Circumscribing Polygon of Disjoint Line Segments *
(extended abstract)

Andranik Mirzaian

Department of Computer Science
York University
North York, Ontario M3J 1P3
Canada

* This work is partly supported by Natural Sciences and Engineering Research Council of Canada grant OGP0005516.

Circumscribing Polygon of Disjoint Line Segments
(extended abstract)

1. Introduction

Throughout this paper the domain of discussion is with respect to the Euclidean plane. It is well known that any finite set of points admits a simple polygon with the given points as its vertices [3]. What is a suitable generalization of this fact from points to (disjoint) line segments? Let $\Sigma = \{ S_1, \cdots, S_n \}$ be a set of n pairwise disjoint line segments. Rappaport [5] defined a *simple circuit* of Σ to be a simple polygon Q whose vertices are the endpoints of the segments in Σ, and every segment in Σ is an edge of Q. He showed that not every such Σ has *a simple circuit*, and to decide whether it does is NP-complete. The set Σ is said to be *extremally situated* if each segment in Σ has at least one of its endpoints on the boundary of the convex hull of Σ. Rappaport et al. [6] showed that if Σ is extremally situated, whether it admits a simple circuit can be decided in $O(n \log n)$ time, and a simple circuit of Σ, if one exists, can be constructed within the same time bound. We define a *circumscribing polygon* of Σ (not necessarily extremally situated) to be a simple polygon P such that vertices of P are the endpoints of the segments in Σ and every segment in Σ is either an edge *or* an internal diagonal of P. Note that any simple circuit of Σ is a circumscribing polygon of Σ (but not conversely). We first propose the following (see Fig. 1):

Conjecture. *Any finite set of pairwise disjoint line segments admits a circumscribing polygon.*

This conjecture, if true, answers the following question of Marcotte and Suri [4]. Let M_o be the weight of an optimum (Euclidean) matching of a set of (an even number of) points, P be a simple polygon that spans the set of points, and M_P be the weight of the optimum matching in which all matching edges are constrained to remain in the polygon. The question is: what polygon P achieves the minimum ratio M_P/M_o? The above conjecture implies P is a circumscribing polygon of an optimum matching of the point set, and $M_P/M_o = 1$ is the minimum ratio.

In the rest of the paper assume $\Sigma = \{ S_1, \cdots, S_n \}$ is a set of n extremally situated pairwise disjoint line segments S_i. The main results of this paper are the following two theorems:

Theorem 1. *Any extremally situated Σ admits a circumscribing polygon.*

Theorem 2. *There is an algorithm that constructs a circumscribing polygon of extremally situated Σ in linear space and $O(n \log n)$ time, and this is optimal.*

In a first attempt we may try to construct simple polygons, in $O(n)$ time, that encapsulate the given segments by going around the convex hull in an Euler tour fashion (see Fig. 2), then use a triangulation of these polygons to proceed further. The apparent difficulty in this approach is in maintaining convexity as we descend down to subproblems.

Our proofs are based on a divide-and-conquer technique and employ a novel structure called the *tournament pseudoforest* of Σ. (The term *pseudoforest* has been used in [2] in a different context.) The convex planar subdivision induced by the tournament pseudoforest of Σ is a linear space data structure and can be constructed in optimal $O(n \log n)$ time using the sweep method. Given this data structure, a circumscribing polygon of Σ can be constructed in $O(n)$ time using divide-and-conquer.

Let $\partial\Sigma$ denote the boundary of the convex hull of Σ. A segment S_i in Σ is called an *edge segment*, a *diagonal segment*, or an *internal segment*, if it is, resp., an edge of $\partial\Sigma$, a diagonal of

$\partial\Sigma$, or has only one endpoint on $\partial\Sigma$. In the latter case the endpoint of S_i on $\partial\Sigma$ is called its *head* (denoted h_i) and the other endpoint is called its *foot* (denoted f_i). Each endpoint of an edge segment or a diagonal segment is considered as both its head and foot! To simplify the discussion, we assume (a) no segment S_i is vertical, (b) no three endpoints of segments in Σ are collinear, and (c) no three segments are concurrent if extended.

