

Improved approximation ratios for two Euclidean maximum spanning tree problems

Ahmad Biniaz

School of Computer Science, University of Windsor
ahmad.biniaz@gmail.com

Abstract

We study the following two maximization problems related to spanning trees in the Euclidean plane. It is not known whether or not these problems are NP-hard. We present approximation algorithms with improved approximation ratios for both problems. To obtain the new approximation ratios we employ the Steiner ratio, which has not been used in this context earlier.

(i) *Longest spanning tree with neighborhoods*: Given a collection of regions (neighborhoods) in the plane, the goal is to select a point in each neighborhood so that the longest spanning tree on selected points has maximum length. We present an algorithm with approximation ratio 0.524 for this problem. This improves the previous best ratio of 0.511 due to Chen and Dumitrescu (Discret. Math. Algorithms Appl. 2020). Our algorithm is fairly simple, its analysis is relatively short, and it takes linear time after computing a diametral pair of points.

(ii) *Longest noncrossing spanning tree*: Given a set of points in the plane, the goal is to find a maximum-length noncrossing spanning tree. Alon, Rajagopalan, and Suri (SoCG 1993) studied this problem for the first time and gave a 0.5-approximation algorithm. Over the years, the approximation ratio has been successively improved to 0.502, 0.503, and to 0.512. We present a 0.519-approximation algorithm for this problem. After the appearance of our algorithm on arXiv the ratio has been further improved to 0.546 by Cabello, Hoffmann, Klost, Mulzer, and Tkadlec (arXiv 2021). Our algorithm still has its own merits in terms of simplicity and brevity of the analysis as that of Cabello et al. involves thorough analysis that numerically approximates the roots of a degree-7 equation.

1 Introduction

The spanning tree is a well-studied and fundamental structure in graph theory and combinatorics. The well-known minimum spanning tree (Min-ST) problem asks for a spanning tree of a weighted graph, with minimum total edge weight. In contrast, the maximum spanning tree (Max-ST) problem asks for a spanning tree with maximum total edge weight. In the context of abstract graphs, the two problems are algorithmically equivalent in the sense that an algorithm for finding a Min-ST can also find a Max-ST within the same time bound (by simply negating the edge weights), and vice versa. The situation is quite different in the context of geometric graphs where vertices are points in the plane and edge weights are Euclidean distances between pairs of points. An algorithm that uses the geometry of the Euclidean plane for finding a Min-ST is not necessarily useful for computing a Max-ST because there is no known geometric transformation for setting up a duality between the “nearest” and the “farthest” relations between points [27]. In fact the existing geometric algorithms for the Min-ST and Max-ST problems exploit different sets of techniques.

This research is supported by NSERC.




Figure 1: (a) minimum spanning tree, (b) maximum spanning tree, (c) longest noncrossing spanning tree, and (d) longest spanning tree with four neighborhoods colored red, green, blue, and purple.

The problems of computing spanning trees with enforced properties (such as having minimum weight, maximum weight, bounded degree, or being noncrossing) have been well-studied in the last decades for both abstract graphs and geometric graphs. We study two problems related to maximum spanning trees in geometric graphs. The maximum spanning tree and related problems, in addition to their fundamental nature, find applications in worst-case analysis of heuristics for various problems in combinatorial optimization [2], and in approximating maximum triangulations [5, pp. 338]. They also have applications in cluster analysis, where one needs to partition a set of entities into well-separated and homogeneous clusters [4, 27]. Maximum spanning trees are directly related to computing diameter and farthest neighbors which are fundamental problems in computational geometry, with many applications [1].

In the classical Euclidean Max-ST problem we are given a set of n points in the plane (as vertices) and we want to find a spanning tree of maximum total edge length, where the length of every edge is the Euclidean distance between its two endpoints; see Figure 1. This problem can be solved in $O(n^2)$ time by Prim’s algorithm [19] for abstract graphs, and in $O(n \log n)$ time by an algorithm that uses the geometry [27]. In contrast to the Euclidean Min-ST which is always noncrossing (because of the triangle inequality), the Euclidean Max-ST is almost always self-crossing. One problem that we study in this paper is the *longest noncrossing spanning tree* (Max-NC-ST) problem which is to compute a noncrossing spanning tree of maximum length, as depicted in Figure 1(c). It is not known whether or not this problem is NP-hard.

Another problem that we study is the *longest spanning tree with neighborhoods* (Max-ST-NB): Given a collection of n neighborhoods (regions) in the plane, we want to find a maximum-length tree that connects n representative points, one point from each neighborhood, as in Figure 1(d). Each *neighborhood* is the union of simple polygons, and the neighborhoods are not necessarily disjoint. The neighborhoods are assumed to be colored by n different colors. The hardness of the Max-ST-NB problem is open. The difficulty lies in choosing the representative points; once these points are selected, the problem is reduced to the Euclidean Max-ST problem.

1.1 Related work on the longest spanning tree with neighborhoods

The Max-ST-NB problem has the same flavor as the Euclidean group Steiner tree problem in which we are given n groups of points in the plane and the goal is to find a shortest tree that contains “at least” one point from each group. The group Steiner tree problem in graphs is NP-hard and cannot be approximated by a factor $O(\log^{2-\epsilon} n)$ for any $\epsilon > 0$ [21]. The Max-ST-NB problem also lies in the concept of imprecision in computational geometry where each input point provided as a region of uncertainty and the exact position of the point may be anywhere in the region; see e.g. [15, 24].

One can show that the longest star, in which one vertex of a polygon is connected to one vertex in all other polygons, achieves a 0.5-approximate solution for the Max-ST-NB problem. Recently, Chen and Dumitrescu [13] present an approximation algorithm with improved ratio 0.511. Although their algorithm is simple, the analysis of its ratio is thoroughly involved. They also show that the approximation ratio of an algorithm that always includes a bichromatic diametral pair as an edge in the solution cannot be better than $\sqrt{2 - \sqrt{3}} \approx 0.517$.

Analogous problem has been studied for other structures with neighborhoods, e.g., minimum spanning tree [10, 15, 31], traveling salesman tour [3, 25, 26], convex hull [24, 30], to name a few.

1.2 Related work on the longest noncrossing spanning tree

Inspired by the seminal work of Alon, Rajagopalan, and Suri in SoCG 1993 [2], the study of long noncrossing configurations in the plane has received considerable attention in recent years. Alon et al. show how to compute constant-factor approximations of longest noncrossing spanning trees, perfect matchings, and Hamiltonian paths. They show that the longest *star*, a tree in which one vertex is connected to all others, always gives a 0.5 approximation of the longest spanning tree (a short proof of this claim is given in [17]). As pointed out by Alon et al. the ratio 0.5 between the lengths of a longest star and a longest (possibly crossing) spanning tree is the best possible (in the limit); this can be verified by placing $n/2$ points at arbitrary small neighborhood around $(0, 0)$ and $n/2$ points at arbitrary small neighborhood around $(1, 0)$. Therefore, to obtain a better approximation ratio one should take into account spanning trees other than stars. The ratio 0.5 remained the best known for almost seventeen years until Dumitrescu and Tóth (STACS 2010) [17] slightly improved it to 0.502, which was then improved to 0.503 by Biniaz et al. [8]. The ratios 0.5, 0.502, and 0.503 are obtained by considering the length of a longest (possibly crossing) spanning tree as the upper bound. Although such a tree provides a safe upper bound, it is not a valid solution for the Max-NC-ST problem. Alon et al. show that if the longest crossing spanning tree were to be used as the upper bound then the approximation ratio cannot be improved beyond $2/\pi < 0.637$ because such a tree can be $\pi/2$ times longer than a longest noncrossing spanning tree. Cabello et al. [11] employed a longest noncrossing spanning tree as the upper bound and obtained the ratio of 0.512.

The survey article by Eppstein [18, pp. 439] lists the hardness of Max-NC-ST as an open problem in the context of geometric network optimization. This problem has also been studied for other structures. Alon et al. show approximation ratios $2/\pi$ and $1/\pi$ for the longest noncrossing perfect matching and Hamiltonian path, respectively. The ratio for the Hamiltonian path was improved to $2/(1+\pi)$ by Dumitrescu and Tóth who also gave the first (though non constant-factor) approximation algorithm for the longest noncrossing Hamiltonian cycle. The longest noncrossing spanning tree is also studied in multipartite geometric graphs [8].

