Hardness and Approximation Results for Packing Steiner Trees

Joseph Cheriyan⋆[1] and Mohammad R. Salavatipour⋆⋆[2]

[1] Department of Combinatorics and Optimization, University of Waterloo, Waterloo, Ontario N2L3G1, Canada
jcheriyan@math.uwaterloo.ca
[2] Department of Computing Science, University of Alberta, Edmonton, Alberta T6G2H1, Canada
mreza@cs.ualberta.ca

Abstract. We study approximation algorithms and hardness of approximation for several versions of the problem of packing Steiner trees. For packing edge-disjoint Steiner trees of undirected graphs, we show APX-hardness for 4 terminals. For packing Steiner-node-disjoint Steiner trees of undirected graphs, we show a logarithmic hardness result, and give an approximation guarantee of $O(\sqrt{n}\log n)$, where n denotes the number of nodes. For the directed setting (packing edge-disjoint Steiner trees of directed graphs), we show a hardness result of $\Omega(m^{\frac{1}{3}-\epsilon})$ and give an approximation guarantee of $O(m^{\frac{1}{2}+\epsilon})$, where m denotes the number of edges. The paper has several other results.

1 Introduction

We study approximation algorithms and hardness (of approximation) for several versions of the problem of packing Steiner trees. Given an undirected graph $G = (V, E)$ and a set of *terminal* nodes $T \subseteq V$, a *Steiner tree* is a connected, acyclic subgraph that contains all the terminal nodes (nonterminal nodes, which are called *Steiner nodes*, are optional). The basic problem of Packing Edge-disjoint Undirected Steiner trees (**PEU** for short) is to find as many edge-disjoint Steiner trees as possible. Besides PEU, we study some other versions (see below for details).

The PEU problem in its full generality has applications in VLSI circuit design (e.g., see [11, 21]). Other applications include multicasting in wireless networks (see [7]) and broadcasting large data streams (such as videos) over the Internet (see [14]). There is significant motivation from the areas of graph theory and combinatorial optimization. Menger's theorem on packing edge-disjoint s, t-paths [5] corresponds to the special case of packing edge-disjoint Steiner trees on two

⋆ Supported by NSERC grant No. OGP0138432.
⋆⋆ Research done while a postdoc at the Dept. of Combinatorics and Optimization, University of Waterloo. Supported by an NSERC postdoctoral fellowship and the Dept. of Comb. and Opt., University of Waterloo.

terminal nodes (i.e., $T = \{s, t\}$). Another special case is when all the nodes are terminals (i.e., $T = V$). Then the problem is to find a maximum set of edge-disjoint spanning trees. This topic was studied in the 1960's by graph theorists, and a min-max theorem was developed by Tutte and independently by Nash-Williams [5]. Subsequently, Edmonds and Nash-Williams derived such results in the more general setting of the matroid intersection theorem. One consequence is that efficient algorithms are available via the matroid intersection algorithm for the case of $T = V$. (Note that most problems on packing Steiner trees are NP-hard, so results from matroid optimization do not apply directly.) A set of nodes S is said to be λ-*edge connected* if there exist λ edge-disjoint paths between every two nodes of S. An easy corollary of the min-max theorem is that if the node set V is $2k$-edge connected, then the graph has k edge-disjoint spanning trees. Recently, Kriesell [19] conjectured an exciting generalization: If the set of terminals is $2k$-edge connected, then there exist k edge-disjoint Steiner trees. He proved this for Eulerian graphs (by an easy application of the splitting-off theorem). Note that a constructive proof of this conjecture may give a 2-approximation algorithm for PEU. Jain, Mahdian, and Salavatipour [14] gave an approximation algorithm with guarantee (roughly) $\frac{|T|}{4}$. Moreover, using a versatile and powerful proof technique (that we will borrow and apply in the design of our algorithms), they showed that the fractional version of PEU has an α-approximation algorithm if and only if the minimum-weight Steiner tree problem has an α-approximation algorithm (Theorem 4.1 in [14]). The latter problem is well studied and is known to be APX-hard. It follows that PEU is APX-hard (Corollary 4.3 in [14]). Kaski [16] showed that the problem of finding two edge-disjoint Steiner trees is NP-hard, and moreover, problem PEU is NP-hard even if the number of terminals is 7. Frank et al. [8] gave a 3-approximation algorithm for a special case of PEU (where no two Steiner nodes are adjacent). Very recently Lau [20], based on the result of Frank et al. [8], has given an $O(1)$ approximation algorithm for PEU (but Kriesell's conjecture remains open).

Using the fact that PEU is APX-hard, Floréen et al. [7] showed that packing Steiner-node-disjoint Steiner trees (see problem PVU defined below) is APX-hard. They raised the question whether this problem is in the class APX. Also, they showed that the special case of the problem with 4 terminal nodes is NP-hard.

Our Results:

We use n and m to denote the number of nodes and edges, respectively. The underlying assumption for most of our hardness results is P$\neq$NP.

