



Prize-Collecting Steiner Tree: A 1.79 Approximation

Ali Ahmadi

University of Maryland
College Park, Maryland, USA
ahmadia@umd.edu

Iman Gholami

University of Maryland
College Park, Maryland, USA
igholami@umd.edu

MohammadTaghi Hajiaghayi

University of Maryland
College Park, Maryland, USA
hajiagha@umd.edu

Peyman Jabbarzade

University of Maryland
College Park, Maryland, USA
peymanj@umd.edu

Mohammad Mahdavi

University of Maryland
College Park, Maryland, USA
mahdavi@umd.edu

ABSTRACT

Prize-Collecting Steiner Tree (PCST) is a generalization of the Steiner Tree problem, a fundamental problem in computer science. In the classic Steiner Tree problem, we aim to connect a set of vertices known as terminals using the minimum-weight tree in a given weighted graph. In this generalized version, each vertex has a penalty, and there is flexibility to decide whether to connect each vertex or pay its associated penalty, making the problem more realistic and practical.

Both the Steiner Tree problem and its Prize-Collecting version had long-standing 2-approximation algorithms, matching the integrality gap of the natural LP formulations for both. This barrier for both problems has been surpassed, with algorithms achieving approximation factors below 2. While research on the Steiner Tree problem has led to a series of reductions in the approximation ratio below 2, culminating in a $\ln(4) + \epsilon$ approximation by Byrka, Grandoni, Rothvoß, and Sanità [STOC'10], the Prize-Collecting version has not seen improvements in the past 15 years since the work of Archer, Bateni, Hajiaghayi, and Karloff [FOCS'09, SIAM J. Comput.'11], which reduced the approximation factor for this problem from 2 to 1.9672. Interestingly, even the Prize-Collecting TSP approximation, which was first improved below 2 in the same paper, has seen several advancements since then (see, e.g., Blauth and Nägele [STOC'23]).

In this paper, we reduce the approximation factor for the PCST problem substantially to 1.7994 via a novel iterative approach.

CCS CONCEPTS

• **Mathematics of computing** → **Graph algorithms; Approximation algorithms.**

KEYWORDS

Steiner tree, prize-collecting, approximation algorithm

ACM Reference Format:

Ali Ahmadi, Iman Gholami, MohammadTaghi Hajiaghayi, Peyman Jabbarzade, and Mohammad Mahdavi. 2024. Prize-Collecting Steiner Tree: A



This work is licensed under a Creative Commons Attribution 4.0 International License.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

© 2024 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0383-6/24/06

<https://doi.org/10.1145/3618260.3649789>

1.79 Approximation. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24)*, June 24–28, 2024, Vancouver, BC, Canada. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3618260.3649789>

1 INTRODUCTION

The Steiner Tree problem is a well-known problem in the field of combinatorial optimization. It involves connecting a specific set of vertices (referred to as terminals) in a weighted graph while aiming to minimize the total cost of the edges used. The problem also allows for the inclusion of additional vertices, known as Steiner points, which can help reduce the overall cost. This problem has a long history and was formally defined mathematically by Hakimi in 1971 [19]. It is recognized as one of the classic NP-hard problems [20]. The Steiner Tree problem finds applications in various domains, including network design [4] and phylogenetics [24], prompting continuous research efforts to develop more efficient approximation algorithms.

Initial algorithmic strategies for the Steiner Tree problem, while heuristic in nature, set the stage for more precise approaches. Zelikovsky's 1993 introduction of a polynomial-time approximation algorithm achieved an 11/6-approximation ratio [25], which was followed by further improvements including Karpinski and Zelikovsky's 1.65-approximation in 1995 [21]. The approach was refined to a 1.55-approximation by Robins and Zelikovsky in 2005 [23], and by 2010, Byrka, Grandoni, Rothvoß, and Sanità advanced this to a 1.39-approximation [12]. An earlier MST-based 2-approximation algorithm, introduced in the early 1980s, also played a crucial role due to its simplicity [22].

The computational complexity of the Steiner Tree problem has been firmly established. Bern and Plassmann showed its MAX SNP-hardness, indicating the absence of a polynomial-time approximation scheme (PTAS) for this problem unless P equals NP [8]. Building on this, Chlebík and Chlebíková in 2008 established a lower bound, demonstrating that approximating the Steiner Tree problem within a factor of 96/95 of the optimal solution is NP-hard. This finding marks a crucial step in understanding the inherent complexity of the problem [13].

In combinatorial optimization, prize-collecting variants are distinct for their detailed decision-making approach. These variants focus not only on building an optimal structure but also on intentionally excluding certain components, which leads to a penalty. This introduces more complexity and makes these problems more applicable to real-world scenarios. The concept of prize-collecting problems in optimization was first brought forward by Balas in

the late 1980s [7]. This pioneering work opened a new research direction, particularly in scenarios where avoiding certain elements results in penalties. Following this, the first approximation algorithms for prize-collecting problems were introduced in the early 1990s by authors including Bienstock, Goemans, Simchi-Levi, and Williamson [9]. Their initial contributions have significantly shaped the research direction in this area, focusing on developing solutions that effectively balance costs against penalties.

The Prize-collecting Steiner Tree (PCST) problem is a key example in this category, as it takes into account both the costs of connectivity and penalties for excluding vertices. In this problem, we consider an undirected graph $G = (V, E)$ where V represents vertices and E represents edges. Each edge $e \in E$ has an associated cost $c(e)$, and each vertex $v \in V$ comes with a penalty $\pi(v)$ that needs to be paid if the vertex is not connected in the solution. The objective is to find a tree $T = (V_T, E_T)$ within G that minimizes the sum of edge costs in T and penalties for vertices not in T . This is mathematically expressed as:

$$\text{Minimize } \sum_{e \in E_T} c(e) + \sum_{v \in V \setminus V_T} \pi(v).$$

This formulation captures the essence of the PCST problem: a trade-off between the infrastructure cost, represented by the sum of the edge costs within the chosen tree, and the penalties assigned to vertices excluded from this connecting structure. This detailed view of the problem applies to various situations, such as network design where not every node needs to be connected, and resource allocation where some demands might not be met, resulting in a cost.

Initial strides in developing approximation algorithms for PCST were made by Bienstock, Goemans, Simchi-Levi, and Williamson with a 3-approximation achieved through linear programming relaxation [9]. Subsequent advancements by Goemans and Williamson, and later by Archer, Bateni, Hajiaghayi, and Karloff, refined the approximation ratio to 2 and 1.967, respectively [6, 15]. Our work contributes to the ongoing research efforts in the field by presenting a 1.7994-approximation algorithm for the PCST problem, improving upon the previous best-known ratio of 1.967 established in 2009 [6]. This achievement marks progress in enhancing the efficiency of solutions for this long-standing open problem.

Besides PCST, the Prize-collecting Steiner Forest (PCSF) problem stands as another open area of research in combinatorial optimization. In PCSF, the objective is to efficiently connect pairs of vertices, each of which has an associated penalty for remaining unconnected. Work on this area began with the work of Agrawal, Klein, and Ravi [1, 2]. Following this, 3-approximation algorithms were developed using cost-sharing and iterative rounding, respectively [16, 17]. Progress continued with Hajiaghayi and Jain's 2.54-approximation algorithm [18], and more recently, the 2-approximation by Ahmadi, Gholami, Hajiaghayi, Jabbarzade, and Mahdavi [3].

Another related problem, the Prize-collecting version of the classic Traveling Salesman Problem (PCTSP), focuses on optimizing the length of the route taken while also accounting for penalties associated with unvisited cities. Although the natural LP formulations for PCTSP and PCST share lots of similarities, PCTSP has experienced considerably more progress. The first breakthrough in breaking the

barrier of 2 for PCST also introduced a 1.98-approximation algorithm for PCTSP [6]. Subsequently, Goemans improved this to a 1.91 approximation factor [14]. The approximation factor was further improved to 1.774 by Blauth and Nägele [11], and most recently, to 1.599 by Blauth, Klein, and Nägele [10]. These advances in PCSF and PCTSP underline the significance and continuous research interest in prize-collecting problems.

1.1 Contribution Overview

In this paper, we focus on rooted PCST where a designated vertex, denoted as *root*, must be included in the solution tree. The objective is to connect other vertices to *root* or pay their penalty. The general PCST and its rooted variant are equivalent. Solving the general PCST involves iterating over all vertices as potential roots and solving the rooted variant for each. Conversely, we can adapt the general version to address rooted PCST by assigning an infinite penalty to the root vertex, ensuring its inclusion in the optimal solution. This two-way equivalence is crucial for our approach, allowing us to concentrate on rooted PCST and extend our findings to the general case. In the rooted version, we define an instance of the PCST problem using a graph $G = (V, E, c)$ with edge weight function $c : E \rightarrow \mathbb{R}_{\geq 0}$, root vertex *root*, and penalty function $\pi : V \rightarrow \mathbb{R}_{\geq 0}$. In the penalty function, while only non-root vertices have actual penalties, we include *root* in the domain of π and assume it has penalty $\pi(\text{root}) = \infty$. This does not affect the actual costs of solutions, but simplifies our statements by adding consistency.

In designing our algorithm, we utilize the recursive approach introduced by [3]. The concept involves running a baseline algorithm with a higher approximation factor on PCST to get an initial solution. We then account for the penalties associated with any vertices identified by the baseline algorithm, paying these penalties, and subsequently removing their penalties from consideration. Next, we apply a Steiner tree algorithm to the remaining vertices to obtain another solution. We then call our algorithm recursively with the adjusted penalties. At each recursive step, two algorithms are executed on the current input, each producing a tree as a solution. Our procedure aggregates all solutions generated during the recursion process and selects the one with the lowest cost as the final output.

We give a quick overview of the major components of our algorithm here.

Goemans and Williamson Algorithm for PCST. We use a slightly modified version of the algorithm introduced by Goemans and Williamson in [15] as the baseline algorithm in the recursive process. We briefly present this algorithm for completeness. Throughout the paper, we refer to this algorithm as PCSTGW and denote the solution found by the algorithm as GW.

