Available online at www.sciencedirect.com
ScienceDirect
Operations Research Letters 35 (2007) 739–742
Operations Research Letters
www.elsevier.com/locate/orl

Polynomiality of sparsest cuts with fixed number of sources

Sung-Pil Hong[a], Byung-Cheon Choi[b,*]

[a]Department of Industrial Engineering, Seoul National University, Seoul, Republic of Korea
[b]Semiconductor Business, Samsung Electronics Co., LTD., Gyeonggi-Do, Republic of Korea

Received 7 May 2006; accepted 24 December 2006
Available online 22 February 2007

Abstract

We show that when the number of sources is constant the sparsest cut problem is solvable in polynomial time.

Keywords: Sparsest cut problem; Fixed number of sources; Polynomiality

1. Introduction

The sparsest cut problem is a well-studied optimization problem arising in various contexts.

Problem 1.1 (*Sparsest cut problem*). Let $G=(V,E)$ be a connected undirected graph with integral edge capacities $c_e \geqslant 0$. Consider a set of source–sink pairs, $T=\{(s_1,t_1),\dots,(s_l,t_l)\} \subseteq V \times V$. An integral demand $d_i \geqslant 0$ is assigned to each pair (s_i,t_i). Let $\mathrm{dem}(E')$ be the demand separated by a cut $E' := (U; V\backslash U)$: $\mathrm{dem}(E') = \sum_{i:|\{s_i,t_i\}\cap U|=1} d_i$. Find a cut E' that minimizes the *sparsity*, $c(E')/\mathrm{dem}(E')$, where $c(E')$ denotes the capacity of E'.

Notice that we may assume there is no zero capacity cut. Otherwise, we can easily construct a trivial solution, or reduce the problem into smaller one(s). The problem is NP-hard [4]. If there are k vertices that cover all the pairs of T, then it will be referred to as a *k-source sparsest cut problem*. In this paper, we show that the k-source sparsest cut problem is solvable in $2^{\mathrm{O}(k)} \times$ a polynomial time of input size. Therefore, the sparsest cut problem is polynomial when the number of distinct sources is fixed. This work was motivated by the following application.

Example 1.2 (*K-server-net vulnerability*). Consider an undirected graph $G=(V,E)$ with k servers, $\{s_1,\dots,s_k\} \subseteq V$. The service of s_i has utility $d_{v,i} \geqslant 0$ to each user $v \in V$. Each edge is assigned a destruction cost by $c_e \geqslant 0$. If an edge set $E' \subseteq E$ is removed from G, some users will be disconnected from servers. The maximum disconnected utility per unit effort is a natural measure of the vulnerability of the network. Then, the problem is to compute maximum ratio of the disconnected utility $\mathrm{dscn}(E')$ to the edge-destruction cost $c(E')$: $\max_{E' \subseteq E}\{\mathrm{dscn}(E')/c(E')\}$.

* Corresponding author.
E-mail address: cbc@optima.snu.ac.kr (B.-C. Choi).

0167-6377/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.orl.2006.12.011

Although there is no upper bound on the number of connected components of $G\backslash E'$ in Example 1.2, the argument of Proposition 5.1 in [8] can be used to infer that the maximum is achieved when $G\backslash E'$ has exactly two connected components. Hence, the k-server-net vulnerability is the special case of the sparsest cut problem where each of k sources has $O(|V|)$ sinks and is polynomially solvable when k is a small constant as usually the case for a server network.

It is well known that the dual of the *maximum concurrent flow* problem [7] is an LP-relaxation of the sparsest cut problem. The worst case ratio between the sparsest cut and the maximum concurrent flow values has been intensively studied in establishing multi-commodity version of max-flow min-cut theorems pioneered by Leighton and Rao [3]. Along the line of research, Günlük [2] studied the k-source sparsest cut problem in improving the previous ratio, $O(\log l)$ to $O(\log k)$ (recall that l is the number of source–sink pairs in T), although he did not discuss how to solve the sparsest cut problem exactly in this case.

We also note that the polynomiality of the 2-source sparsest cut problem follows from that the maximum concurrent flow value provides the exact sparsest cut value when the graph induced by the source–sink pair edges is the union of two stars [5]. Readers are referred to Chapter 72 of [6] for some other special cases of the integrality gap of 1.

2. Polynomiality of case with fixed number of sources

In this section, we consider the k-source sparsest cut problem with the sources $S=\{s_1, s_2, \dots, s_k\} \subseteq V$. Assume that each vertex $v \in V$ has demand $d_{v,i}$ for each $s_i \in S$. We will show that the k-source sparsest cut problem is solvable in $2^{O(k)} \times$ a time polynomial in input size. In doing so, we first argue that for a fixed subset $Q \subseteq S$, the following problem is polynomially solvable via a parametric min-cut problem on an augmented graph.

