Compatible Trees*

Alfredo García [1] Clemens Huemer [2] Ferran Hurtado [3] Javier Tejel [1]

Abstract

Two plane geometric graphs are said to be compatible when their union is a plane geometric graph. Let S be a set of n points in the Euclidean plane in general position and let T be any given plane geometric tree whose vertex set is S. The main problem we consider in this work consists of finding a second plane geometric tree T' on S, such that T' is compatible with T and shares with T a minimum number of edges. We prove, up to additive constants, that there is always a compatible plane geometric tree T' having in common with T at most $n/4$ edges; while for some plane geometric trees T, any other tree T' spanning S, compatible with T, has at least $n/5$ edges in common with T.

1 Introduction

Preliminaries. Throughout this paper let S be a given set of $n \geq 3$ points in the Euclidean plane in general position, in the sense that no three of them are collinear. We denote by $K(S)$ the complete geometric graph on top of S, i.e., the graph with vertex set S whose edges are all the straight-line segments connecting two points in S. A *geometric graph on* S is a subgraph G of $K(S)$. A geometric graph G is *plane* if no two edges of G intersect except possibly at a common vertex. Plane geometric graphs are also known in the literature as *plane straight line graphs* or as *crossing-free* or *non-crossing* subgraphs of $K(S)$. Unless specified otherwise, all geometric graphs considered in this paper are plane and have as vertices the same set of points S.

Two plane geometric graphs are said to be *compatible* if their union is also a plane geometric graph. A geometric graph that is compatible with a given geometric graph G will be called *G-compatible*.

Purpose of this article. In this paper we focus on the following problem: Given a plane geometric tree T on S, find another plane geometric tree T' spanning S, such that T and T' are compatible and T' has a minimum number of edges in common with T. We will denote by $d(T)$ this minimum number of common edges between T and any other T-compatible tree T'.

*Work partially supported by projects MICINN MTM2009-07242, MINECO MTM2012-30951 and ESF EUROCORES programme EuroGIGA, CRP ComPoSe: MICINN Project EUI-EURC-2011-4306. In addition, the first and fourth authors are also supported by project E58-DGA, and the second and third authors by project Gen. Cat. DGR2009SGR1040.

[1]Departamento de Métodos Estadísticos, IUMA, Universidad de Zaragoza, Zaragoza, Spain. `olaverri@unizar.es,jtejel@unizar.es`.

[2]Departament de Matemàtica Aplicada IV, Universitat Politècnica de Catalunya, Barcelona, Spain. `clemens.huemer@upc.edu`.

[3]Departament de Matemàtica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Spain. `ferran.hurtado@upc.edu`.

In general, $d(T)$ depends on the abstract tree T and also on the position of the points of S, and can be seen as a measure of the obstruction caused by T. For example, $d(T) = 0$ would mean that T is no obstruction at all, as a disjoint plane spanning tree compatible with T could still be drawn on top of S; on the contrary, a high value of $d(T)$ would mean that no much room is left for the edges of a second tree. This degree of obstruction is also a measure of the *visibility* between the points of S when the edges of T are considered as obstacles, and hence our work overlaps the area of Art Gallery Problems, as mentioned below in more detail.

The computation of the exact value of $d(T)$, for a given input tree T, appears to be a difficult problem. However, for some special cases $d(T)$ is easily obtained. For example, if T is a star, then $d(T) = 1$, because we can connect the leaves by a path in radial order around the unique internal node, which has to be reached by one of edge already in T. When the n points of the set S are in convex position, it is easy to see that $d(T) = 1$ for any tree T. Situations in which $d(T)$ can be arbitrarily large are shown in Figures 20 and Figure 21. Therefore, it is natural to study the extremal value

$$d_n = \max_{\{S,T\}} d(T),$$

where S is any set of n points in the plane, and T any plane geometric tree spanning S.

In this paper we give combinatorial bounds on d_n, which we describe precisely in the summary of results below, and we also show that the parameter $d(T)$ can be derived from the analysis of the triangulations of S that are compatible with T.

Previous related research. Problems on compatibility of plane graphs, combining visibility issues and augmentation questions, have been receiving substantial attention in computational geometry since the mid eighties to the present. To *augment* a graph G consists of adding edges to G, ideally in minimum number, to obtain some properties or structure in the resulting augmented graph G'. In our context, we focus on the case in which G and G' are plane and compatible; for this topic we refer the interested reader to the thorough survey [17].

One of the oldest problems in this field is, given a set M of segments (i.e., a crossing-free geometric matching of the set S of their endpoints), to find conditions and decision algorithms for the existence of a plane Hamiltonian cycle spanning S and containing M [24, 25]. As this is not always possible, several alternative conjectures and partial results arose [21, 28, 23], culminating with the proof by Hoffmann and Tóth [14] that a compatible Hamiltonian cycle always exists, i.e., a *polygonization* P of S in which all the segments in M are either sides of the polygon P, or internal diagonals of P, or external diagonals of P. A quite similar problem was posed and discussed in [3]: Given a non-crossing perfect matching on a point set S, one wants to find another non-crossing perfect matching that is compatible and edge-disjoint with the given one. A solution in the affirmative has been obtained very recently by Ishaque et al. [19] when there is an even number of segments (disjointness is not always possible when the number is odd, as proved in [3]).

The preceding problems have often been studied as *visibility problems*, considering the input data as "obstacles" for the visibility between segment endpoints. For this viewpoint and related visibility questions, we refer the reader to the surveys [6, 22, 29].

With more emphasis on the *augmentation* aspects for geometric graphs, we can mention the family of results on augmenting a plane perfect matching to become a compatible plane tree with "good properties" [7, 8, 16], or the large set of works on compatibly augmenting given plane graphs to improve on their vertex-connectivity or their edge-connectivity [1, 12, 5, 26, 27]. Problems in which the vertices of the initial graph are colored and the augmentation has associated constraints have also been studied, as for trees in [13, 16] and for compatible matchings (for example in the thesis [30]). Other related problems, leaning towards the topic of simultaneously drawing two graphs, are discussed in [9, 10, 11, 18, 20].

Finally, it is worth mentioning that another family of compatibility problems is related to the idea of *morphing* graphs. For example, given two crossing-free geometric trees T_a and T_b, spanning the same set of n points, there is always a sequence of trees $T_0 = T_a, T_1, \ldots, T_{k-1}, T_k = T_b$, with $k \in O(\log n)$, in which every two consecutive trees are compatible [2]. Similar results have been obtained for other configurations, such as matchings [3, 15], and (in the negative) for pointed pseudotriangulations [4]. A major open problem in the area is whether this kind of sequences exist or not for polygonizations, or for spanning paths, of any given point set.

Spanning trees and compatible triangulations. There is another way, non obvious yet fundamental, to look at the parameter $d(T)$ of a plane geometric tree T. Let us add edges to T, while keeping plane the augmented geometric graph, until this addition of edges is not possible any more. In this way, we obtain a T-compatible triangulation Δ of S. Let $c_\Delta(T)$ be the number of connected components of the subgraph $\Delta - T$, formed by the edges added to T. Let $c(T) = \min_\Delta c_\Delta(T)$ be the minimum of the values $c_\Delta(T)$ taken over all possible T-compatible triangulations Δ. We state in Lemma 1 at the end of this section that $d(T) = c(T) - 1$. This is, $c(T) = d(T) + 1$ gives the minimum number of components that we can obtain in any spanning subgraph of $K(S)$ that uses edges (segments) that neither cross nor coincide with those of T. We will call the edges of the given tree T *black* edges, and the edges added to T for obtaining a compatible triangulation Δ *red* edges. The connected components of $\Delta - T$ will be called the *red components* of Δ. Therefore, our problem translates now to finding a T-compatible triangulation having a minimum number of red components. See Figure 1.

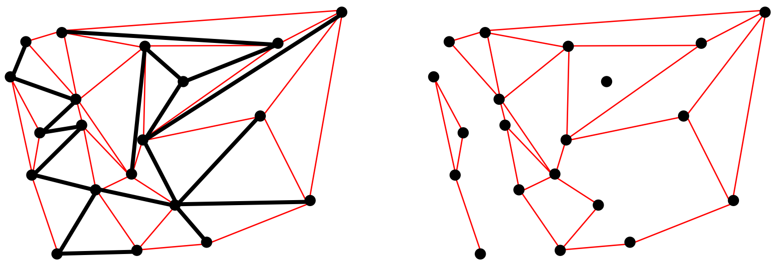

Figure 1: Left: A tree T and a compatible triangulation Δ. Right: The three components of the "red" graph $\Delta - T$ (one of them is an isolated vertex). We use thick lines for black edges, and thin lines for red edges; in this way the difference is still visible when a black and white printer is used.

Results

Our study of $d(T)$ is built up on the relationship between this value and the minimum number $c(T)$ of connected components in any spanning subgraph of $K(S)$ that uses edges (segments) that neither cross nor coincide with those in T, made precise in the following statement, proved in Section 2.2:

Lemma 1. *$d(T) = c(T) - 1$.*

The study of $c(T)$ requires the introduction of several definitions, intermediate results, and the proof of many technical lemmas, some of them with a non-short proof. Therefore, for the ease of understanding, we have made the option of stating in this section the most relevant results, and presenting the corresponding proofs in Section 2.

If u and v are two consecutive vertices of the convex hull of S, $CH(S)$, and $uv \notin T$, there is a unique path in p in T that connects u and v. The polygonal region bounded by uv and p is what we call a *pocket* of the tree T (roughly speaking, the precise definition is given in Section 2). The nature of the pockets, and the way they interact, is the back spine of the series of results that lead to Propositions 1 and 2, stated and proved in Section 2.6, which tell us that there are essentially two types of trees: On one hand, there are the trees that admit a compatible triangulation Δ having one big red component containing at least $\frac{3n}{4}$ of the vertices, while all the other components, if any, are interior isolated vertices; on the other hand, there are the trees in which $\Delta - T$ has exactly two components, each one containing at least one vertex from $CH(S)$. The triangulations of the latter type have $d(T) = c(T) - 1 = 1$, while this parameter clearly depends on the number of isolated vertices for the first type of triangulations. On the light of the preceding considerations, one can see why the following theorem, proved in Section 2.7, plays a crucial role in our analysis of $d(T)$:

Theorem 1. *Let T be a geometric tree. Then there is a T-compatible triangulation Δ' such that $\Delta' - T$ contains at most $\frac{n-3}{4}$ isolated vertices in addition to a unique large component. On the other hand, there are geometric trees T, such that any T-compatible triangulation Δ gives at least $\frac{n-2}{5}$ isolated vertices in $\Delta - T$.*

Combining Theorem 1, the preceding discussion, and Lemma 1, we immediately obtain our main result, which provides combinatorial bounds for d_n:

Theorem 2. *For every integer $n \geq 3$ the following inequalities hold:*

$$\frac{n-2}{5} \leq d_n \leq \frac{n-3}{4}.$$

In other words, for every set S of n points and every plane geometric tree T spanning $K(S)$, there is a compatible tree T' sharing at most $\frac{n-3}{4}$ edges with T, and there are some sets S with $|S| = n$ and plane geometric spanning trees T such that any T-compatible tree has at least $\frac{n-2}{5}$ edges in common with T.

In the section in which we prove Theorem 1 we provide as well some evidence that the tight value of d_n appears to be close to the lower bound in Theorem 2, which leads us to formulate the following conjecture:

Conjecture 1. *For some constant c, we have $d_n = \frac{n}{5} + c$, for every integer $n \geq 3$.*

As for results, in the last section of the paper we prove two additional theorems that are easily derived from the generic framework we construct, but that we consider interesting on their own:

Theorem 3. *Let T be a simple spanning path. Then, we can find a T-compatible triangulation Δ such that the number of components of $\Delta - T$ is at most 2, and hence $d(T) \leq 1$.*

Theorem 4. *Let Γ be a simple polygon with vertex set S. Then, we can always find a Γ-compatible triangulation Δ (which in general will have edges inside and outside Γ), such that the number of connected components of $\Delta - \Gamma$ is at most 3. Moreover, if Γ is non-convex, this minimum achievable number of components is 1 or 2.*

2 Proofs

2.1 Basic definitions. Organization of the section

Let us remind first some common definitions and notations for simple polygons. Let $q_1, \ldots, q_n$ be the vertices of a simple polygon P given in clockwise order; the arithmetic of their indices is done modulo n. We say that the vertex q_i is *convex* if the angle between the vectors $\overrightarrow{q_i q_{i-1}}$ and $\overrightarrow{q_i q_{i+1}}$ is less than π, and that it is *reflex* otherwise. A diagonal $q_i q_j$, when we consider it as oriented from q_i to q_j, divides the polygon P into two polygons: the *left polygon* P_L, whose vertices are the vertices of P from q_i to q_j in clockwise order, and the *right polygon* P_R, whose vertices are the vertices from q_j to q_i in clockwise order. If we place the maximum possible number $(n-3)$ of non-crossing diagonals inside P, we obtain a *triangulation* Δ of P. We will always assign red color to the diagonals of the triangulations we are using.

