Polygonal Reconstruction from Approximate Offsets*

Eric Berberich† Dan Halperin‡ Michael Kerber§ Roza Pogalnikova‡

Abstract

Given a polygonal shape Q with n vertices, can it be expressed, up to a tolerance ε in Hausdorff distance, as the Minkowski sum of another polygonal shape with a disk of fixed radius? If it does, we also seek a preferably simple solution shape P; P's offset constitutes an accurate, vertex-reduced, and smoothened approximation of Q. We give a decision algorithm for fixed radius in $O(n \log n)$ time that handles any polygonal shape. For convex shapes, the complexity drops to $O(n)$, which is also the time required to compute a solution shape P with at most one more vertex than a vertex-minimal one.

1 Introduction

Computing the *offset* of a polygon, namely points at most some fixed distance r away from the polygon, is a fundamental geometric operation recurring in a variety of applications. A standard way to obtain it is via the *Minkowski sum* of the polygon and a disk of radius r, which results in a shape bounded by straight-line segments and circular arcs. Modeling the disk in the Minkowski sum with a (tight) polygon yields an approximate piecewise-linear offset. Often, such an approximation is the legacy data which a program has to deal with – the original shape before offsetting is unknown.

While offset computation and smoothening of shapes have been extensively studied, we address the *(offset-)reconstruction problem*, that seems not to have been addressed in the literature: Given a polygonal shape Q, is it the approximate offset of another polygonal shape? And if so, is there a good such P (say, one with a small number of vertices)? As offsetting blurs small features, a definite reconstruction of the original shape from Q (or even of its topology) is impossible in general. However, a good choice of P could lead to a more compact and smooth representation of the shape given by Q.

In Section 2, we present an algorithm that decides for any given polygonal shape Q with n vertices (possibly unbounded), and two real parameters $r, \varepsilon > 0$, whether Q is within Hausdorff-distance ε to the r-offset of some other (yet unknown) polygonal shape P; if the answer is yes, we also return one such P. It gives the exact answer after $O(n \log n)$ operations in the real-RAM model by constructing offsets with increasing radii three times, exploiting this increase in a particular fashion. For convex Q we reduce the running time to optimal $O(n)$ in Section 3 and also compute a P as above which even minimizes (up to one extra vertex) the number of vertices among all valid choices. Furthermore, P's r-offset constitutes a tangent-continuous arc spline approximation of Q where all circular arcs have the same radius. This work summarizes [2] in which we give more details and full proofs.

Related work LEDA and CGAL contain code to compute Minkowski sums of polygons. The latter implementation also computes the exact or approximate offset of a polygon [5].

Smoothening polygonal shapes is desirable for NC machining. Such aims at tangent-continuous *arc-splines* consisting of segments and circular arcs which enable a uniform and fast processing and often alleviate the problem of overheating of the machine or the material. For purely polygonal input one can distinguish results using single arcs or biarcs (besides segments). Drysdale et al. [3] compute a vertex-minimal solution not adding new vertices, while Held et al. [4] compute approximations with arbitrary vertex placements and their tolerance band might even be asymmetric. Our reconstruction approach constrains the solution by allowing a single radius only. It disables tangent-continuity in general. But this can also be seen as a relaxation: We consider our reconstruction approach as an interesting alternative to existing approaches because on success, it yields an approximation that reflects the construction history of Q.

We also seek a vertex-minimal P whose offset is close to Q. P is actually constrained by a set of shapes. A related problem is to find a *minimal-link polygon* that is nested between two others [1] from which our approach adapts some ideas.

2 The Decision for Polygonally Bounded Sets

For a set $X \subset \mathbb{R}^2$ we denote its boundary by ∂X and its complement by $X^C := \mathbb{R}^2 \setminus X$. For a point p and

*Work by E.B. and D.H. has been supported in part by the German-Israeli Foundation (grant no. 969/07). Work by D.H. has also been supported in part by the Israel Science Foundation (grant no. 236/06) and the Hermann Minkowski–Minerva Center for Geometry at Tel Aviv University.