1.3 Our contributions and approach

We report improved approximation ratios for the Max-NC-ST and the Max-ST-NB problems. Our results are summarized in the following theorems.

Theorem 1. *A 0.524-approximation for the longest spanning tree with neighborhoods can be computed in linear time after computing a bichromatic diameter.*

Theorem 2. *A 0.519-approximation for the longest noncrossing spanning tree can be computed in polynomial time.*

To obtain the new approximation ratios we use the Euclidean Steiner ratio (of three points), which has not been used in this context earlier. We employ the Steiner ratio to handle a situation where some points lie far from a diametral pair from the input; this situation is a bottleneck of previous algorithms for both problems. The employment, which is sketched below, is not straightforward. We first obtain a lower bound for the length of the minimum spanning tree of a small subset of the input. We use this lower bound, along with the Steiner ratio, and obtain a lower bound on the length of the Steiner minimal tree of the subset. Then we use this new lower bound to construct a long spanning tree on the entire input. To see why this employment is nontrivial, one may think of this counterintuitive question: how could a lower bound on the length of the “minimum” spanning tree lead to a constant-ratio approximation of the “maximum” spanning tree? The employment of the Steiner ratio, not only improves the approximation ratios but also simplifies our analysis.

For the Max-ST-NB problem we give an approximation algorithm with improved ratio 0.524. The algorithm is not complicated: we find a bichromatic diameter (farthest pair of input vertices with different colors) and use it to compute three stars and one *double-star* (a tree of diameter 3), and then report the longest one. After computing a bichromatic diameter (which is a well-studied problem) the rest of the algorithm takes linear time. Our analysis of the ratio 0.524 is relatively short, compared to that of Chen and Dumitrescu [13] for the ratio 0.511. As a secondary result we show that the approximation ratio of an algorithm that always includes a bichromatic diameter in the solution cannot be better than 0.5, thereby improving the previous upper bound 0.517 by [13].

For the well-studied Max-NC-ST problem we give a polynomial-time approximation algorithm with ratio 0.519. Very recently, after the appearance of our work on arXiv [7], Cabello, Hoffmann, Klost, Mulzer, and Tkadlec [12] further improved the ratio significantly to 0.546 using thorough analysis. Despite this improvement our algorithm still has its own merits in terms of simplicity and brevity. Our analysis is fairly short compared to the detailed analysis of Cabello et al. [12] which involves numerically approximating the roots of a degree-7 equation. The ratios obtained by our algorithm and that of [11, 12] are with respect to the longest “noncrossing” spanning tree while the ratios of [2, 8, 17] are with respect to the longest spanning tree.

2 Preliminaries for the algorithms

Both our algorithms make extensive use of trees with low diameter, such as stars and double-stars. In fact our solutions for the Max-ST-NB and the Max-NC-ST problems have diameters at most three and six, respectively. A *star*, centered at a vertex p , is a tree in which every edge is incident to p . A *double-star*, centered at two vertices p and q , is a tree that contains the edge pq and its every other edge is incident to either p or q .

We denote the Euclidean distance between two points p and q in the plane by $|pq|$. A *geometric graph* is a graph whose vertices are points in the plane and whose edges are straight line segments. The length of a geometric graph G , denoted by $\text{len}(G)$, is the total length of its edges. We denote by $D(p, r)$ the disk of radius r that is centered at point p .

For a point set P , a *diametral pair* is a pair of points that attain the maximum distance. If the points in P are colored, then a *bichromatic diametral pair* is defined as a pair of points with different colors that attain the maximum distance.

The *Euclidean Steiner tree* problem asks for the shortest connected geometric graph spanning a given set P of points in the plane. The solution takes the form of a tree, called a Steiner minimal tree (SMT), that includes all points of P along with zero or more extra vertices called *Steiner points* [5]. The *Steiner ratio* is defined to be the infimum of the length of the Steiner minimal tree divided

by the length of the minimum spanning tree, over all point sets in the plane:

$$\rho = \inf_{P \subset \mathbb{R}^2} \left\{ \frac{\text{len}(\text{SMT}(P))}{\text{len}(\text{Min-ST}(P))} \right\}.$$

An old conjecture of Gilbert and Pollak [20] states that $\rho = \sqrt{3}/2 \approx 0.866$; this is achieved when P is the vertices of an equilateral triangle. Although the conjecture seems to be still open [22, 23], it has been verified when P has 3 [20], 4 [28], 5 [16], or 6 [29] points. For the purpose of our analysis the original proof of Gilbert and Pollak [20] for $|P| = 3$ is enough.

A simple 0.5-approximation algorithm. Chen and Dumitrescu [13] gave the following simple 0.5-approximation algorithm for the Max-ST-NB (a similar approach was previously used in [17]). Take a bichromatic diametral pair (a, b) from the given n neighborhoods; a and b belong to two different neighborhoods. Each edge of every optimal solution T^* has length at most $|ab|$, and thus $\text{len}(T^*) \leq (n-1)|ab|$. Choose an arbitrary point from each of the other $n-2$ neighborhoods. Let S_a be the star obtained by connecting a to b and to all chosen points. Define S_b analogously. By the triangle inequality $\text{len}(S_a) + \text{len}(S_b) \geq n|ab| \geq \text{len}(T^*)$. Therefore the longer of S_a and S_b is a 0.5-approximate solution for the Max-ST-NB problem. This idea also achieves a 0.5-approximate solution for the Max-NC-ST problem (for which every input point can be viewed as a neighborhood).

3 Maximum spanning tree with neighborhoods

In this section we prove Theorem 1. Put $\delta = 0.524$.

We describe our algorithm for the Max-ST-NB first, as it is easier to understand. It also gives insights on a better understanding of our algorithm for the Max-NC-ST problem. To facilitate comparisons we use the same notation as of Chen and Dumitrescu [13]. Let $\mathcal{X} = \{X_1, X_2, \dots, X_n\}$ be the given collection of n polygonal neighborhoods of N total vertices. We may assume that each X_i is colored by a unique color. We present a δ -approximation algorithm for computing a longest spanning tree with neighborhoods in \mathcal{X} . Our algorithm selects representative points only from boundary vertices of polygonal neighborhoods. Thus, in the algorithm (but not in the analysis) we consider each polygonal neighborhood X_i as the set of its boundary vertices, and consequently we consider \mathcal{X} as a collection of N points colored by n colors. Define the *longest spanning star* centered at a vertex $p \in X_i$ as the star connecting p to its farthest vertex in every other neighborhood.

Algorithm. The main idea of the algorithm is simple: we compute a spanning double-star D and three spanning stars S_1, S_2, S_3 , and then report the longest one.

Let (a, b) be a bichromatic diametral pair of \mathcal{X} . After a suitable relabeling assume that $a \in X_1$ and $b \in X_2$. We compute D as follows. Add the edge ab to D . For each X_i , with $i \in \{3, \dots, n\}$, find a vertex $p_i \in X_i$ that is farthest from a and find a vertex $q_i \in X_i$ that is farthest from b (it might be the case that $p_i = q_i$). If $|ap_i| \geq |bq_i|$ then add ap_i to D otherwise add bq_i to D . Observe that D spans all neighborhoods in \mathcal{X} , and each edge of D has length at least $|ab|/2$. Now we introduce the three stars. Let a' be a vertex of X_1 that is farthest from a , and let b' be a vertex of X_2 that is farthest from b . Notice that a' has the same color as a , and b' has the same color as b . We compute S_1 as the longest spanning star that is centered at a' , and we compute S_2 as the longest spanning star that is centered at b' . Now let c be a vertex in \mathcal{X} that maximizes $|ac| + |bc|$. The vertex c can be in any of neighborhoods X_1, \dots, X_n , also it might be the case that $c = a'$ or $c = b'$. We compute S_3 as the longest spanning star that is centered at c .