• *Packing Edge Disjoint Undirected Steiner Trees (***PEU***)*: For this setting, we show that the maximization problem with only four terminal nodes is APX-hard. This result appears in Section 3. (An early draft of our paper also proved, independently of [16], that finding two edge-disjoint Steiner trees is NP-hard.)

• *Packing Vertex Capacitated Undirected Steiner Trees (***PVCU***):* We are given an undirected graph G, a set $T \subseteq V$ of terminals, and a positive vertex capacity c_v for each Steiner vertex v. The goal is to find the maximum number of Steiner trees such that the total number of trees containing each Steiner vertex v is at most c_v. The special case where all the capacities are 1 is the problem of *Packing Vertex Disjoint Undirected Steiner Trees* (**PVU** for short). Note that here and elsewhere we use "vertex disjoint Steiner trees" to mean trees that are disjoint on the *Steiner nodes* (and of course they contain all the terminals). We show essentially the same hardness results for PVU as for PEU, that is, finding two vertex disjoint Steiner trees is NP-hard and the maximization problem for a constant number of terminals is APX-hard. For an arbitrary number of terminals, we prove an $\Omega(\log n)$-hardness result (lower bound) for PVU, and give an approximation guarantee (upper bound) of $O(\sqrt{n}\log n)$ for PVCU, by an LP-based rounding algorithm. This shows that PVU is significantly harder than PEU, and this settles (in the negative) an open question of Floréen et al. [7] (mentioned above). These results appear in Section 3. Although the gap for PVU between our hardness result and the approximation guarantee is large, we tighten the gap for another natural generalization of PVU, namely *Packing Vertex Capacitated Priority Steiner Trees* (PVCU-priority for short), which is motivated by the Quality of Service in network design problems (see [4] for applications). For this priority version, we show a lower-bound of $\Omega(n^{\frac{1}{3}-\epsilon})$ on the approximation guarantee; moreover, our approximation algorithm for PVCU extends to PVCU-priority to give a guarantee of $O(n^{\frac{1}{2}+\epsilon})$; see Subsection 3.1. (Throughout, we use ϵ to denote any positive real number.)

• *Packing Edge/Vertex Capacitated Directed Steiner Trees (***PECD** *and* **PVCD***):* Consider a directed graph $G(V, E)$ with a positive capacity c_e for each edge e, a set $T \subseteq V$ of terminals, and a specified root vertex r, $r \in T$. A *directed Steiner tree* rooted at r is a subgraph of G that contains a directed path from r to t, for each terminal $t \in T$. In the problem of Packing Edge Capacitated Directed Steiner trees (**PECD**) the goal is to find the maximum number of directed Steiner trees rooted at r such that for every edge $e \in E$ the total number of trees containing e is at most c_e. We prove an $\Omega(m^{\frac{1}{3}-\epsilon})$-hardness result even for unit-capacity PECD (i.e., packing edge-disjoint directed Steiner trees), and also provide an approximation algorithm with a guarantee of $O(m^{\frac{1}{2}+\epsilon})$. Moreover, we show the NP-hardness of the problem of finding two edge-disjoint directed Steiner trees with only three terminal nodes. We also consider the problem of Packing Vertex Capacitated Directed Steiner trees (**PVCD**), where instead of capacities on the edges we have a capacity c_v on every Steiner node v, and the goal is to find the maximum number of directed Steiner trees such that the number of trees containing any Steiner node v is at most c_v. For directed graphs, PECD and PVCD are quite similar and we get the following results on the approximation guarantee for the latter problem: a lower-bound of $\Omega(n^{\frac{1}{3}-\epsilon})$ (even for the special case of unit capacities), and an upper bound of $O(n^{\frac{1}{2}+\epsilon})$. These results appear in Section 2.

In summary, with the exception of PVCU, the approximation guarantees (upper bounds) and hardness results (lower bounds) presented in this paper and the previous works [7, 14, 16, 20] are within the same class with respect to the classification in Table 10.2 in [1].

Comments:

Several of our proof techniques are inspired by results for disjoint-paths problems in the papers by Guruswami et al. [9], Baveja and Srinivasan [2], and Kolliopoulos and Stein [17]. (In these problems, we are given a graph and a set of source-sink pairs, and the goal is to find a maximum set of edge/node disjoint source-sink paths.) To the best of our knowledge, there is no direct relation between Steiner tree packing problems and disjoint-paths problems – neither problem is a special case of the other one. (In both problems, increasing the number of terminals seems to increase the difficulty for approximation, but each Steiner tree has to connect together all the terminals, whereas in the disjoint-paths problems the goal is to connect as many source-sink pairs as possible by disjoint paths.) There is a common generalization of both these problems, namely, the problem of packing Steiner trees with different terminal sets (given ℓ sets of terminals $T_1, T_2, \ldots, T_\ell$, ℓ polynomial in n, find a maximum set of edge-disjoint, or Steiner-node-disjoint, Steiner trees where each tree contains one of the terminal sets T_i $(1 \leq i \leq \ell)$). We chose to focus on our special cases (with identical terminal sets) because our hardness results are of independent interest, and moreover, the best approximation guarantees for the special cases may not extend to the general problems.

