On the Area Bisectors of a Polygon[*]

K.-F. Böhringer,[1] B. R. Donald,[2] and D. Halperin[3]

[1] Department of Electrical Engineering, University of Washington,
Box 352500, Seattle, WA 98195-2500, USA
karl@ee.washington.edu, www.ee.washington.edu/faculty/karl

[2] Department of Computer Science, Dartmouth College,
6211 Sudikoff Laboratory, Hanover, NH 03755-3510, USA
brd@cs.dartmouth.edu, www.cs.dartmouth.edu/~brd

[3] Department of Computer Science, Tel Aviv University,
Tel Aviv 69978, Israel
halperin@math.tau.ac.il, www.math.tau.ac.il/~halperin

Abstract. We consider the family of lines that are area bisectors of a polygon (possibly with holes) in the plane. We say that two bisectors of a polygon P are *combinatorially distinct* if they induce different partitionings of the vertices of P. We derive an algebraic characterization of area bisectors. We then show that there are simple polygons with n vertices that have $\Omega(n^2)$ combinatorially distinct area bisectors (matching the obvious upper bound), and present an output-sensitive algorithm for computing an explicit representation of all the bisectors of a given polygon.

1. Introduction

Let P be a polygon in the plane, possibly with holes, and having n vertices in total. We denote by V the set of vertices of P. For a directed line λ in the plane, we denote by $h_l(\lambda)$ (resp. $h_r(\lambda)$) the open half-plane bounded by λ on the left-hand (resp. right-) side of λ. The line λ is an *area bisector* of P if the area of $P \cap h_l(\lambda)$ is equal to the area of $P \cap h_r(\lambda)$.

A line λ partitions V into three sets (two of which may be empty): $V \cap h_l(\lambda)$, $V \cap \lambda$, and $V \cap h_r(\lambda)$. We say that two area bisectors of P are *combinatorially distinct* if the partitioning of V as above induced by the two bisectors is different. We say that two area bisectors of P are *combinatorially equivalent* if they induce the same partitioning of V.

We assume that the polygon P is connected, and nondegenerate in the sense that its interior is connected.

If instead of a polygon we take a set of n points, we can ask how many combinatorially distinct *halving lines* there are, namely, instead of equating the area on the two sides of a line we now wish to equate the cardinality of the subsets on either side of the line. This is the well-known *k-set* problem (for $k = n/2$) that has been extensively studied in discrete geometry [8]. The best bounds known to date for this problem are the recently obtained upper bound $O(n^{4/3})$ [5] and a lower bound $\Omega(n \log n)$ [10]. An obvious upper bound on the number of distinct *area* bisectors of a polygon with n vertices is $O(n^2)$—see Section 2. In this paper we show that a polygon with n vertices can have $\Omega(n^2)$ distinct area bisectors. (Note that the polygon in our construction is *simple*.)

We devise an output-sensitive algorithm for computing an explicit representation of all the area bisectors of a given polygon. In our analysis we give a more refined distinction between area bisectors (than the distinction by the partitioning of V). We order the area bisectors by slope in the range $[-\pi/2, \pi/2)$, and partition the range into maximal connected intervals with bisectors of the same combinatorial equivalence class. We denote by $\kappa = \kappa(P)$ the number of such slope intervals for a polygon P, and by $\kappa(n)$ the maximum number of such slope intervals over all nondegenerate polygons with n vertices. We denote by $K = K(P)$ the number of combinatorially distinct area bisectors of a polygon P, and by $K(n)$ the maximum number of combinatorially distinct area bisectors over all nondegenerate polygons with n vertices. It is clear that $\kappa(n) \geq K(n)$; we show below that there are polygons for which $\kappa(P) > K(P)$.

It is convenient to consider the problem at a dual setting as well. The *bisector curve* β is defined in a plane dual to the plane containing the polygon: β is the union of points dual to area bisectors. We show that β is *x*-monotone in the dual plane (where the *x*-coordinate corresponds to the slope of the bisector) and that it is a continuous piecewise algebraic curve. Each maximal piece b of β that is defined by the same function describes a contiguous (in slope) set of bisectors that cross the same set of edges of the polygon P, the endpoints of b correspond to bisectors that contain a vertex of P, and any bisector described by an interior point of b does not contain any vertex of P. Hence, $\kappa(P)$ denotes the number of maximal algebraic pieces of β for a polygon P.

A *c*-oriented polygon is a polygon whose edges are all parallel to exactly c directions, where c is a fixed constant. For a *c*-oriented polygon with n vertices, our algorithm for computing an explicit representation of the bisectors by determining the pieces of the curve β runs in time $O((n + \kappa) \log^2 n)$. For a general nondegenerate polygon having

n vertices the algorithm runs in time $O((n + \kappa) \log^2 n + (n + \kappa)\alpha(n)\psi(t))$. Here and throughout the paper, $\alpha(\cdot)$ denotes the extremely slowly growing functional inverse of Ackermann's function, t denotes the maximum number of distinct slopes of edges of P intersected by any line, and $\psi(j)$ denotes the time required to find the roots of a polynomial equation of degree j.

Our algorithm proceeds by constructing the *zone* of the curve β in an *arrangement of lines* in a plane dual to the plane of the polygon. An arrangement $\mathcal{A}(L)$ of a set L of lines is the subdivision of the plane induced by L into vertices, edges, and faces [8]. The zone of a curve in an arrangement of lines is the collection of faces of the arrangement crossed by the curve [12]. In our case the equations defining the curve β depend on the face of the arrangement that β crosses, that is, these equations change from face to face. Hence, we could not use ready-made algorithms for computing the zone of β since the algorithms we are aware of assume that the curve for which the zone is computed is known in advance.

