COLORFUL WORDS AND d-TVERBERG COMPLEXES

Florian Frick*[1] and R. Amzi Jeffs†[2]

[1]*Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA, and Institut für Mathematik, Freie Universität Berlin, Arnimallee 2, 14195 Berlin, Germany.*
frick@cmu.edu
[2]*Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA.*
amzij@cmu.edu

Submitted: May 17, 2022; Accepted: Dec 8, 2023; Published: TBD

Abstract. We give a complete combinatorial characterization of weakly d-Tverberg complexes. These complexes record which intersection combinatorics of convex hulls necessarily arise in any sufficiently large general position point set in $\mathbb{R}^d$. This strengthens the concept of d-representable complexes, which describe intersection combinatorics that arise in at least one point set. Our characterization allows us to construct for every fixed d a graph that is not weakly d'-Tverberg for any $d' \leqslant d$, answering a question of De Loera, Hogan, Oliveros, and Yang.

Keywords. Tverberg's theorem, word representable, d-representable, nerve, general position, strong general position, fully independent

Mathematics Subject Classifications. 52A35, 52C45

1. Introduction

Tverberg's theorem states that any set of $(d+1)(r-1)+1$ points in $\mathbb{R}^d$ can be partitioned into r parts so that the convex hulls of these parts share a common point. Over the last five decades Tverberg's theorem has inspired numerous extensions and variations [BS18]. Concurrently, intersection patterns of convex sets have been investigated through d-representable complexes, whose k-dimensional faces correspond to nonempty $(k+1)$-fold intersections among a family of convex sets in $\mathbb{R}^d$ [Tan13].

Recall that a simplicial complex is a downward-closed set system. A simplicial complex Δ consisting of subsets of $[n] \stackrel{\text{def}}{=} \{1, 2, \ldots, n\}$ is called *weakly d-Tverberg* if for any sufficiently large point set $P \subseteq \mathbb{R}^d$ in general position there are pairwise disjoint sets $P_1, \ldots, P_n \subseteq P$ such that for every $\sigma \subseteq [n]$ we have that $\bigcap_{i\in\sigma} \operatorname{conv}(P_i) \neq \varnothing$ if and only if $\sigma \in \Delta$. Tverberg's theorem

*Supported by NSF grant DMS 1855591, NSF CAREER grant DMS 2042428, and a Sloan Research Fellowship.
†Supported by the National Science Foundation through Award No. 2103206.

states that the $(r-1)$-simplex is weakly d-Tverberg for every $d \geqslant 1$, and $(d+1)(r-1)+1$ points in $\mathbb{R}^d$ are sufficient to guarantee the existence of $P_1, \dots, P_r$ as above.

De Loera, Hogan, Oliveros, and Yang [DLHOY21] initiated the study of *d-Tverberg complexes*, which are a restriction of weakly d-Tverberg complexes in which $P_1, \dots, P_n$ are required to partition all of P. Every (weakly) d-Tverberg complex is d-representable, and in [DLHOY21] it is shown that the converse fails for $d = 2$. Here we show that the converse fails with arbitrarily large dimension gap. More precisely, for every fixed $d \geqslant 1$ we construct graphs (which are 3-representable, as follows from more general results of Wegner and Perel'man [Tan13, Section 3.1]) that are not weakly d'-Tverberg for any $d' \leqslant d$. This answers a question raised by De Loera, Hogan, Oliveros, and Yang.

Theorem 1.1. *For every $d \geqslant 1$, there is a graph that is not (weakly) d'-Tverberg for any $d' \leqslant d$.*

Along the way to constructing these graphs, we give a complete combinatorial characterization of weakly d-Tverberg complexes. This characterization will involve finding words whose letters come from $[n]$, with certain alternating subwords corresponding to faces. Oliveros and Torres [OT21] were the first to use this approach in the study of d-Tverberg complexes, defining *general d-word-representable graphs* and proving that certain families of such graphs—for example, bipartite graphs—are d-Tverberg for appropriate values of d. Inspired by this strategy, we define the class of *d-colorfully representable complexes* (Definition 1.3) and connect their combinatorics to the geometry of finite point sets in $\mathbb{R}^d$. While Oliveros and Torres primarily worked constructively, we proceed in the opposite direction: our proof of Theorem 1.1 is based on combinatorially recognizing obstructions that prevent a graph from being d-Tverberg.

Below, a *word* is simply a finite list of symbols from some chosen alphabet. A *subword* is any word obtained by deleting some (possibly none) of the instances of letters in a word.

Definition 1.2. Let $\sigma \subseteq [n]$ be a set of size $r \geqslant 1$, and let $d \geqslant 1$. A *d-colorful word* on alphabet σ is a word W of length $(d+1)(r-1)+1$ such that for every $i \in [d+1]$ the restriction of W to the indices $(i-1)(r-1)+1$ through $i(r-1)+1$ inclusive contains every letter from σ exactly once. We call these segments of consecutive indices *blocks* in W.

