

ON HOP-CONSTRAINED STEINER TREES IN TREE-LIKE METRICS*

MARTIN BÖHM[†], RUBEN HOEKSMAS[‡], NICOLE MEGOW[§], LUKAS NÖLKE[§], AND
BERTRAND SIMON[¶]

Abstract. We consider the problem of computing a Steiner tree of minimum cost under a hop constraint that requires the depth of the tree to be at most k . Our main result is an exact algorithm for metrics induced by graphs with bounded treewidth that runs in time $n^{O(k)}$. For the special case of a path, we give a simple algorithm that solves the problem in polynomial time, even if k is part of the input. The main result can be used to obtain, in quasi-polynomial time, a near-optimal solution that violates the k -hop constraint by at most one hop for more general metrics induced by graphs of bounded highway dimension and bounded doubling dimension. For nonmetric graphs, we rule out an $o(\log n)$ -approximation, assuming $P \neq NP$ even when relaxing the hop constraint by any additive constant.

Key words. k -hop Steiner tree, hop-constrained, dynamic programming, bounded treewidth

AMS subject classifications. 68Q25, 90C27, 05C12

DOI. 10.1137/21M1425487

1. Introduction. The *minimum-cost Steiner tree* problem is a fundamental network design problem: Given a set of terminals in a finite metric space, the task is to find a tree that spans all terminals and has minimum overall cost. While present in the original list of Karp’s 21 NP-complete problems [34], research on this problem goes back as far as the 1930s [32], and its origins can be traced back even further, as it is named after the 19th century Swiss mathematician Jakob Steiner. When the set of terminals to be interconnected is equal to the ground set of the metric space, we speak of the equally fundamental *minimum-cost spanning tree problem*, which is solvable in polynomial time. Both problems have had major applications in the 20th and 21st centuries, particularly in the design of transportation and communication networks. For a comprehensive discussion of the background on these two problems, we refer the reader to the book [31] and the historical treatise [25].

When considering the efficiency and reliability of a network, it is a common and natural requirement that vertices are not just simply connected, but rather connected with a path that consists of only a few edges. In the literature on network design, this requirement is known as *bounded hop distance*, where *hop* refers to an edge and *hop distance* to the number of edges on a path. A restriction on hop distances aims at reducing transmission delays and packet loss, avoiding the flooding of a network when routing, and increasing reliability of networks by limiting the amplifying effect of link

*Received by the editors June 9, 2021; accepted for publication (in revised form) November 14, 2021; published electronically June 1, 2022.

<https://doi.org/10.1137/21M1425487>

Funding: Partially supported by the German Science Foundation (DFG) under contracts ME 3825/1 and 146371743, TRR 89 Invasive Computing.

[†]Institute of Computer Science, University of Wrocław, 50-383 Wrocław, Poland (boehm@cs.uni.wroc.pl).

[‡]Department of Applied Mathematics, University of Twente, Enschede 7500 AE, The Netherlands (r.p.hoeksma@utwente.nl).

[§]Department of Mathematics and Computer Science, University of Bremen, 28350 Bremen, Germany (nicole.megow@uni-bremen.de, noelke@uni-bremen.de).

[¶]IN2P3 Computing Center, CNRS, 69100 Villeurbanne, France (bertrand.simon@cc.in2p3.fr).

failures. There exists a multitude of applications; see, e.g., [9, 14, 18, 23, 24, 28, 41, 43].

Despite their practical relevance, adding hop constraints makes network design problems substantially harder. The minimum-cost spanning tree problem, for example, is well known to be polynomial-time solvable, whereas its hop-constrained variants do admit constant lower bounds on the approximation ratio [4, 26] in certain metrics. However, network design problems (without hop constraints) often become easier when the underlying metric can be represented as a tree or is somewhat close to a tree such as graphs with bounded treewidth. In this paper, we investigate whether this is the case also in presence of hop constraints.

We now formally define the Steiner tree problem with bounded hop distance. Here, we are given a finite metric space (V, d) with a set V of n points as well as a distance function $d : V \times V \rightarrow \mathbb{Q}_+$, a set of terminals $\mathcal{X} \subseteq V$, a root $r \in \mathcal{X}$, and an integer $k \geq 1$. A k -hop Steiner tree is said to be a tree $\check{S} = (V_{\check{S}}, E_{\check{S}})$ rooted at r that spans all points in \mathcal{X} and has a depth of at most k . That is, $\mathcal{X} \subseteq V_{\check{S}} \subseteq V$, and for $v \in V_{\check{S}}$, the number of edges in the r - v path in \check{S} is at most k . The cost of a Steiner tree refers to the sum of its edge costs, $\sum_{\{u,v\} \in E_{\check{S}}} d(u,v)$, with edge costs given by d . We consider the *minimum-cost k -hop Steiner tree problem* (k -hop MŠT problem¹) that asks for a k -hop Steiner tree of minimum cost. When $\mathcal{X} = V$, this problem is equivalent to the *minimum-cost k -hop spanning tree (k -hop MST) problem*.

In this work, we show how to solve the k -hop MŠT problem in certain tree-like metrics. That is, we consider metrics which are represented by graphs from certain tree-like graph classes using the natural correspondence between metric spaces and weighted complete graphs via the shortest path metric. We say a weighted graph $G = (V, E)$ with a weight function $d : E \rightarrow \mathbb{R}^+$ induces a metric (V, d) if for any two vertices $u, v \in V$ the length of the shortest u - v path in G equals $d(u, v)$. While, in this paper, we do differentiate results based on the structure of the graph G , we do not do so based on the weight function d . Therefore, whenever a graph G is said to induce a metric, this implies that G has an accompanying arbitrary weight function d . A metric is called a *tree (resp., path) metric* if there is a tree (resp., path) inducing it, and it is called a *metric with bounded treewidth* if it is induced by some graph with bounded treewidth. For a given metric, it can be decided in polynomial time whether it is a path metric, a tree metric, or a metric with constant treewidth ω ; details are outlined in section 2. For convenience, we may not always distinguish between a metric and the graph inducing it.

A particular tree metric was studied by Althaus et al. [5]. They design an optimal algorithm tailored to metrics that can be represented as a *hierarchically separated tree* (HST). An HST is a tree with a very regular cost structure, where the costs of the edges on any root-to-leaf path are geometrically decreasing, and where the points of the metric space appear as leaves. In such a structure, subtrees of a solution tree can be described by subintervals of leaf indices and admit a dynamic program. However, general tree metrics exhibit a much more complex structure, and inner nodes of the tree must be handled with extreme care.

In this context, probabilistic tree embedding is a celebrated tool that forms the basis of many algorithms for network design problems in arbitrary metric spaces. Any metric can be approximated within a logarithmic factor by a distribution over trees, as was shown by Fakcharoenphol, Rao, and Talwar [21]. It is a common approach of net-

¹For brevity and as an homage to the work of Jarník and Kössler [32, 35], we use the Czech letter Š to distinguish Steiner trees from spanning trees in MŠT, respectively, MST. The pronunciation of Š is ⟨sh⟩, the same as the German pronunciation of the letter S in Steiner.

work design algorithms to first embed the given metric probabilistically in a tree, then solve the actual problem on this tree (optimally or approximately), and finally project the solution back into the original metric. In fact, this is the approach of Althaus et al. [5] to obtain an $O(\log n)$ -approximation algorithm for k -hop MŠT in general metrics. Notice that it is inherent to the approach that the approximation factor cannot be better than logarithmic. While the mentioned embedding schemes are not capable of preserving hop constraints, very recently, Haeupler, Hershkowitz, and Zuzic [29] proposed a new framework for approximating hop-constrained distances with partial tree metrics. It allows one to reduce the hop-constrained problem to the problem without hop constraints in a tree and admits polylogarithmic bicriteria-approximations with respect to both cost and hop distance. However, exact algorithms, constant-factor, and/or unicriteria approximation results seem out of reach for such methods.

1.1. Further related work. Hop-constrained problems have been studied since the 1980s. Various well-studied problems are in fact special cases of the k -hop MŠT problem, most notably the k -hop MST problem, where $\mathcal{X} = V$, the minimum Steiner tree problem, where $k \geq n - 1$, and the uncapacitated facility location problem, where $k = 2$. Hardness and inapproximability results for any of these problems are therefore valid for k -hop MŠT as well. In particular, k -hop MŠT is NP-hard [5], even for graph metrics, while the minimum Steiner tree problem is polynomial-time solvable on graphs with bounded treewidth [15]. It shall be mentioned that the minimum Steiner tree problem in arbitrary metrics can be solved by the classical dynamic programming-based Dreyfus–Wagner algorithm [19] with running time $O(3^t \cdot n)$, where $t = |\mathcal{X}|$; via fast subset convolution the running time can be improved to $O(2^t \cdot n)$ [12].

When considering k -hop MŠT in general metric spaces as well as nonmetric distance functions, several hardness results are known. Bern and Plassmann [11] show that the Steiner tree problem on a metric induced by a complete graph with edge weights 1 or 2 is MaxSNP-hard. The same is shown for metric 2-hop MST by Alfandari and Paschos [4]. Thus, these problems do not admit a polynomial-time approximation scheme (PTAS) unless $P = NP$ [7]. For general nonmetric distance functions defined by a graph, Manyem and Stallmann [39] show that k -hop MŠT on a graph with unit-weight edges and 2-hop MST cannot admit a constant-factor approximation algorithm. They also show that k -hop MST on a graph with edge weights 1 or 2 cannot admit a PTAS. When weights on the graph edges are unconstrained, Alfandari and Paschos [4] prove that even for 2-hop MST no $(1 - \varepsilon) \log(n)$ -approximation can exist unless $NP \subseteq DTIME[n^{O(\log \log n)}]$.

The following works, which are conceptually closest to our paper, focus on approximation algorithms. Kortsarz and Peleg [36] consider k -hop MŠT on nonmetric graphs obtaining an approximation factor $O(\log n)$ for constant k and $O(n^\varepsilon)$ otherwise. Althaus et al. [5] give an $O(\log n)$ -approximation for metric k -hop MST for arbitrary k that first uses a randomized embedding of the given metric into a hierarchically separated tree and then solves this problem optimally. For constant k , Laue and Matijević [37] derive a PTAS for k -hop MŠT in the Euclidean plane. Their algorithm implies a quasi-polynomial-time approximation scheme (QPTAS) for Euclidean spaces of higher dimensions. While the first constant-factor approximation algorithm for metric k -hop MŠT is due to Kantor and Peleg [33], the attained approximation factor $1.52 \cdot 9^{k-2}$ is prohibitively high. For $k = 2$, a nearly optimal algorithm is known: the best-known approximation ratio of 1.488 for metric uncapacitated facility location [38] and lower bound of 1.463 [26] are valid for metric 2-hop MST as well.

The bounded-diameter minimum Steiner tree problem [24, 33] is also closely re-

lated to the bounded-hop problem yet is neither a generalization nor a special case. Here, for given d we look for a minimum-cost Steiner tree with diameter at most d . For constant d , an $O(1)$ -approximation algorithm is known for graph metrics [33]. For nonmetric cost functions, an $o(\log n)$ -approximation algorithm has been ruled out, assuming $P \neq NP$ [10].

Furthermore, shallow-light and buy-at-bulk Steiner trees [8, 16, 20, 30, 36] are conceptually similar to k -hop MŠTs. However, a key difference is that, here, lengths of paths in the tree are bounded with respect to metric distance instead of the number of edges on the path. Elkin and Solomon [20] additionally bound the number of hops, but do so by $O(\log n)$ to bound other measures of interest. Chimani and Spoerhase [16] consider two different measures for distance and weight and achieve an n^ε -approximation, violating the distance by a factor of $1 + \varepsilon$.

Minimum-cost k -hop spanning and Steiner trees have been studied in the context of random graphs as well. There, the goal is to give estimates on the weight of an optimal tree. In this setting, sharp thresholds for k are known [6].

1.2. Our results. We give polynomial-time algorithms that optimally solve the k -hop MŠT problem in certain tree-like metrics. Our main and most general result is a dynamic program (DP) for metrics with treewidth ω with a running time of $n^{O(\omega k)}$. As stepping stones towards this result, we first present some key techniques by considering algorithms for simpler metrics, namely the path metric and tree metric. Later, we show how to utilize our general algorithm for (bicriteria) approximation algorithms for other metric spaces.

In section 3, we give a quite simple exact algorithm for k -hop MŠT on the path metric with running time $O(kn^5)$, where n is the number of vertices and k is the number of hops. Thus, the algorithm retains the polynomial running time even when k is part of the input.

In section 4, we consider the special case of tree metrics, i.e., graphs with treewidth one, establishing the essential building blocks for our general algorithm in a more accessible setting. The running time of our algorithm for tree metrics is $n^{O(k)}$.

