

The edge rotation graph

Javier Cano · José-Miguel Díaz-Báñez ·
Clemens Huemer · Jorge Urrutia

Received: date / Accepted: date

Abstract Let P be a set of n points on the plane in general position, $n \geq 3$. The edge rotation graph $\mathcal{ERG}(P, k)$ of P is the graph whose vertices are the plane geometric graphs on P with exactly k edges, two of which are adjacent if one can be obtained from the other by an edge rotation. In this paper we study some structural properties of $\mathcal{ERG}(P, k)$, such as its connectivity and diameter. We show that if the vertices of $\mathcal{ERG}(P, k)$ are not triangulations of P , then it is connected and has diameter $O(n^2)$. We also show that the chromatic number of $\mathcal{ERG}(P, k)$ is $O(n)$, and show how to compute an implicit coloring of its vertices. We also study edge rotations in edge-labelled geometric graphs.

Keywords: Geometric graph; Rotation; Diameter.

J.M. Díaz-Báñez partially supported by Project MEC MTM2009-08652.

C. Huemer partially supported by Projects MEC MTM2009-07242 and Gen. Cat. DGR 2009SGR1040.

J. Cano and J. Urrutia partially supported by grants CONACyT CB-2007/80268, and by Project MEC MTM2009-08652.

J. Cano
Posgrado en Ciencia e Ingeniería de la Computación, Universidad Nacional Autónoma de México, Mexico
`j_cano@uxmcc2.iimas.unam.mx`

J.M. Díaz-Báñez
Departamento de Matemática Aplicada II, Universidad de Sevilla, Spain,
`dbanez@us.es`

C. Huemer
Departament de Matemàtica Aplicada IV, Universitat Politècnica de Catalunya, Spain
`clemens.huemer@upc.edu`

J. Urrutia
Instituto de Matemáticas, Universidad Nacional Autónoma de México
`urrutia@matem.unam.mx`

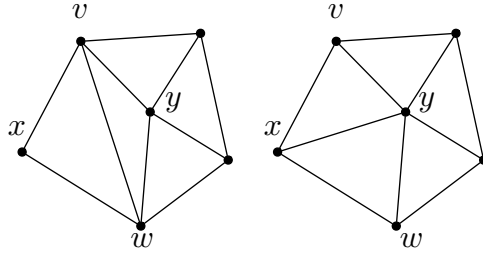


Fig. 1 Flipping edge vw .

1 Introduction

In this paper, P will always denote a set of n points in the plane in general position, $\text{Conv}(P)$ will denote the convex hull of P , and c and i the number of elements of P on the boundary and in the interior of $\text{Conv}(P)$ respectively. A geometric graph G on P is a graph whose vertices are the elements of P , and its edges are straight line segments joining pairs of elements of P . The edge of G joining $x, y \in P$ will be denoted by xy . We say that G is plane if no edges of G intersect except at a common vertex. All geometric graphs considered here will be plane, and thus the term geometric graph will refer to plane geometric graphs.

A geometric graph T on P is called a triangulation if the edges of T partition $\text{Conv}(P)$ into a set of interior disjoint triangles such that they do not contain an element of P in their interior. These triangles are called the triangles of T . It is well known that any triangulation of P contains exactly $2c + 3i - 3$ edges.

An edge e of a triangulation T of P is called flippable if it belongs to the boundary of two triangles t_1 and t_2 of T such that $t_1 \cup t_2$ is a convex quadrilateral Q . By “flipping e ,” we mean the operation of removing e from T and inserting the second diagonal of Q ; see Figure 1. It is known that any triangulation of P can be transformed into any other triangulation of P by executing at most $O(n^2)$ edge flips; see [7, 8].

In this paper we study an operation similar to edge flipping called *edge rotations*. Let G be a plane geometric graph on P that is not a triangulation, and xy an edge of G . An *edge rotation* of xy around point x replaces xy by an edge $xw \notin G$ if:

1. The geometric graph $G - xy + xw$ is plane,
2. the open triangle $\triangle xyw$ with vertices x, y, w contains no element of P , and
3. $\triangle xyw$ does not intersect any edge $e \in G$; see Figure 2.