For example, the word $\underline{1243}\overline{42}\underline{13}\overline{42}\overline{134}$ is a 3-colorful word on the alphabet $\{1,2,3,4\}$. Above, we have under- and overlined the four blocks in which every letter must appear exactly once. We note that d-colorful words also appear in equivalent forms as "rainbow partitions" in [PÍ18] and "colorful Tverberg types" in [BLN17]. The relationship between the combinatorics of these words and the geometry of Tverberg partitions—a topic which both these papers deal with—will allow us to completely characterize weakly d-Tverberg complexes.

Definition 1.3. A simplicial complex $\Delta \subseteq 2^{[n]}$ is called *d-colorfully representable* if there is a word W on alphabet $[n]$ so that for every $\sigma \subseteq [n]$, we have $\sigma \in \Delta$ if and only if W contains a d-colorful subword on alphabet σ.

Theorem 1.4. *A simplicial complex is weakly d-Tverberg if and only if it is d-colorfully representable.*

The forward direction of Theorem 1.4 is a consequence of the existence of linearly ordered point sets in $\mathbb{R}^d$ whose minimal Tverberg partitions correspond to d-colorful words—for example the "stretched diagonal" of Bukh, Loh, and Nivasch [BLN17]. The converse follows from a recent universality result of Pór [PÍ8] on Tverberg partitions.

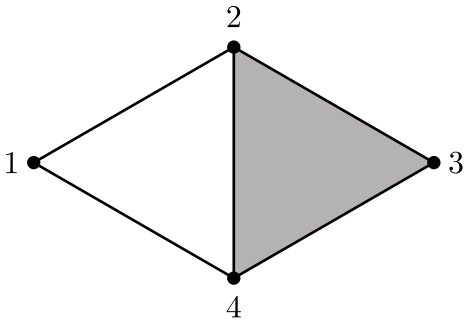

$$W = 121412432432$$
$$121 \quad 2$$
$$141 \quad 4$$
$$2432432$$

Figure 1.1: A simplicial complex and a word that 2-colorfully represents it.

Example 1.5. Consider the simplicial complex with facets 12, 14, and 234 on the vertex set $\{1, 2, 3, 4\}$. This complex is 2-colorfully represented by the word $W = 121412432432$. Figure 1.1 shows this complex and subwords of W which correspond to the faces. In fact, the figure only shows subwords corresponding to facets, but a d-colorful word on alphabet σ contains d-colorful words on every smaller alphabet (see Lemma 5.1), so it suffices to find the maximal d-colorful subwords.

Remark 1.6. We note that d-colorfully representable complexes naturally extend the definition of general d-word-representable graphs formulated by Oliveros and Torres [OT21]. A graph is *general d-word-representable* if one can find a word whose letters are the vertices of the graph, and whose alternating subwords of length $d+2$ correspond to edges of the graph. Hence general d-word-representable graphs are exactly the 1-skeleta of d-colorfully representable complexes. In particular, the two notions coincide for triangle-free graphs.

Depending on one's context, weakly d-Tverberg complexes may be a more natural tool than d-Tverberg complexes. Weakly d-Tverberg complexes are consistent with the usual approach in other Ramsey-type results that seek to classify the emergence of substructures locally instead of global behavior, which is captured by Tverberg complexes in the strong sense.

2. Background

Before proving our main theorems, we first introduce some notation and summarize relevant background material. The *nerve* of a collection $\mathcal{C} = \{C_1, \ldots, C_n\}$ of convex sets is the simplicial complex

$$\operatorname{nerve}(\mathcal{C}) \stackrel{\text{def}}{=} \left\{ \sigma \subseteq [n] \;\middle|\; \bigcap_{i \in \sigma} C_i \neq \varnothing \right\}.$$

We say that a simplicial complex $\Delta \subseteq 2^{[n]}$ is *partition induced* on a finite set P of points in $\mathbb{R}^d$ if there is a partition $\{P_1, \ldots, P_n\}$ of P so that the collection $\{\operatorname{conv}(P_1), \ldots, \operatorname{conv}(P_n)\}$ has

nerve equal to Δ. With this terminology we can rephrase the definition of (weakly) d-Tverberg complexes. Below, a point set $P \subseteq \mathbb{R}^d$ is in *general position* if for every $k \leqslant d+1$, no set of k points from P lie on a common $(k-2)$-dimensional affine subspace.

Definition 2.1 ([DLHOY21])**.** A simplicial complex $\Delta \subseteq 2^{[n]}$ is *d-Tverberg* if there is a constant C so that Δ is partition induced on any point set $P \subseteq \mathbb{R}^d$ in general position with at least C points.

Definition 2.2. A simplicial complex $\Delta \subseteq 2^{[n]}$ is *weakly d-Tverberg* if there is a constant C so that Δ is partition induced on a subset of any point set $P \subseteq \mathbb{R}^d$ in general position with at least C points.

We will often order point sets $P \subseteq \mathbb{R}^d$, and regard them as sequences. This allows us to tie the combinatorics of d-colorful words to the geometry of Tverberg partitions. A *Tverberg partition* of a finite point sequence is a collection of disjoint subsequences whose convex hulls share a common point. Note that we are abusing terminology slightly, since a Tverberg partition is not necessarily a partition of the entirety of our original sequence. A Tverberg partition is called *minimal* if deleting any point yields a collection of subsequences whose convex hulls no longer share a common point. The following multipartite generalization of Kirchberger's theorem, first proved by Pór [PÓ98], implies that every minimal Tverberg partition with r parts in $\mathbb{R}^d$ consists of at most $(d+1)(r-1)+1$ points in total (see also the work of Arocha, Bárány, Bracho, Fabila, and Montejano [ABB+09, Cor. 4]).

