. Note that the hardness result of [29] implies that for every field $\mathbb{F}$, the ORTHO-DIM$_{\mathbb{F}}$ problem parameterized by the solution value d is not fixed-parameter tractable unless $\mathsf{P} = \mathsf{NP}$.

The main parameterization we consider is the vertex cover number of the input graph, which corresponds to the family $\mathcal{G} =$ EMPTY. We start with the following fixed-parameter tractability result.

▶ **Theorem 2.** *Let $\mathbb{F}$ be either a finite field or $\mathbb{R}$. The ORTHO-DIM$_{\mathbb{F}}$ problem on EMPTY$+k\mathrm{v}$ graphs is fixed-parameter tractable.*

In fact, we prove an extension of Theorem 2, showing that if the ORTHO-DIM$_{\mathbb{F}}$ problem over a field $\mathbb{F}$ is decidable, then the corresponding ORTHO-DIM$_{\mathbb{F}}$ problem on EMPTY $+ k\mathrm{v}$ graphs is fixed-parameter tractable (see Theorem 14). While it is easy to see that the ORTHO-DIM$_{\mathbb{F}}$ problem is decidable for any finite field $\mathbb{F}$, the real case relies on a result of Tarski [31] on the decidability of the existential theory of the reals (see Proposition 17).

We next consider the kernelizability of the d-ORTHO-DIM$_{\mathbb{F}}$ problem parameterized by the vertex cover number for a fixed integer d. For finite fields $\mathbb{F}$, one may deduce from a result of [24] that the problem admits a kernel of polynomial size, where the degree of the polynomial grows exponentially with d. We prove the following generalized and stronger result.

▶ **Theorem 3.** *For every field $\mathbb{F}$ and for every integer $d \geq 3$, the d-ORTHO-DIM$_{\mathbb{F}}$ problem on EMPTY$+ k\mathrm{v}$ graphs admits a kernel with $O(k^d)$ vertices and bit-size $O(k^d)$.*

Theorem 3 prompts us to determine the smallest possible kernel size for the d-ORTHO-DIM$_{\mathbb{F}}$ problem on EMPTY $+ k\mathrm{v}$ graphs. The following result furnishes a lower bound, conditioned on the complexity-theoretic assumption $\mathsf{NP} \not\subseteq \mathsf{coNP/poly}$.

▶ **Theorem 4.** *For every field $\mathbb{F}$, every integer $d \geq 3$, and any real $\varepsilon > 0$, the d-ORTHO-DIM$_{\mathbb{F}}$ problem on EMPTY $+ k\mathrm{v}$ graphs does not admit a kernel with bit-size $O(k^{d-1-\varepsilon})$ unless $\mathsf{NP} \subseteq \mathsf{coNP/poly}$.*

The proof of Theorem 4 combines the lower bound on kernels for coloring problems proved in [23, 24] with a novel linear-parameter transformation from those problems to those associated with orthogonality dimension. More specifically, we show that for every field $\mathbb{F}$ and for every integer $d \geq 3$, it is possible to efficiently transform a graph G into a graph G' so that $\chi(G) \leq d$ if and only if $\overline{\xi}_{\mathbb{F}}(G') \leq d$ while essentially preserving the vertex cover number (see Theorem 25). This is in contrast to a reduction of [29], which is appropriate only for $d = 3$ and significantly increases the vertex cover number. The transformation relies on a gadget graph that enforces the vectors assigned to two specified vertices to be either orthogonal or equal up to scalar multiplication (see Lemma 21).

We remark that Theorem 4 implies that, unless NP $\subseteq$ coNP/poly, the degree of the polynomial lower bound on the size of a kernel for the d-Ortho-Dim$_{\mathbb{F}}$ problem on Empty$+k$v graphs can be arbitrarily large when d grows. This yields that the Ortho-Dim$_{\mathbb{F}}$ problem on Empty $+ k$v graphs, in which d constitutes part of the input, is unlikely to admit a kernel of polynomial size.

Theorems 3 and 4 leave a multiplicative gap of roughly k between the upper and lower bounds on the kernel size achievable for the d-Ortho-Dim$_{\mathbb{F}}$ problem parameterized by the vertex cover number. For the real field $\mathbb{R}$, we narrow this gap to a multiplicative term of $k^{o(1)}$, as stated below.

▶ **Theorem 5.** *For every integer $d \geq 3$, the d-Ortho-Dim$_{\mathbb{R}}$ problem on Empty $+ k$v graphs admits a kernel with $O(k^{d-1})$ vertices and bit-size $O(k^{d-1} \cdot \log k)$.*

The proof of Theorem 5 borrows the sparsification technique of [25] (see also [24]). A key technical ingredient in applying this method lies in a construction of a low-degree polynomial, which assesses the feasibility of assigning a vector to a vertex based on the vectors of its neighbors (see Lemma 20). Our construction hinges on the fact that the zero vector is the only self-orthogonal vector over the reals. It would be interesting to decide whether or not a similar upper bound on the kernel size could be obtained for finite fields, where this property does not hold.

We finally turn to the study of kernels for the d-Ortho-Dim$_{\mathbb{F}}$ problem on $\mathcal{G} + k$v graphs for general hereditary graph families $\mathcal{G}$. Our first result in this context offers a sufficient condition on $\mathcal{G}$ for the existence of a polynomial size kernel for d-Ortho-Dim$_{\mathbb{F}}$ on $\mathcal{G} + k$v graphs. This condition is related to a variant of the d-Ortho-Dim$_{\mathbb{F}}$ problem, termed d-Subspace Choosability$_{\mathbb{F}}$, which was previously studied in various forms (see, e.g., [21, 11]) and may be viewed as a counterpart of the q-List Coloring problem for orthogonal representations. In this problem, the input consists of a graph G and an assignment of a subspace of $\mathbb{F}^d$ to each vertex, and the goal is to decide whether G admits an orthogonal representation over $\mathbb{F}$ that assigns to every vertex a vector from its subspace. We show that if the d-Subspace Choosability$_{\mathbb{F}}$ problem on graphs from a family $\mathcal{G}$ is local, in the sense that every NO instance has a NO sub-instance on at most $g(d)$ vertices, then the d-Ortho-Dim$_{\mathbb{F}}$ problem on $\mathcal{G} + k$v graphs admits a kernel with $O(k^{d \cdot g(d)})$ vertices. We demonstrate the applicability of this result for the graph families $\cup$Split and $\cup$Cochordal. On the contrary, for the Path family, we show that it is unlikely that the d-Ortho-Dim$_{\mathbb{F}}$ problem on Path $+ k$v graphs admits a polynomial size kernel even for $d = 3$.

1.2 Outline

The remainder of the paper is structured as follows. In Section 2, we collect several definitions and facts that will be used throughout the paper. In Section 3, we study the fixed-parameter tractability of the Ortho-Dim$_{\mathbb{F}}$ problem parameterized by the vertex cover number and prove Theorem 2. In Section 4, we present polynomial size kernels for the d-Ortho-Dim$_{\mathbb{F}}$ problem parameterized by the vertex cover number and prove Theorems 3 and 5. In Section 5, we complement the results of Section 4 by providing limits on the kernelizability of the d-Ortho-Dim$_{\mathbb{F}}$ problem parameterized by the vertex cover number and prove Theorem 4. For our study on the kernelizability of the d-Ortho-Dim$_{\mathbb{F}}$ problem on $\mathcal{G} + k$v graphs for general hereditary graph families $\mathcal{G}$, we refer the reader to the full version of the paper.