Edge rotations on graphs (not necessarily geometric graphs) have been studied before in the literature; see e.g. [4]. Edge rotations on binary trees have been studied in [9], among other things to their applications in data structures; see also [10, 11]. A similar transformation on plane geometric graphs, called edge slides, was studied by Aichholzer and Reinhardt [1]. They proved that it

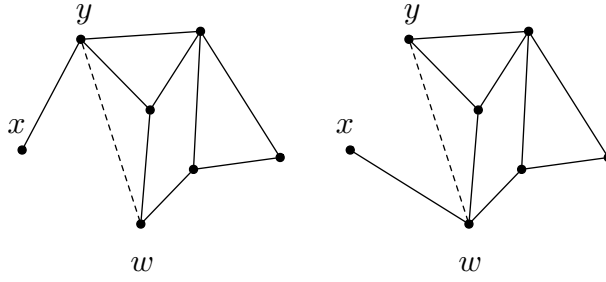


Fig. 2 Rotating edge xy to xw .

is possible to transform any two plane-spanning trees of a point set into each other by using $O(n^2)$ local and *constant-size* edge slides. Edge rotations on labelled trees are studied in [5]. For more details on edge rotations and edge flips, we refer the interested reader to a recent survey on this topic [2].

We are concerned with the question of when and how fast two geometric plane graphs can be transformed into each other by means of edge rotations. Given a set of points P , let $PG(P, k)$ be the set of all plane geometric graphs on P with k edges, $k < 2c + 3i - 3$. The *edge rotation graph* $\mathcal{ERG}(P, k)$ has vertex set $PG(P, k)$. Two vertices of $\mathcal{ERG}(P, k)$ are adjacent if they differ by an edge rotation.

We prove that $\mathcal{ERG}(P, k)$ is connected, and give tight asymptotic bounds on its diameter, and its chromatic number. Our paper is organized as follows: In Sections 2 and 3 we prove that the edge rotation graph is connected and give tight asymptotic bounds on its diameter. In Section 4, we study the chromatic number of $\mathcal{ERG}(P, k)$, and finally in Section 5 we extend our results to plane geometric graphs with labelled edges.

2 Connectivity

In what follows, we will assume that the geometric graphs we study are not triangulations and have at least three vertices. A device that will prove useful to us is the following: Given a geometric graph G on P which is not a triangulation, we will add to it some Steiner edges, which we will call *absent edges*, such that G together with these extra edges is a triangulation of P .

Thus instead of talking about plane geometric graphs, we will refer to *triangulations of P with at least one absent edge*, or to “triangulations with absent edges” for short. The *dual graph* of a triangulation is the graph whose vertices are the triangular faces of T contained in $\text{Conv}(P)$, two of which are adjacent if they share an edge on their boundaries.

Let a and b be an absent and a non-absent edge of a triangulation T of P . Then $T(a, b)$ will denote the graph obtained from T by removing b from T (and thus making it an absent edge), and replacing a by a new non-absent edge, i.e. we can think of this as *switching* a with b .

The following two lemmas will be the main tools for proving the connectivity of $\mathcal{ER}\mathcal{G}(P, k)$.

Lemma 1 *Let T be a triangulation of P with at least one absent edge, and a and b an absent and a non-absent edge of T respectively. Then there is a sequence of edge rotations that transforms T to $T(a, b)$.*

Proof Let T' be the dual graph of T , and Π a shortest path in T' among all paths in T' connecting a triangle of T containing a on its boundary to another triangle whose boundary contains b .

The proof proceeds by induction over the length of Π . If Π is a path whose length is zero, a and b are edges of a triangle of T . By definition, we can rotate b to a around the vertex incident to both of them, thus obtaining $T(a, b)$.

Suppose then that the length of Π is $m > 0$. Let t_1, \dots, t_{m+1} be the triangles in Π , and assume that a belongs to the boundary of t_1 . Let e_i be the edge common to t_i and to t_{i+1} , $i = 1, \dots, m$. Two cases arise. Suppose first that all the edges e_i are non-absent edges of T . Then since a and e_1 belong to a triangle of T , we can rotate e_1 to a in T to obtain $T_1 = T(a, e_1)$. Thus a is now a non-absent edge of T_1 , and e_1 is absent. By induction on the length of $\Pi \setminus \{t_1\}$, we can now exchange e_1 (now an absent edge) with b to obtain $T_1(e_1, b)$. But $T_1(e_1, b) = T(a, b)$, and our result follows.

Suppose next that at least one edge e_i of Π is an absent edge of T . Let j be the largest index such that e_j is absent. By the previous paragraph, we can now transform T to $T_1 = T(e_j, b)$. By induction on $\Pi \setminus \{t_{j+1}, \dots, t_{m+1}\}$ we can now exchange a with e_j (which is now non-absent) to obtain $T_1(a, e_j)$. Again, it is easy to see that $T_1(a, e_j) = T(a, b)$. \square

Recall that an edge e of a triangulation T is flippable if the union of the two triangles of T containing e on their boundaries is a convex quadrilateral; denote it as Q_e . Recall also that flipping d means replacing it in T by the second diagonal of Q_e ; see Figure 1.