Proposition 2.3. *Let $P_1, \dots, P_r$ be disjoint subsets of $\mathbb{R}^d$. Then $\bigcap_{i=1}^r \operatorname{conv}(P_i) \neq \varnothing$ if and only if there exist $P_i' \subseteq P_i$ so that $\bigcap_{i=1}^r \operatorname{conv}(P_i') \neq \varnothing$ and $\bigcup_{i=1}^r P_i'$ contains at most $(d+1)(r-1)+1$ points.*

Given a collection of disjoint subsequences $P_1, \dots, P_r$ of a sequence P, one can form a word on alphabet $[r]$ by associating the points of P_i with the letter i for every $i \in [r]$, and writing down the sequence of these letters as they appear in P from start to end. We say that $P_1, \dots, P_r$ is a *d-colorful partition* if this word is d-colorful. In particular, a d-colorful partition will consist of exactly $(d+1)(r-1)+1$ points in total. Again note the slight abuse of terminology: a colorful partition is only a partition of a subset of the original sequence.

There are arbitrarily large point sequences in $\mathbb{R}^d$ whose minimal Tverberg partitions are exactly the d-colorful partitions. One example is a collection of points on the moment curve, chosen so that their coordinates increase heavily (see [PÍ18, page 4] for a discussion). Another example, provided by work of Bukh, Loh, and Nivasch [BLN17, Section 4], is the diagonal of the stretched grid.

A further important fact, established by Pór [PÍ18, Theorem 1.7], is that for every positive integer m, any sufficiently large point sequence in strong general position in $\mathbb{R}^d$ contains a subsequence of length m whose minimal Tverberg partitions are exactly the colorful ones. Strong general position means, informally, that affine hulls of subsets of the point set intersect generically (see Definition 2.4).

We wish to apply Pór's result to point sequences that are only in general position. To do this, it suffices to argue that for every n, any sufficiently large general position set must contain a

subset of n points in strong general position. The natural approach to proving this is to build up a strong general position subset of a general position set in an iterative, greedy way. However, one can become "stuck" when taking this approach: Doignon and Valette [DV77] showed that in dimension $d = 5$ there exists a finite set in strong general position that is maximal. The following section will circumvent this subtlety.

2.1. Large general position sets contain large strong general position subsets.

To start, we recall the formal definition of strong general position.

Definition 2.4. A finite set $P \subseteq \mathbb{R}^d$ is said to be in *strong general position* if for any $r \geqslant 1$ and any pairwise disjoint subsets $P_1, \ldots, P_r \subseteq S$ we have that

$$\dim \bigcap_{i=1}^{r} \operatorname{aff}(P_i) = \max \left\{ -1, d - \sum_{i=1}^{r} (d - \dim \operatorname{aff}(P_i)) \right\}. \qquad (\star)$$

When $\mathbb{R}^d$ is regarded as a subset of $\mathbb{R}\mathrm{P}^d$ and affine hulls are replaced by projective hulls in the above expression, we obtain the analogous notion of being *fully independent* [DV77].

Perhaps surprisingly, strong general position and full independence are not comparable notions. To see this, first observe that the set of vertices of a parallelogram in $\mathbb{R}^2$ is not in strong general position (opposite sides span parallel lines) but this set is fully independent. On the other hand, Doignon and Valette [DV77] constructed a 5-dimensional set with 12 points which is in strong general position, and is not fully independent. Moreover, they showed that such sets are exactly the maximal sets in strong general position. That is, a strong general position set is not fully independent if and only if it is inclusion-maximal among strong general position sets.

We wish to establish that every sufficiently large general position set contains a relatively large strong general position subset. It turns out that it is more straightforward to establish the analogous result for sets that are both in strong general position *and* fully independent.

Theorem 2.5. *For every $n \geqslant 1$ and $d \geqslant 1$, there exists $N = N(n, d) \geqslant 1$ so that every set of N points in general position in $\mathbb{R}^d$ contains a subset of n points that is in strong general position and fully independent.*

Proof. Fix a set $Q \subseteq \mathbb{R}\mathrm{P}^d$ consisting of d points whose projective hull is the hyperplane at infinity. Our starting point is the observation that if $P \subseteq \mathbb{R}^d$ is such that $P \cup Q$ is fully independent, then P is both fully independent and in strong general position. Indeed, P is fully independent because $P \cup Q$ is, and the only way that P can fail to be in strong general position is for P to contain disjoint subsets $P_1, \ldots, P_r$ so that $\bigcap_{i=1}^{r} \operatorname{proj}(P_i)$ is nonempty and contained in the hyperplane at infinity—if this intersection contains any real points, then the intersection of the corresponding affine hulls would have the same dimension, satisfying ($\star$). But then observe that $\operatorname{proj}(Q) \cap \bigcap_{i=1}^{r} \operatorname{proj}(P_i)$ would be equal to $\bigcap_{i=1}^{r} \operatorname{proj}(P_i)$, when it should in fact have dimension one less, contradicting the fact that $P \cup Q$ is fully independent.

