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

Gill Barequet* Vadim Rogol†

Abstract

In this paper we solve the following optimization 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 solving this problem, analyze its complexity, and describe our implementation of it (for the case of a convex polygon). 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 is $\Theta(n^3)$ in the worst case.

1 Introduction

Containment problems have always held an important role in discrete and computational geometry, as well as in its practical applicative domains. In general, we are given some object and aim to compute a containing or contained object satisfying some additional criteria (see, e.g., [5, §27.5, p. 507] and [6]). In the current work we seek a maximum-area polygon contained by a given simple polygon, with the restriction that the inscribed polygon is axially symmetric. The main motivation for this problem is industrial, originating from cutting shapes from metal and cloth sheets. (The 3D version of the problem is relevant for cutting a diamond from a raw stone.) As we show below, this problem can be reduced to the problem of maximizing the area of overlap between two polygons, one of them is the input polygon and the other is its mirror reflection. From this perspective we should note that a variety of similar problems were studied in the literature [1, 13] with efficient algorithms offered to solve them. In these works the mutual placement of the polygons is restricted much more than with our axial-symmetry constraint. Tuzikov and Heijmans [12] studied the somewhat similar problem of representing a polygon as the Minkowski sum of two polygons P_s and P_a, where P_s is the maximum-area polygon symmetric about some input direction. Böhringer et al. [2] look for area bisectors of polygons. This problem is different from ours, but the presented solution uses similar tools: computing an arrangement of curved segments in the dual plane, induced by the original polygon, and processing that arrangement for solving the problem in hand.

In this paper we present an algorithm for computing the maximum-area axially-symmetric polygon contained in a given simple polygon. In a sense, it follows a host of algorithms that make use of the idea of configuration space of placements of polygons. Although the paper deals mainly with convex polygons, the algorithm can also be applied to general simple polygons.

This paper is organized as follows. In Section 2 we provide some assumptions and definitions. In Section 3 we investigate the representation of the map of symmetry axes in the dual plane, providing a full analysis of its structure and complexity. In Section 5 we analyze the complexity of the map. Section 6 develops the polygon-area function associated with each face of the map. In Section 7 we analyze the running time of the algorithm and describe our implementation of it.

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

†Dept. of Electrical Engineering, The Technion—Israel Institute of Technology, Haifa 32000, Israel. E-mail: rogolv@tx.technion.ac.il

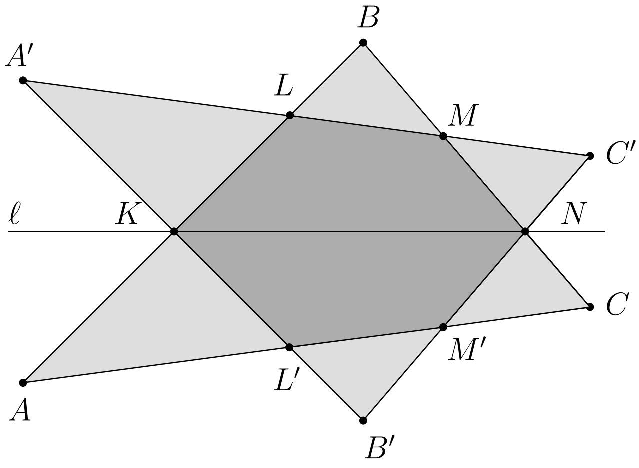

Figure 1: Maximum-area symmetric contained polygon

2 Assumptions and Definitions

Throughout this paper we refer by the term "polygon" to the *interior* of the polygon. When we mean the boundary of a polygon, we will mention this explicitly.

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, we wish to find the one with maximum area. This is easily done by computing the intersection of P and its reflection with respect to ℓ, denoted as P_ℓ. It is easily verified that this is the maximum-area symmetric polygon contained in P and symmetric about ℓ. First, it satisfies the two requirements of Definition 1. Second, the area of $P \bigcap P_\ell$ is larger than any other polygon contained in P and symmetric about ℓ. Suppose for contradiction that there exists another polygon $Q \neq P \bigcap P_\ell$ such that Area(Q) > Area($P \bigcap P_\ell$). By definition, every point in $Q \setminus (P \bigcap P_\ell)$ is outside P and/or P_ℓ, so at least one of the requirements is violated. Therefore the area of Q cannot exceed the area of $P \bigcap P_\ell$, which is a contradiction. 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.

Note that the boundary of $P \bigcap P_\ell$ consists of portions of edges of P and of P_ℓ. See Figure 1 for an example. The polygon $KLMNM'L'$ has the maximum area among all polygons contained in the triangle ABC and symmetric about ℓ. Such a polygon is hereafter called *symmetric inscribed* polygon and denoted by $I_{P,\ell}$.

