

Packing Plane Spanning Graphs with Short Edges in Complete Geometric Graphs*

Oswin Aichholzer[†] Thomas Hackl[†] Matias Korman[‡]
Alexander Pilz[†] Günter Rote[§] André van Renssen[¶]
Marcel Roeloffzen^{||} Birgit Vogtenhuber[†]

August 16, 2019

Abstract

Given a set of points in the plane, we want to establish a connected spanning graph between these points, called connection network, that consists of several disjoint layers. Motivated by sensor networks, our goal is that each layer is connected, spanning, and plane. No edge in this connection network is too long in comparison to the length needed to obtain a spanning tree.

We consider two different approaches. First we show an almost optimal centralized approach to extract two layers. Then we consider a distributed model in which each point can compute its adjacencies using only information about vertices at most a predefined distance away. We show a constant factor approximation with respect to the length of the longest edge in the graphs. In both cases the obtained layers are plane.

1 Introduction

Given a set S of n points in the plane and an integer k , we are interested in finding k edge-disjoint non-crossing spanning graphs H_1, H_2, \dots, H_k on S such that the length $\text{BE}(H_1 \cup H_2 \cup \dots \cup H_k)$ of the *bottleneck edge* (the longest edge which is used) is as short as possible. Each subgraph H_i is referred to as a

*A preliminary version of this paper was presented in the 27th International Symposium on Algorithms and Computation (ISAAC 2016), Sydney, Australia [2].

[†]Institute of Software Technology, Graz University of Technology, Austria. {oaich|thackl|apilz|bvogt}@ist.tugraz.at

[‡]Department of Computer Science, Tufts University, Medford, MA, USA. matias.korman@tufts.edu

[§]Institut für Informatik, Freie Universität Berlin, Germany. rote@inf.fu-berlin.de

[¶]University of Sydney, Sydney, Australia. andre.vanrenssen@sydney.edu.au

^{||}Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands. m.j.m.roeloffzen@tue.nl

layer of the complete graph G on the n points. We require each layer to be non-crossing, but edges from different layers are allowed to cross each other. For $k = 1$, the minimum spanning tree $\text{MST}(S)$ solves the problem: its longest edge $\text{BE}(\text{MST}(S))$ is a lower bound on the bottleneck edge of any spanning subgraph, and it is non-crossing. For larger k , we take $\text{BE}(\text{MST}(S))$ as the yardstick and measure the solution quality in terms of $\text{BE}(\text{MST}(S))$ and k .

Although we find the problem to be of its own (theoretical) interest, this particular variation comes motivated from the field of sensor networks. In sensor networks, the energy consumed in transmission drastically grows as the distance between the two points increases [5, 6]. Since we want to avoid high energy consumption, it is desirable to apply the bottleneck criterion in order to minimize the maximum length of the whole network.

Once the network is built, we want to send messages through it without having to store the network explicitly at each node. One of the most commonly used methods for doing so is called *face routing* [7], which uses only local information and guarantees delivery as long as the underlying network is plane. In fact, most local routing algorithms can only route on plane graphs. Extending these algorithms for non-plane graphs is a long-standing open problem in the field. In this paper, we provide a different way to avoid this obstacle. Rather than one plane graph, we construct several disjoint plane spanning graphs. If we split all the messages among the different layers we can potentially spread the load among a larger number of edges.

Previous Work. This problem falls into the family of *graph packing* problems, where we are given a graph $G = (V, E)$ and a family \mathcal{F} of subgraphs of G . The aim is to pack as many pairwise edge-disjoint subgraphs $H_1 = (V, E_1), H_2 = (V, E_2), \dots$ as possible into G .

A related problem is the *decomposition* of G . In this case, we also look for disjoint subgraphs but require that $\bigcup_i E_i = E$. For example, there are known characterizations of when we can decompose the complete graph of n points into paths [10] (for n even) and stars [9] (for n odd). Dor and Tarsi [4] showed that to determine whether we can decompose a graph G into subgraphs isomorphic to a given graph H is NP-complete. Concerning graph packing, Aichholzer *et al.* [1] showed that we can pack $\Omega(\sqrt{n})$ edge-disjoint plane spanning trees in the complete graph on any set S of n points. This bound has been improved to $\lfloor n/3 \rfloor$ by Biniaz and García [3]. Note that a trivial $\lfloor n/2 \rfloor$ upper bound follows from the number of edges in the complete graph. Thus, the latter result is close to optimal.

In our case, the graph G is the complete graph on a given point set S , and \mathcal{F} consists of all plane spanning graphs of G . In addition to proving results for a large (fixed) number of layers, we are interested in minimizing a geometric constraint (Euclidean length of the longest edge among the selected graphs of \mathcal{F}). To the best of our knowledge, this is the first packing problem of such type.

Results. Recall that both the point set S and the integer k are given and that we aim to find k edge-disjoint connected plane spanning graphs H_1, H_2, \dots, H_k on S such that the length $\text{BE}(H_1 \cup H_2 \cup \dots \cup H_k)$ of the *bottleneck edge* (the longest edge that is used) is minimized.

We give two different approaches to solve the problem. In Section 2 we give a construction for two spanning trees, i.e., $k = 2$. This construction is centralized in a classic model that assumes that the positions of all points are known and computed in a single place. Our construction creates two trees and guarantees that all edges (except possibly one) have length at most $2\text{BE}(\text{MST}(S))$. The remaining edge has length at most $3\text{BE}(\text{MST}(S))$. We complement this construction with a matching worst-case lower bound that shows that for two spanning trees this is the shortest length the longest edge in the graphs can have.

In Section 3 we use a different approach to construct k edge-disjoint connected plane spanning graphs (not necessarily trees). The construction works for any $k \leq n/12$ in an almost local fashion, i.e., using only information about vertices at most a certain maximum distance away. The only global information that is needed is β : $\text{BE}(\text{MST}(S))$ or some upper bound on it. Each point of S can compute its adjacencies by only looking at nearby points, namely, those at distance $O(k\beta)$.

A simple adversary argument shows that it is impossible to construct spanning networks locally without knowing $\text{BE}(\text{MST}(S))$ (or an upper bound). The lower bound of Section 2 shows that a neighborhood of radius $\Omega(k\text{BE}(\text{MST}(S)))$ may be needed for the network, so we conclude that our construction is asymptotically optimal in terms of the neighborhood.

For simplicity, throughout the paper we make the usual general position assumption that no three points are collinear. Without this assumption, it might be impossible to obtain more than a single plane layer (for example, when all points lie on a line). Note however, that if collinear and partially overlapping edges are considered as non-crossing, our algorithms do not require the point set to be in general position.

2 Centralized Construction

In this section we present a centralized algorithm to construct two layers. We start with some properties on the minimum spanning tree of a set of points.

Lemma 1. *If $|uw| > \max\{|uv|, |vw|\}$ for three points $u, v, w \in S$, the edge uw does not belong to $\text{MST}(S)$.*

Proof. This is a special case of the more general well-known statement that the longest edges of any cycle in a graph, if it is unique, does not belong to its minimum spanning tree: The greedy algorithm would first pick all other edges of the cycle unless their endpoints are already connected. Thus, when the algorithm looks at the longest edge, its endpoints are already connected, and the edge is not included in the minimum spanning tree. \square

Lemma 2. *Let S be a finite set of points in the plane and let uv and vw be two edges of $\text{MST}(S)$. Then the triangle uvw does not contain any other point of S .*

Proof. Suppose for the sake of contradiction that the triangle uvw contains a point $p \in S$. Then the sum of the angles vpu and vpw is at least π ; see Figure 2 (a). Hence, one of these angles, say, vpu is at least $\pi/2$. But then vu is the longest edge in the triangle vpu , and by Lemma 1, vu cannot belong to $\text{MST}(S)$, a contradiction. \square




Figure 1: (a) Proof of Lemma 2. (b) Proof of Lemma 3: The neighbors of v_i cannot lie outside the wedge W_i defined by its two siblings in $\text{MST}(S)$.