Running time. It is implied by a result of Biniáz et al. [9] that a bichromatic diametral pair of \mathcal{X} can be found in $O(N \log N \log n)$ time (the algorithm of Bhattacharya and Toussaint [6] also computes a bichromatic diameter, but only for two-colored points). After finding (a, b) , the rest of the algorithm (finding a', b', c , and finding farthest points from a, b, a', b', c) takes $O(N)$ time.

3.1 Analysis of the approximation ratio

For the analysis we consider \mathcal{X} as the initial collection of polygonal neighborhoods. Let T^* denote a longest spanning tree with neighborhoods in \mathcal{X} . It is not hard to see that for any point in the plane, its farthest point in a polygon P must be a vertex of P (see e.g. [14, Chapter 7]). Thus, any bichromatic diameter of \mathcal{X} is introduced by two vertices of polygons in \mathcal{X} . Hence the pair (a, b) , selected in the algorithm, is a bichromatic diametral pair of the initial collection \mathcal{X} . Therefore, $|ab|$ is an upper bound for the length of edges in T^* . Recall our assumption from the algorithm that $a \in X_1$ and $b \in X_2$. After a suitable rotation, translation, and scaling assume that ab is horizontal, $a = (0, 0)$, and $b = (1, 0)$. Since $|ab| = 1$ and T^* has $n - 1$ edges,

$$\text{len}(T^*) \leq (n - 1)|ab| \leq n - 1. \quad (1)$$

Lemma 1. *The double star D is longer than any star that is centered at a or at b .*

Proof. We prove this lemma for any star S_a centered at a ; the proof for stars centered at b is analogous. Recall that D contains the edge ab where $b \in X_2$, and for each $i \in \{3, \dots, n\}$ it contains the longer of ap_i and aq_i where p_i and q_i are the farthest vertices of X_i from a and b , respectively.

For any $i \in \{2, \dots, n\}$ let r_i be the vertex of X_i that is connected to a in S_a . If $i = 2$ then $|ar_i| \leq |ab|$. If $i > 2$ then $|ar_i| \leq \max\{|ap_i|, |bq_i|\}$. Therefore $\text{len}(D) \geq \text{len}(S_a)$. \square

Recall a' and b' as vertices of X_1 and X_2 that are farthest from a and b , respectively.

Lemma 2. *If $|aa'| \geq 2\delta$ or $|bb'| \geq 2\delta$ then $\max\{\text{len}(S_1), \text{len}(S_2), \text{len}(D)\} \geq \delta \cdot \text{len}(T^*)$.*

Proof. First assume that $|aa'| \geq 2\delta$. Choose a vertex p_i in each neighborhood X_i , with $i \neq 1$, and connect it to a and to a' . We obtain two stars S_a and $S_{a'}$ that are centered at a and a' . The total length of these stars is at least

$$\sum_{i \in \{2, \dots, n\}} (|ap_i| + |p_i a'|) \geq \sum_{i \in \{2, \dots, n\}} |aa'| \geq 2\delta(n - 1),$$

and thus the longer star has length at least $\delta(n - 1)$. If $S_{a'}$ is longer then $\text{len}(S_1) \geq \text{len}(S_{a'}) \geq \delta(n - 1) \geq \delta \cdot \text{len}(T^*)$, where the first inequality is implied by the fact that S_1 is the longest spanning star centered at a' and the last inequality is implied by (1). If S_a is longer then $\text{len}(D) \geq \text{len}(S_a) \geq \delta(n - 1) \geq \delta \cdot \text{len}(T^*)$, where the first inequality is implied by Lemma 1.

If $|bb'| \geq 2\delta$, an analogous implication shows that the length of S_2 or D is at least $\delta \cdot \text{len}(T^*)$. \square

Having Lemma 2 in hand, in the rest of this section we turn our attention to the case where $|aa'| < 2\delta$ and $|bb'| < 2\delta$. The intersection of two disks is called a *lens*. Define lens $L = D(a, 1) \cap D(b, 1)$ to be the region of distance at most 1 from a and b . Since (a, b) is a bichromatic diametral pair and $|ab| = 1$, all vertices of X_3, \dots, X_n lie in L , all vertices of X_1 lie in $D(b, 1)$, and all vertices of X_2 lie in $D(a, 1)$. Moreover since $|aa'| < 2\delta$ all vertices of X_1 lie in lens $L_1 = D(b, 1) \cap D(a, 2\delta)$, and since $|bb'| < 2\delta$ all vertices of X_2 lie in lens $L_2 = D(a, 1) \cap D(b, 2\delta)$. See Figure 2(a) for an illustration. In this setting, all vertices of \mathcal{X} lie in $L_1 \cup L_2$.

We fix a parameter $\omega = \frac{6\delta}{\sqrt{3}} - 1 \approx 0.815$. Let E be the set of all points whose total distance from a and b is at most $\omega + 2\delta$, i.e., $E = \{x \in \mathbb{R}^2 : |xa| + |xb| \leq \omega + 2\delta\}$. The boundary ∂E of E is an ellipse with foci a and b . Put $Q = (L_1 \cup L_2) \setminus E$, as depicted in Figure 2(b).

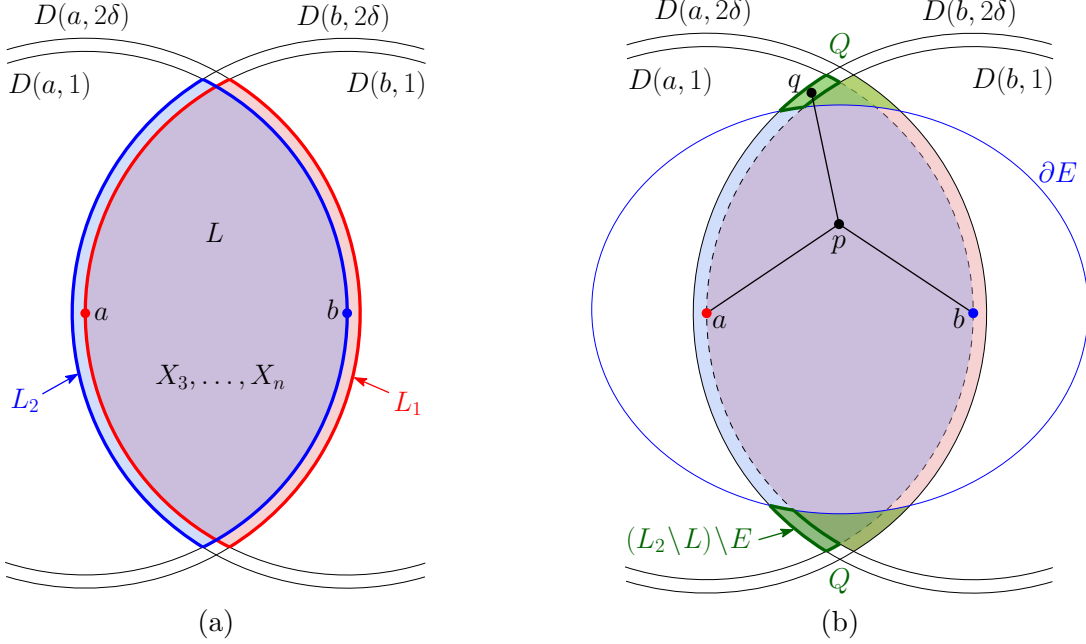


Figure 2: Illustration of (a) lenses L , L_1 , L_2 , and (b) the ellipse ∂E and the region Q .

Lemma 3. *For any point $q \in Q$ it holds that $|aq| > \omega$ and $|bq| > \omega$.*

Proof. The point q lies in $L_1 \cup L_2$, and thus $|aq| \leq 2\delta$ and $|bq| \leq 2\delta$. The point q does not lie in E , and hence $|aq| + |bq| > \omega + 2\delta$. Combining these inequalities yields $|aq| > \omega$ and $|bq| > \omega$. \square

The next “helper lemma” is the place where we use the Steiner ratio to obtain a lower bound on the length of the Steiner minimal tree of a subset of input vertices.