Furthermore, the *order* of edges of P whose portions are the edges of $I_{P,\ell}$ is referred to as the *configuration* of $I_{P,\ell}$.[1] 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 (whether it is P or P_ℓ). Hence we may apply the term *configuration* to only one half of the boundary of $I_{P,\ell}$.

Thus the problem we actually need to solve is the following:

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

[1] Where appropriate, we also apply this term to a chain of segments, part of the boundary of the polygon.

3 The Map of Axes

3.1 Outline

The number of possible axis positions is obviously infinite, but the number of possible configurations is restricted combinatorially by the number of intersections between edges of the original polygon and its mirrored version. The general idea of our solution is to consider all possible configurations of the inscribed polygon, then for each configuration to find the polygon of maximum area, and finally to choose out of all these candidates the largest-area polygon. Thus we may split the problem into two subproblems:

Problem 2 *Given a simple polygon P, find all the possible configurations of its inscribed polygons.*

Problem 3 *Given a configuration of an inscribed polygon, find the instance of that configuration with maximum area.*

As noted above, an inscribed polygon is determined by the axis ℓ, such that a change of ℓ in general causes a change in the inscribed polygon. Only more rarely a small movement of ℓ causes a change in the configuration. Thus, every legal configuration corresponds to a set of axes. These sets are certainly disjoint and some of them have common boundaries, since the space of symmetry axes is continuous. 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 (each set corresponding to the same configuration), induce a subdivision (or map) of the dual plane. The faces in this planar map correspond to configurations of the inscribed polygons. In the sequel we investigate the edges that bound the faces of this map.

4 Geometric Description

To distinguish between edges of the original polygon (and its inscribed polygons) and the edges of the map in the dual plane, we will refer to the latter by the term "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 mutually invertible and in fact represent two aspects of the same event, while appearing or disappearing of an edge on the boundary of the inscribed polygon is determined by the direction in which we move in the dual plane, that is, the direction in which the axis moves in the primal plane (see Figure 2 for an illustration).

Let us now analyze the structure of an inscribed polygon $I_{P,\ell}$. By definition, its boundary is the union of two symmetric 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) clipped by P (or P_ℓ). In other words, the major events occur with a change of some polygon-edge clipping. (As mentioned above, the roles of P and P_ℓ are symmetric.) A clipping configuration is combinatorially altered when a clipped edge changes its position with respect to an edge of the clipping polygon; here we mean only combinatorial changes of the situations "intersection," "no intersection," and "touching." Actually we are interested only in the touching events, in which an endpoint of a clipped edge lies on the clipping edge.[2] 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, in simpler words, when vertices of P lie on edges of P_ℓ. Let us reformulate this crucial observation and refine the problem in question:

Problem 4 *Given a vertex v of the polygon P, find the family of axes reflecting v on edges of P.*

[2]Due to the inherent symmetry between P and P_ℓ, both touching edges are clipping and are clipped by each other at the same time.

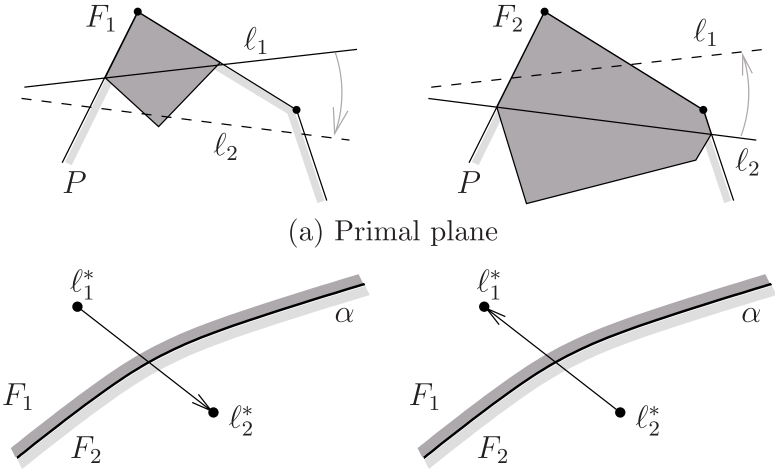

(a) Primal plane

(b) Dual plane

Figure 2: 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 between the faces of F_1 and F_2

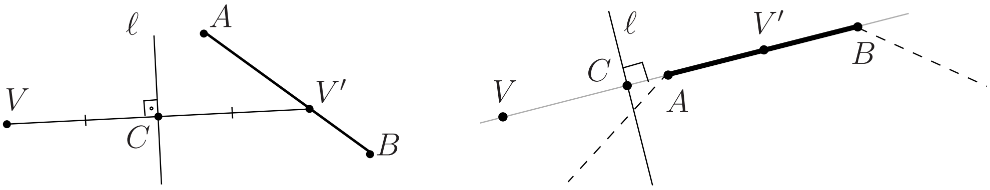

(a)

