

Gathering by Repulsion

Prosenjit Bose¹

School of Computer Science, Carleton University, Canada
jit@scs.carleton.ca

Thomas C. Shermer

School of Computing Science, Simon Fraser University, Canada
shermer@sfu.ca

Abstract

We consider a repulsion actuator located in an n -sided convex environment full of point particles. When the actuator is activated, all the particles move away from the actuator. We study the problem of gathering all the particles to a point. We give an $O(n^2)$ time algorithm to compute all the actuator locations that gather the particles to one point with one activation, and an $O(n)$ time algorithm to find a single such actuator location if one exists. We then provide an $O(n)$ time algorithm to place the optimal number of actuators whose sequential activation results in the gathering of the particles when such a placement exists.

2012 ACM Subject Classification Mathematics of computing → Graph theory, Theory of computation → Design and analysis of algorithms, Theory of computation → Computational geometry

Keywords and phrases polygon, kernel, beacon attraction

Digital Object Identifier 10.4230/LIPIcs.SWAT.2018.13

1 Introduction

In this paper, we consider some basic questions about *movement by repulsion*. Here a point actuator repels particles, or put another way, particles move so as to locally maximize their distance from the actuator. This problem models magnetic repulsion, movement of floating objects due to waves, robot movement (if robots are programmed to move away from certain stimuli), and crowd movement in an emergency or panic situation. It is, in one sense, the opposite of movement by attraction, which has recently been an active topic of research [2, 3, 11, 10, 14, 9, 1, 8].

1.1 Related work

We initiate the study of repulsion in polygonal settings. The closest comparable work is the work on attraction. Although attraction and repulsion have a similar definition, each has a distinct character. Attraction as it has been studied is mainly a two-point relation: a point p attracts a point q if q , moving locally to minimize distance to p , eventually reaches p . In repulsion, p cannot repulse q to itself; p must always repulse q to some other point r . Thus repulsion is a three-point relation.

In attraction, if a particle is attracted onto an edge by a beacon, it is pulled towards the point p where there is a perpendicular from the beacon to the line through the edge. If p is on the edge, this creates a stable minimum at p , and particles accumulate at such minima. As well, particles can accumulate on some convex vertices.

¹ Research supported in part by NSERC.



In repulsion, if a particle is repelled onto an edge by a repulsion actuator, it is pushed away from the point p with the perpendicular to the actuator. This implies that p is an *unstable* maximum. We forbid particles from stopping at unstable maxima, so in repulsion the only accumulation points will be convex vertices. We elaborate further on our model in Subsection 1.2.

In this article, we highlight some of the similarities as well as distinctions between these two concepts. For instance, Biro [2] designed an $O(n^2)$ time algorithm for computing the *attraction kernel* of a simple n -vertex polygon P ; these are all points $p \in P$ that attract all points $q \in P$. The closest counterpart of this for repulsion which we call the *repulsion kernel* of a polygon P , which is all points $p \in P$ such that there exists a point $r \in P$ such that p repels all points in P to r . We give an $O(n^2)$ time algorithm to compute the repulsion kernel of an n -vertex convex polygon, and an $O(n)$ time algorithm to find a single-point in the repulsion kernel or report that the kernel is empty.

Both the attraction kernel and the repulsion kernel are concerned with the problem of gathering particles to a point. When the repulsion kernel is empty, it may be the case that we can still gather all particles to a point using more than one repulsion actuator. In this vein, we prove that this is impossible in a polygon with three acute angles. In a convex polygon with at most two acute angles, two repulsion actuators are always sufficient and sometimes necessary. We then provide an $O(n)$ time algorithm to place the optimal number of actuators.

1.2 The model


We start with an n -vertex convex polygon P , which includes its interior. Before the activation of any repulsion actuator, there is a particle on every point of the polygon, including the boundary. During and after activation, we allow many particles to be on the same point; once two particles reach the same point, they travel identically, so we consider them to be one particle.

We restrict the location of the repulsion actuator to points in P ; allowing the actuator to reside outside P leads to a variation of the problem in which convex polygons are easily dispensed.