2 Preliminaries

2.1 Notations

For an integer n, let $[n] = \{1, 2, \dots, n\}$. All graphs considered in this paper are simple. For a graph $G = (V, E)$ and a set $X \subseteq V$, we let $G[X]$ denote the subgraph of G induced by X. The set X is called a vertex cover of G if every edge of G is incident with a vertex of X. We let $G \setminus X$ denote the graph obtained from G by removing the vertices of X (and the edges that touch them). For a vertex $v \in V$, we let $N_G(v)$ denote the set of neighbors of v in G.

2.2 Linear Algebra

For a field $\mathbb{F}$ and an integer d, two vectors $x, y \in \mathbb{F}^d$ are said to be orthogonal if $\langle x, y\rangle = 0$, where $\langle x, y\rangle = \sum_{i=1}^{d} x_i y_i$ with operations over $\mathbb{F}$. If $\langle x, x\rangle = 0$ then x is self-orthogonal, and otherwise it is non-self-orthogonal. The orthogonal complement of a subspace $W \subseteq \mathbb{F}^d$ is the subspace $W^\perp$ of all vectors in $\mathbb{F}^d$ that are orthogonal to all vectors of W, that is, $W^\perp = \{x \in \mathbb{F}^d \mid \forall y \in W,\ \langle x, y\rangle = 0\}$. Note that the orthogonal complement satisfies $\dim(W) + \dim(W^\perp) = d$ and $W = (W^\perp)^\perp$. The following simple lemma, proved in the full version of the paper, characterizes the subspaces whose orthogonal complement includes a non-self-orthogonal vector. Recall that the characteristic of a field is the smallest positive number of copies of the field's identity element that sum to zero, or 0 if no such number exists.

▶ **Lemma 6.** *Let $\mathbb{F}$ be a field, let d be an integer, and let W be a subspace of $\mathbb{F}^d$.*

1. *If the characteristic of $\mathbb{F}$ is 2, then there exists a non-self-orthogonal vector in $W^\perp$ if and only if the all-one vector does not lie in W.*
2. *If the characteristic of $\mathbb{F}$ is not 2, then there exists a non-self-orthogonal vector in $W^\perp$ if and only if $W^\perp \not\subseteq W$.*

Borrowing the terminology of [25], we say that a field $\mathbb{F}$ is efficient if field operations and Gaussian elimination can be performed in polynomial time in the size of a reasonable input encoding. All finite fields, as well as the real field $\mathbb{R}$ when restricted to rational numbers (to ensure finite representation), are efficient.

▶ **Lemma 7.** *For every efficient field $\mathbb{F}$, there exists a polynomial-time algorithm that given a collection of vectors in $\mathbb{F}^d$, decides whether there exists a non-self-orthogonal vector in $\mathbb{F}^d$ that is orthogonal to all of them.*

Proof. For input vectors $u_1, \dots, u_\ell \in \mathbb{F}^d$, let $W = \mathrm{span}(u_1, \dots, u_\ell)$. Observe that there exists a non-self-orthogonal vector in $\mathbb{F}^d$ that is orthogonal to $u_1, \dots, u_\ell$ if and only if there exists a non-self-orthogonal vector in the orthogonal complement $W^\perp$. By Lemma 6, for a field $\mathbb{F}$ of characteristic 2, this is equivalent to the all-one vector not lying in W, and for every other field $\mathbb{F}$, this is equivalent to $W^\perp \not\subseteq W$ (that is, at least one vector of a basis of $W^\perp$ does not lie in W). Since $\mathbb{F}$ is an efficient field, these conditions can be checked in polynomial time. This completes the proof. ◀

For a field $\mathbb{F}$ and an integer d, if W is a subspace of $\mathbb{F}^d$ of dimension smaller than d, then its orthogonal complement $W^\perp$ has dimension at least 1, hence there exists a nonzero vector orthogonal to W. However, if we require this vector not only to be nonzero but also non-self-orthogonal, its existence is no longer guaranteed. This consideration motivates the following definition.

▶ **Definition 8.** *For a field $\mathbb{F}$ and an integer d, let $m(\mathbb{F}, d)$ denote the largest integer m such that for every subspace W of $\mathbb{F}^d$ with $\dim(W) < m$, there exists a non-self-orthogonal vector in $W^\perp$.*

▶ Remark 9. For every field $\mathbb{F}$ and for every integer $d \geq 1$, it holds that $1 \leq m(\mathbb{F}, d) \leq d$. Indeed, the lower bound holds because there exists a non-self-orthogonal vector orthogonal to the zero subspace of $\mathbb{F}^d$, and the upper bound holds because no nonzero vector is orthogonal to the entire vector space $\mathbb{F}^d$. For a field $\mathbb{F}$ of characteristic 2, it holds that $m(\mathbb{F}, d) = 1$, because no non-self-orthogonal vector is orthogonal to the 1-dimensional subspace spanned by the all-one vector. For every other field $\mathbb{F}$, it holds that $m(\mathbb{F}, d) \geq \lceil d/2 \rceil$. To see this, consider a subspace $W \subseteq \mathbb{F}^d$ of dimension $\dim(W) < \lceil d/2 \rceil$, and observe that it satisfies $\dim(W^\perp) = d - \dim(W) > d - \lceil d/2 \rceil = \lfloor d/2 \rfloor$, and thus $\dim(W^\perp) > \dim(W)$. This implies that $W^\perp \nsubseteq W$, hence by Item 2 of Lemma 6, there exists a non-self-orthogonal vector in $W^\perp$, as required. We also observe that if the vector space $\mathbb{F}^d$ has no nonzero self-orthogonal vectors, then $m(\mathbb{F}, d) = d$, because for every subspace $W \subseteq \mathbb{F}^d$ of dimension smaller than d there exists a nonzero vector in $W^\perp$. In particular, for every integer d, it holds that $m(\mathbb{R}, d) = d$.

2.3 Orthogonality Dimension

The orthogonality dimension of a graph over a given field is defined as follows.

▶ **Definition 10.** *For a field $\mathbb{F}$ and an integer d, a d-dimensional orthogonal representation of a graph $G = (V, E)$ over $\mathbb{F}$ is an assignment of a vector $u_v \in \mathbb{F}^d$ with $\langle u_v, u_v \rangle \neq 0$ to each vertex $v \in V$, such that for every two adjacent vertices v and v' in G, it holds that $\langle u_v, u_{v'} \rangle = 0$. The orthogonality dimension of a graph G over a field $\mathbb{F}$, denoted by $\overline{\xi}_{\mathbb{F}}(G)$, is the smallest integer d for which G admits a d-dimensional orthogonal representation over $\mathbb{F}$.*

▶ Remark 11. Let us emphasize that the definition of an orthogonal representation does not require vectors assigned to non-adjacent vertices to be non-orthogonal. Orthogonal representations that satisfy this additional property are called faithful (see, e.g., [27, Chapter 10]). Note that orthogonal representations of graphs are sometimes defined in the literature as orthogonal representations of the complement, requiring vectors associated with non-adjacent vertices to be orthogonal (with no constraint imposed on vectors of adjacent vertices). We decided to use here the other definition, but one may view the notation $\overline{\xi}_{\mathbb{F}}(G)$ as standing for $\xi_{\mathbb{F}}(\overline{G})$.

▷ Claim 12. For every field $\mathbb{F}$ and for every graph G, it holds that $\overline{\xi}_{\mathbb{F}}(G) \leq \chi(G)$.

