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

Natan Rubin*

June 8, 2018

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.

*Université Pierre & Marie Curie, Institut de Mathématiques de Jussieu (UMR 7586), 4 Place Jussieu, 75252 Paris Cedex, France. Email: `rubinnat.ac@gmail.com`. Work on this paper was partly done while the author was a Minerva Postdoctoral Fellow at the Department of Mathematics and Computer Science of the Free University of Berlin. Work on this paper was partly supported by Minerva Fellowship Program of the Max Planck Society, and by Fondation Sciences Mathématiques de Paris.

1 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 and only 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, 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, 14] for surveys and a textbook on these structures. The problem has also been studied in the *dynamic* setting, where one seeks to maintain $\mathrm{DT}(P)$ and $\mathrm{VD}(P)$ under updates of P (insertion and deletion of points); see, e.g., [7].

The kinetic setting: Previous work. In many applications of Delaunay/Voronoi methods (e.g., mesh generation and kinetic collision detection) the points of the input set P are moving continuously, so these diagrams need to be efficiently updated during the motion. Even though the motion of the points is continuous, the combinatorial structure of the Voronoi and Delaunay diagrams changes only at discrete times when certain critical events occur. Interest in efficient maintenance of geometric structures under simple motion[1] of the underlying point set 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 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* in this paper, 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 maximum possible number of *discrete changes*, that is, Delaunay co-circularities and convex hull collinearities, which $\mathrm{DT}(P)$ experiences during the motion of its points.

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, 14]. 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 (and even more difficult) version of the problem, each point of P moves with so-called pseudo-algebraic motion of constant description complexity. This implies (in particular) 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 on the continuous motion of P, Fu and Lee [15], and Guibas et al. [16] 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 [25]. A substantial gap exists between these near-cubic upper bounds and the best known quadratic lower bound [25]. 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 [23] 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, several alternative structures were

[1] While there are several ways to define this notion, the simplest would be to assume that each coordinate of each point $p = p(t)$ in P is a is fixed-degree polynomial in t.

[2] We assume that the motion of the points is generic, so that no more than four points can become co-circular at any given time.

studied. For example, Chew [8] proved that $\mathrm{VD}(P)$ undergoes a near-quadratic number of discrete changes if it is defined with respect to a "polygonal" distance function. More recent studies [3, 19] show how to maintain a (non-Delaunay) triangulation of P so that it undergoes only a near-quadratic number of changes. Agarwal et al. [2] show how to efficiently maintain a so called α-*stable* subgraph of the Euclidean $\mathrm{DT}(P)$, which experiences only a near-quadratic number of changes, and whose edges are robust with respect to small changes in the underlying norm.

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, 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 (and briefly discussed in Appendix A), 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 the edge pq, for a total of $O(n^2)$ such events. We thus assume that pq is again Delaunay at some moment $t_1 > t_0$.

One of the major observations used in our analysis is 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 the points u or v crosses pq. This fairly simple observation lies at the heart of our charging strategy.

Combinatorial charging. Our goal is to derive a recurrence formula 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 (i.e., their corresponding circumdiscs are penetrated by additional points of P).

In Section 2 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 absent[3] from $\mathrm{DT}(P)$. 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[4], 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 [21]), 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

[3]In fact, the analysis in Section 2 is more general, and applies to any interval (t_0, t_1) with the property that pq is Delaunay at one of its endpoints t_0, t_1.

[4]Each of them would become a Delaunay co-circularity after removal of at most k points of P.

in the convex hull of P (see [25, Section 8.6.1]) yields an additional near-quadratic term in the recurrence.

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 2.2), 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 the Clarkson-Shor argument) in a reduced Delaunay 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 (time-reversed) argument applies if we encounter the *first* co-circularity of p, q, a, b.

As argued in the predecessor paper [23], Delaunay crossings are especially nice objects due to their strict structural properties. In particular, as shown in [23]: (i) The edges pa and aq belong to $\mathrm{DT}(R)$ throughout the above interval I, and (ii) Assuming a hits pq exactly once during I, every other point $w \in R \setminus \{p, q, a\}$ is involved during this interval in a co-circularity with p, q, a.

The roadmap. In Section 3 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 [23]) works only for the first and the last Delaunay co-circularities of the quadruple.

To extend the above reduction to the remaining, "middle" Delaunay co-circularities, we resort in Section 3 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.

In Section 4, we recall (or re-establish) several structural properties of Delaunay crossings, which will be used throughout the rest of the analysis. 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 the number of Delaunay crossings in terms of the maximum number of certain *quadruples* in P. Each such quadruple $\sigma = (p, q, a, r)$ is composed of a pair of "consecutive" Delaunay crossings of p-adjacent edges pq and pa, by the same point r.

In Section 5 we apply the routine "charge-or-refine" strategy (via our Red-Blue Theorem) to analyze the maximum number of the above quadruples. 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 4 a direct quadratic bound on their number.

There are two main types of such 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 only $O(n^2)$.

Each step of the analysis enforces additional constraints on the surviving quadruples. There are two main types of such constraints. The first is to enforce more Delaunay crossings involving sub-triples of the points of the quadruple. The other is to enforce "almost-Delaunayhood" of various pairs of points in the quadruple, for progressively larger time intervals. By this we mean that the corresponding edge is Delaunay if we remove from P a small subset of points. The ultimate goal is to enforce sufficiently many Delaunay crossings, so that some triple of points undergoes *two* distinct Delaunay crossings. As mentioned above, this is the main type of the "terminal" configurations, for which we have a quadratic bound on their number.

Each step of the analysis yileds a recurrence formula that involves "near-quadratic" terms (of the kind mentioned earlier) plus terms involving further-constrained configurations, until we finally bottom out (in Section 7) by reaching the terminal triples mentioned above. In each of the recurrences we make use of the Clarkson-Shor probabilistic argument [9], in order to get rid of the small "obstruction" subset of P that we need to remove; this is done by passing to a random sample of P, the standard style of [9]. The overall collection of recurrences solves to a near-quadratic bound, in a manner similar to many earlier works involving such recurrences (see, e.g., [1, 17, 21, 22, 24] and [25, Section 7.3.2]).

Unfortunately, the analysis is fairly involved and consists of many steps. In addition to the aforementioned type of quadruples (formed by pairs of Delaunay crossings), we use two additional classes of quadruples which are studied in Sections 6 and 7, respectively. Note that only the last kind of configurations, referred to as *terminal quadruples*, can always be traced to some of the above "terminal" triples.

We postpone the rest of this discussion until Section 4.2, where we provide a more detailed summary of the three classes of quadruples, and of the connections between these classes, and the Delaunay crossings.

Finally, we emphasize that the contribution of the paper, and its main ideas, are delivered already in Sections 1 through 4.

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 invaluable help in the preparation and careful reading of this paper.

2 Geometric Preliminaries

Delaunay co-circularities. Let P be a collection of n points moving along 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* during the motion. In addition, we assume, without loss of generality, that the trajectories of the points of P satisfy all the standard general position assumptions; see Appendix B for more details.

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 two types of 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 on the convex hull boundary is known to be at most nearly quadratic; see, e.g., [25, 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 2.1. *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 that can arise in a set of 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 (with respect to P) becomes a Delaunay co-circularity (with respect to 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. Similar notations apply to collinearities of triples of points p, q, r. 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 above probabilistic 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 [25, Section 8.6.1], $\mathsf{H}(m) = O(m^2\beta(m))$, where $\beta(\cdot)$ is an extremely slowly growing function.[5] We thus get that 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 (since, by assumption, p and q never coincide during the motion). Note also that L_{pq} and L_{qp} are oppositely 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 circumdisc $B[p, q, r]$ changes instantly from being all of L^+_{pq} to all of L^-_{pq}. Similarly, when r crosses L_{pq} from L^-_{pq} to L^+_{pq} *outside* pq, the circumdisc changes instantly from L^-_{pq} to L^+_{pq}. Symmetric changes occur when r crosses L_{pq} from L^+_{pq} to L^-_{pq}.

[5] 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 [25].

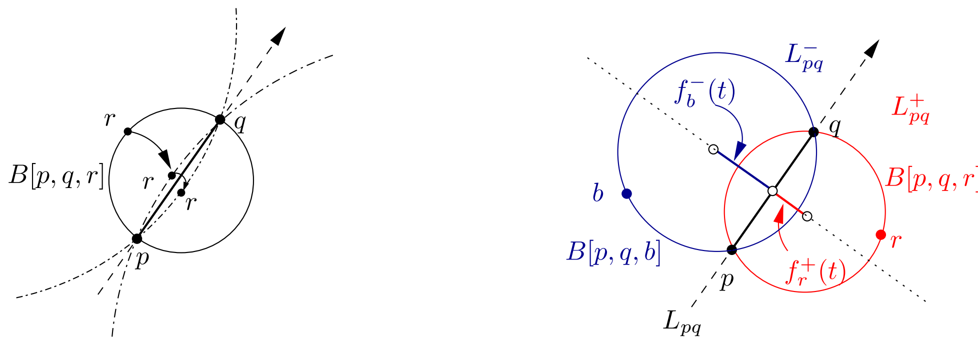

Figure 2: Left: The circumdisc $B[p,q,r]$ of p,q and r moves continuously as long as these three points are not collinear, and then flips over to the other side of the line of collinearity after the collinearity. Right: A snapshot at moment t. In the depicted configuration we have $f^-_b(t) < 0 < f^+_r(t)$.

The red-blue arrangement. As in [16, 23], 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 the sake of completeness, we provide below a formal definition of this arrangement.

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). The common points of discontinuity of f^+_r, f^-_r occur at moments when r crosses L_{pq}. Specifically, f^+_r tends to $+\infty$ before r crosses L_{pq} from L^+_{pq} to L^-_{pq} outside the segment pq, and it tends to $-\infty$ when r does so within pq; the behavior of f^-_r is fully symmetric.

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 $E^-(t) < E^+(t)$. Any disc whose bounding circle passes through p and q which is centered anywhere in the interval $(E^-(t), E^+(t))$ along the perpendicular bisector of pq (with the origin on this line lying at the midpoint of pq) is empty at time t, and thus serves as a witness to pq being Delaunay. If pq is not Delaunay at time t, there is a pair of a red function $f^+_r(t)$ and a blue function $f^-_b(t)$ such that $f^+_r(t) < f^-_b(t)$. For example, we can take f^+_r (resp., f^-_b) to be the function attaining E^+ (resp., E^-) at time t; see Figure 3 (left). In such a case, we say that the Delaunayhood of pq is *violated* by the pair of points $r, b \in P$ that define f^+_r, f^-_b. Note that in general there can be many pairs (r, b) that violate pq (quadratically many in the worst case).

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. Note that our general position assumptions imply that $\mathcal{A}$ is also in general position, so that no three functions pass through a common vertex, and no pair of functions are tangent to each other. As discussed above, the functions forming $\mathcal{A}$ have in general discontinuities, at the corresponding collinearities. At the time t_0 of each such collinearity,

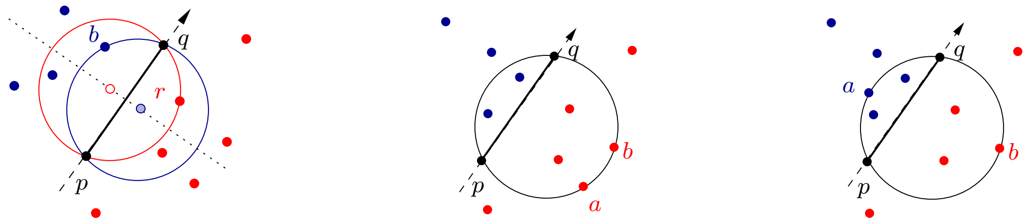

Figure 3: Left: A snapshot at fixed time t. The red and blue envelopes E^+, E^- coincide with the functions f_r^+, f_b^-, respectively. The edge pq is not a Delaunay edge because $E^+(t)$ (the hollow center) is smaller than $E^-(t)$ (the shaded center). Center and right: Red-red and red-blue co-circularities.

a red function f_r^+ tends to ∞ or $-\infty$ on one side of t_0, and is replaced on the other side of t_0 by the corresponding blue function f_r^- which tends to $-\infty$ or ∞, respectively.

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 event involving p, q, a, b where both a and b lie in L_{pq}^-. Also, an intersection of a red fuction f_a^+ and a blue function f_b^- represents a co-circularity of p, q, a, b, where $a \in L_{pq}^+$ and $b \in L_{pq}^-$. We label these co-circularities, as we labeled the vertices of $\mathcal{A}$, as red-red, blue-blue, and red-blue (all with respect to pq), depending on the respective colors of a and b. See Figure 3 (center and right).

It is instructive to note that in any co-circularity of four points of P there are exactly two pairs (the opposite pairs in the co-circularity) 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 $\mathcal{A}_{pq}$, which is one of the major tools in our analysis, was established in [23] by applying routine techniques for analyzing planar arrangements. For the sake of completeness, we provide its proof in Appendix C.

Theorem 2.2 (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.*[6] *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).

Notice that we do not assume that pq leaves $\mathrm{DT}(P)$ at any moment during (t_0, t_1) (in that case, case (iii) holds, with $A = \emptyset$). Note also that, although we do not need this property, the theorem continues to hold in the more general setting of pseudo-algebraic motions of constant description complexity.

[6]The constants is the $O(\cdot)$ and $\Omega(\cdot)$ notations do not depend on k.

3 From Delaunay Co-Circularities to Delaunay Crossings

Let P be a set of n points moving in the plane, so that any four points can be co-circular at most three times, and any triple of points can be collinear more than twice. For the sake of brevity, we will often take these topological restrictions for granted. As before, $N(n)$ denotes the maximum possible number of Delaunay co-circularities that can arise in such a set P.

In this section we introduce the notion of a Delaunay crossing, which plays a central role both in this paper and in its predecessor [23], and express the above quantity $N(n)$ in terms of the maximum numbers 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[7] 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).

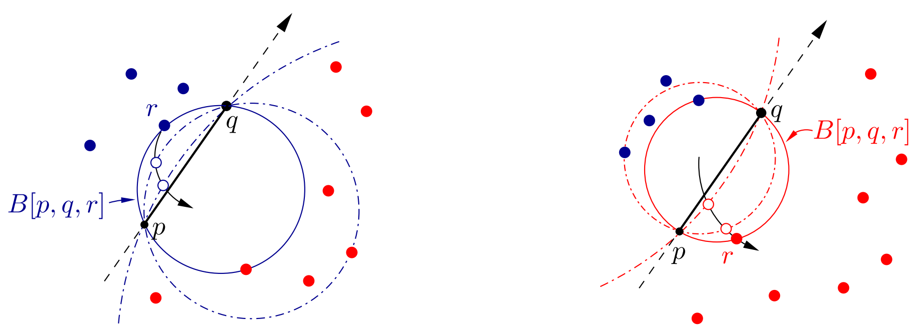

Figure 4: A Delaunay crossing of pq by r from L^-_{pq} to L^+_{pq}. Several snapshots of the continuous motion of $B[p,q,r]$ before and after r crosses pq are depicted (in the left and right figures, respectively). Hollow points specify the positions of r when $pq \notin \mathrm{DT}(P)$. The solid circle in the left (resp., right) figure is the Delaunay co-circularity that destroys (resp., restores) the Delaunayhood of pq.

Note that each of the Delaunay co-circularities that destroys the Delaunayhood of pq at time t_0 and restores it at time t_1 must involve r.

Note that we also allow Delaunay crossings, where the point r hits pq at one (or both) of the times t_0, t_1. In this case, the crossed edge pq leaves the convex hull of P at time t_0, or enters it at time t_1, so the overall number of such "degenerate" crossings does not exceed $O(n^2\beta(n))$, and we may ignore them in what follows.

Assuming $n \geq 5$, 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.[8] 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

[7]And at most twice, by assumption.

[8]If $n = 4$, then, in order for (3) to hold, we also need that the remaining point of P does not cross pq during I.

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 $m \leq n$), 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 [23], 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 $N_E(n/k)$. (Here k is an arbitrary 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, for that matter, also with respect to ab) of level 0 (that is, a co-circularity that is Delaunay). Without loss of generality, assume that the co-circularity of p, q, a, b has index 3 (the case of index 1 is handled fully symmetrically, 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 $t_1 > t_0$. (As reviewed in Section 2, $\mathrm{DT}(P)$ experiences then either a red-blue Delaunay co-circularity with respect to pq, or a hull event, when pq is crossed by a point of $P \setminus \{p, q\}$. In the latter case, pq is not strictly Delaunay at time t_1, and joins $\mathrm{DT}(P)$ right after t_1.) Note that in this case, 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:

Lemma 3.1. *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.

Furthermore, the Delaunayhood of pq is violated by a and b (so, in particular, the segments pq and ab intersect) after time t_0 and until the first time in $(t_0, t_1]$ when at least one of the events in (1)–(3) occurs.

Clearly, the third scenario is not possible if the co-circularity at time t_0 has index 3. A symmetric version of Lemma 3.1 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 $t_1 < t_0$.

Proof. Refer to Figure 5. Clearly, the Delaunayhood of pq remains violated by a and b after time t_0 as long as a remains within the cap $B[p, q, b] \cap L_{pq}^-$, and b remains within the cap $B[p, q, a] \cap L_{pq}^+$ (as depicted in the left figure).

Consider the first time $t^* \in (t_0, t_1]$ when the above state of affairs ceases to hold. Notice that the Delaunayhood of pq is violated by a and b (so, in particular, pq is intersected by ab) throughout the interval (t_0, t^*). Assume without loss of generality that a leaves the the cap $B[p,q,b] \cap L_{pq}^-$. If a crosses pq, then the first scenario holds. Otherwise, a can leave the above cap only through the boundary of $B[p,q,b]$ (as depicted in the right figure), so the third scenario occurs. □

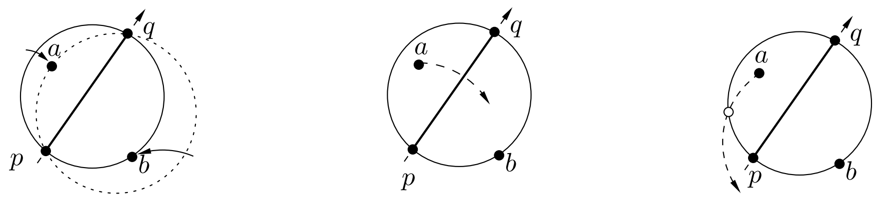

Figure 5: Proof of Lemma 3.1. Left: The setup right after time t_0. Center and right: the point a can leave $B[p,q,b] \cap L_{pq}^-$ (before b leaves the symmetric cap $B[p,q,a] \cap L_{pq}^+$) in two possible ways, corresponding to cases (1) and (3) of the lemma.

Notice, however, that the points of P can define $\Omega(n^3)$ collinearities, so a naive charging of extremal Delaunay co-circularities to collinearities of type (1) or (2) in Lemma 3.1 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 [23].

We fix some sufficiently large constant parameter $k > 12$ and apply Theorem 2.2 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 2, 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,[9] 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 3.1, 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[10] $\mathrm{DT}(P \setminus A \cup \{a\})$, the collinearity of p, q and a can be turned into one or several Delaunay crossings.

We can 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 $[t_0, t_1]$. In addition, we assume that a hits pq during $(t_0, t_1]$, and then $a \in A$.

We sample at random (and without replacement) 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 time t_0 and (right after) time t_1, the sample R induces a 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,

[9] 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 $\mathrm{DT}(P)$.

[10] To simplify the ongoing discourse, we apply slight abuse of notation, where we refer to certain non-Delaunay events as occurring *in* a suitable triangulation. These events are closely related to the changes that the triangulation undergoes, even though they themselves are not part of the Delaunay triangulation.

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 (or to the first such crossing if there are two) 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.

If the Delaunay co-circularity of p, q, a, b has index 1, we reverse the direction of the time axis and argue as above for the edge ab instead of pq. We thus obtain the following recurrence for the maximum possible number $N_E(n)$ of extremal Delaunay co-circularities:

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

Remark. Our analysis will generate many recurrences of similar nature. Informally, each recurrence will have "quadratic" terms (such as the second and the third terms in (1)), which, in themselves, lead to a near-quadratic bound, and "non-quadratic" terms (such as the first one in (1)), which delegate the charging to new quantities. These quantities will generate recurrences of their own, of a similar nature, and the process will bottom out, in Section 7, with recurrences that have only "quadratic" terms. Using known techniques, such as in [17] and [25, Section 7.3.2], the whole system of recurrences will yield a near quadratic bound (for all the involved quantities).

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 2.2 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 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. 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 3.1: Either r hits[11] 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 fairly 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.

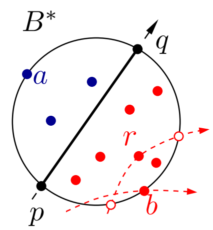

Figure 6: Analysis of middle Delaunay co-circularities. The four points p,q,a,b are involved, during $[t_0,t_1]$, in their third co-circularity, whose respective circumdisc B^* contains at least $7k$ red points. At least k red points r, whose red level ranges between $3k$ and $7k$, do not hit pq during $[t_0,t_1]$.

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 2.2, 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:

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

Informally, and as will be argued rigorously later on, the combination of (1) and (2) implies that the maximum number of extremal Delaunay co-circularities is asymptotically dominated by the maximum number of Delaunay crossings (assuming it is at least quadratic).

4 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) will yield the near-quadratic bound on $N(n)$ that was asserted in Theorem 2.1.

[11] Recall that, by assumption, a does not hit pq in the present case.

4.1 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 holds for both types of Delaunay crossings (see Figure 7).

Lemma 4.1. *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.*

Lemma 4.1, whose explicit proof appears in the predecessor paper [23], is a direct corollary of the following well-known result on *static* Delaunay triangulations:

Lemma 4.2. *Let Q be a finite set of points in $\mathbb{R}^2$, and let r be a point not in Q. Let pq be an edge that is Delaunay in Q, but not in $Q \cup \{r\}$. Then the triangulation $\mathrm{DT}(Q \cup \{r\})$ includes the two edges pr and qr.*

For the sake of completeness, we prove Lemma 4.2 in Appendix E.

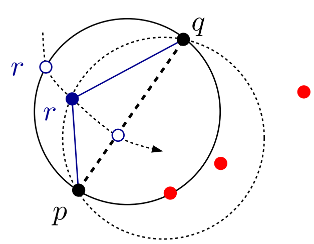

Figure 7: Lemma 4.1. If (pq, r, I) is a Delaunay crossing, then each of pr, rq belongs to $\mathrm{DT}(P)$ throughout I.

In the full version of the predecessor paper [23], we obtain an upper bound of $O(n^2)$ on the number of double Delaunay crossings. Since the argument from [23] holds (as is) also in the setting studied by this paper, we have the following theorem.

Theorem 4.3. *Any set P of n moving points, as above, induces at most $O(n^2)$ double Delaunay crossings.*

For the sake of completeness, we supply the complete analysis of double Delaunay crossings in Appendix D.

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. For the sake of brevity, we shall often refer to single Delaunay crossings simply as Delaunay crossings, and use $C(n)$ to denote the maximum number of single Delaunay crossings.

We next establish several topological properties of (single) Delaunay crossings.

Single Delaunay crossings: notational conventions. Recall from Section 2 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), because there can be at most one collinearity[12] where r crosses the line L_{pq} within pq from L^-_{pq} to L^+_{pq}.

For convenience of reference, 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 4.4. *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. The proof is an adaptation of similar arguments made earlier. 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 $[t_0, t^*)$, forming a co-circularity with p, q, r at the time it entered the disc.[13] See Figure 8 (left). (As mentioned in Section 2, this co-circularity of p, q, r, s is red-blue with respect to pq, that is, the point s enters $B[p, q, r]$ through $\partial B[p, q, r] \cap L^+_{pq}$.) A symmetric argument (in which we reverse the direction of the time axis) shows that the same holds for all the points $s \in P$ that lie in L^-_{pq} at time t^*; see Figure 8 (right). $\square$

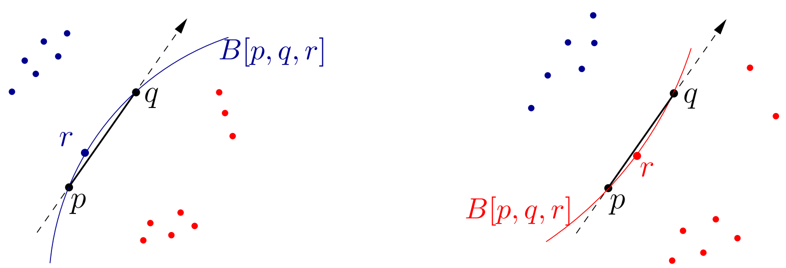

Figure 8: Left: Right before r crosses pq, the circumdisc $B = B[p, q, r]$ contains all points in $P \cap L^+_{pq}$. Right: Right after r crosses pq, B contains all points in $P \cap L^-_{pq}$.

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 4.5 takes care of this easy case.

Lemma 4.5. *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 (otherwise this triple would be collinear

[12]If r hits pq twice, then the other crossing of pq by r is from L^+_{pq} back to L^-_{pq}.

[13]If $t^* = t_0$ then there are no red points when r hits pq, so we consider only the second interval. The case of $t^* = t_2$ is treated symmetrically. As noted in Section 2, in such cases the crossed edge pq either leaves or joins the convex hull of P at the time of the collinearity.

at least three times). Hence, the statement of the lemma holds in full generality. It is easy to check that Lemma 4.5 covers all possible scenarios (up to a permutation of p, q, r and/or reversal of the time axis) where some triple p, q, r is involved two single Delaunay crossings (again, because no three points of P can be collinear more than twice).

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 4.4, 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. If we show that these four co-circularities are distinct, we reach a contradiction to the fact that any four points can be co-circular at most three times.

If the aforementioned triples p, q, r and p, q, s satisfy the first condition, the resulting four crossings of pq happen during pairwise disjoint intervals of time. Hence, the four co-circularities are clearly distinct.

We now proceed to establish the distinctness in the second and the third cases. Assume next that both (p, q, r) and (p, q, s) fall into Case (ii); Case (iii) is handled in a fully symmetric manner. By assumption, we have four points p, q, r, s and four time intervals I_1, I_2, I_3, I_4, such that $(pq, r, I_1), (rq, p, I_2), (pq, s, I_3)$, and (sq, p, I_4) are all Delaunay crossings. I_1 and I_3 are clearly disjoint, and Lemma 4.4 yields two co-circularities of p, q, r, s, one occuring during I_1 and one during I_3, both red-blue with respect to pq. Similarly, Lemma 4.4 yields a co-circularity of p, q, r, s during I_2 which is red-blue with respect to qr, and a co-circularity of the same quadruple during I_4, which is red-blue with respect to qs. Clearly, these two co-circularities are different, and are also different from the former two co-circularities, since the vertex opposite to q is different in each of these co-circularities. This completes the proof of the lemma. ☐

The following lemma defines a natural order on (p, r)-crossings of a given orientation (clockwise or counterclockwise).

Lemma 4.6. *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.*

Proof. In the configuration considered in the main statement of the lemma, r crosses pq from L^-_{pq} to L^+_{pq}, and it crosses pa from L^-_{pa} to L^+_{pa}. We only prove the part of the lemma concerning the ending times of the crossings, because the proof about the starting times is fully symmetric (by reversing the direction of the time axis). The statement clearly holds if I and J are disjoint; the interesting situation is when they partially overlap. Note that r enters L^+_{pq} only once during the Delaunay crossing of pq by r, namely, right after r hits pq. Indeed, by assumption, r cannot exit L^+_{pq} by crossing pq again during I, and it cannot cross $L_{pq} \setminus pq$ because at that time pq, which is Delaunay in $\mathrm{DT}(P \setminus \{r\})$, would be Delaunay also in the presence of r, contrary to the definition of a Delaunay crossing. Hence, we may assume that r still lies in L^+_{pq} when it hits pa during the Delaunay crossing of that edge. Indeed, otherwise the crossing of pq would by then be over, so the claim would hold trivially, as noted above. In particular, $\vec{pa}$ lies clockwise to $\vec{pq}$ at that time.

It suffices to prove that the co-circularity of p, q, r, a, which (by Lemma 4.4) occurs during the Delaunay crossing of pa by r, takes place when the crossing of pq by r is already finished (and, in particular, after the co-circularity of p, q, r, a that occurs during the crossing of pq).

Before the Delaunayhood of pa is restored, we have a co-circularity p, q, r, a in which q leaves $B[p, a, r] \cap L^-_{pa}$. (This is argued in the proof of Lemma 4.4: Right after the crossing, the point q lies in $B[p, a, r] \cap L^-_{pa}$, as in Figure 9 (left), and has to leave that disc before it becomes empty; it cannot cross pa during J, when this edge undergoes the Delaunay crossing by r). Notice that this is a red-blue co-circularity with respect to pa, and a red-red co-circularity with respect to pq; see Figure 9 (right).

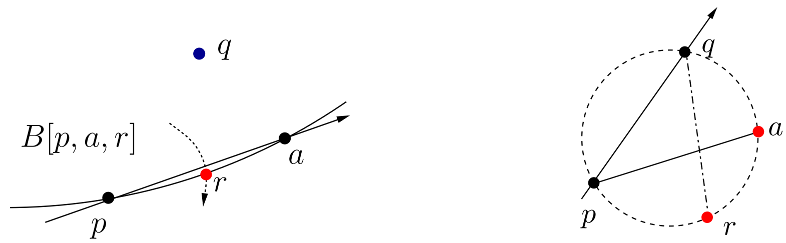

Figure 9: Proof of Lemma 4.6. Left: if r remains in L^+_{pq} after I and before it crosses pa, then q lies in $B[p,a,r] \cap L^+_{pa}$ before that last collinearity. Right: The second co-circularity of p,q,r,a which occurs when q leaves $B[p,a,r] \cap L^+_{pq}$. This is a red-red co-circularity with respect to pq, so the crossing of pq is already over.

Since no red-red or blue-blue co-circularities occur during a Delaunay crossing of an edge, the crossing of pq is already over. □

Consecutive crossings. By Lemma 4.6, 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. We say that clockwise (p,r)-crossings $(pq,r,I), (pa,r,J)$ are *consecutive* if they are consecutive in this order. More generally, we say that these crossings are k*-consecutive* if at most k other clockwise (p,r)-crossings separate them in this order.

Similar notions of consecutiveness and k-consecutiveness apply to pairs of counterclockwise (p,r)-crossings $(qp,r,I), (ap,r,J)$.

4.2 The roadmap

In Section 3 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 4.5.

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 [17], [25, Section 7.3.2], and also [22, Section 4.5] for further details on solving such systems of recurrences.)

In the predecessor paper [23], we used the following fairly direct charging strategy. For each single Delaunay crossing (pq,r,I) in P we first checked whether it (or its immediate neighbor) is near-extremal in the order implied by Lemma 4.6. Notice that (pq,r,I) appears (and thus can be extremal) in two restricted families of crossings: that of the clockwise (p,r)-crossings, and that of the counterclockwise (q,r)-crossings. If this were the case, we could charge (pq,r,I) to one of the edges pr and qr, for an overall quadratic bound. Otherwise, we applied Theorem 2.2 in the arrangements $\mathcal{A}_{pr}$ and $\mathcal{A}_{rq}$, and tried to charge (pq,r,I), within at least one of these two arrangements, either to a shallow collinearity, or to sufficiently many shallow co-circularities. Finally, if none of the previous chargings succeeded, we charged (pq,r,I) to some triple (not necessarily p,q,r) which performed two Delaunay crossings in some sub-sample of P, so our analysis bottomed out via (the weaker analogue in [23] of) Lemma 4.5.

Unfortunately, the above direct approach no longer works in the present setting, where any four points can be co-circular up to three times. Informally, its main weakness 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 3.1). The quadruples that we will shortly introduce more formally, 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, initially, 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

We next review the three types of topological configurations that arise in the course of our analysis, and highlight the intimate relations between these types of configurations, and Delaunay crossings.

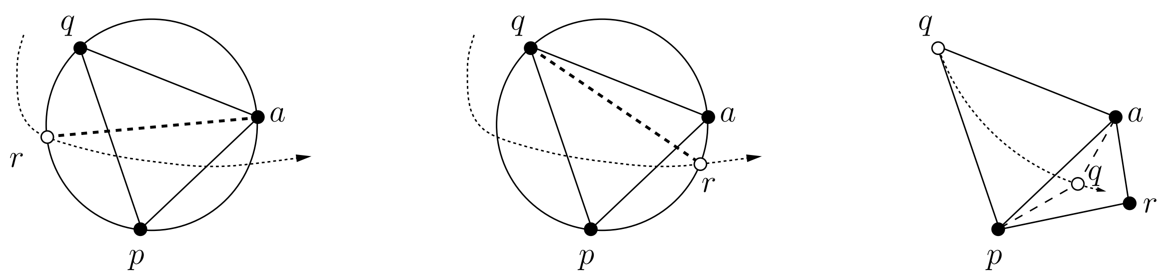

Figure 10: A (clockwise) regular quadruple $\sigma = (p, q, a, r)$, which is composed of clockwise (p, r)-crossings (pq, r, I) and (pa, r, J). Left and center: A possible motion of r, with the two co-circularities of p, q, a, r that occur during $I \setminus J$ and $J \setminus I$, respectively. Right: 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 10. We say that the quadruple is *consecutive* if (pq, r, I) and (pa, r, J) are consecutive.

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 4.6). 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 $\Psi(n)$ of consecutive regular quadruples.

Let $\sigma = (p, q, a, r)$ be a consecutive regular quadruple as above. By Lemma 4.1, 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 Section 5, 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 2.2 in the respective red-blue arrangements of these edges. Each such application of Theorem 2.2 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.

In Section 5.1, we show (via Lemmas 4.1 and 4.4) 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 10 (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 3, 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 in Section 5 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 4.5.

Unfortunately, there may still exist quadruples σ whose four resulting Delaunay crossings (including the two original (p, r)-crossings (pq, r, I) and (pa, r, J)) involve four distinct sub-triples p, q, a, r, so Lemma 4.5 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 10 (right). We refer to this latter crossing as a *special crossing* of pa by q, and pass the analysis of such crossings, each accompanied by a regular quadruple that induces it, to Section 6.

Special quadruples. In Section 6 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 implied by Lemma 4.6. See Figure 11.

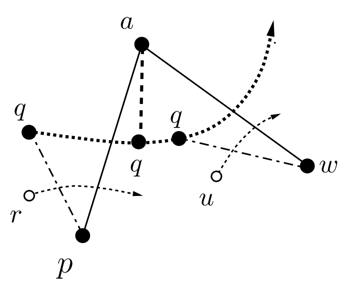

Figure 11: A (counterclockwise) special quadruple $\chi = (a, p, w, q)$, is composed of two special crossings $(pa, q, \mathcal{I})$ and $(wa, q, \mathcal{J})$, which respectively correspond to some (clockwise) regular quadruples (p, q, a, r) and (w, q, a, u).

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 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 $\sigma_2 = (w, q, a, u)$ that induce them.

At the final stage of the analysis (and only there), we use the above correspondence with the regular quadruples in order to charge the surviving special quadruples χ to especially convenient topological 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;[14] see Figure 12. 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. The analysis of these configurations is delegated to Section 7, where we directly bound their number in terms of simpler quantities, introduced in Section 2, and thereby complete the proof of Theorem 2.1. (We again emphasize that the recurrences that bound the number of terminal quadruples must have only "quadratic" terms.)

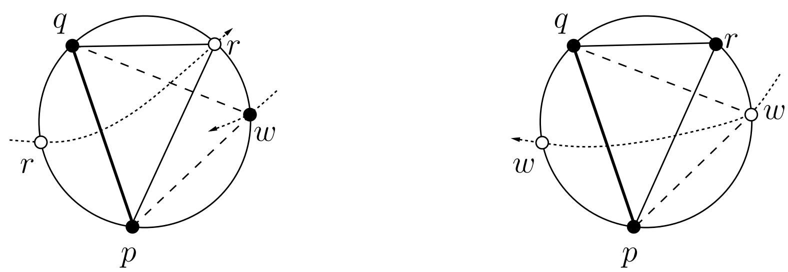

Figure 12: A terminal quadruple $\varrho = (p, q, r, w)$. The points r and w cross pq in opposite directions. The points of ϱ are co-circular three times. The extremal two co-circularities are red-blue with respect to pq, and the middle one is monochromatic with respect to pq. The left figure depicts 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 (in contrast to the one of regular quadruples) 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 12. 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). As a result, some sub-triple among p, q, r, w is involved in two Delaunay crossings, so Lemma 4.5 can always be invoked.

5 Regular Quadruples

5.1 Notation and topology

Definition. Four distinct points $p, q, a, r \in P$ form a *clockwise 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. We say that the quadruple is *consecutive* if (pq, r, I) and (pa, q, J) are consecutive. The definitions of a *counterclockwise quadruple* and of a consecutive counterclockwise quadruple are similar.

Each quadruple σ is equipped with the intervals $I_\sigma = I = [t_0, t_1]$ and $J_\sigma = J = [t_2, t_3]$ during which the corresponding edges pq and pa are absent from $\mathrm{DT}(P)$.

Recall that, by Theorem 4.3, any set of n moving points admits at most $O(n^2)$ double Delaunay crossings. Clearly, every clockwise (resp., counterclockwise) single (p, r)-crossing forms the first part of exactly one clockwise (resp., counterclockwise) consecutive quadruple, unless it is the last such (p, r)-crossing (with respect to the order given by Lemma 4.6). The overall number of these last crossings is clearly bounded by $O(n^2)$. Therefore, using $\Psi(n)$ to denote in maximum possible number of consecutive

[14]The letters p, q, r, w designate the way in which a terminal quadruple is extracted from the 6-point configuration of the surviving special quadruple $\chi = (a, p, w, q)$ and its respective pair of regular quadruples $\sigma_1 = (p, q, a, r)$ and $\sigma_2 = (w, q, a, u)$.

clockwise quadruples in a set of n moving points, we have the following obvious bound on the maximum number $C(n)$ of all Delaunay crossings:

$$C(n) \leq \Psi(n) + O(n^2).$$

The topology of quadruples. According to Lemma 4.4, the points of a clockwise quadruple σ are involved in at least one co-circularity during I_σ, and in at least one co-circularity during J_σ. Specifically, the former co-circularity is red-blue with respect to pq (and monochromatic with respect to pa), so it occurs before the beginning of J_σ, during $I_\sigma \setminus J_\sigma$. Similarly, the latter co-circularity is red-blue with respect to pa (and monochromatic with respect to pq), so occurs after the end of I_σ, during $J_\sigma \setminus I_\sigma$.

Notice that the points p, q, r, a are involved in exactly one co-circularity during each of the intervals I, J. Indeed, recall that the point a lies outside the disc $B[p, q, r]$ right before I_σ begins and right after I_σ ends. Moreover, $B[p, q, r]$ switches instantly from L^+_{pq} to L^-_{pq} only once during I_σ, so a hits the boundary $B[p, q, r]$ an odd number of times during I_σ. A symmetric behaviour takes place during J_σ, so the points p, q, a, r are involved in exactly one co-circularity in each interval.

Lemma 5.1. *Let $\sigma = (p, q, a, r)$ be a clockwise quadruple with the associated Delaunay crossings $(pq, r, I_\sigma = [t_0, t_1])$ and $(pa, r, J_\sigma = [t_2, t_3])$ (occuring in this order). Assume also that the point r hits pq again after I_σ and before r hits pa (and enters L^+_{pa}) during J_σ. Then (with the conventions assumed above) the edge rq is hit during (t_1, t_3) by the point a, which crosses L_{rq} from L^+_{rq} to L^-_{rq}.*

Since the roles of q and a in σ are interchangable (by reversing the direction of the time axis), we also have a symmetric variant of the lemma, which applies if r hits the edge pa before J_σ but after it hits pq during I_σ. Symmetric versions of the lemma and this subsequent also hold if σ is a counterclockwise quadruple.

Proof. Let ζ_1 denote the time in $J_\sigma \setminus I_\sigma$ when the points p, q, a, r are co-circular, and recall that this co-circularity is red-blue with respect to pa. Since any three points can be collinear at most twice, both points r, a lie in L^-_{pq} when r hits pa during J_σ (this is because r must lie in L^-_{pq} at that time, so a also has to lie there when r hits pa). Hence, q lies then in L^+_{pa}. Right before this event, q lies in the cap $B[p, q, r] \cap L^+_{pa}$. Arguing as in the proof of Lemma 4.4, the point q enters the above cap at time ζ_1; see Figure 13 (left). In addition, the point a leaves the cap $B[p, q, r] \cap L^-_{rq}$ at the very same time ζ_1.

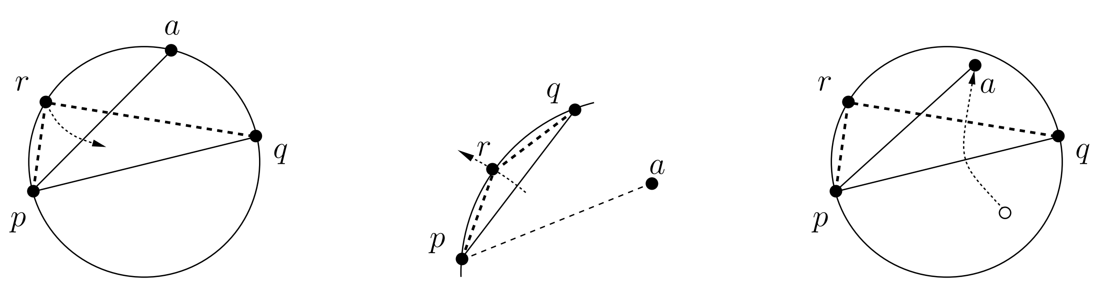

Figure 13: Illustrating the proof of Lemma 5.1. Left: If r hits pq again before crossing pa, then q enters $B[p, a, r]$ during the second co-circularity of p, q, a, r (and a leaves the cap $B[p, q, r] \cap L^-_{rq}$). Center: The case where a lies in the cap $B[p, q, r] \cap L^+_{pq}$ right after r returns to L^-_{pq}. Right: The point a can enter the cap $B[p, q, r] \cap L^-_{rq}$ (without leaving $B[p, q, r]$) only through rq.

In particular, the preceding discussion implies that the second collinearity of p, q, r occurs at some time $\tilde{t}$ before ζ_1. Since r can cross L_{pq} only twice, the motion of $B[p, q, r]$ remains continuous after time $\tilde{t}$ (when $B[p, q, r]$ instantly flips from L^-_{pq} to L^+_{pq}). We distinguish between the following two cases.

(i) Assume first that a lies in L^+_{pq} at time $\tilde{t}$, so it lies in the cap $B[p, q, r] \cap L^+_{pq}$ right afterwards; see Figure 13 (center). The lemma clearly holds if the point a remains in $B[p, q, r]$ during the interval $(\tilde{t}, \zeta_1)$.

Indeed, in this case a lies in $L^+_{rq} = L^+_{pq}$ at time $\tilde{t}$, so it can enter the cap $B[p,q,r] \cap L^+_{rq}$ (without leaving $B[p,q,r]$) only through the edge rq. See Figure 13 (right). Furthermore, a cannot leave $B[p,q,r]$ during $(\tilde{t}, \zeta_1)$, because it would have to re-enter $B[p,q,r]$ before ζ_1 (recall that it leaves $B[p,q,r]$ right after ζ_1). But then the points of σ would have been involved in at least *four* distinct co-circularities, one occuring during I_σ and before time $\tilde{t}$, the two co-circularities just considered, both occuring during $(\tilde{t}, \zeta_1)$, and one at ζ_1 itself. This contradiction establishes the lemma in case (i).

(ii) Now suppose that a lies in L^-_{pq} at time $\tilde{t}$. In this case, as in the proof of Lemma 4.4, a lies in $B[p,q,r]$ right before $\tilde{t}$. Since a lies outside $B[p,q,r]$ right after the end of I_σ (and since the motion of $B[p,q,r]$ is continuous between the two collinearities of p,q,r), the point a has to cross the boundary of $B[p,q,r]$ after I_σ and before $\tilde{t}$. In addition, the point a must now enter $B[p,q,r]$ during $(\tilde{t}, \zeta_1)$, because it lies outside $B[p,q,r]$ right after $\tilde{t}$. Once again, we obtain four distinct co-circularities of p,q,a,r, a contradiction that shows that case (ii) is impossible, and thus completes the proof. $\square$

Overview. In this section we analyze the maximum number of consecutive clockwise quadruples. The underlying intuition behind our (admittedly, faily involved) analysis is the following. We analyze quadruples of four points p,q,a,r. The purpose of the analysis is to charge these quadruples to special restricted configurations that are easier to analyze. Theorem 2.2 allows us to charge some quadruples to shallow co-circularities or collinearities, which forms the basis for various recurrences that the analysis will be deriving. In addition, Theorem 4.3 and Lemma 4.5 yield a quadratic bound for the number of quadruples that can be charged to a double Delaunay crossing of some triple of their points, or to two Delaunay crossings of the same triple.

Our strategy is therefore to filter away quadruples that can be charged by either of these tools, untill all quadruples are exhausted. To do so, we keep enforcing our quadruples to be involved with progressively more Delaunay crossings. Each quadruple is associated with four triples, and our goal is to force at least one triple of points to perform two Delaunay crossings, in which case Theorem 4.3 and Lemma 4.5 will yield the desired quadratic bounds.

Right from the start, a quadruple $\sigma = (p,q,a,r)$ already has, by definition, two Delaunay crossings: of pq by r, and of pa by r. To enforce additional crossings, we need a careful (and involved) analysis of the "topological" changes of the four moving points of σ, where each event is either a collinearity of three of the points (in which case the order type of p,q,a,r changes), or a co-circularity of the four points of σ (in which case the Delaunayhood of a pair of its edges "flips").

The analysis of consecutive clockwise quadruples proceeds through six stages, numbered $0,1,\ldots,5$. At the i-th stage we consider a certain family $\mathcal{F}_i$ of clockwise quadruples, which are defined with respect to an underlying set P of n points moving as above in $\mathbb{R}^2$. (Initially, $\mathcal{F}_0$ consists of all consecutive quadruples in the original point set P. In subsequent stages, P is a smaller sample from the original point set, but we continue, for simplicity, to denote it as P.) We assume that each quadruple σ in $\mathcal{F}_i$ satisfies certain topological conditions, which are formulated in terms of the four points of σ, other points of P (and, possibly, also nearby quadruples in $\mathcal{F}_i$). Our goal is to bound the maximum possible cardinality $\Psi_i(n)$ of $\mathcal{F}_i$. This is achieved by developing a system of recurrences, each expressing Ψ_i in terms of Ψ_{i+1}, except for Ψ_5, which is analyzed in Section 6. The overall solution of this system yields the desired near-quadratic bound.

5.2 Stage 0: Charging events in $\mathcal{A}_{pr}$

Let $\sigma = (p,q,a,r)$ be a consecutive clockwise quadruple, whose two Delaunay crossings occur during the intervals $I = I_\sigma = [t_0, t_1]$ and $J = J_\sigma = [t_2, t_3]$. By Lemma 4.1, the edge pr is Delaunay during each of the intervals I, J, but it may leave $\mathrm{DT}(P)$ during the possible gap between I and J.

Charging events in $\mathcal{A}_{pr}$. We fix a constant $k > 12$ and apply Theorem 2.2 in $\mathcal{A}_{pr}$ over the interval (t_1, t_3) (which covers the aforementioned gap between I and J, if it exists).

First, assume that at least one of the Conditions (i), (ii) of Theorem 2.2 holds. In this case, we charge σ either to a k-shallow collinearity, or to $\Omega(k^2)$ k-shallow co-circularities, that occur in $\mathcal{A}_{pr}$ during (t_1, t_3). We claim that any k-shallow collinearity or co-circularity in $\mathcal{A}_{pr}$ is charged in this manner by at most $O(1)$ quadruples. Indeed, consider the moment t^* when the charged event occurs, and notice that it involves p and r (together with one or two additional points of P). After guessing p and r (in $O(1)$ ways), σ is the unique quadruple (p, q, a, r) for which the interval $[t_1, t_3]$, delimited by the ending times of the two corresponding Delaunay crossing intervals, contains t^* (by definition of consecutive quadruples, the intervals $[t_1, t_3]$ are pairwise openly disjoint, for p and r fixed).

Using the standard bounds on the number of k-shallow collinearities and co-circularities (established in Sections 2 and 3), in combination with the fact that each co-circularity pays only $\Theta(1/k^2)$ units when it is charged, we get that the number of such quadruples σ for which the red-blue arrangement of pr satisfies one of the Conditions (i), (ii) of Theorem 2.2, is $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$.

Assume then that the red-blue arrangement of pr (during (t_1, t_3)) satisfies Condition (iii) of Theorem 2.2. That is, one can restore the Delaunayhood of pr during (t_1, t_3) by removing a set A of at most $3k$ points of P (possibly including q and/or a).[15] We now consider a random subset R of $\Theta(n/k)$ points of P. By the standard probabilistic argument of Clarkson and Shor [9], the following two events occur simultaneously with probability at least $\Theta(1/k^4)$: (1) $p, q, a, r \in R$, and (2) none of the points of $A \setminus \{a, q\}$ belong to R.

Condition (1) guarantees that the smaller set R induces Delaunay crossings $(pq, r, I_R = [t'_0, t'_1])$ and $(pa, r, J_R = [t'_2, t'_3])$, such that $I_R \subseteq I$ and $J_R \subseteq J$. (The latter property follows because the intervals of non-Delaunayhood of pq can only shrink as we pass to the triangulation $\mathrm{DT}(R)$ of the reduced set R.) In particular, both of these crossings are single Delaunay crossings. Clearly, (pq, r, I_R) is followed by (pa, r, J_R) in the order implied by Lemma 4.6. In other words, the four points p, q, a, r define within $\mathrm{DT}(R)$ a clockwise quadruple σ_R. Recall that pr is Delaunay during each of the intervals I, J. Condition (2) guarantees that pr belongs to $\mathrm{DT}(R \setminus \{q, a\})$ throughout the interval $[t_1, t_3]$ which covers the possible gap between I and J. In particular, this edge belongs to $\mathrm{DT}(R \setminus \{q, a\})$ throughout the extended interval $[I_R, J_R] = [t'_0, t'_3]$ which consists of I_R, J_R, and the possible gap between them. See Figure 14 (left). (As a matter of fact, the Delaunayhood of pr in $R \setminus \{q, a\}$ extends (at least) to the bigger interval $[t_0, t_3]$.)

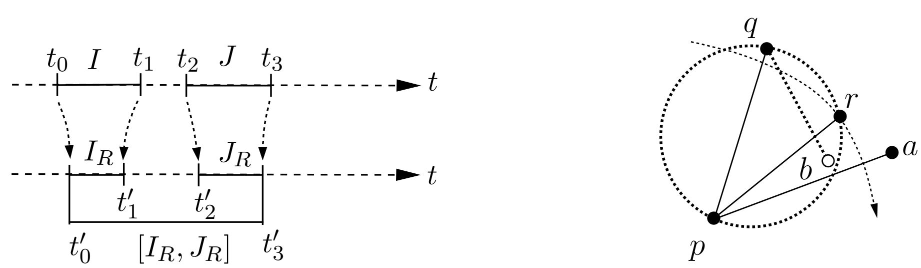

Figure 14: Left: The edge pr of σ_R belongs to $\mathrm{DT}(R \setminus \{q, a\})$ throughout $[I_R, J_R]$, including the gap between I_R and J_R. Right: Any violating pair of pr in R, such as the pair q, b, must involve either q or a.

To recap, we can charge σ to its more refined counterpart σ_R, formed by the pair of crossings (pq, r, I_R) and (pa, r, J_R), which shows up in the smaller triangulation $\mathrm{DT}(R)$, with probability at least $\Theta(1/k^4)$.

Let $\mathcal{F}_R$ denote the family of all such "hereditary" quadruples $\sigma_R = (p, q, a, r)$, each of them corresponding to some consecutive clockwise quadruple $\sigma = (p, q, a, r)$ in P, as defined above. Notice that the quadruples of $\mathcal{F}_R$ are not necessarily consecutive in R, as the set R may induce additional Delaunay

[15] Note that, if the gap between I and J does not exist, then $A = \emptyset$.

crossings that do not show up in $\mathrm{DT}(P)$. Below we introduce a weaker notion of consecutiveness, which holds for the quadruples of $\mathcal{F}_R$. In the definitions below, P stands for a generic set, which in general is a proper subsample of the original input.

Definition. We say that a quadruple $\sigma = (p, q, a, r)$ is *Delaunay* if the edge pr belongs to $\mathrm{DT}(P \setminus \{q, a\})$ throughout the interval $[I_\sigma, J_\sigma] = \mathsf{conv}(I_\sigma \cup J_\sigma)$.

Definition. Let $\mathcal{F}$ be a family of clockwise quadruples. We say that $\mathcal{F}$ is *nonoverlapping* if for any two quadruples $\sigma_1 = (p, q_1, a_1, r)$ and $\sigma_2 = (p, q_2, a_2, r)$, that share their first and last points, the clockwise (p, r)-crossings corresponding to σ_1 and σ_2 are distinct, except for the possibility $a_1 = q_2$ or $a_2 = q_1$, and occur in non-interleaving order. That is, in the order implied by Lemma 4.6, the two crossings (pq_1, r, I_1) and (pa_1, r, J_1) of σ_1 appear either both before or both after the two crossings (pq_2, r, I_2) and (pa_2, r, J_2) of σ_2 (again, with the possible coincidence of the second of one quadruple and the first crossing crossing of the other).

We say that a Delaunay crossing (pq, r, I) is *in* $\mathcal{F}$ if it is either the first or the second crossing for at least one quadruple σ in $\mathcal{F}$. (In total, it may show up in at most two quadruples.)

Notice that, as argued above, the "sampled" subfamily $\mathcal{F}_R$ includes only Delaunay quadruples. Moreover, $\mathcal{F}_R$ is nonoverlapping, as the Delaunay crossings in $\mathcal{F}_R$ (which are defined in terms of R) inherit the order, implied by Lemma 4.6, of their ancestors in P (that is, in $\mathcal{F}$).

In the rest of this section, the underlying family $\mathcal{F}$ is typically fixed at each stage of our analysis, and is assumed to be nonoverlapping, and to consist only of Delaunay quadruples. In particular, by the "nonoverlapping" property, any ordered triple (p, q, r) in P will define the first (resp., second) crossing (pq, r, I_σ) (resp., (pq, r, J_σ)) for at most one quadruple in $\mathcal{F}$. In other words, the following condition holds:

Proposition 5.2. *Let $\mathcal{F}$ be a nonoverlapping family of clockwise quadruples. Then every quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$ is uniquely determined by each of the ordered triples (p, q, r) and (p, a, r) of its points, which specify, respectively, the first crossing (pq, r, I) and the second crossing (pa, r, J) associated with σ.*

Let $\Psi(n)$ be the maximum number of consecutive quadruples that can be defined by a set of n points moving as above in $\mathbb{R}^2$. Let $\Psi_0(n)$ be the maximum cardinality of a nonoverlapping family $\mathcal{F}$ of Delaunay quadruples, which is defined with respect to a set of n such moving points. Then the quantities $\Psi(n)$ and $\Psi_0(n)$ are related by the recurrence

$$\Psi(n) = O\left(k^4 \Psi_0(n/k) + k^2 N(n/k) + kn^2 \beta(n)\right), \tag{3}$$

where $k \le n$ is an arbitrary parameter.

5.3 Stage 1

To bound the above quantity $\Psi_0(n)$, we fix the underlying point set P and the nonoverlapping family $\mathcal{F}$ of Delaunay quadruples. In addition, we fix a pair of constants $k \ll \ell$.

Let $\sigma = (p, q, a, r)$ be a Delaunay quadruple in $\mathcal{F}$ whose two Delaunay crossings occur during the intervals $I = I_\sigma = [t_0, t_1]$ and $J = J_\sigma = [t_2, t_3]$. Recall that (by Lemma 4.4) the points of σ are involved in two co-circularities, one during $I \setminus J$ and one during $J \setminus I$. (The former co-circularity is red-blue with respect to pq, and the latter one is red-blue with respect to pa.) Denote by $\zeta_0 \in I \setminus J$ and $\zeta_1 \in J \setminus I$ the times when these co-circularities occur. Clearly, at least one of these co-circularities of p, q, a, r has to be extremal. Without loss of generality, suppose that the co-circularity at time ζ_0 is the first co-circularity of the points of σ.

Our analysis (at this stage) proceeds by distinguishing between several possible scenarios, and treating each of them separately. In all but the last case, we will obtain a bound in terms of quantities that were already introduced. In the last case (case (e)), the bound will also depend on the cardinality of a more specialized subfamily of quadruples, which is defined over an appropriate subsample of P. Such families are called 1-*refined*, and their analysis is passed on to the subsequent stages.

Case (a). The edge pr is hit during $[t_0, t_3]$ by at least one of the points q, a. In fact, Lemma 4.1 implies that this additional collinearity must occur during the gap (t_1, t_2) (after I and before J), so I and J are disjoint in this case. See Figure 15 (left).

Assume, for instance, that pr is hit by q. Since σ is a Delaunay quadruple, the edge pr belongs to $\mathrm{DT}(P)$ at each of the times t_0, t_3, and it belongs to the pruned triangulation $\mathrm{DT}(P \setminus \{a, q\})$ throughout $[t_0, t_3]$. It thus follows that the edge pr undergoes a Delaunay crossing by q within the triangulation $\mathrm{DT}(P \setminus \{a\})$. That is, the triple p, q, r defines two Delaunay crossings (of distinct order types) within this smaller triangulation. A routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor [9] (in which we sample, say, half of the points) yields an upper bound of $O(n^2)$ on the overall number of such triples p, q, r in P (independently of the fourth point a). Since each Delaunay quadruple (p, q, a, r) in $\mathcal{F}$ is uniquely determined by the respective ordered triple (p, q, r) (as its first crossing), the same upper bound also holds for the overall number of such Delaunay quadruples in $\mathcal{F}$.

A similar counting argument applies if pr is hit by a during $[t_0, t_3]$. Namely, we argue that the edge pr undergoes a Delaunay crossing by a within the triangulation $\mathrm{DT}(P \setminus \{q\})$, so the triple p, a, r defines two Delaunay crossings within that reduced triangulation, and the quadratic bound follows from Lemma 4.5, as above. Hence we may assume, from now on, that pr is not hit by q or a during $[t_0, t_3]$.

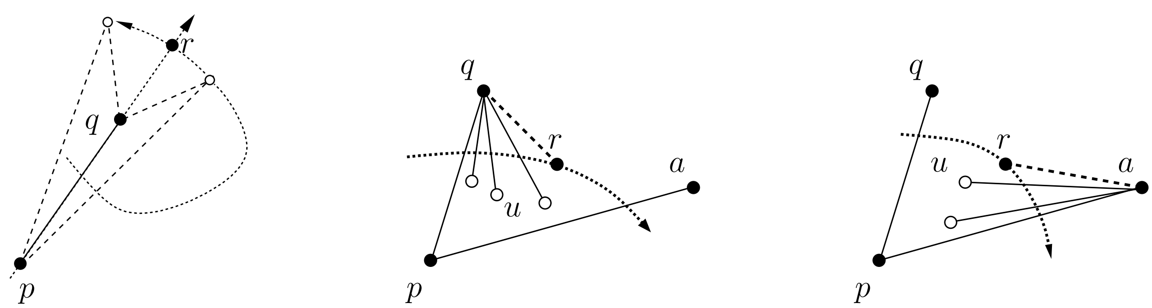

Figure 15: Left: Case (a). The edge pr is hit by q during (t_1, t_2). Center: Case (b). At least k counterclockwise (q, r)-crossins (uq, r, I_u) end during $(t_1, t_3]$. Right: Case (b) – the symmetric scenario. At least k counterclockwise (a, r)-crossings (ua, r, I_u) begin during $[t_0, t_2)$.

Case (b). At least k counterclockwise (q, r)-crossings (uq, r, I_u) end during $(t_1, t_3]$ (see Figure 15 (center)), or at least k counterclockwise (a, r)-crossings (ua, r, I_u) start during $[t_0, t_2)$ (see Figure 15 (right)). To dispose of such quadruples σ, we introduce an auxiliary counting scheme that we will use at several stages of our analysis. We first need a few definitions.

Chargeability. We say that an edge pq is *almost Delaunay* during an interval $\mathcal{I} = [t_0, t_1]$ if there is a set A of at most c_0 points such that pq belong to $\mathrm{DT}(P \setminus A)$ throughout $\mathcal{I}$. Here c_0 is some absolute constant[16] smaller than 8.

We say that a Delaunay crossing $(pq, r, I) = [t_0, t_1]$ is (p, r, k)-*chargeable* if there exists an interval $\mathcal{I} = [\alpha_0, \alpha_1]$ containing I such that the following two conditions hold: (1) the edge pr is Delaunay at times α_0 and α_1, and almost Delaunay during the the rest of $\mathcal{I}$, and (2) at least k counterclockwise

[16]This condition is similar to Condition (iii) in Theorem 2.2, except that here c_0 is a small *absolute* constant, whereas the parameter k in the theorem can be, and is indeed set to, a suitable large value that grows as $\varepsilon \downarrow 0$.

(q,r)-crossings (uq,r,I_u) occur within $\mathcal{I}$ (i.e., we have $I_u \subseteq \mathcal{I}$ for each of these points u). See Figure 16.

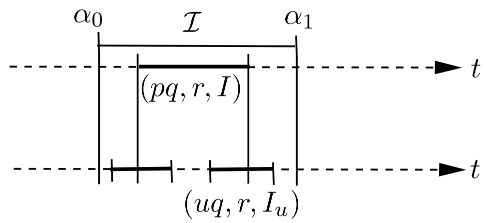

Figure 16: The crossing (pq,r,I) is (p,r,k)-chargeable with reference interval $\mathcal{I} = [\alpha_0, \alpha_1]$. At least k counterclockwise (q,r)-crossings (uq,r,I_u) occur within $\mathcal{I}$. By Lemma 4.6, each of their respective intervals I_u is contained in exactly one of the intervals $[t_0, \alpha_1]$, $[\alpha_0, t_1]$.

Similarly, we say that a Delaunay crossing (pq,r,I) is (q,r,k)-*chargeable* if the edge qr is almost Delaunay throughout the extended interval $\mathcal{I}$ (and Delaunay at the endpoints of $\mathcal{I}$), and at least k clockwise (p,r)-crossings (pu,r,I_u) occur within $\mathcal{I}$.

Several remarks are in order. If (pq,r,I) is a (p,r,k)-chargeable crossing then it need not be the *only* clockwise (p,r)-crossing to occur within the corresponding interval $\mathcal{I} = [\alpha_0, \alpha_1]$. Moreover, the other such (p,r)-crossings (pz,r,I_z), that occur (if at all) within $\mathcal{I}$, are not necessarily (p,r,k)-chargeable (because this notion also depends on the other endpoint z of the edge pz being crossed by r). Note also that, according to (a counterclockwise variant of) Lemma 4.6, each of the clockwise (q,r)-crossings (uq,r,I_u) that contribute to the (p,r,k)-chargeability of (pq,r,I) must satisfy either $I_u \subseteq [\alpha_0, t_1]$ or $I_u \subseteq [t_0, \alpha_1]$, because the intervals I and I_u are either disjoint or partially overlapping (but not nested).

Informally, the (p,r,k)-chargeability allows us to distribute the "weight" of (pq,r,I) over the $\Omega(k)$ arrangements $\mathcal{A}_{ru}$, which correspond to the above counterclockwise (q,r)-crossings (uq,r,I_u) (each of these latter crossings is also a clockwise (u,r)-crossing, and is denoted this way). In Section 8 we use this idea to establish the following theorem:

Theorem 5.3. *Let $k > 12$ be a sufficiently large constant. Then any set P of n points, moving as above in $\mathbb{R}^2$, induces at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$ Delaunay crossings (pq,r,I) that are either (p,r,k)-chargeable or (q,r,k)-chargeable.*

We next return to the setup of the first subcase of Case (b). Since σ is a Delaunay quadruple, the edge pr is almost Delaunay during the interval $[t_0, t_3]$ (it suffices to remove q to a to ensure Delaunayhood). According to Lemma 4.6, each of the (q,r)-crossings (uq,r,I_u) occurs entirely within $I \cup [t_1, t_3] = [t_0, t_3]$, that is, $I_u \subseteq [t_0, t_3]$. Indeed, by definition, each such I_u ends before t_3 and after t_1, the end of I, so it has to start after t_0, where I starts. Thus, (pq,r,I) is (p,r,k)-chargeable (with $\mathcal{I} = [t_0, t_3]$). Hence, by Theorem 5.3, the overall number of the corresponding quadruples σ is at most

$$O(k^2 N(n/k) + kn^2\beta(n)).$$

A symmetric argument applies if at least k counterclockwise (a,r)-crossings (ua,r,I_u) begin in $[t_0, t_2]$. Indeed, arguing as in the preceding paragraph, each of these Delaunay crossings has to occur entirely within $[t_0, t_3] = [t_0, t_2] \cup J$, so (pa,r,J) is (p,r,k)-chargeable.

Hence, we may assume, from now on, that at most k counterclockwise (q,r)-crossings end in $(t_1, t_3]$, and that at most k counterclockwise (a,r)-crossings begin in $[t_0, t_2)$.

Case (c). Either rq is never Delaunay during $[t_3, \infty)$, or ra is never Delaunay during $(-\infty, t_0]$. In the former case, by Lemma 4.1, no counterclockwise (q,r)-crossings can end in $[t_3, \infty)$, because rq has to be Delaunay throughout the interval of such a crossing. Since case (b) is ruled out, (pq,r,I) is among the last $k+1$ counterclockwise (q,r)-crossings (with respect to the order implied by Lemma 4.6). Clearly, this can happen for at most $O(kn^2)$ crossings (pq,r,I) (and their respective quadruples σ). A fully symmetric argument applies if ra never shows up in $\mathrm{DT}(P)$ during $(-\infty,t_0]$, in which case (pa,r,J) is among the first $k+1$ counterclockwise (a,r)-crossings.

Preparing for cases (d) and (e). In the remainder of our analysis we may therefore assume that neither of the situations considered in cases (a)–(c) arises. Let t_{rq} denote the first time in $[t_3,\infty)$ when rq belongs to $\mathrm{DT}(P)$. Namely, we have $t_{rq}=t_3$ if rq is Delaunay also at time t_3, and otherwise t_{rq} is the first time after t_3 when rq enters $\mathrm{DT}(P)$ (recall that rq is Delaunay at time t_1); refer to the schematic Figure 17 (left). Similarly, we let t_{ra} denote the last time in $(-\infty,t_0]$ when ra belongs to $\mathrm{DT}(P)$; see Figure 17 (right).

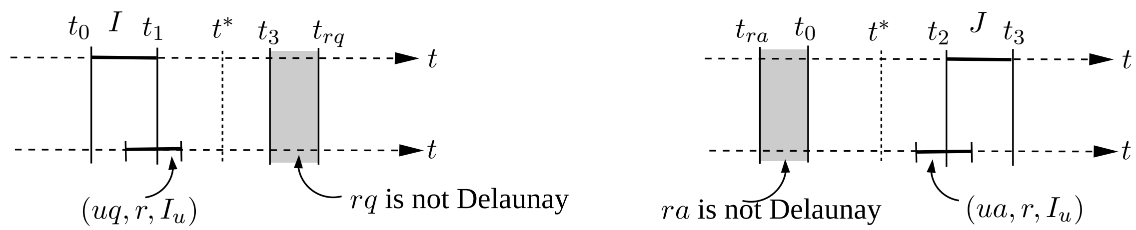

Figure 17: Charging events in $\mathcal{A}_{rq}$ and $\mathcal{A}_{ra}$. Left: t_{rq} is the first time in $[t_3,\infty)$ when rq belongs to $\mathrm{DT}(P)$. Since case (b) is ruled out, (pq,r,I) is among the last $k+1$ counterclockwise (q,r)-crossings to end before any event in (t_1,t_{rq}). Right: t_{ra} is the last time in $(-\infty,t_2]$ when ra belongs to $\mathrm{DT}(P)$. After outruling case (b), (pa,r,J) is among the first $k+1$ counterclockwise (u,r)-crossings to begin after any event in (t_{ra},t_2).

Before proceeding to the cases (d) and (e), we first apply Theorem 2.2 in $\mathcal{A}_{rq}$ over the interval (t_1,t_{rq}), and then apply it in $\mathcal{A}_{ra}$ over (t_{ra},t_2), both times with the second constant parameter ℓ.

Consider the first application of Theorem 2.2. If at least one of its Conditions (i), (ii) holds, we charge the quadruple σ, via its first crossing[17] (pq,r,I), either to $\Omega(\ell^2)$ ℓ-shallow co-circularities, or to an ℓ-shallow collinearity in $\mathcal{A}_{rq}$. We claim that each of these ℓ-shallow co-circularities or collinearities that occurs at some moment $t^*\in(t_1,t_{rq})$, is charged at most $O(k)$ times in this manner. Indeed, such an event must involve the points q and r of σ (together with one or two additional points). To guess the point p, we use the fact that at most k counterclockwise (q,r)-crossings end after I and before t_3. Moreover, assuming $t_{rq}>t_3$ and recalling Lemma 4.1, no (q,r)-crossings can take place (let alone end) during $(t_3,t_{rq}]$ (when the edge rq is not Delaunay). Thus, pq is among the $k+1$ edges whose counterclockwise (q,r)-crossings (by r) are the latest to end before t^*. Therefore, the overall number of quadruples σ in $\mathcal{F}$ for which such a charging applies is at most

$$O\left(k\ell^2N(n/\ell)+k\ell n^2\beta(n)\right).$$

Finally, if Condition (iii) of Theorem 2.2 holds, then the Delaunayhood of rq can be restored, throughout the interval $I\cup[t_1,t_{rq}]=[t_0,t_{rq}]$ (recall that rq is Delaunay during I), by removing a set A of at most 3ℓ points of P (possibly including p and/or a).

The second application of Theorem 2.2 in $\mathcal{A}_{ra}$ over (t_{ra},t_2) is fully symmetric. If at least one of Conditions (i), (ii) is satisfied, we dispose of σ by charging it, via its second crossing (pa,r,J), either to $\Omega(\ell^2)$ ℓ-shallow co-circularities, or to an ℓ-shallow collinearity that occur in $\mathcal{A}_{ra}$ during that interval. Arguing as above, (pa,r,J) is among the first $k+1$ counterclockwise (a,r)-crossings to begin after each charged event, which also involves a and r. Hence, every collinearity or co-circularity is charged at most $O(k)$ times, so, as above, this charging takes place for at most $O\left(k\ell^2N(n/\ell)+k\ell n^2\beta(n)\right)$ quadruples

[17] Recall that, according to Proposition 5.2, σ is uniquely determined by the choice of (p,q,r), which specify its first crossing.

σ. For each of the remaining quadruples we have a set B of at most 3ℓ points (possibly including p and/or q) whose removal restores the Delaunayhood of ra throughout $[t_{ra}, t_2] \cup J = [t_{ra}, t_3]$.

To recap, in each of remaining cases (d) and (e) we may assume the existence of the two sets A and B that satisfy the above properties. See Figure 18 (left) for a summary of what we assume now.

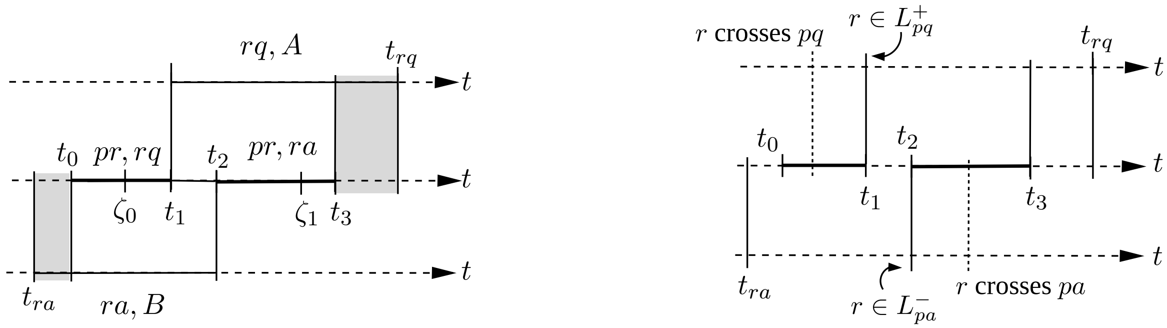

Figure 18: Left: The situation when entering case (d). If we remove $A \cup B$ but retain p, q, a, r, then: (i) During $[t_0, t_1]$, the edges pr and rq are Delaunay. (ii) During $[t_2, t_3]$, the edges pr and ra are Delaunay. (iii) During $[t_0, t_3]$, the edge pr is almost Delaunay. (iv) During $[t_0, t_{rq}]$, the edge rq is almost Delaunay (and will be Delaunay if we remove p and a). (v) During $[t_{ra}, t_3]$, the edge ra is almost Delaunay (and will be Delaunay if we remove p and q). Right: The situation when entering case (e). The point r can leave L^+_{pq} during $(t_1, t_{rq}]$ only through the edge pq. Similarly, r can enter L^-_{pa} during $[t_{ra}, t_2)$ only through the edge pa (and otherwise remains in L^-_{pa}).

Case (d). The point p hits the edge rq during (t_1, t_{rq}), or it hits the edge ra during the symmetric interval (t_{ra}, t_2). Without loss of generality, we focus on the former scenario, and handle the latter one in a fully symmetric manner.

As is easy to check, the edge rq undergoes a Delaunay crossing by p in $\mathrm{DT}((P \setminus A) \cup \{p\})$, with an appropriate interval that contains the time of the actual crossing. Therefore, Lemma 4.5, in combination with the Clarkson-Shor argument [9], provides an upper bound of $O(\ell n^2)$ on the number of such triples p, q, r (and of the corresponding quadruples σ, each of which is uniquely determined by the choice of (pq, r, I) as its first crossing).

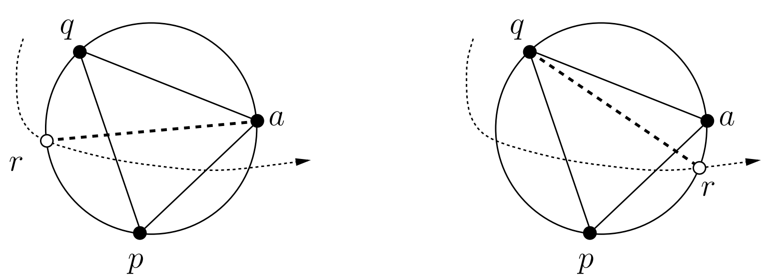

Figure 19: The co-circularities at times $\zeta_0 \in I \setminus J$ (left) and $\zeta_1 \in J \setminus I$ (right). In the depicted scenario, no additional collinearity of p, q, r or p, a, r occurs between the times when r enters L^+_{pq} and L^+_{pa}.

Case (e). None of the preceding cases holds; this is the most involved case in Stage 1. See Figure 18 (left and right) for a schematic summary of the following properties that we assume now. Recall that the points of σ are involved in co-circularities at times $\zeta_0 \in I \setminus J$ and $\zeta_1 \in J \setminus I$ (see Figure 19), and that at least one of these co-circularities has to be extremal. Without loss of generality, suppose, as already assumed earlier, that the co-circularity at time ζ_0 is the *first* co-circularity of the points of σ. In addition, we continue to assume that there exists a set A of cardinality at most 3ℓ, such that rq belongs to $\mathrm{DT}(P \setminus A)$ throughout the interval $[t_0, t_{rq})$. Similarly, we assume the existence of a set B of at most

3ℓ points such that ra belongs to $\mathrm{DT}(P \setminus B)$ throughout the interval $(t_{ra}, t_3]$. Finally, since neither of the preceding cases (a), (d) holds, r can re-enter the halfplane L^-_{pq} during $(t_1, t_{rq}]$ (after leaving it during $I = [t_0, t_1]$) only by crossing pq again; otherwise it remains in L^+_{pq} throughout $(t_1, t_{rq}]$. Similarly, r can enter L^-_{pa} during $[t_{ra}, t_2)$ (before leaving it during $J = [t_2, t_3]$) only through pa; otherwise it remains in L^-_{pa} throughout $[t_{ra}, t_2)$.

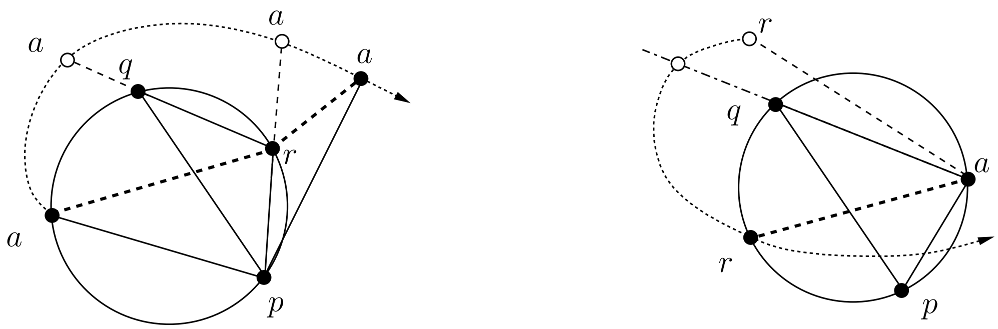

Figure 20: Case (e): proving that ra is hit by q. Left: a lies in L^-_{pq} when r enters L^+_{pq}, so r has to enter L^-_{pa} (through pa) afterwards and before J. The corresponding trajectory of a during (ζ_0, t_2) is depicted. Right: a lies in L^+_{pq} when r enters L^+_{pq}, so the Delaunayhood of ra is violated, right before ζ_0, by p and q.

We next argue[18] that the edge ra must be hit during $[t_{ra}, t_2)$ by the point q. We distinguish between two possible scenarios (see Figure 20).

(i) If a lies in $L^-_{pq} = L^-_{pr}$ when r enters L^+_{pq} (during I), then r has to enter L^-_{pa} before J. As noted above, r can enter L^-_{pa} only through pa, as depicted in Figure 20 (left). Therefore, according to a suitable variant of Lemma 5.1, in which the time is reversed and the points a and q are interchanged, the point q enters the halfplane L^-_{ra} during $[t_0, t_2]$, through ra, as claimed.

(ii) Now suppose that a lies in L^+_{pq} when r enters this halfplane, so the first co-circularity (at time ζ_0) occurs while r still lies in L^-_{pq}. Hence, the Delaunayhood of ra is violated, right before time ζ_0, by the points $q \in L^-_{ra}$ and $p \in L^+_{ra}$; see Figure 20 (right). Since ra is Delaunay at time t_{ra} and throughout $J = [t_2, t_3]$, and since the points p, q, a, r are never co-circular before ζ_0, Lemma 3.1 implies that at least one of the points p, q has to hit ra during the interval $[t_{ra}, \zeta_0)$, which is clearly contained in $[t_{ra}, t_2)$. (Specifically, we apply Lemma 3.1 so that the edge pq in the lemma is ra, the points a, b in the lemma are q, p, respectively, and the direction of the time axis is reversed.) Moreover, since case (d) does not occur, p cannot hit ra during the above interval. Hence, the other point, q, has to cross ra during $[t_{ra}, \zeta_0)$, from L^+_{ra} to L^-_{ra}.

If q hits ra twice during $[t_{ra}, t_2)$, then the triple q, a, r defines either a double Delaunay crossing, or two single crossings, which occur in the smaller triangulation $\mathrm{DT}((P \setminus B) \cup \{q\})$. Therefore, we can use Theorem 4.3, or Lemma 4.5, in combination with the Clarkson-Shor technique, to show that the overall number of such triples in P is at most $O(\ell n^2)$. Moreover, knowing q, a, r allows us to guess p in at most $O(k)$ possible ways, as (pa, r, J) is one of the first $k+1$ counterclockwise (a, r)-crossings to begin after the above collinearity (or collinearities) of q, a, r (this follows since we assume that case (b) does not arise). Hence, this scenario happens for at most $O(k\ell n^2)$ quadruples $\sigma \in \mathcal{F}$.

Assume then that ra is hit by q exactly once during (t_{ra}, t_2). In this only remaining case, the edge ra or, more precisely, its reversely oriented copy ar undergoes (within $[t_{ra}, t_2)$) exactly one (single) Delaunay crossing by q in the smaller triangulation $\mathrm{DT}((P \setminus B) \cup \{q\})$. To handle these latter quadruples σ, we apply a similar analysis to the edge rq (keeping in mind that the co-circularity at time ζ_1 is *not* necessarily extremal).