Lemma 3. *Let S be a finite set of points in the plane. Let $v \in S$ be a point with $k \geq 3$ neighbors v_0, \dots, v_{k-1} in $\text{MST}(S)$ in counterclockwise order. Then for every triple (v_{i-1}, v_i, v_{i+1}) (indices modulo k), the neighbors of v_i in $\text{MST}(S)$ are inside the wedge W_i that is bounded by the rays vv_{i-1} and vv_{i+1} and contains the edge vv_i .*

Proof. Assume for the sake of contradiction that v_i has a neighbor u in $\text{MST}(S)$ that does not lie in W_i ; see Figure 2 (b). Then the edge v_iu intersects one of the boundary rays of W_i , say, vv_{i+1} . As $\text{MST}(S)$ is plane, the edge vv_{i+1} does not intersect the edge v_iu . Hence, the point v_{i+1} lies in the triangle vv_iu . As vv_i and v_iu are in $\text{MST}(S)$, this contradicts Lemma 2. \square

We denote by $\text{MST}^2(S)$ the square of $\text{MST}(S)$, the graph connecting all pairs of points of S that are at distance at most 2 in $\text{MST}(S)$. We call the edges of $\text{MST}(S)$ *short* edges and all remaining edges of $\text{MST}^2(S)$ *long* edges. For every long edge uw , the points u and w have a unique common neighbor v in $\text{MST}(S)$, which we call the *witness* of uw . We define the *wedge* of uw to be the area that is bounded by the rays vu and vw and contains the segment uw .

We now characterize edge crossings in $\text{MST}^2(S)$; see Figure 2.

Lemma 4. *Let S be a finite set of points in the plane. Two edges e and f of $\text{MST}^2(S)$ cross if and only if one of the following two conditions is fulfilled:*




Figure 2: The different cases of Lemma 4. Short edges are solid and long edges are dashed. (a) Condition 1, showing two options e and e' , (b) Condition 2, and (c) the contradiction used in the proof.

1. At least one of e and f is a long edge with witness v and wedge W , and the other edge has v as an endpoint and lies inside W .
2. Both e and f are long edges with the same witness v , their wedges intersect, but none is contained in the other.

Proof. Clearly, if both e and f are short, they cannot cross. Without loss of generality assume that $f = uw$ is a long edge with witness v and wedge W . If e is incident to v , then it must lie in W in order to cross f , and we satisfy Condition 1.

In the remaining case, $e = xz$ with $x, z \in S \setminus \{u, v, w\}$. By Lemma 2, x and z cannot lie in the triangle uvw ; hence, e must cross one of the MST edges uv or vw in addition to the edge $f = uw$. It follows that e cannot be short, and it has some witness y and some wedge W' . We distinguish three possibilities for y :

(i) If $y = v$, we satisfy Condition 2: W' is not contained in W because e crosses uv or vw , and by swapping the roles of e with f , we conclude that W is not contained in W' . The wedges W and W' must overlap because otherwise e and f could not intersect.

(ii) If $y = u$ or $y = w$, we can swap the roles of e and f , thus satisfying Condition 1.

(iii) We are left with the case that all six points u, v, w, x, y, z are distinct. Let $g = uv$ or $g = vw$ be the edge that is intersected by e . By Lemma 2, the triangle xyz is empty; thus, g must intersect a second edge xy or yz of this triangle, in addition to $e = xz$. This is a contradiction, since the edges g , xy , and yz are edges of the MST.

It is easy to see that the two conditions are sufficient for a crossing: In both situations of Condition 1 and Condition 2 (Figure 2), if there were no crossing between e and f , an endpoint of one edge would be contained in the triangle spanned by the other edge and its witness, contradicting Lemma 2. \square

2.1 Constructing two almost disjoint layers

With the above observations we can proceed to show a construction that almost works for two layers: a single edge will be part of both layers, while all other

edges occur in at most one tree. To this end we consider the minimum spanning tree $\text{MST}(S)$ to be rooted at an arbitrary leaf r . For any $v \in S$, we define its *level* $\ell(v)$ as its distance to r in $\text{MST}(S)$. That is, $\ell(v) = 0$ if and only if $v = r$. Likewise, $\ell(v) = 1$ if and only if v is adjacent to r etc.

For any $v \in S \setminus \{r\}$, we define its *parent* $p(v)$ as the first vertex traversed in the unique shortest path from v to r in $\text{MST}(S)$. Similarly, we define its *grandparent* $g(v)$ as $g(v) = p(p(v))$ if $\ell(v) \geq 2$ and as $g(v) = r$ otherwise (i.e., $g(v) = p(v) = r$ if $\ell(v) = 1$). Each vertex q for which $v = p(q)$ is called a *child* of v .




Figure 3: Example of Construction 1. $\text{MST}(S)$ is drawn in solid black, and the red and blue graphs in dashed and dot dashed, respectively. Note that the only common edge between the red and blue trees is the one from the root to its only neighbor in $\text{MST}(S)$.

Construction 1. Let S be a finite set of points in the plane and let $\text{MST}(S)$ be rooted at an arbitrary leaf $r \in S$. We construct two graphs $R = G(S, E_R)$ and $B = G(S, E_B)$ as follows: For any vertex $v_o \in S$ whose level is odd, we add the edge $v_o p(v_o)$ to E_R and the edge $v_o g(v_o)$ to E_B . For any vertex $v_e \in S \setminus \{r\}$ whose level is even, we add the edge $v_e g(v_e)$ to E_R and the edge $v_e p(v_e)$ to E_B .

For simplicity we say that the edges of $R = G(S, E_R)$ are colored red and the edges of $B = G(S, E_B)$ are colored blue. An edge in both graphs is called red-blue. See Figure 3 for a sketch of the construction.

Theorem 1. Let $\text{MST}(S)$ be rooted at r . The two graphs $R = G(S, E_R)$ and $B = G(S, E_B)$ from Construction 1 fulfill the following properties:

1. Both R and B are plane spanning trees.
2. $\max\{\text{BE}(R), \text{BE}(B)\} \leq 2\text{BE}(\text{MST}(S))$.
3. $E_R \cap E_B = \{rs\}$, with $r = p(s)$, that is, $|E_R \cap E_B| = 1$.

Proof. Recall from Construction 1 that r is a leaf of $\text{MST}(S)$. Hence r has a unique neighbor s in $\text{MST}(S)$ and we have $r = p(s) = g(s)$ and $\ell(s) = 1$. Let $S_o \subset S \setminus \{s\}$ be all $v_o \in S$ whose level $\ell(v_o)$ is odd. Likewise, let $S_e \subset S \setminus \{r\}$ be all $v_e \in S$ whose level $\ell(v_e)$ is even. By construction, E_R contains all the edges from odd-leveled nodes to their parents, those from even-leveled nodes to their grandparents and rs . More formally,

$$E_R = \bigcup_{v_o \in S_o} \{v_o p(v_o)\} \cup \bigcup_{v_e \in S_e} \{v_e g(v_e)\} \cup \{rs\}.$$

Similarly, E_B contains edges from odd-leveled nodes to their grandparents, those from even-leveled nodes to their parents and rs , that is

$$E_B = \bigcup_{v_o \in S_o} \{v_o g(v_o)\} \cup \bigcup_{v_e \in S_e} \{v_e p(v_e)\} \cup \{rs\}.$$

Thus, the edge rs is the only shared edge between the sets E_R and E_B , as stated in Property 3 (we call this unique edge the *double-edge*).

As E_R and E_B are subsets of the edge set of $\text{MST}^2(S)$, the vertices of every edge in E_R and E_B have link distance at most 2 in $\text{MST}(S)$, and the bound on $\max\{\text{BE}(R), \text{BE}(B)\}$ stated in Property 2 follows.

Further, both R and B are spanning trees, that is, connected and cycle-free graphs, as each vertex except r is connected either to its parent or grandparent in $\text{MST}(S)$. To prove Property 1, it remains to show that both trees are plane.

Assume for the sake of contradiction that an edge f is crossed by an edge e of the same color. Recall that all edges of E_R and E_B are edges of $\text{MST}^2(S)$ whose endpoints have different levels. By Lemma 4, at least one of $\{e, f\}$ has to be a long edge. Without loss of generality let $f = uw$ be a long edge and let v be the witness of f with $\ell(u) = \ell(v) - 1 = \ell(w) - 2$. First note that v cannot be an endpoint of e : since the level of v has different parity than the one of u and w , then v must be a leaf in this tree. Moreover, its only neighbor must be u and thus the edges uv and $f = uw$ cannot cross (and in particular v cannot be an endpoint of e as claimed). We further claim that v cannot be the witness of e . Any edge that has v as its witness is an edge from a child of v to u and therefore cannot cross $f = uw$. As e is neither incident to v nor has v as a witness, e crossing f contradicts Lemma 4. This proves Property 1 and concludes the proof. \square

The properties of our construction imply a first result stated in the following corollary.

Corollary 2. *For any set S of n points in the plane, there are two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $|E_R \cap E_B| = 1$ and $\max\{\text{BE}(R), \text{BE}(B)\} \leq 2\text{BE}(\text{MST}(S))$.*

Although the construction might seem to generalize to more layers by using edges of $\text{MST}^k(S)$, this is not the case. Already for $k = 3$, we can show that the trees may be non-plane. Take the example of Figure 4, where the full edges denote the minimum spanning tree. However, if a is chosen as the root of the tree, the edge ad will be crossed by the edge from e to either its parent, grandparent or great-grandparent. In this example the problem can be remedied by choosing a different root. But now consider placing a horizontally mirrored copy of this construction to the left so that a and its mirrored version are connected by an edge. Regardless of which root is chosen, in one of the two subtrees the node a or its mirrored equivalent is the root of the respective subtree. Hence, any root will create a crossing.




