Bounds for Point Recolouring in Geometric Graphs

Henk Meijer
Roosevelt Academy, Middelburg, The Netherlands
(h.meijer@roac.nl)
Yurai Núñez-Rodríguez and David Rappaport *
Queen's University, Kingston, ON Canada
(yurai,daver@cs.queensu.ca)

Abstract

We examine a recolouring scheme ostensibly used to assist in classifying geographic data. Given a drawing of a graph with bi-chromatic vertices, a vertex can be recoloured if it is surrounded by neighbours of the opposite colour. The notion of surrounded is defined as a contiguous subset of neighbours that span an angle greater than 180 degrees. The recolouring of surrounded vertices continues in sequence, in no particular order, until no vertex remains surrounded. We show that for some classes of graphs the process terminates in a polynomial number of steps. On the other hand, there are classes of graphs where the process never terminates.

Keywords: computational geometry, point recolouring, triangulations, geometric graphs

1 Introduction

Given a set of planar points partitioned into red and blue subsets, a red-blue separator is a boundary that separates the red points from the blue ones. There has been considerable investigation of methods for obtaining such red-blue separating boundaries.

In his PhD thesis, Seara [12], examines various means for red-blue separation, including all feasible red-blue separations by a line, by a strip, or by a wedge. For the case of red-blue separation with the minimum perimeter polygon the problem is known to be NP-hard [3, 1]. A somewhat related topic is to obtain a balanced subdivision of red and blue points, that is, faces of the subdivision contain a prescribed ratio of red and blue points. Kaneko and Kano [4] give a comprehensive survey of results pertaining to red and blue points in the plane, including results on balanced subdivisions.

For some applications one is willing to reclassify points by recolouring them so as to obtain a more reasonable boundary. For example Chan [2] shows that finding a red-blue separating line with the minimum number of reclassified points takes $O((n + k^2) \log k)$ expected time, where k is the number of recoloured points.

In Reinbacher et al. [11] a heuristic algorithm that recolours points is presented for obtaining a better delineating boundary. The input is a triangulated set of n planar red-blue points. For a point p to be recoloured it needs to be "surrounded" by points of the

*Supported by an NSERC of Canada Discovery Grant

opposite colour. Reinbacher et al. show experimental results on delineating boundaries after recolouring.

A point is *surrounded* when there is a contiguous set of oppositely coloured neighbours of p, in the triangulation, that span a radial angle greater than $180°$. As the recolouring occurs in an iterative sequence it is not clear that the process will ever come to an end. However, Reinbacher et al. show that no sequence that iteratively recolours surrounded points will ever visit the exact same colouring of the points more than once. Thus the maximum number of recolourings is bounded by the total number of possible colourings which is $2^n - 1$.

Recolouring problems have also been studied in other areas, in some cases under different names. For example, a recolouring-like problem applied to distributed systems with fault propagation has been examined by de la Noval et al. [6]. In their work, recolourings occur synchronously (in parallel). Peleg and others have studied recolourings in parallel, looking for initial configurations that make all the vertices (points) converge to a single colour [10]. In his survey, Peleg poses some open problems regarding the study of asynchronous or sequential recolourings. His asynchronous model coincides with the iterative recolouring model discussed in this note; although for most of our results, except for Theorem 5, we rely on the geometry of the input graph rather than purely combinatorial features, such as the degree of the points. Also, we study the length of recolouring sequences instead of initial configurations that converge to a mono-chromatic colouring.

Our results begin where Reinbacher et al. leave off. Using some of their ideas we are able to obtain an $O(n^2)$ upper bound for the number of recolourings of a triangulated set. We also provide bounds for recolouring sequences in other types of geometric graphs (see Table 1 for a summary of our results on different types of geometric graphs). Preliminary versions of this work have appeared in [7, 8].

Table 1: Summary of main results. The column “All recoloured” indicates whether there is a recolouring sequence that recolours all points.

Type of graph	**Lower bound**	**Upper bound**	**All recoloured**	**Section**
Triangulations	$O(n^2)$	$O(n^2)$	No	3
Convex drawings	$O(n^2)$	$O(n^2)$	No	4.1
Max. degree three	$O(n)$	$O(n)$	No	4
Trees	$O(n)$	$O(n)$	No	4.3
Planar	∞	–	Unknown	4.1
Planar with adjacent convex vertices	$O(n^2)$	$O(n^2)$	No	4.1
Non-planar	∞	–	Yes	4.2
Non-planar containing convex drawing	$O(n^2)$	$O(n^3)$	No	4.2
One-bend planar	∞	–	Yes	5.1

In the next section we formally describe the recolouring problem and present some basic results. We also present some preliminary results obtained by Reinbacher et al. for their upper bound proof. Section 3 precisely describe the recolouring problem in triangulations.

We follow this in the subsequent section with our new results ultimately obtaining an $O(n^2)$ upper bound. In Section 4 bounds are given for other types of straight-line drawings of graphs, such as planar graphs, non-planar graphs, and trees. The last section discusses some extensions of our results.

2 Preliminaries

The input of our recolouring problem consists of a bi-chromatic (red and blue) set of points S in the plane and a set of straight-line segments connecting pairs of points of S. Thus, the input defines a graph G, or more specifically, a drawing of G in the plane. We use the term graph drawing, or simply *drawing*, to refer to such straight-line drawing of a graph in the plane. Consider a drawing D of G in the following. We assume throughout, for simplicity of exposition, that the points of D are in general position and no two points share the same x-coordinate. We colour the edges of D red if its two incident points are red, and blue if its two incident points are blue. If one of the incident points is red and the other is blue we mix the colours to obtain a magenta edge.

The following definitions and lemmas follow, with a few minor modifications, the paper of Reinbacher et al. [11].

Definition 1 *Let the edges of a bi-chromatic straight-line drawing D be coloured as above. Then the* ***magenta angle*** *of a point $p \in S$ is:*

- $0°$, *if p has at most one radially consecutive incident magenta edge,*
- $360°$, *if p has degree greater than one and is incident only to magenta edges,*
- *the maximum angle between two or more radially consecutive incident magenta edges, otherwise.*

Notice that, according to the previous definition, a point with only one neighbour in D has magenta angle $0°$ regardless of the colour of its neighbour (see Figure 1). A *surrounded* point is one with magenta angle larger than $180°$. Therefore, a point of degree zero or one is never surrounded nor recoloured.

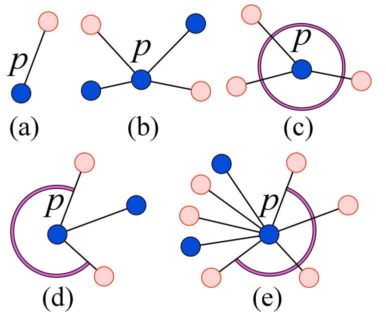

Figure 1: Examples of red-blue point configurations around point p and corresponding magenta angles, α. In a grey-scale output red points appear lighter than blue points. Magenta angles larger than $0°$ are illustrated using arcs of circles. Thus we have angle values (a), (b) $\alpha = 0°$, (c) $\alpha = 360°$, (d), (e) $180° < \alpha < 360°$.

