Ark. Mat., **49** (2011), 335–350
DOI: 10.1007/s11512-009-0119-z

Balanced complexes and complexes without large missing faces

Michael Goff, Steven Klee and Isabella Novik

Abstract. The face numbers of simplicial complexes without missing faces of dimension larger than i are studied. It is shown that among all such $(d-1)$-dimensional complexes with non-vanishing top homology, a certain polytopal sphere has the componentwise minimal f-vector; and moreover, among all such 2-Cohen–Macaulay (2-CM) complexes, the same sphere has the componentwise minimal h-vector. It is also verified that the l-skeleton of a flag $(d-1)$-dimensional 2-CM complex is $2(d-l)$-CM, while the l-skeleton of a flag piecewise linear $(d-1)$-sphere is $2(d-l)$-homotopy CM. In addition, tight lower bounds on the face numbers of 2-CM balanced complexes in terms of their dimension and the number of vertices are established.

1. Introduction

In this paper we study balanced simplicial complexes and complexes without large missing faces. For the latter class of complexes we settle in the affirmative several open questions raised in the recent papers by Athanasiadis [1] and Nevo [14], while for the former class we establish tight lower bounds on their face numbers in terms of dimension and the number of vertices, thus strengthening the celebrated lower bound theorem for spheres.

A simplicial complex Δ on the vertex set $[n]:=\{1,2,...,n\}$ is a collection of subsets of $[n]$ that is closed under inclusion and contains all singletons $\{i\}$ for $i\in[n]$. The elements of Δ are called its *faces*. A set $F\subseteq[n]$ is called a *missing face* of Δ if it is not a face of Δ, but all its proper subsets are. Hence the collection of all missing faces of Δ carries the same information as Δ itself. Thus it is perhaps not very surprising that imposing certain conditions on the allowed sizes of missing faces may result in severe restrictions on the corresponding simplicial complexes.

Novik's research was partially supported by an Alfred P. Sloan Research Fellowship and NSF grant DMS-0801152.

One simple example of this phenomenon is that while a simplicial $(d-1)$-sphere may have as few as $d+1$ vertices, a *flag* $(d-1)$-sphere (that is, a simplicial complex with all its missing faces of size two or, equivalently, 1-dimensional) needs at least $2d$ vertices. In fact, Meshulam [12] proved that among all $(d-1)$-dimensional flag simplicial complexes with non-vanishing top homology, the boundary of the d-dimensional cross-polytope simultaneously minimizes all the face numbers. Similarly, it was recently verified in [1] that among all 2-Cohen–Macaulay (2-CM, for short) flag $(d-1)$-dimensional complexes, the boundary of the d-dimensional cross-polytope simultaneously minimizes all of the h-numbers.

In [14], Nevo considered the more general class of $(d-1)$-dimensional simplicial complexes with no missing faces of dimension larger than i (equivalently, of size larger than $i+1$). He conjectured [14, Conjecture 1.3] that among all such complexes with non-vanishing top homology, a certain polytopal sphere, $S(i, d-1)$ (that for $i=1$ coincides with the boundary of the cross-polytope), simultaneously minimizes all of the face numbers. He also asked [14, Problem 3.1] if the same sphere $S(i, d-1)$ has the componentwise minimal h-vector in the class of all homology $(d-1)$-spheres without missing faces of dimension larger than i. One of our main results, Theorem 3.1, establishes both of these conjectures.

In addition to verifying that the h-numbers of flag spheres are at least as large as those of the cross-polytope, Athanasiadis shows in [1, Theorem 1.1] that the graph of a flag simplicial pseudomanifold of dimension $d-1$ is $2(d-1)$-vertex-connected. This is in contrast to the fact that without the flag assumption one can only guarantee its d-connectedness (for polytopes this is Balinski's theorem, see [18, Theorem 3.14]; the general case is due to Barnette [3]). The above result prompted Athanasiadis to ask [1, Remark 3.2] if, for every $0 \le l \le d-1$, the l-skeleton of a flag homology $(d-1)$-sphere is $2(d-l)$-CM and if the l-skeleton of a flag piecewise linear (PL, for short) $(d-1)$-sphere is $2(d-l)$-homotopy CM. In Theorem 4.1 we settle both of these questions in the affirmative.

The face numbers of flag complexes are closely related to those of balanced complexes. (A simplicial $(d-1)$-dimensional complex is called *balanced* [15] if its 1-skeleton, considered as a graph, is vertex d-colorable.) Indeed, it is a result of Frohmader [10] that for every flag complex Δ there exists a balanced complex Γ with the same f-vector, and it is a conjecture of Kalai [17, p. 100] that if Δ is flag and CM, then one can choose the corresponding balanced Γ to also be CM.

The lower bound theorem for spheres [4], [11] asserts that among all homology $(d-1)$-spheres on n vertices, a stacked sphere has the componentwise minimal f-vector. Here we provide a sharpening of these bounds for the class of balanced homology spheres in Theorem 5.3. In the case of balanced $(d-1)$-spheres whose number of vertices, n, is divisible by d, our result amounts to the statement that the spheres obtained by taking the connected sum of $n/d-1$ copies of the boundary of the d-dimensional cross-polytope have the componentwise minimal f-vector.

The rest of the paper is structured as follows. In Section 2 we review basic facts and definitions related to simplicial complexes and their face numbers. Section 3 is devoted to complexes without large missing faces. Section 4 deals with CM connectivity of skeletons of flag complexes. Finally, in Section 5 we discuss balanced complexes. Sections 3–5 are independent of each other and can be read in any order. We hope that our results will be helpful in attacking additional stronger conjectures proposed in [14].

2. Preliminaries

Here we review basic facts and definitions related to simplicial complexes. An excellent reference to this material is Stanley's book [17].