Area bisectors were considered by Díaz and O'Rourke [6], [7]. However, their focus is on the continuous version of the *ham-sandwich cut* problem, and of a problem they introduce of *orthogonal four-sections*; see [6] and [7] for more details.

Our study is motivated by the development of novel, flexible feeding devices for parts positioning and orienting (such as microelectromechanical systems and transversely vibrating plates). The question of determining all the bisectors of polygonal parts arises in connection with the development of efficient part positioning strategies when using these devices (see Fig. 1). For more details see [1]–[4].

The rest of the paper is organized as follows. In Section 2 we present basic properties of area bisectors that are later necessary for the development of the algorithm. In Section 3 we present a simple polygon with n vertices that has $\Omega(n^2)$ combinatorially distinct area bisectors, and a simple polygon P with $\kappa(P) > K(P)$. The algorithmic results are then presented in Section 4. In Section 5 we give concluding remarks and point to open problems raised by our study.

2.   Properties of Area Bisectors

In this section we show that all combinatorially equivalent bisectors of a fixed equivalence class of any polygon can be described by a rational function. We develop explicit formulas for bisectors for given polygon geometry and investigate their algebraic complexity—this is summarized in Theorems 2 and 3 below.

We extend the definition of *combinatorial equivalence* (see the Introduction) to arbitrary lines that partition a polygon. We define combinatorially equivalent partitioning lines of a polygon P to be all those lines λ that induce the same vertex partition $V \cap h_l(\lambda)$, $V \cap \lambda$, and $V \cap h_r(\lambda)$. A necessary (but not sufficient) condition for combinatorial equivalence is that λ intersects the same ordered set of polygon edges.

It is convenient to study the problem in a dual plane: a line $y = 2\bar{x}x - \bar{y}$ in the primal plane is transformed into the point $(\bar{x}, \bar{y})$ in the dual plane. A point (x, y) in the primal plane is transformed into the line $\bar{y} = 2x\bar{x} - y$ in the dual. The dual of an object o will be denoted by o^*. If O is a set of objects in the plane, O^* will denote the set of dual objects. For details on duality transforms see, e.g., [14].

(a)

(b)

Fig. 1. Positioning of parts in a two-dimensional force vector field. (a) In a *squeeze field*, unit forces act perpendicularly toward a central line (the *squeeze line*). The effective forces acting on the left and right portions of the part are shown as arrows attached to the respective centers of area. (b) The part is in equilibrium if and only if its bisector coincides with the squeeze line, and if the connector between the left and right centers of area is perpendicular to the squeeze line. See www.ee.washington.edu/research/mems/Projects/Video for an animated simulation.

In the following lemma we summarize basic facts on area bisectors of a polygon.

Lemma 1. *Given a polygon P with vertices V, $n := |V|$, and a line λ.*

(i) *There exist $O(n^2)$ combinatorially distinct ways in which a line can partition P.*

(ii) *Let A and B be the intersections of a bisector λ with the boundary of the convex hull of P. As the slope angle of λ grows from $-\pi/2$ to $\pi/2$, A and B progress monotonically counterclockwise on the boundary of the convex hull of P.*

(iii) *If the interior of P is connected, then for every slope $\bar{x}$ there exists a unique bisector $\lambda = \lambda(\bar{x})$ of P.*

Proof. (i) This is a well-known fact. We prove it here to demonstrate how duality is used in our arguments. The dual V^* of the n polygon vertices V form an arrangement of n lines in the dual plane. Each point on a line v^* in the dual plane corresponds to a line through the polygon vertex v in the primal plane, and vice versa. Each point within a face of the dual arrangement corresponds to a line in the primal plane that does not intersect

any polygon vertices. Thus, all points within a particular face of the dual arrangement correspond to combinatorially equivalent line placements. We also have to account for equivalence classes where the bisecting lines contain one or more vertices of the polygon. These correspond to edges and vertices of the dual arrangement respectively. Since an arrangement of n lines has $O(n^2)$ vertices, edges, and faces, there are at most $O(n^2)$ combinatorially distinct ways in which a line can partition P.

(ii) [6, Lemma 5.3.1] Assume we are given $\lambda = \lambda(\bar{x})$ with intersection points A and B, and we wish to find $\lambda' = \lambda'(\bar{x} + \varepsilon)$, with intersection points A' and B'. If A' and B' lie to the same side of λ, then it must be the case that one of the partitions of P determined by λ lies strictly within one of the partitions determined by λ', a contradiction since both must be of equal area. Hence it must be the case that A' and B' lie on opposite sides of λ and since $\bar{x}' > \bar{x}$, the progression must be counterclockwise.

(iii) [6, Section 5.2] The existence and uniqueness of $\lambda(\bar{x})$ is guaranteed since the area behind any sweep line of P increases strictly monotonically from zero to the area of P. □

Now consider a bisector λ of polygon P for varying $\bar{x}$ values, as described in Lemma 1(ii). The intersections of λ with the boundary of the convex hull of P, A and B, progress monotonically about that boundary. In general, this progression corresponds to a simultaneous rotation and translation of λ.

Our goal is to obtain a precise mathematical relationship between bisector slope $\bar{x}$ and intercept $\bar{y}$. To this end, we first analyze the effects of rotation and translation separately. To simplify the analysis, we choose a parametrization of the partitioning line λ different from $\bar{x}$ and $\bar{y}$. The line λ is given by a fixed point s on λ and a vector p that specifies its direction (see Figs. 2 and 3).[1]

The results of this analysis are given in Theorems 2 and 3 below. Note that during pure rotation or pure translation, in general the bisector property is not maintained. By imposing the bisector condition $|P \cap h_l(\lambda)| = |P \cap h_r(\lambda)|$ we obtain the relationship between rotation and translation of λ. This is addressed in Theorem 3.