†Max-Planck-Institut für Informatik, Saarbrücken, Germany, eric@mpi-inf.mpg.de

‡Tel Aviv University, Tel Aviv, Israel, {danha|rozapoga}@post.tau.ac.il

§IST Austria, Klosterneuburg, Austria, mkerber@ist.ac.at

a closed X, letting $d(\cdot,\cdot)$ be the Euclidean distance function, we write $d(p,X) := \min\{d(p,x) \mid x \in X\}$. A *polygonal region* $X \subset \mathbb{R}^2$ has a piecewise-linear (finite number of lines) boundary. The points where these straight-line segments intersect are the *vertices* of the polygonal region. If X is bounded, ∂X is a set of (weakly) simple polygons. For two sets X and Y, we denote their Minkowski sum by $X \oplus Y := \{x+y \mid x \in X, y \in Y\}$. For any $c \in \mathbb{R}^2, v \in \mathbb{R}$, we write $D_v(c) := \{p \in \mathbb{R}^2 \mid d(c,p) \leq v\}$ for the (closed) v-disk around c, and $D_v := D_v(O)$ for the disk centered at the origin. The r-offset of a set X, $\text{offset}(X,r)$, is the Minkowski sum $X \oplus D_r$.

The (symmetric) Hausdorff distance of two closed point sets X and Y is $H(X,Y) := \max\{\max\{d(x,Y) \mid x \in X\}, \max\{d(y,X) \mid y \in Y\}\}$. *We say that X is ε-close to Y (and Y to X) if $H(X,Y) \leq \varepsilon$,* which can also be expressed alternatively:

Proposition 1 *For X,Y closed, X is ε-close to Y if and only if $Y \subseteq \text{offset}(X,\varepsilon)$ and $X \subseteq \text{offset}(Y,\varepsilon)$.*

Decision algorithm From now, we fix $r > 0$, $\varepsilon > 0$, and a polygonal region Q, and consider the following question: Can we find a polygonal region P such that Q and the r-offset of P have Hausdorff-distance at most ε? First of all, we can assume that $r > \varepsilon$; otherwise, we can choose $P := Q$, because $\text{offset}(Q,r)$ and Q have Hausdorff-distance at most ε.

Definition 1 *For $r > 0$, and $X \subset \mathbb{R}^2$, the r-*inset *of X is the set* $\text{inset}(X,r) := \text{offset}(X^C,r)^C = \{x \in \mathbb{R}^2 \mid D_r(x) \subseteq X\}$.

Algorithm 1 Is there any closed polygonal region P such that a given Q is ε-close to $\text{offset}(P,r)$?

(1) $Q_\varepsilon \leftarrow \text{offset}(Q,\varepsilon)$
(2) $\Pi \leftarrow \text{inset}(Q_\varepsilon,r)$
(3) $\tilde{Q} \leftarrow \text{offset}(\Pi,r+\varepsilon)$
(4) return $Q \subseteq \tilde{Q}$

We next prove that Algorithm 1 correctly decides whether Q is ε-close to some r-offset of a polygonal region. A first observation is that for any polygonal region P, $\text{offset}(P,r) \subseteq Q_\varepsilon$ if and only if $P \subseteq \Pi$. This is an immediate consequence of the definition of insets. This shows that for any $\text{offset}(P,r)$ that is ε-close to Q, P must be inside Π. Moreover, it shows that any choice of $P \subseteq \Pi$ already satisfies one of Propositions 1's inclusions. It is only left to check whether $Q \subseteq \text{offset}(\text{offset}(P,r),\varepsilon) = \text{offset}(P,r+\varepsilon)$.

