Stable Delaunay Graphs*

Pankaj K. Agarwal† Jie Gao‡ Leonidas Guibas§ Haim Kaplan¶
Natan Rubin‖ Micha Sharir**

April 29, 2014

Abstract

Let P be a set of n points in $\mathbb{R}^2$, and let $\mathrm{DT}(P)$ denote its Euclidean Delaunay triangulation. We introduce the notion of an edge of $\mathrm{DT}(P)$ being *stable*. Defined in terms of a parameter $\alpha > 0$, a Delaunay edge pq is called α-stable, if the (equal) angles at which p and q see the corresponding Voronoi edge e_{pq} are at least α. A subgraph G of $\mathrm{DT}(P)$ is called $(c\alpha, \alpha)$-*stable Delaunay graph* (SDG in short), for some constant $c \geq 1$, if every edge in G is α-stable and every $c\alpha$-stable of $\mathrm{DT}(P)$ is in G.

We show that if an edge is stable in the Euclidean Delaunay triangulation of P, then it is also a stable edge, though for a different value of α, in the Delaunay triangulation of P under any convex distance function that is sufficiently close to the Euclidean norm, and vice-versa. In particular, a 6α-stable edge in $\mathrm{DT}(P)$ is α-stable in the Delaunay triangulation under the distance function induced by a regular k-gon for $k \geq 2\pi/\alpha$, and vice-versa. Exploiting this relationship and the analysis in [3], we present a linear-size kinetic data structure (KDS) for maintaining an $(8\alpha, \alpha)$-SDG as the points of P move. If the points move along algebraic trajectories of bounded degree, the KDS processes nearly quadratic events during the motion, each of which can processed in $O(\log n)$ time. Finally, we show that a number of useful properties of $\mathrm{DT}(P)$ are retained by SDG of P.

*An earlier version of this paper appeared in *Proc. 26th Annual Symposium on Computational Geometry*, 2010, 127–136.

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

‡Department of Computer Science, Stony Brook University, Stony Brook, NY 11794, USA, jgao@cs.sunysb.edu.

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

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

‖Jussieu Institute of Mathematics, Pierre and Marie Curie University and Paris Diderot University, UMR 7586 du CNRS, Paris 75005, France. rubinnat.ac@gmail.com.

**School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel; and Courant Institute of Mathematical Sciences, New York University, New York, NY 10012, USA. michas@tau.ac.il.

1 Introduction

Let P be a set of n points in $\mathbb{R}^2$. For a point $p \in P$, the (Euclidean) *Voronoi cell* of p is defined as

$$\text{Vor}(p) = \{x \in \mathbb{R}^2 \mid \|xp\| \le \|xp'\| \, \forall p' \in P\}.$$

The Voronoi cells of points in P are nonempty, have pairwise-disjoint interiors, and partition the plane. The planar subdivision induced by these Voronoi cells is referred to as the (Euclidean) *Voronoi diagram* of P and we denote it as $\text{VD}(P)$. The *Delaunay graph* of P is the dual graph of $\text{VD}(P)$, i.e., pq is an edge of the Delaunay graph if and only if $\text{Vor}(p)$ and $\text{Vor}(q)$ share an edge. This is equivalent to the existence of a circle passing through p and q that does not contain any other point of P in its interior—any circle centered at a point of $\partial\text{Vor}(p) \cap \partial\text{Vor}(q)$ and passing through p and q is such a circle. If no four points of P are cocircular, then the planar subdivision induced by the Delaunay graph is a triangulation of the convex hull of P—the well-known (Euclidean) *Delaunay triangulation* of P, denoted as $\text{DT}(P)$. See Figure 1 (a). $\text{DT}(P)$ consists of all triangles whose circumcircles do not contain points of P in their interior. Delaunay triangulations and Voronoi diagrams are fundamental to much of computational geometry and its applications. See [8] for a very recent textbook on these structures.

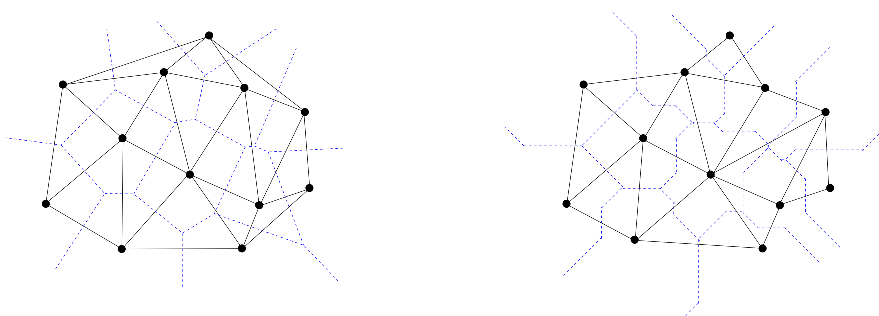

(a) (b)

Figure 1. (a) The Euclidean Voronoi diagram (dotted) and Delaunay triangulation (solid); (b) the Q-Voronoi diagram and Q-Delaunay triangulation for an axis-parallel square Q, i.e., the Voronoi and Delaunay diagrams under the L_∞-metric.

In many applications of Delaunay/Voronoi methods (e.g., mesh generation and kinetic collision detection), the input points are moving continuously, so they need to be efficiently updated as motion occurs. Even though the motion of the points is continuous, the combinatorial and topological structure of $\text{VD}(P)$ and $\text{DT}(P)$ change only at discrete times when certain "critical events" occur. A challenging open question in combinatorial geometry is to bound the number of critical events if each point of P moves along an algebraic trajectory of constant degree.

Guibas et al. [14] showed a roughly cubic upper bound of $O(n^2\lambda_s(n))$ on the number of critical events. Here $\lambda_s(n)$ is the maximum length of an (n,s)-Davenport-Schinzel sequence [21], and s is a constant depending on the degree of the motion of the points. See also Fu and Lee [13]. The best known lower bound is quadratic [21]. Recent works of Rubin [19, 20] establish an almost quadratic bound of $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$, for the restricted cases where any four points of P can be cocircular at most *two* or *three* times. In particular, the latter study [20] covers the case of points moving along lines at common unit speed, which has been highlighted as a major open problem in discrete and computational geometry; see [12]. Nevertheless, no sub-cubic upper bound is known for more general motions, including the case where the points of P are moving along lines at non-uniform, albeit fixed, speeds. It is worth mentioning that the analysis in [19], and even more so in [20], is fairly involved, which results in a huge implicit constant of proportionality.

Given this gap in the bound on the number of critical events, it is natural to ask whether one can define a large subgraph of the Delaunay graph of P so that (i) it provably experiences at most a nearly quadratic number of critical events, (ii) it is reasonably easy to define and maintain, and (iii) it retains useful properties for further applications. This paper defines such a subgraph of the Delaunay graph, shows that it can be maintained efficiently, and proves that it preserves a number of useful properties of $\mathrm{DT}(P)$.

Related work. It is well known that $\mathrm{DT}(P)$ can be maintained efficiently using the so-called kinetic data structure framework proposed by Basch *et al.* [10]. A triangulation $\mathcal{T}$ of the convex hull $\mathrm{conv}(P)$ of P is the Delaunay triangulation of P if and only if for every edge pq adjacent to two triangles $\triangle pqr^+$ and $\triangle pqr^-$ in $\mathcal{T}$, the circumcircle of $\triangle pqr^+$ (resp., $\triangle pqr^-$) does not contain r^- (resp., r^+). Equivalently,

$$\angle pr^+q + \angle pr^-q < \pi. \tag{1}$$

Equality occurs when p, q, r^+, r^- are cocircular, which generally signifies that a combinatorial change in $\mathrm{DT}(P)$ (a so-called *edge flip*) is about to take place. We also extend (1) to apply to edges pq of the hull, each having only one adjacent triangle, $\triangle pqr^+$. In this case we take r^- to lie at infinity, and put $\angle pr^-q = 0$. An equality in (1) occurs when p, q, r^+ become collinear (along the hull boundary), and again this signifies a combinatorial change in $\mathrm{DT}(P)$.

This makes the maintenance of $\mathrm{DT}(P)$ under point motion quite simple: an update is necessary only when the empty circumcircle condition (1) fails for one of the edges, i.e., for an edge pq, adjacent to triangles $\triangle pqr^+$ and $\triangle pqr^-$, p, q, r^+, and r^- become cocircular.[1] Whenever such an event happens, the edge pq is *flipped* with r^+r^- to restore Delaunayhood. Keeping track of these cocircularity events is straightforward, and each such event is detected and processed in $O(\log n)$ time [14]. However, as mentioned above, the best known upper bound on the number of events processed by this KDS (which is the number of topological changes in $\mathrm{DT}(P)$ during the motion), assuming that the points of P are moving along algebraic trajectories of bounded degree, is near cubic [14] (except for the special cases treated in [19, 20]).

So far we have only considered the Euclidean Voronoi and Delaunay diagrams, but a considerable amount of literature exists on Voronoi and Delaunay diagrams under other norms and so-called *convex distance functions*; see Section 2 for details.

Chew [11] showed that the Delaunay triangulation of P under the L_1- or L_∞-metric experiences only a near-quadratic number of events, if the motion of the points of P is algebraic of bounded degree. In the companion paper [3], we present a kinetic data structure for maintaining the Voronoi diagram and Delaunay triangulation of P under a polygonal convex distance function for an arbitrary convex polygon Q. (see Section 2 for the definition) that processes only a near-quadratic number of events, and can be updated in $O(\log n)$ time at each event. Since a regular convex k-gon approximates a circular disk, it is tempting to maintain the Delaunay triangulation under a polygonal convex distance function as a (hopefully substantial) portion of the Euclidean Delaunay graph of P. Unfortunately, the former is not necessarily a subgraph of the latter [8].

[1] We assume the motion of the points to be sufficiently generic, so that no more than four points can become cocircular at any given time, and so that equality in (1) is not a local maximum of the left-hand side.

Many subgraphs of $\mathrm{DT}(P)$, such as the Euclidean minimum spanning tree (MST), Gabriel graph, relative neighborhood graph, and α-shapes, have been used extensively in a wide range of applications (see e.g. [8]). However no sub-cubic bound is known on the number of discrete changes in their structures under an algebraic motion of the points of P of bounded degree. Furthermore, no efficient kinetic data structures are known for maintaining them, for unlike $\mathrm{DT}(P)$, they may undergo a "non-local" change at a critical event; see [1, 9, 18] for some partial results on maintaining an MST.

Our results. ***Stable Delaunay edges:*** We introduce the notion of *α-stable Delaunay edges*, for a fixed parameter $\alpha > 0$, defined as follows. Let pq be a Delaunay edge under the Euclidean norm, and let $\triangle pqr^+$ and $\triangle pqr^-$ be the two Delaunay triangles incident to pq. Then pq is called *α-stable* if its opposite angles in these triangles satisfy

$$\angle pr^+q + \angle pr^-q \leq \pi - \alpha. \tag{2}$$

As above, the case where pq lies on $\mathrm{conv}(P)$ is treated as if r^-, say, lies at infinity, so that the corresponding angle $\angle pr^-q$ is equal to 0. An equivalent and more useful definition, in terms of the dual Voronoi diagram, is that pq is α-stable if *the angles at which p and q see their common Voronoi edge e_{pq} are at least α each.* See Figure 2(a).

