

Approximating Min-Cost Chain-Constrained Spanning Trees: A Reduction from Weighted to Unweighted Problems*

André Linhares ^{**1} and Chaitanya Swamy ^{**1}

Combinatorics and Optimization, Univ. Waterloo, Waterloo, ON N2L 3G1.
{alinhare, cswamy}@uwaterloo.ca

Abstract. We study the *min-cost chain-constrained spanning-tree* (abbreviated MCCST) problem: find a min-cost spanning tree in a graph subject to degree constraints on a nested family of node sets. We devise the *first* polytime algorithm that finds a spanning tree that (i) violates the degree constraints by at most a constant factor *and* (ii) whose cost is within a constant factor of the optimum. Previously, only an algorithm for *unweighted* CCST was known [13], which satisfied (i) but did not yield any cost bounds. This also yields the first result that obtains an $O(1)$ -factor for *both* the cost approximation and violation of degree constraints for any spanning-tree problem with general degree bounds on node sets, where an edge participates in multiple degree constraints. A notable feature of our algorithm is that we *reduce* MCCST to unweighted CCST (and then utilize [13]) via a novel application of *Lagrangian duality* to simplify the *cost structure* of the underlying problem and obtain a decomposition into certain uniform-cost subproblems.

We show that this Lagrangian-relaxation based idea is in fact applicable more generally and, for any cost-minimization problem with packing side-constraints, yields a reduction from the weighted to the unweighted problem. We believe that this reduction is of independent interest. As another application of our technique, we consider the *k-budgeted matroid basis* problem, where we build upon a recent rounding algorithm of [4] to obtain an improved $n^{O(k^{1.5}/\epsilon)}$ -time algorithm that returns a solution that satisfies (any) one of the budget constraints exactly and incurs a $(1 + \epsilon)$ -violation of the other budget constraints.

1 Introduction

Constrained spanning-tree problems, where one seeks a minimum-cost spanning tree satisfying additional ($\{0, 1\}$ -coefficient) packing constraints, constitute an important and widely-studied class of problems. In particular, when the packing constraints correspond to node-degree bounds, we obtain the classical *min-cost bounded-degree spanning tree* (MBDST) problem, which has a rich history of

* A full version of the paper is available on the CS arXiv.

** Research supported partially by NSERC grant 327620-09 and the second author's Discovery Accelerator Supplement Award, and Ontario Early Researcher Award.

study [7, 11, 12, 5, 10, 15] culminating in the work of [15] that yielded an optimal result for MBDST. Such degree-constrained network-design problems arise in diverse areas including VLSI design, vehicle routing and communication networks (see, e.g., the references in [14]), and their study has led to the development of powerful techniques in approximation algorithms.

Whereas the *iterative rounding and relaxation* technique introduced in [15] (which extends the iterative-rounding framework of [9]) yields a versatile technique for handling node-degree constraints (even for more-general network-design problems), we have a rather limited understanding of spanning-tree problems with more-general degree constraints, such as constraints $|T \cap \delta(S)| \leq b_S$ for sets S in some (structured) family \mathcal{S} of node sets.¹ A fundamental impediment here is our inability to leverage the techniques in [10, 15]. The few known results yield: (a) (sub-) optimal cost, but a *super-constant* additive- or multiplicative-factor violation of the degree bounds [3, 1, 6, 2]; or (b) a multiplicative $O(1)$ -factor violation of the degree bounds (when \mathcal{S} is a nested family), but *no cost guarantee* [13]. In particular, in stark contrast to the results known for node-degree-bounded network-design problems, there is no known algorithm that yields an $O(1)$ -factor cost approximation *and* an (additive or multiplicative) $O(1)$ -factor violation of the degree bounds. (Such guarantees are only known when each edge participates in $O(1)$ degree constraints [2]; see however [16] for an exception.)

We consider the *min-cost chain-constrained spanning-tree* (MCCST) problem introduced by [13], which is perhaps the most-basic setting involving general degree bounds where there is a significant gap in our understanding vis-a-vis node-degree bounded problems. In MCCST, we are given an undirected connected graph $G = (V, E)$, nonnegative edge costs $\{c_e\}$, a nested family \mathcal{S} (or *chain*) of node sets $S_1 \subsetneq S_2 \subsetneq \dots \subsetneq S_\ell \subsetneq V$, and integer degree bounds $\{b_S\}_{S \in \mathcal{S}}$. The goal is to find a minimum-cost spanning tree T such that $|\delta_T(S)| \leq b_S$ for all $S \in \mathcal{S}$, where $\delta_T(S) := T \cap \delta(S)$. Olver and Zenklus [13] give an algorithm for *unweighted CCST* that returns a tree T such that $|\delta_T(S)| = O(b_S)$ (i.e., there is *no bound on $c(T)$*), and show that, for some $\rho > 0$, it is *NP*-complete to obtain an additive $\rho \cdot \frac{\log |V|}{\log \log |V|}$ violation of the degree bounds. We therefore focus on bicriteria (α, β) -guarantees for MCCST, where the tree T returned satisfies $c(T) \leq \alpha \cdot OPT$ and $|\delta_T(S)| \leq \beta \cdot b_S$ for all $S \in \mathcal{S}$.

Our contributions. Our main result is the *first* $(O(1), O(1))$ -approximation algorithm for MCCST. Given any $\lambda > 1$, our algorithm returns a tree T with $c(T) \leq \frac{\lambda}{\lambda-1} \cdot OPT$ and $|\delta_T(S)| \leq 9\lambda \cdot b_S$ for all $S \in \mathcal{S}$, using the algorithm of [13] for unweighted CCST, denoted \mathcal{A}_{OZ} , as a black box (Theorem 3). As noted above, this is also the *first* algorithm that achieves an $(O(1), O(1))$ -approximation for any spanning-tree problem with general degree constraints where an edge belongs to a super-constant number of degree constraints.

