Discrete Comput Geom (2009) 42: 631–639
DOI 10.1007/s00454-008-9091-9

Dimension Gaps between Representability and Collapsibility

Jiří Matoušek · Martin Tancer

Received: 13 December 2007 / Revised: 13 May 2008 / Accepted: 14 May 2008 /
Published online: 20 June 2008

Abstract A simplicial complex K is called *d-representable* if it is the nerve of a collection of convex sets in $\mathbb{R}^d$; K is *d-collapsible* if it can be reduced to an empty complex by repeatedly removing a face of dimension at most $d-1$ that is contained in a unique maximal face; and K is *d-Leray* if every induced subcomplex of K has vanishing homology of dimension d and larger.

It is known that d-representable implies d-collapsible implies d-Leray, and no two of these notions coincide for $d \geq 2$. The famous Helly theorem and other important results in discrete geometry can be regarded as results about d-representable complexes, and in many of these results, "d-representable" in the assumption can be replaced by "d-collapsible" or even "d-Leray."

We investigate "dimension gaps" among these notions and construct, for all $d \geq 1$, a $2d$-Leray complex that is not $(3d-1)$-collapsible and a d-collapsible complex that is not $(2d-2)$-representable. In the proofs, we obtain two results of independent interest: (i) The nerve of every finite family of sets, each of size at most d, is d-collapsible; (ii) If the nerve of a simplicial complex K is d-representable, then K embeds in $\mathbb{R}^d$.

Keywords d-Representability · d-Collapsibility · Leray number · Simplicial complex · Convex set · Nerve

1 Introduction

d-Representability Helly's theorem [9] asserts that if $C_1, C_2, \dots, C_n$ are convex sets in $\mathbb{R}^d$, $n \geq d+1$, and every $d+1$ of the C_i have a common point, then $\bigcap_{i=1}^n C_i \neq \emptyset$. This famous theorem and many others in discrete geometry deal with

J. Matoušek (✉) · M. Tancer
Department of Applied Mathematics and Institute of Theoretical Computer Science (ITI), Charles University, Malostranské nám. 25, 118 00 Prague 1, Czech Republic
e-mail: matousek@kam.mff.cuni.cz

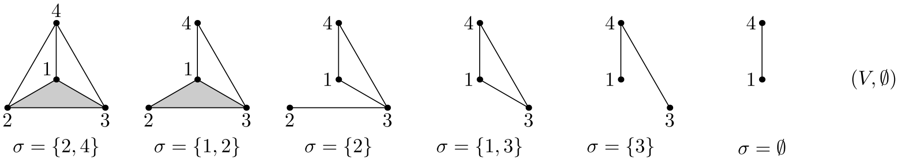

Fig. 1 An example of 2-collapsing

intersection patterns of convex sets in $\mathbb{R}^d$, and they can be restated using the notion of *d-representable* simplicial complexes.

We recall that the *nerve* $\mathsf{N}(\mathcal{S})$ of a family $\mathcal{S}=\{S_1,S_2,\dots,S_n\}$ is the simplicial complex with vertex set $[n]:=\{1,2,\dots,n\}$ and with a set $\sigma\subseteq[n]$ forming a simplex if $\bigcap_{i\in\sigma}S_i\neq\emptyset$. A simplicial complex K is *d-representable* if it is isomorphic to the nerve of a family of convex sets in $\mathbb{R}^d$ (all simplicial complexes throughout this paper are assumed to be finite).

In this language, Helly's theorem implies that a d-representable complex is determined by its d-skeleton. Other examples of theorems that can be seen as statements about d-representable complexes include the fractional Helly theorem of Katchalski and Liu [15], the colorful Helly theorem of Lovász ([16]; also see [4]), the (p,q)-theorem of Alon and Kleitman [1], and the Helly-type result of Amenta [3] (conjectured by Grünbaum and Motzkin). Among the deepest results concerning d-representable complexes, let us mention a complete characterization of their f-vectors[1] conjectured by Eckhoff and proved by Kalai [11, 12]. We also refer to [5, 6, 17] for more examples and background.

d-Collapsibility and d-Leray Complexes Wegner in his seminal 1975 paper [22] introduced d-collapsible simplicial complexes. To define this notion, we first introduce an *elementary d-collapse*. Let K be a simplicial complex, and let $\sigma,\tau\in\mathsf{K}$ be faces (simplices) such that

(i) $\dim\sigma\le d-1$,
(ii) τ is an inclusion-maximal face of K,
(iii) $\sigma\subseteq\tau$, and
(iv) τ is the *only* face of K satisfying (ii) and (iii).

