2013 IEEE 54th Annual Symposium on Foundations of Computer Science

On Kinetic Delaunay Triangulations: A Near Quadratic Bound for Unit Speed Motions

Natan Rubin
Freie Universität Berlin
Email: rubinnat@post.tau.ac.il

***Abstract*—Let P be a collection of n points in the plane, each moving along some straight line at unit speed. We obtain an almost tight upper bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the maximum number of discrete changes that the Delaunay triangulation $\mathrm{DT}(P)$ of P experiences during this motion. Our analysis is cast in a purely topological setting, where we only assume that (i) any four points can be co-circular at most three times, and (ii) no triple of points can be collinear more than twice; these assumptions hold for unit speed motions.**

***Keywords*-Delaunay triangulation; moving points; discrete changes; Voronoi diagram; combinatorial complexity**

I. Introduction

Delaunay triangulations. Let P be a finite set of points in the plane. Let $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ denote the Euclidean Voronoi diagram and Delaunay triangulation of P, respectively. The Delaunay triangulation consists of all triangles spanned by P whose circumcircles do not contain points of P in their interior. A pair of points $p, q \in P$ are connected by a Delaunay edge if there is a circle passing through p and q that does not contain any point of P in its interior.

Delaunay triangulations and their duals, Voronoi diagrams [10], are among the most extensively and longest studied constructs in computational geometry, with a wide range of applications. For a *static* point set P, both $\mathrm{DT}(P)$ and $\mathrm{VD}(P)$ have linear complexity and can be computed in optimal $O(n \log n)$ time. See [6], [12], [13] for surveys and a textbook on these structures. The problem has also been studied in the *dynamic* setting; see, e.g., [7].

The kinetic setting: Previous work. In many applications of Delaunay/Voronoi methods the points of the input set P are moving continuously. Interest in efficient maintenance of geometric structures under simple motion[1] of P goes back at least to Atallah [4], [5].

For the purpose of kinetic maintenance, Delaunay triangulations are nice structures, because, as mentioned above, they admit local certifications associated with individual triangles (namely, that their circumcircles be P-empty). This makes it simple to maintain $\mathrm{DT}(P)$ under point motion: an update is necessary only at critical times when one of these empty circumcircle conditions fails—this (typically) corresponds to co-circularities of certain subsets of four points, where the relevant circumcircle is P-empty. Whenever such an event, referred to as a *Delaunay co-circularity*, happens, a single edge flip easily restores Delaunayhood.[2] In addition, the Delaunay triangulation changes when some triple of points of P become *collinear* on the boundary of the convex hull of P; see below for details. Hence, the performance of any Voronoi- or Delaunay-based kinetic algorithm depends on the number of the above *discrete changes* (or *critical events*).

This paper studies the best-known formulation of the problem, in which each point of P moves along a straight line with unit speed; see [11], [13]. In this case, the previously best-known upper bound on the number of discrete changes in $\mathrm{DT}(P)$ is $O(n^3)$. In the more general setting, each point of P moves with so-called pseudo-algebraic motion of constant description complexity, implying that any four points are co-circular at most s times, and any triple of points can are collinear at most s' times, for some constants $s, s' > 0$. Given these (purely topological) restrictions, Fu and Lee [14], and Guibas et al. [15] established a roughly cubic upper bound of $O(n^2\lambda_{s+2}(n))$, where $\lambda_s(n)$ is the (almost linear) maximum length of an (n,s)-Davenport-Schinzel sequence [23]. A substantial gap exists between these near-cubic upper bounds and the best known quadratic lower bound [23]. Closing this gap has been in the computational geometry lore for many years, and is considered as one of the major open problems in the field. It is listed as Problem 2 in the TOPP project; see [11]. A recent work [21] by the author provides an almost tight bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, for a more restricted version of the problem, in which any four points can be co-circular at most *twice*.

In view of the very slow progress on the above general problem, a number of alternative structures, with (at most) near-quadratically many discrete changes, were studied; see, e.g., [2], [3], [8], [17].

Our result. We study the problem in a purely topological setup, where we assume that (i) any four points of P are co-circular at most three times during their (continuous) motion,

Work on this paper was supported by the Minerval Fellowship Program.

[1]The simplest way to define this motion is assume that each coordinate of each point $p = p(t)$ in P is a is fixed-degree polynomial in t.

[2]We assume, without loss of generality, that the trajectories of the points of P satisfy the standard general position assumptions; see, e.g., [21] for more details. In particular, no five points can become co-circular.

DOI 10.1109/FOCS.2013.62

and (ii) any three points of P can be collinear at most twice. For any point set P whose motion satisfies these two axioms, we derive a nearly tight upper bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, on the overall number of discrete changes experienced by $\mathrm{DT}(P)$. As is well known, these properties hold for points that move along straight lines with a common (unit) speed, so our near-quadratic bound holds in this case.

Proof ingredients. The majority of the discrete changes in $\mathrm{DT}(P)$ occur at moments t_0 when some four points $p, q, a, b \in P$ are co-circular, and the corresponding circumdisc contains no other points of P. We refer to these events as *Delaunay co-circularities*. Suppose that p, a, q, b appear along their common circumcircle in this order, so ab and pq form the chords of the convex quadrilateral spanned by these points. Right before t_0, one of the chords, say pq, is Delaunay and thus admits a P-empty disc whose boundary contains p and q. Right after time t_0, the edge pq is replaced in $\mathrm{DT}(P)$ by ab, an operation known as an *edge-flip*. Informally, this happens because the Delaunayhood of pq is violated by a and b: Any disc whose boundary contains p and q contains at least one of the points a, b. If pq does not re-enter $\mathrm{DT}(P)$ after time t_0, we can charge the event at time t_0 to pq, for a total of $O(n^2)$ such events. We thus assume that pq is again Delaunay at some moment $t_1 > t_0$.

It is insightful to note that one of the following always holds: either the Delaunayhood of pq is interrupted during (t_0, t_1) by at least k^2 pairs $u, v \in P$, or this edge can be made Delaunay throughout (t_0, t_1) by removal of at most $\Theta(k)$ points of P. In the former case, each violating pair u, v contributes during (t_0, t_1) either a co-circularity of p, q, u, v, or a collinearity in which one of u or v crosses pq.

Combinatorial charging. Our goal is to derive a recurrence for the maximum number $N(n)$ of such Delaunay co-circularities induced by any set P of n points (whose motion satisfies the above conditions). Notice that the number of *all* co-circularities, each defined by some four points of P, can be as large as $\Theta(n^4)$. The challenge is thus to show that the vast majority of co-circularity events are *not* Delaunay.

In Section II we study the set of all co-circularities that involve some disappearing Delaunay edge pq and some other pair of points of $P \setminus \{p, q\}$, and occur during the period (t_0, t_1) when pq is absentfrom $\mathrm{DT}(P)$. A co-circularity is called *k-shallow* if its circumdisc contains at most k points of P. If we find at least $\Omega(k^2)$ such k-shallow co-circularities involving p, q, and another pair of points, we can charge them for the disappearance of pq. We use the routine probabilistic argument of Clarkson and Shor [9] to show that the number of Delaunay co-circularities, for which this simple charging works, is $O\left(k^2 N(n/k)\right)$. Informally, this term that such Delaunay co-circularities contribute to the overall recurrence formula (see, e.g., [1] and [19]), yields a near-quadratic bound for $N(n)$. Similarly, if we find a "shallow" collinearity of p, q and another point (one halfplane bounded by the line of collinearity contains at most k points), we can charge the disappearance of pq to this collinearity. A combination of the Clarkson-Shor technique with the known near-quadratic bound on the number of topological changes in the convex hull of P (see [23, Section 8.6.1]) yields an additional near-quadratic term.

