Packing Short Plane Spanning Trees in Complete Geometric Graphs*†

Oswin Aichholzer[1], Thomas Hackl[2], Matias Korman[3], Alexander Pilz[4], Günter Rote[5], André van Renssen[6], Marcel Roeloffzen[7], and Birgit Vogtenhuber[8]

1 **Institute for Software Technology, Graz University of Technology, Austria**
oaich@ist.tugraz.at
2 **Institute for Software Technology, Graz University of Technology, Austria**
thackl@ist.tugraz.at
3 **Tohoku University, Sendai, Japan**
mati@dais.is.tohoku.ac.jp
4 **Department of Computer Science, ETH Zürich, Switzerland**
alexander.pilz@inf.ethz.ch
5 **Institut für Informatik, Freie Universität Berlin, Germany**
rote@inf.fu-berlin.de
6 **National Institute of Informatics, Tokyo, Japan; and JST, ERATO, Kawarabayashi Large Graph Project, Japan**
andre@nii.ac.jp
7 **National Institute of Informatics, Tokyo, Japan; and JST, ERATO, Kawarabayashi Large Graph Project, Japan**
marcel@nii.ac.jp
8 **Institute for Software Technology, Graz University of Technology, Austria**
bvogt@ist.tugraz.at

Abstract

Given a set of points in the plane, we want to establish a connection network between these points that consists of several disjoint layers. Motivated by sensor networks, we want that each layer is spanning and plane, and that no edge is very long (when compared to the minimum length needed to obtain a spanning graph). We consider two different approaches: first we show an almost optimal centralized approach to extract two trees. Then we show a constant factor approximation for a distributed model in which each point can compute its adjacencies using only local information. This second approach may create cycles, but maintains planarity.

1998 ACM Subject Classification I.3.5 Computational Geometry and Object Modeling

Keywords and phrases Geometric Graphs, Graph Packing, Plane Graphs, Minimum Spanning Tree, Bottleneck Edge

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2016.9

* This research was initiated during the 10th European Research Week on Geometric Graphs (GGWeek 2013), Illgau, Switzerland.

† 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. 12H00855 and 15H02665). A. P. is supported by an Erwin Schrödinger fellowship, Austrian Science Fund (FWF): J-3847-N35.

27th International Symposium on Algorithms and Computation (ISAAC 2016).
Editor: Seok-Hee Hong; Article No. 9; pp. 9:1–9:12
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

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 trees $H_1, H_2, \ldots, H_k$ on S such that the length $\mathsf{BE}(H_1 \cup H_2 \cup \cdots \cup H_k)$ of the *bottleneck edge* (the longest edge which is used) is as short as possible. Each tree H_i is referred to as a *layer* of G. 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 $\mathsf{MST}(S)$ solves the problem: its longest edge $\mathsf{BE}(\mathsf{MST}(S))$ is a lower bound on the bottleneck edge of any spanning subgraph, and it is non-crossing. For larger k, we take $\mathsf{BE}(\mathsf{MST}(S))$ as the yardstick and measure the solution quality in terms of $\mathsf{BE}(\mathsf{MST}(S))$ and k.

The particular variation that we consider comes motivated from the field of sensor networks. Imagine one wants to construct a network so that afterwards communication between sensors is possible. One of the most important requirements for such a network is that we can send messages through it easily. Ideally, we want a method that – given the source, destination, information on the current position (and possibly $O(1)$ additional information) – computes the next node to visit in order to reach our destination.

One of the most famous such methods is called *face routing* [7], which guarantees the delivery under the above constraints provided that the underlying graph is plane. Indeed, when considering local routing algorithms in the literature that are guaranteed to succeed, most route deterministically on a plane spanning subgraph of the underlying graph where the plane subgraph can be computed locally. Even though there exist routing strategies for non-plane graphs, in most cases they route through a plane subgraph (for example, Bose *et al.* [2] showed how to locally identify the edges of the Gabriel graph from the unit disk graph). Extending these algorithms for non-plane graphs is a long-standing open problem.