Figure 4: Example graph where choosing a as root creates a crossing when we generalize the above construction to three trees.




Figure 5: Illustration of the various definitions used in Section 2.2. Grey triangles illustrate further subtrees, one shows interior vertices purely for illustrative purposes.

2.2 Avoiding the double edge

Construction 1 is almost valid in the sense that only one edge was shared between both trees. In the following we modify this construction to avoid the shared edge.

Let $N^- \subset (S \setminus \{r\})$ be the set of neighbors $v^- \in (S \setminus \{r\})$ of s in $\text{MST}(S)$ such that the ordered triangle rsv^- is oriented clockwise. Let $N^+ \subset (S \setminus \{r\})$ be the set of neighbors $v^+ \in (S \setminus \{r\})$ of s in $\text{MST}(S)$ such that the ordered triangle rsv^+ is oriented counterclockwise. Let T^- be the subtree of $\text{MST}(S)$ that is connected to s via the vertices in N^- and let T^+ be the subtree of $\text{MST}(S)$ that is connected to s via the vertices in N^+ . Let $S^- \subset S$ consist of r and the set of vertices from T^- and let $S^+ \subset S$ consist of r and the set of vertices from T^+ . Observe that $S^- \cap S^+ = \{r, s\}$ (see Figure 5). Let E_R and E_B be sets of red and blue edges as defined in the Construction 1. Then let $E_R^- \subset E_R$ ($E_B^- \subset E_B$) be the subset of edges that have at least one endpoint in $S^- \setminus \{r, s\}$

and let $E_R^+ \subset E_R$ ($E_B^+ \subset E_B$) be the subset of edges that have at least one endpoint in $S^+ \setminus \{r, s\}$. Note that by this definition $E_R = E_R^- \cup E_R^+ \cup \{rs\}$ and $E_B = E_B^- \cup E_B^+ \cup \{rs\}$. With this we define the subgraphs $R^- = G(S^-, E_R^-)$, $R^+ = G(S^+, E_R^+)$, $B^- = G(S^-, E_B^-)$, and $B^+ = G(S^+, E_B^+)$ and prove a useful non-crossing property between these graphs.

Lemma 5. *For any set S of n points in the plane, let $R = G(S, E_R)$ and $B = G(S, E_B)$ be the graphs from Construction 1. Then no edge in E_R^- crosses an edge in E_B^+ and no edge in E_R^+ crosses an edge in E_B^- .*

Proof. Consider any edge $e \in E_R^-$ that is not incident to r . By Lemma 4, such an edge e can be crossed only by an edge incident to at least one vertex of $S^- \setminus \{r, s\}$. Hence, e does not cross any edge of E_B^+ .

Assume for the sake of contradiction that there is an edge $f \in E_B^+$ incident to r that crosses an edge $e \in E_R^-$. By construction, $e = rz$ is a long edge of $\text{MST}^2(S)$ with witness s and wedge W . By Lemma 4, f has to be incident to s , since s cannot be the witness of any blue edges by construction. If f is a short edge, then f is not in W by our definition of S^- and S^+ , which contradicts Lemma 4. Hence, let $f = sc$ be a long edge of $\text{MST}^2(S)$ with witness b . Following Lemma 4, the witness b must be s , which contradicts the fact that s cannot be a witness of any blue edge. This concludes the proof that no edge in E_R^- is crossed by an edge in E_B^+ . Symmetric arguments prove that no edge in E_R^+ is crossed by an edge in E_B^- . \square

With this observation we can now prove that the two spanning trees (rooted at an arbitrary leaf r) from Construction 1 actually exist in 4 different color combination variants.

Lemma 6. *Let S be a set of n points in the plane. Let $R = G(S, E_R)$ and $B = G(S, E_B)$ be the graphs from Construction 1 and let $R^- = G(S^-, E_R^-)$, $R^+ = G(S^+, E_R^+)$, $B^- = G(S^-, E_B^-)$, and $B^+ = G(S^+, E_B^+)$ be subgraphs as defined above. Then R and B can be recolored to any of the four versions below, where each version fulfills the properties of Theorem 1.*

- (1) $R = G(S, E_R)$ and $B = G(S, E_B)$ (the original coloring)
- (2) $R = G(S, E_B)$ and $B = G(S, E_R)$ (the inverted coloring)
- (3) $R = G(S, E_B^- \cup E_R^+ \cup \{rs\})$ and $B = G(S, E_R^- \cup E_B^+ \cup \{rs\})$ (the $-$ side inverted coloring)
- (4) $R = G(S, E_R^- \cup E_B^+ \cup \{rs\})$ and $B = G(S, E_B^- \cup E_R^+ \cup \{rs\})$ (the $+$ side inverted coloring)

Proof. The statement is trivially true for recolorings (1) and (2). It is easy to observe that this corresponds to a simple recoloring. Hence, Properties 2 and 3 of Theorem 1 are also obviously true. By Lemma 5, both R and B are plane for the recolorings (3) and (4) and thus fulfill Property 1 of Theorem 1 as well. \square

With these tools we can now show how to construct two disjoint spanning trees. For technical reasons we use two different constructions based on the existence of a vertex v in the minimum spanning tree where no two consecutive adjacent edges span an angle larger than π .

Theorem 3. *Let S be a set of n points in the plane, and v a vertex of S . Assume that there is a minimum spanning tree $\text{MST}(S)$ such that the angle between any two consecutive adjacent edges of v in $\text{MST}(S)$ is smaller than π . Then there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $E_R \cap E_B = \emptyset$ and $\max\{\text{BE}(R), \text{BE}(B)\} \leq 2\text{BE}(\text{MST}(S))$.*

Proof. We build the two spanning trees by using the vertex v to decompose the minimum spanning tree into trees where v is a leaf. For each of these subtrees we apply Construction 1 and possibly recolor them in one of the variants from Lemma 6.

Let $S_v = \{v_1, \dots, v_k\}$ be the set of vertices incident to v in $\text{MST}(S)$, labeled in counterclockwise order as they appear around v . Observe that $k \geq 3$ is necessary to fulfill the angle condition from the theorem. By Lemma 2, the convex hull of S_v contains no points of S except v . We start by constructing two plane spanning trees of $S_v \cup \{v\}$. The red spanning tree $R_v = G(S_v \cup \{v\}, E_{vR})$ contains all edges incident to v except vv_1 , plus the edge v_1v_2 , which lies on the convex hull boundary of S_v . The blue spanning tree $B_v = G(S_v \cup \{v\}, E_{vB})$ contains all edges on the convex hull boundary of S_v except v_1v_2 , plus the edge vv_1 . Observe that R_v and B_v are plane spanning trees, $E_{vR} \cap E_{vB} = \emptyset$, and $\max\{\text{BE}(R_v), \text{BE}(B_v)\} \leq 2\text{BE}(\text{MST}(S_v \cup \{v\}))$.

Next consider a vertex v_i of S_v and let M_i be the maximal subtree of $\text{MST}(S)$ that is connected to v by v_i . Let $S_i \subset S$ be the vertex set of M_i . Note that $M_i = \text{MST}(S_i)$ and that v is a leaf in M_i . Thus, we can use Construction 1 to get spanning trees $R_i = G(S_i, E_{iR})$ and $B_i = G(S_i, E_{iB})$, all rooted at v . The graphs R_i and B_i fulfill the three properties of Theorem 1 and the only edge shared between R_i and B_i is vv_i .

Considering Lemma 4 and the fact that for $i \neq j$ the edges of $E_{iR} \cup E_{iB}$ have no point, except for the root v , in common with $E_{jR} \cup E_{jB}$, it is easy to see that no edge of $E_{iR} \cup E_{iB}$ crosses any edge of $E_{jR} \cup E_{jB}$. In order to join the graphs to two plane spanning trees on S , we adapt them slightly, while keeping the properties of Theorem 1. We first state how we combine the edge sets of the different plane spanning trees to get $R = G(S, E_R)$ and $B = G(S, E_B)$ and then reason why the claim in the theorem is true for this construction.

$$E_R = E_{vR} \cup (E_{3R} \setminus \{vv_3\}) \cup \dots \cup (E_{kR} \setminus \{vv_k\}) \cup (E_{1R}^- \cup E_{1B}^+) \cup (E_{2B}^- \cup E_{2R}^+)$$

$$E_B = E_{vB} \cup (E_{3B} \setminus \{vv_3\}) \cup \dots \cup (E_{kB} \setminus \{vv_k\}) \cup (E_{1B}^- \cup E_{1R}^+) \cup (E_{2R}^- \cup E_{2B}^+)$$

First we add the construction for $S_v \cup \{v\}$ to both edge sets. This is the base to which all other trees will be attached. Then the graphs from the subtrees M_i for $1 \leq i \leq k$ are added to this base. The edges vv_i are already used in R_v or B_v , so we add the edges vv_i neither from E_{iR} nor E_{iB} . As both v and v_i are connected to both colors (both spanning trees), the construction stays

connected. As we did not add any additional edges the construction obviously stays cycle-free and the edge length bound is maintained.