See Figure 1 for an illustration of the following definitions. The activation of an actuator will cause all particles to move to locally maximize their distance from the actuator. This means that if a particle is in the interior of P , then it moves in a straight line away from the actuator's location. If a particle is on an edge of the polygon, then it proceeds along the edge in the direction that will further its distance from the active actuator. Once moving, a particle moves until it is stable and can no longer locally increase its distance from the actuator. Stable maxima happen at vertices where neither of the two edges allows movement away from the actuator. We call such vertices the *accumulation points* of the activation.

Unstable maxima happen when a particle is on an edge where one or both directions give no differential change of distance from the actuator; this happens only at the perpendicular projection of the actuator onto the edge (see Figure 1c). A particle at an unstable maximum will move off of it in a direction of no improvement and then will be able to increase the distance from the actuator by continuing in that direction. To maintain a deterministic model, we will assume that particles move counterclockwise around the polygon at unstable maxima if there is a choice of two directions of no improvement. However, the choice of counterclockwise motion is arbitrary, and does not affect our results.

We may activate actuators sequentially from several places inside the polygon. We would like for every activation of an actuator to be from a location without particles, but the particle-on-every-point model forbids this on the first activation. So, when we choose a



■ **Figure 1** (a) An activation at w drives the particle at q away from w . On reaching an edge at q' , it will continue to move away from w , until it reaches a local maximum of distance from w at a . (b) Accumulation points of an active actuator at w . (c) At an unstable maximum, such as q' , particles will turn left.

location for the first actuator, we remove the particle at that location from the problem. For subsequent activations, however, we do require that the actuator's position be chosen from the points of the polygon without particles.

The main question we consider is when can we place a sequence of points such that repulsion from those points gathers all other points in the polygon to one point? When the repulsion kernel is non-empty, one point is sufficient. In general, our goal is to minimize the number of sequential activations performed to gather all the particles to one point. If all the particles in a polygon can be gathered to a point with k sequential activations of actuators, we call the polygon k -gatherable. If this is not possible for any k , then we call the polygon *ungatherable*.


2 Background, notation, and terminology

2.1 General notation


We will use the convention that the vertices of P are v_0, v_1, \dots, v_{n-1} in counterclockwise order around the polygon. Vertex indices are taken modulo n , so $v_{-1} = v_{n-1}, v_0 = v_n, v_1 = v_{n+1}$, etc. Edges are denoted e_0, e_1, \dots, e_{n-1} , with e_i being the edge between v_i and v_{i+1} . The boundary of the polygon P will be denoted ∂P , and by $\partial P(p, q)$ we mean the part of ∂P from p counterclockwise to q . In reference to curves, line segments, or intervals, we use the usual parentheses to denote relatively open ends and square brackets to denote relatively closed ends. Thus $\partial P[p, q)$ is the boundary from p to q , including p but not q . Given three distinct points a, b, c in the plane, by $\angle abc$ we mean the counterclockwise angle between the ray from b to a and the ray from b to c .

2.2 Slabs and the three regions of an edge

Consider a polygon edge with particles covering it. When an actuator is activated, depending on its location relative to the edge, there are three possible effects on the particles: it drives them counterclockwise over the entire edge, it drives them clockwise over the entire edge, or it drives some of them clockwise and some of them counterclockwise (see Figure 2). In the latter case, a perpendicular from the edge to the actuator exists, and the particles clockwise of the perpendicular are driven clockwise, and the particles counterclockwise of the perpendicular are driven counterclockwise. The point where the perpendicular hits the edge is called a *split point*. We allow split points at the endpoint of an edge if a perpendicular from the endpoint to the actuator exists.



■ **Figure 2** We use arrows in the diagrams to show the direction that the particles are driven. (a) The activation drives the particles (on the indicated edge) clockwise. (b) The activation drives the particles counterclockwise. (c) The activation splits the particles at s , driving some clockwise and some counterclockwise.



■ **Figure 3** A flow diagram, showing the direction of particle movement, along with the accumulation points and split points, given an actuator at w .

We divide the inner halfplane of an edge e into three regions depending on what effect an activation in the region has on the particles on the edge. This is done by drawing interior-facing perpendiculars to the edge at each of its vertices. The regions are $R_{cw}(e)$, where an activation drives the particles clockwise, $R_{ccw}(e)$, where an activation drives the particles counterclockwise, and $S(e)$, where an activation drives some particles clockwise and some counterclockwise. We refer to $S(e)$ as the *slab* of e . The slab is closed on its boundaries, and $R_{cw}(e)$ and $R_{ccw}(e)$ are open where they meet $S(e)$.

