Planar Tree Transformation: Results and Counterexample

Selim G Akl, Kamrul Islam, and Henk Meijer
School of Computing, Queen's University
Kingston, Ontario, Canada K7L 3N6

Abstract

We consider the problem of planar spanning tree transformation in a two-dimensional plane. Given two planar trees T_1 *and* T_2 *drawn on a set* S *of* n *points in general position in the plane, the problem is to transform* T_1 *into* T_2 *by a sequence of simple changes called edge-flips or just flips. A flip is an operation by which one edge* e *of a geometric object is removed and an edge* f $(f \neq e)$ *is inserted such that the resulting object belongs to the same class as the original object. We present two algorithms for planar tree transformations. The first technique is an indirect approach which relies on some 'canonical' tree to obtain such transformation results. It is shown that it takes at most* $2n - m - s - 2$ *flips* $(m, s > 0)$ *which is an improvement over the result in [3]. Second, we present a direct approach which takes at most* $n - 1 + k$ *flips* $(k \geq 0)$ *for such transformation when* S *in convex position and also show results when the points are in general position. We provide cases where the second technique performs an optimal number of flips. A counterexample is given to show that if* $|T_1 \setminus T_2| = k$ *then they cannot be transformed to one another by* k *flips.*

Keywords: Computational geometry, Planar tree, Flip, Transformation, Canonical tree.

1 Introduction

The problem of transforming of a certain class of geometric objects consisting of straight line segments and points in the plane, by applying small changes called flips in the objects, has been studied extensively [3, 8, 7, 9, 11]. Given any two objects for a certain set of points, the question is whether the two objects can be transformed to each other by a sequence of flips and how many such flips are required. A flip can be informally defined as the removal of an edge from, and insertion of another edge to, the object given. Originally, triangulations were investigated with positive results by K. Wagner [13]. Since then the problem has been studied for other classes of planar graphs such as tetrahedrons, linked-edge lists, pseudo-triangulations, planar spanning trees, crossing-free Hamiltonian paths and so on. Algorithms for such transformation as well as lower and upper bounds for achieving transformation results can be found in [1, 3, 7, 9].

One of the best-known results in the case of planar tree transformation is by Avis and Fukuda [3] who showed that for n points in general position every planar tree can be transformed into another planar tree by means of at most $2n - 4$ flips. Later, the bound was slightly improved to $2n - m - s - 2$ $(m, s > 0)$ in [10] which is better when the sum of m and s is greater than 2. Both of these approaches rely on the use of some 'canonical' tree. Informally, a canonical tree is a planar tree that has some particular characteristics such as, for example, all the vertices are directly connected to some vertex called the *root*. Surprisingly, most results related to transformations of different classes of graphs are based on the notion of some 'canonical' form of these graphs, as mentioned in [4]. The main idea of these techniques can be stated as follows: Given two objects A and B of a certain class of graph, the technique is to transform A into some canonical object C of that class by a sequence of transformations. Later, the sequence of transformations that transform B to C is reversed to obtain the desired transformation from A to B. This is an indirect approach. The main problem with this approach is that it takes a long sequence of additional flips to obtain the canonical graph even if the two objects are quite similar or they differ only in a few edges.

1.1 The meta graph and its connectedness

One can define a mathematical model of a class of whether any two geometric objects of a certain class (satisfying some property) are reachable from one to another via a finite sequence of flips. Here we define a graph, called the *meta graph*, that serves as the mathematical model to describe the relationships among the objects of that graph. The meta graph can be defined as the graph having the set of objects in the class as its vertex set, and a pair of vertices is connected by an edge if the objects represented by the vertices differ by a small change. In this way, one can develop adjacencies among nodes of the meta graph.

In most of the cases, generating, examining and establishing relationships among all the objects is not feasible due to the amount of time needed to construct a graph containing such a huge number of vertices. There is a way that makes it possible to establish polynomial-size descriptions among the objects even though the size of the graph is exponential in most of the cases. The procedure of incorporating local modifications on some initial object to visit new objects allows us to study the essential characteristics such as connectedness, reachability, diameter of the graph and so on. The existence of a path between two vertices in the graph means transformability of the corresponding trees represented by the vertices into one another by means of repeated application of the local transformations. The length of the shortest path between two vertices corresponds to the distance between the two trees in terms of the number of transformations.

