

On the Integrality Gap of the Prize-Collecting Steiner Forest LP*

Jochen Könemann¹, Neil Olver², Kanstantsin Pashkovich³,
R. Ravi⁴, Chaitanya Swamy⁵, and Jens Vygen⁶

- 1 Department of Combinatorics and Optimization, University of Waterloo,
Waterloo, ON, Canada
jochen@uwaterloo.ca
- 2 Department of Econometrics and Operations Research, Vrije Universiteit
Amsterdam, Amsterdam, The Netherlands; and
CWI, Amsterdam, The Netherlands
n.olver@vu.nl
- 3 Department of Combinatorics and Optimization, University of Waterloo,
Waterloo, ON, Canada
kpashkovich@uwaterloo.ca
- 4 Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA, USA
ravi@andrew.cmu.edu
- 5 Department of Combinatorics and Optimization, University of Waterloo,
Waterloo, ON, Canada
jcswamy@uwaterloo.ca
- 6 Research Institute for Discrete Mathematics, Universität Bonn, Bonn,
Germany
vygen@or.uni-bonn.de

Abstract

In the *prize-collecting Steiner forest (PCSF)* problem, we are given an undirected graph $G = (V, E)$, edge costs $\{c_e \geq 0\}_{e \in E}$, terminal pairs $\{(s_i, t_i)\}_{i=1}^k$, and penalties $\{\pi_i\}_{i=1}^k$ for each terminal pair; the goal is to find a forest F to minimize $c(F) + \sum_{i:(s_i, t_i) \text{ not connected in } F} \pi_i$. The *Steiner forest* problem can be viewed as the special case where $\pi_i = \infty$ for all i . It was widely believed that the integrality gap of the natural (and well-studied) linear-programming (LP) relaxation for **PCSF** (PCSF-LP) is at most 2. We dispel this belief by showing that the integrality gap of this LP is at least $9/4$. This holds even for planar graphs. We also show that using this LP, one cannot devise a Lagrangian-multiplier-preserving (LMP) algorithm with approximation guarantee better than 4. Our results thus show a separation between the integrality gaps of the LP-relaxations for prize-collecting and non-prize-collecting (i.e., standard) Steiner forest, as well as the approximation ratios achievable relative to the optimal LP solution by LMP- and non-LMP- approximation algorithms for **PCSF**. For the special case of *prize-collecting Steiner tree (PCST)*, we prove that the natural LP relaxation admits basic feasible solutions with all coordinates of value at most $1/3$ and all edge variables positive. Thus, we rule out the possibility of approximating **PCST** with guarantee better than 3 using a direct iterative rounding method.

1998 ACM Subject Classification G.1.6 [Optimization] Integer Programming, Linear Programming

Keywords and phrases Integrality gap, Steiner tree, Steiner forest, prize-collecting, Lagrangian-multiplier-preserving

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2017.17

* NO was partially supported by an NWO Veni grant. RR was supported in part by the U.S. National Science Foundation under award number CCF-1527032. CS was supported in part by NSERC grant 327620-09 and an NSERC Discovery Accelerator Supplement Award.



© Jochen Könemann, Neil Olver, Kanstantsin Pashkovich, R. Ravi, Chaitanya Swamy, and Jens Vygen;
licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2017).

Editors: Klaus Jansen, José D. P. Rolim, David Williamson, and Santosh S. Vempala; Article No. 17; pp. 17:1–17:13



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction and Background

In an instance of the well-studied Steiner tree problem one is given an undirected graph $G = (V, E)$, a non-negative cost c_e for each edge $e \in E$, and a set of *terminals* $R \subseteq V$. The goal is to find a minimum-cost tree in G spanning R . In the more general *Steiner forest* problem, terminals are replaced by *terminal pairs* $(s_1, t_1), \dots, (s_k, t_k)$ and the goal now becomes to compute a minimum-cost forest that connects s_i to t_i for all i . Both of the above problems are well-known to be NP- and APX-hard [7, 17]. The best-known approximation algorithm for the Steiner tree problem is due to Byrka et al. [5] (see also [11]) and achieves an approximation ratio of $\ln 4 + \epsilon$, for any $\epsilon > 0$; the Steiner forest problem admits a $(2 - 1/k)$ -approximation algorithm [1, 12].

Our work focuses on the *prize-collecting* versions of the above problems. In the prize-collecting Steiner tree problem (**PCST**) we are given a Steiner-tree instance and a non-negative penalty π_v for each terminal $v \in R$. The goal is to find a tree T that minimizes $c(T) + \pi(T)$, where $c(T)$ denotes the total cost of all edges in T , and $\pi(T)$ denotes the total penalty of all terminals not spanned by T . In the prize-collecting Steiner forest problem (**PCSF**), we are given a Steiner-forest instance and a non-negative penalty π_i for each terminal pair (s_i, t_i) , and the goal is to find forest F that minimizes $c(F) + \pi(F)$ where, similar to before, $c(F)$ is the total cost of forest F , and $\pi(F)$ denotes the total penalty of terminal pairs that are not connected by F . We can view **PCST** as a special case of **PCSF** by guessing a node r in the optimal tree, and then modeling each vertex in $v \in R \setminus \{r\}$ by the terminal pair (v, r) .

