

Constrained Flips in Plane Spanning Trees

Oswin Aichholzer 

Graz University of Technology, Austria

Joseph Dorfer 

Graz University of Technology, Austria

Birgit Vogtenhuber 

Graz University of Technology, Austria

Abstract

A flip in a plane spanning tree T is the operation of removing one edge from T and adding another edge such that the resulting structure is again a plane spanning tree. For trees on a set of points in convex position we study two classic types of constrained flips: (1) Compatible flips are flips in which the removed and inserted edge do not cross each other. We relevantly improve the previous upper bound of $2n - O(\sqrt{n})$ on the diameter of the compatible flip graph to $\frac{5n}{3} - O(1)$, by this matching the upper bound for unrestricted flips by Bjerkevik, Kleist, Ueckerdt, and Vogtenhuber [SODA 2025] up to an additive constant of 1. We further show that no shortest compatible flip sequence removes an edge that is already in its target position. Using this so-called happy edge property, we derive a fixed-parameter tractable algorithm to compute the shortest compatible flip sequence between two given trees. (2) Rotations are flips in which the removed and inserted edge share a common vertex. Besides showing that the happy edge property does not hold for rotations, we improve the previous upper bound of $2n - O(1)$ for the diameter of the rotation graph to $\frac{7n}{4} - O(1)$.

2012 ACM Subject Classification Mathematics of computing → Discrete mathematics; Mathematics of computing → Combinatorics; Mathematics of computing → Combinatoric problems

Keywords and phrases Non-crossing spanning trees, Flip Graphs, Diameter, Complexity, Happy edges

Digital Object Identifier 10.4230/LIPIcs.GD.2025.5

Related Version *arXiv Version*: <https://arxiv.org/abs/2508.15520> [2]

Funding *Joseph Dorfer*: Austrian Science Foundation (FWF) 10.55776/DOC183.

1 Introduction

Let S be a finite point set in the plane in general position, that is, no three points lie on a common line. We call S a *convex* point set if no point in S lies in the interior of the convex hull of S . A plane straight-line graph of S is a graph with vertex set S and whose edges are straight line segments between pairs of points of S such that no two edges intersect, except at a common endpoint. All graphs considered in this paper are straight-line graphs. For brevity we will omit the term straight-line. A *plane spanning tree* is a plane connected graph with $n - 1$ edges. In a convex point set, edges appear in two forms. *Convex hull edges* are edges that lie on the boundary of the convex hull of the point set. Edges that are not convex hull edges are *diagonals*.

Flip Graphs of Plane Spanning Trees. A *flip* between two plane spanning trees of S is the operation that removes an edge from a tree and adds another edge such that the resulting structure is again a plane spanning tree. We also denote this operation as an *unrestricted flip*. A constrained version of flips are so-called *compatible flips*, for which the removed edge and the added edge are not allowed to cross. *Rotations* are a restricted subclass of compatible flips for which we additionally require the removed edge and the added edge to share a common



© Oswin Aichholzer, Joseph Dorfer, and Birgit Vogtenhuber;
licensed under Creative Commons License CC-BY 4.0

33rd International Symposium on Graph Drawing and Network Visualization (GD 2025).


Editors: Vida Dujmović and Fabrizio Montecchiani; Article No. 5; pp. 5:1–5:18

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

vertex. The most restricted and local operation is a so called *slide*. A slide is an operation in which the added edge (u, v) and the removed edge (v, w) are again required to share a common vertex, additionally, the triangle $\Delta(u, v, w)$ must not contain any other vertex and the edge (u, w) has to be contained in the current tree. This can visually be interpreted as sliding the edge (u, v) along the edge (u, w) into the edge (v, w) . The (*compatible*) *flip graph* of plane spanning trees of S has as its vertex set all such trees. Two vertices T_1, T_2 of this flip graph are connected with an edge if and only if T_1 can be transformed into T_2 by a single (*compatible*) flip. See Figure 1 for an example of flips in plane spanning trees.




■ **Figure 1** Four flips in a plane spanning tree. (1) unrestricted flip, (2) compatible flip, (3) rotation, (4) slide.

Given an initial tree T_{in} and a target tree T_{tar} , a (*compatible*) *flip sequence* from T_{in} to T_{tar} is a path from T_{in} to T_{tar} in the (*compatible*) flip graph. The (*compatible*) *flip distance* between T_{in} and T_{tar} is the length of a shortest path between T_{in} and T_{tar} in the flip graph. Any (*compatible*) flip sequence of this length is called a *shortest (compatible) flip sequence*. Similarly, we define the *rotation graph*, *rotation sequences* and the *rotation distance*. See Figure 2 for an illustration of the concepts.

For graph reconfiguration problems three natural questions arise: (1) Is the flip graph connected? (2) What is the diameter of the flip graph? (3) What is the complexity of computing shortest flip sequences between two specific configurations?

Connectedness and Diameter. For any n -point set S , the flip graph of plane spanning trees on S is known to be connected and has radius exactly $n - 2$. The same holds for the compatible flip graph and the rotation graph; see [5, 6]. Hence, the diameter always lies between $n - 2$ and $2n - 4$. Since a lower bound for the diameter of $\lfloor \frac{3}{2}n \rfloor - 5$ for convex n -point sets was shown in [12], the flip graph of plane spanning trees on convex point sets has received considerable attention. However, despite a prevailing conjecture that the diameter is at most $\frac{3}{2}n + O(1)$ [10, 14], even improving the upper bound to $2n - o(n)$ took surprisingly long. In the last few years, the upper bound for its diameter was improved to $2n - \Omega(\log(n))$ in [1] and soon after to $2n - \Omega(\sqrt{n})$ in [10]. The upper bound from [10] is constructive and, though not stated explicitly, only uses compatible flips. Hence it also provides an upper bound for the diameter of the according compatible flip graph. In [8], the diameter was bounded from above by cn with $c = \frac{1}{12}(22 + \sqrt{2}) \approx 1.95$, which marked the first linear improvement over the initial bound from [5]. Very recently, a lower bound of $\frac{14}{9}n - O(1)$ and an upper bound of $\frac{5}{3}n - 3$ on the diameter were achieved in [6], where the upper bound in general requires



■ **Figure 2** Flip graph of plane spanning trees on four vertices in convex position. Dashed edges denote flips that are part of the flip graph, but not the compatible flip graph. The dotted lines are compatible flips, but not rotations and the dash dotted lines are rotations but not slides. One pair of two plane spanning trees is marked in orange. Their shortest flip sequence of length one is marked in red (thicker dashed). A shortest compatible flip sequence of length two is marked in blue (thicker).

non-compatible flips. It is also known that any plane spanning tree can be transformed into any other using $\Theta(n^2)$ many slides [4]. For a survey on further flip operations in trees with multiple flips in parallel as well as flips with labeled edges, we refer to [18].

