

On the Inapproximability of Disjoint Paths and Minimum Steiner Forest with Bandwidth Constraints

Bin Ma¹

*Peking University, Beijing, People's Republic of China and
City University of Hong Kong, Hong Kong*

and

Lusheng Wang²

City University of Hong Kong, Hong Kong

Received July 30, 1998; revised June 3, 1999

In this paper, we study the inapproximability of several well-known optimization problems in network optimization. We show that the *max directed vertex-disjoint paths* problem cannot be approximated within ratio $2^{\log^{1-\epsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$, the *max directed edge-disjoint paths* problem cannot be approximated within ratio $2^{\log^{1-\epsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$, the *integer multicommodity flow* problem in directed graphs cannot be approximated within ratio $2^{\log^{1-\epsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$, the *max undirected vertex-disjoint paths* problem does not have a polynomial time approximation scheme unless $\text{P} = \text{NP}$, and the *minimum Steiner forest with bandwidth constraints* problem cannot be approximated within ratio $\exp(\text{poly}(n))$ unless $\text{P} = \text{NP}$. © 2000 Academic Press

1. INTRODUCTION

Disjoint paths problems are fundamental, extensively studied NP-hard problems. They have applications in areas such as telecommunications, VLSI, and scheduling [8, 13, 14, 16]. Due to the rapid growth of high-speed integrated networks that

¹ Supported in part by RGC 9040297, 9040352, and CityU Strategic 7000693. Current address: Department of Mathematics, Peking University, Beijing 100871, P.R. China and Department of Computer Science, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon, Hong Kong. E-mail: bma@sxx0.math.pku.edu.cn.

² Supported in part by RGC 9040297, 9040352, and CityU Strategic 7000693. Current address: Department of Computer Science, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon, Hong Kong. E-mail: lwang@cs.cityu.edu.hk.

provide vast bandwidth and support heterogeneous applications, considerable effort has been made recently for disjoint paths, bandwidth allocation, and related algorithmic problems. In this paper, we study the inapproximability of several well-known optimization problems including *max directed/undirected vertex/edge-disjoint paths*, *integer multicommodity flow*, and *minimum Steiner forest with bandwidth constraints*.

Maximum Disjoint Paths

Given a graph $G = (V, E)$ and a set of vertex pairs $T = \{\langle s_1, t_1 \rangle, \langle s_2, t_2 \rangle, \dots, \langle s_k, t_k \rangle\}$. The *max undirected vertex/edge-disjoint paths* problem (VDP/EDP) is to find a maximum cardinality collection \mathcal{P} of vertex/edge-disjoint paths in G , each of which connects a distinct pair $\langle s_i, t_i \rangle \in T$. If the graph is directed and the directed paths are requested, the problem is called *max directed vertex/edge-disjoint paths* (DVDP/DEDP).

All the above versions were proved to be NP-complete [6, 7]. A lot of effort has been made in design polynomial time approximation algorithms [8–11]. Most of them are about special classes of graphs, e.g., expander graphs [11], densely embedded graphs [9], high-diameter planar networks [10], and high-bandwidth models [1].

For general graphs, the best-known performance ratios are $O(\sqrt{|E_0|})$ for EDP and $O(\sqrt{|V_0|})$ for VDP, where $|E_0|$ and $|V_0|$ are the numbers of edges and vertices that appear in the paths in an optimal solution, respectively [12]. On the other hand, the only known inapproximability result is that EDP is MAX SNP-hard [2]. In [15], Srinivasan conjectures that the problem does not admit any good ratio; e.g., it cannot be approximated within ratio $O(n^\varepsilon)$ for some ε .

In this paper, we prove that DVDP cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$, DEDP cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$, and VDP does not have a polynomial time approximation scheme (PTAS) unless $\text{P} = \text{NP}$.

After this paper was accepted, the referee informed us that a stronger lower bound $m^{1/2-\varepsilon}$ was given by Guruswami *et al.* [3]. Their results are later and the techniques used are completely different.

Integer Multicommodity Flow

Given a (directed/undirected) graph $G = (V, E)$, a capacity function $c: E \rightarrow \mathbb{N}$, and k pairs of vertices (s_i, t_i) ($i = 1, 2, \dots, k$) indicating s_i as the source and t_i the sink for this commodity, a multicommodity flow is a way of simultaneously routing commodities from their sources to sinks subject to capacity and conservation constraints. The problem is to find an integral multicommodity flow such that the total flow is maximized.

This problem was proved to be MAX SNP-hard for undirected graphs [2]. For directed graphs, it has a randomized approximation algorithm when $c(e)$ is at least $5.2 \log 4 |E|$ [4]. As a corollary of our hardness result for DEDP, we prove that the integer multicommodity flow problem in directed graphs cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$.