Let us give a few high-level insights for our algorithm on tree metrics. A tree metric naturally lends itself to recursive computation based on the structure of the underlying tree, so we compute an optimal partial solution for a subtree rooted at some vertex v (denoted by $T[v]$) by using the solutions of its children. We index a dynamic programming cell by v and beyond that—as more information is required on how to connect the subtree to the remainder of the tree—by $2k$ additional vertices which represent possible parents of v at different depths in a k -hop MŠT. We refer to these $2k$ additional vertices as *anchoring guarantees*, and for each possible depth in the k -hop Steiner tree, there is one anchoring guarantee inside of $T[v]$ and one outside. The crucial decision that the DP needs to make is how to correctly propagate these anchoring guarantees when advancing in the recursion.

Finally, in section 5, we present our main result, the general version of our DP that applies to all graph metrics with bounded treewidth. Specifically, our algorithm computes the optimal k -hop MŠT for metrics with treewidth ω in time $n^{O(\omega k)}$. To this end, we consider a so-called *nice tree decomposition* of the input graph, a specifically structured tree whose nodes are *bags* containing at most $\omega + 1$ input vertices, and index our DP cells by a bag b and the $2k$ possible parents, i.e., anchoring guarantees, of each vertex v in the bag. This determines the DP table size and runtime of $n^{O(\omega k)}$.

Our DP for tree metrics is substantially different from the aforementioned DP by Althaus et al. [5] that is tailored to HSTs. Further, while the DP for the plane, the

two-dimensional Euclidean space, by Laue and Matijević [37] has similarities to our construction for tree metrics, a notable difference lies in the indexing of their cells by distances. In our case, such a strategy does not carry enough information; hence, we resort to indexing by vertices, as explained above, and retain more structure.

Our general algorithm also facilitates a quasi-polynomial-time approximation algorithm for more general metrics induced by graphs of bounded highway dimension. This graph class was introduced by Abraham et al. [3] to model transportation networks. Intuitively, in graphs of bounded highway dimension, locally, there exists a small set of transit vertices such that the shortest paths between two distant vertices pass through some transit vertex. We provide the full definition and all details in section 6.

Using a framework by Feldmann et al. [22], we show that our approach for bounded treewidth metrics and the constant-factor approximation by Kantor and Peleg [33] can be combined to design an algorithm that computes, in quasi-polynomial time, a $(k + 1)$ -hop Steiner tree of cost at most $(1 + \varepsilon)OPT_k$, where OPT_k is the cost of the (minimal) k -hop MST. This seems to be the first result taking advantage of a slight relaxation of the hard hop constraints in network design, a research direction proposed by Althaus et al. [5].

Additionally, we consider the concept of doubling dimension that was proposed by Gupta, Krauthgamer and Lee [27]. A metric space is said to have doubling dimension d if every ball of radius $2r$ can be covered by 2^d balls of radius r . While this concept is closely related to highway dimension and Abraham et al. [1] show that constant highway dimension implies constant doubling dimension, they also show that the converse does not hold. However, we argue that the framework by Feldmann et al. [22] can be applied in both settings, yielding a result analogous to that for bounded highway dimension.

Finally, we also show a limit of usefulness of relaxing the hop constraint by a constant amount, at least in the setting of nonmetric distance functions. Extending a result of Manyem and Stallmann [39], we show that, unless $P = NP$, there is no chance of obtaining a $((1 - c) \cdot \log n)$ -approximation, for some constant $c > 0$, with a Steiner tree that uses $k + \ell$ hops and comparing to an optimal solution that uses only k hops.

2. Preliminaries. Let (V, d) be a metric induced by the graph $G = (V, E)$. We assume without loss of generality that the metric is given by G . If this is not the case, we construct G as the complete graph on V where every edge $\{u, v\}$ has weight equal to the distance from u to v . In order to break ties consistently, we assume that the shortest paths in G are unique. This can be achieved by adding some sufficiently small random noise to the weight of each edge of G . We also assume that G is the *minimal graph inducing* (V, d) . That is, no edge in E can be removed without changing the length of some shortest path. A k -hop MST for this modified instance is also optimal for the original instance. Furthermore, the minimal graph G inducing (V, d) is unique.

Given a metric, we can decide in polynomial time whether it is a path metric, a tree metric, or a metric with treewidth ω for some constant $\omega \geq 1$. To verify that G is a path or a tree, we simply run a depth-first search. Moreover, for constant ω , it can be decided in polynomial time whether G has treewidth ω by computing a treewidth decomposition [13].

We give two alternative representations of Steiner trees that are useful when working with partial solutions. Let \hat{S} be a Steiner tree on (V, d) with terminals $\mathcal{X} \subseteq V$ and root $r \in \mathcal{X}$. Let $V_{\hat{S}} \subseteq V$ with $\mathcal{X} \subseteq V_{\hat{S}}$ be the set of vertices in \hat{S} . The tree \hat{S} can




FIG. 1. A 3-hop MST (thick, bent) with root r and terminals (filled) $u_1, u_2, u_3, u_4, z, r, w, w_1, w_2$ on a tree metric (thin, straight) with unit-weight edges. Its cost is 12. The table on the right describes the corresponding labeling ℓ and anchoring α , where x symbolizes vertices not used by the MST.

be viewed as a function mapping a vertex of $V_{\check{S}} \setminus \{r\}$ to its immediate predecessor, i.e., its parent in \check{S} . More generally, for $U \subseteq V$, call a function $\alpha : U \setminus \{r\} \rightarrow V$ an *anchoring on U* . The *anchor* $\alpha(v)$ of vertex v represents its parent in \check{S} , and we set $\alpha(w) = w$ if $w \notin V_{\check{S}}$.

If \check{S} is of minimum cost, this additionally allows for the following representation. Consider a function assigning to each vertex $v \in V_{\check{S}}$ its depth, i.e., the number of edges on the r - v path in \check{S} . Since a vertex $v \in V_{\check{S}} \setminus \{r\}$ of depth x is anchored to the (uniquely determined) vertex of depth $x - 1$ that is of minimum distance to v with respect to d , this yields a complete representation of \check{S} . Generalizing again to subsets $U \subseteq V$, we call a function $\ell : U \rightarrow \{0, 1, \dots, k\} \cup \{\infty\}$ a *labeling on U* . We call $\ell(w)$ the *label of w* and set $\ell(w) = \infty$ if $w \notin V_{\check{S}}$. Note that this representation automatically enforces the k -hop condition. See Figure 1 for an example of a k -hop MST with the corresponding anchoring and labeling.

When \check{S} is of minimum cost and $U = V$, we can easily compute an anchoring from a labeling or vice versa. However, when considering partial solutions, i.e., when $U \subsetneq V$, this may not be possible. Thus, to retain the essential structural information, we utilize both representations simultaneously in this case. This motivates the following definition.

DEFINITION 2.1. A pair (ℓ, α) is called a *labeling-anchoring pair (LAP)* on U if the labeling ℓ and anchoring α are consistent, i.e., for every $u \in U \setminus \{r\}$ for which $\alpha(u) \in U$ and $\ell(u) \neq \infty$, we have $\ell(u) = \ell(\alpha(u)) + 1$. Moreover, if $\ell(u) = \infty$, then $u \notin \mathcal{X}$ and $\alpha^{-1}(u) = \{u\}$.

The *cost* of a LAP (ℓ, α) is given by $\sum_{u \in U \setminus \{r\}} d(u, \alpha(u))$. In this sum, the term $d(u, \alpha(u))$ is called the *cost to anchor u* . When $U \subsetneq V$, we may say *partial LAP* to emphasize that the LAP only represents a portion of \check{S} , namely the edges between U and its anchors.

The representation as LAP is used to avoid the ambiguity that arises from simultaneously considering a Steiner tree \check{S} and the tree-like graph G that induces the underlying metric space. For example, in section 4, both \check{S} and G are trees. Throughout the paper, we represent Steiner trees as LAPs. Hence, we use the term *anchor* to

refer to a predecessor in \check{S} instead of *parent*. Moreover, when talking about distances or closeness, this always refers to distances in G . Given a point v and a set $U \subseteq V$, denote by $\text{closest}_v(U)$ the (unique) element of U with minimum distance to v . For simplicity, we write $\text{closest}_v(u, w)$ instead of $\text{closest}_v(\{u, w\})$.

In sections 4 and 5, when querying a DP cell, a vertex with a desired label may not exist. To make these queries technically simple, we extend the vertex set of the metric to contain an auxiliary vertex, denoted by v_\emptyset . It is defined to have distance ∞ to all other vertices. In order to avoid the use of k auxiliary vertices (one per label), we slightly overload the notation and assume that the equality $\ell(v_\emptyset) = i$ is correct for all $i \in [k]$, where $[k] := \{1, 2, \dots, k\}$. Note that anchoring v_\emptyset incurs an infinite cost, so it will never be used in a k -hop Steiner tree.

3. The k -hop M \check{S} T problem in path metrics. Our first result is an efficient algorithm for k -hop M \check{S} T on path metrics.

THEOREM 3.1. *On path metrics, k -hop M \check{S} T can be solved exactly in time $O(kn^5)$.*

We view a path metric as a set of vertices $V = \{v_1, v_2, \dots, v_n\}$ placed on the real line from left to right by increasing index, such that edges in the path correspond to consecutive vertices.

On path metrics, we observe that there is no algorithmic difference between a k -hop M \check{S} T and a k -hop MST, since there exists a (uniquely defined) minimum-cost k -hop M \check{S} T $\text{OPT} = (\ell, \alpha)$ rooted at $r \in V$ that only uses terminals. Indeed, if OPT contains a nonterminal vertex v , we may simply replace it by the next vertex on the line in the direction in which v has the most edges (we break ties arbitrarily). This removes a nonterminal vertex without increasing the cost of OPT or violating the k -hop condition. In this section, we therefore assume $\mathcal{X} = V$.

We give a recursive procedure which computes the k -hop MST, and we analyze its complexity via dynamic programming. The goal is to first compute the internal (nonleaf) vertices of the k -hop MST and then add the cost of anchoring the leaves to the closest internal vertices.




FIG. 2. *On path metrics, the optimal k -hop M \check{S} T never anchors j to i if $\ell(i) > \ell(s)$ and $i < s < j$.*

A key observation is the following. Fix an internal vertex s with label $\ell(s) < k$. It partitions the remaining vertex set into the vertices on the left of s and those on the right of s . If a vertex i to the left of s has label $\ell(i) > \ell(s)$, then in OPT , the vertex i is never adjacent to a vertex to the right of s ; see Figure 2. This follows from the fact that such a vertex could be attached to s directly, decreasing the overall cost of OPT without using more hops.

We define a recursive expression $A[p, s, a, b]$ for $p \in \mathbb{N}$ and $s, a, b \in [n]$. It yields the minimum cost p -hop spanning tree \check{S} rooted at v_s that contains all vertices v_i with $i \in [a, b]$ and satisfies $s \notin [a, b]$. If $a > b$, let $[a, b] = \emptyset$.

For $p \in \mathbb{N}$ and $s, a, b \in [n]$, define $A[p, s, a, b]$ as follows.

1. If $a > b$, then $A[p, s, a, b] = 0$.
2. If $a = b$, then $A[p, s, a, a] = d(v_s, v_a)$.
3. If $p = 1$, then $A[1, s, a, b] = \sum_{x \in [a, b]} d(v_s, v_x)$ (all vertices anchored to v_s).
4. If $p > 1$, consider the rightmost child $v_{s'}$ of v_s in \check{S} such that $s' \in [a, b]$.




FIG. 3. Computation of $A[p, s, a, b]$ with three recursive calls.

The subtree of \check{S} rooted at $v_{s'}$ covers all vertices v_i with $i \in [c, b]$ for some $c \in [a, s']$. Thus, $A[p, s, a, b]$ is the sum of the cost of this subtree and that of all remaining subtrees of v_s in $[a, c - 1]$ plus the cost of connecting v_s to $v_{s'}$. That is, $A[p, s, a, b]$ is defined as

$$\min_{s' \in [a, b], c \in [a, s' - 1]} d(v_{s'}, v_s) + A[p, s, a, c - 1] + A[p - 1, s', c, s' - 1] + A[p - 1, s', s' + 1, b].$$

See Figure 3 for an illustration where $b < s$. Note that in the last case, any recursive call can refer to an empty interval and incur zero cost.

Proof of Theorem 3.1. Due to the key observation above, $A[p, s, a, b]$ correctly computes the minimum cost of a p -hop spanning tree \check{S} with root v_s and vertices v_i with $i \in [a, b]$: For s' and c as in OPT, there are no edges in OPT between $[a, c - 1]$, $[c, s' - 1]$, and $[s' + 1, b]$. Also, the recursive procedure only queries intervals $[a, b]$ with $s \notin [a, b]$. The cost of OPT is $A[k, r, 0, r - 1] + A[k, r, r + 1, n]$.

