Reconfiguration of plane trees in convex geometric graphs

Nicolas Bousquet[1], **Lucas De Meyer**[2], **Théo Pierron**[3], **and Alexandra Wesolek**[4]

1 **Université de Lyon, LIRIS, CNRS, Université Claude Bernard Lyon 1, France**
`nicolas.bousquet@cnrs.fr`
2 **Université de Lyon, LIRIS, CNRS, Université Claude Bernard Lyon 1, France**
`lucas.de-meyer@univ-lyon1.fr`
3 **Université de Lyon, LIRIS, CNRS, Université Claude Bernard Lyon 1, France**
`théo.pierron@univ-lyon1.fr`
4 **Technische Universität Berlin, Germany**
`alexandrawesolek@gmail.com`

Abstract

A non-crossing spanning tree of a set of points in the plane is a spanning tree whose edges pairwise do not cross. Avis and Fukuda in 1996 proved that there always exists a flip sequence of length at most $2n-4$ between any pair of non-crossing spanning trees (where n denotes the number of points). Two recent results of Aichholzer et al. and Bousquet et al. improved the upper bound on the length of a flip sequence to $2n - \Omega(\log n)$ and $2n - \Omega(\sqrt{n})$ when the points are in convex position.

We pursue the investigation of the convex case by improving the upper bound by a linear factor for the first time in 30 years. We prove that there always exists a flip sequence between any pair of non-crossing spanning trees T_1, T_2 of length at most cn where $c \approx 1.95$. Our result is actually stronger since we prove that, for any two trees T_1, T_2, there exists a flip sequence from T_1 to T_2 of length at most $c|T_1 \setminus T_2|$.

We give a new lower bound in terms of the symmetric difference by proving that there exists a pair of trees T_1, T_2 such that a minimal flip sequence has length $\frac{5}{3}|T_1 \setminus T_2|$. We generalize this lower bound construction to non-crossing flips (where we close the gap between upper and lower bounds) and rotations.

Related Version arXiv:2310.18518

1 Introduction

Let C be a set of n points in the plane in convex position. A *spanning tree* T on the set of points C is a subset of edges that forms a connected acyclic graph on C. A spanning tree T on C is *non-crossing* if every pair of edges of T (represented by the straight line interval between their endpoints) are pairwise non-crossing. Let us denote by $\mathcal{S}(C)$ the set of all non-crossing spanning trees on the point set C. Let $T \in \mathcal{S}(C)$. A *flip* on T consists of removing an edge e from T and adding another edge f so that the resulting graph $(T \cup f) \setminus e$ is also in $\mathcal{S}(C)$. A *flip sequence* is a sequence of non-crossing spanning trees such that consecutive spanning trees in the sequence differ by exactly one flip.

Avis and Fukuda [2] proved that there always exists a flip sequence between any pair of non-crossing spanning trees of length at most $2n-4$ by showing that there is a star S on C such that T_1 and T_2 can be turned into S with at most $n-2$ flips. In fact, they showed that this flip sequence exists even if the point set C is in general position.

Given two spanning trees T_1, T_2, the *symmetric difference* between T_1 and T_2 is denoted by $\Delta(T_1, T_2) = (T_1 \setminus T_2) \cup (T_2 \setminus T_1)$. We denote by $\delta(T_1, T_2) = |\Delta(T_1, T_2)|/2$ the number of edges in T_1 and not in T_2, which is a trivial lower bound on the length of a flip sequence from

40th European Workshop on Computational Geometry, Ioannina, Greece, March 13–15, 2024.
This is an extended abstract of a presentation given at EuroCG'24. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

T_1 to T_2. The set of spanning trees of a graph G forms a matroid, hence, for any possible pair of spanning trees T_1, T_2, there is a (non geometric) flip sequence that transforms T_1 into T_2 in exactly $\delta(T_1, T_2)$ flips. However, more flips are needed for non-crossing spanning trees. Hernando et al. [5] provided for every n, two non-crossing spanning trees T_1, T_2 on a convex set of n points whose minimal flip sequence needs $\frac{3}{2}n - 5$ flips.

During 30 years, no improvement of the lower or upper bound has been obtained until a recent result of Aichholzer et al. [1]. They showed that the upper bound of Avis and Fukuda can be improved when points are in convex position by proving that there exists a flip sequence between any pair of non-crossing spanning trees of length at most $2n - \Omega(\log n)$. Their result has been further improved by Bousquet et al. [4] who proved that $2n - \Omega(\sqrt{n})$ flips are enough. However, until now, there does not exist any general proof that there always exists a flip sequence of length at most $(2 - \epsilon)n$ for some $\epsilon > 0$. On the other side, Bousquet et al. [4] conjectured that the lower bound of Hernando et al. [5] is essentially tight:

▶ **Conjecture 1.1.** Let C be a set of n points in convex position. There exists a flip sequence between any pair of non-crossing spanning trees of length at most $\frac{3}{2}n$.