Happy Edges and Complexity. For any graph reconfiguration problem, *happy edges* are edges that lie in both the initial and the target graph. A flip graph fulfills the *happy edge property* if there exists a shortest flip sequence between any two graphs that never flips happy edges. The happy edge property often is a good indication for the complexity of a graph reconfiguration problem. For example, for triangulations of simple polygons [3] and general point sets [16, 19], finding shortest flip sequences is computationally hard. The hardness proofs use counterexamples to the happy edge property as a key ingredient. Conversely, the happy edge property is known to hold for triangulations of convex polygons [20]. Though the complexity of finding shortest flip sequences is still open, the happy edge property yields multiple fixed-parameter tractable algorithms [7, 11, 15, 17]. Moreover, the happy edge property holds for plane perfect matchings of convex point sets and shortest flip sequences can be found in polynomial time [13]. For plane spanning trees in convex position the happy edge property was conjectured to hold [1, 14] for different types of flips. Up to now, it has only been disproved for slides [1].

Contribution and Outline. We will start by improving upper bounds on the flip graph for restricted flip types in convex point sets. In Section 2 we improve the upper bound on the diameter of the compatible flip graph from $2n - \sqrt{n}$ in [10] to $\frac{5}{3}(n - 1)$. Section 3 is dedicated to improving the upper bound on the rotation graph from $2n - O(1)$ in [5] to $\frac{7}{4}(n - 1)$.

Afterwards we deal with the happy edge property. In Section 4 we provide a more refined framework for happy edge properties and in Section 5 we prove the happy edge property for compatible flips in convex point sets and provide a counterexample to the happy edge property for rotations.

Finally, in Section 6 we use variants of the happy edge property to show fixed-parameter tractability of the flip distance for both unrestricted and compatible flips.


Due to space restrictions, full technical details and proofs for statements marked with (★) are given in the arXiv Version of our paper [2].

2 Bounding the Diameter of the Compatible Flip Graph

This section is devoted to improving the upper bound on the diameter of the compatible flip graph. All the currently known improvements to the upper bound of the length of shortest flip sequences for unrestricted flips are based on two different ways to handle edges: They either invest two flips to flip an edge from $T_{\text{in}} \setminus T_{\text{tar}}$ to a convex hull edge and later flip that convex hull edge to an edge of $T_{\text{tar}} \setminus T_{\text{in}}$, or perform a single *perfect flip*, that is, a flip from an edge of $T_{\text{in}} \setminus T_{\text{tar}}$ to an edge of $T_{\text{tar}} \setminus T_{\text{in}}$. Upper bounding the length of a flip sequence is then achieved by lower bounding the number of perfect flips. Note that any such flip sequence induces a pairing of the edges of $T_{\text{in}} \setminus T_{\text{tar}}$ with edges of $T_{\text{tar}} \setminus T_{\text{in}}$.

To obtain our improved bound for the diameter of the compatible flip graph, we build on strategies for bounding the diameter of the flip graph with unrestricted flips from [6]. We give an overview of their proof and highlight the parts where we provide adjustments to make the flip sequence compatible.

The pairing of edges in [6] is obtained as follows: The authors consider the linear representation of a tree T on a convex point set as illustrated in Figure 3. The linear representation is obtained by cutting the convex point set between two consecutive vertices on the convex hull and unfolding the tree T such that its vertices are on a straight line.




■ **Figure 3** A tree T on a convex point set (left) and a linear representation of T (right).

The linear representation induces a linear order $v_1, \dots, v_n$ of the vertices of T . With this linear order, the *length of an edge* $e = (v_i, v_j)$, where $i < j$, is defined as $j - i$. Further, the edge e *covers* all vertices $\{v_i, \dots, v_j\}$ as well as any edge (v_k, v_ℓ) with $k, \ell \in \{i, \dots, j\}$.

A *gap* $\{v_i, v_{i+1}\}$ is an edge between two consecutive vertices in the linear order (which might or might not be present in the tree T). Next, the authors define a bijection (Lemma 3.1 in [6]) between *gaps* and edges of T , where each gap is in bijection with the shortest edge of T that covers that gap. The edges are then partitioned into categories based on how many vertices they share with their corresponding gap (v_i, v_{i+1}) ; see Figure 4 for an illustration.

Concretely, an edge (u, v) is a ...

- *short edge* if $\{u, v\} = \{v_i, v_{i+1}\}$
- *near edge* if $|\{u, v\} \cap \{v_i, v_{i+1}\}| = 1$
- *wide edge* if $|\{u, v\} \cap \{v_i, v_{i+1}\}| = 0$




■ **Figure 4** Edges e_1 , e_3 and e_5 are short edges for gaps g_1 , g_3 and g_5 respectively, e_4 is a near edge for g_4 , and e_2 is a wide edge for g_2 .

The authors show that the number of short edges in any tree T is always larger than the number of wide edges. In a next step, they construct an edge-edge-bijection between the edges in T_{in} and T_{tar} by concatenating the edge-gap-bijection for T_{in} with the gap-edge-bijection of T_{tar} . An edge $e \in T_{\text{in}}$ and its image from the edge-edge-bijection $e' \in T_{\text{tar}}$ form a *pair*. If an edge e is in both T_{in} and T_{tar} , and e is in bijection to itself, then e is not flipped at all. Note that any edge can be flipped to its assigned gap by a single (compatible) flip. The authors use this and the relation between the numbers of short and wide edges in each tree to show that all pairs of edges containing at least one short or wide edge can be flipped using an average of $\frac{3}{2}$ flips per pair. This is done by flipping every edge in such a pair to its corresponding gap and later from its gap to its target edge, which takes two flips minus the number of short edges in the pair.

