

Chain-Constrained Spanning Trees

Neil Olver · Rico Zenklusen

the date of receipt and acceptance should be inserted later

Abstract We consider the problem of finding a spanning tree satisfying a family of additional constraints. Several settings have been considered previously, the most famous being the problem of finding a spanning tree with degree constraints. Since the problem is hard, the goal is typically to find a spanning tree that violates the constraints as little as possible.

Iterative rounding has become the tool of choice for constrained spanning tree problems. However, iterative rounding approaches are very hard to adapt to settings where an edge can be part of more than a constant number of constraints. We consider a natural constrained spanning tree problem of this type, namely where upper bounds are imposed on a family of cuts forming a chain. Our approach reduces the problem to a family of independent matroid intersection problems, leading to a spanning tree that violates each constraint by a factor of at most 9.

We also present strong hardness results: among other implications, these are the first to show, in the setting of a basic constrained spanning tree problem, a qualitative difference between what can be achieved when allowing multiplicative as opposed to additive constraint violations.

Keywords network design · spanning trees · approximation algorithms

Mathematics Subject Classification (2000) 68W25 · 90C27

This project was supported by NSF grant CCF-1115849, and an NWO Veni grant.

N.Olver
Vrije Universiteit Amsterdam, The Netherlands E-mail: n.olver@vu.nl

N.Olver
CWI, The Netherlands E-mail: olver@cwi.nl

R. Zenklusen
ETH Zürich, Switzerland E-mail: rico.zenklusen@ifor.math.ethz.ch

1 Introduction

Spanning tree problems with additional $\{0, 1\}$ -packing constraints have spawned considerable interest recently. This development was motivated by a variety of applications, including VLSI design, vehicle routing, and applications in communication networks [9, 5, 16]. Since even finding a feasible solution of a constrained spanning tree problem is typically NP-hard, the focus is on efficient procedures that either certify that the problem has no feasible solution, or find a spanning tree that violates the additional constraints by as little as possible. Often, an objective function to be minimized is also provided; here, however, we focus just on minimizing the constraint violations.

A wide variety of constrained spanning tree problems have been studied. Unfortunately, for most settings, little is known about what violation of the constraints must be accepted in order that a solution can be efficiently obtained. An exception is the most classical problem in this context, the degree-bounded spanning tree problem. Here the goal is to find a spanning tree $T \subseteq E$ in a graph $G = (V, E)$ such that T satisfies a degree bound for each vertex $v \in V$, i.e., $|\delta(v) \cap T| \leq b_v$. For this problem, Fürer and Raghavachari [9] presented an essentially best possible algorithm that either shows that no spanning tree satisfying the degree constraints exists, or returns a spanning tree violating each degree constraint by at most 1. We call this an *additive 1-approximation*, in contrast to an α -approximation, where each constraint can be violated by a *factor* $\alpha > 1$.

Recently, iterative rounding/relaxation algorithms became the tool of choice for dealing with constrained spanning tree problems. A cornerstone for this development was the work of Singh and Lau [18], which extended the iterative rounding framework of Jain [11] with a relaxation step. They obtained an additive 1-approximation even for the *minimum* degree-bounded spanning tree problem, i.e., the cost of the tree they return is upper bounded by the cost of an optimal solution not violating any constraints. This result was the culmination of a long sequence of papers presenting methods with various trade-offs between constraint violation and cost (see [12, 13, 6, 7, 10] and references therein).

Singh and Lau's iterative relaxation technique was later generalized by Bansal, Khandekar and Nagarajan [4], to show that even when upper bounds are given on an arbitrary family of edge sets E_1, \dots, E_k , one can still find a (minimum cost) spanning tree violating each constraint by not more than $\max_{e \in E} |\{i \in [k] \mid e \in E_i\}| - 1$. If each edge is only in a constant number of constraints, this leads to an additive $O(1)$ -approximation. But extending the iterative rounding technique beyond such settings seems to typically be very difficult. Some progress was achieved by Bansal, Khandekar, Könemann, Nagarajan and Peis [3], who used an iterative approach that iteratively replaces constraints by weaker ones, leading to an additive $O(\log n)$ -approximation if the constraints are upper bounds on a laminar family of cuts. They left open whether an additive or multiplicative $O(1)$ -approximation is possible in this setting, even when the cuts form a chain. Recently, Zenklusen [19] presented an additive $O(1)$ -approximation for generalized degree bounds, where for

each vertex an arbitrary matroid constraint on its adjacent edges has to be satisfied. This algorithm differs from previous iterative rounding approaches in that it successively simplifies the problem to reach a matroid intersection problem, rather than attempting to eliminate constraints until only spanning tree constraints remain.

To the best of our knowledge, with the exception of the setting of Zenklusen [19], no $O(1)$ -approximations are known for constrained spanning tree problems where an edge can lie in a super-constant number of (linear) constraints. This seems to be an important difficulty that current techniques have trouble overcoming. Furthermore, in many settings, it is not well understood if finding an additive approximation is any harder than a multiplicative one. In particular, no constrained spanning tree problem was previously known where an $O(1)$ -approximation is possible, but an additive $O(1)$ -approximation is not. The goal of this paper is to address these points by studying chain-constrained spanning trees—a natural constrained spanning tree problem that evades current techniques.

1.1 Our results

The chain-constrained spanning tree problem is the following. We are given an undirected graph $G = (V, E)$ together with a family of cuts $\emptyset \subsetneq S_1 \subsetneq S_2, \dots, \subsetneq S_\ell \subsetneq V$ forming a chain, and bounds $b_1, \dots, b_\ell \in \mathbb{Z}_{>0}$. The goal is to find a spanning tree T that satisfies all of the constraints, i.e.,

$$|T \cap \delta(S_i)| \leq b_i \quad \forall i \in [\ell], \quad (1)$$

if such a spanning tree exists. Here, $\delta(S_i)$ denotes the set of edges with precisely one endpoint in S_i .

Notice that chain constraints allow edges to be in a super-constant number of constraints. It is also a natural candidate problem that captures many of the difficulties faced when trying to construct $O(1)$ -approximations for the laminar case.

Our main algorithmic result is the following.

Theorem 1 *There is an efficient 9-approximation for the chain-constrained spanning tree problem.*

Like most work in the area, we exploit the natural LP relaxation of the problem. This relaxation asks for a point x in the spanning tree polytope P_{ST} which satisfies the constraints. But our method is not based on iterative rounding, which has become the standard tool. Instead, we reduce the problem to a family of independent matroid intersection problems. In order to do this, we decompose the problem into a number of independent subproblems, based on the laminar decomposition induced by a maximal family of independent tight spanning tree constraints of the solution to the LP relaxation. By a judicious choice of objective function, we are able to ensure that each of the resulting subproblems has a convenient structural property, namely, they have

no *rainbows* in their support. A rainbow is a pair of edges e, f such that e is in a proper superset of the chain constraints in which f is contained. Within a subproblem, the lack of rainbows yields a natural “left-to-right” ordering of the edges in its support. This ordering is crucially exploited in order to define a partition matroid with the property that any independent set of this matroid does not contribute much more to any constraint than the fractional solution for the subproblem.

