On Connectivity of the Facet Graphs of Simplicial Complexes

Ilan I. Newman* Yuri Rabinovich†

July 26, 2018

Abstract

The paper studies the connectivity properties of facet graphs of simplicial complexes of combinatorial interest. In particular, it is shown that the facet graphs of d-cycles, d-hypertrees and d-hypercuts are, respectively, $(d+1)$-, d- and $(n-d-1)$-vertex-connected. It is also shown that the facet graph of a d-cycle cannot be split into more than s connected components by removing at most s vertices. In addition, the paper discusses various related issues, as well as an extension to cell-complexes.

1 Introduction

Graphs of convex polytopes have been studied for many decades, starting with the classical Steinitz characterization of the graphs of 3-polytopes [12], experiencing a bust with the advance of the Simplex Method for Linear Programming, and continuing to draw a research effort in the modern era. See, e.g., the books [12, 23], and the survey [17] for many related results and open problems. One of the more well-known results in the area is Balinski Theorem [3] from 1961, stating that the graph (i.e., the 1-skeleton) of a d-polytope is $(d+1)$-vertex connected. This theorem and its various geometrical, topological and algebraic extensions have received a considerable attention, see e.g., [4, 5, 7, 2] for a very partial list of old and new related results. It has been extended to simple d-cycles in simplicial complexes in [14], where it is shown that the geometric realization of such graphs in $\mathbb{R}^{d+1}$ is generically rigid. See also the very recent [1] for an algebraic treatment of graphs of simple d-cycles.

In this paper we study the connectivity properties of *facet* graphs of simplicial complexes and, more generally, of cell complexes. That is, the facet graph $G_d(K)$ of a d-complex K, has a vertex for every d-face of K, and two such vertices are connected by an edge if the corresponding faces share a $(d-1)$-face. Since the facet graph of a convex d-polytope P is isomorphic to the graph of its dual polytope P^*, the facet graphs of convex polytopes do not require a separate study. This is not the case for simplicial complexes, where graphs and facet graphs differ significantly. For example, it is folklore that if a pure d-complex K is *strongly connected*, i.e., its facet graph $G_d(K)$ is connected, then the graph of K is d-connected. Obviously, this implication cannot be reversed, and connectivity of the graph of K implies nothing about the connectivity of its facet graph.

Motivated on one hand by the classical results about graphs of convex polytopes, and on the other hand by the recent surge of interest in the combinatorial theory of simplicial complexes, we study the facet graphs of the basic objects of this theory: simple d-cycles, d-hypertrees and d-hypercuts. These are the higher dimensional analogs of simple cycles, spanning trees and cuts in the complete graph K_n, respectively.

The paper proceeds as follows. We start with building our tools, and show that a simplicial complex induced by at most d simplices of dimension at most d, collapses to its $(d-2)$-skeleton. This lemma will be generalized in the last section, and one of its variants will be shown to be equivalent to the Homological Mixed Connectivity Theorem [13, 7], an elegant topological generalization of Balinski Theorem to higher-dimensional skeletons. We shall also discuss the duality of simple cycles and hypercuts in the complete simplicial complex on n vertices.

*Department of Computer Science, University of Haifa, Haifa, Israel. Email: ilan@cs.haifa.ac.il. This Research was supported by The Israel Science Foundation (grant number 862/10.)

†Department of Computer Science, University of Haifa, Haifa, Israel. Email: yuri@cs.haifa.ac.il. This Research was supported by The Israel Science Foundation (grant number 862/10.). Part of this research was done while this author visited Mittag-Leffler Institute, Stockholm

Next, we address the connectivity of the facet graphs of the basic combinatorial-topological objects. It comes, perhaps, as a little surprise that the facet graph of a simple d-cycle is $(d+1)$-connected. The facet graph of a d-hypertree T turns out to be d-connected, while the facet graph of a d-hypercut $G_d(H)$ is $(n-d-1)$-connected, where n is the number of vertices. All the results are tight.

In Section 4.3, inspired by [18], we study what happens to the facet graph $G_d(Z_d)$ of a simple d-cycle Z_d upon removal of s of its d-simplices. A somewhat unexpected conclusion is that the remaining part of $G_d(Z_d)$ has at most s components.

In the last section we study cell complexes with mild topological assumptions about the structure of the cells, and show, among other things, that the facet graphs of a simple d-cycles are still $(d+1)$-connected.

The paper employs only the very basic notions of Algebraic Topology (defined in the body of the paper), and should hopefully be accessible to anyone interested in Combinatorial Topology.

2 Preliminaries

2.1 Basic Standard Algebraic Topology Notions

We mostly use the basics of Homology Theory, beautifully presented in [19]. Throughout the paper we work over a fixed finite set (universe), identified with $[n]$, and an arbitrary fixed field $\mathbb{F}$. Many of our results hold if $\mathbb{F}$ is replaced by any Abelian group, however, in this paper we shall not pursue this direction.

A d-dimensional simplex, abbreviated as d-simplex, is a set $\sigma \subseteq [n]$. A simplicial complex K is a collection of simplices over $[n]$ closed under containment, i.e., if $\sigma \in K$, then so are all the faces of σ. As before, $\sigma \in K$ is called a face of K. The dimension of K is the largest dimension over all its faces. Some of the complexes discussed in this paper are *pure* d-dimensional complexes, i.e., all the maximal faces of K are all of the same dimension. Such faces are called *facets*.

The complete d-dimensional complex $K_n^d = \{\sigma \subset [n] \mid |\sigma| \leq d+1\}$ consists of all possible simplices over $[n]$ of dimension at most d.

Chains: Let K be a d-complex. We denote by $K^{(d)}$ the set of all d-faces of K. A *d-chain* of K is formal sum on *oriented* d-simplicies, where an orientation of a d-simplex σ is the sign of a permutation of the elements in σ. I.e., for two different ordering π_1, π_2 of the vertices of σ, $\sigma_{\pi_1} \equiv \sigma_{\pi_2}$ if $\pi_1 \equiv \pi_2$, and $\sigma_{\pi_1} \equiv -1 \cdot \sigma_{\pi_2}$ otherwise. For convenience, in this paper we assume that the vertex set is ordered, and fix a standard orientation $\sigma = (s_1, s_2, \ldots, s_{d+1})$, where $s_1 < \ldots < s_{d+1}$. Identifying the simplex σ with the standardly oriented σ, a d-chain C_d is simply $C_d = \sum_{\sigma_i \in K^{(d)}} c_i \sigma_i$ with $c_i \in \mathbb{F}$. The d-chains of K form a linear space over $\mathbb{F}$.

The support $\mathrm{Supp}(C_d)$ is the set of all d-simplices appearing in C_d with not-zero coefficients. The pure simplicial simplex $K(C_d)$ associated with C_d is the downwards closure of $\mathrm{Supp}(C_d)$ with respect to containment. A complex K is said to have *full* d-skeleton if $K^{(d)}$ contains all $\binom{n}{d+1}$ d-simplices.

The Boundary Operator: For a d-simplex $\sigma = (s_1, s_2, \ldots, s_{d+1})$, its d-boundary is defined as a $(d-1)$-chain $\partial_d(\sigma) = \sum_{i=1}^{d+1} (-1)^{i-1} (\sigma \setminus s_i)$, where $(\sigma \setminus s_i)$ is an oriented facet of σ obtained by the deletion of s_i. The case $d = 0$ requires a little special care; the (reduced) boundary of a 0-dimensional simplex, i.e., a single point p, is defined as $\partial_0(p) = 1$. I.e., the setting is that of the reduced homology. Taking a linear extension of ∂_d as above, one obtains a linear boundary operator ∂_d from the d-chains of K_n^d to the $(d-1)$-chains. (Observe that for a specific complex K, the d-chains of K are mapped by ∂_d to the $(d-1)$-chains of K). The key property of the boundary operators (due to the orientation rule above) is that $\partial_{d-1}\partial_d = 0$.

Using the vector form of d-chains, ∂_d is represented by a $\binom{n}{d} \times \binom{n}{d+1}$ matrix M_d whose rows are indexed by all $(d-1)$-simplices, and columns by d-simplices, and $\partial C_d = M_d C_d$. The entries of M_d are given by $M_d(\tau, \sigma) = \mathrm{sign}(\sigma, \tau)$, also written as $[\sigma : \tau]$, where $\mathrm{sign}(\sigma, \tau) = 0$ if τ is not a facet of σ, and $\mathrm{sign}(\sigma, \tau) = (-1)^{i-1}$ if τ is a facet of σ obtained by deletion of the i'th element in the ordered $V[\sigma]$. The requirement $\partial_{d-1}\partial_d = 0$ translates to $M_{d-1}M_d = 0$ for any $d = 1, 2, \ldots n$.