Proof. For a graph $G = (V, E)$, let $q = \chi(G)$, and consider a proper coloring $c : V \to [q]$ of G. Assign to each vertex $v \in V$ the vector $e_{c(v)}$ in $\mathbb{F}^q$, where e_i stands for the vector of $\mathbb{F}^q$ with 1 on the ith entry and 0 everywhere else. The vectors assigned here to the vertices of G are obviously non-self-orthogonal vectors of $\mathbb{F}^q$. Further, for every two adjacent vertices v and v' in G, it holds that $c(v) \neq c(v')$, hence $\langle e_{c(v)}, e_{c(v')} \rangle = 0$. This implies that there exists a q-dimensional orthogonal representation of G over $\mathbb{F}$, and thus $\overline{\xi}_{\mathbb{F}}(G) \leq q$. ◁

2.4 Parameterized Complexity

We present here a few fundamental definitions from the area of parameterized complexity. For a thorough introduction to the field, the reader is referred to, e.g., [14, 17].

A parameterized problem is a set $Q \subseteq \Sigma^* \times \mathbb{N}$ for some finite alphabet Σ. A fixed-parameter algorithm for Q is an algorithm that given an instance $(x, k) \in \Sigma^* \times \mathbb{N}$ decides whether $(x, k) \in Q$ in time $f(k) \cdot |x|^c$ for some computable function f and some constant c. If Q admits a fixed-parameter algorithm, then we say that Q is fixed-parameter tractable.

A compression (also known as generalized kernel and bikernel) for a parameterized problem $Q \subseteq \Sigma^* \times \mathbb{N}$ into a parameterized problem $Q' \subseteq \Sigma^* \times \mathbb{N}$ is an algorithm that given an instance $(x, k) \in \Sigma^* \times \mathbb{N}$ returns in time polynomial in $|x|+k$ an instance $(x', k') \in \Sigma^* \times \mathbb{N}$, such that $(x, k) \in Q$ if and only if $(x', k') \in Q'$, and in addition, $|x'| + k' \leq h(k)$ for some computable function h. The function h is referred to as the size of the compression. If h is polynomial, then the compression is called a polynomial compression. If $|\Sigma| = 2$, the function h is called the bit-size of the compression. When we say that a parameterized problem Q admits a compression of size h, we mean that there exists a compression of size h for Q into some parameterized problem. A compression for a parameterized problem Q into itself is called a kernelization for Q (or simply a kernel). It is well known that a decidable problem admits a kernel if and only if it is fixed-parameter tractable.

A transformation from a parameterized problem $Q \subseteq \Sigma^* \times \mathbb{N}$ into a parameterized problem $Q' \subseteq \Sigma^* \times \mathbb{N}$ is an algorithm that given an instance $(x, k) \in \Sigma^* \times \mathbb{N}$ returns in time polynomial in $|x|+k$ an instance $(x', k') \in \Sigma^* \times \mathbb{N}$, such that $(x, k) \in Q$ if and only if $(x', k') \in Q'$, and in addition, $k' \leq h(k)$ for some computable function h. If h is polynomial, the transformation is called polynomial-parameter, and if h is linear, the transformation is called linear-parameter.

3 Fixed-Parameter Tractability of Ortho-Dim$_{\mathbb{F}}$

In this section, we prove that the ORTHO-DIM$_{\mathbb{F}}$ problem parameterized by the vertex cover number is fixed-parameter tractable for various fields $\mathbb{F}$ (recall Definition 1).

3.1 Finite Fields

We begin with the simple case, where the field $\mathbb{F}$ is finite. The proof resembles the one of the fixed-parameter tractability of the COLORING problem parameterized by the vertex cover number.

▶ **Theorem 13.** *For every finite field $\mathbb{F}$, ORTHO-DIM$_{\mathbb{F}}$ on EMPTY + kv graphs is fixed-parameter tractable.*

Proof. Fix a finite field $\mathbb{F}$. The input of ORTHO-DIM$_{\mathbb{F}}$ on EMPTY + kv graphs consists of a graph $G = (V, E)$, a vertex cover $X \subseteq V$ of G of size $|X| = k$, and an integer d. Consider the algorithm that given such an input acts as follows. If $d > k$ then the algorithm accepts. Otherwise, the algorithm enumerates all possible assignments of non-self-orthogonal vectors from $\mathbb{F}^d$ to the vertices of X. For every such assignment, the algorithm checks for every vertex $v \in V \setminus X$ if there exists a non-self-orthogonal vector in $\mathbb{F}^d$ that is orthogonal to the vectors assigned to the neighbors of v (note that they all lie in X). If there exists an assignment to the vertices of X such that the answer is positive for all the vertices of $V \setminus X$, then the algorithm accepts, and otherwise it rejects.

For correctness, observe first that the input graph G is $(k+1)$-colorable, as follows by assigning k distinct colors to the vertices of the vertex cover X and another color to the vertices of the independent set $V \setminus X$. It thus follows, using Claim 12, that $\overline{\xi}_{\mathbb{F}}(G) \leq \chi(G) \leq k+1$. Therefore, if $d > k$, then it holds that $\overline{\xi}_{\mathbb{F}}(G) \leq d$, hence our algorithm correctly accepts. Otherwise, the algorithm tries all possible assignments of non-self-orthogonal vectors of $\mathbb{F}^d$ to the vertices of X. Since the vertices of $V \setminus X$ form an independent set in G, an assignment to the vertices of X can be extended to the whole graph if and only if for each vertex $v \in V \setminus X$ there exists a non-self-orthogonal vector in $\mathbb{F}^d$ that is orthogonal to the vectors assigned to the neighbors of v (which all lie in X). Since this condition is checked by the algorithm for all possible assignments to the vertices of X, its answer is correct.

We finally analyze the running time of the algorithm. On instances with $d > k$, the algorithm is clearly efficient. For instances with $d \leq k$, the number of assignments of vectors from $\mathbb{F}^d$ to the vertices of X is at most $|\mathbb{F}|^{d \cdot |X|} \leq |\mathbb{F}|^{k^2}$. Further, by Lemma 7, given a collection of vectors of $\mathbb{F}^d$, it is possible to decide in polynomial time whether there exists a non-self-orthogonal vector in $\mathbb{F}^d$ that is orthogonal to all of them. This implies that our algorithm for Ortho-Dim$_{\mathbb{F}}$ on Empty $+ k$v graphs can be implemented in time $|\mathbb{F}|^{k^2} \cdot n^{O(1)}$, where n stands for the input size, hence the problem is fixed-parameter tractable. ◀

3.2 General Fields

We turn to the following generalization of Theorem 13.

▶ **Theorem 14.** *Let $\mathbb{F}$ be a field for which the Ortho-Dim$_{\mathbb{F}}$ problem is decidable. Then the Ortho-Dim$_{\mathbb{F}}$ problem on Empty $+ k$v graphs is fixed-parameter tractable.*

Recall that the algorithm of Theorem 13 for the Ortho-Dim$_{\mathbb{F}}$ problem on Empty $+ k$v graphs enumerates all possible assignments of non-self-orthogonal vectors to the vertices of a given vertex cover. This approach is clearly not applicable when the field $\mathbb{F}$ is infinite. In order to extend the fixed-parameter tractability result to general fields and to obtain Theorem 14, we use the following definition inspired by an idea of [23].

