Distinguishing Cube Orientations Arising from Linear Programs

Walter D. Morris, Jr.*

September 17, 2002

Abstract

An acyclic unique-sink orientation of the 4-dimensional cube is given that satisfies the the Holt-Klee condition but does not come from a linear program. A conjecture of Stickney and Watson is discussed, and a geometric view of the realizability problem is presented.

1 Introduction

An orientation of the graph of the n-cube is called an LP-cube if it is isomorphic to the orientation of the graph of a polytope P, combinatorially equivalent to the n-cube, that orients the edges of P in the direction of increase of a linear function. The classic paper by Klee and Minty [9] started the study of the performance of linear programming algorithms on polytopes combinatorially equivalent to the n-cube. An understanding of the combinatorial structure of these directed graphs would help greatly in improving pivoting algorithms for linear programming. Unfortunately, the few known necessary conditions for orientations of the cube to be LP-cubes do not give much insight into why the simplex method generally performs better on LP-cubes than on other more loosely defined classes of orientations. See [4], [6], [11] for more on this motivating problem.

The main nontrivial necessary condition for LP-cubes is the *Holt-Klee condition:* [8] Every d-dimensional face has d edge-disjoint directed paths from its source to its sink. This condition, together with some trivial conditions, is known to be sufficient [13] for an orientation of the graph of a 3-polytope to be obtainable from a linear program. In Section 2 of this paper, we give an example of an orientation of the 4-cube that satisfies the Holt-Klee condition but is not an LP-cube.

The orientation that we produce is shown to also fail to have the P-cube property, a property weaker than that of being an LP-cube. A conjectured necessary condition for P-cubes (and hence also necessary for LP-cubes) is investigated in Section 3 and shown to be necessary in dimension 4. Geometric considerations that should be useful for studying questions such as those studied in Sections 3 and 4 are developed in Section 5.

*Department of Mathematical Sciences, George Mason University, Fairfax, Virginia, 22030

2 LP-cubes and P-cubes

The n-dimensional cube will be considered as a graph with vertex set $\{0,1\}^n$, and with two vertices connected by an edge if they differ in exactly one component.

DEFINITION **2.1** *An orientation Q of the edges of the n-cube is said to be an LP-cube if there exists a polytope P combinatorially equivalent to the standard n-cube $\{x \in I\!R^n : 0 \leq x_i \leq 1\}$, a linear objective function f on P that is not constant on any edge of P, and an adjacency-preserving bijection ϕ from $\{0,1\}^n$ to the vertex set of P so that (u,v) is an arc of Q only if $f(\phi(u)) < f(\phi(v))$.*

A linear program for which the feasible region is combinatorially equivalent to the n-cube may, without loss of generality, be written as: Maximize $f(x) = c^T x$, subject to $x \in P = \{x \in I\!R^{2n} : Cx = d, x \geq 0\}$, where $c \in I\!R^{2n}$, $d \in I\!R^n$, $C \in I\!R^{n \times 2n}$. We may assume that, for $i = 1, 2, \ldots, n$, the equalities $x_i = 0$ and $x_{n+i} = 0$ define opposite facets of the n-cube. We will assume that the objective function $c^T x$ is not constant on any edge of the cube.

A vector $u \in \{0,1\}^n$ determines a vertex $\phi(u)$ of P by the rule $\phi(u)_i = 0$ if $u_i = 1$ and $\phi(u)_{n+i} = 0$ if $u_i = 0$. A vector $u \in \{0,1\}^n$ also determines a *complementary* submatrix A_u of C by setting the i^{th} column $(A_u)_i$ equal to C_i if $u_i = 1$ and equal to C_{n+i} if $u_i = 0$.

The relative values of f on adjacent vertices of P can be read from the corresponding simplex tableaux. We augment C to create the $(n+1) \times (2n+2)$ matrix

$$\hat{C} = \begin{pmatrix} C & 0 & d \\ c^T & 1 & 0 \end{pmatrix}.$$

Each A_u is augmented to create an $(n+1) \times (n+1)$ matrix $\hat{A}_u$ by appending to the bottom of each column of A_u the corresponding entry of c^T and appending column $2n+1$ of $\hat{C}$. The tableau T_u corresponding to u is then $T_u = \hat{A}_u^{-1} \hat{C}$. If $u \in \{0,1\}^n$ and v is obtained from u by changing the i^{th} component, then we will have $f(\phi(v)) < f(\phi(u))$ if and only if the nonzero entry of T_u that is in row $n+1$ and in column i or $n+i$, is greater than 0.