In the case where pq lies on $\mathrm{conv}(P)$ the corresponding dual Voronoi edge e_{pq} is an infinite ray emanating from some Voronoi vertex x. We define the angle in which a point p *see* such a Voronoi ray to be the angle between the segment px and an infinite ray parallel to e_{pq} emanating from p. With this definition of the angle in which a point of $\mathrm{conv}(P)$ sees a Voronoi ray, it is easy to check that the alternative definition of α-stability is equivalent to the ordinal definition also when pq lies on $\mathrm{conv}(P)$. We call the Voronoi edges corresponding to α-stable Delaunay edges *α-long* (and call the remaining edges *α-short*). See Figure 2. Note that for $\alpha = 0$, when no four points are cocircular, (2) coincides with (1).

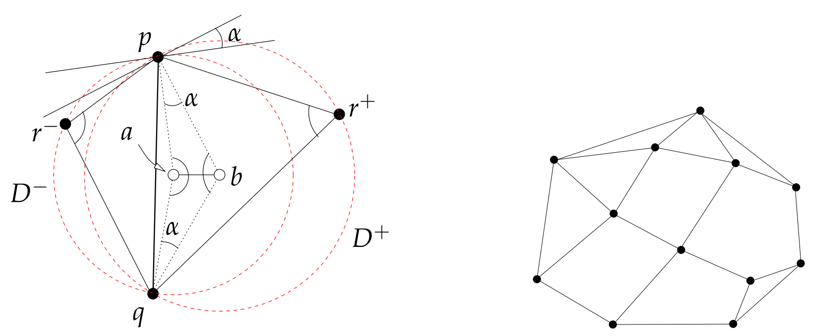

(a) (b)

Figure 2. (a) An α-stable edge pq in $\mathrm{DT}(P)$, and its dual edge $e_{pq} = ab$ in $\mathrm{VD}(P)$; the fact that $\angle r^+ + \angle r^- = \pi - \alpha$ follows by elementary geometric considerations. (b) A SDG of the Delaunay triangulation $\mathrm{DT}(P)$ shown in Figure 1 (a), with $\alpha = \pi/8$.

A justification for calling such edges stable lies in the following observation: If a Delaunay edge pq is α-stable then it remains in $\mathrm{DT}(P)$ during any continuous motion of the points of P for which every angle $\angle prq$, for $r \in P \setminus \{p, q\}$, changes by at most $\alpha/2$. This is clear because, as is easily verified, at any time when pq is α-stable we have $\angle pr^+q + \angle pr^-q \leq \pi - \alpha$ for *any* pair of points r^+,

r^- lying on opposite sides of the line ℓ supporting pq, so, if each of these angles changes by at most $\alpha/2$ we still have $\angle pr^+q + \angle pr^-q \le \pi$ for every such pair r^+, r^-, implying that pq remains an edge of $\mathrm{DT}(P)$.[2] Hence, as long as the "small angle change" condition holds, stable Delaunay edges remain a "long time" in the triangulation. Informally speaking, the non-stable edges pq of $\mathrm{DT}(P)$ are those for which p and q are almost cocircular with their two common Delaunay neighbors r^+, r^-, and hence pq is more likely to get flipped "soon."

Stable Delaunay graph: For two parameters $0 \le \alpha \le \alpha' \le \pi$, we call a subgraph $\mathcal{G}$ of $\mathrm{DT}(P)$ an (α', α)-*stable Delaunay graph* (an (α', α)-SDG for short) if

(S1) every edge of $\mathcal{G}$ is α-stable, and

(S2) every α'-stable edge of $\mathrm{DT}(P)$ belongs to $\mathcal{G}$.

Note that an (α', α)-SDG is not uniquely defined even for fixed α, α' because the edges that are α-stable but not α'-stable may or may not be in $\mathcal{G}$. Throughout this paper, α' will be some fixed (and reasonably small) multiple of α.

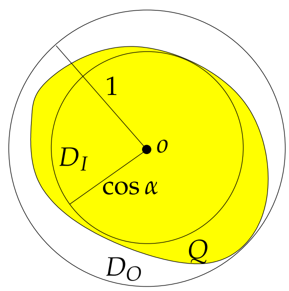

Figure 3. A convex distance function, induced by Q, that is α-close to the Euclidean norm.

Our main result is that a stable edge of the Euclidean Delaunay triangulation appears as stable edge in the Delaunay triangulation under any convex distance function that is sufficiently close to the Euclidean norm (see Section 2 for more details). More precisely, we say that the distance function induced by a compact convex set Q is α-*close* to the Euclidean norm if Q is contained in the unit disk D_O and contains the disk $D_I = (\cos\alpha)D_O$ both centered in the origin.[3] See Figure 3. In particular, for $k = \pi/\alpha$, the regular k-gon Q_k is such a set, as easy trigonometry shows. We prove the following:

Theorem 1.1. *Let P be a set of n points in $\mathbb{R}^2$, $\alpha > 0$ a parameter, and Q a compact, convex set inducing a convex distance function $d_Q(\cdot,\cdot)$ that is α-close to the Euclidean norm. Then the following properties hold.*

(i) Every 11α-stable Delaunay edge under the Euclidean norm is an α-stable Delaunay edge under d_Q.

(ii) Symmetrically, every 11α-stable Delaunay edge under d_Q is also an α-stable Delaunay edge under the Euclidean norm.

[2]This argument also covers the cases when a point r crosses ℓ from side to side: Since each point, on either side of ℓ, sees pq at an angle of $\le \pi - \alpha$, it follows that no point can cross pq itself – the angle has to increase from $\pi - \alpha$ to π. Any other crossing of ℓ by a point r causes $\angle prq$ to decrease to 0, and even if it increases to $\alpha/2$ on the other side of ℓ, pq is still an edge of DT, as is easily checked.

[3]The Hausdorff distance between Q and D_O is at most $1 - \cos\alpha \approx \alpha^2/2$.

In particular, if Q is a regular k-gon for $k \geq 2\pi/\alpha$, then the above theorem holds for Q. In the companion paper [3], we have presented an efficient kinetic data structure for maintaining the Delaunay triangulation and Voronoi diagram of P under a polygonal convex distance function. Using this result, we obtain the second main result of the paper:

Theorem 1.2. *Let P be a set of n moving points in $\mathbb{R}^2$ under algebraic motion of bounded degree, and let $\alpha \in (0, \pi)$ be a parameter. A Euclidean $(8\alpha, \alpha)$-stable Delaunay graph of P can be maintained by a linear-size KDS that processes $O\left(\frac{1}{\alpha^4} n\lambda_r(n)\right)$ events and updates the SDG at each event in $O(\log n)$ time. Here r is a constant that depends on the degree of the motion of the points of P, and $\lambda_r(n)$ is the maximum length of a Davenport-Schnizel sequence of order r.*

For simplicity, we first prove in Section 3 Theorem 1.1 for the case when Q is a regular k-gon, for $k \geq 2\pi/\alpha$, and use the argument to prove Theorem 1.2. Actually, we prove Theorem 1.1 with a slightly better constant using the additional structure possessed by the diagrams when Q is a regular k-gon. Next, we prove in Section 4 Theorem 1.1 for an arbitrary Q. Finally, we prove in Section 5 a few useful properties of $\mathrm{DT}(P)$ that are retained by the stable Delaunay graph of P.

2 Preliminaries

This section introduces a few notations, definitions, and known results that we will need in the paper.

We represent a direction in $\mathbb{R}^2$ as a point on the unit circle $\mathbb{S}^1$. For a direction $u \in \mathbb{S}^1$ and an angle $\theta \in [0, 2\pi)$, we use $u + \theta$ to denote the direction obtained after rotating the vector $\vec{ou}$ by angle θ in clockwise direction. For a point $x \in \mathbb{R}^2$ and a direction $u \in \mathbb{S}^1$, let $u[x]$ denote the ray emanating from x in direction u.

Q-distance function. Let Q be a compact, convex set with non-empty interior and with the origin, denoted by o, lying in its interior. A homothetic copy Q' of Q can be represented by a pair (p, λ), with the interpretation $Q' = p + \lambda Q$; p is the *placement* (location) of the center o of Q', and λ is its *scaling factor* (about its center). Q defines a distance function (also called the *gauge* of Q)

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

Note that, unless Q is centrally symmetric with respect to the origin, d_Q is not symmetric.

Given a finite point set $P \subset \mathbb{R}^2$ and a point $p \in P$, we denote by $\mathrm{Vor}^Q(p)$, $\mathrm{VD}^Q(P)$, and $\mathrm{DT}^Q(P)$ the Voronoi cell of p, the Voronoi diagram of P, and the Delaunay triangulation of P, respectively, under the distance function $d_Q(\cdot, \cdot)$; see Figure 1 (b). To be precise (because of the potential asymmetry of d_Q), we define

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

and then $\mathrm{VD}^Q(P)$ and $\mathrm{DT}^Q(P)$ are defined in complete analogy to the Euclidean case. We refer the reader to [3] for formal definitions and details of these structures. Throughout this paper, we will drop the superscript Q from $\mathrm{Vor}^Q, \mathrm{VD}^Q, \mathrm{DT}^Q$ when referring to them under the Euclidean norm.

For a point $z \in \mathbb{R}^2$, let $Q[z]$ denote the homothetic copy of Q centered at z such that its boundary touches the Q-nearest neighbor(s) of z in P, i.e., $Q[z]$ is represented by the pair (z, λ)

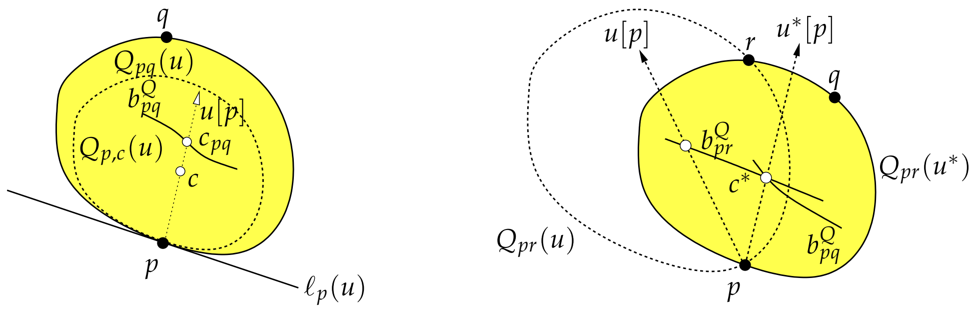

Figure 4. (a) Homothetic copies $Q_p(u)$, centered at $c \in u[p]$ (dashed), and $Q_{pq}(u)$ (solid) of Q. The ray $u[p]$ hits b^Q_{pq} at the center c_{pq} of $Q_{pq}(u)$, which is defined, when Q is smooth at its contact with p, if and only if q and $Q_p(u)$ lie on the same side of $\ell_p(u)$; (b) Sliding $Q_{pr}(u)$ away from q.

where $\lambda = \min_{p \in P} d_Q(z, p)$. In other words, $Q[z]$ is the largest homothetic copy of Q that is centered at z whose interior is P-empty. We also use the notation $Q_p(u)$ to denote a "generic" homothetic copy of Q which touches p and is centered at some point on $u[p]$. See Figure 4 (a). Note that all homothetic copies of $Q_p(u)$ touch p at the same point ζ of ∂Q, and therefore share the same tangent at p. This tangent is unique if Q is smooth at ζ, and we denote it by $\ell_p(u)$. If Q is not smooth at ζ then there is a nontrivial range of tangents (i.e., supporting lines) to Q at ζ; we can take $\ell_p(u)$ to be any of them, and it will be a supporting line to all the copies $Q_p(u)$ of Q.