Minimum Steiner Forest with Bandwidth Constraints

Minimum Steiner forest with bandwidth constraints is an important problem in multicast routing [5]. Let $G=(V, E)$ be a graph. Each edge $e \in E$ has a positive cost $c(e)$ and a bandwidth $b(e)$. We are given a group of processes (a set of vertices) $D \subseteq V$ and the bandwidth requirement of the group connection B ; the problem here is to find $|D|$ Steiner trees such that the total bandwidth used for each edge e is at most $b(e)$ and the total cost of the $|D|$ Steiner trees is minimized. Without loss of generality, here we assume that $B=1$. We prove that this problem cannot be approximated within ratio $\exp(\text{poly}(n))$.

2. MAXIMUM DISJOINT PATHS

In this section, we will discuss the directed and undirected edge/vertex-disjoint paths. We also consider the integer multicommodity flow problem in directed graphs.

2.1. Max Directed Disjoint Paths

In this subsection, we show that DVDP cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ for any $\varepsilon > 0$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$. A similar result also holds for DEDP.

The reduction is from the maximization version of *label cover*.

Max Label Cover (LC)

Instance. A regular bipartite graph $G=(V_1, V_2, E)$ (every node in V_i ($i=1, 2$) has degree d_i), an integer N in unary, and for each edge $e \in E$, a partial function $\Pi_e: [1, N] \rightarrow [1, N]$. A *labeling* has to assign each $v \in V_1 \cup V_2$ a number in $[1, N]$. An edge $e=(v_1, v_2) \in E$, where $v_1 \in V_1$ and $v_2 \in V_2$, is said to be *covered* if v_i is labeled with a_i ($i=1$ or 2) and $\Pi_e(a_1)=a_2$.

Question. Find a labeling that maximizes the number of covered edges.

It was shown that the problem is hard to approximate [4]. That is,

THEOREM 1 [4]. *For any $\varepsilon > 0$, max label cover cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$.*

The main theorem in this subsection is that DVDP is hard to approximate. That is,

THEOREM 2. *If there is an approximation algorithm for DVDP with ratio ρ , then one can find an approximate solution for max label cover with ratio 2ρ in polynomial time.*

Proof. Consider an instance I_{lc} of max label cover that contains a regular bipartite graph $G=(V_1, V_2, E)$, an integer N , and a partial function $\Pi_e: [1, N] \rightarrow [1, N]$ for each edge $e \in E$. Suppose that $V_1 = \{v_1, v_2, \dots, v_m\}$ and $V_2 = \{v_{m+1}, v_{m+2}, \dots, v_n\}$. The degree of nodes in V_1 is d_1 and the degree of nodes in V_2 is d_2 .




FIG. 1. The subgraph $G_1(v_i)$. Each dotted circle corresponds to a label l , where $1 \leq l \leq N$.

Let K be an integer that is related to the size of I_{lc} . (The value of K will be determined later.) In our reduction, each $v_i \in V_1$ corresponds to a directed subgraph $G_1(v_i)$ that can be constructed in a Euclidean plane as follows:

- (1) $d_1 \cdot K$ points denoted as $v_{i,1}^1, v_{i,2}^1, \dots, v_{i,d_1}^1, v_{i,1}^2, v_{i,2}^2, \dots, v_{i,d_1}^2, \dots, v_{i,1}^K, \dots, v_{i,d_1}^K$ are lined up in a horizontal line from left to right;
- (2) another $d_1 \cdot K \cdot N$ points denoted as

$$\begin{aligned}
 & w_{i,1,1}^1, \dots, w_{i,d_1,1}^1, w_{i,1,1}^2, \dots, w_{i,d_1,1}^2, \dots, w_{i,1,1}^K, \dots, w_{i,d_1,1}^K, \\
 & w_{i,1,2}^1, \dots, w_{i,d_1,2}^1, w_{i,1,2}^2, \dots, w_{i,d_1,2}^2, \dots, w_{i,1,2}^K, \dots, w_{i,d_1,2}^K, \\
 & \dots, \\
 & w_{i,1,N}^1, \dots, w_{i,d_1,N}^1, w_{i,1,N}^2, \dots, w_{i,d_1,N}^2, \dots, w_{i,1,N}^K, \dots, w_{i,d_1,N}^K,
 \end{aligned}$$

are lined up in another horizontal line from left to right (the two lines in (1) and (2) are parallel);

- (3) for each possible j, k, l , draw a directed line segment from $v_{i,j}^k$ to $w_{i,j,l}^k$ in the plane (this leads to many intersection points);

(4) the intersection points obtained in (3) as well as the points in (1) and (2) are the nodes in $G_1(v_i)$; every edge in $G_1(v_i)$ corresponds to a pair of nodes which are adjacent in some line segment in the plane and its direction is the same as the line segment.

See Fig. 1. Similarly, for each $v_i \in V_2$, a subgraph $G_2(v_i)$ can be constructed. The differences are that d_1 is replaced with d_2 and that the directions of edges are reversed.

