Tight Bounds for Point Recolouring

Yurai Núñez and David Rappaport *
Queen's University
Kingston, ON Canada
(yurai,daver@cs.queensu.ca)

Abstract

In their paper on delineating boundaries from geographic data, Reinbacher et al. [9] use a recolouring method to reclassify points so that the boundaries that separate points of different colours become more reasonable. They show that an iterative recolouring method converges after a finite, but exponential, number of steps. They also give, as an example, a triangulated point set with a specified re-colouring sequence that uses $\Omega(n^2)$ recolourings. In this note we revisit this recolouring problem and obtain a matching $O(n^2)$ upper bound.

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 [10] , 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, cells 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. [9] a heuristic algorithm is presented for obtaining a better delineating boundary that recolours points. 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 opposite colour. Reinbacher et al. show experimental results on delineating boundaries after recolouring.

A *surrounded point* is realized 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 until no point is surrounded 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 been studied before, in some cases under different names. A recolouring-like problem applied to distributed systems with fault propagation has been examined by de la Noval et al. [6]. Previous work has been done on graphs where vertices alternate between two states, and state changes correspond to connections in the graph. Peleg and others have studied synchronous state-changes that occur in parallel, looking for initial configurations that make all the vertices end up being in the same state [8]. In this survey paper Peleg poses some open problems regarding the study of asynchronous state-changes controlled by a scheduler. This is a recolouring model that is similar to the model discussed in this paper where the vertices get recoloured one at a time following a given strategy.

Our results begin where Reinbacher et al. leave off. Using some of their ideas we are able to obtain a $O(n^2)$ upper bound for the number of recolourings.

In the next section of this paper we precisely describe the recolouring problem. We also present some preliminary results obtained by Reinbacher et al. for their upper bound proof. We follow this in the subsequent section with our new results ultimately obtaining a $O(n^2)$ upper bound. The remainder of our paper discusses some extensions of our results.

2 Preliminaries

We are given as input a set, S, of points in the plane partitioned into, blue points and red points together with a triangulation T of S. Recall that a triangulation of a point set S is a collection of diagonals incident to

*Supported by an NSERC of Canada Discovery Grant

every point that partitions the interior of the convex hull of S into triangles [7]. We assume throughout for simplicity of exposition that the points are in general position and no two points share the same x coordinate. These assumptions pose no loss in generality as they can be handled in a routine manner. We colour the edges of T red if its two incident vertices are red, and blue if its two incident vertices are blue. If one of the incident vertices 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. [9] .

Definition 1 *Let the edges of T be coloured as above. Then the **magenta angle** Φ of a point $p \in S$ is:*

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

We use the notation $\overline{pq}$ to denote an edge of the triangulation T. 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 *Consider an edge $\overline{pq}$ such that q is not on the convex hull. We say that the edge $\overline{qr}$ is an **opposite edge** of $\overline{pq}$ 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.*

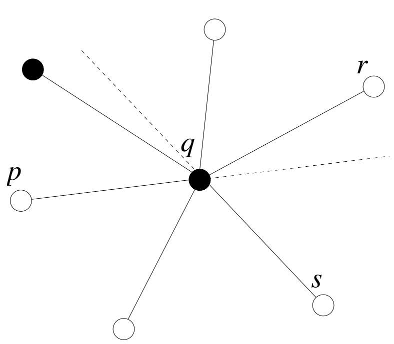

Figure 1: Edges $\overline{qr}$ and $\overline{qs}$ are opposite to $\overline{pq}$ with respect to q. However, $\overline{qp}$ is not opposite to $\overline{sq}$ with respect to q.

Note that a non convex hull edge $\overline{pq}$ has two opposite edges with respect to q. For example, in Figure 1 both $\overline{qr}$ and $\overline{qs}$ are opposite to $\overline{pq}$ with respect to q. Furthermore observe the non symmetry of opposites, as $\overline{qp}$ is not opposite with respect to q for the edge $\overline{sq}$.

The strategy of reclassification by recolouring, recolours a surrounded point p. 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. Alternately we may recolour surrounded points in an arbitrary manner. The recolouring process stops when there are no more surrounded points.

