Information Processing Letters 114 (2014) 365–371

Contents lists available at ScienceDirect

Information Processing Letters

www.elsevier.com/locate/ipl

On mimicking networks representing minimum terminal cuts

Arindam Khan [a,*], Prasad Raghavendra [b,**]

[a] School of Computer Science, Georgia Institute of Technology, Atlanta, USA
[b] Electrical Engineering and Computer Sciences, University of California, Berkeley, USA

ARTICLE INFO

Article history:
Received 16 July 2013
Received in revised form 17 February 2014
Accepted 20 February 2014
Available online 26 February 2014
Communicated by Ł. Kowalik

Keywords:
Algorithms
Analysis of algorithms
Approximation algorithms
Combinatorial problems
Graph algorithms

ABSTRACT

Given a capacitated undirected graph $G = (V, E)$ with a set of terminals $K \subset V$, a *mimicking network* is a smaller graph $H = (V_H, E_H)$ which contains the set of terminals K and for every bipartition $[U, K - U]$ of the terminals, the cost of the minimum cut separating U from $K - U$ in G is exactly equal to the cost of the minimum cut separating U from $K - U$ in H.

In this work, we improve both the previous known upper bound of 2^{2^k} [1] and lower bound of $(k + 1)$ [2] for mimicking networks, reducing the doubly-exponential gap between them to a single-exponential gap as follows:

- Given a graph G, we exhibit a construction of mimicking network with at most k'th Hosten–Morris number ($\approx 2^{\binom{(k-1)}{\lfloor (k-1)/2 \rfloor}}$) of vertices (independent of the size of V).
- There exist graphs with k terminals that have no mimicking network with less than $2^{\frac{k-1}{2}}$ number of vertices.

1. Introduction

Suppose that there are small number of terminals or clients that are part of a huge network such as the internet. Often, it is useful to construct a smaller graph which preserves the properties of the huge network that are relevant to the terminals. For example, if the terminals or the clients are interested in routing flows through the large network, one would want to construct a small graph preserving the routing properties of the original network. The notion of *mimicking networks* introduced by Hagerup et al. [1] is an effort in this direction.

Let G be an undirected graph with edge capacities c_e for all $e \in E$, and a set of k terminals $K (\subset V) := \{v_1, v_2, \ldots, v_k\}$. A *mimicking network* for G is an undirected capacitated graph $H = (V_H, E_H)$ such that $K \subseteq V_H$ and for each subset $U \subset K$ of terminals, the cost of the minimum cut separating U from $K - U$ in H is exactly equal to the cost of the minimum cut separating U and $K - U$ in the graph G. Let us assume G to be connected; otherwise we can consider each component separately. Here, we will use edge costs and edge capacities interchangeably. As a corollary, the set of realizable external flows (possible total flows at terminals) in G is preserved in a mimicking network. The vertices of the mimicking network that are not terminals, namely $(V_H - K)$ will be referred to as *Steiner* vertices.

The work of Hagerup et al. [1] exhibited a construction of mimicking networks with at most 2^{2^k} vertices for every graph with k terminals. Subsequently, Chaudhuri et al. [2] proved that there exist graphs that require at least $(k + 1)$ vertices in its mimicking network. The same work also obtained improved constructions of mimicking networks for

* Principal corresponding author.
** Corresponding author.
E-mail addresses: akhan67@gatech.edu (A. Khan), prasad@cs.berkeley.edu (P. Raghavendra).

http://dx.doi.org/10.1016/j.ipl.2014.02.011
0020-0190/

special classes of graphs, namely, bounded treewidth and outerplanar graphs.

Mimicking networks constituted the main building block in the development of an $O(n)$ time algorithm for computing the maximum $s-t$ flow in a bounded treewidth network [1] and for obtaining an optimal solution for the all-pairs minimum-cut problem in the same class of networks [3].

Closely tied to mimicking networks, is the more general notion of *vertex sparsifiers* [4] which only approximately preserve the cut values. While there has been progress [5, 6] in efficient constructions of vertex sparsifiers without Steiner nodes, the power of vertex sparsifiers with Steiner nodes is poorly understood [7]. The following question originally posed by Moitra [4] remains open: *Do there exist cut sparsifiers with $k^{O(1)}$ additional Steiner nodes that yield a better than $O(\log k/\log\log k)$ approximation?* In fact, Moitra [4] points out that there could exist exact cut sparsifiers with only k additional Steiner nodes.

1.1. Our results

In this paper, we show improved upper and lower bounds for mimicking networks a.k.a. vertex cut sparsifiers with quality 1.