One can easily prove that there exists a flip sequence of length at most $2\delta(T_1, T_2)$ between any pair of non-crossing spanning trees in convex position. The improvement of Aichholzer et al. [1] also improves this upper bound by $\Omega(\log(\delta(T_1, T_2)))$. Since in the example of Hernando et al. the intersection is reduced to two edges, one can wonder if Conjecture 1.1 can be extended to the symmetric difference, namely:

▶ **Conjecture 1.2.** Let C be a set of n points in convex position. There exists a flip sequence between any pair of non-crossing spanning trees T_1, T_2 of length at most $\frac{3}{2}\delta(T_1, T_2)$.

Contributions Our main results first consist in (i) improving the best known upper bound to approximatively $1.95 \cdot \delta(T_1, T_2)$, breaking the linear factor 2 of the threshold on the length of a minimal flip sequence (even in terms of the symmetric difference), and (ii) disproving Conjecture 1.2 by proving that the best upper bound factor we can hope for is $\frac{5}{3}$. We complete these results by providing improved upper and lower bounds on the length of transformations in other models of flips, namely non-crossing flips and rotations. In particular, we close the gap between upper and lower bounds in the case of non-crossing flips.

Due to space constraints, we only sketch some proofs of the main results, and the full proofs can be found in a recent preprint [3].

2 Definitions

Let C be a set of points in convex position and T be a non-crossing spanning tree on C. We say two points of a convex set C are *consecutive* if they appear consecutively on the convex hull of C. A *border edge* (for T) is an edge between consecutive points. An edge of T which is not a border edge is called a *chord*. A *hole* of T is a pair of consecutive points that is not a border edge. We will say that we *fill a hole* when we apply a flip where the created edge joins the pair of points of the hole.

One can remark that, for each chord e of T, the line containing e splits the convex hull of C in two parts. A *side of a chord* e is the subset of points of C contained in one of the two closed half-planes defined by the line containing the two endpoints of e (see Figure 1 for an illustration). A *side of* T is a side of a chord e for some $e \in T$. We say an edge (or a hole) is *in a side* A if both its endpoints are in A.

In the following, for every side A of a chord, we will denote by k_A the number of holes in A, which is also the number of chords of T in A. Since T is acyclic, we also have $k_A > 0$. Note that each chord e of T defines two sides A and B whose intersection is exactly the endpoints of e. Moreover, T has exactly $k_A + k_B$ holes.

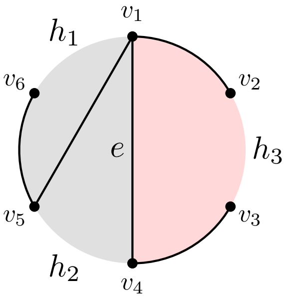

Figure 1 The side A (in grey) of the chord e is is the subset of vertex $\{v_1, v_4, v_5, v_6\}$ and the other side B (in red) of e is $\{v_1, v_4, v_2, v_3\}$. The edges of T in A are the edges v_5v_6, v_1v_5 and v_1v_4. The holes h_1 and h_2 of T are in A and h_3 is in B. So we have $k_A = 2$ and $k_B = 1$.

Let A be a side of a chord e of T. We define the *degree* of a side A in a tree T' as the number of chords of T' crossing e plus twice the number of chords of T' with both endpoints in A (see Figure 2 for an illustration). Note that, if T' has no chords with both endpoints in A, then the degree of A in T' is equal to the number of chords of T' crossing e.

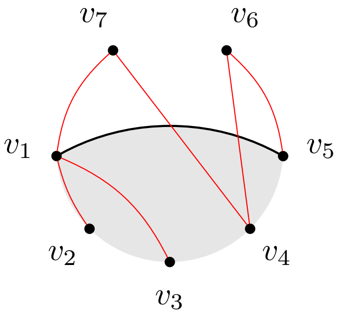

Figure 2 The side A of the edge v_1v_4 highlighted in grey contains v_1, v_2, v_3, v_4 and v_5. The degree of A in the red tree is equal to 4 : v_4v_6 and v_4v_7 cross v_1v_5, and v_1v_3 has both endpoints in A. Note that v_1v_2 is not a chord, thus it does not increase the degree of A in T'.