▶ **Definition 15.** *Let $G = (V, E)$ be a graph, let $X \subseteq V$ be a vertex cover of G, and let $d \geq m \geq 1$ be integers. We define the graph $\mathsf{K} = \mathsf{K}(G, X, m, d)$ as follows. We start with $\mathsf{K} = G[X]$. Then, for every subset $S \subseteq X$ of size $m \leq |S| \leq d$, if there exists a vertex $v \in V \setminus X$ such that $S \subseteq N_G(v)$, then we add to K a new vertex v_S and connect it to all the vertices of S.*

The following lemma lists useful properties of the graph given in Definition 15 (recall Definition 8).

▶ **Lemma 16.** *Let $G = (V, E)$ be a graph, let $X \subseteq V$ be a vertex cover of G of size $|X| = k$, let $d \geq m \geq 1$ be integers, and let $\mathsf{K} = \mathsf{K}(G, X, m, d)$.*

1. *The set X forms a vertex cover of K.*
2. *The number of vertices in K is at most $k + \sum_{i=m}^{d} \binom{k}{i}$.*
3. *The graph K can be encoded in $\binom{k}{2} + \sum_{i=m}^{d} \binom{k}{i}$ bits.*
4. *For every field $\mathbb{F}$ with $m \leq m(\mathbb{F}, d)$, it holds that $\overline{\xi}_{\mathbb{F}}(G) \leq d$ if and only if $\overline{\xi}_{\mathbb{F}}(\mathsf{K}) \leq d$.*

Proof. Consider the graph $\mathsf{K} = \mathsf{K}(G, X, m, d)$ given in Definition 15. Since X is a vertex cover of G, it immediately follows from the definition that every edge of K is incident with a vertex of X, hence X is a vertex cover of K, as required for Item 1. It further follows that the vertex set of K consists of the vertices of X and at most one vertex per every subset $S \subseteq X$ of size $m \leq |S| \leq d$. Since the number of those subsets is $\sum_{i=m}^{d} \binom{k}{i}$, the number of vertices in K is at most $k + \sum_{i=m}^{d} \binom{k}{i}$, as required for Item 2. For Item 3, notice that to encode the graph K, it suffices to specify the adjacencies in $\mathsf{K}[X]$ and the existence of the vertex v_S in K for each $S \subseteq X$ of size $m \leq |S| \leq d$, hence K can be encoded in $\binom{k}{2} + \sum_{i=m}^{d} \binom{k}{i}$ bits.

We turn to the proof of Item 4. Let $\mathbb{F}$ be a field with $m \leq m(\mathbb{F}, d)$. Suppose first that $\overline{\xi}_{\mathbb{F}}(G) \leq d$, that is, there exists a d-dimensional orthogonal representation $(u_v)_{v \in V}$ of G over $\mathbb{F}$. We define a d-dimensional orthogonal representation of K over $\mathbb{F}$ as follows. First, we assign to each vertex $v \in X$ the vector u_v. It clearly holds that every two vertices of X that are adjacent in K are assigned orthogonal vectors. Next, for each vertex v_S of K with $S \subseteq X$ and $m \le |S| \le d$, there exists a vertex $v \in V \setminus X$ such that $S \subseteq N_G(v)$. We assign to v_S the vector u_v of such a vertex v. Notice that such a vector is orthogonal to all the vectors associated with the vertices of S, i.e., the neighbors of v_S in K. This gives us a d-dimensional orthogonal representation of K over $\mathbb{F}$, implying that $\overline{\xi}_{\mathbb{F}}(\mathsf{K}) \le d$.

For the other direction, suppose that $\overline{\xi}_{\mathbb{F}}(\mathsf{K}) \le d$. Letting V' denote the vertex set of K, there exists a d-dimensional orthogonal representation $(u_v)_{v \in V'}$ of K over $\mathbb{F}$. We define a d-dimensional orthogonal representation of G over $\mathbb{F}$ as follows. First, we assign to each vertex $v \in X$ the vector u_v. It clearly holds that every two vertices of X that are adjacent in G are assigned orthogonal vectors. We next extend this assignment to the vertices of the independent set $V \setminus X$ of G. Consider some vertex $v \in V \setminus X$, let $W = \mathrm{span}(\{u_{v'} \mid v' \in N_G(v)\})$, and notice that v may be assigned any non-self-orthogonal vector of $\mathbb{F}^d$ that lies in $W^\perp$. If $\dim(W) < m$, then by $m \le m(\mathbb{F}, d)$, there exists a non-self-orthogonal vector in $W^\perp$, which can be assigned to the vertex v. Otherwise, there exists a set of vertices $S \subseteq N_G(v)$ of size $m \le |S| \le d$ whose vectors form a basis of W, that is, $W = \mathrm{span}(\{u_{v'} \mid v' \in S\})$. By the definition of the graph K, it includes the vertex v_S, and its vector is orthogonal to the vectors $u_{v'}$ with $v' \in S$, and thus lies in $W^\perp$. This yields the existence of the desired vector for v, so we are done. ◀

With Lemma 16 at hand, we are ready to prove Theorem 14.

Proof of Theorem 14. The input of the Ortho-Dim$_{\mathbb{F}}$ problem on Empty $+\, k$v graphs consists of a graph G, a vertex cover X of G of size $|X| = k$, and an integer d. Consider the algorithm that given such an input acts as follows. If $d > k$ then the algorithm accepts. Otherwise, the algorithm calls an algorithm for the Ortho-Dim$_{\mathbb{F}}$ problem on the input (K, d), where $\mathsf{K} = \mathsf{K}(G, X, 1, d)$ is the graph given in Definition 15, and returns its answer. Note that we use here the assumption that the Ortho-Dim$_{\mathbb{F}}$ problem is decidable.

For correctness, observe first that the input graph G is $(k+1)$-colorable, as follows by assigning k distinct colors to the vertices of the vertex cover X and another color to the vertices of the independent set $V \setminus X$. It thus follows, using Claim 12, that $\overline{\xi}_{\mathbb{F}}(G) \le \chi(G) \le k+1$. Therefore, if $d > k$, then it holds that $\overline{\xi}_{\mathbb{F}}(G) \le d$, hence our algorithm correctly accepts. Otherwise, the algorithm calls an algorithm for Ortho-Dim$_{\mathbb{F}}$ on the input (K, d). The correctness of its answer follows from Item 4 of Lemma 16, which guarantees that $\overline{\xi}_{\mathbb{F}}(G) \le d$ if and only if $\overline{\xi}_{\mathbb{F}}(\mathsf{K}) \le d$.

We finally analyze the running time of the algorithm. On instances with $d > k$, the algorithm is clearly efficient. For instances with $d \le k$, by Item 2 of Lemma 16, the number of vertices in K is $O(k^d) \le O(k^k)$. Using the decidability of Ortho-Dim$_{\mathbb{F}}$, this implies that the running time of the algorithm is bounded by $f(k) \cdot n^{O(1)}$ for some computable function f, where n stands for the input size. Therefore, the Ortho-Dim$_{\mathbb{F}}$ problem on Empty $+\, k$v graphs is fixed-parameter tractable. ◀

In order to apply Theorem 14 to the real field $\mathbb{R}$, one has to show that the Ortho-Dim$_{\mathbb{R}}$ problem is decidable. We obtain this result using the problem of the existential theory of the reals, in which the input is a collection of equalities and inequalities of polynomials over the reals, and the goal is to decide whether there exists an assignment of real values to the variables satisfying all the constraints. In 1951, Tarski [31] proved that the problem is decidable. His result was strengthened in 1988 by Canny [10], who proved that it actually lies in the complexity class PSPACE. We derive the following simple consequence.

