

Maintaining the Extent of a Moving Point Set*

P. K. Agarwal,¹ L. J. Guibas,² J. Hershberger,³ and E. Veach⁴

¹Center for Geometric Computing, Department of Computer Science, Duke University,
Box 90129, Durham, NC 27708, USA
pankaj@cs.duke.edu

²Computer Science Department, Stanford University,
Stanford, CA 94305, USA
guibas@cs.stanford.edu

³Mentor Graphics Corporation, 8005 SW Boeckman Road,
Wilsonville, OR 97070, USA
john_hershberger@mentorg.com

⁴Pixar Animation Studios, 1001 West Cutting Boulevard,
Richmond, CA 94804, USA

Abstract. Let S be a set of n moving points in the plane. We give new efficient and compact kinetic data structures for maintaining the diameter, width, and smallest area or perimeter bounding rectangle of S . If the points in S move with algebraic motions, these structures process $O(n^{2+\epsilon})$ events. We also give constructions showing that $\Omega(n^2)$ combinatorial changes are possible for these extent functions even if each point is moving with constant velocity. We give a similar construction and upper bound for the convex hull, improving known results.

1. Introduction

Let S be a set of n moving points in the plane. In this paper we investigate how to maintain various descriptors of the *extent* of S , such as diameter, width, smallest area or

* P.A. was supported in part by National Science Foundation Research Grants EIA-9870724, EIA-997287, and CCR-9732787, by Army Research Office MURI Grant DAAH04-96-1-0013, by a Sloan fellowship, by a National Science Foundation NYI award and matching funds from Xerox Corporation, and by a grant from the U.S.–Israeli Binational Science Foundation. L.G. and E.V. were supported in part by National Science Foundation Grant CCR-9623851, by US Army MURI Grant DAAH04-96-1-0007, and by a grant from Intel Corporation.

perimeter bounding rectangle. These extent measures indicate how spread out the point set S is. They are useful in various virtual reality applications such as clipping, collision checking, etc. As the points move continuously, the extent measure of interest (e.g., diameter) changes continuously as well, though its combinatorial realization (e.g., the pair of points defining the diameter) changes only at certain discrete times. Our approach is to focus on these discrete changes (or *events*) and track through time the combinatorial description of the extent measure of interest.

We do so within the framework of *kinetic data structures* (KDSs for short), as developed by Basch et al. [7] and further elaborated in Section 2. There are two notable and novel aspects of that framework. Firstly, while extensive work has been done on dynamic data structures in computational geometry [9], [10], this is all focused on handling insertions/deletions of objects and not on handling continuous changes. Kinetic data structures by contrast gain their efficiency by exploiting the continuity or coherence in the way the system state changes. Secondly, unlike Atallah's dynamic computational geometry framework [5], which was introduced to estimate the maximum number of combinatorial changes in a geometric configuration under predetermined motions in a certain class, the KDS framework is fully *on-line* and allows each object to change its motion at will, due to interactions with other moving objects, the environment, etc.

Section 3 presents new kinetic algorithms for diameter, width, and smallest enclosing rectangle in both the area and perimeter senses. If we assume that the points of S follow algebraic motions (defined below), then the number of events processed by each of our algorithms is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$; the constant of proportionality hidden in the big-O notation depends on ε . In particular, these bounds prove that all the extent measures mentioned can change combinatorially at most $O(n^{2+\varepsilon})$ times. A quadratic bound is natural for diameter, as it is defined by two points of the set S , but it is somewhat surprising for the other measures, as width is defined by three points, and the minimum bounding rectangles by four or five of the points. The data structures we give are efficient and compact in the KDS sense (see Section 2), though not local.

Section 4 is devoted to giving lower-bound constructions for these extent measures under *linear* motions: we show that diameter, width, and the two flavors of smallest bounding rectangle can all change $\Omega(n^2)$ times as each point in S moves continuously with constant velocity (different points may move with different velocities). Such lower bound constructions are much easier if we allow quadratic or other higher degree motions—the fact that the same lower bounds hold with linear motions is quite interesting. Our constructions employ a key component consisting of cocircular (or nearly cocircular) points that move on straight lines while maintaining their (near-) cocircularity. Finally, in Section 5 we give a similar construction showing that the convex hull of n points moving linearly in the plane can also change $\Omega(n^2)$ times. We also prove a slightly improved upper bound for the number of combinatorial changes to the convex hull. If any three points become collinear at most s times, then we prove an upper bound of $O(n\lambda_s(n))$, where $\lambda_s(n)$ is the maximum length of a Davenport–Schinzel sequence of order s with n symbols [15]; the previously known bound was $O(n\lambda_{s+2}(n))$. The bound improves to $O(n^2)$ for linearly moving points, which is optimal.

2. Kinetic Data Structure Preliminaries

Let $S = \{p_1, \dots, p_n\}$ be a set of n points in the plane, each of which is moving continuously. Let $p_i(t) = (x_i(t), y_i(t))$ denote the point p_i at time t , and let $S(t)$ denote the set S at time t . We say that the motion of S has degree d if, for all $1 \leq i \leq n$, $x_i(t)$ and $y_i(t)$ are polynomials of degree at most d . We call a motion of degree 1 *linear*; in this case each point of S moves along a straight line with fixed velocity. We say that the motion of S is *algebraic* if it is of degree d for some constant $d \geq 0$.

A KDS maintains a *configuration function* of continuously moving data (e.g., the diameter, width, etc., of moving points). It does so by maintaining a set of *certificates* that jointly imply the correctness of the computed configuration function. Each certificate is a geometric predicate on a constant number of data elements, for example, “points A and B are farther apart than points C and D .” The certificates are typically derived from a static algorithm for computing the configuration function. For example, the certificates for maintaining the diameter might include a set of distance comparisons establishing a partial order on the relevant pairwise distances, with a single maximum element.

The certificates are stored in a priority queue, ordered by the next time at which a certificate will be violated. Each data element has a *flight plan* that gives full or partial information about the current motion of the element, and these flight plans are used to compute the next violation time for each certificate. When the next violation time is reached, the algorithm removes the violated certificate from the queue and computes certificates for the new data configuration. Some number of certificates may have to be removed from the queue, and some number of new certificates added; see [7] for details. A KDS is called *responsive* if the time needed to update it after a certificate failure is polylogarithmic in the total number of data elements.

When a data element changes its flight plan, all the certificates in the priority queue that depend on it must have their times of next violation recomputed, and their positions in the queue must be updated. A KDS is called *local* if the number of certificates that depend on a single data element is polylogarithmic in the total number of data elements.

The violation of a certificate is called an *event*. *External events* cause the configuration function to change. *Internal events* do not affect the configuration function, but must be processed for the integrity of the data structure. We evaluate a KDS by counting events under the assumption that the data motions are *algebraic*. A KDS is called *efficient* if the worst-case number of total events (internal plus external) is asymptotically the same as, or only slightly larger than, the worst-case number of external events, under the assumption of algebraic motion.¹ A KDS is called *compact* if the number of certificates stored in the priority queue is roughly linear in the number of data elements.

Efficient, local, and compact KDSs are known for maintaining the convex hull and closest pair of points moving in the plane, and for computing the maximum of points moving along a line [7]. The data structure for computing the maximum is called a *kinetic tournament*. In the last few years much work has been done on KDSs for a wide range of problems; see [1]–[3], [6], [11], [12], and references therein.