2.3 Flow diagrams

Given a polygon P and a location w of an actuator, we may find the accumulation points and the split points, and mark each edge (or portion of a split edge) with the direction of particle movement along that edge, as in Figure 3. We call a diagram of this a *flow diagram for w with respect to P* .

► **Lemma 1.** *In a traversal of ∂P , accumulation and split points alternate.*

Proof. Note that in a flow diagram the only points of the boundary with two opposing directions of particle movement are the accumulation points, where the movement is towards the point, and the split points, where the movement is away from the point. Thus, between any two consecutive split points on the boundary, there must be an accumulation point, and between any two consecutive accumulation points, there must be a split point. This implies the lemma. ◀

► **Theorem 2.** *A convex polygon P is 1-gatherable from w iff w lies in the slab of exactly one of the edges of P .*

Proof. A polygon is 1-gatherable from w iff an actuator at w has one accumulation point. Since accumulation and split points alternate, this holds iff the actuator has exactly one split point. Since an actuator has a single split point in every slab that it is in (and no others), the result follows. ◀

The boundary of the slab for edge e consists of e and two rays perpendicular to e . If we produce these two rays for each edge of P , and intersect all these rays with P , we get a set of at most $2n$ chords that define a decomposition that we call the *slab decomposition* of P . An example slab decomposition is shown in Figure 5. The cells of this decomposition have the property that if two points are in a cell, then these two points are in exactly the same set of slabs of P .

Theorem 2 then immediately implies that the repulsion kernel of P is the union of zero or more cells of the slab decomposition of P . This gives us the basis for an $O(n^2)$ time algorithm for finding the repulsion kernel. We start by constructing the slab decomposition. We can use topological sweep to compute a quad-edge data structure for the slab decomposition in $O(n^2)$ time [5, 4, 6].

► **Theorem 3.** *The repulsion kernel of a convex polygon can be computed in $O(n^2)$ time.*

Proof. We construct the slab decomposition. As we construct the decomposition, we augment each edge with information about which slab or slabs it borders and to which side of the edge said slabs are on. (An edge may border two slabs if the two slabs each have a defining ray that are collinear.). Choose an arbitrary cell c of the decomposition and determine how many slabs it is in. From this cell, perform a graph search on the dual of the decomposition. Each time we step over an edge, from one cell to another, during this search, we update in constant time the number slabs we are in, according to the information on the edge. We maintain a list of all cells where this value is one. At the end of the search, this list is the repulsion kernel. ◀

If we allow actuators to be located outside a polygon P , then every convex polygon is 1-gatherable.

► **Lemma 4.** *Every convex polygon is 1-gatherable from some point in the plane.*


Proof. If you go far enough away, you can always find a point that is not covered by any slab. For this point, there is only one accumulation point. Therefore, an activation of an actuator from this point moves all the particles to the accumulation point. ◀

Given the above, one may be tempted to believe that *every* convex polygon is 1-gatherable when the actuators are restricted to be inside the polygon. However, this is not always the case.

► **Lemma 5.** *For $k \geq 2$, the regular $(2k + 1)$ -gon P_{2k+1} is not 1-gatherable.*

Proof. Assume that the edge length of P_{2k+1} is 2, and that e_0 is oriented with direction 0 (horizontal on the bottom of the polygon). This is illustrated in Figure 4 for P_5 .

By Lemma 12, we need only show that P_{2k+1} is not 1-gatherable from its boundary. By symmetry, we need consider only e_{k+1} . The edge e_{k+1} starts at the top center of the polygon and proceeds downward to the left. The slab $S(e_0)$ contains the upper half of e_{k+1} , as the



■ **Figure 4** (a) $S(e_0)$ covers the top half of e_3 . (b) $S(e_1)$ covers the bottom half.

distance c (see figure) is greater than 1. (It is $1/\sin \alpha$, to be precise, where α is half the vertex angle, or $\frac{(2k-1)\pi}{4k+2}$.) Similarly, the slab $S(e_1)$ contains the bottom half of e_{k+1} .

Thus, each point of e_{k+1} is in $S(e_{k+1})$ and either $S(e_0)$ or $S(e_1)$ or both. Thus, by Theorem 2, the polygon is not 1-gatherable from any point of e_{k+1} .

