

Approximating Bottleneck Spanning Trees on Partitioned Tuples of Points

Ahmad Biniiaz ✉🏠

School of Computer Science, University of Windsor, Canada

Anil Maheshwari ✉🏠

School of Computer Science, Carleton University, Canada

Michiel Smid ✉🏠

School of Computer Science, Carleton University, Canada

Abstract

We present approximation algorithms for the following NP-hard optimization problems related to bottleneck spanning trees in metric spaces.

1. The *disjoint bottleneck spanning tree* problem: Given n pairs of points in a metric space, find two disjoint trees each containing exactly one point from each pair and minimize the largest edge length (over all edges of both trees). It is known that approximating this problem by a factor better than 2 is NP-hard. We present a 4-approximation algorithm for this problem. This improves upon the previous best known approximation ratio of 9. Our algorithm extends to a $(3k - 2)$ -approximation for a more general case where points are partitioned into k -tuples and we seek k disjoint trees.
2. The *generalized bottleneck spanning tree* problem: Given n points in some metric space that are partitioned into clusters of size at most 2, find a tree that contains exactly one point from each cluster and minimizes the largest edge length. We show that it is NP-hard to approximate this problem by a factor better than 2, and present a 3-approximation algorithm.
3. The *partitioned bottleneck spanning tree* problem: Given kn points in some metric space, find k disjoint trees each containing exactly n points and minimize the largest edge length (over all edges of the k trees). We show that it is NP-hard to approximate this problem by a factor better than 2 for any $k \geq 2$. We present an α -approximation algorithm for this problem where $\alpha = 2$ for $k = 2, 3$ and $\alpha = 3$ for $k \geq 4$. Towards obtaining these approximation ratios we present tight upper bounds on the edge lengths of k equal-size disjoint trees that can be obtained from the nodes of a given tree. This result is of independent interest.

Our hardness proofs imply that it is NP-hard to approximate the non-metric version of the above problems within any constant factor. If we seek traveling salesperson tours (instead of trees) then our algorithms simply extend to achieve approximate solutions with factors three times those mentioned above.

Keywords and phrases bottleneck spanning trees, metric space, partitioned point sets, approximation algorithms

Digital Object Identifier 10.57717/cgt.v1i1.17

Funding *Ahmad Biniiaz*: Supported by NSERC

Anil Maheshwari: Supported by NSERC


Michiel Smid: Supported by NSERC

1 Introduction

Spanning trees are fundamental structures in graph theory and combinatorics. The problem of finding spanning trees with enforced properties has received considerable attention from both theoretical and practical points of view. For example, the minimum spanning tree (MST) problem asks for a spanning tree with minimum total edge-length, and the bottleneck



spanning tree (BST) problem asks for a spanning tree whose largest edge-length is minimum. Besides their interesting theoretical properties, these problems find applications in the design of networks, including computer networks, wireless networks, and transportation networks, to name a few. Bottleneck spanning trees in particular are important in designing telecommunications networks with short connections (edges). Short connections are desirable in many ways because they require lower transmission ranges, are more secure, and cause less interference. This paper addresses three closely related bottleneck spanning tree problems (illustrated in Figure 1):



■ **Figure 1** Illustration of the problems for $k = 2$; black and white squares/circles represent tuples.

1. The *disjoint bottleneck spanning tree* (k -DBST) problem: Given kn points in some metric space that are partitioned into k -tuples, find k disjoint trees each containing exactly one point from each tuple and minimize the largest edge length (over all edges of the k trees).
2. The *generalized bottleneck spanning tree* (k -GBST) problem: Given n points in some metric space that are partitioned into clusters of size at most k , find a tree that contains exactly one point from each cluster and minimizes the largest edge length. The term “spanning” refers to span all clusters.
3. The *partitioned bottleneck spanning tree* (k -PBST) problem: Given kn points in some metric space, find k disjoint trees each containing exactly n points and minimize the largest edge length (over all edges of the k trees).

The above problems are natural generalizations of the standard BST problem. For $k = 1$, all above problems are equivalent to the BST problem which can be solved optimally in polynomial time [10]. For $k \geq 2$, all above problems are NP-hard and cannot be approximated by a factor better than 2 unless $P = NP$ (this will become clear shortly). The focus of this paper is on $k \geq 2$. We first present constant-factor approximation algorithms for $k = 2$. Then we extend some of our algorithms for larger k .

Notice that the k -DBST, k -GBST, and k -PBST problems are defined (above) with respect to the metric space. Thus, all our references to previous works and results apply to the metric space (unless it is explicitly mentioned to be for the non-metric space).

1.1 Some related works and applications

The problems introduced above find real-world applications that we put into context together with some related works. In our description we implicitly assume that k is at least 2.

(1) The k -DBST problem is introduced by Arkin et al. [2]. Motivated by the problem of maintaining secure connectivity in networks involving replicated data, Arkin et al. [2] introduced a class of problems that ask for k disjoint structures (trees, cycles, matchings) each containing one point from every given k -tuple. In particular they studied these problems for $k = 2$. Among many interesting results they presented a 9-approximation algorithm for the 2-DBST problem and an 18-approximation algorithm for computing two disjoint

traveling salesperson tours (instead of trees). It is easily seen, from their Lemma 8, that the 9-approximation algorithm can be extended to achieve a $(6k - 3)$ -approximation for the k -DBST problem. Although some of the results of Arkin et al. [2] have been improved by Johnson [24], their ratios 9 and 18 are still the best known. As for the lower bound, Johnson [24] showed that it is NP-hard to approximate the 2-DBST problem by a factor less than 2.

(2) The k -GBST problem is closely related to the k -generalized minimum spanning tree (k -GMST) problem, introduced by Myung et al. [32]. The k -GMST problem asks for a tree that contains exactly one point from each cluster and minimizes the total-edge length. This problem is well studied (see e.g. the recent survey by Pop [36] and references therein). The k -GMST problem is NP-hard even for $k = 2$ in the Euclidean plane. Even a more restricted version where the two points in each cluster have the same x or y coordinates is NP-hard [16, 19, 25]. The metric version of the k -GBST can be approximated by a ratio of $2k$ using linear programming [38] combined with the so-called parsimonious property [21]. Related work [7, 36, 37] also addresses the generalized traveling salesperson problem (TSP) in which the tour must contain exactly one point from each cluster. The group Steiner tree is another related problem which asks for a shortest tree that contains *at least* one point from each cluster. The non-metric versions of both the k -GMST and the group Steiner tree problems are NP-hard and cannot be approximated within any constant factor [23, 32]. Gabow et al. [20] studied the problem of finding a path, from a source to a destination in a graph, that passes through at most one vertex from every given pair of vertices. Arkin et al. [4] studied the multiple-choice minimum-diameter problem which is to select at least one element from each cluster to minimize the diameter of the chosen elements. The k -GBST also lies in the concept of *imprecision* in computational geometry where each input point is provided as a region of uncertainty (also known as *neighborhood*) and the exact position of the point may be anywhere in the region; see e.g. [8, 11, 17, 29, 30, 31].

Both the k -GBST and the k -GMST have real-world applications for example in the field of telecommunications, designing metropolitan area networks, interconnecting local area networks, and determining location of regional service centers (e.g., stores, warehouses, agricultural settings, distribution centers). For a detailed explanation of these applications and for more examples we refer the interested reader to the paper of Myung et al. [32] and the recent survey by Pop [36].

(3) The k -PBST problem falls in the class of partitioning a set into subsets such that the substructures (computed on subsets) are balanced. Balanced partitioning of the input has a long history and gives rise to interesting theoretical problems. For example in the k -partition traveling salesperson problem we are given k salespersons and the goal is to visit every city by exactly one salesperson and minimize the distance traveled by the salesperson making the longest journey [5, 6, 35].