¹ Some KDSs can be shown to be efficient for the larger class of *pseudoalgebraic* motions, which includes all motions such that the total number of events for any constant-size subset of data elements is $O(1)$. However, the KDSs of this paper require algebraic motion for provable efficiency.

3. Algorithms

In this section we present KDSs for maintaining three different versions of the extent of a planar point set: diameter, width, and minimum enclosing box. The diameter and width of a point set S in the plane can be computed by constructing $\text{conv}(S)$, computing the antipodal pairs of vertices,² and choosing an appropriate antipodal pair [16]. The smallest rectangle enclosing S can also be computed using a similar approach, though it requires a little more work (see Section 3.4 below).

Each of these extent versions is therefore maintained using the KDS for maintaining convex hulls of Basch et al. [7]. On top of that data structure we kinetize the rotating calipers algorithm [14], [16] to maintain the antipodal pairs of vertices. We finally store the antipodal pairs in a kinetic tournament, specialized to the desired extent attribute.

3.1. Antipodal Pairs

In this subsection we show how to maintain the set of antipodal pairs in the convex hull of a set of moving points. The boundary of $\text{conv}(S)$ can be partitioned into two convex chains, called *upper* and *lower* hulls, by its leftmost and rightmost vertices. Before we proceed, to simplify the presentation, we dualize the problem. In the dual, a point $p = (a, b)$ maps to a line $p^* : y = ax + b$ and a line $\ell : y = \alpha x + \beta$ maps to a point $\ell^* = (-\alpha, \beta)$. Let $S^* = \{p^* \mid p \in S\}$ be the set of lines dual to S . The *lower* (resp. *upper*) envelope of S^* is the boundary of the cell in the arrangement $\mathcal{A}(S^*)$ lying below (resp. above) all the lines of S^* . The lower (resp. upper) hull of S maps to the lower (resp. upper) envelope of S^* , with each vertex of the hull mapping to an edge of the envelope. Since the slope of a line maps to the x -coordinate of its dual point, the range of slopes of lines supporting a vertex p of the upper (resp. lower) hull maps to the x -interval of the corresponding edge of the upper (resp. lower) envelope.

Let L (resp. U) denote the projection of the lower (resp. upper) envelope of S^* onto the x -axis; the projection partitions the x -axis into intervals (see Fig. 1). Each vertex v of the lower (resp. upper) hull of S corresponds to an interval $I(v)$ of L (resp. U). Let Π denote the partition of the x -axis obtained by overlaying L and U . For each interval $\delta \in \Pi$, we can store the pair of vertices $\mu(\delta) = (a, b)$ such that $\delta \subseteq I(a) \cap I(b)$. It is easy to verify that if (p, q) is an antipodal pair of vertices, then $I(p)$ and $I(q)$ overlap, and therefore they can be obtained from Π .

To kinetize this static algorithm for computing antipodal pairs, besides the kinetic convex hull data structure, we use the vertices of Π as the certificates to guarantee the correctness of the current set of antipodal pairs, i.e., if $\xi_1 < \xi_2 < \dots$ are the vertices of Π , then we store (ξ_i, ξ_{i+1}) as certificates. Certificates are stored in a priority queue so that we can quickly determine the next certificate that will be violated. A certificate is violated when two vertices of Π become coincident (i.e., an interval of Π shrinks to a point) and their order exchanges. A constant number of certificates involving those intervals need to be updated to restore the certificates to correctness.

² A pair (p, q) of vertices of $\text{conv}(S)$ is *antipodal* if there are two parallel lines supporting $\text{conv}(S)$ at p and q .




Fig. 1. Convex hull and its dual.

Theorem 3.1. *The data structure for maintaining antipodal pairs is compact, efficient, and responsive. In particular, its size is $O(n \log n)$, it processes $O(n^{2+\epsilon})$ events in the worst case under algebraic motion, and it can be updated in $O(\log^2 n)$ time at each event.*

Proof. The underlying convex hull data structure is compact, responsive, and efficient [7]—it requires $O(n \log n)$ space, it processes $O(n^{2+\epsilon})$ events in the worst case under algebraic motion, and it can be updated in $O(\log^2 n)$ time at each event. It is also easy to see that the merged list structure is compact: its size is $O(n)$.

To prove that the data structure is efficient, we must bound the number of events in the merged list structure, under the assumption that the points move with algebraic motion. The key quantity to bound is the number of pairs of points that become antipodal over the life of the algorithm, since the list changes only when antipodal pairs change.

We extend the two-dimensional upper and lower envelopes of S^* into \mathbb{R}^3 by considering time as a static third dimension.³ More precisely, let $\ell_i(t) : y = a_i(t) \cdot x + b_i(t)$ be the dual of the point $p_i(t)$. Define $\gamma_i = \bigcup_{t \in \mathbb{R}} \ell_i(t)$; γ_i is an algebraic surface of constant degree. Let $\Gamma = \{\gamma_i \mid 1 \leq i \leq n\}$. The lower (resp. upper) envelope of S^* at time $t = t_0$ is the intersection (two-dimensional slice) of the lower (resp. upper) envelope of Γ with the plane $t = t_0$. Each combinatorial change in the upper (resp. lower) hull of $\text{conv}(S)$ corresponds to a vertex of the upper (resp. lower) envelope of Γ .

Let M (resp. M') be the minimization (resp. maximization) diagram of Γ (onto the xt -plane), and let \mathcal{M} be the overlay of the planar subdivisions M and M' . $L(t_0)$ (resp. $U(t_0)$, $\Pi(t_0)$) is the cross section of M (resp. M' , \mathcal{M}) at $t = t_0$. Each vertex of Π corresponds to an edge of \mathcal{M} . The x -order of two adjacent vertices ξ , ξ' of Π changes at t_0 if the corresponding edges of \mathcal{M} meet at a vertex of \mathcal{M} at $t = t_0$. The number of changes in the x -order of the vertices of Π , and thus the number of changes in the antipodal pairs, is bounded by the number of vertices in \mathcal{M}' . This quantity is bounded by $O(n^{2+\epsilon})$ [4].

³ The *lower envelope* of a set $\{f_1, \dots, f_n\}$ of d -variate functions is the function $F(\mathbf{x}) = \min_{1 \leq i \leq n} f_i(\mathbf{x})$. We do not distinguish between a function and its graph. The *minimization diagram* of a set of d -variate functions is the projection of the graph of its lower envelope onto the plane $x_{d+1} = 0$.

Finally, in order to prove that the antipodal-pair data structure is responsive, we must argue that only $O(1)$ antipodal pairs change at each event. There are two types of events in the antipodal-pair structure: events in the underlying convex hulls and the exchanges of adjacent vertices in Π . An exchange of adjacent vertices of Π affects only $O(1)$ antipodal pairs. If a point ceases to be a vertex of $\text{conv}(S)$, an interval of Π shrinks to a single point. On the other hand, if a new vertex v appears on $\text{conv}(S)$ between p and q , the vertex ξ of Π , corresponding to the common endpoint of $I(p)$ and $I(q)$, splits into two coincident vertices ξ^+ , ξ^- and a new singleton interval $I(v) = [\xi^-, \xi^+] = \xi$ is created. In each case, $O(1)$ certificates and antipodal pairs change. \square

Note that the data structure is *not* local: one point may belong to $O(n)$ antipodal pairs. It may be possible to achieve locality by making the pairing relationship more sophisticated, but this would require some additional insight.

