European Journal of Operational Research 240 (2015) 328–337

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Continuous Optimization

Continuous quadratic programming formulations of optimization problems on graphs ☆

William W. Hager *, James T. Hungerford [1]

Department of Mathematics, University of Florida, PO Box 118105, Gainesville, FL 32611-8105, United States

ARTICLE INFO

Article history:
Received 21 February 2013
Accepted 29 May 2014
Available online 18 June 2014

Keywords:
Vertex separator
Graph partitioning
Maximum clique
Continuous formulation
Quadratic programming

ABSTRACT

Four NP-hard optimization problems on graphs are studied: The vertex separator problem, the edge separator problem, the maximum clique problem, and the maximum independent set problem. We show that the vertex separator problem is equivalent to a continuous bilinear quadratic program. This continuous formulation is compared to known continuous quadratic programming formulations for the edge separator problem, the maximum clique problem, and the maximum independent set problem. All of these formulations, when expressed as maximization problems, are shown to follow from the convexity properties of the objective function along the edges of the feasible set. An algorithm is given which exploits the continuous formulation of the vertex separator problem to quickly compute approximate separators. Computational results are given.

1. Introduction

This paper concerns discrete optimization problems on graphs and their formulation as continuous quadratic programming problems. The paper initially focuses on the vertex separator problem, but later observes that the analytical techniques developed to handle this problem are also applicable to other optimization problems on graphs including the edge separator, maximum clique, and the maximum independent set problems.

Let G be a simple, undirected graph with vertices

$$\mathcal{V} = \{1, 2, \ldots, n\},$$

with edges $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$, and with nonnegative vertex weights $c_1, c_2, \ldots, c_n$, not all zero. Since the graph is undirected, $(i,j) \in \mathcal{E}$ if and only if $(j,i) \in \mathcal{E}$, and since the graph is simple, the elements of $\mathcal{E}$ are distinct and $(i,i) \notin \mathcal{E}$ for any $i \in \mathcal{V}$. A *vertex separator* of G is a set of vertices $\mathcal{S}$ whose removal breaks the graph into two disconnected sets of vertices $\mathcal{A}$ and $\mathcal{B}$. That is, $(\mathcal{A} \times \mathcal{B}) \cap \mathcal{E}$ is empty. The vertex separator problem (VSP) is to minimize the sum of the weights of vertices in $\mathcal{S}$ while requiring that $\mathcal{A}$ and $\mathcal{B}$ satisfy size constraints:

$$\min_{\mathcal{A},\mathcal{B} \subset \mathcal{V}} \sum_{i \in \mathcal{S}} c_i \tag{1.1}$$

$$\text{subject to } \mathcal{S} = \mathcal{V} \setminus (\mathcal{A} \cup \mathcal{B}), \quad \mathcal{A} \cap \mathcal{B} = \emptyset, \quad (\mathcal{A} \times \mathcal{B}) \cap \mathcal{E} = \emptyset,$$
$$\ell_a \leqslant |\mathcal{A}| \leqslant u_a, \quad \ell_b \leqslant |\mathcal{B}| \leqslant u_b.$$

Here $|\mathcal{A}|$ denotes the number of elements in the set $\mathcal{A}$, and ℓ_a, u_a, ℓ_b, and u_b are given integer parameters that describe the flexibility in the size of the sets $\mathcal{A}$ and $\mathcal{B}$. These parameters should be such that $0 \leqslant \ell_a \leqslant u_a \leqslant n \geqslant u_b \geqslant \ell_b \geqslant 0$. We assume that (1.1) is feasible. If $\ell_a, \ell_b \geqslant 1$, then this implies G is not complete; that is, for some $i \neq j \in \mathcal{V}$, we have $(i,j) \notin \mathcal{E}$. Vertex separators have applications to VLSI chip design (Kernighan & Lin, 1970; Leiserson, 1980; Ullman, 1984), to finite element methods (Miller, Teng, Thurston, & Vavasis, 1998), to parallel processing (Evrendilek, 2008), to the computation of fill-reducing orderings for sparse matrix factorizations (George & Liu, 1981), and to network security.

An alternative definition of a vertex separator is sometimes used: $\mathcal{S}$ is a vertex separator with respect to $\mathcal{A}$ and $\mathcal{B}$ if every path from $\mathcal{A}$ to $\mathcal{B}$ passes through a vertex in $\mathcal{S}$. This definition is helpful, since it generalizes to the notion of a *wide* separator, a separator in which every path from $\mathcal{A}$ to $\mathcal{B}$ passes through at least two vertices of $\mathcal{S}$. While some authors give special treatment to wide separators (see for instance Pothen, Simon, & Liou, 1990 and references therein), it is worth noting that any method for finding vertex separators can also be used to find wide separators by simply

☆ January 31, 2014. This material is based upon work supported by the National Science Foundation under Grants 0619080 and 0620286, and the Office of Naval Research under Grant N00014-11-1-0068.