We call *weakly simple polygon* a plane geometric graph consisting of the union of a cycle with trees rooted at the vertices of the cycle, having all their additional vertices in the internal face of the cycle. See Figure 2, left. If we traverse clockwise the boundary of the unique bounded face of the weakly simple polygon, some vertices are visited several times before reaching the starting point of the traversal, as shown in Figure 2, center. However, the list of visited vertices, with the repetitions, behaves a as a simple polygon regarding many geometric properties (in particular regarding its triangulation with internal diagonals), which is obvious if one thinks of a weakly simple polygon as a proper simple polygon by splitting the multiple vertices into points that are infinitesimally close as shown in Figure 2, right.

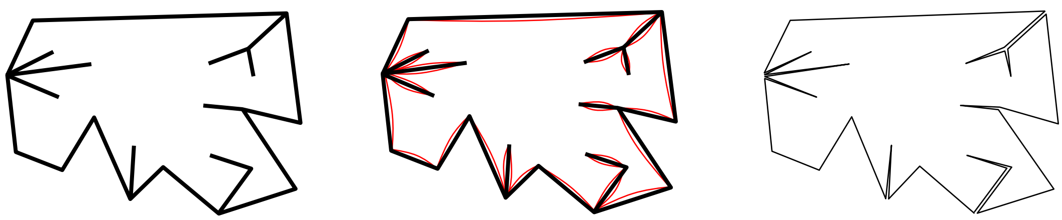

Figure 2: Left: A weakly simple polygon P. Center: traversal of the inner boundary of P. Right: simple polygon equivalent to P.

We will denote by $\{v_1, \ldots, v_n\}$ the set of vertices of a point set S, which will also be the vertices of any crossing-free geometric tree T spanning S. The convex hull of S will be denoted by $CH(S)$. If a vertex of S is also a vertex of $CH(S)$, we will also denote this point by q_i, in such a way that the vertices of $CH(S)$ are $q_1, \ldots, q_k$, in clockwise order. If $e_i = q_i q_{i+1}$ is not an edge of T, adding e_i to T yields a unique bounded face; the red edge e_i and the black interior boundary of the face form a weakly simple polygon. We call these weakly simple polygons the *pockets of* T. Observe that any T-compatible triangulation Δ is obtained by triangulating all the pockets of T.

The pockets of T can be seen as a decomposition of the convex polygon $CH(S)$ into (weakly) simple polygons. In order to simplify some forthcoming proofs, we have to consider some very similar yet slightly more general situation:

Let Q be any simple polygon with vertices $q_1, \ldots, q_k$ in clockwise order, and let T be any non-crossing geometric tree spanning all the vertices in Q and having possibly some other vertices that are points inside Q, satisfying the condition that all the edges of T are either inside Q or are sides of Q. As before, each edge e_i of Q that is not an edge of T defines a pocket of T. If $P_1, \ldots, P_h$ are the pockets of T, we also require that

Property $(\mathcal{W})$: If a vertex $q_i \in Q$ belongs to a pocket P_j, then q_i is a convex vertex of the weakly simple polygon P_j.

In this situation a T-compatible triangulation Δ will be a triangulation of all the pockets of T, that is to say, only edges inside Q or on the boundary of Q will be considered to obtain Δ. See Figure 3. The number of red components of Δ will be $c_\Delta(T)$ and $c(T)$ will be the minimum among all $c_\Delta(T)$. Usually, in the figures we will represent the edges of Q as red curved segments, specially if they coincide with edges of T (always drawn in black). Throughout the paper, when we say that a polygon (in general denoted as $Q, Q', Q_1, \ldots$) *encloses a tree*, we assume that all the vertices placed on that polygon satisfy condition $(\mathcal{W})$.

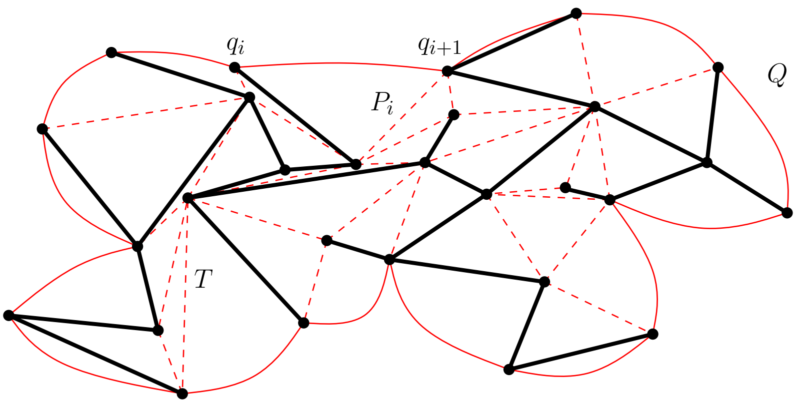

Figure 3: A tree T enclosed in a polygon Q. The vertices of Q have to be convex in each incident pocket.

The remainder of this section is organized as follows. We start with the proof of Lemma 1. Then, in Subsection 2.3, we study the case in which all pockets are convex polygons. The case of all the pockets being simple polygons, or equivalently, that all the leaves of the geometric tree T are on the boundary of Q, and the general case of pockets being weakly simple polygons, are studied in Subsection 2.4. Subsection 2.5 focuses on general properties of all T-compatible triangulations. The main results relating trees and compatible triangulations are given in Subsection 2.6. Finally, we consider in Subsection 2.7 the bounds on the number of red components of a T-compatible triangulation Δ, given a plane geometric tree T, this is, Theorem 1, from which Theorem 2 has been derived in the Introduction, and we conclude in Subsection 2.8 with some observations on special cases, namely the proofs of Theorem 3 and Theorem 4.

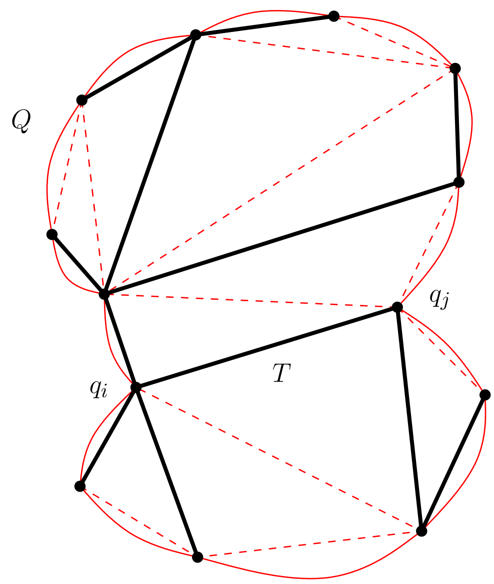

Figure 4: A tree with convex pockets and vertices on Q.

2.2 Proof of Lemma 1

Proof. We have to show that $d(T) = c(T) - 1$. Let T' be a spanning tree having $d(T)$ common edges with T. The $n - 1 - d(T)$ red edges of T' form a spanning forest consisting of $d(T)+1$ red trees, some of which might be singletons, and form together a spanning red subgraph of T', with $d(T)+1$ red components. If we add more red edges until completing $T \cup T'$ to a compatible triangulation, the number of red components cannot increase, hence $c(T) \leq d(T) + 1$.

Now, suppose that Δ is a T-compatible triangulation that has $c(T)$ red components $C_1, \ldots, C_{c(T)}$. Since Δ is connected, and there are no red edges connecting any two red components C_i and C_j, there must exist $c(T) - 1$ black edges between different red components yielding a connected spanning graph M, with $c(T) - 1$ black edges. Therefore, any spanning tree T' of M will be compatible with T and will have at most $c(T) - 1$ black edges, therefore $d(T) \leq c(T) - 1$. □

2.3 Trees with convex pockets

Let T be a tree inside or on the boundary of a simple polygon Q, with all the pockets being convex polygons. We prove in the following lemmas that in this case $c(T)$ is either 1 or 2, and that we can determine the correct value for any such tree.

Lemma 2. *Let Q be a simple polygon with vertices $q_1, \ldots, q_k$, with $k \geq 3$, and let T be a geometric tree having the same set of vertices as Q, whose edges are either sides or internal diagonals of Q, and such that all the pockets of T are convex. See Figure 4. Then:*

- *i) There is no T-compatible triangulation of Q having only one red component.*
- *ii) Given any edge $e = q_i q_j$ of T, there is a T-compatible triangulation Δ of Q that has exactly two red components, one containing vertex q_i, and the other one containing vertex q_j.*

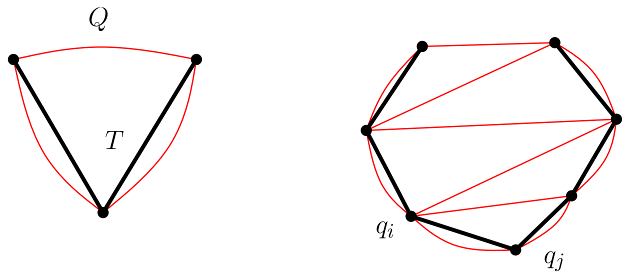

Figure 5: Left: Base of the induction. Right: Case in which T is a convex path.

Proof. *i*) The polygon Q has k vertices and any triangulation of Q consists exactly of $2k-3$ edges, namely $k-3$ internal diagonals and the k sides of Q. Hence, a T-compatible triangulation Δ contains $k-1$ black edges and $k-2$ red edges. Therefore, any spanning tree T' of Δ contains at least one black edge, which implies that $d(T) \geq 1$ and consequently $c(T) \geq 2$.

Let us prove *ii*) by induction on k. For $k = 3$ the result is obvious. See Figure 5. Let us assume that $k \geq 4$, and that the result is true for polygons Q' having less than k vertices. First, suppose that some diagonal $d = q'_i q'_j$ is an edge of T (d might coincide with e). Then d divides Q into two polygons Q_1 and Q_2, each one containing a spanning subtree, in which we can apply induction. Suppose that e is in polygon Q_1. By induction, we can triangulate Q_1 in a way that one component A_1 contains q_i and the other component B_1 contains q_j. In the same way we can triangulate Q_2 having one component A_2 containing q'_i and the other component B_2 containing q'_j. These triangulations of Q_1 and Q_2 form together a compatible triangulation Δ of Q. In addition, if q'_i and q'_j are in the same component of Q_1, say A_1, then Δ has precisely two components, $A_1 \cup A_2 \cup B_2$ containing q_i, and B_1 containing q_j. In the same way, if q'_i and q'_j are in different components of Q_1, say A_1 and B_1, respectively, then Δ has again two components, $A_1 \cup A_2$ containing q_i, and $B_1 \cup B_2$ containing q_j. Finally, if all the edges of T are edges of Q, then there is only one pocket. Therefore, Q must be a convex polygon, and T consists of all sides of Q but one. A compatible triangulation verifying ii) is easily obtained now, as shown in Figure 5. $\square$

Let us now consider the other possible case, in which the tree has vertices inside the polygon Q.

Lemma 3. *Let Q be a simple polygon with vertices $q_1, \ldots, q_k$, $k \geq 3$, and let T be a geometric tree having as set of vertices $\{q_1, \ldots, q_k\} \cup \{p_1, \ldots, p_s\}$, where the p_i's are interior to Q, in such a way that the edges of the tree are in the interior of Q, and all the pockets are convex. See Figure 6. Then:*

i) If T is a star, then $s = 1$ and $c(T) = 2$.

ii) If T is not a star, then $c(T) = 1$.

Proof. *i*) Notice that since all the pockets are simple polygons, the leaves of T must be vertices of Q. Then, if T is a star, the only interior point is the center of the star and the union of Q and T is the only compatible triangulation of Q, and has two components.

ii) If T is not a star every T-compatible triangulation Δ of Q will consist of $2s+k-2$ triangles. Let us prove *ii*) by induction on this number of triangles. If $2s+k-2 = 4$, the

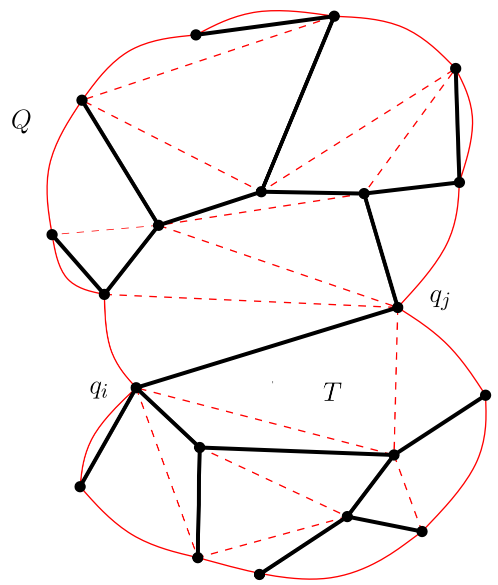

Figure 6: A tree with convex pockets and some vertices inside Q.

base case, the only possibility is shown in Figure 7 (left), and the statement of the lemma holds.