▶ **Proposition 17.** *The* ORTHO-DIM$_{\mathbb{R}}$ *problem lies in* PSPACE.

Proof. It is sufficient to show that the ORTHO-DIM$_{\mathbb{R}}$ problem is reducible in polynomial time to the problem of the existential theory of the reals, which lies in PSPACE [10]. Consider the reduction that given a graph $G = (V, E)$ and an integer d produces a collection P_G of polynomial constraints over the reals defined as follows. For each vertex $v \in V$, let $x_{v,1}, \ldots, x_{v,d}$ denote d variables associated with v. For each vertex $v \in V$, add to P_G the inequality $\sum_{i=1}^{d} x_{v,i}^2 \neq 0$, and for each edge $\{v, v'\} \in E$, add to P_G the equality $\sum_{i=1}^{d} x_{v,i} \cdot x_{v',i} = 0$. The reduction returns the collection P_G, which can clearly be computed in polynomial time. Observe that $\overline{\xi}_{\mathbb{R}}(G) \leq d$ if and only if there exists an assignment over the reals satisfying the constraints of P_G, implying the correctness of the reduction. ◀

Proposition 17 implies that the ORTHO-DIM$_{\mathbb{R}}$ problem is decidable. Using Theorem 14, we obtain the following corollary, which combined with Theorem 13, confirms Theorem 2.

▶ **Corollary 18.** *The* ORTHO-DIM$_{\mathbb{R}}$ *problem on* EMPTY $+ k$v *graphs is fixed-parameter tractable.*

4 Kernelization for d-Ortho-Dim$_{\mathbb{F}}$ Parameterized by Vertex Cover

We consider now the d-ORTHO-DIM$_{\mathbb{F}}$ problem for a fixed constant d and study its kernelizability when parameterized by the vertex cover number (recall Definition 1). We first leverage our discussion from the previous section to derive Theorem 3, namely, to show that for every field $\mathbb{F}$ and for every integer $d \geq 3$, the d-ORTHO-DIM$_{\mathbb{F}}$ problem on EMPTY $+ k$v graphs admits a kernel with $O(k^d)$ vertices and bit-size $O(k^d)$.

Proof of Theorem 3. Fix a field $\mathbb{F}$ and an integer $d \geq 3$. The input of d-ORTHO-DIM$_{\mathbb{F}}$ on EMPTY$+k$v graphs consists of a graph G and a vertex cover X of G of size $|X| = k$. Consider the algorithm that given such an input returns the pair (K, X), where $\mathsf{K} = \mathsf{K}(G, X, 1, d)$ is the graph from Definition 15. Since d is a fixed constant, the algorithm can be implemented in polynomial time. By Lemma 16, the set X forms a vertex cover of K, the graph K has $O(k^d)$ vertices and bit-size $O(k^d)$, and the instances (G, X) and (K, X) are equivalent. This completes the proof. ◀

For the real field $\mathbb{R}$, we prove Theorem 5, which improves on the kernel provided by Theorem 3 to $O(k^{d-1})$ vertices and bit-size $O(k^{d-1} \cdot \log k)$. We start with a couple of auxiliary lemmas.

▶ **Lemma 19.** *For every integer d, if a graph has a d-dimensional orthogonal representation over $\mathbb{R}$, then it has a d-dimensional orthogonal representation over $\mathbb{R}$, all of whose vectors have 1 as their first entry.*

Proof. The proof applies the probabilistic method. Let d be an integer, let $G = (V, E)$ be a graph, and set $n = |V|$. Suppose that there exists a d-dimensional orthogonal representation $(u_v)_{v \in V}$ of G over $\mathbb{R}$. Let $a \in [2n]^d$ be a random d-dimensional vector, such that each entry of a is chosen from $[2n]$ uniformly at random. We observe that for every fixed nonzero vector $u \in \mathbb{R}^d$, it holds that $\langle a, u \rangle = 0$ with probability at most $\frac{1}{2n}$. Indeed, letting $i \in [d]$ be an index with $u_i \neq 0$, for every fixed choice of the values of a_j with $j \in [d] \setminus \{i\}$, there is at most one value of a_i in $[2n]$ for which it holds that $\langle a, u \rangle = 0$. By the union bound, it follows that the probability that there exists a vertex $v \in V$ such that $\langle a, u_v \rangle = 0$ is at most $n \cdot \frac{1}{2n} = \frac{1}{2}$. In particular, there exists a vector $a \in [2n]^d$ satisfying $\langle a, u_v \rangle \neq 0$ for all $v \in V$. Let us fix such a vector a.

Now, let $M \in \mathbb{R}^{d \times d}$ be some orthonormal matrix (i.e., a matrix satisfying $M \cdot M^t = I_d$) whose first row is the vector a scaled to have Euclidean norm 1, i.e., $a/\|a\|$. We assign to each vertex $v \in V$ of the graph G the vector $M \cdot u_v$. Since M is orthonormal, it preserves inner products, hence this assignment forms a d-dimensional orthogonal representation of G over $\mathbb{R}$. Additionally, for every vertex $v \in V$, the first entry of the vector $M \cdot u_v$ is nonzero, because $\langle a, u_v \rangle \neq 0$. By scaling, one can obtain a d-dimensional orthogonal representation of G over $\mathbb{R}$, all of whose vectors have 1 as their first entry, as required. ◀

Before we state the next lemma, we need a brief preparation. For a field $\mathbb{F}$, a polynomial in $\mathbb{F}[x_1, \dots, x_n]$ is called homogeneous of degree d if each of its monomials has degree d. Note that the zero polynomial is homogeneous of degree d for every $d \geq 0$. A monomial is called multilinear if it forms a product of distinct variables, and a polynomial is called multilinear if it forms a linear combination of multilinear monomials. For example, the determinant of $d \times d$ matrices over a field $\mathbb{F}$, viewed as a polynomial on d^2 variables, is multilinear and homogeneous of degree d. Moreover, it is a linear combination of $d!$ monomials, each of which forms a product of d variables, one taken from each row of the matrix. Note that the dimension over $\mathbb{F}$ of the vector space of multilinear homogeneous polynomials of degree d in $\mathbb{F}[x_1, \dots, x_n]$ is $\binom{n}{d}$.

► **Lemma 20.** *For every integer d, there exists a multilinear homogeneous polynomial $p : \mathbb{R}^{d \times d} \to \mathbb{R}$ of degree $d-1$, defined on d^2 variables corresponding to the entries of a $d \times d$ matrix, such that for every matrix $M \in \mathbb{R}^{d \times d}$ whose first row is the all-one vector, it holds that $p(M) = 0$ if and only if there exists a nonzero vector in $\mathbb{R}^d$ that is orthogonal to all columns of M.*

Proof. For an integer d, consider the determinant polynomial $\det : \mathbb{R}^{d \times d} \to \mathbb{R}$. It is well known that for every matrix $M \in \mathbb{R}^{d \times d}$, it holds that $\det(M) = 0$ if and only if the columns of M span a subspace of dimension smaller than d, and that this condition is equivalent to the existence of a nonzero vector in $\mathbb{R}^d$ that is orthogonal to all columns of M. Recall that det is a multilinear polynomial, with each monomial being a product of d variables, each selected from a different row of the matrix. Let $p : \mathbb{R}^{d \times d} \to \mathbb{R}$ be the polynomial obtained from det by substituting 1 for the variables that correspond to the first row of the matrix, and observe that p is a multilinear homogeneous polynomial of degree $d-1$. Note that although p is defined on d^2 variables, it actually depends on only $d^2 - d$ of them. We finally observe that for every matrix $M \in \mathbb{R}^{d \times d}$ whose first row is the all-one vector, it holds that $p(M) = 0$ if and only if there exists a nonzero vector in $\mathbb{R}^d$ that is orthogonal to all columns of M. This completes the proof. ◀

We are ready to prove Theorem 5, providing a kernel with $O(k^{d-1})$ vertices and bit-size $O(k^{d-1} \cdot \log k)$ for the d-Ortho-Dim$_{\mathbb{R}}$ problem on Empty $+ k$v graphs for all integers $d \geq 3$.