(b)

Figure 3: A reflection of the vertex V on the edge AB: (a) the vertices V, A and B are noncollinear, (b) the vertices V, A and B are collinear

Let AB be an edge of P, such that the axis $\ell : y = kx+b$ reflects a vertex V of P to some point V' that lies on AB (see Figure 3(a)). Assume for now that $V \notin AB$ and that A, B, V are noncollinear. (We will treat the special cases later.) Obviously, the axis of symmetry ℓ passes through the midpoint C of the segment VV', and it is perpendicular to VV'. Using these facts we can use elementary geometry and calculate the parameters of ℓ:

$$k = -\frac{1}{\text{slope}(VV')} = -\frac{V'_x - V_x}{V'_y - V_y}, \qquad b = C_y - kC_x = \frac{1}{2}\Big(V_y + V'_y - (V_x + V'_x)k\Big). \tag{1}$$

Omitting intermediate calculations we express:

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

where

$$\begin{aligned}
N_1 &= -((V_x + A_x)(B_y - A_y) + (B_x - A_x)(V_y - A_y),\\
N_2 &= 2((B_y - A_y)V_y - (B_x - A_x)V_x),\\
N_3 &= (V_y + A_y)(B_x - A_x) + (B_y - A_y)(V_x - A_x),\\
D_1 &= 2(B_y - A_y), \quad \text{and}\\
D_2 &= 2(B_x - A_x).
\end{aligned}$$

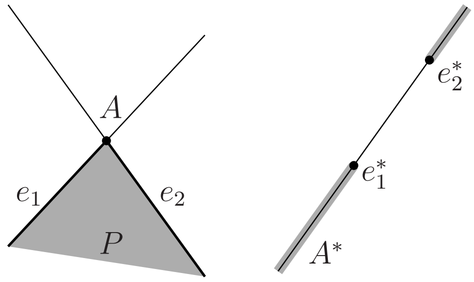

(a) Primal plane (b) Dual plane

Figure 4: An axis passing through a vertex of the polygon

The domain of $b(k)$ is determined by the positions of V, A, and B, and may consist of one or two closed intervals. The arcs described by Equation (2) are hereafter referred to as arcs of *type I*.

The case in which A, B, V are collinear (see Figure 3(b)) requires a 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 we easily see that the arc in the dual plane subdivision is a vertical line segment with the parameters:

$$k = -\frac{B_x - A_x}{B_y - A_y}, \; A_y \neq B_y, \tag{3}$$

$$b \in \left[A_y + \frac{B_x - A_x}{B_y - A_y} A_x, \quad B_y + \frac{B_x - A_x}{B_y - A_y} B_x\right]. \tag{4}$$

We hereafter refer to such edges as arcs of *type II*. If AB is horizontal (that is, if $A_y = B_y$), then the slope of ℓ is infinite. In such case we virtually move the arc to infinity. (Or rotate a bit the input polygon, or use homogeneous coordinates.)

Let us now handle the case in which the vertex V coincides with either point A or B. (Assume without loss of generality that V coincides with A.) Then any axis passing through $V = A$ maps it to itself. The general equation of these axes (in the dual plane) is $b = -A_x k + A_y$. This line has a special role in the planar subdivision in the dual plane. Actually it may be related to two different arc types, which are not obvious from the previous observations. An axis passing through a vertex may cross the polygon (which we call an arc of *type III*), or it may just support the polygon from the outside (*type IV*). In the latter case the inscribed polygon degenerates to a single point (the supported vertex). Slight movements of the axis (that will make it not passing through the vertex) will either cause the vanishing of the inscribed polygon, or will turn it into a regular nondegenerate polygon. In the dual plane we thus consider the arcs of type IV as the "boundaries" of the planar map, beyond which the inscribed polygon is empty.

Despite their essential distinction, 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. Figure 4 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^*. As can be seen in the figure, types III and IV alternate at e_1^* and e_2^* as we go along A^*.

So far we have examined all the possible cases of axes reflecting a vertex of the polygon on all edges of the polygon, thus we have the tool for solving Problem 4. Figure 5(a) shows a convex polygon, while Figure 5(b) shows the planar subdivision induced by this polygon in the dual plane. [4, §5.2.2, pp. 122–124].