As a first case, let us assume that some diagonal $d = q_iq_j$ is an edge of T. Then, d divides Q into two subpolygons Q_1 and Q_2. If both Q_1 and Q_2 contain points in their interior, as d is an edge of both Q_1 and Q_2, neither the subtree T_1 placed in Q_1 nor the subtree T_2 placed in Q_2 can be a star, and the result follows by induction. On the contrary, if for example Q_1 does not contain points in its interior, we can triangulate the points in Q_2, using the hypothesis induction, with only one red component A_2. By the previous lemma, we can triangulate Q_1 with two components: A_1, containing q_i, and B_1, containing q_j. Since the points q_i and q_j are in $Q_1 \cap Q_2$, all the points can be connected using red edges, and therefore there is only one red component.

The second and last case happens when T does not contain diagonals of Q as edges. Then, let us consider an arbitrary interior point p_1. Since T is not a star, p_1 belongs to a convex pocket P_i that is not a triangle. Suppose pocket P_i is defined by two consecutive points q_i and q_{i+1} of Q and that the point on P_i next to q_{i+1} is p (see Figure 7, center and right). We can suppose, changing the orientation of Q if necessary, that $p_1 \neq p$.

Then, if $p \in Q$, it has to be the point q_{i+2}. In this case, the tree $T' = T - q_{i+1}q_{i+2}$ is contained in the polygon Q' obtained by skipping the point q_{i+1}. Therefore, if T' is not a star, by induction we can build a T'-compatible triangulation Δ' with only one red component. Adding to Δ' the black edge $q_{i+1}q_{i+2}$ and the red edge q_iq_{i+1}, we obtain a T-compatible triangulation Δ of Q having only one component. On the contrary, if T' is a star, there is only one T'-compatible triangulation Δ'. In this case a T-compatible triangulation Δ is obtained by removing the edge q_iq_{i+2} from Δ', and inserting the red edges q_iq_{i+1}, p_1q_{i+1} and the black edge $q_{i+1}q_{i+2}$.

Finally, if $p \notin Q$, let us consider the polygon $Q' = q_1, \ldots, q_i, p, q_{i+1}, q_{i+2}, \ldots$. Then, Q' contains T, has one triangle less than Q and contains the interior point p_1, so the result follows by induction. $\square$

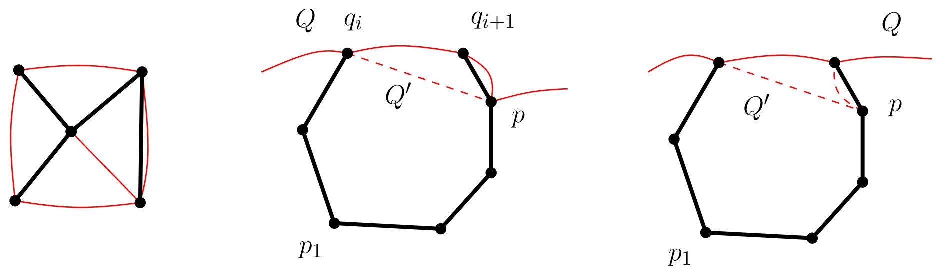

Figure 7: Left: Base of the induction. Center and right: How to obtain Q' and T'.

2.4 Triangulating general pockets

As already stated, in order to find a T-compatible triangulation Δ, where T is a geometric tree included in a polygon Q, we have to triangulate the pockets of T. In the previous section we have studied the case in which all pockets are convex polygons; here we tackle the general case: We consider first pockets that are simple polygons, and finally pockets whose boundary is a weakly simple polygon. Each pocket P of T is defined by some red side $e_j = q_j q_{j+1}$; the other vertices of the pocket will denoted as $v_1, v_2, \cdots$ throughout this section. Observe that all the red diagonals in a triangulation of a pocket always are in one component of $\Delta - T$ (this is easily seen: The internal diagonals in any triangulation of a simple polygon Q' form a connected graph if the edges of Q' are removed).

2.4.1 Simple Pockets

In this subsection a pocket P will always be a simple polygon with all its edges being colored black with one exception, a distinguished edge joining consecutive convex vertices q_j and q_{j+1}, whose color is red. Abusing slightly the terminology, the special red edge $e_j = q_j q_{j+1}$ is also considered to be a red diagonal, and hence when we later talk about the set of red diagonals, e_j is included.

Given a pocket P, suppose that we delete all the sides of P linking consecutive convex vertices (see Figure 8). This operation in general splits P into several connected components, some of which can be isolated vertices, namely convex vertices of P such that its two neighbors on the boundary P are also convex; we call them *convex isolated vertices* of P. The other components are paths beginning and ending in convex vertices. If one of these paths alternates convex and reflex vertices we say that it is a *zigzag*. If the path is not a zigzag, i.e., it contains at least two consecutive reflex vertices, we say that it is an *ordinary path*.

Definition 1. *We say that a pocket P is* Q-convex *if it is either a convex polygon, or for each reflex vertex v_i of P, the rays $\overrightarrow{v_{i-1}v_i}$ and $\overrightarrow{v_{i+1}v_i}$ hit the polygon at the red side $q_j q_{j+1}$. A pocket that is not Q-convex will be called* ordinary *pocket.*

Figure 9 shows a Q-convex pocket and an ordinary pocket. Notice that a Q-convex pocket P cannot have two consecutive reflex vertices and, therefore, its convex vertices are either isolated convex vertices of P or they lie in zigzags.

Let us first observe an elementary property of convex polygons with a special red side $q_j q_{j+1}$ (see Figure 10 (left)): Given any vertex $v_i \neq q_j, q_{j+1}$, there is a triangulation of the polygon using red diagonals in which all vertices but v_i lie all in a unique red component. Furthermore, this can be done with every diagonal having as endpoint one of the vertices v_{i-1} or q_{j+1}. The following lemma gives a similar result for Q-convex pockets.

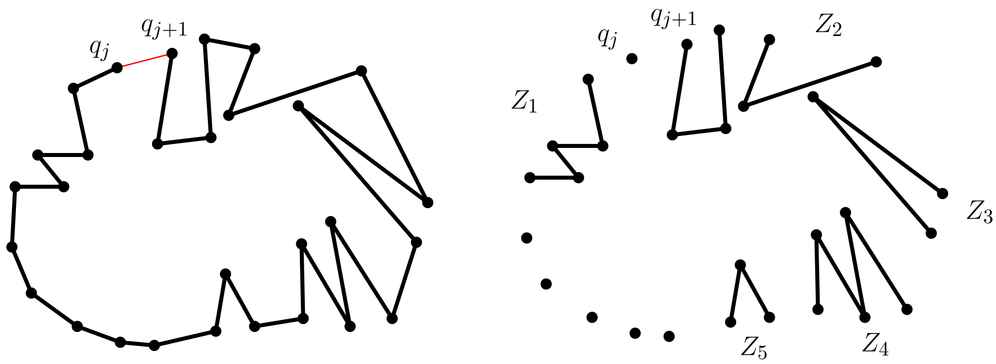

Figure 8: A pocket containing 5 zigzags, 6 isolated vertices and 1 ordinary path.

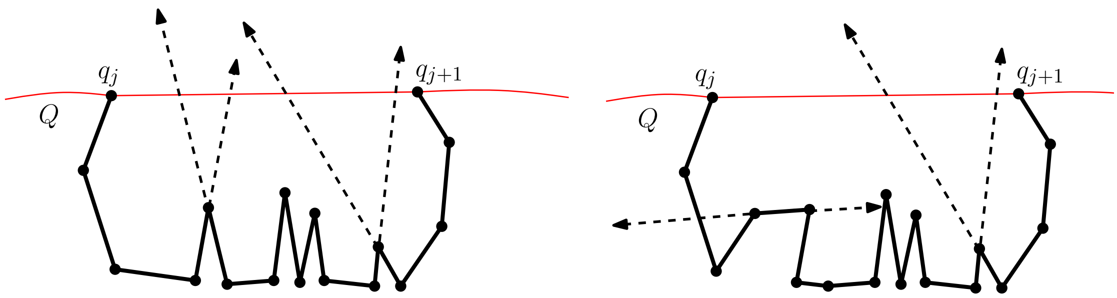

Figure 9: A Q-convex pocket (left) and an ordinary pocket (right).

Lemma 4. *Let P be a Q-convex pocket with red side q_jq_{j+1}, and let $r_1, \dots, r_s$ be the reflex vertices of P (in counterclockwise order from q_j, as in Figure 10). Then*

i) If we draw the diagonals $q_jr_1, r_1r_2, \dots, r_{s-1}r_s, r_sq_{j+1}$ we obtain $s+1$ convex polygons with a red side, letting aside the polygon that contains the edge q_jq_{j+1}.

ii) We can triangulate P in such a way that except for either one chosen isolated convex vertex u, or one chosen zigzag Z, every isolated convex vertex different from u is an endpoint of some diagonal and at least one convex vertex of each zigzag different from Z is an endpoint of some diagonal.

iii) In every triangulation of P either there is at least one isolated convex vertex that is not an endpoint of any diagonal, or there is at least one zigzag such that none of its convex vertices is an endpoint of any diagonal.

Proof. $i)$ Let us define $r_0 = q_j$ and $r_{s+1} = q_{j+1}$ and let u_k be the vertex next counterclockwise to r_k. If for $k = 0, \dots, s$ we rotate counterclockwise the ray r_ku_k we will keep finding convex vertices until we reach vertex r_{k+1}; therefore the region swept by the ray forms a convex polygon.

ii) Let us prove it by induction on the number of reflex vertices. If $s = 0$, then P is a convex polygon and it can be triangulated as before (Figure 10 (left)), skipping only a chosen isolated vertex u. If $s > 0$, let r be the reflex vertex of P closest to the line q_jq_{j+1}. Necessarily, the triangle q_jrq_{j+1} is empty and the two subpolygons, P_L and P_R, defined by the sides q_jr and rq_{j+1}, respectively, are Q-convex. As r lies in a zigzag Z of P and Z has points in both P_L and P_R, then $Z_L = Z \cap P_L$ and $Z_R = Z \cap P_R$ are shortest zigzags or isolated vertices in P_L and P_R, respectively. Assume that the isolated convex vertex u

(or the zigzag Z' different from Z) we want to skip belongs to P_L. Then, by induction, we can triangulate P_L skipping u (or Z') and P_R skipping Z_R. In this way, since Z_L has not been skipped in P_L, we can obtain a triangulation of P with only u (or Z') skipped. If we want to skip Z, then we triangulate P_L skipping Z_L and P_R skipping Z_R. In Figure 10 (right) we show how to achieve a triangulation leaving only one zigzag (Z_3) without red edges incident to any of its convex vertices.

iii) Let us prove that it is not possible to connect all the isolated convex vertices and at least one convex vertex for each zigzag using red edges. We proceed by induction. The base case, in which the pocket is a triangle is obvious. In the general case, in any triangulation Δ of the pocket P, the side $q_j q_{j+1}$ belongs to some triangle $q_j v q_{j+1}$. Let us consider the subpolygons defined by the sides $q_j v$ and $v q_{j+1}$. By the induction hypothesis, in each subpolygon there will be at least one isolated convex vertex without any incident red diagonal, or one zigzag in which all its convex vertices have no incident red diagonal. If v is reflex in P, perhaps these two unconnected zigzags or isolated vertices of each subpolygon produce only one unconnected zigzag in P, yet still at least one zigzag (this is, its convex vertices) remains unconnected with red edges.

□

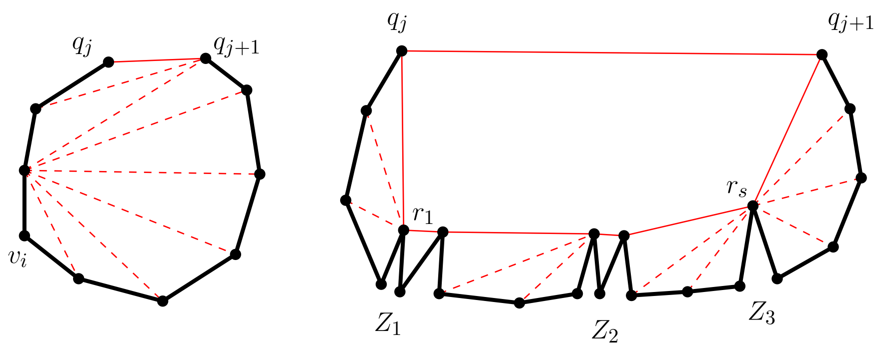

Figure 10: Triangulating a convex and a Q-convex pocket.

Let us observe that the triangulation obtained using the method explained in the proof of *ii)* has the property stated right after, which will be used in Lemma. For each diagonal xx' having one endpoint x in an isolated convex vertex or in a zigzag, there are three possibilities for the other endpoint x': (a) If a zigzag is skipped, then x' is a reflex vertex; (b) If the vertex u_0, counterclockwise next to q_j, is skipped, then x' is either a reflex vertex or q_j; and (c) If an isolated vertex v_i different from u_0 is skipped, then x' is either a reflex vertex or q_{j+1} or v_{i-1}.

