Encompassing Colored Crossing-Free Geometric Graphs

Ferran Hurtado* Mikio Kano† David Rappaport‡ Csaba D. Tóth§

Abstract

Given n red and n blue points in the plane and a planar straight line matching between the red and the blue points, the matching can be extended into a bipartite planar straight line spanning tree. That is, any red-blue planar matching can be completed into a crossing-free red-blue spanning tree. Such a tree can be constructed in $O(n \log n)$ time.

keywords: geometric graph, spanning tree, color

1 Introduction

Interconnection graphs among disjoint objects in the plane are fundamental in computational geometry, the geometric TSP being a flagship example. Since a minimum length TSP tour of points in the plane has no self-crossing, interconnection graphs are often thought of as *planar straight line graphs (PSLGs)*. Numerous variants of interconnection graph problems were studied in recent years, including Hamiltonian tours, Hamiltonian paths, and spanning trees satisfying various constraints.

This paper addresses two problems on connecting disjoint components of a planar straight line graph. The first problem involves color conforming augmentation of colored graphs into connected PSLGs. A second problem is concerned with the augmentation of 2-edge connected (but monochromatic) PSLGs. A connected graph is *2-edge connected* if at least two edges need to be removed to split the graph into two or more connected components. We have the following results.

• Consider a PSLG G and suppose it has k connected components. Furthermore, the vertices of G are colored so that no two adjacent have the same color. See Fig. 1. We show that one can add $k-1$ straight line edges to G so that we obtain a connected PSLG that conforms to the coloring.

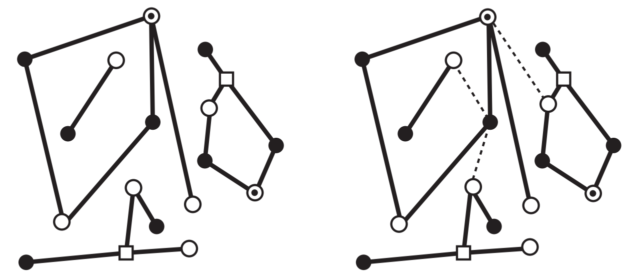

Figure 1: Augmenting a colored disconnected PSLG.

• In particular, if we are given a set of n bi-chromatic line segments, we can find a set of $n-1$ edges so that we are left with a color conforming planar straight line spanning tree. See Fig. 2.

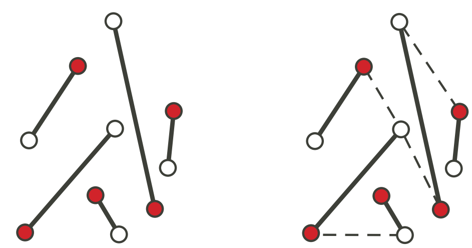

Figure 2: Augmenting disjoint bi-chromatic segments.

• Suppose G is a PSLG consisting of k 2-edge connected components. We can add $2(k-1)$ edges to G so that the result is a 2-edge connected PSLG.

• In particular, we can augment a set of k disjoint triangles with $2(k-1)$ line segments leaving a 2-edge connected PSLG such that every bounded face is a triangle. See Fig. 3.

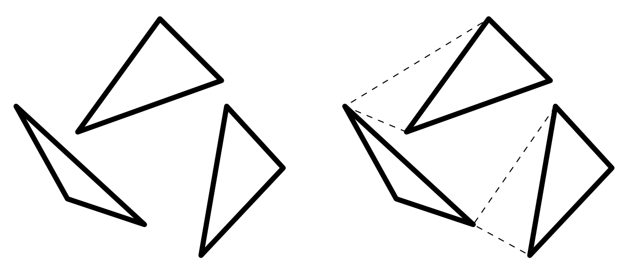

Figure 3: Augmenting a set of triangles to obtain a 2-edge connected PSLG such that every bounded face is a triangle.

We offer a constructive proof for all the above problems based on the following theorem.

*Dept. Matèmatica Aplicada II, Universitat Politècnica de Catalunya, Barcelona, Ferran.Hurtado@upc.es supported by projects DURSI 2001SGR00224 and MCYT BFM2003-0368.

†Department of Computer and Information Sciences, Ibaraki University, Hitachi, kano@cis.ibaraki.ac.jp.

‡School of Computing, Queen's University, Kingston, ON, daver@cs.queensu.ca supported by NSERC of Canada Discovery Grant 9204.

