Discrete Comput Geom (2009) 42: 155–165
DOI 10.1007/s00454-009-9181-3

On the Graph Connectivity of Skeleta of Convex Polytopes

Christos A. Athanasiadis

Received: 23 June 2008 / Revised: 27 February 2009 / Accepted: 8 April 2009 /
Published online: 12 May 2009
© Springer Science+Business Media, LLC 2009

Abstract Given a d-dimensional convex polytope P and nonnegative integer k not exceeding $d-1$, let $\mathcal{G}_k(P)$ denote the simple graph on the node set of k-dimensional faces of P in which two such faces are adjacent if there exists a $(k+1)$-dimensional face of P which contains them both. The graph $\mathcal{G}_k(P)$ is isomorphic to the dual graph of the $(d-k)$-dimensional skeleton of the normal fan of P. For fixed values of k and d, the largest integer m such that $\mathcal{G}_k(P)$ is m-vertex-connected for all d-dimensional polytopes P is determined. This result generalizes Balinski's theorem on the one-dimensional skeleton of a d-dimensional convex polytope.

Keywords Convex polytope · Balinski's theorem · Incidence graph · Skeleton · Vertex connectivity · Cell complex

1 Introduction

The combinatorial theory of convex polytopes has provided mathematicians with interesting problems since antiquity. Some of these problems concern skeleta of polytopes and their connectivity; see for instance [8].

Balinski's theorem [1] (see also Theorem 2.2) gives a sharp lower bound on the (vertex) connectivity of the abstract graph $\mathcal{G}(P)$ defined by the one-dimensional skeleton of a d-dimensional convex polytope P. This graph can also be understood as the dual graph of the normal fan $\mathcal{N}_P$ of P or, equivalently, as the dual graph of the $(d-1)$-dimensional skeleton of any polytope Q which is polar-dual to P. By definition, these are the simple graphs on the node set of maximal faces of $\mathcal{N}_P$ and of

Supported by the 70/4/8755 ELKE Research Fund of the University of Athens.

C.A. Athanasiadis (✉)
Department of Mathematics (Division of Algebra-Geometry), University of Athens, Panepistimioupolis, 15784 Athens, Greece
e-mail: caath@math.uoa.gr

$(d-1)$-dimensional faces (facets) of Q, respectively, in which two such faces are adjacent if they share a common codimension one face.

It is natural to inquire about the connectivity of the dual graphs of other skeleta of $\mathcal{N}_P$ or Q. These graphs correspond to the simple graphs $\mathcal{G}_k(P)$ defined for nonnegative integers k 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. In other words, $\mathcal{G}_k(P)$ is the graph on the node set of rank $k+1$ elements in the face lattice of P, two such elements being adjacent if they have a common cover in this lattice. For $k=0$, the graph $\mathcal{G}_k(P)$ reduces to the graph $\mathcal{G}(P)$ which appears in Balinski's theorem.

It is folklore that the graphs $\mathcal{G}_k(P)$ are connected; see [8, Theorem 19.5.2]. Their higher (vertex) connectivity is the subject of this paper. Let m be a positive integer and recall that an abstract graph $\mathcal{G}$ is said to be m-connected if $\mathcal{G}$ has at least $m+1$ nodes and every graph obtained from $\mathcal{G}$ by deleting $m-1$ or fewer nodes and their incident edges is connected. Our main result is as follows.

Theorem 1.1 *For fixed nonnegative integers k and d satisfying $0 \le k \le d-1$, let $m_k(d)$ denote the largest integer m such that the graph $\mathcal{G}_k(P)$ is m-connected for all convex polytopes P of dimension d. We have*

$$m_k(d) = \begin{cases} d, & \text{if } k = d-2 \\ (k+1)(d-k), & \text{otherwise.} \end{cases}$$

A few remarks on Theorem 1.1 are in order. Every node of an m-connected graph $\mathcal{G}$ must have at least m neighbors, since the graph obtained from $\mathcal{G}$ by removing all neighbors of a given node either has a single node or else is disconnected. Theorem 1.1 is made plausible by the fact that every k-dimensional face F of a d-dimensional polytope P has at least $(k+1)(d-k)$ neighbors in $\mathcal{G}_k(P)$. Indeed, such a face F is contained in at least $d-k$ faces of P of dimension $k+1$, each one of those has at least $k+1$ faces of dimension k other than F and all these k-dimensional faces of P are pairwise distinct and are neighbors of F in $\mathcal{G}_k(P)$. On the other hand, if P is a d-dimensional simplex, then every k-dimensional face of P has exactly $(k+1)(d-k)$ neighbors in $\mathcal{G}_k(P)$. Therefore, $\mathcal{G}_k(P)$ is not m-connected for any value of m which exceeds $(k+1)(d-k)$. This example shows that $m_k(d) \le (k+1)(d-k)$ for all k and d.