We are ready now for proving the main result of this subsection:

Lemma 5. *Let P be an ordinary pocket with red side $q_j q_{j+1}$. Then, we can triangulate P in such a way that all the isolated convex vertices and at least one convex vertex of each zigzag are endpoints of diagonals. In addition, for each of these diagonals, the other endpoint is either a reflex vertex, or q_j, or q_{j+1}.*

Proof. First of all, notice that we only have to prove the lemma in the case in which P only contains zigzags and isolated convex vertices, because if P contains ordinary paths, we can transform each of them into a "smooth zigzag" by adding a dummy vertex infinitesimally near the midpoint of each side joining reflex vertices, as shown in Figure 11. Then, once we have obtained a set of diagonals satisfying the conditions of the lemma, we can delete

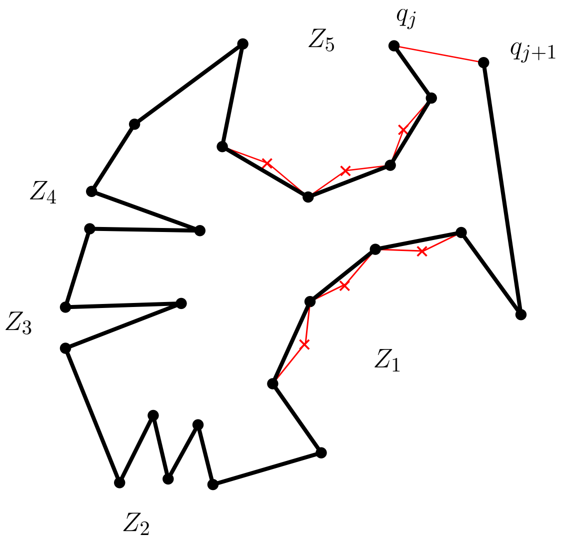

Figure 11: Transforming ordinary paths into zigzags.

the diagonals having as endpoints the dummy points, and there will still exist diagonals reaching the zigzags and isolated vertices of the original pocket.

Since P is not Q-convex, there is a reflex vertex r such that a ray emanating from r hits a black side $e = v_k v_{k+1} \neq q_j q_{j+1}$ in some point v', splitting the polygon into two subpolygons P_L and P_R (see Figure 12). By construction, the points r and v' belong to both subpolygons and are convex in both of them. Notice also that r lies in a zigzag Z of P, and that Z has points in both P_L and P_R. Therefore $Z_L = Z \cap P_L$ and $Z_R = Z \cap P_R$ are shorter zigzags, or isolated convex vertices, in each subpolygon, respectively. In general, we can choose several reflex vertices for the role of r and there are several ways of shooting the rays, but we can always choose the vertex r and the ray rv' in such a way that the subpolygon (P_L or P_R) that does not contain the red edge $q_j q_{j+1}$ has a minimum number of vertices. We can suppose, without loss of generality, that this minimum subpolygon is P_R and, therefore, $q_j q_{j+1} \in P_L$.

Now, taking in P_R for red edge the side rv', the minimality of P_R implies that P_R is a Q-convex pocket. This minimality also implies that if one endpoint of the edge $v_k v_{k+1}$ is reflex, then this reflex endpoint belongs to P_L, because the extension of the side $v_k v_{k+1}$ cannot cross the edge rv'.

Suppose first that one endpoint of $v_k v_{k+1}$ is reflex and belongs to some zigzag Z' of P, see Figure 12 (left). Consider then the zigzag Z'_L of P_L formed by the points of $Z' \cap P_L$ plus point v', and the zigzag (or isolated convex vertex) Z'_R of P_R formed by the points of $Z' \cap P_R$. Then, if P_L is an ordinary pocket, it contains less reflex vertices than P. Using induction on the number of reflex vertices we can draw diagonals reaching each zigzag and each isolated vertex. If we delete the diagonal reaching Z'_L, we still have diagonals for all the other zigzags and all the isolated convex vertices. Similarly, if P_L is Q-convex, we can use the same construction: By Lemma 4, we can triangulate P_L, skipping the zigzag Z'_L, in such a way that all the other isolated convex vertices and zigzags (including Z_L) are connected with red edges, while the opposite endpoint of each edge is a reflex vertex of P. On the other hand, as P_R is Q-convex, we can add diagonals (without using vertex v') for all the zigzags and all the isolated vertices except for the zigzag or isolated convex vertex Z_R. In particular, there will be a diagonal reaching a convex vertex of Z'_R. In summary, Z

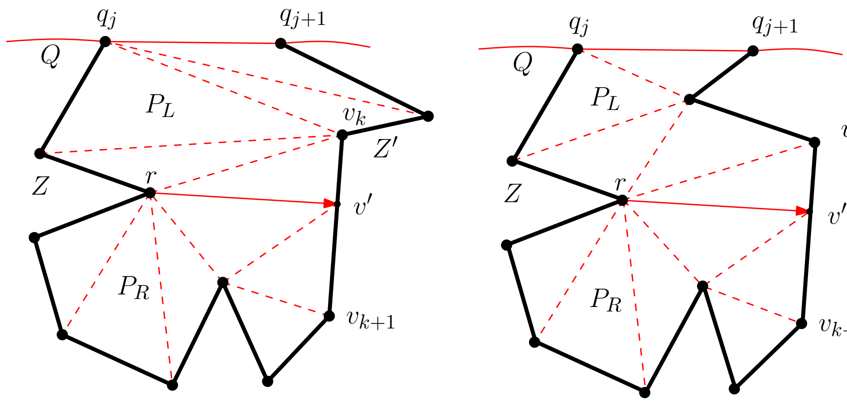

Figure 12: Illustration of Lemma 5: Vertex v_k is reflex (left). Both vertices of v_k and v_{k+1} are convex (right).

is reached by a diagonal of P_L, Z' by a diagonal of P_R, and the other zigzags and isolated convex vertices are reached by diagonals in the subpolygon in where they lie.

The second and last case, when both endpoints of the edge $v_k v_{k+1}$ are convex, are dealt with in a similar way. The only difference is that now we consider Z' to consist only of the vertex v'. Again, we apply induction in P_L, or we apply Lemma 4, skipping the vertex v' if P_L is Q-convex. In the last case, notice that, since in P_L the reflex vertex r is adjacent to v', the endpoints of diagonals are still either reflex vertices, or q_j, or q_{j+1}. In P_R we can use the same construction as in the previous case, just skipping Z_R, see Figure 12 (right). □

2.4.2 Weakly simple pockets

We consider here the most general case, in which the pockets are bounded by weakly simple polygons. In this case, the proper terminology would be to call them *weakly simple pockets*, but for conciseness we call them simply *pockets* throughout this section. On the other hand, we have already observed in Subsection 2.1 (see also Figure 2) that weakly simple polygons behave as simple polygons regarding their internal diagonals, therefore one would expect the results to be the same than in the preceding section. This is in fact the case, yet we briefly provide the statements and proofs for the sake of completeness.

Given a tree T inside a polygon Q, let $q_j q_{j+1}$ be a red edge of Q defining a pocket whose boundary is a weakly simple polygon P. Recall that this polygon P is obtained by visiting geometrically (beginning clockwise with the edge $q_j q_{j+1}$, say) all the edges of the pocket, and that a vertex of a pocket can appear several times in P. See Figure 13. In this case, for a zigzag, ordinary path, or isolated convex vertex we mean respectively a zigzag, ordinary path or isolated vertex of P. In the figure, the three zigzags, the two ordinary paths and the five isolated convex vertices defined by P are shown. Observe that a same vertex of the pocket can appear several times in the same zigzag, can belong to different paths or zigzags, and even can define several isolated convex vertices of P. For simplicity, we also denote the pocket by P.

Definition 2. *We say that a pocket is* Q-convex *if the weakly simple polygon* P *is* Q-*convex, that is, if it is either a convex polygon, or for each reflex vertex* v_i *of* P*, the rays* $\overrightarrow{v_{i-1}v_i}$ *and* $\overrightarrow{v_{i+1}v_i}$ *hit* P *at the red side* $q_j q_{j+1}$*. A pocket that is not* Q*-convex will be called* ordinary *pocket. Figure 14 shows a* Q*-convex pocket and an ordinary pocket.*

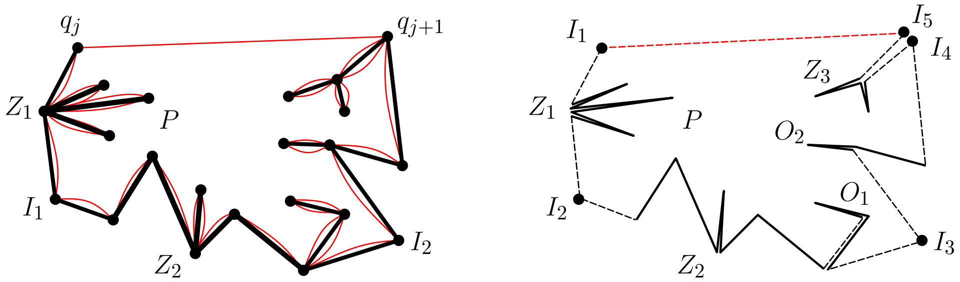

Figure 13: The zigzags, ordinary paths and isolated convex vertices of P.

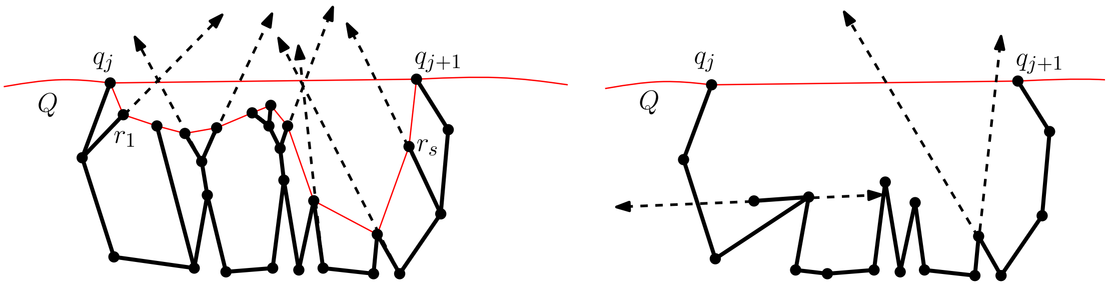

Figure 14: A Q-convex pocket (left) and an ordinary pocket (right).

Lemma 6. *Let P be a Q-convex pocket with red side q_jq_{j+1}, and let $r_1, \ldots, r_s$ be the reflex vertices of P in counterclockwise order from q_j, as in Figure 14. Then:*

i) If we draw the diagonals $q_jr_1, r_1r_2, \ldots, r_{s-1}r_s, r_sq_{j+1}$, we obtain $s+1$ convex polygons with a red side, letting aside the polygon containing the edge q_jq_{j+1}. See Figure 14.

ii) We can triangulate P in such a way that except for either one chosen isolated convex vertex u, or one chosen zigzag Z, every isolated convex vertex different from u is an endpoint of some diagonal and at least one convex vertex of each zigzag different from Z is an endpoint of some diagonal.

Proof. The proofs are identical to those of Lemma 4. However, we must notice that two of the polygons obtained in *i)* can have more than one vertex in common, because consecutive polygons can share some vertices. Therefore, in some cases, all isolated convex vertices and all zigzags can be endpoints of diagonals. In other words, sometimes part *iii)* of Lemma 4 is not true. □

A version of Lemma 5, for the case of weakly simple pockets, also holds:

Lemma 7. *Let P be an ordinary pocket. Then, we can triangulate P in such a way that all isolated convex vertices and at least one convex vertex of each zigzag are endpoints of diagonals.*

Proof. The proof is identical to that of Lemma 5. □

2.5 Additional properties of T-compatible triangulations

Let again T be a geometric tree enclosed in a simple polygon Q. In this section we deep more in the properties of any arbitrary T-compatible triangulation Δ, in order to obtain some technical results that will be used in the following section.

Definition 3. *Let Δ be a T-compatible triangulation of the interior of Q. A sequence of adjacent triangles $B = F_1, F_2, \ldots, F_l$ of Δ, is called a* separating band *if it satisfies the following conditions:*

1) *Each triangle F_i has a black edge e_i in common with F_{i+1}, $i = 1, \ldots l-1$.*

2) *Additionally, F_1 contains another black edge e_0 and F_l contains another black edge e_l, and these two edges e_0 and e_l are sides of Q.*

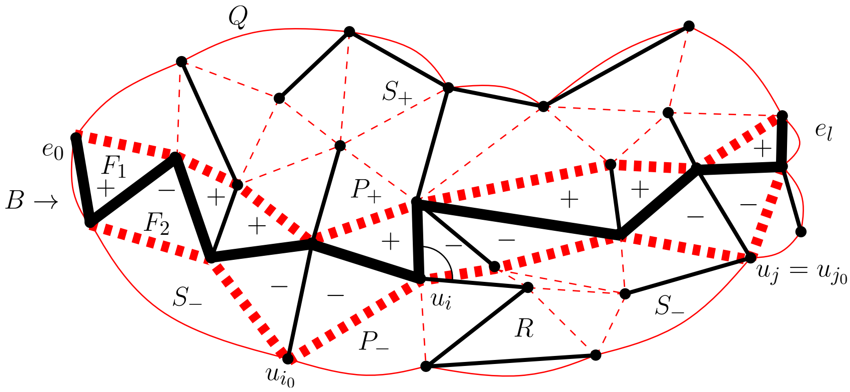

Figure 15: A separating band B of a compatible triangulation.

