Largest Placement of One Convex Polygon inside Another*

Pankaj K. Agarwal[†] Nina Amenta[‡] Micha Sharir[§]

December 5, 1995

Abstract

We show that the largest similar copy of a convex polygon P with m edges inside a convex polygon Q with n edges can be computed in $O(mn^2 \log n)$ time. We also show that the combinatorial complexity of the space of all similar copies of P inside Q is $O(mn^2)$, and that it can also be computed in $O(mn^2 \log n)$ time.

Let P be a convex polygon with m edges and Q a convex polygon with n edges. Our goal is to find the largest similar copy of P inside Q (allowing translation, rotation, and scaling of P); see Figure 1. A restricted version of this problem, in which we just determine whether P can be placed inside Q without scaling, was solved by Chazelle [4], in $O(mn^2)$ time. See also [1, 6, 12] for other approaches to the more general problem, in which Q is an arbitrary polygonal region. (We remark that the complexity of the algorithms for the general case is considerably higher, about $O(m^2n^2)$ in [1], $O(m^3n^2)$ in [12], and $O(m^4n^2)$ in [6].)

Problems concerning the placement of one polygon inside another are important in robotics and manufacturing. This restricted problem is also applicable to an approach to object recognition recently proposed by Basri and Jacobs [3], based on matching two-dimensional faces of polyhedral objects. The transformation which places the largest similar copy of a polygon P derived from a face of an object model inside a polygon Q derived from an image is a candidate for a transformation which matches the entire model to the image.

*Pankaj Agarwal and Micha Sharir have been supported by a grant from the U.S.-Israeli Binational Science Foundation. Pankaj Agarwal has also been supported by National Science Foundation Grant CCR-93–01259, an NYI award, and by matching funds from Xerox Corp. Nina Amenta was employed by the Geometry Center, which is officially the NSF Center for Computation and Visualization of Geometric Structures, supported by NSF/DMS-8920161. Micha Sharir has also been supported by NSF Grants CCR-91-22103 and CCR-93-11127, by a Max-Planck Research Award, and by grants from the Israel Science Fund administered by the Israeli Academy of Sciences, and the G.I.F., the German-Israeli Foundation for Scientific Research and Development.

[†]Department of Computer Science, Box 90129, Duke University, Durham, NC 27708-0129, USA.

[‡]Xerox PARC, 3333 Coyote Hill Road, Palo Alto, CA 94304, USA.

[§]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, ISRAEL and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA.

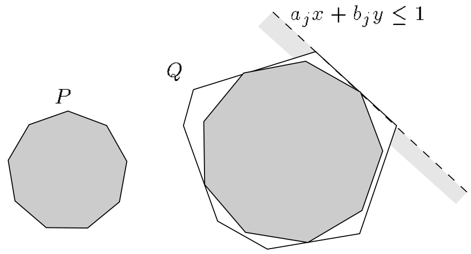

Figure 1: (i) The polygon P; (ii) The polygon Q and a largest copy of P inside Q

The geometric setup of the problem is as follows. We observe, following Baird [2], that similar placements of P can be parameterized nicely by referring to an arbitrarily chosen reference point $p \in P$. A placement π is represented by a quadruple (s, t, u, v), where (u, v) is a translation of p in the plane, and $s = \rho \cos\theta$, $t = \rho \sin\theta$, where P is rotated by θ and scaled by ρ, around p. Let P_π denote the similar copy of P corresponding to the placement π. The standard placement puts p at the origin, with $\rho = 1$, $\theta = u = v = 0$. Thus if (x, y) is a vertex of P in the standard placement, its position at the placement (s, t, u, v) is $(sx - ty + u, tx + sy + v)$. Such a placement of P lies fully within Q if and only if every vertex (x_i, y_i) of P lies in every halfspace $a_j x + b_j y \leq 1$ containing Q and bounded by the line supporting an edge of Q; see Figure 1. That is, the placement (s, t, u, v) must satisfy the following system of mn linear inequalities:

$$a_j(sx_i - ty_i + u) + b_j(tx_i + sy_i + v) \leq 1$$

or

$$L_{i,j}: \quad (a_j x_i + b_j y_i)s + (-a_j y_i + b_j x_i)t + a_j u + b_j v \leq 1\,.$$

In other words, the space $\mathcal{C}$ of all similar placements of P inside Q is a 4-dimensional convex polyhedron formed by the intersection of mn halfspaces. This already implies that the combinatorial complexity of $\mathcal{C}$ is $O(m^2n^2)$, and that it can be constructed in $O(m^2n^2)$ time [11]. However, we will improve this bound in what follows, exploiting the fact that $\mathcal{C}$ is highly degenerate.

