Operations Research Letters 42 (2014) 156–160

Contents lists available at ScienceDirect

Operations Research Letters

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

A semidefinite approach to the K_i-cover problem

João Gouveia [a], James Pfeiffer [b,*]

[a] CMUC, Department of Mathematics, University of Coimbra, 3001-454 Coimbra, Portugal

[b] Department of Mathematics, University of Washington, Seattle, WA 98195, United States

ARTICLE INFO

Article history:
Received 23 March 2013
Received in revised form
20 January 2014
Accepted 20 January 2014
Available online 30 January 2014

Keywords:
Sums of squares relaxations
Theta bodies
Minimal K_i-cover problem
Maximal K_i-free subgraph problem
Tuza's conjecture

ABSTRACT

We apply theta body relaxations to the K_i-cover problem and show polynomial time solvability for certain classes of graphs. In particular, we give an effective relaxation where all K_i-p-hole facets are valid, and study its relation to an open question of Conforti et al. For the triangle free problem, we show for K_n that the theta body relaxations do not converge by $(n-2)/4$ steps; we also prove for all G an integrality gap of 2 for the second theta body.

1. Introduction

A common way to model a combinatorial optimization problem is as the optimization of a function over the set $S \subseteq \{0, 1\}^n$ of characteristic vectors of the objects in question. When the objective function is linear, we may replace S by its convex hull conv(S). The problem can be solved efficiently if we can find a small description of this polytope. Since for NP hard problems we cannot expect this, we look instead for approximations to conv(S). One possibility is to use semidefinite approximations, as introduced by Lovász [9] with the construction of the *theta body* of the stable set polytope of a graph. Another famous example is the approximation algorithm for the max cut problem due to Goemans and Williamson [2]. In this paper we will use the semidefinite relaxations introduced by Gouveia, Parrilo and Thomas [4] to analyze the *K_i-cover problem*.

Recall that K_i denotes the complete graph, or clique, on i vertices. Given a graph G, let $\mathbf{K}_j(G)$ be the collection of cliques in G of size j (usually, the graph is clear from context, and we write $\mathbf{K}_j$). A (possibly empty) collection $C \subset \mathbf{K}_{i-1}$ is said to be a K_i-cover if for each $K \in \mathbf{K}_i$, there is some $H \in C$ with $H \subset K$. In this case we say that H covers K. The K_i-cover problem is, given a graph G and a set of weights on $\mathbf{K}_{i-1}$, to compute the minimum weight K_i-cover. The case $i = 2$ is more commonly known as the vertex cover problem, in which we seek a collection C of vertices such that each edge in G contains at least one vertex from C. However, note that the usage of "cover" is reversed here: the vertex cover problem is the K_2-cover problem, not the K_1-cover problem.

A closely related problem, and the setting in which we will prove our results, is the *K_i-free problem*. As before, we are given a graph and a collection of weights on $\mathbf{K}_{i-1}$. But now we seek the maximum weight collection $C \subseteq \mathbf{K}_{i-1}$ such that C is K_i-free. That is, for each $K \in \mathbf{K}_i$, there is some $H \in \mathbf{K}_{i-1}$, with $H \subset K$ and $H \notin C$. Again, the case $i = 2$ of this problem is well-known as the stable set problem: we seek a maximum weight *stable set* C, where C is stable if no two of its vertices are connected by an edge.

The vertex cover and stable set problems are related in the following sense: let $G = (V, E)$ be a graph. Then a subset C of vertices is a vertex cover if and only if $V \setminus C$ is a stable set. The same is true for the K_i-cover and K_i-free problems: a subset $C \subset \mathbf{K}_{i-1}$ is a K_i-cover if and only if $\mathbf{K}_{i-1} \setminus C$ is K_i-free. Therefore, for a given set of weights on $\mathbf{K}_{i-1}$, optimal solutions to the two problems are complementary, and so solving one solves the other.

In this paper, we consider the polytope associated with the K_i-free problem. Let $P_i(G) = \text{conv}(\{\chi_S : S \subset \mathbf{K}_{i-1}(G) \text{ and } S \text{ is } K_i\text{-free}\})$, the convex hull of the incidence vectors of the K_i-free sets. Note that $P_i(G) \subseteq [0, 1]^{\mathbf{K}_{i-1}(G)}$.

As the K_i-free problem is NP-complete (see [1]), we cannot expect a small description of $P_i(G)$ for general graphs G. However, for certain classes of facets of $P_i(G)$, Conforti, Corneil, and Mahjoub [1] show that we can solve the separation problem in polynomial time, allowing us to optimize efficiently over a relaxation of $P_i(G)$. We provide a strictly tighter relaxation of $P_i(G)$, improving their optimization result, but without proving the existence of polynomial separation oracles for any new family of facets.