Probabilistic refinement. It thus remains to bound the number of the above Delaunay co-circularities, for which p and q participate in fewer shallow co-circularities and in no shallow collinearity during (t_0, t_1). In this case, we show, in what follows we refer to as the *Red-Blue Theorem* (or Theorem II.2), that one can restore the Delaunayhood of pq throughout (t_0, t_1) by removal of some subset A of at most $3k$ points of P. To bound the maximum number of such "non-chargeable" events, we incorporate them into more structured topological configurations (or, more precisely, processes), which are likely to show up (in the style of Clarkson and Shor) in a reduced triangulation $\mathrm{DT}(R)$, defined over a random sample $R \subset P$ of $\Theta(n/k)$ points.

For example, suppose that the above co-circularity at time t_0, is the *last* co-circularity of p, q, a, b. Then (at least) one of the points a or b must hit the edge pq before it re-enters $\mathrm{DT}(P)$ at time t_1. Clearly, the point which crosses pq, let it be a, must belong to A. Notice that the following two events occur simultaneously, with probability $\Omega\left(1/k^3\right)$: (1) the random sample R contains the crossing triple p, a, q, and (2) none of the points of $A \setminus \{a\}$ belong to R. In such case, we say that the edge pq undergoes a *Delaunay crossing by a* in the *refined* triangulation $\mathrm{DT}(R)$, which takes place during a certain subinterval $I \subset [t_0, t_1]$ (such that (i) a hits pq during I, (ii) $pq \in \mathrm{DT}(R)$ at the beginning and the end of I, and (iii) $pq \notin \mathrm{DT}(R)$ in the interior of I, but belongs to $\mathrm{DT}(R \setminus \{a\})$ throughout I). A symmetric argument applies if we encounter the *first* co-circularity of p, q, a, b. As argued in the predecessor paper [21] (and reviewed in Section IV-A), Delaunay crossings are especially nice objects due to their strict structural properties.

The roadmap. In Section III we show that the number of Delaunay co-circularities is dominated by the maximum possible number of Delaunay crossings. Notice the previously sketched argument (which appears in [21]) works only for the "first" and the "last" Delaunay co-circularities.

To extend the above reduction to the remaining, "middle" Delaunay co-circularities, we resort in Section III to a fairly simple argument, expressing the maximum possible number of such co-circularities in terms of the numbers of extremal Delaunay co-circularities and Delaunay crossings that arise in smaller-size subsets of P.

Informally, our goal is to show that, for an average pair (p, r), the point r is involved in "few" crossings of p-incident edges. To do so, we express, in Section IV, the number of Delaunay crossings in terms of the maximum number of certain *quadruples* $\sigma = (p, q, a, r)$, each composed of a pair of "consecutive" Delaunay crossings of p-adjacent edges pq and pa, by the same point r.

Unfortunately, the analysis of quadruples is fairly involved, so we only overview it in Section IV-B; the missing details are delegated to the full version of this paper.

To bound the number of quadruples, we resort to the routine "charge-or-refine" strategy (via our Red-Blue Theorem). This is done in several steps. At each stage we first try to dispose of as many quadruples as possible by charging each of them either to sufficiently many "shallow" co-circularities (or collinearities), or to one of the several kinds of "terminal" triples, for which we provide back in Section IV a direct quadratic bound on their number.

There are two main types of terminal triples (p, q, a). In one of them, we have a *double Delaunay crossing*—the point a crosses pq twice during the interval I. In the other the same triple performs two distinct "single" Delaunay crossings, where, say, a crosses pq during one crossing, and q crosses pa during the second one. In both cases the number of such triples is shown to be only $O(n^2)$.

Acknowledgements. I would like to thank my former Ph.D. advisor Micha Sharir whose dedicated support made this work possible. In particular, I would like to thank him for the insightful discussions, and, especially, for his help in the preparation and careful reading of this paper.

II. Geometric Preliminaries

Delaunay co-circularities. Let P be a collection of n points moving along (generic) pseudo-algebraic trajectories in the plane, so that any four points are co-circular at most *three* times, and any three points can be collinear at most *twice*.

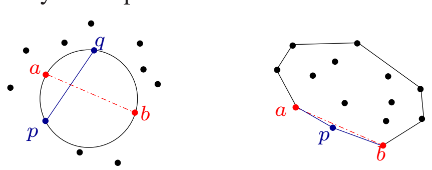

Figure 1. Left: A Delaunay co-circularity of a, b, p, q. An old Delaunay edge pq is replaced by the new edge ab. Right: A collinearity of a, p, b right before p ceases being a vertex on the boundary of the convex hull.

The Delaunay triangulation $\mathrm{DT}(P)$ changes at discrete time moments t_0 when one of the following events occurs:

(i) Some four points a, b, p, q of P become co-circular, so that the cicrumdisc of p, q, a, b is *empty*, i.e., does not contain any point of P in its interior. We refer to such events as *Delaunay co-circularities*. See Figure 1 (left). At each such co-circularity $\mathrm{DT}(P)$ undergoes an *edge-flip*, where an old Delaunay edge pq is replaced by the "opposite" edge ab.

(ii) Some three points a, b, p of P become collinear on the boundary of the convex hull of P. Assume that p lies between a and b. In this case, if p moves into the interior of the hull then the triangle abp becomes a new Delaunay triangle, and if p moves outside and becomes a new vertex, the old Delaunay triangle abp shrinks to a segment and disappears. See Figure 1 (right). The number of such collinearities is known to be at most nearly quadratic; see, e.g., [23, Section 8.6.1] and below.

In view of the above, it suffices to obtain a near-quadratic bound on the number of Delaunay co-circularities. Hence, the rest of this paper is devoted to proving the following main result:

Theorem II.1. *Let P be a collection of n points moving along pseudo-algebraic trajectories in the plane, so that (i) any four points of P are co-circular at most three times, and (ii) no triple of points can be collinear more than twice. Then P admits at most $O(n^{2+\varepsilon})$ Delaunay co-circularities, for any $\varepsilon > 0$.*

In what follows, we use $N(n)$ to denote the maximum possible number of Delaunay co-circularities among n points whose motion satisfies the above assumptions.

Shallow co-circularities. We say that a co-circularity event, where four points of P become co-circular, has *level* k if its corresponding circumdisc contains exactly k points of P in its interior. In particular, the Delaunay co-circularities have level 0. The co-circularities having level at most k are called *k-shallow*.

We can bound the maximum possible number of k-shallow co-circularities (for $k \geq 1$) in terms of the maximum number of Delaunay co-circularities in smaller-size point sets using the following fairly general argument of Clarkson and Shor [9]. Consider a random sample R of $\Theta(n/k)(< n/2)$ points of P and observe that any k-shallow co-circularity in P becomes a Delaunay co-circularity in R with probability $\Theta(1/k^4)$. (For this to happen, the four points of the co-circularity have to be chosen in R, and the at most k points of P inside the circumdisc must not be chosen; see [9] for further details.) Hence, the overall number of k-shallow co-circularities is $O(k^4 N(n/k))$.

