

Reconstruction of the Crossing Type of a Point Set from the Compatible Exchange Graph of Noncrossing Spanning Trees*

Marcos Oropeza[†]

Csaba D. Tóth^{†‡}

Abstract

Let P be a set of n points in the plane in general position. The *order type* of P specifies, for every ordered triple, a positive or negative orientation; and the *x-type* (a.k.a. *crossing type*) of P specifies, for every unordered 4-tuple, whether they are in convex position. Geometric algorithms on P typically rely on primitives involving the order type or x-type (i.e., triples or 4-tuples). In this paper, we show that the x-type of P can be reconstructed from the *compatible exchange* graph $\mathcal{G}_1(P)$ of noncrossing spanning trees on P . This extends a recent result by Keller and Perles (2016), who proved that the x-type of P can be reconstructed from the *exchange graph* $\mathcal{G}_0(P)$ of noncrossing spanning trees, where $\mathcal{G}_1(P)$ is a subgraph of $\mathcal{G}_0(P)$. A crucial ingredient of our proof is a structure theorem on the maximal sets of pairwise noncrossing edges (MSNES) between two components of a planar straight-line graph on the point set P .

1 Introduction

Let P be a set of n points in the plane in general position (i.e., no three points are collinear), and let $K(P)$ be the straight-line drawing of the complete graph on P . The crossings between the edges of $K(P)$ have a rich combinatorial structure that is not fully understood. For example, the rectilinear crossing number, $\overline{\text{cr}}(K_n)$, is the minimum number of crossings between the edges of $K(P)$ over all n -element points sets P in general position [1, 2]. For a given graph G , computing $\overline{\text{cr}}(G)$ is known to be NP-hard; although approximation algorithms are available [8, 11].

The *order type* of P specifies, for every ordered triple in P , a positive or negative orientation. The *crossing type* (for short, *x-type*) of P specifies, for every unordered pair of edges in $K(P)$, whether they properly cross; which in turn determines, for every unordered

*Research supported in part by the NSF award CCF-1423615.

[†]Department of Mathematics, California State University Northridge, Los Angeles, CA, Email: marcos.oropeza.774@my.csun.edu and csaba.toth@csun.edu

[‡]Department of Computer Science, Tufts University, Medford, MA.

4-tuple in P , whether the four points are in convex position (i.e., whether they span two crossing edges). Clearly, the order type determines the x-type of a point set, but an x-type may correspond to up to $O(n)$ distinct order types [4, 18]. There are $\exp(\Theta(n \log n))$ order types realized by n points in the plane [12]. Nevertheless, there are efficient algorithms to decide whether two (unlabeled) point sets have the same order type [6] or x-type [18]. The x-type also uniquely determines the *rotation system* of a point set, which specifies the cyclic order of edges of $K(P)$ incident to each point in P [16, Proposition 6]. Geometric algorithms typically rely on elementary predicates involving a constant number of points (such as orientations or convexity). It is a fundamental problem to understand the relations between these predicates [15].

In this paper, we wish to reconstruct the x-type of a point set P from relations between noncrossing subgraphs of $K(P)$. Let $\mathcal{T} = \mathcal{T}(P)$ be the set of all noncrossing spanning trees of $K(P)$. There are n^{n-2} abstract spanning trees on $n \geq 3$ labeled vertices [9]; but the number of noncrossing straight-line spanning trees on n points is in $O(141.07^n)$ [13] and $\Omega(6.75^n)$ [3, 10].

Let $\mathcal{G}_0 = \mathcal{G}_0(P)$ be the graph on the vertex set \mathcal{T} where a pair of trees $\{T_1, T_2\}$ forms an edge if we can obtain T_2 from T_1 by exchanging one edge for another (i.e., the symmetric difference of the edge sets of T_1 and T_2 has cardinality 2). The graph \mathcal{G}_0 is called the *exchange graph* of the trees in \mathcal{T} . It is derived from the well-known exchange property of graphic matroids.

Recently, Keller and Perles [14] proved that the x-type of P can be reconstructed from the (unlabeled) graph $\mathcal{G}_0(P)$. In particular, they showed that the set system of maximal cliques in $\mathcal{G}_0(P)$ already provides enough information to reconstruct the x-type of P .

Over the last few decades, several graphs have been introduced on the noncrossing spanning trees \mathcal{T} , besides the classic exchange graph $\mathcal{G}_0(P)$. A *compatible exchange* between T_1 and T_2 exchanges two edges that do not cross each other. See [17] for the hierarchy of five relations. Each of these relations defines a graph on \mathcal{T} . In particular, we denote by $\mathcal{G}_1 = \mathcal{G}_1(P)$ the *compatible exchange graph*. By definition, $\mathcal{G}_1(P)$ is a subgraph of $\mathcal{G}_0(P)$. Both $\mathcal{G}_0(P)$ and $\mathcal{G}_1(P)$ are known to be connected, and their diameters are bounded above by $2n - 4$ [7].

Our results. Our main result (Theorem 3) is that the compatible exchange graph $\mathcal{G}_1(P)$ determines the exchange graph $\mathcal{G}_0(P)$, and hence the x-type of P . The key ingredient of our proof is Theorem 1, which is a structural theorem about maximal sets of pairwise noncrossing edges (for short, MSNEs) between two geometric graphs. An MSNE is a bipartite variant of a triangulation, but (unlike triangulations) MSNEs do not necessarily have the same cardinality, and they need not be connected under edge flip operations.

Organization. In Section 2, we distinguish two types of cliques of size 3 or higher in the compatible exchange graph $\mathcal{G}_1(P)$, and prove some of their basic properties. Since $\mathcal{G}_1(P)$ is a subgraph of $\mathcal{G}_0(P)$, many of these properties hold for the cliques of both $\mathcal{G}_0(P)$ and $\mathcal{G}_1(P)$. In Section 3, we characterize the *maximal* cliques of $\mathcal{G}_1(P)$, which allows us to determine the

type of these cliques in many, but not all cases. Section 4 describes an algorithm for finding all edges of the exchange graph $\mathcal{G}_0(P)$ that are not present in the compatible exchange graph $\mathcal{G}_1(P)$, and hence reconstruct both $\mathcal{G}_o(P)$ and the x-type of the point set P . This is the main result of our paper. We conclude in Section 5 with related open problems.