* Corresponding author.
E-mail addresses: jgouveia@mat.uc.pt (J. Gouveia), jamesrpfeiffer@gmail.com, jpfeiff@math.washington.edu (J. Pfeiffer).

http://dx.doi.org/10.1016/j.orl.2014.01.003
0167-6377/

The structure of this paper is: in Section 2, we outline the main algebraic machinery, *theta bodies*, a semidefinite relaxation hierarchy. In Section 3 we show that the K_i-p-hole facets are valid on $\lceil i/2 \rceil$ level of the theta body hierarchy. Finally, in Section 4 we focus on the triangle free problem. We show that in the case of $G = K_n$, the theta body relaxations cannot converge in less than $(n-2)/4$ steps. We also use a result of Krivelevich [7] to show an integrality gap of 2 for the second theta body.

2. Theta bodies

Theta bodies are semidefinite approximations to the convex hull of an algebraic variety. For background, see [5,4]. Here we state the necessary results for this paper without proofs.

Let $V \subseteq \mathbb{R}^n$ be a finite point set. One description of the convex hull of V is as the intersection of all affine half spaces containing V (recall that $f|_V$ is the restriction of f to V):

$$\text{conv}(V) = \{x \in \mathbb{R}^n : f(x) \geq 0 \text{ for all linear } f \text{ such that } f|_V \geq 0\}.$$

Since it is computationally intractable to find whether $f|_V \geq 0$, we relax this condition. Let I be the vanishing ideal of V, i.e., the set of all polynomials vanishing on V. Recall that $f \equiv g \mod I$ means $f - g \in I$, and implies that $f(x) = g(x)$ for all $x \in V$. A function f is said to be a sum of squares of degree at most k mod I, or *k-sos mod I*, if there exist functions g_j, $j = 1, \ldots, m$ with degree at most k, such that $f \equiv \sum_{j=1}^m g_j^2 \mod I$. If f is k-sos mod I for any k, it is clear that $f|_V \geq 0$ since g_j^2 is visibly nonnegative on V. Therefore, we make the following definition of $\text{TH}_k(I)$, the k-th theta body of I:

$$\text{TH}_k(I) = \{x \in \mathbb{R}^n : f(x) \geq 0 \text{ for all linear } f \equiv k\text{-sos mod } I\}.$$

The reason why the theta bodies $\text{TH}_k(I)$ provide a computationally tractable relaxation of $\text{conv}(V)$ is that the membership problem for $\text{TH}_k(I)$ can be expressed as a semidefinite program, using *moment matrices* that are reduced mod I.

For what follows, we will restrict ourselves to a special class of varieties, and suppose that our variety $V \subseteq \{0, 1\}^n$ and is *lower-comprehensive*; i.e., if $x \leq y$ componentwise, and $y \in V$, then $x \in V$. Additionally, we will always assume that V contains the canonical basis of $\mathbb{R}^n$, $\{e_1, \ldots, e_n\}$, as otherwise we could restrict ourselves to a subspace. All combinatorial optimization problems of avoiding certain finite list of configurations, such as stable set, K_i-free, etc., have lower-comprehensive varieties. The restriction to this class is not necessary, but makes the theta body exposition simpler. In particular, the ideal of a lower-comprehensive variety has the following simple description.

Lemma 2.1. *Let V be a lower-comprehensive subset of $\{0, 1\}^n$. Then its vanishing ideal is given by*

$$I = \langle x_j^2 - x_j : j = 1, \ldots, n; x^S : S \notin V \rangle,$$

and a basis for $\mathbb{R}[V] = \mathbb{R}[x]/I$ is given by $B = \{x^S : S \in V\}$, where $x^S := \prod_{i \in S} x_i$ is a shorthand used throughout the paper.

Another important fact about $\text{TH}_k(I)$ in this setting (when I is a real ideal) is that a linear inequality $f(x) \geq 0$ is valid on $\text{TH}_k(I)$ if and only if f is actually k-sos modulo I. In Section 3, we will prove that certain facet-defining inequalities of $P_i(G)$ are also valid on its theta relaxations $\text{TH}_k(I)$ by presenting a sum of squares representation modulo the ideal. For now, we observe that by considering degrees, we can get a bound on which theta bodies are trivial; that is, equal to the hypercube $[0, 1]^n$.

