Kinetic Voronoi Diagrams and Delaunay Triangulations under Polygonal Distance Functions*

Pankaj K. Agarwal† Haim Kaplan‡ Natan Rubin§ Micha Sharir¶

May 13, 2015

Abstract

Let P be a set of n points and Q a convex k-gon in $\mathbb{R}^2$. We analyze in detail the topological (or discrete) changes in the structure of the Voronoi diagram and the Delaunay triangulation of P, under the convex distance function defined by Q, as the points of P move along prespecified continuous trajectories. Assuming that each point of P moves along an algebraic trajectory of bounded degree, we establish an upper bound of $O(k^4 n\lambda_r(n))$ on the number of topological changes experienced by the diagrams throughout the motion; here $\lambda_r(n)$ is the maximum length of an (n, r)-Davenport-Schinzel sequence, and r is a constant depending on the algebraic degree of the motion of the points. Finally, we describe an algorithm for efficiently maintaining the above structures, using the kinetic data structure (KDS) framework.

*Work by P.A. and M.S. was supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation. Work by P.A. was also supported by NSF under grants CCF-09-40671, CCF-10-12254, and CCF-11-61359, by an ARO contract W911NF-13-P-0018, and by an ERDC contract W9132V-11-C-0003. Work by H.K. has been supported by grant 822/10 from the Israel Science Foundation, grant 1161/2011 from the German-Israeli Science Foundation, and by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11). Work by N.R. was partially supported by Grants 975/06 and 338/09 from the Israel Science Fund, by Minerva Fellowship Program of the Max Planck Society, by the Fondation Sciences Mathématiques de Paris (FSMP), and by a public grant overseen by the French National Research Agency (ANR) as part of the Investissements d'Avenir program (reference: ANR-10-LABX-0098). Work by M. S. has also been supported by NSF Grant CCF-08-30272, by Grants 338/09 and 892/13 from the Israel Science Foundation, by the Israeli Centers for Research Excellence (I-CORE) program (center no. 4/11), and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.

†Department of Computer Science, Duke University, Durham, NC 27708-0129, USA; pankaj@cs.duke.edu.

‡School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; haimk@tau.ac.il.

§Department of Computer Science, Ben-Gurion University of the Negev, Beer-Sheva, Israel 84105, rubinnat.ac@gmail.com.

¶School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; michas@tau.ac.il.

1 Introduction

Let P be a set of n points in $\mathbb{R}^2$, and let Q be a compact convex (not necessarily polygonal) set in $\mathbb{R}^2$ with nonempty interior and with the origin lying in its interior. For an ordered pair of points $x, y \in \mathbb{R}^2$, the *Q-distance* from x to y is defined as

$$d_Q(x, y) = \min\{\lambda \mid y \in x + \lambda Q\};$$

d_Q is a metric if and only if Q is centrally symmetric with respect to the origin (otherwise d_Q need not be symmetric). For a point of P, the *Q-Voronoi cell* of p is defined as

$$\mathrm{Vor}_Q(p) = \{x \in \mathbb{R}^2 \mid d_Q(x, p) \leq d_Q(x, p')\ \forall p' \in P\}.$$

If the points of P are in general position with respect to Q (see Section 2 for the definition), the Voronoi cells of points in P are nonempty, have pairwise-disjoint interiors, and partition the plane (see Figure 1(b)). The planar subdivision induced by these Voronoi cells is referred to as the *Q-Voronoi diagram* of P and we denote it as $\mathrm{VD}_Q(P)$.

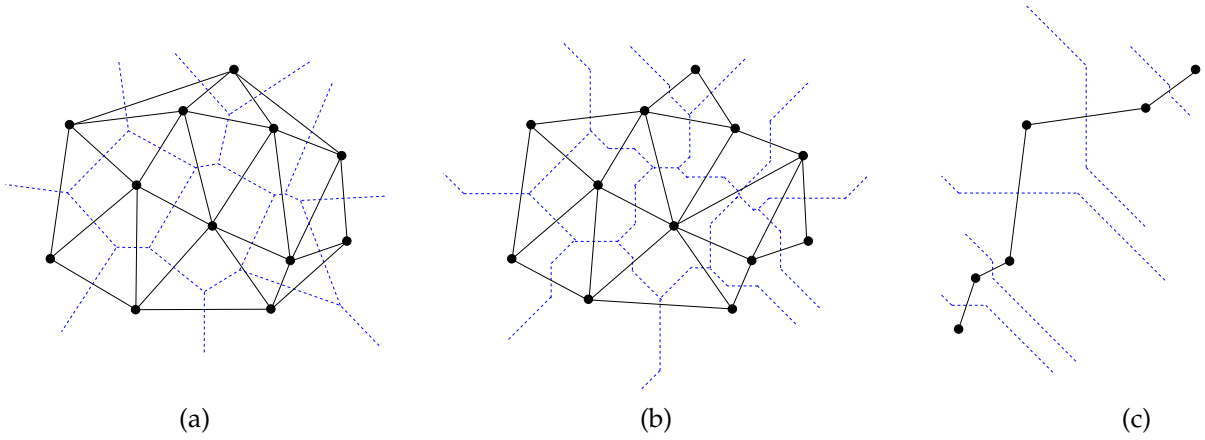

(a) (b) (c)

Figure 1. (a) The Euclidean Voronoi diagram (dotted) and Delaunay triangulation (solid). (b) $\mathrm{VD}_Q(P)$ and $\mathrm{DT}_Q(P)$ for an axis-parallel square Q, i.e., the diagrams VD_Q and DT_Q under the L_∞-metric. (c) $\mathrm{VD}_Q(P)$ and $\mathrm{DT}_Q(P)$, for the same Q as in (b), with an empty-interior support hull ($\mathrm{VD}_Q(P)$ has no vertices in this case).

The *Q-Delaunay triangulation* of P, denoted by $\mathrm{DT}_Q(P)$, is the dual structure of $\mathrm{VD}_Q(P)$. Namely, a pair of points $p, q \in P$ are connected by an edge in $\mathrm{DT}_Q(P)$ if and only if the boundaries of their respective Q-Voronoi cells $\mathrm{Vor}_Q(p)$ and $\mathrm{Vor}_Q(q)$ share a Q-Voronoi edge, given by

$$e_{pq} = \{x \in \mathbb{R}^2 \mid d_Q(x, p) = d_Q(x, q) \leq d_Q(x, p')\ \forall p' \in P\}.$$

$\mathrm{DT}_Q(P)$ can be defined directly as well: it is composed of all edges pq, with $p, q \in P$, for which there exists a homothetic placement of Q whose boundary touches p and q and whose interior contains no other points of P.[1] Placements of Q whose interior contains no point of P are called

[1] We remark that $\mathrm{Vor}_Q(p)$ is often defined in the literature as the set $\mathrm{Vor}_Q(p) = \{x \in \mathbb{R}^2 \mid d_Q(p, x) \leq d_Q(q, x)\ \forall q \in P\}$ [4, 9, 26]. If Q is not centrally symmetric, then this definition of $\mathrm{Vor}_Q(p)$ is not the same as the one given above. Furthermore, under this definition, pq is an edge of $\mathrm{DT}_Q(P)$ if there exists a P-empty homothetic placement of $-Q$ (and not of Q) whose boundary touches p and q.

P-empty. If Q is a circular disk then $\mathrm{DT}_Q(P)$ (resp., $\mathrm{VD}_Q(P)$) is the well-known Euclidean Delaunay triangulation (resp., Voronoi diagram) of P.

If P is in general position with respect to Q, then $\mathrm{DT}_Q(P)$ is spanned by so called Q-Delaunay triangles. Each of these triangles $\triangle pqr$ corresponds to the (unique) P-empty homothetic placement Q_{pqr} of Q whose boundary touches p, q, and r. That is, $\triangle pqr$ corresponds to a Q-Voronoi vertex v_{pqr} that lies at equal Q-distances to p, q, and r, so that v_{pqr} is the center of Q_{pqr} (that is, v_{pqr} is the image of the origin under the homothetic mapping of Q into Q_{pqr}). If Q is smooth (e.g., as in the Euclidean case), then $\mathrm{DT}_Q(P)$ is a triangulation of the convex hull of P; otherwise it is a triangulation of a simply-connected polygonal subregion of $\mathrm{conv}(P)$, sometimes referred to as the *support hull* of P with respect to Q (see [26] and Figure 1 (b)). The interior of the support hull may be empty, as shown in Figure 1 (c).

In many applications of Delaunay/Voronoi methods (e.g., mesh generation and kinetic collision detection), the points in P move continuously, so these structures need to be updated efficiently as motion takes place. Even though the motion of the points of P is continuous, the topological structures of $\mathrm{VD}_Q(P)$ and $\mathrm{DT}_Q(P)$ change only at discrete times when certain *events* occur.[2] Assume that each point of P moves independently along some known trajectory. Let $p_i(t) = (x_i(t), y_i(t))$ denote the position of point p_i at time t, and set $P(t) = \{p_1(t), \ldots, p_n(t)\}$. We call the motion of P *algebraic* if each $x_i(t), y_i(t)$ is a polynomial function of t, and the *degree* of the motion of P is the maximum degree of these polynomials.[3]

In this paper we focus on the case when Q is a convex k-gon and study the resulting Q-Voronoi and Q-Delaunay structures as each point of P moves continuously along an algebraic trajectory whose degree is bounded by a constant. Since Q will be either fixed or obvious from the context, we will use the simplified notations $\mathrm{Vor}(p)$, $\mathrm{VD}(P)$, and $\mathrm{DT}(P)$ to denote $\mathrm{Vor}_Q(p)$, $\mathrm{VD}_Q(P)$, and $\mathrm{DT}_Q(P)$, respectively.

Related work. There has been extensive work on studying the geometric and topological structure of Voronoi diagrams and Delauany triangulations under convex distance functions; see e.g. [4] and the references therein. In the late 1970s, $O(n \log n)$-time algorithms were proposed for computing the Voronoi diagram of a set of n points in $\mathbb{R}^2$ under any L_p-metric [14, 18, 19]. In the mid 1980s, Chew and Drsydale [9] and Widmayer *et al.* [26] showed that if Q is a convex k-gon, $\mathrm{VD}(P)$ has $O(nk)$ size and that it can be computed in $O(kn \log n)$ time. Motivated by a motion-planning application, Leven and Sharir [20] studied Voronoi diagrams under a convex polygonal distance function for the case where the input sites are convex polygons. Efficient divide-and-conquer, sweep-line, and edge-flip based incremental algorithms have been proposed to compute $\mathrm{DT}(P)$ directly [10, 21, 25]. Several recent works study the structure of $\mathrm{VD}(P)$ under a convex polyhedral distance function in $\mathbb{R}^3$ [7, 15, 17].

One of the hardest and best-known open problems in discrete and computational geometry is to determine the asymptotic behavior of the maximum possible number of discrete changes

[2]The *topological structures* of $\mathrm{DT}_Q(P)$ and $\mathrm{VD}_Q(P)$ are the graphs that they define. More specifically, the topological structure of $\mathrm{DT}_Q(P)$ and $\mathrm{VD}_Q(P)$ consists of the set of triples of points of P defining the Voronoi vertices, and the sets of Voronoi and Delaunay edges and faces, where each edge is associated with a pair of points of P and with features of Q that define it. As we will see later each Voronoi edge is a sequence of one or more edgelets. Each such edgelet is defined by a pair of edges of Q. The sequences of pairs of edges of Q defining the edgelet structures of the Voronoi edges are also part of the topological structure of $\mathrm{VD}_Q(P)$.

[3]This assumption can be somewhat relaxed to allow more general motions, as can be inferred from the analysis in the paper.

experienced by the *Euclidean* Delaunay triangulation during an algebraic motion of constant degree of the points of P, where the prevailing conjecture is that this number is nearly quadratic in n. A near-cubic bound was proved in [13]. After almost 25 years of no real progress, two recent works by one of the authors [22, 23] substantiate this conjecture, and establish an almost tight upper bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, for restricted motions where any four points of P can become cocircular at most *twice* (in [22]) or at most *three* times (in [23]). In particular, the latter result [23], involving at most three cocircularities of any quadruple, applies to the case of points moving along lines at common (unit) speed. Still, only near-cubic bounds are known so far for more general motions.

Chew [8] showed that the number of topological changes in the Delaunay triangulation under the L_1 or L_∞ metric is $O(n\lambda_r(n))$, where $\lambda_r(n)$ is the almost-linear maximum length of a Davenport-Schinzel sequence of order r on n symbols, and r is a constant that depends on the algebraic degree of the motions of the points. Chew's result also holds for any convex quadrilateral Q. He focuses on bounding the number of changes in the Delaunay triangulation and not how it changes at each "event," so his analysis omits some critical details of how the Delaunay triangulation and the Voronoi diagram change at an event; changes in the topological structure of $\mathrm{VD}(P)$ are particularly subtle. Chew remarks, without supplying any details, that his technique can be extended to general convex polygons.

Later, Basch *et al.* [5] introduced the *kinetic data structure* (KDS in short) framework for designing efficient algorithms for maintaining a variety of geometric and topological structures of mobile data. Several algorithms have been developed in this framework for kinetically maintaining various geometric and topological structures; see [12]. The crux in designing an efficient KDS is finding a set of *certificates* that, on one hand, ensure the correctness of the configuration currently being maintained, and, on the other hand, are inexpensive to maintain as the points move. When a certificate fails during the motion of the objects, the KDS fixes the configuration, replaces the failing certificate(s) by new valid ones, and computes their failure times. The failure times, called *events*, are stored in a priority queue, to keep track of the next event that the KDS needs to process. The performance of a KDS is measured by the number of events that it processes, the time taken to process each event, and the total space used. If these parameters are small (in a sense that may be problem dependent and has to be made precise), the KDS is called, respectively, *efficient*, *responsive*, and *compact*. See [5, 12] for details.

Delaunay triangulations and Voronoi diagrams are well suited for the KDS framework because they admit local certifications associated with their individual features. These certifications fail only at the events when the topological structure of the diagrams changes. The resulting KDS is compact ($O(n)$ certificates suffice) and responsive (each update takes $O(\log n)$ time, mainly to update the event priority queue), but its efficiency, namely, the number of events that it has to process, depends on the number of topological changes in $\mathrm{DT}(P)$, so a near quadratic bound on the number of events for the Euclidean case holds only when each point moves along some line with unit speed (or in similar situations when only three co-circularities can exist for any quadruple of points). A KDS for $\mathrm{DT}(P)$ when Q is a convex quadrilateral was presented by Abam and de Berg [1], but it is not straightforward to extend their KDS for the case where Q is a general convex k-gon. Furthermore, it is not clear how to use their KDS for maintaining $\mathrm{VD}(P)$.

Our contribution. First, we establish a few key topological properties of $\mathrm{VD}(P)$ and $DT(P)$ when P is a set of n stationary points in $\mathbb{R}^2$ and Q is a convex k-gon (Section 2). Although these properties follow from earlier work on this topic (see [4, Chapter 7]), we include them here because they are

important for the kinetic setting and most of them have not been stated in earlier work in exactly the same general and detailed form as here.