Lemma 2 *Given a triangulation T with at least one absent edge, any edge flip can be attained by a sequence of edge rotations.*

Proof Let xy be a flippable edge of T . If xy is an absent edge of T , we simply flip it to the second diagonal of Q_{xy} which is also an absent edge in the resulting triangulation. Suppose then that xy is a flippable non-absent edge of T , and that the other two vertices of Q_{xy} are w and z ; see Figure 3.

Let e be an absent edge of T . By Lemma 1, we can exchange e and xw by performing some edge rotations (at this point xw is an absent edge). Now rotate xy to xw , then xy —which is now absent—can be flipped to wz , and then rotate wx to wz . At this point we have flipped xy , and xw is again an absent edge of T . Applying Lemma 1 again, move xw back to its original position. Our result follows. \square

Let T be a triangulation of P with r absent edges, $r \geq 1$. The underlying triangulation \mathcal{T} of T is the triangulation of P obtained by replacing each

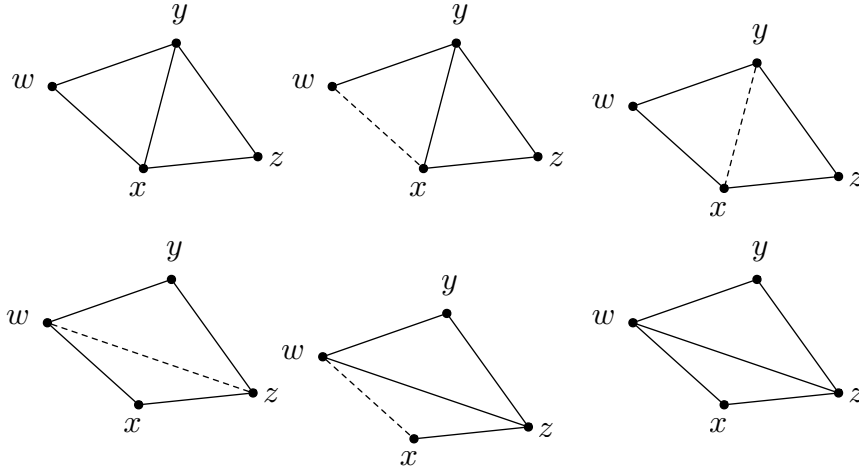


Fig. 3 Flipping edge xy to wz .

absent edge of T by a non-absent edge. Let T and T' be two triangulations of P each with r absent edges and such that they have the same underlying triangulation \mathcal{T} , but their sets of absent edges are different. We prove:

Lemma 3 T can be transformed to T' using at most $O(n^2)$ edge rotations.

Proof Suppose that the absent edges of T and T' are not the same. Then there are two different edges xy and uv such that:

- a) xy is absent in T , and uv is non-absent in T , and
- b) xy is non-absent in T' , and uv is absent in T' .

By Lemma 1, we can transform T to $T(xy, uv)$ with a linear number of flips. Observe that $T_1 = T(xy, uv)$ and T' have now two more absent edges in common. Our result follows by induction on the number of absent edges not common to T and T' . Clearly the above procedure is performed at most a linear number of times, until T is transformed into T' . Each iteration costs linear time. Our result follows. \square

Next we prove:

Theorem 1 $\mathcal{ERG}(P, k)$ is connected.

Proof Let G_1 and G_2 be two plane geometric graphs on P with k edges. Complete G_1 and G_2 to two triangulations T_1 and T_2 by adding some absent edges to each of them. Let \mathcal{T}_1 , and \mathcal{T}_2 be the underlying triangulations of T_1 and T_2 . Since \mathcal{T}_1 , and \mathcal{T}_2 are triangulations of P , we can transform \mathcal{T}_1 into \mathcal{T}_2 by using at most $O(n^2)$ edge flips [6, 7]. By Lemma 2, each edge flip can be achieved with a sequence of edge rotations, and thus \mathcal{T}_1 can be transformed to \mathcal{T}_2 by a sequence of edge rotations. We must now be careful, as the absent edges of \mathcal{T}_1 were not necessarily mapped to the absent edges of \mathcal{T}_2 . We apply Lemma 3 to fix this, and transform G_1 into G_2 . \square

We observe that the proof of the previous result gives us an $O(n^3)$ upper bound on the diameter of $\mathcal{ER}\mathcal{G}(P, k)$. We must perform $O(n^2)$ flips, each potentially using a linear number of edge rotations. In the next section, we will improve on this by showing that the diameter of $\mathcal{ER}\mathcal{G}(P, k)$ is $O(n^2)$.

3 The diameter of $\mathcal{ER}\mathcal{G}(P, k)$

In this section we show tight asymptotic bounds on the diameter of the edge rotation graph. For the sake of completeness, we recall first some well known facts about Delaunay triangulations.

