Bounded-Angle Minimum Spanning Trees

Ahmad Biniaz* Prosenjit Bose† Anna Lubiw‡ Anil Maheshwari†

April 21, 2020

Abstract

Motivated by the connectivity problem in wireless networks with directional antennas, we study bounded-angle spanning trees. Let P be a set of points in the plane and let α be an angle. An α-ST of P is a spanning tree of the complete Euclidean graph on P with the property that all edges incident to each point $p \in P$ lie in a wedge of angle α centered at p. We study the following closely related problems for $\alpha = 120°$ (however, our approximation ratios hold for any $\alpha \geqslant 120°$).

1. The *α-minimum spanning tree* problem asks for an α-ST of minimum sum of edge lengths. Among many interesting results, Aschner and Katz (ICALP 2014) proved the NP-hardness of this problem and presented a 6-approximation algorithm. Their algorithm finds an α-ST of length at most 6 times the length of the minimum spanning tree (MST). By adopting a somewhat similar approach and using different proof techniques we improve this ratio to 16/3.
2. To examine what is possible with non-uniform wedge angles, we define an $\overline{\alpha}$-ST to be a spanning tree with the property that incident edges to all points lie in wedges of average angle α. We present an algorithm to find an $\overline{\alpha}$-ST whose largest edge-length and sum of edge lengths are at most 2 and 1.5 times (respectively) those of the MST. These ratios are better than any achievable when all wedges have angle α. Our algorithm runs in linear time after computing the MST.

1 Introduction

A wireless network can be represented by disks in the plane, where a wireless node at point p with transmission range r is represented by a disk of radius r centered at point p. An edge of the network connects two nodes if each one is inside the disk centered at the other one. The question of assigning transmission ranges to the nodes to ensure a well-connected network of low interference has been widely studied [5, 6, 10, 13, 15, 16, 20, 23]. If different nodes may have different transmission ranges then we obtain "power assignment" problems, which have also been heavily studied. The minimum transmission range to ensure network connectivity is the *bottleneck* of the minimum Bottleneck Spanning Tree (BST)—equivalently, the maximum edge-length in a Minimum Spanning Tree (MST).

In recent years, the idea of replacing omni-directional antennas with *directional antennas* has received considerable attention (see, e.g., [1, 4, 8, 11, 12, 14, 15, 24]). In this model, the full disk at each point p is restricted to a circular wedge with apex p that has some angle α and is oriented in some direction. Directional antennas are desirable in many ways, for example, they require

*School of Computer Science, University of Windsor; part of this work was done while the author was an NSERC postdoctoral fellow at University of Waterloo. ahmad.biniaz@gmail.com

†School of Computer Science, Carleton University, Supported by NSERC, {jit, anil}@scs.carleton.ca

‡Cheriton School of Computer Science, University of Waterloo, Supported by NSERC, alubiw@uwaterloo.ca

lower transmission power, cause less interference, and provide more secure communication (see [5, 24] and references therein). The *symmetric communication network* [5] has an edge between two nodes if each one is inside the other's wedge.

When every node has the same transmission range r and angle α, there is still freedom to orient the directional antennas. The question of whether r and α permit a connected network is NP-hard (see Further Background below for details). Most previous work has concentrated on the case where α is some fixed value. Aschner and Katz [4] formulated this in terms of an *α-Spanning Tree* (α-ST): a spanning tree of the complete Euclidean graph on a point set P in the plane such that for each point $p \in P$ all the edges incident to p lie in a wedge of angle α centered at p (see Figure 1-left). For any $\alpha < \pi/3$, an α-ST may not exist, for example, if P is the set of vertices of an equilateral triangle. However, for any $\alpha \geqslant \pi/3$, an α-ST always exists (see [1], [2], and [12] for three different and somewhat involved proofs). There is a relationship between α-STs and the well-studied concept of restricted degree spanning trees, since d edges at a vertex always lie in some wedge of angle at most $2\pi(1 - 1/d)$.

Figure 1: Left: A 120°-ST ($\frac{2\pi}{3}$-ST). Right: a degree-5 minimum spanning tree which is a $\frac{8\pi}{5}$-ST.

To evaluate edge lengths of α-STs two concepts are useful. An *α-Bottleneck Spanning Tree* (α-BST) is an α-ST that minimizes the maximum edge length, and an *α-Minimum Spanning Tree* (α-MST) is an α-ST that minimizes the sum of the edge lengths. Both are NP-hard to compute. Although any MST is also a BST [10], this is not necessarily true for α-MST and α-BST—we give an example later on. For $\alpha \geqslant 120°$ Aschner and Katz [4, 22] gave a simple polynomial time 5-approximation algorithm for the α-BST.

The main result of Aschner and Katz [4] is an involved 6-approximation algorithm for the α-MST for $\alpha \geqslant 120°$. In particular, they construct an α-ST of length at most 6 times the length of an MST of the points. Our main technical result is to improve this approximation factor to 16/3. An interesting aspect of their proof is that they orient the wedges at small subsets of the points so that the network on these points is connected *and* the wedges cover the whole plane. The problem of orienting wedges at fixed points to cover (or "light up") the plane is called the "Floodlight Problem". Bose et al. [9] showed that for any n points and any set $\{\alpha_1, \ldots, \alpha_n\}$ of angles that are at most π and sum to 2π, there is a way to assign angles to points and orient the wedges to light up the plane. To obtain our main result, we prove a strengthened version of the Floodlight result for $n = 3$ where we obtain a connected symmetric communication network even if the angles are pre-assigned to the points.

The new problem we study in this paper is the directional antenna problem when different nodes are allowed to have different wedge angles. The goal is to minimize the sum of the angles. For consistency with previous notation, we define an $\overline{\alpha}$-ST to be a spanning tree of the point set P such that for each point $p \in P$ all the edges incident to p lie in a wedge of angle α_p centered at p and the average of the α_p's is α, or in other words, $\sum_p \alpha_p = n\alpha$. One might hope that, as with the Floodlight Problem, a constant angle sum would suffice to construct a good network. Indeed, taking a star as a spanning tree achieves constant angle sum (with π at an extreme point taken as the center, and 0 at the leaves of the star)—but this uses a large transmission range. We show that it is not possible in general to achieve constant angle sum and transmission ranges bounded by a constant times the BST bottleneck.

In the positive direction, we show that allocating an angle sum of $n \cdot 120°$ non-uniformly does help, in particular, we can achieve smaller maximum edge-length and sum of edge lengths (compared to the MST) than is possible with uniform angles. Now we give more details about our results and techniques, and more details on background.

1.1 Our Results

We obtain the following results for angle $\alpha = 120°$, however, our approximation ratios hold for any angle $\alpha \geqslant 120°$.

(1) Aschner and Katz [4] gave an elegant algorithm that finds a 120°-ST of length at most 6 times the MST length, thus providing a 6-approximation algorithm for the 120°-MST. Our main technical result (Theorem 1 in Section 3) is to improve the approximation ratio to 16/3.

Aschner and Katz prove their result using (in their words) a "surprising theorem", proved via a long case-analysis, that there is a way to assign 120° wedges to any 3 points (a "triplet") such that: the union of the 3 wedges covers the plane; the resulting graph on the triplet is connected; and the resulting graph on any pair of triplets (whose wedges are assigned independently) is also connected. After that, their approach is to take an MST of the points, double its edges and take short-cuts to obtain a Hamiltonian path of length at most 2 times the MST length (as in the standard TSP approximation), and then apply their result to successive triplets of points on the path. Their theorem guarantees that successive pairs of triplets are connected. Finally, they get an approximation factor of 6 using the triangle inequality, and a judicious choice of whether to start the partition into triplets at the first, second, or third point of the path. One limitation of their approach is that the two edges used in each triplet may be the longest and second longest edges of the triplet.

We follow the same approach. We take a non-crossing Hamiltonian path and consider successive triplets of points along the path. The fact that edges of the Hamiltonian path are non-crossing ensures that the triplets are also "non-crossing" in some sense. This allows us to assign wedges to the triplets in a different way which leads to a shorter proof of the"surprising theorem" for non-crossing triplets, as well as a better length guarantee within each triplet.

Our algorithm is slower than that of Aschner and Katz because of the extra work of uncrossing the edges of the Hamiltonian path.

(2) We give an algorithm to find a $\overline{120°}$-ST whose largest edge-length is at most 2 times that of the MST, and whose sum of edge lengths is at most 1.5 times that of the MST (Theorem 6 in Section 4). The idea of our algorithm is to start with an MST that has maximum degree 5 (which is known to exist) and then re-assign parts of the wedge angles from leaves to inner vertices. Our algorithm runs in linear time after computation of the MST. The ratios 2 and 1.5 improve the best known ratios for 120°-BST and 120°-MST (5 and 16/3, respectively). In fact, our ratios for non-uniform angles are better than any possible with uniform angles, as we prove by designing an infinite class of point sets such that every 120°-ST has maximum edge length at least 3 times that of the MST, and sum of edge lengths at least 2 times that of the MST.

(3) We present the following lower bounds for approximating the above problems with respect to the MST. These lower bounds are proved in Section 2. Although the lower bounds 2 and 3 (in Proposition 1) for the 120°-MST and the 120°-BST seem to be common knowledge [4, 22], for the sake of completeness we provide a proof of them.

Proposition 1. *The* $120°$*-MST, the* $\overline{120}^\circ$*-MST, and the* $120°$*-BST problems cannot be approximated by ratios smaller than* 2, 4/3, *and* 3, *respectively, given the MST length and the MST largest edge-length as lower bounds.*

Proposition 2. *Let* $A(n)$ *be the smallest value that suffices to construct, for any set of* n *points in the plane, a connected symmetric network with angle sum* $A(n)$ *and with transmission ranges bounded by a constant times the BST largest edge-length. Then* $A(n) = \Omega(\sqrt{n})$.

Proposition 3. *For any* $\alpha < \pi$ *there exists a point set for which no* α*-MST is an* α*-BST.*

1.2 Further Background

As mentioned above, there is a connection between α-STs and restricted degree spanning trees, due to the fact that d edges at a vertex always lie in some wedge of angle at most $2\pi(1 - 1/d)$. The Minimum degree-k spanning tree (degree-k MST) problem has been well-studied (see, e.g., [3, 13, 21, 23]). It is easy to compute a degree-2 spanning tree (a Hamiltonian path) of length at most twice the length of the Euclidean MST by doubling the MST edges, taking an Euler tour, short-cutting repeated vertices, and then removing an edge. It is also possible to compute in polynomial time degree-3, degree-4, and degree-5 spanning trees of lengths at most 1.402 [13], 1.1381 [13, 21], and 1 [25] times the MST length, respectively. This immediately implies the existence of 180°-ST, 240°-ST, 270°-ST, and 288°-ST ($\frac{8\pi}{5}$-ST) of lengths at most 2, 1.402, 1.1381, and 1 times the MST length, respectively. See Figure 1-right for an illustration of a degree-5 MST which is a 288°-ST.