The remaining pairs then only consist of near edges. These pairs can further be partitioned into three sets as illustrated in Figure 5.

- *above pairs*, where e and e' share a vertex and e is longer than e' .
- *below pairs*, where e and e' share a vertex and e is shorter than e' .
- *crossing pairs*, where e and e' cross, or equivalently, e and e' share a different vertex with their gap.



■ **Figure 5** Illustration of above, below and crossing pairs of edges. Edges of T_{in} are drawn above the vertices, while edges of T_{tar} are drawn below the vertices. The respective gap is denoted by g .

The sets of above, below and crossing pairs are then denoted by A , B , and C , respectively. The authors provide a flip sequence that flips all pairs in A (resp. B , C) directly [6, Lemma 3.2]. Since a direct flip of an above or below pair is compatible, this implies the following lemma.

► **Lemma 1.** *Assume that in the linear representation, the trees T_{in} and T_{tar} consist only of short edges and pairs of edges in A (resp. B). Then there is a compatible flip sequence from T_{in} to T_{tar} of length $|A|$ (resp. $|B|$).*

In contrast, no direct flip between two edges of a pair in C is compatible, since any such pair is crossing. In the arXiv version of our paper [2] we develop an approach to flip these edges such that the resulting flip sequence is compatible while using only one additional flip.

► **Lemma 2 (★).** *Assume that in the linear representation, the trees T_{in} and T_{tar} consist only of short edges and pairs of edges in C . Then there is a compatible flip sequence from T_{in} to T_{tar} of length $|C| + 1$.*

To obtain their upper bound on the flip distance between two trees T_{in} , T_{tar} , the authors cut the original instances T_{in} , T_{tar} along happy diagonals, handle each of the resulting subinstances separately, and then add up the flips needed in each of the subinstances. Let b denote the number of happy convex hull edges, c the number of happy diagonals, and d the cardinality of $T_{in} \setminus T_{tar}$. Using that $b + c + d = n - 1$, their resulting bound [6, Theorem 4.1] is $\frac{5}{3}d + \frac{2}{3}b - \frac{4}{3} \leq \frac{5}{3}(n - 1) - \frac{4}{3}$. With similar calculations, we show the following theorem for compatible flips in the arXiv version of our paper [2].

► **Theorem 3 (★).** *Let T_{in} and T_{tar} be two plane spanning trees in convex point sets with n vertices. Then, there exists a compatible flip sequence from T_{in} to T_{tar} of length at most $\frac{5}{3}d + \frac{2}{3}b + c - \frac{1}{3} \leq \frac{5}{3}(n - 1) - \frac{1}{3}$, where $d = |T_{in} \setminus T_{tar}| = \frac{|T_{in} \Delta T_{tar}|}{2}$ and where b and c denote the number of happy convex hull edges and happy non-convex-hull edges, respectively.*

We remark that our bound with respect to the parameters d , b , and c can be relevantly higher than the one in [6, Theorem 4.1] if the two trees share many happy diagonals. This behaviour was not unexpected due to the following result: In [9, Theorem 1.6], the authors provide examples of pairs of plane spanning trees with (unrestricted) flip distance d and compatible flip distance $2d$. The according tree pairs have $c \approx \frac{n}{2}$ happy diagonals. However, any pair of trees that could reach the upper bounds in [6, Theorem 4.1] and Theorem 3 must have $c = 0$. So while for certain instances the flip distance and the compatible flip distance may differ strongly, extremal examples that determine the diameter are not affected.

3 Bounding the Diameter of the Rotation Graph

Next, we show that for the even more constrained rotation as flip operation, the diameter of the resulting flip graph is linear with leading coefficient less than two.

For this purpose we again construct a pairing of edges such that a perfect flip can be performed for many of them. The challenge that arises, compared to unrestricted or compatible flips, is that rotations require the removed and the added edge to share a vertex. This rules out the gap based pairing from the last section, since in crossing pairs, edges do not share vertices and the same holds for most pairs that contain a wide edge.


► **Lemma 4.** *For a plane spanning tree T on a point set S and a vertex $v \in S$ there exists a bijection $\rho_{T,v}$ between vertices in $S \setminus \{v\}$ and edges in T such that $\rho_{T,v}(w)$ is incident to w for every vertex $w \in S \setminus \{v\}$.*

Proof. We root the tree T at v and orient every edge away from v such that v forms the unique source of the oriented graph; see Figure 6 for an illustration. Now map each vertex in $S \setminus \{v\}$ to its incoming edge. ◀

► **Corollary 5.** Let T_1 and T_2 be two plane spanning trees on a point set S and let v be a vertex in S . Then there exists a pairing of edges from T_1 with edges in T_2 such that:

- Two edges that are paired together share a vertex.
- Different pairs of edges share a different vertex.
- No pair of edges has v as its shared vertex.

Proof. The pairing $\{(\rho_{T_1,v}(w), \rho_{T_2,v}(w)) \mid w \in S \setminus \{v\}\}$ induced by the bijection in Lemma 4 has all the desired properties. ◀




■ **Figure 6** Illustration of the bijection from Lemma 4. All edges are oriented away from the root v . Each vertex in $S \setminus \{v\}$ is in bijection with its unique incoming edge.

For any two given plane spanning trees T_{in} and T_{tar} , we apply the bijection from Corollary 5 with $v := v_n$ to obtain a pairing of the edges of the two trees with all the mentioned properties. For an edge e that is in bijection with a vertex v_i we say that e is *attached* to v_i . Note that each vertex v_i with $i \in \{1, \dots, n-1\}$ has an attached edge from each of T_{in} and T_{tar} , while v_n has no attached vertices.