Let Δ be a simplicial complex on the vertex set $[n]$. For $F\in\Delta$, set $\dim F:=|F|-1$ and define the *dimension* of Δ, $\dim\Delta$, as the maximal dimension of its faces. We say that Δ is *pure* if all of its facets (maximal faces under inclusion) have the same dimension. The *f-vector* of Δ is $f(\Delta)=(f_{-1},f_0,\dots,f_{d-1})$, where $d-1=\dim\Delta$ and f_j is the number of j-dimensional faces of Δ. Thus $f_{-1}=1$ (unless Δ is the empty complex) and $f_0=n$. We also consider the *f-polynomial* of Δ,

$$f(\Delta,x):=\sum_{j=0}^{d} f_{j-1}x^j.$$

It is sometimes more convenient to work with the *h-vector*,

$$h(\Delta)=(h_0,h_1,\dots,h_d),$$

(or the h-polynomial, $h(\Delta,x):=\sum_{j=0}^{d}h_jx^j$) instead of the f-vector (f-polynomial, resp.). It carries the same information as the f-vector and is defined by the following relation:

$$h(\Delta,x)=(1-x)^d f\left(\Delta,\frac{x}{1-x}\right).$$

In particular, $h_0=1$, $h_1=n-d$, and the f-numbers of Δ are non-negative linear combinations of its h-numbers.

Let Δ_1 and Δ_2 be simplicial complexes on disjoint vertex sets V_1 and V_2. Then their *join* is the following simplicial complex on $V_1\cup V_2$,

$$\Delta_1*\Delta_2:=\{F_1\cup F_2 : F_1\in\Delta_1 \text{ and } F_2\in\Delta_2\}.$$

Therefore,

$$f(\Delta_1 * \Delta_2, x) = f(\Delta_1, x) f(\Delta_2, x) \quad \text{and} \quad h(\Delta_1 * \Delta_2, x) = h(\Delta_1, x) h(\Delta_2, x).$$

Also, a set $F \subseteq V_1 \cup V_2$ is a missing face of $\Delta_1 * \Delta_2$ if and only if it is a missing face of either Δ_1 or Δ_2. Thus if both complexes have no missing faces of dimension larger than i, then so does their join.

Similarly, if Δ_1 and Δ_2 are pure simplicial $(d-1)$-dimensional complexes on disjoint vertex sets, and $F_1 = \{v_1, ..., v_d\} \in \Delta_1$ and $F_2 = \{w_1, ..., w_d\} \in \Delta_2$ are facets, then the complex obtained from Δ_1 and Δ_2 by identifying F_1 and F_2 via the bijection $\rho(v_i) = w_i$, and then removing this identified face, is called the *connected sum* of Δ_1 and Δ_2 along F_1 and F_2, and is denoted $\Delta_1 \#_\rho \Delta_2$. While the combinatorics of the resulting complex depends on F_1, F_2, and ρ, its f- and h-vectors do not:

$$h_i(\Delta_1 \# \Delta_2) = \begin{cases} h_i(\Delta_1) + h_i(\Delta_2) - 1, & \text{if } i = 0 \text{ or } d, \\ h_i(\Delta_1) + h_i(\Delta_2), & \text{if } 0 < i < d. \end{cases}$$

If Δ is a simplicial complex and F is a face of Δ, then the *link* of F in Δ is $\operatorname{lk}_\Delta F = \operatorname{lk} F := \{G \in \Delta : F \cup G \in \Delta \text{ and } F \cap G = \varnothing\}$, the *star* of F in Δ is $\operatorname{st}_\Delta F = \operatorname{st} F := \{G \in \Delta : F \cup G \in \Delta\}$, and the *antistar* of F in Δ is $\operatorname{ast}_\Delta F = \operatorname{ast} F = \{G \in \Delta : F \not\subseteq G\}$. Also, for $W \subseteq [n]$, let $\Delta_{-W} := \{F \in \Delta : F \subseteq [n] \setminus W\}$ denote the *restriction* of Δ to $[n] \setminus W$. The links, stars, antistars, and restrictions are simplicial complexes in their own right. If Δ is a complex without missing faces of dimension larger than i, then so are links, stars, and restrictions of Δ; furthermore this property is preserved under taking antistars of faces of dimension at most i.

We say that a $(d-1)$-dimensional complex Δ is *Cohen–Macaulay* over $\mathbf{k}$ (CM, for short) if $\widetilde{H}_i(\operatorname{lk} F; \mathbf{k}) = 0$ for all $F \in \Delta$ and all $i < d - |F| - 1$. Here $\mathbf{k}$ is either a field or $\mathbb{Z}$ and $\widetilde{H}_i(\,\cdot\,, \mathbf{k})$ denotes the ith reduced simplicial homology with coefficients in $\mathbf{k}$. If in addition, $\widetilde{H}_{d-|F|-1}(\operatorname{lk} F; \mathbf{k}) \cong \mathbf{k}$ for every $F \in \Delta$, then Δ is a $\mathbf{k}$*-homology sphere*. We say that Δ is q*-CM* if for all $W \subset [n]$, $|W| \le q - 1$, the complex Δ_{-W} is CM and has the same dimension as Δ. 2-CM complexes are also known as doubly CM complexes. Every simplicial sphere (that is, a simplicial complex whose geometric realization is homeomorphic to a sphere) is a homology sphere (over any $\mathbf{k}$), and every $\mathbf{k}$-homology sphere is doubly CM over $\mathbf{k}$. Moreover, joins and connected sums of (homology) spheres are (homology) spheres.

Similarly, we say that Δ is *homotopy Cohen–Macaulay* (homotopy CM, for short) if $\operatorname{lk} F$ is $(d - |F| - 2)$-connected for all $F \in \Delta$, and that Δ is q*-homotopy CM* if Δ_{-W} is homotopy CM and has the same dimension as Δ for all $W \subset [n]$, $|W| \le q - 1$. (Recall that a complex, or more precisely its geometric realization, is i*-connected* if all of its homotopy groups from 0th to the ith one vanish.) Unlike the usual Cohen–Macaulayness, homotopy Cohen–Macaulayness is not a topological property: there exist simplicial spheres that are not homotopy CM. It is however worth pointing out that all PL simplicial spheres are homotopy CM (in fact, 2-homotopy CM).

Two simplicial complexes are said to be *PL homeomorphic* if there exists a piecewise linear map between their geometric realizations that is also a homeomorphism. A simplicial complex is a *PL* $(d-1)$-*sphere* if it is PL homeomorphic to the boundary of the d-simplex. The importance of PL spheres is that all their links are also PL spheres (see e.g. [5, Section 12(2)]).