The strategy of reclassification by recolouring, recolours a surrounded point p at each step. Recall that p is surrounded when p has an associated magenta angle that is greater than 180°. The sequence in which surrounded points are recoloured can be driven by a greedy approach, such as recolouring a point with the largest magenta angle. Alternatively we may recolour surrounded points in an arbitrary manner. The recolouring process stops when there are no more surrounded points.

We use the notation $\overline{pq}$ to denote an edge of the drawing D. For descriptive reasons we sometimes write $\overline{qp}$ to denote the same edge, however, the edges are not directed so both $\overline{pq}$ and $\overline{qp}$ denote the same edge.

Definition 2 *We say that the edge $\overline{qr}$ is an **opposite edge** of $\overline{pq}$ $(p \neq r)$ with respect to q if there is no edge between $\overline{qr}$ and the ray from q that goes in the direction opposite to p.*

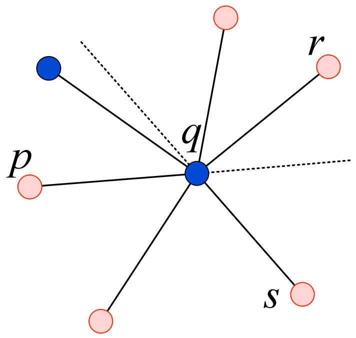

Figure 2: Edges $\overline{qr}$ and $\overline{qs}$ are opposite to $\overline{pq}$. However, $\overline{qp}$ is not opposite to $\overline{sq}$.

For simplification, we omit the "with respect to" qualifier whenever it is clear from the notation. Note that an edge has at most two opposite edges with respect to each endpoint. For example, in Figure 2, both $\overline{qr}$ and $\overline{qs}$ are opposite to $\overline{pq}$. Furthermore, observe the non symmetry of opposites, as $\overline{qp}$ is not opposite to $\overline{sq}$ in the figure.

Definition 3 *Let q be a surrounded point at the beginning of iteration j of the recolouring sequence, then there exists a maximal consecutive sequence of magenta edges incident to q which we denote by $C_q(j) = (\overline{qp_1}, \overline{qp_2}, \ldots, \overline{qp_k})$. We say that the edges $\overline{qp_1}$ and $\overline{qp_k}$ are **extremal** in $C_q(j)$.*

Lemma 1 *For an edge $\overline{pq}$ and any one of its opposite edges with respect to q, $\overline{qr}$, if point q is recoloured, then q receives either the colour of p or the colour of r.*

Proof. [11] Observe that either $\overline{pq}$, $\overline{qr}$, or both are in $C_q(j)$. Thus, if q is recoloured it receives the colour of p or the colour of r. □

We continue with an analogue of Lemma 1 when applied to an edge that is extremal in $C_q(j)$.

Lemma 2 *Let q be a surrounded point, at the beginning of iteration j, such that $\overline{pq}$ is extremal in $C_q(j)$ and $\overline{qr}$ is any opposite edge of $\overline{pq}$. Then both p and r are the same colour, that is not the colour of q.*

Proof. Since q is surrounded and $\overline{pq}$ is extremal in $C_q(j)$ there is a radial span of more than 180° of magenta edges incident to q beginning at $\overline{pq}$ and containing $\overline{qr}$. Thus $\overline{qr}$ is in $C_q(j)$ and both p and r are the same colour, that is not the colour of q. □

Definition 4 *Let D be a straight-line drawing. A* ***monotone chain*** *is a maximal simple path $P = (p_0, \ldots, p_m)$ in D such that*

- $p_0 \neq p_m$
- *the path is ordered from left (p_0) to right (p_m) by x-coordinate*
- *either $\overline{p_i p_{i+1}}$ is opposite to $\overline{p_{i-1} p_i}$ or $\overline{p_i p_{i-1}}$ is opposite to $\overline{p_{i+1} p_i}$, for all $1 < i < m-1$.*

In the previous definition, P is not necessarily the longest monotone path from p_0 to p_m, but it is maximal in the sense that no opposite edge can be added to either end of P such that a longer monotone chain is obtained.

3 Recolouring of Triangulations

Recall that a triangulation T of a point set S is a collection of diagonals incident to every point that partitions the interior of the convex hull of S into triangles [9]. Given a coloured triangulation of n points, Reinbacher et al. [11] show that the number of recolourings is finite, in fact at most $2^n - 1$, independent of the recolouring strategy that is used. They also give an example that following a greedy recolouring scheme results in $\Omega(n^2)$ recolourings. We reproduce this example in Figure 3.

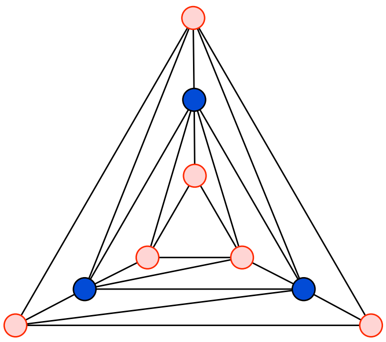

Figure 3: A sequence of recolourings that always recolours a point that is the most nested surrounded point uses $\Omega(n^2)$ recolourings.

Reinbacher et al. also present a number of strategies that always lead to $O(n)$ recolourings. In particular, there is a 2-phase strategy that recolours at most $O(n)$ points: (phase 1) all the surrounded red points are recoloured, (phase 2) all the surrounded blue points are recoloured. The resulting triangulation does not accept any more recolourings since all surrounded red and blue points have already been recoloured. Notice that recolouring blue points (to red) in phase 2 does not create new red surrounded points. However, for red-blue separation purposes, this strategy is not "fair" in the sense that it favours blue. Symmetrically, a red bias strategy exists. The previous strategy and the $O(n)$ bound also apply to other types of geometric graphs referred to throughout this note. However, in the following we do not refer to a specific strategy but to any, and all, strategies that recolour one surrounded point at a time.

Lemma 3 *A surrounded point q on the convex hull of S can be recoloured at most once.*

Proof. Both convex hull neighbours of q must be in $C_q(j)$. Thus q takes on their colour. Such neighbours of q can no longer become surrounded. This implies that q can be recoloured at most once. □

We define an *internal point* as a point that is not on the boundary of the convex hull of T. Similarly, an *internal edge* is an edge with at least one internal endpoint. Note that the convexity of the faces in the triangulation ensures that for any internal edge $\overline{pq}$ that is incident to an internal point q there always exists at least one opposite edge with respect to q, $\overline{qr}$, such that the points p, q, r appear in x-coordinate order.

We can cover T using a set C of monotone chains as follows. For every edge $\overline{pq}$ in T we can obtain a monotone chain $P_{\overline{pq}}$ starting with edge $\overline{pq}$, then adding edges $\overline{rp}$ and $\overline{qs}$, opposite to $\overline{pq}$. This process is repeated iteratively by inserting opposite edges in both directions from $\overline{pq}$ until the points with smallest and largest x-coordinates are reached. These are obviously points on the convex hull of T. See Figure 4 for an example. Planarity implies that we have $O(n)$ edges in T and, therefore, $O(n)$ chains in C. A similar type of covering of a planar subdivision with monotone chains has been used before by Lee and Preparata [5] with a different purpose.