From this point on we consider the two trees presented in their linear representation, see Figure 3, such that v_n is the rightmost vertex. The pairs of edges $(e_{\text{in}} = (v_i, v_j), e_{\text{tar}} = (v_i, v_k))$, where v_i is the vertex where edges e_{in} and e_{tar} are attached to, can now be subdivided into multiple classes. We say that a pair of edges is ...

- *right-attached above* if $j \leq k < i$
- *left-attached above* if $i < k \leq j$
- *right-attached below* if $k \leq j < i$
- *left-attached below* if $i < j \leq k$
- *diving* if $j < i < k$
- *jumping* if $k < i < j$

We denote the respective sets of pairs in that order by RA , LA , RB , LB , D and J . Additionally, we set $L = LA \cup LB$ and $R = RA \cup RB$. See Figure 7 for an illustration of the six types of edge pairs, as well as the two groups L and R . We remark that happy edges that are left-attached fulfill both, the definition for left-attached above and the one left-attached below. This does not pose any problem, since in the parts of the proof where this is relevant, we will only consider the whole set L . By symmetry, the same can be said about right-attached happy edges. Another noteworthy observation is that we do not handle convex hull edges any differently than diagonals. There might even be happy convex hull edges that belong to different pairs. This will lead to cases in which we rotate happy (convex-hull) edges. We allow this on purpose, since we will see in Section 5 that the happy edge property does not hold for rotations.



■ **Figure 7** Illustration of pairs of edges in the sets RA, RB, R, LA, LB, L, D and J. Edges of T_{in} are drawn above the vertices, while edges of T_{tar} are drawn below the vertices. Each edge points to the vertex which it is attached to.

The new upper bound on the diameter of the rotation graph is described by the following Theorem.

► **Theorem 6.** *There exists a rotation sequence from T_{in} to T_{tar} of length at most $2(n-1) - \max\{|L|, |R|, |J|, |D|\} \leq \frac{7}{4}(n-1)$*

We will use Lemma 7 below to get the edges from three of the sets L , R , J , and D “out of the way”. Then we will show that all edges from the remaining set can be perfectly flipped.

We pair the edges and vertices of T via the bijection ρ_{T, v_n} . Further, for an arbitrary $j \in \{1, \dots, n-1\}$, we assign a gap to every vertex except v_n by the following function (indices taken modulo n).

$$\delta_j : v_k \mapsto \begin{cases} (v_k, v_{k+1}) & \text{if } k > j \\ (v_{k-1}, v_k) & \text{if } k \leq j \end{cases}$$

In other words, we choose one particular gap (v_j, v_{j+1}) that is not assigned to any vertex and the other gaps are assigned such that they contain their assigned vertex.

In the following lemma, j is an arbitrary index, I is a subset of the convex hull edges in T that behave “nicely” with respect to δ_j , and K are all edges of T that are incident to v_n . I^* and K^* are according subsets that are required to be present in the resulting tree T^* .

► **Lemma 7.** *For any $j \in \{1, \dots, n-1\}$, let $I = \{k \in \{1, \dots, n-1\} \mid \rho_{T, v_n}(v_k) = \delta_j(v_k)\}$ and $K = \{k \in \{1, \dots, n-1\} \mid \rho_{T, v_n}(v_k) = (v_n, v_k)\}$. Further, let I^*, K^* be a partition of $\{1, \dots, n-1\}$, that is, $I^* \cap K^* = \emptyset$ and $I^* \cup K^* = \{1, \dots, n-1\}$. Then there exists a rotation sequence from T to T^* of length at most $n-1 - |I \cap I^*| - |K \cap K^*|$ that does not flip edges in $\rho_{T, v_n}(\{v_k \mid k \in (I \cap I^*) \cup (K \cap K^*)\})$ and T^* has the following properties:*


- *For all $k \in I^*$, the tree T^* contains the convex hull edges $\delta_j(v_k)$. That is, the convex hull edges attached to v_k directed away from v_n with respect to the order of vertices along the path on the convex hull from v_n to (v_j, v_{j+1}) ¹.*
- *The tree T^* contains the edges (v_n, v_k) for all $k \in K^*$. That is, all vertices that are attached to v_k are connected directly to v_n via edges that form a fan at v_n .*

¹ We remark that this direction in general differs from the orientation induced by the tree that is used in the proof of Lemma 4.

We refer to Figure 8 for an illustration of Lemma 7 applied to the tree from Figure 6.

Proof. We set $E_0 = \rho_{T,v_n}(\{v_k \mid k \in (I \cap I^*) \cup (K \cap K^*)\})$ and $T_0 = T \setminus E_0$. Now consider the forest $T_{0,vis}$ of edges that are *visible* from v_n in T_0 . These are all edges for which no straight-line segment from v_n to the edge is intersected by any other edge. Further, let v_k be a vertex of degree one in $T_{0,vis}$. Then the only edge incident to v_k in $T_{0,vis}$ is $\rho_{T,v_n}(v_k)$. In particular, if $k \in I^*$ then $\delta_j(v_k)$ does not contain any edge and is visible from v_n , and if $k \in K^*$ then the straight-line segment between v_n and v_k is uncrossed.

We can execute a flip that exchanges the edge $\rho_{T,v_n}(v_k)$ with the edge $\delta_j(v_k)$ (if $k \in I^*$) or (v_n, v_k) if $k \in K^*$ to get a valid tree. We replace $E_0 \leftarrow E_0 \cup \{\delta_j(v_k)\}$ (resp. $E_0 \leftarrow E_0 \cup \{(v_n, v_k)\}$) and $T_0 \leftarrow T_0 \setminus \{\delta_j(v_k)\}$ (resp. $T_0 \leftarrow T_0 \setminus \{(v_n, v_k)\}$) and repeat the process. We repeat until $E_0 = T$. This takes at most $n - 1 - |I \cap I^*| - |K \cap K^*|$ rotations, since the cardinality of E_0 increases by one in every iteration. ◀



■ **Figure 8** Illustration of Lemma 7. Left: The tree from Figure 6 with $n = 8$. Middle: Target structure of Lemma 7 with $K^* = \emptyset$ and chosen index $j = 3$. The edges in $\rho_{T,v_n}(I \cap I^*) = \{(v_8, v_1), (v_1, v_2), (v_7, v_6)\}$ that do not get rotated when rotating from the left to the middle tree are marked in red. Right: Target structure of Lemma 7 with $K^* = \{2, 5\}$ and the same index $j = 3$ as before. The edges in $\rho_{T,v_n}(I \cap I^*) = \{(v_8, v_1), (v_7, v_6)\}$ that do not get rotated when rotating from the left to the right tree are marked in red.