3. Counting face numbers

The goal of this section is to prove the following result conjectured in [14]. Throughout this section we fix positive integers i and d and write $d=qi+r$, where q and r are (uniquely defined) integers satisfying $1\leq r\leq i$. Let σ^j denote the j-dimensional simplex, $\partial\sigma^j$ its boundary complex, and $(\partial\sigma^j)^{*q}$ the join of q copies of $\partial\sigma^j$. Define

$$S(i,d-1):=(\partial\sigma^i)^{*q}*\partial\sigma^r. \tag{1}$$

We remark that $S(1,d-1)$ coincides with the boundary of the d-dimensional cross-polytope.

Theorem 3.1. *Let Δ be a $(d-1)$-dimensional simplicial complex without missing faces of dimension larger than i, and let $\mathbf{k}$ be a field or $\mathbb{Z}$.*

1. *If Δ has a non-vanishing top homology (with coefficients in $\mathbf{k}$), then $f_j(\Delta)\geq f_j(S(i,d-1))$ for all j. Moreover, if $f_0(\Delta)=f_0(S(i,d-1))$ and either d is divisible by i or $f_r(\Delta)=f_r(S(i,d-1))$, then $\Delta=S(i,d-1)$.*

2. *If Δ is 2-CM over $\mathbf{k}$, then $h_j(\Delta)\geq h_j(S(i,d-1))$ for all j. Moreover, if $h_1(\Delta)=h_1(S(i,d-1))$ and either d is divisible by i or $h_{r+1}(\Delta)=h_{r+1}(S(i,d-1))$, then $\Delta=S(i,d-1)$.*

Several cases of Theorem 3.1 are known: Nevo [14, Theorem 1.1] verified the inequalities in part 1 for all j assuming that i divides d, and for all $j\leq r$ if i does not divide d; the $i=1$ case of part 2 is due to Athanasiadis [1, Theorem 1.3].

Throughout the proof, the inequality $P(x)\geq Q(x)$ between two polynomials means that the polynomial $P(x)-Q(x)$ has non-negative coefficients. The proof of both parts relies on the following simple property of the h-numbers of $S(i,d-1)$.

Lemma 3.2. *For every $1\leq s\leq i$, one has*

$$h(S(i,d-1),x)\leq h(\partial\sigma^s,x)h(S(i,d-1-s),x), \tag{2}$$

and hence also

$$f(S(i,d-1),x)\leq f(\partial\sigma^s,x)f(S(i,d-1-s),x).$$

Proof. Since the f-numbers are non-negative combinations of the h-numbers, it is enough to verify the first inequality. Express d as $d=s+(q'i+r')$, where q' and r' are integers satisfying $1\le r'\le i$. Then $q-q'\in\{0,1\}$ and $s+r'=(q-q')i+r$. Since $h(\partial\sigma^j,x)=\sum_{l=0}^{j}x^l$, the inequality in (2) divided by $(\sum_{l=0}^{i}x^l)^{q'}$ reads

$$\left(\sum_{l=0}^{i}x^l\right)^{q-q'}\left(\sum_{l=0}^{r}x^l\right)\le\left(\sum_{l=0}^{r'}x^l\right)\left(\sum_{l=0}^{s}x^l\right).$$

If $q=q'$, then $r=r'+s$, and the above inequality holds without equality. Otherwise, $q-q'=1$ and $i+r=r'+s$ with $i=\max(r,r',s,i)$, and the assertion follows by comparing coefficients. □

Proof of Theorem 3.1, *part* 1. We first prove the inequalities on the f-numbers of Δ by induction on d. If $d\le i$, then $S(i,d-1)=\partial\sigma^d$, and the result follows from the well-known and easy-to-prove fact that among all simplicial complexes of dimension $d-1$ with non-vanishing top homology, $\partial\sigma^d$ has the componentwise minimal f-vector. So assume that $d>i$ and that the statement holds for all $d'<d$.

If F is a face with $0\le\dim F\le i$ and $\widetilde{H}_{d-1-|F|}(\mathrm{lk}\,F;\mathbf{k})=0$, then consider $\Delta':=\mathrm{ast}\,F$ and $\Delta'':=\mathrm{st}\,F$, so that $\Delta=\Delta'\cup\Delta''$ and $\Delta'\cap\Delta''=\partial\overline{F}*\mathrm{lk}\,F$. (Here $\overline{F}$ denotes the simplex F together with all its faces.) Since

$$\widetilde{H}_{d-2}(\Delta'\cap\Delta'';\mathbf{k})\cong\widetilde{H}_{d-1-|F|}(\mathrm{lk}\,F;\mathbf{k})=0$$

and since Δ'' is a cone, and hence acyclic, the Mayer–Vietoris sequence [6, p. 229] yields that $\widetilde{H}_{d-1}(\Delta';\mathbf{k})\cong\widetilde{H}_{d-1}(\Delta;\mathbf{k})\ne0$. Therefore, by considering Δ' instead of Δ, we may assume without loss of generality that every face F of Δ with $\dim F\le i$ satisfies $\widetilde{H}_{d-1-|F|}(\mathrm{lk}\,F;\mathbf{k})\ne0$.

Let G be a missing face of Δ and consider $G'\subsetneq G$. Define $\Delta^{G'}$ to be the collection of faces of Δ of the form $G'\cup F$, where $F\cap G=\varnothing$. Note that $\Delta^{G'}$ is not generally a simplicial complex. Since $H\cap G=G'$ for all $H\in\Delta^{G'}$, the collections $\Delta^{G'}$ are pairwise disjoint as G' ranges over all proper subsets of G. For $G'\subset G$, choose $G'\subseteq G''\subset G$ satisfying $|G''|=|G|-1$. Since G is a missing face in Δ, $\mathrm{lk}\,G''$ does not contain any vertices from G, and therefore $F\cup G'\in\Delta^{G'}$ for all $F\in\mathrm{lk}\,G''$. Hence $f(\Delta^{G'},x)\ge x^{|G'|}f(\mathrm{lk}\,G'',x)\ge x^{|G'|}f(S(i,d-|G|),x)$ by the inductive hypothesis. As the collections $\Delta^{G'}$ are pairwise disjoint for $G'\subsetneq G$, by summing over all such G', we obtain

$$\begin{aligned}f(\Delta,x)&\ge\sum_{G'\subsetneq G}x^{|G'|}f(S(i,d-|G|),x)=f(S(i,d-|G|),x)\sum_{G'\subsetneq G}x^{|G'|}\\&=f(S(i,d-|G|),x)f(\partial\overline{G},x)\ge f(S(i,d-1),x),\end{aligned}$$

where the last step is by Lemma 3.2.

