

# A Continuous Refinement Strategy for the Multilevel Computation of Vertex Separators

William W. Hager<sup>1</sup>, James T. Hungerford<sup>1</sup>, and Ilya Safro<sup>2</sup>

<sup>1</sup> Department of Mathematics, University of Florida, Gainesville, FL

<sup>2</sup> School of Computing, Clemson University, Clemson, SC

**Abstract.** The Vertex Separator Problem (VSP) on a graph is the problem of finding the smallest collection of vertices whose removal separates the graph into two disjoint subsets of roughly equal size. Recently, Hager and Hungerford [1] developed a continuous bilinear programming formulation of the VSP. In this paper, we reinforce the bilinear programming approach with a multilevel scheme for learning the structure of the graph.

## 1 Introduction

Let  $G = (\mathcal{V}, \mathcal{E})$  be an undirected graph with vertex set  $\mathcal{V}$  and edge set  $\mathcal{E}$ . Vertices are labeled  $1, 2, \dots, n$ . We assign to each vertex a non-negative weight  $c_i \in \mathbb{R}_{\geq 0}$ . If  $\mathcal{Z} \subset \mathcal{V}$ , then we let  $\mathcal{W}(\mathcal{Z}) = \sum_{i \in \mathcal{Z}} c_i$  be the total weight of vertices in  $\mathcal{Z}$ . Throughout the paper, we assume that  $G$  is simple; that is, there are no loops or multiple edges between vertices.

The Vertex Separator Problem (VSP) on  $G$  is to find the smallest weight subset  $\mathcal{S} \subset \mathcal{V}$  whose removal separates the graph into two roughly equal sized subsets  $\mathcal{A}, \mathcal{B} \subset \mathcal{V}$  such that there are no edges between  $\mathcal{A}$  and  $\mathcal{B}$ ; that is,  $(\mathcal{A} \times \mathcal{B}) \cap \mathcal{E} = \emptyset$ . We may formulate the VSP as

$$\begin{aligned} & \min_{\mathcal{A}, \mathcal{B}, \mathcal{S} \subset \mathcal{V}} \mathcal{W}(\mathcal{S}) \\ \text{subject to } & \mathcal{A} \cap \mathcal{B} = \emptyset, \quad (\mathcal{A} \times \mathcal{B}) \cap \mathcal{E} = \emptyset, \quad \mathcal{S} = \mathcal{V} \setminus (\mathcal{A} \cup \mathcal{B}), \\ & \ell_a \leq |\mathcal{A}| \leq u_a, \quad \text{and} \quad \ell_b \leq |\mathcal{B}| \leq u_b. \end{aligned} \quad (1)$$

Here, the size constraints on  $\mathcal{A}$  and  $\mathcal{B}$  take the form of upper and lower bounds. Since the weight of an optimal separator  $\mathcal{S}$  is typically small, in practice the lower bounds on  $\mathcal{A}$  and  $\mathcal{B}$  are almost never attained at an optimal solution, and may be taken to be quite small. In [2], the authors consider the case where  $\ell_a = \ell_b = 1$  and  $u_a = u_b = \frac{2n}{3}$  for the development of efficient divide and conquer algorithms. The VSP has several applications, including parallel computations [3], VLSI design [4, 5], and network security. Like most graph partitioning problems, the VSP is NP-hard [6]. Heuristic methods proposed include vertex swapping algorithms [5, 7], spectral methods [3], continuous bilinear programming [1], and semidefinite programming [8].

For large-scale graphs, heuristics are more effective when reinforced by a multilevel framework: first coarsen the graph to a suitably small size; then, solve the

problem for the coarse graph; and finally, uncoarsen the solution and refine it to obtain a solution for the original graph [9]. Many different multilevel frameworks have been proposed in the past two decades [10]. One of the most crucial parameters in a multilevel algorithm is the choice of the refinement scheme. Most multilevel graph partitioners and VSP solvers refine solutions using variants of the Kernighan-Lin [5] or Fiducia-Matheyses [7, 11] algorithms. In these algorithms, a low weight edge cut is found by making a series of vertex swaps starting from an initial partition, and a vertex separator is obtained by selecting vertices incident to the edges in the cut. One disadvantage of using these schemes is that they assume that an optimal vertex separator lies near an optimal edge cut. As pointed out in [8], this assumption need not hold in general.