The problem of k -balanced partitioning of a graph asks for partitioning the vertices of the graph into k subsets such that the induced subgraph on each subset is connected and the maximum cardinality of the subsets is minimized. Dyer and Frieze [18] showed that this problem is NP-hard; they also showed the hardness of many variations of this problem. Chlebíková [14] presented constant-factor approximations for $k = 2, 3$, and Chen et al. [13] presented a $k/2$ -approximation for $k \geq 4$. The max-min version of this problem is also studied [14, 40].

Motivated by a problem from the shipbuilding industry, Andersson et al. [1] studied the k -partition minimum spanning tree (k -PMST) problem where the goal is to partition an input point set into k subsets such that the length of the longest MST on the subsets

is minimized. As noted in [27] (and references therein) the k -PMST problem also arises in multi-vehicle scheduling, task sequencing, and political districting. Andersson et al. [1] showed that the k -PMST problem is NP-hard even for $k = 2$ in the Euclidean metric in the plane, and presented $(4/3 + \epsilon)$ and $(2 + \epsilon)$ approximations for $k = 2$ and $k \geq 3$, respectively. Karakawa et al. [27] studied this problem in higher dimensions. The k -PMST problem has also been studied in trees and cactus graphs under the name “minmax subtree cover” problem [33, 34].

1.2 Our contributions

We study the k -DBST, k -GBST, and k -PBST problems in metric spaces (where distances satisfy the triangle inequality). We show the hardness as well as approximation algorithms for these problems. We present our results for the simplest version where $k = 2$ (as it is easier to understand) and then extend them for larger k . Some of our results (such as Theorems 2, 7, and 11) are of independent interest and present self-contained combinatorial arguments.

- The 2-DBST problem is known [24] to be NP-hard and inapproximable by a factor better than 2. We present a 4-approximation algorithm for this problem. This improves the previous best known ratio of 9 due to Arkin et al. [2]. We extend our algorithm and achieve a $(3k - 2)$ -approximation for the k -DBST for any $k \geq 2$ (Theorem 3). To achieve these algorithms we demonstrate a set labeling technique (Theorem 2) as an implication of Hall’s marriage theorem and a theorem of Kőnig [22].
- The difficulty of the 2-GBST problem lies in choosing representative points from clusters; once these points are selected, the problem is reduced to the standard BST problem. We show that it is NP-hard to approximate the 2-GBST problem by a factor better than 2 using a reduction from 3-SAT (Theorem 4), and present a 3-approximation algorithm for this problem (Theorem 6). In some part of our algorithm we show the following result which is of independent interest (Theorem 7): Given a tree T_1 and a partitioning of its nodes into clusters of size at most two, we can obtain a tree T_2 that contains exactly one node from each cluster and the length of its edges is at most 3 in the metric¹ of T_1 ; the upper bound 3 is the best achievable.
- We show that it is NP-hard to approximate the k -PBST problem by a factor less than 2 for any $k \geq 2$ (Theorem 8) using a reduction from the 2-balanced partitioning of a graph [18]. We present an α -approximation algorithm for this problem (Theorem 10) where $\alpha = 2$ for $k = 2, 3$ and $\alpha = 3$ for $k \geq 4$. Towards obtaining these approximation ratios we present tight upper bounds on the edge lengths of k equal-size disjoint trees that can be obtained from the nodes of a given tree (Theorem 11); the edge lengths are measured in the metric of the given tree. This result is of independent interest.

A straightforward implication of our hardness proofs and that of Johnson [24] is that the non-metric versions of the above problems cannot be approximated within any constant factor.

Extension to bottleneck TSP tours. If instead of trees in the above problems we seek TSP tours that minimize the largest edge length, then our algorithms simply extend to obtain approximate solutions with factors that are three times those for bottleneck trees. This can

¹ In this metric the distance between two nodes u and v in a tree T is the number of edges in the unique path between them in T .

be done via a known result that the *cube*² of every connected graph has a Hamiltonian cycle, and such a cycle can be computed in polynomial time [26, 28]; this is also hinted in [15, Exercise 37.2.3]. To use this result, we first obtain an α -approximate solution, namely \mathcal{B} , for the corresponding BST problem (using our BST algorithms) and then we find TSP tours, namely \mathcal{T} , in the cube of \mathcal{B} . By the triangle inequality the largest edge-length in the cube graph, and in particular in \mathcal{T} , is at most three times the largest edge-length in \mathcal{B} . Notice that in all above problems the largest edge length in any optimal BST solution is a lower bound for the largest edge length in any optimal TSP solution. Thus \mathcal{T} would be a 3α -approximate solution for the TSP. For example our 4-approximation algorithm for the 2-DBST can be extended to obtain a 12-approximation for two disjoint TSP tours that minimize the largest edge length; this improves the previous approximation ratio of 18 due to Arkin et al. [2].

Notation. The largest edge length in a tree T is referred to as the *bottleneck* of T and is denoted by $\lambda(T)$. We denote the distance between two points p and q in a metric space by $|pq|$. Conceptually, a point set P in a metric space can be viewed as a *metric graph*, i.e., as a complete edge-weighted graph with vertex set P where the weight $w(e)$ of each edge $e = (p, q)$ is equal to the distance between p and q , that is $w(e) = |pq|$.

2 The k -DBST problem

Let $k \geq 2$ be an integer. In this section we present an approximation algorithm for the k -DBST problem: Given kn points in some metric space that are partitioned into k -tuples, we want to find k disjoint trees each containing exactly one point from each tuple and minimize the largest edge length (over all the k trees). We first present our approximation algorithm for $k = 2$ as it is easier to understand. Then we extend the algorithm to larger k . Our algorithm benefits from the following remarkable result of König which is stated in [22].

▶ **Theorem 1** (König, 1916). *Let S be any set with kn elements that is partitioned, in two different ways, into n subsets each with k elements, namely A_1, \dots, A_n and B_1, \dots, B_n . Then there exist n elements of S , namely r_1, \dots, r_n , and a permutation π of $\{1, \dots, n\}$ such that $r_i \in A_i \cap B_{\pi(i)}$ for all $i \in \{1, \dots, n\}$.*

Example. Let $k = 3$, $n = 4$, $S = \{1, 2, \dots, 12\}$, and consider two partitions of S

$$A_1 = \{1, 2, 3\}, A_2 = \{4, 5, 6\}, A_3 = \{7, 8, 9\}, A_4 = \{10, 11, 12\}$$

$$B_1 = \{4, 9, 12\}, B_2 = \{2, 8, 11\}, B_3 = \{1, 3, 5\}, B_4 = \{6, 7, 10\}.$$

Then by taking $r_1 = 1$, $r_2 = 6$, $r_3 = 8$, $r_4 = 12$, and $\pi = (3, 4, 2, 1)$ we get that

$$r_1 \in A_1 \cap B_3, r_2 \in A_2 \cap B_4, r_3 \in A_3 \cap B_2, r_4 \in A_4 \cap B_1.$$

Hall (1935) showed a more general version of König's theorem (where subsets can have different sizes) as an implication of his famous result [22]—today known as the Hall's marriage theorem. The set $R = \{r_1, \dots, r_n\}$ in Theorem 1 is called a *complete system of representatives* for subsets A_i (and also for subsets B_i). The following theorem (which is a generalized version of Lemma 8 in [2]) is an implication of König's theorem.