The simplicial complex $K(\partial_d C_d)$ will be often denoted by ΔC_d.

Cycles and Boundaries: A d-chain in $\ker(\partial_d)$ is called a *d-cycle*. The fact $\partial_d\partial_{d+1} = 0$ implies that if $C_d = \partial_{d+1}K$ then C_d is a d-cycle. Such a cycle is called a *d-boundary* of K. The boundary of any $(d+1)$-simplex is a d-cycle of size $d+2$, which is the smallest possible size of any non-zero d-cycle. The space of d-cycles supported on K is denoted by $\mathcal{Z}_d(K)$, and the space of d-boundaries supported on K is denoted by $\mathcal{B}_d(K)$. The factor space $\mathcal{Z}_d(K) \;/\; \mathcal{B}_d(K) \;=\; \tilde{H}_d(K)$ is called the *d-th (reduced) homology group of K*. The dimension of $\tilde{H}_d(K)$ is the d-th Betti number of K, denoted by $\tilde{\beta}_d(K)$.

A d-cycle Z is called *simple* if no other (non-zero) d-cycle is supported on $\mathrm{Supp}(Z)$. Sometimes, slightly abusing the notation, the supports of d-cycles will also be called d-cycles.

Cocycles, Coboundaries and Hypercuts: The coboundary operator δ^{d-1} is a linear operator adjoint to ∂_d. It is described by the left action of M_d, or, equivalently, by the right action of M_d^T. For standard terminological reasons, both the range and the domain of δ^{d-1} are called *cochains*, and denoted C^d and C^{d-1} respectively. In this paper, while retaining the standard notation, we shall not make any distinction whatsoever between d-chains and d-cochains[1]. Analogously to boundaries, $\delta^{d-1}C_{d-1}$ is the coboundary of the $(d-1)$-chain C_{d-1}.

A member in the kernel of δ^d is called *cocycle*. Thus, $\ker(\delta^d)$, is the space of d-cocycles. A d-cocycle Z^* is called a *d-hypercut* if it is simple, i.e., no other non-zero cocycle is supported on $\mathrm{Supp}(Z^*)$. Since $M_d^T M_{d-1}^T = (M_{d-1}M_d)^T = 0$, it holds that $\delta^d\delta^{d-1} = 0$, i.e., any coboundary $\delta^{d-1}C_{d-1}$ is a cocycle. Observe that in K_n^d every d-cocycle is also a d-coboundary (e.g., by a simple rank argument).

Hypertrees: A set A of d-simplices over $[n]$ is called *acyclic* if there are no d-cycles supported on A. Equivalently, A is acyclic if the columns vectors of M_n^d corresponding to its elements are linearly independent over $\mathbb{F}$. Thus, it immediately follows that all maximal acyclic sets $A \subseteq K_n^d$ have the same cardinality. A maximal acyclic set of d-simplices in K_n^d is called *d-hypertree*. Hypertrees were first introduced and studied by Kalai [16]. The cardinality of every d-hypertree is $\binom{n-1}{d}$ over any field.

For any d-simplex $\sigma \in K_n^d$ and any d-hypertee $T \subset K_n^d$, there exists a unique d-cycle of the form $Z_d = \sigma - \mathrm{Cap}_T(\sigma)$, where $\mathrm{Cap}_T(\sigma) = \sum_{\zeta_i \in T} c_i\zeta_i$ for some c_i's in $\mathbb{F}$. Observe that $\partial_d\sigma = \partial_d \mathrm{Cap}_T(\sigma)$.

A complex K with a full $(d-1)$-skeleton has $\tilde{H}_{d-1}(K) = 0$ if and only if K contains a d-hypertree. Indeed, containing a d-hypertree means that all $(d-1)$-cycles of the kind $\partial_d\sigma$, $\sigma \in K_n^d$ are $(d-1)$-boundaries of K; since they also span the cycle space of K, this is equivalent to $\mathcal{B}_{d-1}(K) = \mathcal{Z}_{d-1}(K)$.

Given the definitions above, it is clear that K_n^d defines a linear matroid $\mathcal{M}_d$ over $\mathbb{F}$, whose cycles correspond to supports of simple d-cycles as above, and whose maximal d-acyclic sets correspond to d-hypertrees. With slightly more effort one can show that the supports of the cycles of the *dual matroid* of $\mathcal{M}_d$, i.e., the cocycles of $\mathcal{M}_d$, correspond to d-hypercuts.

Using either the basic Matroid Theory (c.f., [21]), or the fact that d-hypecuts are simple d-coboundaries, one concludes that a d-subcomplex of K_n^d contains a d-hypertree iff it has a nonempty intersection with every d-hypercut. Moreover, for any d-hypercut H and $\sigma \in H$, there exists a d-hypertree T such that $T \cap H = \{\sigma\}$. Conversely, for any d-hypertree T and $\sigma \in T$, there exists a (unique) d-hypercut H such that $T \cap H = \{\sigma\}$. For a more detailed discussion on this see [20].

2.2 Facet Graphs

The facet graph $G(K) = G_d(K)$ of K, where K is a d-complex (or, with a slight abuse of notation, just a set of d-simplices, or even a d-chain), is a simple graph whose vertices correspond to the d-simplices in K, and two vertices form an edge if the corresponding d-simplices have a common $(d-1)$-dimensional face. Thus, each edge of $G(K)$ corresponds to a unique $(d-1)$-face of K. However, a $(d-1)$-face of K may correspond to many or none of the edges of $G(K)$.

With a slight abuse of notation, we shall speak of facet graph of d-cycles and d-hypercuts, although technically they are not complexes but chains.

[1]Usually, a d-chain C_d is regarded as a free $\mathbb{F}$-weighed sum of the elements of $K^{(d)}$, and a d-cochain C^d as a mapping from $K^{(d)}$ to $\mathbb{F}$.

3 Tools

3.1 The Basic Lemma

The following simple lemma will be at the core of many arguments to come.

Lemma 3.1 *Let D be a collection of d simplices on the same underlying set of vertices, and let $K(D)$ be the simplicial complex defined by D. Then, $\tilde{H}_{d-1}(K(D)) = 0$. Equivalently, every $(d-1)$-cycle supported on $K(D)$ is a $(d-1)$-boundary of $K(D)$.*

As it was pointed out to us by the anonymous referee, the Lemma is an immediate consequence of the Nerve Theorem (see, e.g, [8] for a combinatorial version). Assume, w.l.o.g., that no simplex in D contains another. Let $\mathcal{N}(K(D))$ denote the *nerve complex* of $K(D)$, whose vertices correspond to the elements of D, and faces correspond to the intersecting subsets of D. The Nerve Theorem asserts in that $K(D)$ and $\mathcal{N}(K(D))$ (or, rather, their geometric realizations) are homotopy equivalent. In particular, since homology groups are homotopy invariant, the two complexes have the same homology groups. The latter complex, however, has at most d vertices, and so either it is a $(d-1)$-simplex, or it has dimension $< d-1$. In either case, its $(d-1)$'th homology vanishes.

Here is an alternative combinatorial argument establishing the Lemma. We shall use the notion of d-*collapsibility* and its basic properties. Let K be a simplicial complex and $\tau \in K$. An elementary d-collapse of K with respect to such τ can be performed provided that $\dim(\tau) < d$, and that there exists a *unique* maximal (with respect to inclusion) face ξ of K containing τ. Unlike the more common *simplicial collapse* (see, e.g., [11]), τ and ξ are allowed to be identical. The corresponding elementary d-collapse is the action of removing from K all the faces ζ such that $\tau \subseteq \zeta \subseteq \xi$. The complex K is said to be d-collapsible to its subcomplex K' if it is possible to get from K to K' by a series of elementary d-collapses. Finally, K is called d-collapsible if it is possible to d-collapse it to $\emptyset$.

It is well known (see, e.g., [22]) and is easy to verify that K and K' as above have isomorphic homology groups in dimension $\geq d$. (Observe, however, that this is not true for dimension $< d$, as any $(d-1)$-complex is d-collapsible.) Thus, the following combinatorial lemma strengthens Lemma 3.1:

Lemma 3.2 *A simplicial complex $K(D)$ as in Lemma 3.1 is $(d-1)$-collapsible.*

The proof of the Lemma appears in the Appendix.

3.2 Duality between Cycles and co-Cycles in the Complete Complex K_n^{n-1}

In order to discuss the structure of the facet graphs of hypercuts, it will be useful to establish a duality between hypercuts and simple cycles. Such duality exists in Matroid Theory [21], and in a related, but a slightly more sophisticated form in the Algebraic Topology. It is at the core of the important Poincaré Duality and Alexander Duality. For a relevant combinatorial exposition of the latter see [9] and the references therein. In fact, Claim 3.3 below is a special case of the main result of this paper.