We dynamically compute the values $A[p, s, a, b]$ by iterating in an increasing manner over p in an outer loop and the set of intervals $[a, b]$ in an inner loop, with shorter intervals having precedence. This is feasible, as a call of $A[p, s, a, b]$ recursively only queries values $A[p', s', a', b']$ with $p' < p$ or $(b' - a')^+ < (b - a)^+$. Assuming that all previous values are precomputed, the value of a cell $A[p, s, a, b]$ can be computed in time $O(n^2)$. Since there are only kn^3 possible values of (p, s, a, b) to be queried, the total running time is bounded by $O(kn^5)$. \square

4. The k -hop MŠT problem in tree metrics. In this section, we construct a DP for the k -hop MŠT problem on tree metrics, formally proving the following.

THEOREM 4.1. *On tree metrics, k -hop MŠT can be solved exactly in time $n^{O(k)}$.*

Consider an instance of k -hop MŠT with root $r \in \mathcal{X}$ and metric (V, d) induced by a tree $T = (V, E)$. Without loss of generality, we consider T to be rooted at r . For $v \in V$, denote by $T[v]$ the set of vertices in the subtree of T rooted at v .

We start by giving a high-level overview of our approach for computing the minimum cost k -hop Steiner tree $\text{OPT} = (\ell, \alpha)$. We use a DP with cells $\check{A}[v, \rho, \phi]$ indexed by a node $v \in V$ and vectors ρ and ϕ of k vertices each. Intuitively, ρ and ϕ represent *anchoring guarantees* that convey information about the structure of OPT in relation to v and serve as possible points to which v is anchored in α . Specifically, for each possible label i , there is an anchoring guarantee inside (ϕ_i) and one outside (ρ_i) of $T[v]$ acting as candidates for anchoring v in OPT. We show that a cell $\check{A}[v, \rho, \phi]$ computes a partial labeling anchoring pair (LAP) (recall Definition 2.1) on $T[v]$ that is of minimum cost and respects the given anchoring guarantees. The cells are filled up in a bottom-to-top manner, starting at the leaves of the underlying tree T . Doing this consistently, while filling in correct anchoring guarantees, finally yields OPT.




FIG. 4. Anchoring guarantees $\rho(v)$, $\phi(v)$ for vertex v . The subtree $T[v]$ with respect to the underlying metric (thin, straight) is encircled. In this example, the depicted anchoring guarantees are indeed the correct choices for an optimal 3-hop Steiner tree (thick, bent). To satisfy the anchoring guarantees, v must be anchored to either $\phi_2(v)$ or $\rho_2(v)$.

A key property of a k -hop MST that we implicitly use in the DP recursion is that all nodes with label $i + 1$ of a subtree $T[v]$ that are not anchored to a node of $T[v]$ are anchored to the same node—the closest node of label i outside $T[v]$. Locating this node is the motivation of ρ_i . More generally, the objective of the anchoring guarantees ϕ and ρ is to constrain some node labels in order to convey this information consistently to the children v_j of v : the ρ_i associated to a v_j is either the ρ_i or the ϕ_i of v .

4.1. Anchoring guarantees. Fix a vertex $v \in V \setminus \{r\}$. Formally, anchoring guarantees of v are of the form $\phi(v) = (\phi_1(v), \dots, \phi_{k-1}(v))$ and $\rho(v) = (\rho_1(v), \dots, \rho_{k-1}(v))$ with $\phi_i(v) \in T[v]$ and $\rho_i(v) \in V \setminus T[v]$ for all $i \in [k - 1]$. Additionally, we allow the $\phi_i(v)$ and $\rho_i(v)$ to take the value v_\emptyset and let $\rho_0(v) = r$ and $\phi_0(v) = v_\emptyset$. Call anchoring guarantees of a vertex v correct if $\phi_i(v)$ and $\rho_i(v)$ are the closest vertices to v with label i in a k -hop MST inside, respectively, outside of $T[V]$, or equal v_\emptyset if no such vertex exists. See Figure 4 for an example where this is the case. If v is part of the Steiner tree, then for an anchoring guarantee to be correct, the following must hold.

OBSERVATION 4.2. Given a nonroot vertex v in $\text{OPT} = (\ell, \alpha)$ and its correct anchoring guarantees $\rho(v)$, $\phi(v)$, there exists a unique label $i_v \in [k]$ such that $\phi_{i_v}(v) = v$. Moreover, this implies $\ell(v) = i_v$ and $\alpha(v) = \text{closest}_v(\rho_{i_v-1}(v), \phi_{i_v-1}(v))$.

This observation is of crucial importance for choosing the relevant anchoring guarantees when working with partial solutions. Indeed, it means that correct anchoring guarantees completely determine the label and the anchor of v ; hence there is no freedom or complexity left to account for the cost to anchor v in the DP process. Specifically, we are interested in partial LAPs on $T[v]$. Given a LAP, denote by $\lambda_i(v)$ the vertex in $T[v]$ with label i closest to v (or v_\emptyset if no such vertex exists).

DEFINITION 4.3. We define the set $\mathcal{P}(v, \rho(v), \phi(v))$ to be the (possibly empty) set of LAPs on $T[v]$ respecting the anchoring guarantees. That is, its elements (ℓ, α) satisfy the following:

- (i) For all i , we have $\phi_i(v) = \lambda_i(v)$. In particular, if $\phi_i(v) = \phi_j(v)$ and $i \neq j$, then it holds that $\phi_i(v) = v_\emptyset$.
- (ii) A vertex $w \in T[v]$ with $\ell(w) \neq \infty$ is anchored to a vertex of $T[v]$ with label

$\ell(w) - 1$ or to $\rho_{\ell(w)-1}(v)$. Recall that $\ell(w) = \infty$ implies $\alpha(w) = w$ (and $w \notin \mathcal{X}$).

Intuitively, $\mathcal{P}(v, \rho(v), \phi(v))$ represents all relevant ways to extend a partial LAP (ℓ', α') on $V \setminus T[v]$ to V while respecting the anchoring guarantees: all $\phi_i(v)$ are consistent with $\lambda_i(v)$, and all anchors outside of $T[v]$ are some $\rho_i(v)$, respecting labels. Note that vertices of $T[v]$ are anchored either to another vertex in $T[v]$ or to some $\rho_i(v)$. Therefore, if $\rho_i(v)$ is used, it should be the closest vertex to v outside of $T[v]$ for which $\ell'(\rho_i(v)) = i$. Assume (ℓ', α') is extended with minimum cost, and consider the subtree $T[v_j]$ of a child v_j of v . Its vertices are anchored either to a vertex of $T[v_j]$, or to a $\phi_i(v)$ (which may be in the subtree of a different child), or to a $\rho_i(v)$. The anchoring guarantees $\phi_i(v)$ are then necessary to determine the anchoring guarantees $\rho_i(v_j)$ for the children of v . Note that when defining \mathcal{P} , we do not require any constraint on the values of ρ .

4.2. The dynamic program. For $v \neq r$, denote by $A[v, \rho(v), \phi(v)]$ the minimum cost of a LAP on $T[v]$ in $\mathcal{P}(v, \rho(v), \phi(v))$, or ∞ if none exists. Assuming that these values for the children v_1, v_2, \dots, v_p of r and all possible anchoring guarantees are known, we observe that the cost of OPT can be computed by evaluating

$$(4.1) \quad \sum_{j=1}^p \min_{\phi_i(v_j) \in T[v_j], \forall i \in [k-1]} A[v_j, \rho_\emptyset, \phi(v_j)],$$

where $\rho_\emptyset := \{r, v_\emptyset, \dots, v_\emptyset\}$. This expression combines the minimum-cost LAPs on all subtrees of children of r , which, in an optimal solution, only have r as an outer anchor. It is crucial for the algorithm complexity to note that these partial LAPs can be optimized independently as there is no need to have an edge going through r in a Steiner tree solution: r would be a better anchor as it is closer and less deep.

We now describe a DP that computes the values $A[v, \rho(v), \phi(v)]$. Fix some vertex $v \neq r$, and denote by v_1, v_2, \dots, v_p the children of v in T . Keeping Observation 4.2 in mind, we fill the cells $\bar{A}[v, \rho(v), \phi(v)]$ of our dynamic programming table according to the following recursive relation. For a vertex v that is a leaf of T , we define

$$(4.2) \quad \bar{A}[v, \rho(v), \phi(v)] := \begin{cases} 0 & \text{if } v \notin \mathcal{X} \text{ and } \phi_i(v) = v_\emptyset \text{ for all } i; \\ d(v, \rho_{i_v-1}(v)) & \text{if } \exists \text{ unique } i_v, \text{ s.t. } \phi_{i_v}(v) = v; \\ \infty & \text{otherwise.} \end{cases}$$

The three cases correspond, respectively, to a nonterminal node with no need to get a particular label; a node required by anchoring guarantees to get a label; and inconsistent anchoring guarantees. For nonleaf vertices, we define

$$(4.3) \quad \bar{A}[v, \rho(v), \phi(v)] := c_v + \sum_{j=1}^p \min_{\phi_i(v_j) \in \Phi_i(v_j), \forall i} \bar{A}[v_j, \rho(v_j), \phi(v_j)].$$

Here, c_v denotes the cost of anchoring v , while $\Phi_i(v_j)$ and $\rho(v_j)$ encode which of the n^{2k-2} possible anchoring guarantees of v_j are consistent with that of v . The cells of each child are queried independently, which is crucial for the algorithm complexity: the choice of some $\phi(v_j)$ does not impact the choice for other children of v . Precise definitions of $\Phi_i(v_j)$, $\rho(v_j)$, and c_v follow.

Let $\Phi_i(v_j)$ be the subset of $T[v_j]$ consisting of all feasible choices for $\phi_i(v_j)$. Specifically, if $\phi_i(v) \in T[v_j]$, then $\Phi_i(v_j) = \{\phi_i(v)\}$. Indeed, as the shortest v - $\phi_i(v)$ path passes through v_j , node $\phi_i(v)$ must be the closest vertex to v_j in $T[v_j]$ with

(already guaranteed) label i . If $\phi_i(v) = v_\emptyset$, we must have $\Phi_i(v_j) = \{v_\emptyset\}$ or contradict property (i). Otherwise, if $v_\emptyset \neq \phi_i(v) \notin T[v_j]$, then $\Phi_i(v_j)$ contains all $w \in T[v_j]$ with $d(v, w) \geq d(v, \phi_i(v))$ and the auxiliary vertex v_\emptyset . A distance $d(v, w) < d(v, \phi_i(v))$ would contradict the choice of $\phi_i(v)$ as the vertex in $T[v]$ with label i closest to v .

As for $\rho_i(v_j)$, we define it to be the feasible choice for $\rho_i(v_j)$, which is (uniquely) determined as follows. If $\phi_i(v) \in T[v_j]$, then $\rho_i(v_j) = \rho_i(v)$ since the shortest v_j - $\rho_i(v_j)$ path passes through v . Otherwise, we have $\rho_i(v_j) = \text{closest}_v(\rho_i(v), \phi_i(v))$.

We now define c_v . If $v \notin \mathcal{X}$ and no $\phi_i(v)$ equals v , then $c_v := 0$. Next, if there exists a unique i_v such that $\phi_{i_v}(v) = v$, let $c_v := d(v, \text{closest}_v(\rho_{i_v-1}(v), \phi_{i_v-1}(v)))$. In all other cases set $c_v := \infty$, as the values of $\phi(v)$ are contradictory.

4.3. Analysis of the dynamic program. The complexity to evaluate (4.1) and (4.3) is linear in the size of the table, i.e., $n^{O(k)}$, and the complexity of (4.2) is $O(k)$. Thus, in order to prove Theorem 4.1, it remains to show the correctness of the DP. In this analysis, properties (i) and (ii) refer to Definition 4.3.

Proof of Theorem 4.1. By mathematical induction, we prove that $\bar{A}[v, \rho(v), \phi(v)]$, as defined in (4.2) and (4.3), is equal to $A[v, \rho(v), \phi(v)]$ for any node $v \neq r$ and all possible anchoring guarantees $\rho(v)$ and $\phi(v)$.

For the base step, i.e., when v is a leaf of T , we consider the three cases of (4.2). If $v \notin \mathcal{X}$ and $\phi_i(v) = v_\emptyset$ for all i , then clearly properties (i) and (ii) are satisfied for the LAP that excludes v from the Steiner tree, so $\bar{A}[v, \rho(v), \phi(v)] = A[v, \rho(v), \phi(v)] = 0$. Otherwise, there is at most one LAP that satisfies (i) and (ii)—namely the one that anchors v to $\rho_{i_v-1}(v)$ if i_v is defined. It incurs a cost of $d(v, \rho_{i_v-1}(v))$, as desired. If no such LAP exists, then $\mathcal{P}(v, \rho(v), \phi(v)) = \emptyset$ and both $A[v, \rho(v), \phi(v)]$ and $\bar{A}[v, \rho(v), \phi(v)]$ are infinite. This concludes the base step.