2. Proof Sketch of Theorem 1

The proof is by induction on the size plus the number of internal segments of Σ. We need to prove a slightly stronger version in which we allow a pair of edge segments in Σ to have a common head. If all segments of Σ are edge segments, then $\partial\Sigma$ is a circumscribing polygon of Σ. If there is a diagonal segment S_i, then S_i divides the problem into two smaller subproblems one on each of its sides. Otherwise proceed as follows: Consider a sequence $\sigma = (\sigma_1, \sigma_2, \cdots, \sigma_i)$ of segments in Σ constructed as follows. The initial segment σ_1 is any internal segment in Σ. Let $\bar{\sigma}_j$ denote the *extension* of σ_j obtained by the following process: extend σ_j along the direction of its supporting line from the side of its foot until it hits either $\partial\Sigma$, one of $\bar{\sigma}_k$ $(k<j)$, or a new segment in $\Sigma - \partial\Sigma$. In the latter case the new segment becomes σ_{j+1}. Furthermore, let $t_j = closure(\bar{\sigma}_j - \sigma_j)$. See Fig. 3. We construct the sequence σ until $\bar{\sigma}_i$ hits either (i) $\partial\Sigma$, or (ii) some $\bar{\sigma}_k$, $k<i$. (In case (ii) we have completed a cycle and w.l.o.g. we may assume k=1.)

Case (i.1). $\bar{\sigma}_i$ hits an edge of $\partial\Sigma - \Sigma$: In this case divide Σ in two parts one on each side of $\bar{\sigma}_i$ with σ_i belonging to both sides (as in Fig. 4) and proceed inductively on each side, then paste together the two resulting polygons along σ_i.

Case (i.2). $\bar{\sigma}_i$ hits an edge segment S_j at a point p: Split S_j in two segments S'_j and S''_j at point p. Divide the problem in two, as in Fig. 5. One contains t_i, S'_j, and every other segment of Σ on the same side. The second contains $\bar{\sigma}_i$, S''_j, and every remaining segment of Σ. Note that p is a "new" endpoint in both subproblems, the head of σ_i is not an end point in the first subproblem, and the foot of σ_i is not an endpoint in the second subproblem. Paste together the two resulting polygons along t_i.

Case (ii). We use a similar argument as in case (i.2). See Fig. 6. The induced cycle produces a convex "hole" in the middle, which will become part of the eventual circumscribing polygon, plus subproblems (that are similar to the first subproblem of case (i.2)) one for each of the segments that induce the cycle.

3. The Tournament Pseudoforest

The idea is motivated from the proof of Theorem 1. Imagine the segments in Σ play a tournament as follows: Extend each segment in Σ along its line of support from the side of its foot until it hits either $\partial\Sigma$, or the extension of another segment. If two segment extensions intersect, they play a match. The extension of the looser ends at the intersection point, while the winner continues to be extended. If the intersection point is in the relative interior of one of the segments, then that segment is declared the winner of that match, otherwise the winner is chosen arbitrarily (so the structure is not unique). If a segment extension intersects $\partial\Sigma$, it looses to $\partial\Sigma$ and its extension terminates at the intersection point. Now let $\bar{\sigma}_i$ denote the extension of σ_i, and $t_i = closure(\bar{\sigma}_i - \sigma_i)$ be called the *trail* of σ_i. The tournament pseudoforest of Σ, denoted $TP(\Sigma)$, is the plane graph that is the union of the segments and their trails. $TP(\Sigma)$ forms a convex partitioning of the convex hull of Σ. We refer to this convex planar subdivision as the tournament pseudoforest subdivision, denoted *TP–subdivision*. Each connected component of $TP(\Sigma)$ is a tree plus possibly an extra edge which creates a unique cycle. Thus the name pseudoforest. See Fig. 7.

We can construct the *TP–subdivision* in $O(n \log n)$ time by sweeping Σ with a vertical sweep line along the direction of the x-axis, where the event points are the left ends of the segments and some "tournament intersection points". There are a total of $O(n)$ event points and each contributes $O(\log n)$ to the time complexity. Some important details are omitted here. We also maintain each connected component of $TP(\Sigma)$ as a rooted pseudotree (with appropriate bidirectional pointers with the *TP–subdivision*), where the root corresponds to the cyclic list of segments forming its unique cycle, if it has one; otherwise, the root of the tree is its only segment that is either an edge segment, a diagonal segment, or a segment whose trail hits an edge of $\partial\Sigma - \Sigma$. We refer to this entire data structure collectively as the tournament pseudoforest data structure, denoted $\Psi(\Sigma)$.