The natural integer program (IP) for **PCSF** (see e.g. [3]) uses a binary variable x_e for every edge $e \in E$ whose value is 1 if e is part of the forest corresponding to x . The IP also has a variable z_i for each pair (s_i, t_i) whose value is 1 if s_i and t_i are *not* connected by the forest corresponding to x . We use $i \odot S$ for the predicate that is *true* if $S \subseteq V$ contains exactly one of s_i and t_i , and *false* otherwise. We use $\delta(S)$ to denote the set of edges with exactly one endpoint in S . In any integer solution to the LP relaxation below, the constraints insist that every cut separating pair (s_i, t_i) must be crossed by the forest unless we set z_i to 1 and pay the penalty for not connecting the terminals.

$$\begin{aligned} \min \quad & c^\top x + \pi^\top z && \text{(PCSF-LP)} \\ \text{s.t.} \quad & x(\delta(S)) + z_i \geq 1 \quad \forall S \subseteq V, i \odot S \\ & x, z \geq 0. \end{aligned}$$

Bienstock et al. [3] first presented a 3-approximation for **PCST** via a natural *threshold rounding* technique applied to this LP relaxation. This idea also works for **PCSF**, and proceeds as follows. First, we compute a solution (x, z) to the above LP. Let R' be the set of terminal pairs (s_i, t_i) with $z_i < 1/3$. Note that $\frac{3}{2} \cdot x$ is a feasible solution for the standard Steiner-forest cut-based LP (obtained from (PCSF-LP) by deleting the z variables) on the instance restricted to R' . Thus, applying an LP-based 2-approximation for Steiner forest [1, 12] to terminal pairs R' yields a forest F' of cost at most $2 \cdot \frac{3}{2} c^\top x = 3c^\top x$. The total penalty of the disconnected pairs is at most $3 \cdot \pi^\top z$. Hence, $c(F') + \pi(F')$ is bounded by $3(c^\top x + \pi^\top z)$, and the algorithm is a 3-approximation. Goemans showed that by choosing a random threshold (instead of the value $1/3$) from a suitable distribution, one can obtain an improved performance guarantee of $1/(1 - e^{-1/2}) \approx 2.5415$ (see page 136 of [20], which attributes the corresponding randomized algorithm for PCST in Section 5.7 of [20] to Goemans).

Goemans and Williamson [12] later presented a primal-dual 2-approximation for **PCST** based on the Steiner tree special case (PCST-LP) of (PCSF-LP). In fact, the algorithm gives even a slightly better guarantee; it produces a tree T such that

$$c(T) + 2\pi(T) \leq 2 \cdot \text{opt}_{\text{PCST-LP}},$$

where $\text{opt}_{\text{PCST-LP}}$ is the optimum value of (PCST-LP). Algorithms for prize-collecting problems that achieve a performance guarantee of the form

$$c(F) + \beta \cdot \pi(F) \leq \beta \cdot \text{opt}$$

are called *β -Lagrangian-multiplier preserving* (β -LMP) algorithms. Such algorithms are useful, for instance, for obtaining approximation algorithms for the *partial* covering version of the problem, which in the case of Steiner tree and Steiner forest translates to connecting at least a desired number of terminals (e.g., see [4, 16, 8, 9, 18]). Archer et al. [2] later used the strengthened guarantee of Goemans and Williamson's LMP algorithm for **PCST** to obtain a 1.9672-approximation algorithm for the problem.

The best known approximation guarantee for **PCSF** is 2.5415 obtained, as noted above, via Goemans' random-threshold idea applied to the threshold-rounding algorithm of Bienstock et al. This also shows that the integrality gap of (PCSF-LP) is at most 2.5415. The only known lower bound prior to this work was 2.

Our contributions

We demonstrate some limitations of (PCSF-LP) for designing approximation algorithms for **PCSF** and its special case, **PCST**, and in doing so dispel some widely-held beliefs about (PCSF-LP) and its specialization to **PCST**.

The integrality gap of (PCSF-LP) has been widely believed to be 2 since the work of Hajiaghayi and Jain [13], who devised a primal-dual 3-approximation algorithm for **PCSF** and pose the design of a primal-dual 2-approximation based on (PCSF-LP) as an open problem. However, as we show here, this belief is incorrect. Our main result is as follows.

► **Theorem 1.** *The integrality gap of (PCSF-LP) is at least $9/4$, even for planar instances of **PCSF**. Furthermore, any β -LMP approximation algorithm for the problem via (PCSF-LP) must have $\beta \geq 4$.*

When restricted to the non-prize-collecting Steiner forest problem, by setting $\pi_i = \infty$ for all i , (PCSF-LP) yields the standard LP for Steiner forest, which has an integrality gap of 2 [1]. Our result thus gives a clear separation between the integrality gaps of the prize-collecting and standard variants. It also shows a gap between the approximation ratios achievable relative to $\text{opt}_{\text{PCSF-LP}}$ by LMP and non-LMP approximation algorithms for **PCSF**. To the best of our knowledge, no such gaps were known previously for an LP for a natural network design problem. For example, for Steiner tree, there are no such gaps relative to the natural undirected LP obtained by specializing (PCSF-LP) to **PCST**. (There are however gaps in the current best approximation ratios known for Steiner tree and **PCST**, and approximation ratios achievable for **PCST** via LMP and non-LMP algorithms.)

In order to prove Theorem 1 we construct an instance on a large layered planar graph. Using a result of Carr and Vempala [6] it follows that (PCSF-LP) has a gap of α iff $\alpha \cdot (x, z)$ dominates a convex combination of integral solutions for any feasible solution (x, z) . We show that this can only hold if $\alpha \geq 9/4$.

In his groundbreaking paper [15] introducing the iterative rounding method, Jain showed that extreme points x of the Steiner forest LP (and certain generalizations) have an edge e

with $x_e = 0$ or $x_e \geq 1/2$. This then immediately yields a 2-approximation algorithm for the underlying problem, by iteratively deleting an edge of value zero or rounding up an edge of value at least half to one and proceeding on the residual instance. Again, it was long believed that a similar structural result holds for **PCST**: extreme points of (PCST-LP) have an edge variable of value 0, or a variable of value at least $1/2$. In fact, there were even stronger conjectures that envisioned the existence of a z -variable with value 1 in the case where all edge variables had positive value less than $1/2$. We refute these conjectures.