2 Preliminaries

Let P be a set of points in general position in the plane. A clique in $\mathcal{G}_0(P)$ can be formed in the following two ways (refer to Fig. 1):

1. Let F be a noncrossing spanning forest on P consisting of two trees A and B . Let E_{AB} be a set of all edges in $K(P)$ between A and B that do not cross any edge of A and B . Let C be a set of trees $T = (P, E)$, where $E = A \cup B \cup \{e\}$ and $e \in E_{AB}$. The intersection of the edge sets of any two trees in C is the forest F . The set C induces a clique in $\mathcal{G}_0(P)$. We call C a clique of **type 1**.
2. Let $G = (P, E_c)$ be a noncrossing connected spanning graph with n edges. Then G contains a unique cycle that we denote by U . Let C be a set of trees $T = (P, E)$, where $E = E_c \setminus \{e\}$ and e is an edge of U . The union of any two trees in the clique C is the graph G . The set C induces a clique in $\mathcal{G}_0(P)$. We call C a clique of **type 2**.

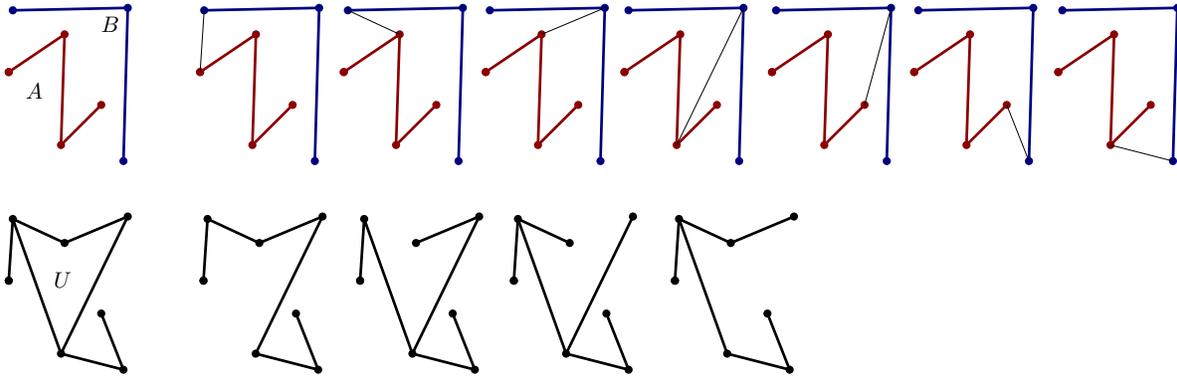


Figure 1: Top row: A noncrossing spanning forest $F = A \cup B$; and the noncrossing spanning trees obtained by adding an edge from E_{AB} . Bottom row: A connected noncrossing spanning graph with a unique cycle U ; and the noncrossing spanning trees obtained by deleting an edge from U .

It is easy to see that *every* clique of size three or higher in $\mathcal{G}_0(P)$ is of type 1 or 2. We provide a proof for completeness.

Lemma 1. *Every clique of size 3 in $\mathcal{G}_0(P)$ is either of type 1 or of type 2.*

Proof. Let C_0 be a clique of size 3 in \mathcal{G}_0 , whose vertices are the trees $T_1 = (P, E_1)$, $T_2 = (P, E_2)$, and $T_3 = (P, E_3)$. Since the three trees are pairwise adjacent, the symmetric difference of the edge sets of any two trees has size two. Since $|E_1| = |E_2| = |E_3| = n - 1$, we have

$$n \leq |E_1 \cup E_2 \cup E_3| \leq n + 1.$$

Since T_1 and T_2 are adjacent, there are edges $e_1 \in E_1$ and $e_2 \in E_2$ such that $E_2 = E_1 - e_1 + e_2$. Similarly, there are edges $e'_1 \in E_1$ and $e_3 \in E_3$ such that $E_3 = E_1 - e'_1 + e_3$. If e_1 , e'_1 , e_2 , and e_3 are pairwise distinct, then all four edges would be in the symmetric difference $E_2 \Delta E_3$, contradicting the assumption that T_2 and T_3 are adjacent in $\mathcal{G}_0(P)$. We arrive at two possibilities:

If $e_1 = e'_1$, then $e_2 \neq e_3$ (since $T_2 \neq T_3$). In this case, $|E_1 \cap E_2 \cap E_3| = |E_1 - e_1| = n - 2$. The intersection of the three trees gives a forest of two trees, hence C_0 is a clique of type 1.

If $e_2 = e_3$, then $e_1 \neq e'_1$ (since $T_2 \neq T_3$). In this case, $|E_1 \cup E_2 \cup E_3| = |E_1 + e_2| = n$. Then the graph $G = (P, E_1 \cup E_2 \cup E_3)$ is connected and has a unique cycle, hence C_0 is a clique of type 2. \square

Lemma 2. *If a clique C_1 in $\mathcal{G}_0(P)$ contains a clique C_2 of size 3, then C_1 is of the same type as C_2 .*

Proof. Let C_k be a clique in $\mathcal{G}_0(P)$ induced by k trees $T_1, \dots, T_k \in \mathcal{T}$. We have seen that every triple $\{T_1, T_2, T_i\}$, $i = 3, \dots, k$, induces a clique of type 1 or 2. It is enough to show that all these triples are of the same type, and then all k trees are either supergraphs of the forest $F = (P, E_1 \cap E_2)$ and hence C_k is of type 1, or subgraphs of $G_c = (P, E_1 \cup E_2)$ and hence C_k is of type 2.