With this observation in hand, it will suffice to argue that for every $n \geqslant 1$ and $d \geqslant 1$, there exists a constant $N = N(n, d)$ so that every set of N points in general position in $\mathbb{R}^d$ contains a subset P of size n so that $P \cup Q$ is fully independent. To establish this, we argue that there

is a constant $C(n,d)$ so that if $P \subseteq \mathbb{R}^d$ contains $n-1$ points and $P \cup Q$ is fully independent, then any general position set $X \subseteq \mathbb{R}^d$ of size $C(n,d)$ has $p \in X$ so that $P \cup \{p\} \cup Q$ is also fully independent. The choice $N(n,d) = \sum_{i=1}^{n} C(i,d)$ is then sufficient to prove the theorem inductively. From the first $N(n-1,d) = \sum_{i=1}^{n-1} C(i,d)$ points in general position, we extract a point set P of size $n-1$, and the remaining $C(n,d)$ points allow us to extract one further point so that we have n points in total whose union with Q is fully independent.

To prove the existence of $C(n,d)$, it suffices to argue that the set of points in $\mathbb{R}^d$ that can be added to $P \cup Q$ to obtain a larger fully independent set are those in the complement of a union of subspaces with positive codimension, where the number of such subspaces depends only on n and d. Doignon and Valette [DV77, Lemma 1] showed that if $Z \subseteq \mathbb{R}\mathrm{P}^d$ is fully independent and a point p avoids all proper subspaces of the form

$$\operatorname{proj}\left(Z_r \cup \bigcap_{i=1}^{r-1} \operatorname{proj}(Z_i)\right)$$

where $Z_1, \dots, Z_r$ are disjoint subsets of Z, then $Z \cup \{p\}$ is also fully independent. Setting $Z = P \cup Q$, we see that the number of subspaces that must be avoided is at most the number of tuples of disjoint subsets of $P \cup Q$, which depends only on n and d, establishing the result. □

The following theorem was proved for point sets in strong general position by Pór [PÍ18], and Theorem 2.5 allows us to state the result for point sets in general position.

Theorem 2.6 ([PÍ18, Thm. 1.7])**.** *Given $d, m, r \in \mathbb{N}$ with $r \geqslant 2$, there is $N = N(d,m,r) \in \mathbb{N}$ such that every sequence of length N in $\mathbb{R}^d$ in general position contains a subsequence of length m whose minimal Tverberg partitions are exactly the colorful ones.*

3. Characterizing Weakly d-Tverberg Complexes

Lemma 3.1. *Let $P = (p_1, p_2, \dots, p_m)$ be a sequence in $\mathbb{R}^d$ whose minimal Tverberg partitions are exactly the d-colorful ones. A simplicial complex $\Delta \subseteq 2^{[n]}$ is partition induced on P if and only if Δ is d-colorfully represented by a word W of length m.*

Proof. First suppose that Δ is partition induced on P, say by a partition $\{P_1, \dots, P_n\}$. Let W be the word on $[n]$ whose i-th letter is the unique j so that $p_i \in P_j$. In other words, W is the word obtained by labeling the points in P in sequence according to which part of the partition they belong to. We claim that W d-colorfully represents Δ.

To prove this, let $\sigma \subseteq [n]$. We aim to show that σ is a face of Δ if and only if W contains a d-colorful subword on alphabet σ. Note that σ is a face of Δ if and only if $\{P_j \mid j \in \sigma\}$ is a Tverberg partition. Since the minimal Tverberg partitions of P are exactly the colorful ones, we conclude that σ is a face of Δ if and only if there are $P'_j \subseteq P_j$ so that $\{P'_j \mid j \in \sigma\}$ is a d-colorful Tverberg partition. The letters corresponding to points in the various P'_j will form a d-colorful subword of W on alphabet σ, and conversely such a d-colorful subword allows us to construct appropriate P'_j. This proves the first half of the lemma.

For the converse, suppose that Δ is d-colorfully represented by some word W with length m. Define a partition $\{P_1, \ldots, P_n\}$ of P, where P_j consists of the points p_i so that the i-th letter of W is j. A similar analysis shows that Δ is partition-induced on P by this partition, proving the result. □

Theorem 1.4. *A simplicial complex is weakly d-Tverberg if and only if it is d-colorfully representable.*

Proof. Let Δ be a simplicial complex on the vertex set $[n]$. First suppose that Δ is weakly d-Tverberg. Choose a large sequence of points P in $\mathbb{R}^d$ whose minimal Tverberg partitions are the d-colorful partitions. By choosing P large enough, we can guarantee that Δ is partition-induced on a subsequence of P. A subsequence also has the property that its minimal Tverberg partitions are exactly the d-colorful ones, and so Lemma 3.1 implies that Δ is d-colorfully representable.