In this article, we present a new refinement strategy for multilevel separator algorithms which computes vertex separators directly. Refinements are based on solving the following continuous bilinear program (CBP):

$$\max_{\mathbf{x}, \mathbf{y} \in \mathbb{R}^n} \quad \mathbf{c}^\top(\mathbf{x} + \mathbf{y}) - \gamma \mathbf{x}^\top(\mathbf{A} + \mathbf{I})\mathbf{y} \quad (2)$$

subject to  $\mathbf{0} \leq \mathbf{x} \leq \mathbf{1}$ ,  $\mathbf{0} \leq \mathbf{y} \leq \mathbf{1}$ ,  $\ell_a \leq \mathbf{1}^\top \mathbf{x} \leq u_a$ , and  $\ell_b \leq \mathbf{1}^\top \mathbf{y} \leq u_b$ .

Here,  $\mathbf{A}$  denotes the adjacency matrix for  $G$  (defined by  $a_{ij} = 1$  if  $(i, j) \in \mathcal{E}$  and  $a_{ij} = 0$  otherwise),  $\mathbf{I}$  is the  $n \times n$  identity matrix,  $\mathbf{c} \in \mathbb{R}^n$  stores the vertex weights, and  $\gamma := \max \{c_i : i \in \mathcal{V}\}$ . In [1], the authors show that (2) is equivalent to (1) in the following sense: Given any feasible point  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$  of (2), one can find a piecewise linear path to another feasible point  $(\mathbf{x}, \mathbf{y})$  such that

$$f(\mathbf{x}, \mathbf{y}) \geq f(\hat{\mathbf{x}}, \hat{\mathbf{y}}), \quad \mathbf{x}, \mathbf{y} \in \{0, 1\}^n, \quad \text{and} \quad \mathbf{x}^\top(\mathbf{A} + \mathbf{I})\mathbf{y} = 0. \quad (3)$$

(see the proof of Theorem 2.1, [1]). In particular, there exists a global solution to (2) satisfying (3), and for any such solution, an optimal solution to (1) is given by

$$\mathcal{A} = \{i : x_i = 1\}, \quad \mathcal{B} = \{i : y_i = 1\}, \quad \mathcal{S} = \{i : x_i = y_i = 0\}. \quad (4)$$

(Note that the fact that (4) is a partition of  $\mathcal{V}$  with  $(\mathcal{A} \times \mathcal{B}) \cap \mathcal{E} = \emptyset$  follows from the last property of (3).)

In the next section, we outline a multilevel algorithm which incorporates (2) in the refinement phase. Section 3 concludes the paper with some computational results comparing the effectiveness of this refinement strategy with traditional Kernighan-Lin refinements.

## 2 Algorithm

The graph  $G$  is coarsened by visiting each vertex and matching [10] it with an unmatched neighbor to which it is most strongly coupled. The strength of the coupling between vertices is measured using a heavy edge distance: For the finest graph, all edges are assigned a weight equal to 1; as the graph is coarsened, multiple edges arising between any two vertex aggregates are combined into a single edge which is assigned a weight equal to the sum of the weights of the

constituent edges. This process is applied recursively: first the finest graph is coarsened, then the coarse graph is coarsened again, and so on. When the graph has a suitably small size, the coarsening stops and the VSP is solved for the coarse graph using any available method (the bilinear program (2), Kernighan-Lin, etc.) The solution is stored as a pair of incidence vectors  $(\mathbf{x}^{\text{coarse}}, \mathbf{y}^{\text{coarse}})$  for  $\mathcal{A}$  and  $\mathcal{B}$  (see (4)).

When the graph is uncoarsened,  $(\mathbf{x}^{\text{coarse}}, \mathbf{y}^{\text{coarse}})$  yields a vertex separator for the next finer level by assigning components of  $\mathbf{x}^{\text{fine}}$  and  $\mathbf{y}^{\text{fine}}$  to be equal to 1 whenever their counterparts in the coarse graph were equal to 1, and similarly for the components equal to 0. This initial solution is refined by alternately holding  $\mathbf{x}$  or  $\mathbf{y}$  fixed, while solving (2) over the free variable and taking a step in the direction of the solution. (Note that when  $\mathbf{x}$  or  $\mathbf{y}$  is fixed, (2) is a linear program in the free variable, and thus can be solved efficiently.) When no further improvement is possible in either variable, the refinement phase terminates and a separator is retrieved by moving to a point  $(\mathbf{x}, \mathbf{y})$  which satisfies (3).

Many multilevel algorithms employ techniques for escaping false local optima encountered during the refinement phase. For example, in [12] simulated annealing is used. In the current algorithm, local maxima are escaped by reducing the penalty parameter  $\gamma$  from its initial value of  $\max \{c_i : i \in \mathcal{V}\}$ . The reduced problem is solved using the current solution as a starting guess. If the current solution is escaped, then  $\gamma$  is returned to its initial value and the refinement phase is repeated. Otherwise,  $\gamma$  is reduced in small increments until it reaches 0 and the escape phase terminates.