Let Σ be an $(n-1)$-simplex (seen as a complex) on the underlying space $[n]$. I.e., $\Sigma = K_n^{n-1}$. Define a correspondence between the $(k-1)$-chains and the $(r-1)$-cochains of Σ, where $k + r = n$, in the following way.

For $\sigma = \langle p_1, p_2, \ldots, p_k \rangle$ where $1 \leq p_1 < p_2 < \ldots < p_k \leq n$, let $\bar{\sigma} = \langle q_1, q_2, \ldots, q_r \rangle$, where $1 \leq q_1 < q_2 < \cdots q_r < n$, and q_j appears in $\bar{\sigma}$ iff it does not appear in σ. Set $s(\sigma) = \prod_{p_i \in \sigma} (-1)^{p_i - 1}$. The dual (signed) $(r-1)$-simplex of σ is defined by $\sigma^* = s(\sigma) \cdot \bar{\sigma}$. Extending this definition to chains and cochains, the dual of a $(k-1)$-chain (or cochain) $C = \sum c_\sigma \sigma$ is defined as a $(r-1)$-cochain (respectively, chain) $C^* = \sum c_\sigma \sigma^*$. The key fact about this correspondence is:

Claim 3.3 *[9]* $(\partial_{k-1} C)^* = \delta^{r-1} C^*$.

This leads to the following lemma, to be used in the Section 4.3, dedicated to hypercuts. Let k be a natural number in the range $[1, n]$, and let $k + r = n$.

Lemma 3.4 *The operator $*$ defines a 1-1 correspondence between simple $(k-1)$-cycles Z_k and $(r-1)$-hypercuts H_{r-1} of K_n^{n-1}, given by $Z_{k-1} \mapsto Z^*_{k-1} = H_{r-1}$. Moreover, the corresponding facet graphs $G_{k-1}(Z_{k-1})$ and $G_{r-1}(H_{r-1})$ are isomorphic.*

Proof. Observe that for any chain or cochain C of K_n^{n-1}, $C^{**} = (-1)^{\binom{n}{2}} C$, and hence the duality map $*$ is a 1-1 correspondence between the $(k-1)$-chains and the $(r-1)$-cochains. Since by Claim 3.3, it maps cycles to co-cycles, and co-cycles to cycles, it yields a 1-1 correspondence between $(k-1)$-cycles and $(r-1)$-co-cycles. Moreover, since it reverses containment, it yields a 1-1 correspondence between the minimal, i.e., simple, $(k-1)$-cycles and the minimal $(r-1)$-co-cycles, i.e., the $(r-1)$-hypercuts.

The isomorphism between the facet graphs of Z_{k-1} and $H_{r-1} = Z^*_{k-1}$ is given by the mapping $v_\sigma \mapsto v_{\bar{\sigma}}$ from $V[G_{k-1}(Z_{k-1})]$ to $V[G_{r-1}(H_{r-1})]$. Since a pair of $(k-1)$-simplices $\sigma, \zeta \in K_n^{n-1}$ share an $(k-2)$-face (i.e., are adjacent), if and only if they are both contained in a k-simplex $\xi \in K_n^{n-1}$, one concludes that σ, ζ are adjacent iff $\bar{\sigma}, \bar{\zeta}$ are. ■

4 Central Results

4.1 Connectivity of Cycles

We start with the $(d+1)$-connectivity of the simple d-cycles.

Theorem 4.1 *Let Z be a simple d-cycle, $d \geq 1$. Then, its facet graph $G(Z) = G_d(Z)$ is $(d+1)$-connected.*

Proof. Assume by contradiction that $G(Z)$ is not $(d+1)$-connected. Then, there exists a subset D of d-simplexes in $\mathrm{Supp}(Z)$, $|D| \leq d$, such that the removal of the vertices corresponding to D in $G(Z)$ disconnects this graph. Let $V_1, \ldots, V_r \subset V$, $r > 1$, be the vertex sets of the resulting connected components, and let $S_1, \ldots, S_r$ be the corresponding sets of d-simplices in $\mathrm{Supp}(Z)$. Finally, given that $Z = \sum c_j \sigma_j$, define d-chains $Z_i = \sum_{\sigma_j \in S_i} c_j \sigma_j$.

By definition of $G(Z)$, different Z_i's have disjoint $(d-1)$-supports. Since Z is a d-cycle, all the $(d-1)$-boundaries $C_i = \partial Z_i$ are supported on D. Since every $(d-1)$-boundary is a $(d-1)$-cycle, Lemma 3.1 applies to C_i's, implying, in particular, that there exist d-chain B_1 supported on D such that $\partial B_1 = C_1$. Consequently, the d-chain $Z_1 - B_1$ is a d-cycle, as $\partial(Z_1 - B_1) = C_1 - C_1 = 0$. Since Z_1 and B_1 have disjoint d-supports, $Z_1 - B_1 \neq 0$. Also, $Z_1 - B_1$ is supported on $S_1 \cup D$, a strict subset of d-faces of Z. This contradicts the fact that Z is simple cycle, concluding the proof. ■

Remark 4.2 *The above argument yields, in fact, a slightly more robust type of connectivity than stated. Recall that Lemma 3.1 applies not only to D as in the statement of Theorem 4.1, but also to a mixed set of d-simplices and $(d-1)$-simplices of a total size $\leq d$. Thus, the graph $G(Z)$ remains connected after removal of any r vertices and q edges (or even the edges of any q cliques induced by $(d-1)$-faces of Z), as long as $r + q \leq d$.*

Theorem 4.1 is tight, e.g., for d-pseudomanifolds, i.e., simple d-cycles, where every $(d-1)$-face is included in exactly two d-faces. In this case G_d is $(d+1)$-regular, and thus at most $(d+1)$-connected. For $d = 1$, all simple cycles are pseudomatifolds, and thus they are exactly 2-connected. For $d > 2$, other simple cycles exist, and it is not presently clear to us whether such cycles can be more than $(d+1)$-connected, and if yes, by how much.

Theorem 4.1 has an immediate implication on connectivity of the facet graphs of d-*biconnected* sets S of d-complexes. This interesting notion originates in Matroid Theory, and generalizes the 2-(edge)-connectivity in graphs. Let S be a set of d-simplices. Define the following relation on S:

Definition 4.3 *$\sigma \sim \zeta$ if there is a* simple *cycle Z supported on S containing both σ and ζ. Treating $\{\sigma, -\sigma\}$ as a simple cycle, it is also postulated that $\sigma \sim \sigma$. It is known from Matroid Theory [21] that $\sim$ is an equivalence relation. Call S* bi-connected *if all its d-simplices are $\sim$ equivalent.*

Corollary 4.4 *For biconnected S as above, $G_d(S)$ is $(d+1)$-connected.*

4.2 Connectivity of Hypertrees

Next, we establish the d-connectivity of d-hypertrees.

Theorem 4.5 *Let T be a d-hypertree in K_n^d, $d \geq 1,\ n \geq d+2$. Then, its facet graph $G(T) = G_d(T)$ is d-connected.*

Proof. As before, it suffices to show that for any subset X of d-simplices of T, $|X| \leq d-1$, the removal of the vertices corresponding to the X from $G(T)$ does not disconnect the graph. Assume by contradiction that $X \subseteq T,\ |X| \leq d-1$ and that removing the vertices corresponding to X from $G(T)$ results in r connected components with vertex sets $V_1, \ldots, V_r \subset V$, respectively. Let $\mathcal{S}_1, \ldots, \mathcal{S}_r$ be the corresponding sets of d-simplices in T. Let also $\mathcal{S}_0 = X$. We shall prove that $r = 1$, and thus X is non-separating.

For $i = 1, \ldots, r$ let us color all d-faces of $\mathcal{S}_i$, by color i. In particular, every d-face of T $\sigma \notin X$ has a (unique) associated color, while X is colorless.

The first step is to extend this coloring of T to all d-simplices in $K_n^d \setminus X$ in the following manner. Let $\sigma \in K_n^d \setminus T$ be a d-simplex. As explained in Section 2, there is a (unique) d-chain $\mathrm{Cap}_T(\sigma)$ supported on T satisfying $\partial(\mathrm{Cap}(\sigma)) = \partial(\sigma)$, namely $Z_\sigma = \mathrm{Cap}_T(\sigma) - \sigma$ is a simple d-cycle. Since any non-empty d-cycle is of size $\geq d+2$, and $|X| \leq d-1$, it follows that $\mathrm{Supp}(Z_\sigma) \setminus \{X \cup \sigma\} \subseteq T$ is not empty, and so $\mathrm{Cap}(\sigma)$ must contain some colored d-simplices in T. We claim that *all* such d-simplices must have the same color. This color will be assigned to σ.