Theorem 2. *Let P be a polygon with n vertices. Let s be a point in $\mathbb{R}^2$, let p be a vector in $\mathbb{R}^2$, and let λ be a line through s in direction p that intersects r edges of P with $t \leq r$ distinct slopes.*

(i) *Let λ' be a line through s in direction p', where $p' = p + \alpha q$ for some fixed vector q and some parameter $\alpha \in \mathbb{R}$. If λ' is combinatorially equivalent to λ, then the area of $P \cap h_l(\lambda')$ is given by a rational function $a_l(\alpha)$ of numerator and denominator degree $O(t)$. In particular, the combinatorially equivalent area bisectors of P passing through s are determined by the roots of a polynomial equation of degree t.*

(ii) *Let λ'' be a line parallel to λ' with distance α' to λ'. If λ'' is combinatorially equivalent to λ', then the area of $P \cap h_l(\lambda'')$ is given by a quadratic polynomial $b_l(\alpha')$.*

---

[1] The transformation to $\bar{x}$ and $\bar{y}$ is straightforward: given $s = (s_x, s_y)$ and $p = (p_x, p_y)$, we get $\bar{x} = p_y/2p_x$ and $\bar{y} = p_y s_x/p_x - s_y$.

Fig. 2. Two nonparallel lines λ and λ' in combinatorially equivalent intersection with polygon P.

Note that the rational functions a_l and b_l express the changing area of $P \cap h_l(\lambda)$ for pure rotation and pure translation of λ, respectively. If λ is a bisector, this area needs to remain constant, which results in a coupling between the rotational parameter α and the translational parameter α'. This relationship is characterized in the following theorem:

Theorem 3. *Let λ be a bisector of P, and consider the combinatorially equivalent line λ'' parametrized by $\alpha, \alpha' \in \mathbb{R}$ as described in Theorem 2 above. There exists a function $f(\alpha, \alpha')$ such that λ'' is a bisector of P iff $f(\alpha, \alpha') = 0$. f is a polynomial of degree t in α and a quadratic polynomial in α', where $t \leq r$ is the number of distinct slopes of the r polygon edges intersected by λ and λ''.*

Proof (Theorems 2 and 3). Consider a line λ, and a point s that lies on λ (Fig. 2). The direction of λ is given by a vector p. Assume for now that the line λ intersects two edges e_1 and e_2 of the polygon P in points p_1 and p_2. Also assume that these edges have direction vectors q_1 and q_2.

We use the following notation throughout this proof. For two vectors $v_1 = (x_1, y_1)$ and $v_2 = (x_2, y_2)$, $v_1 \times v_2$ denotes the signed real whose value is $x_1 y_2 - y_1 x_2$. This value constitutes the area of the parallelogram defined by v_1 and v_2.

Rotating the Bisector. Consider another line λ' with direction p' that intersects λ in s (λ' can be understood as a rotation of λ around s). Assume that λ and λ' have combinatorially equivalent intersections with polygon P, and that λ' intersects the polygon edges e_1 and e_2 in p'_1 and p'_2 respectively. We write $p_i = s + \varrho_i p$ and $p'_i = s + \varrho'_i p'$. Then the polygon area a between λ and λ' is

$$a = \tfrac{1}{2}(\varrho'_2 p' \times \varrho_2 p - \varrho'_1 p' \times \varrho_1 p)$$
$$= \tfrac{1}{2}(\varrho'_2 \varrho_2 - \varrho'_1 \varrho_1)(p' \times p).$$

In the general case where λ and λ' intersect multiple edges with directions q_i of some arbitrary polygon P at points $p_1, p_2, \ldots, p_r$ and $p'_1, p'_2, \ldots, p'_r$ (r even), we can compute the polygon area a between λ and λ' as a sum of signed areas of triangles $sp_i p'_i$:

$$a = \tfrac{1}{2} \sum_{i=1}^{r} (-1)^i w_i$$

$$= \tfrac{1}{2}(p' \times p) \sum_{i=1}^{r} (-1)^i \varrho'_i \varrho_i.$$

Without loss of generality assume $\varrho_r \neq 0$. Then p' can be written as $p' = p + \alpha q_r$ for some $\alpha \in \mathbb{R}$, and the above equation becomes

$$a = \tfrac{1}{2}((p + \alpha q_r) \times p) \sum_{i=1}^{r} (-1)^i \varrho'_i \varrho_i$$

$$= \frac{\alpha}{2}(q_r \times p) \sum_{i=1}^{r} (-1)^i \varrho'_i \varrho_i. \tag{1}$$

From the two vector equations $p'_i = s + \varrho'_i p'$ and $p'_i = s + \varrho_i p + \mu q_i$, $\mu \in \mathbb{R}$, we can determine ϱ'_i as

$$\varrho'_i = \frac{\varrho_i (q_i \times p)}{(q_i \times p) + \alpha(q_i \times q_r)}. \tag{2}$$

Note that the denominator in (2) can be written as $q_i \times (p + \alpha q_r) = q_i \times p' \neq 0$ since q_i cannot be parallel to p' as long as the intersections p_i lie in the same equivalence class. If we also choose the edge direction vectors q_i such that $(q_i \times p) = 1$, then (2) and (1) simplify respectively to the following rational functions in α:

$$\varrho'_i = \frac{\varrho_i}{1 + \alpha(q_i \times q_r)}, \tag{3}$$

$$a(\alpha) = \frac{\alpha}{2} \sum_{i=1}^{r} (-1)^i \frac{\varrho_i^2}{1 + \alpha(q_i \times q_r)}. \tag{4}$$

We look at the denominator $d_i(\alpha) = 1 + \alpha(q_i \times q_r)$ in more detail. This is important because we shall see that in the formulas that we obtain, the denominators consist only of $d_i(\alpha)$'s. For an arbitrary polygon, d_i is a linear function of α. If all q_i are parallel, then $d_i = 1$.

