AFFINE ORIENTATIONS OF POLYTOPES WITH FEW VERTICES

JAMES MIHALISIN

Abstract. Given a combinatorial class of d-polytopes with either $d+2$ or $d+3$ vertices, we provide a good characterization telling which sequences of vertices (as well as which directed graphs) can be induced by an affine function acting on some member of the combinatorial class.

Introduction

In an earlier paper, we defined a *d-polytopal* digraph in a fashion analogous to the definition of a *d-polytopal* graph. 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. 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 acting on P [MK00]. (The orientation is produced by directing each edge from the vertex that takes a smaller value for the affine function to the vertex that takes a larger value.)

A closely related question is the determination of which vertex sequences may be induced by some affine function acting on some member of a fixed combinatorial class of polytopes. (Every vertex sequence determines a digraph, but certain digraphs determine only a partial order of the vertices.) In its dual formulation, the question asks "For a fixed combinatorial class and a given sequence of its facets, is there some representative of the class for which this sequence is actually a line shelling?" If there exists such a representative, the sequence is called a *geometric shelling* [IT00].

Both of the above references attack the question for the case of 3-polytopes. In this paper, we seek the easiest higher-dimensional cases. The case of the simplex is trivial — any vertex sequence is affinely inducible on any (fixed) simplex. Below we provide a simple characterization for any combinatorial class of d-polytopes with only $d+2$ vertices and a tractable description for any combinatorial class of d-polytopes with only $d+3$ vertices. Our primary tool is the Gale transform. Recent work by (ETH Reference), utilizes the Gale transform of a polytope's dual to examine this question for simple d-polytopes with only $d+2$ facets.

Date: 15 May 2001.

1991 *Mathematics Subject Classification.* 52-02, 52B10, 52B11, 52B12, 52B30, 52B55, 05-12, 05C20, 05C40, 05C75, 658Q25, 68U05.

Key words and phrases. polytope, graph, digraph, Gale transform, orientation, linear programming, algorithm, convex.

The author is indebted to Victor Klee for invaluable guidance and to the National Science Foundation for partial support of this research.

1. The Gale Transform

We define the Gale transform in the simplest way that suits our purposes. The concept is defined in more generality in Grünbaum's text [Grü67], where several properties are derived. See Ziegler's text [Zie94] for the connection between the Gale transform and Oriented Matroids.

Let γ be an affine function on $\mathbb{R}^d$. Let $c_0 = \gamma(0)$, the value of γ at the origin, and let $c_i = \gamma(e_i)$ where e_i is the unit vector in the ith coordinate direction for $i \in \{1, \ldots, d\}$. Then for any $v \in \mathbb{R}^d$, written as a column vector, the value of γ at v may be calculated as the product of a row vector and a column vector:

$$\gamma(v) = (c_0, c_1, \ldots, c_d) \begin{pmatrix} 1 \\ v \end{pmatrix}$$

Any d-polytope, whose n vertices are written as column vectors $p_1 \ldots p_n$, may be given the following matrix representation:

$$P = \begin{pmatrix} 1 & 1 & 1 & \ldots & 1 \\ p_1 & p_2 & p_3 & \ldots & p_n \end{pmatrix}$$

The matrix P determines a *range space*, defined as the set of all the row vectors that may be obtained by multiplying P by a row vector on the left. As above, the row vector on the left may be interpreted as representing an affine function. From the construction of the matrix P (with ones along the top row), each vector in range(P) is the *vertex-value vector* for some affine function, γ, acting on the polytope — i.e., the row vector's ith coordinate gives $\gamma(p_i)$.

We define a *Gale transform* of the polytope P to be any full rank matrix G such that $PG = 0$. To see that such a matrix always exists, note that if $p_1, p_2, \ldots p_{d+1}$ are affinely independent then $P = \begin{pmatrix} A & B \end{pmatrix}$ with A invertible. Let I_k be the $k \times k$ identity matrix, with $k = n - (d+1)$. The matrix G defined as

$$G = \begin{pmatrix} -A^{-1}B \\ I_k \end{pmatrix} = \begin{pmatrix} g_1 \\ g_2 \\ \vdots \\ g_n \end{pmatrix}.$$

is then a Gale transform for P.

There is a natural correspondence between the ith column of P and the ith row of G. Since the labelling of the vertices of P is arbitrary, it is natural to interpret a Gale transform matrix G simply as a set of k-dimensional row vectors in one-to-one correspondence with the vertices of P. This will be our perspective for the remainder of the paper. Since the dimension k depends only on the difference between n and d, the Gale transform associates any polytope with "few" vertices (relative to its dimension) with a set of n vectors living in a space of "small" dimension.

By definition, $PG = 0$;this implies range$(P) \subseteq$ null(G). Since G and P are each of full rank, a dimension argument implies the two subspaces are equal. The equality of the vector spaces range(P) and null(G) means the Gale transform preserves the set of vertex-value vectors. In particular, any $(\gamma(p_1), \gamma(p_2), \ldots, \gamma(p_n))$ is in range(P) if and only if it satisfies $\sum_{i=1}^{n} \gamma(p_i) g_i = 0$.

The face-lattice of a polytope is determined by the set of affine functions that are zero at some vertex and non-negative for all the vertices. These functions correspond to supporting hyperplanes. A set S of vertices represents a face of the polytope if and only if there is a vector in range(P) with strictly non-negative coordinates such that the zero coordinates correspond exactly to S. The equality of range(P) and null(G) implies the polytope's face-lattice can also be determined from the vectors in null(G) with strictly non-negative coordinates.

Note that only the signs of the coordinates of vectors in null(G) are used to determine the face-lattice. Let P be a polytope and G be a Gale transform for it. If we scale each row in G by a different positive constant to create G', then even though G' will no longer be a Gale transform for P it will be a Gale transform for some P' that is necessarily in the same combinatorial class as P. (If we start with P and G and apply a projective transformation to P to create $\tilde{P}$, then $\tilde{G}$ will be simply a row rescaling of G. However, not all row rescalings correspond to a projective transformation of the original polytope.)

For a fixed combinatorial class, we wish to determine whether a given sequence of vertices is affinely inducible. Throughout the paper, we use the vertex subscripts to directly indicate the sequence currently in question. (I.e., we relabel indices so that the vertex sequence under consideration is always $\{p_1, p_2, p_3, \ldots, p_n\}$.) Under this convention, the sequence is affinely inducible if and only if there is a vector in range(P) with strictly increasing coordinates. The equality of range(P) and null(G) as well as the correspondence between the ith vertex of P and the ith row vector in G imply that the sequence is affinely inducible if and only if there is a row vector in null(G) with strictly increasing coordinates. I.e., the sequence is affinely inducible if and only if there exists a strictly increasing sequence of coefficients $\{c_i : i \in \{1, \ldots, n\}\}$ such that $\sum c_i g_i = 0$.

Since we will often start with a set of row vectors and work backward to get a polytope, the term *Gale configuration* will refer to a set of row vectors that is the Gale transform of some polytope. Both Grünbaum and Ziegler, [Grü67] and [Zie94], derive criteria for when a set of row vectors constitutes a Gale configuration but these criteria are not needed here since our investigations implicitly begin with a combinatorial class of polytopes. It may easily be shown that if $\{g_i\}$ is any set of vectors whose corresponding matrix has an appropriate null space (i.e. a null space whose non-negative elements correspond to the face-lattice of a polytope), then the set $\{g_i\}$ meets the criteria for being a Gale configuration.