Next, we characterize the topological changes that $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ can undergo when the points of P move along continuous trajectories (Section 3). These changes occur at critical moments when the points of P are not in Q-general position, so that some $O(1)$ points of P are involved in a *degenerate* configuration with respect to Q. The most ubiquitous type of such events is when four points of P become *Q-cocircular*, in the sense that there exists a P-empty homothetic placement of Q whose boundary touches those four points.

We provide the first comprehensive and rigorous asymptotic analysis of the maximum number of topological changes that $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ can undergo during the motion of the points of P (Section 4). Specifically, if Q has k vertices, then $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ experience $O(k^4 n \lambda_r(n))$ such changes, where $\lambda_r(n)$ is the almost-linear maximum length of a Davenport-Schinzel sequence of order r on n symbols, and r is some constant that depends on the algebraic degree of the motions of the points. Some of these changes occur as components of so-called *singular sequences*, in which several events that affect the structure of $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ occur simultaneously, and their collective effect might involve a massive change in the topological structures of these diagrams. These compound effects are a consequence of the non-strict convexity of Q, and their analysis requires extra care. Nevertheless, the above near-quadratic bound on the number of changes also holds when we count each of the individual critical events in any such sequence separately.

Finally, we describe an efficient algorithm for maintaining $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ during an algebraic motion of P, within the standard KDS framework (Section 5). Here we assume an algebraic model of computation, in which algebraic computations, including solving a polynomial equation of constant degree, can be performed in an exact manner, in constant time. The precise sense of this assumption is that comparisons between algebraic quantities that are defined in this manner can be performed exactly in constant time. This is a standard model used widely in theory [24, Section 6.1] and nowadays also in practice (see, e.g., [11]). This model allows us to perform in constant time the various computations that are needed by our KDS, the most ubiquitous of which are the calculation of the failure times of the various certificates being maintained; see Section 5 for details.

This paper is a stand-alone study of a fairly basic problem in the theory and practice of Delaunay and Voronoi diagrams. In contrast with earlier studies, it analyzes the problem in full generality, and addresses all the intricate aspects of the kinetic changes that the diagrams undergo. In addition, the results of this paper form a major ingredient of the analysis in the companion paper [3] (see also [2]), concerning *stable* Delaunay graphs. Details concerning this construct, its properties, and its connection to the results of this paper are presented in the the companion paper.

2 The topology of $\mathrm{VD}(P)$

In this section we state and prove a few geometric and topological properties of the Q-Voronoi diagram of a set of stationary points when Q is a convex polygon.

Some notations. Let Q be a convex k-gon with vertices $v_0, \ldots, v_{k-1}$ in clockwise order, whose interior contains the origin. For each $0 \leq i < k$, let e_i denote the edge $v_i v_{i+1}$ of Q, where index addition is modulo k (so $v_k = v_0$). We refer to the origin as the *center* of Q and denote it by o. A *homothetic placement* (or *placement* for short) Q' of Q is represented by a pair (p, λ), with $p \in \mathbb{R}^2$ and

$\lambda \in \mathbb{R}^+$, so that $Q' = p + \lambda Q$; p is the location of the center of Q', and λ is the *scaling factor* of Q' (about its center). The homothets of Q thus have three degrees of freedom.

There is an obvious bijection between the edges (and vertices) of Q' and of Q, so, with a slight abuse of notation, we will not distinguish between them and use the same notation to refer to an edge or vertex of Q and to the corresponding edge or vertex of Q'. For a point $u \in \mathbb{R}^2$, let $Q[u]$ denote the homothetic copy of Q centered at u such that its boundary touches the d_Q-nearest neighbor(s) of u in P, i.e., $Q[u]$ is represented by the pair (u, λ) where $\lambda = \min_{p \in P} d_Q(u, p)$. In other words, $Q[u]$ is the largest homothetic copy of Q that is centered at u whose interior is P-empty.

Q-general position. To simplify the presentation, we assume our point set P to be in *general position* with respect to the underlying polygon Q. Specifically, this means that

(Q1) no pair of points of P lie on a line parallel to a boundary edge or a diagonal of Q,

(Q2) no four points of P lie on the boundary of the same homothetic copy Q' of Q, and

(Q3) if some three points in P lie on the boundary of the same homothetic copy Q' of Q, then each of them is incident to a *relatively open edge* of $\partial Q'$ (and all the three edges are distinct, due to (Q1)), as opposed to one or more of these points touching a vertex of Q'.

The above conditions can be enforced by an infinitesimally small rotation of Q or of P.

Bisectors, corner placements, and edgelets. The *bisector* between two points p and q, with respect to the distance function d_Q induced by Q, denoted by b_{pq} or b_{qp}, is the set of all points $x \in \mathbb{R}^2$ that satisfy $d_Q(x, p) = d_Q(x, q)$. Equivalently, b_{pq} is the locus of the centers of all homothetic placements Q' of Q that touch p and q on their boundaries; Q' does not have to be P-empty, so it may contain additional points of $P \setminus \{p, q\}$. If p and q are not parallel to an edge of Q (assumption (Q1)), then b_{pq} is a one-dimensional polygonal curve, whose structure will be described in detail momentarily.

A homothetic placement Q' centered along b_{pq} that touches one of p and q, say, p, at a vertex, and touches q at the relative interior of an edge (as must be the case in general position) is called a *corner placement at p*; see Figure 2 (a). Note that a corner placement at which a vertex v_i of (a copy Q' of) Q touches p has the property that the center o' of Q' lies on the fixed ray emanating from p in direction $\vec{v_i o}$.

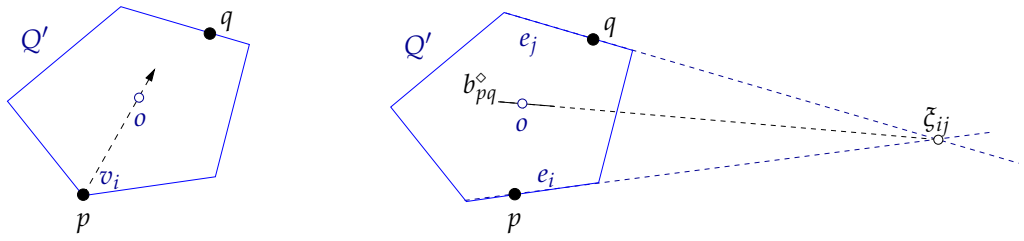

Figure 2. Possible placements of Q' on b_{pq}. (a) A corner placement at p. The center of Q' lies on the fixed ray emanating from p in direction $\vec{v_i o}$. (b) A placement of Q' with edge contacts of e_i and e_j at p and q, respectively. (The centers of) such placements trace an edgelet of b_{pq} with label (e_i, e_j).

A non-corner placement Q' centered on b_{pq} can be classified according to the pair of edges of Q', say, e_i and e_j, that touch p and q, respectively. We may assume (by (Q1)) that $e_i \neq e_j$. Slide Q' so that its center o' moves along b_{pq} and its size expands or shrinks to keep it touching p and q at the edges e_i and e_j, respectively. If e_i and e_j are parallel, then the center o' of Q' traces a line segment in the direction parallel to e_i and e_j; otherwise o' traces a segment in the direction that connects it to the intersection point ξ_{ij} of the lines containing (the copies on $\partial Q'$ of) e_i and e_j. See Figure 2 (b) for the latter scenario. We refer to such a segment g as an *edgelet* of b_{pq}, and label it by the pair $\psi(g) = (e_i, e_j)$ (or by (i, j) for brevity). The orientation of the edgelet depends only on the corresponding edges e_i, e_j, and is independent of p and q. The structure of b_{pq} is fully determined by the following proposition, with a fairly straightforward proof that is omitted from here.

Lemma 2.1. *An edgelet g with the label $\psi(g) = (e_i, e_j)$ appears on b_{pq} if and only if there is an oriented line parallel to $\vec{pq}$ that crosses ∂Q at (the relative interiors of) e_i and e_j, in this order.*

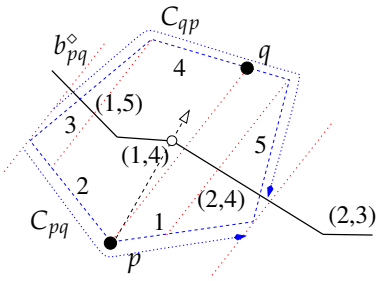

Figure 3. The edgelets of b_{pq}. The breakpoints of b_{pq} correspond to corner placements of Q. We have $C_{pq} = \langle 2, 1\rangle$ and $C_{qp} = \langle 3, 4, 5\rangle$. The terminal edgelets of b_{pq} are the rays with labels $(1, 5)$ and $(2, 3)$.

See Figure 3 for an illustration. The endpoints of edgelets are called the *breakpoints* of b_{pq}. Each breakpoint is the center of a corner placement of Q; If e_i and e_j are adjacent, then the edgelet labeled (i, j) is a ray and the common endpoint of e_i, e_j is one of the two vertices of Q extremal in the direction orthogonal to $\vec{pq}$ (i.e., these vertices have a supporting line parallel to $\vec{pq}$).

Assuming Q-general position of P, Lemma 2.2 below implies that b_{pq} is the concatenation of exactly $k - 1$ edgelets. Let C_{pq} and C_{qp} denote the two chains of ∂Q, delimited by the vertices that are extremal in the direction orthogonal to $\vec{pq}$, such that p lies on an edge of C_{pq} and q on an edge of C_{qp} at all placements of Q touching p and q and centered along the bisector b_{pq}. We orient both C_{pq}, C_{qp} so that they start (resp., terminate) at the vertex of Q that is furthest to the left (resp., to the right) of $\vec{pq}$; see Figure 3.

Our characterization of b_{pq} is completed by the following lemma, which follows from Lemma 2.1 and the preceding discussion.

Lemma 2.2. *Let $(e_{11}, e_{21}), \ldots, (e_{1h}, e_{2h})$ be the sequence of labels of the edgelets of b_{pq} in their order along b_{pq} when we trace it so that p lies on its left side and q on its right side. Then $e_{11}, \ldots, e_{1h}$ appear (with possible repetitions as consecutive elements) in this order along C_{pq}, and $e_{21}, \ldots, e_{2h}$ appear (again, with possible repetitions) this order along C_{qp}. Furthermore, the following additional properties also hold:*

(a) All the edges of C_{pq} (resp., C_{qp}) appear, possibly with repetitions, in the first (resp., second) sequence.

(b) The elements of C_{qp} appear in clockwise order and the elements of C_{pq} in counterclockwise order along ∂Q.

(c) *Assuming general position, the passage from a label* (e_{1i}, e_{2i}) *to the next label* $(e_{1(i+1)}, e_{2(i+1)})$ *is effected by either replacing* e_{1i} *by the following edge on* C_{pq} *or by replacing* e_{2i} *by the following edge on* C_{qp}*. In the former (resp., latter) case, the common endpoint of the two edgelets corresponds to the corner placement of* p *(resp.,* q*) at the common vertex of* e_{1i} *and* $e_{1(i+1)}$ *(resp.,* e_{2i} *and* $e_{2(i+1)}$*).*

The proof of the lemma, whose details are omitted, proceeds by sweeping a line parallel to $\vec{pq}$, and keeping track of the pairs of edges of Q that are crossed by the line, mapping each position of the line to a homothetic placement of Q that touches p and q at the images of the two intersection points.

For $0 \leq i < j < k$, let θ_{ij} be the orientation of the line passing through the vertices v_i and v_j of Q, and let Θ be the set of these orientations. Θ partitions the unit circle $\mathbb{S}^1$ into a collection $\mathcal{I}$ of $O(k^2)$ angular intervals (for a regular k-gon, the number of intervals is only $\Theta(k)$). Lemmas 2.1 and 2.2 implies the following corollary:

Corollary 2.3. *The sequence of edgelet labels along* b_{pq} *is the same for all the ordered pairs of points* p, q *such that the orientation of the vector* pq *lies in the same interval of* $\mathcal{I}$*.*

The following additional property of bisectors is crucial for understanding the topological structure of $\mathrm{VD}(P)$.

Lemma 2.4. *Let* p, q_1, q_2 *be three distinct points of* P*. The bisectors* b_{pq_1}*,* b_{pq_2} *can intersect at most once, assuming that* p, q_1 *and* q_2 *are in* Q*-general position.*

Proof. Suppose to the contrary that b_{pq_1}, b_{pq_2} intersect at two points. Then there exist two homothetic copies Q_1 and Q_2 of Q such that $p, q_1, q_2 \in \partial Q_1 \cap \partial Q_2$. However, it is well known that homothetic placements of Q behave like pseudo-disks, in the sense that the portion of the boundary of each of them outside the other homothetic placement is connected; see, e.g., [16]. Therefore, ∂Q_1 and ∂Q_2 intersect in at most two connected portions, each of which is either a point or a segment parallel to some edge of Q. Clearly, one of these connected components of $\partial Q_1 \cap \partial Q_2$ must contains two out of the three points p, q_1, and q_2, in contradiction to the fact that the points are in Q-general position. □

The following lemma provides additional details concerning the structure of the breakpoints of the bisectors in case Q is a regular k-gon.

Lemma 2.5. *Let* Q *be a regular* k*-gon, and let* p *and* q *be two points in general position with respect to* Q*. The breakpoints along the bisector* b_{pq} *correspond alternatingly to corner placements at* p *and corner placements at* q*.*

Proof. Refer to Figure 4. Suppose that two consecutive breakpoints of b_{pq} correspond to corner placements at p. From Lemmas 2.1 and 2.2, we obtain that these corner placements are formed by two adjacent vertices, say v_0 and v_1 of Q, and q lies in the relative interior of (the homothetic copies of) the same edge e of Q at these placements. This implies that the projections of v_0 and v_1 in direction $\vec{pq}$ lie in the interior of the projection of the edge e of Q in direction $\vec{pq}$, which is impossible if Q is a regular k-gon. Indeed, the convex hull of $e_0 = v_0v_1$ and of e is an isosceles trapezoid τ, which implies that, for any other strip σ bounded by two parallel lines through v_0 and v_1, e cannot cross both boundary lines of σ. We note that if σ is the strip spanned by τ, then e touches both lines bounding σ but does not cross any of them. This completes the proof of the lemma. □

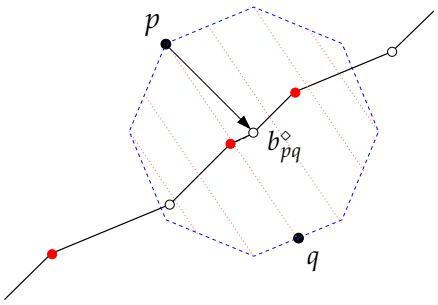

Figure 4. The bisector b_{pq} for a regular octagon Q; it has seven edgelets and the centers of the corner placements along b_{pq} alternate between p (hollow circles) and q (filled circles).

