Some combinatorial properties of flag simplicial pseudomanifolds and spheres

Christos A. Athanasiadis

Dedicated to Anders Björner on the occasion of his sixtieth birthday.

Abstract. A simplicial complex Δ is called *flag* if all minimal nonfaces of Δ have at most two elements. The following are proved: First, if Δ is a flag simplicial pseudomanifold of dimension $d-1$, then the graph of Δ (i) is $(2d-2)$-vertex-connected and (ii) has a subgraph which is a subdivision of the graph of the d-dimensional cross-polytope. Second, the h-vector of a flag simplicial homology sphere Δ of dimension $d-1$ is minimized when Δ is the boundary complex of the d-dimensional cross-polytope.

1. Introduction

We will be interested in finite simplicial complexes. Such a complex Δ is called *flag* if every set of vertices which are pairwise joined by edges in Δ is a face of Δ. For instance, every order complex (meaning the simplicial complex of all chains in a finite partially ordered set) is a flag complex. According to [11, p. 100], flag complexes form a fascinating class of simplicial complexes which deserves further study. The class of flag complexes coincides with that of clique complexes of finite graphs.

Much of the combinatorial structure of flag complexes seems to be significantly different from that of general simplicial complexes. For instance, a simplicial sphere of dimension $d-1$ can have as few as $d+1$ vertices and the minimum is attained by the boundary complex of the d-dimensional simplex. In contrast, as observed in [4, Lemma 2.1.14], every flag simplicial (homology) sphere (or, more generally, flag simplicial pseudomanifold) of dimension $d-1$ has at least $2d$ vertices and the minimum is attained by the boundary complex of the d-dimensional cross-polytope.

This paper proves some analogous (stronger) statements related to the graph structure and face enumeration of flag complexes, which demonstrate further that these complexes have a special position within the class of all simplicial complexes.

We will denote by $\mathcal{G}(\Delta)$ the one-dimensional skeleton of a simplicial complex Δ, called the *graph* of Δ. Recall that, given a positive integer m, an abstract graph $\mathcal{G}$ is said to be *m-connected* if $\mathcal{G}$ has at least $m+1$ nodes and any graph obtained from $\mathcal{G}$ by deleting $m-1$ or fewer nodes and their incident edges is connected (necessarily with at least one edge). It follows from Balinski's theorem [13, Theorem 3.14] that $\mathcal{G}(\Delta)$ is d-connected if Δ is the boundary complex of a d-dimensional simplicial polytope and from [2, Corollary 5] that the same statement holds for every simplicial pseudomanifold Δ of dimension $d-1$. Our first result is the following.

Theorem 1.1. *For every flag simplicial pseudomanifold Δ of dimension $d-1$, the graph $\mathcal{G}(\Delta)$ is $(2d-2)$-connected.*

A *subdivision* of an abstract graph $\mathcal{G}$ is any graph which can be obtained from $\mathcal{G}$ by selecting some of the edges of $\mathcal{G}$ and replacing each selected edge e by a path with the same endpoints as e, so that the interiors of these paths are pairwise disjoint and do not intersect the set of nodes of $\mathcal{G}$. Clearly, the graph of any simplicial pseudomanifold of dimension $d-1$ has a subgraph which is a subdivision of the complete graph on $d+1$ nodes (the same property was proved by Grünbaum [5] for graphs of d-dimensional convex polytopes and by Barnette [2] for a more general class of graphs of cell decompositions of manifolds). Our second result is the following.

Theorem 1.2. *For every flag simplicial pseudomanifold Δ of dimension $d-1$, the graph $\mathcal{G}(\Delta)$ has a subgraph which is a subdivision of the graph of the d-dimensional cross-polytope.*

It is not hard to show (see Proposition 2.2) that for every integer k, the number of faces of dimension k of a $(d-1)$-dimensional flag simplicial pseudomanifold Δ is minimized when Δ is the boundary complex of the d-dimensional cross-polytope. The same conclusion was proved by Meshulam [6] for the class of flag simplicial complexes with nonzero reduced $(d-1)$-homology. Our third result asserts that the analogous (stronger) statement for the h-vector (see Section 2 for definitions) of Δ is also valid when one is restricted to a certain class of simplicial complexes which includes all flag homology spheres, namely that of doubly Cohen–Macaulay flag complexes. We refer to recent work of Nevo [7] for conjectured lower bounds on face numbers of flag (and more general) homology spheres with fixed dimension and number of vertices.

Theorem 1.3. *The h-vector $(h_0(\Delta), h_1(\Delta), ..., h_d(\Delta))$ of any doubly Cohen–Macaulay flag simplicial complex Δ of dimension $d-1$ satisfies the inequalities*

$$(1) \qquad\qquad\qquad h_i(\Delta) \geq \binom{d}{i}$$

for $0 \leq i \leq d$. In particular, these inequalities are valid for all flag simplicial homology spheres of dimension $d-1$.