As we have defined the Gale transform, the matrix P has each entry in the top row equal one. This corresponds to a particular embedding of the d-polytope in $(d+1)$-dimensional space. It is clear that requiring this embedding places no restrictions on the affinely induced value vectors of the polytope, but it does imply that for $\{g_i\}$ to be a Gale configuration, we need $\sum g_i = 0$. In other literature, this constraint often is not imposed. For situations where only the face-lattice is needed, it may be omitted. However, we require information about all the affine value vectors — not just the supporting hyperplanes. In our case, the constraint is crucial.

A Gale transform may contain one or more row vectors with all entries zero. As described in Ziegler [Zie94], each of these points corresponds to the apex of a pyramid taken over the polytope represented by the remaining points. (Multiple rows of zeros imply that the process of "taking a pyramid" has been performed multiple times — each time the dimension is increased by one.) It is clear that if v is a vertex that corresponds to a zero vector in the Gale transform, i.e. if v is the apex of a pyramid, then its position in the vertex sequence does not affect whether the sequence is affinely realizable. The apex of a pyramid may be located anywhere in the sequence as long as the rest of the sequence is affinely realizable for the base. (This may be seen either directly from the definition of a pyramid or from the Gale transform reformulation given above.)

For clarity, for the remainder we will restrict attention to Gale configurations which are free of zero vectors.

2. The $d+2$ Case

The Gale transform of a d-polytope with exactly $d+2$ vertices gives vectors on a line. As described in [Grü67], the combinatorial class of a d-polytope with $d+2$ vertices is characterized by three integers – the respective numbers of vectors with positive, zero, or negative coordinate in the Gale transform. As stated in the last section, the effect of adding zero vectors is easily characterized, so for simplicity we will assume that the combinatorial class in question has no zero vectors in its Gale transform.

A vector sequence *mixes signs* for the Gale transform if some vector with negative coordinate precedes some vector with positive coordinate and some vector with positive coordinate precedes some vector with negative coordinate. (The sequence is sign-mixing if and only if it does not begin with all those vectors that have negative coordinate followed by all those with positive coordinate, or vice-versa. Equivalently, if the vertex sequence is represented as a sequence of the symbols "+" and "−", according to the sign of their corresponding vector in the Gale transform, the vertex sequence mixes signs if and only if it possesses either "$+-+$" or "$-+-$" as a subsequence.)

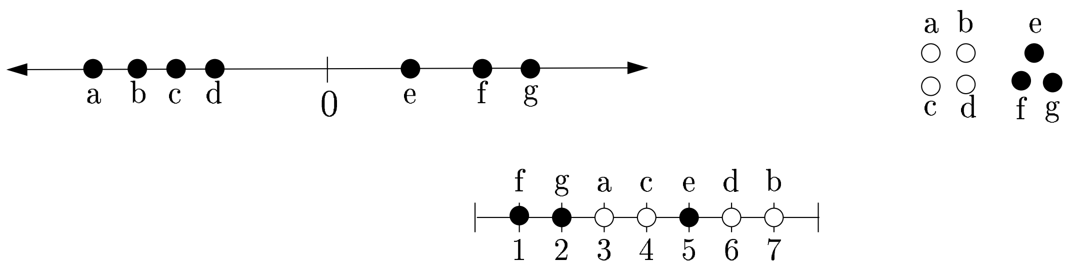

Figure 1. The combinatorics of a polytope with $d+2$ vertices are determined by the number of vectors with positive coordinate in the Gale transform. This implies we can represent the combinatorial class by a "cloud" of positive and negative points (shown by color). The particular vertex sequence shown $\{f, g, a, c, e, d, b\}$ mixes signs, hence by Theorem 2.1 it is linearly inducible.

Theorem 2.1. *A sequence of the vertices of a d-polytope with exactly $d+2$ vertices arises from an affine function acting on some member of the combinatorial class if and only if the sequence mixes the signs of the corresponding vectors in the Gale transform.*

Proof. To see necessity, we start with an affinely inducible order. Let $\{g_i\}$ be the (one-dimensional) vectors of the Gale transform. Affine inducibility implies $\sum c_i g_i = 0$ for some set of strictly increasing $\{c_i\}$. By our definition of the Gale transform $\sum g_i = 0$, which implies that for any fixed α the set $\{\tilde{c}_i\}$ defined as $\tilde{c}_i = c_i - \alpha$ also has $\sum \tilde{c}_i g_i = 0$.

Assume the sequence does not mix signs. This implies the sequence begins with all the positive (or negative) coordinates. By subtracting a constant, we may assure that only the vectors with positive (negative) x coordinate have a negative $\tilde{c}_i$. This leads to a contradiction since $\sum \tilde{c}_i g_i = 0$ then gives a sum of strictly negative (positive) quantities adding to zero.

To prove sufficiency, given a sign-mixing vertex sequence for a chosen combinatorial class, we shall explicitly construct a Gale configuration, G, that represents a member of the class and whose null space possesses a row vector with strictly increasing coordinates. In fact, we shall prove a stronger result. Given *any* increasing sequence $\{c_i\}$, we shall construct a Gale configuration for a polytope that has $(c_1, \ldots, c_n)$ as a vertex value vector.

We begin with a fixed combinatorial class of polytope. This determines the sign of each one-dimensional vector, g_i. Partition $\{1, \ldots n\}$ into two sets, $A = \{i \in \{1, \ldots, n\} : g_i > 0\}$ and $B = \{j \in \{1, \ldots, n\} : g_j < 0\}$. Consider the interval $[c_1, c_n]$. (Recall that $\{c_i\}$ is an arbitrary strictly increasing sequence.) The mixing signs requirement implies that the convex hull of $\{c_i\}_{i \in A}$ overlaps the convex hull of $\{c_j\}_{j \in B}$. Choose any point, p, in the intersection and write it as two different convex combinations, $p = \sum_{i \in A} a_i c_i$ and $p = \sum_{j \in B} b_j c_j$. By perturbing the coefficients, we may assume that each $a_i > 0$ and each $b_j > 0$. We will use the coefficients for these convex combinations as the actual coordinates for our Gale configuration. Set $g_i = a_i$ for all $i \in A$ and $g_j = -b_j$ for all $j \in B$. This defines g_k for all $k \in 1, \ldots, n$.

