Edge Operations on Non-Crossing Spanning Trees*

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

1 Introduction

Let S be a set of n points in the Euclidean plane, and consider the set $\mathcal{T}_S$ of all non-crossing spanning trees of S. (For each tree, its edges are straight line segments which pairwise do not cross.) A *tree graph* $\mathcal{TG}_{\mathsf{op}}(S)$ is the graph that has $\mathcal{T}_S$ as its vertex set and that connects vertex (tree) T to vertex T' iff $T' = \mathsf{op}(T)$, where op is some operation that exchanges two tree edges following a specific rule. The existence of a path between two vertices in $\mathcal{TG}_{\mathsf{op}}(S)$ means transformability of the corresponding trees into each other by repeated application 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 combinatorial versions of the problem have been largely studied; see e.g. [6, 8], or for graphs more general than trees, [5]. The geometric version has received attention in recent years [1], especially for the case where S is a set of points in convex position [3, 4]. In this paper, we prove new results on $\mathcal{TG}_{\mathsf{op}}(S)$ for two classical operations op, namely the *(improving) edge move* (Section 4) and the *edge slide* (Section 5). Applications to morphing of trees and to the continuous deformation of sets of line segments seem reasonable. Our results mainly rely on a fact of interest in its own right: Let $MST(S)$ and $DT(S)$ be the minimum spanning tree and the Delaunay triangulation of S, respectively. Then any pair (T, Δ), for $T \in \mathcal{T}_S$ and Δ being $T's$ constrained Delaunay triangulation, can be transformed into the pair $(MST(S), DT(S))$ in a canonical manner. We explain this fact in Sections 2 and 3.

*Part of this work was carried out while the third author visited the Institute for Theoretical Computer Science, Graz University of Technology

[†]Research partially supported by the Austrian Spezialforschungsbereich F003, Optimierung und Kontrolle, and the Austrian Programme for Advanced Research and Technology

[‡]Research partially supported by Proyectos DGES-MEC PB96-0005-C02-02 and DGES-MEC PB98-0933

2 Fixed tree theorem

Let $S \times S$ be the set of all (straight line) edges spanned by S. For a connected edge set $G \subset S \times S$ let $MST(G)$ denote the minimum spanning tree of G.

It is well known that $MST(G)$ can be constructed in a greedy fashion, by adding G's edges in increasing length order, and discarding edges which close a cycle in the graph built so far. We assume throughout that all edges in $S \times S$ have different lengths. Then, by construction, $MST(G)$ is unique for each choice of G. Another property of minimum spanning trees is evident from the greedy construction:

Property 2.1 *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.*

Rather than forcing a spanning tree of S to use only prescribed edges, we may force the tree 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 is called *visible (under G)* otherwise.

Note that G itself thus consists of visible edges. Let V_G be the set of all edges visible under G. Then $MST(V_G)$ is termed the *minimum spanning tree of S constrained by G*, or $MST|_G$ for short.

As an auxiliary structure, we will need the *constrained Delaunay triangulation*, $CD(G)$, of G; see [7]. This triangulation includes 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 $MST|_G = MST(CD(G))$.*

Proof. All edges of $CD(G)$ are visible under G. Therefore it is sufficient to prove that $MST|_G$ is a subset of $CD(G)$. Let e be an edge of $MST|_G$. 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 diametrical circle of e. Then px and py are shorter than e and, as both being visible edges, they yield a path which excludes e from $MST|_G$ by Property 2.1, a contradiction. □

By Lemma 2.1, we can construct $MST|_G$ (in time $O(n \log n)$) by a greedy algorithm that inserts the edges of $CD(G)$ in length-increasing order. 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. Note further that $MST|_G$ is a non-crossing tree, that is, $MST|_G \in \mathcal{T}_S$.

The following assertion on constrained minimum spanning trees is a central result of this paper.

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

The if part is trivial as the fixed tree property $T = MST|_T$ is surely fulfilled by $MST(S)$. We prove the only-if part below.

Consider some triangulation Δ of S, let e be some inner edge of Δ, and denote with $t_1(e)$ and $t_2(e)$ the two triangles of Δ that have e in common. Then the circumcircle of $t_1(e)$ encloses $t_2(e)$ iff the circumcircle of $t_2(e)$ encloses $t_1(e)$. Edge e is termed a *flippable edge of Δ* in this case.

