Discrete Comput Geom 24:421–435 (2000)
DOI: 10.1007/s004540010046

Discrete & Computational
Geometry

Convex and Linear Orientations of Polytopal Graphs*

J. Mihalisin and V. Klee

Department of Mathematics, University of Washington,
Box 354350, Seattle, WA 98195-4350, USA
{mihalisi,klee}@math.washington.edu

Abstract. This paper examines directed graphs related to convex polytopes. For each fixed d-polytope and any acyclic orientation of its graph, we prove there exist both convex and concave functions that induce the given orientation. For each combinatorial class of 3-polytopes, we provide a good characterization of the orientations that are induced by an affine function acting on some member of the class.

Introduction

A **graph** is *d-polytopal* if it is isomorphic to the graph $G(P)$ formed by the vertices and edges of some (convex) d-polytope P. As the term is used here, a **digraph** is *d-polytopal* if it is isomorphic to a digraph that results when the graph $G(P)$ of some d-polytope P is oriented by means of some affine function on P.

3-Polytopes and their graphs have been objects of research since the time of Euler. The most important result concerning 3-polytopal graphs is the theorem of Steinitz [SR], [Gr1], asserting that a graph is 3-polytopal if and only if it is planar and 3-connected. Also important is the related fact that the combinatorial type (i.e., the entire face-lattice) of a 3-polytope P is determined by the graph $G(P)$. Steinitz's theorem has been very useful in studying the combinatorial structure of 3-polytopes because it makes it easy to recognize the 3-polytopality of a graph and to construct graphs that represent 3-polytopes without producing an explicit geometric realization.

During the 1900s a theory of higher-dimensional polytopes and their graphs gradually emerged. The most familiar property of d-polytopal graphs is that they are d-connected [Ba], and the most striking property is that for a d-polytope P whose graph is d-regular, the entire face-lattice is determined by the graph $G(P)$ [BML], [Ka1]. Grünbaum wrote two survey articles dealing with polytopal graphs [Gr2], [Gr3], and these graphs also

* The authors are indebted to the National Science Foundation for partial support of this research.

played an important role in his book on convex polytopes [Gr1] and in Ziegler's book on the same subject [Zi].

In contrast to the broad study of polytopal graphs, attention to polytopal digraphs has been much narrower. Except for some very recent papers [Ed], [ADRS], [AER], research on polytopal digraphs has been largely confined to establishing bounds on the lengths of monotone paths for various sorts of polytopes. (See [Kl], [KM], [To], [Ka2], [Ka3], [Zi], [Gä], [GZ], [AZ], and [HK], and see [AZ] for additional references.)

We hope that the present article will open the door to a broader study of polytopal digraphs. Its main accomplishment is adapting a proof of Steinitz's theorem due to Barnette and Grünbaum [BG], and thereby obtaining a characterization of 3-polytopal digraphs that is easily testable and thus facilitates the study, without constructing an explicit geometric realization, of the digraphs that can arise when the edges of some 3-polytope are oriented by means of some affine function. We also include a short discussion of the digraphs that result when the graph $G(P)$ of a d-polytope P is oriented by means of a function that is convex or concave.

1. Statement of Main Result

An *orientation* of a graph G is a way of assigning a direction to each edge of G, thus turning G into a digraph. When f is a real-valued function whose domain includes the vertex set of a graph G, we say that f is *admissible* for G provided that f does not attain the same value at any two adjacent vertices of G. Each function that is admissible for G *induces* an orientation by replacing each undirected edge xy by the directed edge $\overrightarrow{xy}$ or $\overrightarrow{yx}$ according as $f(x) < f(y)$ or $f(y) < f(x)$. The resulting digraph is denoted by $D(G, f)$. If a function does not attain the same value at any two vertices (whether or not they are connected by an edge), it is called *strongly admissible*.

When C is a convex set, a function $\varphi\colon C \longrightarrow \mathbb{R}$ is *convex* if $\varphi(\alpha x + \beta y) \leq \alpha\varphi(x) + \beta\varphi(y)$ for each choice of $x, y \in C$ and $\alpha, \beta > 0$ with $\alpha + \beta = 1$. If this holds with $\leq$ replaced by $\geq$, the function φ is *concave*. The function φ is *affine* if it is both convex and concave, and in this case φ can be extended to a function φ^* on the entire ambient space $\mathbb{R}^d$ such that $\varphi^*(\alpha x + \beta y) = \alpha\varphi^*(x) + \beta\varphi^*(y)$ for all $x, y \in \mathbb{R}^d$ and $\alpha, \beta \in \mathbb{R}$ with $\alpha + \beta = 1$. Of course the extension φ^* is equal to a linear function plus a constant.

When P is a polytope in real d-space $\mathbb{R}^d$, an *LP orientation* of P's graph $G(P)$ is one that is induced by some affine function on P that corresponds to an admissible function for $G(P)$. (Since adding a constant to a function will not change the orientation it induces, any orientation induced by an affine function is also induced by a linear function. "LP" is intended to suggest "linear programming.") Thus a digraph is d-polytopal if and only if it is isomorphic to an LP orientation of the graph of some d-polytope.

A *path* in a graph is a sequence $(v_0, \ldots, v_k)$ of distinct vertices such that, for $1 \leq i \leq k$, the unordered pair $v_{i-1}v_i$ is an edge of the graph. With respect to a given orientation of the graph, the path is *monotone* provided that each of the ordered pairs $\overrightarrow{v_{i-1}v_i}$ is a directed edge of the digraph, and a path is *antitone* provided that each of the ordered pairs $\overrightarrow{v_i v_{i-1}}$ is a directed edge of the digraph. The *length* of a path is the number of edges it uses. A set of paths from u to v is called *independent* if no vertex other than u or v appears in more than one member of the set.

For an acyclic digraph, we define a *detour* to be a monotone path of length greater than one whose endpoints are also connected by a single edge. If M is a monotone path, then *adding a detour* to M means replacing a single edge of M with a longer monotone path having the same endpoints as the edge.