The previous theorem provides some (although only weak) evidence for the truth of a conjecture of Kalai (see [11, p. 100]), stating that the h-vector of any flag Cohen–Macaulay simplicial complex is also equal to the f-vector of a (balanced) simplicial complex.

This paper is structured as follows. Section 2 reviews basic definitions and background on simplicial complexes, as well as graph-theoretic terminology. Theorems 1.1 and 1.2 are proved in Section 3. Theorem 1.3 is proved in Section 4. Finally, a higher-dimensional analogue of Theorem 1.1 is discussed in Section 5.

2. Preliminaries

We will use the notation $[d] = \{1, 2, ..., d\}$, when d is a positive integer, and write $|S|$ for the cardinality of a finite set S.

Simplicial complexes. Let E be a finite set. An (abstract) *simplicial complex* on the ground set E is a collection Δ of subsets of E such that $\sigma \subseteq \tau \in \Delta$ implies $\sigma \in \Delta$. The elements of Δ are called *faces*. The *dimension* of a face σ is defined as one less than the cardinality of σ. The dimension of Δ is the maximum dimension of a face and is denoted by $\dim \Delta$. Faces of Δ of dimension zero and one are called *vertices* and *edges*, respectively. A *facet* of Δ is a face which is maximal with respect to inclusion. The complex Δ is *pure* if all its facets have the same dimension. The *k-skeleton* $\Delta^{\leq k}$ of Δ is the subcomplex formed by the faces of Δ of dimension at most k. The *simplicial join* $\Delta_1 * \Delta_2$ of two simplicial complexes Δ_1 and Δ_2 on

disjoint ground sets has as its faces the sets of the form $\sigma_1 \cup \sigma_2$, where $\sigma_1 \in \Delta_1$ and $\sigma_2 \in \Delta_2$.

The *closed star* of $v \in E$ in Δ is the subcomplex of Δ consisting of all subsets of those faces of Δ which contain v. The *antistar* of v in Δ is defined as the restriction $\{\tau \in \Delta : v \notin \tau\}$ of Δ on the set $E \setminus \{v\}$ and is denoted by $\Delta \setminus v$. More generally, for $\sigma \subseteq E$ we denote by $\Delta \setminus \sigma$ the restriction $\{\tau \in \Delta : \tau \cap \sigma = \varnothing\}$ of Δ on the set $E \setminus \sigma$. The *link* of a face σ in Δ is defined as $\Delta / \sigma = \{\tau \setminus \sigma : \tau \in \Delta$ and $\sigma \subseteq \tau\}$. For simplicity, we write Δ / v instead of $\Delta / \{v\}$ for $v \in E$.

A sequence $(\tau_0, \tau_1, ..., \tau_n)$ of facets of Δ is said to be a *strong chain* if $\tau_{i-1} \cap \tau_i$ is a codimension-one face of both τ_{i-1} and τ_i for $1 \le i \le n$. One can define an equivalence relation $\sim$ on the set of facets of Δ by letting $\sigma \sim \tau$ if there exists a strong chain $(\tau_0, \tau_1, ..., \tau_n)$ of facets of Δ such that $\tau_0 = \sigma$ and $\tau_n = \tau$. The simplicial complex formed by all subsets of the facets of Δ in an equivalence class of $\sim$ is called a *strong component* of Δ. A simplicial complex is *strongly connected* if it has a unique strong component. In particular, such a complex must be pure. A strongly connected $(d-1)$-dimensional simplicial complex Δ is said to be a (simplicial) *pseudomanifold* if each face of Δ of dimension $d-2$ is contained in exactly two facets. It is an easy observation that if Δ is a pseudomanifold, then $\Delta \setminus v$ is pure for every vertex v of Δ. The following stronger statement is a special case of [2, Lemma 2].