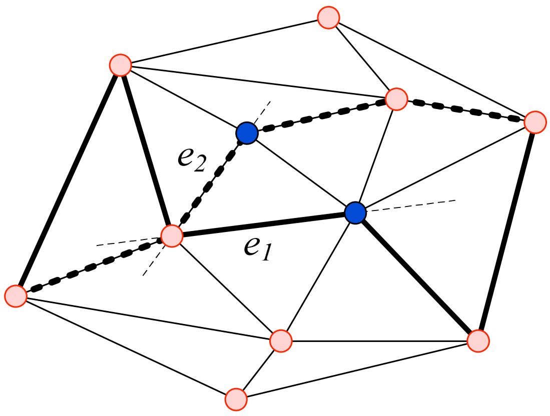

Figure 4: Monotone chains corresponding to the edges e_1 (thick lines) and e_2 (thick dashed lines)

Definition 5 *Let $P = (p_0, \ldots, p_m)$ be a monotone chain. The **colour-change** number of the chain P, $\chi(P)$, is the number of magenta edges in P. Similarly, the colour-change number of a monotone chain cover C, $\chi(C)$, is defined as the total colour-change number over all chains of C.*

Theorem 4 *Given a set of red and blue points, S, together with a triangulation T of S, with $|S| = n$, any recolouring sequence on T will consist of at most $O(n^2)$ recolourings.*

Proof. Consider a cover C of monotone chains for the set of edges of T, and let P be a monotone chain in C. We observe how the colour-change number of C changes with every point recolouring. The analysis is divided into recolourings that occur at the convex hull points, and recolourings that occur at internal points of T. From Lemma 3 it follows that convex hull recolourings can occur at most once per convex hull point. Lemma 1 implies that the colour change number $\chi(P)$ cannot increase when any internal point is recoloured.

We will see that for every recolouring of an internal point, there is always at least one monotone chain whose colour-change number decreases. Suppose that at some step of the recolouring sequence the internal point q changes colour. If the magenta angle of q is less than $360°$ we take an extremal magenta edge $\overline{pq}$ in $C_q(j)$; otherwise all edges incident to q are magenta and we choose $\overline{pq}$ arbitrarily. An opposite edge of $\overline{pq}$ with respect to q, $\overline{qr}$, is in the monotone chain $P_{\overline{pq}}$. Furthermore, point r must be the same colour as p and different from q by Lemma 2. Thus, the colour-change number of $P_{\overline{pq}}$ must decrease by two. See Figure 5.

Obviously, the colour-change number of any monotone chain is at most $n-1$ and can increase (by two) only at its endpoints (points of minimum and maximum x-coordinates) during the entire recolouring sequence. Since the number of chains in C is $O(n)$, $\chi(C)$ is $O(n^2)$. This number is not significantly affected (asymptotically) by the possible linear increase of the colour-change number of C at the points of minimum and maximum x-coordinates. On the other hand, every recolouring of an internal point q produces a decrease in $\chi(P_{\overline{pq}})$ for at least one chain $P_{\overline{pq}} \in C$, while the colour-change number for all other chains in C either decreases or remains unchanged. Thus, $\chi(C)$ decreases (by at least two) with the recolouring of an internal point. This proves that at most $O(n^2)$ internal point recolourings can occur. This together with the $O(n)$ number of convex hull point recolourings add up to $O(n^2)$ recolourings, which completes our proof. □

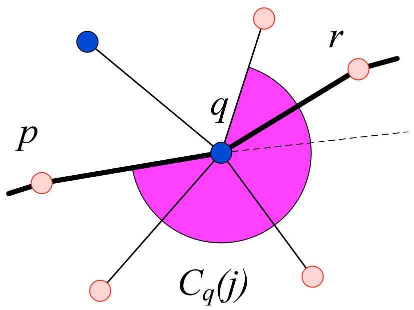

Figure 5: When q is re-coloured, the colour change number of the monotone chain $P_{\overline{pq}}$ (thick line) decreases by two. $C_q(j)$ is represented by the shaded area.

4 Recolouring of Straight Line Drawings

In this section we consider more general graphs and their drawings. At all times the graphs are assumed to be connected because, in general, connected component can be considered independently.

There are families of drawings for which every recolouring sequence is finite and others that allow infinite recolouring sequences. We characterize some of these families. A family of drawings with finite recolouring sequence comes from graphs with maximum degree 3.

Theorem 5 *Let D be a straight-line drawing with n bi-chromatic points. If D has maximum point degree 3, the length of any recolouring sequence of D is $O(n)$.*

Proof. The proof follows from an observation on how the number of magenta edges decreases with each recolouring. Let p be a point that is being recoloured at a certain step along a recolouring sequence. At least two edges incident to p need to be magenta, according to the definition of magenta angle. These edges change to a solid colour after the recolouring of p. Before p is recoloured, at most one edge of solid colour can be incident to it given that $deg(p) \leq 3$. This edge, if it exists, becomes magenta. Therefore, the number of magenta edges decreases by at least one with each recolouring. As the initial number of magenta edges is $O(n)$, the number of recolourings is also $O(n)$. $\square$

The proof of Theorem 5 relies only on the degree of a point. Thus, this bound also holds for non-planar drawings with maximum degree 3. Also, in Section 5 we further extend the scope of this theorem to include non-straight edge drawings.

One may ask whether a recolouring sequence recolours every point at least once. This question is of particular interest for the case of infinite recolouring sequences in planar and non-planar graphs (see Subsection 4.2). The following theorem answers this question, in the negative sense, for drawings with maximum point degree three.

Proposition 6 *Let D be a straight-line drawing with n bi-chromatic points. If D has maximum point degree 3, then no recolouring sequence of D recolours all points.*

Proof. Recall that a point with degree 0 or 1 never gets recoloured. Thus, we assume all points have degree 2 or 3, or else the theorem trivially follows. We continue by giving orientations to the edges. If a point has two consecutive incident edges that form an angle larger than 180°, these two edges are directed "inward" and the third edge, if any, is directed "outward". Otherwise, the three edges are directed inward. Notice that a point is surrounded if and only if its incoming edges are magenta. This process may produce edges being directed inward with respect to its two endpoints. In this case, a dependency between the two points exists: each point needs the other to be of the opposite colour for it to be surrounded. Therefore, at most one of them can change colour to become the same colour of its neighbour along the entire recolouring sequence. The other point, never changes colour. Since there are more inward edges than outward edges, there are edges directed inward with respect to their two endpoints, by the pigeon-hole principle. This completes our proof. $\square$

4.1 Planar Drawings

One may think of obtaining recolouring bounds for planar graphs based on the fact that a planar graph is a subgraph of a triangulation. However, this is not the case. There are simple examples of a graph G, and a subgraph of G, S, where S has either a larger or a smaller number of recolourings than G in the worst case over all initial colourings.

In fact, a recolouring sequence of a planar graph can be infinite. Figure 6 shows an example of a graph and a colour configuration that may lead to an infinite recolouring sequence (see Appendix A, Figures 10 and 11 for a sequence that repeats a colour configuration). The graph in Figure 6 can be slightly modified and made 2-connected, or the minimum point degree increased, for instance, by carefully adding more edges incident to the points of degree one, without affecting the recolouring sequence.

Notice that the example shown in Figure 6 can also be augmented by attaching other points and edges on the "outside" of any of the convex points in the outer face without affecting the infinite recolouring sequence. The recolouring lower bound for planar drawings is generalized in the following theorem.