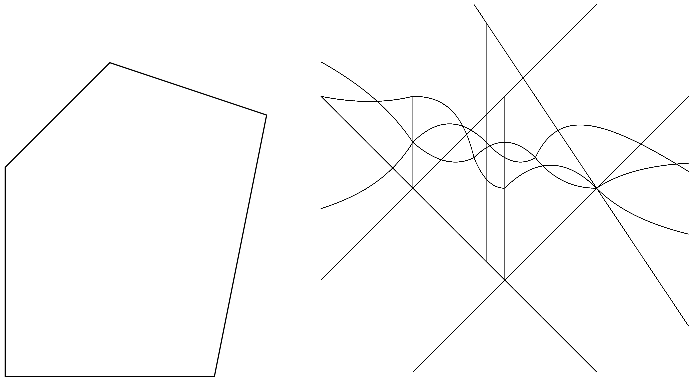

(a) Polygon (b) Map

Figure 5: A sample convex polygon and its corresponding planar map in the dual plane

5 Map Complexity

The planar subdivision in the dual plane is induced by a set of Jordan arcs, any pair of which intersect in a constant number of points (to be specified below). Therefore, the combinatorial complexity of the arrangement that they form is $O(m^2)$ [5, §21.1.2, pp. 391–392] , where m is the total number of arcs. 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 Equation (2), 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 the underlying polygon P): each vertex of P generates at most $n-1$ arcs of type I, two arcs of type II extending one another and looking like 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. We will show below that for a convex polygon the complexity of the map is only $\Theta(n^3)$, which is a tight bound in the worst case.

To compute the complexity of the arrangement of arcs we will count its vertices. This is sufficient since this is a planar map. To this aim we will perform a case analysis of intersections of arcs of all types.

Intersections of arcs of type I. As mentioned above, a straightforward counting will give us a trivial bound of $O(n^4)$. 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). In other words, a chain is a concatenation of all the type-I arcs of the same vertex. See Figure 6 for an example. Note that when VV' is horizontal, the chain will split into two chains. Asymptotically this does not change the complexity of the map. The convexity of P ensures that a chain is k-monotone. Consider two chains made of arcs of type I. Each such chain consists of $n-2$ arcs,[3] 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. As mentioned above, two such arcs can intersect at most three times, yielding a total of at most $3(2n-1)$ intersections between chains. Since there are $n(n-1)/2$ pairs of chains, we obtain the upper bound $3n(n-1)(2n-1)/2 = O(n^3)$ on the total number of intersections.

[3] Actually, it may consists of $n-1$ curves, since one curve may split into two. This occurs when the vertex is reflected onto nonincident edge by a vertical symmetry axis.

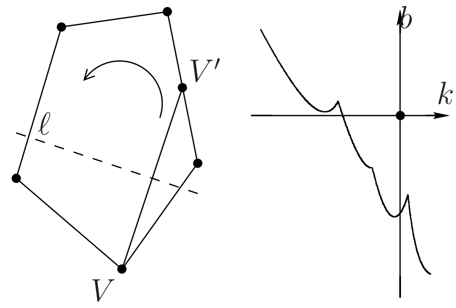

Figure 6: Sliding V' (the reflection of V) along the boundary of P

Intersections of 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 intersect it at most once. There are n such chains and n such vertical segments, giving a total of $O(n^2)$ intersections.
Intersection of arcs of types I and {III,IV}. The 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.
Intersections of arcs of types II, III, and IV. All these together are $2n$ straight lines or line segments which intersect in at most $O(n^2)$ points.

Summing up, we obtain that the total number of arc intersection points (and hence the total complexity of the map) is $O(n^3)$.

We will now give a matching lower bound in the worst case. Namely, we will show that there exists an n-gon whose respective map in the dual plane has complexity $\Theta(n^3)$. The following observation will set the ground for our example.

Observation 1 *Let the point A be mirrored into the point A' by the line ℓ. The points A and A' lie on the same circle centered at the origin if and only if ℓ passes through the origin.*