Theorem 1.1 reduces to Balinski's theorem in the case $k=0$ and is trivial for $k=d-1$, since $\mathcal{G}_{d-1}(P)$ is the complete graph on the set of facets of P and every d-dimensional polytope has at least $d+1$ facets.

The connectivity of skeleta of polytopes and more general cell complexes was studied by Fløystad [6] from a homological point of view and, more recently, by Björner [4] from a homotopy-theoretic point of view. Earlier, Barnette [2] had shown that Balinski's theorem can be extended to the one-dimensional skeleta of certain cell decompositions of manifolds (as well as more general objects). Motivated by the results of [6], we conjecture (see Conjecture 6.2) that for $0 \le k \le d-2$, Theorem 1.1 generalizes to the class of finite Cohen–Macaulay regular cell complexes with the intersection property.

This paper is structured as follows. Section 2 reviews graph-theoretic terminology and basic background on convex polytopes. The special cases $k = d - 2$ and $k = 1$ of Theorem 1.1 are proved in Sects. 3 and 4, respectively. The proof of the theorem is completed in Sect. 5. Except for the case $k = d - 2$, in which a direct combinatorial argument is given, the proof proceeds by induction on k. For the inductive step, the induction hypothesis is applied to appropriate vertex figures of P. In particular, Balinski's theorem is used in the proof of the special case $k = 1$ (and elsewhere). Section 6 discusses possible generalizations to classes of cell complexes. As mentioned earlier, it is conjectured that for $0 \le k \le d - 2$, Theorem 1.1 generalizes to all finite Cohen–Macaulay regular cell complexes with the intersection property. This conjecture is verified (i) for finite Cohen–Macaulay polyhedral complexes and (ii) for $k = d - 2$. The proof relies on Theorem 1.1 in the former case and proceeds by adapting the argument of Sect. 3 in the latter.

Remark 1.2 After this paper was written, it was pointed out to the author by Ronald Wotzlaw that the graphs $\mathcal{G}_k(P)$ have been previously studied by Sallee [10]. In the notation of Theorem 1.1, it is proved in [10] (see (7.18) on p. 495) that

$$(k+1)(d-k) - k \le m_k(d) \le (k+1)(d-k).$$

More generally, given integers $0 \le r < s \le d - 1$, upper and lower bounds are given in [10] for the connectivity of the graph on the node set of r-dimensional faces of a d-dimensional convex polytope P, in which two such faces are adjacent if there exists an s-dimensional face of P which contains them both (thus our setting corresponds to the case $r = k$ and $s = k+1$). Upper and lower bounds are obtained in [10] for various other notions of connectivity for the incidence graphs between faces of dimension r and faces of dimension s of d-dimensional polytopes.

2 Preliminaries

All graphs considered in this paper will be simple (without loops or multiple edges) and finite. Thus every edge of a graph $\mathcal{G}$ connects two distinct nodes of $\mathcal{G}$, called its *endpoints*, and is said to be *incident* to each of these nodes. Two nodes of $\mathcal{G}$ are said to be *adjacent* if they are connected by an edge in $\mathcal{G}$. A *walk* of length n in $\mathcal{G}$ is an alternating sequence $w = (v_0, e_1, v_1, \ldots, e_n, v_n)$ of nodes and edges, such that v_{i-1} and v_i are the endpoints of e_i for $1 \le i \le n$. We say that w *connects* nodes v_0 and v_n, which are the *endpoints* of w. Thus $\mathcal{G}$ is connected if any two nodes can be connected by a walk in $\mathcal{G}$. Given a subset V of the set of nodes of $\mathcal{G}$, we denote by $\mathcal{G}\backslash V$ the graph obtained from $\mathcal{G}$ by deleting the nodes in V and all edges of $\mathcal{G}$ incident to these nodes. Given a positive integer m, the graph $\mathcal{G}$ is said to be *m-connected* if it has at least $m + 1$ nodes and $\mathcal{G}\backslash V$ is connected for all subsets V of the set of nodes of $\mathcal{G}$ with cardinality at most $m - 1$.

A *convex polytope* P is defined as the convex hull of a finite set of points in $\mathbb{R}^N$. The *dimension* of P is the dimension of the affine hull of P, as an affine subspace of $\mathbb{R}^N$. A *face* of P is either a subset of P on which some linear functional on $\mathbb{R}^N$ achieves its minimum on P or else the empty set. Every face of P is a polytope in $\mathbb{R}^N$

and the intersection of two faces of P is again a face of P. Faces of P of dimension zero or one are called *vertices* or *edges*, respectively, and faces of codimension one are called *facets*. Every edge has exactly two vertices, also called its *endpoints*. The following lemma is well known; see, for instance [9, Sect. 1.2].