In order to find the largest similar copy of P inside Q, we need to find a point of $\mathcal{C}$ that maximizes $s^2 + t^2 = \rho^2$. Unfortunately, maximizing a convex function over a convex polyhedral domain is not an *LP-type problem* (in the setup of [8], where a linear-time randomized solution for such problems is described), so it appears that the algorithm of choice is to examine each vertex of $\mathcal{C}$ and select the one with the largest value of $s^2 + t^2$ (the maximum of such a convex function is clearly attained at a vertex of $\mathcal{C}$). Moreover, since $s^2 + t^2$ depends only on s and t, it suffices to project $\mathcal{C}$ onto the st-plane, and examine only the vertices of that projection.

The main result of the paper is

Theorem 1 (a) *The total number of vertices of $\mathcal{C}$ is $O(mn^2)$, and they can all be computed in time $O(mn^2 \log n)$.*

(b) *The vertices of the projection of $\mathcal{C}$ onto the st-plane can all be computed in time $O(mn^2 \log n)$.*

Remark. Although part (b) follows immediately from part (a), we will give a direct proof of (b), which is somewhat simpler and provides more geometric insight into the structure of the problem.

Proof of Theorem 1: We prove both parts by applying the standard duality transform that maps a point $(\xi_1, \xi_2, \xi_3, \xi_4)$ to the hyperplane $\xi_1 s + \xi_2 t + \xi_3 u + \xi_4 v = 1$ and vice versa. We denote the coordinates in the dual space by s^*, t^*, u^*, v^*. For $1 \le i \le m$ and $1 \le j \le n$, let $w_{i,j}$ denote the point dual to the hyperplane bounding the halfspace $L_{i,j}$, i.e.,

$$w_{i,j} = (a_j x_i + b_j y_i,\ -a_j y_i + b_j x_i,\ a_j,\ b_j).$$

The convex hull of the points in $\{w_{i,j} \mid 1 \le i \le m, 1 \le j \le n\}$, denoted by $\mathcal{D}$, is the dual polytope of $\mathcal{C}$. It is easy to verify that all the points $w_{i,j}$ are extreme points of $\mathcal{D}$ (or, equivalently, that all the hyperplanes bounding the halfspaces $L_{i,j}$ contain facets of $\mathcal{C}$). Note that, for each fixed j corresponding to an edge of Q, the convex hull G_j of $\{w_{i,j}\}_{i=1}^m$ is a similar copy of P that lies in the 2-plane $\pi_j :\ u^* = a_j,\ v^* = b_j$. The dual polytope $\mathcal{D}$, then, is the convex hull of n similar copies of P, placed in parallel 2-planes in 4-space. Each facet of $\mathcal{D}$ corresponds to a placement π of P inside Q such that $P_\pi \subseteq Q$ and there are at least four vertex-edge incidences between the vertices of P_π and the edges of Q.

We begin with the proof of part (b). We exploit the well-known fact that projection in the primal is slicing in the dual. In more detail, let $\mathcal{C}_2$ denote the projection of $\mathcal{C}$ onto the st-plane $u = 0$, $v = 0$, as effected by the mapping $(s, t, u, v) \mapsto (s, t, 0, 0)$. Then a line $\alpha s + \beta t = 1$ in the st-plane is a supporting line of $\mathcal{C}_2$ if and only if the hyperplane $\alpha s + \beta t = 1$ is a supporting hyperplane of $\mathcal{C}$ in $\mathbb{R}^4$. This is equivalent, in the dual, to having the point $(\alpha, \beta, 0, 0)$ belong to the boundary of $\mathcal{D}$. Thus, computing $\mathcal{C}_2$ is equivalent to computing the cross section $\mathcal{D}_2$ of $\mathcal{D}$ with the 2-plane $u^* = 0$, $v^* = 0$.