It seems counter-intuitive that having additional edges cannot help in the delivery of messages. In this paper, we provide a different way to avoid this obstacle. Rather than limiting considerations to one plane graph, we aim to construct several disjoint plane spanning graphs. If we split all the messages among the different layers (and route through each layer with routing strategies that work on plane graphs) we can potentially spread the load among a larger number of edges. Another important feature to consider when creating networks is energy consumption. The required energy for sending a message increases with the distance between the two points (usually with the third or fourth power) [4]. Since we want to avoid high energy consumption at one particular node, it is desirable to apply the bottleneck criterion and to minimize the longest edge [6].

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 pairwise disjoint subgraphs $H_1 = (V, E_1), H_2 = (V, E_2), \ldots$ into G. A related problem is the *decomposition* of G. In this case, we also look for disjoint subgraphs but require that $\cup_i E_i = E$. For example, there are known characterizations of when we can decompose the complete graph of n points into paths [9] and stars [8]. Dor and Tarsi [3] showed that to determine whether we can decompose a graph G into subgraphs isomorphic to a given graph H is NP-complete. Aichholzer *et al.* [1] showed that any set of n points contains $\Omega(\sqrt{n})$ disjoint plane spanning trees. This bound has been improved to $\lfloor n/3 \rfloor$ by Garcia [5].

In our case, the graph G consists of the complete graph on S, and $\mathcal{F}$ consists of all plane spanning trees of G. 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. We give two different approaches to solve the problem. In Section 2 we give a construction for $k = 2$ trees. This construction is centralized in a classical model that assumes that the position of all points are known and computed in a single place. Our construction guarantees that all edges (except possibly one) have length at most $2\mathsf{BE}(\mathsf{MST}(S))$. The remaining edge has length at most $3\mathsf{BE}(\mathsf{MST}(S))$. We complement this construction with a matching worst-case lower bound.

Following the spirit of sensor networks, in Section 3 we use a different approach to construct k disjoint plane graphs (not necessarily trees). The construction works for any $k \leq n/12$ in an almost local fashion. The only global information that is needed is β: $\mathsf{BE}(\mathsf{MST}(S))$ or some upper bound. Each point of S can compute its adjacencies by only looking at nearby points: those at distance $O(k\beta)$.

A simple adversary argument shows that it is impossible to construct spanning networks locally without knowing $\mathsf{BE}(\mathsf{MST}(S))$ (or an upper bound). The lower bound of Section 2 shows that a neighborhood of radius $\Omega(k\mathsf{BE}(\mathsf{MST}(S)))$ 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 colinear. Without this assumption, it might be impossible to obtain more than a single plane layer (for example, when all points lie on a line).

2 Centralized Construction

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

▶ **Lemma 1.** *Let S be a set of points in the plane and let uv and vw be two edges of $\mathsf{MST}(S)$. Then the triangle uvw does not contain any other point of S.*

Proof. Observe that, as v is adjacent to both u and w in $\mathsf{MST}(S)$, uw is the longest edge of the triangle uvw (otherwise one could locally shorten $\mathsf{MST}(S)$).

Suppose for the sake of contradiction that there is a point $p \in S$ in the interior of uvw. We split uvw into two sub-triangles by the line ℓ through v perpendicular to the supporting line of u and w. Let Δ_u be the sub-triangle that has u as a vertex, and assume w.l.o.g. that p lies in Δ_u. Note that the edge uv is the hypotenuse of the right-angled triangle Δ_u and hence $\max\{|pu|, |pv|\} < |uv|$.

Consider the paths in $\mathsf{MST}(S)$ from p to u and v, respectively. Since $\mathsf{MST}(S)$ is a tree, one of the two paths must use the edge uv (as otherwise there would be a cycle). Suppose first that this edge is used in the path to u. By removing the edge uv and adding the edge pu to $\mathsf{MST}(S)$ we would obtain a connected (not necessarily plane) tree whose overall weight is smaller, a contradiction. If the edge used is in the path to v, the addition of edge pv yields a similar contradiction. ◀