Suppose, to the contrary, that $\{T_1, T_2, T_i\}$ is of type 1 and $\{T_1, T_2, T_j\}$ is of type 2, for some $i \neq j$. Then there exist edges $e_i \in E_i$ and $e_j \in E_1 \cup E_2$ such that $E_i = (E_1 \cap E_2) + e_i$ and $E_j = (E_1 \cup E_2) - e_j$. Since the unique cycle in $E_1 \cup E_2$ contains at least two edges between the two components of $E_1 \cap E_2$, the set of these two edges is exactly $E_1 \Delta E_2$. Consequently, $E_i \Delta E_j = (E_1 \Delta E_2) \cup \{e_i, e_j\}$, contradicting the assumption that T_i and T_j are adjacent in $\mathcal{G}_0(P)$. \square

Corollary 1. *For every $k \geq 3$, every clique of size k in $\mathcal{G}_0(P)$ is either of type 1 or of type 2.*

Remark 1. *Since $\mathcal{G}_1(P)$ is a subgraph of $\mathcal{G}_0(P)$, Lemmas 1 and 2 as well as Corollary 1 hold for the cliques of $\mathcal{G}_1(P)$, as well.*

Lemma 3. *If cliques C_1 and C_2 in $\mathcal{G}_0(P)$ intersect in a clique C_3 of size 3 or more, then their union induces a clique in $\mathcal{G}_0(P)$.*

Proof. By Lemma 1, C_3 is a clique of either type 1 or type 2. By Lemma 2, C_1 and C_2 are of the same types as C_3 .

If C_1 , C_2 , and C_3 are of type 1, then the intersection of the trees in C_3 is a noncrossing forest F of two trees, A and B , such that every tree in C_3 consists of F and an edge between A and B . The same holds for every tree in C_1 and in C_2 (with the same forest F). Hence any two trees in $C_1 \cup C_2$ are adjacent in $\mathcal{G}_0(P)$.

If C_1 , C_2 , and C_3 are of type 2, then the union of the trees in C_3 form a *plane* graph G with n edges and a unique *noncrossing* cycle U such that all trees in C_3 are obtained by deleting an edge from U . The same holds for every tree in C_1 and in C_2 (with the same graph G). Hence any two trees in $C_1 \cup C_2$ are adjacent in $\mathcal{G}_0(P)$. \square

Remark 2. *Keller and Perles [14] show that in the exchange graph $\mathcal{G}_0(P)$, every edge T_1T_2 is part of at most two maximal cliques: at most one clique of type 1 and at most one clique of type 2. The maximal cliques of type 1 (resp., 2) in \mathcal{G}_0 are called I-cliques (resp., U-cliques). This observation plays a crucial role in their method for determining the types of all maximal cliques in $\mathcal{G}_0(P)$, and ultimately reconstructing the x -type of P . The compatible exchange graph $\mathcal{G}_1(P)$, however, does not have this property: An edge T_1T_2 in $\mathcal{G}_1(P)$ may be part of many maximal cliques of type 2. In particular, if two maximal cliques in $\mathcal{G}_1(P)$ intersect in a single edge T_1T_1 , then both cliques might be of type 2.*

For this reason, the proof strategy in [14] is no longer viable for compatible exchange graphs. Instead, we analyse the interactions of maximal cliques of type 2, and show how to reconstruct the maximal cliques of $\mathcal{G}_0(P)$ from the maximal cliques of $\mathcal{G}_1(P)$.

3 Properties of the Compatible Exchange Graph

Let P be set of n points in the plane in general position, and F a spanning forest on P consisting of two trees, A and B . Let E_{AB} be a maximal set of pairwise noncrossing edges (for short, MSNE) between A and B that do not cross any edges in F . Let C be a set of trees $T = (P, E)$, where $E = A \cup B \cup \{e\}$ and $e \in E_{AB}$. This set C of trees induces a clique in $\mathcal{G}_1(P)$, that we call a **T-clique**.

Lemma 4. *Every T-clique of size 3 or higher is a maximal clique in $\mathcal{G}_1(P)$.*

Proof. Suppose, to the contrary, that C is a T-clique that contains at least 3 trees but is not maximal. Then C is part of a larger clique C' . Since C is of type 1, C' is also of type 1 by Lemma 1. Then the intersection of any two trees in C' is a forest F composed of two trees A and B , and the trees in C' are obtained by augmenting F with an edge from a set E'_{AB} of segments between A and B . Then the T-clique C is defined by the same forest F , and an MSNE E_{AB} , $E_{AB} \subset E'_{AB}$.

Let $T_0 \in C' \setminus C$ be a tree $T_0 = (P, E_0)$, $E_0 = A \cup B \cup \{e_0\}$, where $e_0 \in E'_{AB} \setminus E_{AB}$. Because of the maximality of E_{AB} , e_0 crosses some edge $e_1 \in E_{AB}$. Then there is a tree $T_1 = (P, E_1)$ where $E_1 = A \cup B \cup \{e_1\}$. Since e_0 and e_1 cross, there is no compatibility exchange between T_0 and T_1 . This contradicts the assumption that both T_0 and T_1 are in a clique of the compatible exchange graph $\mathcal{G}_1(P)$. \square

Lemma 5. *Every maximal clique in $\mathcal{G}_1(P)$ of size three or higher is either a U-clique (i.e., a maximal clique of type 2 in \mathcal{G}_0) or a T-clique.*

Proof. Let C be a maximal clique in $\mathcal{G}_1(P)$. By Corollary 1, then C is either of type 1 or type 2. We distinguish between two cases.

Case 1: C is of type 1. The clique C is the set of trees that contain a forest F of two trees, A and B , and an edge from a set E_{AB} . Since any two trees in C are compatible, no two edges in E_{AB} cross each other. If E_{AB} is not a maximal set of noncrossing edges, then it is a proper subset of some MSNE between A and B , hence C is proper subset of a T-clique, contradicting the maximality of C .

Case 2: C is of type 2. The noncrossing spanning graph G is the union of all trees in the clique C . Note that G determines a U-clique in \mathcal{G}_0 , which contains C . That is, G has a unique (noncrossing) cycle U , and all trees in C are obtained by deleting an edge from U . Since U is noncrossing, the exchange between any two of its edges is a compatible exchange, and so C equals this U-clique of \mathcal{G}_0 . \square

Definition 1. (*Special Configurations*) A noncrossing straight-line graph $G = (P, E)$ is *special* if G consists of two connected components, A and B , that contain edges $a_1a_2 \in A$ and $b_1b_2 \in B$ such that the vertices (a_1, b_1, b_2, a_2) are on the boundary of $\text{conv}(P)$ in this counterclockwise order, both a_1b_1 and a_2b_2 are edges of $\text{conv}(P)$, and the interior of $\text{conv}\{a_1, b_1, b_2, a_2\}$ does not contain any point of P (see Fig. 2(left)). Note that $T_1 = (P, A \cup B \cup \{a_1b_1\})$ and $T_2 = (P, A \cup B \cup \{a_2b_2\})$ are noncrossing spanning trees; and T_1 and T_2 are adjacent in $\mathcal{G}_1(P)$. We call the edge T_1T_2 of $\mathcal{G}_1(P)$ a *special edge*.