A gap (v_i, v_{i+1}) in a tree T is called *empty* if the edge (v_i, v_{i+1}) is not part of T . Note that in the target tree T^* from Lemma 7, the gap (v_j, v_{j+1}) is empty. Moreover, by setting $K^* = \emptyset$, a flip sequence as in Lemma 7 can be used to rotate all edges to the convex hull such that all convex hull edges that point away from v_n do not have to be rotated.

3.1 The Sets L and R

► **Proposition 8.** *There exists a flip sequence from T_{in} to T_{tar} that uses $2(n - 1) - \max\{|L|, |R|\}$ rotations.*

Without loss of generality, let $|R| > |L|$. For the sets R , we will (1) rotate all edges except the ones in R to the convex hull such that every edge in R has an empty gap to the left of its attachment point, (2) rotate all edges in pairs in R to their target position, and (3) obtain the final tree. Steps (1) and (3) are described in Lemma 9, Step (2) happens in Lemma 10. An illustration of the strategy can be seen in Figure 9.

► **Lemma 9** (*). *There exists a rotation sequence of length at most $n - 1 - |R|$ that rotates all edges except the ones in R to the convex hull such that each edge is rotated into the gap in the convex hull that is left of its assigned vertex. In particular, for an edge $e = (v_i, v_r)$ with $i < r$ that belongs to a pair in R , the gap (v_{r-1}, v_r) in the convex hull is empty. A symmetric argument holds for pairs in L .*



■ **Figure 9** Steps of the proof of Proposition 8. First, rotate all edges except the ones in R (marked in red) to the convex hull, then rotate the pairs in R . At last, rotate all remaining edges from the convex hull back in.

Lemma 9 is an iterative application of Lemma 7. We apply the lemma with $K^* = \emptyset$ to every subtree that results from cutting the tree along edges in R (resp. L). Let T_1 and T_2 be the trees obtained from applying Lemma 9 to T_{in} and T_{tar} .

► **Lemma 10** (*). *There exists a rotation sequence from T_1 to T_2 of length at most $|R|$ (resp. $|L|$).*

The proof of Lemma 10 reduces to the notion of conflict graphs from [6]. In Section 2, left-attached and right attached pairs have been combined to above and below pairs, but we also get an acyclic conflict graph, if we pair them according to whether they are right-attached or left-attached.


3.2 The Sets D and J

► **Proposition 11.** *There exists a flip sequence from T_{in} to T_{tar} that uses $2(n-1) - \max\{|D|, |J|\}$ rotations.*

As the first rotation of the rotation sequence we rotate $\rho_{T_{\text{in}}, v_n}(v_1)$ into the edge (v_1, v_n) . From now on, we say that a *gap is visible from above* if the gap is empty and shares a *face* in the tree (more specifically a face of the drawing of T_{in} combined with the convex hull) with (v_1, v_n) . The concept is illustrated in Figure 10. Note that there is always a unique gap that is visible from above and that it has an edge e_ℓ that is attached to its left vertex and (maybe) an edge e_r that is attached to its right vertex. e_r may not exist if g is the rightmost gap. Further, we show that e_ℓ and e_r share a face with (v_1, v_n) .

► **Lemma 12.** *e_ℓ and e_r share a face with (v_1, v_n) .*


Proof. Assume v_1 and v_n share a face with k additional vertices. Then it also has to share a face with k further edges in order for T to be a tree. Those k edges are exactly the edges that are attached to the k vertices and, thus, include e_ℓ and e_r . ◀



■ **Figure 10** Gap that is visible from above.

► **Lemma 13** (★). *There exists a flip sequence of length at most $n - 1$ that transforms T_{in} into a tree T^* that only contains the following three groups of edges:*

- (A) *some convex hull edges attached to their right vertex,*
- (B) *all target edges of jump pairs, and*
- (C) *some edges that are attached to their left vertex and have as a second vertex a vertex of a jump pair or the vertex v_n .*



■ **Figure 11** An intermediate tree T^* as produced by Lemma 13. The target edges from jump pairs are marked in red. The edges of type C, that from the set K^* in Lemma 14 are marked in blue.

A tree T^* as obtained from Lemma 13 is depicted in Figure 11. The edges that are of special interest for the further application of Lemma 14 are highlighted with colors.

► **Lemma 14** (★). *We can transform T_{tar} into T^* as in Lemma 13 in at most $n - 1 - |J|$ rotations.*

To obtain Lemma 14, we apply Lemma 7 to every face that results from cutting the tree along edges in pairs in J .

Combining the rotation sequence from Lemma 13 (going from T_{in} to T^*) and the reverse rotation sequence from Lemma 14 (going from T^* to T_{tar}), we obtain a rotation sequence between T_{in} and T_{tar} of length at most $2(n - 1) - |J|$, which proves Proposition 11.

4 Relations of Happy Edge Properties

In this section, we introduce a more refined distinction of happy edge properties.

- **Definition 15.** *A graph reconfiguration problem, in which flips exchange edges, fulfills the*
- (weak) happy edge property, *if, from any initial graph G_{in} to any target graph G_{tar} , there exists a shortest flip sequence that does not flip happy edges.*
 - strong happy edge property, *if, from any initial graph G_{in} to any target graph G_{tar} , any shortest flip sequence does not flip happy edges.*
 - perfect flip property *if, whenever we can perform a flip f in G_{in} such that the resulting graph G_1 has one edge more in common with G_{tar} than G_{in} has in common with G_{tar} , then there exists a shortest flip sequence from G_{in} to G_{tar} that has G_1 as its first intermediate graph (that is, the flip sequence starts with f and f is called a perfect flip).*

Further, we provide a framework, that allows us to compare the different properties.