For the converse, suppose that Δ is d-colorfully representable by a word W of length m. By Theorem 2.6 every sufficiently large point sequence in general position in $\mathbb{R}^d$ contains a subsequence of length m whose Tverberg partitions are exactly the colorful ones. Lemma 3.1 then implies that Δ is partition-induced on this subsequence, so Δ is weakly d-Tverberg. This proves the result. □

4. A Graph That Is Not Weakly *d*-Tverberg

We begin by introducing some additional notation. Given a word W on alphabet σ, let $\Delta^d(W)$ be the simplicial complex

$$\Delta^d(W) \stackrel{\text{def}}{=} \{\tau \subseteq \sigma \mid W \text{ contains a } d\text{-colorful subword on } \tau\}.$$

In other words, $\Delta^d(W)$ is the simplicial complex[1] that is d-colorfully represented by W. If W is any word, let $\mathrm{red}(W)$ denote the *reduced word* formed by deleting any consecutive occurrences of the same letter. Observe that $\Delta^d(W) = \Delta^d(\mathrm{red}(W))$. Finally, if W is a word on alphabet σ, and $\tau \subseteq \sigma$, let $W(\tau)$ denote the *restriction* of W to τ, which is obtained by deleting all letters not in τ. Note that $\Delta^d(W(\tau))$ is the induced subcomplex of $\Delta^d(W)$ on vertex set τ.

Fix $d \geqslant 1$. For a positive integer n, let $K(n) \stackrel{\text{def}}{=} n(n-1)(d+2)^2+1$. Choose a fixed n large enough that $2^n > \binom{K(n)}{(d+2)^2}$, noting that this is possible because the latter quantity is polynomial in n. From here on, we let K denote $K(n)$ for our fixed choice of n. Now define a bipartite graph G_d on a vertex set $A \sqcup B$ as follows. First, A is simply a set of size n. The second part B consists of $2^n(2^K+1)$ vertices: for every $\sigma \subseteq A$, the part B contains 2^K+1 copies of a vertex whose neighborhood is exactly σ.

We aim to prove the following result, from which Theorem 1.1 will follow.

Theorem 4.1. *The graph G_d is not d'-colorfully representable for $d' \leqslant d$.*

[1] It is not a priori obvious that $\Delta^d(W)$ is a simplicial complex, but this follows from Lemma 5.1, which we explain in our concluding remarks.

The remainder of this section is dedicated to proving the theorem above. Since G_d contains an isomorphic copy of $G_{d'}$ as an induced subgraph for every $d' \leqslant d$, it will suffice to prove that G_d is not d-colorfully representable. For contradiction, assume that G_d is d-colorfully representable, and let W be a word that d-colorfully represents it. For the remainder of this section we treat W as fixed.

Note that W is a word whose letters come from $A \sqcup B$, where A and B are the two parts of the vertex set of G_d as defined above. Throughout this section we typically refer to elements of $A \sqcup B$ as "letters" since we are considering instances of them in the word W, though they are also vertices of the graph G_d.

Chunks. We will often examine $W(A)$, the restriction of W to the set of letters in A. A *chunk* in $W(A)$ is a maximal subword consisting of consecutive instances of the same letter. Note that a chunk could consist of only a single letter. We will regard the letters in a chunk interchangeably as letters in W and letters in $W(A)$.

Claim 1: $W(A)$ contains no more than $n(n-1)(d+2)$ distinct chunks. Hence $\mathrm{red}(W(A))$ has length at most $n(n-1)(d+2)$.

Proof of Claim 1: If $W(A)$ has more than $n(n-1)(d+2)$ chunks, then $\mathrm{red}(W(A))$ has length larger than $n(n-1)(d+2)$. Then some letter a appears more than $(n-1)(d+2)$ times in $\mathrm{red}(W(A))$. Since consecutive letters in $\mathrm{red}(W(A))$ are distinct, there are at least $(n-1)(d+2)$ gaps between the instances of a, filled by the remaining $n-1$ letters. But then by pigeonhole principle some remaining letter $a' \neq a$ appears in at least $d+2$ distinct gaps. Thus $\mathrm{red}(W(A))$ (and hence $W(A)$, and hence W) contains a d-colorful subword on alphabet $\{a, a'\}$. Hence aa' is an edge in G_d, contradicting the fact that A is an independent set in G_d.

Insertion patterns. Let $b \in B$. Observe that $W(A \cup \{b\})$ is a d-colorful representation of the induced subgraph of G_d on vertex set $A \cup \{b\}$. Let I_b be a word obtained by deleting as many instances of b from $W(A \cup \{b\})$ as possible without changing the graph that it d-colorfully represents. There could be many possibilities for I_b, but we will fix one from here on. We call our fixed I_b the *insertion pattern* of b.

Note that I_b determines the neighborhood of the vertex b in G_d. Also note that when forming I_b we never delete letters from A, and consequently we can regard the chunks of $W(A)$ canonically as subwords of I_b. Hence we will speak below of "chunks in I_b." Finally, observe that at most one instance of b occurs between two consecutive chunks in I_b—any further occurrences would be redundant.

Chunk containment. An instance of b that occurs in I_b is *contained in* a chunk if there are letters from that chunk both to its left and to its right. Note that each instance of b that occurs in I_b is contained in at most one chunk, and could be contained in no chunks if it lies between two adjacent chunks, or at the beginning or end of I_b.