Even though the high-level approach is quite clean, there are several difficulties we have to address. In particular, it turns out to be impossible to obtain a multiplicative guarantee within each subproblem separately. Instead we must use a more relaxed target for each subproblem that nevertheless yields a constant multiplicative guarantee overall.

It is interesting to compare our approach to the one taken by Goemans [10] in work giving an additive 2-approximation to the minimum degree-bounded spanning tree problem with no loss in cost. Like our result, this result is not based on iterative rounding. Instead, local sparsity of an extreme point solution is exploited to argue that the edges of G can be oriented so that each node has indegree at most 2. The degree bound at a vertex v is then relaxed to involve only the edges which are oriented away from v , so that each edge occurs in only one constraint; the degree bounds are thus replaced by a single partition matroid. The result then follows by applying matroid intersection.

Our approach can be seen in a somewhat similar light. Using the rainbow-free structure in the subproblems, we eventually end up with a partition matroid in each subproblem. One can of course combine these resulting partition matroids over all the subproblems, to obtain a single global partition matroid. So at a high level, both algorithms proceed by relaxing (or in our case, replacing) the given constraints by a matroid constraint, and then applying matroid intersection.

We complement our $O(1)$ -approximation result by showing that an additive $O(1)$ -approximation is impossible (assuming $P \neq NP$). As mentioned, this is the first result showing such a separation between what can be achieved additively and multiplicatively for a constrained spanning tree problem.

Theorem 2 *For the chain-constrained spanning tree problem it is NP-hard to distinguish between the cases where a spanning tree satisfying the chain constraints exists, and the case that every spanning tree violates some degree bound by $\Omega(\log n / \log \log n)$ units.*

This result is given in Section 3. Previously, the only hardness result of a similar nature was presented by Bansal et al. [3] for a very general constrained spanning tree problem, where constraints $|T \cap E_i| \leq b_i \forall i \in [k]$ are given for an arbitrary family of edge sets $E_1, \dots, E_k \subseteq E$. They showed that unless NP has quasi-polynomial time algorithms, there is no additive $(\log^c n)$ -approximation for this case, for some small constant $c \in (0, 1)$. Notice that our hardness result is stronger in terms of the approximation ratio, the underlying constrained spanning tree model, and the complexity assumption. Furthermore, Theorem 2

shows that the additive $O(\log n / \log \log n)$ -approximation of Bansal et al. [3] for the laminar-constrained spanning tree problem is close to optimal.

1.2 Thin trees

Given a graph G and a point x in the spanning tree polytope of G , a spanning tree T is called α -thin with respect to x if

$$|T \cap \delta(S)| \leq \alpha \cdot x(\delta(S)) \quad \forall S \subseteq V.$$

The problem of finding an α -thin tree can be interpreted as a constrained spanning tree problem, where an upper bound $b_S := x(\delta(S))$ is imposed on every cut $\delta(S)$ of the graph. By construction, this exponentially sized LP has a feasible solution, and an α -approximate solution is exactly an α -thin spanning tree.

The concept of thin spanning trees gained considerably in relevance when Asadpour et al. [2] showed that an efficient algorithm for finding an α -thin spanning tree leads to an $O(\alpha)$ -approximation for the Asymmetric Traveling Salesman Problem (ATSP)¹. In a very recent tour de force, Anari and Oveis Gharan [1] gave a nonconstructive proof of the existence of polyloglog n -thin trees, which implies the same bound on the integrality gap of ATSP. The best constructive result for thin spanning trees (and ATSP) yield $O(\log n / \log \log n)$ -thin spanning trees [2, 8]. It is open whether $O(1)$ -thin spanning trees exist, which (if shown constructively) would immediately imply an $O(1)$ -factor approximation for ATSP.

2 The algorithm

To simplify the exposition, we assume that we are dealing with a maximal chain of constraints imposed on the spanning tree. So we may label the vertices $V = \{v_1, \dots, v_n\}$ of G such that $S_i = \{v_1, \dots, v_i\} \forall i \in [n-1]$, the constraints being $|T \cap \delta(S_i)| \leq b_i$ for all $i \in [n-1]$. This is clearly not restrictive since by choosing a large right-hand side, any constraint can be made redundant.

Recall that the natural LP for the problem asks for a point x in the spanning tree polytope P_{ST} of G satisfying $x(\delta(S_i)) \leq b_i$ for all $i \in [n-1]$. Our algorithm starts with a fractional solution of this relaxation (if the relaxation is infeasible, this provides a certificate that the given instance has no solution). But we do not begin with an arbitrary feasible solution; we require an optimal solution with respect to a carefully chosen objective. More precisely, we take a solution that minimizes the total length of the edges, where the length of an edge $\{v_i, v_j\} \in E$ is $|i - j|$. Equivalently, the length of an edge is the number of chain constraints to which the edge contributes. This leads to the LP (2) shown below.

¹ Strictly speaking, Asadpour et al.'s approach required the spanning tree not only to be thin, but also to be of low cost. However this second requirement is not necessary for the mentioned statement to be true (see [15]).

Let x^* be an optimal solution to (2), which can be computed by standard techniques (see [17]). Notice that the objective function of (2) is the same as the total load on all cuts: $\sum_{i=1}^{n-1} x(\delta(S_i))$.

$$\begin{aligned} \min \quad & \sum_{\{v_i, v_j\} \in E} |i - j| \cdot x(\{v_i, v_j\}) \\ & x \in P_{ST} \\ & x(\delta(S_i)) \leq b_i \quad \forall i \in [n - 1] \end{aligned} \tag{2}$$

The above objective function is motivated by a subprocedure we use to find a spanning tree in an instance that does not contain what we call a *rainbow*. A rainbow consists of two edges $\{v_i, v_j\}, \{v_p, v_q\}$ with $i \leq p < q \leq j$ and either $i < p$ or $q < j$, i.e., the first edge is in a proper superset of the chain constraints in which the second edge is in. Even though the above objective function does not necessarily lead to an LP solution x^* whose support $\text{supp}(x^*) = \{e \in E \mid x^*(e) > 0\}$ does not contain rainbows—a feasible rainbow-free solution may not even exist—it eliminates rainbows in subproblems we are interested in, as we will see later. Clearly, if LP (2) is not feasible, we know that the reference problem has no feasible solution.

In all what follows, we only work on edges in $\text{supp}(x^*)$. Therefore, to simplify the exposition, we assume from now on that $E = \text{supp}(x^*)$. This can easily be achieved by deleting all edges $e \in E$ with $x^*(e) = 0$ from G .