For the problems PVCU, PVCU-priority, PECD, and PVCD, we are not aware of any previous results on approximation algorithms or hardness results other than [7], although there is extensive literature on approximation algorithms for the corresponding *minimum-weight* Steiner tree problems (e.g., [3] for minimum-weight directed Steiner trees and [12] for minimum-node-weighted Steiner trees).

Due to space constraints, most of our proofs are deferred to the full version of the paper. Very recently, we have obtained new results that substantially decrease the gap between the upper bounds and the lower bounds for PVU; these results will appear elsewhere.

2 Packing Directed Steiner Trees

In this section, we study the problem of packing directed steiner trees. We show that Packing Vertex Capacitated Directed Steiner trees (PVCD) and the edge capacitated version (PECD) are as hard as each other in terms of approximability (similarly for the unit capacity cases). Then we present the hardness results for PECD with unit capacities (i.e., edge-disjoint directed case), which immediately implies similar hardness results for PVCD with unit capacities (i.e., vertex-disjoint directed case). We also present an approximation algorithm for PECD which implies a similar approximation algorithm for PVCD. The proof of the following theorem is easy. The idea for the first direction is to take the line graph, and for the second one is to split every vertex into two adjacent vertices.

Theorem 1. *Given an instance* $I = (G(V, E), T \subseteq V, k)$ *of PECD (of PVCD), there is an instance* $I' = (G'(V', E'), T' \subseteq V, k)$ *of PVCD (of PECD) with*

$|G'| = poly(|G|)$, such that I has k directed Steiner trees satisfying the capacities of the edges (vertices) if and only if I' has k directed Steiner trees satisfying the capacities of the vertices (edges).

2.1 Hardness results

First we prove that PVCD with unit capacities is NP-hard even in the simplest non-trivial case where there are only three terminals (one root r and two other terminals) and we are asked to find only 2 vertex-disjoint Steiner trees. The problem becomes trivially easy if any of these two conditions is tighter, i.e., if the number of terminals is reduced to 2 or the number of Steiner trees that we have to find is reduced to 1. If the number of terminals is arbitrary, then we show that PVCD with unit capacities is NP-hard to approximate within a guarantee of $O(n^{\frac{1}{3}-\epsilon})$ for any $\epsilon > 0$. The proof is relatively simple and does *not* rely on the PCP theorem. Also, as we mentioned before, both of these hardness results carry over to PECD. For both reductions, we use the following well-known NP-hard problem (see [10]):
PROBLEM: **2DIRPATH:**
INSTANCE: A directed graph $G(V, E)$, distinct vertices $x_1, y_1, x_2, y_2 \in V$.
QUESTION: Are there two vertex-disjoint directed paths, one from x_1 to y_1 and the other from x_2 to y_2 in G?

Theorem 2. *Given an instance I of PVCD with unit capacities and only three terminals, it is NP-hard to decide if it has 2 vertex-disjoint directed Steiner trees.*

From Theorems 1 and 2, it follows that:

Theorem 3. *Given an instance I of PECD with unit capacities and only three terminals, it is NP-hard to decide if it has 2 edge-disjoint directed Steiner trees.*

Now we show that, unless P=NP, any approximation algorithm for PVCD with unit capacities has a guarantee of $\Omega(n^{\frac{1}{3}-\epsilon})$.

Theorem 4. *Given an instance of PVCD with unit capacities, it is NP-hard to approximate the solution within $O(n^{\frac{1}{3}-\epsilon})$ for any $\epsilon > 0$.*

Proof. We use a reduction from the 2DIRPATH problem. Our proof is inspired by a reduction used in [9] for the edge-disjoint path problem. Assume that $I = (G, x_1, y_1, x_2, y_2)$ is an instance of 2DIRPATH and let $\epsilon > 0$ be given. We construct a directed graph H. First we construct a graph G' whose underlying structure is shown in Figure 1. For $N = |V(G)|^{1/\epsilon}$, create two sets of vertices $A = \{a_1, \ldots, a_N\}$ and $B = \{b_1, \ldots, b_N\}$. In the figure, all the edges are directed from top to bottom and from left to right. For each grey box, there is a vertex at each of the four corners, and there are two edges, from left to right and from top to bottom. This graph may be viewed as the union of N vertex-disjoint directed trees $T_1, \ldots, T_N$, where T_i is rooted at a_i and has paths to all the vertices in $B - \{b_i\}$. Each tree T_i consists of one horizontal path H^i, which is essentially the