We now prove the statement on equality by induction on d. Assume that $f_0(\Delta)=f_0(S(i,d-1))$ and, if $r<i$, that $f_r(\Delta)=f_r(S(i,d-1))$. Then

$$f_j(\Delta)=\binom{f_0(\Delta)}{j+1}=f_j(S(i,d-1)) \quad \text{for all } j<r.$$

Furthermore, Δ has a missing face of dimension r. In the case that $r<i$ this follows from $f_r(\Delta)=f_r(S(i,d-1))=\binom{f_0(\Delta)}{r+1}-1$. When $r=i$, this follows from the fact that Δ has a complete $(r-1)$-dimensional skeleton and no missing face of Δ has dimension greater than i. Finally, $\widetilde{H}_{d-1-|G|}(\operatorname{lk} G;\mathbf{k})\neq 0$ for all $G\in\Delta$ with $\dim G<r$: otherwise

$$\widetilde{H}_{d-1}(\operatorname{ast} G;\mathbf{k})\neq 0 \quad \text{and} \quad f_{\dim G}(\operatorname{ast} G)<f_{\dim G}(S(i,d-1)),$$

a contradiction.

Let F be a missing face of Δ of dimension r and G be a maximal proper subset of F. We claim that if F' is a missing face in $\operatorname{lk} G$ of dimension i, then F' is a missing face in Δ as well. Let G' be a minimal subface of G such that $\operatorname{lk} G'$ does not contain F' as a face. Then every proper subface of $G'\cup F'$ is a face in Δ, but not $G'\cup F'$ itself. Since $\dim F'=i$, we infer that $G'=\varnothing$ and F' is a missing face in Δ.

We have that $f_0(\operatorname{lk} G)\leq f_0(\Delta)-r-1$, since $\operatorname{lk} G$ contains no vertex of F; and, in fact, equality holds here by the inductive hypothesis since $\operatorname{lk} G$ has non-vanishing top homology. Also $\dim(\operatorname{lk} G)+1=\dim(\Delta)+1-r=d-r$ is divisible by i, and so it follows by the inductive hypothesis that $\operatorname{lk} G=S(i,d-1-r)$. Label the missing faces of $\operatorname{lk} G$ by $F_1,...,F_q$. Every missing face of $\operatorname{lk} G$ has dimension i, and hence every missing face of $\operatorname{lk} G$ is also a missing face of Δ by the previous paragraph. Thus Δ has $F,F_1,...,F_q$ as disjoint missing faces with $\dim F=r$ and $\dim F_1=...=\dim F_q=i$. These are precisely the missing faces of $S(i,d-1)$, and so Δ is contained in $S(i,d-1)$. Since $S(i,d-1)$ has componentwise minimal face numbers, $\Delta=S(i,d-1)$. □

The proof of part 2 utilizes the following results in addition to Lemma 3.2. The first of them is due to Stanley [17, Corollary II.3.2], the second appears in works of Adin, Kalai, and Stanley, see e.g. [16], and the third one is [1, Lemma 4.1].

Lemma 3.3. *If Δ is a $(d-1)$-dimensional CM complex, then $h(\Delta,x)\geq 0$. Moreover, if Δ has a non-vanishing top homology (which happens, for instance, if Δ is a homology sphere, or more generally, if Δ is 2-CM), then $h(\Delta,x)\geq\sum_{l=0}^{d} x^l=h(\partial\sigma^d,x)$.*

Lemma 3.4. *Let Δ be a simplicial complex and Γ be a subcomplex of Δ. If Δ and Γ are both CM (over the same* $\mathbf{k}$*) and have the same dimension, then $h(\Delta,x)\geq h(\Gamma,x)$.*

Lemma 3.5. *Let Δ be a pure simplicial complex and v be a vertex of Δ. If $\operatorname{ast}_\Delta v$ has the same dimension as Δ, then $h(\Delta,x)=xh(\operatorname{lk}_\Delta v,x)+h(\operatorname{ast}_\Delta v,x)$.*

We are now in a position to prove part 2 of the theorem. It follows the same general outline as the proof of [1, Theorem 1.3], but requires a bit more bookkeeping.

Proof of Theorem 3.1, *part* 2. The proof is by induction on d. If $d\leq i$, then $S(i,d-1)=\partial\sigma^d$, and the statement follows from Lemma 3.3. So assume that $d>i$ and that the statement holds for all $d'<d$.

Let $F=\{v_0,v_1,\dots,v_s\}$ be a missing face of Δ (in particular, $s\leq i$). Then $F_j:=\{v_0,v_1,\dots,v_j\}$ is a face for every $-1\leq j\leq s-1$, and so is $F\backslash\{v_j\}:=\{v_0,\dots,\hat{v}_j,\dots,v_s\}$ for every $0\leq j\leq s$. Repeatedly applying Lemma 3.5 and using the fact that $\operatorname{lk}_{\operatorname{lk}G} H=\operatorname{lk}_\Delta(H\cup G)$ for all $G\in\Delta$ and $H\in\operatorname{lk}_\Delta G$ (here and below, lk without a subscript refers to the link in Δ), we obtain

$$\begin{aligned} h(\Delta,x) &= xh(\operatorname{lk}_\Delta v_0,x)+h(\operatorname{ast}_\Delta v_0,x)\\ &= x(xh(\operatorname{lk}_{\operatorname{lk} v_0} v_1,x)+h(\operatorname{ast}_{\operatorname{lk} v_0} v_1,x))+h(\operatorname{ast}_\Delta v_0,x)\\ &= x^2(xh(\operatorname{lk}_{\operatorname{lk} F_1} v_2,x)+h(\operatorname{ast}_{\operatorname{lk} F_1} v_2,x))+xh(\operatorname{ast}_{\operatorname{lk} F_0} v_1,x)+h(\operatorname{ast}_\Delta v_0,x)\\ &= \dots\\ &= x^s h(\operatorname{lk}_\Delta F_{s-1},x)+\sum_{j=0}^{s-1} x^j h(\operatorname{ast}_{\operatorname{lk} F_{j-1}} v_j,x). \end{aligned} \tag{3}$$

Since Δ is 2-CM, all its links are also 2-CM [2], and so all the complexes appearing in (3) are CM. We now show that the h-polynomial of each of these complexes is (componentwise) at least as large as $h(S(i,d-s-1),x)$, and hence

$$\begin{aligned} h(\Delta,x) &\geq \sum_{j=0}^{s} x^j h(S(i,d-s-1),x)=\left(\sum_{j=0}^{s} x^j\right)h(S(i,d-s-1),x)\\ &= h(\partial\sigma^s,x)h(S(i,d-s-1),x)\geq h(S(i,d-1),x) \end{aligned}$$

(by Lemma 3.2), as required.

