The Translation-Scale Diagram for Point-Containing Placements of a Convex Polygon*

Gill Barequet† Matthew T. Dickerson‡

Abstract

We present a diagram that captures containment information for scalable translations of a convex polygon. For a given polygon P and contact point q in a set S, the diagram parameterizes possible translations and scales of the polygon in order to represent containment regions for each other point v in S. We present geometric and combinatorial properties for this diagram, and describe how it can be computed and used in the solution of several geometric problems.

1 Introduction

One set of problems that has received considerable attention in the computational-geometry literature is the placement of a polygon (or a polygonal annulus region) so that it contains a given point set (or a subset of it); see, e.g., [BBD, BBDG, BDP, BDG, DGR, DS, EE, ESZ, GMR]. Problem variants include placement of a polygon by translation only, placement by translation and rotation, or placement allowing some geometric transformation (such as scaling, offsetting, or perspective transformations). Optimization variants of the problem include maximization of the number of contained points as well as minimization of the size of the polygon (or polygonal annulus) by either scaling or offsetting. Solutions to these problems have numerous applications in areas such as pattern matching, geometric tolerancing, and robot localization [DGR, Fo, GMR]. Related problems have been studied for convex and nonconvex polygons, and in two dimensions as well as three and higher dimensions.

One such problem, that of finding a translation and rotation of a convex polygon P that maximizes the number of contained points from a given input set S, was studied by Dickerson and Scharstein [DS]. As part of their solution they presented the so-called *rotation diagram*.[1] The rotation diagram $R_{P,q}$ represents all possible placements of a convex polygon P in contact with a particular point $q \in S$. This two-dimensional diagram parameterizes translations along one axis and rotations along the other axis. For every other point $v \in S$, the diagram has a region R_v of all placements of P containing v. The paper describes the combinatorial and geometric properties of rotation diagrams. The complete diagram with all such regions contains the necessary information to solve several placement problems for the polygon P and associated annulus regions. For example, the placement of P that contains the maximum number of points from S and is in contact with a particular point $q \in S$ can be found in $O(n^2m^2 \log(mn))$ time by searching $R_{P,q}$ for a densest point in the arrangement of regions.

In [DS], and also in the current work, the diagrams emphasize placements of a polygon that are *in contact* with some point of the point set. This is because any not-in-contact placement of the polygon that optimizes some point-containment problem can be modified to an in-contact placement without altering the set (or subset) of points contained in the polygon.

In this paper we explore polygon placements that allow scaling. Our placements are limited to translation-only. In particular, we present a two-dimensional containment diagram of a similar nature to that of [DS], but representing translation and scale instead of translation and rotation. We describe several important combinatorial and geometric properties of the containing regions, and show that the complete translation-scale diagram has complexity $O(n(n+m))$ and can be computed in $O(n(n+m)\log(nm))$ time.

*Work on this paper by the first author has been supported in part by an Abraham and Jennie Fialkow Academic Lectureship. Work by the second author has been supported by the National Science Foundation under Grant CCR-99-02032. This work was done while both authors visited the Center for Geometric Computing at Johns Hopkins University, supported by the U.S. Army Research Office under MURI Grant DAAH04-96-1-0013.

†Faculty of Computer Science, The Technion—Israel Institute of Technology, Haifa 32000, Israel. E-mail: barequet@cs.technion.ac.il

‡Dept. of Mathematics and Computer Science, Middlebury College, Middlebury, VT 05753. E-mail: dickerso@middlebury.edu

[1] The *rotation diagram* of [DS] might more appropriately have been called the *translation-rotation diagram*.