Theorem 7 *There exist bi-chromatic planar straight-line drawings with 80 or more points that have infinite recolouring sequences.*

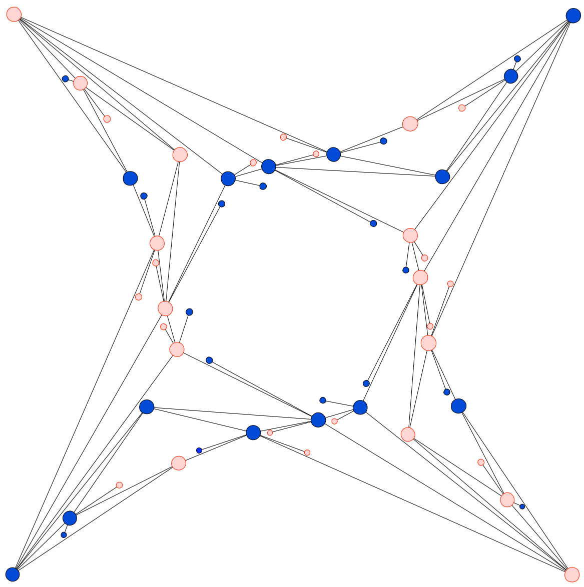

Figure 6: 80 point drawing and initial colouring that leads to an infinite recolouring sequence. See Appendix A for a sequence of recolourings that repeats the initial configuration. The points represented by small circles never change colour.

Our example in Figure 6 was drawn for clarity of exposition. In fact, we have an example of a planar graph with smaller number of points that has an infinite recolouring sequence[1].

Definition 6 *A **convex point** p of a drawing is a point with two consecutive incident edges, the **convex edges** with respect to p, that defines an angle larger than* $180°$.

Lemma 8 *Let p and q be two convex points that share a convex edge $\overline{pq}$. Edge $\overline{pq}$ cannot change from a solid colour to magenta.*

Proof. Figure 7 depicts the only two different scenarios that comply with the hypothesis of the lemma. It is obvious that a span of an angle larger than π around p or q needs to include edge pq given the hypothesis of the lemma. Therefore, if pq is a solid colour, neither p nor q can be surrounded. □

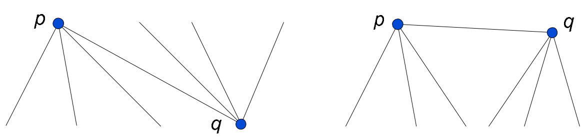

Figure 7: Two cases of adjacent convex points p and q with a convex edge in common.

Theorem 9 *Let D be a planar straight-line drawing with n bi-chromatic points, such that any convex point p_i is connected to another convex point p_j by a convex edge e with respect to both p_i and p_j. A recolouring sequence of D has length $O(n^2)$.*

[1]A 68 point planar drawing with infinite recolouring sequence can be obtained from the example in Figure 6 by merging together points that never change colour and are on the same face.

Proof. The proof of this theorem is similar to the proof of Theorem 4; the only difference is the construction of the cover. In this case, the chains in the cover C do not necessarily end at the points with minimum and maximum x-coordinates. Monotone chains are built, likewise, starting with an edge of the drawing and adding opposite edges in both directions. No more edges can be added to the chains (in each direction) once one of the following conditions is met:

- A point of degree one is reached.
- A convex point is reached.

The first case is benign since, by definition, a point of degree one is never surrounded and the colour-change number of an opposite chain can not increase at this point.

For the second case we show that the colour-change number can increase by one at most per endpoint. Let $P = (p_0, \ldots, p_m)$ be an opposite chain in C and assume p_0 is a convex point –the following analysis similarly applies to p_m. By the conditions of the theorem, there exists a convex point p adjacent to p_0 such that $\overline{p_0 p}$ is a convex edge. By Lemma 8, p_0 can change colour at most once. Therefore, P can increase its colour-change number by one at most at endpoint p_0.

Thus we have shown that in both cases the colour-change number of a monotone chain does not increase significantly. Using this fact we conclude that the number of recolourings is $O(n^2)$ as argued in Theorem 4.

□

Definition 7 *A* ***convex drawing*** *is a planar straight-line drawing where all internal faces are convex and the outer face is defined by the convex hull of the set of points.*

Corollary 10 *(of Theorem 9) A convex drawing with n bi-chromatic points has $O(n^2)$ recolourings.*

For completeness we show that when convex points are adjacent to each other as in the hypothesis of Theorem 9, not all the points can be included in the longest recolouring sequence.

Proposition 11 *Let D be a planar straight-line drawing with n bi-chromatic points, such that any convex point p_i is connected to another convex point p_j by a convex edge e with respect to both p_i and p_j. No recolouring sequence of D recolours all points.*

Proof. Notice that D is finite and, therefore, bounded. Thus, there is at least one convex point. By the hypothesis of the theorem, such a convex point is adjacent to another convex point through a convex edge. Lemma 8 shows that two convex points depend on each other: if one of the points changes colour, the other one does not. Thus, one can always find a point that does not change colour. □

4.2 Non-Planar Drawings

At this point, it is obvious that one can also construct non-planar drawings with infinite recolouring sequences since planar drawings allow so. Nevertheless, the examples shown for infinite recolouring sequences on planar drawings include points that never change colour. In Figure 8 we show an example of a non-planar drawing that has an infinite recolouring sequence in which every point changes colour infinitely many times (see Appendix A, Figure 12 for an example infinite recolouring sequence). If similar examples can be built for planar drawings, these have not yet been found.

Notice that the example shown in Figure 8 can be augmented by attaching additional points and edges to any of the existing points towards the inside of the largest angles between consecutive incident edges without affecting the infinite sequence.

Theorem 12 *There exist bi-chromatic non-planar straight-line drawings with 16 or more points that have infinite recolouring sequences in which every point changes colour infinitely many times.*

As in the planar example, the non-planar example shown in Figure 8 is not of minimal size[2]. For clarity, we do not show a smaller example.

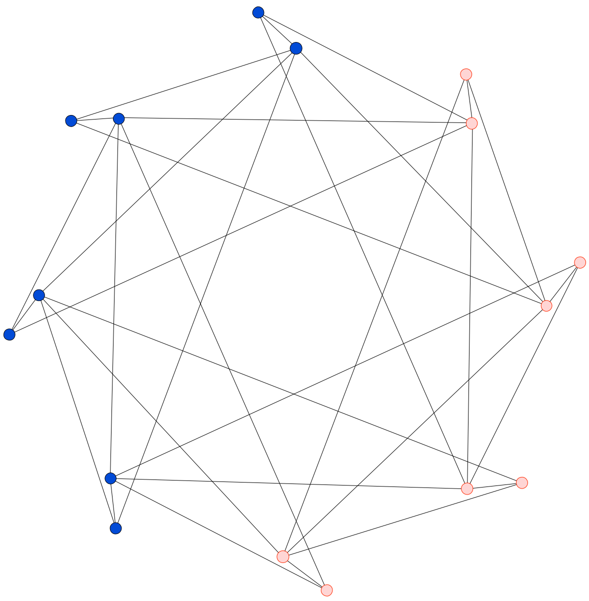

Figure 8: Non-planar drawing and initial colouring that leads to an infinite recolouring sequence. See Appendix A for a sequence of recolourings that repeats the initial configuration.