4. Proof Sketch of Theorem 2

We use the data structure $\Psi(\Sigma)$ and a *careful* refinement of the proof of Theorem 1. The entire process takes $O(n)$ time, given $\Psi(\Sigma)$. We work on the pseudotrees of $TP(\Sigma)$ one by one. We omit most details here and sketch only the case corresponding to case (ii) in the proof of Theorem 1 (the most involved case): The root of the pseudotree provides the cycle, the convex "hole", and the segments that form the cycle (cycle-segments). Then, for each cycle-segment, we use depth-first-search to go down the pseudotree (akin to the idea of topological ordering of vertices around a plane tree as described in [1]) and *clip off* maximally connected trails starting from those that intersect the cycle-segment (but not its trail) or intersect the appropriate edge of $\partial\Sigma$ adjacent to the cycle-segment. In this way, we cut the problem into *many* subproblems, *all extremally situated.* See Fig. 8. Each step of the clipping process is charged to one of the trails that is clipped off, $O(1)$ time to each. But the number of such trails is in the order of (almost equal to) the number of leaves of the trees that are cut off, and that is the number of feet of segments in Σ that were internal but have become edge segments in the smaller subproblems. This is crucial in proving the linear time complexity, given $\Psi(\Sigma)$. The optimality of the algorithm is implied by an easy reduction from sorting.

References

1. AGGARWAL, A., L.J. GUIBAS, J. SAXE, AND P.W. SHORE, "A linear-time algorithm for computing the Voronoi diagram of a convex polygon," *Discrete & Computational Geometry*, vol. 4, no. 6, pp. 591-604, 1989.
2. GABOW, H.N. AND R.E. TARJAN, "A linear-time algorithm for finding a minimum spanning pseudoforest," *Inform. Process. Lett.*, vol. 27, pp. 259-263, 1988.
3. GRAHAM, R.L., "An efficient algorithm for determining the convex hull of a finite planar set," *Inform. Process. Lett.*, vol. 1, pp. 132-133, 1972.
4. MARCOTTE, O. AND S. SURI, "Fast matching for points on a polygon," in *Proc. 30th IEEE Symp. on Foundations of Computer Science*, pp. 60-65, 1989.
5. RAPPAPORT, D., "Computing simple circuits from a set of line segments is NP-complete," in *Proc. 3rd ACM Symp. on Computational Geometry*, pp. 322-330, 1987.
6. RAPPAPORT, D., H. IMAI, AND G.T. TOUSSAINT, "Computing simple circuits from a set of line segments," *Discrete & Computational Geometry*, vol. 5, no. 3, pp. 289-304, 1990.

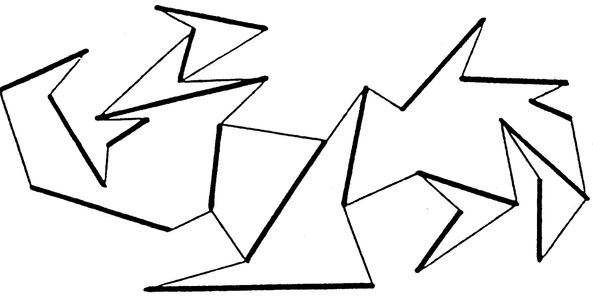

Fig. 1.

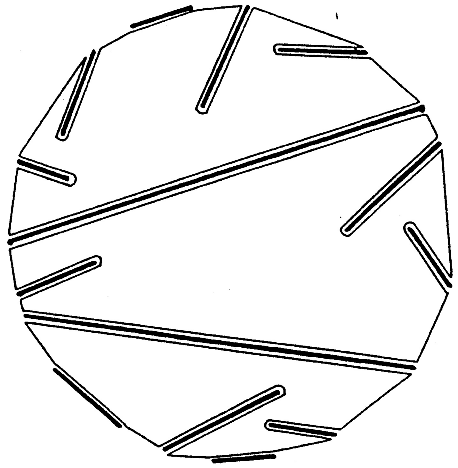

Fig. 2.

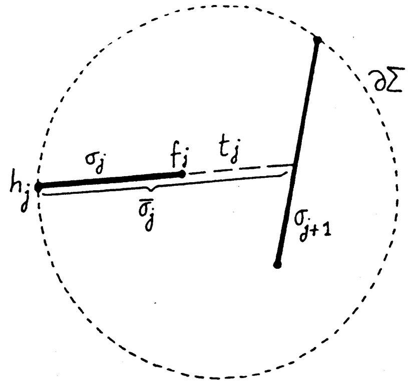

Fig. 3.