3 Upper bound

The first result of the paper is to improve the best upper bound of [4] by a linear factor by proving that the following holds:

▶ **Theorem 3.1.** *Let C be a set of n points in convex position. There exists a flip sequence between any pair of non-crossing spanning trees T_1 and T_2 of length at most $c \cdot \delta(T_1, T_2)$ with $c = \frac{1}{12}(22 + \sqrt{2}) \approx 1.95$.*
In particular, there exists a flip sequence of length at most $cn \approx 1.95n$ between any pair of non-crossing spanning trees.

We prove Theorem 3.1 by induction. Let T_I, T_F be two trees on a convex point set C, and assume that Theorem 3.1 holds for every pair of trees T'_I and T'_F, which are either defined on the same set of points and $\delta(T'_I, T'_F) < \delta(T_I, T_F)$ or on a smaller set of points. The goal to prove Theorem 3.1 is to match pairs of chords in T_I, T_F using few flips, i.e. less than c flips per pair of chords matched. Indeed, if we manage to apply at most ck flips on T_I and T_F to obtain T'_I and T'_F with k more edges in common, we get $\delta(T'_I, T'_F) = \delta(T_I, T_F) - k$ and we can conclude by induction.

Basic properties of T_I and T_F Since common chords and non-common border edges can be trivially reduced, we first observe that, if T_I, T_F share a common chord or do not have the same border edges, we can conclude by induction. Hence, we may assume for the rest of the proof that T_I and T_F form a *nice pair* of trees, i.e. the two trees have no common chord and have the same border edges. Note that for a nice pair of trees, every pair of consecutive points is either a common hole or a common border edge. Thus, for a nice pair of trees (T, T'), we will refer to a hole of T or T' simply as a hole.

The rest of the proof describes a transformation from the nice pair (T_I, T_F) to a pair (T^*_I, T^*_F) which matches k pairs of chords of T_I and T_F using at most $c \cdot k$ flips. The main steps of the proof are illustrated in Figure 3. We define a τ-extremal side, which is a side which always exists in a nice pair of trees (and hence which we can also find in T_I and T_F). We then transform a τ-extremal side to what we call a very good side without using too many flips. At the end we observe that in very good sides, we can match the k_A chords in the side using at most $\frac{5}{3}k_A$ flips in total. Analysing the steps then gives us the desired bound.

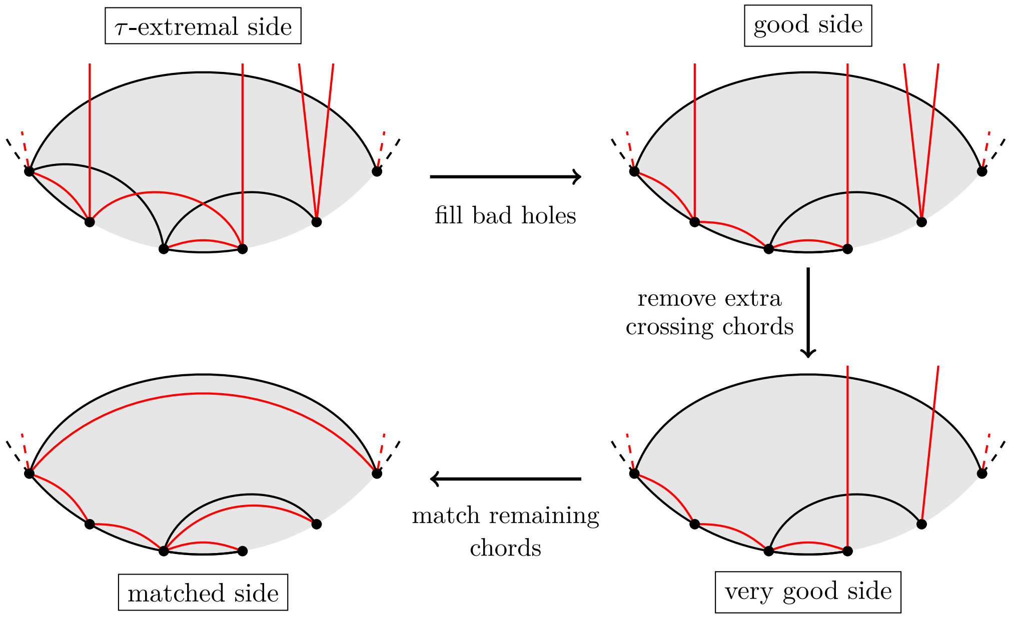

Figure 3 The main steps in the proof of Theorem 3.1. The goal is to match the chords in a side using few steps. (We define τ-extremal, bad, good and very good later).