Lemma 2.1 *Any d-dimensional convex polytope has at least $\binom{d+1}{i+1}$ faces of dimension i for all nonnegative integers i.*

The *graph of* P, denoted by $\mathcal{G}(P)$, is the abstract graph which has as nodes the vertices of P and in which two nodes are adjacent if they are the endpoints of an edge of P. We will need the following slightly stronger version of Balinski's theorem, which follows, for instance, from the proof given in [12, Sect. 3.5].

Theorem 2.2 (Balinski [1]) *Given a d-dimensional convex polytope $P \subset \mathbb{R}^N$, the graph $\mathcal{G}(P)\backslash V$ is connected for every subset V of the vertex set of P which is contained in some $(d-2)$-dimensional affine subspace of $\mathbb{R}^N$. In particular, $\mathcal{G}(P)$ is d-connected.*

For every convex polytope P, there exists a polytope Q (necessarily of the same dimension) whose set of faces is in inclusion-reversing bijection with the set of faces of P. Such a polytope is said to be *polar-dual* to P. Given an r-dimensional face F of a d-dimensional convex polytope P, the set of faces of P which contain F is in inclusion-preserving bijection with the set of faces of a $(d-r-1)$-dimensional polytope P_F, called a *face figure* (or *vertex figure* if F is a vertex) of P at F. It is clear that walks in the graph $\mathcal{G}_k(P_F)$ (as defined in the Introduction) correspond bijectively to walks in $\mathcal{G}_{k+r+1}(P)$ which involve only faces of P containing F. For more information on convex polytopes and their combinatorial structure we refer the interested reader to [7, 9, 12].

3 The Case $k = d - 2$

In this section we restate and prove Theorem 1.1 in the case $k = d - 2$.

Proposition 3.1 *The graph $\mathcal{G}_{d-2}(P)$ is d-connected for all convex polytopes P of dimension d. Moreover, in any dimension $d \geq 2$ there exist convex polytopes for which $\mathcal{G}_{d-2}(P)$ is not $(d+1)$-connected.*

It is perhaps easier to visualize the graph $\mathcal{G}_{d-2}(P)$ in terms of a polytope Q which is polar-dual to P. Since faces of P of dimension $d-2$ and $d-1$ correspond to edges and vertices of Q, respectively, $\mathcal{G}_{d-2}(P)$ is isomorphic to the graph $\Gamma(Q)$ on the node set of edges of Q in which two nodes are adjacent if they have a vertex of Q as a common endpoint. The graph $\Gamma(Q)$ is commonly referred to as the *line graph* of the one-skeleton of Q.

Proof of Proposition 3.1 Let Q be a polytope polar-dual to P, so that $\mathcal{G}_{d-2}(P)$ may be replaced by the graph $\Gamma(Q)$ on the node set of edges of Q. To show that $\Gamma(Q)$ is d-connected, let $\mathcal{E}$ be any subset of the set of edges of Q of cardinality at most $d-1$. We

observe first that given any $e \in \mathcal{E}$, it is possible to choose some two-dimensional face F of P containing e, so that e is the unique edge of F which belongs to $\mathcal{E}$. Indeed, in view of our assumption on the cardinality of $\mathcal{E}$, this is so because e is contained in at least $d-1$ faces of P of dimension 2 and any two such faces share no edge other than e in common. Since $\mathcal{G}(Q)$ is connected, so is $\Gamma(Q)$. Our previous remark shows that any walk w in $\Gamma(Q)$ connecting two nodes not in $\mathcal{E}$ can be transformed to one connecting the same nodes that does not involve elements of $\mathcal{E}$. This can be done by replacing any node $e \in \mathcal{E}$ that may appear in w with the sequence of edges other than e (and vertices other than the endpoints of e), ordered appropriately, of a two-dimensional face of P which contains e but no other edge in $\mathcal{E}$. It follows that $\Gamma(Q)\backslash\mathcal{E}$ is connected and hence that $\Gamma(Q)$ is d-connected.

To prove the second statement in the proposition, let I be a line segment, let $Q = \Delta \times I$ be the prism over a $(d-1)$-dimensional simplex Δ and denote by $\mathcal{E}$ the set of edges of Q of the form $v \times I$, where v is a vertex of Δ. It is clear that the set $\mathcal{E}$ has d elements and that the graph $\Gamma(Q)\backslash\mathcal{E}$ is disconnected. This implies that $\Gamma(Q)$ is not $(d+1)$-connected. □

4 The Case $k = 1$

In this section we restate and prove Theorem 1.1 in the case $k = 1$.