▶ **Lemma 2.** *Let S be a set of points in the plane. Let $v \in S$ be a point of degree $k \geq 3$ in $\mathsf{MST}(S)$, with $\{v_0, \ldots, v_{k-1}\}$ being the neighbors of v in $\mathsf{MST}(S)$ in counterclockwise order around v. Then for every triple (v_{i-1}, v_i, v_{i+1}) (indices modulo k), the neighbors of v_i in $\mathsf{MST}(S)$ 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. Let $u \in S \setminus \{v\}$ be a neighbor of v_i in $\mathsf{MST}(S)$, and assume for the sake of contradiction that u is not in W_i. Then the edge v_iu intersects the boundary of W_i and hence one of the rays starting at v and going through v_{i-1} and v_{i+1}, respectively. Assume without loss of generality that v_iu intersects the ray from v through v_{i+1}. As $\mathsf{MST}(S)$ is plane, the edge v_iu does not intersect the edge vv_{i+1}. Hence, the triangle (v, v_i, u) contains the point v_{i+1} in its interior. As the path vv_iu is a subgraph of $\mathsf{MST}(S)$, this contradicts Lemma 1. ◀

We denote by $\mathsf{MST}^2(S)$ the square of $\mathsf{MST}(S)$, the graph connecting all pairs of points of S that are at distance at most 2 in $\mathsf{MST}(S)$. We call the edges of $\mathsf{MST}(S)$ *short* edges and all remaining edges of $\mathsf{MST}^2(S)$ *long* edges. For every long edge uw, the points u and w have a unique common neighbor v in $\mathsf{MST}(S)$, 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. Next we state a simple fact on crossings of the edges in $\mathsf{MST}^2(S)$.

▶ **Lemma 3.** *Let S be a set of points in the plane. Two edges e and f of $\mathsf{MST}^2(S)$ cross if and only if one of the following two conditions is fulfilled:*

1. *At least one of $\{e, f\}$ 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 of $\{e, f\}$ are long edges with the same witness v, and their wedges are intersecting but none is contained in the other.*

Proof. Clearly, if both e and f are short edges, i.e., edges of $\mathsf{MST}(S)$, then they do not cross. Let $f = uw$ be a long edge with witness v and wedge W. Every edge $e = vz$ of $\mathsf{MST}^2(S)$, $z \in S \setminus \{u, v, w\}$ that lies inside W either crosses f or has z inside the triangle $\Delta = (u, v, w)$. The latter is a contradiction to Lemma 1. Obviously, f is neither crossed by any edge incident to u or w, nor crossed by any edge incident to v but not lying inside W.

It remains to prove that every long edge $e = xz$ of $\mathsf{MST}^2(S)$, $x, z \in S \setminus \{u, v, w\}$ that crosses f fulfills Condition 2. Note that for e to cross f, either e has an endpoint inside Δ or e is also crossing one edge out of $\{uv, vw\} \in \mathsf{MST}(S)$. The former is a contradiction to Lemma 1. If e is a short edge (i.e., an edge of $\mathsf{MST}(S)$), then the latter is a contradiction to the planarity of $\mathsf{MST}(S)$. Hence, e is a long edge (with wedge W') and is also crossing one edge g out of $\{uv, vw\} \in \mathsf{MST}(S)$. This also implies that the wedges W and W' intersect in their interiors but none of W, W' is contained in the other. Finally, if e has witness $y \neq v$, then either g has an end point in the triangle xyz or g crosses one edge out of $\{xy, yz\} \in \mathsf{MST}(S)$. Again, the former is a contradiction to Lemma 1 and the latter is a contradiction to the planarity of $\mathsf{MST}(S)$. Hence the witness of e must be v. ◀

With the above observations we can proceed to show a construction that almost works for two layers. To this end we consider the minimum spanning tree $\mathsf{MST}(S)$ to be rooted at a leaf r. For any $v \in S$, we define its *level* $\ell(v)$ as its distance to r in $\mathsf{MST}(S)$. 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 $\mathsf{MST}(S)$. 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.

▶ Construction 4. Let S be a set of points in the plane and let $\mathsf{MST}(S)$ be rooted at one of its leaves, $r \in S$. 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 $v_op(v_o)$ to E_R and the edge $v_og(v_o)$ to E_B. For any vertex $v_e \in S \setminus \{r\}$ whose level is even, we add the edge $v_eg(v_e)$ to E_R and the edge $v_ep(v_e)$ 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.