1.2 Flips in trees

A considerable amount of research in this area has been conducted for general graphs [1, 3, 4, 5]. The meta graph of trees was introduced in [5], in connection with the study of electrical networks. A characteristic of the meta graph, specifically that the graph of trees is Hamiltonian is proved in [5]. Graph-theoretical versions of the problem of geometric object transformations have been largely studied in [12] for tree graphs and was shown that tree graphs have maximum connectivity (a directed graph is said to be maximally connected if it is p-connected and p is the minimum in or out degree of all vertices). Motivated by the question of enumerating the set of all planar spanning trees for points in general position, Avis and Fukuda [3] showed that the corresponding tree graph is connected and has a diameter bounded by $2n-4$. To the best of our knowledge this is the only known result for a general point set. We begin by presenting a bound of $2n-m-s-2$ (where $m, s \geq 1$) which is better ($m+s>2$) or at least equal to theirs. We then propose a second approach that avoids relying on canonical trees. With this approach, trees could be transformed in a more direct manner. We determine the bounds on the number of transformations needed and show that the upper bound on the number of flips using this transformation is $n-1+k$, where k is the number of edges of one planar tree crossed by edges of the other planar tree drawn on S. We provide counterexamples where this direct approach cannot apply. We also show that our algorithm obtains an optimal bound on the number of flips when there are no such intersections.

The organization of the paper is as follows. In section 2, we provide definitions and terminologies that will be used throughout the paper. The first technique for planar tree transformation and the results are presented in Section 3. In Section 4, we present the second approach, analyze its implications on points in general position and in convex position and provide a counterexample. We conclude in Section 5.

2 Preliminaries

Let $S=\{v_0, v_1, v_2, \cdots, v_{n-1}\}$ be a set of n points in general position (no three points are collinear) in the two-dimensional Euclidean plane. A geometric **graph** $G=(V,E)$ consists of a set of vertices V representing the point set S, and an edge set $E=\{(v_i,v_j)|v_i,v_j \in V\}$ connects vertices of V. G is called **planar** if it can be drawn in the plane so that no two edges intersect, except at their common vertex. If $(v_i,v_j)\in E$, then v_i and v_j are **adjacent**.

We designate the root of a tree $T=(V,E)$

(drawn in the plane) as the vertex with minimal $x-$coordinate which is represented by v_0. The vertex with the smallest $y-$coordinate will be designated as the root if there are ties. A **canonical tree** $T^*_{v_0}$ is a planar rooted tree where all vertices are adjacent to v_0, that is, in $T^*_{v_0}$, $v_i \in N(v_0), \forall_i, i \neq 0$. Two vertices v_i and v_j, $v_i \neq v_j$ in an embedding of G are **visible** to each other if the straight line segment $(v_i, v_j) \in E$ between them does not intersect any of the edges in G. A **flip** in a tree T_1 is the operation of removal an edge e and addition of a edge f so that $T_2 = T_1 \backslash \{e\} \cup \{f\}$ is a tree.

Let $\mathcal{T}(S)$ denote the set of all trees of S and the geometric tree graph $T_G(S)$ denote the graph having $\mathcal{T}(S)$ as vertex set. Two trees $T_1, T_2 \in \mathcal{T}(S)$ are adjacent if $T_2 = T_1 \backslash \{e\} \cup \{f\}$ for some edges e and f. In the rest of the paper, it is assumed that a tree is planar unless otherwise mentioned.

3 Tree transformation

Let $T' = (V, E')$ and $T'' = (V, E'')$ be any two trees belonging to $\mathcal{T}(S)$. It is required to construct T'' by applying a sequence of flips one by one to T'. In general, we say that T'' can be transformed from T' by p flips if there is a set of trees $T_0, T_1, \cdots, T_p$ where $T' = T_0$ and $T'' = T_p$ such that T_{t+1} can be obtained from T_t by a single flip. This implies that for any t, T_t and T_{t+1} are adjacent in $T_G(S)$.

Consider Fig. 1 where the tree T'' is obtained from T' by a sequence of transformation.