Lemma 2 *Q is ε-close to $\text{offset}(P,r)$ if and only if $P \subseteq \Pi$ and $Q \subseteq \text{offset}(P,r+\varepsilon)$.*

To prove correctness of the algorithm, we have to show that $Q \subseteq \text{offset}(\Pi,r+\varepsilon)$ already implies that there also exists a polygonal region $P \subseteq \Pi$ with $Q \subseteq \text{offset}(P,r+\varepsilon)$. Indeed, Π is not polygonal in general; we have to study its shape closer to prove that we can approximate it by a polygonal region, maintaining the property that the offset remains ε-close to Q.

The shape of offsets and insets For a polygonal region Q, it is not hard to figure out the shape of $Q_\varepsilon = \text{offset}(Q,\varepsilon)$: It is a 2-manifold with boundary that is bounded by straight-line segments and by circular arcs, belonging to a circle of radius ε. It is important to remark that all circular arcs are *convex*:

Definition 2 *Let $X \subset \mathbb{R}^2$ be a 2-manifold with boundary with some circular arc γ bounding it. Then, γ is called* concave *with respect to X, if any segment connecting two distinct points on γ is not fully contained in X. Otherwise, the arc is called* convex.

We call X a convexly *(resp.* concavely*)* bounded *region with radius r, if ∂X consists of finitely many straight-line segments and convex (resp. concave) circular arcs that are all of radius r, interlinked at the* vertices *of the region.*

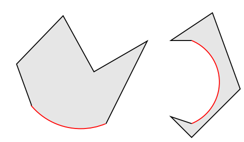

Note that a convexly bounded region (left) is not necessarily convex. The r-offset of a polygonal region P is a convexly bounded region with radius r. The heart of this section is a proof that the same also holds if P is concavely bounded (right) with radius smaller than r:

Theorem 3 *Let P be a concavely bounded region with radius r_1, and let $r_2 > r_1$. Then, there is a polygonal region $P_L \subseteq P$ such that $\text{offset}(P,r_2) = \text{offset}(P_L,r_2)$. In particular, $\text{offset}(P,r_2)$ is a convexly bounded region with radius r_2.*

Note that the correctness of Algorithm 1 already follows by noticing that Q_ε is a convexly bounded region with radius ε, and we can apply Theorem 3 to all constructed offsets, since $\varepsilon < r < r+\varepsilon$:

Corollary 4 *Algorithm 1 returns true if and only if there exists a polygonal region P such that $\text{offset}(P,r)$ is ε-close to Q.*

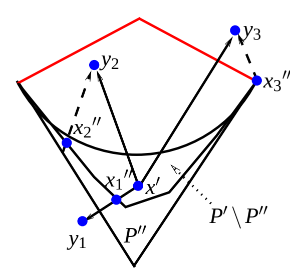

We now give a sketch of the proof of Theorem 3. W.l.o.g., we assume that each concave circular arc γ spans less than half a circle. The arc's *linear cap* is the region enclosed by γ and the two lines tangent to the circle through the endpoints of γ. The *extended linear cap* is the region spanned by the two tangents just mentioned and the two normals at the endpoints.

We iteratively replace an arc γ of a concavely bounded region P' (starting with P) by a polyline ending in the endpoints of γ, such that the polyline does neither leave P' nor γ's linear cap, and such that other boundary parts of P' are not intersected. This yields a concavely bounded region P'' with one arc less.

We show that in each iteration, the r_2-offsets of P' and P'' are the same. For that we consider any point $x' \in P' \setminus P''$, in the region that is cut off by P'', and consider $y = x' + v'$ for an arbitrary $v' \in D_{r_2}$. We show that in each case, y can also be written by $y = x'' + v''$, with $x'' \in P''$, and $v'' \in D_{r_2}$.

The proof then proceeds by studying several cases based on the location of the point y with respect to the extended linear cap of γ; see y_1, y_2, y_3 in the previous figure and [2] for full details of the proof.