Voronoi cells, edges, and vertices. Each bisector b_{pq} partitions the plane into open regions $H_{pq} = \{x \mid d_Q(p,x) < d_Q(q,x)\}$ and $H_{qp} = \{x \mid d_Q(q,x) < d_Q(p,x)\}$. Hence, for each point $p \in P$, its Q-Voronoi cell $\mathrm{Vor}(p)$ can be described as $\bigcap_{q \in P \setminus \{p\}} H_{pq}$.

By Q-general position of P, for any $p \in P$, $\partial\mathrm{Vor}(p)$ is composed of *Q-Voronoi edges*, where each such edge is a maximal connected portion of the bisector b_{pq}, for some other point $q \in P$, that lies within $\partial\mathrm{Vor}(p) \cap \partial\mathrm{Vor}(q)$. The portion of b_{pq} within this common boundary can be described as

$$b_{pq} \cap \bigcap_{r \neq p,q} H_{pr} = b_{pq} \cap \bigcap_{r \neq p,q} H_{qr}.$$

That is, this portion is the locus of all centers x of placements of Q for which the equal distances $d_Q(x,p) = d_Q(x,q)$ are the smallest among the distances from x to the points of P. Note that the homothetic copy $x + d_Q(x,p)Q$ of Q touches p and q and is P-empty.

Since P is in Q-general position, Lemma 2.4 guarantees that this portion of b_{pq} is either connected or empty. Therefore, any bisector b_{pq} contains *at most one* Q-Voronoi edge, which we denote by e_{pq}. This edge is called a *corner edge* if it contains a breakpoint (i.e., a center of a corner placement); otherwise it is a *non-corner edge*—a line segment.

The endpoints of Q-Voronoi edges e_{pq} are called *Q-Voronoi vertices*. By the Q-general position of P, each such vertex is incident on exactly three Voronoi cells $\mathrm{Vor}(p)$, $\mathrm{Vor}(q)$, and $\mathrm{Vor}(r)$. This vertex, denoted by v_{pqr}, can be described as the center of the unique homothetic P-empty placement $Q' = Q[v_{pqr}]$ of Q, whose boundary contains only the three points p, q, and r of P. From the Delaunay point of view, $\mathrm{DT}(P)$ contains the triangle $\triangle pqr$.

We say that g is an *edgelet of* e_{pq} if (i) g is an edgelet of b_{pq}, and (ii) the Voronoi edge e_{pq} either contains or, at least, overlaps g. We refer to an edgelet g of e_{pq} as *external* if it contains one of the endpoints of e_{pq}, namely, a vertex of $\mathrm{VD}(P)$, and as *internal* otherwise. In general position, an external edgelet of e_{pq} is always properly contained in an edgelet of b_{pq}.

We conclude this section by making the following remarks. If assumption (Q3) does not hold, then a Voronoi vertex may coincide with a breakpoint of an edge adjacent to it; if (Q2) does not hold then a Voronoi vertex may have degree larger than three; if the segment pq connecting a pair $p, q \in P$ is parallel to a diagonal of Q, then an edgelet of a Voronoi edge may degenerate to a single point; and if such a segment pq is parallel to an edge of Q then b_{pq} may be a two-dimensional region (Figure 7 middle). These degenerate configurations are discussed in detail in the next section.

3 Kinetic Voronoi and Delaunay diagrams

As the points of P move along continuous trajectories, $\mathrm{VD}(P)$ also changes continuously, namely, vertices of $\mathrm{VD}(P)$ and breakpoints of edgelets trace continuous trajectories, but, unless the motion is very degenerate, the topological structure of $\mathrm{VD}(P)$ changes only at discrete times, at which an edgelet in a Voronoi edge appears/disappears, a Voronoi vertex moves from one edgelet to another, or two adjacent Voronoi cells cease to be adjacent or vice versa (equivalently, an edge appears or disappears in $\mathrm{DT}(P)$), because of a Q-cocircularity of four points of P. In this section we discuss when do these changes occur and how does $\mathrm{VD}(P)$ change at such instances. To simplify the presentation, (i) we assume that the orientations of the edges and diagonals v_iv_j, for all pairs of vertices v_i, v_j of Q, are distinct, and that they are different from those of ov_i, for any vertex v_i; (ii) we make certain general-position assumptions on the trajectories of P; and (iii) we augment P with some points at infinity. At the end of the section, we remark what happens if we do not make these assumptions or do not augment P in this manner.

Augmenting P. We add points to P so that the convex hull of the augmented set does not change as the (original) points move, and the boundary of $\mathrm{DT}(P)$ is this stationary convex hull at all times. Specifically, for each vertex v_i of Q, we add a corresponding point q_i at infinity, so that q_i lies in the direction $\vec{v_io}$. Let P_∞ denote the set of these k new points. We maintain $\mathrm{VD}(P \cup P_\infty)$ and $\mathrm{DT}(P \cup P_\infty)$. It can be checked that $\mathrm{DT}(P \cup P_\infty)$ contains all edges of $\mathrm{DT}(P)$, some "unbounded" edges (connecting points of P to points of P_∞), and k edges at infinity (forming the convex hull of $P \cup P_\infty$). Furthermore, every edge of $\mathrm{DT}(P \cup P_\infty)$ incident on at least one point of P is adjacent to two triangles; only the edges at infinity are "boundary" edges of the triangulation. During the motion of the points of P, the points of P_∞ remain stationary.

Let $\triangle$ be a triangle of $\mathrm{DT}(P \cup P_\infty)$. There is a $(P \cup P_\infty)$-empty homothetic copy $Q_\triangle$ of Q associated with $\triangle$, whose boundary touches the three vertices of $\triangle$. If two vertices of $\triangle$ belong to P and one vertex of $\triangle$ is a point q_i at infinity, then $Q_\triangle$ is a wedge formed by the two corresponding consecutive edges e_{i-1} and e_i of Q, each touching a vertex of $\triangle$ not in P_∞ (e.g., C_1 in Figure 5). If only one vertex of $\triangle$, say p, belongs to P, then $\triangle$ is of the form $\triangle pq_iq_{i+1}$ (for some $0 \leq i < k$), and there are arbitrarily large empty homothetic copies of Q incident on p at the edge $e_i = v_iv_{i+1}$ (e.g., C_2 in Figure 5). The number of triangles of the latter kind is only k, one for each edge of Q. Abusing the notation slightly, we will use P to denote $P \cup P_\infty$ from now on.

Q-general position for trajectories. We assume that the trajectories of the points of P are in *Q-general position*, which we define below. Informally, if the motion of each point of P is algebraic of bounded degree, as we assume, then the time instances at which degenerate configurations occur, namely, configurations violating one of the assumptions (Q1)–(Q3), can be represented as the roots of certain constant-degree polynomials in t. The present "kinetic" general-position assumption for the trajectories says that none of these polynomials is identically zero (so each of them has $O(1)$ roots), that each root has multiplicity one (so the sign of the polynomial changes in the neighborhood of each root), and that the roots of all polynomials are distinct. We now spell out these conditions in more detail and make them geometrically concrete.

(T1) For any pair of points $p, q \in P$, $p(t) \neq q(t)$ for all t, namely, p and q do not collide during the motion.

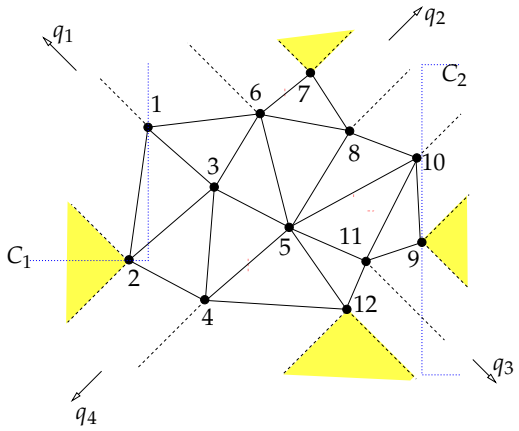

Figure 5. The extended DT(P) under the L_∞-metric (where Q is an axis-parallel square) for the set of points in Figure 1, with four points $q_1, \ldots, q_4$ added at infinity in directions $(\pm 1, \pm 1)$. Each of the four shaded triangles has two vertices at infinity, and the unbounded half-strips between them represent triangles with one vertex at infinity. The empty wedge C_1 corresponds to $\triangle 12q_1$ (the half-strip right above the left shaded triangle), and the arbitrarily large empty square C_2 (a halfplane in the limit) corresponds to $\triangle 9q_2q_3$ (the right shaded triangle).

(T2) For any pair of points $p, q \in P$, there exist at most $O(1)$ times when the segment pq is parallel to any given edge or diagonal v_iv_j of Q, and at each of these times p *properly* crosses the line through q parallel to v_iv_j, which moves continuously, together with q (and q does the same for the parallel line through p).

(T3) For any ordered set of three points $p, q, r \in P$ and for any vertex v_p and a pair of edges e_q, e_r of Q, there exist at most $O(1)$ times when q touches e_q and r touches e_r at a corner placement Q' of Q at p in which p touches v_p. Furthermore, given that e_q is not adjacent to v_p, at each such time the point r either enters or leaves the interior of the unique P-empty homothetic copy of Q that touches p at v_p and q at e_q.

(T4) For any four points $p, q, a, b \in P$ and any ordered quadruple e_p, e_q, e_a, e_b of edges of Q, such that at least three of these edges are distinct, there are only $O(1)$ times at which there exists a placement Q' of Q such that p, q, a, b touch the respective relative interiors of e_p, e_q, e_a, e_b. We say that p, q, a, b are Q-*cocircular* at these $O(1)$ times. At any such Q-cocircularity, the four points p, q, a, b are partitioned into two pairs, say, (p, q) and (a, b), so that right before the cocircularity there exists a homothetic copy of Q that is disjoint from a and b and whose boundary touches p and q, and right afterwards there exists a homothetic copy of Q that is disjoint from p and q and whose boundary touches a and b.

(T5) Events of type (T2)–(T4) do not occur simultaneously, except when two points p and q become parallel to an edge of Q. In this case there could be many events of type (T3) and (T4) that occur simultaneously, each of which involves p, q; see below for more details.

Events. Since the motion of P is continuous, the topological changes in VD(P) occur only when some points of P are involved in a degenerate configuration, i.e., they violate one of the assumptions (Q1)–(Q3). However not every degenerate configuration causes a change in VD(P). We define an *event* to be the occurrence of a P-empty placement of a homothetic copy of Q whose boundary contains two, three, or four points of P that are in a Q-degenerate configuration. The center of such

a placement lies on an edge or at a vertex of $\mathrm{VD}(P)$. The subset of points involved in the degenerate configuration is referred to as the subset *involved in the event*. The event is called a *bisector, corner*, or *flip* event if assumption (Q1), (Q2), or (Q3), respectively, is violated.

An event is called *singular* if some pair among the (constantly many) points involved in the event span a line parallel to an edge of Q. Otherwise, we say that the event is *generic*. The Q-general-position assumption on the trajectories of the points of P implies, in particular, that (i) no generic event can occur simultaneously with any other event, and (ii) all singular events that occur at a given time, must involve the *same* pair of points p, q that span a line parallel to an edge of Q.

The changes in $\mathrm{VD}(P)$ are simple and local at a generic event, but $\mathrm{VD}(P)$ can undergo a major change at a singular event. We therefore first discuss the changes at a generic event and then discuss singular events.

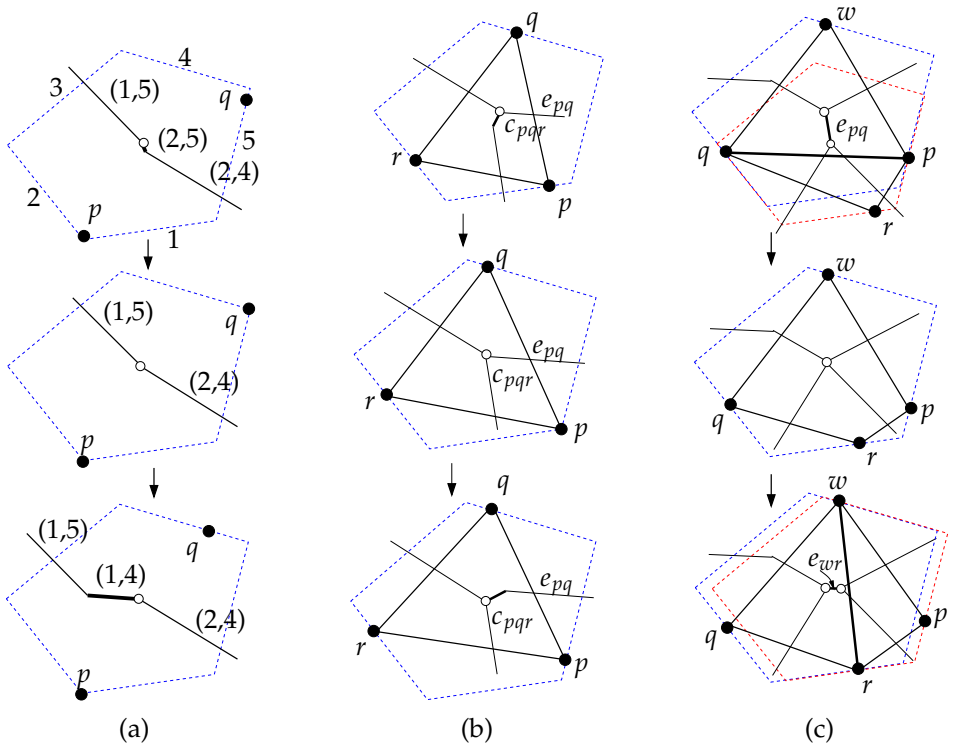

Figure 6. Different types of generic events for a pentagon Q: (a) a bisector event, (b) a corner event, and (c) a flip event. Thick segments denote an appropriate portion of $\mathrm{DT}(P)$ and the elements of $\mathrm{VD}(P)$ that change at the event.

3.1 Generic events

Recall that the orientation of pq, for every pair $p, q \in P$, at a generic event is different from that of any edge of Q, which implies that no two points of P lie on the same edge of a homothetic copy of Q at a generic event.

Bisector event. A pair of points $p, q \in P$ are incident to the vertices v_i and v_j, respectively, of a P-empty homothetic copy of Q so that the vertices v_i and v_j are *not consecutive* along ∂Q. In particular, e_{pq} is an edge in $\mathrm{VD}(P)$.

Recall that in our notation, v_i is adjacent to the consecutive edges e_{i-1}, e_i, and v_j is adjacent to the consecutive edges e_{j-1}, e_j; in the present scenario, these four edges are all distinct. Without loss of generality, we may assume that before the event, there is an oriented line parallel to $\vec{pq}$ that intersects e_{i-1} and e_{j-1} in this order, and there is no such line after the event. Similarly, after the event there is an oriented line parallel to $\vec{pq}$ that intersects e_i and e_j in this order, and there is no such line before the event. Hence, Lemma 2.1 and assumption (T2) imply that e_{pq} loses a bounded edgelet with label $(i-1, j-1)$, which is replaced by a new bounded edgelet with label (i, j). Our assumption (T5) implies $(i-1, j-1)$ cannot be an external edgelet of e_{pq}. Hence, $(i-1, j-1)$ appears shortly before the event as an internal edgelet of e_{pq}, shrinks to a point and is replaced by the new internal edgelet (i, j); see Figure 6 (a). This is the only topological change in $\mathrm{VD}(P)$ at this event.