Consider a regular polygon P centered at the origin, and assume without loss of generality that it has an even number of vertices. Figure 7(a) shows such a polygon inscribed by a circle. One of the vertices of P, V_i, is reflected by ℓ to another vertex V_i' of P. Recall that moving ℓ so as to sweep the reflection of V_i along P generates a chain of arcs of type I in the dual plane. Observe the case in which V_i' coincides with another vertex V_j of P ($j \neq i$). This is equivalent to the coincidence of V_j' and V_i. According to Observation 1, the reflection axis must pass through the origin, that is, it has the form $y = k_0x$. Now let us slightly change the slope of the axis and make it $\ell_\Delta : y = (k_0 + \Delta k)x$. The axis will slightly turn around the origin, so that the points V_i' and V_j' will stay on the circle circumscribing P (see Figure 7(b)). Indeed, the new line ℓ_Δ is no longer a valid axis. In order to fix that (by returning V_i' to the boundary of P) without changing the slope of the axis, we need to translate ℓ_Δ towards V_i. Now the axis ℓ_{Δ_i} is valid but it does not pass through the origin. Now do the same for V_j and V_j'. This time we translate ℓ_Δ towards V_j and obtain a valid axis ℓ_{Δ_j}, see Figure 7(c). Since the vertices V_i, V_j and their proximities lie on different sides of ℓ_Δ, its two translations ℓ_{Δ_i} and ℓ_{Δ_j} are also on its two sides. The b-value of ℓ_Δ is 0. Consequently, at $k = k_0 + \Delta k$ the values of $b(k)$ have different signs and the respective chains of V_i and V_j (in a sufficiently small neighborhood of $k = k_0$) touch, but do not intersect (forming an "X"-like configuration, see Figure 7(d)). This fact is crucial for the sequel.

Observation 2 *P is a regular polygon with an even number n of vertices. If a line ℓ reflects a vertex V_i of P to another vertex V_j, then ℓ is an axis of symmetry of P, that is, every vertex V_k of P is reflected to another vertex V_l ($l \neq k$). Furthermore, there are n such axes of symmetry ($n/2$ of which passing through $n/2$ vertices of P, and $n/2$ of which passing through n midpoints).*

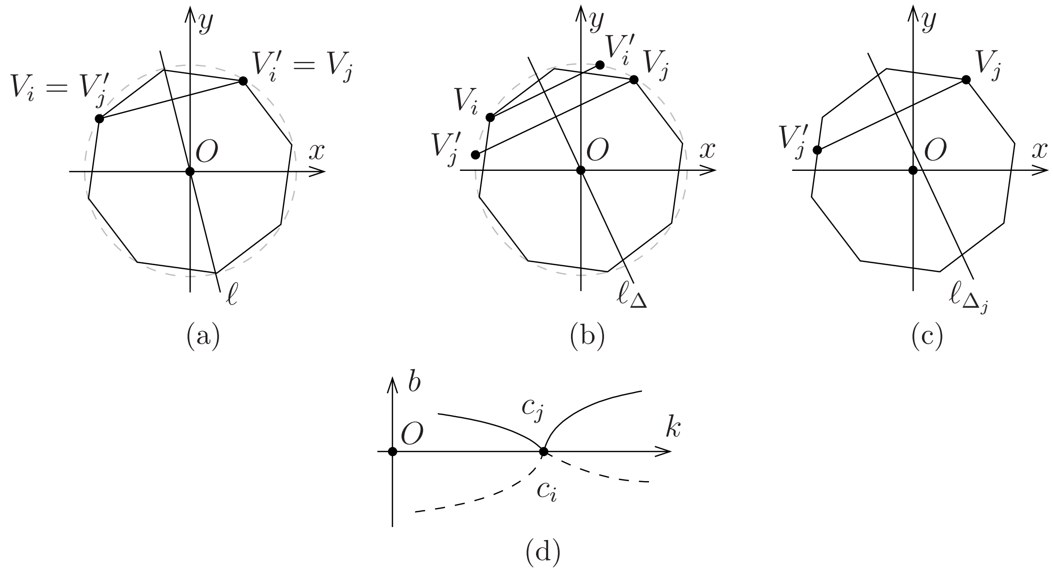

Figure 7: Touching chains

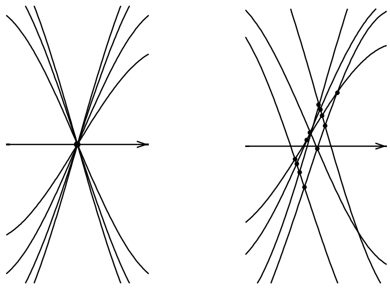

(a) Before (b) After

Figure 8: Arcs at the vicinity of a touching point: the effect of displacing polygon vertices

The symmetry axes subject of Observation 2 are touching points of chains of type I in the dual plane. There are exactly n such touching points, one of which can lie in infinity (representing a vertical axis of symmetry). In addition, $\Theta(n)$ chains pass through each of these touching points. More precisely, $n-2$ (respectively, $n-4$) chains meet at touching points that correspond to symmetry axes that pass through antipodal vertices (respectively, midpoints of antipodal edges) of the underlying polygon. In fact, arcs of other types pass through the meeting points as well, completing the number of meeting arcs to n.