Problem 2.1. Given $Q \subseteq S$, compute an optimal Q-*cut* that minimizes the sparsity $c(E')/\mathrm{dem}(E')$, where by a Q-cut we mean a cut $E' := (U; V\backslash U)$ with $Q \subseteq U$ and $(S\backslash Q) \subseteq (V\backslash U)$. The sparsity of an optimal Q-cut will be denoted by Φ_Q.

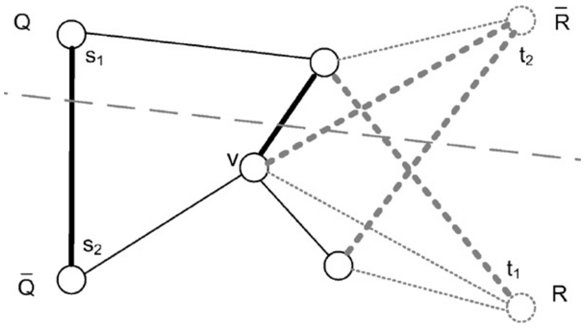

Fig. 1. The cut separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$. Solid, dotted and thick lines, respectively, represent the original graph, augmented part, and cut edges $E' \cup F'$.

Then, as there are 2^k possible Q's that U of an optimal solution $(U; V\backslash U)$ of a k-source sparsest cut problem can contain, the polynomiality of the case with a fixed k will follow.

The augmented graph $\tilde{G}$ and parametric min-cut problem are defined for a fixed $Q \subseteq S$ as follows. Augment k super-sinks $\{t_1, t_2, \dots, t_k\}$ to G and connect each $v \in V$ to each of the super-sinks with a new edge $\{v, t_i\}$. Assign each edge $\{v, t_i\}$ the capacity of $\beta d_{v,i}$ with a parameter $\beta > 0$. Also define the new demand between s_i and t_i to be $D_i = \sum_{v\in V} d_{v,i}$. Denote $\bar{Q} = S\backslash Q$, $R = \{t_i | s_i \in Q\}$, and by $\bar{R}$ the rest of the super-sinks. Consider the problem to find a minimum capacity cut that separates $Q \cup \bar{R}$ from $\bar{Q} \cup R$ in $\tilde{G}$. Throughout this paper, such a cut will be expressed as $E' \cup F'$ where $E' \subseteq E$ while the edges of F' are new. See Fig. 1.

Lemma 2.2. *The sparsity of a cut* $E' \cup F'$ *in* $\tilde{G}$ *is*

$$\frac{c(E') + \sum_{v,i:\{v,t_i\}\in F'} \beta d_{v,i}}{\sum_i D_i} = \frac{c(E') + \sum_{v,i:\{v,t_i\}\in F'} \beta d_{v,i}}{\sum_{\{s_i,v\}\ \textit{separated by}\ E'} d_{v,i} + \sum_{v,i:\{v,t_i\}\in F'} d_{v,i}}. \tag{1}$$

Proof. By the construction, the numerator is the capacity of the cut $E' \cup F'$. As every source–super-sink pair is disconnected, the separated demand is $\sum_{i=1}^{k} D_i$ and hence it suffices to show the equality of the two denominators. But, due to the cut separating $Q \cup \bar{R}$

from $\bar{Q} \cup R$, for each i, every $v \in V$ should be disconnected from either s_i or t_i, but never both. (See Fig. 1.) Hence the equality follows. □

As $\sum_i D_i$ is constant, Lemma 2.2 implies the equivalence between minimization of the cut-capacity and the sparsity (1) over the cuts separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$ in $\tilde{G}$. Therefore, for each β and Q, a cut $E' \cup F'$ both separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$ and minimizing the sparsity (1) can be computed in polynomial time using a max-flow algorithm.

Lemma 2.3. *For every $\beta \geqslant \Phi_Q$, in computing a cut $E' \cup F'$ separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$ that minimizes the sparsity* (1), *we can always choose one satisfying $\sum_{\{s_i,v\} \text{ separated by } E'} d_{v,i} > 0$.*

Proof. From $\beta \geqslant \Phi_Q$, G has a Q-cut $E'' := (U; V \backslash U)$ with $c(E'') \leqslant \beta \sum_{\{s_i,v\} \text{ separated by } E''} d_{v,i}$. Define $F'' = \{(v, t_i) : v \in U \text{ and } t_i \in R, \text{ or } v \in V \backslash U \text{ and } t_i \in \bar{R}\}$. Then, $E'' \cup F''$ is a cut separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$ whose sparsity (1) is $\leqslant \beta$, as easily seen. Let $E' \cup F'$ be a cut separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$ that minimizes (1).

Suppose that Q is a proper subset of S. Then, as $E' \neq \emptyset$, we have $c(E') > 0$. Therefore, if $\sum_{\{s_i,v\} \text{ separated by } E'} d_{v,i} = 0$, then the sparsity (1) of $E' \cup F'$ is $> \beta$, which contradicts the optimality of $E' \cup F'$.