► **Theorem 2.** *There exists an instance of **PCST** where (PCST-LP) has an extreme point with all edge variables positive and all variables having value at most $1/3$.*

In [14] it was shown, that for every vertex (x, z) of (PCSF-LP) (and hence also (PCST-LP)) where x is positive, there is at least one variable of value at least $1/3$. Moreover for (PCSF-LP) this result is tight, i.e. there are instances of **PCSF** such that for some vertex (x, z) of (PCSF-LP), we have $x > 0$ and all coordinates are at most $1/3$. However, no such example was known for (PCST-LP). We provide such an example for **PCST**, showing that the $1/3$ upper bound on variable values is tight also for (PCST-LP).

2 The Integrality Gap for PCSF

2.1 Lower Bound on the Integrality Gap

We start proving Theorem 1 by describing the graph for our instance. Let P be a planar n -node 3-regular 3-edge-connected graph (for some large enough n to be determined later). Note that such graphs exist for arbitrarily large n ; e.g., the graphs of simple 3-dimensional polytopes (such as planar duals of triangulations of a sphere) have these properties; they are 3-connected by Steinitz's theorem [19].

We obtain H from P by subdividing every edge e of P , so that e is replaced by a corresponding path with n internal nodes. Let r denote an arbitrary degree-3 node in H , and call it the root. Define $H^{(0)} := H$ and obtain $H^{(i)}$ from $H^{(i-1)}$ by attaching a copy of H to each degree-2 node v in $H^{(i-1)}$, identifying the root node of the copy with v ; we call this the *copy of H with root v* . We also define the *parent* of any node $u \neq v$ in this copy to be v . In the end, we let $G := H^{(k)}$ for some large k , and we let r_0 be the node corresponding to the root of $H^{(0)}$. Figure 1 gives an example of this construction. Note that each copy of H can be thought of as a subgraph of G .


Next, let us define the source-sink pairs. We introduce a source-sink pair s, t whenever s and t are degree-3 nodes in the same copy of H . We also introduce a source-sink pair r_0, t whenever t is a degree-2 node in G .

Now let $x_e := 1/3$, for all $e \in E$, $z_{uv} := 0$ if u and v are degree-3 nodes in the same copy of H , and $z_{uv} := 1/3$ otherwise. (Here and henceforth, we abuse notation slightly and index z by the source-sink terminal pair that it corresponds to.) Clearly, (x, z) is a feasible solution for (PCSF-LP) by the 3-edge-connectivity of G .

Let α be the integrality gap of (PCSF-LP). By [6] there is a collection of forests F_1, \dots, F_q in G (the same forest could appear multiple times in the collection) such that picking a forest F uniformly at random from F_1, \dots, F_q satisfies

- (a) $\mathbb{P}[e \in F] \leq \frac{\alpha}{3}$ for all $e \in E$, and
- (b) Letting $u \sim_F v$ denote the event that u and v are connected in F , for all $u, v \in V(G)$, we have

$$\mathbb{P}[u \sim_F v] \geq (1 - \alpha z_{uv}) = \begin{cases} 1 & \text{if } u, v \text{ are degree-3 nodes in the same copy of } H, \\ 1 - \frac{\alpha}{3} & \text{if } u = r_0 \text{ and } v \text{ is a degree-2 node in } G. \end{cases}$$



■ **Figure 1** Taking $n = 4$, and hence P to be the complete graph on 4 vertices, the resulting graph $H^{(1)}$ is shown.

We begin by observing that we may assume that each forest F_1, \dots, F_q induces a tree when restricted to any of the copies of H in G . For consider any F_i , and a copy of H with root v ; call this H' . Every degree-3 node in H' is connected to v in F_i , by requirement (b). So consider any degree-2 node u in H' . If u is not connected to a degree-3 node of H' (and hence to v) in F_i , then any edges of F_i adjacent to u can be safely deleted without destroying any connectivity amongst the source-sink pairs of the instance.

The argument will show that if α is too small, not all degree-2 nodes can be connected to r_0 with high enough probability. More precisely, we will show a geometrically decreasing probability, in k . The intuition is roughly as follows. Consider a copy H' of H with root u , where $u \neq r_0$. Almost all of the degree-2 nodes of H' that are connected to u in F will have degree 2 in F , since $F[H']$ is a tree and H' is made up of long paths. This is rather wasteful, since both edges adjacent to a typical degree-2 node v are used to connect; as each edge appears with probability $\alpha/3$, v can only be part of F (and hence connected to u) with probability about $\alpha/3$. Moreover, we will show that even conditioned on the event that u is *not* connected to r_0 , there will be some choice of v such that v is connected to u in F with probability around $2/3$ (see (2) in Claim 3). This is again a waste in terms of connectivity to r_0 . If p_i denotes the worst connectivity probability amongst nodes in $H^{(i)}$ in the construction, we have

$$p_{i+1} \lesssim \frac{\alpha}{3} - \frac{2}{3}(1 - p_i). \tag{5} \text{ is a more precise version of this inequality}$$

If $\alpha < 9/4$, this decreases geometrically, providing a counterexample for n large enough.

For now, let us introduce an abstract event I (that the reader may think of as “an ancestor of node v is not connected to r_0 ” motivated by the above discussion).