Finally, let $G = \begin{pmatrix} g_1 & \ldots & g_n \end{pmatrix}^T$. The combinations above are convex, so $\sum_{i \in A} a_i = \sum_{j \in B} b_j = 1$; thus $\sum_{k=1}^{n} g_k = 0$. Further, $\sum_{i \in A} a_i c_i = \sum_{j \in B} b_j c_j = p$; thus $\sum_{k=1}^{n} c_k g_k = 0$. To see that G is a Gale configuration that represents the appropriate combinatorial class, we need only verify each coordinate has the appropriate sign. This is clear from the definition of A and B. By construction, G has $(c_1, c_2, \ldots, c_n)$ in its null space, hence the polytope corresponding to G has $(c_1, c_2, \ldots, c_n)$ as a (strictly increasing) affinely induced vertex value vector. □

The choice of a Gale configuration does not uniquely specify the polytope of which it is the transform. For example, any affine tranformation may be applied to the polytope or any linear transformation may be applied to the Gale configuration while the other is kept fixed. The polytope may be uniquely specified by imposing additional constraints. For example, by employing an affine tranformation, we may require that the first $d+1$ vertices consist of the origin and canonical points along the coordinate axes. This gives P the form:

$$P = \begin{pmatrix} 1 & 1 & 1 & 1 & \ldots & 1 & 1 \\ 0 & 1 & 0 & 0 & \ldots & 0 & x_2 \\ 0 & 0 & 1 & 0 & \ldots & 0 & x_3 \\ \vdots & & & & & & \\ 0 & 0 & 0 & 0 & \ldots & 1 & x_{d+1} \end{pmatrix}$$

In this case, P will have G as its transform if and only if $x_i = -g_i/g_{d+2}$.

3. The $d+3$ Case

When the number of vertices in a d-polytope is $d+3$, the Gale transform produces 2-dimensional row vectors. In this entire section, all vectors will be assumed to be in the plane.

Since the magnitudes of the vectors of the Gale transform do not affect the face-lattice of the polytope, we often wish to consider "directions" rather than actual vectors. For our purposes, a *direction* is simply a vector of unit length. A vector v is *in the direction of* w if $v = \lambda w$ for some positive λ. A set of directions will be called *spanning* if each vector in the plane may written as a sum of vectors in these directions. A set of vectors is *spanning* if their associated directions are spanning. (A set of vectors is spanning if and only if the cone they generate is the entire plane.) Given a set of vectors, define a *subconfiguration* to be a set of vectors in one-to-one correspondence with a subset of the original vectors, where the directions of corresponding vectors are identical. (Magnitudes may vary.)

Each nonzero vector determines a ray issuing from the origin. For this paragraph, we will call the ray in the opposite direction the vector's *tail ray*. Given a Gale configuration in the plane, the face-lattice of its associated polytope is not changed by rotating a vector as long as its tail ray does not cross over any other vector. The heads of the vectors are free to cross over each other without changing the face-lattice.

Note that if a Gale configuration contains a pair or more of vectors in opposite directions then there are constraints on how such vectors may be rotated. If g_1 and g_2 are in opposite directions then g_1 may be rotated only if g_2 is rotated in tandem, keeping the two opposite. If greater numbers of vectors are involved, say five in one direction and two in the opposite direction, then the entire set must be rotated in tandem. (In the example, all seven would have to move in synch.) However, if five vectors lie in one direction without any vector in the direction opposite then each of the five could be rotated individually by small amounts without changing the face-lattice of the associated polytope.

This freedom to rotate allows us to coalesce the vectors into groups, then rotate the groups into "canonical directions" (i.e., those with angle $(2\pi)/k$ for some minimal k). This process creates a multiset — we assign to each canonical direction an integer to indicate how many vectors have been coalesced into that direction. Details are given in "Convex Polytopes", [Grü67]. The result, a multiset of vectors in canonical directions, scaled to be of unit length, is called a *contracted Gale diagram*. Theorem 3.1 is also described in [Grü67].

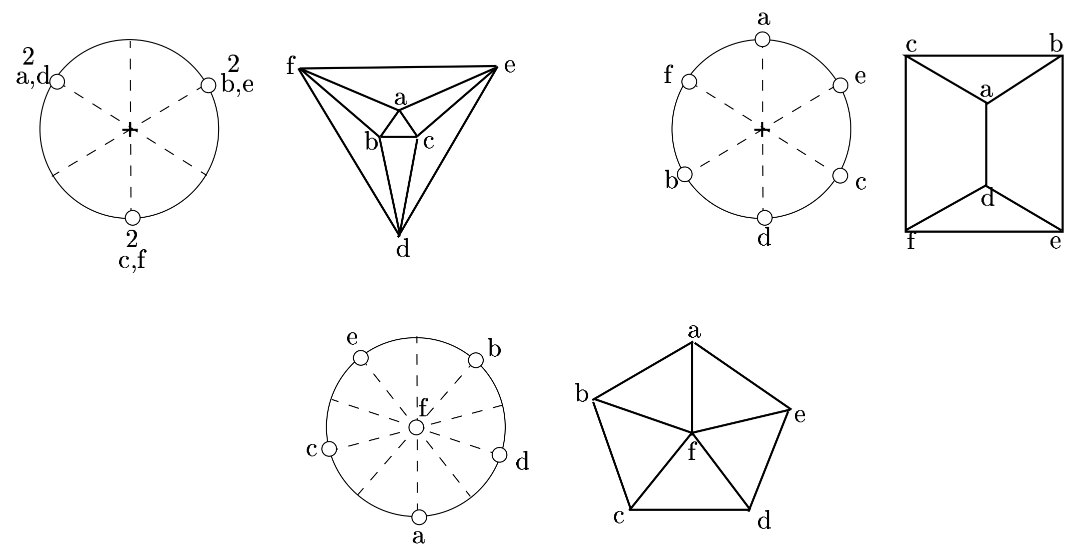

FIGURE 2. Here are the contracted Gale diagrams and graphs for some familiar 3-polytopes. In the first diagram, each 2 indicates that two vectors have been coalesced into that direction.

Theorem 3.1. *Two d-polytopes having exactly $d+3$ vertices are combinatorially equivalent if and only if their contracted Gale diagrams are congruent (i.e. related by a rotation and/or reflection).*

Certain simple rules dictate when a set of unit vectors in canonical directions constitutes a contracted Gale diagram (i.e., when the set actually represents a combinatorial class of polytopes). Again, we do not need to consider these criteria explicitly since we begin with a chosen combinatorial class. We refer the reader to Grünbaum's text, [Grü67], for the list of these rules.

Just as the enumeration of combinatorial classes can be accomplished using only contracted Gale diagrams, the main result of this section will show that the question of linear inducibility for vertex sequences also can be answered by considering just the contracted diagrams.

Our plan of attack is as follows: we first fix a set of directions — this is more restrictive than simply specifying a combinatorial class. We then introduce the notion of a "target cone" and show that it suffices to seek subconfigurations of vectors whose sum is in the target cone. After investigating weighted sums of vector dependencies for triples or quadruples, we show that it suffices to consider just these small sets. Finally, we show that the choices made in initially fixing the set of directions were largely irrelevant, which allows us to characterize the results in terms of contracted Gale diagrams.

Given a set of vectors or directions, $\{g_i : 1 \leq i \leq n\}$, define its *open cone*, written as $\text{cone}(\{g_i\})$, to be the set of all possible sums $\{\sum_{i=1}^{n} \lambda_i g_i$: where each $\lambda_i > 0\}$. Later we will need the *closed cone*, written as written as $\text{cone}[\{g_i\}]$, and defined to be the set of all possible sums $\{\sum_{i=1}^{n} \lambda_i g_i$: where each $\lambda_i \geq 0\}$. (If $\{g_i\}$ has just two elements, we will omit the brackets and write either $\text{cone}(g_1, g_2)$ or $\text{cone}[g_1, g_2]$.)