Our strategy for computing $\mathcal{D}_2$ is first to compute $\mathcal{D}_3$, the cross-section of $\mathcal{D}$ with the hyperplane $u^* = 0$, and then to slice $\mathcal{D}_3$ with the plane $v^* = 0$. Since it is trivial to intersect a three-dimensional polytope with a plane, in time proportional to the complexity of the polytope, we only consider the construction of $\mathcal{D}_3$.

Without loss of generality, we can assume that none of the a_j's is 0. Then each of the polygons G_j lies outside the hyperplane $u^* = 0$. Hence, any vertex w of $\mathcal{D}_3$ must be an intersection of $u^* = 0$ with an edge of $\mathcal{D}$, connecting two vertices of a pair of distinct polygons, G_i and G_j, where G_i lies above $u^* = 0$ and G_j lies below. Moreover, w must also be a vertex of the intersection of the convex hull of $G_i \cup G_j$ with $u^* = 0$. So we can construct

$\mathcal{D}_3$ by taking the convex hull, in $\mathbb{R}^4$, of every pair of polygons G_i, G_j, intersecting all of these sub-hulls with $u^* = 0$, and then taking the convex hull of the resulting intersections.

Let us consider the geometry of one such sub-hull. The two parallel 2-planes $u^* = a_i, v^* = b_i$ and $u^* = a_j, v^* = b_j$ lie in the common 3-plane $F_{i,j}$ defined by

$$(b_j - b_i)u^* + (a_i - a_j)v^* + (b_i a_j - b_j a_i) = 0$$

and so does the sub-hull determined by G_i, G_j. The three-dimensional geometry of $conv(G_i \cup G_j)$ in $F_{i,j}$ is as shown in Figure 2.

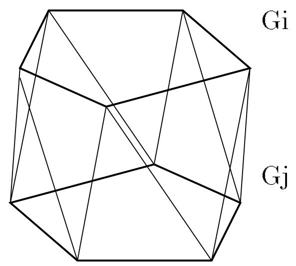

Figure 2: Convex hull of parallel polygons

The intersection of $F_{i,j}$ with $u^* = 0$ is the 2-plane

$$u^* = 0, \quad v^* = (b_i a_j - b_j a_i)/(a_i - a_j)$$

which is also parallel to the two polygons G_i, G_j. Slicing the convex hull of the two parallel polygons with a parallel plane, we get a third parallel polygon $G_{i,j}$ which is the Minkowski sum of appropriately scaled copies of G_i and G_j. This polygon has at most $2m$ vertices, and it is easy to compute directly from the vertices of G_i and G_j. Note that $G_{i,j}$ lies in both $F_{i,j}$ and in $u^* = 0$.

The 3-polytope $\mathcal{D}_3$ in $u^* = 0$ is the convex hull of all these polygons $G_{i,j}$. There are $O(n^2)$ such polygons, each with at most $2m$ vertices, so the total complexity of $\mathcal{D}_3$ is $O(mn^2)$ (which of course is also a consequence of the bound for the overall complexity of $\mathcal{D}$, as asserted in part (a) and proven below).

The algorithm is simply to form the polygons $G_{i,j}$, take their three-dimensional convex hull, and intersect it with $v^* = 0$. Since the Minkowski sum of two convex polygons can be computed in linear time [7], we spend $O(mn^2)$ time in computing the polygons $G_{i,j}$. Their convex hull can be computed in $O(mn^2 \log n)$ time, using the divide-and-conquer algorithm of [10] (which has now only $O(\log n)$ recursive levels, because we start with the already available polygons $G_{i,j}$). Hence, the total running time is $O(mn^2 \log n)$.