► **Claim 3.** *Let a forest F be picked uniformly at random from F_1, \dots, F_q , let I be an event with $\mathbb{P}[I] > 0$ and let H' be a copy of H in G . Then there exists a degree-2 node v in H' such that*

$$\mathbb{P}[\deg_{F[H']}(v) = 1] \leq \frac{2}{n} \tag{1}$$

and

$$\mathbb{P}[Q_v \subseteq F \mid I] \geq \frac{2(n-1)}{3n}, \tag{2}$$

where Q_v is the path in H' corresponding to the edge of P containing v .

Proof. The event I corresponds to a nonempty multiset $\mathcal{F} \subseteq \{F_1, \dots, F_q\}$ of the forests. Each of $F_1[H']$, \dots , $F_q[H']$ is a tree, by our earlier assumption, and so each of them naturally induce a spanning tree of P . More precisely, for each $e \in E(P)$, let Q_e denote the corresponding path in H' ; then $\{e \in E(P) : Q_e \subseteq F_i[H']\}$ is a spanning tree for each i . Thus

$$\sum_{e \in E(P)} |\{F' \in \mathcal{F} : Q_e \subseteq F'\}| = \sum_{F' \in \mathcal{F}} |\{e \in E(P) : Q_e \subseteq F'\}| = \sum_{F' \in \mathcal{F}} (n-1) = |\mathcal{F}|(n-1).$$

So there is an edge $f \in E(P)$ for which

$$\mathbb{P}[Q_f \subseteq F | I] = \frac{|\{F' \in \mathcal{F} : Q_f \subseteq F'\}|}{|\mathcal{F}|} \geq \frac{(n-1)}{|E(P)|} = \frac{2(n-1)}{3n}.$$

At most two of the nodes on Q_f are leaves in any of $F_1[H']$, \dots , $F_q[H']$ (again since they are all trees). The total number of degree-2 nodes in H' lying on Q_f is n , so there exists a degree-2 node v in H' such that $v \in Q_f$ and $\mathbb{P}[\deg_{F[H']}(v) = 1] \leq \frac{2}{n}$. \blacktriangleleft

► Claim 4. *Let $\epsilon > 0$ be given. Then for n and k chosen sufficiently large, there exists a degree-2 node u in G such that*

$$\mathbb{P}[u \sim_F r_0] \leq \alpha - 2 + \epsilon,$$

where F is a uniformly random forest from F_1, \dots, F_q .

Proof. Consider the root copy $H^{(0)}$ of H , with root r_0 . Set $H_0 = H^{(0)}$. Pick a degree-2 node v in H_0 that satisfies (1) in Claim 3 for the trivial event $I := \{r_0 \sim_F r_0\}$ and $H' := H_0$. Let $r_1 := v$. Note that

$$\mathbb{P}[r_1 \sim_F r_0] = \mathbb{P}[\deg_{F[H_0]}(r_1) = 2] + \mathbb{P}[\deg_{F[H_0]}(r_1) = 1] \leq \frac{\alpha}{3} + \frac{2}{n} < 1.$$

The first inequality follows from (a) and (1), and the second since $\alpha \leq 2.5415$. Therefore $\mathbb{P}[r_1 \not\sim_F r_0] > 0$.

Suppose that we have defined $(H_0, r_1), (H_1, r_2), \dots, (H_{i-1}, r_i)$ for some i with $1 \leq i \leq k$, such that the following hold for all $1 \leq j \leq i$: (i) r_j is a degree-2 node in H_{j-1} , and r_{j-1} is the root of H_{j-1} ; (ii) $\mathbb{P}[r_{j-1} \not\sim_F r_0] > 0$ if $j \geq 2$; (iii) if $j \geq 2$, then (1) and (2) hold in Claim 3 for $H' = H_{j-1}$, $I = \{r_{j-1} \not\sim_F r_0\}$ and $v = r_j$. We now show how to define H_i and r_{i+1} such that the above properties continue to hold for $j = i + 1$.

First, set H_i to be the copy of H whose root is r_i . We have $\mathbb{P}[r_i \not\sim_F r_0] \geq \mathbb{P}[r_{i-1} \not\sim_F r_0] > 0$, so property (ii) continues to hold. Given this, pick a degree-2 node v in H_i that satisfies (1) and (2) in Claim 3 for the event $I := \{r_i \not\sim_F r_0\}$ and $H' := H_i$. Set $r_{i+1} := v$. Thus, properties (i) and (iii) continue to hold as well.

For $j \in \{0, \dots, k\}$, due to the choice of r_{j+1} and (1), we have $\mathbb{P}[\deg_{F[H_j]}(r_{j+1}) = 1] \leq \frac{2}{n}$ and thus

$$\mathbb{P}[r_{j+1} \sim_F r_j] = \mathbb{P}[\deg_{F[H_j]}(r_{j+1}) = 2] + \mathbb{P}[\deg_{F[H_j]}(r_{j+1}) = 1] \leq \frac{\alpha}{3} + \frac{2}{n}. \quad (3)$$

For $j \in \{1, \dots, k\}$, due to (2) and the choice of r_{j+1} , we get

$$\begin{aligned} \mathbb{P}[Q_{r_{j+1}} \subseteq F \wedge r_j \not\sim_F r_0] &= \mathbb{P}[Q_{r_{j+1}} \subseteq F | r_j \not\sim_F r_0] \cdot \mathbb{P}[r_j \not\sim_F r_0] \\ &\geq \frac{2(n-1)}{3n} \mathbb{P}[r_j \not\sim_F r_0] \\ &\geq \frac{2}{3} (1 - \mathbb{P}[r_j \sim_F r_0]) - \frac{2}{3n}. \end{aligned} \quad (4)$$