A separating band in a triangulation of a tree inside a polygon is shown in Figure 15. By construction, the separating band $B = F_1, F_2, \ldots, F_l$ has $l+1$ black edges and l red edges given by the constituent triangles. Let us describe some elementary properties of a separating band. First of all, we give a rule for assigning a $+$ or $-$ sign to each triangle in the band. We start by giving sign $+$ to F_1. In the generic step, let us assume that signs have been defined for $F_1, \ldots, F_i$. Then, if the common vertex between e_{i-1} and e_i is the same as the common vertex between e_i and e_{i+1}, we assign to F_{i+1} the same sign as F_i, while we assign the opposite sign if those common vertices are different. By construction all the red edges of the positive triangles form a path P_+, and similarly a red path P_- with the red edges of the negative triangles. These two paths together with the edges e_0 and e_l form a simple polygon, the *boundary of the band*. When all the triangles have the same sign (say $+$) or when the sequence is formed by only one triangle with two black edges on Q, then we take P_- as a path of length 0 formed by the unique vertex of the band not in P_+.

A separating band B partitions the vertices of Δ into two sets (the *sides of the band*), those that lie in the same zone than the points of P_+, which will denote it by S_+, and those on the side of the points of P_-, which we denoted by S_-. As there are no red edges between S_+ and S_-, any component C_i of Δ is completely included in one of the two sides of a separating band. Let us call *spine of the band* the black path formed by first the edge e_0, then the common edges of consecutive triangles having different sign, and finally edge e_l. The spine is a path that alternates between points of each side.

The following lemma makes clear why separating bands are a useful tool.

Lemma 8. *Let Δ be a T-compatible triangulation of the interior of Q, and let $e = v_iv_j$ be an edge of T, such that v_i and v_j are in different red components C_i and C_j of Δ, respectively. Let S_{ij} be the set of edges of Δ with an endpoint in C_i and the other in C_j. Then, either one of the components, say C_i, consists of only one vertex v_i, and S_{ij}*

is formed by all the edges incident to v_i, or S_{ij} consists of the black edges in a separating band of Δ.

Proof. Let us label the points of S as *type* i when they belong to the component C_i of Δ. Any edge linking in Δ points with different labels must belong to T, and since T is a tree, every triangle of Δ has at least two vertices with the same label. Therefore, given any triangle $v_1v_2v_3$ of Δ, if v_1 has label i_1 different from the label i_2 of v_2, then the label of v_3 must be either i_1 or i_2, being v_2v_3 black and v_1v_3 red in the first case, and v_1v_3 black and v_2v_3 red in the second case.

The edge $e = v_iv_j$ belongs to at least one triangle F_1 of Δ. The other two edges of F_1 cannot be both black because T is a tree, and cannot be both red either because v_i and v_j are in different red components. Thus, if v' is the other vertex of F_1, one of the edges v_iv' and v_jv' has to be red, and the other one black. Suppose, for example, that the black edge is $e_1 = v_iv'$, and let us give sign $+$ to this triangle. If e_1 is not an edge of Q, it belongs to other triangle F_2 with another additional black edge e_2. If this edge e_2 is not in Q, we can again obtain a new triangle F_3 containing e_2 and another black edge e_3, and so on. In the same way that we have described after the definition of separating bands, we assign to F_{i+1} same or the opposite sign than F_i, depending on whether the common vertex between e_i and e_{i+1} is the same or different to the previous common vertex. This process of incrementally gluing triangles can finish only in two ways: Either in the last explored triangle F_l the last edge e_l is on Q, in which case triangle F_l can be glued only with F_{l-1}, or F_l can be glued with F_1 by the edge e_l. No other option is possible because a triangle with two black edges can be glued to at most two other triangles.

Let us analyze first the case in which after F_l we encounter F_1 again. This case may happen when $F_1, \dots, F_l$ are all the triangles placed around the interior vertex v_i; in fact, it is the only situation in which it can arise: If in the cyclic sequence $F_1, \dots, F_l, F_1$ there are triangles with different signs, consider the edges $e_{i_1}, e_{i_2}, \dots e_{i_s}$ adjacent to consecutive triangles with different sign, with the arithmetic of indices modulo s, so that $e_{i_{s+1}} = e_{i_1}$. Let us consider a visit of these triangles in the order, $F_1, \dots, F_l, F_1$, and let us orient the edges e_{i_r} from the vertex in C_i to the vertex in C_j, if e_{i_r} when crossing the edge means going from a positive triangle to a negative triangle, and in the opposite direction for the reverse transition of signs. Since e_{i_r} and $e_{i_{r+1}}$ share a vertex and alternately visit points from one component and the other, the edges $e_{i_1}, e_{i_2}, \dots e_{i_s}$ have to form a cycle, contradicting that T is a tree.

Notice that if v_i is a vertex of Q, some of the triangles adjacent to v_i have only one neighbor, therefore this cyclic case cannot happen if the initial edge $e = v_iv_j$ belongs to Q.

Let us conclude with the case in which the incremental gluing of triangles finishes because e_l is an edge of Q. If the initial edge $e = v_iv_j$ is also an edge of Q, then the sequence $B = F_1, \dots, F_l$ is clearly a separating band, and since all the vertices of C_i have to be on one side, S_+ or S_-, of the band, and the vertices of C_j on the other side, there are no other black edges connecting points of C_i and C_j. If the initial edge $e = v_iv_j$ does not belong to Q, we can repeat the process starting with the last edge e_l, which is on Q, and obtain in this way a separating band. $\square$

From the previous lemma, we can infer that a T-compatible triangulation Δ can contain two types of red components, those formed by exactly one isolated vertex placed inside Q, and those containing vertices of Q, separated one from each other by separating bands. We will call these latter components *the big components* of Δ (despite the name, notice that a big component might consist of only one vertex of Q).

We present next two more lemmas that we are using in the forthcoming sections. Let Δ be a triangulation containing the geometric tree T, and let v be a vertex of T having degree at least 2 (in T). Suppose that e is a red edge incident to v, and that rotating e around v the first black edges we find are b_1 and b_2, clockwise and counterclockwise, respectively. We will say that point v is *convex with respect to edge* e if the clockwise angle from b_2 to b_1 is $< \pi$.

Lemma 9. *Let $B = F_1, F_2, \ldots, F_l$ be a separating band of Δ. Suppose that the path P_- consists of the vertices $u_1, u_2, \ldots, u_m$, and that not all of them are on Q. Then, there are two vertices u_i and u_j, $i < j$, of P_- satisfying:*

i) For $i < k < j$, u_k is not on Q and all the edges of T incident to u_k are in the band B.

ii) If u_i is not on Q, then it is convex with respect to $u_i u_{i+1}$, and if u_j is not on Q, then it is convex with respect to $u_{j-1} u_j$.

iii) If u_i and u_j are both on Q, then they are not consecutive on P_-.

See Figure 15. An equivalent statement can be formulated taking the path P_+ instead of the path P_-.

Proof. Since u_1 and u_m belong to Q and not all the vertices of P_- are on Q, there are at least two nonconsecutive vertices of P_-, u_{i_0} and u_{j_0}, such that $i_0 < j_0$, they are on Q and none of the vertices of P_- between them belongs to Q. Let R be the region defined by the two polygonal paths from u_{i_0} to u_{j_0}, one on Q and the other one on P_-. See Figure 15.

Then, consider the subsequence $u_{i_0}, u_{i_1}, \ldots, u_{i_s} = u_{j_0}$ beginning in u_{i_0}, ending in u_{j_0}, consisting of all the intermediate vertices u_{i_k} that are incident to some edge of T inside R. Note that s can be 1, and this subsequence consists of only two vertices, u_{i_0} and u_{j_0}. Let us visit these vertices u_{i_k} of P_- in order from 1 to $s-1$, and create a pre-label C or N before each such vertex u_{i_k}, according to whether the vertex is convex or non-convex with respect to edge $u_{i_k-1} u_{i_k}$, and create as well a post-label C or N after each vertex, according to whether it is convex or non-convex with respect to edge $u_{i_k} u_{i_k+1}$. By convenience, let us place a label C after vertex u_{i_0} and before u_{j_0}. Observe that, when P_- is traversed, if a vertex u_{i_k} is not convex with respect to edge $u_{i_k-1} u_{i_k}$, then it has to be convex with respect to edge $u_{i_k} u_{i_k+1}$. Therefore, we are necessarily finding two consecutive vertices u_{i_k} and $u_{i_{k+1}}$ (in the subsequence $u_{i_0}, u_{i_1}, \ldots u_{j_0}$), such that the pattern of labels placed after u_{i_k} and before $u_{i_{k+1}}$ is C-C. Taking these two last vertices as u_j and u_j, clearly they satisfy $i), ii)$ and $iii)$. □

Lemma 10. *Let B be a separating band of a T-compatible triangulation Δ, and let us suppose that not all the vertices of P_- are on Q. Then, there is a pocket P_i with at least four vertices, with red side $q_i q_{i+1} \in S_-$, such that either P_i contains an isolated convex vertex in P_+, or there is a zigzag of P_i with all its convex vertices in P_+.*

Proof. Figure 16 shows an example of the situation described in the statement, in which the zigzag starting at v_i and finishing at v_{j-1} is contained in the pocket P_i defined by the red edge $q_i q_{i+1}$.

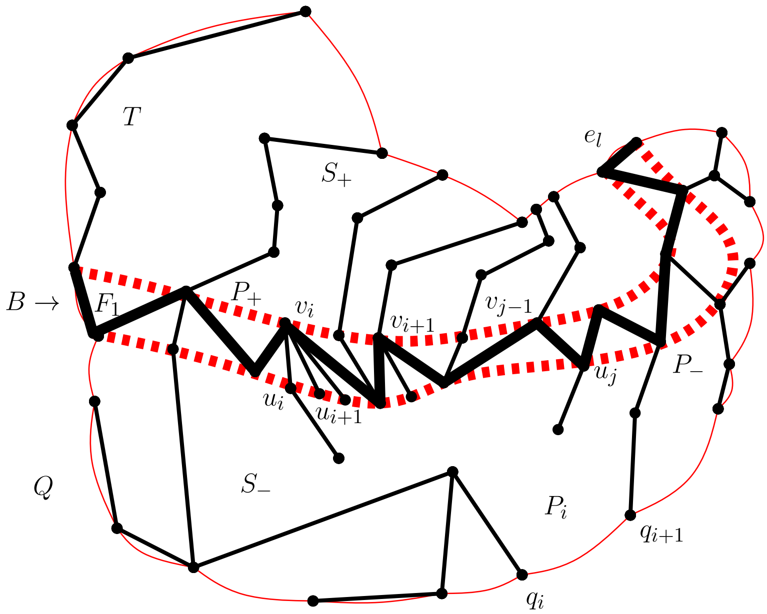

Figure 16: Illustration of Lemma 10.

As a first observation for the proof, let us consider any arbitrary vertex u_k of P_-, yet different from the last one. Let e'_k be the red edge $u_k u_{k+1}$, let F_k be the triangle of B containing e'_k, and let e_k be the black edge of F_k with endpoint u_k. Then, the endpoint e_k different from u_k must be a vertex $v_k \in P_+$. Moreover, if P_k is a pocket of T containing a set of consecutive triangles of B, $F_k, \ldots, F_{k+k'}$, with an edge on P_- and sharing vertex v_k, then v_k is convex in that pocket, because it is a vertex of all these triangles.

Now, due to Lemma 9, there are two vertices u_i and u_j, $i < j$ on P_- satisfying the conditions i), ii) and iii) of that lemma. Suppose that u_i and u_j are not on Q. Hence, u_i is convex with respect to the following red edge $e'_i = u_i u_{i+1}$, u_j is convex with respect to its predecessor edge $u_{j-1}u_j$ and all other vertices in P_- between u_i and u_j cannot have any incident edge placed outside of B. See Figure 16.

In the pocket P_i that contains the edge e'_i, clockwise after the convex vertex u_i, comes the convex vertex v_i, then vertex u_{i+1}, which in its turn is followed by another vertex v'_i that may coincide with v_i or may be vertex v_{i+1}. Again after v'_i comes vertex u_{i+2}, which is followed by another vertex of P_+, and so on. See Figure 16. Consider this sequence of vertices $u_i, v_i, u_{i+1}, v'_i, \ldots, v'_{j-1}, u_j$. Notice that in this sequence vertices of P_- and P_+ appear intertwined, and that the vertices of P_+ are convex in P_i. Therefore, since u_j has to be convex in P_i, the sequence contains a first vertex u_l of P_- (different from u_i) that is convex in P_i. Then, if $l = i+1$, v_i is an isolated convex vertex of P_i, and if $l \neq i+1$ (l can be j), the sequence $v_i, u_{i+1}, \ldots, v_{l-1}$ is a zigzag of P_i with all its convex vertices on P_+. Moreover, P_i has at least four vertices: q_i, q_{i+1}, u_i and u_j.