The α-MST problem is also related to the problem of computing angle-restricted Hamiltonian paths and cycles on points in the plane. One can compute a Hamiltonian path with angles of at most 90° by starting from an arbitrary point and iteratively connecting the current point to its farthest among the remaining points (see [19]); the angle 90° is tight in the sense that there are point sets for which every Hamiltonian path has an angle larger than $90° - \varepsilon$ for any $\varepsilon > 0$ (see [12, 17]). Fekete and Woeginger [19] conjectured that for any even-size point set of at least 8 elements there exists a Hamiltonian cycle with angles at most 90°. Dumitrescu et al. [17] gave a partial solution by constructing a cycle with angles of at most 120°. The conjecture remains open.

Aschner and Katz [4] studied the α-MST problem for $\alpha \in \{90°, 120°, 180°\}$. They proved the NP-harness of this problem for $\alpha = 120°$ and $\alpha = 180°$ by reductions from the problem of finding a Hamiltonian path in hexagonal grid graphs and in square grid graphs of degree at most three, respectively. In addition to the result mentioned above for a 120°-ST of length at most 6 times the MST length, they also presented an algorithm for computing a 90°-ST of length at most 16 times the MST length.

The problem of constructing bounded-angle networks with no long edges (the α-BST problem) has been studied extensively. The NP-hardness reduction of Aschner and Katz for the 120°-MST problem also implies the NP-hardness of the 120°-BST problem and its inapproximability by a factor smaller than $\sqrt{3}$. Carmi et al. [12] construct 90°-Hamiltonian paths with edges that are shorter than that of the construction by Fekete and Woeginger [19]. Aschner and Katz [4] construct 120°-hop-spanners with hop-ratio 6 and edge lengths at most 7 for unit disk graphs. Dobrev et al. [15, 16] and Caragiannis et al. [11] construct strongly connected directed networks with short edges and nodes of bounded out-degree.

1.3 Notation and Preliminaries

Let w_p be a wedge in the plane with apex p. We denote by $\overrightarrow{w_p}$ the clockwise (right) boundary ray of w_p, and by $\overleftarrow{w_p}$ the counterclockwise (left) boundary ray of w_p. Let w_q be another convex wedge in the plane with apex q. If q lies in w_p then we say that p *sees* q and denote this by $p \rightarrow q$. We use $p \leftrightarrow q$ to denote that p and q are *mutually visible*, that is, p and q see each other. In other words, $p \leftrightarrow q$ denotes $p \rightarrow q$ and $q \rightarrow p$. In Figure 2 the point p sees q but q does not see p, and thus they are not mutually visible. Let P be a set of points in the plane and assume that wedges, possibly of different angles, are placed at every point of P. Then the *induced mutual-visibility graph* on P is a geometric graph with vertex set P that has a straight-line edge between two points p and q if and only if p and q are mutually visible. The underlying non-geometric graph is the *symmetric communication network* on P.

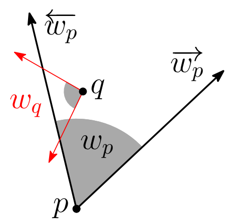

Figure 2: The point p sees q but q does not see p.

We denote the sum of edge lengths of a geometric graph G by $w(G)$. We need the following basic fact about triangles in the plane.

Lemma 1. *Let a, b, and c be three points in the plane, and let $E = \{ab, ac, bc\}$ be the set of edges between them. Then the total length of the shortest and longest edges in E is at most 1.5 times the total length of any two edges of E.*

Proof. After a suitable relabeling we may assume that $|ab| \leqslant |ac| \leqslant |bc|$. To prove the lemma it suffices to show that (i) $|ab| + |bc| \leqslant 1.5(|ac| + |bc|)$ and (ii) $|ab| + |bc| \leqslant 1.5(|ab| + |ac|)$. Statement (i) holds because $|ab| \leqslant |ac|$. To verify statement (ii) observe that by the triangle inequality we have $|bc| \leqslant (|ab| + |ac|)$, and by our assumption that $|ab|$ is not larger than $|ac|$ we have $|ab| \leqslant 0.5(|ab| + |ac|)$. The sum of the two inequalities implies statement (ii). ☐

2 Lower bounds

In this section we prove Propositions 1, 2, and 3. First we prove Proposition 1 that is: *The* $120°$*-MST, the* $\overline{120}^{\circ}$*-MST, and the* $120°$*-BST problems cannot be approximated by ratios smaller than* 2, 4/3*, and* 3*, respectively, given the MST length and the MST largest edge-length as lower bounds.*

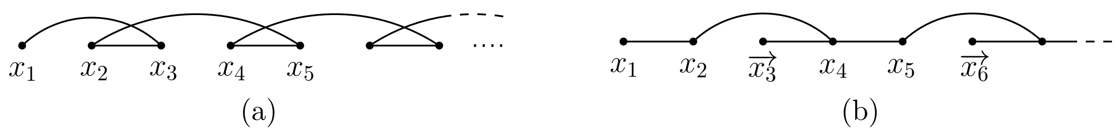

Figure 3: (a) A $120°$-ST of length $2n - 3$, and (b) a $\overline{120}^{\circ}$-ST of length $4n/3 - 3$.

Proof. Consider a sequence $X = (x_1, x_2, \ldots, x_n)$ of n points on the x-axis with coordinates $1, 2, \ldots, n$, respectively, as in Figure 3. We show that X satisfies the statement of the proposition for the three problems. The MST of X is a path with edges of length 1 and total length $n - 1$.

First we prove the lower bound 2 for the $120°$-MST problem (this lower bound is also mentioned in [4]). We show by induction that any $120°$-minimum spanning tree T on X has length at least $2n - 3$. This is trivial if $n \in \{1, 2\}$, thus assume that $n \geqslant 3$. There are $n - 1$ intervals on the x-axis between consecutive points of X. If every interval is covered by two edges of T then $w(T) \geqslant 2n - 2$. Assume that some interval $[i, i + 1]$ is covered by only one edge of T, say edge e. By removing e we obtain two subtrees T_1 and T_2 where T_1 spans the points $x_1, \ldots, x_i$ while T_2 spans $x_{i+1}, \ldots, x_n$. Let n_1 and n_2 be the number vertices of T_1 and T_2, respectively.

Assume that both n_1 and n_2 are at least 2. Since T_1 spans $x_1, \ldots, x_i$, the vertex x_i is oriented to the left (and does not see any point to its right). Analogously, x_{i+1} is oriented to the right. Therefore e is not incident to x_i or x_{i+1}, and thus its length is at least 3. By the induction hypothesis we get $w(T) = w(T_1)+w(T_2)+w(e) \geqslant (2n_1-3)+(2n_2-3)+3 = 2(n_1+n_2)-3 = 2n-3$. Now assume that $n_1 = 1$, and thus $n_2 \geqslant 2$. Then T_1 has only vertex $x_i = x_1$ which is an endpoint of e. Since T_2 spans $x_2, \ldots, x_n$, the vertex x_2 is oriented to the right. Therefore e is not incident to x_2, and thus its length is at least 2. Thus $w(T) = w(T_1) + w(T_2) + w(e) \geqslant 0 + (2n_2 - 3) + 2 = 2n - 3$. The case where $n_2 = 1$ is handled similarly.

To verify the lower bound 3 for the 120°-BST problem, assume that $n = 5$. Consider the edge-maximal graph on X that has edges of length at most 2. In any spanning tree in this graph, at least one of x_2, x_3, x_4 has incident edges in both directions (left and right). Thus, no matter how we place wedges of angle 120° on vertices of X, we cannot get a 120°-ST of edge lengths at most 2. Therefore, any 120°-ST on X has an edge of length at least 3, as in Figure 3(a). This argument can be generalized for any n larger than 5.

Now consider any $\overline{120}^{\circ}$-minimum spanning tree $\overline{T}$ on X. We show that $w(\overline{T}) \geqslant 4n/3 - 3$. Partition the vertices of $\overline{T}$ into X_1 and X_2 where X_1 is the set of vertices with wedges of angle strictly less than 180° and X_2 is the set of vertices with wedges of angle at least 180°. Since the total available angle is $120n$ degrees, $|X_2| \leqslant 120n/180 = 2n/3$. Thus $|X_1| = n - |X_2| \geqslant n/3$. Observe that every interval (between consecutive vertices of X) is covered by an edge of $\overline{T}$. Every vertex $x_i \in X_1$ sees the vertices that are either to its left or to its right. We denote x_i by $\overleftarrow{x_i}$ if it sees the vertices to its left, and by $\overrightarrow{x_i}$ otherwise (see Figure 3(b)). For every $\overleftarrow{x_i}$ the interval $[i-1, i]$ is covered by at least two edges otherwise connectivity is lost: one edge is incident to $\overleftarrow{x_i}$ and another edge connects a point to the left of $\overleftarrow{x_i}$ with a point to the right (assuming $i \neq n$). Similarly for every $\overrightarrow{x_i}$ the interval $[i, i+1]$ is covered by an edge that is incident to $\overrightarrow{x_i}$ and by an edge that connects a point to the right of $\overrightarrow{x_i}$ with a point to the left (assuming $i \neq 1$), as in Figure 3(b). Thus, for every $\overleftarrow{x_i}$ (except possibly $\overleftarrow{x_n}$) there exists a unique interval that is covered by two edges of $\overline{T}$. Similarly, for every $\overrightarrow{x_i}$ (except possibly $\overrightarrow{x_1}$) there exists a unique interval that is covered by two edges of $\overline{T}$. (If x_i is oriented to the left and x_{i+1} is oriented to the right then—by the minimality of the tree—(x_i, x_{i+1}) is an edge of $\overline{T}$ and the interval $[i, i+1]$ is covered by three edges.) Therefore the length of $\overline{T}$ is at least $(n-1) + (|X_1| - 2) \geqslant 4n/3 - 3$. □

Now we prove Proposition 2 that is: *Let $A(n)$ be the smallest value suffices to construct, for any set of n points in the plane, a connected symmetric network with angle sum $A(n)$ and with transmission ranges bounded by a constant times the BST largest edge-length. Then $A(n) = \Omega(\sqrt{n})$.*

Proof. Consider a set of n points on the vertices of a regular $\sqrt{n} \times \sqrt{n}$ grid of side length $\sqrt{n}-1$. The largest edge-length of any BST for this point set is 1. Now consider any connected symmetric network T for this point set that satisfies the statement of the proposition, and let the constant c be the largest transmission range (edge-length) in T. Let L denote the set of supporting lines for all edges of T. Since any line in the plane contains at most $\sqrt{n}$ points of the grid, L has at least $\sqrt{n}$ lines. Therefore, T has at least one vertex, say v, where the neighbors of v lie on more than one line of L. Let V be the set of all such vertices of T. Observe that every line in L passes through a vertex in V. For every $v \in V$ the neighbors of v lie in a square of side-length $2c$ that is centered at v. Thus v has at most $(2c)^2$ neighbors (that can lie on at most $(2c)^2$ lines in L). Therefore $|V| \geqslant |L|/(2c)^2 \geqslant \sqrt{n}/(2c)^2$.