And indeed, $\operatorname{lk}_\Delta F_{s-1}$ is $(d-s-1)$-dimensional, 2-CM, and has no missing faces of dimension larger than i. Hence $h(\operatorname{lk}_\Delta F_{s-1},x)\geq h(S(i,d-s-1),x)$ by the inductive hypothesis. For all other complexes appearing in (3), observe that since F is a missing face, the complex $v_s * v_{s-1} * \ldots * v_{j+1} * \mathrm{lk}_\Delta(F - v_j)$ is well defined, does not contain v_j, and is contained in $\mathrm{lk}_\Delta F_{j-1}$. In other words,

$$\mathrm{ast}_{\mathrm{lk}\, F_{j-1}}\, v_j \supseteq v_s * v_{s-1} * \ldots * v_{j+1} * \mathrm{lk}_\Delta(F - v_j).$$

As both of these complexes are CM of dimension $d-j-1$, Lemma 3.4 yields that

$$h(\mathrm{ast}_{\mathrm{lk}\, F_{j-1}}\, v_j, x) \geq h(v_s * \ldots * v_{j+1} * \mathrm{lk}(F - v_j), x)$$
$$= h(\mathrm{lk}(F - v_j), x) \geq h(S(i, d-s-1), x),$$

where the last step is by the inductive hypothesis. This completes the proof. The treatment of equality follows from the first part and the observation that $S(i, d-1)$ has a complete $(r-1)$-dimensional skeleton. □

4. Cohen–Macaulay connectivity of flag complexes

This section is devoted to the proof of the following theorem. Recall that the l-skeleton of a simplicial complex Δ, $\mathrm{Skel}_l(\Delta)$, consists of all faces of Δ of dimension at most l.

Theorem 4.1. *Let Δ be a flag simplicial complex of dimension $d-1$.*
1. *If Δ is 2-CM over* $\mathbf{k}$, *then $\mathrm{Skel}_l(\Delta)$ is $2(d-l)$-CM over* $\mathbf{k}$ *for all $0 \leq l \leq d-1$.*
2. *Moreover, if Δ is a simplicial PL sphere, then $\mathrm{Skel}_l(\Delta)$ is $2(d-l)$-homotopy CM for all $0 \leq l \leq d-1$.*

Throughout the proof, $\|\Delta\|$ stands for the geometric realization of Δ; for $W \subset [n]$, $\overline{W}$ denotes the simplex on the vertex set W together with all its faces, and p_W denotes the barycenter of $\|\overline{W}\|$. If Γ is a subcomplex of Δ, and W is a subset of $[n]$ (but not necessarily a subset of $V(\Gamma)$—the vertex set of Γ), we write Γ_{-W} to denote the restriction of Γ to $V(\Gamma) \setminus W$. We make use of the following observation: for $F \in \Delta$ and $W \subseteq [n] \setminus F$,

$$\mathrm{lk}_{(\mathrm{Skel}_l(\Delta))_{-W}}\, F = (\mathrm{lk}_{\mathrm{Skel}_l(\Delta)}\, F)_{-W} = (\mathrm{Skel}_{l-|F|}(\mathrm{lk}_\Delta F))_{-W}. \tag{4}$$

Proof of part 1. In the following $\mathbf{k}$ is fixed and is suppressed from our notation. The proof is by induction on d. Since Δ is flag and 2-CM, we already know that it has at least $2d$ vertices, and hence that $\mathrm{Skel}_0(\Delta)$ is $2d$-CM. This implies the assertion for $d \leq 2$ as well as for $l=0$ and any d.

Assume now that the statement holds for all $d' < d$. In particular, it holds for all links of non-empty faces of Δ since they are also 2-CM and have dimension strictly smaller than $d-1$. Thus for a non-empty face $F\in\Delta$, the complex $\mathrm{Skel}_{l-|F|}(\mathrm{lk}\, F)$ is $2((d-|F|)-(l-|F|))=2(d-l)$-Cohen–Macaulay. Putting this together with (4) and using that for $j<l$ the jth simplicial homology of $\mathrm{Skel}_l(\Delta)$ coincides with that of Δ, to complete the proof it only remains to show that (i) for every $W\subset[n]$ of size $2(d-l)-1$, Δ_{-W} is at least l-dimensional, and (ii) for all $j<l\leq d-1$ and any subset $W=\{v_1,...,v_k\}\subset[n]$ of size $1\leq k\leq 2(d-l)-1$, the homology $\widetilde{H}_j(\Delta_{-W})$ vanishes.

To verify (i) consider $F\in\Delta_{-W}$ of dimension at most $l-1$. We need to show that F is not a maximal (under inclusion) face in Δ_{-W}. Since the link of F in Δ is a flag 2-CM complex of dimension $\geq d-l-1$, it has at least $2(d-l)>|W|$ vertices. Thus, at least one of these vertices, say v, is not in W, yielding that $F\cup\{v\}\in\Delta_{-W}$ is a larger face.

To prove (ii) we induct on k. There are two possible cases to consider.

Case 1. Every two vertices of W are connected by an edge in Δ (this, for instance, happens if $k=1$).

Since Δ is flag, this condition implies that $W\in\Delta$. Then $\|\Delta_{-W}\|$ is a strong deformation retract of $\|\Delta\|\backslash\|\overline{W}\|$ (see e.g. [5, Lemma 11.15]) which in turn is a strong deformation retract of $\|\Delta\|\backslash p_W$. Since Δ is 2-CM, the latter complex is $(d-2)$-acyclic (this is essentially due to Walker, see [17, Proposition III.3.7]), and the statement follows.