Shallow collinearities. A collinearity of p, q, r is called *k-shallow* if the number of points of P to the left, or to the right, of the line through p, q, r is at most k. The argument of Clarkson and Shor implies, in a similar manner, that the number of such events, for $k \geq 1$, is $O(k^3 \mathsf{H}(n/k))$, where $\mathsf{H}(m)$ denotes the maximum number of discrete changes of the convex hull of an m-point subset of P. As shown, e.g., in [23, Section 8.6.1], $\mathsf{H}(m) = O(m^2\beta(m))$, where $\beta(\cdot)$ is an extremely slowly growing function.[3] Thus, the number of k-shallow collinearities is $O(kn^2\beta(n/k)) = O(kn^2\beta(n))$.

For every ordered pair (p, q) of points of P, denote by L_{pq} the line passing through p and q and oriented from p to q. Define L_{pq}^- (resp., L_{pq}^+) to be the halfplane to the left (resp., right) of L_{pq}. Notice that L_{pq} moves continuously with p and q. Note also that L_{pq} and L_{qp} are oppositely

[3]Specifically, $\beta(n) = \frac{\lambda_{s+2}(n)}{n}$, where s is the maximum number of collinearities of any fixed triple of points, and where $\lambda_{s+2}(n)$ is the maximum length of $(n, s+2)$-Davenport-Schinzel sequences [23].

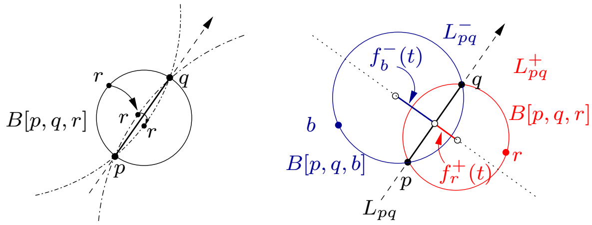

Figure 2. Left: The circumdisc $B[p,q,r]$ of p,q and r moves continuously as long as these three points are not collinear, and then flips over to the other side of the line of collinearity after the collinearity. Right: A snapshot at moment t. In the depicted configuration we have $f_b^-(t) < 0 < f_r^+(t)$.

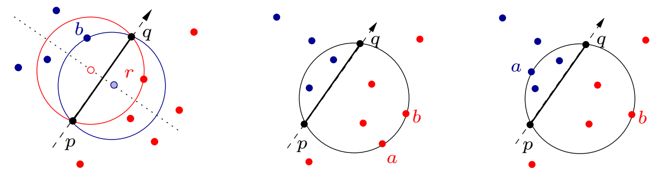

Figure 3. Left: A snapshot at time t. The red and blue envelopes E^+, E^- coincide with the functions f_r^+, f_b^-, respectively. The edge pq is not a Delaunay edge because $E^+(t)$ (the hollow center) is smaller than $E^-(t)$ (the shaded center). Center and right: Red-red and red-blue co-circularities.

oriented and that $L_{pq}^+ = L_{qp}^-$ and $L_{pq}^- = L_{qp}^+$. We also orient the edge pq connecting p and q from p to q, so that the edges pq and qp have opposite orientations.

Any three points p, q, r span a circumdisc $B[p,q,r]$ which moves continuously with p, q, r as long as p, q, r are not collinear. See Figure 2 (left). When p, q, r become collinear, say, when r crosses pq from L_{pq}^- to L_{pq}^+, the disc $B[p,q,r]$ changes instantly from being all of L_{pq}^+ to all of L_{pq}^-.

The red-blue arrangement. Following [15] and [21], we use the so called red-blue arrangement to facilitate the analysis of co-circularities whose corresponding discs touch the same two points $p, q \in P$.

For a fixed ordered pair $p, q \in P$, we call a point a of $P \setminus \{p,q\}$ *red* (with respect to the oriented edge pq) if $a \in L_{pq}^+$; otherwise it is *blue*.

We define, for each $r \in P \setminus \{p,q\}$, a pair of partial functions f_r^+, f_r^- over the time axis as follows. If $r \in L_{pq}^+$ at time t then $f_r^-(t)$ is undefined, and $f_r^+(t)$ is the signed distance of the center c of $B[p,q,r]$ from L_{pq}; it is positive (resp., negative) if c lies in L_{pq}^+ (resp., in L_{pq}^-). A symmetric definition applies when $r \in L_{pq}^-$. Here too $f_r^-(t)$ is positive (resp., negative) if the center of $B[p,q,r]$ lies in L_{pq}^+ (resp., in L_{pq}^-). We refer to f_r^+ as the *red function* of r (with respect to pq) and to f_r^- as the *blue function* of r. Note that at all times when p, q, r are not collinear, exactly one of f_r^+, f_r^- is defined. See Figure 2 (right).

Let E^+ denote the lower envelope of the red functions, and let E^- denote the upper envelope of the blue functions. The edge pq is a Delaunay edge at time t if and only if $E^-(t) < E^+(t)$. Indeed, any disc whose bounding circle passes through p and q which is centered anywhere in the interval $(E^-(t), E^+(t))$ along the bisector of pq is empty at time t. If pq is not Delaunay at time t, there is a "bichromatic" pair $r, b \in P$ such that $f_r^+(t) < f_b^-(t)$. In such a case, we say that the Delaunayhood of pq is *violated* by r and b.

Hence, at any time when the edge pq joins or leaves $\mathrm{DT}(P)$, via a Delaunay co-circularity involving p, q, and two other points of P, we have $E^-(t) = E^+(t)$. In this case the two other points, a, b, are such that one of them, say a, lies in L_{pq}^+ and b lies in L_{pq}^-, and $E^+(t) = f_a^+(t), E^-(t) = f_b^-(t)$.

Let $\mathcal{A} = \mathcal{A}_{pq}$ denote the arrangement of the $2n-4$ functions $f_r^+(t), f_r^-(t)$, for $r \in P \setminus \{p,q\}$, drawn in the parametric (t,ρ)-plane, where t is the time and ρ measures signed distance to the midpoint of pq along the perpendicular bisector of pq. We label each vertex of $\mathcal{A}$ as red-red, blue-blue, or red-blue, according to the colors of the two functions meeting at the vertex. An intersection between two red functions f_a^+, f_b^+ corresponds to a co-circularity event which involves p, q, a and b, occurring when both a and b lie in L_{pq}^+. Similarly, an intersection of two blue functions f_a^-, f_b^- corresponds to a co-circularity where both a and b lie in L_{pq}^-, and an intersection of a red fuction f_a^+ and a blue function f_b^- represents a co-circularity where $a \in L_{pq}^+$ and $b \in L_{pq}^-$. We label these co-circularities as red-red, blue-blue, and red-blue, depending on the respective colors of a and b. See Figure 3 (center and right).

Note that in any co-circularity of four points of P there are exactly two pairs (namely, the opposite pairs) with respect to which the co-circularity is red-blue, and four pairs (the adjacent pairs) with respect to which the co-circularity is "monochromatic". When the co-circularity is Delaunay, the two pairs for which the co-circularity is red-blue are those that enter or leave the Delaunay triangulation $\mathrm{DT}(P)$ (one pair enters and one leaves). The Delaunayhood of pairs for which the co-circularity is monochromatic is not affected by the co-circularity, which appears in the corresponding arrangement as a *breakpoint* of either $E^+(t)$ or $E^-(t)$.

The following useful result on $\mathcal{A}_{pq}$ was established in [21]. (Note that the theorem holds for all pseudo-algebraic motions of constant description complexity, and the constants in the $O(\cdot)$ and $\Omega(\cdot)$ notations do not depend on k.)

Theorem II.2 (Red-blue Theorem)**.** *Let P be a collection of n points moving in the plane as described above. Suppose that an edge pq belongs to $\mathrm{DT}(P)$ at (at least) one of the two moments t_0 and t_1, for $t_0 < t_1$. Let $k > 12$ be some sufficiently large constant. Then one of the following conditions holds:*

(i) There is a k-shallow collinearity which takes place during (t_0, t_1), and involves p, q and another point r.