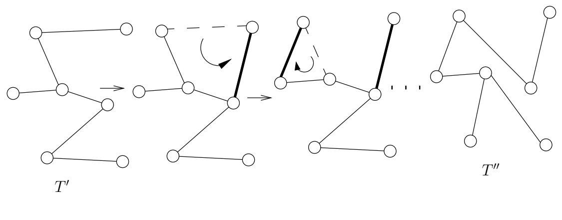

Fig. 1: Transformations (shown with thick edges) applied on T' to construct T''.

We improve the upper bound $(2n-4)$ for tree transformation proposed in [3]. In the following section, it is proved that two trees $T', T'' \in \mathcal{T}(S)$ can be transformed to each other with at most $2n-m-s-2$ flips, where $m > 0$ and $s > 0$ are the number of neighbors of the roots of T' and T'', respectively. Instead of proving directly that a tree can be transformed into another tree, our strategy is to show that it is always possible to transform the given tree into a unique canonical tree. Then, one can perform the reverse operations that transforms the target tree into the same canonical tree.

3.1 Transformation via canonical tree

The following lemma provides a useful direction towards the main result:

Lemma 3.1 *At least one vertex of $T \in \mathcal{T}(S)$ and $T \neq T^*_{v_0}$ is visible to the root, v_0 of T.*

Proof Since the tree $T = (V, E)$ is not in its canonical form, $\exists i$ such that $v_i \notin N(v_0)$. Let N_1 and N_2 be the set of neighbors of v_0 such that $deg(v_p) = 1$ for $v_p \in N_1$ and $deg(v_q) > 1$ for $v_q \in N_2$ with $0 < p, q \leq n-1$. According to the definition of visibility, v_0 can see neither the vertices of N_1 nor those of N_2.

Let ℓ_1 be a line segment connecting v_0 to any of the vertices $v_r \notin N(v_0)$. Denote by $\angle \ell_1$ the angle made by ℓ_1 with v_0 along the vertical line that goes up through v_0. If ℓ_1 does not cross any edge of E, this means v_r is visible to v_0 and the proof is done. So we may assume that ℓ_1 crosses some edges of T. Denote the set of intersected edges E_{int_1}.

Select the edge $e_1 \in E_{int_1}$ whose intersection with ℓ_1 is nearest to v_0.

Now at least one of the vertices of $e_1 = (v_1, v_2)$ must not be a neighbor of v_0. Assume v_1,$v_2 \notin N(v_0)$. In this case, we can arbitrarily choose v_1 or v_2 to connect to v_0. If one of them is a neighbor then we select the other vertex which is not a neighbor of v_0 and connect it to v_0. Let ℓ_2 denote a line segment joining v_0 to one of the vertices v_1 or v_2 . Now check whether the edge ℓ_2 crosses any edges of T. If there is no crossing then the other vertex on ℓ_2 is visible to v_0. If this is not the case, i.e., $|E_{int_2}| > 0$, then choose the closest intersected edge e_2 from E_{int_2} from the root and follow the above procedure of connecting v_0 to the vertex of e_2 which is not neighbor of v_0. Repeatedly applying the procedure described above gives rise to the following cases:

Case 1: Case 1 deals with the situation where $\angle \ell_1 < \angle \ell_2 < \angle \ell_3 \cdots < \angle \ell_a$ (here $a \leq n-2$). Since there is a finite number of vertices in the graph, by following the above procedure we can reach a vertex v_s where the edge ℓ_a whose endpoints are v_s and v_0 will not cross any edges. This ensures us that in the worst case ($a = n-2$) the line segment ℓ_a makes the largest angle joining the vertex v_s to v_0. Then v_s is the vertex visible to v_0.

Case 2: Case 2 is the reverse of case 1 ($\angle\ell_1 > \angle\ell_2 > \angle\ell_3 \cdots > \angle\ell_a$) and can be similarly resolved.

Case 3: In Case 3 we consider the following:

There exists some i such that (i) $\angle\ell_i < \angle\ell_{i+1} > \angle\ell_{i+2}$ or

(ii) $\angle\ell_i > \angle\ell_{i+1} < \angle\ell_{i+2}$

Without loss of generality consider (i) $\angle\ell_i < \angle\ell_{i+1} > \angle\ell_{i+2}$. Now the line segment ℓ_{i+2} must lie between ℓ_{i+1} and ℓ_i and e_i. Let v_i and v_{i+1} be the end vertices of the line segments ℓ_i and ℓ_{i+1} respectively. Any ℓ_j (for $j > i+2$) must lie in triangle formed by the line segments ℓ_i,ℓ_{i+1} and e_i because at each step we are taking the closest intersected edges to v_0. Also this implies that at every iteration the number of vertices (falling inside the triangles) to be considered for visibility to v_0 are reduced as the area of successive triangles decreases gradually. Since the number of vertices are finite and the number is becoming smaller and smaller at each step, we must end up with a single vertex v_s inside the smallest of such triangles (in the worst case) that will be visible to v_0.