Lemma 4. *For any two points $q \in Q$ and $p \in \mathbb{R}^2$ it holds that $|pa| + |pb| + |pq| > 3\delta$.*

Proof. Put $P = \{a, b, q\}$. We are going to obtain a lower bound for the Steiner minimal tree of P by the minimum spanning tree of P . The point q lies in $(L_2 \setminus L) \setminus E$ or in $(L_1 \setminus L) \setminus E$ or in $L \setminus E$. If $q \in (L_2 \setminus L) \setminus E$, then bq is the longest side of the triangle $\triangle(a, b, q)$; this case is depicted in Figure 2(b). In this case $\text{Min-ST}(P)$ has edges aq and ab . Thus, $\text{len}(\text{Min-ST}(P)) = |aq| + |ab| > \omega + 1$ where the inequality holds by Lemma 3 and the fact that $|ab| = 1$. If $q \in (L_1 \setminus L) \setminus E$, then analogously we get $\text{len}(\text{Min-ST}(P)) > \omega + 1$. If $q \in L \setminus E$, then ab is the longest side of the triangle $\triangle(a, b, q)$, and hence $\text{Min-ST}(P)$ has edges aq and bq . Thus, $\text{len}(\text{Min-ST}(P)) = |aq| + |bq| > \omega + 2\delta > \omega + 1$ where the first inequality holds because $q \notin E$. Therefore, in all cases we have $\text{len}(\text{Min-ST}(P)) > \omega + 1$.

The union of the three segments pa , pb , and pq form a tree that connects the points of P , as in Figure 2(b). The length of this tree cannot be smaller than the length of Steiner minimal tree of P , and thus $|pa| + |pb| + |pq| \geq \text{len}(\text{SMT}(P))$. By using the Steiner ratio for three points (proved by Gilbert and Pollak [20]) we get

$$|pa| + |pb| + |pq| \geq \text{len}(\text{SMT}(P)) \geq \frac{\sqrt{3}}{2} \cdot \text{len}(\text{Min-ST}(P)) > \frac{\sqrt{3}}{2} \cdot (\omega + 1) = 3\delta. \quad \square$$

The following two lemmas consider two cases depending on whether or not a vertex of \mathcal{X} lies in Q . Both lemmas benefit from our helper Lemma 4. In Lemma 5 we use the helper lemma directly to obtain a lower bound on the maximum length of S_3 and D . In Lemma 6 we use the helper lemma indirectly to obtain a better upper bound on the length of T^* .

Lemma 5. *If at least one vertex of \mathcal{X} lies in Q then $\max\{\text{len}(S_3), \text{len}(D)\} \geq \delta \cdot \text{len}(T^*)$.*

Proof. Let q be any vertex of \mathcal{X} in Q . We have three cases: (i) $q \notin X_1$ and $q \notin X_2$, (ii) $q \in X_1$, (iii) $q \in X_2$. First consider case (i). After a suitable relabeling assume that $q \in X_3$. Consider an arbitrary representative vertex p_i from each X_i with $i \in \{4, \dots, n\}$. It is implied by Lemma 4 that

$$\sum_{i=4}^n |p_i a| + |p_i b| + |p_i q| > 3\delta(n-3).$$

Let x denote the point in $\{a, b, q\}$ that has the largest total distance to all points p_i . By bounding the maximum with the average, the total distance of p_i 's to x is at least $\delta(n-3)$. If $x = q$, then the star that connects q to p_4, \dots, p_n , a , and b has length

$$|qp_4| + \dots + |qp_n| + |qa| + |qb| > \delta(n-3) + \omega + \omega > \delta(n-1) \geq \delta \cdot \text{len}(T^*),$$

where the inequalities hold by Lemma 3, the fact that $w > \delta$, and (1). In this case $\text{len}(S_3) \geq \delta \cdot \text{len}(T^*)$, where the vertex c in the algorithm plays the role of q in the analysis. If $x = a$, then the star that connects a to p_4, \dots, p_n , q , and b has length

$$|ap_4| + \dots + |ap_n| + |aq| + |ab| > \delta(n-3) + \omega + 1 > \delta(n-1) \geq \delta \cdot \text{len}(T^*).$$

The length of this star is not larger than the length of D (by Lemma 1), and thus $\text{len}(D) \geq \delta \cdot \text{len}(T^*)$. If $x = b$, an analogous argument implies that D is a desired tree, proving the lemma for case (i).

Now consider case (ii) where $q \in X_1$. Our proof of this case is somewhat similar to that of case (i). Consider an arbitrary representative vertex p_i from each X_i with $i \in \{3, \dots, n\}$. It is implied by Lemma 4 that $\sum_{i=3}^n |p_i a| + |p_i b| + |p_i q| > 3\delta(n-2)$. Let x denotes the point in $\{a, b, q\}$ that has the largest total distance to all points p_i . The total distance of p_i 's to x is at least $\delta(n-2)$. If $x = q$, then the star connecting q to p_3, \dots, p_n , and b has length at least $\delta(n-2) + \omega > \delta(n-1) \geq \delta \cdot \text{len}(T^*)$. In this case $\text{len}(S_3) \geq \delta \cdot \text{len}(T^*)$, where c plays the role of q . If $x = a$, then the star connecting a to p_3, \dots, p_n , and b has length at least $\delta(n-2) + 1 > \delta \cdot \text{len}(T^*)$. The length of this star is not larger than that of D , and thus $\text{len}(D) \geq \delta \cdot \text{len}(T^*)$. If $x = b$, by an analogous argument D is a desired tree, proving the lemma for case (ii). Our proof of case (iii) is analogous to that of case (ii). \square

Lemma 6. *If no vertex of \mathcal{X} lies in Q then $\text{len}(D) \geq \delta \cdot \text{len}(T^*)$.*

Proof. In this case all vertices of \mathcal{X} lie in the region $R = (L_1 \cup L_2) \setminus Q$, as illustrated in Figure 3. Define the lens $L' = D(a, \delta) \cap D(b, \delta)$. Denote by l the lowest point of L' . Disregarding symmetry, denote by f the topmost intersection point of boundaries of $D(b, 2\delta)$ and E as in Figure 3. Consider the smallest disk D_l with center l that contains the entire region R . By basic geometry of circle-circle intersection (that boundaries of D_l and $D(b, 2\delta)$ intersect at two points) and circle-ellipse intersection (that boundaries of D_l and E intersect at four points—these four points are marked in Figure 3) we can verify that D_l passes through f . Thus f is the farthest point in R from l . Conversely, it follows from circle-circle intersection that l is the farthest point of L' from f . Therefore $|lf|$ is an upper bound for the distance between any point in L' and any point in R , and so is an upper bound on the length of any edge of T^* with an endpoint in L' .

Since f lies on ∂E we have $|af| + |bf| = \omega + 2\delta$, and since it lies on the boundary of $D(b, 2\delta)$ we have $|bf| = 2\delta$. Thus, $|af| = \omega$. Therefore f is an intersection point of boundaries of $D(a, \omega)$ and $D(b, 2\delta)$ that are centered at $(0, 0)$ and $(1, 0)$, respectively. The point l is an intersection point

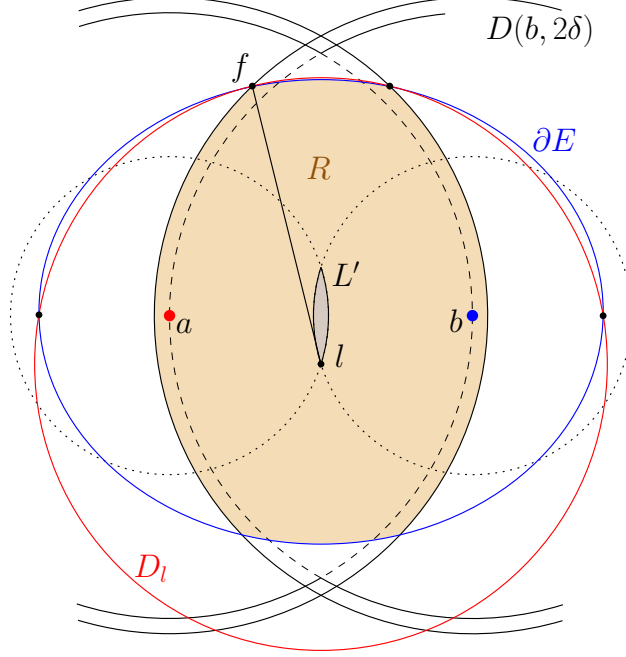


Figure 3: The length $|lf|$ is the largest possible for any edge having an endpoint in L' .

