The General Steiner Tree-Star Problem

Samir Khuller* An Zhu†

Abstract

The Steiner tree problem is defined as follows - given a graph $G = (V, E)$ and a subset $X \subset V$ of terminals, compute a minimum cost tree that includes all nodes in X. Furthermore, it is reasonable to assume that the edge costs form a metric. This problem is NP-hard and has been the study of many heuristics and algorithms. We study a generalization of this problem, where there is a "switch" cost in addition to the cost of the edges. Switches are placed at internal nodes of the tree (essentially, we may assume that all non-leaf nodes of the Steiner tree have a switch). The cost for placing a switch may vary from node to node. A restricted version of this problem, where the terminal set X cannot be connected to each other directly but only via the Steiner nodes $V \setminus X$, is referred to as the Steiner Tree-Star problem. The General Steiner Tree-Star problem does not require the terminal set and Steiner node set to be disjoint. This generalized problem can be reduced to the node weighted Steiner tree problem, for which algorithms with performance guarantees of $\Theta(\ln n)$ are known. However, such approach does not make use of the fact that the edge costs form a metric. In this paper we derive approximation algorithms with small constant factors for this problem. We show two different polynomial time algorithms with approximation factors of 5.16 and 5.

Keywords: approximation algorithms, facility location, Steiner trees.

*Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742. This research was supported by an NSF CAREER Award CCR-9501355 and NSF Award CCR-9820965. E-mail : samir@cs.umd.edu.

†Department of Computer Science, Stanford University, Stanford, CA 94305. This work was done while this author was at the University of Maryland. Currently supported by a GRPW fellowship from Bell Labs, Lucent Technologies. E-mail : anzhu@cs.stanford.edu.

1 Introduction

In this paper we study the General Steiner Tree-Star problem [16, 12, 13]. A description of this problem is as follows: we are given a complete graph $G = (V, E)$ with the edge weights W_{ij} forming a metric, and a non-negative cost function W_i defined on the vertices. We are also given a subset X of vertices, called terminal nodes, and another subset Y of vertices, called Steiner nodes. The goal is to compute a minimum cost Steiner tree T that includes all terminal nodes, so that only Steiner nodes can be the non-leaf nodes in the tree. The cost of the tree is the total cost of the edges in the tree and the cost of the non-leaf vertices in the tree.

The restricted version of this problem - the Steiner Tree-Star problem arises in private line data network design. In this setting the sets X and Y are disjoint. This problem is NP-complete, and various heuristics have been developed. In previous works, this problem is reduced to the node weighted Steiner tree problem [12] or the directed Steiner tree problem [14]. Several heuristics and algorithms have been developed for the node weighted Steiner tree problem [16, 12, 13, 11, 7], for computing approximate solutions, as well as for computing optimal solutions.

In this paper we focus on approximation algorithms. The node weighted Steiner tree problem generalizes the edge weighted Steiner tree problem to the case when the nodes and edges have weights [5]. The approximation algorithms developed by Klein and Ravi, and Guha and Khuller yield approximation factors of $2 \log n$ and $1.35 \log n$ respectively [11, 7]. The directed Steiner tree problem is a significantly harder problem to develop approximation algorithms for, and the approximation factors are significantly worse [3]. Both the node weighted and directed Steiner tree problem have a lower bound of $\Omega(\log n)$ in approximation ratios due to a reduction from the Set Cover problem. However, the reduction from the General Steiner Tree-Star problem to either problem above does not use the fact that the edge weights form a metric. Thus one may hope for a better approximation ratio using a different approach.

Indeed our problem is a special case of the connected facility location problem. This problem is first studied by Karger and Minkoff [10]. It generalizes our problem by adding a constant scaling factor M to the costs of all edges in the solution that connect two non-leaf nodes[1]. Later Gupta et. al. improved the approximation ratio to 12 [8].

We show that the General Steiner Tree-Star problem (where the sets X and Y need not be disjoint, and edge costs form a metric) can be approximated to within a small constant factor of the optimal solution[2]. We use techniques that have been developed for the Uncapacitated Facility Location (UFL) problem to develop polynomial time approximation algorithms with a worst case approximation guarantee of 5.

2 Algorithms

The basic idea behind the algorithms is very simple. For completeness we give a brief description of the Uncapacitated Facility Location Problem (UFL) below. We are given a set of locations $N = \{1, 2, \ldots, n\}$, and distances between them c_{ij}, $i, j = 1, 2, \ldots, n$. There is a subset $F \subseteq N$ of locations at which we may open a facility, and there is a subset $C \subseteq N$ of locations that must be assigned to an open facility. For each location $i \in F$, the non-negative cost of opening a facility at i is f_i. The cost of assigning each $j \in C$ to $i \in F$ is c_{ij}. The distance metric satisfies the triangle inequality and is symmetric. The objective function is to find a feasible assignment of each location

[1] Our problem is simply the version where $M = 1$.