Our algorithm decomposes the problem of finding an $O(1)$ -approximate spanning tree $T \subseteq E$ into an independent family of a special type of spanning tree problem on rainbow-free graphs. To decompose the problem, we consider tight spanning tree constraints. More precisely, let $\mathcal{L} \subseteq 2^V$ be any maximal laminar family of vertex-sets corresponding to spanning tree constraints that are tight with respect to x^* . In other words, \mathcal{L} is maximal laminar family chosen from the sets $L \subseteq V$ satisfying $x^*(E[L]) = |L| - 1$, where, $E[L] \subseteq E$ denotes the set of edges with both endpoints in L . In particular, \mathcal{L} contains all singletons. We say that $L_2 \in \mathcal{L}$ is a child of $L_1 \in \mathcal{L}$ if $L_2 \subsetneq L_1$ and there is no set $L_3 \in \mathcal{L}$ with $L_2 \subsetneq L_3 \subsetneq L_1$. For $L \in \mathcal{L}$, we denote by $\mathcal{C}(L) \subset \mathcal{L}$ the set of all children of L . Notice that $\mathcal{C}(L)$ forms a partition of L .

To construct a spanning tree T in G we will determine for each $L \in \mathcal{L}$ a set of edges T_L in

$$E_L := E[L] \setminus (\cup_{C \in \mathcal{C}(L)} E[C]),$$

that form a spanning tree in the graph G_L obtained from the graph (L, E_L) by contracting all children of L . Hence, the vertex set of G_L is $\mathcal{C}(L)$, and an original edge $\{u, v\} \in E_L$ is simply interpreted as an edge between the two children $C_u, C_v \in \mathcal{C}(L)$ that contain u and v , respectively. For singletons $L \in \mathcal{L}$, we set $T_L = \emptyset$. One can easily observe that a family $\{T_L\}_{L \in \mathcal{L}}$ of spanning trees in $\{G_L\}_{L \in \mathcal{L}}$ leads to a spanning tree $T = \cup_{L \in \mathcal{L}} T_L$ in G . Constructing “good” spanning trees T_L in G_L , for each $L \in \mathcal{L}$, will be our independent subproblems. As we will argue more formally later, the main benefit of this division is that the edge set E_L used in the subproblem to find T_L does not contain any rainbows.




Fig. 1 The situation that makes (3) unattainable in general. Shown is a subproblem L with two children L_1 and L_2 ; for the chain constraints shown as dashed boxes, $x^*(\delta(S_i) \cap E_L) = \epsilon$. This can occur even though there must be at least one unit of weight connecting L_1 and L_2 .

Our goal is to define constraints that the spanning trees T_L have to satisfy, that allow us to conclude that the resulting spanning tree $T = \cup_{L \in \mathcal{L}} T_L$ does not violate the chain constraints by more than a constant factor.

One of the arguably most natural constraint families to impose would be to require that the contribution of T_L to any cut S_i is within a constant factor of the contribution of x^* on S_i when only considering edges in E_L , i.e.,

$$|T_L \cap \delta(S_i)| \leq O(x^*(\delta(S_i) \cap E_L)). \quad (3)$$

If the above inequality holds for each $L \in \mathcal{L}$ and $i \in [n-1]$, then the final spanning tree T will indeed not violate any chain constraint by more than a constant factor: it suffices to sum up the inequalities for a fixed i over all sets L and observe that $\{T_L\}_{L \in \mathcal{L}}$ partitions T , and $\{E_L\}_{L \in \mathcal{L}}$ is a partition of E_L :

$$\begin{aligned} |T \cap \delta(S_i)| &= \sum_{L \in \mathcal{L}} |T_L \cap \delta(S_i)| \leq O\left(\sum_{L \in \mathcal{L}} x^*(\delta(S_i) \cap E_L)\right) \\ &= O(x^*(\delta(S_i))) = O(1)b_i. \end{aligned} \quad (4)$$

Unfortunately, it turns out that it is in general impossible to find spanning trees T_L that satisfy (3). This is because there can be many constraints S_i for which $x^*(\delta(S_i) \cap E_L) = o(1)$, in a setting where one has to include at least one edge in T_L that crosses one of these constraints; see Fig. 1.

We therefore introduce a weaker condition on T_L . For $L \in \mathcal{L}$ and $i \in [n-1]$, let $\mathcal{C}_i(L) \subseteq \mathcal{C}(L)$ be the family of all children $C \in \mathcal{C}(L)$ of L that cross the cut S_i , i.e., $S_i \cap C \neq \emptyset$ and $L \setminus S_i \neq \emptyset$. We want the sets T_L to satisfy the following:

$$|T_L \cap \delta(S_i)| \leq 7 \cdot x^*(\delta(S_i) \cap E_L) + 2 \cdot \mathbf{1}_{\{|\mathcal{C}_i(L)| \geq 2\}} \quad \forall i \in [n-1]. \quad (5)$$

Here, $\mathbf{1}_{\{|\mathcal{C}_i(L)| \geq 2\}}$ is the indicator that is equal to 1 if $|\mathcal{C}_i(L)| \geq 2$ and 0 otherwise.

We first show in Section 2.1 that satisfying the above condition indeed leads to a good spanning tree T .

Theorem 3 *For $L \in \mathcal{L}$, let T_L be a spanning tree in G_L that satisfies (5). Then $T = \cup_{L \in \mathcal{L}} T_L$ is a spanning tree in G satisfying*

$$|T \cap \delta(S_i)| \leq 9x^*(\delta(S_i)) \leq 9b_i \quad i \in [n-1].$$

We then show in Section 2.2 that such spanning trees can indeed be found efficiently.

Theorem 4 *For each $L \in \mathcal{L}$, we can efficiently find a spanning tree T_L in G_L satisfying (5).*

Combining the above two theorems immediately leads to an efficient algorithm to find a spanning tree in G that violates each chain constraint by at most a factor of 9 whenever LP (2) is feasible, and thus proves Theorem 1. For convenience, a summary of our algorithm is provided below.

Algorithm to find spanning tree T that violates chain constraints by a factor of at most 9.

1. Compute an optimal solution x^* to the linear program (2).
2. Independently for each $L \in \mathcal{L}$, invoke Theorem 4 to obtain a spanning tree T_L in G_L satisfying (5).
3. Return $T = \cup_{L \in \mathcal{L}} T_L$.

2.1 Analysis of algorithm (proof of Theorem 3)

For each $L \in \mathcal{L}$, let T_L be a spanning tree in G_L that satisfies (5), let $T = \cup_{L \in \mathcal{L}} T_L$, and let $i \in [n-1]$. Using the same reasoning as in (4) we can bound the load on chain constraint i as follows:

$$\begin{aligned} |T \cap \delta(S_i)| &= \sum_{L \in \mathcal{L}} |T_L \cap \delta(S_i)| \stackrel{(5)}{\leq} 7 \sum_{L \in \mathcal{L}} x^*(\delta(S_i) \cap E_L) + 2 \sum_{L \in \mathcal{L}} \mathbf{1}_{\{|\mathcal{C}_i(L)| \geq 2\}} \\ &= 7x^*(\delta(S_i)) + 2 \sum_{L \in \mathcal{L}} \mathbf{1}_{\{|\mathcal{C}_i(L)| \geq 2\}}, \end{aligned}$$

using the fact that $\{E_L\}_{L \in \mathcal{L}}$ partitions E . To prove Theorem 3, it thus suffices to show

$$\sum_{L \in \mathcal{L}} \mathbf{1}_{\{|\mathcal{C}_i(L)| \geq 2\}} \leq x^*(\delta(S_i)), \quad (6)$$

which then implies

$$|T \cap \delta(S_i)| \leq 9x^*(\delta(S_i)) \leq 9b_i,$$

where the last inequality follows from x^* being feasible for (2). We complete the analysis by showing the following result, which is a stronger version of (6).