It remains to argue the planarity of the resulting graphs. By Lemma 4, edges of E_{iR} or E_{iB} that cross any edge of E_{vR} and E_{vB} have to be incident to v . By Lemma 3, only the edges $e_i^- = v_{i-1}v_i$ and $e_i^+ = v_iv_{i+1}$ (indices modulo k) are crossed by edges of $E_{iR} \setminus \{vv_i\}$ and $E_{iB} \setminus \{vv_i\}$.

Using the original coloring (see Lemma 6) for R_i and B_i only red edges (edges of $E_{iR} \setminus \{vv_i\}$) cross e_i^- and e_i^+ . For any $3 \leq i \leq k$, e_i^- and e_i^+ are blue, i.e., $e_i^-, e_i^+ \in E_{vB}$.

For $i = 1$, the edge e_i^+ is red. In this case, we use the $+$ side inverted coloring (see Lemma 6) for R_i and B_i (and exclude the edge vv_i): $E_{iR} = E_{iR}^- \cup E_{iB}^+$ and $E_{iB} = E_{iB}^- \cup E_{iR}^+$. Using this coloring, all shown properties remain valid (see Lemma 6), all edges from R_i and B_i that cross the blue edge e_i^- remain red, and all edges from R_i and B_i that cross the red edge e_i^+ are now blue.

In a similar manner we fix the case of $i = 2$, where the edge e_i^- is red. We use the $-$ side inverted coloring (see Lemma 6) for R_i and B_i (and exclude the edge vv_i): $E_{iR} = E_{iB}^- \cup E_{iR}^+$ and $E_{iB} = E_{iR}^- \cup E_{iB}^+$. Again, all shown properties remain valid (see Lemma 6). All edges from R_i and B_i that cross the red edge e_i^- are now blue, and all edges from R_i and B_i that cross the blue edge e_i^+ remain red.

Hence, with this slightly adapted construction (and coloring), $R = G(S, E_R)$ and $B = G(S, E_B)$ are plane spanning trees that solely use edges of $\text{MST}^2(S)$ and have no edge in common. \square

In the remaining case, for every vertex in an $\text{MST}(S)$ there are two consecutive adjacent edges that span an angle larger than π . In such an $\text{MST}(S)$, every vertex has degree at most three, since the angle between adjacent edges is at least $\pi/3$.

Theorem 4. *Consider a set S of $n \geq 4$ points in the plane for which every vertex in the minimum spanning tree $\text{MST}(S)$ has two consecutive adjacent edges spanning an angle larger than π . Then there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $E_R \cap E_B = \emptyset$ and $\max\{\text{BE}(R), \text{BE}(B)\} \leq 3\text{BE}(\text{MST}(S))$. In addition, at most one edge of $E_R \cup E_B$ is longer than $2\text{BE}(\text{MST}(S))$.*

Proof. As before in Theorem 3 we will use our construction scheme for trees rooted at a leaf for the majority of the points and use a small local construction that avoids double edges. In this case, instead of removing a single vertex v to decompose the tree we use a set of four vertices as follows. We start at a leaf of $\text{MST}(S)$ to generate a connected graph P with four vertices that is a subgraph of $\text{MST}(S)$. Then we show how to construct R and B for S for the different cases of P in combination with the remainder of $\text{MST}(S)$.

The construction of $P = \{v_3, v_2, v_1, v_0\}$: Let v_3 be a leaf of $\text{MST}(S)$. For the construction of P , we root $\text{MST}(S)$ at v_3 . We call the number of vertices in the (sub)tree of which that vertex is a root of (including itself) the *weight* of

a vertex. Hence, the weight of v_3 is n . Further, the angle between two successive incident edges at a vertex of $\text{MST}(S)$ that is larger than π is called the *big angle*.




Figure 6: Case 1 for P and the connections to the rest of $\text{MST}(S)$. The gray triangles indicate possible subtrees of $\text{MST}(S)$ and where and how they might be connected. Dotted edges are from $\text{MST}^2(P)$. Note that the subtree with root v_0 can be on either side of the supporting line of v_1v_0 and even on both sides as indicated in the figure.

Let v_2 be the unique child of v_3 in $\text{MST}(S)$ (i.e., $v_3 = p(v_2)$). To define v_1 and v_0 we use a case distinction. Consider the set C of the children of v_2 that are not spanning the big angle with v_3 at v_2 . (Note that v_3 may or may not be spanning the big angle at v_2 and that C contains 0, 1, or 2 vertices).

1. If v_2 has only a single child (i.e., C is empty), or if C contains a vertex that is not a leaf in $\text{MST}(S)$, we choose it (or one of them) as v_1 . We assume w.l.o.g. that v_1 is the successor of v_3 in clockwise order around v_2 . Further, we choose v_0 as a child of v_1 such that v_2 and v_0 are consecutive around v_1 and do not span the big angle at v_1 . If v_1 has two children, and this is true for both, then we choose v_0 such that it is the successor of v_2 in counterclockwise order around v_1 . See Figure 6 (a)–(f) for the six different variations of this case, i.e., all possible distributions of the positions of the big angles. Explicit definitions for these cases can be found below. W.l.o.g., we require the subtree at v_2 to be nonempty in cases (d)–(f) and the subtree of v_1 to be nonempty in cases (c) and (f). For all cases we assume without loss of generality that the angle $v_3v_2v_1$ is clockwise less than π .

- (1a) The edge v_3v_2 is **adjacent** to the big angle at v_2 and the angle $v_2v_1v_0$ is clockwise **less** than π .
- (1b) The edge v_3v_2 is **adjacent** to the big angle at v_2 , the angle $v_2v_1v_0$ is clockwise **greater** than π , and the edge v_2v_1 is **adjacent** to the big angle at v_1 .

- (1c) The edge v_3v_2 is **adjacent** to the big angle at v_2 , the angle $v_2v_1v_0$ is clockwise **greater** than π , and the edge v_2v_1 is **not adjacent** to the big angle at v_1 .
 - (1d) The edge v_3v_2 is **not adjacent** to the big angle at v_2 and the angle $v_2v_1v_0$ is clockwise **less** than π .
 - (1e) The edge v_3v_2 is **not adjacent** to the big angle at v_2 , the angle $v_2v_1v_0$ is clockwise **greater** than π , and the edge v_2v_1 is **adjacent** to the big angle at v_1 .
 - (1f) The edge v_3v_2 is **not adjacent** to the big angle at v_2 , the angle $v_2v_1v_0$ is clockwise **greater** than π , and the edge v_2v_1 is **not adjacent** to the big angle at v_1 .
2. C contains at least one vertex and all vertices in C are leaves in $MST(S)$. Note that this implies that v_2 has degree exactly three in $MST(S)$, as the degree of any vertex with a big angle in a minimum spanning tree is at most three and thus C would be empty if v_2 had degree at most two. We choose a vertex of C as v_1 (assuming w.l.o.g. that v_1 is the successor of v_3 in clockwise order around v_2), and choose the other child of v_2 as v_0 . Note that if $n \geq 5$ then v_0 cannot be a leaf in $MST(S)$ and hence v_0 spans a big angle with v_3 at v_2 . Therefore, taking the location of the big angle at v_2 into account, there are two possibilities for the counterclockwise order of incident edges around v_2 . See Figure 7 (a) and (b).
- (2a) The angle $v_1v_2v_0$ is clockwise **less** than π (the edge v_3v_2 is **adjacent** to the big angle at v_2 .)
 - (2b) The angle $v_1v_2v_0$ is clockwise **greater** than π (the edge v_3v_2 is **not adjacent** to the big angle at v_2). Here, both v_1 and v_0 must be leaves implying that $n = 4$ and we do not have any subtrees.




Figure 7: Case 2 for the selection of P and the connections to the rest of $MST(S)$. Vertices drawn with squares are leaves of $MST(S)$. The gray triangles indicate possible subtrees of $MST(S)$ and where they might be connected. Dotted edges are from $MST^2(P)$. Note that the subtree with root v_0 can be on either side of the supporting line of v_1v_0 and even on both sides as indicated in the figure. Further note that case (b) can appear only if S contains exactly 4 vertices.