Our final goal is to show that we can perturb the original vertices of P in such a way that will make each of the n touching points (incident to n chains) split in a nondegenerate manner into $\Theta(n^2)$ intersection points. When we slightly displace each vertex of the polygon, all the proper intersections are maintained at the vicinity of the original intersection points (see Figure 8(b)). We thus perturb one vertex at a time. While handling the ith vertex, there are infinitely-many possible ways to perturb it, while only a finite number of moves will modify its chain (in the dual plane) so that its intersections with the first $i-1$ chains will overlap the intersections among these $i-1$ chains. Consequently, we can perturb the original vertices of the polygon such that each of the n touching points in the dual map (with n chains passing through it) is broken into $\binom{n}{2}$ intersection points. This gives us a total of $\Omega(n^3)$ intersection points, which matches the upper bound proven above. In conclusion, we have:

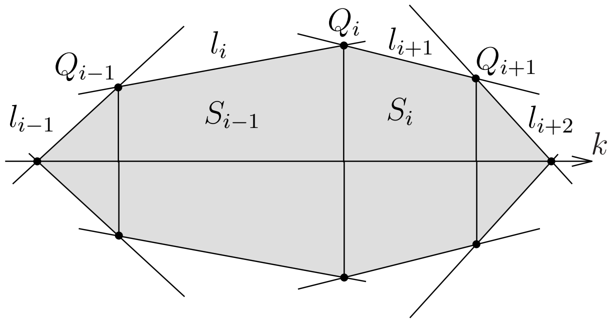

Figure 9: Calculating the area of $I_{P,\ell}$

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

6 Maximizing the Area of the Inscribed Polygon

6.1 Area Function

We may now assume that the configuration of the inscribed polygon $I_{P,\ell}$ is given as a sequence of edges of P: $\{l_1, l_2, ..., l_m\}$ $(m \leq 2n)$. 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 specifies whether the edge belongs to the original polygon or to its reflected version. Note, if l_i contributes to the boundary of $I_{P,\ell}$, then its reflection $\overline{l_i}$ contributes as well. The indicators μ_i are defined as follows: $\mu_i = 1$ if l_i supports an original edge of P and $\mu_i = -1$ if l_i supports an reflected edge.

We briefly describe the computation of the area of $I_{P,\ell}$. We rotate the plane so as to make the axis of symmetry ℓ parallel to the X-axis (see Figure 9). The endpoints Q_i of the edges of $I_{P,\ell}$ can be determined as intersections of pairs of neighboring lines l_i and l_{i+1}. The vertices Q_i and Q_{i+1} are projected onto the X-axis and form a trapezoid whose area is

$$S_i = \begin{vmatrix} Q_{x_{i+1}} & Q_{y_{i+1}} \\ Q_{x_i} & Q_{y_i} \end{vmatrix}. \qquad (5)$$

The coordinates of the point Q_i are

$$Q_{x_i} = \frac{b'_i - b'_{i-1}}{k'_i - k'_{i-1}} \quad \text{and} \quad Q_{y_i} = \frac{k'_{i-1}b'_i - k'_i b'_{i-1}}{k'_i - k'_{i-1}},$$

where

$$k'_i(k,b) = \mu_i \frac{k_i - k}{1 + k_i k} \quad \text{and} \quad b'_i(k,b) = \mu_i \frac{b_i - b}{1 + k_i k}\sqrt{1+k^2}.$$

Note that all the expressions above depend on k and b, which are line coefficients of the axis of symmetry ℓ. Summing up all the partial areas $S_i(b,k)$ gives the total area $S(b,k)$:

$$S(b,k) = \sum_{i=1}^{m} \left(2\frac{b_i b_{i+1}}{k_{i+1} - k_i} + b_i^2 \frac{k_{i-1} - k_{i+1}}{(k_{i-1} - k_i)(k_i - k_{i+1})} \right) \qquad (6)$$

Our convention is that $k_0 = b_0 = k_{m+1} = b_{m+1} = 0$.

The area function for the case of a general simple polygon is identical. The only difference is that the inscribed symmetric polygon may contain few disconnected components, so the configuration can consist of more than one sequence of lines.

Figure 10 plots the area of the inscribed polygon as a function of b and k (the parameters of the axis of symmetry), for a simple inscribing polygon (a 2×2 square).