In the particular case $u = \mathbf{0} = (0, 0, \ldots, 0)$, the tableau $T_{\mathbf{0}}$ has the form

$$T_{\mathbf{0}} = \begin{pmatrix} M^T & I & 0 & p \\ q^T & 0 & 1 & k \end{pmatrix}$$

for some $n \times n$ matrix M, $n \times 1$ vectors q, p and scalar k.

It is clear that the following conditions are satisfied by every LP-cube Q:

(1) Q is acyclic.

(2) Every face of Q has a unique sink.

A third necessary condition is the *Holt-Klee condition:*

(3) For every d-dimensional face F of Q there are d edge-disjoint directed paths from the unique source of F to its unique sink.

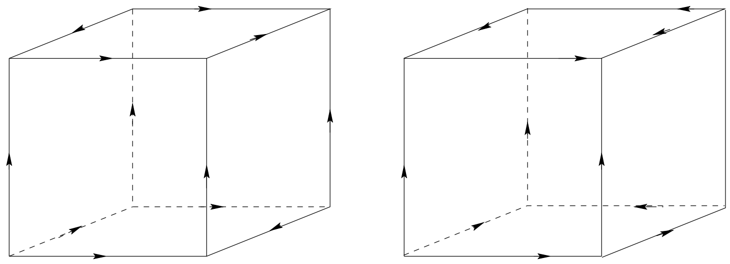

Figure 1: Two cubes that are not LP-cubes

Condition (3) is a directed version of Balinski's theorem [1], that the graph of a d-polytope is d-connected. Conditions (1) - (3) are not only necessary for LP-cubes, but for every orientation of the graph of a polytope that is determined by the direction of increase of a linear function. Mihalisin and Klee [13] proved that conditions (1) - (3) are also *sufficient* for an orientation of the graph of a 3-dimensional polytope to be of this type. For the case of 3-dimensional cubes, this implies that all orientations satisfying (1) and (2) are LP-cubes except the ones in Figure 1.

These orientations were discovered by Stickney and Watson [21] and independently by Gärtner and Kaibel [5]. The first of the two orientations was also found, in a dual form, by Smilansky [20] as a shelling of the octahedron that is not a straight-line shelling.

We will also be interested in a weakening of the concept of an LP-cube. Note that if we have a tableau $T_{\mathbf{0}}$ as above realizing an LP-cube, then it follows from Cramer's rule that M must be a *P-matrix*, that is, all principal minors of M must be positive. For a P-matrix M and a vector q, let us define the tableau

$$R_{\mathbf{0}} = \left(\begin{array}{ccc} M^T & I & 0 \\ q^T & 0 & 1 \end{array} \right).$$

For each vertex u of the cube we can define the submatrix $\hat{A}_u$ of $R_{\mathbf{0}}$ as above, and define $R_u = \hat{A}_u^{-1} R_{\mathbf{0}}$. Assume that the vector q is in sufficiently general position relative to the rows of M^T and I so that the bottom row of every tableau R_u has $n+1$ nonzero entries. If $u \in \{0,1\}^n$ and v is obtained from u by changing the i^{th} component, then we can orient the edge connecting them from v to u if and only if the nonzero entry of R_u that is in row $n+1$ and in column i or $n+i$, is greater than 0. It follows from [18] and [21] that this orientation is well defined and satisfies (2).

An orientation of the cube obtained from a P-matrix M and a vector q will be called a P-cube. Note that every LP-cube is a P-cube. Given a P-matrix M and a vector q, the problem of finding the sink of the associated P-cube is a linear complementarity problem. See [2], [16] for more on the linear complementarity problem. The complexity of the linear complementarity problem with a P-matrix is not known, see [12],[14]. A P-cube may be thought of as coming from a "linear program with no right hand side." The 2-dimensional P-cubes are exactly the LP-cubes. The 3-dimensional *acyclic* P-cubes are exactly the LP-cubes, as shown by the enumerations of [21] and [5]. No example of an acyclic P-cube that is not an LP-cube has yet been found. The problem of recognizing P-matrices that