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

Proof. Let edge $e \in \Delta$ be flippable. Then the two angles of $t_1(e)$ and $t_2(e)$ opposite to e, respectively, sum up to at least π. So e is the longest edge in at least one of the triangles $t_1(e)$ and $t_2(e)$. But then $e \notin MST(\Delta)$ by Property 2.1. □

The tree $T = MST(\Delta)$ thus minimizes the number of flippable edges over all spanning trees of Δ. Conversely, Δ minimizes the number of flippable edges in T over all triangulations that contain T. We further observe:

Lemma 2.3 *Let $T \in \mathcal{T}_S$. All flippable edges of $CD(T)$ are in T.*

Hence the tree T maximizes the number of 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 flippable edges in T over all triangulations that contain T.

Let us complete the proof of Theorem 2.1. $T = MST|_T$ is equivalent to $T = MST(CD(T))$ by Lemma 2.1. From Lemma 2.3 we know that all flippable edges of $CD(T)$ are in T. Furthermore, $MST(CD(T))$ contains no such edge, by Lemma 2.2. So $T = MST(CD(T))$ implies that no edge of $CD(T)$ is flippable. Consequently, $CD(T) = DT(S)$ and $T = MST(S)$, as claimed.

3 Canonical sequences

Given an arbitrary non-crossing spanning tree $T \in \mathcal{T}_S$, we now may look at the sequence of trees $T_1 = T$, $T_2 = MST|_{T_1}$, $T_3 = MST|_{T_2}$, 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 these two trees do not cross each other (and of course they both belong to $\mathcal{T}_S$). T_{i+1} is, by definition, minimum in $CD(T_i)$ but T_i is not unless both trees are identical. That is, T_{i+1} has to be shorter than T_i. So the sequence produces pairwise different trees until it 'converges' after a finite number k of steps. Theorem 2.1 tells us that $T_{k+1} = MST(S)$, regardless of the choice of the initial tree T. Let us call $T_1, \ldots, T_{k+1}$ the *canonical sequence* of the non-crossing spanning tree T.

Canonical sequences will prove useful in our study of edge operations on non-crossing spanning trees. Our intention here is to bound the length k of a canonical sequence, with respect to the number n of points in S. An exponential upper bound is obvious from the number of possible spanning trees of S. A bound $k = O(n^2)$ directly follows from the lemma below. Less immediate is a linear bound.

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

Proof. Let V_i be the set of edges visible under T_i. An edge $e = xy$ of V_i does not appear in T_{i+1} because there exist paths between x and y which contain only edges of V_i that are shorter than e (Property 2.1). One such path is part of T_{i+1} and thus of V_{i+1}, whose edges T_{i+2} is constructed from. So, even if e is not blocked by T_{i+1}, it will not be present in T_{i+2}. By induction, e can never appear in further trees of the sequence. □

Lemma 3.2 *For every tree $T \in \mathcal{T}_S$ the length k of its canonical sequence is bounded by $O(n)$, where $n = |S|$.*

Proof. Suppose $k \geq 2$ and put $j = k - 1$. Then T_j contains some edge, e_j say, whose diametrical circle $C(e_j)$ encloses some point p of S. (Otherwise, T_j is part of the Delaunay triangulation $DT(S)$ such that $T_{j+1} = T_k = MST(S)$, a contradiction.)

Let $e_j = xy$ and consider the triangle $e_j p$. If this triangle is empty of points in S then put $p_j = p$ else choose p_j within $e_j p$ such that the triangle $e_j p_j$ is empty. As $p_j \in C(e_j)$, both edges $p_j x$ and $p_j y$ 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 Property 2.1, otherwise. So there must exist some edge e_{j-1} of T_{j-1} that splits the triangle $e_j p_j$ but does not cross e_j. Moreover, $C(e_{j-1})$ encloses p_j. So we can apply the same argument inductively to the triangle $e_{j-1} p_j$ which gives a sequence $e_j, e_{j-1}, \ldots, e_1$ of edges.