The construction of R and B : First we show how to construct the trees $R_P = G(\{v_3, v_2, v_1, v_0\}, E_{PR})$ and $B_P = G(\{v_3, v_2, v_1, v_0\}, E_{PB})$. The

vertices of P can either be in convex position or form a triangle with one interior point, with v_1 interior for the cases shown in Figure 6 (b) and (e), and v_2 interior for the cases shown in Figure 6 (e) and (f); there are no other non-convex versions: otherwise either the path v_3, v_2, v_1, v_0 could not be in $\text{MST}(S)$, or one of the vertices of P could not be incident to a big angle (recall Lemma 2). As in any non-convex case the complete graph on $\{v_3, v_2, v_1, v_0\}$ is crossing-free, any construction of R_P and B_P for the convex cases is also valid for the non-convex cases. Let us thus assume the points of P to be in convex position. For Cases 1a and 1d the vertices $\{v_0, v_1, v_2, v_3\}$ must appear in this order on their convex hull and we set $R_P = \{v_2v_3, v_3v_1, v_1v_0\}$ and $B_P = \{v_3v_0, v_0v_2, v_2v_1\}$, as illustrated in Figure 8 (a). For Cases 1b, 1c, 1e, and 1f the vertices appear in the order $\{v_3, v_1, v_0, v_2\}$ on their convex hull. (The order $\{v_0, v_1, v_2, v_3\}$ violates the fact that the clockwise angle $v_0v_1v_2$ is not large and the order $\{v_0, v_3, v_1, v_2\}$ has edges v_0v_1 and v_2v_3 crossing, which are both edges of $\text{MST}(S)$; no other orderings exist when accounting for symmetry.) For these cases we set $R_P = \{v_3v_2, v_2v_1, v_1v_0\}$ and $B_P = \{v_1v_3, v_3v_0, v_0v_2\}$, as in Figure 8 (b). All edges except the edge v_3v_0 (in E_{PB}) are from $\text{MST}^2(S)$ and have endpoints with different levels in $\text{MST}(S)$ rooted at v_3 . In contrast, v_3v_0 is an edge of $\text{MST}^3(S)$, which could be crossed by other edges of the construction. We will later discuss how to handle this.

For Case 2 the placement must be convex as v_0, v_1 and v_3 are adjacent to v_2 and the clockwise $v_0v_2v_3$ and $v_1v_2v_0$ are convex for Cases 2a and 2b, respectively. Since the ordering around v_2 is fixed (modulo symmetry) by the case definition, the vertices appear in the order $\{v_2, v_3, v_1, v_0\}$ for Case 2a and $\{v_2, v_0, v_3, v_1\}$ for Case 2b. For Case 2a we set $R_P = \{v_2v_3, v_3v_0, v_0v_1\}$ and $B_P = \{v_3v_1, v_1v_2, v_2v_0\}$ and for Case 2b we set $R_P = \{v_2v_0, v_0v_1, v_1v_3\}$ and $B_P = \{v_1v_2, v_2v_3, v_3v_0\}$ as illustrated in Figure 8 (c) and (d). For both cases, all edges are from $\text{MST}^2(S)$.




Figure 8: The different colorings for P : (a) for the cases from Figure 6 (a) and (d), (b) for the cases from Figure 6 (b), (c), (e), and (f), (c) for the case from Figure 7 (a), and (d) for the case from Figure 7 (b). The bold edge in (a) and (b) is an edge of $\text{MST}^3(S)$, that is, an edge with link distance 3 in $\text{MST}(S)$.

With R_P and B_P as a base, we now create red and blue trees for all remaining subtrees of $\text{MST}(S)$ and “attach” them to the base. For Case 1 we define three possible subtrees. Let $T'_0 = G(S'_0, E'_0)$ be the subtree (i.e., connected component) of $\text{MST}(S)$ that contains v_0 when removing v_1 (and its incident edges) from $\text{MST}(S)$. Likewise, let $T'_1 = G(S'_1, E'_1)$ be the subtree of $\text{MST}(S)$ that contains v_1 when removing v_0 and v_2 from $\text{MST}(S)$, and let $T'_2 = G(S'_2, E'_2)$ be the subtree of $\text{MST}(S)$ that contains v_2 when removing v_1 and v_3 from $\text{MST}(S)$. For Case 2(a)

there is one possible subtree $T'_0 = G(S'_0, E'_0)$, which is the subtree of $\text{MST}(S)$ that contains v_0 when removing v_2 from $\text{MST}(S)$. (Case 2(b) appears only if $n = 4$ and hence the construction is already completed.) The subtrees T'_0 , T'_1 , and T'_2 are shown as (pairs of) gray triangles in Figure 6 and 7. To connect these subtrees to the bases R_P and B_P , we create corresponding trees T_0 , T_1 , and T_2 depending on the different shown cases, then apply Construction 1 to them, and possibly recolor them using Lemma 6.

We first consider the different subtrees for Case 1. In essence, for each tree we pick a neighbor from $\{v_3, v_2, v_1, v_0\}$ to add to T'_0, T'_1, T'_2 , which we then use as root for Construction 1. When there is a choice we pick a root that is adjacent to the outgoing edge from v_i into the subtree T'_i as defined more precisely below.

T_0 : For all cases, we consider the subtree $T_0 = G(S_0, E_0)$ of $\text{MST}(S)$, with $S_0 = S'_0 \cup \{v_1\}$, $E_0 = E'_0 \cup \{v_1 v_0\}$. We root T_0 at $r = v_1$ (observe, v_1 is a leaf in T_0 with unique child $s = v_0$) and apply Construction 1 to get $R_0 = G(S_0, E_{0R})$ and $B_0 = G(S_0, E_{0B})$, with the double-edge rs removed from both E_{0R} and E_{0B} .

T_1 : For the cases depicted in Figure 6 (a), (b), (d), and (e), we define $T_1 = G(S_1, E_1)$ of $\text{MST}(S)$, such that $S_1 = S'_1 \cup \{v_0\}$, $E_1 = E'_1 \cup \{v_1 v_0\}$. We root T_1 at $r = v_0$ (observe, v_0 is a leaf in T_1 with unique child v_1) and apply Construction 1 to get $R_1 = G(S_1, E_{1R})$ and $B_1 = G(S_1, E_{1B})$, with the double-edge rs removed from both E_{1R} and E_{1B} .

In the cases shown in Figure 6 (c) and (f), we define $T_1 = G(S_1, E_1)$ of $\text{MST}(S)$, such that $S_1 = S'_1 \cup \{v_2\}$, $E_1 = E'_1 \cup \{v_1 v_2\}$. We root T_1 at $r = v_2$ (observe, v_2 is a leaf in T_1 with unique child v_1) and apply Construction 1 to get $R_1 = G(S_1, E_{1R})$ and $B_1 = G(S_1, E_{1B})$, with the double-edge rs removed from both E_{1R} and E_{1B} .

T_2 : For the cases depicted in Figure 6 (a)–(c), let $T_2 = G(S_2, E_2)$ be a subtree of $\text{MST}(S)$, such that $S_2 = S'_2 \cup \{v_1\}$, $E_2 = E'_2 \cup \{v_1 v_2\}$. We root T_2 at $r = v_1$ (observe, v_1 is a leaf in T_2 with unique child v_2) and apply Construction 1 to get $R_2 = G(S_2, E_{2R})$ and $B_2 = G(S_2, E_{2B})$, with the double-edge rs removed from both E_{2R} and E_{2B} .

For the cases depicted in Figure 6 (d)–(f), $T_2 = G(S_2, E_2)$ is the subtree of $\text{MST}(S)$, such that $S_2 = S'_2 \cup \{v_3\}$, $E_2 = E'_2 \cup \{v_3 v_2\}$. We root T_2 at $r = v_3$ (observe, v_3 is a leaf in T_2 with unique child v_2) and apply Construction 1 to get $R_2 = G(S_2, E_{2R})$ and $B_2 = G(S_2, E_{2B})$, with the double-edge rs removed from both E_{2R} and E_{2B} .

It is easy to see that the edge sets E_{PR} , E_{0R} , E_{1R} , E_{2R} , E_{PB} , E_{0B} , E_{1B} , and E_{2B} are all individually edge disjoint. In the following, we describe how these edge sets are combined to form the two plane spanning trees R and B in the different cases; see Figure 9 for the convex versions and Figure 10 for the non-convex versions of the corresponding cases illustrated in Figure 6. For the non-convex cases only points v_1 and v_2 can be in the interior as v_3 or v_0 being in the interior would violate Lemma 2. Furthermore, in some of the cases (a)–(f), further restrictions apply as listed below.