Lemma 2.1. (Cf. [2, Lemma 2]) *For every simplicial pseudomanifold Δ and vertex v, the complex $\Delta \setminus v$ is strongly connected.*
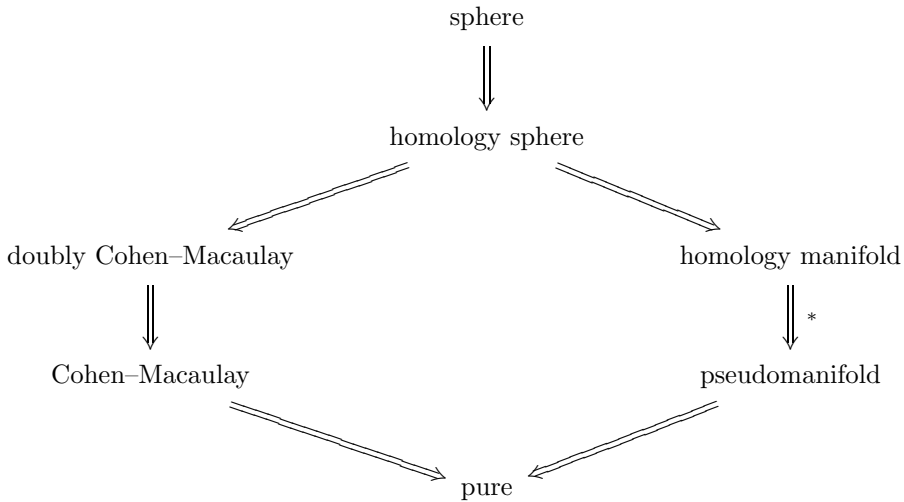
The link Δ / σ of a face σ in a pseudomanifold Δ may not be a pseudomanifold, since it may fail to be strongly connected. However, any strong component of such a link is also a pseudomanifold.

The *boundary complex* of a simplicial convex polytope P is the abstract simplicial complex on the vertex set of P for which a subset σ of the vertex set of P is a face if and only if the convex hull of σ is a face of P, other than P itself. For instance, the boundary complex of a d-dimensional simplex consists of all subsets of a $(d+1)$-element set of cardinality at most d. When we talk about topological properties of an abstract simplicial complex Δ, we implicitly refer to those of its geometric realization $\|\Delta\|$ [3, Section 9]. For instance, Δ is said to be a $(d-1)$-*sphere* if $\|\Delta\|$ is homeomorphic to a sphere of dimension $d-1$. We call Δ a *homology sphere* (over some fixed field $\mathbb{K}$) if for all $\sigma \in \Delta$ (including $\sigma = \varnothing$) we have

$$\widetilde{H}_i(\Delta / \sigma, \mathbb{K}) = \begin{cases} 0, & \text{if } i < \dim(\Delta / \sigma), \\ \mathbb{K}, & \text{if } i = \dim(\Delta / \sigma), \end{cases}$$

where $\widetilde{H}_*(\Gamma, \mathbb{K})$ denotes reduced simplicial homology of Γ with coefficients in the field $\mathbb{K}$. We call Δ a *homology manifold* (over $\mathbb{K}$) if Δ/σ is a homology sphere for every nonempty face σ of Δ. We note that every homology manifold which is connected and has at least two vertices must be a pseudomanifold.

Suppose that the ground set E of Δ has n elements, say $E = \{v_1, v_2, ..., v_n\}$. The *face ring* (or *Stanley–Reisner ring*) associated with Δ is defined as the quotient $\mathbb{K}[\Delta] = \mathbb{K}[x_1, x_2, ..., x_n]/I_\Delta$ of the polynomial ring over $\mathbb{K}$ in the variables $x_1, x_2, ..., x_n$ by the ideal I_Δ generated by the square-free monomials $x_{i_1} x_{i_2} ... x_{i_r}$ for which $\{i_1, i_2, ..., i_r\} \notin \Delta$. The complex Δ is said to be *Cohen–Macaulay* or *Gorenstein* (over $\mathbb{K}$) if $\mathbb{K}[\Delta]$ is a *Cohen–Macaulay* or *Gorenstein* ring, respectively. We refer to [11] for a thorough discussion of these concepts. By Reisner's theorem [11, Corollary II.4.2], Δ is Cohen–Macaulay if and only if $\widetilde{H}_i(\Delta/\sigma, \mathbb{K}) = 0$ for all $\sigma \in \Delta$ and $i < \dim(\Delta/\sigma)$. Given a positive integer m, the complex Δ is said to be *m-Cohen–Macaulay* (or *doubly Cohen–Macaulay* for $m=2$) if $\Delta \setminus \sigma$ is Cohen–Macaulay of the same dimension as Δ for all subsets σ of E of cardinality less than m (including $\sigma = \varnothing$). We have the hierarchy of properties

sphere

$\Downarrow$

homology sphere

doubly Cohen–Macaulay   homology manifold

Cohen–Macaulay   pseudomanifold

pure

(where the implication $\overset{*}{\Rightarrow}$ assumes connectivity and positive dimension). Moreover, boundary complexes of simplicial polytopes are spheres. The classes of boundary complexes of simplicial polytopes, homology spheres, homology manifolds, m-Cohen–Macaulay complexes (for any fixed $m \geq 1$) and pure complexes are all closed under taking links of faces. The class of homology spheres coincides with that of nonacyclic Gorenstein complexes.