of boundaries of $D(a, \delta)$ and $D(b, \delta)$. By using the Pythagorean theorem and the circle equation, we can obtain the coordinates of f and l , and give the following expression for $|lf|$:

$$|lf| = \sqrt{\left(\frac{\omega^2 - 4\delta^2}{2}\right)^2 + \left(\sqrt{\omega^2 - \left(\frac{1 + \omega^2 - 4\delta^2}{2}\right)^2} + \sqrt{\delta^2 - \frac{1}{4}}\right)^2} \approx 0.9464 < 0.95.$$

We use $|lf|$ to obtain a better upper bound on the length of T^* . Let m be the number of neighborhoods X_i that lie entirely in the interior of L' . Notice that $0 \leq m \leq n - 2$ (X_1 and X_2 do not lie in L'). The number of edges of T^* , that are incident to these neighborhoods, is at least m . Since each such edge is of length at most $|lf|$ (< 0.95), we get

$$\text{len}(T^*) \leq (n - m - 1) + 0.95m = n - 0.05m - 1. \quad (2)$$

Now we compute a lower bound on the length of D in terms of m . One vertex of each neighborhood in L' is connected to a or b (whichever is the farthest) in D . Each such connection has length at least 0.5. There are $n - m - 2$ neighborhoods that have vertices outside L' (excluding X_1 and X_2). One vertex of each such neighborhood is connected to a or b (whichever is the farthest) in D . Each such connection has length at least δ . The points a and b are connected to each other in D . Thus

$$\begin{aligned} \text{len}(D) &\geq \delta(n - m - 2) + 0.5m + 1 \\ &= 0.524n - 0.024m - 0.048 \\ &> 0.524(n - 0.05m - 1) \\ &\geq 0.524 \cdot \text{len}(T^*), \end{aligned}$$

where the equality holds by plugging $\delta = 0.524$, and the last inequality holds by (2). \square

The cases considered in Lemmas 2, 5, and 6 ensure that the length of one of S_1, S_2, S_3 , and D is at least $\delta \cdot \text{len}(T^*)$. This concludes our analysis and proof of Theorem 1.

Remark. Although the stars S_1, S_2, S_3 are noncrossing, the double-star D can have crossing edges. Thus the algorithm of this section cannot be used for the Max-NC-ST problem.

Inclusion of bichromatic diameter. The simple 0.5-approximation algorithm of Chen and Dumitrescu [13] (described in Section 2) always includes the bichromatic diametral pair (a, b) in the solution. Chen and Dumitrescu argue that the approximation ratio of an algorithm that always includes a bichromatic diametral pair in the solution cannot be larger than $\sqrt{2 - \sqrt{3}} \approx 0.517$. We present an input instance that improves this upper bound to 0.5, thereby showing that the ratio of the simple algorithm is tight, in this sense.




Figure 4: Illustration of the upper bound 0.5 for inclusion of a bichromatic diametral pair.

Consider four points $p_0 = (0, 0)$, $p_1 = (1, 0)$, $p_2 = (2, 0)$, and $p_3 = (3 - 2\varepsilon, 0)$ for arbitrary small $\varepsilon > 0$, e.g. $\varepsilon = 1/n$. Our input instance consists of neighborhoods X_1, \dots, X_n where $X_1 = \{p_0, p_3\}$, $X_2 = \{p_2\}$, and each of X_3, \dots, X_n has exactly one point that is placed at distance at most ε from p_1 ; see Figure 4. In this setting, (p_0, p_2) is the unique bichromatic diametral pair. Consider any tree T that contains the bichromatic diameter p_0p_2 (this means that p_3 is not in T). Any edge of T incident to X_3, \dots, X_n has length at most $1 + \varepsilon$. Therefore $\text{len}(T) \leq 2 + (1 + \varepsilon)(n - 2) < n + 1$. Now consider the tree T^* that does not contain p_0p_2 but connects each of X_2, \dots, X_n to p_3 . The length of T^* is at least $(1 - 2\varepsilon) + (2 - 3\varepsilon)(n - 2) > 2n - 6$. This would establish the upper bound 0.5 on the approximation ratio because $\frac{\text{len}(T)}{\text{len}(T^*)} \leq \frac{n+1}{2n-6}$ which tends to $1/2$ in the limit.

4 Maximum noncrossing spanning tree

In this section we prove Theorem 2. Put $\delta = 0.519$.

We present a δ -approximation algorithm for the Max-NC-ST problem. To facilitate comparisons, we give a proper citation for each part of the proof that shares similarities with previous algorithms.

Let $P \subset \mathbb{R}^2$ be the given point set of size n and let (u, v) be a diametral pair of P , throughout this section. After a suitable scaling assume that $|uv| = 1$. Let T^* be a longest noncrossing spanning tree of P . If ab is a longest edge of T^* then $|ab| \leq |uv| = 1$, and thus

$$\text{len}(T^*) \leq (n - 1)|ab| \leq (n - 1)|uv| \leq n - 1. \quad (3)$$

Our plan is to construct a noncrossing spanning tree for P of length at least $\delta \cdot \text{len}(T^*)$. For a point $p \in P$, we denote by S_p the star that connects p to all other points of P . We start with the following simple lemma, proved in [17], which comes in handy in our construction.

Lemma 7. *For any two points p and q in P it holds that $\max\{\text{len}(S_p), \text{len}(S_q)\} \geq \frac{n}{2}|pq|$.*

Proof. By bounding the maximum with the average and then using the triangle inequality we get:

$$\max\{\text{len}(S_p), \text{len}(S_q)\} \geq \frac{1}{2}(\text{len}(S_p) + \text{len}(S_q)) = \frac{1}{2} \sum_{r \in P} (|pr| + |rq|) \geq \frac{1}{2} \sum_{r \in P} |pq| = \frac{n}{2}|pq|. \quad \square$$

From Lemma 7 and (3) we have $\max\{\text{len}(S_u), \text{len}(S_v)\} \geq \frac{n}{2}|uv| = \frac{n}{2} > \frac{1}{2} \cdot \text{len}(T^*)$. Therefore, the longer of S_u and S_v is a 0.5 approximation of T^* . As pointed out by Alon et al. [2] the longest star may not give an approximation ratio better than 0.5. To establish better ratios we need to consider trees that are not necessarily stars, and we need to incorporate more powerful ingredients.

Now we describe our δ -approximation algorithm. Our approximation ratio is with respect to the length of the longest noncrossing spanning tree T^* .

Algorithm approach. Guess a longest edge ab of T^* , say by trying all pairs of points in P . For each guess ab , construct seven noncrossing spanning trees as described in the rest of this section. Then report the longest tree, over all guesses for ab .

From now on we assume that ab is a longest edge of T^* . In the rest of this section we describe how to construct the seven noncrossing spanning trees in such a way that the length of the longest tree is at least $\delta \cdot \text{len}(T^*)$. Fix a parameter $d = \frac{1}{2\delta}$.

Lemma 8. *If $|ab| \leq d$ then $\max\{\text{len}(S_u), \text{len}(S_v)\} \geq \delta \cdot \text{len}(T^*)$.*

Proof. By Lemma 7, the fact that $|uv| = 1 = 2\delta \cdot d$, our assumption that $|ab| \leq d$, and (3) we get

$$\max\{\text{len}(S_u), \text{len}(S_v)\} \geq \frac{n}{2}|uv| = \delta \cdot nd \geq \delta \cdot n|ab| > \delta \cdot \text{len}(T^*). \quad \square$$

Having Lemma 8 in hand, in the rest of this section we assume that $d \leq |ab| \leq 1$. After a suitable rotation and a translation assume that $a = (0, 0)$ and $b = (|ab|, 0)$. Define the lens $L = D(a, 1) \cap D(b, 1)$; see Figure 5(a). Since the diameter of P is 1, all points of P lie in L . Define $\lambda = \frac{6\delta}{\sqrt{3}} + 1 - |ab|$. We use this parameter along with the Steiner ratio to take care of the situation where some input points lie far from a and b —this situation is a bottleneck for previous algorithms.