- (a) Neither v_1 nor v_2 can be the middle point as both clockwise angles $v_3 v_2 v_1$ and $v_2 v_1 v_0$ have angle less than π .

- (b) Only v_1 can be in the center (the subtree of v_2 is nonempty and hence v_2 would not be incident to a big angle).
- (c) Neither v_1 nor v_2 can be in the center (both subtrees are nonempty).
- (d) Similar to (a).
- (e) Both v_1 and v_2 may be the middle point.
- (f) Only v_2 can be in the center (the subtree of v_1 is non-empty).




Figure 9: The different plane spanning trees R and B for case 1 when P is in convex position. The bold blue edges v_3v_0 are edges of $\text{MST}^3(S)$, i.e., edges with link distance 3 in $\text{MST}(S)$, that might still be crossed. Note that v_3v_0 can only be crossed in cases (a) and (b), as in the other cases this edge is surrounded by “uncrossable” $\text{MST}^2(S)$ -edges.




Figure 10: The different plane spanning trees R and B for case 1 when P is not in convex position, with v_1 or v_2 in the interior. The case numberings are the same as the ones in Figure 6.

First we add the red and blue trees for T_0 . By construction, only edges of E_{0R} connect to v_1 (only crossing edges of E_{PB}) and the edges of E_{0B} do not cross any edge outside T_0 . For the cases (a), (b), (d), and (e) we use the inverted coloring (see Lemma 6) for the two trees of T_1 . For the remaining cases (c)

and (f) we use the original coloring (see Lemma 6) for the two trees of T_1 . For adding the red and blue trees for T_2 we use the original coloring for the cases (a-c), and the inverted coloring for the cases (d-f). For the sake of simplicity, we exchange the names of the according sets E_{iR} and E_{iB} whenever we use the inverted coloring for a T_i , $i \in \{1, 2\}$.

Since the edge sets E_{0R}, E_{1R} , and E_{2R} form spanning trees connecting the nodes of T_0, T_1 , and T_2 and since E_{PR} connects their roots and v_3 with a spanning tree, it follows that $R = G(S, E_{PR} \cup E_{0R} \cup E_{1R} \cup E_{2R})$ is a spanning tree for S . The same argument applies to show that $B = G(S, E_{PB} \cup E_{0B} \cup E_{1B} \cup E_{2B})$ is a spanning tree. All edges in $E_{0R}, E_{1R}, E_{2R}, E_{0B}, E_{1B}$, and E_{2B} are from $\text{MST}^2(S)$ and all edges of E_{PR} and E_{PB} are from $\text{MST}^2(S) \cup \{v_3v_0\}$, it follows that $\max\{\text{BE}(R), \text{BE}(B)\} \leq 3\text{BE}(\text{MST}(S))$, with only the edge v_3v_0 , which may occur in E_{PB} being possibly larger than $2\text{BE}(\text{MST}(S))$. What remains is to show that both R and B are non-crossing.

The $\text{MST}^3(S)$ -edge v_3v_0 from E_{PB} is also the only edge that could cause a crossing, see Lemma 4 and Theorem 1. Hence, R is a plane spanning tree. If v_3v_0 is not crossed by any other edge of E_B then also B is a plane spanning tree and we are done. Otherwise, note first that by Lemma 2, the triangles $v_3v_2v_1$ and $v_2v_1v_0$ cannot contain any points of S . Then observe that for the cases in Figure 9 (b), (c), (e) and (f) any edge crossing v_3v_0 that does not have v_0, v_1, v_2 or v_3 as an endpoint must cross an MST -edge between v_0, v_1, v_2 and v_3 . This implies that any MST^2 -edge that crosses v_3v_0 must have v_0, v_1, v_2 or v_3 as its witness by Lemma 4. By construction, v_0, v_1, v_2 and v_3 are not a witness to any blue edge between vertices in $S \setminus \{v_0, v_1, v_2, v_3\}$. The edges in E_{PB} are incident to v_0 or v_3 so they also cannot cross v_3v_0 .

For cases from Figure 9 (a) and (d), as well as all cases from Figure 10, observe first that v_0v_1, v_1v_2 , and v_2v_3 cannot be crossed, again due to Lemma 4 and the fact that v_0, v_1, v_2 and v_3 cannot be a witness to any long edge in the grey subtrees by construction. So any edge that crosses v_3v_0 must have an endpoint in the interior of the convex hull of P or connect directly to v_1 or v_2 . The latter however cannot happen: The only points connecting with blue edges to v_1 or v_2 are direct neighbors of these vertices, which reside in the large-angled wedge $v_1v_2v_3$ or $v_0v_1v_2$, respectively. (Here, the *large-angled* wedge $v_1v_2v_3$ is the wedge spanned by the a ray from v_2 to v_1 and a ray from v_2 to v_3 so that the wedge has an opening angle larger than π . The large-angled wedge $v_0v_1v_2$ is defined analogously.) Hence, if v_0v_3 is crossed by some blue edge, there must be a nonempty set $X \subset S \setminus P$ that resides in the interior of the convex hull of P . In the cases depicted in Figure 9 (a) and (d), X lies in the triangle Δ spanned by v_3, v_0 , and the intersection of v_3v_1 and v_2v_0 . In the cases depicted in Figure 10, the X lies in the triangle Δ spanned by v_3, v_0 , and the vertex of P in the interior of the convex hull of P . Further, removing the edge v_0v_3 from B splits B into two connected components that are each plane trees, where v_3 is in one and v_0 is in the other component. Now consider the convex hull of $X \cup \{v_0, v_3\}$, and the path along the boundary of that convex hull between v_0 and v_3 that contains at least one vertex of X . This path contains exactly one edge e that connects the two components of B . Due to the construction of B

and R , e can neither be part of R (as the two endpoints of e must reside in two different subtrees of v_0, v_1 or v_2) nor cross any edge of B (as the only possibly intersected segment of the convex hull boundary of X was v_0v_3). Further, the length of e must be less than $3\text{BE}(\text{MST}(P))$, as e lies inside the triangle Δ , and as all sides of Δ are bounded from above by $3\text{BE}(\text{MST}(P))$. Hence, as v_3v_0 was the only edge that could be crossed by others, the replacement of v_3v_0 by e in B results in two edge-disjoint plane spanning trees R and B with maximum edge length less than $3\text{BE}(\text{MST}(P))$.

As for Case 2 (b), S consists only of four vertices and hence Figure 8 (d) already shows all of the two trees R and B , it remains to consider the subtree for Case 2 (a).




Figure 11: The plane spanning trees R and B for Case 2.

T_0 : Consider the subtree $T_0 = G(S_0, E_0)$ of $\text{MST}(S)$, with $S_0 = S'_0 \cup \{v_2\}$, $E_0 = E'_0 \cup \{v_2v_0\}$. We root T_0 at $r = v_2$ (observe, v_2 is a leaf in T_0 with unique child $s = v_0$) and apply Construction 1 to obtain edge sets E'_{0R} and E'_{0B} , since we will add connectivity between $r = v_2$ and $s = v_0$ using E_{PR} and E_{PB} we remove the edge rs to obtain $R_0 = G(S_0, E_{0R})$ and $B_0 = G(S_0, E_{0B})$, with $E_{0R} = E'_{0R} \setminus \{rs\}$ and $E_{0B} = E'_{0B} \setminus \{rs\}$. We use the inverted coloring as defined in Lemma 6 for the two trees of T_0 , implying that the edges connecting to v_2 and crossing red edges of E_{PR} are all blue. Hence the graphs $R = G(S, E_{PR} \cup E_{0R})$ and $B = G(S, E_{PB} \cup E_{0B})$ are plane spanning trees, $E_R \cap E_B = \emptyset$, and $\max\{\text{BE}(R), \text{BE}(B)\} \leq 2\text{BE}(\text{MST}(S))$.

This concludes the proof. \square

Corollary 5. *For any set S of $n \geq 4$ points in the plane, there are two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $E_R \cap E_B = \emptyset$ and $\max\{\text{BE}(R), \text{BE}(B)\} \leq 3\text{BE}(\text{MST}(S))$.*

We now show that the above construction is worst-case optimal.

Theorem 6. *For any $n > 3$ and $k > 1$ there is a set S of n points such that for any k disjoint spanning trees, at least one has a bottleneck edge larger than $(k + 1)\text{BE}(\text{MST}(S))$.*

Proof. A counterexample simply consists of n points equally distributed on a line segment. The points can be slightly perturbed to obtain general position (similar to Figure 3). In this problem instance there are $kn - (k(k + 1)/2)$ edges whose distance is strictly less than $(k + 1)\text{BE}(\text{MST}(S)) = k + 1$. However, we need $kn - k$ edges for k disjoint trees and thus it is impossible to construct that many trees with sufficiently short edges. \square

3 Distributed Approach

The previous construction relies heavily on the minimum spanning tree of S . It is well known that this tree cannot be constructed locally, thus we are implicitly assuming that the network is constructed by a single processor that knows the location of all other vertices. In ad-hoc networks, it is often desirable that each vertex can compute its adjacencies using only local information, i.e., using only information about vertices at most a certain maximum distance away.