Lemma 5

$$\sum_{L \in \mathcal{L}} (|\mathcal{C}_i(L)| - 1)^+ \leq x^*(\delta(S_i)),$$

where $(\cdot)^+ = \max(0, \cdot)$.

Proof Let $\mathcal{L}_i \subseteq \mathcal{L}$ be the family of all sets in \mathcal{L} that cross S_i , and let $\mathcal{L}_i^{\min} \subseteq \mathcal{L}_i$ be all minimal sets of \mathcal{L}_i . We will show that

$$\sum_{L \in \mathcal{L}} (|\mathcal{C}_i(L)| - 1)^+ = |\mathcal{L}_i^{\min}| - 1. \quad (7)$$

Let us first see that this is a strengthening of the lemma. Since all sets $W \in \mathcal{L}_i$ correspond to tight spanning tree constraints with respect to x^* , we have that the restriction $x^*|_{E[W]}$ of x^* to the edges in the graph $G[W]$ is a point in the spanning tree polytope of $G[W]$. In particular, at least one unit of $x^*|_{E[W]}$ crosses any cut in $G[W]$. Since $W \in \mathcal{L}_i$, the set S_i induces a cut $(S_i \cap W, W \setminus S_i)$ in $G[W]$. Hence

$$x^*(\delta(S_i) \cap E[W]) \geq 1 \quad \forall W \in \mathcal{L}_i.$$

Now observe that due to minimality of the sets in \mathcal{L}_i^{\min} , all sets in \mathcal{L}_i^{\min} are disjoint. Thus

$$x^*(\delta(S_i)) \geq \sum_{W \in \mathcal{L}_i^{\min}} x^*(\delta(S_i) \cap E[W]) \geq |\mathcal{L}_i^{\min}|,$$

which, together with (7), implies Lemma 5. Hence, it remains to show (7).

Let $\mathcal{L}_i^{\text{nm}} = \mathcal{L}_i \setminus \mathcal{L}_i^{\min}$ be all sets in \mathcal{L}_i that are not minimal. Notice that only sets $L \in \mathcal{L}^{\text{nm}}$ can have a strictly positive contribution to the left-hand side of (7) since these are precisely the sets $L \in \mathcal{L}$ with $|\mathcal{C}_i(L)| \geq 1$: for any other set $L \in \mathcal{L}$, either (i) $L \notin \mathcal{L}_i$, in which case none of its children can cross S_i since not even L crosses S_i , or (ii) $L \in \mathcal{L}_i^{\min}$, in which case we again get $|\mathcal{C}_i(L)| = 0$ since L has no children in \mathcal{L}_i due to minimality. We thus obtain

$$\sum_{L \in \mathcal{L}} (|\mathcal{C}_i(L)| - 1)^+ = \sum_{L \in \mathcal{L}_i^{\text{nm}}} (|\mathcal{C}_i(L)| - 1). \quad (8)$$

Observe that $\sum_{L \in \mathcal{L}_i^{\text{nm}}} |\mathcal{C}_i(L)|$ counts each set in \mathcal{L}_i precisely once, except for the set $V \in \mathcal{L}_i$ which is the only set in \mathcal{L}_i that is not a child of some other set in \mathcal{L}_i . Hence

$$\sum_{L \in \mathcal{L}_i^{\text{nm}}} |\mathcal{C}_i(L)| = |\mathcal{L}_i| - 1. \quad (9)$$

Finally, combining (8) with (9) we obtain

$$\sum_{L \in \mathcal{L}} (|\mathcal{C}_i(L)| - 1)^+ = \sum_{L \in \mathcal{L}_i^{\text{nm}}} (|\mathcal{C}_i(L)| - 1) = |\mathcal{L}_i| - 1 - |\mathcal{L}_i^{\text{nm}}| = |\mathcal{L}_i^{\min}| - 1,$$

thus proving (7). \square

2.2 Main step of algorithm (proof of Theorem 4)

Let $L \in \mathcal{L}$. We now consider the problem of finding a spanning tree T_L in G_L that satisfies (5). Recall that G_L is obtained from the graph (L, E_L) by contracting all children of L . For simplicity, we again interpret an edge $\{v_i, v_j\} \in E_L$ as an edge in G_L between the two vertices corresponding to the sets $C_i, C_j \in \mathcal{L}$ that contain v_i and v_j , respectively.

We start by showing that there are no rainbows in E_L , which is a crucial assumption in the algorithm to be presented in the following.

Lemma 6 *For $L \in \mathcal{L}$, E_L does not contain any rainbows.*

Proof Since \mathcal{L} is a maximal laminar family of tight spanning tree constraints in G with respect to x^* , the spanning tree constraints imposed by the sets in \mathcal{L} define the minimal face of P_{ST} on which x^* lies (this is a well-known result that can be proven through combinatorial uncrossing, see e.g. [10]). From this we can conclude that any other tight spanning tree constraint $x^*(E[W]) = |W| - 1$, for some set $W \subseteq V, W \notin \mathcal{L}$, is implied by the spanning tree constraints given by \mathcal{L} .

Now assume by the sake of contradiction that there are two edges $e_1, e_2 \in E_L$ that form a rainbow, and let e_1 be the edge contained in strictly more chain constraints than e_2 . We will argue that one could slightly increase the fractional value $x^*(e_2)$ by some small value $\delta > 0$ and decrease $x^*(e_1)$ by δ to get a new feasible solution $y^\delta = x^* + \delta \cdot \chi(e_2) - \delta \cdot \chi(e_1)$ to (2) with strictly smaller objective value than x^* , thus violating that x^* is an optimal solution of (2). Clearly, since x^* does not violate any chain constraint, also y does not violate any chain constraint. Furthermore y^δ , for any $\delta > 0$, has indeed a lower objective value than x^* . The only point that remains to show is that there is a small $\delta > 0$ such that $y^\delta \in P_{ST}$. Since all tight spanning tree constraints are implied by the constraints that correspond to sets in \mathcal{L} , it suffices to check that y^δ does not violate any of the spanning tree constraints induced by \mathcal{L} . This indeed holds since any set $W \in \mathcal{L}$ either satisfies $e_1, e_2 \in E[W]$ —this happens if W is a set containing L —or $e_1, e_2 \notin E[W]$, which happens for any other set $W \in \mathcal{L}$. Hence, for any $W \in \mathcal{L}$ and $\delta > 0$, we have

$$y^\delta(E[W]) = x^*(E[W]) = |W| - 1.$$

Thus, none of the spanning tree constraints that are tight with respect to x^* will be violated by y^δ for any $\delta > 0$. Hence, by choosing a sufficiently small $\delta > 0$ that makes sure that $y^\delta \geq 0$, and that no other (non-tight) spanning tree constraints are violated, we obtain $y^\delta \in P_{ST}$. \square