Then we say that σ is a *d-collapsible face* of K and that the simplicial complex $\mathsf{K}':=\mathsf{K}\setminus\{\eta\in\mathsf{K}:\sigma\subseteq\eta\subseteq\tau\}$ arises from K by an elementary d-collapse. A simplicial complex K is *d-collapsible* if there exists a sequence of elementary d-collapses that reduces K to the empty complex $\emptyset$. Figure 1 shows an example of 2-collapsing.

Another related notion is a *d-Leray* simplicial complex, where K is d-Leray if every induced subcomplex of K (i.e., a subcomplex of the form $\mathsf{K}[X]:=\{\sigma\cap X:\sigma\in\mathsf{K}\}$ for some subset X of the vertex set $V(\mathsf{K})$) has zero homology (over $\mathbb{Q}$) in dimension d and larger.

[1] The *f-vector* of a d-dimensional simplicial complex K is the integer vector $(f_0,f_1,\dots,f_d)$, where f_i is the number of i-dimensional simplices in K.

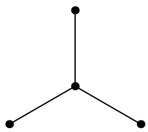

Fig. 2 A complex that is 1-collapsible but not 1-representable

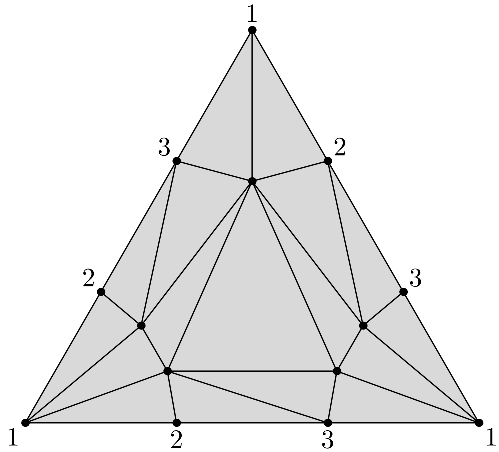

Fig. 3 The triangulation of the dunce hat [22]; vertices with the same numbers should be identified

Wegner [22] proved that every d-representable complex is d-collapsible and every d-collapsible complex is d-Leray. By inspecting proofs of several theorems about intersection patterns of convex sets in $\mathbb{R}^d$, i.e., about d-representable complexes, one can sometimes see that they actually use only the d-collapsibility, and thus they are valid for all d-collapsible complexes (good examples, among those mentioned earlier, are the fractional Helly theorem and the colorful Helly theorem).

With more work it has been shown that all of the results mentioned above and some others also hold for d-Leray complexes. For example, for Helly's theorem, this follows essentially from Helly's own topological generalization [10], for the (p,q)-theorem, this was proved in [2], and for the colorful Helly theorem and for Amenta's theorem, this was shown recently by Kalai and Meshulam [13, 14]. Kalai's characterization of f-vectors of d-representable complexes is also valid for the f-vectors of d-Leray complexes, showing that f-vectors cannot distinguish these classes.

These results indicate that the notions of d-representable, d-collapsible, and d-Leray are similar in some important respects. However, no two of them coincide. Figure 2 shows an example of a 1-collapsible complex that is not 1-representable. Wegner [22] noted that the well-known examples of 2-dimensional complexes that are contractible but not collapsible, such as suitable triangulations of the "dunce hat" (Fig. 3) or Bing's house (see, e.g., [8]), are 2-Leray but not 2-collapsible.

Results The goal of the present paper is to exhibit stronger differences among these notions; more precisely, to investigate "dimension gaps." We set

$$\begin{aligned}
\rho(\mathsf{K}) &:= \min\{d : \mathsf{K} \text{ is } d\text{-representable}\} && \text{("representability")},\\
\gamma(\mathsf{K}) &:= \min\{d : \mathsf{K} \text{ is } d\text{-collapsible}\} && \text{("collapsibility")},\\
\lambda(\mathsf{K}) &:= \min\{d : \mathsf{K} \text{ is } d\text{-Leray}\} && \text{("Leray number")}.
\end{aligned}$$

Theorem 1.1 (a) *For every $d \geq 1$, there exists a complex* K *with* $\gamma(\mathsf{K}) = d$ *and* $\rho(\mathsf{K}) = 2d - 1$ (*i.e.*, *d-collapsible and not* $(2d-2)$*-representable*).