Notice that whenever the direction of $\vec{pq}$ coincides with that of a diagonal v_iv_j of Q, p and q are incident to the vertices v_i and v_j, respectively, of a unique copy of Q. If this copy contains further points of P, then the bisector b_{pq} still loses an edgelet $(i-1, j-1)$, which is replaced by a new edgelet (i, j). However, both of these edgelets now belong to the portion of b_{pq} outside e_{pq}, so the discrete structure of $\mathrm{VD}(P)$ does not change (and, therefore, no bisector event is recorded).

Corner event. A corner event occurs when there is a P-empty homothetic copy $Q' = Q[\nu_{pqr}]$ of Q with a corner placement of a vertex of $v \in Q$ at p and two other points q and r lie on two distinct edges of Q, none of which is incident to v. We refer to such an event as a *generic corner event of* p.

This event corresponds to a vertex ν_{pqr} of $\mathrm{VD}(P)$, an endpoint of an edge e_{pq}, coinciding with a breakpoint of b_{pq}. Then ν_{pqr}, also an endpoint of the Voronoi edge e_{pr}, coincides with a breakpoint of b_{pr} as well. By assumption (T3), one of the Q-Voronoi edges e_{pq} and e_{pr} gains a new edgelet and the other loses an edgelet at this event; see Figure 6 (b).

Flip event. A flip event occurs when there is a P-empty homothetic copy Q' of Q that touches four points p, p', q, q' at four distinct edges of Q', in this circular order along $\partial Q'$. By assumption (T4), up to a cyclic relabeling of the points, the Voronoi edge e_{pq} *flips* to a new Voronoi edge $e_{p'q'}$ at this event; see Figure 6 (c). Note that e_{pq} (resp., $e_{p'q'}$) is a non-corner edge immediately before (resp., after) the flip event, as both the vanishing edge e_{pq} and the newly emerging edge $e_{p'q'}$ are "too short" to have breakpoints near the event (this is a consequence of the kinetic Q-general position assumptions).

This completes the description of the changes in $\mathrm{VD}(P)$ at a generic event. We remark that a Voronoi edge newly appears or disappears only at a flip event, so, by definition, $\mathrm{DT}(P)$ changes only at a flip event. Suppose the Voronoi edge e_{pq} flips to the edge $e_{p'q'}$ at a flip event. Then p, q, p', q' are vertices of two adjacent triangles $\triangle pqp'$ and $\triangle pqq'$ immediately before the event, the edge pq of $\mathrm{DT}(P)$ flips to $p'q'$ at the event, and the Delaunay triangles $\triangle pqp', \triangle pqq'$ flip to $\triangle p'q'p, \triangle p'q'q$ at the event; again, see Figure 6 (c).

3.2 Singular events

Recall that a singular event occurs, at time t_0, if two points $p, q \in P$ lie on an edge $e_i = v_iv_{i+1}$ of a P-empty homothetic copy of Q. Hence, a *singular bisector event* (involving p and q) occurs at t_0. We may assume that neither p nor q is in P_∞ since in such a case the orientation of pq remains fixed

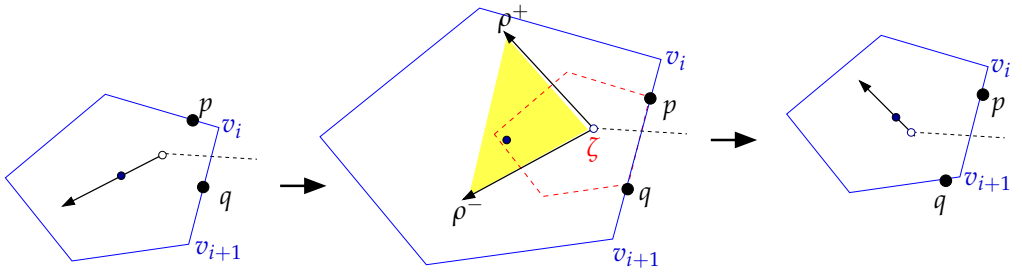

Figure 7. A singular bisector event. The ray ρ^- of b_{pq} is replaced instantly by the ray ρ^+, and the entire shaded wedge is part of b_{pq} at the event itself.

throughout the motion (namely, it is $o\vec{v}_i$) and, as we have assumed, different from the orientations of the edges of Q.

Changes in bisectors. Assume that pq becomes parallel to the edge e_i, and, without loss of generality, assume that $\vec{pq}$ and $v_i\vec{v}_{i+1}$ have the same orientation, as in Figure 7 (center). When this occurs, the set of placements Q' of Q at which both p and q touch e_i is a wedge W_0 whose boundary rays ρ^- and ρ^+ have respective directions $\vec{v_i o}$ and $\vec{v_{i+1} o}$, and whose apex ζ corresponds to the placement at which p and q touch v_i and v_{i+1} respectively; see Figure 7 (center).

Let t_0^- (resp., t_0^+) denote an instance of time immediately before (resp., after) t_0, so that no event occurs in the interval $[t_0^-, t_0)$ (resp., $(t_0, t_0^+]$). Then the terminal ray of b_{pq} that becomes the wedge W_0 at time t_0 is either in direction $\vec{v_i o}$ or $\vec{v_{i+1} o}$ at time t_0^-. Without loss of generality, throughout the present discussion of the singular event, we assume that this ray is in the direction $\vec{v_i o}$, i.e., it consists of all placements with e_i touching q and e_{i-1}, the other edge adjacent to v_i, touching p. This ray is parallel to ρ^- and approaches ρ^- as time approaches t_0; see Figure 7 (left). By assumption (T2), the bisector b_{pq} at time t_0^+ contains a terminal ray parallel to ρ^+, which consists of all placements with e_i touching p, and with the other edge e_{i+1} adjacent to v_{i+1} touching q. At time t_0 this ray coincides with ρ^+, which is clearly different from ρ^-. See Figure 7 (right). That is, the terminal ray of b_{pq} *instantly* switches from ρ^- to ρ^+ at time t_0.

Changes in $\mathrm{VD}(P)$**.** All topological changes in $\mathrm{VD}(P)$ at the time t_0 of a singular event occur on the boundaries of the Q-Voronoi cells $\mathrm{Vor}(p)$ and $\mathrm{Vor}(q)$. Since p and q are not in P_∞, both of these cells are *bounded*.

Refer to the state of $\mathrm{VD}(P)$ at times t_0^-, t_0, and t_0^+, as illustrated in Figure 8. For $t \in [t_0^-, t_0)$, let $\rho^-(t)$ be the edgelet of the bisector b_{pq} that is parallel to ρ^- at time t, and let $\eta^-(t)$ be the Voronoi vertex which is incident to $e_{pq} \cap \rho^-(t)$ and to some other pair of edges e_{pr^-} and e_{qr^-}. Similarly, for $t \in (t_0, t_0^+]$, let $\rho^+(t)$ be the edgelet of the bisector b_{pq} that is parallel to ρ^+ at time t, and let $\eta^+(t)$ be the Voronoi vertex which is incident to $e_{pq} \cap \rho^+(t)$ and to some other pair of edges e_{pr^+} and e_{qr^+}.

Let η^- (resp., η^+) be the limit of $\eta^-(t)$ (resp., of $\eta^+(t)$) as $t \uparrow t_0$ (resp., $t \downarrow t_0$). Alternatively, η^- is the center of a P-empty corner placement of Q that touches, at time t_0, p at v_i, and also touches q (at e_i) and r^-. As is easily checked, the edgelets e_{pr^-} and e_{qr^-} incident on η^- are collinear at time t_0, but just before that time they were separated by another short edgelet that has shrunk to a point; see Figure 8 (a). Similarly, η^+ is the center of a P-empty corner placement of Q that touches, at time t_0, q at v_{i+1}, and also touches p (at e_i) and r^+. Here too the edgelets of e_{pr^+} and e_{qr^+} incident on

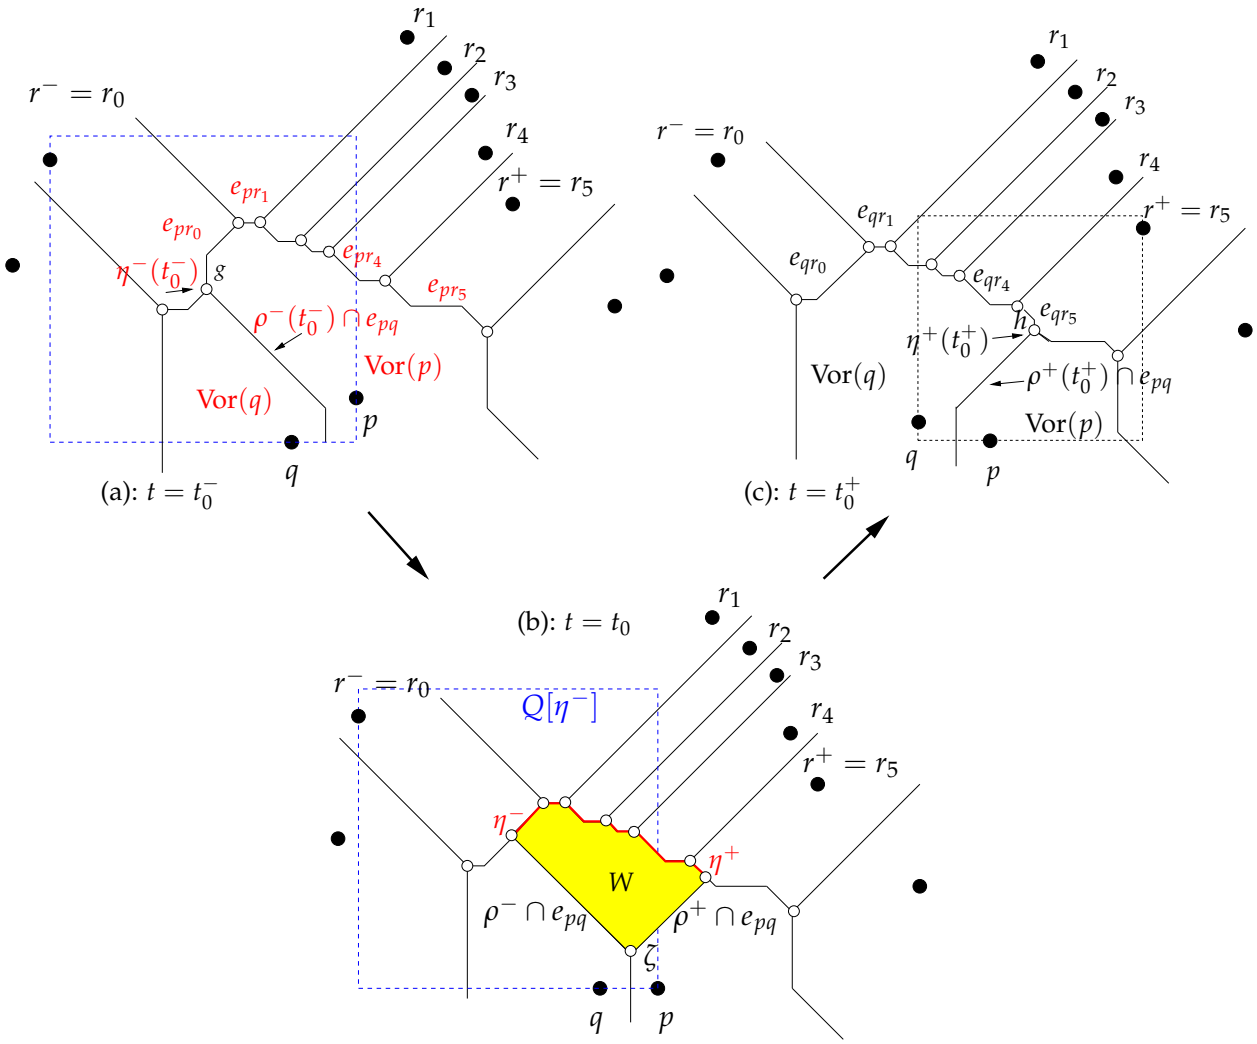

Figure 8. Changes in VD(P) at the time t_0 of a singular bisector event. (a): VD$(P(t_0^-))$; (b):VD$(P(t_0))$: the terminal ray of b_{pq} instantly switches from ρ^- to ρ^+ and the cell Vor(p) loses the entire region W to Vor(q); (c): VD$(P(t_0^+))$.

η^+ are collinear at t_0, and get separated from each other by a newly emerging short edgelet; see Figure 8 (c).

Let γ be the polygonal chain connecting η^- and η^+ at time t_0, consisting of all centers of placements touching p and q (at e_i) and some other point r_i of P. At time t_0 the degenerate Voronoi edge e_{pq} includes a two-dimensional star-shaped polygonal region, denoted by W, bounded by $\zeta\eta^-$, $\zeta\eta^+$, and γ. (It is a portion of the wedge W_0 discussed earlier.)

At time t_0 the terminal edgelet of e_{pq} instantly switches from $\zeta\eta^-$ to $\zeta\eta^+$, the Voronoi cell Vor(p) loses the entire region W to Vor(q). As a result, VD(P) can experience $\Omega(n)$ topological changes at time t_0 (in addition to the obvious change of the edgelet structure of e_{pq}, as just discussed).

Specifically, let $e_{pr^-} = e_{pr_0}, e_{pr_1}, \ldots, e_{pr_s} = e_{pr^+}$ be the edges of Vor(p) at times $t \in [t_0^-, t_0)$ that also overlap γ at t_0 (listed in the order their appearance along ∂Vor(p)). Right after time t_0, the point $\eta^+ = \rho^+ \cap e_{pr^+}$ becomes a new Voronoi vertex instead of $\eta^- = \rho^- \cap e_{pr^-}$.

Assuming $s \geq 1$ (or alternatively, $r^- \neq r^+$), every old edge e_{pr_j} of Vor(P), that is completely

contained in γ (such edges exist only if $s \geq 2$), is instantly *relabeled* as the new edge e_{qr_j} of $\mathrm{Vor}(q)$. The edge e_{pr^-}, which was incident at times $t \in [t_0^-, t_0)$ to $\eta^-(t)$, loses its portion within γ to the adjacent edge e_{qr^-}. Symmetrically, the old edge e_{pr^+} of $\mathrm{Vor}(p)$, which is hit by $\rho^+(t)$, for $t \in (t_0, t_0^+]$ is split at η^+ into e_{qr^+} (its portion within γ) and e_{pr^+} (its portion outside γ).

Changes in $\mathrm{DT}(P)$. If $s \geq 1$ then each of the old Delaunay edges pr_j, for $j = 0, \ldots, s-1$, flips in $\mathrm{DT}(P)$ to the new edge qr_{j+1}; see Figure 9. These are the only changes that $\mathrm{DT}(P)$ experiences at time t_0.