We begin by fixing a multiset of directions that corresponds to our chosen combinatorial class of polytope. Let d_1 be the first direction in the sequence and let d_n be the last direction in the sequence. Let $\{o_m\}_{m \in M}$ be the set of all directions for which the opposite direction occurs later in the sequence. (The existence of pairs of opposite directions is determined by the combinatorial class. Each pair corresponds to a cofacet consisting of exactly two vertices.) Define T to be $\text{cone}(\{d_1, -d_n\} \cup \{o_m\}_{m \in M})$. T is called the *target cone*. (Note: if d_1 and d_n are opposite directions and M contains no other directions then T will be a single ray. We'll refer to this degenerate case as a *ray target cone*. For the non-degenerate case, T is non-empty and open.)

The next lemma reformulates the criteria for being linearly inducible. It is sufficient to consider spanning subconfigurations of the direction set and to look for weighted sums which are not required to be zero but which are merely required to lie in T.

Lemma 3.2. *Given a multiset of directions $\{d_i\}$, let T, the target cone, be defined as above. The sequence $\{1, 2, \dots n\}$ is linearly inducible if and only if there exists a spanning subconfiguration $\{\tilde{g}_i\}$ of $\{d_i\}$ that satisfies $\sum \tilde{g}_i = 0$ as well as a strictly increasing sequence of coefficients $\{\tilde{c}_i\}$ such that $\sum \tilde{c}_i \tilde{g}_i \in T$.*

Proof. The sequence $\{1, 2, \dots n\}$ is linearly inducible for our choice of directions, if and only if there exists a set of vectors $\{g_i\}$ with each g_i in the direction of d_i and $\sum g_i = 0$, as well as a strictly increasing sequence of coefficients $\{c_i\}$ such that $\sum c_i g_i = 0$.

To prove the "only if" direction for Lemma 3.2, we may choose $\{\tilde{g}_i\} = \{g_i\}$ (a valid subconfiguration), choose $\{\tilde{c}_i\}_{i=2}^{n-1} = \{c_i\}_{i=2}^{n-1}$, choose any $\tilde{c}_1$ such than $c_1 < \tilde{c}_1 < c_2$ and any $\tilde{c}_n$ such that $c_{n-1} < \tilde{c}_n < c_n$. This guarantees $\sum \tilde{c}_i \tilde{g}_i \in T$.

To prove the "if" direction, we begin with a subconfiguration $\{\tilde{g}_j\}_{j \in J}$ that satisfies the hypotheses above. We must first expand it to include all the directions of $\{d_i\}$. Then we must adjust coefficients and/or vectors to produce a set $\{g_i\}$ such that $\sum c_i g_i = 0$. For simplicity, assume the index set, J, is the subsequence of $\{1, \ldots, n\}$ such that for each $j \in J$, $\tilde{g}_j$ is in the same direction as g_j.

First, any direction d_l that is not represented (including multiplicities) in $\{\tilde{g}_j\}_{j\in J}$ (i.e. any l such that $1 \le l \le n$ and $l \notin J$) may be represented by a vector $\bar{g}_l$ with very small magnitude. Since $\{\tilde{g}_j\}_{j\in J}$ is a spanning set, the magnitude of each $\tilde{g}_j$ may be adjusted slightly to produce a $\bar{g}_j$ in such a way that $\sum_{i=1}^n \bar{g}_i = 0$. Set $\tilde{c}_j = \bar{c}_j\ \forall j \in J$. Then arbitrarily extend $\{\tilde{c}_j\}_{j\in J}$ to create an increasing sequence $\{\bar{c}_i\}_{1\le i\le n}$.

By taking the $\bar{g}_l$ for $l \notin J$ to be sufficiently small, we can make $\sum_{i=1}^n \bar{c}_i\bar{g}_i$ arbitrarily close to $\sum_{i\in J} \tilde{c}_i\tilde{g}_i$. Since $\{\tilde{g}_j\}_{j\in J}$ includes a spanning set and since the $\{\bar{c}_i\}$ are strictly increasing, we may adjust the $\{\bar{c}_i\}$ slightly to erase any small difference between the two sums, thus we may ensure that $\sum_{i=1}^n \bar{c}_i\bar{g}_i = \sum_{i\in J} \tilde{c}_i\tilde{g}_i$.

Finally, $\bar{c}_n$ may be increased (or $\bar{c}_1$ may be decreased) without changing the fact that $\{\bar{c}_i\}$ is increasing. The effect on $\sum \bar{c}_i\bar{g}_i$ is to add a positive multiple of $\bar{g}_n$ (or $-\bar{g}_1$) to the sum. Additionally, for any vector $\bar{g}_k$ in the direction of some o_m, both $\bar{g}_k$ and a vector in the opposite direction may be lengthened additively without changing the fact that $\sum \bar{g}_i = 0$. The effect on $\sum \bar{c}_i\bar{g}_i$ is to add a negative multiple of o_m. Thus if $\sum \bar{g}_i = 0$ and $\sum \bar{c}_i\bar{g}_i \in T$, these two operations allow us to find a set of vectors $\{g_i\}$ in the same directions as the $\{\bar{g}_i\}$ and an increasing sequence $\{c_i\}$ such that $\sum g_i = 0$ and $\sum c_ig_i = 0$. □

We now consider some simple lemmas about weighted sums of small sets of vectors.

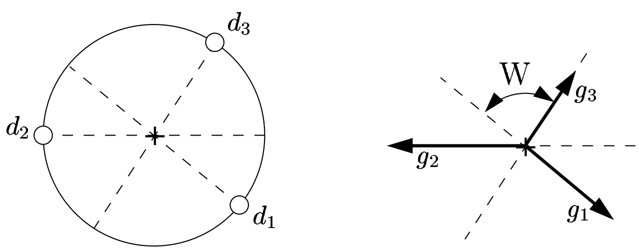

FIGURE 3. The left figure simply indicates the three directions chosen. The three vectors on the right sum to zero. In the text we show that any weighted sum with strictly increasing coefficients must be in the cone W.

Lemma 3.3. *Given three vectors $\{g_i\}_{i=1}^3$ whose sum is zero, for any strictly increasing sequence $c_1 < c_2 < c_3$, the weighted sum $\sum_{i=1}^3 c_ig_i$ lies in the open cone formed by $-g_1$ and g_3. Any vector in* $\text{cone}(-g_1, g_3)$ *may be obtained by an appropriate choice of $\{c_i\}$.*

Proof. Since $\sum g_i = 0$, for any constant λ we may subtract $\lambda \sum g_i$ from $\sum c_ig_i$ without changing the value of the sum. Thus we may assume $c_2 = 0$. This implies that $\sum c_ig_i$ lies in $\text{cone}(-g_1, g_3)$.