[2] The metric requirement is essential, otherwise the problem becomes hard to approximate due to a reduction from the Set Cover problem.

in C to an open facility so as to minimize the total cost incurred. This problem is NP-hard, and several constant factor approximation algorithms have been developed recently [15, 6, 4, 9, 2]. We map the General Steiner Tree-Star problem $G = (V, E)$ to UFL as follows: V corresponds to N, X corresponds to C, Y corresponds to F, W_{ij} corresponds to c_{ij}, and W_i corresponds to f_i. The extra requirement for our problem is that we require the facilities to be connected as well. This gives rise to our main approach. Namely, we first solve the UFL on the given input and then we simply connect the open facilities together with a Minimum Spanning Tree. However, we need to prove that the ratio of the cost of this solution to the optimal solution can be bounded by a (small) constant. We will exploit a common property of the algorithms for UFL based on LP-rounding [15] and the Primal-Dual method [9].

We use two different methods to solve the UFL problem. The first approach is the method proposed by Shmoys, Tardos and Aardal [15] that has an approximation factor of 3.16 for UFL. In this method we first solve the linear program that is a relaxation to the natural integer program for the UFL problem. The second method uses a primal-dual algorithm to solve the UFL problem without requiring an optimal solution to the corresponding linear program [9].

We first describe an Integer Program for the UFL problem as follows [1]:

Minimize

$$\sum_{i \in F, j \in C} c_{ij} x_{ij} + \sum_{i \in F} f_i y_i$$

Subject to:

$$\forall j \in C : \sum_{i \in F} x_{ij} \geq 1$$
$$\forall i \in F, j \in C : y_i - x_{ij} \geq 0$$
$$\forall i \in F, j \in C : x_{ij} \in \{0, 1\}$$
$$\forall i \in F : y_i \in \{0, 1\}$$

The LP-relaxation of this problem is obtained by replacing the last two constraints with

$$\forall i \in F, j \in C : x_{ij} \geq 0$$
$$\forall i \in F : y_i \geq 0.$$

The dual program for this relaxation is:

Maximize

$$\sum_{j \in C} \alpha_j$$

Subject to:

$$\forall i \in F, j \in C : \alpha_j - \beta_{ij} \leq c_{ij}$$
$$\forall i \in F : \sum_{j \in C} \beta_{ij} \leq f_i$$
$$\forall i \in F, j \in C : \beta_{ij} \geq 0$$
$$\forall i \in F : \alpha_i \geq 0$$

2.1 Algorithm 1

The approach in [15] is to solve the LP, and then to round the obtained solution. Let $0 \leq \alpha \leq 1$ be a constant. Define $N_j(\alpha)$ to be the neighborhood of customer j with total supply at least α [15]. (In other words, fix a demand location $j \in C$ and order the facilities of F in non-decreasing distance order. $N_j(\alpha)$ is the closest set of facilities such that $\sum_{i \in N_j(\alpha)} x_{ij} \geq \alpha$.) Let $\tau(\alpha, j)$ denote the radius of $N_j(\alpha)$. The algorithm, followed by a simple MST step, is outlined as follows:

General Steiner Tree-Star(G,α) —

1 Solve the LP for the UFL problem
2 All customers are unmarked
3 **while** unmarked customers exist
4 Pick unmarked customer j with the smallest $\tau(\alpha, j)$
5 Open a facility at the **cheapest** center $F(j)$ in $N_j(\alpha)$.
6 Let $F(j)$ be the direct supplier for customer j and j is the witness for $F(j)$.
7 Add j to the witness set $\mathcal{D}$
8 Mark all unmarked customers j' if $N_j(\alpha) \cap N_{j'}(\alpha) \neq \phi$ and assign them to $F(j)$
9 Connect all open facilities with an MST

2.2 Algorithm 2

The second approach to UFL, proposed by Jain and Vazirani [9], is based on the primal-dual method. It does not directly solve the LP, even though it uses the framework described above. The approximation factor of this method (for UFL) is 3. We replace lines 1 through 8 in the above algorithm, with the primal dual method. The last step remains the same. The primal-dual method also uses the notion of witness customers, they are referred to as directly connected cities in [9]. Furthermore, each dual variable α_j is decomposed into 2 parts, $\alpha_j = \alpha_j^f + \alpha_j^e$. For witness customers, α_j^f is the partial facility cost that the customer contributes, while α_j^e is exactly the cost to connect the customer to its open facility.