² The cube of a graph G has the same vertices as G , and has an edge between two distinct vertices if and only if there exists a path, with at most three edges, between them in G .

► **Theorem 2.** *Let S be a set with kn elements that is partitioned, in two different ways, into n subsets each with k elements, namely A_1, \dots, A_n and B_1, \dots, B_n . Then, it is possible to label all elements of S with k distinct labels such that the k elements in each of $A_1, \dots, A_n, B_1, \dots, B_n$ have k distinct labels. Moreover, such a labeling can be found in polynomial time.*

Proof. By König's theorem there exists a subset $R = \{r_1, \dots, r_n\}$ of S that is a complete system of representatives for subsets A_i and for subsets B_i . Such a system R can be found as follows. Construct a bipartite graph $G = (V, E)$ with $2n$ vertices such that $V = \{A_1, \dots, A_n, B_1, \dots, B_n\}$ and there is an edge between A_i and B_j if and only if $A_i \cap B_j \neq \emptyset$. According to Hall's marriage theorem [9, 22] G has a perfect matching M (with n edges) which can be found in polynomial time. For every edge (A_i, B_j) in M pick an arbitrary representative element in $A_i \cap B_j$. These n representatives form R .

Label all elements of R by l_1 . Then remove the vertices of R from S and from corresponding subsets A_i and B_j . As a result we obtain a new set S with $(k-1)n$ elements and two distinct partitions of S each with n subsets of size $k-1$. By applying König's and Hall's theorems we can find another complete system of representatives, and label them l_2 . Repeating the above process achieves a desired labeling l_1, \dots, l_k . ◀

In the example above we can label the elements of S by $k (= 3)$ labels l_1, l_2, l_3 where (with a slight abuse of notation) $l_1 = \{1, 6, 8, 12\}$, $l_2 = \{2, 5, 9, 10\}$, and $l_3 = \{3, 4, 7, 11\}$ such that all elements in each A_i and B_i have different labels.


2.1 A 4-approximation for the 2-DBST

In this section we present a 4-approximation algorithm for the 2-DBST problem. Let P be a set of $2n$ points in a metric space that is partitioned into n tuples A_1, \dots, A_n each with two points. Let λ^* denote the bottleneck of a fixed optimal solution (consisting of two trees). We show how to find two disjoint trees R and B with edges of length at most $4\lambda^*$. To simplify our description we assume that the nodes of R and B are colored red and blue, respectively.

We start by computing a minimum spanning tree T of P , which is also a bottleneck spanning tree. Let e be a longest edge of T , that is $\lambda(T) = w(e)$. Let T_1 and T_2 be the two trees obtained by removing e from T . Notice that $\max\{\lambda(T_1), \lambda(T_2)\} \leq w(e)$. If each A_i has a point in T_1 and a point in T_2 , then we claim that $R = T_1$ and $B = T_2$ form an optimal solution because if the fixed optimal solution contains an edge between a node of T_1 and a node of T_2 then the length of that edge is at least $w(e)$ which implies that $\lambda^* \geq w(e)$. Therefore $\max\{\lambda(R), \lambda(B)\} \leq \lambda^*$.

Now assume that both points of some tuple A_i belong to say T_1 . In any feasible solution, one point of A_i is red and the other is blue. Then regardless of the coloring of the nodes of T_2 , the optimal solution should contain an edge between a node of T_1 and a node of T_2 . Thus $\lambda^* \geq w(e)$. We are going to color the nodes of T (which are the points of P) red and blue and then obtain R and B in such a way that $\max\{\lambda(R), \lambda(B)\} \leq 4 \cdot \lambda(T)$. This will imply that $\max\{\lambda(R), \lambda(B)\} \leq 4\lambda^*$.

We root T at a leaf q . Then we partition the nodes of T into n buckets B_1, \dots, B_n each with two vertices. The partitioning is done iteratively in a bottom-up fashion as follows. Consider a deepest leaf l and let v be its parent. Let u_1, u_2, \dots, u_j be the children of v where $u_1 = l$ as in Figure 2(a). If j is even then we create $j/2$ buckets $\{u_1, u_2\}, \{u_3, u_4\}, \dots, \{u_{j-1}, u_j\}$, and then remove u_1, \dots, u_j from T . If j is odd then we create $(j+1)/2$ buckets $\{v, u_1\}, \{u_2, u_3\}, \{u_4, u_5\}, \dots, \{u_{j-1}, u_j\}$, and then remove v, u_1, \dots, u_j from T . Then we repeat the above process until q and its only child form a bucket. We denote this last bucket by B_n .



■ **Figure 2** (a) Creating buckets. (b) Construction of R ; dashed edges represent δ_i .

The total number of buckets is n because T has $2n$ nodes initially. Between any two nodes in the same bucket there exists a path of length at most 2 in T , because the two nodes are either siblings or a child and its parent.

Now that we have two partitions A_1, \dots, A_n and B_1, \dots, B_n of P , we color (or label) the points of P by two colors, red and blue, as in Theorem 2. Thus in each A_i and each B_i we get a red point and a blue point. We construct the tree R by interconnecting the red points of buckets as follows; see Figure 2(b): Consider each bucket B_i with $i \in \{1, \dots, n - 1\}$ and let r_i denote its red point.

- (i) If the parent of r_i is not in B_i , then we connect r_i to the red point of its parent's bucket.
- (ii) If the parent of r_i is in B_i , then we connect r_i to the red point of its grandparent's bucket.

Since each r_i is connected to the red point in its parent's or grandparents's bucket, there is a unique path between r_i and r_n (the red point of the bucket B_n which contains the root q). Thus R is a tree. We construct the tree B on the blue points in a similar fashion. We claim that R and B are the desired trees. Since each A_i contains a red point and a blue point (by Theorem 2), each of R and B contains exactly one point from A_i . Thus R and B form a feasible solution for the problem.

Analysis of the approximation ratio. We show that $\lambda(R) \leq 4 \cdot \lambda(T)$; an analogous argument holds for B . Root R at the red point r_n of B_n . Consider any red node r_i in R where $i \in \{1, \dots, n - 1\}$. Recall that $r_i \in B_i$. Let p_i be the parent of r_i in R . It suffices to show that $|r_i p_i| \leq 4 \cdot \lambda(T)$. Consider the unique path δ_i between r_i and p_i in T . See Figure 2(b). If r_i was connected to p_i in case (i) then δ_i has at most 3 edges. If r_i was connected to p_i in case (ii) then δ_i has at most 4 edges. Therefore $|r_i p_i| \leq w(\delta_i) \leq 4 \cdot \lambda(T)$.

2.2 A $(3k - 2)$ -approximation for the k -DBST


Here we extend our 4-approximation algorithm of the previous section to get a $(3k - 2)$ -approximation for the k -DBST problem. We should note that (although it is not mentioned explicitly in their paper) Theorem 7 from Arkin et al. [2] combined with their Lemma 8 already gives a $(6k - 3)$ -approximation for the k -DBST problem.

Let P be a set of kn points that is partitioned into n tuples A_1, \dots, A_n each with k points. Let λ^* denote the bottleneck of a fixed optimal solution (consisting of k trees). We show how to color the points in each A_i by k colors c_1, \dots, c_k , and to obtain a tree T_i on all points with color c_i such that $\lambda(T_i) \leq (3k - 2)\lambda^*$.

Let T be a minimum spanning tree of P . Root T at a leaf q . We partition the nodes of T into n buckets B_1, \dots, B_n each with k nodes. The partitioning is done iteratively in a bottom-up fashion. We describe it for obtaining bucket B_j . For each node v in the current tree T , let $N(v)$ denote the number of nodes in the subtree rooted at v , including v itself.