Hence, for $j \in \{0, \dots, k\}$,

$$\begin{aligned}
 \mathbb{P}[r_{j+1} \sim_F r_0] &= \mathbb{P}[r_{j+1} \sim_F r_j \wedge r_j \sim_F r_0] \\
 &= \mathbb{P}[r_{j+1} \sim_F r_j] - \mathbb{P}[r_{j+1} \sim_F r_j \wedge r_j \not\sim_F r_0] \\
 &\leq \mathbb{P}[r_{j+1} \sim_F r_j] - \mathbb{P}[Q_{r_{j+1}} \subseteq F \wedge r_j \not\sim_F r_0] \\
 &\leq \frac{\alpha}{3} + \frac{2}{n} - \frac{2}{3} + \frac{2}{3} \mathbb{P}[r_j \sim_F r_0] + \frac{2}{3n}, \tag{5}
 \end{aligned}$$

where the first inequality follows from the fact that $r_{j+1} \sim_F r_j$ holds whenever $Q_{r_{j+1}} \subseteq F$ holds, and the second inequality follows from (3) and (4).

Expanding the recursion, we get

$$\mathbb{P}[r_{k+1} \sim_F r_0] \leq \left(\frac{\alpha - 2}{3} + \frac{8}{3n} \right) \sum_{i=0}^k \left(\frac{2}{3} \right)^i + \left(\frac{2}{3} \right)^{k+1},$$

so for n and k large enough we obtain

$$\mathbb{P}[r_{k+1} \sim_F r_0] \leq \left(\frac{\alpha - 2}{3} + \frac{\epsilon}{6} \right) \sum_{i=0}^{\infty} \left(\frac{2}{3} \right)^i + \frac{\epsilon}{2} = \alpha - 2 + \epsilon.$$

Since $u := r_{k+1}$ is a degree-2 node in G , the proof is complete. ◀

Now, we can prove the first part of Theorem 1. By Claim 4 and property (b) of the collection of forests, we get the inequality

$$\alpha - 2 \geq 1 - \alpha/3,$$

leading to $\alpha \geq 9/4$.

2.2 The Integrality Gap is Tight for the Construction

We note that for any n and k , the **PCSF** instance given by our construction has integrality gap at most $9/4$. More generally, we show that the integrality gap over **PCSF** instances which admit a feasible solution (x, z) to (PCSF-LP) with $z_i \in \{0, 1/3\}$ for all i , is at most $9/4$. (That is, the maximum ratio between the optimal values of the IP and the LP for such instances is at most $9/4$.) This nicely complements our integrality-gap lower bound, and shows that our analysis above is tight (for such instances).

To show the first statement, we simply provide a distribution over forests F_1, \dots, F_q satisfying (a) and (b). (The next paragraph, which proves the second claim above, gives another proof.) Since $(2(n-1)/(3n)) \cdot \mathbf{1}$ is in the spanning tree polytope of P , there is a list of spanning trees such that every edge is contained in less than $2/3$ of them. Consider the following distribution of forests. With probability $3 - \alpha$ we pick one of these spanning trees of P uniformly at random and subdivide it to obtain a tree in H ; we take this tree in each copy of H to obtain a (non-spanning) tree in G . With probability $\alpha - 2$ we pick an arbitrary spanning tree of G . This random forest F satisfies

$$\mathbb{P}[e \in F] \leq (\alpha - 2) \cdot 1 + (3 - \alpha) \cdot \frac{2}{3} = \frac{\alpha}{3}.$$

Thus (a) holds for the above distribution. To see that (b) holds, note that for every degree-2 node v in G we have $\mathbb{P}[v \sim_F r_0] \geq \alpha - 2 = 1 - \alpha/3$.

For the second claim, we utilize threshold rounding to show that the integrality gap is at most $9/4$ for such instances. Consider an instance of **PCSF** and a feasible point (x, z) for (PCSF-LP) such that the values of z -variables are 0 or γ for some fixed γ with $0 < \gamma < 1/2$. Using [1, 12], we can obtain an integer solution of cost at most $2c^\top x + \pi^\top z/\gamma$ by paying the penalties for all pairs with a non-zero z value. We can also obtain a solution of cost at most $2c^\top x/(1-\gamma)$ by connecting all pairs. Therefore, for any $p \in [0, 1]$, we can obtain an integer solution of cost at most

$$p \left(2c^\top x + \frac{\pi^\top z}{\gamma} \right) + (1-p) \left(\frac{2c^\top x}{1-\gamma} \right) \leq \max \left\{ \frac{2-2p\gamma}{1-\gamma}, \frac{p}{\gamma} \right\} (c^\top x + \pi^\top z)$$

showing that the integrality gap is at most

$$\mu := \min_{0 \leq p \leq 1} \max \left\{ \frac{2-2p\gamma}{1-\gamma}, \frac{p}{\gamma} \right\}.$$

The number μ is at most $2/(2\gamma^2 - \gamma + 1)$, which is equal to $9/4$ for $\gamma = 1/3$. Note that for $\gamma = 1/4$ the $2/(2\gamma^2 - \gamma + 1)$ achieves its maximum value of $16/7$.