We show in Section 4 that our techniques are applicable more generally. We give a *reduction* showing that for *any* cost-minimization problem with packing

¹ Such general degree constraints arise in the context of finding *thin trees* [1], where \mathcal{S} consists of all node sets, which turn out to be a very useful tool in devising approximation algorithms for *asymmetric TSP*.

side-constraints, if we have an algorithm for the *unweighted* problem that returns a solution with an $O(1)$ -factor violation of the packing constraints and satisfies a certain property, then one can utilize it to obtain an $(O(1), O(1))$ -approximation for the cost-minimization problem. Furthermore, we show that if the algorithm for the unweighted counterpart satisfies a stronger property, then we can utilize it to obtain a $(1, O(1))$ -approximation (Theorem 9).

We believe that our reductions are of independent interest and will be useful in other settings as well. Demonstrating this, we show an application to the *k-budgeted matroid basis* problem, wherein we seek to find a basis satisfying k budget constraints. Grandoni et al. [8] devised an $n^{O(k^2/\epsilon)}$ -time algorithm that returned a $(1, 1+\epsilon, \dots, 1+\epsilon)$ -solution: i.e., the solution satisfies (any) one budget constraint exactly and violates the other budget constraints by a $(1+\epsilon)$ -factor (if the problem is feasible). Very recently, Bansal and Nagarajan [4] improved the running time to $n^{O(k^{1.5}/\epsilon)}$ but return only a $(1+\epsilon, \dots, 1+\epsilon)$ -solution. Applying our reduction (to the algorithm in [4]), we obtain the *best of both worlds*: we return a $(1, 1+\epsilon, \dots, 1+\epsilon)$ -solution in $n^{O(k^{1.5}/\epsilon)}$ -time (Theorem 12).

The chief novelty in our algorithm and analysis, and the key underlying idea, is an unorthodox use of *Lagrangian duality*. Whereas typically Lagrangian relaxation is used to drop complicating constraints and thereby simplify the constraint structure of the underlying problem, in contrast, we use Lagrangian duality to simplify the *cost structure* of the underlying problem by equalizing edge costs in certain subproblems. To elaborate (see Section 3.1), the algorithm in [13] for unweighted CCST can be viewed as taking a solution x to the natural linear-programming (LP) relaxation for MCCST, converting it to another feasible solution x' satisfying a certain structural property, and exploiting this property to round x' to a spanning tree. The main bottleneck here in handling costs (as also noted in [13]) is that $c^\top x'$ could be much larger than $c^\top x$ since the conversion ignores the c_e s and works with an alternate “potential” function.

Our crucial insight is that *we can exploit Lagrangian duality to obtain perturbed edge costs $\{c_e^{y^*}\}$ such that the change in perturbed cost due to the conversion process is bounded*. Loosely speaking, if the conversion process shifts weight from x_f to x_e , then we ensure that $c_e^{y^*} = c_f^{y^*}$ (see Lemma 5); thus, $(c^{y^*})^\top x = (c^{y^*})^\top x'$! The perturbation also ensures that applying \mathcal{A}_{OZ} to x' yields a tree whose perturbed cost is equal to $(c^{y^*})^\top x' = (c^{y^*})^\top x$. Finally, we show that for an optimal LP solution x^* , the “error” $(c^{y^*} - c)^\top x^*$ incurred in working with the c^{y^*} -cost is $O(OPT)$; this yields the $(O(1), O(1))$ -approximation.

We extend the above idea to an arbitrary cost-minimization problem with packing side-constraints as follows. Let x^* be an optimal solution to the LP-relaxation, and \mathcal{P} be the polytope obtained by dropping the packing constraints. We observe that the same Lagrangian-duality based perturbation ensures that all points on the minimal face of \mathcal{P} containing x^* have the same perturbed cost. Therefore, if we have an algorithm for the unweighted problem that rounds x^* to a point \hat{x} on this minimal face, then we again obtain that $(c^{y^*})^\top \hat{x} = (c^{y^*})^\top x^*$, which then leads to an $(O(1), O(1))$ -approximation (as in the case of MCCST).

Related work. Whereas node-degree-bounded spanning-tree problems have been widely studied, relatively few results are known for spanning-tree problems with general degree constraints for a family \mathcal{S} of node-sets. With the exception of the result of [13] for unweighted CCST, these other results [3, 1, 6, 2] all yield a tree of cost at most the optimum with an $\omega(1)$ additive- or multiplicative-factor violation of the degree bounds. Both [3] and [2] obtain additive factors via iterative rounding and relaxation. The factor in [3] is $(r - 1)$ for an arbitrary \mathcal{S} , where r is the maximum number of degree constraints involving an edge (which could be $|V|$ even when \mathcal{S} is a chain), while [2] yields an $O(\log |V|)$ factor when \mathcal{S} is a laminar family (the factor does not improve when \mathcal{S} is a chain). The dependent-rounding techniques in [1, 6] yield a tree T satisfying $|\delta_T(S)| \leq \min\{O(\frac{\log |S|}{\log \log |S|})b_S, (1 + \varepsilon)b_S + O(\frac{\log |S|}{\varepsilon})\}$ for all $S \in \mathcal{S}$, for any family \mathcal{S} .

For MBDST, Goemans [10] obtained the first $(O(1), O(1))$ -approximation; his result yields a tree of cost at most the optimum and at most +2 violation of the degree bounds. This was subsequently improved to an (optimal) additive +1 violation by [15]. Zenklusen [16] considers an orthogonal generalization of MBDST, where there is a matroid-independence constraint on the edges incident to each node, and obtains a tree of cost at most the optimum and “additive” $O(1)$ violation (defined appropriately) of the matroid constraints. To our knowledge, this is the only prior work that obtains an $O(1)$ -approximation to both the cost and packing constraints for a constrained spanning-tree problem where an edge participates in $\omega(1)$ packing constraints (albeit this problem is quite different from spanning tree with general degree constraints).

Finally, we note that our Lagrangian-relaxation based technique is somewhat similar to its use in [11]. However, whereas [11] uses this to reduce uniform-degree MBDST to the problem of finding an MST of minimum maximum degree, which is another *weighted* problem, we utilize Lagrangian relaxation in a more refined fashion to reduce the weighted problem to its *unweighted* counterpart.