Fig. 2. shows such a scenario.

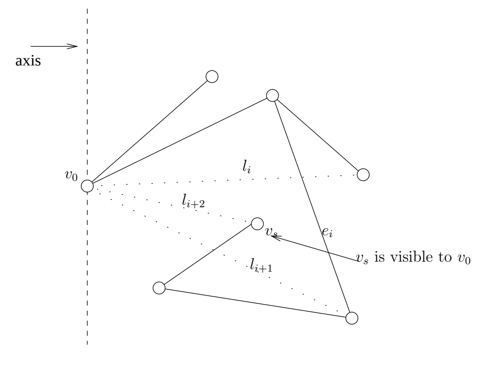

Fig. 2: A tree T depicting case 3. Only the solid line segments represent the edges of the tree.

Lemma 3.2 *Any tree $T \in \mathcal{T}(S)$ and $T \neq T^*_{v_0}$ can be transformed into its canonical form $T^*_{v_0}$ through flipping of edges at most $(n-m-1)$ steps, where $m = |N(v_0)|$.*

Proof Begin with any vertex visible to v_0. Assume v_0 can see v_j. Since adding a line segment (v_0, v_j) makes a unique cycle $v_0 \cdots v_i v_j v_0$ in T, break the cycle by eliminating the edge (v_i, v_j). This flip increases the degree of v_0 by one. Then continue the same process with another vertex visible to v_0 until $\forall v_i \in N(v_0)$ and $v_i \neq v_0$. An illustration is Fig. 3 shows the transformation of T into canonical tree $T^*_{v_0}$ through successive flips.

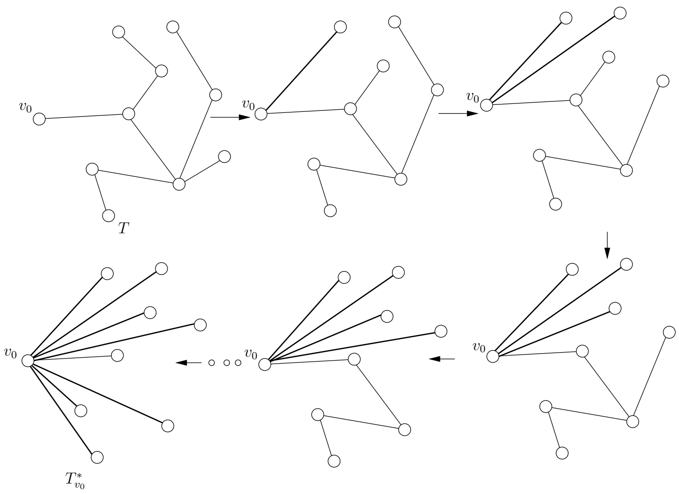

Fig. 3: Showing the transformation of tree T into canonical tree $T^*_{v_0}$. Thick edges represent the edges that are flipped.

Since for any tree there can be $n-m-1$ edges possible which are not incident to v_0, the total number of flips required to produce $T^*_{v_0}$ is $n-m-1$.

Given the tree, T in Fig. 3 with $n = 10$ and $m = 1$ note that the number of flips to obtain $T^*_{v_0}$ from T is, $n - m - 1 = 10 - 1 - 1 = 8$.

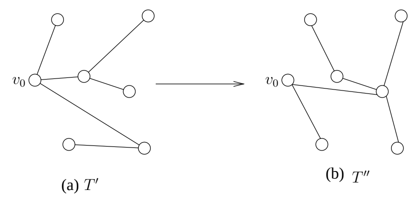

Fig. 4: Showing an example for the Corollary 4.