(ii) There are $\Omega(k^2)$ k-shallow red-red, red-blue, or blue-blue co-circularities (with respect to pq) which occur during (t_0, t_1).

(iii) There is a subset $A \subset P$ of at most $3k$ points whose removal guarantees that pq belongs to $\mathrm{DT}(P \setminus A)$

throughout (t_0, t_1).

III. From Delaunay Co-Circularities to Delaunay Crossings

As before, $N(n)$ denotes the maximum possible number of Delaunay co-circularities that can arise in a set P of n points moving as above in $\mathbb{R}^2$. In this section we introduce the notion of a Delaunay crossing, which plays a central role both in this paper and in its predecessor [21], and express the above quantity $N(n)$ in terms of the number of Delaunay crossings that can arise in smaller sets of moving points.

Delaunay crossings. A *Delaunay crossing* is a triple $(pq, r, I = [t_0, t_1])$, where $p, q, r \in P$ and I is a time interval, such that

1) pq leaves $\mathrm{DT}(P)$ at time t_0, and returns at time t_1 (and pq does not belong to $\mathrm{DT}(P)$ during (t_0, t_1)),
2) r crosses the segment pq *at least* once (and at most twice, by assumption) during I, and
3) pq is an edge of $\mathrm{DT}(P \setminus \{r\})$ during I (i.e., removing r restores the Delaunayhood of pq during the entire time interval I).

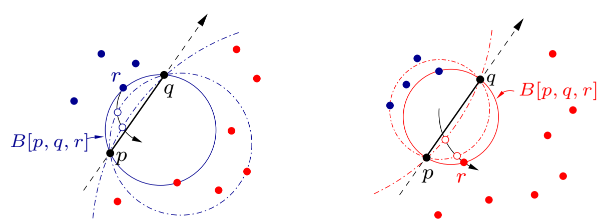

Figure 4. A Delaunay crossing of pq by r from L^-_{pq} to L^+_{pq}. Several snapshots of the continuous motion of $B[p,q,r]$ before and after r crosses pq are depicted (in the left and right figures, respectively).

It is easy to see that the third condition is equivalent to the following condition, expressed in terms of the red-blue arrangement $\mathcal{A}_{pq}$ associated with pq: The point r participates only in red-blue co-circularites during the interval I, and these are the only red-blue co-circularities that occur during I. More specifically, note that r is red during some portion of I and is blue during the complementary portion (both portions are not necessarily connected). During the former portion the graph of f^+_r coincides with the red lower envelope E^+ (otherwise $E^+(t) < E^-(t)$ would hold sometime during I even after removal of r), so it can only meet the graphs of blue functions. Similarly, during the latter portion f^-_r coincides with the blue upper envelope E^-, so it can only meet the graphs of red functions. When passing from the former portion to the latter, f^+_r goes down to $-\infty$, meeting all blue functions below it, and then it is replaced by f^-_r, which goes down from ∞. See Figure 4 for a schematic illustration of this behavior.

Notice that no points, other than r, cross pq during I (any such crossing would clearly contradict the third condition at the very moment when it occurs). Moreover, r does not cross L_{pq} outside pq during I; otherwise pq would belong to $\mathrm{DT}(P)$ when r belongs to $L_{pq} \setminus pq$.

Types of Delaunay co-circularities. We say that a co-circularity event at time t_0 involving a, b, p, q has *index* $1, 2$, or 3 if this is, respectively, the first, the second, or the third co-circularity involving a, b, p, q. A co-circularity is *extremal* if its index is 1 or 3, and the co-circularities with index 2 are referred to as *middle* co-circularities.

Let $C(n)$ denote the maximum possible number of Delaunay crossings that can arise in a set of n moving points $\mathbb{R}^2$. To bound $N(n)$ in terms of $C(n)$ (or, more precisely, in terms of $C(m)$, for some $m \le n$), we first develop a recurrence which expresses the maximum possible number $N_E(n)$ of extremal Delaunay co-circularities in P in terms of $C(n/k)$. (In [21], there were no middle co-circularities, so the same argument worked for *all* Delaunay co-circularities.) We then express the maximum possible number $N_M(n)$ of middle Delaunay co-circularities in P in terms of $C(n/k)$ and $N_E(n/k)$ (where k is any sufficiently large parameter.)

The number of extremal co-circularities. Consider a Delaunay co-circularity event at time t_0 at which an edge pq of $\mathrm{DT}(P)$ is replaced by another edge ab, because of an extremal red-blue co-circularity with respect to pq and ab. Without loss of generality, assume that the co-circularity of p, q, a, b has index 3 (the symmetric case of index 1 is handled by reversing the direction of the time axis).

There are at most $O(n^2)$ such events for which the vanishing edge pq never reappears in $\mathrm{DT}(P)$, so we focus on the Delaunay co-circularities (of index 3) whose corresponding edge pq rejoins $\mathrm{DT}(P)$ at some future moment $t_1 > t_0$. Note that at least one of the two other points a, b involved in the co-circularity at time t_0 must cross pq at some time between t_0 and t_1. Indeed, otherwise p, q, a and b would have to become co-circular again, in order to "free" pq from its non-Delaunayhood, which is impossible since our co-circularity has index 3. More generally, we have the following lemma, whose easy proof is illustrated in Figure 5:

Lemma III.1. *Assume that the Delaunayhood of pq is violated at time t_0 (or rather right after it) by the points $a \in L^-_{pq}$ and $b \in L^+_{pq}$. Furthermore, suppose that pq re-enters $\mathrm{DT}(P)$ at some future time $t_1 > t_0$. Then at least one of the followings occurs during $(t_0, t_1]$:*

(1) The point a crosses pq from L^-_{pq} to L^+_{pq}.
(2) The point b crosses pq from L^+_{pq} to L^-_{pq}.
(3) The four points p, q, a, b are involved in a red-blue co-circularity.

Clearly, the third scenario is not possible if the co-circularity at time t_0 has index 3. A symmetric version of Lemma III.1 applies if the Delaunayhood of pq is violated right *before* time t_0 by a and b, and this edge is Delaunay at an *earlier* time $t_1 < t_0$.

Notice, however, that the points of P can define $\Omega(n^3)$

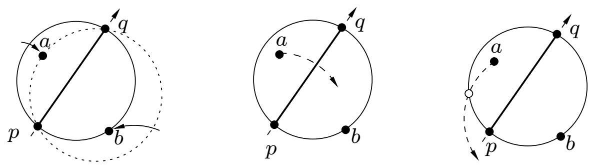

Figure 5. Proof of Lemma III.1. The Delaunayhood of pq remains violated by a and b after time t_0 as long as none of a, b hits L_{pq}, and a remains in $B[p,q,b] \cap L^-_{pq}$ (left). Hence, pq can become free from its violation only after being hit by a and/or b (center), or after an additional co-circularity of p, q, a, b (right).

collinearities, so a naive charging of extremal Delaunay co-circularities to collinearities of type (1) or (2) in Lemma III.1 will not lead to a near-quadratic upper bound. Before we get to this (major) issue in our analysis, we begin by laying down the infrastructure of our charging scheme, similar to the one used in [21].