2 An LP-relaxation for MCCST and preliminaries

We consider the following natural LP-relaxation for MCCST. Throughout, we use e to index the edges of the underlying graph $G = (V, E)$. For a set $S \subseteq V$, let $E(S)$ denote $\{uv \in E : u, v \in S\}$, and $\delta(S)$ denote the edges on the boundary of S . For a vector $z \in \mathbb{R}^E$ and an edge-set F , we use $z(F)$ to denote $\sum_{e \in F} z_e$.

$$\min \sum_e c_e x_e \tag{P}$$

$$\text{s.t.} \quad x(E(S)) \leq |S| - 1 \quad \forall \emptyset \neq S \subsetneq V \tag{1}$$

$$x(E) = |V| - 1 \tag{2}$$

$$x(\delta(S)) \leq b_S \quad \forall S \in \mathcal{S} \tag{3}$$

$$x \geq 0. \tag{4}$$

For any $x \in \mathbb{R}_+^E$, let $\text{supp}(x) := \{e : x_e > 0\}$ denote the support of x . It is well known that the polytope, $P_{\text{ST}}(G)$, defined by (1), (2), and (4) is the convex hull of

spanning trees of G . We call points in $P_{\text{ST}}(G)$ *fractional spanning trees*. We refer to (1), (2) as the *spanning-tree constraints*. We will also utilize (P_λ) , the modified version of (P) where we replace (3) with $x(\delta(S)) \leq \lambda b_S$ for all $S \in \mathcal{S}$, where $\lambda \geq 1$. Let $OPT(\lambda)$ denote the optimal value of (P_λ) , and let $OPT := OPT(1)$.

Preliminaries. A family $\mathcal{L} \subseteq 2^V$ of sets is a *laminar family* if for all $A, B \in \mathcal{L}$, we have $A \subseteq B$ or $B \subseteq A$ or $A \cap B = \emptyset$. As is standard, we say that $A \in \mathcal{L}$ is a child of $L \in \mathcal{L}$ if L is the minimal set of \mathcal{L} such that $A \subsetneq L$. For each $L \in \mathcal{L}$, let $G_L^\mathcal{L} = (V_L^\mathcal{L}, E_L^\mathcal{L})$ be the graph obtained from $(L, E(L))$ by contracting the children of L in \mathcal{L} ; we drop the superscript \mathcal{L} when \mathcal{L} is clear from the context.

Given $x \in P_{\text{ST}}(G)$, define a *laminar decomposition* \mathcal{L} of x to be a (inclusion-wise) maximal laminar family of sets whose spanning-tree constraints are tight at x , so $x(E(A)) = |A| - 1$ for all $A \in \mathcal{L}$. Note that $V \in \mathcal{L}$ and $\{v\} \in \mathcal{L}$ for all $v \in V$. A laminar decomposition can be constructed in polytime (using network-flow techniques). For any $L \in \mathcal{L}$, let $x_L^\mathcal{L}$, or simply x_L if \mathcal{L} is clear from context, denote x restricted to E_L . Observe that x_L is a fractional spanning tree of G_L .

3 An LP-rounding approximation algorithm

3.1 An overview

We first give a high-level overview. Clearly, if (P) is infeasible, there is no spanning tree satisfying the degree constraints, so in the sequel, we assume that (P) is feasible. We seek to obtain a spanning tree T of cost $c(T) = O(OPT)$ such that $|\delta_T(S)| = O(b_S)$ for all $S \in \mathcal{S}$, where $\delta_T(S)$ is the set of edges of T crossing S .

In order to explain the key ideas leading to our algorithm, we first briefly discuss the approach of Olver and Zenklus [13] for unweighted CCST. Their approach *ignores* the edge costs $\{c_e\}$ and instead starts with a feasible solution x to (P) that minimizes a suitable (linear) potential function. They use this potential function to argue that if \mathcal{L} is a laminar decomposition of x , then (x, \mathcal{L}) satisfies a key structural property called *rainbow freeness*. Exploiting this, they give a rounding algorithm, hereby referred to as \mathcal{A}_{OZ} , that for every $L \in \mathcal{L}$, rounds x_L to a spanning tree T_L of G_L such that $|\delta_{T_L}(S)| \approx O(x_L(\delta(S)))$ for all $S \in \mathcal{S}$, so that concatenating the T_L s yields a spanning tree T of G satisfying $|\delta_T(S)| = O(x(\delta(S))) = O(b_S)$ for all $S \in \mathcal{S}$ (Theorem 2). However, as already noted in [13], a fundamental obstacle towards generalizing their approach to handle the weighted version (i.e., MCCST) is that in order to achieve rainbow freeness, which is crucial for their rounding algorithm, one needs to *abandon the cost function c and work with an alternate potential function*.

We circumvent this difficulty as follows. First, we note that the algorithm in [13] can be equivalently viewed as rounding an *arbitrary* solution x to (P) as follows. Let \mathcal{L} be a laminar decomposition of x . Using the same potential-function idea, we can convert x to another solution x' to (P) that admits a laminar decomposition \mathcal{L}' refining \mathcal{L} such that (x', \mathcal{L}') satisfies rainbow freeness (see Lemma 1), and then round x' using \mathcal{A}_{OZ} . Of course, the difficulty noted above remains, since the move to rainbow freeness (which again ignores c and

uses a potential function) does not yield *any* bounds on the cost $c^\top x'$ relative to $c^\top x$. We observe however that there is one simple property (*) under which one can guarantee that $c^\top x' = c^\top x$, namely, if for every $L \in \mathcal{L}$, all edges in $\text{supp}(x) \cap E_L$ have the same cost. However, it is unclear how to utilize this observation since there is no reason to expect our instance to have this rather special property: for instance, all edges of G could have very different costs!