A vertex v of a digraph K is a *source* if all edges incident to v are directed away from v, and a *sink* if all edges incident to v are directed toward v. Now our main result can be stated as follows:

Theorem 1.1. *Suppose that K is an orientation of a 3-polytopal graph G. Then the digraph K is 3-polytopal if and only if it is acyclic, has a unique source and a unique sink, and admits three independent monotone paths from source to sink.*

Note that, since the graph determines the combinatorial type for 3-polytopes, Theorem 1.1 specifies all LP orientations possible within a combinatorial type. It is easy to find examples consisting of two combinatorially equivalent polytopes P and P' and an LP orientation of $G(P)$ such that the corresponding orientation of $G(P')$ is not LP. (In two dimensions, for example, any LP orientation of a rectangle must have source and sink diagonally opposite, but a more general quadrilateral may possess an LP orientation with source and sink adjacent.)

The "only if" in Theorem 1.1 follows from a previous result [HK] which we restate in Section 3. The proof of "if," given in Section 3, closely follows the "edge deletion" proof of Steinitz's theorem given by Barnette and Grünbaum [BG]. Other proofs of Steinitz's theorem appear in [SR], [Gr1], [BG], [Zi], and [RG], but since they utilize the dual of the given polytope they do not seem to be adaptable to deal with polytopal digraphs.

The proof of Theorem 1.1 appears in Section 3, but we first devote a short Section 2 to orientations induced by convex functions. The following proposition plays a role in both Sections 2 and 3.

Proposition 1.2. *If G is a finite graph, then each orientation of G that is induced by an admissible function is acyclic, and each acyclic orientation of G is induced by some strongly admissible function.*

Proof. The first assertion is obvious. For the second, suppose that an acyclic orientation of the graph $G = (V, E)$ is given. Call the resulting digraph K. Let n be the number of vertices in the graph and let T denote the subset of V that consists of all sinks in K. Acyclicity implies that T is not empty. Choose any $v \in T$ and set $l(v) = n$. Now consider the acyclic digraph $K\backslash\{v\}$. Take any sink w in this smaller digraph and set $l(w) = n - 1$. Iterating this process, we eventually define the function l on all vertices of the finite graph G. As defined, l is an injection hence it is strongly admissible. For any directed edge $\overrightarrow{vw}$, the vertex v will not be a sink until w is removed, hence we must have $l(v) < l(w)$, so $K = D(G, l)$. □

The function l is called a *level function* of the digraph K.

2. Orientations Induced by Convex or Concave Functions

It is well known [HH], [GK] that a convex function f on a polytope P attains its maximum at some vertex of P.

Further, if f is strictly convex, then each of its local maxima is attained at a vertex. The following remark shows that knowledge along P's edges is sufficient to recognize which vertices of P constitute local maxima for a given convex function. This reduces the recognition problem to a finite number of one-dimensional tests.

Proposition 2.1. *Suppose that P is a polytope, f is a convex function on P, and v is a vertex of P. Then f attains a local maximum at v if and only if, for each edge E of P incident to v, the restriction of f to E attains a local maximum at v.*

Proof. "Only if" is obvious. For "if," let $x_1, \ldots, x_k$ be the other ends of the edges of P that are incident to v and suppose that, for each i, the restriction of f to the edge vx_i has a local maximum at v. Then each edge vx_i contains a segment vw_i such that f's restriction to vw_i has a maximum at v. The rays that issue from v and pass through the various points x_i are the extreme rays of a pointed convex cone C that has apex v and contains P. From this it follows that the convex hull Q of the set $\{v, w_1, \ldots, w_k\}$ is a neighborhood of v relative to P. Each point $q \in Q$ can be expressed in the form $q = \sum_{i=1}^{k} \lambda_i q_i$ with all $\lambda_i \geq 0$, $\sum_{i=1}^{k} \lambda_i = 1$, and $q_i \in vw_i$. Since f is convex, it then follows that

$$f(q) \leq \sum_{i=1}^{k} \lambda_i f(w_i) \leq \sum_{i=1}^{k} \lambda_i f(v) = f(v). \qquad \square$$

If the objective function f is affine rather than merely convex, then either or both occurrences of "local" can be replaced by "global" in the statement of Proposition 2.1. This is an essential aspect of edge-following algorithms for linear programming.

The main purpose of this paper is to provide, in the next section, a good characterization of the edge-orientations that can be induced on 3-polytopes (within a given combinatorial class) by means of affine functions. First we prove Theorem 2.2, which shows that, for an arbitrary d and an arbitrary *fixed* d-polytope P, each acyclic orientation of $G(P)$ is induced by some convex and also by some concave function on P.

Now suppose that, for an admissible convex objective function f on a polytope P, the resulting orientation $K = D(P, f)$ is known. Then it follows from Proposition 2.1 that each sink in K corresponds to a local maximum for f. Clearly all the global maxima must be among the sinks. The following result shows that nothing more can be said about local and global maxima if we know only $D(P, f)$ and the fact that f is convex.

Theorem 2.2. *Suppose that P is a d-polytope in $\mathbb{R}^d$, V is P's vertex-set, and K is an acyclic orientation of P's graph. Let T denote the set of all K-sinks in V, and let W and M be any sets satisfying $M \subset T \subset W \subset V$. Then $\mathbb{R}^d$ admits two convex functions f and $\overline{f}$ and a concave function g such that the following conditions are satisfied*:

the functions f, g, and $\overline{f}$ are all piecewise affine and admissible for $G(P)$;

each function induces the orientation K of $G(P)$;
for each $p \in P$, $f(p) \leq g(p)$;
over each face of P that is a simplex (in particular, along each edge), f and g agree and are both affine;
for the restriction of $\overline{f}$ to P, global maxima are attained precisely at the members of M and local maxima are attained precisely at the members of W.

Proof. By Proposition 1.2, there is an admissible function l on V that induces the orientation K of $G(P)$. Let

$$C = \text{conv}\{(v, l(v)): v \in V\} \subset P \times \mathbb{R} \subset \mathbb{R}^{d+1}.$$

Then C is a convex polytope in $\mathbb{R}^{d+1}$. Unless P is a simplex (in which case the result is trivial), we may assure that C is $(d+1)$-dimensional by perturbing the values of l while preserving the orientation K induced by l. For each $p \in P$, let $f(p)$ and $g(p)$ denote respectively the smallest and the largest real number η such that $(p, \eta) \in C$. It is evident that, for each $v \in V$, C intersects the line $\{v\} \times \mathbb{R}$ in the single point $(v, l(v))$. Hence $f(v) = l(v) = g(v)$.