Indeed, by Theorem 3.1, the graph $G(Z_\sigma)$ is $(d+1)$-connected. Since $|X| \leq d-1$, it remains connected after the removal of d vertices corresponding to $\{\sigma\} \cup X$. I.e., for any two colored d-simplices $\zeta, \tau \in \mathrm{Cap}(\sigma) \setminus X \subset T$, there exists a path $\xi_1, \xi_2, \ldots, \xi_s$ of (colored) d-simplices in $\mathrm{Cap}(\sigma) \setminus X \subset T$ where $\xi_1 = \zeta$, $\xi_s = \tau$, and every two consecutive ξ_i, ξ_{i+1} share a $(d-1)$-dimensional face. By definition of $\mathcal{S}_i$'s, if a pair of colored d-faces of T share a $(d-1)$-dimensional face, then they have the same color. Hence, all ξ_i's, and in particular ζ and τ, have the same color.

Having constructed a consistent extension of the coloring of $T \setminus X$ to the entire $K_n^d \setminus X$, it will be convenient to extend the definition of $\mathcal{S}_i$'s to contain all d-simplices of K_n^d colored i. The set $\mathcal{S}_0 = X$ remains unaffected.

The second step is to show that any two adjacent (i.e., sharing a $(d-1)$-face) colored d-simplices $\sigma_i, \sigma_j \in K_n^d \setminus X$ have the same color. While in $T \setminus X$ this is immediately implied by the definition of the color classes, in $K_n^d \setminus X$ a proof is required.

Assume by contradiction that σ_i and σ_j have different colors. Let τ_{ij} be the $(d-1)$-face they share, and let $\zeta_i \in \mathrm{Cap}(\sigma_i)$, $\zeta_j \in \mathrm{Cap}(\sigma_j)$ be d-simplices in T so that $\sigma_i \cap \zeta_i = \sigma_j \cap \zeta_j = \tau_{ij}$. Obviously there are such ζ_i, ζ_j by the definition of a cap. Now, on one hand, ζ_i and ζ_j are adjacent, and so, if both are colored, it must be the same color. In addition, this color must be the same as this of σ_i, σ_j by consistency of the color extension. Thus, if σ_i and σ_j differ in color, then at least one of ζ_i, ζ_j must belong to X. In particular, τ_{ij} is a $(d-1)$-facet of some d-simplex in X.

Let ψ be the (unique) $(d+1)$-simplex containing both σ_i and σ_j, and let $\Delta\psi$ denote the support of its boundary. In particular, $\sigma_i, \sigma_j \in \Delta\psi$. Consider $G_d(\Delta\psi)$, whose vertices are colored according to the colors of the corresponding d-simplices, and the vertices and the edges corresponding respectively to d- and $(d-1)$-faces of $K(X)$, are marked. As explained above, any two colored vertices of $G_d(\Delta\psi)$ connected by an unmarked edge must be of the same color. Thus, showing that any two colored vertices of this graph are connected by an unmarked path, will imply that there is only one color, contrary to the assumption.

Observe that $G_d(\Delta\psi)$ is isomorphic to K_{d+2}. Observe also that any d-simplex of X may cause the marking of a single vertex, or, alternatively, of a single edge of $G_d(\Delta\psi)$. Since $|X| \leq d-1$, this amounts to at most $d-1$ vertices and edges altogether. However, by an argument already used in the proof of Claim 3.1, removing all marked vertices and edges is not enough to disconnect K_{d+2}. Thus, any two colored vertices are indeed connected by an unmarked path, and thus the adjacent σ_i and σ_j must be of same color.

To sum up, we have shown so far that each d-simplex $\sigma \in K_n^d \setminus X$ has a well defined color, and that every two adjacent colored d-simplices have the same color. Recall that the goal is to show that there is only one color. Hence, to conclude the proof, it suffices to show that the facet graph $G_d(K_n^d)$ remains connected after removal of the vertices corresponding to the d-faces of X, i.e., that $G_d(K_n^d)$ is d-connected. One way of doing it is by observing that the d-skeleton of K_n^d is biconnected, and then applying Corollary 4.4. Since $G_d(K_n^d)$ is obviously connected, by transitivity of $\sim$, it suffices to check that any two adjacent σ, ζ are contained in a simple d-cycle. And indeed, they are contained in the boundary of the (unique) $(d+1)$-simplex containing both. ∎

Remark 4.6 *The graph $G_d(K_n^d)$ discussed above is known in Combinatorics as the graph of the hypersimplex polytope $\Delta_d(n)$, or as the graph of the $(d+1)$'th slice of n-hypercube, where two strings are adjacent iff they are at Hamming distance 2. See [10] for a relevant discussion of some connectivity properties of this important graph. It turns out that it is $(d+1)(n-d-1)$-connected. This can be deduced from a general result of Athanatos [2] stating (adapted to our terminology via duality of polytopes) that the facet graph of the d-skeleton of a convex polytope in $\mathbb{R}^m$ is $(d+1)(m-d)$-connected for $d > 1$. The proof in [2] involves rather heavy geometric arguments. It is an interesting open question whether an analogous result holds for simple m-cycles.*

To establish the tightness of Theorem 4.5, consider first the star $T = \{\sigma \in K_n^d \mid n \in \sigma\}$. This is a d-hypertree: it obviously spans all the d-simplices in K_n^d. On the other hand, it is acyclic, as every $\sigma \in T$ contains an exposed face, namely $(\sigma \backslash n)$. Now, consider, e.g., the hypertree $T' = T \backslash \{\sigma\} \cup \{\zeta\}$ where $\sigma = (1, \ldots, d, n)$, and $\zeta = (1, \ldots, d+1)$. It is easy to verify that T' is indeed a hypertree. Observe that the $(d-1)$-face $(1, \ldots, d)$ of ζ is exposed in T', while every other $(d-1)$-face of ζ is shared with a single d-simplex in T'. Hence, the vertex corresponding to ζ in $G_d(T')$ has degree d, implying that this graph is not $(d+1)$-connected.

Let us remark that the facet graph of a d-hypertree can be more than d-connected. E.g., when T is a star as above, $G_d(T)$ is obviously isomorphic to $G_d(K_{n-1}^{d-1})$ as in Remark 4.6, and thus it is $d(n-d-1)$-connected.

Consider the notion of r-edge-connectivity in graphs, i.e., G is r-edge connected if and only if every cut (of the complete graph) intersects $E[G]$ in at least r edges. This can be extended to higher dimensions as follows. Call a d-complex K *r-facet-connected* if for every d-hypercut H, $|H \cap K^{(d)}| \geq r$. It immediately follows from Theorem 4.5 that:

Corollary 4.7 *The facet graph $G_d(K)$ of a r-facet-connected d-complex K is $(d+r-1)$-connected.*

Proof. Assume by contradiction that there is a set of d-simplices $D = \{\sigma_1, \ldots, \sigma_{d+r-2}\}$ whose removal disconnect $G_d(K)$. Remove first $D' = \{\sigma_1, \ldots, \sigma_{r-1}\}$ from K. Since by assumption every hypercut has size at least r in K, every hypercut in $K \setminus D'$ is still nonempty, implying that $K \setminus D'$ contains a d-hypertree. But then $G_d(K \setminus D')$ is d-connected by Theorem 4.5, and hence it remains connected after the removal of the next $d-1$ simplices in $D \setminus D'$. ∎

Another implication of Theorem 4.5 is about the connectivity of complements of d-hypercuts.

Corollary 4.8 *Let $H \subset K_n^d$ be a d-hypercut, $d \geq 1$, and let $\overline{H} = K_n^d \setminus H$. Then, $G_d(\overline{H})$ is $(d-1)$-connected.*

Proof. Recall that H, being a hypercut, is critical with respect to hitting d-hypertrees, i.e., it hits every such T. Moreover, for any $\sigma \in H$ there exists a d-hypertree T_σ such that $T_\sigma \cap H = \{\sigma\}$. Hence, augmenting $\overline{H}$ by any $\sigma \notin \overline{H}$ makes it contain a d-hypertee T_σ. Hence by Theorem 4.5, the corresponding facet graph $G_d(T_\sigma)$ is d-connected. Removing the extra vertex corresponding to σ form this graph leaves us with $G_d(\overline{H})$, that must be $(d-1)$-connected. ∎

We conjecture that for $d \geq 2$ the above bound is not tight, and that it should involve n as well.