Singular sequences While we can regard the overall change in $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$ at a singular event as a single compound event, it is more convenient for our analysis, for the KDS implementation presented in Section 5, and perhaps also for a better intuitive perception of this change, to consider it as a sequence of individual separate singular corner and flip events, as we describe next.

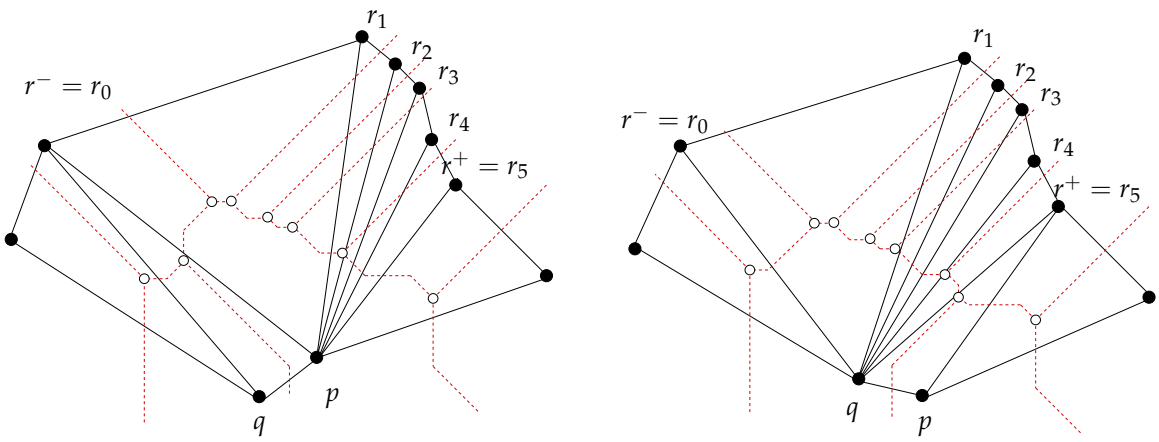

Figure 9. Topological changes in $\mathrm{DT}(P)$ at the time t_0 of a singular bisector event; five singular flip events occur at t_0, where $pr_0, \ldots, pr_4$ (left) flip to $qr_1, \ldots, qr_5$ (right).

To do so, we "stretch" the time t_0 and *continuously* rotate the terminal ray ρ of b_{pq} from ρ^- to ρ^+. Hence, the intersection point $\eta = \rho \cap \gamma$ traces γ from η^- to η^+; see Figure 10. At any given moment during this virtual rotation, ρ hits some old Voronoi edge e_{pr_j}, for $0 \leq j \leq s$, which is split by the current η into e_{pr_j} (the portion not yet swept by ρ) and e_{qr_j} (the swept portion). That is, we can interpret η as an instantaneous Q-Voronoi vertex v_{pqr_j} which is incident to three Voronoi edges, namely, e_{pq}, e_{pr_j} and e_{qr_j} (since e_{pq} is two-dimensional, this does not fix the vertex and indeed leaves it with one degree of freedom, of moving along the appropriate portion of γ). The topological structure of $\mathrm{VD}(P)$ changes (during the above rotation) only at the following two kinds of events, which closely resemble their generic counterparts in Section 3.1.

Singular corner event. This occurs when the Voronoi vertex $\eta = v_{pqr_j}$ coincides with a breakpoint of $\gamma \subseteq b_{pr_j} = b_{qr_j}$, occurring along the corresponding Voronoi edge. There are three types of singular corner events, depending on whether the corner placement occurs at p, q, or r_j for some $0 \leq j \leq s$. Each of the first two types occurs just once at a singular bisector event, whereas the third type may occur multiple times.

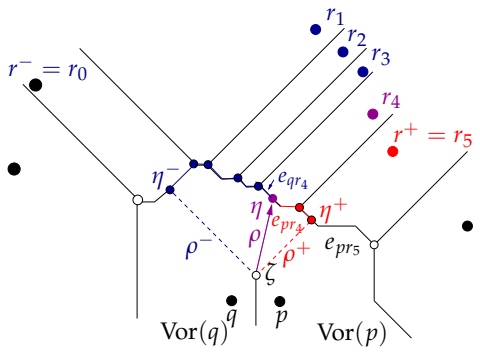

Figure 10. "Stretching" the time t_0. As the terminal ray ρ of b_{pq} continuously rotates from ρ^- to ρ^+, the vertex $\eta = \rho \cap \gamma$ traces γ from η^- to η^+. We only show the intersection of the rays ρ^-, ρ, and ρ^+ with e_{pq}. In the depicted snapshot ρ hits γ within the old edge e_{pr_4}, splitting it into e_{pr_4} and e_{qr_4}. Hence, η is the Voronoi vertex adjacent to Vor(p), Vor(q), and Vor(r_4).

(i) $\eta = \nu_{pqr_j}$ is the center of a corner placement at p along e_{pr_j}. This occurs at $\eta = \eta^-$, i.e., at the starting point of the rotation, for $Q[\eta^-]$ is indeed a corner placement at p. We refer to this event as the *initial* corner event of the singular sequence.

As already discussed, $Q[\eta^-]$ touches p at the vertex v_i, q on the edge v_iv_{i+1}, and the third point r^- at some other edge e^- (see Figure 8 (a)—the figure depicts what happens just before t_0; the limiting situation is depicted in Figure 8 (b)). Immediately after the singular corner event of p (in the sense of stretching the time during the virtual rotation of ρ), the Q-Voronoi edge e_{pr^-} loses an edgelet (denoted as g in that figure), namely the edgelet corresponding to p touching e_{i-1} and r^- touching e^-. It is interesting to note that, unlike at a generic corner event, e_{pq} does not gain an edgelet. (The only edgelet that it could have gained is the one that encodes the double contact of p and q at e_i, which is not a real edgelet.) Note also that the external edgelet of e_{qr^-} becomes aligned with the next edgelet of e_{pr^-}, and as η starts rotating, begins to "annex" it.

(ii) $\eta = \nu_{pqr_j}$ is the center of a corner placement at q along e_{qr_j}. This event occurs only at the end of the rotation, when $\rho = \rho^+$ and η coincides with the Voronoi vertex $\eta^+ = \nu_{pqr^+}$ newly created at this singular event. We refer to this event as the *final* corner event of the singular sequence.

The situation is symmetric to that in (i). Specifically, $Q[\eta^+]$ touches q at the vertex v_{i+1}, p at the edge v_iv_{i+1}, and the third point r^+ at some other edge e^+. Immediately after the singular corner event at q (as the "real" time t increases past t_0), the Q-Voronoi edge e_{qr^+} gains an edgelet (marked by h in Figure 8 (c)), but e_{pq} does not lose an edgelet. The external edges of e_{pr^+} and of e_{qr^+}, which were aligned as η approaches η^+, begin to shift apart from each other, with h in between them.

(iii) $\eta = \nu_{pqr_j}$ coincides with the center of a corner placement at r_j along the currently traced edge e_{pr_j} (see Figure 11). Let $Q[\eta] = Q[\nu_{pqr_j}]$ be the resulting corner placement of Q at r_j, which touches also p and q, both on the edge $e_i = v_iv_{i+1}$ of $Q[\eta]$. We refer to this event as an *intermediate* corner event of the singular sequence. Immediately after a singular corner event of r_j (again, in the sense of the stretched time), the Q-Voronoi edge e_{pr_j} loses an edgelet, and the adjacent Q-Voronoi edge e_{qr_j} gains an edgelet.

Singular flip event. The moving vertex $\eta = \nu_{pqr_j}$ coincides with a Q-Voronoi vertex $\nu_{pr_jr_{j+1}}$, which is the common endpoint of the Voronoi edges e_{pr_j} and $e_{pr_{j+1}}$ along γ (thinking of the scenario just

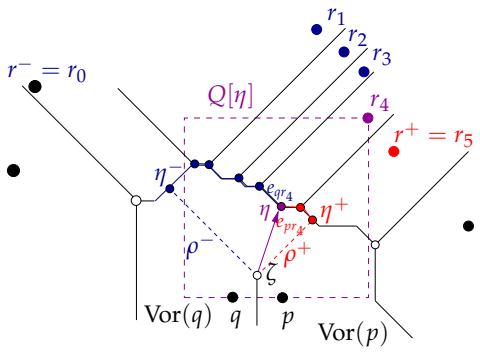

Figure 11. A singular corner event at r_4. A corner of $Q[\eta]$ touches r_4, so η coincides with a breakpoint on each of the overlapping bisectors b_{pr_4} and b_{qr_4}.

before t_0; see Figure 12). In the stretched time during the rotation of ρ, the Q-Voronoi edge e_{pr_j} shrinks (or, more precisely, "overtaken" by e_{qr_j}) and disappears from VD(P). The new edge $e_{qr_{j+1}}$ is born, as η starts moving along the old edge $e_{pr_{j+1}}$ of Vor(p), annexing a portion of it for Vor(q). The growing portion of that edge between $v_{qr_jr_{j+1}}$ and η becomes $e_{qr_{j+1}}$, and $e_{pr_{j+1}}$ is the shrinking remainder of that edge. Accordingly, the edge pr_j of DT(P) flips to qr_{j+1}.

Analogous to a generic flip event, e_{pr_j} is a non-corner edge when the above flip event occurs. It shrinks to a point at the event, and the four points p, q, r_j, r_{j+1}, which are Q-cocircular, are vertices of the two adjacent triangles $\triangle pr_jq$ and $\triangle pr_jr_{j+1}$ of DT(P) (sharing the edge pr_j). The event is singular because p and q are on the same edge of $Q[\eta]$ (with label e_i), and the remaining two points r_j and r_{j+1} are incident to some pair of other distinct edges of $Q[\eta]$.

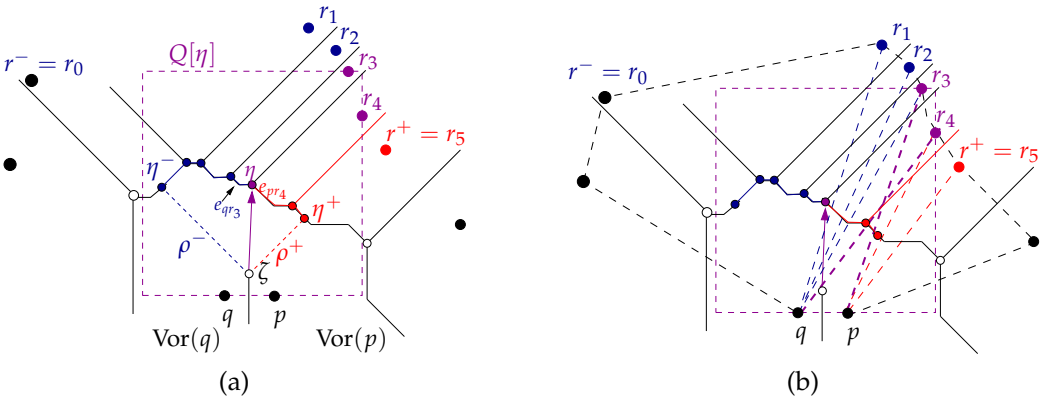

Figure 12. A singular flip event involving p, q, r_3 and r_4. (a): The Voronoi perspective. The vertex η coincides with the common endpoint $v_{pr_3r_4}$ of e_{pr_3} and e_{pr_4}. The corresponding copy $Q[\eta]$ touches p, q, r_3 and r_4 in this clockwise order. (b): The Delaunay perspective. The previously Delaunay edge pr_3 flips into qr_4.

We remark that, unless $j = 0$, the edge qr_j does not belong to DT(P) at time t_0^-. It becomes an edge of DT(P) only after executing the previous singular flip, which adds it to DT(P). In other words, the rotational order in which ρ generates these events is such that each flip event facilitates the next one, by introducing the appropriate edges qr_j into the diagram.

We refer to the entire sequence of events that are triggered by a singular bisector event (including the singular bisector event itself) as a *singular sequence*. The above order, in which the events of

a singular sequence are encountered during the continuous rotation of the terminal edgelet ρ, is called the *rotational order* of this sequence.

A crucial property of this interpretation of a singular event is that, in the stretched time, the various singular corner and flip events occur in the order at which η encounters the various breakpoints and vertices along γ. At each instance of the rotation, only one singular event, namely the next one in the sequence, is locally consistent with the current structure of the diagrams, in the sense that it can be effected by the appropriate local rule—a flip of Delaunay and Voronoi edges, or the passage of an edgelet from one Voronoi edge to another.

Remark. (i) If we allow Q to have several mutually parallel diagonals, then several interior edgelets (up to $\Theta(k)$ in the worst case) of b_{pq} are simultaneously replaced by new ones when pq becomes parallel to all of them. We treat each of them as a separate bisector event. Similarly, at a singular event, besides the changes discussed above, some interior edgelets of b_{pq} may also change (if the relevant edge of Q is parallel to some diagonals).

(ii) If the trajectories of P are not in general position, as defined above, we can perform symbolic perturbation, using the framework proposed by Yap [27], to simulate general position.

(iii) If we do not augment P with P_∞, the set of points at infinity, some of the edges of $\mathrm{DT}(P)$ are adjacent to only one triangle, i.e., the edges of the support hull of P (defined in the introduction, cf. Figure 1 (b,c)), and this set can change over time with the motion of P. Consequently, besides edge flips, edges need to be inserted into or deleted from $\mathrm{DT}(P)$. From the Voronoi diagram perspective, some of the Voronoi cells of P will be unbounded, and their status from bounded to unbounded will change over time, as their defining sites become or cease to be vertices of the support hull of P. Consequently, at some singular biscetor events involving a pair p, q, the two-dimensional region of b_{pq} (the polygonal region W in Figure 8) may be unbounded. Therefore extra care is needed while performing the rotational sweep to process the singular sequence of these events. Adding P_∞ simplifies and unifies all these special handlings.

4 Bounding the number of events

In this section we show that the overall number of events, at which the topological structure of $\mathrm{VD}(P)$ changes, during the motion of the points is only nearly quadratic in n. The bounds on the number of events hold even if the trajectories of points in P are not in Q-general position, but for simplicity, the proofs are presented under the general-position assumption. We can use a symbolic-perturbation based argument to extend the proof when trajectories are not in general position; see for example [24, Chapter 5].

The number of bisector events. The number of generic (resp., singular) bisector events is $O(k^2n^2)$ (resp., $O(kn^2)$) since the orientation of the segment pq for a given pair $p, q \in P$, can become parallel to a fixed diagonal v_iv_j (resp., a fixed edge $e_i = v_iv_{i+1}$) of Q only $O(1)$ times, as implied by our kinetic general position assumption.

Each time the orientation of pq becomes parallel to a fixed diagonal or edge of Q there are $O(1)$ topological changes in the structure of e_{pq}. Therefore, all the Q-Voronoi edges e_{pq} undergo a total of $O(k^2n^2)$ changes in their edgelet structure. Note that we only count the bisector events themselves in the $O(kn^2)$ bound on the number of singular bisector events, and not the singular sequences of

corner and flip events associated with them; these other events will be bounded separately in what follows.