The vertices are sometimes special cases, but here the vertex v_{k+1} (the top vertex of the polygon) is in $S(e_0)$, $S(e_k)$, and $S(e_{k+1})$, and thus the polygon is not 1-gatherable from there. By symmetry, it is not 1-gatherable from any vertex. ◀

In fact, some convex polygons may be ungatherable. It turns out that acute angles are a major impediment to gathering.

► **Lemma 6.** *A particle that is at an acute vertex v of P cannot be moved by an actuator activated at any point in $P \setminus v$.*

Proof. Given any point $p \in P \setminus v$, the acute vertex v is a local maximum with respect to distance since any point in P that is infinitesimally close to v is closer to p than v . ◀

This immediately implies the following.

► **Theorem 7.** *A convex polygon with three acute vertices is not k -gatherable for any $k > 0$.*

For the remainder of the paper, we only consider convex polygons with at most two acute vertices.

3 1-Gatherability

We have shown so far that not all convex polygons are 1-gatherable. We have also given a complete characterization of when a convex polygon is 1-gatherable by computing the repulsion kernel of a polygon in $O(n^2)$ time. This begs the question whether it is possible to find a point from which the polygon is 1-gatherable more efficiently, without having to compute the repulsion kernel. We answer this question in the affirmative by providing an $O(n)$ time algorithm. Before presenting the algorithm, we highlight some useful geometric properties.

► **Lemma 8.** *Let a be an accumulation point of an actuator activated at w in P . The line L that goes through a and is perpendicular to wa is a line of support of the polygon.*

Proof. Since a is an accumulation point, it is a local maximum of distance from w . Thus, the circle C with center w and radius aw encloses the polygon in the neighborhood of a . The line L is tangent to (outside of) C at a and thus locally supports the polygon at a . Since the polygon is convex, L also globally supports the polygon. ◀

We now show that we can restrict our attention to particles starting only on the boundary of P .

► **Lemma 9.** *An actuator in P that 1-gathers all the particles on ∂P also 1-gathers all particles in P .*

Proof. The activation of an actuator in P forces a particle p in the interior of P to move directly away from the actuator until it hits the boundary at some point b . Since there was a particle p' whose initial position is b , the particle p will follow the path of p' and stop at the same place p' stops. Thus, the location of p will always be accounted for by the position of p' . In other words, p is redundant and can be removed from the problem. ◀

We can take this a step further and show that particles located on the interior of edges are redundant.

► **Lemma 10.** *An actuator in P that 1-gathers all the particles on the vertices P also 1-gathers all particles on ∂P .*

Proof. The activation of an actuator in P forces a particle p in the interior of an edge of P to move along along the edge until it reaches a vertex v . There was a particle p' that started at v , and we can follow the proof of Lemma 9. ◀

The above lemmas show that particle movement can be restricted to the boundary. In fact, to solve the general problem, we only need to consider the problem where particles are only on vertices. We show a relationship between self-approaching paths and the path on the boundary followed by a particle under the influence of an actuator. Recall that a directed path Π is self-approaching if for any three consecutive points x, y, z on the path, we have the property that $|xz| \geq |yz|$ [7].

► **Lemma 11.** *If $\partial P(x, y)$ is self-approaching from x to y then activating an actuator at y sends all the particles on $\partial P(x, y)$ to x along the boundary.*


Proof. Let z be an arbitrary point on $\partial P(x, y)$. We observed that activating an actuator at y will move z along the boundary. We need to establish in which direction the particle will move. Since $\partial P(x, y)$ is self-approaching from x to y , we have that $|yz| \leq |yx|$. Therefore, the particle z will move to x since particles move in a direction to increase their distance from an actuator. ◀

Next, we show that if the repulsion kernel is not empty, then there is at least one point on the boundary that is in the repulsion kernel.


► **Lemma 12.** *Let P be a convex polygon that is 1-gatherable from a point w in the interior of P , and let a be the accumulation point for w . Let R be the ray from a through w , not including the point a . Then P is 1-gatherable from the point $w' = R \cap \partial P$, with a as its accumulation point (see Figure 5).*

Proof. By Theorem 2, the point of gatherability w is in one edge e 's perpendicular slab. Without loss of generality, we assume that e is horizontal at or below w (by rotation), that a is not to the right of w (by reflection), and that e is $e_0 = v_0v_1$ (by labelling). Let m be such that $a = v_m$. See Figure 6. Let p be the point on e which has a perpendicular through w . Note that p is a split point for w .