Define $E_1 = \{x \in \mathbb{R}^2 : |xa| + |xb| \leq \lambda\}$. The boundary ∂E_1 is an ellipse with foci a and b . Put $Q = L \setminus E_1$, as depicted in Figure 5(a). Lemma 10, below, plays an important role in improving the approximation ratio. We keep its proof short as it is somewhat similar to our proof of Lemma 4.

Lemma 9. *For any point $q \in Q$ it holds that $|aq| > \lambda - 1$ and $|bq| > \lambda - 1$.*

Proof. Since $q \in L$, $|aq| \leq 1$ and $|bq| \leq 1$. Since $q \notin E_1$, $|aq| + |bq| > \lambda$. Combining these inequalities yields $|aq| > \lambda - 1$ and $|bq| > \lambda - 1$. \square

The next “helper lemma” is the place where we use the Steiner ratio.

Lemma 10. *For any two points $q \in Q$ and $p \in \mathbb{R}^2$ it holds that $|pa| + |pb| + |pq| > 3\delta$.*

Proof. Put $P' = \{a, b, q\}$, and define $L' = D(a, |ab|) \cap D(b, |ab|)$ as in Figure 5(a). If $q \in L' \setminus E_1$ then $\text{Min-ST}(P')$ has edges aq and bq , and thus $\text{len}(\text{Min-ST}(P')) = |aq| + |bq| > \lambda \geq 6\delta/\sqrt{3}$. If $q \in L \setminus L' \setminus E_1$ then $\text{Min-ST}(P')$ has the edge ab together with the shorter of aq and bq , which we may assume it is aq by symmetry; this case is depicted in Figure 5(a). Thus $\text{len}(\text{Min-ST}(P')) = |aq| + |ab| > \lambda - 1 + |ab| = 6\delta/\sqrt{3}$, where the inequality is implied by Lemma 9. Therefore in all cases we have $\text{len}(\text{Min-ST}(P')) > 6\delta/\sqrt{3}$. It follows from the Steiner ratio (for three points) that

$$|pa| + |pb| + |pq| \geq \text{len}(\text{SMT}(P')) \geq \frac{\sqrt{3}}{2} \cdot \text{len}(\text{Min-ST}(P')) > \frac{\sqrt{3}}{2} \cdot \frac{6\delta}{\sqrt{3}} = 3\delta. \quad \square$$

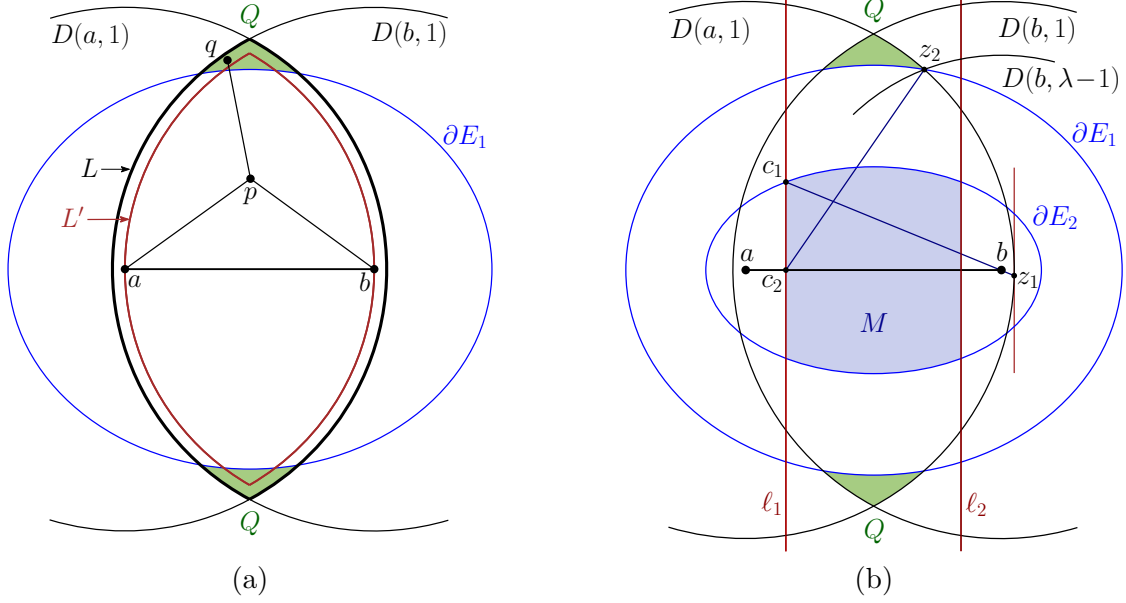


Figure 5: (a) Illustration of the case where $P \cap Q \neq \emptyset$ and $q \in L \setminus L' \setminus E_1$. (b) Illustration of the longest edges starting from c_1 and c_2 when $P \cap Q = \emptyset$.

Lemmas 11 and 12 distinguish between two cases where $P \cap Q \neq \emptyset$ and $P \cap Q = \emptyset$. Both lemmas benefit from our helper Lemma 10. A combination of these three lemmas lead to a significant improvement on the approximation ratio. In Lemma 11 we use the helper lemma directly to obtain a long tree for the case $P \cap Q \neq \emptyset$ which is a bottleneck case for previous algorithms. In Lemma 12 we use the helper lemma indirectly to obtain a better upper bound for the length of T^* .

Lemma 11. *If $P \cap Q \neq \emptyset$ then there is a noncrossing spanning tree for P of length at least $\delta \cdot \text{len}(T^*)$.*

Proof. Consider any point $q \in P \cap Q$. By Lemma 9, $|aq| \geq \lambda - 1 = 6\delta/\sqrt{3} - |ab|$ and $|bq| \geq \lambda - 1 = 6\delta/\sqrt{3} - |ab|$. Recall that $1/2\delta = d \leq |ab| \leq 1$ and $\delta = 0.519$. Then

$$\min\{|ab|, |aq|, |bq|\} \geq \min\left\{\frac{1}{2\delta}, \frac{6\delta}{\sqrt{3}} - |ab|, \frac{6\delta}{\sqrt{3}} - |ab|\right\} \geq \min\left\{\frac{1}{2\delta}, \frac{6\delta}{\sqrt{3}} - 1, \frac{6\delta}{\sqrt{3}} - 1\right\} > \delta. \quad (4)$$

Let p_4, \dots, p_n denote the points in $P \setminus \{a, b, q\}$. It is implied by Lemma 10 that

$$\sum_{i=4}^n |p_i a| + |p_i b| + |p_i q| > 3\delta(n - 3). \quad (5)$$

Denote by x the point in $\{a, b, q\}$ that has the largest total distance to p_4, \dots, p_n , and denote by y and z the other two points. It is implied from (5) that the total distance of p_i 's to x is at least $\delta(n - 3)$. In this setting, the star S_x is a desired noncrossing tree because

$$\text{len}(S_x) = |xp_4| + \dots + |xp_n| + |xy| + |xz| > \delta(n - 3) + \delta + \delta = \delta(n - 1) \geq \delta \cdot \text{len}(T^*),$$

where the first inequality is implied by (4) and the second inequality is implied by (3). \square

Fix parameters $\omega = 0.16$ and $\hat{\beta} = 0.44$ and define

$$\hat{\alpha} = \frac{2\delta + 3\omega - 2}{\omega - 1} - \hat{\beta} \quad \gamma = \frac{(2\delta + \hat{\alpha} - 1)|ab|}{\hat{\alpha}}.$$

Subdivide L into three parts by two vertical lines ℓ_1 and ℓ_2 at $\omega|ab|$ and $(1-\omega)|ab|$, respectively (a similar subdivision is used in [8, 11, 17]). Define $E_2 = \{x \in \mathbb{R}^2 : |xa| + |xb| \leq \gamma\}$. Let M be the part of $L \cap E_2$ between ℓ_1 and ℓ_2 . See Figure 5(b). Let α be the fraction of points in $L \setminus E_2$, and let β be the fraction of points in M (the denominator of the fractions is n). Observe that $1 - (\alpha + \beta)$ fraction of points lie in parts of $L \cap E_2$ that are to the left of ℓ_1 and to the right of ℓ_2 .