For each node $v \in V_1 (V_2)$, we number the edges incident to v from 1 to d_1 (from 1 to d_2). Let

$$\begin{aligned}
 E_e = \{ & (w_{i_1, j_1, l_1}^k, w_{i_2, j_2, l_2}^k) \mid e = (v_{i_1}, v_{i_2}) \text{ be the } j_1 \text{th edge of } v_{i_1} \text{ and} \\
 & \text{the } j_2 \text{th edge of } v_{i_2}, \Pi_e(l_1) = l_2, 1 \leq k \leq K \}
 \end{aligned} \tag{1}$$

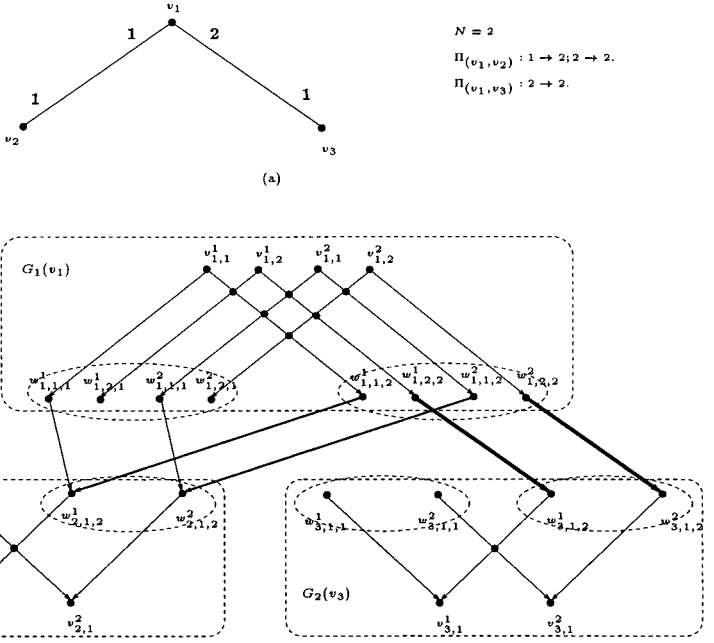


FIG. 2. (a) An instance of maximum label cover, where the bold numbers are the numbering of edges in the reduction. (b) The constructed instance of DVDP from the graph in (a), where $K=3$. Note that, in (b) there are three types of edges not in any $G_i(v_j)$ (outside the dashed rectangles). The thin edges correspond to the mapping $\Pi_{(v_1, v_2)} : 1 \rightarrow 2$ for edge (v_1, v_2) . The bolder edges correspond to the mapping $\Pi_{(v_1, v_2)} : 2 \rightarrow 2$ for edge (v_1, v_2) . The boldest edges correspond to the mapping $\Pi_{(v_1, v_3)} : 2 \rightarrow 2$ for edge (v_1, v_3) .

and

$$T_e = \{ \langle v_{i_1, j_1}^k, v_{i_2, j_2}^k \rangle \mid e = (v_{i_1}, v_{i_2}) \text{ be the } j_1 \text{ the edge of } v_{i_1} \text{ and} \\ \text{the } j_2 \text{ th edge of } v_{i_2}, 1 \leq k \leq K \}. \quad (2)$$

Now, we can define a new graph G' as follows: (1) the set of vertices for G' is the set of all the nodes in $G_1(v_i)$'s and $G_2(v_j)$'s for $v_i \in V_1$ and $v_j \in V_2$; (2) the set of edges for G' contains the edges in $\bigcup_{e \in E} E_e$ and the edges in $G_1(v_i)$'s and $G_2(v_j)$'s for $v_i \in V_1$ and $v_j \in V_2$. We define T to be the set of vertex pairs $\bigcup_{e \in E} T_e$. G' and T form an instance $I_{dv dp}$ for DVDP. An example is given in Fig. 2.

LEMMA 3. *If there is an optimal solution of I_{lc} that covers c edges then the optimal solution of $I_{dv dp}$ contains at least $c \cdot K$ vertex-disjoint paths.*

Proof. Suppose that v_i is labeled with $l(i)$ in the solution for I_{lc} . For each covered edge $e = (v_{i_1}, v_{i_2})$, there are K paths,

$$P_{i_1, i_2}^k : v_{i_1, j_1}^k \rightarrow w_{i_1, j_1, l(i_1)}^k \rightarrow w_{i_2, j_2, l(i_2)}^k \rightarrow v_{i_2, j_2}^k,$$

where $k = 1, 2, \dots, K$, e is the j_1 th edge of v_{i_1} and the j_2 th edge of v_{i_2} , the second \rightarrow is the edge $(w_{i_1, j_1, l(i_1)}^k, w_{i_2, j_2, l(i_2)}^k)$, and the first \rightarrow and the third \rightarrow are the paths corresponding to the directed line segments from v_{i_1, j_1}^k to $w_{i_1, j_1, l(i_1)}^k$ and from $w_{i_2, j_2, l(i_2)}^k$ to v_{i_2, j_2}^k , respectively. Since there are c covered edges, there are $c \cdot K$ paths. It is easy to see that the $c \cdot K$ paths are vertex-disjoint. Thus, we have a solution of $I_{dv dp}$ with $c \cdot K$ paths. Therefore, the optimal solution of $I_{dv dp}$ contains at least $c \cdot K$ vertex-disjoint paths. ■