To show that P is 1-gatherable from the point w' on ∂P , by Lemma 10, it suffices to show that the particles located on the vertices of P move to one accumulation point with



■ **Figure 5** P is 1-gatherable from w' . Each slab in P is shown, with areas darkness corresponding to the number of slabs overlapping there.




■ **Figure 6** Some relevant points on the polygon.


the activation of an actuator at w' . We will show that this accumulation point is a . We assume without loss of generality that w' is located on the edge $e_k = [v_k v_{k+1}]$. Recall that if w' happens to be on v_k , then the placement of the actuator on w' means the particle located at w' is removed from consideration.

We begin with the claim that an accumulation point for w' is a . If this were not the case, then there would be a way to increase the distance from w' on the boundary in the neighborhood of a . By Lemma 8, there is a line L perpendicular to wa that is a line of support of P at a . By construction, L is perpendicular to $w'a$. Therefore, a is a local maximum with respect to w' , and thus is an accumulation point for w' . We will now show that particles located at all other vertices move to a when an actuator is activated at w' .

Since p is a split point for w , we have that upon activation of w , the particles on the vertices on $\partial P(a, p)$ move clockwise along the boundary to a . Similarly, the particles on the vertices on $\partial P(p, a)$ move counterclockwise along the boundary to a . By Theorem 2, this means that w is in all of the regions $R_{ccw}(e_1), R_{ccw}(e_2), \dots, R_{ccw}(e_{m-1})$ and w is in $R_{cw}(e_m), R_{cw}(e_{m+1}), \dots, R_{cw}(e_{n-1})$. See Figure 8.



■ **Figure 7** w is in the regions $R_{ccw}(e_0)$ to $R_{ccw}(e_{m-1})$.




■ **Figure 8** w is in the regions $R_{cw}(e_m)$ to $R_{cw}(e_n)$.

Since all of the slabs $S(e_m), S(e_{m+1}), \dots, S(e_{n-1})$ cross the chord aw' between a and w , we have that w' is also in $R_{cw}(e_m), R_{cw}(e_{m+1}), \dots, R_{cw}(e_{n-1})$. Thus, the vertices v_{m+1}, \dots, v_n move in a clockwise direction to a .

Now, we must show that the particles on vertices v_1, \dots, v_{m-1} also move to a . We first consider the vertices v_{k+1}, \dots, v_{m-1} . Again, since these vertices move counterclockwise when the actuator is activated at w , the slabs $S(e_j)$ for $k \leq j \leq m - 1$ cross the chord aw' between a and w . Therefore, none of them can contain w' . This implies that w' is in $R_{ccw}(e_{k+1}), \dots, R_{ccw}(e_{m-1})$.

We now show that the vertices v_1, \dots, v_k move in a clockwise direction to a . Consider the circle C centered at w and going through w' . This circle contains $\partial P[v_1, v_k]$ since particles on v_1, v_2, \dots, v_k move in a counterclockwise direction to a when an actuator is activated at w . It is strict containment as the particles always move away from w . Now consider the circle C' that has the chord aw' as diameter. Since a is the accumulation point for w , it is the farthest point from w . This implies that the center c of C' lies on the segment aw , with radius $|cw'|$. C' contains C since $|cw'| > |ww'|$. (Figure 9).

Let q be an arbitrary point in $\partial P(p, w')$. Since q is in the interior of C , we have that $\angle w'qa > \pi/2$. By convexity, we have that $\angle w'qp > \angle w'qa$. Consider the cone formed by the ray from q to w' and the ray at q that is an extension of the line through a and q . Since $\angle w'qp > \pi/2$, we have that the angle formed at this ray is strictly less than $\pi/2$ and $\partial P[q, w')$ is contained in the cone. Lemma 3 in [7] states that when $\partial P[q, w')$ is contained in a cone



■ **Figure 9** The circle C' contains the circle C and thus contains the boundary from p to w' .

at q with angle at most $\pi/2$ for every $q \in \partial P(p, w')$ then $\partial P(p, w')$ is self-approaching from p to w' . By Lemma 11, we have that an activation of an actuator at w' sends q clockwise around the boundary to p since $|pw'| \geq |qw'|$. Therefore, the vertices v_1, \dots, v_k move in a clockwise direction to a .