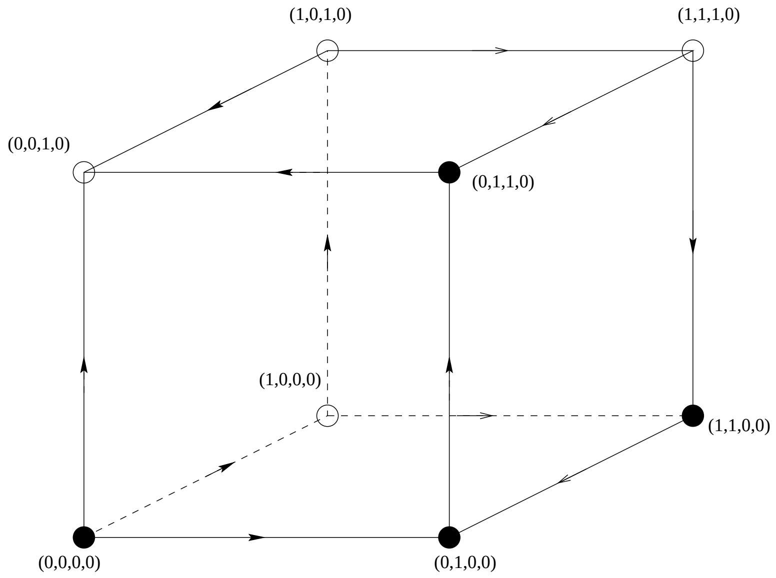

Figure 2: A face of a 4-cube

can be used to define LP-cubes is studied in [10], [15], [17]. To determine if all P-cubes satisfy condition (3) is an interesting open problem.

The problem of finding the sink in orientations of the cube that satisfy (2) has received attention recently in [19], [22].

3 Nonsufficiency of the Holt-Klee condition for P-cubes

Figure 2 gives a face of a 4-cube that satisfies conditions (1) - (3), but is not a P-cube and hence not an LP-cube.

The figure gives the orientation of the edges between vertices that have a 0 in the last coordinate. One can copy the figure and change the fourth coordinates from 0 to 1 to get the orientation of the edges between vertices that have a 1 in the last coordinate. Finally, an edge connecting the pictured face to the opposite face will leave the vertex with coordinate 0 if the vertex is solid and enter it if the vertex is clear. The resulting orientation satisfies (1) and (2). To see that it is acyclic, note that the orientation is consistent with the following ordering of the vertices, in that if vertices v and u are adjacent in the cube and v precedes u in the ordering, then the cube edge between them is oriented from v to u.
$((0,0,0,0),(0,0,0,1),(1,0,0,1),(1,0,0,0),(1,0,1,1),(1,1,1,1),(1,0,1,0),(1,1,1,0),$
$(1,1,0,0),(1,1,0,1),(0,1,0,0),(0,1,0,1),(0,1,1,0),(0,1,1,1),(0,0,1,1),(0,0,1,0))$

It is also easy to check that condition (2) is satisfied. One can check that condition (3) is satisfied for all 3-dimensional faces by noting that for each of the 3 dimensional faces there is a direction for which three of the edges are pointing one way and one in the other way. A look at Figure 1 shows that this is sufficient to exclude the examples of that Figure. To show that the cube itself satisfies condition (3), we display four edge-disjoint directed paths from $(0,0,0,0)$ to $(0,0,1,0)$. They are

1. $(0,0,0,0) \to (0,0,1,0)$
2. $(0,0,0,0) \to (1,0,0,0) \to (1,0,1,0) \to (0,0,1,0)$
3. $(0,0,0,0) \to (0,1,0,0) \to (0,1,1,0) \to (0,0,1,0)$
4. $(0,0,0,0) \to (0,0,0,1) \to (0,0,1,1) \to (0,0,1,0)$

THEOREM **3.1** *The cube defined above is not a P-cube*

Proof: Suppose that the above cube were a P-cube. Then we would have a 4×4 matrix M^T and a vector q determining the orientation. The orientation would be determined by sign patterns of vectors in the row space of the matrix

$$R_{\mathbf{0}} = \begin{pmatrix} M^T & I & 0 \\ q^T & 0 & 1 \end{pmatrix}.$$

Let V be the affine space consisting of those vectors in the row space of $R_{\mathbf{0}}$ that have a 1 in the last coordinate. This space is closed under convex combinations.