Let c be the number of covered edges for an optimal solution of a max label cover instance I_{lc} . Let $I_{dv dp}$ be the instance constructed from I_{lc} . From Lemma 3, the optimal solution for $I_{dv dp}$ contains at least $c \cdot K$ paths. Thus, the approximation algorithm gives a solution \mathcal{P} for $I_{dv dp}$ connecting more than cK/ρ pairs. To prove the theorem, we only have to show that one can find a solution of I_{lc} that covers at least $c/(2\rho)$ edges in polynomial time based on \mathcal{P} .

CLAIM 4. *For each path $p \in \mathcal{P}$ connecting the pair $(v_{i_1, j_1}^k, v_{i_2, j_2}^k)$, p can be divided into three segments: (i) $v_{i_1, j_2}^k \rightarrow w_{i_1, j_1, l_1}^{k'}$; (ii) the edge $(w_{i_1, j_1, l_1}^{k'}, w_{i_2, j_2, l_2}^{k'})$; and (iii) $w_{i_2, j_2, l_2}^{k'} \rightarrow v_{i_2, j_2}^k$.*

Proof. Note that for any i such that $1 \leq i \leq n$ and $a = 1, 2$, the subgraph $G_a(v_i)$ connects the other parts of G' through nodes $w_{i, j, l}^k$'s for $1 \leq k \leq K$. Since p connects the two components $G_1(v_{i_1})$ and $G_2(v_{i_2})$, p contains two nodes $w_{i_1, j, l_1}^{k'}$ and $w_{i_2, j', l_2}^{k''}$. Thus, p is divided into three segments by $w_{i_1, j, l_1}^{k'}$ and $w_{i_2, j', l_2}^{k''}$.

From the construction of the directed graph G' , the segment between $w_{i_1, j, l_1}^{k'}$ and $w_{i_2, j', l_2}^{k''}$ is the edge $(w_{i_1, j, l_1}^{k'}, w_{i_2, j', l_2}^{k''})$. From formula (1), $k' = k''$ and the edge (v_{i_1}, v_{i_2}) in I_{lc} is the j th edge of v_{i_1} and j' th edge of v_{i_2} . From formula (2) and the fact that $\langle v_{i_1, j_1}^k, v_{i_2, j_2}^k \rangle \in T$, (v_{i_1}, v_{i_2}) is the j_1 th edge of v_{i_1} and the j_2 th edge of v_{i_2} . Thus, $j = j_1$ and $j' = j_2$. This completes the proof. ■

For each path $p: v_{i_1, j_1}^k \rightarrow w_{i_1, j_1, l_1}^{k'} \rightarrow w_{i_2, j_2, l_2}^{k'} \rightarrow v_{i_2, j_2}^k$ in \mathcal{P} connecting $\langle v_{i_1, j_1}^k, v_{i_2, j_2}^k \rangle$, denote $i_1(p)$, $i_2(p)$, $l_1(p)$, and $l_2(p)$ to be the subscripts i_1 , i_2 , l_1 , and l_2 , respectively. We divide the paths in \mathcal{P} into $m \cdot K$ groups \mathcal{P}_i^k , where $i = 1, 2, \dots, m$, $k = 1, 2, \dots, K$, and

$$\mathcal{P}_i^k = \{p \mid p \text{ starts at some } v_{i, j}^k\}.$$

Define

$$\mathcal{P}^k = \bigcup_{i=1}^m \mathcal{P}_i^k. \quad (3)$$

A group \mathcal{P}_i^k is well defined if for any two paths p and p' in \mathcal{P}_i^k , $l_1(p) = l_1(p')$. A group \mathcal{P}^k is well defined if for any two paths p and p' in \mathcal{P}^k , $i_1(p) = i_1(p')$ implies $l_1(p) = l_1(p')$ and $i_2(p) = i_2(p')$ implies $l_2(p) = l_2(p')$. Well-defined \mathcal{P}^k 's are important to our proof since each well-defined \mathcal{P}^k induces a solution for the instance of the max label cover problem.