Finally, the argument is very similar for the other cases –u_i or u_j or both being on Q– because these vertices are convex in any pocket (property ($\mathcal{W}$) of the vertices of Q). □

2.6 The two main types of trees

Throughout this section we consider a geometric tree T included in a polygon Q, and assume that we are in the most general situation, this is, that all the pockets of T are

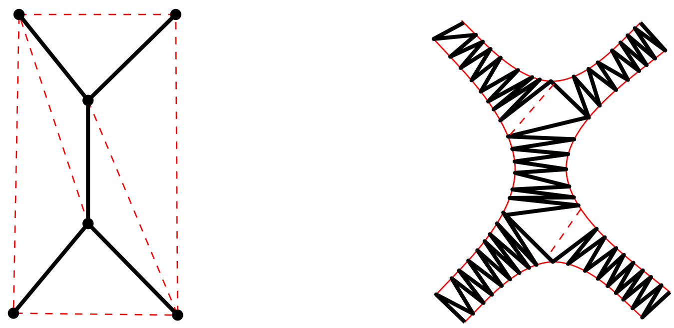

Figure 17: Examples of trees with Q-convex pockets, having one and two components (left and right, respectively).

weakly simple polygons, which in some case may be proper simple polygons. For simplicity we are calling them just *pockets* in all cases.

Proposition 1. *Let T be a given geometric tree included in a polygon Q. If all the pockets of T are Q-convex, then we can obtain a T-compatible triangulation of the interior of Q having either one or two red components.*

Proof. The result is a direct consequence of Lemmas 2, 3, and 6. We can use Lemma 6 to transform each non-convex pocket into several convex pockets, hence obtaining a tree with convex pockets placed inside a polygon Q. If all the vertices of T are placed on Q or if T is a star, by Lemma 2 we can triangulate the interior of Q obtaining two red components. Otherwise, using Lemma 3, we obtain only one red component. $\square$

Figure 17 shows two geometric trees with Q-convex pockets, and corresponding compatible triangulations for these examples with the minimum number of components.

The following corollary tells us when two red components are unavoidable in any T-compatible triangulation, for trees T with Q-convex pockets:

Corollary 1. *Let T be a given geometric tree included in a polygon Q. Suppose that all the pockets of T are Q-convex and that T is not a star. Then:*

- *i) If T contains a vertex v, not placed on Q, and convex in all its incident pockets, then there is a T-compatible triangulation with only one red component.*
- *ii) Otherwise, any T-compatible triangulation has at least two red components. Furthermore, given any edge $e = v_iv_j$ of T, there is a T-compatible triangulation with two red components, one of them containing v_i, and the other one containing v_j.*

Proof. i) We transform the pockets of Q into convex pockets of a new polygon Q' using Lemma 6. Then, vertex v is in the interior of Q'. Therefore, the result follows from Lemma 3.

ii) If T does not contain any vertex v satisfying condition *i)*, once we have obtained the transformation of pockets using Lemma 6, we obtain a T-compatible triangulation with two red components by Lemma 2, as claimed. In fact, we should also prove that if a Q-convex pocket is not triangulated decomposing this pocket into convex pockets in the usual way, then still we have two or more components, yet this result is easily obtained by induction. $\square$

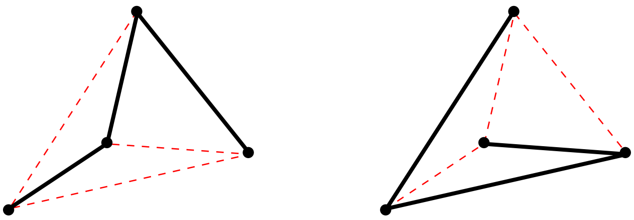

Figure 18: Base of the induction in Proposition 2.

Proposition 2. *Let T be a given geometric tree included in a polygon Q, having at least one ordinary pocket. Then, we can obtain a T-compatible triangulation Δ of the interior of Q, such that Δ only has one big red component, and all the other components, if any, consist of isolated interior vertices.*

Proof. If T has Q-convex pockets we can transform them as indicated in Lemma 6. Therefore, we can assume hereafter that all the pockets of T are either convex or ordinary pockets. We are proving by induction that it is possible to triangulate the pockets of T in such a way that all the vertices placed on Q are in the same red component. In this way, we are gluing all the big red components, because, by Lemma 8, any big component has a vertex on Q. For $n = 4$, Figure 18 shows the two only possibilities for a tree to have ordinary pockets and their corresponding triangulations.

Let us consider first the case in which there is an edge $e = v_1v_2$ of T such that v_1 and v_2 are nonconsecutive vertices of Q. This edge divides Q into two subpolygons Q_1 and Q_2 enclosing two subtrees T_1 and T_2, such that $T = T_1 \cup T_2$ and $T_1 \cap T_2 = e$. Since T contains at least one ordinary pocket, at least one of the trees T_1 and T_2 has an ordinary pocket. If both subtrees T_1 and T_2 contain ordinary pockets, then the result follows by induction. If one of the trees, say T_1, contains an ordinary pocket and the pockets of T_2 are convex, then by induction all the vertices of Q_1 (including v_i and v_j) can be placed in the same red component. By Lemma 2, all the vertices of Q_2 can be placed in two components, one containing v_i, the other one containing v_j. Therefore, in the T-compatible triangulation all the vertices of Q will be in the same red component, and we are done.

As a second possibility, let us consider the case in which there are no edges of T linking nonconsecutive points of Q, but T has a convex pocket P_i, with red side q_iq_{i+1}, which contains another vertex q_j of Q. If $q_j = q_{i+2}$, then necessarily q_{i+1} is leaf of T and we can apply induction in the subtree $T' = T - q_{i+1}q_{i+2}$ enclosed in the polygon $Q' = q_1, q_2, \ldots, q_i, q_{i+2}, \ldots$. Adding to this triangulation of Q' the red edge q_iq_{i+1} and the black edge $q_{i+1}q_{i+2}$, we obtain a triangulation of Q with all the vertices of Q in the same red component. The same argument applies if $q_j = q_{i-1}$. Suppose now that $q_j \neq q_{i-1}, q_{i+2}$, and that $i < j$. In this situation we consider the subpolygons $Q_1 = q_1, \ldots, q_i, q_j, q_{j+1}, \ldots, q_1$ and $Q_2 = q_j, q_{i+1}, q_{i+2}, \ldots, q_j$, enclosing two subtrees T_1, T_2, having in common only the point q_j, with $T = T_1 \cup T_2$ (see Figure 19, left). As in the previous case, if both T_1 and T_2 contain ordinary pockets, then the result follows by induction on T_1 and T_2 (enclosed in Q_1 and Q_2, respectively). If the assumption doesn't hold, we have to consider the case in which T_1 contains an ordinary pocket and T_2 contains only convex pockets. Now, all the vertices of Q_1 can be placed in one red component, and observe that the vertex v following q_{i+1} in the pocket P_i cannot be a vertex of Q_2 (otherwise, $q_{i+1}v$ would be a diagonal of Q). Since all the pockets incident to v must be convex, by Corollary 1 the vertices of T_2 can also be placed in one red component, except if T_2 is a star centered at v. But in this last case we can connect all the leaves of T_2, including q_{i+1} and q_j, and hence the result follows again.

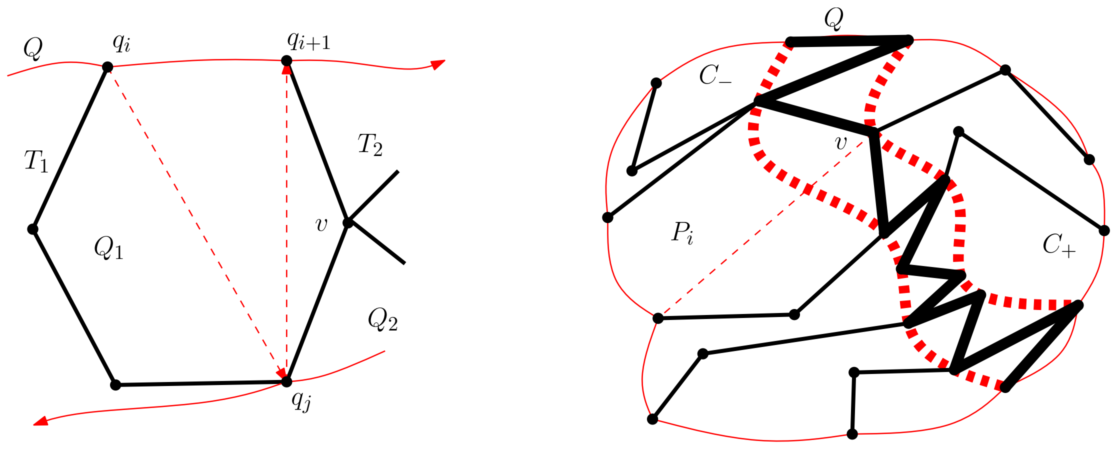

Figure 19: Illustration of cases for Proposition 2.

As a third and last case, suppose that T neither contains an edge linking two non-consecutive points of Q nor any convex pocket with at least three vertices of Q. Using Lemmas 5 and 7, we can built a T-compatible triangulation Δ in which all the ordinary pockets are triangulated in such a way that all the isolated convex vertices and all the zigzags (this is, at least one convex vertex of each zigzag) are linked by some diagonal. Suppose that not all the vertices of Q are in the same red component of Δ. Therefore, there is a separating band $B = F_1, \ldots, F_l$ such that each component is placed either in S_+ or in S_-. In particular we have a red component C_+ containing the vertices of the path P_+ of red edges and another red component C_- containing the vertices of P_-.

Now, suppose that not all the vertices of P_- are on Q. Then, by Lemma 10, there is a pocket P_i placed on S_- that contains either an isolated convex vertex in P_+ or all the convex vertices of a zigzag in P_+. Now, since in each ordinary pocket all the zigzags and all the isolated convex vertices are endpoints of red edges, the only possibility for P_i is to be a convex pocket: $P_i \cap S_+$ consists of only one vertex v, and P_i has been triangulated such that it contains two red components, one of them consisting only of the point v. In addition, notice that P_i must contain at least four vertices, and as P_i cannot contain three points of Q, v must be an interior vertex. See Figure 19 (right).

We can now change the triangulation of P_i in such a way that v lies in a bigger red component. In this way, we obtain new red components and new separating bands only retriangulating a convex pocket. Notice that all the vertices belonging to the component C_+, are in a new component C'_+ containing at least two more vertices. Again C'_+ can be separated from another component C'_- by a separating band B', and if not all the vertices of P'_- are on Q, we can build once again a new bigger red component C''_+.

We can continue doing the preceding process that increases the size of the component C_+, until either only one big component is obtained, or we have obtained a separating band B such that all the vertices of P_- are on Q. But this last case can never happen, because then not all the vertices of P_+ would be on Q and, therefore, again by Lemma 10, there would be a convex pocket P_i containing one vertex v in $S_- \subset Q$ and the remaining ones in S_+, that is, a pocket with at least three vertices of Q, which is not possible. $\square$

Notice that the method of the proof provides implicitly an algorithm for obtaining a compatible triangulation having all the vertices of Q in the same red component. The complexity of the algorithm depends on the specific way of computing and updating the red components of the triangulations, but even with a naive approach for these steps the complexity of the method is $O(n^2)$.

2.7 Proof of Theorem 1

Let us call *ordinary trees* those trees inside a polygon Q that have some ordinary pocket. We have seen that for ordinary trees T, there is always a T-compatible triangulation Δ such that all the vertices of Q are in the same red component of $\Delta - T$, the *big component*, and that other component, if any, would consist each of a single isolated vertex. We study closely here the number of isolated vertices that a T-compatible triangulation may admit as components in the red graph.

Notice first that for some trees components consisting of only one vertex are unavoidable. Figures 20 and 21 show two examples of that situation. In the tree of Figure 20 there are several *totally isolated vertices*, i.e., vertices of T not "seeing" any other vertex of T. Therefore, there are no red edges incident to these vertices in any compatible triangulation. In Figure 21 we show a tree in which either vertex v_1 or vertex v_2 have to be isolated in any compatible triangulation Δ, in spite of the fact that each of them can see some vertex of T.