The number of corner events. The number of corner events (which can be either generic or singular) is bounded by the following lemma.

Lemma 4.1. *Let Q be a convex k-gon, and let P be a set of n moving points in $\mathbb{R}^2$ along algebraic trajectories of bounded degree. The overall number of corner events (generic and singular) in $\mathrm{DT}(P)$ is $O(k^2 n \lambda_r(n))$, where r is a constant that depends on the degree of the motion of the points of P.*

Proof. Fix a point p and a vertex v_i of Q, and consider all the corner events in which (for an appropriate homothetic copy of Q) the vertex (with label) v_i touches p. As noted above, at any such event the center c of Q lies on the ray $u_i[p]$ emanating from p in direction $v_i o$ (where o denotes the center of Q). Note that, since p is moving, $u_i[p]$ is a moving ray, but its orientation remains fixed. For every other point $q \in P \setminus \{p\}$, let $\varphi_i[p,q](t)$ denote the distance, at time t, from p along $u_i[p]$ to the center of a homothetic copy of Q that touches p (at v_i) and q.

Each function $\varphi_i[p,q](t)$ is well defined only when q lies strictly in the wedge between two rays emanating from p, one in the direction $v_i\vec{v}_{i-1}$ and the other in the direction $v_i\vec{v}_{i+1}$. Since the segment pq becomes parallel to e_{i-1} or to e_i only $O(1)$ times, the domain in which $\varphi_i[p,q](t)$ is defined consists of $O(1)$ open intervals.

The value $\min_q \varphi_i[p,q](t)$ represents the intersection of $\partial\mathrm{Vor}(p)$ with $u_i[p]$ at time t (recall that $\mathrm{Vor}(p)$ is the Voronoi cell of p in $\mathrm{VD}(P)$). The point q that attains the minimum defines the Voronoi edge e_{pq} of $\mathrm{Vor}(p)$ that $u_i[p]$ intersects. At appropriate discrete times, the minimum may be attained by more than one point q, and then $u_i[p]$ hits a vertex of $\mathrm{Vor}(p)$.

In other words, we have a collection of $n-1$ partially defined functions $\varphi_i[p,q]$, for $q \in P \setminus \{p\}$, and the breakpoints of their lower envelope (where two different functions attain the envelope simultaneously) represent the corner events that involve the contact of v_i with p.

By our assumption on the motion of P, each function $\varphi_i[p,q]$, within each interval of its domain, is piecewise algebraic, with $O(k)$ pieces. Each piece encodes a continuously varying contact of q with a specific edge of Q, and has constant description complexity. Hence (see, e.g., [24, Corollary 1.6]) the complexity of the envelope is at most $O(k\lambda_r(n))$, for an appropriate constant r that depends on the degree of the motion of the points of P. Repeating the analysis for each point p and each vertex v_i of Q, we obtain the bound asserted in the lemma.

To complete the proof, we also need to take into account the initial and the final corner events of each singular sequence. The number of these events, though, is subsumed by the bound asserted in the lemma, and the proof is completed. □

The number of flip events. We distinguish between generic flips and their singular counterparts, handling the singular flip events first.

Lemma 4.2. *Let Q be a convex k-gon, and let P be a set of n moving points in $\mathbb{R}^2$ along algebraic trajectories of bounded degree. The number of singular flip events in $\mathrm{DT}(P)$ is $O(k^3 n \lambda_r(n))$, where r is the same constant as in the statement of Lemma 4.1.*

Proof. We choose three arbitrary edges e_1, e_2, e_3 of Q, and analyze the number of singular flips at which two points of P touch e_1 and one point of P lies on each of e_2 and e_3. The overall bound

on the number of singular flips is obtained by repeating our analysis for the $O(k^3)$ such triples of edges.

Let $\tilde{Q}$ be the convex hull of e_1, e_2 and e_3, with at most six vertices. Every singular flip event with respect to Q that involves e_1, e_2, and e_3 in the manner stated in the preceding paragraph is also a singular flip event with respect to $\tilde{Q}$. To bound the number of singular flip events for $\tilde{Q}$, we first show how to charge each of them to a singular corner event (again, with respect to $\tilde{Q}$) so that each corner event is charged only $O(1)$ times, and then make use of Lemma 4.1.

Consider the homothetic copy of $\tilde{Q}$ associated with a singular flip involving four points $p, q, r, s \in P$; we abuse the notation slightly, as we did earlier, and refer to this copy also as $\tilde{Q}$. Suppose, more concretely, that p and q touch e_1, r touches e_2, and s touches e_3. We continuously slide $\tilde{Q}$ so that it maintains its contacts with p, q at e_1, and with r at e_2, and moves[4] away from s. Specifically, the segments pr, qr partition $\tilde{Q}$ into three parts: the portion that touches s, the triangle $\triangle pqr$, and the third complementary portion (bounded by pr). During the continuous motion of $\tilde{Q}$, the part that was previously incident to s shrinks, the triangle $\triangle pqr$ does not change (as a portion of $\tilde{Q}$), and the third part expands; see Figure 13 (a). The motion of $\tilde{Q}$ stops when one of the following situations is reached:

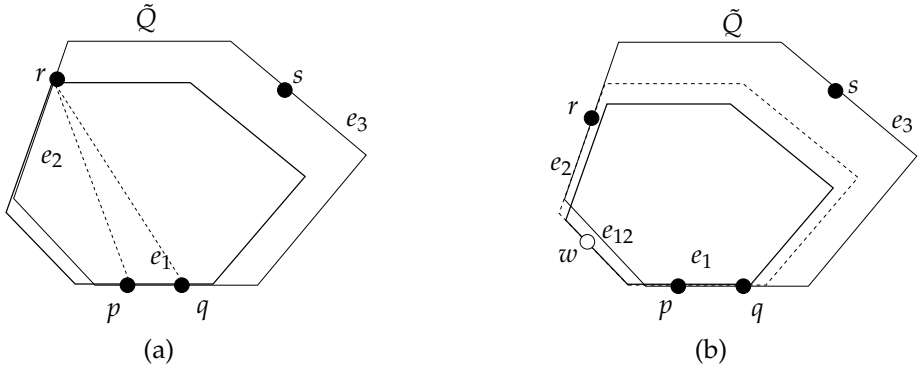

Figure 13. Charging a singular flip event. We slide $\tilde{Q}$ while maintaining the contacts with p, q, r: (a) one of the points runs into a vertex, (b) $\partial\tilde{Q}$ touches a new point. Thick polygons denote the corner event to which the flip event is charged.

(i) *One of the points p, q, r runs into a vertex of $\tilde{Q}$*; see Figure 13 (a) (where r is that point): This is a singular corner event of p, q or r, which can be charged for our singular flip event. This is because both events occur simultaneously, and any corner event is charged only $O(1)$ times in this manner (reversing the sliding process, the fourth point s is the first new point that the expanding part of $\tilde{Q}$ hits).

(ii) *The boundary of the expanding third portion of $\tilde{Q}$ touches a new point w of P*; see Figure 13 (b). In this case e_2 cannot be adjacent to e_1, because then there would be no expanding portion of $\tilde{Q}$, and no new point could be reached. Let e_{12} be the (unique) edge of $\tilde{Q}$ joining e_1 and e_2. Since e_3 belongs to the shrinking portion of $\tilde{Q}$, the boundary of the expanding portion of $\tilde{Q}$ consists of e_{12} and of portions of e_1, e_2, and thus w must lie on e_{12}; see Figure 13 (b). We now slide $\tilde{Q}$ along its contacts with p, q, and w and away from r. Since e_{12} is adjacent to e_1, the motion now stops, as in case (i), when one of p, q, w reaches a vertex of $\tilde{Q}$. We charge the flip event to this corner event. There are $O(1)$ possible ways to backtrack from the resulting corner event to the original singular flip (first

[4]Note that our continuous motion argument refers to the *stationary* point set $P = P(t_0)$, where t_0 is the time of the singular flip event. Thus we move $\tilde{Q}$ (for the purpose of the analysis) but the points of P remain fixed.

by expanding until r is reached, and then by expanding again until s is reached), which occurs simultaneously with it. Hence, each corner event is charged in this manner only $O(1)$ times.

By Lemma 4.1, the number of singular corner events with respect to $\tilde{Q}$ is $O(n\lambda_r(n))$, so the number of singular flip events with respect to $\tilde{Q}$ is also $O(n\lambda_r(n))$. Multiplying this bound by the number $O(k^3)$ of triples of edges e_1, e_2, e_3 yields the bound asserted in the lemma. □

Lemma 4.3. *Let Q be a convex k-gon, and let P be a set of n moving points in $\mathbb{R}^2$ along algebraic trajectories of bounded degree. The number of generic flip events in $\mathrm{DT}(P)$ is $O(k^4 n\lambda_r(n))$, where r is the same constant as in the statement of Lemma 4.1.*

Proof. Each generic flip event involves a placement of an empty homothetic copy Q' of Q that touches simultaneously four points p_1, p_2, p_3, p_4 of P, in this counterclockwise order along $\partial Q'$, so that the Voronoi edge $e_{p_1p_3}$, which was a non-corner edge before the event, shrinks to a point and is replaced by the newly emerging non-corner edge $e_{p_2p_4}$ right after the event. Let e_i denote the edge of Q' that touches p_i, for $i = 1, 2, 3, 4$. Since this is a generic flip, the e_i's are distinct.

We fix the quadruple of edges e_1, e_2, e_3, e_4, analyze the number of flip events involving a quadruple contact with these edges, and sum the bound over all $O(k^4)$ choices of four edges of Q. For a fixed quadruple of edges e_1, e_2, e_3, e_4, we replace Q by the convex hull $\tilde{Q}$ of these edges, which is at most octagonal, and note that any flip event of Q involving these four edges is also a flip event for $\tilde{Q}$. We therefore restrict our attention to $\tilde{Q}$, which is a convex k_0-gon, for some $k_0 \leq 8$, and argue that the number of generic flip events for $\tilde{Q}$, with contacts at the above four edges $e_1, \ldots, e_4$, is $O(n\lambda_r(n))$.

We note that if pq is a Delaunay edge containing an edgelet which is the locus of centers of placements of copies $\tilde{Q}'$ of $\tilde{Q}$ where p and q touch two *adjacent* edges of $\tilde{Q}'$, then pq must be a corner edge. Indeed, shrinking $\tilde{Q}'$ towards the vertex common to the two edges, so that it continues to touch p and q, will keep it empty, and eventually reach a placement where either p or q touches a corner of $\tilde{Q}'$.

Consider the situation just before the critical event takes place, as depicted in Figure 14 (a). The $\tilde{Q}$-Voronoi edge $e_{p_1p_3}$ (to simplify the notation, we write this edge as e_{13}, and similarly for the other edges and vertices in this analysis) is delimited by two $\tilde{Q}$-Voronoi vertices ν_{123} and ν_{143}, so that $\tilde{Q}[\nu_{123}]$ touches p_1, p_2, p_3 at the respective edges e_1, e_2, e_3, and $\tilde{Q}[\nu_{143}]$ touches p_1, p_4, p_3 at the respective edges e_1, e_4, e_3.

Let T be the maximum connected component of the union of the non-corner $\tilde{Q}$-Voronoi edges that contains e_{13}. We claim that T is a tree with at most five edges; specifically, it can include only e_{13} itself, e_{12}, e_{23} (the edges adjacent to ν_{123}), and e_{14}, e_{34} (the edges adjacent to ν_{143}); see Figure 14(b). To see this, consider, for specificity, the neighboring edge e_{12} incident to the vertex ν_{123}, and assume that e_{12} is a non-corner edge (otherwise it does not belong to T, by construction), so e_{12} is part of T. As we move the center of $\tilde{Q}$ along that edge away from ν_{123}, $\tilde{Q}$ loses the contact with p_3, shrinks on the side of p_1p_2 that contains p_3 (and p_4, already away from $\tilde{Q}$), and expands on the other side. Being a non-corner edge, the other endpoint of e_{12} is a placement at which the (artificial) edge e_{12} of $\tilde{Q}$ between e_1 and e_2 touches another point p_5 (while p_1 and p_2 continue to touch e_1 and e_2, respectively). Since e_{12} is adjacent to both edges e_1, e_2, the new Voronoi edges e_{15} and e_{25} must both be corner edges. The same argument applies for the other three Voronoi edges adjacent to e_{13}, and we conclude that T is a tree consisting of at most five edges.

We note that as long as no discrete change occurs at any of the surrounding corner edges of T, it can undergo only $O(1)$ discrete changes because all its edges are defined by a total of $O(1)$ points

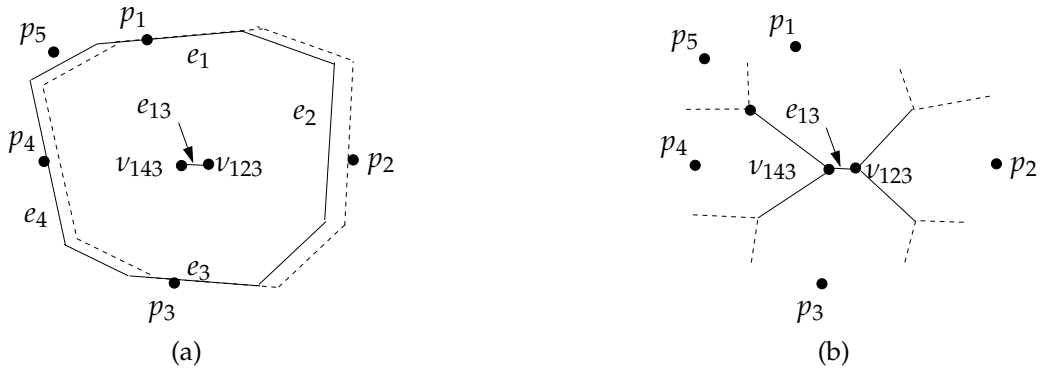

Figure 14. (a) The edge e_{13} in the diagram $\mathrm{VD}(P)$ before it disappears. The endpoint ν_{123} (resp., ν_{143}) of e_{13} corresponds to the homothetic copy of $\tilde{Q}$ whose edges e_1, e_2, e_3 (resp., e_1, e_4, e_3) are incident to the respective points p_1, p_2, p_3 (resp., p_1, p_4, p_3). (b) The tree of non-corner edges that contains e_{13}. In the "worst-case" scenario depicted here, all five solid edges belong to the tree, and the dashed edges are all corner edges.

of P (and the size of $\tilde{Q}$ is independent of k). Furthermore when a corner edge undergoes a discrete change, this can affect only $O(1)$ adjacent non-corner trees of the above kind.

Hence, the number of changes in non-corner edges, and in particular the number of flip events they are involved in, is proportional to the number of corner events, because only corner events affect the actual number of breakpoints on an edge. More specifically, if e is an edge adjacent to T, its interaction with T can change only when its last edgelet (the one nearest to T) changes, and that can happen only at a corner event. By Lemma 4.1 and the reasoning preceding it (applied to $\tilde{Q}$), the number of corner events is $O(n\lambda_r(n))$ (the same r as in the previous lemmas). Multiplying by the $O(k^4)$ choices of quadruples of edges of Q, we conclude that the total number of generic flip events is $O(k^4 n\lambda_r(n))$. □

Combining the above three lemmas, we obtain the following summary result.