CLAIM 5. *If \mathcal{P}^k is well defined and \mathcal{P}^k contains h paths, one can get a solution for I_{lc} that covers h edges.*

Proof. Suppose that \mathcal{P}^k is well defined. For every path p in \mathcal{P}^k , labeling $v_{i_1(p)} \in V_1$ with $l_1(p)$ and $v_{i_2(p)} \in V_2$ with $l_2(p)$ one can get a solution for I_{lc} . From Claim 4, the edge $(w_{i_1(p), j_1}^{k'}, w_{i_2(p), j_2}^{k'})$ is in p . So, formula (1) indicates that $\Pi_{(v_{i_1(p)}, v_{i_2(p)})}(l_1(p)) = l_2(p)$. Thus, the labeling covers the set of edges $\{(v_{i_1(p)}, v_{i_2(p)}) \mid p \in \mathcal{P}^k\}$ in I_{lc} . Furthermore, if p and p' are two different paths in \mathcal{P}^k , then they connect different pairs $\langle v_{i_1(p), j_1}^k, v_{i_2(p), j_2}^k \rangle$ and $\langle v_{i_1(p'), j_1'}^k, v_{i_2(p'), j_2'}^k \rangle$. Consider formula (2). The pair (v_{i_1}, v_{i_2}) uniquely determines both j_1 and j_2 . Thus, the fact that $\langle v_{i_1(p), j_1}^k, v_{i_2(p), j_2}^k \rangle$ and $\langle v_{i_1(p'), j_1'}^k, v_{i_2(p'), j_2'}^k \rangle$ are different implies either $i_1(p) \neq i_1(p')$ or $i_2(p) \neq i_2(p')$. So, different paths in \mathcal{P}^k correspond to different covered edges of I_{lc} . Therefore, we get a labeling that contains at least h covered edges. ■

Unfortunately, a solution for I_{dvdp} cannot ensure that every \mathcal{P}^k is well defined. However, we can show that *lots* of \mathcal{P}^k 's are well defined and thus there exists a \mathcal{P}^k containing *enough* number of paths.

CLAIM 6. *Among the K groups $\mathcal{P}^1, \mathcal{P}^2, \dots, \mathcal{P}^K$, there are at most $n \times N$ groups that are not well defined.*

Proof. Fix a node $v_i \in V_1$. Consider two paths p and p' in \mathcal{P} starting at v_{ij}^k and $v_{ij'}^{k'}$, respectively. Note that p and p' are vertex-disjoint and the subgraph $G_1(v_i)$ is constructed in the plane. Thus, if $k < k'$, then $l_1(p) \leq l_1(p')$. For example, if p starts at $v_{i,2}^1$, $l_1(p) = 2$ (See Fig. 1, the dark path.), and if p' starts at $v_{i,1}^2$, then $l_1(p')$ is at least 2. Otherwise, p and p' are not vertex-disjoint.

For a fixed i , $l_1(p)$'s can take at most N values. Thus, there are at most N $\mathcal{P}_i^{k'}$'s for a fixed i that are not well defined. Applying a similar argument to $v_i \in V_2$ and counting all $v_i \in V_1 \cup V_2$, we can conclude that there are at most $n \times N$ \mathcal{P}^k 's that are not well defined. ■

It is easy to see that $|\mathcal{P}^k| \leq d_1 m$. From Claim 6, the number of paths in \mathcal{P} but not in any well-defined \mathcal{P}^k is at most $nNd_1 m$. Thus, the average number of paths in each well defined \mathcal{P}^k is at least $(cK/\rho - nNd_1 m)/K$. Setting $K = \lceil 2\rho nNd_1 m \rceil$, we have $(cK/\rho - nNd_1 m)/K \geq c/(2\rho)$. So, at least one of the well-defined \mathcal{P}^k 's contains at least $c/(2\rho)$ paths.

From Claim 5, by choosing a well-defined group \mathcal{P}^k which contains the largest number of paths, we can get a labeling that contains at least $c/(2\rho)$ covered edges. ■

THEOREM 7. *For any $\varepsilon > 0$, DVDP cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$.*

Proof. From Theorem 2, if there is an approximation algorithm for DVDP with ratio $\rho = 2^{\log^{1-\varepsilon} n}$, then there is an approximation algorithm for the max label cover problem with ratio $2\rho = 2^{\log^{1-\varepsilon'} n}$ for some small positive number ε' . This contradicts Theorem 1. ■

THEOREM 8. *For any $\varepsilon > 0$, DEDP cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$.*




FIG. 3. A node v_i is replaced with two nodes v'_i and v''_i .

Proof. We reduce DVDP to DEDP. Given an instance $I_{dv dp} = \langle G, T \rangle$ of DVDP, where $G = (V, E)$ is a directed graph and $T = \{ \langle v_{i_1}, v_{j_1} \rangle, \dots, \langle v_{i_m}, v_{j_m} \rangle \}$ is a set of vertex pairs. For each node $v_i \in V$, we introduce two new nodes v'_i and v''_i . Define $V' = \{ v'_i, v''_i \mid v_i \in V \}$, $E' = \{ (v'_i, v''_i) \mid v_i \in V \} \cup \{ (v'_i, v'_j) \mid (v_i, v_j) \in E \}$, $G' = (V', E')$, and $T' = \{ \langle v'_{i_1}, v'_{j_1} \rangle, \dots, \langle v'_{i_m}, v'_{j_m} \rangle \}$. We then have an instance $I_{dedp} = \langle G', T' \rangle$ of DEDP. Figure 3 gives part of the reduction.