(b) *For every $d \geq 1$, there exists a complex* K *with* $\lambda(\mathsf{K}) = 2d$ *and* $\gamma(\mathsf{K}) = 3d$ (*i.e.*, *$2d$-Leray and not* $(3d-1)$*-collapsible*).

In part (a), our example is the *nerve* of a d-dimensional simplicial complex L that is not embeddable in $\mathbb{R}^{2d-2}$. A well-known example of such L is the d-skeleton of the $(2d+2)$-dimensional simplex, due to Van Kampen [20, 21] and Flores [7]. The proof of Theorem 1.1(a) then follows immediately from the two propositions below, which may be of independent interest.

Proposition 1.2 *Let* L *be a simplicial complex such that the nerve* $\mathsf{N}(\mathsf{L})$ *is d-representable. Then* L *embeds in* $\mathbb{R}^d$, *even linearly.*

Proposition 1.3 *Let $\mathcal{F}$ be a finite family of sets, each of size at most d. Then the nerve* $\mathsf{N}(\mathcal{F})$ *is d-collapsible.*

For part (b) of Theorem 1.1, our example is a d-fold join of the dunce hat triangulation from Fig. 3.

Open Problems The main questions, which we unfortunately have not solved, are: Can representability be bounded in terms of collapsibility (formally, is there a function f_1 such that $\rho(\mathsf{K}) \leq f_1(\gamma(\mathsf{K}))$ for all K)? Can collapsibility be bounded in terms of the Leray number (formally, is there a function f_2 such that $\gamma(\mathsf{K}) \leq f_2(\lambda(\mathsf{K}))$ for all K)? Theorem 1.1 shows that $f_1(d) \geq 2d - 1$ and $f_2(2d) \geq 3d - 1$.

It is clear that our method cannot give a better lower bound for f_1 than Theorem 1.1(a), since every d-dimensional complex embeds in $\mathbb{R}^{2d-1}$. A 2-collapsible complex whose representability might perhaps be unbounded was noted by Alon et al. [2], namely, a finite projective plane (regarded as a simplicial complex, where the lines of the projective plane are the maximal simplices). More generally, any *almost-disjoint* set system is easily seen to be 2-collapsible, and it would be interesting to decide whether all almost-disjoint systems are d_0-representable for some constant d_0.

2 Representability of the Nerve and Embeddability

In this section, we will prove Proposition 1.2. First, we recall a classical lemma of Radon ([19]; also see, e.g., [6] or [17]) in the following form:

Lemma 2.1 *Let P be a set of affinely dependent points in $\mathbb{R}^d$. Then there exist two disjoint affinely independent subsets $A, B \subset P$ with* $\operatorname{conv}(A) \cap \operatorname{conv}(B) \neq \emptyset$.

We will also need the following result of a similar flavor:

Lemma 2.2 *Let A and B be finite subsets of $\mathbb{R}^d$. Suppose that there is a point* $x \in (\operatorname{conv}(A) \cap \operatorname{conv}(B)) \setminus \operatorname{conv}(A \cap B)$. *Then there exist disjoint affinely independent sets $A' \subseteq A$ and $B' \subseteq B$ such that* $\operatorname{conv}(A') \cap \operatorname{conv}(B') \neq \emptyset$.

Proof The proof is similar to the usual proof of Radon's lemma, only slightly more complicated.

We can write x as a convex combination of points of A:

$$x = \sum_{a \in A} \alpha_a a, \tag{1}$$

where $\alpha_a \geq 0$ for all $a \in A$ and $\sum_{a \in A} \alpha_a = 1$. Similarly,

$$x = \sum_{b \in B} \beta_b b, \tag{2}$$

where $\beta_b \geq 0$ for all $b \in B$ and $\sum_{b \in B} \beta_b = 1$. Let $K := A \cap B$, and let

$$K^+ := \{p \in K : \alpha_p > \beta_p\}, \qquad K^- := K \setminus K^+.$$

We define the sets $A_0 := A \setminus K^-$ and $B_0 := B \setminus K^+$ and note that $A_0 \cap B_0 = \emptyset$ and $A_0 \cup B_0 = A \cup B$. We claim that $\operatorname{conv}(A_0) \cap \operatorname{conv}(B_0) \neq \emptyset$; this will imply the lemma, since the desired affinely independent A' and B' can be obtained from A_0 and B_0 by removing redundant points.