Proposition 4.1 *The graph $\mathcal{G}_1(P)$ is $(2d-2)$-connected for all convex polytopes P of dimension $d \geq 4$.*

Throughout this section P will be a convex polytope of dimension $d \geq 4$, $\mathcal{E}$ will be a subset of the edge set of P of cardinality less than $2d-2$ and f and g will be distinct edges of P not in $\mathcal{E}$. For any vertex v and any edge e of P, we denote by $s(v)$ the number of edges in $\mathcal{E}$ which have v as an endpoint and by $t(e)$ the number of edges in $\mathcal{E}$ which have at least one common endpoint with e. The following technical lemma will be used in the proof.

Lemma 4.2 *There exists a walk $w = (v_0, e_1, v_1, \ldots, e_n, v_n)$ in $\mathcal{G}(P)$ such that v_0 is an endpoint of f, v_n is an endpoint of g and the following hold*:

(a) $s(v_i) \leq d-2$ *for* $0 \leq i \leq n$ *and*
(b) $t(e_i) \leq d-1$ *for all* $1 \leq i \leq n$ *with* $e_i \in \mathcal{E}$.

Proof Let us call a vertex of P *bad* if it is an endpoint of at least $(d+1)/2$ edges in $\mathcal{E}$ and *good* otherwise. We will also call an edge e of P *bad* if it violates condition (b), meaning that $e \in \mathcal{E}$ and $t(e) \geq d$. Otherwise, we call e *good*. Clearly, every good vertex satisfies condition (a) and every bad edge has at least one bad endpoint. As a result, every walk in $\mathcal{G}(P)$ which only goes through good vertices satisfies conditions (a) and (b).

Let p denote the number of bad vertices of P and let q denote the cardinality of $\mathcal{E}$. Since each edge of P has only two endpoints, we have $p\lceil(d+1)/2\rceil \leq 2q \leq 2(2d-3)$. From this and our assumption $d \geq 4$ it follows that $p \leq d-1$, that is,

there exist at most $d-1$ bad vertices. Therefore, by Theorem 2.2, deleting all bad vertices of P from $\mathcal{G}(P)$ and their incident edges results in a connected graph. Thus, it suffices to show that each of the edges f and g either has a good endpoint or else one of its endpoints, say v, satisfies $s(v) \le d-2$ and is connected to a good vertex by a good edge of P. We will prove this statement for f, the same arguments applying for g. Let a and b be the endpoints of f. We distinguish two cases.

Case 1: At least one of the endpoints of f, say a, satisfies $s(a) \ge d-1$. Since $\mathcal{E}$ has at most $2d-3$ elements and f is not in $\mathcal{E}$, we must have $s(b) \le d-2$. As a result, there exists at least one edge e of P other than f which has b as an endpoint and does not belong to $\mathcal{E}$. Since $s(a) \ge d-1$, at most $d-1$ edges in $\mathcal{E}$ have a common endpoint with e. Hence, at least one of the endpoints of e is good. Since e is a good edge, we are done in this case.

Case 2: We have $s(a) \le d-2$ and $s(b) \le d-2$. As a result, each of a, b is an endpoint of at least one edge of P not in $\mathcal{E}$, other than f. There is nothing to prove if at least one of a, b is good, so we may assume that both of them are bad.

Suppose first that there exist distinct vertices a' and b' of P other than a, b which are connected to a and b, respectively, by edges of P not in $\mathcal{E}$. Since $\mathcal{E}$ has at most $2d-3$ elements and at most three of them have both their endpoints in $\{a, b, a', b'\}$, we must have $s(a)+s(b)+s(a')+s(b') \le 2d$. Since a and b are bad, this implies that $s(a')+s(b') \le d-1$. As a result, at least one of a' and b', say a', is good. Since $s(a) \le d-2$ and a is connected to a' by an edge of P not in $\mathcal{E}$, the desired statement follows.

Otherwise, a and b must be connected to a vertex c of P with edges not in $\mathcal{E}$ and we must have $s(a) = s(b) = d-2$. It follows that $s(c) \le 1$ and that c is good, so we are done in this case as well. □

Remark 4.3 The conclusion in the previous proof that there exist at most $d-1$ bad vertices can fail in the case $d=3$. For example, if P is the bipyramid over a triangle Δ and $\mathcal{E}$ is the set of edges of Δ, then $q = 2d-3 = 3$ and all three vertices of Δ are bad, so that $p = 3 = d$.