3.2. Diameter

The diameter of a point set S is the maximum pairwise separation of two points in S . It is realized by a pair of antipodal vertices of $\text{conv}(S)$. A standard way of computing the diameter of S is to compute the convex hull, find all pairs of antipodal vertices, and then identify the pair with the maximum separation [14].

To kinetize this algorithm, we maintain the set of antipodal pairs using the algorithm described in the previous subsection. We then construct a kinetic tournament on these antipodal pairs (p, q) , with $d(p, q)$ as the key. Whenever an antipodal pair changes, we insert or delete an item from this tournament. Since we perform $O(n^{2+\epsilon})$ insertions and deletions in the tournament, the result by Basch et al. [8] shows that the diameter can be maintained efficiently. Theorem 4.1 of Section 4.1 shows that the total number of different diametral pairs is $\Omega(n^2)$, and hence our KDS is efficient. Finally, each event in the antipodal-pair structure affects $O(\log n)$ nodes of the kinetic tournament, so our data structure is responsive. We thus obtain the following.

Theorem 3.2. *The data structure for maintaining the diameter is compact, efficient, and responsive. In particular, its size is $O(n \log n)$, it processes $O(n^{2+\epsilon})$ events under algebraic motion, and it can be updated in $O(\log^2 n)$ time at each event.*

3.3. Width

The width of a point set S in \mathbb{R}^2 is the minimum separation of two parallel lines so that S lies in the strip bounded by the lines. It is well known that one of the lines contains an edge ab of the convex hull, that the other passes through a vertex v of the convex hull, and that (a, v) and (b, v) are antipodal pairs. We refer to v and ab as an antipodal edge–vertex pair, and say that v is antipodal to ab .

If (ab, v) is an antipodal edge–vertex pair, then the common endpoint of $I(a)$ and $I(b)$ lies in the interval $I(v)$. We can therefore find all antipodal edge–vertex pairs using Π , the overlay of L and U , defined in Section 3.1. To kinetize the data structure, we store all antipodal edge–vertex pairs (ab, v) in a tournament with $d(ab, v)$ as the key.

Theorem 3.3. *The data structure for maintaining the width is compact, efficient, and responsive. In particular, its size is $O(n \log n)$, it processes $O(n^{2+\varepsilon})$ events, and it can be updated in $O(\log^2 n)$ time at each event.*

Proof. Because the width data structure is virtually identical to the diameter structure, the bounds are the same—the worst-case time spent processing events is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. Theorem 4.7 of Section 4.2 shows that the number of combinatorial changes to the width is $\Omega(n^2)$ in the worst case, so the KDS is efficient. \square

Remarks. The similarity of the diameter and width data structures masks a rather surprising difference. We expect the number of combinatorial changes to the diameter to be $O(n^{2+\varepsilon})$, because there are only $O(n^2)$ pairs of points. Because the width is determined by *triples* of points—two edge endpoints and an opposing vertex—the naive bound on the number of combinatorial changes to the width is $O(n^{3+\varepsilon})$. However, the points of the triples are not independent, and our algorithm shows that the actual number of changes is only $O(n^{2+\varepsilon})$.

3.4. Extremal Boxes

A common way of reducing the complexity of spatial algorithms is to approximate a complex geometric object by a rectilinear box enclosing the object. Queries (e.g., intersection tests) are first performed on the box, then on the actual object only if the approximate test shows it to be necessary. In this way many queries on the complex object may be avoided. The approximating box is often chosen to be axis-aligned, but in situations in which a better approximation is desired, an arbitrarily oriented box may be computed.

Several criteria are used to choose the approximating box, depending on the application. For a set S of points in the plane, we may wish to compute a rectangle of minimum area or perimeter that encloses S ; we may wish to find the smallest square enclosing the set; or we may even wish to compute a rectangle of the maximum area or perimeter, each of whose sides touches at least one point of S . We can maintain an optimal enclosing box for any of these criteria with a single technique. For concreteness, we focus on maintaining the minimum-area rectangle enclosing a point set S in the plane. The following well-known lemma gives a simple algorithm for computing the minimum-area enclosing rectangle.

Lemma 3.4. *There exists a minimum-area rectangle R enclosing S so that each edge of R contains at least one point of S , and at least one edge of R contains an edge of $\text{conv}(S)$.*

In view of this lemma we maintain antipodal vertex–edge pairs of $\text{conv}(S)$. For each such pair (v, e) , we also maintain the vertices v_L, v_R that support the lines perpendicular to e . In the dual setting, we compute the lower and upper envelopes of S^* and their projections L and U . We rotate S by $\pi/2$ in the clockwise direction, and let L', U' be the x -projections of the lower and upper envelopes of the lines dual to this rotated set. We

merge L , U , L' , and U' ; let Π' be the resulting partition of the x -axis. An interval in Π' corresponds to a slope range in which the four convex hull vertices supported by lines parallel and perpendicular to the slope remain fixed. A rectangle satisfying the condition of Lemma 3.4 corresponds to an endpoint of Π' . For a vertex $a \in \Pi'$, let $R(a)$ be the rectangle defined by a . We compute $R(a)$ for each vertex a of Π' and choose the one with the minimum area.

To kinetize this algorithm, we construct a tournament on the set

$$\mathcal{R} = \{R(a) \mid a \text{ is a vertex of } \Pi'\}$$

and update this set whenever any antipodal pair changes. The KDS is essentially similar to those described above. The only difference is that instead of maintaining the overlay of two lists, namely, L and U , we now maintain the overlay of four lists.

Theorem 3.5. *The data structure for maintaining the minimum-area bounding rectangle is compact, efficient, and responsive. In particular, its size is $O(n \log n)$, it processes $O(n^{2+\varepsilon})$ events, and it can be updated in $O(\log^2 n)$ time at each event.*

Proof. Compactness follows as in Theorem 3.2. In order to prove the efficiency, we need to argue that $O(n^{2+\varepsilon})$ combinatorially distinct rectangles ever appear in the set \mathcal{R} . Let Γ be the same as defined in the proof of Theorem 3.1. Let S' be the set obtained by rotating S by $\pi/2$ in the clockwise direction with respect to the origin. Let Γ' be the corresponding set of surfaces for S' . We can argue that every change in \mathcal{R} corresponds to a vertex in the overlay \mathcal{M} of $M(\Gamma)$, $M(\Gamma')$, $M'(\Gamma)$, and $M'(\Gamma')$ —the minimization and maximization diagrams of Γ and Γ' , respectively. Since every vertex in this overlay is a vertex of the overlay of two of these four planar subdivisions, the total number of vertices in \mathcal{M} , and thus the number of changes in the set \mathcal{R} , is $O(n^{2+\varepsilon})$. \square

Remarks. (i) A similar approach can maintain the minimum-perimeter bounding rectangle of moving points.

(ii) The extremal box is determined by four or five points of the set, so a naïve bound on the number of combinatorially different boxes for points in algebraic motion would be $O(n^5)$. However, the preceding theorem shows that the actual number is only $O(n^{2+\varepsilon})$.

4. Lower Bounds with Linear Motion

In this section we give a collection of lower bounds on the number of combinatorial changes to the extent of a point set when each point moves linearly. Each of our constructions uses cocircular points whose linear motion maintains cocircularity. However, the lower bounds hold even if we perturb the points slightly to place them in general position.

Let C be the unit-radius circle centered at the origin o . Let $\bar{c}(\theta) = (\cos \theta, \sin \theta)$ be the point on C at angle θ from the x -axis. Suppose a point p moves linearly along a chord of C whose endpoints are $\bar{c}(\alpha)$ (at $t = 0$) and $\bar{c}(\alpha + \varphi)$ (at $t = 1$). That is,

$$p(t) = (1 - t)\bar{c}(\alpha) + t\bar{c}(\alpha + \varphi).$$



Fig. 2. Linear motion along equal-length chords preserves cocircularity.