In the following, we provide an alternative construction. Although, for any fixed k , the length of the edges is increased by a constant factor of $12\sqrt{2}k$ (see Theorem 9 below for details), it has the benefit that it can be constructed locally and that it can be extended to compute k layers, for $k \leq n/12$. The only global property that is needed is a value β that should be at least $\text{BE}(\text{MST}(S))$. We also note that these plane disjoint graphs are not necessarily trees, as large cycles cannot be detected locally.

Before we describe our approach, we report the result of Biniaz and García [3] that every point set of at least $3k$ points contains k layers. Since the details of this construction are important for our construction, we add a proof sketch.

Theorem 7 ([3]). *Every finite point set that consists of at least $3k$ points contains k layers.*

Proof. First, recall that for every set of n points, there is a *center point* c such that every line through c splits the point set into two parts that each contain at least $\lfloor n/3 \rfloor$ points, see e.g. Chapter 1 in [8] (note that c need not be one of the initial n points). For ease of explanation, we assume that every line through c contains at most one point. Number the points v_0, v_1, \dots, v_{n-1} in clockwise circular order around c . We split the plane into three angular regions by the three rays originating from c and passing through $v_0, v_{\lfloor \frac{n}{3} \rfloor}$, and $v_{\lfloor \frac{2n}{3} \rfloor}$; see Figure 12. Since every line through the center contains at least $n/3$ points on each side, the smaller angle of the two rays defining a region is at most π and thus the three angular regions are convex. We declare v_0 to be the representative of the angular region between the rays through v_0 and $v_{\lfloor \frac{n}{3} \rfloor}$ and connect the vertices $v_1, \dots, v_{\lfloor \frac{n}{3} \rfloor}$ in this region to v_0 . Similarly, we assign $v_{\lfloor \frac{n}{3} \rfloor}$ to be the representative of angle between the rays center through $v_{\lfloor \frac{n}{3} \rfloor}$ and $v_{\lfloor \frac{2n}{3} \rfloor}$ and connect vertices $v_{\lfloor \frac{n}{3} \rfloor+1}, \dots, v_{\lfloor \frac{2n}{3} \rfloor}$ to $v_{\lfloor \frac{n}{3} \rfloor}$. Finally, we connect vertices $v_{\lfloor \frac{2n}{3} \rfloor+1}, \dots, v_{n-1}$ to $v_{\lfloor \frac{2n}{3} \rfloor}$. This results in a non-crossing spanning tree.

For the second tree, we rotate the construction and we use $v_1, v_{\lfloor \frac{n}{3} \rfloor+1}$, and $v_{\lfloor \frac{2n}{3} \rfloor+1}$ to define the three regions, and so on. \square

While this construction provides a simple method of constructing the k layers, it does not give any guarantee on the length of the longest edge in this construction. To give such a guarantee, we combine it with a bucketing approach: we partition the point set using a grid (whose size will depend on k and β), solve the problem in each box with sufficiently many points independently, and then combine the subproblems to obtain a global solution (see Figure 13).




Figure 12: Extracting one layer: (a) The three sectors defined by v_0 , $v_{\lfloor \frac{2n}{3} \rfloor}$, and $v_{\lfloor \frac{n}{3} \rfloor}$. (b) Connecting the points to the representative of their sector. The red edges connect the representatives.

We place a grid with cells of height and width $6k\beta$ and classify the points according to which grid cell contains them (if a point lies exactly on a separating line, we assign it an arbitrary adjacent cell). We say that a grid cell is a *dense box* if it contains at least $3k$ points of S . Similarly, it is a *sparse box* if it contains points of S but is not dense. Two boxes are adjacent if they share some boundary or vertex. Hence, each box has 8 neighbors. This is also referred to as the 8-neighbor topology. We observe that dense and sparse boxes satisfy the following properties.

Lemma 7. *Given two non-adjacent boxes B and B' , the points in B and B' cannot be connected by edges of length at most β using only points from sparse boxes.*

Proof. Suppose the contrary and let B and B' be two boxes such that there is a path that uses edges of length at most β between a point in B and a point in B' visiting only points in sparse boxes. This path crosses the sides of a certain number of boxes in a given order; let σ be the sequence of these sides, after repeatedly removing adjacent duplicates. Observe first that horizontal and vertical sides alternate in σ , as otherwise the path would traverse the cell width of $6k\beta$ using at most $3k - 1$ points connected by edges of length at most β . Since B and B' are non-adjacent, w.l.o.g., there is a vertical side s that has two adjacent horizontal sides in σ with different y -coordinates. Hence, between the two horizontal sides, the corresponding part of the path has length at least $6k\beta$, and may use only the points in the two boxes adjacent to s . But since any sparse box contains at most $3k - 1$ points and the distance between two consecutive points along the path is at most β , that part of the path can have length at most $(6k - 1)\beta$, a contradiction. \square

Corollary 8. *Dense boxes are connected by the 8-neighbor topology.*

Lemma 8. *Any finite set S of at least $4 \cdot (3k - 1) + 1$ points with $\beta \geq \text{BE}(\text{MST}(S))$ contains at least one dense box.*




Figure 13: The distributed approach: a grid is placed over the point set and different representatives construct different graphs ((a) and (b)). The red and black edges form the tree in each dense cell, blue edges connect the dense cells, and orange edges connect the vertices in sparse cells.

Proof. Assume S and β induce only sparse boxes. This implies that the points are distributed over at least five boxes, and thus, there is a pair of boxes that is non-adjacent. Using Lemma 7, this means that these boxes cannot be connected using edges of length at most $\text{BE}(\text{MST}(S))$, a contradiction. \square

Lemma 9. *In any finite set S of at least $4 \cdot (3k - 1) + 1$ points with $\beta \geq \text{BE}(\text{MST}(S))$, all sparse boxes are adjacent to a dense box.*

Proof. This follows from Lemma 7, since any sparse box that is not adjacent to a dense box cannot be connected to any dense box using edges of length at most $\beta \geq \text{BE}(\text{MST}(S))$. \square

Next, we assign all points to dense boxes. In order to do this, let c_B be the center of a dense box B . Note that c_B is not necessarily the center point of the points in this box. We consider the Voronoi diagram of the centers of all dense boxes and assign a point p to B if p lies in the Voronoi cell of c_B . Let S_B be the set of points of S that are associated with a dense box B . We note that each dense box B gets assigned at least all points in its own box, since in the case of adjacent dense boxes, the boundary of the Voronoi cell coincides with the shared boundary of these boxes (see Figure 14).

Furthermore, we can compute the points assigned to each box locally. By Lemma 9 all sparse boxes are adjacent to a dense box, and hence for any point p in a sparse box B its distance to its nearest center is at most $3\ell/\sqrt{2}$, where $\ell = 6k\beta$. It follows that only the centers of cells of neighbors and neighbors of neighbors need to be considered.




Figure 14: The Voronoi cells of the centers of the dense boxes.

Lemma 10. *For any two dense boxes B and B' , we have that the convex hulls of S_B and $S_{B'}$ are disjoint.*

Proof. We observe that the convex hull of S_B is contained in the Voronoi cell of c_B . Hence, since the Voronoi cells of different dense boxes are disjoint, the convex hulls of the points assigned to them are also disjoint. \square

For each dense box B , we apply Theorem 7 on the points inside the dense box to compute k disjoint layers of S_B . Next, we connect all sparse points in S_B to the representative of the sector that contains them in each layer. Since all points in the same sector connect to the same representative and the sectors of the same layer do not overlap, we obtain a plane graph for each layer within the convex hull of each S_B .

Hence, we obtain k pairwise disjoint layers such that in each layer the points associated to each dense box are connected. Moreover, since the created edges stay within the convex hull of each subproblem and by Lemma 10 those hulls are disjoint, each layer is plane. Thus, to assure that each layer is connected, we must connect the construction between dense boxes.

We connect adjacent dense boxes in a tree-like manner using the following rules:

- Connect every dense box to any dense box below it.
- Always connect every dense box to any dense box to the left of it.
- If neither the box below nor the one to the left of it is dense, connect the box to the dense box diagonally below and to the left of it.
- If neither the box above nor the one to the left of it is dense, connect the box to the dense box diagonally above and to the left of it.

To connect two dense boxes, we find and connect two representatives p and q (one from each dense box) such that p lies in the sector of q and q lies in the sector of p ; see Figure 15 (a).