For notational convenience, we extend the definition of α_p and β_p to all $p \in A \cup B$ by letting $\alpha_p = 0$ for $p \notin A$ and $\beta_p = 0$ for $p \notin B$. By subtracting (2) from (1) and rearranging we get

$$\sum_{p \in A_0} (\alpha_p - \beta_p) p = \sum_{p \in B_0} (\beta_p - \alpha_p) p.$$

All coefficients on both sides of this equation are nonnegative. Let us set $S := \sum_{p \in A_0} (\alpha_p - \beta_p)$. Since $\sum_{p \in A} \alpha_p = \sum_{p \in B} \beta_p = 1$, we also have $S = \sum_{p \in B_0} (\beta_p - \alpha_p)$. Moreover, since $x \notin \operatorname{conv}(K)$, at least one α_p with $p \in A \setminus K$ is nonzero, and thus $S \neq 0$. We set

$$y := \frac{1}{S} \sum_{p \in A_0} (\alpha_p - \beta_p) p = \frac{1}{S} \sum_{p \in B_0} (\beta_p - \alpha_p) p;$$

thus, y is expressed as a convex combination of points of A_0 and also as a convex combination of points of B_0. Hence, $\operatorname{conv}(A_0) \cap \operatorname{conv}(B_0) \neq \emptyset$, as claimed. □

Proof of Proposition 1.2 Let L be a simplicial complex such that $\mathsf{N}(\mathsf{L})$ is d-representable. This means that there exists a system $(C_\sigma : \sigma \in \mathsf{L})$ of convex sets in $\mathbb{R}^d$ such that, for every collection $\mathcal{M} \subseteq \mathsf{L}$ of simplices, we have $\bigcap_{\sigma \in \mathcal{M}} C_\sigma = \emptyset$ iff $\bigcap \mathcal{M} = \emptyset$.

For every $v \in V(\mathsf{L})$, we fix a point $p(v) \in \bigcap_{\tau \in \mathsf{L}: v \in \tau} C_\tau$ (this intersection is nonempty since $v \in \bigcap \{\tau \in \mathsf{L} : v \in \tau\}$).

This defines a mapping p: $V(\mathsf{L}) \to \mathbb{R}^d$. For every $\sigma \in \mathsf{L}$, we set

$$D_\sigma := \operatorname{conv}\bigl(p(\sigma)\bigr).$$

We claim that each D_σ is a simplex in $\mathbb{R}^d$ and that the D_σ form a geometric representation of K in $\mathbb{R}^d$. To this end, it suffices to verify that the set $p(\sigma)$ is affinely independent for every $\sigma \in \mathsf{L}$ and that $D_\sigma \cap D_\tau = D_{\sigma\cap\tau}$ for every two simplices $\sigma, \tau \in \mathsf{L}$.

First, let us suppose for contradiction that $p(\sigma)$ is affinely dependent for some $\sigma \in \mathsf{L}$. Then by Radon's lemma (Lemma 2.1) there are two disjoint affinely independent subsets $A, B \subset p(\sigma)$ with intersecting convex hulls. Then we have $A = p(\alpha)$ and $B = p(\beta)$ for disjoint simplices $\alpha, \beta \in \mathsf{L}$. But we have $p(v) \in C_\alpha$ for all $v \in \alpha$, hence $D_\alpha = \mathrm{conv}(p(\alpha)) \subseteq C_\alpha$, and similarly $D_\beta \subseteq C_\beta$. Then $C_\alpha \cap C_\beta \supseteq D_\alpha \cap D_\beta \neq \emptyset$, and this contradicts the assumption that the C_σ form a representation of $\mathsf{N}(\mathsf{L})$. So each $p(\sigma)$ is affinely independent.

Next, let $\sigma, \tau \in \mathsf{L}$. We clearly have $D_{\sigma\cap\tau} \subseteq D_\sigma \cap D_\tau$. To prove the reverse inclusion, we assume for contradiction that there is some $x \in (D_\sigma \cap D_\tau) \setminus D_{\sigma\cap\tau}$. Lemma 2.2 provides disjoint $\sigma' \subseteq \sigma$ and $\tau' \subseteq \tau$ with $D_{\sigma'} \cap D_{\tau'} \neq \emptyset$, and this is a contradiction as above. □

3 *d*-Collapsibility of the Nerve

Here we prove Proposition 1.3.