Let

$$\theta(t) = \tan^{-1} \left((2t - 1) \tan \frac{\varphi}{2} \right).$$

For $t \in [0, 1]$, $\theta(t) \in [-\varphi/2, \varphi/2]$. The position of $p(t)$, for $t \in [0, 1]$, can be expressed in polar coordinates, $p(t) = (r_p(t), \theta_p(t))$, as follows:

$$r_p(t) = \frac{\cos(\varphi/2)}{\cos \theta(t)}, \quad \theta_p(t) = \alpha + \frac{\varphi}{2} + \theta(t). \tag{1}$$

We say that the motion of p is *clockwise* (resp. *counterclockwise*) if the distance from $p(0)$ to $p(1)$ along C in the clockwise (resp. *counterclockwise*) direction is at most π . Note that the initial position $p(0)$ does not appear in the expression for $r_p(t)$. If multiple points start on C and then move with the same speed along chords of C of the same length, then they remain cocircular throughout the whole motion. If all the motions are clockwise (or all counterclockwise), then the angular separation of each pair of points is constant: $\theta_p(t) - \theta_q(t)$ is just the difference of the initial angular positions of p and q . See Fig. 2.

4.1. Diameter

In this section we present an $\Omega(n^2)$ lower bound on the number of distinct diametral pairs that can appear in a set of n points under linear motion. We first discuss diametral pairs for points lying on two concentric circles, then specify a particular set of linearly moving points, and finally argue that our set has $\Omega(n^2)$ diametral pairs over time.

Suppose the points in S lie on two concentric circles C_1 and C_2 of radii r_1 and r_2 and centered at the origin o . Suppose that the points on C_i lie on an arc of length at most $\pi r_i/4$, for each $i \in \{1, 2\}$. If a point $p \in S \cap C_1$, another point $q \in S \cap C_2$, and the origin o are collinear, and o lies between p and q , then $\text{diam}(S) = d(p, q) = r_1 + r_2$. If there is only one such pair, then it is the unique diametral pair.

Suppose n is even and set $m = n/2$. S consists of m stationary points and m moving points. The set $P = \{p_0, \dots, p_{m-1}\}$ of stationary points is defined by

$$p_i = \bar{c} \left(\frac{\pi}{8m^2} i \right) \quad \text{for } 0 \leq i < m.$$

The set of moving points is $Q = \{q_0, \dots, q_{m-1}\}$. Each q_j moves linearly along a chord of C with an angular span of $\pi/4$:

$$q_j(t) = (1 - t) \bar{c} \left(\frac{7}{8} \pi + \frac{\pi}{8m} j \right) + t \bar{c} \left(\frac{9}{8} \pi + \frac{\pi}{8m} j \right) \quad \text{for } 0 \leq j < m.$$




Fig. 3. Lower bound construction for diameter.

Since all points in Q move with the same speed in the counterclockwise direction, they remain cocircular for all $t \in [0, 1]$ (Fig. 3).

We will prove that for every $0 \leq i < m$, there exist a time $t_{ij} \in [0, 1]$ at which (p_i, q_j) is the only diametral pair of S .

Theorem 4.1. *The diameter of the set of n linearly moving points described above is defined by $\Omega(n^2)$ different pairs of points during the time interval $t \in [0, 1]$.*

Proof. As noted above, the points of Q lie on a common circle whose radius varies with time. The angular position of $q_j(t)$ is

$$\theta_j(t) = \pi + \frac{\pi}{8m}j + \theta(t) \quad \text{for } \theta(t) \in \left[-\frac{\pi}{8}, \frac{\pi}{8}\right].$$

Thus Q lies in a constant-size angular range

$$\theta_{m-1}(t) - \theta_0(t) = \frac{\pi(m-1)}{8m} < \frac{\pi}{8}.$$

The angular range of P is less than $\pi/8m$. Points $p_i, q_j(t)$, and the origin are collinear if

$$\frac{\pi}{8m^2}i + \pi = \theta_j(t) = \pi + \frac{\pi}{8m}j + \theta(t),$$

that is, if

$$\theta(t) = \frac{\pi}{8m^2}(i - jm).$$

Since $0 \leq i < m$, each (i, j) pair determines a unique value of $\theta(t)$ in the interval $[-\pi/8, \pi/8]$, which corresponds to a unique value of $t \in [0, 1]$; call this value t_{ij} . Thus there are $m^2 = \Theta(n^2)$ distinct values $t_{ij} \in [0, 1]$ such that $(p_i, q_j(t_{ij}))$ is the unique diametral pair of the point set at time t_{ij} . \square

4.2. Width

In this section we present an $\Omega(n^2)$ lower bound on the number of distinct vertex triples that determine the width of a set of n linearly moving points. As before, the construction uses two sets of cocircular points, one stationary and one moving.

We define the *slab* of a line segment s to be the set of all points that project perpendicularly onto s . The basis of our construction is the following observation:

Observation 4.2. *Suppose that the width of a point set is determined by a convex hull edge e and a hull vertex v . Then v lies in the slab of e .*

Without loss of generality assume that $m = 4k$ for some integer $k \geq 2$. The stationary points of our set are

$$\begin{aligned}
 a &= \bar{c}\left(\frac{\pi}{8}\right), \\
 b &= \bar{c}\left(-\frac{\pi}{8}\right), \\
 p_i &= \bar{c}\left(-\pi + \frac{\pi(i - m/2)}{64m^2}\right) + \left(2 \cos \frac{\pi}{8}, 0\right) \quad \text{for } 0 \leq i \leq m.
 \end{aligned}$$

Let $P = \{p_0, \dots, p_m\}$. The set $Q = \{q_0, \dots, q_m\}$ of moving points is defined by

$$q_j(t) = (1 - t) \bar{c}\left(-\frac{3\pi}{32} + \frac{\pi}{8m}j\right) + t \bar{c}\left(-\frac{\pi}{32} + \frac{\pi}{8m}j\right) \quad \text{for } 0 \leq j \leq m.$$

See Fig. 4. The total number of points is $n = 2m + 4$. Q lies on the unit circle C centered at $(0, 0)$, and P lies on the unit circle C' centered at $(2 \cos(\pi/8), 0)$. The two circles intersect in arcs of length $\pi/4$. The intersection points are a and b ; P lies in a tight clump at the center of the left arc; Q lies in a $\pi/8$ sector of the right arc. Each point in Q moves along a chord of C of angular span $\pi/16$. All points in Q move with the same speed in the counterclockwise direction, so they remain cocircular for all $t \in [0, 1]$.




Fig. 4. Lower bound construction for width.

As in the previous subsection, $q_j(t)$ can be expressed in polar form as

$$r_j(t) = \frac{\cos(\pi/32)}{\cos \theta(t)},$$

$$\theta_j(t) = -\frac{\pi}{16} + \frac{\pi}{8m}j + \theta(t),$$

for all $\theta(t) \in [-\pi/32, \pi/32]$.

We prove below that for all $0 \leq i \leq m$ and $m/4 \leq j \leq 3m/4$, there is a time t_{ij} at which p_i, q_j, q_{j+1} is the only triple determining the width of S , which implies that the triple defining the width of S changes $\Omega(n^2)$ times. We need a sequence of technical lemmas to prove this claim.