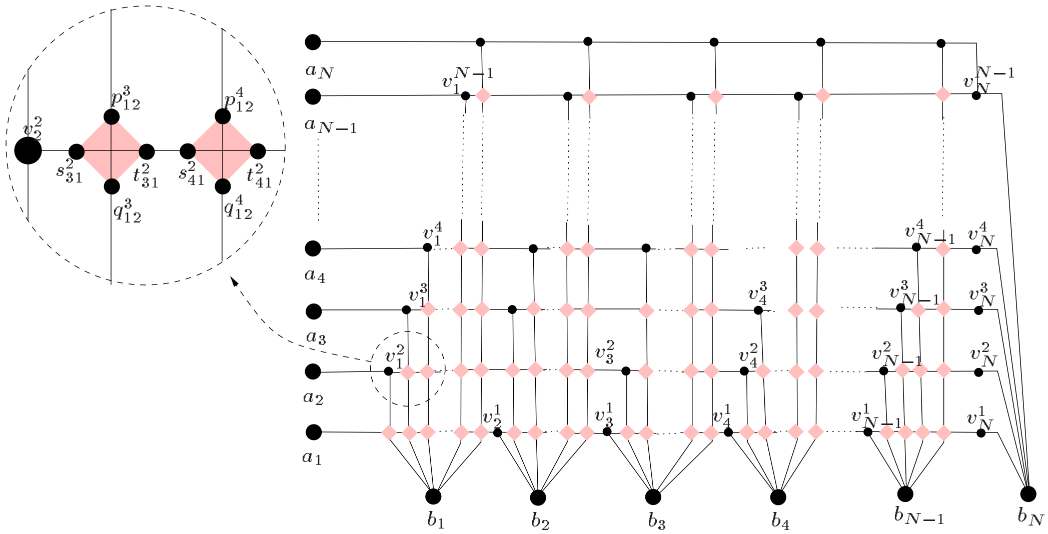

Fig. 1. Construction of H: each grey box will be replaced with a copy of G

ith horizontal row above b_j's, and starts with $a_i, v_1^i, \ldots$ and ends in v_N^i, together with $N-1$ vertical paths P_j^i ($1 \leq j \neq i \leq N$), such that each of these vertical paths branches out from the horizontal path, starting at vertex v_j^i and ending at vertex $b_j \in B$. Each vertex v_j^i is in the ith horizontal row, and is the start vertex of a vertical path that ends at b_j; note that there are no vertices v_i^i. Also, note that each grey box corresponds to a triple (i, j, ℓ) where the box is in the ith horizontal line and is in the vertical path that starts at v_j^ℓ and ends at b_j; the corner vertices of the grey box are labeled $s_{\ell j}^i, t_{\ell j}^i, p_{ji}^\ell, q_{ji}^\ell$ for the left, right, top, and bottom corners, respectively. More specifically, for T_1 the horizontal and vertical paths are: $H^1 = a_1, s_{21}^1, t_{21}^1, \ldots, s_{N1}^1, t_{N1}^1, v_2^1, s_{32}^1, t_{32}^1, \ldots, v_3^1, \ldots, v_N^1$, and $P_j^1 = v_j^1, b_j$, for $2 \leq j \leq N$. For T_2 the horizontal and vertical paths are: $H^2 = a_2, v_1^2, s_{31}^2, t_{31}^2, \ldots, s_{N2}^2, t_{N2}^2, v_3^2, \ldots, v_N^2$, and $P_j^2 = v_j^2, p_{j1}^2, q_{j1}^2, b_j$ (for $1 \leq j \neq 2 < N$), and $P_N^2 = v_N^2, b_N$. In general, for T_i:

- $H^i = a_i, v_1^i, s_{(i+1)1}^i, t_{(i+1)1}^i, \ldots, s_{N1}^i, t_{N1}^i, v_2^i, s_{(i+1)2}^i, t_{(i+1)2}^i, \ldots, s_{N2}^i, t_{N2}^i, v_3^i, \ldots,$ $v_{N-1}^i, s_{(i+1)(N-1)}^i, t_{(i+1)(N-1)}^i, \ldots, s_{N(N-1)}^i, t_{N(N-1)}^i, v_N^i$,
- For $j \neq i < N$: $P_j^i = v_j^i, p_{j(i-1)}^i, q_{j(i-1)}^i, p_{j(i-2)}^i, q_{j(i-2)}^i, \ldots, p_{j1}^i, q_{j1}^i, b_j$, and $P_N^i = v_N^i, b_N$.

Graph H is obtained from G' by making the following modifications:

- For each grey box, with corresponding triple say (i, j, ℓ) and with vertices $s_{\ell j}^i, t_{\ell j}^i, p_{ji}^\ell, q_{ji}^\ell$, we first remove the edges $s_{\ell j}^i t_{\ell j}^i$ and $p_{ji}^\ell q_{ji}^\ell$, then we place a copy of graph G and identify vertices x_1, y_1, x_2, and y_2 with $s_{\ell j}^i, t_{\ell j}^i, p_{ji}^\ell$, and q_{ji}^ℓ, respectively.

– Add a new vertex r (root) and create directed edges from r to $a_1, \ldots, a_N$.
– Create N new directed edges $a_i b_i$, for $1 \le i \le N$.