We have now shown the polygon is 1-gatherable from w' . ◀

As a consequence of the previous lemma, in order to tell if a polygon is 1-gatherable, it suffices to determine if it is 1-gatherable from the boundary. To do this in linear time, we employ an approach that resembles the rotating calipers algorithm to compute the diameter of a convex polygon [13]. In essence, for every point x on ∂P , we want to compute the first clockwise and first counterclockwise accumulation point. We do this in two steps. We compute all the counterclockwise accumulation points then compute the clockwise accumulation points. The algorithm to compute the counterclockwise accumulation points proceeds as follows. We start at the lowest point x of P and place the first horizontal caliper at x . We then walk around the boundary in counterclockwise direction until we find the counterclockwise accumulation point y for x . We place the second caliper at y such that it is perpendicular to xy . As x moves counterclockwise around P , there are two types of events. Either x moves to a new vertex or the caliper at y becomes coincident to an edge of P in which case y moves from one vertex to the next. There are a linear number of events that occur and by recording these events, when the calipers returns to its starting positions, we know the counterclockwise accumulation point for every point on the boundary of P . By repeating this in the clockwise direction, we find the clockwise accumulation points. For any point on the boundary of P , if its clockwise accumulation point is the same as its counterclockwise accumulation point, then the polygon is 1-gatherable from that point. We conclude this section with the following:

► **Theorem 13.** *We can determine if a convex n -vertex polygon is 1-gatherable in $O(n)$ time.*

Proof. Follows from Lemma 12 and the discussion above. ◀

4 2-Gatherability

In this section we prove that a convex polygon with at most two acute vertices is 2-gatherable. We then give an $O(n)$ algorithm to determine the location of the two actuators and the sequence of activation.

► **Theorem 14.** *If a convex polygon has two or fewer acute vertices, then it is 2-gatherable.*

Proof. Let $D(P)$ be the smallest disk enclosing polygon P with centre c . Either there are two vertices v_i and v_j of P that form a diameter of $D(P)$ or there are three vertices v_i, v_j , and v_k on $\partial D(P)$ such that c is in the interior of the triangle formed by the three vertices [12, 15]. We consider each case separately. Recall that by Lemma 9, we can assume that the particles are only located on the boundary of P .

Case 1: Two vertices v_i and v_j of P form a diameter of $D(P)$. In this case, we show that an actuator activated at vertex v_i results in all particles accumulating at v_j . Assume, without loss of generality, that v_i and v_j lie on a vertical line L with v_i below v_j . The two vertices partition the polygon boundary into two chains, $\partial P[v_i, v_j]$ which is to the right of L and $\partial P[v_j, v_i]$ which is to the left. We also assume that each chain consists of at least two edges, since otherwise, one of the chains is the edge $v_i v_j$ and trivially any particle on this edge moves to v_j when an actuator at v_i is activated. To complete the proof in this case, by Lemma 11, it suffices to show that both $\partial P[v_i, v_j]$ and $\partial P[v_j, v_i]$ are self-approaching curves from v_j to v_i .

Consider any point $x \in \partial P(v_i, v_j)$. Since x is in $D(P)$ strictly to the right of L we have that $\pi > \angle v_j x v_i \geq \pi/2$. Consider the cone formed by the intersection of the half-space bounded by the line through v_j and x that contains v_i and the half-space bounded by the line through v_i and x that does not contain v_j . This cone has angle at most $\pi/2$ and contains $\partial P[v_i, x]$. Since x is an arbitrary point on $\partial P(v_i, v_j)$, by Lemma 3 in [7], we have that $\partial P[v_i, v_j]$ is self-approaching from v_j to v_i . A similar argument shows that $\partial P[v_j, v_i]$ is also self-approaching from v_j to v_i .

Case 2: There are three vertices v_i, v_j , and v_k appearing in counter-clockwise order on $\partial D(P)$ such that c is in the interior of the triangle formed by the three vertices. Since there are at most two acute vertices, without loss of generality, assume that v_j is a polygon vertex with interior angle at least $\pi/2$. Reorient the polygon such that v_i is the lowest point. The polygonal chains $\partial P[v_i, v_j]$, $\partial P[v_j, v_k]$ and $\partial P[v_k, v_i]$ are self-approaching from v_j to v_i , v_k to v_j and v_k to v_i , respectively, by the same argument as the one used in Case 1. In fact, since c is strictly in the interior of the triangle formed by the three vertices, we have that the cones used to prove that the chains are self-approaching have an angle that is strictly less than $\pi/2$.

