Sequences of Spanning Trees and a Fixed Tree Theorem *

Oswin Aichholzer†
Franz Aurenhammer‡
Institute for Theoretical Computer Science
Graz University of Technology
Graz, Austria
e-mail: {oaich,auren}@igi.tu-graz.ac.at

Ferran Hurtado§
Departament de Matematica Aplicada II
Universitat Politecnica de Catalunya
Barcelona, Spain
e-mail: hurtado@ma2.upc.es

Abstract

Let $\mathcal{T}_S$ be the set of all crossing-free spanning trees of a planar n-point set S. We prove that $\mathcal{T}_S$ contains, for each of its members T, a length-decreasing sequence of trees $T_o, \ldots, T_k$ such that $T_o = T$, $T_k = MST(S)$, T_i does not cross T_{i-1} for $i = 1, \ldots, k$, and $k = O(\log n)$. Here $MST(S)$ denotes the Euclidean minimum spanning tree of the point set S.

As an implication, the number of length-improving and planar edge moves needed to transform a tree $T \in \mathcal{T}_S$ into $MST(S)$ is only $O(n \log n)$. Moreover, it is possible to transform any two trees in $\mathcal{T}_S$ into each other by means of a local and constant-size edge slide operation. Applications of these results to morphing of simple polygons are possible by using a crossing-free spanning tree as a skeleton description of a polygon.

1 Introduction

Let S be a set of n points in the Euclidean plane. A *crossing-free spanning tree* of S is a tree whose edges connect all points in S (and no others) with straight line segments that pairwise do not cross. In this paper, we investigate the questions of whether, and how fast, two crossing-free spanning trees of S can be transformed into each other by means of predefined rules.

---

*Part of this work was carried out while the first author visited the Departament de Matematica Aplicada II, Universitat Politecnica de Catalunya, and while the third author visited the Institute for Theoretical Computer Science, Graz University of Technology

†Research supported by the Austrian Programme for Advanced Research and Technology, and the EU Improving Human Potential Programme HPRI-1999-CT-00071

‡Research supported by the Austrian SFB F003 Optimierung und Kontrolle

§Research supported by Proyectos DGES-MEC PB98-0933 and Gen. Cat. 1999SGR00356

More formally, consider the set $\mathcal{T}_S$ of all crossing-free spanning trees of the point set S. A *tree graph* $\mathcal{TG}_{\mathtt{op}}(S)$ is a directed graph that has $\mathcal{T}_S$ as its set of nodes and that realizes an arc from node (tree) T to node T' iff $T' = \mathtt{op}(T)$, where op is some predefined operation. The existence of a path between two nodes in $\mathcal{TG}_{\mathtt{op}}(S)$ implies transformability of the corresponding trees into each other by means of the operation op. The length of a shortest path corresponds to the distance between the two trees with respect to the operation op. Distances of this kind provide a measure of similarity between trees.

Purely graph-theoretical versions of the problem (i.e., without the restriction to crossing-free trees) have been largely studied; see e.g. Holzmann and Harary [14], Liu [18], or for graphs more general than trees, Goddard and Swart [11]. For the geometric version which is the topic of the present paper, results are recent and relatively sparse. They exclusively treat the case where the allowed operation is what is called an *edge move*, which relates two trees in the set $\mathcal{T}_S$ if they have all but one edge in common. Motivated by the question of enumerating the set $\mathcal{T}_S$, Avis and Fukuda [6] showed that the corresponding tree graph is connected and has a diameter bounded by $2n-4$. To the knowledge of the authors, this is the only known result for general point sets S. The special case where S is a set of points in convex position has been treated in García et al. [10], who showed that the tree graph has maximum connectivity and is Hamiltonian in this case. A lower bound of $3n/2-5$ on its diameter is also given there. Hernando et al. [12] proved that the number of crossing-free spanning trees is minimum in the convex case. See Károlyi et al. [15] for results of a different flavour on crossing-free geometric spanning trees.

The present paper provides a thorough study of path lengths in tree graphs $\mathcal{TG}_{\mathtt{op}}(S)$, for a general planar point set S, and for three operations op which we prove to be successive refinements of each other. The strongest operation, min, maps a given tree $T \in \mathcal{T}_S$ to the tree $\mathtt{min}(T)$ of minimum Euclidean length that does not cross T. Its refinement, the operation move, is the classical edge move mentioned above, with the additional requirement of reducing the tree length. We will talk of an *improving edge move* in this case. The weakest operation, slide, is a local edge move that keeps one endpoint of the moved edge fixed and moves the other one along an adjacent tree edge. Following [11], we will call this operation an *edge slide*.

Our results mainly rely on a fact which might be of interest in its own right. It can be stated as follows. Let $MST(S)$ be the Euclidean minimum spanning tree of the point set S. For a tree $T \in \mathcal{T}_S$, let $CD(T)$ be the constrained Delaunay triangulation of T. Then, for every tree $T \in \mathcal{T}_S$, there exists a unique (and logarithmically short) sequence of trees $T_o, \ldots, T_k$ that starts at T, ends at $MST(S)$, and fulfills $T_i \subset CD(T_{i-1})$ for $i = 1, \ldots, k$. This fact is explained in Section 2, where we rigerously define the operation min and prove a basic yet seemingly unknown fixed point-type property: a tree $T \in \mathcal{T}_S$ fulfills the property $T = \mathtt{min}(T)$ if and only if $T = MST(S)$. As a consequence, the tree graph $\mathcal{TG}_{\mathtt{min}}(S)$ is a tree all whose paths are directed towards its root, $MST(S)$. In Section 3, an upper bound of $O(n)$ and a lower bound of $\Omega(\log n)$ for path lengths in $\mathcal{TG}_{\mathtt{min}}(S)$ are given. This gap is closed in Section 4 where the upper bound is reduced to $O(\log n)$ by more involved methods. Logarithmically small path lengths seem surprising in view of the fact that consecutive trees on any path of $\mathcal{TG}_{\mathtt{min}}(S)$ do not cross each other.

In Section 5, we turn to edge operations on crossing-free spanning trees. We show that any tree $T \in \mathcal{T}_S$ can be tranformed into the (shorter) tree $T' = \mathtt{min}(T)$ by at most $n-1$ improving edge moves. An algorithm that constructs such a sequence of moves in $O(n \log n)$ time is provided. By the results on the operation `min` stated above, the number of improving edge moves needed to transform a given tree $T \in \mathcal{T}_S$ into $MST(S)$ now is $O(n \log n)$, and $O(n \log^2 n)$ time suffices for finding such a sequence of moves. That is, the tree graph $\mathcal{TG}_{\mathtt{move}}(S)$ is a directed acyclic graph with unique sink $MST(S)$ and with shortest-path lengths bounded by $O(n \log n)$. This compares well (by a factor of $O(\log n)$) to the result of Avis and Fukuda [6] for unrestricted edge moves. Our result gains in importance from the fact that – unlike the unrestricted case – there is in general no path in $\mathcal{TG}_{\mathtt{move}}(S)$ from a given tree to a given and shorter tree. Our result should also be seen in contrast to the situation for triangulations of a point set S, where it is well known (see e.g. Fortune [9]) that $\Theta(n^2)$ improving edge operations (Delaunay flips) may be necessary to reach the global optimum, the Delaunay triangulation $DT(S)$.

Section 6 demonstrates that edge moves can always be simulated by a sequence of edge slides. This implies that the tree graph $\mathcal{TG}_{\mathtt{slide}}(S)$ is connected. To our knowledge, this is the first result showing that any two crossing-free spanning trees of a point set S can be transformed into each other by local and constant-size operations. As is reflected by the informal definition of an edge slide above, this operation can be carried out in a (geometrically) continuous way. Thus our results allow for a continuous deformation of crossing-free spanning trees into each other. Deriving satisfactory bounds on shortest path lengths in $\mathcal{TG}_{\mathtt{slide}}(S)$ seems to require techniques different from the ones used in this paper, however. We decided to elaborate on this question in a forthcoming paper. Our conjecture is that any two trees in $\mathcal{T}_S$ that do not cross each other can be joined by $O(n^2)$ edge slides. An affirmative answer would imply a path length bound of $O(n^2 \log n)$ by the results of Sections 2 and 3. Edge slide operations could also prove useful in enumerating all simple polygons on a point set S via constant-size local transformations. This question is still unsettled; see e.g. Hernando et al. [13].