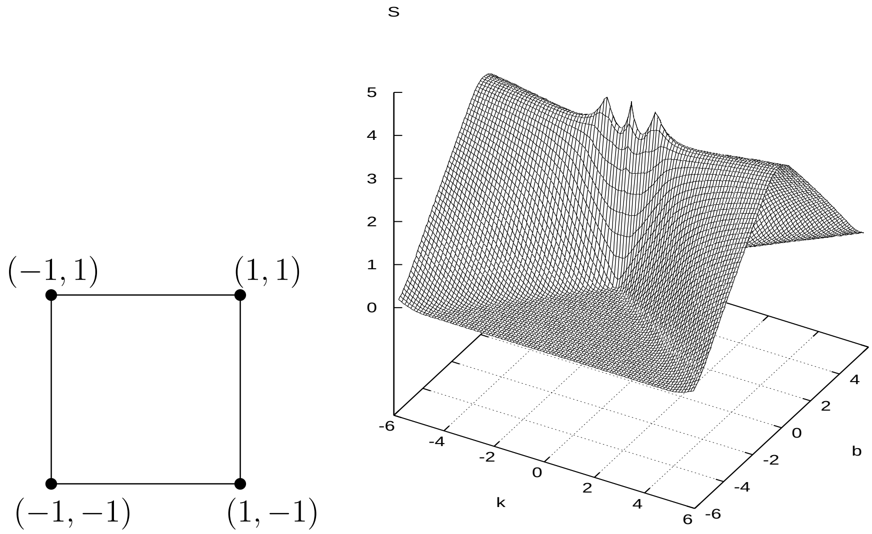

(a) A square (b) Area function

Figure 10: The area $S(k, b)$ as a function of the axis of symmetry

6.2 Maximizing the Area Function

The rest of the effort is to find a global maximum of $S(k, b)$ within each cell of the planar subdivision in the dual plane. This optimization problem is analytically intractable [11] because of the complexity of the objective function, hence we have to resort to numerical methods. Formally, we need to optimize (maximize) a low-dimensional rational polynomial (but with many terms). The objective function is unconstrained (in the sense that the solution does not have to fulfill any other constraints) but the optimum is sought within a bounded region.

A variety of strategies have been proposed to solve exactly such problems. Some methods rely on an a priori knowledge of how rapidly the function varies (e.g., the Lipshitz constant) or on the availability of an analytic formulation of the objective function (e.g., interval methods). Statistical methods use domain partitioning to decompose the search space and a priori information (or assumptions) about how the objective function can be modeled. A wide variety of other methods have been proposed for solving the problems inexactly, including simulated annealing, genetic algorithms, and clustering methods.

A description of a host of algorithms for a global optimization and their available implementations is found in [10]. We note three systems that are suitable for our application: INTGLOB [9] (using an interval technique), ASA [7] (a simulated-annealing algorithm), and Derivative-Free GLOBAL [3] (based on a clustering stochastic multistart method). For our purposes we implemented a straightforward method that works well in practice. We evaluate $S(k, b)$ in regularly-scattered points within the current cell, and choose the best point (with respect to 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 standard steepest-descent and simulated-annealing heuristic methods. We stop when no sufficiently-improving direction is found any more.

7 Running-Time Analysis and Implementation

In this paper we have proposed the following algorithm for computing the maximum-area axis-symmetric polygon inscribed by another polygon P:

1. For each vertex $V_i \in P$ compute the arcs of types I-IV of the planar subdivision M in the dual plane. Construct M.

2. For each face of M compute the associated area function (of symmetric inscribed polygons) and find its maximum (within the face).

3. Report the global maximum (of all faces) as the answer.

We have proven that the combinatorial complexity of M is $\Theta(n^3)$ in the worst case, where n is the complexity of a convex polygon P. Constructing M can easily be done by a plane-sweep procedure whose running time is $O(n^3 \log n)$ (taking into account the k-monotonicity of the chains). In each face we need (a) to compute the area function; and (b) to find its maximum. Computing the area function of the first face takes $O(n)$ time.[4] However, updating the area function while moving from a face to a neighboring face can simply 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 (see details below), for a total of $O(n^3)$ for all the cells. For the analysis we denote by $T(n)$ the average time complexity of the optimization step in a single cell, and note that in practice $T(n) = O(n)$. To conclude, the computation and processing of M requires $O(n^3 \log n)$ time for all stages except the optimization steps. For the latter steps we spend $O(n^3 T(n))$ time. In total, the running time of the algorithm is $O(n^3(\log n + T(n)))$.

Similarly, for a simple polygon, constructing the map and maximizing the area functions require $O(n^4 \log n)$ and $O(n^4 T(n))$ time, respectively. Thus, the total running time is $O(n^4(\log n + T(n)))$.