Let's assume that each edge of the input graph G is a curve with a length equal to its cost. We want to build a spanning tree F , which starts as a forest during our algorithm and transforms into a tree by the end of the algorithm. We then remove certain edges from this tree to obtain our final tree T and pay penalties for every vertex outside T .

To run our algorithm, we define C as the connected components of F , and active sets $ActS$ as subsets of C . Initially, both C and $ActS$ consist of single-member sets, with each vertex belonging to exactly

one set. We assign a unique color to each vertex of the graph, with the value $\pi(v)$ representing the total duration that color v can be used. As $\pi(\text{root}) = \infty$, the color of root can be used without any limitation.

At any moment, each active set colors its adjacent edges (edges with exactly one endpoint in that set) with the color of one of its vertices that still has available color.

Every time an edge becomes fully colored, it will be added to F , and subsequently, the connected components of F and active sets will be updated. Moreover, if all vertices in an active set run out of color, the active set becomes deactivated and will be considered a dead set, along with all the vertices inside it. We continue this process until all vertices are connected to the root. Note that this is ensured since the root has an infinite amount of color.

After the completion of this process, we remove some edges from F to obtain T . We will select every dead set S that cuts exactly one edge of F and remove all vertices in S from F to obtain T . Every live vertex, which refers to vertices not marked as dead, will be connected to the root in T , along with some dead vertices. In fact, the tree T is the *smallest* subtree of F that contains all live vertices, including root , and every vertex whose color has been used in T .

Steiner Tree Algorithm for PCST. Here we want to construct a new solution ST based on the outcome of PCSTGW. During the execution of the PCSTGW algorithm, certain active sets and their vertices may reach a dead state, leaving them incapable of coloring edges as their vertices have used all of their colors. In such cases, it is reasonable to pay their penalties and subsequently remove them from consideration. This decision makes sense, as connecting these vertices to other vertices requires excessive costs compared to their penalties.

In the GW solution, some of these dead vertices may eventually connect to the root when other active sets link to them, and we utilize these dead vertices to connect live vertices to root . However, in ST , we pay the penalties of all dead vertices and seek a tree that efficiently connects other vertices to root . The problem of finding a minimum tree that connects a set of vertices to root is known as the Steiner Tree problem, and we employ the best-known algorithm for this, assuming it has a p approximation factor, which currently is $\ln(4) + \epsilon$ [12].

Improving the approximation factor of the Steiner Tree algorithm would consequently enhance the approximation factor of our PCST algorithm. It's worth noting that one might suggest paying penalties only for vertices that the GW solution pays penalties for, rather than all dead vertices. However, the GW solution may connect all vertices to the root and influence the Steiner Tree algorithm to establish connections for every vertex. This constraint restricts the algorithm's flexibility in exploring alternative tree structures.

Iterative algorithm. Now, let's explore our iterative algorithm. Our aim is to create an iterative procedure that results in a α -approximation algorithm for PCST. We will discuss the value of α in the future.

At the initiation of our algorithm, we divide the vertex penalties by a constant factor β to obtain π_β . The idea of altering penalties has been used in [5], but they focus on increasing penalties, while we decrease them. The specific value of β will be determined towards the conclusion of our paper. This determination will be based

on the value of p , representing the best-known approximation factor for the Steiner Tree problem, with the goal of minimizing the approximation factor α .

Now, we execute PCSTGW using the modified penalties π_β . Running PCSTGW on π_β provides us with a tree T_{GW} , and paying the penalty of vertices outside T_{GW} yields one solution for the input. Subsequently, we pay the penalty of every vertex that becomes dead during the execution of PCSTGW, set their penalty to zero for the remainder of our algorithm, and connect the remaining vertices using the best-known algorithm for the Steiner tree problem, denoted as STEINERTREE . The tree generated by STEINERTREE , denoted as T_{ST} , presents another solution for the input.

Then, if no vertices with a non-zero penalty become inactive in PCSTGW, indicating that we haven't altered the penalties of vertices at this step, we terminate our algorithm by returning the minimum cost solution between T_{GW} and T_{ST} . Otherwise, we recursively apply this algorithm to the new penalties, and refer the tree of the best solution found by the recursive approach as T_{IT} .

Finally, we select the best solution among T_{GW} , T_{ST} , and T_{IT} . It's important to note that our algorithm essentially identifies two solutions at each iteration and, in the end, selects the solution with the minimum cost among all these alternatives.

In analyzing our algorithm, we focus on its initial step, specifically the first invocation of PCSTGW and STEINERTREE . We categorize vertices based on their status in PCSTGW, distinguishing between those marked as dead or live, and whether their penalties have been paid in both PCSTGW and the optimal solution. Additionally, we classify active sets based on whether they color only one edge or more than one edge of the optimal solution. Through this partitioning, we derive lower bounds for the optimal solution and upper bounds for the solutions T_{GW} and T_{ST} . Leveraging the recursive nature of our algorithm, we establish an upper bound for the solution T_{IT} using induction. Following that, we evaluate how much these solutions deviate from $\alpha \cdot \text{cost}_{\text{OPT}}$.

Next, we show that for $\beta = 1.252$ and $\alpha = 1.7994$, a weighted average of the cost of the three solutions is at most $\alpha \cdot \text{cost}_{\text{OPT}}$. This shows that our algorithm when using this value of β is a 1.7994 approximation of the optimal solution since the minimum cost is lower than any weighted average. We note that throughout our analysis, we do not know the value of α . Instead, we obtain a system of constraints involving α , β , p , and the weights in the weighted average which needs to be satisfied in order for our proof steps to be valid. Then, we find a solution to this system minimizing α to find our approximation guarantee. In this solution, we use $p = \ln(4) + \epsilon$, using the current best approximation factor for the Steiner tree [12]. Finally, we explain the intuition behind certain parts of our algorithm, including why we need to consider all three solutions that we obtain.

Outline. In Section 2, we explain Goemans and Williamson's 2-approximation algorithm for PCST [15], using the coloring schema effectively utilized by [3] for PCSF. Then, in Section 3, we present our iterative algorithm along with its analysis. Finally, in Section 4, we highlight the importance of employing both algorithms in conjunction with the iterative approach to improve the approximation factor.

1.2 Preliminaries

Throughout our paper, we assume without loss of generality that the given graph is connected.

Let T be a subgraph, then $c(T)$ denotes the total cost of edges in T , i.e., $c(T) = \sum_{e \in T} c(e)$.

For a subgraph T , we use $V(T)$ to represent the set of vertices in T , and $\bar{V}(T)$ denotes the set of vertices outside T .

Given a subset of vertices $S \subseteq V$, we define $\pi(S) = \sum_{v \in S} \pi(v)$ as the sum of penalties associated with vertices in S .

For a PCST solution X , we denote its corresponding tree as T_X . Furthermore, we use $cost_X$ to represent the total cost of X , defined as $c(T_X) + \pi(\bar{V}(T_X))$.

2 GOEMANS AND WILLIAMSON ALGORITHM

Here we define a slightly modified version of the algorithm initially proposed by Goemans and Williamson in [15] (hereinafter the GW algorithm) for the sake of completeness of our algorithm. Then we use it as a building block in our algorithm in the next section. We introduce several lemmas stating the properties of the algorithm and its output. We defer the proofs of these lemmas to the full version of our paper.

The algorithm consists of two phases. In the first phase, we simulate a continuous process of vertices growing components around themselves and coloring the edges adjacent to these components at a constant rate. In this process, we imagine each edge e with weight $c(e)$ as a curve of length $c(e)$. Each vertex v has a potential coloring duration equal to its penalty $\pi(v)$. We assume that the root vertex $root$ has $\pi(root) = \infty$, indicating infinite coloring potential. This process of coloring will give us a spanning tree, which we will then trim in the second phase to get a final tree.

During the algorithm, we keep a forest F of tentatively selected edges, a set C of connected components of this forest, and a subset $ActS$ of active sets in C . For each component S in C , we will also store its coloring duration y_S . Initially, the forest F is empty, every vertex is an active set in C , and all y_S values are 0.

At any moment in the process, all active sets color their adjacent edges using the coloring potential of their vertices at the same rate. So, the amount of color on each edge is the total duration its endpoints have been in active sets. We define an edge as fully colored if the combined active time of its endpoints totals at least the length of the edge while they belong to different components. When such an edge between two sets becomes fully colored, it is added to F , and the two sets containing its endpoints are merged, with their coloring potentials summed together. An active set becomes inactive if it runs out of coloring potential. This means that this set and its subsets have used the coloring potential of all the vertices in the set. It may be possible for multiple of these events to happen at the same time, and we would handle them one by one in an arbitrary order. The addition of one edge in the order may prevent the addition of other fully colored edges. However, this can only happen if the latter edge forms a cycle in F , and therefore, the resulting components are independent of the order in which we handle the events. As the component containing $root$ remains active and edges are only added between different components, F will eventually become a spanning tree of G . This marks the completion of the coloring phase.




Figure 1: Illustration of dead sets in the final tree of GW algorithm. The dead sets colored in blue cut multiple edges of F , and removing them would disconnect other vertices so they are not removed. On the other hand, the dead sets colored in red can be safely removed without affecting other vertices.

In the second phase, we will select a subset of F as our Steiner tree and pay the penalties for the remaining vertices. We refer to any active set that becomes inactive as a dead set. Throughout the first phase, we maintain dead sets in DS to utilize them in the second phase. We categorize vertices into dead and live, where a dead vertex is any vertex contained in at least one dead set, and all other vertices are considered live. We store dead vertices in K and return them at the end of PCSTGW since they are used in our iterative algorithm in the next section. For any dead set S , if there is exactly one edge of F cut by S (i.e., $|\delta(S) \cap F| = 1$), we remove this edge and all the edges in F that have both endpoints in S . This effectively removes S from the tree and disconnects its vertices from the root. We repeat this process until no dead set with this property can be found. Figure 1 illustrates how dead sets may be removed.

As each operation in the second phase disconnects only the selected dead set from the root, the final result will be a tree T that contains all the live vertices, including $root$. We pay the penalties for the vertices outside the tree, which are all dead vertices belonging to the dead sets we removed in the second phase. Algorithm 1 provides a pseudocode that implements this process.