* Corresponding author. Tel.: +1 (352) 294 2308; fax: +1 (352) 392 8357.
E-mail addresses: hager@ufl.edu (W.W. Hager), freerad@ufl.edu (J.T. Hungerford).
URLs: http://www.people.clas.ufl.edu/hager/ (W.W. Hager), http://www.people.clas.ufl.edu/freerad/ (J.T. Hungerford).
[1] Tel.: +1 (352) 294 2350; fax: +1 (352) 392 8357.

http://dx.doi.org/10.1016/j.ejor.2014.05.042
0377-2217/

adding edges to the graph. In particular, define an edge set $\mathcal{E}_w$ be letting $(i,j) \in \mathcal{E}_w$ if and only if i and j are connected by a path of length 1 or 2 in G. Note that if $\mathbf{A}$ is the adjacency matrix for G defined by $a_{ij} = 1$ when $(i,j) \in \mathcal{E}$, and $a_{ij} = 0$ otherwise, then $\mathcal{E}_w$ corresponds to the nonzero off-diagonal elements of $\mathbf{A} + \mathbf{A}^2$. And if $\mathcal{E}$ is replaced by $\mathcal{E}_w$ in (1.1), then $(\mathcal{A} \times \mathcal{B}) \cap \mathcal{E}_w = \emptyset$ if and only if every path from $\mathcal{A}$ to $\mathcal{B}$ passes through at least two vertices in $\mathcal{S}$.

The difficulty of solving the VSP and the size of the optimal solution are strongly tied to the structure of the graph. For instance, every tree has an optimal vertex separator consisting of exactly one vertex. For planar graphs, Lipton and Tarjan showed that a separator of size $O(\sqrt{n})$ can be found in linear time (Lipton & Tarjan, 1979). However, for general graphs (and even planar graphs), the VSP is NP-hard (Bui & Jones, 1992; Fukuyama, 2006). Hence, heuristic algorithms have been developed for obtaining approximate solutions; for example, see (Benlic & Hao, 2013; Djidjev, 2000; Evrendilek, 2008; Feige, Hajiaghayi, & Lee, 2008; Karypis & Kumar, 1995).

In Balas and de Souza (2005), Cavalcante and de Souza (2007) and de Souza and Balas (2005), the authors studied the following exact integer programming formulation of the VSP:

$$\max_{\mathbf{x},\mathbf{y} \in \mathbb{B}^n} \mathbf{c}^{\mathsf{T}}(\mathbf{x} + \mathbf{y}) \tag{1.2}$$

subject to $\mathbf{x} + \mathbf{y} \leqslant \mathbf{1}, \quad x_i + y_j \leqslant 1$ for every $(i,j) \in \mathcal{E}$,

$$\ell_a \leqslant \mathbf{1}^{\mathsf{T}}\mathbf{x} \leqslant u_a, \quad \text{and} \quad \ell_b \leqslant \mathbf{1}^{\mathsf{T}}\mathbf{y} \leqslant u_b.$$

Here $\mathbb{B}^n = \{0,1\}^n$ is the collection of binary vectors with n components, c_i is the weight of vertex i, $\mathbf{1}$ is the vector whose entries are all 1, and $\mathbf{x}$ and $\mathbf{y}$ are the incidence vectors for $\mathcal{A}$ and $\mathcal{B}$ respectively; that is, $x_i = 1$ if $i \in \mathcal{A}$ and $x_i = 0$ otherwise. Minimizing the weight of the separator is equivalent to maximizing the weight $\mathbf{c}^{\mathsf{T}}(\mathbf{x} + \mathbf{y})$ of the vertices outside the separator. The componentwise inequality $\mathbf{x} + \mathbf{y} \leqslant \mathbf{1}$ is the partition constraint, which ensures that $\mathcal{A}$ and $\mathcal{B}$ are disjoint. The condition $x_i + y_j \leqslant 1$ when $(i,j) \in \mathcal{E}$ is the separation constraint, which ensures that $(\mathcal{A} \times \mathcal{B}) \cap \mathcal{E} = \emptyset$. Finally, the balance constraints $\ell_a \leqslant \mathbf{1}^{\mathsf{T}}\mathbf{x} \leqslant u_a$ and $\ell_b \leqslant \mathbf{1}^{\mathsf{T}}\mathbf{y} \leqslant u_b$ restrict the size of the sets $\mathcal{A}$ and $\mathcal{B}$.

In Balas and de Souza (2005) and de Souza and Balas (2005) the authors studied the program (1.2) in the case where $\ell_a = \ell_b = 1$. Valid inequalities for the integer polytope of (1.2) were obtained and the program was solved on a variety of small ($n \leqslant 200$) problem instances using a branch and cut algorithm. In Cavalcante and de Souza (2007) an improved algorithm was presented which made use of Lagrangian relaxation, improving the pool of cutting planes and providing better primal bounds for the nodes in the search tree.