Theorem 4.4. *Let P be a set of n moving points in $\mathbb{R}^2$ along algebraic trajectories of bounded degree, and let Q be a convex k-gon. Then the number of topological changes in $\mathrm{VD}(P)$ (and in $\mathrm{DT}(P)$) with respect to Q, is $O(k^4 n\lambda_r(n))$, where r is a constant that depends on the degree of the motion of the points of P.*

5 A KDS for kinetic maintenance of $\mathrm{VD}(\mathbf{P})$ and $\mathrm{DT}(\mathbf{P})$

In this section we turn the detailed analysis of the kinetic behavior of $\mathrm{VD}(P)$ and $\mathrm{DT}(P)$, as studied in Section 3, into an efficient algorithm for their kinetic maintenance, using the KDS framework of Basch *et al.* [5]. In particular, as the points of P move continuously, $\mathrm{DT}(P)$ and $\mathrm{VD}(P)$ are maintained using a collection of local *certificates* (Boolean predicates) that assert the correctness of the topological structure of the current Voronoi diagram (and thus also of the Delaunay triangulation). When one of these certificates fails, one of the events discussed in Section 3 takes place and the KDS repair mechanism updates $\mathrm{VD}(P)$, $\mathrm{DT}(P)$, and the certificates. A global *event queue* is used to determine when the next certificate fails. For simplicity, the KDS is described under the assumptions that the orientations of all edges and diagonals of Q are distinct, the trajectories of points in P are in Q-general position (i.e., satisfy (T1)–(T5)), and P is augmented with the points of P_∞. In cases where these assumptions do not hold, we apply an infinitisimal

symbolic perturbation of the vertices of Q, thereby putting it in general position, and ensuring the distinctness of the orientations of all the edgese and diagonals of Q. With suitable (minor) adjustments, this will allow our KDS to operate in such degenerate situations too.

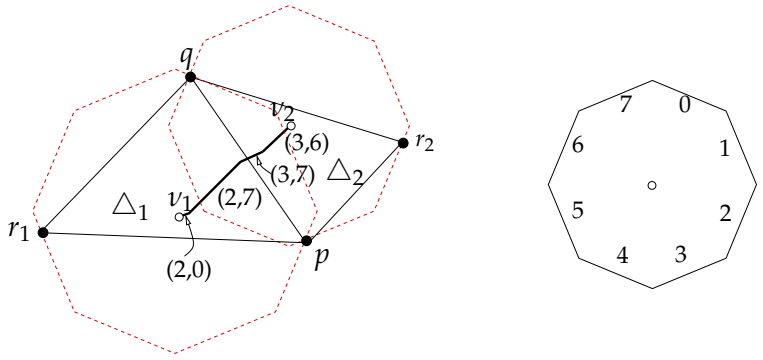

Figure 15. A Delaunay edge pq and its two adjacent triangles $\triangle_1, \triangle_2$. We have $\delta(v_1, p) = 2$, $\delta(v_1, q) = 0$, $\delta(v_2, p) = 3$, and $\delta(v_2, q) = 6$, using the edge numbering given on the right, where v_i is the Voronoi vertex corresponding to $\triangle_i$, for $i = 1, 2$. The edgelets appearing in e_{pq} are labeled $(2, 0)$, $(2, 7)$, $(3, 7)$, and $(3, 6)$.

We need the following definitions and notation to describe the KDS. For a Voronoi vertex ν of some Voronoi region $\mathrm{Vor}(p)$, where p is a point of P not lying at infinity (so p is a vertex of the Delaunay triangle $\triangle_\nu$ corresponding to ν), we define $\delta(\nu, p)$ to be the edge of $Q[\nu]$ that touches p (at a generic time instance, p does not lie at a vertex of $Q[\nu]$); see Figure 15.

Let e_{pq} be a Voronoi edge with endpoints ν_1 and ν_2. Let g_1 (resp., g_2) be the edgelet of b_{pq} that contains the endpoint v_1 (resp., v_2). Writing $\psi(g_1) = (i_1, j_1)$ and $\psi(g_2) = (i_2, i_2)$, we have $\delta(\nu_1, p) = i_1$, $\delta(\nu_2, p) = i_2$, $\delta(\nu_1, q) = j_1$, and $\delta(\nu_2, q) = j_2$. That is, $\delta(\nu_1, p)$ and $\delta(\nu_1, q)$ together encode the edgelet of b_{pq} containing ν_1, which is an external edgelet of e_{pq}, and an analogous property holds for ν_2, so the mapping $\delta(\cdot, \cdot)$ encodes the external edgelets for all edges of $\mathrm{VD}(P)$. If e_{pq} is a non-corner edge of $\mathrm{VD}(P)$, then $g_1 = g_2$, and therefore $\delta(\nu_1, p) = \delta(\nu_2, p)$ and $\delta(\nu_1, q) = \delta(\nu_2, q)$.

We focus on the more detailed representation of $\mathrm{VD}(P)$; $\mathrm{DT}(P)$ is easier to represent, and for the ongoing analysis we regard it merely as the dual of $\mathrm{VD}(P)$. We assume for now that $\mathrm{VD}(P)$ is represented *explicitly*; that is, each Voronoi edge is represented as a sequence of its edgelets. We represent $\mathrm{VD}(P)$ (and $\mathrm{DT}(P)$ too, if so desired) using the so-called *doubly connected edge list* (DCEL) data structure [6], a commonly used structure for representing planar subdivisions. It stores a record for each face, edgelet, and vertex (a Voronoi vertex or a breakpoint) of the subdivision, along with some adjacency information. The basic entities in this representation are *half-edges*, where each half-edge is an oriented copy of some Voronoi edgelet. Each half-edge stores pointers to its endpoints (each of which is a breakpoint or a Voronoi vertex), the Voronoi cells adjacent to it, its neighboring edgelets, and its label (i, j). Each Voronoi cell $\mathrm{Vor}(p)$ is associated with its input point p, which stores its trajectory. With each breakpoint or vertex ν we can either store its trajectory, or compute it on the fly from the trajectories of the two or three points defining ν, and the edges of $Q[\nu]$ that they touch (we can determine the identity of these edges of $Q[\nu]$ from the labels of the edgelets adjacent to ν). Note that as long as no event happens, the trajectories of each vertex and breakpoint are algebraic of constant complexity. See [6] for more details on the DCEL data structure. This explicit representation of $\mathrm{VD}(P)$ requires $O(nk)$ space.

5.1 Certificates

We maintain three types of certificates, each corresponding to a different type of event. Each certificate is stored at the event queue, with its next future failure time as its priority. In addition, we store with each certificate information that will be needed to update the diagram and the event queue when the certificate fails.

Bisector certificates. For each internal edgelet g of each Voronoi edge e_{pq}, with $\psi(g) = (i, j)$, we maintain a certificate asserting that the orientation of pq lies between the two extreme orientations of lines that cross both edges e_i and e_j of Q (these are the orientations of $v_i v_j$ and $v_{i+1} v_{j+1}$). We recall that i and j are not consecutive for internal edgelets. This certificate fails when $\vec{pq}$ attains one of these extreme orientations; See Figure 16. The failure causes g to shrink and disappear and be replaced by a new internal edgelet g' with $\psi(g') = (i-1, j-1)$ or $\psi(g') = (i+1, j+1)$, as appropriate. To efficiently perform the required updates when this certificate fails, the certificate also stores a direct pointer to the edgelet itself (in the DCEL). We call this certificate a *generic bisector certificate*.

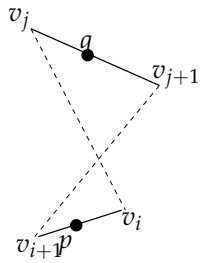

Figure 16. The generic bisector certificate corresponding to an internal edgelet g such that $\psi(g) = (i, j)$. It fails when $\vec{pq}$ becomes parallel to $\vec{v_i v_j}$ or $\vec{v_{i+1} v_{j+1}}$

External edgelets of an edge e_{pq} come in two types: edgelets that are portions of non-terminal (bounded) edgelets of the bisector b_{pq}, and edgelets that are portions of the terminal rays of b_{pq}; see Figure 3. Our general position assumption (T5) implies that we do not have to associate bisector events with edgelets of the former type, because before such an edgelet shrinks to a point, as an edgelet of b_{pq}, it will disappear as an edgelet of e_{pq}, at a corresponding corner event, if e_{pq} is a corner edge, or at a flip event, if e_{pq} is a non-corner edge.

Consider then external edgelets of the latter type. Each such edgelet g has a label of the form $\psi(g) = (i, i+1)$, and its bisector certificate fails when the corresponding points p and q are such that $\vec{pq}$ becomes parallel to either e_i or e_{i+1}. This is a singular bisector event. We refer to such a certificate as a *singular bisector certificate*.

(Voronoi) Vertex certificates. For each vertex ν of $\mathrm{VD}(P)$ not lying at infinity and for each point $p \in P$ such that $\mathrm{Vor}(p)$ is adjacent to ν and p is not at infinity, we maintain a certificate asserting that $\delta(\nu, p) = v_i v_{i+1}$ for a suitable index i, i.e., p lies on the edge $v_i v_{i+1}$ of the homothetic copy $Q[\nu]$ of Q. Because the motion of p is continuous, the certificate fails only when p coincides with v_i or v_{i+1} in $Q[\nu]$; that is, the failure occurs at a corner event. This certificate stores pointers to ν, p. The index i and $\delta(\nu, p)$ can be identified using the labels of the edgelets adjacent to ν in the DCEL.

Let q and r be the other two points of P defining ν; i.e. $\nu = \nu_{pqr}$. The corner event is either generic, when the vertex of Q with which p coincides is not adjacent to $\delta(\nu, q)$ or to $\delta(\nu, r)$, or singular, when this vertex is adjacent to one of these edges, say to $\delta(\nu, q)$. In the latter case, this is

the first singular corner event in the singular sequence associated with the corresponding singular bisector event involving p and q at which $\vec{pq}$ is parallel to $\delta(\nu, q)$. When the singular corner event happens the placement $Q[\nu]$ is a placement at which p lies at the common vertex of $\delta(\nu, p)$ and $\delta(\nu, q)$, and q lies on the adjacent edge $\delta(\nu, q)$. In the terminology of Section 3.2, $\nu = \eta^-$; see there for more details. We refer to a vertex certificate that fails at a generic corner event as a *generic vertex certificate* and to a vertex certificate that fails at a singular corner event as a *singular vertex certificate*.

For any sequence of singular events, only the initial corner event is detected through the failure of a singular vertex certificate; the remaining corner events are detected and treated separately, as described below.

(Voronoi) Edge certificates. For each non-corner Voronoi edge e_{pq}, we maintain the following certificate. Let ν_{pqr} and ν_{pqw} be the two endpoints of e_{pq}, i.e., $\triangle pqr$ and $\triangle pqw$ are the two Delaunay triangles adjacent to the edge pq. We maintain the certificate asserting that e_{pq} exists, namely, that $\nu_{pqr} \neq \nu_{pqw}$; since the trajectories of these vertices are available, and each of them is of constant complexity (as long as no other event affecting these vertices occurs), we can compute the failure time of this certificate in $O(1)$ time. The edge certificate fails when $\nu_{pqr} = \nu_{pqw}$ (while maintaining $\delta(\nu_{pqr}, p) = \delta(\nu_{pqw}, p) \neq \delta(\nu_{pqr}, q) = \delta(\nu_{pqw}, q) \neq \delta(\nu_{pqr}, r) \neq \delta(\nu_{pqw}, w)$), that is when p,q,r, and w become Q-circular while touching distinct edges of Q. This is a generic flip event.

Singular flip events, occurring as part of a compound singular bisector event, as described in Section 3.2, are detected and treated separately, as described below.

5.2 Handling certificate failures

We now describe the KDS repair mechanism that updates $\mathrm{VD}(P)$, $\mathrm{DT}(P)$, the certificates, and the event queue, when a certificate fails. The queue contains the failure times of generic bisector and vertex certificates, edge certificates, and singular bisector and vertex certificates. Note that singular certificates come in pairs, each consisting of a singular bisector certificate and a singular vertex certificate, both failing *at the same time*. When we extract from the queue one certificate of such pair, we immediately extract its sibling too, and process them together, as will be described later. We note that, except for this kind of simultaneous failure of two distinct certificates, assumption (T5) implies that there are no other duplications of failure times in the queue.

The processing of the failure of a certificate, when it is extracted from the queue, is handled according to its type, and proceeds as follows.

Generic bisector certificates. Consider such a certificate involving the shrinking of some edgelet g of some Voronoi edge e_{pq}. We remove g from $\mathrm{VD}(P)$ and replace it by a new (internal) edgelet g'. We compute $\psi(g')$ from $\psi(g)$, by adding or subtracting 1 from its two indices, as appropriate, generate the new bisector certificate for g', compute its failure time, and insert it into the queue.

Generic vertex certificates. Let $\nu := \nu_{pqr}$ be a Voronoi vertex such that one of its certificates, say $\delta(\nu, p) = e_i$, fails, and suppose that this is a generic vertex certificate. This happens when p coincides with one of the endpoints of e_i, say, v_i, in the homothetic copy $Q[\nu]$ of Q (and $\delta(\nu, q)$ and $\delta(\nu, r)$ are not adjacent to v_i). Then ν lies at a breakpoint of e_{pq} and at a breakpoint of e_{pr}, and we detect a generic corner event, at which, without loss of generality, e_{pq} loses an edgelet g and e_{pr} gains an edgelet g'. We remove g from the DCEL, including its adjacent breakpoint (which is "overtaken" by ν), create the new edgelet g' on e_{pr} and the corresponding new breakpoint delimiting it, and update the DCEL accordingly. We next update the certificates as follows.

(i) We find the new edge e' of Q (which is either e_{i-1} or e_{i+1}) that supports p after the event, replace the old certificate with a new certificate asserting that $\delta(\nu, p) = e'$, compute the failure time of the new certificate, and update the event queue.

(ii) Since p lies on a new edge of $Q[\nu]$ after the event, the algebraic representation of the trajectory of ν changes, and thus we also update the failure times of the vertex certificates that specify $\delta(\nu_{pqr}, q)$ and $\delta(\nu_{pqr}, r)$.

(iii) The new edgelet g' is an external edgelet of e_{pr}. If its label consists of two consecutive edges of Q then g' generates a singular bisector certificate, which we process into the queue (otherwise, as noted earlier, no action is needed here).

(iv) The edgelet of e_{pr} adjacent to g' may now become internal, and then we generate the corresponding generic bisector certificate. Otherwise e_{pr} was a non-corner edge before the event and becomes now a corner edge, so we remove the edge certificate associated with e_{pr} from the queue.