4.3 Connectivity of Hypercuts and Cocycles

Theorem 4.9 *Let $H \subset K_n^d$ be a d-hypercut, $d \geq 1$. Then, its facet graph $G_d(H)$ is $(n-d-1)$-connected.*

Proof. This is an immediate consequence of the duality result of Lemma 3.4, claiming that $H^* \subset K_n^{n-d-2}$ is a simple $(n-d-2)$-cycle with $G_d(H) \cong G_{n-d-2}(H^*)$, and an application of Theorem 4.1 to H^*. ∎

For tightness, consider the following example. Let $\tau = (1, 2, \ldots, d) \in K_n^d$ be a $(d-1)$-face. Then, the d-cochain $H_\tau = \sum_{i \in [n] \setminus \tau} \text{sign}(\tau \cup \{i\}; \tau) \cdot (\tau \cup \{i\})$ is a d-hypercut. Its graph $G_d(H_\tau)$ is an $(n-d)$-clique, which by convention is $(n-d-1)$-connected.

The facet graphs of cocycles that are not hypercuts behave very differently. For $d = 1$, the cocycles are precisely the graph-theoretic cuts, and so Theorem 4.9 applies. For $d \geq 3$, the facet graph of cocycle can be disconnected as exemplified by $H_\tau + H_{\tau'}$ as above, where the Hamming distance between τ and τ' *as sets* is at least 3. For $d = 2$, the answer is given by the following theorem:

Theorem 4.10 *The facet graph of a (non-empty) 2-cocycle Z^* of K_n^2, is 2-connected.*

Proof. It is immediate to verify the claim for $n \leq 4$, and thus we assume $n \geq 5$. To simplify the discussion, we use the duality between cocycles and cycles, as stated in Claim 3.3 and Lemma 3.4. Let Z be the dual chain of Z^*. Then, Z is a d-cycle, $d = n - 4 \geq 1$, of $K_n^{n-4} = K_{d+4}^d$, and $G_2(Z^*) \cong G_d(Z)$.

Now, Z, being a d-cycle, can be represented as $Z = \sum Z_i$, where each Z_i is a simple d-cycle, and $\text{Supp}(Z_i) \subseteq \text{Supp}(Z)$, and thus $G_d(Z_i) \subseteq G_d(Z)$. The key point of the argument is that K_{d+4}^d is a "narrow" place for d-cycles. We claim that for any two Z_i, Z_j as above, there exists a $(d-1)$-simplex τ belonging to $K(Z_i) \cap K(Z_j)$. The proof is by induction on d. The simplicity of Z_i, Z_j will not be used. For $d = 1$, one needs to show that any two cycles in K_5^1 have a common vertex. This is obvious. For general $d \geq 2$, let Z_1, Z_2 be two d-cycles. Since each of Z_1, Z_2 contains at least $(d+2)$ vertices in its support, and since $2(d+2) > d+4$, they share a common vertex v. Assuming $Z_1 = \sum c_\ell \sigma_\ell$, let $C_1 = \text{link}_v(Z_i) = \partial_d(\sum_{\sigma_\ell \ni v} c_\ell \sigma_\ell)$. Then, C_1 is a nonempty $(d-1)$-boundary, and hence a nonempty $(d-1)$-cycle. Moreover, keeping in mind that $\partial(Z_1) = 0$, we conclude that the vertex v does not appear in the vertex set $V(C_1)$. The same applies to the similarly defined C_2. Thus, C_1, C_2 are $(d-1)$-cycles on $(d+4)-1$ vertices. By induction hypothesis, they share a $(d-2)$-face τ'. The desired $(d-1)$-face τ is given by $\tau = (\tau' \cup v)$.

To conclude the proof of the Theorem, consider two d-simplices $\sigma, \zeta \in \text{Supp}(Z)$. It suffices to present a 2-connected (not K_2) subgraph of $G_d(Z)$ that contains the corresponding vertices v_σ, v_ζ. If σ, ζ fall in the same Z_i, we are done by the $(d+1)$-connectivity of $G_d(Z_i)$. Else, $\sigma \in Z_i$ and $\zeta \in Z_j$. If Z_i and Z_j share a common d-simplex ξ, then, using the equivalence relation of Def. 4.3 on $\text{Supp}(Z)$, $\sigma \sim \xi$, $\zeta \sim \xi \implies \sigma \sim \zeta$, and by Cor. 4.4, we are done again. Finally, even if Z_i and Z_j have no common d-simplices, by the above claim they still have a common $(d-1)$-simplex τ. Then, $V(G_d(Z_i)) \cap V(G_d(Z_j)) = \emptyset$, and there are 2 vertices in $V(G_d(Z_i))$ (corresponding to d-faces of Z_i that contain τ), and 2 vertices in $V(G_d(Z_j))$, that induce a K_4 in $G_d(Z)$. Since each of $G_d(Z_i), G_d(Z_j)$ is $(d+1)$-connected, the conclusion follows.

For tightness, consider the 2-cocycle of K_n^2 supported on all triangles containing either the edge e or the edge f, where e and f are disjoint. ■

Remark 4.11 *We have chosen to present this proof, because it provides more information about the structure of $G_2(Z^*)$. A simpler alternative proof would first reduce the problem to $n \leq 6$, by using the following argument.*

By definition of d-cocycles, the restriction of Z^* to any subset $S \subset [n]$ is a d-cocycle of $K_{|S|}^d$. Thus, for any pair of 2-simplices $\sigma, \zeta \in Z^*$, instead of considering the paths between the corresponding vertices in $G_2(Z^*)$, it suffices to consider them in $G_2(Z^*|_S)$, where S is the union of the vertex sets of σ and ζ.

5 Extensions and Refinements

In this section we further develop the results obtained in the previous section. We shall make a wider use of the homology-related notions of Algebraic Topology, which are luckily well suited for the discussion. This will make the presentation slightly more advanced, but the benefits will be apparent.

5.1 Facet Graphs of Simple Cycles: Beyond Connectivity

What follows is a direct continuation of Theorem 4.1.

Analogously to Klee's question about vertex graphs of convex polytopes [18][2], we ask what happens to the facet graph $G = G_d(Z)$ of a simple d-cycle Z, $d \geq 1$, after removal of a set of vertices $V_D \subseteq V[G]$ corresponding to a set of d-simplices $D \subseteq \mathrm{Supp}(Z)$. The following localization lemma provides an answer to this question in terms of the topological structure of $K(D)$.

Lemma 5.1 *Let Z be a simple d-cycle and $D \subset \mathrm{Supp}(Z)$. Assume that removal of V_D from $G_d(Z)$ creates $m > 1$ connected components. Then, $K(D)$ contains m $(d-1)$-cycles so that:* **(a)** *any two of them have disjoint $(d-1)$-supports;* **(b)** *any $m-1$ of them are linearly independent modulo the space of $(d-1)$-boundaries $\mathcal{B}_{d-1}(K(D))$, while all m of them are dependent.*

Proof. Proceeding as in the proof of Theorem 4.1, let $V_1, \ldots, V_m \subset V$, be the vertex sets of the resulting connected components, and let $S_1, \ldots, S_m$ be the corresponding sets of d-simplices in $\mathrm{Supp}(Z)$. Given that $Z = \sum c_j \sigma_j$, define the d-chains $Z_i = \sum_{\sigma_j \in S_i} c_j \sigma_j$, and $Z_D = \sum_{\sigma_j \in D} c_j \sigma_j$. Finally, define the $(d-1)$-cycles $C_i = \partial Z_i$ supported on $K(D)$. This will be the set of the desired cycles.

By definition of $G_d(Z)$, different $K(S_i)$'s have disjoint $(d-1)$-supports, and since C_i is supported on $K(S_i)$, the same applies to C_i's. This establishes **(a)**.

To prove **(b)**, consider, e.g., the first $m-1$ cycles, and assume by contradiction that for some d-chain Q_d on D, and for some (not all zero) coefficients $k_i \in \mathbb{F}$, it holds that $k_0 \cdot \partial Q_d + \sum_{i=1}^{m-1} k_i \cdot C_i = 0$. Let $Z' = k_0 \cdot Q_d + \sum_{i=1}^{m} k_i \cdot Z_i$ be a d-chain on $K(Z)$. Using the disjointness of $\mathrm{Supp}(Z_i)$'s, and the disjointness of $\cup_{i=1}^{m-1} \mathrm{Supp}(Z_i)$ and D, it is easily verified that Z' is neither 0 nor Z; the latter since $\mathrm{Supp}(Z') \cap \mathrm{Supp}(Z_m) = \emptyset$. Moreover, by definition of Z', it holds that $\partial Z' = 0$. This contradictions the simplicity of Z.