We fix some sufficiently large constant parameter $k > 12$ and apply Theorem II.2 to the edge pq over the interval (t_0, t_1) of its absense from $\mathrm{DT}(P)$. Assume first that one of the conditions (i) or (ii) of the theorem holds, so we can charge the co-circularity of $p, q, a,$ and b either to $\Omega(k^2)$ k-shallow co-circularities (each involving $p, q,$ and some two other points of P), or to a k-shallow collinearity (involving $p, q,$ and some third point of P). As argued in Section II, the overall number of k-shallow co-circularities is $O(k^4 N(n/k))$. Each k-shallow co-circularity is charged by only $O(1)$ Delaunay co-circularities in this manner,[4] and it has to "pay" only $O(1/k^2)$ units every time it is charged. Similarly, as already argued, the number of k-shallow collinearities is $O(kn^2\beta(n))$, and each such collinearity is charged by at most $O(1)$ Delaunay co-circularities. Hence, there are at most $O(k^2 N(n/k) + kn^2\beta(n))$ Delaunay co-circularities for which one of the conditions (i) or (ii) holds.

Assume then that condition (iii) holds for our co-circularity. By assumption, there is a set A of at most $3k$ points (necessarily including at least one of a or b) whose removal ensures the Delaunayhood of pq throughout (t_0, t_1). By Lemma III.1, at least one the two points a, b, let it be a, crosses pq during (t_0, t_1). As we will shortly show, in the reduced triangulation $\mathrm{DT}(P \setminus A \cup \{a\})$, the collinearity of p, q and a can be turned into a Delaunay crossing.

We now express the number of remaining Delaunay co-circularities of index 3 in terms of the maximum possible number of Delaunay crossings. Recall that for each such co-circularity there is a set A of at most $3k$ points whose removal restores the Delaunayhood of pq throughout $[t_0, t_1]$. We can assume that a hits pq during $(t_0, t_1]$, so $a \in A$.

We sample at random a subset $R \subset P$ of n/k points, and notice that the following two events occur simultaneously with probability at least $\Omega(1/k^3)$: (1) the points p, q, a belong to R, and (2) none of the points of $A \setminus \{a\}$ belong to R. Since a crosses pq during $[t_0, t_1]$, and pq is Delaunay at times t_0 and time t_1, the sample R induces at least one Delaunay crossing (pq, a, I), for some time interval $I \subset [t_0, t_1]$. (If a crosses pq twice, we have either two separate Delaunay crossings, which occur at disjoint sub-intervals of (t_0, t_1), or only one Delaunay crossing, during which a crosses pq twice. This depends on whether pq manages to become Delaunay in $\mathrm{DT}(R)$ in between these crossings.) We charge the disappearance of pq from $\mathrm{DT}(P)$ to this crossing and note that the charging is unique (i.e., every Delaunay crossing (pq, a, I) in $\mathrm{DT}(R)$ is charged by at most one disappearance t_0 of the respective edge pq from $\mathrm{DT}(P)$, which is *last* such disappearance of pq before a hits pq in I). Hence, the number of Delaunay co-circularities of this kind is bounded by $O(k^3 C(n/k))$, where $C(n)$ denotes, as above, the maximum number of Delaunay crossings induced by any collection P of n points whose motion satisfies the above assumptions.

We thus obtain the following recurrence for the maximum possible number $N_E(n)$ of extremal Delaunay co-circularities:

$$N_E(n) = O\left(k^3 C(n/k) + k^2 N(n/k) + kn^2\beta(n)\right). \quad (1)$$

The number of middle Delaunay co-circularities. We now develop a recurrence that expresses the number of middle Delaunay co-circularities in terms of $C(n/k)$, $N_E(n/k)$, and $N(n/k)$, for an appropriate constant parameter k.

Consider such a middle co-circularity event at time t_0, when an edge pq of $\mathrm{DT}(P)$ is replaced by another edge ab. As in the previous case, there are at most $O(n^2)$ such events for which the vanishing edge pq never reappears in $\mathrm{DT}(P)$, so we focus on middle Delaunay co-circularities whose corresponding edge pq rejoins $\mathrm{DT}(P)$ at some future moment $t_1 > t_0$.

Once again, we fix a sufficiently large constant $k > 12$ and apply Theorem II.2 to the red-blue arrangement of pq over the interval (t_0, t_1). Assume first that one of the Conditions (i) and (ii) is satisfied, or that one of the points a, b hits pq during $(t_0, t_1]$. Then the preceding analysis (used for extremal Delaunay co-circularities) can be applied, essentially verbatim, in this case too, and it implies that the number of such middle co-circularities is $O\left(k^3 C(n/k) + k^2 N(n/k) + kn^2\beta(n)\right)$.

Assuming that the above scenario does not occur, the four points p, q, a, b are involved in an additional red-blue co-circularity during $(t_0, t_1]$, which "frees" pq from its violation by a and b. Moreover, there is a set A of at most $3k$ points whose removal restores the Delaunayhood of pq throughout $[t_0, t_1]$. Let $t_0 \le t^* \le t_1$ be the time of the additional (third) co-circularity of p, q, a, b, and let B^* be the corresponding circumdisc of p, q, a, b at time t^*.

If B^* contains at most $14k$ points, we can charge the disappearance of pq to the resulting $14k$-shallow extremal co-circularity. Clearly, any such co-circularity of

[4]Indeed, there are at most $O(1)$ ways to guess p and q among the four points of the charged co-circularity, and then the charging co-circularity corresponds to the latest previous disappearance of pq from $\mathrm{DT}(P)$.

index 3 is charged for at most one middle Delaunay co-circularity. Moreover, the number of $14k$-shallow extremal co-circularities is bounded by $O\left(k^4 N_E(n/k)\right)$ using the standard probabilistic argument of Clarkson and Shor [9]. Hence, this scenario arises for at most $O\left(k^4 N_E(n/k)\right)$ middle Delaunay co-circularities.

Now assume that B^* contains at least $14k$ points of P. Without loss of generality, assume that the cap $B \cap L_{pq}^+$ contains at least $7k$ points of P. That is, the corresponding red function, say f_b^+, has level at least $7k$ in the red arrangement at time t^*. Refer to Figure 6 (left). Let r be a red point whose respective function f_r^+ lies, at time t^*, at red level between $3k$ and $7k-1$. That is, the number of red points in the circumdisc $B[p,q,r]$ ranges from $3k$ to $7k-1$. Then the number of blue points in $B[p,q,r]$ is at most $3k$. Indeed, if there were more that $3k$ blue points in $B[p,q,r]$ then after removing A this disc would still contain at least one blue point and at least one red point (possibly r itself), so pq could not be Delaunay at time t^*. Since $f_r^+ < f_b^+$, this disc also contains a (which is still a blue point on the boundary of $B[p,q,b]$), so the Delaunayhood of pq is violated at time t^* by r and a. Before pq re-enters $\mathrm{DT}(P)$ at time t_1, one of the following must happen, according to Lemma III.1: Either r hits pq or the points p,q,r,a are involved in a red-blue co-circularity (when a leaves $B[p,q,r]$ and before r hits L_{pq}). A fully symmetric argument shows that either r hits pq, or p,q,r,a are involved in a red-blue co-circularity during (t_0,t^*) (when a enters $B[p,q,r]$). Note, however, that pq is hit by at most $3k$ points during $(t_0,t_1]$, all of them in A. Thus, at least k such points r do not hit pq during $(t_0,t_1]$, so each of them is involved in two co-circularities with p,q,a during $(t_0,t_1]$: one before t^*, and another afterwards.