To substantiate the practical relevance of our theoretical results, we briefly mention a situation where a controllable deformation of crossing-free spanning trees might be of advantage: *morphing of simple polygons*. The idea is to use a crossing-free tree as a skeleton description of a polygon. Now, given two simple polygons P and Q, a morph between them can be computed as follows. First, a suitable overlay of P and Q has to be found, by using translation, scaling, and rotation. Algorithms for computing an overlay are known for various optimality measures, like Hausdorff distance, Fréchet distance, and area of symmetric difference (see [1], [4], [3], respectively, and references therein). The next step is the selection of a point set S in the intersection $P \cap Q$, as well as of two spanning trees T_P and T_Q of S which approximate a skeleton description of the respective polygon. The medial axis and the straight skeleton (see e.g. the survey article [5]) are appropriate descriptions because they allow for recovering the original polygon shape. Computing a morph of P into Q now amounts to deforming tree T_P into tree T_Q. Our results offer a trade-off between efficiency and accuracy for this task, ranging from the very local edge slide – which is likely to give a long sequence – to a much stronger operation which gives a sequence of only logarithmically many trees that still might shape the set adequately in many cases. Also, our approach for generating tree sequences seems preferable to the approach in [6] because it uses the minimum spanning tree – rather than the star centered at some convex hull point – for an intermediate 'standard form' which T_P and T_Q are transformed into.

2 Fixed tree theorem

In this section, a certain minimum operation on crossing-free spanning trees is introduced and investigated. In particular, the operation is shown to fulfill a fixed point-type property which is central to the further investigations in this paper. We start by reviewing some basics on geometric minimum spanning trees.

Let $S \times S$ be the set of all (straight line) edges spanned by the point set S. For a connected subgraph $G \subset S \times S$ let $MST(G)$ denote a spanning tree of G which minimizes the total (Euclidean) edge length.

It is well known that $MST(G)$ can be constructed in a greedy fashion, by applying Kruskal's [16] widely used algorithm. This algorithm adds G's edges in increasing length order, and discards edges which close a cycle in the graph built so far. We assume throughout that all edges in $S \times S$ have different lengths.[1] Then, by construction, $MST(G)$ is unique for each choice of G. Another property of minimum spanning trees is evident from the greedy construction. We will use this property in many proofs later in this paper, and therefore state it explicitely.

Property 2.1 *(P(ath)-property) An edge $e \in G$ is not present in $MST(G)$ if and only if there is some path in G between e's endpoints which solely consists of edges shorter than e.*

We proceed by introducing the concept of *constrained* minimum spanning tree. Rather than forcing a spanning tree of S to use only prescribed edges, the tree may be forced not to cross them. Consider some planar straight line graph G on S. An edge $e \in S \times S$ is called *blocked by G* if e crosses some edge of G. (Edge e crosses

[1] This assumption simplifies the exposition. It may be removed by imposing some fixed order on edges of the same length. All results of the paper remain valid.

edge f iff $e \cap f$ is a point which lies in the relative interior of both edges.) Edge e is called *visible under* G otherwise. Note that G itself consists of visible edges because G does not cross itself. If now E_G is the set of all edges visible under G then $MST(E_G)$ is termed the *minimum spanning tree of* S *constrained by* G, or $\mathtt{min}(G)$ for short.

A special case important to us arises when the constraining graph is a tree itself. Then $\mathtt{min}$ can be interpreted as an operation that transforms a spanning tree T of S into another one, $T' = \mathtt{min}(T)$. Our interest is in how $\mathtt{min}$ operates on the set $\mathcal{T}_S$ of all crossing-free spanning trees of S, that is, in the properties of the tree graph $\mathcal{TG}_{\mathtt{min}}(S)$. In the remainder of this section, we will show that the set $\mathcal{T}_S$ is closed under the operation $\mathtt{min}$, and that the directed graph $\mathcal{TG}_{\mathtt{min}}(S)$ has a unique sink, $MST(S)$, the (global) minimum spanning tree of the point set S. Both properties are by no means trivial. The former guarantees that – in spite of the visibility constraints – no self-crossing trees are generated, and the latter excludes the possibility of local minima (concerning tree length) in $\mathcal{TG}_{\mathtt{min}}(S)$.

In the investigation of the operation $\mathtt{min}$ a well-known structure defined by visibility constraints will be utilized: the *constrained Delaunay triangulation*, $CD(G)$, of a planar straight line graph G of S, introduced in Lee and Lin [17]. This triangulation of S includes all the edges of G, and all edges $xy \in S \times S$ for which there is some circle that encloses x and y but no other point $p \in S$ for which both edges px and py are visible under G.

Lemma 2.1 *Let G be any planar straight line graph on S. Then* $\mathtt{min}(G) = MST(CD(G))$.

Proof As $CD(G)$ includes the graph G, all edges of $CD(G)$ are visible under G. It therefore suffices to prove that $\mathtt{min}(G)$ is a subset of $CD(G)$. Let e be an edge of $\mathtt{min}(G)$. Edge e is visible under G, so in order to have $e \notin CD(G)$ every circle C which encloses e also has to enclose some point $p \in S$ which is visible from both endpoints x and y of e. Choose for C the circle with diameter e. Then px and py are both shorter than e and, as being visible edges, they yield a path which excludes e from $\mathtt{min}(G)$ by the P-property – a contradiction. $\square$

Lemma 2.1 is a generalization of the well-known fact that $MST(S)$ is a subset of the Delaunay triangulation, $DT(S)$, of S. One of the consequences of Lemma 2.1 is that the tree $\mathtt{min}(G)$ is crossing-free because it is part of a triangulation, that is, $\mathtt{min}(G) \in \mathcal{T}_S$ for any planar straight line graph G of S. We further note that $\mathtt{min}(G)$ can be constructed by a greedy algorithm that inserts the edges of $CD(G)$ in length-increasing order, implementable in time $O(n \log n)$ and space $O(n)$. These are just the requirements for computing $CD(G)$; see e.g. Chew [8].

We proceed by recalling the notion of D(elaunay)-flippable edges and by studying their relevance to constrained Delaunay triangulations and minimum spanning trees.

Consider an arbitrary triangulation Δ of S. Let e be an edge of Δ which does not lie on the convex hull of S. Then e is common to two triangles $t_1(e)$ and $t_2(e)$ of Δ. Edge e is termed *D-flippable in* Δ if there exists some circle that has e as a chord and that encloses both of $t_1(e)$ and $t_2(e)$. (Equivalently, the circumcircle of $t_1(e)$ encloses $t_2(e)$ and vice versa.) By definition of a Delaunay triangulation, we have $\Delta = DT(S)$ iff no edge of Δ is D-flippable. The following can be asserted on constrained Delaunay triangulations.

Observation 2.1 *Consider an arbitrary tree $T \in \mathcal{T}_S$. All D-flippable edges of $CD(T)$ are in T.*

Proof Let e be a D-flippable edge of $CD(T)$. Then every circle which encloses e also has to enclose at least one triangle of $CD(T)$ incident to e, say $t_1(e)$. All edges of $t_1(e)$ are visible under T because $T \subset CD(T)$. So, by the definition of a constrained Delaunay triangulation, edge e can be D-flippable only if $e \in T$. $\square$

Hence the tree T maximizes the number of D-flippable edges that are contained in a spanning tree of $CD(T)$. Conversely, it is easy to show that $CD(T)$ maximizes the number of D-flippable edges in T over all triangulations of S that conform with T. Observation 2.1 remains valid if T is replaced by a general planar straight-line graph on S. We further have the following property of minimum spanning trees.