Then we look at all nodes v for which $N(v)$ is at least k . Among those, pick a node v for which $N(v)$ is minimum. Then $N(v)$ is at least k and each of its children has a subtree of size at most $k - 1$. Now we construct B_j : Take a leaf in the subtree of v , add it to B_j , and remove it from the tree. Repeat this until B_j has size k .

Notice that between any two points in the subtree of v there is a path in T with at most $2k - 2$ edges. We say that v is the *representative* of B_j . Moreover, we define the *parent* of B_j to be the bucket containing v (if $v \notin B_j$) or the bucket containing v 's parent (if $v \in B_j$). With the two partitions A_1, \dots, A_n and B_1, \dots, B_n in hand, we color the points of P by k colors c_1, \dots, c_k as in Theorem 2. Thus in each A_i and in each B_i we get k distinct colors. For each color c_i we construct T_i as follows: for each bucket B_j we connect its point with color c_i (say point r_i) to the point with color c_i in B_j 's parent bucket (say point p_i). To prove the approximation ratio it suffices to show that between r_i and p_i there is a path of length at most $3k - 2$ in T . There is a path of length at most $k - 1$ from r_i to the representative of B_j , say v . If $v \in B_j$ then there is an edge from v to a node u in B_j 's parent bucket, and there is a path of length at most $2k - 2$ between u and p_i in the parent bucket. If $v \notin B_j$ then v is in B_j 's parent bucket in which case there is a path of length at most $2k - 2$ between v and p_i in the parent bucket. The following theorem summarizes our result.



► **Theorem 3.** *There exists a polynomial-time $(3k - 2)$ -approximation algorithm for the k -disjoint bottleneck spanning tree problem on points in a metric space.*

Remark. The length $2k - 2$ within each bucket of size k is the best achievable. For example consider a tree rooted at a node v with $k + 1$ subtrees each is a path with $k - 1$ nodes. This tree has k^2 nodes in total which will be partitioned into k buckets of size k . Since there are $k + 1$ leaves at least two of them lie in the same bucket (by the pigeonhole principle), and thus the distance between them will be $2k - 2$.

3 The 2-GBST problem


In this section we study the 2-GBST problem: Given a set P of n points in some metric space that is partitioned into clusters of size at most 2, find a tree that contains exactly one point from each cluster and minimizes the largest edge length. The 2-GBST problem is NP-hard; this is implied from the hardness of the so-called *rainbow minmax gap problem* [3] (given a set of pairs of points on a line, select exactly one point from each pair such that the maximum Euclidean distance between consecutive selected points is minimized).

In this section we first prove the inapproximability of the metric 2-GBST problem and then present an approximation algorithm.

► **Theorem 4.** *Unless $P = NP$, there is no polynomial-time algorithm that approximates the metric 2-generalized bottleneck spanning tree problem by a factor better than 2.*

Proof. We use a reduction from the 3-SAT problem: given a boolean expression E as the conjunction of clauses, each of which is the disjunction of three distinct literals (a variable or its negation), decide whether E is satisfiable.

Given any instance of the 3-SAT problem consisting of an expression E with m clauses C_1, \dots, C_m and n variables x_1, \dots, x_n we construct an instance of the 2-GBST problem consisting of a metric graph G as follows (the vertices of G represent points in a metric space). For each clause C_j create a cluster with one vertex p_j . For each variable x_i create a cluster with two *literal vertices* v_i and \bar{v}_i that correspond to positive literal x_i and negative literal \bar{x}_i , respectively. Create a cluster with one vertex r . To simplify our description we use vertices and their corresponding clauses or literals interchangeably. Connect each literal vertex, by edges of weight 1, to vertices p_j of all clauses C_j that they appear in. Connect r to all literal vertices by edges of weight 1. The edges that are drawn in the figure have weight 1. All other edges of G have weight 2 (notice that G is a complete graph). With this weight assignment, G is a metric graph with $m + 2n + 1$ vertices. We show that E is satisfiable if and only if G has a generalized spanning tree with edges of weight 1. This would imply the statement of the theorem because (by contraposition) any approximation algorithm with factor less than 2 would give a tree with edges of weight 1, and thus could solve the 3-SAT problem.



First suppose that E is satisfiable, and consider a truth assignment of variables that satisfies E . We obtain a tree T as follows. For the vertex set of T we select r , all vertices p_j , and each v_i (if x_i is true) or \bar{v}_i (if x_i is false). For the edge set of T we connect r to every selected literal vertex, and we connect each p_j to exactly one selected literal vertex that satisfies C_j . The tree T is a feasible solution for the 2-GBST problem on G (as it contains exactly one vertex from each cluster) and all its edges have weight 1.

For the other direction assume that T is a generalized spanning tree of G with edges of weight 1. The tree T should contain r and all vertices p_j because they are the only vertices in their clusters. For each p_j only edges of G that connect p_j to literal vertices have weight 1. Thus each p_j is connected to at least one literal vertex in T . Moreover T contains exactly one vertex from each cluster $\{v_i, \bar{v}_i\}$ of literal vertices. Therefore, by setting x_i as true (if T contains v_i) or false (if T contains \bar{v}_i) we obtain a satisfying assignment for E . ◀

If in the proof of Theorem 4 we replace all edge-weights of 2 with an arbitrary large constant, we obtain the following corollary.

► **Corollary 5.** *It is NP-hard to approximate the non-metric 2-generalized bottleneck spanning tree problem within any constant factor.*

If we were interested in generalized minimum spanning trees, then our reduction in the proof of Theorem 4 would also give a short proof for the NP-hardness of the metric 2-GMST problem: It can be verified that E is satisfiable if and only if G has a generalized spanning tree of total weight $m + n$. We note the existence of (somewhat involved) proofs for the hardness of the Euclidean 2-GMST problem; see the thesis of Fraser [19, page 140] (reduction from maximum 2-SAT), the paper of Ataei et al. [25] (reduction from planar 3-SAT), and a recent result of Dey et al. [16] (reduction from maximum 2-SAT).

3.1 A 3-approximation for the 2-GBST

Here we present our 3-approximation algorithm for the 2-GBST problem on a set P of n points in a metric space that is partitioned into m clusters C_1, \dots, C_m , each of size at most 2.

3:10 Approximating Bottleneck Spanning Trees on Partitioned Tuples of Points

Notice that $n/2 \leq m \leq n$. Let λ^* be the bottleneck of a fixed optimal solution. In a nutshell, our algorithm works as follows. First we compute a tree T_1 that contains at least one point from each cluster and its bottleneck is at most λ^* . Then we obtain a tree T_2 from T_1 that contains exactly one point from each cluster and its bottleneck is at most three times $\lambda(T_1)$. Therefore

$$\lambda(T_2) \leq 3 \cdot \lambda(T_1) \leq 3 \cdot \lambda^*,$$

which means that T_2 is a 3-approximate solution for the 2-GBST problem. In the rest of this section we show how to construct T_1 and T_2 . Our algorithm for computing T_2 from T_1 is of independent interest. The running time of our algorithm is dominated by the computation of a minimum spanning tree. The following theorem summarizes our result.

► **Theorem 6.** *There exists a polynomial-time 3-approximation algorithm for the 2-generalized bottleneck spanning tree problem on points in a metric space.*

3.1.1 Construction of T_1

First we make an empty graph G over the n points of P . Then we add edges between the points of G in a non-decreasing order of the pairwise distances, and stop as soon as G has a connected component, say C , that contains at least one point from each cluster. All edges of C are of length at most λ^* . Now we compute T_1 as an arbitrary spanning tree of C .