There are also families of non-planar drawings where recolourings always end after a finite number of steps. One such class has already been characterized in Theorem 5. Another class is formally described in the following theorem.

Theorem 13 *Let D be a bi-chromatic non-planar straight-line drawing with set of points S, $|S| = n$, and let CD be a convex drawing also with set of points S. If $CD \subset D$, the length of any recolouring sequence of D is $O(n^3)$.*

Proof. It is known from Corollary 10 that CD has recolouring sequences of length $O(n^2)$ and this can be proved by using a cover with monotone chains. Since $CD \subset D$, it is always possible to cover D with monotone chains as well. This stems from the availability of opposite edges that maintain the monotonicity of the opposite chains at all times. Each

[2]A 10 point non-planar drawing with infinite recolouring sequence can be obtained from the example in Figure 8 if only 5 pairs of points are used, instead of 8, and edges are slightly changed.

monotone chain has $O(n)$ length. The $O(n^2)$ edges of D are assigned one monotone chain each. Therefore, a cover of D with monotone chains has an overall complexity of $O(n^3)$. Thus, the total colour-change number on all monotone chains is also $O(n^3)$. Based on this fact, we conclude that any recolouring sequence of D has $O(n^3)$ length as argued in Theorems 4 and 9. $\square$

4.3 Trees

In this subsection we use the term *tree drawing* to refer to a straight-line drawing of a tree (not necessarily planar). A trivial example of a tree drawing that has a $O(n)$ recolouring sequence is a "jigsaw" path with points alternately coloured. In such example, all blue points can be coloured to red, leading to approximately $n/2$ recolourings.

An argument similar to the proof of Theorem 9 can be made in order to prove a quadratic upper bound on the number of recolourings of tree drawings, with the difference that opposite chains need not be monotone. Opposite chains will always end at points of degree one (leaves). Consequently, the colour change of one such opposite chain never increases, thus the overall number of recolourings is bounded by $O(n^2)$.

As a corollary of Theorem 5 we have that binary tree drawings have a linear number of recolourings. In this section we prove that the number of recolourings of tree drawings is also linear.

In order to prove a tight bound (Theorem 16) on the number of recolourings of tree drawings, we define a partial order on the recolourings involved in a recolouring sequence. We then bound the total number of recolourings based on the number of minimal elements (sinks) of such partial order. This idea is explained and formalized in the remainder of this section.

We denote a recolouring event r, or simply a recolouring, as the event of a certain point p being recoloured. Let $R = (r_1, \ldots, r_k)$ be a recolouring sequence in which r_i denotes the recolouring at step i, $1 \leq i \leq k$, $k > 0$. We also denote $p(r_i)$ as the point that changes colour at recolouring r_i, and $N(r_i)$ the number of times that $p(r_i)$ has changed colour in R prior to event r_i.

Definition 8 *Let T be a tree drawing and let R be a recolouring sequence of T. The **history graph** of R is a directed graph $H = (R, I)$, $I \subset R \times R$ such that $(r_j, r_i) \in I$ if and only if $\overline{p(r_i)p(r_j)} \in C_{p(r_j)}(j)$ and $i = max(\{l : 1 \leq l < j, p(r_l) = p(r_i)\})$.*

In simple words, the history graph of a given recolouring sequence contains all the recolourings of the sequence as vertices and certain dependencies among the recolourings as edges. More specifically, there exists an edge from recolouring r_j to recolouring r_i if r_i is the last recolouring of a neighbour $p(r_i)$ of $p(r_j)$ prior to r_j, such that $p(r_i)$ is in the magenta angle associated to r_j. See Figure 9 for an example.

Observation 1 *By the definition of history graph all the edges are directed from later recolourings to earlier ones. Therefore, a history graph is a directed acyclic graph (DAG) and defines a partial order on the elements of the recolouring sequence.*

The following lemma proves that there are at least two neighbours of a point p that are recoloured at least once between two consecutive recolourings of p.

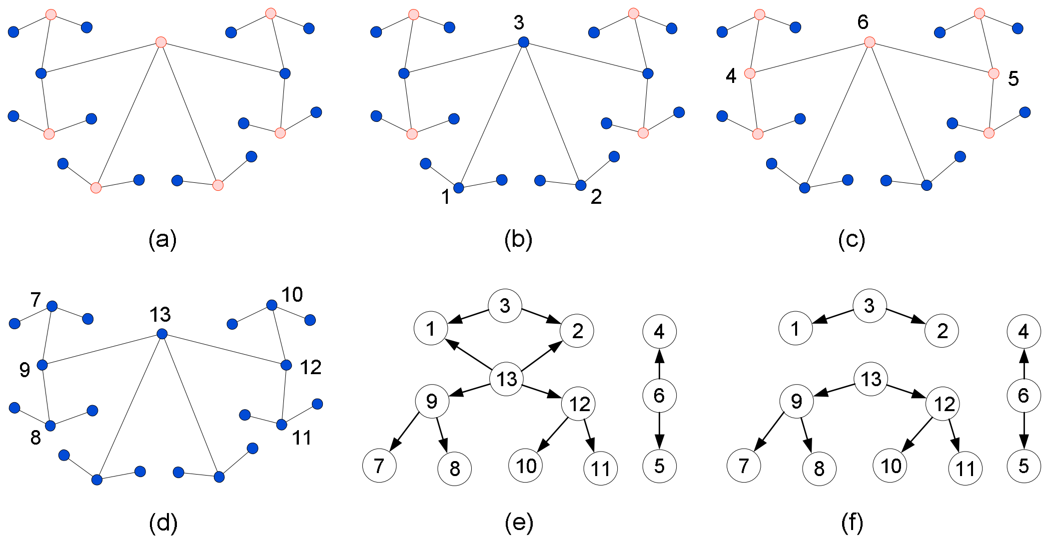

Figure 9: Example tree with recolouring sequence $R = \{r_1, \ldots, r_{13}\}$. (a) original configuration; (b), (c), (d) different stages along the recolouring sequence –the indices of the recolourings appear as labels at the corresponding points–; (e) history graph of R; (f) binary history graph of R.

Lemma 14 *Let T be a tree drawing with n bi-chromatic points, let R be a recolouring sequence of T, and let $H = (R, I)$ be the history graph of R. Consider a recolouring $r \in R$ with $N(r) > 0$. Then the outdegree of r is at least 2. Moreover, there exist two distinct neighbours of $p(r)$, $p_1, p_2 \in T$, with recolourings $s_1, s_2 \in R$, respectively, such that $(r, s_1) \in I$ and $(r, s_2) \in I$.*

Proof. Obviously, if a point is recoloured red (or similarly blue) and was recoloured earlier in the sequence, the previous recolouring was to blue (red). The intersection between the magenta angles at the time it is surrounded by red (blue) and previously by blue (red) contains at least two edges since the corresponding magenta angles are larger than π. Therefore, there are at least two neighbours of p, $p_1, p_2 \in T$, that are recoloured at least once between two consecutive recolourings of p. $\square$

In the light of Lemma 14, we can state the following definition.