Let Δ be any simplicial complex of dimension $d-1$. The number of k-dimensional faces of Δ will be denoted by $f_k(\Delta)$, so that $f_{-1}(\Delta)=1$ unless Δ is the void complex $\varnothing$. The sequence $f(\Delta)=(f_{-1}(\Delta),f_0(\Delta),...,f_{d-1}(\Delta))$ is called the *f-vector* of Δ. The *h-vector* of Δ is the sequence $h(\Delta)=(h_0(\Delta),h_1(\Delta),...,h_d(\Delta))$ defined by the equality

$$(2) \qquad \sum_{i=0}^{d} h_i(\Delta)\,x^i = \sum_{i=0}^{d} f_{i-1}(\Delta)\,x^i(1-x)^{d-i}.$$

The *reduced Euler characteristic* of Δ is defined as

$$(3) \qquad \widetilde{\chi}(\Delta) = (-1)^{d-1}\,h_d(\Delta) = \sum_{i=0}^{d} (-1)^{i-1} f_{i-1}(\Delta).$$

The polynomial which appears in either side of (2) is called the *h-polynomial* of Δ and is denoted by $h_\Delta(x)$. It is a fundamental property of Cohen–Macaulay complexes that $h_i(\Delta)\geq 0$ holds for every index i.

A simplicial complex Δ is called *flag* if all its minimal nonfaces have at most two elements. The simplicial join of two flag complexes and the link of any face of a flag complex are also flag complexes. For instance, the boundary complex of the d-dimensional cross-polytope [13, Example 0.4] is isomorphic to the simplicial join of d copies of the 0-sphere (the zero-dimensional complex with just two vertices) and hence it is a flag complex. The next proposition asserts that among all $(d-1)$-dimensional flag simplicial pseudomanifolds, the boundary complex of the d-dimensional cross-polytope has the minimum number of faces in each dimension.

Proposition 2.2. *Any $(d-1)$-dimensional flag simplicial pseudomanifold has no fewer than $2^i\binom{d}{i}$ faces of dimension $i-1$ for all $0\leq i\leq d$.*

Proof. Let Δ be a flag simplicial pseudomanifold of dimension $d-1$. We proceed by induction on d. The result is easily verified for $i=0$ or $d=1$, so we assume that $i\geq 1$ and $d\geq 2$. Since Δ is a flag simplicial complex which is not a simplex, there exist two vertices, say u and v, of Δ such that $\{u,v\}$ is not an edge of Δ. Among the $(i-1)$-dimensional faces of Δ, there exist $f_{i-2}(\Delta/u)$ faces which contain u, $f_{i-2}(\Delta/v)$ faces which contain v and $f_{i-1}(\Delta/u)$ faces which belong to Δ/u. Since these three sets of $(i-1)$-dimensional faces of Δ are pairwise disjoint, we conclude that

$$f_{i-1}(\Delta) \ \geq \ f_{i-1}(\Delta/u)+f_{i-2}(\Delta/u)+f_{i-2}(\Delta/v).$$

Since every strong component of Δ/u or Δ/v is a flag pseudomanifold of dimension $d-2$, it follows from the previous inequality and the induction hypothesis that

$$f_{i-1}(\Delta) \geq 2^i \binom{d-1}{i} + 2 \cdot 2^{i-1} \binom{d-1}{i-1} = 2^i \binom{d}{i}.$$

This completes the induction and the proof. $\square$

Graphs. A graph is a simplicial complex of dimension zero or one. We will refer to the vertices of a graph as *nodes*, to avoid possible confusion with vertices of other simplicial complexes considered simultaneously. Two nodes u and v of a graph $\mathcal{G}$ are said to be *adjacent* (or joined by an edge) in $\mathcal{G}$ if $\{u,v\}$ is an edge of $\mathcal{G}$. A *walk* of length n in $\mathcal{G}$ is an alternating sequence $w=(v_0,e_1,v_1,...,e_n,v_n)$ of nodes and edges, such that $e_i=\{v_{i-1},v_i\}$ for $1 \leq i \leq n$. We say that w *connects* nodes v_0 and v_n, which are the *endpoints* of w. The walk w is said to be a *path* if $v_0,v_1,...,v_n$ are pairwise distinct; in this case $v_1,...,v_{n-1}$ are the *interior nodes* of w. We say that $\mathcal{G}$ is *connected* if any two nodes can be connected by a walk in $\mathcal{G}$. Given a positive integer m, the graph $\mathcal{G}$ is said to be *m-connected* if it has at least $m+1$ nodes and $\mathcal{G} \backslash \sigma$ is connected for all subsets σ of the set of nodes of $\mathcal{G}$ with cardinality less than m. Equivalently, $\mathcal{G}$ is m-connected if it is one-dimensional and m-Cohen–Macaulay over some field (equivalently, over all fields) as a simplicial complex. A *subgraph* of $\mathcal{G}$ is any graph which can be obtained by deleting some of the edges of $\mathcal{G} \backslash \sigma$, for some subset σ of the set of nodes of $\mathcal{G}$. A *subdivision* of $\mathcal{G}$ is any graph which can be obtained from $\mathcal{G}$ by selecting some of the edges of $\mathcal{G}$ and replacing each selected edge e by a path with the same endpoints as e, so that the interiors of these paths are pairwise disjoint and do not intersect the set of nodes of $\mathcal{G}$.