Now let x^* be an optimal solution to (P) (we will later modify this somewhat) and \mathcal{L} be a laminar decomposition of x^* . The crucial insight that allows us to leverage property (*), and a key notable aspect of our algorithm and analysis, is that *one can leverage Lagrangian duality to suitably perturb the edge costs so that the perturbed costs satisfy property (*)*. More precisely, letting $y^* \in \mathbb{R}_+^{\mathcal{S}}$ denote the optimal values of the dual variables corresponding to constraints (3), if we define the perturbed cost of edge e to be $c_e^{y^*} := c_e + \sum_{S \in \mathcal{S}: e \in \delta(S)} y_S^*$, then *the c^{y^*} -cost of all edges in $\text{supp}(x^*) \cap E_L$ are indeed equal, for every $L \in \mathcal{L}$* (Lemma 5). In essence, this holds because for any $e' \in \text{supp}(x^*)$, by complementary slackness, we have $c_{e'} = (\text{dual contribution to } e' \text{ from (1),(2)}) - \sum_{S \in \mathcal{S}: e' \in \delta(S)} y_S^*$. Since any two edges $e, f \in \text{supp}(x^*) \cap E_L$ appear in the *same sets of \mathcal{L}* , one can argue that the dual contributions to e and f from (1), (2) are *equal*, and thus, $c_e^{y^*} = c_f^{y^*}$.

Now since (x^*, \mathcal{L}^*) satisfies (*) with the perturbed costs c^{y^*} , we can convert (x^*, \mathcal{L}^*) to (x', \mathcal{L}') satisfying rainbow freeness without altering the perturbed cost, and then round x' to a spanning tree T using \mathcal{A}_{OZ} . This immediately yields $|\delta_T(S)| = O(b_S)$ for all $S \in \mathcal{S}$. To bound the cost, we argue that $c(T) \leq c^{y^*}(T) = \sum_e c_e^{y^*} x_e^* = c^\top x^* + \sum_{S \in \mathcal{S}} b_S y_S^*$ (Lemma 6), where the last equality follows from complementary slackness. (Note that the perturbed costs are used only in the analysis.) However, $\sum_{S \in \mathcal{S}} b_S y_S^*$ need not be bounded in terms of $c^\top x^*$. To fix this, we modify our starting solution x^* : we solve (P_λ) (which recall is (P) with inflated degree bounds $\{\lambda b_S\}$), where $\lambda > 1$, to obtain x^* , and use this x^* in our algorithm. Now, letting y^* be the optimal dual values corresponding to the inflated degree constraints, a simple duality argument shows that $\sum_{S \in \mathcal{S}} b_S y_S^* \leq \frac{OPT(1) - OPT(\lambda)}{\lambda - 1}$ (Lemma 7), which yields $c(T) = O(OPT)$ (see Theorem 3).

A noteworthy feature of our algorithm is the rather unconventional use of Lagrangian relaxation, where we use duality to simplify the *cost structure* (as opposed to the constraint-structure) of the underlying problem by equalizing edge costs in certain subproblems. This turns out to be the crucial ingredient that allows us to utilize the algorithm of [13] for unweighted CCST *as a black box* without worrying about the difficulties posed by (the move to) rainbow freeness. In fact, as we show in Section 4, this Lagrangian-relaxation idea is applicable more generally, and yields a novel reduction from weighted problems to their unweighted counterparts. We believe that this reduction is of independent interest and will find use in other settings as well.

3.2 Algorithm details and analysis

To specify our algorithm formally, we first define the rainbow-freeness property and state the main result of [13] (which we utilize as a black box) precisely.

For an edge e , define $\mathcal{S}_e := \{S \in \mathcal{S} : e \in \delta(S)\}$. Note that \mathcal{S}_e could be empty. We say that two edges $e, f \in E$ form a *rainbow* if $\mathcal{S}_e \subseteq \mathcal{S}_f$ or $\mathcal{S}_f \subseteq \mathcal{S}_e$. (This definition is slightly different from the one in [13], in that we allow $\mathcal{S}_e = \mathcal{S}_f$.) We say that (x, \mathcal{L}) is a *rainbow-free decomposition* if \mathcal{L} is a laminar decomposition of x and for every set $L \in \mathcal{L}$, no two edges of $\text{supp}(x) \cap E_L$ form a rainbow. (Recall that $G_L = (V_L, E_L)$ denotes the graph obtained from $(L, E(L))$ by contracting the children of L .) The following lemma shows that one can convert an arbitrary decomposition (x, \mathcal{L}) to a rainbow-free one; the proof mimics the potential-function argument in [13] and is deferred to the full version.

Lemma 1. *Let $x \in P_{ST}(G)$ and \mathcal{L} be a laminar decomposition of x . We can compute in polytime a fractional spanning tree $x' \in P_{ST}(G)$ and a rainbow-free decomposition (x', \mathcal{L}') such that: (i) $\text{supp}(x') \subseteq \text{supp}(x)$; (ii) $\mathcal{L} \subseteq \mathcal{L}'$; and (iii) $x'(\delta(S)) \leq x(\delta(S))$ for all $S \in \mathcal{S}$.*

Theorem 2 ([13]). *There is a polytime algorithm, \mathcal{A}_{OZ} , that given a fractional spanning tree $x' \in P_{ST}(G)$ and a rainbow-free decomposition (x', \mathcal{L}') , returns a spanning tree $T_L \subseteq \text{supp}(x')$ of G_L for every $L \in \mathcal{L}'$ such that the concatenation T of the T_L s is a spanning tree of G satisfying $|\delta_T(S)| \leq 9x'(\delta(S))$ for all $S \in \mathcal{S}$.*

We can now describe our algorithm compactly. Let $\lambda > 1$ be a parameter.

1. Compute an optimal solution x^* to (P_λ) , a laminar decomposition \mathcal{L} of x^* .
2. Apply Lemma 1 to (x^*, \mathcal{L}) to obtain a rainbow-free decomposition (x', \mathcal{L}') .
3. Apply \mathcal{A}_{OZ} to the input (x', \mathcal{L}') to obtain spanning trees $T_L^{\mathcal{L}'}$ of $G_L^{\mathcal{L}'}$ for every $L \in \mathcal{L}'$. Return the concatenation T of all the $T_L^{\mathcal{L}'}$ s.

Analysis. We show that the above algorithm satisfies the following guarantee.

