Maximizing the Area of an Axially-Symmetric Polygon Inscribed by a Simple Polygon*

Gill Barequet† Vadim Rogol‡

Abstract

In this paper we resolve the following problem: Given a simple polygon P, what is the maximum-area polygon that is axially symmetric and is contained by P? We propose an algorithm for answering this question, analyze the algorithm's complexity, and describe our implementation of it (for convex polygons). The algorithm is based on building and investigating a planar map, each cell of which corresponds to a different configuration of the inscribed polygon. We prove that the complexity of the map is $O(n^4)$, where n is the complexity of P. For a convex polygon the complexity, in the worst case, is $\Theta(n^3)$.

1 Introduction

Containment problems have always held an important role in discrete and computational geometry. In general, a containment problem means that we are given some object and then are asked to compute a containing or contained object satisfying some additional criteria (see, e.g., [2, 1]), as it relates to the first object. In the current work we seek a maximum-area polygon contained by a given simple (or convex) polygon, with the restriction that the inscribed polygon is axially symmetric. The main motivation for this problem is industrial, originating from the need to calculate shapes that are to be cut from metal and cloth sheets.

Definition 1 *Given a simple polygon P in the plane, another polygon $C_{P,\ell}$ is called* symmetric contained *in P if (1) $C_{P,\ell} \subset P$; and (2) $C_{P,\ell}$ is symmetric about some line (axis) ℓ.*

Among all symmetric polygons contained in P with some axis of symmetry ℓ, the one with the maximum area is the intersection of P and its reflection with respect to ℓ, denoted as P_ℓ. The axially-symmetric polygon contained in a simple (nonconvex) polygon may consist of several disconnected components. In the case of a convex polygon, the axially-symmetric contained polygon is always connected and convex. The boundary of $P \bigcap P_\ell$ consists of portions of edges of P and of P_ℓ. Such a polygon is hereafter called a *symmetrically inscribed* polygon and denoted by $I_{P,\ell}$. The *order* of edges of P whose portions are the edges of $I_{P,\ell}$ is referred to as the *configuration* of $I_{P,\ell}$. Since every pair of symmetric edges of $I_{P,\ell}$ is contributed by some edge of P and by its reflection P_ℓ, the configuration of the two halves of the boundary of $I_{P,\ell}$ (delimited by ℓ) are identical with respect to the edge identities, and have the opposite "origin" for each edge (P or P_ℓ). Thus, the problem we actually solve is:

Problem 1 *Given a simple polygon P, find the axis ℓ^{opt} whose respective symmetrically inscribed polygon I_P^{opt} is of maximum area.*

2 The Map of Axes

The number of possible configurations is restricted by the number of intersections of edges of the original polygon and its mirrored version. We consider all possible configurations of the inscribed polygon; then for each configuration we find the polygon of maximum area, and finally choose the largest-area polygon. The problem is thus split into two subproblems:

Problem 1 *Given a simple polygon, find all the possible configurations of its inscribed polygons.*

Problem 2 *Given a configuration C of an inscribed polygon, find the instance of C with maximum area.*

An inscribed polygon is determined by the axis ℓ, s.t. a change of ℓ causes a change in the inscribed polygon. Only more rarely will a small movement of ℓ cause a change in the configuration. Thus, every legal configuration corresponds to a set of axes. To alleviate the consideration of the sets of lines, we use a *duality transform* that maps lines of the form $\ell : y = kx + b$ in the primal plane (XY) into points $\ell^*(k, b)$ in the dual plane. Thus, the sets of legal axes induce a subdivision of the dual plane. The faces in this planar map correspond to configurations of the inscribed polygons.

3 Geometric Description

To distinguish between edges of the original polygon and the edges of the map in the dual plane, we will refer to the latter as "arcs."

While we move in the dual plane, crossing an arc means a change in the combinatorial structure of the inscribed polygon. There are two basic types of such changes: 1. A new edge emerges in the boundary of $I_{P,\ell}$ between two existing edges; and 2. An edge disappears from the boundary. Both events are invertible, and in fact, represent two aspects of the same event. (Fig. 1).

Let us analyze the structure of $I_{P,\ell}$. By definition, its boundary is the union of two symmetrical chains: one containing edges of P clipped by P_ℓ, and the other containing edges of P_ℓ clipped by P. A new edge emerges (resp., vanishes) between two existing edges of $I_{P,\ell}$'s boundary only when some portion of an edge of P_ℓ (or P) becomes (resp., ceases to be) clipped by P (or P_ℓ). In other words, the major

*Work on this paper has been supported in part by Grant 01-06-00032 under the Korean-Israeli Research Cooperation and the European FP6 Network of Excellence Grant 506766 (AIM@SHAPE).