If $w \in W = \text{cone}(-g_1, g_3)$ then $w = -\alpha g_1 + \beta g_3$, where α and β are positive. By taking $c_1 = 1 - \alpha < c_2 = 1 < c_3 = 1 + \beta$, we get $\sum c_ig_i = w$. □

With just three directions, requiring a set of vectors in these directions to sum to zero fixes the relative magnitudes of the vectors. With four directions, we have more freedom in choosing a set of the vectors that gives a zero sum. However, there is still much we can say about $\sum_{i=1}^4 c_ig_i$ when $\sum_{i=1}^4 g_i = 0$ and $c_1 < c_2 < c_3 < c_4$.

Lemma 3.4. *Let $\{d_1, d_2, d_3, d_4\}$ be a spanning quadruple of directions and let $\{c_i\}$ be an increasing sequence. For any set of vectors $\{g_i\}$ in the directions $\{d_i\}$ with $\sum g_i = 0$, the weighted sum $\sum c_i g_i$ may be written as the sum of two vectors w and v, with $w \in W = \text{cone}(-d_1, d_4)$ and $v \in V = \text{cone}(-d_1, -d_2) \cap \text{cone}(d_3, d_4)$. Any vector in $\text{cone}(W \cup V)$ may be obtained by an appropriate choice of $\{c_i\}$ and $\{g_i\}$.*

Proof. Again, we may subtract a constant times $\sum g_i$ from the $\sum c_i g_i$ without changing the result. Thus we may assume that $-c_2 = c_3 = c$. Note that $c > 0$.

We rewrite the sum as:

$$\begin{aligned}\sum c_i g_i &= (c_1 + c)g_1 - cg_1 - cg_2 + cg_3 + cg_4 + (c_4 - c)g_4 \\ &= -\alpha g_1 + v + v + \delta g_4\end{aligned}$$

where $\alpha = -(c_1 + c) > 0$ and $\delta = c_4 - c > 0$ and $v = -c(g_1 + g_2) = c(g_3 + g_4)$.

This implies that $\sum c_i g_i$ is the sum of a vector $2v \in V = \text{cone}(-g_1, -g_2) \cap \text{cone}(g_3, g_4)$ and a vector $w = -\alpha g_1 + \delta g_4 \in W = \text{cone}(-g_1, g_4)$.

To prove the reverse containment, take any $x = (2v + w) \in \text{cone}(W \cup V)$. By the definition of V it is clear that we may find some set of $\{g_i\}$ in the directions $\{d_i\}$ such that $g_1 + g_2 = -v$ and $g_3 + g_4 = v$. If $w = -\alpha g_1 + \delta g_4$ (with $\alpha, \delta > 0$), then we may take $c_1 = -1 - \alpha, c_2 = -1, c_3 = 1, c_4 = 1 + \delta$ to get $\sum c_i g_i = x$. $\square$

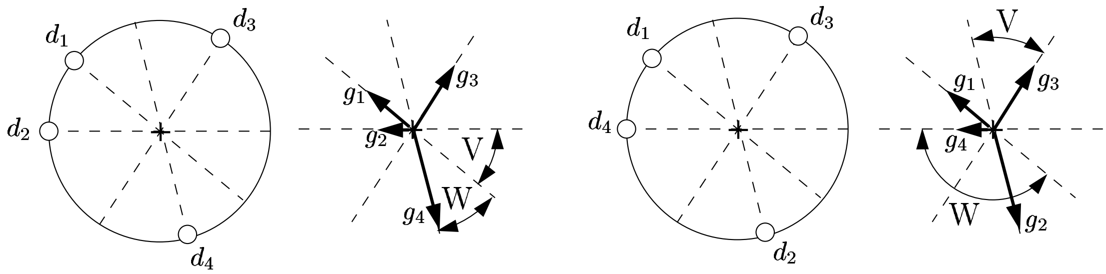

FIGURE 4. There are different possibilities for V and W depending on the configuration. On the left, V and W touch. On the right, the $\text{cone}(W \cup V)$ is the entire plane.

The following lemma shows that understanding subconfigurations of quadruples will be sufficient for understanding larger configurations.

Lemma 3.5. *Let T, the target cone, be defined as above. If $\{g_i\}_{i=1}^m$ is a configuration of more than four vectors with $\sum_{i=1}^m g_i = 0$ and $\sum_{i=1}^m c_i g_i \in T$ for some increasing $\{c_i\}$ then there exists a subconfiguration of four vectors $\{\tilde{g}_j, \tilde{g}_k, \tilde{g}_p, \tilde{g}_q\}$ and increasing coefficients $\tilde{c}_j < \tilde{c}_k < \tilde{c}_p < \tilde{c}_q$ such that $\sum \tilde{g}_i = 0$ and $\sum \tilde{c}_i \tilde{g}_i \in T$.*

Proof. We first check the case where the first three directions and the last three directions $\{g_i\}_{i=1}^m$ are each a spanning triple.

Choose any $t \in T$. Since the first triple and the last triple are spanning sets, the vector t must lie in $\text{cone}[-g_j, -g_k]$ for some $j, k \in \{1, 2, 3\}$ and it must also lie in $\text{cone}[g_p, g_q]$ for some $p, q \in \{m-2, m-1, m\}$. Take the four directions to be those corresponding to j, k, p and q. The conditions above then imply the quadruple is a spanning set (note the presence of negative signs for g_i and g_j). Assuming that t does not equal $-g_j$, $-g_k$, g_p or g_q, we see that t lies in the intersection cone V defined in Lemma 3.4. The Lemma then implies there exists vectors

$\tilde{g}_i$ in the appropriate directions and increasing coefficients $\tilde{c}_i$ such that $\sum \tilde{c}_i \tilde{g}_i = t \in T$. If t does equal one of the four vectors above, it is straightforward to show that it is still possible possible to choose a quadruple that possesses a weighted sum equal to t. (In this situation, t will be on the border of V. There are several cases to check, but a quadruple of directions may always be chosen so that cone$(W \cup V)$ includes that border.)

This leaves the case where either the first three directions or the last three directions of $\{g_i\}$ do not span the plane. In each of these cases, we shall create a subconfiguration that uses one less vector hence we'll be done by induction. (Note that the proof allows $\{g_i\}$ to be a subconfiguration of the original configuration used to define T, so induction is valid.) The case involving the last three directions is symmetric to that involving the first three (modulo a minus sign), so we shall check only the situation where the first three directions do not span the plane.

Again, we may use the fact that $\sum g_i = 0$ to add or subtract a constant from the $\{c_i\}$. This allows us to assume that $c_1, c_2, c_3 < 0$ while all the other c_i are strictly positive. Up to symmetries, there are only three configurations possible. However, in each of these configurations, it is possible for the situation to be "degenerate", meaning that two of the vectors lie in opposite directions. We will check one of the three configurations in detail (including the degenerate situation). The verification for the other two configurations follows the same pattern.