Observation 2.2 *For every triangulation Δ of S, no edge of $MST(\Delta)$ is D-flippable in Δ.*

Proof Let edge e be D-flippable in Δ. Then the circle with diameter e enclosed a triangle $t_1(e)$ of Δ which is incident to e. This implies that e is the longest edge in $t_1(e)$. But then $e \notin MST(\Delta)$ by the P-property. $\square$

The tree $T = MST(\Delta)$ thus minimizes the number of D-flippable edges over all spanning trees of Δ. Conversely, Δ minimizes the number of D-flippable edges in T over all triangulations of S that conform with T.

We are now ready to state and prove the following fixed point-type theorem.

Theorem 2.1 *(Fixed tree theorem) Consider some tree $T \in \mathcal{T}_S$. We have $T = \mathtt{min}(T)$ if and only if $T = MST(S)$.*

Proof The if part is, of course, trivial as the fixed tree property $T = \mathtt{min}(T)$ is fulfilled by $MST(S)$. Let us prove the only-if part. $T = \mathtt{min}(T)$ is equivalent to $T = MST(CD(T))$ by Lemma 2.1. From Observation 2.1 we know that all D-flippable edges of $CD(T)$ are in T. Furthermore, $MST(CD(T))$ contains no such edge, by Observation 2.2. So $T = MST(CD(T))$ implies that no edge of $CD(T)$ is D-flippable. Consequently, $CD(T) = DT(S)$ and $T = MST(S)$, as claimed. $\square$

Let us consider the consequences of Theorem 2.1 to the tree graph $\mathcal{TG}_{\texttt{min}}(S)$. Consider an arbitrary tree $T \in \mathcal{T}_S$ and repeatedly apply to T the operation min. This yields a sequence of trees $T_o = T$, $T_1 = \texttt{min}(T_o)$, $T_2 = \texttt{min}(T_1)$, and so on. By Lemma 2.1, consecutive trees T_i and T_{i+1} are part of the same triangulation of S, namely $CD(T_i)$. So both trees belong to $\mathcal{T}_S$ and, in addition, they do not cross each other. Tree T_{i+1} is, by definition, minimum in $CD(T_i)$ but tree T_i is not unless both trees are identical. That is, T_{i+1} has to be shorter than T_i in the former case. So the operation min produces a length-decreasing sequence of trees which reaches a 'fixed point' $T_k = \texttt{min}(T_k)$ after a finite number k of steps. Theorem 2.1 now tells us that $T_k = MST(S)$. As a consequence, for each tree $T \in \mathcal{T}_S$ there is a unique path in $\mathcal{TG}_{\texttt{min}}(S)$ leading from T to $MST(S)$. We conclude:

Theorem 2.2 *The tree graph $\mathcal{TG}_{\texttt{min}}(S)$ is a tree all whose paths are directed towards the root $MST(S)$.*

3 Canonical sequences

The purpose of this section (and of Section 4) is in deriving bounds on the path lengths in the tree graph $\mathcal{TG}_{\texttt{min}}(S)$. In the light of Theorem 2.2, we define the *canonical sequence* of a crossing-free spanning tree T of S as the unique path in $\mathcal{TG}_{\texttt{min}}(S)$ from T to $MST(S)$.

Canonical sequences will prove useful in our study of more refined operations on crossing-free spanning trees in Sections 5 and 6. Sequences of *edge operations* (certain edge moves and edge slides) can be constructed on this basis. Let, here and in the rest of this paper, $T_o, \ldots, T_k$ denote an arbitrary but fixed canonical sequence. Determining its length k with respect to the number n of points in S is – apart from the theoretical interest – of influence to the number of such edge operations.

Figure 1 illustrates a canonical sequence of length two. For each tree T_i drawn in bold, the edges of $T_{i+1} \setminus T_i$ are shown dashed.

An exponential upper bound on k is evident from the maximal number of planar straight-line graphs for an n-point set S; see Ajtai et al. [2]. A bound $k = O(n^2)$ follows form the lemma below which asserts that, intuitively speaking, at least one edge of $S \times S$ is lost forever when going from a tree to the next one. Less immediate is the linear bound in Lemma 3.2.

Lemma 3.1 *For some $i \geq 0$, let e be an edge visible under tree T_i but not appearing in tree T_{i+1}. Then e belongs to none of $T_{i+2}, \ldots, T_k$.*

Proof Let E_i be the set of all edges visible under T_i. Recall that $T_{i+1} = MST(E_i)$. An edge $e = xy$ of E_i does not appear in T_{i+1} because there are paths between x and y which contain only edges of E_i that are shorter than e (by the P-property). Clearly, T_{i+1} has to contain such a (unique) path. Consequently, this path is part of E_{i+1} whose edges T_{i+2} is constructed from, in turn. So, even if not being blocked

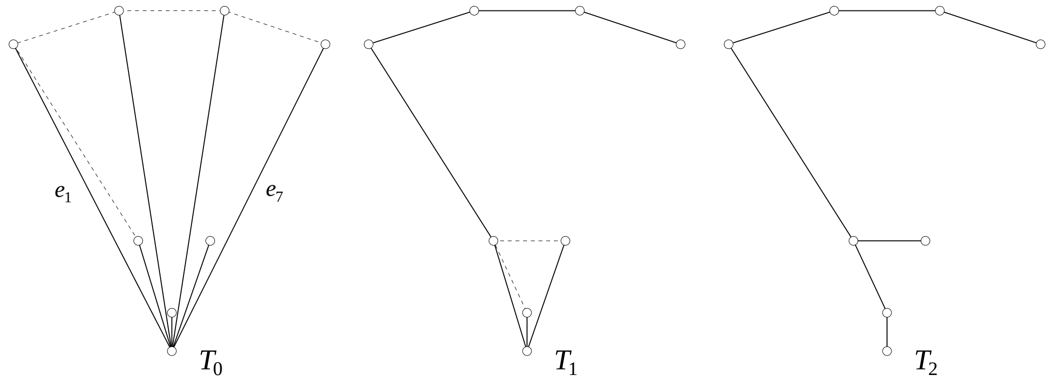

Figure 1: Canonical sequence for a tree on 8 points

by T_{i+1}, edge e will not be present in T_{i+2}. By induction, e can never reappear in further trees of the sequence. □

Lemma 3.2 *The length of a canonical sequence $T_o, \ldots, T_k$ is at most $n-1$.*

Proof We establish an injective mapping from the trees in $T_o, \ldots, T_k$ to the points in the set S.

Suppose $k \geq 2$ and put $j = k-2$. Then the tree T_j contains at least one edge, e_j say, which is D-flippable in $CD(T_j)$ because $CD(T_j) = DT(S)$ otherwise (by Observation 2.1), implying $T_{j+1} = T_{k-1} = MST(DT(S)) = MST(S)$ – a contradiction. Denote with $diam(e_j)$ the circle with diameter e_j. Since e_j is D-flippable, $diam(e_j)$ encloses some point p of S.

Consider the triangle e_jp. If this triangle is empty of points in S then put $p_j = p$ else choose p_j within e_jp such that the triangle e_jp_j is empty. Call the latter triangle Δ_j and let $e_j = xy$. As $p_j \in diam(e_j)$, both edges p_jx and p_jy are shorter than e_j. But not both of them can be visible under the preceding tree T_{j-1}, as $e_j \notin T_j$ by the P-property, otherwise. So there must exist some edge e_{j-1} of T_{j-1} that splits the triangle Δ_j but does not cross e_j. The crucial consequence thereof is that $diam(e_{j-1})$ encloses p_j. So we can repeat the construction above for the triangle $e_{j-1}p_j$. This yields a sequence $\Delta_j, \Delta_{j-1}, \ldots, \Delta_o$ of empty triangles.