†Center for Graphics and Geometric Computing, Dept. of Computer Science, The Technion—Israel Institute of Technology, Haifa 32000, Israel, barequet@cs.technion.ac.il

‡Dept. of Electrical Engineering, The Technion—IIT, Haifa 32000, Israel, rogolv@tx.technion.ac.il

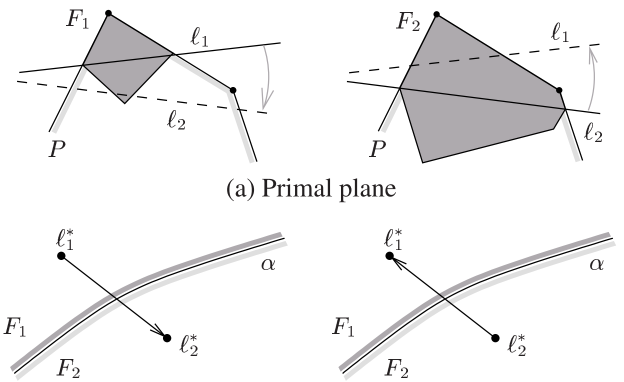

(a) Primal plane

(b) Dual plane

Figure 1: Change of the combinatorial structure: ℓ_1 and ℓ_2 are two lines, and ℓ_1^* and ℓ_2^* are their respective dual points. F_1 and F_2 are the respective configurations of ℓ_1 and ℓ_2; the dual-plane view shows the arc α that separates the faces of F_1 and F_2.

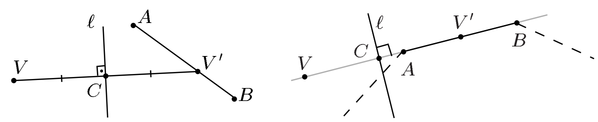

(a) V, A, B noncollinear (b) V, A, B collinear

Figure 2: A reflection of the vertex V on the edge AB

events occur in conjunction with a change of some polygon-edge clipping. A clipping configuration is combinatorially altered when a clipped edge changes its position with respect to an edge of the clipping polygon. Actually we are interested only in the touching events, in which an endpoint of a clipped edge lies on the clipping edge. The moment of touching corresponds to the appearance (or disappearance) of an edge in the boundary of the inscribed polygon. Thus, arcs of the map in the dual plane correspond to such axis positions, where edges of P touch edges of P_ℓ (and vice versa), or, simply, when vertices of P lie on edges of P_ℓ. This leads to the following question, which we must now answer:

Problem 1 *Given a vertex v of the polygon P, find the family of axes reflecting v on edges of P.*

Let AB be an edge of P, s.t. the axis $\ell : y = kx + b$ reflects a vertex V of P to some point V' that lies on AB (Fig. 2(a)). Assume first that $V \notin AB$ and that A, B, V are noncollinear. Obviously, ℓ passes through the midpoint C of the segment VV' and is perpendicular to VV'. Using these facts we can calculate the parameters of ℓ: $k = -\frac{1}{\mathrm{slope}(VV')} = -\frac{V'_x - V_x}{V'_y - V_y}$ and $b = C_y - kC_x = \frac{1}{2}\Big(V_y + V'_y - (V_x + V'_x)k\Big)$. We obtain

$$b(k) = \frac{N_1 k^2 + N_2 k + N_3}{D_1 k + D_2},\ k \neq -\frac{B_x - A_x}{B_y - A_y}, \qquad (1)$$

where $N_1 = -((V_x + A_x)(B_y - A_y) + (B_x - A_x)(V_y - A_y)$, $N_2 = 2((B_y - A_y)V_y - (B_x - A_x)V_x)$, $N_3 = (V_y + A_y)(B_x - A_x) + (B_y - A_y)(V_x - A_x)$, $D_1 = 2(B_y - A_y)$, and $D_2 = 2(B_x - A_x)$. The domain of $b(k)$ is determined by the positions of V, A, B, and may consist of one or two closed intervals. The arcs described by Eq. (1) are hereafter referred to as arcs of *type I*.