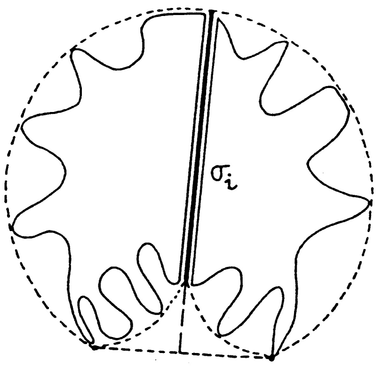

Fig. 4.

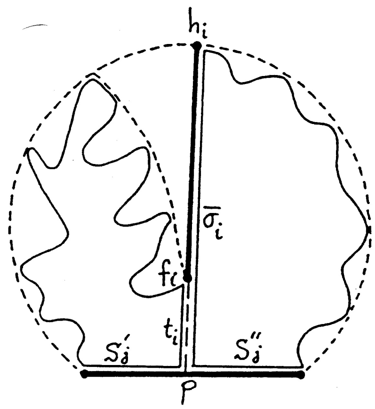

Fig. 5.

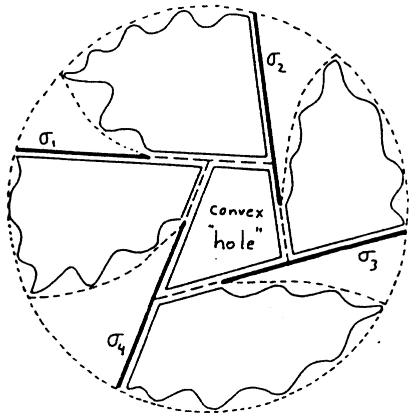

Fig. 6.

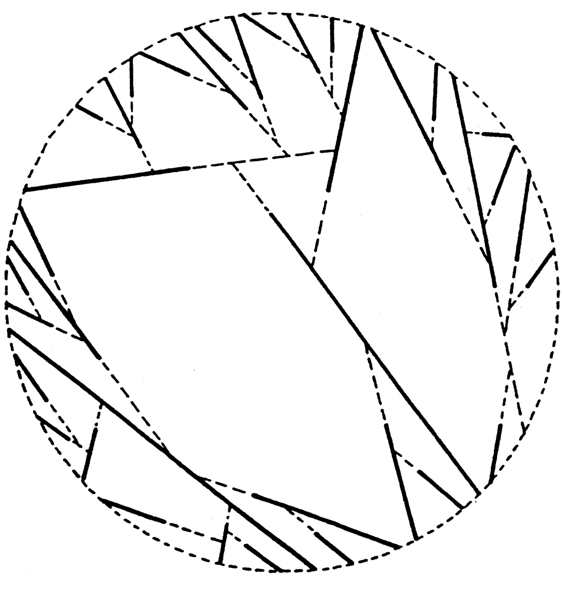

Fig. 7.

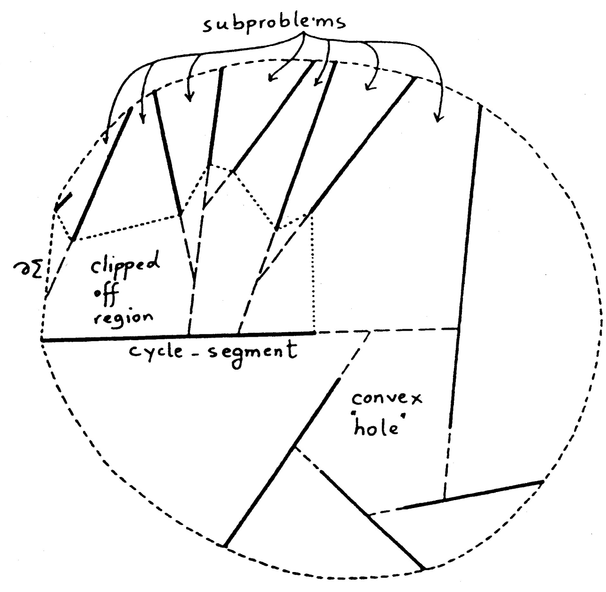

Fig. 8.