To see that $\{C_i\}_{i=1}^m$ are linearly dependent over $\mathcal{B}_{d-1}(K(D))$, recall that $C_i = \partial Z_i$, and that $Z_D + \sum_{i=1}^m Z_i = Z$. Therefore, $\partial Z_D + \sum_{i=1}^m C_i = \partial Z = 0$. ■

As an immediate corollary to the part **(b)** of the Lemma one gets:

Corollary 5.2 *Let Z be a simple d-cycle, and $D \subseteq \mathrm{Supp}(Z)$. The number of the connected components of $G_d(Z) \setminus D$ is at most $1 + \dim \tilde{H}_{d-1}(K(D)) = 1 + \tilde{\beta}_{d-1}(K(D))$.* ■

Part **(a)** of the lemma leads to:

Theorem 5.3 *Let $G_d(Z)$ be as above. Then, removing from $G_d(Z)$ any s vertices, may create at most s connected components.*

Proof. Recall that a $(d-1)$-cycle is of size at least $d+1$. Since the m different $(d-1)$-cycles in Lemma 5.1 are disjoint, they contain, altogether, at least $(d+1)m$ different $(d-1)$-simplices. On the other hand, let D be the set of d-simplices in $\mathrm{Supp}(Z)$ corresponding to the removed vertices. Then, the number of $(d-1)$ faces of $K(D)$ is at most $(d+1) \cdot |D| = (d+1)s$. Thus, $(d+1)s \geq (d+1)m$, and the conclusion follows. ■

The inequality $s \geq m$ is tight (for some s's), as shown by the following construction achieving $s = m$. Take a triangulation T of a d-pseudomanifold over $\mathbb{F}$, (i.e., every $(d-1)$-face of $K(T)$ is contained in exactly two d-simplices of T, and T supports a unique nonempty d-cycle), with the property that its facet graph $G_d(T)$ is bipartite.

An example of such a triangulation of the sphere is the d-cross-polytope, also known as the d-cocube, the dual polytope of the d-cube. Many more examples of such triangulations exists. See, e.g., [15] for a related discussion, proving among other things that if the graph of a triangulation T of a combinatorial d-cycle is $(d+1)$-colorable, then the facet graph of T is bipartite.

Obviously, for such T, taking D to be the set of all d-simplices corresponding to one color class of $G_d(T)$, results is decomposing the resulting $G_d(T) \setminus D$ into singletons.

The graph-theoretic property stated in the above theorem is called *toughness*. It has implications. E.g., using Tutte's criterion for existence of a perfect matching in a graph, one concludes via toughness that if a d-cycle Z is of even size, then $G(Z)$ has a perfect matching. For a survey of toughness see [6].

[2]Klee studied the following question: what is the maximum possible number of the connected components in the vertex-graph of a convex d-polytope, after the removal of m vertices? The answer: it at most 1 for $m \leq d$, 2 for $m = d+1$, and for a general m, at most the maximum possible number of facets in a convex d-polytope on m vertices.

5.2 Cycles in Cell Complexes

So far, we have discussed structures in simplicial complexes. In this section, we would like to discuss a class of axiomatically defined *cell complexes* that includes, in particular, simplicial complexes and convex polytopes (more precisely, the combinatorial abstraction preserving the structure of their faces). The methods and results obtained for simplicial complexes will be re-examined and generalized. We are mostly interested in the generalizations of Lemma 3.1 and Theorem 4.1, in particular we will generalize Balinski Theorem for such complexes.

Replacing simplices by *cells* with a specified combinatorial structure, and equipped with a boundary operator ∂, gives rise to cell complexes and their homology groups. The following axioms describe the structure of the cells. Notably, the standard assumption that the boundary of a cell is a pseudomanifold will be replaced here by a significantly weaker assumption that it is a simple cycle.

Formally, an *abstract cell complex* is a graded poset (partially ordered set) $\mathcal{P}$ whose elements will be called *open cells*. The order represents the cell-subcell relation. Let $\mathcal{C}^o \in \mathcal{P}$ be an open cell. The (closed) cell $K(\mathcal{C}^o) \subseteq \mathcal{P}$ is defined as the set of all elements that are dominated by $\mathcal{C}^o$ in $\mathcal{P}$, including $\mathcal{C}^o$. As $\mathcal{P}$ is a graded poset, the *rank* or *dimension* of its elements is well defined. Define also $\Delta\mathcal{C}^o \subset \mathcal{P}$, as the set of the proper faces of $\mathcal{C}^o$ in $\mathcal{P}$. $\mathcal{P}$ is formally extended to contain a unique minimal element of dimension -1, associated with the empty cell $\emptyset$.

A d-chain is a formal sum of $\mathbb{F}$-weighted open d-cells.

The axioms satisfied by $\mathcal{P}$ are as follows:

A1: *The poset $\mathcal{P}$ is a meet-semilattice. I.e., every two elements in it have a unique maximal element dominated by both.*

A2: *For every dimension $d \geq 0$, there is a boundary operator ∂_d mapping every open d-cell $\mathcal{C}^o$ to a nontrivial $(d-1)$-chain supported on $\Delta\mathcal{C}^o$. The (reduced) boundary of a 0-cell is 1, as before. It is further required that for any open d-cell $\mathcal{C}^o$, $\partial_{d-1}\partial_d\mathcal{C}^o = 0$. The operator ∂_d linearly extends to a mapping from d-chains of cells to $(d-1)$-chains.*

d-cycles, d-boundaries and the reduced homology groups $\tilde{H}_d$ are defined accordingly in a standard manner.

A3: *For every open d-cell $\mathcal{C}^o$, there are no nontrivial $(d-1)$-cycles in $K(\mathcal{C}^o)$ besides $\partial_d\mathcal{C}^o$.*

Observe that the axiom **A3** is equivalent to $\tilde{H}_{d-1}(K(C^o)) = 0$. For some needs, it will be convenient to have a stronger assumption:

A3*: *For every open cell $\mathcal{C}^o$, the corresponding $K(C^o)$ has $\tilde{H}_k(K(C^o)) = 0$ for every $k \geq 0$.*

The facet graph $G_d(K)$ of d-dimensional cell complex K has a vertex for each d-cell in K, and has an edge between two vertices if the corresponding d-cells have a common $(d-1)$-cell. As one would expect,

Theorem 5.4 *Under* $\{$**A1, A2, A3**$\}$*, the facet graph $G_d(Z)$ of a simple d-cycle Z, $d \geq 1$, is $(d+1)$-connected.*[3]

The argument previously used in deriving the analoguos Theorem 4.1 from Lemma 3.1 carries over to the present setting *verbatim*, employing the following generalization of Lemma 3.1:

Lemma 5.5 *Let D be a set of at most d open cells, and let $K(D)$ be the corresponding cell complex. Then, it holds that $\tilde{H}_{d-1}(K(D)) = 0$, under either*

(a) *the axioms* $\{$**A1, A2, A3**$\}$*, and the additional assumption that the cells in D are of dimension $< d$;*

(b) *the axioms* $\{$**A1, A2, A3***$\}$.

Part **(a)** of Lemma is fully sufficient for the needs of Theorem 5.4. Part **(b)**, a homological version of **(a)**, is a stronger statement, and, as we shall see in a moment, has some interesting consequences. An elementary self-contained proof of the Lemma appears in the Appendix.

[3]The assumption that $\mathcal{P}$ is a meet-semilattice is essential. Consider, e.g., the following set of open 2-cells (originating from faces of convex 2-polytopes): $C_1 = (1,2,3,4,5)$ with $\partial_2 C_1 = (1,2) + (2,3) + (3,4) + (4,5) - (1,5)$, $C_2 = (1,2,3)$ with $\partial_2(1,2,3) = (1,2) + (2,3) - (1,3)$, $C_3 = (1,3,4)$ with $\partial_2 C_3 = (1,3) + (3,4) - (1,4)$, and $C_4 = (1,4,5)$ with $\partial C_4 = (1,4) + (4,5) - (1,5)$. The meet-semilattice property fails, and the facet graph $G_2(Z_2)$ of the 2-cycle $Z_2 = C_1 - C_2 - C_3 - C_4$ is not 3-connected.

To conclude this paper, we would like to close the circle and return to where we have started, the Balinski's Theorem. The following elegant a generalization of it is due to [13], further strengthened by Björner in [7].

Let P be a convex $(d+1)$-polytope viewed as a cell complex, and let $B = \Delta(P)$ be the boundary of P, i.e., the subcomplex of P obtained by removing its single $(d+1)$-face. Then:

Theorem 5.6 (Homological Mixed-Connectivity Theorem) *For any $r = 0, \ldots, d-1$, B remains homologically r-connected with non-empty r-skeleton after removal* [4] *of any set F of at most $d - r$ of its open faces.*

Clearly, B is homologically k-connected (i.e., $\tilde{H}_k(B) = 0$) if and only if so is its $(k+1)$-skeleton, the subcomplex of B obtained by retaining only the faces of dimension $\leq k+1$. Thus, the case $r = 0$ of the the Mixed-Connectivity Theorem is indeed the Balinski's Theorem.

We would like to show that Theorem 5.6 is in fact equivalent (via polytope duality, and Alexander duality) to Lemma 5.5 **(b)**:

Claim 5.7 *For $\mathcal{P}$ whose elements correspond to the faces of a $(d+1)$-convex polytope under the containment relation, Lemma 5.5* **(b)** $\iff$ The Homological Mixed-Connectivity Theorem.

