Discrete Applied Mathematics 163 (2014) 155–164

Contents lists available at ScienceDirect

Discrete Applied Mathematics

journal homepage: www.elsevier.com/locate/dam

Counting unique-sink orientations

Jan Foniok [a,*], Bernd Gärtner [b], Lorenz Klaus [a], Markus Sprecher [b]

[a] Institute for Operations Research, ETH Zurich, 8092 Zurich, Switzerland

[b] Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland

ARTICLE INFO

Article history:
Received 6 December 2010
Received in revised form 5 July 2013
Accepted 22 July 2013
Available online 21 August 2013

Keywords:
Unique-sink orientation
Linear complementarity problem
Pivoting
P-matrix
K-matrix

ABSTRACT

Unique-sink orientations (USOs) are an abstract class of orientations of the n-cube graph. We consider some classes of USOs that are of interest in connection with the linear complementarity problem. We summarize old and show new lower and upper bounds on the sizes of some such classes. Furthermore, we provide a characterization of K-matrices in terms of their corresponding USOs.

1. Introduction

Unique-sink orientations (USOs) are an abstract class of orientations of the n-cube graph. A number of concrete geometric optimization problems can be shown to have the combinatorial structure of a USO. Examples are the linear programming problem [11], and the problem of finding the smallest enclosing ball of a set of points [11,30], or a set of balls [7]. In this paper, we count the USOs of the n-cube that are generated by P-matrix linear complementarity problems (P-USOs). This class covers many of the "geometric" USOs. We show that the number of P-USOs is $2^{\Theta(n^3)}$. The lower bound construction is the interesting contribution here, and it even yields USOs from the subclass of K-USOs, whose combinatorial structure is known to be very rigid [8]. In contrast, the number of all n-cube USOs is doubly exponential in n [17].

1.1. Unique-sink orientations

We follow the notation of [8]. Let $[n] := \{1, 2, \dots, n\}$. For a bit vector $v \in \{0, 1\}^n$ and $I \subseteq [n]$, let $v \oplus I$ be the element of $\{0, 1\}^n$ defined by

$$(v \oplus I)_j := \begin{cases} 1 - v_j & \text{if } j \in I, \\ v_j & \text{if } j \notin I. \end{cases}$$

Instead of $v \oplus \{i\}$ we write $v \oplus i$.

Under this notation, the (undirected) *n-cube* is the graph $G = (V, E)$ with

$$V := \{0, 1\}^n, \qquad E := \{\{v, v \oplus i\} : v \in V,\ i \in [n]\}.$$

* Correspondence to: Department of Mathematics and Statistics, Queen's University, Jeffery Hall, 48 University Avenue, Kingston ON K7L 3N6, Canada. Tel.: +1 6135332418.

E-mail addresses: jan@foniok.ch, foniok@mast.queensu.ca (J. Foniok), gaertner@inf.ethz.ch (B. Gärtner), lklaus@ifor.math.ethz.ch (L. Klaus), markussp@student.ethz.ch (M. Sprecher).

0166-218X/$ – see front matter
http://dx.doi.org/10.1016/j.dam.2013.07.017

A *subcube* of G is a subgraph $G' = (V', E')$ of G where $V' = \{v \oplus I : I \subseteq C\}$ for some vertex v and some set $C \subseteq [n]$, and $E' = E \cap \binom{V'}{2}$. The *dimension* of such a subcube is $|C|$.

Let ϕ be an orientation of the n-cube (a digraph with underlying undirected graph G). If ϕ contains the directed edge $(v, v \oplus i)$, we write $v \xrightarrow{\phi} v \oplus i$, or simply $v \to v \oplus i$ if ϕ is clear from the context. If V' is the vertex set of a subcube, then the directed subgraph of ϕ induced by V' is denoted by $\phi[V']$. For $F \subseteq [n]$, let $\phi^{(F)}$ be the orientation of the n-cube obtained by reversing all edges in coordinates contained in F; formally

$$v \xrightarrow{\phi^{(F)}} v \oplus i :\Leftrightarrow \begin{cases} v \xrightarrow{\phi} v \oplus i & \text{if } i \notin F, \\ v \oplus i \xrightarrow{\phi} v & \text{if } i \in F. \end{cases}$$

An orientation ϕ of the n-cube is a *unique-sink orientation* (*USO*) if every subcube $G' = (V', E')$ has a unique sink (that is, vertex of outdegree zero) in $\phi[V']$. It is not difficult to show that in a unique-sink orientation, every subcube also has a unique source (that is, vertex of indegree zero). More generally, if ϕ is a unique-sink orientation and $F \subseteq [n]$, then $\phi^{(F)}$ is a unique-sink orientation as well [30, Lemma 2.1].