Now observe that, for each such $\Delta_i = e_ip_i$, the edge e_i blocks the visibility between point p_i and all the preceding edges $e_{i+1}, \ldots, e_j$. So we can uniquely charge each index i, for $0 \leq i \leq j-1$, to some endpoint of e_{i+1} which does not lie on p_o's side of e_o. Clearly, e_o's endpoints and p_o itself are not among the charged points. But this implies $j \leq n-3$, that is, $k \leq n-1$. □

Let us now turn to a lower bound example. See Figure 1 again, which actually displays a tree $T = T_o$ whose canonical sequence has length $\Theta(\log n)$. All $n-1$ edges of T concur in one point and cover an angle of at most $\frac{\pi}{3}$. Let $e_1, \ldots, e_{n-1}$ be their cyclic order and let $n = 2^m$. Starting with e_1, every other edge has the same length ℓ. For $i = 1, \ldots, m$, starting with e_{2^i} every other of the remaining edges has length $\frac{\ell}{3^i}$. Now observe that the longest edges of T, and only these, are not present in the tree T_1, because their neighbor edges are short enough for the P-property to apply as well as for keeping themselves in T_1. Similarly, the i-longest edges of T are not present in the tree T_i, for $1 < i < m$. Thus the tree undergoes $m - 1 = \log_2 n - 1$ changes.

In fact, this lower bound is asymptotically tight. As will be demonstrated in Section 4, the upper bound can be reduced to $k = O(\log n)$ by more sophisticated methods than being used in the proof of Lemma 3.2. We state a main theorem of this paper.

Theorem 3.1 *Let T be any crossing-free spanning tree of the point set S. The length k of the canonical sequence of T is bounded by $O(\log n)$, where $n = |S|$. This bound is asymptotically optimal.*

Theorem 3.1 reveals that the minimum operation `min` is strong enough to quickly transform any crossing-free spanning tree of S into the global minimum, $MST(S)$. This result seems surprising in view of the fact that consecutive trees during this transformation never cross each other.

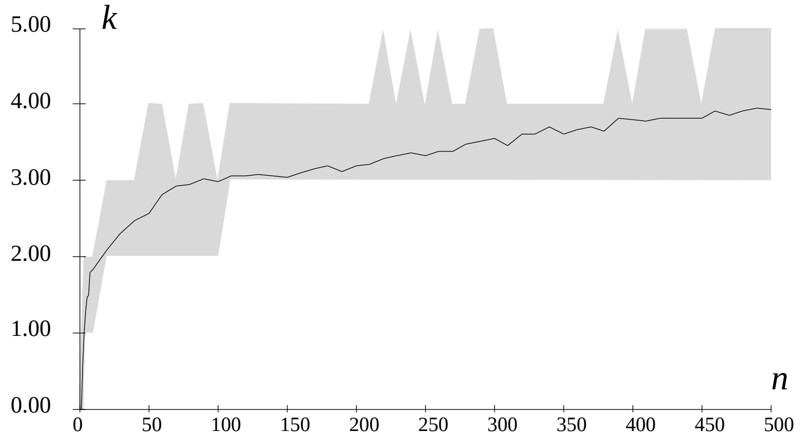

Figure 2: Behavior of k for uniformly distributed points

Figure 2 displays the best, worst, and average empirical behavior of the value of k for n points uniformly distributed in a square. As the starting tree T_o, the star of a randomly chosen point was taken. The extreme and average values of k were obtained from 100 runs. Observe that, even for a quite large number of points, the value of k never exceeded 5.

4 A logarithmic bound

This section settles the question on the length of canonical sequences and proves $k = O(\log n)$. The proof uses the two ingredients below. Here and in the remainder of this section, let t denote an arbitrary but fixed edge of $MST(S)$.

(1) Once being visible, edge t shows up in all further trees of the sequence.

(2) Edge t is blocked by only $O(\log n)$ trees of the sequence.

Whereas (1) is an immediate consequence of Lemma 3.1, (2) is the main property to be shown in the sequel. One of the key observations is that the edges blocking t have a rather restricted position. This is made explicit in Observation 4.1 and Corollary 4.1 below, and is exploited to make Lemma 4.2 (the S-Lemma) hold. Repeated application of the S-Lemma finally yields an exponential growth of edges blocking t, when viewing the canonical sequence in reverse direction $T_k, \ldots, T_o$.

4.1 Aligned edges

For an edge e, let $lune(e)$ denote the intersection of the two disks centered at the endpoints of e and having radius e. As t is a minimum spanning tree edge, no point of S lies in $lune(t)$ by the P-property. In particular, the circle $diam(t)$ with diameter t is empty of points in S.

Observation 4.1 *Consider two or more pairwise non-crossing edges each of which blocks t. Let $\alpha_i \leq \pi$ denote the angle between the i-th and $(i+1)$-st edge (in their order along t). Then $\sum \alpha_i \leq \pi$.*

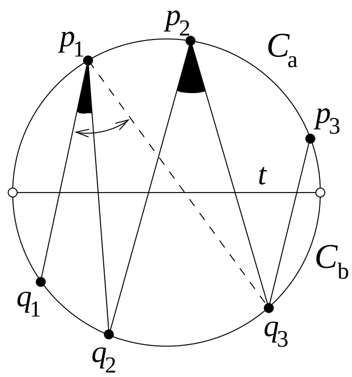

Figure 3: Angles sum up to at most π

Proof Consider, for each blocking edge above, the chord it defines for $diam(t)$. The sum of angles for any set of chords for $diam(t)$ which cross t is maximum if the chords form a path P. Let C_a and C_b, respectively, be the two semi-circles the boundary of $diam(t)$ is divided into by t. Then path P connects points $p_1, \ldots, p_a$ on C_a to points $q_1, \ldots, q_b$ on C_b; see Figure 3. By Thales' theorem, the sum of the angles at the points p_i is just the angle between $p_1 q_1$ and $p_1 q_b$. This angle is at most $\frac{\pi}{2}$ because $p_1 \in C_a$ and $q_1, q_b \in C_b$. Applying a symmetric argument to the angles at the points q_i and adding up gives the lemma. $\square$

Let e and e' be two non-crossing edges both of which cross t. Consider the convex quadrilateral $Q(e, e')$ spanned by $e \cap lune(t)$ and $e' \cap lune(t)$. Call e and e' *α-aligned* (for t) if the two diagonals of $Q(e, e')$ intersect at an angle of at most α; see Figure 4, left side.

Corollary 4.1 *Let $e_1, \ldots, e_m$ be pairwise non-crossing edges which cross t in this order. Then all but $\lceil \frac{\pi}{\alpha} - 1 \rceil$ neighbor pairs (e_i, e_{i+1}) are α-aligned.*

Proof Assume the pair (e_i, e_{i+1}) is not α-aligned. Let d and d' be the two diagonals of $Q(e_i, e_{i+1})$. Define $\beta(d)$ as the angle between e_i and d plus the angle between d and e_{i+1} (same for $\beta(d')$). Then $\beta(d) + \beta(d') > 2\alpha$, because d and d' intersect at an angle of more than α. So $\max\{\beta(d), \beta(d')\} > \alpha$ which proves that (e_i, e_{i+1}) contributes more than α to the sum of angles between possible edges that cross t. But this sum is at most π by Observation 4.1. $\square$

4.2 Separation lemma

It is well known that any two edges of $MST(S)$ that share an endpoint have to form an angle of at least $\frac{\pi}{3}$. The following discussion, in particular the S-lemma below, is a generalization of this fact to constrained minimum spanning trees and $\frac{\pi}{3}$-aligned edges.

Any two non-crossing edges e and e' may be completed to a 4-cycle such that the longest edge added is minimized. Define $Z(e, e')$ to be this cycle. ($Z(e, e')$ degenerates to a 3-cycle if e and e' share an endpoint.) Let $\tilde{Z}(e, e')$ be the region enclosed by $Z(e, e')$. To ease notation, let us write $e > e'$ if edge e is longer than edge e'. We first need a technical lemma.

Lemma 4.1 *Let (e, e') be $\frac{\pi}{3}$-aligned for t. Let R be the portion of $\tilde{Z}(e, e') \setminus lune(t)$ on a fixed side of t. (R may be disconnected.) Then the diameter of R is less than* $\max\{e, e'\}$.