[18]Here the symmetry between q and a breaks down, because the co-circularity at ζ_0 is extremal, but the one at ζ_1 is not.

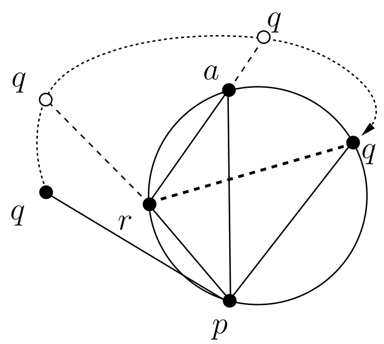

Figure 21: Case (e): The proposed trajectory of q if r re-enters L^-_{pq} before crossing pa. According to Lemma 5.1, the point a must hit the edge rq during $(t_1, t_3) \subseteq (t_1, t_{rq}]$.

If rq is hit by a during $(t_1, t_{rq}]$, then the points q, a, r define two single[19] Delaunay crossings in the triangulation $\mathrm{DT}([P \setminus (A \cup B)] \cup \{q, a\})$. A routine combination of Lemma 4.5 with the probabilistic arugment of Clarkson and Shor shows that the overall number of such triples q, a, r is at most $O(\ell n^2)$. Moreover, (pq, r, I) is among the $k + 1$ last counterclockwise (q, r)-crossings to end before the second collinearity of q, a, r. Thus, one can guess σ, based on q, a, r, in at most $O(k)$ possible ways. In conclusion, the above scenario happens for at most $O(k\ell n^2)$ Delaunay quadruples of $\mathcal{F}$.

To recap, the previous chargings account for

$$O\left(k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right)$$

Delaunay quadruples σ in $\mathcal{F}$. Hence, recalling that case (d) has been ruled out, we may assume, from now on, that none of the points p, a hits rq during the interval $[t_1, t_{rq}]$ (which contains $[t_1, t_3]$). In particular, this implies that q lies in $L^-_{pa} = L^-_{pr}$ at the moment when r enters L^+_{pa} during J (i.e., r lies then in L^+_{pq}). Indeed, otherwise r would have to first leave L^+_{pq} after I, necessarily through the edge pq (because cases (a) and (d) do not occur), which is now impossible according to Lemma 5.1. See Figure 21.

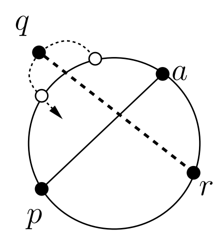

Figure 22: Case (e). The last two co-circularities of p, q, a, r that occur at times $\zeta_1 \in J \setminus I$ and $\zeta_2 \in (\zeta_1, t_{rq}] \setminus J$. The edges pa and rq intersect throughout (ζ_1, ζ_2); that is, the order type of p, q, a, r does not change there.

Since q lies in L^-_{pa} when r crosses pa (during J) from L^-_{pa} to L^+_{pa}, the Delaunayhood of rq is violated right after time ζ_1 by the points $p \in L^-_{rq}$ and $a \in L^+_{rq}$, as depicted in Figure 19 (right). (In other words, ζ_1 must occur *after* r enters L^+_{pa}, when q leaves the cap $B[p, a, r] \cap L^-_{pa}$.) Since neither of p, a can cross rq during the interval $(\zeta_1, t_{rq}]$ (which is clearly contained in $(t_1, t_{rq}]$), Lemma 3.1 implies that the points p, q, a, r are involved during this interval in a third co-circularity, at some time $\zeta_2 > \zeta_1$, and the Delaunayhood of rq is violated by p and a throughout the interval (ζ_1, ζ_2); see Figure 22. As a matter of fact, the discussion preceding Lemma 5.1 also implies that ζ_2 occurs after J.

Recall that each of the remaining quadruples σ is accompanied by a pair of subsets $A, B \subset P$, whose properties are detailed above. To facilitate the subsequent stages of our analysis, we augment

[19] Since any three points can be collinear at most twice, a can hit rq at most once.

the above "obstruction sets" A and B as follows. We add to A every point u for which there exists a counterclockwise (q, r)-crossing (uq, r, I_u) that ends in (t_1, t_{rq}). (In fact, Lemma 4.1 implies that none of these (q, r)-crossings end after t_3.) This is done to ensure that in the sampled configurations that we reach no such crossings take place. Similarly, we add to B every point u for which there exists a counterclockwise (a, r)-crossing (ua, r, I_u) that begins in (t_{ra}, t_2). (Again, Lemma 4.1 implies that none of these (a, r)-crossings begin before t_0.) Since we assume that case (b) does not hold, the above augmentation increases the cardinality of each of the sets A, B by at most $k \le \ell$.

Remark. We may assume that a is not among the (at most k) points lately added to A, and that q is not among the (at most k) points lately added to B. Indeed, if the edge qa (or its reversely oriented copy aq) undergoes a Delaunay crossing by r then the triple q, a, r defines two Delaunay crossings within $\mathrm{DT}((P \setminus A) \cup \{a\})$. By Lemma 4.5, the overall number of such triples is at most $O(\ell n^2)$. Furthermore, each of these triples is shared by at most $O(k)$ quadruples that fall into case (e), so the above scenario occurs for at most $O(k\ell n^2)$ quadruples of $\mathcal{F}$.

Probabilistic refinement. To proceed, we consider a subset R of $\lceil n/\ell \rceil$ points chosen at random from P. We fix a Delaunay quadruple σ as above (i.e., σ was not disposed of by the chargings of the previous cases, or by the previous chargings of case (e)), and notice that the following two events occur simultaneously, with probability at least $\Omega(1/\ell^4)$: (1) R includes the four points of σ, and (2) none of the points of $(A \cup B) \setminus \{p, q, a\}$ (for the augmented sets A, B) belong to R.

Consider the triangulation $\mathrm{DT}(R)$ which is induced by a "successful" sample R (satisfying (1) and (2)). Notice that the four points of σ still define a Delaunay quadruple, now with respect to R. We continue to denote this new quadruple by σ. (Note, however, that the suitably re-defined intervals $I = I_\sigma$ and $J = J_\sigma$ may shrink.)

Let $\mathcal{F}_R$ denote the family of all such "hereditary" Delaunay quadruples σ in R (such that the sample R is successful for their ancestors in $\mathcal{F}$). Clearly, $\mathcal{F}_R$ is nonoverlapping.

Fix a quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}_R$, whose two Delaunay crossings occur (within $\mathrm{DT}(R)$) during the intervals $I = [t_0, t_1]$, and $J = [t_2, t_3]$, and whose first two co-circularities occur at times $\zeta_0 \in I \setminus J$ and $\zeta_1 \in J \setminus I$. As before, let t_{ra} denote the last time in $(-\infty, t_0]$ when ra belongs to $\mathrm{DT}(R)$, and let t_{rq} denote the first time in $[t_3, \infty)$ when rq belongs the same triangulation $\mathrm{DT}(R)$. (Notice that, as we replace P by R, t_{ra} either remains unchanged or moves ahead, towards (the new) t_0. Symmetrically, t_{rq} stays the same or moves back, towards (the new) t_3. Hence, the extended intervals $[t_{ra}, t_3]$ and $[t_0, t_{rq}]$ can only shrink as we pass from $\mathrm{DT}(P)$ to $\mathrm{DT}(R)$.) The preceding analysis implies that the following conditions hold for σ:

(R1) No counterclockwise (a, r)-crossings in $\mathcal{F}_R$ begin during $[t_{ra}, t_2)$. Moreover, the edge ra belongs to $\mathrm{DT}(R \setminus \{p, q\})$ throughout the interval $[t_{ra}, t_3]$. See Figure 23 (left).

(R2) No counterclockwise (q, r)-crossings in $\mathcal{F}_R$ end during $(t_1, t_{rq}]$. Moreover, the edge rq belongs to $\mathrm{DT}(R \setminus \{p, a\})$ throughout the interval $[t_0, t_{rq}]$.

(R3) The set $R \setminus \{p\}$ induces a Delaunay crossing (ar, q, H), whose respective interval $H_\sigma = H$ is contained in $[t_{ra}, t_2]$. In addition, we encounter a third co-circularity of p, q, a, r at some time $\zeta_2 \in [t_3, t_{rq}]$, so that the Delaunayhood of rq is violated by $p \in L^-_{rq}$ and $a \in L^+_{rq}$ throughout (ζ_1, ζ_2). See Figures 22 and 23 (right). Finally, none of the points a, p crosses rq during $(\zeta_2, t_{rq}]$.

We say that a nonoverlapping family $\mathcal{F}$ of Delaunay quadruples in a set P is 1-*refined* if its quadruples satisfy the following modified three conditions, restated with respect to $\mathcal{F}$ and its underlying set P.

(Q1) No counterclockwise (a, r)-crossings in $\mathcal{F}$ begin during $[t_{ra}, t_2)$. Moreover, the edge ra belongs to $\mathrm{DT}(P \setminus \{p, q\})$ throughout the interval $[t_{ra}, t_3]$.

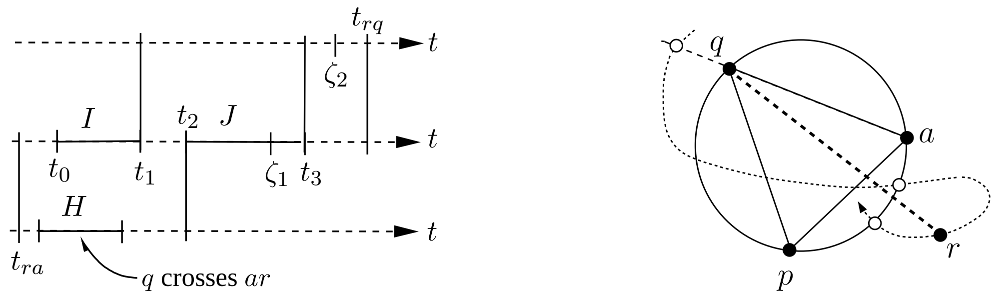

Figure 23: Left: The edge ar is crossed by q during $[t_{ra}, t_2)$. The interval $(t_3, t_{rq}]$ contains the third co-circularity ζ_2. The edges ar and rq are almost Delaunay during, respectively, $[t_{ra}, t_2) \cup J = [t_{ra}, t_3]$ and $I \cup (t_1, t_{rq}] = [t_0, t_{rq}]$. Right: A schematic description of the trajectory of r.

(Q2) No counterclockwise (q, r)-crossings in $\mathcal{F}$ end during $(t_1, t_{rq}]$. Moreover, the edge rq belongs to $\mathrm{DT}(P \setminus \{p, a\})$ throughout the interval $[t_0, t_{rq}]$.

(Q3) The set $P \setminus \{p\}$ induces a Delaunay crossing (ar, q, H), whose respective interval H is contained in $[t_{ra}, t_2]$. In addition, we encounter a third co-circularity of p, q, a, r at some time $\zeta_2 \in [t_3, t_{rq}]$, so that the Delaunayhood of rq is violated by $p \in L_{rq}^-$ and $a \in L_{rq}^+$ throughout (ζ_1, ζ_2). Finally, none of the points point a, p crosses rq during $(\zeta_2, t_{rq}]$.

Let $\Psi_1(m)$ denote the maximum possible cardinality of a 1-refined family of Delaunay quadruples, that is defined over a set of m points moving in $\mathbb{R}^2$ as above. The preceding discussion implies that the maximum cardinality $\Psi_0(n)$ of any nonoverlapping family $\mathcal{F}$ of Delaunay quadruples in a set of n moving points satisfies the recurrence:

$$\Psi_0(n) = O\left(\ell^4 \Psi_1(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right),$$

for any pair of parameters $k \ll \ell < n$.

Proposition 5.4. *Let $\mathcal{F}$ be a 1-refined family of Delaunay quadruples. Then each quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$ is uniquely determined by the ordered triple q, a, r. (That is, there is no other quadruple in $\mathcal{F}$ that shares its last three points with σ.)*

Proof. By Conditions (Q1) and (Q3), (pa, r, J_σ) is the first counterclockwise (a, r)-crossing (in $\mathcal{F}$) to begin after q hits ar during the corresponding interval $H = H_\sigma$. $\square$

The subsequent chargings — Overview. To bound the above quantity $\Psi_1(n)$, we fix an underlying set P of n moving points and a 1-refined family $\mathcal{F}$ of nonoverlapping Delaunay quadruples. In addition, we fix a quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$, whose Delaunay crossings occur during the intervals $I = I_\sigma = [t_0, t_1]$ and $J = J_\sigma = [t_2, t_3]$ (in this order). Recall that the points p, q, a, r are involved in three co-circularities, at times $\zeta_0 \in I \setminus J$, $\zeta_1 \in J \setminus I$, and $\zeta_2 > t_3$, and that the Delaunayhood of edge rq is violated during (ζ_1, ζ_2) by the points p and a. Furthermore, since the co-circularities at times ζ_1 and ζ_2 have the same order type, the Delaunayhood of pa is violated right after time ζ_2 by the points q and r.

Informally, the remainder of this section (except for Stage 5) is devoted to showing that the co-circularity at time ζ_2 yields a Delaunay crossing of pa by q. Similarly to the crossing of ar by q in Condition (Q3), this crossing occurs in an appropriately reduced triangulation, and only if σ is not previously disposed of by one of the standard chargings (using Theorems 2.2 and 5.3).

The above implication is relatively easy to establish if pa undergoes only few Delaunay crossings after (pa, r, J) and before time ζ_2, when it is violated by q and r. Indeed, following the general strategy demonstrated in Section 3 (and at Stage 1), we consider three possible scenarios.

If pa never re-enters $\mathrm{DT}(P)$ after time ζ_2, then (pa, r, J) (and, thereby, σ) can be charged to the edge pa, because it is then among the few last Delaunay crossings of this edge. Otherwise, we consider the first time t_{pa} after ζ_2 when pa enters $\mathrm{DT}(P)$ and apply Theorem 2.2 in $\mathcal{A}_{pa}$ over the interval $[t_3, t_{pa}]$. Notice that, according to Lemma 3.1, pa is crossed during this interval (or, more precisely, during its proper subinterval $(\zeta_2, t_{ra}]$) by at least one of r and q. (This follows because no further co-circularities of p, q, a, r can occur after ζ_2.)

If at least one of the Conditions (i), (ii) of Theorem 2.2 holds, we dispose of σ by charging it within $\mathcal{A}_{pa}$ (and, again, via its second crossing (pa, r, J)) either to sufficiently many shallow co-circularities, or to a shallow collinearity. As in the previous similar cases, the charging of each event in $\mathcal{A}_{pa}$ is almost unique, as (pa, r, J) is among the few last Delaunay crossings of pa to end before it.

Finally, if Condition (iii) of Theorem 2.2 holds, then we end up with a "small" subset A of P (including at least one of r, q) whose removal restores the Delaunayhood of pa throughout $[t_3, t_{pa}]$. Hence, pa undergoes, within a suitably sampled triangulation $\mathrm{DT}(R)$, a Delaunay crossing by one of the points q, r. If pa is crossed by r during $[t_3, t_{pa}]$, then we can again dispose of such quadruples σ using Lemma 4.5. Otherwise, we say that the edge pa undergoes within $\mathrm{DT}(R)$ a *special crossing* by the point q. By our assumption, each special crossing is charged by only a small number of triples p, a, r (and quadruples σ). In Section 6 we derive a recurrence for the maximum possible number of these special crossings, which, combined with the recurrences derived in this section, and in the preceding ones, yield the asserted near-quadratic bound on the number of Delaunay co-circularities.

Unfortunately, the above argument does not work if the edge pa of σ undergoes "too many" Delaunay crossings during (t_3, ζ_2). In this case, we cannot easily trace the events that occur in $\mathcal{A}_{pa}$, back to (pa, r, J) (and to σ); that is, there are too many ways to guess r. At Stage 4 we use Theorems 2.2 and 5.3 to dispose of such quadruples. To facilitate the fairly involved analysis of that stage, we first extend the almost-Delaunayhood of ra and rq from, respectively, $[t_{ra}, t_3]$ and $[t_0, t_{rq}]$, to their superinterval $[t_{ra}, t_{rq}]$, which covers $\zeta_0, \zeta_1, \zeta_2$ together with the aforementioned crossing of ar by q. These extensions are performed at the auxiliary Stages 2 and 3, and they also involve the sampling argument of Clarkson and Shor. (Hence, the instants t_{ra} and t_{rq} are each time redefined with respect to the underlying, progressively reduced subset of P.)

5.4 Stage 2: Charging events in $\mathcal{A}_{pr}$ (again)

Before extending the almost-Delaunayhood of ra and rq, as promised in the previous paragraph, we first tackle the edge pr, and extend its almost-Delaunayhood. Handling ra and rq will be done in the next Stage 3.

Let $\sigma = (p, q, a, r)$ be a quadruple in the 1-refined family $\mathcal{F}$. Recall that the edge pr is almost Delaunay during $[I, J] = [t_0, t_3]$ (and that it is in fact Delaunay if q and a are removed). We extend the almost-Delaunayhood of pr to a (potentially) larger interval $[\zeta_{pr}^-, \zeta_{pr}^+]$, which covers $[t_{ra}, t_{rq}]$. To do so, we fix a (new) pair of constants $k \ll \ell$.

Stage 2a. First, we consider the interval $[t_{ra}, t_3]$, where, by assumption, the edge ra is almost Delaunay. Refer to Figure 24 (left).

If at least k clockwise (p, r)-crossings (pu, r, J_u) begin in (t_{ra}, t_2), then the Delaunay crossing (pa, r, J) is (a, r, k)-chargeable with $\mathcal{I} = [t_{ra}, t_3]$. Indeed, according to Lemma 4.6, each of the corresponding intervals J_u has to be contained in $[t_{ra}, t_3] = [t_{ra}, t_2] \cup J$ (since J_u starts before t_2, the starting time of (pa, q, J), it has to end before t_3). Hence, and according to Theorem 5.3, the overall number of such crossings (pa, r, J) is at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$. Clearly, this also bounds the overall number of such quadruples σ. Therefore, we can assume, from now on, that at most k clockwise (p,r)-crossings (pu,r,I_u) begin during (t_{ra},t_2).

Figure 24: Left: Extending the almost-Delaunayhood of pr from $[t_0,t_3]$ to (ζ_{pr}^-,t_0) (left) and $(t_3,\zeta_{pr}^+]$ (right).

If the edge pr is never Delaunay during $(-\infty,t_{ra}]$, then (pq,r,I) and (pa,r,J) are among the first $k+1$ clockwise (p,r)-crossings, so there are only $O(kn^2)$ such crossings (and quadruples σ). Otherwise, let ζ_{pr}^- denote the last time in $(-\infty,t_{ra}]$ when pr belongs to $\mathrm{DT}(P)$.

We now apply Theorem 2.2 in $\mathcal{A}_{pr}$ over the interval (ζ_{pr}^-,t_2), with the threshold ℓ. Note that pr is Delaunay at times ζ_{pr}^- and t_2 (in addition to its being Delaunay throughout $I\subseteq[\zeta_{pr}^-,t_2)$). If at least one of the Conditions (i), (ii) of that theorem is satisfied, we charge σ (via (pa,r,J)) either to $\Omega(\ell^2)$ ℓ-shallow co-circularities, or to an ℓ-shallow collinearity. As in the previous such chargings, the crucial observation is that (pa,r,J) is among the first $k+1$ clockwise (p,r)-crossings to begin after each charged event in $\mathcal{A}_{pr}$. Hence, any ℓ-shallow co-circularity or collinearity is charged, as above, by at most $O(k)$ quadruples σ. Clearly, the above charging succeeds for at most $O\left(k\ell^2N(n/\ell)+k\ell n^2\beta(n)\right)$ quadruples σ in $\mathcal{F}$.

Finally, if Condition (iii) of Theorem 2.2 holds, we end up with a set A of at most 3ℓ points so that pr belongs to $\mathrm{DT}(P\setminus A)$ throughout the interval $[\zeta_{pr}^-,t_3]$. (Note that A can include one, or both of the points q,a.) For each (p,r)-crossing (pu,r,J_u) that begins in (t_{ra},t_0) we add the respective point u to the "obstruction set" A, whose cardinality then increases by at most $k\ll\ell$. (Informally, as earlier, this allows us to assume that, in the refined configuration, no such (p,r)-crossings occur.)

Stage 2b. We next consider the interval $[t_0,t_{rq}]$ where, by assumption, edge rq is almost Delaunay. Refer to Figure 24 (right). The argument is fully symmetric to the one in Stage (2a), but we repeat it for the sake of completeness.

If at least k clockwise (p,r)-crossings (pu,r,I_u) end in (t_1,t_{rq}), then the crossing (pq,r,I) is clearly (q,r,k)-chargeable with $\mathcal{I}=[t_0,t_{rq}]$, as each of the corresponding intervals I_u begins after t_0 (by Lemma 4.6). As before, this scenario happens for at most $O\left(k^2N(n/k)+kn^2\beta(n)\right)$ quadruples σ in $\mathcal{F}$. Hence, we may assume, from now on, that the above scenario does not happen for σ.

If pr is never Delaunay during $[t_{rq},\infty)$, then the crossings (pq,r,I) and (pa,r,J) are among the last $k+1$ clockwise (p,r)-crossings; as above, the number of these situations is $O(kn^2)$. Otherwise, let ζ_{pr}^+ denote the first time after t_{rq} when pr is Delaunay.

We now apply Theorem 2.2 in $\mathcal{A}_{pr}$ over the interval (t_1,ζ_{pr}^+), with the threshold ℓ (noting that pr is Delaunay at times t_0 and ζ_{pr}^+). If at least one of the Conditions (i), (ii) holds, we dispose of σ by charging it either to $\Omega(\ell^2)$ ℓ-shallow co-circularities, or to an ℓ-shallow collinearity. As before, each event in $\mathcal{A}_{pr}$ is charged at most $O(k)$ times, as (pq,r,I) and (pa,r,J) are among the last $k+1$ clockwise (p,r)-crossings to end before this event. Hence, the overall number of such quadruples is at most $O\left(k\ell^2N(n/\ell)+k\ell n^2\beta(n)\right)$.

Finally, if Condition (iii) of Theorem 2.2 holds, we end up with a set B of at most 3ℓ points (possibly including and/or a) so that pr belongs to $\mathrm{DT}(P\setminus B)$ throughout $[t_0,\zeta_{pr}^+]$. For each (p,r)-crossing (pu,r,J_u) that ends in (t_1,t_{rq}) we add the respective point u to B, whose cardinality then increases by at most $k\ll\ell$.

To recap, we may assume the existence of sets A,B, each of size at most $3\ell+k\leq 4\ell$, for which the edge pr belongs to $\mathrm{DT}(P\setminus(A\cup B))$ throughout the interval $\mathcal{I}_{pr}=[\zeta_{pr}^-,\zeta_{pr}^+]$, which covers $[t_{ra},t_{rq}]$. In addition, pr belongs to $\mathrm{DT}(P)$ at times ζ_{pr}^- and ζ_{pr}^+.

Probabilistic refinement. Consider a subset R of $\lceil n/\ell \rceil$ points, chosen at random from P. Fix a quadruple σ in $\mathcal{F}$, and note that, with probability at least $\Omega(1/\ell^4)$, (1) R contains the four points p, q, a, r of σ, and (2) none of the points of $A \cup B \setminus \{q, a\}$ belong to R.

Assuming that the sample R is successful for the chosen σ, the four points p, q, a, r define a Delaunay quadruple, now with respect to R. We continue to denote this new quadruple by σ. As is easy to check, the family $\mathcal{F}_R$ of all such "hereditary" quadruples σ (such that the sample R is successful for their ancestors in $\mathcal{F}$) is 1-refined with respect to the new point set R. Moreover, each quadruple in $\mathcal{F}_R$ satisfies the following new condition:

(Q4) The edge pr belongs to $\mathrm{DT}(R \setminus \{q, a\})$ throughout an interval $\mathcal{I}_{pr} = [\zeta_{pr}^-, \zeta_{pr}^+]$ which covers[20] $[t_{ra}, t_{rq}]$, and it belongs to $\mathrm{DT}(R)$ at times ζ_{pr}^- and ζ_{pr}^+. Moreover, no clockwise (p, r)-crossings (in $\mathcal{F}_R$) begin in (t_{ra}, t_0) or end in (t_3, t_{rq}).

Definition. Let $\mathcal{F}$ be a 1-refined family of Delaunay quadruples. We say that $\mathcal{F}$ is 2-*refined* if its quadruples also satisfy the above condition (Q4) with respect to the underlying point set P (instead of R).

Without loss of generality, we can put ζ_{pr}^- to be the last time in $(-\infty, t_{ra}]$ when pr belongs to $\mathrm{DT}(R)$. Similarly, we can put ζ_{pr}^+ to be the first time in $[t_{rq}, \infty)$ when the edge pr belongs to $\mathrm{DT}(R)$.

Let $\Psi_2(n)$ denote the maximum cardinality of a 2-refined family $\mathcal{F}$, which is defined over a set P of n moving points. The preceding discussion implies the following relation between the quantities $\Psi_1(n)$ and $\Psi_2(n)$:

$$\Psi_1(n) = O\left(\ell^4 \Psi_2(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right). \tag{4}$$

5.5 Stage 3

To bound the above quantity $\Psi_2(n)$, we fix a 2-refined family $\mathcal{F}$ which is defined over a set P of n points moving as above in $\mathbb{R}^2$, and a Delaunay quadruple σ in $\mathcal{F}$.

By assumption, the edges rq and ra of σ are almost Delaunay during the respective intervals $[t_0, t_{rq}]$ and $[t_{ra}, t_3]$. The goal of this stage is to extend the almost-Delaunayhood of these two edges to the interval $[t_{ra}, t_{rq}]$. For the purpose of our analysis, we fix new constants k and ℓ such that $k \ll \ell$.

Figure 25: Left: Extending the almost-Delaunayhood of rq from $[t_0, t_{rq}]$ to $[t_{ra}, t_{rq}]$. Right: Extending the almost-Delaunayhood of ra from $[t_{ra}, t_3]$ to $[t_{ra}, t_{rq}]$.

Charging events in $\mathcal{A}_{rq}$. Refer to Figure 25 (left). If at least k Delaunay counterclockwise (q, r)-crossings (uq, r, I_u) begin in (t_{ra}, t_0), then the crossing (pq, r, I) is again (p, r, k)-chargeable. Indeed, according to Lemma 4.6, each of these crossings occurs within the larger interval $[\zeta_{pr}^-, t_0] \cup I = [\zeta_{pr}^-, t_1]$,

[20] As in the previous step, the times t_{ra} and t_{rq} must be appropriately redefined with respect to the set R at hand, and the interval $[t_{ra}, t_{rq}]$ may shrink. The same applies to the times ζ_{pr}^- and ζ_{pr}^+.

where, by property (Q4), the edge pr is assumed to be almost Delaunay. Moreover, pr belongs to $\mathrm{DT}(R)$ at times ζ_{pr}^- and t_1. Therefore, Theorem 5.3 provides an upper bound of $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$ on the overall number of such crossings (pq, r, I) (and, hence, of their corresponding quadruples σ, as implied by Proposition 5.2). Thus, we can assume, from now on, that the above scenario does not happen for σ. (Notice that the above application of Theorem 5.3 has been prepared by the previous Stage 2, which has extended the almost-Delaunayhood of pr from $[t_0, t_3]$ to $[\zeta_{pr}^-, \zeta_{pr}^+]$.)

We now apply Theorem 2.2 in $\mathcal{A}_{rq}$ over the interval (t_{ra}, t_0), with the threshold ℓ (noting that rq is Delaunay at time t_0, and recalling that Theorem 2.2 also holds if rq is Delaunay at only one endpoint of the interval under consideration). If one of the Conditions (i), (ii) holds, we dispose of σ by charging it (via (pq, r, I)) either to $\Omega(\ell^2)$ ℓ-shallow co-circularities or to an ℓ-shallow collinearity. As in the previous such chargings, each event in $\mathcal{A}_{rq}$ is charged at most $O(k)$ times, as (pq, r, I) is among the $k+1$ first counterclockwise (q, r)-crossings to begin after it. Hence, this charging is applicable for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2\beta(n)\right)$ quadruples σ in $\mathcal{F}$.

Finally, if Condition (iii) of Theorem 2.2 holds, we end up with a set A of at most 3ℓ points such that the edge rq belongs to $\mathrm{DT}(P \setminus A)$ throughout the interval $[t_{ra}, t_1]$.

Charging events in $\mathcal{A}_{ra}$. We now apply a symmetric analysis to the edge ra, spelling it out for the sake of completeness. Refer to Figure 25 (right).

If at least k counterclockwise (a, r)-crossings (ua, r, J_u) end during (t_3, t_{rq}) then the crossing (pa, r, J) is (p, r, k)-chargeable, as each of the respective intervals I_u is the contained in $(t_2, \zeta_{pr}^+]$ (when the edge pr is almost Delaunay). By Theorem 5.3 (and since pr is Delaunay at times t_2 and ζ_{pr}^+), the overall number of such crossings (pa, r, J) (and of their corresponding quadruples σ) is at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$.

Otherwise, we apply Theorem 2.2 in $\mathcal{A}_{ra}$ over the interval (t_3, t_{rq}) (noting that ra is Delaunay at time t_3). If one the Conditions (i), (ii) of that theorem holds, we dispose of σ by charging it (now via (pa, r, J)) either to $\Omega(\ell^2)$ ℓ-shallow co-circularities, or to an ℓ-shallow collinearity. Once again, each event in $\mathcal{A}_{ra}$ is charged at most $O(k)$ times, as (pa, r, J) is among the $k+1$ last counterclockwise (a, r)-crossings to end before it.

Finally, if Condition (iii) of Theorem 2.2 holds, we end up with a set B of at most 3ℓ points such that the edge ra belongs to $\mathrm{DT}(P \setminus B)$ throughout the interval $[t_1, t_{rq}]$.

To recap, we may assume, in what follows, that there exist sets A, B as above, each of cardinality at most 3ℓ.

Probabilistic refinement. We consider a subset R of $\lceil n/\ell \rceil$ points chosen at random from P. We fix a quadruple σ, not disposed of by the previous chargings, and notice that the following two events occur simultaneously, with probability at least $\Omega(1/\ell^4)$: (1) R contains the four points p, q, a, r of σ, and (2) none of the points of $A \cup B \setminus \{q, a, r\}$ belong to R.

Let $\mathcal{F}_R$ denote the family of all hereditary quadruples σ (such that R is successful for their ancestors in $\mathcal{F}$). As is easy to check, $\mathcal{F}_R$ is 2-refined (in R). Moreover, the following new conditions hold for every quadruple σ in $\mathcal{F}$:

(Q5) The edge ra belongs to $\mathrm{DT}(R \setminus \{p, q\})$ throughout the interval $[t_{ra}, t_{rq}]$.

(Q6) The edge rq belongs to $\mathrm{DT}(R \setminus \{p, a\})$ throughout the interval $[t_{ra}, t_{rq}]$.

We say that a family $\mathcal{F}$ of Delaunay quadruples is 3-*refined* if (1) it is 2-refined, and (2) its quadruples satisfy Conditions (Q5) and (Q6) with respect to the underlying point set. Let $\Psi_3(n)$ denote the maximum cardinality of a 3-refined family of Delaunay quadruples that is defined over a set of n moving points (that we keep denoting as P, replacing R in these conditions). The preceding discussion implies the following relation between the quantities $\Psi_2(n)$ and $\Psi_3(n)$:

$$\Psi_2(n) = O\left(\ell^4\Psi_3(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2\beta(n)\right). \quad (5)$$

5.6 Stage 4

To bound the above quantity $\Psi_3(n)$, we fix a 3-refined family $\mathcal{F}$ which is defined over an underlying set P of n moving points. (That is, $\mathcal{F}$ satisfies all the six condtions (Q1)–(Q6).) Proposition 5.3 implies that every quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$ is uniquely determined by the ordered triple (q, a, r).

For the purpose of our analysis, we also fix three new constants k, ℓ, h such that $12 < k \ll \ell \ll h$.

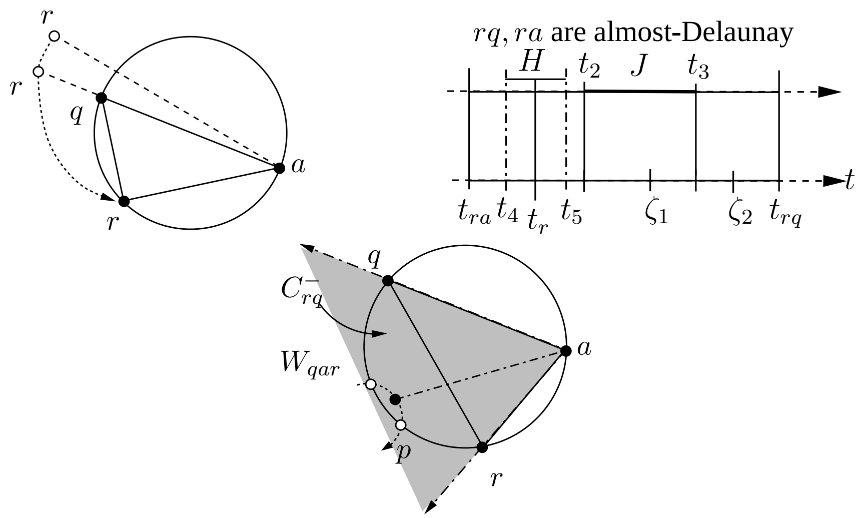

Figure 26: The topological setup during the interval $[t_{ra}, t_{rq}]$. Left: The edge ar is hit at some time $t_r \in [t_{ra}, t_2]$ by q. Center: we have $t_{ra} \le t_4 \le t_r \le t_5 < t_2 < \zeta_1 < t_3 < \zeta_2$. Right: The motion of $B[q, a, r]$ is continuous throughout $(t_r, \zeta_2]$ (the hollow circles represent the co-circularities at times ζ_1 and ζ_2).

Topological setup. We fix a quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$, whose two Delaunay crossings take place during the intervals $I = [t_0, t_1]$ and $J = [t_2, t_3]$ (in this order). Refer to Figure 26.

Since $\mathcal{F}$ is 1-refined, there exists a time $t_{ra} \le t_0$ which is the last time before[21] t_0 when the edge ra belongs to $\mathrm{DT}(P)$, and a symmetric first time $t_{rq} \ge t_3$ when rq belongs to $\mathrm{DT}(P)$. Moreover, by Conditions (Q5) and (Q6), the edge ra belongs to $\mathrm{DT}(P \setminus \{p, q\})$, and the edge rq belongs to $\mathrm{DT}(P \setminus \{p, a\})$, throughout the interval $[t_{ra}, t_{rq}]$.

Let us summarize what we know so far about σ. By Condition (Q3), the points p, q, a, r of σ are co-circular at times $\zeta_0 \in I \setminus J$, $\zeta_1 \in J \setminus I$, and $\zeta_2 \in (t_3, t_{rq}]$. Moreover, the Delaunayhood of pa is violated, throughout (ζ_1, ζ_2), by the points $q \in L^-_{pa}$ and $r \in L^+_{pa}$. In particular, p lies throughout that interval within the wedge $W_{qar} = L^+_{qa} \cap L^-_{ra}$ and inside the cap $C^-_{rq} = B[q, a, r] \cap L^-_{rq}$; see Figure 26 (right). We emphasize that the order type of the quadruple (p, q, a, r) remains unchanged during (ζ_1, ζ_2), and is exactly as depicted in this figure.

In addition, by the same Condition (Q3), the smaller set $P \setminus \{p\}$ induces a (single) Delaunay crossing (ar, q, H_σ), whose interval $H = H_\sigma = [t_4, t_5]$ is contained in $[t_{ra}, t_2)$; see Figure 26 (left and center). In particular, q hits ar at some moment[22] $t_r \in H$, and crosses L_{ar} from L^-_{ar} to L^+_{ar}. Since q lies in L^+_{ar} at times $\zeta_1 > t_2$ and ζ_2, no further collinearities of q, a, r can occur during $(t_r, \zeta_2]$. (Otherwise, the point q would have to re-enter L^+_{ar}, after previously crossing L_{ar} back to L^-_{ar}, and then the triple q, a, r would

[21]If ra is Delaunay at time t_0 then we put $t_0 = t_{ra}$.

[22]Recall from Section 5.3 that q can cross ar either before or after ζ_0, depending on the location of a when r crosses pq. Our analysis only relies on the fact that $t_r < \zeta_1 < \zeta_2$, which follows because $\zeta_r < t_2$ and $\zeta_1 \ge t_2$.

be collinear three times, contrary to our assumptions.) To recap, the disc $B[q, a, r]$ moves continuously throughout the interval $(t_r, \zeta_2]$, which is obviously contained in $[t_{ra}, t_{rq}]$.

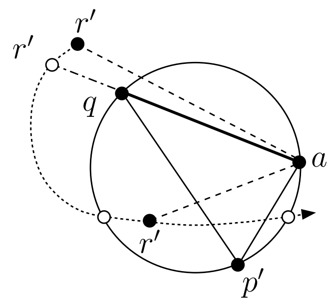

Figure 27: A quadruple $\sigma' = (p', q, a, r')$ in $\mathcal{F}_{qa}$. The edge ar' undergoes an (a, q)-crossing $(ar', q, H_{\sigma'})$ within the triangulation $\mathrm{DT}(P \setminus \{p'\})$.

Let $\mathcal{F}_{qa}$ denote the subfamily of all quadruples $\sigma' = (p', q, a, r')$ in $\mathcal{F}$, whose middle points q and a are fixed and equal to those of σ. (In particular, $\mathcal{F}_{qa}$ contains σ.) For each $\sigma' = (p', q, a, r')$ in $\mathcal{F}_{qa}$, the appropriately pruned set $P \setminus \{p'\}$ induces the (a, q)-crossing $(ar', q, H_{\sigma'})$; see Figure 27. In what follows, we keep σ and $\mathcal{F}_{qa}$ fixed and distinguish between several cases.

Case (a). The family $\mathcal{F}_{qa}$ contains at least k quadruples σ' whose respective crossings $(ar', q, H_{\sigma'})$ end during (t_5, t_{rq}). Refer to Figure 28. Recall that, according to Proposition 5.3, the point p' is uniquely determined by the choice r'.

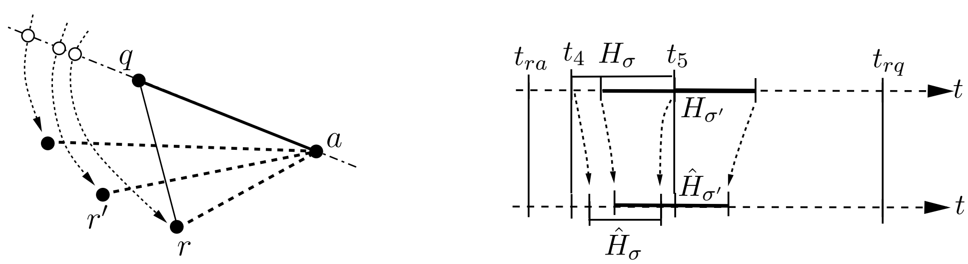

Figure 28: Case (a). Left: At least k of the crossings $(ar', q, H_{\sigma'})$ end during (t_5, t_{rq}). Right: A successful sample $\hat{P}$ yields Delaunay crossings $(ar, q, \hat{H}_\sigma)$ and $(ar', q, \hat{H}_{\sigma'})$, which occur within $[t_4, t_{rq}]$.

Informally, we would like to dispose of σ using Theorem 5.3, by showing that the counterclockwise (r, q)-crossing (ar, q, H) is $(r, q, \Theta(k))$-chargeable (for the interval $\mathcal{I} = [t_4, t_{rq}]$). Unfortunately, the (a, q)-crossings $(ar', q, H_{\sigma'})$ to be charged are defined with respect to (potentially) distinct sets $P \setminus \{p'\}$, and thus do not fit the definition of chargeability.

To free sufficiently many crossings $(ar', q, H_{\sigma'})$ from their violating points p', we pass from P to a sample $\hat{P}$ of $\lceil n/2 \rceil$ points chosen at random from P. Notice though that $\mathcal{F}_{qa}$ can potentially include quadruples $\sigma' = (p', q, a, r')$ with $p' = r$, which cannot be freed without destroying rq and (ar, q, H).

Fortunately, by Proposition 5.3, for any quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}_{qa}$ there is at most one other quadruple $\sigma = (p', q, a, r')$, also in $\mathcal{F}_{qa}$, with $r' = p$. The pigeonhole principle then implies that at least half of the quadruples $\sigma = (p, q, a, r)$ in $\mathcal{F}_{qa}$ satisfy the following converse condition:

(PH) *There is at most one quadruple $\sigma' = (p', q, a, r')$ in $\mathcal{F}_{qa}$ with $p' = r$.*

In more detail, consider the (possibly partial) map $\lambda : \mathcal{F}_{qa} \to \mathcal{F}_{qa}$, so that λ maps each quadruple $\sigma = (p, q, a, r) \in \mathcal{F}_{qa}$ to the unique quadruple $\lambda(\sigma) = (\omega, q, a, p) \in \mathcal{F}_{qa}$ if it exists, and otherwise λ is undefined at σ. Put $\mu_\sigma = |\{\sigma' \mid \lambda(\sigma') = \sigma\}|$, for each $\sigma \in \mathcal{F}_{qa}$. Then $\sum_{\sigma \in \mathcal{F}_{qa}} \mu_\sigma \le M = |\mathcal{F}_{qa}|$, so the number of quadruples σ with $\mu_\sigma \ge 2$ is at most $M/2$. All the remaining quadruples satisfy (PH).

Since q and a are arbitrary points of P, (PH) holds for at least half of all quadruples in $\mathcal{F}$; hence we may assume that it holds for the quadruple σ under consideration.

Let σ' be a quadruple in $\mathcal{F}_{qa} \setminus \{\sigma\}$ whose crossing $(ar', q, H' = H_{\sigma'})$ ends in (t_5, t_{rq}). We further assume that $p' \neq r$ and $r' \neq p$. Then we have the following relaxed version of Lemma 4.6, which can be established by observing that its original proof holds also in the new setup. (An alternative proof of Lemma 5.5 can be obtained through examining the two co-circularities that are performed by a, q, r, r', according to Lemmas 4.1 and 4.4, during the intervals $H \setminus H'$ and $H' \setminus H$, and then applying Lemma 4.6 for the reduced set $P \setminus \{p, p'\}$.)

Lemma 5.5. *Let P be a set of points moving as above in $\mathbb{R}^2$, and let (ar, q, H) and (ar', q, H') be a pair of clockwise (a, q)-crossings that occur in the respective reduced triangulations $\mathrm{DT}(P \setminus \{p\})$ and $\mathrm{DT}(P \setminus \{p'\})$, for $p, p' \in P$.*[23] *Furthermore, assume that $r \neq p'$ and $r' \neq p$. Then the statement of Lemma 4.6 holds for (ar, q, H) and (ar', q, H'). That is, q hits ar (during H) before it hits ar' (during H') if and only if H begins (resp., ends) before the beginning (resp., end) of H'.*

Clearly, the above restriction on p' and r' is now satisfied by at least $k - 2 \geq k/2$ of the quadruples $\sigma' = (p', q, a, r')$ that are assumed to exist in the current case (a). Since their intervals H' end in (t_5, t_{rq}), Lemma 5.5 implies that, for each of them, H' starts after t_4, and the point q hits ar' (during H') after time t_r.

We now return to the sample $\hat{P}$ and observe that the following two events occur simultaneously, with at least some fixed constant probability:

(1) The sample $\hat{P}$ includes the three points q, a, r, but not p. Hence, $\hat{P}$ induces a single Delaunay crossing $(ar, q, \hat{H} = \hat{H}_\sigma)$ of ar by q.

(2) The sample $\hat{P}$ includes the point r', but not p', for at least $k/16$ of the above quadruples $\sigma' = (p', q, a, r')$. For each of these $k/16$ quadruples, the sample $\hat{P}$ yields a Delaunay (a, q)-crossing $(ar', q, \hat{H}_{\sigma'})$ with $\hat{H}_{\sigma'} \subseteq H_{\sigma'}$.

(To see (2), note that this property holds for any single quadruple with probability at least $1/4$, so the expected number of successful quadruples is at least $k/8$. By a variant of Markov's bound, the probability of having at least $k/16$ successful quadruples is at least $14/15$.)

Suppose that the sample $\hat{P}$ is indeed successful for σ. Recall that, for each quadruple σ' in (2), q hits the respective edge ar' (during $H_{\sigma'}$) after it hits ar (during H_σ).

We now pass to the sampled triangulation $\mathrm{DT}(\hat{P})$. Lemma 4.6 implies, in combination with the containment $\hat{H}_{\sigma'} \subseteq H_{\sigma'}$, that all the Delaunay crossings $(ar', q, \hat{H}_{\sigma'})$ in (2) end after $\hat{H}$ and before t_{rq}; see Figure 28 (right). Therefore, all of them must occur within the interval $H_\sigma \cup [t_5, t_{rq}] \subseteq [t_4, t_{rq}]$, where the edge rq is assumed to be almost Delaunay.[24] In addition, the edge rq belongs to $\mathrm{DT}(\hat{P})$ at *both* times t_4 and t_{rq}, because $\hat{P}$ does not include p. Since σ' and $(ar', q, \hat{H}_{\sigma'})$ can be chosen in at least $k/16$ distinct ways, the crossing $(ar, q, \hat{H})$ is $(r, q, k/16)$-chargeable (with respect to $\hat{P}$).

By Theorem 5.3, the overall number of such triples (q, a, r) in $\hat{P}$ is $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$. Clearly, the same bound must hold for the overall number of quadruples σ that fall into case (a).

Preparing for cases (b), (c): Charging events in $\mathcal{A}_{qa}$. We can assume, from now on, that the family $\mathcal{F}_{qa}$ contains at most k quadruples σ' whose "almost Delaunay" crossings $(ar', q, H_{\sigma'})$ end during (t_5, t_{rq}).

Before proceeding to the subsequent cases, we apply Theorem 2.2 in $\mathcal{A}_{qa}$ over the interval (t_5, t_{rq}), now with the second constant ℓ. Notice that the edge qa belongs to $\mathrm{DT}(P \setminus \{p\})$ at time t_5, so we omit p and apply the theorem with respect to that smaller triangulation.

If at least one of the Conditions (i), (ii) of Theorem 2.2 is satisfied, we charge σ either to an $(\ell+1)$-shallow collinearity, or to $\Omega(\ell^2)$ $(\ell+1)$-shallow co-circularities. (Each of these events is ℓ-shallow with

[23] We do not require that p and p' be distinct.

[24] Notice that the times t_{rq}, t_4 and t_5 are defined with respect to the original point set P.

respect to $P \setminus \{p\}$, and its depth can go up by 1 when p is added back.) It remains to check that each $(\ell+1)$-shallow event, which occurs in $\mathcal{A}_{qa}$ at some time $t^* \in (t_5, t_{rq})$, is charged by at most $O(k)$ quadruples σ. Indeed, q and a are among the three or four points involved in the event. We guess q and a (in $O(1)$ possible ways) and consider all "almost Delaunay" crossings of the form $(ar', q, H_{\sigma'})$, each of them associated with some (unique) "candidate" quadruple $\sigma' = (p', q, a, r')$ in $\mathcal{F}_{qa}$. Since case (a) is ruled out (and since t^* belongs to (t_5, t_{rq})), $(ar, q, H = H_\sigma)$ is among the k last such "almost Delaunay" crossings to end before time t^*. Since p is uniquely determined by the choice of q, a and r, we can guess σ in $O(k)$ possible ways. Hence, this scenario happens for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2\beta(n)\right)$ quadruples.

Now assume that Condition (iii) of Theorem 2.2 holds. Then P contains a subset A of at most 3ℓ points such that the edge qa belongs to $\mathrm{DT}(P \setminus (A \cup \{p\}))$ throughout the interval $H \cup [t_5, t_{rq}] = [t_4, t_{rq}]$. In particular, the following property must hold:

At most 3ℓ points $s \in P \setminus \{p\}$ hit qa during the interval (t_r, ζ_2) $(\subseteq (t_{rq}, t_{rq}))$.

Case (b). There exist at least ℓ points, distinct from p, that enter the cap $C^-_{rq} = B[q,a,r] \cap L^-_{rq}$ during (t_r, ζ_2). We refer to Figure 29 and let s be any of these points. By Condition (Q6), s cannot hit rq during the interval (t_r, ζ_2) (which is covered by $[t_{ra}, t_{rq}]$). Note also that C^-_{rq} is contained in the wedge $W_{qar} = L^+_{qa} \cap L^-_{ra}$. Therefore, and since the wedge W_{qar} is empty immediately after time t_r (when q, a and r are collinear), the above point s has to enter W_{qar}, through one its rays $\vec{ar}, \vec{aq}$, during (t_r, ζ_2) and before it enters C^-_{rq}.

Furthermore, Condition (Q6) implies that s can enter the cap C^-_{rq} only through the boundary of $B[q,a,r]$, which results in a co-circularity of q, a, r, s. (Recall also that s enters each halfplane L^+_{qa} and L^-_{ra} at most once, so it crosses the ray $\vec{ar}$ or $\vec{aq}$ *outside* the respective edge ar or aq when entering W_{qar} as above. Indeed, otherwise s would be able to access C^-_{rq}, after crossing one of these two edges, only through the interior of rq.)

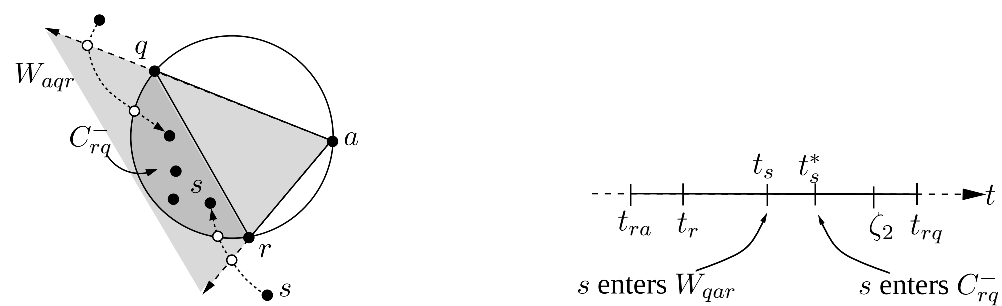

Figure 29: Case (b). At least ℓ points $s \neq p$ enter the cap C^-_{rq} during (t_r, ζ_2) (p is not shown). Each of the first ℓ of these points causes an $(\ell+1)$-shallow co-circularity with q, a, r. Each of them must first enter the wedge W_{qar}, which is empty at time t_r, through one of the rays $\vec{aq}, \vec{ar}$ (outside the edges aq and ar), because none of them can cross rq.

Assume that s is among the first ℓ points to enter C^-_{rq} during (t_r, ζ_2). Let t^*_s denote the time of the corresponding co-circularity of q, a, r, s, which occurs when s enters C^-_{rq}. Since σ satisfies Condition (Q6) (and t^*_s belongs to (t_{ra}, t_{rq})), the opposite cap $C^+_{rq} = B[q,a,r] \cap L^+_{rq}$ contains no points of $P \setminus \{p\}$ at time t^*_s. (Otherwise, the Delaunayhood of rq would then be violated by s and another point of $P \setminus \{p\}$, contrary to (Q6).) Therefore, and since the motion of $B[q,a,r]$ is continuous during (t_r, ζ_2), the co-circularity at time t^*_s has to be $(\ell-1)$-shallow in $P \setminus \{p\}$, and thus ℓ-shallow in P.

Note also that the crossing (ar, q, H) has to end before t^*_s (that is, $t_5 < t^*_s$). Indeed, the Delaunayhood of qr is violated, right after time t^*_s, by s and a, which is forbidden by Lemma 4.1 during H.

We distinguish between two possible subcases. In each of them we dispose of σ by charging it, within one of the arrangements $\mathcal{A}_{ra}, \mathcal{A}_{qa}$, either to $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities, or to a (2ℓ)-

shallow collinearity.

Case (b1). At least half of the above points s cross the line L_{ra}, from L_{ra}^+ to L_{ra}^-, during (t_r, t_s^*). Since s lies in L_{ra}^- at time t_s^*, s enters L_{ra}^- exactly once during (t_r, t_s^*), and it does not return to L_{ra}^+ before t_s^*; see the motion of the marked point s in Figure 29 (left). Moreover, by Condition (Q5), each of these crossings occurs outside ra (i.e., within one of the outer rays of L_{ra}).

To dispose of σ, we again fix one of the aforementioned points s and argue as in Section 3. If the halfplane L_{ra}^- contains at most 2ℓ points of P when s enters it, then we encounter a (2ℓ)-shallow collinearity of a, r, s. Otherwise, the disc $B[a, r, s]$ contains at least 2ℓ points right after the crossing, so the three points r, a, s are involved in at least ℓ (2ℓ)-shallow co-circularities before time t_s^* (when the open disc $B[a, r, s]$, equal to $B[a, r, q]$ at that time, contains ℓ or fewer points of P). After repeating the above argument for each of the (at least) $\ell/2$ possible choices of s, we encounter in $\mathcal{A}_{ra}$ (during (t_r, ζ_2)) either $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities, or a (2ℓ)-shallow collinearity. In both cases, we charge σ to these events.

We claim that each (2ℓ)-shallow event, which occurs in $\mathcal{A}_{ra}$ at some time $t^* \in (t_r, \zeta_2)$, is charged by at most $O(1)$ quadruples σ. Indeed, r and a are among the three or four points involved in every charged event. Moreover, according to Condition (Q5) and the argument in case (e) of Stage 1, q is among the last two points to hit the edge ra before time t^*. Hence, knowing t^* allows us to guess the three points q, a, r (which uniquely determine σ) in at most $O(1)$ ways. In conclusion, the above scenario happens for at most $O\left(\ell^2 N(n/\ell) + \ell n^2 \beta(n)\right)$ quadruples σ in $\mathcal{F}$.

Case (b2). At least half of the above points s remain in L_{ra}^- throughout the respective intervals (t_r, t_s^*). Each of these points must enter W_{qar} (during (t_r, t_s^*)) through the ray emanating from q in direction $\vec{aq}$, thereby crossing L_{qa} from L_{qa}^- to L_{qa}^+. (Recall that such a collinearity of q, a, s can occur only once during (t_r, t_s^*).)

Once again, we fix one of the above points s and let t_s denote the time in (t_r, ζ_2) when s enters W_{qar} through the ray emanating from q in direction $\vec{aq}$. Arguing as in the previous case, we conclude that the three points q, a, s are involved (during $(t_s, t_s^*) \subset (t_r, \zeta_2)$) either in a (2ℓ)-shallow collinearity, or in $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities. Below we prove that each of the (2ℓ)-shallow events, that occur in $\mathcal{A}_{qa}$ during (t_r, ζ_2), can be traced back to σ in at most $O(k)$ ways.[25] Hence, it is charged at most $O(k)$ times. We then repeat the same argument for each of the remaining $\ell/2 - 1$ choices of s, and use (as in case (b1)) the standard bounds on the number of (2ℓ)-shallow events of each type. As a result, we obtain an upper bound of $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ of the number of such quadruples σ.

To conclude, the overall number of quadruples σ that fall into Case (b) is at most

$$O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right).$$

To complete the analysis of Case (b), we show that each (2ℓ)-shallow event that occurs in $\mathcal{A}_{aq}$ during (t_r, ζ_2) is charged as above by at most $O(k)$ quadruples σ that fall into case (b2). Let t^* be the time of such an event. First, we guess the points q, a, in $O(1)$ possible ways, from among the three or four points involved in the event. Recall that, in the charging scheme of case (b2), each (2ℓ)-shallow co-circularity or collinearity that we charge in $\mathcal{A}_{qa}$ is "obtained" via some point s, which is also involved in this event and enters L_{qa}^+ at some prior time t_s. We, therefore, guess s among the remaining one or two points that participate in the event under consideration. To guess the remaining points r and p of σ, we examine all "candidate" quadruples $\sigma' \in \mathcal{F}_{qa}$ whose two "middle" points are shared with σ. Recall that each of these quadruples $\sigma' = (p', q, a, r')$ is accompanied by an "almost Delaunay" crossing $(ar', q, H_{\sigma'})$, where r' enters L_{qa}^+ at some time $t_{r'} \in H_{\sigma'}$. Also recall that σ' is uniquely determined by the choice of r' (as long as q and a remain fixed).

It suffices to consider only quadruples $\sigma' = (p', q, a, r')$, in $\mathcal{F}_{qa}$, with the following properties: (1) $s \neq p', r'$, (2) $t_{r'} < t_s$, and (3) s lies in $L_{ar'}^+$ during the second portion of $H_{\sigma'}$ (after $t_{r'}$). This is because

[25] Note the difference between the two subcases: Here we only know q, a, and then guessing r is not immediate.

each of these conditions holds for $\sigma' = \sigma$ (and for s) in the charging scheme of case (b2). For example, (3) follows because we assume that case (b1) does not occur (and since $t_5 < t_s^*$). The corresponding points r', which determine the above quadruples σ', are called *candidates* (for r).

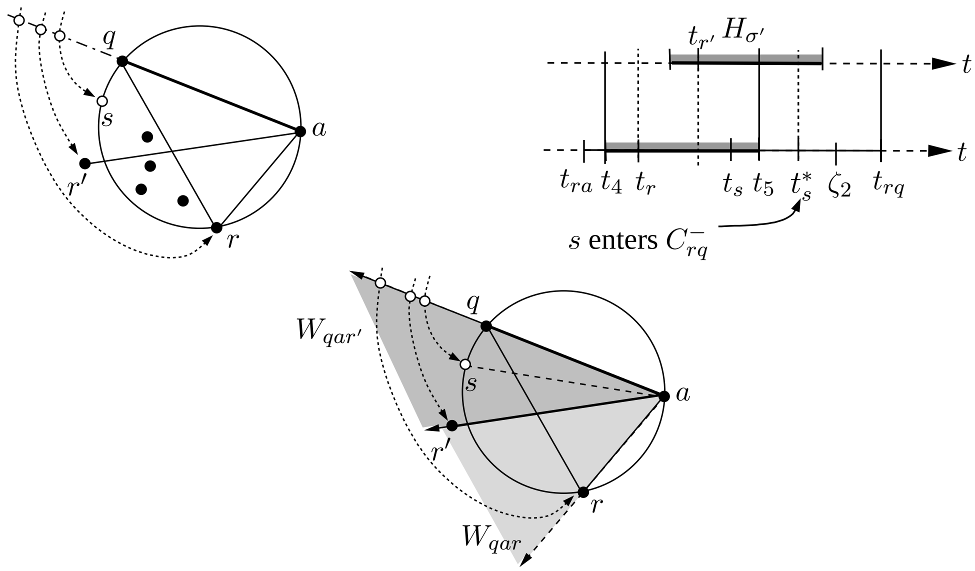

Figure 30: Top: Proposition 5.6: r is among the last $k+3$ candidates r' to enter L_{qa}^+ before time t_s; the various critical events occur in the depicted order. Bottom: Proof of Proposition 5.6. The candidate r' remains in W_{qar} throughout $(t_{r'}, t_s^*)$. If $H_{\sigma'}$ ends after t_s^*, then the point s remains in $W_{qar'} (\subset W_{qar})$ throughout (t_s, t_s^*).

Proposition 5.6. *With the above assumptions, the point r is among the last $k+3$ candidates r' to enter the halfplane L_{qa}^+ before t_s (each candidate at the respective time $t_{r'}$).*

Proof. Refer to Figure 30. Assume to the contrary that the proposition does not hold (for σ and $s \neq p, q, a, r$ as above). Hence, we have at least k candidates r' such that $t_r < t_{r'} < t_s$ and $r' \notin \{p, r\}$, and such that the points p' of their respective quadruples $\sigma' = (p', q, a, r')$ are distinct from r. (We continue to assume that σ satisfies property (PH), introduced in case (a), so the last two restrictions on p' and r' exclude from our consideration at most three candidates r', with their quadruples σ'.)

To establish the proposition, we fix a candidate r' and its corresponding quadruple $\sigma' = (p', q, a, r')$, as above, and argue that the respective interval $H_{\sigma'}$ ends during (t_5, t_{rq}). Repeating the same argument for the remaining $k-1$ possible choices of r' will imply that the quadruple σ falls into case (a) and thereby reach a contradiction.

Indeed, since $t_r < t_{r'}$, Lemma 5.5 shows that the interval $H_{\sigma'}$ ends after $H = [t_4, t_5]$. (As in case (a), the lemma relies on the assumption that $p \neq r'$ and $r \neq p'$.) It remains to check that $H_{\sigma'}$ ends before t_{rq}.

If $H_{\sigma'}$ ends before t_s^*, then we are done (as $t_s^* < t_{rq}$). Hence, we may also assume that both times $t_{r'}$ and t_s^* belong to the interval $H_{\sigma'}$ (as depicted in Figure 30 (top-right)). This, and the choice of r' as a candidate for r, implies that r' remains in the halfplanes L_{qa}^+, L_{sa}^+ throughout the interval $(t_{r'}, t_s^*) \subseteq H_{\sigma'}$. Indeed, r' cannot re-enter L_{qa}^- during the second portion of $H_{\sigma'}$, after entering L_{qa}^+ at time $t_{r'} \in H_{\sigma'}$. (This is because q, a, r' perform only one collinearity during the crossing $(ar', q, H_{\sigma'})$.) Similarly, since σ' satisfies property (3), the point s remains in $L_{ar'}^+$ throughout $(t_{r'}, t_s^*)$ (so r' remains in L_{sa}^+). We thus conclude that s lies inside $W_{qar'} = L_{qa}^+ \cap L_{r'a}^-$ throughout the interval (t_s, t_s^*); see Figure 31 (bottom).

Also notice that, with the above assumptions, r' must lie, throughout the longer interval $(t_{r'}, t_s^*) \subseteq H_{\sigma'}$, inside the wedge $W_{qar} = L_{qa}^+ \cap L_{ra}^-$. Indeed, r' enters W_{qar} at time $t_{r'} \in (t_r, t_s^*) (\subseteq (t_r, \zeta_2) \cap H_{\sigma'})$ and cannot again cross the ray $\vec{aq}$ during $H_{\sigma'}$. Moreover, if r' leaves W_{qar} (during $(t_{r'}, t_s^*)$) through the other ray $\vec{ar}$, then the edge ar' is hit by r, or the edge ar is hit by r'. Clearly, the former crossing is forbidden by Lemma 4.1 during the interval $H_{\sigma'}$ (where ar' experiences a Delaunay crossing by q), and the latter one is ruled out by Condition (Q5). (As a matter of fact, in the second case r' must also cross rq, thereby entering $\triangle qar$, before it reaches ra. This collinearity is also impossible by Condition (Q6).)

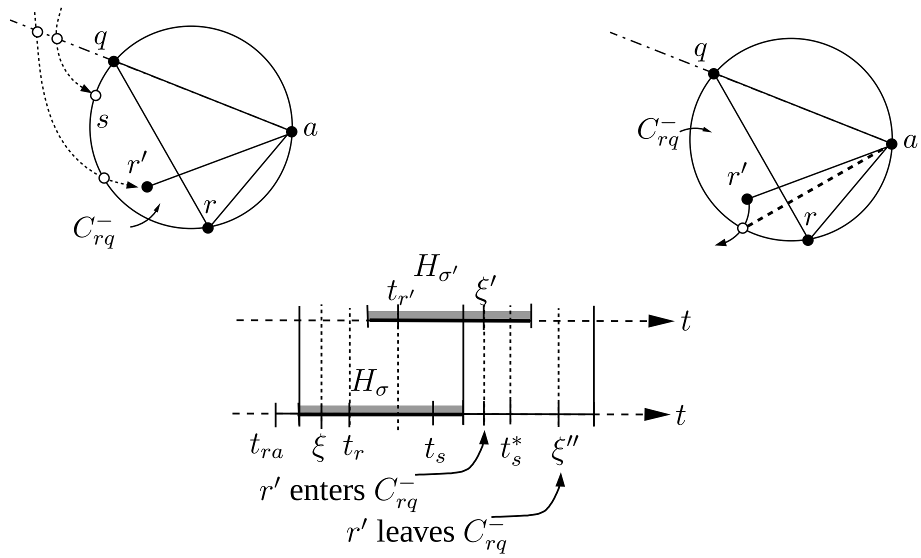

Figure 31: Proof of Proposition 5.6: The scenario where r' lies within $B[q,a,r]$ at time t^*_s. Left: r' enters C^-_{rq} during (t_r, t^*_s) through the arc $C[q,a,r] \cap L^-_{rq}$, at some time ξ' (left). Right: r' must leave C^-_{rq} before t_{rq} (and after $H_{\sigma'}$). Below: The points q,a,r,r' are co-circular at times $\xi \in H_\sigma \setminus H_{\sigma'}, \xi' \in H_{\sigma'} \setminus H_\sigma$ and $\xi'' \in (\xi', t_{rq}]$. The interval $H_{\sigma'}$ ends before ξ'' (and, thus, before t_{rq}).

To recap, we can assume that $H_{\sigma'}$ ends *after* t^*_s, and that the edges aq, as, ar' and ar appear, at time t^*_s, in counterclockwise order around a. To show that $H_{\sigma'}$ ends before t_{rq}, we distinguish between two possible cases.

(1) If r' lies outside $B[q,a,s] = B[q,a,r]$ at time t^*_s, then the Delaunayhood of the edge ar' is violated, at that very moment, by the points $s \in L^+_{ar'}$ and $r \in L^-_{ar'}$ (as depicted in Figure 30 (left)). Since $p' \notin \{s,r\}$, the crossing $(ar', q, H_{\sigma'})$ (occurring in $\mathrm{DT}(P \setminus \{p'\})$) has to end before t^*_s, which is contrary to our assumptions.

(2) Now suppose that r' lies at time t^*_s within $B[q,a,r]$, as depicted in Figure 31 (left). Since r' remains in W_{qar} throughout $(t_{r'}, t^*_s]$ (and since r' lies outside $B[q,a,r]$ at time $t_{r'}$, when it enters W_{qar}), it can enter $B[q,a,r]$ (or, more precisely, its cap C^-_{rq}) during $(t_{r'}, t^*_s)$ only through the circular arc $C[q,a,r] \cap L^-_{rq}$. When that happens, we encounter a co-circularity of q,a,r,r' at some time $\xi' \in (t_{r'}, t^*_s] \subseteq H_{\sigma'}$, right after which the Delaunayhood of rq is violated by $r' \in L^-_{rq}$ and $a \in L^+_{rq}$. Since $p \neq r'$ and $r \neq p'$, this co-circularity occurs after $H = H_\sigma$.

Applying Lemma 4.4 to (ar, q, H) shows that another co-circularity of q,a,r,r' (red-blue with respect to ar and thus monochromatic with respect to ar') must occur at some time $\xi < \xi'$ during the symmetric interval $H_\sigma \setminus H_{\sigma'}$. As is easy to check[26], ξ and ξ' are the only co-circularities of q,a,r,r' to occur during H_σ and $H_{\sigma'}$.

To complete our analysis, we apply Lemma 3.1 for the edge rq, with the reference interval $(\xi', t_{rq}]$. By Conditions (Q3) and (Q6), neither of a, r' can cross rq during the larger interval $[t_r, t_{rq}]$. Therefore, we encounter a third co-circularity of q,a,r,r' at some time ξ'' in $(\xi', t_{rq}]$, which occurs when r' leaves the cap C^-_{rq}. See Figure 31 (right). Since ξ and ξ' are the only co-circularities to occur during $H_\sigma \cup H_{\sigma'}$,

[26]Note, for instance, that (a, r, r', q) is a counterclockwise quadruple in $\mathrm{DT}(P \setminus \{p, p'\})$, so the argument preceding Lemma 5.1 applies to it.

the third co-circularity ξ'' must occur *after* $H_{\sigma'}$. (See Figure 31 (bottom).) Hence, $H_{\sigma'}$ has to end before t_{rq} also in this last case. $\square$

Case (c). Assume that none of the previous cases or preliminary chargings applies to σ. In particular, since the charging within $\mathcal{A}_{qa}$ following case (a) does not apply, at most 3ℓ points of $P \setminus \{p\}$ cross qa during (t_r, ζ_2). Furthermore, since case (b) does not occur, at most ℓ points of $P \setminus \{p\}$ enter the cap $C^-_{rq} = B[q,a,r] \cap L^-_{rq}$, during the interval (t_r, ζ_2). See Figure 32 (left).

We again emphasize that, by condition (Q5), no point in $P \setminus \{p, q\}$ can hit the edge ra during the interval $[t_{ra}, t_{rq}]$ (which contains $[t_r, \zeta_2]$). Similarly, condition (Q6) implies that no point in $P \setminus \{p, a\}$ can hit the edge rq during that interval.

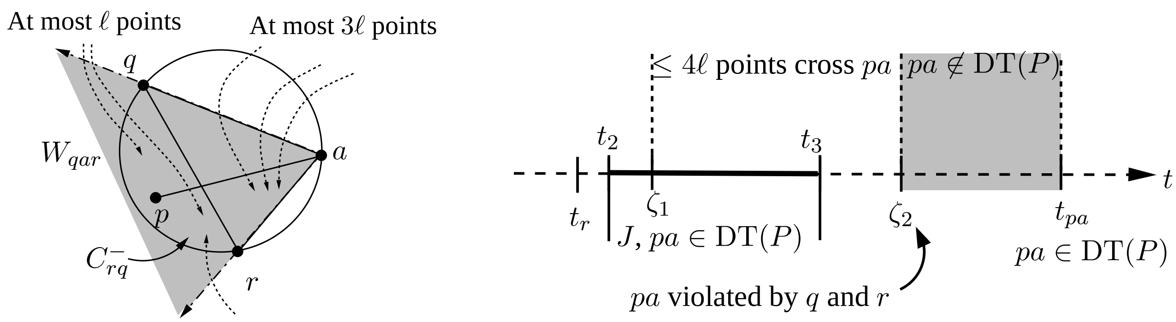

Figure 32: Left: Case (c). At most ℓ points of $P \setminus \{p\}$ enter C^-_{rq}, and at most 3ℓ points of $P \setminus \{p\}$ cross qa, during (t_r, ζ_2). Hence, at most 4ℓ points cross pa during (ζ_1, ζ_2). Right: a schematic summary of our setup in case (c).

We claim that at most 4ℓ points of $P \setminus \{p, a\}$ can hit the edge pa during the interval (t_3, ζ_2) ($\subseteq (\zeta_1, \zeta_2)$). Indeed, fix any of these points s. Recall the edge pa is contained during the interval (ζ_1, ζ_2) in the region $B[q,a,r] \cap W_{qar}$; see Figures 26 (right) and 32 (left). Hence, s has to lie in $B[q,a,r] \cap W_{qar}$ when it hits pa, as well. Since W_{qar} contains no points of P at time t_r, the point s has to enter this wedge during (t_r, ζ_2) through one of the rays $\vec{ar}, \vec{ap}$. If s crosses pa within L^-_{rq} then, in particular, it has to enter the cap C^-_{rq} during (t_r, ζ_2). Otherwise, if s hits pa within L^+_{rq}, then it must have previously entered the triangle $\triangle qar$ through the edge qa. (By Conditions (Q5) and (Q6), s cannot crosses either of the edges ra, rq during (t_r, ζ_2).) We thus conclude that the overall number of points in P that cross pa during (t_3, ζ_2) cannot exceed $\ell + 3\ell = 4\ell$.

Charging events in $\mathcal{A}_{pa}$. The above analysis implies, in particular, that the edge pa undergoes at most 4ℓ Delaunay crossings within (t_3, ζ_2). If the edge pa never re-enters $\mathrm{DT}(P)$ after time ζ_2, then (pa, r, J) is among the last $4\ell + 1$ Delaunay crossings of pa. Clearly, this scenario happens for at most $O(\ell n^2)$ quadruples σ.

Otherwise, let t_{pa} be the first time after ζ_2 when pa re-enters $\mathrm{DT}(P)$. Refer to the schematic Figure 32 (right). Since the co-circularity at time ζ_2 is the *last* co-circularity of the points of σ, Lemma 3.1 implies that the edge pa is hit during $(\zeta_2, t_{pa}] \subseteq (t_3, t_{pa}]$ by at least one of the remaining two points q and r.

We apply Theorem 2.2 in $\mathcal{A}_{pa}$ over the interval (t_3, t_{pa}), with the third constant parameter h (noting that pa is Delaunay at both endpoints of that interval). If one of the Conditions (i), (ii) holds, we charge σ (via (pa, r, J)) either to an h-shallow collinearity, or to $\Omega(h^2)$ h-shallow co-circularities (where each charged event occurs during (t_3, t_{pa}) and involves p and a, together with one or two additional points of P). Any such h-shallow event is charged by at most $O(\ell)$ quadruples. Indeed, the two points p, a can be guessed in at most $O(1)$ possible ways out of the three or four points involved in it, and (pa, r, J) is among the last $4\ell + 1$ Delaunay crossings of pa to end before the respective time of the event. Therefore, the above charging accounts for at most $O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ quadruples σ.

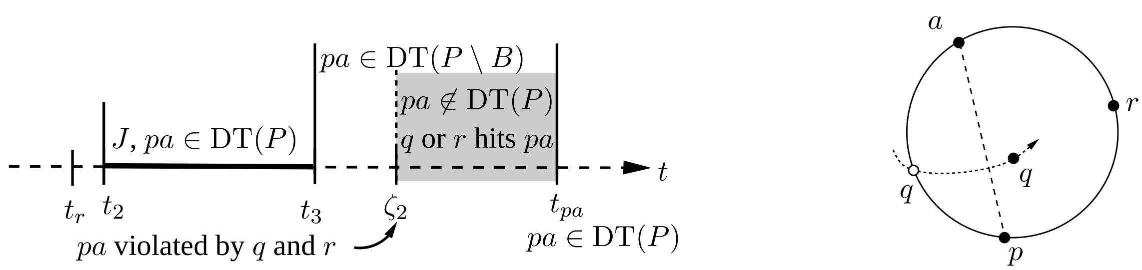

Figure 33: Left: If Condition (iii) of Theorem 2.2 holds, then we have a subset B of at most $3h$ points whose removal restores the Delaunayhood of pa throughout $[t_2, t_{pa}] = J \cup [t_3, t_{pa}]$. Right: If q hits pa during $[t_3, t_{pa}]$, then $(P \setminus B) \cup \{q\}$ induces a Delaunay crossing of pa by q.

Assume then that Condition (iii) of Theorem 2.2 holds. That is, P contains a subset B of at most $3h$ points (possibly including one, or both of the points q, r) such that the edge pa belongs to $\mathrm{DT}(P \setminus B)$ throughout the interval $J \cup [t_3, t_{pa}] = [t_2, t_{pa}]$. See Figure 33 (left).

If pa is crossed by r during $[t_3, t_{pa}]$, then the smaller set $(P \setminus B) \cup \{r\}$ yields two Delaunay crossings of pa by the same point r. The routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that the overall number of such triples p, a, r in P is at most $O(hn^2)$. Clearly, this also bounds the overall number of such quadruples σ.

Assume then that pa is hit by q, as depicted in Figure 33 (right). If this happens twice during (t_3, t_{pa}) then the smaller set $(P \setminus B) \cup \{q\}$ induces either two single Delaunay crossings or one double Delaunay crossing, of the edge pa by q. In each of these cases, we can show, as usual, that the overall number of such triples p, q, a in P is at most $O(hn^2)$ by combining Lemma 4.5 or Theorem 4.3 with the probabilistic argument of Clarkson and Shor. Furthermore, (pa, r, J) is among the last $4\ell + 1$ Delaunay crossings that the edge pa undergoes before being hit by q. Hence, this scenario occurs for at most $O(\ell h n^2)$ Delaunay quadruples σ in $\mathcal{F}$.

To recap, we may assume that q hits the edge pa only once during (t_3, t_{pa}), so this edge undergoes a single Delaunay crossing by q within $(P \setminus B) \cup \{q\}$.

Probabilistic refinement. Consider a random sample R of $\lceil n/h \rceil$ points chosen at random from P. Notice that the following two conditions hold simultaneously, with probability at least $\Omega(1/h^4)$: (1) the four points of σ belong to R, and (2) R includes none of the points of $B \setminus \{q, r\}$.

If the sample R is indeed successful, the four points p, q, a, r define a Delaunay quadruple with respect to R. Let $\mathcal{F}_R$ be the resulting family of such hereditary Delaunay quadruples in R. Clearly, $\mathcal{F}_R$ is 3-refined (with respect to the underlying set R). In addition, each quadruple σ in $\mathcal{F}_R$ satisfies the following new condition:

(Q7) The edge pa belongs to the triangulation $\mathrm{DT}(R \setminus \{q, r\})$ throughout the interval (t_2, t_{pa}), where t_{pa} denotes the first time after ζ_2 when the edge pa re-enters $\mathrm{DT}(R)$. Moreover, pa is hit in $(t_3, t_{pa}]$ by q, but not by r, and this occurs only once during $(t_3, t_{pa}]$. In particular, the point set $R \setminus \{r\}$ induces a single Delaunay crossing $(pa, q, \mathcal{I}_r)$, whose interval $\mathcal{I}_r$ is contained in $(t_3, t_{pa}]$.

We say that a family $\mathcal{F}$ of quadruples is 4-*refined* if (1) it is 3-refined, and (2) its quadruples satisfy the above condition (Q7) with respect to the underlying point set P (i.e., with R replaced by P). For each quadruple σ in such a 4-refined family $\mathcal{F}$, we refer to the corresponding crossing $(pa, q, \mathcal{I}_r)$ (which figures in condition (Q7)) as the *special crossing* of pa by q in $\mathcal{F}$.

As in the previous conditions, when regarding R as an underlying point set, some of the critical times (e.g., t_{pa}) may shift. As is easy to check Condition (Q7), we have the following analogue of Propositions 5.2 and 5.3, showing that the notion of a special crossing is well defined:

Proposition 5.7. *Let $\mathcal{F}$ be a 4-refined family of Delaunay quadruples. Then every quadruple $\sigma = (p,q,a,r)$ in $\mathcal{F}$ is uniquely determined by its triple (p,a,q). Hence, there is one-to-one correspondence between Delaunay quadruples of $\mathcal{F}$ and their special crossings, so it remains to bound the number of the latter.*

Proof. We are to show that the fourth point, r, of σ, is uniquely determined by the first three points p, a, r. Indeed, by condition (Q7), r the *last* point of P to cross pa, from L^-_{pq} to L^+_{pq}, before q performs this same type of crossing. $\square$

Let $\Psi_4(m)$ denote the maximum cardinality of a 4-refined family $\mathcal{F}$ of Delaunay quadruples that is defined with respect to a set of m moving points. The preceding discussion implies the following recurrence:

$$\Psi_3(n) = O\left(h^4\Psi_4(n/h) + \ell h^2 N(n/h) + k\ell^2 N(n/\ell) + k^2 N(n/k) + \ell hn\beta(n)\right), \qquad (6)$$

for any triple of parameters $12 \ll k \ll \ell \ll h$.

By the above Proposition 5.7, there is one-to-one correspondence between Delaunay quadruples $\sigma = (p,q,a,r)$ of a 4-refined family $\mathcal{F}$, and their respective triples (p,q,a), which yield the corresponding special crossings, so it suffices to bound the number of the latter configurations. This is indeed done in Section 6, whose analysis is formulated mainly in the terms of *special* crossings. However, before we proceed in that direction, one last refinement is in order.

5.7 Stage 5: Extending the almost-Delaunayhood of pq

Let $\mathcal{F}$ be a 4-refined family of Delaunay quadruples, which is defined over a set P of n moving points. Let $\sigma = (p,q,a,r)$ be a Delaunay quadruple in $\mathcal{F}$, which satisfies all the seven conditions (Q1)–(Q7) that were enforced in the course of the preceding four stages.

Note that the edge pq belongs to $\mathrm{DT}(P \setminus \{r\})$ throughout the interval I of its Delaunay crossing by r. Furthermore, by condition (Q7), the edge pa undergoes in $P \setminus \{r\}$ a Delaunay crossing $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$. Hence, Lemma 3.1 implies that pq belongs to $\mathrm{DT}(P \setminus \{r\})$ also during $\mathcal{I}_r$. We next extend the almost-Delaunayhood of pq from I and $\mathcal{I}_r$ to the rest of $[I, \mathcal{I}_r] = \mathsf{conv}(I \cup \mathcal{I}_r)$.