The graph of f is the lower boundary of C, and $f(p)$ is the height at p of this lower boundary above the hyperplane $\mathbb{R}^d \times \{0\}$ in $\mathbb{R}^{d+1}$. Similarly, the graph of g is the upper boundary of C. Then f and g are respectively a convex and a concave function on P, and each is an extension of l so each induces the orientation K of $G(P)$.

Now suppose that $p \in P$, that F is the smallest face of P containing p, and that F is a k-simplex with vertices $v_0, \ldots, v_k$. Then there is a unique expression of p as a convex combination of the v_i—say $p = \sum_{i=0}^{k} \lambda_i v_i$—and it is easily verified that the intersection of C with the line $\{p\} \times \mathbb{R}$ consists of the single point $(p, \sum_{i=0}^{k} \lambda_i l(v_i))$. Hence $f(p) = g(p)$ and it follows that, for each face G of P that is a simplex, f and g agree on G. Since f is convex while g is concave, it follows that both functions have affine restriction to G.

Now we want to extend each of f and g to all of $\mathbb{R}^d$ in such a way that the extensions are piecewise affine and are respectively convex and concave. The argument is essentially the same for the two cases, so we consider only f. The graph of f is the union of certain "lower" facets of C. Let these be $F_1, \ldots, F_m$. Each facet F_i lies in a hyperplane H_i in $\mathbb{R}^{d+1}$ such that the set $C \backslash F_i$ lies entirely above H_i. That is, there is a linear function $\Phi_i: \mathbb{R}^{d+1} \longrightarrow \mathbb{R}$ and a constant γ_i such that $F_i \subset \{(x, \alpha): \Phi_i(x, \alpha) = \gamma_i\}$ and $\Phi_i(y, \beta) > \gamma_i$ for all $(y, \beta) \in C \backslash F_i$. Now we can define a function $\varphi_i: \mathbb{R}^d \longrightarrow \mathbb{R}$ by the condition that, for each $x \in \mathbb{R}^d$, $(x, \varphi_i(x)) \in H_i$; equivalently, $\Phi_i(x, \varphi_i(x)) = \gamma_i$. A routine exercise shows that each function φ_i is affine. Finally, for each $x \in \mathbb{R}^d$ set $f^*(x) = \max\{\varphi_i(x): 1 \leq i \leq m\}$. This is an extension of f (hence of φ) to all of $\mathbb{R}^d$ and, as the maximum of a collection of convex functions, it is convex. It is clearly piecewise affine.

It remains only to produce the convex function $\overline{f}$. We begin by constructing a function f' in a fashion analogous to the construction of f, except that, for all vertices $v \in M$, we choose some sufficiently large integer μ and use (v, μ) instead of $(v, l(v))$ in defining the polytope C. Since $M \subset T$, it is clear that f' still induces K. For large enough μ, each $v \in M$ is a global maximum. Note that, since f''s restriction to each edge is affine, it follows from Proposition 2.1 that the local maxima for f' are precisely the members

of T. Now for each vertex $w \in W$, let H_w be a hyperplane that strictly separates w from the set of all midpoints of edges of P incident to w, and let the affine function $\psi_w\colon \mathbb{R}^d \longrightarrow \mathbb{R}$ be defined by the condition that $\psi_w(w) = 1$ and $\psi_w(y) = 0$ for all $y \in H_w$. Take $\psi_w^+ = \max\{\psi_w, 0\}$. Finally, set

$$\overline{f} = f' - \tfrac{1}{2} + \tfrac{1}{2} \sum_{w \in W} \psi_w^+.$$

Then $\overline{f}$ is convex. It induces the orientation K since f' is integer valued on V, f' and $\overline{f}$ agree on the vertices in W, and $\overline{f} = f' - \frac{1}{2}$ on all other vertices of P (hence their relative order is unchanged). If each H_w is chosen sufficiently close to w, then each vertex $w \in W$ provides a strict local maximum for the restriction of $\overline{f}$ to P. However, at each vertex $v \in V \backslash W$, the local behavior of $\overline{f}$ agrees with that of f' and hence (since $v \notin T$) v does not produce a local maximum for $\overline{f}$. □

3. Linear Orientations and 3-Polytope Combinatorial Classes

If a d-polytope P is presented as the convex hull of a finite subset of $\mathbb{R}^d$, the vertices and edges of P can be produced by routine applications of linear programming. When an orientation K of P's graph is given, linear programming can be used to decide whether the orientation is LP. For this decision, let W denote the set of all points of the form $y - x$ where $\overrightarrow{xy}$ is a directed edge of P. Then the orientation is LP if and only if the origin does not belong to the convex hull of W—and this can be tested by means of linear programming. It follows that when P is presented as the convex hull of finitely many points with integer coordinates, it can be decided in polynomial time whether a given orientation of P's graph is LP.

Now suppose that we are merely given an abstract graph or digraph and a positive integer d, and we wonder whether the given object is d-polytopal. By reasoning similar to that used by Grünbaum for a different but related purpose [Gr1, pp. 91–92], there is a (very slow) recognition algorithm based on Tarski's decision method for real-closed fields [Ta], [Re]. On the other hand, when $d = 3$ the theorem of Steinitz provides a characterization that is testable in polynomial time for the undirected case. We obtain such a characterization for the directed case by using (for $d = 3$) the following directed analogue of d-connectedness established in [HK].

Theorem 3.1. *Suppose that f is an admissible affine function for a d-polytope P, and x and y are vertices of P such that $f(x) = \min f(P)$ and $f(y) = \max f(P)$. Then in the orientation of P's graph induced by f there are d independent monotone paths from x to y.*

(Since f is affine each of the sets $\{x \in P\colon f(x) = \min f(P)\}$ and $\{x \in P\colon f(x) = \max f(P)\}$ is actually a face of P. Admissibility implies that each of these faces must simply be a vertex. Hence the hypotheses imply that f attains a unique minimum and a unique maximum on P.)