Lemma 2.2. *Let $V \subseteq \{0, 1\}^n$ be lower-comprehensive, and suppose that all elements $x \notin V$ have $\sum_j x_j \geq k$. Let I be the vanishing ideal of V. Then for $l < k/2$, $TH_l(I) = [0, 1]^n$.*

Proof. Let f be linear with $f \equiv \sum_j g_j^2 \mod I$ with each g_j of degree at most l. Then $f - \sum_j g_j^2 =: F \in I$, and F has degree at most $2l < k$, since $\deg(g_j^2) \leq 2l$. But the basis from Lemma 2.1 is a Groebner basis, and the only elements with degree less than k are $x_j^2 - x_j$, so $F \in I' := \langle x_j^2 - x_j; j = 1, \ldots, n \rangle$. Thus $\text{TH}_l(I) \supseteq \text{TH}_l(I') = [0, 1]^n$. □

Let V_k be the subset of V whose elements have at most k entries equal to one. For convenience, we will often identify the elements of V, characteristic vectors χ_S for $S \subseteq \{1, \ldots, n\}$, with their supports, via $S \leftrightarrow \chi_S$. Given $y \in \mathbb{R}^{V_{2k}}$ we denote the *reduced moment matrix* of y with respect to I to be the matrix $M_{V_k}(y) \in \mathbb{R}^{V_k \times V_k}$ defined by

$$[M_{V_k}(y)]_{X,Y} = \begin{cases} y_{X \cup Y} & \text{if } X \cup Y \in V, \\ 0 & \text{otherwise.} \end{cases}$$

With these matrices we can finally give a semidefinite description of $\text{TH}_k(I)$.

Proposition 2.3. *With I and V as before, $TH_k(I)$ is the canonical projection onto $\mathbb{R}^n$ via the coordinates $(y_{e_1}, \ldots, y_{e_n})$ of the set*

$$\{y \in \mathbb{R}^{V_{2k}} : M_{V_k}(y) \succeq 0 \text{ and } y_0 = 1\}.$$

In particular, optimizing to arbitrary fixed precision over $TH_k(I)$ can be done in time polynomial in n, for fixed k.

Now we can consider the specific case of the K_i-free problem. Here the variety $V \subseteq \mathbb{R}^{\mathbf{K}_{i-1}(G)}$ is the set of characteristic vectors of K_i-free subsets of $\mathbf{K}_{i-1}(G)$, V_k is the subset of V of elements of size at most k, and I is the vanishing ideal of V, described by Lemma 2.1. Since the K_is in G are the minimal elements not in V, by Lemma 2.1 we can write the ideal I as follows.

$$I = \left\langle x_j^2 - x_j : j \in \mathbf{K}_{i-1}(G); \prod_{j \subseteq K} x_j : K \in \mathbf{K}_i(G) \right\rangle.$$

For example, let G be a triangle, with edges A, B, C, and consider the triangle free problem on G. Then the ideal is

$$I = \langle x_A^2 - x_A, x_B^2 - x_B, x_C^2 - x_C, x_A x_B x_C \rangle,$$

and the variety V is as follows.

$$V = \{\emptyset, \{A\}, \{B\}, \{C\}, \{A, B\}, \{A, C\}, \{B, C\}\} \equiv \{0, 1, 2, 3, 4, 5, 6\}.$$

Note that here, we again use our identification of sets with their characteristic vectors. To avoid writing, e.g., $y_{\{A,C\}}$ or even $y_{\chi_{\{A,C\}}}$, we label the elements of V by numbers as above. Then the moment matrix $M_{V_2}(y)$ is as follows:

$$M_{V_2}(y) = \begin{bmatrix} y_0 & y_1 & y_2 & y_3 & y_4 & y_5 & y_6 \\ y_1 & y_1 & y_4 & y_5 & y_4 & y_5 & 0 \\ y_2 & y_4 & y_2 & y_6 & y_4 & 0 & y_6 \\ y_3 & y_5 & y_6 & y_3 & 0 & y_5 & y_6 \\ y_4 & y_4 & y_4 & 0 & y_4 & 0 & 0 \\ y_5 & y_5 & 0 & y_5 & 0 & y_5 & 0 \\ y_6 & 0 & y_6 & y_6 & 0 & 0 & y_6 \end{bmatrix}.$$

Projecting the set $\{y : y_0 = 1, M_{V_2}(y) \succeq 0\}$ onto (y_1, y_2, y_3) gives $\text{TH}_2(I)$ for this graph.

3. Polynomial-time algorithm

A graph H is a K_i-p-hole if H is the union of $G_1, \ldots, G_p$, each a copy of K_i, where G_j and G_l share a common K_{i-1} if and only if $j - l = \pm 1 \mod p$; see Fig. 1. Theorem 3.5 in [1] establishes that for $i \geq 3$ and odd p, the inequality $\sum_{\mathbf{K}_{i-1}(H)} x_j \leq (\frac{p-1}{2})(2i-3)+i-2$ defines a facet of $P_i(G)$ for each induced K_i-p-hole H of G. We will show that the facets corresponding to induced K_i-p-holes are valid