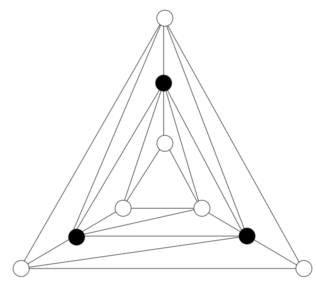

Figure 2: A sequence of recolourings that always recolours a surrounded point that is in the most nested triangle uses $\Omega(n^2)$ recolourings.

Given a coloured triangulation of n points, Reinbacher et al. [9] 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 using a greedy recolouring scheme that uses at most $\Omega(n^2)$ recolourings. We reproduce this example in Figure 2.

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

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

Lemma 5 *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. [9] 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. □

3 New Results

We begin with an analogue of Lemma 5 when applied to an edge that is extremal in $C_q(j)$.

Lemma 6 *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 $\overline{pq}$ with respect to q. 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. □

We introduce the notion of covering the triangulation T by a set of simple paths, or chains.

Definition 7 *A **monotone chain** is a path in the triangulation, $P = (p_0, \ldots, p_m)$ ordered from left to right by x-coordinate, such that*

- *p_0 and p_m are convex hull points*
- *either $\overline{p_i p_{i+1}}$ is opposite to $\overline{p_{i-1} p_i}$ with respect to p_i or $\overline{p_i p_{i-1}}$ is opposite to $\overline{p_{i+1} p_i}$ with respect to p_i, for all $0 < i < m$.*

Note that the convexity of the cells in the triangulation ensures that for any non-convex hull triangulation edge $\overline{pq}$, there is always at least one opposite edge with respect to an internal point q, $\overline{qr}$, such that the vertices p, q, r appear in x-coordinate order.

We can cover T using a set C of monotone chains so that for every edge $\overline{pq}$ in T there is at least one monotone chain in C that contains $\overline{pq}$ and $\overline{qr}$ an opposite edge of $\overline{pq}$ with respect to q and $\overline{ps}$ an opposite edge of $\overline{qp}$ with respect to p. It is easy to obtain such a cover of T using $O(n)$ monotone chains. For every edge e we can obtain a monotone chain P_e that contains e, and then by an iterative process inserts opposite edges in both directions from e until a convex hull point is reached. See Figure 3. Planarity implies that we have $O(n)$ chains in C. A similar type of covering of a planar subdivision with monotone chains has been used before by Lee and Preparata [5].

Definition 8 *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 times an adjacent pair of points on P have different colours.*

Theorem 9 *Given a set of red and blue points, S, together with a triangulation T of S, with $|S| = n$, any recolouring sequence 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 count the number of recolourings that can occur to points contained in P. We separate the count into recolourings that occur at the convex hull points, and

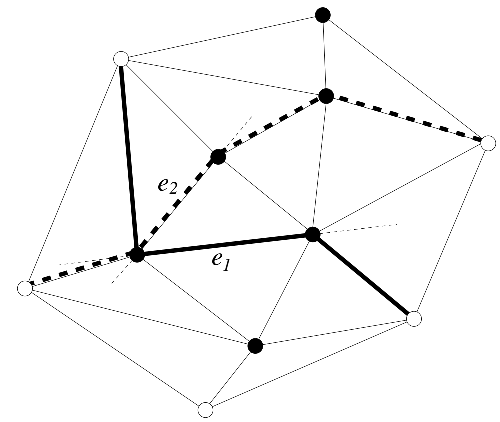

Figure 3: Monotone chains corresponding to the edges e_1 (thick lines) and e_2 (thick dashed lines)

recolourings that occur at internal points of P. From Lemma 4 it follows that convex hull recolourings can occur at most once per convex hull point. Lemma 5 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 vertex 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, vertex r must be the same colour as p and different from q by Lemma 6. Thus, the colour-change number of $P_{\overline{pq}}$ must decrease by two. See Figure 4.