We classify chain constraints S_i into two types, depending on the right-hand side of (5). Call a cut S_i *bad* if one can include at most one edge that crosses S_i in T_L without violating (5), i.e.,

$$7x^*(\delta(S_i) \cap E_L) + 2 \cdot \mathbf{1}_{\{|C_i(L)| \geq 2\}} < 2.$$

Otherwise, call the cut S_i *good*. Notice that for a cut S_i to be bad, we need to have $|\mathcal{C}_i(L)| = 1$ because of the following. Clearly, if $|\mathcal{C}_i(L)| \geq 2$, then S_i cannot be bad due to the term $2 \cdot \mathbf{1}_{\{|\mathcal{C}_i(L)| \geq 2\}}$. If $|\mathcal{C}_i(L)| = 0$, then we use the fact that all edges in $E[L]$ that cross S_i are part of E_L , hence

$$x^*(\delta(S_i) \cap E_L) = x^*(\delta(S_i) \cap E[L]) \geq 1,$$

where the last inequality follows from the fact that $x^*|_{E[L]}$ is in the spanning tree polytope of the graph $(L, E[L])$. Hence a cut S_i is bad if and only if the following two conditions hold simultaneously:

1. $|\mathcal{C}_i(L)| = 1$,
2. $x^*(\delta(S_i) \cap E_L) < \frac{2}{7}$.

An edge $e \in E_L$ is called *bad* if e crosses at least one bad cut S_i , otherwise it is called *good*. We denote by $A_L \subseteq E_L$ the sets of all good edges.

The procedure we use to find a tree T_L satisfying (5) constructs a tree T_L that consists of only good edges, i.e., $T_L \subseteq A_L$. We determine T_L using a matroid intersection problem that asks to find a spanning tree in G_L satisfying an additional partition matroid constraint.

To define the partition matroid we first number the edges $A_L = \{e_1, \dots, e_k\}$ as follows. For $e \in A_L$, let $\alpha(e) < \beta(e)$ be the lower and higher index of the two endpoints of e , hence, $e = \{v_{\alpha(e)}, v_{\beta(e)}\}$. (Notice that $\alpha(e) = \beta(e)$ is not possible since $x^*(e) > 0 \forall e \in E$ and $x^* \in P_{ST}$.) The edges $e \in A_L$ are numbered lexicographically, first by increasing value of $\alpha(e)$ and then by increasing value of $\beta(e)$, i.e., for any $p \in [k-1]$ either $\alpha(e_p) < \alpha(e_{p+1})$, or $\alpha(e_p) = \alpha(e_{p+1})$ and $\beta(e_p) \leq \beta(e_{p+1})$. Ideally, we would like to group the edges in A_L into consecutive blocks $\{e_p, e_{p+1}, \dots, e_q\}$ each having a total weight of exactly $x^*(\{e_p, \dots, e_q\}) = 3/7$. Since this is in general not possible, we will split some of the edges by creating two parallel copies. More precisely, to define the first set P_1 of our partition, let $p \in [k]$ the largest index for which $x^*(\{e_1, \dots, e_p\}) \leq 3/7$. If $x^*(\{e_1, \dots, e_p\}) = 3/7$ then $P_1 = \{e_1, \dots, e_p\}$. Otherwise, we replace the edge e_{p+1} by two parallel copies e'_{p+1}, e''_{p+1} of e_{p+1} , and we distribute the weight of $x^*(e_{p+1})$ on e'_{p+1}, e''_{p+1} as follows:

$$\begin{aligned} x^*(e'_{p+1}) &= \frac{3}{7} - x^*(\{e_1, \dots, e_p\}), \\ x^*(e''_{p+1}) &= x^*(e_{p+1}) - x^*(e'_{p+1}). \end{aligned}$$

This splitting operation does not violate any previous assumptions: the weight x^* on the new edge set $\{e_1, \dots, e_p, e'_{p+1}, e''_{p+1}, e_{p+2}, \dots, e_k\}$ is still a point in the spanning tree polytope of the graph over the vertices $\mathcal{C}(L)$ with the new edge set. By applying this splitting operation whenever necessary, we can assume that $A_L = \{e_1, \dots, e_k\}$ can be partitioned into sets $P_1 = \{e_1, \dots, e_{p_1}\}$, $P_2 = \{e_{p_1+1}, \dots, e_{p_2}\}$, \dots , $P_s = \{e_{p_{s-1}+1}, \dots, e_k\}$ satisfying:

- (i) $x^*(P_h) = 3/7 \quad \forall h \in [s-1]$,
- (ii) $x^*(P_s) \leq 3/7$.

Using this partition we define a unitary partition matroid $M = (A_L, \mathcal{I})$ on the good edges A_L , with independent sets

$$\mathcal{I} = \{U \subseteq A_L \mid |U \cap P_h| \leq 1 \forall h \in [s]\}.$$

The tree spanning T_L in G_L that our algorithm selects is any spanning tree $T_L \subseteq A_L$ in G_L that is independent in the partition matroid M . Notice that if there exists a spanning tree in G_L that is independent in M , then such a spanning tree can be found in polynomial time by standard matroid intersection techniques (see [17, Volume B] for more details about matroids in general and techniques to find common independent sets in the intersection of two matroids). Hence to complete the description and analysis of our algorithm, all that remains is to show the existence of a spanning tree in G_L that is independent in M , and that such a spanning tree satisfies (5). We address these two statements in the following.

The theorem below shows the feasibility of the matroid intersection problem.

Theorem 7 *There exists a spanning tree $T_L \subseteq A_L$ in G_L that is independent in M , i.e., $T_L \in \mathcal{I}$.*

Proof Let $y \in [0, 1]^{A_L}$ be defined by $y(e) = \frac{7}{3}x^*(e)$ for $e \in A_L$. To prove the theorem, we show that y is simultaneously in the matroid polytope of M and in the dominant of the spanning tree polytope of G_L^A , where G_L^A is the graph obtained from G_L by deleting all bad edges. This implies that the intersection of $P_{ST}(G_L^A)$ and the matroid polytope P_M of M is nonempty. Since $P_{ST}(G_L^A) \cap P_M$ is a face of the matroid intersection polytope corresponding to intersecting the matroid M with the graphic matroid on G_L^A , it is therefore integral [17]. Hence, if $P_{ST}(G_L^A) \cap P_M$ is nonempty, it contains an integral point, and this corresponds to a spanning tree that is independent in M .

The vector y is clearly in the matroid polytope of the partition matroid M , since for any partition P_h with $h \in [s]$ we have $y(P_h) = \frac{7}{3}x^*(P_h) \leq 1$.