Proof. Formally, the Homological Mixed Connectivity Theorem is about vanishing of the lower homology groups of the cell complex $B \setminus U(F)$, where $U(F)$ (not a complex!) is the *upper* closure of F in B with respect to containment. Let P^* be the dual polytope of P, and, respectively, let F^* be the set faces dual to F in P^*. Then, by Combinatorial Alexander duality for polytopes (see e.g., [9] [5], in particular the discussion towards the end of the Introduction section, and the reference therein),

$$\tilde{H}_r(B \setminus U(F)) \;\cong\; \tilde{H}^{d-r-1}(K(F^*))\,.$$

Since the cohomology groups are isomorphic to the homology groups, $\tilde{H}^{d-r-1}(K(F^*)) \cong \tilde{H}_{d-r-1}(K(F^*))$. Therefore, the cell complex $B \setminus U(F)$ has a vanishing r'th homology group if and only if $\tilde{H}_{d-r-1}(K(F^*)) = 0$. As the cells in F^* satisfy the assumptions of Lemma 5.5 **(b)** for any d, and $|F^*| = |F| \leq d - r$, this is precisely the content of this lemma. Thus, Lemma 5.5 **(b)** implies Theorem 5.6.

Observing that the argument is completely reversible, the reverse implication follows as well.

The fact that the r-skeleton of $B \setminus U(F)$ is not empty, can be shown by induction. Clearly, the most "destructive" set F is the set of $d - r$ 0-cells, i.e., points. Consider such F, and remove its points from B one by one. The link of the first point p_1 is a nonempty $(d-1)$-cycle Z_{d-1} in the cell complex B, and it survives the removal of p_1. Similarly, the removal of the second point in F either misses Z_d, or reduces it to a nonempty $(d-2)$-cycle in B, etc. After the removal of the entire F from B, a nonempty $(d-r)$-cycle Z_{d-r} survives. ∎

Acknowledgments: We are grateful to Roy Meshulam and Eran Nevo for enlightening discussions. Many thanks also to the anonymous referee for careful reading and numerous improvement suggestions.

References

[1] K. A. Adiprasito, A. Goodarzi, M. Varbaro, *Connectivity of Pseudomanifold Graphs From an Algebraic Point of View*, Comptes Rendus Mathematique Vol. 353, Issue 12, December 2015, pp. 1061-1065

[2] Ch. A. Athanasiadis, *On the Graph Connectivity of Skeleta of Convex Polytopes*, Discrete & Computational Geometry Vol. 42 No. 2 (2009), pp. 155-165

[3] M. L. Balinski, *On the Graph Structure of Convex Polyhedra in n-Space*, Pacific Journal of Mathematics 11 (2) (1961), pp. 431-434,

[4] Removing open faces means removing the faces themselves and their super-faces, but not their subfaces.

[5] The point of [9] is to provide a self-contained exposition based on the first principles. Alternatively, the Combinatorial Alexander duality for polytopes can be derived from the topological Alexander duality, see, e.g., [19], after providing a geometric argument that $|B \setminus U(F)|$, the geometric realization of $B \setminus U(F)$, is homotopy equivalent to $|B^* \setminus K(F^*)|$.

[4] David Barnette, *Graph Theorems for Manifolds*, Israel Journal of Mathematics Vol. 16, Issue 1 (1973), pp. 62-72

[5] David Barnette, *Decompositions of Homology Manifolds and Their Graphs*, Israel Journal of Mathematics Vol. 41, Issue 3 (1982), pp. 203-212

[6] D. Bauer, H. Broersma, E.Schmeichel, *Toughness in Graphs - A Survey*, Graphs and Combinatorics Vol. 22 (2006), pp. 1-35

[7] A. Björner, *Mixed Connectivity, Polytope Boundaries, and Matroid Basis Graphs*, a talk at Billerafest 2008, `http://emis.matem.unam.mx/journals/SLC/wpapers/s61vortrag/bjoerner3.pdf` Also, *Connectivity of Polytopes and Cohen-Macaulay Rigidity*, in preparation.

[8] A. Björner, *Topological methods*, in: Handbook of Combinatorics (R. Graham, M. Groetschel, and L. Lovasz, Eds.), pp. 18191872, North-Holland, Amsterdam, 1995

[9] A. Björner, M. Tancer, *Combinatorial Alexander Duality - a Short and Elementary Proof*, Discrete and Computational Geometry 42(4) (2009), pp. 586-593

[10] B. Bollobas, I. Leader, *Isoperimetric Problems for r-sets*, Combinatorics, Probability and Computing Vol. 13/2 (2004), pp. 277-279

[11] M.M. Cohen, *A Course in Simple-Homotopy Theory*, Springer-Verlag, New York, 1973

[12] Branko Gruenbaum, *Convex Polytopes*, 2nd ed., Springer (2003)

[13] G. Floystad, *Cohen-Macaulay cell complexes, Algebraic and geometric combinatorics*, Contemp. Math. 423, Amer. Math. Soc., Providence, RI, 2006, pp. 205-220.

[14] A. Fogelsanger, *The Generic Rigidity of Minimal Cycles*, Ph.D. Thesis, Cornell University, 1988. `http://www.armadillodanceproject.com/af/cornell/rigidityintro.pdf`

[15] M. Joswig, *Projectivities in Simplicial Complexes and Colorings of Simple Polytopes*, Mathematische Zeitschrift 240 vol.2 (2002) pp. 43–259.

[16] G. Kalai, *Enumeration of Q-acyclic Simplicial Complexes*, Israel Journal of Mathematics 45 (1983), pp.337-351.

[17] G. Kalai, *Polytope Skeletons And Paths*, Handbook of Discrete and Computational Geometry (2nd Edition), Chapter 20, CRC Press, 2004.

[18] V. Klee, *A property of polyhedral graphs*, J. Math. Mech. 13 (1964), pp. 1039-1042.

[19] J. R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley, 1984.

[20] I. Newman and Y. Rabinovich, *On Multiplicative λ-Approximations and Some Geometric Applications*, SIAM J. Comput. 42(3), 855-883 (2013).

[21] J. G. Oxley, *Matroid Theory*, Oxford University Press, 2006.

[22] G. Wegner, *d-Collapsing and Nerves of Families of Convex Sets* Archives of Mathematics 26 (1975), pp. 317-321

[23] G. M. Ziegler, *Lecture on Polytopes*, Springer-Verlag, 1995.

6 Appendix

6.1 A Proof of Lemma 3.2

Recall the statement of the Lemma:

Lemma *Let D be a collection of d simplices on the same underlying set of vertices, and let $K(D)$ be the simplicial complex defined by D. Then, $K(D)$ is $(d-1)$-collapsible.*

Observe that for $d = 1$, allowing a somewhat uncommon but well-defined 0-collapse with respect to the (-1)-dimensional $\emptyset$, the statement is trivial. For $d \geq 2$, we shall need the following closely related claim:

Claim 6.1 *Let σ be a simplex, and let T be a collection of at most r faces of σ. Then, $K(\sigma)$ r-collapses to $K(T)$.*

Proof. We proceed by an induction on r, the size of T, and on the dimension of σ for a fixed r. The base cases are $r = 0$, and for a fixed $r > 1$, when $\dim(\sigma) \leq r - 1$. In the former case the statement holds true, as its suffices to perform a single elementary 0-collapse with respect to the (-1)-dimensional $\emptyset$. In the latter case, one simply performs a series of r-collapses with respect to the set of all proper faces of σ that are not in $K(T)$, in a descending inclusion order.

For the inductive step, obverse that if T contains σ, then we are done. Otherwise, let ζ be a maximal face in T, and let v be a vertex of σ such that $v \notin \zeta$. Let $\sigma' = \sigma - \{v\}$, be the facet of σ not containing v.

For a collection of simplices $W \subseteq \sigma$ define $W_v = \{\tau \in W \,|\, v \in \tau\}$, and $W_v^- = \{\tau \setminus \{v\} \,|\, \tau \in W_v\}$. Observe that $T_v^- \subseteq K(\sigma')$, and $|T_v^-| \leq r - 1$ (as $\zeta \notin T_v^-$). Hence, by induction on r, $K(\sigma')$ is $(r-1)$-collapsible to T_v^-. Let $\tau_1', \ldots, \tau_k'$, be the sequence of faces of σ' on which the corresponding elementary $(r-1)$-collapses are performed. It is easy to see that the sequence $(\tau_1' \cup \{v\}), \ldots, (\tau_k' \cup \{v\})$ defines a legal series of elementary r-collapses reducing $K(\sigma)$ to $K(T_v) \cup K(\sigma')$ without ever touching a simplex in $K(\sigma)$ that does not contain v. This is due the inclusion-preserving bijection between the faces of σ' and σ given by $S \mapsto S \cup \{v\}$. To sum up, $K(\sigma)$ r-collapses to $K(T_v) \cup K(\sigma')$.