Let us assume that the ground set of $\mathcal{F}$ is $[n]$. Let us fix an arbitrary linear ordering $\leq$ on $\mathcal{F}$. The nerve $\mathbf{K} = \mathsf{N}(\mathcal{F})$ consists of all intersecting subfamilies of $\mathcal{F}$. For $i = 1, 2, \dots, n$, let $\mathbf{K}_i$ consist of all intersecting families $\mathcal{G} \in \mathbf{K}$ with $\min \bigcap \mathcal{G} = i$ (so the $\mathbf{K}_i$ form a partition of $\mathbf{K}$).

Let us consider a $\mathcal{G} \in \mathbf{K}_i$. Each of the elements $1, 2, \dots, i-1$ is excluded from $\bigcap \mathcal{G}$ by at least one $G \in \mathcal{G}$. Let us define the *minimal exclusion sequence* $\mathrm{mes}(\mathcal{G}) = (G_1, G_2, \dots, G_{i-1})$ as follows. First, we choose G_1 as the smallest set of $\mathcal{G}$ with $1 \notin G$. Having already defined sets $G_1, \dots, G_{j-1} \in \mathcal{G}$ (not necessarily all distinct), we define G_j as follows: If at least one of the sets among $G_1, \dots, G_{j-1}$ avoids the element j, we let G_j be such a G_k with the smallest possible k. In this case, we call G_j *old at* j. On the other hand, if all of $G_1, \dots, G_{j-1}$ contain j, then we let G_j be the smallest set of $\mathcal{G}$ not containing j, and we call it *new at* j.

Let $M(\mathcal{G}) \in \mathbf{K}_i$ be the family consisting of all sets G_j that occur in $\mathrm{mes}(\mathcal{G})$. In particular, for $i = 1$, we have $M(\mathcal{G}) = \emptyset$ for all $\mathcal{G} \in \mathbf{K}_1$. It is easily seen that we always have $|M(\mathcal{G})| \leq d$ (indeed, G_1 covers at most $d-1$ elements among $1, 2, \dots, i-1$, and only these elements may contribute G_j's distinct from G_1). We also note that $\mathrm{mes}(M(\mathcal{G})) = \mathrm{mes}(\mathcal{G})$.

We let $\mathbf{M}_i = \{M(\mathcal{G}) : \mathcal{G} \in \mathbf{K}_i\}$ and $\mathbf{M} = \bigcup_{i=1}^n \mathbf{M}_i$. The families in $\mathbf{M}$ will be the d-collapsible simplices we will use for d-collapsing the simplicial complex $\mathbf{K} = \mathsf{N}(\mathcal{F})$. We order them first by *decreasing* i, i.e., $\mathbf{M}_n$ comes first, then $\mathbf{M}_{n-1}$, etc., and within each $\mathbf{M}_i$, we order the families lexicographically by their minimal exclusion sequences. Let $\preceq$ denote this linear ordering of $\mathbf{M}$. This defines the sequence of elementary collapses.

Clearly, each simplex $\mathcal{G} \in \mathbf{K}$ contains at least one simplex of $\mathbf{M}$, namely, $M(\mathcal{G})$. It remains to verify that each $\mathcal{M} \in \mathbf{M}$ is contained in a unique maximal simplex in the simplicial complex obtained from $\mathbf{K}$ by collapsing all $\mathcal{N} \prec \mathcal{M}$.

We inductively define

$$\mathbf{K}_{\mathcal{M}} = \left\{\mathcal{H} \in \mathbf{K} \Big\backslash \bigcup_{\mathcal{N} \prec \mathcal{M}} \mathbf{K}_{\mathcal{N}} : \mathcal{M} \subseteq \mathcal{H}\right\};$$

as we will see, this is the set of all simplices removed from the current simplicial complex by collapsing $\mathcal{M}$.