Equation (4) can be used to find a bisector of P that passes through a given point s. Given q_i and ϱ_i ($i = 1 \cdots r$), a is chosen such that $|P \cap h_l(\lambda')| = |P \cap h_r(\lambda')| = \tfrac{1}{2}|P|$. Then (4) can be transformed into a polynomial in α of degree r. More specifically, the degree of the polynomial α is equal to the number t of distinct slopes of polygon edges that the bisector crosses, which in the worst case is r, and c for a c-oriented polygon. This result was summarized in Theorem 2(i).

Translating the Bisector. We now consider the case where λ' shifts parallel (Fig. 3) and show that the change in polygon area b is a quadratic function of the translation

Fig. 3. Two parallel lines λ' and λ'' in combinatorially equivalent intersection with polygon P.

parameter α'. Intuitively, as long as the line $\lambda'(\alpha')$ remains in the same combinatorial class, the total length of the intersection $P \cap \lambda'(\alpha')$ is a linear function of α'. Using an integral argument, it follows that the change in area $b(\alpha')$ is a quadratic function of α'.

We now derive the exact formula for $b(\alpha')$. Analogously to the rotation case, let $p'_i = s' + \varrho'_i p'$ and $p''_i = s'' + \varrho''_i p'$. Also s' and s'' are chosen such that $s'' - s' = \alpha' q_2$. Then the polygon area between λ' and λ'' is

$$
\begin{aligned}
b &= \alpha' q_2 \times \tfrac{1}{2}((p'_2 + p''_2) - (p'_1 + p''_1)) \\
&= \frac{\alpha'}{2}(\varrho'_2 + \varrho''_2 - \varrho'_1 - \varrho''_1)(q_2 \times (p + \alpha q_2)) \\
&= \frac{\alpha'}{2}(\varrho'_2 + \varrho''_2 - \varrho'_1 - \varrho''_1).
\end{aligned}
\tag{5}
$$

Note that the edge direction vectors q_i were chosen such that $q_i \times p = 1$. In the general case λ' and λ'' intersect multiple edges of some arbitrary polygon P at points $p'_1, p'_2, \ldots, p'_r$ and $p''_1, p''_2, \ldots, p''_r$. We choose s' and s'' such that $s'' - s' = \alpha' q_r$. Now the ϱ''_i can be determined from the two vector equations $p''_i = p'_i + \mu q_i$, $\mu \in \mathbb{R}$, and $p''_i = s'' + \varrho''_i p'$:

$$
\begin{aligned}
\varrho''_i &= \varrho'_i - \alpha' \frac{q_i \times q_r}{q_i \times p'} \\
&= \varrho'_i - \alpha' \frac{q_i \times q_r}{1 + \alpha(q_i \times q_r)} \\
&= \frac{\varrho_i - \alpha'(q_i \times q_r)}{1 + \alpha(q_i \times q_r)}.
\end{aligned}
\tag{6}
$$

Then the polygon area between λ' and λ'' is

$$
\begin{aligned}
b &= \frac{\alpha'}{2} \sum_{i=1}^{r} (-1)^i (\varrho_i' + \varrho_i'') \\
&= \frac{\alpha'}{2} \sum_{i=1}^{r} (-1)^i \frac{2\varrho_i - \alpha'(q_i \times q_r)}{1 + \alpha(q_i \times q_r)}.
\end{aligned}
\tag{7}
$$

This is a quadratic polynomial in α' (unless all q_i are parallel, in which case it simplifies to the linear equation $b = \alpha' \sum_{i=1}^{r} (-1)^i \varrho_i$). Equation (7) can be used to find a bisector of P with a given direction p, i.e., with a given slope. This result was summarized in Theorem 2(ii).

Maintaining the Bisector Property.   From the analysis so far we see that if the bisector λ is rotated to λ', then the left and right areas are changed by a value a ($\neq 0$ in general) as described in (4). Hence a subsequent shift of λ' is necessary to restore the bisector property, by changing the areas by a value b, as described in (7).

This implies the condition $a + b = 0$, with a and b given by (4) and (7):

$$
\begin{aligned}
a + b &= \tfrac{1}{2} \sum_{i=1}^{r} (-1)^i \frac{\alpha \varrho_i^2 + 2\alpha' \varrho_i - (\alpha')^2 (q_i \times q_r)}{1 + \alpha(q_i \times q_r)} \\
&= 0.
\end{aligned}
\tag{8}
$$

This equation ensures that λ is a bisector of P. Equation (8) describes $a + b$ as a rational function in α, and a quadratic polynomial in α'. Hence for all combinatorially equivalent bisectors, we can obtain an explicit formula to describe α' as a function of α. We conclude that all combinatorially equivalent bisectors can be represented by an explicit expression parametrized by α.