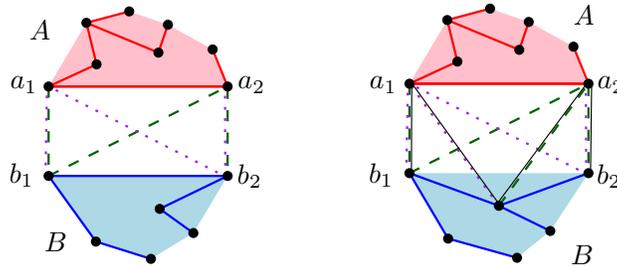


Figure 2: Left: A special edge is T_1T_2 , where $T_1 = (P, A \cup B \cup \{a_1b_1\})$ and $T_2 = (P, A \cup B \cup \{a_2b_2\})$. Right: The edge T_1T_2 is in three T-cliques: C_1 (solid edges), C_2 (dotted edges), and C_3 (dashed edges), where $C_2 \cap C_3 = \{T_1, T_2\}$.

Remark 3. For the pair of trees, A and B , in the configuration of a special edge, there are precisely two MSNEs between A and B : $E'_{AB} = \{a_1b_1, a_1b_2, a_2b_2\}$ and $E''_{AB} = \{a_1b_1, a_2b_1, a_2b_2\}$. The sets E'_{AB} and E''_{AB} define two T-cliques, each of size 3. The edge T_1T_2 is a part of both T-cliques.

If T_1T_2 is a special edge of the compatible exchange graph \mathcal{G}_1 , then it is part of precisely two T-cliques. We next prove the converse: if an edge of \mathcal{G}_1 is part of precisely two T-cliques, then it is special (Corollary 2). We first prove a crucial property of MSNEs, which may be of independent interest, since MSNEs are considered a bipartite variant of triangulations.

Theorem 1. Let $G = (P, E)$ be a noncrossing straight-line graph with two connected components, A and B . Then one of the following holds:

1. there is a unique MSNE between A and B ,
2. G is special (and there are precisely two MSNEs between A and B , cf. Definition 1), or
3. for every MSNE E_{AB} between A and B , there exists another MSNE E'_{AB} such that E_{AB} and E'_{AB} share at least three edges (i.e., $|E_{AB} \cap E'_{AB}| \geq 3$).

Proof. if there is a unique MSNE between A and B , or if G is special, then the proof is complete. We may assume that neither A nor B is a singleton (otherwise there would be a unique MSNE), consequently every MSNE between A and B has at least three edges. Assume that E'_{AB} and E''_{AB} are two distinct MSNEs between A and B . It suffices to show that there exists a MSNE E'''_{AB} such that $|E'_{AB} \cap E'''_{AB}| \geq 3$. If $|E'_{AB} \cap E''_{AB}| \geq 3$, then we can take $E'''_{AB} = E''_{AB}$. Hence we may assume that $|E'_{AB} \cap E''_{AB}| \leq 2$. We distinguish between two cases.

Case 1: A and B each have a vertex on the boundary of $\text{conv}(P)$. Since both A and B are connected, the boundary of $\text{conv}(P)$ contains precisely two edges between A and B . These two edges are contained in every MSNE between A and B (since convex hull edges do not cross any other edges). We may assume that $E'_{AB} \cap E''_{AB} = \{e_1, e_2\}$, where both e_1 and e_2 are convex hull edges.

We **claim** that there exist edges $e_3 \in E'_{AB} \setminus \{e_1, e_2\}$ and $e_4 \in E''_{AB} \setminus \{e_1, e_2\}$ (possibly $e_3 = e_4$) such that e_3 and e_4 do not cross each other, and they cross neither e_1 nor e_2 . Assuming the claim is true, there is a set $\{e_1, e_2, e_3, e_4\}$ of at least three pairwise noncrossing edges between A and B . We can augment this set to an MSNE, denoted by E'''_{AB} , as required.

Suppose, to the contrary, that the claim is false. This implies the following:

$$\text{Every edge in } E'_{AB} \setminus \{e_1, e_2\} \text{ crosses all edges in } E''_{AB} \setminus \{e_1, e_2\}. \quad (\star)$$

Let $e_1 = a_1b_1$, where $a_1 \in A$ and $b_1 \in B$. We claim that there is an edge in $E'_{AB} \setminus \{e_1, e_2\}$ incident to a_1 or b_1 . (Similarly, there is an edge in $E''_{AB} \setminus \{e_1, e_2\}$ incident to a_1 or b_1 .)

Consider $G_0 = (P, A \cup B \cup E'_{AB})$, which is a noncrossing graph on P . In a triangulation of G_0 , the edge a_1b_1 is incident to some triangle $\Delta a_1b_1c'$. If $c' \in A$, then $b_1c' \in E'_{AB}$ because it connects A and B and it does not cross any edge in E'_{AB} . Otherwise $c' \in B$, and $a_1c' \in E'_{AB}$. Assume w.l.o.g. that $c' \in B$, hence $a_1c' \in E'_{AB}$. Since A is not a singleton, a_1c' cannot be a convex hull edge. Therefore $a_1c' \in E'_{AB} \setminus \{e_1, e_2\}$. By assumption (\star) , a_1c' crosses all edges in $E''_{AB} \setminus \{e_1, e_2\}$.

Similarly, a_1b_1 is incident to some triangle $\Delta a_1b_1c''$ in the triangulation of $(P, A \cup B \cup E''_{AB})$, where a_1c'' or b_1c'' is in $E''_{AB} \setminus \{e_1, e_2\}$. Since a_1c' and a_1c'' are adjacent, they do not cross. By assumption (\star) , $b_1c'' \in E''_{AB} \setminus \{e_1, e_2\}$, consequently $c'' \in A$. In particular $c' \neq c''$. Let d be a vertex in A such that $b_1d \in E''_{AB} \setminus \{e_1, e_2\}$ and $\angle a_1b_1d$ is minimal (Fig. 3).