We need to express $\mathbf{K}_{\mathcal{M}}$ as the set of all simplices of $\mathbf{K} \setminus \bigcup_{\mathcal{N} \prec \mathcal{M}} \mathbf{K}_{\mathcal{N}}$ that contain $\mathcal{M}$ and are contained in a suitable maximal simplex $T(\mathcal{M})$. As we will see, the desired $T(\mathcal{M})$ can be described as follows (here $\mathcal{M} \in \mathbf{K}_i$ and $\mathrm{mes}(\mathcal{M}) = (G_1, \ldots, G_{i-1})$):

$$T(\mathcal{M}) = \mathcal{M} \cup \{F \in \mathcal{F} : i \in F \text{ and } F > G_j \text{ for all } j \notin F \text{ such that } G_j \text{ is new at } j\}.$$

That is, $T(\mathcal{M})$ consists of $\mathcal{M}$ plus those sets of $\mathcal{F}$ that contain i and satisfy $\mathrm{mes}(\mathcal{M} \cup \{F\}) = \mathrm{mes}(\mathcal{M})$. Clearly $T(\mathcal{M}) \in \mathbf{K}$ and $\mathcal{M} \subseteq T(\mathcal{M})$. We also have $M(T(\mathcal{M})) = \mathcal{M}$.

We let $\mathbf{K}'_{\mathcal{M}} = \{\mathcal{H} \in \mathbf{K} \setminus \bigcup_{\mathcal{N} \prec \mathcal{M}} \mathbf{K}_{\mathcal{N}} : \mathcal{M} \subseteq \mathcal{H} \subseteq T(\mathcal{M})\}$, and by induction we prove that $\mathbf{K}'_{\mathcal{M}} = \mathbf{K}_{\mathcal{M}}$. This will show that $T(\mathcal{M})$ is indeed the unique maximal simplex containing $\mathcal{M}$. So we consider some $\mathcal{M} \in \mathbf{M}_i$ and assume that $\mathbf{K}'_{\mathcal{N}} = \mathbf{K}_{\mathcal{N}}$ for all $\mathcal{N} \prec \mathcal{M}$.

By just comparing the definitions we immediately obtain $\mathbf{K}'_{\mathcal{M}} \subseteq \mathbf{K}_{\mathcal{M}}$. For the reverse inclusion, let us consider an $\mathcal{H} \in \mathbf{K}_{\mathcal{M}}$ and, for contradiction, suppose that $\mathcal{H}$ contains a set $F \notin T(\mathcal{M})$. We will exhibit an $\mathcal{N} \in \mathbf{M}$ with $\mathcal{N} \subseteq \mathcal{H}$ and $\mathcal{N} \prec \mathcal{M}$; this will lead to the desired contradiction, since then we either have $\mathcal{H} \in \mathbf{K}_{\mathcal{N}}$ or $\mathcal{H}$ has been collapsed even earlier.

By the definition of $T(\mathcal{M})$, there can be two reasons for $F \notin T(\mathcal{M})$. First, it might happen that $i \notin F$. But then $\min \bigcap \mathcal{H} > i$, and therefore $\mathcal{N} := M(\mathcal{H}) \prec \mathcal{M}$. This leads to a contradiction as explained above, and hence we may assume $i \in F$.

Second, letting $\mathrm{mes}(\mathcal{M}) = (G_1, G_2, \ldots, G_{i-1})$, there may be some $j < i$, $j \notin F$ such that G_j is new at j and $F < G_j$ (we cannot have $F = G_j$ since we assumed that $F \notin \mathcal{M}$). Let j be the smallest possible with this property. We consider the family $\mathcal{G} = \{G_1, G_2, \ldots, G_{j-1}, F, G_{j+1}, \ldots, G_{i-1}\}$, which is in $\mathbf{K}_i$. Hence $\mathcal{N} = M(\mathcal{G})$ is in $\mathbf{M}_i$, and we also have $\mathcal{N} \subseteq \mathcal{H}$. It now suffices to verify that $\mathcal{N} \prec \mathcal{M}$. To this end, we check that the first j terms of the sequence $\mathrm{mes}(\mathcal{N})$ are $G_1, G_2, \ldots, G_{j-1}, F$, since then $\mathrm{mes}(\mathcal{N})$ is indeed lexicographically smaller than $\mathrm{mes}(\mathcal{M}) = (G_1, G_2, \ldots, G_{j-1}, G_j, \ldots)$.

Let us suppose that $\mathrm{mes}(\mathcal{N})$ agrees with $(G_1, G_2, \ldots, G_{j-1}, F)$ in the first $k-1$ terms, and we want to check that the kth terms agree as well. This is clear if $k < j$ and G_k is old at k in $\mathrm{mes}(\mathcal{M})$. If $k < j$ and G_k is new at k in $\mathrm{mes}(\mathcal{M})$, then $k \in F$ or $F > G_k$, for otherwise, we should have taken k instead of j, and hence the kth terms also agree in this case. Finally, for $k = j$, G_j is new at j in $\mathrm{mes}(\mathcal{M})$ by the assumption, so G_j is the smallest set in $\mathcal{M}$ not containing j. Then F is an even smaller such set, and so it is in the jth position of $\mathrm{mes}(\mathcal{N})$. This concludes the proof of Proposition 1.3.