Our induction hypothesis is that

$$(IH) \quad \bar{A}[v', \rho(v'), \phi(v')] = A[v', \rho(v'), \phi(v')] \quad \text{for all } v' \neq r, \rho(v'), \text{ and } \phi(v').$$

Now, for some nonleaf v , we assume that (IH) holds for all descendants $v' \in T[v] \setminus \{v\}$ of v and prove that (IH) holds for v as well. For v , the recursive equation (4.3) becomes

$$\bar{A}[v, \rho(v), \phi(v)] := c_v + \sum_{j=1}^p \min_{\phi_i(v_j) \in \Phi_i(v_j), \forall i} A[v_j, \rho(v_j), \phi(v_j)].$$

If $c_v = \infty$, then $\mathcal{P}(v, \rho(v), \phi(v)) = \emptyset$, so both $A[v, \rho(v), \phi(v)] = \infty = \bar{A}[v, \rho(v), \phi(v)]$. From now on, assume that c_v is finite. We prove (IH) for v by showing the two inequalities $A[v, \rho(v), \phi(v)] \geq \bar{A}[v, \rho(v), \phi(v)]$ and $A[v, \rho(v), \phi(v)] \leq \bar{A}[v, \rho(v), \phi(v)]$. Once these two inequalities are proven, the proof of Theorem 4.1 is complete.

Claim. We have $A[v, \rho(v), \phi(v)] \geq \bar{A}[v, \rho(v), \phi(v)]$.

Consider the LAP (ℓ, α) which yields the value $A[v, \rho(v), \phi(v)]$. In particular, properties (i) and (ii) are satisfied. If no such LAP exists, then $A[v, \rho(v), \phi(v)] = \infty$, and the inequality holds. For each child v_j of v , set $\phi_i(v_j) = \lambda_i(v_j)$, which respects $\phi_i(v_j) \in \Phi_i(v_j)$. Also, set $\rho(v_j) = \rho(v_j)$ as defined above. We show that for each v_j , the restriction of the LAP (ℓ, α) to $T[v_j]$ belongs to $\mathcal{P}(v_j, \rho(v_j), \phi(v_j))$.

Property (i) follows directly from the choice of $\phi_i(v_j) = \lambda_i(v_j)$.

For property (ii), consider a vertex $w \in T[v_j]$ which is not anchored to a vertex of $T[v_j]$. We show that α anchors w to $\rho_{\ell_w}(v_j)$, with $\ell_w := \ell(w) - 1$. Note that by definition of $\rho(v_j)$, we have that $\rho_{\ell_w}(v_j)$ equals $\rho_{\ell_w}(v)$ or $\phi_{\ell_w}(v)$, so $\ell(\rho_{\ell_w}(v_j)) = \ell_w$. Since α is an anchoring of minimal cost (with respect to the given anchoring

guarantees), w is anchored to the vertex $\alpha(w) = \text{closest}_w\{x \in T[v] \cup \{\rho_{\ell_w}(v)\} \mid \ell(x) = \ell_w\}$, so $\alpha(w) = \text{closest}_{v_j}(\rho_{\ell_w}(v), \phi_{\ell_w}(v))$. If $\phi_{\ell_w}(v) \in T[v_j]$, then $\rho_{\ell_w}(v_j) = \rho_{\ell_w}(v) = \alpha(w)$ as w is not anchored to a vertex in $T[v_j]$. If $\phi_{\ell_w}(v) \notin T[v_j]$, then $\rho_{\ell_w}(v_j) = \text{closest}_{v_j}(\rho_{\ell_w}(v), \phi_{\ell_w}(v)) = \alpha(w)$, by definition of $\rho(v_j)$.

Therefore, the LAP (ℓ, α) restricted to $T[v_j]$ belongs to $\mathcal{P}(v_j, \rho(v_j), \phi(v_j))$, so its cost is at least $A[v_j, \rho(v_j), \phi(v_j)]$. If $\ell(v) \neq \infty$, then $\alpha(v) = \text{closest}_v(\rho_{i_v-1}(v), \phi_{i_v-1}(v))$ with cost c_v , since the anchoring cost is minimized. If $\ell(v) = \infty$, then $c_v = 0$, so

$$A[v, \rho(v), \phi(v)] = c_v + \sum_{j=1}^p A[v_j, \rho(v_j), \phi(v_j)] \geq \bar{A}[v, \rho(v), \phi(v)].$$

Claim. We have $A[v, \rho(v), \phi(v)] \leq \bar{A}[v, \rho(v), \phi(v)]$.

We assume $\bar{A}[v, \rho(v), \phi(v)]$ to be finite; otherwise the inequality trivially holds. Consider the LAPs that correspond to the values $A[v_j, \rho(v_j), \phi(v_j)]$ for which the value $\bar{A}[v, \rho(v), \phi(v)]$ is attained. We extend these LAPs on the subtrees $T[v_j]$ to (ℓ, α) on $T[v]$ in the following way. If $v \notin \mathcal{X}$ and no $\phi_i(v)$ equals v , we let $\ell(v) = \infty$ and $\alpha(v) = v$. Otherwise, as $c_v \neq \infty$ by our assumption at the start of the proof, there exists a unique i_v such that $\phi_{i_v}(v) = v$. We then let $\ell(v) = i_v$ and anchor v to $\text{closest}_v(\rho_{i_v-1}(v), \phi_{i_v-1}(v))$. We show that this yields an element of $\mathcal{P}(v, \rho(v), \phi(v))$.

We first show property (i). If i_v is defined, $\phi_{i_v}(v) = v = \lambda_{i_v}(v)$ since $\ell(v) = i_v$. Consider $\phi_i(v)$ for $i \neq i_v$. If $\phi_i(v) = v_\emptyset$, then all $\phi_i(v_j) = v_\emptyset$, too, by definition of $\Phi_i(v_j)$. Thus, $\lambda_i(v_j) = v_\emptyset$ for all j and $\ell(v) \neq i$, so $\lambda_i(v) = v_\emptyset = \phi_i(v)$. Otherwise, if $\phi_i(v) \neq v_\emptyset$, there exists a j_i with $\phi_i(v) \in T[v_{j_i}]$. Then, we have $\lambda_i(v_{j_i}) = \phi_i(v)$, and for all j , we have $d(v, \lambda_i(v_j)) = d(v, \phi_i(v_{j_i})) \geq d(v, \phi_i(v))$. Since $\ell(v) \neq i$, we obtain $\lambda_i(v) = \phi_i(v)$.

It is easy to see that property (ii) holds as well. If we set $\alpha(v) = v$, then $v \notin \mathcal{X}$ and $\ell(v) = \infty$. Otherwise, we define $\alpha(v)$ to be either $\rho_{i_v-1}(v)$ or $\phi_{i_v-1}(v) \in T[v]$. Furthermore, any vertex w of $T[v_j]$ is anchored either to a vertex in $T[v_j] \subseteq T[v]$ or to $\rho_{\ell(w)-1}(v_j)$, since the partial anchorings fulfill property (ii). That means w is anchored either to a vertex of $T[v]$ or, by definition of $\rho(v_j)$, to $\rho_{\ell(w)-1}(v)$.

In conclusion, $(\ell, \alpha) \in \mathcal{P}(v, \rho(v), \phi(v))$, so its cost is at least $A[v, \rho(v), \phi(v)]$. \square

5. Metrics of bounded treewidth. We proceed to present our algorithm in full generality, building upon ideas from section 4 to obtain our main result on metrics with bounded treewidth, which are defined as follows.

DEFINITION 5.1. A graph $G = (V, E)$ is said to have treewidth ω if there exists a tree $T_G = (B, E_B)$ whose nodes $b \in B$ are identified with subsets $S_b \subseteq V$, called bags, satisfying the following:

- (i) for each edge in E , there is a bag containing both endpoints;
- (ii) for each vertex in V , the bags containing it form a connected subtree of T_G ;
and
- (iii) each bag contains at most $\omega + 1$ vertices.

The tree T_G is called a tree decomposition of G .

We say that a metric (V, d) has treewidth ω if there exists a graph $G = (V, E)$ with treewidth ω that induces it. Our main result is as follows.

THEOREM 5.2. There exists an algorithm that solves the k -hop MST problem exactly on metrics with treewidth ω and has a running time of $n^{O(\omega k)}$.

Figure 5 spotlights the main difference between sections 4 and 5. In the former, we consider a subtree $T[v]$ of a tree metric (V, d) and the question of how to extend a LAP on $V \setminus T[v]$ to $T[v]$. Here, we could exploit the fact that shortest paths between




FIG. 5. Illustration of the key difference between tree metrics (left) and metrics of treewidth ω (right). While for tree metrics, we extend LAPs to subgraphs separated by only one cut-vertex, v , in the bounded treewidth case, cuts are formed by the bags S_b of the tree decomposition and have size at most ω .

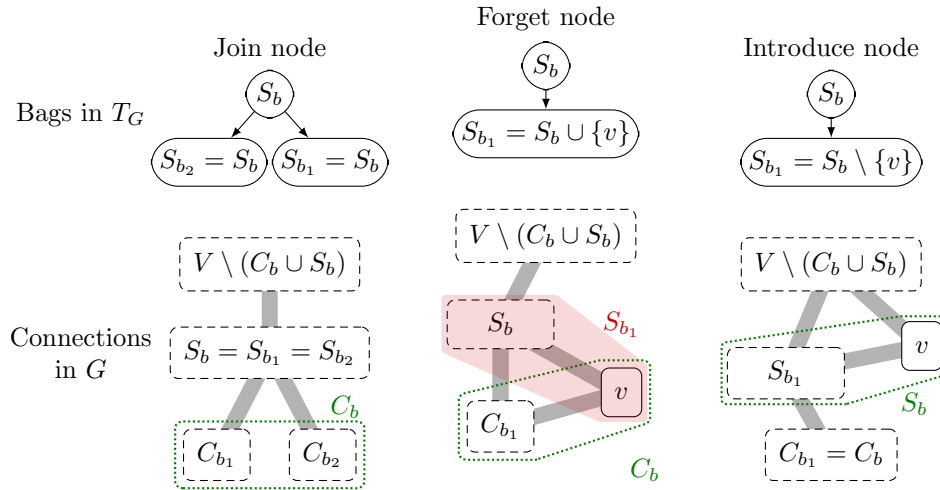


FIG. 6. Types of bag nodes in a nice tree decomposition and possible edges in G .

these sets pass through the root of the subtree, v , which forms a vertex cut of size 1. In the bounded treewidth case, however, we consider the cuts formed by the bags S_b of the tree decomposition. These contain up to ω vertices. While the approach conceptually is the same, the setting requires more care and is technically involved.

Therefore, it will be useful to consider a more restricted variant of tree decompositions (see also [17]), which we now define.

DEFINITION 5.3. A nice tree decomposition is a tree decomposition $T_G = (B, E_B)$ in which, with respect to a designated root $b_r \in B$, every node $b \in B$ is of exactly one of the following four types:

- Leaf: Its bag is empty, that is, $S_b = \emptyset$.
- Join node: It has two children b_1 and b_2 with $S_b = S_{b_1} = S_{b_2}$.
- Forget node of v : It has one child b_1 with $S_{b_1} = S_b \cup \{v\}$ and $v \notin S_b$.
- Introduce node of v : It has one child b_1 with $S_{b_1} = S_b \setminus \{v\}$ and $v \in S_b$.

For an illustration, we refer to Figure 6. By Definition 5.1, property (ii), a vertex in V may have several introduce nodes but at most one forget node. Let C_b be the

union of the bags $S_{b'}$ for all descendants b' of b , excluding vertices in S_b . Property (ii) implies that there is no edge between C_b and $V \setminus (S_b \cup C_b)$, and that, for a join node, $C_{b_1} \cap C_{b_2} = \emptyset$.

Given a graph with treewidth ω , it is possible to compute a nice tree decomposition with $|B| = O(n\omega)$ in polynomial time [17]. Without loss of generality our input is a nice tree decomposition T_G .

Choose a root node b_r whose bag contains the root r of the k -hop MŠT which we aim to compute. To extend the dynamic programming approach from section 4 to nice tree decompositions, we again compute cells in a bottom-up fashion, now in T_G . As mentioned above, a key difference lies in the fact that, here, a node b in T_G corresponds to several vertices in G , so we require anchoring guarantees for every vertex in S_b . A DP cell, indexed by a bag b and $O(\omega k)$ anchoring guarantees, computes a minimum cost LAP on C_b that respects these guarantees. Thankfully, the structure of the nice tree decomposition enables us to recurse in an organized manner and construct the cells consistently. Join nodes combine previous results. Forget nodes decide the label and anchoring of the corresponding vertex and possibly new anchoring guarantees needed due to forgetting it. Introduce nodes deduce anchoring guarantees about the introduced vertex from previous knowledge.