To show that y is in the dominant of the spanning tree polytope of G_L^A , we first discuss some structural properties of G_L^A that allow us to decompose the problem. Let $S_{i_1}, S_{i_2}, \dots, S_{i_p} \in \mathcal{S}$ be all bad cuts, where $1 \leq i_1 < \dots < i_p \leq \ell$. For $j \in [p]$, let $C_j \in \mathcal{C}(L)$ be the child that crosses S_{i_j} . Notice that since S_{i_j} is bad, there is indeed precisely one child of L that crosses S_{i_j} , and furthermore, there are no good edges crossing S_{i_j} . Hence, every C_j for $j \in [p]$ is a cut vertex in G_L^A , whose removal splits G_L^A into a part where all vertices are contained in S_{i_j} and a part where all vertices are outside of S_{i_j} .

We define the following vertex sets of G_L^A , which correspond to the vertex sets between those cut vertices, including the cut vertices themselves:

$$\begin{aligned} \mathcal{B}^1 &= \{C \in \mathcal{C}(L) \mid C \cap S_{i_1} \neq \emptyset\}, \\ \mathcal{B}^j &= \{C \in \mathcal{C}(L) \mid C \cap (S_{i_j} \setminus S_{i_{j-1}}) \neq \emptyset\} \quad \forall j \in \{2, \dots, p\}, \\ \mathcal{B}^{p+1} &= \{C \in \mathcal{C}(L) \mid C \cap (V \setminus S_{i_p}) \neq \emptyset\}. \end{aligned}$$

Hence, the above sets contain all vertices of G_L^A precisely once except for the cut vertices, which appear in at least two sets.

For $j \in [p+1]$, let $G_L^j = G_L^A[\mathcal{B}^j]$, i.e., G_L^j is the induced subgraph of G_L^A over the vertices \mathcal{B}^j . Also let A_L^j denote the edge set of G_L^j . To show that y is in the dominant of the spanning tree polytope of G_L^A , we show that the restriction of y to the edges of any of the graphs G_L^j for $j \in [p+1]$ is in the dominant of the spanning tree polytope of G_L^j .

Fix $j \in [p+1]$. Let $y^j = y|_{A_L^j}$ be the restriction of y to the edges in G_L^j . To show that y^j is in the dominant of the spanning tree polytope of G_L^j , we use the following well-known partition-based description of the dominant of the spanning tree polytope (see [17]). Let $\mathcal{W} = \{W_1, \dots, W_q\}$ be a partition of the sets in \mathcal{B}^j , i.e., the vertices of G_L^j , and we denote by $A_L^j(\mathcal{W})$ all edges of A_L^j with endpoints in two different sets of the partition \mathcal{W} . To show that y^j is in the dominant of the spanning tree polytope of G_L^j , we have to show that the following inequality holds for any partition $\mathcal{W} = \{W_1, \dots, W_q\}$ of \mathcal{B}^j :

$$y(A_L^j(\mathcal{W})) \geq q - 1. \quad (10)$$

Given a partition $\mathcal{W} = \{W_1, \dots, W_q\}$ of \mathcal{B}^j with $q \geq 2$, we define a partition $\mathcal{Z} = \{Z_1, \dots, Z_q\}$ of the set $\mathcal{C}(L)$ as follows. We start by setting $Z_r = W_r$ for $r \in [q]$. If $j \in \{2, \dots, p+1\}$, we add $\cup_{s=1}^{j-1} \mathcal{B}^s$ to the set $Z \in \mathcal{Z}$ that contains C_{j-1} . Additionally, if $j \in [p]$, we add $\cup_{s=j+1}^{p+1} \mathcal{B}^s$ to the set $Z \in \mathcal{Z}$ that contains C_j . Hence, \mathcal{Z} is identical to \mathcal{W} with the possible exception of up to two sets. Let $E_L(\mathcal{Z})$ be the set of all edges in E_L that cross the partition \mathcal{Z} . Notice that $E_L(\mathcal{Z})$ consists of all edges in $A_L^j(\mathcal{W})$ together with all bad edges that cross either of the bad cuts C_{j-1} or C_j . Since the x^* -weight of the set of all edges crossing any bad cut is bounded by $\frac{2}{7}$ we obtain

$$x^*(E_L(\mathcal{Z})) \leq x^*(A_L^j(\mathcal{W})) + 2 \cdot \frac{2}{7}. \quad (11)$$

Furthermore, since x^* is in the spanning tree polytope of G_L , it fulfills the partition-constraints that define the dominant of the spanning tree polytope of G_L . For the partition \mathcal{Z} this leads to

$$x^*(E_L(\mathcal{Z})) \geq q - 1. \quad (12)$$

Using the definition of y and combining (11) and (12) we obtain

$$y(A_L^j(\mathcal{W})) = \frac{7}{3}x^*(A_L^j(\mathcal{W})) \stackrel{(11)}{\geq} \frac{7}{3}x^*(E_L(\mathcal{Z})) - \frac{4}{3} \stackrel{(12)}{\geq} \frac{7}{3}(q-1) - \frac{4}{3} \geq q-1,$$

since $q \geq 2$. This shows (10) and therefore completes the proof. \square

The following theorem finishes the analysis of our algorithm.

Theorem 8 *Let $T_L \subseteq A_L$ be a spanning tree in G_L that is independent in M . Then T_L satisfies (5).*

Proof Consider a cut $S_i \in \mathcal{S}$ for some fixed $i \in [n-1]$. We consider the partition P_1, \dots, P_s of A_L used to define the partition matroid M . We are interested in all sets in this partition that contain edges crossing S_i . The definition of the partition P_1, \dots, P_s , together with the fact that A_L has no rainbows, implies that the sets of the partition containing edges crossing S_i are consecutively numbered P_a, P_{a+1}, \dots, P_b , for some $1 \leq a \leq b \leq s$. Since T_L contains at most one edge in each partition, we have

$$|T_L \cap \delta(S_i)| \leq b - a + 1. \quad (13)$$

We first consider the case $b - a \geq 2$. Notice that all edges in any set P_h for $a < h < b$ cross S_i . Hence,

$$x^*(\delta(S_i) \cap E_L) \geq \sum_{h=a+1}^{b-1} x^*(P_h) = (b - a - 1) \cdot \frac{3}{7},$$

where we used $x^*(P_h) = \frac{3}{7}$ for $1 \leq h \leq s - 1$. Combining the above inequality with (13), and using that $b - a \geq 2$ in the second inequality, we obtain that

$$|T_L \cap \delta(S_i)| \leq b - a + 1 \leq 3(b - a - 1) \leq 7x^*(\delta(S_i) \cap E_L).$$

Thus T_L satisfies (5).

Assume now $b - a \leq 1$. If S_i is bad, then $|T_L \cap \delta(S_i)| = 0$ since T_L only contains good edges and no good edge crosses any bad cut. Hence, T_L trivially satisfies (5). So assume that S_i is good, i.e., either $|\mathcal{C}(L)| \geq 2$ or $x^*(\delta(S_i) \cap E_L) \geq \frac{2}{7}$. If $|\mathcal{C}(L)| \geq 2$, then beginning again from (13) we have

$$|T_L \cap \delta(S_i)| \leq b - a + 1 \leq 2 = 2 \cdot \mathbf{1}_{|\mathcal{C}(L)| \geq 2}.$$

Otherwise, if $x^*(\delta(S_i) \cap E_L) \geq \frac{2}{7}$, then

$$|T_L \cap \delta(S_i)| \leq 2 \leq 7x^*(\delta(S_i) \cap E_L).$$

Either way, T_L satisfies (5). \square

3 Hardness and integrality gaps

In this section, we provide the proof of Theorem 2.