### 3 Computational Results

The algorithm was implemented in C++. Graph structures such as the adjacency matrix and the vertex weights were stored using the LEMON Graph Library [13]. For our preliminary experiments, we used several symmetric matrices from the University of Florida Sparse Matrix Library having dimensions between 1000 and 5000. For all problems, we used the parameters  $\ell_a = \ell_b = 1$ ,  $u_a = u_b = \lfloor 0.503n \rfloor$ , and  $c_i = 1$  for each  $i = 1, 2, \dots, n$ . We compared the sizes of the separators obtained by our algorithm with the routine METIS\_ComputeVertexSeparator available from METIS 5.1.0. Comparisons are given in Table 1.

Problem	$ \mathcal{V} $	Sparsity	CBP	METIS	Problem	$ \mathcal{V} $	Sparsity	CBP	METIS
BCSPWR09	1723	.0016	8	7	G42	2000	.0059	498	489
NETZ4504	1961	.0013	17	20	LSHP3466	3466	.0017	61	61
SSTMODEL	3345	.0017	26	23	MINNESOTA	2642	.0009	17	21
JAGMESH7	1138	.0049	14	15	YEAST	2361	.0024	196	229
CRYSTM01	4875	.0042	65	65	SHERMAN1	1000	.0028	28	32

**Table 1.** Illustrative comparison between separators obtained using either METIS or CBP (2)

Both our algorithm and the METIS routine compute vertex separators using a multilevel scheme. Moreover, both algorithms coarsen the graph using a heavy edge distance. Therefore, since initial solutions obtained at the coarsest level are typically exact, the algorithms differ primarily in how the solution is refined during the uncoarsening process. While our algorithm refines using the CBP (2), METIS employs Kernighan-Lin style refinements. In half of the problems tested, the size of the separator obtained by our algorithm was smaller than that of METIS. No correlation was observed between problem dimension and the quality of the solutions obtained by either algorithm. Current preliminary implementation of our algorithm is not optimized, so the running time is not compared. (However, we note that both algorithms are of the same linear complexity.) Nevertheless, the results in Table 1 indicate that the bilinear program (2) can serve as an effective refinement tool in multilevel separator algorithms. We compared our solvers on graphs with heavy-tailed degree distributions and the results were very similar. We found that in contrast to the balanced graph partitioning [10], the practical VSP solvers are still very far from being optimal. We hypothesize that the breakthrough in the results for VSP lies in the combination of KL/FM and CBP refinements reinforced by a stronger coarsening scheme that introduces correct reductions in the problem dimensionality (see some ideas related to graph partitioning in [10]).

## References

1. Hager, W.W., Hungerford, J.T.: A continuous quadratic programming formulation of the vertex separator problem. *European J. Oper. Res.* (2013, submitted)
2. Balas, E., de Souza, C.C.: The vertex separator problem: a polyhedral investigation. *Math. Program.* **103** (2005) 583–608
3. Pothen, A., Simon, H.D., Liou, K.: Partitioning sparse matrices with eigenvectors of graphs. *SIAM J. Matrix Anal. Appl.* **11**(3) (1990) 430–452
4. Ullman, J.: *Computational Aspects of VLSI*. Computer Science Press (1984)
5. Kernighan, B.W., Lin, S.: An efficient heuristic procedure for partitioning graphs. *Bell System Tech. J.* **49** (1970) 291–307
6. Bui, T., Jones, C.: Finding good approximate vertex and edge partitions is NP-hard. *Information Processing Letters* **42** (1992) 153–159
7. Fiduccia, C.M., Mattheyses, R.M.: A linear-time heuristic for improving network partitions. In: *Proc. 19th Design Automation Conf., Las Vegas, NV* (1982) 175–181
8. Feige, U., Hajiaghayi, M., Lee, J.: Improved approximation algorithms for vertex separators. *SIAM J. Comput.* **38** (2008) 629–657
9. Ron, D., Safro, I., Brandt, A.: Relaxation-based coarsening and multiscale graph organization. *Multiscale Modeling & Simulation* **9**(1) (2011) 407–423
10. Buluc, A., Meyerhenke, H., Safro, I., Sanders, P., Schulz, C.: Recent advances in graph partitioning. (2013) arXiv:1311.3144
11. Leiserson, C., Lewis, J.: Orderings for parallel sparse symmetric factorization. In: *Third SIAM Conference on Parallel Processing for Scientific Computing*. (1987)
12. Safro, I., Ron, D., Brandt, A.: A multilevel algorithm for the minimum 2-sum problem. *J. Graph Algorithms Appl.* **10** (2006) 237–258
13. Dezső, B., Jüttner, A., Kovács, P.: Lemon - an open source C++ graph template library. *Electronic Notes in Theoretical Computer Science* **264**(5) (2011) 23–45