§Department of Computer Science, University of California at Santa Barbara, Santa Barbara, toth@cs.ucsb.edu supported by NSF grants CCR-9901958 and IIS-0121562.

Theorem 1 *For any two finite PSLGs whose planar drawings are disjoint, one of the graphs has a vertex that sees an entire edge of the other graph.*

Note that the roles of two PSLGs, A and B, in Theorem 1 are not symmetric: It is possible that A has no vertex that would see an entire edge of B; in this case a vertex of B sees a full edge of A. In Section 2, we show that if the convex hull of (drawing of) B does not contain A then a vertex of A sees an entire edge of B.

Theorem 1 leads to an $O(n^2)$ time algorithm to construct a color conforming spanning tree in the first problem and a 2-edge connected augmented graph in the second problem for an input of size n. In Section 4, we provide an alternative proof for the first problem that can be turned into a $O(n \log n)$ time algorithms.

Theorem 2 *Every set of bi-chromatic line segments, where any two segments are either disjoint or share an endpoint, can be extended to a color conforming and connected PSLG in $O(n \log n)$ time.*

1.1 Related previous results

Colored PSLGs. Geometric graphs on red-blue points have received increasing attention recently. For a set R of red and B of blue points in the plane, $K(R, B)$ denotes the geometric bipartite graph whose vertex set is $R \cup B$ and whose edges are the red-blue line segments. A path in $K(R, B)$ is necessarily *alternating* between red and blue points. It is well known that for n red and n blue points in the plane, there is always a crossing free perfect red-blue matching (e.g., by repeated application of the *ham sandwich* theorem [11]).

For n red and n blue points in the plane, $K(R, B)$ does not always contain a crossing-free Hamiltonian tour [1]. Kaneko, Kano, and Yoshimoto [10] proved that such a Hamiltonian tour have $n - 1$ self-crossings in the worst case. Kaneko and Kano [9] showed that if $|R| = \Theta(|B|^2)$ then there is an alternating path containing all *red* points. Kaneko [7] proved that for any n red and n blue points in the plane, there is a color conforming connected PLSG of maximal degree three.

These and many other interesting results on geometric red-blue graphs can be found in a recent survey paper of Kaneko and Kano [8].

Encompassing graphs. Given a set of pairwise disjoint line segments in the plane, an *encompassing tree* is a PSLG whose vertex set is the set of segment endpoints and contains every input segment as an edge.

Notice that every encompassing path consists of input segments and non-input segments alternately. Not every set of segments admits a Hamiltonian encompassing path. Pach and Rivera-Campo [12] showed that every set of n segments have a subset of size $\Omega(n^{1/3})$ for which an Hamiltonian encompassing tree exists. The longest alternating path not crossing any of the initial n segments has size $\Theta(\log n)$ in the worst case [5]. Bose, Houle, and Toussaint [3] proved that every set of disjoint line segments in the plane can be augmented to a connected PSLG of maximal degree three. They can construct such a tree for n segments in $O(n \log n)$ time. Later, Hoffmann and Tóth [6] proved that there is also an Hamiltonian encompassing graph of maximal degree three.

2 Proof of Theorem 1

In order to prove Theorem 1 we first establish the following two lemmas.

The proof of the first lemma is based on a result by Avis and Fukuda [2]. It shows that a point external to the convex hull of PSLG sees every point of (the drawing of) at least one edge. The proofs of both lemmas are available in the full version of this paper.

Lemma 3 *Let a be a point external to the convex hull $CH(B)$ of (the drawing of) a PSLG B. Then there exists an edge e in B such that a sees all of e.*

Proof. See the full paper. □

Our second lemma uses a similar argument.

Lemma 4 *Let $\alpha, a_0, a_1, \ldots, a_k, \omega$ denote the boundary of the relative convex hull of B relative to A so that α and ω are in B and $a_0, a_1, \ldots, a_k$ are a reflex chain in A. Then there is a vertex $a_i, i \in \{0 \ldots k\}$ that sees an entire segment in B.*

Proof. See the full paper. □

These two lemmas now establish Theorem 1.

Proof of Theorem 1. Consider two PSLGs A and B whose planar drawings are disjoint. Assume that a vertex of the convex hull of $A \cup B$ is a vertex of A (in other words, the convex hull of B does not contain that of A). We show that a vertex of A sees an entire edge of B.