τ-extremal side Our process starts from a τ-extremal side S of T_I. Let T, T' be a nice pair of trees and $\tau > 2$. We say a side A of a chord e of T is *τ-extremal* for a tree T' if the degree of A in T' is at most $\tau \cdot k_A$, and, for every side $A' \subsetneq A$ of T', the degree of A' in T is more than $\tau \cdot k_{A'}$. To prove that such a side always exists in T_I or T_F (say T_I by symmetry), we start from an arbitrary side and use an iterative greedy argument until we get a τ-extremal side.

The next step is to refine S until we can show it can be matched using few flips. This refinement will start by removing *bad holes*. A hole h in a side A of T is *bad* w.r.t T' if it is also in a side $B \subsetneq A$ of T', see Figure 4. For our process to yield the desired number of flips, we first need to show that S contains few bad holes. In particular, we prove that S contains $m \leqslant \frac{2}{\tau} k_S$ bad holes w.r.t. T_F.

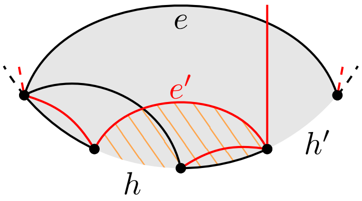

Figure 4 Let T_1 be the black tree and T_2 the red tree. The hole h is a bad hole of the side A (in grey) w.r.t T_2 since it is inside the side of e' included in A.

Refining a τ-extremal side We describe in this paragraph how we refine S into a *very good side*. Let T and T' be a pair of trees. A *good* side A of T with respect to T' is a side of T containing no chord of T' (see Figure 5 for an illustration). A *very good* side A of T (w.r.t. T') is a good side w.r.t T' whose degree in T' is at most k_A.

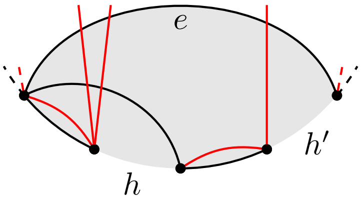

Figure 5 Let T_1 be the black tree and T_2 the red tree. The side A (in grey) of e is a good side of T_1 w.r.t. T_2 since there is no chord of T_2 inside A, but A is not very good w.r.t T_2 since the degree of A in T_2 is $3 > k_A = 2$.

As we already said, we first obtain a good side from S by filling its bad holes with chords of T_I and T_F that have both endpoints in S, therefore we perform $2m$ flips to match m pairs of chords. In the resulting pair of trees (T'_I, T'_F), the size of S is now $k'_S = k_S - m$ and its degree is $d_S - 2m$ with d_S the degree of S in T_F.

Observe that S being good but not very good simply means that there are too many chords crossing the unique chord e on the boundary of S. Hence, we now remove these $d_S - 2m - k'_S$ extra chords crossing e by matching them with a chord of T'_I on a hole which is not in S.

To refine S from a good side into a very good side, we use $2(d_S - 2m - k'_S) = 2(d_S - 2m - k'_S)$ flips to match $d_S - 2m - k'_S = d_S - 2m - k'_S$ pairs of chords.

Very good side Now that S is a very good side, we can match the k'_S chords in S with few flips using the following:

▶ **Lemma 3.2.** *Let T_1 and T_2 be a nice pair of trees, e be a chord of T_1, and A be a very good side of e (w.r.t. T_2). Then, we can match the k_A chords of T_1 in A with chords of T_2 using at most $\frac{5}{3}k_A$ flips in total.*

Bounding the number of flips Let (T_I^*, T_F^*) be the pair of trees obtained after refining S then applying Lemma 3.2 to it. We are now ready to conclude the proof of Theorem 3.1. Our transformation is as follows: first, we transform (T_I, T_F) into (T_I^*, T_F^*) by matching $m+(d_S-2m-k'_S)+k'_S = d_S-m$ pairs using $2m+2(d_S-2m-k'_S)+\frac{5}{3}k'_S = 2d_S-k_S/3-5m/3$ flips. Then, we apply induction on (T_I^*, T_F^*) and get a transformation from T_I^* to T_F^* using at most $c\delta(T_I^*, T_F^*)$ flips.