Claim 2: Let $\mathbb{W}$ be the word obtained from $W(A)$ by shortening every chunk to have length at most $d+2$. Then for every $b \in B$, $\mathrm{red}(I_b)$ can be obtained by inserting non-consecutive copies of b into $\mathbb{W}$, and reducing the resulting word.

Proof of Claim 2: We first claim that for any $b \in B$, no chunk in I_b contains more than $d+1$ copies of b. Suppose for contradiction that some chunk consisting of copies of a letter $a \in A$ contains more than $d+1$ copies of b. None of these copies of b are consecutive (otherwise we could delete one without changing $\Delta^d(I_b)$) and hence this chunk contains an alternating subword on $\{a, b\}$ of length at least $2d+3 > d+2$. We now observe that we could delete a copy of b from this chunk without changing the complex $\Delta^d(I_b)$. Indeed, the remaining copies of b in the chunk would still form an alternating subword of length $d+2$ with a, and any alternating subwords involving b and a letter $a' \neq a$ can use at most one copy of b from this chunk, and so deleting one of the copies of b does not disrupt these alternating subwords either. There are no d-colorful subwords on alphabets of size three or more in I_b, since $\Delta^d(I_b)$ is a graph. We chose I_b so that deleting any copy of b would change $\Delta^d(I_b)$, and so we have arrived at a contradiction.

Since no chunk in I_b contains more than $d+1$ copies of b, at most $d+2$ letters from each chunk are preserved in $\mathrm{red}(I_b)$. In particular, $\mathrm{red}(I_b)$ restricted to A is a subword of $\mathbb{W}$, and $\mathrm{red}(I_b)$ can thus be obtained from $\mathbb{W}$ by inserting copies of b and then reducing. Since $\mathrm{red}(I_b)$ is reduced, we need not insert any consecutive copies of b. Hence the claim follows.

Claim 3: There exists $\sigma \subseteq A$ so that for every $b \in B$ with neighborhood σ, $\mathrm{red}(I_b)$ contains at least $(d+2)^2$ instances of the letter b.

Proof of Claim 3: Suppose for contradiction that every $\sigma \subseteq A$ admits some $b \in B$ with neighborhood σ and $\mathrm{red}(I_b)$ containing fewer than $(d+2)$ copies of b. Now, consider the process described in Claim 2. The word $\mathbb{W}$ has length at most $n(n-1)(d+2)^2$ by Claim 1. Hence each instance of b has at most $K = n(n-1)(d+2)^2+1$ places it can be inserted, and by Claim 2 the choices of where to insert b determine $\mathrm{red}(I_b)$. With at most $(d+2)^2$ instances of b to insert, we obtain at most $\binom{K}{(d+2)^2}$ distinct possibilities for $\mathrm{red}(I_b)$, even as σ varies over the $2^n > \binom{K}{(d+2)^2}$ possible subsets of A. This contradicts the fact that σ can be recovered from $\mathrm{red}(I_b)$, and the claim follows.

Claim 4: For any $\sigma \subseteq A$, there exist distinct letters $b, b' \in B$ whose neighborhood is σ, and which have $\mathrm{red}(I_b)$ and $\mathrm{red}(I_{b'})$ identical, up to replacing b by b'.

Proof of Claim 4: We again use Claim 2. There are at most 2^K ways to insert copies of a letter b into $\mathbb{W}$, but we constructed G_d so that there are more than 2^K vertices with neighborhood exactly σ. Hence at least two of them have the same reduced insertion patterns.

Claim 5: Let $\sigma \subseteq A$ be as in Claim 3, and let b and b' be as in Claim 4 with this choice of σ. Then W contains a d-colorful subword on the alphabet $\{b, b'\}$.

Proof of Claim 5: We have two relevant facts. First, $\mathrm{red}(I_b)$ and $\mathrm{red}(I_{b'})$ are the same up to replacing b by b'. Second, each of these reduced insertion patterns contains at least $(d+2)^2$ instances of b and b' respectively. By Claim 2, the longest consecutive sequence in $\mathrm{red}(I_b)$ which consists of b alternating with some letter $a \in A$ has length at most $d+2$. For each such sequence, let us delete all alternations except for one, obtaining a new word W_b. Note that W_b still has at least $d+2$ copies of b in it, and instances of letters from A in W_b are in one-to-one correspondence with chunks in $W(A)$.

Let $W_{b'}$ be obtained by performing the corresponding deletions in $\mathrm{red}(I_{b'})$. Each instance of b in W_b (respectively, b' in $W_{b'}$) can be associated to an instance of b (respectively, b') in our

original word W. Let W' be the subword of W on these instances of b and b'. Note that W' has length at least $2(d+2)$, and does not contain three consecutive instances of b (respectively b') since each instance of b is paired with an instance of b' which appears within or between the same chunks in W. Hence $\mathrm{red}(W')$ has length at least $d+2$, and is a d-colorful subword of W on alphabet $\{b, b'\}$. This proves the claim.