Remark. When the running time is a concern, one can guess λ^* in a binary search fashion to speed up the algorithm. Also, it is possible to compute T_1 as a subtree of the minimum spanning tree of P . In this case, the total running time is dominated by the computation of the minimum spanning tree; the details are skipped as we are not concerned about the running time here.


3.1.2 Construction of T_2

In this section we prove the following theorem.

► **Theorem 7.** *Given a tree T_1 and a partitioning of its nodes into clusters of size at most two, we can obtain a tree T_2 that contains exactly one node from each cluster and the length of its edges is at most 3 in the metric of T_1 . The upper bound 3 is the best achievable.*

First we show that the distance 3 (in the metric of T_1) is the best achievable upper bound. Figure 3 illustrates a tree T_1 as a path with eight nodes. The nodes of T_1 are partitioned into five clusters $\{a\}$, $\{b_1, b_2\}$, $\{c_1, c_2\}$, $\{d_1, d_2\}$, $\{e\}$. To obtain T_2 we have to choose points a and e because they are the only points in their clusters. Due to symmetry we may choose b_1 from cluster $\{b_1, b_2\}$. In this case if we do not choose d_2 then the distance of e to its closest point in T_2 would be at least 3, thus we may assume d_2 is chosen. In this setting, if we choose c_1 (as depicted in Figure 3) then the distance between c_1 and d_2 will be 3, and if we choose c_2 then the distance between b_1 and c_2 will be 3. Thus, in all cases we get an edge of length 3 in T_2 .

Now we present an algorithm for obtaining T_2 . Our algorithm consists of two phases: In the first phase we select the nodes of T_2 and in the second phase we define its edges. To select the nodes of T_2 , we *visit* the nodes of T_1 in some order and *select* exactly one node from each cluster. While visiting the nodes of T_1 we refer to an unvisited node by *open node*, to a visited node that is selected by *selected node*, and to a visited node that is not selected by *burned node*.




■ **Figure 3** Illustration of the lower bound 3. Dashed lines represent two nodes belonging to the same cluster. The black vertices are chosen for T_2 .

Node selection. See Figure 4(a) for an illustration of this phase. At the beginning all nodes of T_1 are open. First we visit and select all nodes of clusters of size one (which must be in T_2). Now we are going to select exactly one node from each cluster of size two. We root T_1 at an arbitrary node. Then we repeat the following process until all nodes of T_1 are visited. The process starts from an open node. At the beginning if the root is open then we start from the root, otherwise start from an arbitrary open node. In Figure 4(a) the nodes are labeled by the order they have been visited; the nodes of clusters of size one (which are already visited) are labeled with 0s.

Process: Let a_1 denote the starting open node (which belongs to a cluster of size two). Select a_1 and burn its twin say a_2 . If the parent of a_2 is open then repeat the process starting from the parent. If the parent is not open (selected or burned) then check the children of a_2 . If a_2 has some open child then repeat the process starting from an open child. If a_2 has no open child (or if a_2 does not have any child at all) then repeat the process starting from an arbitrary open node if such a node exists otherwise terminate the node selection phase.

Defining edges. The node selection algorithm selects exactly one node from each cluster. At the end of the selection algorithm, every node is either selected or burned (there is no open node). We claim that for each selected node a at any level of T_1 (except for the root) there exists a selected node b at a higher level such that the path between a and b in T_1 has at most three edges, i.e. the distance between a and b is at most 3 in the metric of T_1 . For each selected node a , we add the edge (a, b) to T_2 . As each a is connected to a node in a higher level, all nodes of T_2 are connected (via the root) and hence it is a tree.



■ **Figure 4** (a) Node selection (dashed lines represent two nodes in the same cluster): red squares (belong to clusters of size one) and black nodes (belong to clusters of size two) are selected whereas the white nodes (paired with black nodes) are burned. (b) Illustration for the edge length of T_2 .

To verify our claim let a_1 be the parent of a , as in Figure 4(b). If a_1 is selected then set $b = a_1$ and we are done. Assume that a_1 is burned. Let a_2 be the parent of a_1 . If a_2 is selected then set $b = a_2$ and we are done. Assume that a_2 is also burned. Notice that a_2 was burned before a_1 was, because otherwise the selection process would select a_2 right after burning a_1 . Right after burning a_2 the process has checked the parent of a_2 which we denote by a_3 . If a_3 was open then it would have been selected, and thus we set $b = a_3$ and we are done. If a_3 was burned then the process would have checked the children of a_2 and have selected a child a'_1 because a_2 had an open child which was a_1 ; this case is depicted in Figure 4(b). In this case we set $b = a'_1$ and we are done.

Remark. It might be tempting to use our 3-approximation algorithm for the 2-GBST problem to obtain a 3-approximation for the 2-DBST problem, say by coloring the selected nodes red and the burned nodes blue. This may not be an easy task because each time the process starts by selecting an *arbitrary* open node, these selected nodes could form a long path between burned nodes.

4 The k -PBST problem

Let $k \geq 2$ be an integer. In this section we study the k -PBST problem: Given kn points in some metric space, find k disjoint trees each containing exactly n points and minimize the largest edge length (over all trees). First we prove the hardness of this problem, and then we present an approximation algorithm. For $n = 2$ the problem is equivalent to the bottleneck matching problem which can be solved in polynomial time; see for example [12]. In the rest of this section we assume that $n \geq 3$.

► **Theorem 8.** *Unless $P = NP$, there is no polynomial-time algorithm that approximates the metric k -partition bottleneck spanning tree problem by a factor better than 2, for any $k \geq 2$.*

Proof. We use a reduction from the NP-hard problem of partitioning the vertex set of a graph $G = (V, E)$ into k ($2 \leq k \leq |V|/3$) equal-size subsets V_1, \dots, V_k such that the subgraph induced by each V_i is connected [18]. Let G' be the complete edge-weighted graph obtained by adding edges to G and then assigning weight 1 to every edge of E and weight 2 to every other edge. Notice that G' is a metric graph with $|V|$ vertices. It is easily seen that the partition problem on G has a solution if and only if G' contains k equal-size disjoint spanning trees with edges of weight 1. The inapproximability claim follows because any approximation algorithm with factor less than 2 would give spanning trees with edges of weight 1, which would solve the partitioning problem on G . ◀

If in the proof of Theorem 8 we replace all edge-weights of 2 with an arbitrary large constant, we obtain the following corollary.

► **Corollary 9.** *It is NP-hard to approximate the non-metric k -partition bottleneck spanning tree problem within any constant factor, for any $k \geq 2$.*

4.1 Approximating the k -PBST

Now we present an $\alpha(k)$ -approximation algorithm for the k -PBST problem, where $\alpha(k) = 2$ for $k = 2, 3$ and $\alpha(k) = 3$ for $k \geq 4$. In view of Theorem 8 the factor 2 is the best achievable for $k = 2, 3$. Given kn points in a metric space, we show how to construct k trees each containing exactly n points and their largest edge length is at most $\alpha(k) \cdot \lambda^*$, where λ^* is the bottleneck of a fixed optimal solution.

We start by computing a minimum spanning tree T of all points. Let e be a longest edge of T , that is $\lambda(T) = w(e)$. Let T' and T'' be the two trees obtained by removing e from T . If the number of nodes in T' and in T'' are multiples of n , say $i \cdot n$ and $j \cdot n$ where $i + j = k$, then we recursively construct i trees on the nodes of T' and j trees on the nodes of T'' .