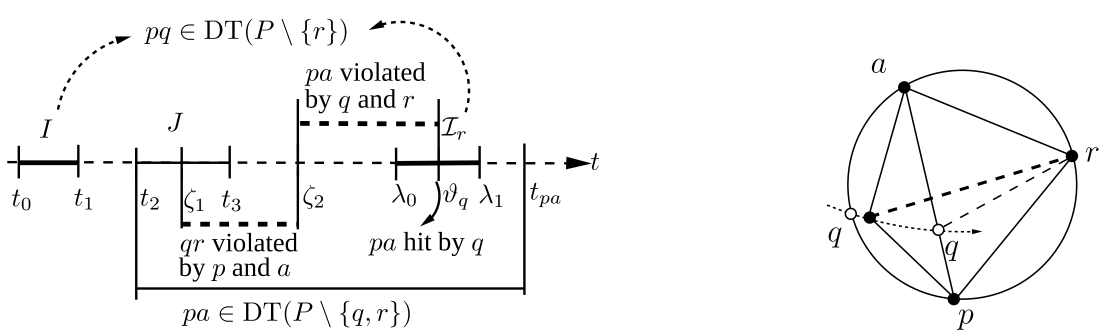

Figure 34: Left: The setup at the beginning of Stage 5. Note that the edge pq belongs to $\mathrm{DT}(P \setminus \{r\})$ throughout each of the intervals I and $\mathcal{I}_r$. The Delaunayhood of rq is violated by p and a between the last two co-circularities ζ_1, ζ_2. The edge pa is hit by q at some time $\vartheta_q \in (\zeta_2, t_{pa})$, and its Delaunayhood is violated by q and r throughout the interval (ζ_2, ϑ_q). Right: A possible motion of q during (ζ_2, ϑ_q).

Setup. Refer to Figure 34. By condition (Q3), the Delaunayhood of rq is violated by $p \in L^-_{rq}$ and $a \in L^+_{rq}$ between the last two co-circularities $\zeta_1 \in J$ and $\zeta_2 > t_3$ of p, q, a, r (both of them red-blue with respect to pa and rq). Right after time ζ_2 (when rq is freed from the above violation by p and a), the Delaunayhood of pa is violated by $q \in L^-_{pa}$ and $r \in L^+_{pa}$. By condition (Q7), pa re-enters $\mathrm{DT}(P)$ at some time $t_{pa} > \zeta_2$ (which is the first such time after ζ_2), and belongs to $\mathrm{DT}(P \setminus \{r, q\})$ throughout $(t_3, \lambda_{pa}]$. Finally, pa is hit at some time in $(t_3, t_{pa}]$ by q but not by r. Hence, applying Lemma 3.1 from time ζ_2, we conclude that q crosses pa from L^-_{pa} to L^+_{pa} at some moment $\vartheta_q \in (\zeta_2, t_{pa}]$, with the property that the Delaunayhood of pa is violated by $q \in L^-_{pq}$ and $r \in L^+_{pq}$ throughout (ζ_2, ϑ_q). In particular, the aforementioned special crossing $(pa, q, \mathcal{I}_r)$ in $P \setminus \{r\}$ occurs entirely during $(t_3, t_{pa}]$, and its interval $\mathcal{I}_r$ contains the above time ϑ_q when q enters L^+_{pa}. (However, $\mathcal{I}_r$ need not necessarily contain ζ_2.)

The preceding discussion implies that the intervals $I = [t_0, t_1]$ and $\mathcal{I}_r = [\lambda_0, \lambda_1]$ (where pq is known to be almost Delaunay) are indeed disjoint. We also emphasize that the edges pa and rq intersect throughout $(\zeta_1, \vartheta_q) = (\zeta_1, \zeta_2) \cup (\zeta_2, \vartheta_q)$.

To enforce the almost-Delaunayhood of pq in the resulting gap (t_1, λ_0), we fix a pair of constants $12 < k \ll \ell$ and proceed in two steps.

Charging events in $\mathcal{A}_{rq}$. As a preparation, we first extend the almost-Delaunayhood of rq. Recall that, by condition (Q6), rq belongs to $\mathrm{DT}(P \setminus \{p, a\})$ throughout the interval (t_{ra}, t_{rq}). Here t_{rq} denotes the first time after t_3 when rq is Delaunay, and t_{ra} denotes the last time before (or at) t_0 when ra is Delaunay. Note that (t_{ra}, t_{rq}) contains the respective times ζ_0, ζ_1 and ζ_2 of the three co-circularities co-circularities of p, q, a, r. Recall also that ζ_2 occurs after the ending time t_3 of J. Hence, the inequality $t_{rq} > t_3$ is strict, so rq is not Delaunay right before time t_{rq}.

We next extend the almost-Delaunayhood of rq to a potentially larger interval (t_{ra}, ϑ_q) (where, as above, ϑ_q denotes the time in $\mathcal{I}_r$ when q enters the halfplane L^+_{pa} through pa). We can assume, with no loss of generality, that $t_{rq} < \vartheta_q$. (Otherwise, we are done.) Therefore, and since $\zeta_2 < t_{rq}$, the Delaunayhood of pa is violated by $q \in L^-_{pa}$ and $r \in L^+_{pa}$ throughout the interval $(t_{rq}, \vartheta_q) \subset (\zeta_2, \vartheta_q)$.

We apply Theorem 2.2 in $\mathcal{A}_{rq}$ over the interval (t_{rq}, ϑ_q), and with the first constant k. (This is possible because rq is Delaunay at time t_{rq}.) In the first two cases of Theorem 2.2, we charge σ (via (pq, r, I)) either to a k-shallow collinearity, or to $\Omega(k^2)$ k-shallow co-circularities. Below we argue that any event in $\mathcal{A}_{rq}$ is charged as above by at most $O(1)$ quadruples σ.

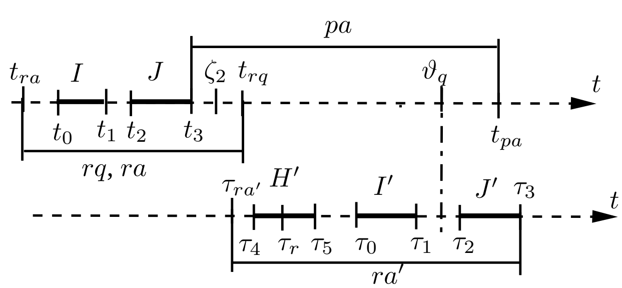

Figure 35: Proposition 5.8: The subfamily Γ_{qr} contains at most 3 quadruples $\sigma' = (p', q, a', r)$ whose respective crossings $(p'q, r, I')$ end in (t_1, ϑ_q). To establish the proposition, we fix such a quadruple σ', with $p' \neq a$ and $a' \neq p$, and argue that the second crossing $(p'a', r, J')$ of σ' ends after ϑ_q.

Note that the respective points q and r of σ can be chosen in $O(1)$ possible ways from among the three or four points involved in the event. Now consider the subfamily Γ_{qr} of all quadruples $\sigma' = (p', q, a', r) \in \mathcal{F}$ whose second and fourth points are equal to q and r, respectively. (In particular, Γ_{qr} includes the quadruple $\sigma = (p, q, a, r)$ under consideration.) Notice that each $\sigma' \in \Gamma_{qr}$ is composed of two clockwise (p', r)-crossings $(p'q, r, I' = [\tau_0, \tau_1])$, $(p'a', r, J' = [\tau_2, \tau_3])$, and comes with a counterclockwise (r, q)-crossing $(a'r, q, H' = [\tau_4, \tau_5])$ (which occurs in the smaller set $P \setminus \{p'\}$, and before

J' begins). Note also that the first crossing $(p'q, r, I')$ of σ' is also a counterclockwise (q, r)-crossing.

Proposition 5.8 below implies that the first crossing (pq, r, I) of σ is among the last four such (q, r)-crossings $(p'q, r, I')$ to end before any event that occurs in $\mathcal{A}_{rq}$ during (t_{rq}, ϑ_q). (See Figure 35 for a schematic illustration.)

Proposition 5.8. *With the above notation, the family* Γ_{qr} *contains at most* 3 *quadruples* $\sigma' = (p', q, a', r)$ *whose respective first crossings* $(p'q, r, I' = [\tau_0, \tau_1])$ *end in* (t_1, ϑ_q).

Hence, any k-shallow co-circularity or k-shallow collinearity is charged as above by at most $O(1)$ quadruples of Γ_{qr}, so the above charging accounts for at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$ quadruples $\sigma \in \mathcal{F}$.

We can assume, then, that Condition (iii) of Theorem 2.2 holds, so there is a set A_{rq} of at most $3k$ point whose removal restores the Delaunayhood rq throughout (t_{rq}, ϑ_q).

Proof of Proposition 5.8. Propositions 5.2, 5.3 and 5.7 imply that (i) there exist at most 2 quadruples $\sigma' = (p', q, a', r) \in \Gamma_{qr}$ with $p'' = a$ or $a'' = p$, and (ii) for any other choice of $\sigma' \in \Gamma_{qr} \setminus \{\sigma\}$, all the six points p, q, a, r, p', a' are distinct.

Consider all the quadruples quadruples $\sigma' \in \Gamma_{qr}$ that fall into the second category, and whose first crossings $(p'q, r, I')$ end in (t_{rq}, ϑ_q). Let σ' be the unique quadruple of this kind whose respective first crossing $(p'q, r, I' = [\tau_0, \tau_1])$ ends *first*. (That is, there is no other quadruple $\sigma'' = (p'', q, a'', r) \in \Gamma_{qr}$ that satisfies $\{p'', a''\} \cap \{p, a\} = \emptyset$, and whose first crossing $(p''q, r, I'')$ ends in (t_{rq}, τ_1).) Refer to Figure 35.

Let $\tau_{ra'}$ denote the last time before (or at) the beginning τ_0 of I' when the edge ra' is Delaunay. Since σ' is 4-refined, the respective intervals $I' = [\tau_0, \tau_1]$, $J' = [\tau_2, \tau_3]$, and H', of σ', are all contained in $[\tau_{ra'}, \tau_3]$. Condition (Q6) on σ' implies that rq belongs to $\mathrm{DT}(P \setminus \{p', a'\})$ throughout $[\tau_{ra'}, \tau_3]$. Therefore, and since both I' and J' end after t_{rq}, we get that $\zeta_2 < \tau_{ra'}$. (Otherwise, we would get $\tau_{ra'} < \zeta_2 < t_{rq} < \tau_1 < \tau_3$, so the above interval $[\tau_{ra'}, \tau_3]$ would contain the time ζ_2, right before which the Delaunayhood of rq is violated by p and a).

By the choice of σ', any quadruple $\sigma'' = (p'', q, r, a'') \in \Gamma_{qr}$ whose respective (q, r)-crossing $(p''q, r, I'')$ ends in (t_{rq}, τ_1), must satisfy $p'' = a$ or $a'' = p$. Furthermore, Condition (Q2) (on σ') implies that there exist no quadruples $\sigma'' \in \Gamma_{qr}$ whose respective (q, r)-crossings $(p''q, r, I'')$ end in (τ_1, τ_3). It, therefore, suffices to show that $\tau_3 > \vartheta_q$ (that is, that the second crossing $(p'a', r, J')$ of σ' ends after q enters L_{pa}^+).

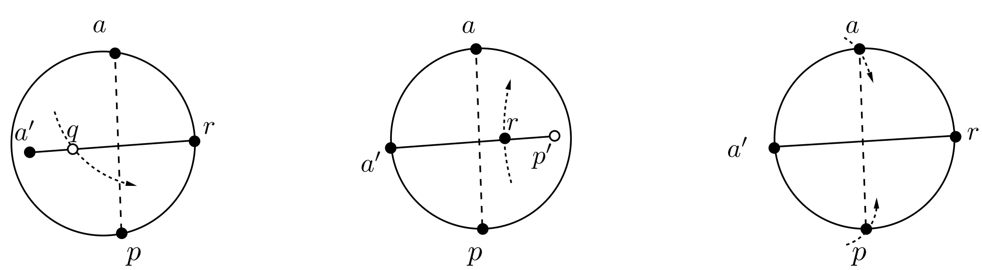

Figure 36: Proof of Proposition 5.8. We assume, for a contradiction, that $\tau_3 < \vartheta_q$, so both crossings $(a'r, q, H')$ and $(p'a', r, J')$ occur within (ζ_2, ϑ_q). Left: At the time $\tau_q \in H'$ when q hits $a'r$, the Delaunayhood of pa is violated by a and r'. Center: If ar' and pa still intersect at the time in J' when r hits $p'a'$, then the Delaunayhood of pa is violated by p' and a' at some moment during $(\zeta_2, \vartheta_q) \subset (t_3, t_{pa})$. Right: The last scenario, where pa recovers from its previous violation by a' and r through a co-circularity.

Indeed, assume for a contradiction that $\tau_3 < \vartheta_q$. Then, recalling that $\tau_{ra'} > \zeta_2$, we conclude that $[\tau_{ra'}, \tau_3]$ is contained in the interval (ζ_2, ϑ_q), where the Delaunayhood of pa is violated by q and r.

By condition (Q1) on σ', its edge ra' belongs to $\mathrm{DT}(P \setminus \{p', q\})$ throughout $[\tau_{ra'}, \tau_3]$. Hence, at the time $\tau_r \in H' \subset [\tau_{ra'}, \tau_3]$ when q enters $L^+_{a'r}$, the edge pa is intersected by $a'r = a'q \cup qr$, so the Delaunayhood of pa is violated then by r and a'. See Figure 36 (left). (Otherwise, the Delaunayhood of ra' would be violated by p and a, which is impossible during $[\tau_{ra'}, \tau_3]$.)

If ra' still intersects pa at the time in $J' \subset (\zeta_2, \vartheta_q)$ when r hits $p'a'$ during the second crossing of σ', then the same argument shows that Delaunayhood of pa is violated then by p' and a', contrary to condition (Q7) on σ. (See Figure 36 (center).) Otherwise, there is a time in $(\tau_{r'}, \tau_3)$ when the edge pa recovers from its previous violation by r and a'. Notice that, by condition (Q7), none of r, a' can hit pa during the above interval (which is contained in $(\zeta_2, \vartheta_q) \subset (t_3, t_{pa})$). Applying Lemma 3.1 for $\{p, a, r, a'\}$, we get that the four points p, a, r, a' are involved during $(\tau_{ra'}, \tau_3)$ in a red-blue co-circularity with respect to pa and ra' (as depicted in Figure 36 (right)), contrary to the almost-Delaunayhood of ra' in $(\tau_{ra'}, \tau_3)$. This final contradiction completes the proof of Proposition 5.8. $\Box$

We thus can assume, in what follows, that there is a subset A_{rq} of at most $3k$ points whose removal restores the Delaunayhood of rq throughout (t_{rq}, ϑ_q).

Charging events in $\mathcal{A}_{pq}$. We apply Theorem 2.2 in $\mathcal{A}_{pq}$ over the interval (t_1, ϑ_q), which covers the gap (t_1, λ_0) between I and $\mathcal{I}_r$.

In cases (i) and (ii) of Theorem 2.2, we charge σ within $\mathcal{A}_{pq}$ either to an ℓ-shallow collinearity or to $\Omega(\ell^2)$ ℓ-shallow co-circularities. We claim that any such event, which occurs in $\mathcal{A}_{pq}$ during (t_1, ϑ_q), is charged in this manner by at most $O(k)$ quadruples $\sigma = (p, q, a, r)$.

Indeed, the points p and q of σ can be guessed in $O(1)$ possible ways among the three or four points involved in the event. Let $\mathcal{Q}_{pq}$ denote the sub-family of all quadruples $\sigma' = (p, q, a', r') \in \mathcal{F}$ whose first two points are equal to p and q, respectively. Note that $\mathcal{Q}_{pq}$ includes the quadruple σ under consideration, and that, for each $\sigma' \in \mathcal{Q}_{pq}$, its first crossing is of the form (pq, r', I'). Proposition 5.9 (below) implies that the first crossing (pq, r, I) of σ is among the last $6k + 3$ such crossings to end before any ℓ-shallow event that we charge in $\mathcal{A}_{pq}$. Hence, the above charging applies to at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ quadruples.

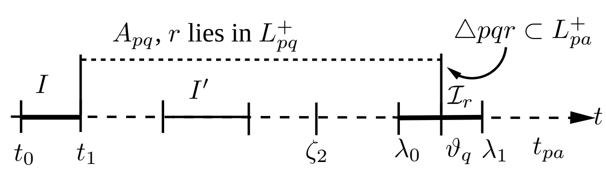

Figure 37: Extending the almost-Delaunayhood of pq to (t_1, λ_0). We apply Theorem 2.2 over the larger interval (t_1, ϑ_q). By Proposition 5.9, the family $\mathcal{Q}_{pq}$ contains at most $6k+2$ quadruples $\sigma' = (p, q, a', r')$ whose respective first crossings (pq, r', I') end in (t_1, ϑ_q).

Finally, if Condition (iii) of Theorem 2.2 is satisfied, we end up with a subset A_{pq} of at most 3ℓ points whose removal restores the Delaunayhood of pq throughout (t_1, ϑ_q). In this case, we can "free" σ from the points of $A_{pq} \setminus \{a, r\}$ (thereby extending the almost-Delaunayhood of pq to $(t_1, \lambda_0) \subset (t_1, \vartheta_q)$) through the standard probabilistic argument.

Proposition 5.9. *The family $\mathcal{Q}_{pq}$ contains at most $6k + 2$ quadruples $\sigma' = (p, q, a', r')$ whose respective crossings (pq, r', I') end in (t_1, ϑ_q).*

Proof. Since p and q are fixed, Propositions 5.2 and 5.7 imply that any $\sigma' \in \mathcal{Q}_{pq}$ is uniquely determined by each of its respective points a', r'. Hence, we have at most two $\sigma' = (p, q, a', r') \in \mathcal{Q}_{pq}$ that satisfy $a' = r$ or $r' = a$, and, for any other quadruple in $\mathcal{Q}_{pq}$, none of its respective points a' and r' is equal to a or r.

Fix $\sigma' = (p, q, a', r') \in \mathcal{Q}_{pq}$ whose respective first crossing (pq, r', I') ends in (t_1, ϑ_q), and with the property that $\{a, r\} \cap \{a', r'\} \neq \emptyset$. To establish the proposition, it suffices to show that, for any such σ', at least one of its points a', r' belongs to the set A_{rq} (obtained at the end of the preceding step) of cardinality is at most $3k$. Indeed, we have at most $3k$ quadruples σ' with $a' \in A_{rq}$, and at most $3k$ quadruples σ' with $r' \in A_{rq}$, and each $\sigma' \in \mathcal{Q}_{pq}$ is uniquely determined by each of its respective points a' and r'.

We, therefore, proceed to establishing the latter property. Notice that the intervals I and I' are disjoint, so we have $I' \subset (t_1, \vartheta_q)$. Note also that r lies in L^+_{pq} right after time t_1, and also at the later time ϑ_q when q hits pa (thereby freeing pa from its previous violation by $q \in L^-_{pa}$ and $r \in L^+_{pa}$). Hence, r has to remain in L^+_{pq} throughout (t_1, ϑ_q) (or, else, it would cross L_{pq} three times during $I \cup (t_1, \vartheta_q)$); see Figure 37. We thus conclude that, at the moment in I' when r' hits pq, r' enters the triangle $\triangle pqr$ (whose order type remains fixed throughout (t_1, ϑ_q)).

Claim 5.10. *Let t' be the time in I' when the above point $r' \in P \setminus P_\sigma$ enters $\triangle pqr$ through the interior of pq. Then r' must leave $\triangle pqr$ during (t', ϑ_q).*

Proof. Assume for a contradiction that r' remains in $\triangle pqr$ throughout (t', ϑ_q). Recall that pa is intersected by rq throughout $(t_3, \vartheta_q) \subset (\zeta_1, \zeta_2) \cup (\zeta_2, \vartheta_q)$, with $q \in L^-_{pa}$ and $r \in L^+_{pa}$. Observe that there is a time in $[t_3, \vartheta_q)$ when r' lies within $\triangle pqr \cap L^-_{pa}$. Indeed, this property clearly holds if r' enters L^+_{pa} in the interval (t_3, ϑ_q), where pq is contained in L^-_{pa}; see Figure 38 (center). Assume then that r' enters $\triangle pqr$ before t_3 (i.e., $t_3 \in (t', \vartheta_q)$). However, in this case r' has to lie at time t_3 within $\triangle pqr \cap L^-_{pa}$, as depicted in Figure 38 (left). (Otherwise, r' would lie at that moment in the cap $B[p, a, r] \cap L^+_{pa} \supset \triangle pqr \cap L^+_{pa}$, which is known to be P-empty throughout the second portion of $J = [t_2, t_3]$.)

To see a contradiction, notice that $\triangle pqr$ lies at time ϑ_q entirely within the closure of L^+_{pa}; see Figure 38 (right). Therefore, r' too has to enter L^+_{pa} during (t_3, ϑ_q). However, r' cannot cross L_{pa} during (t_3, ϑ_q) through one of its rays outside pa and while remaining inside the triangle $\triangle pqr$ (because the segments pa and rq intersect there), and condition (Q7) on σ implies that r' cannot hit pa during $(t_3, \vartheta_q) \subset [t_3, t_{pa}]$. This contradiction completes the proof of Claim 5.10. $\square$

Consider the first time in (t', ϑ_q) when r' leaves $\triangle pqr$, through one of the edges pr, pq, rq. (Here, as before, t' denotes the time when r' hits pq during the first Delaunay crossing (pq, r', I') of $\sigma' = (p, q, a', r')$.) Recall that r' cannot cross pr during (t_1, t_3), because σ is a Delaunay quadruple (that is, pr belongs to $\mathrm{DT}(P \setminus \{q, a\})$ throughout $[I, J] = [t_0, t_3]$). Furthermore, r' cannot cross pr in (t_3, ϑ_q) either: otherwise r' would first have to enter L^+_{pa} through the relative interior of pa, contrary to condition (Q7) on σ. We, thereby, conclude that r' can leave $\triangle pqr$ during $(t', \vartheta_q) \subset (t_1, \vartheta_q)$ only through one of the remaining edges qr and pq.

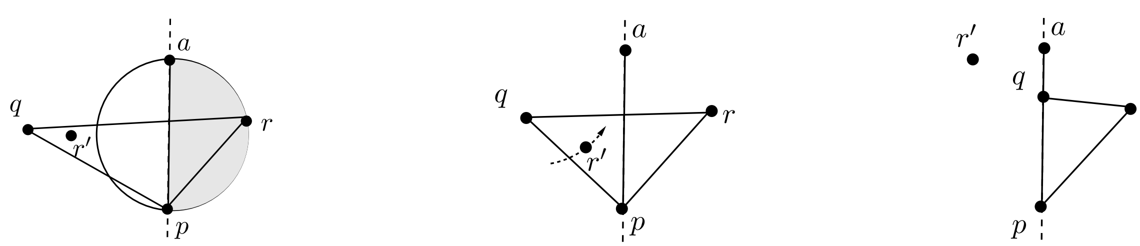

Figure 38: Proof of Claim 5.10. Left: If r' enters $\triangle pqr$ during $[t_3, \vartheta_q)$, this can happen only within L^-_{pa}. Center: If $t' < t_3$ then r' lies in $\triangle pqr \cap L^-_{pa}$ at time t_3 (because the rest of $\triangle pqr$ lies inside the P-empty cap $B[p, a, r] \cap L^+_{pa}$). Right: In both cases, r' must exit $\triangle pqr$ before time ϑ_q (at which $\triangle pqr$ passes entirely to L^+_{pa}).

If r' exits $\triangle pqr$ during $(t', \vartheta_q) \subset (t_1, \vartheta_q)$ through the relative interior rq, then, by condition (Q2), this can occur only in the smaller interval (t_{rq}, ϑ_q) (and only if $\vartheta_q > t_{rq}$). Hence, in this case q belongs to A_{rq}, and we are done.

Assume, then, that r' leaves $\triangle pqr$ through the edge pq, as depicted in Figure 39. Consider the second Delaunay crossing (pa', r', J') of $\sigma' = (p, q, a', r')$. Recall that I' begins after t_1 and before the beginning of J', so (pa', r', J') occurs too after the end of I. Since $\sigma' \in \mathcal{Q}_{pq}$ is 4-refined, the point r' remains L_{pq}^+ after t' and until the end of J' (or, else, r' would cross L_{pq} three times). Therefore, J' ends before r' exits $\triangle pqr$ through pq (and, in particular, before ϑ_q). To conclude, the second crossing (pa', r', J') of σ' occurs entirely within (t_1, ϑ_q). To complete our analysis, we distinguish between the following two sub-cases:

If a' lies in L_{rq}^+ at the time in J' when r' hits pa', then rq is intersected at that moment by the Delaunay edge $r'a'$; see Figure 39 (left). Hence, Delaunayhood of rq is violated at some moment in $J' \subset (t_1, \vartheta_q)$ by r' and a'. Furthermore, condition (Q2) on σ implies that the above violation is possible only during (t_{rq}, ϑ_q), so at least one of a', r' must belong to A_{rq}.

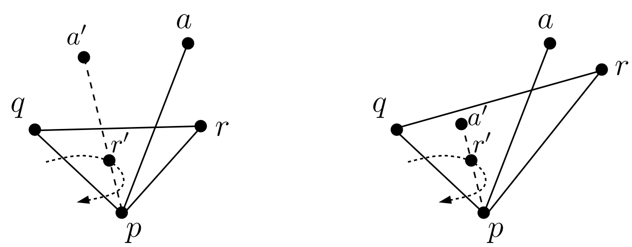

Figure 39: Proof of Proposition 5.9. The second (p, r')-crossing (pa', r', J') of σ' ends before r' hits pq again. The two possible scenarios are depicted.

Assume, then, that a' lies in L_{rq}^- when r' hits pa' during J'. Hence, both points r', a' lie at that time inside the triangle $\triangle pqr$; see Figure 39 (right). Arguing as before, we conclude that a' leaves $\triangle pqr$ before ϑ_q through one of the edges rq and pq. However, condition (Q7) on σ' implies that a' cannot leave $\triangle pqr$ through the edge pq: otherwise q would enter the halfplane L_{pa}^+ twice (once during the respective special crossing of σ', and another time through one of the outer rays of $L_{pa} \setminus pa$). Therefore, in this case a' can leave $\triangle pqr$ before ϑ_q only through the relative interior of rq. Arguing as before, we conclude that a' again belongs to A_{rq}. $\square$

To recap, the previous chargings within $\mathcal{A}_{pq}$ and $\mathcal{A}_{rq}$ altogether account for at most $O(k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n))$ quadruples in our 4-refined family $\mathcal{F}$. Each surviving quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$ comes with a subset A_{pq} of at most 3ℓ points so that pq is Delaunay in $P \setminus A_{pq}$ throughout the gap $(t_1, \lambda_0) \subset (t_1, \vartheta_q)$ between the respective intervals I and $\mathcal{I}_r$ of σ.

Probabilistic refinement. We apply the probabilistic argument of Clarkson and Shor [9] one more time.

We say that a family $\mathcal{F}$ of Delaunay quadruples is *5-refined* or, simply, *refined* if it is 4-refined with respect to the underlying point set P, and each quadruple σ in $\mathcal{F}$ satisfies the following new condition:

(Q8) The edge pq belongs to $\mathrm{DT}(P \setminus \{a, r\})$ throughout the respective interval $[I, I_r] = [t_0, \lambda_1]$. (Here, as above, $I = [t_0, t_1]$ is the interval of the first (p, r)-crossings of σ, and $\mathcal{I}_r = [\lambda_0, \lambda_1]$ is the interval of the special crossing of pa by q.)

That is, we require that the family $\mathcal{F}$ is nonoverlapping, and that its quadruples are Delaunay and satisfy all the 8 conditions (Q1) – (Q8).

Let $\Psi_5(n)$ denote the maximum cardinality of a refined family of Delaunay quadruples, that can be defined over an underlying set of n moving points.

The routine sampling argument of Clarkson and Shor [9] leads to the following recurrence:

$$\Psi_4(n) \le O\left(\ell^4\Psi_5(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + kn^2\beta(n)\right).$$

As argued in the previous section, there is one-to-one correspondence between (1) quadruples $\sigma = (p,q,a,r)$ in a refined family $\mathcal{F}$, (2) their respective triples (p,q,a), and (3) the special crossings $(pa,q,\mathcal{I}_r)$ performed by these triples.

As reviewed in the beginning of this section, the analysis of $\Psi_5(m)$ is delegated to Section 6, which primarily deals with the third type of configurations.

6 Special Crossings and Special Quadruples

In the preceding section we have established a sequence of recurrences implying that the maximum number $\Psi(n)$ of consecutive quadruples (and, hence, the maximum number $N(n)$ of Delaunay co-circularities) in a set P of n moving points is (asymptotically) dominated by the maximum possible cardinality $\Psi_5(m)$ of a *refined* family $\mathcal{F}$ of Delaunay quadruples that is defined over of a certain m-size subsample $R \subset P$.

To bound the above quantity $\Psi_5(n)$, for any $n > 0$, we fix a set P, and a refined family $\mathcal{F}$ of (clockwise) Delaunay quadruples that is defined over P. That is, $\mathcal{F}$ is nonoverlapping, and each of its quadruples $\sigma = (p,q,a,r)$ satisfies the eight conditions (Q1) – (Q8) (stated in terms of p,q,a,r, $\mathcal{F}$ and P).

In particular, every triple of points of $\sigma = (p,q,a,r) \in \mathcal{F}$ yield a Delaunay crossing, which sometimes occurs within a *reduced* triangulation obtained by omitting from P the remaining fourth point of σ. Indeed, recall that σ, as any clockwise quadruple, is formed by a pair of clockwise (p,r)-crossings (pq,r,I) and (pa,r,J). The two additional crossings (ar,q,H) and $(pa,q,\mathcal{I}_r)$ have been enforced at Stages 1 and 4 of Section 5, as parts of the respective conditions (Q3) and (Q7), and they occur within the respective appropriately *reduced* triangulations $\mathrm{DT}(P\setminus\{p\})$ and $\mathrm{DT}(P\setminus\{r\})$.

Recall also that, according to Propositions 5.2, 5.3, and 5.7, each quadruple σ in $\mathcal{F}$ is *uniquely* determined by *each* of the four ordered triples (p,q,r), (p,a,r), (a,r,q), and (p,a,q), which realize its four Delaunay crossings. (That is, in each triple the third point performs a clockwise Delaunay crossing of the edge connecting the first two points.)

To bound the cardinality of $\mathcal{F}$, we focus, for each quadruple $\sigma = (p,q,a,r)$ in $\mathcal{F}$, on the last type of crossing $(pa,q,\mathcal{I}_r)$, realized by its first three points p,q,a, and referred to as the special crossing of pa by q. We emphasize that $(pa,q,\mathcal{I}_r)$ is also a regular Delaunay crossing which occurs in the smaller triangulation $\mathrm{DT}(P\setminus\{r\})$. For convenience of notation, we refer to r as the *outer point* of $(pa,q,\mathcal{I}_r)$.

We further label each special crossing $(pa,q,\mathcal{I}_r)$ as a *clockwise (special)* (p,q)*-crossing*, and as a *counterclockwise (special)* (a,q)*-crossing*. Notice that Lemma 4.6 need not hold for *special* (p,q)-crossings of the same type (that is, either clockwise or counterclockwise), because these are defined with respect to reduced point sets, each omitting the respective outer point r. As a matter of fact, the respective outer points of any two such (p,q)-crossings are always distinct, because, as noted above, their ancestor quadruples in $\mathcal{F}$ are uniquely determined by the respective triples (p,q,r). Hence, any two (p,q) crossings (of the same type) are always defined with respect to *distinct* point sets. Instead, we use Lemma 5.5, which imposes certain restrictions on the almost-Delaunay crossings that can be compared by it. For example, two counterclockwise special (a,q)-crossings $(pa,q,\mathcal{I}_r)$ and $(wa,q,\mathcal{I}_u)$, with respective outer points r and u, become incompatible if and only if $r = w$ or $p = u$.

We first perform a preliminary pruning step that will ensure, in particular, that Lemma 5.5 indeed applies to any pair of surviving counterclockwise special (a,q)-crossings. This will be done by considering all possible pairs of *distinct* such (a,q)-crossings $(pa,q,\mathcal{I}_r)$ and $(wa,q,\mathcal{I}_u)$, and by omitting from $\mathcal{F}$ their corresponding quadruples $\sigma = (p,q,a,r)$ and $\sigma' = (w,q,a,u)$ if they share one or more additional points, apart from q and a. A similar pruning step will ensure that any two clockwise special (p,q)-crossings $(pa,q,\mathcal{I}_r)$ and $(pw,q,\mathcal{I}_u)$ share only the pair (p,q).

The crucial observation is that the overall number of quadruples that we omit from $\mathcal{F}$, at both steps, does not exceed $O(n^2)$. Indeed, assume, for instance, that a pair $(pa,q,\mathcal{I}_r)$ and $(wa,q,\mathcal{I}_u)$ of counterclockwise special (a,q)-crossings share an additional, third point (again, apart from q and a). Recalling that each quadruple σ in $\mathcal{F}$ is uniquely determined by any ordered sub-triple of its points, we conclude that $p \neq w$ and $r \neq u$. That is, we have $r = w$ or $p = u$. Assume, with no loss of generality, that $r = w$. Recall that each ordered sub-triple in σ or in σ' performs a Delaunay crossing (perphaps within a suitably reduced triangulation). We therefore get from σ the crossing (ar,q,H), within $P \setminus \{p\}$, and we get from σ' the crossing $(wa,q,\mathcal{I}_u) = (ra,q,\mathcal{I}_u)$, within $P \setminus \{u\}$. We thus obtain two *distinct*[27] Delaunay crossings which are performed by the *same* triple $(a, r = w, q)$ and within the same reduced triangulation $\mathrm{DT}(P \setminus \{p,u\})$. Hence, a routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that the underlying point set P contains at most $O(n^2)$ such triples (a,q,r). Clearly, this also bounds the overall number of such quadruples $\sigma = (p,q,a,r)$ and $\sigma' = (w = r, q, a, u)$ that we omit. A symmetric analysis is peformed for pairs $(pa,q,\mathcal{I}_r)$ and $(pw,q,\mathcal{I}_u)$ of clockwise (p,q)-crossings that have a third point in common, and their respective quadruples $\sigma = (p,q,a,r)$ and $\sigma' = (p,q,w,u)$.

To conclude, we can assume, from now on, that any pair which consists of any two counterclockwise special (a,q)-crossings, or of any two special clockwise (p,q)-crossings, involves six distinct points (including the two outer points) and, therefore, satisfies the conditions of Lemma 5.5. Therefore, all the remaining counterclockwise special (a,q)-crossings, with a,q-fixed, can be linearly ordered by the starting times of their intervals, or by the ending times of their intervals, or by the times when q hits the corresponding a-edge, and all three orders are identical. Furthermore, Lemma 5.5 imposes a similar order on the remaining clockwise special (p,q)-crossings, with p,q fixed.

Special quadruples. We say that two counterclockwise special (a,q)-crossings are *consecutive* if they are consecutive with respect to the natural order induced by Lemma 5.5. That is, no other counterclockwise special (a,q)-crossings appear in this order between them.

Four points a,p,w,q form a *special quadruple* $\chi = (a,p,w,q)$ if we encounter two (distinct) counterclockwise special (a,q)-crossings $(pa,q,\mathcal{I}_r)$ and $(wa,q,\mathcal{J}_u)$, with the respective outer points r and u, that occur in this order (that is, q crosses pa before wa); these crossings need not be consecutive. Refer to Figure 40. We then use P_χ to denote the set which consists of the four points a,p,w,q of χ, and of the two outer points r and u.

Remark. Our notation requires some understanding from the reader: Whenever we talk about a special quadruple $\chi = (a,p,w,q)$, we also need to specify the two outer points r and u. We generally do so, but do not consider them as an integral part of the quadruple, because, until Stage 4, they do not play any role in the topological changes that the quadruple undergoes. However, the outer points will "return to life" in Stage 4, and then their presence will lead to so called *terminal quadruples* which we will use to finish up the analysis. See also the overview below.

Fix a special quadruple $\chi = (a,p,w,q)$, as above. Lemma 4.4 implies[28] that the four points a,p,w,q are involved in at least one co-circularity during $\mathcal{I}_r$, and in at least one co-circularity during $\mathcal{J}_u$. Specifically, the former co-circularity is red-blue with respect to the edges pa and qw, so it must occur before the beginning of $\mathcal{J}_u$, during $\mathcal{I}_r \setminus \mathcal{J}_u$. (See Figure 40 (center).) Similarly, the latter co-circularity is red-blue with respect to the edges wa and pq, so it must occur after the end of $\mathcal{I}_r$, during $\mathcal{J}_u \setminus \mathcal{I}_r$. (See Figure 40 (right).) Furthermore, the same argument as in Section 5.1 shows that the points of χ are involved in

[27] Indeed, recall that, in our notation, q crosses ar (during H) from L^-_{ar} to L^+_{ar}, and it crosses the reversely oriented copy ra of ar (during $\mathcal{I}_u$) from $L^+_{ra} = L^-_{ar}$ to L^-_{ra}.

[28] Since the crossings of χ are defined with respect to reduced points sets $P \setminus \{r\}$ and $P \setminus \{u\}$, this implication critically relies on the assumption that $p \neq u$ and $w \neq r$.

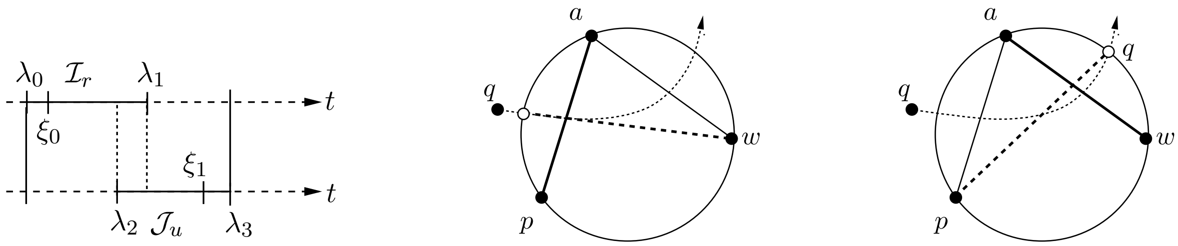

Figure 40: The special quadruple $\chi = (a, p, w, q)$. The respective intervals $\mathcal{I}_r$ and $\mathcal{J}_u$ of the two special crossings associated with χ are either disjoint, or partially overlapping (left). The points of χ are co-circular at times $\xi_0 \in \mathcal{I}_r \setminus \mathcal{J}_u$ (center) and $\xi_1 \in \mathcal{J}_u \setminus \mathcal{I}_r$ (right).

exactly one co-circularity during each of the intervals $\mathcal{I}_r$ and $\mathcal{J}_u$, and we denote the respective times of these co-circularities as $\xi_0 \in \mathcal{I}_r \setminus \mathcal{J}_u$ and $\xi_1 \in \mathcal{J}_u \setminus \mathcal{I}_r$.

It is also instructive to note that the triangulation $\mathrm{DT}(P \setminus \{r, u\})$ contains an ordinary counterclockwise quadruple (a, p, w, q), with the associated Delaunay crossings $(pa, q, \mathcal{I})$ and $(wa, q, \mathcal{J})$, such that $\mathcal{I} \subseteq \mathcal{I}_r$ and $\mathcal{J} \subseteq \mathcal{J}_u$. This immediately implies that the statement of Lemma 5.1 (or, more precisely, of its counterclockwise variant) must hold also for the counterclockwise *special* quadruples.

Consecutive special quadruples. We say that the special quadruple $\chi = (a, p, w, q)$, as above, is *consecutive* if its counterclockwise (a, q)-crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$ are consecutive in the previously established order (implied by Lemma 5.5). In this case, $\chi = (a, p, w, q)$ is uniquely determined by each of its crossings $(pa, q, \mathcal{I}_r), (wa, q, \mathcal{J}_u)$. This, combined with Propositions 5.2, 5.3 and 5.7, implies that χ is uniquely determined by every (ordered) triple of points that are chosen from the *same* quadruple (p, q, a, r) or (w, q, a, u); see Figure 41. That is, the following statement holds (with the above assumptions):

Proposition 6.1. *Let $\chi = (a, p, w, q)$ be a consecutive special quadruple, and let $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$ be the special crossings associated with χ, with respective outer points r and u. Then χ is uniquely determined by each of the following eight triples:* $(p, q, a), (p, q, r), (p, a, r), (a, r, q), (w, q, a), (w, q, u), (w, a, u)$, *and* (a, u, q).

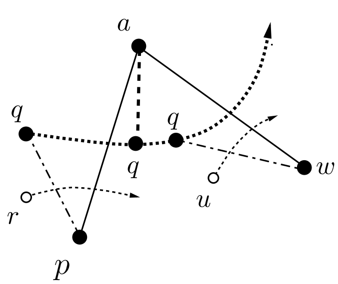

Figure 41: A consecutive counterclockwise special quadruple $\chi = (a, p, w, q)$, composed of two special crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$, with respective outer points r and u. The special crossings of χ correspond to regular Delaunay quadruples (p, q, a, r) and (w, q, a, u) in $\mathcal{F}$.

Let $\Phi(n)$ denote the maximum number of consecutive special quadruples that can be induced by a set of n moving points and a refined family $\mathcal{F}$ of Delaunay quadruples. The preceding discussion implies the following relation between the maximum possible numbers of special crossings (identified with their respective *ordinary* quadruples in $\mathcal{F}$) and consecutive *special* quadruples:

$$\Psi_5(n) = \Phi(n) + O(n^2).$$

Overview. The analysis of special consecutive (counterclockwise) quadruples proceeds through five stages, numbered $0, 1, \ldots, 4$.

At the i-th stage we consider a certain subclass of consecutive (counterclockwise) special quadruples, defined with respect to a refined family $\mathcal{F}$, which is constructed over the underlying set P of n moving points. We assume that each quadruple $\chi = (a, p, w, q)$ under consideration satisfies certain topological conditions, which are formulated in terms of the extended set P_χ (including the outer points r and u of the two special crossings associated with χ), $\mathcal{F}$, and P. At each new stage we enforce one, or several new conditions, so our special quadruples become progressively constrained.

The first four stages $i = 0, \ldots, 3$ are almost identical to the corresponding stages described in Section 5. Informally, we put the outer points r and u aside and then gradually enforce upon our quadruples χ the counterclockwise variants of the six conditions (Q1)–(Q6), which arise in the similar stages of Section 5. As noted above, this requires some caution, as the corresponding special (a, q)-crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$ are defined in terms of the (distinct) reduced point sets $P \setminus \{r\}$ and $P \setminus \{u\}$.

At each of these four stages, we first invoke Theorems 2.2, 4.3 and 5.3, and Lemma 4.5, in order to dispose of all special quadruples that fail to satisfy the newly enforced conditions, even after removal of a small-size subset of P. The surviving quadruples are passed on to the next stage, after an appropriate probabilistic refinement.

At the last Stage 4 we follow the same strategy and first apply a sequence of preparatory chargings, similar to those described in Section 5.6. To handle the remaining quadruples χ (that are not disposed of by these chargings) we re-introduce the corresponding outer points r, u of their special crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$ to our analysis. This allows us to charge such quadruples χ to especially convenient topological configurations, referred to as *terminal quadruples*.

Informally, each terminal quadruple is formed by an edge, say $e = pq$, and by a pair of points that cross e in *opposite* directions (i.e., one of them crosses e from L^-_{pq} to L^+_{pq}, and the other crosses e from L^+_{pq} to L^-_{pq}). The analysis of these configurations is delegated to Section 7, where we directly bound their number in terms of simpler quantities, introduced in Section 2 (and thereby complete the proof of Theorem 2.1). To do so, we show that, for "most" terminal quadruples (if their number is at least superquadratic), some three of their four points perform two Delaunay crossings, again allowing us to use Lemma 4.5, to obtain a quadratic bound on their number.

The emergence of terminal quadruples can be attributed to the following interplay between special quadruples and their outer points. Fix a consecutive (counterclockwise) special quadruple $\chi = (a, p, w, q)$, as above. Recall that the four points of χ are co-circular at some times $\xi_0 \in \mathcal{I}_r \setminus \mathcal{J}_u$ and $\xi_1 \in \mathcal{J}_u \setminus \mathcal{J}_u$. Assume, with no loss of generality, that the co-circularity at time ξ_0 is the first co-circularity of a, p, w, q, and has index 1. (A similar assumption was made for ordinary quadruples in Section 5.) At Stage 1 we enforce upon such special quadruples χ a suitable counterclockwise variant of condition (Q3), according to which the edge qw undergoes a Delaunay crossing by p (where it crosses wq from L^+_{wq} to L^-_{wq}). Recall, however, that the underlying family $\mathcal{F}$ includes the ordinary quadruple (w, q, a, u), so the reversely-oriented copy wq of qw undergoes a Delaunay crossing by the outer point u (which then crosses L_{wq} from L^-_{wq} to L^+_{wq}). This makes (w, q, u, p) an obvious candidate for a terminal quadruple that χ can charge. A symmetric behaviour occurs if the co-circularity at time ξ_1 has index 3.

6.1 Stage 0: Charging events in $\mathcal{A}_{qa}$

Fix a consecutive special quadruple $\chi = (a, p, w, q)$, whose two special (a, q)-crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$, with respective outer points r and u, correspond to quadruples (p, q, a, r) and (w, q, a, u) in the underlying refined family $\mathcal{F}$. See Figure 41. Recall that, according to Proposition 6.1, χ is uniquely determined by each of the ordered triples (p, a, q), (w, a, q), which perform its two special (a, q)-crossings $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1]$ and $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$. Our goal is to extend the almost-Delaunayhood of qa to the possible gap $[\lambda_1, \lambda_2]$ between $\mathcal{I}_r$ and $\mathcal{J}_u$. To do so, we fix a suitable constant

k and apply Theorem 2.2 in $\mathcal{A}_{qa}$ over the interval (λ_1, λ_3), which covers the aforementioned gap (if it exists). Notice that the edge qa is not necessarily Delaunay (in $\mathrm{DT}(P)$) at times λ_1 and λ_3, so we apply this theorem with respect to the smaller set $P \setminus \{r\}$ (where, by Lemma 4.1, $\mathrm{DT}(P \setminus \{r\})$ clearly contains qa at time λ_1.

If at least one of the Conditions (i) or (ii) of Theorem 2.2 holds, we can charge χ either to a k-shallow collinearity, or to $\Omega(k^2)$ k-shallow co-circularities, which are encountered in the reduced red-blue arrangement $\mathcal{A}_{qa}^{(r)}$, defined over $P \setminus \{r\}$, during (λ_1, λ_3). (Each of these events is $(k+1)$-shallow in in $\mathcal{A}_{qa}$ when defined over the entire set P.) It suffices to check that each $(k+1)$-shallow collinearity or co-circularity, that occurs in the larger arrangement $\mathcal{A}_{qa}$ at some time $t^* \in (\lambda_1, \lambda_3)$, is charged by at most $O(1)$ special quadruples χ. Indeed, the points q and a of χ can be guessed in at most $O(1)$ ways among the three or four points involved in the shallow event. Furthermore, no counterclockwise special (a, q)-crossings $(p'a, q, \mathcal{I}_{r'})$ end in (λ_1, λ_3), so $(pa, q, \mathcal{I}_r)$ is the last such (a, q)-crossing to end before time t^*. This gives us the third point p, and Proposition 6.1 then completes the proof of the claim. To conclude, the Clarkson-Shor probabilistic argument implies that the above scenario happens for at most

$$O\left(k^2 N(n/k) + kn^2\beta(n)\right)$$

special quadruples χ.

Now suppose that Condition (iii) of Theorem 2.2 is satisfied. Then there is a subset A of at most $3k$ points (not including r) such that the edge qa belongs to $\mathrm{DT}(P \setminus (A \cup \{r\}))$ throughout the interval $\mathcal{I}_r \cup [\lambda_1, \lambda_3] = [\lambda_0, \lambda_3]$.

To proceed, we consider a random subset R of $\lceil n/k \rceil$ points of P. Let $\mathcal{F}_R$ denote the induced family of surviving (regular) Delaunay quadruples. Namely, a (regular) quadruple σ in $\mathcal{F}$ yields a counterpart in $\mathcal{F}_R$ if and only if R includes the four points of σ. As is easy to check, $\mathcal{F}_R$ is also refined with respect to its underlying set R. Furthermore, it can be viewed as a subset of $\mathcal{F}$, because each of its quadruples has a (unique) ancestor in $\mathcal{F}$. Therefore, $\mathcal{F}_R$ yields no new Delaunay crossings, whose counterparts did not arise already in the context of $\mathcal{F}$.

Note that the following two events occur simultaneously with probability at least $\Omega(1/k^6)$: (1) R includes the six points of P_χ, and (2) none of the points of $A \setminus P_\chi$ belongs to R.

Assume that the sample R is indeed successful (for the chosen special quadruple χ). Then the family $\mathcal{F}_R$ still contains the quadruples (p, q, a, r) and (w, q, a, u). Hence, $\mathcal{F}_R$ still yields the special crossings of pa and wa by q (with the same outer points r and u). We continue to denote these crossings by $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$ but observe that the corresponding intervals $\mathcal{I}_r = [\lambda_0, \lambda_1]$ and $\mathcal{J}_u = [\lambda_2, \lambda_3]$ may shrink in the process. (See Section 5.2 for more details.) Therefore, R and $\mathcal{F}_R$ also yield the (counterclockwise) special quadruple (a, p, w, q), which we continue to denote by χ. Furthermore, χ is again consecutive with respect to R and $\mathcal{F}_R$ (because the underlying family $\mathcal{F}_R$ induces no new special crossings, which did not arise in the context of $\mathcal{F}$). Moreover, since R' contains none of the points r, u, the edge qa now belongs to $\mathrm{DT}(R \setminus \{r, u, p, w\})$ throughout the extended interval $[\mathcal{I}_r, \mathcal{J}_u] = [\lambda_0, \lambda_3]$; see Figure 42.

Definition. Let P be a (finite) set of moving points, and let $\mathcal{F}$ be a refined family constructed over P. We say that a consecutive special quadruple $\chi = (a, p, w, q)$, formed by counterclockwise special (a, q)-crossings $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ and $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$ (both of them in $\mathcal{F}$) is *Delaunay* (again, with respect to P and $\mathcal{F}$), if its edge aq belongs to $\mathrm{DT}(P \setminus \{p, w, r, u\})$ throughout the extended interval $[\mathcal{I}_r, \mathcal{J}_u] = [\lambda_0, \lambda_3]$.

Let $\Phi_0(m)$ denote the maximum number of consecutive Delaunay special quadruples that can be induced by a refined family $\mathcal{F}$ defined over n moving points. The preceding discussion implies the following recurrence.

$$\Phi(n) \leq O\left(k^6 \Phi_0(n/k) + k^2 N(n/k) + kn^2\beta(n)\right), \tag{7}$$

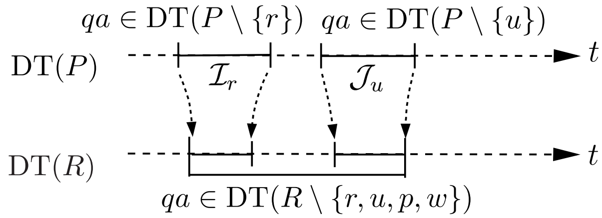

Figure 42: After replacing the underlying set P by its subsample R, the edge qa belongs to $\mathrm{DT}(R \setminus \{r, u, p, w\})$ throughout $[\mathcal{I}_r, \mathcal{J}_u] = \mathsf{conv}(\mathcal{I}_r \cup \mathcal{J}_u)$, including the gap between $\mathcal{I}_r$ and $\mathcal{J}_u$. (The intervals $\mathcal{I}_r$ and $\mathcal{J}_u$ may shrink in the process.)

for any constant parameter $k \geq 12$.

6.2 Stage 1

To bound the above quantity $\Phi_0(n)$, we fix an underlying set P of n moving points, a refined family $\mathcal{F}$, and a consecutive Delaunay special quadruple $\chi = (a, p, w, q)$, obtained from the corresponding special crossings $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ and $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$; r and u are the corresponding outer points. See Figure 41. By definition, the edge qa belongs to $\mathrm{DT}(P \setminus \{p, w, r, u\})$ throughout the interval $[\lambda_0, \lambda_3]$.

As in Section 5, we fix constants $12 < k \ll \ell$ and distinguish between five possible scenarios, where the roles of the edges pq and wq are mostly symmetric. In all but the last case, we will be able to bound the number of (the relevant) Delaunay special quadruples in terms of quantities that were already introduced in Section 2. In the last case (case (e)), our bound will also depend on the number of special quadruples of a more restricted type, which are defined over an appropriate subsample of R of P. Such quadruples will be called 1-*restricted*, and their analysis will be passed on to the subsequent stages.

Case (a). The edge qa is hit during $[\lambda_0, \lambda_3]$ by at least one of the points p, w. Clearly, this collinearity can happen only during the gap between $\mathcal{I}_r$ and $\mathcal{J}_u$ (if it exists).

If qa is hit by p then the triple p, a, q defines two distinct (single) Delaunay crossings within the smaller triangulation $\mathrm{DT}(P \setminus \{w, r, u\})$. (Here we exploit the fact that the crossed edge qa is almost Delaunay throughout $[\lambda_0, \lambda_3]$.) According to Lemma 4.5, combined with the Clarkson-Shor argument, where we use a sample of size $n/2$, the overall number of such triples (p, q, a) (and, hence, of such special quadruples $\chi = (a, p, w, q)$, each of them uniquely determined by its corresponding triple (p, q, a)) is at most $O(n^2)$.

If qa is hit by w then we similarly argue that the triple (w, a, q) defines two distinct Delaunay crossings within $\mathrm{DT}(P \setminus \{p, r, u\})$, so the number of such special quadruples χ (each of them uniquely determined by the corresponding triple (w, q, a)) is at most $O(n^2)$ too.

Case (b). At least k clockwise special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$ end in $(\lambda_1, \lambda_3]$, or at least k clockwise special (w, q)-crossings $(wa', q, \mathcal{J}_{u'})$ begin in $[\lambda_0, \lambda_2)$; each of these crossings comes with its respective outer point r' or u'.

Without loss of genarality, we consider only the former scenario, and handle the latter one in a fully symmetric manner. Recall that a special (p, q)-crossing $(pa', q, \mathcal{I}_{r'})$ is uniquely determined by each of the triples (p, a', q) and (p, q, r'). Hence, at most one of these special crossings has a' equal to u. Moreover, the preliminary pruning (applied to clockwise special (p, q)-crossings) guarantees that none of them can have $r' = a$ or $a' = r$.

We apply Theorem 5.3, in combination with the standard argument of Clarkson and Shor, in order to dispose of such special quadruples χ. To do so, we consider a random subset R of $\lceil n/4 \rceil$ points of P and notice that the following two conditions hold simultaneously with probability $\Omega(1)$: (1) R

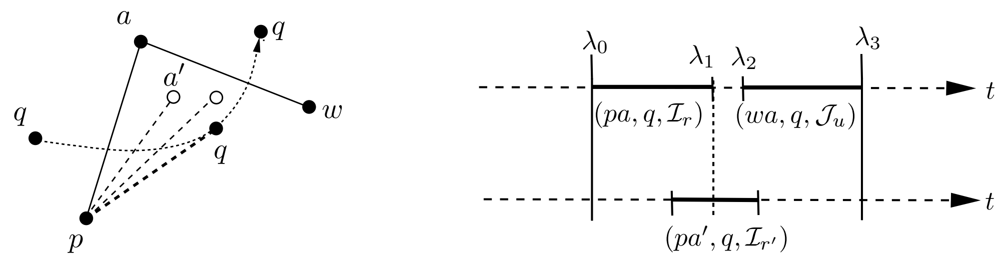

Figure 43: Case (b): at least k clockwise special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$ end in $(\lambda_1, \lambda_3]$.

includes the points p, q and a, but none of r, u, and (2) for at least a constant fraction of the above special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$, the set R includes the point a' but not r'.

Specifically, (1) holds with some constant probability close to $(1/4)^3(3/4)^2$. Concerning (2), assume without loss of generality that the number of relevant crossings $(pq', q, \mathcal{I}_{r'})$ is exactly k (so at least $k-1$ of them satisfy $a' \neq u$). Then, conditioning on the success of (1), the expected number of these crossings that satisfy the property in (2) is very close to $(k-1)(3/16)$, or larger. Hence, Markov's inequality implies that, with an appropriate choice of parameters, the probability of (2), conditioned on the success of (1), is also some fixed constant. Hence, the probability that both (1) and (2) hold is also $\Omega(1)$, as claimed.

If the sample R is successful (for the given χ), then it clearly yields an (ordinary) Delaunay crossing $(pa, q, \mathcal{I})$, whose respective interval $\mathcal{I}$ is contained in $[\lambda_0, \lambda_1]$ (as $R \subseteq P \setminus \{r, u\}$). It remains to check that this crossing is $(a, q, \Theta(k))$-chargeable, with respect to the interval $[\lambda_0, \lambda_3]$.

To see the latter property, note that each of the above special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$, for which the sample R includes a' but not r', yields the Delaunay crossing $(pa', q, \mathcal{I}')$ in R, with $\mathcal{I}' \subseteq \mathcal{I}_{r'}$. Therefore, Lemma 4.6 implies that $(pa', q, \mathcal{I}')$ occurs within $[\lambda_0, \lambda_1] \cup \mathcal{I}_{r'} \subseteq [\lambda_0, \lambda_3]$. Moreover, aq belongs to $\mathrm{DT}(R)$ at times λ_0 and λ_3 (in addition to its almost-Delaunayhood in $\mathrm{DT}(R)$, with only two points p, w removed, during $[\lambda_0, \lambda_3]$).

Theorem 5.3 implies, then, that the overall number of such triples (p, q, a) in R is only

$$O\left(k^2 N(n/k) + kn^2\beta(n)\right).$$

Clearly, this also bounds the overall number of Delaunay special quadruples χ falling into case (b).

To conclude, we can assume, from now on, that case (b) does not occur. That is, fewer than k clockwise special (p, q)-crossings end in $(\lambda_1, \lambda_3]$, and fewer than k clockwise special (w, q)-crossings begin in the symmetric interval $[\lambda_0, \lambda_2)$.

Case (c). No clockwise special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$, with $r' \notin \{w, u\}$, end during $[\lambda_3, \infty)$, or no clockwise special (w, q)-crossings $(wa', q, \mathcal{J}_{u'})$, with $u' \notin \{p, r\}$, begin during $(-\infty, \lambda_0]$.

Without loss of generality, we consider only the first subcase and handle the other one in a fully symmetric manner. Note that the preliminary pruning (combined with the fact that $(pa, q, \mathcal{I}_r)$ is uniquely determined by the triple (p, q, r)) guarantees that no clockwise special (p, q)-crossing $(pa', q, \mathcal{I}_{r'})$ can have r' in $\{r, a\}$.

Since case (b) does not occur, $(pa, q, \mathcal{I}_r)$ is among the $k+3$ last clockwise special (p, q)-crossings (in the standard order provided by Lemma 5.5). Indeed, at most k such special crossings $(pa', q, \mathcal{I}_{r'})$ end in $(\lambda_1, \lambda_3]$, and at most two of them can end after λ_3, namely those whose outer point is either u or w (recalling that this outer point, together with p, q, uniquely determines the crossing). Therefore, we can charge $(pa, q, \mathcal{I}_r)$ and χ to the edge pq, so this situation happens for at most $O(kn^2)$ special quadruples χ.

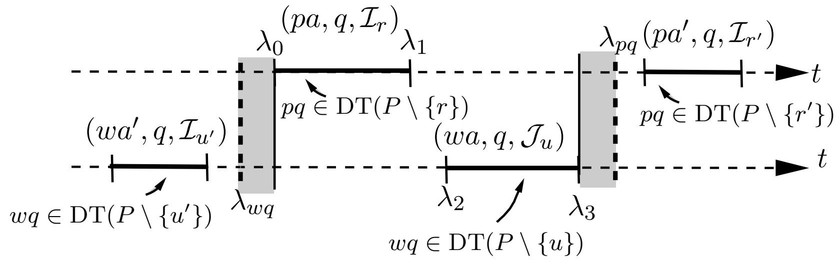

Figure 44: Assuming (c) does not hold, we put λ_{pq} to be the first time in $[\lambda_3, \infty)$ when pq belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, for $r' \notin P_\chi$. Similarly, we put λ_{wq} to be the last time in $(-\infty, \lambda_0]$ when wq belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{u'\})$, for $u' \notin P_\chi$.

Preparing for cases (d) and (e). For the remainder of this stage, we assume that none of the cases (a), (b) or (c) occurs. In particular, there is a special (p, q)-crossing $(pa', q, \mathcal{I}_{r'})$, whose outer point r' satisfies $r' \notin \{w, u\}$, that ends after λ_3. (Refer to Figure 44.) Therefore, and according to Lemma 4.1, pq belongs to $\mathrm{DT}(P \setminus \{r'\})$ either at time λ_3 or at some later time. Moreover, r' does not belong to P_χ because, after the preliminary pruning, there remain no clockwise special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$ with $r' \in \{a, r\}$. Let λ_{pq} be the first time in $[\lambda_3, \infty)$ when pq belongs to some triangulation $\mathrm{DT}(P \setminus \{r'\})$, for some $r' \notin P_\chi$. More precisely, we put $\lambda_{pq} = \lambda_3$ if pq belongs to such a triangulation at time λ_3, and otherwise we set λ_{pq} to be the first time after λ_3 when pq enters $\mathrm{DT}(P \setminus \{r'\})$ (for some $r' \notin P_\chi$).

A symmetric argument (adapted for clockwise special (w, q)-crossings) shows that there is a special (w, q)-crossing $(wa', q, \mathcal{J}_{u'})$, with an outer point $u' \notin \{a, r\}$, that begins before λ_0 (so $wq \in \mathrm{DT}(P \setminus \{u'\})$ at some time before or at λ_0). We define λ_{wq} to be the last time in $(-\infty, \lambda_0]$ when the edge wq belongs to some triangulation $\mathrm{DT}(P \setminus \{u'\})$, for some $u' \notin P_\chi$. In what follows, we use r' and u' to denote a fixed[29] pair of points, outside P_χ, whose removal restores the Delaunayhood of pq and wq at respective times λ_{pq} and λ_{wq}, and for which λ_{pq} is smallest and λ_{wq} is largest.

Before proceeding to the cases (d) and (e), we first apply Theorem 2.2 in $\mathcal{A}_{pq}$ over the interval $(\lambda_1, \lambda_{pq})$, and then apply it in $\mathcal{A}_{wq}$ over the symmetric interval $(\lambda_{wq}, \lambda_2)$, both times with the second constant $\ell \gg k$.

Consider the first application of Theorem 2.2. It is performed with respect to the *reduced* triangulation $\mathrm{DT}(P \setminus \{r, r'\})$, which contains pq at time λ_{pq}. If (at least) one of the first two conditions of Theorem 2.2 holds, we charge χ, via $(pa, q, \mathcal{I}_r)$, either to $\Omega(\ell^2)$ $(\ell + 2)$-shallow co-circularities, or to an $(\ell+2)$-shallow collinearity. (Each of these events is ℓ-shallow with respect to $P \setminus \{r, r'\}$.) As before, the crucial observation is that each co-circularity or collinearity, which occurs at some time $t^* \in (\lambda_1, \lambda_{pq})$, is charged in the above manner by at most $O(k)$ special quadruples χ. Indeed, the points p and q of χ can be chosen in $O(1)$ ways among the three or four points involved in the event. Furthermore, recall that χ is uniquely determined by the triple (a, p, q), so it suffices to guess a (for the chosen p, q and t^*).

Since case (b) has been ruled out, at most k clockwise special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$ end in (λ_1, λ_3). Moreover, assuming $\lambda_{pq} > \lambda_3$, no such crossing can end in $(\lambda_3, \lambda_{pq}]$ unless its respective outer point r' belongs to $\{w, u\}$ (which happens for at most two special (p, q)-crossings). Therefore, $(pa, q, \mathcal{I}_r)$ is among the last $k + 3$ clockwise special (p, q)-crossings to end before time t^*.

To conclude, the above charging accounts for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ special quadruples χ.

Finally, if Condition (iii) of Theorem 2.2 holds, then the Delaunayhood of pq can be restored throughout the interval $[\lambda_1, \lambda_{pq}]$ by removing a subset A of at most $3\ell + 2$ points of P (including r and r'); see Figure 45.

The second application of Theorem 2.2 in $\mathcal{A}_{wq}$ is fully symmetric, and it is done with respect to the set $P \setminus \{u, u'\}$ in the interval $(\lambda_{wq}, \lambda_2)$. If at least one of the Conditions (i), (ii) is satisfied, we dispose of

[29]Notice that we do not claim that the choice r' and u' is unique.

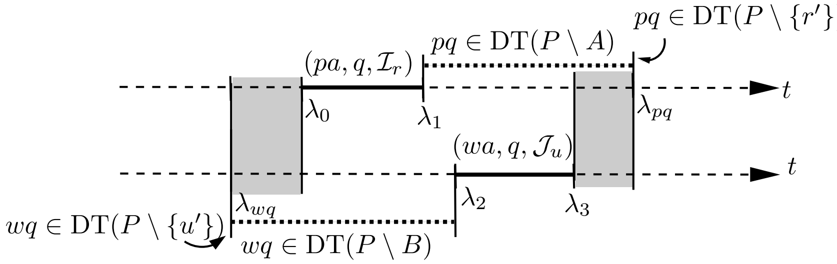

Figure 45: Extending the almost-Delaunayhood of pq and wq, in preparation for cases (d) and (e), respectively, from $\mathcal{I}_r = [\lambda_0, \lambda_1]$ to $[\lambda_0, \lambda_{pq}]$, and from $\mathcal{J}_u = [\lambda_2, \lambda_3]$ to $[\lambda_{wq}, \lambda_3]$.

χ by charging it (via $(wa, q, \mathcal{J}_u)$) to $(\ell+2)$-shallow collinearities and co-circularities that occur in $\mathcal{A}_{wq}$ during $(\lambda_{wq}, \lambda_2)$. (Since case (b) has been ruled out, $(wa, q, \mathcal{J}_u)$ is among the first $k+3$ special counterclockwise (w, q)-crossings to begin after each charged event. Hence, every collinearity or co-circularity is charged at most $O(k)$ times.) As before, this accounts for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ special quadruples χ.

For each of the remaining special quadruples we have a set B of at most $3\ell+2$ points (including u and u') whose removal restores the Delaunayhood of wq throughout $[\lambda_{wq}, \lambda_2]$; see Figure 45 again.

To recap, in each of the remaining cases (d) and (e), we may assume the existence of the first time $\lambda_{pq} \geq \lambda_3$ when pq belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, and of the symmetric last time $\lambda_{wq} \leq \lambda_0$ when wq belongs to a similarly reduced triangulation $\mathrm{DT}(P \setminus \{u'\})$, where u' and r' are fixed points outside P_χ. In addition, there exist sets A (including r and r') and B (including u and u'), both of cardinality at most $3\ell+2$, whose removal restores the Delaunayhood of pq and wq throughout the respective intervals $[\lambda_1, \lambda_{pq}]$ and $[\lambda_{wq}, \lambda_2]$ (and, therefore, extends the almost-Delaunayhood of these edges to the respective larger intervals $[\lambda_0, \lambda_{pq}] = I_r \cup [\lambda_1, \lambda_{pq}]$ and $[\lambda_{wq}, \lambda_3] = [\lambda_{wq}, \lambda_2] \cup \mathcal{J}_u$).

Case (d). The point a hits the edge pq during $[\lambda_1, \lambda_{pq}]$, or it hits the edge wq during the symmetric interval $[\lambda_{wq}, \lambda_2]$.

In the former scenario, the triple (a, p, q) defines two Delaunay crossings within $\mathrm{DT}((P \setminus A) \cup \{p\})$, and, in the latter, the symmetric triple (a, w, q) defines two Delaunay crossings within $\mathrm{DT}((P \setminus B) \cup \{w\})$. In both cases, we can use Lemma 4.5, in combination with the sampling argument of Clarkson and Shor, to show that the overall number of the relevant triples in P is at most $O(\ell n^2)$. As in case (a), this also bounds the overall number of such special quadruples $\chi = (a, p, w, q)$.

Case (e). None of the previous cases (a)–(d) occurs, and none of the preliminary charging arguments apply to χ.

In particular, since cases (a) and (d) have been ruled out, either the point q either remains in L^+_{pa} after the end λ_1 of $\mathcal{I}_r$ and until crossing wa (during $\mathcal{J}_u$), or else it re-enters L^-_{pa} during that period, through the relative interior of pa. Similarly, q must remain in L^-_{wa} after crossing pa (during $\mathcal{I}_r$) and until the beginning λ_2 of $\mathcal{J}_u$, unless it crosses wa (from L^+_{wa} to L^-_{wa}) during that period.

In addition, we assume the existence of the sets A and B, as above, whose removal restores the Delaunayhood of pq and wq throughout the respective intervals $[\lambda_{wq}, \lambda_3]$ and $[\lambda_0, \lambda_{pq}]$.

Recall that, according to Lemma 4.4, the four points of χ are co-circular at times $\xi_0 \in \mathcal{I}_r \setminus \mathcal{J}_u$ and $\xi_1 \in \mathcal{J}_u \setminus \mathcal{I}_r$ (see, e.g., Figure 40). Clearly, at least one of these co-circularities is extremal. We therefore distinguish between two subcases (whose treatment remains fully symmetric untill the beginning of Stage 4).

Case (e1): The co-circularity at time ξ_1 has index 3. In this case, we say that χ is a *right special quadruple*. We claim that in this case the edge pq is hit during $(\lambda_1, \lambda_{pq}]$ by the point w, which crosses it from L^+_{pq} to L^-_{pq}. To show this, we distinguish between two sub-scenarios.

(i) If p lies in L_{wa}^- when q enters the opposite halfplane L_{wa}^+ (during $\mathcal{J}_u$), then the Delaunayhood of pq is violated, right after time ξ_1, by $w \in L_{pq}^+$ and $a \in L_{pq}^-$. See Figure 46 (left). Hence, pq is hit by at least one of these two points during $(\xi_1, \lambda_{pq}] \subseteq (\lambda_2, \lambda_{pq}]$, as prescribed in cases (i) and (ii) of Lemma 3.1 (case (iii) thereof cannot arise since ξ_1 has index 3). Since case (d) has been ruled out, a cannot hit pq during $(\lambda_0, \lambda_{pq}]$. Hence, pq must be hit by w, which then crosses it from L_{pq}^+ to L_{pq}^- (this crossing direction is also prescribed by the lemma).

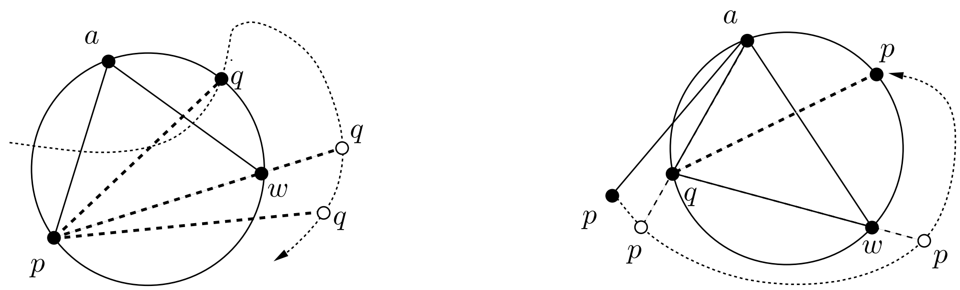

Figure 46: Case (e1): ξ_1 is the last co-circularity of a, p, w, q. Arguing that the edge qp is crossed by w during $(\lambda_1, \lambda_{pq}]$. Left: A possible motion of q if $p \in L_{wa}^-$ when q crosses wa (during $\mathcal{J}_u$). Right: A possible motion of p (after $\mathcal{I}_r$) if q re-enters L_{pa}^- through pa.

(ii) If p lies in L_{wa}^+ when q enters this halfplane, then q must re-enter L_{pa}^- after $\mathcal{I}_r$ and before it reaches L_{wa}^+. Hence, the co-circularity at time ξ_1 is as depicted in Figure 46 (right); that is, it occurs with $p \in L_{wa}^+$ and $q \in L_{wa}^-$. Since none of the preceding cases (a), (d) holds, q can re-enter L_{pa}^- during this interval only through the edge pa. Therefore, the counterclockwise variant of Lemma 5.1 (adapted for special quadruples, as described in the introduction to this section) implies that in this case too w crosses pq from L_{pq}^+ to L_{pq}^-, during $(\lambda_1, \lambda_3] \subseteq (\lambda_1, \lambda_{pq}]$; see Figure 46 (right). (As a matter of fact, this collinearity must occur during (λ_1, ξ_1).)

To conclude, in both sub-scenarios the edge qp undergoes a Delaunay crossing by w within the smaller triangulation $\mathrm{DT}((P \setminus A) \cup \{w\})$, and the respective interval $\mathcal{H} = [\lambda_4, \lambda_5]$ of that crossing is contained in $[\lambda_1, \lambda_{pq}]$. (We again emphasize that A includes both points $r, r' \neq w$, so the edge pq belongs to $\mathrm{DT}((P \setminus A) \cup \{w\})$ throughout $\mathcal{I}_r = [\lambda_0, \lambda_1]$ and at time λ_{pq}.)

If w hits pq *twice* during $(\lambda_1, \lambda_{pq}]$, then pq undergoes within $\mathrm{DT}((P \setminus A) \cup \{w\})$ either two (single) Delaunay crossings, or a double Delaunay crossing, by the same point w. We thus charge χ to the respective triple (p, q, w) and use Theorem 4.3 or Lemma 4.5, in combination with the probabilistic argument of Clarkson and Shor, to show that the overall number of such triples (p, q, w) is at most $O(\ell n^2)$. Since case (b) does not occur, $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ is among the last $k+3$ special clockwise (p, q)-crossings to end before the above crossings of pq by w. (Namely, at most k such (p, q)-crossings end during $(\lambda_1, \lambda_3]$, and at most two of them can end in $(\lambda_3, \lambda_{pq}]$, if $\lambda_3 \neq \lambda_{pq}$; see the analysis preceding case (e) for more details.) In particular, any triple (p, q, w) is shared by at most $k+3$ charging quadruples χ. Hence, the above additional collinearities of p, q, w are encountered for at most $O(k\ell n^2)$ special quadruples.

A similar argument applies if the edge wq is hit by p during[30] $[\lambda_{wq}, \lambda_2)$. In this case, the triple (p, q, w) performs two distinct single Delaunay crossings within the triangulation $\mathrm{DT}((P \setminus (A \cup B)) \cup \{w, p\})$ (namely, the crossing of qp by w, and the crossing of wq by p). The same bound $O(k\ell n^2)$ holds in this case too.

We thus assume, from now on, that w hits pq exactly once during $(\lambda_1, \lambda_{pq}]$, and that p does not cross wq during the symmetric interval $[\lambda_{wq}, \lambda_2)$. In particular, this implies that q lies in L_{wa}^- when it enters

[30] Since $p \neq u$, the point p cannot cross wq during $\mathcal{J}_u = [\lambda_2, \lambda_3]$, as wq belongs to $\mathrm{DT}(P \setminus \{u\})$ during that interval.

L_{pa}^{+} during $\mathcal{I}_r$. Indeed, otherwise q would have to cross L_{wa} (thereby leaving L_{wa}^{+}) between the times when it enters the halfplanes L_{pa}^{+} and L_{wa}^{+} (both times during the respective special crossings). Since neither of the cases (a), (d) holds, q can cross L_{wa}, for the first time, only within wa. However, in this latter case the counterclockwise variant Lemma 5.1 would imply that p hits wq during $[\lambda_{wq}, \lambda_2)$ (which has been ruled in the previous paragraph).

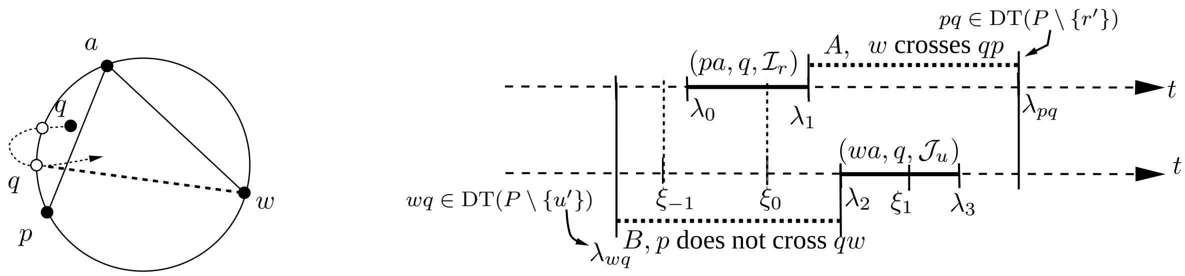

Figure 47: Case (e1). Left: A possible motion of q before and during $\mathcal{I}_r$. The Delaunayhood of wq is violated, right before ξ_0, by p and a. The points of χ are involved, at some time $\xi_{-1} \in [\lambda_{wq}, \xi_0)$ in another co-circularity (of index 1). The order type of χ remains fixed throughout $[\xi_{-1}, \xi_0]$. Right: A schematic summary of what we eventually assume at the end of case (e1).

We may therefore assume that w lies in $L_{pa}^{+} = L_{qa}^{+}$ when q crosses pa (during $\mathcal{I}_r$). See Figure 47 (left). Arguing as in the previous similar situations, we conclude that the Delaunayhood of wq is violated, right before ξ_0, by $p \in L_{wq}^{-}$ and $a \in L_{wq}^{+}$. (That is, w enters the cap $B[p, q, a] \cap L_{pa}^{+}$ at time ξ_0.) By Lemma 3.1 (applied with respect to $\mathrm{DT}(P \setminus \{u, u'\})$), and since none of the points a, p is allowed to cross wq during $[\lambda_{wq}, \lambda_2]$, the four points p, q, a, w must be co-circular at some time $\xi_{-1} \in [\lambda_{wq}, \xi_0)$, right before which the Delaunayhood of pa is violated by $q \in L_{pa}^{-}$ and $w \in L_{pa}^{+}$. (We must have $\lambda_{wq} \leq \xi_{-1} < \xi_0 < \lambda_2 < \xi_1$.) Moreover, wq is intersected by pa throughout $[\xi_{-1}, \xi_0]$. (In other words, the order type of a, p, w, q remains fixed there.)

A schematic summary of what we assume in case (e1) (by the end of its analysis) is given in Figure 47 (right).

Case (e2): The co-circularity at time ξ_0 has index 1. In this case, we say that χ is a *left special quadruple*. We apply a fully symmetric topological analysis (in which we switch the roles of pq and wq, and reverse the direction of the time axis).

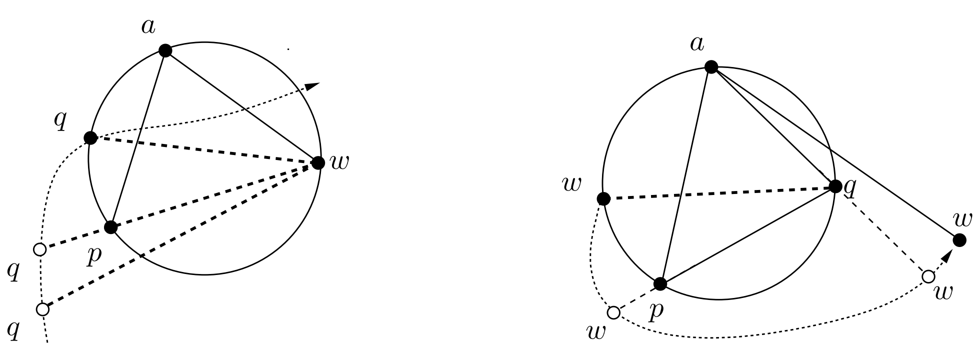

Figure 48: Case (e2): ξ_0 is the *first* co-circularity of a, p, w, q. Arguing that p hits qw in $[\lambda_{wq}, \lambda_2)$. Left: A possible motion of q if w lies in $L_{pa}^{+} = L_{qa}^{+}$ when q hits pa (during $\mathcal{I}_r$). Right: A possible motion of w if q hits wa also before $\mathcal{J}_u$ (and after its hits pa in $\mathcal{I}_r$).