Proof of Theorem 5. Fix an integer $d \geq 3$. The input of d-Ortho-Dim$_{\mathbb{R}}$ on Empty $+ k$v graphs consists of a graph $G = (V, E)$ and a vertex cover $X \subseteq V$ of G of size $|X| = k$. Consider the algorithm that given such an input acts in two phases, as described next.

In the first phase, the algorithm constructs the graph $G' = \mathsf{K}(G, X, d, d)$ given in Definition 15. Let V' denote the vertex set of G', and recall that every vertex $v_S \in V' \setminus X$ is associated with some set $S \subseteq X$ of size $|S| = d$ such that $N_{G'}(v_S) = S$. By Lemma 16, the set X is a vertex cover of G', and it holds that $|V'| \leq k + \binom{k}{d}$.

In the second phase, the algorithm constructs a graph G''. To do so, the algorithm first associates with each vertex $v \in X$ a d-dimensional vector x_v of variables over $\mathbb{R}$. Note that the total number of variables is $k \cdot d$. For each vertex $v_S \in V' \setminus X$, we apply Lemma 20 to obtain a multilinear homogeneous polynomial p_S of degree $d-1$, defined on the d^2 variables associated with the d neighbors of v_S in G' (which all lie in X). The polynomial p_S satisfies that for every assignment $M \in \mathbb{R}^{d\times d}$ to its variables with first row equal to the all-one vector, it holds that $p_S(M) = 0$ if and only if there exists a nonzero vector in $\mathbb{R}^d$ that is orthogonal to all columns of M. Let $P = \text{span}(\{p_S \mid v_S \in V' \setminus X\})$ denote the subspace spanned by the polynomials associated with the vertices of $V' \setminus X$. The algorithm proceeds by finding a set $Y \subseteq V' \setminus X$, such that the polynomials associated with the vertices of Y form a basis for P. Note that P is contained in the vector space of multilinear homogeneous polynomials of degree $d-1$ on $k \cdot d$ variables. Since the dimension of the latter is $\binom{k\cdot d}{d-1}$, recalling that d is a fixed constant, it follows that $|Y| \leq \binom{k\cdot d}{d-1} \leq O(k^{d-1})$. Letting $V'' = X \cup Y$, the algorithm returns the graph $G'' = G'[V'']$ and the set X, which forms a vertex cover of G'' because it forms a vertex cover of G'.

The number of vertices in G'' is $|V''| = |X| + |Y| \leq k + O(k^{d-1}) = O(k^{d-1})$. The number of edges in $G''[X]$ is at most $\binom{k}{2}$, and since the degree of each vertex of Y is d, the number of edges in G'' that involve vertices of Y is $d \cdot |Y|$. It follows that the total number of edges in G'' is at most $\binom{k}{2} + d \cdot O(k^{d-1}) \leq O(k^{d-1})$. Therefore, the number of bits required to encode the edges of G'' is at most $O(k^{d-1} \cdot \log |V''|) \leq O(k^{d-1} \cdot \log k)$, as required.

It is not difficult to verify that the algorithm can be implemented in polynomial time. Note that the set Y can be calculated in polynomial time by applying Gaussian elimination with $\binom{k\cdot d}{d-1}$ variables.

For the correctness of the algorithm, we shall prove that $\overline{\xi}_{\mathbb{R}}(G) \leq d$ if and only if $\overline{\xi}_{\mathbb{R}}(G'') \leq d$. By Item 4 of Lemma 16, using $m(\mathbb{R}, d) = d$ (see Remark 9), it holds that $\overline{\xi}_{\mathbb{R}}(G) \leq d$ if and only if $\overline{\xi}_{\mathbb{R}}(G') \leq d$. It thus suffices to show that $\overline{\xi}_{\mathbb{R}}(G') \leq d$ if and only if $\overline{\xi}_{\mathbb{R}}(G'') \leq d$.

It obviously holds that if $\overline{\xi}_{\mathbb{R}}(G') \leq d$ then $\overline{\xi}_{\mathbb{R}}(G'') \leq d$, because G' contains G'' as a subgraph. For the converse, suppose that $\overline{\xi}_{\mathbb{R}}(G'') \leq d$, that is, there exists a d-dimensional orthogonal representation of G'' over $\mathbb{R}$. By Lemma 19, it further follows that there exists a d-dimensional orthogonal representation $(u_v)_{v\in V''}$ of G'' over $\mathbb{R}$, such that every vector u_v has 1 as its first entry. For each vertex $v \in X$, assign the vector u_v to the vertex v as well as to the variables of the vector x_v associated with v. We will show that this assignment to the vertices of X can be extended to an orthogonal representation of G' over $\mathbb{R}$. Indeed, for every vertex $v_S \in Y$ of G'', the nonzero vector u_{v_S} is orthogonal to the vectors of the vertices of S. This implies, using Lemma 20 and the fact that the first entries of the vectors x_v with $v \in X$ are all 1, that the polynomial p_S vanishes on this assignment. Since the polynomials p_S with $v_S \in Y$ form a basis of the subspace P, it follows that all the polynomials p_S associated with the vertices $v_S \in V' \setminus X$ vanish on this assignment as well. Using Lemma 20 again, we obtain that for each vertex $v_S \in V' \setminus X$, there exists a nonzero vector that is orthogonal to the vectors of the vertices of S, and these are precisely the neighbors of v_S in G'. This gives us a d-dimensional orthogonal representation of G' over $\mathbb{R}$, which yields that $\overline{\xi}_{\mathbb{R}}(G') \leq d$, concluding the proof. ◀

5 Lower Bound

In this section, we prove our lower bound on the kernel size of the d-Ortho-Dim$_{\mathbb{F}}$ problem parameterized by the vertex cover number. We first present the gadget graph that will be used in the proof.

5.1 Gadget Graph

A key ingredient in the proof of our lower bound is the following lemma, which generalizes a construction of [29] (corresponding to the case of $d = 3$). Here, two nonzero vectors $u_1, u_2 \in \mathbb{F}^d$ are said to be proportional if there exists some $\alpha \in \mathbb{F}$ such that $u_1 = \alpha \cdot u_2$. The proof can be found in the full version of the paper.

▶ **Lemma 21.** *For an integer $d \geq 3$, let C_{2d} denote the cycle graph on $2d$ vertices, let x_0 and x_1 denote two adjacent vertices in the cycle, and let $H = \overline{C_{2d}}$ denote its complement graph.*
1. *There exists a proper d-coloring of H that assigns to x_0 and x_1 the same color.*
2. *There exists a proper d-coloring of H that assigns to x_0 and x_1 distinct colors.*
3. *For every field $\mathbb{F}$ and for every d-dimensional orthogonal representation of H over $\mathbb{F}$, the vectors assigned to x_0 and x_1 are either orthogonal or proportional.*

5.2 The d-Ortho-Dim$_{\mathbb{F}}$ Problem on Empty + kv Graphs

We prove the following lower bound on the size of any compression for the d-Ortho-Dim$_{\mathbb{F}}$ problem parameterized by the vertex cover number.