Let T be a triangulation of P . A triangle t_i of T is called a *Delaunay triangle* if the circumcircle of t_i (the circle passing through the vertices of t_i) contains no element of P in its interior. A triangulation T is a Delaunay triangulation of P if all of its triangles are Delaunay; see Figure 4. When P contains no four co-circular points, the Delaunay triangulation of P is unique, and thus well defined; we denote it as $DT(P)$. We will suppose that such is the case. (This condition can easily be dropped, leaving our results unchanged.) Let e be a flippable edge of a triangulation T , and t_1 and t_2 the triangles of T containing e on their boundaries. Since e is flippable, $Q_e = t_1 \cup t_2$ is a convex quadrilateral. We say that flipping e to the second diagonal of Q_e is a *Delaunay flip* if there is a circle passing through the endpoints of e that contains Q_e ; see Figure 5. It is well known that when no Delaunay flips can be performed on a triangulation T , it is the Delaunay triangulation of P [7]. An edge e of a triangulation T is called a *Delaunay edge* if there is a circle passing through its endpoints and containing no elements of P in its interior.

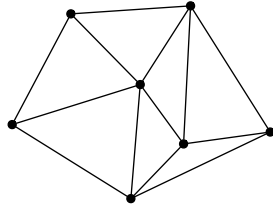


Fig. 4 A Delaunay triangulation.

A classical result on edge flipping and the Delaunay triangulation asserts that any triangulation of P can be transformed to $DT(P)$ by performing at most $O(n^2)$ Delaunay edge flips. To achieve this bound, we never flip an edge which is already a Delaunay edge of any of the triangulations obtained while transforming T to $DT(P)$ [7]. To prove that $O(n^2)$ Delaunay edge flips are necessary, a weight is associated to any triangulation T of P as follows: Each triangle of T receives a weight equal to the number of points in the interior of its circumcircle. The weight of T is the sum of the weights of its triangles.

Observe that the weight of a triangulation is in the worst case quadratic, and that the weight of the Delaunay triangulation is $DT(p) = 0$. In [7], it is shown that each time a Delaunay flip is performed on T , its weight decreases by at least two, and thus the number of Delaunay flips required to reach $DT(P)$ is bounded above by n^2 . We are now ready to prove:

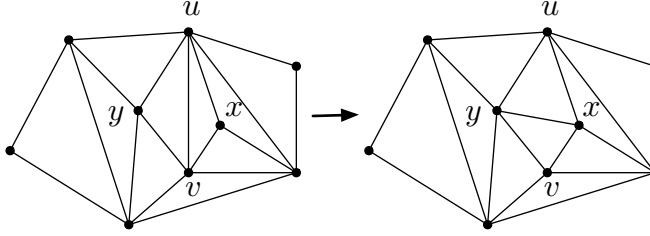


Fig. 5 Flipping uv to xy is a Delaunay flip.

Theorem 2 *The diameter of $\mathcal{ERG}(P, k)$ is at most $O(n^2)$.*

Proof We will show that any triangulation T of P with absent edges can be transformed to $DT(P)$ with $O(n^2)$ edge rotations. We will show how to perform a linear number of iterations, each using a linear number of edge rotations, plus a sequence of edge flips, each achieved with a constant number of edge rotations. Each iteration stops when an edge is flipped into a Delaunay edge.

Each iteration proceeds as follows:

- 1) Find a flippable edge e of T which is not a Delaunay edge of T (it is easy to see that if we flip e , we perform a Delaunay edge flip).
- 2) Consider the quadrilateral Q_e . If no edge of Q_e is an absent edge, exchange an edge of Q_e , say e' , with an absent edge of T .
- 3) Perform a sequence of edge rotations which exchange e for the second diagonal f of Q_e . As shown in Figure 6, there are two ways of achieving this. We will choose one of them according to which edge will be flipped next.
- 4) If f is a Delaunay edge of P , stop, and start the next iteration.
- 5) If f is not a Delaunay edge, then it is easy to see that at least one of the edges of Q_e , call it e'' , is flippable and is not a Delaunay edge. We now observe that we can rotate e to f in 3) in such a way that the absent edge of Q_e is incident to e'' , or e'' is absent; see Figure 6. Choose e'' , and go to 1).