Briefly, we use one of the Lemmas 3.1 or 5.1 to show that p crosses wq, from L_{wq}^{+} to L_{wq}^{-}, during the interval $[\lambda_{wq}, \lambda_2]$. As in case (e1), we distinguish between two possible scenarios, now depending on the location of w when q crosses pa (during $\mathcal{I}_r$).

(i) If w lies in L^+_{pa} when q crosses pa during $\mathcal{I}_r$ then the Delaunayhood of wq is violated, right before ξ_0, by $p \in L^-_{wq}$ and $a \in L^+_{wq}$. Hence, the promised crossing follows from the *time-reversed* variant of Lemma 3.1 (and because case (d) has been ruled out); see Figure 48 (left).

We again emphasize that, in this subscenario of case (e2), the crossing of wq by p occurs after λ_{wq} and *before* ξ_0. (Note that Figure 48 (left) depicts a possible trajectory of q in the standard time direction. In the time-reversed application of Lemma 3.1, the point q moves *backwards*, so p crosses wq from L^-_{wq} to L^+_{wq}. In the standard time direction, the crossing is from L^+_{wq} to L^-_{wq}, as asserted.)

(ii) If w lies in $L^-_{pa} = L^-_{pq}$ (i.e., q and p lie in L^+_{wa}) when q crosses pa, then q will have to enter L^-_{wa} before the beginning of $\mathcal{J}_u$ (and only through the interior of wa, as cases (a) and (d) have been ruled out). Therefore, the asserted crossing of wq by p now follows from a suitable (counterclockwise and time-reversed) variant of Lemma 5.1; see Figure 48 (right).

(Again, Figure 48 (right) depicts a possible trajectory of w in the standard time direction. In the time-reversed application of Lemma 5.1, the roles of p and w in the statement of the lemma are switched, and p crosses wq in the opposite direction, from L^-_{wq} to L^+_{wq}.)

If p hits wq twice during $[\lambda_{wq}, \lambda_2)$, or if w hits pq during $(\lambda_1, \lambda_{pq}]$, then we can dispose of χ using Theorem 4.3 or Lemma 4.5. Namely, we then argue that the triple (p, w, q) is involved within $\mathrm{DT}((P \setminus (B \cup A)) \cup \{q, w\})$ either in two distinct single Delaunay crossings, or in one double Delaunay crossing. Hence, the overall number of such triples in P is at most $O(\ell n^2)$. Furthermore, any triple (p, w, q) is shared by at most $k + 3$ special quadruples χ (namely, such special quadruples $\chi = (a, p, w, q)$ whose second crossings $(wa, q, \mathcal{J}_u)$ are among the first $k + 3$ clockwise special (w, q)-crossings to begin after p crosses qw from L^-_{qw} to L^+_{qw}); see case (e1) for a fully symmetric argument.

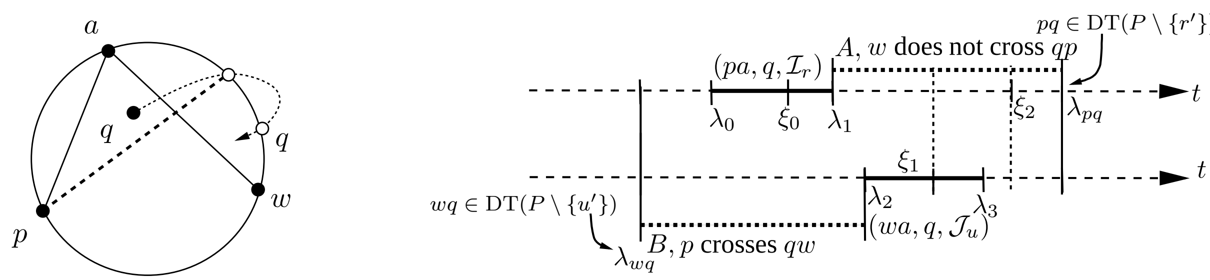

Figure 49: Case (e2). Left: A possible motion of q during $\mathcal{J}_u$, and afterwards. The Delaunayhood of pq is violated, right after ξ_1, by a and w. The points of χ are involved, at some time $\xi_2 \in (\xi_1, \lambda_{pq}]$ in another co-circularity (of index 3). The order type of χ remains fixed throughout $[\xi_1, \xi_2]$. Right: A schematic summary of what we eventually assume at the end of case (e2).

To conclude, we may assume that p hits wq only once during $[\lambda_{wq}, \lambda_2)$ (crossing it from L^+_{wq} to L^-_{wq}), and that w does not cross pq during $[\lambda_1, \lambda_{pq}]$. Lemma 3.1 then implies that the points of χ are co-circular at some time $\xi_2 \in (\lambda_1, \lambda_{pq}]$, and that pq is intersected by aw throughout $[\xi_1, \xi_2]$; see Figure 49 (left). A schematic summary of what we assume by the end of case (e2) is given in Figure 49 (right).

Probabilistic refinement. For each clockwise special (p, q)-crossing $(pa', q, \mathcal{I}_{r'})$ that ends during $(\lambda_1, \lambda_{pq})$ we add the corresponding point a' to the obstruction set A of pq. Similarly, for each clockwise special (w, q)-crossing $(wa', q, \mathcal{J}_{u'})$ that begins during $(\lambda_{wq}, \lambda_2)$ we add the point a' to the obstruction set B of wq. As in Section 5, this is done in order to dispose of the corresponding special (p, q)- and (w, q)-crossings. Since we add at most $k + 2$ elements to each set, and since $k \ll \ell$, each of the sets A, B still contains at most 4ℓ points of P.

Consider a subset R of $\lceil n/\ell \rceil$ points chosen at random from P. Let $\mathcal{F}_R$ denote the refined family induced by $\mathcal{F}$ over R. Notice that the following two conditions hold simultaneously with probability at least $\Omega(1/\ell^6)$: (1) The 6 points of P_χ belong to R, and (2) R includes none of the points of $(A \cup B) \setminus P_\chi$.

Assume that the above sample R is indeed successful for the chosen $\chi = (a, p, w, q)$. Then the points of P_χ still yield a Delaunay consecutive special quadruple (of the same topological type, which can be either right or left) with respect to R and $\mathcal{F}_R$. We continue to denote this new quadruple as χ but note that the respective intervals $\mathcal{I}_r$ and $\mathcal{J}_u$ of the special crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$ may shrink as we pass from $\mathrm{DT}(P)$ to $\mathrm{DT}(R)$. We next review the additional properties gained by χ in $\mathrm{DT}(R)$.

First, recall that the old time λ_{pq} (defined after case (c) in terms of P) was accompanied by a point $r' \notin P_\chi$, whose removal restored the Delaunayhood of pq at that time. Since r' is among the omitted points of A, we can redefine λ_{pq} as the first time in $[\lambda_3, \infty)$ when pq belongs to $\mathrm{DT}(R)$. Similarly, we redefine λ_{wq} as the last time in $(-\infty, \lambda_0]$ when wq belongs to $\mathrm{DT}(R)$. (In both cases, we refer to the new values of λ_0 and λ_3.) By what has just been noted, the new value of λ_{pq} (resp., of λ_{wq}) decreases (resp., increases) from its old value.

Second, the following three conditions hold with respect to R and $\mathcal{F}_R$, and with the new values of $\lambda_0, \lambda_1, \lambda_2, \lambda_3, \lambda_{pq}$ and λ_{wq} (see Figure 50 for a schematic summary):

(S1) The edge pq belongs to $\mathrm{DT}(R \setminus \{a, r, w, u\})$ throughout the interval $[\lambda_0, \lambda_{pq}]$. Furthermore, no clockwise special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$ end during $(\lambda_1, \lambda_{pq})$ (except perhaps for the special crossings of pu and pw by q).

(S2) The edge wq belongs to $\mathrm{DT}(R \setminus \{p, a, r, u\})$ throughout the interval $[\lambda_{wq}, \lambda_3]$. Furthermore, no clockwise special (w, q)-crossings $(wa', q, \mathcal{I}_{u'})$ begin during $(\lambda_{wq}, \lambda_2)$ (except perhaps for the special crossings of wr and wp by q).

(S3a) If χ is a *right* quadruple, then the set $P \setminus \{a, r, u\}$ induces a Delaunay crossing $(qp, w, \mathcal{H})$ which occurs within $(\lambda_1, \lambda_{pq}]$. Furthermore, w hits pq only once during $(\lambda_1, \lambda_{pq}]$, so this is a single Delaunay crossing. Moreover, the points of χ are co-circular at some time $\xi_{-1} \in [\lambda_{wq}, \xi_0)$, and the edge qw is violated by $a \in L^+_{qw}$ and $p \in L^-_{qw}$ throughout the interval (ξ_{-1}, ξ_0). Finally, p does not cross qw in $[\lambda_{wq}, \lambda_2]$.

(S3b) If χ is a *left* quadruple, then the set $P \setminus \{a, r, u\}$ induces a Delaunay crossing $(qw, p, \mathcal{H})$, which occurs within $[\lambda_{wq}, \lambda_2)$. Furthermore, p hits wq only once during $[\lambda_{wq}, \lambda_2)$, so this is a single Delaunay crossing. Moreover, the points of χ are co-circular at some time $\xi_2 \in (\xi_1, \lambda_{pq}]$, and the edge pq is violated by $a \in L^-_{pq}$ and $w \in L^+_{pq}$ throughout the interval (ξ_1, ξ_2). Finally, w does not cross qp in $[\lambda_1, \lambda_{pq}]$.

Definition. Assume that we are given a set P of moving points, and a refined family $\mathcal{F}$. Let $\chi = (a, p, w, q)$ be a consecutive Delaunay special quadruple that is defined with respect to $\mathcal{F}$ and P. We say that χ is 1-*restricted* if it satisfies the above three conditions (S1), (S2), and (S3a) or (S3b), where the reference sets R and $\mathcal{F}_R$ are replaced by P and $\mathcal{F}$, respectively. (We also implicitly require that the values λ_{pq} and λ_{wq}, mentioned in conditions (S1) and (S2), actually exist.)

Let $\Phi_1(m)$ denote the maximum number of 1-restricted special quadruples that can be defined over a set of n moving points (and a refined family of regular Delaunay quadruples). Then the following recurrence holds:

$$\Phi_0(n) \le O\left(\ell^6 \Phi_1(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right).$$

Proposition 6.2. *With the above assumptions, any ordered triple (p, q, w) can be shared by at most three 1-restricted special quadruples $\chi = (a, p, w, q)$ of each topological type (i.e., right or left).*

Proof. Let $\chi = (a, p, w, q)$ be a 1-restricted right special quadruple. By Conditions (S2) and (S3a), $(pa, q, \mathcal{I}_r)$ is among the 3 last special counterclockwise (p, q)-crossings to end before w enters L^+_{pq} (during $\mathcal{H}$). Hence, a is determined, up to three possible values, by the choice of (p, q, w). A fully symmetric argument applies if χ is a left special quadruple. □

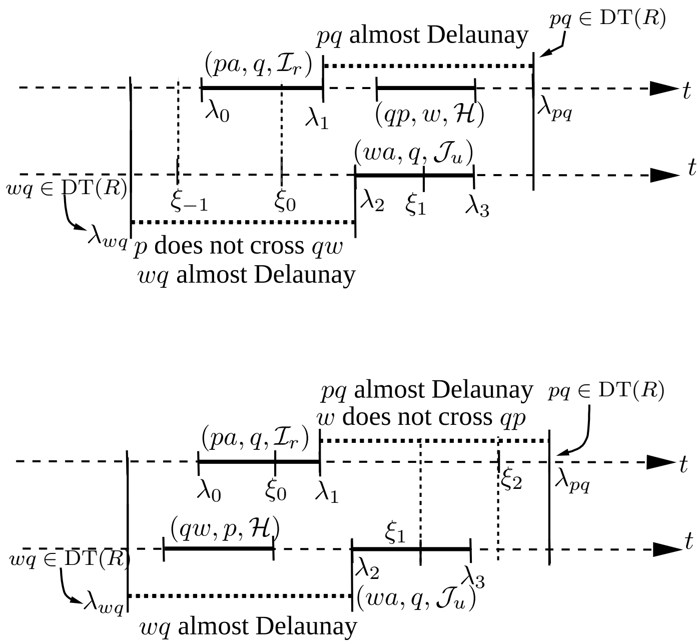

Figure 50: A schematic summary of the properties of χ within $\mathrm{DT}(R)$. The edge pq is Delaunay at time λ_{pq}, and it is almost Delaunay (with the omission of only a, r, u) throughout $[\lambda_0, \lambda_{pq}]$. The edge wq is Delaunay at time λ_{wq}, and it is almost Delaunay (with the same omission) throughout $[\lambda_{wq}, \lambda_3]$. Top: If χ is a right quadruple, then qp undergoes the crossing $(qp, w, \mathcal{H})$ within $(\lambda_1, \lambda_{pq}]$, and we encounter an additional co-circularity of a, p, w, q at some time $\xi_{-1} \in [\lambda_{wq}, \xi_0)$. Bottom: If χ is a left quadruple, then qw undergoes the crossing $(qw, p, \mathcal{H})$ within $[\lambda_{wq}, \lambda_2)$, and the additional co-circularity occurs at some time $\xi_2 \in (\xi_1, \lambda_{pq}]$ (below).

The subsequent stages — Overview. Fix a refined family $\mathcal{F}$, defined with respect to an underlying set P of n moving points. Let $\chi = (a, p, w, q)$ be a 1-refined Delaunay quadruple, consistent with P and $\mathcal{F}$ and induced by special crossings $(pa, q, \mathcal{I}_r)$ and $(wa, q, \mathcal{J}_u)$. The correspondence between special crossings and their ordinary quadruples in $\mathcal{F}$ implies that the edges pq and wq undergo Delaunay crossings by the respective outer points r and u; see Figure 41. Furthermore, if χ is a right special quadruple, then condition (S3a) implies that pq or, more precisely, its reversely oriented copy qp, undergoes a Delaunay crossing (in the reduced triangulation $\mathrm{DT}(P \setminus \{a, r, u\})$) by w, so the points r and w cross pq *in opposite directions*. Similarly, if χ is a left special quadruple, then the edge wq undergoes two oppositely oriented Delaunay crossings, by u and p (the latter occuring within $\mathrm{DT}(P \setminus \{a, r, u\})$, as above).

Our general strategy is to charge χ to one of the above configurations (p, q, r, w) or (w, q, u, p) (depending on the right or left nature of χ), which will be referred to as *terminal quadruples*. Notice that each of those configurations involves one of the outer points r and u, in addition to some three regular points of χ. Nevertheless, several preparatory restrictions need to be enforced upon our special quadruples before actually charging them to terminal quadruples. Informally, this is done to further restrict the arising terminal quadruples and, consequently, to facilitate their eventual treatment at Stage 4 and in Section 7.

At the subsequent Stages 2 and 3, we do not distinguish between left and right special quadruples $\chi = (a, p, w, q)$. The topological restrictions enforced during these stages on special quadruples are fairly analogous to the ones enforced on ordinary quadruples during the parallel stages in Section 5. Namely, for each χ as above we extend the almost-Delaunayhood of its three edges aq, pq, and wq from, respectively, $[\lambda_0, \lambda_3]$, $[\lambda_0, \lambda_{pq}]$, and $[\lambda_{wq}, \lambda_3]$ to larger intervals, which cover $[\lambda_{wq}, \lambda_{pq}]$. The intimate correspondence between special crossings and ordinary quadruples is largely ignored throughout these technical stages, and the outer points r and u do not play any meaningful role.

At the last Stage 4, we finally distinguish between left and right special quadruples. In both cases, we exploit the interplay between our quadruples and their respective outer points r and u, which re-enter the analysis and finally give rise to terminal quadruples. (As noted above, for each of the two types, only one outer point is used.) This analysis is preceded by several preparatory charging arguments, analogous to the ones described in Section 5.6.

6.3 Stage 2

Let $\chi = (a, p, w, q)$ be a 1-restricted (Delaunay) special quadruple. Our next goal is to extend the almost Delaunayhood of qa from $[\lambda_0, \lambda_3] = [\mathcal{I}_r, \mathcal{J}_u]$ to some larger interval $[\xi^-_{qa}, \xi^+_{qa}]$, which covers $[\lambda_{wq}, \lambda_{pq}]$. As in the parallel Section 5.4, we proceed in two steps, after fixing the constant parameters $12 < k \ll \ell$.

Stage (2a). First, we consider the interval $[\lambda_{wq}, \lambda_3]$ where, by assumption, wq is almost Delaunay. (It is in fact Delaunay in $P \setminus \{u\}$ throughout $\mathcal{J}_u = [\lambda_2, \lambda_3]$ and at time λ_{wq}.) Refer to Figure 51 (left). If at least k special counterclockwise (a, q)-crossings $(w'a, q, \mathcal{J}_{u'})$ (in $\mathcal{F}$) begin during $[\lambda_{wq}, \lambda_2)$, then we can bound the overall number of such special quadruples χ via the already routine combination of Theorem 5.3 with random sampling.

Note, as a preparation, that the preliminary pruning (described in at beginning of this section) ensures that each of the above special (a, q)-crossings $(w'a, q, \mathcal{J}_{u'})$, where u' is its respective outer point, satisfies $\{w', u'\} \cap P_\chi = \emptyset$. Therefore, Lemma 5.5 implies that q hits each of the respective edges $w'a$ (during $\mathcal{J}_{u'}$) before it hits wa (during $\mathcal{J}_u$).

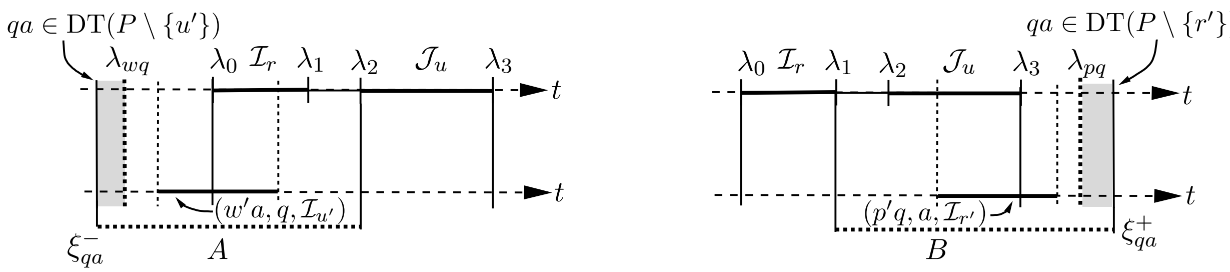

Figure 51: Extending the almost-Delaunayhood of qa from $[\lambda_0, \lambda_3]$ to $[\xi^-_{qa}, \lambda_0]$ (left) and to $[\lambda_3, \xi^+_{qa}]$ (right).

To set the stage for an application of Theorem 5.3, we consider a random subset $\hat{P} \subset P$ of $\lceil n/2 \rceil$ points, and argue that, with some fixed positive probability, $(wa, q, \mathcal{J}_u)$ becomes a $(w, q, \Theta(k))$-chargeable Delaunay crossing in $\hat{P}$ (with a potentially shrunk interval $\mathcal{J}_u$), with the reference interval $[\lambda_{wq}, \lambda_3]$ (the proof of this property is identical to that given in Sections 5.6 and 6.2). Briefly, this follows because $\hat{P}$ satisfies the following two conditions with probability $\Omega(1)$: (1) $\hat{P}$ includes a, w, q but not u, and (2) for at least a constant fraction of the above special (a, q)-crossings $(w'a, q, \mathcal{J}_{u'})$, $\hat{P}$ includes w' but not u'. The former condition guarantees that $\hat{P}$ yields a Delaunay crossing $(wa, q, \mathcal{J})$, for some interval $\mathcal{J} \subseteq \mathcal{J}_u$, and that qa belongs to $\mathrm{DT}(\hat{P})$ at times λ_{wq} and λ_3. The latter condition implies that $\Omega(k)$ (ordinary) counterclockwise (a, q)-crossings occur within $[\lambda_{wq}, \lambda_2) \cup \mathcal{J} \subseteq [\lambda_{wq}, \lambda_3]$.

Theorem 5.3 now implies that the overall number of the above triples (w, q, a) in $\hat{P}$ is at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$. By Proposition 6.1, this yields the same bound on the maximum number of the special quadruples χ that fall into the present scenario. Assume, then, that at most k clockwise special (a, q)-crossings $(w'a, q, \mathcal{J}_{u'})$ begin during $[\lambda_{wq}, \lambda_2)$.

If no clockwise special (a, q)-crossings begin in $(-\infty, \lambda_{wq}]$, then $(wa, q, \mathcal{J}_u)$ is among the first $k + 1$ such special (a, q)-crossings $(w'a, q, \mathcal{J}_{u'})$, so it can be charged to the pair (a, q). (After the preliminary pruning, there remain no counterclockwise special (a, q)-crossings $(w'a, q, \mathcal{J}_{u'})$ with $u' \in P_\chi$. Furthermore, Lemma 4.1 implies that no other such (a, q)-crossings, with $u' \notin P_\chi$, can begin in $[\xi_{qa}^-, \lambda_{wq})$.) Therefore, and because of Proposition 6.1, this happens for at most $O(kn^2)$ special quadruples χ.

Assume next that some clockwise special (a,q)-crossing $(w'a, q, \mathcal{J}_{u'})$ begins in $(-\infty, \lambda_{wq}]$. Therefore, using Lemma 4.1, there is a last time ξ_{qa}^- in $(-\infty, \lambda_{wq}]$ when the edge qa belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{u'\})$, for $u' \notin P_\chi$. In what follows, we use u' to denote such a (fixed) point whose removal restores the Delaunayhood of qa at (the last possible) time ξ_{qa}^-.

To proceed, we apply Theorem 2.2 in $\mathcal{A}_{qa}$ over the interval (ξ_{qa}^-, λ_2). We do this for the above, reduced triangulation $\mathrm{DT}(P \setminus \{u'\})$, and with the second constant ℓ. If at least one of the Conditions (i), (ii) of that theorem holds, we charge χ (via $(wa, q, \mathcal{J}_u)$) either to an $(\ell+1)$-shallow collinearity or to $\Omega(\ell^2)$ $(\ell+1)$-shallow co-circularities. (Each of these events is ℓ-shallow in $\mathrm{DT}(P \setminus \{u'\})$.) The choice of ξ_{qa}^- implies that no special (a,q)-crossing $(w'a, q, \mathcal{J}_{u'})$ begins in $[\xi_{qa}^-, \lambda_{wq})$, and therefore, arguing as above, it guarantees that any event in $\mathcal{A}_{qa}$ is charged by at most $O(k)$ quadruples. Hence, this charging accounts for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ quadruples χ.

Finally, if Condition (iii) of Theorem 2.2 holds, then there is a set A of at most $3\ell+1$ points (including u') whose removal restores the Delaunayhood of qa throughout $[\xi_{qa}^-, \lambda_3]$.

Stage (2b). We similarly use Theorem 5.3 to extend the almost-Delaunayhood of qa from $\mathcal{I}_r = [\lambda_0, \lambda_1]$ to the interval $(\lambda_1, \lambda_{pq}]$ where, by assumption, the edge pq is almost Delaunay. (It is Delaunay in $P \setminus \{r\}$ throughout $\mathcal{I}_r = [\lambda_0, \lambda_1]$ and at time λ_{pq}.) The argument is fully symmetric to the one in Stage (2a), but we briefly repeat it for the sake of completeness.

Refer to Figure 51 (right). If at least k special (a,q)-crossings $(p'a, q, \mathcal{I}_r)$ end in $(\lambda_1, \lambda_{pq}]$ then we again use Theorem 5.3 to show that the number of such special quadruples is at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$. In short, we argue that a random subset of $\lceil n/2 \rceil$ points yields a $(p, q, \Theta(k))$-chargeable Delaunay crossing of pa by q, with probability $\Theta(1)$.) Hence, we can assume that at most k special (a,q)-crossings, as above, end during $(\lambda_1, \lambda_{pq}]$.

If no clockwise special (a,q)-crossings begin in $[\lambda_{pq}, \infty)$, then $(wa, q, \mathcal{J}_u)$ is among the last $k+1$ such special (a,q)-crossings $(p'a, q, \mathcal{I}_{r'})$, so it can be charged to the pair (a,q). Clearly, that scenario occurs for at most $O(kn^2)$ special quadruples χ.

Otherwise, we choose the first time ξ_{qa}^+ in $[\lambda_{pq}, \infty)$ when the edge qa belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, with $r' \notin P_\chi$. In what follows, we use r' to denote such a (fixed) point whose removal restores the Delaunayhood of qa at time ξ_{qa}^+. We then apply Theorem 2.2 in $\mathcal{A}_{qa}$ over the interval (λ_1, ξ_{qa}^+). This is done with respect to the point set $P \setminus \{r'\}$, and with the constant ℓ.

If at least one of the Conditions (i), (ii) is satisfied, we dispose of χ by charging it (via $(pa, q, \mathcal{I}_r)$) to $(\ell+1)$-shallow events in $\mathcal{A}_{qa}$, and argue, as above, that each event is charged by at most $O(k)$ quadruples. Hence, the above charging accounts for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ special quadruples.

Finally, if none of the preceding scenarios occur, we end up with a subset B of at most $3\ell+1$ points (including r') whose removal restores the Delaunayhood of aq throughout $[\lambda_0, \xi_{qa}^+]$.

Probabilistic refinement. We say that a special quadruple $\chi = (a, p, w, q)$ is 2-*restricted* if (1) it is 1-restricted with respect to the underlying set P and refined family $\mathcal{F}$, and (2) it satisfies the following new condition:

(S4) The edge qa belongs to $\mathrm{DT}(P \setminus \{p, w, u, r\})$ throughout the interval $[\xi_{aq}^-, \xi_{aq}^+]$, where ξ_{aq}^- (resp., ξ_{aq}^+) denotes the last time in $(-\infty, \lambda_{wq}]$ (resp., first time in $[\lambda_{pq}, \infty)$) when the edge aq is Delaunay (and where we assume that the times ξ_{qa}^-, ξ_{qa}^+ exist).

Let $\Phi_2(m)$ denote the maximum number of 2-restricted special quadruples that can be defined over a set of m moving points (and a refined family $\mathcal{F}$). The preceding analysis, combined with the standard sampling argument of Clarkson and Shor, leads to the following recurrence:

$$\Phi_1(n) = O\left(\ell^6 \Phi_2(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right). \tag{8}$$

6.4 Stage 3

To bound the above quantity $\Phi_2(n)$, we fix a set P of n moving points, and a refined family $\mathcal{F}$. In addition, we fix a 2-restricted special quadruple $\chi = (a, p, w, q)$ (with outer points r and u), which is defined with respect to P and $\mathcal{F}$.

Recall that the edge pq is Delaunay at time λ_{pq}, and that it is almost Delaunay during $[\lambda_0, \lambda_{pq}]$ (it is Delaunay with the omission of a, w, r and u). Similarly, wq is Delaunay at time λ_{wq}, and it is almost Delaunay during $[\lambda_{wq}, \lambda_3]$ (it is Delaunay with the omission of a, p, r, u). Our goal in this stage is (i) to extend the almost-Delaunayhood of pq to a (possibly) larger interval $[\xi_{pq}, \lambda_{pq}]$, for some $\xi_{pq} \le \xi_{qa}^- \le \lambda_{wq}$, and (ii) to extend the almost-Delaunayhood of wq to an interval $[\lambda_{wq}, \xi_{wq}]$, for some $\xi_{wq} \ge \xi_{qa}^+ \ge \lambda_{pq}$.

Our analysis consists of two symmetric arguments, similar to the ones used in Section 6.2 (cases (b) and (c)). Both arguments use Theorem 5.3 (in combination with the almost-Delaunayhood of qa in $[\xi_{qa}^-, \xi_{qa}^+]$) and refer to the same pair of constant parameters $12 < k \ll \ell$.

Extending the almost-Delaunayhood of pq. Refer to Figure 52 (left). If at least k special (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$ begin during $[\xi_{qa}^-, \lambda_0)$, then we can invoke Theorem 5.3 to show that the number of such special quadruples χ is at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$.

Specifically, recall that the edge qa is Delaunay at times ξ_{qa}^-, ξ_{qa}^+, and that it is almost Delaunay (with only four potentially obstructing points p, w, u, r) during $[\xi_{qa}^-, \xi_{qa}^+] \supset [\xi_{qa}^-, \lambda_0) \cup \mathcal{I}_r = [\xi_{qa}^-, \lambda_1]$. Hence, a random subset $\hat{P} \subset P$ of $\lceil n/2 \rceil$ points would make $(pa, q, \mathcal{I}_r)$, with some fixed positive probability, an $(a, q, \Theta(k))$-chargeable crossing in $\hat{P}$ with $[\xi_{qa}^-, \lambda_1]$ as a reference interval (where ξ_{qa}^- and λ_1 are still defined with respect to P, and $\mathcal{I}_r$ is possibly shrunk in $\hat{P}$).

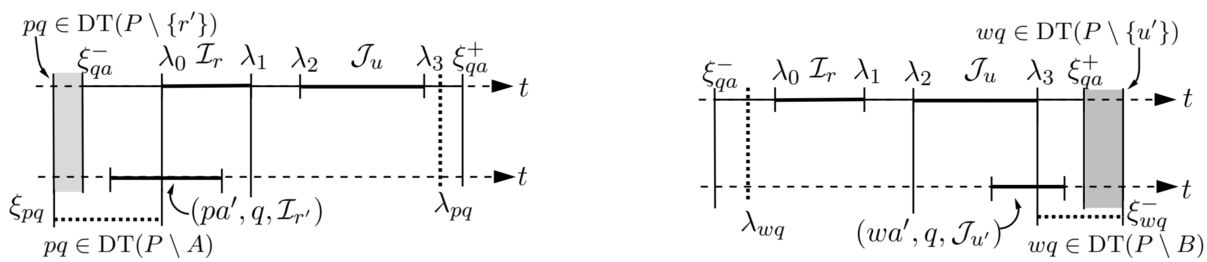

Figure 52: Left: Extending the almost-Delaunayhood of pq from $[\lambda_0, \lambda_{pq}]$ to $[\xi_{pq}, \lambda_{pq}]$. Right: Extending the almost-Delaunayhood of wq from $[\lambda_{wq}, \lambda_3]$ to $[\lambda_{wq}, \xi_{wq}]$.

Assume then that at most k clockwise special (p, q)-crossings begin during $[\xi_{qa}^-, \lambda_0)$. If no such (p, q)-crossings $(pa', q, \mathcal{I}_{r'})$, with $r' \notin P_\chi$, begin before λ_0, then $(pa, q, \mathcal{I}_r)$ is among the first $k + 3$ clockwise special (p, q)-crossings (including such crossings whose respective outer point r' belongs to P_χ).[31] Clearly, the overall number of such quadruples χ is at most $O(kn^2)$.

We may therefore assume that the previous sub-scenario does not occur. In particular, there exists ξ_{pq} which is the last time in $(-\infty, \xi_{qa}^-]$ when pq belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, for $r' \notin P_\chi$. In what follows, we use r' to denote such a fixed point, whose removal restores the Delaunayhood of pq at (the latest possible) time ξ_{pq}.

To proceed, we apply Theorem 2.2 in $\mathcal{A}_{pq}$ over the interval (ξ_{pq}, λ_0). We do so with the second constant ℓ, and with respect to the reduced point set $P \setminus \{r'\}$.

If at least one of the Conditions (i), (ii) of Theorem 2.2 holds, we charge χ (via $(pa, q, \mathcal{I}_r)$) either to $\Omega(\ell^2)$ $(\ell + 1)$-shallow co-circularities, or to an $(\ell + 1)$-shallow collinearity. As above, the choice of ξ_{pq} guarantees that $(pa, q, \mathcal{I}_r)$ is among the first $k + 3$ special (p, q)-crossings to begin after any event that

[31] Recall from Section 6.2 that at most two such crossings have $r' \in P_\chi$.

we charge within $\mathcal{A}_{pq}$, so any event is charged as above by at most $O(k)$ quadruples χ. Hence, the above charging is applicable for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2\beta(n)\right)$ special quadruples.

Finally, if Condition (iii) of Theorem 2.2 is satisfied, we have a set A of at most $3\ell+1$ points (including r', and perhaps some of a, w, r, u) whose removal restores the Delaunayhood of pq throughout $[\xi_{pq}, \lambda_0]$. We further add to our conflict set A every point a' whose respective (p,q)-crossing $(pa', q, \mathcal{I}_{r'})$ begins in $[\xi_{pq}, \lambda_0)$. (This is done to ensure that these (p,q)-crossings do not arise in the following Stage 4. Note that at most 2 such crossings $(pa', q, \mathcal{I}_{r'})$ begin in $[\xi_{pq}, \xi_{qa}^-)$, and each of them satisfies $r' \in P_\chi$.) Since there are at most $k+2$ crossings of this kind, and since $k \ll \ell$, the cardinality of the augmented set A does not exceed 4ℓ.

Extending tha almost-Delaunayhood of wq. The argument is fully symmetric to the one that was used for pq, but we briefly repeat it for the sake of completeness.

Refer to Figure 52 (right). If at least k special (w,q)-crossings $(wa', q, \mathcal{J}_{u'})$ end during $(\lambda_3, \xi_{qa}^+]$, we consider a random subset of $\lceil n/2 \rceil$ points and argue as before that $(wa, q, \mathcal{J}_u)$ becomes, with some fixed positive probability, an $(a, q, \Theta(k))$-chargeable special crossing (now with $[\lambda_2, \xi_{qa}^+]$ as the reference interval). Therefore, Theorem 5.3 implies that the number of such special quadruples χ is at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$.

Assume then that at most k clockwise special (w,q)-crossings $(wa', q, \mathcal{J}_{u'})$ end during $(\lambda_3, \xi_{qa}^+]$. Furthermore, we may assume that there exists ξ_{wq}, which is the first time in $[\xi_{qa}^+, \infty)$ when the edge wq belongs to some triangulation $\mathrm{DT}(P \setminus \{u'\})$, for $u' \not\in \{a, r, w, u\}$. (Otherwise, $(wa, q, \mathcal{J}_u)$ would be among the last $k+3$ clockwise special (w,q)-crossings, which can happen for at most $O(kn^2)$ special quadruples of the kind considered here.) In what follows, we use u' to denote a fixed point whose removal restores the Delaunayhood of wq at time ξ_{wq}.

To proceed, we apply Theorem 2.2 in $\mathcal{A}_{wq}$ over the interval (λ_3, ξ_{wq}), with the second parameter ℓ and respect to the point set $P \setminus \{u'\}$.

If at least one of the Conditions (i), (ii) of Theorem 2.2 holds, we dispose of χ by charging it (via $(wa, q, \mathcal{J}_u)$) to $(\ell+1)$-shallow events in $\mathcal{A}_{wq}$. The choice of ξ_{wq} guarantees that each collinearity or co-circularity is charged in this manner by at most $O(k)$ quadruples χ. Hence, the above charging is applicable for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2\beta(n)\right)$ special quadruples.

Finally, if Condition (iii) of Theorem 2.2 is satisfied, we end up with a subset B of at most $3\ell+1$ points (including u' and perhaps some of a, p, r, u) whose removal restores the Delaunayhood of wq throughout the interval $[\lambda_3, \xi_{wq}]$. We add to B every point a' whose respective crossing $(wa', q, \mathcal{J}_{u'})$ ends in $(\lambda_3, \xi_{qa}^+]$. (As before, this is done to ensure that these (w,q)-crossings do not arise in the following Stage 4.) As above, the cardinality of the augmented set B does not exceed 4ℓ.

Probabilistic refinement. We say that a special quadruple χ is 3-*restricted* if (1) it is 2-restricted, and (2) it satisfies the following additional conditions:

(S5) The edge pq belongs to $\mathrm{DT}(P \setminus \{a, w, u, r\})$ throughout the interval $[\xi_{pq}, \lambda_{pq}]$, where ξ_{pq} denotes the last time in $(-\infty, \xi_{qa}^-]$ when the edge pq is Delaunay (and we assume the existence of such a time ξ_{pq}). In addition, at most two special (p,q)-crossings $(pa', q, \mathcal{I}_{r'})$ begin during $[\xi_{pq}, \lambda_0)$ (namely, the possible crossings of pw and pu by q).

(S6) The edge wq belongs to $\mathrm{DT}(P \setminus \{a, p, u, r\})$ throughout the interval $[\lambda_{wq}, \xi_{wq}]$, where ξ_{wq} denotes the first time in $[\xi_{qa}^+, \infty)$ when the edge wq is Delaunay (and we assume the existence of such a time ξ_{wq}). In addition, at most two special (w,q)-crossings $(wa', q, \mathcal{J}_{u'})$ end during $(\lambda_3, \xi_{wq}]$ (namely, the possible crossings of wp and wr by q).

Let $\Phi_3^R(m)$ (resp., $\Phi_3^L(m)$) denote the maximum number of 3-restricted right (resp., left) special quadruples that can be defined over a set of m moving points (and a fixed refined family $\mathcal{F}$). The preceding analysis, in combination with the routine sampling argument of Clarkson and Shor, implies the following recurrence:

$$\Phi_2(n) = O\left(\ell^6 \Phi_3^R(n/\ell) + \ell^6 \Phi_3^L(n/\ell) + k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right) \tag{9}$$

6.5 Stage 4: The number of right quadruples

To bound the maximum possible number $\Phi_3^R(n)$ of 3-restricted right special quadruples, we fix the underlying set P of n moving points, and a refined family $\mathcal{F}$.

Topological setup. According to Proposition 6.2, any 3-restricted quadruple $\chi = (a, p, w, q)$ shares its triple (p, q, w) with at most two other such quadruples. (In other words, it suffices to bound the overall number of the corresponding triples (p, q, w).) We strengthen the above property, by considering, without loss of generality, at most *one* 3-restricted right quadruple for each triple (p, q, w). Therefore, in what follows every special quadruple $\chi = (a, p, w, q)$ under consideration will be uniquely determined by its triple (p, q, w).

To proceed, we fix a 3-restricted right special quadruple $\chi = (a, p, w, q)$, with respect to P and $\mathcal{F}$, whose two special (a, q)-crossings take place during the intervals $\mathcal{I}_r = [\lambda_0, \lambda_1]$ and $\mathcal{J}_u = [\lambda_2, \lambda_3]$ (in this order), where r and u are the respective outer points. Recall that the original "regular" family $\mathcal{F}$ includes the quadruples $\sigma_1 = (p, q, a, r)$ and $\sigma_2 = (w, q, a, u)$.

Refer to Figure 53. Since χ is 3-restricted, there exist a time $\lambda_{wq} \leq \lambda_0$ which is the last time before[32] λ_0 when the edge wq belongs to $\mathrm{DT}(P)$, and a symmetric first time $\lambda_{pq} \geq \lambda_3$ when pq belongs to $\mathrm{DT}(P)$. By Condition (S4), there exist the first time ξ_{qa}^+ in $[\lambda_{pq}, \infty)$, and the last time ξ_{qa}^- in $(-\infty, \lambda_{wq}]$ when the edge qa is Delaunay, so that this edge is almost-Delaunay during the interval $[\xi_{qa}^-, \xi_{qa}^+]$ (with only p, w, u, r as the possible obstructing points). Moreover, by Conditions (S5) and (S6), there exist the first time $\xi_{wq} \in [\xi_{qa}^+, \infty)$, and the symmetric last time $\xi_{pq} \in (-\infty, \xi_{qa}^-]$ when the respective edges wq and pq are Delaunay. Moreover, wq and pq are almost Delaunay during, respectively, $[\lambda_{wq}, \xi_{wq}]$ and $[\xi_{pq}, \lambda_{pq}]$ (each with four obstructing points, as specified in these conditions).

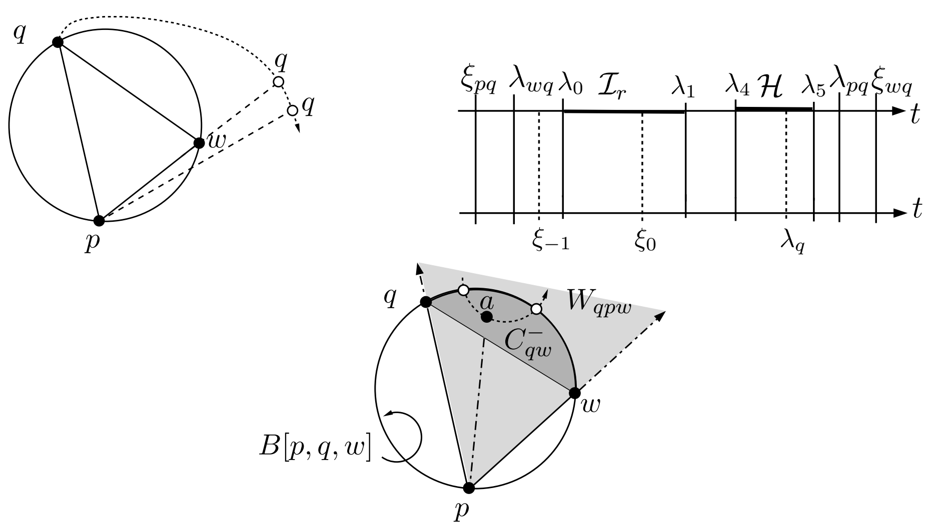

Figure 53: The topologt76ical setup during the interval $(\xi_{-1}, \lambda_q) \subseteq [\lambda_{wq}, \lambda_{pq}]$. Left: The edge qp is hit at some time $\lambda_q \in [\lambda_1, \lambda_{pq}]$ by w, so it undergoes a Delaunay crossing $(qp, w, \mathcal{H} = [\lambda_4, \lambda_5])$ within $\mathrm{DT}(P \setminus \{a, r, u\})$. Right: We have $\lambda_{wq} \leq \xi_{-1} < \xi_0 < \lambda_4 < \lambda_q < \lambda_5 < \xi_{wq}$. Bottom: The motion of $B[p, q, w]$ is continuous throughout $[\xi_{-1}, \lambda_q)$ (the hollow circles represent the co-circularities at times ξ_{-1} and ξ_0).

[32]If wq is Delaunay at time λ_0 then we put $\lambda_{wq} = \lambda_0$.

Let us summarize what we know so far about the motion of a, p, w, q. By Condition (S3a), these points are co-circular at times $\xi_{-1} \in [\lambda_{wq}, \lambda_0)$, $\xi_0 \in \mathcal{I}_r \setminus \mathcal{J}_u$, and $\xi_1 \in \mathcal{J}_u \setminus \mathcal{I}_r$. Moreover, the Delaunayhood of wq is violated, throughout (ξ_{-1}, ξ_0), by the points $a \in L^+_{wq}$ and $p \in L^-_{wq}$ (so, in particular, neither of these points crosses wq during this period). Hence, a lies throughout that interval within the wedge $W_{qpw} = L^+_{pq} \cap L^-_{pw}$ and inside the cap $C^-_{qw} = B[p, q, w] \cap L^-_{qw}$.We emphasize that the order type of the quadruple (q, p, w, a) remains unchanged during (ξ_{-1}, ξ_0).

In addition, by the same Condition (S3a), the smaller set $P \setminus \{a, r, u\}$ yields a (single) Delaunay crossing $(qp, w, \mathcal{H}_\chi)$, whose interval $\mathcal{H} = \mathcal{H}_\chi = [\lambda_4, \lambda_5]$ is contained in $(\lambda_1, \lambda_{pq}]$. In particular, w hits pq at some moment[33] $\lambda_q \in \mathcal{H}$, when w crosses L_{pq} from L^+_{pq} to L^-_{pq}. Since w lies in L^+_{pq} at times ξ_{-1} and ξ_0, no further collinearities of p, w, q can occur during $[\xi_{-1}, \lambda_q)$. (Otherwise, the point w would have to re-enter L^+_{pq} before λ_q, and then the triple p, q, w would be collinear three times, contrary to our assumptions.) To conclude, the disc $B[p, q, w]$ moves continuously throughout the interval $[\xi_{-1}, \lambda_q)$, which is obviously contained in $[\xi_{pq}, \lambda_{pq}] \cap [\lambda_{wq}, \xi_{wq}] = [\lambda_{wq}, \lambda_{pq}]$.

Overview. We fix three constant parameters k, ℓ, h, such that $12 < k \ll \ell \ll h$, and distinguish between four possible cases. The first two cases (a)–(b) are fairly similar to the cases (a)–(b) that we encountered in Section 5.6 when handling ordinary quadruples, and case (c) is very similar to the preceding case (b). In case (a) we bound the number of right special quadruples χ, that fall into it, using Theorem 5.3. In each of the subsequent cases (b) and (c), we manage to bound the number of special quadruples χ, that fall into that case, by charging them within the arrangements $\mathcal{A}_{pw}$, $\mathcal{A}_{pq}$ and $\mathcal{A}_{wq}$. (The crucial difference between the two setups is that the extremal co-circularity among ξ_0 and ξ_1 now occurs during the *second* crossing $(wa, q, \mathcal{J}_u)$, so the topological analysis of Section 5.6 must be performed in a "time-reversed" manner.)

In the final, most involved, case (d), we re-introduce at last the outer point r. (The other outer point u is not used in the analysis of right special quadruples.) The correspondence between $(pa, q, \mathcal{I}_r)$ and its ancestor quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$ implies that the points r and w cross the same edge pq in opposite directions. Hence, χ can be charged to the resulting so-called *terminal quadruple* (p, q, r, w). In Section 7 we express the number of these terminal quadruples in terms of more elementary quantities, that were introduced in Section 2. This, combined with a parallel (and mostly symmetric, although considerably simplified) analysis of 3-restricted *left* special quadruples, finally produces a complete system of recurrences whose solution is $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$.

In what follows, we consider the family $\mathcal{G}^R_{pw}$ of all 3-restricted right special quadruples of the form $\chi' = (a', p, w, q')$, which share their middle pair with χ. We may assume that each $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw}$ is uniquely determined by the choice of q' (as the only "free" point in the triple (p, q', w)). Note that the set $P_{\chi'}$ of each χ' includes, in addition to the four points a', p, w, q' of χ', the respective outer points r' and u' of its special crossings $(pa', q', \mathcal{I}_{r'})$ and $(wa', q', \mathcal{J}_{u'})$. Furthermore, each of these quadruples $\chi' \in \mathcal{G}^R_{pw}$ is accompanied by a counterclockwise (p, w)-crossing $(q'p, w, \mathcal{H}_{\chi'} = \mathcal{H}')$, which occurs within the smaller triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$. See Figure 54. We use $\lambda_{q'}$ to denote the time in $\mathcal{H}'$ when the respective point q' of χ' enters the halfplane L^+_{pw} (or, equivalently, when w crosses $q'p$ from $L^+_{pq'} = L^-_{q'p}$ to $L^+_{q'p}$).

Notice that Lemma 5.5 readily generalizes to the above (p, w)-crossings. Namely, a pair of such crossings $(qp, w, \mathcal{H}_\chi)$ and $(q'p, w, \mathcal{H}_{\chi'})$, which occur within the respective triangulations $\mathrm{DT}(P \setminus \{a, r, u\})$ and $\mathrm{DT}(P \setminus \{a', r', u'\})$, are *compatible*, provided that $q' \neq a, r, u$ and $q \neq a', r', u'$, in the sense that the orders in which the intervals $\mathcal{H}_\chi$ and $\mathcal{H}_{\chi'}$ begin or end are both consistent with the time stamps λ_q and $\lambda_{q'}$.

[33]Recall from Section 6.2 that w can cross qp either before or after ξ_2, depending on the location of p when q crosses wa. Our analysis only relies on the fact that $\lambda_q > \xi_0 > \xi_{-1}$.

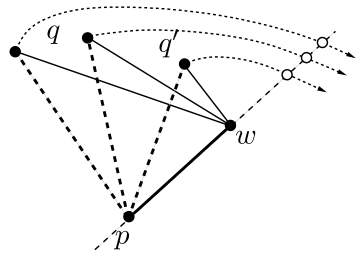

Figure 54: Each right special quadruple $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw}$ (with respective outer points r' and u') comes with a counterclockwise (p, w)-crossing $(q'p, w, \mathcal{H}_{\chi'})$, which occurs within $\mathrm{DT}(P \setminus \{a', r', u'\})$.

To proceed, we distinguish between four possible cases.

Case (a). For at least k of the above quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw}$, their respective (p, w)-crossings $(q'p, w, \mathcal{H}')$ either begin in $[\lambda_{wq}, \lambda_4]$, or end in $[\lambda_5, \xi_{wq}]$. Refer to Figure 55. Recall that, by condition (S6), the edge qw is Delaunay at each of the times λ_{wq} and ξ_{wq}, and that it is almost Delaunay during the entire interval $[\lambda_{wq}, \xi_{wq}]$.

To bound the number of such quadruples χ, we wish to argue that the crossing $(qp, w, \mathcal{H})$ is $(q, w, \Theta(k))$-chargeable, for the reference interval $[\lambda_{wq}, \xi_{wq}]$. Unfortunately (and we have already encountered this technical issue before, e.g., in Section 5.6), the crossing $(qp, w, \mathcal{H})$ occurs within the reduced triangulation $\mathrm{DT}(P \setminus \{a, r, u\})$, whereas each of the above crossings $(q'p, w, \mathcal{H}')$ occurs within a possibly different (and also reduced) triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$.

As in the previous similar situations (including the matching scenario (a) in Section 5.6), we can free sufficiently many crossings $(q'p, w, \mathcal{H}')$ from their "violators" a', r' and u' by passing to a smaller triangulation $\mathrm{DT}(\hat{P})$, which is induced by a random subset $\hat{P} \subset P$ of $\lceil n/4 \rceil$ points. Note though that $\mathcal{G}^R_{pw}$ can potentially include many quadruples χ' with $q \in \{a', r', u'\}$, which cannot be freed without destroying $(qp, w, \mathcal{H})$.

Fortunately, for any special quadruple $\chi = (a, p, w, q) \in \mathcal{G}^R_{pw}$ (with outer points r and u) the family $\mathcal{G}^R_{pw}$ includes at most three other quadruples $\chi' = (a', p, w, q')$ whose respective points q' are equal to one of a, r or u. The pigeonhole principle then implies that at least *one quarter* of all quadruples $\chi = (a, p, w, q)$ in $\mathcal{G}^R_{pw}$ satisfy the following condition:

(PHR1) *There exist at most three quadruples $\chi' \in \mathcal{G}^R_{pw}$ with $q \in \{a', r', u'\}$.*

(See Section 5.6 for the short proof of a similar claim, with the matching condition (PH).)

Since p and w are arbitrary points of P, (PHR1) holds for at least a quarter of all 3-restricted right special quadruples under consideration; hence we may assume that it holds for the special quadruple χ at hand. Therefore, at least $k - 6 \geq k/2$ of the relevant quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw} \setminus \{\chi\}$ (with respective outer points r' and u', and with $(q'p, w, \mathcal{H}')$ starting in $[\lambda_{wq}, \lambda_4]$ or ending in $[\lambda_5, \xi_{wq}]$) satisfy (i) $q \notin \{a', r', u'\}$, and (ii) $q' \notin \{a, r, u\}$.

A suitable extension of Lemma 5.5 then implies that at least $k/2$ of the above crossings $(q'p, w, \mathcal{H}')$ fully occur within $[\lambda_{wq}, \xi_{wq}]$. Returning to the sampled triangulation $\mathrm{DT}(\hat{P})$, it is easy to check that the following two conditions hold simultaneously with some fixed probability (see Stage 1 of this section for a similar argument): (1) the set $\hat{P}$ includes p, q and w, but none of a, r, u, and (2) for at least $\Theta(k)$ of the above quadruples χ' (with $\mathcal{H}_{\chi'}$ starting in $[\lambda_{wq}, \lambda_4)$ or ending in $(\lambda_5, \xi_{wq}]$), the sample $\hat{P}$ includes their respective points q', but none of a', r', u'.

In the case of success, $\hat{P}$ yields a $(q, w, \Theta(k))$-chargeable (ordinary) Delaunay crossing of qp by w, for the reference interval $[\lambda_{wq}, \xi_{wq}]$. To see this, recall that wq is Delaunay at both times λ_{wq} and ξ_{wq}, and that it is almost Delaunay in (λ_{wq}, ξ_{wq}) (it is Delaunay with the omission of a, p, r, u). Then, according to condition (1), the sample $\hat{P}$ yields some single Delaunay crossing $(qp, w, \hat{\mathcal{H}}_\chi)$, whose respective interval

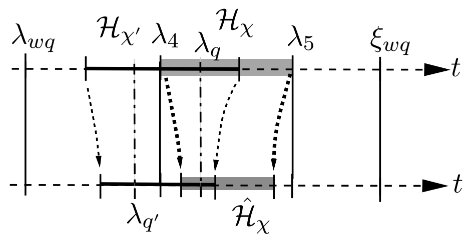

Figure 55: Case (a): At least k counterclockwise (p,w)-crossings $(q'p,w,\mathcal{H}_{\chi'})$ either begin in $[\lambda_{wq},\lambda_4)$ or end in $(\lambda_5,\xi_{wq}]$ (one such crossing of the former type is depicted). Then, with some fixed and positive probability, the sample $\hat{P}$ yields a Delaunay crossing $(qp,w,\hat{\mathcal{H}}_\chi)$ that is $(q,w,\Theta(k))$-chargeable with respect to $[\lambda_{wq},\xi_{wq}]$.

$\hat{\mathcal{H}}_\chi$ is contained in $\mathcal{H}_\chi$ (as depicted in Figure 55). Similarly, according to condition (2), $\hat{P}$ yields at least $\Theta(k)$ counterclockwise Delaunay (p,w)-crossings that occur within $[\lambda_{wq},\xi_{wq}]$.

To conclude, Theorem 5.3 implies that the overall number of such triples (p,q,w) in P does not exceed

$$O\left(k^2N(n/k)+kn^2\beta(n)\right),$$

which immediately also bounds the overall number of the corresponding 3-restricted quadruples χ.

Preparing for cases (b) and (c): Charging events in $\mathcal{A}_{pw}$. We may assume, from now on, that there exist at most k special quadruples $\chi'\in\mathcal{G}_{pw}^R$ whose respective (p,w)-crossings $(q'p,w,\mathcal{H}')$ either begin in $[\lambda_{wq},\lambda_4]$, or end in $[\lambda_5,\xi_{wq}]$.

Before proceeding to the following cases, we apply Theorem 2.2 in $\mathcal{A}_{pw}$ in order to extend the almost-Delaunayhood of pw from $\mathcal{H}=[\lambda_4,\lambda_5]$ to $[\lambda_{wq},\xi_{wq}]$. Notice that $[\lambda_{wq},\xi_{wq}]\setminus\mathcal{H}$ consists of two, possibly empty, intervals $[\lambda_{wq},\lambda_4)$ and $(\lambda_5,\xi_{wq}]$, and we consider each of them separately. Note also that the edge pw belongs during $\mathcal{H}$ to the reduced triangulation $\mathrm{DT}(P\setminus\{a,r,u\})$ (but not necessarily to $\mathrm{DT}(P)$), so Theorem 2.2 must be applied with respect to this smaller set.

Consider, for instance, the interval $[\lambda_{wq},\lambda_4)$. We apply Theorem 2.2 within $\mathcal{A}_{pw}$ over (λ_{wq},λ_4), with our second parameter ℓ, and with respect to the reduced set $P\setminus\{a,r,u\}$, noting that pw belongs to $\mathrm{DT}(P\setminus\{a,r,u\})$ at the end of this interval.

If at least one of the Conditions (i), (ii) holds, we charge χ within $\mathcal{A}_{pw}$, via $(qp,w,\mathcal{H})$, either to an $(\ell+3)$-shallow collinearity, or to $\Omega(\ell^2)$ $(\ell+3)$-shallow co-circularities in P. (Each of these events is ℓ-shallow with respect to the reduced set $P\setminus\{a,r,u\}$.) Notice that the points p and w are involved in each of these events, and since case (a) has been ruled out, at most k other (p,w)-crossings $(q'p,w,\mathcal{H}')$ of this kind begin after the respective time t^* of any charged event and before $(qp,w,\mathcal{H})$. That is, $(qp,w,\mathcal{H})$ is among the first $k+1$ such (p,w)-crossings to begin after t^*. Hence, any $(\ell+3)$-shallow collinearity or co-circularity is charged in the above manner by at most $O(k)$ special quadruples χ. To conclude, the above scenario occurs for at most $O\left(k\ell^2N(n/\ell)+k\ell n^2\beta(n)\right)$ quadruples χ.

Otherwise, if Condition (iii) holds, one can restore the Delaunayhood of pw throughout $[\lambda_{wq},\lambda_4]$ by removing at most $3\ell+3$ points of P (including a,r,u).

A fully symmetric argument can be used to extend the almost-Delaunayhood of pw to the symmetric interval $(\lambda_5,\lambda_{wq}]$. At the end, we have either disposed of χ through conditions (i), (ii) of Theorem 2.2 or ended up with a set A_{pw} of at most $6\ell+3$ points (including a,r,u) whose removal restores the Delaunayhood of pw throughout $[\lambda_{pq},\xi_{wq}]$. Hence, we may assume, in what follows, that the above set A_{pw} exists.

Case (b). There exist a total of at least ℓ points of P, distinct from a,r,u, such that each of them appears in the cap $C_{qw}^-=B[p,q,w]\cap L_{qw}^-$ at some time during the interval (ξ_{-1},λ_q). (Note that some of these points may belong to A_{pw}.) Recall that λ_q denotes the time in $\mathcal{H}$ when w enters L_{pq}^-, through pq, and

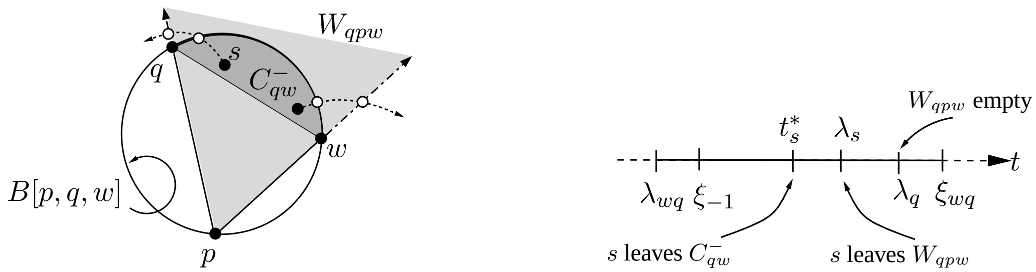

Figure 56: Case (b). A total of at least ℓ points $s \neq a, r, u$ appear in the cap C^-_{qw} during (ξ_{-1}, λ_q). Each of them must leave the cap C^-_{qw} (through the boundary of $B[p, q, w]$) and then leave the wedge W_{qpw} (through one of the rays $\vec{pq}, \vec{pw}$, outside the respective edges pq and pw) before time λ_q. Left: The geometric scenario. Right: A symbolic summary of the corresponding events.

that no additional collinearities of p, q, w can occur during (ξ_{-1}, λ_q), so the motion of $B[p, q, w]$ is fully continuous there.

Refer to Figure 56. Let $s \in P \setminus \{a, r, u\}$ be one of the points that visit C^-_{qw} during (ξ_{-1}, λ_q). Since the above cap C^-_{qw} is fully contained in the wedge $W_{qpw} = L^+_{pq} \cap L^-_{pw}$ during that interval, s must leave W_{qpw} before time λ_q (when W_{qpw} shrinks to the single ray $\vec{pq} = \vec{pw}$) through one of the rays $\vec{pw}, \vec{pq}$. We also note that, by condition (S6) (and since $(\xi_{-1}, \lambda_q) \subseteq [\lambda_{wq}, \xi_{wq}]$), $wq \in \mathrm{DT}(P \setminus \{a, p, r, u\})$ throughout (ξ_{-1}, λ_q), so s, which has to leave C^-_{qw} before it leaves W_{qpw}, can do so only through the boundary of $B[p, q, w]$. This results in a co-circularity of p, q, w, s, and is easily seen to imply that s leaves W_{qpw} by crossing one of the rays $\vec{pw}$ or $\vec{pq}$ *outside* the respective edge pw or pq.

In what follows, we assume that s is among the last ℓ points to leave C^-_{qw} during (ξ_{-1}, λ_q). Let t^*_s denote the time of the corresponding co-circularity of p, q, w, s, which occurs when s leaves C^-_{qw} through the boundary of $B[p, q, w]$. Since χ satisfies condition (S6), the opposite cap $C^+_{qw} = B[p, q, w] \cap L^+_{qw}$ contains no points of $P \setminus \{a, r, u\}$ at time t^*_s. (Otherwise, the Delaunayhood of wq would be violated, at time t^*_s, by s and any of these points.) Therefore, the co-circularity at time t^*_s has to be $(\ell - 1)$-shallow in $P \setminus \{a, r, u\}$, and thus $(\ell + 2)$-shallow in P.

Note also that the co-circularity at time t^*_s is red-blue with respect to the edge wq, which is violated right before it by p and s. Lemma 4.1, together with the choice of $s \neq a, p, r, u$, imply that this co-circularity cannot occur during the crossing $(qp, w, \mathcal{H}_\chi = [\lambda_4, \lambda_5])$ (which occurs in $P \setminus \{a, r, u\}$), so $t^*_s < \lambda_4$. (However, condition (S6) does not rule out the violation of wq by p and s during the larger interval $[\lambda_{wq}, \xi_{wq}] \setminus \mathcal{H}$, because the Delaunayhood of wq is assumed to hold there only under the omission of a, r, u, and of p.)

To proceed, we distinguish between two possible subcases. In each of them we manage to dispose of χ by charging it, within one of the arrangements $\mathcal{A}_{pq}, \mathcal{A}_{pw}$, either to $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities, or to a (2ℓ)-shallow collinearity.

Case (b1). At least half of the above points s cross the line L_{pq}, from L^+_{pq} to L^-_{pq}, during (t^*_s, λ_q). (This also includes points s that possibly cross L_{pq} outside the ray $\vec{pq}$, after leaving W_{qpw} through the other ray $\vec{pw}$.) By Condition (S5) (and since $(t^*_s, \lambda_q) \subseteq (\xi_{-1}, \lambda_q) \subseteq [\xi_{pq}, \lambda_{pq}]$), each of these crossings occurs outside pq, within one of the corresponding outer rays of L_{pq}.

For each s we argue, exactly as in Section 5.6, that the points p, q, s are involved during $(t^*_s, \lambda_q) \subseteq (\xi_{-1}, \lambda_q)$ either in a (2ℓ)-shallow collinearity, or in $\Omega(\ell)$ (2ℓ)-shallow co-circularities. That is, as s approaches L_{pq}, the disc $B[p, q, s]$ "swallows" the entire halfplane L^+_{pq}. If the disc, which contains at most $\ell + 2$ points at the beginning of the process, "swallows" at least $\ell - 2$ points in this process, then each of the first $\ell - 2$ resulting co-circularities are (2ℓ)-shallow (in P). Otherwise, the collinearity of q, p, s is (2ℓ)-shallow.

Since s can be chosen in at least $\Omega(\ell)$ different ways, the points p and q are involved during (ξ_{-1}, λ_q)

either in $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities, or in a (2ℓ)-shallow collinearity. In both cases, we charge χ to these events.

Note that each (2ℓ)-shallow event, which occurs in $\mathcal{A}_{pq}$ at some time $t^* \in (\xi_{-1}, \lambda_q)$, can be traced back to $(qp, w, \mathcal{H})$ (and, by Proposition 6.2, also to χ) in at most $O(1)$ possible ways because w is among the first four points to hit the edge pq after time t^*, according to condition (S5). Hence, the above scenario happens for at most $O(\ell^2 N(n/\ell) + \ell n^2 \beta(n))$ special quadruples χ.

Case (b2). At least half of the above points $s \neq a, r, u$ remain in L_{pq}^+ throughout the respective intervals (t_s^*, λ_q). Each of these points must leave $W_{qpw} = L_{pq}^+ \cap L_{pw}^-$, also during (t_s^*, λ_q), through the ray emanating from w in direction $\vec{pw}$, thereby crossing L_{pw} from L_{pw}^- to L_{pw}^+. (Recall that s can cross L_{pw} from L_{pw}^+ to L_{pw}^- at most once, because the triple p, w, s can be collinear at most twice.)

We again fix one of these points s, and use λ_s to denote the corresponding time in (t_s^*, λ_q) when s leaves W_{qpw} through the ray emanating from w in direction $\vec{pw}$. As in the previous case, we conclude that either the collinearity of p, w, s at time λ_s is (2ℓ)-shallow, or the points p, w, s are involved in $\Omega(\ell)$ (2ℓ)-shallow co-circularities during the preceding interval (t_s^*, λ_s). As in Section 5.6, the main challenge is to argue that each of the above (2ℓ)-shallow events, which occur in $\mathcal{A}_{pw}$ during $(t_s^*, \lambda_s] \subseteq (\xi_{-1}, \lambda_q)$, can be traced back to χ in at most $O(k)$ ways.[34]

To show this, let $t^* \in (\xi_{-1}, \lambda_q)$ be the time of a (2ℓ)-shallow collinearity or co-circularity that occurs in $\mathcal{A}_{pw}$. First, we guess the points p and w of χ in $O(1)$ possible ways among the three or four points involved in the event. We next recall that, in the charging scheme of case (b2), each (2ℓ)-shallow co-circularity or collinearity that we charge in $\mathcal{A}_{pw}$ is obtained via some point s, which is also involved in the event, that leaves L_{pw}^- at the respective time λ_s. We therefore guess s among the remaining one or two points involved in the event. To guess the remaining points a and q of χ, we examine all "candidate" special quadruples $\chi' \in \mathcal{G}_{pw}^R$ whose two middle points (p, w) are shared with χ. Recall that each of these quadruples is accompanied by the (p, w)-crossing $(q'p, w, \mathcal{H}' = \mathcal{H}_{\chi'})$, where q' enters L_{pw}^+ at the respective time $\lambda_{q'} \in \mathcal{H}'$. Recall also that χ' is uniquely determined by the choice of q' (as long as p and w remain fixed).

Clearly, with s fixed, it suffices to consider only special quadruples $\chi' = (a', p, w, q')$ in $\mathcal{G}_{pw}^R$ with the following properties: (1) $s \neq a', r', u'$, where r' and u' are the outer points of χ', (2) $\lambda_{q'} > \lambda_s$, and (3) s lies in $L_{pq'}^+$ during the first portion of $\mathcal{H}_{\chi'}$ (before $\lambda_{q'}$). This is because each of these conditions holds for χ and s in the charging scheme of case (b2). For example, (3) follows because case (b1) does not occur for s (and since $t_s^* < \lambda_4$).

If a special quadruple $\chi' = (a', p, w, q') \in \mathcal{G}_{pw}^R$ satisfies the above three conditions (1)–(3), we say that the respective point q' (which uniquely determines χ') is a *candidate* (for being q).

Proposition 6.3 below guarantees that each (2ℓ)-shallow event, which occurs in $\mathcal{A}_{pw}$ at some time $t^* \in (\xi_{-1}, \lambda_q)$, is charged by at most $k + 7$ quadruples in $\chi' \in \mathcal{G}_{pw}^R$, because the corresponding points q' of these quadruples are among the first $k + 7$ candidates to leave L_{pw}^- after time λ_s. Repeating the same argument for each of the $\Omega(\ell)$ possible choice of s shows that at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n)\right)$ special quadruples can fall into case (b2).

Proposition 6.3. *With the above assumptions, the point q is among the first $k + 6$ candidates q' to leave the halfplane L_{pw}^- after λ_s.*

Proof. The fairly technical proof of this proposition is symmetric to the one of Proposition 5.6, so we only briefly review it.

Assume to the contrary that the proposition does not hold (for χ and $s \neq a, r, u$ as above). Hence, we have at least k candidates q' such that $\lambda_s < \lambda_{q'} < \lambda_q$ and $q' \notin \{a, r, u\}$, and such that the first points a', and the outer points r' and u', of their quadruples $\chi' = (a', p, w, q')$ are all distinct from q. (We continue to assume that χ satisfies property (PHR1), introduced in case (a), so the last two restrictions

[34] As in Section 5.6, the multiplicity of the chargings is the major difference between case (b1) and the present case (b2).

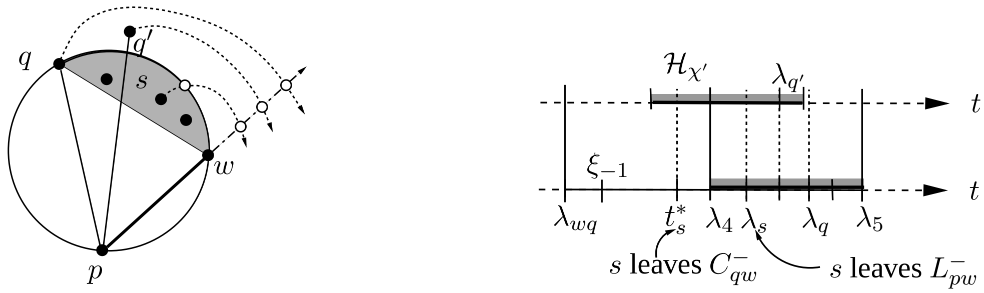

Figure 57: Proposition 6.3. Left: q is among the first $k+7$ candidates q' to leave L_{pw}^- after time λ_s. The figure depicts a point q' lying outside $B[p,q,w]$ at the time t_s^* when s leaves the cap C_{qw}^-. Right: The various critical events occur in the depicted order. Note that λ_s occurs either before λ_4, or in (the first part, preceding λ_q, of) $\mathcal{H} = [\lambda_4, \lambda_5]$.

$q' \neq \{q, a, r\}$ and $q \neq \{a', r', u'\}$ (the latter using (PHR1)) exclude from our consideration at most six candidates $q' \neq q$ together with their quadruples χ'.)

To establish the proposition, we fix a candidate q' and its corresponding quadruple $\chi' = (a', p, w, q')$ (with outer points r' and u'), as above, and argue that the respective interval $\mathcal{H}_{\chi'}$ begins during $(\lambda_{wq}, \lambda_4)$. See Figure 57 (right). Repeating the same argument for the remaining $k-1$ possible choices of q' will imply that the quadruple χ falls into case (a) and we would thereby reach a contradiction.

Indeed, since $\lambda_{q'} < \lambda_q$ (and $q' \neq a, r, u$ and $q \neq a', r', u'$), a suitable variant of Lemma 5.5 shows that the interval $\mathcal{H}_{\chi'}$ begins before $\mathcal{H}_\chi = [\lambda_4, \lambda_5]$. It thus remains to check that $\mathcal{H}_{\chi'}$ begins after λ_{wq}.

If $\mathcal{H}_{\chi'}$ begins after t_s^*, then we are done (as $t_s^* > \lambda_{wq}$). Hence, we may also assume that both times t_s^* and $\lambda_{q'}$ belong to the interval $\mathcal{H}_{\chi'}$. (More precisely, t_s^* belongs to the first part of $\mathcal{H}_{\chi'}$, before $\lambda_{q'}$; this is the situation considered in Figure 57 (right).) This, and the above conditions (2)–(3) (which hold for χ' because q' is a candidate point), imply that q' remains in the halfplanes L_{pw}^-, L_{ps}^- throughout the interval $(t_s^*, \lambda_{q'})$. Therefore, s lies inside $W_{q'pw} = L_{pq'}^+ \cap L_{pw}^-$ throughout the interval (t_s^*, λ_s).

In addition, the standard properties of χ and χ' as 3-restricted special quadruples imply that q' must lie, throughout the longer interval $(t_s^*, \lambda_{q'}) \subseteq \mathcal{H}_{\chi'} \cap (\xi_{-1}, \lambda_q)$, inside the wedge $W_{qpw} = L_{pq}^+ \cap L_{pw}^-$. (Otherwise either the points q', p and w would be collinear more than once during $\mathcal{H}_{\chi'}$, or the edge $q'p$ would be hit by q, or the edge qp would be hit by q'. The first two cases are impossible by the definition of $(q'p, w, \mathcal{H}_{\chi'})$, and the last one is ruled out by condition (S5).)

To recap, we may assume that $\mathcal{H}_{\chi'}$ begins *before* t_s^*, and that the edges pq, pq', ps, and pw appear, at time t_s^*, in this clockwise order around p. To show that $\mathcal{H}_{\chi'}$ begins after λ_{wq}, we distinguish between two possible cases.

(1) If q' lies outside $B[p,w,s] = B[p,q,w]$ at time t_s^* (as depicted in Figure 57 (left)), then the Delaunayhood of pq' is violated, at that very moment, by s and q. Hence, the crossing $(q'p, w, \mathcal{H}_{\chi'})$ (occurring in $\mathrm{DT}(P \setminus \{a', r', u'\})$) has to begin after t_s^*, contrary to our assumptions.

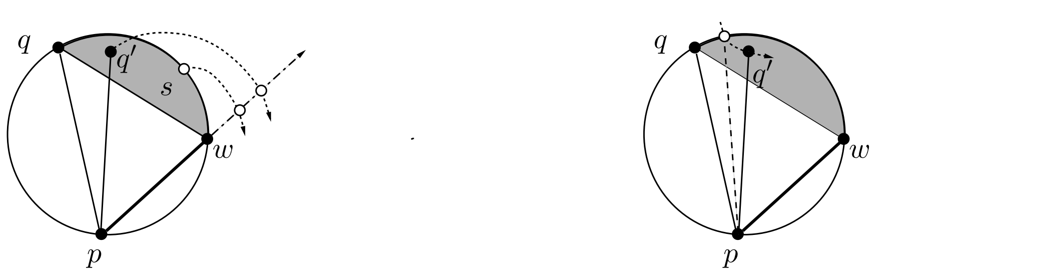

Figure 58: Proof of Proposition 6.3: Left: The scenario where q' lies within $B[p,q,w]$ at time t_s^*. Right: The candidate q' must have entered C_{qw}^-, through $B[p,q,w]$, after λ_{wq} (and before t_s^*, $\mathcal{H} = \mathcal{H}_\chi$ and $\mathcal{H}_{\chi'}$).

(2) If q' lies at time t_s^* within $B[p,q,w]$ (as depicted in Figure 58 (left)), then the interplay between the (p,w)-crossings $(qp,w,\mathcal{H}_\chi)$ and $(q'p,w,\mathcal{H}_{\chi'})$ yields three co-circularities of the points p,w,q,q'. Namely, the last two co-circularities occur during $\mathcal{H}_{\chi'}\setminus\mathcal{H}_\chi$ and $\mathcal{H}_\chi\setminus\mathcal{H}_{\chi'}$. The first co-circularity occurs when q' enters C_{qw}^- after time λ_{wq}, when wq is fully Delaunay, and before t_s^*, when the Delaunayhood of wq is violated by $q'\in C_{qw}^-$ and $p\in B[p,q,w]\cap L_{qw}^+$. (Briefly, this follows since, by conditions (S2) and (S5), none of p,q' can cross wq in the interval $[\lambda_{wq},\lambda_q]$; see the proof of Proposition 6.3 for a fully symmetric argument.) As is easy to check, the points p,w,q and q' are co-circular only once during each of the intervals $\mathcal{H}_{\chi'}$ and $\mathcal{H}_\chi$, so their first co-circularity occurs before $\mathcal{H}_{\chi'}$; see Figure 58 (right). Hence, to allow room for the first co-circularity to occur, $\mathcal{H}_{\chi'}$ has to begin after λ_{wq} also in this case. As noted above, this completes the proof of the proposition. $\square$

Case (c). A total of at least ℓ points $s\in P\setminus A_{pw}$ appear in the cap $C_{pw}^+=B[p,q,w]\cap L_{pw}^+$ at some time during (ξ_{-1},λ_q). Here A_{pw} continues to denote the subset of at most $6\ell+3$ points, including a,r and u, whose removal restores the Delaunayhood of pw throughout the interval $[\lambda_{wq},\xi_{wq}]$. (Recall that A_{pw} was obtained by applying Theorem 2.2 in $\mathcal{A}_{pw}$, after ruling out case (a).)

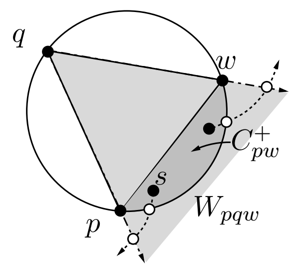

Figure 59: Case (c). A total of at least ℓ points $s\in P\setminus A_{pw}$ appear in the cap C_{pw}^+ during (ξ_{-1},λ_q). Each of them must leave the cap C_{pw}^+ (through the boundary of $B[p,q,w]$) and then exit the wedge W_{pqw} (through one of the rays $\vec{qp},\vec{qw}$, outside the respective edges pq and wq) before time λ_q.

Clearly, C_{pw}^+ is contained in the wedge $W_{pqw}=L_{pq}^+\cap L_{wq}^-$, which shrinks at time λ_q to the ray $\vec{qp}=\vec{qw}$. Hence, each of these points s has to leave C_{pw}^+ and W_{pqw} (in this order) before time λ_q. Furthermore, s can leave C_{pw}^+ only through the boundary of $B[p,q,w]$, at a co-circularity of p,q,w,s. (Otherwise s would have to hit pw and, therefore, belong to A_{pw}.) In addition, s can leave W_{pqw} only through one of the rays $\vec{qp}$ and $\vec{qw}$ (outside the respective segments qp,qw). See Figure 59.

As in the previous case (b), we restrict our attention to the last ℓ such points s of $P\setminus A_{pw}$ to leave C_{pw}^+ during (ξ_{-1},λ_q), and use t_s^* to denote the time of the respective co-circularity. Clearly, the opposite cap $C_{pw}^-=B[p,q,w]\cap L_{pw}^-$ contains then no points of $P\setminus A_{pw}$. Indeed, otherwise the Delaunayhood of pw would be violated by s and any one of these points (contrary to our assumption that $pw\in \mathrm{DT}(P\setminus A_{pw})$ throughout $[\lambda_{wq},\xi_{wq}]\supset(\xi_{-1},\lambda_q)$). Hence, the resulting co-circularity of p,q,w,s at time t_s^* is $(7\ell+2)$-shallow in P, because, at the time of co-circularity, the circumdisc $B[p,q,w]=B[p,s,w]$ can contain in its interior at most the $6\ell+3$ points of A_{pw} and at most $\ell-1$ points of $P\setminus A_{pw}$.

Case (c1). If at least half of the above points s cross the line L_{pq} (from L_{pq}^+ to L_{pq}^-) during their respective intervals (t_s^*,λ_q), then we argue exactly as in subcase (b1). Namely, we fix one of these points s and notice that the points p,q,s are involved, during (t_s^*,λ_q), either in an (8ℓ)-shallow common collinearity, or in $\Omega(\ell)$ (8ℓ)-shallow co-circularities, occuring within the whole set P. That is, as s approaches L_{pq}, the disc $B[p,q,s]$ "swallows" the entire halfplane L_{pq}^+. If the disc, which contains at most $7\ell+2$ points at the beginning of the process, "swallows" at least $\ell-2$ points in this process, then each of the first $\ell-2$ resulting co-circularities are (8ℓ)-shallow (in P). Otherwise, the collinearity of q,p,s is (8ℓ)-shallow.

We thus repeat the above argument for each of the (at least) $\ell/2$ possible choices of s and charge χ within $\mathcal{A}_{pq}$ (via $(qp,w,\mathcal{H})$) either to $\Omega(\ell^2)$ (8ℓ)-shallow co-circularities, or to an (8ℓ)-shallow collinearity. As in case (b1), each (8ℓ)-shallow collinearity or co-circularity occurs during (ξ_{-1}, λ_q), and involves p and q, so it is charged by at most $O(1)$ special quadruples χ (because χ is uniquely determined by (p, q, w) and w is among the first four points to hit pq after the respective time t^* of the event, because of condition (S5)).