In the case where the convex hulls $of A$ and that of B are disjoint, we find a vertex $a \in A$ for which we can apply Lemma 3. Let $CH(B)$ denote the boundary of the convex hull of B and let a be a point on $CH(A)$. We say that a *bridges* $CH(B)$, if there are two distinct support lines of $CH(B)$, L_1 and L_2 passing through a and points α and ω in B such that the the region bounded by L_1, L_2 and $CH(B)$ contains no point of A except for a. One way to find this bridge is to "inflate" the convex hull of B until it hits a vertex in A. By Lemma 3, a sees an entire edge of B and the visibility is not occluded by edges of A either.

In the case where the convex hull of B relative to A is incident to A let $\alpha, a_0, a_1, ..., a_k, \omega$ denote the boundary of the convex hull of B relative to A so that α and ω are in B and $a_0, a_1, ..., a_k$ are in A. It follows from Lemma 4 that there is an $a_i, i \in \{0...k\}$ that sees and entire edge in B.

3 Applications

For our first application we consider a plane drawing of a graph with k connected components and with vertices colored so that no edge of the graph is monochromatic. We want to add $k-1$ edges so that we are left with one single connected component with no monochromatic edges.

We proceed by induction on the number of components. If there is only one component, then the input graph is connected. Otherwise we partition the input in two disjoint parts which we call A and B. It follows from Theorem 1 that a vertex v of A or B sees at least one entire edge of B or A, respectively. Since no edge is monochromatic, either $\{a, w\}$, or $\{a, u\}$ is a color conforming connection between A and B. Augment the input graph by this edge: the number of connected components drops by one—induction completes the proof.

For our second result, assume that we have a planar drawing of a graph such that each component of the graph is 2-edge connected. One example of such an input is a set of disjoint triangles. Suppose that there are k connected components in the input. We want to augment this drawing with $2(k-1)$ edges so that we have a single component that is 2-edge connected. This problem closely resembles the problem above where we augment a colored straight line drawing of a planar graph. The only difference is that we connect a vertex to both endpoints of the visible line segment.

4 Proof of Theorem 2

Our second proof for the first problem is similar in some sense to that of Bose et al. [3]: We construct a convex partition of the free space around the line segments, and then add non-crossing edges in each of the convex faces. Since the number of vertices lying along a single face can be arbitrary and we have no control over the distribution of red and blue segment endpoints incident to a single cell, we cannot give a bound on the maximal degree of the resulting spanning graph.

4.1 Convex partitioning

Assume that we are given n pairwise disjoint segments in the plane. We assume, for simplicity, that no segment is vertical and there are no two collinear segments. The free space around the segments can be partitioned into $n+1$ convex cells by the following two-phase partitioning algorithm (Bose et al. [3] used a similar partition):

In the first phase, we sweep the plane from left to right. We extend every input segment beyond its *right* endpoint simultaneously to the sweep line. If an extension hits an input segment, then it stops there. If two extensions meet then they are merged into one extension as follows: if their slopes have opposite signs then both extension continue as a horizontal extension; if they have the same sign then the extension whose slope has smaller absolute value continues and the other stops. In the second phase, every segment is extended beyond its *left* endpoint in a right-to-left plane sweep. An extension stops if it hits an input segment or a previous extension. We apply the same rules as in the first phase if two extensions meet. (See Fig. 4 for an example.)

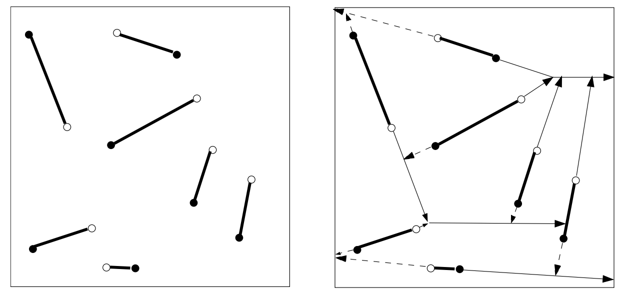

Figure 4: Disjoint segments and a convex partition.

The segments and their extensions form a cell complex in the plane. Every cell is convex and the number of cells is exactly $n+1$. We say that a portion of an input segment (or an extension) is an *edge of the complex* if it lies on the common boundary of two cells. The *vertices of the complex* are the segment endpoints and points lying on the common boundary of three cells. Observe that none of the edges of the cell complex is vertical.