For each vertex u of the cube, define $x_u \in V$ to be the bottom row of tableau R_u. The sign pattern of each x_u is given by the cube orientation. We first consider the vectors $x_{(0,0,0,0)}$ and $x_{(0,1,0,0)}$. The sign patterns of these vectors must be $(-,-,-,-,0,0,0,0,+)$ and $(+,0,-,-,0,+,0,0,+)$. Because the vectors have opposite signs in the first coordinate, there must be a unique convex combination y of $x_{(0,0,0,0)}$ and $x_{(0,1,0,0)}$ that is 0 in the first coordinate. The sign pattern of y must be $(0,-,-,-,0,+,0,0,+)$. The sign pattern of the vector $x_{(1,1,0,0)}$ is $(0,0,+,-,-,+,0,0,+)$. There must be a unique convex combination z of $x_{(1,1,0,0)}$ and y in which the third coordinate is 0. The sign pattern of z must be $(0,-,0,-,-,+,0,0,+)$. The sign patterns of $x_{(1,0,1,0)}$ and $x_{(1,1,1,0)}$ are $(0,-,0,+,-,0,+,0,+)$ and $(0,0,0,+,-,+,-,0,+)$. There must be a convex combination w of these two with sign pattern $(0,-,0,+,-,+,0,0,+)$. Note that z and w have the same support, but differ in sign on coordinate 4.

If S is a subset of $\{1,2,\ldots,8\}$ that does not contain $\{i, 4+i\}$ for any $i \in \{1,2,3,4\}$, then the subspace of V consisting of the points with 0 in all of the entries indexed by S is of dimension $4-|S|$, by the P-matrix property of M. In particular, the set of elements of V with entries 1,7, and 8 all equal to 0 is a 1-dimensional subspace W of V. The vectors z and w are distinct vectors in W. Since both z and w have a 0 in entry 3, all elements of W are 0 on the third coordinate. This is contradicted by vector $x_{(1,1,0,0)}$, which is in W but has a nonzero third entry. □

The cube constructed is thus the smallest example of an orientation of a cube that satisfies (1),(2), and (3) but is not a P-cube (and hence not an LP-cube.) It would be interesting to know how many different orientations of the 4-cube satisfy (1), (2), and (3) but are not LP-cubes or P-cubes. Finally, it should be pointed out that the the same contradiction can be derived if one assumes that the given orientation is derived from an "oriented matroid complementarity problem," as studied in [23].

M. Develin [3] has recently proved in a non-constructive way that almost all cube orientations satisfying (1), (2), and (3) are not LP-cubes. There is therefore ample motivation to look for stronger combinatorial conditions that are necessary for LP-cubes.

4 A conjecture of Stickney and Watson

Stickney and Watson [21] made the following conjecture:

Conjecture **4.1** *No P-cube of dimension $n \geq 3$ may have all of its vertices of even in-degree on the same $(n-1)$-dimensional face.*

The requirement that no $(n-1)$-dimensional face of a cube may contain all of the vertices of even in-degree seems too weak for this conjecture to be relevant to the study of linear programming. Nevertheless, the condition characterizes 3-dimensional P-cubes among the orientations satisfying (2). The main difficulty in proving necessity in the 4-dimensional case lies in showing that the example of the previous section is not a P-cube.

Lemma **4.2** *If Q is an n-dimensional P-cube and F and F' are two opposite $(n-1)$-dimensional faces of Q, then the orientation Q' obtained from Q by replacing the orientation of F' by a copy of the orientation of F is also a P-cube.*

Proof: Assume that Q is given by

$$R_{\mathbf{0}} = \begin{pmatrix} M^T & I & 0 \\ q^T & 0 & 1 \end{pmatrix}.$$

We can assume that F consists of the vertices u with $u_n = 0$ and F' consists of vertices u with $u_n = 1$. Continuing with the notation of the previous section, this implies that coordinate n of x_u is nonzero for $u \in F$ and coordinate $2n$ of x_u is nonzero for $u \in F'$. Note that if we replace column n of M by column n of I, the resulting matrix M' will still be a P-matrix. Moreover, if u and v differ on only coordinate n, then x_u and x_v will only differ on coordinates n and $2n$. Thus we will have a new P-cube for which the orientation of face F' is replaced with a copy of the orientation of F. $\square$

The Lemma remains true if we replace "P-cube" by "LP-cube." Assuming enough nondegeneracy, one can also replace column n of M by a sequence of convex combinations of column n of M and column n of I so that the edges of F' that differ in orientation from the corresponding edges of F get reversed one at a time, with the P-cube property preserved after each reversal.