Case (c2). We may assume, then, that at least half of the above points s leave W_{pqw} through the ray $\vec{qw}$ (outside the segment qw). For each of these points s, a symmetric variant of the argument in case (c1) implies that the points q, w, s are involved during (t_s^*, λ_q) either in an (8ℓ)-shallow collinearity, or in $\Omega(\ell)$ (8ℓ)-shallow co-circularities. As before, we repeat the above argument for the (at least) $\ell/2$ eligible choices of s and charge χ, within $\mathcal{A}_{wq}$, either to $\Omega(\ell^2)$ (8ℓ)-shallow co-circularities or to an (8ℓ)-shallow collinearity.

We claim that each of the resulting (8ℓ)-shallow events, which occur in $\mathcal{A}_{wq}$ during (ξ_{-1}, λ_q), can be traced back to χ in at most $O(1)$ possible ways. Indeed, fix any of the above events, which occurs in $\mathcal{A}_{wq}$ at some time $t^* \in (\xi_{-1}, \lambda_q)$. We first guess w and q in $O(1)$ possible ways among the three or four points involved in the event. To guess the point a (which would then uniquely determine $(wa, q, \mathcal{J}_u)$ and thereby also χ), we consider all special (w, q)-crossings $(wa', q, \mathcal{J}_{u'})$ (in $\mathcal{F}$) and recall that, according to conditions (S2) and (S6), at most $O(1)$ such crossings can begin during $[\lambda_{wq}, \lambda_2)$ or end during $(\lambda_3, \xi_{wq}]$. Notice also that the interval $[\lambda_{wq}, \xi_{wq}]$, which covers (ξ_{-1}, λ_q), is the union of $[\lambda_{wq}, \lambda_2)$, $\mathcal{J}_u = [\lambda_2, \lambda_3]$, and $(\lambda_3, \xi_{wq}]$.

To guess a (based on t^*, q and w), we distinguish between two possible situations.
(i) If t^* belongs to $(\lambda_3, \lambda_q) \subseteq (\lambda_3, \xi_{wq}]$ then $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$ is among the last three special clockwise (w, q)-crossings to end before t^*, because χ satisfies condition (S6). See Figure 60 (left).

Figure 60: Case (c2): Guessing a based on t^*, w and q. Left: If $t^* \in (\lambda_3, \lambda_q)$, then $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$ is among the last three special clockwise (w, q)-crossings to end before t^*. Right: If $t^* \in (\xi_{-1}, \lambda_3]$, then $(wa, q, \mathcal{J}_u)$ is among the first $O(1)$ special clockwise (w, q)-crossings to end after t^*. Any other such (w, q)-crossing $(wa', q, \mathcal{J}_{u'})$ (with $u' \neq p$), that ends in $(t^*, \lambda_3) \subset (\xi_{-1}, \lambda_3)$, must begin after ξ_{-1} (and, therefore, in $[\lambda_{wq}, \lambda_2)$).

(ii) If t^* belongs to the interval $(\xi_{-1}, \lambda_3]$, which is contained in $[\lambda_{wq}, \lambda_2) \cup \mathcal{J}_u$, then we resort to a more subtle argument, in which we show that $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$ is among the first $O(1)$ special clockwise (w, q)-crossings to end after t^*. See Figure 60 (right).

Our goal is to bound the number of special clockwise (w, q)-crossings that end in (t^*, λ_3). Note that the preliminary pruning (peformed before the definition of special quadruples) guarantees that each of these crossings $(wa', q, \mathcal{J}_{u'})$ satisfies $a' \neq u$ and $u' \neq a$, and therefore begins before $\mathcal{J}_u = [\lambda_2, \lambda_3]$ (by Lemma 5.5). Furthermore, note that we have $u' = p$ for at most one of these crossings $(wa', q, \mathcal{J}_{u'})$, because each of them is uniquely determined by the respective triple w, q, u'. We claim that each of the remaining (w, q)-crossings $(wa', q, \mathcal{J}_{u'})$ under consideration (satisfying also $u' \neq p$) must begin in $(\xi_{-1}, \lambda_2) \subseteq [\lambda_{wq}, \lambda_2)$. This, together with condition (S2), implies that their number is $O(1)$ too.

To see this final claim, note that if a (w, q)-crossing $(wa', q, \mathcal{J}_{u'})$, as above, begins before ξ_{-1}, then its respective interval $\mathcal{J}_{u'}$ contains the time ξ_{-1} (because it ends after $t^* > \xi_{-1}$), right after which the Delaunayhood of wq is violated by p and a. This, however, is impossible because, by Lemma 4.1, wq is Delaunay throughout $\mathcal{J}_{u'}$ in $P \setminus \{u'\}$, and $u' \neq p, a$.

To recap, in each of the cases (c1) and (c2) we charge χ either to $\Omega(\ell^2)$ (8ℓ)-shallow co-circularities, or to an (8ℓ)-shallow collinearity, which occur in one of the arrangements $\mathcal{A}_{pq}, A_{wq}$ during the interval

(ξ_{-1}, λ_q). Furthermore, each (8ℓ)-shallow event is charged by at most $O(1)$ special quadruples. Hence, at most $O\left(\ell^2 N(n/\ell) + \ell n^2 \beta(n)\right)$ special quadruples χ fall into case (c).

Case (d). Assume that none of the preceding cases occurs. In particular, there is a subset A_{pw} of at most $6\ell + 3$ points (including a, r and u) whose removal restores the Delaunayhood of pw throughout the interval $[\lambda_{wq}, \xi_{wq}]$. Furthermore, a total of fewer than ℓ points of $P \setminus \{a, r, u\}$ ever appear in the cap C^-_{qw} during (ξ_{-1}, λ_q), and a total of fewer than ℓ points of $P \setminus A_{pw}$ points ever appear in the cap C^+_{pw} during that interval.[35]

In this last remaining scenario, we finally consider the interplay between the special quadruple χ under consideration and the ordinary Delaunay quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$, which corresponds to the *first* special (a, q)-crossing $(pa, q, \mathcal{I}_r)$ of χ. At the end of this section, we shall charge χ to the terminal quadruple $\varrho = (p, q, r, w)$, which is composed of the edge pq, and of the two points r and w that cross pq in opposite directions. (The outer point u of the second special (a, q)-crossing $(wa, q, \mathcal{J}_u)$ is not used for right quadruples; it will be used in the mostly symmetric analysis of left quadruples, given in Section 6.6.)

Before charging χ to the above terminal quadruple ϱ, we enforce a Delaunay crossing of one of the edges pr, qr by the point w. In addition, we shall have to enforce two more crossings performed by the points of ϱ in order to ensure that at least two of the resulting *five* crossings are performed by the same sub-triple of ϱ (so as to allow us to apply our cornerstone Lemma 4.5 and thereby obtain a quadratic bound on the number of such quadruples).

To facilitate the forthcoming analysis, we first establish several auxiliary claims.

Lemma 6.4. *With the above assumptions, a total of at most $8\ell + 1$ points of P appear in the cap $C^+_{pq} = B[p, q, w] \cap L^+_{pq}$ during (ξ_{-1}, λ_q).*

Proof. Refer to Figure 61. Recall that the motion of $B[p, q, w]$ is continuous throughout (ξ_{-1}, λ_q). Notice that the above cap $C^+_{pq} = B[p, q, w] \cap L^+_{pq}$ (which contains w on its boundary) is empty right before time λ_q, when the edge pq is crossed by w. Hence, any point s that appears in this cap during (ξ_{-1}, λ_q) has to leave it before λ_q. Furthermore, condition (S5) (together with the inclusion $(\xi_{-1}, \lambda_q) \subseteq [\xi_{pq}, \lambda_{pq}]$) implies that s cannot escape C^+_{pq} through the edge pq, unless it is equal to one of a, r, u. Therefore, any such point $s \neq a, r, u$ has to leave C^+_{pq} through one of the circular arcs bounding the earlier caps C^-_{qw}, C^+_{pw}, so it must first appear in one of the caps C^-_{qw} or C^+_{pw}. Since cases (b) and (c) have been ruled out, and since a, r, u belong to the set A_{pw}, the overall number of such points cannot exceed $(\ell - 1) + (\ell - 1) + (6\ell + 3) = 8\ell + 1$. $\square$

We next consider the ordinary quadruple $\sigma = (p, q, a, r)$ in $\mathcal{F}$, which corresponds to the first special crossing $(pa, q, \mathcal{I}_r)$ of χ. Refer to Figure 62. We continue to denote the two Delaunay crossings of σ by $(pq, r, I = [t_0, t_1])$ and $(pa, r, J = [t_2, t_3])$. Recall that the points of σ are co-circular at times $\zeta_0 \in I \setminus J, \zeta_1 \in J \setminus I$ and $\zeta_2 > t_3$. By condition (Q3) on σ, the last two co-circularities of σ (at times ζ_1 and ζ_2) have the same order type, and the Delaunayhood of rq is violated by $p \in L^-_{rq}$ and $a \in L^+_{rq}$ throughout the interval (ζ_1, ζ_2) (see Figure 62 (left)). Therefore, the Delaunayhood of pa is violated right after time ζ_2 by r and q.

Remark: Note that σ and χ have "opposite" topological behaviour, in the sense that the additional co-circularity of σ (outside I and J) occurs at time ζ_2, after the respective second interval J of σ, whereas the corresponding additional co-circularity of χ (outside $\mathcal{I}_r$ and $\mathcal{J}_u$) occurs at time ξ_{-1}, before the respective first interval $\mathcal{I}$ of χ.

[35]Note the built-in asymmetry between qw and pw in the analysis: The former is almost Delaunay in the interval $[\lambda_{wq}, \xi_{wq}]$ (and Delaunay at both endpoints λ_{wq}, ξ_{wq}]), whereas the latter becomes Delaunay there only after the removal of A_{pw} (which includes a, p, r, u).

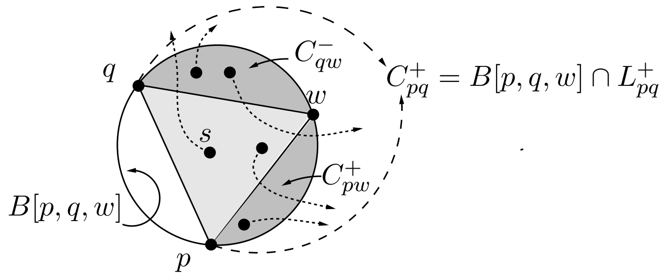

Figure 61: Lemma 6.4: A total of at most $8\ell+1$ points s of P appear in the cap $C^+_{pq} = B[p,q,w] \cap L^+_{pq}$ (consisting of all the shaded portions) during (ξ_{-1}, λ_q). All of them must leave C^+_{pq} before λ_q. None of these points s can leave C^+_{pq} through pq, unless it is one of a, r, u.

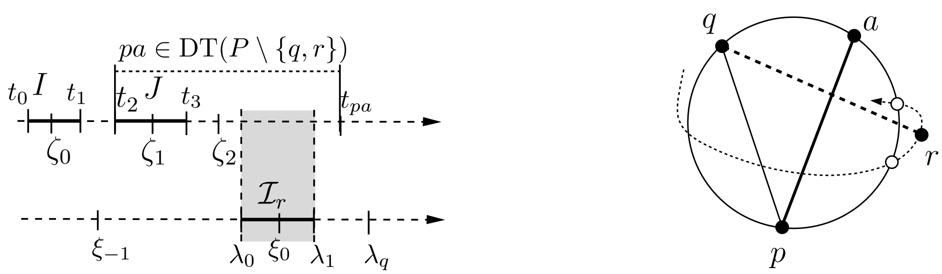

Figure 62: The (regular) clockwise quadruple $\sigma = (p,q,a,r)$ of $(pa, q, \mathcal{I}_r)$ is composed of two (p,r)-crossings $(pq, r, I = [t_0, t_1])$, $(pa, r, J = [t_2, t_3])$. The points p, q, a, r are co-circular at times $\zeta_0 \in I \setminus J, \zeta_1 \in J \setminus I$, and $\zeta_2 > t_3$ (left). The last two co-circularities have the same order type, and the Delaunayhood of rq is violated by p and a throughout (ζ_1, ζ_2) (right).

By condition (Q7), the edge pa re-enters $\mathrm{DT}(P)$ at some time $t_{pa} \geq \zeta_2 > t_2$. Furthermore, pa belongs to $\mathrm{DT}(P \setminus \{r,q\})$ throughout the interval $[t_2, t_{pa}] = J \cup [t_3, t_{pa}]$, which covers J (including $\zeta_1 \in J \setminus I$) and ζ_2. Moreover, we recall that (using the Delaunayhood of pa at time t_3, and the extremality of ζ_2, via Lemma 3.1), q crosses pa from L^-_{pa} to L^+_{pa} during $(t_3, t_{pa}]$. As argued in Section 5.6, this yields the Delaunay crossing $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ in $P \setminus \{r\}$ as the unique special crossing of σ, with $\mathcal{I}_r \subset (t_3, t_{pa}]$.

To conclude, the second crossing $(pa, r, J = [t_2, t_3])$ of σ and the first crossing $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ of χ occur during disjoint intervals and in this order.[36]

Finally, by condition (Q8), the edge pq belongs to $\mathrm{DT}(P \setminus \{r,a\})$ throughout the interval $[t_0, \lambda_1] = [I, \mathcal{I}_r] (= \mathsf{conv}(I \cup \mathcal{I}_r))$. Therefore, the almost-Delaunayhood of pq extends from $[\xi_{pq}, \lambda_{pq}]$ to the potentially larger interval $[t_0, \lambda_{pq}]$ (assuming $\xi_{pq} > t_0$, that is, $I = [t_0, t_1]$ is not contained in $[\xi_{pq}, \lambda_{pq}]$).

The following claim is crucial for understanding the interplay between σ and χ.

Lemma 6.5. *With the above assumptions, we have* $\zeta_1 \in (\xi_{-1}, \xi_0)$.

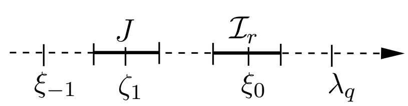

Figure 63: Lemma 6.5 claims that $\zeta_1 \in (\xi_{-1}, \xi_0)$.

[36]Note that, even though q hits pa after ζ_2, during the above special crossing, it is not known whether the last co-circularity ζ_2 of σ occurs in $\mathcal{I}_r$ or beforehand, in $(t_3, \lambda_0]$.

Proof. The inequality $\zeta_1 < \xi_0$ follows because ζ_1 occurs during the second crossing (pa, r, J) of σ, whereas ξ_0 occurs during the first special crossing $(pa, q, \mathcal{I}_r)$ of χ (which begins after J). See Figure 63 (left) and Figure 63.

To establish the inequality $\zeta_1 > \xi_{-1}$, let us assume for a contradiction that $\zeta_1 \le \xi_{-1}$; see Figure 64. Since $\sigma = (p, q, a, r)$ belongs to the refined family $\mathcal{F}$ (and, therefore, satisfies condition (Q3)), its point q remains in $B[p, a, r] \cap L^-_{pa}$ after r enters L^+_{pa} during $J = [t_2, t_3]$ and until time $\zeta_1 \in J$ (when q leaves the cap $B[p, a, r] \cap L^-_{pa}$). Also note that, with the above assumption that $\zeta_1 < \xi_{-1}$, the point q cannot leave L^-_{pa} during (ζ_1, ξ_{-1}). Indeed, q lies in L^-_{pa} at both endpoints of that interval, because the quadruples σ and χ satisfy the respective conditions (Q3) and (S3a), and it can enter the halfplane L^+_{pa} only once (which occurs during $\mathcal{I}_r$ and after ξ_{-1}).

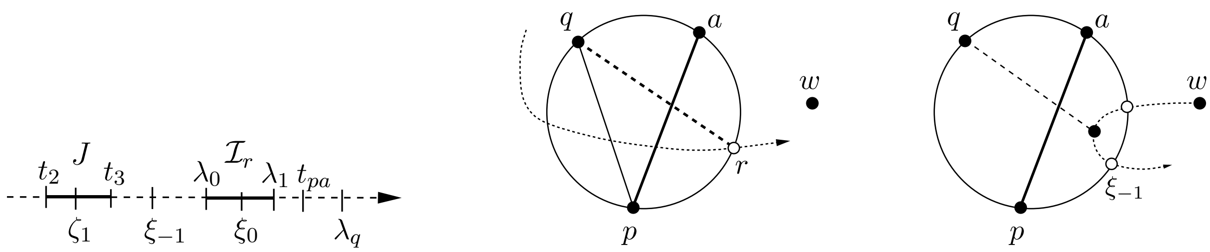

Figure 64: Proof of Lemma 6.5. If $\zeta_1 < \xi_{-1}$ (left) then w has to enter $B[p, q, a] \cap L^+_{pa}$, which is empty at time ζ_1 (center), before leaving it at time ξ_{-1} (right). By Condition (Q7), w can enter $B[p, q, a] \cap L^+_{pa}$ during $(\zeta_1, \xi_{-1}) \subset (t_2, t_{pa})$ only through the boundary of $B[p, q, a]$.

The above reasoning implies that the motion of $B[p, q, a]$ is continuous throughout (ζ_1, ξ_{-1}). Furthermore, w lies outside the cap $B[p, q, a] \cap L^+_{pa}$ at time ζ_1, for otherwise the Delaunayhood of pa would be violated by q and w (which cannot happen during the interval J, where pa belongs to $\mathrm{DT}(P \setminus \{r\})$); see Figure 64 (center). By condition (S3a), w leaves the cap $B[p, q, a] \cap L^+_{pa}$ at time ξ_{-1}. Therefore, w must have previously entered that cap, in the interval (ζ_1, ξ_{-1}). Note that, since $\xi_{-1} < \lambda_0$, the latter interval is contained in (ζ_1, t_{pa}), where t_{pa} denotes the first time after ζ_1 and ζ_2 when pa again belongs to $\mathrm{DT}(P)$.

Since σ satisfies condition (Q7), w cannot enter $B[p, q, a] \cap L^+_{pa}$ during $(\zeta_1, \xi_{-1}) \subseteq (t_2, \lambda_1) \subseteq (t_2, t_{pa})$ through the edge pa. Furthermore, w cannot enter $B[p, q, a] \cap L^+_{pa}$ during that interval through the boundary of $B[p, q, a]$, as that would cause a forbidden fourth co-circularity of p, q, a, w; see Figure 64 (right). Hence, we have reached a contradiction, and the claim follows. $\square$

By Lemma 4.1, none of the co-circularities ξ_{-1}, ξ_0 can occur during J, so we have $J \subset (\xi_{-1}, \xi_0)$. This, combined with the properties (S1)–(S3a) of χ, implies that $(\xi_{pq} <)\xi_{-1} < t_2 < \zeta_1 < t_3 < \lambda_0 < \xi_0 < \lambda_1 < \lambda_q (< \lambda_{pq})$. See Figure 65 (left).

Since σ satisfies condition (Q3), r cannot return to L^-_{pq} (after leaving it during I) before time ζ_1 (when r leaves the cap $B[p, q, a] \cap L^+_{pq}$), for otherwise the triple p, q, r would be collinear at least three times.

Furthermore, if r re-enters L^-_{pq} through pq during the subsequent interval $(\zeta_1, \lambda_{pq}]$, then the edge pq undergoes two Delaunay crossings by r within the triangulation $\mathrm{DT}(P \setminus \{a, u, w\})$. Indeed, Lemma 6.5 implies that (ζ_1, λ_{pq}) is contained in $[\xi_{-1}, \lambda_{pq}] \subset [\xi_{pq}, \lambda_{pq}]$, and the edge pq belongs to $\mathrm{DT}(P \setminus \{a, r, u, w\})$ throughout the latter interval by condition (S5) (in addition to its being Delaunay at the endpoints ξ_{pq} and λ_{pq}). By Lemma 4.5 and Proposition 5.2, this happens for at most $O(n^2)$ special quadruples χ.

To conclude, ignoring the favourable quadruples just considered, we may assume that the above scenario does not occur, so r does not cross pq in the interval $(t_1, \lambda_{pq}]$. (However, r can still return to L^-_{pq} during $(t_1, \lambda_{pq}]$, or, more precisely, during $(\zeta_2, \lambda_{pq}]$, by crossing one of the outer rays of L_{pq}, outside pq.) See Figure 65 (right).

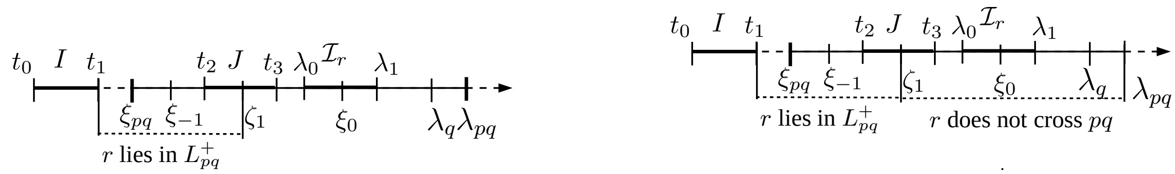

Figure 65: Left: The setup implied by Lemma 6.5. We have $\xi_{pq} < \xi_{-1} < \zeta_1 < \xi_0 < \lambda_q < \lambda_{pq}$, and $J \subset (\xi_{-1}, \xi_0)$. The point r remains in L^+_{pq} throughout (t_1, ζ_1). Right: If r were to hit pq also in (ζ_1, λ_{pq}), then pq would undergo two Delaunay crossings by r within $\mathrm{DT}(P \setminus \{a, w, u\})$. Hence, we can assume that no such collinearity occurs.

The three co-circularities of p, q, r, w**.** We now argue that the four points p, q, r, w are involved in exactly three co-circularities, and characterize the order types of these co-circularities. First, recall that one such co-circularity occurs at some time $\delta_0 \in I$, according to Lemma 4.4. Since this co-circularity is induced by the crossing of pq by r, it is red-blue with respect to pq and to rw. Moreover, as will follow from the subsequent analysis, this is the first co-circularity of this quadruple; see Figure 66.

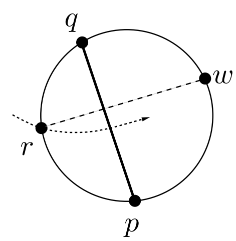

Figure 66: The co-circularity of p, q, r, w occurring at some time $\delta_0 \in I$. It is red-blue with respect to the edges pq and rw.

To obtain the second co-circularity of p, q, r, w, we recall that (as reviewed at the beginning of this section, and depicted in Figure 53 (bottom)) the Delaunayhood of wq is violated by $p \in L^-_{wq}$ and $a \in L^+_{wq}$ throughout the interval (ξ_{-1}, ξ_0), and the order type of p, q, w remains fixed (i.e, w lies in L^+_{pq}) throughout the larger interval (ξ_{-1}, λ_q).

By Lemma 6.5, the interval (ξ_{-1}, ξ_0) contains ζ_1, so a lies at that time in the cap $C^-_{qw} \subset C^+_{pq}$ (after it enters C^-_{qw} at time ξ_{-1}, and before escaping it at time ξ_0). Since the points p, q, a, r are involved at time ζ_1 in a red-red co-circularity with respect to pq (as prescribed by condition (Q3) on σ), both a and r lie at time ζ_1 within the cap $C^+_{pq} = B[p, q, w] \cap L^+_{pq}$; see Figure 67 (left).

Since w remains in L^+_{pq} throughout the longer interval (ξ_{-1}, λ_q), the four points p, q, r, w are involved, during (ζ_1, λ_q), in a co-circularity, which occurs when r leaves the above cap $B[p, q, w] \cap L^+_{pq}$. (Otherwise r would have to escape $B[p, q, w] \cap L^+_{pq}$ through the interior of pq before this cap shrinks to pq at time λ_q, which cannot happen during $(t_1, \lambda_q] \subseteq (t_1, \lambda_{pq}]$ by condition (Q8).) Clearly, this co-circularity is red-red with respect to the edge pq, and occurs after I and between $\zeta_1 \in J \setminus I$ and λ_q. We denote by δ_1 the time of the *first* such co-circularity event in (ζ_1, λ_q), at which r leaves $B[p, q, w] \cap L^+_{pq}$. (As will soon turn out, this is the second co-circularity of p, q, r, w.)

Remark. We again emphasize that $\delta_1 \in (\zeta_1, \lambda_q) \subset (\xi_{-1}, \lambda_q)$, and that r remains in the cap $B[p, q, w] \cap L^+_{pq}$ throughout the interval $[\zeta_1, \delta_1)$. (However, the order between δ_1 and ξ_0 is not known, and is immaterial for our analysis.)

We next claim that the points p, q, r, w are involved in a third co-circularity, red-blue with respect to pq, at some time $\delta_2 \in (\delta_1, \lambda_{pq}]$. Notice that the desired co-circularity cannot be obtained by simply applying Lemma 4.4 to the crossing $(qp, w, \mathcal{H})$, because it is defined only with respect to the reduced point set $P \setminus \{a, r, u\}$.

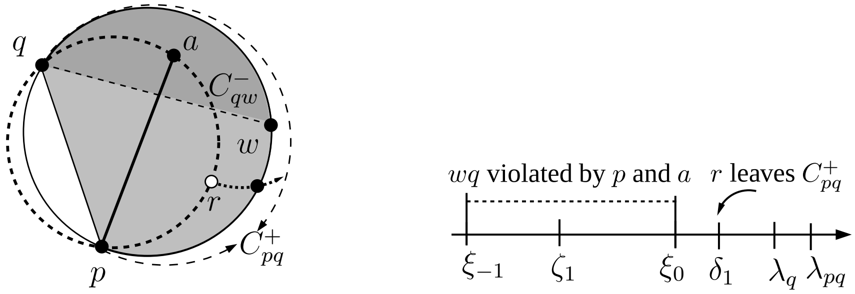

Figure 67: Obtaining the second co-circularity δ_1 of p, q, r, w. The co-circularity of p, q, a, r at time ζ_1 is red-red with respect to pq, and belongs to the interval (ξ_{-1}, ξ_0), during which a lies in $C^-_{qw} (\subset C^+_{pq})$. Hence, r lies at that time within the cap $C^+_{pq} = B[p, q, w] \cap L^+_{pq}$, so δ_1 necessarily occurs in (ζ_1, λ_q), when r escapes the above cap C^+_{pq} (without crossing pq). This is also a red-red co-circularity with respect to pq.

Instead, we consider the four-point triangulation $\mathrm{DT}(\{p, q, r, w\})$, and observe that the edge qp undergoes there a Delaunay crossing by w, which takes place during some sub-interval of $(\delta_1, \lambda_{pq}]$ that contains λ_q (the time of the actual collinearity of the three points). Indeed, pq is Delaunay in $\{p, q, r, w\}$ at times $\delta_1 < \lambda_q$ and $\lambda_{pq} \geq \lambda_q > \delta_1$, and it is Delaunay in $\{p, q, r\}$ throughout $(\delta_1, \lambda_{pq}]$ (because r is assumed not to cross pq in the even larger interval $(t_1, \lambda_{pq}]$).

Furthermore, the above crossing in $\mathrm{DT}(\{p, q, r, w\})$ must be single. Indeed, since w lies in $C^-_{wq} \subset L^+_{pq}$ throughout the interval (ξ_{-1}, ξ_0) which contains ζ_1, it has to remain in L^+_{pq} throughout $[\zeta_1, \lambda_q] \supset [\delta_1, \lambda_q]$ (or, else, w would cross L_{pq} three times). Furthermore, w does not cross pq again in $(\lambda_q, \lambda_{pq}]$ (by condition (S3b)). We hence apply Lemma 4.4 to this single crossing, which gives us the desired third co-circularity (see Figure 68).

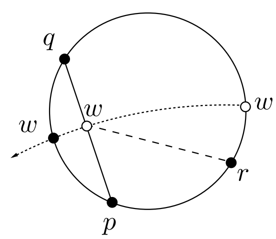

Figure 68: The third co-circularity of p, q, r, w occurs at some time $\delta_2 \in (\delta_1, \lambda_{pq}]$, and is red-blue with respect to the edges pq and rw. This co-circularity is part of a Delaunay crossing of qp by w, which occurs within the four-point triangulation $\mathrm{DT}(\{p, q, r, w\})$, during some subinterval of $(\delta_1, \lambda_{pq}]$ that contains λ_q.

To conclude, the four points p, q, r, w are involved in three co-circularities, which occur at times $\delta_0 \in I = [t_0, t_1]$, $\delta_1 \in (\zeta_1, \lambda_q) (\subset (\xi_{-1}, \lambda_q))$, and $\delta_2 \in (\delta_1, \lambda_{pq})$. The two extremal co-circularities (which occur at times δ_0 and δ_2) are red-blue with respect to the edges pq and wr, and thus monochromatic with respect to pr, qr, pw, qw. The middle co-circularity (at time δ_1) is red-red with respect to pq.[37]

We are now ready to establish the following important consequence of Lemma 6.4.

Lemma 6.6. *With the above assumptions, at most $8\ell+1$ clockwise (Delaunay) (p, r)-crossings (pq', r, I') in $\mathcal{F}$, and at most $8\ell + 1$ counterclockwise (Delaunay) (q, r)-crossings $(p'q, r, I')$ in $\mathcal{F}$, can end in the interval (t_1, δ_1).*

Recall that an (ordinary) Delaunay crossing is in $\mathcal{F}$ if it is either the first or the second crossings of some Delaunay quadruple in $\mathcal{F}$. In Section 5 we have already enforced comparable restrictions (via

[37]This alternation in the order type is crucial of the forthcoming analysis.

conditions (Q2) and (Q4)), which imply that no clockwise (p,r)-crossings (pq',r,I') in $\mathcal{F}$, and no counterclockwise (q,r)-crossings $(p'q,r,I')$ in $\mathcal{F}$, end after time t_1 and before time $t_{rq} > t_3(> \zeta_1)$, which is first such time after ζ_1 when the edge rq belongs to $\mathrm{DT}(P)$. See Figure 69. (In addition, conditions (Q2) and (Q3) imply that rq belongs to $\mathrm{DT}(P \setminus \{a,p\})$ throughout $(t_1, t_{rq}]$, and neither a nor p can hit rq in that interval.) Unfortunately, the order of t_{rq} and δ_1 is not known, so condition (Q4) does not immediately imply the above property.

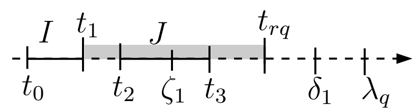

Figure 69: Preparing for the proof of Lemma 6.6. By conditions (Q2) and (Q4), no clockwise (p,r)-crossings, and no counterclockwise (q,r)-crossings in $\mathcal{F}$ end in the shaded interval between t_1 and $t_{rq} > t_3(> \zeta_1)$, where t_{rq} is the first such time after ζ_1 when rq belongs to $\mathrm{DT}(P)$. Unfortunately, the order of δ_1 and t_{rq} is not known.

Proof of Lemma 6.6. We first consider clockwise (p,r)-crossings. Let (pq',r,I') be such a Delaunay crossing that ends in (t_1,δ_1). Note that the point q' has to be distinct from a (for, otherwise, (pq',r,I') would co-incide with (pa,r,J)), and that the points p,q,q',r form an (ordinary, not necessarily consecutive quadruple) clockwise quadruple. Recall also that r remains in L_{pq}^+ after entering that halfplane during $I = [t_0,t_1]$ and until time δ_1 (when r escapes $C_{pq}^+ = B[p,q,w] \cap L_{pq}^+$). In particular, q lies in $L_{pr}^- = L_{pq'}^-$ when r enters $L_{pq'}^+$ (during I'). Hence, the points p,q,r,q' are involved in a co-circularity at some time $\zeta' \in I' \setminus I$, right after which the Delaunayhood of rq is violated by p and q'. See Figure 70 (left).

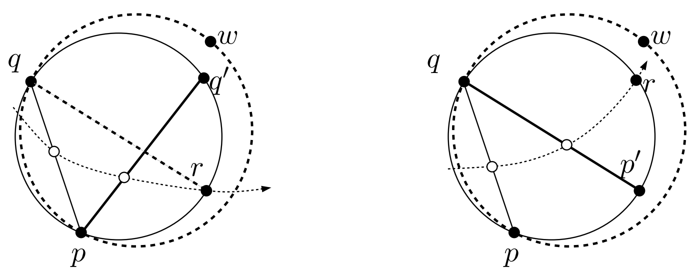

Figure 70: Proof of Lemma 6.6. Left: (pq',r,I') is a clockwise (p,r)-crossing that ends (t_1,δ_1). The points p,q,r,p' are co-circular at some time $\zeta' \in I' \setminus I$. If ζ' occurs in (ζ_1,δ_1), then p' lies in $C_{pq}^+ = B[p,q,w] \cap L_{pq}^+$ at that moment. Right: $(p'q,r,I')$ is a counterclockwise (q,r)-crossing that occurs within $(\zeta_1,\delta_1]$. The points p,p',q,r are co-circular at some time $\zeta' \in I' \setminus I$, when both r and p' lie inside C_{pq}^+.

We first argue that ζ' cannot occur before ζ_1. Indeed, otherwise, applying Lemma 3.1 for the edge rq, from time ζ', would imply that at least one of the following events must occur between ζ' and t_{rq} (which is the first time after ζ_1 when rq belongs $\mathrm{DT}(P)$): (1) q' hits rq, (2) p hits rq, or (3) the four points p,q,q',r are involved in an additional co-circularity of the same order type.

However, cases (1), (2) are impossible by conditions (Q2) and (Q3) on σ (using that $\zeta_1 < t_{rq}$). Moreover, the co-circularity in (3) can occur only after the end of both I and I' (because p,q,q' and r form a regular clockwise quadruple; see Section 4.1), in which case (pq',r,I') has to end before t_{rq}, contrary to condition (Q4) on σ. Hence, ζ' must occur after ζ_1.

We may thus assume that ζ' belongs to the interval (ζ_1,δ_1) which, by Lemma 6.5, is contained in (ξ_{-1},λ_q), so both q' and r' lie at time ζ' within the cap $C_{pq}^+ = B[p,q,w] \cap L_{pq}^+$. According to Lemma 6.4, the overall number of such points q' is at most $8\ell + 1$.

The treatment of counterclockwise (q,r)-crossings (also in $\mathcal{F}$) is similar (but somewhat simpler). Indeed, let $(p'q,r,I')$ be such a crossing. Condition (Q2) implies that it cannot end in the interval $(t_1,\zeta_1]$

(because ζ_1 belongs to $J \setminus I \subset (t_1, \lambda_{rq})$). Furthermore, Lemma 4.1 implies that any counterclockwise (q, r)-crossing $(p'q, r, I')$ that ends after ζ_1 has to begin also after ζ_1. (Otherwise, its respective interval I' would contain the time ζ_1 of a red-blue co-circularity with respect to rq, contrary to the Delaunayhood of rq during I'.) We consider the co-circularity of p, p', q, r, which must occur at some time $\zeta' \in I' \setminus I$ and notice, as in the previous case, that both r and p' lie at that moment in the cap C^+_{pq}; see Figure 70 (right). Therefore, the overall number of such points p' does not exceed $8\ell + 1$. □

Cases (d1) and (d2): Overview. To proceed, we distinguish between two possible subcases. In subcase (d1), we assume that the middle co-circularity, which occurs at time δ_1, is red-blue with respect to the edges pr and wq (see Figure 71 (left)), and then use it to enforce (via Lemma 3.1) the following two additional crossings: (i) a Delaunay crossing of pr by at least one the points w, q, and (ii) a Delaunay crossing of wq by at least one of the points p, r. (For the second crossing, it will suffice to argue that wq is hit by one of the points p, r in the interval $[\lambda_{wq}, \lambda_{pq}] \subseteq [\lambda_{wq}, \xi_{wq}]$.) However, this can easily be established by applying Lemma 3.1 to wq *backwards* from the second co-circularity $\delta_1 \in [\lambda_{wq}, \lambda_{pq}]$ of p, q, r, w.)

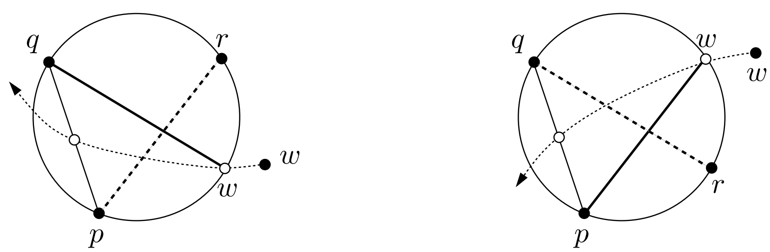

Figure 71: Left: Case (d1). The co-circularity at time δ_1 is red-blue with respect to the edges pr and wq. Right afterwards, the Delaunayhood of pr is violated by q and w. Right: Case (d2). The co-circularity at time δ_1 is red-blue with respect to the edges rq and pw. Right afterwards, the Delaunayhood of rq is violated by p and w.

Therefore, the points p, q, r, w (or, more precisely, their sub-triples) will perform four distinct Delaunay crossings—the two new crossings just promised and the two "old" ones, of pq by r and by w. If a pair of these crossings is performed by the *same* triple, we will use Lemma 4.5 to bound the overall number of such special quadruples χ. Otherwise we will charge χ to the (probabilistically refined) terminal quadruple $\varrho = (p, q, r, w)$, whose four possible sub-triples are involved in *four* Delaunay crossings, namely, the crossings of pq by r and w, the crossing of pr by w, and the crossing of wq by r.

In Section 7 we will use the third co-circularity δ_2 to enforce, for each terminal quadruple $\varrho = (p, q, r, w)$ of the above kind, an additional, fifth crossing (namely, a crossing of rw by p or q). As a result, some sub-triple of p, q, r, w will be involved in two Delaunay crossings, which will allow us to obtain a "quadratic" recurrence for the number of such quadruple, via Lemma 4.5.

In subcase (d2), we assume the co-circularity at time δ_1 to be red-blue with respect to the edges rq and pw (see Figure 71 (right)), and then use it to enforce a Delaunay crossing of rq by at least one of p and w. If rq is crossed by p, we can dispose of χ via Lemma 4.5. Otherwise, we charge χ to the (probabilistically refined) terminal quadruple $\varrho = (p, q, r, w)$ (whose points are known, so far, to perform only *three* crossings).

In Section 7 we will enforce, for each terminal quadruple $\varrho = (p, q, r, w)$ of the latter type, *two* additional crossings, namely, a crossing of pw by one of r, q, and a crossing of rw by one of p, q. Hence, once again we will be able to use Lemma 4.5 to handle such terminal quadruples too.

Case (d1). The co-circularity at time δ_1 is red-blue with respect to the edge pr whose Delaunayhood is violated right afterwards by q and w (see Figure 71 (left)).

Note that the above violation of pr does not hold either right before, or right after time λ_q. More precisely, it does not hold for that side of λ_q when w and r lie in the same side of L_{pq}, in which case the segments pq and rw do not even intersect; see Figure 68.

Therefore, and since δ_1 is the *only* red-red co-circularity of p, q, r, w with respect to pq, applying Lemma 3.1 over the interval (δ_1, λ_q), within the triangulation $\mathrm{DT}(\{p, q, r, w\})$, shows that pr is hit during (δ_1, λ_q) by at least one of q or w. See Figure 72 (top).

A very similar argument shows that the edge wq is hit by one of p or r after r enters L^+_{pq} (during I) and before δ_1. Indeed, let υ_{pq} denote the time in I when r hits pq. Note that the edge wq is violated right before δ_1 by p and r, and that the above violation did not hold at time υ_{pq}. Therefore, another application of Lemma 3.1 in $\mathrm{DT}(\{p, q, r, w\})$, from time δ_1 backwards, shows that the edge wq is hit during $(\upsilon_{pq}, \delta_1)$ by at least one of the two points p or r. See Figure 72 (bottom).

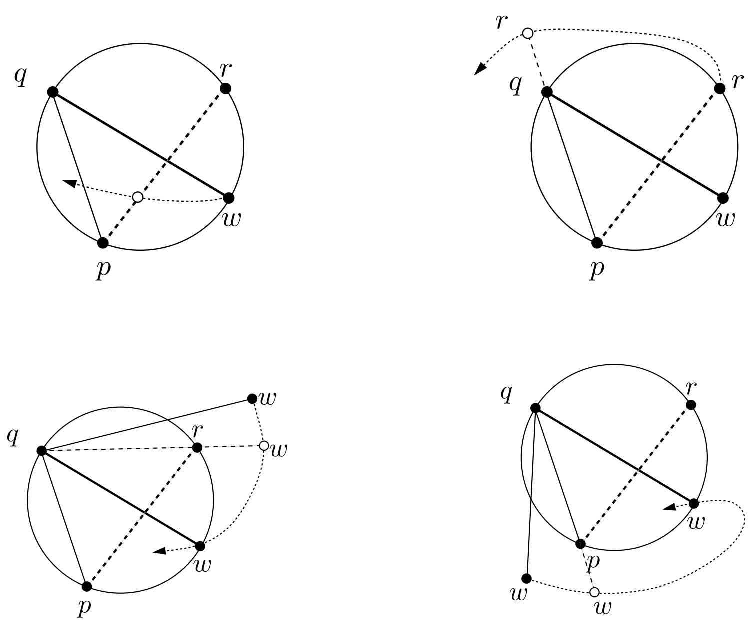

Figure 72: Lemma 6.7. Top: Possible trajectories of w (left) or r (right) during (δ_1, λ_q), which realize the crossing of pr by the respective point. Bottom: Possible trajectories of w during $(\upsilon_{pq}, \delta_1)$, which realize the crossing of wq by r (left) or by p (right).

To conclude, we have established the following claim.

Lemma 6.7. *With the above notation, the following two properties hold in case (d1):*

(i) The edge wq is hit in $(\upsilon_{pq}, \delta_1)$ by at least one of the points p, r. Namely, either r crosses wq from L^-_{wq} to L^+_{wq}, or p crosses wq in the reverse direction. Moreover, the Delaunayhood of wq is violated by p and r after the last such crossing and until δ_1.

(ii) The edge pr is hit in (δ_1, λ_q) by at least one of the points w, q. Namely, either w crosses pr from L^+_{pr} to L^-_{pr}, or q crosses pr in the reverse direction. Moreover, the Delaunayhood of pr is violated by w and q after δ_1 and until the first such crossing.

Case (d1) – the crossing of wq by p or r. We next turn the crossing in Lemma 6.7 (i) into a Delaunay crossing of wq by r. Recall that δ_1 belongs to the interval (λ_{wq}, ξ_{wq}). Therefore, and since wq is Delaunay at time λ_{wq} (and at time ξ_{wq}), the crossing in Lemma 6.7 (i) has to occur in the interval $[\lambda_{wq}, \delta_1)$; see Figure 73. Therefore, and since wq is Delaunay in $\mathrm{DT}(P \setminus \{a, p, r, u\})$ during $[\lambda_{wq}, \xi_{wq}]$ (by condition (S6)), wq undergoes within that latter interval a Delaunay crossing by p or r within a suitably reduced triangulation $\mathrm{DT}(P \setminus \{a, r, u\})$ or $\mathrm{DT}(P \setminus \{a, p, u\})$.

If wq is hit by p during $[\lambda_{wq}, \delta_1]$, then the points p, q, w define two Delaunay crossings within the reduced triangulation $\mathrm{DT}(P \setminus \{r, a, u\})$. A routine combination of Lemma 4.5 with the Clarkson-Shor

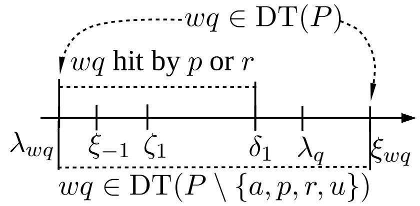

Figure 73: Case (d1)–obtaining a crossing of wq by at least one of p, r. The edge wq is Delaunay at times λ_{wq} and ξ_{wq}. Since the Delaunayhood of wq is violated by p and r right before time $\delta_1 \in [\lambda_{wq}, \xi_{wq}]$, it is hit by one of these points during $[\lambda_{wq}, \delta_1)$.

probabilistic argument implies that the overall number of such triples (p,q,w) in P is $O(n^2)$. By Proposition 6.2, this also bounds the overall number of such special quadruples χ.

We may therefore assume that wq is hit during $[\lambda_{wq}, \delta_1)$ by the point r, in which case the smaller set $P \setminus \{a,p,u\}$ induces a Delaunay crossing of wq by r. Note that each such triple (q,w,r) is shared by at most $O(1)$ special quadruples χ as above. Indeed, by Lemma 4.1, r cannot hit wq during the crossing $(wa, q, \mathcal{J}_u)$ (which is defined with respect to $P \setminus \{u\}$). If r hits wq in $[\lambda_{wq}, \lambda_2)$ then, by condition (S2), $(wa, q, \mathcal{J}_u)$ is among the first three clockwise special (w,q)-crossings to begin after that collinearity. Otherwise, if r hits wq in $(\lambda_3, \xi_{wq}]$, then condition (S6) similarly implies that $(wa, q, \mathcal{J}_u)$ is among the last three such (w,q)-crossings to end before this collinearity. We thus have established the following claim:

Lemma 6.8. *With the above assumptions, for any given triple (q,w,r) there remain at most six 3-restricted special quadruples $\chi = (a',p',w',q')$, with respective outer points r' and u', that satisfy $(q',w',r') = (q,w,r)$.*

In other words, any triple (q,w,r) is shared by at most *six* special quadruples that have survived the previous chargings (after falling into case (i)). Hence, the special quadruple χ under consideration is almost-uniquely determined by the choice of (q,w,r).

In what follows, we therefore assume that the edge wq undergoes (within a suitably reduced triangulation $\mathrm{DT}(P \setminus \{a,p,u\})$) a Delaunay crossing by r, and that χ and ϱ are almost uniquely determined by this additional crossing triple (q,w,r).

Case (d1)–the crossing of pr by q or w. We next turn the crossing in Lemma 6.7 (ii) into a Delaunay crossing of pr by w. If pr does not re-enter $\mathrm{DT}(P)$ after time δ_1 then, by Lemma 6.6, (pq, r, I) is among the $O(\ell)$ last (regular) (p,r)-crossings (because pr is Delaunay during each of these crossings). By Proposition 6.1, this can happen for at most $O(\ell n^2)$ special quadruples χ. Therefore, we may assume that pr re-enters $\mathrm{DT}(P)$ after δ_1.

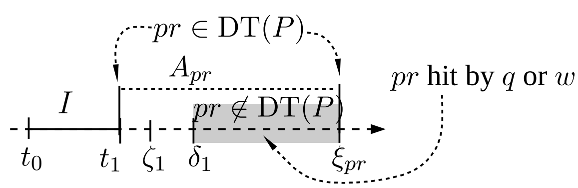

Figure 74: Case (d1)–enforcing a crossing of pr by one of the points q, w. The edge pr is Delaunay throughout $I = [t_0, t_1]$ and at time $\xi_{pr} > \delta_1$, which is the first such time after δ_1 when pr re-enters $\mathrm{DT}(P)$. The Delaunayhood of pr is violated by q and w right after $\delta_1 \in (t_1, \xi_{pr}]$, so it is hit by one of these points during $(\delta_1, \xi_{pr}]$.

Let ξ_{pr} denote the first time in $[\delta_1, \infty)$ when the edge pr is again Delaunay (in P); see Figure 74. Clearly, the time when pr is hit by one of q, w (as prescribed by Lemma 6.7 (ii)) belongs to the interval $(\delta_1, \xi_{pr}]$, which is contained in $(\zeta_1, \xi_{pr}] \subseteq (t_1, \xi_{pr}]$. To turn this crossing into a Delaunay crossing, we apply Theorem 2.2 in $\mathcal{A}_{pr}$ over the interval (t_1, ξ_{pr}), with the third constant $h \gg \ell$.

If at least one of the Conditions (i), (ii) of Theorem 2.2 holds, we can charge χ, within $\mathcal{A}_{pr}$, either to an h-shallow collinearity or to $\Omega(h^2)$ h-shallow co-circularities. Lemma 6.6 ensures that each h-shallow event, that occurs in $\mathcal{A}_{pr}$ at some time $t^* \in (t_1, \xi_{pr})$, is charged in this manner by at most $O(\ell)$ special quadruples. Indeed, the corresponding points p and r are involved in the event, so we can guess them in $O(1)$ possible ways, and (pq, r, I) is among the last $8\ell + 2$ clockwise (p, r)-crossings to end before time t^*. Therefore, the above charging accounts for at most $O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples χ.

We may assume, then, that Condition (iii) of Theorem 2.2 holds. That is, there is a subset A_{pr} of at most $3h$ points (perhaps including some of q, a, u and w) whose removal restores the Delaunayhood of pr throughout the interval $[t_1, \xi_{pr}]$.

If pr is crossed during $(\delta_1, \xi_{pr}]$ by q (from L^-_{pr} to L^+_{pr}), then the triple p, q, r performs two Delaunay crossings within the triangulation $\mathrm{DT}((P \setminus A_{pr}) \cup \{q\})$. A routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that P contains at most $O(hn^2)$ triples p, q, r of this kind. By Proposition 6.1, this also bounds the overall number of such special quadruples χ.

To conclude, we are left with the case where the edge pr is crossed during $(\delta_1, \xi_{pr}]$ by w (from L^+_{pr} to L^-_{pr}). Hence, the reversely oriented copy rp of pr undergoes within the smaller triangulation $\mathrm{DT}((P \setminus A_{pr}) \cup \{w\})$ a Delaunay crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$, where $\mathcal{T} \subseteq [t_1, \xi_{pr}]$ (the crossing must begin after t_1, since pr is Delaunay during I, by Lemma 4.1).

Lemma 6.9. *With the above assumptions, for any given triple (p, r, w) there remain at most $8\ell + 2$ 3-restricted special quadruples $\chi = (a', p', w', q')$, with respective outer points r' and u', that fall into case (d1) and satisfy $(p', r', w') = (p, r, w)$.*

Proof. By Proposition 6.1, each χ as above is uniquely determined by (pq, r, I) which, according to Lemma 6.6, is among the last $8\ell + 2$ clockwise (p, r)-crossings to end before w hits pr (as prescribed by Lemma 6.7). $\square$

If the above Delaunay crossing of rp by w, which occurs within the reduced triangulation $\mathrm{DT}((P \setminus A_{pr}) \cup \{w\})$, is a double Delaunay crossing, then we can charge χ to this crossing. A standard combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that the overall number of such triples (p, r, w) in P is only $O(hn^2)$, so the overall number of such special quadruples χ does not exceed $O(\ell h n^2)$. Therefore, we may assume, in what follows, that the above crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$ is a *single* Delaunay crossing.

To facilitate the subsequent steps of the analysis, we augment the above conflict set A_{pr} as follows. For each clockwise (p, r)-crossing (pq', r, I') (in $\mathcal{F}$) that ends during (t_1, δ_1) we add the respective point q' to A_{pr}. Informally, this is done to get rid of these (p, r)-crossings (pq', r, I') (see below for details). Since there are only at most $8\ell + 1$ such points q' (and since $\ell \ll h$), the overall cardinality of A_{pr}, after the augmentation, is at most $3h + 8\ell + 1 \leq 4h$.

To conclude, in case (d1), after disposing of $O\left(N(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples, we may assume that the four points of $\varrho = (p, q, r, w)$ perform at least four Delaunay crossings, namely, (pq, r, I), $(qp, w, \mathcal{H})$, the crossing of wq by r (which occurs in $P \setminus \{a, p, u\}$ and within $[\lambda_{wq}, \xi_{wq}]$), and the lately enforced single Delaunay crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$ (which occurs in $(P \setminus A_{pr}) \cup \{w\}$).

Case (d1) – converging. In Section 7.1, we will exploit the third co-circularity of p, q, r, w, which occurs at time $\delta_2 \in (\delta_1, \lambda_{pq}]$ and is red-blue with respect to pq and rw, to enforce the crossing of rw by at least one of p and q. As a result, one of the triples (p, r, w) of (q, r, w) will perform two Delaunay crossings in an appropriately refined triangulation, and our analysis will bottom out into a quadratic bound via Lemma 4.5.

To obtain the above crossing of rw, we will first apply Theorem 2.2 in the red-blue arrangement of this edge, so as to extend the (almost-)Delaunayhood interval of rw from $\mathcal{T} = [\tau_0, \tau_1]$ (where rp

undergoes an almost-Delaunay crossing by w) to a larger interval that will contain both δ_2, and a time when rw is hit by p or q. At the end of the analysis, we will manage either to charge the quadruple (p, q, r, w) within $\mathcal{A}_{rw}$ (in cases (i) and (ii) of Theorem 2.2), or else to extract the desired Delaunay crossing of rw.

The above use of Theorem 2.2 will be prepared by applying Theorem 5.3 for the clockwise (r, w)-crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$, so as to ensure that each event in $\mathcal{A}_{rw}$ be charged by only few other such terminal quadruples $\varrho' = (p', q', r, w)$, via the respective (r, w)-crossings $(rp', w, \mathcal{T}')$. That is, if we encounter too many (r, w)-crossings $(rp', w, \mathcal{T}')$ that can charge such an event, the crossing $(rp, w, \mathcal{T})$ will become (p, w)-chargeable, and can thus be accounted for by Theorem 2.2.

In order for the crossing $(rp, w, \mathcal{T})$ to be (p, w)-chargeable, we need an appropriate time ξ_{pw} after δ_2 when the edge pw is Delaunay (or, at least, almost Delaunay, with none of the obstruction points equal to r, p, w). In addition, the edge pw must be almost Delaunay throughout the entire interval where Theorem 5.3 is applied. We next proceed to accomplish all these steps in more detail.

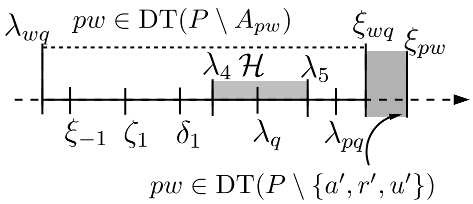

Figure 75: In the preparation for cases (b) and (c), we have extended the Delaunayhood of pw from $\mathcal{H} = [\lambda_4, \lambda_5]$ (where it belongs to $\mathrm{DT}(P \setminus \{a, r, u\})$) to the larger interval $[\lambda_{wq}, \xi_{wq}]$. We next extend the almost-Delaunayhood of pw beyond ξ_{wq}, until some time ξ_{pw} when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$ (for $a', r', u' \notin \{q, r\}$).

Charging even more events in $\mathcal{A}_{pw}$. Our first step is to extend the almost-Delaunayhood of pw. Refer to Figure 75. Recall that, in preparation for cases (b) and (c), we have already extended the almost-Delaunayhood of pw from $\mathcal{H} = \mathcal{H}_\chi = [\lambda_4, \lambda_5]$ (where qp is crossed by w) to the interval $[\lambda_{wq}, \xi_{wq}]$, which covers $\mathcal{H} = [\lambda_4, \lambda_5], (\xi_{-1}, \lambda_q)$ and λ_{pq}. (In particular, $[\lambda_{wq}, \xi_{wq}]$ contains $\delta_1 \in (\zeta_1, \lambda_q) \subset (\xi_{-1}, \lambda_q)$ and $\delta_2 \in (\delta_1, \lambda_{pq}]$.) This has been achieved at the cost of removing a certain subset A_{pw}, which consists of at most $6\ell + 3$ points, including a, r, u. Unfortunately, the above obstruction set A_{pw} contains r (and perhaps also w), so removing A_{pw} in its entirety would destroy the Delaunay crossing $(rp, w, \mathcal{T})$ (instead of facilitating its (p, w)-chargeability in a smaller triangulation).

We next obtain a time $\xi_{pw} > \xi_{wq}$ when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, for $a', r', u' \notin \{q, r\}$, and extend the almost-Delaunayhood of pw from λ_5 beyond ξ_{wq}, until ξ_{pw}.

To do so, we return to the family $\mathcal{G}^R_{pw}$ of 3-restricted right special quadruples $\chi' = (a', p, w, q')$ that share their middle points p, w with χ. (In particular, $\mathcal{G}^R_{pw}$ includes χ.)

Recall that each special quadruple $\chi' \in \mathcal{G}^R_{pw}$ is accompanied by a (p, w)-crossing $(q'p, w, \mathcal{H}')$, which is defined with respect to the corresponding set $P \setminus \{a', r', u'\}$. Without loss of generality, we assume that all quadruples in $\mathcal{G}^R_{pw}$ fall into case (d1), and that none of them have been disposed of by the previous chargings within $\mathcal{A}_{pr}$. (In addition, we continue to assume that the special quadruple χ under consideration satisfies condition (PHR1).)

By Lemma 6.9, any triple (p, r', w) can be shared by at most $8\ell + 2$ special quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw}$ under consideration (each with its respective outer points r' and u'). Therefore, the pigeonhole principle implies that at least some fixed fraction of all 3-restricted quadruples $\chi = (a, p, w, q) \in \mathcal{G}^R_{pw}$ under consideration (again, with respective outer points r and u) satisfy the following condition:

(PHR2) *At most $O(\ell)$ other 3-restricted quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw}$ (each with its respective* *outer points r' and u') can satisfy $r \in \{a', r', u'\}$.*

(Briefly, this can be shown by considering the *multi-function* $\mu : \mathcal{G}^R_{pw} \to \mathcal{G}^R_{pw}$ mapping each special quadruple $\chi = (a, p, w, q)$, with respective outer points r and u, to at most $(8\ell + 2) \times 3 = O(\ell)$ other quadruples $\chi' = (a', p, w, q')$, whose respective outer points r' are chosen from a, r, u. Hence, average "in-degree" of each quadruple $\chi \in \mathcal{G}^R_{pw}$, which is exactly the number of quadruples χ' so that at least one of their respective points a', r', u' is equal to the first outer point r of χ, is also $O(\ell)$.)

Therefore, we can assume, in what follows, that the above condition holds for χ at hand. Combining this[38] with (PHR1) shows that all but at most $3 + O(\ell) = O(\ell)$ special quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^R_{pw}$, with respective outer points r' and u', have $\{q, r\} \cap \{a', r', u'\} = \emptyset$. Recall also that, since case (a) has been ruled out, $\mathcal{G}^R_{pw}$ contains at most k quadruples χ' whose respective (p, w)-crossings $(q'p, w, \mathcal{H}_{\chi'})$ end in $(\lambda_5, \xi_{wq}]$. See Figure 76.

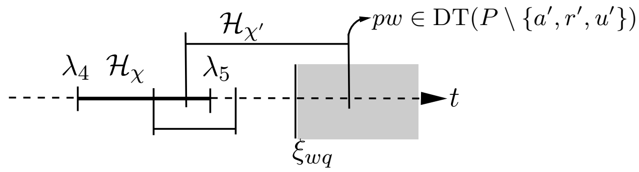

Figure 76: The family $\mathcal{G}^R_{pw}$ contains at most $O(\ell)$ quadruples χ' with non-empty intersection $\{a', r', u'\} \cap \{q, r\}$, and at most k quadruples χ' whose respective (p, w)-crossings end in $(\lambda_5, \xi_{wq}]$. If $\mathcal{G}^R_{pw}$ contains no special quadruples χ' that satisfy $\{a', r', u'\} \cap \{q, r\} = \emptyset$, and whose respective (p, w)-crossings $(q'p, w, \mathcal{H}_{\chi'})$ end after ξ_{wq}, then $(qp, w, \mathcal{H})$ is among the last $O(\ell)$ such (p, w)-crossings.

Assume first that $\mathcal{G}^R_{pw}$ contains no special quadruples $\chi' = (a', p, w, q')$ (with respective outer points r' and u') that satisfy $\{a', r', u'\} \cap \{q, r\} = \emptyset$, and whose respective (p, w)-crossings $(q'p, w, \mathcal{H}_{\chi'})$ end after ξ_{wq}. Therefore, $\mathcal{G}^R_{pw}$ contains at most $k + O(\ell) = O(\ell)$ such quadruples χ' whose (p, w)-crossings $(q'p, w, \mathcal{H}_{\chi'})$ end after the ending time λ_5 of $\mathcal{H} = \mathcal{H}_\chi$ (including the at most k such quadruples whose (p, w)-crossings χ' end in $(\lambda_5, \xi_{wq}]$, and the at most $O(\ell)$ such quadruples χ' with non-empty intersection $\{a', r', u'\} \cap \{q, r\}$). Hence, we can charge χ, via its respective (p, w)-crossing $(qp, w, \mathcal{H}_\chi = \mathcal{H})$, to the edge pw, so the above scenario occurs for at most $O(\ell n^2)$ special quadruples χ under consideration.

Assume, then, that, for some $\chi' \in \mathcal{G}^R_{pw}$, with $\{q, r\} \cap \{a', r', u'\} = \emptyset$, its respective (p, w)-crossing $(q'p, w, \mathcal{H}_{\chi'})$ ends after ξ_{wq}. By Lemma 4.1, pw belongs to $\mathrm{DT}(P \setminus \{a', r', u'\})$ throughout $\mathcal{H}_{\chi'}$. In particular, we can choose a time $\xi_{pw} \in [\xi_{wq}, \infty)$, which is the first such time when the edge pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, where $a', r', u' \in P \setminus \{r, q\}$. In what follows, we use a', r' and u' to denote the above three points a', r', u', whose removal restores the Delaunayhood of pw at time ξ_{pw}.

The preceding discussion implies that at most $O(\ell)$ of the above (p, w)-crossings $(q'p, w, \mathcal{H}_{\chi'})$ can end in $(\xi_{wq}, \xi_{pw}]$ (and that, for each of those crossings, its respective obstruction set $\{a', r', u'\}$ intersects $\{q, r\}$). Therefore, and since case (a) has been ruled out, at most $k + O(\ell) = O(\ell)$ of the above (p, w)-crossings can end in (λ_5, ξ_{pw}).

We are finally ready to apply Theorem 2.2 in $\mathcal{A}_{pw}$ over the interval (λ_5, ξ_{pw}) (see Figure 77). This is done with the third constant $h \gg \ell$ and with respect to the smaller set $P \setminus \{a', r', u'\}$. If at least one of the first two conditions of Theorem 2.2 holds, we charge χ within $\mathcal{A}_{pw}$ either to an $(h + 3)$-shallow collinearity, or to $\Omega(h^2)$ $(h + 3)$-shallow co-circularities (as in the previous chargings, these events are h-shallow in $P \setminus \{a', r', w'\}$, and $(h + 3)$-shallow in P). Clearly, each $(h + 3)$-shallow event in $\mathcal{A}_{pw}$ is charged as above by at most $O(\ell)$ special quadruples χ, because $(qp, w, \mathcal{H}_{\chi'})$ is among the last $O(\ell)$ such (p, w)-crossings to end before the event. Hence, the above charging accounts for

[38] As a matter of fact, our previous inability to enforce (PHR2) was the only reason why the present analysis in $\mathcal{A}_{pw}$ had not been applied right after handling case (a), in a more general context.

$O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples χ.

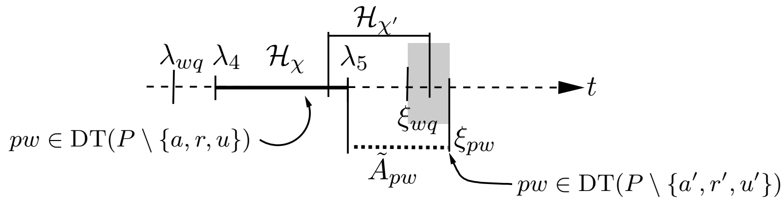

Figure 77: Extending the almost-Delaunayhood of pw to $[\lambda_5, \xi_{pw}]$. ξ_{pw} is the first time in $[\xi_{wq}, \infty)$ when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, for $\{a', r', u'\} \cap \{q, r\} = \emptyset$. We apply Theorem 2.2 within $\mathcal{A}_{pw}$ over the interval (λ_5, ξ_{pw}), noting that $(qp, w, \mathcal{H}_\chi)$ is among the last $O(\ell)$ such (p, w)-crossings $(q'p, w, \mathcal{H}_{\chi'})$ to end before any charged event.

We can therefore assume that Condition (iii) of Theorem 2.2 holds. Hence, there is a subset $\tilde{A}_{pw}$ of at most $3h + 3$ points (including the above three points $a', r', u' \in P \setminus \{p, q, r, w\}$), whose removal restores the Delaunayhood of pw throughout (λ_5, ξ_{pw}). Therefore, pw belongs to $\mathrm{DT}(P \setminus (A_{pw} \cup \tilde{A}_{pw}))$ throughout the entire interval $[\lambda_{wq}, \xi_{pw}] = [\lambda_{wq}, \xi_{wq}] \cup (\lambda_5, \xi_{pw}]$ (where A_{pw} denotes the set of at most $6\ell + 3$ points, including a, r, u, whose removal restores the Delaunayhood of pw throughout $[\lambda_{wq}, \xi_{wq}]$).

Case (d1)–Wrap up. We again emphasize that the times of the various events discussed so far appear in the order

$$\xi_{pq} < \lambda_{wq} < \xi_{-1} < \zeta_1 < \delta_1 < \lambda_q < \lambda_{pq} < \xi_{wq} < \xi_{pw},$$

that $\delta_2 \in (\delta_1, \lambda_{pq}]$, and that w crosses rp from L^-_{rp} to L^+_{rp} in the interval (δ_1, λ_q), as part of a single Delaunay crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$ (which occurs in $(P \setminus A_{pr}) \cup \{w\}$). Refer to Figure 78.

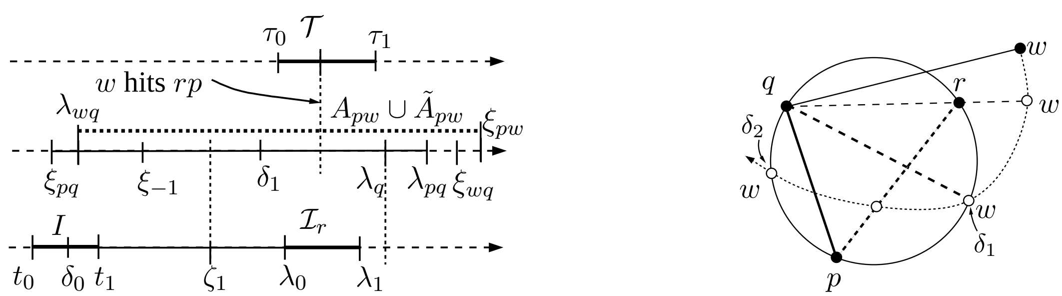

Figure 78: Case (d1): A (partial) summary of what we assume at the end of the analysis. Left: Various events occur in the depicted order (and δ_2 lies in $(\delta_1, \lambda_{pq}]$). Right: A possible motion of w after r enters L^+_{pq} (during I).

By the definition of A_{pw} and $\tilde{A}_{pw}$ (of total cardinality $6\ell + 3 + 3h + 3 = O(h)$), the edge pw belongs to $\mathrm{DT}(P \setminus (A_{pw} \cup \tilde{A}_{pw}))$ throughout the interval $[\delta_1, \xi_{pw}] \subseteq [\lambda_{wq}, \xi_{pw}]$. Furthermore, pw belongs at time ξ_{pw} to the triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, where $a', r', u' \in \tilde{A}_{pw} \setminus \{q, r\}$.

Recall also that, since ζ_1 belongs to both intervals $[t_0, \lambda_1] = \mathsf{conv}(I \cup \mathcal{I}_r)$ and $(\xi_{-1}, \xi_0) \subset [\xi_{pq}, \lambda_{pq}]$, the combination of conditions (Q8) and (S6) (on, respectively, σ and χ) implies that the edge pq belongs to $\mathrm{DT}(P \setminus \{a, w, r, u\})$ throughout the interval $[t_0, \lambda_{pq}] \subseteq [t_0, \lambda_1] \cup [\xi_{pq}, \lambda_{pq}]$.

Finally, we continue to assume that the edge wq undergoes a Delaunay crossing by r within $P \setminus \{a, p, u\}$. (The precise interval of this crossing is immaterial for our future analysis.)

In what follows, we use A^+_{pq} to denote the set of all points of P that appear in the cap C^+_{pq} at some time in (ξ_{-1}, λ_q). By Lemma 6.4, the cardinality of A^+_{pq} does not exceed $8\ell + 1$.

Case (d1) – charging terminal quadruples. To proceed, we draw a random sample R of $\lceil n/h \rceil$ points of P. Notice that the following two events occur simultaneously with probability at least $\Omega(1/h^4)$: (1)

The four points p, q, w, r belong to R, and (2) R includes none of the points of

$$(A^+_{pq} \cup \tilde{A}_{pw} \cup A_{pw} \cup A_{pr} \cup \{a, u\}) \setminus \{p, q, r, w\}.$$

Suppose that the sample R is indeed successful for the 3-restricted right special quadruple $\chi = (a, p, w, q)$ at hand, with respective two outer points r and u. Then we can charge χ to the quadruple $\varrho = (p, q, r, w)$, which satisfies the following conditions with respect to the sample R (see Figure 106 in Section 7.1 for a schematic summary, with R replaced by P).

(A1) The edge pq undergoes (in R) a Delaunay crossing $(pq, r, I = [t_0, t_1])$ and is crossed by w, from L^+_{pq} to L^-_{pq}, at some later time $\lambda_q > t_1$. In addition, pq is again Delaunay at some time λ_{pq} which is the first such time after t_q, and it belongs to $\mathrm{DT}(P \setminus \{r, w\})$ throughout (t_1, λ_{pq}). Hence, its reversely oriented copy qp undergoes in $R \setminus \{r\}$ (and entirely within $(t_1, \lambda_{pq}]$) a Delaunay crossing by w.

(A2) The points p, q, w, r are co-circular at times $\delta_0 \in I$, $\delta_1 \in [t_1, \lambda_q]$, and $\delta_2 \in (\delta_1, \lambda_{pq}]$, and the following properties hold:
(i) The co-circularity at time δ_0 is red-blue with respect to pq.
(ii) The co-circularity at time δ_1 is red-red with respect to pq and red-blue with respect to the edge pr, whose Delaunayhood is violated right after time δ_1 by q and w. Furthermore, the open cap $C^+_{pq} = B[p, q, w] \cap L^+_{pq}$ contains no points of P at time δ_1.
(iii) The co-circularity at time δ_2 is again red-blue with respect to pq. It arises during a single Delaunay crossing of qp by w, which occurs in $\mathrm{DT}(\{p, q, r, w\})$ during some sub-interval of $(\delta_1, \lambda_{pq}]$.

(A3) The set $R \setminus \{q\}$ induces a (single) Delaunay crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$, where w crosses rp from L^-_{rp} to L^+_{rp} during (δ_1, λ_q).

Similarly, the set $R \setminus \{p\}$ induces a Delaunay crossing of wq by r, where r crosses wq before δ_1, and from L^-_{wq} to L^+_{wq}.

(A4) There exists a time $\xi_{pw} > \lambda_{pq}$ so that (i) the edge pw is Delaunay (in R) at time ξ_{pw}, and (ii) pw belongs to $\mathrm{DT}(R \setminus \{q, r\})$ throughout the interval $[\delta_1, \xi_{pw}]$.

In Section 7.1 we show that pw is Delaunay also at time δ_1. In addition, Lemma 4.1 implies that pw belongs to $\mathrm{DT}(R \setminus \{q\})$ throughout the interval $\mathcal{T} = [\tau_0, \tau_1]$, which obviously intersects $[\delta_1, \xi_{pw}] (\supset [\delta_1, \lambda_q])$.

Notice that any such quadruple $\varrho = (p, q, r, w)$ in R is charged as above by at most one 3-restricted right special quadruple $\chi = (a, p, w, q)$ in $\mathcal{F}$ (with outer points r and u), because the latter quadruple is uniquely determined by each of the triples (p, q, r) and (p, q, w).

We say that a quadruple $\varrho = (p, q, r, w)$ is *terminal of type A* if it satisfies the above four conditions (A1)–(A4) with respect to the underlying set R. (In Section 7, we shall again use P to denote the underlying point set of our terminal quadruples. See Figure 106 in that section for a partial summary of the properties of terminal quadruples of type A.)

Let Σ^A_R denote the resulting family of terminal quadruples $\varrho = (p, q, r, w)$ (of type A) in R that are charged by 3-restricted right special quadruples in P through the above probabilistic argument.

Lemma 6.10. *With the above assumptions, each terminal quadruple $\varrho = (p, q, r, w)$ in Σ^A_R is uniquely determined by each of its sub-triples $(p, q, r), (p, q, w)$, (p, r, w). Furthermore, any triple (q, r, w) is shared by at most six terminal quadruples of Σ^A_R.*

Proof. Clearly, the second part of the lemma is directly implied by Lemma 6.8, so it suffices to establish the first part of it.

By condition (A1), w is the first point of P to hit the edge pq after its Delaunay crossing $(pq, r, I = [t_0, t_1])$ by r. Hence, $\varrho = (p, q, r, w)$ is uniquely determined by the choice of p, q and r. A similar agrument implies that $\varrho = (p, q, r, w)$ is uniquely determined by the triple (p, q, w).

To see that ϱ is uniquely determined by (p, r, w), let us assume for a contradiction that Σ_R^A contains another such quadruple $\varrho' = (p, q', r, w)$ (of type A and with $q' \neq q$). Furthermore, assume with no loss of generality that the respective (p, r)-crossing (pq', r, I') of ϱ' ends after $I = [t_0, t_1]$. Note though that I' must end before w enters L_{pr}^- through pr (as prescribed by condition (A3)). However, in that case I' would end in (t_1, δ_1), so q' would have been included in the respective set A_{pr} of ϱ, and, therefore, omitted[39] from R, contrary to the choice of $\varrho' \in \Sigma_R^A$. □

To simplify the presentation, in what follows we only consider a subfamily $\Sigma^A = \Sigma_R^A$ of terminal quadruples of type A whose members $\varrho = (p, q, r, w)$ are *uniquely determined* by each one of their respective four sub-triples $(p, q, r), (p, q, w), (p, r, w)$, and (q, r, w). This stronger uniqueness condition can be enforced by prunning Σ_R^A (without affecting its asymptotic cardinality), so that, for each triple (q, r, w), we keep in Σ_R^A only one terminal quadruple (p, q, r, w), if such quadruples exist at all in Σ_R^A.

Let $T^A(m)$ denote the maximum cardinality of a family Σ^A of terminal quadruples of type A (with the above uniqueness property) that can be defined over a set of m moving points. The preceding analysis implies that the overall number of special quadruples that fall into Case (d1) is at most

$$O\left(h^4 T^A(n/h) + \ell h^2 N(n/h) + \ell h n^2 \beta(n)\right).$$

Case (d2). The co-circularity at time δ_1 is red-blue with respect to the edge qr, whose Delaunayhood is violated right after that by p and w. We continue to assume that r does not cross pq again during $(t_1, \lambda_{pq}]$.

As in case (d1), we use υ_{pq} to denote the time in $I = [t_0, t_1]$ when r enters the halfplane L_{pq}^+. We have the following lemma, whose proof is fully symmetric to that of Lemma 6.7.

Lemma 6.11. *With the above notation, the following two properties hold in case (d2):*

(i) The edge pw is hit in $(\upsilon_{pq}, \delta_1)$ by at least one of the points q, r. Namely, either r crosses pw from L_{pw}^- to L_{pw}^+, or q crosses pw in the reverse direction. Moreover, the Delaunayhood of pw is violated by q and r after the last such crossing and until δ_1.

(ii) The edge rq is hit in (δ_1, λ_q) by at least one of the points p, w. Namely, either w crosses rq from L_{rq}^+ to L_{rq}^-, or p crosses rq in the reverse direction. Moreover, the Delaunayhood of rq is violated by w and q after δ_1 and until the first such crossing.

Refer to Figure 79. To prove part (i) of Lemma 6.11, we note that, right before time δ_1, the Delaunayhood of pw is violated by $q \in L_{pw}^-$ and $r \in L_{pw}^+$, and that this violation does not hold either right before, or right after the time υ_{pq} when r crosses pq. Hence, to obtain the desired crossing of pw, we can apply the time reversed variant Lemma 3.1 for the triangulation $\mathrm{DT}(\{p, q, r, w\})$, over the interval $(\upsilon_{pq}, \delta_1)$.

To prove part (ii) of Lemma 6.11, we apply (the regular variant of) Lemma 3.1 in $\mathrm{DT}(\{p, q, r, w\})$ over the interval (δ_1, λ_q), noting that the violation of rq by $p \in L_{rq}^-$ and $w \in L_{rq}^+$, which holds right after time δ_1, no longer exists either right before, or right after, the time λ_q when w hits pq.

Case (d2) – enforcing the crossing of rq by p or w. Our argument is fully symmetric to the one used in case (d1) to enforce a Delaunay crossing of pr by q or w.

Recall that, according to Lemma 6.6, at most $8\ell + 1$ counterclockwise (q, r)-crossings can end in the interval (t_1, δ_0). If rq never re-enters $\mathrm{DT}(P)$ after time δ_1, then (pq, r, I) is among the last $8\ell + 2$ counterclockwise Delaunay (q, r)-crossings in $\mathcal{F}$ (with respect to the standard order implied by Lemma 4.6). Clearly, this scenario happens for at most $O(\ell n^2)$ special quadruples χ, because each of them is uniquely determined by the respective triple (p, q, r) (according to Proposition 6.1). Therefore, we may

[39] Clearly, we have $q' \neq w$ (i.e., there is no crossing (pw, r, I')), because r can enter the halfplane L_{pw}^+ only once, and it is already assumed to cross the line L_{pw}, from L_{pw}^- to L_{pw}^+, and *outside* pw (as prescribed by Lemma 6.7 (ii)).

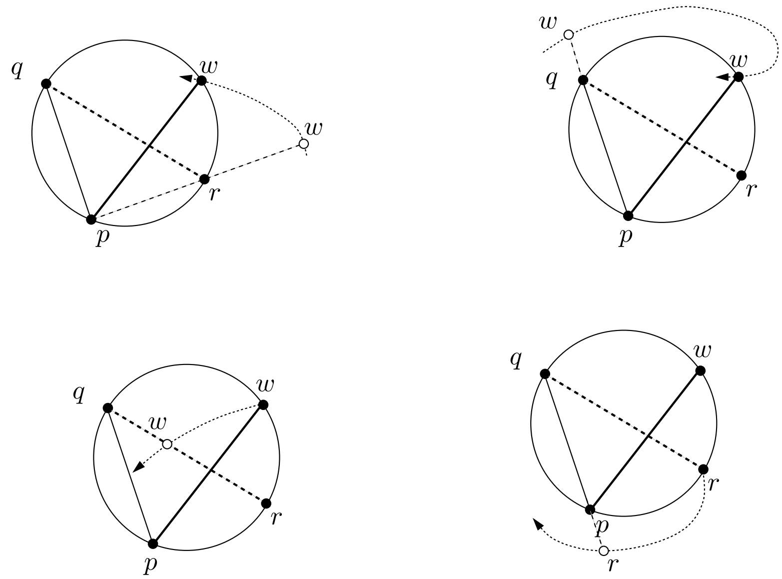

Figure 79: Lemma 6.11. Top: Possible trajectory of w during $(\upsilon_{pq}, \delta_1)$, which realize the crossing of pw by r (left) or q (right). Bottom: Possible trajectories of w (left) and r (right) during $(\upsilon_{pq}, \delta_1)$, which realize the crossing of rq by the respective point.

assume, in what follows, that rq re-enters $\mathrm{DT}(P)$ at some future time $\xi_{rq} > \delta_1$ (which is the *first* such time when rq is Delaunay); see Figure 80. By Lemma 6.11 (ii), rq is hit during $(\delta_1, \xi_{rq}] \subset (t_1, \xi_{rq}]$ by p or w. Furthermore, Lemma 6.6 (combined with Lemma 4.1) implies at most $8\ell + 1$ counterclockwise (Delaunay) (q, r)-crossings in $\mathcal{F}$ can end during $(t_1, \xi_{rq}]$.

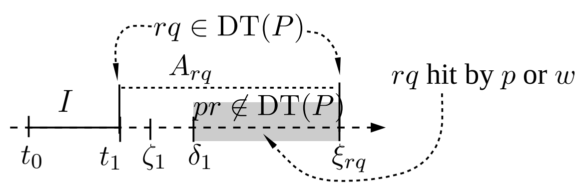

Figure 80: Case (d2)–enforcing a crossing of rq by at least one of the points p, w. The edge rq is Delaunay throughout $I = [t_0, t_1]$ and at time $\xi_{rq} > \delta_1$, which is the first such time after δ_1 when rq re-enters $\mathrm{DT}(P)$. The Delaunayhood of rq is violated by q and w right after $\delta_1 \in (t_1, \xi_{rq}]$, so it is hit by one of these points during $(\delta_1, \xi_{rq}]$.