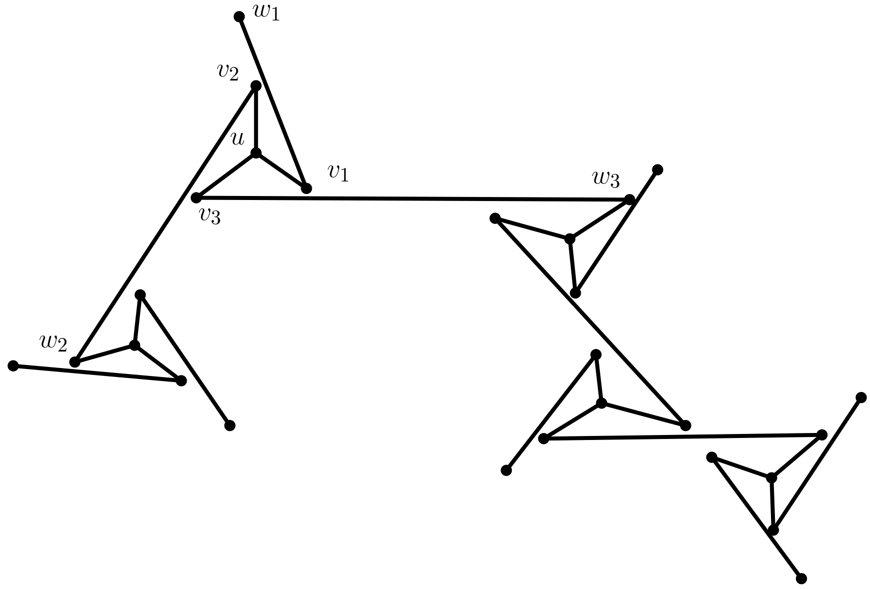

Figure 20: A tree with many totally isolated vertices.

Observe that, given an ordinary tree T and a T-compatible triangulation Δ with only a big component, each isolated vertex u_i of the red graph must be in the interior of Q; therefore, u_i must be incident to at least three edges of T. Edges incident to an isolated vertex u_i are different from edges incident to any other isolated vertex u_j, because if u_iu_j were an edge of T this edge would have to belong to a triangle of Δ with at least a red side, contradicting that both u_i and u_j are isolated. Therefore, the number of isolated vertices is at most $(\frac{n-1}{3})$. In fact, we can slightly improve on this upper bound.

To that goal, let us see first that we can assume that T satisfies the following property:

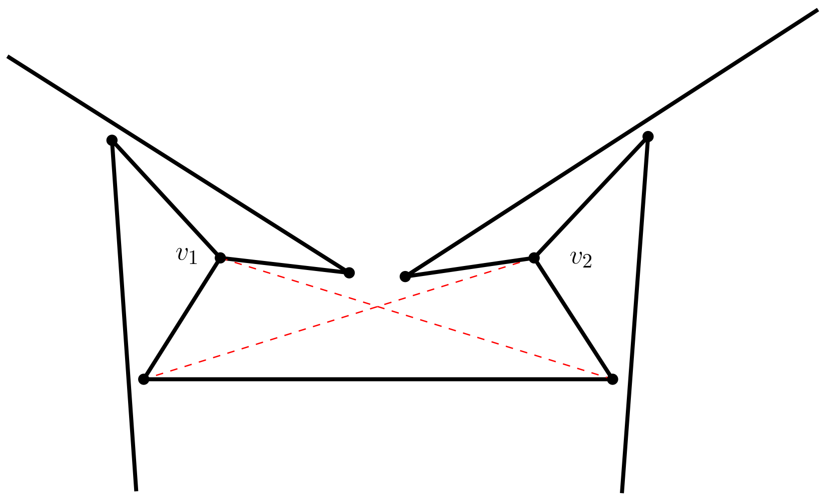

Figure 21: Either v_1 or v_2 has to be isolated in any compatible triangulation.

Property ($\widehat{\mathcal{W}}$): For any two adjacent edges uv and uw of T, if the geometric triangle uvw contains some vertices $v'_1, \ldots, v'_k$, then not all of them are leaves of T adjacent to u.

If property ($\widehat{\mathcal{W}}$) is not satisfied, and hence triangle uvw contains only the leaves $v'_1, v'_2, \ldots, v'_k$ of T, with $k \geq 1$ and each uv'_i being an edge of T, consider the tree T' obtained by the removal of $v'_1, v'_2, \ldots, v'_k$ from T. Given a T'-compatible triangulation Δ' of T', we can insert in Δ' the edges uv'_i without producing any crossing. Therefore, we can complete Δ' to a T-compatible triangulation Δ of T with the same number of isolated points as Δ', because vertices v'_i cannot be isolated once added. Therefore, any upper bound for the minimum number of isolated points for trees satisfying ($\widehat{\mathcal{W}}$) is also valid for arbitrary trees.

As a second step, let us see that we can increase the number of vertices of the big component by doing the following process:

Suppose that Δ is a T-compatible triangulation having one big component C, and u is a vertex not in C. Let $v_1, \ldots, v_k$ be the clockwise neighbors of u in T. Notice that $k \geq 3$ and that u is inside Q. We can also assume that T is not a star. Therefore, some of the triangles uv_iv_{i+1} has an adjacent triangle $v_iv_{i+1}w_i$. Then, if the quadrilateral $uv_iv_{i+1}w_i$ is convex, we can replace the red edge v_iv_{i+1} by the red edge uw_i. Notice that the vertices $v_1, \ldots, v_k, v_1$ form a red cycle and, hence, when the edge v_iv_{i+1} is removed, the red components do not change, and then, adding the edge uw_i, vertex u becomes part of the big component C.

We are now ready for proving the first claim in Theorem 1:

Let T be a geometric tree. Then there is a T-compatible triangulation Δ' such that $\Delta' - T$ contains at most $\frac{n-3}{4}$ isolated vertices, in addition to one large component.

Proof. Let Q be the boundary of the convex hull of T. If T is Q-convex, then, since there is a triangulation with one or two components, the maximum number of isolated vertices is 1.

If T is an ordinary tree, let Δ be a triangulation containing all the vertices of Q in one component, and let Δ' be the triangulation obtained from Q following the process described in the paragraph above, preceding the claim.

Leu u be an isolated vertex in Δ', and let $v_1, \ldots, v_k$ be its neighbors in T. At least one triangle uv_iv_{i+1} has an adjacent triangle $v_iv_{i+1}w_i$, and the quadrilateral $uv_iv_{i+1}w_i$ is not convex. Then, either the ray $\overrightarrow{uv_{i+1}}$ hits first the segment v_iw_i, or the ray $\overrightarrow{uv_i}$ hits first the segment $v_{i+1}w_i$.

Assume that we are the first case: we will say that w_i is *clockwise in relation to* uv_iv_{i+1}. Observe that

i) Vertex v_{i+1} does not belong to Q, because the clockwise angle between uv_{i+1} and $v_{i+1}w_i$ is greater than π. Therefore, the triangle $uv_{i+1}v_{i+2}$ has an adjacent triangle $v_{i+1}v_{i+2}w_{i+1}$.

ii) Vertex w_i cannot coincide with vertex v_{i+2}, because then property $(\widehat{\mathcal{W}})$ is not satisfied. Therefore, since the quadrilateral $uv_{i+1}v_{i+2}w_{i+1}$ cannot be convex, the ray $\overrightarrow{uv_{i+2}}$ hits first the segment $v_{i+1}w_{i+1}$, and, therefore, again w_{i+1} is clockwise in relation to $uv_{i+1}v_{i+2}$.

Repeating the same argument, we will find that, for $i = 1, \ldots, k$, the triangle uv_iv_{i+1} has an adjacent triangle $v_iv_{i+1}w_i$, with w_i being clockwise in relation to uv_iv_{i+1}, i.e., the ray $\overrightarrow{uv_{i+1}}$ hits first the segment v_iw_i, and therefore, in the tree T, clockwise around v_{i+1}, after the edge uv_{i+1} comes an edge forming an angle greater than π (or v_{i+1} is a leaf of T). Let us call such an isolated vertex a *clockwise isolated vertex* of Δ'. See Figure 22. Notice that edges uv_i have to be black but edges v_iw_i can be red or black.

Similarly, in the second case, $\overrightarrow{uv_i}$ hits first the edge $v_{i+1}w_i$, all the triangles have an adjacent triangle, and in the tree T, counterclockwise around v_i, after the edge uv_i comes an edge forming an angle greater than π. Let us call such an isolated vertex a *counterclockwise isolated vertex* of Δ'. See Figure 22.

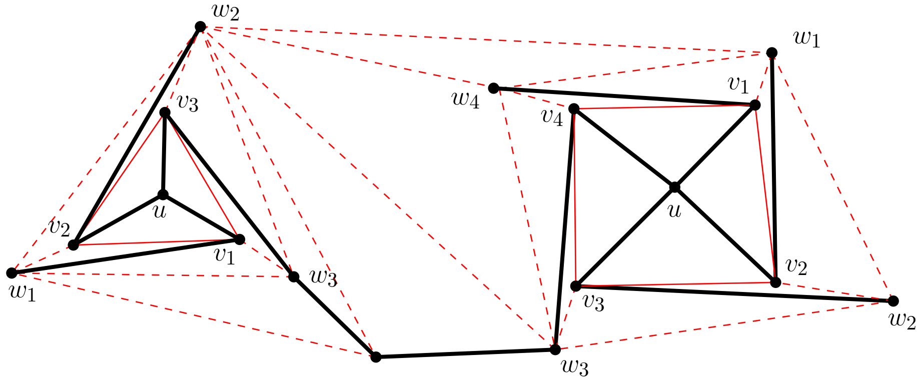

Figure 22: Clockwise and counterclockwise isolated vertices in a triangulation Δ'.

Let us now suppose that in Δ' the isolated vertices u and u' have a common neighbor v_i in T, this is, uv_i and $u'v_i$ are edges of T. Only two consecutive edges of T around v_i can form an angle greater than π. Therefore, one among u and u' has to be clockwise isolated, say u, and the other one, u', counterclockwise isolated. See Figure 23 left. Let us denote the clockwise neighbors of u by $v_i, v_{i+1}, \ldots$ and the clockwise neighbors of u' by $v'_i = v_i, v'_{i+1}, \ldots$. Since ray $\overrightarrow{uv_i}$ first has to cut the edge $v_{i-1}w_{i-1}$, then, ray $\overrightarrow{u'v_i}$ has to enter first into either the triangle uv_iv_{i-1}, or into $v_iv_{i-1}w_{i-1}$, hitting first edge $v_{i-1}w_{i-1}$ or edge uv_{i-1} instead of edge $v'_{i+1}w'_i$. Therefore, this situation is impossible except in the case that $w_{i-1} = v'_{i+1}$ and $w'_i = v_{i-1}$; this last situation is shown in Figure 23, right. In this case, edge w'_iw_{i-1} belongs to two triangles of Δ', namely $v_iw'_iw_{i-1}$ and $w'_iw_{i-1}w$. Consider the polygon $P = w'_iww_{i-1}u'v_iu$. In this polygon, the diagonal w'_iv_i can be deleted without changing the red components because the edges $v_j v_{j+1}$ are red and form a cycle. By the same reason we can delete the edge $v_i w_{i-1}$ and the edge $w'_i w_{i-1}$. Finally, since the clockwise angle from $w'_i u$ to $w'_i w_{i-2}$ is bigger than π, vertex w'_i is convex in the polygon P, and by the same reason w_{i-1} is convex in that polygon. Therefore, we can retriangulate the interior of P using either the diagonal uw or the diagonal $u'w$ (or both), obtaining in this way a new triangulation where either vertex u or vertex u', or maybe both of them, have been connected to the big component C. Repeating this same construction for all the isolated vertices with a common neighbor, we eventually obtain a T-compatible triangulation Δ'' in which any two isolated vertices cannot have a common neighbor.

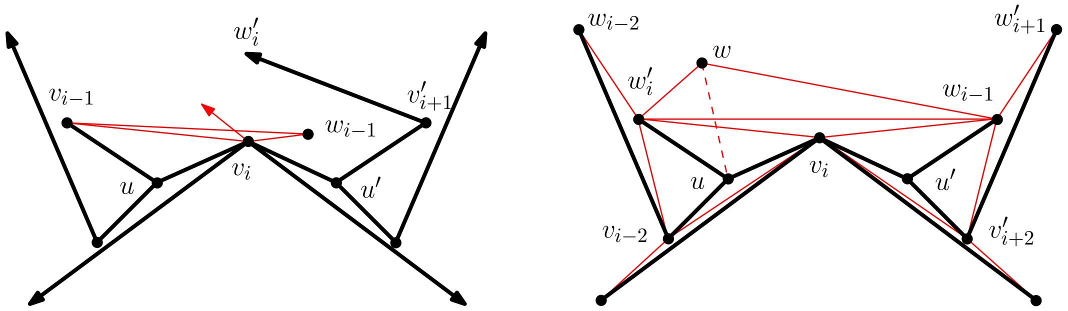

Figure 23: Isolated vertices with common neighbors.

Therefore, if Q contains h vertices, and there are s isolated vertices in Δ'', as each isolated convex vertex u has at least three neighbors in T, we get

$$4s \leq n - h$$

This proves the first claim in Theorem 1. □

We can proceed now to prove the second claim in Theorem 1:

There are geometric trees T, such that any T-compatible triangulation Δ gives at least $\frac{n-2}{5}$ isolated vertices in $\Delta - T$.