▶ **Theorem 5.** *Let MST(S) be rooted at r. The two graphs $R = G(S, E_R)$ and $B = G(S, E_B)$ from Construction 4 fulfill the following properties:*
1. *Both R and B are plane spanning trees.*
2. $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.
3. $E_R \cap E_B = \{rs\}$*, with* $r = p(s)$*, i.e.,* $|E_R \cap E_B| = 1$.

Proof. Recall from Construction 4 that r is a leaf of $\mathsf{MST}(S)$. Hence r has a unique neighbor s in $\mathsf{MST}(S)$ 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 the construction, the set of red edges is $E_R = \bigcup_{v_o \in S_o} \{v_o p(v_o)\} \cup \bigcup_{v_e \in S_e} \{v_e g(v_e)\} \cup \{rs\}$ and the set of blue edges 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 single shared edge between the sets E_R and E_B, as stated in Property 3.

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

Further, both R and B are spanning trees, i.e., connected and cycle free graphs, as each vertex except r is connected either to its parent or grandparent in $\mathsf{MST}(S)$. 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 $\mathsf{MST}^2(S)$ whose endpoints have different levels. By Lemma 3, 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 due to its level. That is, uv is not crossing f (common endpoint) and all other edges incident to v in E_R and E_B are either blue if f is red, or red if f is blue. Further, v cannot be the witness of e due to its level. All edges E_R and E_B with witness v have u as one of its endpoints (as for all other edges with witness v in $\mathsf{MST}^2(S)$, both endpoints have the same level). With u as a shared vertex, the edges e and f cannot cross. As e is neither incident to v nor has v as a witness, e crossing f is a contradiction to Lemma 3. This proves Property 1 and concludes the proof. ◀

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

▶ **Corollary 6.** *For any set S of n points in the plane, there exist two plane spanning trees $R = G(S, E_R)$ and $B = G(S, E_B)$ such that $|E_R \cap E_B| = 1$ and $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.*

Construction 4 is almost valid in the sense that only one edge was shared between both trees. In the following we enhance this construction so as to avoid the shared edge.

Let $N^- \subset (S \setminus \{r\})$ be the set of neighbors $v^- \in N^-$ of s in $\mathsf{MST}(S)$ such that the ordered triangle rsv^- is oriented clockwise. Let $N^+ \subset (S \setminus \{r\})$ be the set of neighbors $v^+ \in N^+$ of s in $\mathsf{MST}(S)$ such that the ordered triangle rsv^+ is oriented counter-clockwise. Let T^- be the subtree of $\mathsf{MST}(S)$ that is connected to s via the vertices in N^- and let T^+ be the subtree of $\mathsf{MST}(S)$ 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 $S^- \cap S^+ = \{r, s\}$. 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 $B^+ = G(S^+, E_B^+)$. The following property follows from Lemma 3.

▶ **Lemma 7.** *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 4. Then no edge in E_R^- crosses an edge in E_B^+ and no edge in E_R^+ crosses any edge in E_B^-.*

Proof. Consider any edge $e \in E_R^-$ that is not incident to r. By Lemma 3, 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 exists an edge $f \in E_B^+$ that crosses an edge $e \in E_R^-$ that is incident to r. By construction, $e = rz$ is a long edge of $\mathsf{MST}^2(S)$ with witness s and wedge W. By Lemma 3, 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 is a contradiction to Lemma 3. Hence, let $f = sc$ be a long edge of $\mathsf{MST}^2(S)$ with witness b. Following Lemma 3, the witness b must be s, which is in contradiction to 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^-. ◀

With this observation we can now prove that the two spanning trees from Construction 4 actually exist in 4 different color combination variants.