Case 2. Not every two vertices of W form an edge.

By reordering the vertices, if necessary, assume that $\{v_{k-1},v_k\}\notin\Delta$. Consider complexes $\Delta_{k-1}:=\Delta_{-(W-v_k)}$, $\mathrm{st}_{\Delta_{k-1}}\, v_k$, and the intersection $\Delta_{-W}\cap\mathrm{st}_{\Delta_{k-1}}\, v_k=\mathrm{lk}_{\Delta_{k-1}}\, v_k$. The first two complexes have vanishing jth homology: indeed, the star is contractible and for Δ_{k-1} this holds by our inductive hypothesis on k. Also, since v_{k-1} and v_k are not connected by an edge, v_{k-1} is not in the link of v_k. Hence

$$\mathrm{lk}_{\Delta_{k-1}}\, v_k=(\mathrm{lk}_\Delta\, v_k)_{-\{v_1,...,v_{k-2}\}}.$$

But $\dim(\mathrm{lk}_\Delta\, v_k)=d-2$ and $k-2\leq 2(d-l)-3=2((d-1)-l)-1$, so our inductive hypothesis on d applies to $\mathrm{lk}_\Delta\, v_k$ and shows that $\widetilde{H}_{j-1}(\mathrm{lk}_{\Delta_{k-1}}\, v_k)=0$. Finally, since $\Delta_{k-1}=\Delta_{-W}\cup\mathrm{st}_{\Delta_{k-1}}\, v_k$, the appropriate portion of the Mayer–Vietoris sequence [6, p. 229] yields that $\widetilde{H}_j(\Delta_{-W})=0$, and the assertion follows. □

We now turn to part 2 of the theorem. A PL sphere is 2-CM over $\mathbb{Z}$, so part 1 implies vanishing of relevant homology groups computed with coefficients in $\mathbb{Z}$. In particular, all the spaces involved are (path) connected, and this allows us to suppress the base point when discussing homotopy groups. We also write $\pi_j(\Delta)$ instead of $\pi_j(\|\Delta\|)$.

The Hurewicz theorem [6, p. 479] asserts that if Δ is j-connected, $j\geq 1$, then $\pi_{j+1}(\Delta)\cong\widetilde{H}_{j+1}(\Delta;\mathbb{Z})$. In particular, if Δ is simply connected and $\widetilde{H}_i(\Delta;\mathbb{Z})=0$ for all $0\leq i\leq j$, then $\widetilde{H}_{j+1}(\Delta;\mathbb{Z})\cong\pi_{j+1}(\Delta)$. Also, PL spheres are simply connected and their links are PL spheres in their own right. Thus part 2 will follow from part 1 if we can show that for a PL $(d-1)$-sphere Δ and an arbitrary $W\subset[n]$ of size $1\leq k\leq 2(d-2)-1=2d-5$, $\pi_1(\Delta_{-W})=0$. This is done exactly as in the proof of part 1: except that in case 2 one needs to use the Seifert–van Kampen theorem [6, p. 161] instead of the Mayer–Vietoris sequence. It asserts (using notation of case 2 in the proof of part 1) that

$$\pi_1(\Delta_{k-1})\cong\pi_1(\Delta_{-W})*_{\pi_1(\mathrm{lk}_{\Delta_{k-1}}v_k)}\pi_1(\mathrm{st}_{\Delta_{k-1}}v_k).$$

Since by the inductive hypothesis all groups, except possibly $\pi_1(\Delta_{-W})$, in this equation are trivial, it follows that $\pi_1(\Delta_{-W})$ is trivial as well. As for case 1, just notice that a topological sphere with a point removed is a topological ball, and hence contractible.

We close this section with several remarks.

1. In part 2 of the theorem the ‘PL sphere’ condition cannot be relaxed to the ‘triangulated sphere’ one. This can be seen by considering the double suspension of the Poincaré sphere. According to Edwards, see [7], the resulting space is a topological sphere. Now start with any triangulation of the Poincaré sphere, and let Γ be its barycentric subdivision. Then $\Delta=(\partial\sigma^1)^{*2}*\Gamma$ is a flag complex that triangulates Edwards’ sphere. But Δ is not homotopy CM: indeed, some of the edges of Δ have Γ as their link, and Γ is not simply connected.

2. Let $n\geq 2d$ be any integer, and let C_k denote the graph-theoretical cycle on k vertices. Then the complex $S(1,d-1,n):=(\partial\sigma^1)^{*(d-2)}*C_{n-2d+4}$ is a $(d-1)$-dimensional polytopal sphere on n vertices. It is flag, and for all $1\leq l\leq d-1$, its l-skeleton is $2(d-l)$-CM, but not $(2(d-l)+1)$-CM. Thus part 1 of Theorem 4.1 is as strong as one can hope for. This example together with the theorem also adds plausibility to Conjecture 1.4 from [14] asserting that among all flag homology $(d-1)$-spheres on n vertices, $S(1,d-1,n)$ has the smallest face numbers.

3. An immediate consequence of part 1 of the theorem is that if Δ is a $(d-1)$-dimensional flag complex that is k-CM for some $k\geq 2$, then $\mathrm{Skel}_l(\Delta)$ is $(2(d-l-1)+k)$-CM for $0\leq l\leq d-1$. Indeed to show that $\mathrm{Skel}_l(\Delta)_{-W}$ is CM and of dimension l for any $|W|\leq 2(d-l)+k-3$, consider a subset W' of W of size $\min(k-2,|W|)$ and its complement $W''=W\backslash W'$. Then $|W''|\leq 2(d-l)-1$. Since $\Delta_{-W}=(\Delta_{-W'})_{-W''}$, and since $\Delta_{-W'}$ is 2-CM, Theorem 4.1 applied to $\Delta_{-W'}$ and W'' completes the proof.

It is interesting to compare this result with a theorem of Fløystad [8] asserting that if Δ is an arbitrary $(d-1)$-dimensional k-CM simplicial complex, then its l-skeleton is $((d-l-1)+k)$-CM.