The case in which A, B, V are collinear (Fig. 2(b)) requires special treatment. The axis ℓ is perpendicular to AB, and consequently, the line $\ell : y = kx + b$ has a constant slope. The reflected point V' sweeps along AB between the segment endpoints, so the arc in the dual plane subdivision is a vertical line segment with the parameters $k = -\frac{B_x - A_x}{B_y - A_y}$, $A_y \neq B_y$ and $b \in \Big[A_y + \frac{B_x - A_x}{B_y - A_y}A_x,\ B_y + \frac{B_x - A_x}{B_y - A_y}B_x\Big]$. We hereafter refer to such edges as arcs of *type II*. If AB is horizontal (i.e., $A_y = B_y$), then the slope of ℓ is infinite. In such a case we virtually move the arc to infinity.

Let us now handle the case in which V is either A or B (assume w.l.o.g. A). Here any axis passing through $V = A$ maps it to itself. The general equation of these axes (in the dual plane) is $b = -A_x k + A_y$. It may be related to two different arc types. An axis passing through a vertex may cross the polygon (an arc of *type III*), or it may support the polygon from the outside (*type IV*). In the latter case the inscribed polygon degenerates to a single point (the supported vertex). In the dual plane we consider the arcs of type IV as the "boundaries" of the map, beyond which the inscribed polygon is empty.

Both arc types (III and IV) are described by the same formula of a straight line. If we rotate the axis around the vertex, both types meet when the axis passes through one of the polygon edges incident to the vertex. Fig. 3 illustrates this situation. In the primal plane, the axis passes through either e_1 or e_2, which are two polygon edges that share the vertex A. In the dual plane, these axis positions are two points e_1^* and e_2^* lying on the line A^*. Types III and IV alternate at e_1^* and e_2^* as we go along A^*.

Fig. 4(a) shows a convex polygon, while Fig. 4(b) shows the planar subdivision induced by this polygon in the dual

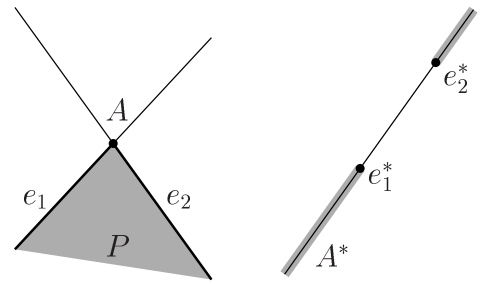

(a) Primal plane (b) Dual plane

Figure 3: An axis passing through a vertex

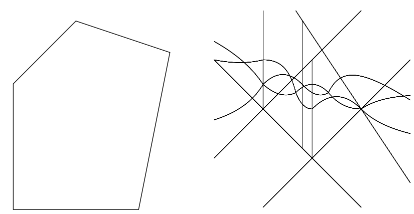

(a) Polygon (b) Map

Figure 4: A sample convex polygon and its corresponding planar map in the dual plane

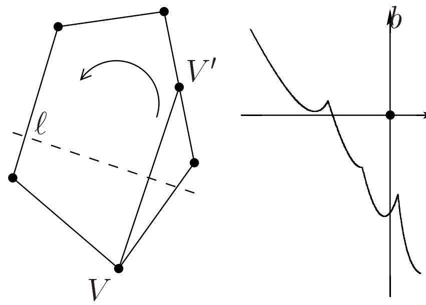

(a) Primal plane (b) Dual plane

Figure 5: Sliding V' (the reflection of V) along P

plane.

4 Map Complexity

The planar subdivision in the dual plane is induced by a set of m Jordan arcs, any pair of which intersects in at most a constant number of points (see below). Therefore, the combinatorial complexity of the arrangement that they form is $O(m^2)$. First, we need to show that any two arcs intersect a constant number of times. Second, we will show that for a convex polygon the complexity of the map can be considerably improved.

The number of intersection points of two arcs is indeed a small constant. Consider first two arcs of type I. By comparing two terms of the form as in Eq. (1), we obtain a cubic equation that has at most three real solutions. Intersections with other arc types are even simpler. In all cases we get equations of degree at most 2, which have at most two solutions. The number of arcs m is quadratic in n (the complexity of P): each vertex of P generates at most $n-1$ arcs of type I, two arcs of type II extending one another and resembling one segment, and a few arcs (up to six) of types III and IV, which are collinear and are considered as a single unbounded arc. In total there are $n+1$ arcs per vertex and $n(n+1)$ for all vertices, thus $m = \Theta(n^2)$ and the subdivision complexity is $O(m^2) = O(n^4)$. Note that this is true for *any* simple polygon P.

However, for a convex polygon the map complexity is $O(n^3)$, which is attainable in the worst case. To compute the complexity of the arrangement we will count its vertices, which is sufficient since this is a planar map.