Now we will prove one of the main results of this paper, i.e., two trees $T', T'' \in \mathcal{T}(S)$ can be transformed into one another with only a linear number of flips. Instead of directly proving that a tree can be transformed into another tree our strategy is that we transform T' into the unique canonical tree $T^*_{v_0}$. And then we can reverse the operations that transform T'' into $T^*_{v_0}$. This is shown in the following lemma:

Lemma 3.3 *Two trees $T', T'' \in \mathcal{T}(S)$ can be transformed to each other with at most $2n-m-s-2$ flips where $m \geq 1$ and $s \geq 1$ are the number of neighbors of the roots of T' and T'' respectively.*

Proof Given the tree T', one can transform it into canonical tree $T^*_{v_0}$ using at most $n-m-1$ steps where m is the number of neighbors of v_0. Similarly, the tree T'' can be transformed into canonical tree $T^*_{v_0}$ with at most $n-s-1$ flips where s is the number of neighbors of v_0 in T''. Since the canonical trees for T' and T'' are same, it takes $2n-m-s-2$ flips to transform T' into T''. This completes the proof. ∎

Corollary 3.4 *The number of flips required to transform a tree T' into another tree T'', where $T', T'' \in \mathcal{T}(P)$ is $2n-p-2$ where $p = max_{v \in CH(V)}$(degree v in T' + degree v in T'') where $CH(V)$ is the convex hull of V.*

Suppose that T' and T'' are two trees with $n = 7$ as shown in Fig. 4 where the neighbors of the roots of T' and T'' are 3 and 2 respectively. It takes 3 flips for T' to be modified to $T^*_{v_0}$ and 4 flips for T'' to be transformed to $T^*_{v_0}$. Hence, we need $2n-m-s-2 = 14-3-2-2 = 7$ flips to transform T' to T''.

The above result can be expressed in terms of the meta graph $T_G(S)$ whose vertices are the non-crossing set of trees, $\mathcal{T}(S)$. Two vertices of the meta graph are adjacent if the trees represented by them can be obtained from the other by a single edge replacement. According to the proof of the above lemma, we can state that $T_G(S)$ is connected and the diameter of $T_G(S)$ which is the shortest distance between any two furthest vertices, can be at most $2n-m-s-2$. We improve the result of [3] and show that our result produces a better result when the values for m or s are greater than one. Even in the worst case when both the parameters m and s are equal to 1 our result is as good as that in [3].

4 Direct tree transformation

In the following, we outline the main idea of our algorithm which does not rely on any form of canonical tree but obtains the desired transformation.

We draw two trees T' and T'' on S in the plane and obtain the graph $G = (V, E' \cup E'')$ where E' and E'' denote the edge-sets of T' and T'', respectively. Let $G_0 = G = (V, E_0 \cup E'')$ (where $E_0 = E'$). If there are p flips that transform T' into T'', our idea is to apply the sequence of flips on the edges of E' on G such that the resulting graphs are represented by $G_1 = (V, E_1 \cup E'')$, $G_2 = (V, E_2 \cup E'')$, $G_3 = (V, E_3 \cup E''), \cdots, G_p = T'' = (V, E'')$ where G_{i+1} is obtained from G_i by a single flip. In $G_i = (V, E_i \cup E'')$, E_i represents the edge set of $T_i = (V, E_i)$ being transformed into T''. Note that after the pth flip, the graph G_p turns into tree T'', since we expect that as flips are applied on the edges of T', gradually T' is turned into T'' and each instance of the intermediate trees $T_i = (V, E_i)$ along with $T'' = (V, E'')$ is reflected in G_p. In other words, we remember the order and the set of flips carried on G_i to produce G_p, then we apply these sequence of flips on T' in order to obtain T''.

To identify the edges of E_i from the edges of E'' in G_i, we color them with different colors. Edges $(u, v) \in E_i \backslash E''$ are colored in *red*, edges $(u', v') \in E'' \backslash E_i$ in *blue*, and edges $(u'', v'') \in E_i \cap E''$ in *purple*. Observe that, in graph G_i, only red edges can intersect blue edges and there will be no intersections between red and purple or blue and purple edges since T' and T'' are planar. If a red edge is intersected by one or more blue edges, then we call it an *intersected red edge*. We count the total number of such intersected red edges after forming $G_0 = G = (V, E' \cup E'')$ at the beginning of our algorithm and denote it by k. We use intersections and crossings interchangeably by which we mean two edges intersect (cross) at exactly one point but not at an endpoint.