In general, (8) is equivalent to a polynomial in α and α' whose degree depends on the number $t \leq r$ of distinct slopes of polygon edges intersected by the bisectors λ, λ', or λ'', provided that the $d_i(\alpha)$'s are all nonzero; see the remark above. The polynomial is at most quadratic in α'. In the case that all q_i are parallel, (8) simplifies to the linear equation $\sum_{i=1}^{r} (-1)^i \left(\alpha(\varrho_i/2) + \alpha' \right) \varrho_i = 0$. This result was summarized in Theorem 3.   □

With Theorem 3 we obtain a parametric description for every set of combinatorially equivalent bisectors of P: Given a polygon, we can generate (8) for any particular equivalence class, which yields the polynomial equation $f(\alpha, \alpha') = 0$. Solving f for α', we finally obtain $\bar{x}$ and $\bar{y}$ parametrized by α.

While Theorem 3 gives an algebraic characterization of the area bisectors, in subsequent sections we only use an intermediate result, summarized in Theorem 2(i), describing the area bisectors through a given point in the plane. In this case we have a univariate (in α) polynomial equation. We use this equation in the time analysis of the algorithm in Section 4, and for a combinatorial upper bound on κ in Section 5. In the following section on lower bounds we show that there exist polygons P for which $\kappa(P)$ is strictly greater than $K(P)$.

Fig. 4. A simple polygon with n vertices that has $\Omega(n^2)$ combinatorially distinct bisectors.

3. Lower Bounds

As argued above, a polygon with n vertices can have at most $O(n^2)$ combinatorially distinct area bisectors. Here we give an example where the bound $\Omega(n^2)$ is attained.

Theorem 4. *There exist simple polygons P with n vertices and $\Omega(n^2)$ combinatorially distinct bisectors.*

Proof. Consider Fig. 4. All the vertices v_i, v_i', u_i, and u_i' lie on a circle whose center is at c. The vertices w_j lie very close to c on a small circle whose center is c as well, along two convex polygonal chains.

We fix an integer m (that we determine later; for the polygon in the figure $m = 3$). The distance between the vertices v_i and v_{i+1} is the same for $i = 1, \ldots, m$, and it is the same as the distance between v_i' and v_{i+1}' for $i = 1, \ldots, m$. The area of all triangles $v_i u_i v_{i+1}$ for $i = 1, \ldots, m$ is the same and is equal to the area of all triangles $v_i' u_i' v_{i+1}'$ for $i = 1, \ldots, m$. There are $2m$ vertices w_j near c and they are equally spaced on a small circle centered at c. As can be easily verified, for every pair of vertices v_i and v_i', there is a bisector passing through these points that passes also through the center point c. We

next claim that as we rotate the bisector from v_i to v_{i+1} it will move off the center c and sweep m vertices w_j.

The reason that the bisector will move off center as it rotates from v_i to v_{i+1} is that the angle $\angle u_i v_i v_{i+1}$ is greater than the angle $\angle u_i' v_i' v_{i+1}'$. Hence, as the bisector rotates, it will proceed "faster" on the bottom part of our polygon than on the top part and therefore will sweep half of the vertices w_j on its way.

Finally, m is chosen such that (roughly) $n = 6m + 8$. It follows that the number of distinct area bisectors is $\Omega(m^2) = \Omega(n^2)$. $\qquad \square$

Recall from the Introduction that for a polygon P, $K(P)$ is the number of combinatorially distinct bisectors of P, and $\kappa(P)$ is the number of maximal bisector slope intervals such that within each interval the bisectors are combinatorially equivalent. We conclude this section by showing a polygon P for which $K(P)$ and $\kappa(P)$ are not equal. This example is due to Lovett.[2]

Theorem 5. *There exist simple polygons P such that $\kappa(P) > K(P)$.*

Proof. We show that for a given polygon with n vertices and two combinatorially equivalent area bisectors λ_1 and λ_2, as we vary the slope between the slope of λ_1 and λ_2 the corresponding area bisectors exit and enter the equivalence class of λ_1 (and λ_2) $\Omega(n)$ times.

Let abc be a triangle. It is easily verified that a median in a triangle is an area bisector of the triangle. We denote the median incident to c by λ_1, and the median incident to a by λ_2. Consider the family F of all the area bisectors obtained when we vary the slope between the slopes of λ_1 and λ_2, in the interval where all area bisectors intersect the edges ab and bc. We take the upper envelope of the lines in F, which is a convex curve E.[3] See Fig. 5 for an illustration. For a given n, we take a set Q of $\lfloor (n-1)/3 \rfloor$ evenly spaced points along E, and draw the tangent line to E at each of these points. Each tangent line is an area bisector. Consider two consecutive points and their tangents: In between the points, slightly below the envelope and above the two tangent lines we draw the apex of a thin spike extending from the edge ac of the triangle. We carve this spike out of the triangle. We repeat this process for each pair of consecutive points of $Q \cup \{a, c\}$. Finally we pull the edge ac outward so that the area of the polygon above λ_1 (which is half the area of the original triangle) will be equal to the area below λ_1, compensating for the area of the spikes.

All the tangent lines induced by the set Q are area bisectors of the same equivalence class—they have the single vertex b on one side of the bisector. Evidently as we let the slope vary, every time we cross a spike the equivalence class changes. Hence we obtained a polygon with n vertices and an equivalence class of area bisectors that is entered and exited $\Omega(n)$ times as we let the slope vary. $\qquad \square$

---

[2] S. Lovett, personal communication, Tel Aviv University, 1997.
[3] For example, if we choose the coordinates of T's vertices to be $a = (0, 1)$, $b = (0, 0)$, and $c = (2, 0)$, then E is the hyperbola $f(x) = 1/4x$.

Fig. 5. A simple polygon with n vertices for which $\kappa > K$: by continuously changing the slope of λ, a particular equivalence class of bisectors is entered and exited $\Omega(n)$ times.