Proof of Proposition 4.1 Recall that the set of nodes of the graph $\mathcal{G}_1(P)$ coincides with the edge set of P. Let $\mathcal{E}$ be any subset of this set of cardinality less than $2d-2$, as before, and let f and g be any two edges of P not in $\mathcal{E}$. We need to show that f and g can be connected by a walk in the graph $\mathcal{G}_1(P)\backslash\mathcal{E}$.

Let $w = (v_0, e_1, v_1, \ldots, e_n, v_n)$ be a walk as in Lemma 4.2. For any index $1 \le i \le n$ with $e_i \in \mathcal{E}$, there exist at least $d-1$ two-dimensional faces of P which contain e_i and any two of them have no edge other than e_i in common. Because of condition (b), in at least one of these faces none of the two edges which share exactly one common endpoint with e_i belongs to $\mathcal{E}$. Therefore, we may choose edges g_i and f_i of P not in $\mathcal{E}$ which are adjacent nodes in the graph $\mathcal{G}_1(P)$, so that v_{i-1} is an endpoint of g_i and v_i is an endpoint of f_i. We set $g_i = f_i = e_i$ for those indices $1 \le i \le n$ for which e_i is not an element of $\mathcal{E}$. We also set $f_0 = f$ and $g_{n+1} = g$.

For $0 \le i \le n$ we observe that f_i and g_{i+1} are edges of P not in $\mathcal{E}$ which have v_i as a common endpoint. We denote by P_i the vertex figure of P at v_i and by V_i the set of vertices of P_i which correspond to elements of $\mathcal{E}$ having v_i as an endpoint.

Since, by Theorem 2.2, the graph $\mathcal{G}(P_i)$ is $(d-1)$-connected and, by condition (a), the set V_i has no more than $d-2$ elements, the vertices of P_i which correspond to f_i and g_{i+1} can be connected by a walk in $\mathcal{G}(P_i)\backslash V_i$. This implies that f_i and g_{i+1} can be connected by a walk in the graph $\mathcal{G}_1(P)\backslash\mathcal{E}$ for $0 \le i \le n$. Therefore, $f_0 = f$ and $g_{n+1} = g$ can also be connected by a walk in this graph. □

5 Proof of Theorem 1.1

In this section we complete the proof of our main theorem. The structure of the proof is similar to that of Proposition 4.1, given in Sect. 4.

Proof of Theorem 1.1 Let us write $n_k(d) = (k+1)(d-k)$. We have already remarked in the Introduction that there exist d-dimensional convex polytopes P such that $\mathcal{G}_k(P)$ is not m-connected for $m > n_k(d)$ and that Theorem 1.1 is trivial for $k = d-1$. In view of Proposition 3.1, it remains to show that for $k \ge 1$, the graph $\mathcal{G}_k(P)$ is $n_k(d)$-connected for all convex polytopes P of dimension $d \ge k+3$.

We proceed by induction on k, where the case $k = 1$ was treated by Proposition 4.1. Assume that $k \ge 2$. Let U be any subset of the set of k-dimensional faces of P of cardinality less than $n_k(d)$ and let F and G be two k-dimensional faces of P not in U. We need to show that F and G can be connected by a walk in $\mathcal{G}_k(P)\backslash U$.

Claim: There exists a walk w in $\mathcal{G}(P)$ which connects a vertex of F to a vertex of G and has the following properties:

(a) no edge of w is contained in $\binom{d-1}{k-1}$ or more faces in U and
(b) no node of w belongs to $n_{k-1}(d-1)$ or more faces in U.

Given the claim, the proof proceeds as follows. Let $w = (v_0, e_1, v_1, \ldots, e_n, v_n)$ be a walk as in the claim and set $F_0 = F$ and $F_{n+1} = G$. It follows from Lemma 2.1 that each edge of P is contained in at least $\binom{d-1}{k-1}$ faces of P of dimension k. Therefore, in view of our condition (a), for each index $1 \le i \le n$ we may choose a k-dimensional face F_i of P not in U which contains the edge e_i. Note that v_i is a vertex of both F_i and F_{i+1} for $0 \le i \le n$. Let P_i denote the vertex figure of P at v_i and let U_i denote the set of $(k-1)$-dimensional faces of P_i which correspond to the faces of U containing v_i. By the induction hypothesis on k, the graph $\mathcal{G}_{k-1}(P_i)$ is $n_{k-1}(d-1)$-connected. By condition (b), this implies that $\mathcal{G}_{k-1}(P_i)\backslash U_i$ is connected and hence that F_i and F_{i+1} can be connected by a walk in the graph $\mathcal{G}_k(P)\backslash U$. Since this holds for all $0 \le i \le n$, we conclude that $F_0 = F$ and $F_{n+1} = G$ can be connected by a walk in $\mathcal{G}_k(P)\backslash U$ as well. It follows that $\mathcal{G}_k(P)\backslash U$ is connected and hence that $\mathcal{G}_k(P)$ is $n_k(d)$-connected, as desired. It therefore suffices to prove the claim. We distinguish two cases.