Theorem 5 *Let P be concavely bounded with radius r_1 having n vertices, and assume $r_2 > r_1$. Then,* $\text{offset}(P, r_2)$ *has $O(n)$ vertices and it can be computed in $O(n \log n)$ time.*

Proof. By Theorem 3, it suffices to consider a polygonally bounded P_L instead of P; a trapezoidal decomposition leads to a P_L with $O(n)$ vertices. The Voronoi diagram of P_L's vertices and (open) edges can be computed in $O(n \log n)$ and has size $O(n)$ [6]. The r_2-offset boundary inside a Voronoi cell is formed by the intersection of the cell with a parallel line (for the cell of an edge of P_L) or a circle (for the cell of a vertex). Because the offset boundary intersects any Voronoi edge only a constant number of times, the number of vertices (and edges) of the offset is proportional to the number of Voronoi edges. The offset is constructed by sweeping the collection of all the boundary curves from all Voronoi cells, which runs in $O(n \log n)$ because of the absence of interior intersections. □

The running time of Algorithm 1 follows by applying Theorem 5 for the first three steps. The fourth step is easily seen to run in $O(n \log n)$ time as well.

3 Convex Polygons

Lemma 6 *If Q is a convex polygonal region, then Π, as computed by Algorithm 1, is also a convex polygonal region, and it can be computed in $O(n)$ time.*

Proof. Q is the intersection of the halfplanes bounded by lines that support the polygon edges. Observe that Π can be constructed by shifting each such line by $r - \varepsilon$ inside the polygon, which shows that Π is convex. For the time complexity, we compute the lower (upper) envelope for the lines supporting upper (lower) edges of Q by dualizing the lines supporting the edges to points and computing their upper (lower) hull by Graham's scan. We exploit the fact that we already know the x-order of these points. □

The next step of Algorithm 1 would be to check $Q \subseteq \text{offset}(\Pi, r + \varepsilon)$. Let $q_1, \ldots, q_n$ be the vertices of Q (in counterclockwise order) and define $K_i = D_{r+\varepsilon}(q_i)$. The following lemma together with Lemma 6 implies that Algorithm 2 runs in linear time.

Lemma 7 *Q is ε-close to* $\text{offset}(\Pi, r)$ *if and only if* Π *intersects each of the K_i.*

Algorithm 2 Is there any closed polygonal region P such that a given *convex* Q is ε-close to $\text{offset}(P, r)$?

(1) $Q_\varepsilon \leftarrow \text{offset}(Q, \varepsilon)$
(2) $\Pi \leftarrow \text{inset}(Q_\varepsilon, r)$
(3) return $\bigwedge_{i=1}^{n} (K_i \cap \Pi \neq \emptyset)$

Reducing the number of vertices We assume in the remainder of this section that $\text{offset}(\Pi, r)$ is ε-close to Q. Since our goal is to find a possibly simple approximation of Q, we look for a $P \subseteq \Pi$ whose offset is ε-close to Q, but with fewer vertices than Π. Any such P intersects each of the convex (convexly bounded) regions $\kappa_i := K_i \cap \Pi, i = 1, \ldots, n$, of radius $r + \varepsilon$, which we call *eyelets* from now on. The converse is also true: Any *convex* polygonal manifold $P \subseteq \Pi$ that intersects all eyelets $\kappa_1, \ldots, \kappa_n$ has an r-offset that is ε-close to Q.

Proposition 8 *If* $\text{offset}(P, r)$ *is ε-close to Q, and $P \subseteq P' \subseteq \Pi$, then* $\text{offset}(P', r)$ *is ε-close to Q.*

We call a polygonal region P *(vertex-)minimal*, if its r-offset is ε-close to Q, and there exists no other such region with fewer vertices. Necessarily, a minimal P must be convex – otherwise, its convex hull $\text{CH}(P)$ has fewer vertices and it can be seen by Proposition 8 that $\text{offset}(\text{CH}(P), r)$ is also ε-close to Q.