Lemma 4.1 *Suppose $G_i = (V, E_i \cup E'')$ is not planar. The removal of an intersected red edge, $e \in E_i \backslash E''$ from G_i splits E_i into two edge sets E'_i (and vertex set V'_i) and E''_i (and vertex set V''_i). Assume $CH(V) \cap V'_i \neq \emptyset$ and $CH(V) \cap V''_i \neq \emptyset$ where $CH(V)$ is the convex hull of V. There exists an edge $f \in V \times V$ such that $G_{i+1} = (V, E_i \backslash \{e\} \cup E'' \cup \{f\})$*

where f is an edge between an incident vertex of some edge E_i' and an incident vertex of some edge E_i'' such that f does not cross any edges of G_{i+1}.

Proof Begin by removing an intersected red edge $e = (v_k, v_\ell)$ from $G_i = (V, E_i \cup E'')$ and obtain two edge sets E_i' and E_i''. Let $V_i' \subset V$ and $V_i'' \subset V$ denote the incident vertices of E_i' and E_i'' respectively. The aim is to connect $v_i \in V_i'$ and $v_j \in V_i''$ ($(v_i, v_j) = f$) so that the edge f does not cross any edges in G_{i+1}.

Color the vertices of V_i' and V_i'' black and white, respectively. It suffices now to connect a black vertex to a white one without yielding any crossing. Select any of the black vertices v_i, $v_i \in V_i'$, at random on the convex hull of $CH(V)$ and start walking along the boundary of $CH(V)$ in some order. Once a walk is complete (that is, we reach the same vertex from which we started), we get a sequence of white and black vertices. We can insert an edge f by connecting any two consecutive white and black vertices in the sequence. This does not generate any crossing in G_{i+1} since the edge is drawn on the boundary of $CH(V)$. ∎

An illustration of the above lemma is shown in Fig.5.

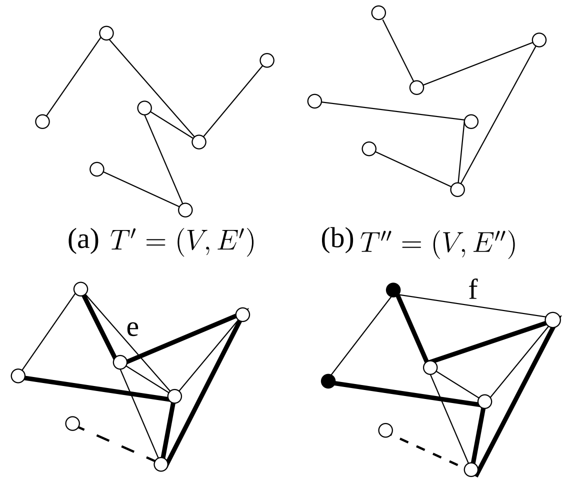

(a) $T' = (V, E')$ (b) $T'' = (V, E'')$ (c) $G_0 = (V, E' \cup E'')$ (d) $G_1 = (V, E_1 \cup E'')$

Fig. 5: (a) Tree $T' = (V, E')$ is to be transformed into (b) tree $T'' = (V, E'')$. (c) Shows the graph G formed by $E' \cup E''$ where thin edges represent red edges, thick edges denote blue edges and dashed edges denote purple edges. (b) Edge e is removed and edge f is inserted without making any crossing in the graph. Black vertices belong to V_1 while the rest belong to V_2.

We identify a case where we can obtain an optimal number of flips for the desired transformation; this is given in the following lemma.

Lemma 4.2 *Any tree $T' = (V, E')$ can be transformed into another tree $T'' = (V, E'')$ with at most $n-1$ flips when the number of intersected red edges is zero.*

Proof Obtain the graph $G_0 = (V, E_0 \cup E'')$, where $E_0 = E'$. Since there are no intersected red edges, G_0 is planar. Begin in the following way. At each step, remove an arbitrary red edge $(u, v) \in E_i \backslash E''$ from G_i and color the vertices black and white as in the proof of Lemma 4.1. Insert a purple edge between a black and a white vertex, otherwise the purple edge will make a cycle if the two incident vertices are of the same color. Since at every step a flip is carried out, we get a new graph $G_{i+1} = (V, E_{i+1} \cup E'')$, where $|E_{i+1} \cap E''| = |E_i \cap E''| + 1$ $(0 \leq i < p)$. The procedure stops when $|E_p \cap E''| = n-1$, meaning that T' has been transformed into T'' and G_p becomes $G_p = (V, E'')$. Since there can be zero purple edge in $G_0 = (V, E_0 \cup E'')$, the number of flips is at most $n-1$. ∎

The above two lemmas allow us to formulate the following theorem.