4. Output-Sensitive Algorithm

Let P be a polygon with n vertices as defined in the Introduction, namely, connected, nondegenerate, and possibly with holes. Throughout this section we assume that no three vertices of P are collinear (see also the remark at the end of the section). In the dual plane every vertex v of P is transformed into a line v^* which is the collection of all points dual to lines in the primal plane that pass through v.

For any given direction there is a unique area bisector. We denote the oriented bisector of P that makes an angle θ with the positive x-axis by $B(\theta)$, and (because of symmetry) confine ourselves to the range $[-\pi/2, \pi/2)$ for θ. We denote the collection of points dual to area bisectors of P in that range by β. Note that any θ (besides $-\pi/2$) corresponds to an $\bar{x}$-coordinate in the dual plane.

The curve β is a piecewise algebraic and $\bar{x}$-monotone curve (see Theorem 3 and Lemma 1(iii)). We call β the *bisector curve* of P, as it gives a complete specification of all the area bisectors of the polygon P. The number κ as defined in the Introduction corresponds to the number of maximal connected algebraic pieces in β, where the function describing each piece is defined by the fixed set of edges that the corresponding set of bisectors cross. In this section we describe an output-sensitive algorithm to compute β.

Since we aim for output-sensitivity, we cannot afford to compute the entire arrangement $\mathcal{A}(V^*)$ whose complexity can be $\Omega(n^2)$. We will discover the maximal pieces of β in their order along β by dynamically exploring the faces of the arrangement through which β passes, and by computing the intersection points of β with the edges bounding these faces.

Fig. 6.   (a) The face f containing p is the intersection of half-planes containing p; (b) the maximal pieces of β inside f and its exit points from f.

We choose an arbitrary direction $\theta_0 \in [-\pi/2, \pi/2)$ and look for the area bisector of P in that direction. This can be done in $O(n \log n)$ time for an arbitrary polygon [7], and in $O(n)$ time for a simple polygon [16]. We assume that the bisector $B(\theta_0)$ does not pass through a vertex of P, otherwise we change θ_0 slightly so that $B(\theta_0)$ does not pass through a vertex.

Next, we obtain the set of edges crossed by $B(\theta_0)$ in $O(n)$ time (this could also be obtained as a by-product of the algorithm for finding $B(\theta_0)$). We denote by $E(\theta)$ the set of edges crossed by $B(\theta)$. The set $E(\theta_0)$ determines the function describing the bisector curve β in a neighborhood of θ_0. More precisely, the set $E(\theta_0)$ induces a function that describes the bisectors as long as the set of edges crossed by the bisector does not change (see Section 2). In the dual plane this function describes the curve β as long as we do not leave the face of $\mathcal{A}(V^*)$ which contains the point $p := (B(\theta_0))^*$.

Our next step is to construct the face $f = f(p)$ that contains the point p in $\mathcal{A}(V^*)$. We do a little more than just exploring the face f, in preparation for the rest of the algorithm. For each line $v_i^* \in V^*$, let $h_i(f)$ denote the half-plane bounded by v_i^* and containing the face f. Let $H = H(f)$ denote the collection of n half-planes $h_i(f)$. The face f equals the intersection of all half-planes in $H(f)$; see Fig. 6 for an illustration. We explore the faces of the zone of β in the arrangement $\mathcal{A}(V^*)$ by dynamically maintaining the intersection of half-planes H [15]. We denote the necessary dynamic data structure by D. The time to construct the face f and prepare D is $O(n \log^2 n)$ [15].

Now we determine the maximal pieces of $\beta \cap f$. We split the edges bounding f into an upper chain and a lower chain, each set being ordered from left to right. For each edge on the upper chain we compute the intersection of its supporting line with β. Computing the intersection of β with the line v^* supporting an edge is equivalent to finding the bisectors that pass through the vertex v and intersect a fixed set of edges of P. Hence we can use (4) (see also Theorem 2(i)). The roots of the polynomial determine the intersection points; we may have to filter out roots that correspond to bisectors with slope outside the given slope interval. We order the resulting intersections along the $\bar{x}$-axis. We repeat the same for the edges of the lower chain, and merge the two lists of intersection points into a single list I. Since the curve β is $\bar{x}$-monotone, the list I provides a description of the curve β inside f. Moreover, the list I indicates what are the neighboring faces that

β crosses. We mark each of these additional faces by the point where β crosses out of f. We call each such point an *exit point*. See Fig. 6(b).

Since f has already been constructed, we know for each exit point of β the line that contains it (in fact we know a little more, we also know the two lines that bound the edge from which β exits). We next explore one of the faces, say f', into which the curve β crosses out of f through an edge that we denote by e. Let v_i^* be the line containing the edge e. In terms of intersection of half-planes, to obtain f' the only half-plane that needs to be modified is the half-plane supported by v_i^*. More specifically to obtain the new set $H(f')$, we delete $h_i(f)$ from H and add $h_i(f')$ which is the half-plane supported by v_i^* on the other side of v_i^*. (Some caution is needed if β crosses out of f through a vertex. However, since only two lines intersect at that vertex, the operations are similar to those described for the case of crossing out of an edge, and their asymptotic running time is the same.)

Thus we need to delete a half-plane from D and insert a half-plane into D. As a result we obtain the new face f'. The cost of the operation is $O(\log^2 n + C(f'))$, where $C(f')$ is the number of edges on the boundary of f'. The insertion and deletion each costs $O(\log^2 n)$, and the cost of reporting the edges on the boundary of the new face f' is proportional to the number of these edges [15]. We keep a data structure, say a quad-edge structure [11], that describes all the faces of $\mathcal{A}(V^*)$ that have already been constructed so that we do not construct the same face twice. The cost of updating the quad-edge structure with the new face is proportional to $C(f')$.