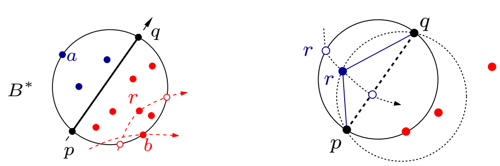

Figure 6. Left: Analysis of middle Delaunay co-circularities. The four points p,q,a,b are involved, during $[t_0,t_1]$, in their third co-circularity, whose respective circumdisc B^* contains at least $7k$ red points. At least k red points r, whose red level ranges between $3k$ and $7k$, do not hit pq during $[t_0,t_1]$. Right: Lemma IV.1. If (pq,r,I) is a Delaunay crossing, then each of pr, rq belongs to $\mathrm{DT}(P)$ throughout I.

Fix a point r, as above, which does not cross pq. Notice that at least one of the two promised co-circularities of p,q,r,a is extremal. If the above extremal co-circularity of p,q,r,a, occuring at some $t^{**} \in (t_0,t_1)$, is $(11k)$-shallow, we charge it for the disappearance of pq. As before, this charging is unique, and the number of charged co-circularities is $O(k^4 N_E(n/k))$. Otherwise, the boundary of $B[p,q,r]$ is crossed during the interval (t^*,t^{**}) (or (t^{**},t^*)) by at least k points, so the triple p,q,r defines $\Omega(k)$ $(11k)$-shallow co-circularities involving p,q during (t_0,t_1).

Repeating the same argument for the (at least) k possible choices of r, we obtain $\Omega(k^2)$ $(11k)$-shallow co-circularities, each involving p,q and some other pair of points and occurring during $(t_0,t_1]$. As in Case (ii) of Theorem II.2, we charge these co-circularities for the disappearance of pq.

We have thus established the following recurrence for the maximum possible number $N_M(n)$ of middle Delaunay co-circularities for a set of n moving points:

$$N_M(n) = \tag{2}$$
$$O\left(k^4 N_E(n/k) + k^2 N(n/k) + kn^2\beta(n) + k^3 C(n/k)\right).$$

IV. The Number of Delaunay crossings

The remainder of the paper is devoted to deriving a recurrence relation for the maximum number $C(n)$ of Delaunay crossings induced by any set P of n moving points as above. In this section we establish several basic properties of Delaunay crossings, and outline the forthcoming stages of their analysis. The eventual system of recurrences that we will derive will express $C(n)$ in terms of the maximum number of Delaunay co-circularities of smaller-size sets, plus a nearly quadratic additive term. Plugging that relation into (1) and (2) will yield the near-quadratic bound on $N(n)$ that was asserted in Theorem II.1.

A. *Delaunay crossings: the key properties*

Consider a Delaunay crossing (pq,r,I). Recall that p,q,r can be collinear at most twice. Moreover, both collinearities can (but do not have to) occur during the interval I of the same Delaunay crossing of pq by r. Clearly, r cannot hit L_{pq} *outside* pq during I because, at such an "outer" collinearity, pq, which is Delaunay when r is removed, would also be Delaunay in the presence of r.

The Delaunay crossing of pq by r is called *single* (resp., *double*) if r hits pq exactly once (resp., twice) during the corresponding interval I of pq's absence from $\mathrm{DT}(P)$.

The following lemma, whose explicit proof appears in the predecessor paper [21], holds for both types of Delaunay crossings (see Figure 6 (right)).

Lemma IV.1. *If $(pq,r,I=[t_0,t_1])$ is a Delaunay crossing then each of the edges pr,rq belongs to $\mathrm{DT}(P)$ throughout I.*

In [21], we obtain an upper bound of $O(n^2)$ on the number of double Delaunay crossings. Since the argument from [21] holds (as is) also in the setting studied by this paper, we have the following theorem.

Theorem IV.2. *Any set P of n moving points, as above, induces at most $O(n^2)$ double Delaunay crossings.*

It therefore suffices to establish a suitable recurrence for the maximum possible number of single Delaunay crossings, and this is what is undertaken in the the remainder of the paper is devoted to the study of the latter crossings.

Single Delaunay crossings: notational conventions. Recall from Section II that every edge pq is oriented from p to q, and its corresponding line L_{pq} splits the plane into the left halfplane L_{pq}^- and the right halfplane L_{pq}^+.

Without loss of generality, we assume in what follows that, for any single Delaunay crossing $(pq, r, I = [t_0, t_1])$, the point r crosses pq from L_{pq}^- to L_{pq}^+ during I. Recall that r cannot cross L_{pq} outside pq during I, so this is the *only* collinearity of p, q, r in I. If r crosses pq in the opposite direction, we denote this crossing as $(qp, r, I = [t_0, t_1])$.

Note that every such Delaunay crossing (pq, r, I) is uniquely determined by the respective ordered triple (p, q, r), since there can be at most one collinearity where r crosses the line L_{pq} from L_{pq}^- to L_{pq}^+ (or, else, r would cross L_{pq} three times). We label each such crossing (pq, r, I) as *a clockwise* (p, r)*-crossing*, and as *a counterclockwise* (q, r)*-crossing*, with an obvious meaning of these labels.

The following lemma lies at the heart of our analysis.

Lemma IV.3. *Let $(pq, r, I = [t_0, t_1])$ be a single Delaunay crossing. Then, with the above conventions, for any $s \in P \setminus \{p, q, r\}$ the points p, q, r, s define a red-blue co-circularity with respect to pq, which occurs during I when the point s either enters the cap $B[p, q, r] \cap L_{pq}^+$, or leaves the opposite cap $B[p, q, r] \cap L_{pq}^-$.*

Proof: By definition, r crosses pq at some (unique) time $t_0 < t^* < t_1$ from L_{pq}^- to L_{pq}^+. The disc $B[p, q, r]$ is P-empty at t_0 and at t_1 and moves continuously throughout $[t_0, t^*)$ and $(t^*, t_1]$. Just before t^*, $B[p, q, r]$ is the entire L_{pq}^+, so every point $s \in P \cap L_{pq}^+$ at time t^* must have entered $B[p, q, r]$ during $[t_0, t^*)$, thus forming a co-circularity with p, q, r. See Figure 7 (left). Clearly, this co-circularity of p, q, r, s is red-blue with respect to pq, since s can enter $B[p, q, r]$ only through $\partial B[p, q, r] \cap L_{pq}^+$. A symmetric argument applies to the points that lie in L_{pq}^- at time t^*; see Figure 7 (right). ∎

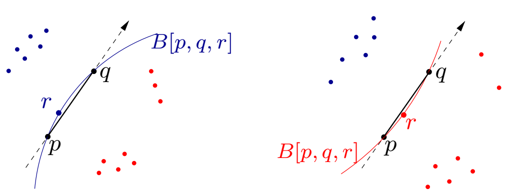

Figure 7. Proof of Lemma IV.3. Left: Right before r crosses pq, the circumdisc $B = B[p, q, r]$ contains all points in $P \cap L_{pq}^+$. Right: Right after r crosses pq, B contains all points in $P \cap L_{pq}^-$.

Our local charging schemes "bottom out" when a carefully chosen triple of points defines two Delaunay crossings (again, possibly in a triangulation of some smaller-size sample). Lemma IV.4 takes care of this easy case.

Lemma IV.4. *The number of triples of points $p, q, r \in P$ for which there exist two time intervals I_1, I_2 such that either (i) both (pq, r, I_1) and (qp, r, I_2) are Delaunay crossings, (ii) both (pq, r, I_1) and (rq, p, I_2) are Delaunay crossings, or (iii) both (pq, r, I_1) and (pr, q, I_2) are Delaunay crossings, is at most $O(n^2)$.*

Notice that, if some triple of points p, q, r in P performs two distinct Delaunay crossings, both of these crossings must necessarily be single Delaunay crossings.