Assume that the number of nodes in T' and T'' are not multiples of k . Then the optimal solution must have an edge between a node of T' and a node of T'' . The length of this edge is at least $w(e)$, and thus $\lambda^* \geq \lambda(T)$. Then by Theorem 11 we obtain k trees on the nodes of T such that their edge lengths are at most $\alpha(k) \cdot \lambda(T)$. We emphasize that the edge lengths are measured in the metric of T . The following theorem summarizes our result in this section.

► **Theorem 10.** *There exists a polynomial-time α -approximation algorithm for the k -partition bottleneck spanning tree problem on kn points in a metric space where $\alpha = 2$ for $k = 2, 3$ and $\alpha = 3$ for $k \geq 4$. The approximation factor 2 for $k = 2, 3$ is the best achievable in polynomial time.*

4.2 Balanced tree partitioning theorem

In this section we prove the following theorem. We denote the number of nodes of a tree T by $|T|$.

► **Theorem 11.** *Given a tree T with kn nodes we can obtain k disjoint trees T_1, \dots, T_k each containing exactly n nodes of T such that*


1. *If $k = 2$ or $k = 3$ then the length of edges in each T_i is at most 2 in the metric of T .*
2. *If $k \geq 4$ then the length of edges in each T_i is at most 3 in the metric of T .*

The upper bounds 2 and 3 for the edge lengths are the best achievable.

For the proof we first show that the upper bounds 2 and 3 are the best achievable. Then we present algorithms that achieve *desirable* trees T_1, \dots, T_k with the claimed edge lengths. The lengths mentioned in our proof are in the metric of T .

Upper bounds. It is easily seen that the upper bound of 2 is the best achievable (for $k = 2, 3$) for example when T is a star with 3 and 5 leaves, respectively.

To verify that 3 is the best achievable upper bound (for $k \geq 4$) consider a tree T rooted at a node v with $k + 1$ subtrees each is a path with $k - 1$ nodes; see the figure to the right for $k = 4$. The tree T has k^2 nodes. Let R be the set of $k + 1$ nodes that are at distance 1 from v . Each node of R represents a path connected to v . Now consider any set of k disjoint trees T_1, \dots, T_k each consisting of k nodes of T . We show, by contradiction, that the length of an edge in some T_i is at least 3. After a suitable relabeling assume that v belongs to T_1 . Then each tree T_i with $i \in \{2, \dots, k\}$ should have nodes from at least two of the paths connected to v because each path itself has $k - 1$ nodes. In particular T_i should contain the representatives of these paths because otherwise T_i would have an edge of length at least 3. Thus each T_i contains at least two distinct nodes from R . This implies that $|R| \geq 2(k - 1)$. Combining this inequality with the fact that $|R| = k + 1$, implies that $k \leq 3$ which is a contradiction.




Algorithm for $k \geq 4$. Our algorithm for $k \geq 4$ uses the fact that the cube of T is Hamiltonian. It is implied from a result of Karaganis [26] and independently from a result of Lesniak [28] that in polynomial time we can find a Hamiltonian path on nodes of T with

3:14 **Approximating Bottleneck Spanning Trees on Partitioned Tuples of Points**

edges of length at most 3. By cutting this path into k equal-size pieces we obtain k desired trees.

Remark. One could simply obtain a 2-approximation if the *square*³ of T has a Hamiltonian path. However, this property holds only for a very restricted class of trees called *horsetail* [39].



■ **Figure 5** Obtaining trees R (in red) and B (in blue) from T (in black). The red nodes marked with \times have been recolored blue to ensure that the tree has the same number of red and blue nodes.

Algorithm for $k = 2$. We show how to find two disjoint trees R and B each containing exactly n nodes of T and the length of their edges is at most 2. To simplify our description we assume that the nodes of R and B are colored red and blue, respectively.

We root T at a leaf r , as in Figure 5. Then r has only one child which we denote by b . Assume that r is at level 1, its child b is at level 2, the children of b are at level 3, and so on. Color all nodes at odd levels red and color all nodes at even levels blue. If the number of nodes of one color, say red, is more than n then we iteratively recolor the red nodes from the bottommost level by blue. In the end we get n red nodes and n blue nodes. We compute a rooted tree R on red nodes by connecting each red node to its grandparent, and compute a rooted tree B on blue nodes by connecting each blue node to its parent or grandparent that is blue, as in Figure 5. Notice that R is rooted at r and B is rooted at b . Since each red and blue node is connected to its parent or grandparent (in T), every edge of R and B has length at most 2. Therefore $\{R, B\}$ is a 2-approximate solution.

Remark. It is easily seen that the above algorithm can be extended to obtain trees R and B of different sizes (as long as $|R| + |B| = |T|$) with the same upper bound of 2 on their edge lengths.

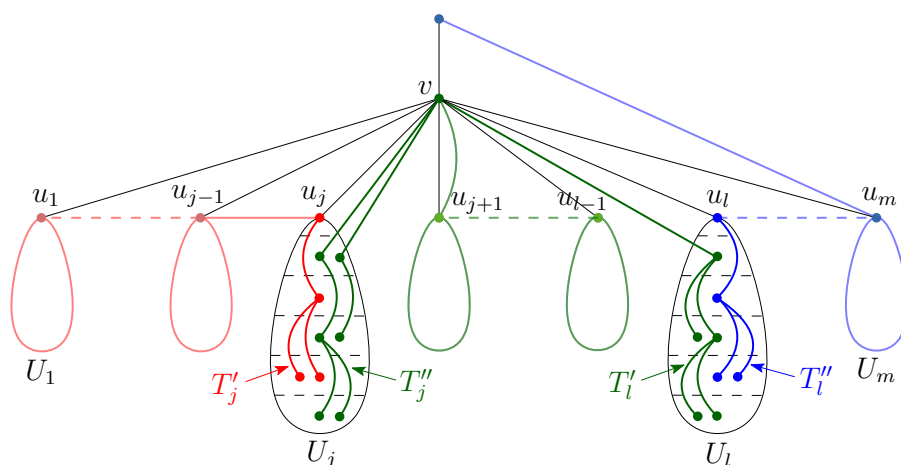
Algorithm for $k = 3$. Notice that T has $3n$ nodes. We show how to find three disjoint trees R , G , and B each containing exactly n nodes of T and the length of their edges is at most 2.

We root T at a leaf node. For each node v in T , let $N(v)$ denote the number of nodes in the subtree rooted at v ; the node v is counted. Then we look at all nodes v for which $N(v)$ is at least n . Among those, pick a node v for which $N(v)$ is minimum. Then $N(v)$ is at least n and each of its children has a subtree of size at most $n - 1$. Observe that v is not the root.

If $N(v) = n$ then we take the subtree rooted at v as R , remove R from T , and then obtain two trees G and B from the new tree T (which now has $2n$ nodes) using our algorithm for $k = 2$.

³ The square of a graph G has the same vertices as G , and has an edge between two distinct vertices if and only if there exists a path, with at most two edges, between them in G .

Assume that $N(v) > n$. Then v has at least two children which we denote by u_1, u_2, \dots, u_m where $m \geq 2$. Let U_i denote the subtree rooted at u_i . Take the smallest index j in $\{1, \dots, m\}$ for which $|U_1| + \dots + |U_j| \geq n$. Then $|U_1| + \dots + |U_{j-1}| < n$. Let $n'_1 = n - (|U_1| + \dots + |U_{j-1}|)$. Let U_j^v be the subtree consisting of U_j and the node v together with the edge connecting v to u_j . We use our algorithm for $k = 2$ to obtain from U_j^v two trees T'_j and T''_j of sizes n'_1 and $|U_j^v| - n'_1 = |U_j| + 1 - n'_1$, respectively, such that T'_j is rooted at u_j , T''_j is rooted at v , and their edge lengths are at most 2; see Figure 6. Now we obtain R by taking the trees U_1, \dots, U_{j-1} , and T'_j and interconnecting their roots to form one tree. Notice that R has n nodes and its edge lengths are at most 2. We remove the nodes of R from T . We also remove all edges of T that lie in U_j , and add the edges of T''_j (which are of length at most 2) to T . Notice that $|T''_j| < n$ because it does not have u_j (although it contains v). To obtain G and B we consider the following cases depending on the number $N(v)$ in the new tree T which has $2n$ nodes:



■ **Figure 6** Obtaining trees R (in red), G (in green), and B (in blue) from T (in black). The trees T'_j , T''_j , T'_l , and T''_l are shown with bold edges.