Definition 9 *Let R be a recolouring sequence of a tree drawing T and $H = (R, I)$ be the corresponding history graph. The **binary history graph** of R, $BH = (R, BI)$, $BI \subseteq I$, is a subgraph of the history graph where nodes have outdegrees 2 or 0: nodes with outdegree 0 correspond to first time recolourings; nodes with outdegree 2 correspond to subsequent recolourings. Consider a node r_k of degree 2 in the binary history tree. From Lemma 14 we know that there are two distinct neighbours of $p(r_k)$ that have been previously recoloured. Thus we choose the two outgoing edges of r_k $\overline{(r_k, r_i)}$, $\overline{(r_k, r_j)}$, such that i and j are the largest indices smaller than k for neighbours of r_k in the history graph where $p(r_i) \neq p(r_j)$.*

The motivation to define the binary history graph is to obtain a cycle free subgraph of the history graph that involves all the recolourings (see Figure 9 (f)). In the next lemma we prove that the binary history graph is, in fact, cycle free.

Lemma 15 *Let T be a tree drawing, and R a recolouring sequence of T with binary history graph BH. BH has no directed or undirected cycles. Therefore, BH is a forest of trees.*

Proof. Since BH is a subgraph of the history graph of R, BH is also a DAG, by Observation 1. Therefore, there are no directed cycles in BH. Any undirected cycle in BH would have at least one node r with two outgoing edges and one node s with two incoming edges. For the sake of contradiction, we assume that there exists such a cycle, C, in BH.

Consider the function $f : R^* \to V(T)^*$ such that $f(r_1, r_2, \ldots, r_k) = p(r_1), p(r_2), \ldots, p(r_k)$, $r_i \in R, 1 \leq i \leq k$. In particular, f maps a path in BH to a path in T. Let P_1 and P_2 be the two undirected paths that connect r and s in C. By the definition of binary history graph, the outgoing edges of r are incident to nodes t_1 and t_2 such that $p(t_1) \neq p(t_2)$. Thus, $|C| > 2$. Without loss of generality, let t_1 be in P_1 and t_2 be in P_2. Then paths $f(P_1)$ and $f(P_2)$ are different at points $p(t_1)$ and $p(t_2)$. This implies that there are two different paths in T connecting $p(r)$ and $p(s)$. Thus, we establish a contradiction. $\square$

To obtain a bound on the size of binary history trees we show that the number of nodes is linear in the size of the corresponding tree drawing. This will lead us to conclude the results of the following theorem.

Theorem 16 *Let T be a tree drawing with n bi-chromatic points. The length of any recolouring sequence of T is $O(n)$.*

Proof. Let R be any recolouring sequence of T, and let BH be the binary history graph of R. In order to prove this theorem we show that $|V(BH)| = |R|$ is $O(n)$. Let $V_k(BH)$ denote the set of nodes of degree k in BH, and $V_{k^+}(BH)$ be the set of nodes of degree at least k in BH. For accounting purposes, we split the nodes of BH into four classes: $V_0(BH)$, $V_1(BH)$, $V_2(BH)$, and $V_{3^+}(BH)$.

Nodes of degree 0 and 1 are all first-time recolourings (sinks) according to the definition of binary history graph, since these have 0 outgoing edges. Also, nodes of degree 2 are either sinks or sources because internal nodes have degree at least 3, that is, one or more incoming edges and two outgoing edges. The following transformation removes the sources of BH such that, in the resulting graph, all nodes of degree 2 are guaranteed to be sinks.

Let H' be a copy of BH, except that every source r and outgoing edges (r, t_1) and (r, t_2) in BH are replaced by the edge (t_1, t_2) in H'. We already know, from Lemma 15, that there are no undirected cycles in BH. Therefore, edges (t_1, t_2) or (t_2, t_1) could not have existed in BH. Notice that one edge is added in H' for each node removed. Thus,

$$|V(BH)| \leq |V(H')| + |E(H')| \leq 2|V(H')| - m, \tag{1}$$

where m is the number of connected components of H', given that H' is a forest of trees. The degrees of all the nodes remaining in H' is preserved. Therefore, $V_0(H') = V_0(BH)$, $V_1(H') = V_1(BH)$, $V_{3^+}(H') = V_{3^+}(BH)$, and $V_2(H')$ consist of sink nodes only. At most n nodes can be sinks since in the worst case all nodes are recoloured for the first time. Consequently,

$$|V_0(H')| + |V_1(H')| + |V_2(H')| \leq n. \tag{2}$$

Thus, a linear bound on $|V_{3^+}(H')|$ entails a linear bound on $|V(H')|$. We derive such bound in what follows. From properties of graphs and, in particular, of forests of trees,

$$\sum_{r \in V(H')} deg(r) = 2|E(H')| = 2|V(H')| - 2m = 2(|V_0(H')| + |V_1(H')| + |V_2(H')| + |V_{3^+}(H')|) - 2m. \tag{3}$$

According to the definitions of V_k and V_{k^+},

$$\sum_{r \in V(H')} deg(r) \geq |V_1(H')| + 2|V_2(H')| + 3|V_{3^+}(H')|. \tag{4}$$

Equations (3) and (4) lead to

$$|V_{3^+}(H')| \leq 2|V_0(H')| + |V_1(H')| - 2m \leq 2(|V_0(H')| + |V_1(H')|) \tag{5}$$

Combining this with (2) we obtain

$$|V_{3^+}(H')| \leq 2n. \tag{6}$$

Finally, from (1), (2), and (6) we have

$$|V(BH)| \leq 2|V(H')| - m \leq 2(|V_0(H')| + |V_1(H')| + |V_2(H')| + |V_{3^+}(H')|) \leq 6n \tag{7}$$

□

5 Extensions

5.1 Non-straight Edges

Recolourings may also occur in non-straight line drawings. One can consider a point as surrounded whenever the magenta angle defined by the tangents of all the edges leaving the point is greater than 180°. Theorem 5 also holds in this case, that is, any non-straight line drawing with minimum point degree 3 has a recolouring sequence of length at most $O(n)$. However, for more general graphs the results seem to differ from straight line drawings. Figure 13 in Appendix A shows a simple example of a planar graph where the edges have at most one bend and there is a recolouring sequence in which all the points change colour infinitely many times. The infinite recolouring sequence shown involves all the points.

5.2 More Than Two Colours

Suppose that the points come in more than two colours. We define the colour of an edge as the mixture of the colours of its endpoints. In a multicoloured scenario we say that p is *surrounded* by a set of edges of a single mixed colour if the edges define a continuous angle greater than 180°. As we may intuitively observe, increasing the number of colours only lowers the chances of a point being surrounded without changing the fundamental nature of the problem. In fact, inspection shows that all of our previous definitions and results hold in a multicoloured scenario. Thus, our recolouring bounds for a bi-chromatic set of points carries over to multicoloured point sets.

6 Conclusions and Future Work

We have re-examined a point recolouring method useful for reclassifying points to obtain reasonable subdividing boundaries. We show tight finite bounds for point recolouring on triangulations, trees, and graphs of maximum degree 3 for the longest possible recolouring sequences. In contrast, there are examples of finite planar graphs that allow for infinitely many recolourings.

Some interesting questions remain open. First, can planar drawings have sequences of recolourings in which all the points change colour infinitely many times? Another interesting issue is the existing gap between the lower bound ($\Omega(n^2)$) for the number of recolourings in non-planar graphs that contain a convex drawing of the set of points –this bound is provided by the example in Figure 3– and the upper bound of $O(n^3)$ provided in Theorem 13.