We call a digraph 3-*monotone* if its underlying graph is 3-polytopal and the digraph is acyclic with a unique source, a unique sink, and three independent monotone paths from

source to sink. It follows from Theorem 3.1 that each 3-polytopal digraph is 3-monotone. The main result of this paper is that the converse is also true. That is, we want to prove the "if" part of Theorem 1.1.

It is worth noting that Theorem 3.1 actually implies a stronger property. Since f determines an admissible affine function in every face-defining hyperplane, the theorem implies that every j-face has j independent monotone paths from the face's (relative) source to its (relative) sink. Our definition of 3-monotone makes no explicit restrictions on the 2-faces of the polytope.

The proof of Theorem 1.1 works by induction on the number of edges and consists of two parts. The first part shows that each 3-monotone digraph has a "deletable" edge. The second part shows that if the digraph that results from deleting an edge is 3-polytopal, then the original digraph is also 3-polytopal.

To help clarify some details in what follows, we introduce some definitions and notation. A digraph J is *contained in* a digraph K if there exist two injections

$$\psi\colon \operatorname{vert}(J) \to \operatorname{vert}(K) \quad \text{and}$$
$$\phi\colon \operatorname{edge}(J) \to \text{ the set of all } \textbf{monotone} \text{ paths in } K$$

satisfying the following conditions:

for each edge $\overrightarrow{ab}$ of J, $\phi(\overrightarrow{ab})$ is a monotone path in K from $\psi(a)$ to $\psi(b)$;
the interior of each path in image(ϕ) is disjoint from all other paths in image(ϕ).

A particular choice of ψ and ϕ is called an *embedding* of J in K. With each embedding (ψ, ϕ), we associate two subsets of vert(K). The first is image(ψ). For any ψ, this set has cardinality equal to that of vert(J); its members are called the *active vertices* of the embedding. The second set consists of all vertices in the interiors of those paths in image(ϕ); its members are called the *dormant vertices* of the embedding. Even with ψ fixed, the cardinality of this set can vary with different ϕ.

If an edge contained in an element of image(ϕ) possesses a detour that intersects the embedding only at the detour's endpoints, then we may add the detour to increase the number of dormant vertices. Starting with any embedding, we may successively modify it by adding available detours (ones whose interiors do not intersect the current embedding) to arrive at an embedding that does not admit any detours.

Note that up to symmetry, the diagram on the left in Fig. 1 represents the unique acyclic orientation of the complete graph on four vertices. Let D_4 denote this orientation.

Lemma 3.2. *Each* 3*-monotone digraph contains* D_4*. Further, the embedding may be chosen so that the source and sink of* D_4 *are sent to the source and sink of the digraph.*

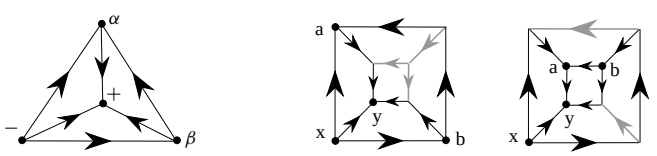

Fig. 1. Here are two different embeddings of D_4 in a digraph on the cube. $x = \psi(-)$, $y = \psi(+)$, $a = \psi(\alpha)$, and $b = \psi(\beta)$. The injection ϕ may be inferred from the dark edges. The left cube shows an embedding with three dormant vertices. The right cube shows an embedding with four dormant vertices.

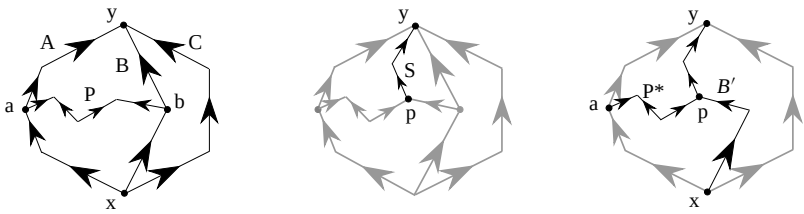

Fig. 2. Labels refer to vertices or darkened paths. From the nonmonotone connecting path P we find either a monotone connecting path (not shown) or a new quadruple (A, B', C, P^*) where P^* has smaller length than P.

Proof. Call the source of the 3-monotone digraph x and the sink y. We must find four monotone paths A, B, C, M that satisfy the following conditions: each of A, B, and C goes from x to y; M goes from an interior vertex of A to an interior vertex of B; and except for the intersections implied by the preceding two conditions, the paths A, B, C, and M are pairwise vertex-disjoint.

Let $\mathcal{Q}$ denote the set of all quadruples (A, B, C, P) that satisfy all the above conditions except that P is just a path in the graph (not necessarily monotone). We show first that $\mathcal{Q}$ is nonempty.

Since K is 3-monotone with source x and sink y, there exist three independent monotone paths from x to y. Start with any such triple and note that, since our definitions do not permit parallel edges, at least two of the three paths must have interior vertices. Let u and v be interior vertices of different paths. Since G's graph is 3-connected, G contains a path P from u to v that misses x and y. We may take a subpath of P, and combine it with the original triple (appropriately labeled), to satisfy the above conditions and produce an element of $\mathcal{Q}$.

Now let (A, B, C, P) be a quadruple in $\mathcal{Q}$ for which the length of P is a minimum (Fig. 2). Assume that P connects $b \in B$ to $a \in A$. Starting at b, let p be the first vertex of P where the "direction changes" (i.e., the subpath of P from b to p is the longest monotone or antitone subpath that includes b).

If this first segment is monotone, let S be a monotone path from p to y. Since K is acyclic with unique sink y, such a monotone path must exist. From planarity, S must encounter A or B before it encounters C. We claim that S encounters A before reaching y. The subpath of P from b to p followed by the initial segment of S up to A can then serve as the monotone path M connecting B and A to satisfy the conditions of the lemma.

To justify the claim, note that if S goes all the way to y before it encounters A, then the minimality of P in the quadruple $(A, B, C, P) \in \mathcal{Q}$ is contradicted by the quadruple (A, B', C, P^*) where B' is formed by following B from x to b then P to p followed by S to y, and where P^* is the subpath of P connecting p to a.