We again implicitly rely on structural properties of Steiner trees to make use of no more than $O(\omega k)$ anchoring guarantees describing the structure outside of C_b . Consider a k -hop MŠT where two nodes v_1 and v_2 of C_b with the same label $i + 1$ are anchored to nodes outside of C_b , and the shortest paths from v_1 to $\alpha(v_1)$ and from v_2 to $\alpha(v_2)$ go through the same node $u \in S_b$. Then, we must have $\alpha(v_1) = \alpha(v_2)$. Locating this node is the objective of the anchoring guarantee $\rho_i(u)$, and the anchoring guarantees $\phi_i(u)$ allow us to convey this information consistently through the DP recursion.

5.1. Anchoring guarantees. Fix a bag $b \in B$. Its *anchoring guarantees* are of the form $\phi(b) = \{\phi_i^u(b) \mid i \in [k-1] \wedge u \in S_b\}$ and $\rho(b) = \{\rho_i^u(b) \mid i \in [k-1] \wedge u \in S_b\}$, with all $\phi_i^u(b) \in \{v_\emptyset\} \cup C_b$ and $\rho_i^u(b) \in \{v_\emptyset\} \cup V \setminus C_b$. We additionally set $\phi_0^u(b) = v_\emptyset$ and $\rho_0^u(b) = r$ for all u ; see Figure 7. We use these anchoring guarantees to define a subset of partial LAPs on C_b analogously to Definition 4.3.

DEFINITION 5.4. *We define the set $\mathcal{P}(b, \rho(b), \phi(b))$ to be the (possibly empty) set of partial LAPs on C_b respecting the anchoring guarantees. That is, its elements (ℓ, α) satisfy the following:*

- (i') $\phi_i^u(b)$ is the closest vertex to u in C_b with label i (or v_\emptyset if no such vertex exists);
- (ii') each vertex u of C_b with $\ell(u) \neq \infty$ is anchored either to a vertex of C_b with label $\ell(u) - 1$ or to $\rho_{\ell(u)-1}^w(b)$ for some $w \in S_b$;
- (iii') for all i and $u, w \in S_b$, we have $d(u, \rho_i^u(b)) \leq d(u, \rho_i^w(b))$.

Intuitively, $\mathcal{P}(b, \rho(b), \phi(b))$ represents all relevant ways to extend a partial LAP on $V \setminus C_b$ to V . Vertices of C_b are anchored either to a vertex of C_b or to some $\rho_i^u(b)$. An addition compared to section 4 is the last item, which requires that the $\rho_i^u(b)$ are consistent between each other. If the inequality were reversed, then $\rho_i^w(b)$ would have been a better choice for $\rho_i^u(b)$. Indeed, the intuitive objective of $\rho_i^u(b)$ is to select the vertex of $V \setminus C_b$ with label i closest to u . As this is a guarantee too strong to be maintained in some cases where $\rho_i^u(b)$ is not useful as it is too far compared to $\phi_i^u(b)$, we only require this weaker guarantee of consistency among the other anchoring guarantees.




FIG. 7. Possible values of $\rho_i(b)$ and $\phi_i(b)$ with respect to two vertices u and w . Note that $\phi_3^u = \phi_3^w = v_0$.

5.2. The dynamic program. Let $A[b, \rho(b), \phi(b)]$ be the minimum cost of a partial LAP on C_b in $\mathcal{P}(b, \rho(b), \phi(b))$, or ∞ if this set is empty. Note that if property (iii') is satisfied and b is a leaf, then $A[b, \rho(b), \phi(b)] = 0$.

First, let us argue how knowledge of the values $A[b, \rho(b), \phi(b)]$ allows us to compute the cost of the optimal k -hop Steiner tree OPT.

We have already fixed one bag b_r of the nice tree decomposition containing r to serve as the root of the decomposition. We can imagine OPT being computed in a two-phase process: First, the optimal labeling and anchoring on vertices of the set S_{b_r} are selected, with anchoring choices even outside S_{b_r} being taken into consideration. Next, we compute the optimal extension of this anchoring and labeling preselection on the rest of the vertices—namely $V \setminus S_{b_r}$. Such a process clearly yields the cost of OPT.

As any LAP in $\mathcal{P}(b_r, \rho(b_r), \phi(b_r))$ by definition covers the entire $C_{b_r} = V \setminus S_{b_r}$ and $A[b_r, \rho(b_r), \phi(b_r)]$ is the minimum cost of the LAPs therein, we can repeat the above two-phase process formally.

Assume we are given a partial labeling $\bar{\ell}$ on $S_{b_r} \setminus \{r\}$ and values $\phi(b_r)$. For all $v \in S_{b_r}$, we define $\rho_i^v(b_r)$ as the vertex w closest to v for which $\bar{\ell}(w) = i$. The minimum cost of a LAP extending $\bar{\ell}$ that respects $\phi(b_r)$ is equal to

$$(5.1) \quad A[b_r, \rho(b_r), \phi(b_r)] + \sum_{v \in S_{b_r} \setminus \{r\}} d(v, \text{closest}_v(\phi_{\bar{\ell}(v)-1}^v(b_r), \rho_{\bar{\ell}(v)-1}^v(b_r))).$$

As argued above, picking the $\bar{\ell}$ and ϕ that minimize this value gives us the cost of OPT.

In the following, we define in a recursive procedure how to fill the cells $\bar{A}[b, \rho(b), \phi(b)]$ of our dynamic programming table \bar{A} for each node $b \in B$. The goal will be to again show that $\bar{A} = A$. First, as an easy special case, if property (iii') is not respected by $\rho(b)$, we set \bar{A} to infinity, which matches with the fact that $\mathcal{P}(b, \rho(b), \phi(b))$ is empty and $A = \infty$. Next, we describe how to compute \bar{A} depending on the type of the node b when (iii') is respected.

Leaves: Node b has no child and $S_b = \emptyset$. We set $\bar{A}[b, \rho(b), \phi(b)] = 0$.

Join nodes: Node b has children b_1, b_2 with $S_{b_1} = S_{b_2} = S_b$ and $C_{b_1} \cup C_{b_2} = C_b$ and $C_{b_1} \cap C_{b_2} = \emptyset$. Intuitively, the objective is to independently query partial solutions on each C_{b_j} . We compute sets of possible values for $\rho_i^u(b_j)$ and $\phi_i^u(b_j)$ which define sets of partial LAPs on each C_{b_j} respecting such guarantees. These possible values are determined such that the minimum cost of a combination of any two partial LAPs on C_{b_1} and C_{b_2} in these sets equals $A[b, \rho(b), \phi(b)]$. Here, the $\rho_i^u(b_j)$ need to be equal to the closest anchoring guarantee outside of C_{b_j} . The $\phi_i^u(b_j)$ may take any value not contradicting $\phi(b)$. Specifically, for both $j \in \{1, 2\}$, $i \in [k - 1]$ and $u \in S_b$, we do the following:

- We set $\rho_i^u(b_j) = \text{closest}_u\{\{\rho_i^u(b)\} \cup \{\phi_i^w(b) \mid w \in S_b \wedge \phi_i^w(b) \notin C_{b_j}\}\}$.
- If $\phi_i^u(b) \in C_{b_j}$, then we set $\Phi_i^u(b_j) = \{\phi_i^u(b)\}$. Otherwise, we set

$$\Phi_i^u(b_j) = \{x \in C_{b_j} \cup \{v_\emptyset\} \mid \text{for all } z \in S_b, \text{ we have } d(z, \phi_i^z(b)) \leq d(z, x) \quad (\star)\},$$

where (\star) ensures that $\phi_i^z(b)$ is the vertex in C_b that is closest to z .

We then define

$$(5.2) \quad \bar{A}[b, \rho(b), \phi(b)] = \sum_{j \in \{1,2\}} \min_{\phi_i^u(b_j) \in \Phi_i^u(b_j), \forall i \in [k-1], u \in S_{b_j}} \bar{A}[b_j, \rho(b_j), \phi(b_j)].$$

Forget node of v: We have $S_{b_1} = S_b \cup \{v\}$ and $C_{b_1} = C_b \setminus \{v\}$. There is no edge between v and $V \setminus (C_b \cup S_b)$. In this node, we want to define the label i_v of v and the corresponding anchoring of v . However, $\phi(b)$ may not contain sufficient information for deciding i_v since v can be far away from S_b . We therefore need to consider all possible values for i_v , that is, all values that are consistent with the guarantees $\phi(b)$. We first define the set \mathbf{I}_v of possible labels of v that do not contradict $\phi(b)$, then proceed to define possible values of $\phi(b_1)$, $\rho(b_1)$, and finally express the cost to anchor v .

Let \mathbf{I}_v be the set of labels i such that for all $u \in S_b$, we have $d(u, v) \geq d(u, \phi_i^u(b))$, and for all $i' \neq i$, we have $\phi_{i'}^u(b) \neq v$. In other words, if there is a label i and $u \in S_b$ with $\phi_i^u(b) = v$, then \mathbf{I}_v cannot contain any other label: In order to respect the guarantee $\phi_i^u(b) = v$, we must have $i_v = i$. Moreover, if there exist some $u \in S_b$ and i such that $\phi_i^u(b)$ is further from u than v , then \mathbf{I}_v cannot contain i as it would contradict the definition of $\phi_i^u(b)$. If $v \notin \mathcal{X}$ and no $\phi_i^u(b)$ equals v , we include ∞ in \mathbf{I}_v as v does not need to have a finite label in order to respect the guarantees $\phi_i(b)$. If \mathbf{I}_v is empty, set $\bar{A}[b, \rho(b), \phi(b)]$ to be infinite since it is impossible to label v while respecting the guarantees $\phi_i(b)$. Assume now that \mathbf{I}_v is not empty.

The values $\phi_i^u(b_1)$ can take any value in C_{b_1} not contradicting $\phi(b)$. Specifically, for $u \in S_b$, if $\phi_i^u(b) \neq v$, let $\Phi_i^u(b_1) = \{\phi_i^u(b)\}$, and if $\phi_i^u(b) = v$, let

$$\Phi_i^u(b_1) = \{x \in C_{b_1} \cup \{v_\emptyset\} \mid d(u, v) \leq d(u, x)\}.$$

Indeed, if $\phi_i^u(b) = v$, then we need to provide a new guarantee for $\phi_i^u(b_1)$, as $v \in S_{b_1}$, which must be further from u than v . We also define

$$\Phi_i^v(b_1) = \{x \in C_{b_1} \cup \{v_\emptyset\} \mid \text{for all } u \in S_b, \text{ we have } d(u, \phi_i^u(b)) \leq d(u, x) \quad (\star)\}.$$

Again, (\star) must be satisfied since $\phi_i^u(b)$ is the vertex in C_b which is closest to u .

For the remainder, fix some $i_v \in \mathbf{I}_v$. In the case where $i_v = \infty$, we need not consider ρ_{i_v} 's. Otherwise, any path from v to a vertex in $V \setminus C_b$ passes through S_b . Therefore, $\rho_{i_v}^v(b_1)$ is determined by $\rho_{i_v}^v(b_1) = \text{closest}_v\{\rho_{i_v}^u(b) \mid u \in S_b\}$ for $i \neq i_v$, and $\rho_{i_v}^v(b_1) = v$. Similarly, for $u \in S_b$, let $\rho_i^u(b_1) = \rho_i^u(b)$ for $i \neq i_v$, and $\rho_{i_v}^u(b_1) = \text{closest}_u\{\rho_{i_v}^u(b), v\}$.

Additionally, we charge a cost of c_{i_v} for anchoring v . If $i_v = \infty$, then set $c_{i_v} := 0$. Otherwise, set $c_{i_v} := d(v, \text{closest}_v\{\phi_{i_v-1}^v(b_1), \rho_{i_v-1}^v(b_1)\})$.

We then define, with $\rho(b_1)$ depending on i_v and c_{i_v} depending on $\phi_i^u(b_1)$,

$$(5.3) \quad \bar{A}[b, \rho(b), \phi(b)] = \min_{i_v \in \mathbf{I}_v} \min_{\phi_i^u(b_1) \in \Phi_i^u(b_1), \forall i \in [k-1], u \in S_{b_j}} \left(c_{i_v} + \bar{A}[b_1, \rho(b_1), \phi(b_1)] \right).$$

Introduce node of v: In this case, b has one child b_1 with $S_{b_1} = S_b \setminus \{v\}$ and $C_{b_1} = C_b$. There is no edge between v and $C_b = C_{b_1}$ as $v \notin S_{b_1} \cup C_{b_1}$; see Figure 6. If

there is an i with $\phi_i^v(b) \neq \text{closest}_v\{\phi_i^u(b) \mid u \in S_{b_1}\}$, then $\bar{A}[b, \rho(b), \phi(b)]$ is infinite since the shortest v - $\phi_i^v(b)$ path has to pass through a vertex of S_b by the above observation. Otherwise, the guarantees do not change, so we define $\rho(b_1) = \rho(b)$, $\phi(b_1) = \phi(b)$, and we set

$$(5.4) \quad \bar{A}[b, \rho(b), \phi(b)] = \bar{A}[b_1, \rho(b_1), \phi(b_1)].$$

5.3. Analysis of the dynamic program. Given the above definitions and the size of the dynamic programming table, one can check that \bar{A} as well as (5.1) can be computed in time $n^{O(\omega^k)}$. It remains to show the correctness of the DP.