Intersections of arcs of type I. A trivial bound is $O(n^4)$. Instead of isolated arcs we will consider *chains* of arcs generated in the dual plane by continuously sweeping a mirrored vertex along the boundary of P (except on the two edges incident to it). That is, a chain is a concatenation of all the type-I arcs of the same vertex (Fig. 5). When VV' is horizontal, the chain splits into two. The convexity of P ensures that a chain is k-monotone. Consider two such chains. Each chain consists of $n-2$ arcs, thus it contains $n-1$ arc-transition points. Hence, for any two chains we have $2n-1$ k-intervals in which each chain is represented by a single arc. The two arcs can intersect at most three times for a total of at most $3(2n-1)$ intersections between chains. Since there are $n(n-1)/2$ pairs of chains, we have in total $3n(n-1)(2n-1)/2 = O(n^3)$ intersections.

Arcs of types I and II. Due to the k-monotonicity of the chains of type I, any vertical segment (arc of type II) can

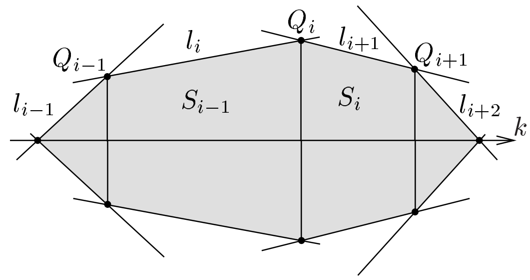

Figure 6: Calculating the area of $I_{P,\ell}$

intersect it at most once. There are n chains and n vertical segments, giving a total of $O(n^2)$ intersections.

Arcs of types I and $\{$III,IV$\}$. Arcs of types III and IV can be handled as a special case of a chain that intersects a regular chain in $O(n)$ points. There are $\Theta(n)$ such special chains and $\Theta(n)$ regular chains of type I, yielding $O(n^3)$ intersection points.

Arcs of types II, III, and IV. All these together are $2n$ straight lines or line segments that intersect in at most $O(n^2)$ points.

The total number of arc intersection points (and hence the total map complexity) is $O(n^3)$. In the full version of the paper we give a matching lower bound in the worst case by showing that there exists an n-gon whose respective map in the dual plane has complexity $\Theta(n^3)$. In conclusion, we have:

Theorem 1 *For a convex polygon the complexity of the planar subdivision in the dual plane is $\Theta(n^3)$ in the worst case.*

5 Maximizing the Area of the Inscribed Polygon

5.1 Area Function

We assume that the configuration of $I_{P,\ell}$ is given as a sequence of edges of P: $\{l_1, l_2, ..., l_m\}$ $(m \le 2n)$. The edges of $I_{P,\ell}$ are represented by lines l_i, for each of which we store a triple (k_i, b_i, μ_i), where k_i and b_i are line coefficients and μ_i equals 1 (resp., -1) if the edge belongs to P (resp., its reflection).

We rotate the plane so as to make ℓ X-parallel (Fig. 6). The vertices Q_i of $I_{P,\ell}$ can be determined as intersections of pairs of neighboring lines l_i and l_{i+1}. Simple analytic geometry yields the area function

$$S(b,k) = \sum_{i=1}^{m} \Big(2\frac{b_i b_{i+1}}{k_{i+1}-k_i} + b_i^2 \frac{k_{i-1}-k_{i+1}}{(k_{i-1}-k_i)(k_i-k_{i+1})} \Big),$$

with the convention $k_0 = b_0 = k_{m+1} = b_{m+1} = 0$. The area function for a simple polygon is identical. The only difference is that the inscribed symmetric polygon may contain several disconnected components.

Fig. 7 plots the area of the inscribed polygon (as a function of b and k) for a square.

5.2 Maximizing the Area Function

Finally we find a global maximum of $S(k,b)$ within each cell of the map in the dual plane. This optimization problem