For a pair of points $p, q \in P$, let b^Q_{pq} denote the *Q-bisector* of p and q—the locus of all placements of the center of any homothetic copy Q' of Q that touches p and q. If Q is strictly convex or if Q is not strictly convex but no two points are collinear with a straight segment on ∂Q, then b^Q_{pq} is a one-dimensional curve and any ray $u[p]$ that hits b^Q_{pq} does so at a unique point. For such a direction u and a pair of points $p, q \in P$, let $Q_{pq}(u)$ denote the homothetic copy of Q that touches p and q, whose center is $u[p] \cap b^Q_{pq}$.

If Q is not strictly convex and p, q are points in P such that $\vec{pq}$ is parallel to a straight portion e of ∂Q then b^Q_{pq} is not one-dimensional. In this case $Q_{pq}(u)$ is not well defined when u is a direction that connects e to the center c of Q. As is easy to check, in any other case the ray $u[p]$ either hits b_{pq} at a unique point which determines $Q_{pq}(u)$, or entirely misses b^Q_{pq}. See the companion paper [3] for a detailed discussion of this phenomenon.

A useful property of the Q-bisectors is that any two bisectors b^Q_{pq}, b^Q_{pr} with a common generating point p, intersect exactly once, namely, at the center c^* of the unique homothetic copy of Q that simultaneously touches p, q and r [17]. For this property to hold, though, we need to assume that (i) the points p, q, r are not collinear, and (ii) either Q is strictly convex, or, otherwise, that none of the directions $\vec{pq}, \vec{pr}$ is parallel to a straight portion of ∂Q. (The precise condition is that b_{pq} and b_{pr} be one-dimensional in a neighborhood of c^*.) The local topology of the restricted Q-Voronoi diagram $\mathrm{VD}^Q(\{p, q, r\})$ near c^* is largely determined by the orientation of the triangle $\triangle pqr$. Specifically, assume with no loss of generality that $\vec{pr}$ is counterclockwise to $\vec{pq}$, and let u^* be the direction of the ray $\vec{pc^*}$. Refer to Figure 4 (b). If we continuously rotate a ray $u[p]$, for $u \in \mathbb{S}^1$, in counterclockwise direction from $u^*[p]$, the corresponding copy $Q_{pr}(u)$ will slide away from its contact with q because the portion of $Q_{pr}(u)$ to the right of $\vec{pr}$ shrinks during the rotation. Therefore, the rotating ray $u[p]$ either misses b^Q_{pq} entirely or hits b^Q_{pq} after b^Q_{pr}. A symmetric phenomenon, with q and r interchanged, takes place if we rotate the ray $u[p]$ in clockwise direction from $u^*[p]$.

It is known that $\mathrm{Vor}^Q(p)$, for every point $p \in P$, is star-shaped [8], which implies that each Voronoi edge e^Q_{pq} is fully enclosed between the two rays that emanate from p, or from q, through its endpoints. We remark that, unlike the Euclidean case, the angles $\angle xpy$, $\angle xqy$ need not be equal in general.

Finally, we extend the notion of stable edges to $\mathrm{DT}^Q(P)$. We call an edge $pq \in \mathrm{DT}^Q(P)$ *α-stable* if the following property holds for the dual edge e^Q_{pq} in $\mathrm{VD}^Q(P)$:

> *Each of the points p, q sees their common Q-Voronoi edge e^Q_{pq} at angle at least α. That is, if x and y are the endpoints of e^Q_{pq}, then* $\min\{\angle xpy, \angle xqy\} \geq \alpha$.

This definition coincides with the definition of α-stability under the Euclidean norm when Q is the unit disk (and in this case both angles are equal).

Remark. If Q is not strictly convex (and pq is parallel to a straight portion of ∂Q), the endpoints of e^Q_{pq} may be not well defined. In this case, we resort to the following, more careful definition of α-stability. A ray $u[p]$ is said to (properly) *cross* e^Q_{pq} only if the copy $Q_{pq}(u)$ is uniquely defined. The center of such a copy $Q_{pq}(u)$ necessarily lies within the one-dimensional portion $\tilde{e}^Q_{pq}$ of $Q_{pq}(u)$, which is easily seen to be non-empty and connected. We say that the edge e^Q_{pq} is α-stable if the set of rays $u[p]$ properly crossing e^Q_{pq} spans an angle of at least α, and a symmetric condition holds for the rays emanating from q. In other words, our notion of α-stability ignores the two-dimensional regions of e^Q_{pq} (if these exist).[4]

Polygonal convex distance function. As mentioned in the introduction, we will be considering the case when Q is a regular k-gon, for some even integer $k \geq 2\pi/\alpha$, centered at the origin. Let $v_0, \ldots, v_{k-1}$ be its sequence of vertices arranged in clockwise direction. For each $0 \leq j < k$, let u_j be the direction of the vector that connects v_j to the center of Q (see Figure 5 (a)). We will use $b^\diamond_{pq}, \mathrm{Vor}^\diamond(p), \mathrm{VD}^\diamond(P), \mathrm{DT}^\diamond(P)$ to denote $b^Q_{pq}, \mathrm{Vor}^Q(p), \mathrm{VD}^Q(P), \mathrm{DT}^Q(P)$, respectively, when Q is a regular k-gon.

We say that P is in *general position* (with respect to Q) if no three points of P lie on a line, no two points of P lie on a line parallel to an edge or a diagonal of Q, and no four points of P are *Q-cocircular*, i.e., no four points of P lie on the boundary of a common homothetic copy of Q.

The placements on $b^\diamond_{pq}$ at which (at least) one of p and q, say, p, touches Q' at a vertex is called a *corner placement* (or a *corner contact*) at p; see Figure 5 (b). We also refer to these points on $b^\diamond_{pq}$ as *breakpoints*. We call a homothetic copy of Q whose vertex v_j touches a point p, a *v_j-placement of Q at p*.

The following property of $b^\diamond_{pq}$ is proved in [3, Lemma 2.5]:

Lemma 2.1. *Let Q be a regular k-gon, and let p and q be two points in general position with respect to Q. Then $b^\diamond_{pq}$ is a polygonal chain with $k-2$ breakpoints and the breakpoints along $b^\diamond_{pq}$ alternate between corner contacts at p and corner contacts at q.*

[4] As is easy to check, the one-dimensional portion $\tilde{e}^Q_{pq}$ of e^Q_{pq} varies continuously (in Hausdorff sense) with any sufficiently small perturbation of p and q within P. Furthermore, it is the *only* such portion: If a ray $u[p]$ hits e^Q_{pq} outside $\tilde{e}^Q_{pq}$ (i.e., within its two-dimensional portion), there is a symbolic perturbation of p and q causing $u[p]$ to completely miss e^Q_{pq}.

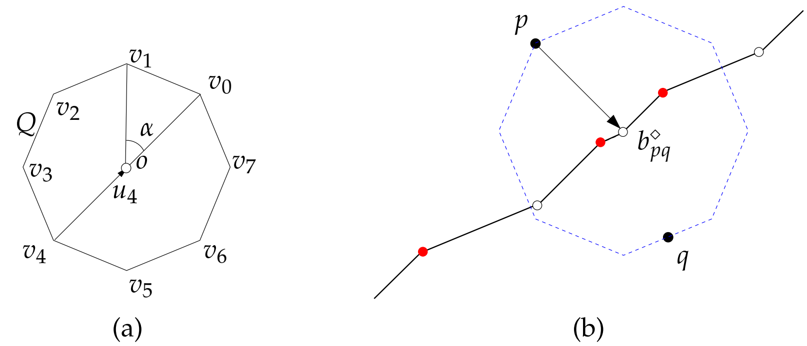

Figure 5. (a) A regular octagon Q centered at the origin. (b) The bisector $b^\diamond_{pq}$ for the regular octagon Q; it has six breakpoints and the corner contacts along $b^\diamond_{pq}$ alternate between contacts at p (hollow circles) and contacts at q (filled circles).

3 $\mathbf{DT}^\diamond(\mathbf{P})$ and Euclidean SDG's

In this section we first prove a slightly stronger version of Theorem 1.1 for the case when Q is a regular k-gon for some even integer $k \geq 2\pi/\alpha$, and then prove Theorem 1.2. We follow the notation in Section 2, and, for simplicity, we assume that $\alpha = 2\pi/k$.

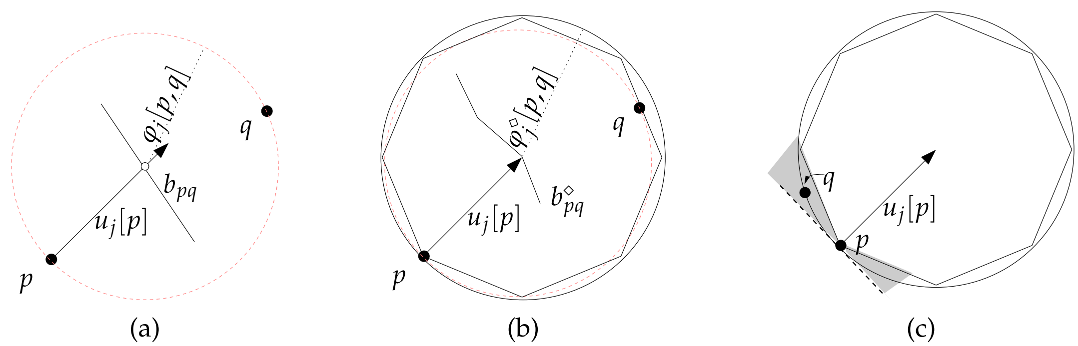

Figure 6. (a) The function $\varphi_j[p,q]$, which is equal to the radius of the circle that pass through p and q and whose center lies on $u_j[p]$. (b) The bisector $b^\diamond_{pq}$ and the function $\varphi^\diamond_j[p,q]$, which is equal to the radius of the circle that circumscribes the v_j-placement of Q at p that also touches q. (c) The case when $\varphi^\diamond_j[p,q] = \infty$ while $\varphi_j[p,q] < \infty$. In this case q lies in one of the shaded wedges.

For any pair $p,q \in P$, let $\varphi_j[p,q]$ denote the distance from p to the point $u_j[p] \cap b_{pq}$; we put $\varphi_j[p,q] = \infty$ if $u_j[p]$ does not intersect b_{pq}. See Figure 6 (a). The point q minimizes $\varphi_i[p,q']$, among all points q' for which $u_i[p]$ intersects $b_{pq'}$, if and only if the intersection between b_{pq} and $u_i[p]$ lies on the Voronoi edge e_{pq}. We call q the *neighbor of p in direction u_i*, and denote it by $N_i(p)$.

Similarly, let $\varphi^\diamond_j[p,q]$ denote the distance from p to the point $u_j[p] \cap b^\diamond_{pq}$; we put $\varphi^\diamond_j[p,q] = \infty$ if $u_j[p]$ does not intersect $b^\diamond_{pq}$. If $\varphi^\diamond_j[p,q] < \infty$ then the point $b^\diamond_{pq} \cap u_j[p]$ is the center of the v_j-placement Q' of Q at p that also touches q, and there is a unique such point. The value $\varphi^\diamond_j[p,q]$ is equal to the circumradius of Q'. See Figure 6 (b). The *neighbor* $N^\diamond_j[p]$ of p in direction u_j is defined to be the point $q \in P \setminus \{p\}$ that minimizes $\varphi^\diamond_j[p,q]$.