3 Analysis

The proofs for both approaches are similar. We will begin with the 5.16 bound from the LP based approach. We decompose the cost of an optimal integral solution for the General Steiner Tree-Star as follows:

$$OPT = OPT_f + OPT_s + OPT_c$$

Where OPT_f denotes the total cost of open facilities, OPT_s denotes the cost to serve all customers, and OPT_c denotes the cost to connect all open facilities together. OPT_s and OPT_c are also referred to as service cost and connection cost respectively. Figure 1 shows the optimal solution and corresponding assignments. Similarly, let the cost of our solution be of the following:

$$C = C_f + C_s + C_c$$

Let OPT_{LP} denote the optimal LP solution of the UFL problem. Obviously $OPT_{LP} \leq OPT_f + OPT_s \leq OPT$.

Lemma 3.1 *The MST connection cost for our algorithm can be bounded as follows:* $C_c \leq 2(OPT_s + OPT_c) + 2\sum_{j \in \mathcal{D}} d_j$, *where* d_j *is the distance of the witness customer* j *to its open facility. In the first approach* $d_j = \tau(\alpha, j)$. *In the primal-dual approach* $d_j = \alpha_j^e$.

Proof. Figure 1 illustrates the fact that there exists a Steiner tree connecting all the open facilities. The Steiner tree first connects all open facilities to their witness customers, then the witness customers are connected via the Steiner tree from OPT. The total cost of this Steiner tree is upper bounded by $OPT_s + OPT_c + \sum_{j \in \mathcal{D}} d_j$. We use MST to connect all open facilities, whose cost is no more than twice that of any Steiner tree. □

Lemma 3.2 *Using the first approach,* $C_f \leq \frac{1}{\alpha} \sum_{i \in F} f_i y_i$ *and* $C_s \leq \sum_{j \in \mathcal{D}} \tau(\alpha, j) + 3 \sum_{j' \notin \mathcal{D}} \tau(\alpha, j)$

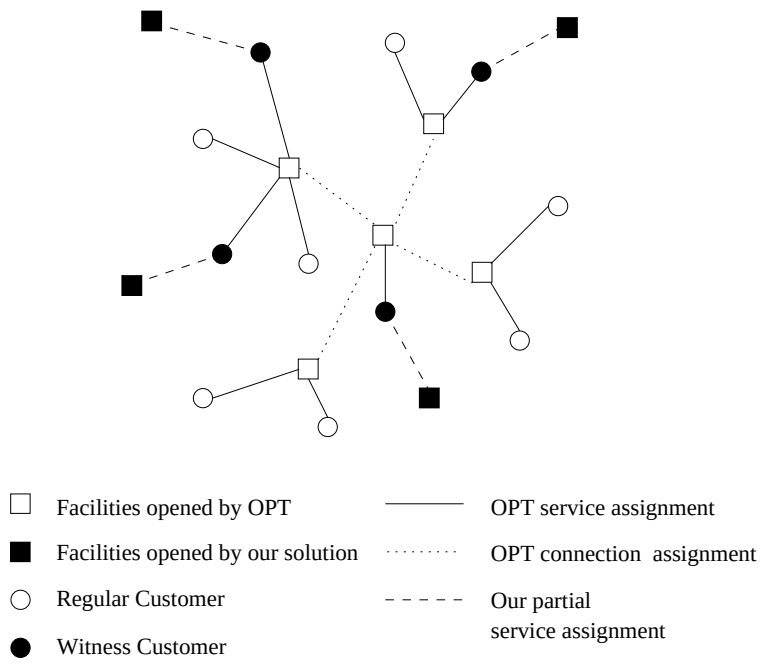

Figure 1: Connection Cost

Proof. The proof is a slight modification of the proof in [15]. The facility cost of our solution is no more than $\frac{1}{\alpha}\sum_{i\in F} f_i y_i$. The service cost for witness customer j is no more than $\tau(\alpha, j)$. The service cost for any other customer j' is no more than $\tau(\alpha, j') + 2\tau(\alpha, j) \leq 3\tau(\alpha, j')$, where j is the witness customer for the open facility that j' is assigned to. This is because at the time customer j was chosen, $\tau(\alpha, j)$ is the smallest among the rest of the customers. In particular this implies $\tau(\alpha, j) \leq \tau(\alpha, j')$. □

Combining the above lemmas, we conclude that the total cost is no more than $\frac{1}{\alpha}\sum_{i\in F} f_i y_i + 3\sum_{j\in C}\tau(\alpha, j) + 2(OPT_s + OPT_c)$. If we ignore the term $2(OPT_s + OPT_c)$, the rest of the cost is bounded by the following theorem in [15].

Theorem 3.3 *For the metric uncapacitated facility location problem, randomized filtering and rounding yields an algorithm that finds a solution whose expected total cost is within a factor of* $3/(1 - e^{-3}) < 3.16$ *of the optimum.*

This expected bound is achieved by randomly pick α from range $[1/e^3, 1]$. Using this theorem, we have the following fact.