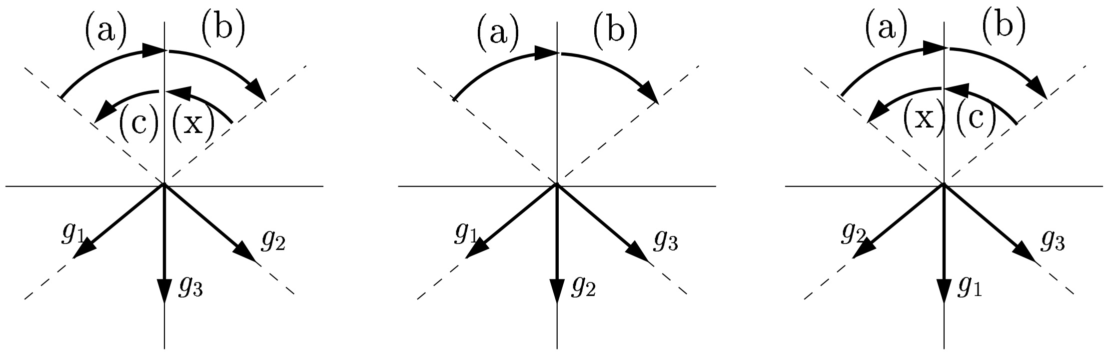

FIGURE 5. The figure shows the three non-degenerate configurations possible in the case that $\{g_i\}_{i=1}^3$ do not span the plane. The arrowed arcs indicate the rotation of u relative to $-w$, where u is the weighted sum of the $\{g_i\}_{i=1}^3$ and w is their unweighted sum. For example, (a) in the left indicates w in cone(g_2, g_3) and u rotated clockwise relative to $-w$.

The left diagram of Figure 5 shows the situation where g_3 is contained in $C = \text{cone}(g_1, g_2)$ (or it represents the degenerate case of g_1 and g_2 opposite, in this case let $C = \text{cone}(g_1, g_2, g_3)$). Clearly, $w = \sum_{i=1}^3 g_i \in C$ and $u = \sum_{i=1}^3 c_i g_i \in -C$. Comparing u with $-w$, we may find that the direction of u is rotated either clockwise or counter-clockwise relative to w. This depends on the relative magnitudes of the g_i and c_i. Note that adding a multiple of $-g_1$ to any vector in $-C$ will rotate its direction clockwise. Adding a multiple of $-g_2$ to any vector in $-C$ will rotate its direction counter-clockwise.

If the pair $\{g_1, g_2\}$ is not degenerate and u is rotated clockwise relative to $-w$, we will be able to replace the triple $\{g_1, g_2, g_3\}$ with a subconfiguration $\{\tilde{g}_1, \tilde{g}_2\}$ where each $\tilde{g}_i$ is in the same direction as g_i. Clearly we can find some $\tilde{g}_1$ and $\tilde{g}_2$ to give $\tilde{g}_1 + \tilde{g}_2 = w$. Since u is clockwise relative to $-w$, for some positive μ the vector $-w - \mu\tilde{g}_1$ is in the direction of u. Thus for some positive λ, $u = \lambda(-w - \mu\tilde{g}_1)$. Setting the coefficients $c_1 = \lambda(-1-\mu)$ and $c_2 = -\lambda$ gives $\tilde{c}_1 < \tilde{c}_2 < 0$ with $\tilde{c}_1\tilde{g}_1 + \tilde{c}_2\tilde{g}_2 = u$.

If u is rotated counter-clockwise relative to $-w$, we claim that w must lie in $\mathrm{cone}(g_2, g_3)$. In this case we may replace the triple $\{g_1, g_2, g_3\}$ with some $\{\tilde{g}_2, \tilde{g}_3\}$ in a similar fashion. (The situation where u is rotated counter-clockwise relative to $-w$, yet w lies in $\mathrm{cone}(g_1, g_3)$ is labelled "Situation (x)" in Figure 5. The claim is that this situation cannot occur.)

The claim may be intuitively obvious (or sufficiently plausible), in which case the reader may skip the tedious details of this paragraph. To prove the claim, let $\hat{y}$ be the direction perpendicular to w that is clockwise from $-w$. If the direction of u is counter-clockwise relative to $-w$, then u has a negative $\hat{y}$ coordinate. This means $u - c_3w$ also has negative $\hat{y}$ coordinate. (Please see the figure below.) In our orientation, g_1 has negative $\hat{y}$ coordinate. Since $c_1 < c_2$, $u - c_3w + (c_2 - c_1)g_1$ also has negative $\hat{y}$ coordinate, but this sum is just $-(c_3 - c_2)(g_1 + g_2)$ whence $g_1 + g_2$ has positive $\hat{y}$ coordinate. Since $g_1 + g_2$ has positive $\hat{y}$ coordinate and $-w$ has zero $\hat{y}$ coordinate, their sum must have positive $\hat{y}$ coordinate but their sum simply equals $-g_3$. If w were in $\mathrm{cone}(g_1, g_3)$, then $-g_3$ would necessarily be counter-clockwise to $-w$ and hence would have negative $\hat{y}$ coordinate. Therefore if u is counter-clockwise relative to $-w$, we cannot have $w \in \mathrm{cone}(g_1, g_3)$.

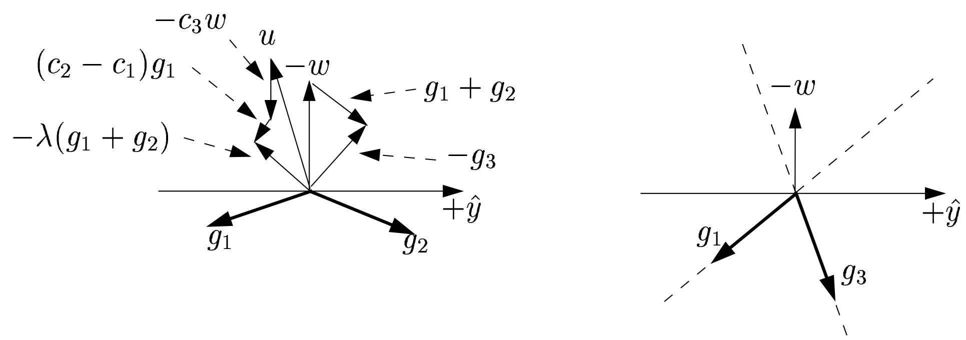

FIGURE 6. The left diagram illustrates the construction to prove that $g_1 + g_2$ has positive $\hat{y}$ coordinate. This implies that $-g_3$ has positive $\hat{y}$ coordinate as well. The right diagram illustrates that w in $\mathrm{cone}(g_1, g_3)$ implies that $-g_3$ has negative $\hat{y}$ coordinate.

Finally, we consider the degenerate case. Here we cannot replace the triple with the pair $\{g_1, g_2\}$ since these two are in opposite directions. If w lies in $\mathrm{cone}(g_1, g_3)$, then we may replace the triple with a subconfiguration $\{\tilde{g}_1, \tilde{g}_3\}$ as above. If w lies in $\mathrm{cone}(g_2, g_3)$, we may find some positive λ such that $u + \lambda g_1$ is counter-clockwise to w. Then the paragraphs above show us how to replace the triple with some $\{\tilde{g}_2, \tilde{g}_3\}$, such that $\tilde{g}_2 + \tilde{g}_3 = w$ and $\tilde{c}_2\tilde{g}_2 + \tilde{c}_3\tilde{g}_3 = u + \lambda g_1$. The net effect would be to change the weighted sum from $\sum_{i\in I} c_i g_i = t$ to $\sum_{j \in J} \tilde{c}_j \tilde{g}_j = t + \lambda g_1$ where $|J| = |I| - 1$. In the degenerate case, g_1 is one of the defining vectors for T (not necessarily as the first vector in the original configuration but as a member of M since it is opposite to g_2 and comes earlier in the sequence). Thus $t \in T$ implies $t + \lambda g_1 \in T$ for any positive λ. Since we are able to reduce the number of vectors while keeping the weighted sum in T, the induction step is complete.