Proof: We claim that every pair $p, q \in P$ participates in at most one triple of each type. Indeed, fix $p, q \in P$ and assume that there exist two points r, s such that the triples p, q, r and p, q, s are involved in two (single) Delaunay crossings of the same prescribed order type (i), (ii), or (iii). By Lemma IV.3, we encounter at least one co-circularity of p, q, r, s during each of the two Delaunay crossings induced by p, q, r and the two induced by p, q, s. In the full version, we argue that these four co-circularities are distinct, contrary to the fact that any four points can be co-circular at most three times. ∎

In [21] we establish the following lemma:

Lemma IV.5. *Let (pq, r, I) and (pa, r, J) be clockwise (p, r)-crossings, and suppose that r hits pq (during I) before it hits pa (during J). Then I begins (resp., ends) before the beginning (resp., end) of J. Clearly, the converse statements hold too. Similar statements hold for pairs of counterclockwise (p, r)-crossings.*

Lemma IV.5 implies that, for any pair of points p, r, all the clockwise (p, r)-crossings can be linearly ordered by the starting times of their intervals, or by the ending times of their intervals, or by the times when r hits the corresponding p-edge, and all three orders are indentical.

B. Quadruples

In Section III we have established a pair of recurrences (1) and (2), whose combination allows to express the maximum number $N(n)$ of Delaunay co-circularities in terms of the maximum number of Delaunay crossings $C(m)$ in smaller-size subsets, plus the maximum number of Delaunay co-circularities in smaller-size sets, plus a nearly quadratic additive term. Furthermore, we have seen that there can be at most quadratically many double Delaunay crossings, and quadratically many of pairs of single Delaunay crossings of the kinds considered in Lemma IV.4.

It therefore suffices to obtain a suitable recurrence, or a system of such recurrences, that express the maximum possible number $C(n)$ of (single) Delaunay crossings only in terms of the maximum number of Delaunay co-circularities in smaller-size sets, plus a nearly quadratic additive term. (In order for the solution of such a recurrence to be near-quadratic, the respective coefficient of each recursive term of the form $N(n/k)$ must be roughly equal to k^2. See, e.g., [16], [23, Section 7.3.2], and [20, Section 4.5].)

Informally, the main weakness of Delaunay crossings stems from the fact that Delaunay crossings involve triples of points, whereas our primary topological restriction refers to quadruples of points of P. Thus, Delaunay crossings are not "rich" enough to capture the underlying combinatorial structure of the problem. We therefore consider several additional types of topological configurations that involve *quadruples* of moving points, obtained by combining two Delaunay crossings with two common points, such as (pq, r, I) and (pa, r, J). Recall that, for each Delaunay crossing (pq, r, I), its edge pq is almost Delaunay in $I = [t_0, t_1]$ (and fully Delaunay at the endpoints t_0, t_1), and the other two edges pr and rq are fully Delaunay in I (by Lemma III.1). The quadruples that we will shortly introduce, inherit all these properties of their Delaunay crossings, but will have a rich structure, due to additional interactions between their edges and subtriples. These quadruples can be viewed as an extension of Delaunay crossings, in the sense that their edges are forced to be either Delaunay, or almost Delaunay, during various intervals whose endpoints are defined "locally", in terms of the points and the edges of the configuration at hand. Furthermore, by construction, the points of each quadruple perform at least two Delaunay crossings. The major goal of the analysis is to obtain configurations with progressively many Delaunay crossings.

Due to the lack of space, we only review the three types of quadruples that arise in the course of our analysis, and highlight the intimate relations between them and Delaunay crossings. The rest of the details can be found in the full version of the paper.

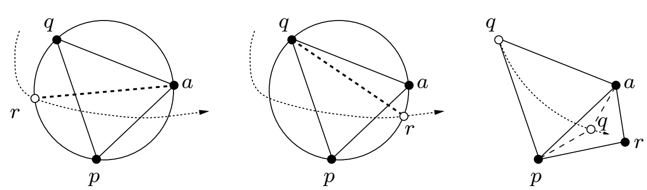

Figure 8. A (clockwise) regular quadruple $\sigma = (p, q, a, r)$, which is composed of clockwise (p, r)-crossings (pq, r, I) and (pa, r, J). Left and center: A possible motion of r, with the two co-circularities of p, q, a, r that occur during $I \setminus J$ and $J \setminus I$, respectively. Right: The special crossing of pa by q which we enforce at the end of the analysis of regular quadruples.

Regular quadruples. Four distinct points $p, q, a, r \in P$ form a clockwise *regular quadruple* (or, simply, a *quadruple*) $\sigma = (p, q, a, r)$ in $\mathrm{DT}(P)$ if there exist clockwise (p, r)-crossings $(pq, r, I), (pa, r, J)$ that appear in this order in the sequence of clockwise (p, r)-crossings; refer to Figure 8. We say that the quadruple is *consecutive* if (pq, r, I) and (pa, r, J) are consecutive in the order of Lemma IV.5.

Clearly, every clockwise (p, r)-crossing (pq, r, I) forms the first part of exactly one (clockwise) consecutive quadruple, unless it is the last such (p, r)-crossing (with respect to the order given by Lemma IV.5). The overall number of these last crossings is clearly bounded by $O(n^2)$. Hence, the maximum number $C(n)$ of single Delaunay crossings is asymptotically dominated by the maximum possible number of consecutive regular quadruples.

Let $\sigma = (p, q, a, r)$ be a consecutive regular quadruple as above. By Lemma IV.1, edge pr of σ is Delaunay during the respective intervals I and J of its two (p, r)-crossings, whereas each of the edges rq and ra is (provably) Delaunay in only one of these two intervals. In addition, the edges pq and pa are almost Delaunay during their respective Delaunay crossings by r.

Regular quadruples are studied extensively in the full version of this paper, where we gradually extend the corresponding (almost-)Delaunayhood intervals of the respective edges pr, rq, ra, pa and pq of each quadruple σ until most of them cover $[I, J] = \mathsf{conv}(I \cup J)$, including the possible gap between I and J. This is achieved by applying Theorem II.2 in the respective red-blue arrangements of these edges. Each such application of Theorem II.2 is done over a carefully chosen interval, which guarantees that any shallow collinearity or co-circularity, that we encounter in the first two cases of the theorem, is charged by only few quadruples.

We show (via Lemmas IV.1 and IV.3) that the points of each regular quadruple $\sigma = (p, q, a, r)$ are co-circular exactly once in each of the intervals $I \setminus J$ and $J \setminus I$; see Figure 8 (left and center). Specifically, the former co-circularity is red-blue with respect to the edges pq and ra, and the latter co-circularity is red-blue with respect to pa and rq. Notice that at least one of these co-circularities, let it be the one in $I \setminus J$, is extremal.

Arguing similarly to Section III, we use the above co-circularities of p, q, a, r (together with the additional constraints on the Delaunayhood of rq, ra and pa) to enforce a pair of additional Delaunay crossings which occur in smaller-size point sets (which are random samples of P, needed for the application of the Clarkson-Shor argument [9]) and involve various sub-triples of p, q, a, r. Thr analysis is fairly involved, due to the fact that neither of the above two co-circularities of σ has to be Delaunay, or even shallow. If some sub-triple of σ performs two Delaunay crossings, we immediately bottom out via Lemma IV.4.