3.1 The chain-constrained partition problem

We will begin by considering a different problem, where we replace the spanning tree constraint by a unitary partition matroid. We will show integrality gaps and hardness for this problem first, and then show how this can be leveraged to the spanning tree setting via a gadget.

So consider the following problem, that we call the *chain-constrained partition problem*. We are given a graph $G = (V, E)$ and a chain $S_1 \subsetneq S_2 \subsetneq \dots \subsetneq S_\ell$




Fig. 2 First step of the reduction to the chain-constrained partition problem. Edges E' are shown as wavy; the dashed sets are the chain constraints.

of cuts, with associated upper bounds b_i for $i \in [\ell]$. We are also given a partition $\{B_1, B_2, \dots, B_q\}$ of E . The goal is to pick precisely one edge from each part, while satisfying the chain constraints. Without loss of generality, we can take E to be a matching (by splitting vertices as needed); we will always assume this in what follows.

The problem has a completely equivalent formulation as follows. We are given a $\ell \times m$ consecutive-ones matrix A (i.e., entries of A are 0-1, and in any column of A , the 1's are all consecutive) and an integral right-hand side vector b of length ℓ . We are also given a partition $\mathcal{B} = \{B_1, \dots, B_q\}$ of $[m]$. The goal is to pick one column from each part so that the sum of the chosen columns does not exceed b , if such a choice exists; or in other words, to find a vector $x \in \{0, 1\}^m$ such that $Ax \leq b$ and for $i \in [q]$ there is precisely one index $j \in B_i$ with $x_j = 1$. The correspondence to the other formulation is that each row corresponds to a chain constraint, and each column of 1's to an edge.

Relation to the chain-constrained spanning tree problem. Here we will show that additive integrality gaps and hardness results for the chain-constrained partition problem transfer to the chain-constrained spanning tree problem. Note that only *additive* results will transfer; our gadgets will require increasing the right hand side, and constant-factor multiplicative hardness (which is an easy consequence for the partition version) will not carry over.

So let an instance of the chain-constrained partition problem be given, with $G = (V, E)$ being the graph. For each edge e , let $\alpha(e)$ be the innermost endpoint of edge e (with respect to the chain) and $\beta(e)$ the outermost endpoint. Now construct the set of edges E' as the union of an arbitrary spanning tree on $\{\alpha(e) \mid e \in E\}$, and for each $i \in [q]$, an arbitrary spanning tree on $\{\beta(e) \mid e \in B_i\}$ (see Fig. 2). Keep the same set of chain constraints, but modify the upper bounds by setting $b'_i = b_i + |\delta(S_i) \cap E'|$.

Now consider the chain-constrained spanning tree problem on $G' = (V, E \cup E')$, but subject to the extra restriction that *every edge in E' must be picked*.




Fig. 3 Edge splitting ensures that there is no advantage to not taking all edges in E'' .

This is clearly precisely the same problem as the original chain-constrained partition problem. In order to eliminate this extra restriction, we need one further trick. Let $s = |E'| + 1$. For each edge $e = \{v_i, v_j\} \in E'$, with $i < j$, we split e into a path of length $s\ell$. Let $v_e^{(r)}$ denote the vertices of this path in order, with $v_e^{(0)} = v_i$ and $v_e^{(s\ell)} = v_j$. Call the resulting graph G'' , and let E'' be the set of edges which replaced E' . Note that for each edge in E' , all but at most one of the edges in the subsequent path must be chosen in any spanning tree of G'' . We now define, for $i \in [\ell]$ and $j \in [s]$,

$$S_i^{(j)} = S_i \cup \{v_e^{(k)} : e \in E'[S_i], k \in [s\ell]\} \cup \{v_e^{(k)} : e \in \delta_{G'}(S_i), k < is + j\}.$$

The sets $S_i^{(j)}$ are chosen so that

- (i) they form a chain;
- (ii) for any spanning tree in G'' , and for any $i \in [\ell]$, there is some $j \in [s]$ so that every edge of $\delta(S_i^{(j)}) \cap E''$ is in the spanning tree; and
- (iii) $\delta(S_i^{(j)}) \cap E = \delta(S_i) \cap E$ for all $i \in [\ell], j \in [s]$.

We prescribe an upper bound of b'_i for set $S_i^{(j)}$, for all $i \in [\ell], j \in [s]$. See Fig. 3.

Now consider any cut S_i in G' , and observe that for any spanning tree T of G'' , there must be some $j \in [s]$ so that $\delta(S_i^{(j)}) \cap E'' \subseteq T$. Observe that this construction ensures that if T is a spanning tree of G'' respecting these new chain constraints, then any spanning tree T' such that $T' \supseteq E''$ and $T' \cap E \subseteq T \cap E$ also satisfies the constraints. This concludes the reduction from the chain-constrained partition problem.

Note that the size of the instance G'' is polynomial in the size of the starting partition problem instance. If the starting instance has m edges and ℓ constraints, then the final instance will have $O(m^2\ell)$ edges and $O(m\ell^2)$ constraints.




Fig. 4 Reduction from independent set to a column selection problem; here $k = 2$.

3.2 NP-hardness

We will now consider only the chain-constrained partition problem from this point, and we will use the equivalent consecutive-ones formulation discussed above. Our first modest goal will be to show that it is NP-complete to decide whether there is a solution that does not violate any constraints at all.

First note that the same problem, but without the requirement that the ones in any column of A must be consecutive, is clearly hard, even if all parts of the partition have size 2. For example, we could reduce from the independent set problem, as shown in Fig. 4. Let $H = (W, E)$ be a given graph, the question being if H has an independent set of size k . Let $t = |W|$, and w_1, \dots, w_t a labelling of W . The partition is $\{\{2i - 1, 2i\} \mid i \in [t]\}$; if column $2i$ is picked, this will represent that vertex i is chosen in the independent set. The first row of C contains a 1 in all columns of odd index, and $b_1 = t - k$; this ensures that at least k vertices are chosen. Each other row, one per edge $\{w_i, w_j\} \in E$, represents the constraint that at most one of w_i and w_j can be picked, and so contains a 1 in columns $2i$ and $2j$, with a right-hand side of 1.