By placing a first active actuator on v_i , we have that all the particles on $\partial P(v_i, v_j]$ and all the particles on $\partial P[v_k, v_i)$ move onto $\partial P[v_j, v_k]$. Since $\partial P[v_j, v_k]$ is self-approaching from v_k to v_j , if we activated a second actuator at v_j then all the particles on this chain move to v_j 's accumulation point which would complete the proof. However, even though v_j is not acute, it may be the case that v_j is the counterclockwise accumulation point for v_i . This would prevent us from placing an actuator on v_j since after the activation of the first actuator on v_i , particles have accumulated on v_j . Recall that all subsequent placements of actuators must be on points in P that are free of particles. Since for every point x on $\partial P(v_j, v_k)$, $\angle v_j x v_k > \pi/2$, there must exist a point y on the edge $v_j v_{j-1}$ infinitesimally close to v_j such

that the $\angle yzv_k$ is still strictly greater than $\pi/2$ for every $z \in \partial P[v_j, v_k]$. This implies that $\partial P[y, v_k]$ is self-approaching from v_k to y . Thus, by Lemma 11, activating a second actuator at y , which is free of particles after the first activation, moves all the particles that have accumulated on $\partial P[v_j, v_k]$ to the counterclockwise accumulation point of y . ◀

References

- 1 Sang Won Bae, Chan-Su Shin, and Antoine Vigneron. Improved bounds for beacon-based coverage and routing in simple rectilinear polygons. *arXiv preprint arXiv:1505.05106*, 2015.
- 2 Michael Biro. *Beacon-based routing and guarding*. PhD thesis, State University of New York at Stony Brook, 2013.
- 3 Michael Biro, Justin Iwerks, Irina Kostitsyna, and Joseph SB Mitchell. Beacon-based algorithms for geometric routing. In *WADS*, pages 158–169. 2013.
- 4 Herbert Edelsbrunner and Leonidas J. Guibas. Topologically sweeping an arrangement. *J. Comput. Syst. Sci.*, 38(1):165–194, 1989.
- 5 Herbert Edelsbrunner and Leonidas J. Guibas. Corrigendum: Topologically sweeping an arrangement. *J. Comput. Syst. Sci.*, 42(2):249–251, 1991.
- 6 Leonidas J. Guibas and Jorge Stolfi. Primitives for the manipulation of general subdivisions and computation of voronoi diagrams. *ACM Trans. Graph.*, 4(2):74–123, 1985.
- 7 Christian Icking, Rolf Klein, and Elmar Langetepe. Self-approaching curves. *Math. Proc. Camb. Phil. Soc.*, 123(3):441–453, 1999.
- 8 Irina Kostitsyna, Bahram Kouhestani, Stefan Langerman, and David Rappaport. An optimal algorithm to compute the inverse beacon attraction region. In *SoCG*, 2018.
- 9 Bahram Kouhestani, David Rappaport, and Kai Salomaa. On the inverse beacon attraction region of a point. In *CCCG*, 2015.
- 10 Bahram Kouhestani, David Rappaport, and Kai Salomaa. The length of the beacon attraction trajectory. In *CCCG*, pages 69–74, 2016.
- 11 Bahram Kouhestani, David Rappaport, and Kai Salomaa. Routing in a polygonal terrain with the shortest beacon watchtower. *Comput. Geom.*, 68:34–47, 2018.
- 12 Nimrod Megiddo. Linear-time algorithms for linear programming in \mathbb{R}^3 and related problems. *SIAM J. Comput.*, 12(4):759–776, 1983.
- 13 Michael Shamos. *Computational Geometry*. PhD thesis, Yale Univeristy, 1978.
- 14 Thomas C Shermer. A combinatorial bound for beacon-based routing in orthogonal polygons. *arXiv preprint arXiv:1507.03509*, 2015.
- 15 Emo Welzl. Smallest enclosing disks (balls and ellipsoids). In *New Results and New Trends in Computer Science*, pages 359–370, 1991.