Next we claim that every edge in $E'_{AB} \setminus \{e_1, e_2\}$ is incident to a_1 . Let uv be an edge in $E'_{AB} \setminus \{e_1, e_2\}$. Suppose for the sake of contradiction that neither u nor v is equal to a_1 . By assumption, uv crosses b_1d . Let H_1 be the halfplane bounded by a_1b_1 that contains all points in P , and let H_2 be the halfplane bounded by b_1d that contains a_1 . Let $W = H_1 \cap H_2$ and note that it is a wedge with apex b_1 . Since uv crosses b_1d , we can assume without loss of

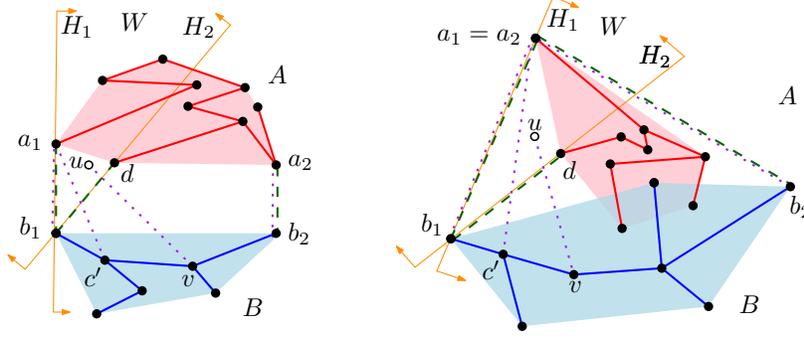


Figure 3: The configurations described in the proof of Theorem 1, Case 1. Left: an instance where $a_1, a_2, b_1,$ and b_2 are distinct. Right: another instance where $a_1 = a_2$.

generality that $u \in H_2$. Note that $u \in H_1$ since all vertices are in H_1 . We have that $u \in W$ because it is in both H_1 and H_2 .

Let $P(a_1, d)$ be a unique path in A between a_1 and d . Let C_0 be the cycle $P(a_1, d) \cup (d, b_1, a_1)$. Let C be the set of points in the interior or on the boundary of C_0 . We want to show that $u \in A$. Suppose, to the contrary, that $u \in B$. Then u is in the interior of C_0 . That means that there exists a path $P(b_1, u)$ in B . Since C_0 has only one vertex in B , namely b_1 , this path lies inside C_0 . Among all vertices of B in the interior of C_0 , let $u' \in B$ be a vertex u' farthest from the line b_1d (in Euclidean distance). Augment E'_{AB} to a triangulation arbitrarily. Since u' lies in the interior of C_0 , it has a neighbor v' farther from the line b_1d . By the choice of u' , we have $v' \in A$, hence $u'v' \in E'_{AB}$. Since the edge $u'v'$ lies inside C_0 , it does not cross b_1d , contradicting the assumption that every edge in $E'_{AB} \setminus \{e_1, e_2\}$ crosses b_1d . We have shown that every edge in $E'_{AB} \setminus \{e_1, e_2\}$ is incident to a_1 . Symmetrically, every edge in $E'_{AB} \setminus \{e_1, e_2\}$ is incident to b_2 .

Overall, every edge in $E'_{AB} \setminus \{e_1, e_2\}$ equals to a_1b_2 . Hence $E'_{AB} = \{a_1b_2, e_1, e_2\}$. Analogously, $E''_{AB} = \{a_2b_1, e_1, e_2\}$. Since these are MSNES between A and B , we have $a_1a_2 \in A$ and $b_1b_2 \in B$. In this case, the edge T_1T_2 is a special edge. This completes the proof in Case 1.

Case 2: A or B lies in the interior of $\text{conv}(P)$. Assume without loss of generality that B lies in the interior of $\text{conv}(P)$. Let $e_1 = a_1b_1$ and $e_2 = a_2b_2$ be two arbitrary edges in E'_{AB} , with $a_1, a_2 \in A$ and $b_1, b_2 \in B$. Let $D = \text{conv}(B \cup \{a_1, a_2\})$. (See Fig. 4.) Note that $B \cup \{a_1, a_2\}$ has at least three vertices that are not collinear since P is in general position. Consequently, D has at least three extremal vertices. One of these vertices, denote it by b_3 , is in B . Create a triangulation of $A \cup B \cup E'_{AB}$. Since B is in the interior of the convex hull P , vertex b_3 is incident to triangles whose angles add up to 2π . Since b_3 is an extremal vertex of D , it is incident to an edge in the triangulation in the exterior of D . Let one such edge be $e_3 = b_3a'$. Since B is contained in D , $a' \notin B$, and therefore $a' \in A$. It is also clear that $b_3a' \in E'_{AB}$ as b_3a' does not cross any edge in E'_{AB} because it is from a triangulation. Note that also $b_3a' \neq e_1, e_2$ since both e_1 and e_2 are contained in D .

Similarly, there exists an $a'' \in A$ such that $b_3a'' \in E''_{AB}$, and $b_3a'' \neq e_1, e_2$. Let $e_4 = b_3a''$. Since e_3 and e_4 have a common endpoint, namely b_3 , they do not cross (possibly $e_3 = e_4$).

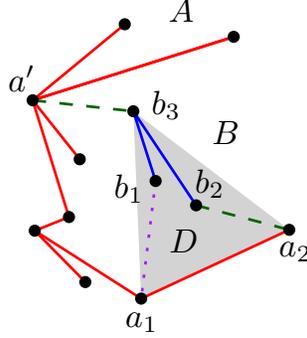


Figure 4: The configurations described in the proof of Theorem 1, Case 2.