Note that $\varphi_j[p,q] < \infty$ if and only if the angle between $\vec{pq}$ and $u_j[p]$ is smaller than $\pi/2$. In contrast, $\varphi^\diamond_j[p,q] < \infty$ if and only if the angle between $\vec{pq}$ and $u_j[p]$ is at most $\pi/2 - \pi/k =$

$\pi/2 - \alpha/2$. Moreover, we have $\varphi_j[p,q] \le \varphi_j^\diamond[p,q]$ (see Figure 6). Therefore, $\varphi_j^\diamond[p,q] < \infty$ always implies $\varphi_j[p,q] < \infty$, but not vice versa; see Figure 6 (c). Note also that in both the Euclidean and the polygonal cases, the respective quantities $N_j[p]$ and $N_j^\diamond[p]$ may be undefined.

Lemma 3.1. *Let $p, q \in P$ be a pair of points such that $N_j(p) = q$ for $h \ge 3$ consecutive indices, say $0 \le j \le h-1$. Then for each of these indices, except possibly for the first and the last one, we also have $N_j^\diamond[p] = q$.*

Proof. Let w_1 (resp., w_2) be the point at which the ray $u_0[p]$ (resp., $u_{h-1}[p]$) hits the edge e_{pq} in $\mathrm{VD}(P)$. (By assumption, both points exist.) Let D_1 and D_2 be the disks centered at w_1 and w_2, respectively, and touching p and q. By definition, neither of these disks contains a point of P in its interior. The angle between the tangents to D_1 and D_2 at p or at q (these angles are equal) is $\beta = (h-1)\alpha$; see Figure 7 (a).

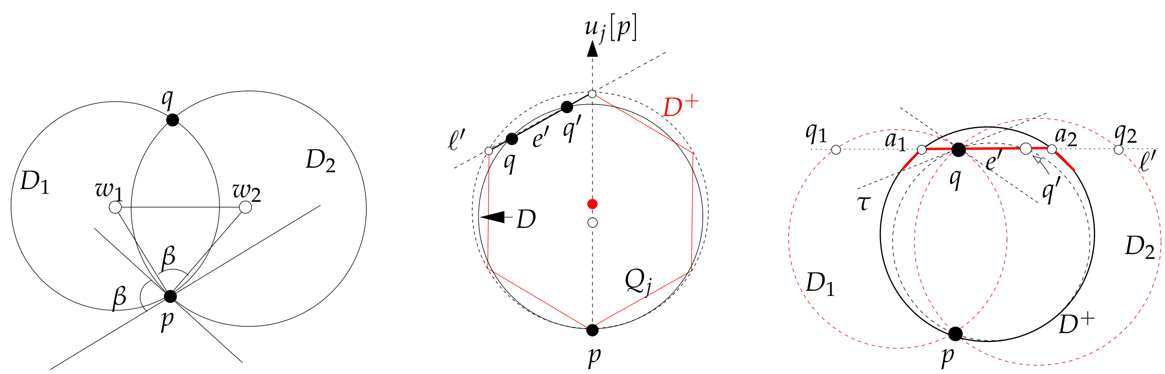

Figure 7. (a) The angle between the tangents to D_1 and D_2 at p (or at q) is equal to $\angle w_1 p w_2 = \beta = (h-1)\alpha$. (b) The disks D and D^+ and the homothetic copy Q_j of Q; $\ell' \cap D = qq' \subseteq e'$. (c) ℓ' forms an angle of at least $\alpha/2$ with each of the tangents to D_1, D_2 at q, and the edge $e' = a_1a_2 \subset D_1 \cup D_2$.

Fix an arbitrary index $1 \le j \le h-2$, so $u_j[p]$ intersects e_{pq} and forms an angle of at least α with each of pw_1, pw_2. Let $Q_j := Q_{pq}(u_j)$ be the v_j-placement of Q at p that touches q. To see that such a placement exists, we note that, by the preceding remark, it suffices to show that the angle between $\vec{pq}$ and $u_j[p]$ is at most $\pi/2 - \alpha/2$; that is, to rule out the case where q lies in one of the shaded wedges in Figure 6 (c). This case is indeed impossible, because then one of $u_{j-1}[p], u_{j+1}[p]$ would form an angle greater than $\pi/2$ with $\vec{pq}$, contradicting the assumption that both of these rays intersect the (Euclidean) b_{pq}. The claim now follows from the next lemma, which shows that $Q_j \subset D_1 \cup D_2$, which implies that $\mathrm{int}\, Q_j \cap P = \emptyset$ and thus $N_j^\diamond[p] = q$, as claimed. □

Lemma 3.2. *In the notation in the proof of Lemma 3.1, $Q_j \subset D_1 \cup D_2$, for $1 \le j \le h-2$.*

Proof. Fix a value of $1 \le j \le h-2$. Let e' be the edge of Q_j passing through q; see Figure 7 (b). Let D be the disk whose center lies on $u_j[p]$ and which passes through p and q, and let D^+ be the circumscribing disk of Q_j. Since $p \in \partial D \cap \partial D^+$, $q \in \partial D \cap \mathrm{int}\, D^+$, and D and D^+ are centered on the ray $u_j[q]$ emanating from p, it follows that $D \subset D^+$. The line ℓ' containing e' crosses D in a chord qq' that is fully contained in e', as $qq' = D \cap \ell' \subseteq D^+ \cap \ell' = e'$.

The angle between the tangent to D at q, denoted by τ, and the chord qq' is equal to the angle at which p sees qq'. This angle is smaller than the angle at which p sees e', which in turn is equal to

$\alpha/2$. Recall that $u_j[p]$ makes an angle of at least α with each of pw_1 and pw_2, therefore τ forms an angle of at least α with each of the tangents to D_1, D_2 at q. Combining this with the fact that the angle between τ and e' is at most $\alpha/2$, we conclude that e' forms an angle of at least $\alpha/2$ with each of these tangents; see Figure 7 (c).

The line ℓ' marks two chords q_1q, qq_2 within the respective disks D_1, D_2. We claim that e' is fully contained in their union q_1q_2. Indeed, the angle q_1pq is equal to the angle between ℓ' and the tangent to D_1 at q, so $\angle q_1pq \geq \alpha/2$. On the other hand, the angle at which p sees e' is $\alpha/2$, which is no larger. This, and the symmetric argument involving D_2, are easily seen to imply the claim.

Now consider the circumscribing disk D^+ of Q_j. Denote the endpoints of e' as a_1 and a_2, where a_1 lies in q_1q and a_2 lies in qq_2. Since the ray $\vec{pa}_1$ hits ∂D^+ before hitting ∂D_1, and the ray $\vec{pq}$ hits these circles in the reverse order, it follows that the second intersection of ∂D_1 and ∂D^+ (other than p) must lie on a ray from p which lies between the rays $\vec{pa}_1, \vec{pq}$ and thus crosses e'. See Figure 7 (c). Symmetrically, the second intersection point of ∂D_2 and ∂D^+ also lies on a ray which crosses e'. It follows that the arc of ∂D^+ delimited by these intersections and containing p is fully contained in $D_1 \cup D_2$. Hence all the vertices of Q_j (which lie on this arc) lie in $D_1 \cup D_2$. This, combined with the fact, established in the preceding paragraph, that $e' \subseteq q_1q_2$ implies that $Q_j \subset D_1 \cup D_2$. □

Next, we use Lemma 3.1 to prove its converse. Specifically, we prove the following lemma.

Lemma 3.3. *Let $p, q \in P$ be a pair of points such that $N_j^\diamond[p] = q$ for at least five consecutive indices $j \in \{0, \ldots, k-1\}$. Then for each of these indices, except possibly for the first two and the last two indices, we have $N_j[p] = q$.*

Proof. Again, assume with no loss of generality that $N_j^\diamond[p] = q$ for $0 \leq j \leq h-1$, with $h \geq 5$. Suppose to the contrary that, for some $2 \leq j \leq h-3$, we have $N_j[p] \neq q$. By assumption, $N_i^\diamond[p] = q$, for each $0 \leq i \leq h-1$, and therefore we have $\varphi_i[p,q] \leq \varphi_i^\diamond[p,q] < \infty$, for each of these indices. In particular, we have $\varphi_j[p,q] \leq \varphi_j^\diamond[p,q] < \infty$, so there exists $r \in P$ for which $\varphi_j[p,r] < \varphi_j[p,q]$. Assume with no loss of generality that r lies to the left of the line from p to q. In this case we claim that $\varphi_{j-1}[p,r] < \varphi_{j-1}[p,q] < \infty$ and $\varphi_{j-2}[p,r] < \varphi_{j-2}[p,q] < \infty$.

Indeed, the boundedness of $\varphi_{j-1}[p,q]$ and $\varphi_{j-2}[p,q]$ has already been noted. Moreover, because r lies to the left of the line from p to q, the orientation of b_{pr} lies counterclockwise to that of b_{pq}. This, and our assumption that $u_j[p]$ hits b_{pr} before hitting b_{pq}, implies that the point $b_{pr} \cap b_{pq}$ lies to the right of the (oriented) line through $u_j[p]$; see Figure 8. Hence, any ray ρ emanating from p counterclockwise to $u_j[p]$ that intersects b_{pq} must also hit b_{pr} before hitting b_{pq}, so we have $\varphi_{j-1}[p,r] < \varphi_{j-1}[p,q]$ and $\varphi_{j-2}[p,r] < \varphi_{j-2}[p,q]$ (since $j \geq 2$, both $u_{j-1}[p]$ and $u_{j-2}[p]$ intersect b_{pq}), as claimed. Now applying Lemma 3.1 to the point set $\{p,q,r\}$ and to the index set $\{j-2, j-1, j\}$, we get that $\varphi_{j-1}^\diamond[p,r] < \varphi_{j-1}^\diamond[p,q]$. But this contradicts the fact that $N_{j-1}^\diamond[p] = q$. The case where r lies to the right of $\vec{pq}$ is handled in a fully symmetric manner, using the indices $\{j, j+1, j+2\}$. □

Combining Lemmas 3.1 and 3.3, we obtain the following stronger version of Theorem 1.1.

Theorem 3.4. *Let P be a set of n points in $\mathbb{R}^2$, $\alpha > 0$ a parameter, and Q a regular k-gon with $k \geq 2\pi/\alpha$. Then the following properties hold.*

(i) Every 4α-stable edge in $\mathrm{DT}(P)$ is an α-stable edge in $\mathrm{DT}^\diamond(P)$.

(ii) Every 6α-stable edge in $\mathrm{DT}^\diamond(P)$ is also an α-stable edge in $\mathrm{DT}(P)$.

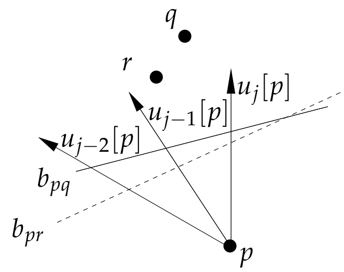

Figure 8. Proof of Lemma 3.3.