If instead, the first segment of P is antitone, the proof is similar with the role of S being played by an antitone path from p to x. □

When a digraph K is 3-monotone we call an edge E of K *deletable* if there exists a 3-monotone digraph J contained in K with an embedding (ψ, ϕ) such that image(ϕ)

includes all edges of K except E. As described in [BG], there are three possibilities for the change in cardinality from vert(J) to vert(K). If the deletable edge connects two vertices that are active in (ψ, ϕ), then the two cardinalities are the same. Otherwise, the new edge "activates" one or two dormant vertices and thus increases the cardinality of vert(K) by one or two. Correspondingly, K has one, two, or three more edges than J since each time a dormant vertex is activated it splits a J edge into two K edges.

Lemma 3.3. *For each* 3*-monotone digraph* K *there exists a sequence of* 3*-monotone digraphs* $J_0, \ldots, J_k$ *such that* J_0 *is* D_4, $J_k = K$, *and each* J_{n-1} *is obtained from* J_n *by deleting one edge.*

Proof. Lemma 3.2 proves the existence of at least one embedding (ψ, ϕ) of D_4 in K, where ψ takes the source and sink of D_4 to x (the source) and y (the sink) of K. By adding detours, if necessary, we arrive at an embedding (ψ_0, ϕ_0) which has no detours available.

Now assume we have a sequence $J_0, \ldots, J_n$ of digraphs, where each J_{m-1} is obtained from J_m by deleting one edge, and where J_n is contained in K with an embedding (ψ_n, ϕ_n) which has no detours available. If $J_n \neq K$, we will show that it is always possible to "add an edge" to J_n to create a J_{n+1}. (That is, we will find a monotone K-path which may be added to the embedding of J_n to create an embedding of a larger digraph, J_{n+1}, which differs from J_n by a single deletable edge.) By assumption, J_n is a 3-monotone digraph, but care must be taken to ensure that J_{n+1} is a 3-monotone digraph as well.

Since K is planar and acyclic, any digraph contained in it is planar and acyclic. The process of "adding an edge" cannot create new sources or sinks. Further, since J_0 is D_4 and (by composing embeddings) each J_n contains J_0, each J_{n+1} will still have three independent paths from the unique source to the unique sink. Finally, it is routine to verify that "adding an edge" cannot destroy the 3-connectedness of a graph.

Our one concern will be to ensure that the new object is actually a digraph. If the proposed "new edge" (i.e., a monotone path in K whose endpoints belong to nth embedding but whose interior is disjoint from it) is not properly chosen, J_{n+1} will have parallel edges. It is also routine to check that a proposed "new edge" will create parallel edges in J_{n+1} if and only if a single element of image(ϕ_n) contains both of the path's endpoints (Fig. 3). (Note that dormant vertices are contained in a unique element of image(ϕ_n), active vertices are contained in at least three.)

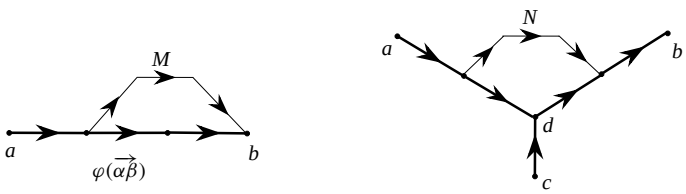

Fig. 3. Thicker lines indicate edges contained in the embedding (ψ, ϕ). All active vertices are labeled, with $a = \psi(\alpha)$ and $b = \psi(\beta)$. Adding the monotone path M creates two parallel edges, but adding N does not create any parallel edges since no element of image(ϕ) contains both of the path's endpoints.

Given the sequence $J_0, J_1, \ldots, J_n$ as described above, if $J_n \neq K$, then either:

(a) the nth embedding has no dormant vertices and one of J_n's vertices is "missing a K edge" (i.e., there is an α in J_n such that $a = \psi_n(\alpha)$ has greater valence than α), or
(b) the nth embedding includes a dormant vertex.

Note: A greek letter always signifies a vertex of J_n, its corresponding vertex in K is given the roman counterpart.

Case (a): *When There Are No Dormant Vertices*. If the nth embedding has no dormant vertices, then ϕ_n must take edges of J_n to single edges of K. Let α be a vertex of J_n that is missing a K edge and let $a = \psi_n(\alpha)$. Let b be the other endpoint of the missing K edge, i.e., $\overrightarrow{ab}$ or $\overrightarrow{ba}$ is an edge of K that is not contained in image(ϕ_n).

If b happens to be in image(ψ_n), we may simply take J_{n+1} to be the digraph that results from adding a directed edge between α and $\beta = \psi_n^{-1}(b)$. Since these are active vertices, this would create a parallel edge only if an edge between α and β already exists in J_n. By hypothesis, this is not the case. Thus adding $\overrightarrow{\alpha\beta}$ or $\overrightarrow{\beta\alpha}$ creates a 3-monotone J_{n+1}.

To produce the $(n+1)$st embedding, take $\psi_{n+1} = \psi_n$ and begin with $\phi_{n+1}(\overrightarrow{\alpha\beta}) = \overrightarrow{ab}$ (or $\phi_{n+1}(\overrightarrow{\beta\alpha}) = \overrightarrow{ba}$, respectively) then add detours to arrive at an $(n+1)$st embedding with no detours available.

If b is not in image(ψ_n) and the K edge is directed from a to b, then consider any monotone K path from b to the sink y. (The case where the K edge points from b to a is similar, with "monotone" replaced by "antitone" and "sink" replaced by "source.") Let c be the first vertex of this path that is contained in the nth embedding. Call M the monotone K path from a through b to c. Let $\alpha = \psi_n^{-1}(a)$ and $\gamma = \psi_n^{-1}(c)$. We now form J_{n+1} by adding a directed edge to J_n from α to γ. Note that $\overrightarrow{\alpha\gamma}$ does not already exist in J_n since ϕ_n takes edges to edges, the nth embedding has no detours available, and M is a detour to $\overrightarrow{ac}$ which only intersects the nth embedding at its endpoints.

To determine the $(n+1)$st embedding, again take $\psi_{n+1} = \psi_n$ and begin with $\phi_{n+1}(\overrightarrow{\alpha\gamma}) = M$ then add detours to arrive at an $(n+1)$st embedding with no detours available.