Suppose, on the other hand, that $Q = S$. Then, from the min-cuts separating $Q \cup \bar{R} = Q$ and $\bar{Q} \cup R = R$, we can always choose one, say $E' \cup F'$ with the smallest Q-side (for instance, by a search algorithm on the residual capacity network obtained after the termination of a max-flow algorithm). Since $\beta \geqslant \Phi_Q$ and the sparsity (1) is equal to β when $E' = \emptyset$, this will guarantee $E' \neq \emptyset$ and hence, $c(E') > 0$. Thus, again the optimality of $E' \cup F'$ implies $\sum_{\{s_i,v\} \text{ separated by } E'} d_{v,i} > 0$. □

Now, we state some easy but useful facts without proof.

Lemma 2.4. *Let x be a nonnegative variable, y a positive variable, c and d constants. Then, $(x+c)/(y+d)$ is minimized if and only if x/y is. Hence, if z is another nonnegative variable, then $(x+cz)/(y+dz)$ is minimized if and only if x/y is.*

Proposition 2.5. *Problem* 2.1 *is solvable by finding the smallest $\beta \geqslant 0$ such that an optimal solution $E' \cup F'$ of the min-cut problem has the sparsity* (1) *equal to β. Furthermore, in that case we have $\beta = \Phi_Q$ and E' is an optimal solution of Problem* 2.1.

Proof. Suppose $\beta < \Phi_Q$. Then any Q-cut, E', satisfies $c(E') > \beta \sum_{\{s_i,v\} \text{ separated by } E'} d_{v,i}$. The sparsity (1) is then greater than β for any cut separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$.

Suppose, on the other hand, that $\beta \geqslant \Phi_Q$. By Lemma 2.3, any min-cut $E' \cup F'$ separating $Q \cup \bar{R}$ from $\bar{Q} \cup R$ has $\sum_{\{s_i,v\} \text{ separated by } E'} d_{v,i} > 0$. Note that the second terms in the numerator and denominator of (1) are the same up to the factor of β. By Lemma 2.4, for a fixed β, (1) is minimized exactly when

$$\frac{c(E')}{\sum_{\{s_i,v\} \text{ separated by } E'} d_{v,i}} \tag{2}$$

is minimized to be Φ_Q. Thus, E' is an optimal solution of Problem 2.1. Since $\beta \geqslant \Phi_Q$, this also means that $E' \cup F'$ gives the sparsity (1) $\leqslant \beta$. If, in particular, $\beta = \Phi_Q$, then it is easy to see that the sparsity (1) is equal to β. □

Such a β can be found by a binary search on an interval containing Φ_Q. Let $c_{\max}$ ($d_{\max}$) be the maximum value of an edge cost (a demand, respectively). An easy upperbound on Φ_Q is then $c_{\max}|E|$. As the denominator of any sparsity is no greater than $d_{\max} l$, an interval with length less than $1/d_{\max} l$ cannot contain more than one sparsity value. Hence, we can find Φ_Q in $\log(l|E|c_{\max} d_{\max})$ binary search queries for each $Q \subseteq S$. The total computation for the k-source sparsest cut problem is, therefore, $2^k \times$ the time for solving a parametric max-flow problem on $\tilde{G}$.

Remark 2.6. The increasing monotonicity of capacities in β of the arcs to the super-sinks of R does not satisfy the assumption requiring them monotone non-increasing, for a parametric max-flow problem to be solvable in time of a single max-flow algorithm [1]. However, the simple structure leaves an open question if a more efficient method than the straightforward binary search, a strongly polynomial method, for instance, is possible.

Acknowledgement

The authors are indebted to an anonymous referee for the excellent comments to which the paper has been intensively revised. The research was partially supported by KOSEF research fund R01-2005-000-10271-0.

References

[1] G. Gallo, M.D. Grigoriadis, R.E. Tarjan, A fast parametric maximum flow algorithm and applications, SIAM J. Comput. 18 (1) (1989) 30–55.

[2] O. Günlük, A new min-cut max-flow ratio for multi-commodity flows, in: Integer programming and combinatorial optimization, Lecture Notes in Computer Science, vol. 2337, Springer, Berlin, 2002, pp. 54–66.

[3] F.T. Leighton, S.B. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, J. ACM 46 (6) (1999) 787–832 (Preliminary version in 29th FOCS, 1988, pp. 422–431).

[4] D.W. Matula, F. Shahrokhi, Sparsest cuts and bottleneck in graphs, Discrete Appl. Math. 27 (1990) 113–123.

[5] B. A. Papernov, Feasibility of multicommodity flows, Stud. Discrete Optim. (1976) 230–261 (in Russian).

[6] A. Schrijver, Combinatorial optimization. Polyhedra and Efficiency, Algorithms and Combinatorics, vol. 24, Springer, Berlin, 2003.

[7] F. Shahrokhi, D.W. Matula, The maximum concurrent flow problem, J. Assoc. Comput. Mach. 37 (2) (1990) 318–334.

[8] D. Shmoys, Cut problems and their application to divide-and-conquer, in: D.S. Hochbaum (Ed.), Approximation Algorithms for NP-hard Problems, PWS, Boston, 1997, pp. 192–235.