To facilitate our analysis throughout the paper, we assume that each vertex is associated with a specific color. During the coloring process of an active set S , we assign each moment of coloring to a vertex $v \in S$ with non-zero remaining coloring potential and utilize its color on the adjacent edges. For consistency, we choose vertex v based on a fixed ordering of the vertices in V where $root$ comes first. So, a set S containing $root$ will always assign its coloring to $root$. We note that a set S can not use the color of a vertex that is already dead. Based on this assignment, we define the following values:

Definition 1. For each vertex v , we define its total coloring duration y_v , and the coloring duration assigned to it by a set S as y_{Sv} :

- y_{Sv} = total coloring duration using color v in set S
- $y_v = \sum_{S \subseteq V; v \in S} y_{Sv}$

Note that for every vertex $v \in V$, we have $\sum_{v \in S} y_{Sv} = y_S$.

We bound the cost of both the chosen tree and the penalty of the dead vertices in the following lemmas. Proofs of these lemmas are in the full version of our paper given their similarity to [15].

Lemma 2. Let T be the tree returned by Algorithm 1. We can bound the total weight of this tree by

$$c(T) \leq 2 \cdot \sum_{\substack{S \subseteq V - \{\text{root}\}; \\ S \cap V(T) \neq \emptyset}} y_S = 2 \cdot \sum_{v \in V(T) - \{\text{root}\}} y_v.$$

Lemma 3. For any vertex $v \in V$, we have $y_v \leq \pi(v)$. Furthermore, if $v \in K$ which means it is a dead vertex, we have $y_v = \pi(v)$.

Lemma 4. Any vertex $v \notin V(T)$ is a dead vertex.

Lemmas 2, 3, and 4 immediately conclude the following lemma.

Lemma 5. The total cost of the GW algorithm is bounded by

$$\text{cost}_{\text{GW}} = c(T) + \pi(\overline{V(T)}) \leq 2 \cdot \sum_{v \in V(T) - \{\text{root}\}} y_v + \sum_{v \notin V(T)} y_v.$$

We note that Lemma 5 can be used to prove that the GW algorithm achieves a 2-approximation by showing that the optimal solution has cost at least $\sum_{v \in V - \{\text{root}\}} y_v$. We prove a stronger version of this fact in Lemma 13.

In addition to the above lemmas on the cost of the solution and its connection to the coloring, we also prove the following lemma. This lemma will help in our analysis in Section 3.1, where we use it to introduce an upper bound for the cost of the optimal Steiner tree connecting all the live vertices in a call to the GW algorithm. The proof of this lemma is also provided in the full version of our paper.

Lemma 6. Let $I = (G, \text{root}, \pi)$ and $I' = (G', \text{root}, \pi)$ be instances of PCST, where G' is obtained from G by adding a set of edges E_0 with weight 0 from root to a set of vertices U . Let y_v be the coloring duration for vertex v in a run of the GW algorithm on I , and let K be the set of dead vertices in this run. Let y'_v and K' be the corresponding values when running the GW algorithm on instance I' using the same order to assign coloring duration to vertices. We have

$$\begin{aligned} y'_v &\leq y_v, \\ y'_v &= 0 \text{ if } v \in U, \\ &\text{and} \\ K' &\subseteq K. \end{aligned}$$

Algorithm 1 PCSTGW($I = (G, \text{root}, \pi)$): GW Algorithm

Input: undirected graph $G = (V, E, c)$ with edge costs $c : E \rightarrow \mathbb{R}_{\geq 0}$, root root , and penalties $\pi : V \rightarrow \mathbb{R}_{\geq 0}$.

Output: Subtree T of G containing root , alongside a set K of dead vertices.

```

1: Initialize  $F$  as an empty forest
2: Initialize  $\text{Act}S$  and  $C$  as  $\{\{v\} \mid v \in V\}$ 
3: Set  $y_S \leftarrow 0$  for all  $S \in \text{Act}S$ 
4:  $K \leftarrow \emptyset$ 
5:  $DS \leftarrow \emptyset$ 
6: while  $|C| > 1$  do
7:    $\Delta_1 \leftarrow \min_{S \in \text{Act}S} (\sum_{v \in S} \pi(v) - \sum_{S' \subseteq S} y_{S'})$ 
8:    $\Delta_2 \leftarrow \min_{e=(u,v) \in E; e \cup F \text{ is a forest}} (\frac{c_e - \sum_{S: e \in \delta(S)} y_S}{|\{S \in \text{Act}S \mid u \in S \vee v \in S\}|})$ 
9:    $\Delta \leftarrow \min(\Delta_1, \Delta_2)$ 
10:  for  $S \in \text{Act}S$  do
11:     $y_S \leftarrow y_S + \Delta$ 
12:  end for
13:  if  $\Delta_1 \leq \Delta_2$  then
14:    Find a set  $S$  minimizing  $\Delta_1$ 
15:     $\text{Act}S \leftarrow \text{Act}S - \{S\}$ 
16:     $K \leftarrow K \cup S$ 
17:     $DS \leftarrow DS \cup \{S\}$ 
18:  else
19:    Find an edge  $e = (u, v)$  minimizing  $\Delta_2$ 
20:     $F \leftarrow F \cup e$ 
21:    Update  $C$  and  $\text{Act}S$  accordingly
22:  end if
23: end while
24: Extract  $T$  from  $F$  by repeatedly removing dead sets in  $DS$  that cut a single edge in  $F$ 
25: return  $(T, K)$ 

```

3 THE ITERATIVE ALGORITHM

In this section, we present our iterative algorithm which is described in Algorithm 2. In Section 3.1 we give an analysis for this algorithm.

Our algorithm makes use of the PCSTGW procedure from Algorithm 1 as a fundamental component. Additionally, we employ an approximation algorithm for the Steiner tree problem to improve the approximation factor. This can be any approximation algorithm for the Steiner tree problem. We denote the approximation factor for this algorithm as p . Whenever we require this p -approximation solution for the Steiner tree, we invoke the procedure named STEINERTREE. As our final approximation factor will depend on p , we will use the current best approximation algorithm for Steiner Tree [12] with $p = \ln(4) + \epsilon$ in our analysis. In addition, our algorithm depends on a constant β which we will fix later in Section 3.2 to optimize the approximation ratio.

Our algorithm, as described in Algorithm 2, identifies three solutions for the given PCST instance $I = (G, \text{root}, \pi)$. Subsequently, we opt for the solution with the minimum cost as the final solution.

First, we construct the instance $I_\beta = (G, \text{root}, \pi_\beta)$ from I by replacing π_v with $\frac{\pi_v}{\beta}$ for all vertices. One solution named ‘‘GW’’ for instance I , with tree T_{GW} , can be obtained by invoking procedure PCSTGW (Line 3) on instance I_β , buying edges of T_{GW} and paying penalties for vertices in $\overline{V(T_{\text{GW}})}$. From the definition of I_β , we can

conclude that $\pi(\overline{V(T_{GW})}) = \beta\pi_\beta(\overline{V(T_{GW})})$. As stated in Section 2, in addition to T_{GW} , procedure PCSTGW also returns a set of vertices, K , which represents dead vertices during the coloring process.

Another solution for instance I named “ST” is obtained by retrieving a Steiner tree T_{ST} in graph G for the set of terminals $L := V \setminus K$ which are the live vertices in the output of the GW algorithm. This solution is found using the procedure STEINERTREE and is therefore a p -approximation of the minimum Steiner tree on this terminal set. We pay the penalties for the vertices outside T_{ST} , which will be a subset of K .

If K is empty, the algorithm immediately returns the solution with the lower total cost between the two obtained solutions. Otherwise, a third solution named “IT”, denoted as T_{IT} , is obtained through a recursive call on a simplified instance R . The simplified instance is formed through a process of adjusting penalties. We set the penalties for the vertices in K , which are the dead vertices in the result of the PCSTGW procedure, to zero while maintaining the penalty for the live vertices L , as indicated in Lines 11 through 12.

As a final step, the algorithm simply selects and returns the solution with the lowest cost. To help with the comparison of these three solutions, the algorithm calculates the values $cost_{GW} = c(T_{GW}) + \pi(\overline{V(T_{GW})})$, $cost_{ST} = c(T_{ST}) + \pi(\overline{V(T_{ST})})$, and $cost_{IT} = c(T_{IT}) + \pi(\overline{V(T_{IT})})$, representing the costs of the solutions (as indicated in Lines 4, 7, and 14).

Algorithm 2 IPCST($I = (G, root, \pi)$): Iterative PCST algorithm

Input: Undirected graph $G = (V, E, c)$ with edge costs $c : E \rightarrow \mathbb{R}_{\geq 0}$, root $root$, and penalties $\pi : V \rightarrow \mathbb{R}_{\geq 0}$.

Output: Subtree T of G containing $root$.

- 1: Construct π_β by dividing all penalties by β .
 - 2: Construct the PCST instance $I_\beta = (G, root, \pi_\beta)$.
 - 3: $T_{GW}, K \leftarrow \text{PCSTGW}(I_\beta)$
 - 4: $cost_{GW} \leftarrow c(T_{GW}) + \pi(\overline{V(T_{GW})})$
 - 5: $L \leftarrow \{v : v \in V, v \notin K\}$
 - 6: $T_{ST} \leftarrow \text{STEINERTREE}(G, L)$
 - 7: $cost_{ST} \leftarrow c(T_{ST}) + \pi(\overline{V(T_{ST})})$
 - 8: **if** $\pi(K) = 0$ **then**
 - 9: **return** T_X where $cost_X$ is minimum among $X \in \{GW, ST\}$
 - 10: **end if**
 - 11: Construct π' by adjusting π through the assignment of penalties for vertices in K to 0.
 - 12: Construct the PCST instance $R = (G, root, \pi')$.
 - 13: $T_{IT} \leftarrow \text{IPCST}(R)$
 - 14: $cost_{IT} \leftarrow c(T_{IT}) + \pi(\overline{V(T_{IT})})$
 - 15: **return** T_X where $cost_X$ is minimum among $X \in \{GW, ST, IT\}$
-

3.1 Analysis

For an arbitrary instance $I = (G, root, \pi)$ in PCST, our aim is to analyze the approximation factor achieved by Algorithm 2. We compare the output of IPCST on I with an optimal solution OPT for the instance I . We denote the tree selected in OPT as T_{OPT} . Then, the cost of OPT is given by $cost_{OPT} = c(T_{OPT}) + \pi(\overline{V(T_{OPT})})$.