This completes the proof of part (b). Note that in practical terms, the implementation of this algorithm is a straightforward setup followed by a three-dimensional convex hull computation, which can be performed efficiently with publicly available software. [1]

We now return to the proof of part (a). We first consider the facets of $\mathcal{D}$ whose supporting hyperplanes are parallel to the 2-plane $u^* = 0, v^* = 0$. The equation of such a hyperplane h_F of a facet F has the form $\beta u^* + \gamma v^* + \delta = 0$. Hence, if h_F contains a vertex of some G_j, it must contain the entire polygon G_j. It then follows that F must be the convex hull of the union of two polygons G_i, G_j (as in the proof of part (b) given above). The facet F is dual to the placement of P in which it is shrunk to a point and all its vertices are incident to the vertex of Q where edge i meets edge j (so that these two edges must be consecutive edges of Q). The number of such placements is n, and the complexity of each of the corresponding facets is $O(m)$, since it is the 3-dimensional convex hull of $2m$ points. (It is easily verified that each of these hulls is indeed a facet of $\mathcal{D}$.) It follows that the overall complexity of these facets of $\mathcal{D}$ is $O(mn)$. Constructing all these facets is easy to do in $O(mn)$ time.

Next, consider the facets of $\mathcal{D}$ whose supporting hyperplanes are not parallel to the 2-plane $u^* = 0$, $v^* = 0$. Let F be such a facet of $\mathcal{D}$, and let h be the hyperplane supporting F. The equation of h can be written as $t^* = \alpha s^* + \beta u^* + \gamma v^* + \delta$ (for simplicity we assume, without loss of generality, that α is never infinite). Then, for each $j = 1, \ldots, n$, the line ℓ_j of intersection between h and the 2-plane π_j containing G_j either touches or is disjoint from G_j. The equation of ℓ_j is $t^* = \alpha s^* + \beta a_j + \gamma b_j + \delta$, $u^* = a_j$, $v^* = b_j$. Note that the coefficient α uniquely determines the vertex of G_j nearest to ℓ_j, for every j, unless α is a 'critical' value equal to the slope of an edge of some G_j. There are $\nu = mn$ such critical slopes α, corresponding to the orientations at which an edge of P is parallel to an edge of Q, and it is easy to compute them, in order, in time $O(mn \log n)$. Let $\alpha_1 < \alpha_2 < \cdots < \alpha_\nu$ be these critical slopes.

Let K be an open interval of α-coefficients between two successive critical slopes. Then, for each $j = 1, \ldots, n$, there exists a unique vertex $w_{i(K),j}$ of G_j, such that if h is any supporting hyperplane of $\mathcal{D}$ whose α-coefficient lies in K, then h can touch G_j, if at all, only at $w_{i(K),j}$. In other words, such an h is also a supporting hyperplane of $S_K = \{w_{i(K),j}\}_{j=1}^n$ (h must of course touch at least one of these vertices, and at least four if it contains a facet of $\mathcal{D}$). For two adjacent intervals K and K', the set $S_{K'}$ is obtained from S_K by replacing one vertex w by another vertex w' (both being adjacent vertices of some G_j). It easily follows that every facet F of $\mathcal{D}$ not parallel to $u^* = 0$, $v^* = 0$ is either a facet of $conv(S_K)$, for some interval K, or, if the α-coefficient of F is a critical value, a facet of $conv(S_K \cup S_{K'})$, for some pair of consecutive intervals K and K'. If the vertices of P and Q are in general position, these latter facets correspond to placements in which an edge of P is incident to

[1] For example Ken Clarkson's `hull` program, at `http://netlib.att.com/netlib/voronoi/hull.html`, or Ioannis Emiris' `chD`, available by `ftp` from `robotics.eecs.Berkeley.edu` in `/pub/ConvexHull`. These and other convex hull programs are listed on the computational geometry software Web page at `http://www.geom.umn.edu/software/cglist`. Using either of these programs gives a randomized algorithm which runs in time $O(mn^2 \log mn)$, slightly worse than our theoretical result.

an edge of Q. In fact, we can prove the following stronger claim. Assuming $\alpha_0 = -\infty$ and $\alpha_{\nu+1} = +\infty$, let K_i be the open interval (α_i, α_{i+1}), for $0 \le i \le \nu$. With a slight abuse of notation, let $S_i = S_{K_i}$ and let ξ_i denote the unique element of $S_i \setminus S_{i-1}$, for $1 \le i \le \nu$.