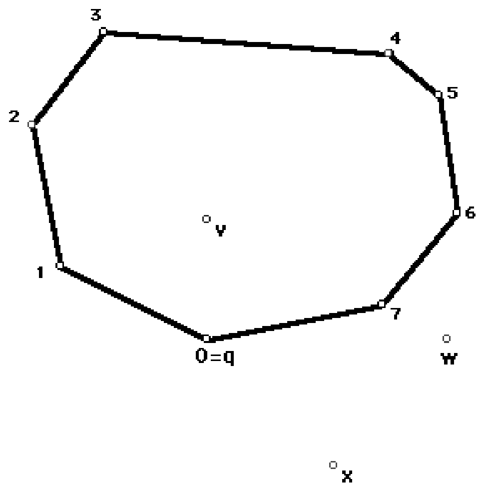

(a) A point set and a polygon

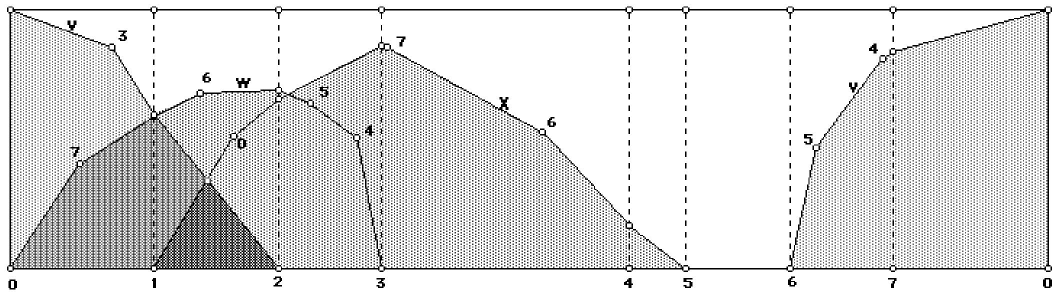

(b) The Diagram

Figure 1: The translation-scale diagram of an 8-gon and a 4-point set

2 The Diagram

We begin with a description of our diagram. We want to create a diagram $D_{P,q}$ that represents translations and scales of a given polygon P. Specifically, we represent all translations that keep P in contact with a given point $q \in S$. For scalings of the polygon, the point q is also used as the center of scaling. That is, we have a two-dimensional diagram that represents along one axis all translations of P in contact with q, and represents along the other axis the factors by which P can be scaled (larger or smaller) keeping the same contact point with q.

For each other point $v \in S$, $v \neq q$, we have a region R_v in $D_{P,q}$ that corresponds to those scales and translations of P that contain v. Actually, rather than parameterizing the scaling factor α of the polygon P, the diagram $D_{P,q}$ parameterizes the inverse $1/\alpha$, which is equivalent to scaling of the entire plane by α while leaving P unscaled. Parameterizing the scaling of the plane rather than of the polygon P results in regions R_v that are X-monotone polygons, as we show in the next section.

Furthermore, although the polygon (and conversely the plane) can be scaled in size by any factor from 0 to ∞, we would like our diagram to be finite along both axes. This is achievable without loss of information. We simply note that if the polygon P is scaled by a small-enough factor α (equivalently, if the plane is scaled by a large enough-factor $1/\alpha$), no point of S (other than q) will be contained by any translation of P in contact with q. Thus all but a finite range of scales can be ignored with no loss of information.

To summarize more formally, we parameterize the two dimensions of the diagram as follows:

1. The horizontal axis of $D_{P,q}$ is a parameterization of the circumference of the polygon: each point on the horizontal axis represents a translation of P putting that particular point of P on q. (The polygon P is keeping the same rotational orientation while sliding around the point q.) For consistency, the horizontal axis parameterizes the polygon in clockwise order, which corresponds to sliding the polygon in counter-clockwise order around point q. In this way our diagram represents all translations of P in contact with q.

2. The vertical axis of $D_{P,q}$ is parameterized from 0 to A, where A is the maximum for all $v \in S$ (with $v \neq q$) of the ratio $d(v,q)/w_{\overrightarrow{vq}}(P)$. Here $d(v,q)$ is the Euclidean distance from v to q and $w_{\overrightarrow{vq}}(P)$ is the width of polygon P in the direction $\overrightarrow{vq}$. So a point on the vertical axis corresponds equivalently to a scaling factor α of the plane or to the scaling factor $1/\alpha$ of P.