Case (b): *Activating a Dormant Vertex*. Let u be a dormant vertex in the nth embedding. Let U be the element of image(ϕ_n) that contains u. (U is a monotone K path.) Choose any other vertex v contained in the nth embedding but not contained in U. As in Lemma 3.2, the 3-connectedness of K implies there exists a (not necessarily monotone) path P from u to v in K that avoids U's endpoints. We may relabel v and take a subpath of P to ensure that no vertex in the interior of P is contained in the nth embedding. (We may need to replace u with a different dormant vertex in U.) Set $p_0 = u$. Within P, label vertices where the "direction changes" as $p_1, p_2, \ldots, p_l$. (That is, the portion of P from p_{i-1} to p_i is monotone; the portion from p_i to p_{i+1} is antitone; etc.)

For each p_i that is a sink relative to P, consider any monotone path from p_i up to the sink y. (Acyclicity with a unique sink implies there is at least one such path.) Let s_i be the first vertex of this path that is contained in the nth embedding. Call S_i the monotone subpath leading from p_i to s_i (Fig. 4). Similarly, for each p_i that is a source relative to P, consider any antitone path from p_i down to the source x. Let s_i be the first vertex in

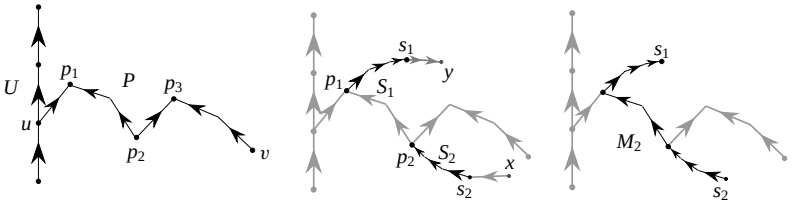

Fig. 4. Labels refer to vertices or darkened paths. Starting with a nonmonotone path P between u and v, we define two families of points and two families of monotone paths. S_i connects p_i to s_i and M_i connects s_i to s_{i-1}.

this antitone path that is contained in the nth embedding. Call S_i the antitone subpath leading from p_i to s_i. Define s_0 to be u and S_0 to be the trivial path from u to u. Define both s_{l+1} and p_{l+1} to be v and S_{l+1} to be trivial. Finally, call M_i the monotone or antitone path that begins at s_i, follows S_i to p_i, follows the portion of P from p_i to p_{i-1}, then follows S_{i-1} to end at s_{i-1}.

Let s_{j+1} be the first element of $\{s_0, \ldots, s_{l+1}\}$ not contained in U. (Recall that $v = s_{l+1}$ is not in U.) Unless s_j and s_{j+1} belong to a single element of ϕ_n, the path M_{j+1} may be used to create a "new edge" for J_{n+1}, as we shall describe. If s_j is active in the nth embedding, let α be the corresponding endpoint of $\phi_n^{-1}(U)$. ($\phi_n^{-1}(U)$ is an edge of J_n.) If s_j is dormant in the nth embedding, let α be a new vertex created in the interior of $\phi_n^{-1}(U)$. Let W be an element of image(ϕ_n) which contains s_{j+1}. If s_{j+1} is active, then let β be the appropriate endpoint of $\phi_n^{-1}(W)$. If s_{j+1} is dormant, let β be a new vertex created in the interior of $\phi_n^{-1}(W)$. We form J_{n+1} by adding to J_n a directed edge between α and β (with orientation determined by the orientation of M_{j+1}).

To find the $(n+1)$st embedding, we begin by using the same assignments as the nth embedding, by setting $\psi_{n+1}(\alpha) = s_j$ and $\psi_{n+1}(\beta) = s_{j+1}$ and by sending the edge between α and β to M_{j+1}. (Note that "using the same assignments" implies that if E is an edge in J_n which gets "split" into two J_{n+1} edges, then ϕ_{n+1} must send the two edges thus created to subpaths of $\phi_n(E)$.) We then successively add detours to ensure that (ψ_{n+1}, ϕ_{n+1}) has no detours available.

If s_j and s_{j+1} do belong to a single element of image(ϕ_n), call this element W. (W is a monotone K path.) In this case s_j must be an active vertex. Let α be a new vertex created in the interior of $\phi_n^{-1}(U)$ and let β be a new vertex created in the interior of $\phi_n^{-1}(W)$. We form J_{n+1} by adding to J_n a directed edge between α and β. Since α and β are in the interiors of different edges of J_n, the new edge will not be parallel to an existing one. As described below, the $(n+1)$st embedding will send α to p_{j-1} and β to p_j. The subpath of P from p_{j-1} to p_j determines the orientation of the new edge between α and β.

We determine the $(n+1)$st embedding by first finding a new embedding for J_n. Call W' the monotone path that results from replacing the section of W from s_{j+1} to s_j by the path M_{j+1} (Fig. 5). If $j > 1$, call U' the monotone path that results from replacing the section of U from s_{j-1} to s_{j-2} by the path M_{j-1}. If $j = 1$, set $U' = U$. (Since s_j is an active vertex, $j > 0$.) The interiors of the M_i contain no vertices—neither active nor dormant—from the nth embedding. This implies that by replacing U with U' and W with W' we create a new embedding (ψ_n, ϕ_n') of J_n. To find the $(n+1)$st embedding, we begin by using the same assignments as (ψ_n, ϕ_n'), by sending α to p_{j-1} and β to p_j,

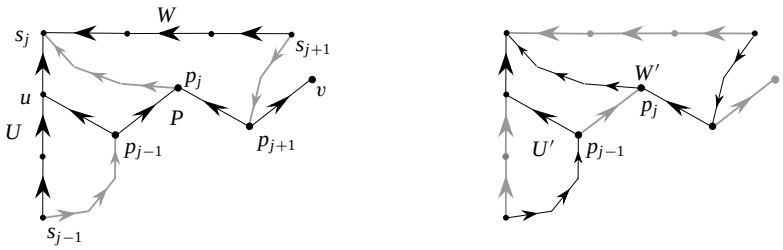

Fig. 5. Labels refer to vertices or darkened paths. We define U' by replacing the subpath of U from s_{j-1} to s_{j-2} with M_{j-1}. We define W' by replacing the subpath of W from s_{j+1} to s_j with M_{j+1}. The figure shows the case where $j = 2$.

and by sending the new edge (between α and β) to the subpath of P from p_{j-1} to p_j. We then successively add detours to ensure that (ψ_{n+1}, ϕ_{n+1}) has no detours available.