Lemma 2 *Every facet F of $\mathcal{D}$ that is not parallel to $u^* = 0$, $v^* = 0$ is either a facet of the convex hull $conv(S_0)$ or a facet of the convex hull $conv(S_{i-1} \cup \{\xi_i\})$ incident to ξ_i for some $1 \le i \le \nu$.*

Proof: Let F be a facet of $\mathcal{D}$ that is not parallel to $u^* = 0$, $v^* = 0$ and that is not a facet of $conv(S_0)$. Let W be the set of vertices of F, and let $i \le \nu$ be the index such that the α-coefficient of the hyperplane supporting F lies in the (semi-open) interval $(\alpha_{i-1}, \alpha_i]$. Then, by the above argument, $W \subseteq S_{i-1} \cup \{\xi_i\}$. Suppose $j \le i$ is the largest index such that $\xi_j \in W$ (i.e., S_j is obtained from S_{j-1} by inserting one of the points of W and deleting a point of S_{j-1}.) Then it is easily seen that $W \subseteq S_{j-1} \cup \{\xi_j\}$. Hence, F is a facet of $conv(S_{j-1} \cup \{\xi_j\})$ incident to ξ_j, as asserted. $\square$

This lemma suggests that we should compute $conv(S_0)$ and, for each $1 \le i \le \nu$, we compute the facets of $conv(S_{i-1} \cup \{\xi_i\})$ incident to ξ_i. Since the hyperplanes containing the facets of $conv(S_{i-1} \cup \{\xi_i\})$ incident to ξ_i have only three degrees of freedom, this problem can be formulated as a three-dimensional convex hull problem, and can be solved in $O(n \log n)$ time; the number of these facets, as well as their overall complexity, is $O(n)$. Notice that the set S_0 and the vertices ξ_i, for $1 \le i \le \nu$, can be computed in $O(mn \log n)$ time. Repeating this algorithm for all $1 \le i \le \nu$ and computing $conv(S_0)$, the algorithm produces a total of $O(mn^2)$ facets, of $O(mn^2)$ overall complexity, in time $O(mn^2 \log n)$.

These arguments prove that the total number of facets of $\mathcal{D}$ is $O(mn^2)$, and that their overall complexity, and hence the overall complexity of $\mathcal{C}$, is $O(mn^2)$. Unfortunately, the algorithm might produce additional *spurious* facets, which are not facets of $\mathcal{D}$. Indeed, a facet F of $conv(S_{i-1} \cup \{\xi_i\})$ corresponds to a placement π of P such that there are at least 4 vertex-edge incidences between the vertices of P_π and the edges of Q, and F is spurious if $P_\pi \not\subseteq Q$. If the α-coefficient of F lies in the interval $K_{i-1} \cup K_i$, then it follows by definition that F cannot be spurious. However, if this α-coefficient lies in another interval K_j, for some $j \notin \{i-1, i\}$, then F may be spurious, because P_π may violate a constraint $L_{u,v}$ corresponding to some vertex $w_{u,v} \in S_j \setminus (S_{i-1} \cup S_i)$. See Figure 3 for an example: Let α_i be the critical slope at which the edge p_1p_2 of P is parallel to the edge e_5 of Q. Then, by construction, $S_{i-1} = \{w_{6,1}, w_{5,2}, w_{4,3}, w_{3,4}, w_{1,5}, w_{7,6}\}$, and $\xi_i = w_{2,5}$. It is easy to verify that $conv(\{w_{5,2}, w_{4,3}, w_{3,4}, w_{2,5}\})$ is a facet of $S_{i-1} \cup \{\xi_i\}$ incident to $\xi_i = w_{2,5}$, but, as shown in Figure 3, the corresponding copy of P does not lie inside Q (this facet is 'violated' by $w_{7,1}$).