Theorem 4.3 *Any tree $T' = (V, E')$ can be transformed into another tree $T'' = (V, E'')$ with at most $n-1+k$ flips where k is the number of intersected red edges, provided that for any flip $1 \leq i \leq k$, $CH(V) \cap V_i' \neq \emptyset$ and $CH(V) \cap V_i'' \neq \emptyset$.*

Proof Consider the graph $G_0 = (V, E_0 \cup E'')$, where $E_0 = E'$. The graph can be made planar by removing all the intersections between red and blue edges, as previously shown. Thus we need at most k flips to make the graph planar provided for any flip $1 \leq i \leq k$, $CH(V) \cap V_i' \neq \emptyset$ and $CH(V) \cap V_i'' \neq \emptyset$. As the graph is made planar, we can now follow Lemma 4.2 to obtain T''. It takes at most $n-1+k$ flips to transform T' into T''. ∎

If the set of points are in convex position, then each flip must reduce the number of intersections between red and blue edges by at least one, since there will always be two consecutive black and white points available to make the flip successful. Now we have the following corollary:

Corollary 4.4 *When the set of points is in convex position we need at most $n-1+k$ flips for the above transformation since for any flip $1 \leq i \leq k$, $CH(V) \cap V_i' \neq \emptyset$ and $CH(V) \cap V_i'' \neq \emptyset$.*

4.1 Counterexample

In this section, we show that there exist two trees defined on the same point set such that there does not exist any flip in one of the two trees that reduces the total number of intersections by at least 1 in G_i. Such an example is shown in Fig. 6 where the tree, T_1 in 6(a) has three edges different from the tree, T_2 in 6(b), that is, $|T_1 \setminus T_2| = 3$. However, there is no way (as evident from 6(c)) that any of the trees can be transformed to the other by three flips. This means that the direct transformation would fail after looking for all possible removal of edge crossings where the searching would take time proportional to the number of crossings. However, this kind of problem can be resolved by resorting to the technique of using a canonical tree which guarantees to take at most $2n-m-s-2$ flips for such transformation.

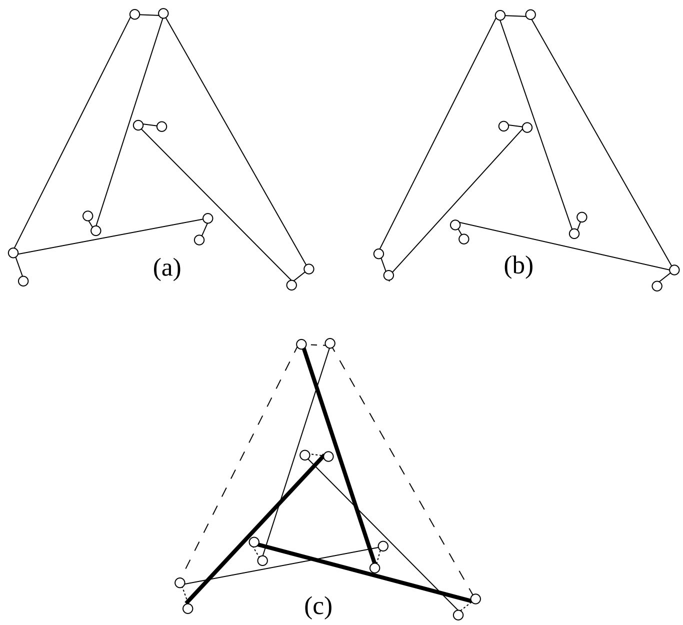

Fig. 6: A counterexample with two trees defined on the same point where there does not exist any flip in one of the two trees such that the total number of intersections is reduced by at least one in G_i.

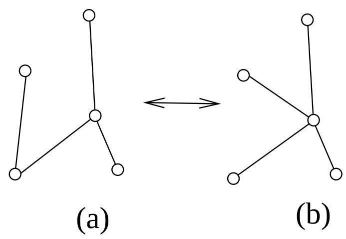

Fig. 7: Transforming one tree into another takes only one flip optimally, but 5 flips through a canonical tree.