The space complexity of the algorithm is affected by the construction of the planar map and by the optimization algorithm. The complexity of the planar map for a convex polygon is $\Theta(n^3)$ in the worst case, so the optimal amount of space used for its description is $\Theta(n^3)$, too. However, we can modify the map-construction algorithm and reduce this bound to superquadratic. The idea is to find the combinatorial structures, the area functions, and immediately their local maxima while building the map from a set of arcs (whose cardinality is $\Theta(n^2)$). After a cell of the planar map was discovered and interpreted, and the maximum area within this cell was computed, we do not need this cell any more and may discard the memory used for storing it. That means that at all times only the *zone* of the sweep line is stored in memory, while the complexity of a zone of a vertical line in the map is slightly over $\Theta(n^2)$. Indeed, all arc chains are k-monotone, and there are $\Theta(n)$ such chains, so a vertical sweep line intersects $\Theta(n)$ arcs and hence it intersects $\Theta(n)$ cells of the map. This can also be understood from another point of view: moving along a vertical line corresponds to a vertical translation of the axis of symmetry. During such a translation only $\Theta(n)$ combinatorial changes in a structure of the symmetric inscribed polygon may occur. The complexity of a single cell is superlinear [4, §5.2.2, pp. 122–124]. Thus, the complexity of a zone of the sweep line is superquadratic at any time. The space complexity of the optimization routine depends on the underlying algorithm and its parameters. Regarding these parameters as fixed, we can state that the space complexity of the optimization step is linear in n. Hence, the space complexity of the whole algorithm is superqudratic.

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 which visualizes the objects, concepts, and relations presented in the paper. The user is able to draw a polygon, visualize the respective planar subdivision in the dual plane, move an axis in the primal plane (and see the induced inscribed polygon), and visualize the effect also in the dual plane, or vice versa. The user can also change continuously the geometry of the input polygon and visualize how this affects the inscribed polygon and the planar map in the dual plane. Finally, a three-dimensional graph of the area function can be plotted on screen. All display windows can be saved into graphic files. Figure 11 shows a screen snapshots of our system.

[4] This time reduces to $O(1)$ if the first face corresponds to an empty inscribed polygon (where the axes do not cross the input polygon).

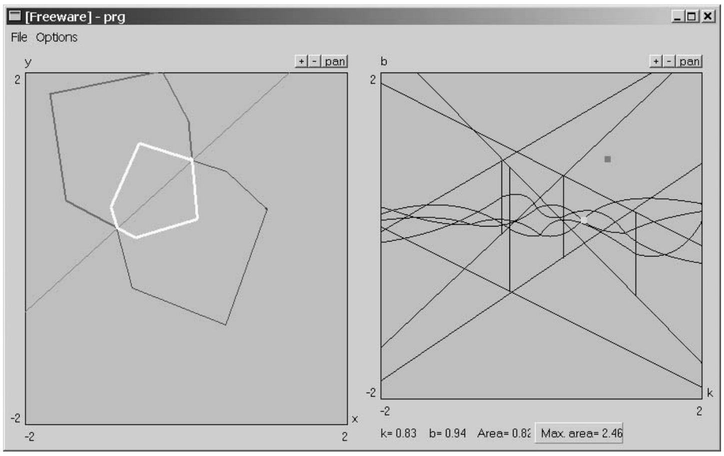

Figure 11: Screen snapshots of the system

References

[1] M. de Berg, O. Devillers, M.J. van Kreveld, O. Schwarzkopf, and M. Teillaud, Computing the maximum overlap of two convex polygons under translations, *Int. Symp. on Algorithms and Computation*, 126–135, 1996.

[2] K.-F. Böhringer, B.R. Donald, and D. Halperin, On the area bisectors of a polygon, *Discrete & Computational Geometry*, 22 (1999), 269–285.

[3] C.G.E. Boender, A.H.G. Rinnooy Kan, G.T. Timmer, and L. Stougie, A stochastic method for global optimization, *Mathematical Programming*, 22 (1982), 125–140.

[4] Micha Sharir, Pankaj K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, *Cambridge University Press, Cambridge-New York-Melbourne*, 1995.

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

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

[7] L. Ingber, Simulated annealing: Practice versus theory, *Computer Modelling*, 18 (1993), 29–57, 1993.

[8] S. Kirkpatrick, C. Gellat Jr. and M. Vecchi, Optimization by simulated anealing, *Science*, 220 (1983), 671–680.

[9] M. Mongeau, H. Karsenty, V. Rouzé, and J. Hiriart-Urruty, Comparison of public-domain software for black box global optimization, *Optimization Methods & Software*, 13 (2000), 203–226.

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

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

[12] A. Tuzikov and H.J.A.M. Heijmans, Minkowski decomposition of convex polygons into their symmetric and asymmetric parts, *Pattern Recognition Letters*, 19 (1998), 247–254.

[13] R. Venkatasubramanian, On the area of intersection between two closed 2-d objects, *Information Science*, 82 (1995), 25–44.