Theorem 3. *The above algorithm run with parameter $\lambda > 1$ returns a spanning tree T satisfying $c(T) \leq \frac{\lambda}{\lambda-1} \cdot OPT$ and $|\delta_T(S)| \leq 9\lambda b_S$ for all $S \in \mathcal{S}$.*

The bound on $|\delta_T(S)|$ follows immediately from Lemma 1 and Theorem 2 since x^* , and hence x' obtained in step 2, is a feasible solution to (P_λ) . So we focus on bounding $c(T)$. This will follow from three things. First, we define the perturbed c^{y^*} -cost precisely. Next, Lemma 5 proves the key result that for every $L \in \mathcal{L}$, all edges in $\text{supp}(x^*) \cap E_L$ have the same perturbed cost. Using this it is easy to show that $c(T) \leq c^{y^*}(T) = \sum_e c_e^{y^*} x_e^* = OPT(\lambda) + \lambda \sum_{S \in \mathcal{S}} b_S y_S^*$ (Lemma 6). Finally, we show that $\sum_{S \in \mathcal{S}} b_S y_S^* \leq \frac{OPT - OPT(\lambda)}{\lambda - 1}$ (Lemma 7), which yields the bound stated in Theorem 3.

To define the perturbed costs, we consider the Lagrangian dual of (P_λ) obtained by dualizing the (inflated) degree constraints $x(\delta(S)) \leq \lambda b_S$ for all $S \in \mathcal{S}$:

$$\max_{y \in \mathbb{R}_+^{\mathcal{S}}} \left(g_\lambda(y) := \min_{x \in P_{ST}(G)} \left(\sum_e c_e x_e + \sum_{S \in \mathcal{S}} (x(\delta(S)) - \lambda b_S) y_S \right) \right). \quad (\text{LD}\lambda)$$

For $x \in \mathbb{R}^E$ and $y \in \mathbb{R}^{\mathcal{S}}$, let $\mathcal{G}_\lambda(x, y) := \sum_e c_e x_e + \sum_{S \in \mathcal{S}} (x(\delta(S)) - \lambda b_S) y_S = \sum_e c_e^y x_e - \lambda \sum_{S \in \mathcal{S}} b_S y_S$ denote the objective function of $g_\lambda(y)$, where $c_e^y := c_e + \sum_{S \in \mathcal{S}: e \in \delta(S)} y_S$.

Let y^* be an optimal solution to (LD_λ) . One can show via LP-duality that this holds iff there exist dual multipliers $\mu^* = (\mu_S^*)_{\emptyset \neq S \subseteq V}$ corresponding to constraints (1), (2) of (P_λ) such that (μ^*, y^*) is an optimal solution to the LP dual of (P_λ) . This also implies that $g_\lambda(y^*) = \mathcal{G}_\lambda(x^*, y^*) = OPT(\lambda)$. Our *perturbed costs* are $\{c_e^{y^*}\}$. We defer the proof of Lemma 4 to the full version and use it to show that $c_e^{y^*} = c_f^{y^*}$ for any $L \in \mathcal{L}$ and any edges $e, f \in \text{supp}(x^*) \cap E_L$.

Lemma 4. *We have $g_\lambda(y^*) = \mathcal{G}_\lambda(x^*, y^*) = OPT(\lambda)$. Further, there exists $\mu^* = (\mu_S^*)_{\emptyset \neq S \subseteq V}$ such that (μ^*, y^*) is an optimal solution to the dual (D_λ) of (P_λ) .*

Lemma 5. *For any $L \in \mathcal{L}$, all edges of $\text{supp}(x^*) \cap E_L$ have the same c^{y^*} -cost.*

Proof sketch. Consider any two edges $e, f \in \text{supp}(x^*) \cap E_L$. Suppose for a contradiction that $c_e^{y^*} < c_f^{y^*}$. Obtain \hat{x} from x^* by increasing x_e^* by ϵ and decreasing x_f^* by ϵ (so $\hat{x}_{e'} = x_{e'}^*$ for all $e' \notin \{e, f\}$). We argue that one can choose sufficiently small $\epsilon > 0$ such that $\hat{x} \in \text{P}_{\text{ST}}(G)$. This follows since any spanning-tree constraint that is tight at x^* can be expressed as a linear combination of the spanning-tree constraints for the sets in \mathcal{L} . Since $c_e^{y^*} < c_f^{y^*}$, we have $g_\lambda(y^*) \leq \mathcal{G}_\lambda(\hat{x}, y^*) < \mathcal{G}_\lambda(x^*, y^*) = g_\lambda(y^*)$, which is a contradiction. \square

Lemma 6. *We have $c(T) \leq \sum_e c_e^{y^*} x_e^* = \sum_e c_e x_e^* + \lambda \sum_{S \in \mathcal{S}} b_S y_S^*$.*

Proof. Observe that $c(T) \leq c^{y^*}(T)$ since $c_e \leq c_e^{y^*}$ for all $e \in E$ as $y^* \geq 0$. We now bound $c^{y^*}(T)$. To keep notation simple, we use $G_L = (V_L, E_L)$ and x_L^* to denote $G_L^\mathcal{L}$ and $(x^*)_L^\mathcal{L}$ (which recall is x^* restricted to $E_L^\mathcal{L}$) respectively, and $G'_L = (V'_L, E'_L)$ and x'_L to denote G'_L and $(x^*)'_L$ respectively.

We have $c^{y^*}(T) = \sum_{L \in \mathcal{L}} c^{y^*}(T \cap E_L)$ since the sets $\{E_L\}_{L \in \mathcal{L}}$ partition E . Fix $L \in \mathcal{L}$. Note that x_L^* is a fractional spanning tree of $G_L = (V_L, E_L)$ since for any $\emptyset \neq Q \subseteq V_L$, if R is the subset of V corresponding to Q , and A_1, \dots, A_k are the children of L whose corresponding contracted nodes lie in Q , we have $x_L^*(E_L(Q)) = x^*(E(R)) - \sum_{i=1}^k x^*(E(A_i)) \leq |R \setminus (A_1 \cup \dots \cup A_k)| + k - 1 = |Q| - 1$ with equality holding when $Q = V_L$. Note that $T \cap E_L$ is a spanning tree of G_L since T is obtained by concatenating spanning trees for the graphs $\{G'_L\}_{L' \in \mathcal{L}'}$, and \mathcal{L}' refines \mathcal{L} . Also, all edges of $\text{supp}(x^*) \cap E_L$ have the same c^{y^*} -cost by Lemma 5. So we have $c^{y^*}(T \cap E_L) = \sum_{e \in E_L} c_e^{y^*} x_e^*$. It follows that

$$\begin{aligned} c^{y^*}(T) &= \sum_e c_e^{y^*} x_e^* = \sum_e \left(c_e x_e^* + \sum_{S \in \mathcal{S}: e \in \delta(S)} y_S^* x_e^* \right) \\ &= \sum_e c_e x_e^* + \sum_{S \in \mathcal{S}} y_S^* x^*(\delta(S)) = \sum_e c_e x_e^* + \lambda \sum_{S \in \mathcal{S}} b_S y_S^*. \quad \square \end{aligned}$$