Case 1: $k = 2$. We are given that $d \ge 5$ and that U is a set of two-dimensional faces of P of cardinality less than $n_k(d) = 3d-6$ and note that $n_{k-1}(d-1) = 2d-4$.

Let us call an edge or vertex of P *bad* if this edge or vertex is contained in at least $d-1$ or $2d-4$, respectively, elements of U. The following hold: (i) there exist at most two bad edges of P, (ii) there exist at most two bad vertices of P and (iii) if v is a bad vertex of P and e is a bad edge, then v is an endpoint of e. Indeed, the existence

of three bad edges of P would require at least $3d-6$ elements of U, since given any two edges of a polytope P, there exists at most one two-dimensional face of P which contains both of these edges. In view of our assumption on the cardinality of U, this proves (i). We next observe that if u and v are distinct bad vertices of P, then there exist at least $d-1$ elements of U which contain both u and v. Therefore, any two bad vertices are connected by a bad edge and (ii) follows from (i). Finally, if v is a bad vertex of P and e is a bad edge, then there must be at least two elements of U which contain both v and e. The intersection of these has to equal e and contains v. This proves (iii).

It follows from facts (i)–(iii) that one can choose a set V consisting of at most two vertices of P such that $\mathcal{G}(P)\backslash V$ contains no bad vertex or edge. This completes the proof of the claim in this case since $\mathcal{G}(P)\backslash V$ is connected, by Theorem 2.2, and any walk in this graph connecting a vertex of F to a vertex of G satisfies conditions (a) and (b).

Case 2: $k \geq 3$. Since we have $\binom{d-1}{k-1} \geq k(d-k) = n_{k-1}(d-1)$ for $d \geq k+3 \geq 6$, condition (a) follows from (b) and can thus be ignored.

Let V be the set of vertices of P which belong to at least $n_{k-1}(d-1) = k(d-k)$ faces in U. We will show that any $k+1$ vertices, say $v_0, v_1, \ldots, v_k$, in V are affinely dependent. Indeed, since U has less than $n_k(d) = (k+1)(d-k)$ elements and each v_i belongs to at least $k(d-k)$ of them, there must be at least k elements of U which contain all of $v_0, v_1, \ldots, v_k$. Clearly, the intersection of any two if these k elements contains the v_i and has affine dimension at most $k-1$, so $v_0, v_1, \ldots, v_k$ must be affinely dependent. It follows that the dimension of the affine span of V is at most $k-1$. As a consequence, each one of the k-dimensional faces F and G has at least one vertex not in V and, by Theorem 2.2, the graph $\mathcal{G}(P)\backslash V$ is connected. These two facts imply the existence of a walk w in $\mathcal{G}(P)$ with the claimed properties. □

6 Cell Complexes

In this section we discuss possible generalizations of Theorem 1.1 to certain classes of regular cell complexes. We will assume some familiarity with regular cell complexes and standard notions in topological combinatorics; excellent sources on these topics are [5, Sect. 4.7] and the article [3].

A *regular cell complex* is a finite collection $\mathcal{C}$ of subspaces of a Hausdorff space X, called *cells* or *faces*, such that each cell is homeomorphic to the closed unit ball in some finite-dimensional Euclidean space and the following conditions hold:

(i) $\varnothing \in \mathcal{C}$
(ii) The relative interiors of the nonempty cells partition X and
(iii) The boundary of any cell in $\mathcal{C}$ is a union of cells in $\mathcal{C}$.

Cells of dimension zero or one are called *vertices* or *edges*, respectively, and cells which are maximal with respect to inclusion are called *facets*. The dimension of $\mathcal{C}$ is the maximum dimension of a cell. The (loop-free) abstract graph $\mathcal{G}(\mathcal{C})$ defined by the vertices and edges of $\mathcal{C}$ is called the *graph* of $\mathcal{C}$.

A *polyhedral complex* in $\mathbb{R}^N$ is a regular cell complex each of whose cells is a polytope in $\mathbb{R}^N$. A *simplicial complex* is a polyhedral complex in which every cell

is a simplex. A regular cell complex $\mathcal{C}$ is said to have the *intersection property* if the intersection of any two cells in $\mathcal{C}$ is also a cell in $\mathcal{C}$ (in particular, the graph $\mathcal{G}(\mathcal{C})$ is simple). For instance, polyhedral complexes have the intersection property.