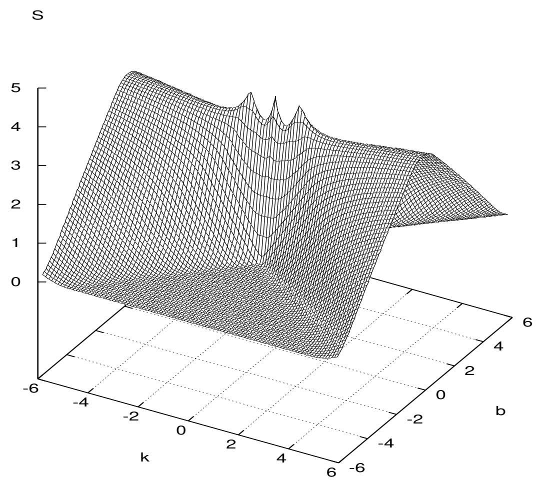

Figure 7: The area function $S(k,b)$

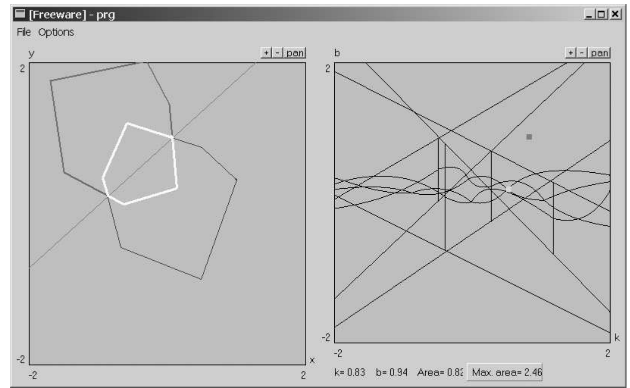

Figure 8: Screen snapshot of the system

is analytically intractable [4] because of the complexity of the objective function, hence we resort to numerical methods. Formally, we need to maximize a low-dimensional rational polynomial with many terms. The objective function is unconstrained (i.e., the solution does not have to fulfill any other constraints) but the optimum is sought within a bounded region.

A description of a host of algorithms for a global optimization and available implementations is found in [3]; we provide more details in the full version of this paper. We implemented a simple method that works well in practice. We evaluate $S(k,b)$ in regularly-scattered points within the current cell and choose the best point (w.r.t. the objective function) as the first approximation of the optimum. Then we iteratively resample the function at the vicinity of the current optimum and vary the sampling resolution, combining steepest-descent and simulated-annealing heuristics. We stop when no sufficiently-improving point is found any more.

6 Running-Time Analysis and Implementation

Our algorithm computes the maximum-area axis-symmetric polygon inscribed by another polygon P using the following steps: 1. For each vertex $V_i \in P$ compute the arcs of the planar subdivision M in the dual plane. Construct M. 2. For each face of M compute the associated area function (of symmetrically inscribed polygons) and find its maximum (within the face). 3. Report the global maximum as the answer.

If P is convex, the combinatorial complexity of M is $\Theta(n^3)$ in the worst case, where n is the complexity of P. Constructing M can be done by a plane-sweep procedure whose running time is $O(n^3 \log n)$. In each face we need to compute the area function and find its maximum. Computing the area function of the first face takes $O(n)$ time. Updating the area function while moving from a face to a neighboring face can be done in constant time by adding and subtracting only a few terms. Thus, the amount of time needed for computing all the area functions is proportional to the number of faces, that is, $O(n^3)$. Maximizing the area function within a face is done by a numerical method. In theory the optimization problem is intractable. In practice the running time of the “black box” that solves the optimization problem depends linearly on the number of terms in the objective function (n, in our case), linearly on the complexity of the cell's boundary, and on the convergence parameter, to which we refer as a constant. On average the complexity of a single cell is constant,[1] for a total of $O(n^3)$ for all the cells. We denote by $T(n)$ the average time complexity of the optimization step in a single cell; in practice $T(n) = O(n)$. In total, the running time of the algorithm is $O(n^3(\log n + T(n)))$. For a simple polygon the running time is $O(n^4(\log n + T(n)))$.

We implemented the entire algorithm for convex inscribing polygons. The software was written in C++ under the Windows operating system. It consists of about 6,500 lines of code, and it also uses the geometric package CGAL, the GUI toolkit Qt, and an Open Inventor compatible toolkit Coin3D. Our system offers an interactive tool that visualizes the objects, concepts, and relations presented in the paper. Fig. 8 shows a screen snapshot of our system.

References

[1] J.E. Goodman and J. O'Rourke, *Handbook of Discrete and Computational Geometry*, CRC Press LLC, 1997.

[2] V. Gritzmann and P. Klee, On the complexity of some basic problems in computational convexity: I. Containment problems, *Discrete Mathematics*, 136 (1994), 129–174.

[3] J.D. Pintèr, Continuous global optimization software: A brief review, *Optima*, 52 (1996), 1–8.

[4] A.H.G. Rinnooy and G.T. Timmer, New Methods in Optimization and their Industrial Uses, Birkhäuser Verlag, Basel, 1989.

[1]However, the complexity of a single cell can be superlinear in the worst case.