Lemma 7. *We have $\sum_{S \in \mathcal{S}} b_S y_S^* \leq \frac{OPT(1) - OPT(\lambda)}{\lambda - 1}$.*

Proof. By Lemma 4, there exists μ^* such that (μ^*, y^*) is an optimal solution to the dual (D_λ) of (P_λ) . Also, (μ^*, y^*) is a feasible solution to (D_1) . Therefore,

$$OPT(1) \geq - \sum_{\emptyset \neq S \subseteq V} (|S| - 1) \mu_S^* - \sum_{S \in \mathcal{S}} b_S y_S^*, \quad OPT(\lambda) = - \sum_{\emptyset \neq S \subseteq V} (|S| - 1) \mu_S^* - \lambda \sum_{S \in \mathcal{S}} b_S y_S^*.$$

Hence $OPT(1) - OPT(\lambda) \geq (\lambda - 1) \sum_{S \in \mathcal{S}} b_S y_S^*$. \square

Proof of Theorem 3. As noted earlier, the bounds on $\delta_T(S)$ follow immediately from Lemma 1 and Theorem 2: for any $S \in \mathcal{S}$, we have $|\delta_T(S)| \leq 9x'(\delta(S)) \leq 9x^*(\delta(S)) \leq 9\lambda b_S$. The bound on $c(T)$ follows from Lemmas 6 and 7 since $\sum_e c_e x_e^* = OPT(\lambda)$. \square

4 A reduction from weighted to unweighted problems

We now show that our ideas are applicable more generally, and yield bicriteria approximation algorithms for any cost-minimization problem with packing side-constraints, provided we have a suitable approximation algorithm for the *unweighted* counterpart. Thus, our technique yields a *reduction* from weighted to unweighted problems, which we believe is of independent interest.

To demonstrate this, we first isolate the key properties of the rounding algorithm \mathcal{B} used above for unweighted CCST that enable us to use it as a black box to obtain our result for MCCST; this will yield an alternate, illuminating explanation of Theorem 3. Note that \mathcal{B} is obtained by combining the procedure in Lemma 1 and \mathcal{A}_{OZ} (Theorem 2). First, we of course utilize that \mathcal{B} is an approximation algorithm for unweighted CCST, so it returns a spanning tree T such that $|\delta_T(S)| = O(x^*(\delta(S)))$ for all $S \in \mathcal{S}$. Second, we exploit the fact that \mathcal{B} returns a tree T that *preserves tightness of all spanning-tree constraints that are tight at x^** . This is the crucial property that allows us to bound $c(T)$, since this implies (as we explain below) that $c^{y^*}(T) = \sum_e c_e^{y^*} x_e^*$, which then yields the bound on $c(T)$ as before. The equality follows because the set of optimal solutions to the LP $\min_{x \in P_{ST}(G)} \mathcal{G}_\lambda(x, y^*)$ is a face of $P_{ST}(G)$; thus *all* points on the *minimal* face of $P_{ST}(G)$ containing x^* are optimal solutions to this LP, and the stated property implies that the characteristic vector of T lies on this minimal face. In other words, while \mathcal{A}_{OZ} proceeds by exploiting the notions of rainbow freeness and laminar decomposition, these notions are not essential to obtaining our result; *any* rounding algorithm for unweighted CCST satisfying the above two properties can be utilized to obtain our guarantee for MCCST.

We now formalize the above two properties for an arbitrary cost-minimization problem with packing side-constraints, and prove that they suffice to yield a bicriteria guarantee. Consider the following abstract problem:

$$\min c^\top x \quad \text{s.t.} \quad x \text{ is an extreme point of } \mathcal{P}, \quad Ax \leq b, \quad (\text{Q}^\mathcal{P})$$

where $\mathcal{P} \subseteq \mathbb{R}_+^n$ is a fixed polytope, $c, b \geq 0$, and $A \geq 0$. Observe that we can cast MCCST as a special case of $(\text{Q}^\mathcal{P})$, by taking $\mathcal{P} = P_{ST}(G)$ (whose extreme points are spanning trees of G), c to be the edge costs, and $Ax \leq b$ to be the degree constraints. Moreover, by taking \mathcal{P} to be the convex hull of a bounded set $\{x \in \mathbb{Z}_+^n : Cx \leq d\}$ we can use $(\text{Q}^\mathcal{P})$ to encode a discrete-optimization problem.

We say that \mathcal{B} is a *face-preserving rounding algorithm* (FPRA) for unweighted $(\text{Q}^\mathcal{P})$ if given any point $x \in \mathcal{P}$, \mathcal{B} finds in polytime an extreme point \hat{x} of \mathcal{P} such that:

$$(P1) \quad \hat{x} \text{ belongs to the minimal face of } \mathcal{P} \text{ that contains } x.$$

We say that \mathcal{B} is a β -approximation FPRA (where $\beta \geq 1$) if we also have:

$$(P2) \quad A\hat{x} \leq \beta Ax.$$

Let $(R_\lambda^{\mathcal{P}})$ denote the linear program: $\min\{c^\top x : x \in \mathcal{P}, Ax \leq \lambda b\}$; note that $(R_1^{\mathcal{P}})$ is the LP-relaxation of $(Q^{\mathcal{P}})$. Let $\text{opt}(\lambda)$ denote the optimal value of $(R_\lambda^{\mathcal{P}})$, and let $\text{opt} := \text{opt}(1)$. We say that an algorithm is a (ρ_1, ρ_2) -approximation algorithm for $(Q^{\mathcal{P}})$ if it finds in polytime an extreme point \hat{x} of \mathcal{P} such that $c^\top \hat{x} \leq \rho_1 \text{opt}$ and $A\hat{x} \leq \rho_2 b$.