▶ **Lemma 8.** *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 4 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 be (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"), and (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"), such that the properties from Theorem 5 hold for all versions.*

Proof. The statement is trivially true for recolorings (1) and (2). It is easy to observe that this really is corresponding to a simple recoloring. Hence, Properties 2 and 3 of Theorem 5 are also obviously true. By Lemma 7, both R and B are plane for the recolorings (3) and (4) and thus fulfill Property 1 of Theorem 5 as well. ◀

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

▶ **Theorem 9.** *Consider a set S of n points in the plane for which the minimum spanning tree $\mathsf{MST}(S)$ has a vertex v where between any two consecutive adjacent edges the angle 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 $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 2\mathsf{BE}(\mathsf{MST}(S))$.*

Proof (sketch). When removing v from the tree, we obtain up to five connected components (assuming general position). For each of these, we individually re-add v and apply Construction 4 with v as the root using one of the variants of Lemma 8. This leaves some components of the tree disconnected, but this is resolved by adding additional edges from v to its neighbors and between adjacent neighbors of v. Full details of the construction and a proof can be found in the full version of this paper. ◀

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

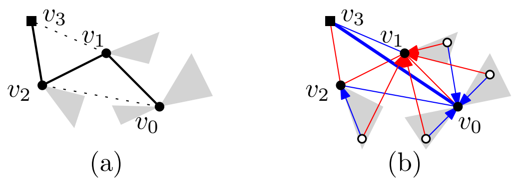

Figure 1 Illustration of one of the cases for the proof of Theorem 10. Grey subtrees indicate potential continuations of the MST and dashed edges indicate edges from $\mathsf{MST}^2(S)$. In (b) colored arrows indicate how the subtrees connect to P. Note that half of these arrows are from Construction 4.

▶ **Theorem 10.** *Consider a set S of $n \geq 4$ points in the plane for which every vertex in the minimum spanning tree $\mathsf{MST}(S)$ 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 $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 3\mathsf{BE}(\mathsf{MST}(S))$ (where at most one edge of $E_R \cup E_B$ is larger than $2\mathsf{BE}(\mathsf{MST}(S))$).*

Proof (sketch). This case is dealt with using similar ideas as for Theorem 9. The main difference is that we now use a cluster $P = \{v_0, v_1, v_2, v_3\}$ of 4 points that are connected in $\mathsf{MST}(S)$ to serve as roots for up to three subtrees. The exact choice of P depends on the exact embedding of the tree, which leads to several potential embeddings of P and the subtrees of $\mathsf{MST}(S)$ attached to P. For this proof sketch we focus on one specific case shown in Figure 1, where v_3 is a leaf and v_2, v_1, v_0 form a path that, starting from v_3, takes a left and right turn. For ease of description we root the entire MST at v_3, creating parent and child relations between nodes. The subtrees we consider are T_0, T_1, T_2 defined as follows:

- T_0, consisting of v_1, v_0, and the subtrees rooted at the children of v_0, rooted at v_1.
- T_1, consisting of v_1, v_0 and the subtrees rooted at children of v_1, rooted at v_0.
- T_2, consisting of v_2, v_1 and the subtrees rooted at children of v_2, rooted at v_1.

Each of these trees is colored using one of the variants of Lemma 8, but we remove all edges going to the roots of the respective subtrees, leaving its children disconnected from P. We then re-attach them as follows. The roots of disconnected subtrees of T_0 are connected to v_1, those from T_1 are attached to v_0 and those from T_2 to v_1. By construction, the red and blue trees then form spanning trees with a maximum edge length of $3\mathsf{BE}(\mathsf{MST}(S))$ as all edges except v_0v_3 are part of $\mathsf{MST}^2(S)$, and v_0v_3 is part of $\mathsf{MST}^3(S)$. For planarity, non-crossing of edges that are not v_0v_3 follows relatively easily from Lemma 3 and Theorem 5. To see that v_0v_3 cannot be crossed, one can observe that by Lemma 1 the convex hull of P must be empty and from Lemma 2 and 3 it follows that no edge can cross the convex hull through v_3v_1 to v_1 or v_2. Full details of the construction and correctness arguments can be found in the full version of this paper. ◀

▶ **Corollary 11.** *For any set S of $n \geq 4$ points in the plane, 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 $\max\{\mathsf{BE}(R), \mathsf{BE}(B)\} \leq 3\mathsf{BE}(\mathsf{MST}(S))$.*