▶ **Theorem 22.** *For every field $\mathbb{F}$, every integer $d \geq 3$, and any real $\varepsilon > 0$, the d-Ortho-Dim$_{\mathbb{F}}$ problem on Empty + kv graphs does not admit a compression with bit-size $O(k^{d-1-\varepsilon})$ unless* NP $\subseteq$ coNP/poly.

Note that Theorem 22 confirms Theorem 4. Another immediate corollary is the following.

▶ **Corollary 23.** *For every field $\mathbb{F}$, the Ortho-Dim$_{\mathbb{F}}$ problem on Empty + kv graphs does not admit a polynomial compression unless* NP $\subseteq$ coNP/poly.

The starting point of the proof of Theorem 22 is the following theorem, which summarizes the lower bounds proved in [23, 24] on the size of any compression for the q-Coloring problem parameterized by the vertex cover number (see [24, Corollary 2]).

▶ **Theorem 24** ([23, 24])**.** *For every integer $q \geq 3$, the q-Coloring problem on Empty + kv graphs does not admit a compression with bit-size $O(k^{q-1-\varepsilon})$ unless* NP $\subseteq$ coNP/poly.

Equipped with Lemma 21, we relate the d-Coloring and d-Ortho-Dim$_{\mathbb{F}}$ problems parameterized by the vertex cover number, as stated below.

▶ **Theorem 25.** *For every field $\mathbb{F}$ and for every integer $d \geq 3$, there exists a linear-parameter transformation from d-Coloring on Empty + kv graphs to d-Ortho-Dim$_{\mathbb{F}}$ on Empty + kv graphs.*

Proof. Fix a field $\mathbb{F}$ and an integer $d \geq 3$. Consider an instance of the d-Coloring problem on Empty + kv graphs, namely, a graph $G = (V, E)$ and a vertex cover $X \subseteq V$ of G of size $|X| = k$. Our goal is to construct in polynomial time a graph $G' = (V', E')$ and a vertex cover $X' \subseteq V'$ of G' of size $|X'| = O(k)$, such that $\chi(G) \leq d$ if and only if $\overline{\xi}_{\mathbb{F}}(G') \leq d$.

To do so, we start with the graph G and add to it a clique of size d whose vertices are denoted by $z_1, \ldots, z_d$. Then, for each index $i \in [d]$ and each vertex $v \in X$, we add to the graph a copy $H_{i,v}$ of the complement $\overline{C_{2d}}$ of the cycle graph on $2d$ vertices, where two consecutive vertices of the cycle are identified with the vertices z_i and v. Note that we add here $d \cdot k$ such gadgets to the graph and that each of them involves $2d - 2$ new vertices. Let $G' = (V', E')$ denote the obtained graph, and let X' denote the set that consists of the vertices of X and the vertices that were added to G in the construction. The transformation returns the pair (G', X'). Since X is a vertex cover of G, the set $V \setminus X$ is an independent set of G. It follows that $V' \setminus X'$ is an independent set of G', hence X' is a vertex cover of G'. Its size satisfies $|X'| = k + d + d \cdot k \cdot (2d-2) = O(k)$, hence the transformation is linear-parameter. The transformation can clearly be implemented in polynomial time. For correctness, we shall prove that $\chi(G) \leq d$ if and only if $\overline{\xi}_{\mathbb{F}}(G') \leq d$.

Suppose first that $\chi(G) \leq d$, and consider some proper d-coloring of G with color set $[d]$. We extend this coloring to a d-coloring of G' as follows. First, for each $i \in [d]$, we assign the color i to the vertex z_i. Clearly, no edge that connects two of the vertices $z_1, \ldots, z_d$ is monochromatic. Next, for each $i \in [d]$ and $v \in X$, consider the vertices of the component $H_{i,v}$. The only vertices of $H_{i,v}$ that already received colors are z_i and v. By Lemma 21, this partial coloring of $H_{i,v}$ can be extended to a proper d-coloring of the whole gadget. Indeed, if z_i and v are assigned the same color then this follows from Item 1 of the lemma, and if z_i and v are assigned distinct colors then this follows from Item 2 of the lemma. This gives us a proper d-coloring of G', which implies using Claim 12 that $\overline{\xi}_{\mathbb{F}}(G') \leq \chi(G') \leq d$.

For the converse direction, suppose that $\overline{\xi}_{\mathbb{F}}(G') \leq d$, and consider a d-dimensional orthogonal representation $(u_v)_{v \in V'}$ of G' over $\mathbb{F}$. Since the vertices $z_1, \ldots, z_d$ form a clique in G', it follows that their vectors $u_{z_1}, \ldots, u_{z_d}$ are pairwise orthogonal. Since they are non-self-orthogonal, it follows that they are linearly independent, and thus span the entire vector space $\mathbb{F}^d$. For each $i \in [d]$ and $v \in X$, consider the vectors assigned by the given orthogonal representation to the vertices of the component $H_{i,v}$ in G', and apply Item 3 of Lemma 21 to obtain that the vectors u_{z_i} and u_v are either orthogonal or proportional. However, the vectors $u_{z_1}, \ldots, u_{z_d}$ span the vector space $\mathbb{F}^d$, hence for each $v \in X$, the nonzero vector u_v cannot be orthogonal to all of them. This yields that for each vertex $v \in X$, the vector u_v is proportional to exactly one of the vectors $u_{z_1}, \ldots, u_{z_d}$.

We define a d-coloring of G as follows. To each vertex $v \in X$, assign the color $i \in [d]$ for which u_v is proportional to u_{z_i}. Since the given orthogonal representation assigns orthogonal vectors to adjacent vertices, it follows that this coloring assigns distinct colors to adjacent vertices in X. Next, to each vertex $v \in V \setminus X$, assign a color from $[d]$ that does not appear on its neighbors. Notice that all the neighbors of v lie in X and were already colored, because X is a vertex cover of G. To see that such a color exists, recall that the vector u_v is nonzero and orthogonal to the vectors associated with its neighbors in X by the given orthogonal representation of G'. Since every such vector is proportional to one of $u_{z_1}, \ldots, u_{z_d}$, it follows that there exists some $i \in [d]$ for which no neighbor of v is associated with a vector proportional to u_{z_i}, yielding the existence of the desired color for v. This gives us a proper d-coloring of G and implies that $\chi(G) \leq d$, so we are done. ◀

We finally combine Theorems 24 and 25 to derive Theorem 22.

Proof of Theorem 22. Fix a field $\mathbb{F}$, an integer $d \geq 3$, and a real $\varepsilon > 0$. By Theorem 25, there exists a linear-parameter transformation from d-Coloring on Empty $+ k$v graphs to d-Ortho-Dim$_{\mathbb{F}}$ on Empty$+k$v graphs. Therefore, if d-Ortho-Dim$_{\mathbb{F}}$ on Empty$+k$v graphs admits a compression with bit-size $O(k^{d-1-\varepsilon})$, then by composing this compression with the given transformation, it follows that d-Coloring on Empty$+k$v graphs admits a compression with bit-size $O(k^{d-1-\varepsilon})$ as well. By Theorem 24, this implies that $\mathsf{NP} \subseteq \mathsf{coNP/poly}$, and we are done. ◀

References

1 Ziv Bar-Yossef, Yitzhak Birk, T. S. Jayram, and Tomer Kol. Index coding with side information. *IEEE Trans. Inform. Theory*, 57(3):1479–1494, 2011. Preliminary version in FOCS'06. doi:10.1109/TIT.2010.2103753.