Lemma 11. *For any layer and any two adjacent dense boxes B and B' , there are two representatives p and q in B and B' , respectively, such that p lies in the sector of q and q lies in the sector of p .*




Figure 15: Connecting two dense boxes by means of p and q . The half-circles in (a) indicate which sector each representative covers. The red edges connect the dense boxes internally and the blue edge connects the two dense cells. (b) illustrates the sectors involved in connecting two neighboring dense boxes.

Proof. Consider two boxes B and B' with center points (of their respective point sets) c and c' . Now let W_1 and W'_1 with representatives r_1 and r'_1 denote the sectors containing c' and c , respectively; see Figure 15. The other sectors W_2 and W_3 of B with representatives r_2 and r_3 are ordered clockwise. We use ℓ_i to denote the ray from c containing r_i . If $r_1 \in W'_1$ and $r'_1 \in W_1$ we are done. So assume that $r'_1 \notin W_1$, the case when $r_1 \notin W'_1$ (or when both $r_1 \notin W'_1$ and $r'_1 \notin W_1$) is symmetric. It follows that r'_1 is in sector W_2 if the line segment $c'r'_1$ intersects ℓ_2 or sector W_3 if the segment intersects ℓ_2 and ℓ_3 . Assume that r'_1 is in sector W_2 (again the argument is symmetric when r'_1 is in sector W_3). Now r_2 can be positioned on ℓ_2 between c and the intersection point with $c'r'_1$ or behind this intersection point when viewed from c . In the former case r'_1 is in W_2 and r_2 is in W'_1 and we are done. In the latter case the segments cr_2 and $c'r'_1$ cross. Since $c, r_2 \in B$ and $c', r'_1 \in B'$ this crossing would imply that B and B' are not disjoint, a contradiction. \square

Now that we have completed the description of the construction, we show that each layer of the resulting graph is plane and connected, and that the length of the longest edge is bounded.

Lemma 12. *Each layer is plane.*

Proof. Since dense boxes are internally plane and the addition of edges to the sparse points do not violate planarity, it suffices to show that the edges between dense boxes cannot cross any previously inserted edges and that these edges cannot intersect other edges used to connect dense boxes.

We first show that the edge used to connect boxes B and B' is contained in the union of the Voronoi cells of these two boxes. If B and B' are horizontally or vertically adjacent, the connecting edge stays in the union of the two dense boxes,

which is contained in their Voronoi cells. If B and B' are diagonally adjacent, we connect them only if their shared horizontal and vertical neighbors are not dense. This implies that at least the two triangles defined by the sides of B and B' that are adjacent to their contact point are part of the union of the Voronoi cells of these boxes. Hence, the edge used to connect B and B' cannot intersect the Voronoi cell of any other dense box. Since all points of a dense box in a sector connect to the same representative and these edges lie entirely inside the sector, the edge connecting two adjacent boxes can intersect only at one of the two representatives, but does not cross them. Therefore, an edge connecting two adjacent dense boxes by connecting the corresponding representatives cannot cross any previously inserted edge.

Next, we show that edges connecting two pairs of dense boxes cannot cross. Since any edge connecting two dense boxes stays within the union of the Voronoi cells of B and B' , the only way for two edges to intersect is if they connect to the same box B and intersect in the Voronoi cell of B . If the connecting edges lie in the same sector of B , they connect to the same representative and thus they cannot cross. If they lie in different sectors of B , the edges lie entirely inside their respective sectors. Since these sectors are disjoint, this implies that the edges cannot intersect. \square

Lemma 13. *Each layer is connected.*

Proof. Since the sectors of the representatives of the dense boxes cover the plane, each point in a sparse box is connected to a representative of the dense box it is assigned to. Hence, showing that the dense boxes are connected completes the proof.

By Corollary 8, the dense boxes are connected using the 8-neighbor topology. This implies that there is a path between any pair of dense boxes where every step is one to a horizontally, vertically, or diagonally adjacent box. Since we always connect horizontally or vertically adjacent boxes and we connect diagonally adjacent boxes when they share no horizontal and vertical dense neighbor, the layer is connected after adding edges as described in the proof of Lemma 11. \square

Lemma 14. *The distance between a representative in a dense box B and any point connecting to it is at most $12\sqrt{2}k\beta$.*

Proof. Since the representatives of B are connected only to points from dense and sparse boxes adjacent B , the distance between a representative and a point connected to it is at most the length of the diagonal of the 2×2 grid cells with B as one of its boxes. Since a box has width $6k\beta$, this diagonal has length $2\sqrt{2} \cdot 6k\beta = 12\sqrt{2}k\beta$. \square

Theorem 9. *For all finite point sets with at least $4(3k-1)+1$ points, we can extract k plane layers with the longest edge having length at most $12\sqrt{2}k\text{BE}(\text{MST}(S))$.*

4 Conclusions

We presented two algorithms for constructing k edge-disjoint non-crossing plane spanning graphs on a given point set such that the length of the bottleneck edge is minimized. The first algorithm uses global properties in order to keep all edges as small as possible. We also give matching worst-case lower bounds, making the algorithm tight. The main drawback is that this method can only be used to construct two layers, and it is unlikely that a similar approach can work for more. Our second algorithm works for a large number of layers (up to $n/12$ layers). It uses only local information, thus it can be executed in a distributed manner. The drawback of this approach is that the length of the bottleneck edge grows considerably: for two layers, the $24\sqrt{2}\text{BE}(\text{MST}(S))$ implied by this method is far larger than the $3\text{BE}(\text{MST}(S))$ of the first approach.

So far, there is no centralized method to construct more than two trees. Finding such a method and comparing it to the distributed method presented here is an interesting direction of future research. Another direction would be to lower the length of the longest edge in the distributed construction, though from a purely worst-case theoretical point of view this is likely to require a different approach from the one used in this paper.

Acknowledgments

This research was initiated during the 10th European Research Week on Geometric Graphs (GGWeek 2013), Illgau, Switzerland. We would like to thank all participants for fruitful discussions. O.A., A.P., and B.V. were partially supported by the ESF EUROCORES programme EuroGIGA - ComPoSe, Austrian Science Fund (FWF): I 648-N18. T.H. was supported by the Austrian Science Fund (FWF): P23629-N18 ‘Combinatorial Problems on Geometric Graphs’. M.K. was supported in part by the ELC project (MEXT KAKENHI No.17K12635) and NSF award CCF-1423615.. A.P. is supported by an Erwin Schrödinger fellowship, Austrian Science Fund (FWF): J-3847-N35. A.v.R. and M.R. were supported by JST ERATO Grant Number JPMJER1201, Japan.

References

- [1] Aichholzer, O., Hackl, T., Korman, M., van Kreveld, M.J., Löffler, M., Pilz, A., Speckmann, B., Welzl, E.: Packing plane spanning trees and paths in complete geometric graphs. *Inf. Process. Lett.* 124, 35–41 (2017)
- [2] Aichholzer, O., Hackl, T., Korman, M., Pilz, A., Rote, G., van Renssen, A., Roeloffzen, M., Vogtenhuber, B.: Packing short plane spanning trees in complete geometric graphs. In: *ISAAC. LIPIcs*, vol. 64, pp. 9:1–9:12. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2016)

- [3] Biniaz, A., García, A.: Packing plane spanning trees into a point set. In: Proc. 30th Canadian Conference on Computational Geometry, (CCCG). pp. 49–53 (2018)
- [4] Dor, D., Tarsi, M.: Graph decomposition is NP-complete: A complete proof of Holyer’s conjecture. SIAM J. Comput. 26(4), 1166–1187 (1997)
- [5] Fussen, M., Wattenhofer, R., Zollinger, A.: Interference arises at the receiver. In: Proc. International Conference on Wireless Networks, Communications, and Mobile Computing (WIRELESSCOM). pp. 427–432 (2005)
- [6] Korman, M.: Minimizing interference in ad-hoc networks with bounded communication radius. Inf. Process. Lett. 112(19), 748–752 (2012)
- [7] Kranakis, E., Singh, H., Urrutia, J.: Compass routing on geometric networks. In: Proc. 11th Canadian Conference on Computational Geometry (CCCG). pp. 51–54 (1999)
- [8] Matoušek, J.: Lectures on Discrete Geometry. Graduate Texts in Mathematics, Springer (2002)
- [9] Priesler, M., Tarsi, M.: Multigraph decomposition into stars and into multistars. Discrete Mathematics 296(2-3), 235–244 (2005)
- [10] Tarsi, M.: Decomposition of a complete multigraph into simple paths: Nonbalanced handcuffed designs. J. Comb. Theory, Ser. A 34(1), 60–70 (1983)