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

▶ **Theorem 12.** *For any $n > 3$ and $k > 1$ there exists a set of n points such that for any k disjoint spanning trees, at least one has a bottleneck edge larger than $(k+1)\mathsf{BE}(\mathsf{MST}(S))$.*

Proof. A counterexample simply consists of n points equally distributed on a line segment. (The points can be slightly perturbed to obtain general position.) In this problem instance

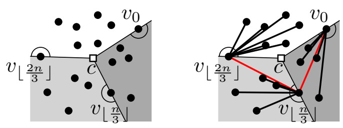

Figure 2 Extracting one layer: (a) The three sectors defined by v_0, $v_{\lfloor \frac{n}{3} \rfloor}$, and $v_{\lfloor \frac{2n}{3} \rfloor}$. (b) Connecting the points to the representative of their sector. The red edges connect the representatives.

there are $kn-(k(k+1)/2)$ edges whose distance is strictly less than $(k+1)\mathsf{BE}(\mathsf{MST}(S)) = k+1$. However, we need $kn-k$ edges for k disjoint trees and thus it is impossible to construct that many trees with sufficiently short edges. ◀

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.

In the following, we provide an alternative construction. Although the length of the edges is increased by a constant factor, it has the benefit that it can be constructed locally and that it can be extended to compute k layers. The only global property that is needed is a value β that should be at least $\mathsf{BE}(\mathsf{MST}(S))$. 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 García [5] that states that every point set of at least $3k$ points contains k layers. Since the details of this construction are important for our construction and the manuscript is not yet available, we add a proof sketch.

▶ **Theorem 13** ([5]). *Every 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 $n/3$ points. For ease of explanation, we assume that every line through c contains at most one point. Number the points $v_0, v_1, \ldots, v_{n-1}$ 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 2. Since every line through the center contains at least $n/3$ points on each side, 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 $v_1, \ldots, v_{\lfloor \frac{n}{3} \rfloor}$ 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 $v_{\lfloor \frac{n}{3} \rfloor+1}, \ldots, v_{\lfloor \frac{2n}{3} \rfloor}$ to $v_{\lfloor \frac{n}{3} \rfloor}$. Finally, we connect vertices $v_{\lfloor \frac{2n}{3} \rfloor+1}, \ldots, v_{n-1}$ to $v_{\lfloor \frac{2n}{3} \rfloor}$. This results in a non-crossing spanning tree.

For the second tree, we use v_1, $v_{\lfloor \frac{n}{3} \rfloor+1}$, and $v_{\lfloor \frac{2n}{3} \rfloor+1}$, and so on. ◀

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 3).

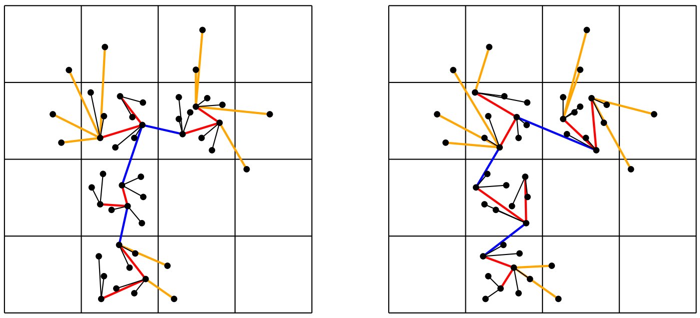

(a) (b)

Figure 3 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.

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 the separating lines, pick 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. We observe that dense and sparse boxes satisfy the following properties.