Lemma 9 *There exists a minimal polygonal region $P \subseteq \Pi$ whose vertices are all on $\partial\Pi$.*

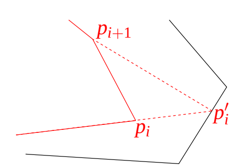

Proof. We pull each vertex $p_i \notin \partial\Pi$ in direction of the ray emanating from p_{i-1} towards p_i until it intersects $\partial\Pi$ in the point p'_i (dragging p_i's incident edges along with it); see the enclosed illustration. For $P' = (p_1, \ldots, p_{i-1}, p'_i, p_{i+1}, \ldots, p_m)$: $P \subseteq P' \subseteq \Pi$, $\text{offset}(P', r)$ is ε-close to Q by Proposition 8. □

Thus, we can restrict our search to polygons with vertices on $\partial\Pi$. We call a polygonal region P *good*, if $P \subseteq \Pi$, all vertices of P lie on $\partial\Pi$, and P intersects each eyelet $\kappa_1, \ldots, \kappa_n$.

Definition 3 *For two points $p_1, p_2 \in \partial\Pi$, we denote by $[p_1, p_2] \subset \partial\Pi$ all points that are met when travelling along $\partial\Pi$ from p_1 to p_2 in counterclockwise order. Likewise, we define half-open and open intervals $[p_1, p_2)$, $(p_1, p_2]$, (p_1, p_2).*

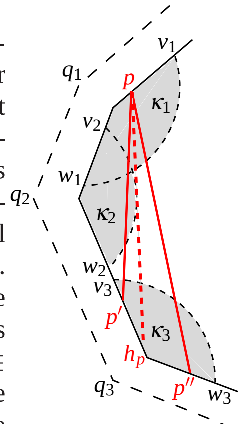

Let $\kappa_i = K_i \cap \Pi$ be q_i's eyelet as before. We consider $\kappa_i \cap \partial\Pi$. The portion of that intersection set that is visible from q_i (considering Π as an obstacle) defines an interval $[v_i, w_i] \subset \partial\Pi$. We call v_i the *spot* of the eyelet κ_i. Finally, for $p_1, p_2 \in \partial\Pi$, we say that the segment $\overline{p_1 p_2}$ is *good*, if for all spots $v_i \in (p_1, p_2)$, $\overline{p_1 p_2}$ intersects the corresponding eyelet κ_i. The figure on the right illustrates these definitions: The segment $\overline{pp'}$ is good, whereas $\overline{pp''}$ is not good, because $v_2 \in (p, p'')$, but it does not intersect κ_2.

Theorem 10 *Let P be a convex polygonal region with all its vertices on $\partial\Pi$. Then, P is good if and only if all its bounding edges are good.*

Proof. Any spot v_i of an eyelet κ_i either corresponds to some vertex p_ℓ of P, or lies inside some interval $(p_\ell, p_{\ell+1})$. Since $\overline{p_\ell p_{\ell+1}}$ is good, it intersects κ_i. For the converse, assume that $\overline{p_\ell p_{\ell+1}}$ is not good, which encloses with the interval $(p_\ell, p_{\ell+1})$ a polygonal region $R \subseteq \Pi \setminus P$. Hence, there is a spot $v_i \in R$ such that $\overline{p_\ell p_{\ell+1}}$ does not intersect the eyelet κ_i. It follows that the entire κ_i is inside R (see the above illustration, considering $\overline{pp''}$ and κ_2). Thus, $P \cap \kappa_i = \emptyset$, and so P cannot be good. □

For $p \in \partial\Pi$, we define its *horizon* $h_p \in \partial\Pi$ as the maximal point (when travelling from p in counter-clockwise order on $\partial\Pi$) such that the segment $\overline{ph_p}$ is good. An example is depicted in the previous figure: The segment $\overline{ph_p}$ is tangential to κ_2, so if going any further than h_p from p, the segment would miss κ_2 and thus become non-good.