Now consider any $v \in V$. Let v_1 and v_2 be two neighbors of v such that edges $(v, v_1) \in T$ and $(v, v_2) \in T$ lie on two different lines of L. Since v, v_1, and v_2 are grid points and lie on a square of constant side-length $2c$, the angle between v_1 and v_2 at v is a constant number. Let

α be the smallest such constant over all vertices in V. Then, the total angle at vertices in V is at least $\alpha \cdot |V| \geqslant \alpha\sqrt{n}/(2c)^2 = \Omega(\sqrt{n})$. □

The following is a proof of Proposition 3 that is: *For any $\alpha < \pi$ there exists a point set for which no α-MST is an α-BST.*

Proof. Consider the point set P in Figure 4 consisting of $n \geqslant 9$ points partitioned into $t = n/3$ triplets (a_i, b_i, c_i) where $|c_i a_{i+1}| = 1$ and $|a_i b_i| = |b_i c_i| = \epsilon$ for some $\epsilon < 1/(2n-3)$. We refer to an edge of length at least 1 as a *long edge*. The red bold tree is an α-ST of total length $(t-1)\left((1+3\epsilon)+2\epsilon+\epsilon\right)+2\epsilon+\epsilon = (t-1)+(6t-3)\epsilon < t$, where the last inequality holds by our choice of ϵ. Therefore any α-MST for P has at most $t-1$ long edges because otherwise it would have a length larger than t. This implies that each interval $[c_i, a_{i+1}]$ is covered by at most 1 edge of any α-MST. On the other hand, any spanning tree for P has at least $t-1$ long edges as all triplet should be connected to the rest of the tree by a long tree edge. It turns out that any α-MST T of P has exactly $t-1$ long edges each connecting a point in triplet (a_i, b_i, c_i) to a point in triplet $(a_{i+1}, b_{i+1}, c_{i+1})$. In the rest of the proof we show that T has an edge of length at least $1+3\epsilon$. This immediately proves our claim because P admits an α-ST of edge lengths at most $1+2\epsilon$, see for example the thin blue tree in Figure 4.

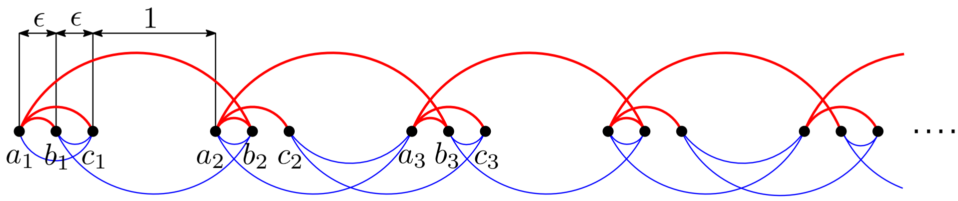

Figure 4: The red bold spanning tree has a total length of less than t. The blue thin spanning tree has edges of lengths at most $1+2\epsilon$.

Consider any $i \in \{1, \ldots, t-1\}$. Let e_i be the long edge of T between triplets (a_i, b_i, c_i) and $(a_{i+1}, b_{i+1}, c_{i+1})$. The point c_i cannot be an endpoint of e_i because otherwise there must be a long edge e'_i in T that connects a point to the left of c_i to a point to the right of c_i; this contradicts the fact that the interval $[c_i, a_{i+1}]$ is covered by exactly one edge of T (by the degree constraint c_i cannot be connected to any point to the left). By symmetry, a_{i+1} cannot be an endpoint of e_i either. If a_i or c_{i+1} is an endpoint of e_i then $|e_i| \geqslant 1+3\epsilon$ and we are done. Assume that for every $i \in \{1, \ldots, t-1\}$ we have $e_i = (b_i, b_{i+1})$. Then (b_1, b_2) and (b_2, b_3) are edge of T. Thus the angle at b_2 is larger than α which contradicts T being an α-ST. □

3 Approximating the 120°-MST

Let $\alpha = 120°$ in this section. Let P be a set of points in the plane. Aschner and Katz [4] showed a construction of an α-ST on P of length at most 6 times the MST length. In this section we present an alternate construction that achieves an α-ST of length at most 16/3 times the MST length, thereby proving the following theorem.

Theorem 1. *Given a set of points in the plane and an angle $\alpha \geqslant 120°$, there is an α-spanning tree of length at most 16/3 times the length of the MST. Furthermore, there is a polynomial time algorithm to find such an α-ST, thus providing a 16/3-approximation algorithm for the α-MST problem.*

3.1 The construction of Aschner and Katz

To facilitate comparisons, we briefly describe the algorithm of Aschner and Katz [4]. They use the following theorems to compute an α-ST.

Theorem 2 (Aschner and Katz, 2014)**.** *Given a set P of three points in the plane, one can place at each point of P a wedge of angle $120°$ such that the three wedges cover the plane and the induced mutual-visibility graph on P is connected, and hence it contains an $120°$-ST.*

A placement that satisfies the conditions of Theorem 2 is given in the proof of Claim 2.1 [4]. This placement has an interesting property which is given in the following theorem whose proof is very involved.

Theorem 3 (Aschner and Katz, 2014)**.** *Let P_1 and P_2 be two disjoint sets each containing three points in the plane. Assume that a wedge of angle $120°$ is placed at each point of P_1 and at each point of P_2 according to the placement of Theorem 2. Then, the induced mutual-visibility graph on $P_1 \cup P_2$ is connected, and hence it contains a $120°$-ST.*

Let H be a Hamiltonian path on P of length at most 2 times the MST length. The constant 2 is tight, as Fekete et al. [18] showed that for any fixed $\varepsilon > 0$ there exists a point set for which any Hamiltonian path has length at least $2 - \varepsilon$ times the MST length.

Let $(p_0, \ldots, p_{n-1})$ be the sequence of points of P from one endpoint of H to the other. Let $(h_0, h_1, \ldots, h_{n-2})$ be the sequence of the edges of H where $h_i = (p_i, p_{i+1})$. Partition the edges of H into three sets $H_0 = \{h_0, h_3, h_6, \ldots\}$, $H_1 = \{h_1, h_4, h_7, \ldots\}$, and $H_2 = \{h_2, h_5, h_8, \ldots\}$, as in Figure 5. The length of one of these sets, say H_2, is at least $w(H)/3$. Therefore $w(H_0)+w(H_1) \leqslant \frac{2}{3}w(H)$. Partition P into a sequence of triplets $(p_0, p_1, p_2), (p_3, p_4, p_5), \ldots$ such that the edges of H, that lie between consecutive triplets, are in H_2. Then place three wedges on the points of each triplet according to Theorem 2, and let G_α be the induced mutual-visibility graph. The points in each triplet are connected (by Theorem 2) and any pair of triplets are connected (by Theorem 3), and thus G_α is connected.

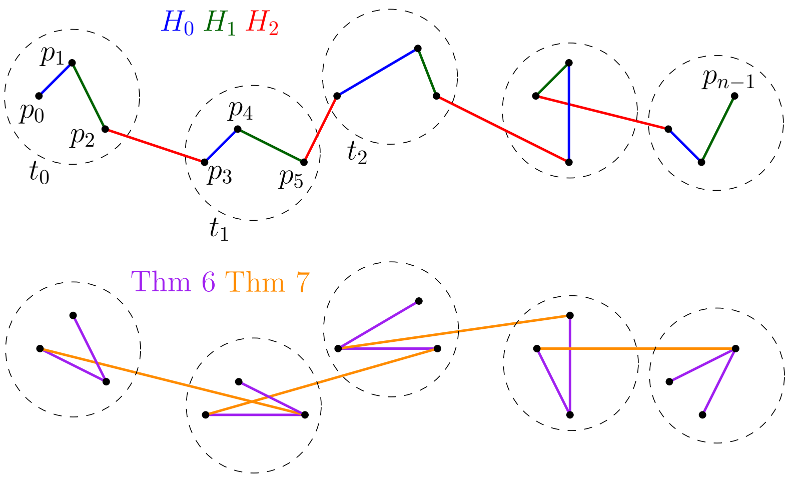

Figure 5: Top: path H where H_0, H_1, H_2 are colored blue, green, and red, respectively. Bottom: spanning tree T where edges obtained by Theorems 2 and 3 are colored purple and orange, respectively.

Now let T be a spanning tree of G_α computed as follows (see Figure 5): between the points in each triplet take two edges that are obtained from Theorem 2 (these edges are called *inner edges*), and between every two consecutive triplets take an edge that is obtained by Theorem 3 (these edges are called *connecting edges*). To bound the length of T, we charge edges of H for the edges of T. Every edge of H that belongs to H_2 lies between two consecutive triplets, and thus is charged only once for the connecting edge between the two triplets. Every edge of H that belongs to $H_0 \cup H_1$ lies inside a triplet, say t. Such an edge is charged 4 times: twice for the two inner edges of t and twice for the two edges that connect t to adjacent triplets. Therefore,

$$w(T) \leqslant w(H_2) + 4(w(H_0) + w(H_1)) = w(H) + 3(w(H_0) + w(H_1)) \leqslant 3w(H) \leqslant 6w(MST).$$

3.2 The new construction

Our approach, for the construction of an α-ST of length at most 16/3 times the MST length, is similar to that of Aschner and Katz, however we use a different placement of wedges on points of triplets which in turn requires different proof techniques.

Let $P = \{p_1, p_2, p_3\}$ be any set of three points in the plane. The complete graph on P has edge set $E = \{p_1p_2, p_1p_3, p_2p_3\}$ and contains exactly three spanning trees. Let T_1 be the tree that has the two shortest edges of E, T_3 be the tree that has the two longest edges of E, and T_2 be the tree that has the shortest and the longest edges of E. Observe that $w(T_1) \leqslant w(T_2) \leqslant w(T_3)$. We refer to T_1, T_2, and T_3 as the *shortest*, *intermediate*, and *longest* trees on P, respectively.

The α-ST obtained by the wedge placement in (the proof of) Theorem 2 contains the two longest edges of E. In other words, this placement gives the spanning tree T_3. Due to the objective of minimizing the sum of edge lengths, one may ask for a wedge placement that covers the plane using T_1 or T_2. We answer this question in the affirmative in Theorem 4. This theorem improves Theorem 2 in the sense that it uses T_1 or T_2 and also allows more flexibility on the angles of the wedges.

Consider three points in the plane, and three angles $\alpha_1, \alpha_2, \alpha_3$ where each α_i is at most π and $\alpha_1 + \alpha_2 + \alpha_3 = 2\pi$. The Floodlight result of Bose et al. [9] implies that one can cover the plane by placing three wedges of angles $\alpha_1, \alpha_2, \alpha_3$ at the three points. Our Theorem 4 also improves this result in two ways: (i) one can cover the plane with three wedges even if the assignment of angles to points is specified in advance, and (ii) the induced mutual-visibility graph is connected, in particular it contains the shortest or the intermediate tree on points.

Theorem 4 (proved in Section 3.3)**.** *Given a set $P = \{p_1, p_2, p_3\}$ of three points in the plane and three angles $\alpha_1, \alpha_2, \alpha_3 \leqslant \pi$ where $\alpha_1 + \alpha_2 + \alpha_3 = 2\pi$, one can place at each p_i a wedge of angle α_i such that the three wedges cover the plane and the induced mutual-visibility graph contains the shortest or the intermediate tree on P.*

If we set each α_i equal to $120°$ in Theorem 4, then we get the following corollary.