   We have a special case at the bottom line of the diagram, which corresponds to the plane being scaled to 0 and the polygon P being scaled to ∞. Here, we must deal in the limit as the scale of the polygon approaches ∞. The definition of containment in this case is as follows.

   Let P be placed such that q is on an edge e of P but not on a vertex. The line through e defines a pair of halfplanes, one of which, call it $H_{e,P}$, contains P. In this case, we say that P contains a point v when the plane is scaled to 0 if and only if v is in $H_{e,P}$.

   Similarly, if P is placed such that a *vertex* p of P is on q, then we consider the half planes $H_{e_1,P}$ and $H_{e_2,P}$ defined by both edges e_1 and e_2 of P incident to p and containing P. A point v is contained by this placement of P when the plane is scaled to 0 if and only if v is in both half planes $H_{e_1,P}$ and $H_{e_2,P}$.

Figure 1 shows an example of the translation-scale diagram. Figure 1(a) shows a set of four points and a convex 8-gon sliding in contact with a point q which belongs to the set. Figure 1(b) shows the diagram that corresponds to the polygon and to the point set. The diagram contains a region R_v for each point $v \in S$ other than q. Note that the left side of the diagram associates with the right side. That is, a translation of 0 is equivalent to a translation of C, where C is the circumference of P. Thus the diagram is really a cylinder.

3 Properties of the Diagram

We now present some geometric and combinatorial properties of the translation-scale diagrams $D_{P,q}$. In the following, we let the polygon P have m vertices, S have n points, and the right and left edges of $D_{P,q}$ are associated. Some of the containment regions may be split in two by the arbitrary choice of which a point on P is mapped to translation 0 by our parameterization. In Figure 1(b), the region corresponding to point v is split into two. By treating the diagram as wrapping around at the right end, we also treat these split regions as being continuous.

Pictorially, each region R_v is a "mountain" with no overhanging cliffs (see Figure 1(b)). As you go from left to right, the trail first climbs up until it reaches the summit, then it descends down. That is, the upper boundary of R_v is always seen from the bottom of the region. The following lemma states this more formally.

Lemma 1 *Every region R_v in $D_{P,q}$ is an X-monotone polygon and exactly two vertices of R_v lie on the horizontal axis (at scale 0) at locations corresponding to vertices of P.*

Proof. Let t be any point on the boundary of P, and equivalently the corresponding translation on the horizontal axis of the diagram $D_{P,q}$. For any vector $\overrightarrow{qv}$ and point t on P, define $w_{\overrightarrow{qv},t}(P)$ to be the width of P at location t in parallel to $\overrightarrow{qv}$. That is, if the vector $\overrightarrow{qv}$, placed at t, points toward the interior of the polygon, then $w_{\overrightarrow{qv},t}(P)$ is the length of the segment obtained as the intersection of the line through q, v and the polygon P. Otherwise, if the vector $\overrightarrow{qv}$, placed at t, points toward the exterior of P, then $w_{\overrightarrow{qv},t}(P) = 0$.

For any point $v \in S$, we can think of the boundary of P as divided into two chains. The *lower chain* corresponds to those points of ∂P from which $\overrightarrow{qv}$ points inward P, and the *upper chain* corresponds to those points of ∂P from which $\overrightarrow{vq}$ points inward P. The vertices of the polygon that are antipodal in the direction orthogonal to $\overrightarrow{qv}$ are the endpoints of both chains. In case P has an edge (or two edges) parallel to $\overrightarrow{qv}$, the endpoints of the chains are the vertices on those edges that are the farthest in direction $\overrightarrow{vq}$. For any point t on the upper chain of P, $w_{\overrightarrow{qv},t}(P) = 0$. Figure 2 shows the upper and lower chains (U and L, respectively) of a polygon, that correspond to the direction $\overrightarrow{qv}$.