2.3 Lagrangian-Multiplier Preserving Approximation Algorithms for PCSF

Recall that a β -Lagrangian-multiplier-preserving (LMP) approximation algorithm for **PCSF** is an approximation algorithm that returns a forest F satisfying

$$c(F) + \beta \cdot \pi(F) \leq \beta \cdot \text{opt}.$$

We show that we must have $\beta \geq 4$ in order to obtain a β -LMP algorithm relative to the optimum of the LP-relaxation (PCSF-LP), that is, to obtain the guarantee $c(F) + \beta \cdot \pi(F) \leq \beta \cdot \text{opt}_{\text{PCSF-LP}}$. To obtain this lower bound, we modify our earlier construction slightly. We construct $G = H^{(k)}$ in a similar fashion as before, but we now choose P (the “base graph”) to be an n -node l -regular l -edge-connected graph. Let $x_e := 1/l$ for all $e \in E$, and let $z_{uv} := 0$ if u and v are degree- l nodes in the same copy of H , and $z_{uv} := 1 - 2/l$ otherwise.

By arguments similar to [6] (see, e.g., the proof of Theorem 7.2 in [10], and Theorem 8 in the Appendix), one can show that if there exists a β -LMP approximation algorithm for **PCSF** relative to (PCSF-LP) then there are forests F_1, \dots, F_q in G (the same forest could appear multiple times) such that picking a forest F uniformly at random from F_1, \dots, F_q satisfies

(a') $\mathbb{P}[e \in F] \leq \frac{\beta}{l}$ for all $e \in E$, and

(b') $\mathbb{P}[u \sim_F v] \geq (1 - z_{uv}) = \begin{cases} 1 & u, v \text{ are degree-}l \text{ nodes in the same copy of } H \\ \frac{2}{l} & \text{if } u = r_0 \text{ and } v \text{ is a degree-2 node in } G \end{cases}$

for all $u, v \in V(G)$.

It is straightforward to obtain the analogues of Claim 3 and Claim 4.


► **Claim 5.** *Let a forest F be picked uniformly at random from F_1, \dots, F_q , let I be an event with $\mathbb{P}[I] > 0$ and let H' be a copy of H in G . There exists a degree-2 node v in H' , such that*

$$\mathbb{P} \left[\deg_{F[H']}(v) = 1 \right] \leq \frac{2}{n} \tag{6}$$

and

$$\mathbb{P}[Q_v \subseteq F \mid I] \geq \frac{2(n-1)}{ln}, \tag{7}$$

where Q_v is the path in H' that contains v and corresponds to an edge of P .



■ **Figure 2** Here, each of the nodes v_1, \dots, v_k corresponds to the gadget in Figure 3. Additionally, a cut $\{r\}$ is marked as a tight constraint in (PCST-LP) for the constructed point (x, z) . $x_e = 1/k$ for all edges e .

► **Claim 6.** *Let $\epsilon > 0$ be given. Then for n and k sufficiently large, and choosing F uniformly at random from F_1, \dots, F_q , there exists a degree-2 node u in G such that*

$$\mathbb{P}[u \sim_F r_0] \leq \frac{\beta - 2}{l - 2} + \epsilon.$$

For the node u from Claim 6, we have $(\beta - 2)/(l - 2) \geq \mathbb{P}[v \sim_F r_0] \geq 2/l$. Thus, β is at least $4 - 4/l$, which approaches 4 as l increases. This completes the proof of the second part of Theorem 1.

Moreover, the analysis is tight for the above construction. For a solution to (PCSF-LP) where z takes on only two distinct values, say 0 and γ , threshold rounding shows that for $\beta = 2 + 2\gamma < 4$ the desired collection of forests exists. However, for an unbounded number of distinct values of z , no constant-factor upper bound is known.


3 An Extreme Point for PCST with All Values at most $\frac{1}{3}$

In this section we present a proof of Theorem 2. Take an integer $k \geq 4$ and consider the graph $G = (V, E)$ in Figure 2. Here, the nodes v_1, \dots, v_k represent the gadgets shown in Figure 3. The gadget consists of ten nodes, and there are precisely four edges incident to a node in the gadget. We let r to be the root node and introduce a source-sink node pair (v, r) for every node $v \in V \setminus \{r\}$.

In the case $k = 6$, the next claim proves Theorem 2.

► **Claim 7.** *The following is an extreme point of (PCST-LP) for this instance: $z_s = 0$ and $z_u = 1 - 4/k$ for every node u in $V \setminus \{r, s\}$. For the wavy edges in Figure 3, we have $x_{u_1 u_2} := x_{u_3 u_4} := x_{u_5 u_6} := x_{u_7 u_8} := x_{u_9 u_{10}} := 2/k$, and $x_e = 1/k$ for all the other edges e .*

Proof. It is straightforward to check that the defined point (x, z) is feasible. Let us show that the defined point (x, z) is a vertex of (PCST-LP). To show this, it is enough to provide a set of tight constraints in (PCST-LP) which uniquely define the above point (x, z) .



■ **Figure 3** A gadget used for the construction in Figure 2. Additionally, the cuts are marked as tight constraints in (PCST-LP) for the constructed point (x, z) . For an edge e , $x_e = 2/k$ if e is a wavy edge, and $x_e = 1/k$ if it is a straight edge.

Let us consider the gadget in Figure 3. For each such gadget, the set of tight inequalities from (PCST-LP) contains the following constraints:

$$x(\delta(u_i)) + z_{u_i} = 1 \quad \forall i \in \{1, \dots, 10\} \quad (8)$$

$$x(\delta(\{u_1, \dots, u_{10}\})) + z_{u_i} = 1 \quad \forall i \in \{1, \dots, 10\} \quad (9)$$

$$x(\delta(\{u_i, u_{i+1}\})) + z_{u_i} = 1 \quad \forall i \in \{1, 3, 5, 7, 9\} \quad (10)$$

$$x(\delta(\{u_1, \dots, u_4\})) + z_{u_1} = 1 \quad (11)$$

$$x(\delta(\{u_7, \dots, u_{10}\})) + z_{u_7} = 1. \quad (12)$$

There are two more tight constraints which we use in the proof:

$$x(\delta(r)) + z_s = 1 \quad (13)$$

$$z_s = 0. \quad (14)$$

Let us prove that the constraints (8)–(14) define the point (x, z) from Claim 7. First, let us consider a gadget in Figure 3. It is clear that (9) implies $z_{u_1} = \dots = z_{u_{10}}$. By (8) and (10), we get

$$2x_{u_1 u_2} = x(\delta(u_1)) + x(\delta(u_2)) - x(\delta(\{u_1, u_2\})) = (1 - z_{u_1}) + (1 - z_{u_1}) - (1 - z_{u_1}) = (1 - z_{u_1}),$$

and hence $x_{u_1 u_2} = (1 - z_{u_1})/2$. Similarly, we obtain $x_{u_1 u_2} = x_{u_3 u_4} = \dots = x_{u_9 u_{10}} = (1 - z_{u_1})/2$.

Now, we have

$$x_{u_3 u_5} + x_{u_2 u_3} = x(\delta(u_3)) - x_{u_3 u_4} = (1 - z_{u_1})/2$$

$$x_{u_2 u_3} + x_{u_2 u_5} = x(\delta(u_2)) - x_{u_1 u_2} = (1 - z_{u_1})/2$$

$$x_{u_2 u_5} + x_{u_3 u_5} = x(\delta(u_5)) - x_{u_5 u_6} = (1 - z_{u_1})/2,$$

implying $x_{u_2 u_3} = x_{u_2 u_5} = x_{u_3 u_5} = (1 - z_{u_1})/4$. Similarly, $x_{u_6 u_7} = x_{u_6 u_{10}} = x_{u_7 u_{10}} = (1 - z_{u_1})/4$.

By (10) and (11), we get

$$2x_{u_1 u_4} = x(\delta(\{u_1, u_2\})) + x(\delta(\{u_3, u_4\})) - x(\delta(\{u_1, \dots, u_4\})) - 2x_{u_2 u_3} = (1 - z_{u_1})/2,$$

showing $x_{u_1 u_4} = (1 - z_{u_1})/4$. Similarly, we get $x_{u_8 u_9} = (1 - z_{u_1})/4$. From here, it is straightforward to show that all straight edges in Figure 3 have value $(1 - z_{u_1})/4$ and all wavy edges have value $(1 - z_{u_1})/2$.

Consider the graph in Figure 2. Due to the edge v_1v_2 , the straight edges in the gadget associated to v_1 have the same x value as the straight edges in the gadget associated to v_2 . Thus, due to the cycle $v_1v_2 \dots v_k$ the straight edges in all gadgets have the same x value. To finish the proof use (13) and (14). ◀

Acknowledgement. We would like to thank Hausdorff Research Institute for Mathematics. This research was initiated during the Hausdorff Trimester Program “Combinatorial Optimization”.

References

- 1 A. Agrawal, P. Klein, and R. Ravi. When trees collide: an approximation algorithm for the generalized Steiner problem on networks. *SIAM Journal on Computing*, 24(3):440–456, 1995.
- 2 A. Archer, M. Bateni, M. Hajiaghayi, and H. Karloff. Improved approximation algorithms for prize-collecting Steiner tree and TSP. *SIAM Journal on Computing*, 40(2):309–332, 2011.
- 3 D. Bienstock, M. X. Goemans, D. Simchi-Levi, and D. Williamson. A note on the prize collecting traveling salesman problem. *Mathematical Programming*, 59(1):413–420, 1993.
- 4 A. Blum, R. Ravi, and S. Vempala. A constant-factor approximation algorithm for the k -MST problem. *Journal of Computer and System Sciences*, 58(1):101–108, 1999.
- 5 J. Byrka, F. Grandoni, T. Rothvoss, and L. Sanità. Steiner tree approximation via iterative randomized rounding. *Journal of the ACM*, 60(1):6, 2013.
- 6 R. D. Carr and S. Vempala. On the Held-Karp relaxation for the asymmetric and symmetric traveling salesman problems. *Mathematical Programming A*, 100:569–587, 2004.
- 7 M. Chlebík and J. Chlebíková. The Steiner tree problem on graphs: Inapproximability results. *Theoretical Computer Science*, 406(3):207–214, 2008. doi:10.1016/j.tcs.2008.06.046.
- 8 F. A. Chudak, T. Roughgarden, and D. P. Williamson. Approximate k -MSTs and k -Steiner trees via the primal-dual method and Lagrangean relaxation. *Mathematical Programming*, 100(2):411–421, 2004.
- 9 N. Garg. Saving an epsilon: a 2-approximation algorithm for the k -MST problem in graphs. In *Proceedings of the 37th ACM Symposium on Theory of Computing*, pages 396–402, 2005.
- 10 K. Georgiou and C. Swamy. Black-box reductions for cost-sharing mechanism design. *Games and Economic Behavior*, 2013. doi:10.1016/j.geb.2013.08.012.
- 11 M. X. Goemans, N. Olver, T. Rothvoß, and R. Zenklusen. Matroids and integrality gaps for hypergraphic Steiner tree relaxations. In *Proceedings of the 44th ACM Symposium on Theory of Computing*, pages 1161–1176, 2012.
- 12 M. X. Goemans and D. P. Williamson. A general approximation technique for constrained forest problems. *SIAM Journal on Computing*, 24(2):296–317, 1995.
- 13 M. Hajiaghayi and K. Jain. 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*, pages 631–640, 2006.
- 14 M. Hajiaghayi and A. Nasri. Prize-collecting Steiner networks via iterative rounding. In *Theoretical Informatics: LATIN 2010*, pages 515–526. Springer, 2010.
- 15 K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21(1):39–60, 2001. doi:10.1007/s004930170004.
- 16 K. Jain and V. V. Vazirani. Approximation algorithms for metric facility location and k -median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM*, 48(2):274–296, 2001.