Similarly, by induction, now on the dimension, $K(\sigma')$ r-collapses to $K(T) \cap K(\sigma')$. Let $\tau_1, \ldots, \tau_\ell$ be the corresponding sequence of faces of σ' used to perform this collapse. Then, the following sequence of faces $(\tau_1' \cup \{v\}), \ldots, (\tau_k' \cup \{v\}), \tau_1, \ldots, \tau_\ell$ r-collapses $K(\sigma)$ to $K(T_v) \cup (K(T) \cap K(\sigma')) = K(T)$. ∎

We return to the proof of the Lemma. The proof is by induction on d, the size of D. The base case $d = 1$ has already been discussed. Assume $d \geq 2$, and that w.l.o.g., no simplex in D is contained in another one. Let σ be any simplex in D, and define $D' = D \backslash \{\sigma\}$.

Let $T = \{\sigma \cap \xi \mid \xi \in D'\}$; then $|T| \leq d - 1$. Hence, by Claim 6.1, $K(\sigma)$ $(d-1)$-collapses to $K(T)$. Since the corresponding elementary $(d-1)$-collapses do not involve any faces of $K(D')$, this implies that $K(D)$ $(d-1)$-collapses to $K(D')$. The latter complex, by inductive hypothesis, is $(d-2)$-collapsible, hence also $(d-1)$-collapsible, and the conclusion follows.

6.2 A Proof of Lemma 5.5

Let us formulate, for every $d \geq 1$, the following two statements, analogous to Claim 6.1 and Lemma 3.1, respectively:

$(\mathbf{I}^*)_\mathbf{d}$: *Let $\mathcal{C}$ be a closed cell, and T be a set of at most $d-1$ faces of $\mathcal{C}$. Let $K(T) \subseteq \mathcal{C}$ be the cell complex corresponding to T. Then $\tilde{H}_{d-1}(\mathcal{C}, K(T)) = 0$.* [6]

$(\mathbf{II}^*)_\mathbf{d}$: *Let D be a set of at most d open cells,and let $K(D)$ be the corresponding cell complex. Then, $\tilde{H}_{d-1}(K(D)) = 0$.*

Let also $(\mathbf{I})_\mathbf{d}$ and $(\mathbf{II})_\mathbf{d}$ be the corresponding versions of these statements, where $\mathcal{C}$ and the elements of D are assumed to be of dimension $\leq d$. We need to prove:

[6] Meaning that for every relative $(d-1)$-cycle X_{d-1} over $\mathcal{C}$, i.e., a $(d-1)$-chain such that $\partial X_{d-1} = 0$ outside of $K(T)$, there exists a $(d-1)$-boundary B_{d-1} over $\mathcal{C}$, such that $B_{d-1} = X_{d-1}$ outside of $K(T)$.

Lemma *Under the axioms* $\{$**A1** ,**A2, A3***$\}$*, the statements* $(\mathbf{I}^*)_\mathbf{d}$ *and* $(\mathbf{II}^*)_\mathbf{d}$ *hold for any* $d \geq 1$.
Similarly, under the axioms $\{$**A1** ,**A2, A3**$\}$*, the statements* $(\mathbf{I})_\mathbf{d}$ *and* $(\mathbf{II})_\mathbf{d}$ *hold for any* $d \geq 1$.

We shall present here the proof of the first part of the Lemma. The proof of the second part is almost identical, and the required small modifications are obvious.

We start with a preliminary statement about the size of d-cycles.

Claim 6.2 *For* $d \geq 1$*, any nontrivial* d*-cycle* Z *contains at least* $d+2$ d*-cells.*

Proof. The proof is by induction on d. For $d = 0$, one needs at least two 0-dimensional cells for the sum of coefficients to cancel out. Assume correctness for $d-1$. Let $\mathcal{C}^o$ be any open d-cell in the support of Z. Since $\partial\mathcal{C}^o$ is a $(d-1)$-cycle, by inductive assumption $K(\mathcal{C}^o)$ has $r \geq d+1$ facets $\Upsilon_1, \Upsilon_2, \ldots \Upsilon_r$ of dimension $d-1$. Since $\partial Z = 0$, for every Υ_i there exists at least one additional d-cell $\mathcal{C}_i^o \in Z$ besides $\mathcal{C}^o$ that contains Υ_i. Observe that $\mathcal{C}_i^o$ may not contain any other Υ_j, since otherwise $\mathcal{C}^o$ and $\mathcal{C}_i^o$ would have more than one common maximal subcell, contrary to **A1**. Thus, all $\mathcal{C}_i^o$ are distinct, and the support Z contains at least $d+2$ d-cells: $\mathcal{C}^o$ and $\{\mathcal{C}_i^o\}_{i=1}^r$. ■

The Lemma will be proven by induction on d, with the base case $(\mathbf{II}^*)_\mathbf{1}$, whose correctness is implied by Axiom **A3***. The inductive argument is provided by the following claim:

Claim 6.3 $(\mathbf{II}^*)_\mathbf{d-1} \Longrightarrow (\mathbf{I}^*)_\mathbf{d} \Longrightarrow (\mathbf{II}^*)_\mathbf{d}$.

Proof.

$(\mathbf{II}^*)_\mathbf{d-1} \Longrightarrow (\mathbf{I}^*)_\mathbf{d}$: Let $\mathcal{C}$, T and $K(T)$ be as in the premises of $(\mathbf{I}^*)_\mathbf{d}$. We need to show that any relative $(d-1)$-cycle in $(\mathcal{C}, K(T))$ is a relative $(d-1)$-boundary. Consider a relative $(d-1)$-cycle X_{d-1}, i.e., $\partial_{d-1}X_{d-1}$ supported on $\Delta(K(T))$. By $(\mathbf{II}^*)_\mathbf{d-1}$, $\tilde{H}_{d-2}(K(T)) = 0$. Therefore, there exists a $(d-1)$-chain Y_{d-1} supported on $K(T)$, such that $\partial X_{d-1} = \partial Y_{d-1}$ (by **A2**, ∂X_{d-1} is a $(d-2)$-cycle in $K(T)$). Then, $X_{d-1} - Y_{d-1}$ is a $(d-1)$-cycle in $\mathcal{C}$. By **A3***, this $(d-1)$-cycle is of the form $\partial_d C_d$ for some d-chain C_d. Thus, $X_{d-1} = Y_{d-1} + \partial_d C_d$. Keeping in mind that Y_{d-1} is supported on $K(T)$, the conclusion follows.

$(\mathbf{I}^*)_\mathbf{d} \Longrightarrow (\mathbf{II}^*)_\mathbf{d}$: Let D be a set of cells as in the premises of $(\mathbf{II}^*)_\mathbf{d}$. We proceed by induction on the number of *large* cells in D, i.e., the cells of dimension $\geq d$. Let Z_{d-1} be a $(d-1)$-cycle on $K(D)$; we need to show that it is also a $(d-1)$-boundary. If D contains no large cells, then $K(D)$ contains at most $|D| \leq d$ different $(d-1)$-cells. However, by Claim 6.2, a nontrivial $(d-1)$-cycle has support of size $\geq d+1$. Thus, in this case there are no nontrivial Z_{d-1} must be trivial, and we are done.

Assume now that there exists a large cell $\mathcal{C}^o \in D$. Let $D' = D \backslash \{\mathcal{C}^o\}$. Let $T = \{\ \mathcal{Q}^o \wedge \mathcal{C}^o\ \}_{\mathcal{Q}^o \in D'}$, where $\mathcal{Q}^o \wedge \mathcal{C}_d^o$ denotes the maximal element in $\mathcal{P}$ dominated by both $\mathcal{Q}^o$ and $\mathcal{C}_d^o$, as in $\mathbf{A1}$. Clearly, $|T| \leq d-1$.

Applying $(\mathbf{I}^*)_\mathbf{d}$ to $\mathcal{C}$, the closure of $\mathcal{C}^o$, and T, one concludes that Z_{d-1} is a relative boundary on $\mathcal{C}$ with respect to $K(D')$. Namely, there exists a boundary B_{d-1} in $\mathcal{C}$ such that $Z'_{d-1} = Z_{d-1} - B_{d-1}$ is supported on $K(D')$. However, D' has one less large cell than D, and by induction hypothesis, the $(d-1)$-cycle Z'_{d-1} is also a boundary B'_{d-1} in $K(D')$. Hence, $Z_{d-1} = B_{d-1} + B'_{d-1}$ is a $(d-1)$-boundary in $K(D)$. ■