We continue exploring the faces through which β passes by moving monotonically in the positive x-direction. When the move in that direction has been completed, namely, there are no more exit points to the right of the last exit point handled, we return to the starting point p and repeat the process in the negative x-direction. The algorithm stops when we have identified all the intersection points of β with lines in V^*, and so we have also identified the *zone* of β in $\mathcal{A}(V^*)$, namely, all the faces of $\mathcal{A}(V^*)$ crossed by β.

Recall that κ is the sum, over all faces f in the zone of β in $\mathcal{A}(V^*)$, of the number of maximal connected components $\beta \cap f$. The overall running time of the algorithm consists of the following components:

1. The initial construction of the structure D takes $O(n \log^2 n)$.
2. For each exit point we pay $O(\log^2 n)$ to update D, for a total of $O(\kappa \log^2 n)$.
3. For each newly visited face f we also pay $O(C(f))$ to report its bounding edges and keep them in the quad-edge structure.
4. For each edge on each explored face we compute its intersection with β.

To estimate the cost of components 3 and 4, we need a bound on the sum of $C(f)$ over all faces f of the zone of β in the arrangement $\mathcal{A}(V^*)$. The following lemma bounds the complexity of the zone of an arbitrary curve in an arrangement of lines. We make no assumptions on the curve besides that it consists of a single connected component.

Lemma 6. *Given a collection $\mathcal{L}$ of n lines in the plane and a connected curve γ that intersects the lines of $\mathcal{L}$ in N points in total, then the complexity of the zone of γ in $\mathcal{A}(\mathcal{L})$ is $O((n + N)\alpha(n))$.*

Proof. We modify the arrangements $\mathcal{A}(\mathcal{L})$ as follows (by an idea borrowed from [9]).

Wherever γ crosses a line $\ell \in \mathcal{L}$, we make a small gap in ℓ, splitting ℓ into two segments such that γ passes between them. As a result we get a collection S of $n + N$ segments, and γ lies in a single face of the arrangement $\mathcal{A}(S)$. Since the complexity of a single face in an arrangement of t segments is $O(t\alpha(t))$ a bound of $O((n + N)\alpha(n + N))$ follows. Har-Peled [13] showed that if all the segments lie on only n lines, then the complexity of a single face is $O((n + N)\alpha(n))$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Since the number of intersection points between β and V^* is $\kappa - 1$ we have the following immediate corollary.

Corollary 7. *The maximum combinatorial complexity of the zone of β in $\mathcal{A}(V^*)$ is $O((n + \kappa)\alpha(n))$.*

Thus component 3 of the running time is $O((n + \kappa)\alpha(n))$ which is subsumed by components 1 and 2. It remains to bound component 4, namely, the time to compute all the intersection points of β with the edges of the zone of β. Let t denote the maximum number of distinct slopes of edges of P intersected by a single line. Also, let $\psi(j)$ denote the time to find the roots of a polynomial equation of degree j. (Asymptotically, the term $\psi(j)$ subsumes the time to construct the polynomial.) For general polygons (of the type defined in the Introduction) the time to compute all the intersection points of β with V^* is obviously bounded by $O((n + \kappa)\alpha(n)\psi(t))$.

If P is c-oriented, then by Theorem 2(i) the equation describing the bisectors through a given vertex and intersecting a fixed set of edges is a polynomial of constant maximum degree and it takes constant time to find its roots. It follows that, in this case, the time to compute the intersections of β with the lines V^* is asymptotically dominated by the time to compute the zone of β in $\mathcal{A}(V^*)$.

We summarize the discussion above in the following theorem.

Theorem 8. *Given a nondegenerate polygon P (possibly with holes) with n vertices, we can find a complete specification of its bisectors in time $O((n + \kappa) \log^2 n + (n + \kappa)\alpha(n)\psi(t))$, where κ is the number of maximal connected algebraic pieces of the bisector curve β. If P is c-oriented, then the algorithm runs in time $O((n + \kappa) \log^2 n)$.*

Remark. We assumed above that no three vertices of the polygon P are collinear. If we relax this assumption the following situation may arise. As we explore the zone of β in the arrangement, β crosses out of a face f through a vertex u at which $n_u > 2$ lines are incident. We could reach the face into which β crosses out of f by $O(n_u)$ update operations to the structure D. However, the overall cost of operations at such degenerate crossings is no longer as sharply related to κ as in the bounds above.

5. Conclusions and Open Problems

In this paper we have characterized the family of area bisectors of a polygon. We have shown that a polygon with n vertices can have $\Theta(n^2)$ combinatorially distinct area bisectors, namely, $\Theta(n^2)$ equivalence classes of bisectors, where two bisectors are equivalent

if the lines containing them split the set of polygon vertices in the same fashion. We have also presented an output-sensitive algorithm that produces an explicit representation of all bisectors of a polygon P.

Our algorithm produces the bisector families ordered by slope. We have shown at the end of Section 3 that there exist polygons for which, as we vary the slope of the bisector, the same equivalence class can be entered and exited $\Omega(n)$ times. This raises the following problem: How many equivalence classes of bisectors can be encountered when we let the slope of the bisector vary in the range $[-\pi/2, \pi/2)$? In the paper's notation the question is to obtain a bound on the value $\kappa(n)$. The lower bound $\Omega(n^2)$ for $K(n)$ shown in Section 3 applies here as well. The best upper bound we can show at the moment is $O(n^3)$ derived by the following argument, using the dual arrangement formulation. Consider one edge e of the arrangement $\mathcal{A}(V^*)$. The bisectors whose dual points lie on e are determined by the roots of a polynomial of degree at most n (equation (4)), so their number is at most n, thus the bisector curve cannot cross an edge of the arrangement more than n times. Since the arrangement has $O(n^2)$ edges, the cubic bound follows. In summary there is a gap of an order of magnitude between the lower bound and upper bound on $\kappa(n)$ and the problem is to tighten this gap.