Lemma 4.3. *For all $t \in [0, 1]$, all points of S appear on the boundary of the convex hull $\text{conv}(S)$ in the same order, namely $a, p_0, \dots, p_m, b, q_0, \dots, q_m$ form the vertices of $\text{conv}(S)$ in a counterclockwise order.*

Proof. By construction a, b, p_0, \dots, p_m appear on $\text{conv}(S)$ in that order for all $t \in [0, 1]$. Since the points in Q are always cocircular and remain in the same order, it suffices to show that the lines through the pairs $(b, q_0(t))$ and $(a, q_m(t))$ are supporting lines of $\text{conv}(S)$ for all $t \in [0, 1]$. Let $C(t)$ be the circle containing Q at time t . By construction, the radius $r(t)$ of $C(t)$ is at least $\cos(\pi/32)$. Let $\xi(t)$ be a point on $C(t)$ such that the segment $b\xi(t)$ is tangent to $C(t)$. Then the angle $\angle bo\xi(t) = \cos^{-1}(r(t)) \leq \pi/32$. (See Fig. 5(i).) Since $\angle boq_0(t) \geq \pi/32$, the segment from b to $q_0(t)$ passes through the interior of $C(t)$, which means that the line passing through b and $q_0(t)$ supports $\text{conv}(S)$. Similarly, we can show that the line containing the segment $aq_m(t)$ is a supporting line. \square

For each j , let

$$\beta_j = \frac{\pi}{16} - \frac{\pi}{8m}j,$$




Fig. 5. (i) The tangent covers a sector of at most $\pi/32$. (ii) Intersection segment of a slab with a vertical line.

so $\theta_j(t) = \theta(t) - \beta_j$. For $m/4 \leq j \leq 3m/4$, the point q_j intersects the x -axis when $\theta(t) = \beta_j$. Therefore the convex hull edge (q_j, q_{j+1}) intersects the x -axis in the interval $\theta(t) \in [\beta_{j+1}, \beta_j]$. That is,

$$\frac{\pi}{16} - \frac{\pi}{8m}(j+1) \leq \theta(t) \leq \frac{\pi}{16} - \frac{\pi}{8m}j. \tag{2}$$

The restriction on j ensures that this is a valid interval of $\theta(t)$. We define the interval $I_j = [\beta_{j+1} + \pi/32m, \beta_j - \pi/32m]$, i.e.,

$$\frac{\pi}{16} - \frac{\pi}{8m}\left(j + \frac{3}{4}\right) \leq \theta(t) \leq \frac{\pi}{16} - \frac{\pi}{8m}\left(j + \frac{1}{4}\right). \tag{3}$$

A simple calculation implies the following.

Lemma 4.4. *For any $m/4 \leq j < 3m/4$, $I_j \subseteq [-\pi/32, \pi/32]$.*

During the interval I_j , the edge q_jq_{j+1} intersects the x -axis, but $\theta_j(t) \leq -\pi/32m$ and $\theta_{j+1}(t) \geq \pi/32m$, i.e., neither q_j nor q_{j+1} lies “very close” to the x -axis. We will show that for any $\theta(t) \in I_j$, only q_jq_{j+1} and its antipodal vertex satisfy the condition of Observation 4.2, and that q_jq_{j+1} has $m + 1$ different antipodal vertices during the interval I_j .

We need the following simple lemma (see Fig. 5(ii)) to prove the first claim, whose proof is omitted.

Lemma 4.5. *Let C be the unit circle centered at the origin, and let c be a chord of C subtending an angle of $\theta < \pi$, whose midpoint lies in direction α from the origin. Then the slab of c intersects the vertical line $x = d$ in the y -interval*

$$\left[d \tan \alpha - \frac{\sin(\theta/2)}{\cos \alpha}, d \tan \alpha + \frac{\sin(\theta/2)}{\cos \alpha} \right].$$

Lemma 4.6. *For any $\theta(t) \in I_j$, $m/4 \leq j < 3m/4$, and an edge e in $\text{conv}(S)$, if $\text{slab}(e)$ contains a vertex of S antipodal to e , then $e = q_jq_{j+1}$.*

Proof. We first show that $W = \text{slab}(ap_0)$ does not contain any vertex of $\text{conv}(S)$ antipodal to edge ap_0 . By construction, W cannot contain b or any point of P , so it suffices to argue that W does not contain any point of Q antipodal to ap_0 . For convenience, we define $\bar{x} = 2 \cos(\pi/8) - 1$. Let ℓ and ℓ' be the vertical lines $x = 1$ and $x = \bar{x}$, respectively. Let y_0 be the y -coordinate of the lower endpoint of the intersection segment $W \cap \ell$. Let $p'_0 = \bar{c}(\pi/128m)$ be the reflection of p_0 , and let C' be the reflection of C with respect to the line ab . Then y_0 is the same as the y -coordinate of the lower endpoint of the segment $\text{slab}(ap'_0)$. See Fig. 6(i). Hence, by applying Lemma 4.5 with

$$\alpha = \frac{\pi/8 + \pi/128m}{2} \quad \text{and} \quad \theta = \left(\frac{\pi}{8} - \frac{\pi}{128m} \right),$$

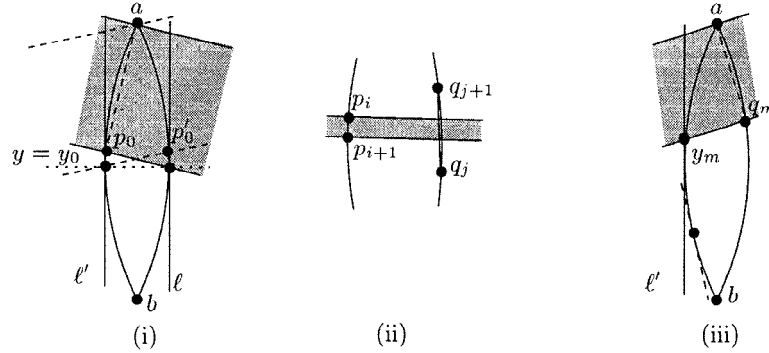


Fig. 6. Different cases for the proof of Lemma 4.6.

we obtain that

$$\begin{aligned}
 y_0 &\geq \bar{x} \tan \alpha - \frac{\sin(\theta/2)}{\cos \alpha} \\
 &= \bar{x} \tan \left(\frac{\pi}{16} + \frac{\pi}{256m} \right) - \frac{\sin(\pi/16 - \pi/(256m))}{\cos(\pi/16 + \pi/(256m))} \\
 &\geq (\bar{x} - 1) \tan \left(\frac{\pi}{16} \right) \\
 &= \left(2 \cos \frac{\pi}{8} - 2 \right) \tan \left(\frac{\pi}{16} \right) \\
 &\geq -0.0303.
 \end{aligned}$$

A point $q_k \in Q$ is antipodal to the edge ap_0 at time τ only if

$$\theta_{k-1}(\tau) < -\left(\frac{\pi}{16} + \frac{\pi}{256m} \right) < \theta_{k+1}(\tau).$$