The graph $\mathcal{G}(\Delta)$. The 1-skeleton of a simplicial complex Δ is called the *graph* of Δ and is denoted by $\mathcal{G}(\Delta)$. We are primarily interested in $\mathcal{G}(\Delta)$ when Δ is a pseudomanifold. A walk $(v_0,e_1,v_1,...,e_n,v_n)$ in $\mathcal{G}(\Delta)$ will be called a Δ-*strong walk* if for every index $1 \leq i \leq n$ there exists a facet τ_i of Δ containing e_i, so that τ_i and τ_{i+1} lie in the same strong component of the closed star of v_i in Δ for all $1 \leq i \leq n-1$. This condition allows for arguments which use induction on the dimension of a pseudomanifold, by considering the links of its vertices. It implies, for instance, that for every index $1 \leq i \leq n-1$, some strong component of Δ/v_i contains both v_{i-1} and v_{i+1}.

Part (i) of the following proposition is implicit in the proof of [2, Theorem 4]. We include the proof for the convenience of the reader.

Proposition 2.3. *Let Δ be a simplicial pseudomanifold of dimension $d-1$.*

(i) *(Cf. [2, Theorem 4]) If σ is any subset of the vertex set of Δ of cardinality less than d, then any two vertices of Δ not in σ can be connected by a Δ-strong walk in $\mathcal{G}(\Delta)\backslash\sigma$. In particular, the graph $\mathcal{G}(\Delta)$ is d-connected.*

(ii) *If σ is a face of Δ, then any two vertices of Δ not in σ can be connected by a Δ-strong walk in $\mathcal{G}(\Delta)\backslash\sigma$. In particular, the graph $\mathcal{G}(\Delta)\backslash\sigma$ is connected.*

Proof. Let a and b be any two vertices of Δ not in σ. We will show that a and b can be connected by a Δ-strong walk in $\mathcal{G}(\Delta)\backslash\sigma$. Pick any element v of σ. By Lemma 2.1, there exists a strong chain $\mathcal{C}=(\tau_0, \tau_1, ..., \tau_n)$ of facets of $\Delta\backslash v$ such that $a\in\tau_0$ and $b\in\tau_n$. We claim that the intersection $\tau_{i-1}\cap\tau_i$ is not a subset of σ for any index $1\leq i\leq n$. Indeed, this is clear in part (i) since $\tau_{i-1}\cap\tau_i$ has $d-1$ elements and does not contain v. If σ is a face of Δ, as in part (ii), of cardinality at least d, so that σ is a facet, then the claim holds because σ does not appear in the chain $\mathcal{C}$ and τ_{i-1} and τ_i are the only facets of Δ which contain $\tau_{i-1}\cap\tau_i$. For $1\leq i\leq n$ we pick any vertex $v_i\in\tau_{i-1}\cap\tau_i$ not contained in σ and observe that $\{a, v_1, ..., v_n, b\}$ is the set of nodes of a Δ-strong walk in $\mathcal{G}(\Delta)\backslash\sigma$ connecting a and b. This completes the proof. $\square$

3. Connectivity of $\mathcal{G}(\Delta)$

In this section we prove Theorems 1.1 and 1.2.

Proof of Theorem 1.1. We will first prove the theorem for connected homology manifolds and then indicate how this proof can be modified in the case of pseudomanifolds.

Let Δ be a connected flag simplicial homology manifold of dimension $d-1$ and let τ be any subset of the vertex set of Δ of cardinality less than $2d-2$. Since Δ has at least $2d$ vertices, we only need to show that any two vertices, say a and b, of Δ not in τ can be connected by a walk in $\mathcal{G}(\Delta)\backslash\tau$. We proceed by induction on d. Clearly, we may assume that $d\geq 3$. Let σ denote the set of elements of τ which are adjacent to at least $2d-4$ elements of τ in $\mathcal{G}(\Delta)$. Observe that, in view of our assumption on the cardinality of τ, each element of σ is adjacent in $\mathcal{G}(\Delta)$ to all other elements of τ. In particular, the elements of σ are pairwise adjacent in $\mathcal{G}(\Delta)$. As a result, since Δ is a flag complex, σ is a face of Δ. Therefore, by Proposition 2.3(ii), any two vertices of Δ not in σ can be connected by a walk in $\mathcal{G}(\Delta)\backslash\sigma$. Let $w=(v_0, e_1, v_1, ..., e_n, v_n)$ be such a walk connecting a and b.