- **Proposition 16** (*). *Let P be a graph reconfiguration problem where flips exchange one edge at a time.*
- (i) *If any flip sequence in P that flips at least one happy edge can be shortened by at least two flips, then P fulfills the perfect flip property.*
 - (ii) *If for any flip in P from G to $(G \setminus \{e_1\}) \cup \{e_2\}$, e_1 and e_2 can never be in the same configuration, then the reverse direction in (i) holds.*

We remark that all the introduced properties may or may not hold for certain graph reconfiguration problems. While the condition in Proposition 16(ii) is not fulfilled for trees, it does hold for example for triangulations.

5 Happy Edges in Plane Spanning Trees on Convex Sets

In [1], the authors formulated the *Weak Happy Edge Conjecture* for trees on convex point sets:

- **Conjecture 17** (Conjecture 17 in [1]). *For any two plane spanning trees T_{in} and T_{tar} on a convex point set, there is a shortest flip sequence from T_{in} to T_{tar} that does not flip happy edges.*


Based on an example from [1] for a different context, we first observe that the perfect flip property holds neither for unrestricted flips nor for compatible flips on trees.

In contrast, we will show in this section that the strong happy edge property does hold for compatible flips on trees. Our first proof ingredient is the following lemma, which is an extension of [1, Proposition 18] for compatible flips.

- **Lemma 18** (*). *Consider any point set S and any two plane spanning trees T_{in} and T_{tar} on S and any shortest compatible flip sequence from T_{in} to T_{tar} . If some edge e is removed and later added back, then some flip during that subsequence must add an edge f that crosses e .*

In [1], a *parking edge* is defined as an edge that appears in a flip sequence and that is not contained in $T_{in} \cup T_{tar}$. A second ingredient of our proof is Lemma 20, which verifies the compatible flip analogue of the following conjecture from [1].


- **Conjecture 19** (Parking Edge Conjecture, Conjecture 21 in [1]). *For any convex point set S and any two plane spanning trees T_{in} and T_{tar} on S , there is a shortest flip sequence from T_{in} to T_{tar} that only uses parking edges from the boundary of the convex hull of S .*



■ **Figure 12** Counterexample to the perfect flip property based on [1, Figure 7]. The top shows the shortest flip sequence. The bottom sequence starts with two perfect flips, but reaches a point where no more perfect flips are possible. Since there are two edges in the tree T^* that both cross two edges from the target tree, at least three additional flips or four additional compatible flips are needed.

We remark that in [1, Claim 22] it is shown that for unrestricted flips Conjecture 19 implies Conjecture 17.

► **Lemma 20** (*). *For any convex point set S and any two plane spanning trees T_{in} and T_{tar} on S , there is a shortest compatible flip sequence from T_{in} to T_{tar} that only uses parking edges from the boundary of the convex hull of S .*



■ **Figure 13** Reordering flips in Lemma 20. The top path is the part of the original flip sequence that contains f , the bottom path is the part of the reordered flip sequence that contains h .

Figure 13 shows the intuition behind the proof: One by one, we replace a parking edge f that is not on the boundary of the convex hull with a parking edge h on the convex hull boundary. To make this replacement possible, we change the order of the flips that happen while f is part of the tree. The edge f splits the convex point set into two sides, say A and B . Flips in either of the two sides can be executed independently from flips in the other side. Assume f is added when flipping from the tree T_{i1-1} to T_{i1} and removed in the flip from T_{i2} to T_{i2+1} . Exactly one of the sides, say side B , of T_{i1-1} entirely contains a path that connects the two endpoints of f . The flips in the other side, say side A , can be executed before f gets added. Afterwards, we close a cycle in side A by adding the convex hull parking edge h and execute all the flips in side B . We conclude the new subsequence of flips by removing h and obtain the tree T_{i2} .

We now show the strong happy edge property for compatible flips on trees.


► **Theorem 21.** *For any convex point set S and any two plane spanning trees T_{in} and T_{tar} on S , any compatible flip sequence from T_{in} to T_{tar} that removes (and adds) a happy edge is at least one step longer than the shortest compatible flip sequence from T_{in} to T_{tar} .*

Proof. Let $T_{in} = T_0, T_1, \dots, T_k = T_{tar}$ be a compatible flip sequence that removes a happy edge e in a flip from T_i to T_{i+1} . By Lemma 20, there exists a shortest compatible flip sequence from $T_{i+1} = T'_{i+1}, T'_{i+2}, \dots, T'_{tar} = T_{tar}$ that only uses parking edges from the convex hull boundary. Since e crosses neither an edge of $T_{in} \cap T_{tar}$ nor one on the convex hull boundary, e crosses no edge in the flip sequence $T_0, \dots, T_i, T'_{i+1}, \dots, T_{tar}$. Hence, by Lemma 18, we can construct a shorter flip sequence that preserves e . ◀

By Proposition 16(i) and the example in Figure 12, Theorem 21 is best possible in the sense that there are flip sequences that flip a happy edge and cannot be shortened by two flips.

When we restrict the flip even further by only allowing edge rotations instead of all compatible exchanges, the happy edge property no longer holds. Also for the more restricted edge slides, the happy edge property does not hold, see [4] for the general case, and [1] for convex point sets.

Figure 14 shows a counterexample to the happy edge property for convex point sets with rotations. Assume the happy edge property would hold. Then the happy edge that passes through the middle of our convex polygon would split the tree into two proper subtrees. In each such subtree, the unhappy edge has a fixed position where it has to go. Since the initial position and the target position of the unhappy edge do not share any vertex, we need at least two rotations to get each unhappy edge from one position to the other. This gives us a minimum length of four for any flip sequence that preserves happy edges.