The set of terminals (for the directed Steiner trees) is $B \cup \{r\} = \{b_1, \ldots, b_N, r\}$. The proof follows from the following two claims (we omit their proofs).
Claim 1: If I is a "Yes" instance of 2DIRPATH, then H has N vertex-disjoint directed Steiner trees.
Claim 2: If I is a "No" instance, then H has exactly 1 vertex-disjoint directed Steiner tree.

The number of copies of G in the construction of H is $O(N^3)$ where $N = |V(G)|^{1/\epsilon}$. So the number of vertices in H is $O(N^{3+\epsilon})$. By Claims 1 and 2 it is NP-hard to decide if H has at least N or at most one directed Steiner trees. This creates a gap of $\Omega(n^{\frac{1}{3}-\epsilon})$. □

For PECD with unit capacities we use a reduction very similar to the one we presented for PVCD. The only differences are: (i) the instance that we use as the building block in our construction (corresponding to graph G above) is an instance of another well-known NP-hard problem, namely edge-disjoint 2DIRPATH (instead of vertex-disjoint), (ii) the parameter N above is $|E(G)|^{1/\epsilon}$. Using this reduction we can show:

Theorem 5. *Given an instance of PECD with unit capacities, it is NP-hard to approximate the solution within $O(m^{\frac{1}{3}-\epsilon})$ for any $\epsilon > 0$.*

2.2 Approximation algorithms

In this section we show that, although PECD is hard to approximate within a ratio of $O(m^{1/3-\epsilon})$, there is an approximation algorithm with a guarantee of $O(m^{\frac{1}{2}+\epsilon})$ (details in Theorem 7). The algorithm is LP-based with a simple rounding scheme similar to those in [2, 17]. The main idea of the algorithm is to start with one of the known approximation algorithms for finding a Minimum-weight Directed Steiner Tree. Using this and an extension of Theorem 4.1 in [14], we obtain an approximate solution to the fractional version of PECD. After that, a simple randomized rounding algorithm yields an integral solution. A similar method yields an approximation algorithm for PVCD that has a guarantee of $O(n^{\frac{1}{2}+\epsilon})$.

We may formulate PECD as an integer program (IP). In the following, $\mathcal{F}$ denotes the collection of all directed Steiner trees in G.

$$\begin{array}{ll} \text{maximize} & \sum_{F \in \mathcal{F}} x_F \\ \text{subject to} & \forall e \in E : \sum_{F : e \in F} x_F \le c_e \\ & \forall F \in \mathcal{F} : \quad x_F \in \{0, 1\} \end{array} \tag{1}$$

The *fractional packing edge capacitated directed Steiner tree* problem (fractional PECD, for short) is the linear program (LP) obtained by relaxing the integrality condition in the above IP to $x_F \ge 0$. For any instance I of the (integral) packing problem, we denote the fractional instance by I_f. The proof of Theorem 4.1 in [14] may be adapted to prove the following:

Theorem 6. *There is an α-approximation algorithm for fractional PECD if and only if there is an α-approximation algorithm for the minimum (edge weighted) directed Steiner tree problem.*

Charikar et al. [3] gave an $O(n^\epsilon)$-approximation algorithm for the minimum-weight directed Steiner tree problem. This, together with Theorem 6 implies:

Corollary 1. *There is an $O(n^\epsilon)$-approximation algorithm for fractional PECD.*

The key lemma in the design of our approximation algorithm for PECD is as follows.

Lemma 1. *Let I be an instance of PECD, and let φ^* be the (objective) value of a (not necessarily optimal) feasible solution $\{x^*_F : F \in \mathcal{F}\}$ to I_f such that the number of non-zero x^*_F's is polynomially bounded. Then, we can find in polynomial time, a solution to I with value at least $\Omega(\max\{\varphi^*/\sqrt{m},\ \min\{{\varphi^*}^2/m,\ \varphi^*\}\})$.*

Theorem 7. *Let I be an instance of PECD. Then for any $\epsilon > 0$, we can find in polynomial time a set of directed Steiner trees (satisfying the edge capacity constraints) of size at least $\Omega(\max\{\varphi_f/m^{\frac{1+\epsilon}{2}},\ \min\{\varphi_f^2/m^{1+\epsilon},\ \varphi_f/m^{\epsilon/2}\}\})$, where φ_f is the optimal value of the fractional instance I_f.*

Proof. Let I_f be the fractional instance. By Corollary 1, we can find an approximate solution φ^* for I_f such that $\varphi^* \geq c\varphi_f/m^{\epsilon/2}$ for some constant c and the given $\epsilon > 0$. Furthermore, the approximate solution contains only a polynomial number of Steiner trees with non-zero fractional values (this follows from the proof of Theorem 6 which is essentially the same as Theorem 4.1 in [14]). If we substitute φ^* in Lemma 1 we obtain an approximation algorithm that finds a set $\mathcal{F}'$ of directed Steiner trees such that $\mathcal{F}'$ has the required size. □