A regular cell complex $\mathcal{C}$ with the intersection property is said to be *Cohen–Macaulay* (over a field $\mathbb{K}$) [6, Sect. 2] if, under the inclusion partial order, it is a Cohen–Macaulay poset (over $\mathbb{K}$); see [3, Sect. 11] for the notion of Cohen–Macaulayness for simplicial complexes and posets. Such a complex $\mathcal{C}$ is pure, meaning that all facets of $\mathcal{C}$ have the same dimension, and strongly connected, meaning that for any two facets τ and τ' of $\mathcal{C}$ there exists a sequence $\tau = \tau_0, \tau_1, \ldots, \tau_n = \tau'$ of facets of $\mathcal{C}$ such that τ_{i-1} and τ_i intersect on a common face of codimension one for all $1 \le i \le n$. The following generalization of Balinski's theorem is an easy consequence of the results of Barnette [2] and follows also from [6, Corollary 2.7].

Theorem 6.1 *For every d-dimensional Cohen–Macaulay regular cell complex with the intersection property, the graph $\mathcal{G}(\mathcal{C})$ is d-connected.*

Let $\mathcal{G}_k(\mathcal{C})$ denote the simple graph on the node set of k-dimensional cells of $\mathcal{C}$ in which two such cells are adjacent if there exists a $(k+1)$-dimensional cell of $\mathcal{C}$ which contains them both. Theorem 6.1 is the special case $k = 0$ of the following statement.

Conjecture 6.2 *For every d-dimensional Cohen–Macaulay regular cell complex $\mathcal{C}$ with the intersection property, the graph $\mathcal{G}_k(\mathcal{C})$ is*

- *$(k+1)(d-k)$-connected if $0 \le k \le d-3$ and*
- *d-connected if $k = d-2$.*

The d-dimensional simplicial complex which has two d-dimensional simplices as facets, intersecting on a common codimension one face, shows that Theorem 1.1 does not extend to the setup of Conjecture 6.2 for $k = d-1$. We can verify this conjecture in the special cases which appear in the following statement.

Proposition 6.3 *Under the assumptions of Conjecture 6.2, the graph $\mathcal{G}_k(\mathcal{C})$ is*

(i) *$(k+1)(d-k)$-connected if $0 \le k \le d-3$ and $\mathcal{C}$ is a polyhedral complex,*
(ii) *d-connected if $k = d-2$.*

Proof Suppose first that $\mathcal{C}$ is a polyhedral complex and that $0 \le k \le d-3$ (a similar argument works for $k = d-2$). The result in this case follows easily from Theorem 1.1, as we now explain.

Let U be any subset of the set of k-dimensional faces of $\mathcal{C}$ of cardinality less than $(k+1)(d-k)$ and let F and G be two k-dimensional faces of $\mathcal{C}$ not in U. We need to show that these two faces can be connected by a walk in the graph $\mathcal{G}_k(\mathcal{C})\backslash U$. Let P and Q be facets of $\mathcal{C}$ (necessarily of dimension d) containing F and G, respectively. By strong connectivity of $\mathcal{C}$, we may choose a sequence $P = P_0, P_1, \ldots, P_n = Q$ of facets of $\mathcal{C}$ such that P_{i-1} and P_i intersect on a common $(d-1)$-dimensional face for all $1 \le i \le n$. By Lemma 2.1, each intersection $P_{i-1} \cap P_i$ has at least $\binom{d}{k+1}$ faces of dimension k. Since $\binom{d}{k+1} \ge (k+1)(d-k)$ for $k \le d-3$, there exists at least

one k-dimensional face, say F_i, of $P_{i-1} \cap P_i$ which is not an element of U. We let $F_0 = F$ and $F_{n+1} = G$ and note that, by Theorem 1.1, F_i and F_{i-1} can be connected with a walk in $\mathcal{G}_k(P_{i-1})\backslash U$ for every $1 \le i \le n+1$. It follows that F and G can be connected with a walk in $\mathcal{G}_k(\mathcal{C})\backslash U$. This proves (i).

Suppose now that $k = d-2$ and that $\mathcal{C}$ is as in Conjecture 6.2. We can proceed as in the proof of Proposition 3.1 with no need to pass to an object dual to $\mathcal{C}$. Indeed, let $\mathcal{E}$ be any subset of the set of $(d-2)$-dimensional faces of $\mathcal{C}$ of cardinality at most $d-1$. We observe first that every $e \in \mathcal{E}$ has at least $d-1$ codimension one faces and that no $(d-2)$-dimensional face of $\mathcal{C}$ other than e contains two or more of those. Therefore, it is possible to choose a codimension one face σ of e so that e is the unique element of $\mathcal{E}$ which contains σ.