Corollary 1. *Given a set P of three points in the plane, one can place wedges of angle $120°$ at points of P such that the three wedges cover the plane and the induced mutual-visibility graph contains the shortest or the intermediate tree on P.*

Corollary 1 ensures covering of the plane and also uses T_1 or T_2. Unfortunately, we cannot directly use this corollary in the approach of Aschner and Katz because their Theorem 3 ensures a connection between two triplets t_1 and t_2 only if they are oriented by Theorem 2; such a connection may not exist if the triplets are oriented by Corollary 1. However, we show in Theorem 5 that if the relative positions of the

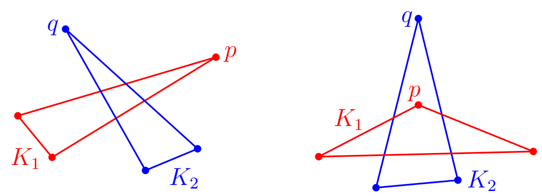

Figure 6: Two configurations of a forbidden pair $\{P_1, P_2\}$ forming triangles K_1 and K_2.

points of t_1 and t_2 is not among the two configurations shown in Figure 6, then the orientation of Corollary 1 ensures a connection between them. Our final proof relies on the fact that the configurations in Figure 6 never arise from triplets of a non-crossing Hamiltonian path.

To be more precise, consider disjoint point sets P_1 and P_2, each of size three, in the plane. Let K_1 and K_2 be the complete graphs (triangles) on points of P_1 and P_2, respectively. We say that the (unordered) pair $\{P_1, P_2\}$ is *forbidden* if there exist vertices $p \in K_1$ and $q \in K_2$ such that (i) both incident edges of p intersect both incident edges of q, or (ii) p lies inside the triangle K_2 and the non-incident edge of p intersects both incident edges of q. See Figure 6.

Theorem 5 (proved in Section 3.4)**.** *Let P_1 and P_2 be two disjoint sets, each containing three points in the plane, such that $\{P_1, P_2\}$ is not forbidden. Assume that a wedge of angle $120°$ is placed at each point of P_1 and at each point of P_2 according to the placement algorithm of Corollary 1. Then, the induced mutual-visibility graph on $P_1 \cup P_2$ is connected, and hence it contains a $120°$-ST.*

Remark 1. Intuitively, it would seem that our proof of Theorem 5 should follow the same approach as that of Aschner and Katz's Theorem 3. We note, however, that their proof of Theorem 3 is highly involved and uses a combination of nontrivial ideas. In order to cope with the large number of cases, they classify the number of edges in the induced mutual-visibility graph of each triplet. Their proof ensures the existence of an edge between P_1 and P_2 only if the wedges are oriented according to the placement in the proof of Theorem 2. Such an edge may not exist for other wedge placements with similar properties, i.e., coverage of the plane and connectivity of each set P_1 and P_2; for example see Figure 7 which is borrowed from [4] (notice that the pair of triplets in this figure is forbidden). Thus, there is no straightforward way of adjusting their proof to work for our Theorem 5 where the wedges in P_1 and P_2 are oriented according to Corollary 1 instead. We give a relatively short proof of Theorem 5 in Appendix 3.3. Instead of classifying the number of edges of visibility graphs, we introduce a "representative" point for each triplet; this facilitates a shorter presentation of our proof.

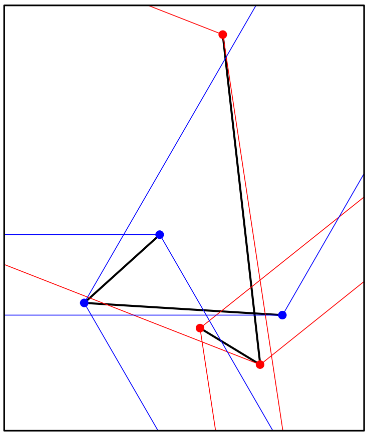

Figure 7: There is no connection between the red and blue triplets.

The remaining ingredient we need in order to apply Theorem 5 is the following observation that forbidden triplets do not arise if we begin with a Hamiltonian path H that is non-crossing.

Observation 1. *Let H be a non-crossing Hamiltonian path and let t_1 and t_2 be disjoint triplets, each of which is obtained by consecutive vertices of H. Then, $\{t_1, t_2\}$ is not forbidden.*

Proof. Each of t_1 and t_2 must contain two edges of H (because they come from consecutive vertices of H), and the union of these edges is non-crossing because H is non-crossing. But if $\{t_1, t_2\}$ is forbidden (see Figure 6), then there is no way to choose two edges in each of t_1 and t_2 such that their union is non-crossing. □

To use this observation, we need a non-crossing Hamiltonian path. Such a path can be obtained by iteratively flipping crossing edges of H. It is known that this iterative process terminates after $O(n^3)$ edge flips [26], where n is the number of path vertices. (See [7] for some recent results on obtaining non-crossing configurations by edge flips.) Since the edge flip operation does not increase the total edge length (this can be verified by the triangle inequality), the length of the resulting non-crossing path is not more than that of the original path. Therefore, we assume from now on that H is non-crossing and $w(H) \leqslant 2w(MST)$.

Now we have the necessary tools (a non-crossing path H and Theorem 5) to use the orientation of the wedges as in Corollary 1. Let $(p_0, \dots, p_{n-1})$ be the sequence of points from one endpoint of H to the other. Construct edge sets H_0, H_1, H_2, and as before assume that $w(H_2) \geqslant w(H)/3$, which implies $w(H_0) + w(H_1) \leqslant \frac{2}{3}w(H)$. Construct triplets $t_0 = (p_0, p_1, p_2), t_1 = (p_3, p_4, p_5), \dots$, and orient the wedges according to Corollary 1. Since H is non-crossing, no pair of triplets is forbidden (by Observation 1). The induced mutual-visibility graph, G_α, is connected because the points of each triplet are connected (by Corollary 1) and every two triplets are connected (by Theorem 5). We obtain a spanning tree T from G_α as before. Every edge of H_2 is charged only once for the connecting edge between two consecutive triplets. Every pair (h_i, h_{i+1}) of edges in each triplet t (where $h_i \in H_0$ and $h_{i+1} \in H_1$) is charged 3.5 times: twice for connecting edges between t and its adjacent triplets, and 1.5 times for inner edges of t (by Lemma 1 the length of the tree obtained by Corollary 1 is at most $1.5(w(h_i) + w(h_{i+1}))$). Therefore

$$\begin{aligned} w(T) &\leqslant w(H_2) + 3.5(w(H_0) + w(H_1)) = w(H) + 2.5(w(H_0) + w(H_1)) \\ &\leqslant w(H) + \tfrac{5}{3}w(H) = \tfrac{8}{3}w(H) \leqslant \tfrac{16}{3}w(MST). \end{aligned}$$

This finishes the proof of Theorem 1.

3.3 Proof of Theorem 4

In this section we prove Theorem 4: *Given a set $P = \{p_1, p_2, p_3\}$ of three points in the plane and three angles $\alpha_1, \alpha_2, \alpha_3 \leqslant \pi$ where $\alpha_1 + \alpha_2 + \alpha_3 = 2\pi$, one can place at each p_i a wedge of angle α_i such that the three wedges cover the plane and the induced mutual-visibility graph contains the shortest or the intermediate tree on P.*

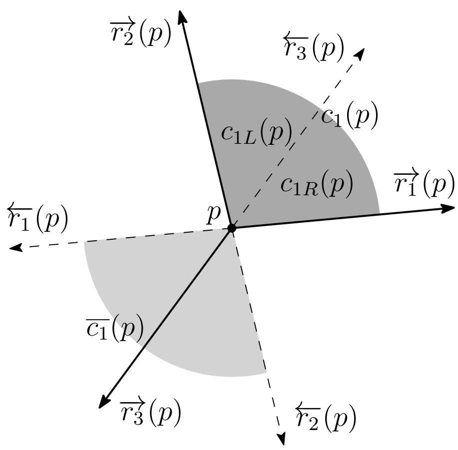

First we provide some preliminaries for the proof. We use similar notation also for the proof of Theorem 5 in Section 3.4. Let p be a point in the plane and let $[\overrightarrow{r_1}(p), \overrightarrow{r_2}(p), \overrightarrow{r_3}(p)]$ be the cyclic counterclockwise permutation of three rays emanating from p where the angle between any two consecutive rays is less than π. See the figure to the right for an illustration. These rays partition the plane into three cones with apices at p. For an index $i \in \{1, 2, 3\}$ we denote by $\overrightarrow{r_{i+1}}(p)$ the ray after $\overrightarrow{r_i}(p)$ in the cyclic permutation, by $\overleftarrow{r_i}(p)$ the ray emanating from p in the opposite direction of $\overrightarrow{r_i}(p)$, and by $l_i(p)$ the line through $\overrightarrow{r_i}(p)$. We denote by $c_i(p)$ the convex cone with boundary rays $\overrightarrow{r_i}(p)$ and $\overrightarrow{r_{i+1}}(p)$, and by $\overline{c_i}(p)$ the reflection of $c_i(p)$ with respect to p. Moreover, we denote by $c_{iR}(p)$ the portion of $c_i(p)$ that is between $\overrightarrow{r_i}(p)$ and $\overleftarrow{r_{i+2}}(p)$, and by $c_{iL}(p)$ the of portion $c_i(p)$ of that is between $\overleftarrow{r_{i+2}}(p)$ and $\overrightarrow{r_{i+1}}(p)$.

Now we proceed with the proof. Consider the triangle with vertices p_1, p_2, and p_3. Let $\beta_1, \beta_2, \beta_3$ denote the interior angles of this triangle at p_1, p_2, p_3 respectively, and note that $\beta_1 + \beta_2 + \beta_3 = \pi$. Without loss of generality assume that $\beta_1 \leqslant \beta_2 \leqslant \beta_3$, and thus $|p_2p_3| \leqslant |p_1p_3| \leqslant |p_1p_2|$. After a suitable reflection assume that p_3 appears to the left of the ray from p_1 towards p_2. We consider two cases where $\alpha_3 \geqslant \beta_3$ and $\alpha_3 < \beta_3$. For each case we show how

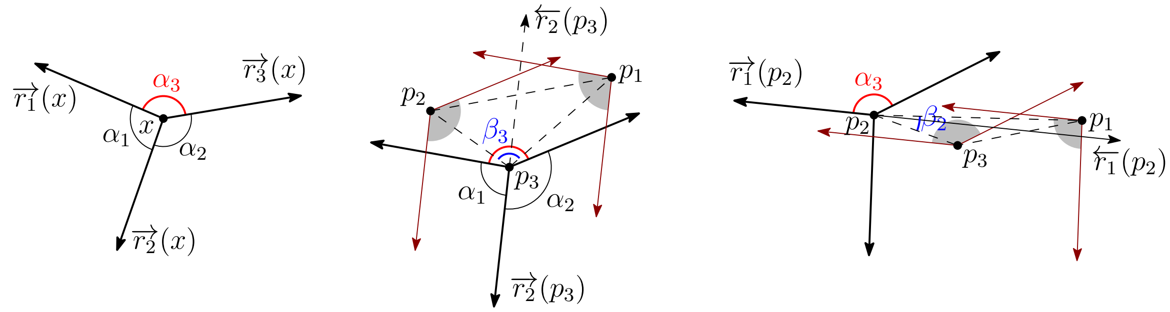

(a) The embedding $R(x)$ (b) $\alpha_3 \geqslant \beta_3$ (c) $\alpha_3 < \beta_3$ and $\alpha_2 \geqslant \beta_2$

Figure 8: Illustration of the proof of Theorem 4.

to place three wedges w_1, w_2, w_3 (with angles $\alpha_1, \alpha_2, \alpha_3$) at points p_1, p_2, p_3, respectively, to satisfy the conditions of the theorem.