- $N(v) = n$. In this case we take the subtree rooted at v as G , remove G from T , and then take the resulting tree T (which now has n nodes) as B .
- $N(v) < n$. We walk up the tree T from v and stop at the first node w for which $N(w) \geq n$. Similar to the process of obtaining R from the subtree of v , we obtain G from the subtree of w . This time we denote the subtree of w that contains v by U_1 . This ensures that the edges of T''_j will appear in G without getting longer, because the subtree of a potential child w' of w with $N(w') \geq n$ must be disjoint from T''_j . After obtaining G , the remaining part of T will form the tree B .
- $N(v) > n$. See Figure 6. In this case we repeat a procedure similar to what we did to obtain R . Let $l \in \{j + 1, \dots, m\}$ be the smallest index for which $|T''_j| + |U_{j+1}| + \dots + |U_l| \geq n$. Notice that U_{j+1} exists because $N(v) > n$. Then $|T''_j| + |U_{j+1}| + \dots + |U_{l-1}| < n$. Let $n'_2 = n - (|T''_j| + |U_{j+1}| + \dots + |U_{l-1}|) + 1$ (the addition of 1 will become clear shortly). Let U_l^v be the subtree consisting of U_l and the node v together with the edge connecting v to u_l (notice that v also belongs to T''_j). We use our algorithm for $k = 2$ one more time to obtain from U_l^v two trees T'_l and T''_l of sizes n'_2 and $|U_l^v| - n'_2 = |U_l| + 1 - n'_2$, respectively, such that T'_l is rooted at v , T''_l is rooted at u_l and their edge lengths are at most 2. Now we obtain G by taking the trees T''_j , U_{j+1}, \dots, U_{l-1} , and T'_l and then

interconnecting their roots to form one tree. The tree G has n nodes (without double counting v which is in both T_j'' and T_l') and its edge lengths are at most 2. We obtain the third tree, i.e. B , as follows. We remove the nodes of G from T . By interconnecting the roots of $T_l'', U_{l+1}, \dots, U_m$ together and then connecting u_m to the parent of v (which exists) we obtain the tree B .

Remark. To see why the above procedure cannot be extended to the case of $k = 4$, assume that $N(v) > n$ after the removal of G from T . As v is already used for G we cannot reuse it to construct another tree, and hence we will be forced to introduce longer edges.

5 Conclusions

A natural open problem is to improve the presented approximation ratios further. Most of our approximation ratios consider the largest edge length of the standard BST as the lower bound. A better lower bound for the largest edge length of an optimal solution (not the standard BST) would improve the approximation ratios. It would be interesting to explore whether our algorithm for the 2-GBST problem could be extended to an $O(k)$ -approximation algorithm for the k -GBST problem. Also it would be interesting to verify whether the approximation ratio of 3 for the k -PBST problem ($k \geq 4$) is tight, knowing that 2 is a lower bound.

References

- 1 Mattias Andersson, Joachim Gudmundsson, Christos Levcopoulos, and Giri Narasimhan. Balanced partition of minimum spanning trees. *International Journal of Computational Geometry & Applications*, 13(4):303–316, 2003. <https://doi.org/10.1142/S0218195903001190>.
- 2 Esther M. Arkin, Aritra Banik, Paz Carmi, Gui Citovsky, Su Jia, Matthew J. Katz, Tyler Mayer, and Joseph S. B. Mitchell. Network optimization on partitioned pairs of points. In *Proceedings of the 28th International Symposium on Algorithms and Computation (ISAAC)*, pages 6:1–6:12, 2017. <http://arxiv.org/abs/1710.00876>.
- 3 Esther M. Arkin, Aritra Banik, Paz Carmi, Gui Citovsky, Matthew J. Katz, Joseph S. B. Mitchell, and Marina Simakov. Selecting and covering colored points. *Discret. Appl. Math.*, 250:75–86, 2018. <https://doi.org/10.1016/j.dam.2018.05.011>.
- 4 Esther M. Arkin and Refael Hassin. Minimum-diameter covering problems. *Networks*, 36(3):147–155, 2000. [https://doi.org/10.1002/1097-0037\(200010\)36:3<147::AID-NET1>3.0.CO;2-M](https://doi.org/10.1002/1097-0037(200010)36:3<147::AID-NET1>3.0.CO;2-M).
- 5 Igor Averbakh and Oded Berman. A heuristic with worst-case analysis for minimax routing of two travelling salesmen on a tree. *Discrete Applied Mathematics*, 68(1-2):17–32, 1996. [https://doi.org/10.1016/0166-218X\(95\)00054-U](https://doi.org/10.1016/0166-218X(95)00054-U).
- 6 Igor Averbakh and Oded Berman. $(p - 1)/(p + 1)$ -approximate algorithms for p -traveling salesmen problems on a tree with minmax objective. *Discrete Applied Mathematics*, 75(3):201–216, 1997. [https://doi.org/10.1016/S0166-218X\(97\)89161-5](https://doi.org/10.1016/S0166-218X(97)89161-5).
- 7 Binay K. Bhattacharya, Ante Custic, Akbar Rafiey, Arash Rafiey, and Vladyslav Sokol. Approximation algorithms for generalized MST and TSP in grid clusters. In *Proceedings of the 9th International Conference on Combinatorial Optimization and Applications (COCOA)*, pages 110–125, 2015. https://doi.org/10.1007/978-3-319-26626-8_9.
- 8 Víctor Blanco, Elena Fernández, and Justo Puerto. Minimum spanning trees with neighborhoods: Mathematical programming formulations and solution methods. *European Journal of Operational Research*, 262(3):863–878, 2017. <https://doi.org/10.1016/j.ejor.2017.04.023>.