We then check that in any part of a walk in $\mathcal{G}_{d-2}(\mathcal{C})$ of the form (τ, e, τ'), where $e \in \mathcal{E}$ and τ, τ' are faces of dimension $d-1$ containing e, the node e can be replaced by a walk that does not involve elements of $\mathcal{E}$ as follows. Since the inclusion poset of faces of $\mathcal{C}$ which strictly contain e is Cohen–Macaulay of rank one, and hence connected, we may assume that some facet ρ of $\mathcal{C}$ contains both τ and τ'. We pick a codimension one face σ of e as above and observe that the set of faces of $\mathcal{C}$ containing σ and contained in ρ is in inclusion-preserving bijection with the set of faces of a polygon Π (this holds more generally for Gorenstein* posets of rank 3) and that e and τ, τ' correspond to a vertex and its two incident edges in Π. Hence we can deviate our given walk in $\mathcal{G}_{d-2}(\mathcal{C})$ around e through the boundary of Π, thus avoiding nodes in $\mathcal{E}$ by the defining property of σ.

To complete the proof of (ii) it remains to comment that the graph $\mathcal{G}_{d-2}(\mathcal{C})$ is connected. This holds because the inclusion order on the set of faces of $\mathcal{C}$ of dimension $d-2$ or $d-1$ inherits the Cohen–Macaulay property from $\mathcal{C}$. □

Note Added in Revision The proof of Theorem 1.1 given in this paper has been extended by Ronald Wotzlaw [11] to the graph manifold setting of [2]. Combined with the argument given in the proof of part (i) of Proposition 6.3, this shows that Conjecture 6.2 holds even if the Cohen–Macaulay condition is relaxed to the requirement that the complex $\mathcal{C}$ is pure and strongly connected.

Acknowledgements The author thanks Bernd Sturmfels for encouraging discussions, Vic Reiner for providing an example of a regular cell complex with a nonregular face figure and Ronald Wotzlaw for the content of Remark 1.2 and for useful pointers to the literature. He also thanks the anonymous referees for helpful suggestions.

References

1. Balinski, M.L.: On the graph structure of convex polyhedra in n-space. Pac. J. Math. **11**, 431–434 (1961)
2. Barnette, D.: Graph theorems for manifolds. Israel J. Math. **16**, 62–72 (1973)
3. Björner, A.: Topological methods. In: Graham, R.L., Grötschel, M., Lovász, L. (eds.) Handbook of Combinatorics, pp. 1819–1872. North Holland, Amsterdam (1995)
4. Björner, A.: Mixed connectivity properties of polytopes, cell complexes and posets, in preparation
5. Björner, A., Las Vergnas, M., Sturmfels, B., White, N., Ziegler, G.M.: Oriented Matroids. Encyclopedia of Mathematics and Its Applications, vol. 46, Cambridge University Press, Cambridge (1993). 2nd printing (1999)
6. Fløystad, G.: Cohen-Macaulay cell complexes. In: Athanasiadis, C.A., Batyrev, V.V., Dais, D.I., Henk, M., Santos, F. (eds.) Algebraic and Geometric Combinatorics, Proceedings of a Euroconference in Mathematics, Anogia, Crete, Greece, August 2005. Contemporary Mathematics, vol. 423, pp. 205–220. American Mathematical Society, Providence (2007)

7. Grünbaum, B.: Convex Polytopes. Interscience, London (1967). Graduate Texts in Mathematics, 2nd edn., vol. 221. Springer, New York (2003)
8. Kalai, G.: Polytope skeletons and paths. In: Goodman, J.E., O'Rourke, J. (eds.) Handbook of Discrete and Computational Geometry, pp. 331–344. CRC Press, Boca Raton (1997)
9. Kalai, G.: Some aspects of the combinatorial theory of convex polytopes. In: Polytopes: Abstract, Convex and Computational, Scarborough, ON, 1993, pp. 205–229. Kluwer Academic, Dordrecht (1994)
10. Sallee, G.T.: Incidence graphs of convex polytopes. J. Comb. Theory **2**, 466–506 (1967)
11. Wotzlaw, R.: Incidence graphs and unneighborly polytopes. Doctoral Dissertation, Technical University of Berlin (2009)
12. Ziegler, G.M.: Lectures on Polytopes. Graduate Texts in Mathematics, vol. 152. Springer, New York (1995)