The next lemma is the place where we use the noncrossing property of T^* . Since Lemma 11 takes care of points in Q , we turn our attention to the case where no point of P lies in Q . This constraint helps us to obtain a better upper bound on the length of T^* , which in turn leads to a better lower bound on the maximum length of the stars S_u and S_v . Recall that (u, v) is a diametral pair of P and that $|uv| = 1$.

Lemma 12. *If $P \cap Q = \emptyset$ and $\beta \geq \hat{\beta}$ then $\max\{\text{len}(S_u), \text{len}(S_v)\} \geq \delta \cdot \text{len}(T^*)$.*

Proof. Recall that ab is a longest edge of T^* . Since $P \cap Q = \emptyset$, all points of P lie in $L \cap E_1$. Let c_1 be the top intersection point of ℓ_1 and ∂E_2 , and let c_2 be the intersection point of ℓ_1 and ab , as in Figure 5(b). If we ignore symmetry then it follows from convexity that the longest possible edge of T^* (not intersecting ab) with an endpoint in M has either c_1 or c_2 as an endpoint. The longest edge starting from c_1 and not intersecting ab ends at the intersection point z_1 of the line through c_1 and b with the boundary of L . The longest edge starting from c_2 ends at the intersection point z_2 of the boundaries of L and E_1 . See Figure 5(b) for an illustration of these edges. Observe that z_2 is also an intersection point of the boundaries of $D(a, 1)$ and $D(b, \lambda - 1)$. By using the Pythagorean theorem and the ellipse and circle equations, we can obtain the coordinates of c_1, c_2, z_1, z_2 , and give the following expressions for $|c_1 z_1|$ and $|c_2 z_2|$:

$$|c_2 z_2| = \sqrt{\left(\frac{1 + |ab|^2 - (\lambda - 1)^2}{2|ab|} - \omega|ab|\right)^2 + 1 - \left(\frac{1 + |ab|^2 - (\lambda - 1)^2}{2|ab|}\right)^2}$$

$$|c_1 b| = \sqrt{(1 - \omega)^2 |ab|^2 + \frac{(\gamma^2 - |ab|^2)}{\gamma^2} \left(\left(\frac{\gamma}{2}\right)^2 - \left(\frac{|ab|}{2} - \omega|ab|\right)^2\right)}$$

$$|c_1 z_1| \leq |c_1 b| + \frac{(1 - |ab|)|c_1 b|}{(1 - \omega)|ab|}$$

where the upper bound on $|c_1 z_1|$ is obtained by extending $c_1 z_1$ to intersect the vertical line through the rightmost point of L ; this line is depicted in Figure 5(b).

Both $|c_1 z_1|$ and $|c_2 z_2|$ depend only on $|ab|$, and thus we denote them by functions $f_1(|ab|)$ and $f_2(|ab|)$, respectively. By considering the plots of these functions (Figure 6) it follows that for any $|ab|$ in interval $[d, 1]$ we have $f_1(d) \geq f_1(|ab|)$ and $f_1(|ab|) > f_2(|ab|)$. Therefore, the longest possible edge starting in M has length at most $f_1(d)$. By plugging the chosen constants in the above formula we get $f_1(d) \approx 0.913117 < 0.914$. Hence we obtain the following upper bound for the length of T^* :

$$\text{len}(T^*) \leq (1 - \hat{\beta})n + 0.914\hat{\beta}n = (1 - 0.086\hat{\beta})n < 0.963n.$$

It is implied from Lemma 7 that $\max\{\text{len}(S_u), \text{len}(S_v)\} \geq n/2$. Therefore,

$$\frac{\max\{\text{len}(S_u), \text{len}(S_v)\}}{\text{len}(T^*)} \geq \frac{0.5n}{0.963n} > 0.519 = \delta. \quad \square$$

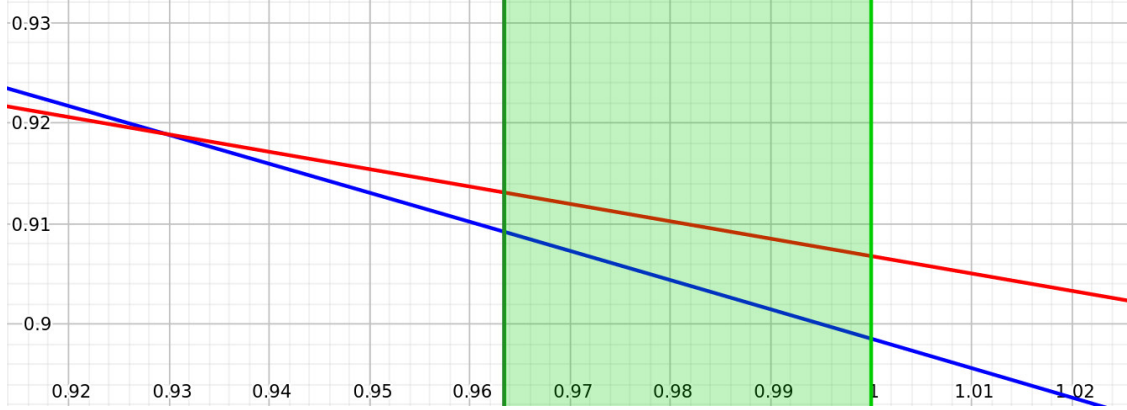


Figure 6: Plots of $|c_1z_1|$ (in red) and $|c_2z_2|$ (in blue) over $|ab|$ in interval $[d, 1]$ (in green).

The following two lemmas do not use the constraint $P \cap Q = \emptyset$. Lemma 13 is adapted from one by Cabello et al. [11], and Lemma 14 is adapted from one by Biniiaz et al. [8].

Lemma 13. *If $\alpha \geq \hat{\alpha}$ then $\max\{\text{len}(S_a), \text{len}(S_b)\} \geq \delta \cdot \text{len}(T^*)$.*

Proof. The proof is somewhat similar to that of Lemma 7, except now we have a better lower bound for points in $L \setminus E_2$, which is γ . Thus

$$\max\{\text{len}(S_a), \text{len}(S_b)\} \geq \frac{\text{len}(S_a) + \text{len}(S_b)}{2} \geq \frac{\alpha n \cdot \gamma + (1 - \alpha)n \cdot |ab|}{2} = \frac{n(|ab| + \alpha(\gamma - |ab|))}{2}.$$

This and the fact that $\text{len}(T^*) < n|ab|$ by (3), imply that


$$\frac{\max\{\text{len}(S_a), \text{len}(S_b)\}}{\text{len}(T^*)} > \frac{n(|ab| + \alpha(\gamma - |ab|))}{2n|ab|} \geq \frac{|ab| + \hat{\alpha}(\gamma - |ab|)}{2|ab|} = \delta. \quad \square$$

Lemma 14. *If $\alpha \leq \hat{\alpha}$ and $\beta \leq \hat{\beta}$ then there is a noncrossing spanning tree for P of length at least $\delta \cdot \text{len}(T^*)$.*

Proof. We keep the description short. We construct two trees T_a and T_b such that the longer one is a desired tree. We describe the construction for T_a ; the construction of T_b is analogous.

See the figure to the right for an illustration. Start by connecting all points to the right of ℓ_2 to a (red edges). Each such edge has length at least $(1 - \omega)|ab|$. Let ap_1, \dots, ap_m be these edges in radial order around a as in the figure; notice that $m \geq 1$ because b is to the right of ℓ_2 . Now we connect the points to the left of ℓ_1 : connect all points below ap_1 (that are not above ab) to p_1 , connect all points above ap_m (that are not below ab) to p_m , and connect all points between two consecutive edges ap_i and ap_{i+1} to p_i for $1 \leq i \leq m - 1$ (blue edges). Each new edge has length at least $(1 - 2\omega)|ab|$.