Observe that the height of a region R_v in $D_{P,q}$ at translation t is simply the ratio $w_{\overrightarrow{qv},t}(P)/d(q,v)$, as its inverse $d(q,v)/w_{\overrightarrow{qv},t}(P)$ corresponds to the amount by which the polygon needs to be scaled so that v is on the boundary of P, i.e., this is the largest scale of the plane such that v is still contained in P when P is translated to t. (That is, t identifies with q.)

Now take any such vector $\overrightarrow{qv}$ and sweep it across P in a direction perpendicular to $\overrightarrow{qv}$. In particular, consider moving t in clockwise direction along the lower chain of P, and consider vector $\overrightarrow{qv}$ at t. The function $w_{\overrightarrow{qv},t}(P)$ is first monotonically increasing and then monotonically decreasing as t moves along the lower chain ($\overrightarrow{qv}$ is swept across P). In fact, this width function of P in direction $\overrightarrow{qv}$ at location t on the boundary of P gives us precisely the region R_v in $D_{P,q}$. All of the properties of the lemma follow from this observation. $\square$

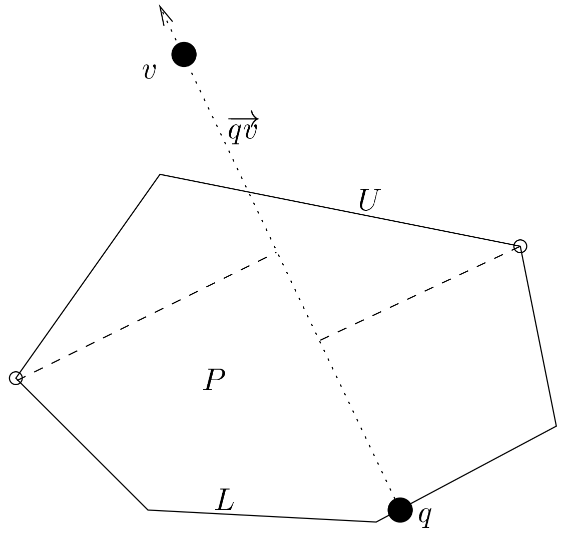

Figure 2: Upper and lower chains of a polygon with respect to a direction

Lemma 2 *For all points $v \in S$, R_v has $O(m)$ edges.*

Proof. This follows from the proof of Lemma 1 with the additional observation that as $\overrightarrow{qv}$ sweeps across P, the only *critical events* where the width does not change linearly are when $\overrightarrow{qv}$ sweeps through a vertex of P. Note that the vertices of R_v are all the points on the lower chain of P where the line parallel to $\overrightarrow{qv}$ passes through a vertex of P on either chain of P. □

In fact, from the proof of Lemma 1, we can be more specific.

Lemma 3 *For each vertex of P, there is exactly one corresponding vertex of R_v. The vertices corresponding to those of the upper chain of P appear in counter-clockwise order from left to right along the top of R_v, while the vertices corresponding to the lower chain of P appear in clockwise order along the top of R_v. The two vertex lists that correspond to the two chains of P are interleaved in R_v.*

This lemma is illustrated in Figure 1(b). In the diagram, we have marked with a vertical line each location on the horizontal axis corresponding to a vertex of P. The upper boundary of region R_v has a vertex at each vertical line that it intersects. These vertices correspond to vertices on the lower chain of P with respect to v and are not labeled. The additional (labeled) vertices are those corresponding to vertices on the upper chain of P.

The previous three lemmas describe what a single region R_v looks like. The following lemmas describes the complexity of the intersections of regions.