Let $[\overrightarrow{r_1}(x), \overrightarrow{r_2}(x), \overrightarrow{r_3}(x)]$ be the cyclic counterclockwise permutation of three rays emanating from some point x in the plane such that the angle at each cone $c_i(x)$ is α_i for $i \in \{1, 2, 3\}$. Let $R(x)$ denote a fixed embedding of these rays in the plane, as in Figure 8(a).

- $\alpha_3 \geqslant \beta_3$. Translate $R(x)$ so that x lies on p_3, i.e., $x = p_3$. Rotate $R(x)$ so that p_1 and p_2 lie in $c_3(p_3)$ and $\overleftarrow{r_2}(p_3)$ intersects the segment p_1p_2 (observe the existence of such rotation). Place w_1, w_2, and w_3 at cones $c_1(p_3)$, $c_2(p_3)$, and $c_3(p_3)$, respectively. The union of these wedges covers the entire pane. Now translate w_1 and w_2 so that their apices lie on p_1 and p_2, respectively, as in Figure 8(b). The union of the three wedges still covers the plane. In this setting, p_3 and p_1 are mutually visible and so are p_3 and p_2. Thus, the mutual-visibility graph contains the shortest tree on P with edges $T_1 = \{p_3 \leftrightarrow p_1, p_3 \leftrightarrow p_2\}$.

- $\alpha_3 < \beta_3$. We claim that $\alpha_2 \geqslant \beta_2$. For the sake of contradiction assume that $\alpha_2 < \beta_2$. Then $\alpha_1 + \beta_2 + \beta_3 > \alpha_1 + \alpha_2 + \alpha_3 = 2\pi$. Since β_2 and β_3 are interior angles of a triangle, $\beta_2 + \beta_3 \leqslant \pi$. By combining these two inequalities we get $\alpha_1 > \pi$, which contradicts an assumption in the statement of the theorem. Thus our claim follows.

  Translate $R(x)$ so that x lies on p_2. Rotate $R(x)$ so that p_1 and p_3 lie in $c_2(p_2)$ and $\overleftarrow{r_1}(p_2)$ intersects the segment p_1p_3. Place w_1, w_2,w_3 at cones $c_1(p_2)$, $c_2(p_2)$, $c_3(p_2)$. Translate w_1 and w_3 so that their apices lie on p_1 and p_3, respectively, as in Figure 8(c). Again the three wedges cover the entire plane, and the mutual-visibility graph contains the intermediate tree on P with edges $T_2 = \{p_2 \leftrightarrow p_1, p_2 \leftrightarrow p_3\}$.

3.4 Proof of Theorem 5

In this section we prove Theorem 5: *Let P_1 and P_2 be two disjoint sets, each containing three points in the plane, such that $\{P_1, P_2\}$ is not forbidden. Assume that a wedge of angle $120°$ is placed at each point of P_1 and at each point of P_2 according to the placement algorithm of Corollary 1. Then, the induced mutual-visibility graph on $P_1 \cup P_2$ is connected, and hence it contains a $120°$-ST.*

We use the notation introduced in Section 3.3. Let $P_1 = \{a, b, c\}$ and $P_2 = \{a', b', c'\}$. Orient P_1 and P_2 according to Corollary 1 (in fact according to the proof of Theorem 4 with angles $120°$). Let w_a, w_b, w_c, $w_{a'}$, $w_{b'}$ and $w_{c'}$ be the wedges of angle $120°$ that are place at these points, respectively. Recall, from the proof of Theorem 4, three rays $[\overrightarrow{r_1}(x), \overrightarrow{r_2}(x), \overrightarrow{r_3}(x)]$ with cones of angles $120°$ that are placed at a point x. In the orientation of P_1 we may assume that these rays are placed at a, i.e., $x = a$. Thus b and c lie in the same cone $c_i(a)$ for some

$i \in \{1,2,3\}$; in particular one of them lies in $c_{iL}(a)$ and the other lies in $c_{iR}(a)$. Moreover w_a covers $c_i(a)$ and we have $a \leftrightarrow b$ and $a \leftrightarrow c$. In the orientation of P_2 assume that these rays are placed at a'. Thus one of b' and c' lies in $c_{iL}(a')$ and the other lies in $c_{iR}(a')$ for some $i \in \{1,2,3\}$, $w_{a'}$ covers $c_i(a')$, and we have $a' \leftrightarrow b'$ and $a' \leftrightarrow c'$. Also recall that each cone $\overline{c_i}(a)$ contains a point of P_1 whose wedge covers $c_i(a)$, and similarly each cone $\overline{c_i}(a')$ contains a point of P_2 whose wedge covers $c_i(a')$.

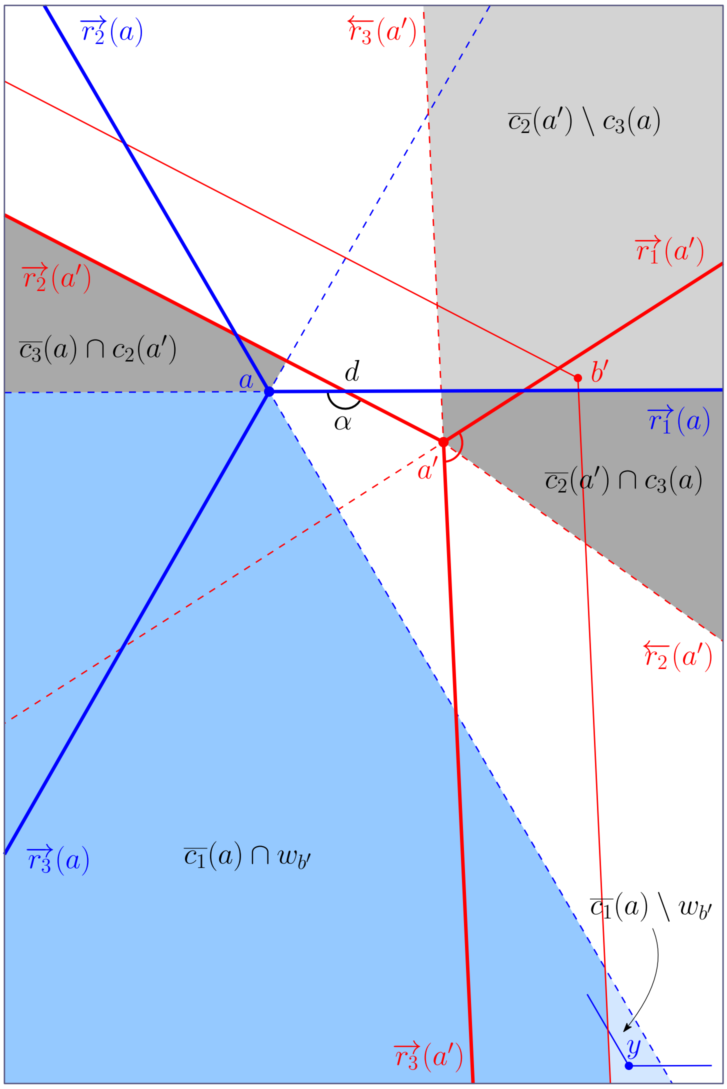

Figure 9: $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ intersect, $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ do not intersect, and $b' \in c_{3L}(a') \cap c_1(a)$.

Since the three cones $c_1(a)$, $c_2(a)$, and $c_3(a)$ cover the plane, a' lies in one of them, say $c_3(a)$. Similarly, a lies in one of the three cones at a', say $c_2(a')$. In this setting one of the following three configurations holds:

(A) $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ intersect, but $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ do not intersect. See Figures 9 and 10.

(B) $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ do not intersect, but $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ intersect.

(C) $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ intersect, and $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ intersect. See Figures 11 and 12.

We consider each configuration separately. Since configurations (A) and (B) are symmetric, we describe only (A) and (C).

Configuration (A). Let d be the intersection point of $\overrightarrow{r_1}(a)$ with $\overrightarrow{r_2}(a')$, and let α denote the convex angle $\angle ada'$, as in Figure 9. Since $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ do not intersect, $\alpha \geqslant 2\pi/3$, and consequently a' lies in $c_{3L}(a)$ and a lies in $c_{2R}(a')$. Recall that each cone at a (resp. a') is covered by a point of P_1 (resp. P_2). Let $x \in P_1$ be the point in $\overline{c_3}(a)$ whose wedge w_x covers $c_3(a)$, and let $x' \in P_2$ be the point in $\overline{c_2}(a')$ whose wedge $w_{x'}$ covers $c_2(a')$. If $x \in \overline{c_3}(a) \cap c_2(a')$ and $x' \in \overline{c_2}(a') \cap c_3(a)$ (the dark-gray regions in Figure 9) then $x \leftrightarrow x'$ and we are done. Therefore we may assume that $x \in \overline{c_3}(a) \setminus c_2(a')$ or $x' \in \overline{c_2}(a') \setminus c_3(a)$. By symmetry we assume that $x' \in \overline{c_2}(a') \setminus c_3(a)$ (the light-gray region in Figure 9). This and the fact that $a' \in c_3(a)$ imply that $x' \neq a'$, and thus $x' \in \{b', c'\}$. After a suitable relabeling we assume that $x' = b'$. Therefore, $w_{b'}$ covers $c_2(a')$.

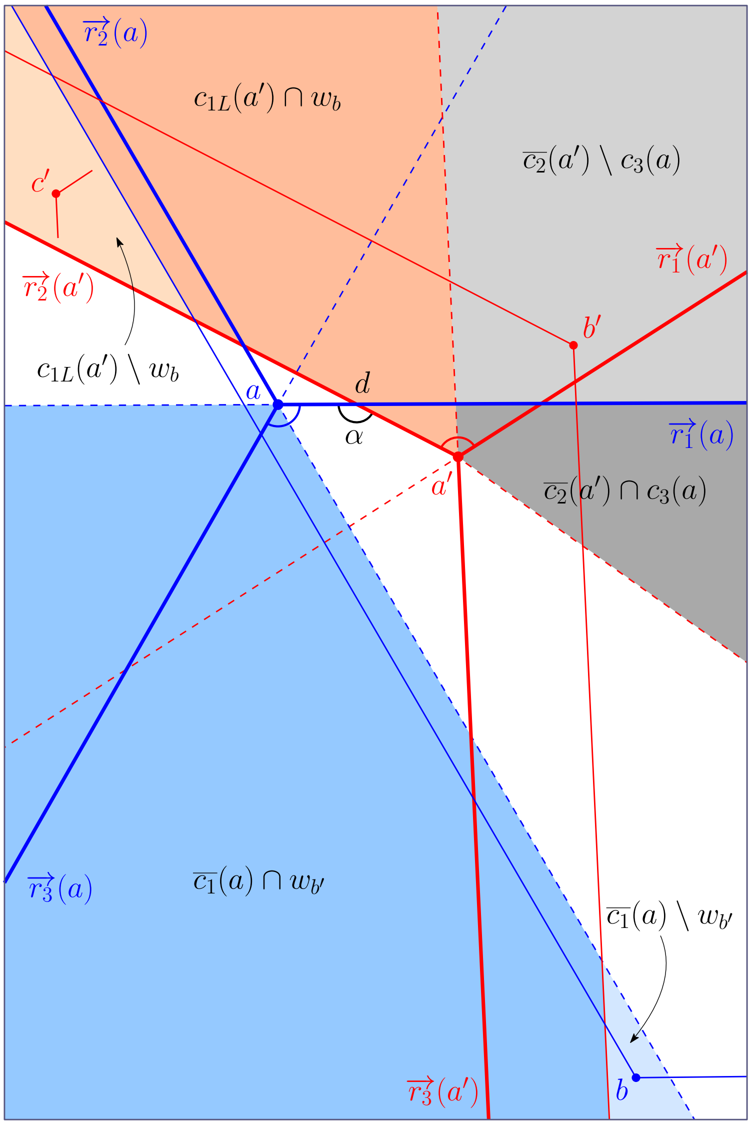

Figure 10: $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ intersect, $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ do not intersect, and $b' \in c_{1R}(a') \cap c_1(a)$.