Finally we connect the points in the region between ℓ_1 and ℓ_2 . Let β' be the fraction of points in this region. This region is subdivided into subregions by the current (red and blue) edges of T_a . Each subregion is bounded by one or two edges of T_a in such a way that at least one edge is fully visible from the interior of the subregion. Connect all points in



each subregion to the endpoint (of the visible edge) that gives a larger total distance; these new edges are shown in green. By an argument similar to the proof of Lemma 7 one can show that the total length of the new edges is at least $\beta'n(1-2\omega)|ab|/2$.

By the definition of α , β , and β' it holds that $\beta' \leq \alpha + \beta$. Since $\alpha \leq \hat{\alpha}$ and $\beta \leq \hat{\beta}$, we have $\beta' \leq \hat{\alpha} + \hat{\beta}$. The total fraction of points to the left of ℓ_1 and to the right of ℓ_2 is $1 - \beta'$. Thus

$$\begin{aligned} \text{len}(T_a) + \text{len}(T_b) &\geq (1 - \beta')n(2 - 3\omega)|ab| + \beta'n(1 - 2\omega)|ab| \\ &= (2 - 3\omega + (\omega - 1)\beta')n|ab| \\ &\geq (2 - 3\omega + (\omega - 1)(\hat{\beta} + \hat{\alpha}))n|ab|, \end{aligned}$$

where the last inequality holds because $\omega - 1 < 0$. Therefore

$$\max \{\text{len}(T_a) + \text{len}(T_b)\} \geq \frac{2 - 3\omega + (\omega - 1)(\hat{\beta} + \hat{\alpha})}{2} \cdot n|ab| = \delta \cdot n|ab| > \delta \cdot \text{len}(T^*). \quad \square$$

The cases considered in Lemmas 8, 11, 12, 13, and 14 ensure that at least one of S_u , S_v , S_a , S_b , T_a , T_b , and S_q (introduced in Lemma 11) has length at least $\delta \cdot \text{len}(T^*)$. This concludes our analysis and proof of Theorem 2.

5 Conclusions

A natural open problem is to further improve the approximation ratios for the Max-ST-NB and the Max-NC-ST problems. We believe there is a possibility to slightly improve both bounds by discretization. For example in our Max-ST-NB (resp. Max-NC-ST) algorithm if the points in the lens L' (resp. the region M) are close to ab then one could obtain a better upper bound on the length of T^* otherwise obtain a better lower bound on the length of the double-star D (resp. the star S_a or S_b). We conjecture that our algorithms have better approximation guarantees, however this requires tight analysis.

We obtained the improved approximation ratios by employing the Steiner ratio which was not used in this context earlier. It would be interesting to see if the Steiner ratio could be used to design better algorithms for other related problems.

References

- [1] P. K. Agarwal, J. Matousek, and S. Suri. Farthest neighbors, maximum spanning trees and related problems in higher dimensions. *Comput. Geom.*, 1:189–201, 1991.
- [2] N. Alon, S. Rajagopalan, and S. Suri. Long non-crossing configurations in the plane. *Fundam. Inform.*, 22(4):385–394, 1995. Also in *SoCG* 1993.
- [3] E. M. Arkin and R. Hassin. Approximation algorithms for the geometric covering salesman problem. *Discret. Appl. Math.*, 55(3):197–218, 1994.
- [4] T. Asano, B. K. Bhattacharya, J. M. Keil, and F. F. Yao. Clustering algorithms based on minimum and maximum spanning trees. In *Proceedings of the 4th Annual Symposium on Computational Geometry (SoCG)*, pages 252–257, 1988.

- [5] M. Bern and D. Eppstein. Approximation algorithms for geometric problems. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*, page 296–345. PWS Publishing, 1996.
- [6] B. K. Bhattacharya and G. T. Toussaint. Efficient algorithms for computing the maximum distance between two finite planar sets. *J. Algorithms*, 4(2):121–136, 1983.
- [7] A. Biniáz. Improved approximation ratios for two euclidean maximum spanning tree problems. *arXiv:2010.03870*.
- [8] A. Biniáz, P. Bose, K. Crosbie, J. D. Carufel, D. Eppstein, A. Maheshwari, and M. Smid. Maximum plane trees in multipartite geometric graphs. *Algorithmica*, 81(4):1512–1534, 2019.
- [9] A. Biniáz, P. Bose, D. Eppstein, A. Maheshwari, P. Morin, and M. Smid. Spanning trees in multipartite geometric graphs. *Algorithmica*, 80(11):3177–3191, 2018.
- [10] V. Blanco, E. Fernández, and J. Puerto. Minimum spanning trees with neighborhoods: Mathematical programming formulations and solution methods. *Eur. J. Oper. Res.*, 262(3):863–878, 2017.
- [11] S. Cabello, A. Choudhary, M. Hoffmann, K. Klost, M. M. Reddy, W. Mulzer, F. Schröder, and J. Tkadlec. A better approximation for longest noncrossing spanning trees. In *36th European Workshop on Computational Geometry (EuroCG)*, 2020.
- [12] S. Cabello, M. Hoffmann, K. Klost, W. Mulzer, and J. Tkadlec. Long plane trees. *arXiv:2101.00445*.
- [13] K. Chen and A. Dumitrescu. On the longest spanning tree with neighborhoods. *Discrete Mathematics, Algorithms and Applications*, doi:10.1142/S1793830920500676, 2020. Also in Proceedings of the 12th International Workshop on Frontiers in Algorithmics (FAW), pages 15–28, 2018.
- [14] M. de Berg, O. Cheong, M. J. van Kreveld, and M. H. Overmars. *Computational geometry: algorithms and applications, 3rd Edition*. Springer, 2008.
- [15] R. Dorrigiv, R. Fraser, M. He, S. Kamali, A. Kawamura, A. López-Ortiz, and D. Seco. On minimum- and maximum-weight minimum spanning trees with neighborhoods. *Theory Comput. Syst.*, 56(1):220–250, 2015.
- [16] D. Du, F. Hwang, and E. Yao. The Steiner ratio conjecture is true for five points. *Journal of Combinatorial Theory, Series A*, 38(2):230–240, 1985.
- [17] A. Dumitrescu and C. D. Tóth. Long non-crossing configurations in the plane. *Discret. Comput. Geom.*, 44(4):727–752, 2010. Also in *STACS 2010*.
- [18] D. Eppstein. Spanning trees and spanners. In J. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 425–461. North Holland / Elsevier, 2000.
- [19] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34(3):596–615, 1987. Also in *FOCS 1984*.
- [20] E. N. Gilbert and H. O. Pollak. Steiner minimal trees. *SIAM Journal on Applied Mathematics*, 16(1):1–29, 1968.

- [21] E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 585–594, 2003.
- [22] N. Innami, B. H. Kim, Y. Mashiko, and K. Shiohama. The Steiner ratio conjecture of Gilbert-Pollak may still be open. *Algorithmica*, 57(4):869–872, 2010.
- [23] A. O. Ivanov and A. A. Tuzhilin. The Steiner ratio Gilbert-Pollak conjecture is still open - clarification statement. *Algorithmica*, 62(1-2):630–632, 2012.
- [24] M. Löffler and M. J. van Kreveld. Largest and smallest convex hulls for imprecise points. *Algorithmica*, 56(2):235–269, 2010.
- [25] J. 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.
- [26] J. 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.
- [27] C. L. Monma, M. Paterson, S. Suri, and F. F. Yao. Computing Euclidean maximum spanning trees. *Algorithmica*, 5(3):407–419, 1990.
- [28] H. Pollak. Some remarks on the Steiner problem. *Journal of Combinatorial Theory, Series A*, 24(3):278–295, 1978.
- [29] J. Rubinstein and D. Thomas. The Steiner ratio conjecture for six points. *Journal of Combinatorial Theory, Series A*, 58(1):54–77, 1991.
- [30] M. J. van Kreveld and M. Löffler. Approximating largest convex hulls for imprecise points. *J. Discrete Algorithms*, 6(4):583–594, 2008.
- [31] Y. Yang, M. Lin, J. Xu, and Y. Xie. Minimum spanning tree with neighborhoods. In *Proceedings of the 3rd International Conference Algorithmic Aspects in Information and Management (AAIM)*, pages 306–316, 2007.