Theorem 8. *Let \mathcal{B} be a β -approximation FPRA for unweighted $(Q^{\mathcal{P}})$. Then, given any $\lambda > 1$, one can obtain a $(\frac{\lambda}{\lambda-1}, \beta\lambda)$ -approximation algorithm for $(Q^{\mathcal{P}})$ using a single call to \mathcal{B} .*

Proof sketch. We dovetail the algorithm for MCCST and its analysis. We simply compute an optimal solution x^* to $(R_\lambda^{\mathcal{P}})$ and round it to an extreme point \hat{x} of \mathcal{P} using \mathcal{B} . By property (P2), it is clear that $A\hat{x} \leq \beta(Ax^*) \leq \beta\lambda b$.

Let m be the number of rows of A . For $y \in \mathbb{R}_+^m$, define $c^y := c + A^\top y$. To bound the cost, as before, we consider the Lagrangian dual of $(R_\lambda^{\mathcal{P}})$ obtained by dualizing the side-constraints $Ax \leq \lambda b$.

$$\max_{y \in \mathbb{R}_+^m} \left(h_\lambda(y) := \min_{x \in \mathcal{P}} \mathcal{H}_\lambda(x, y) \right), \quad \text{where } \mathcal{H}_\lambda(x, y) := (c^y)^\top x - \lambda y^\top b.$$

Let $y^* = \text{argmax}_{y \in \mathbb{R}_+^m} h_\lambda(y)$. We can mimic the proof of Lemma 4 to show that x^* is an optimal solution to $\min_{x \in \mathcal{P}} \mathcal{H}_\lambda(x, y^*)$. The set of optimal solutions to this LP is a face of \mathcal{P} . So all points on the minimal face of \mathcal{P} containing x^* are optimal solutions to this LP. By property (P1), \hat{x} belongs to this minimal face and so is an optimal solution to this LP. So $(c^{y^*})^\top \hat{x} = (c^{y^*})^\top x^* = c^\top x^* + (y^*)^\top Ax^* = \text{opt}(\lambda) + \lambda(y^*)^\top b$, where the last equality follows by complementary slackness. Also, by the same arguments as in Lemma 7, we have $(y^*)^\top b \leq \frac{\text{opt}(1) - \text{opt}(\lambda)}{\lambda - 1}$. Since $c \leq c^{y^*}$, we have $c^\top \hat{x} \leq (c^{y^*})^\top \hat{x} \leq \frac{\lambda}{\lambda - 1} \cdot \text{opt}$. \square

5 Towards a $(1, O(1))$ -approximation algorithm for $(Q^{\mathcal{P}})$

A natural question that emerges from Theorems 3 and 8 is whether one can obtain a $(1, O(1))$ -approximation, i.e., obtain a solution of *cost at most* opt that violates the packing side-constraints by an (multiplicative) $O(1)$ -factor. Such results are known for degree-bounded spanning tree problems with various kinds of degree constraints [10, 15, 3, 16], so, in particular, it is natural to ask whether such a result also holds for MCCST. (Note that for MCCST, the dependent-rounding techniques in [1, 6] yield a tree T with $c(T) \leq OPT$ and $|\delta_T(S)| \leq \min\{O(\frac{\log |S|}{\log \log |S|})b_S, (1 + \varepsilon)b_S + O(\frac{\log |S|}{\varepsilon})\}$ for all $S \in \mathcal{S}$. We show that our approach is versatile enough to yield such a guarantee provided we assume a stronger property from the rounding algorithm \mathcal{B} for unweighted $(Q^{\mathcal{P}})$.

Let A_i denote the i -th row of A , for $i = 1, \dots, m$. We say that \mathcal{B} is an (α, β) -approximation FPRA for unweighted $(Q^{\mathcal{P}})$ if *in addition* to properties (P1), (P2), it satisfies:

(P3) it rounds a feasible solution x to $(R_\alpha^{\mathcal{P}})$ to an extreme point \hat{x} of \mathcal{P} satisfying $A_i^\top \hat{x} \geq \frac{A_i^\top x}{\alpha}$ for every i such that $A_i^\top x = \alpha b_i$.

(For MCCST, property (P3) requires that $|\delta_T(S)| \geq b_S$ for every set $S \in \mathcal{S}$ whose degree constraint (in (P_α)) is tight at the fractional spanning tree x .)

Theorem 9. *Let \mathcal{B} be an (α, β) -approximation FPRA for unweighted $(Q^{\mathcal{P}})$. Then, one can obtain a $(1, \alpha\beta)$ -approximation algorithm for $(Q^{\mathcal{P}})$ using a single call to \mathcal{B} .*

Proof. We show that applying Theorem 8 with $\lambda = \alpha$ yields the claimed result. It is clear that the extreme point \hat{x} returned satisfies $A\hat{x} \leq \alpha\beta b$. As in the proof of Theorem 8, let y^* be an optimal solution to $\max_{y \in \mathbb{R}_+^m} h_\lambda(y)$ (where $\lambda = \alpha$). In Lemma 6 and the proof of Theorem 8, we use the weak bound $c^\top \hat{x} \leq (c^{y^*})^\top \hat{x}$. We tighten this to obtain the improved bound on $c^\top \hat{x}$. We have $c^\top \hat{x} = (c^{y^*})^\top \hat{x} - (y^*)^\top A\hat{x}$, and

$$(y^*)^\top A\hat{x} = \sum_{i: A_i^\top x^* = \lambda b_i} y_i^* (A_i^\top \hat{x}) \geq \sum_{i: A_i^\top x^* = \lambda b_i} \frac{y_i^* A_i^\top x^*}{\alpha} = \sum_{i: A_i^\top x^* = \lambda b_i} y_i^* b_i = (y^*)^\top b.$$

The first and last equalities above follow because $y_i^* > 0$ implies that $A_i^\top x^* = \lambda b_i$. The inequality follows from property (P3). Thus, following the rest of the arguments in the proof of Theorem 8, we obtain that

$$c^\top \hat{x} \leq (c^{y^*})^\top \hat{x} - (y^*)^\top b = c^\top x^* + (\lambda - 1)(y^*)^\top b \leq \text{opt}(1). \quad \square$$

We also obtain the following variant of Theorem 9 with two-sided *additive* guarantees (which can be proved by essentially the same arguments).