For PVCD we do the following. Given an instance I of PVCD with graph $G(V,E)$ and $T \subseteq V$ (with $|V| = n$ and $|E| = m$), we first use the reduction presented in Theorems 1 to produce an instance I' of PECD with graph $G'(V',E')$ and $T' \subseteq V'$. By the construction of G' we have $|V'| = 2|V| = 2n$ and there are only n edges in E' with bounded capacities (corresponding to the vertices of G). Therefore, if we use the algorithm of Theorem 7, the number of bad events will be n, rather than m. Using this observation we have the following:

Theorem 8. *Let I be an instance of PVCD. Then for any $\epsilon > 0$ we can find in polynomial time a set of directed Steiner trees (satisfying the vertex capacity constraints) of size at least $\Omega(\max\{\varphi_f/n^{\frac{1+\epsilon}{2}},\ \min\{\varphi_f^2/n^{1+\epsilon},\ \varphi_f/n^{\epsilon/2}\}\})$, where φ_f is the optimal value of the fractional instance I_f.*

3 Packing Undirected Steiner Trees

For packing edge-disjoint undirected Steiner trees (PEU), Jain et al. [14] showed that the (general) problem is APX-hard, and Kaski [16] showed the special case of the problem with only 7 terminal nodes is NP-hard. Here we show that PEU is APX-hard even when there are only 4 terminals. In an early draft of this paper we also showed, independently of [16], that finding two edge-disjoint Steiner trees is NP-hard. Both of these hardness results carry over (using similar constructions) to PVU. The following observation will be used in our proofs:

Observation 2. *For any solution to any of our Steiner tree packing problems, we may assume that: (1) In any Steiner tree, none of the leaves is a Steiner node (otherwise we simply remove it). (2) Every Steiner node with degree 3 belongs to at most one Steiner tree.*

Using this observation, the proof of NP-hardness for finding two edge-disjoint Steiner trees implies the following theorem.

Theorem 9. *Finding 2 vertex-disjoint undirected Steiner trees is NP-hard.*

Theorem 10. *PEU is APX-hard even if there are only 4 terminals.*

Proof. We use a reduction from Bounded 3-Dimensional Matching (B3DM). Assume that we are given three disjoint sets X, Y, Z (each corresponding to one part of a 3-partite graph G), with $|X| = |Y| = |Z| = n$, and a set $E \subseteq X \times Y \times Z$ containing m triples. Furthermore, we assume that each vertex in $X \cup Y \cup Z$ belongs to at most 5 triples. It is known [15] that there is an absolute constant $\epsilon_0 > 0$ such that it is NP-hard to distinguish between instances of B3DM where there is a perfect matching (i.e., n vertex-disjoint triples) and those in which every matching (set of vertex-disjoint triples) has size at most $(1 - \epsilon_0)n$. Assume that $x_1, \ldots, x_n$, $y_1, \ldots, y_n$, and $z_1, \ldots, z_n$ are the nodes of X, Y, and Z, respectively. We construct a graph H which consists of:

- 4 terminals t_x, t_y, t_z, and t_{yz}.
- Non-terminals $x_1, \ldots, x_n$, $y_1, \ldots, y_n$, and $z_1, \ldots, z_n$ (corresponding to the nodes in X, Y, Z), $x'_1, \ldots, x'_{m-n}$, $y'_1, \ldots, y'_{m-n}$, and $z'_1, \ldots, z'_n$.
- Two non-terminals U and W.
- Edges $t_x x_i$, $t_y y_i$, $t_z z'_i$, $z'_i z_i$, and $t_{yz} z'_i$, for $1 \le i \le n$.
- Edges $t_x x'_i$, $x'_i y'_i$, $x'_i U$, $y'_i t_y$, and $y'_i t_{yz}$, for $1 \le i \le m - n$.
- $m - n$ parallel edge from W to t_z.
- For each triple $e_q = x_i y_j z_k \in E$, 3 non-terminals v^q_c, v^q_x, v^q_z and the following edges: $v^q_x x_i$, $v^q_z z_k$, $v^q_c v^q_x$, $v^q_c y_j$, $v^q_c v^q_z$, $v^q_x U$, and $v^q_z W$.

See Figure 2. We can show that (a) [completeness] if G has a perfect matching then H has m edge-disjoint Steiner trees and (b) [soundness] if every matching in G has size at most $(1 - \epsilon_0)n$ then H has at most $(1 - \epsilon_1)m$ edge-disjoint Steiner trees, for $\epsilon_1 \ge \epsilon_0/110$. □

The constant 110 in the above theorem is not optimal. We can find explicit lower bounds for the hardness of PEU with constant number of terminals using the known hardness results for kDM, for higher values of k. For instance, Hazan et al. [13] proved that 4DM (with upper bounds on the degree of each vertex)