Now we want to show that $I_{dv dp}$ has a solution connecting K pairs in T if and only if I_{dedp} has a solution connecting K pairs in T' .

(if) Suppose \mathcal{P} is a solution of $I_{dv dp}$. For each $p \in \mathcal{P}$, if p connects $\langle v_{i_k}, v_{j_k} \rangle$ and $p = v_{t_1}, v_{t_2}, \dots, v_{t_l}$, where $v_{t_1} = v_{i_k}$ and $v_{t_l} = v_{j_k}$, one can construct a path $p' = v'_{t_1}, v''_{t_1}, v'_{t_2}, v''_{t_2}, \dots, v'_{t_l}, v''_{t_l}$ in G' connecting the pair $\langle v'_{i_k}, v'_{j_k} \rangle \in T'$. Since paths p in G is vertex-disjoint, it is easy to see that p' 's in G' are edge-disjoint.

(only if) Suppose \mathcal{P}' is a solution of I_{dedp} . For each path $p' \in \mathcal{P}'$, if p' connects the pair $\langle v'_{i_k}, v'_{j_k} \rangle$, then the construction of G' ensures that p' has the form $v'_{t_1}, v''_{t_1}, v'_{t_2}, v''_{t_2}, \dots, v'_{t_l}, v''_{t_l}$, where $v'_{t_1} = v'_{i_k}$ and $v''_{t_l} = v''_{j_k}$. Then $p = v_{t_1}, v_{t_2}, \dots, v_{t_l}$ is a path connecting the pair $\langle v_{i_k}, v_{j_k} \rangle$ in G . Since p' 's are edge-disjoint in G' , p 's are also vertex-disjoint in G .

From Theorem 7, we know Theorem 8 is true. ■

COROLLARY 9. For any $\varepsilon > 0$, the integer multicommodity flow problem in directed graphs cannot be approximated within ratio $2^{\log^{1-\varepsilon} n}$ unless $\text{NP} \subseteq \text{DTIME}[2^{\text{polylog } n}]$.

Proof. Given an instance of DEDP with graph $G = (V, E)$ and the set of vertex pairs $T = \{ \langle s_1, t_1 \rangle, \langle s_2, t_2 \rangle, \dots, \langle s_k, t_k \rangle \}$, we construct a new graph G' from G . For each pair (s_i, t_i) in T , we introduce two new nodes s'_i and t'_i . Define $V' = V \cup \{ s'_i, t'_i \mid 1 \leq i \leq k \}$, $E' = E \cup \{ (s'_i, s_i) \mid 1 \leq i \leq k \} \cup \{ (t_i, t'_i) \mid 1 \leq i \leq k \}$, and $G' = (V', E')$. Obviously, G is a subgraph of G' . Let $c(e) = 1$ for each $e \in E'$ and the k pairs of the commodity sources and sinks be (s'_i, t'_i) , $i = 1, 2, \dots, k$. Then we have an instance of the integer multicommodity flow problem.

Since there is only one edge out from each s'_i and $c(e) = 1$ for any e , an integer flow in G' with total capacity f is exactly f edge-disjoint paths connecting f pairs of (s'_i, t'_i) . Thus, one can easily find f paths connecting $\langle s_i, t_i \rangle$'s in G . From Theorem 8, we have this corollary. ■

2.2. Max Undirected Edge/Vertex-Disjoint Paths

It is proved in [2] that the integer multicommodity flow problem in undirected graphs is MAX SNP-hard. The same reduction implies that the max undirected




FIG. 4. A node v_i is changed to k nodes $v_{i,1}, \dots, v_{i,k}$. An edge e is changed to a node v_e .

edge-disjoint paths problem (EDP) is also MAX SNP-hard. In this subsection, we prove that

THEOREM 10. *The max undirected vertex-disjoint paths problem (VDP) is MAX SNP-hard.*

Proof. We give an L-reduction from EDP to VDP. Suppose that we have an instance I_1 of EDP containing a graph $G=(V, E)$ and a set of vertex pairs $T = \{\langle v_{s_r}, v_{t_r} \rangle \mid 1 \leq r \leq k\}$. A new graph $G'=(V', E')$ is constructed as follows:

$$V' = \{v_{i,1}, v_{i,2}, \dots, v_{i,k} \mid v_i \in V\} \cup \{v_e \mid e \in E\}$$

$$E' = \{(v_{i,j}, v_e) \mid v_i \text{ is adjacent to } e, 1 \leq j \leq k\}.$$

Figure 4 gives part of the construction. For each pair $\langle v_{s_r}, v_{t_r} \rangle$, where $r = 1, 2, \dots, k$, we construct a new pair $\langle v_{s_r,r}, v_{t_r,r} \rangle$. Let T' be the set of k new pairs. Then $G'=(V', E')$ and T' form an instance I_2 of VDP.