Another open problem is raised by the motivating application. As explained in the Introduction, area bisectors correspond to force equilibria of polygonal parts put on certain part-orienting devices. We are in fact interested in *total* equilibria which are simultaneously force and *moment* equilibria of polygonal parts. A line λ crossing a polygon P induces total equilibrium if (i) λ is an area bisector of P, and (ii) the line connecting the center of mass (area) of $P \cap h_l(\lambda)$ with the center of area of $P \cap h_r(\lambda)$ is perpendicular to λ.[4] The problems we propose are to bound the number of total equilibria of a polygon, and to compute them efficiently. For the combinatorial problem we have only trivial bounds: a lower bound $\Omega(n)$ for a regular n-gon and an upper bound $O(n^3)$ which is obtained in a way similar to the upper bound for $\kappa(n)$, namely, for every equivalence class of area bisectors, the total equilibria are defined by the roots of a polynomial of degree at most n, where n is the number of vertices of the polygon. As for an algorithm, our way to obtain total equilibria is first to find all the distinct bisectors and then solve the problem for each class of bisectors. Could the stage of computing all bisectors be avoided when computing total equilibria? For example, for a sufficiently large n, the part in Fig. 4 (Section 3) has only one stable (and one unstable) total equilibrium, even though there exist $\Omega(n^2)$ distinct area bisectors.[5] This equilibrium is attained when the connecting bar between the outer circle segments and the inner circle is parallel to the squeeze line, and for the unstable equilibrium, the squeeze line is perpendicular. In both cases the squeeze line passes through the center point c. For more details on total equilibria and the motivating application, see [3].

---

[4] The locations of the centers of area can be computed by an analysis similar to the derivation of β in the proof of Theorem 2.

[5] Small disturbances from a *stable* equilibrium cause the part to return to this equilibrium. Slight disturbances in an *unstable* equilibrium are sufficient to move the part further away from equilibrium.

Acknowledgments

References

1. K.-F. Böhringer. Programmable Force Fields for Distributed Manipulation, and Their Implementation Using Microfabricated Actuator Arrays. Ph.D. thesis, Department of Computer Science, Cornell University, Ithaca, NY, Aug. 1997.

2. K.-F. Böhringer, V. Bhatt, and K. Y. Goldberg. Sensorless manipulation using transverse vibrations of a plate. In *Proc. IEEE International Conference on Robots and Automation* (*ICRA*), pages 1989–1996, Nagoya, May 1995. www.ee.washington.edu/research/mems/Projects/VibratoryPlate.

3. K.-F. Böhringer, B. R. Donald, and N. C. MacDonald. Upper and lower bounds for programmable vector fields with applications to MEMS and vibratory plate parts feeders. In J.-P. Laumond and M. Overmars, editors, *Algorithms for Robotic Motion and Manipulation*, pages 255–276. A. K. Peters, Wellesley, MA, 1997. www.ee.washington.edu/research/mems/Projects/PFF.

4. K.-F. Böhringer, B. R. Donald, N. C. MacDonald, G. T. A. Kovacs, and J. W. Suh. Computational methods for design and control of MEMS micromanipulator arrays. *IEEE Comput. Sci. Engrg*, 4(1):17–29, January–March 1997.

5. T. K. Dey. Improved bounds on planar k-sets and k-levels. In *Proc.* 38*th Annual Symposium on Foundations of Computer Science*, pages 156–161, Miami Beach, FL, 20–22 Oct. 1997.

6. M. G. Díaz. Algorithms for Balanced Partitioning of Polygons and Point Sets. Ph.D. thesis, Johns Hopkins University, Baltimore, MD, 1990.

7. M. Díaz and J. O'Rourke. Ham-sandwich sectioning of polygons. In *Proc.* 2*nd Canadian Conference on Computational Geometry*, pages 98–101, Ottawa, 1990.

8. H. Edelsbrunner. *Algorithms in Combinatorial Geometry*, volume 10 of EATCS Monographs on Theoretical Computer Science. Springer-Verlag, Heidelberg, 1987.

9. H. Edelsbrunner, L. Guilbas, J. Pach, R. Pollack, R. Seidel, and M. Sharir. Arrangements of curves in the plane: topology, combinatorics, and algorithms. *Theoret. Comput. Sci.*, 92:319–336, 1992.

10. H. Edelsbrunner and E. Welzl. On the number of line separations of a finite set in the plane. *J. Combin. Theory Ser. A*, 31(1):15–29, 1985.

11. L. J. Guibas and J. Stolfi. Primitives for the manipulation of general subdivisions and the computation of Voronoi diagrams. *ACM Trans. Graphs*, 4:74–123, 1985.

12. D. Halperin. Arrangements. In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, chapter 21, pages 389–412. CRC Press, Boca Raton, FL, 1997.

13. S. Har-Peled. The complexity of many cells in the overlay of many arrangements. M.Sc. thesis, Department of Computer Science, Tel Aviv University, Tel Aviv, 1995.

14. J. O'Rourke. *Computational Geometry in C*. Cambridge University Press, Cambridge, 1994.

15. M. H. Overmars and J. van Leeuwen. Maintenance of configurations in the plane. *J. Comput. Systems Sci.*, 23:166–204, 1981.

16. T. C. Shermer. A linear algorithm for bisecting a polygon. *Inform. Process. Lett.*, 41:135–140, 1992.