■ **Figure 14** Counterexample to the happy edge property for the reconfiguration of plane spanning trees in convex point sets using rotations.


The flip sequence in Figure 14 only uses three rotations. Therefore, the happy edge property does not hold for the reconfiguration of plane spanning trees with rotations.

6 Fixed-Parameter Tractable Algorithms

Next, we show that the flip distance k and the compatible flip distance k_c are fixed-parameter tractable in k (resp. k_c). The core ingredient for all proofs is the happy edge property for convex hull edges, which is shown to hold for unrestricted flips in Proposition 18 in [1]. We obtain further improvements on the runtime built on the happy edge property for all edges and the parking edge property.

The upcoming lemma will describe a crucial step in the first two algorithms. Even if we never flip happy convex hull edges, the number of happy edges that are on the convex hull can still be arbitrarily large when compared to the number of unhappy edges. Therefore, we have to reduce that number in order to obtain a reasonable instance size that is bounded by

a function in k . We do this as follows: We mark all vertices that are incident to an unhappy edge that is in the initial configuration or the target configuration. All non-marked vertices are now only incident to convex hull edges. If we have a path between two marked vertices that only travels along convex hull edges incident to only non-marked vertices (except the first and/or last edge that start or end in a marked vertex), then we contract that path into one single edge (if both endpoints are marked) or a single vertex (if only one endpoint is marked). All the non-marked vertices get removed in the process. We call the new reduced instances $T_{in,r}$ and $T_{tar,r}$. The following Lemma states that in this contraction step no information about the length of the flip sequence is lost.



■ **Figure 15** An example of a tree before (left) and after (right) contracting the path from v_2 to v_6 along the convex hull.

► **Lemma 22** (★). *The (compatible) flip distance between T_{in} and T_{tar} has the same length as the flip distance between $T_{in,r}$ and $T_{tar,r}$.*

Now, we are ready to obtain our first FPT-algorithm. We remark that the notation O^* hides terms that are polynomial in n .

► **Theorem 23.** *The flip distance k for unrestricted flips is fixed-parameter tractable in k with runtime $O^*((ck^3)^k)$, where $c = 91.125$. The constant c can be improved to 8 if the happy edge property holds for unrestricted flips.*

Proof. For the unrestricted flip we define *good happy edges*. All happy edges on the convex hull are good happy edges. Further, a happy edge e that does not lie on the convex hull splits the point set into two parts. If one of the two parts contains only happy edges, then e is a good happy edge as well.

▷ **Claim 24** (★). No good happy edge is flipped in any shortest flip sequence.


Now consider the tree $N(T_{in})$ that is obtained from T_{in} from exhaustively contracting happy leaves and paths of happy edges along the convex hull as in Lemma 22. We consider the dual tree of $N(T_{in})$, whose vertices are the faces of $N(T_{in}) \cup \{\text{convex hull edges}\}$ except the outer face and which contains edge between two vertices if and only if their corresponding faces share an edge of $N(T_{in})$.

▷ **Claim 25** (★). Let P be a path in the dual tree that connects two faces containing distinct unhappy edges such that all the edges on traversed faces are happy edges. Then any shortest flip sequence either flips all those happy edges or none of them.

Let there be a total of t unhappy edges. Then there are a total of $t - 1$ paths $P_1, \dots, P_{t-1}$ that fit the form of Claim 25. Let $l_1, \dots, l_{t-1}$ denote their number of happy edges in some order. Any flip sequence of length k gets to spend a total of at most $k - t$ flips on all these

happy edges combined. We enumerate all $2^{t-1} < 2^k$ subsets of $\{P_1, \dots, P_{t-1}\}$ and if their number of happy edges sums up to less than $\frac{k-t}{2}$, then we consider an instance that contains all the faces that contain unhappy edges and the faces along the paths in the subset. Then the total remaining tree contains at most k unhappy edges from T_{in} and k unhappy edges from T_{tar} . That gives a total of at most $2k$ unhappy edges and a total of at most $4k$ vertices. Additionally, we have a total of less than $0.5k$ edges from faces with only happy edges. This gives another $0.5k$ vertices. So we have at most $4.5k$ vertices in our reduced forest. Brute forcing the optimal solution gives us a total of $4.5k - 1$ possible edges to remove and a total of at most $\binom{4.5k}{2}$ positions to add edges for every flip. This gives us a total runtime of $O^*\left(2^k \left(4.5k \binom{4.5k}{2}\right)^k\right) = O^*((91.125k^3)^k)$ based on the number of flip sequences that we have to check.

▷ **Claim 26** (*). If Conjecture 17 is true, then the constant c in Theorem 23 improves to 8.



■ **Figure 16** We split the initial tree T_{in} along its happy edges (colored in green) into three parts, from there we flip every part individually into its corresponding counterpart in T_{tar} .

Figure 16 gives an intuitive idea, where the improvement comes from. We cut the point set along all happy diagonals such that the new instances have all their happy edges on the convex hull. ◀

Next, we prove the existence of a faster FPT-algorithm for compatible flips that is based on the parking edge property.

► **Theorem 27.** *The compatible flip distance k_c is fixed-parameter tractable in k_c with runtime $O^*((2 + \epsilon)k_c^2)^{k_c}$, for any fixed $\epsilon > 0$.*

Proof. We divide the point set into components by cutting along happy edges and then find the shortest flip sequence for every remaining component individually, see Figure 16. Observe that all the happy edges in these components are convex hull edges. If all edges in a component are happy, do not flip any edge in it. By Lemma 20 there exists a shortest flip sequence that only uses edges from T_{in} , T_{tar} or the convex hull.