The remaining two diagrams in Figure 5 may be summarized as follows. In the middle diagram, for situation (a) $\{g_2, g_3\}$ can replace the triple. For situation (b), $\{g_1, g_2\}$ can replace the triple. No counter-clockwise rotations are possible.In the right diagram, for (a) $\{g_1, g_3\}$ replaces the triple. For (b), a non-degenerate $\{g_2, g_3\}$ replaces the triple, otherwise the situation may be converted to (c). For (c), $\{g_1, g_2\}$ replaces the triple. Situation (x) (as indicated in Figure 5) cannot occur. □

The final step is to relax our restriction of fixed directions. Since the contracted diagram coalesces all vectors that can be rotated to cross over each other, surprisingly little can be gained by rotating vectors away from their canonical directions. With the exception of one powerful special case, the only effect of rotating vectors is to allow certain cones to be closed rather than open.

We begin with the special case. If d_i and d_j lie in the same canonical direction and the contracted Gale diagram has no vector in the opposite direction, then any cone that includes both d_i and $-d_j$ can be assumed to contain either of the half-planes bounded by the line through d_i. This follows since in such a case we are free to rotate d_i by a small amount in either direction. This small change in d_i changes cone$(d_i, -d_j)$ dramatically, giving us the freedom to include any vector from either open half-plane. It is clear that apart from this special case, any small rotation of vectors will produce only a small change in cones that include them.

The above special case can arise in just two ways. First, if the initial and final vectors of the entire sequence lie in the same canonical direction with no vector in the opposite direction, then the target cone can include either of a pair of half-planes, hence ANY choice of sequence on the remaining vertices must be linearly inducible. Second, if some spanning quadruple has its lowest and highest vectors in the same canonical direction with no vector of the configuration lying in the opposite direction, then the cone of possible sums for the quadruple may be taken to be the entire plane. In this case as well, the sequence must be linearly inducible irrespective of the order for the remaining $n-4$ vectors. (To verify these are the only two ways, first note that such vectors can only arise in the target cone if they are the initial and final vectors in the sequence. Second, notice that any spanning triple cannot include two vector in the same direction. Finally, a case by case analysis of the possibilities for spanning quadruples shows that the only significant effect is when the special pair are lowest and highest vectors of the quadruple.)

We previously defined open cones and closed cones. We define the *half-open cone* of two vectors to be cone$(g_i, g_j] = \{\lambda g_i + \mu g_j : \lambda \geq 0 \text{ and } \mu > 0\}$. Since any closed or half-open cone contains the origin, an intersection of two cones is called *non-trivial* if it contains a vector other than the zero vector.

Lemma 3.3 motivates the following definition. (Below, PS is short for "possible sums".) Assume $j < k < l$. Let $\{d_j, d_k, d_l\}$ be a spanning triple of canonical directions. If the contracted diagram contains vectors in the directions opposite to d_j and d_l., then the *PS cone for the triple* is defined to be cone$(-d_j, d_l)$, the open cone. If there are no vectors in directions opposite to d_j or d_l, then the *PS cone* is defined to be cone$[-d_j, d_l]$, the closed cone. If just one of d_j or d_l has a vector in the opposite direction then the *PS cone* is defined as half-open, either cone$(-d_j, d_l]$ if a vector is opposite d_j, or cone$[-d_j, d_l)$ if a vector is opposite d_l.

Similarly, Lemma 3.4 allows us to define PS cones for quadruples of canonical directions. In the special case described above, the PS cone is the entire plane. Otherwise, the lemma shows that for fixed directions the weighted sum will lie in the cone$(V \cup W)$. This cone will be either the entire plane or cone$(\pm d_i, \pm d_j)$ for some choice of signs and some i and j taken from the original four directions. If d_i has no vector in the opposite direction, the *PS cone for the quadruple* is defined to include the ray in the appropriate d_i or $-d_i$ direction. Similarly for d_j. Thus the PS cone is defined as open, half-closed, or closed.

In an analogous fashion, the target cone is now defined as open, half-closed or closed under the appropriate circumstances.

Theorem 3.6. *A sequence of vertices of a d-polytope with exactly $d+3$ vertices is affinely inducible if and only if its contracted Gale diagram possesses a spanning triple or a spanning quadruple of directions whose PS cone has a non-trival intersection with T, the target cone.*

Proof. Lemmas 3.2 through 3.5 prove the equivalence under the constraint of fixed directions. To prove the theorem, we must show that: (a) the flexibility to rotate vectors from their canonical directions allows certain cones to be considered as closed or half-closed, (b) other than the special cases described above, this is the only advantage to be gained by rotating vectors.

We noted earlier that the heads of vectors were free to cross over each other as long as those heads *do not* cross over the tail direction of any other vector. Such rotations will not change the face-lattice of the polytope. These conditions imply that the vectors of a Gale configuration are partitioned into equivalence classes, whereby two vectors are free to cross over each other if and only if they are in the same class. In the contracted Gale diagram, each such class has been coalesced into a single canonical direction. Moreover, all vectors in a canonical direction form a single class if and only if the diagram does not contain any vector in the opposite direction. If the opposite direction is occupied, then none of the vectors in the original direction may cross over each other.

The previous lemmas show that for fixed directions, the question of linear inducibility has been reduced to whether certain open cones overlap. It was necessary to use open cones since to give each direction a vector with non-zero magnitude we required the flexibility to erase small perturbations.

We have already handled all cases where rotating vectors by small amounts causes dramatic changes in the cones they produce. The only other way that rotating vectors can cause two cones to overlap is if the heads of their extremal vectors are allowed to cross over.

We now can see immediately both (a) and (b) from above. The cones were defined to be half-closed or closed precisely when an extremal vector had the freedom cross over other vectors sharing its canonical direction. At a minimum, such freedom clearly allows the cones in these situations to be considered half-closed or closed, thus proving (a). Moreover, in the process of creating the contracted Gale diagram, all vectors that have the freedom to cross over have been coalesced into a single direction. This implies (b) — beyond including certain rays (to create a half-closed or closed cone) no further overlap of cones can be accomplished by rotating vectors without changing the face-lattice. □

We calculate the cones in Theorem 3.6 as if the vectors were constrained to lie in canonical directions. This considerably simplifies the situation since it makes the problem discrete. Simply checking all spanning triples and quadruples requires at most $O((d+3)f)$ operations where f is the number of facets. By the Upper Bound Theorem applied to d-polytopes with $d+3$ vertices, f is less than or equal to $\sum_{i=0}^{d/2*}(i+2)(i+1)$ (where the asterisk means that we take only half the last term for $i = d/2$ if d is even, or we take the whole last term for

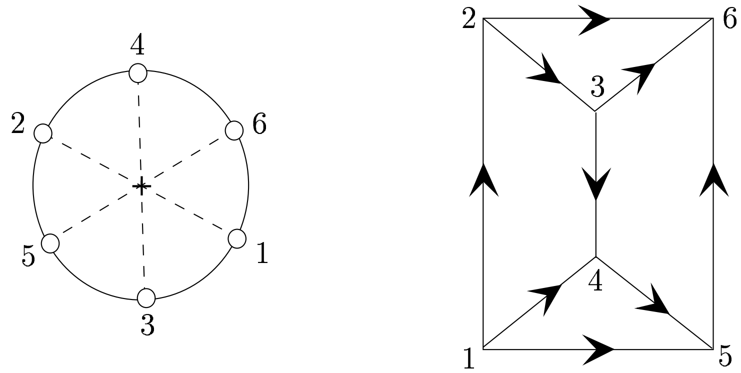

FIGURE 7. This example demonstrates the critical distinction between open and closed cones. Here the boundary of the target cone overlaps the boundary of each PS cone. For this configuration, these cones are all defined to be open hence there is no intersection. Note that this sequence cannot possibly be affinely inducible since there are only two independent monotone paths from source to sink in the orientation.