We use an inductive approach to analyze the algorithm, where we focus on a single call of the algorithm and find upper bounds

for each of our three solutions and a lower bound for the optimal solution OPT. To find these lower and upper bounds, we make use of the coloring done by the GW algorithm on instance I_β and the values y_S, y_{Sv} , and y_v relating to this coloring process. In addition, we establish an upper bound for the solution obtained from the recursive call based on the induction hypothesis. In our inductive analysis, we only consider one individual call to the procedure at each time, to analyze either the induction base or the induction step. So, all the variables used in the analysis will relate to the algorithm’s variables in the specific call we are analyzing. This includes the trees T_{GW}, T_{ST} , and T_{IT} , and the live and dead vertices L and K .

We note that in our induction, we do not initially know the value of the approximation factor α which we want to prove the algorithm achieves. Instead, we use α as a variable in our inequalities, and this leads to a system of constraints involving α that need to be satisfied for our induction to prove an α approximation guarantee. These inequalities involve not only the approximation factor α which we seek to find but also the parameter β which defines the behavior of our algorithm. Throughout the analysis, we assume that $\beta \leq 2$. We justify this assumption in Subsection 4.1 by showing that values of $\beta > 2$ cannot lead to a better than 2 approximation. To determine our approximation factor α , we consider the range $p \leq \alpha \leq 2$. This range is chosen because we cannot assume that our algorithm performs better than the Steiner tree algorithm, which we use as a component. Additionally, our solution is guaranteed to be at least as good as the 2-approximation provided by the GW algorithm.

In the first step, we categorize non-root vertices based on the output of PCSTGW(I_β) and OPT. This categorization helps us establish more precise bounds for the solutions by enabling a more tailored analysis within each category.

Definition 7. For an instance I , OPT partition vertices into two sets: $V(T_{OPT})$ and $\overline{V(T_{OPT})}$. PCSTGW(I_β) also partitions vertices into two sets: L and K . We define four sets to categorize the vertices, excluding $root$, based on these two partitions:

$$\begin{aligned} \mathcal{A} &= V(T_{OPT}) \cap L - \{root\} & \mathcal{B} &= V(T_{OPT}) \cap K \\ \mathcal{C} &= \overline{V(T_{OPT})} \cap L & \mathcal{D} &= \overline{V(T_{OPT})} \cap K \end{aligned}$$

Table 1: This table illustrates the categories of vertices excluding $root$.

		PCSTGW(I_β)	
		Live vertices	Dead vertices
Optimal Solution	Connected	\mathcal{A}	\mathcal{B}
	Penalty paid	\mathcal{C}	\mathcal{D}

Using the coloring scheme of PCSTGW(I_β), we introduce the following values to represent the total duration of coloring with vertices in these sets:

$$\begin{aligned} r_{\mathcal{A}} &= \sum_{v \in \mathcal{A}} y_v & r_{\mathcal{B}} &= \sum_{v \in \mathcal{B}} y_v \\ r_{\mathcal{C}} &= \sum_{v \in \mathcal{C}} y_v & r_{\mathcal{D}} &= \sum_{v \in \mathcal{D}} y_v \end{aligned}$$

Definition 8 (Connected and unconnected dead vertices). For an instance I , based on Definition 7, the sets \mathcal{B} and \mathcal{D} represent dead vertices in the output of PCSTGW(I_β). We further divide set \mathcal{B} into \mathcal{B}' and \mathcal{B}'' , and set \mathcal{D} into \mathcal{D}' and \mathcal{D}'' , based on whether they are connected to the root at the end of the PCSTGW(I_β) procedure. Let \mathcal{B}' and \mathcal{D}' be the subsets of \mathcal{B} and \mathcal{D} , respectively, representing the vertices connected to the root. Similarly, \mathcal{B}'' and \mathcal{D}'' are the subsets of \mathcal{B} and \mathcal{D} , respectively, indicating the vertices not connected to the root at the end of the procedure.

$$\begin{aligned}\mathcal{B}' &= \mathcal{B} \cap V(T_{\text{GW}}) = V(T_{\text{OPT}}) \cap K \cap V(T_{\text{GW}}) \\ \mathcal{B}'' &= \mathcal{B} \cap \overline{V(T_{\text{GW}})} = V(T_{\text{OPT}}) \cap K \cap \overline{V(T_{\text{GW}})} \\ \mathcal{D}' &= \mathcal{D} \cap V(T_{\text{GW}}) = \overline{V(T_{\text{OPT}})} \cap K \cap V(T_{\text{GW}}) \\ \mathcal{D}'' &= \mathcal{D} \cap \overline{V(T_{\text{GW}})} = \overline{V(T_{\text{OPT}})} \cap K \cap \overline{V(T_{\text{GW}})}\end{aligned}$$

Subsequently, we define $r_{\mathcal{B}'}$, $r_{\mathcal{B}''}$, $r_{\mathcal{D}'}$, and $r_{\mathcal{D}''}$ as the total duration of coloring with vertices in sets \mathcal{B}' , \mathcal{B}'' , \mathcal{D}' , and \mathcal{D}'' , respectively.

$$\begin{aligned}r_{\mathcal{B}'} &= \sum_{v \in \mathcal{B}'} y_v & r_{\mathcal{B}''} &= \sum_{v \in \mathcal{B}''} y_v \\ r_{\mathcal{D}'} &= \sum_{v \in \mathcal{D}'} y_v & r_{\mathcal{D}''} &= \sum_{v \in \mathcal{D}''} y_v\end{aligned}$$

It is trivial to see that $r_{\mathcal{D}} = r_{\mathcal{D}'} + r_{\mathcal{D}''}$ as $\mathcal{D}' \cup \mathcal{D}'' = \mathcal{D}$ and $\mathcal{D}' \cap \mathcal{D}'' = \emptyset$. Similarly, $r_{\mathcal{B}} = r_{\mathcal{B}'} + r_{\mathcal{B}''}$.

Definition 9 (Single-edge and multi-edge sets). For an instance I , we call a set $S \subseteq V$ a single-edge set if $|\delta(S) \cap T_{\text{OPT}}| = 1$ and a multi-edge set if $|\delta(S) \cap T_{\text{OPT}}| > 1$ (Illustrated in Figure 2). We assign each moment of coloring with colors of vertices in \mathcal{B} which are inside a single-edge set to b_1 , and those in a multi-edge set to b_2 . These definitions are as follows:

$$\begin{aligned}b_1 &= \sum_{v \in \mathcal{B}} \sum_{|\delta(S) \cap T_{\text{OPT}}|=1} y_{Sv} \\ b_2 &= \sum_{v \in \mathcal{B}} \sum_{|\delta(S) \cap T_{\text{OPT}}|>1} y_{Sv}\end{aligned}$$

Note that $r_{\mathcal{B}} = b_1 + b_2$, as every vertex in \mathcal{B} is connected to *root* in the optimal solution. Therefore, with each moment of coloring involving vertices in \mathcal{B} , the corresponding active set cuts an edge belonging to the path from that vertex to *root* in the optimal solution.

Lemma 10. For any vertex $v \in V$, we have $\beta y_v \leq \pi(v)$. Furthermore, if $v \in \mathcal{B} \cup \mathcal{D}$, which means is a dead vertex, we have $\beta y_v = \pi(v)$.

PROOF. Since we run PCSTGW on π_β in Line 3, we can use Lemma 3 using penalties π_β . That means, for any vertex $v \in V$, we have $y_v \leq \pi_\beta(v)$, and if v is a dead vertex, we have $y_v = \pi_\beta(v)$. Since in Line 1, we set $\pi_\beta(v) = \frac{\pi(v)}{\beta}$, we can conclude the lemma. \square

Now for a given instance I , we derive lower bounds on the optimal solution using terms defined earlier. We use a similar approach that is used in [3] to bound the optimal solution.




Figure 2: Illustration of single-edge set vs. multi-edge set in T_{OPT} . The red set is a single-edge set, but the blue one is a multi-edge set.

Lemma 11. We can bound the cost of the optimal solution in terms of the cost of its tree as follows:

$$\text{cost}_{\text{OPT}} \geq c(T_{\text{OPT}}) + \beta r_{\mathcal{C}} + \beta r_{\mathcal{D}}.$$

PROOF. According to the definition of cost in PCST, we can determine the cost of the optimal solution by separately calculating the weight of its tree and the penalties it pays. Additionally, based on Definition 7, we have $\overline{V(T_{\text{OPT}})} = \mathcal{C} \cup \mathcal{D}$. Utilizing these two observations, we can establish an upper bound for cost_{OPT} as follows:

$$\begin{aligned}\text{cost}_{\text{OPT}} &= c(T_{\text{OPT}}) + \pi(\overline{V(T_{\text{OPT}})}) \\ &= c(T_{\text{OPT}}) + \sum_{v \in \overline{V(T_{\text{OPT}})}} \pi(v) \\ &= c(T_{\text{OPT}}) + \sum_{v \in \mathcal{C}} \pi(v) + \sum_{v \in \mathcal{D}} \pi(v) \quad (\mathcal{C} \cap \mathcal{D} = \emptyset) \\ &\geq c(T_{\text{OPT}}) + \sum_{v \in \mathcal{C}} \beta y_v + \sum_{v \in \mathcal{D}} \beta y_v \quad (\text{Lemma 10}) \\ &= c(T_{\text{OPT}}) + \beta r_{\mathcal{C}} + \beta r_{\mathcal{D}} \quad (\text{Definition 7})\end{aligned}$$

\square

Based on Lemma 11, we can easily conclude the following corollary which bounds the weight of the optimal solution tree using the cost of the optimal solution.

Corollary 12. We can bound the cost of optimal solution's tree as follows:

$$c(T_{\text{OPT}}) \leq \text{cost}_{\text{OPT}} - \beta r_{\mathcal{C}} - \beta r_{\mathcal{D}}.$$

Now we use Lemma 11, to expand the bound of the optimal solution.