Summary. We argued that $\mathrm{red}(W(A))$ has bounded length, and used this fact to show that there must be some $\sigma \subseteq A$ so that every letter in B with neighborhood σ appears a large number of times, even when we forget the other letters in B and delete duplicate copies. Since there are many letters in B with neighborhood σ, we were able to argue that two of them must appear with the same patterns. In Claim 5 we showed that these identical patterns imply a d-colorful subword of W on these two letters. But this is a contradiction, since there is not an edge between these two letters in G_d. Hence G_d is not d-colorfully representable. This proves Theorem 4.1, and hence Theorem 1.1.

5. Concluding Remarks and Questions

We conclude by establishing some basic facts about d-colorfully representable complexes and (weakly) d-Tverberg complexes, and raising several questions. First, we show that any d-colorful word on alphabet σ contains a d-colorful subword on every alphabet $\tau \subseteq \sigma$. This implies that every word d-colorfully represents *some* simplicial complex; that is, given a word W on alphabet $[n]$, the set of σ so that W contains a d-colorful subword on σ actually does form a simplicial complex, as we would hope. One could prove this fact geometrically by way of Lemma 3.1—a colorful word on σ would give a Tverberg partition of a point set whose minimal Tverberg partitions are colorful, and deleting a part would yield Tverberg partition with fewer parts, then the lemma applied in converse would yield a colorful subword on a smaller alphabet. Below, we give a purely combinatorial proof.

Lemma 5.1. *If W is a d-colorful word on alphabet σ, then W contains a d-colorful subword on τ for any $\tau \subseteq \sigma$.*

Proof. It suffices to prove the result when $\tau = \sigma \setminus \{i\}$ for some $i \in \sigma$. Consider the $(d+1)$-many blocks of length r in W, which overlap at d-many indices. Whenever i occurs at an overlap index, let j be the letter immediately preceding it. We delete i and the occurrence of j in the block starting at the overlap index. Further, we delete all other occurrences of i (these did not appear at overlap indices in the original word).

We have deleted $d+1$ letters in total (one for each block that contains letter i in its "interior", and two for every pair of blocks which overlap with letter i) so we have obtained a word W' of length $(d+1)(r-2)+1$. We claim that W' is a d-colorful word on alphabet $\sigma \setminus \{i\}$. Consider any two consecutive blocks in the original word. After deletions, these become consecutive blocks in W'. We consider several cases.

If the original blocks overlapped with letter i, then the new blocks overlap at the letter immediately preceding i in the original word, and our choice to delete the second occurrence of this letter guarantees that the new blocks contain each letter from $\sigma \setminus \{i\}$ exactly once. Otherwise,

the original blocks did not overlap at letter i. In this case, the second block is altered only by deleting the letter i, and so contains every letter from $\sigma \setminus \{i\}$ exactly once. The first block is altered by deleting letter i, and possibly another letter j, if i is the first letter in the block. If we delete the letter j, our choice of j guarantees that it now appears at the first index in the new block. In either case, the first block contains every letter from $\sigma \setminus \{i\}$ and the result follows. □

Lemma 5.1 lets us quickly prove that every complex is d-colorfully representable for d large enough. When combined with Theorem 1.4 we obtain the corollary that every complex is weakly d-Tverberg for d large enough. The proof below is analogous to Oliveros and Torres' proof that every graph is general d-word-representable for d large enough [OT21, Proposition 2].

Proposition 5.2. *Every simplicial complex Δ with m facets is $(m+1)$-colorfully representable.*

Proof. Let $\sigma_1, \ldots, \sigma_m$ be the facets of Δ, and for each $i \in [m]$ let W_i be an $(m+1)$-colorful word on alphabet σ_i. Consider the word $W \stackrel{\text{def}}{=} W_1 W_2 \cdots W_m$ obtained by concatenating these colorful words. We claim that Δ is $(m+1)$-colorfully represented by W.

Clearly W contains an $(m+1)$-colorful word on alphabet σ for every $\sigma \in \Delta$. For the reverse containment, let F be any non-face of Δ. Suppose for contradiction that F is a face of $\Delta^{m+1}(W)$. Then W contains an $(m+1)$-colorful subword W_F on alphabet F. The word W_F has $m+1$ blocks of length $|F|$, each overlapping at their endpoints and containing the letters of F exactly once. Since F is not a face of Δ, none of the blocks in W_F can consist of letters from a single W_i. Thus the rightmost letter of the first block in W_F must appear after W_1 in W, and more generally the rightmost letter of the i-th block in W_F must appear after W_i in W. But W is comprised of only m-many W_i's, while W_F has $m+1$ blocks. Thus we have arrived at a contradiction, and we conclude that Δ is d-colorfully representable as desired. □

Corollary 5.3. *Every simplicial complex with m facets is weakly $(m+1)$-Tverberg.*

Recall that a cone is a simplicial complex whose facets all share at least one common vertex. We note that the classes of d-Tverberg complexes and weakly d-Tverberg complexes coincide for cones. Combined with the corollary above, we obtain that every cone is d-Tverberg for d large enough.