We define an orientation on the edges: The input segments have *no* orientation; an extension edge is oriented left-to-right (resp., right-to-left) if it was created in the left-to-right (resp., right-to-left) plane sweep of the partition algorithm. The orientation of an edge e on the boundary of a cell C is clockwise or counter-clockwise *with respect to cell* C. We use a simple but key property of the cell complex in our main argument:

Lemma 5 *If the boundary of a cell C contains edges of both clockwise and counter-clockwise orientation w.r.t. C, then the boundary of C must contain an entire input segment.*

Proof. See the full paper. □

Using the terminology of [3], we call every connected component of oriented edges an *extension tree*. From every point of an extension tree, the orientations lead to a common point (*root*), lying on an input segment or at infinity.

4.2 Two phase algorithm

We construct the required graph in two phases.

First phase. Consider a set of bi-chromatic segments and convex partition obtained by the above algorithm. We construct a PSLG G_1 by augmenting the input matching with edges between segment endpoints incident to a common cell. If a cell C is incident to red (resp., blue) vertices only, then we add no edges in C. If a cell C is incident to both red and blue segment endpoints then we connect them by a spanning tree within C:

Lemma 6 *For a set P of a red points and b blue points in convex position, $a, b \geq 1$, one can construct a red-blue planar straight line spanning tree in $O(a+b)$ time.*

Proof. See the full paper. □

The resulting bipartite PSLG G_1 is not necessarily connected yet. See Fig. 5 for an example. We can establish connectivity of the points lying on a common *extension trees*, though.

Lemma 7 *Every segment endpoint incident to the same extension tree of the cell complex belong to the same connected component of the graph G_1.*

Proof. See the full paper. □

Similarly, vertices of extension trees whose roots are on the same input segment belong to the same component of G_1.

Lemma 8 *All segment endpoint whose extension tree hits an input segment q_1q_2 on the the same side are in the same connected component of G_1.*

Proof. See the full paper. □

Second phase. We add one more edge to the graph for each connected component. Let us denote the connected components of G_1 by $L_1, L_2, \ldots, L_h$, ordered according to the x-coordinate of the right-most segment endpoint of each component (i.e., L_1 contains the overall right-most segment endpoint). We describe how to connect the components L_1 and L_2 by a red-blue straight line edge while maintaining a PSLG. Iterating this step leads to the required graph G_2.

Let p denote the right-most segment endpoint of L_2. Consider the extension tree T of p. The root of the extension tree cannot be at infinity, because then T would be incident to a segment endpoint s which lies to the right of p, that is, $s \in L_1$, and by Lemma 7, L_1 and L_2 were connected. Assume that the root of T is r and r lies on a segment q_1q_2. Since $q_1q_2 \in L_1$, the endpoints q_1 and q_2 are not connected to L_2 in the graph G_2.

Let C_1 and C_2 be the cells incident to q_1 and the q_2, resp., on the left side of q_1q_2. Both C_1 and C_2 are incident to segment endpoints whose extension tree hits the left side of q_1q_2, and by Lemma 8 belong to L_2. Let p_1 and p_2 denote the segment endpoints of L_2 incident to C_1 and C_2, respectively.

Every segment endpoint incident to C_1 (resp, C_2) have the same color, otherwise L_1 and L_2 would be connected by a subgraph within C_1 (resp., C_2). We conclude that the graph G_1 has no edges within C_1 or C_2.

Consider the triangle Δ formed by the lines q_1q_2, q_1p_1, and q_2p_2. Let p' be the right-most segment endpoint in $\Delta \setminus \{q_1, q_2\}$. We argue that $p' \in L_2$: If $p' = p_1$ or $p' = p_2$, then obviously $p' \in L_2$. Otherwise p' lies to the right of both p_1 and p_2. Clearly, p' is a left segment endpoint and its extension (an x-monotone curve), must hit either q_1q_2 or the extension tree of p_1 or p_2. In any case, the root of the extension tree of p' lies on q_1q_2.

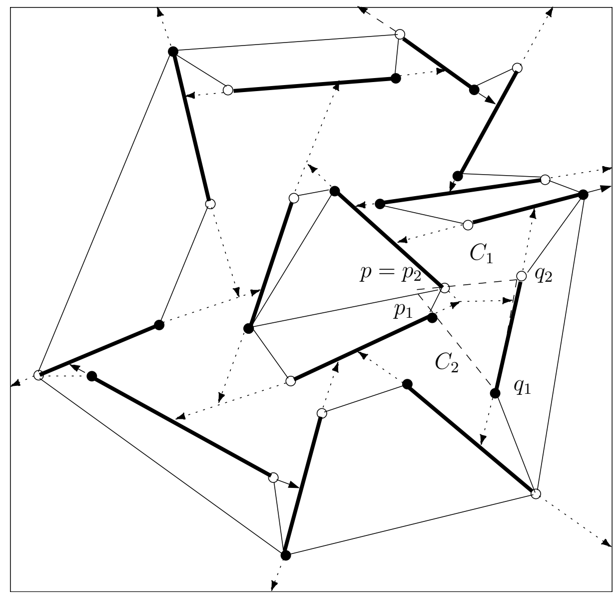

Figure 5: Connecting the components of G_1 in the second phase.