Observe now that at the beginning of each iteration, the first time we execute Step 2) we may need a linear number of edge flips to exchange an edge of Q_e with an absent edge. However, we can rotate e to f in such a way that the next edge to be rotated, namely e'' , is such that one of the edges of

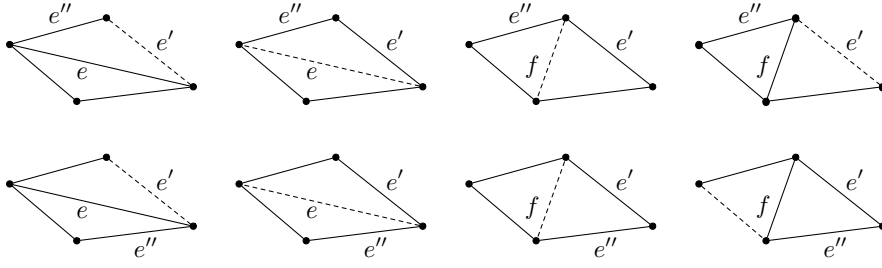


Fig. 6 Two ways to flip e to f . We can choose one of them such that when finished, the dotted absent edge is the next edge to be flipped, or is incident to edge e'' .

$Q_{e''}$ is absent. Thus all edge flips take at most two edge rotations, except for the first edge flip of each iteration.

The number of iterations is linear, since the number of Delaunay edges is linear, and Step 3) of our procedure is executed at most a quadratic number of times. This follows from the fact that each time Step 3) is performed, the weight of the current triangulation decreases. Our result follows, \square

We now prove that our bound is tight.

Theorem 3 *There exists a point set P such that the diameter of $\mathcal{ERG}(P, k)$ is $\Omega(n^2)$ for k equal the number of edges of a triangulation on P minus a constant.*

Proof Consider the point set P with $2m = n$ points, and the triangulations of P shown in Figure 7 (for now, all edges, both solid and dotted, are in both triangulations). It was proved in [6] that to transform the triangulation on the left of Figure 7 to the triangulation on the right of the same figure takes $2m^2$ edge flips.

Remove a constant number of edges from each of these triangulations, some from the the left and some from the right (the dashed lines in Figure 7). It is now easy to verify that the same arguments used in [6] to show that $2m^2$ edge flips are required to transform one triangulation into the other can be used to show that we need a quadratic number of edge rotations to transform these geometric graphs into each other. We omit the details. \square

4 The Chromatic Number of $\mathcal{ERG}(P, k)$

In this section we prove that the chromatic number of $\mathcal{ERG}(P, k)$ is at most n . Let G be a graph. An edge coloring of G is an assignment of colors to its edges in such a way that any two edges with a common vertex receive different colors. The smallest integer r such that the edges of G can be colored with r colors is called the chromatic index of G , and is usually denoted as $\chi'(G)$.

We recall a well known result in graph theory:

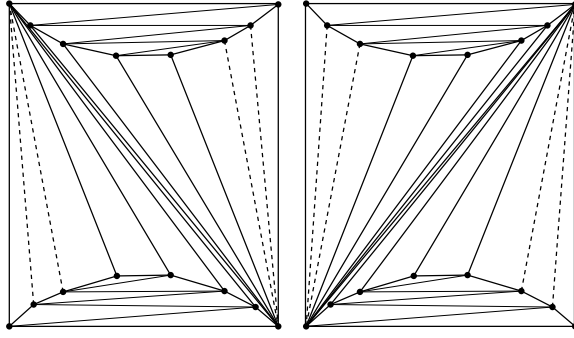


Fig. 7 An example that requires $\Omega(n^2)$ edge rotations.

Theorem 4 *The chromatic index of the complete graph K_n on n vertices is $n - 1$ for n even, or n if n is odd, $n \geq 2$.*

Color the edges of the complete geometric graph \mathcal{K}_n on P with the integers $\{0, \dots, \chi'(\mathcal{K}_n) - 1\}$, that is with the integers $\{0, \dots, n - 2\}$ for n even, and with $\{0, \dots, n - 1\}$ for n odd. We now use a similar idea to that used in [3] to obtain colorings of tree graphs.

Let G be a plane geometric graph on P with k edges. To each edge of G , assign a weight equal to the integer it received in the coloring of \mathcal{K}_n obtained above. Assign to G the color obtained by adding the weights of its edges, mod $\chi'(\mathcal{K}_n)$.

Suppose that G' is obtained from G by an edge rotation that takes an edge, say xy , to xw . Since both of these edges are adjacent to x , they receive different colors in the edge coloring of \mathcal{K}_n . Suppose that xy has color i , and xz color j , $0 \leq i, j \leq \chi'(\mathcal{K}_n) - 1$, $i \neq j$. Then if the color of G is r , the color of G' is $s = r - i + j$, mod $\chi'(\mathcal{K}_n)$. Clearly $r \neq s$. Thus we have just obtained a good coloring of $\mathcal{ERG}(P, k)$ with $\chi'(\mathcal{K}_n)$ colors.