Proof Let E and E', respectively, be the portions of e and e' in R. If the diameter of R is E or E' then we are done. Else the diameter is assumed between points $x \in E$ and $x' \in E'$. Let p and p', respectively, be the endpoint of e and e' not on R's side of t. Further, let s be the point where the diagonals of $Q(e, e')$ intersect; consult Figure 4, left side.

Consider the two triangles pxs and $p'x's$. Their angles at s exceed $\frac{2\pi}{3}$ each, by the $\frac{\pi}{3}$-alignment of (e, e'). Therefore $xs < xp < e$ and $x's < x'p' < e'$. Now consider the triangle $xx's$. Its angle at s is smaller than $\frac{\pi}{3}$. This gives $xx' < \max\{xs, x's\} < \max\{e, e'\}$. □

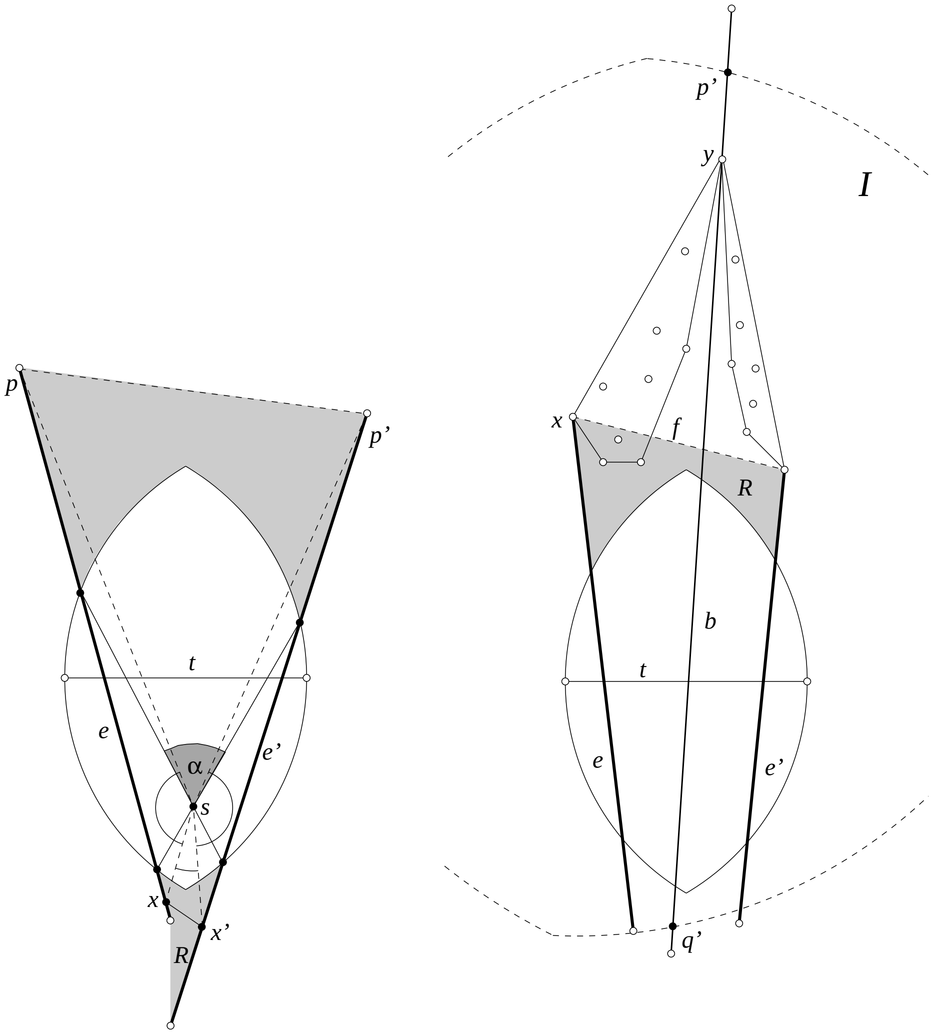

Figure 4: Illustration of the proofs for Lemma 4.1 (left) and the S-lemma (right)

By Lemma 4.1, either e or e' has to be the longest edge in the cycle $Z(e, e')$, provided e and e' are $\frac{\pi}{3}$-aligned. We now will exploit this fact systematically. For expository purposes, colors are assigned to consecutive trees in a canonical sequence. For the moment, let us color an edge e *red* if it belongs to tree T_i and *black* if it belongs to its predecessor T_{i-1}. Note that e may belong to more than one tree and thus may have several colors. Observe further that two edges of the same color or of consecutive colors do not cross.

Lemma 4.2 *(S(eparation)-lemma) Let* (e, e') *be a pair of red and* $\frac{\pi}{3}$*-aligned edges for* t*. There exists some black edge* b *that crosses* t *between* e *and* e'*. If* b *is unique then* $lune(b)$ *covers an edge of* $Z(e, e')$ *other than* e *and* e'*.*

Proof Suppose $e > e'$ without loss of generality. Lemma 4.1 together with the P-property implies that not all edges of $Z(e, e')$ are visible under the black tree because e, the longest edge of $Z(e, e')$, would not be part of the red tree then. Black edges do not cross red ones, so some edge f of $Z(e, e')$ other than e and e' must be blocked by some black edge b. If b crosses t then b does this between e and e'. Else b has an endpoint y in the portion R of $\tilde{Z}(e, e') \setminus lune(t)$ on f's side of t, because b would have to enter $lune(t)$ without leaving it, otherwise. Consider the two convex edge paths $P(y)$ and $P'(y)$ from y to the endpoints of f, defined such that the region bounded by $P(y)$, $P'(y)$, e, e', and t is empty of points in S. By Lemma 4.1, all edges of $P(y)$ and $P'(y)$ are shorter than e. So $P(y)$ or $P'(y)$ must be blocked by some other black edge, which now cannot have an endpoint in R and thus crosses t.

Refer to Figure 4. Assume there is exactly one black edge b that crosses t. Denote by I the intersection of all disks with center in R and radius e. By Lemma 4.1, I covers R. Let y be any endpoint of b. We first prove $y \notin I$. Suppose $y \in I$ and define convex paths $P(y)$ and $P'(y)$ as above. Let g be any edge of these paths. If g has at least one endpoint in R then $g < e$ by definition of I. Else g lies in the triangle yf all whose edges have lengths less than e, such that $g < e$ again. Only black edges crossing t (none except b exist) could block these paths, which gives a contradiction.

As a consequence, b intersects the boundary of I at points p' and q'. Observe $p'q' \geq 2e - d > e$, where d is the diameter of R. Consider any endpoint x of the edge f. Because $x \in R$ we have $xp' \leq e$ and $xq' \leq e$. So in the triangle $xp'q'$, the largest angle is at x. This clearly remains true for the triangle xb. We get $x \in lune(b)$ which proves that $lune(b)$ covers f. □

4.3 Growth of blocking edges

Define $A_i(t)$ to be the number of edges of tree T_i which block the minimum spanning tree edge t. We are now ready to prove that $A_i(t)$ decreases exponentially with i. The main tools used are a repeated application of the S-Lemma, and of the simple fact below which we state without proof.