In order to conclude the proof of Theorem 3.1, we need to make sure that we save enough using Lemma 3.2 to compensate for the expensive refinement process. More precisely, we need that the total number of flips we used to get (T_I^*, T_F^*), namely $2d_S - k_S/3 - 5m/3$, is at most $c(d_S - m)$. Using that $d_S \leqslant \tau k_S$ and that $m \leqslant \frac{2}{\tau}k_S$ (since S is τ-extremal), this boils down to an inequality implying only c and τ, which is satisfied when plugging in the values $\tau = 2 + \sqrt{2}$ and $c = \frac{1}{12}(22 + \sqrt{2})$.

4 Lower bounds

Our second set of results consists in proving stronger lower bounds in terms of the symmetric difference of the two trees. In particular, we disprove Conjecture 1.2:

▶ **Theorem 4.1.** *For every $k > 0$, there exist two trees T_k and T'_k such that $\delta(T_k, T'_k) = 3k$ and every flip sequence between T_k and T'_k has length at least $5k = \frac{5}{3}\delta(T_k, T'_k)$.*

The proof of Theorem 4.1 consists in first providing two spanning trees T_1, T'_1 on 8 vertices for which $\delta(T_1, T'_1) = 3$ and such that the minimal flip sequence between T_1 and T'_1 needs 5 flips (see Figure 6). We can prove that if we glue many instances of (T_1, T'_1) appropriately, we can obtain a similar example with arbitrarily large value of k.

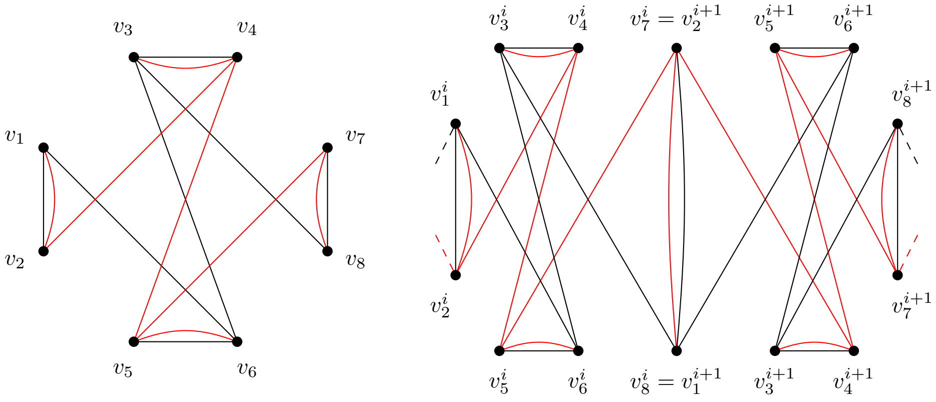

Figure 6 On the left, the tree T_1 in black and the tree T'_1 in red. On the right, an example of two copies of T_1 and T'_1 glued together.

We have not found any example for trees T_1, T_2 for which a flip sequence of length more than $\frac{5}{3}\delta(T_1, T_2)$ is necessary. We therefore leave the following as an open problem:

▶ **Question 4.2.** Let C be a set of points in convex position and T_1, T_2 two non-crossing spanning trees on C. Does there always exist a flip sequence between T_1 and T_2 of length at most $\frac{5}{3}\delta(T_1, T_2)$?

References

1 Oswin Aichholzer, Brad Ballinger, Therese Biedl, Mirela Damian, Erik D Demaine, Matias Korman, Anna Lubiw, Jayson Lynch, Josef Tkadlec, and Yushi Uno. Reconfiguration of non-crossing spanning trees. *arXiv preprint arXiv:2206.03879*, 2022.

2 David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Appl. Math.*, 65(1-3):21–46, 1996. First International Colloquium on Graphs and Optimization (GOI), 1992 (Grimentz).

3 Nicolas Bousquet, Lucas De Meyer, Théo Pierron, and Alexandra Wesolek. Reconfiguration of plane trees in convex geometric graphs. *arXiv preprint arXiv:2310.18518*, 2023.

4 Nicolas Bousquet, Valentin Gledel, Jonathan Narboni, and Théo Pierron. A note on the flip distance between non-crossing spanning trees. *arXiv preprint arXiv:2303.07710*, 2023.

5 M.C. Hernando, F. Hurtado, A. Márquez, M. Mora, and M. Noy. Geometric tree graphs of points in convex position. *Discrete Applied Mathematics*, 93(1):51–66, 1999. 13th European Workshop on Computational Geometry CG '97. URL: `https://www.sciencedirect.com/science/article/pii/S0166218X99000062`, `doi:https://doi.org/10.1016/S0166-218X(99)00006-2`.