In the current paper, we develop conditions under which the VSP is equivalent to the following continuous bilinear quadratic program for some choice of the parameter $\gamma > 0$:

$$\max_{\mathbf{x},\mathbf{y} \in \mathbb{R}^n} \mathbf{c}^{\mathsf{T}}(\mathbf{x} + \mathbf{y}) - \gamma \mathbf{x}^{\mathsf{T}}(\mathbf{A} + \mathbf{I})\mathbf{y} \tag{1.3}$$

subject to $\mathbf{0} \leqslant \mathbf{x} \leqslant \mathbf{1}, \ \mathbf{0} \leqslant \mathbf{y} \leqslant \mathbf{1}, \ \ell_a \leqslant \mathbf{1}^{\mathsf{T}}\mathbf{x} \leqslant u_a,$

and $\ell_b \leqslant \mathbf{1}^{\mathsf{T}}\mathbf{y} \leqslant u_b.$

We will show that the term $-\gamma \mathbf{x}^{\mathsf{T}}(\mathbf{A} + \mathbf{I})\mathbf{y}$ in the objective function amounts to a penalty term for enforcing the separation constraint that there are no edges connecting $\mathcal{A}$ and $\mathcal{B}$.

We show that (1.3) is equivalent to the VSP if the following conditions are satisfied:

(C1) $\gamma \geqslant c_i$ for all i.

(C2) The total weight of an optimal vertex separator is less than or equal to

$$\left(\sum_{i=1}^{n} c_i\right) - \gamma(\ell_a + \ell_b). \tag{1.4}$$

The first condition is satisfied by taking $\gamma = \max\{c_i : 1 \leqslant i \leqslant n\}$. In practice, the second condition is often easily satisfied. In the common case where $c_i = 1$ for all i and $\ell_a = \ell_b = 1$, the expression (1.4) reduces to $n - 2$. Hence, in this case, (C2) is satisfied as long as (1.1) is feasible, since $\mathcal{A}$ and $\mathcal{B}$ must each contain at least one vertex.

The equivalence between the VSP and (1.3) is in the following sense: For any solution of (1.3), there is an associated, easily constructed binary solution. Moreover, when (C1) and (C2) hold, there exists a binary solution for which the penalty term vanishes and the separation constraint is satisfied. For such a solution, an optimal separator for the VSP is given by

$$\mathcal{A} = \{i : x_i = 1\}, \quad \mathcal{B} = \{i : y_i = 1\}, \quad \text{and} \quad \mathcal{S} = \{i : x_i = y_i = 0\}. \tag{1.5}$$

In some applications such as finite element methods, parallel processing, and sparse matrix factorizations, it is important to obtain an approximate solution to the VSP quickly. We show in Section 7 how the continuous formulation may be incorporated into a multilevel algorithm for finding approximate solutions in a reasonable amount of CPU time. Multilevel algorithms have recently been shown to produce fast and high quality solutions to a variety of graph problems (Hendrickson & Leland, 1995; Karypis & Kumar, 1995; Safro, Ron, & Brandt, 2006a, Safro, Ron, & Brandt, 2006b). Other options for finding approximate separators include the use of standard optimization algorithms applied to (1.3), such as the gradient projection algorithm (see Bertsekas, 1999).

In other applications where we need to solve (1.3) exactly, branch and bound techniques can be applied. For illustration, in Hager, Phan, and Zhang (2013) we develop a branch and bound algorithm for the closely related edge separator problem. The continuous formulation of the edge separator problem is the same as (1.3), but with the additional constraint $\mathbf{x} + \mathbf{y} = \mathbf{1}$. In Hager et al. (2013) we show that a branch and bound algorithm applied to the continuous formulation of the edge separator problem is particularly effective for sparse graphs.

As noted earlier, our continuous formulation of the VSP is in some sense an exact penalty method. In most exact penalty methods for solving binary maximization problems, the penalty function is chosen both to make the objective function convex, guaranteeing an extreme point solution (Bauer, 1958), and to force the extreme solution to be binary (Giannessi & Niccolucci, 1976; Raghavachari, 1969). Our penalty formulation differs in these two crucial aspects. In particular, if f is the objective function in (1.3):

$$f(\mathbf{x},\mathbf{y}) = \mathbf{c}^{\mathsf{T}}(\mathbf{x} + \mathbf{y}) - \gamma \mathbf{x}^{\mathsf{T}}(\mathbf{A} + \mathbf{I})\mathbf{y},$$

then the Hessian is

$$\nabla^2 f = -\gamma \begin{pmatrix} \mathbf{0} & \mathbf{B} \\ \mathbf{B} & \mathbf{0} \end{pmatrix}, \quad \mathbf{B} = \mathbf{A} + \mathbf{I}.$$

If $\lambda_i, \ 1 \leqslant i \leqslant n$, are the eigenvalues of $\mathbf{B}$, then $\pm\gamma\lambda_i, \ 1 \leqslant i \leqslant n$, are the eigenvalues of the Hessian. Hence, f is neither convex nor concave, and the number of positive eigenvalues of the Hessian is equal to the number of negative eigenvalues. Nonetheless, we will show that f *is* convex along the directions parallel to the edges of the feasible set in (1.3). Consequently, the existence of an extreme point maximizer follows from results of Tardella (1990). Furthermore, we show that every extreme point of the constraint polyhedron in (1.3) is binary, and if (C1) and (C2) hold, then there exists a binary maximizer of (1.3) such that $\mathbf{x}^{\mathsf{T}}(\mathbf{A} + \mathbf{I})\mathbf{y} = 0$.