To enforce the desired crossing of rq, we apply Theorem 2.2 in $\mathcal{A}_{rq}$ over the interval (t_1, ξ_{rq}), with the third threshold $h \gg \ell$.

If one of the Conditions (i), (ii) holds, we charge χ (via (pq, r, I)) either to an h-shallow collinearity or to $\Omega(h^2)$ h-shallow co-circularities. Clearly, each of these h-shallow events is charged at most $O(\ell)$ times in the above manner, because (pq, r, I) is among the last $8\ell + 2$ counterclockwise (q, r)-crossings (in $\mathcal{F}$) to end before the time of the event. Hence, the above charging accounts for at most $O(\ell h^2 N(n/h) + \ell h n^2 \beta(n))$ special quadruples χ.

Assume, then, that Condition (iii) of Theorem 2.2 holds, so we have a subset A_{rq} of at most $3h$ points (possibly including p or w, or both) whose removal restores the Delaunayhood of rq throughout the entire interval $[t_0, \xi_{rq}] = I \cup [t_1, \xi_{rq}]$. To facilitate the subsequent analysis, we augment the set A_{pr} as follows. For each crossing $(p'q, r, I')$ (in $\mathcal{F}$) that ends in the interval (t_1, δ_1) we add the respective point p' to A_{rq}.

If rq is hit during $(\delta_1, \xi_{rq}]$ by p, then the triple (p, q, r) is involved in two Delaunay crossings, which occur within the smaller triangulation $\mathrm{DT}((P \setminus A_{rq}) \cup \{p\})$. According to Lemma 4.5, the overall number of such triples in P does not exceed $O(hn^2)$. By Proposition 6.1, this also bounds the overall number of the respective special quadruples χ.

To conclude, we may assume, in what follows, that rq is hit during $(\delta_1, \xi_{rq}] \subset (t_1, \xi_{rq})$ by w (which crosses it from L_{rq}^+ to L_{rq}^-). Therefore, the reversely oriented copy qr of rq undergoes, within the reduced triangulation $\mathrm{DT}((P \setminus A_{rq}) \cup \{w\})$, a Delaunay crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$.

Notice that (pq, r, I) is among the last $8\ell + 2$ (q, r)-crossings in $\mathcal{F}$ to end before w crosses rq from L_{rq}^+ to L_{rq}^-, which implies the following symmetric analogue of Lemma 6.9:

Lemma 6.12. *Any triple (q, r, w) is shared by at most $8\ell + 2$ 3-restricted special quadruples $\chi = (a, p, w, q)$ (with respective outer points r and u) of the above kind.*

If the above crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$ is a double Delaunay crossing, we apply Lemma 4.5 (in combination with the Clarkson-Shor argument) to establish an upper bound of $O(hn^2)$ on the overall number of such triples (q, r, w) in P, which immediately yields an upper bound of $O(\ell hn^2)$ on the number of special quadruples χ of this kind. Hence, we may assume, in what follows, that the above crossing of qr by w in $\mathrm{DT}((P \setminus A_{rq}) \cup \{w\})$ is a *single* Delaunay crossing.

We again emphasize that $\lambda_{wq} < \delta_1 < \lambda_q < \lambda_{pq} < \xi_{wq}$ and $\delta_2 \in (\delta_1, \lambda_{pq}]$, and that w hits qr (during $\mathcal{T} = [\tau_0, \tau_1]$) in the interval (δ_1, λ_q). Furthermore, by condition (S6), wq belongs to $\mathrm{DT}(P \setminus \{a, p, r, u\})$ throughout $[\lambda_{wq}, \xi_{wq}] \subset (\delta_1, \xi_{wq})$ (and is Delaunay at times λ_{wq} and ξ_{wq}). See Figure 81.

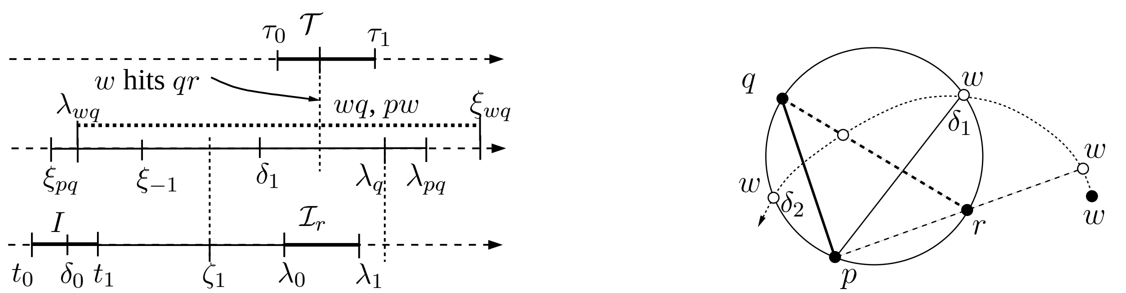

Figure 81: Case (d2): A (partial) summary of what we assume at the end of the analysis. Left: Various events occur in the depicted order (and δ_2 lies in $(\delta_1, \lambda_{pq}]$). Right: A possible motion of w after r enters L_{pq}^+ (during I).

Case (d2) – charging terminal quadruples. As in case (d1), let A_{pq}^+ denote the set of at most $8\ell + 1$ points that show up in the cap $C_{pq}^+ = B[p, q, w] \cap L_{pq}^+$ at some time in (ξ_{-1}, λ_q) (see Lemma 6.4).

To proceed, we draw a random sample of R of $\lceil n/h \rceil$ points of P. Notice that the following two events occur simultaneously with probability at least $\Omega(1/h^4)$: (1) The four points p, q, w, r belong to R, and (2) R includes none of the points of

$$(A_{pw} \cup A_{rq} \cup A_{pq}^+ \cup \{a, u\}) \setminus \{p, q, r, w\}$$

Suppose that the sample R is indeed successful for the 3-restricted right special $\chi = (a, p, w, q)$ at hand (with respective two outer points r and u). Then we can charge χ to the quadruple $\varrho = (p, q, r, w)$, which satisfies the following conditions with respect to the sample R:

(B1) The edge pq undergoes a Delaunay crossing $(pq, r, I = [t_0, t_1])$ and is crossed by w, from L_{pq}^+ to L_{pq}^-, at some later time $\lambda_q > t_1$. In addition, pq is again Delaunay at some time λ_{pq} which is the first such time after t_q, and it belongs to $\mathrm{DT}(R \setminus \{r, w\})$ throughout (t_1, λ_{pq}). Hence, its reversely oriented copy qp undergoes in $R \setminus \{r\}$ (and entirely within $(t_1, \lambda_{pq}]$) a Delaunay crossing by w. Finally, r does not cross pq in $(t_1, \lambda_{pq}]$.

(B2) The points p, q, w, r are co-circular at times $\delta_0 \in I$, $\delta_1 \in [t_1, \lambda_q]$, and $\delta_2 \in (\delta_1, \lambda_{pq}]$, and the following properties hold:

(i) The co-circularity at time δ_0 is red-blue with respect to pq.
(ii) The co-circularity at time δ_1 is red-red with respect to pq and red-blue with respect to the edge rq, whose Delaunayhood is violated right after time δ_1 by p and w. (In particular, this implies that r remains in L_{pq}^+ throughout (t_1, δ_1), after entering this halfplane during I.) Furthermore, the open cap $C_{pq}^+ = B[p,q,w] \cap L_{pq}^+$ contains no points of P at time δ_1.
(iii) The co-circularity at time δ_2 is again red-blue with respect to pq. It arises during a single Delaunay crossing of qp by w, which occurs in $\mathrm{DT}(\{p,q,r,w\})$ during some sub-interval of $(\delta_1, \lambda_{pq}]$.

(B3) The set $R \setminus \{p\}$ induces a (single) Delaunay crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$, where w crosses rq, from L_{rq}^+ to L_{rq}^-, during (δ_1, λ_q).

(B4) There exists a time $\xi_{qw} > \lambda_{pq}$ so that (i) the edge qw is Delaunay at time ξ_{qw}, and (ii) the edges qw and pw belong to, respectively, $\mathrm{DT}(R \setminus \{p,r\})$ and $\mathrm{DT}(R \setminus \{q.r\})$ throughout the interval $[\delta_1, \xi_{qw}]$.

Notice that any such quadruple $\varrho = (p,q,r,w)$ in R is charged as above by at most one 3-restricted right special quadruple $\chi = (a,p,w,q)$ in $\mathcal{F}$ (with respective outer points r and u), because the latter quadruple is uniquely determined by each of the triples (p,q,r) and (p,q,w).

We say that a quadruple $\varrho = (p,q,r,w)$ is *terminal of type B* if it satisfies the above four conditions with respect to the underlying set R. (In Section 7, we shall again use P to denote the underlying set of our terminal quadruples.)

Let Σ_R^B denote the resulting family of terminal quadruples $\varrho = (p,q,r,w)$ (of type B) in R that are charged by 3-restricted right special quadruples in P through the above probabilistic argument.

Lemma 6.13. *With the above assumptions, each terminal quadruple $\varrho = (p,q,r,w)$ in Σ_R^B is uniquely determined by each of its sub-triples $(p,q,r), (p,q,w), (q,r,w)$.*

Proof. By condition (B1), w is the first point of P to hit the edge pq after its Delaunay crossing $(pq, r, I = [t_0, t_1])$ by r. Hence, $\varrho = (p,q,r,w)$ is uniquely determined by the choice of p, q and r. A similar agrument implies that $\varrho = (p,q,r,w)$ is uniquely determined by the triple (p,q,w).

To see that ϱ is uniquely determined by (q,r,w), let us assume for a contradiction that Σ_R^B contains another such quadruple $\varrho' = (p', q, r, w)$ (of type B and with $p' \neq p$). Furthermore, assume with no loss of generality that the respective counterclockwise (q,r)-crossing $(p'q, r, I')$ of ϱ' ends after $I = [t_0, t_1]$. Note though that I' must end before w enters L_{pr}^- through pr (as prescribed by condition (A3)). However, in that case I' would end in (t_1, δ_1), so q' would be included in the respective set A_{rq} of ϱ, and, thereby, omitted[40] from R, contrary to the choice of ϱ' in Σ_R^B. $\square$

Let $T^B(m)$ denote the maximum cardinality of any family Σ^B of terminal quadruples of type B (with the uniqueness property stated in Lemma 6.13) that can be defined over a set P of m moving points. The preceding discussion implies that the number of special quadruples that fall into case (d2) is at most

$$O\left(h^4 T^B(n/h) + \ell h^2 N(n/h) + \ell h n^2 \beta(n)\right).$$

We delegate the analysis of terminal quadruples of type B to Section 7.2. Note that the points of each such terminal quadruple $\varrho = (p,q,r,w)$ perform at least three Delaunay crossings (namely, the crossings of pq by r and w, and the crossing of qr by w). Hence, it suffices to enforce two more crossings in order to ensure that some sub-triple of ϱ be involved in *two* distinct Delaunay crossings.

As in the case of terminal quadruples of type A, we shall exploit the co-circularity at time δ_2, which is red-blue with respect to rw, in order to enforce a Delaunay crossing of that edge by at least one of the two points p, q. In addition, we shall enfore a Delaunay crossing of pw by at least one of r, q (during which pw will be hit by r or q, as suggested by Lemma 6.11 (i) and depicted in Figure 79 (top)).

[40] Clearly, we have $p' \neq w$ (i.e., there is no crossing (wq, r, I')), because r is already assumed to cross the line L_{wq}, from L_{wq}^- to L_{wq}^+, *outside* pw (as prescribed by Lemma 6.11 (ii)).

3-**restricted right special quadruples–wrap up.** Putting together the previously established bounds on the maximum possible numbers of 3-restricted right special quadruples that fall into cases (a), (b), (c), (d1) and (d2) yields the following recurrence:

$$\Phi_3^R(n) =$$
$$O\left(h^4 T^A(n/h) + h^4 T^B(n/h) + \ell h^2 N(n/h) + k\ell^2 N(n/\ell) + k^2 N(n/k) + \ell h n^2 \beta(n)\right). \qquad (10)$$

Discussion. Notice that the roles of p and q in subcases (d1) and (d2) are largely symmetric, which enables us to enforce a Delaunay crossing of the respective edge pr or rq by at least one of the remaining two points of p, q, r, w. In both scenarios, we first apply Theorem 2.2 (with threshold $h \gg \ell$) in order to extend the (almost-)Delaunayhood of pr or qr from $I = [t_0, t_1]$ (where pq undergoes the Delaunay crossing by r) to a larger interval. Lemma 6.6 implies that each event, that arises within the respective red-blue arrangement $\mathcal{A}_{pr}$ or $\mathcal{A}_{rq}$ during the gap interval, can be traced back to χ (via (pq, r, I)) in only $O(\ell)$ possible ways.

The main difference between the two subcases stems from condition (S6), according to which wq is almost-Delaunay in the interval $[\lambda_{wq}, \xi_{wq}]$, and is *fully Delaunay* at the endpoints λ_{wq}, ξ_{wq}. Since the latter interval contains δ_1, in subcase (d1) the corresponding Lemma 6.7 (i) immediately yields a Delaunay crossing of wq by (at least) one of the points p, r.

In subcase (d2), however, we only know that pw belongs throughout $[\lambda_{wq}, \xi_{wq}]$ to some reduced triangulation $\mathrm{DT}(P \setminus A_{pw})$, where A_{pw} is a subset of cardinality at most $6\ell + 3$ which includes a, r, u, and perhaps also q. That is, we are not necessarily able to restore the Delaunayhood of pw at times λ_{wq} and ξ_{wq} without removing some of r, q, and thereby destroying $\varrho = (p, q, r, w)$. In fact, it is not even known whether the collinearity mentioned in Lemma 6.11 (i) occurs in $[\lambda_{wq}, \xi_{wq}]$ or *before* λ_{wq}. In Section 7.2 we use conditions (B1)–(B4) obtained above, to enforce the long-awaited crossing of pw by q or r.

6.6 Stage 4: The number of left special quadruples

To bound the maximum possible number $\Phi_3^L(n)$ of 3-restricted right special quadruples, we fix the underlying set P of n moving points, and a refined family $\mathcal{F}$.

Topological setup. According to Proposition 6.2, any 3-restricted left special quadruple $\chi = (a, p, w, q)$ shares its triple (p, q, w) with at most two other such quadruples. (In other words, it suffices to bound the overall number of the corresponding triples (p, q, w).) We strengthen the above property, by considering at most *one* 3-restricted left quadruple for each triple (p, q, w). Therefore, in what follows every special quadruple $\chi = (a, p, w, q)$ under our consideration will be uniquely determined by its triple (p, q, w).

To proceed, we fix a 3-restricted left special quadruple $\chi = (a, p, w, q)$, with respect to P and $\mathcal{F}$, whose two special (a, q)-crossings take place during the intervals $\mathcal{I}_r = [\lambda_0, \lambda_1]$ and $\mathcal{J}_u = [\lambda_2, \lambda_3]$ (in this order), where r and u are the respective outer points. Recall that the original "regular" family $\mathcal{F}$ includes the quadruples $\sigma_1 = (p, q, a, r)$ and $\sigma_2 = (w, q, a, u)$.

By assumption, χ satisfies the six conditions (S1)–(S2), (S3b), and (S4)–(S6). We emphasize that all these conditions, except for (S3b), are common to *all* 3-restricted special quadruples, including the right special quadruples studied in Section 6.5. Moreover, one can switch the roles p and w by reversing the direction of the time axis, so our condition (S3b) of left special quadruples is fully symmetric to condition (S3a) on right special quadruples (which has been assumed throughout the analysis Section 6.5). See below for details.

Refer to Figure 82. As reviewed in the preceding Section 6.6, the 3-restrictedness of χ implies that there exist times $\lambda_{wq} \leq \lambda_0$, $\xi_{pq} \leq \lambda_{wq}$, $\lambda_{pq} \geq \lambda_3$ and $\xi_{wq} \geq \lambda_{pq}$, whose properties have been summarized in the beginning of that section. In particular, pq is Delaunay at times λ_{pq} and ξ_{pq}, and wq

is Delaunay at the symmetric times λ_{wq} and ξ_{wq}. Furthermore, pq and wq are almost Delaunay during, respectively, $[\xi_{pq}, \lambda_{pq}]$ and $[\lambda_{wq}, \xi_{wq}]$.

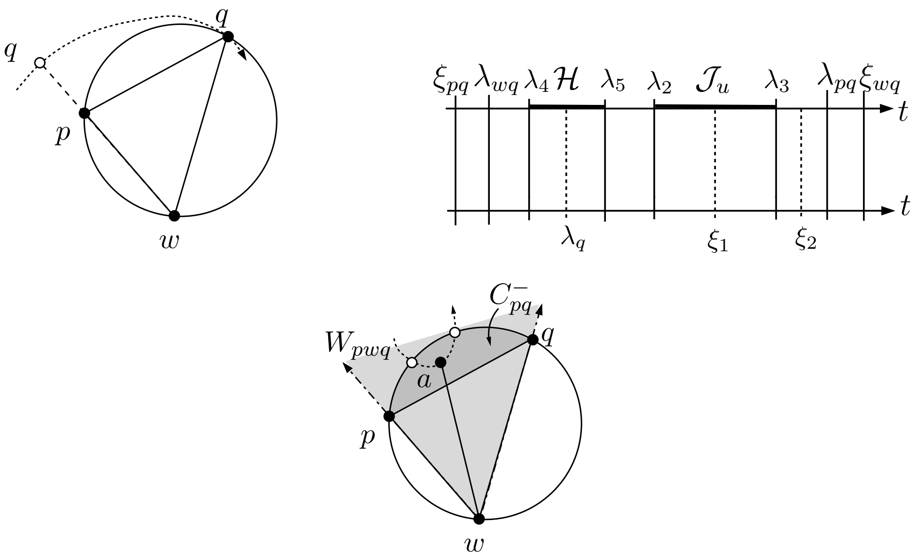

Figure 82: The topological setup during the interval $(\lambda_q, \xi_2) \subseteq [\xi_{pq}, \xi_{wq}]$. Left: The edge qw is hit at some time $\lambda_q \in [\lambda_{wq}, \lambda_2)$ by p, so it undergoes a Delaunay crossing $(qw, p, \mathcal{H} = [\lambda_4, \lambda_5])$ within $\mathrm{DT}(P \setminus \{a, r, u\})$. Right: We have $\xi_{pq} \leq \lambda_4 \leq \lambda_q < \lambda_5 < \xi_1 < \xi_2 \leq \lambda_{pq}$. Bottom: The motion of $B[p, q, w]$ is continuous throughout $(\lambda_q, \xi_2]$ (the hollow circles represent the co-circularities at times ξ_1 and ξ_2).

Let us summarize what we know so far about the motion of a, p, w, q if $\chi = (a, p, w, q)$ is a 3-restricted left special quadruple. By Condition (S3b), these points are co-circular at times $\xi_0 \in \mathcal{I}_r \setminus \mathcal{J}_u$, and $\xi_1 \in \mathcal{J}_u \setminus \mathcal{I}_r$, and $\xi_2 \in (\lambda_3, \lambda_{pq}]$. Moreover, the Delaunayhood of pq is violated, throughout (ξ_1, ξ_2), by the points $a \in L^-_{pq}$ and $w \in L^+_{pq}$. In particular, a lies throughout that interval within the wedge $W_{pwq} = L^+_{wp} \cap L^-_{wq}$ and inside the cap $C^-_{pq} = B[p, q, w] \cap L^-_{pq}$.We emphasize that the order type of the quadruple (q, p, w, a) remains unchanged during (ξ_1, ξ_2).

In addition, by the same Condition (S3b), the smaller set $P \setminus \{a, r, u\}$ yields a (single) Delaunay crossing $(qw, p, \mathcal{H}_\chi)$, whose interval $\mathcal{H} = \mathcal{H}_\chi = [\lambda_4, \lambda_5]$ is contained in $[\lambda_{wq}, \lambda_2)$. Specifically, w hits pq at some moment[41] $\lambda_q \in \mathcal{H}$, when p crosses L_{wq} from L^+_{wq} to L^-_{wq}. Since p lies in L^-_{wq} at times ξ_1 and ξ_2, no further collinearities of p, w, q can occur during $[\lambda_q, \xi_2)$. (Otherwise, the point p would have to re-enter L^+_{wq} before ξ_2, and then the triple p, q, w would be collinear three times, contrary to our assumptions.) To conclude, the disc $B[p, q, w]$ moves continuously throughout the interval $(\lambda_q, \xi_2]$, which is obviously contained in $[\xi_{pq}, \lambda_{pq}] \cap [\lambda_{wq}, \xi_{wq}] = [\lambda_{wq}, \lambda_{pq}]$.

Overview. We fix three constant parameters k, ℓ, h, such that $12 < k \ll \ell \ll h$, and distinguish between four possible cases. The first three cases (a)–(c) are fully symmetric to the cases (a)–(c) that we encountered in Section 6.5 when handling right quadruples. (Moreover, the first two cases (a) and (b) are very similar to the the corresponding cases (a) and (b) in Section 5.6.)

In the final, most involved, case (d), we re-introduce at last the outer point u. (The other outer point r is not used in the analysis of left special quadruples.) The correspondence between $(wa, q, \mathcal{J}_u)$ and its ancestor quadruple $\sigma_2 = (w, q, a, u)$ in $\mathcal{F}$ implies that we have a single Delaunay crossing $(wq, u, I = [t_0, t_1])$ (which is the first among the two (w, u)-crossings of σ_2). Since the points u and p

[41] Recall from Section 6.2 that p can cross qw either before or after ξ_0, depending on the location of w when q crosses pa. Our analysis only relies on the fact that $\lambda_q < \xi_1 < \xi_2$.

cross the same edge wq in opposite directions, χ can again be charged to the resulting terminal quadruple (w, q, u, p).

After ruling out cases (a)–(c), we may assume, in the last remaining case (d), that a total of at most $8\ell + 1$ points of P appear in the cap $C^-_{wq} = B[p, q, w] \cap L^-_{wq}$ during (λ_q, ξ_2). (Notice that this condition is fully symmetric to the one in Lemma 6.4.)

As in Section 6.5, we use the interplay between $\chi = (a, p, w, q)$ and $\sigma_2 = (w, q, a, u)$ to enforce as many Delaunay crossings as possible among w, q, u, p before charging χ to this terminal quadruple. Our analysis is largely simplified[42] by the property that the interval $I = [t_0, t_1]$ of the first crossing of σ_2 is entirely contained in the above interval (λ_q, ξ_2); see below for details.

We establish symmetric variants of Lemmas 6.7 and 6.11. Namely, we argue that (i) the edge wu is hit in (λ_q, t_0) by at least one of p, q, or else (ii) the edge uq is hit in (λ_q, t_0) by at least one of p, r.[43] In the first case (denoted as (d1)), we also show that u hits qp in (λ_q, t_0). In the second case (denoted as (d2)) we similarly show that u also hits pw in (λ_q, t_0). In both scenarios, we invoke Theorem 2.2 to amplify the above two additional collinearities into full-fledged Delaunay crossings. Therefore, by the time we charge χ to the terminal quadruple ϱ, its various sub-triples among w, q, u, p perform four Delaunay crossings (where some of these crossings occur in appropriately reduced subsets of P).

In Section 7 we express the number of such terminal quadruples, which arise in the analysis of left special quadruples, in terms of more elementary quantities, that were introduced in Section 2. To do so, we enforce an additional, fifth crossing among w, q, u, p (namely, the crossing of pu by w or q). As a result, some sub-triple among w, q, u, p is involved in two Delaunay crossings, so our analysis bottoms out via Lemma 4.5.

In what follows, we consider the family $\mathcal{G}^L_{pw}$ of all 3-restricted left special quadruples of the form $\chi' = (a', p, w, q')$, which share their middle pair with χ. We may assume that each $\chi' = (a', p, w, q') \in \mathcal{G}^L_{pw}$ is uniquely determined by the choice of q' (as the only "free" point in the triple (p, q', w)). Note that the set $P_{\chi'}$ of each χ' includes, in addition to the four points a', p, w, q' of χ', the respective outer points r' and u' of its special crossings $(pa', q', \mathcal{I}_{r'})$ and $(wa', q', \mathcal{J}_{u'})$. Furthermore, each of these quadruples $\chi' \in \mathcal{G}^L_{pw}$ is accompanied by a counterclockwise (w, p)-crossing $(q'w, p, \mathcal{H}_{\chi'} = \mathcal{H}')$, which occurs within the smaller triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$. See Figure 83. We use $\lambda_{q'}$ to denote the time in $\mathcal{H}'$ when the respective point q' of χ' enters the halfplane L^+_{wp} (or, equivalently, when p crosses $q'w$ from $L^+_{wq'} = L^-_{q'w}$ to $L^+_{q'w}$).

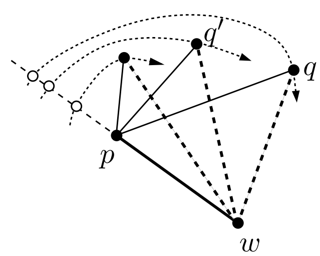

Figure 83: Each left special quadruple $\chi' = (a', p, w, q') \in G^L_{pw}$ (with respective outer points r' and u') comes with a counterclockwise (w, p)-crossing $(q'w, p, \mathcal{H}_{\chi'})$, which occurs within $\mathrm{DT}(P \setminus \{a', r', u'\})$.

Notice that Lemma 5.5 readily generalizes to the above (w, p)-crossings. Namely, a pair of such crossings $(qw, p, \mathcal{H}_\chi)$ and $(q'w, p, \mathcal{H}_{\chi'})$, which occur within the respective triangulations $\mathrm{DT}(P \setminus \{a, r, u\})$ and $\mathrm{DT}(P \setminus \{a', r', u'\})$, are *compatible*, provided that $q' \neq a, r, u$ and $q \neq a', r', u'$, in the sense that

[42] In contrast, in the *almost*-symmetric case of right special quadruples we did not know whether the first crossing (pq, r, I) of $\sigma_1 = (p, q, a, r)$ at all overlaps (ξ_{-1}, λ_q).

[43] These collinearities are fairly symmetric to the crossings of pr and rq that we enforced in Section 6.5.

the orders in which the intervals $\mathcal{H}_\chi$ and $\mathcal{H}_{\chi'}$ begin or end are both consistent with the time stamps λ_q and $\lambda_{q'}$.

Clearly, for any special quadruple $\chi = (a, p, w, q) \in \mathcal{G}^L_{pw}$ (with outer points r and u) the family $\mathcal{G}^R_{pw}$ includes at most three other quadruples $\chi' = (a', p, w, q')$ whose respective points q' are equal to one of a, r or u. The pigeonhole principle then implies that at least *one quarter* of all quadruples $\chi = (a, p, w, q)$ in $\mathcal{G}^L_{pw}$ satisfy the following condition:

(PHL1) *There exist at most three quadruples $\chi' \in \mathcal{G}^L_{pw}$ with $q \in \{a', r', u'\}$.*

Since p and w are arbitrary points of P, (PHL1) holds for at least a quarter of all 3-restricted left special quadruples under consideration; hence we may assume that it holds for the special quadruple χ at hand. Therefore, for all but 6 quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^L_{pw} \setminus \{\chi\}$ (with respective outer points r' and u') their respective (w, p)-crossings $(q'w, p, \mathcal{H}_{\chi'})$ are compatible with $(qw, p, \mathcal{H})$ via a suitable extension of Lemma 5.5.

With the above preparations, we can now proceed with our case analysis.

Case (a). For at least k of the above quadruples $\chi' = (a', p, w, q') \in \mathcal{G}^L_{pw}$, their respective (w, p)-crossings $(q'w, p, \mathcal{H}')$ either begin in $[\xi_{pq}, \lambda_4)$, or end in $(\lambda_5, \lambda_{pq}]$. Refer to Figure 84. Recall that, by condition (S5), the edge pq is Delaunay at each of the times ξ_{pq} and λ_{pq}, and that it is almost Delaunay during the entire interval $[\xi_{pq}, \lambda_{pq}]$.

To bound the number of such quadruples χ that fall into case (a), we pass to a random sub-sample $\hat{P}$ of $n/4$ points in P, and argue that, with some fixed positive probability, the crossing $(qw, p, \mathcal{H})$ becomes $(q, p, \Theta(k))$-chargeable there, for the reference interval $[\xi_{pq}, \lambda_{pq}]$. Therefore, Theorem 5.3 implies that the overall number of such triples (p, q, w) in P does not exceed

$$O\left(k^2 N(n/k) + kn^2\beta(n)\right),$$

which also bounds the overall number of the corresponding 3-restricted left special quadruples χ.

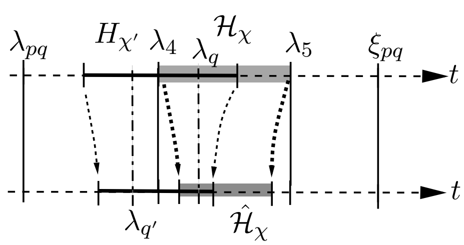

Figure 84: Case (a): At least k counterclockwise (w, p)-crossings $(q'w, p, \mathcal{H}_{\chi'})$ either begin in $[\xi_{pq}, \lambda_4)$ or end in $(\lambda_5, \lambda_{pq}]$ (one such crossing of the former type is depicted). Then, with some fixed and positive probability, the sample $\hat{P}$ yields a Delaunay crossing $(qw, p, \hat{\mathcal{H}}_\chi)$ that is $(q, p, \Theta(k))$-chargeable with respect to $[\xi_{pq}, \lambda_{pq}]$.

Preparing for cases (b) and (c): Charging events in $\mathcal{A}_{pw}$. We may assume, from now on, that there exist at most k special quadruples $\chi' \in \mathcal{G}^L_{pw}$ whose respective (w, p)-crossings $(q'w, p, \mathcal{H}')$ either begin in $[\xi_{pq}, \lambda_4)$, or end in $(\lambda_5, \lambda_{pq}]$.

Before proceeding to the following cases, we apply Theorem 2.2 in $\mathcal{A}_{pw}$ in order to extend the almost-Delaunayhood of pw from $\mathcal{H} = [\lambda_4, \lambda_5]$ to $[\xi_{pq}, \lambda_{pq}]$. We emphasize that $[\xi_{pq}, \lambda_{pq}] \setminus \mathcal{H}$ consists of two intervals $[\xi_{pq}, \lambda_4)$ and $(\lambda_5, \lambda_{pq}]$ (where the former interval can be empty), which we consider separately. Note also that the edge pw belongs during $\mathcal{H}$ to the reduced triangulation $\mathrm{DT}(P \setminus \{a, r, u\})$ (but not necessarily to $\mathrm{DT}(P)$), so Theorem 2.2 must be applied, for each of these two intervals, with respect to $P \setminus \{a, r, u\}$.

In each of these applications, in cases (i) and (ii) we charge χ (via its respective (w,p)-crossing $(qw,p,\mathcal{H})$) to $(\ell+3)$-shallow collinearities and co-circularities that occur in the full red-blue arrangement $\mathcal{A}_{pw}$. Since case (a) has been ruled out, the charging is almost unique, and accounts for at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2\beta(n)\right)$ left special quadruples.

At the end, we have either disposed of χ through (conditions (i), (ii) of) Theorem 2.2 or ended up with a set A_{pw} of at most $6\ell+3$ points whose removal restores the Delaunayhood of pw throughout $[\xi_{pq}, \lambda_{pq}]$. Namely, A_{pw} is composed of a,r,u, and of the two sets of at most 3ℓ points each, which are obtained by separately applying Theorem 2.2, within $\mathcal{A}_{pw}$, over the intervals (ξ_{pq},λ_4) and (λ_5,λ_{pq}). Hence, we may assume, in what follows, that the above set A_{pw} exists.

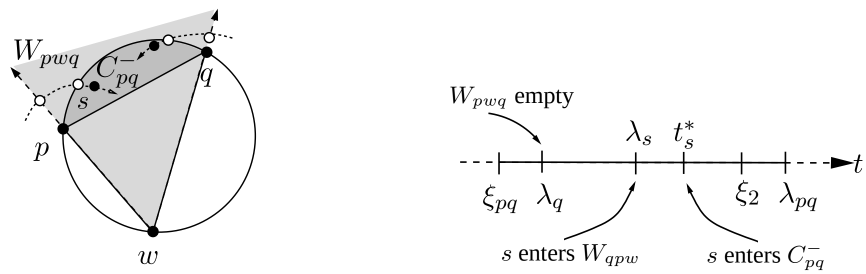

Figure 85: Case (b). At least ℓ points $s \neq a,r,u$ visit the cap C^-_{pq} during (λ_q,ξ_2). Each of them must enter the wedge W_{pwq} (through one of the rays $\vec{wp},\vec{wq}$, outside the respective edges pw and wq) after time λ_q and then enter the cap C^-_{pq} (through the boundary of $B[p,q,w]$).

Case (b). A total of at least ℓ points of P, distinct from a,r,u, appear in the cap $C^-_{pq} = B[p,q,w] \cap L^-_{pq}$ at some time during the interval (λ_q,ξ_2). (Note that some of these points s may belong to A_{pw}.) Recall that λ_q denotes the time in $\mathcal{H}$ when p enters L^-_{wq}, through wq, and that no additional collinearities of p,q,w can occur during (λ_q,ξ_2), so the motion of $B[p,q,w]$ is fully continuous in that interval.

Refer to Figure 85. Let $s \in P \setminus \{a,r,u\}$ be one of the points that visit C^-_{pq} during (λ_q,ξ_2). Since the above cap C^-_{pq} is fully contained there in the wedge $W_{pwq} = L^+_{wp} \cap L^-_{wq}$, s must enter W_{pwq} after time λ_q (when W_{pwq} co-incides with the single ray $\vec{wp} = \vec{wq}$) through one of the rays $\vec{wp},\vec{wq}$. We also note that, by condition (S5) (and since $(\lambda_q,\xi_2) \subseteq [\xi_{pq},\lambda_{pq}]$), the edge pq is Delaunay in $P \setminus \{a,w,r,u\}$ throughout (λ_q,ξ_2), so s, which has to enter C^-_{pq} before it enters W_{pwq}, can do so only through the boundary of $B[p,q,w]$. This results in a co-circularity of p,q,w,s, and is easily seen to imply that s enters W_{pwq} by crossing one of the rays $\vec{wp}$ or $\vec{wq}$ *outside* the respective edges wp or wq.

In what follows, we assume that s is among the last ℓ points to leave C^-_{pq} during (λ_q,ξ_2). Let t^*_s denote the time of the corresponding co-circularity of p,q,w,s, which occurs when s leaves C^-_{pq} through the boundary of $B[p,q,w]$. Since χ satisfies condition (S5), the opposite cap $C^+_{pq} = B[p,q,w] \cap L^+_{qp}$ contains no points of $P \setminus \{a,r,u\}$ at time t^*_s. (Otherwise, the Delaunayhood of wq would be violated by s and any of these points.) Therefore, the co-circularity at time t^*_s has to be $(\ell-1)$-shallow in $P \setminus \{a,r,u\}$, and thus $(\ell+2)$-shallow in P.

Note also that the co-circularity at time t^*_s is red-blue with respect to the edge pq, which is violated right before it by w and s. Lemma 4.1, together with the choice of $s \neq a,r,u$, imply that this co-circularity cannot occur during the crossing $(qw,p,\mathcal{H}_\chi = [\lambda_4,\lambda_5])$ (which occurs in $P \setminus \{a,r,u\}$), so $t^*_s > \lambda_5$.

As in the symmetric case (b) of Section 6.5, we distinguish between two possible subcases. In each of them we manage to dispose of χ by charging it, within one of the arrangements $\mathcal{A}_{wq},\mathcal{A}_{pw}$, either to $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities, or to a (2ℓ)-shallow collinearity.

Case (b1). At least half of the above points s cross the line L_{wq}, from L^+_{wq} to L^-_{wq}, during (λ_q,t^*_s). (This also includes points s that possibly cross L_{wq} outside the ray $\vec{wq}$, before entering W_{pwq} through the other ray $\vec{wp}$.) By Condition (S6) (and since $(\lambda_q, t_s^*) \subseteq (\lambda_q, \xi_2) \subseteq [\lambda_{wq}, \xi_{wq}]$), each of these crossings occurs outside wq, within one of the outer rays of L_{wq}.

For each s we argue, exactly as in Section 5.6, that the points w, q, s are involved during $(\lambda_q, t_s^*) \subseteq (\lambda_q, \xi_2)$ either in a (2ℓ)-shallow collinearity, or in $\Omega(\ell)$ (2ℓ)-shallow co-circularities. That is, right after s enters L_{wq}^- at time λ_q (outside wq), the disc $B[w, q, s]$ "swallows" the entire halfplane L_{wq}^+. (In addition, s must remain in L_{wq}^- until time t_s^*, for otherwise the points w, q, s would be collinear more than twice.) If this disc, which contains at most $\ell + 2$ points at the end of the process, contains at least 2ℓ points at time λ_q, then each of the last $\ell - 2$ resulting co-circularities are (2ℓ)-shallow (in P). Otherwise, the collinearity of q, p, s is (2ℓ)-shallow.

Since s can be chosen in at least $\Omega(\ell)$ different ways, the points w and q are involved during (λ_q, ξ_2) either in $\Omega(\ell^2)$ (2ℓ)-shallow co-circularities, or in a (2ℓ)-shallow collinearity. In both cases, we charge χ to these events.

Note that each (2ℓ)-shallow event, which occurs in $\mathcal{A}_{pq}$ at some time $t^* \in (\lambda_q, \xi_2)$, can be traced back to $(qw, p, \mathcal{H})$ (and, by Proposition 6.2, also to χ) in at most $O(1)$ possible ways because p is among the last four points to hit the edge wq before time t^*, according to condition (S6). Hence, the above scenario happens for at most $O(\ell^2 N(n/\ell) + \ell n^2 \beta(n))$ special quadruples χ.

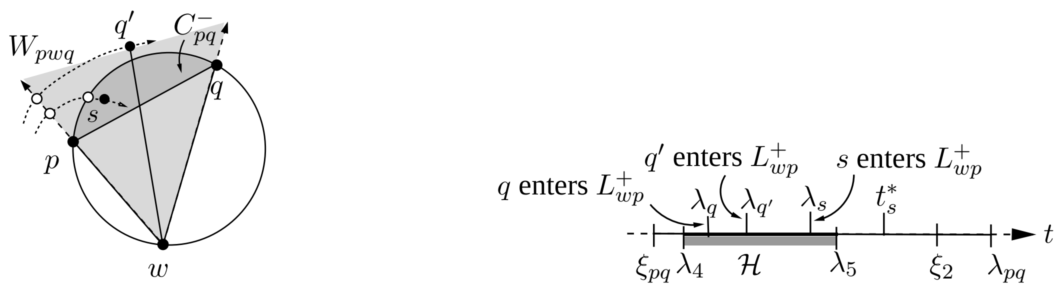

Figure 86: Proposition 6.14. Left: q is among the last $k+7$ candidates q' to enter L_{wp}^+ before time λ_s. Right: The various critical events occur in the depicted order. Note that λ_s may occur in (the second part of) $\mathcal{H} = [\lambda_4, \lambda_5]$.

Case (b2). At least half of the above points $s \neq a, r, u$ remain in L_{wq}^- throughout the respective intervals (λ_q, t_s^*). Each of these points must enter W_{pwq}, also during (λ_q, t_s^*), through the ray emanating from p in direction $\vec{wp}$, thereby crossing L_{pw} from L_{wp}^- to L_{wp}^+. (See Figure 86 (left). Recall that such a collinearity can occur at most once, because the triple p, w, s can be collinear at most twice.)

We again fix one of these points s, and use λ_s to denote the corresponding time in (λ_q, t_s^*) when s enters W_{pwq} through the ray emanating from w in direction pw. As in the previous case, we conclude that either the collinearity of p, w, s at time t_s is (2ℓ)-shallow, or the points p, w, s are involved in $\Omega(\ell)$ (2ℓ)-shallow co-circularities during the preceding interval (λ_s, t_s^*). As in the matching scenarios (b2) in Sections 5.6 and 6.5, the main challenge is to argue that each of the above (2ℓ)-shallow events, which occur in $\mathcal{A}_{pw}$ during $(\lambda_s, t_s^*] \subseteq (\lambda_q, \xi_2)$, can be traced back to χ in at most $O(k)$ ways.

To show this, let $t^* \in (\lambda_q, \xi_2)$ be the time of a (2ℓ)-shallow collinearity or co-circularity that occurs in $\mathcal{A}_{pw}$. First, we guess the points p and w of χ in $O(1)$ possible ways among the three or four points involved in the event. We next recall that, in the charging scheme of case (b2), each (2ℓ)-shallow co-circularity or collinearity or collinearity that we charge in $\mathcal{A}_{pw}$ is obtained via some point s, which is also involved in the event, that enters L_{wp}^+ at the respective time λ_s. We, therefore, guess s among the remaining one or two points involved in the event at time t^*. To guess the remaining points a and q of χ, we examine all "candidate" special quadruples $\chi' \in \mathcal{G}_{pw}^L$ whose two "middle" points (p, w) are shared with χ. Recall that each of these quadruples is accompanied by the (w, p)-crossing $(q'w, p, \mathcal{H}' = \mathcal{H}_{\chi'})$, where q' enters L_{wp}^+ at the respective time $\lambda_{q'} \in \mathcal{H}'$. Recall also that χ' is uniquely determined by the choice of q' (as long as p and w remain fixed).

Clearly, it suffices to consider only special quadruples $\chi' = (a', p, w, q')$ in $\mathcal{G}^L_{pw}$ with the following properties: (1) $s \neq a', r', u'$, where r' and u' are the outer points of χ', (2) $\lambda_{q'} < \lambda_s$, and (3) s lies in $L^-_{wq'}$ during the second portion of $\mathcal{H}_{\chi'}$ (after $\lambda_{q'}$). This is because each of these conditions holds for χ and s in the charging scheme of case (b2). For example, (3) follows because case (b1) does not occur for s (and $t^*_s > \lambda_5$).

If a special quadruple $\chi' = (a', p, w, q') \in \mathcal{G}^L_{pw}$ satisfies the above three conditions (1)–(3), we say that the respective point q' (which uniquely determines χ') is a candidate (for q).

The following symmetric variant of Proposition 6.3 guarantees that each (2ℓ)-shallow event, which occurs in $\mathcal{A}_{pw}$ at some fixed time $t^* \in (\lambda_q, \xi_2)$, is charged by at most $k+7$ quadruples in $\chi' \in \mathcal{G}^L_{pw}$, because its points q is among the last $k+7$ similar candidates q' to enter L^+_{wp} before time λ_s. See Figure 86.

Proposition 6.14. *With the above assumptions, the point q is among the last $k+7$ candidates q' to enter the halfplane L^+_{wp} before λ_s.*

We omit the fairly technical proof of Proposition 6.14, noting that it is fully symmetric to the proof of Proposition 6.3, and very similar to the proof of Proposition 5.6.

Repeating the same charging argument for each of the $\Omega(\ell)$ possible choice of s shows that at most $O\left(k\ell^2 N(n/\ell) + k\ell n^2\beta(n)\right)$ special quadruples can fall into case (b2).

Case (c). A total of at least ℓ points $s \in P \setminus A_{pw}$ appear in the cap $C^-_{wp} = B[p,q,w] \cap L^-_{wp}$ at some time during (λ_q, ξ_2). Here A_{pw} continues to denote the subset of at most $6\ell+3$ points, including a, r and u, whose removal restores the Delaunayhood of pw throughout the interval $[\xi_{pq}, \lambda_{pq}]$. (Recall that A_{pw} was obtained by applying Theorem 2.2 in $\mathcal{A}_{pw}$, after ruling out case (a).)

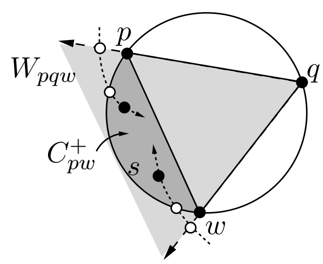

Figure 87: Case (c). A total of at least ℓ points $s \in P \setminus A_{pw}$ enter in the cap C^-_{wp} during (λ_q, ξ_2). Each of them must enter the wedge W_{pqw} (through one of the rays $\vec{qp}, \vec{qw}$, outside the respective edges pq and wq), and only then cap C^-_{wp} (through the boundary of $B[p,q,w]$).

Clearly, C^-_{wp} is contained in the wedge $W_{pqw} = L^+_{pq} \cap L^-_{wq}$, which shrinks at time λ_q to the ray $\vec{qp} = \vec{qw}$. Hence, each of these points s has to enter W_{pqw} and C^-_{wp} (in this order) before time λ_q. Furthermore, s can leave C^+_{pw} only through the boundary of $B[p,q,w]$, at a co-circularity of p, q, w, s. (Otherwise s would have to hit pw and, therefore, belong to A_{pw}.) In addition, s can leave W_{pqw} only through one of the rays $\vec{qp}$ and $\vec{qw}$ (outside the respective segments qp, qw). See Figure 87.

As in the previous case (b), we may assume that each s under consideration is among the first ℓ such points of $P \setminus A_{pw}$ to enter C^-_{wp} during (λ_q, ξ_2), and use t^*_s to denote the time of the respective co-circularity. Clearly, the opposite cap $C^+_{wp} = B[p,q,w] \cap L^+_{wp}$ contains then no points of $P \setminus A_{pw}$. Indeed, otherwise the Delaunayhood of pw would be violated by s and any one of these points (contrary to our assumption that $pw \in \mathrm{DT}(P \setminus A_{pw})$ throughout $[\xi_{pq}, \lambda_{pq}] \supset (\lambda_q, \xi_2)$). Hence, the resulting co-circularity of p, q, w, s at time t^*_s is $(7\ell+2)$-shallow in P, because, at the time of co-circularity, the circumdisc $B[p,q,w] = B[p,s,w]$ can contain in its interior at most $6\ell+3$ points of A_{pw} and at most $\ell - 1$ points of $P \setminus A_{pw}$.

Case (c1). If at least half of the above points s cross the line L_{wq} (from L_{wq}^+ to L_{wq}^-) during their respective intervals (λ_q, t_s^*), then we argue exactly as in subcase (b1).

Namely, we fix one of the these points s and notice that, right after s enters L_{wq}^- outside wq, the disc $B[w,q,s]$ contains the entire halfplane L_{wq}^-. Therefore, the points p, q, s are involved, during (λ_q, t_s^*), either in an (8ℓ)-shallow common collinearity (which occurs when s enters L_{wq}^-), or in $\Omega(\ell)$ (8ℓ)-shallow co-circularities.

We repeat the above argument for each of the $\ell/2$ possible choices of s and charge χ within $\mathcal{A}_{wq}$ (via $(qw, p, \mathcal{H})$) to the above (8ℓ)-shallow events. As in case (b1), each (8ℓ)-shallow collinearity or co-circularity occurs during (λ_q, ξ_2), and involves w and q, so it is charged by at most $O(1)$ special quadruples χ (because χ is uniquely determined by (p,q,w) and p is among the last four points to hit wq before the respective time t^* of the event).

Case (c2). We may assume, then, that at least half of the above points s enter W_{pqw} through the ray $\vec{qp}$. For each of these points s, the triple q, p, s are involved during (λ_q, t_s^*) either in an (8ℓ)-shallow collinearity, or in $\Omega(\ell)$ (8ℓ)-shallow co-circularities. As before, we repeat the above argument for the $\ell/2$ eligible choices of s and charge χ, within $\mathcal{A}_{pq}$, either to $\Omega(\ell^2)$ (8ℓ)-shallow co-circularities or to an (8ℓ)-shallow collinearity.

We claim that each of the resulting (8ℓ)-shallow events, which occur in $\mathcal{A}_{pq}$ during (λ_q, ξ_2), can be traced back to χ in at most $O(1)$ possible ways. Indeed, fix any of the above events, at some time $t^* \in (\lambda_q, \xi_2)$. We first guess p and q in $O(1)$ possible ways among the three or four points involved in the event.

To guess the point a (which would immediately determine $(pa, q, \mathcal{I}_r)$ and thereby also χ), we consider all special (p,q)-crossings $(pa', q, \mathcal{I}_{r'})$ (in $\mathcal{F}$) and recall that, according to conditions (S1) and (S5), at most $O(1)$ such crossings can begin during $[\xi_{pq}, \lambda_0)$ or end during $(\lambda_1, \lambda_{pq}]$. Notice also that the interval $[\lambda_{pq}, \xi_{pq}]$, which covers (λ_q, ξ_2), is contained in the union of $[\xi_{pq}, \lambda_0)$, $\mathcal{I}_r = [\lambda_0, \lambda_1]$, and $(\lambda_1, \lambda_{pq}]$.

To guess a (based on t^*, q and p), we distinguish between two possible situations. As before, our analysis is fully symmetric to that given in case (c2) of Section 6.5, so we only briefly review it.

(i) If t^* belongs to $(\lambda_q, \lambda_0) \subseteq [\xi_{pq}, \lambda_0)$ then $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ is among the last $O(1)$ special clockwise (p,q)-crossings to begin after t^*, because χ satisfies condition (S5). See Figure 88 (left).

Figure 88: Case (c2): Guessing a based on t^*, p and q. Left: If $t^* \in (\lambda_q, \lambda_0)$, then $(pa, q, \mathcal{I}_r = [\lambda_0, \lambda_1])$ is among the first $O(1)$ special clockwise (p,q)-crossings to begin after t^*. Right: If $t^* \in [\lambda_0, \xi_2)$, then $(pa, q, \mathcal{I}_r)$ is among the last $O(1)$ special clockwise (p,q)-crossings to begin before (or at) t^*.

(ii) If t^* belongs to the interval $[\lambda_0, \xi_2]$, which is contained in $\mathcal{I}_r \cup (\lambda_1, \lambda_{pq}]$, then we resort to a more subtle argument (which is fully symmetic to the one given in case (c1) of Section 6.5) to show that $(pa, q, \mathcal{I}_r)$ is among the last $O(1)$ special clockwise (p,q)-crossings to begin before t^*. See Figure 88 (right).

To recap, in each of the cases (c1) and (c2) we charge χ (via $(pa, q, \mathcal{I}_r)$) either to $\Omega(\ell^2)$ (8ℓ)-shallow co-circularities, or to an (8ℓ)-shallow collinearity, which occur in one of the arrangements $\mathcal{A}_{pq}, A_{wq}$ during the interval (λ_q, ξ_2). Furthermore, each (8ℓ)-shallow event is charged by at most $O(1)$ special quadruples. Hence, at most $O\left(\ell^2 N(n/\ell) + \ell n^2 \beta(n)\right)$ special quadruples χ fall into case (c).

Case (d). Assume that none of the preceding cases occurs. In particular, there is a subset A_{pw} of at most $6\ell + 3$ points (including a, r and u) whose removal restores the Delaunayhood of pw throughout the interval $[\xi_{pq}, \lambda_{pq}]$. Furthermore, a total of fewer than ℓ points of $P \setminus \{a, r, u\}$ appear in the cap C^-_{pq} during (λ_q, ξ_2), and a total of fewer than ℓ points of $P \setminus A_{pw}$ points appear in the cap C^-_{pw} during that interval.

The above assumptions imply the following symmetric variant of Lemma 6.4, whose proof is also fully symmetric to its predecessor (see Figure 89 (left)).

Lemma 6.15. *With the above assumptions, a total of at most $8\ell + 1$ points of P appear in the cap $C^-_{wq} = C[p, q, w] \cap L^-_{wq}$ during (λ_q, ξ_2).*

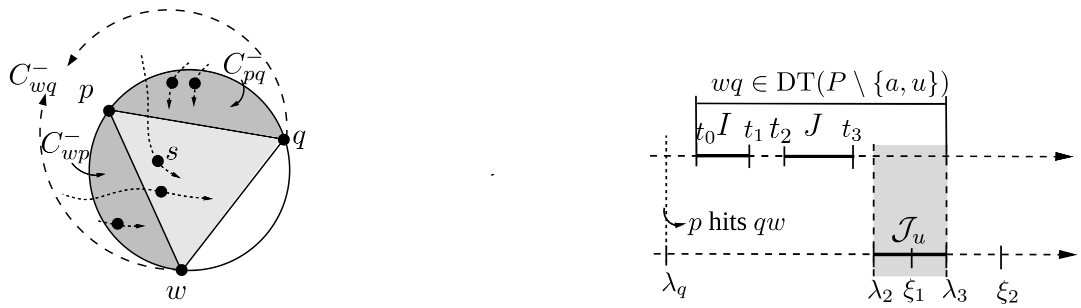

Figure 89: Left: Lemma 6.15. A total of at most $8\ell + 1$ points s of P appear in the cap $C^-_{wq} = B[p, q, w] \cap L^-_{wq}$ (consisting of all the shaded portions) during (λ_q, ξ_2). All of these points must enter C^-_{wq} after λ_q, and none of them can enter C^-_{wq} through wq, unless it is one of a, r, u. Right: The regular quadruple σ_2 of $(wa, q, \mathcal{J}_u)$ is composed of two (w, u)-crossings (wq, u, I), (wa, u, J), which end before the beginning time λ_2 of $\mathcal{J}_u$. By condition (Q8), the edge wq belongs to $\mathrm{DT}(P \setminus \{a, u\})$ throughout $[t_0, \lambda_3]$, implying that $\lambda_q < t_0$.

With the above preparations, we can finally describe the interplay between the special quadruple χ under consideration and the ordinary Delaunay quadruple $\sigma_2 = (w, q, a, u)$ in $\mathcal{F}$, which corresponds to the *second* special (a, q)-crossing $(wa, q, \mathcal{J}_u = [\lambda_2, \lambda_3])$ of χ. At the end of this section, we shall charge χ to the terminal quadruple $\varrho = (w, q, u, p)$, which is composed of the edge wq, and of the two points u and p that cross pq in opposite directions. As in the case of right special quadruples, we first try to enforce as many Delaunay crossings as possible among w, q, u, p, before charging χ to this terminal quadruple.

Recall that the quadruple $\sigma_2 = (w, q, a, u)$ belongs to the refined family $\mathcal{F}$, so it satisfies the eight properties (Q1)-(Q8). (Refer to Figure 89 (right).) Specifically, σ_2 is composed of two clockwise (w, u)-crossings $(wq, u, I = [t_0, t_1])$ and $(wa, u, J = [t_2, t_3])$, where I ends before the end t_3 of J, and J ends before the beginning λ_2 of $\mathcal{J}_u$. (In particular, $\mathcal{J}_u$ is disjoint from both of I, J.) Since σ_2 satisfies condition (Q8), the edge wq belongs to $\mathrm{DT}(P \setminus \{a, u\})$ throughout the interval $[I, \mathcal{J}_u] = \mathsf{conv}(I \cup \mathcal{J}_u) = [t_0, \lambda_3]$. Therefore (and since $\lambda_q < \lambda_2 < \lambda_3$), the point p can cross wq (at time λ_q, and from L^+_{wq} to L^-_{wq}) only before the beginning t_0 of I, and the entire crossing $(qw, p, \mathcal{H} = [\lambda_4, \lambda_5])$ occurs in $P \setminus \{a, r, u\}$ before I. We thus obtain the following important property of 3-restricted left special quadruples (see Figure 90 (left)):[44]

Proposition 6.16. *With the above assumptions, the first Delaunay crossing $(wq, u, I = [t_0, t_1])$ occurs entirely within $(\lambda_q, \lambda_2) \subset (\lambda_q, \xi_1) \subset (\lambda_q, \xi_2)$. In particular, p crosses wq at time λ_q (from L^+_{wq} to L^-_{wq}) before u does so in the opposite direction (during I, from L^-_{wq} to L^+_{wq}).*

Recall that p remains in L^-_{wq} throughout the interval (λ_q, ξ_2), which contains I; see Figure 90 (left). Note that the open cap $B[w, q, u] \cap L^-_{wq}$ contains no points of P at time t_0 (when the Delaunay crossing of wq by u begins). Hence, u lies at that moment within the cap $C^-_{wq} = B[p, q, w] \cap L^-_{wq}$; see Figure 90 (right). Since C^-_{wq} is empty right after time λ_q, the point u has to enter C^-_{wq} in the interval (λ_q, t_0).

[44]Though it is not necessary for our analysis, Proposition 6.16 holds for *all* 1-restricted left special quadruples $\chi = (a, p, w, q)$, with respective outer points r and u.

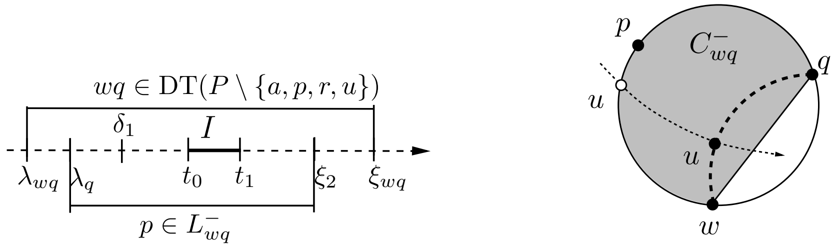

Figure 90: Left: Proposition 6.16: The interval $I = [t_0, t_1]$ (where wq undergoes a crossing by u) is fully contained in the interval (λ_q, ξ_2), during which p lies in L^-_{wq}. Right: The cap C^-_{wq} is empty right after time λ_q, so u must enter C^-_{wq} before the beginning t_0 of I. Unless u crosses qw in $(\lambda_{wq}, t_0) \supset (\lambda_q, t_0)$, u must enter C^-_{wq} at a blue-blue co-circularity of w, q, u, p with respect to wq, at some time $\delta_1 \in (\lambda_q, t_1)$.

Assume first that u hits wq during (λ_{wq}, t_0). (In particular, this includes the scenario where u enters C^-_{wq} during (λ_q, t_0) through the relative interior of wq.) Recall that wq is Delaunay in $P \setminus \{a, p, r, u\}$ throughout $[\lambda_{wq}, t_0] \subset [\lambda_{wq}, \xi_{wq}]$ (in addition to its Delaunayhood in P at times λ_{wq} and t_0). Hence, in the reduced set $P \setminus \{a, p, r\}$, the edge wq or, more, precisely, its reversely oriented copy qw, undergoes a Delaunay crossing by u during some sub-interval of $[\lambda_{wq}, t_0)$. Therefore, together with the crossing (wq, u, I), the triple w, q, u performs two single Delaunay crossings in $P \setminus \{a, p, r\}$. Combining Lemma 4.5 with the probabilistic argument of Clarkson and Shor, we obtain that the number of such triples q, w, u in P cannot exceed $O(n^2)$. By Proposition 6.1, the same quadratic bound must also hold for the overall number of such left special quadruples χ.

Case (d): The three co-circularities of w, q, u, p**.** Assume, then, that u does not cross wq in $[\lambda_{wq}, t_0)$. In particular, u enters $B[p, q, w]$ in (λ_q, t_0) through the boundary of $B[p, q, w]$, at a blue-blue co-circularity of w, q, u, p with respect to wq (as depicted in Figure 91 (right)). We claim that this is the second co-circularity of w, q, u, p, denoting its time by δ_1.

Indeed, by Lemma 4.4, another co-circularity of w, q, u, p occurs at some time $\delta_2 \in I = [t_0, t_1]$ (where wq undergoes a single Delaunay crossing by u), and is red-blue with respect to wq. Refer to Figure 91 (left). Furthermore, since u does not hit wq during $[\lambda_{wq}, \delta_1) \subset [\lambda_{wq}, t_0)$ (and wq belongs to $\mathrm{DT}(\{w, q, u, p\})$ at times λ_{wq} and δ_1), the edge qw undergoes a Delaunay crossing by p in the triangulation of $\{w, q, u, p\}$ too. This crossing occurs during some sub-interval of $[\lambda_{wq}, \delta_1)$ so, by Lemma 4.4, w, q, u, p are involved in another co-circularity at some time $\delta_0 \in [\lambda_{wq}, \delta_1)$; see Figure 91 (center).

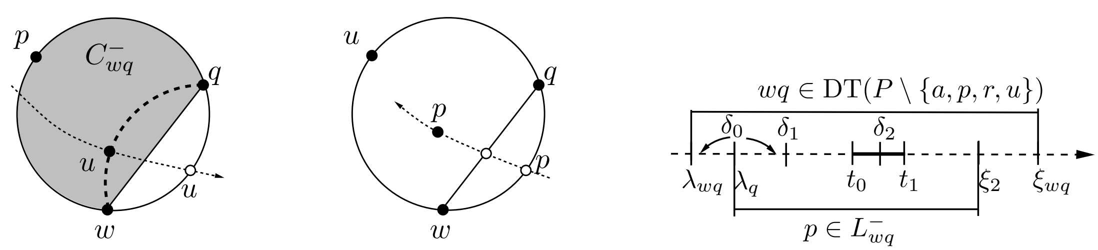

Figure 91: Left: The red-blue co-circularity of w, q, u, p with respect to wq, which must occur at some time $\delta_2 \in I$. Center: The points w, q, u, p are involved at some time $\delta_0 \in [\lambda_{wq}, \delta_1)$ in their first co-circularity, which is also red-blue with respect to wq. Right: A schematic summary of the motion of w, q, u, p (assuming that u does not cross wq during (λ_q, t_0)).

To conclude, the four points w, q, u, p are co-circular at times $\delta_0 \in [\lambda_{wq}, t_0)$, $\delta_1 \in (\delta_0, t_0)$, and $\delta_2 \in I = [t_0, t_1]$. (See Figure 91 (right) for a schematic summary.) Here the two extremal co-circularities, which occur at times δ_0 and δ_1, are red-blue with respect to the edges, and the middle co-circularity at time δ_1, is blue-blue with respect to wq (and occurs when u enters the cap C_{wq}^-). We emphasize that u remains in C_{wq}^- throughout (δ_1, t_0).

Furthermore, the order type of the third co-circularity (at time $\delta_2 \in I$) is completely determined by Proposition 6.16 and the fact that p lies in L_{wq}^- throughout (λ_q, ξ_2). Hence, this co-circularity occurs during the second portion of I (i.e., after u enters L_{wq}^+), when p leaves the cap $B[w, q, u] \cap L_{wq}^-$.

Notice that the caps $B[w, q, u] \cap L_{wq}^-$ and C_{wq}^- coincide at time $\delta_2 \in I \subset (\lambda_q, \xi_2)$. Therefore, Lemma 6.15, together with P-emptiness of $B[w, q, u] \cap L_{wq}^+$ during the second portion of I, imply that the co-circularity at time δ_2 is $(8\ell + 1)$-shallow.

Recall that (wq, u, I) is a clockwise (w, q)-crossing, and a counterclockwise (q, u)-crossing. Lemma 6.15 yields the following symmetric analogue of Lemma 6.6 (with somewhat simpler proof, due to Proposition 6.16).

Lemma 6.17. *With the above assumptions, at most $8\ell + 1$ clockwise (w, u)-crossings (wq', u, I'), and at most $8\ell + 1$ counterclockwise (q, u)-crossings $(w'q, u, I')$, can begin in the interval (δ_1, t_0).*

Proof. Let (wq', u, I') be a clockwise (w, u)-crossing that begins in (δ_1, t_0). By Lemma 4.4, the four points w, q, u, q' are co-circular at some moment $\zeta' \in I' \setminus I \subset (\delta_1, t_0)$, and this co-circularity is red-blue with respect to the edges wq', uq, and monochromatic with respect to wq. Furthermore, since p remains in C_{wq}^- throughout $(\delta_1, t_0) \subset (\lambda_q, \xi_2)$, the above co-circularity is, in fact, blue-blue with respect to wq (see Figure 92). Hence, both points u, q' lie at time ζ' inside the cap C_{wq}^-. Lemma 6.15 now implies that the overall number of such points q' (and, therefore, of their respective (w, u)-crossings (wq', u, I')) cannot exceed $8\ell + 1$.

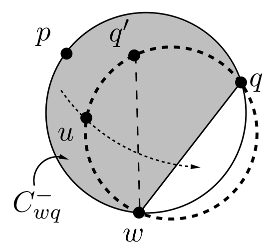

Figure 92: Lemma 6.17: Proving that at most $8\ell + 1$ clockwise (w, u)-crossings (wq', u, I') begin in (δ_1, t_0). For each of these crossings, the four points w, q, u, q' are involved in a blue-blue co-circularity with respect to wq at some time $\zeta \in I' \setminus I \subset (\delta_1, t_0)$, so their respective points q' enter C_{wq}^- during (λ_q, ξ_2).

A fully symmetric argument shows that at most $8\ell + 1$ counterclockwise (q, u)-crossings $(w'q, u, I')$ can begin in the interval (δ_1, t_0), because their respective points w' must appear in C_{wq}^- at some moment during $(\delta_1, t_0) \subset (\lambda_q, \xi_2)$. □

To proceed, we distinguish between two possible cases depicted in Figure 93.

Case (d1). The co-circularity at time δ_1 is red-blue with respect to the edge wu whose Delaunayhood is violated right before δ_1 by $p \in L_{wu}^-$ and $q \in L_{wu}^+$ (see Figure 93 (left)).

Note that the above violation of wu does not hold at time λ_q, when the segments pq and wu do not even intersect. Therefore, and since δ_1 is the *only* blue-blue co-circularity of w, q, u, p with respect to wq, applying (the time-reversed variant of) Lemma 3.1 in $\mathrm{DT}(\{w, q, u, p\})$ over the interval (λ_q, δ_1) shows that wu is hit in that interval by at least one of p or q (see Figure 94).

A very similar argument shows that the edge pq, whose Delaunayhood is violated right after time δ_1 by $u \in L_{pq}^-$ and $w \in L_{pq}^+$, is hit by u after δ_1 and before u enters L_{wq}^+ (during I). Indeed, let υ_{wq} denote the time in I when u hits wq. Note that the above violation of pq does not hold at time υ_{wq}. Therefore, another application of Lemma 3.1 in $\mathrm{DT}(\{p, q, w, u\})$ shows that the edge pq is hit during $(\delta_1, \upsilon_{wq})$ by at least one of the two points w or u. Recall, however, that $I \subset (\lambda_q, \xi_2)$ (by Proposition 6.16). Hence,

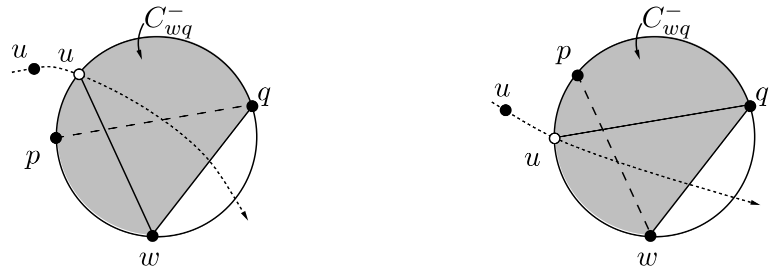

Figure 93: Left: Case (d1). The co-circularity at time δ_1 is red-blue with respect to the edges wu and pq. Right before time δ_1, the Delaunayhood of wu is violated by p and q. Right: Case (d2). The co-circularity at time δ_1 is red-blue with respect to the edges uq and wp. Right before time δ_1, the Delaunayhood of uq is violated by p and w.

both times $\delta_1 \in (\lambda_q, t_0)$ and $\upsilon_{wq} \in I = [t_0, t_1]$ belong to the interval (λ_q, ξ_2) (during which p lies in L^-_{wq}), ruling out the crossing of pq by w in $(\delta_1, \upsilon_{wq})$. Hence, it must be the case that pq is by u, as depicted in Figure 94.

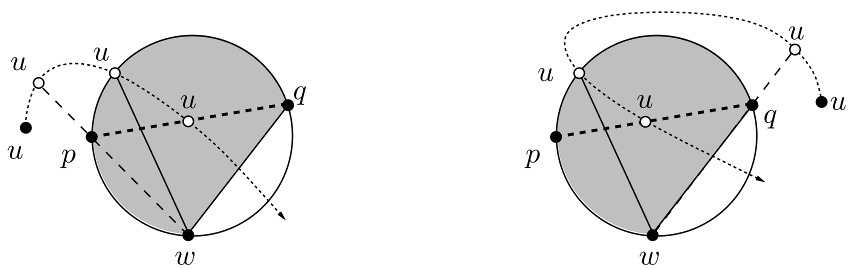

Figure 94: The two possible trajectories of u according to Lemma 6.7. The edge uw is hit in (t_q, δ_1) by p (left) or q (right). In both scenarios, u hits the edge pq after δ_1 and before the time $\upsilon_{wq} \in I$ when u hits wq.

To conclude, we have established the following lemma.

Lemma 6.18. *With the above notation, the following two claims hold in case (d1):*

(i) The edge pq is hit in $(\delta_1, \upsilon_{wq})$ by u, which crosses pq from L^-_{pq} to L^+_{pq}.

(ii) The edge wu is hit in (λ_q, δ_1) by at least one of the points p, q. Namely, either p crosses wu from L^+_{wu} to L^-_{wu}, or q crosses wu in the reverse direction. Moreover, the Delaunayhood of wu is violated by p and q right after the last such crossing and until δ_1.

Case (d1) – the crossing of pq by u. Refer to Figure 95. Recall that both λ_q and δ_1 belong to the interval $(\lambda_q, \xi_2) \subset (\xi_{pq}, \lambda_{pq})$ where, by condition (S5), pq belongs to $\mathrm{DT}(P \setminus \{a, w, r, u\})$ (in addition to its Delaunayhood in P at times ξ_{pq}, λ_{pq}).

By Lemma 6.18 (i), pq is hit by u in $(\lambda_q, \delta_1) \subset [\xi_{pq}, \lambda_{pq}]$. Therefore, and since pq is Delaunay at times ξ_{pq} and λ_{pq}, this edge (or its reversely oriented copy qp) undergoes a Delaunay crossing by u within a suitably reduced triangulation $\mathrm{DT}(P \setminus \{a, w, r\})$.

Case (d1)–enforcing the crossing of wu by p or q. If the edge wq is never Delaunay in P before time δ_1 then, by Lemma 6.17, (wq, u, I) is among the first $O(\ell)$ clockwise (w, u)-crossings (because wu is Delaunay during each of these crossings). Proposition 6.1 implies that this can occur for at most $O(\ell n^2)$ special quadruples χ. Therefore, we may assume that wu has appeared in $\mathrm{DT}(P)$ also before δ_1.

Let ξ_{wu} denote the last time in $(-\infty, \delta_1)$ when the edge wu belongs to $\mathrm{DT}(P)$; see Figure 96. Notice that the time when wu is hit by one of p, q, as prescribed by Lemma 6.18 (ii), must belong to the interval $[\xi_{wu}, \delta_1)$, which is contained in $[\xi_{wu}, t_0)$. To enforce the desired Delaunay crossing of wu, we apply Theorem 2.2 in $\mathcal{A}_{wu}$ over the interval (ξ_{wu}, t_0), with the third constant $h \gg \ell$.

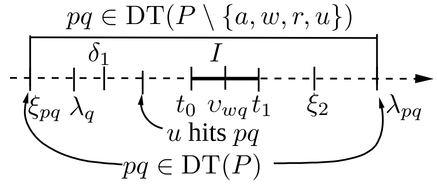

Figure 95: Case (d1)–obtaining a Delaunay crossing of pq by u. The edge pq is Delaunay at times ξ_{pq} and λ_{pq}, and almost Delaunay in (ξ_{pq}, λ_{pq}). Since u hits pq in $(\delta_1, \upsilon_{wq}) \subset (\xi_{pq}, \lambda_{pq})$, pq undergoes a Delaunay crossing by u in $P \setminus \{a, w, r\}$.

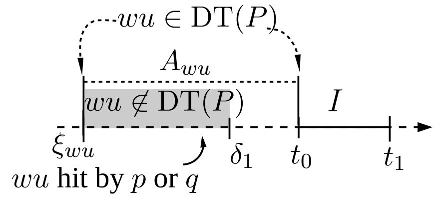

Figure 96: Case (d1)–enforcing a crossing of wu by at least one of the points p, q. The edge wu is Delaunay throughout $I = [t_0, t_1]$ and at time $\xi_{wu} < \delta_1$ (which is the last such time before δ_1). The Delaunayhood of wu is violated by p and q right before $\delta_1 \in (\xi_{wu}, t_0]$, so the promised crossing of pq, by at least one of p, q, must occur in $[\xi_{wu}, \delta_1)$.

If at least one of the Conditions (i), (ii) holds, we can charge χ, within $\mathcal{A}_{wu}$, either to an h-shallow collinearity or to $\Omega(h^2)$ h-shallow co-circularities. Lemma 6.6 ensures that each h-shallow event, that occurs in $\mathcal{A}_{wu}$ at some time $t^* \in (\xi_{wu}, t_0)$, is charged in this manner by at most $O(\ell)$ left special quadruples. Indeed, the corresponding points w and u are involved in the event, so we can guess them in $O(1)$ possible ways, and (wq, u, I) is among the first $8\ell + 2$ clockwise (w, u)-crossings to begin after time t^*. Therefore, the above charging accounts for at most $O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples χ.

We may assume, then, that Condition (iii) of Theorem 2.2 holds. That is, there is a subset A_{wu} of at most $3h$ points (perhaps including some of p, q, a, and r) whose removal restores the Delaunayhood of wu throughout the interval $[\xi_{wu}, t_0]$.

If wu is crossed during $[\xi_{wu}, t_0)$ by q (from L^-_{wu} to L^+_{wu}), then, together with (wq, u, I), the triple w, q, u performs two Delaunay crossings in $(P \setminus A_{wu}) \cup \{q\}$. A routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that P contains at most $O(hn^2)$ triples w, q, u of this kind. By Proposition 6.1, this also bounds the overall number of such left special quadruples χ.

To conclude, we may assume that the edge wu (or its reversely oriented copy uw) undergoes a Delaunay crossing by p in the smaller set $(P \setminus A_{wu}) \cup \{p\}$. In addition, we have shown that the edge pq (or its reversely oriented copy qp) undergoes a Delaunay crossing by u in $P \setminus \{a, r, w\}$. (Note that one, or both of these crossings can be a double Delaunay crossing.) Therefore, together with the crossings (wq, u, I) and $(qw, p, \mathcal{H})$, each of the four possible sub-triples of w, q, u, p performs a Delaunay crossing within a suitably refined triangulation.

Finally, recall that the four points w, q, u, p are involved at some time $\delta_2 \in I \subset (\lambda_q, \xi_2)$ in their third (and last) co-circularity, which is red-blue with respect to the edges wq and pu. Moreover, this co-circularity is $(8\ell + 1)$-shallow in P (because of Lemma 6.15), and the Delaunayhood of up is violated right after it by $q \in L^-_{pu}$ and $p \in L^+_{pu}$. Let A_{δ_2} be the set of at most $8\ell + 1$ points that lie at time δ_2 within the circumdisc of w, q, u, p.

Case (d1): Charging terminal quadruples. We consider a subset R of $\lceil n/h \rceil$ points chosen at random from P. Notice that the following two events occur simultaneously, with probability at least $\Omega(1/h^4)$:

(1) R contains the four points w, q, u, p, and (2) none of the points of $(A_{wu} \cup A_{\delta_2} \cup \{a, r\}) \setminus \{w, q, u, p\}$ belong to R.

In the case of success, we charge χ to the quadruple $\varrho = (w, q, u, p)$, which satisfies the following two conditions with respect to R (see Figure 97):

(C1) The edge pq (or qp) undergoes a Delaunay crossing by u in $R \setminus \{w\}$. Similarly, the edge wu (or uw) undergoes a Delaunay crossing by p in $R \setminus \{q\}$.

(C2) The four points of ϱ are involved in a Delaunay co-circularity, right after which the Delaunayhood of pu is violated by $q \in L^-_{pu}$ and $w \in L^+_{pu}$. Furthermore, this is the last co-circularity of w, q, u, p.

Note that χ is uniquely determined by ϱ.

Definition. Let P be a finite set of moving points in $\mathbb{R}^2$. We say that a quadruple $\varrho = (w, q, u, p)$ in P is *terminal of type C* if it satisfies the above conditions (C1) and (C2), with R replaced by P.

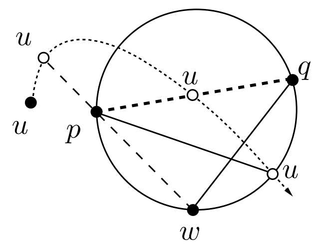

Figure 97: A possible trajectory of u if $\varrho = (w, q, u, p)$ is a terminal quadruple of type C. The points of ϱ are involved in an extremal (last) Delaunay co-circularity, right after which the Delaunayhood of pu is violated by $q \in L^-_{pu}$ and $w \in L^+_{pu}$.

Let $T^C(m)$ denote the maximum possible number of terminal quadruples of type C that can arise in an underlying set of m moving points. Then the overall number of 3-restricted left special quadruples that fall into case (d1) is at most

$$O\left(h^4 T^C(n/h) + \ell h^2 N(n/h) + \ell h n^2 \beta(n)\right).$$

In Section 7.3 we will use the corresponding extremal Delaunay co-circularity of w, q, u, p of each terminal quadruple ϱ to enforce a Delaunay crossing of pu by at least one of the remaining two points w, q of ϱ. Together with the Delaunay crossings in condition (C1), at least one of the triples p, u, w or p, u, q will perform two (single) Delaunay crossings. Therefore, our analysis will again bottom up via Lemma 4.5.

Remark. Notice that in condition (C1) we omit the crossings (wq, u, I) and $(qw, p, \mathcal{H})$ which gave rise to the terminal quadruple $\varrho = (w, q, u, p)$, after having used them to enforce the crossings of p, u, w and p, u, q.

Case (d2). The co-circularity at time δ_1 is red-blue with respect to the edge uq whose Delaunayhood is violated right before δ_1 by $p \in L^-_{uq}$ and $w \in L^+_{uq}$ (see Figure 93 (right)).

Using υ_{wq} as before to denote the unique time in $I = [t_0, t_1]$ when u hits wq, we have the following symmetric variant of Lemma 6.18, which can be established by switching the roles of w and q in the argument that implied Lemma 6.18.

Lemma 6.19. *With the above notation, the following two properties hold in case (d2):*

(i) The edge wp is hit in $(\delta_1, \upsilon_{wq})$ by u, which crosses wp from L^-_{wp} to L^+_{wp}.

(ii) The edge uq is hit in (λ_q, δ_1) by at least one of the points p, w. Namely, either p crosses uq from L^+_{uq} to L^-_{uq}, or w crosses uq in the reverse direction. Moreover, the Delaunayhood of uq is violated by p and w right after the last such crossing and until δ_1.

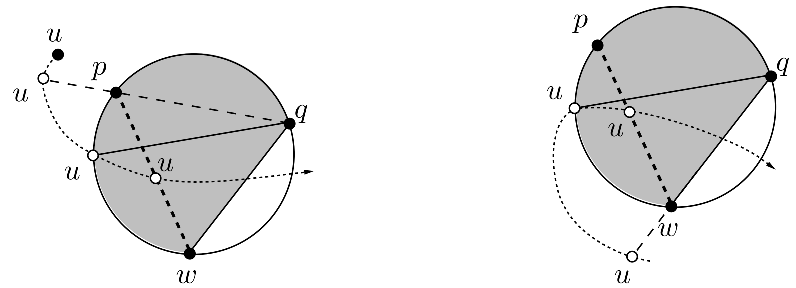

Figure 98: The two possible trajectories of u according to Lemma 6.19. The edge uq is hit in (t_q, δ_1) by p (left) or w (right). In both scenarios, u hits the edge wp after δ_1 and before the time $\upsilon_{wq} \in I$ when u hits wq.

We next amplify the collinearities in Lemma 6.19 into full-fledged Delaunay crossings. We again emphasize that (wq, u, I) is a clockwise (w, u)-crossing, and a counterclockwise (u, q)-crossing, so the role of uq in the present case (d2) is fully symmetric to the role of wu in case (d1). In particular, the crossing of uq (or of its reversed copy qu) by p or w will be enforced using essentially the same argument as was used in case (d1) to enforce the Delaunay crossing of wu by p or q.

In contrast, the properties of pq (in case (d1)) and wp (in case (d2)) are *not* symmetric. Indeed, the edge pq (which was crossed by u in case (d1)) is almost Delaunay throughout the interval $[\xi_{pq}, \lambda_{pq}]$ (which covers $(\lambda_q, \xi_2) \supset (\delta_1, \upsilon_{wq})$, where wp is hit by u or q), and Delaunay at both times ξ_{pq}, λ_{pq}. However, the edge wp (which is crossed by u in the present case (d2)) becomes Delaunay in $[\xi_{pq}, \lambda_{pq}]$ only after removal of a subset A_{pw} of at most $6\ell + 3$ points (including u), which is not enough to obtain a Delaunay crossing of wp by u.