7 Acknowledgements

We would like to thank the referees for their helpful comments and suggestions.

References

[1] E. M. Arkin, S. Khuller, and J. S. B. Mitchell, Geometric knapsack problems, *Algorithmica*, 10:399–427, 1993.

[2] T. M. Chan, Low-dimensional linear programming with violations, *Proc. 43rd IEEE Symposium on Foundations of Computer Science (FOCS'02)*, 570–579, 2002.

[3] P. Eades and D. Rappaport, The complexity of computing minimum separating polygons, *Pattern Recogn. Lett.*, 14:715–718, 1993.

[4] A. Kaneko and M. Kano, Discrete geometry on red and blue points in the plane —a survey, *Discrete and Computational Geometry, The Goodman-Pollack Festschrift*, Algorithms and Combinatorics, Springer, vol. 25, 551–570, 2003.

[5] D. T. Lee, and F. P. Preparata, Location of a point in a planar subdivision and its applications, *SIAM J. Comput.*, 6(3):594–606, 1977.

[6] Y. de la Noval, F. Luccio, Y. Núñez-Rodríguez, and L. Pagli, Dynamos in three-dimensional meshes, *Proc. 7th. International Congress on Computer Science Research*, Tampico, Mexico, 2000.

[7] Y. Núñez-Rodríguez, and D. Rappaport, Tight bounds for point recolouring, *Proc. 18th Canadian Conf. in Comp. Geometry (CCCG'06)*, Kingston, ON, Canada, 67–70, 2006.

[8] H. Meijer, Y. Núñez-Rodríguez, and D. Rappaport, On the Complexity of Point Recolouring in Geometric Graphs, *Proc. 20th Canadian Conf. in Comp. Geometry (CCCG'08)*, Montreal, QC, Canada, 143–146, 2008.

[9] Joseph O'Rourke, Computational Geometry in C, *Cambridge University Press*, 1998.

[10] D. Peleg, Local majority voting, small coalitions and controlling monopolies in graphs: A review, *Technical Report CS96-12, Department Of Mathematics & Computer Science, Weizmann Institute Of Science*, 1996.

[11] Iris Reinbacher, Marc Benkert, Marc van Kreveld, Joseph S.B. Mitchell, Jack Snoeyink, and Alexander Wolff, Delineating boundaries for imprecise regions, *Algorithmica*, 50(3):386–414, 2008.

[12] Carlos Seara, On Geometric Separability, *PhD. Thesis. Universitat Politècnica De Catalunya*, 2002.

A Appendix (Infinite Recolouring Sequence on Different Types of Straight-Line Drawings)

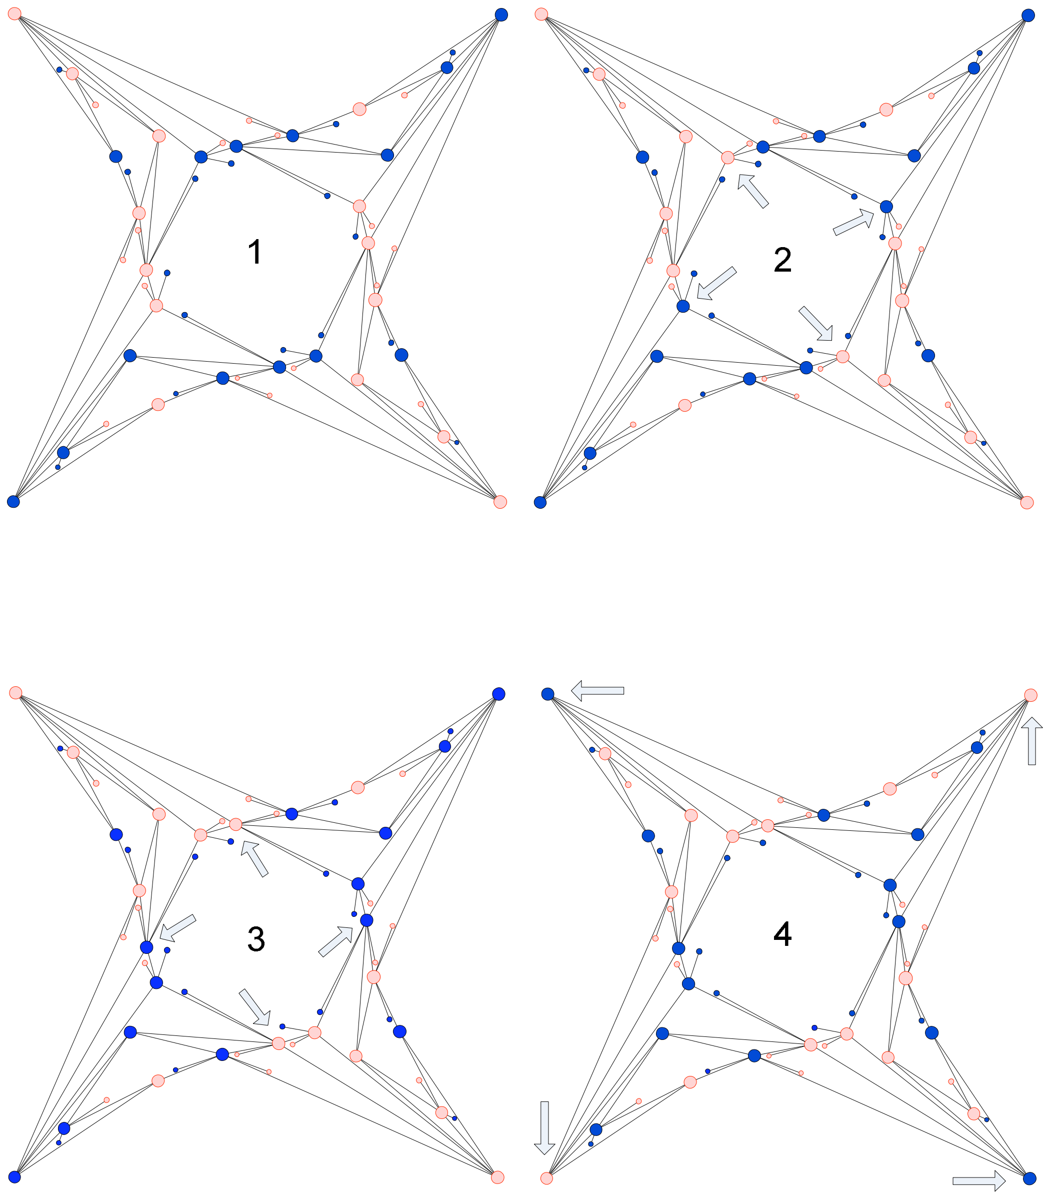

Figure 10: Infinite recolouring sequence of a planar drawing (Part I). Small points never change colour. Recoloured points are pointed out by arrows. The recolourings occur one at a time. The recolouring sequence continues in Figure 11.