So our goal will be to simulate this pairwise column-selection problem, where there is no consecutive-ones restriction. We describe the reduction for a matrix C with two columns; the extension to more columns will be clear. Let r_i be the number of ones in column i , for $i \in \{1, 2\}$, and let $r' = r_1 + r_2$. First expand out the matrix C horizontally, by replacing column i by r_i columns, each with only a single one. Call the resulting matrix \tilde{C} ; the first r_1 columns of \tilde{C} correspond to the first column of C , and the last r_2 columns to the second column of C . Let A_n denote the $n \times 1$ matrix consisting of all ones, and I_n the

$n \times n$ identity matrix. Then define the final matrix A by

$$\left(\begin{array}{c|c|c|c|c} & & & & \\ & \tilde{C} & & & \\ \hline & & I_{r_1} & & A_{r_1} \\ \hline & & & I_{r_2} & A_{r_2} \end{array} \right),$$

and choose the partition

$$\{\{i, r' + i\} \mid i \in [r']\} \cup \{\{2r' + 1, 2r' + 2\}\}.$$

For all new rows, i.e., the rows not corresponding to \tilde{C} , the packing constraint is chosen to be 1 (the constraints for all original rows are preserved). This means that if the final column is chosen, then all of the r_2 columns corresponding to the second column of C must be picked, and similarly if the penultimate column is chosen, all r_1 columns corresponding to the first column of C must be picked. This precisely mimics the choice of at least one of these columns in the original matrix C . Notice that the original problem asked to choose precisely one column, however, this modification is clearly equivalent to the original problem since nothing can be gained by choosing both columns out of a partition. Hence, this completes the reduction.

To extend this to further columns, we simply repeat this procedure on each pair of columns in C , creating separate new rows and columns. Once again, the blowup is only polynomial; if C is $\ell \times m$, then A will be $O(m\ell) \times O(m\ell)$.

3.3 Boosting to an additive $\Omega(\log n / \log \log n)$ hardness

First, consider the following simple integrality gap construction, which will motivate the hardness construction. Fix an integer k . We will construct a sequence of matrices A^1, \dots , with A^i having k^i rows and $\Theta(k^i)$ columns, as well as a partition \mathcal{B}^i of the columns of A^i , as follows. Let A^1 be a $k \times k$ identity matrix, and \mathcal{B}^1 be the trivial partition with only one part. We construct A^{i+1} inductively from A^i as demonstrated by the block matrix diagram of Fig. 5; the left k columns consist of a vertically stretched $k \times k$ identity matrix (i.e., the Kronecker product $A^1 \otimes \Lambda_k$, where recall Λ_k is a column vector of k ones), and the remaining columns consist of the Kronecker product $I_k \otimes A^k$. The partition \mathcal{B}^{i+1} contains one part consisting of the first k columns, and then a copy of \mathcal{B}^i for each set of columns corresponding to a copy of A^i .

Inductively, in the same fashion as in the integrality gap construction, we see that if we cannot find a solution to the starting problem that doesn't violate constraints at all, then we cannot find a solution to A^i that avoids violating some constraint by i units.

Notice that A^i has $k^i \leq (2m)^i = O(m)^i$ rows. Furthermore, the number of columns $\text{cols}(A^i)$ of A^i is equal to $m + k \cdot \text{cols}(A^{i-1})$. Since $\text{cols}(A^1) = m$, we obtain

$$\text{cols}(A^i) = m \sum_{j=0}^{i-1} k^j = m \frac{k^i - 1}{k - 1} = O(m)^i.$$

Choosing $i = m$, we obtain a matrix A^m with $O(m)^m$ columns and $O(m)^m$ rows. This we can reduce to a chain-constrained spanning tree problem on $m^{O(m)}$ nodes; writing m in terms of the number of nodes n , we find that $m = \Theta(\log n / \log \log n)$, and we obtain the required additive hardness of Theorem 2.

4 Conclusions

We would like to close with several interesting directions for future research. One very natural question is whether there is an $O(1)$ -approximation for laminar cut constraints; we believe this to be true. Although it seems non-trivial to directly generalize our procedure for the chain-constrained case to the laminar case, we hope that they can be useful in combination with insights from $O(1)$ -approximations for the degree-bounded case.

Another natural extension would be to a cost version of the problem, where edges are weighted and the goal is to return a spanning tree of minimum cost satisfying the chain constraints. The main reason our approach does not generalize easily to this setting is that we use a particular objective function to eliminate rainbows in the subproblems. Recently, Linhares and Swamy [14] have shown how to produce a spanning tree that violates all chain constraints by a constant multiplicative factor, and in addition has cost within a constant factor of the optimum (the optimum being the cost of a cheapest spanning tree that does not violate any of the constraints at all). Their algorithm relies explicitly on the results of this paper, in particular Theorem 3 and Theorem 4. It remains an open question to efficiently find a spanning tree that violates the chain constraints by a constant multiplicative factor and has cost no larger than the optimum.

References

1. N. Anari and S. Oveis Gharan. Effective-resistance-reducing flows and asymmetric TSP. ArXiv:1411.4613, 2014.
2. A. Asadpour, M. X. 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 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2010.

3. N. Bansal, R. Khandekar, J. Könemann, V. Nagarajan, and B. Peis. On generalizations of network design problems with degree bounds. *Mathematical Programming*, pages 1–28, 2012.
4. N. Bansal, R. Khandekar, and V. Nagarajan. Additive guarantees for degree-bounded directed network design. *SIAM Journal on Computing*, 39(4):1413–1431, 2009.
5. F. Bauer and A. Varma. Degree-constrained multicasting in point-to-point networks. In *Proceedings of the Fourteenth Annual Joint Conference of the IEEE Computer and Communication Societies (INFOCOM)*, pages 369–376, 1995.
6. K. Chaudhuri, S. Rao, S. Riesenfeld, and K. Talwar. A push-relabel approximation algorithm for approximating the minimum-degree MST problem and its generalization to matroids. *Theoretical Computer Science*, 410:4489–4503, 2009.
7. 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.
8. C. Chekuri, J. Vondrák, and R. Zenklusen. Dependent randomized rounding via exchange properties of combinatorial structures. In *Proceedings of the 51st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 575–584, 2010.
9. M. Fürer and B. Raghavachari. Approximating the minimum-degree Steiner Tree to within one of optimal. *Journal of Algorithms*, 17(3):409–423, 1994.
10. M. X. Goemans. Minimum bounded degree spanning trees. In *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 273–282, 2006.
11. K. Jain. A factor 2 approximation algorithm for the generalized Steiner Network Problem. *Combinatorica*, 21:39–60, 2001.
12. J. Könemann and R. Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. *SIAM Journal on Computing*, 31:1783–1793, 2002.
13. J. Könemann and R. Ravi. Primal-dual meets local search: approximating MST’s with nonuniform degree bounds. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 389–395, 2003.
14. A. Linhares and C. Swamy. Personal communication, 2015.
15. S. Oveis Gharan and A. Saberi. The asymmetric traveling salesman problem on graphs with bounded genus. In *Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 967–975, 2011.
16. R. Ravi, M. V. Marathe, S. S. Ravi, D. J. Rosenkrantz, and H. B. Hunt III. Approximation algorithms for degree-constrained minimum-cost network-design problems. *Algorithmica*, 31(1):58–78, 2001.
17. A. Schrijver. *Combinatorial Optimization, Polyhedra and Efficiency*. Springer, 2003.
18. M. Singh and L. C. Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 661–670, 2007.
19. R. Zenklusen. Matroidal degree-bounded minimum spanning trees. In *Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1512–1521, 2012.