Proof of Theorem 5.2. By mathematical induction on the nice tree decomposition T_G , we prove that the value of the DP cell $\bar{A}[b, \rho(b), \phi(b)]$ as defined in subsection 5.2 is equal to the minimum cost $A[b, \rho(b), \phi(b)]$ of a LAP in $\mathcal{P}(b, \rho(b), \phi(b))$.

If b is a leaf, then depending on property (iii'), both values are either ∞ or zero, and the result holds. Otherwise, we assume by the induction hypothesis that for all children b_j of b and all possible anchoring guarantees $\rho(b_j)$ and $\phi(b_j)$, we have $A[b_j, \rho_j(b), \phi_j(b)] = \bar{A}[b_j, \rho_j(b), \phi_j(b)]$. We now show this equality for node b by proving both inequalities separately and according to the type of b , concluding the proof of Theorem 5.2.

Claim. We have $A[b, \rho(b), \phi(b)] \geq \bar{A}[b, \rho(b), \phi(b)]$.

Consider a bag b and any values $\rho(b)$ and $\phi(b)$. Consider a LAP $(\ell, \alpha) \in \mathcal{P}(b, \rho(b), \phi(b))$ for which the value $A[b, \rho(b), \phi(b)]$ is attained. If no such LAP exists, then $A[b, \rho(b), \phi(b)] = \infty$, and the inequality holds. There are three different cases depending on the type of b . For each case, we focus on a bag node b_j child of b and define values $\rho(b_j)$ and $\phi(b_j)$. We show that the restriction of (ℓ, α) to C_{b_j} belongs to $\mathcal{P}(b_j, \rho(b_j), \phi(b_j))$. We then show that the value of cell $\bar{A}[b_j, \rho(b_j), \phi(b_j)]$ was considered in the computation of $\bar{A}[b, \rho(b), \phi(b)]$, i.e., each $\phi_i^u(b_j)$ belongs to the corresponding $\Phi_i^u(b_j)$ and $\rho(b_j) = \rho(b_j)$.

We define $\phi_i^u(b_j)$ as the closest vertex to u in C_{b_j} with label i (with respect to ℓ), and $\rho_i^u(b_j) = \rho_i^u(b_j)$, which is defined in section 5 according to the type of bag b . This way, $\phi_i^u(b_j)$ automatically satisfies property (i') for (ℓ, α) restricted to C_{b_j} . In order to prove that the restriction of (ℓ, α) belongs to $\mathcal{P}(b_j, \phi(b_j), \rho(b_j))$, it therefore remains to show that this LAP also respects properties (ii') and (iii') regarding $b_j, \phi(b_j), \rho(b_j)$.

Once all three requirements, $\phi_i^u(b_j) \in \Phi_i^u(b_j)$, property (ii'), and property (iii'), are verified, the definition for each bag type of $\bar{A}[b, \rho(b), \phi(b)]$ (equations (5.2)–(5.4)) and the induction hypothesis on the children of b lead to the desired inequality $A[b, \rho(b), \phi(b)] \geq \bar{A}[b, \rho(b), \phi(b)]$.

- *Join nodes:* For a join node b with children b_1, b_2 , we focus on a single child b_j . We first show that each $\phi_i^u(b_j)$ belongs to $\Phi_i^u(b_j)$. If $\phi_i^u(b) \in C_{b_j}$, then $\phi_i^u(b_j) = \phi_i^u(b)$ as desired. Otherwise, $\phi_i^u(b_j)$ cannot be closer to u than $\phi_i^u(b)$, satisfying condition (\star) from the definition of $\Phi_i^u(b_j)$.

Regarding property (ii'), consider a vertex $v_1 \in C_{b_j}$ anchored to $v_2 = \alpha(v_1) \notin C_{b_j}$ and a vertex $u \in S_b$ on the shortest path from v_1 to v_2 . The objective is to show that $v_2 = \rho_{\ell(v_1)-1}^u(b_j)$. As α is an anchoring of minimum cost that respects property (ii'), we must have $v_2 = \text{closest}_u\{\rho_{\ell(v_1)-1}^u(b), \phi_{\ell(v_1)-1}^u(b)\}$. If $\phi_{\ell(v_1)-1}^u(b) \in C_{b_j}$, we must have $v_2 = \rho_{\ell(v_1)-1}^u(b)$, as $v_2 \notin C_{b_j}$. By definition of $\phi_{\ell(v_1)-1}^u(b)$ and $\rho_{\ell(v_1)-1}^u(b_j)$, we obtain $\rho_{\ell(v_1)-1}^u(b_j) = \rho_{\ell(v_1)-1}^u(b) = v_2$. If $\phi_{\ell(v_1)-1}^u(b) \notin C_{b_j}$, then by definition $v_2 = \rho_{\ell(v_1)-1}^u(b_j)$.

For property (iii'), consider i and $u, v \in S_{b_j}$. Since $(\ell, \alpha) \in \mathcal{P}(b, \rho(b), \phi(b))$,

we have $d(u, \rho_i^u(b)) \leq d(u, \rho_i^w(b))$. From the definition of $\rho_i(b)$, we get $d(u, \rho_i^u(b_j)) \leq d(u, \rho_i^w(b_j))$.

- *Forget nodes:* Recall that for a forget node b with respect to v , we have $S_{b_1} = S_b \cup \{v\}$ and $C_{b_1} = C_b \setminus \{v\}$. Let $i_v = \ell(v)$, which can be ∞ . Using the fact that $\phi(b)$ satisfies property (i'), it is easy to see that $i_v \in \mathbf{I}_v$. Clearly, the cost to anchor v in ℓ equals c_{i_v} .

For a given i and $u \in S_b$ (implying $u \neq v$), assume first that $\phi_i^u(b) \neq v$. Then, $\phi_i^u(b) \in C_{b_1}$, so we must have $\phi_i^u(b_1) = \phi_i^u(b) \in \Phi_i^u(b_1)$. If, on the other hand, $\phi_i^u(b) = v$, then we must have $d(u, v) \leq d(u, \phi_i^u(b_1))$, since $\alpha \in P(b, \rho(b), \phi(b))$, and so $\phi_i^u(b_1) \in \Phi_i^u(b_1)$.

We now want to show property (ii'). Consider a vertex $v_1 \in C_{b_1}$ with label $\ell(v_1)$ anchored to $\alpha(v_1) = v_2 \notin C_{b_1}$. If $v_2 = v$, then $v_2 = \rho_{i_v}^v(b_1)$, so the property holds for this case. If $v_2 \neq v$, then $v_2 \notin C_b$, so there exists $u \in S_b$ such that $v_2 = \rho_{\ell(v_1)-1}^u(b)$ as α belongs to $\mathcal{P}(b, \rho(b), \phi(b))$. By definition, if $\rho_{\ell(v_1)-1}^u(b) \neq \rho_{\ell(v_1)-1}^u(b_1)$, then $\rho_{\ell(v_1)-1}^u(b_1) = v$ and $d(u, v) < d(u, \rho_{\ell(v_1)-1}^u(b_1))$, which contradicts α being a minimum cost anchoring.

Property (iii') follows from the definition of $\rho(b_1)$ and from the fact that (LAP, α) belongs to $\mathcal{P}(b, \rho(b), \phi(b))$.

We thus obtain $A[b, \rho(b), \phi(b)] = c_{i_v} + A[b_1, \rho(b_1), \phi(b_1)]$, which proves the inequality.

- *Introduce nodes:* If b is an introduce node with respect to v , it has one child b_1 with $S_{b_1} = S_b \setminus \{v\}$ and $C_{b_1} = C_b$. Since $(\ell, \alpha) \in P(b, \rho(b), \phi(b))$, property (i') implies that for all i , we have $\phi_i^v(b) = \text{closest}_v\{\phi_i^u(b) \mid u \in S_{b_1}\}$.

Thus, the conditions that would cause $\bar{A}[b, \rho(b), \phi(b)]$ to be infinite are not met.

Consider a vertex $v_1 \in C_{b_1}$ with label $\ell(v_1)$ anchored to $\alpha(v_1) = v_2 \notin C_{b_1}$. There exists $u \in S_b$ such that $v_2 = \rho_{\ell(v_1)-1}^u(b)$, since α belongs to $\mathcal{P}(b, \rho(b), \phi(b))$. If $u \neq v$, we have the result. If $u = v$, then there exists a vertex $w \in S_{b_1}$ on the shortest path from v_1 to v . By property (iii') for $\rho(b)$, we know that $d(w, \rho_{\ell(v_1)-1}^w(b)) \leq d(w, \rho_{\ell(v_1)-1}^v(b))$. As α is a minimum cost anchoring, this inequality must be an equality, so we obtain property (ii').

Property (iii') holds for $\rho(b_j)$ as it is a subset of $\rho(b)$.

Claim. We have $A[b, \rho(b), \phi(b)] \leq \bar{A}[b, \rho(b), \phi(b)]$.

Consider a bag b and values $\rho(b)$ and $\phi(b)$. If $\bar{A}[b, \rho(b), \phi(b)]$ is infinite, the inequality trivially holds. We therefore consider the remaining cases. In particular, property (iii') is respected regarding $\rho(b)$, and the bag b has either one child b_1 or two children b_1 and b_2 . Consider, for $j = 1$ or $j \in \{1, 2\}$, the anchoring guarantees $\rho_j^u(b_j)$ and $\phi_j^u(b_j) \in \Phi_j^u(b_j)$ yielding the value $\bar{A}[b, \rho(b), \phi(b)]$. In particular, $\bar{A}[b, \rho(b), \phi(b)]$ depends on the value(s) $\bar{A}[b_j, \rho(b_j), \phi(b_j)]$. By the induction hypothesis, we have $\bar{A}[b_j, \rho(b_j), \phi(b_j)] = A[b_j, \rho(b_j), \phi(b_j)]$, so there exists a partial LAP $(\ell_j, \alpha_j) \in \mathcal{P}(b_j, \rho(b_j), \phi(b_j))$ on each C_{b_j} of cost $\bar{A}[b_j, \rho(b_j), \phi(b_j)]$.

Define (ℓ, α) to be the union of these LAPs in case of a join node. If b is a forget node, then extend (ℓ, α) to v , by choosing the label $\ell(v) = i_v \in \mathbf{I}_v$ (possibly ∞) which minimizes $\bar{A}[b, \rho(b), \phi(b)]$, as well as $\alpha(v) = v$ if $i_v = \infty$ and $\alpha(v) = \text{closest}_v\{\phi_{i_v-1}^v(b_1), \rho_{i_v-1}^v(b_1)\}$ otherwise. And lastly, if b is an introduce node, then simply keep the LAP since $C_b = C_{b_1}$. Thus, the cost of (ℓ, α) is precisely $\bar{A}[b, \rho(b), \phi(b)]$. We show that in all cases, the LAP (ℓ, α) belongs to $\mathcal{P}(b, \rho(b), \phi(b))$. Then, by defini-

tion of $A[b, \rho(b), \phi(b)]$, the inequality holds.

- *Join node:* Consider a join node b with children b_1, b_2 . Since $C_{b_1} \cap C_{b_2} = \emptyset$, the union of the LAPs is well defined.

Consider the anchoring guarantee $\phi_i^u(b)$ for some $u \in S_b$. We want to show that $\ell(\phi_i^u(b)) = i$ and that no vertex of C_b with label i is closer to u . If $\phi_i^u(b) \in C_{b_1}$, then by definition, $\phi_i^u(b_1) = \phi_i^u(b)$. Therefore, $\ell(\phi_i^u(b)) = i$, and no vertex in C_{b_1} closer to u is labeled i . We also know that $d(u, \phi_i^u(b)) \leq d(u, \phi_i^u(b_2))$, by condition (\star) in the definition of $\Phi_i^u(b_1)$. Therefore, by symmetry, property (i') holds for $\phi_i^u(b) \in C_{b_2}$ as well.

Regarding property (ii'), consider a vertex v_1 of either of the C_{b_j} 's anchored to a vertex $v_2 = \alpha(v_1) \notin C_b$. Then, there exists some $u \in S_b$ such that $v_2 = \rho_{\ell(v_1)-1}^u(b_j)$ by definition of $\mathcal{P}(b_j, \rho(b_j), \phi(b_j))$, and since $v_2 \notin C_b$, we must have $v_2 = \rho_{\ell(v_1)-1}^u(b)$.