Lemma **4.3** *Suppose that the conjecture is true for dimension $n-1$. Suppose that Q is an orientation of the n-cube, and F is the face of Q containing all vertices on which coordinate n is 0. Suppose that all vertices of F have even indegree in Q. If there is a coordinate $i, 1 \leq i \leq n-1$, so that all edges of F connecting vertices that differ on coordinate i are either all going from the $(n-2)$-dimensional face $E \subset F$ for which $u_i = 0$ or all going toward E, then Q is not a P-cube.*

Proof: The facet of Q for which $u_i = 0$, or the opposite facet, will be an $(n-1)$-dimensional cube with all of its vertices of even indegree in an $(n-2)$-dimensional face. Since we are assuming that the conjecture is true in dimension $(n-1)$, this facet will not be a P-cube. Therefore, Q will not be a P-cube. □

PROPOSITION **4.4** *The conjecture is true in dimension 4.*

Proof: Suppose that Q is a 4-dimensional P-cube, and that all of the vertices of Q of even indegree are in the 3-dimensional face F for which $u_4 = 0$. Lemma 4.3 states that F may not have any direction in which all the edges point the same way. From the enumeration of [21], it follows that F is either the orientation of Figure 1 or a cyclic 3-cube obtained by reversing all of the vertical edges in Figure 1. It was shown in [21] that the class of P-cubes is closed under the operation of reversing all of the edges of a given direction. If Q is to have all of the vertices of even indegree in F, and F is as in Figure 1, then the edges leaving F must be exactly those connected to the solid vertices. Since we can assume that the orientation of the face opposite to F is identical to F, we have the orientation of the previous section, which we know is not a P-cube. □

5 The set of edges leaving a facet

A common theme of the previous two sections was an attempt to understand what kind of set of vertices of a facet F of a P-cube can be the set of vertices of F that are on edges directed away from F. In the example of section 2, vertices $(0,0,0,0)$ and $(0,1,0,0)$ are on edges directed away from the pictured face, and vertices $(1,0,1,0)$ and $(1,1,1,0)$ are on edges directed toward the face. The contradiction obtained derived from an inability to assign a sign to the fourth coordinate of a vector that was in some sense "between" the segments $[x_{(0,0,0,0)}, x_{(0,1,0,0)}]$ and $[x_{(1,0,1,0)}, x_{(1,1,1,0)}]$. Stickney and Watson's conjecture can be restated as follows: Let F be an $(n-1)$-dimensional face of a P-cube Q, where $n \geq 3$. Then the set of vertices of F on edges leaving F may not be the same as the set of vertices of even indegree within F.

Suppose that we have an $(n-1)$-dimensional P-cube F and a set X of vertices of F. We would like to know if there exists an n-dimensional P-cube Q that has F as a face, so that X is the set of vertices of F on edges leaving F. If there is such a Q, we may assume that the face opposite to F is identical to F, and that the orientation of Q is given by a matrix of the form

$$\hat{R}_{\mathbf{0}} = \begin{pmatrix} M^T & b & I & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ q^T & r & 0 & 0 & 1 \end{pmatrix}.$$

Here the orientation of the face F is determined by the submatrix

$$R_{\mathbf{0}} = \begin{pmatrix} M^T & I & 0 \\ q^T & 0 & 1 \end{pmatrix},$$

obtained from $\hat{R}_{\mathbf{0}}$ by deleting row n, column n and column $2n$.

As in section 2, let V be the set of vectors in the row space of $R_{\mathbf{0}}$ with last coordinate 1, and let the vectors x_u in V correspond to vertices of F. Let $\hat{V}$ be the set of vectors in the row space of $\hat{R}_{\mathbf{0}}$ that have 1 in the last coordinate. The subspace of $\hat{V}$ consisting of vectors in which coordinates n and $2n$ are 0 can be identified with a hyperplane through V. A subset X of vertices of F will be the set of vertices on edges leaving F if and only if this hyperplane separates $\{x_u : u \in X\}$ from $\{x_u : u \notin X\}$.

PROPOSITION **5.1** *If Q is and n-dimensional P-cube, for $n \geq 3$, and F is a facet of Q, then the set of vertices of F that are on edges leaving F may not be the set consisting of the source and the sink of F.*

Proof: Suppose that we have a matrix $\hat{R}_{\mathbf{0}}$ determining Q. Let $\hat{V}$ be the set of vectors in the row space of $\hat{R}_{\mathbf{0}}$ that have 1 in the last coordinate. Assume that F is the facet of Q for which $u_n = 0$. For each $u \in Q$ let $\hat{x}_u$ be the last row of the tableau $\hat{R}_{\mathbf{0}}$. We can assume that the source of F is at $\mathbf{0} = (0, 0, \ldots, 0)$ and that the sink v of F has $v_i = 1, i = 1, 2, \ldots, k$ and $v_i = 0, i = k+1, k+2, \ldots, n$ for some $1 \leq k \leq n-1$. We will assume that there are edges leaving F from $\mathbf{0}$ and v.