Hence, to complete our algorithm, we need to detect and discard the facets of the hulls $conv(S_K)$ which are not facets of $\mathcal{D}$. This is accompished as follows. We triangulate each computed facet F into $O(|F|)$ tetrahedra, using the bottom-vertex triangulation scheme described in [5]. Let $\mathbf{\Delta}$ denote the set of resulting tetrahedra; $|\mathbf{\Delta}| = O(mn^2)$. Let $\mathcal{D}^*$ be

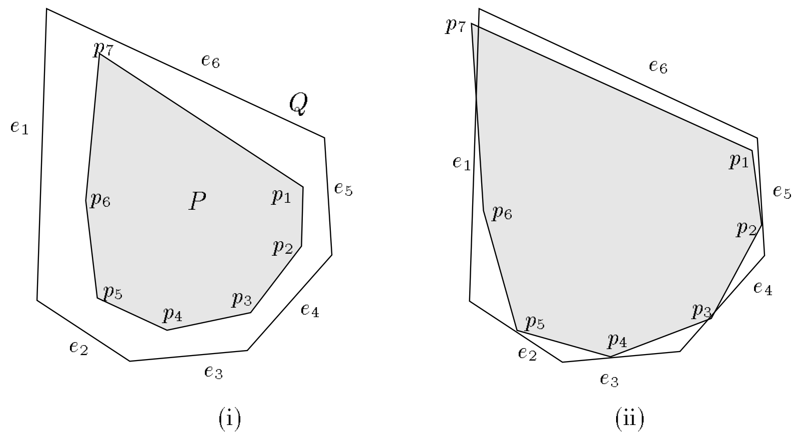

Figure 3: Spurious facets generated by the algorithm: (i) The orientation of P lies in K_{i-1}, where $\xi_i = w_{2,5}$; (ii) A placement of P corresponding to a spurious facet of $S_{i-1} \cup \{w_{2,5}\}$.

the bottom-vertex triangulation of the boundary of $\mathcal{D}$. We want to discard those tetrahedra of $\mathbf{\Delta}$ that are not facets of $\mathcal{D}^*$. For a vertex w, let $\mathbf{\Delta}_w \subseteq \mathbf{\Delta}$ be the subset of tetrahedra incident to w, and let V_w be the set of vertices of the tetrahedra in $\mathbf{\Delta}_w$. It is easily verified that a tetrahedron $\Delta \in \mathbf{\Delta}_w$ is a facet of $\mathcal{D}^*$ if and only if Δ is a tetrahedron in the bottom-vertex triangulation of the boundary of $conv(V_w)$, which is necessarily incident to w. We therefore compute the facets of $conv(V_w)$ that are incident to w, by the reduction, noted above, to a 3-dimensional convex hull construction, and then compute the bottom-vertex triangulation of each such facet. Note that these facets can be computed in $O(|V_w| \log n)$ time, since the vertices of V_w lie on only n 2-planes, so that the convex hull computation requires only $O(\log n)$ recursive levels; we omit the easy details. We can now discard those tetrahedra in $\mathbf{\Delta}_w$ that do not lie on the boundary of $conv(V_w)$. Repeating this procedure for all vertices w of $\mathcal{D}$ gets rid of all spurious facets computed by the algorithm.

The running time of this step is $\sum_w O(|V_w| \log n)$, where the sum extends over all vertices w of $\mathcal{D}$. Since $\sum_w |V_w| = 4|\mathbf{\Delta}| = O(mn^2)$, the total time spent is $O(mn^2 \log n)$.

This completes the proof of part (a). $\Box$

An immediate corollary of Theorem 1(b) is the following.

Corollary 3 *The largest similar copy of P inside Q can be computed in $O(mn^2 \log n)$ time.*

We conclude this paper by constructing a pair of polygons P and Q, with m and n vertices, respectively, such that there are $\Omega(mn^2)$ placements of P inside Q, each of which induces four incidences of the form (p, e), where p is a vertex of P and e is an edge of Q. This implies that the combinatorial bound of Theorem 1(a) is tight in the worst case.