Therefore, we have a set $\{e_1, e_2, e_3, e_4\}$ of at least three pairwise noncrossing edges between A and B . We can augment this set into an MSNE, and denote it E'_{AB} . By construction, $E'_{AB} \cap E'''_{AB} = \{e_1, e_2, e_3\}$, as required. \square

Corollary 2. *If an edge T_1T_2 of $\mathcal{G}_1(P)$ is contained in two or more T-cliques, then it is either a special edge of $\mathcal{G}_1(P)$ or every T-clique containing T_1T_2 intersects some other T-clique containing T_1T_2 in at least three vertices.*

Proof. Let T_1T_2 be an edge in $\mathcal{G}_1(P)$ contained in two or more T-cliques. Let $T_1 = (P, E_1)$ and $T_2 = (P, E_2)$ such that $E_2 = E_1 - e_1 + e_2$. Let $F = (P, E_1 \cap E_2)$, which is a forest of two trees that we denote by A and B . Every T-clique containing T_1T_2 corresponds to a set of trees that contains the forest F and one edge from some MSNE E_{AB} between A and B . The T-cliques containing T_1T_2 differ in the sets E_{AB} , but share the same forest F . Theorem 1, applied for $G = F$, completes the proof: Either T_1T_2 is a special edge of $\mathcal{G}_1(P)$, or every T-clique containing T_1T_2 intersects some other T-clique containing T_1T_2 in at least three vertices. \square

Cliques of Type 2. It is easy to show that $\mathcal{G}_0(P)$ and $\mathcal{G}_1(P)$ have the same maximal cliques of type 2 of size 3 or higher. Adopting terminology from [14], we call these cliques U-cliques.

Note, however, that $\mathcal{G}_0(P)$ may have a maximal clique of size 2, where the unique cycle U has one crossing. Such a degenerate U-clique is no longer possible in $\mathcal{G}_1(P)$.

Proposition 1. *If T_1T_2 is an edge of $\mathcal{G}_1(P)$, then T_1T_2 is part of a unique U-clique.*

Proof. The union of T_1 and T_2 is a connected noncrossing graph with n edges. Therefore it contains a unique cycle u . If the edge T_1T_2 belongs to any U-clique generated by a plane graph $G = (P, E_u)$ with a unique cycle u , then we have $E_1 \subseteq E_u$ and $E_2 \subseteq E_u$. Therefore, $E_u = E_1 \cup E_2$. \square

Proposition 2. *Every U-clique has size at least three.*

Proof. Every cycle in a graph has at least three edges. Therefore the U-clique constructed from any cycle has size at least three. \square

Propositions 1 and 2 immediately imply the following.

Corollary 3. *Every edge in $\mathcal{G}_1(P)$ is contained in at least one U-clique of size at least three.*

4 Reconstruction of the Exchange Graph

Recall that the compatible exchange graph $\mathcal{G}_1(P)$ is a subgraph of the exchange graph $\mathcal{G}_0(P)$. In this section, we explore how the maximal cliques of the two graphs are related, and then show how to find all edges of $\mathcal{G}_0(P)$ that are not present in $\mathcal{G}_1(P)$.

Lemma 6. *For every edge T_1T_2 in $\mathcal{G}_1(P)$, the maximal cliques that contain T_1T_2 satisfy precisely one of the following conditions:*

1. T_1T_2 is contained in exactly one maximal clique, which is a U-clique.
2. T_1T_2 is contained in precisely two maximal cliques: one U-clique and one T-clique.
3. T_1T_2 is contained in precisely three maximal cliques, two of which are of size three and the third of size four. The clique of size four must be a U-clique and the other two are T-cliques. In particular, T_1T_2 is a special edge.
4. T_1T_2 is contained in three or more maximal cliques, one of which intersects the other cliques in only T_1T_2 and, for all other cliques, there is another clique of size at least three that intersects with them.

Proof. Let T_1T_2 be an edge in $\mathcal{G}_1(P)$. It is contained in at least one clique of size 3 or higher by Corollary 3. If it is contained in precisely one or two maximal cliques, then it satisfies condition 1 or 2, respectively.

Assume T_1T_2 is contained in three or more maximal cliques. One of them is a U-clique by Propositions 1 and 2. Since T_1T_2 is contained in a unique U-clique, all other cliques that contain it are T-cliques. By Corollary 2, either T_1T_2 is a special edge or each of these T-cliques intersects some other T-clique that contain T_1T_2 in three or more vertices. If T_1T_2 is a special edge, then the U-clique is of size 4, hence T_1T_2 satisfies condition 3. Otherwise, condition 4 is satisfied. \square

Lemma 7. *If T_1T_2 is an edge of the exchange graph $\mathcal{G}_0(P)$, but not an edge of the compatible exchange graph $\mathcal{G}_1(P)$, then:*

1. T_1 and T_2 are elements of two T-cliques, C_1 and C_2 respectively, that intersect in three or more vertices; or
2. T_1 and T_2 are elements of two T-cliques, C_1 and C_2 respectively, that each intersect some T-clique C_3 in three or more vertices; or
3. $T_1 = (P, A \cup B \cup \{e_1\})$ and $T_2 = (P, A \cup B \cup \{e_2\})$ where $e_1 = a_1b_2$ and $e_2 = a_2b_1$ are the two crossing diagonal edges of the special configuration (Fig. 2).

Proof. Let T_1T_2 be an edge in $\mathcal{G}_0(P)$. Then $T_1 = (P, E_1)$ and $T_2 = (P, E_2)$ and there exist edges $e_1 \in E_1$ and $e_2 \in E_2$ such that $E_2 = E_1 - e_1 + e_2$ and $E_1 = E_2 - e_2 + e_1$. Since T_1T_2 is not an edge in the compatibility exchange graph, e_1 and e_2 cross. The edge T_1T_2 of $\mathcal{G}_0(P)$ is contained in a unique maximal clique of type 1, which is called an *I-clique* (see [14, Section 2.3]). Denote by C_0 this I-clique in $\mathcal{G}_0(P)$, and note that it corresponds to a forest F of two trees, A and B , and the set E_{AB} of all edges between A and B (which may cross each other, but do not cross edges in F). An MSNE between A and B is a T-clique contained in C_0 . Let $e_1 = a_1b_1$ and $e_2 = a_2b_2$, where $a_1, a_2 \in A$ and $b_1, b_2 \in B$. Let x be the intersection point of e_1 and e_2 . Let P_1 be the shortest path between a_1 and b_2 homotopic to (a_1, x, b_2) w.r.t. F (possibly, $P_1 = (a_1, b_2)$). Let P_2 be the shortest path between a_2 and b_1 homotopic to (a_2, x, b_1) w.r.t. F (possibly $P_2 = (a_2b_1)$). The paths P_1 and P_2 each contain an edge between A and B . Denote two such edges by f_1 and f_2 , respectively. By construction, neither f_1 nor f_2 crosses e_1 and e_2 . (See Fig. 5.)