Thus we have shown that if J_n is contained in K but $J_n \neq K$, we may always "add an edge" to create a J_{n+1}. Since K is finite, we must eventually arrive at $J_k = K$. This completes the proof of Lemma 3.3. □

We omit the proof of the following lemma, since it concerns only the undirected graph and it appears in detail in [BG].

Lemma 3.4. *The totality of all vertices and faces of a* 3*-connected planar graph, G, may be arranged in a sequence in such a manner that each element is incident with at most three of the elements that precede it in the sequence. Moreover, such an arrangement exists even under the additional requirement that for any chosen edge E of G the two vertices of E and the two faces containing E be placed at the beginning of the sequence.*

The following result also concerns only undirected graphs. While it is true for directed graphs as well, it is not convenient to prove it separately from the main result. It is proved in detail in [BG], but we sketch the proof here to help clarify the directed case.

Lemma 3.5. *If a graph G is* 3*-polytopal and is obtained from a planar graph G′ by deleting an edge* (*in the sense of the definition preceding Lemma* 3.3), *then G′ is also* 3*-polytopal.*

Proof. Intuitively, if we have a polytope that realizes G, then we first imagine "drawing an edge" across the appropriate face to give us an object with the combinatorial structure of G'. However, this object is not a polytope since the two newly created faces are coplanar. We then "bend one face" slightly to produce a polytope. It is a nontrivial result that such a bending may always take place without disturbing the combinatorial structure. (In fact, such an operation may be impossible in higher dimensions.) The fact that this is always possible for 3-polytopes follows from Lemma 3.4.

Lemma 3.4 implies that after "bending" one face, we may iteratively adjust the remaining faces to maintain the same combinatorial structure. As we "propagate" the perturbation according to the sequence given by the lemma, each vertex is specified to

be in at most three of the preceding planes (which have been only slightly perturbed from previously intersecting in a point) and each plane is required to contain at most three of the preceding points, hence we cannot run into a geometric impossibility from our small perturbation. □

Armed with the preceding lemmas, we may now restate the theorem of Steinitz in combination with its directed analogue.

Theorem 3.6. *A graph is* 3*-polytopal if and only if it is planar and* 3*-connected. If K is an orientation of a* 3*-polytopal graph, then the digraph K is* 3*-polytopal if and only if K is acyclic, has a unique source and a unique sink, and admits three independent monotone paths from the source to the sink.*

Proof. The undirected case would follow simply from an induction using the previous lemmas. One minor complication arises in the digraph case. We must take pains to ensure that each new edge receives the proper orientation. To do this, we use the entire J_n sequence from Lemma 3.3. Let a "sequence function" for the digraph be any integer-valued function on the vertices such that each directed edge points from the vertex with lower function value to the vertex with higher value. Proposition 1.2 implies that for any acyclic digraph at least one such function exists. Fix any sequence function. By composing the embeddings implied by the containment sequence, we may identify the vertices of each J_n with vertices of K in a consistent manner. When choosing a polytope and function to realize $J_0 = D_4$, we take the linear function to be the z-coordinate (height) and take the four vertices to have heights equal to the value of the sequence function at the corresponding vertices in K. Additionally, we may ensure that each bending is so small that it has a negligible effect on the (integer-valued) heights of the vertices. Now, whenever bending creates a new vertex, we can choose the bend to intersect the existing edges at heights equal to the sequence function of the corresponding vertices in K. This ensures that each new edge will have the proper orientation.

Again, we would like to emphasize the difference from the undirected case. If J is a 3-polytopal digraph that results from deleting an edge of K, we do not claim that ANY polytope plus linear function that realizes J may be bent to realize K. The new edge might necessarily receive the wrong orientation. Rather, the proof shows that by starting with a particular tetrahedron and linear function we may follow the entire sequence from Lemma 3.3 to arrive at a polytope and linear function that induces the orientation K.

This completes the proof. □

The undirected proof in [BG] shows that arbitrarily close (in the Hausdorff metric) to any 3-simplex in $\mathbb{R}^3$ there are realizations of every 3-polytopal graph. Since all 3-simplices are affinely equivalent, the above proof shows the same is true of realizations of 3-polytopal digraphs.

Grünbaum [Gr1] showed that each combinatorial type of 3-polytope can be realized by a 3-polytope whose vertices belong to the integer lattice in $\mathbb{R}^3$. It follows from our argument that this extends to the lattice realization of an arbitrary combinatorial type along with an arbitrary 3-monotone orientation of its graph.

Gil Kalai asked us which linear orderings of the vertices of a given 3-polytope P can be induced, on some member P' of P's combinatorial class, by some linear function that is strongly admissible for P'. From the above proof of Theorem 3.6 it follows that the linear orderings so inducible are precisely those for which the induced edge-orientation satisfies the condition stated in the theorem.

Theorem 3.6 also yields an immediate corollary characterizing the line shellings of 3-polytopes. Any strongly admissible function determines a sequence of a polytope's vertices. A sequence of a polytope's facets is called a *shelling* if it possesses certain desirable properties, see [Zi]. There is a natural identification between the strongly admissible linear orientations of a polytope and a special class of shellings (called *line shellings*) of its polar, see [DK]. The corollary below follows from this identification and from the definition of a general shelling.

Corollary 3.7. *For a fixed combinatorial class of* 3*-polytopes, a shelling can be realized as a line shelling on some member of the class if and only if there are three independent monotone ridge-paths from the initial facet to the final facet of the shelling.*

Finally, we state without proof a characterization of linear orientations for d-polytopes with $d+2$ vertices. As described in [Gr1], the Gale transform of such a polytope lies in $\mathbb{R}^1$. Further, the combinatorial class is determined by the respective numbers of points with positive, zero, or negative coordinate. Each vertex of the polytope corresponds to a point in the Gale transform, so a sequence of vertices naturally corresponds to a sequence of points in the Gale transform. We say that a sequence *mixes signs* if some point with negative coordinate precedes some point with positive coordinate and some point with positive coordinate precedes some point with negative coordinate.