4 Collapsibility Versus Leray Number

In this section we prove Theorem 1.1(b).

First, we recall the notion of join of two simplicial complexes K and L. First, assuming that the vertex sets $V(\mathsf{K})$ and $V(\mathsf{L})$ are disjoint, the *join* is the simplicial complex $\mathsf{K} * \mathsf{L} := \{\sigma \cup \tau : \sigma \in \mathsf{K}, \tau \in \mathsf{L}\}$ on the vertex set $V(\mathsf{K}) \cup V(\mathsf{L})$. If the vertex sets are not disjoint (as will be the case in our application below), we first take isomorphic copies of K and L with disjoint vertex sets and then we form the join as above.

The next lemma shows that the Leray number behaves nicely with respect to joins.

Lemma 4.1 *For every two nonempty simplicial complexes* K *and* L, *we have* $\lambda(\mathsf{K} * \mathsf{L}) = \lambda(\mathsf{K}) + \lambda(\mathsf{L})$.

Proof This is a simple consequence of the Künneth formula for joins

$$\tilde{H}_k(\mathsf{X} * \mathsf{Y}) = \bigoplus_{i+j=k-1} \tilde{H}_i(\mathsf{X}) \otimes \tilde{H}_j(\mathsf{Y}) \tag{3}$$

for any two simplicial complexes X and Y, where $\tilde{H}_k(\cdot)$ denotes the k-dimensional reduced homology group (over $\mathbb{Q}$). The Künneth formula in this form can easily be derived from [18], Example 4 (p. 349) and Exercise 3 (p. 373).

For notational convenience, we assume that $V(\mathsf{K}) \cap V(\mathsf{L}) = \emptyset$ and let $\lambda(\mathsf{K}) = k$ and $\lambda(\mathsf{L}) = \ell$. Then there exist $A \subseteq V(\mathsf{K})$ and $B \subseteq V(\mathsf{L})$ such that $\tilde{H}_{k-1}(\mathsf{K}[A]) \neq 0 \neq \tilde{H}_{\ell-1}(\mathsf{L}[B])$. Since $(\mathsf{K} * \mathsf{L})[A \cup B] = \mathsf{K}[A] * \mathsf{L}[B]$, (3) shows that $\tilde{H}_{k+\ell-1}((\mathsf{K} * \mathsf{L})[A \cup B]) \neq 0$, and thus $\lambda(\mathsf{K} * \mathsf{L}) \geq k + \ell$.

On the other hand, every induced subcomplex of $\mathsf{K} * \mathsf{L}$ has the form $\mathsf{K}[A] * \mathsf{L}[B]$ for some $A \subseteq V(\mathsf{K})$ and $B \subseteq V(\mathsf{L})$, and if $\tilde{H}_i(\mathsf{K}[A]) = 0$ for all $i \geq k$ and $\tilde{H}_j(\mathsf{L}[B]) = 0$ for all $j \geq \ell$, (3) gives $\tilde{H}_s(\mathsf{K}[A] * \mathsf{L}[B]) = 0$ for all $s \geq k + \ell$, thus showing $\lambda(\mathsf{K} * \mathsf{L}) \leq k + \ell$. □

We cannot say how $\gamma(\cdot)$ behaves under joins but we can do so for the following related quantity:

$$\gamma_0(\mathsf{K}) := \min\{d : \mathsf{K} \text{ has a } d\text{-collapsible face}\}.$$

Lemma 4.2 *For every two simplicial complexes* K, L, *we have* $\gamma_0(\mathsf{K} * \mathsf{L}) = \gamma_0(\mathsf{K}) + \gamma_0(\mathsf{L})$.