Theorem 3.4 *For the General Steiner Tree-Star problem, an LP based method yields an algorithm that finds a solution whose expected total cost is within a factor of* 5.16 *of the optimum.*

Proof. Our cost is no more than $3.16OPT_{LP} + 2(OPT_s + OPT_c) \leq 3.16(OPT_f + OPT_s) + 2(OPT_s + OPT_c) \leq 5.16OPT$. □

The 5-approximation proof for the primal-dual method is very similar. We again exploit the fact that the extra term $2\sum_{j\in\mathcal{D}} d_j$ in C_c can be absorbed in the original 3-approximation proof for the UFL problem in [9].

Corollary 3.5 *Using the primal-dual method,* $C_f + C_s + 2\sum_{j\in\mathcal{D}} d_j \leq 3\sum_{j\in C}\alpha_j \leq 3OPT_{LP}$.

Proof. From Theorem 7 in [9], the service cost of the witness customers is upper bounded by $3\alpha_j^e$, instead of the tighter bound α_j^e. Thus we can add the extra term at the left hand side and still get the same 3-approximation. □

Theorem 3.6 *The primal-dual approach will give an approximation factor of* 5 *for the General Steiner Tree-Star problem.*

Proof. Our total cost is then $C \leq 3OPT_{LP} + 2(OPT_s + OPT_c) \leq 3(OPT_f + OPT_s) + 2(OPT_s + OPT_c) \leq 5OPT$. □

Notice that in both algorithms, the extra $2\sum_{j\in\mathcal{D}} d_j$ term in C_c can essentially be included for free in the argument of approximation results for the UFL problem. Although there are even better results for the UFL problem in terms of approximation ratios, employing them in a similar fashion directly does not give better approximation ratios because of this extra $2\sum_{j\in\mathcal{D}} d_j$ term.

4 Conclusion

In this paper we presented polynomial-time algorithms with approximation ratios 5.16 and 5 for the General Steiner Tree-Star problem, exploiting the connection to the Uncapacitated Facility Location problem. An interesting open problem is the possibility of utilizing existing better approximation algorithms designed for the UFL problem to obtain better approximation ratios. Another open problem is to improve the 12 approximation ratio for the connected facility location problem.

References

[1] M. L. Balinksi. On finding integer solutions to linear programs. In *Proceedings of the IBM Scientific Computing Symposium on Combinatorial Problems*, pages 225-248. IBM, 1966.

[2] M. Charikar and S. Guha. Improved combinatorial algorithms for the facility location and k-Median problem. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 378–388, (1999).

[3] M. Charikar, C. Chekuri, T. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed Steiner problems. *Journal of Algorithms*, Vol 33:73–91 (1999).

[4] F. Chudak. Improved approximation algorithms for uncapacitated facility location problem. In *Proceedings of Integer Programming and Combinatorial Optimization*, pages 180–194, 1998.

[5] F. Hwang, D. Richards and P. Winter. *The Steiner tree problem*, Elsevier Science Publishers B.V. (1992).

[6] S. Guha and S. Khuller. Greedy strikes back: Improved facility location algorithms. *Journal of Algorithms*, Vol 31(1):228–248 (1999).

[7] S. Guha and S. Khuller. Improved methods for approximating node weighted steiner trees and connected dominating sets. *Information and Computation*, Vol 150:57–74, (1999).

[8] A. Gupta, J. Kleinberg, A. Kumar, R. Rastogi, and B. Yener. Provisioning a virtual private network: A network design problem for multicommodity flow. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing*, pages 389-398, (2001).

[9] K. Jain and V. Vazirani. Primal-Dual approximation algorithms for metric facility location and k-median problems. In *Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science*, pages 2-13 (1999).

[10] D. R. Karger and M. Minkoff. Building Steiner trees with incomplete global knowledge. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science*, pages 613-623, (2000).

[11] P. N. Klein and R. Ravi. A nearly best-possible approximation algorithm for node-weighted Steiner trees. *Journal of Algorithms*, Vol 19(1):104–114, (1995).

[12] Y. Lee, S. Chiu, and J. Ryan. A branch and cut algorithm for the Steiner tree/star problem. *INFORMS Journal on Computing*, Vol 8:194-201, (1996).

[13] Y. Lee, L. Lu, Y. Qiu and F. Glover. Strong formulations and cutting planes for designing digital data networks. *Telecommunication Systems*, Vol 2:261-274, (1994).

[14] S. Raghavan, personal communication (1999).

[15] D. Shmoys, E. Tardos and K. Aardal. Approximation algorithms for facility location problems. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing*, pages 265-274, (1997).

[16] J. Xu, S. Chiu and F. Glover. Using Tabu search to solve the Steiner tree-star problem in telecommunication network design. *Telecommunications systems*, Vol 6:117-125 (1996).