Theorem 1. *There exist graphs with k terminals, for which every mimicking network has at least $2^{k-1}-1$ edges and $2^{(k-1)/2}$ vertices.*

Theorem 2. *For every graph G, there exists a mimicking network that has at most k'th Hosten–Morris number ($\approx 2^{\binom{(k-1)}{\lfloor (k-1)/2\rfloor}}$) of vertices.*

Related work In a concurrent and independent work, Krauthgamer et al. [8] showed a slightly weaker lower bound of $\binom{k}{k/2}$ ($< 2^{(k-1)}$) for the number of *edges* of mimicking networks. Chambers et al. had mentioned an upper bound of Dedekind number of vertices for mimicking networks, without an elaborate proof in [9]. Dedekind number is the number of antichains in the partial order $\subseteq$ induced on the subsets of a $(k-1)$-element set by containment. Hosten–Morris number is the number of *intersecting* antichains in this partial order and thus gives a slight improvement over this bound.

2. Preliminaries

In this section, we set up the notation and present formal definitions of the terms related to mimicking networks. Let $c: E \to \mathbb{R}_0^+$ be the capacity function of the graph. Let $h_G: 2^V \to \mathbb{R}_0^+$ denote the cut function of G:

$$h_G(A) = \sum_{e\in\delta(A)} c(e)$$

where $\delta(A)$ denotes the set of edges crossing the cut $[A, V \setminus A]$. Now we define the terminal cut function $h_K^G: 2^K \to \mathbb{R}_0^+$ on K as

$$h_K^G(U) = \min_{A\subset V,\, A\cap K=U} h_G(A)$$

In words, $h_K^G(U)$ is the cost of the minimum cut separating U from $K \setminus U$ in G. Let $S(U)$ be the smallest subset of V such that $h_G(S(U)) = h_K^G(U)$, $S(U) \cap K = U$ i.e., $S(U)$ is the partition containing U in the minimum terminal cut separating U from $K - U$. For any fixed $U \subset K$, the minimum cut $h_K^G(U)$ can be computed efficiently. We will sometimes abuse this notation and use $h_K^G(U)$ to denote both the size of the minimum terminal cut and the set of edges belonging to the minimum terminal cut.

Contraction of edges will be our main tool to construct mimicking networks. Note that given a graph G and an edge e whose endpoints are not both terminals, contracting the edge e in the graph G will not decrease the value of any minimum terminal cut.

Definition 1. A graph $H = (V_H, E_H)$ is a *contraction-based* mimicking network of graph $G = (V, E)$ with terminal set K if there exists a function $f: V \to V_H$ such that the edge cost function of H is defined as follows: $c_H(y,z) = \sum_{u,v|f(u)=y,f(v)=z} c(u,v)$ where $(y,z) \in E(H)$ and $(u,v) \in E(G)$.

3. Exponential lower bound

In this section we will exhibit the lower bound on the size of mimicking networks using a subtle rank argument. For a set of k terminals K, there are $2^{k-1}-1$ minimum terminal cuts. Let us enumerate these cuts by $[U_i, K \setminus U_i]$ for $i \in \{1, 2, \dots, p(=2^{k-1}-1)\}$. Fix $p = 2^{k-1}-1$ for the remainder of the section. Let $h_K^G(U_i)$ be the minimum terminal cut separating U_i from the rest of the terminals for $i \in \{1, 2, \dots, p(=2^{k-1}-1)\}$.

Definition 2. A *minimum terminal cut vector (MTCV)* $m^{G,K}$ for graph G with terminal set K, is a p-dimensional vector where i'th coordinate $m_i^{G,K} = h_K^G(U_i)$.

Let M_k be the set of all possible minimum terminal cut vectors with k terminals. Not all vectors $v \in \mathbb{R}^{2^{k-1}-1}$ can be minimum terminal cut vectors. The submodularity of the cut function introduces constraints on the coordinates of the minimum terminal cut vector. For example there are 3 possible terminal cuts for graphs with terminal set size 3. However $[\mathbf{0.1}, \mathbf{0.1}, \mathbf{0.8}]$ is not a valid MTCV. First we prove that these minimum terminal cut vectors form a convex set.

Lemma 1. *M_k is a convex cone in $\mathbb{R}^{2^{k-1}-1}$.*

Proof. Note that by scaling the edges of a graph G, the corresponding minimum terminal cut vector also scales. Therefore, it is sufficient to show the convexity of the set M_k.

Let G_1 and G_2 be graphs with terminal set K of size k. Let N_1 and N_2 be their set of non-terminals respectively i.e., $N_i \cup K = V(G_i)$ for $i = 1, 2$. Note that these graphs might have different edge weights or different number of vertices. So depending on the edge values minimum terminal cuts will have different values. Let us assume that t_1 and t_2 be the minimum terminal cut vectors for graphs G_1