We may assume that no two consecutive nodes of w are in τ. Indeed, suppose that v_{i-1} and v_i are both elements of τ for some index i. The link of e_i in Δ is

a flag homology sphere of dimension $d-3$ and hence it has at least $2d-4$ vertices. Since τ has at most $2d-5$ elements other than v_{i-1} and v_i, there exists a vertex u of Δ/e_i not in τ. Then $\{u, v_{i-1}, v_i\}$ is a two-dimensional face of Δ and inserting u between v_{i-1} and v_i in w results in a walk in $\mathcal{G}(\Delta)\backslash\sigma$ having a smaller number of pairs of consecutive nodes in τ. Repeating this process for every pair of consecutive nodes of w in τ results in a walk in $\mathcal{G}(\Delta)\backslash\sigma$ connecting a and b with the desired property.

Consider any node v_i of w which is an element of τ, so that $1 \leq i \leq n-1$ and neither v_{i-1} nor v_{i+1} is an element of τ. To complete the proof, it suffices to show that v_{i-1} and v_{i+1} can be connected by a walk in $\mathcal{G}(\Delta)\backslash\tau$ for any such index i. Since Δ/v_i is a flag simplicial homology sphere of dimension $d-2$, by our induction hypothesis the graph $\mathcal{G}(\Delta/v_i)$ is $(2d-4)$-connected. By construction v_i is not in σ and hence at most $2d-5$ vertices of Δ/v_i are in τ. As a result, v_{i-1} and v_{i+1} can be connected by a walk in $\mathcal{G}(\Delta/v_i)\backslash\tau$ and hence by a walk in $\mathcal{G}(\Delta)\backslash\tau$.

Finally, suppose that Δ is a flag simplicial pseudomanifold of dimension $d-1$. By Proposition 2.3(ii), we may choose the walk w in the previous argument to be Δ-strong. Let $\tau_1, ..., \tau_n$ be facets of Δ as in the definition of a Δ-strong walk. When inserting a vertex u not in τ between two consecutive nodes v_{i-1} and v_i of w which are in τ, we can guarantee that the new walk will also be Δ-strong. Indeed, let Γ be the strong component of Δ/e_i which contains $\tau_i\backslash e_i$. Since Γ is a flag pseudomanifold of dimension $d-3$, it has at least $2d-4$ vertices and we may choose u to be in Γ. By construction, $\{u, v_{i-1}, v_i\}$ is contained in a facet of Δ which can be connected to τ_i by a strong chain of facets, each of which contains both v_{i-1} and v_i. From this fact it follows that our new walk is also Δ-strong. In the final part of the argument we only need to replace the link of v_i in Δ with its strong component which contains $\tau_i\backslash\{v_i\}$ and $\tau_{i+1}\backslash\{v_i\}$. $\square$

The next statement follows from the proof of Theorem 1.1.

Corollary 3.1. *Let Δ be a flag simplicial pseudomanifold of dimension $d-1$. If τ is a subset of the vertex set of Δ of cardinality less than $2d-2$, then any two vertices of Δ not in τ can be connected by a Δ-strong walk in $\mathcal{G}(\Delta)\backslash\tau$.*

Remark 3.2. Responding to a question posed by the author in a previous version of this paper, Novik [8] has shown that the k-skeleton of every $(d-1)$-dimensional flag simplicial homology sphere is $2(d-k)$-Cohen–Macaulay (this statement generalizes Theorem 1.1 in the case of flag homology spheres). The proof uses the Stanley–Reisner ring of Δ, [9, Lemma 5.1] and the Taylor resolution for quadratic monomial ideals.

Proof of Theorem 1.2. Let $\sigma = \{v_1, v_2, ..., v_d\}$ be a facet of Δ. Since Δ is a pseudomanifold, for each $1 \leq i \leq d$ there exists a unique vertex u_i of Δ other than v_i such that $(\sigma \backslash \{v_i\}) \cup \{u_i\}$ is also a facet of Δ. Since Δ is flag, the u_i are pairwise distinct. We set $\tau = \{u_1, u_2, ..., u_d\}$ and recall that the graph of the d-dimensional cross-polytope can be obtained from the complete graph on $2d$ nodes by removing d edges which are mutually disjoint. Since we have $\{v_i, v_j\} \in \Delta$ and $\{v_i, u_j\} \in \Delta$ for distinct indices i and j, it suffices to show that any two elements of τ can be connected by a path in $\mathcal{G}(\Delta)$ so that the sets of interior nodes of all $\binom{d}{2}$ resulting paths are mutually disjoint and each such set intersects neither σ nor τ.