5. The lower bound theorem for balanced complexes

In this section we establish tight lower bounds on the face numbers of balanced 2-CM complexes in terms of their dimension and the number of vertices. Recall that a $(d-1)$-dimensional complex Δ on the vertex set V is (*completely*) *balanced* if its 1-dimensional skeleton is d-colorable: that is, there exists a coloring $\varkappa: V \to [d]$ such that for all $F \in \Delta$ and distinct $v, w \in F$, $\varkappa(v) \neq \varkappa(w)$. We assume that a balanced complex Δ comes equipped with such a coloring $\varkappa$. The order complex of a rank-d graded poset is one example of a balanced simplicial complex.

If Δ is a balanced complex and $T \subseteq [d]$, then the T*-rank selected subcomplex* of Δ is $\Delta_T := \{F \in \Delta : \varkappa(F) \subseteq T\}$. We make use of the following basic facts from [15].

Lemma 5.1. *Let Δ be a $(d-1)$-dimensional balanced CM complex. Then for any $T \subseteq [d]$, Δ_T is also CM, and $h_i(\Delta) = \sum_{|T|=i} h_i(\Delta_T)$ for all $0 \leq i \leq d$.*

Since deleting a vertex commutes with taking a rank-selected subcomplex: $(\Delta_T)_{-v} = (\Delta_{-v})_T$ for any v with $\varkappa(v) \in T$, one consequence of the above lemma is that a rank-selected subcomplex of a 2-CM complex is 2-CM as well.

The lower bound theorem for simplicial spheres [4], [11] asserts that among all $(d-1)$-dimensional homology spheres with n vertices, a stacked sphere, $\mathcal{ST}(n, d-1)$, has the componentwise minimal f-vector. A *stacked sphere*, $\mathcal{ST}(n, d-1)$, is defined as the connected sum of $n-d$ copies of the boundary of the d-simplex. Since $h_1(\partial\sigma^d) = h_2(\partial\sigma^d) = 1$ if $d \geq 2$, it follows that for $d \geq 3$,

$$h_1(\mathcal{ST}(n, d-1)) = h_2(\mathcal{ST}(n, d-1)).$$

Therefore, via a well-known reduction due to McMullen, Perles, and Walkup (see [4, Theorem 1] or [11, Section 5]), the proof of the lower bound theorem for $d \geq 3$ reduces to showing that the h-vector of a homology sphere of dimension at least 2 satisfies $h_2 \geq h_1$. Recently, Nevo [13] extended this result to all 2-CM simplicial complexes.

Lemma 5.2. *If Δ is a simplicial 2-CM complex of dimension at least 2, then $h_2(\Delta) \geq h_1(\Delta)$.*

It follows easily from the results of [15] that the boundary of the d-dimensional cross-polytope has the componentwise minimal h-vector among all balanced $(d-1)$-dimensional spheres. This motivates us to define a *stacked cross-polytopal sphere*, $\mathcal{ST}^{\times}(n,d-1)$, for n that is a multiple of d, as the connected sum of $n/d-1$ copies of the boundary complex of the d-dimensional cross polytope. At each step the vertices of the same colors are identified to guarantee that the resulting complex is balanced as well.

What are the h-numbers of $\mathcal{ST}^{\times}(n,d-1)$? Since the h-numbers of the d-dimensional cross-polytope are given by $h_j=\binom{d}{j}$, it follows that for $0<j<d$,

$$h_j(\mathcal{ST}^{\times}(n,d-1))=\left(\frac{n}{d}-1\right)\binom{d}{j},$$

that is,

$$(j+1)h_{j+1}=(d-j)h_j \quad \text{for } 0<j<d-1.$$

In particular, $(d-1)h_1=2h_2$ if $d\geq 3$. Similarly, a direct computation shows that

$$\begin{aligned}\psi_{j-1}(n,d-1) &:= jf_{j-1}(\mathcal{ST}^{\times}(n,d-1))\\ &=\begin{cases}(2^j-1)\binom{d-1}{j-1}(n-d)+d\binom{d-1}{j-1}, & 1\leq j\leq d-1,\\ (2^d-2)(n-d)+2d, & j=d.\end{cases}\end{aligned}$$

One advantage of the last expression is that it is defined for *all* n rather than just multiples of d. This allows us to state and prove the main theorem of this section—the lower bound theorem for balanced spheres and, more generally, balanced 2-CM complexes.

Theorem 5.3. *Let Δ be a balanced 2-CM simplicial complex of dimension $d-1$. If $d\geq 3$, then $2h_2(\Delta)\geq(d-1)h_1(\Delta)$. In particular, if $d\geq 2$ and $f_0(\Delta)=n$, then $jf_{j-1}(\Delta)\geq\psi_{j-1}(n,d-1)$ for all $2\leq j\leq d$.*

Proof. Repeatedly applying Lemma 5.1, we see that

$$\sum_{|T|=3} h_2(\Delta_T)=\sum_{|T|=3}\sum_{\substack{S\subset T\\|S|=2}} h_2(\Delta_S)=\sum_{|S|=2}(d-2)h_2(\Delta_S)=(d-2)h_2(\Delta),$$

and

$$\sum_{|T|=3} h_1(\Delta_T)=\sum_{|T|=3}\sum_{\substack{S\subset T\\|S|=1}} h_1(\Delta_S)=\sum_{|S|=1}\binom{d-1}{2}h_1(\Delta_S)=\binom{d-1}{2}h_1(\Delta).$$

Since Δ is balanced and 2-CM, its rank-selected subcomplexes share the same properties. In particular, when $|T|=3$, Δ_T is a 2-dimensional 2-CM complex, and so by Lemma 5.2, $h_2(\Delta_T)\geq h_1(\Delta_T)$. Thus we infer that $(d-2)h_2(\Delta)\geq\binom{d-1}{2}h_1(\Delta)$, and the inequality $2h_2(\Delta)\geq(d-1)h_1(\Delta)$ (for $d\geq3$) follows.

The proof of the "in particular" part is a routine computation similar in spirit to the McMullen–Perles–Walkup reduction. We sketch it here for completeness. We use induction on d. For $d=2$ we need only show that $2f_1(\Delta)\geq 2n$. This indeed holds, since Δ is a 2-CM graph, hence it is 2-connected, and so every vertex of Δ has degree at least 2.