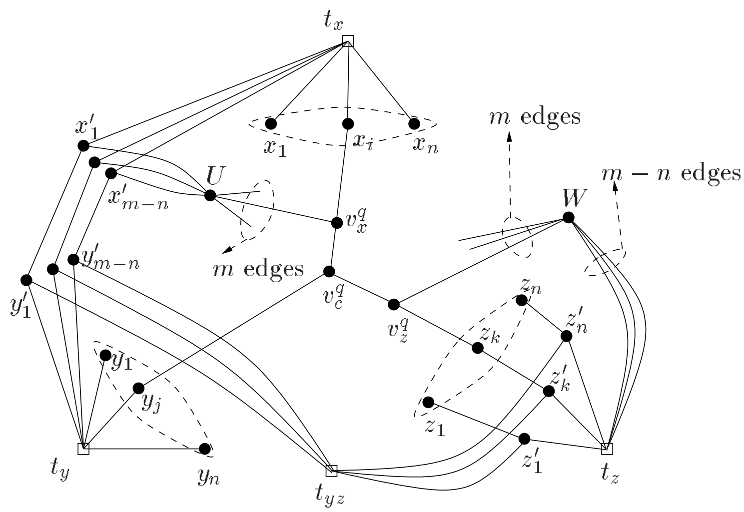

Fig. 2. Construction with 4 terminals from B3DM

is hard to approximate within a factor $\frac{53}{54}$. Using this and a reduction similar to the one presented in Theorem 10 it seems possible to show that PEU with 5 terminals is hard to approximate within a factor $(1 + \frac{1}{2000})$. The proof of Theorem 10 extends to give the next result.

Theorem 11. *PVU is APX-hard even with only 6 terminals.*

These results may seem to suggest that PEU and PVU have the same approximation guarantees. But the next theorem shows that PVU is significantly harder than PEU. We show this by a reduction from the set-cover packing problem (or domatic number problem). Given a bipartite graph $G(V_1 \cup V_2, E)$, a *set-cover* (of V_2) is a subset $S \subseteq V_1$ such that every vertex of V_2 has a neighbor in S. A *set-cover packing* is a collection of pairwise disjoint set-covers of V_2. The goal is to find a packing of set-covers of maximum size. Feige et al. [6] show that, unless NP $\subseteq$ DTIME($n^{\log \log n}$), there is no $(1-\epsilon)\ln n$ approximation algorithm for set-cover packing, where $n = |V_1| + |V_2|$. We have the following theorem.

Theorem 12. *PVU cannot be approximated within ratio* $(1-\epsilon)\log n$*, for any* $\epsilon > 0$*, unless NP* $\subseteq$ *DTIME*($n^{\log \log n}$)*.*

On the other hand, we can obtain an $O(\sqrt{n}\log n)$ algorithm for PVCU (which contains PVU as a special case). To do so, consider the fractional version of PVCU obtained by relaxing the integrality condition in the IP formulation. The separation problem for the dual of this LP is the minimum node-weighted Steiner tree problem. For this problem, Guha and Khuller [12] give an $O(\log n)$ approximation algorithm. Using the following analog of Theorem 6 (or Theorem 4.1 in [14]) we obtain a polytime $O(\log n)$ approximation for fractional PVCU.

Lemma 3. *There is an α-approximation for fractional PVCU if and only if there is an α-approximation for the minimum node weighted Steiner tree problem.*

Remark: Lemma 3 and the fact that the minimum node weighted Steiner tree problem is hard to approximate within $O(\log k)$ (with k being the number of terminals) yields an alternative proof for the $\Omega(\log k)$ hardness of PVCU.

The algorithm for PVCU is similar to the ones we presented for PECD and PVCD. That is, we apply randomized rounding to the solution of the fractional PVCU instance. Skipping the details, this yields the following:

Theorem 13. *Given an instance of PVCU and any $\epsilon > 0$, we can find in polynomial time a set of Steiner trees (satisfying the vertex capacity constraints) of size at least $\Omega(\max\{\varphi_f/\sqrt{n}\log n,\ \min\{\varphi_f^2/n\log^2 n,\ \varphi_f/\log n\}\})$, where φ_f is the optimal value of the instance of fractional PVCU.*

3.1 Packing vertex-disjoint priority Steiner trees

The priority Steiner problem has been studied by Charikar et al. [4]. Here, we study the problem of packing vertex-disjoint priority Steiner trees of undirected graphs. (One difference with the earlier work in [4] is that weights and priorities are associated with vertices rather than with edges). Consider an undirected graph $G = (V, E)$ with a set of terminals $T \subseteq V$, one of which is distinguished as the root r, every vertex v has a nonnegative integer p_v as its priority, and every Steiner vertex $v \in V - T$ has a positive capacity c_v. A *priority Steiner* tree is a Steiner tree such that for each terminal $t \in T$ every Steiner vertex v on the r, t path has priority $p_v \geq p_t$. In the problem PVCU-priority (Packing Vertex Capacitated Undirected Priority Steiner Trees) the goal is to find a maximum set of priority Steiner trees obeying vertex capacities (i.e., for each Steiner vertex $v \in V - T$ the number of trees containing v is $\leq c_v$). The algorithm we presented for PVCU extends to PVCU-priority, giving roughly the same approximation guarantee.