Case (d2): Enforcing a Delaunay crossing of wp by u. We emphasize that the third co-circularity of w, q, u, p is $(8\ell + 1)$-shallow and occurs at some time δ_2 during the second portion of I, starting right after the unique time υ_{wq} in I when u hits wq. Recall also that I begins after δ_1 and is contained in the nested intervals (λ_q, ξ_2) and (ξ_{pq}, λ_{pq}) (where the Delaunayhood of pw can be restored by removing the above set A_{pw} of at most $6\ell + 3$ points).

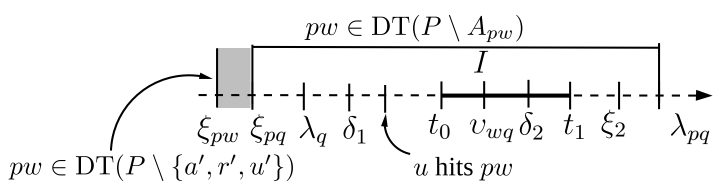

Figure 99: Case (d2)–enforcing the crossing of wp by u. The edge pw is Delaunay in $P \setminus A_{pw}$ throughout the interval $[\xi_{pq}, \lambda_{pq}]$, which contains δ_1 and I (inclduing υ_{wq} and δ_2). We first obtain a time $\xi_{pw} \le \xi_{pq}$ when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, so that none of the obstruction points a', r', u' is equal to u. Note that u hits pw in the interval (ξ_{pw}, δ_2).

Notice that the Delaunayhood of pw at time δ_2 can be enforced by removing the subset A_{δ_2} of at most $8\ell + 1$ points that lie at time δ_2 in the interior of the circumdisc of w, q, u, p. Since A_{δ_2} does not include u, its removal does not destroy the crossing triple w, p, u.

We first obtain a time $\xi_{pw} \le \xi_{pq}$ when the edge pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, for some $a', r', u' \in P \setminus \{w, p, u\}$. In particular, (ξ_{pw}, δ_2) contains the above time in $(\delta_1, \upsilon_{wq})$ when u crosses wp from L^-_{wp} to L^+_{wp}. We then use Theorem 2.2 to extend the almost-Delaunayhood of pw to $[\xi_{pw}, \xi_{pq})$, so as to cover the entire $[\xi_{pw}, \lambda_{pq}]$. As a result, wp will undergo a Delaunay crossing by u during some sub-interval of $[\xi_{pw}, \delta_2]$ (and in an appropriately reduced subset of P).

To obtain the above time $\xi_{pw} \le \xi_{pq}$, we return to the subfamily $\mathcal{G}^L_{pw}$ of all 3-restricted left special quadruples $\chi' = (a', p, w, q')$ (each coming with respective outer points r' and u') whose two middle points are equal to p and w, respectively. Recall that each quadruple in $\mathcal{G}^L_{pw}$ is uniquely determined by its respective point q'. In addition, we can assume that all quadruples in $\mathcal{G}^L_{pw}$ fall into case (d2) of the present analysis (because the remaining quadruples in $\mathcal{G}^L_{pw}$ are handled using previous charging arguments). In particular, $\mathcal{G}^L_{pw}$ contains the quadruple $\chi = (a, p, w, q)$ under consideration.

Our analysis relies on the following uniqueness property:

Lemma 6.20. *With the above assumptions, the family $\mathcal{G}^L_{pw}$ contains at most $3\ell + 1$ other 3-restricted left special quadruples $\chi' = (a', p, w, q')$, with respective outer points r' and u', that fall into case (d2) and satisfy $u' = u$.*

In other words, any triple w, p, u can be shared by at most $3\ell + 2$ 3-restricted left special quadruples χ under consideration.

Proof. Notice that, for each terminal quadruple $\chi' = (a', p, w, q') \in \mathcal{G}^L_{pw}$ under consideration, with respective outer points r' and $u' = u$, the four points w, u, p, q' are involved in their third co-circularity at some time δ'_2 during the respective regular crossing of wq' by u. Right after time δ'_2, the Delaunayhood of pu is violated by $q' \in L^-_{pu}$ and $w \in L^+_{pu}$. Clearly, the lemma will follow if we show that δ_2 is among the first $8\ell + 2$ such times δ'_2 to occur after u crosses wp from L^-_{wp} to L^+_{wp} (as prescribed in Lemma 6.19 (i)). (See Figure 100.)

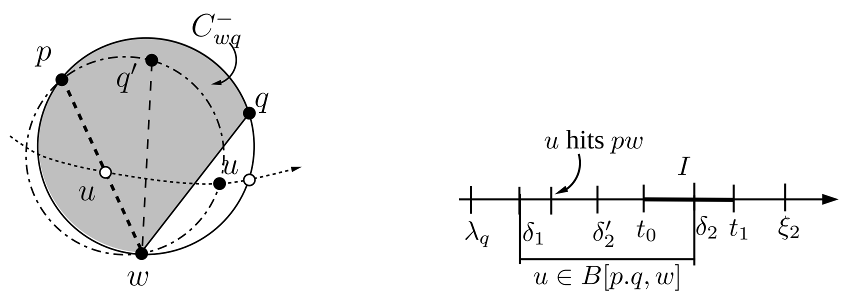

Figure 100: Proof of Lemma 6.20. We fix a quadruple $\chi' = (a', p, w, q') \in \mathcal{G}^L_{pw}$ whose second outer point u' equal to u, so that the third co-circularity of w, p, u, q' occurs at some time δ'_2 after u crosses wp (from L^-_{wp} to L^+_{wp}) and before δ_2. We claim that q' lies at time δ'_2 in the cap C^-_{wq}. The two hollow circles in the left figure represent the location of u when it hits wp, and at time $\delta_2 > \delta'_2$ (when u leaves $B[p, q, w]$).

To establish the last claim, let $\chi' = (a', p, w, q')$ be a 3-restricted left quadruple, with respective outer points r' and $u' = u$, and such that the corresponding third co-circularity of w, u, p, q' occurs at some time δ'_2 after w enters L^+_{wp} through wp and before δ_2. We claim that q' lies at time δ_2 within the cap $C^-_{wq} = B[p, q, w] \cap L^-_{wq}$ (as depicted in Figure 100 (left)), so, by Lemma 6.15, the overall number of such points q' (and, therefore, also of their respective quadruples χ') cannot exceed $8\ell + 1$.

Indeed, recall that the motion of $B[p, q, w]$ is continuous in the interval (λ_q, ξ_2), which contains $\delta_1 \in (\lambda_q, t_0)$ and $\delta_2 \in I = [t_0, t_1]$. Therefore, and since δ_2 is (the time of) the last co-circularity of w, q, u, p, the point u must remain in $B[p, w, q]$ after time δ_1, when u enters that disk, and until the time δ_2, when u leaves $B[p, q, w]$. Therefore, both u and q' lie in $B[p, q, w] \cap L^+_{wp}$ at time δ'_2, when we encounter a red-red co-circularity of p, w, u, q' with respect to wp. It hence suffices to show that q' lies in L^-_{wq} at time δ_2.

Assume for a contradiction that q' lies at time δ'_2 in the opposite cap $C^+_{wq} = B[p, q, w] \cap L^+_{wq}$. This readily implies that the four edges wp, wq, wq', and wu, appear around w in this clockwise order at time δ_2; see Figure 101 (left). In particular, the point u too lies at time δ'_2 in C^+_{wq}, so δ'_2 belongs to the second portion of I (which starts at time υ_{wq}, when u hits wq); see Figure 101 (right). Notice that the Delaunayhood of wq is violated in $(\upsilon_{wq}, \delta_2)$ by $p \in L^-_{wq}$ and $u \in L^+_{wq}$, so p must lie in the cap

$B[w,q,u] \cap L^-_{wu}$ at time $\delta_2' \in (v_{wq}, \delta_2)$. However, since the co-circularity of w, u, p, q' at time δ_2' is blue-blue with respect to wu, the above cap $B[w,q,u] \cap L^-_{wu}$ must then contain also the point q'. In particular, q' lies at time δ_2' within the disk $B[w,q,u]$, contrary to the P-emptiness of $B[w,q,u] \cap L^+_{wq}$ during the second portion of I.

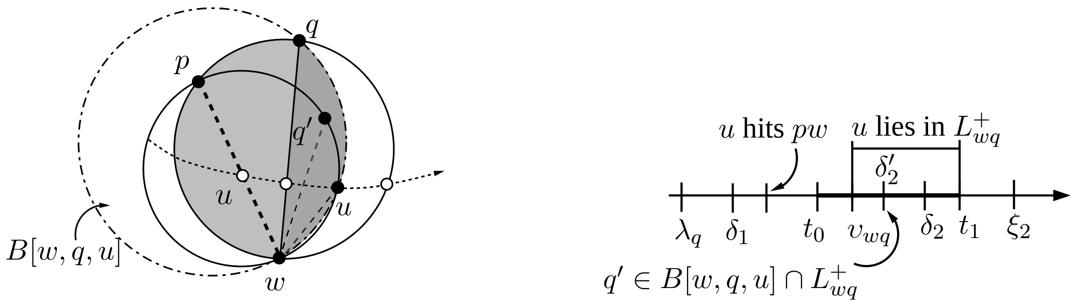

Figure 101: If q' lies at time δ_2' in the opposite cap C^+_{wq}, then this co-circularity occurs during the second portion $(v_{wq}, t_1]$ of I. In this hypothetic case, q' lies at time δ_2 within the disk $B[w,q,u]$, contrary to the P-emptiness of $B[w,q,u] \cap L^+_{wq}$ during $(v_{wq}, t_1]$.

To conclude, the above contradiction implies that q' lies at time δ_2' in the cap $B[p,q,w] \cap L^-_{wq}$. Hence, Lemma 6.15 implies the overall number of such points q' cannot exceed $8\ell + 1$. Therefore, the family $\mathcal{G}^L_{pw}$ contains at most $8\ell + 1$ 3-restricted left special quadruples $\chi' = (a', p, w, q')$, with respective outer points r' and $u' = u$, that fall into case (d2), and whose respective third co-circularities δ_2' occur after u crosses wp from L^-_{wp} to L^+_{wp} and before δ_2. In other words, δ_2 is among the first $8\ell + 2$ such times δ_2' to occur after u crosses wp as above. $\square$

Lemma 6.20 implies, through the standard pigeonhole argument, that at least some constant positive fraction of all 3-restricted left special quadruples $\chi = (a, p, w, q)$ under consideration (with respective outer points r and u) satisfy the following condition:

(PHL2) *There exist at most $O(\ell)$ quadruples $\chi' \in \mathcal{G}^L_{pw}$, with respective outer points r' and u', so that $u \in \{a', r', u'\}$.*

We may assume, with no loss of generality, that (PHL2) holds for χ under consideration. With these preparations, we can proceed to the main argument in $\mathcal{A}_{pw}$. Recall that each such quadruple $\chi' \in \mathcal{G}^L_{pw}$ is uniquely determined by the respective point q', and is accompanied by a counterclockwise (w,p)-crossing $(q'w, p, \mathcal{H}_{\chi'})$ which occurs in the reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$.

Figure 102: Left: If there exists no quadruple χ' (with respective outer points r' and u') in $\mathcal{G}^L_{pw}$ that satisfies $a', r', u' \neq u$, and whose respective (w,p)-crossing $(q'w, p, \mathcal{H}_{\chi'})$ begins before ξ_{pq}, then $\mathcal{G}^L_{pw}$ contains a total of at most $O(\ell)$ quadruples χ' whose respective (w,p)-crossings $(q'w, p, \mathcal{H}_{\chi'})$ begin before the starting time λ_4 of $\mathcal{H} = \mathcal{H}_\chi$. Right: Otherwise, there is a time $\xi_{pw} \leq \xi_{pq}$ which is the last such time when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, for $a', r', u' \neq u$.

Refer to Figure 102. Assume first that there is no quadruple $\chi = (a', p, w, q') \in \mathcal{G}^L_{pw}$ (with respective outer points r' and u') such that $a', r', u' \neq u$, and whose respective (w,p)-crossing $(q'w, p, \mathcal{H}_{\chi'})$ begins before ξ_{pq}. (See Figure 102 (left).) Since case (a) has been ruled out, $\mathcal{G}^L_{pw}$ contains at most k special quadruples whose respective (w,p)-crossings begin in $[\xi_{pq}, \lambda_4)$. Thus, G^L_{pw} contains a total of at most $O(k + \ell)$ quadruples whose respective (w,p)-crossing $(q'w, p, \mathcal{H}_{\chi'})$ begin before the starting λ_4 of $\mathcal{H} = \mathcal{H}_\chi$ (including the at most $O(\ell)$ such (w,p)-crossings that begin before ξ_{pq} and have one of their respective obstruction points a', r', u' equal to u). We charge χ to the edge pw, noting that the above scenario can occur for at most $O(\ell n^2)$ 3-restricted left special quadruples under consideration.

We thus can assume, in what follows, that there is at least one quadruple $\chi' = (a', p, w, q')$, with respective outer points r' and u', that satisfies $a', r', u' \neq u$, and whose respective (w,p)-crossing $(q'w, p, \mathcal{H}')$ in $P \setminus \{a', r', u'\}$ begins before (or at) ξ_{pq}. (See Figure 102 (right).) In particular, Lemma 4.1 implies that there is a time before (or at) ξ_{pq} when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{a', r', u'\})$, for some three points a', r', u' distinct from u. We choose ξ_{pw} as the *last* such time in $(\infty, \xi_{pq}]$.

Notice that the above choice of ξ_{pw} guarantees that there exist at most $O(\ell)$ quadruples $\chi' \in \mathcal{G}^L_{pw}$ whose respective (w,p)-crossings begin in $[\xi_{pw}, \lambda_4)$. In what follows, we will use a', r', u' to denote some three fixed points whose removal restores the Delaunayhood of pw at time ξ_{pw}.

We next apply Theorem 2.2 in $\mathcal{A}_{pw}$ over the interval (ξ_{pw}, λ_4). This is done with respect to the reduced set $P \setminus \{a', r', u'\}$ (which ensures the Delaunayhood of pw at the endpoint ξ_{pw}), and with the third constant $h \gg \ell$.

In cases (i), (ii) of Theorem 2.2 we charge χ within the reduced arrangement $\mathcal{A}_{pw}$ either to $\Omega(h^2)$ h-shallow co-circularities, or to an h-shallow collinearity. Notice that each of the charged events is $(h+3)$-shallow with respect to the original set P, and is charged by at most $O(\ell)$ left special quadruples χ. (The latter holds because the respective (w,p)-crossing $(qw, p, \mathcal{H} = [\lambda_4, \lambda_5])$ of χ is among the first $O(\ell)$ such (p,w)-crossings to begin after the time of the event.) Therefore, the above charging accounts for at most $O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples χ.

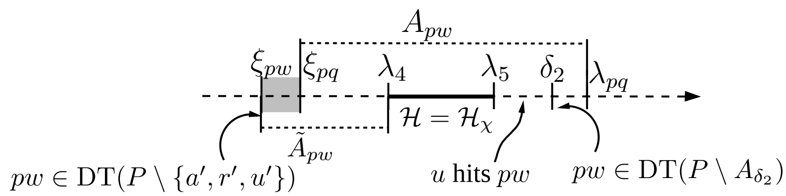

Figure 103: In case (iii) of Theorem 2.2 we end up with a subset $\tilde{A}_{pw}$ of at most $3h+3$ points (including a', r', u') whose removal restores the Delaunayhood of pw throughout $[\xi_{pw}, \lambda_4]$. In addition, pw is Delaunay in $P \setminus A_{pw}$ throughout the interval $[\xi_{pq}, \lambda_{pq}]$ (which contains $\mathcal{H} = [\lambda_4, \lambda_5]$, δ_2, and the time before δ_2 when u crosses wp from L^-_{wp} to L^+_{wp}), and it is Delaunay in $P \setminus A_{\delta_2}$ at time δ_2. Hence, if we omit the $O(h)$ points of $(\tilde{A}_{pw} \cup A_{pw} \cup A_{\delta_2}) \setminus \{u\}$, the edge wp (or pw) undergoes a Delaunay crossing by u.

Finally, in case (iii) of Theorem 2.2 we end up with a subset $\tilde{A}_{pw}$ of at most $3h+3$ points (including the three points a', r', u' which were put aside) whose removal restores the Delaunayhood of pw throughout $[\xi_{pw}, \lambda_4]$; see Figure 103. In particular, pw is Delaunay in $P \setminus (A_{pw} \cup \tilde{A}_{pw})$ throughout the interval $[\xi_{pq}, \lambda_{pq}] = [\xi_{pw}, \lambda_4) \cup [\xi_{pq}, \lambda_{pq}]$, which contains δ_2 and the above time in $(\delta_1, \upsilon_{wq}) \subset (\delta_1, \delta_2)$ when u crosses wp from L^-_{wp} to L^+_{wp}. Furthermore, recall that the co-circularity of p, q, w, u at time δ_2 is a Delaunay co-circularity in $P \setminus A_{\delta_2}$, where $A_{\delta_2} \subset P$ is a subset of cardinality at most $8\ell+1$. In particular, pw is Delaunay in $P \setminus A_{\delta_2}$ at time δ_2. Hence, in the even more reduced set $(P \setminus (A_{\delta_2} \cup A_{pw} \cup \tilde{A}_{pw})) \cup \{u\}$, the edge wp (or its reversely oriented copy pw) undergoes a Delaunay crossing by u during some subinterval of $[\xi_{pw}, \delta_2)$. (Specifically, the Delaunayhood of wp at times ξ_{pw} and δ_2 is guaranteed by removal of $a', r', u' \in \tilde{A}_{pw} \setminus \{u\}$ and $A_{\delta_2} \subset P \setminus \{p, q, w, u\}$.)

Case (d2): enforcing the crossing of qu by p or w. If the edge uq is never Delaunay in P before time δ_1, Lemma 6.17 implies that (wq, u, I) is among the first $O(\ell)$ counterclockwise (q,u)-crossings (because uq is Delaunay during each of these crossings). Proposition 6.1 implies that this can occur for at most $O(\ell n^2)$ special quadruples χ. Therefore, we may assume that uq appears in $\mathrm{DT}(P)$ also before δ_1.

Let ξ_{uq} denote the last time before δ_1 when the edge uq belongs to $\mathrm{DT}(P)$; see Figure 104. Notice that the time when uq is hit by one of p, w, as prescribed by Lemma 6.18 (ii), must belong to the interval

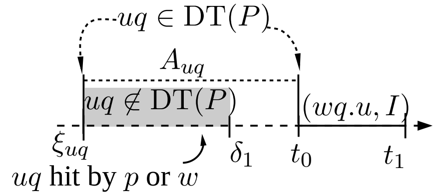

Figure 104: Case (d2)–enforcing a Delaunay crossing of qu by at least one of the points p, w. The edge qu is Delaunay throughout $I = [t_0, t_1]$ and at time $\xi_{uq} < \delta_1$ (which is the last such time before δ_1). The Delaunayhood of uq is violated by p and w right before $\delta_1 \in (\xi_{uq}, t_0)$, so it is hit in $[\xi_{uq}, \delta_1)$ by at least one of p, w.

$[\xi_{uq}, \delta_1)$, which is contained in $[\xi_{uq}, t_0)$. To enforce the desired Delaunay crossing of qu, we apply Theorem 2.2 in $\mathcal{A}_{qu}$ over the interval (ξ_{qu}, t_0), with the third constant $h \gg \ell$.

If at least one of the Conditions (i), (ii) holds, we can charge χ, within $\mathcal{A}_{uq}$, either to an h-shallow collinearity or to $\Omega(h^2)$ h-shallow co-circularities. Lemma 6.6 ensures that each h-shallow event, that occurs in $\mathcal{A}_{uq}$ at some time $t^* \in (\xi_{uq}, t_0)$, is charged in this manner by at most $O(\ell)$ left special quadruples. Indeed, the corresponding points u and q are involved in the event, so we can guess them in $O(1)$ possible ways, and (wq, u, I) is among the first $8\ell + 2$ (regular) counterclockwise (q, u)-crossings to begin after time t^*. Therefore, and since χ is uniquely determined by (wq, u, I) (see Proposition 6.1), the above charging accounts for at most $O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples χ.

We may assume, then, that Condition (iii) of Theorem 2.2 holds. That is, there is a subset A_{uq} of at most $3h$ points (perhaps including some of p, q, a, and r) whose removal restores the Delaunayhood of uq throughout the interval $[\xi_{uq}, t_0]$.

If uq is crossed during $[\xi_{uq}, t_0)$ by w (from L^-_{uq} to L^+_{uq}), then the triple w, q, u performs two Delaunay crossings in $(P \setminus A_{uq}) \cup \{w\}$. A routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that P contains at most $O(hn^2)$ triples w, q, u of this kind. By Proposition 6.1, this also bounds the overall number of such left special quadruples χ.

Case (d2): Converging. To recap, after excluding $O\left(\ell h^2 N(n/h) + \ell h n^2 \beta(n)\right)$ special quadruples χ, we may assume that the edge uq is hit in (ξ_{uq}, t_0) by p, so it (or its reversely oriented copy qu) undergoes a Delaunay crossing by p in the smaller set $(P \setminus A_{uq}) \cup \{p\}$.

In addition, the four points w, q, u, p are involved at some time $\delta_2 \in I$ in their third (and last) co-circularity, which is red-blue with respect to the edges wq and pu. Moreover, this co-circularity is $(8\ell+1)$-shallow in P, and the Delaunayhood of up is violated right after it by $q \in L^-_{pu}$ and $p \in L^+_{pu}$. As before, we use A_{δ_2} to denote the set of at most $8\ell + 1$ points that lie at time δ_2 within the circumdisc of w, q, u, p.

Finally, there exist sets A_{pw} and $\tilde{A}_{pw}$ that contain at most $O(\ell+h) = O(h)$ points in total, so that wp (or its reversely oriented copy pw) undergoes a Delaunay crossing by u in $(P \setminus (A_{\delta_2} \cup A_{pw} \cup \tilde{A}_{pw})) \cup \{u\}$. (Note that one, or both of these crossings can be a double Delaunay crossing.)

Case (d2): Charging terminal quadruples. We consider a subset R of $\lceil n/h \rceil$ points chosen at random from P. Notice that the following two events occur simultaneously, with probability at least $\Omega(1/h^4)$: (1) R contains the four points w, q, u, p, and (2) none of the points of $(A_{uq} \cup A_{pw} \cup \tilde{A}_{pw} \cup A_{\delta_2}) \setminus \{w, q, u, p\}$ belong to R.

In the case of success, we charge χ to the quadruple $\varrho = (w, q, u, p)$, which satisfies the following two conditions with respect to R (noting that χ is uniquely determined by ϱ); see Figure 105:

(D1) The edge wp (or wp) undergoes a Delaunay crossing by u in $R \setminus \{q\}$. Similarly, the edge qu (or uq) undergoes a Delaunay crossing by p in $R \setminus \{q\}$.

(D2) The four points of ϱ are involved in a Delaunay co-circularity, right after which the Delaunayhood of pu is violated by $q \in L_{pu}^-$ and $w \in L_{pu}^+$. Furthermore, this is the last co-circularity of w, q, u, p.

Definition. Let P be a finite set of moving points in $\mathbb{R}^2$. We say that a quadruple $\varrho = (w, q, u, p)$ in P is *terminal of type D* if it satisfies the above conditions (D1) and (D2), with R replaced by P.

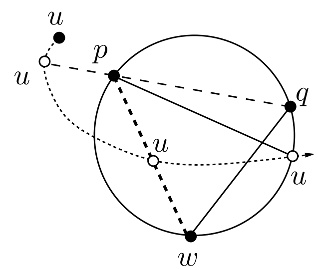

Figure 105: A possible trajectory of u if $\varrho = (w, q, u, p)$ is a terminal quadruple of type D. The points of ϱ are involved in an extremal (last) Delaunay co-circularity, right after which the Delaunayhood of pu is violated by $q \in L_{pu}^-$ and $w \in L_{pu}^+$.

Let $T^D(m)$ denote the maximum possible number of terminal quadruples of type D that can arise in an underlying set of m moving points. Then the overall number of 3-restricted left special quadruples that fall into case (d1) is at most

$$O\left(h^4 T^D(n/h) + \ell h^2 N(n/h) + \ell h n^2 \beta(n)\right).$$

In Section 7.3 we will use the corresponding extremal Delaunay co-circularity of w, q, u, p of each terminal quadruple ϱ to enforce a Delaunay crossing of pu by at least one of the remaining two points w, q of ϱ. Together with the Delaunay crossings in condition (D1), at least one of the triples p, u, w or p, u, q will perform two (single) Delaunay crossings. Therefore, our analysis will again bottom up via Lemma 4.5.

3-**restricted left special quadruples–wrap up.** Putting together the previously established bounds on the maximum possible numbers of 3-restricted left special quadruples that fall into cases (a), (b), (c), (d1) and (d2) yields the following recurrence:

$$\Phi_3^L(n) = O\left(h^4 T^C(n/h) + h^4 T^D(n/h) + \ell h^2 N(n/h) + k\ell^2 N(n/\ell) + k^2 N(n/k) + \ell h n^2 \beta(n)\right). \tag{11}$$

7 The number of terminal quadruples

In this section we obtain "quadratic" recurrences for the maximum numbers $T^A(n), T^B(n), T^C(n)$, and $T^D(n)$, of terminal quadruples of the respective types A, B, C, and D, which arise at the last stage of the analysis in Section 6. Each of these four quantities is expressed only in terms of the maximum number of Delaunay co-circularities in smaller-size sets, plus a nearly quadratic additive term. In other words, our analysis bottoms out. Combining these four new recurrences with the ones, obtained in Sections 3, 5, and 6, we finally get a complete system of "quadratic" recurrences, whose solution is $N(n) = O(n^{2+\varepsilon})$, for any $\varepsilon > 0$. This completes the proof of Theorem 2.1.

7.1 Terminal quadruples of type A

In this section we finally express the maximum possible cardinality $T^A(n)$ of a family Σ^A of terminal quadruples of type A (where each quadruple in Σ^A is uniquely determined by each of its four sub-triples) in terms of more elementary quanitities that were introduced in Section 2.

To do so, we fix the underlying set P of n moving points, a family Σ^A as above, and a terminal quadruple $\varrho = (p, q, r, w)$ in Σ_A. We emphasize that ϱ, as well as any other quadruple $(p, q, r, w) \in \Sigma^A$, is uniquely determined by each of its four sub-triples $(p, q, r), (p, q, w), (p, r, w), (q, r, w)$.

Recall that the four points of ϱ perform *four* Delaunay crossings, namely the crossing of pq by r, the crossing of qp by w, the crossing of rp by w, and the crossing of wq by r. Here only the first crossing, $(pq, r, I = [t_0, t_1])$, is defined with respect to the compelete point set P. Each of the remaining three crossings of ϱ occurs within a reduced point set, which is obtained by omitting from P the fourth point of ϱ (not directly involved in the crossing).

In this section, we shall enforce on the points of ϱ an additional *fifth* crossing, namely the crossing of rw (or its reversely oriented copy wr) by one of p, q. As a result, at least one of the triples p, r, w or q, r, w will perform two Delaunay crossings (within an appropriately reduced triangulation). We thus shall charge ϱ to that triple and complete our analysis by invoking Lemma 4.5.

Topological setup. Refer to Figure 106. By condition (A1), the edge pq is crossed by r (during $I = [t_0, t_1]$, as part of the corresponding Delaunay crossing) and w (at some later time $\lambda_q > t_1$), in opposite directions. Furthermore, pq re-enters $\mathrm{DT}(P)$ at some later time λ_{pq} after λ_q, and it belongs to $\mathrm{DT}(P \setminus \{r, w\})$ throughout $[t_0, \lambda_{pq}]$.

By condition (A2), the four points p, q, r and w are co-circular at some times $\delta_0 \in I$, $\delta_1 \in (t_1, \lambda_{pq}]$ and $\delta_2 \in (\delta_1, \lambda_{pq}]$, where the two extremal co-circularities (at times δ_0 and δ_2) are red-blue with respect to pq, and the middle co-circularity (at time δ_1) is red-red with respect to pq (and red-blue with respect to pr).

As a matter of fact, δ_2 arises as part of a single Delaunay of qp by w, which occurs in the triangulation $\mathrm{DT}(\{p, q, r, w\})$ within the interval $(\delta_1, \lambda_{pq}]$. Therefore, if w lies at that moment in L_{pq}^- (so r lies then in L_{pq}^+), the Delaunayhood of rw is violated right after δ_2 by p and q, and otherwise the Delaunayhood of pq is violated right after δ_2 by r and w.

Furthermore, the open cap $C_{pq}^+ = B[p, q, w] \cap L_{pq}^+$ contains no points of P at time δ_1, which is easily seen to imply the following property:

Claim 7.1. *With the above assumptions, both edges pw and rw are Delaunay at time δ_1.*

Proof. If the the opposite cap $C_{pq}^- = B[p, q, w] \cap L_{pq}^-$ contains no points of P at time δ_1, then this co-circularity of p, q, r, w is Delaunay, and we are done. Otherwise, pq is not Delaunay even in $P \setminus \{w\}$, and each of its violating pairs in $P \setminus \{w\}$ must involve r (because $\delta_1 \in (t_0, \lambda_{pq})$). Hence, applying Lemma 4.2 to pq and r in $P \setminus \{w\}$ shows that both edges pr and rq belong at that moment to the triangulation $\mathrm{DT}(P \setminus \{w\})$. Furthermore, since pr does not belong to $\mathrm{DT}(P)$ (as $C_{pq}^- \subseteq B[p, q, r] \cap L_{pr}^-$ is not empty), the claim now follows by another application of Lemma 4.2, this time to pr and w. $\square$

By condition (A4), we have a time $\xi_{pw} > \lambda_{pq} > \lambda_q$ so that pw belongs to $\mathrm{DT}(P \setminus \{r, q\})$ throughout the interval (δ_1, ξ_{pw}), and it is Delaunay at time ξ_{pw} (in addition to its Delaunayhood at time δ_1).

Finally, by condition (A3), the edge rp undergoes in $P \setminus \{q\}$ a single Delaunay crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$, where w enters $L_{rp}^+ = L_{pr}^-$ in the interval (δ_1, λ_q). Hence, Lemma 4.1 implies that pw belongs $\mathrm{DT}(P \setminus \{q\})$ throughout the interval $\mathcal{T} = [\tau_0, \tau_1]$, which clearly intersects $(\delta_1, \xi_{pw}) \supset (\delta_1, \lambda_q)$. Similarly, the edge wq undergoes in $P \setminus \{p\}$ a Delaunay crossing by r.

In what follows, we consider a subfamily Σ_{rw}^A of all terminal quadruples $\varrho' = (p', q', r, w)$ in Σ^A whose last two points are equal to, respectively, r and w. In particular, Σ_{rw}^A includes the terminal quadruple $\varrho = (p, q, r, w)$ under consideration. Note that each $\varrho' = (p', q', r, w) \in \Sigma_{rw}^A$ is accompanied by a clockwise (r, w)-crossing $(rp', w, \mathcal{T}')$ (which occurs within an appropriately reduced triangulation $\mathrm{DT}(P \setminus \{q'\})$.

To enforce a Delaunay crossing of by rw by p or q, we fix a pair of constants $k \ll \ell$ and distinguish between three possible cases, treating each in turn.

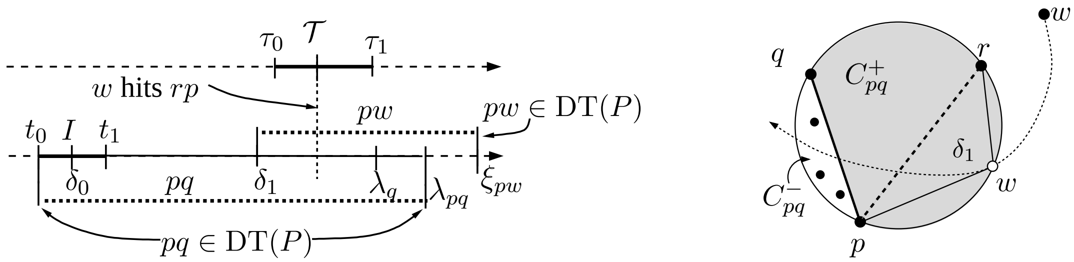

Figure 106: A partial summary of the properties of a terminal quadruple $\varrho = (p, q, r, w)$ of type A. Left: Various events occur in the depicted order (and δ_2 occurs in $(\delta_1, \lambda_{pq}]$). Notice that w hits rp in $[\tau_0, \tau_1] \cap (\delta_1, \lambda_q)$. Right: The edges pw and rw are Delaunay at time δ_1, because the open cap C^+_{pq} contains then no points of P.

Case (a) The crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$ begins after δ_1 and Σ^A_{rw} contains at least k terminal quadruples $\varrho' = (p', q', r, w)$ whose respective clockwise (r, w)-crossings $(rp', w, \mathcal{T}')$ begin in $[\delta_1, \tau_0)$, or $[\tau_0, \tau_1]$ ends before ξ_{pw} and Σ^A_{rw} contains at least k terminal quadruples $\varrho' = (p', q', r, w)$ whose respective clockwise (r, w)-crossings $(rp', w, \mathcal{T}')$ end in $(\tau_1, \xi_{pw}]$.

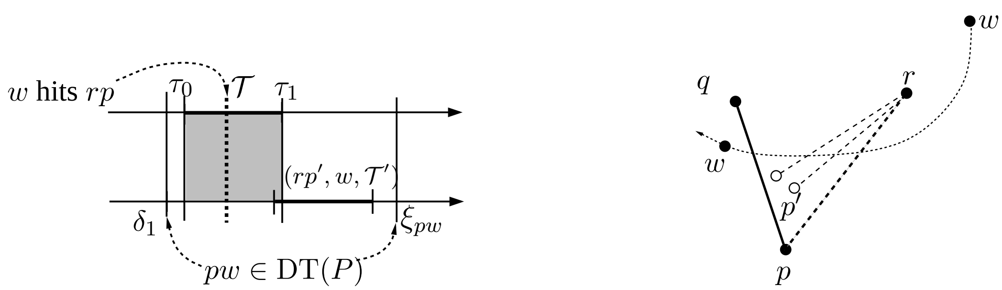

Figure 107: Case (a): The scenario where $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$ ends before ξ_{pw}, and the family Σ^A_{rw} contains at least k terminal quadruples $\varrho' = (p', q', r, w)$ whose respective (r, w)-crossings $(rp', w, \mathcal{T}')$ end in $(\tau_1, \xi_{pw}]$. At at least $k - 2$ of these quadruples satisfy $p' \neq q$ and $q' \neq p$, so their respective intervals $\mathcal{T}'$ are entirely contained in $[\tau_0, \xi_{pw}]$.

Assume without loss of generality that the latter scenario occurs, so at least k clockwise (r, w)-crossings $(rp', w, \mathcal{T}')$ end in $(\tau_1, \xi_{pw}]$; see Figure 107. Notice that each of them occurs within a smaller triangulation $\mathrm{DT}(P \setminus \{q'\})$ which is, in general, distinct from the ambient triangulation $\mathrm{DT}(P \setminus \{q\})$ of $(rp, w, \mathcal{T})$. Fortunately, any terminal quadruple $\varrho' = (p', q', r, w) \in \Sigma^A_{rw}$ is uniquely determined by each of its respective points p' and q'. Hence, at least $k - 2$ of the above quadruples ϱ' satisfy $p' \neq q$ and $q' \neq p$, in which case their respective (r, w)-crossings are compatible with $(rp, w, \mathcal{T})$ (through Lemma 5.5) and, therefore, occur entirely within $[\tau_0, \xi_{pw}] = \mathcal{T} \cup (\tau_1, \xi_{pw}]$.

We sample a subset $\hat{P}$ of $n/4$ points and argue that, with some positive fixed probability, $(rp, w, \mathcal{T})$ becomes a $(p, w, \Theta(k))$-chargeable Delaunay crossing within $\mathrm{DT}(\hat{P})$. Namely, we notice that the following two events occur simultaneously with some fixed positive probability: (1) $\hat{P}$ includes the three points p, r, w, but not q, and (2) $\hat{P}$ includes p' but not q' for at least some constant fraction of the above quadruples $\varrho' = (p', q', r, w) \in \Sigma^A_{rw}$ (whose respective (r, w)-crossings $(rp', w, \mathcal{T}')$ end in (τ_1, ξ_{pw})). In the case of success, condition (1) implies that rp still undergoes a single Delaunay crossing by w in $\hat{P}$, which occurs in some sub-interval of $\mathcal{T} = [\tau_0, \tau_1] \subset [\tau_0, \xi_{pw}]$. Similarly, condition (2) implies that at least $\Omega(k)$ clockwise (r, w)-crossings in R occur within $[\tau_0, \xi_{pw}]$.

By Theorem 5.3, the overall number of such triples (p, w, r) in $\hat{P}$ (and, thereby, in P) cannot exceed $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$. Clearly, this also bounds the overall number of the corresponding terminal quadruples $\varrho = (p, q, r, w)$ in P. If $(rp, w, \mathcal{T})$ ends before ξ_{pw}, and Σ^A_{rw} contains at least k terminal quadruples ϱ' whose respective (r,w)-crossings $(rp', w, \mathcal{T}')$ end in $(\tau_1, \xi_{pw}]$, we argue in a fully symmetrical manner, for the same upper bound on the number of such terminal quadruples ϱ.

We thus can assume, in what follows, that either the crossing $(rp, w, \mathcal{T} = [\tau_0, \tau_1])$ ends after ξ_{pw}, or the sub-family Σ_{rw}^A contains at most k other quadruples $\varrho' = (p, q, r, w)$ whose respective (r,w)-crossings $(rp', w, \mathcal{T}')$ end in $(\tau_1, \xi_{pw}]$. Similarly, we can assume that either $[\tau_0, \tau_1])$ begins before δ_1, or the sub-family Σ_{rw}^A contains at most k other quadruples $\varrho' = (p', q', r, w)$ whose respective (r,w)-crossings $(rp', w, \mathcal{T}')$ begin in $[\delta_1, \tau_0)$.

Case (b) The family Σ_{rw}^A contains no terminal quadruple $\varrho' = (p', q', w, r) \neq \varrho$ that satisfies $q' \neq p$, and whose respective (r,w)-crossing $(p'r, w, \mathcal{T}')$ ends in $[\xi_{pw}, \infty)$.

Since case (a) has been ruled out (and Σ_{rw}^A contains at most one quadruple $\varrho' = (p', q', r, w)$ with $q' = p$), we conclude that there exist at most $k+1$ terminal quadruples $\varrho' \in \Sigma_{rw}^A$ whose respective (r,w)-crossings $(p'r, w, \mathcal{T}')$ end after $\mathcal{T} = [\tau_0, \tau_1]$. Hence, we charge (p, q, r, w) (via its respective (r,w)-crossing $(pr, w, \mathcal{T} = [\tau_0, \tau_1])$) to the edge rw and notice that any edge can be charged in this manner by at most $k+2$ terminal quadruples.

To conclude, the above scenario is encountered for at most $O(kn^2)$ terminal quadruples ϱ.

Case (c) None of the previous cases occurs. In particular, since case (b) has been ruled out, the family Σ_{rw}^A contains at least one quadruple $\varrho' = (p', q', r, w) \neq \varrho$, with $q' \neq p$, and whose respective (r,w)-crossing $(rp', w, \mathcal{T}')$ ends in $[\xi_{pw}, \infty)$. (Clearly, we have $q' \neq q$, for otherwise ϱ would coincide with ϱ'.)

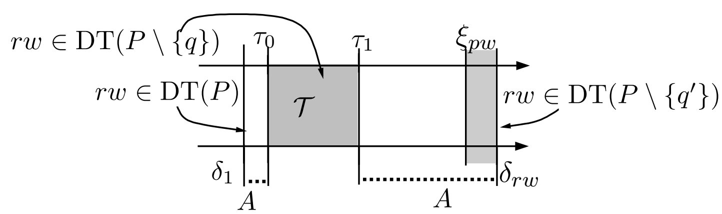

Figure 108: Case (c): Extending the almost-Delaunayhood of rw to $[\delta_1, \delta_{rw}]$. Here δ_{rw} is the first time in $[\xi_{pw}, \infty)$ when rw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{q'\})$, for some $q' \neq p, q$.

Applying Lemma 4.1 to the crossing $(rp', w, \mathcal{T}')$ (in its ambient set $P \setminus \{q'\}$) implies, then, there is a time $\delta_{rw} \geq \xi_{pw}$ which is the first such time when the edge rw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{q'\})$, where $q' \neq p, q$. In what follows, we use q' to denote a fixed point in $P \setminus \{p, q, r, w\}$ whose removal restores the Delaunayhood at time δ_{rw}; see Figure 108.

Note that we have $\delta_{rw} > \lambda_{pq} > \delta_1$. Since case (a) has been ruled out, the choice of δ_{rw} guarantees that, unless δ_{rw} belongs to $\mathcal{T} = [\tau_0, \tau_1]$, there exist at most $k+1$ quadruples $\varrho' \in \Sigma_{rw}^A$ whose respective (r,w)-crossings $(rp', w, \mathcal{T}')$ end in $(\tau_1, \delta_{rw}]$. (In particular, by the choice of δ_{rw}, there is at most one quadruple $\varrho' = (p', q', r, w)$ whose respective (r,w)-crossing ends in $(\xi_{pw}, \delta_{rw}]$, and it must satisfy $q' = p$.)

Charging events in $\mathcal{A}_{rw}$. We next invoke Theorem 2.2 in order to extend the almost-Delaunayhood of rw, which already belongs to $\mathrm{DT}(P \setminus \{q\})$ throughout $\mathcal{T} = [\tau_0, \tau_1]$ (by Lemma 4.1), to the interval $[\delta_1, \delta_{rw}]$, which clearly intersects $\mathcal{T}$.

Note that $[\delta_1, \delta_{rw}] \setminus \mathcal{T}$ is composed of two disjoint (and possibly empty) sub-intervals $[\delta_1, \tau_0)$ and $(\tau_1, \delta_{rw}]$. We apply Theorem 2.2 separately over each of these sub-intervals (and only if they are non-empty). In both cases, we use the second threshold parameter $\ell \gg k$.

The first application of Theorem 2.2 in $\mathcal{A}_{rw}$, over (δ_1, τ_0), can be done with respect to the complete point set P (using the Delaunayhood of rw at time δ_1, given in Claim 7.1). It is necessary only if $\delta_1 < \tau_0$.

If at least one of the conditions (i), (ii) of that theorem is satisfied, we charge ϱ within $\mathcal{A}_{rw}$, via $(rp, w, \mathcal{T})$, either to an ℓ-shallow collinearity, or to $\Omega(\ell^2)$ ℓ-shallow co-circularities during (δ_1, τ_0). Since case (a) has been ruled out, $(rp, w, \mathcal{T})$ is among the first $k+1$ such (r, w)-crossings to begin after any event that we charge. Hence, the above charging accounts for at most $O(k\ell^2 N(n/\ell) + k\ell n^2\beta(n))$ quadruples $\varrho \in \Sigma^A$. Otherwise, we end up with a subset of at most 3ℓ points (perhaps including p or q, or both) whose removal restores the Delaunayhood of rw throughout $[\delta_1, \tau_0]$.

The similar second application of Theorem 2.2 (over (τ_1, δ_{rw})) is done with respect to the reduced point set $P \setminus \{q'\}$ (where q' denotes the point whose removal restores the Delaunayhood of rw at time δ_{rw}). It is necessary only if $\tau_1 < \delta_{rw}$.

If at least one of the conditions (i), (ii) of that theorem holds, we charge ϱ (via $(rp, w, \mathcal{T})$) within $\mathcal{A}_{rw}$ either to an $(\ell+1)$-shallow collinearity, or to $\Omega(\ell^2)$ $(\ell+1)$-shallow co-circularities (which are ℓ-shallow with respect to $P \setminus \{q'\}$). By the choice of δ_{rw}, $(rp, w, \mathcal{T})$ is among the last $k+2$ such (r, w)-crossings to end after the event, so any $(\ell+1)$-shallow event in $\mathcal{A}_{rw}$ is charged by at most $O(k)$ quadruples ϱ. Otherwise, we end up with a subset of at most $3\ell+1$ points (inclding q', and perhaps also some of p, q) whose removal restores the Delaunayhood of rw throughout $[\tau_1, \delta_{rw}]$.

To conclude, we may assume that there is a subset A_{rw} of at most $6\ell+1$ points (including q') whose removal restores the Delaunayhood of rw throughout $[\delta_1, \delta_{rw}]$. To obtain the crossing of rw by p or q (which would occur in, respectively, $\mathrm{DT}((P \setminus A_{rw}) \cup \{p\})$ or $\mathrm{DT}((P \setminus A_{rw}) \cup \{q\})$), it suffices to show that rw is hit by one of these two points during the interval $[\delta_1, \delta_{rw}]$. Notice that the latter interval contains $\delta_2 \in (\delta_1, \lambda_{pq}) \subset (\delta_1, \xi_{pw}]$. See Figure 109. To do so, we distinguish between two possible sub-scenarios, depending on the precise order type of the co-circularity (at time) δ_2, which is red-blue with respect to pq and rw.

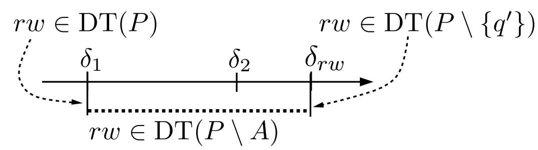

Figure 109: Case (c): The edge rw belongs to $\mathrm{DT}(P \setminus A)$ throughout the interval $[\delta_1, \delta_{rw}]$, which contains the last co-circularity δ_2 of p, q, r, w. In addition, rw belongs to $\mathrm{DT}(P)$ and $\mathrm{DT}(P \setminus \{q'\})$ at times δ_1 and δ_{rw}, respectively.

If r lies in L^-_{pq} when w enters L^-_{pq} (through pq), then the Delaunayhood of rw is violated right *after* δ_2 by $p \in L^-_{rw}$ and $q \in L^+_{rw}$, as depicted in Figure 110 (left). Since δ_2 is the *last* co-circularity of p, q, r, w, Lemma 3.1 implies that rw is hit during $(\delta_2, \delta_{rw}]$ by at least one of p, q (because $q' \neq p, q, r, w$), so we are done.

Assume, then, that r lies in L^+_{pq} when w enters L^-_{pq}, so the Delaunayhood of rw is violated right *before* δ_2 by $p \in L^+_{rw}$ and $q \in L^-_{rw}$, as depicted in Figure 110 (right). Notice that this violation does not hold at time δ_1. Hence, we can obtain the desired crossing of rw in (δ_1, δ_2) by applying the time-reversed variant of Lemma 3.1 (from δ_2). The crucial observation is that δ_1, δ_2 have different order types, which rules out the last case in Lemma 3.1.

If rw is hit during $(\delta_1, \delta_{rw}]$ by the point p, then the triple p, r, w performs two Delaunay crosings within the triangulation $\mathrm{DT}((P \setminus A_{rw}) \cup \{p\})$, namely, $(rp, w, \mathcal{T})$ and the just established crossing of wr by p. Otherwise, if rw is hit during $(\delta_1, \delta_{rw}]$ by q, the other triple q, r, w performs two Delaunay crossings within the triangulation $\mathrm{DT}((P \setminus A) \cup \{q\})$, namely, the crossing of wq by r (as prescribed by condition (A3)) and the just established crossing of rw by q.

In both cases, a standard combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that the overall number of the corresponding triples (p, r, w) or (q, r, w) in P cannot exceed $O(\ell n^2)$. Since the quadruple ϱ at hand is uniquely determined by each of its four sub-triples, this also bounds the overall number of such quadruples in Σ^A.

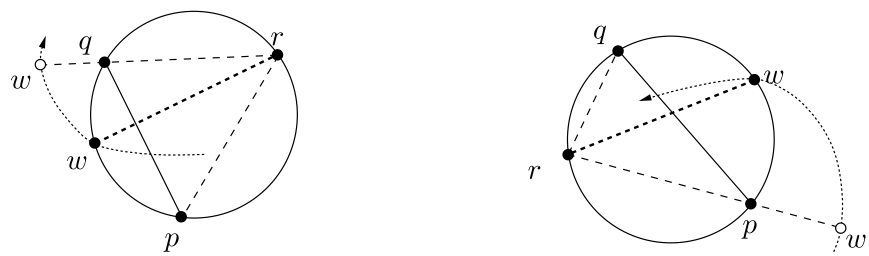

Figure 110: Case (c): Left: A possible motion of w if it lies at time δ_2 in the halfplane L_{pq}^- (so r lies then in L_{pq}^+). The Delaunayhood of rw is violated right after this event by p and q, so at least one of them must cross rw during (δ_2, δ_{rw}). Right: A possible motion of w if it lies at time δ_2 in the halfplane L_{pq}^- (so r lies then in L_{pq}^+). The Delaunayhood of rw is violated right after this event by p and q, so at least one of them must cross rw during (δ_1, δ_2).

To conclude, we have established the following bound on the maximum possible cardinality of Σ^A:

$$T^A(n) = O\left(k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right). \tag{12}$$

Notice that we have expressed the maximum possible number of terminal quadruples of type A in terms of more elementary quantities which were introduced in Section 2.

7.2 Terminal quadruples of type B

In this section we at last express the maximum possible cardinality $T^B(n)$ of a family Σ^B of terminal quadruples of type B (where each quadruple $(p, q, r, w) \in \Sigma^B$ is uniquely determined by each of the respective sub-triples (p, q, r), (p, q, w) and (q, r, w)) in terms of more elementary quanitities that were introduced in Section 2. To do so, we fix the underlying set P of n moving points, a family Σ^B as above, and a terminal quadruple $\varrho = (p, q, r, w)$ of type B in Σ^B.

Recall that the four points of ϱ perform (at least) three Delaunay crossings, namely the crossing of pq by r, the crossing of qp by w, and the crossing of qr by w. Here only the first crossing, namely, $(pq, r, I = [t_0, t_1])$, is defined with respect to the compelete point set P. Each of the remaining three crossings of ϱ occurs within a reduced point set, which is obtained from P by removing the fourth point of ϱ (not directly involved in the crossing).

In the course of this section, we will enforce on the points of ϱ two additional crossings, namely the crossing of pw by one of q, r, and, finally, the crossing of rw by one of p, q. As a result, at least one of the triples (p, q, w), (p, r, w) or (q, r, w) will perform two Delaunay crossings (within an appropriately reduced triangulation). We will thus charge ϱ to that triple and bottom out by invoking Lemma 4.5.

Topological setup. Refer to Figure 111. By condition (B1), the edge pq is crossed by r (during $I = [t_0, t_1]$, as part of the corresponding Delaunay crossing) and w (at some later time $\lambda_q > t_1$), in opposite directions. Furthermore, pq re-enters $\mathrm{DT}(P)$ at some later time $\lambda_{pq} \geq \lambda_q$, and it belongs to $\mathrm{DT}(P \setminus \{r, w\})$ throughout $[t_0, \lambda_{pq}]$.

By condition (B2), the four points of p, q, r and w are co-circular at some three times $\delta_0 \in I$, $\delta_1 \in (t_1, \lambda_{pq}]$ and $\delta_2 \in (\delta_1, \lambda_{pq}]$, where the two extremal co-circularities (at times δ_0 and δ_2) are red-blue with respect to pq, and the middle co-circularity (at time δ_1) is red-red with respect to pq (and red-blue with respect to rq). Clearly, r remains in L_{pq}^+ throughout (t_1, δ_1) after entering this halfplane during I (for otherwise r would have to cross L_{pq} three times).

As a matter of fact, the last co-circularity at time δ_2 arises as part of a single Delaunay of qp by w, which occurs in the triangulation $\mathrm{DT}(\{p, q, r, w\})$ within the interval $(\delta_1, \lambda_{pq}]$. Therefore, if w lies at that moment in L_{pq}^- (so r lies then in L_{pq}^+), the Delaunayhood of rw is violated right after δ_2 by p and q, and otherwise the Delaunayhood of pq is violated right after δ_2 by r and w.

Furthermore, the open cap $C_{pq}^+ = B[p,q,w] \cap L_{pq}^+$ contains no points of P at time δ_1. Using Lemma 4.2, we obtain the following property:

Claim 7.2. *With the above assumptions, both edges wq and rw belong to $\mathrm{DT}(P)$ at time δ_1. Furthermore, the edge pw belongs then to $\mathrm{DT}(P \setminus \{r\})$.*

Proof. The first part of the claim is fully symmetric to Claim 7.1, and can be established using a fully symmetric argument (switching the roles of p and q). We thus proceed to proving the Delaunayhood of pw in $P \setminus \{w\}$. Indeed, if the opposite cap $C_{pq}^- = B[p,q,w] \cap L_{pq}^-$ contains no points of P at time δ_1, then this co-circularity of p,q,r,w is Delaunay, and we are done. Otherwise, pq is not Delaunay even in $P \setminus \{r\}$, and each of its violating pairs in $P \setminus \{r\}$ must involve w (because $\delta_1 \in (t_0, \lambda_{pq})$). Hence, Lemma 4.2 implies that pw belongs at that moment to the triangulation $\mathrm{DT}(P \setminus \{r\})$. $\square$

By condition (B4), we have a time $\xi_{wq} > \lambda_{pq} > \lambda_q$ so that wq belongs to $\mathrm{DT}(P \setminus \{p,r\})$ throughout the interval $[\delta_1, \xi_{wq}]$, and it is Delaunay at time ξ_{wq} (in addition to its Delaunayhood at time δ_1).

Finally, by condition (B3), the edge qr undergoes in $P \setminus \{p\}$ a single Delaunay crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$, where w enters $L_{qr}^+ = L_{rq}^-$ in the interval (δ_1, λ_q). Hence, Lemma 4.1 implies that wq belongs $\mathrm{DT}(P \setminus \{p\})$ throughout the interval $\mathcal{T} = [\tau_0, \tau_1]$, which clearly intersects $[\delta_1, \xi_{wq}] \supset [\delta_1, \lambda_q]$.

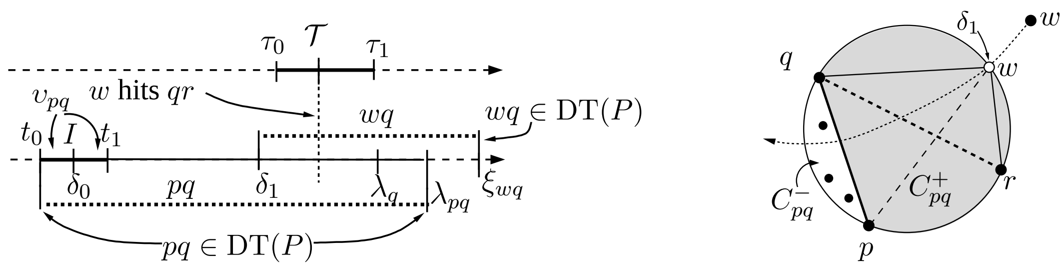

Figure 111: A partial summary of the properties of a terminal quadruple $\varrho = (p,q,r,w)$ of type B. Left: Various events occur in the depicted order (and δ_2 occurs in $(\delta_1, \lambda_{pq}]$). v_{pq} is the time in I at which r hits pq. Notice that w hits qr in $[\tau_0, \tau_1] \cap (\delta_1, \lambda_q)$. Right: The edges wq and rw are Delaunay at time δ_1, and pw belongs to $\mathrm{DT}(P \setminus \{r\})$, because the open cap C_{pq}^+ contains then no points of P.

Overview. Clearly, the motion of p,q,r and w still obeys Proposition 6.11. In particular, using v_{pq} to denote the time[45] in I when r enters L_{pq}^+ through pq, the edge pw is hit in (v_{pq}, δ_1) by at least one of the points q, r. Namely, q crosses pw from L_{pw}^+ to L_{pw}^-, or r crosses pw in the reverse direction. Our analysis proceeds in two steps. At the first step, we refine this collinearity into a full-fledged Delaunay crossing of pw. If pw (or, more precisely, its reversely oriented copy wp) is crossed by q, then our analysis bottoms out through Lemma 4.5. If pw is crossed by r, we proceed to the second step, at which we enforce a Delaunay crossing of rw by at least one of p,q. At this step, our analysis is fully symmetric to the one that was used in Section 7.1 to enforce the same type of crossing.

Part 1: Enforcing the crossing of pw by q or r. We consider the subfamily Σ_{pw}^B of all terminal quadruples $\varrho' = (p, q', r', w) \in \Sigma^B$ of type B whose first and last points are equal to, respectively, p and w. (By the definition of Σ^B, each $\varrho' \in \Sigma_{pw}^B$ is uniquely determined by its respective point q'.) In particular, Σ_{pw}^B includes the quadruple ϱ under consideration.

[45]Note that the order of δ_0 and v_{pr} is unknown, and it is determined by the location of w at time v_{pq}.

For each $\varrho' \in \Sigma^B_{pw}$ we use δ'_1 to denote the respective time when p, q', r', w are involved in a red-red co-circularity with respect to pq', as prescribed by Condition (B2). We emphasize that, by condition (B2), the open cap $C^+_{pq'} = B[p, q', w] \cap L^+_{pq}$ contains no points of P at time δ'_1, so the edge pw belongs at that moment to the triangulation $\mathrm{DT}(P \setminus \{r'\})$.

Let υ_{pw} denote the time when pw is hit by r or q, as prescribed in Proposition 6.11. Namely, we assume that $\upsilon_{pw} < \delta_1$, and that the Delaunayhood of pw is violated by $r \in L^-_{pw}$ and $q \in L^+_{pw}$ throughout the interval $(\upsilon_{pw}, \delta_1)$.

Proposition 7.3. *With the above notation, there exist no terminal quadruples $\varrho' \in \Sigma^B_{pw}$ whose respective second co-circularities δ'_1 occur in $(\upsilon_{pw}, \delta_1)$.*

Proof. Assume for a contradiction that there is a terminal quadruple $\varrho' = (p, q', r', w)$ whose respective time δ'_1 belongs to $(\upsilon_{pw}, \delta_1)$, where the Delaunayhood of pw is violated by $q \in L^-_{pw}$ and $r \in L^+_{pw}$. Note that $q \neq q'$. By Claim 7.2, pw belongs to $\mathrm{DT}(P \setminus \{r'\})$ at time δ'_1. Therefore, and since both r and r' lie then in L^+_{pw}, we obtain $r = r'$. (Otherwise, the Delaunayhood of pw would be violated at time δ'_1 by the points r and q, none of them equal to r'.) In other words, ϱ and ϱ' differ only in their second points. Hence, q lies at time δ'_1 within the disc $B[p, q', r] = B[p, q', w]$.

Since q cannot lie at time δ'_1 inside the cap $C^+_{pw} = B[p, q', w] \cap L^+_{pq'}$, it has to lie inside the complementary cap $C^-_{pq} = B[p, q', w] \cap L^-_{pq'}$, which coincides with $B[p, q', r] \cap L^-_{pq'}$. In other words, the Delaunayhood of both pq' and pw is violated at time δ'_1 by $q \in L^-_{pq'}$ and $r \in L^+_{pq'}$. See Figure 112.

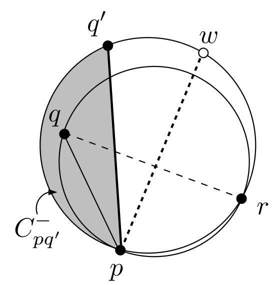

Figure 112: Proof of Proposition 7.3. We assume that $\varrho' = (p, q', r, w)$ is a terminal quadruple in Σ^B_{pw}, whose second co-circularity occurs at time $\delta'_1 \in (\upsilon_{pw}, \delta_1)$. The point q must lie at time δ'_1 in the cap $C^-_{pq} = B[p, q', w] \cap L^-_{pq'}$, which coincides with $B[p, q', r] \cap L^-_{pq'}$.

Recall that δ'_1 occurs after the end of the respective (p, r)-crossing $(pq', r, I' = [t'_0, t'_1])$ of ϱ' (which is prescribed by condition (B1)). Since the disc $B[p, q', r]$ contains no points of P right after time t'_1 (and the motion of $B[p, q', r]$ is continuous throughout (t'_1, δ'_1)), the point q must enter the cap $B[p, q', r] \cap L^-_{pq'}$ in (t'_1, δ'_1). Furthermore, conditions (B1) and (B2) imply that q cannot hit pq' in (t'_1, δ'_1), so q can enter $B[p, q', r] \cap L^-_{pq'}$ only through the boundary of $B[p, q', r]$, at a common co-circularity of p, q, q', r. See Figure 113 (left). In what follows, we use δ' to denote the time of (the last) such co-circularity in (t'_1, δ'_1), noting that q remains in $B[p, q', r] \cap L^-_{pq'}$ throughout (δ', δ'_1).

We claim that δ' occurs after $I = [t_0, t_1]$; see Figure 113 (right). Indeed, since q lies in $L^-_{pr} (\supset B[p, q', r] \cap L^-_{pq'})$ throughout (δ', δ'_1), and since $\delta'_1 > \upsilon_{pw} > \upsilon_{pq}$, we obtain that $\upsilon_{pq} < \delta'$ (for, otherwise, υ_{pq} would belong to (δ', δ'_1)). Furthermore, by Lemma 4.1, δ' cannot occur during $I = [t_0, t_1]$, because it is (the time of) a red-blue co-circularity with respect to rq. Therefore, we have $\delta' > t_1$.

To conclude, q enters $B[p, q', r] \cap L^-_{pq'}$ at a common co-circularity of p, q, q', r, and only after the ends of I and I'. According to Lemma 4.4, the points p, q, q', r are involved in at least two previous co-circularities in the intervals $I \setminus I'$ and $I' \setminus I$. Hence, the co-circularity at time δ' has index 3. Note that the Delaunayhood of pq' is violated by $q \in L^-_{pq'}$ and $r \in L^+_{pq'}$ throughout the interval (δ', δ'_1). Moreover, since ϱ' satisfies condition (B1), the edge pq' re-enters $\mathrm{DT}(P)$ at some time $\lambda_{pq'} > \delta'_1$. Since δ'_1 is the

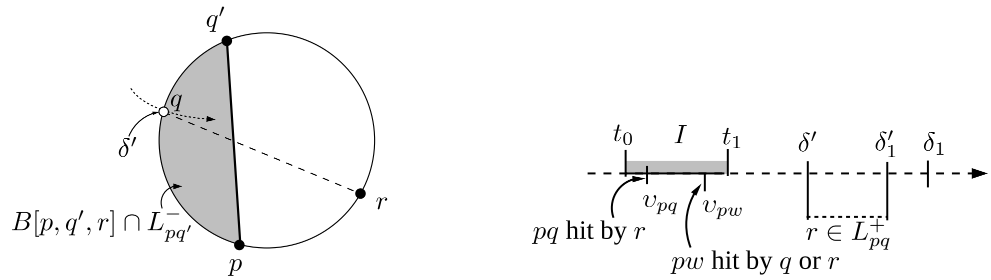

Figure 113: Proof of Proposition 7.3. Left: q must enter $B[p, q', r] \cap L^-_{pq'}$ at some time $\delta' \in (t'_1, \delta'_1)$. The Delaunayhood of pq' is violated by q and r throughout (δ', δ'_1). Right: Arguing that δ' occurs after I. Various events occur in the depicted order. The co-circularity at time δ' occurs after I, so it is the third co-circularity of p, q, r, q'.

last co-circularity of p, q', r, w, Lemma 3.1 implies that the edge pq' is hit during $(\delta'_1, \lambda_{pq'}]$ by at least one of the points q, r, contrary to Condition (B1) on ϱ'. This last contradiction completes the proof of Proposition 7.3. □

Note that the subfamily Σ^B_{pw} can contain at most one quadruple $\varrho' = (p, q', r', w)$ with $q' = r$. Applying the pigeonhole principle (as this was done in Section 5.6) we get that at least half of all terminal quadruples $\varrho = (p, q, r, w) \in \Sigma^B_{pw}$ satisfy the following condition:

(PHT) *There is at most one quadruple $\varrho' = (p, q', r', w) \in \Sigma^B_{pw}$ that satisfies $r' = q$.*

With no loss of generality, we can assume, in what follows, that (PHT) holds for the terminal quadruple $\varrho = (p, q, r, w)$ under consideration. To proceed, we distinguish between two possible cases.

Case (1a). The edge pw is hit at time υ_{pw} by q, which crosses pw from L^+_{pw} to L^-_{pw}.

Assume first that there exist no terminal quadruples $\varrho' = (p, q', r', w)$ in Σ^B_{pw}, with $r' \neq q$, and whose respective second co-circularities δ'_1 occur before υ_{pw}. In this scenario, Proposition 7.3 together with condition (PHT) imply that δ_1 is among the first two such second co-circularities δ'_1 of terminal quadruples $\varrho' \in \Sigma^B_{pw}$, so we can charge ϱ (via δ_1) to the edge pw. Clearly, this can happen for $O(n^2)$ terminal quadruples $\varrho \in \Sigma^B_{pw}$.

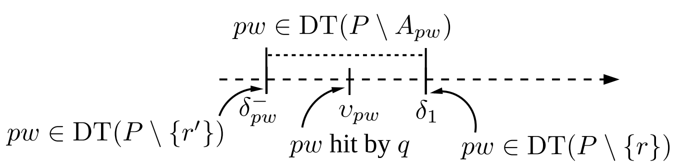

Figure 114: Case (1a): pw is hit by q at time υ_{pw}. We choose δ^-_{pw} as the last time before υ_{pw} when pw belongs to a reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, for some $r' \neq q$, and apply Theorem 2.2 over $(\delta^-_{pw}, \delta_1)$.

To conclude, we may assume in what follows that the above scenario does not occur. Recall that, for each $\varrho' \in \Sigma^B_{pw}$, the edge pw belongs to $\mathrm{DT}(P \setminus \{r'\})$ at the respective time δ'_1. Hence, there is a time $\delta^-_{pw} < \upsilon_{pw}$ which is the last such time when pw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, for $r' \neq p, q, w$.

We apply Theorem 2.2 for pw in the interval $(\delta^-_{pw}, \delta_1)$, with the first threshold parameter k. This is done with respect to the reduced set $P \setminus \{r, r'\}$ (to ensure the Delaunayhood of pw at times δ^-_{pw} and δ_1). Refer to Figure 114.

In cases (i), (ii) of Theorem 2.2, we encounter in the appropriately reduced red-blue arrangement $\mathcal{A}^{(r,r')}_{pw}$ of pw (defined with respect to $P \setminus \{r, r'\}$) either a k-shallow collinearity or $\Omega(k^2)$ k-shallow co-circularities, and charge ϱ to these events, which are $(k+2)$-shallow with respect to the original set P. Notice that each $(k+2)$-shallow event in the full arrangement $\mathcal{A}_{pw}$ is charged in this manner by at most $O(1)$ terminal quadruples $\varrho \in \Sigma_{pw}^B$, whose respective second co-circularities δ_1 are among the first two such co-circularities to occur in $\mathcal{A}_{pw}$ after the time of the event. Hence, the above charging accounts for at most $O(k^2 N(n/k) + kn^2\beta(n))$ terminal quadruples $\varrho \in \Sigma^B$.

Assume, then, that condition (iii) of Theorem 2.2 is satisfied. That is, there is a subset A_{pw} of at most $3k+2$ points (including r and r') whose removal restores the Delaunayhood of pw in $[\delta_{pw}^-, \delta_1]$. Moreover, since $q \neq r, r'$, the edge pw belongs to $\mathrm{DT}((P \setminus A_{pw}) \cup \{q\})$ at both times δ_{pw}^- and δ_{pw}^+. Therefore, the triple p, q, w performs two (single) Delaunay crossings in the reduced set $(P \setminus A_{pw}) \cup \{q\}$, namely, the crossing of qp by w, and the crossing of wp by q. A rountine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor shows that the overall number of such triples (p, q, w) (and, therefore, of their corresponding terminal quadruples $\varrho \in \Sigma^B$) is at most $O(kn^2)$.

In conclusion, we have shown that at most $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$ terminal quadruples fall into case (1a).

Case (1b). The edge pw is hit at time υ_{pw} by the point r, which crosses pw from L_{pw}^- to L_{pw}^+.

Notice that, by Proposition 7.3, each terminal quadruple $\varrho = (p, q, r, w)$ falling into case (1a) is uniquely determined by the choice of (p, r, w), because the second co-circularity δ_1 of ϱ is the first co-circularity of this kind (over all $\varrho' = (p, q', r, w) \in \Sigma^B$) to occur after the *unique* time when r enters the halfplane L_{pw}^+ through pw.

If there exists no terminal quadruple $\varrho' = (p, q', r', w) \in \Sigma_{pw}^B$ whose respective second co-circularity δ_1' occurs before υ_{pw}, Proposition 7.3 implies that δ_1 is the first such co-circularity, so ϱ can be charged to the edge pw. Clearly, this accounts for at most $O(n^2)$ terminal quadruples ϱ.

For each of the remaining quadruples $\varrho \in \Sigma_{pw}^B$ (that fall into case (1b)), Σ_{pw}^B contains another quadruple $\varrho' = (p, q', r', w)$, necessarily with $r' \neq r$, so that the edge pw is Delaunay in $P \setminus \{r'\}$ at the time $\delta_1' < \upsilon_{pw}$ of the respective second co-circularity of ϱ'. In particular, we can choose a time $\delta_{pw}^- < \upsilon_{pw}$ which is the last such time when pw belongs to a reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, for some $r' \neq p, w, r$.

Similarly, if there exists no quadruple $\varrho' = (p, q', r', w) \in \Sigma_{pw}^B$ whose respective second co-circularity δ_1' occurs after δ_1, we can charge ϱ (via its respective time stamp δ_1) to pw. Otherwise, there is a time δ_{pw}^+ which is the first such time when pw belongs to a reduced triangulation $\mathrm{DT}(P \setminus \{r''\})$, for some $r'' \neq p, w, r$.

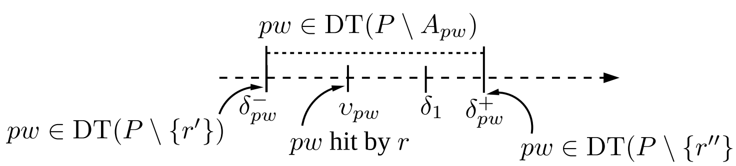

Figure 115: Case (1b): pw is hit by r at time υ_{pw}. We choose δ_{pw}^- as the last time before υ_{pw} when pw belongs to a reduced triangulation $\mathrm{DT}(P \setminus \{r'\})$, for some $r' \neq r$, and apply Theorem 2.2 over $(\delta_{pw}^-, \delta_1)$. In addition, we choose δ_{pw}^+ as the first time after δ_1 when pw belongs to a reduced triangulation $\mathrm{DT}(P \setminus \{r''\})$, for some $r'' \neq r$, and apply Theorem 2.2 over $(\delta_1, \delta_{pw}^+)$.

For each of the remaining quadruples $\varrho \in \Sigma^B$ (that fall into case (1b)) there exist times $\delta_{pw}^- < \upsilon_{pw} < \delta_1$ and $\delta_{pw}^+ > \delta_1$ as above, with respective obstruction points $r', r'' \notin \{p, w, r\}$; refer to Figure 115. We can now apply Theorem 2.2 for the edge pw, over the interval $(\delta_{pw}^-, \delta_{pw}^+)$ (containing δ_1). This is done with the threshold k, and with respect to the reduced point set $P \setminus \{r', r''\}$.

In cases (i) and (ii) of Theorem 2.2, we charge ϱ within $\mathcal{A}_{pw}$ (via δ_1) either to a $(k+2)$-shallow collinearity, or to $\Omega(k^2)$ $(k+2)$-shallow co-circularities. Note that each $(k+2)$-shallow event, that occurs in $\mathcal{A}_{pw}$ during $(\delta_{pw}^+, \delta_{pw}^+)$, is charged by at at most $O(1)$ terminal quadruples in Σ_{pw}^B (that fall into case (1b)), because the second co-circularity δ_1 of ϱ is either the last such co-circularity to occur before the time t^* of the event, or the first such co-circularity to occur after t^*. Therefore, the above charging accounts for at most $O(k^2 N(n/k) + kn^2\beta(n))$ terminal quadruples.

Now assume that Condition (iii) of Theorem 2.2 holds. That is, we have a subset A_{pw} of at most $3k+2$ points (including r' and r'') whose removal restores the almost Delaunayhood of pw throughout the interval $[\delta_{pw}^-, \delta_{pw}^+]$. Moreover, since $r \neq r', r''$, the edge pw belongs to $\mathrm{DT}((P \setminus A_{pw}) \cup \{r\})$ at both times δ_{pw}^- and δ_{pw}^+. Therefore, the edge pw undergoes a Delaunay crossing by r within the reduced triangulation $\mathrm{DT}((P \setminus A_{pw}) \cup \{r\})$.

Part 2: Enforcing a Delaunay crossing of rw. To conclude, we may assume, from now on, that each terminal quadruple $\varrho = (p, q, r, w) \in \Sigma^B$ under consideration is uniquely determined by each of its four sub-triples $(p, q, r), (p, q, w), (p, r, w)$ and (q, r, w). Moreover, each of these triples defines a Delaunay crossing (in an appropriately reduced subset of P).

We now exploit the last co-circularity of p, q, r, w (at time $\delta_2 \in [\delta_1, \lambda_{pq}]$) to enforce a fifth such crossing, namely the Delaunay crossing of rw by one of p, q. Here our argument is symmetric to the one that was used in Section 7.1. (Namely, we now switch the roles of p and q). In the case of success, at least one of the triples $(p, r, w), (q, r, w)$ performs two (single) Delaunay crossings, so Lemma 4.5 can be invoked. Otherwise, we dispose of ϱ either through Theorem 5.3, or by charging it within $\mathcal{A}_{rw}$.

Before proceeding with our case analysis, we emphasize that wq is Delaunay at times δ_1 and $\xi_{wq} > \lambda_{pq} (> \delta_2)$, and that the single Delaunay crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$ is defined with respect to a smaller point set $P \setminus \{p\}$. In addition, both $[\delta_1, \lambda_{wq}]$ and $[\tau_0, \tau_1]$ contain the time when w crosses rq from L_{rq}^+ to L_{rq}^-.

We keep $\varrho = (p, q, r, w) \in \Sigma^B$ fixed and consider a subfamily Σ_{rw}^B of all such terminal quadruples $\varrho' = (p', q', r, w) \in \Sigma^B$ whose last two points are equal to, respectively, to r and w. (In particular, Σ_{rw}^B includes the terminal quadruple $\varrho = (p, q, r, w)$ at hand.) As in the symmetric setup of Section 7.1, we distinguish between three possible scenarios (a)–(c), ruling them out one by one.

Case (a) The crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$ begins after δ_1 and Σ_{rw}^B contains at least k terminal quadruples $\varrho' = (p', q', r, w)$ whose respective counterclockwise (r, w)-crossings $(q'r, w, \mathcal{T}')$ begin in $[\delta_1, \tau_0)$, or $[\tau_0, \tau_1]$ ends before ξ_{wq} and Σ_{rw}^B contains at least k terminal quadruples $\varrho' = (p', q', r, w)$ whose respective counterlclockwise (r, w)-crossings $(q'r, w, \mathcal{T}')$ end in $(\tau_1, \xi_{wq}]$.

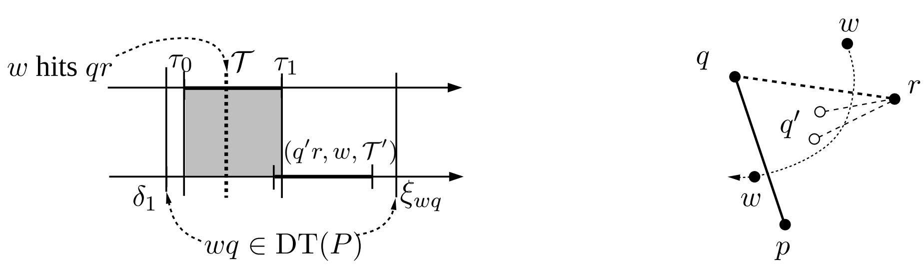

Figure 116: Case (a): The scenario where $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$ ends before ξ_{wq}, and the family Σ_{rw}^B contains at least k terminal quadruples $\varrho' = (p', q', r, w)$ whose respective (r, w)-crossings $(q'r, w, \mathcal{T}')$ end in $(\tau_1, \xi_{wq}]$. At at least $k-2$ of these quadruples satisfy $p' \neq q$ and $q' \neq p$, so their respective intervals $\mathcal{T}'$ are entirely contained in $[\tau_0, \xi_{wq}]$.

Assume without loss of generality that the latter scenario occurs, so at least k counterclockwise (r, w)-crossings $(q'r, w, \mathcal{T}')$ end in $(\tau_1, \xi_{wq}]$; see Figure 116. Notice that each of them occurs within a smaller triangulation $\mathrm{DT}(P \setminus \{p'\})$ which, in general, is distinct from the ambient triangulation $\mathrm{DT}(P \setminus \{p\})$ of $(qr, w, \mathcal{T})$. Fortunately, any terminal quadruple $\varrho' = (p', q', r, w) \in \Sigma_{rw}^B$ is uniquely determined by each of its respective points p' and q'. Hence, at least $k-2$ of the above quadruples ϱ' satisfy $p' \neq q$ and $q' \neq p$, in which case their respective (r, w)-crossings are compatible with $(qr, w, \mathcal{T})$ through Lemma 5.5, and, therefore, occur entirely within $[\tau_0, \xi_{wq}] = \mathcal{T} \cup (\tau_1, \xi_{wq}]$.

We sample a subset $\hat{P}$ of $n/4$ points and argue that, with some positive fixed probability, $(qr, w, \mathcal{T})$ becomes a $(q, w, \Theta(k))$-chargeable Delaunay crossing within $\mathrm{DT}(\hat{P})$. Namely, we notice that the following two events occur simultaneously with some fixed positive probability: (1) $\hat{P}$ includes the three points q, r, w, but not p, and (2) $\hat{P}$ includes q' but not p' for at least some constant fraction of the above quadruples $\varrho' = (p', q', r, w) \in \Sigma^B_{rw}$ (whose respective (r, w)-crossings $(q'r, w, \mathcal{T}')$ end in $(\tau_1, \xi_{wq}]$). In the case of success, condition (1) implies that qr still undergoes a single Delaunay crossing by w in $\hat{P}$, which occurs in some sub-interval of $\mathcal{T} = [\tau_0, \tau_1] \subset [\tau_0, \xi_{wq}]$. Similarly, condition (2) implies that at least $\Omega(k)$ counterclockwise (r, w)-crossings in R occur within $[\tau_0, \xi_{wq}]$.

By Theorem 5.3, the overall number of such triples (q, r, w) in $\hat{P}$ (and, thereby, in P) cannot exceed $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$, which also bounds the overall number of the corresponding terminal quadruples $\varrho = (p, q, r, w)$ in P.

We thus can assume, in what follows, that either the crossing $(qr, w, \mathcal{T} = [\tau_0, \tau_1])$ ends after ξ_{pw}, or the sub-family Σ^A_{rw} contains at most k other quadruples $\varrho' = (p, q, r, w)$ whose respective (r, w)-crossings $(q'r, w, \mathcal{T}')$ end in $(\tau_1, \xi_{wq}]$. Similarly, we can assume that either $[\tau_0, \tau_1])$ begins before δ_1, or the sub-family Σ^B_{rw} contains at most k other quadruples $\varrho' = (p', q', r, w)$ whose respective (r, w)-crossings $(q'r, w, \mathcal{T}')$ begin in $[\delta_1, \tau_0)$.

Case (b) The family Σ^B_{rw} contains no terminal quadruple $\varrho' = (p', q', w, r) \neq \varrho$ that satisfies $p' \neq q$, and whose respective (r, w)-crossing $(rq', w, \mathcal{T}')$ ends in $[\xi_{wq}, \infty)$.