We claim that these edges form a planar graph, which finally implies $j = k - 1 = O(n)$. By construction, each such edge e_i completly blocks the visibility between point p_i and all the preceding edges $e_j, \ldots, e_{i+1}$. Therefore the next edge e_{i-1}, which does not cross e_i but splits the triangle $e_i p_i$, crosses none of $e_j, \ldots, e_i$, as claimed. □

We do not consider the bound in Lemma 3.2 to be tight, as the best lower length bound we obtained is $k = \Omega(\log n)$.

Conjecture 3.1 *$O(\log n)$ is an upper bound for the length of canonical sequences.*

4 Improving edge moves

Consider two spanning trees $A, B \in \mathcal{T}_S$. 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 talk of a *crossing-free* edge move if e and f do not cross. It is known that any two non-crossing spanning trees of S can be transformed into each other by at most $2n - 4$ crossing-free edge moves; see [1]. Let us call an edge move *improving* if it reduces the total tree length. Then the following holds:

Lemma 4.1 *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 crossing-free edge moves.*

Proof. Imagine the greedy construction of T_{i+1} by adding edges of $CD(T_i)$ in increasing length order. Initialize a tree $A = T_i$ and modify A during this construction as follows. Whenever an edge $e \in T_{i+1} \setminus T_i$ is encountered, 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 there would exist a path of edges shorter than e in $CD(T_i)$ that connects e's endpoints, so that $e \notin T_{i+1}$ by Property 2.1.) 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 crossing-free as only edges of $CD(T_i)$ are involved. $\square$

Lemma 4.1 in conjunction with Lemma 3.2 implies:

Theorem 4.1 *Any non-crossing spanning tree of S can be transformed into $MST(S)$ by a sequence of at most $k(n-1) = O(n^2)$ crossing-free and improving edge moves.*

To put it another way, the tree graph $\mathcal{TG}_{\mathtt{op}}(S)$, for $\mathtt{op}$ being the crossing-free and improving edge move, is a directed acyclic graph with unique sink $MST(S)$ and with shortest-path lengths bounded by $O(n^2)$. The connectivity of $\mathcal{TG}_{\mathtt{op}}(S)$ has also been shown in [2] who use a special kind of improving edge move, the improving rotation flip.

5 Sliding edges

We now investigate a more restrictive edge operation on spanning trees. Let $A \in \mathcal{T}_S$. An *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 introducing edge crossings and without sweeping across points in S. 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 crossing-free edge move: B is obtained by closing with f a cycle C of length three in A and removing e from C, in a way such that A avoids the interior of the triangle C. We are able to prove the following:

Lemma 5.1 *Any crossing-free edge move can be simulated by a sequence of edge slides.*

Proof. Let $B = A \cup \{e_{in}\} \setminus \{e_{out}\}$ be a crossing-free edge move which transforms tree $A \in \mathcal{T}_S$ into tree $B \in \mathcal{T}_S$. 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 crossing-free edge move which, by induction assumption, can be simulated by sliding 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 void. Otherwise, R_2 is empty and steps (2) and (4) are void. $\square$

As any two non-crossing spanning trees of S can be transformed into each other by crossing-free edge moves (see [1], or Theorem 4.1), we immediately obtain:

Theorem 5.1 *Let T and T' be any two non-crossing spanning trees of S. Then T can be transformed into T' by a sequence of edge slides.*

In other words, the tree graph $\mathcal{TG}_{\mathtt{op}}(S)$ is connected, for $\mathtt{op}$ being the edge slide operation. Note that the slides (or moves) applied in the proof of Theorem 5.1 need not be improving. (In fact, insisting on improving slides disconnects the tree graph.) Concerning the diameter of this graph, we conjecture:

Conjecture 5.1 *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 $O(n)$ edge slides.*

Conjectures 5.1 and 3.1 together would give a diameter of $O(n \log n)$ for the corresponding tree graph.

Acknowledgement We acknowledge discussions with Emo Welzl on Theorem 2.1.

References

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

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

[3] A.García, M.Noy, J.Tejel, *Lower bounds on the number of crossing free subgraphs of K_n*, Computational Geometry: Theory and Applications (to appear).

[4] 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.

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

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

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

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