Theorem 10. *Let \mathcal{B} be an FPRA for unweighted $(Q^{\mathcal{P}})$ that given $x \in \mathcal{P}$ returns an extreme point \hat{x} of \mathcal{P} such that $Ax - \Delta \leq A\hat{x} \leq Ax + \Delta$. Using a single call to \mathcal{B} , we can obtain an extreme point \tilde{x} of \mathcal{P} such that $c^\top \tilde{x} \leq \text{opt}$ and $A\tilde{x} \leq b + 2\Delta$.*

Application to k -budgeted matroid basis. Here, we seek to find a basis S of a matroid M satisfying k budget constraints $\{d_i(S) \leq B_i\}_{1 \leq i \leq k}$, where $d_i(S) = \sum_{e \in S} d_i(e)$. Note that this can be cast a special case of $(Q^{\mathcal{P}})$, where \mathcal{P} is the basis polytope of M , the objective function encodes (say) the first budget constraint and $Ax \leq b$ encode the remaining budget constraints. Applying Theorem 10 to a recent randomized algorithm of [4], we obtain a (randomized) algorithm that, for any $\epsilon > 0$, returns in $n^{O(k^{1.5}/\epsilon)}$ time a basis that (exactly) satisfies a chosen budget constraint, and violates the other budget constraints by (at most) a $(1 + \epsilon)$ -factor, where n is the size of the ground-set of M . This *matches* the current-best approximation guarantee of [8] (who give a deterministic algorithm) *and* the current-best running time of [4].

Theorem 11 ([4]). *There exists a randomized FPRA \mathcal{B}_{BN} for unweighted $(Q^{\mathcal{P}})$ that rounds any $x \in \mathcal{P}$, to an extreme point \hat{x} of \mathcal{P} such that $Ax - O(\sqrt{k})a_{\max} \leq A\hat{x} \leq Ax + O(\sqrt{k})a_{\max}$ where $a_{\max} = \max_{1 \leq i \leq k, 1 \leq j \leq n} a_{ij} = \max_{1 \leq i \leq k, 1 \leq j \leq n} d_i(j)$.*

Applying Theorem 10 with $\mathcal{B}=\mathcal{B}_{\text{BN}}$, we obtain a basis S of M such that $d_1(S) \leq B_1$, and $d_i(S) \leq B_i + O(\sqrt{k})a_{\text{max}}$ for all $2 \leq i \leq k$. We combine this with a partial-enumeration step, where we “guess” all elements of an optimal solution having $d_i(j) = \Omega\left(\frac{\epsilon}{\sqrt{k}}\right) \cdot B_i$ for any $2 \leq i \leq k$, update M and the budget constraints, and then apply Theorem 10. This yields the following result.

Theorem 12. *There exists a randomized algorithm that, given any $\epsilon > 0$, finds in $n^{O(k^{1.5}/\epsilon)}$ time a basis S of M such that $d_1(S) \leq B_1$ and $d_i(S) \leq (1 + \epsilon)B_i$ for all $2 \leq i \leq k$.*

References

1. A. Asadpour, M. Goemans, A. Madry, S. Oveis Gharan, and A. Saberi. An $O(\log n / \log \log n)$ -approximation algorithm for the asymmetric traveling salesman problem. In *Proceedings of the 20th SODA*, pages 379–389, 2010.
2. N. Bansal, R. Khandekar, J. Könemann, V. Nagarajan, and B. Peis. On generalizations of network design problems with degree bounds. *Mathematical Programming*, 141(1-2):479–506, 2013.
3. N. Bansal, R. Khandekar, and V. Nagarajan. Additive guarantees for degree-bounded directed network design. *SICOMP*, 39(4):1413–1431, 2009.
4. N. Bansal, V. Nagarajan. Approximation-friendly discrepancy rounding. To appear in *Proc., IPCO 2016*. Also appears as arXiv:1512.02254, 2015.
5. K. Chaudhuri, S. Rao, S. Riesenfeld, and K. Talwar. What would Edmonds do? Augmenting paths and witnesses for degree-bounded MSTs. *Algorithmica* 55:157–189, 2009.
6. C. Chekuri, J. Vondrak, and R. Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *51st FOCS*, 2010.
7. M. Fürer and B. Raghavachari. Approximating the minimum-degree Steiner Tree to within one of optimal. *Journal of Algorithms* 17(3):409–423, 1994.
8. F. Grandoni, R. Ravi, M. Singh, R. Zenklusen. New approaches to multi-objective optimization. *Mathematical Programming*, 146(1-2):525–554, 2014.
9. K. Jain. A factor 2 approximation algorithm for the generalized Steiner network problem. *Combinatorica*, 21:39–60, 2001.
10. M. Goemans. Minimum bounded degree spanning trees. In *47th FOCS*, 2006.
11. J. Könemann and R. Ravi. A matter of degree: improved approximation algorithms for degree-bounded minimum spanning trees. *SICOMP*, 31:1783–1793, 2002. *SIAM Journal on Computing*, 31:1783–1793, 2002.
12. J. Könemann and R. Ravi. Primal-dual meets local search: approximating MST’s with nonuniform degree bounds. In *Proceedings of the 35th STOC*, 389–395 2003.
13. N. Olver and R. Zenklusen. Chain-constrained spanning trees. In *Proceedings of the 16th IPCO*, pages 324–335, 2013.
14. R. Ravi, M. Marathe, S. Ravi, D. Rosenkrantz, and H. Hunt III. Approximation algorithms for degree-constrained minimum-cost network-design problems. *Algorithmica* 31(1):58–78, 2001.
15. M. Singh and L. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proceedings of the 39th STOC*, pages 661–670. 2007.
16. R. Zenklusen. Matroidal degree-bounded minimum spanning trees. In *Proceedings of the 23rd SODA*, pages 1512–1521, 2012.