▶ **Lemma 14.** *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 dense boxes s.t. there is a path that uses edges of length at most β between a point in B to 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, with adjacent duplicates removed. Observe first that horizontal and vertical sides alternate in σ, as otherwise the path would have to use at least $6k-1$ points to traverse a sparse box, but there are only at most $3k-1$. 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. ◀

▶ **Corollary 15.** *Dense boxes are connected by the 8-neighbor topology.*

▶ **Lemma 16.** *Any set S of at least $4 \cdot (3k-1)+1$ points with $\beta \geq \mathsf{BE}(\mathsf{MST}(S))$ contains at least one dense box.*

Proof. Assume S consists of 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 14, this means that these boxes cannot be connected using edges of length at most $\mathsf{BE}(\mathsf{MST}(S))$, a contradiction. ◀

▶ **Lemma 17.** *In any set S of at least $4 \cdot (3k-1)+1$ points with $\beta \geq \mathsf{BE}(\mathsf{MST}(S))$, all sparse boxes are adjacent to a dense box.*

Proof. This follows from Lemma 14, 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 $\beta \geq \mathsf{BE}(\mathsf{MST}(S))$. ◀

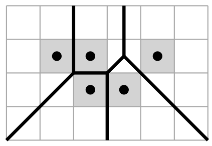

Figure 4 The Voronoi cells of the centers of the dense boxes.

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 4).

Furthermore, we can compute the points assigned to each box locally. By Lemma 17 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.

▶ **Lemma 18.** *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. ◀

For each dense box B, we apply Theorem 13 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 18 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:

- Always connect a dense box to the dense box below it.
- Always connect a dense box to the 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 5 (a).

▶ **Lemma 19.** *For any layer and any two adjacent dense boxes B and B', there are two representatives p and q in B and B', respectively, s.t. p lies in the sector of q and q lies in the sector of p.*

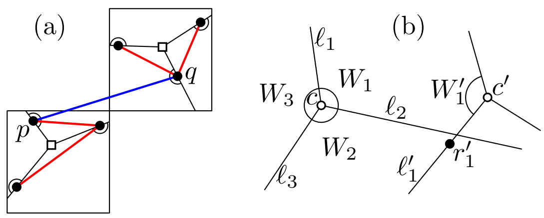

Figure 5 Connecting two dense boxes by means of p and q. The half-circles in (a) indicate which sector each representative covers. The red edges connect the dense boxes internally and the blue edge connects the two dense cells. (b) illustrates the sectors involved in connecting two neighboring dense 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 5. 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. ◀

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 20.** *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 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 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. ◀

▶ **Lemma 21.** *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 15, 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 19. ◀

▶ **Lemma 22.** *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 with B as one of its boxes. Since a box has width $6k\beta$, this diagonal has length $2\sqrt{2} \cdot 6k\beta = 12\sqrt{2}k\beta$. ◀

▶ **Theorem 23.** *For all point sets with at least $4(3k-1)+1$ points, we can extract k plane layers with the longest edge having length at most $12\sqrt{2}k\mathsf{BE}(\mathsf{MST}(S))$.*

Acknowledgements. We would like to thank all participants of the 10th European Research Week on Geometric Graphs for fruitful discussions.

References

1. O. Aichholzer, T. Hackl, M. Korman, M. J. van Kreveld, M. Löffler, A. Pilz, B. Speckmann, and E. Welzl. Packing plane spanning trees and paths in complete geometric graphs. In *Proc. Canadian Conference on Computational Geometry (CCCG)*, pages 233–238, 2014.
2. P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
3. D. Dor and M. Tarsi. Graph decomposition is NP-complete: A complete proof of Holyer's conjecture. *SIAM Journal on Computing*, 26(4):1166–1187, 1997.
4. M. Fussen, R. Wattenhofer, and A. Zollinger. Interference arises at the receiver. In *Proc. International Conference on Wireless Networks, Communications, and Mobile Computing (WIRELESSCOM)*, pages 427–432, 2005.
5. A. García. Personal Communication, 2015.
6. M. Korman. Minimizing interference in ad-hoc networks with bounded communication radius. *Information Processing Letters*, 112(19):748–752, 2012.
7. E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Proc. Canadian Conference on Computational Geometry (CCCG)*, pages 51–54, 1999.
8. M. Priesler and M. Tarsi. Multigraph decomposition into stars and into multistars. *Discrete Mathematics*, 296(2-3):235–244, 2005.
9. M. Tarsi. Decomposition of a complete multigraph into simple paths: Nonbalanced handcuffed designs. *Journal of Combinatorial Theory, Series A*, 34(1):60–70, 1983.