(v) The edgelet g^+ of e_{pq} adjacent to g becomes external. We first remove the generic bisector certificate of g^+. Next, if e_{pq} becomes a non-corner edge, then we generate a new edge certificate associated with e_{pq}.

(vi) If the third Voronoi edge e_{qr} adjacent to ν is a non-corner edge then its certificate failure time has to be updated, because of the new trajectory of ν.

Edge certificates. If a non-corner edge e_{pq}, with endpoints ν_{pqr} and ν_{pqw}, shrinks to a point, i.e., p, q, r, w become Q-cocircular, then the edge certificate associated with e_{pq} fails. We replace in $\mathrm{VD}(P)$ the non-corner edge e_{pq} with the non-corner edge e_{rw} and the vertices ν_{pqr} and ν_{pqw} with ν_{rwp} and ν_{rwq} (these are the endpoints of the new edge). The edge pq in $\mathrm{DT}(P)$ is flipped to rw, and $\triangle pqr, \triangle pqw$ are replaced with $\triangle rwp$ and $\triangle rwq$.

Next, we remove the six vertex certificates associated with the vertices ν_{pqr} and ν_{pqw}. We add the edge certificate corresponding to the edge e_{rw} and the six vertex certificates corresponding to the vertices ν_{rwp} and ν_{rwq}. Finally, since the endpoints of the edges e_{pr}, e_{qr}, e_{pw}, and e_{qw} have changed, if any of them is a non-corner edge then we also update the certificate corresponding to that edge, recompute its failure time, and update the queue.

Singular bisector and vertex certificates. Finally we turn to the processing of a pair of a singular bisector certificate and a singular initial corner certificate that fail simultaneously. Let p and q be the points generating the bisector event, such that at the corresponding corner event p touches Q at a vertex v_i and q lies on, say, the edge v_iv_{i+1}. Following the notations of Section 3, we denote by $\eta^- = \nu_{pqr^-}$ the Voronoi vertex at which the corner event occurs, where r^- is the third point defining this event, and by ρ^- the terminal ray of b_{pq} that contains η^-.

Let g be the edgelet of e_{pr^-} that disappears at the event. We remove g from the DCEL structure, and remove its two endpoints, one of which is η^-. We also remove the external edgelet of e_{pq} contained in ρ^- and the certificates associated with these features. Recall (see Section 3.2) that the next edgelet g^+ of e_{pr^-} and the external edgelet g^- of e_{qr^-} adjacent to η^- were parallel before the event and become aligned after the event.

We now execute a process that effectively simulates the rotational sweep of ρ from ρ^- to ρ^+, as described in Section 3.2, continuously traversing the corresponding portion γ of $\partial\mathrm{Vor}(p)$ between η^- and η^+ with the endpoint $\eta = \rho \cap \gamma$ of e_{pq}, and transferring this portion to $\mathrm{Vor}(q)$. Specifically, we iterate over the edges and edgelets of γ in order, starting with the edgelet g^+ following g on e_{pr^-}, and do the following. Let f be the edgelet currently traversed by η, of some old Voronoi edge e_{pr}. Notice that e_{pr} is now split by $\eta = v_{pqr}$ into e_{qr}, the portion already traversed by ρ, and the

remainder e_{pr}, still labeled as such. We test whether $\eta^+ \in f$, that is, whether the terminal corner event, where q lies at v_{i+1}, occurs before f ends.

Suppose first that $\eta^+ \notin f$. In this case the entire f becomes an edgelet of the new edge e_{qr} when η coincides with the other endpoint ν of f. If ν is a breakpoint, this corresponds to a singular corner event. Otherwise, ν is an old Voronoi vertex $\nu_{prr'}$, so its coincidence with η indicates a singular flip event of e_{pr} to $e_{qr'}$. In both cases, we replace the old bisector certificate associated with f, if there was one, by a corresponding certificate in which q replaces p. In case of a singular flip event, we replace it by $\nu_{qrr'}$ and update its vertex certificates. In addition, if the old edge e_{pr} was non-corner, we also replace its edge certificate by the corresponding edge certificate in which q replaces p again.

A somewhat special treatment applies to the first edgelet g^+ in the sequence. If $\eta^+ \notin g^+$ then we remove g^+, as it now merges with the last edget g^- of e_{pr^-} into a common larger edgelet, and we replace the endpoint of g^- (which was η^-) by the endpoint of g^+. The handling of this endpoint, if it is a vertex, is done exactly as just described. Finally, if e_{qr^-} was a non-corner edge, we remove its certificate from the queue as it is now obsolete. If in addition g^+ ends at a Voronoi vertex, e_{qr^-} continues to be non-corner, and we generate a new edge certificate for it.

Finally, consider the case where η^+ lies on the current edgelet f. In this case we create (i) a new breakpoint w, for which q touches $Q[w]$ at v_{i+1} (indicating the final corner event), (ii) the new Voronoi vertex $\eta^+ = \nu_{pqr}$, and (iii) three new edgelets f_1, f_2, and f_3 that replace f, where f_1 is the portion of f preceding w, f_2 is the new edgelet between w and η^+ (which is denoted by h in Figure 8(c)), and f_3 is the remainder of f; f_1 and f_2 are now part of e_{qr} while f_3 is still an (external) edgelet of e_{pr}. We also add a new external edgelet e_{pq}, which is contained in ρ^+ and is adjacent to η^+. We generate new certificates associated with these new features of the diagram and delete the old ones, similar to the preceding discussion.

This completes the description of the KDS repair mechanism.

The KDS as we described it maintains $O(nk)$ certificates, so its event queue takes $O(nk)$ space (in addition to $\mathrm{VD}(P)$ which also takes $O(nk)$ space). When a generic certificate fails we make $O(1)$ changes to the diagram in $O(1)$ time and $O(1)$ changes to the event queue in $O(\log n)$ time. The simulation of the rotational sweep when we process a singular bisector event may take a long time. But it makes $O(1)$ changes to the diagram in $O(1)$ time and $O(1)$ changes to the event queue in $O(\log n)$ time, for each singular corner or flip event along γ, and the analysis of Section 4 bounds the overall number of these singular sub-events. The formulation of the following theorem (and of Theorem 5.2) assumes this point of view.

Putting everything together, obtain the following main result of this paper.

Theorem 5.1. *Let P be a set of n moving points in $\mathbb{R}^2$ along algebraic trajectories of bounded degree, and let Q be a convex k-gon. Then both $\mathrm{DT}(P)$ and $\mathrm{VD}(P)$, including the sequence of edgelets of each Voronoi edge, can be maintained using a KDS that requires $O(nk)$ storage, encounters $O(k^4 n \lambda_r(n))$ events, and processes each event in $O(\log n)$ time, in a total of $O(k^4 n \lambda_r(n) \log n)$ time. Here r is a constant that depends on the degree of the motion of the points of P.*

5.3 Implicit representation of $\mathrm{VD}(P)$

The drawback of the maintenance of the explicit representation of $\mathrm{VD}(P)$, including the edgelets comprising each Voronoi edge, is that it uses $O(nk)$ storage. We can reduce the storage to $O(n)$ by completely ignoring the edgelets, and storing each Voronoi edge as a single entity in the DCEL connecting the two Voronoi vertices that delimit it (breakpoints are not maintained either).

This in fact simplifies the KDS, because there is no need anymore to maintain bisector certificates or detect bisector events. That is, we only maintain generic and initial singular corner certificates, and generic edge certificates. We also store $\delta(v,p)$, $\delta(v,q)$, and $\delta(v,r)$ explicitly with each vertex $v = v_{pqr}$. As discussed before, $\delta(v,p)$, $\delta(v,q)$, and $\delta(v,r)$ encode the labels of the external edgelets adjacent to v.[5] Suppose that a generic vertex certificate asserting that $\delta(v,p) = e_i = v_iv_{i+1}$ fails when p reaches one of the endpoints v_i or v_{i+1} of e_i. We compute using $\delta(v,p)$, $\delta(v,q)$, $\delta(v,r)$, which among the edges e_{pr} and e_{qr} gains an edgelet and which loses an edgelet, and the labels of the new external edgelets of e_{pr} and e_{qr} that touch v. We compute the new vertex certificates of v and insert them into the queue replacing the old vertex certificates associated with v. We check, by comparing the external edgelets of e_{pr} and e_{qr}, if any of these edges changes its type from a corner edge to a non-corner edge or vice versa. If indeed the type of such edge changes, we remove or add the corresponding edge certificate, as appropriate. In addition we update the edge certificate of e_{pq} if it is a non-corner edge and update $\delta(v,p)$. The processing of a flip event is similar to its processing when the diagram is explicit.

When an initial singular corner event is detected, we run the same process as in the previous subsection, simulating the rotational sweep of Section 3.2. We perform this process by traversing the corresponding portion γ of $\partial\mathrm{Vor}(p)$ between η^- and η^+, essentially as before, but here we compute the sequence of edgelets of each edge e_{pr} along γ on the fly, starting and ending at its external edgelets (which are encoded by the δ values of the vertices adjacent to e_{pr}). The tracing of the edgelets is needed only for finding the edgelet containing η^+, so that we can generate the relevant vertex certificates involving $\delta(\eta^+,p)$, $\delta(\eta^+,q)$, and $\delta(\eta^+,r^+)$ (where r^+ is the third point defining η^+).

Putting everything together, we obtain the following theorem.

Theorem 5.2. *Let P be a set of n moving points in $\mathbb{R}^2$ along algebraic trajectories of bounded degree, and let Q be a convex k-gon. Then both $\mathrm{DT}(P)$ and $\mathrm{VD}(P)$ (without the edgelet subdivision of its edges) can be maintained using a KDS that requires $O(n)$ storage, encounters $O(k^4 n\lambda_r(n))$ events, and processes each event in $O(\log n)$ time, in a total of $O(k^4 n\lambda_r(n)\log n)$ time. Here r is a constant that depends on the degree of the motion of the points of P.*

6 Conclusion

In this paper we have presented a detailed and comprehensive analysis of kinetic Voronoi diagrams and Delaunay triangulations under a convex distance function induced by a convex polygon. We have shown that the number of topological changes, when the input points move along algebraic trajectories of constant degree, is nearly quadratic in the input size, and that the diagrams can be maintained kinetically by a simple KDS, which is efficient, compact, and responsive, in the terminology of [5]. As stated in the introduction, our result, along with the result in the companion paper [2], lead to a compact, responsive, and efficient KDS for maintaining the stable edges of the Euclidean Delaunay triangulation of P, as defined earlier, that processes a near-quadratic number of events if the motion of P is algebraic of bounded degree. See [2] for further details.

An interesting open problem is to improve the dependence on k of the bounds on the number of events, as obtained in Section 4. We tend to conjecture a near-quadratic dependence on k (down from the current factor k^4). A more challenging question is to develop an efficient KDS for Voronoi

[5] We cannot compute these values now from the labels of the external edgelets as we do not maintain the edgelets.

diagrams under polyhedral distance functions in $\mathbb{R}^3$. Of course, the major open problem on this topic is to extend the recent result by Rubin [23] to the case when the points of P move with different speeds.

References

[1] M. A. Abam, M. de Berg, Kinetic spanners in R^d, *Discrete Comput. Geom.* 45 (2011), 723–736.

[2] P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, V. Koltun, N. Rubin and M. Sharir, Kinetic stable Delaunay graphs, *Proc. 26th Annu. Sympos. Comput. Geom.*, 2010, pp. 127–136. Also in CoRR abs/1104.0622: (2011).

[3] P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, N. Rubin and M. Sharir, Stable Delaunay graphs, CoRR abs/1404.4851: (2015).

[4] F. Aurenhammer, R. Klein, and D.-T. Lee, *Voronoi Diagrams and Delaunay Triangulations,* World Scientific, 2013.

[5] J. Basch, L. J. Guibas and J. Hershberger, Data structures for mobile data, *J. Algorithms* 31 (1999), 1–28.

[6] M. de Berg, O. Cheong, M. van Kreveld, and M. Overmars, *Computational Geometry: Algorithms and Applications*, 3rd ed., Springer-Verlag, Berlin, 2008.

[7] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec, Voronoi diagrams in higher dimensions under certain polyhedral distance functions, *Discrete Comput. Geom.* 19 (1998), 485–519.

[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] L. P. Chew and R. L. Drysdale, Voronoi diagrams based on convex distance functions, *Proc. 1st Annu. Sympos. Comput. Geom.*, 1985, pp. 235–244.

[10] R. L. Drysdale III, A practical algorithm for computing the Delaunay triangulation for convex distance functions, *Proc. 1st Annu. ACM-SIAM Sympos. Discrete Algos.*, pp. 159–168.

[11] E. Fogel, D. Halperin, and R. Wein, *CGAL arrangements and their applications: A step-by-step guide*, Springer-Verlag, Heidelberg, 2012.

[12] L. J. Guibas, Modeling motion, *Handbook of Discrete and Computational Geometry,* 2nd edition, (J. E. Goodman and J. O'Rourke, eds.), CRC Press, Boca Raton, 2004, pp. 1117–1134.

[13] 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.*, 1992, pp. 113–125.

[14] F. K. Hwang: An $O(n \log n)$ algorithm for rectilinear minimal spanning trees, *J. ACM* 26 (1979), 177–182.

[15] 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.

[16] K. Kedem, R. Livne, J. Pach, and M. Sharir, On the union of Jordan regions and collision-free translational motion amidst polygonal obstacles, *Discrete Comput. Geom.* 1 (1986), 59–70.

[17] V. Koltun and M. Sharir, Polyhedral Voronoi diagrams of polyhedra in three dimensions, *Discrete Comput. Geom.* 31 (2004), 83–124.

[18] D.-T. Lee, Two-dimensional Voronoi diagrams in the L_p-metric, *J. ACM* 27 (1980), 604–618.

[19] D.-T. Lee and C. K. Wong, Voronoi diagrams in L_1 (L_∞) metrics with 2-dimensional storage applications, *SIAM J. Comput.* 9 (1980), 200–211.

[20] D. Leven and M. Sharir, Planning a purely translational motion for a convex object in two–dimensional space using generalized Voronoi diagrams, *Discrete Comput. Geom.* 2 (1987), 9–31.

[21] L. Ma, *Bisectors and Voronoi Diagrams for Convex Distance Functions*, PhD Thesis, Fern University, Hagen, 2000.

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

[23] N. Rubin, On kinetic Delaunay triangulations: A near quadratic bound for unit speed motions, *Proc. 54th Annual IEEE Sympos. Foundat. Comput. Science*, 2013, 519–528.

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

[25] S. Skyum, A sweepline algorithm for generalized Delaunay triangulations, Tech. Rept. DAIMI PB-373, Computer Science Department, Aarhus University, 1991.

[26] P. Widmayer, Y.-F. Wu, and C. K. Wong, On some distance problems in fixed orientations, *SIAM J. Comput.* 16 (1987), 728–746.

[27] C. K. Yap, A geometric consistency theorem for a symbolic perturbation scheme, *J. Comput. Syst. Sci.* 40 (1990), 2–18.