We now show that for n even, the chromatic number of some $\mathcal{ERG}(P, k)$ is $n - 1$. Let $n = 2m$, and let P contain $2m - 1$ points p_1, \dots, p_{2m-1} equally spaced on a circle C together with the center p of C . Choose $k = n$, and consider the $n - 1$ geometric graphs G_i such that p_1, \dots, p_{2m-1} form a cycle, and p is adjacent to p_i , $i = 1, \dots, 2m - 1$; see Figure 8. Clearly G_i is adjacent to G_j in $\mathcal{ERG}(P, k)$, $i \neq j$. Then $\mathcal{ERG}(P, k)$, contains a clique (complete subgraph) of size $2m - 1$. It follows that the chromatic number of $\mathcal{ERG}(P, k)$ is exactly $2m - 1$.

Thus we have proved:

Theorem 5 *The chromatic number of $\mathcal{ERG}(P, k)$ is at most the chromatic index $\chi'(K_n)$. This bound is sometimes achieved for n even, and $k = n$. For n odd, the chromatic number of $\mathcal{ERG}(P, k)$ is sometimes at least $\chi'(K_n) - 1$.*

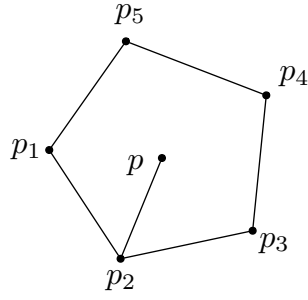


Fig. 8 G_2 .

5 Edge-labelled graphs

In this section we consider graphs in which each edge has a different label or *identity*. As we perform edge rotations, the edges change their endpoints but keep their identities. Suppose then that the edges of a plane geometric graph G with k edges are labelled e_1, \dots, e_k .

In Figure 9, we show how to exchange the labels of two edges of a triangle with an absent edge. This is accomplished using three edge rotations.

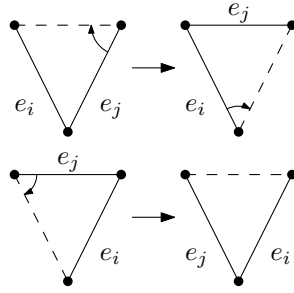


Fig. 9 Exchanging the labels of e_i and e using three edge rotations.

Let $PG_L(P, k)$ be the set of all plane geometric graphs on P with k edges such that in all of them, their edges are labelled with the labels e_1, \dots, e_k . Observe that each geometric graph on P such that its edges are unlabelled generates $k!$ vertices in $PG_L(P, k)$, and thus $|PG_L(P, k)| = k! |PG(P, k)|$. The *labelled-edge rotation graph* $\mathcal{ERG}_L(P, k)$ is the graph whose vertex set is $PG_L(P, k)$ and two vertices are joined by an edge if and only if they differ by an edge rotation.

A plane geometric graph W is called a k -*wheel* if it contains a cycle \mathcal{C} with k vertices and an extra vertex c , called the center of W , adjacent to all the vertices of \mathcal{C} . The edges of W split the interior of \mathcal{C} into k triangles which we call the triangles of W ; see Figure 10. The following result will be needed.

Lemma 4 *Let W be a k -wheel such that its edges incident to c are labelled e_1, \dots, e_k in the clockwise direction. Suppose further that e_k is an absent edge of W . Then we can exchange the labels of e_1 and e_2 with a linear number of edge rotations, leaving the rest of W unchanged.*

Proof Suppose that W is as in the statement of the lemma, and that the third edge of the triangle of W containing e_1 and e_2 is labelled a ; see Figure 10.

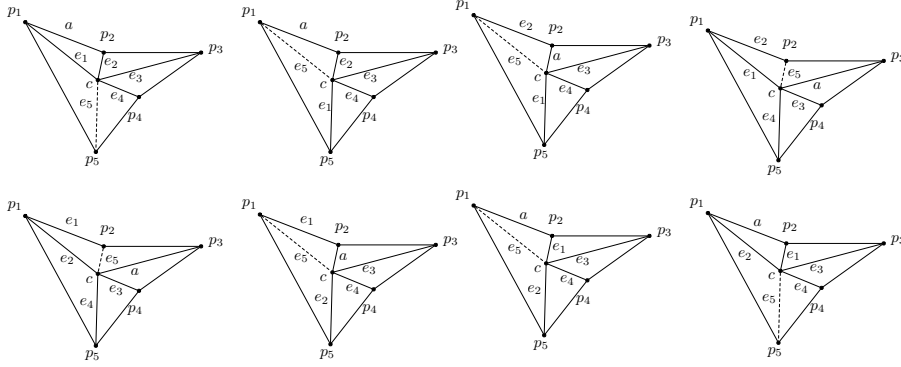


Fig. 10 Exchanging the labels of two edges in a wheel.