Since case (a) has been ruled out (and Σ^B_{rw} contains at most one quadruple $\varrho' = (p', q', r, w)$ with $q' = p$), we conclude that there exist at most $k+1$ terminal quadruples $\varrho' \in \Sigma^B_{rw}$ whose respective (r, w)-crossings $(rq', w, \mathcal{T}')$ end after $\mathcal{T} = [\tau_0, \tau_1]$. Hence, we charge (p, q, r, w) (via its respective (r, w)-crossing $(pr, w, \mathcal{T} = [\tau_0, \tau_1])$) to the edge rw and notice that any edge can be charged in this manner by at most $k+2$ terminal quadruples.

To conclude, the above scenario happens for at most $O(kn^2)$ terminal quadruples ϱ.

Case (c) None of the previous cases occurs. In particular, since case (b) has been ruled out, the family Σ^B_{rw} contains at least one quadruple $\varrho' = (p', q', r, w) \neq \varrho$, with $p' \neq q$, and whose respective (r, w)-crossing $(q'r, w, \mathcal{T}')$ ends in $[\lambda_{wq}, \infty)$. (Clearly, we have $p' \neq p$, for otherwise ϱ would coincide with ϱ'.)

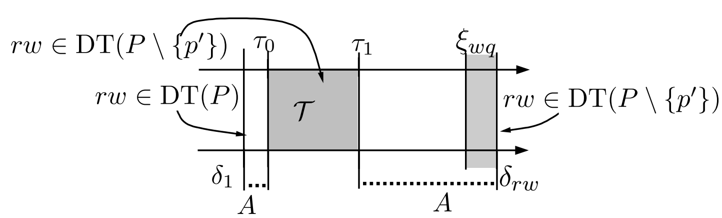

Figure 117: Case (c): Extending the almost-Delaunayhood of rw to $[\delta_1, \delta_{rw}]$. Here δ_{rw} is the first time in $[\xi_{wq}, \infty)$ when rw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{p'\})$, for some $p' \neq p, q$.

Lemma 4.1 implies, then, there is a time $\delta_{rw} \geq \xi_{wq}$ which is the first such time when the edge rw belongs to some reduced triangulation $\mathrm{DT}(P \setminus \{p'\})$, for $p' \neq p, q$. In what follows, we use p' to denote a fixed point in $P \setminus \{p, q, r, w\}$ whose removal restores the Delaunayhood at time δ_{rw}; see Figure 117.

Note that we have $\delta_{rw} > \xi_{wq} > \lambda_{pq} > \delta_1$. Since case (a) has been ruled out, the choice of δ_{rw} guarantees that, unless δ_{rw} belongs to $\mathcal{T} = [\tau_0, \tau_1]$, there exist at most $k+1$ quadruples $\varrho' \in \Sigma^B_{rw}$ whose respective (r, w)-crossings $(q'p, w, \mathcal{T}')$ end in $(\tau_1, \delta_{rw}]$.

Charging events in $\mathcal{A}_{rw}$. We are now ready to invoke Theorem 2.2 in order to extend the almost-Delaunayhood of rw, which already belongs to $\mathrm{DT}(P \setminus \{p\})$ throughout $\mathcal{T} = [\tau_0, \tau_1]$ (by Lemma 4.1), to the interval $[\delta_1, \delta_{rw}]$, which clearly intersects $\mathcal{T}$.

Note that $[\delta_1, \delta_{rw}] \setminus \mathcal{T}$ is composed of two disjoint (and possibly empty) sub-intervals $[\delta_1, \tau_0)$ and $(\tau_1, \delta_{rw}]$. We apply Theorem 2.2 separately over each of these sub-intervals (and only if they are non-empty). In both cases, we use the second threshold parameter $\ell \gg k$.

The first application of Theorem 2.2, over (δ_1, τ_0), is done with respect to the complete point set P (using the Delaunayhood of rw at time δ_1). It is necessary only if $\delta_1 < \tau_0$.

If at least one of the conditions (i), (ii) of that theorem is satisfied, we charge ϱ within $\mathcal{A}_{rw}$, via $(rp, w, \mathcal{T})$, either to an ℓ-shallow collinearity, or to $\Omega(\ell^2)$ ℓ-shallow co-circularities during (δ_1, τ_0). Since case (a) has been ruled out, $(qr, w, \mathcal{T})$ is among the first $k+1$ such (r, w)-crossings to begin after any event that we charge. Hence, the above charging accounts for at most $O(k\ell^2 N(n/\ell) + k\ell n^2 \beta(n))$ quadruples $\varrho \in \Sigma^B$. Otherwise, we end up with a subset of at most 3ℓ points (perhaps including p or q, or both) whose removal restores the Delaunayhood of rw throughout $[\delta_1, \tau_0]$.

The similar second application of Theorem 2.2 (over (τ_1, δ_{rw})) is done with respect to the reduced point set $P \setminus \{p'\}$ (where p' denotes the point in $P \setminus \{p, q\}$ whose removal restores the Delaunayhood of rw at time δ_{rw}). It is necessary only if $\tau_1 < \delta_{rw}$.

If at least one of the conditions (i), (ii) of that theorem holds, we charge ϱ (via $(rp, w, \mathcal{T})$) within $\mathcal{A}_{rw}$ either to an $(\ell+1)$-shallow collinearity, or to $\Omega(\ell^2)$ $(\ell+1)$-shallow co-circularities (which are ℓ-shallow with respect to $P \setminus \{p'\}$). By the choice of δ_{rw}, $(qr, w, \mathcal{T})$ is among the last $k+2$ such (r, w)-crossings to end after the event, so any $(\ell+1)$-shallow event in $\mathcal{A}_{rw}$ is charged by at most $O(k)$ quadruples ϱ. Otherwise, we end up with a subset of at most $3\ell + 1$ points (inclding p', and perhaps also some of p, q) whose removal restores the Delaunayhood of rw throughout $[\tau_1, \delta_{rw}]$.

To conclude, we have a subset A_{rw} of at most $6\ell + 1$ points, *including* p', and perhaps also some of p, q, whose removal restores the Delaunayhood of rw throughout $[\delta_1, \delta_{rw}]$. To obtain the crossing of rw by p or q (which would occur in, respectively, $\mathrm{DT}((P \setminus A_{rw}) \cup \{p\})$ or $\mathrm{DT}((P \setminus A_{rw}) \cup \{q\})$), it suffices to show that rw is hit by one of these two points during the interval $[\delta_1, \delta_{rw}]$. Notice that this interval contains $\delta_2 \in (\delta_1, \lambda_{pq}] \subset (\delta_1, \xi_{wq}]$. See Figure 118. To do so, we distinguish between two possible sub-scenarios, depending on the precise order type of δ_2, which is red-blue with respect to pq and rw.

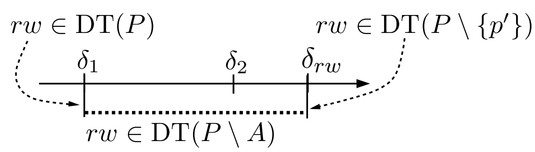

Figure 118: Case (c): The edge rw belongs to $\mathrm{DT}(P \setminus A)$ throughout the interval $[\delta_1, \delta_{rw}]$, which contains the last co-circularity δ_2 of p, q, r, w. In addition, rw belongs to $\mathrm{DT}(P)$ and $\mathrm{DT}(P \setminus \{p'\})$ at times δ_1 and δ_{rw}, respectively.

If the Delaunayhood of rw is violated right *after* δ_2 by $p \in L_{rw}^-$ and $q \in L_{rw}^+$, then, since δ_2 is the *last* co-circularity of p, q, r, w, Lemma 3.1 implies that rw is hit during $(\delta_2, \delta_{rw}]$ by at least one of p, q (because $p' \neq p, q, r, w$), so we are done. (See Figure 110 (left).)

Assume, then, that the Delaunayhood of rw is violated right *before* δ_1 by p and q. Notice that this violation does not hold at time δ_1. Hence, we can obtain the desired crossing of rw in (δ_1, δ_2) by applying the time-reversed variant of Lemma 3.1 (for the point set $P = \{p, q, r, w\}$, backwards from δ_2). The crucial observation is that δ_1 and δ_2 have different order types, which rules out the last case in Lemma 3.1. (See Figure 110 (right).)

If rw is hit during $(\delta_1, \delta_{rw}]$ by the point p, then, together with the crossing of pw by r (enforced in Part 1 by omitting $A_{pw} \setminus \{r\}$, where the subset A_{pw} was obtained by applying Theorem 2.2 in $\mathcal{A}_{pw}$), the triple p, r, w now performs two Delaunay crosings within the triangulation $\mathrm{DT}\left((P \setminus (A_{rw} \cup A_{pw})) \cup \{p, r\}\right)$.

Otherwise, if rw is hit during $(\delta_1, \delta_{rw}]$ by q, the other triple q, r, w performs two Delaunay crossings within the triangulation $\mathrm{DT}((P \setminus (A_{rw} \cup \{p\})) \cup \{q\})$, namely, the crossing of qr by w (prescribed by condition (B3)), and the just obtained crossing of rw by q.

In both cases, a standard combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor implies that the overall number of the corresponding triples (p, r, w) or (q, r, w) in P cannot exceed $O(\ell n^2)$. Since the quadruple $\varrho = (p, q, r, w)$ at hand is uniquely determined by each of its four sub-triples, this also bounds the overall number of such quadruples in Σ^B.

To conclude, we have established the following bound on the maximum possible cardinality of Σ^B:

$$T^B(n) = O\left(k\ell^2 N(n/\ell) + k^2 N(n/k) + k\ell n^2 \beta(n)\right). \tag{13}$$

That is, we have expressed the maximum possible number of terminal quadruples of type B in terms of more elementary quantities which were introduced in Section 2. Informally, here the system of our recurrences bottoms out, in the sense that no new quantities appear in the righ-hand side.

7.3 Terminal quadruples of types C and D

We next establish near-quadratic recurrences for the maximum possible numbers $T^C(n)$ and $T^D(n)$ of terminal quadruples of types C and D, respectively, that can arise in an underlying set P of n moving points. See Section 6.6 for precise definitions of these two types of configurations.

Let $\varrho = (w, q, u, p)$ be a terminal quadruple of type C or D. Notice that each of the (unordered) triples u, p, w and u, p, q is involved in a Delaunay crossing (see Figure 119).

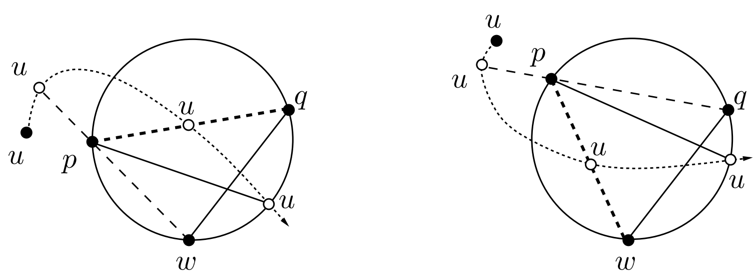

Figure 119: Possible trajectories of u in a terminal quadruple $\varrho = (w, q, u, p)$ of type C or D (resp., left and right). In both types, each of the unordered triples p, u, q and p, u, w is involved in a Delaunay crossing.

Specifically, if ϱ is of type C, we have a Delaunay crossing of wu (or uw) by p in $P \setminus \{q\}$, and a Delaunay crossing of wq (or of qw) by u in $P \setminus \{p\}$. Similarly, if ϱ is of type D, we have a Delaunay crossing of qu (or of uq) by p in $P \setminus \{p\}$. and a Delaunay crossing of wp (or of pw) by u in $P \setminus \{q\}$.

In both types, the four points of ϱ are involved in a Delaunay co-circularity, right after which the Delaunayhood of pu is violated by $w \in L^-_{pu}$ and $q \in L^+_{pu}$, and this is the last co-circularity of w, q, u, p. We will use the above co-circularity to enforce a Delaunay crossing of up by at least one of w, q. As a result, one of the triples u, p, w and u, p, q will perform two single Delaunay crossings in a suitably refined subset of P, so our analysis will bottom out via Lemma 4.5.

The desired crossing of up can be enforced using exactly the same analysis as was used in Section 3 to express the maximum possible number $N_E(n)$ of extremal Delaunay co-circularities in P in terms of the maximum possible number $C(n/k)$ of Delaunay crossings that can arise in a subset (of P) of cardinality n/k. Nevertheless, we briefly review the argument of Section 3 for the sake of completeness.

Let t_0 denote the time of the above extremal Delaunay co-circularity of w, q, u, p. If the edge pu never re-enters $\mathrm{DT}(P)$ (leaving $\mathrm{DT}(P)$ at time t_0), then we can charge ϱ to this last disappearance of pu from $\mathrm{DT}(P)$, which occurs for at most $O(n^2)$ terminal quadruples ϱ under consideration. Otherwise, let t_1 be the first time after t_0 when up re-enters $\mathrm{DT}(P)$. By Lemma 3.1, pu is hit in $(t_0, t_1]$ by at least one of w, q. Namely, either q crosses up from L^-_{pu} to L^+_{pu}, or w crosses pu in the opposite direction. Furthermore, this is the second and last collinearity of p, q, u or p, w, u (and, therefore, the *only* such collinearity of this order type to occur in $(t_0, t_1]$).

In both cases, we invoke Theorem 2.2 to amplify the above second collinearity of p, u, q or p, u, w into an additional Delaunay crossing. Specifically, we fix a constant threshold $k > 12$ and apply Theorem 2.2 in $\mathcal{A}_{pu}$ over the interval (t_0, t_1).

In cases (i) and (ii) of Theorem 2.2, we can charge $\varrho = (w, q, p, u)$ within $\mathcal{A}_{pu}$ either $\Omega(k^2)$ k-shallow co-circularities, or a k-shallow collinearity. Furthermore, each shallow event is charged at most $O(1)$ times, because it involves p and u, and t_0 is the last disappearance of pu from $\mathrm{DT}(P)$. Hence, the overall number of such terminal quadruples does not exceed $O\left(k^2 N(n/k) + kn^2\beta(n)\right)$.

Finally, in case (iii) of Theorem 2.2, we end up with a subset A of at most $3k$ points so that pu belongs to $\mathrm{DT}(P \setminus A)$ throughout $[t_0, t_1]$. Thus, either pu undergoes a single Delaunay crossing by q in $(P \setminus A) \cup \{q\}$, or its reversed copy up undergoes a single Delaunay crossing by w in $(P \setminus A) \cup \{w\}$.

Therefore, we can charge ϱ to the corresponding triple p, q, u or p, w, u which performs two Delaunay crossings in a suitable subset of P. Lemma 4.5 together with the Clarkson-Shor argument imply that the overall number of such triples in P cannot exceed $O(kn^2)$. Furthermore, each of them can be charged at most once, because t_0 is the last time when pu disappears from $\mathrm{DT}(P)$ before being hit as above by q or w.

To conclude, we have established the following recurrences for the above quantities $T^C(n)$ and $T^D(n)$:

$$T^C(n) = O\left(k^2 N(n/k) + kn^2\beta(n)\right) \tag{14}$$

and

$$T^D(n) = O\left(k^2 N(n/k) + kn^2\beta(n)\right). \tag{15}$$

8 Proof of Theorem 5.3

Let $(pq, r, I = [t_0, t_1])$ be a (p, r, k)-chargeable Delaunay crossing, and let $\mathcal{I} = [t_2, t_3]$ be the corresponding interval which certifies the (p, r, k)-chargeability of (pq, r, I). In particular, at least k counterclockwise (q, r)-crossings (uq, r, I_u) occur within $\mathcal{I}$ (in the sense that $I_u \subseteq \mathcal{I}$). In additon, the edge pr belongs to $\mathrm{DT}(P)$ when $\mathcal{I}$ begins or ends, and there is a subset $A_0 \subset P$ of $c_0 = O(1)$ points whose removal restores the Delaunayhood of pr throughout $\mathcal{I}$.

By Lemma 4.6, each of the above (q, r)-crossings (uq, r, I_u) occurs within one of the intervals $\mathcal{I}^+ = (t_0, t_3]$ or $\mathcal{I}^- = [t_2, t_1)$. In particular, we have $I_u \subseteq (t_0, t_3]$ if and only if r enters L^+_{uq} after entering L^+_{pq}; see Figure 120. Without loss of generality, we assume that at least $\lceil k/2 \rceil$ of these crossings occur within $(t_0, t_3]$. Again, Lemma 4.6 implies that each such crossing must end within $(t_0, t_3]$.

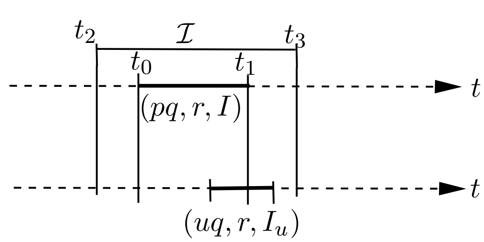

Figure 120: The setup in the proof of Theorem 5.3. The crossing $(pq, r, I = [t_0, t_1])$ is (p, r, k)-chargeable, for $\mathcal{I} = [t_2, t_3]$. We fix a counterclockwise (q, r)-crossing (uq, r, I_u), which ends in $(t_1, t_3]$ (so $I_u \subseteq (t_0, t_3]$). The (q, r)-crossings (pq, r, I) and (uq, r, I_u) form a counterclockwise quadruple (q, p, u, r).

Overview. To establish Theorem 5.3, we distribute the "weight" of (pq, r, I) over the above $\Omega(k)$ (q, r)-crossings (uq, r, I_u) or, more precisely, over their respective arrangements $\mathcal{A}_{ur}$. (Recall that each counterclockwise (q, r)-crossing (uq, r, I_u) is also a clockwise (u, r)-crossing.) In what follows, we fix one of the first $\lceil k/2 \rceil$ counterclockwise (q, r)-crossings that ends after time t_1 (and before t_3), and assume that its respective point u does not belong to A_0. Our charging strategy is to make each such u pay $\Theta(1/k)$ units of charge to (pq, r, I), so that (pq, r, I) receives a total of at least 1 unit. The charging will be performed in one of two possible ways (depending on the structure of $\mathcal{A}_{ur}$ and on the motion of p, q, u, and r).

We shall first try to charge (uq, r, I_u) (rather than (pq, r, I)) to events within $\mathcal{A}_{ur}$ using the standard techniques of Section 5 (involving Lemma 4.5 and Theorem 2.2). In case of success, (uq, r, I_u) will be declared as *heavy* for (pq, r, I) and will pay $\Theta(1/k)$ units of charge to (pq, r, I). As we will show, the overall number of such crossings (uq, r, I_u), that will be declared as heavy for *at least one* of their neighboring (q, r)-crossings, does not exceed $O(k^2 N(n/k) + kn^2\beta(n))$. Moreover, any crossing (uq, r, I_u) will be charged (as heavy) by at most $\lceil k/2 \rceil$ neighboring (q, r)-crossings (pq, r, I), due to the $\lceil k/2 \rceil$-proximity of the crossings (pq, r, I) and (uq, r, I_u), and will pay $\Theta(1/k)$ units of charge to each. Therefore, at most $O(k^2 N(n/k) + kn^2\beta(n))$ units of charge will be transferred in this fashion.

If the above strategy fails, we shall resort to a more subtle type of charging. In that case, we shall charge (pq, r, I) (again within $\mathcal{A}_{ur}$) to $\Theta(k)$ $(4k)$-shallow co-circularities that involve u, r and p (together with some fourth point, not necessarily q), and each of these co-circularities will pay $\Theta(1/k^2)$ units of charge to (pq, r, I). Moreover, we shall argue that each $(4k)$-shallow co-circularity can be charged in this latter manner by at most $O(1)$ crossings (pq, r, I). Hence, at most $O(k^2 N(n/k))$ units will be transferred in the second scheme. The theorem then follows from these two charging schemes.

Before proceeding with the above general strategy, we fix one such crossing (uq, r, I_u) and establish several essential properties of it.

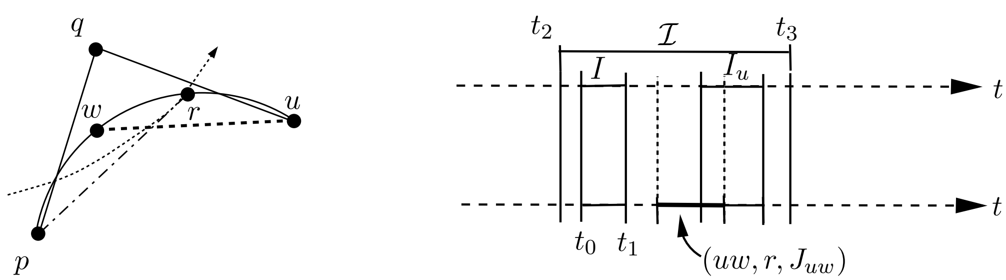

Figure 121: Proof of Proposition 8.1. Assuming $p \neq w$, the four points w, u, r, p are involved in a red-blue co-circularity during the crossing (uw, r, J_{uw}). Since the Delaunayhood of pr is then violated by w and u, and u is chosen outside A_0, the set A_0 must contain w.

Proposition 8.1. *With the above assumptions, and with (uq, r, I_u) fixed, at most $c_0 + 1$ clockwise (u, r)-crossings (uw, r, J_{uw}) occur within $[t_0, t_3]$.*

Proof. Fix a clockwise (u, r)-crossing (uw, r, J_{uw}), such that $w \neq p$ and $J_{uw} \subset [t_0, t_3]$. Refer to Figure 121.

By Lemma 4.4, the points w, u, r, p are involved during J_{uw} in a co-circularity which is red-blue with respect to the edges uw and pr. Hence, the Delaunayhood of pr is violated by u and w either right before or right after this co-circularity. Since $[t_0, t_3] \subseteq \mathcal{I}$, the set A_0 must include at least one of the points u, w. Since, by assumption, $u \notin A_0$, we must have $w \in A_0$, so there can be at most c_0 such crossings. Adding the possible crossing (up, r, J_{up}) yields the asserted bound. $\square$

Notice that the set P induces a counterclockwise quadruple $\sigma_u = (q, p, u, r)$ whose respective interval $[I, I_u]$ is contained in $[t_0, t_3]$. The following proposition is stated in full generality and applies to *all* counterclockwise quadruples (i.e., not necessarily the ones that arise in the course of the present proof of Theorem 5.3). It can be viewed simply as an extension of Lemma 5.1.

Proposition 8.2. *Let $\sigma_u = (q,p,u,r)$ be a counterclockwise quadruple, with associated crossings (pq,r,I) and (uq,r,I_u). Suppose that the edge rq is hit by the point p, and that this happens in the interval after r enters L^+_{pq} and before r enters L^+_{uq}. Then pr is also hit, during that same interval, by the point u.*

Remarks. (1) Clearly, a symmetric statement holds if rq is hit by u. Namely, in that case the edge ru is hit by the point p. As a matter of fact, the proof of Proposition 8.2 implies that the two scenarios coincide: The edge rq is hit by p between the times when r crosses pq and uq if and only if rq is hit there by u too.

(2) The reader might be tempted to use Lemma 4.5 in order to bound the number of such crossings (uq,r,I_u), whose respective counterclockwise quadruples $\sigma_u = (q,p,u,r)$ satisfy the conditions of Proposition 8.2 (as was done, e.g., for clockwise Delaunay quadruples in case (a) of Section 5.3). However, since we do not assume the edge rq to be almost-Delaunay during $[I,I_u]$, the argument of Section 5.3 does not immediately apply to such instances.

Proof. Refer to Figure 122. Notice that, according to Lemma 4.1, p can hit rq (as prescribed in the proposition) only during the gap between the intervals I and I_u of the two (q,r)-crossings of σ_u (a gap that we therefore assume to exist).

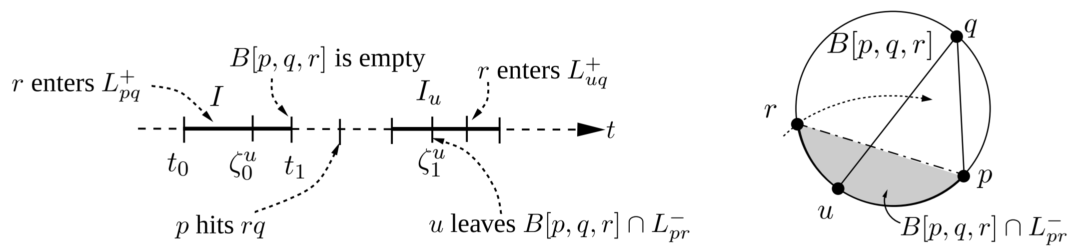

Figure 122: Proof of Proposition 8.2. Left: The summary of events that are assumed to occur during $[I,I_u]$. Right: The point u leaves the cap $B[p,q,r] \cap L^-_{pr}$ at time $\zeta^u_1 \in I_u \setminus I = I_u$.

Since the points p,q,r can be collinear at most twice, the halfplane L^+_{uq} contains p when r enters it during I_u. Therefore, and according to Lemma 4.4, the four points p,q,u,r are involved at some time $\zeta^u_1 \in I_u \setminus I = I_u$ in a co-circularity, occurring before r crosses uq; see Figure 122 (right). Right after this co-circularity the Delaunayhood of uq is violated by $r \in L^-_{uq}$ and $p \in L^+_{uq}$. Note that at that very moment the point u leaves the cap $B[p,q,r] \cap L^-_{pr}$. Note also that, according to Lemma 4.4, the points p,q,u,r are also involved in an earlier co-circularity which occurs at some time $\zeta^u_0 \in I \setminus I_u = I$ (and before p hits rq, which occurs between I and I_u). We distinguish between the following two scenarios.
(i) If u lies in L^-_{pq} when p hits rq (and r re-enters L^-_{pq}), then u lies within the cap $B[p,q,r] \cap L^-_{rq}$ right after this collinearity, as depicted in Figure 123 (top-left). Right after this event and before ζ^u_1, u must move from this cap to the disjoint cap $B[p,q,r] \cap L^-_{pr}$ (which it exits at time ζ^u_1) either[46] through pr (and through rq) or through the boundary of $B[p,q,r]$. See Figure 123 (top-right). However, in the latter case u would first have to leave its present cap through $\partial B[p,q,r]$, so the points p,q,u,r would be co-circular at least twice during (ζ^u_0, ζ^u_1), contradicting the assumption that any four points are co-circular at most three times. Hence, u can enter $B[p,q,r] \cap L^-_{pr}$ only through pr and rq, as claimed in the proposition.
(ii) If u lies in L^+_{pq} when p hits rq, then u lies within the disc $B[p,q,r]$ right before this event; see Figure 123 (bottom-left). By the definition of Delaunay crossings, the disc $B[p,q,r]$ contains no points of P

[46]Here we implicitly rely on the fact that the motion of $B[p,q,r]$ is continuous after the second collinearity of p,q,r.

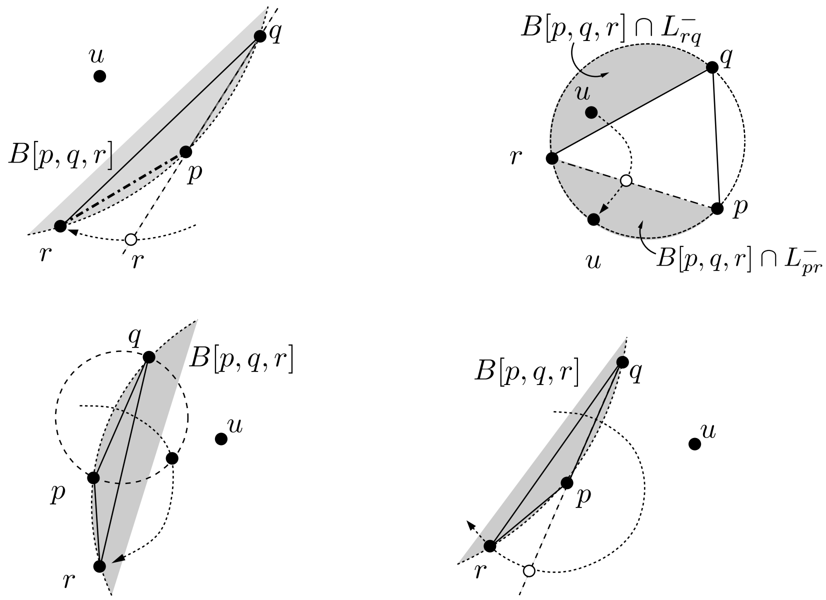

Figure 123: Proof of Proposition 8.2. Top: If u lies in L_{pq}^- when p hits rq (top-left), then u can exit the cap $B[p,q,r] \cap L_{pr}^-$ only after crossing pr (top-right). Bottom: The hypothetic scenario where u lies in L_{pq}^+ when p hits rq. Right before p hits rq, the disc $B[p,q,r]$ contains u, which must have entered it after I (bottom-left). Right after that collinearity, u lies outside $B[p,q,r]$, so it will have to re-enter $B[p,q,r]$ before ζ_1^u (bottom-right).

right after the end of I, as depicted in Figure 123 (bottom-right). Hence, u enters $B[p,q,r]$ at the end of I and before p hits rq. We also note that u lies outside $B[p,q,r]$ right after the second collinearity of p,q,r, so u must enter $B[p,q,r]$ (through its boundary) afterwards and before ζ_1^u (in order to exit it after ζ_1^u). Similar to the preceding scenario, we obtain four impossible co-circularities of p,q,u,r, showing that the present scenario cannot occur. $\square$

Back to the proof of Theorem 5.3. With these preparations, we are finally ready to establish Theorem 5.3. Recall that we have fixed a counterclockwise (q,r)-crossing (uq,r,I_u) that ends in $(t_1,t_3]$, and which is among the first $\lceil k/2 \rceil$ such (q,r)-crossings to end after t_1. Recall also that u does not belong to the set A_0 (of size c_0, appearing in the definition of the (p,r,k)-chargeability of (pq,r,I)), and that the (q,r)-crossings (pq,r,I) and (uq,r,I_u) form a (not necessarily consecutive) counterclockwise (q,r)-quadruple $\sigma_u = (q,p,u,r)$.

We first claim that r cannot cross pq again between the times when it enters the halfplanes L_{pq}^+ and L_{uq}^+ (during the two respective Delaunay crossings). Indeed, otherwise a counterclockwise variant of Lemma 5.1 would imply that the edge pr is hit by u during the interval $[I,I_u]$. As the latter interval is contained in $[t_0,t_3]$, this is a clear contradiction to the assumed choice of u outside A_0. Similarly, p cannot hit rq between the times when r enters the halfplanes L_{pq}^+, L_{uq}^+, for otherwise we would invoke Proposition 8.2 to show that pr is again hit by u during $[I,I_u] \subseteq [t_0,t_3]$, and reach the same contradiction as above.

If the edge pr is hit during $[t_1,t_3]$ by q (which is the only remaining way in which p,q,r can be collinear again), then the set $(P \setminus A_0) \cup \{q\}$ induces a Delaunay crossing of pq by r, and a Delaunay crossing of pr by q. A routine combination of Lemma 4.5 with the probabilistic argument of Clarkson and Shor shows that this scenario happens for at most $O(n^2)$ Delaunay crossings (pq,r,I).

It therefore suffices to focus on the scenarios where r does not re-enter L_{pq}^- after I and before it enters L_{uq}^+ (through uq, during I_u). As noted in Section 5.1 (see also the proof of Proposition 8.2), the four points q,p,u,r are involved in co-circularities at some times $\zeta_0^u \in I \setminus I_u$ and $\zeta_1^u \in I_u \setminus I$; see Figure

124. Moreover, these are the only co-circularities of p, q, u, r to occur during I and I_u.

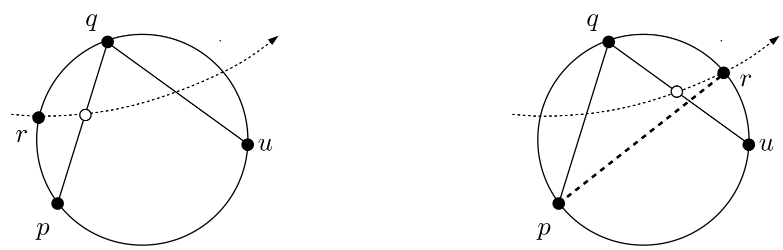

Figure 124: The two co-circularities of q, p, u, r which occur at times $\zeta_0^u \in I \setminus I_u$ (left) and $\zeta_1^u \in I_u \setminus I$ (right).

Consider the latter co-circularity, occurring at some time $\zeta_1^u \in I_u \setminus I$, which is red-blue with respect to the edges pr, uq. Since r does not return to L_{pq}^-, p lies in L_{uq}^- when r hits uq during I_u. (See Figure 124 (right).) Arguing as in Section 4.1 (see, e.g., the proofs of Lemmas 4.4 and 4.6), we can conclude that the Delaunayhood of pr is violated right after time ζ_1^u by the points u and q.

We first argue that the above co-circularity at time ζ_1^u cannot be the *last* co-circularity of q, p, u, r. Indeed, otherwise Lemma 3.1 (combined with the assumption that q does not hit pr during $[t_1, t_3]$) would imply that the edge pr is hit by u during the interval (ζ_1^u, t_3). However, in that case u would belong to A_0, contrary to the choice of u.

To conclude, we can assume, from now on, that the co-circularity at time ζ_1^u is the *middle* co-circularity of the points q, p, u, r. Hence, the preceding co-circularity, which occurs at time $\zeta_0^u \in I \setminus I_u$, must be the *first* co-circularity of these four points.

To proceed, we distinguish between several topological scenarios, treating each in turn. In each of them, (pq, r, I) receives $\Theta(1/k)$ units of charge (via u alone, as reviewed in the beginning of this section). Recall that, with (pq, r, I) fixed, u and (uq, r, I_u) can be chosen in $\Theta(k)$ possible ways. Hence, with an appropriate choice of the constants of proportionality, each (p, r, k)-chargeable crossing (pq, r, I) will eventually receive at least one unit of charge.

Case (a). The edge ru is never Delaunay during $(-\infty, t_0]$. In this case, we classify the crossing (uq, r, I_u) as *heavy* (for (pq, r, I)), and we make it pay $\Theta(1/k)$ units of charge to (pq, r, I).

Notice that (uq, r, I_u) is one of the first $c_0 + 2$ clockwise (u, r)-crossings (according to the standard order provided by Lemma 4.6). Indeed, by Lemma 4.1, no such crossings begin before time t_0, when the edge ru is not even Delaunay. In addition, by Proposition 8.1, at most $c_0 + 1$ clockwise (u, r)-crossings can begin after t_0 and before the beginning of (uq, r, I_u), as each of them has to occur within the interval $[t_0, t_3]$. In conclusion, the overall number of such crossings (uq, r, I_u), that are classified as heavy for at least one of their neighboring (q, r)-crossings (pq, r, I) (upon falling into case (a)), is at most $O(n^2)$.

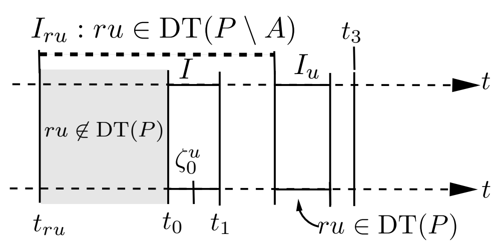

Figure 125: Preparing for cases (b), (c), and (d): We pick the last time t_{ru} in $(-\infty, t_0]$ when ru is Delaunay and apply Theorem 2.2 over the interval I_{ru} (containing ζ_0^u).

Preparing for cases (b), (c) and (d). In each of the subsequent three cases, we assume that ru appeared in $\mathrm{DT}(P)$ also before (or at) t_0. Let t_{ru} be the last time in $(-\infty, t_0]$ when ru belongs to $\mathrm{DT}(P)$, and let I_{ru} denote the subsequent interval that lasts from t_{ru} to the beginning of I_u. Note that I_{ru} contains $I \setminus I_u$, and therefore includes the time ζ_0^u of the first co-circularity of p, q, u, r. Refer to Figure 125.

As a preparation, we apply Theorem 2.2 in $\mathcal{A}_{ru}$ over I_{ru} (with the same constant parameter k, and keeping in mind that ru is Delaunay at both endpoints of I_{ru}), and then proceed depending on the outcome.

Case (b). If one of the Conditions (i), (ii) of Theorem 2.2 is satisfied (i.e., $\mathcal{A}_{ru}$ contains either $\Omega(k^2)$ k-shallow co-circularities or a k-shallow collinearity, all of them occurring in I_{ru}), the crossing (uq, r, I_u) is again classified as heavy for (pq, r, I), and pays $\Theta(1/k)$ units of charge to it.

We claim that the overall number of such crossings (uq, r, I_u), that are classified as heavy for at least one of their neighbors (pq, r, I) (within the present case (b)), is at most $O(k^2 N(n/k) + kn^2\beta(k))$. To show this, we keep the crossing (pq, r, I) fixed and charge (uq, r, I_u) within $\mathcal{A}_{ru}$ either to $\Omega(k^2)$ k-shallow co-circularities, or to a k-shallow collinearity, which are assumed to occur during the respective interval I_{ru}.

We emphasize that the first endpoint t_{ru} of I_{ru} might depend on the choice of (pq, r, I) from among those crossings that expect to receive $\Theta(1/k)$ units from (uq, r, I_u). Furthermore, an event in $\mathcal{A}_{ru}$ might be charged by the same (uq, r, I_u) in the context of several (p, r, k)-chargeable crossings (pq, r, I) that charge (uq, r, I_u) (for various values of p). Nevertheless, for each choice of an event in $\mathcal{A}_{ru}$ and each clockwise (u, r)-crossing (uq, r, I_u), all such episodes cause only one charging of this event by (uq, r, I).

We next show that each event in $\mathcal{A}_{ru}$ is charged in the above manner by at most $O(1)$ crossings (uq, r, I_u). Indeed, let t^* be the time of a k-shallow event that we charge within $\mathcal{A}_{ru}$. Clearly, one can guess the points u and r of (uq, r, I_u) in at most $O(1)$ ways, as they are involved in the event. Thus, it suffices to guess the third point q of (uq, r, I_u) (armed only with the knowledge of t^*, r and u), which is done as follows.

Let q be a potential third point, and let (pq, r, I) be *any* (p, r, k)-chargeable crossing that receives $\Theta(1/k)$ units of charge from the corresponding crossing (uq, r, I_u) (after the latter crossing is classified as heavy for (pq, r, I), by the rule of case (b)). By Lemma 4.1, no clockwise (u, r)-crossing (uq, r, I_u) can begin during the respective interval $I_{ru} \cap (-\infty, t_0]$ (when ru is not even Delaunay). Moreover, Proposition 8.1 implies that at most $c_0 + 1$ such (u, r)-crossings begin in the interval that lasts from t_0 to the beginning of I_u (which is contained in $[t_0, t_3]$). Hence, (uq, r, I_u) is among the first $c_0 + 2$ clockwise (u, r)-crossings to begin after t^*, so knowing t^*, r, and u enables us to guess (uq, r, I_u) in at most $O(1)$ ways (irrespective of the choice of p and (pq, r, I)).

The number of k-shallow co-circularities in $\mathcal{A}_{ru}$, over all r, u, is at most $O(k^4 N(n/k))$. Similarly, the number of k-shallow collinearities is $O(kn^2\beta(n))$. Each such event is charged by only $O(1)$ (u, r)-crossings (uq, r, I_u) (which are declared as heavy in case (b), for at least one of their (p, r, k)-chargeable neighbors (pq, r, I)). Furthermore, each such crossing (uq, r, I_u) charges either $\Omega(k^2)$ k-shallow co-circularities, or a k-shallow collinearity. All these considerations imply that the number of charging crossings (uq, r, I_u) of this kind is $O(k^2 N(n/k) + kn^2\beta(n))$, as claimed.

Recall that, in the rest of the analysis, each of these (u, r)-crossings will pay $\Theta(1/k)$ units of charge to $O(k)$ "neighboring" crossings (pq, r, I), so these latter crossings will recieve in total $O(k^2 N(n/k) + kn^2\beta(n))$ units of charge in this manner.

Preparing for cases (c) and (d). Now suppose that Condition (iii) of Theorem 2.2 holds. That is, the Delaunayhood of ru can be restored throughout I_{ru} by removing a subset A of cardinality at most $3k$. To handle this more difficult scenario, we first establish the following proposition.

Proposition 8.3. *With the above assumptions, the edge ru is hit during I_{ru} by at least one of the points p, q.*

Proof. The proof proceeds (essentially) along the same lines as in case (e) of Section 5.3. (The main difference is that the quadruple σ_u under consideration is *counterclockwise*.)

We first get rid of the instances where r crosses L_{uq} between the times when it enters the halfplanes L_{pq}^+ and L_{uq}^+ (in the respective intervals I and I_u). Note that if ru is hit there by q then we are done (as it can happen only during the gap between I and I_u, which is obviously covered by I_{ru}).

If rq is hit by u, then a symmetric version of Proposition 8.2 (see Remark (1) following the proposition), in which we switch the roles of p and u and reverse the direction of the time axis, implies that p hits uq between the times when r enters the halfplanes L_{pq}^+ and L_{uq}^+ (during the respective intervals I and I_u). In particular, this latter collinearity of u, r, p occurs after $t_0 > t_{ru}$ and before I_u, and, therefore, also during I_{ru}. (As previously noted, this scenario is not only symmetric to the one assumed in Proposition 8.2, but, in fact, coincides with it.)

Finally, if r hits uq, then a counterclockwise and time-reversed variant of Lemma 5.1 similarly implies that ru is hit during I_{ru} by p; see Figure 126 (left). (As in the previous case, this collinearity occurs during $[I, I_u]$, between the times when r enters the halfplanes L_{pq}^+ and L_{uq}^+.)

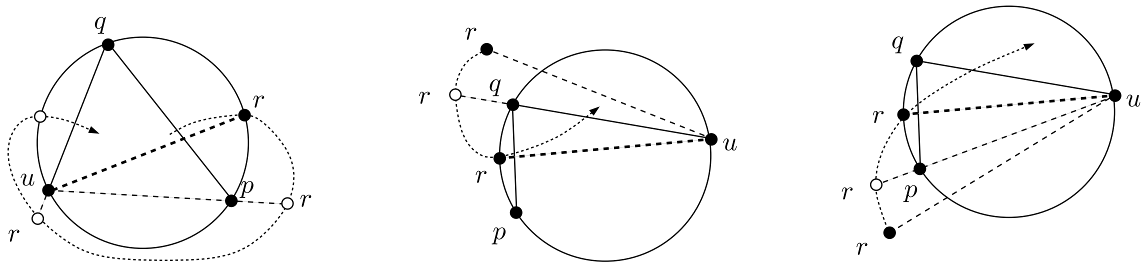

Figure 126: Proof of Proposition 8.3: Arguing that ru is hit, during I_{ru}, by at least one of p or q. Left: The edge rq is hit by u between the times when r crosses pq and uq. Hence, the asserted crossing of ur by p follows from Proposition 8.2. Center and right: The point r remains in L_{uq}^- after entering L_{pq}^+ and till the beginning of I_u. The Delaunayhood of ru is violated, right before ζ_0^u, by p and q, so the asserted collinearity follows from Lemma 3.1.

Let us then assume that r remains in L_{uq}^- between the times when it enters the halfplanes L_{pq}^+ and L_{uq}^+. In particular, u lies in L_{pq}^+ when r enters this halfplane, so the Delaunayhood of ru is violated, right before time ζ_0^u, by the points p and q, as depicted in Figure 126 (center and right). By (a time-reversal version of) Lemma 3.1, and since the co-circularity at time ζ_0^u is the *first* co-circularity of q, p, u, r, the edge ru is hit during I_{ru}, and before ζ_0^u, by at least one of the points p, q. Hence, the proposition holds also in this last remaining scenario. $\square$

Case (c). If ru is hit by q during I_{ru} then the triple q, u, r defines two single Delaunay crossings within the triangulation $\mathrm{DT}((P \setminus A) \cup \{q\})$. In this case, the crossing (uq, r, I_u) is again declared as a heavy and pays $\Theta(1/k)$ units of charge to (pq, r, I). A combination of Lemma 4.5 with the standard probabilistic argument of Clarkson and Shor yields an upper bound of $O(kn^2)$ on the overall number of such crossings (uq, r, I_u), that are declared as heavy for at least one choice of (pq, r, I) (upon falling into case (c)).

Case (d). We can, therefore, assume that ru is hit during I_{ru} by p, so the reduced set $(P \setminus A) \cup \{p\}$ induces at least one Delaunay crossing of ru by p. In this case, we say that the crossing (uq, r, I_u) is *light* for (pq, r, I), and distinguish between the following two subcases.

Case (d1). If at least one of the collinearities of u, r, p that occur during I_{ru} is $(4k)$-shallow, we directly charge (pq, r, I) to it. In other words, in this case (pq, r, I) receives 1 unit of charge via u alone, and it does not have to charge any other neighboring (q, r)-crossings.

We next argue that each $(4k)$-shallow collinearity, which occurs at some time t^*, is charged in the above manner by at most $O(1)$ (p, r, k)-chargeable crossings (pq, r, I). Indeed, the points p and r of (pq, r, I) can be guessed in $O(1)$ possible ways from among the three points involved in the charged collinearity, and their choice immediately determines the third point u (which figures in the charging scenario of case (d1)). The guessing of q, which is the last unknown point of (pq, r, I), is done exactly as in case (b), and it requires only the knowledge of t^*, r and u. (As before, we use the property that (uq, r, I_u) is among the first $c_0 + 2$ such clockwise (u, r)-crossings to begin after t^*.) To conclude, the above charging accounts for at most $O(kn^2\beta(n))$ crossings (pq, r, I).

Case (d2). It thus remains to handle the scenario where all collinearities of u, p, r that occur during I_{ru} are $(4k)$-deep.

We first argue that $\mathcal{A}_{ru}$ contains at least k $(4k)$-shallow co-circularities, each occurring within the previously defined interval I_{ru} and involving p, u, r and some fourth point of P. Indeed, the open disc $B[u, r, p]$ contains no points of $P \setminus A$ when the above crossing of ru by p begins, within the reduced triangulation $\mathrm{DT}((P\setminus A)\cup\{p\})$. (If ru undergoes more than one crossing by p within $\mathrm{DT}((P\setminus A)\cup\{p\})$, we consider the first such crossing.) Since the corresponding collinearity of u, r, p is not $(4k)$-shallow (and the cardinality of A is at most $3k$), the disc $B[u, r, p]$ "swallows" at least k points of $P \setminus A$ before ru is hit by p, which can enter $B[u, r, p]$ only through its boundary. Since at the beginning of the process $B[u, r, p]$ contains only (at most $3k$) points of A, the first k points that $B[u, r, p]$ "swallows" form with u, r and p k co-circularities, all of which are $(4k)$-shallow.

Each of the above $(4k)$-shallow co-circularities pays $\Theta(1/k^2)$ units of charge to (pq, r, I). Therefore, (pq, r, I) still receives at least $\Theta(1/k)$ units of charge via (uq, r, I_u). To complete our analysis, we argue, almost exactly as in the previous case (d1), that each $(4k)$-shallow co-circularity, which occurs at some fixed time t^*, is charged in this manner by at most $O(1)$ crossings (pq, r, I). Indeed, the points p, r and u can be chosen in at most $O(1)$ possible ways from among the four points that are co-circular at time t^*. Moreover, the knowledge of t^*, r and u enables us to guess the last unknown point q of (pq, r, I) in at most $c_0 + 2$ possible ways, as was done in cases (b) and (d1).

To conclude, in case (d2) the crossing (pq, r, I) receives a total $\Theta(1/k)$ units of charge from $\Theta(k)$ $(4k)$-shallow co-circularities within $\mathcal{A}_{ru}$ (each involving p, r and u), where each co-circularity is charged by at most $O(1)$ crossings.

Wrap up. To finish the proof of Theorem 5.3, it remains to check that all the (p, r, k)-chargeable crossings (pq, r, I) (over all possible $p, r \in P$) receive a total of at most $O(k^2N(n/k) + kn^2\beta(n))$ units of charge from neighboring heavy (q, r)-crossings (uq, r, I_u) and from $(4k)$-shallow collinearities and co-circularities in appropriate arrangements $\mathcal{A}_{ru}$.

Indeed, the overall number of crossings (uq, r, I_u) that are classified as heavy (upon falling into one of the cases (a)–(c)), for at least one of their neighbors (pq, r, I), is at most $O(k^2N(n/k) + kn^2\beta(n))$. Moreover, a heavy crossing (uq, r, I_u) pays $\Theta(1/k)$ units of charge to (pq, r, I) only if these crossings are $\lceil k/2\rceil$-consecutive (as (q, r)-crossings), so it pays at most $O(1)$ units of charge in total.

Furthermore, we have shown that any $(4k)$-shallow co-circularity or collinearity is charged, through the mechanism of case (d), by $O(1)$ crossings (pq, r, I). Namely, in case (d1) each $(4k)$-shallow collinearity pays 1 unit of charge to each of the $O(1)$ possible charging crossing (pq, r, I), so the total charge paid by these collinearities is $O(kn^2\beta(n))$. In contrast, in case (d2) each $(4k)$-shallow co-circularity pays each time only $\Theta(1/k^2)$ units of charge, so the total charge paid by these co-circularities is $O\left(\frac{1}{k^2}k^4N(n/k)\right) = O\left(k^2N(n/k)\right)$.

Finally, each (p, r, k)-chargeable crossing (pq, r, I) charges $\lceil k/2\rceil$ neighboring (q, r)-crossings (uq, r, I_u). Except for case (d1), where (pq, r, I) receives via (uq, r, I_u) one unit of charge (from a $(4k)$-shallow collinearity of u, r and p), (pq, r, I) recieves each time $\Theta(1/k)$ units of charge, either directly from (uq, r, I_u) (when that last crossing is heavy), or from certain $(4k)$-shallow events within the corresponding arrangement $\mathcal{A}_{ru}$ (when (uq, r, I_u) is light). In either case, (pq, r, I) receives one at least one unit of charge, and the proof of Theorem 5.3 is now complete. $\square$

Remark. It is instructive to compare the arguments used in cases (b) and (d) of the above analysis. Notice that both of them proceed by charging events that occur in $\mathcal{A}_{ru}$ during I_{ru}.

In case (b), each k-shallow event under consideration is only known to involve r and u (but not necessarily p or q). This information appears to be sufficient for guessing q and (uq, r, I_u), but not necessarily p and (pq, r, I). Hence, we cannot directly charge (pq, r, I) to such events in $\mathcal{A}_{ru}$, so the charging is performed indirectly, via the crossing (uq, r, I_u), which is then classified as heavy for (pq, r, I). (Note, though, that the same crossing (uq, r, I_u) can be heavy for $\Omega(k)$ neighboring (q, r)-crossings (pq, r, I). This is compensated by the fact that u and (uq, r, I_u) can be chosen in $\Theta(k)$ possible ways.)

In case (d), the $(4k)$-shallow events under consideration are more restricted and involve *three* fixed points u, r, p. As in case (b), the knowledge u, r, and the time t^*, of each event, enables us to guess q and (uq, r, I_u) in $O(1)$ possible ways. However, since the point p is now also involved in the event, we can now guess it too in $O(1)$ possible ways. This enables direct charging of such events by (pq, r, I).

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] H. Edelsbrunner and E. P. Mücke, Simulation of simplicity: a technique to cope with degenerate cases in geometric algorithms, *ACM Transactions on Graphics* 9 (1990), 66–104.

[14] 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.

[15] 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.

[16] 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.

[17] 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.

[18] C. Icking, R. Klein, N.-M. Lê and L. Ma, Convex distance functions in 3-space are different, *Fundam. Inform.* 22 (4) (1995), 331–352.

[19] 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.

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

[21] 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.

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

[23] N. Rubin, On topological changes in the Delaunay triangulation of moving points, *Proc. 28th Annu. Sympos. Comput. Geom.*, 2012, pp. 1–10.

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

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

A On Co-circularities and Collinearities of Points Moving at Unit Speeds

Lemma A.1. *Let P be a finite collection of points in the plane, each moving along some straight line at unit speed. Then (i) any four points of P can be co-circular at most three times, and (ii) no triple of points can be collinear more than twice.*

Proof. To see (i), we note that each co-circularity of a quadruple $\{p_i = (x_i(t), y_i(t)) \mid 1 \le i \le 4\}$ (in P) occurs at a time t when the following determinant is equal to zero (see, e.g., [12, 13]):

$$D(t) = \begin{vmatrix} 1 & 1 & 1 & 1 \\ x_1(t) & x_2(t) & x_3(t) & x_4(t) \\ y_1(t) & y_2(t) & y_3(t) & y_4(t) \\ x_1^2(t)+y_1^2(t) & x_2^2(t)+y_2^2(t) & x_3^2(t)+y_3^2(t) & x_4^2(t)+y_4^2(t) \end{vmatrix}$$

Since each p_i is moving along some line in $\mathbb{R}^2$, its respective location $(x_i(t), y_i(t))$ can be represented as $(x_i + u_i t, y_i + v_i t)$, where (x_i, y_i) is the location of p_i at the time $t = 0$. Furthermore, since each p_i is moving at unit speed, we obtain $u_i^2 + v_i^2 = 1$.

Substituting $x_i(t) = x_i + u_i t$ and $y_i(t) = y_i + v_i t$ into the previous expression for $D(t)$, and cancelling the equal terms $(u_i^2 + v_i^2)t^2 = t^2$ in the bottom row of the determinant, we can replace the equation $D(t) = 0$ with its cubic equivalent, with at most *three solutions*.

To see (ii), we note that each collinearity of a triple $\{p_i(t), | \ 1 \le i \le 3\}$ occurs at a time t when the following determinant is equal to zero:

$$F(t) = \begin{vmatrix} 1 & 1 & 1 \\ x_1(t) & x_2(t) & x_3(t) \\ y_1(t) & y_2(t) & y_3(t) \end{vmatrix}$$

Substituting $x_i(t) = x_i + u_i t$ and $y_i(t) = y_i + v_i t$, for $1 \le i \le 3$, we get that the equation $F(t) = 0$ is quadratic (for any choice of u_i and v_i), with at most *two solutions*. □

B The General Position Assumption

In our analysis we assume that no five points can become co-circular during the motion, no four points can become collinear, no two points can coincide, and no two events of either a co-circularity of four points or of collinearity of three points can occur simultaneously. In addition, we assume that in every co-circularity event involving some four points $a, b, p, q \in P$, each of the points, say a, crosses the circumcircle of the other three points b, p, q; that is, it lies outside the circle right before the event and inside right afterwards, or vice versa. Similarly, we assume that in every collinearity event involving some triple of points of P, each of the points crosses the line through the remaining two points. Degeneracies in the point trajectories of the above kinds can be handled, both algorithmically and combinatorially, by any of the standard symbolic perturbation techniques, such as simulation of simplicity [13]; for combinatorial purposes, a sufficiently small generic perturbation of the motions will get rid of any such degeneracy, without decreasing the number of topological changes in the diagram.

C Proof of Theorem 2.2

In this section we establish Theorem 2.2. Without loss of generality, we assume that the edge pq is Delaunay at time t_0. (If pq is Delaunay at time t_1 then we can argue in a fully symmetrical fashion.)

Consider the portion of the red-blue arrangement associated with pq within the time interval (t_0, t_1). As above, refer to the parametric plane in which this arrangement is represented as the $t\rho$-plane, where t is the time axis and ρ measures signed distances from L_{pq}. We define the *red* (resp., *blue*) *level* of a point $x = (t, \rho)$ in this parametric $\mathbb{R}^2$ as the number of red (resp., blue) functions that lie below (resp., above) x (in the ρ-direction). See Figure 127. It is easily checked that the level of a co-circularity event at time t, with circumcenter at distance ρ from L_{pq}, is the sum of the red and the blue levels of (t, ρ).

We distinguish between the following (possibly overlapping) cases:

(a) p and q participate in a k-shallow collinearity with a third point r at some moment during I. That is, Condition (i) is satisfied. (Note that here we do not care whether r crosses pq or $L_{pq} \setminus pq$.)

Suppose that this does not happen. That is, each time when a point $r \in P$ changes its color from red to blue or vice versa, the number of points on each side of L_{pq} is larger than k. Hence, either the number of points on each side of L_{pq} is always larger than k (during (t_0, t_1)), or the sets of red and blue points remain fixed throughout (t_0, t_1) (no crossing takes place), and the size of one of them is at most k. More concretely, either one of the sets contains fewer than k points at the start of I, and then no crossing can ever occur during I, or both sets contain at least k points at the start of I, and this property is maintained during I, by assumption. In the latter case Condition (iii) trivially holds, since removal of all points in $P \cap L_{pq}^+$ or in $P \cap L_{pq}^-$ guarantees that pq is a hull edge throughout (t_0, t_1), and thus belongs to the

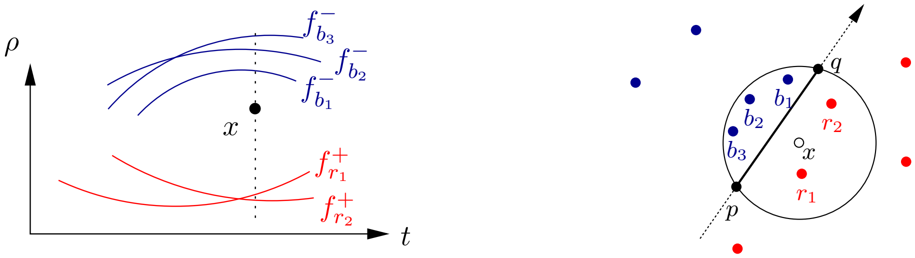

Figure 127: Left: The point $x = (t, \rho)$ lies below three blue functions and above two red functions, so its blue and red levels are 3 and 2, respectively. Right: The circumdisc centered at (signed) distance ρ from L_{pq} and touching p and q at time t contains the three corresponding blue points and two red points.

Delaunay triangulation. Hence, we may assume that the number of red points, and the number of blue points, are always both larger than k during (t_0, t_1).

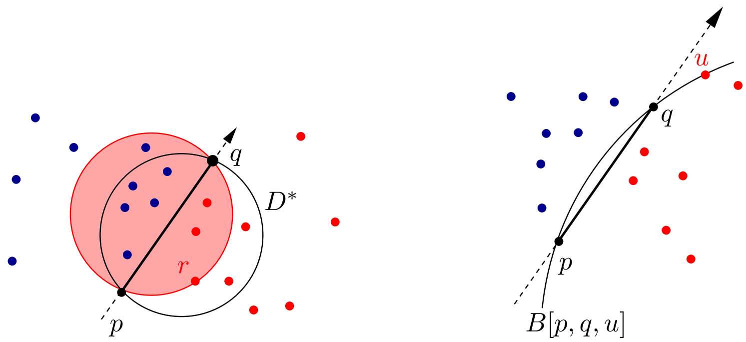

Figure 128: Left: Case (b). The disc D^* contains at least $\lceil k/3 \rceil = 5$ red points, and at least $\lceil k/3 \rceil$ blue points. If r lies at red level at most $\lceil k/3 \rceil$, it belongs to D^*. Hence, the circumdisc $B[p, q, r]$ contains at least $\lceil k/3 \rceil$ blue points, so the blue level of f_r^+ is at least $\lceil k/3 \rceil$. Right: Case (c). The setup right after time t' when u crosses $L_{pq} \setminus pq$. $B[p, q, u]$ contains at least k red points and no blue points.

(b) At some moment $t_0 \le t^* \le t_1$ there is a disc D^* that touches p and q, and contains at least $\lceil k/3 \rceil$ red points and at least $\lceil k/3 \rceil$ blue points. In particular, for each of the $\lceil k/3 \rceil$ shallowest red functions f_r^+ at time t^*, its respective red point r belongs to D^*. and similarly for the $\lceil k/3 \rceil$ shallowest blue functions. See Figure 128 (left). Before we use the existence of D^* we first conduct the following structural analysis.

Let f_r^+ be a red function which is defined at time t_0, and whose red level is then at most $\lfloor k/6 \rfloor$. (Recall that, at time t_0, the blue level of any red function is 0 since pq belongs to $\mathrm{DT}(P)$.) We claim that either f_r^+ is defined and continuous throughout (t_0, t_1) and its red level is always at most $\lceil k/3 \rceil$, or r participates in at least $\lceil k/6 \rceil$ red-red and/or red-blue co-circularities, all of which are $\lceil k/3 \rceil$-shallow.

Indeed, the circumdisc $B[p, q, r]$ contains at most $\lfloor k/6 \rfloor$ red points (and no blue points) at time t_0, and it moves continuously as long as r remains in L_{pq}^+. By the time at which either (the graph of) f_r^+ reaches red level $\lceil k/3 \rceil$ or r hits L_{pq}, this disc "swallows" either at least $\lceil k/6 \rceil$ red points (either in the former case or in the latter case when r crosses $L_{pq} \setminus pq$) or at least $\lceil k/6 \rceil$ blue points (in the latter case when r crosses pq). (Recall that, by assumption, the number of red points and the number of blue points is always larger than k during I.) We thus obtain at least $\lceil k/6 \rceil$ $\lceil k/3 \rceil$-shallow red-red or red-blue co-circularities involving p, q, r, and a fourth (red or blue) point.

To recap, if at least $\lfloor k/12 \rfloor$ red functions, which at time t_0 are among the $\lceil k/6 \rceil$ shallowest red functions, reach red level at least $\lceil k/3 \rceil + 1$, or have a discontinuity at $\rho = -\infty$ or $+\infty$ (at a crossing of L_{pq} by the corresponding point), then we encounter $\Omega(k^2)$ co-circularities (involving p and q) which are k-shallow, so Condition (ii) holds.

Hence, we may assume that at least $\lceil k/12 \rceil$ red functions f_r^+ that are among the $\lceil k/6 \rceil$ shallowest red functions at time t_0, are defined throughout (t_0, t_1), and their red level always remains at most $\lceil k/3 \rceil$. Fix any such red function f_r^+. Clearly, the red point r that defines f_r^+ belongs to D^* at time t^*, and the circumdisc $B[p,q,r]$ contains at least $\lceil k/3 \rceil$ blue points. See Figure 128 (left). This implies that the blue level of f_r^+ reaches $\lceil k/3 \rceil$ so (since the blue level was 0 at time t_0) r participates in at least $\lfloor k/6 \rfloor$ $\lceil k/3 \rceil$-shallow co-circularities during (t_0, t^*). Repeating this argument for each of the remaining $\lceil k/12 \rceil$ such red functions, we conclude that Condition (ii) is again satisfied.

(c) Suppose that neither of the two cases (a), (b) holds. Let A_R (resp., A_B) be the subset of all points u whose red (resp., blue) functions f_u^+ (resp., f_u^-) appear at red (resp., blue) level at most $\lceil k/3 \rceil$ at some moment during (t_0, t_1).

Since the situation in (b) does not occur, we can restore the Delaunayhood of pq, throughout the entire interval (t_0, t_1), by removing all points in $A_R \cup A_B$. To see this, suppose that pq is not Delaunay (in $\mathrm{DT}(P \setminus (A_R \cup A_B))$) at some time $t_0 < t^* < t_1$. This is witnessed by a disc D^* whose boundary passes through p and q and which contains a red point $r \notin A_R$ and a blue point $b \notin A_B$. Since the red level of f_r^+ is greater than $\lceil k/3 \rceil$ at time t^*, D^* must also contain the $\lceil k/3 \rceil$ red points corresponding to the $\lceil k/3 \rceil$ shallowest red functions at time t^*, and, symmetrically, also the $\lceil k/3 \rceil$ blue points corresponding to the $\lceil k/3 \rceil$ shallowest blue functions at time t^*. But then the disc D^* satisfies the conditions of Case (b), contrary to assumption.

Let A_R^o (resp., A_B^o) be the set of k points whose red (resp., blue) functions are shallowest at time t_0. It remains to consider the case where at least k points u in $A_R \cup A_B$ belong to neither of A_R^o, A_B^o, for otherwise Condition (iii) is trivially satisfied, with a removed set of size at most $3k$. Fix such a point u and consider the first time $t^* \in (t_0, t_1)$ when its red function f_u^+ has red level at most $\lceil k/3 \rceil$, or its blue function f_u^- has blue level at most $\lceil k/3 \rceil$. Without loss of generality, suppose that at time t^* the red function f_u^+ has red level at most $\lceil k/3 \rceil$. We claim that u does not cross pq during $(t_0, t^*]$. Indeed, if there were such a crossing from L_{pq}^- to L_{pq}^+ then the blue function f_u^- would tend to ∞ right before the crossing, and its blue level would then be 0 even before t^*, contrary to the choice of t^*. Similarly, if the crossing were from L_{pq}^+ to L_{pq}^- then the red level of f_u^+ would be 0 just before the crossing, again contradicting the choice of t^*.

First, assume that u does not cross L_{pq} during (t_0, t^*), so the graph of f_u^+ is continuous during this time interval. Hence, the motion of the circumdisc $B[p,q,u]$ is also continuous. Since $u \notin A_R^o$, at time t_0 the circumdisc $B[p,q,u]$ contains at least k red points and no blue points. At time t^*, $B[p,q,u]$ contains $\lceil k/3 \rceil$ red points and fewer than $\lceil k/3 \rceil$ blue points (otherwise Case (b) would occur). Hence, we encounter at least $\lfloor k/3 \rfloor$ k-shallow co-circularities during (t_0, t^*), each involving p, q, u and some other point of P.

Now, suppose u crosses $L_{pq} \setminus pq$ during (t_0, t^*), and consider the last time t' when this happens. We can use exactly the same argument as in the "continuous" case but now starting from t'. Indeed, f_u^+ is continuous during $(t', t^*]$ and, right after t', the circumdisc $B[p,q,u]$ contains (all the red points and thus) at least k red points, and no blue points. See Figure 128 (right).

Repeating this argument for all such points $u \in A_R \cup A_B \setminus (A_R^o \cup A_B^o)$, we get $\Omega(k^2)$ k-shallow co-circularities which occur during (t_0, t_1) and involve p and q. Hence, Condition (ii) is again satisfied. $\square$

D The number of double Delaunay crossings

In this subsection we show that any set P of n points moving as above in $\mathbb{R}^2$ admits at most $O(n^2)$ double Delaunay crossings. Since double Delaunay crossings are not possible if no ordered triple of points can be collinear more than once (i.e., if for any p, q, r the third point r can hit the segment pq at most once), we may assume throughout this subsection that no triple of points in P can be collinear more than twice.

Without loss of generality, we only bound the number of such double Delaunay crossings (pq, r, I) whose point r crosses through pq from L_{pq}^- to L_{pq}^+ during the first collinearity of p, q, r (and then returns back to L_{pq}^- during the second collinearity). Indeed, if the crossing (pq, r, I) does not satisfy the above condition then they are satisfied by (qp, r, I). Our goal is to show that (on average) a point r of P is involved in only few Delaunay crossings of edges that share the same endpoint p.

The following theorem provides certain structural properties of two double crossings that share the same crossing point (r) and one endpoint (p) of the crossed edges.

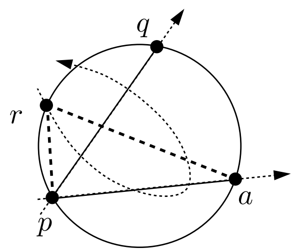

Figure 129: The trace of r according to Theorem D.1. The four points p, q, a, r are involved during I in two co-circularities, which are red-blue with respect to the edges pq and ra.

Theorem D.1. *Let (pq, r, I) and (pa, r, J) be two double Delaunay crossings of p-edges (that is, edges incident to p) pq, pa by the same point r. Assume that the first collinearity of p, q, r occurs before the first collinearity of p, a, r. Then the following properties hold (with the conventions assumed above):*

(i) a lies in L_{pq}^+ at both times when r hits pq.

(ii) q lies in L_{pa}^- at both times when r hits pa.

(iii) The points p, q, a, r are involved during $I \setminus J$ in two co-circularities, both of them red-blue with respect to pq and occurring when $r \in L_{pq}^-$ and $a \in L_{pq}^+$.

(iv) One of the two co-circularities in (iii) occurs before the beginning of J; right before it the Delaunayhood of ra is violated by p and q. A symmetric such co-circularity occurs after the end of J; right after it the Delaunayhood of ra is again violated by p and q. In particular, $J \subset I$.

The schematic description of the motion of r during I, according to the above theorem, is depicted in Figure 132 (right). Clearly, a suitable variant of Theorem D.1 exists also for similar pairs of double crossings of incoming p-edges qp, ap that are oriented *towards* p (again, by the same point r).

Proof. We first establish Part (ii) of the theorem. The crucial observation is that the first collinearity of p, a, r occurs when r lies in L_{pq}^+ (i.e., during the interval between the two collinearities of p, q, r). Indeed, otherwise the point a must lie in $L_{pq}^+ = L_{pr}^+$ at both collinearities of p, a, r, and q must lie in L_{pa}^+ at both collinearities of p, a, r. We shall prove that, in this hypothetical setup, the points p, q, a, r are involved in two co-circularities during I which are red-blue with respect to pq, and in a symmetric pair of co-circularities during J, both of them red-blue with respect to pa. That will clearly contradict the assumption that any four points can be co-circular at most three times.

Indeed, in the above situation the point a lies in the cap $B[p, q, r] \cap L_{pq}^+$ shortly before the first collinearity of p, q, r, and shortly after their second collinearity. Since $B[p, q, r]$ contains no points at the beginning of I, the point a must have entered this cap before the first collinearity of p, q, r. Moreover, a can enter this cap only through the boundary of $B[p, q, r]$, for otherwise it would hit pq during I, and no point of $P \setminus \{p, q, r\}$ can hit pq during its Delaunay crossing by r. This argument gives us the first of the promised two red-blue co-circularities that p, q, a, r define with respect to pq. The second such co-circularity is symmetric to the first one, and occurs when a leaves the cap $B[p, q, r] \cap L_{pq}^+$ (and after r returns to L_{pq}^- through pq). See Figure 130 (left). The other pair of co-circularities, both red-blue with respect to pa, is obtained by applying a fully symmetric argument to the cap $B[p, a, r] \cap L_{pa}^+$ and the point r. See Figure 130 (center). (For example, we can switch the roles of q and a by reversing the direction of the time axis.) Finally, all four co-circularities are distinct, because the same co-circularity cannot be red-blue with respect to two edges pq, pa with a common endpoint.

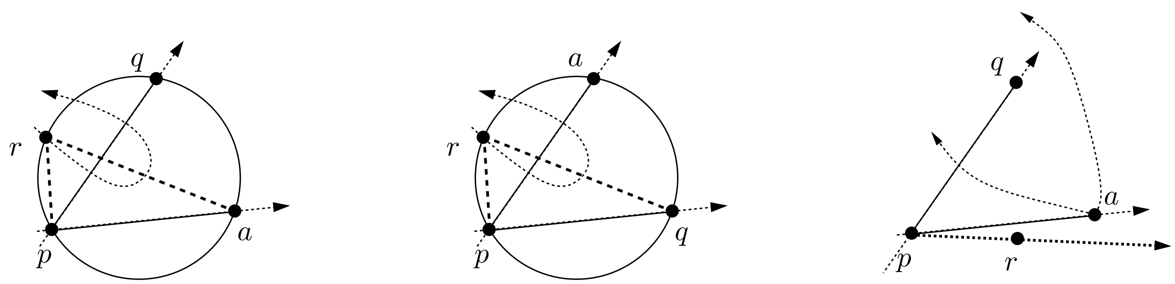

Figure 130: Proof of Theorem D.1. Left and center: The hypothetical case where r first hits pa within L_{pq}^-, after twice hitting pq. The points p, q, a, r are involved in a pair of co-circularities during I, and in a symmetric pair of co-circularities during J. Right: The hypothetical traces of a if it enters L_{pq}^+ before r (and before the second collinearity of p, a, r occurs).

Hence, we can assume, from now on, that the first time when r hits pa occurs when both points lie in L_{pq}^+. To complete the proof of Part (ii), it suffices to show that the points a and r still remain in L_{pq}^+ during the second collinearity of the triple p, a, r. Indeed, otherwise a must lie in L_{pq}^- when r hits pq for the second time, because, untill it crosses pa again, a lies in L_{pr}^- which coincides with L_{pq}^- at the second crossing of pq by r. See Figure 130 (right). That is, a must cross L_{pq} from L_{pq}^+ to L_{pq}^- while r still remains in L_{pq}^+, and before r hits the edges pq, pa for the second time. In particular, the above collinearity of p, q, a must occur during $I \cap J$. Clearly, the point a can potentially cross L_{pq} in three ways. If a crosses L_{pq} within pq, this contradicts the definition of I as the interval of the Delaunay crossing of pq by r. If a hits $L_{pq} \setminus pq$ within the ray emanating from q then (at that very moment) q hits pa, which contradicts the definition of J. Finally, a cannot hit $L_{pq} \setminus pq$ within the outer ray emanating from p before an additional (and forbidden) collinearity of p, a, r takes place. This establishes part (ii), and the analysis given above immediately implies part (i) two.

Part (i) follows immediately from Part (ii), because a lies in L_{pr}^+ during both collinearities of p, q, r.

Parts (iii) and (iv) follow from Parts (i) and (ii). Indeed, recall that the open disc $B[p, q, r]$ contains no points of P at the beginning of I. Right before r hits pq for the first time, the right cap $B[p, q, r] \cap L_{pq}^+$ of this disc contains a. Clearly, a first enters this cap through the corresponding portion of $\partial B[p, q, r]$. This determines the first red-blue co-circularity with respect to pq, right before which the Delaunayhood of ra is violated by p and q. The symmetric such co-circularity occurs during I when the point a leaves the cap $B[p, q, r] \cap L_{pq}^+$, after the second collinearity of p, q, r. Clearly, the Delaunayhood of ra is violated right after that co-circularity by p and q. By Lemma 4.1, neither of these co-circularities can occur during J, because ra remains Delaunay throughout J. Hence, the former one occurs, according to the previously established Parts (i) and (ii), before J, and the latter one occurs after J. This establishes parts (iii) and (iv), and completes the proof. $\square$

Theorem D.2. *Let P be a set of n points, whose motion in $\mathbb{R}^2$ respects the following conventions: (i) any four points can be co-circular at most three times, and (ii) no three points can be collinear more than twice. Then P admits at most $O(n^2)$ double Delaunay crossings.*

Proof. We fix a pair of points p, r in P. Our strategy is to show that, for an average such pair, there is at most a constant number of double Delaunay crossings of p-edges by r. Indeed, let $(pq_1, r, I_1), (pq_2, r, I_2), \ldots, (pq_k, r, I_k)$ be the complete list of such double Delaunay crossings of p-edges by r, and assume that r hits the edges $pq_1, pq_2, \ldots, p_{q_k}$, for the first time, in this same order. By Theorem D.1, the respective intervals of the above double crossings form a nested sequence $I_1 \supset I_2 \supset \ldots \supset I_k$.

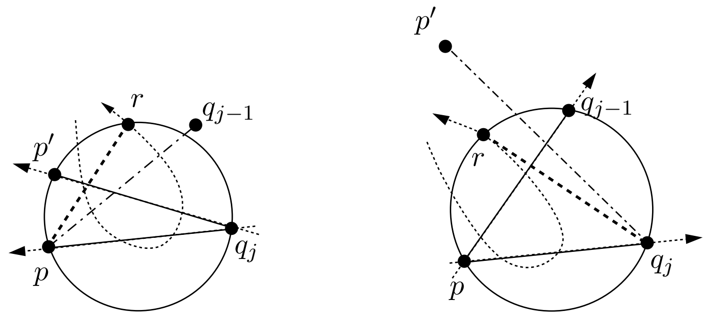

Figure 131: Proof of Theorem D.2. Left: If the double crossing $(p'q_j, r, I')$ ends before the end of I_{j-1} then the second co-circularity of q_j, p, p', r occurs during I_{j-1}. Right: If the double crossing $(p'q_j, r, I')$ ends after I_{j-1} then the second co-circularity of p, q_{j-1}, q_j, r occurs during I'.

Clearly, the first crossing (pq_1, r, I_1) can be uniquely charged to the pair p, r. Now assume that $k > 1$. We show that each of the additional double Delaunay crossings (pq_j, r, I_j), for $2 \le j \le k$, can be uniquely charged to the corresponding pair q_j, r. Specifically, we show that no double Delaunay crossing of incoming q_j-edges $p'q_j$ (that is, p-edges that are oriented towards p), by r, can end after I_j. In other words, (pq_j, r, I_j) is the "last" such double crossing.

Indeed, fix $2 \le j \le k$ as above. We first show that no double crossing of the form $(p'q_j, r, I')$ can end during the interval which lasts from the end of I_j and to the end of I_{j-1}. Indeed, suppose to the contrary that such a situation occurs, and apply a suitable variant of Theorem D.1 to the double Delaunay crossings of q_j-edges $p'q_j$ and pq_j by r. By Part (iv) of that theorem, I_j is contained in I', and the four points q_j, p, p', r are involved in a red-blue co-circularity with respect to $p'q_j$ during the second portion of $I' \setminus I_j$. See Figure 131 (left). Right after that co-circularity, the Delaunayhood of pr is violated by q_j and p'. If I' ends before the end of I_{j-1}, the above co-circularity must occur during I_{j-1} (as $I_{j-1} \supset I_j$), which contradicts Lemma 4.1 (applied to the crossing of pq_{j-1} by r).

It remains to show that no double Delaunay crossing $(p'q_j, r, I')$, as above, can end after the end of I_{j-1}. Indeed, by Part (iv) of Theorem D.1 (now applied to the double crossings of the p-edges pq_{j-1} and of pq_j, by r), the points p, q_{j-1}, q_j, r are involved in a co-circularity during the second portion of $I_{j-1} \setminus I_j$. Right after this co-circularity, the Delaunayhood of $q_j r$ is violated by p and q_{j-1}. If the interval I' (which contains I_j) ends after the end of I_{j-1}, the aforementioned co-circularity must occur during I'; see Figure 131 (right). However, this is another contradiction to Lemma 4.1 (now applied to the crossing of $p'q_j$ by r, which takes place during I').

We have shown that every double Delaunay crossing can be uniquely charged to an (ordered) pair of points of P, so their number is $O(n^2)$, as asserted. $\square$

E Proof of Lemma 4.2

Assume with no loss of generality that r lies in L^-_{pq}. Clearly, it is sufficient to establish only the Delaunayhood of the edge rq; the Delaunayhood of pr follows in a fully symmetrical manner.

The crucial observation is that the cap $B[p,q,r] \cap L^-_{pq}$ has Q-empty interior (or, else, pq would be Delaunay also in $Q \cup \{r\}$). That is, in terms of the static red-blue arrangement of pq, the corresponding blue function f^+_r of r coincides with the blue upper envelope E^-.

Assume for a contradiction that rq is not Delaunay in $Q \cup \{r\}$. We now consider the static red-blue arrangement of rq. Let $x \in Q \cap L^+_{rq}$ be the point whose function f^+_x (all functions in this argument are from the red-blue arrangement of rq) coincides with the red lower envelope E^+ (again, with respect to rq). In particular, we have $f^+_x \leq f^+_p$ (as is easily checked, $p \in L^+_{rq}$, when $r \in L^-_{pq}$). Clearly, x cannot be equal to p, for then the disc $B[p,q,r]$ would have Q-empty interior. Indeed, we argued that $B[p,q,r] \cap L^-_{pq}$ is Q-empty, and a similar argument shows that $B[p,q,r] \cap L^+_{rq}$ would also have to be empty if x and p coincide, from which the emptiness of the whole interior follows. It follows that pq is Delaunay in $Q \cup \{r\}$, contradicting the definition of a Delaunay crossing. See Figure 132 (left). Moreover, x cannot lie in L^-_{pq}, for it would then have to lie in $B[p,q,r] \cap L^-_{pq}$ (because $f^+_x < f^+_p$), which is impossible since this portion of $B[p,q,r]$ is Q-empty. Thus, $p \in L^-_{xq}$.

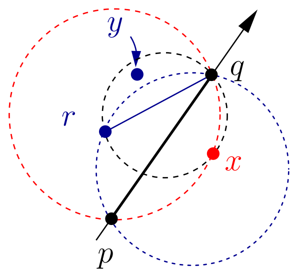

Figure 132: Left: Proof of Lemma 4.2.

Since rq is not Delaunay, the disc $B = B[q,r,x]$ contains another point $y \in Q \cap L^-_{rq}$, which is easily seen to lie in L^-_{pq} and in L^-_{xq}. We can move B so that its boundary continues to touch x and q and its portion within L^-_{xq} expands, until its boundary touches p, q and x, and its interior contains y. This implies that pq does not belong to $\mathrm{DT}(Q)$, which contradicts the definition of a Delaunay crossing. $\square$