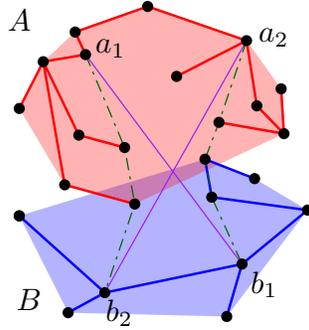


Figure 5: The shortest paths P_1 and P_2 , respectively, homotopic to (a_1, x, b_2) and (a_2, v, b_1) with respect to F in the proof of Lemma 7.

Similar to the proof of Theorem 1, we distinguish between two cases based on the position of A and B relative to $\text{conv}(P)$.

Case 1: A and B each have a vertex on the boundary of $\text{conv}(P)$. Every MSNE contains two convex hull edges between A and B . Since e_1 and e_2 cross, they are not convex hull edges. If E'_{AB} is an MSNE that contains e_1 and E''_{AB} is an MSNE that contains e_2 , then $|E'_{AB} \cap E''_{AB}| \geq 2$. We would like to find edge sets E'_{AB} and E''_{AB} such that they share at least three edges.

Let E'_{AB} be an MSNE that contains e_1 , f_1 , and f_2 . Let E''_{AB} be an MSNE that contains e_2 , f_1 , and f_2 . If f_1 or f_2 is not a convex hull edge, then $|E'_{AB} \cap E''_{AB}| = 3$ because their intersection is f_1 , f_2 , and a convex hull edge. This implies that T_1 and T_2 are elements of two T-cliques, C_1 and C_2 , respectively, that intersect in three or more vertices.

Assume that both $f_1 = a_1b_1$ and $f_2 = a_2b_2$ are convex hull edges. If a_1a_2 and b_1b_2 are edges of A and B , respectively, then we are in the case that $e_1 = a_1b_2$ and $e_2 = a_2b_1$ are the two crossing diagonal edges of the special configuration.

Assume without loss of generality that a_1a_2 is not an edge in A . We want to choose a point $a_3 \in A$ that is not equal to a_1 or a_2 such that the line segment a_3x does not cross any

edge in A . There is a unique path $P(a_1, a_2)$ between a_1 and a_2 in the tree A . Let R be the region enclosed by $P(a_1, a_2) \cup (a_2, x, a_1)$. Triangulate R . Let a_3 be the third vertex adjacent to a_1x . Let P_3 be the shortest path between a_3 and b_1 homotopic to (a_3, x, b_1) (possibly $P_3 = a_3b_1$). Let P_4 be the shortest path between a_3 and b_2 , homotopic to (a_3, x, b_2) (possibly $P_4 = (a_3, b_2)$). These paths each contain an edge between A and B . Denote two such edges by f_3 and f_4 , respectively. By construction, f_3 does not cross e_1 (but may cross e_2) and f_4 does not cross e_2 (but may cross e_1). These edges are distinct, i.e., $f_3 \neq f_4$, as they lie on opposite sides of a_1b_1 and a_2b_2 .

Let $S_1 = \{e_1, f_1, f_2, f_3\}$, $S_2 = \{e_2, f_1, f_2, f_4\}$, and $S_3 = \{f_1, f_2, f_3, f_4\}$. Note that each set contains pairwise noncrossing edges. Let E'_{AB} , E''_{AB} , and E'''_{AB} be MSNEs that contain S_1 , S_2 , and S_3 , respectively. We then have that T_1 and T_2 are elements of two T-cliques C' and C'' , respectively, that each intersect some T-clique C''' in three or more vertices in $\mathcal{G}_1(P)$.

Case 2: A or B lies in the interior of $\text{conv}(P)$. Assume without loss of generality that B lies in the interior of $\text{conv}(P)$. Similar to the proof of Case 2 of Theorem 4, we find vertices $a' \in A$ and $b_3 \in B$ such that the edge $a'b_3$ does not cross any of the edges in A and B as well as does not cross the edges e_1, e_2, f_1 , and f_2 . Let E'_{AB} and E''_{AB} be MSNEs that contain $\{e_1, f_1, f_2, a'b_3\}$ and $\{e_2, f_1, f_2, a'b_3\}$, respectively. Then T_1 and T_2 are elements of two T-cliques, C' and C'' respectively, that intersect in three or more vertices. \square

Theorem 2. *The compatible exchange graph $\mathcal{G}_1(P)$, for a point set P , determines the exchange graph $\mathcal{G}_0(P)$.*

Proof. Our proof is constructive. We are given $\mathcal{G}_1(P)$, which is an unlabeled graph. We can find all special edges T_1T_2 in $\mathcal{G}_1(P)$ based on the characterization in Lemma 6(3), that is, T_1T_2 is contained in precisely 3 maximal cliques, two of which are of size 3 and the third of size 4.

1. Find all maximal cliques of $\mathcal{G}_1(P)$.
2. Find all special edges T_1T_2 in $\mathcal{G}_1(P)$.
3. Put $\mathcal{H} := \mathcal{G}_1(P)$.
4. For every special edge $e = T_1T_2$ in $\mathcal{G}_1(P)$, let $T_3(e)$ and $T_4(e)$ be the 3rd vertices of the two cliques of size 3 in $\mathcal{G}_1(P)$ that contain T_1T_2 ; and augment \mathcal{H} with the edge $T_3(e)T_4(e)$ (if it is not already present in \mathcal{H}).
5. While \mathcal{H} contains two maximal cliques that intersect in three or more vertices, augment \mathcal{H} with the edge to merge them into a single clique.
6. Return \mathcal{H} .

The “dual” special edges added in step 2 are in $\mathcal{G}_0(P)$ by Lemma 6. Edges added in the merge steps are in $\mathcal{G}_0(P)$ by Lemma 3. We see that the algorithm returns a subgraph of $\mathcal{G}_0(P)$. Conversely, the algorithm adds all edges in $\mathcal{G}_0(P) \setminus \mathcal{G}_1(P)$ by Lemma 7. Therefore it returns $\mathcal{G}_0(P)$, as required. \square

The combination of Theorem 2 with [14] readily implies our main result.