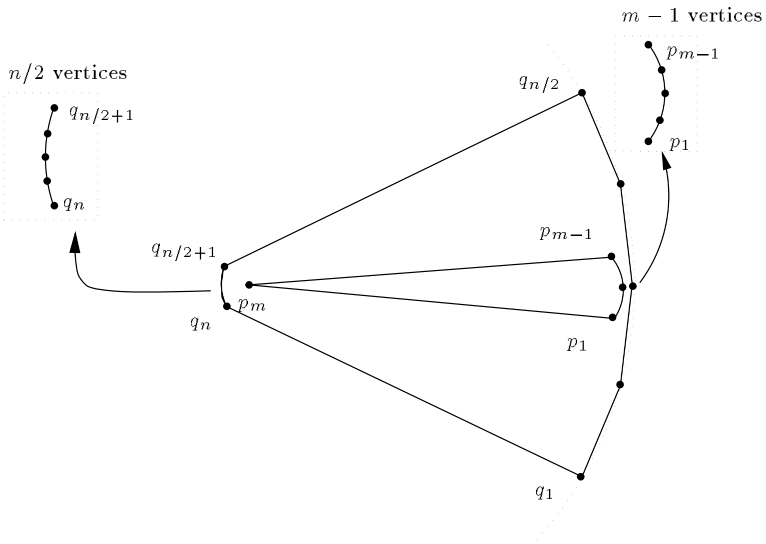

Figure 4: Polygons P and Q for which there exist $\Omega(mn^2)$ similar placements of P in Q with four vertex-edge incidences per placement

The construction is depicted in Figure 4. Let n be of the form $2l+2$, for some positive integer l, m an even integer, and o the origin. The first $n/2$ vertices $q_1, \ldots, q_{n/2}$ of Q are evenly distributed along the arc of the unit-radius circle, centered at o, which goes from $-\pi/6$ to $\pi/6$ (in counterclockwise direction). The vertices $q_{n/2+1} \cdots q_n$ are evenly distributed along a tiny arc of a larger circle, say the circle with radius $10+\varepsilon$ and center $(10,0)$, and we let the tiny arc span the orientations between $\pi - \frac{\varepsilon}{2(10+\varepsilon)}$ and $\pi + \frac{\varepsilon}{2(10+\varepsilon)}$, so that its arc length is ε. The value of ε will be chosen sufficiently small, in a manner to be detailed in a moment.

We place one vertex p_m of P at the origin o and the remaining $m-1$ vertices, equally spaced, on a circular arc of radius $1/4$, centered at $(3/4, 0)$, that spans the orientations between $-\frac{\pi}{40l}$ and $+\frac{\pi}{40l}$.

Claim: *If ε is chosen sufficiently small then the following holds. For every triple $n/2+1 \leq i \leq n$, $1 \leq j < n/2$, and $1 \leq k \leq m-2$, there is a placement of P inside Q, using translation, rotation, and scaling, such that the vertex p_m of P coincides with the vertex q_i of Q, and such that the edge $p_k p_{k+1}$ of P coincides with the edge $q_j q_{j+1}$ of Q.*

Notice that every such placement of P induces four vertex-edge incidences between P and Q, and is thus a vertex of $\mathcal{C}$.

Proof: We consider the scaling, rotation, and translation of P that places $p_k p_{k+1}$ on the line ℓ supporting $q_j q_{j+1}$ and also places p_m at q_i.

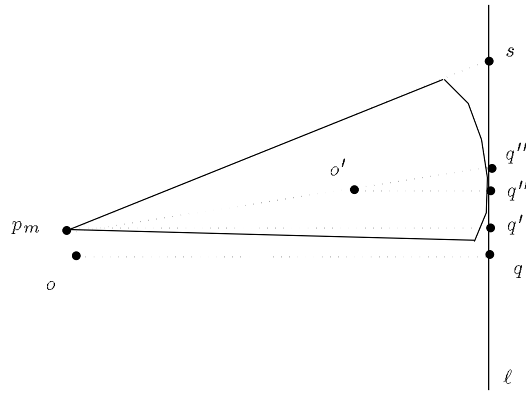

Figure 5: Proof of claim

As in Figure 5, let q be the center of edge $q_j q_{j+1}$; q is also the orthogonal projection of the origin o onto the line ℓ supporting $q_j q_{j+1}$. Let q' be the projection of p_m, which is placed at q_i, onto ℓ. Let q'' be the projection onto ℓ of o', the center of the small circle whose boundary contains the points $p_1, \ldots, p_{m-1}$, which is appropriately shifted with P. Let q''' be the intersection of the line from $p_m = q_i$ through o' with ℓ. Finally, let s be the intersection of the line supporting $p_m p_{m-1}$ (at this placement of P) with ℓ.