- 9 J. Adrian Bondy and Uppaluri S. R. Murty. *Graph Theory with Applications*. Macmillan Education UK, 1976.
- 10 Paolo M. Camerini. The min-max spanning tree problem and some extensions. *Information Processing Letters*, 7(1):10–14, 1978. [https://doi.org/10.1016/0020-0190\(78\)90030-3](https://doi.org/10.1016/0020-0190(78)90030-3).
- 11 Erin W. Chambers, Alejandro Erickson, Sándor P. Fekete, Jonathan Lenchner, Jeff Sember, S. Venkatesh, Ulrike Stege, Svetlana Stolpner, Christophe Weibel, and Sue Whitesides. Connectivity graphs of uncertainty regions. *Algorithmica*, 78(3):990–1019, 2017. <https://doi.org/10.1007/s00453-016-0191-2>.
- 12 Maw-Shang Chang, Chuan Yi Tang, and Richard C. T. Lee. Solving the Euclidean bottleneck matching problem by k -relative neighborhood graphs. *Algorithmica*, 8(3):177–194, 1992. <https://doi.org/10.1007/BF01758842>.
- 13 Yong Chen, Zhi-Zhong Chen, Guohui Lin, Yao Xu, and An Zhang. Approximation algorithms for maximally balanced connected graph partition. *Algorithmica*, 83(12):3715–3740, 2021. <https://doi.org/10.1007/s00453-021-00870-3>.
- 14 Janka Chlebíková. Approximating the maximally balanced connected partition problem in graphs. *Information Processing Letters*, 60(5):223–230, 1996. [https://doi.org/10.1016/S0020-0190\(96\)00175-5](https://doi.org/10.1016/S0020-0190(96)00175-5).
- 15 Thomas H. Cormen, Charles E. Leiserson, and Ronald L. Rivest. *Introduction to Algorithms*. McGraw-Hill, Cambridge, 1989. <http://mitpress.mit.edu/books/introduction-algorithms>.
- 16 Sanjana Dey, Ramesh K. Jallu, and Subhas C. Nandy. On the minimum spanning tree problem in imprecise set-up. 2021. <https://arxiv.org/abs/2104.04229>.
- 17 Reza Dorrigiv, Robert Fraser, Meng He, Shahin Kamali, Akitoshi Kawamura, Alejandro López-Ortiz, and Diego Seco. On minimum- and maximum-weight minimum spanning trees with neighborhoods. *Theory of Computing Systems*, 56(1):220–250, 2015. <https://doi.org/10.1007/s00224-014-9591-3>.
- 18 Martin E. Dyer and Alan M. Frieze. On the complexity of partitioning graphs into connected subgraphs. *Discrete Applied Mathematics*, 10(2):139–153, 1985. [https://doi.org/10.1016/0166-218X\(85\)90008-3](https://doi.org/10.1016/0166-218X(85)90008-3).
- 19 Robert Fraser. *Algorithms for geometric covering and piercing problems*. Phd thesis, University of Waterloo, 2012. <https://uwspace.uwaterloo.ca/handle/10012/7190>.
- 20 Harold N. Gabow, Shachindra N. Maheshwari, and Leon J. Osterweil. On two problems in the generation of program test paths. *IEEE Transactions on Software Engineering*, 2(3):227–231, 1976. <https://doi.org/10.1109/TSE.1976.233819>.
- 21 Michel X. Goemans and Dimitris Bertsimas. Survivable networks, linear programming relaxations and the parsimonious property. *Mathematical Programming*, 60:145–166, 1993. <https://doi.org/10.1007/BF01580607>.
- 22 Philip Hall. On representatives of subsets. *Journal of the London Mathematical Society*, 10(1):26–30, 1935. <https://doi.org/10.1112/jlms/s1-10.37.26>.
- 23 Eran Halperin and Robert Krauthgamer. Polylogarithmic inapproximability. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 585–594, 2003. <https://doi.org/10.1145/780542.780628>.
- 24 Matthew P. Johnson. Red-blue-partitioned MST, TSP, and matching. In *Proceedings of the 30th Canadian Conference on Computational Geometry (CCCG)*, pages 259–264, 2018. <http://www.cs.umanitoba.ca/~7Ecccg2018/papers/session5B-p1.pdf>.
- 25 Homa Ataei Kachooei, Mansoor Davoodi, and Dena Tayebi. On the generalized minimum spanning tree in the Euclidean plane. In *Proceedings of the 1st Iranian Conference on Computational Geometry (ICCG)*, pages 19–23, 2018. <https://iasbs.ac.ir/~mdmonfared/uploads/papers/on%20the%20gmst%20in%20the%20euclidean%20plane.pdf>.
- 26 Jerome Karaganis. On the cube of a graph. *Canadian Mathematical Bulletin*, 11(2):295–296, 1968. <https://doi.org/10.4153/CMB-1968-037-0>.

- 27 Seigo Karakawa, Ehab Morsy, and Hiroshi Nagamochi. Minmax tree cover in the Euclidean space. *Journal of Graph Algorithms and Applications*, 15(3):345–371, 2011. https://doi.org/10.1007/978-3-642-00202-1_18.
- 28 Linda Lesniak. Graphs with 1-Hamiltonian-connected cubes. *Journal of Combinatorial Theory, Series B*, 14(2):148–152, 1973. [https://doi.org/10.1016/0095-8956\(73\)90058-0](https://doi.org/10.1016/0095-8956(73)90058-0).
- 29 Maarten Löffler. *Data Imprecision in Computational Geometry*. Phd thesis, Utrecht University, 2009. <http://dspace.library.uu.nl/handle/1874/36022>.
- 30 Joseph S. B. Mitchell. A PTAS for TSP with neighborhoods among fat regions in the plane. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 11–18, 2007. <https://dl.acm.org/doi/10.5555/1283383.1283385>.
- 31 Joseph S. B. Mitchell. A constant-factor approximation algorithm for TSP with pairwise-disjoint connected neighborhoods in the plane. In *Proceedings of the 26th ACM Symposium on Computational Geometry (SoCG)*, pages 183–191, 2010. <https://doi.org/10.1145/1810959.1810992>.
- 32 Young-Soo Myung, Chang-ho Lee, and Dong-wan Tcha. On the generalized minimum spanning tree problem. *Networks*, 26(4):231–241, 1995. <https://doi.org/10.1002/net.3230260407>.
- 33 Hiroshi Nagamochi and Taizo Kawada. Minmax subtree cover problem on cacti. *Discrete Applied Mathematics*, 154(8):1254–1263, 2006. <https://doi.org/10.1016/j.dam.2005.10.013>.
- 34 Hiroshi Nagamochi and Kohei Okada. Polynomial time 2-approximation algorithms for the minmax subtree cover problem. In *Proceedings of the 14th International Symposium on Algorithms and Computation (ISAAC)*, pages 138–147, 2003. https://doi.org/10.1007/978-3-540-24587-2_16.
- 35 Hiroshi Nagamochi and Kohei Okada. A faster 2-approximation algorithm for the minmax p -traveling salesmen problem on a tree. *Discrete Applied Mathematics*, 140(1-3):103–114, 2004. <https://doi.org/10.1016/j.dam.2003.06.001>.
- 36 Petrica C. Pop. The generalized minimum spanning tree problem: An overview of formulations, solution procedures and latest advances. *European Journal of Operational Research*, 283(1):1–15, 2020. <https://doi.org/10.1016/j.ejor.2019.05.017>.
- 37 Petrica C. Pop, Walter Kern, Georg Still, and Ulrich Faigle. Relaxation methods for the generalized minimum spanning tree problem. *Electronic Notes in Discrete Mathematics*, 8:76–79, 2001. [https://doi.org/10.1016/S1571-0653\(05\)80085-X](https://doi.org/10.1016/S1571-0653(05)80085-X).
- 38 Petrica C. Pop, Georg Still, and Walter Kern. An approximation algorithm for the generalized minimum spanning tree problem with bounded cluster size. In *Proceedings of the 1st Workshop on Algorithms and Complexity in Durham*, pages 115–121, 2005. <https://research.utwente.nl/en/publications/an-approximation-algorithm-for-the-generalized-minimum-spanning-t>.
- 39 Jakub Radoszewski and Wojciech Rytter. Hamiltonian paths in the square of a tree. In *Proceedings of the 22nd International Symposium on Algorithms and Computation (ISAAC)*, pages 90–99, 2011. https://doi.org/10.1007/978-3-642-25591-5_11.
- 40 Yoshiko Wakabayashi, Frédéric Chataigner, and Liliane Benning Salgado. Approximation and inapproximability results on balanced connected partitions of graphs. *Discrete Mathematics & Theoretical Computer Science*, 9(1), 2007. <http://dmtcs.episciences.org/384>.