Proof. Let pq be a 4α-stable edge of $\mathrm{DT}(P)$. Then the corresponding edge e_{pq} in $\mathrm{VD}(P)$ stabs at least four rays $u_j[p]$ emanating from p, and, by Lemma 3.1, $N_j^\diamond[p] = q$ for at least two of these values of j. Therefore, p sees the edge $e_{pq}^\diamond$ in $\mathrm{VD}^\diamond(P)$ at an angle at least α. Similarly q sees the edge $e_{pq}^\diamond$ at an angle at least α. Conversely, if pq is 6α-stable in $\mathrm{DT}^\diamond(P)$ then $e_{pq}^\diamond$ meets at least six rays $u_j[p]$, and then Lemma 3.3 is easily seen to imply that p (and, symmetrically, q too) sees e_{pq} at an angle at least α. □

The next lemma gives a slightly different characterization of stable edges, which is more algorithmic and which will be useful in maintaining a SDG under a constant-degree algebraic motion of the points of P.

Lemma 3.5. *Let $\mathcal{G}$ be the subgraph of* $\mathrm{DT}^\diamond(P)$ *composed of the edges whose dual Q-Voronoi edges contain at least eleven breakpoints. Then $\mathcal{G}$ is an* $(8\alpha, \alpha)$-SDG *of P (in the Euclidean norm).*

Proof. Let $p, q \in P$ be two points. If (p, q) is an 8α-stable edge in $\mathrm{DT}(P)$ then the dual Voronoi edge e_{pq} stabs at least eight rays $u_j[p]$ emanating from p, and at least eight rays $u_j[q]$ emanating from q. Lemma 3.1 implies that $\mathrm{VD}^\diamond(P)$ contains the edge $e_{pq}^\diamond$ with at least six breakpoints corresponding to corner placements of Q at p that touch q, and at least six breakpoints corresponding to corner placements of Q at q that touch p. Therefore, $e_{pq}^\diamond$ contains at least twelve breakpoints, so $(p, q) \in \mathcal{G}$.

Conversely, suppose $p, q \in P$ define an edge $e_{pq}^\diamond$ in $\mathrm{VD}^\diamond(P)$ with at least eleven breakpoints. By the interleaving property of breakpoints, stated in Lemma 2.1, we may assume, without loss of generality, that at least six of these breakpoints correspond to empty corner placements of Q at p that touch q. Lemma 3.3 implies that $\mathrm{VD}(P)$ contains the edge e_{pq}, and that this edge is hit by at least two consecutive rays $u_j[p]$. But then the (p, q) is α-stable in $\mathrm{DT}(P)$. □

In a companion paper [3], we describe a kinetic data structure (KDS) for maintaining $\mathrm{DT}^\diamond(P)$. As shown in that paper, it can also keep track of the number of breakpoints for each edge of $\mathrm{DT}^\diamond(P)$. If each point of P moves along an algebraic trajectory of bounded degree, then the KDS processes $O(k^4 n \lambda_r(n))$ events, where r is a constant depending on the complexity of motion of P. A change in the number of breakpoints in a Voronoi edge is an event that the KDS can detect and process. As discussed in detail in [3], many events, so-called *singular* events, that occur when an edge of $\mathrm{DT}^\diamond(P)$ becomes parallel to an edge of Q, can occur simultaneously. Nevertheless, each of the events can be processed in $O(\log n)$ time, and their overall number is within the bound cited above. We maintain the subgraph $\mathcal{G}$ of $\mathrm{DT}^\diamond(P)$, consisting of the edges of $\mathrm{DT}^\diamond(P)$ that have at least *eleven* breakpoints, which, by Lemma 3.5, is an $(8\alpha, \alpha)$-Euclidean SDG. Putting everything together, we obtain a KDS that maintains an $(8\alpha, \alpha)$-SDG of P. It uses linear storage, it processes $O\left(\frac{1}{\alpha^4} n\lambda_r(n)\right)$ events, where r is a constant that depends on the degree of the motion of the points of P, and it updates the SDG at each event in $O(\log n)$ time. This proves Theorem 1.2.

Remarks. (i) We remark that $\mathrm{DT}^{\diamond}(P)$ can undergo $\Omega(n)$ changes at a time instance when $\Omega(n)$ singular events occur simultaneously, say, when pq becomes parallel to an edge of Q, but all these changes occur at the edges incident to p or q in $\mathrm{DT}(P)$. However, only $O(1/\alpha)$ edges among them can have at least eleven breakpoints, before or after the event. Hence, $O(1/\alpha)$ edges can simultaneously enter or leave the $(8\alpha, \alpha)$-SDG $\mathcal{G}$ of Theorem 1.2.

(ii) Note that there is a slight discrepancy between the value of k that we use in this section ($k \geq 2\pi/\alpha$), and the value needed to ensure that the regular k-gon Q_k is α-close to the Euclidean disk, which is $k \geq \pi/\alpha$. This is made for the convenience of presentation.

(iii) An interesting open problem is whether the dependence on α can be improved in the above KDS. We have developed an alternative scheme for maintaining stable (Euclidean) Delaunay graphs, which is reminiscent of the kinetic schemes used by Agarwal *et al.* [4] for maintaining closest pairs and nearest neighbors. It extracts a nearly linear number of pairs of points of P that are candidates for being stable Delaunay edges and then sifts the stable edges from these candidate pairs using the so-called *kinetic tournaments* [10]. Although the overall structure is not complicated, the analysis is rather technical and lengthy, so we omit this KDS from this version of the paper; it can be found in the arXiv version [2]. In summary, the resulting KDS is of size $O((n/\alpha^2)\log n)$, it processes a total of $O((n/\alpha)^2\beta_r(n/\alpha)\log^2 n \log(\log(n)/\alpha))$ events, and it takes a total of $O((n/\alpha)^2 \log^2 n(\log^2 n + \beta_r(n/\alpha)\log^2 n \log^2(\log(n)/\alpha)))$ time to process them; here $\beta_r(n) = \lambda_r(n)/n$ is an extremely slowly growing function for any fixed r. The worst-case time of processing an event is $O(\log^4(n)/\alpha))$. Another advantage of this data structure is that, unlike the above KDS, it is *local* in the terminology of [10]. Specifically, each point of P is stored, at any given time, at only $O(\log^2(n)/\alpha^2)$ places in the KDS. Therefore the KDS can efficiently accommodate an update in the trajectory of a point.

4 Stability under Nearly Euclidean Distance Functions

In this section we prove Theorem 1.1 for an arbitrary convex distance function that is α-close to the Euclidean norm (see the Introduction for the definition).

Let Q be a compact, convex set that contains the origin in its interior, and let d_Q denote the distance function induced by Q. Assume that Q is α-close to the Euclidean norm.

For sake of brevity, we carry out the proof assuming that Q is strictly convex, i.e., the relative interior of the chord connecting two point x and y on ∂Q is strictly contained in the interior of Q (there are no straight segments on the boundary of Q). The proof also holds verbatim when Q is not strictly convex, provided that no pair of points $p, q \in P$ is such that $\vec{pq}$ is parallel to a straight portion of ∂Q.

We assume that P is in general position with respect to Q, in the sense that no three points of P lie on a line, and no four points of P are *Q-cocircular*, i.e., no four points of P lie on the boundary of a common homothetic copy of Q.

Recall that for a direction u and for a point $p \in P$, $Q_p(u)$ denotes a "generic" homothetic copy of Q that touches p and is centered at some point on $u[p]$. See Figure 4 (a). As mentioned in Section 2, all homothetic copies of $Q_p(u)$ touch p at (points corresponding to) the same point $\zeta \in \partial Q$ and therefore share the same tangent $\ell_p(u)$ at p. If Q is not smooth at ζ, there is a range of possible orientations of such tangents. In this situation we let $\ell_p(u)$ denote an arbitrary tangent of this kind. In addition, if Q is smooth at ζ then $Q_{pq}(u)$ exists for any point q that lies in the same side of $\ell_p(u)$ as $Q_p(u)$. Otherwise, this has to hold for every possible tangent $\ell_p(u)$ at ζ, which is equivalent to requiring that q lies in the wedge formed by the intersection of the two halfplanes bounded by the two extreme tangents at ζ and containing $Q_p(u)$.

Remarks. (1) An important observation is that, when q satisfies these conditions, $Q_{pq}(u)$ is unique, unless all the three following conditions hold: (i) Q is not strictly convex, (ii) $\vec{pq}$ is parallel to straight portion e of ∂Q, and (iii) u is a direction connecting some point on e to the center c of Q. We leave the straightforward proof of this property to the reader. The proof of the theorem exploits the uniqueness of $Q_{pq}(u)$, and breaks down when it is not unique. In fact, this is the only way in which the assumptions concerning strict convexity are used in the proof.

(2) With some care, our analysis applies also if (the directions of) some pairs pq are parallel to straight portions of ∂Q, in which case $Q_{pq}(u)$ is not uniquely defined for certain directions u. This extension requires the more elaborate notion of α-stability, which ignores the possible two-dimensional portions of e^Q_{pq}; see Section 2 for more details. Informally, this allows us to avoid the "problematic" directions u in which $Q_{pq}(u)$ is not unique. (The latter happens exactly when $u[p]$ hits b^Q_{pq} within one of its two-dimensional portions.) We note, though, that the loss in the amount of stability caused by ignoring a two-dimensional portion of e^Q_{pq} is at most 2α. which is an upper bound on the angular span of directions that connect a straight portion of ∂Q to its center. This latter property holds since Q is α-close to the Euclidean disk; see below for more details.

The proof of Theorem 1.1 relies on the following three simple geometric properties. Recall that the α-closeness of Q to the Euclidean norm means that $D_I = (\cos\alpha)D_O \subseteq Q \subseteq D_O$.

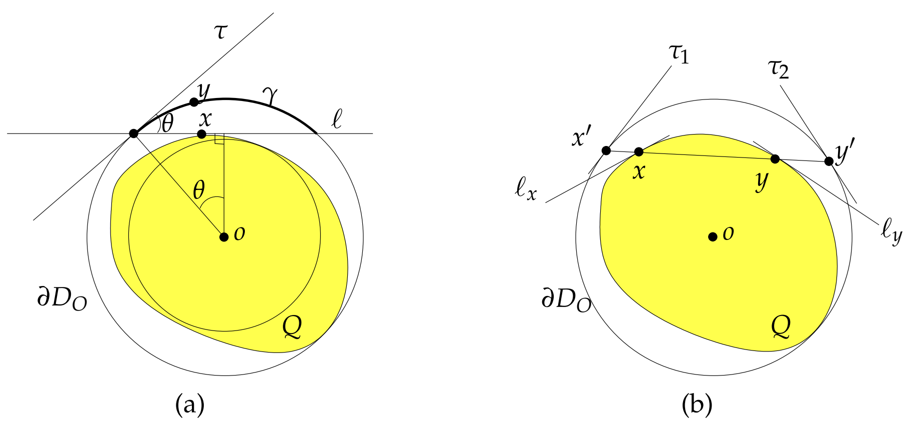

Figure 9. Illustrations for: (a) Claim 4.1; (b) Claim 4.2.