Lemma 4 *Discounting the intersections along the bottom line of the diagram corresponding to scale 0, the boundaries of any two regions R_{v_i} and R_{v_j} in $D_{P,q}$ (for $i \neq j$) have at most one intersection which can be either a point or a continuous segment.*

Proof. An intersection of R_{v_i} and R_{v_j} in $D_{P,q}$ corresponds to a translation and scale of P in which v_i, v_j and q are all on the boundary. Not counting the anomaly of scaling the entire plane to size 0, for general-position points there can be at most one such translation-scale pair for which this is true. If two of v, w, and q are collinear on a line parallel to an edge of P, then this intersection may be a continuous segment. □

Lemma 5 *Let edge e_i of region R_i intersect edge e_j of region R_j at a point x in $D_{P,q}$. Let R_i be the region whose boundary is below the boundary of R_j to the right of x. Then the edge e_i corresponds to an edge of P between two vertices of the upper chain of P farther clockwise than those vertices of e_j.*

Proof. Omitted in this short version of the paper. □

Theorem 6 *The intersection of any subset of regions in $D_{P,q}$ is X-monotone and its complexity is $O(m)$. If it is nonempty, then one of its edges lies on the base of the diagram.*

Proof. There are two types of vertices describing the boundary of an intersection regions. Vertices which are part of one of the regions being intersected, and introduced vertices which are the intersections of the boundaries of two (or more) regions. At each point on the horizontal axis corresponding to a vertex of P, there will be a vertex in the intersection region. These correspond to vertices on the lower chain and contribute $O(m)$ vertices to the intersection. By Lemmas 2 and 5, each vertex in an upper chain appears at most once as a vertex in the intersection, with at most one new intersection vertex between each pair, for a total of $O(m)$ vertices. Since every region is X-monotone and has an edge on the base of the diagram (by Lemma 1), so does the intersection of the regions (unless it is empty). □

Refer again to Figure 1(b). The intersection of the three regions in the diagram that correspond to all the points in S but q, is shaded in the figure. Three of the regions contribute edges to the boundary of this intersection region.

We are now ready to state the complexity of the entire diagram, that is, the total number of vertices and intersections between regions.

Theorem 7 *For any convex polygon P and point set S, $D_{P,q}$ has complexity $O(n(m+n))$.*

Proof. There are n regions, each of complexity m for a total of nm features. By the previous lemma, a pair of regions can intersect at most once, so the total number of intersections between regions is $O(n^2)$. In total, we have $O(n(m+n))$ features in the diagram. □

4 Applications of the Diagram

We now highlight some algorithms for computing and using the translation-scale diagram. Detailed proofs of the running times are omitted from this version.

Theorem 8 *For a convex polygon P and point $q \in S$, where $|P| = m$ and $|S| = n$, we can construct a full (explicit) version of $D_{P,q}$ in $O(n(n+m)\log(nm))$ time.*

Standard line-sweep techniques suffice to compute the explicit diagram in the claimed time. We assume that the diagram is represented as a doubly-connected edge list (abbreviated DCEL; see [BKOS]) or by an equivalent data structure, and we store depth information for each region of the diagram. It follows immediately that diagrams for all n points in S can be computed in $O(n^2(n+m)\log(nm))$ time.

Theorem 9 *Given a precomputed diagram $D_{P,q}$, we can compute in $O(m)$ time the smallest scale of P, if it exists, such that there is a translation of P in contact with q and containing the entire point set S.*

We are looking in $D_{P,q}$ for $\cap_{v\in S\setminus\{q\}} R_v$ (the area of maximum depth in the arrangement of regions). This intersection is nonempty if and only if q is on the convex hull of S. To find the intersection, or to determine it is empty, we traverse the bottom of the diagram (the horizontal axis corresponding to scale $\alpha = 0$) looking for a region of depth $n-1$. By Theorem 6, we know that the intersection region of maximum depth has a segment along the bottom of the diagram. This search can be done in $O(m)$ time as the only vertices along the bottom of the diagram correspond to the m vertices of P in the parameterization of the circumference. If the maximum depth of the deepest region is less than $n-1$, then we report that there is no scaled placement of P in contact with q containing the entire set S. If we find the region of depth $n-1$, we then traverse the boundary of that region in $D_{P,q}$ to find the highest point of the region which corresponds to the smallest scale of P containing all the points. By Lemma 1 and Theorem 6 this region has complexity $O(m)$ and can be traversed in that amount of time.