Theorem 3.8. *A sequence of the vertices of a d-polytope with $d+2$ vertices arises from a linear function acting on some member of the combinatorial class if and only if the sequence mixes the signs of the corresponding points in the Gale transform.*

We can use this characterization to find examples which illustrate that, even in four dimensions, the "directed d-connectedness" of Theorem 3.1 is insufficient to ensure that an orientation arises from a linear function acting on some member of a fixed combinatorial class. (Consider neighborly polytopes.)

Applying the Gale transform to d-polytopes with $d+3$ vertices leads to a characterization of which orientations are inducible by a linear function acting on some member of the combinatorial class. This will appear elsewhere, along with the proof of Theorem 3.8.

Acknowledgments

The authors thank Asa Packer for useful feedback, and the referee for several helpful suggestions including a substantial simplification of the proof of Lemma 3.3.

References

[ADRS] C.A. Athanasiadis, J. DeLoera, V. Reiner, and F. Santos, Fiber polytopes for projections between cyclic polytopes, *European J. Combin.* **21** (2000), 19–47.

[AER] C.A. Athanasiadis, P.H. Edelman, and V. Reiner, Monotone paths on polytopes, *Math. Z.*, to appear.

[AZ] N. Amenta and G.M. Ziegler, Deformed products and maximal shadows of polytopes, in *Advances in Discrete and Computational Geometry*, Contemporary Mathematics, vol. 223, American Mathematical Society, Providence, R.I., 1999 (B. Chazelle, J.E. Goodman, and R. Pollack, eds.), pp. 57–90.

[Ba] M.L. Balinski, On the graph structure of convex polyhedra in *n*-space, *Pacific J. Math.* **11** (1961), 431–434.

[BG] D. Barnette and B. Grünbaum, On Steinitz's theorem concerning convex 3-polytopes and on some properties of 3-connected graphs, *The Many Facets of Graph Theory*, Lecture Notes in Mathematics, vol. 110, Springer-Verlag, Berlin, 1969, pp. 27–40.

[BML] R. Blind and P. Mani-Levitska, On puzzles and polytope isomorphisms, *Aequationes Math.* **34** (1987), 287–297.

[DK] G. Danaraj and V. Klee, Shellings of spheres and polytopes, *Duke Math. J.* **41** (1974), 443–451.

[Ed] P.H. Edelman, Ordering points by linear functionals, *European J. Combin.* **21** (2000), 145–152.

[Gä] B. Gärtner, Combinatorial linear programming: geometry can help, *Proc. 2nd Workshop on Randomization and Approximation Techniques in Computer Science* (*RANDOM*), Springer-Verlag, Berlin, 1998, pp. 82–96.

[GK] P. Gritzmann and V. Klee, Mathematical programming and convex geometry, in *Handbook of Convex Geometry*, North-Holland, Amsterdam, 1993 (P.M. Gruber and J.M. Wills, eds.), pp. 627–674.

[Gr1] B. Grünbaum, *Convex Polytopes*, Interscience/Wiley, London, 1967.

[Gr2] B. Grünbaum, Polytopes, graphs, and complexes, *Bull. Amer. Math. Soc.* **76** (1970), 1131–1202.

[Gr3] B. Grünbaum, Polytopal graphs, in *Studies in Graph Theory, Part II*, MAA Studies in Mathematics, vol. 12, Mathematic Association of America, Washington, D.C., 1975 (D.R. Fulkerson, ed.), pp. 201–224.

[GZ] B. Gärtner and G. Ziegler, Randomized simplex algorithms on the Klee–Minty cubes, *Combinatorica* **18** (1998), 349–372.

[HH] W.M. Hirsch and A.J. Hoffman, Extreme varieties, concave functions and the fixed charge problem, *Comm. Pure Appl. Math.* **14** (1961), 355–369.

[HK] F. Holt and V. Klee, A proof of the strict monotone 4-step conjecture, in *Advances in Discrete and Computational Geometry*, Contemporary Mathematics, vol. 223, American Mathematical Society, Providence, R.I., 1999 (B. Chazelle, J.E. Goodman, and R. Pollack, eds.), pp. 201–216.

[Ka1] G. Kalai, A simple way to tell a simple polytope from its graph, *J. Combin. Theory Ser. A* **49** (1988), 381–383.

[Ka2] G. Kalai, A subexponential randomized simplex algorithm, *Proc. 24th ACM Symposium on the Theory of Computing* (*STOC*), ACM Press, New York, 1992, pp. 475–482.

[Ka3] G. Kalai, Upper bounds for the diameter and height of graphs of convex polyhedra, *Discrete Comput. Geom.* **8** (1992), 363–372.

[Kl] V. Klee, Heights of convex polytopes, *J. Math. Anal. Appl.* **11** (1965), 176–190.

[KM] V. Klee and G. Minty, How good is the simplex algorithm?, in *Inequalities* III, Academic Press, New York, 1972 (O. Shisha, ed.), pp. 159–175.

[Re] J. Renegar, On the computational complexity of the first-order theory of the reals. I, II, III, *J. Symbolic Comput.* **13** (1992), 255–352.

[RG] J. Richter-Gebert, *Realization Spaces of Polytopes*, Springer-Verlag, Berlin, 1996.

[SR] E. Steinitz and H. Rademacher, *Vorlesungen über die Theorie der Polyeder*, Springer-Verlag, Berlin, 1934.

[Ta] A. Tarski, *A Decision Method for Elementary Algebra and Geometry*, University of California Press, Berkeley, Calif., 1951.

[To] M.J. Todd, The monotonic bounded Hirsch conjecture is false for dimension at least 4, *Math. Oper. Res.* **5** (1980), 599–601.

[Zi] G.M. Ziegler, *Lectures on Polytopes*, Springer-Verlag, New York, 1994.

Received March 10, 1999, *and in revised form May* 30, 1999. *Online publication May* 16, 2000.