2 Alexander Barg and Gilles Zémor. High-rate storage codes on triangle-free graphs. *IEEE Trans. Inform. Theory*, 68(12):7787–7797, 2022. doi:10.1109/TIT.2022.3191309.

3 Yitzhak Birk and Tomer Kol. Coding on demand by an informed source (ISCOD) for efficient broadcast of different supplemental data to caching clients. *IEEE Trans. Inform. Theory*, 52(6):2825–2830, 2006. Preliminary version in INFOCOM'98. doi:10.1109/TIT.2006.874540.

4 Hans L. Bodlaender, Rodney G. Downey, Michael R. Fellows, and Danny Hermelin. On problems without polynomial kernels. *J. Comput. Syst. Sci.*, 75(8):423–434, 2009. Preliminary version in ICALP'08. doi:10.1016/J.JCSS.2009.04.001.

5 Hans L. Bodlaender, Bart M. P. Jansen, and Stefan Kratsch. Kernelization lower bounds by cross-composition. *SIAM J. Discret. Math.*, 28(1):277–305, 2014. Preliminary version in STACS'11. doi:10.1137/120880240.

6 Jop Briët, Harry Buhrman, Debbie Leung, Teresa Piovesan, and Florian Speelman. Round elimination in exact communication complexity. In *Proc. of the 10th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC'15)*, pages 206–225, 2015. doi:10.4230/LIPICS.TQC.2015.206.

7 Jop Briët and Jeroen Zuiddam. On the orthogonal rank of Cayley graphs and impossibility of quantum round elimination. *Quantum Information & Computation*, 17(1&2):106–116, 2017.

8 Leizhen Cai. Parameterized complexity of vertex colouring. *Discret. Appl. Math.*, 127(3):415–429, 2003. doi:10.1016/S0166-218X(02)00242-1.

9 Peter J. Cameron, Ashley Montanaro, Michael W. Newman, Simone Severini, and Andreas J. Winter. On the quantum chromatic number of a graph. *Electron. J. Comb.*, 14(1):R81, 2007. doi:10.37236/999.

10 John F. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. of the 20th Annual ACM Symposium on Theory of Computing (STOC'88)*, pages 460–467, 1988. doi:10.1145/62212.62257.

11 Dror Chawin and Ishay Haviv. On the subspace choosability in graphs. *Electron. J. Comb.*, 29(2):P2.20, 2022. Preliminary version in EuroComb'21. doi:10.37236/10765.

12 Dror Chawin and Ishay Haviv. Improved NP-hardness of approximation for orthogonality dimension and minrank. *SIAM J. Discret. Math.*, 37(4):2670–2688, 2023. Preliminary version in STACS'23. doi:10.1137/23M155760X.

13 Bruno Codenotti, Pavel Pudlák, and Giovanni Resta. Some structural properties of low-rank matrices related to computational complexity. *Theor. Comput. Sci.*, 235(1):89–107, 2000. doi:10.1016/S0304-3975(99)00185-1.

14 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms.* Springer, 2015. doi:10.1007/978-3-319-21275-3.

15 Ronald de Wolf. Quantum Computing and Communication Complexity. Ph.D. Thesis, Universiteit van Amsterdam, 2001.

16 Fedor V. Fomin, Bart M. P. Jansen, and Michal Pilipczuk. Preprocessing subgraph and minor problems: When does a small vertex cover help? *J. Comput. Syst. Sci.*, 80(2):468–495, 2014. Preliminary version in IPEC'12. doi:10.1016/J.JCSS.2013.09.004.

17 Fedor V. Fomin, Daniel Lokshtanov, Saket Saurabh, and Meirav Zehavi. *Kernelization: Theory of Parameterized Preprocessing.* Cambridge University Press, 2019.

18 Alexander Golovnev and Ishay Haviv. The (generalized) orthogonality dimension of (generalized) Kneser graphs: Bounds and applications. *Theory of Computing*, 18(22):1–22, 2022. Preliminary version in CCC'21. URL: https://theoryofcomputing.org/articles/v018a022/.

19 Alexander Golovnev, Oded Regev, and Omri Weinstein. The minrank of random graphs. *IEEE Trans. Inform. Theory*, 64(11):6990–6995, 2018. Preliminary version in RANDOM'17. `doi:10.1109/TIT.2018.2810384`.

20 Ishay Haviv. Approximating the orthogonality dimension of graphs and hypergraphs. *Chic. J. Theor. Comput. Sci.*, 2022(2):1–26, 2022. Preliminary version in MFCS'19.

21 Gerald Haynes, Catherine Park, Amanda Schaeffer, Jordan Webster, and Lon H. Mitchell. Orthogonal vector coloring. *Electron. J. Comb.*, 17(1):R55, 2010. `doi:10.37236/327`.

22 Bart M. P. Jansen. The Power of Data Reduction: Kernels for Fundamental Graph Problems. Ph.D. Thesis, Utrecht University, 2013.

23 Bart M. P. Jansen and Stefan Kratsch. Data reduction for graph coloring problems. *Inf. Comput.*, 231:70–88, 2013. Preliminary version in FCT'11. `doi:10.1016/J.IC.2013.08.005`.

24 Bart M. P. Jansen and Astrid Pieterse. Optimal data reduction for graph coloring using low-degree polynomials. *Algorithmica*, 81(10):3865–3889, 2019. Preliminary version in IPEC'17. `doi:10.1007/S00453-019-00578-5`.

25 Bart M. P. Jansen and Astrid Pieterse. Optimal sparsification for some binary CSPs using low-degree polynomials. *ACM Trans. Comput. Theory*, 11(4):28:1–26, 2019. Preliminary version in MFCS'16. `doi:10.1145/3349618`.

26 László Lovász. On the Shannon capacity of a graph. *IEEE Trans. Inform. Theory*, 25(1):1–7, 1979. `doi:10.1109/TIT.1979.1055985`.

27 László Lovász. *Graphs and Geometry*, volume 65. Colloquium Publications, 2019.

28 László Lovász, Michael Saks, and Alexander Schrijver. Orthogonal representations and connectivity of graphs. *Linear Algebra Appl.*, 114–115:439–454, 1989. Special Issue Dedicated to Alan J. Hoffman.

29 René Peeters. Orthogonal representations over finite fields and the chromatic number of graphs. *Combinatorica*, 16(3):417–431, 1996. `doi:10.1007/BF01261326`.

30 Søren Riis. Information flows, graphs and their guessing numbers. *Electron. J. Comb.*, 14(1):R44, 2007. Preliminary version in NETCOD'06. `doi:10.37236/962`.

31 Alfred Tarski. *A Decision Method for Elementary Algebra and Geometry*. Number R-109 in RAND Research & Commentary Reports. RAND, Santa Monica, CA, 1951.

32 Chee-Keng Yap. Some consequences of non-uniform conditions on uniform classes. *Theor. Comput. Sci.*, 26(3):287–300, 1983. `doi:10.1016/0304-3975(83)90020-8`.