Lemma 11 *Let P be a good polygonal region, and $u \in \partial\Pi$. Then, P has a vertex $p \in (u, h_u]$.*

Proof. Assume that P has no such vertex, and let $p_1, \ldots, p_\ell$ be its vertices on $\partial\Pi$. Let p_j be the vertex of P such that $u \in (p_j, p_{j+1})$. Then, also $h_u \in (p_j, p_{j+1})$, because otherwise, $p_{j+1} \in (u, h_u]$. Since P is good, the segment $\overline{p_j p_{j+1}}$ is good, too. It is not hard to see that, consequently, both $\overline{p_j u}$ and $\overline{u p_{j+1}}$ are good. However, the latter contradicts the maximality of the horizon h_u. □

For an arbitrary initial vertex $s \in \partial\Pi$, we finally specify a polygonal region P^s by iteratively defining its vertices. Set $p_1 := s$. For any $j \geq 1$, if the segment $\overline{p_j s}$, which would close P^s, is good, stop. Otherwise, set $p_{j+1} := h_{p_j}$. Informally, we always jump to the next horizon until we can reach s again without missing any of the eyelets. By construction, all segments of P^s are good, so P^s itself is good.

Theorem 12 *Let P be a minimal polygonal region for Q, having* OPT *vertices. Then, for any $s \in \partial\Pi$, P^s has at most* $\mathrm{OPT}+1$ *vertices*

Proof. We first prove that P^s has the minimal number of vertices among all good polygonal regions that have s as a vertex. Let $s := p_1, \ldots, p_m$ be the vertices of P^s. There are $m-1$ segments of the form $\overline{p_\ell h_{p_\ell}}$. By Lemma 11, any good polygonal region has a vertex inside each of the intervals $(p_\ell, h_{p_\ell}]$. Together with the vertex at s, this yields at least m vertices, thus P^s is indeed minimal among these polygonal regions.

Next, consider any minimal polygonal region $P^\star$. We can assume that all its vertices are on $\partial\Pi$ by Lemma 9. If s is not a vertex of $P^\star$, we add it to the vertex set and obtain a polygonal region P' with at most $\mathrm{OPT}+1$ vertices that has s as a vertex. P^s has at most as many vertices as P', so $m \leq \mathrm{OPT}+1$. □

As each visit of an eyelet requires constant time, the construction of a horizon is proportional to the number of visited eyelets, and there are only linearly many eyelets, we can state:

Theorem 13 *For an arbitrary initial vertex s, P^s can be computed in $O(n)$ time.*

Additional Material In the extended version of this paper [2] we also present a new approximation of a polygonal shape's r-offset by line segments and circular arcs aiming at an accurate and compact description. Its vertices are rational and the Hausdorff distance to the exact offset is at most some prescribed ε. In addition we discuss there some immediate extensions of the algorithms presented here.

References

[1] A. Aggarwal, H. Booth, J. O'Rourke, S. Suri, and C.-K. Yap. Finding minimal convex nested polygons. *Inf. Comput.*, 83(1):98–110, 1989.

[2] E. Berberich, D. Halperin, M. Kerber, and R. Pogalnikova. Polygonal reconstruction from approximate offsets. Manuscript, 2009.

[3] R. Drysdale, G. Rote, and A. Sturm. Approximation of an open polygonal curve with a minimum number of circular arcs and biarcs. *Computational Geometry*, 41(1-2):31 – 47, 2008.

[4] M. Heimlich and M. Held. Biarc approximation, simplification and smoothing of polygonal curves by means of voronoi-based tolerance bands. *Int. J. Comput. Geometry Appl.*, 18(3):221–250, 2008.

[5] R. Wein. Exact and approximate construction of offset polygons. *Computer-Aided Design*, 39(6):518–527, 2007.

[6] C.-K. Yap. An $O(n \log n)$ algorithm for the Voronoi diagram of a set of simple curve segments. *Discrete & Computational Geometry*, 2:365–393, 1987.