Since $\theta_{k+1}(\tau) - \theta_{k-1}(\tau) = \pi/4m$, we have $\theta_{k+1}(\tau) < -\pi/16 + \pi/4m$, which implies that $\sin(\theta_{k+1}(\tau)) < 0$. The y -coordinate of q_k is smaller than that of q_{k+1} , and

$$\begin{aligned}
 y(q_{k+1}(\tau)) &= r_{k+1}(\tau) \sin(\theta_{k+1}(\tau)) \\
 &= \frac{\cos(\pi/32)}{\cos \theta(\tau)} \sin(\theta_{k+1}(\tau)) \\
 &< \cos \left(\frac{\pi}{32} \right) \sin(\theta_{k+1}(\tau)) \\
 &< \cos \left(\frac{\pi}{32} \right) \sin \left(-\frac{\pi}{16} + \frac{\pi}{4m} \right) \\
 &\leq \cos \left(\frac{\pi}{32} \right) \sin \left(-\frac{\pi}{32} \right) \\
 &\quad \text{(because } m \geq 8 \text{)} \\
 &= -\frac{1}{2} \sin \left(\frac{\pi}{16} \right) = -0.0980.
 \end{aligned}$$

This implies that when q_k is antipodal to the edge ap_0 , it does not lie in the strip W . Hence, $\text{slab}(ap_0)$ does not contain the vertex of $\text{conv}(S)$ antipodal to ap_0 during I_j . A similar argument proves that $\text{slab}(bp_m)$ does not contain the vertex antipodal to the edge bp_m .

Next, consider an edge $p_i p_{i+1}$ for any $0 \leq i < m$. The projection of each edge $p_i p_{i+1}$ onto the line ℓ is contained in the y -interval

$$\left[-\bar{x} \left(\tan \frac{\pi}{128m} + O \left(\frac{1}{m^2} \right) \right), \bar{x} \left(\tan \frac{\pi}{128m} + O \left(\frac{1}{m^2} \right) \right) \right];$$

see Fig. 6(ii).

More generally, the intersection of $\text{slab}(p_i p_{i+1})$ with the convex hull lies inside the y -interval

$$\left[-\tan \frac{\pi}{128m}, \tan \frac{\pi}{128m} \right].$$

For any $\theta(t) \in I_j$ and $1 \leq k \leq j$, $\theta_k(t) < -\pi/32m$. Therefore the y -coordinate

$$y(q_k(t)) = r_k(t) \sin(\theta_k(t)) \leq -\sin \left(\frac{\pi}{32m} \right) \cos \left(\frac{\pi}{32} \right) < -\tan \frac{\pi}{128m}.$$

Hence, q_1, \dots, q_j cannot lie in $\text{slab}(p_i p_{i+1})$ for any $\theta(t) \in I_j$. On the other hand, for all $k > j$, $\theta_k(t) \geq \pi/32m$ during $\theta(t) \in I_j$. A similar argument shows that $y(q_k(t)) > \tan(\pi/128m)$ during the interval I_j . Hence, none of q_{j+1}, \dots, q_m lies in $\text{slab}(p_i p_{i+1})$.

Next, consider the edge aq_m . The minimum possible y -coordinate, y_m , of the segment $\text{slab}(a, q_m(t)) \cap \ell'$ for all $t \in [0, 1]$ is at $t = 0$. At $t = 0$, the direction of the midpoint of aq_m from the origin is $5\pi/64$. Since aq_m subtends an angle of $3\pi/32$, by Lemma 4.5,

$$y_m = \bar{x} \tan \frac{5\pi}{64} - \frac{\sin(3\pi/64)}{\cos(5\pi/64)} > 0.0611.$$

None of the points in $Q \cup \{a, b\}$ can be antipodal to aq_m . Since the slope of the line supporting aq_m is always negative, the vertex p_i of $\text{conv}(S)$ antipodal to the edge aq_m has to lie below the x -axis. However, then $p_i \notin \text{slab}(aq_m)$. Similarly, the vertex antipodal to bq_0 does not lie in $\text{slab}(bq_0)$.

Finally consider an edge $q_k q_{k+1}$ for some $k \neq j$. For all $\theta(t) \in I_j$ and for any $k > j$, the edge $q_k q_{k+1}$ lies above the x -axis and the y -coordinate of the lower endpoint of the segment $\ell' \cap \text{slab}(q_k q_{k+1})$ is

$$y_k > \bar{x} \tan \frac{3\pi}{32m} - \frac{\sin(\pi/16m)}{\cos(3\pi/32m)} \geq \frac{0.0533}{m} > \sin \frac{\pi}{128m} > \frac{0.0245}{m}.$$

As in the previous case, any point of S antipodal to $q_k q_{k+1}$, for $k > j$, has to lie below the x -axis and thus does not lie in $\text{slab}(q_k q_{k+1})$. Likewise, none of $\text{slab}(q_k q_{k+1})$, for $k < j$, contains the point antipodal to the edge $q_k q_{k+1}$. Hence, $q_j q_{j+1}$ is the only edge of $\text{conv}(S)$ for which the vertex antipodal to the edge lies in its slab. This completes the proof of the lemma. \square

Since the orientation of the normal to edge q_jq_{j+1} at time t is $\theta_{j+1}(t) - \pi/16m$, by (3), it varies from $-\pi/32m$ to $\pi/32m$ during I_j . Therefore each p_i becomes antipodal to q_jq_{j+1} during the interval I_j . Combining Lemmas 4.4 and 4.6, we can prove that there are $m/2 = n/4 - 1$ convex hull edges, each of which in turn determines the width with $m + 1 = n/2 - 1$ different hull vertices during the time interval $[0, 1]$. Thus the triple of hull vertices determining the width changes $\Omega(n^2)$ times. We therefore conclude the following.

Theorem 4.7. *The width of the set of n linearly moving points described above is defined by $\Omega(n^2)$ different triples of points during the time interval $t \in [0, 1]$.*

4.3. Extremal Boxes

This section exhibits a configuration of n points in linear motion such that the minimum-area (or minimum-perimeter) enclosing rectangle undergoes $\Omega(n^2)$ combinatorial changes.

We assume $n = 8i$ for some integer $i \geq 1$. The construction involves $m = n/2$ closely spaced points $p_1, \dots, p_{n/2}$ that always lie on a circular arc of large radius r ($r \approx n^2/2$), rotating counterclockwise around the origin. There are an additional $n/2$ points $q_1, \dots, q_{n/2}$ near the origin so that, for $j = 1, \dots, n/8$, the convex hull of q_{4j-3}, \dots, q_{4j} is always a square Q_j . The side length of each Q_j is between 2 and 3, and each Q_j is slightly bigger than Q_{j-1} . The squares also have different orientations: the base of Q_j makes an angle of $j\theta$ with respect to the x -axis (where θ is a function of n). See Fig. 7.

The idea is that during $t \in [j - 1/8, j + 1/8]$, each p_i will cross L_j , the line through the origin with angle $j\theta$. We will show that the bounding box B has sides parallel to Q_j in this interval, and thus its combinatorial description depends on which of the p_i is farthest from the origin in the direction of L_j . Each of the p_i will become the farthest point in turn, as the points cross L_j , thereby producing $n/2$ combinatorial changes to the bounding box. This is repeated at times $t = 1, \dots, n/8$, yielding $\Theta(n^2)$ changes in total.

The following arguments make this construction more precise.

Lemma 4.8. *Let $P = \{p_1, \dots, p_m\}$ be a set of points on a circle of radius $r \geq 2$, centered at the origin such that $x(p_i) > 0$ and $y(p_i) \in [-1, 1]$ for all i , i.e., they lie on*




Fig. 7. The lower bound for minimum boxes.