Observation 4.2 *Let* e *and* p *be an edge and a point, respectively. If* $p \in lune(e)$*, and if some edge* f *cuts across* $lune(e)$ *between* p *and* e*, then* $p \in lune(f)$*.*

Lemma 4.3 $A_{i-3}(t) \geq c_1 \cdot A_i(t) - c_2$*, where* $c_1 > 1$ *and* c_2 *is sufficiently large.*

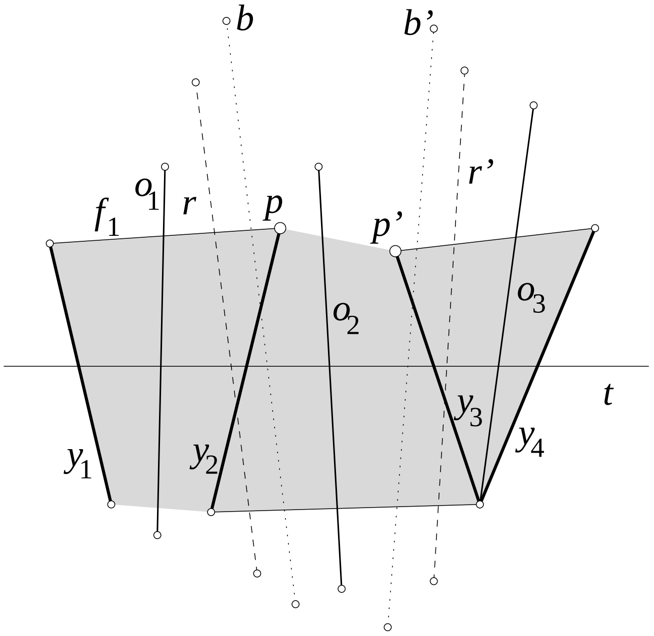

Figure 5: An empty black pair (b, b') for situation ABA

Proof We use the color code yellow, orange, red, and black for the trees $T_i, \ldots, T_{i-3}$ viewed in reverse order of their generation. Assume there are yellow edges $y_1, y_2, \ldots, y_m$ that cross t in this order. By Corollary 4.1, all but two pairs (y_i, y_{i+1}) are $\frac{\pi}{3}$-aligned. So, by the S-lemma, each such pair has some orange edge o_i that intersects t in between. Repeating this argument implies $A_{i-3}(t) \geq A_i(t) - c$ for some constant c. However, a substantial increase of black edges has to occur, as will be proven below.

An uni-colored pair of edges crossing t is said to be *empty* if no edge of the preceding color (i.e., subsequent tree) crosses t in between. Consider any four consecutive yellow edges $y_1, \ldots, y_4$ which define three $\frac{\pi}{3}$-aligned pairs for t. We will show that there exists an empty pair of color orange or red or black which arises between y_1 and y_4. So we gain $c' \cdot A_i(t)$ additional orange or red or black edges, for $c' > 0$. This already verifies the theorem.

Refer to Figure 5 and consider any of the three yellow pairs (y_i, y_{i+1}) above. Recall there exists some orange edge o_i that intersects t in between. If there is another such orange edge, or if y_i or y_{i+1} are orange themselves, then we have found an empty orange pair.

Assume now that, for each pair (y_i, y_{i+1}), the orange edge o_i above is unique. Then, by the S-lemma, $lune(o_i)$ covers an edge f_i of $Z(y_i, y_{i+1})$ other than y_i and y_{i+1}. Assume edge t to be horizontal and call f_i of type A (or type B) if f_i lies above (or below) t. Without loss of generality, f_1 is of type A. Assume first that f_2 is of type A, too. Then f_1 and f_2 have a common endpoint p. (Note that p is an endpoint of y_2.) As $p \in lune(o_1)$, o_1 is the longest edge of the triangle $o_1 p$. So there exists some red edge r that blocks this triangle but does not cross o_1. Edge r enters $lune(o_1)$ and has to leave it again because, by arguments very similar to the proof of the S-lemma, there exists a path of shorter edges between the endpoints of o_1, otherwise. But r cannot leave $lune(o_1)$ without entering $lune(t)$, which implies that r crosses t. Likewise, $p \in lune(o_2)$ implies a red edge r' on o_2's side of p that crosses t. This gives an empty red pair (r, r').

Finally, assume f_2 is of type B. If f_3 also is of type B then we have the case above, so let this edge be of type A. As before, let p be the endpoint of y_2 we know to be in $lune(o_1)$. Let p' be the endpoint of y_3 we know to be in $lune(o_3)$. Then there are red edges r and r' which cut across $lune(o_1)$ and $lune(o_3)$, respectively, and which both cross t. Moreover, we have $p \in lune(r)$ and $p' \in lune(r')$ by Observation 4.2. This, in turn, implies black edges b and b' which cut across $lune(r)$ and $lune(r')$, respectively, and which cross t. If $p = p'$ then b and b' are not identical, which gives the empty black pair (b, b'). On the other hand, $p \neq p'$ and $b = b'$ forces p and p' to lie on different sides of b, giving $p \in lune(y_3)$ (or $p' \in lune(y_2)$). Consequently, o_2 has to separate p from y_3, which implies an additional red edge r^+ that separates p from o_2. We get the empty red pair (r^+, r') in this case. $\square$

Since $A_o(t) < n$, Lemma 4.3 implies $i = O(\log n)$ provided $A_i(t) > \frac{c_2}{c_1 - 1}$. In other words, after a logarithmic number of trees in the canonical sequence, the current tree blocks each edge of $MST(S)$ with at most constantly many edges. We complete the argument by proving that blocking edges then disappear after a constant number of additional steps. Notice that we cannot apply Lemma 3.2 here, because the bound there is linear in n rather than in $A_i(t)$.

Lemma 4.4 $A_{i-3}(t) \geq A_i(t) + 1$.

Proof Let i be an index such that the tree T_i blocks t, and color T_i yellow in the color code used above. It is of convenience to interpret t's endpoints as multi-colored edges of length zero. We further agree on an α-alignment between a point p and an edge e provided p sees e under an angle of at least $\pi - \alpha$. (This is consistent with the limit $f \to p$ for α-aligned edges f and e.) It is easy to check that the S-lemma remains valid in this special case. Whenever the S-lemma applies below, we assume that the induced separating edge is unique. This is no loss of generality as even more edges are gained, otherwise.

Put $m = A_i(t) \geq 1$ and assume first that all $m + 1$ yellow edge pairs (two of which stem from t's endpoints) are $\frac{\pi}{3}$-aligned. By the S-lemma this gives $m + 1$ separating orange edges that cross t. That is, $A_{i-1}(t) \geq m + 1$.

Let now exactly one yellow pair fail to be $\frac{\pi}{3}$-aligned. This gives m orange edges each of which contains one yellow endpoint on either side in its lune, according to the S-lemma. This, in turn, implies red edges separating the yellow endpoints on both sides of each orange edge. We obtain $A_{i-2}(t) \geq m + 1$.

Suppose now the existence of two yellow pairs which are not $\frac{\pi}{3}$-aligned, the maximum possible by Corollary 4.1. Consider the case $m = 1$ first. The corresponding yellow edge y has to cross the diameter of $lune(t)$, as y would be $\frac{\pi}{3}$-aligned with an endpoint of t, otherwise. Consequently, both endpoints of t lie in $lune(y)$. So there exist two orange edges that cross t and separate t's endpoints from y, that is, $A_{i-1}(t) \geq 2$ in this case.

The case $m \geq 2$ is more elaborate. Let (e_1, f_1) and (e_2, f_2) be the two yellow pairs in question and assume first they are not adjacent. Obviously, any two edges which both lie either to the left of e_1, or between f_1 and e_2, or to the right of f_2, are $\frac{\pi}{3}$-aligned. We obtain $m-1$ orange edges and $m-2$ red edges by application of the S-lemma. In addition, for (e_1, f_1) there must exist some red edge r which separates an endpoint p of e_1 from the orange neighbor edge o_1 of e_1, and some red edge r' which separates an endpoint p' of f_1 from the orange neighbor edge o_1' of f_1. If $r \neq r'$ then $A_{i-2}(t) \geq m+1$ because one additional red edge must exist for the other pair (e_2, f_2). If $r = r'$ then observe that r cannot have an endpoint in $lune(o_1)$ or $lune(o_1')$. Therefore, Observation 4.2 implies $p \in lune(r)$ and $p' \in lune(r)$. Finally, recall that each of the $m-2$ red edges gained above from the S-lemma has two orange endpoints in its lune, one on either side. This implies at least $m+1$ black edges which separate points from red edges, that is, $A_{i-3}(t) \geq m+1$.