- *Forget node:* Let b be a forget node with respect to v . That is, $S_{b_1} = S_b \cup \{v\}$ and $C_{b_1} = C_b \setminus \{v\}$. Clearly, (ℓ, α) is well defined.

Consider i and $u \in S_b$ (implying $u \neq v$), and assume first that $\phi_i^u(b) \neq v$. Then, we have $\phi_i^u(b) = \phi_i^u(b_1)$, so $\phi_i^u(b)$ is the closest vertex to u with label i in C_{b_1} . In order to show property (i') for this case, it remains to see that, if $i = i_v$, we have $d(u, v) \geq d(u, \phi_i^u(b))$. This directly follows from the definition of I_v . Assume now that $\phi_i^u(b) = v$. By the definition of $\Phi_i^u(b_1)$, we have $d(u, \phi_i^u(b_1)) \geq d(u, v)$. Then, by the definition of $\mathcal{P}(b_j, \rho(b_j), \phi(b_j))$, there is no vertex in C_{b_1} with label i closer to u than v . This implies that $\phi_i^u(b) = v$ is the closest vertex to u with label i in C_b .

We now prove property (ii'). Consider a vertex v_1 of C_{b_1} (so $v_1 \neq v$) with label $\ell(v_1)$ anchored to $v_2 = \alpha(v_1) \notin C_b$. There exists some $u \in S_{b_1}$ such that $v_2 = \rho_{\ell(v_1)-1}^u(b_1)$ by definition of $\mathcal{P}(b_1, \rho(b_1), \phi(b_1))$ and $v_2 \notin C_b$ so $v_2 \neq v$. Therefore, by definition of $\rho_{\ell(v_1)-1}^u(b_1)$, there must exist some $w \in S_b$ such that $v_2 = \rho_{\ell(v_1)-1}^w(b)$. Consider now v . If $i_v = \infty$ (which can only be the case if $v \notin \mathcal{X}$), we defined its anchor to be $\alpha(v) = v$. If $i_v \neq \infty$, we defined $\alpha(v) = \text{closest}_v\{\phi_{i_v-1}^v(b_1), \rho_{i_v-1}^v(b_1)\}$. If $\alpha(v) \notin C_b$, then $\alpha(v) = \text{closest}_v\{\rho_{i_v-1}^u(b) \mid u \in S_b\}$ by definition of $\rho_{i_v-1}^v(b_1)$, which completes the proof of property (ii').

- *Introduce node:* Consider an introduce node b with respect to v . That is, b has one child b_1 such that $S_{b_1} = S_b \setminus \{v\}$, $C_{b_1} = C_b$. Again, (ℓ, α) is obviously well defined. For each $i < k$ and $u \in S_{b_1}$, we have $\phi_i^u(b) = \phi_i^u(b_1)$, and we have $\phi_i^v(b) \neq \text{closest}_v\{\phi_i^u(b) \mid u \in S_{b_1}\}$. As any path between v and a vertex in S_b contains a vertex in S_{b_1} , property (i') is satisfied.

Consider a vertex v_1 of C_b with label $\ell(v_1)$ anchored to $v_2 = \alpha(v_1) \notin C_b$. There exists some $u \in S_{b_1}$ with $v_2 = \rho_{\ell(v_1)-1}^u(b_1)$ by definition of $\mathcal{P}(b_1, \rho(b_1), \phi(b_1))$. Since $S_{b_1} \subset S_b$, property (ii') holds.

Thus, for all types of bags, (ℓ, α) is a member of $\mathcal{P}(b, \rho(b), \phi(b))$. Since $A[b, \rho(b), \phi(b)]$ gives the minimum cost of all such LAPs, the claim follows. \square

6. The k -hop M \check{S} T problem with relaxed hop constraints. In this section, we consider the k -hop M \check{S} T problem in the relaxed model where an algorithm may output a Steiner tree that uses more than k hops while still comparing its performance to the optimal k -hop M \check{S} T. On the positive side, we look at metrics of bounded highway dimension and present an M \check{S} T of near-optimal cost that violates the hop

constraint by at most one hop. We further show that our reasoning yields an analogous result for metrics of bounded doubling dimension. Note that this is by no means a given as Abraham et al. [1] showed that constant doubling dimension does not imply constant highway dimension (whereas the converse is true). We also comment on the relaxation of the hop constraint of k -hop MST in the case of general distance functions induced by some graph, which are not guaranteed to correspond to a metric. This setting does not admit a polynomial-time constant-factor approximation for k -hop MST even with all distances being equal to 1. Extending the hardness reduction from [39], we show that the problem does not admit a constant-factor approximation even when the Steiner tree can use $k + \ell$ hops for a constant ℓ .

6.1. Bounded highway dimension. Feldmann et al. [22] defined the highway dimension of a graph as follows. Let $B_r(v) = \{u \in V \mid d(u, v) \leq r\}$. Given a universal constant $c > 4$, the *highway dimension* of a graph G is the smallest integer h such that for every $r \geq 0$ and $v \in V$, there is a set of h vertices in $B_{cr}(v)$ that hits all shortest paths of length more than r that lie entirely in $B_{cr}(v)$. Before stating the results, we first define a δ -net of a graph, which is informally a subset of vertices which are far from each other, while every vertex in the graph is close to this subset. Formally, a δ -net of a graph G is a subset U of V such that for all $u \in V$, there exists $v \in U$ with $d(u, v) \leq \delta$ and for all $u, v \in U$, we have $d(u, v) > \delta$.

Note that in the literature, there are definitions of highway dimensions different from the one above, both more general [2] and more restricted [1]. For further discussion, see [22].

Allowing the use of an additional hop, we can extend our algorithm for metrics with bounded treewidth to metrics of bounded highway dimensions and obtain the following result.

THEOREM 6.1. *For a metric induced by a graph of bounded highway dimension and a constant k , let OPT_k be the cost of a k -hop MST. A $(k + 1)$ -hop Steiner tree of cost at most $(1 + \varepsilon)OPT_k$, for $\varepsilon > 0$, can be computed in quasi-polynomial time.*

Feldmann et al. [22] prove the following theorem, which gives sufficient conditions for a problem to admit a QPTAS on graphs of constant highway dimension. (We use the generic form of the theorem, which is not explicitly stated in [22] but fully argued in the text.)

THEOREM 6.2 (reformulation of [22, Theorem 8.1]). *For a graph G of constant highway dimension and a problem \mathcal{P} satisfying conditions 1–6 below, a $(1 + \varepsilon)$ -approximation can be computed in quasi-polynomial time.*

1. *An optimum solution of \mathcal{P} can be computed in time $n^{O(\omega)}$ for graphs with treewidth ω .*
2. *A constant-approximation of \mathcal{P} on metric graphs can be computed in polynomial time.*
3. *The diameter of the graph can be assumed to be $O(n \cdot OPT_G)$, where OPT_G is the cost of an optimal solution in G .*
4. *An optimum solution for \mathcal{P} on a δ -net U has cost at most $OPT_G + O(n\delta)$.*
5. *The objective function of \mathcal{P} is linear in the edge cost.*
6. *A solution for \mathcal{P} on a δ -net U can be converted to a solution on V for an additional cost of $O(n\delta)$.*

We now show that our main result, Theorem 5.2, together with a previously known result, leads to Theorem 6.1, by using a slight variation of Theorem 6.2 to allow the extra hop.

Proof of Theorem 6.1. Applying Theorem 6.2, it remains to verify that the k -hop MŠT problem satisfies its six conditions, for k constant, if we allow the algorithm to use one more hop (i.e., computing a $(k + 1)$ -hop Steiner tree) than the optimal solution of cost OPT_k to which we compare it. The conditions and the explanation of why they are fulfilled are detailed below.

1. Theorem 5.2 states precisely this condition for k -hop MŠT.
2. For k -hop MŠT in metric graphs, Kantor and Peleg [33] presented an algorithm with approximation factor $(1.52 \cdot 9^{k-2})$.
3. The diameter of G can be assumed to be $O(n \cdot OPT_k)$ since edges of cost larger than $1.52 \cdot 9^{k-2} \cdot OPT_k$ can be deleted after computing the approximation of OPT_k from [33].
4. Consider an optimum k -hop MŠT on V , and move each vertex not in U to the closest vertex in U . This induces an extra cost of $O(n\delta)$ and is a solution on U .
5. The objective function of k -hop MŠT is indeed linear in the edge cost.
6. This condition requires an additional hop. We claim that a solution for k -hop MŠT on a δ -net U can be converted to a $(k + 1)$ -hop Steiner tree on V for an additional cost of $O(n\delta)$. Indeed, given a k -hop Steiner tree on U , we can anchor all vertices from $V \setminus U$ to their closest vertex in U for an additional cost of $O(n\delta)$ and obtain a $(k + 1)$ -hop Steiner tree. This procedure of extending a solution is performed exactly once in the underlying algorithm. Therefore, we can allow the algorithm to use one more hop on G than the solution on U . Note that this property is not stated explicitly in [22]. □

6.2. Bounded doubling dimension. The *doubling dimension* d of a graph G refers to the smallest integer d such that any ball $B_{2r}(v)$ of radius $2r$ is contained in the union of $2d$ balls of radius r .

To show Theorem 6.2, Feldmann et al. [22] construct a probabilistic embedding of metrics of bounded highway dimension into metrics with bounded treewidth that maintains distances approximately in expectation. They build upon the work of Talwar [42], who gives an analogous result for metrics of bounded doubling dimension, and use it as a building block for their embedding algorithm. However, as discussed in subsection 1.2 and [22], the two dimensional parameters are not directly linked.

The proof of Theorem 6.2 fundamentally consists of replacing the given metric space by an appropriately chosen δ -net, to limit the aspect ratio of the instance, and then using an algorithm for bounded treewidth graphs on the corresponding probabilistic embedding. Thus, we may use Talwar’s embedding instead and replace “highway dimension” by “doubling dimension” in the statement of Theorem 6.2. This yields the following.

THEOREM 6.3. *For a metric of bounded doubling dimension and a constant k , let OPT_k be the cost of a k -hop MŠT. A $(k + 1)$ -hop Steiner tree of cost at most $(1 + \varepsilon)OPT_k$, for $\varepsilon > 0$, can be computed in quasi-polynomial time.*

6.3. Nonmetric distances. We have shown that relaxing the hop constraint only by one hop leads to efficient algorithms for some rather general metrics, such as metrics with bounded highway dimension or bounded doubling dimension.

We contrast these results by showing that even permitting ℓ additional hops, for any constant ℓ , is not enough for the most general setting, which comprises non-metric graphs with no constraints on the nonnegative edge weights. Specifically, the distance between two vertices may be infinite, so a direct connection may not




FIG. 8. Reduction from UNWEIGHTED SET COVER to k -hop MŠT with relaxed hop constraints. The corresponding set cover universe is $X = \{x_1, \dots, x_5\}$ with sets $Y_1 = \{x_1, x_2, x_4\}$, $Y_2 = \{x_1, x_3, x_4\}$, $Y_3 = \{x_3, x_5\}$ and $Y_4 = \{x_4, x_5\}$. The vertices x_1, \dots, x_5 are set to be terminals. All edges in the graph have length 1.

exist. We generalize a construction by Manyem and Stallmann [39] for showing an inapproximability result for k -hop MŠT and prove the following result for k -hop MŠT with relaxed hop constraints.

THEOREM 6.4. *For any constant c and ℓ , and given a weighted graph G as input, it is NP-hard to find a $(k + \ell)$ -hop Steiner tree in G of cost $(1 - c) \cdot \log n \cdot OPT$, where OPT is the minimal cost of a k -hop Steiner tree in G . The above statement is true even when all edges in the graph G have weight equal to 1.*

Proof. We reduce from UNWEIGHTED SET COVER. There, we are given a family of sets $\mathcal{Y} = \{Y_1, \dots, Y_m\}$ over a shared universe $X = \{x_1, \dots, x_n\}$, and we are tasked with deciding whether there is a selection $\mathcal{S} \subseteq \mathcal{Y}$ of at most s sets such that $\bigcup_{Y_j \in \mathcal{S}} Y_j = X$. We start our reduction by representing the problem as a bipartite graph, with edges $\{x_i, y_j\}$ corresponding to $x_i \in Y_j$. Next, we add a new vertex r to serve as our root, and we set every vertex x_i to be a terminal of the Steiner tree. We connect r to each vertex y_j by a unique path of t edges of weight 1 each. Finally, we subdivide each edge $\{x_i, y_j\}$ to form a path of ℓ edges, again each single edge having weight exactly 1. See Figure 8 for an example.

Let $k = t + \ell$. We first observe that if there is a solution \mathcal{S} to the UNWEIGHTED SET COVER problem of size s , then there exists a k -hop Steiner tree of total cost $st + n\ell$. Such a tree can be created by including the y_j corresponding to $Y_j \in \mathcal{S}$ in the Steiner tree, and connecting any x_i to exactly one of the y_j for which $x_i \in Y_j \in \mathcal{S}$ paying a cost of ℓ each. Finally, connect the selected y_j via the direct paths of cost t to the root r .