Fig. 8. Illustration for Lemma 4.8.

a short arc near the positive x -axis. Let Q be the square whose vertices q_1, \dots, q_4 are the points $(\pm 1, \pm 1)$. Then the minimum-area (or minimum-perimeter) rectangle that encloses the P and Q is given by $B = [-1, x_{\max}] \times [-1, 1]$, where $x_{\max} = \max_{1 \leq i \leq n} x(p_i)$. See Fig. 8.

Proof. Let $p_{i_0} \in P$ be the point with the maximum x -coordinate. B contains p_{i_0} on its right side and its area is $2(1 + x_{\max})$. We will show that the smallest rectangle containing Q and p_{i_0} is B . Suppose, on the contrary, the smallest-area rectangle containing Q and p_{i_0} is B' . Let $0 \leq \alpha \leq \pi/2$ be the orientation of the long side of B' with respect to the x -axis. By Lemma 3.4, each side of B' contains at least one point among p_{i_0}, q_1, \dots, q_4 , and one of them contains two of these points. If $B \neq B'$, no edge of B' contains two q_i 's. Hence, one of them, say, the bottom one, contains a q_i and p_{i_0} . Then it is easily seen that p_{i_0} is one of the vertices of B' . A simple trigonometric calculation shows that the area of B' is

$$2(\sin \alpha + \cos \alpha) \left(\frac{x_{\max} - 1}{\cos \alpha} + 2 \cos \alpha \right) = 2(1 + \tan \alpha)(x_{\max} + \cos(2\alpha)).$$

Since $x_{\max} \geq 2$, the above expression is minimum when $\alpha = 0$, which shows that $B' = B$, as claimed. A similar argument shows that B is also the minimum-perimeter rectangle enclosing P and Q . \square

The construction now proceeds as follows. Set $\theta = 8/n^2$. For each $i = 1, \dots, n/8$, the vertices of Q_i are obtained by taking the square Q whose vertices are $(\pm 1, \pm 1)$, and rotating it by an angle of $i\theta$ counterclockwise around the origin. The size of Q_i varies with time, according to the scale factor

$$s_i(t) = 1 + 4\theta(2it - i^2) = 1 + 4\theta(t^2 - (i - t)^2). \tag{4}$$

In other words, each vertex of Q_i is on a linear trajectory through the origin, such that its distance from the origin at time t is $\sqrt{2}s_i(t)$. By our choice of θ , for every $1 \leq i \leq n/8$, $s_i(t) \in [1/2, 3/2]$ for $t \in [0, n/8]$. Note that, by the second equality in (4), Q_i is the largest square at time $t = i$. The following lemma is a slightly stronger version of this observation.

Lemma 4.9. For all $1 \leq i \neq j \leq n/8$ and for $t \in [j - 1/4, j + 1/4]$, $Q_i \subseteq Q_j$.




Fig. 9. Two nested squares.

Proof. Assume $i < j$. Let $\alpha = (j - i)\theta \leq 1/n$. The orientation of Q_i and Q_j differ by $(j - i)\theta = \alpha$. For $t \geq j - 1/4$, we have

$$\begin{aligned}
 \frac{s_j(t)}{s_i(t)} &= \frac{1 + 4\theta(2jt - j^2)}{1 + 4\theta(2it - i^2)} \\
 &= 1 + \frac{4\theta[2t(j - i) - (j^2 - i^2)]}{1 + 4\theta(2it - i^2)} \\
 &\geq 1 + \frac{8}{3}(j - i)\theta[2t - (j + i)] \\
 &> 1 + \frac{8}{3}\alpha \left[2\left(j - \frac{1}{4}\right) - (2j - 1)\right] \\
 &= 1 + \frac{4}{3}\alpha \\
 &\geq 1 + \sin \alpha \\
 &\geq \sin \alpha + \cos \alpha.
 \end{aligned}$$

Since Q_i and Q_j are centered at the origin and their orientation differs by α , a simple trigonometric calculation shows that if $s_j(t) \geq s_i(t)(\cos \alpha + \sin \alpha)$, then $Q_i(t) \subseteq Q_j(t)$; see Fig. 9. Hence, for $t \geq j - 1/4$, $Q_i(t) \subseteq Q_j(t)$. A similar argument shows that for $i > j$, $Q_i(t) \subseteq Q_j(t)$ for $t \leq j + 1/4$. This completes the proof of the lemma. \square

Finally, the set $P = \{p_1, \dots, p_{n/2}\}$ is placed on a circle C of radius $r = 4/\theta = n^2/2$, equally spaced along an arc of length $1/2$. At time $t = 0$, their y -coordinates lie in the range $[-1/4, 1/4]$. All points in P move counterclockwise along chords that subtend an angle of $\theta^* = (n/8)\theta = 1/n$, such that they intersect C at times $t = 0$ and $t = n/8$. The length of the chord is $2r \sin(\theta^*/2)$, which lies in the range $n/2 + O(1/n^3)$. For simplicity, we ignore the term $O(1/n^3)$. A closer look at the following argument shows that this extra term does not affect the analysis. So we assume that the chord length is $n/2$, which implies that the speed of each p_i is 4, and the y -coordinates of the points lie in the range $[4t - 1/4, 4t + 1/4]$ (noting that the chords are all nearly vertical for large n).

Let L_j be the line that makes an angle of $j\theta$ with the x -axis. Consider the time interval $[j - 1/8, j + 1/8]$. By Lemma 4.9, the smallest rectangle enclosing S is determined by Q_j and the point of P that is farthest from the origin along L_j . Now, at time $t = j - 1/8$, all points in P have y -coordinates in the range $[4j - 3/4, 4j - 1/4]$, and they lie below the intersection of L_j with C (at $y \approx 4j$). As each p_i crosses L_j , it is clearly the farthest point in direction L_j . By time $t = j + 1/8$, all the p_i have crossed L_j . We therefore

have had $n/2$ distinct minimum-area enclosing rectangles of S in this interval. We have established the following theorem.

Theorem 4.10. *The combinatorial description of the minimum-area (or minimum-perimeter) bounding box of n points moving linearly in the plane can change $\Omega(n^2)$ times.*

5. Kinetic Convex Hulls

In this section we give tight bounds on the number of combinatorial changes that may occur in the convex hull of points moving linearly in the plane. The lower bound construction is an easy application of the linear-motion-on-circles technique of Section 4. The upper bound is an improvement on the known bounds for points in general algebraic motion; when specialized to the case of linear motion, it shows that the convex hull may undergo $\Theta(n^2)$ combinatorial changes.

5.1. Lower Bound

We exhibit a configuration of $2n$ points in linear motion for which the convex hull undergoes $\Omega(n^2)$ combinatorial changes. This improves the lower bound example given by Sharir and Agarwal [15], which uses quadratic motions.

We define two convoys of oppositely moving points. The points always lie on a common circle (which varies in size), so all are on the convex hull, but their order along the circle changes.

Let

$$p_i(t) = (1 - t) \bar{c} \left(\frac{\pi}{4n} i \right) + t \bar{c} \left(\frac{\pi}{4} + \frac{\pi}{4n} i \right),$$

$$q_j(t) = (1 - t) \bar{c} \left(\frac{\pi}{4} + \frac{\pi}{8n^2} j \right) + t \bar{c} \left(\frac{\pi}{8n^2} j \right),$$

for $1 \leq i, j \leq n$. At any time $t \in [0, 1]$, all the p_i and q_j lie on a common circle with radius $r(t) = \cos(\pi/8) / \cos \theta(t)$, where $\theta(t) = \tan^{-1}((2t - 1) \tan(\pi/8))$. The angular position of $p_i(t)$ is

$$\theta(p_i, t) = \frac{\pi}{4n} i + \frac{\pi}{8} + \theta(t)$$

and the angular position of $q_j(t)$ is

$$\theta(q_j, t) = \frac{\pi}{8n^2} j + \frac{\pi}{8} - \theta(t).$$

Point p_i coincides with q_j at

$$\theta(t) = \frac{\pi}{8n} \left(\frac{j}{2n} - i \right).$$

Thus each (i, j) pair determines a unique $\theta(t) \in [-\pi/8, 0]$ at which p_i and q_j exchange on the convex hull. We have established the following theorem.