Lemma 13. We can establish a lower bound for the optimal solution as follows:

$$\text{cost}_{\text{OPT}} \geq r_{\mathcal{A}} + b_1 + 2b_2 + \beta r_{\mathcal{C}} + \beta r_{\mathcal{D}}$$

PROOF. First, we demonstrate that $r_{\mathcal{A}} + b_1 + 2b_2$ is a lower bound for $c(T_{\text{OPT}})$. To achieve this, for any set S , we define $d_{\text{OPT}}(S)$ as the number of edges of T_{OPT} that are colored by S . Given that each portion of an edge will be colored at most once, and each set $S \subseteq V$ colors $d_{\text{OPT}}(S) \cdot y_S$ of the optimal solution, we can derive a lower

bound for $c(T_{\text{OPT}})$ based on the proportion of the colored edges in T_{OPT} .

$$\begin{aligned}
c(T_{\text{OPT}}) &\geq \sum_{S \subseteq V} d_{\text{OPT}}(S) \cdot y_S \\
&= \sum_{S \subseteq V} \sum_{v \in S} d_{\text{OPT}}(S) \cdot y_{Sv} \quad (y_S = \sum_{v \in S} y_{Sv}) \\
&= \sum_{v \in V} \sum_{\substack{S \subseteq V \\ v \in S}} d_{\text{OPT}}(S) \cdot y_{Sv} \\
&\geq \sum_{v \in \mathcal{A} \cup \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S}} d_{\text{OPT}}(S) \cdot y_{Sv} \quad (\mathcal{A} \cap \mathcal{B} = \emptyset, \mathcal{A} \cup \mathcal{B} \subseteq V) \\
&\geq \sum_{v \in \mathcal{A}} \sum_{\substack{S \subseteq V \\ v \in S}} d_{\text{OPT}}(S) \cdot y_{Sv} + \sum_{v \in \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S}} d_{\text{OPT}}(S) \cdot y_{Sv}.
\end{aligned}$$

Furthermore, for any vertex v in \mathcal{A} or \mathcal{B} , based on Definition 7, there exists a path from v to $root$ in T_{OPT} . Also, for every set $S \subseteq V$ where $y_{Sv} > 0$, we know $root \notin S$ otherwise all coloring of set S would be assigned to $root$. Using these two observations, we can infer that at least one edge of T_{OPT} is colored by S , resulting in $d_{\text{OPT}}(S) \geq 1$.

Therefore, for vertices in \mathcal{A} , we have:

$$\begin{aligned}
\sum_{v \in \mathcal{A}} \sum_{\substack{S \subseteq V \\ v \in S}} d_{\text{OPT}}(S) \cdot y_{Sv} &\geq \sum_{v \in \mathcal{A}} \sum_{\substack{S \subseteq V \\ v \in S}} y_{Sv} \\
&= \sum_{v \in \mathcal{A}} y_v \quad (\text{Definition 1}) \\
&= r_{\mathcal{A}}. \quad (\text{Definition 7})
\end{aligned}$$

Also, for vertices in \mathcal{B} , we have:

$$\begin{aligned}
\sum_{v \in \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S}} d_{\text{OPT}}(S) \cdot y_{Sv} &= \sum_{v \in \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S \\ d_{\text{OPT}}(S)=1}} d_{\text{OPT}}(S) \cdot y_{Sv} \\
&\quad + \sum_{v \in \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S \\ d_{\text{OPT}}(S)>1}} d_{\text{OPT}}(S) \cdot y_{Sv} \\
&\geq \sum_{v \in \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S \\ d_{\text{OPT}}(S)=1}} y_{Sv} + \sum_{v \in \mathcal{B}} \sum_{\substack{S \subseteq V \\ v \in S \\ d_{\text{OPT}}(S)>1}} 2y_{Sv} \\
&= b_1 + 2b_2. \quad (\text{Definition 9})
\end{aligned}$$

Combining all together, we obtain:

$$c(T_{\text{OPT}}) \geq r_{\mathcal{A}} + b_1 + 2b_2.$$

By using this bound along with Lemma 11, we can bound $cost_{\text{OPT}}$.

$$\begin{aligned}
cost_{\text{OPT}} &\geq c(T_{\text{OPT}}) + \beta r_{\mathcal{C}} + \beta r_{\mathcal{D}} \quad (\text{Lemma 11}) \\
&\geq r_{\mathcal{A}} + b_1 + 2b_2 + \beta r_{\mathcal{C}} + \beta r_{\mathcal{D}}.
\end{aligned}$$

□

In the next lemma, we establish a bound for the GW solution. The proof is included in the full version of our paper.

Lemma 14. The following bound holds for the cost of the solution returned by the output of PCSTGW(I_β) for instance I :

$$cost_{\text{GW}} \leq 2r_{\mathcal{A}} + 2r_{\mathcal{B}} + 2r_{\mathcal{C}} + 2r_{\mathcal{D}}.$$

We restate this upper bound in terms of the variable α and the cost of the optimal solution $cost_{\text{OPT}}$ using Lemma 13.

Lemma 15. The following bound holds for the cost of the solution returned by the output of PCSTGW(I_β) for instance I :

$$\begin{aligned}
cost_{\text{GW}} &\leq \alpha \cdot cost_{\text{OPT}} + (2 - \alpha)r_{\mathcal{A}} + (2 - \alpha)b_1 + (2 - 2\alpha)b_2 \\
&\quad + (2 - \alpha\beta)r_{\mathcal{C}} + (2 - \alpha\beta)r_{\mathcal{D}}.
\end{aligned}$$

PROOF. We can directly apply Lemma 13 to the previous bound obtained in the preceding Lemma 14.

$$\begin{aligned}
cost_{\text{GW}} &\leq 2r_{\mathcal{A}} + 2r_{\mathcal{B}} + 2r_{\mathcal{C}} + 2r_{\mathcal{D}} \quad (\text{Lemma 14}) \\
&\leq 2r_{\mathcal{A}} + 2(b_1 + b_2) + 2r_{\mathcal{C}} + 2r_{\mathcal{D}} \quad (\text{Lemma 13}) \\
&\quad + \alpha \cdot (cost_{\text{OPT}} - r_{\mathcal{A}} - b_1 - 2b_2 - \beta r_{\mathcal{C}} - \beta r_{\mathcal{D}}) \\
&\leq \alpha \cdot cost_{\text{OPT}} + (2 - \alpha)r_{\mathcal{A}} + (2 - \alpha)b_1 + (2 - 2\alpha)b_2 \\
&\quad + (2 - \alpha\beta)r_{\mathcal{C}} + (2 - \alpha\beta)r_{\mathcal{D}}.
\end{aligned}$$

□

Next, we bound the cost of the ST solution. For a set S , let $T_{\text{OPT}'_L}$ denote the minimum cost Steiner tree on this set. In the following lemma, we relate the cost of the ST solution to the cost of $T_{\text{OPT}'_L}$.

Lemma 16. For instance I , we can bound the cost of the solution returned by the output of ST as follows:

$$cost_{\text{ST}} \leq p \cdot c(T_{\text{OPT}'_L}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}}.$$

PROOF. Since in T_{ST} , we are connecting every vertex in L to $root$, using a Steiner tree algorithm with an approximation factor of p , the cost of the tree T_{ST} can be bounded by

$$c(T_{\text{ST}}) \leq p \cdot c(T_{\text{OPT}'_L}).$$

Moreover, as all vertices in L are connected to $root$, the vertices for which we need to pay penalties for this solution form a subset of K , i.e., $\overline{V(T_{\text{ST}})} \subseteq K$. Furthermore, by Definition 7 we have:

$$\begin{aligned}
\mathcal{B} \cup \mathcal{D} &= (V(T_{\text{OPT}}) \cap K) \cup \overline{(V(T_{\text{OPT}}) \cap K)} \quad (\text{Definition 7}) \\
&= K
\end{aligned}$$

Now, we can bound the penalty paid by the ST solution.

$$\begin{aligned}
\pi(\overline{V(T_{\text{ST}})}) &\leq \pi(K) \\
&= \pi(\mathcal{B} \cup \mathcal{D}) \\
&= \sum_{v \in \mathcal{B} \cup \mathcal{D}} \pi(v) \\
&= \sum_{v \in \mathcal{B} \cup \mathcal{D}} \beta y_v \quad (\text{Lemma 10}) \\
&\leq \sum_{v \in \mathcal{B}} \beta y_v + \sum_{v \in \mathcal{D}} \beta y_v \\
&= \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} \quad (\text{Definition 7})
\end{aligned}$$

Finally, we use these bounds to complete the proof

$$cost_{\text{ST}} = c(T_{\text{ST}}) + \pi(\overline{V(T_{\text{ST}})}) \leq p \cdot c(T_{\text{OPT}'_L}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}}.$$

□

We now provide an upper bound for the cost of $T_{\text{OPT}'_L}$ based on the cost of T_{OPT} to obtain our main upper bound for ST.

Lemma 17. For the minimum cost Steiner tree $T_{\text{OPT}}'_L$ on L , we have

$$c(T_{\text{OPT}}'_L) \leq c(T_{\text{OPT}}) + 2r_C + 2r_{\mathcal{D}}.$$

PROOF. We construct a new instance $I'_\beta = (G', \text{root}, \pi_\beta)$ where G' is obtained from G by adding a set E_0 of edges of weight 0 from root to every vertex in $U = \mathcal{A} \cup \mathcal{B} = V(T_{\text{OPT}}) - \{\text{root}\}$. Let T'_{GW} be the resulting tree and y'_v be the coloring duration for the vertices in this process assuming we assign the colors in the same way as we did when running the GW algorithm on I_β . By Lemma 6, $y'_v \leq y_v$ for all vertices in $C \cup \mathcal{D}$. In addition, we have $y'_v = 0$ for all vertices in $U = \mathcal{A} \cup \mathcal{B}$. Then, using Lemma 2 we can bound the cost of T'_{GW} as

$$\begin{aligned} c(T'_{\text{GW}}) &\leq 2 \sum_{v \in V(T'_{\text{GW}}) - \{\text{root}\}} y'_v && \text{(Lemma 2)} \\ &\leq 2 \sum_{v \in V - \{\text{root}\}} y'_v && (V(T'_{\text{GW}}) \subseteq V) \\ &\leq 2 \sum_{v \in \mathcal{A} \cup \mathcal{B}} y'_v + 2 \sum_{v \in C \cup \mathcal{D}} y'_v && (\mathcal{A} \cup \mathcal{B} \cup C \cup \mathcal{D} = V - \{\text{root}\}) \\ &\leq 2 \sum_{v \in C \cup \mathcal{D}} y'_v && (y'_v = 0 \text{ if } v \in \mathcal{A} \cup \mathcal{B} \text{ by Lemma 6}) \\ &\leq 2 \sum_{v \in C \cup \mathcal{D}} y_v && (y'_v \leq y_v \text{ by Lemma 6}) \\ &= 2r_C + 2r_{\mathcal{D}}. && \text{(Definition 7)} \end{aligned}$$

Let K' be the set of dead vertices returned by the GW algorithm on I'_β . Based on Lemma 6, we have $K' \subseteq K$. Therefore, as vertices in $\mathcal{A} \cup C \cup \{\text{root}\} = L$ are not part of K , they cannot be part of K' either and must be live vertices in this run. Lemma 4 means that these vertices are connected by T'_{GW} .