Claim 4.1. *Let x be a point on ∂Q, and let ℓ be a supporting line to Q at x. Let y be the point on ∂D_O closest to x (x and y lie on the same radius from the center o), and let γ be the arc of ∂D_O that contains y and is bounded by the intersection points of ℓ with ∂D_O. Then the angle between ℓ and the tangent, τ, to D_O at any point along γ, is at most α.*

Proof. Denote this angle by θ. Clearly θ is maximized when τ is tangent to D_O at an intersection of

ℓ and ∂D_O; see Figure 9 (a). For this value of θ, it is easy to verify that the distance from o to ℓ is $\cos\theta$. But this distance has to be at least $\cos\alpha$, because ℓ fully contains $D_I = (\cos\alpha)D_O$ on one side. Hence $\cos\theta \geq \cos\alpha$, and thus $\theta \leq \alpha$, as claimed. $\square$

Remark. An easy consequence of this claim is that the angle in which the center of Q sees any straight portion of ∂Q (when Q is not strictly convex) is at most 2α.

Claim 4.2. *Let x and y be two points on ∂Q, and let ℓ_x and ℓ_y be supporting lines of Q at x and y, respectively. Then the difference between the (acute) angles that ℓ_x and ℓ_y form with xy is at most 2α.*

Proof. Denote the two angles in the claim by θ_x and θ_y, respectively. Continue the segment xy beyond x and beyond y until it intersects D_O at x' and y', respectively. Let τ_1 and τ_2 denote the respective tangents to D_O at x' and at y'. See Figure 9 (b). Clearly, the respective angles θ_1, θ_2 between the chord $x'y'$ of D_O and τ_1, τ_2 are equal. By Claim 4.1 (applied once to τ_1 and ℓ_x and once to τ_2 and ℓ_y) we get that $|\theta_1 - \theta_x| \leq \alpha$ and $|\theta_2 - \theta_y| \leq \alpha$, and the claim follows. $\square$

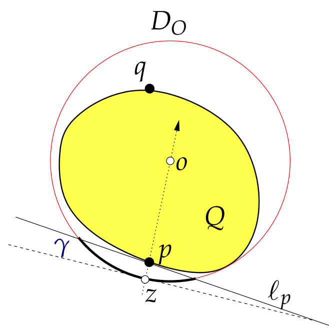

Figure 10. Proof of Claim 4.3.

Claim 4.3. *For a point $p \in \partial Q$, any tangent ℓ_p to Q at p forms an angle at most α with any line orthogonal to $\vec{op}$.*

Proof. See Figure 10. Consider the chord $\xi = \ell_p \cap D_O$, and let γ denote the arc of ∂D_O determined by ξ and containing the intersection point z of ∂D_O and the ray $\vec{op}$. By Claim 4.1, the angle between ℓ_p and the tangent to D_O at z is at most α. Since this tangent is orthogonal to $\vec{op}$, the claim follows. $\square$

Remark. Clearly, Claims 4.1–4.3 continue to hold for any homothetic copy of Q, with a corresponding translation and scaling of D_O and D_I.

Let $Q^-_{pq}(u)$ (resp., $Q^+_{pq}(u)$) denote the portion of $Q_{pq}(u)$ that lies to the left (resp., right) of the directed line from p to q. Let $D_{pq}(u)$ denote the disk that touches p and q, and whose center lies on $u[p]$.

We next establish the following lemma, whose setup is illustrated in Figure 11 (a). It provides the main geometric ingredient for the proof of Theorem 1.1.

Lemma 4.4. *(i) Let $u \in \mathbb{S}^1$ be a direction such that both $Q_{pq}(u)$ and $D_{pq}(u + 5\alpha)$ are defined. Then the region $Q^+_{pq}(u)$ is fully contained in the disk $D_{pq}(u + 5\alpha)$.*

(ii) Let $u \in \mathbb{S}^1$ be a direction such that both $Q_{pq}(u)$ and $D_{pq}(u - 5\alpha)$ are defined. Then the region $Q^-_{pq}(u)$ is fully contained in the disk $D_{pq}(u - 5\alpha)$.

Proof. It suffices to establish Part (i) of the lemma; the proof of the other part is fully symmetric.

Refer to Figure 11 (b). Let $\ell_p = \ell_p(u)$ be any supporting line of $Q_{pq}(u)$ at p, as defined above, and let $\tau_p = \tau_p(u)$ be the line through p that is orthogonal to u (which is also the tangent to $D_{pq}(u)$). By Claim 4.3, the angle between ℓ_p and τ_p is at most α. We next consider the tangent τ_p^+ to $D_{pq}(u+5\alpha)$ at p. Since the angle between τ_p and τ_p^+ is 5α, it follows that the angle between ℓ_p and τ_p^+ is at least 4α. The preceding arguments imply that, when oriented into the right side of $\vec{pq}$, ℓ_p lies between $\vec{pq}$ and τ_p^+, and the angle between ℓ_p and τ_p^+ is at least 4α.

This implies that, locally near p, $\partial Q_{pq}^+(u)$ penetrates into $D_{pq}(u+5\alpha)$. This also holds at q. To establish the claim for q (which is not symmetric to the claim for p, because the center c of $Q_{pq}(u)$ lies on the ray $u[p]$ emanating from p, and there is no control over the orientation of the corresponding ray $\vec{qc}$ emanating from q), we note that, by Claim 4.2, the angles between pq and any pair of tangents ℓ_p, ℓ_q to $Q_{pq}(u)$ at p, q, respectively, differ by at most 2α, whereas the angles between pq and the two tangents τ_p^+, τ_q^+ to $D_{pq}(u+5\alpha)$ at p, q, respectively, are equal. This, and the fact that the angle between τ_p^+ and ℓ_p is at least 4α, imply that, when oriented into the right side of $\vec{pq}$, ℓ_q lies between $\vec{qp}$ and τ_q^+, which thus implies the latter claim. Note also that the argument just given ensures that the angle between τ_q^+ and ℓ_q is at least 2α.

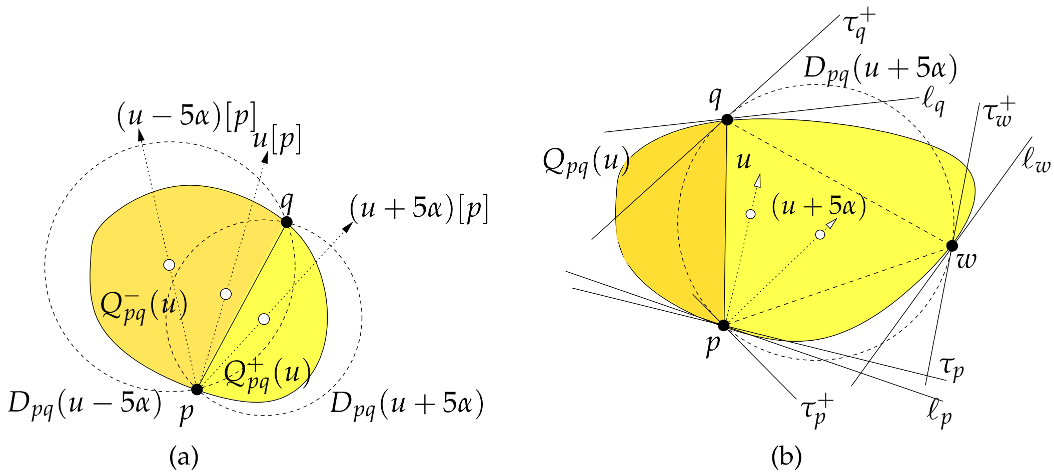

Figure 11. (a) The setup of Lemma 4.4; (b) Proof of Lemma 4.4 (i): $\partial Q_{pq}^+(u)$ cannot cross $\partial D_{pq}(u+5\alpha)$ at any third point w.

It therefore suffices to show that $\partial Q_{pq}^+(u)$ does not cross $\partial D_{pq}(u+5\alpha)$ at any third point (other than p and q). Suppose to the contrary that there exists such a third point w, and consider the tangents ℓ_w to $Q_{pq}(u)$ at w, and τ_w^+ to $D_{pq}(u+5\alpha)$ at w. Consider the two points p and w, and apply to them an argument similar to the one used above for p and q. Specifically, we use the facts that (i) the angles between pw and τ_p^+, τ_w^+ are equal, (ii) the angles between pw and ℓ_p, ℓ_w, for any tangent ℓ_w to $Q_{pq}(u)$ at w, differ by at most 2α, and (iii) the angle between ℓ_p and τ_p^+ is at least 4α, to conclude that, when oriented into the left side of $\vec{wp}$, ℓ_w lies strictly between $\vec{wp}$ and τ_w^+. See Figure 11. Similarly, applying the preceding argument to q and w, we now use the facts that (i) the angles between wq and τ_q^+, τ_w^+ are equal, (ii) the angles between qw and ℓ_q, ℓ_w differ by at most 2α, and (iii) the angle between τ_q^+ and ℓ_q is at least 2α, to conclude that, when oriented into the right side of $\vec{wq}$, ℓ_w lies between $\vec{wq}$ and τ_w^+ or coincides with τ_w^+. This impossible configuration shows that w cannot exist, and consequently that $Q_{pq}^+(u) \subset D_{pq}(u+5\alpha)$. $\square$

Proof of Theorem 1.1 – Part (i). Let pq be an 11α-stable edge in the Euclidean Delaunay triangulation $\mathrm{DT}(P)$. That is, the Euclidean Voronoi edge e_{pq} is hit by two rays $u^-[p], u^+[p]$ which form an angle of at least 11α between them (where u^+ is assumed to lie counterclockwise to u^-). Clearly, e_{pq} is also hit by any ray $u[p]$ whose direction u belongs to the interval $(u^-, u^+) \subset \mathbb{S}^1$. Let $u[p]$ be such a ray whose direction u belongs to the interval $(u^- + 5\alpha, u^+ - 5\alpha)$ (of span at least α). That is, $u[p]$ hits e_{pq} "somewhere in the middle", so all the three disks $D_{pq}(u - 5\alpha), D_{pq}(u)$ and $D_{pq}(u + 5\alpha)$ are defined and contain no points of P in their respective interiors. (Actually, $D_{pq}(u)$ is contained in $D_{pq}(u - 5\alpha) \cup D_{pq}(u + 5\alpha)$, as is easily checked.)

We next consider the Q-Voronoi diagram $\mathrm{VD}^Q(P)$. We claim that the corresponding edge e^Q_{pq} exists and is also hit by $u[p]$. Since this holds for every $u \in (u^- + 5\alpha, u^+ - 5\alpha)$, it follows that (p, q) is α-stable in $\mathrm{VD}^Q(P)$.

To establish this claim, we prove the following two properties.

(i) the homothetic copy $Q_{pq}(u)$ exists, and

(ii) it contains no points of P in its interior.

Proof of (i): Assume to the contrary that the copy $Q_{pq}(u)$ is undefined. Consider the respective tangents $\ell_p(u)$ and $\tau_p(u)$ to $Q_p(u)$ and $D_{pq}(u)$ at p, where $\ell_p(u)$ is any tangent to $Q_p(u)$ at p that separates q from $Q_p(u)$; such a tangent exists if and only if $Q_{pq}(u)$ is undefined. (As noted before, $\ell_p(u)$ does not depend on the location of the center c of Q on $u[p]$.) By Claim 4.3, the angle between $\ell_p(u)$ and $\tau_p(u)$ is at most α. Since $Q_{pq}(u)$ is undefined, the choice of $\ell_p(u)$ guarantees that q lies inside the open halfplane $h_p(u)$ bounded by $\ell_p(u)$ and disjoint from $u[p]$.