Suppose now that $d\geq3$. Then

$$\sum_{v\in\Delta} h_1(\operatorname{lk} v) = 2h_2(\Delta)+(d-1)h_1(\Delta)\geq 2(d-1)h_1(\Delta)$$

by the first part. Inductively, for $3\leq j\leq d-1$, we have

$$\begin{aligned}
jf_{j-1}(\Delta) &= \sum_{v\in\Delta} f_{j-2}(\operatorname{lk} v)\\
&\geq \sum_{v\in\Delta}\frac{1}{j-1}\left[(2^{j-1}-1)\binom{d-2}{j-2}h_1(\operatorname{lk} v)+(d-1)\binom{d-2}{j-2}\right]\\
&\geq (2^j-2)\binom{d-1}{j-1}h_1(\Delta)+\binom{d-1}{j-1}f_0(\Delta)\\
&= (2^j-1)\binom{d-1}{j-1}h_1(\Delta)+d\binom{d-1}{j-1}\\
&= \psi_{j-1}(n,d-1).
\end{aligned}$$

The proof for $j=d$ is similar and is omitted. $\square$

It is worth remarking that at present we do not know whether the assertion of Theorem 5.3 is tight when n is not divisible by d. We also do not know if the stacked cross-polytopal spheres are the only balanced 2-CM complexes satisfying $2h_2=(d-1)h_1$ when d divides n.

In the case when Γ is a $2j$-dimensional homology sphere, the Dehn–Sommerville relations assert that $h_j(\Gamma)=h_{j+1}(\Gamma)$. If we knew that every balanced 2-CM complex Γ of dimension $2j$ satisfies $h_j(\Gamma)\leq h_{j+1}(\Gamma)$, a proof similar to that of Theorem 5.3 would imply that for a balanced 2-CM complex Δ of dimension $d-1\geq 2j$,

$$(j+1)h_{j+1}(\Delta)\geq(d-j)h_j(\Delta).$$

Finally, we observe that the lower bound theorem [3], [11] holds not only for simplicial spheres, but also for triangulations of connected manifolds, and even normal pseudomanifolds of dimension at least two. (The latter result is due to Fogelsanger [9].) Does Theorem 5.3 hold for balanced triangulations of such spaces? Using results of [13] and standard tools from rigidity theory, one can show that any connected pure 3-dimensional simplicial complex all of whose vertex links are 2-CM, satisfies $h_2 \geq h_1$. The proof analogous to that of Theorem 5.3 then implies that if Δ is a balanced triangulation of a manifold of dimension at least 3, then $3h_2(\Delta) \geq (d-1)h_1(\Delta)$. This inequality, however, is weaker than that of Theorem 5.3.

Acknowledgements. We are grateful to Eran Nevo for introducing us to some of the problems considered in the paper. Our thanks also go to Christos Athanasiadis, Eran Nevo, and Ed Swartz for helpful discussions. Michael Goff's research was partially supported by a graduate fellowship from NSF Grant DMS-0801152 and Steven Klee's research by a graduate fellowship from VIGRE NSF Grant DMS-0354131.

Added in proof. Theorem 5.3 continues to hold for balanced triangulations of orientable manifolds. This was recently verified by J. Browder and the second author.

References

1. Athanasiadis, C. A., Some combinatorial properties of flag simplicial pseudo-manifolds and spheres, to appear in *Ark. Mat.*
2. Baclawski, K., Cohen–Macaulay connectivity and geometric lattices, *European J. Combin.* **3** (1982), 293–305.
3. Barnette, D., Graph theorems for manifolds, *Israel J. Math.* **16** (1973), 62–72.
4. Barnette, D., A proof of the lower bound conjecture for convex polytopes, *Pacific J. Math.* **46** (1973), 349–354.
5. Björner, A., Topological methods, in *Handbook of Combinatorics* **2**, pp. 1819–1872, Elsevier, Amsterdam, 1995.
6. Bredon, G., *Topology and Geometry*, Graduate Texts in Mathematics **139**, Springer, New York, 1993.
7. Daverman, R., *Decompositions of Manifolds*, Pure and Applied Mathematics **124**, Academic Press, Orlando, FL, 1986.
8. Fløystad, G., Cohen–Macaulay cell complexes, in *Algebraic and Geometric Combinatorics*, Contemp. Math. **423**, pp. 205–220, Amer. Math. Soc., Providence, RI, 2006.
9. Fogelsanger, A., *The Generic Rigidity of Minimal Cycles*, Ph.D. Thesis, Cornell University, Ithaca, NY, 1988.
10. Frohmader, A., Face vectors of flag complexes, *Israel J. Math.* **164** (2008), 153–164.
11. Kalai, G., Rigidity and the lower bound theorem I, *Invent. Math.* **88** (1987), 125–151.
12. Meshulam, R., Domination numbers and homology, *J. Combin. Theory Ser. A* **102** (2003), 321–330.

13. Nevo, E., Rigidity and the lower bound theorem for doubly Cohen–Macaulay complexes, *Discrete Comput. Geom.* **39** (2008), 411–418.
14. Nevo, E., Remarks on missing faces and generalized lower bounds on face numbers, *Electron. J. Combin.* **16** (2009/2010), Research Paper #8, 11 pp.
15. Stanley, R. P., Balanced Cohen–Macaulay complexes, *Trans. Amer. Math. Soc.* **249** (1979), 139–157.
16. Stanley, R. P., A monotonicity property of h-vectors and h^*-vectors, *European J. Combin.* **14** (1993), 251–258.
17. Stanley, R. P., *Combinatorics and Commutative Algebra*, Progress in Mathematics **41**, Birkhäuser, Boston, MA, 1996.
18. Ziegler, G. M., *Lectures on Polytopes*, Graduate Texts in Mathematics **152**, Springer, New York, 1995.

Michael Goff
Department of Mathematics
P.O. Box 354350
University of Washington
Seattle, WA 98195-4350
U.S.A.
mgoff@math.washington.edu

Steven Klee
Department of Mathematics
P.O. Box 354350
University of Washington
Seattle, WA 98195-4350
U.S.A.
klees@math.washington.edu

Isabella Novik
Department of Mathematics
P.O. Box 354350
University of Washington
Seattle, WA 98195-4350
U.S.A.
novik@math.washington.edu

Received July 10, 2009
published online February 23, 2010