The process followed to exchange e_1 and e_2 is illustrated in Figure 10. We first rotate e_1 to the position of e_k , and then exchange a with e_2 . We now rotate e_k counterclockwise $k - 1$ times until it reaches the position of a . We then exchange e_1 and e_2 , and rotate e_k $k - 1$ times, but now clockwise until it reaches the position of e_2 . We now exchange a with e_1 , and finally rotate e_2 to the position of e_k . At this point the labels of e_1 and e_2 have been exchanged, and the rest of the labels in W are back to their original position; see Figure 10. \square

Lemma 5 *Let G be a graph in $PG_L(P, k)$, and e_i and e_j two edges in G . Then the labels of e_i and e_j can be exchanged, leaving the labels of the remaining edges in G unchanged.*

Proof As in the proof of Lemma 1, we consider the dual graph of a triangulation T of P obtained by adding some absent edges to G . We now find the shortest path Π connecting a triangle t_r containing e_i on its boundary, to another triangle t_s containing e_j on its boundary. We prove the lemma using induction on the length of Π .

Suppose first that $t_r = t_s$, and that the third edge of t_r is an absent edge. Then we can exchange the labels of e_i and e_j by executing three edge rotations as shown in Figure 9. Suppose then that t_r contains no absent edge. Two cases arise. Let a be the third edge of t_r , and suppose that by using Lemma 1, we can exchange it with an absent edge e of T , but leaving e_i and e_2 on the

boundary of t_r (this is not possible, for example, if a is an edge of the convex hull of P).

If we can do this, leaving e_i and e_j as edges of t_r , then we proceed to exchange e_i and e_j as before. Next, we return the absent edge of t_r to its original position by executing the sequence of rotations used to move e to a in reverse order. This guarantees that any labels that may have changed while moving e to a are restored to their original positions.

Suppose then that whenever we try to exchange an absent edge e of T with a , we always have to move e_i or e_j away from t_r . Then it is easy to see that before we move e_i or e_j away from t_r , e is incident to e_i or e_j . Suppose then that e is incident to e_i , and that the endpoints of t_r are x, y, z as shown in Figure 13. Observe that if x is in the interior of $\text{Conv}(P)$, we can now rotate e clockwise around x until it reaches a , leaving both e_i and e_j in t_r , which is a contradiction. Then x belongs to the convex hull of P . If z also belongs to the boundary of the convex hull of P , then the quadrilateral containing the two triangles containing e_i on their boundaries is convex, and it is straightforward to verify that we can exchange e_i with e_j in a constant number of edge flips, see Figure 11.

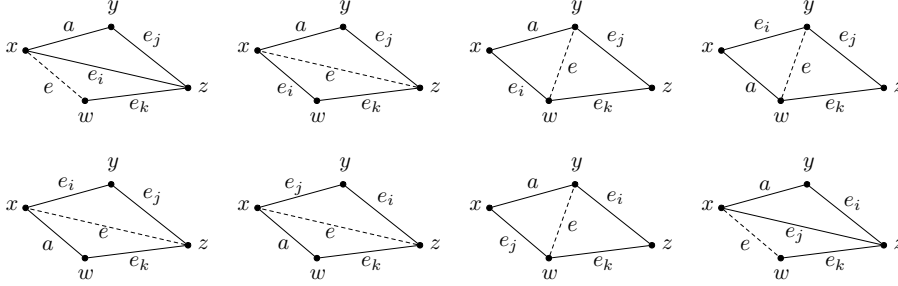


Fig. 11 Exchanging labels using edge rotations.

Suppose then that z belongs to the interior of $\text{Conv}(P)$. Then if e is not incident to z , then with a single rotation, we can make it incident to z ; see Figure 13, and thus the subgraph of T obtained by the union of the triangles of T having c as a vertex is a wheel satisfying the conditions of Lemma 4, and we can now exchange e_i with e_j . As before, return e to its original position to restore the labels that may have shifted. This completes the base case of the induction.

If the path Π between t_r and t_s has length $\ell > 0$, we first exchange the label of e_i with that of the next edge, say e_m , of the triangle strip induced by Π . Then we have a shorter path between e_i and e_j and by induction, we can exchange the labels of e_i and e_j . We then exchange the labels of e_j and e_m . \square

We now show:

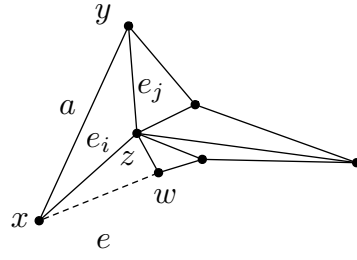


Fig. 12 Exchanging labels using edge rotations.

Theorem 6 $\mathcal{ERG}_{\mathcal{L}}(S, k)$ is connected.