Let $\mu_p(u + 5\alpha)$ (resp., $\mu_p(u - 5\alpha)$) denote the open halfplane bounded by $\tau_p(u + 5\alpha)$ (resp., $\tau_p(u - 5\alpha)$) and disjoint from the disk $D_{pq}(u + 5\alpha)$ (resp., $D_{pq}(u - 5\alpha)$). Since each of the lines $\tau_p(u + 5\alpha)$, $\tau_p(u - 5\alpha)$ makes an angle of at least 5α with $\tau_p(u)$, the halfplane $h_p(u)$ supported by $\ell_p(u)$ is contained in the union $\mu_p(u + 5\alpha) \cup \mu_p(u - 5\alpha)$. Since q is contained in $h_p(u)$, at least one of these latter halfplanes, say $\mu_p(u + 5\alpha)$, must contain q. However, if $q \in \mu_p(u + 5\alpha)$, the corresponding copy $D_{pq}(u + 5\alpha)$ is undefined, a contradiction that establishes (i).

Proof of (ii): Since both $Q_{pq}(u)$ and $D_{pq}(u + 5\alpha)$ are defined, Lemma 4.4(i) implies that $Q^+_{pq}(u) \subset D_{pq}(u + 5\alpha)$. Moreover the interior of $D_{pq}(u + 5\alpha)$ is P-empty, so the interior of $Q^+_{pq}(u)$ is also P-empty. A symmetric argument (using Lemma 4.4(ii)) implies that the interior of $Q^-_{pq}(u)$ is also P-empty.

This completes the proof of part (i) of Theorem 1.1.

Proof of Theorem 1.1 – Part (ii). We fix a direction $u \in \mathbb{S}^1$ for which all the three copies $Q_{pq}(u)$, $Q_{pq}(u - 5\alpha)$, and $Q_{pq}(u + 5\alpha)$ are defined and have P-empty interiors. Again, $Q_{pq}(u) \subset Q_{pq}(u - 5\alpha) \cup Q_{pq}(u + 5\alpha)$. Since pq is 11α-stable under d_Q, there is an arc on $\mathbb{S}^1$ of length at least α, so that every u in this arc has this property. We need to show that, for every such u,

(i) the copy $D_{pq}(u)$ is defined, and

(ii) its interior is P-empty.

Similar to the proof of Part (i), this would imply that the ray $u[p]$ hits the edge e_{pq} of $\mathrm{VD}(P)$ for every u in an arc of length α, so pq is α-stable in $\mathrm{DT}(P)$, as claimed.

Proof of (i): Assume to the contrary that $D_{pq}(u)$ is undefined, so the angle between the vectors $\vec{pq}$ and u is at least $\pi/2$. Let $\ell_p(u-5\alpha)$, $\ell_p(u)$, and $\ell_p(u+5\alpha)$ be any triple of respective tangents to $Q_{pq}(u-5\alpha)$, $Q_{pq}(u)$, and $Q_{pq}(u+5\alpha)$ at p. Let $h_p(u-5\alpha)$ (resp., $h_p(u+5\alpha)$) be the open halfplane supported by $\ell_p(u-5\alpha)$ (resp., $\ell_p(u+5\alpha)$) and disjoint from $Q_{pq}(u-5\alpha)$ (resp., $Q_{pq}(u+5\alpha)$). Claim 4.3 implies that each of the lines $\ell_p(u-5\alpha)$, $\ell_p(u+5\alpha)$ makes with $\tau_p(u)$, the line orthogonal to $u[p]$ at p, an angle of at least 4α (and at most 6α). Indeed, the claim implies that the angle between $\ell_p(u-5\alpha)$ and the line $\tau_p(u-5\alpha)$, which is orthogonal to $(u-5\alpha)[p]$ at p, is at most α. Since the angle between $\tau_p(u-5\alpha)$ and $\tau_p(u)$ is 5α, the claim for $\ell_p(u-5\alpha)$ follows. A symmetric argument establishes the claim for $\ell_p(u+5\alpha)$. Therefore, the halfplane $\mu_p(u)$, supported by $\tau_p(u)$ and containing q, is covered by the union of $h_p(u-5\alpha)$ and $h_p(u+5\alpha)$. We conclude that at least one of these latter halfplanes must contain q. However, this contradicts the assumption that both copies $Q_{pq}(u-5\alpha)$, $Q_{pq}(u+5\alpha)$ are defined, and (i) follows.

Proof of (ii): Assume to the contrary that $D_{pq}(u)$, whose existence has just been established, contains some point r of P in its interior. That is, the ray $u[p]$ hits b_{pr} before b_{pq}. In this case $D_{pr}(u)$ also exists. With no loss of generality, we assume that r lies to the left of the oriented line from p to q.

We claim that the homothetic copy $Q_{pr}(u-5\alpha)$ exists and contains q. Indeed, since $Q_{pq}(u-5\alpha)$ exists and is P-empty, it follows that $(u-5\alpha)[p]$ either hits b^Q_{pr} after b^Q_{pq} (in which case the claim obviously holds) or misses b^Q_{pr} altogether. Suppose that $(u-5\alpha)[p]$ misses b^Q_{pr}. As argued earlier, this means that there exists a tangent $\ell_p(u-5\alpha)$ to $Q_p(u-5\alpha)$ at p, such that r lies in the open halfplane $h_p(u-5\alpha)$ supported by $\ell_p(u-5\alpha)$ and disjoint from $Q_p(u-5\alpha)$.

By applying Claim 4.3 as before we get that the tangent $\tau_p(u)$ to $D_{pq}(u)$ (at p) to the left of $\vec{pq}$ is between $\vec{pq}$ and $\ell_p(u-5\alpha)$ and makes with $\ell_p(u-5\alpha)$ an angle of at least 4α. It follows that the wedge formed by the intersection of $\ell_p(u-5\alpha)$ and the halfplane to the left of $\vec{pq}$ is fully contained in the halfplane $\mu_p(u)$ that is supported by $\tau_p(u)$ and disjoint from $D_{pq}(u)$; see Figure 12 (a). But then $D_{pr}(u)$ is undefined, a contradiction that implies the existence of $Q_{pr}(u-5\alpha)$.

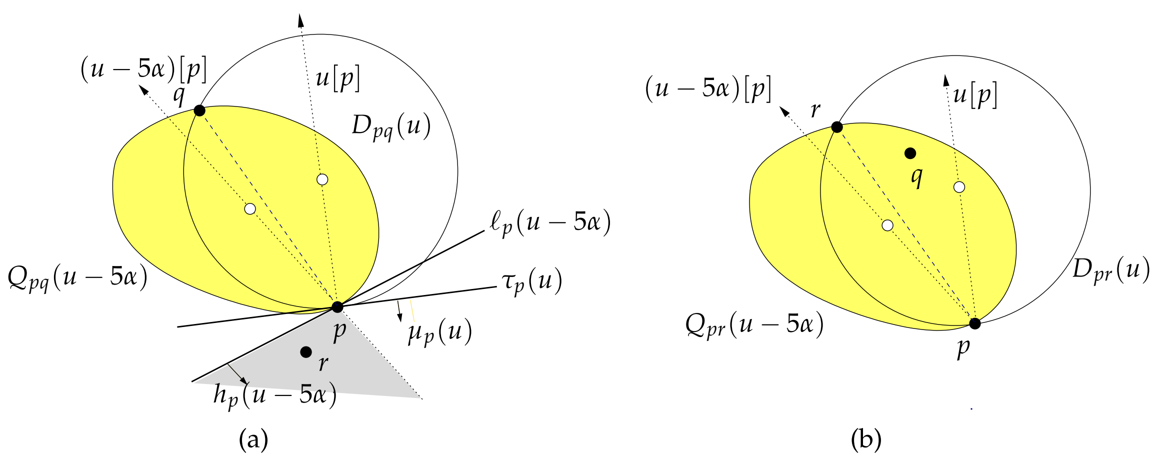

Figure 12. Proof of Theorem 1.1(ii): (a) Arguing (by contradiction) that $Q_{pr}(u-5\alpha)$ exists; (b) the copy $Q_{pr}(u-5\alpha)$ is defined and contains q. Hence, the disk $D_{pr}(u)$, which covers $Q^+_{pr}(u-5\alpha)$, must contain q as well.

We can now assume that $Q_{pr}(u-5\alpha)$ is defined and contains q. More precisely, q lies in the portion $Q^+_{pr}(u-5\alpha)$ of $Q_{pr}(u-5\alpha)$, since q lies to the right of the oriented line from p to r. However, Lemma 4.4(i), applied to $u - 5\alpha$, implies that $Q^{+}_{pr}(u - 5\alpha)$ is contained in $D_{pr}(u)$, so $D_{pr}(u)$ also contains q; see Figure 12 (b). It is however impossible for $D_{pr}(u)$ to contain q and for $D_{pq}(u)$ to contain r. This contradiction concludes the proof of part (ii) of Theorem 1.1. $\square$

5 Properties of stabe Delaunay graphs

We establish a few useful properties of stable Delaunay graphs in this section.

Near cocircularities do not show up in an SDG. Consider a critical event during the kinetic maintenance of $\mathrm{DT}(P)$, in which four points a, b, c, d become cocircular, in this order, along their circumcircle, with this circle being P-empty. Just before the critical event, the Delaunay triangulation contained two triangles formed by this quadruple, say, abc, acd. The Voronoi edge e_{ac} then shrinks to a point (namely, to the circumcenter of $abcd$ at the critical event), and, after the critical cocircularity, is replaced by the Voronoi edge e_{bd}, which expands from the circumcenter as time progresses.

Our algorithm will detect the possibility of such an event before the criticality occurs, when ac ceases to be α-stable (or even before this happens). It will then remove this edge from the stable subgraph, so the actual cocircularity will not be recorded. The new edge bd will then be detected by the algorithm only when it becomes at least α-stable (if this happens at all), and will then enter the stable Delaunay graph. In short, critical cocircularities do not arise *at all* in our scheme.

As noted in the introduction, a Delaunay edge ab (interior to the hull) transitions from being α-stable to not being stable, or vice-versa, when the sum of the opposite angles in its two adjacent Delaunay triangles is $\pi - \alpha$ (see Figure 2). This shows that changes in the stable Delaunay graph occur when the "cocircularity defect" of a nearly cocircular quadruple (i.e., the difference between π and the sum of opposite angles in the quadrilateral spanned by the quadruple) is between α and 8α. Note that a degenerate case of cocircularity is a collinearity on the convex hull. Such collinearities also do not show up in the stable Delaunay graph. A hull collinearity between three nodes a, b, c is detected before it happens, when (or before) the corresponding Delaunay edge is no longer α-stable, in which case the angle $\angle acb$, where c is the middle point of the (near-)collinearity becomes $\pi - \alpha$ (see Figure 13(a)). Therefore a hull edge is removed from the SDG if the Delaunay triangle is almost flat. The edge (or any new edge about to replace it) re-appears in the SDG when it becomes $c\alpha$-stable, for some $1 \leq c \leq 8$.