Consider the face $\sigma_{ij} = \sigma \backslash \{v_i, v_j\}$ of Δ, where $1 \leq i < j \leq d$. The link Δ / σ_{ij} is a one-dimensional simplicial complex, each vertex of which belongs to exactly two edges. As a result, Δ / σ_{ij} is a disjoint union of one-dimensional spheres. Moreover, it contains the edges $\{v_i, v_j\}$, $\{v_i, u_j\}$ and $\{v_j, u_i\}$. From these facts it follows that there exists a path p_{ij} in $\mathcal{G}(\Delta)$ connecting u_i and u_j, each interior node of which is a vertex of Δ / σ_{ij} other than v_i and v_j.

We claim that (i) the set of nodes of p_{ij} does not intersect σ, (ii) the set of interior nodes of p_{ij} does not intersect τ and (iii) no vertex of Δ is an interior node of two or more of the paths p_{ij}. Indeed, (i) is clear by construction. Consider $u_r \in \tau \backslash \{u_i, u_j\}$. Since Δ is $(d-1)$-dimensional and flag and $(\sigma \backslash \{v_r\}) \cup \{u_r\} \in \Delta$, we have $\{u_r, v_r\} \notin \Delta$. Since $v_r \in \sigma_{ij}$, we conclude that u_r is not a vertex of Δ / σ_{ij} and hence that u_r cannot be one of the nodes of p_{ij}. This proves (ii). Finally, suppose that u is a vertex of Δ which is an interior node of two paths p_{ij} and p_{kl}. Then u belongs to both links Δ / σ_{ij} and Δ / σ_{kl} and hence we have $\{u, v\} \in \Delta$ for every $v \in \sigma_{ij} \cup \sigma_{kl}$. Since Δ is flag, it follows that $\sigma_{ij} \cup \sigma_{kl} \cup \{u\} \in \Delta$. Since $\sigma_{ij} \cup \sigma_{kl}$ is equal to σ, if $\{i, j\}$ and $\{k, l\}$ are disjoint, and to $\sigma \backslash \{v_r\}$ for some $r \in \{i, j, k, l\}$ otherwise, it follows that $\sigma \cup \{u\} \in \Delta$ or $(\sigma \backslash \{v_r\}) \cup \{u\} \in \Delta$. By our choice of u_r and since σ is a facet of Δ, we conclude that either $u \in \sigma$ or $u = u_r$, contradicting (i) and (ii). This contradiction proves (iii). It follows from facts (i)–(iii) that the paths p_{ij} have the desired properties. $\square$

4. Lower bound for the h-vector

In this section we prove Theorem 1.3. We will make use of the following elementary lemma.

Lemma 4.1. *Let Δ be a pure simplicial complex and v be a vertex of Δ. We have*

$$h_\Delta(x) = \begin{cases} h_{\Delta \backslash v}(x) + x h_{\Delta / v}(x), & \text{if } \dim(\Delta \backslash v) = \dim \Delta, \\ h_{\Delta \backslash v}(x), & \text{otherwise.} \end{cases}$$

Proof. Let $d-1=\dim\Delta$, so that $\dim(\Delta/v)=d-2$ and either $\dim(\Delta\backslash v)=d-1$ or else $\Delta\backslash v=\Delta/v$. By considering those $(i-1)$-dimensional faces of Δ which contain v and those which do not separately, we see that

$$f_{i-1}(\Delta)=f_{i-1}(\Delta\backslash v)+f_{i-2}(\Delta/v)$$

for $0\leq i\leq d$ (where $f_{i-1}(\Gamma)=0$ for negative integers i by convention). In either case, multiplying this equation with $x^i(1-x)^{d-i}$, summing and using (2) we arrive at the proposed equality expressing the h-polynomial of Δ in terms of those of $\Delta\backslash v$ and Δ/v. $\square$

Proof of Theorem 1.3. Let Δ be a doubly Cohen–Macaulay flag simplicial complex of dimension $d-1$. We need to show that

$$(4) \qquad h_\Delta(x)\geq(1+x)^d,$$

where such an inequality will be meant to hold coefficientwise. We proceed by induction on d. The statement holds for $d=1$ since then Δ consists of $q\geq2$ vertices, having no other nonempty faces, and $h_\Delta(x)=1+(q-1)x$. Suppose that $d\geq2$. Since Δ is a flag simplicial complex which is not a simplex, there exist two vertices, say u and v, of Δ such that $\{u,v\}$ is not an edge of Δ. Since the link Δ/v is doubly Cohen–Macaulay and flag of dimension $d-2$, our induction hypothesis implies that