Theorem 5.1. *There is a set of n linearly moving points whose convex hull undergoes $\Omega(n^2)$ combinatorial changes as the points move.*

5.2. Upper Bound

We bound the number of combinatorial changes to the convex hull in terms of the number of times any three points become collinear. It is well known that if the point trajectories are algebraic of degree k , then three points become collinear at most $s = 2k$ times. The theorem below shows that in this case there are $O(n\lambda_{2k}(n))$ changes to the convex hull. This improves the bound of $O(n\lambda_{2k+2}(n))$ given in [15]. In particular, it implies that for linear motion the number of changes is $O(n^2)$, matching the lower bound of the preceding section.

Theorem 5.2. *Given n points moving in the plane such that no three points become collinear more than s times, the combinatorial description of their convex hull changes at most $O(n\lambda_s(n))$ times.*

Proof. Let the points be identified by integers, $P = \{1, \dots, n\}$, and define the *left-neighbor* function $l_i(t)$ as follows. If i does not belong to the convex hull at time t , then $l_i(t) = \varepsilon$. Otherwise, $l_i(t)$ is the point j on the convex hull that is adjacent to i in the counterclockwise direction.

For each i , let L_i be the sequence of values assumed by $l_i(t)$ as t ranges from $-\infty$ to ∞ . We remove all occurrences of ε from L_i , and replace any strings of identical symbols by a single occurrence, to yield a reduced sequence L_i^* .

First, we show that $\sum |L_i^*|$ is an upper bound on the number of changes to the convex hull (where $|S|$ denotes the length of a sequence S). We do this by charging each change to a unique symbol in some L_i^* .

The convex hull can change in only two ways: either a current vertex ceases to be a vertex of the hull, or a new vertex appears on the hull. Suppose that a current vertex i_1 is being deleted, and let i_0 and i_2 be its counterclockwise and clockwise neighbors just before the deletion. In this case, L_{i_2} will contain the substring i_1i_0 , and we charge the deletion to the symbol i_0 . Similarly, if i_1 was just inserted, then L_{i_2} contains the substring i_0i_1 , and we charge the insertion to the symbol i_1 . It is clear that no symbol is charged twice in this way, and that all charged symbols are present in the reduced sequences L_i^* (since each one is preceded by a different non- ε symbol).

Furthermore, each L_i^* is an $(n - 1, s)$ Davenport–Schinzel sequence. To see this, observe that when symbol j appears in L_i^* , it means that ij is an edge of the convex hull, and thus all triangles ijk have a signed area that is positive. Similarly, if k appears in L_i^* , then all triangles ikj have positive area (and all triangles ijk have negative area). Thus, if the alternation $j \cdots k$ appears in L_i^* , then the signed area of triangle ijk is zero at some intermediate time, implying a collinearity of i , j , and k .

Given that any three points are collinear at most s times, there are at most s alternations between any two symbols j and k . Thus, each L_i^* is a $(n - 1, s)$ Davenport–Schinzel sequence (L_i^* contains no repeated symbols by construction), and we have $\sum |L_i^*| \leq n\lambda_s(n)$. \square

6. Conclusions

In this paper we presented efficient and compact KDSs for maintaining the diameter, width, and a smallest enclosing rectangle of a planar point set. We also gave constructions showing that $\Omega(n^2)$ combinatorial changes for these extent functions are possible even under linear motion. We believe that our construction of linear motion that maintains cocircularity has other applications. For example, it was recently used for proving lower bounds on the number of changes in a triangulation of a planar point set [1]. We conclude by mentioning a few open problems:

- (i) Design an efficient KDS for maintaining the smallest enclosing disk of a point set in the plane.
- (ii) Design an efficient KDS for maintaining the convex hull of a point set in 3-space. We believe that the randomized KDSs proposed in [3] might be useful for this problem.
- (iii) Can the data structures presented in this paper be made local without affecting their efficiency or size?

Acknowledgments

The authors thank two anonymous referees for several useful comments.

References

1. P. K. Agarwal, J. Basch, M. de Berg, L. J. Guibas, and J. Hershberger. Lower bounds for kinetic planar subdivisions. *Discrete Comput. Geom.*, **24** (2000), 721–733.
2. P. Agarwal, D. Eppstein, L. J. Guibas, and M. Herzinger. Parametric and kinetic minimum spanning trees. In *Proc. 39th IEEE Sympos. Found. Comput. Sci.*, 1998, pp. 596–605.
3. P. K. Agarwal, J. Erickson, and L. J. Guibas. Kinetic binary space partitions for triangles. In *Proc. 9th ACM–SIAM Sympos. Discrete Algorithms*, 1998, pp. 107–116.
4. P. K. Agarwal, O. Schwarzkopf, and M. Sharir. The overlay of lower envelopes and its applications. *Discrete Comput. Geom.*, **15** (1996), 1–13.
5. M. J. Atallah. Some dynamic computational geometry problems. *Comput. Math. Appl.*, **11** (1985), 1171–1181.
6. J. Basch, J. Erickson, L. J. Guibas, J. Hershberger, and L. Zhang. Kinetic collision detection between two simple polygons. In *Proc. 10th ACM–SIAM Sympos. Discrete Algorithms*, 1999, pp. 327–336.
7. J. Basch, L. Guibas, and J. Hershberger. Data structures for mobile data. *J. Algorithms* **31** (1999), 1–28.
8. J. Basch, L. Guibas, and G. Ramkumar. Sweeping lines and line segments with a heap. *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, 1997, pp. 469–472.
9. Y.-J. Chiang and R. Tamassia. Dynamic algorithms in computational geometry. *Proc. IEEE*, **80** (1992), 1412–1434.
10. D. Eppstein. Average case analysis of dynamic geometric optimization. *Comput. Geom. Theory Appl.* **6** (1996), 45–68.
11. J. Erickson, L. J. Guibas, J. Stolfi, and L. Zhang. Separation-sensitive collision detection for convex objects. In *Proc. 10th ACM–SIAM Sympos. Discrete Algorithms*, 1999, pp. 102–111.
12. L. J. Guibas. Kinetic data structures: a state of the art report. In *Robotics: The Algorithmic Perspective* (P. K. Agarwal, L. E. Kavvaki, and M. T. Mason, eds.), A. K. Peters, Natick, MA, 1998, pp. 191–210.

13. M. E. Houle and G. T. Toussaint. Computing the width of a set. *IEEE Trans. Pattern Anal. Mach. Intell.* **10** (1988), 761–765.
14. F. P. Preparata and M. I. Shamos. *Computational Geometry*. Springer-Verlag, New York, 1985.
15. M. Sharir and P. K. Agarwal. *Davenport–Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
16. G. Toussaint. Solving geometric problems with the “rotating calipers.” In *Proc. IEEE MELECON '83*, 1983.

Received April 25, 2000, and in revised form September 25, 2000. Online publication May 4, 2001.