Suppose first that $k = 1$. The vector $\hat{x}_{\mathbf{0}}$ has a negative second coordinate, and the vector $\hat{x}_v$ has a positive second coordinate. Let z be the convex combination of $\hat{x}_{\mathbf{0}}$ and $\hat{x}_v$ that has second coordinate 0. By the argument used to prove Theorem 3.1, z is the unique vector of V with support indexed by the set $\{1\} \cup \{3, 4, \ldots, n+1\} \cup \{2n+1\}$. Because both $\hat{x}_{\mathbf{0}}$ and $\hat{x}_v$ have negative nth coordinates, so does z. Let G be the 2-dimensional face of F consisting of vertices u with $u_3 = u_4 = \cdots = u_n = 0$. The vectors $\hat{x}_{(0,1,0,0,\ldots,0)}$ and $\hat{x}_{(1,1,0,0,\ldots,0)}$ must differ in sign on coordinate $2n+2$, because neither of them is the sink nor the source of G. The convex combination of $\hat{x}_{(0,1,0,0,\ldots,0)}$ and $\hat{x}_{(1,1,0,0,\ldots,0)}$ that is 0 on coordinate $2n+2$ must be z, because this combination has support in the set $\{1\} \cup \{3, 4, \ldots, n+1\} \cup \{2n+1\}$. Because z has negative nth coordinate, so does at least one of $\hat{x}_{(0,1,0,0,\ldots,0)}$ and $\hat{x}_{(1,1,0,0,\ldots,0)}$. Therefore there is an edge leaving F from $(0, 1, 0, 0, \ldots, 0)$ or $(1, 1, 0, 0, \ldots, 0)$.

Next, suppose that $k \geq 2$. Let P be the polyhedron in $\hat{V}$ consisting of those vectors with the first k entries nonpositive and entries $n+1$ through $n+k$ nonnegative, and entries $n+k+1, n+k+2, \ldots, 2n$ equal to 0. As in the simplex method, we can move from $\hat{x}_{\mathbf{0}}$ in P by allowing entry $n+1$ to increase, while keeping entries $n+2$ through $2n-1$ equal to 0. Let z be the first vector we meet that has a 0 among the entries indexed by $\{1, 2, \ldots, k\} \cup \{n\}$. The P-matrix property assures that there will be such a z, see for example [23]. Let a nonzero entry of $\hat{x}_{\mathbf{0}}$ that has been turned to 0 in z be in coordinate $i \in \{1, 2, \ldots, k\} \cup \{n\}$.

If $i = 1$ then $z = \hat{x}_{(1,0,0,\ldots,0)}$, and then $(1, 0, 0, \ldots, 0)$ will be a vertex of F other than v on an edge leaving F, because entry n of $\hat{x}_{(1,0,0,\ldots,0)}$ will be negative. If $2 \leq i \leq k$, then we can assume $i = 2$. If the vector $\hat{x}_{(1,0,0,\ldots,0)}$ has a negative second coordinate, then the convex combination y of $x_{\mathbf{0}}$ and $-\hat{x}_{(1,0,0,\ldots,0)}$ that has 0 in the second coordinate will agree in sign on entries 1 and $n+1$. But y must be a multiple of z, where entries 1 and $n+1$ disagree in sign. Thus the second entry of $\hat{x}_{(1,0,0,\ldots,0)}$ is positive and $(1, 0, 0, \ldots, 0)$ is the sink of the 2-face G. As in the argument for the $k = 1$ case, it follows that there is an edge leaving F from $(0, 1, 0, 0, \ldots, 0)$ or $(1, 1, 0, 0, \ldots, 0)$. Note that $(1, 1, 0, 0, \ldots, 0)$ is not v, because it is not the sink of G.