Finally, note that q_1 and q_2 have different colors, and so $p'q_1$ or $p'q_2$ is a bi-chromatic edge. It does not cross any edge of G_1, because the interior of Δ is disjoint from G_1. We connect L_1 and L_2 by augmenting G_1 with either $p'q_1$ or $p'q_2$.

Computational complexity. We can compute our color-conforming planar straight line spanning tree in $O(n \log n)$ time. We sort the right (resp., left) endpoints of the segments in $O(n \log n)$ time. Each sweep-line algorithm is completed in $O(n \log n)$ time. The size of the resulting cell complex (together with orientation of segment extensions) is $O(n)$. We can add edges in all bi-colored cells in $O(n)$ total time. We can detect connected components of G_1 find complete the second phase of the algorithm in $O(n)$ time.

5 Open problems

We have shown that a color conforming spanning tree of a set of bi-chromatic line segments is always obtainable. What about a color conforming spanning tree with the minimum weight where the weight is computed as the sum of the Euclidean distances of the added edges? Given a set of points in the plane it is well known that a greedy algorithm always provides an optimal solution and the solution has no crossings. Bose and Toussaint showed that the minimum spanning tree that augments a set of line segments does not have any crossings [4]. However the minimum spanning tree of bi-chromatic line segments may introduce crossings, as is illustrated by the small example in Fig. 6. It would be interesting to explore methods for determining a color conforming minimum weight spanning tree of a set of bi-chromatic edges.

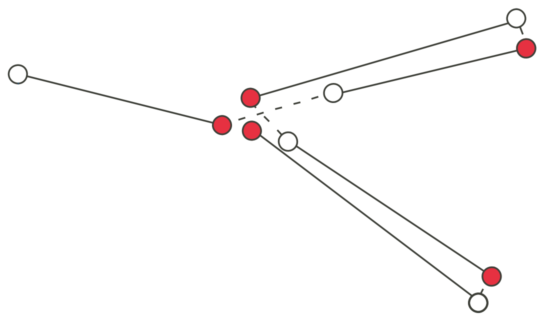

Figure 6: A color conforming minimum spanning tree for this example is not planar.

References

[1] J. Akiyama and J. Urrutia, Simple alternating path problem, *Discrete Math.* **84** (1) (1990), 101–103.

[2] D. Avis and K. Fukuda, A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra, *Discrete Comput. Geom.* **8** (3) (1992), 295–313.

[3] P. Bose, M. E. Houle, and G.T. Toussaint, Every set of disjoint line segments admits a binary tree, *Discrete Comput Geom.* **26** (2001), 387–410.

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

[5] M. Hoffmann and Cs. D. Tóth, Alternating paths through disjoint line segments, *Inform. Proc. Letts.* **87** (2003), 287–294.

[6] M. Hoffmann and Cs. D. Tóth, Segment endpoint visibility graphs are Hamiltonian, *Comput. Geom. Theory Appl.* **26** (2003), 47–68.

[7] A. Kaneko, On the maximum degree of bipartite embeddings of trees in the plane, in *Discrete and Computational Geometry, JCDCG'98*, vol. 1763 of LNCS, Springer, 2098, 2000, pp. 166–171.

[8] 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, 2003, 551–570.

[9] A. Kaneko and M. Kano, On paths in a complete bipartite geometric graph, in *Discrete and Computational Geometry, JCDCG'00*, vol. 2098 of LNCS, Springer, 2001, pp 187–191.

[10] A. Kaneko, M. Kano, and K. Yoshimoto, Alternating Hamiltonian cycles with minimum number of crossings in the plane, *Internat. J. Comput. Geom. Appl.* **10** (2000), 73–78.

[11] C. Larson, *Problem solving through problems*, Springer, 1983, pp. 200–201.

[12] J. Pach and E. Rivera-Campo, On circumscribing polygons for line segments, *Comput. Geom. Theory Appl.* **10** (1998), 121–124.