Since $\alpha \geqslant 2\pi/3$ and $a' \in c_{3L}(a)$, the rays $\overleftarrow{r_3}(a')$ and $\overleftarrow{r_2}(a')$ do not intersect the line $l_2(a)$, and consequently the cone $\overline{c_2}(a')$ does not intersect $l_2(a)$ and hence this cone is disjoint from $c_2(a)$. Therefore, b'—which is in $\overline{c_2}(a') \setminus c_3(a)$—lies in $c_1(a)$. Let $y \in P_1$ be the point in $\overline{c_1}(a)$ whose wedge w_y covers $c_1(a)$. Since $b' \in c_1(a)$, the point y sees b', as in Figure 9. If $y \in \overline{c_1}(a) \cap w_{b'}$ (the dark-blue region) then b' sees y, and hence $y \leftrightarrow b'$ and we are done. Assume that $y \in \overline{c_1}(a) \setminus w_{b'}$, and hence it is not in $c_2(a')$. Since $y \in \overline{c_1}(a)$ and $\overline{c_1}(a)$ is disjoint from $c_1(a')$, the point y is not in $c_1(a')$ either. Thus y lies in $c_3(a')$. In this setting, w_y covers the region $c_{3L}(a) \cap w_{b'}$ which contains a'. Thus, y sees a'. Recall that $b' \in \overline{c_2}(a') \cap c_1(a)$. Since $\overline{c_2}(a') = c_{3L}(a') \cup c_{1R}(a')$ we consider two cases where $b' \in c_{3L}(a') \cap c_1(a)$ and $b' \in c_{1R}(a') \cap c_1(a)$.

- $b' \in c_{3L}(a') \cap c_1(a)$. This case is depicted in Figure 9. Since $b' \in c_3(a')$ and $w_{a'}$ covers the cone that contains b' (and also c'), it covers $c_3(a')$. Therefore, a' sees y which lies in $c_3(a')$, and hence $y \leftrightarrow a'$ and we are done.

- $b' \in c_{1R}(a') \cap c_1(a)$. This case is depicted in Figure 10. Since $b' \in c_1(a')$ and $w_{a'}$ covers the cone that contains b', it covers $c_1(a')$. Since $y \in c_3(a')$ and $a \in c_2(a')$, we have $y \neq a$ and thus $y \in \{b, c\}$. After a suitable relabeling we assume that $y = b$. Therefore w_b covers $c_1(a)$, and w_a covers the cone $c_3(a)$ which contains b. Since b lies in $c_{3R}(a)$, the point c lies in $c_{3L}(a)$. Notice that w_c covers $c_2(a)$ (which is not covered by $w_a \cup w_b$).

  Since $b' \in c_{1R}(a')$, the points c' lies in $c_{1L}(a')$ and $w_{c'}$ covers $c_3(a')$, and hence sees b. If $c' \in c_{1L}(a') \cap w_b$ (the dark-orange region in Figure 10) then b sees c' and we get $c' \leftrightarrow b$. Assume that $c' \in c_{1L}(a') \setminus w_b$ (the light-orange region). This region is to the left of $\overleftarrow{w_b}$ which is parallel to $\overrightarrow{r_2}(a)$; in particular this region is a subset of $c_{2R}(a)$ which is covered by w_c. Thus c sees c'. We are going to show that c' also sees c. Notice that $c_{2R}(a)$ (which contains c') is the reflection of $c_{3L}(a)$ (which contains c) with respect to a. This and the fact that $\overleftarrow{w_{c'}}$ is parallel to $\overrightarrow{r_1}(a')$ which intersects $\overrightarrow{r_1}(a)$ (because $\alpha \geqslant 2\pi/3$) imply that $\overleftarrow{w_{c'}}$ does not intersect $\overrightarrow{r_1}(a)$. Moreover, since $\overrightarrow{w_{c'}}$ is parallel to $\overrightarrow{r_3}(a')$ which intersects $\overleftarrow{r_2}(a)$ (because $\alpha < \pi$), $\overrightarrow{w_{c'}}$ does not intersect $\overleftarrow{r_2}(a)$. Thus $\overrightarrow{w_{c'}}$ and $\overleftarrow{w_{c'}}$ do not intersect the boundary rays $\overrightarrow{r_1}(a)$ and $\overleftarrow{r_2}(a)$ of $c_{3L}(a)$, and hence $w_{c'}$ covers $c_{3L}(a)$ which contains c'. Therefore c' sees c and hence $c' \leftrightarrow c$.

Configuration (C). In this configuration $\overrightarrow{r_1}(a')$ does not intersect any of $\overrightarrow{r_1}(a)$, $\overrightarrow{r_2}(a)$, $\overrightarrow{r_3}(a)$, and $\overrightarrow{r_2}(a)$ does not intersect any of $\overrightarrow{r_1}(a')$, $\overrightarrow{r_2}(a')$, $\overrightarrow{r_3}(a')$. Let d be the intersection point of $\overrightarrow{r_1}(a)$ with $\overrightarrow{r_2}(a')$, and e be the intersection point of $\overrightarrow{r_3}(a)$ with $\overrightarrow{r_3}(a')$, as in Figures 11 and 12. Let α denote the convex angle $\angle ada'$, and β denote the convex angle $\angle aea'$. After a suitable relabeling we assume that $\alpha \geqslant \beta$, and thus $\alpha \geqslant \pi/3$ (notice that the sum of the interior angles of the convex quadrilateral with vertices a, d, a', e is 2π). This and the fact that $\angle ead = 2\pi/3$ imply that $\overrightarrow{r_3}(a)$ and $\overleftarrow{r_2}(a')$ do not intersect. Since $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ intersect, their opposite rays $\overleftarrow{r_3}(a)$ and $\overleftarrow{r_3}(a')$ do not intersect.

Let $x \in P_1$ be the point in $\overline{c_3}(a)$ whose wedge w_x covers $c_3(a)$, and let $x' \in P_2$ be the point in $\overline{c_2}(a')$ whose wedge $w_{x'}$ covers $c_2(a')$. As in configuration (A) if $x \in \overline{c_3}(a) \cap c_2(a')$ and $x' \in \overline{c_2}(a') \cap c_3(a)$ (the dark-gray regions in Figures 11 and 12) then $x \leftrightarrow x'$. Therefore we may assume by symmetry that $x' \in \overline{c_2}(a') \setminus c_3(a)$. This and the fact that $a' \in c_3(a)$ imply that $x' \neq a'$ and thus $x' \in \{b', c'\}$. After a suitable relabeling we assume that $x' = b'$, and thus $w_{b'}$ covers $c_2(a')$, as in Figures 11 and 12. The region $\overline{c_2}(a') \setminus c_3(a)$ (shown by light-gray color in Figures 11 and 12) which contains b' is in fact equal to $c_{1R}(a') \cap c_{1R}(a)$.

The wedge $w_{a'}$ covers $c_1(a')$ which contains b'. Since b' lies in $c_{1R}(a')$, the point c' lies in $c_{1L}(a')$ and $w_{c'}$ covers $c_3(a')$. Let $y \in P_1$ be the point in $\overline{c_1}(a)$ whose wedge w_y covers $c_1(a)$. Since $b' \in c_1(a)$, y sees b'. If $y \in \overline{c_1}(a) \cap w_{b'}$ (the dark-blue regions in Figures 11 and 12) then b' sees y and thus $y \leftrightarrow b'$ and we are done. Assume that $y \in \overline{c_1}(a) \setminus w_{b'}$ (the light-blue regions). This and the fact that a lies in $c_2(a')$ (which is covered by $w_{b'}$) imply that $y \neq a$ and thus $y \in \{b, c\}$. After a suitable relabeling we assume that $y = b$. Our assumption that b is not in $w_{b'}$ implies that $b \notin c_2(a')$. Recall that a' is in $c_3(a)$ and thus it lies either in $c_{3L}(a)$ or in $c_{3R}(a)$. We describe each case separately.

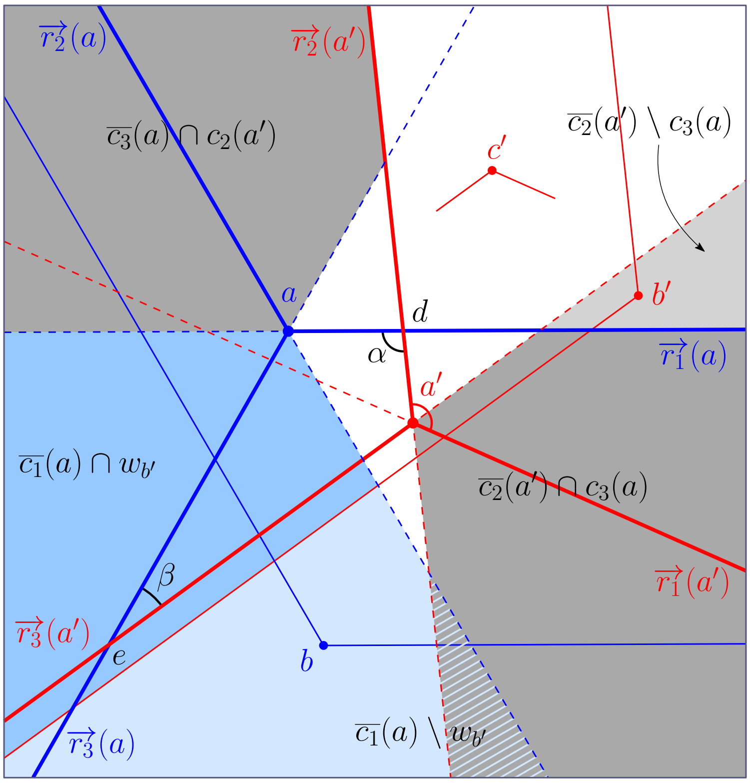

Figure 11: $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ intersect, $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ intersect, and $a' \in c_{3L}(a)$.

- $a' \in c_{3L}(a)$. This case is depicted in Figure 11. In this case $\overrightarrow{r_1}(a')$ and $\overleftarrow{r_2}(a)$ do not intersect because $\alpha \geqslant \pi/3$. Thus $\overline{c_1}(a)$ and $c_1(a')$ are disjoint, and hence $b \notin c_1(a')$. Therefore $b \in c_3(a')$ and thus c' sees b. The boundary rays $\overleftarrow{w_b}$ and $\overrightarrow{w_b}$ of w_b do not intersect the boundary rays $\overrightarrow{r_2}(a')$ and $\overleftarrow{r_3}(a')$ of cone $c_{1L}(a')$ which contains c'. Therefore w_b covers $c_{1L}(a')$. This implies that b sees c', and hence $b \leftrightarrow c'$.