$i = (d-1)/2$ if d is odd). This bounds the number of operations by $O(d^4)$. It seems likely that an even better checking algorithm exists.

The "local nature" of this characterization is somewhat surprising — for a given initial and final vertex, the relative sequence on a single spanning quadruple can guarantee affine inducibility *regardless* of the sequence on the remaining vertices! One might expect that such a restricted examination could only forbid inducibility. (E.g., the order of a subset of vertices might indicate two sources, a cycle or some other impossibility.) It seems counter-intuitive that a restricted examination can guarantee inducibility. (Note: in the $d+2$ characterization the "local nature" can also be seen since sign mixing is the equivalent to the existence of "$+-+$" or "$-+-$" as a subsequence.)

The following points might make this counter-intuitive behavior more plausible:

(a) The triples and quadruples are not arbitrary. They must include an entire co-facet. This implies that the remaining vertices all lie in a hyperplane.

(b) Changing a single vector in the Gale configuration can affect more than one vertex of the polytope (by the nature of matrix multiplication), hence what seems "local" in the Gale picture could have global consequences for the polytope.

(c) The original question allows the polytopes to vary over the entire combinatorial class. Our construction yields special polytopes whose geometry is somehow dominated by the single triple or quadruple that defines the relevant PS cone. (Recall that the vectors comprising the triple or quadruple are chosen to have magnitudes much larger than the other vectors.)

4. Difficulties Encountered with More General Polytopes

Some of the concepts of the $d+3$ proof certainly carry over to the cases with more vertices, but various difficulties make a complete characterization unlikely. (The complexity of d-polytopes with more than $d+3$ vertices has already been amply demonstrated in other contexts.)

In some sense, the sufficiency part of the "local nature" is not restricted to polytopes with few vertices.

Lemma 4.1. *For any set of vectors* $\{g_i\}$ *with* $\sum_{i=1}^{k} g_i = 0$, *and any increasing sequence of coefficients* $\{c_i\}$, *we must have* $\sum_{i=1}^{k} c_i g_i$ *lying in the cone formed by the vectors* $\{g_k, (g_k + g_{k-1}), (g_k + g_{k-1} + g_{k-2}), \ldots, (g_k + \cdots + g_2) = -g_1\}$. *Moreover, any vector in the cone may be obtained by an appropriate choice of the* c_i.

Proof. To see that any vector in the cone may be realized by some choice of c_i, let λ_k be the coefficient of the vector g_k, let λ_{k-1} be the coefficient of the vector $(g_k + g_{k-1})$, let λ_{k-2} be the coefficient of the vector $(g_k + g_{k-1} + g_{k-2})$, ... We then choose $c_0 = 0$ and $c_j = \sum_{i=1}^{j} \lambda_i$.

To see that any sum with increasing c_i must lie in the cone, note that taking $\lambda_j = c_j - c_{j-1}$ will give the appropriate coefficients to create a positive linear combination of the vectors $\{g_k, (g_k + g_{k-1}), (g_k + g_{k-1} + g_{k-2}), \ldots, (g_k + \cdots + g_2) = -g_1\}$ which is equal to $\sum_{i=1}^{k} c_i g_i$. □

Note that the magnitudes and directions of the vectors are fixed in Lemma 4.1. However, the lemma allows us to find the PS cone for a set of vectors (in fixed directions) that corresponds to a cofacet (i.e. a minimal positively dependent set of vectors in the Gale transform). This follows since requiring $\sum_{i=1}^{k} g_i = 0$ fixes the relative magnitudes of the g_i whenever the set $\{g_i\}$ is a cofacet. However, for non-minimal sets of vectors, continuously varying combinations of the magnitudes can yield a zero sum. This makes the determination of the PS cone difficult for larger sets of vectors, even under the assumption of fixed directions. In general, the PS cone of a larger set is *not* the union of the PS cones of the minimal sets it contains — the need for an increasing sequence of coefficients does not allow us to consider these minimal sets independently. (The example of a triangular prism in the figure above shows that PS cones may not be considered independently. In that case, the cone of the union of the PS cones is the entire plane.)

In higher dimensions we can define a target cone using vectors such as $\{g_1, (g_1 + g_2), \ldots\}$ and $\{-g_n, -(g_n + g_{n-1}), \ldots\}$.

For any polytope, a sufficient condition for a vertex sequence to be affinely inducible is for the PS cone of a spanning set of directions to intersect the target cone. However, a set of vectors is spanning if and only if it corresponds to the coface for a simplicial face.

Moreover, for d-polytopes with more than $d+3$ vertices, fixing the vector directions places undue restrictions on which Gale configurations might represent a combinatorial class. For d-polytopes with more than $d+3$ vertices there is no analogue to the choice of "canonical directions" in the plane. Simply determining which Gale configurations can correspond to a particular combinatorial class is a difficult question. (Note: performing a projective transformation on a polytope simply rescales the vectors in its Gale transform without changing their directions. Hence the class of all vertex sequences possible under the restriction of fixed vector directions includes all sequences that might become possible after applying a projective transformation to the original polytope. [Grü67])

5. Conclusions

Since the Gale transform preserves the entire space of affinely induced value vectors, it is an invaluable tool for investigating which vertex sequences are affinely inducible. The characterizations for d-polytopes with either $d+2$ or $d+3$ vertices are fairly direct. While it is unlikely that similar characterizations will emerge for cases involving d-polytopes with $d+4$

or more vertices, many of the simple vector calculations used in the above characterizations may still provide useful information in cases involving greater numbers of vertices.

References

[Grü67] B. Grünbaum, *Convex Polytopes*, Interscience/Wiley, London, 1967.

[IT00] T. Ishizeki and F. Takeuchi, *Geometric Shellings of 3-Polytopes*, preprint (2000).

[MK00] J. Mihalisin and V. Klee, *Convex and linear orientations of polytopal graphs*, Discrete Comput. Geom. **24** (2000).

[Zie94] G.M. Ziegler, *Lectures on Polytopes*, Springer-Verlag, New York, 1994.

Department of Mathematics, University of Washington, Box 354350, Seattle, Washington 98195-4350, U.S.A.

E-mail address: `mihalisi@math.washington.edu`