Proof. Let us see that there are trees T containing at least $\frac{n-2}{5}$ totally isolated vertices, as the example shown in Figure 20. Basically the figure is formed from the simplest tree having a totally isolated vertex, a tree T_u formed by the six edges uv_i and $v_i w_i, i = 1, 2, 3$ placed in such a way that u cannot see any other vertex. Now, observe that an edge $v_i w_i$ can be an edge of type $v'w'$ corresponding to another similar tree $T_{u'}$. Therefore, we can glue several trees T_u along edges $v_i w_i$, as in Figure 20. Initially, for tree T_u we have seven vertices, one of them isolated. Then we can add five more vertices, one of them being isolated. We repeat this operation, say, $s-1$ times, obtaining a tree with s totally isolated vertices, and having a total of $n = 5s + 2$ vertices. □

As a complementary result, let us conclude this section proving that a tree T cannot have more than $\frac{n-2}{5}$ totally isolated vertices in the red graph. This is the reason supporting Conjecture 1, which we have stated in Section 1, as we believe that the tight value of d_n is $n/5$, up to an additive constant.

To prove the result, suppose that T is not a star, that u is a totally isolated point of T, and that $v_1, v_2, \ldots, v_{k'}$ are the clockwise neighbors of T. We can assume that vertex v_1 is a neighbor of u with degree at least two. Consider the ray with origin in u in the direction uv_1, and suppose that rotating clockwise this ray the first edge it hits is v_1w_1 of T, the *active edge*. If we continue rotating the ray, we have to find some vertex v_2 before the ray reaches w_1. If we keep rotating the ray, either it hits first some edge v_2w_2 or, if v_2 is a leaf, it continues hitting first the edge v_1w_1. In any case, before the ray reaches w_1, some other active edge $v_{i_2}w_{i_2}$ (the first edge found by the ray) must appear. Similarly, edges $v_{i_3}w_{i_3}, v_{i_4}w_{i_4}, \ldots, v_1w_1$, have to appear clockwise as first edges found by the rotating ray. Notice that as the clockwise angle between uv_1 and uw_1 is less than π, at least three different active edges v_iw_i appear in this process. Now, consider the subtree T_u formed by the $2k$ edges uv_i and v_iw_i found in the preceding process ($k \geq 3$). If $T_{u'}$ is the subtree corresponding to another isolated point u', it suffices to observe that if edge $e = vv'$ is common to the trees T_u and $T_{u'}$, then v has to be a neighbor of u and v' must be a neighbor of u'. Therefore, an edge can be shared by at most two subtrees T_u. Then, consider the combinatorial graph having as nodes the subtrees T_u, two of them being adjacent if they share at least one edge. If T has s totally isolated vertices, the s subtrees T_u share a maximum of $s-1$ edges, because T contains no cycles. Therefore, T has at least $6s - (s-1) = 5s + 1$ edges, which proves the claim.

2.8 Proof of Theorems 3 and 4

Let us finish this article with the proof of two results that can be derived from Propositions 1 and 2, and that we find interesting in their own.

We recall the statement of Theorem 3, which deals with the special case in which T is a path:

Let T be a simple path, spanning the point set S. Then, we can find a T-compatible triangulation Δ, such that the number of red components of $\Delta - T$ is at most 2, and hence $d(T) \leq 1$.

Proof. Let Q be the boundary of the convex hull $CH(S)$ of S. Clearly, T lies in the polygon Q. If T is Q-convex, then the result is implied by Proposition 1. If T is not Q-convex, since it does not have any vertices with degree three or more, we can obtain triangulations without isolated vertices and, therefore, we obtain only one red component. In particular, if the endpoints of the path T are not vertices in Q, we can always obtain a triangulation with only one red component. □

Notice that, since a vertex not placed on Q cannot be convex in all of its pockets, Corollary 1 implies that, for paths, $d(T) = 1$ if and only if T is Q-convex.

Finally, as a slight departure from our work on trees, we recall that Theorem 4 considers simple polygons:

Let Γ be a simple polygon with vertex set S. Then, we can always find a Γ-compatible triangulation Δ (which in general will have edges inside and outside Γ), such that the number of red connected components of $\Delta - \Gamma$ is at most 3. Moreover, if Γ is non-convex, this minimum achievable number of components is 1 or 2.

Proof. Let Q be, as in the preceding proof, the boundary of the convex hull $CH(S)$ of the point set S spanned by Γ. If Γ does not contain any edge e placed on Q, then Figure 24 shows an example with two red components and another example with only one component. In fact, in this case, every Γ-compatible triangulation has at most 2 components. To prove this result formally, we only need a slight modification of Lemma 8. Observe that now

Γ is not a tree, and it contains a unique cycle (Γ itself). Therefore, if v_i and v_j are in different red components, doing the same construction than in Lemma 8, it is possible now to find a cyclic sequence $F_1, F_2, \ldots, F_l, F_1$ of triangles with different signs, the spine of that sequence being the cycle Γ. In this case we have exactly two red components (as in Figure 24, left). Since a proper separating band cannot appear, because it should begin and end on edges of Γ placed on Q, this is the only case in which there are two components, in all other cases there can be only one red component.

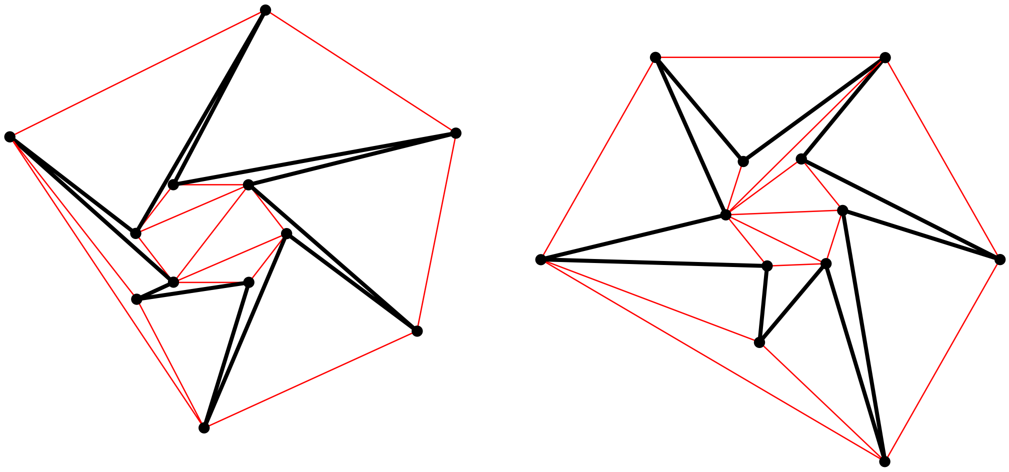

Figure 24: Two simple polygons without consecutive convex vertices.

If Γ contains an edge e on Q, then $\Gamma' = \Gamma - e$ is a path, and considering edge e for a moment as a red edge, we can obtain a compatible triangulation having one or two red components. Now, if we give back to edge e its black color, we see that the number of previous red components can increase only by one, hence giving at most 3 components. Finally, notice that if Γ is non-convex, then there is an edge e of Γ on Q such that its next consecutive edge e' of Q is not an edge of Γ. Taking this edge e as red, either the interior of Γ is transformed into an ordinary pocket with red lid e or the pocket with lid e' is ordinary. Therefore, the path $\Gamma' = \Gamma - e$ gives rise to one red component, and hence Γ generates at most two components. $\square$

References

[1] M. Abellanas, A. García, F. Hurtado, J. Tejel, J. Urrutia, Augmenting the connectivity of geometric graphs, Comput. Geom. 40 (3) (2008) 220–230.

[2] O. Aichholzer, F. Aurenhammer, F. Hurtado, Sequences of spanning trees and a fixed tree theorem, Comput. Geom. 21 (1–2) (2002) 3–20.

[3] O. Aichholzer, S. Bereg, A. Dumitrescu, A. García, C. Huemer, F. Hurtado, M. Kano, A. Márquez, D. Rappaport, S. Smorodinsky, D. Souvaine, J. Urrutia, D. Wood, Compatible geometric matchings, Comput. Geom. 42(6–7) (2009) 617–626.

[4] O. Aichholzer, T. Hackl, B. Vogtenhuber, Compatible pointed pseudo-triangulations, in: Proc. 22^{nd} Annual Can. Conf. Comp. Geom., CCCG 2010, pp. 91–94.

[5] M. Al-Jubeh, M. Ishaque, K. Rédei, D. L. Souvaine, C. D. Tóth, Augmenting the edge connectivity of planar straight line graphs to three, Algorithmica 61 (4) (2011) 971–999.

[6] T. Asano, S. K. Ghosh, T. C. Shermer, Visibility in the plane, in: J.R. Sack, J. Urrutia (Eds.), Handbook on Computational Geometry, Elsevier, Amsterdam, 2000, pp. 829–876.

[7] P. Bose, M. Houle, G. T. Toussaint, Every Set of Disjoint Line Segments Admits a Binary Tree, Discrete Comput. Geom. 26 (3) (2001) 387–410.

[8] P. Bose, G. T. Toussaint, Growing a tree from its branches, J. Algorithms 19 (1995) 86–103.

[9] P. Brass, E. Cenek, C. Duncan, A. Efrat, C. Erten, D. Ismailescu, S. Kobourov, A. Lubiw, J. Mitchell. On simultaneous planar graph embeddings, Comput. Geom. 36 (2) (2007) 117–130.

[10] F. Frati, M. Geyer, M. Kaufmann, Planar packings of trees and spider trees. Inform. Process. Lett. 109 (6) (2009) 301–307.

[11] E. Di Giacomo, W. Didimo, G. Liotta, H. Meijer, S. Wismath, Point-set embeddings of trees with given partial drawings, Comput. Geom.42 (6–7) (2009) 664–676.

[12] A. García, F. Hurtado, M. Noy, J. Tejel, Augmenting the connectivity of outerplanar graphs, Algorithmica 56 (2) (2010) 160–179.

[13] M. Hoffmann, B. Speckmann, C. D. Tóth, Pointed binary encompassing trees: Simple and optimal, Comput. Geom. 43 (1) (2010) 35–41.

[14] M. Hoffmann, C. D. Tóth, Segment endpoint visibility graphs are Hamiltonian, Comput. Geom. 26 (1) (2003) 47–68.

[15] M. E. Houle, F. Hurtado, M. Noy, E. Rivera-Campo, Graphs of triangulations and perfect matchings, Graph Combinator. 21 (3) (2005) 325–331.

[16] F. Hurtado, M. Kano, D. Rappaport, C. D. Tóth, Encompassing Colored Crossing-Free Geometric Graphs, Comput. Geom. 39 (1) (2008) 14–23.

[17] F. Hurtado, C. D. Tóth, Plane geometric graph augmentation: a generic perspective, in: J. Pach (Ed.), Thirty Essays on Geometric Graph Theory, Springer, to appear (2012).

[18] Y. Ikebe, M. Perles, A. Tamura, S. Tokunaga, The rooted tree embedding problem into points in the plane, Discrete Comput. Geom. 11 (1) (1994) 51–63.

[19] M. Ishaque, D. L. Souvaine, C. D. Tóth, Disjoint compatible geometric matchings, Discrete Comput. Geom., to appear (2012).

[20] A. Kaneko, M. Kano, Straight-Line embeddings of two rooted trees in the plane, Discrete Comput. Geom. 21 (4) (1999) 603–613.

[21] A. Mirzaian, Hamiltonian triangulations and circumscribing polygons of disjoint line segments, Comput. Geom. 2 (1) (1992) 15–30.

[22] J. O'Rourke, Visibility, in: J. E. Goodman, J. O'Rourke (ds.), Handbook of Discrete and Computational Geometry, 2nd edition, CRC Press LLC, Boca Raton, FL, 2004, pp. 643–665.

[23] J. Pach, E.Rivera-Campo, On circumscribing polygons for line segments, Comput. Geom. 10 (2) (1998) 121–124.

[24] D. Rappaport, Computing simple circuits from a set of line segments is NP-complete, SIAM J. Comput 18 (1989) 1128–1139.

[25] D. Rappaport, H. Imai, G. T. Toussaint, Computing simple circuits from a set of line segments, Discrete Comput. Geom. 5 (3) (1990) 289–304.

[26] I. Rutter, A. Wolff, Augmenting the Connectivity of Planar and Geometric Graphs, Electronic Notes in Discrete Mathematics 31 (2008) 53–56.

[27] C. D. Tóth, Connectivity augmentation in planar straight line graphs, Eur. J. Comb. 33 (3) (2012) 408–425.

[28] M. Urabe, M. Watanabe, On a counterexample to a conjecture of Mirzaian, Comput. Geom. 2 (1) (1992) 51–53.

[29] J. Urrutia, Art gallery and illumination problems, in: J. R. Sack, J. Urrutia (Eds.), Hand- book on Computational Geometry, Elsevier, Amsterdam, 2000, pp. 973–1127.

[30] B. Vogtenhuber, Combinatorial Aspects of [Colored] Point Sets in the Plane, PhD thesis, Institute for Software Technology, Graz University of Technology, Graz, Austria, 2011.