- $a' \in c_{3R}(a)$. See Figure 12. If c' is to the left side of $\overleftarrow{r_2}(a)$ then $b \leftrightarrow c'$ and we are done. Assume that c' is to the right side of $\overleftarrow{r_2}(a)$, and thus $c' \in c_{3R}(a)$. If c' is to the right side of $\overleftarrow{w_b}$ then again $b \leftrightarrow c'$ and we are done. Assume that c' is to the left side of $\overleftarrow{w_b}$, which implies that $b \in c_{3R}(a)$ (we already knew that b is in $\overline{c_1}(a) = c_{2L}(a) \cup c_{3R}(a)$). Therefore w_a covers $c_3(a)$ which contains b, and thus c lies in $c_{3L}(a)$. This setting is depicted in Figure 12. We show that this is an invalid setting for points in P_1 and P_2.

  Consider the line through a and b. In the current setting a' and c' lie on the same side of this line, but b' lies on the other side. The points a' and c' are to the left side of the ray $\overrightarrow{b'a}$ (because $\overrightarrow{b'a}$ does not intersect the interior of $c_3(a)$ which contains a' and c') and to the right side of the ray $\overrightarrow{b'b}$ (because b is to the left side of $\overleftarrow{w_{b'}}$ while a' and c' are to its right side); see also Figure 13. Therefore a' and c' lie in the convex cone with boundary rays $\overrightarrow{b'a}$ and $\overrightarrow{b'b}$. Since c is in $c_{3L}(a)$, it is also in the convex cone with boundary rays $\overrightarrow{ab'}$ and $\overrightarrow{ab}$. As depicted in Figure 13, no matter where—in this cone—c lies, the two triangles defined by the vertices of P_1 and P_2 form one of the forbidden configurations of Figure 6. This is impossible because $\{P_1, P_2\}$ is not forbidden by the statement of the theorem.

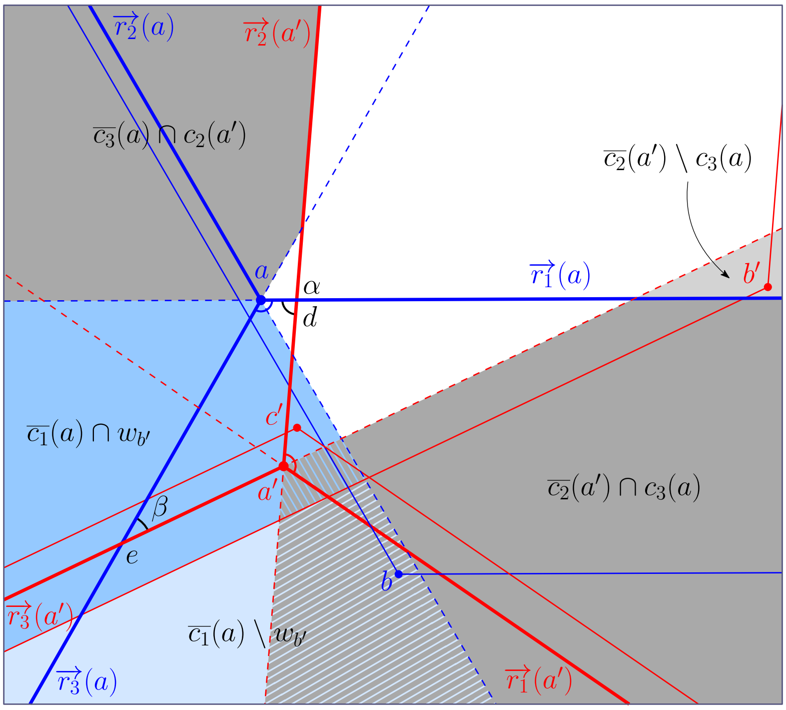

Figure 12: $\overrightarrow{r_1}(a)$ and $\overrightarrow{r_2}(a')$ intersect, $\overrightarrow{r_3}(a)$ and $\overrightarrow{r_3}(a')$ intersect, and $a' \in c_{3R}(a)$.

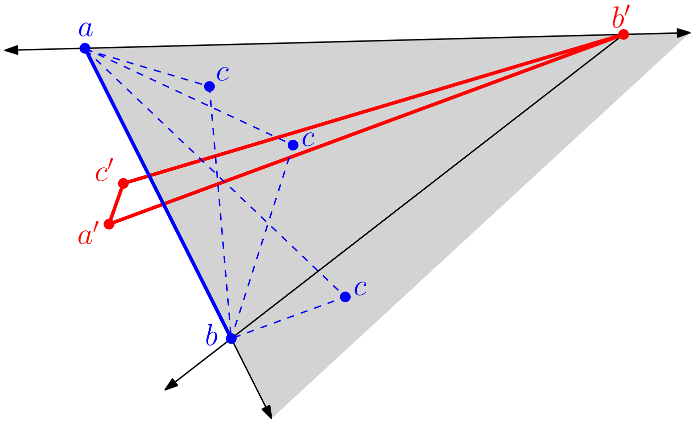

Figure 13: The point c lies in the shaded cone with boundary rays $\overrightarrow{ab'}$ and $\overrightarrow{ab}$.

4 Approximating the $\overline{120}^\circ$-MST

In this section we prove the following theorem.

Theorem 6. *Given a set of points in the plane and an angle $\alpha \geqslant 120^\circ$, there is an $\overline{\alpha}$-spanning tree of length at most 1.5 times the length of the MST and with edges of length at most 2 times the BST largest edge-length. Furthermore, there is an algorithm to find such an $\overline{\alpha}$-ST that runs in linear time after computing the MST.*

Let P be a set of points in the plane. As in previous sections we assume that $\alpha = 120^\circ$. Let T be a degree-5 MST of P. We assign to each vertex of T an initial angle α which we refer to by "charge". The initial charge of a vertex may not cover all its incident edges. The idea is to modify the tree locally and transfer charges between nodes to make sure that all vertices have enough charges to cover their incident edges in the new tree. Mainly we transfer charges from leaves to internal vertices, because edges incident to leaves can be covered by $0°$ wedges (however at the end of this section we assign to all leaves positive wedges).

Assume that T is not a path and thus has a vertex of degree at least 3 (we describe the case where T is a path at the end of this section). A *maximal path* in T is a path with at least two edges where its internal-node degrees are 2 and its end-node degrees are not 2.

Our algorithm has two phases. Phase 1 works as follows. Contract every maximal path of T to an edge (this is done by removing the internal nodes of the path and connecting its endpoints by an edge). This results in a tree T' that has no vertex of degree 2, and moreover, the degree of each vertex in T' is the same as its degree in T. Let ℓ denote the number of leaves of T' and let n_3, n_4, and n_5 denote the number of vertices of degree three, four, and five in T' respectively. Since each vertex of degree three, four, and five introduces 1, 2, and 3 new leaves respectively, we have $\ell = 2 + n_3 + 2n_4 + 3n_5$. We consider the $\ell \cdot \alpha$ charges of all leaves as a "pool of charges" that are available to be distributed among other vertices. From this pool, we give the charge α, 2α, and 3α to each vertex of degree three, four, and five respectively. After this redistribution, every vertex of degree three, four, and five holds the charge $240°$, $360°$, and $480°$ respectively (including its initial $120°$ charge), which is sufficient to cover all its incident edges. Moreover, the pool is left with 2α charges.

Based on the above discussion if T has no degree-2 vertices, then it is an $\overline{\alpha}$-ST and we are done. Now we describe Phase 2 which takes care of contracted paths; here is the place where our tree gains an extra $0.5w(T)$ length. Consider every contracted path and denote it by $(p_1, p_2, \ldots, p_m)$ if it has an even number of edges and by $(p_1, p_2, \ldots, p_m, p_{m+1})$ if it has an odd number of edges. One endpoint of this path, say p_1, has degree at least 3 in T and the other endpoint either is a leaf or it has degree at least 3 (see Figure 14). The charges of the two endpoints of the path have already been considered in Phase 1, but the charges of its internal nodes are untouched. Let σ denote the path $(p_1, p_2, \ldots, p_m)$; σ does not contain p_{m+1} if this point exists. Observe that $m \geqslant 3$. Partition the edges of σ into matchings $\sigma_1 = \{(p_1, p_2), (p_3, p_4), \ldots\}$ and $\sigma_2 = \{(p_2, p_3), (p_4, p_5), \ldots\}$, as in Figure 14. Let $\sigma_{\max}$ denote the heavier of σ_1 and σ_2 (i.e., the one with larger total length) and $\sigma_{\min}$ denote the lighter one. Thus $w(\sigma_{\min}) \leqslant w(\sigma_{\max})$ and hence $w(\sigma_{\min}) \leqslant 0.5w(\sigma)$. We replace the edge set $\sigma_{\max}$ in T by the edge set $\delta = \{(p_1, p_3), (p_3, p_5), \ldots, (p_{m-2}, p_m)\}$, as in Figure 14. Let T'' be the tree obtained after performing this replacement for all contracted paths.

We claim that T'' is a desired $\overline{\alpha}$-ST. First we show that $w(T'') \leqslant 1.5w(T)$. To do so, it suffices to show that, for each contracted path, the length of edges after replacement (i.e., edges of $\sigma_{\min}$ and δ) is not more than 1.5 times the length of original edges (i.e., edges of σ). By the triangle inequality, $w(\delta) \leqslant w(\sigma)$. Therefore $w(\delta) + w(\sigma_{\min}) \leqslant w(\sigma) + 0.5w(\sigma) \leqslant 1.5w(\sigma)$; this proves the total-length constraint of Theorem 6. Moreover, the length of every edge of δ is at most twice the largest edge-length of σ (again by the triangle inequality); this proves the edge-length constraint of Theorem 6. It remains to ensure the coverage of incident edges for all vertices of T''. To that end we distribute the charges of $p_2, p_4, \ldots, p_{m-1}$ (which are new leaves) among other vertices. We consider two cases depending on the existence of p_{m+1}. First assume that p_{m+1} does not exist. Now we consider two sub-cases depending on which of σ_1 and σ_2 is heavier (i.e., is $\sigma_{\max}$).