$$(5) \qquad h_{\Delta/v}(x)\geq(1+x)^{d-1}.$$

For the same reason we have

$$h_{\Delta/u}(x)\geq(1+x)^{d-1}.$$

Let Γ denote the closed star of u in Δ. Then Γ is a subcomplex of $\Delta\backslash v$ and both Γ and $\Delta\backslash v$ are Cohen–Macaulay of dimension $d-1$. The monotonicity property of h-vectors [10, Theorem 2.1] implies that $h_{\Delta\backslash v}(x)\geq h_\Gamma(x)$ and Lemma 4.1 implies that $h_\Gamma(x)=h_{\Gamma\backslash u}(x)=h_{\Delta/u}(x)$, so that

$$(6) \qquad h_{\Delta\backslash v}(x)\geq h_\Gamma(x)=h_{\Delta/u}(x)\geq(1+x)^{d-1}.$$

The desired inequality (4) follows by combining (5) and (6) with Lemma 4.1. $\square$

We conjecture that if equality holds in (1) for some $1\leq i\leq d-1$, then Δ is isomorphic to the boundary complex of the d-dimensional cross-polytope. This statement does not follow immediately from the previous proof.

5. A higher-dimensional analogue

Balinski's theorem on the one-dimensional skeleton $\mathcal{G}(P)$ of a convex polytope P was generalized in [1] to the graphs $\mathcal{G}_k(P)$ defined as follows. The nodes of $\mathcal{G}_k(P)$ are the k-dimensional faces of P and two such faces are adjacent if there exists a $(k+1)$-dimensional face of P which contains them both. Theorem 1.1 can also be generalized in this direction. Given any $(d-1)$-dimensional simplicial complex Δ and integer $0 \le k \le d-2$, we denote by $\mathcal{G}_k(\Delta)$ the graph with the k-dimensional faces of Δ as nodes, in which two such faces are adjacent if there exists a $(k+1)$-dimensional face of Δ which contains them both. The graphs $\mathcal{G}_k(P)$ and $\mathcal{G}_k(\Delta)$ reduce to $\mathcal{G}(P)$ and $\mathcal{G}(\Delta)$, respectively, for $k=0$.

Theorem 5.1. *If $n_k(d)=2(k+1)(d-k-1)$, then the graph $\mathcal{G}_k(\Delta)$ is $n_k(d)$-connected for every connected flag simplicial homology manifold Δ of dimension $d-1$ and all integers $0 \le k \le d-2$.*

The proof of Theorem 5.1 is similar to that of the main result of [1] and is omitted. The value of $n_k(d)$ in Theorem 5.1 cannot be improved, as the example of the d-dimensional cross-polytope shows. It is likely that Theorem 5.1 can be extended to the class of all flag simplicial pseudomanifolds.

References

1. ATHANASIADIS, C. A., On the graph connectivity of skeleta of convex polytopes, *Discrete Comput. Geom.* **42** (2009), 155–165.
2. BARNETTE, D., Graph theorems for manifolds, *Israel J. Math.* **16** (1973), 62–72.
3. BJÖRNER, A., Topological methods, in *Handbook of Combinatorics*, pp. 1819–1872, North-Holland, Amsterdam, 1995.
4. GAL, S. R., Real root conjecture fails for five- and higher-dimensional spheres, *Discrete Comput. Geom.* **34** (2005), 269–284.
5. GRÜNBAUM, B., On the facial structure of convex polytopes, *Bull. Amer. Math. Soc.* **71** (1965), 559–560.
6. MESHULAM, R., Domination numbers and homology, *J. Combin. Theory Ser. A* **102** (2003), 321–330.
7. NEVO, E., Remarks on missing faces and generalized lower bounds on face numbers, *Electron. J. Combin.* **16** (2009), Research Paper 8, 11 pp.
8. NOVIK, I., *Personal communication*, February 4, 2009.
9. NOVIK, I. and SWARTZ, E., Face ring multiplicity via CM-connectivity sequences, *Canad. J. Math.* **61** (2009), 888–903.
10. STANLEY, R. P., A monotonicity property of h-vectors and h^*-vectors, *European J. Combin.* **14** (1993), 251–258.
11. STANLEY, R. P., *Combinatorics and Commutative Algebra*, 2nd ed., Progress in Mathematics **41**, Birkhäuser, Boston, 1996.

12. WOTZLAW, R. F., *Incidence Graphs and Unneighborly Polytopes*, Doctoral Dissertation, TU Berlin, 2009.

13. ZIEGLER, G. M., *Lectures on Polytopes*, Graduate Texts in Mathematics **152**, Springer, New York, 1995.

Christos A. Athanasiadis
Department of Mathematics
University of Athens
Athens 15784
Greece
caath@math.uoa.gr