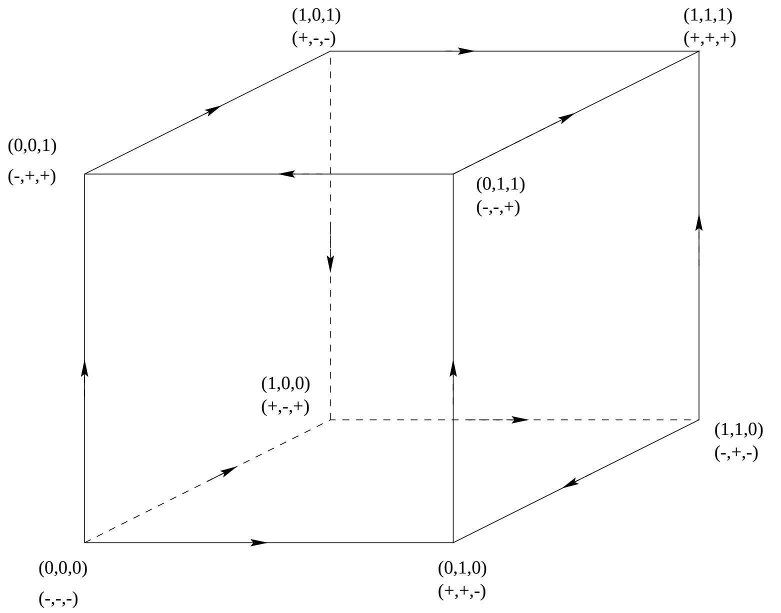

Figure 3: A P-cube with a cycle

There remains the case $i = n$. For $j = 1, \ldots, k$ we can define a vector z^j to be the vector obtained as z was above by increasing entry $n+j$ of $\hat{x}_{\mathbf{0}}$ and holding the other entries in $\{n+1, n+2, \ldots, 2n\}$ at 0. If each of the z^j has a 0 in entry n, then the hyperplane of $\hat{V}$ consisting of points with entry n equal to 0 cuts $\hat{x}_{\mathbf{0}}$ off from the rest of the vertices of P. But $\hat{x}_v$ is a vertex of P that has the same sign as $\hat{x}_{\mathbf{0}}$ on coordinate n, so we have a contradiction. □

The condition that the source and sink are not the only vertices of a facet that are on edges leaving it is much weaker than condition (3), but we do not know if (3) holds for all P-cubes. This weaker condition is still strong enough to characterize 3-dimensional P-cubes among orientations satisfying (2).

We include a 3-dimensional example which illustrates the large variety of sets that can be separated by hyperplanes. Suppose we start with the matrix

$$R_{\mathbf{0}} = \begin{pmatrix} 1 & 0 & 2 & 1 & 0 & 0 & 0 \\ 2 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 2 & 1 & 0 & 0 & 1 & 0 \\ -1 & -1 & -1 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

The P-cube F determined by this matrix is in Figure 3. Next to each vertex u we have displayed the vector σ_u in which coordinate i is $-$ or $+$ depending on whether the edge connecting u to a vertex differing from u in coordinate i is directed away from u or toward u.

The space V is 3-dimensional, and we project it to $\mathbb{R}^3$ by ignoring the last 4 coordinates. The image of x_u under this projection is denoted $\bar{x}_u$, for each u in the cube. Figure 4 shows the $\bar{x}_u$ under this identification.

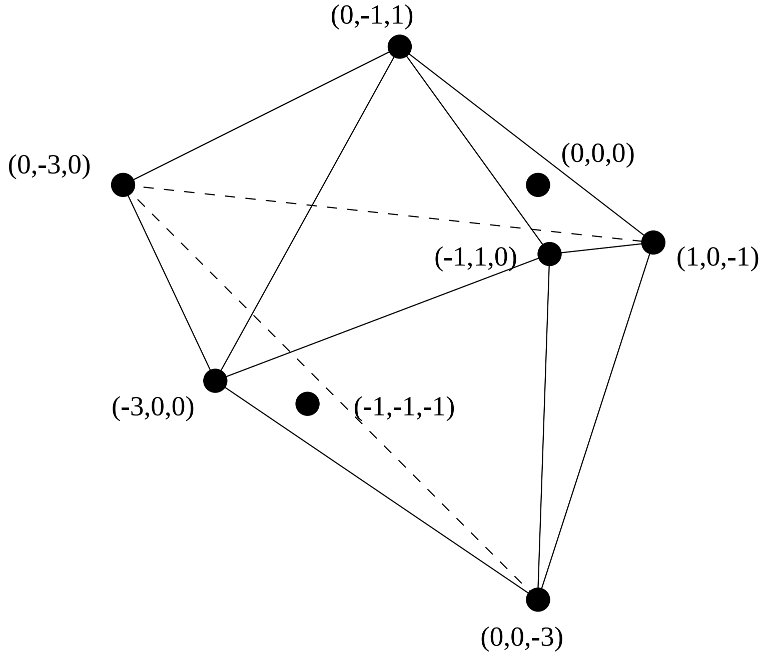

Figure 4: Convex hull of the $\bar{x}_u$

The convex hull of the $\bar{x}_u$ is an octahedron. Two of the $\bar{x}_u$ are in the interiors of 2-dimensional faces of the convex hull. Given any face G of this octahedron, there is a hyperplane separating the $\bar{x}_u$ on G from the rest of the points. One can therefore extend F to a 4-dimensional P-cube by letting the edges leaving F be the ones on vertices u corresponding to $\bar{x}_u$ on G. For example, all of the $\bar{x}_u$ for u that have indegree at most 1 are on the same face. Note also that any open halfspace containing two $\bar{x}_u$ for which the σ_u vectors differ on all three coordinates must contain some other points of the configuration, as follows from Proposition 5.1. On the other hand, the edge from $(-3, 0, 0)$ to $(0, 0, -3)$ contains two points for which the corresponding σ_u differ on two coordinates.