- 17 R. M. Karp. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, pages 85–103. Springer, 1972.
- 18 J. Könemann, O. Parekh, and D. Segev. A unified approach to approximating partial covering problems. *Algorithmica*, 59(4):489–509, 2011.
- 19 E. Steinitz. Polyeder und Raumeinteilungen. In *Enzyklopädie der Mathematischen Wissenschaften, vol. 3, Geometrie, erster Teil, zweite Hälfte*, pages 1–139. Teubner, 1922.
- 20 D. P. Williamson and D. B. Shmoys. *The Design of Approximation Algorithms*. Cambridge University Press, 2010.

A Implications of an LMP Approximation Algorithm for PCSF

We adapt the arguments in [6] to show that a β -LMP approximation relative to (PCSF-LP) implies that any fractional solution (x, z) to (PCSF-LP) can be translated to a distribution over integral solutions to (PCSF-LP) satisfying certain properties; this implies the existence of the forests F_1, \dots, F_q in Section 2.3. The arguments below are known (see, e.g., the proof of Theorem 7.2 in [10]); we include them for completeness.

Let $G = (V, E)$, $\{c_e \geq 0\}_{e \in E}$, $\{(s_i, t_i, \pi_i)\}_{i=1}^k$ be a **PCSF**-instance. Let $\{(x^{(q)}, z^{(q)})\}_{q \in \mathcal{I}}$ be the set of all integral solutions to (PCSF-LP), where \mathcal{I} is simply an index set.

► **Theorem 8.** *Let \mathcal{A} be a β -LMP approximation algorithm for **PCSF** relative to (PCSF-LP). Given any fractional solution (x^*, z^*) to (PCSF-LP), one can obtain nonnegative multipliers $\{\lambda^{(q)}\}_{q \in \mathcal{I}}$ such that $\sum_q \lambda^{(q)} = 1$, $\sum_q \lambda^{(q)} x^{(q)} \leq \beta x^*$, and $\sum_q \lambda^{(q)} z^{(q)} \leq z^*$. Moreover, the $\lambda^{(q)}$ values are rational if (x^*, z^*) is rational.*

Proof. Consider the following pair of primal and dual LPs.

$$\begin{array}{ll}
 \max & \sum_q \lambda^{(q)} \quad (P) \\
 \text{s.t.} & \sum_q \lambda^{(q)} x_e^{(q)} \leq \beta x_e^* \quad \forall e \\
 & \sum_q \lambda^{(q)} z_i^{(q)} \leq z_i^* \quad \forall i \\
 & \sum_q \lambda^{(q)} \leq 1 \\
 & \lambda \geq 0.
 \end{array}
 \quad \left| \quad \begin{array}{ll}
 \min & \sum_e \beta x_e^* d_e + \sum_i z_i^* \rho_i + \gamma \quad (D) \\
 \text{s.t.} & \sum_e x_e^{(q)} d_e + \sum_i z_i^{(q)} \rho_i + \gamma \geq 1 \quad \forall q \\
 & d, \rho, \gamma \geq 0.
 \end{array}$$

It suffices to show that the optimal value of (P) is 1. The rationality of the $\lambda^{(q)}$ values when (x^*, z^*) is rational then follows from the fact that an LP with rational data has a rational optimal solution. (The proof below also yields a polynomial-time algorithm to solve (P) by showing that \mathcal{A} can be used to obtain a separation oracle for the dual.)

Note that both (P) and (D) are feasible, so they have a common optimal value. We show that $\text{opt}_D = 1$. Setting $\gamma = 1$, $d = \rho = 0$, we have that $\text{opt}_D \leq 1$. Suppose (d, ρ, γ) is feasible to (D) and $\sum_e \beta x_e^* d_e + \sum_i z_i^* \rho_i + \gamma < 1$. Consider the **PCSF** instance given by G , edge costs $\{d_e\}_{e \in E}$, and terminal pairs and penalties $\{(s_i, t_i, \rho_i/\beta)\}_{i=1}^k$. Running \mathcal{A} on this instance, we can obtain an integral solution $(x^{(q)}, z^{(q)})$ such that

$$\sum_e d_e x_e^{(q)} + \sum_i \rho_i z_i^{(q)} + \gamma \leq \beta \left(\sum_e d_e x_e^* + \sum_i z_i^* \rho_i / \beta \right) + \gamma < 1$$

which contradicts the feasibility of (d, ρ, γ) . Hence, $\text{opt}_D = 1$. ◀

Note that if (x^*, z^*) is rational, then since the $\lambda^{(q)}$ values are rational, we can multiply them by a suitably large number to convert them to integers; thus, we may view the distribution specified by the $\lambda^{(q)}$ values as the *uniform distribution* over a *multiset* of integral solutions to (PCSF-LP).

We remark that the converse of Theorem 8 also holds in the following sense. If for every fractional solution (x^*, z^*) to (PCSF-LP), we can obtain $\lambda^{(q)}$ values (or equivalently, a distribution over integral solutions to (PCSF-LP)) satisfying the properties in Theorem 8, then we can obtain a β -LMP approximation algorithm for **PCSF** relative to (PCSF-LP): this follows, by simply returning the integral solution $(x^{(q)}, z^{(q)})$ with $\lambda^{(q)} > 0$ that minimizes $\sum_e c_e x_e^{(q)} + \beta \sum_i \pi_i z_i^{(q)}$.