Proposition 5.4. *Let Δ be a simplicial complex that is a cone. Then Δ is d-Tverberg if and only if Δ is weakly d-Tverberg.*

Proof. Every d-Tverberg complex is also weakly d-Tverberg, so we only need to argue that if Δ is weakly d-Tverberg then it is d-Tverberg. Let P be a point set in $\mathbb{R}^d$ which is large enough that Δ is partition induced on some subset S of P. By adding the points of $P \setminus S$ to the part corresponding to the cone vertex in Δ, we obtain a partition of the entire point set P inducing the same nerve. □

Corollary 5.5. *Every cone with m facets is $(m+1)$-Tverberg.*

We conclude by explaining that being weakly d-Tverberg is a monotone property in the parameter d. We are grateful to Attila Pór for pointing out the proof of this result, which we originally posed as a question.

Proposition 5.6. *If $\Delta \subseteq 2^{[n]}$ is d-colorfully representable, then it is also $(d+1)$-colorfully representable. In particular, if Δ is weakly d-Tverberg then Δ is also weakly $(d+1)$-Tverberg.*

Proof. Let W be a word on $[n]$ which d-colorfully represents Δ. By applying a permutation, we may assume that if $i < j$, then the first instance of i in W appears before the first instance of j. Then consider the word W' which is obtained from W by concatenating all letters from $[n]$ in reverse order to the beginning of W. That is, $W' \stackrel{\text{def}}{=} n(n-1)(n-2)\cdots 321W$.

Observe that if W contains a d-colorful word on alphabet σ, then W' must also contain a $(d+1)$-colorful word on σ, since we may use our newly concatenated letters to add one more block to any d-colorful subword of W. Moreover, we claim that the converse holds. If W'' is a $(d+1)$-colorful subword of W' on alphabet σ, then the only letters of W'' that are not from W appear in the first block of W''. If the letter where the first and second blocks of W'' overlap comes from W, then we immediately obtain a d-colorful subword of W on alphabet σ. Otherwise, the letter where this overlap occurs is equal to the minimum of σ. But then an instance of this letter occurs in W before all instances of other letters in σ. Hence we also obtain a d-colorful subword of W on σ in this case. This proves the result. □

Our results above motivate several further questions, which we pose below. For the third question, note that it is at least a decidable problem to determine whether or not a simplicial complex is weakly d-Tverberg, since a word that d-colorfully represents it requires a bounded number of letters to represent each face. However, it is not known if one can algorithmically decide whether or not a simplicial complex is d-Tverberg. If recognizing d-Tverberg complexes is an undecidable problem, this would additionally provide a negative answer to our second question below.

Our construction of bipartite graphs that are not d-Tverberg used a very large number of vertices. It is likely that there are much smaller bipartite graphs that are not d-Tverberg, and for $d = 3$ Oliveros and Torres [OT21] proposed a candidate on ten vertices, motivating our fourth question.

1. Is every d-Tverberg complex also $(d+1)$-Tverberg?
2. Do the classes of d-Tverberg and weakly d-Tverberg complexes coincide?
3. For a fixed $d \geqslant 2$, what is the computational complexity of deciding whether or not a given simplicial complex is (weakly) d-Tverberg?
4. How many vertices are needed to construct a bipartite graph that is not 3-Tverberg?

Acknowledgements

We are grateful to Jesús De Loera, Deborah Oliveros, and Attila Pór for helpful discussions. We also extend our thanks to the anonymous referees for their feedback.

References

[ABB+09] Jorge L. Arocha, Imre Bárány, Javier Bracho, Ruy Fabila, and Luis Montejano. Very colorful theorems. *Discrete Comput. Geom.*, 42(2):142–154, 2009. doi:10.1007/s00454-009-9180-4.

[BLN17] Boris Bukh, Po-Shen Loh, and Gabriel Nivasch. Classifying unavoidable Tverberg partitions. *J. Comput. Geom.*, 8(1):174–205, 2017. doi:10.20382/jocg.v8i1a9.

[BS18] Imre Bárány and Pablo Soberón. Tverberg's theorem is 50 years old: a survey. *Bull. Amer. Math. Soc.*, 55(4):459–492, 2018.

[DLHOY21] Jesús A. De Loera, Thomas A. Hogan, Deborah Oliveros, and Dominic Yang. Tverberg-type theorems with altered intersection patterns (nerves). *Discrete Comput. Geom.*, 65(3):916–937, 2021. doi:10.1007/s00454-020-00241-9.

[DV77] J.-P. Doignon and G. Valette. Radon partitions and a new notion of independence in affine and projective spaces. *Mathematika*, 24(1):86–96, 1977. doi:10.1112/S0025579300008949.

[OT21] Deborah Oliveros and Antonio Torres. From word-representable graphs to altered Tverberg-type theorems. arXiv:2111.10038, November 2021.

[Pór98] Attila Pór. *Diploma Thesis*. PhD thesis, Eötvös University, Budapest, 1998.

[Pór18] Attila Pór. Universality of vector sequences and universality of Tverberg partitions. arXiv:1805.07197, May 2018.

[Tan13] Martin Tancer. Intersection patterns of convex sets via simplicial complexes: a survey. In *Thirty essays on geometric graph theory*, pages 521–540. Springer, 2013.