It is easy to see that I_1 has a solution connecting c pairs if and only if I_2 has a solution connecting c pairs. Thus, we have an L-reduction. ■

3. MINIMUM STEINER FOREST WITH BANDWIDTH CONSTRAINTS

Minimum Steiner forest with bandwidth constraints is an important problem in multicast routing [5].

Minimum Steiner Forest with Bandwidth Constraints

Let $G=(V, E)$ be a graph. Each edge $e \in E$ has a positive cost $c(e)$ and a bandwidth $b(e)$. We are given a group of processes (a set of regular points) $D \subseteq V$ and the bandwidth requirement of the group connection B (without loss of generality, assume $B = 1$); the problem is to find $|D|$ Steiner trees such that the total bandwidth used for each edge e is at most $b(e)$, and the total cost of the $|D|$ Steiner trees is minimized.

First, we discuss the decision version of the problem.

Steiner Forest with Bandwidth Constraints

Given a graph $G=(V, E)$, a vertex set $D \subset V$ with $|D| = K$, we want to find K Steiner trees for the set of regular points D such that any two Steiner trees do not share any common edge.

THEOREM 11. *Steiner forest with bandwidth constraints is NP-hard.*

Proof. The reduction is from undirected edge-disjoint paths that was shown to be NP-complete [6].

Undirected Edge Disjoint Paths (UEDP)

Given a graph G and a set T of $K \geq 3$ vertex pairs, decide whether there are K edge-disjoint paths, one for each vertex pair in T .

Given an instance of undirected edge-disjoint paths G and T , where $G = (V, E)$ is a graph and $T = \langle u_i, v_i \rangle \mid i = 1, 2, \dots, K \rangle$ is a set of vertex pairs, we define a new graph $G' = (V', E')$ as follows:

$$\begin{aligned} V' &= V \cup \{u'_i, v'_i \mid 1 \leq i \leq K\}, \quad \text{and} \\ E' &= E \cup \{(u'_i, v'_j) \mid i \neq j, 1 \leq i, j \leq K\} \cup \{(u_i, u'_i) \mid 1 \leq i \leq K\} \\ &\quad \cup \{(v'_i, v_i) \mid 1 \leq i \leq K\}. \end{aligned}$$

Let $D = \{u'_i \mid 1 \leq i \leq K\}$ be the set of regular points. Then $G' = (V', E')$ and D form an instance of Steiner forest with bandwidth constraints. Figure 5 gives an example of the construction when $K = 3$.

If the instance of undirected edge-disjoint paths has a solution $\{p_i \mid p_i \text{ connects } \langle u_i, v_i \rangle, i = 1, 2, \dots, K\}$, then the instance of Steiner forest with bandwidth constraints has a solution

$$p_i \cup \{(u'_i, u_i), (v_i, v'_i)\} \cup \{(v'_i, u'_j) \mid j \neq i, 1 \leq j \leq K\}, \quad i = 1, 2, \dots, K.$$

Suppose that the instance of Steiner forest with bandwidth constraints has a solution T_1, T_2, \dots, T_K . Since the degree of each u'_i is exactly K , we have:

FACT 1. *Every u'_j must be a leaf in each T_i .*

Now, we want to show that

FACT 2. *Each T_i contains exactly one of the edges (v_j, v'_j) .*




FIG. 5. A construction with $K = 3$ in Theorem 11.

Proof. Suppose that a Steiner tree T_1 does not contain any (v_j, v'_j) . The regular points in $D = \{u'_i \mid 1 \leq i \leq K\}$ can be divided into two disjoint subsets $S_1 = \{u'_i \mid (u'_i, u_i) \in T_1\}$ and $S_2 = \{u'_i \mid (u'_i, v'_j) \in T_1 \text{ for some } j\}$.

Since no (v_j, v'_j) is in T_1 , these two subsets S_1 and S_2 cannot be connected in T_1 . So, one of them is empty. If S_1 is empty, then T_1 contains neither (u_j, u'_j) nor (v_j, v'_j) . In other words, T_1 does not contain any edge in E . Combining with Fact 1, T_1 cannot be connected. The only remaining case is that $S_1 = D$. In this case, T_1 contains every (u_j, u'_j) . Thus, every other Steiner tree contains no (u_j, u'_j) . Therefore, each of the $K-1$ other Steiner trees contains at least two (v_j, v'_j) -type edges. (Otherwise, Fact 1 implies that v'_j 's are isolated. Therefore, only one v'_j is in a Steiner tree and thus u_j cannot be included in the Steiner tree.) However, there are totally K such edges (v_j, v'_j) , where $j = 1, 2, \dots, K$. This is a contradiction. Thus, each T_i contains at least one of the (v_j, v'_j) -type edges. Since there are K such edges, each T_i contains exactly one (v_j, v'_j) . ■