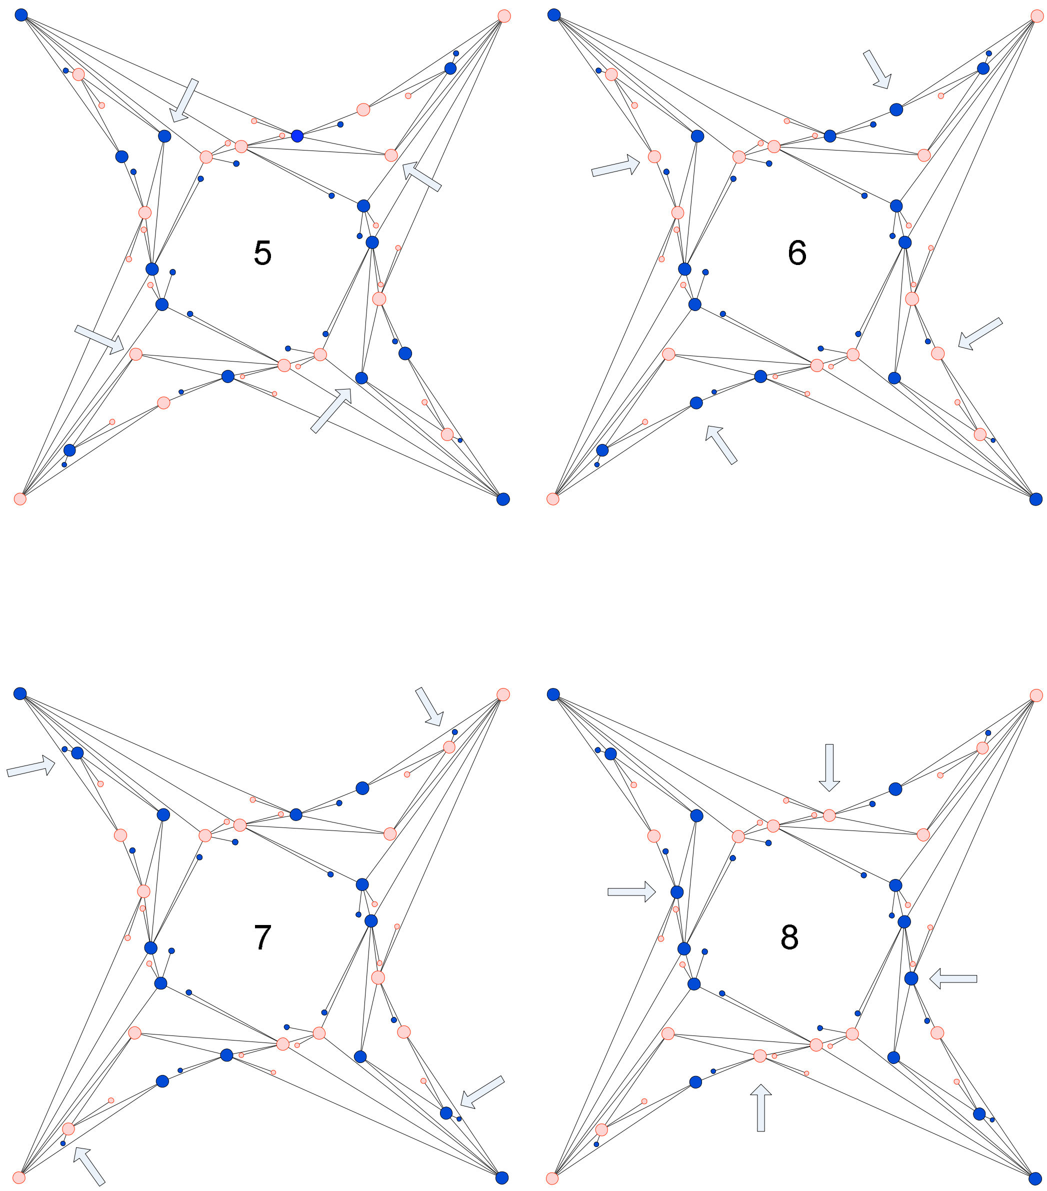

Figure 11: Infinite recolouring sequence of a planar drawing (Part II). Small points never change colour. Recoloured points are pointed out by arrows. The recolourings occur one at a time. Notice that drawing 8 is a rotation of drawing 1 in Figure 10.

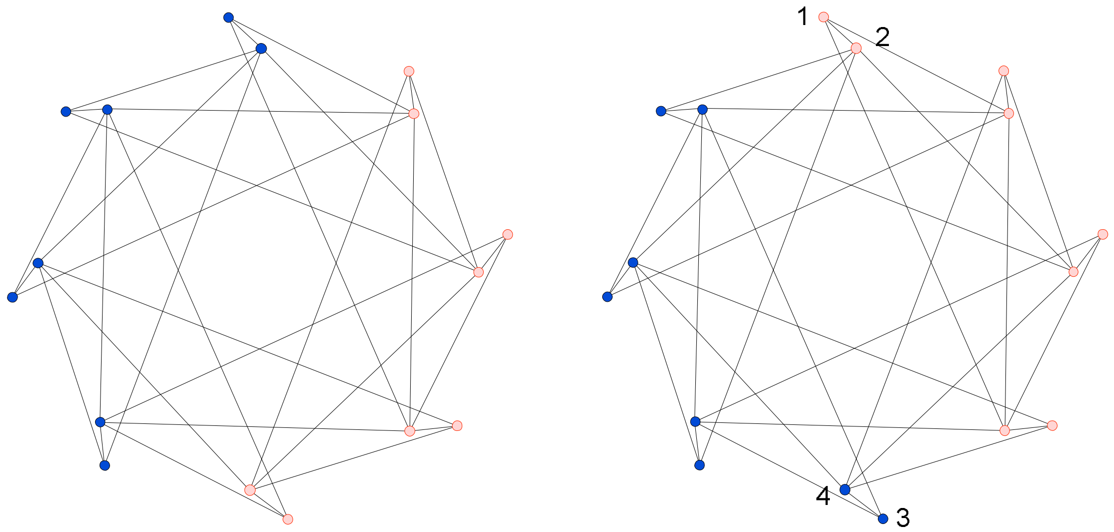

Figure 12: Infinite recolouring sequence of a non-planar drawing. Recolourings occur in the order indicated by the numbers. Notice that the figure on the right is a rotation of the figure on the left.

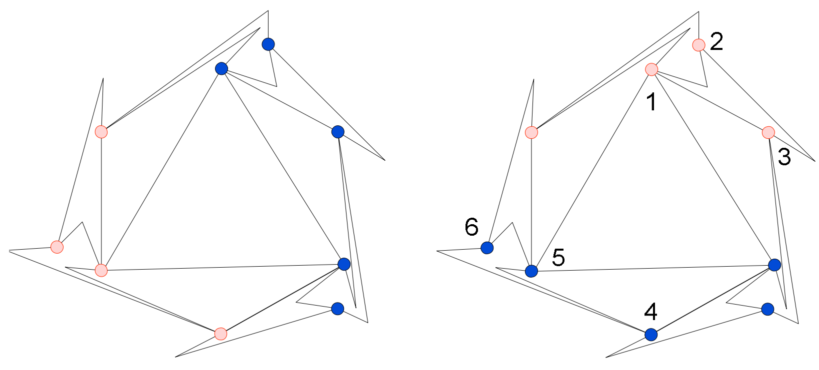

Figure 13: Infinite recolouring sequence of a planar graph with a 1-bend drawing. Recolourings occur in the order indicated by the numbers. Notice that the figure on the right is a rotation of the figure on the left.