Next, we consider any $k + (2\ell - 1)$ -hop Steiner tree T on this instance and a node y_j that is included in the Steiner tree. In this case, y_j must be connected directly to the root by the path of weight t . This is true as the shortest hop distance between any two vertices $y_j, y_{j'}$ is at least 2ℓ . Since the number of hops to connect some $y_{j'}$

to the root is at least t , connecting y_j via $y_{j'}$ would contradict the assumption that we are given a $t + (2\ell - 1)$ -hop Steiner tree. It follows that $st + n\ell$ indeed is also the minimum cost of a Steiner tree with at most $k + (2\ell - 1)$ hops that contains s of the y_j .

With the above, we see that any Steiner tree with at least k hops and at most $k + (2\ell - 1)$ hops directly corresponds to a solution \mathcal{S} of the UNWEIGHTED SET COVER problem that consists of exactly the sets Y_j for which vertex y_j is included in the Steiner tree. Furthermore, any such Steiner tree with at most $k + (2\ell - 1)$ hops and of cost at most $st + n\ell$ corresponds to a set cover with at most s sets.

Assume we can obtain a $k + \ell$ -hop Steiner tree of cost at most $(1 - c) \log(n) \cdot \text{OPT}$. Denoting by s^* the size of an minimum set cover, we have that this equals $(1 - c) \log(n) \cdot (ts^* + \ell n) = (1 - \frac{c}{2}) \log(n) \cdot ts^* + (1 - c) \log(n) \cdot \ell n - \frac{c}{2} \log(n) \cdot ts^*$, which is at most $(1 - \frac{c}{2}) \log(n) \cdot ts^* + \ell n$ when substituting $t = \lceil \frac{\ell n((1-c)\log(n)-1)}{\frac{c}{2}s^* \log n} \rceil$, which is polynomial in n . Thus, the Steiner tree corresponds to a set cover of size at most $(1 - \frac{c}{2}) \log(n) \cdot s^*$.

Summing up, we gave a polynomial-size reduction that implies that given a $(1 - c) \log(n)$ approximation for the k -hop MST problem, we obtain a $(1 - c')$ approximation for set cover, with $c, c' > 0$. The latter, however, is known to be NP-hard [40], concluding the proof. \square

7. Conclusions. In this work, we show how to solve the k -hop MST problem exactly and efficiently in tree-like metrics and extend our results for other metrics, such as bounded highway dimension and doubling metrics. In our extensions, we have relaxed the constraint on the number of hops k to $k + 1$; it remains open whether this relaxation is needed for the latter two classes. While soft hop constraints, at least with an additive constraint, do not help in the nonmetric setting, it would be interesting to know whether one can obtain a constant-factor approximation for the k -hop MST problem in arbitrary metrics under a soft hop constraint, or even an approximation scheme, which is ruled out for hard hop constraints [4, 26]. Further, our exact algorithms raise the question of whether the problem of finding a k -hop MST in a metric induced by a tree is fixed-parameter tractable by the parameter k , i.e., whether there exists an algorithm with running time $O(\text{poly}(n) \cdot f(k))$ for some function f .

Acknowledgment. We thank Jiří Sgall for fruitful discussions on the k -hop MST problem in path metrics.

REFERENCES

[1] I. ABRAHAM, D. DELLING, A. FIAT, A. V. GOLDBERG, AND R. F. WERNECK, *Highway dimension and provably efficient shortest path algorithms*, J. ACM, 63 (2016), 41.
 [2] I. ABRAHAM, D. DELLING, A. FIAT, A. V. GOLDBERG, AND R. F. F. WERNECK, *VC-dimension and shortest path algorithms*, in ICALP 2011: Automata, Languages and Programming, Lecture Notes in Comput. Sci. 6755, Springer, Berlin, 2011, pp. 690–699.
 [3] I. ABRAHAM, A. FIAT, A. V. GOLDBERG, AND R. F. WERNECK, *Highway dimension, shortest paths, and provably efficient algorithms*, in Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, ACM, New York, SIAM, Philadelphia, 2010, pp. 782–793, <https://doi.org/10.1145/2985473>.
 [4] L. ALFANDARI AND V. T. PASCHOS, *Approximating minimum spanning tree of depth 2*, Internat. Trans. Oper. Res., 6 (1999), pp. 607–622.
 [5] E. ALTHAUS, S. FUNKE, S. HAR-PELED, J. KÖNEMANN, E. A. RAMOS, AND M. SKUTELLA, *Approximating k-hop minimum-spanning trees*, Oper. Res. Lett., 33 (2005), pp. 115–120.

- [6] O. ANGEL, A. D. FLAXMAN, AND D. B. WILSON, *A sharp threshold for minimum bounded-depth and bounded-diameter spanning trees and Steiner trees in random networks*, *Combinatorica*, 32 (2012), pp. 1–33.
- [7] S. ARORA, C. LUND, R. MOTWANI, M. SUDAN, AND M. SZEGEDY, *Proof verification and the hardness of approximation problems*, *J. ACM*, 45 (1998), pp. 501–505.
- [8] S. ARYA, G. DAS, D. M. MOUNT, J. S. SALOWE, AND M. SMID, *Euclidean spanners: Short, thin, and lanky*, in *Proceedings of the Twenty-Seventh Annual ACM Symposium on Theory of Computing*, ACM, New York, 1995, pp. 489–498.
- [9] A. BALAKRISHNAN AND K. ALTINKEMER, *Using a hop-constrained model to generate alternative communication network design*, *INFORMS J. Comput.*, 4 (1992), pp. 192–205.
- [10] J. BAR-ILAN, G. KORTSARZ, AND D. PELEG, *Generalized submodular cover problems and applications*, *Theoret. Comput. Sci.*, 250 (2001), pp. 179–200.
- [11] M. BERN AND P. PLASSMANN, *The Steiner problem with edge lengths 1 and 2*, *Inform. Process. Lett.*, 32 (1989), pp. 171–176.
- [12] A. BJÖRKLUND, T. HUSFELDT, P. KASKI, AND M. KOIVISTO, *Fourier meets Möbius: Fast subset convolution*, in *Proceedings of STOC*, ACM, New York, 2007, pp. 67–74.
- [13] H. L. BODLAENDER, *A linear-time algorithm for finding tree-decompositions of small treewidth*, *SIAM J. Comput.*, 25 (1996), pp. 1305–1317, <https://doi.org/10.1137/S0097539793251219>.
- [14] P. CARMİ, L. CHAITMAN-YERUSHALMI, AND O. TRABELSI, *Bounded-hop communication networks*, *Algorithmica*, 80 (2018), pp. 3050–3077.
- [15] M. CHIMANI, P. MUTZEL, AND B. ZEY, *Improved Steiner tree algorithms for bounded treewidth*, *J. Discrete Algorithms*, 16 (2012), pp. 67–78.
- [16] M. CHIMANI AND J. SPOERHASE, *Network design problems with bounded distances via shallow-light Steiner trees*, in *Proceedings of the 32nd International Symposium on Theoretical Aspects of Computer Science (STACS 2015)*, LIPIcs. Leibniz Int. Proc. Inform. 30, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, Germany, 2015, pp. 238–248.
- [17] M. CYGAN, F. V. FOMIN, Ł. KOWALIK, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Treewidth*, in *Parameterized Algorithms*, Springer, Cham, 2015, pp. 151–244.
- [18] G. DAHL, L. GOUVEIA, AND C. REQUEJO, *On formulations and methods for the hop-constrained minimum spanning tree problem*, in *Handbook of Optimization in Telecommunications*, M. G. C. Resende and P. M. Pardalos, eds., Springer, Boston, MA, 2006, pp. 493–515.
- [19] S. E. DREYFUS AND R. A. WAGNER, *The Steiner problem in graphs*, *Networks*, 1 (1971), pp. 195–207.
- [20] M. ELKIN AND S. SOLOMON, *Narrow-shallow-low-light trees with and without Steiner points*, *SIAM J. Discrete Math.*, 25 (2011), pp. 181–210, <https://doi.org/10.1137/090776147>.
- [21] J. FAKCHAROENPHOL, S. RAO, AND K. TALWAR, *A tight bound on approximating arbitrary metrics by tree metrics*, *J. Comput. System Sci.*, 69 (2004), pp. 485–497.
- [22] A. E. FELDMANN, W. S. FUNG, J. KÖNEMANN, AND I. POST, *A $(1 + \epsilon)$ -embedding of low highway dimension graphs into bounded treewidth graphs*, *SIAM J. Comput.*, 47 (2018), pp. 1667–1704, <https://doi.org/10.1137/16M1067196>.
- [23] L. GOUVEIA, *Using the Miller-Tucker-Zemlin constraints to formulate a minimal spanning tree problem with hop constraints*, *Comput. Oper. Res.*, 22 (1995), pp. 959–970.
- [24] L. GOUVEIA, L. SIMONETTI, AND E. UCHOA, *Modeling hop-constrained and diameter-constrained minimum spanning tree problems as Steiner tree problems over layered graphs*, *Math. Program.*, 128 (2011), pp. 123–148.
- [25] R. L. GRAHAM AND P. HELL, *On the history of the minimum spanning tree problem*, *IEEE Ann. Hist. Comput.*, 7 (1985), pp. 43–57.
- [26] S. GUHA AND S. KHULLER, *Greedy strikes back: Improved facility location algorithms*, *J. Algorithms*, 31 (1999), pp. 228–248.
- [27] A. GUPTA, R. KRAUTHGAMER, AND J. R. LEE, *Bounded geometries, fractals, and low-distortion embeddings*, in *Proceedings of FOCS*, IEEE, Washington, DC, 2003, pp. 534–543.
- [28] M. HAENGGI, *Twelve reasons not to route over many short hops*, in *Proceedings of the 60th IEEE Vehicular Technology Conference (VTC2004-Fall)* (Los Angeles, CA), IEEE, Washington, DC, 2004, pp. 3130–3134.
- [29] B. HAEUPLER, D. E. HERSHKOWITZ, AND G. ZUZIC, *Tree embeddings for hop-constrained network design*, in *Proceedings of the 53rd Annual ACM SIGACT STOC*, ACM, New York, 2021, pp. 356–369.
- [30] M. T. HAJIAGHAYI, G. KORTSARZ, AND M. R. SALAVATIPOUR, *Approximating buy-at-bulk and shallow-light k -Steiner trees*, *Algorithmica*, 53 (2009), pp. 89–103.
- [31] F. K. HWANG, D. S. RICHARDS, AND P. WINTER, *The Steiner Tree Problem*, *Ann. Discrete Math.* 53, North-Holland, Amsterdam, 1992.

- [32] V. JARNÍK AND M. KÖSSLER, *O minimálních grafech, obsahujících n daných bodů*, Časopis pro pěstování matematiky a fyziky, 63 (1934), pp. 223–235.
- [33] E. KANTOR AND D. PELEG, *Approximate hierarchical facility location and applications to the bounded depth Steiner tree and range assignment problems*, J. Discrete Algorithms, 7 (2009), pp. 341–362.
- [34] R. M. KARP, *Reducibility among combinatorial problems*, in Complexity of Computer Computations, IBM Research Symposia Series, Plenum Press, New York, 1972, pp. 85–103.
- [35] B. KORTE AND J. NEŠETŘIL, *Vojtěch Jarník's work in combinatorial optimization*, Discrete Math., 235 (2001), pp. 1–17.
- [36] G. KORTSARZ AND D. PELEG, *Approximating shallow-light trees*, in Proceedings of the Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, ACM, New York, SIAM, Philadelphia, 1997, pp. 103–110.
- [37] S. LAUE AND D. MATIJEVIĆ, *Approximating k -hop minimum spanning trees in Euclidean metrics*, Inform. Process. Lett., 107 (2008), pp. 96–101.
- [38] S. LI, *A 1.488 approximation algorithm for the uncapacitated facility location problem*, Inform. and Comput., 222 (2013), pp. 45–58.
- [39] P. MANYEM AND M. F. STALLMANN, *Some Approximation Results in Multicasting*, Tech. report, North Carolina State University at Raleigh, Raleigh, NC, 1996.
- [40] D. MOSHKOVITZ, *The projection games conjecture and the NP-hardness of \ln n -approximating set-cover*, in APPROX-RANDOM, Lecture Notes in Comput. Sci. 7408, Springer, 2012, pp. 276–287.
- [41] V. R. SAKSENA, *Topological analysis of packet networks*, IEEE J. Selected Areas Commun., 7 (1989), pp. 1243–1252.
- [42] K. TALWAR, *Bypassing the embedding: Algorithms for low dimensional metrics*, in Proceedings of STOC, ACM, New York, 2004, pp. 281–290.
- [43] S. VOSS, *The Steiner tree problem with hop constraints*, Ann. Oper. Res., 86 (1999), pp. 321–345.