If we remove any edges in E_0 from T'_{GW} , and instead add T_{OPT} , which is a spanning tree on $\mathcal{A} \cup \mathcal{B} \cup \{\text{root}\}$, all the vertices in $V(T'_{\text{GW}})$ will remain connected. So, we get a connected subgraph of G that connects L . The cost of this subgraph is at most

$$\begin{aligned} c((T'_{\text{GW}} - E_0) \cup T_{\text{OPT}}) &\leq c(T_{\text{OPT}}) + c(T'_{\text{GW}}) \\ &\leq c(T_{\text{OPT}}) + 2r_C + 2r_{\mathcal{D}}. \end{aligned}$$

As this subgraph connects L , its cost gives us an upper bound on the cost of the minimum Steiner tree on these vertices. So we have

$$c(T_{\text{OPT}}'_L) \leq c(T_{\text{OPT}}) + 2r_C + 2r_{\mathcal{D}}. \quad \square$$

We combine the last two lemmas to introduce an upper bound for the ST solution. We again state this upper bound in terms of cost_{OPT} and α . Here, we rely on the fact that $\alpha \geq p$ to add a non-negative value to an initial upper bound based on Lemmas 16 and 17.

Lemma 18. For instance I , we can bound the cost of the solution returned by the output of ST as follows:

$$\begin{aligned} \text{cost}_{\text{ST}} &\leq \alpha \cdot \text{cost}_{\text{OPT}} + (p - \alpha)r_{\mathcal{A}} + (p + \beta - \alpha)b_1 \\ &\quad + (2p + \beta - 2\alpha)b_2 + (2p - \alpha\beta)r_C + (2p + \beta - \alpha\beta)r_{\mathcal{D}}. \end{aligned}$$

PROOF. By combining Lemma 16 with Lemma 17, we can derive a new bound for cost_{ST} .

$$\begin{aligned} \text{cost}_{\text{ST}} &\leq p \cdot c(T_{\text{OPT}}'_L) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} && \text{(Lemma 16)} \\ &\leq p(c(T_{\text{OPT}}) + 2r_C + 2r_{\mathcal{D}}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} && \text{(Lemma 17)} \\ &\leq p(\text{cost}_{\text{OPT}} - \beta r_C - \beta r_{\mathcal{D}} + 2r_C + 2r_{\mathcal{D}}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} && \text{(Corollary 12)} \\ &\leq p(\text{cost}_{\text{OPT}} - \beta r_C - \beta r_{\mathcal{D}} + 2r_C + 2r_{\mathcal{D}}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} \\ &\quad + (\alpha - p)(\text{cost}_{\text{OPT}} - r_{\mathcal{A}} - b_1 - 2b_2 - \beta r_C - \beta r_{\mathcal{D}}) && \text{(Lemma 13, } \alpha - p \geq 0) \\ &= \alpha \cdot \text{cost}_{\text{OPT}} + (p - \alpha)r_{\mathcal{A}} + (p + \beta - \alpha)b_1 \\ &\quad + (2p + \beta - 2\alpha)b_2 + (2p - \alpha\beta)r_C + (2p + \beta - \alpha\beta)r_{\mathcal{D}} \quad \square \end{aligned}$$

Now, assume that we want to show that the algorithm achieves an approximation factor of α . Then, to prove this by induction, we need to show two things. First, we need to show that in the base case where the dead set K returned by the GW algorithm has penalty 0 and we do not make a recursive call, our solution is an α approximation. Secondly, we have to demonstrate the induction step. This means that we have to show that if our recursive call on instance R returns an α approximation for this instance, the final returned solution will also be an α approximation. If these two steps are accomplished, then by induction on the number of vertices with non-zero penalties (which decreases with every recursive call), we can prove that our algorithm achieves an α approximation.

So far, we do not know the value of α so we cannot prove the induction steps directly. Instead, we will show that if α satisfies certain constraints then both the base case and the step of induction can be proven for that value of α and therefore our algorithm will give us an α approximation. These constraints are obtained by thinking of α as a variable and then trying to prove the induction base and the induction step for α . Minimizing α in this system of constraints will give us an upper bound on the approximation factor of our algorithm.

In the following, we first assume that the recursive call on R is an α approximation, and bound the iterative solution using this assumption. Then, in Section 3.2 we combine the bounds for the different solutions to find a system of constraints that restrict α . We also consider the constraints that arise from the base case being an α approximation, which turn out to form a subset of the former constraints. Finally, we find the minimum value of α that can satisfy these constraints to obtain our approximation guarantee.

We start with the next lemma, which bounds the cost of the iterative solution's output, assuming that the recursive call returns an α approximate solution for instance R . Here, OPT_R denotes the optimal solution for the PCST instance R .

Lemma 19. For instance I , the cost of the iterative solution, denoted as cost_{IT} , can be bounded as follows:

$$\text{cost}_{\text{IT}} \leq \alpha \cdot \text{cost}_{\text{OPT}_R} + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}},$$

assuming that the recursive call on instance R returns an α approximate solution.

PROOF. Based on our assumption, $\text{IPCST}(R)$ will return a solution that is an α -approximate of the optimal solution of instance R

which we indicate by OPT_R . This gives us the following bound:

$$c(T_{\text{IT}}) + \pi'(\overline{V(T_{\text{IT}})}) \leq \alpha \cdot \text{cost}_{\text{OPT}_R}.$$

However, as $\text{cost}_{\text{IT}} = c(T_{\text{IT}}) + \pi(\overline{V(T_{\text{IT}})})$, we need to establish the relationship between $\pi(\overline{V(T_{\text{IT}})})$ and $\pi'(\overline{V(T_{\text{IT}})})$. The only difference between these functions lies in setting the penalty for vertices in $K = \mathcal{B} \cup \mathcal{D}$ to zero in π' , as indicated in Line 11. Thus, we can conclude that

$$\begin{aligned} \pi(\overline{V(T_{\text{IT}})}) &\leq \pi'(\overline{V(T_{\text{IT}})}) + \pi(\mathcal{B} \cup \mathcal{D}) \\ &= \pi'(\overline{V(T_{\text{IT}})}) + \beta \sum_{v \in (\mathcal{B} \cup \mathcal{D})} y_v \quad (\text{Lemma 10}) \\ &= \pi'(\overline{V(T_{\text{IT}})}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}}. \quad (\text{Definition 7}) \end{aligned}$$

By combining these inequalities, we get

$$\begin{aligned} \text{cost}_{\text{IT}} &= c(T_{\text{IT}}) + \pi(\overline{V(T_{\text{IT}})}) \\ &\leq c(T_{\text{IT}}) + \pi'(\overline{V(T_{\text{IT}})}) + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} \\ &\leq \alpha \cdot \text{cost}_{\text{OPT}_R} + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}}. \end{aligned}$$

□

Lemma 20. For an instance I , we can remove a set of edges with a total length of b_1 from T_{OPT} in such a way that the vertices in \mathcal{A} remain connected to *root*.

PROOF. Consider a moment of coloring with the color of a vertex $v \in \mathcal{B}$ in a single-edge set $S \subseteq V$. Given that we are coloring with v at this moment, the vertex is still a live vertex. However, since v is in \mathcal{B} , it will become dead at some moment of the algorithm. Since all the vertices in S will remain in the same component until the end of the algorithm, the moment v becomes dead, all vertices in S will also become dead. That means, every vertex in S is either in \mathcal{B} or \mathcal{D} , i.e. $S \subseteq \mathcal{B} \cup \mathcal{D} = K$.

Since S is a single-edge set, there is only one edge from T_{OPT} that cuts this set. Let assume that this edge is e , i.e. $\delta(S) \cap T_{\text{OPT}} = \{e\}$. Removing edge e from T_{OPT} , will only disconnect vertices in S from *root*, since S is a single-edge set and paths in T_{OPT} from *root* to vertices outside of S will not pass through e .

If we remove all such edges from T_{OPT} , the total cost of the removed edges will be at least b_1 . This is due to the fact that the coloring on these edges from single-cut sets assigned to the vertices in \mathcal{B} is equal to b_1 , and the coloring on each edge is at most its weight. Note that, each single-edge set is coloring exactly one edge of the optimal solution at each moment. So, we can remove edges with a total length of at least b_1 from T_{OPT} without disconnecting vertices in \mathcal{A} from *root*. □

Lemma 21. For an instance I , we can bound the cost of the optimal solution for instance R by

$$\text{cost}_{\text{OPT}_R} \leq \text{cost}_{\text{OPT}} - \beta r_{\mathcal{D}} - b_1,$$

where R is created at Line 12 of IPCST(I).

PROOF. To prove this lemma, we start by showing that there is a solution for instance R that costs at most $\text{cost}_{\text{OPT}} - \beta r_{\mathcal{D}} - b_1$. Since OPT_R is the optimal solution of instance R , its cost would not exceed the cost of the instance we are constructing. This will complete the proof of the lemma. To construct the mentioned instance, we

take the optimal solution of instance I , which we indicate by OPT , and remove extra edges from its tree T_{OPT} . Additionally, we do not need to pay penalties for pairs in OPT whose penalty is set to zero in π' at Line 11 for instance R .