Theorem 14. *Given an instance of PVCU-priority and any $\epsilon > 0$, we can find in polynomial time a set of priority Steiner trees (satisfying the vertex capacity constraints) of size at least $\Omega(\max\{\varphi_f/n^{\frac{1+\epsilon}{2}},\ \min\{\varphi_f^2/n^{1+\epsilon},\ \varphi_f/n^{\epsilon/2}\}\})$, where φ_f is the optimal value of the instance of fractional PVCU-priority.*

On the other hand, we prove an $\Omega(n^{\frac{1}{3}-\epsilon})$ hardness result for PVCU-priority by adapting the proof of Theorem 4 (thus improving on our logarithmic hardness result for PVCU). The main difference from the proof of Theorem 4 is that we use instances of the Undir-Node-USF problem (*Undirected Node capacitated Unsplittable Flow*) – which is shown to be NP-complete in [9] – instead of instances of 2DIRPATH as the modules that are placed on the "grey boxes" in Figure 1.

Theorem 15. *Given an instance of PVCU-priority, it is NP-hard to approximate the solution within $O(n^{\frac{1}{3}-\epsilon})$ for any $\epsilon > 0$.*

References

1. S.Arora and C.Lund, *Hardness of approximations*, in *Approximation Algorithms for NP-hard Problems*, Dorit Hochbaum Ed., PWS Publishing, 1996.
2. A.Baveja and A.Srinivasan, *Approximation algorithms for disjoint paths and related routing and packing problems*, Mathematics of Operations Research 25:255-280, 2000. Earlier version in FOCS 1997.
3. M.Charikar, C.Chekuri, T.Cheung, Z.Dai, A.Goel, S.Guha, and M.Li, *Approximation algorithms for directed Steiner problem*, J. Algorithms 33(1):73-91, 1999. Earlier version in SODA 1998.
4. M.Charikar, J.Naor, and B.Schieber, *Resource optimization in QoS multicast routing of real-time multimedia*, Proc. 19th Annual IEEE INFOCOM (2000).
5. R.Diestel, *Graph Theory*, Springer, New York, NY, 2000.
6. U.Feige, M.Halldorsson, G.Kortsarz, and A.Srinivasan, *Approximating the domatic number*, Siam J. Computing 32(1):172-195, 2002. Earlier version in STOC 2000.
7. P.Floréen, P.Kaski, J.Kohonen, and P.Orponen, *Multicast time maximization in energy constrained wireless networks*, in Proc. 2003 Joint Workshop on Foundations of Mobile Computing (DIALM-POMC 2003), San Diego, CA, 2003.
8. A.Frank, T.Király, M.Kriesell, *On decomposing a hypergraph into k connected sub-hypergraphs*, Discrete Applied Mathematics 131(2):373-383, 2003.
9. 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*, J. Computer and System Sciences 67(3):473-496, 2003. Earlier version in STOC 1999.
10. S.Fortune, J.Hopcroft, and J.Wyllie, *The directed subgraph homeomorphism problem*, Theoretical Computer Science 10(2):111-121, 1980.
11. M.Grötschel, A.Martin, and R.Weismantel, *The Steiner tree packing problem in VLSI design*, Mathematical Programming 78:265-281, 1997.
12. S.Guha and S.Khuller, *Improved methods for approximating node weighted Steiner trees and connected dominating sets*, Information and Computation 150:57-74, 1999. Earlier version in FST&TCS 1998.
13. E.Hazan, S.Safra, and O.Schwartz, *On the hardness of approximating k-dimensional matching*, Electronic Colloqium on Computational Complexity, Rep. No. 20, 2003.
14. K.Jain, M.Mahdian, M.R.Salavatipour, *Packing Steiner trees*, Proc. SODA 2003.
15. V.Kann, *Maximum bounded 3-dimensional matching is MAX SNP-complete*, Information Processing Letters 37:27-35, 1991.
16. P.Kaski, *Packing Steiner trees with identical terminal sets*, Information Processing Letters 91(1):1-5, 2004.
17. S.G.Kolliopoulos and C.Stein, *Approximating disjoint-path problems using packing integer programs*, Mathematical Programming 99:63-87, 2004. Earlier version in Proc. IPCO 1998.
18. J.Kratochvil and Z.Tuza, *On the complexity of bicoloring clique hypergraphs of graphs*, J. Algorithms 45:40-54, 2002. Earlier version in SODA 2000.
19. M.Kriesell, *Edge-disjoint trees containing some given vertices in a graph*, J. Combinatorial Theory (B) 88:53-65, 2003.
20. L.C.Lau, *An approximate max-Steiner-tree-packing min-Steiner-cut theorem*, manuscript, 2004.
21. A.Martin and R.Weismantel, *Packing paths and Steiner trees: Routing of electronic circuits*, CWI Quarterly 6:185-204, 1993.