Without loss of generality, assume that T_j contains (v_j, v'_j) . Consider the path connecting u'_j and v'_j in T_j ; it must contain a subpath from u_j to v_j in G . All these subpaths give a solution of the original problem. ■

THEOREM 12. *Minimum Steiner forest with bandwidth constraints (MSF) cannot be approximated in ratio $\exp(\text{poly}(n))$ unless $\text{P} = \text{NP}$.*

Proof. Suppose that there is a polynomial time approximation algorithm for MSF within ratio $\rho \geq 1$, where $\log(\rho)$ might be a polynomial function of the instance size. For any instance $\langle G, D \rangle$ of Steiner forest with bandwidth constraints, where $G = (V, E)$ is a graph and D is a set of regular points, we extend G to a weighted complete graph G' by adding edges with weight ρn^2 , where n is the number of vertices in V , and assigning weight 1 to the edges in E .

If $\langle G, D \rangle$ has a solution for Steiner forest with bandwidth constraints, then $\langle G', D \rangle$ has a solution for minimum Steiner forest with bandwidth constraints with cost at most $n(n-1)$. If $\langle G, D \rangle$ does not have a solution, the optimal solution of $\langle G', D \rangle$ must contain an edge with weight ρn^2 ; i.e., the cost is at least ρn^2 .

Applying the ratio- ρ approximation algorithm to solve G' and testing if the cost of the obtained solution is less than ρn^2 will give a polynomial time algorithm to Steiner forest with bandwidth constraints. This contradicts Theorem 11. ■

ACKNOWLEDGMENTS

We thank the referees for their helpful suggestions. Also, we thank a referee for bringing the work of Guruswami *et al.* [3] to our attention.

REFERENCES

1. B. Awerbuch, Y. Azar, and S. Plotkin, Throughput-competitive online routing, in "Proc. 34th IEEE Symposium on Foundation of Computer Science, 1993," pp. 32–40.
2. N. Garg, V. Vazirani, and M. Yannakakis, Primal-dual approximation algorithms for integral flow and multicut in trees, with applications to matching and set cover, in "Proc. International Colloq.

- on Automata, Languages, and Programming,” Lecture Notes in Computer Science, Vol. 700, pp. 64–75, Springer-Verlag, Berlin, 1993.
3. V. Guruswami, S. Khanna, R. Rajaraman, B. Shepherd, and M. Yannakakis, Near-optimal hardness results and approximation algorithms for edge-disjoint paths and related problems, in “Proc. 31st ACM Symposium on Theory of Computing, 1999,” pp. 19–28.
 4. D. S. Hochbaum, “Approximation Algorithms for NP-hard Problems,” PWS, Kent/Boston, 1997.
 5. X. Jia and L. Wang, A group multicast routing algorithm by using multiple minimum Steiner trees, *Comput. Commun.* **20** (1997), 750–758.
 6. R. M. Karp, Reducibility among combinatorial problems, in “Complexity of Computer Computations” (R. E. Miller and J. W. Thatcher, Eds.), Plenum Press, New York, 1972.
 7. R. M. Karp, On the computational complexity of combinatorial problems, *Networks* **5** (1975), 45–68.
 8. J. Kleinberg, “Approximation Algorithms for Disjoint Paths Problems,” Ph.D. Thesis, Department of EECS, MIT, 1996.
 9. J. Kleinberg and É. Tardos, Disjoint paths in densely embedded graphs, in “Proc. 36th IEEE Symposium on Foundations of Computer Science, 1995,” pp. 52–61.
 10. J. Kleinberg and É. Tardos, Approximations for the disjoint paths problem in high-diameter planar networks, in “Proc. 27th ACM Symposium on Theory of Computing, 1995,” pp. 26–35.
 11. J. Kleinberg and R. Rubinfeld, Short paths in expander graphs, in “Proc. 37th IEEE Symposium on Foundations of Computer Science, 1996,” pp. 86–95.
 12. S. G. Kolliopoulos and C. Stein, Approximating disjoint-path problems using greedy algorithms and packing integer programs, in “Proc. 6th Conference on Integer Programming and Combinatorial Optimization, 1998,” pp. 153–168.
 13. M. E. Kramer and J. van Leeuwen, The complexity of wire routing and finding the minimum area layouts for arbitrary VLSI circuits, in “Advances in Computing Research 2: VLSI Theory” (F. P. Preparata, Ed.), JAI Press, London, 1984.
 14. D. B. Shmoys, C. Stein, and J. Wein, Improved approximation algorithms for shop scheduling problems, *SIAM J. Comput.* **23** (1994), 617–632.
 15. A. Srinivasan, Improved approximations for edge-disjoint paths, unsplittable flow and related routing problems, in “Proc. 38th Ann. Symposium on Foundations of Computer Science, 1997,” pp. 416–425.
 16. C. Stein, “Approximation Algorithms for Multicommodity Flow and Shop Scheduling Problems,” Ph.D. Thesis, MIT, Cambridge, MA, August 1992.