In the "trivial" case where $M = I$ and q is the vector of all ones, it is easy to see that the points $\bar{x}_u$ will be the vertices of the standard cube $[0, 1]^n$. The problem of recognizing sets of vertices of this cube that can be separated by a hyperplane is hardly well understood, see [7]. In that case, however, some strong combinatorial conditions are known.

References

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

[2] R. W. Cottle, J.-S. Pang and R. E. Stone, *The Linear Complementarity Problem*, Academic Press, Boston, 1992.

[3] M. Develin, *On LP-orientations of cubes and crosspolytopes*, (2002) preprint, Dept. of Mathematics, UC Berkeley.

[4] B. Gärtner, *Combinatorial linear programming: geometry can help*, Lecture Notes in Computer Science **1518**, (1998) Springer-Verlag, 82 - 96.

[5] B. Gärtner and V. Kaibel, *Abstract objective function graphs on the 3-cube - a classification by realizability* Tech. Report 296, (1998) Dept. of Computer Science, ETH Zürich.

[6] B. Gärtner, M. Henk and G. Ziegler, *Randomized simplex algorithms on Klee-Minty cubes*, Combinatorica **18** (1998)

[7] K. Hoke, *Extending shelling orders and a hierarchy of functions of unimodal simple polytopes*, Disc. Appl. Math. **60** (1995), 211 - 217.

[8] F. Holt and V. Klee, *A proof of the strict monotone 4-step conjecture*, In J. Chazelle, J. B. Goodman and R. Pollack, eds., *Advances in Discrete and Computational Geometry*, Contemporary Mathematics, American Mathematical Society (1998).

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

[10] O. L. Mangasarian, *Linear complementarity problems solvable by a single linear program*, Mathematical Programming **10** (1976) 263 - 270.

[11] J. Matoušek, *Lower bounds for a subexponential optimization algorithm*, Random Structures and Algorithms **5** (1994) 591 - 607.

[12] N. Megiddo, *A note on the complexity of P-matrix LCP and computing an equilibrium*, Research Report (1988), IBM Almaden Research Center, San Jose.

[13] J. Mihalisin and V. Klee, *Convex and linear orientations of polytopal graphs*, Discrete Comput. Geom. **24** (2000), 421–435.

[14] W. D. Morris, Jr., *Randomized pivot algorithms for P-matrix linear complementarity problems*, Math. Program. Ser. A **92** (2002) 285 - 296.

[15] W. D. Morris, Jr. and J. Lawrence, *Geometric properties of hidden Minkowski matrices*, SIAM Journal on Matrix Analysis and its Applications **10** (1988) 229 - 232.

[16] K. G. Murty, *Linear Programming, Linear and Nonlinear Programming*, Heldermann Verlag, Berlin, 1988.

[17] J.-S. Pang and R. Chandrasekaran, *Linear complementarity problems solvable by a polynomially bounded pivoting algorithm*, Mathematical Programming Study **25** (1985), 13 - 27.

[18] H. Samelson, R. M. Thrall and O. Wesler, *A partition theorem for Euclidean n-space*, Proceedings of the American Mathematical Society, **9** (1958), 805 - 807.

[19] I. Schurr, T. Szabó, *Finding the sink takes some time*, Proc. 10th European Symposium on Algorithms (ESA) (2002), to appear.

[20] Z. Smilansky, *A non-geometric shelling of a 3-poytope*, Israel J. Math., **71** (1990), 29-32.

[21] A. Stickney and L. T. Watson, *Digraph models of Bard-type algorithms for the linear complementarity problem*, Mathematics of Operations Research **3** (1978), 322-333.

[22] T. Szabó and E. Welzl, *Unique sink orientations of cubes*, Proc. 42nd Ann. IEEE Symposium of Foundations of Computer Science (FOCS) (2001), 547 - 555.

[23] M. J. Todd, *Complementarity in oriented matroids*, SIAM J. Algebraic and Discrete Methods **5** (1984), 467 - 487.