The argument is similar if (e_1, f_1) and (e_2, f_2) happen to be adjacent, that is, $f_1 = e_2$. Whereas the S-lemma now yields one red edge more, only one additional red edge r instead of two must exist now. Edge r then crosses e_1, f_1, and f_2. Again, r has two yellow endpoints in its lune, which implies the same number of black edges as before. $\square$

Examples confirm what is also revealed in the proof above: $A_i(t)$ does not decrease *monotonically* with i. This is the main reason that makes the proofs somewhat involved. Lemma 4.3 in conjunction with Lemma 4.4 immediatedly implies Theorem 3.1 in Section 3.

Elaboration on the constants in Lemma 4.3 makes explicit that actually the situation AABB... (rather than ABAB...) is worst, in which case we obtain $A_{i-2}(t) \geq \frac{3}{2} \cdot A_i(t) - c'$. This gives the rough estimate $k < 4 \cdot \log_2 n + c$, for a small constant c, and thus is less than a factor 4 off the lower bound example in Section 3.

5 Improving edge moves

We now turn our attention to the edge operations on spanning trees mentioned in the introduction. Recall that $\mathcal{T}_S$ denotes the set of all crossing-free spanning trees of the point set S.

Let A and B be any two spanning trees of S. Suppose that $B = A \cup \{e\} \setminus \{f\}$ is obtained from A by adding some edge $e \notin A$ and removing some edge $f \neq e$ from the induced cycle. This operation is called an *edge move*. We are interested in *planar* edge moves, where $A, B \in \mathcal{T}_S$ and edges e and f do not cross. (This definition is slightly more restrictive than the one used in Hernando et al. [12] who only require the former condition.)

It is known that any two crossing-free spanning trees of S can be transformed into each other by at most $2n - 4$ planar edge moves; see Avis and Fukuda [6]. Let us call an edge move *improving* if it reduces the total tree length. Clearly, any spanning tree of S can be transformed into $MST(S)$ by a sequence of at most $n-1$ improving (but possibly non-planar) edge moves. However, it is in general not possible to transform a tree $T \in \mathcal{T}_S$ into a given and shorter tree $T' \in \mathcal{T}_S$ by means of improving *and* planar edge moves. (A simple counterexample on four points exists.) Interestingly, the situation changes for trees in a canonical sequence.

Lemma 5.1 *Let T_i and T_{i+1} be consecutive trees in some canonical sequence. Then T_i can be transformed into T_{i+1} by at most $n-1$ improving and planar edge moves.*

Proof Initialize a tree $A = T_i$ and modify A as follows. For each edge $e \in T_{i+1} \setminus T_i$ perform an edge move that inserts e and removes the longest edge, f, of the cycle C induced by e in A. Note that C always contains edges longer than e. (Else the edge set $T_i \cup T_{i+1}$, which is visible under T_i, would contain a path of edges shorter than e that connects e's endpoints, so that $e \notin T_{i+1} = \mathtt{min}(T_i)$, by the P-property.) In particular, f is longer than e and, as being the longest edge in C, f cannot belong to T_{i+1}. So the edge move is improving, and it reduces the symmetric difference of A and T_{i+1} by two. The edge move is planar as only edges of the triangulation $CD(T_i)$ are involved. □

Note that no order of insertion has been imposed on the edges in $T_{i+1} \setminus T_i$ in the proof above. This enables us to compute a corresponding sequence of edge moves as follows. Color an edge *red* or *black* or *purple* if it belongs to $T_{i+1} \setminus T_i$ or $T_i \setminus T_{i+1}$ or $T_i \cap T_{i+1}$, respectively. The union of red, black, and purple edges partitions the plane into simply connected faces plus the outer face. Define a graph $G(T_i, T_{i+1})$ whose nodes are these faces and which connects two nodes by an arc iff the corresponding faces share a red edge.

Observation 5.1 *The graph $G(T_i, T_{i+1})$ is a tree.*

Proof Assume nodes x and y in different connected components of $G(T_i, T_{i+1})$. Then, for either x or y, the corresponding face is enclosed by some cycle which contains no red edge. That is, T_i contains a cycle – a contradiction. This implies that $G(T_i, T_{i+1})$ is connected. On the other hand, removal of any arc a disconnects this graph: Consider the red edge e that corresponds to arc a. Edge e induces a cycle in T_i which obviously contains no red edge except e. This implies that $G(T_i, T_{i+1})$ is cycle-free. □

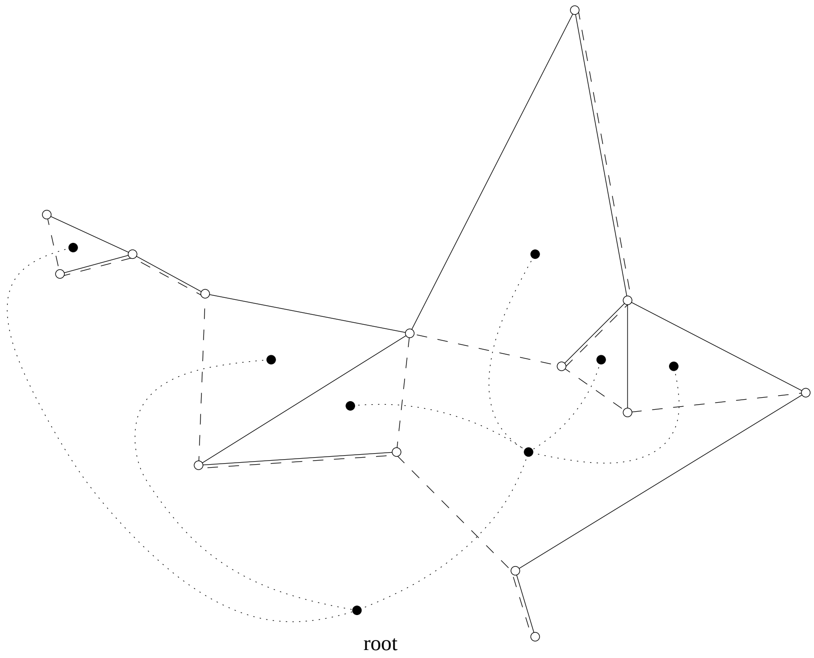

Figure 6: T_i (solid), T_{i+1} (dashed), and the graph $G(T_i, T_{i+1})$

The proof above does not use the fact that $T_{i+1} = \mathtt{min}(T_i)$. So Observation 5.1 remains true if T_{i+1} is replaced by any straight-line graph on S which does not cross T_i.

We now choose the root of the tree $G(T_i, T_{i+1})$ to be the outer face and process the faces in postorder of $G(T_i, T_{i+1})$ as follows; see Figure 6. Let F be the current face. By construction, F contains exactly one red edge e. Let b be the longest black edge of F. Move b to e (this is an improving edge move) and recolor e as purple. If F' is the face that shared edge b with F before the move then join F and F' into a single face F'.

This algorithm can be implemented in $O(n \log n)$ time and $O(n)$ space, by keeping for each face a list of red edges, and a priority queue of black edges organized by decreasing length. From Theorem 3.1 and Lemma 5.1 we conclude:

Theorem 5.1 *Any crossing-free spanning tree of S can be transformed into $MST(S)$ by a sequence of at most $k(n-1) = O(n \log n)$ planar and improving edge moves. Such a sequence can be computed in time $O(n \log^2 n)$ and space $O(n)$.*

To put it another way, the tree graph $\mathcal{TG}_{\mathtt{move}}(S)$, for `move` being the planar and improving edge move, is a directed acyclic graph with unique sink $MST(S)$ and with shortest-path lengths bounded by $O(n \log n)$. The connectivity of $\mathcal{TG}_{\mathtt{move}}(S)$ has also been shown in Bose et al. [7] who use a special kind of edge move, the improving rotation flip, until a spanning tree of $DT(S)$ is reached.

6 Sliding edges

We now investigate a more restrictive edge operation on spanning trees which is both constant-size and local. Consider a tree $A \in \mathcal{T}_S$. A *(planar) edge slide* on A takes some edge $e \in A$ and moves one of its endpoints along some edge adjacent to e in A, without generating any edge crossings. This gives a new edge f and a new tree $B = A \cup \{f\} \setminus \{e\}$ such that $B \in \mathcal{T}_S$. An edge slide is a special kind of planar edge move: B is obtained by closing with f a 3-cycle C in A and by removing e from C, in a way such that A avoids the interior of the triangle C. Intuitively speaking, an edge slide is an edge operation as local as it can be. We are able to prove the following.