A broad analysis shows that initially the colour-change number of any chain is at most n, and can only be increased by two, at convex hull points, during the entire recolouring sequence. As the total number of recolourings for all internal vertices decreases with every recolouring, it is at most $O(n^2)$. That, together with the linear number of possible recolourings for vertices on the convex hull leads to $O(n^2)$ recolourings overall. □

4 Extensions

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 multi-coloured 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

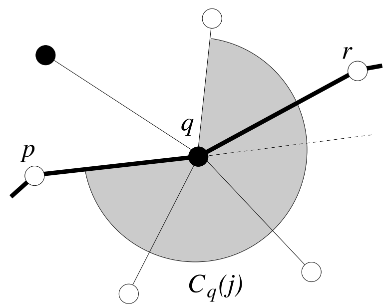

Figure 4: When q is re-coloured, the colour change number of the monotone chain $P_{\overline{pq}}$ (thick line) decreases by two. $C_q(j)$ is represented by the shaded area.

that all of our previous definitions and results hold in a multi-coloured scenario. Thus our recolouring bound for a bi-chromatic set of points carries over to a multi-coloured point set.

A key to our result lies in the fact that we cover the triangulation with chains that are well behaved. Consider the following definition that loosens the requirements of monotone chains slightly by removing the x-coordinate ordering condition.

Definition 10 *An* ***opposite chain*** *is a simple path in the triangulation, $P = (p_0, \ldots, p_m)$, such that*

- *p_0 and p_m are the only convex hull points in P, $p_0 \neq p_m$*
- *either $\overline{p_i p_{i+1}}$ is opposite to $\overline{p_{i-1} p_i}$ with respect to p_i or $\overline{p_i p_{i-1}}$ is opposite to $\overline{p_{i+1} p_i}$ with respect to p_i, for all $0 < i < m$.*

We now make a more general statement regarding our upper bound.

Theorem 11 *Given a set of multi-coloured points, S, together with a connected geometric plane graph G on the points S, with $|S| = n$. If every edge of G can be covered with opposite chains, and all the edges of the convex hull of S are in G then any recolouring sequence will consist of at most $O(n^2)$ recolourings.*

The proof for this theorem is similar to the proof of Theorem 9. Figure 5 provides an example of a graph and its opposite chain cover for which this result applies.

5 Conclusions

We have re-examined a point recolouring method useful for reclassifying points to obtain reasonable subdividing boundaries. We show that no matter what sequence is used to recolour a surrounded point the process enters a state where no point is surrounded in at most $O(n^2)$ recolourings.

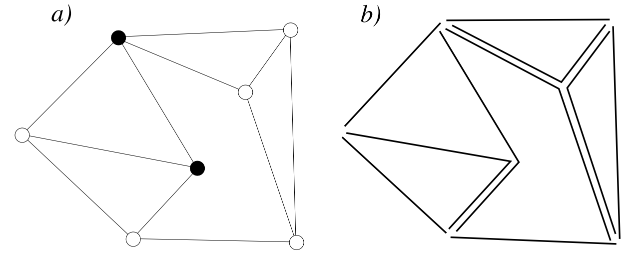

Figure 5: a) A graph G with $O(n^2)$ recolourings. b) An opposite chain cover of G (each chain is a solid line).

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. in *Proceedings of the 43rd IEEE Symposium on Foundations of Computer Science*, 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. in *Discrete and Computational Geometry, The Goodman-Pollack Festschrift*, vol. 25 of Algorithms and Combinatorics, Springer, 551–570, 2003.

[5] D. T. Lee, and F. P. Preparata. Location of a point in a planar subdivision and its applications. in *SIAM J. Comput.*, vol. 6, 235-239, 1997.

[6] Y. de la Noval, F. Luccio, Y. Nuñez Rodriguez, and L. Pagli. Dynamos in Three-Dimensional Meshes. *Proc. 7th. International Congress on Computer Science Research*, Tampico, Mexico, 2000.

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

[8] 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. http://wisdomarchive.wisdom.weizmann.ac.il/archive/00000025

[9] Iris Reinbacher, Marc Benkert, Marc van Kreveld, Joseph S.B. Mitchell, Jack Snoeyink, and Alexander Wolff. Delineating boundaries of imprecise regions. *Proceedings of the EWCG, Eindhoven*, March 9–11, 2005. A longer version of this paper to appear in *Algorithmica* and is obtainable from http://i11www.iti.uni-karlsruhe.de/~awolff/.

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