Let's start with the tree T_{OPT} . Using Lemma 20, we can remove a set of edges from T_{OPT} with a total length of at least b_1 without disconnecting vertices in set \mathcal{A} from *root*.

Moreover, the optimal solution pays penalties for vertices in set $C \cup \mathcal{D}$. However, instance R has been constructed by assigning zero to the penalty of vertices in set K , which includes vertices in set \mathcal{D} . Therefore, the penalty that we pay for vertices in \mathcal{D} in the optimal solution is not required to be paid in OPT_R . This deducts $\pi(\mathcal{D})$ from the cost of the optimal solution, which is equal to $\beta r_{\mathcal{D}}$ according to Lemma 10. This completes the proof of this lemma. □

Lemma 22. For instance I , the output of the iterative solution can be bounded as follows:

$$\text{cost}_{\text{IT}} \leq \alpha \cdot \text{cost}_{\text{OPT}} + (\beta - \alpha)b_1 + \beta b_2 + (\beta - \alpha)r_{\mathcal{D}}$$

assuming that the recursive call on instance R returns an α approximate solution.

PROOF. We utilize Lemma 21 to modify the terms of the bound in Lemma 19.

$$\begin{aligned} \text{cost}_{\text{IT}} &\leq \alpha \text{cost}_{\text{OPT}_R} + \beta r_{\mathcal{B}} + \beta r_{\mathcal{D}} \quad (\text{Lemma 19}) \\ &\leq \alpha(\text{cost}_{\text{OPT}} - \beta r_{\mathcal{D}} - b_1) + \beta(b_1 + b_2) + \beta r_{\mathcal{D}} \quad (\text{Lemma 21 and Definition 9}) \\ &\leq \alpha \cdot \text{cost}_{\text{OPT}} + (\beta - \alpha)b_1 + \beta b_2 + (\beta - \alpha)r_{\mathcal{D}} \end{aligned}$$

□

3.2 Finding The Approximation Factor

Now that we have bounded cost_{GW} , cost_{ST} , and cost_{IT} , we can determine an appropriate value for α such that, during each call of IPCST on instance I , the minimum of cost_{GW} , cost_{ST} , and cost_{IT} is at most $\alpha \cdot \text{cost}_{\text{OPT}}$.

To achieve this, we assign weights to each solution in a way that the weighted average of these three bounds is at most $\alpha \cdot \text{cost}_{\text{OPT}}$. This completes our proof and demonstrates that the minimum among them is at most $\alpha \cdot \text{cost}_{\text{OPT}}$ since any weighted average of a set of values is greater than or equal to their minimum.

Denoting w_{GW} , w_{ST} , and w_{IT} as the weights of solutions GW, ST, and IT respectively, let cost_{WAG} represent their weighted average cost. As we are taking an average, we assume $w_{\text{GW}} + w_{\text{ST}} + w_{\text{IT}} = 1$ to simplify the calculation. We also have $w_{\text{GW}}, w_{\text{ST}}, w_{\text{IT}} \geq 0$. The bound for the weighted average is then given by

$$\begin{aligned} \text{cost}_{\text{WAG}} &\leq (\alpha \cdot w_{\text{GW}} + \alpha \cdot w_{\text{ST}} + \alpha \cdot w_{\text{IT}}) \cdot \text{cost}_{\text{OPT}} \\ &\quad + ((2 - \alpha) \cdot w_{\text{GW}} + (p - \alpha) \cdot w_{\text{ST}}) \cdot r_{\mathcal{A}} \\ &\quad + ((2 - \alpha) \cdot w_{\text{GW}} + (p + \beta - \alpha) \cdot w_{\text{ST}} + (\beta - \alpha) \cdot w_{\text{IT}}) \cdot b_1 \\ &\quad + ((2 - 2\alpha) \cdot w_{\text{GW}} + (2p + \beta - 2\alpha) \cdot w_{\text{ST}} + \beta \cdot w_{\text{IT}}) \cdot b_2 \\ &\quad + ((2 - \alpha\beta) \cdot w_{\text{GW}} + (2p - \alpha\beta) \cdot w_{\text{ST}}) \cdot r_{\mathcal{C}} \\ &\quad + ((2 - \alpha\beta) \cdot w_{\text{GW}} + (2p + \beta - \alpha\beta) \cdot w_{\text{ST}} + (\beta - \alpha\beta) \cdot w_{\text{IT}}) \cdot r_{\mathcal{D}} \end{aligned}$$

Given that $w_{\text{GW}} + w_{\text{ST}} + w_{\text{IT}} = 1$, we have $(\alpha \cdot w_{\text{GW}} + \alpha \cdot w_{\text{ST}} + \alpha \cdot w_{\text{IT}}) \cdot \text{cost}_{\text{OPT}} = \alpha \cdot \text{cost}_{\text{OPT}}$. Thus, the first term in the expression is $\alpha \cdot \text{cost}_{\text{OPT}}$.

To ensure $\text{cost}_{\text{WAG}} \leq \alpha \cdot \text{cost}_{\text{OPT}}$, we aim to make the rest of the expression non-positive. Since $r_{\mathcal{A}}$, b_1 , b_2 , r_C , and $r_{\mathcal{D}}$ are non-negative values, it suffices to make their coefficients non-positive by assigning suitable values to α , β , and the weights w_{GW} , w_{ST} , and w_{IT} . This leads to finding values that satisfy the following inequalities, with each inequality corresponding to one of the coefficients.

$$(2 - \alpha) \cdot w_{\text{GW}} + (p - \alpha) \cdot w_{\text{ST}} \leq 0 \quad (r_{\mathcal{A}})$$

$$(2 - \alpha) \cdot w_{\text{GW}} + (p + \beta - \alpha) \cdot w_{\text{ST}} + (\beta - \alpha) \cdot w_{\text{IT}} \leq 0 \quad (b_1)$$

$$(2 - 2\alpha) \cdot w_{\text{GW}} + (2p + \beta - 2\alpha) \cdot w_{\text{ST}} + \beta \cdot w_{\text{IT}} \leq 0 \quad (b_2)$$

$$(2 - \alpha\beta) \cdot w_{\text{GW}} + (2p - \alpha\beta) \cdot w_{\text{ST}} \leq 0 \quad (r_C)$$

$$(2 - \alpha\beta) \cdot w_{\text{GW}} + (2p + \beta - \alpha\beta) \cdot w_{\text{ST}} + (\beta - \alpha\beta) \cdot w_{\text{IT}} \leq 0 \quad (r_{\mathcal{D}})$$

We can also use a weighted average to ensure that our solution in the induction base has $\text{cost} \leq \alpha \cdot \text{cost}_{\text{OPT}}$. In this case, the IT solution cannot be employed as it represents the final step of recursion. So, we must have $w_{\text{IT}} = 0$. Additionally, it's essential to note that in this step, $\pi(K) = \pi(\mathcal{B} \cup \mathcal{D}) = 0$, resulting in $b_1 = b_2 = r_{\mathcal{D}} = 0$. Thus, only the inequalities for the coefficients of $r_{\mathcal{A}}$ and r_C remain relevant, which already do not contain w_{IT} :

$$(2 - \alpha) \cdot w_{\text{GW}} + (p - \alpha) \cdot w_{\text{ST}} \leq 0 \quad (r_{\mathcal{A}})$$

$$(2 - \alpha\beta) \cdot w_{\text{GW}} + (2p - \alpha\beta) \cdot w_{\text{ST}} \leq 0 \quad (r_C)$$

We can see that if a solution for the system of constraints used for the induction step is found, setting $w_{\text{IT}} = 0$ and scaling w_{GW} and w_{ST} by a factor of $\frac{1}{1 - w_{\text{IT}}}$ gives us a solution for these two new constraints with $w_{\text{GW}} + w_{\text{ST}} = 1$ and $w_{\text{IT}} = 0$. So, whatever values of α and β we find by solving the initial system of inequalities will give us a valid solution and an approximation guarantee of α .

Considering the best-known approximation factor for the Steiner tree problem, which is $p = \ln(4) + \epsilon$ [12], we determine that choosing the values $\alpha = 1.7994$, $\beta = 1.252$, $w_{\text{GW}} = 0.385$, $w_{\text{ST}} = 0.187$, and $w_{\text{IT}} = 0.428$ satisfies all the inequalities for a small enough value of ϵ . This provides a valid proof for both the induction base and induction step, leading to the conclusion of the following theorem.

Theorem 23. The minimum cost among GW, ST, and IT is a 1.7994-approximate solution for the Prize-Collecting Steiner Tree instance I . Therefore, IPCST is an 1.7994 approximation for PCST.

Finally, we note that our algorithm runs in polynomial time.

Theorem 24. The procedure $\text{IPCST}(I)$ runs in polynomial time.

PROOF. The procedure $\text{IPCST}(I)$ calls the $\text{PCSTGW}(I_\beta)$ which runs in polynomial time and a polynomial time algorithm STEINERTREE for the Steiner tree problem. Then it recursively calls itself on a new instance such that the new instance has more vertices with a penalty of 0. The construction of this instance involves a simple loop on the vertices and is done in polynomial time. Since the number of vertices is $|V|$, and each time the number of vertices with non-zero penalty decreases by one, the recursion depth is at most $|V|$. So, in IPCST we have a polynomial number of recursive steps, and each step takes a polynomial amount of time. Therefore, the total running time of the algorithm is polynomial in the size of the input. \square

4 NECESSITIES IN OUR ALGORITHM

In this section, we demonstrate the necessity of utilizing all three solutions in IPCST and selecting the minimum among them. Table 2 is completed based on the constraints $1 < p < \alpha < 1.8$, derived from the NP-hardness of finding an exact algorithm for Steiner tree, the fact that Steiner tree is a special case of PCST, and the goal of achieving an approximation factor better than 1.8. Additionally, we select β such that $2/\alpha \leq \beta \leq \alpha$ because if $2/\alpha > \beta$, both coefficients in r_C will be positive. Also, if $\beta > \alpha$, all coefficients of b_1 become positive.

Table 2: Sign of coefficients for each solution. The sign "?" indicates that its sign cannot be determined with the current assumptions, and the sign "+" means it is positive for the current value of p but could potentially turn negative with future improvements in the approximation factor of the Steiner tree problem.