Proof Suppose we have two labelled graphs G and H in $\mathcal{ERG}_{\mathcal{L}}(S, k)$. By Theorem 1, ignoring the labels on the edges of H and G , transform G to H . At this point the labels on the edges of H are permuted. Now using Lemma 5, we move each labelled edge to its final position. \square

It is easy to see that the number of edge flips used to exchange e_i with e_j in Lemma 5 could be quadratic. Thus the diameter of $\mathcal{ERG}_{\mathcal{L}}(S, k)$ is $O(n^3)$.

We remark that when considering *edge flips* in labelled triangulations, the *labelled edge flip graph* is not necessarily connected any more, as can be seen by assigning labels to the graph shown in Figure 7.

We conclude this section by pointing out that the bounds established earlier for the chromatic number of $\mathcal{ERG}_{\mathcal{L}}(P, k)$ are the same as for $\mathcal{ERG}(P, k)$. The proof of the next result is identical to that of Theorem 5.

Theorem 7 The chromatic number of $\mathcal{ERG}_{\mathcal{L}}(P, k)$ is at most the chromatic index $\chi'(K_n)$. This bound is sometimes achieved for n even, and $k = n$. For n odd, the chromatic number of $\mathcal{ERG}_{\mathcal{L}}(P, k)$ is sometimes at least $\chi'(K_n) - 1$.

6 Concluding remarks

An open problem for labelled graphs is that of determining the diameter of $\mathcal{ERG}_{\mathcal{L}}(P, k)$. We believe that for some point sets, the diameter of $\mathcal{ERG}_{\mathcal{L}}(P, k)$ is cubic.

Other types of edge rotations arise in a natural way. For example:

1. We can allow an edge xy to be rotated to an edge xw not in G if $G - xy + xw$ is plane, see Figure 13 a).
2. A more restrictive rotation allows us to replace xy by xw only if y and w are consecutive vertices in the cyclic order of visible (as seen from x) vertices around x , see Figure 13 b).

It is easy to see that these types of edge rotations, can be achieved by performing a sequence of edge rotations as defined in the introduction of our paper. Thus the rotation graphs $\mathcal{ERG}(P, k)$ generated by these new rotations are connected. The bounds proved on the chromatic number of $\mathcal{ERG}(P, k)$ and $\mathcal{ERGL}(P, k)$ remain unchanged.

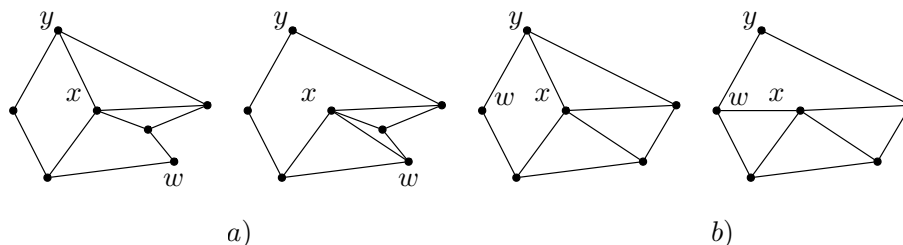


Fig. 13 Different types of edge rotations.

References

1. O. Aichholzer, and K. Reinhardt, A quadratic distance bound on sliding between crossing-free spanning trees. *Computational Geometry Theory and Applications*, 37(3), pp. 155–161 (2007)
2. P. Bose, and F. Hurtado. Flips in planar graphs. *Computational Geometry Theory and Applications*, 42(1), pp. 60–80 (2009)
3. V. Estivill-Castro, M. Noy, and J. Urrutia, On the chromatic number of tree graphs. *Discrete Mathematics*, 223, no. 1–3, pp. 363–366 (2000)
4. W. Goddard, and H. C. Swart. Distances between graphs under edge operations. *Discrete Mathematics*, 161, pp. 121–132 (1996)
5. C. Hernando, F. Hurtado, M. Mora, and E. Rivera-Campo. Grafos de árboles etiquetados y grafos de árboles geométricos etiquetados. in: *Proc. X Encuentros de Geometría Computacional*, pp. 13–19 (2003)
6. F. Hurtado, M. Noy, and J. Urrutia. Flipping edges on triangulations. *Discrete and Computational Geometry*, 22, pp. 333–346 (1999)
7. C. L. Lawson. Transforming triangulations. *Discrete Mathematics*, 3, pp. 365–372 (1972)
8. C. L. Lawson. Software for C^1 surface interpolation. In *Mathematical Software III*, pp. 161–194 (1977)
9. J. Lucas. Untangling binary trees via rotations. *The Computer Journal*, 47, (2), pp. 259–269 (2004)
10. J. Lucas. An improved kernel size for rotation distance in binary trees. *Information Processing Letters*, 110, pp. 481–484 (2010)
11. J. Pallo, An efficient upper bound of the rotation distance of binary trees, *Information Processing Letters* 73 (3-4), pp. 87–92 (2000)