Some other optimization problems that can be solved using these diagrams include:

Theorem 10 *For a given scale α and point q, we can determine the maximum number of points that can be contained by a copy of P in contact with q and scaled by α in $O(n\log(nm))$ time.*

To do this, we shoot a horizontal line L_α through $D_{P,q}$ at scale α looking for the densest point in the arrangement through which L_α passes. In particular, we compute the intersection of L_α with each region in $O(\log m)$ time per region, for a total of $O(n\log m)$ time for all regions. We sort these at most $2n$ intersections in $O(n\log n)$ time, and then find the deepest point along the line. The total time is $O(n\log n + n\log m) = O(n\log(nm))$ if the diagram is precomputed.

Theorem 11 *We can compute the smallest scale of P in contact with q and containing at least k points in $O(n(n+m)\log(nm))$ time.*

To do this, we perform a line sweep down through $D_{P,q}$ from the top to the bottom, looking for the first place of depth k. Since the diagram has complexity $O(n(n+m))$, we can do the sweep in time $O(n(n+m)\log(nm))$ time for one diagram.

5 Conclusions

In this paper we present the translation-scale diagram and show how to use it for solving several optimization problems of placing a convex polygon, while translation and scaling of it are allowed, so that a point set or a subset of it is contained in the polygon.

We are currently working on searching the diagram without explicitly computing it. Our next goal is to extend this diagram to the three-dimensional diagram representing translation, rotation, and scale, representing point-containing placements of the polygon.

Acknowledgment

We thank Cristian Dima and Greg Parent who participated in writing the Java code for the translation-scale diagram.

References

[BBD] G. Barequet, P. Bose, and M. Dickerson, Optimizing constrained offset and scaled polygonal annuli, *Proc. 6th Workshop on Algorithms and Data Structures*, Vancouver, British Columbia, Canada, *Lecture Notes in Computer Science*, 1663, Springer-Verlag, 62–73, 1999.

[BBDG] G. Barequet, A.J. Briggs, M.T. Dickerson, and M.T. Goodrich, Offset- polygon annulus placement problems, *Computational Geometry: Theory and Applications*, 11 (1998), 125–141.

[BDG] G. Barequet, M. Dickerson, and M.T. Goodrich, Voronoi Diagrams for Polygon-Offset Distance Functions, *Proc. 5th Workshop on Algorithms and Data Structures*, Halifax, Nova Scotia, Canada, *Lecture Notes in Computer Science*, 1272, Springer-Verlag, 200–209, 1997. Full version to appear in *Discrete & Computational Geometry.*

[BDP] G. Barequet, M. Dickerson, and P. Pau, Translating a convex polygon to contain a maximum number of points, *Computational Geometry: Theory and Applications*, 8 (1997), 167–179.

[BKOS] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, Computational Geometry: Algorithms and Applications, Springer-Verlag, Germany, 1997.

[DGR] C.A. Duncan, M.T. Goodrich, and E.A. Ramos, Efficient approximation and optimization algorithms for computational metrology, *Proc. 8th Ann. ACM-SIAM Symp. on Discrete Algorithms*, New Orleans, LA, 1997, 121–130.

[DS] M. Dickerson and D. Scharstein, Optimal placement of a convex polygon to maximize point containment, *Computational Geometry: Theory and Applications*, 11 (1998), 1–16.

[ESZ] A. Efrat, M. Sharir, and A. Ziv, Computing the smallest k-enclosing circle and related problems, *Computational Geometry: Theory and Applications*, 4 (1994), 119–136.

[EE] D. Eppstein and J. Erickson, Iterated nearest neighbors and finding minimal polytopes, *Discrete & Computational Geometry*, 11 (1994), 321–350.

[Fo] L.W. Foster, GEO-METRICS II: The application of geometric tolerancing techniques, Addison-Wesley, 1982.

[GMR] L. Guibas, R. Motwani, and P. Raghavan, The robot localization problem, in: *Algorithmic Foundations of Robotics*, A.K. Peters, Ltd., 1995, 269–282.