Proof Again we assume that $V(\mathsf{K}) \cap V(\mathsf{L}) = \emptyset$. It is easily checked that if σ is a k-collapsible face of K and τ is an ℓ-collapsible face of L, then $\sigma \cup \tau$ is a $(k+\ell)$-collapsible face of $\mathsf{K} * \mathsf{L}$, which shows that $\gamma_0(\mathsf{K} * \mathsf{L}) \leq \gamma_0(\mathsf{K}) + \gamma_0(\mathsf{L})$. On the other hand, *every* d-collapsible face of $\mathsf{K} * \mathsf{L}$ is of the form $\sigma \cup \tau$, $\sigma \in \mathsf{K}$, $\tau \in \mathsf{L}$, and one can check that σ is k-collapsible and τ is ℓ-collapsible for some k, ℓ with $k + \ell = d$. This gives the reverse inequality. □

Proof of Theorem 1.1(b) We let K_0 be the triangulation of the dunce hat in Fig. 3. We have $\lambda(\mathsf{K}_0) = 2 < \gamma_0(\mathsf{K}_0)$ according to Wegner [22], and actually $\gamma_0(\mathsf{K}_0) = 3$ because $\dim \mathsf{K}_0 = 2$. Then the join K of d copies of K_0 satisfies $\lambda(K) = 2d$ and $\gamma(\mathsf{K}) \geq \gamma_0(\mathsf{K}) = 3d$ by Lemmas 4.1 and 4.2. □

References

1. Alon, N., Kleitman, D.: Piercing convex sets and the Hadwiger Debrunner (p, q)-problem. Adv. Math. **96**(1), 103–112 (1992)
2. Alon, N., Kalai, G., Matoušek, J., Meshulam, R.: Transversal numbers for hypergraphs arising in geometry. Adv. Appl. Math. **130**, 2509–2514 (2002)
3. Amenta, N.: A short proof of an interesting Helly-type theorem. Discrete Comput. Geom. **15**, 423–427 (1996)
4. Bárány, I.: A generalization of Carathéodory's theorem. Discrete Math. **40**, 141–152 (1982)
5. Danzer, L., Grünbaum, B., Klee, V.: Helly's theorem and its relatives. In: Convexity. Proc. Symp. Pure Math., vol. 7, pp. 101–180. American Mathematical Society, Providence (1963)
6. Eckhoff, J.: Helly, Radon and Carathéodory type theorems. In: Gruber, P.M., Wills, J.M. (eds.) Handbook of Convex Geometry. North-Holland, Amsterdam (1993)
7. Flores, A.: Über n-dimensionale Komplexe die im R_{2n+1} absolut selbstverschlungen sind. Ergeb. Math. Kolloq. **6**, 4–7 (1932/1934)
8. Hatcher, A.: Algebraic Topology. Cambridge University Press, Cambridge (2001). Electronic version available at http://math.cornell.edu/~hatcher#AT1
9. Helly, E.: Über Mengen konvexer Körper mit gemeinschaftlichen Punkten. Jahresber. Dtsch. Math. Ver. **32**, 175–176 (1923)
10. Helly, E.: Über Systeme von abgeschlossenen Mengen mit gemeinschaftlichen Punkten. Monaths. Math. Phys. **37**, 281–302 (1930)
11. Kalai, G.: Characterization of f-vectors of families of convex sets in R^d. I: Necessity of Eckhoff's conditions. Isr. J. Math. **48**, 175–195 (1984)
12. Kalai, G.: Characterization of f-vectors of families of convex sets in R^d. II: Sufficiency of Eckhoff's conditions. J. Comb. Theory Ser. A **41**, 167–188 (1986)
13. Kalai, G., Meshulam, R.: A topological colorful Helly theorem. Adv. Math. **191**(2), 305–311 (2005)
14. Kalai, G., Meshulam, R.: Intersections of Leray complexes and regularity of monomial ideals. J. Comb. Theory Ser. A **113**(7), 1586–1592 (2006)
15. Katchalski, M., Liu, A.: A problem of geometry in R^n. Proc. Am. Math. Soc. **75**, 284–288 (1979)
16. Lovász, L.: Problem 206. Mat. Lapok **25**, 181 (1974)
17. Matoušek, J.: Lectures on Discrete Geometry. Springer, New York (2002)
18. Munkres, J.R.: Elements of Algebraic Topology. Addison-Wesley, New York (1984)
19. Radon, J.: Mengen konvexer Körper, die einen gemeinsamen Punkt enthalten. Math. Ann. **83**, 113–115 (1921)
20. Van Kampen, R.E.: Komplexe in euklidischen Räumen. Abh. Math. Sem. Hamburg **9**, 72–78 (1932)
21. Van Kampen, R.E.: Berichtigung dazu. Abh. Math. Sem. Hamburg **9**, 152–153 (1932)
22. Wegner, G.: d-collapsing and nerves of families of convex sets. Arch. Math. **26**, 317–321 (1975)