- $\sigma_{\max} = \sigma_1$. In this case σ_1 has been replaced by δ. This replacement has not changed the degree of p_1 and thus it holds enough charge (from Phase 1) to cover its incident edges. We move the charges of leaves p_2, p_4, $\ldots, p_{m-1}$ to vertices $p_3, p_5, \ldots, p_m$ respectively (see Figure 14). Each of $p_3, p_5, \ldots, p_{m-2}$ has degree three and now holds $240°$ charge (including its own $120°$ charge) which is sufficient to cover its incident edges. Now consider p_m and

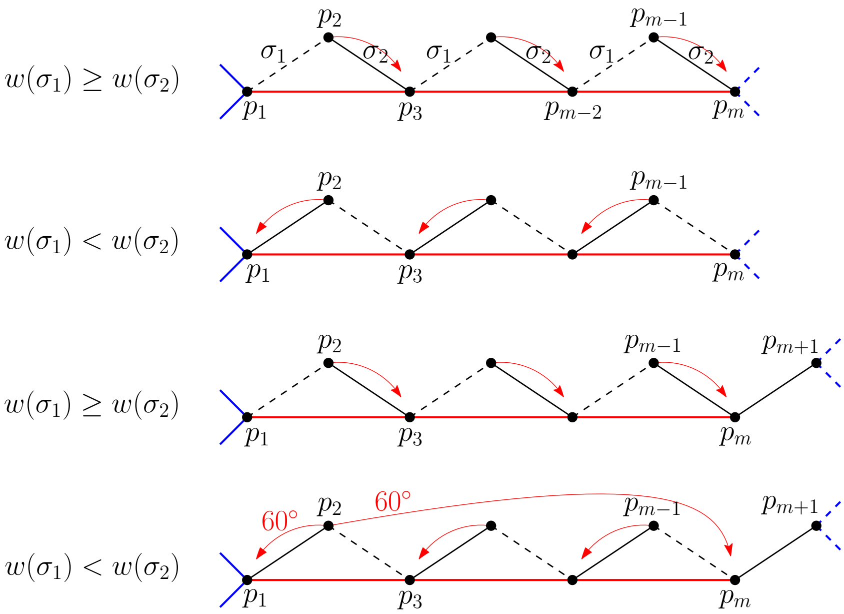

Figure 14: The contracted path is shown by black segments. The dashed-black edges belong to $\sigma_{\max}$ and the red edges belong to δ.

notice that its degree has been increased by one. If p_m is of degree at least 3 in T, then it now holds at least 360° charge (120° from p_{m-1} and at least 240° from Phase 1) which covers all its incident edges. Assume that p_m is a leaf in T, and thus it has degree 2 in T''. The original charge of p_m has been distributed among other vertices in Phase 1, but it now has 120° charge coming from p_{m-1}. Consider the triangle $\triangle p_m p_{m-1} p_{m-2}$. The segment $p_m p_{m-2}$ is the largest side of this triangle because otherwise we could add this edge to the MST and remove the larger of the other two sides to obtain a smaller tree. Thus the angle at p_{m-1} is the largest angle of this triangle, and hence the other two angles (including the one at p_m) are at most 90°. Therefore the two edges incident to p_m can be covered by its 120° charge.

- $\sigma_{\max} = \sigma_2$. Then σ_2 has been replaced by δ. We move the charges of $p_2, p_4, \ldots, p_{m-1}$ to $p_1, p_3, \ldots, p_{m-2}$ respectively (see Figure 14). The replacement does not change the degree of p_m and thus it holds enough charge from Phase 1 to cover its incident edges. Each of $p_3, p_5, \ldots, p_{m-2}$ has degree three and now holds 240° charge which covers its incident edges. The vertex p_1 now has at least 360° charge (120° from p_2 and at least 240° from Phase 1) which covers all its incident edges.

Now assume that p_{m+1} exists. Again we consider two sub-cases.

- $\sigma_{\max} = \sigma_1$. Move the charges of $p_2, p_4, \ldots, p_{m-1}$ to $p_3, p_5, \ldots, p_m$ (see Figure 14). The degrees of p_1 and p_{m+1} have not changed in Phase 2, and thus they hold enough charge from Phase 1 to cover their incident edges. Each of $p_3, p_5, \ldots, p_m$ has degree three and now holds 240° charge which covers its incident edges.

- $\sigma_{\max} = \sigma_2$. Move the charges of $p_4, p_6 \ldots, p_{m-1}$ to $p_3, p_5, \ldots, p_{m-2}$. Split the 120° charge of p_2 evenly between p_1 and p_m as in Figure 14. The degree of p_{m+1} has not changed and thus it holds enough charge from Phase 1. Each of $p_3, p_5, \ldots, p_{m-2}$ has degree three and now holds enough charge 240°. The vertex p_m has degree two and now holds 180° charge (including its original 120° charge) which is sufficient to cover its two incident edges. Now consider p_1. If it has degree 4 or 5 in T then after Phase 1 it holds at least 360° charge which covers all its incident edges. Assume that it has degree 3 in T, and now it has degree 4 in T''. Then it holds 300° charge (60° from p_2 and 240° from Phase 1) which is enough to cover its four incident edges (in fact 270° is enough). It might be the case that p_1 was also incident to another contracted path and hence has received another incident edge which increases p_1's degree to five. In this case p_1 also receives 60° charge from the other path. This would increase p_1's charge to 360° which covers all its incident edges.

Remark 2. If T is a path then we run only Phase 2 of the above algorithm on this path. Since there is no Phase 1, the charges of path endpoints are still available. The charge redistribution of Phase 2 guarantees the coverage of incident edges of all vertices. See Figure 14. If p_{m+1} does not exist then in the first sub-case (resp. the second sub-case) the charge of p_1 (resp. p_m) remains unused. If p_{m+1} exists, then its charge remains unused.

Remark 3. At the end of the algorithm we have at least 120° unused charge (from the pool or from a path endpoint). We split the unused charge evenly between all vertices such that every vertex has a positive charge and its incident edges lie strictly inside the assigned wedge.

5 Conclusions

The obvious open problem is to improve our 16/3 approximation ratio for the 120°-MST problem further by designing better algorithms. Our proof rely on the fact that the orientation of the wedges of every triplet covers the entire plane. Such orientations are a bottleneck for our ratio. It might be possible to get better ratio with orientations that do not necessarily cover the entire plane. Another bottleneck is the use of a Hamiltonian path which forces a factor of 2 in the ratio. It might be possible to get better ratios by using the original MST instead of the path. However, this has to be done in a clever way as the number of connecting edges incident to a triplet may increase.

Acknowledgements. Matthew Katz's invited talk on bounded-angle spanning tree problems at CCCG 2018 was the inspiration for this work. We thank Therese Biedl for helpful discussions.

References

[1] E. Ackerman, T. Gelander, and R. Pinchasi. Ice-creams and wedge graphs. *Computational Geometry: Theory and Applications*, 46(3):213–218, 2013.

[2] O. Aichholzer, T. Hackl, M. Hoffmann, C. Huemer, A. Pór, F. Santos, B. Speckmann, and B. Vogtenhuber. Maximizing maximal angles for plane straight-line graphs. *Computational Geometry: Theory and Applications*, 46(1):17–28, 2013. Also in *WADS'07*.

[3] S. Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *Journal of the ACM*, 45(5):753–782, 1998.

[4] R. Aschner and M. J. Katz. Bounded-angle spanning tree: Modeling networks with angular constraints. *Algorithmica*, 77(2):349–373, 2017. Also in *ICALP'14*.

[5] R. Aschner, M. J. Katz, and G. Morgenstern. Do directional antennas facilitate in reducing interferences? In *Proceedings of the 13th Scandinavian Symposium and Workshops on Algorithm Theory* (*SWAT*), pages 201–212, 2012.

[6] A. Biniaz. Euclidean bottleneck bounded-degree spanning tree ratios. In *ACM-SIAM Symposium on Discrete Algorithms* (*SODA*), 2020.

[7] A. Biniaz, A. Maheshwari, and M. Smid. Flip distance to some plane configurations. *Computational Geometry: Theory and Applications*, 81:12–21, 2019. Also in *SWAT'18*.

[8] P. Bose, P. Carmi, M. Damian, R. Y. Flatland, M. J. Katz, and A. Maheshwari. Switching to directional antennas with constant increase in radius and hop distance. *Algorithmica*, 69(2):397–409, 2014. Also in *WADS'11*.

[9] P. Bose, L. J. Guibas, A. Lubiw, M. H. Overmars, D. L. Souvaine, and J. Urrutia. The floodlight problem. *International Journal of Computational Geometry & Applications*, 7(1/2):153–163, 1997. Also in *CCCG'93*.

[10] P. M. Camerini. The min-max spanning tree problem and some extensions. *Information Processing Letters*, 7(1):10–14, 1978.

[11] I. Caragiannis, C. Kaklamanis, E. Kranakis, D. Krizanc, and A. Wiese. Communication in wireless networks with directional antennas. In *Proceedings of the 20th Annual ACM Symposium on Parallelism in Algorithms and Architectures* (*SPAA*), pages 344–351, 2008.

[12] P. Carmi, M. J. Katz, Z. Lotker, and A. Rosén. Connectivity guarantees for wireless networks with directional antennas. *Computational Geometry: Theory and Applications*, 44(9):477–485, 2011.

[13] T. M. Chan. Euclidean bounded-degree spanning tree ratios. *Discrete & Computational Geometry*, 32(2):177–194, 2004. Also in *SoCG'03*.

[14] M. Damian and R. Y. Flatland. Spanning properties of graphs induced by directional antennas. *Discrete Mathematics, Algorithms and Applications*, 5(3), 2013.

[15] S. Dobrev, E. Kranakis, D. Krizanc, J. Opatrny, O. M. Ponce, and L. Stacho. Strong connectivity in sensor networks with given number of directional antennae of bounded angle. *Discrete Mathematics, Algorithms and Applications*, 4(3), 2012. Also in *COCOA'10*.

[16] S. Dobrev, E. Kranakis, O. M. Ponce, and M. Plzík. Robust sensor range for constructing strongly connected spanning digraphs in UDGs. In *Proceedings of the 7th International Computer Science Symposium in Russia* (*CSR*), pages 112–124, 2012.

[17] A. Dumitrescu, J. Pach, and G. Tóth. Drawing Hamiltonian cycles with no large angles. *Electronic Journal of Combinatorics*, 19(2):P31, 2012. Also in *GD'94*.

[18] S. P. Fekete, S. Khuller, M. Klemmstein, B. Raghavachari, and N. E. Young. A network-flow technique for finding low-weight bounded-degree spanning trees. *Journal of Algorithms*, 24(2):310–324, 1997. Also in *IPCO* 1996.

[19] S. P. Fekete and G. J. Woeginger. Angle-restricted tours in the plane. *Computational Geometry: Theory and Applications*, 8:195–218, 1997.

[20] M. M. Halldórsson and T. Tokuyama. Minimizing interference of a wireless ad-hoc network in a plane. *Theoretical Computer Science*, 402(1):29–42, 2008.

[21] R. Jothi and B. Raghavachari. Degree-bounded minimum spanning trees. *Discrete Applied Mathematics*, 157(5):960–970, 2009.

[22] M. J. Katz. Personal communication.

[23] S. Khuller, B. Raghavachari, and N. E. Young. Low-degree spanning trees of small weight. *SIAM Journal on Computing*, 25(2):355–368, 1996. Also in *STOC'94*.

[24] E. Kranakis, F. MacQuarrie, and O. M. Ponce. Connectivity and stretch factor trade-offs in wireless sensor networks with directional antennae. *Theoretical Computer Science*, 590:55–72, 2015.

[25] C. L. Monma and S. Suri. Transitions in geometric minimum spanning trees. *Discrete & Computational Geometry*, 8:265–293, 1992. Also in *SoCG'91*.

[26] J. van Leeuwen and A. A. Schoone. Untangling a travelling salesman tour in the plane. In *Proceedings of the 7th Conference Graphtheoretic Concepts in Computer Science* (*WG*), pages 87–98, 1981.