The distance from q to q' is at most ε. The angle $q''' p_m q'$ is the same as the angle $q''' o' q''$, which, by the construction of P, is at most $\frac{\pi}{40l}$. The angle $s p_m q'''$ is exactly $\frac{\pi}{40l}$. Since the distance from p_m to q' is at most $1 + \varepsilon$, the distance from q to s is

$$d(q, s) \le \varepsilon + (1 + \varepsilon) \tan \frac{\pi}{20l} .$$

Since the distance from q to q_{j+1} is $\sin \frac{\pi}{6l}$, ε can be chosen small enough so that

$$\varepsilon + (1 + \varepsilon) \tan \frac{\pi}{20l} < \sin \frac{\pi}{6l} ,$$

which then implies that this placement of P fully lies below the segment $p_m q_{j+1}$. An analogous argument shows that P lies above the segment $p_m q_j$, so P lies inside Q, as claimed.

We therefore obtain the following result.

Theorem 4 *There exist a convex m-gon P and another convex n-gon Q such that there are $\Omega(mn^2)$ placements of similar copies of P inside Q, each of which induces four vertex-edge incidences between P and Q.*

Remarks. (1) A weakness of the above lower bound construction is that it only yields placements of P with 'degenerate' vertex-edge contacts, including a vertex-vertex contact and an edge-edge containment. Is there another construction, in which there are $\Omega(mn^2)$ similar placements of P inside Q, such that at each of them four distinct vertices of P touch four distinct edges of Q? This extends a similar open problem, asking for $\Omega(mn^2)$ *congruent* placements of P inside Q, each with three contacts of distinct vertices of P with distinct edges of Q; see [9] for details.

(2) Another open problem is whether the algorithm for finding the largest similar placement of P inside Q can be improved. Such an improvement could be by at most a logarithmic factor if we have to compute the entire space $\mathcal{C}$, as is implied by the above lower bound. Can we do better if we only need to compute the largest placement of P?

Acknowledgments

We are grateful to Boris Aronov and Emo Welzl for helpful discussions, and to David Jacobs and Ronen Basri for bringing Baird's representation to our attention.

References

[1] P.K. Agarwal, B. Aronov and M. Sharir, Motion planning for a convex polygon in a polygonal environment, in preparation.

[2] H. S. Baird, *Model-Based Image Matching Using Location*, Distinguished Dissertation Series, MIT Press, Cambridge, MA, 1984.

[3] R. Basri and D. Jacobs, Recognition using region correspondences, *Proc. 5th Int. Conf. Comput. Vision*, 1995, pp. 8-13.

[4] B. Chazelle, The polygon containment problem, in *Advances in Computing Research, Vol. 1: Computational Geometry* (F. P. Preparata, Ed.), JAI Press, London, England, 1983, pp. 1–33.

[5] B. Chazelle and J. Friedman, A deterministic view of random sampling and its use in geometry, *Combinatorica* 10 (1990), 229–249.

[6] L.P. Chew and K. Kedem, A convex polygon among polygonal obstacles: placement and high-clearance motion, *Comput. Geom. Theory Appls.* 3(2) (1993), 59–89.

[7] L. Guibas, L. Ramshaw, and J. Stolfi, A kinetic framework for computational geometry, *Proc. 24th Annu. IEEE Sympos. Found. Comput. Sci.*, 1983, pp. 100–111.

[8] J. Matoušek, M. Sharir and E. Welzl, A subexponential bound for linear programming and related problems, to appear in *Algorithmica.*

[9] D. Parson and C. Torras, The combinatorics of overlapping convex polygons in contact, *Proc. 4th Canad. Conf. Comput. Geom.*, 1992, pp. 83–92.

[10] F. Preparata and S. Hong, Convex hulls of finite sets of points in two and three dimensions, *Commun. ACM* 20 (1977), 87–93.

[11] F. P. Preparata and M. I. Shamos, *Computational Geometry: An Introduction*, Springer-Verlag, New York, 1985.

[12] M. Sharir and S. Toledo, Extremal polygon containment problems, *Comput. Geom. Theory Appls.* 4 (1994), 99–118.