	$r_{\mathcal{A}}$	b_1	b_2	r_C	$r_{\mathcal{D}}$
GW	+	+	-	-	-
ST	-	+	+	?	+
IT	0	-	+	0	-

Table 2 demonstrates the sign of the coefficient for each variable in every algorithm. We refer to this table to explain why all three algorithms are essential. We need to find a combination of these algorithms such that the weighted average of these coefficients adds up to zero. Since each row associated with an algorithm has at least one positive value, achieving this balance is not possible if we use only one of the algorithms. Moreover, omitting IT results in a positive coefficient for b_1 , making the iterative approach necessary. Similarly, using GW and IT together leads to a positive coefficient for $r_{\mathcal{A}}$, emphasizing the need for ST to offset it. Lastly, if we drop GW, the coefficient of b_2 constrains our approximation factor, as its coefficient in the IT algorithm is positive, and in the ST algorithm, it is $2p + \beta - 2\alpha$. Given that the best-known approximation factor for the Steiner tree is $\ln(4) + \epsilon$ [12], replacing p with $\ln(4) + \epsilon$ results in a positive value for the coefficient of b_2 in the ST algorithm. Therefore, the GW algorithm is necessary to decrease the coefficient of b_2 .

4.1 Bad example for $\beta > 2$

Let $\beta = 2(1 + \epsilon)$ for some $\epsilon > 0$. We consider a star graph G with $n + 1$ vertices as shown in Figure 3, where one vertex is a central vertex and all other vertices are connected to this vertex with edges of length 1 for some value of n such that $\frac{1}{n-1} < \epsilon$. We construct an instance of PCST on this graph where one of the non-central vertices is the root, the central vertex has penalty 0, and any other vertex has penalty $2(1 + \frac{1}{n-1})$.

When we run the GW algorithm on this instance, the center vertex dies instantly as it has 0 coloring potential. Additionally, as $\frac{1}{n-1} < \epsilon$, each non-root leaf has coloring potential

$$\frac{2(1 + \frac{1}{n-1})}{\beta} = \frac{(1 + \frac{1}{n-1})}{1 + \epsilon} < 1$$

and therefore dies before reaching the central vertex. So, the GW solution will pay penalty $(n - 1)(2(1 + \frac{1}{n-1})) = 2n$. This is twice




Figure 3: A star graph with $n+1$ vertices. We construct a PCST instance on this graph with vertex r as the root, the central vertex c having penalty 0, and all other vertices with having penalty $2(1 + \frac{1}{n-1})$.

the cost of the optimal solution, which can be obtained by taking all n edges of length 1. The other solutions we consider will also have the same cost as the GW solution, as they will aim to connect only the *root* and will pay the penalties for all the dead vertices. So, using any $\beta > 2$ will lead to an approximation factor of at least 2.

ACKNOWLEDGMENTS

The work is partially supported by DARPA QuICC, ONR MURI'2024 award on Algorithms, Learning, and Game Theory, Army-Research Laboratory (ARL): #W911NF2410052, NSF AF:Small #2218678, and NSF AF:Small #2114269.

REFERENCES

- [1] Ajit Agrawal, Philip Klein, and R. Ravi. 1991. When Trees Collide: An Approximation Algorithm for the Generalized Steiner Problem on Networks. In *Proceedings of the Twenty-Third Annual ACM Symposium on Theory of Computing* (New Orleans, Louisiana, USA) (STOC '91). Association for Computing Machinery, New York, NY, USA, 134–144. <https://doi.org/10.1145/103418.103437>
- [2] Ajit Agrawal, Philip N. Klein, and R. Ravi. 1995. When Trees Collide: An Approximation Algorithm for the Generalized Steiner Problem on Networks. *SIAM J. Comput.* 24, 3 (1995), 440–456. <https://doi.org/10.1137/S0097539792236237>
- [3] Ali Ahmadi, Iman Gholami, MohammadTaghi Hajiaghayi, Peyman Jabbarzade, and Mohammad Mahdavi. 2024. 2-Approximation for Prize-Collecting Steiner Forest. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM, 669–693. <https://doi.org/10.1137/1.9781611977912.25>
- [4] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. 1993. *Network flows - theory, algorithms and applications*. Prentice Hall.
- [5] Aaron Archer, MohammadHossein Bateni, MohammadTaghi Hajiaghayi, and Howard J. Karloff. 2011. Improved Approximation Algorithms for Prize-Collecting Steiner Tree and TSP. *SIAM J. Comput.* 40, 2 (2011), 309–332. <https://doi.org/10.1137/090771429>
- [6] Aaron Archer, MohammadHossein Bateni, Mohammad Taghi Hajiaghayi, and Howard J. Karloff. 2009. Improved Approximation Algorithms for PRIZE-COLLECTING STEINER TREE and TSP. In *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*. IEEE Computer Society, 427–436. <https://doi.org/10.1109/FOCS.2009.39>
- [7] Egon Balas. 1989. The prize collecting traveling salesman problem. *Networks* 19, 6 (1989), 621–636. <https://doi.org/10.1002/net.3230190602>
- [8] Marshall W. Bern and Paul E. Plassmann. 1989. The Steiner Problem with Edge Lengths 1 and 2. *Inf. Process. Lett.* 32, 4 (1989), 171–176. [https://doi.org/10.1016/0020-0190\(89\)90039-2](https://doi.org/10.1016/0020-0190(89)90039-2)
- [9] Daniel Bienstock, Michel X. Goemans, David Simchi-Levi, and David P. Williamson. 1993. A note on the prize collecting traveling salesman problem. *Math. Program.* 59 (1993), 413–420. <https://doi.org/10.1007/BF01581256>
- [10] Jannis Blauth, Nathan Klein, and Martin Nägele. 2023. A Better-Than-1.6-Approximation for Prize-Collecting TSP. *CoRR* abs/2308.06254 (2023). <https://doi.org/10.48550/ARXIV.2308.06254> arXiv:2308.06254
- [11] Jannis Blauth and Martin Nägele. 2023. An Improved Approximation Guarantee for Prize-Collecting TSP. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, Barna Saha and Rocco A. Servedio (Eds.). ACM, 1848–1861. <https://doi.org/10.1145/3564246.3585159>
- [12] Jaroslav Byrka, Fabrizio Grandoni, Thomas Rothvoß, and Laura Sanità. 2010. An improved LP-based approximation for steiner tree. In *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, Leonard J. Schulman (Ed.). ACM, 583–592. <https://doi.org/10.1145/1806689.1806769>
- [13] Miroslav Chlebík and Janka Chlebíková. 2008. The Steiner tree problem on graphs: Inapproximability results. *Theor. Comput. Sci.* 406, 3 (2008), 207–214. <https://doi.org/10.1016/j.tcs.2008.06.046>
- [14] Michel X. Goemans. 2009. Combining Approximation Algorithms for the Prize-Collecting TSP. *CoRR* abs/0910.0553 (2009). arXiv:0910.0553 <http://arxiv.org/abs/0910.0553>
- [15] Michel X. Goemans and David P. Williamson. 1995. A General Approximation Technique for Constrained Forest Problems. *SIAM J. Comput.* 24, 2 (1995), 296–317. <https://doi.org/10.1137/S0097539793242618>
- [16] Anupam Gupta, Jochen Könemann, Stefano Leonardi, R. Ravi, and Guido Schäfer. 2007. An efficient cost-sharing mechanism for the prize-collecting Steiner forest problem. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7-9, 2007*, Nikhil Bansal, Kirk Pruhs, and Clifford Stein (Eds.). SIAM, 1153–1162. <http://dl.acm.org/citation.cfm?id=1283383.1283507>
- [17] MohammadTaghi Hajiaghayi and Arefeh A. Nasri. 2010. Prize-Collecting Steiner Networks via Iterative Rounding. In *LATIN 2010: Theoretical Informatics, 9th Latin American Symposium, Oaxaca, Mexico, April 19-23, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6034)*, Alejandro López-Ortiz (Ed.). Springer, 515–526. https://doi.org/10.1007/978-3-642-12200-2_45
- [18] Mohammad Taghi Hajiaghayi and Kamal Jain. 2006. The prize-collecting generalized steiner tree problem via a new approach of primal-dual schema. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*. ACM Press, 631–640. <http://dl.acm.org/citation.cfm?id=1109557.1109626>
- [19] S. Louis Hakimi. 1971. Steiner's problem in graphs and its implications. *Networks* 1, 2 (1971), 113–133. <https://doi.org/10.1002/NET.3230010203>
- [20] Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations, held March 20-22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, USA (The IBM Research Symposia Series)*, Raymond E. Miller and James W. Thatcher (Eds.). Plenum Press, New York, 85–103. https://doi.org/10.1007/978-1-4684-2001-2_9
- [21] Marek Karpinski and Alexander Zelikovskiy. 1995. New Approximation Algorithms for the Steiner Tree Problems. *Electron. Colloquium Comput. Complex.* TR95-030 (1995). ECCC:TR95-030 <https://eccc.weizmann.ac.il/eccc-reports/1995/TR95-030/index.html>
- [22] Lawrence T. Kou, George Markowsky, and Leonard Berman. 1981. A Fast Algorithm for Steiner Trees. *Acta Informatica* 15 (1981), 141–145. <https://doi.org/10.1007/BF00288961>
- [23] Gabriel Robins and Alexander Zelikovskiy. 2005. Tighter Bounds for Graph Steiner Tree Approximation. *SIAM J. Discret. Math.* 19, 1 (2005), 122–134. <https://doi.org/10.1137/S0895480101393155>
- [24] F. James Rohlf. 2005. J. Felsenstein, *Inferring Phylogenies*, Sinauer Assoc., 2004, pp. xx + 664. *J. Classif.* 22, 1 (2005), 139–142. <https://doi.org/10.1007/S00357-005-0009-4>
- [25] Alexander Zelikovskiy. 1993. An 11/6-Approximation Algorithm for the Network Steiner Problem. *Algorithmica* 9, 5 (1993), 463–470. <https://doi.org/10.1007/BF01187035>

Received 13 November 2023; accepted 8 February 2024