A special USO is the *uniform orientation*, in which $v \to v \oplus i$ if and only if $v_i = 0$.

Unique-sink orientations enable a graph-theoretic description of simple principal pivoting algorithms for linear complementarity problems. They were introduced by Stickney and Watson [29] and have recently received much attention [10–12,17,21,27,28,30].

1.2. Linear complementarity problems

A *linear complementarity problem* (LCP (M, q)) is for a given matrix $M \in \mathbb{R}^{n\times n}$ and a vector $q \in \mathbb{R}^n$, to find vectors $w, z \in \mathbb{R}^n$ such that

$$w - Mz = q, \quad w, z \geq 0, \qquad w^T z = 0. \tag{1}$$

A P-*matrix* is a square real matrix whose principal minors are all positive. If M is a P-matrix, the appertaining LCP is called a P-*LCP*; in this case there exists a unique solution for any q [26].

Let $B \subseteq [n]$, and let A_B be the $n \times n$ matrix whose jth column is the jth column of $-M$ if $j \in B$, and the jth column of the $n \times n$ identity matrix I_n if $j \notin B$. If M is a P-matrix, then A_B is invertible for every set B. We call B a *basis*. If $A_B^{-1} q \geq 0$, let

$$w_i := \begin{cases} 0 & \text{if } i \in B \\ (A_B^{-1} q)_i & \text{if } i \notin B, \end{cases} \qquad z_i := \begin{cases} (A_B^{-1} q)_i & \text{if } i \in B \\ 0 & \text{if } i \notin B. \end{cases} \tag{2}$$

The vectors w, z are then a solution to the LCP (1).

A problem P-LCP (M, q) is *nondegenerate* if $(A_B^{-1} q)_i \neq 0$ for all B and i. Following [29], a nondegenerate P-LCP (M, q) induces a USO: for $v \in \{0, 1\}^n$, let $B(v) := \{j \in [n] : v_j = 1\}$. Then the unique-sink orientation ϕ induced by P-LCP (M, q) is given by

$$v \xrightarrow{\phi} v \oplus i :\Leftrightarrow (A_{B(v)}^{-1} q)_i < 0. \tag{3}$$

The run of a simple principal pivoting method (see [23, Chapter 4]) for the P-LCP then corresponds to following a directed path in the orientation ϕ. Finding the sink of the orientation is equivalent to finding a basis B with $A_B^{-1} q \geq 0$, and thus via (2) to finding the solution to the P-LCP.

In this paper, we are primarily interested in establishing bounds for the number of n-dimensional USOs satisfying some additional properties (for instance, USOs induced by P-LCPs), which we introduce in the next section.

2. Matrix classes and USO classes

It is NP-complete to decide whether a solution to an LCP exists [2]. If the matrix M is a P-matrix, however, a solution always exists. The problem of finding it is unlikely to be NP-hard, because if it were, then NP = co-NP [18]. Even so, no polynomial-time algorithms for solving P-LCPs are known. Hence our motivation to study some special matrix classes and investigate what combinatorial properties their USOs have. The ultimate goal is then to try and exploit these combinatorial properties in order to find an efficient algorithm for the corresponding LCPs.

A Z-*matrix* is a square matrix whose off-diagonal entries are all non-positive. A K-*matrix* is a matrix which is both a Z-matrix and a P-matrix. A *hidden*-K-*matrix* is a P-matrix M such that there exist Z-matrices X and Y and non-negative vectors r and s with $MX = Y$, $r^T X + s^T Y > 0$. Taking X to be the identity matrix and $Y = M$, $s = 0$ and r any positive vector shows that every K-matrix is a hidden-K-matrix as well.

The importance of these matrix classes is due to the fact that polynomial-time algorithms are known for solving the LCP (M, q) if the matrix M is a Z-matrix [1,25], a hidden-K-matrix [16], or the transpose of a hidden-K-matrix [24].