4.2 Remarks

This technique (direct transformation) has the obvious advantage that in some cases it leads to the optimal number of flips to complete the transformation. It is well known that an approach for transforming a given tree (in general, it is true for other planar graphs of some class, e.g., planar paths, pseudotriangulations, etc.) into another via a flip operation which depends on a canonical tree might not lead to the computation of the optimal number of flips. This is because the two objects may differ only in a very small number of edges, whereas to transform them into a canonical form may take a large number of flips. As can be seen from Fig. 7, transforming any of the trees into the other takes 5 flips via a canonical tree based approach whereas only one flip suffices.

Finally, we provide a simple average case analysis of the number of flips of our algorithm. First, we determine the average number of intersected red edges. The number of intersected red edges varies from 0 to $n-3$. Thus, the total number of intersected red edges is $\sum_{k=0}^{n-3} k$ yielding the average number of intersected red edges, $\sum_{k=0}^{n-3} k/(n-2) = \frac{1}{n-2}(1+2+3+\cdots+n-3) = (n-3)/2$, where k is assumed to be uniformly distributed in $[0, n-3]$. Then the average number of flips required by our algorithm is $(n-1)+(n-3)/2 = 1.5n-2.5$. The above analysis is based on the fact that for any flip and for $1 \leq i \leq k$, $CH(V) \cap V_i' \neq \emptyset$ and $CH(V) \cap V_i'' \neq \emptyset$.

However, the average-case analysis is based on the simplifying assumption that the number of intersections is uniformly distributed over a given interval. It is an interesting open problem to derive a more sophisticated value for the average number of flips required by our algorithm.

5 Conclusion

In this paper, we present two techniques for tree transformation through flips when the points are in general position and also investigate the results when the points are in convex position. The first method relies on canonical tree for such transformation and provides a slight improvement over the previous results [3]. In the second approach, we avoid the use of the canonical tree and directly transform one tree into another and show that it takes at most $n-1+k$ flips ($k \geq 0$) for such transformation when the points are in convex position. We also show results when the points are in general position and provide an upper bound on the number of flips. If there are no intersections in the union of edges of the given trees, it is shown that the second technique performs an optimal number of flips. Finally, a counterexample is given to show that if two planar trees on the same point set differ by k edges, they cannot be transformed to one another by at most k flips.

References

[1] S. Akl, K. Islam, and H. Meijer On planar path transformation, 18th Canadian Conference of Computational Geometry (CCCG), Kingston, Ontario, August 14-16, 2006.

[2] D. Avis, Generating rooted triangulations without repetitions, Algorithmica 16: 618-632, 1996.

[3] D. Avis and K. Fukuda, Reverse search for enumeration, Discrete Applied Mathematics 65: 21-46, 1996.

[4] P. Bose and F. Hurtado, Flips in planar graphs, Computational Geometry: Theory and Applications, accepted, 2007.

[5] R. Cummings, Hamilton circuits in tree graphs, IEEE Transaction Circuit Theory 13: 82-90, 1966.

[6] A. Garcia, M.Noy, and J. Teigel, Lower bounds on the number of crossing free subgraphs of K_n, Computational Geometry: Theory and Applications 16: 211-221, 2000.

[7] C. Hernando, M. E. Houle, and F. Hurtado, On local transformation of polygons with visibility properties, Theoretical Computer Science 289(2): 919-937, 2002.

[8] C. Hernando, F. Hurtado, A. Marquez, M. Mora, and M. Noy, Geometric tree graphs on points in convex position, Discrete Applied Mathematics, 93: 51-66, 1999.

[9] B. Joe, Construction of three dimensional Delaunay triangulations using local transformation, Computer Aided Geometric Design, 16: 419-453, 1991.

[10] K. Islam, Bounds on some geometric transforms, M.Sc. Thesis, 2005.

[11] C.L. Lawson, Transforming triangulations, Discrete Mathematics, 3: 365-372, 1972.

[12] G. Liu, On connectivities of tree graphs, Journal of Graph Theory 12: 453-459, 1988.

[13] K. Wagner, Bemerkungen zum vierfarbenpmblem, Jber. Deutsch. Math.-Verein. 46 (1936), 26-32.