Lemma 6.1 *Any planar edge move can be simulated by a sequence of edge slides.*

Proof Consult Figure 7. Consider some planar edge move which transforms trees $A, B \in \mathcal{T}_S$ into each other, and let $B = A \cup \{e_{in}\} \setminus \{e_{out}\}$. Complete $A \cup B$ to an arbitrary triangulation Δ of S. Let R be the portion of Δ enclosed by the cycle C of edges that corresponds to the edge move in question. Then e_{in} and e_{out} lie on the boundary of R. We show that e_{out} can be slided to e_{in} such that each slide has some triangle in R as its corresponding 3-cycle. We use induction on the number m of triangles in R. The case $m = 1$ is trivial, so let $m \geq 2$.

Let $e_{in} = p_1p_2$ and let p_3 be the third vertex of the triangle $t \in R$ based on e_{in}. Suppose first that neither p_1p_3 nor p_2p_3 belongs to tree A. Then A splits $R \setminus \{t\}$ into two parts R_1 and R_2 such that p_1p_3 bounds R_1 and p_2p_3 bounds R_2. Either part contains less than m triangles. Note that e_{out} can only lie on the boundary of one of R_1 and R_2, say R_1. Let e^* be an arbitrary edge on the boundary of R_2.

Now do the following (in this order): (1) Move e_{out} to p_1p_3. This is a planar edge move which, by induction assumption, can be simulated by slides in R_1. (2) By the same argument, it is possible to move e^* to p_2p_3 by sliding in R_2. (3) Now p_1p_3 can be slided along p_2p_3 to e_{in}. (4) Finally, p_2p_3 is moved back to e^* by reversing the sliding operations in (2). In summary, e_{out} has been moved to e_{in} by sliding in R.

It is easy to modify the proof if one of p_1p_3 and p_2p_3 belongs to A. (Inclusion of both edges is ruled out by the assumption $m \geq 2$.) If the included edge just is e_{out} then R_1 is empty and step (1) is unnecessary. Otherwise, R_2 is empty and steps (2) and (4) are unnecessary. □

Recall that any two crossing-free spanning trees of S can be transformed into each other by planar edge moves (see [6], or Theorem 5.1). We immediately obtain from Lemma 6.1:

Theorem 6.1 *Let T and T' be any two crossing-free spanning trees of S. Then T can be transformed into T' by a sequence of edge slides.*

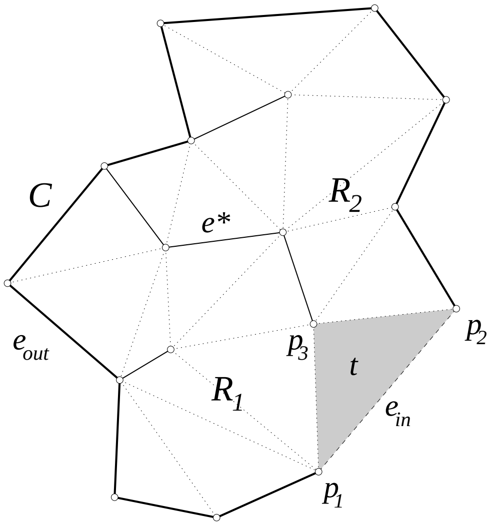

Figure 7: Illustration of the proof of Lemma 6.1

In other words, we have proved that the tree graph $\mathcal{TG}_{\texttt{slide}}(S)$ is connected, for `slide` being the edge slide operation. Note that the slides applied in the proof of Lemma 6.1 need not all be length-improving. Insisting on improving slides in order to reach $MST(S)$ disconnects the tree graph, as can be shown by examples.

The obvious question on the length of the diameter of $\mathcal{TG}_{\texttt{slide}}(S)$ arises. Unfortunately, the proof of Lemma 6.1 translates to an algorithm exponential in n. To see this, let $f(m)$ be the number of triangles that have to be slided across in order to simulate the edge move (where m was the number of triangles in the region R.) Then $f(m) = f(m_1) + 2 \cdot f(m_2) + 1$, for the regions R_1 and R_2 containing m_1 and m_2 triangles, respectively. The factor 2 arises because edge e^* is restored in step (4) of the algorithm in order to make induction apply.

Examples show that the number of edge slides needed is typically much smaller. We plan to elaborate on this question in a separate paper and conjecture:

Conjecture 6.1 *If two trees $T, T' \in \mathcal{T}_S$ are part of the same triangulation of S then they can be transformed into each other by $O(n^2)$ edge slides.*

Conjecture 6.1 combined with Theorem 3.1 would give a diameter of $O(n^2 \log n)$ for the corresponding tree graph.

Acknowledgement We acknowledge discussions on the presented topic with Emo Welzl and Günter Rote. Thanks go to Raimund Leitner for an implementation of canonical sequences, available at http://www.cis.tu-graz.ac.at/igi/oaich/trees/Welcome.html.

References

[1] O.Aichholzer, H.Alt, G.Rote, *Matching shapes with a reference point*, Int'l J. Computational Geometry & Applications 7 (1997), 349-363.

[2] M.Ajtai, V.Chvátal, M.M.Newborn, E.Szemerédi, *Crossing-free subgraphs*, Ann. Discrete Math. 12 (1982), 9-12.

[3] H.Alt, U.Fuchs, G.Rote, G.Weber, *Matching convex shapes with respect to the symmetric difference*, Algorithmica 21 (1998), 89-103.

[4] H.Alt, M.Godau, *Computing the Fréchet distance between two polygonal curves*, Int'l J. Computational Geometry & Applications 5 (1995), 75-91.

[5] F.Aurenhammer, R.Klein, *Voronoi diagrams*, in: Handbook of Computational Geometry, J.-R.Sack and J.Urrutia (eds), Elsevier Science B.V., 2000, 201-290.

[6] D.Avis, K.Fukuda, *Reverse search for enumeration*, Discrete Applied Mathematics 65 (1996), 618-632.

[7] P.Bose, J.Czyzowicz, F.Hurtado, *Rotation flips and non-crossing trees*, Manuscript.

[8] L.P.Chew, *Constrained Delaunay triangulations*, Algorithmica 4 (1989), 97-108.

[9] S.Fortune, *Voronoi diagrams and Delaunay triangulations*, in: Computing in Euclidean Geometry (2nd edition), D.Z.Du and F.K.Hwang (eds.), World Scientific (1995), 225-265.

[10] A.García, M.Noy, J.Tejel, *Lower bounds on the number of crossing free subgraphs of K_n*, Computational Geometry: Theory and Applications 16 (2000), 211-221.

[11] W.Goddard, H.C.Swart, *Distances between graphs under edge operations*, Discrete Mathematics 161 (1996), 121-132.

[12] M.C.Hernando, F.Hurtado, A.Marquez, M.Mora, M.Noy, *Geometric tree graphs of points in convex position*, Discrete Applied Mathematics 93 (1999), 51-66.

[13] M.C.Hernando, M.E.Houle, F.Hurtado, *On local transformation of polygons with visibility properties*, Manuscript, 2000.

[14] C.A.Holzmann, F.Harary, *On the tree graph of a matroid*, SIAM J. Applied Mathematics 22 (1972), 187-193.

[15] G.Károlyi, J.Pach, G.Tóth, *Ramsey-type results for geometric graphs, I*, Discrete & Computational Geometry 18 (1997), 247-255.

[16] J.B.Kruskal, *On the shortest spanning subtree of a graph and the traveling salesman problem*, Proc. Amer. Math. Soc. 7 (1956), 48-50.

[17] D.T.Lee, A.K.Lin, *Generalized Delaunay triangulation for planar graphs*, Discrete & Computational Geometry 1 (1986), 201-217.

[18] G.Liu, *On connectivities of tree graphs*, J. Graph Theory 12 (1988), 453-459.