Unfortunately, there may still exist quadruples σ whose four resulting Delaunay crossings (including (pq, r, I) and (pa, r, J)) involve four distinct sub-triples p, q, a, r, so Lemma IV.4 cannot yet be applied. As our analysis shows, in this only remaining scenario, the edge pa of σ undergoes a Delaunay crossing $(pa, q, \mathcal{I})$ by q; see Figure 8 (right). We refer to this latter crossing as a *special crossing* of pa by q.

Special quadruples. We analyze the number of special (counterclockwise) crossings by first arranging them into *special quadruples*. Informally, each special quadruple $\chi = (a, p, w, q)$ is composed of two special (a, q)-crossings $(pa, q, \mathcal{I})$ and $(wa, q, \mathcal{J})$ which are consecutive in the order of Lemma IV.5. See Figure 9.

The treatment of (counterlockwise) special quadruples is fairly symmetric to that of (clockwise) regular quadruples, in the manner in which we extend the Delaunayhood or

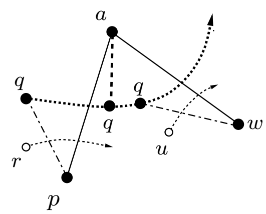

Figure 9. A special quadruple $\chi = (a, p, w, q)$, is composed of two special crossings $(pa, q, \mathcal{I})$ and $(wa, q, \mathcal{J})$, which respectively correspond to some (clockwise) regular quadruples (p, q, a, r) and (w, q, a, u).

almost-Delaunayhood of their edges, and enforce additional (almost-)Delaunay crossings on some of their sub-triples. However, here we have a richer topological structure, because the two special crossings $(pa, q, \mathcal{I})$ and $(wa, q, \mathcal{J})$ of each special quadruple χ are accompanied by two respective regular quadruples $\sigma_1 = (p, q, a, r)$ and $\sigma_2 = (w, q, a, u)$.

At the final stage of the analysis, we use the above correspondence with the regular quadruples in order to charge the surviving special quadruples χ to especially convenient sub-configurations, referred to as *terminal quadruples*.

Terminal quadruples. Each terminal quadruple $\varrho = (p, q, r, w)$ is formed by an edge pq, and by a pair of points r and w that cross pq in *opposite* directions; see Figure 10. In addition, ϱ must satisfy several "local" restrictions on the Delaunayhood of its various edges, and on the co-circularities and collinearities among p, q, r, w. We directly bound the number of such quadruples in terms of simpler quantities, introduced in Section II, and thereby complete the proof of Theorem II.1. (The recurrences that bound the number of terminal quadruples have only "quadratic" terms.)

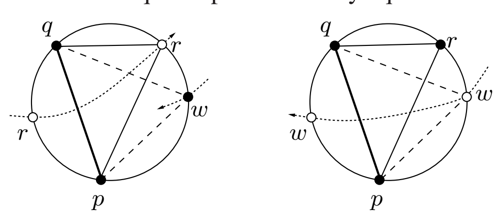

Figure 10. A terminal quadruple $\varrho = (p, q, r, w)$. The points r and w cross pq in opposite directions. The points of ϱ are co-circular three times. The extremal two co-circularities are red-blue with respect to pq, and the middle one is monochromatic with respect to pq. The left figure depicts the first and second co-circularities, and the right figure depicts the second and third co-circularities.

Informally, the analysis of terminal quadruples manages to bottom out because each terminal quadruple comes with *three* "well-behaved" co-circularities. Specifically, the two extremal co-circularities are red-blue with respect to the crossed edge pq (and thus also with respect to rw), and the middle one is mononochromatic with respect to pq; see Figure 10. These patterns allow us to use these co-circularities to enforce *three* additional Delaunay crossings among p, q, r, w (in addition to the crossings of pq by r and w). Thus, some sub-triple among p, q, r, w is involved in two Delaunay crossings, so Lemma IV.4 can always be invoked.

References

[1] P. K. Agarwal, O. Cheong and M. Sharir, The overlay of lower envelopes in 3-space and its applications, *Discrete Comput. Geom.* 15 (1996), 1–13.

[2] P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, V. Koltun, N. Rubin and M. Sharir, Kinetic stable Delaunay graphs, *Proc. 26th Annu. Symp. on Comput. Geom.* (2010), 127–136.

[3] P. K. Agarwal, Y. Wang and H. Yu, A 2D kinetic triangulation with near-quadratic topological changes, *Discrete Comput. Geom.* 36 (2006), 573–592.

[4] M. J. Atallah, Dynamic computational geometry, *Proc. 24th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 92–99, 1983.

[5] M. J. Atallah, Some dynamic computational geometry problems, *Comput. Math. Appl.* 11 (12) (1985), 1171–1181.

[6] F. Aurenhammer and R. Klein, Voronoi diagrams, in *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds., Elsevier, Amsterdam, 2000, pages 201–290.

[7] T. M. Chan, A dynamic data structure for 3-D convex hulls and 2-D nearest neighbor queries, *J. ACM* 57 (3) (2010), Article 16.

[8] L. P. Chew, Near-quadratic bounds for the L_1 Voronoi diagram of moving points, *Comput. Geom. Theory Appl.* 7 (1997), 73–80.

[9] K. Clarkson and P. Shor, Applications of random sampling in computational geometry, II, *Discrete Comput. Geom.* 4 (1989), 387–421.

[10] B. Delaunay, Sur la sphère vide. A la memoire de Georges Voronoi, *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskih i Estestvennyh Nauk* 7 (1934), 793–800.

[11] E. D. Demaine, J. S. B. Mitchell, and J. O'Rourke, The Open Problems Project, `http://www.cs.smith.edu/~orourke/TOPP/`.

[12] H. Edelsbrunner, *Geometry and Topology for Mesh Generation*, Cambridge University Press, Cambridge, 2001.

[13] S. Fortune, Voronoi diagrams and Delaunay triangulations, In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*, CRC Press, Inc., Boca Raton, FL, USA, second edition, 2004, pages 513–528.

[14] J. Fu and R. C. T. Lee, Voronoi diagrams of moving points in the plane, *Int. J. Comput. Geometry Appl.* 1 (1) (1991), 23–32.

[15] L. J. Guibas, J. S. B. Mitchell and T. Roos, Voronoi diagrams of moving points in the plane, *Proc. 17th Internat. Workshop Graph-Theoret. Concepts Comput. Sci.*, volume 570 of *Lecture Notes Comput. Sci.*, pages 113–125. Springer-Verlag, 1992.

[16] D. Halperin and M. Sharir, New bounds for lower envelopes in three dimensions, with applications to visbility in terrains, *Discrete Comput. Geom.* 12 (1994), 313–326.

[17] H. Kaplan, N. Rubin and M. Sharir, A kinetic triangulation scheme for moving points in the plane, *Comput. Geom. Theory Appl.* 44 (2011), 191–205.

[18] V. Koltun, Ready, Set, Go! The Voronoi diagram of moving points that start from a line, *Inf. Process. Lett.* 89(5) (2004), 233–235.

[19] V. Koltun and M. Sharir, 3-dimensional Euclidean Voronoi diagrams of lines with a fixed number of orientations, *SIAM J. Comput.* 32 (3) (2003), 616–642.

[20] V. Koltun and M. Sharir, The partition technique for overlays of envelopes, *SIAM J. Comput.* 32 (4) (2003), 841–863.

[21] N. Rubin, On topological changes in the Delaunay triangulation of moving points, *Discrete Comput. Geom.*, 49 (4) (2013), 710–746.

[22] M. Sharir, Almost tight upper bounds for lower envelopes in higher dimensions, *Discrete Comput. Geom.* 12 (1994), 327–345.

[23] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.