Theorem 3. *The compatible exchange graph $\mathcal{G}_1(P)$, for a point set P in the plane in general position, determines the x -type of P .*

5 Future Directions

We have shown (Theorem 3) that the compatible exchange graph $\mathcal{G}_1(P)$ determines the exchange graph $\mathcal{G}_0(P)$, and hence the x -type of the point set P . Our results and our proof techniques raise several interesting problems. We list a few of them here.

1. Can the x -type of a point set P be reconstructed from any of the more restrictive relations, e.g., the *rotation graph* or the *edge slide graph* on $\mathcal{T}(P)$ (see [17] for precise definitions)? While all these graphs are known to be connected, the edge slide graph of $\mathcal{T}(P)$, for example, may have vertices of degree 1 [5], hence its clique complex is a 1-dimensional simplicial complex.
2. By Corollary 3, the compatible exchange graph $\mathcal{G}_1(P)$ is biconnected for $|P| \geq 3$ (in fact, the clique complex of $\mathcal{G}_1(P)$ is a 2-dimensional simplicial complex). However, the minimum degree is $\Omega(n)$. It remains an open problem to find tight bounds for the diameter, the minimum degree, and the vertex- and edge-connectivity of $\mathcal{G}_1(P)$ over all n -element point sets P .
3. T-cliques played a crucial role in the proof of our main result (cf. Theorem 2), but the underlying geometric structures raise intriguing open problems. Every maximal T-clique corresponds to a MSNE E_{AB} between A and B , which are two trees in a noncrossing spanning forest $F = (P, A \cup B)$. The set E_{AB} can be thought of as a bipartite analogue of straight-line triangulations on a point set P (which is a maximal set of pairwise noncrossing edges in $K(P)$). It is well known that every triangulation on P has the same number of edges, and the space of triangulations on P is connected under the so-called *edge flip* operation. For a forest $F = (P, A \cup B)$, we can define the collection \mathcal{E}_{AB} of all maximal sets of pairwise noncrossing edges between A and B . Two sets in \mathcal{E}_{AB} may have different cardinalities, and there is no clear operation that would generate all sets in \mathcal{E}_{AB} . It remains an open problem to understand the combinatorial structure of \mathcal{E}_{AB} .

References

- [1] Bernardo M. Ábrego, Mario Cetina, Silvia Fernández-Merchant, Jesús Leaños, and Gelasio Salazar. On $\leq k$ -edges, crossings, and halving lines of geometric drawings of K_n . *Discrete & Computational Geometry*, 48(1):192–215, 2012.
- [2] Bernardo M. Ábrego and Silvia Fernández-Merchant. The rectilinear local crossing number of K_n . *Journal of Combinatorial Theory, Series A*, 151:131–145, 2017.

- [3] Oswin Aichholzer, Thomas Hackl, Clemens Huemer, Ferran Hurtado, Hannes Krasser, and Birgit Vogtenhuber. On the number of plane geometric graphs. *Graphs and Combinatorics*, 23(1):67–84, 2007.
- [4] Oswin Aichholzer, Vincent Kusters, Wolfgang Mulzer, Alexander Pilz, and Manuel Wettstein. An optimal algorithm for reconstructing point set order types from radial orderings. *Int. J. Comput. Geometry Appl.*, 27(1-2):57–84, 2017.
- [5] Oswin Aichholzer and Klaus Reinhardt. A quadratic distance bound on sliding between crossing-free spanning trees. *Comput. Geom.*, 37(3):155–161, 2007.
- [6] Greg Aloupis, John Iacono, Stefan Langerman, Özgür Ozkan, and Stefanie Wührer. The complexity of order type isomorphism. In *Proc. 25th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 405–415. SIAM, 2014.
- [7] David Avis and Komei Fukuda. Reverse search for enumeration. *Discrete Applied Mathematics*, 65(1):21–46, 1996.
- [8] Samuel Bald, Matthew P. Johnson, and Ou Liu. Approximating the maximum rectilinear crossing number. In Thang N. Dinh and My T. Thai, editors, *Computing and Combinatorics*, pages 455–467, Cham, 2016. Springer International Publishing.
- [9] Arthur Cayley. A theorem on trees. *Quart. J. Math.*, 23:376–378, 1889.
- [10] Philippe Flajolet and Marc Noy. Analytic combinatorics of non-crossing configurations. *Discrete Mathematics*, 204(1):203–229, 1999.
- [11] Jacob Fox, János Pach, and Andrew Suk. Approximating the rectilinear crossing number. In Yifan Hu and Martin Nöllenburg, editors, *Proc. 24th Symposium on Graph Drawing and Network Visualization (GD)*, pages 413–426, Cham, 2016. Springer.
- [12] Jacob E. Goodman and Richard Pollack. The complexity of point configurations. *Discrete Applied Mathematics*, 31(2):167–180, 1991.
- [13] Michael Hoffmann, André Schulz, Micha Sharir, Adam Sheffer, Csaba D. Tóth, and Emo Welzl. Counting plane graphs: Flippability and its applications. In János Pach, editor, *Thirty Essays on Geometric Graph Theory*, pages 303–325. Springer, New York, NY, 2013.
- [14] Chaya Keller and Micha A. Perles. Reconstruction of the geometric structure of a set of points in the plane from its geometric tree graph. *Discrete & Computational Geometry*, 55(3):610–637, 2016.
- [15] Donald E. Knuth. *Axioms and Hulls*, volume 606 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, 1992.

- [16] Jan Kynčl. Enumeration of simple complete topological graphs. *European Journal of Combinatorics*, 30(7):1676–1685, 2009.
- [17] Torrie L. Nichols, Alexander Pilz, Csaba D. Tóth, and Ahad N. Zehmakan. Transition operations over plane trees. In Michael A. Bender, Martín Farach-Colton, and Miguel A. Mosteiro, editors, *Proc. 13th Latin American Symposium on Theoretical Informatics (LATIN)*, pages 835–848, Cham, 2018. Springer.
- [18] Alexander Pilz and Emo Welzl. Order on order types. *Discrete & Computational Geometry*, 59(4):886–922, 2018.