Now let $\epsilon > 0$ be fixed and $t = \lceil \frac{1}{\epsilon} \rceil$. Let $C_1, \dots, C_m$ be all components with $\leq t$ unhappy edges. Determine the respective flip distances k_j from $T_{\text{in}}^{C_j}$ to $T_{\text{tar}}^{C_j}$ via brute force and discard all the instances. For checking whether the remaining instances have a combined flip distance of $k' = k_c - k_1 - \dots - k_m$ one has to check $((2 + \epsilon)k'^2)^{k'}$ different flip sequences. There are k' choices to remove an edge and every component with $l \geq n$ diagonals has $l + 1 \leq l(1 + \frac{1}{t}) \leq l(1 + \epsilon)$ holes in the convex hull where we can add edges. ◀

7 Conclusion

In this work, we have considered the compatible flip graph and the rotation graph, by improving the upper bound on the respective diameter from $2n - o(n)$ to $\frac{5}{3}(n - 1)$ and to $\frac{7}{4}(n - 1)$, respectively. The constructive upper bounds align with our findings about the happy edge property. On the one hand, we showed that any shortest compatible flip sequence preserves happy edges and that there is a shortest compatible flip sequence that only uses parking edges from the convex hull. The same holds for the compatible flip sequences that we used to obtain our upper bound. On the other hand, shortest rotation sequences may rotate happy edges. This may also happen in our upper bound construction. We further showed an immediate application of happy edge properties in trees, namely, that they can be used to derive fixed-parameter tractable algorithms. Open questions that are related to our work are:

1. Can we close the gap between upper and lower bounds for the flip distance, for each of the flip types? Do the diameters for different flip types differ or do they coincide?
2. Does the happy edge property still hold for flipping plane spanning trees if we drop one of the restrictions that either the point set is convex or the flips are all compatible?
3. What is the time complexity of finding shortest flip sequences for plane spanning trees? Again, this is interesting for each of the flip types.

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. *Journal of Computational Geometry*, 15:224–253, 2024. doi:10.20382/jocg.v15i1a9.
- 2 Oswin Aichholzer, Joseph Dorfer, and Birgit Vogtenhuber. Constrained flips in plane spanning trees, 2025. arXiv:2508.15520.
- 3 Oswin Aichholzer, Wolfgang Mulzer, and Alexander Pilz. Flip distance between triangulations of a simple polygon is NP-complete. *Discrete & Computational Geometry*, 54(2):368–389, June 2015. doi:10.1007/s00454-015-9709-7.
- 4 Oswin Aichholzer and Klaus Reinhardt. A quadratic distance bound on sliding between crossing-free spanning trees. *Computational Geometry*, 37(3):155–161, 2007. Special Issue on the 20th European Workshop on Computational Geometry. doi:10.1016/j.comgeo.2004.12.010.
- 5 David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1):21–46, 1996. First International Colloquium on Graphs and Optimization. doi:10.1016/0166-218X(95)00026-N.
- 6 Håvard Bakke Bjerkevik, Linda Kleist, Torsten Ueckerdt, and Birgit Vogtenhuber. Flipping non-crossing spanning trees. *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2313–2325, 2025. doi:10.1137/1.9781611978322.77.
- 7 Miguel Bosch-Calvo and Steven Kelk. An improved kernel for the flip distance problem on simple convex polygons. *Inf. Process. Lett.*, 182(C), 2023. doi:10.1016/j.ipl.2023.106381.
- 8 Nicolas Bousquet, Lucas de Meyer, Théo Pierron, and Alexandra Wesolek. Reconfiguration of Plane Trees in Convex Geometric Graphs. In Wolfgang Mulzer and Jeff M. Phillips, editors, *40th International Symposium on Computational Geometry (SoCG 2024)*, volume 293 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 22:1–22:17, Dagstuhl, Germany, 2024. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.SoCG.2024.22.
- 9 Nicolas Bousquet, Lucas de Meyer, Théo Pierron, and Alexandra Wesolek. Reconfiguration of plane trees in convex geometric graphs. *CoRR*, 2023. doi:10.48550/arXiv.2310.18518.

- 10 Nicolas Bousquet, Valentin Gledel, Jonathan Narboni, and Théo Pierron. A note on the flip distance between non-crossing spanning trees. *Computing in Geometry and Topology*, 2023. doi:10.57717/cgt.v2i1.36.
- 11 Sean Cleary and Katherine St. John. Rotation distance is fixed-parameter tractable. *Inf. Process. Lett.*, 109(16):918–922, 2009. doi:10.1016/J.IPL.2009.04.023.
- 12 Carmen Hernando, Ferran Hurtado, Alberto Márquez, Mercè Mora, and Marc Noy. Geometric tree graphs of points in convex position. *Discrete Applied Mathematics*, 93(1):51–66, 1999. doi:10.1016/S0166-218X(99)00006-2.
- 13 M. Carmen Hernando, Ferran Hurtado, and Marc Noy. Graphs of non-crossing perfect matchings. *Graphs and Combinatorics*, 18:517–532, January 2002. doi:10.1007/s003730200038.
- 14 Ferran Hurtado. CCCG 2003; Paul Erdős Memorial Lecture: Flip, 2003.
- 15 Haohong Li and Ge Xia. An $O(3.82^k)$ time FPT algorithm for convex flip distance. *Discrete & Computational Geometry*, pages 1–17, 2023.
- 16 Anna Lubiw and Vinayak Pathak. Flip distance between two triangulations of a point set is NP-complete. *Computational Geometry*, 49:17–23, 2015. doi:10.1016/j.comgeo.2014.11.001.
- 17 Joan Lucas. An improved kernel size for rotation distance in binary trees. *Information Processing Letters*, 110:481–484, June 2010. doi:10.1016/j.ipl.2010.04.022.
- 18 Torrie L. Nichols, Alexander Pilz, Csaba D. Tóth, and Ahad N. Zehmakan. Transition operations over plane trees. *Discrete Mathematics*, 343(8):111929, 2020. doi:10.1016/j.disc.2020.111929.
- 19 Alexander Pilz. Flip distance between triangulations of a planar point set is APX-hard. *Computational Geometry*, 47(5):589–604, 2014. doi:10.1016/j.comgeo.2014.01.001.
- 20 Daniel D. Sleator, Robert E. Tarjan, and William P. Thurston. Rotation distance, triangulations, and hyperbolic geometry. *Journal of the American Mathematical Society*, 1(3):647–681, July 1988. doi:10.1090/S0894-0347-1988-0928904-4.