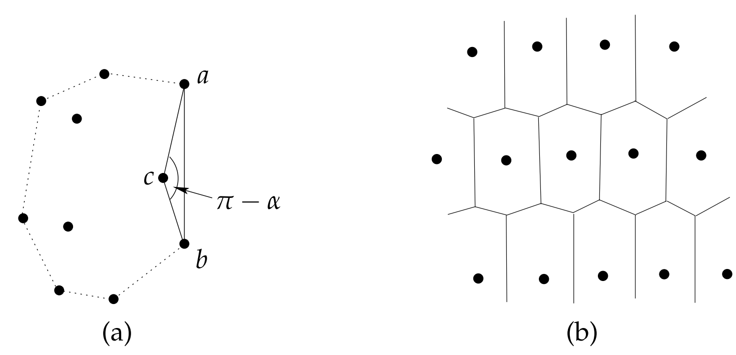

Figure 13. (a) The near collinearity that corresponds to a Delaunay edge ceasing to be α-stable. (b) A set of points for which the number of α-stable edges in $\mathrm{DT}(P)$ (those corresponding to the vertical Voronoi edges) is close to n.

SDGs are not too sparse. Let P be a set of n points in the plane. We give a lower bound on the number of α-stable Delaunay edges in the Delaunay triangulation of P. Our lower bound approaches n as α decreases to zero.

Let n_0 be the number of points with no incident α-stable edges in $\mathrm{DT}(P)$ and let n_1 be the number of points with a single incident α-stable edge in $\mathrm{DT}(P)$. Clearly the total number of α-stable edges in $\mathrm{DT}(P)$ is at least

$$\frac{2(n - n_0 - n_1) + n_1}{2} = n - \left(\frac{2n_0 + n_1}{2}\right) . \tag{3}$$

We now derive an upper bound on $\frac{2n_0+n_1}{2}$. Consider a vertex v with no incident α-stable edges. If v is not a vertex of the convex hull then its degree in $\mathrm{DT}(P)$ must be at least $2\pi/\alpha$ (the boundary of its cell in $\mathrm{VD}(P)$ contains at least $2\pi/\alpha$ α-short edges). If v is a vertex of the convex hull then its degree must be at least $(2\pi - d(v))/\alpha$ where $d(v)$ is the angle between the two infinite rays bounding the Voronoi cell of v. Similarly, consider a vertex v with one incident α-stable edge. If v is not a vertex of the convex hull then its degree must be at least $\pi/\alpha + 1$ and if v is a vertex of the hull then its degree is at least $(\pi - d(v))/\alpha + 1$. Since $\sum d(v)$ over all hull vertices is 2π, we get that the sum of the degrees of the vertices in $\mathrm{DT}(P)$ is at least

$$n_1\left(\frac{\pi}{\alpha} + 1\right) + 2n_0\frac{\pi}{\alpha} - \frac{2\pi}{\alpha} . \tag{4}$$

On the other hand, the Delaunay triangulation of any set with h vertices on the convex hull consists of $3n - h - 3$ edges so the sum of the degrees is $6n - 2h - 6$. Combining this with the lower bound in (4) we get that

$$n_1\left(\frac{\pi}{\alpha} + 1\right) + 2n_0\frac{\pi}{\alpha} - \frac{2\pi}{\alpha} \le 6n,$$

which implies that

$$\left(\frac{2n_0 + n_1}{2}\right) \le \frac{6n}{\pi/\alpha} + 2 .$$

Substituting this upper bound in Equation (3) we get that the number of α-stable edges in $\mathrm{DT}(P)$ is at least

$$n\left(1 - \frac{6\alpha}{\pi}\right) - 2 .$$

This is nearly tight, since, for any α, there exist sets of n points for which the number of α-stable edges is roughly n; see Figure 13(b).

Closest pairs, crusts, β-skeleta, and the SDG. Let $\beta \ge 1$, and let P be a set of n points in the plane. The *β-skeleton* of P is a graph on P that consists of all the edges pq such that the union of the two disks of radius $(\beta/2)d(p,q)$, touching p and q, does not contain any point of $P \setminus \{p,q\}$. See, e.g., [7, 16] for properties of the β-skeleton, and for its applications in surface reconstruction. We claim that the edges of the β-skeleton are α-stable in $\mathrm{DT}(P)$, provided $\beta \ge 1 + \Omega(\alpha^2)$. Indeed, let pq be an edge of the β-skeleton of P, for $\beta > 1$. Let c_1 and c_2 be the centers of the two empty disks of radius $(\beta/2)d(p,q)$ touching p and q; see Figure 14(a). Clearly $\angle c_1pq = \angle c_2pq$. Denote $\theta = \angle c_1pq = \angle c_2pq$. Each of p, q sees the Voronoi edge e_{pq} at an angle at least 2θ, so it is 2θ-stable.

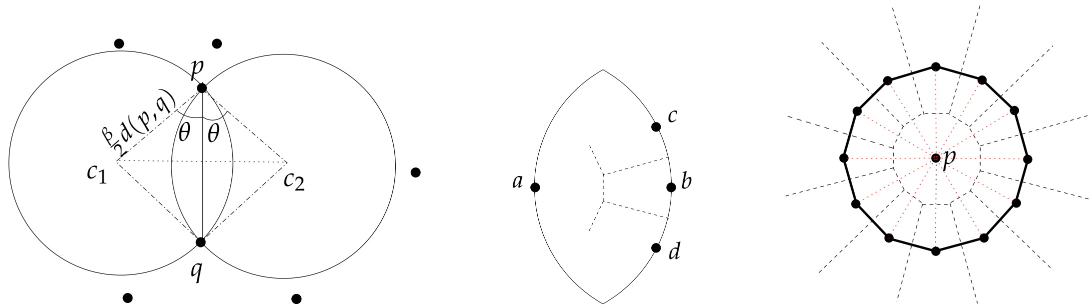

(a) (b) (c)

Figure 14. (a) The skeleton edge pq is a stable edge. (b) An edge ab of the Relative Neighborhood Graph that is not stable. (c) Point p is disconnected in the SDG. VD is drawn with dashed lines, SDG with solid lines, and the remaining DT edges with dotted lines.

We have $1/\beta = \cos\theta \approx 1 - \theta^2/2$ or $\beta = 1 + \Theta(\theta^2)$. That is, for $\beta \geq 1 + \Omega(\alpha^2)$, with an appropriate (small) constant of proportionality, pq is α-stable.

A similar argument shows that the stable Delaunay graph contains the closest pair in $P(t)$ as well as the *crust* of a set of points sampled sufficiently densely along a 1-dimensional curve (see [6, 7] for the definition of crusts and their applications in surface reconstruction). We sketch the argument for closest pairs: If (p, q) is a closest pair then $pq \in \mathrm{DT}(P)$, and the two adjacent Delaunay triangles $\triangle pqr^+, \triangle pqr^-$ are such that their angles at r^+, r^- are at most $\pi/3$ each, so e_{pq} is $(\pi/3)$-long, ensuring that pq belongs to any stable subgraph for α sufficiently small; see [4] for more details. We omit the proof for crusts, which is fairly straightforward too.

Remark. Stable Delaunay graphs need not contain all the edges of several other important subgraphs of the Delaunay triangulation, including the Euclidean minimum spanning tree, the Gabriel graph, the relative neighborhood graph, and the all-nearest-neighbors graph. An illustration for the relative neighborhood graph (RNG) is given in Figure 14 (b). Recall that an edge pq is in RNG if there is no point $r \in P$ such that $\max\{\|pr\|, \|qr\|\} < \|pq\|$. As shown in figure that pq is an edge of RNG, but the angular extent of the dual Voronoi edge e_{pq} can be arbitrarily small. As a matter of fact, the stable Delaunay graph need not even be connected, as is illustrated in Figure 14(c).

6 Conclusion

In this paper we introduced the notion of a stable Delaunay graph (SDG), a large subgraph of the Delaunay triangulation, which retains several useful properties of the full Delaunay triangulation. We proved that a 4α-stable edge in (the Euclidean) $\mathrm{DT}(P)$ is α-stable in $\mathrm{DT}^Q(P)$, where Q is a regular k-gon for any $k \geq 2\pi/\alpha$ and α is the stability parameter, and that the dual Voronoi edge $e^\diamond_{pq}$ in $\mathrm{VD}^\diamond(P)$ contains at least eleven breakpoints. Using these properties and the kinetic data structure for $\mathrm{DT}^\diamond(P)$ developed in the companion paper [3], we presented a linear-size KDS for maintaining a Euclidean $(8\alpha, \alpha)$-SDG as the input points move. The KDS processes only a nearly quadratic number of events if the points move along algebraic trajectories of bounded degree, and each event can be processed in $O(\log n)$ time. We also showed that if an edge is stable in the

Delaunay triangulation under the Euclidean norm, it is also stable in the Delaunay triangulation under any convex distance function sufficiently close to the Euclidean norm, and vice versa.

Proving a subcubic upper bound on the number of topological changes in the Euclidean Delaunay triangulation for a set of moving points still remains elusive (in spite of the recent progress in [19, 20]), but our result implies that if the true bound is really close to cubic, or just significantly super-quadratic, then the overwhelming majority of these changes involve edges appearing and disappearing while the four vertices of the two triangles adjacent to each such edge remain nearly cocircular throughout the entire time in which the edge exists.

We conclude by mentioning two open problems:

(i) Is there a KDS for maintaining a triangulation of the *entire* convex hull of a set of moving points in the plane, which is an *approximate* Delaunay triangulation, defined appropriately, and which processes only a near-quadratic number of events? In particular, can the SDG maintained by our KDS be extended to a triangulation scheme of $\mathrm{conv}(P)$ (recall Figure 1(b)), e.g., using the ideas from the kinetic triangulation schemes presented in [5, 15], which also undergoes only a near-quadratic number of topological changes during the motion? Perhaps a deeper analysis of the structure of the "holes" in the stable sub-diagram may yield a solution to this problem, using the fact that for every missing edge, the two incident triangles form a near-cocircularity in the diagram. This may lead to a scheme that fills in the holes by near-Delaunay edges and that has the desired properties.

(ii) What are the other large and interesting subgraphs of $\mathrm{DT}(P)$ that undergo only a near-quadratic number of topological changes under a motion of the points of P of the above kind? For instance, can one prove that there are only a near-quadratic number of changes in the α-shape or the relative neighborhood graph of P if the points of P move along algebraic trajectories of bounded degree.

References

[1] P. K. Agarwal, D. Eppstein, L. J. Guibas, and M. R. Henzinger, Parametric and kinetic minimum spanning trees, *Proc. 39th Annual IEEE Sympos. Found. Comp. Sci.*, 1998, 596–605.

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

[3] P. K. Agarwal, H. Kaplan, N. Rubin and M. Sharir, Kinetic Voronoi diagrams and Delaunay triangulations under polygonal distance functions, manuscript, 2014.

[4] P. K. Agarwal, H. Kaplan and M. Sharir, Kinetic and dynamic data structures for closest pair and all nearest neighbors, *ACM Trans. Algorithms* 5 (2008), Art. 4.

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

[6] N. Amenta and M. Bern, Surface reconstruction by Voronoi filtering, *Discrete Comput. Geom.* 22 (1999), 481–504.

[7] N. Amenta, M. W. Bern and D. Eppstein, The crust and beta-skeleton: combinatorial curve reconstruction, *Graphic. Models and Image Processing* 60 (2) (1998), 125–135.

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

[9] J. Basch, L. J. Guibas, L. Zhang, Proximity problems on moving points, *Proc. 13th Annu. Sympos. Comput. Geom.*, 1997, 344–351.

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

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

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

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

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

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

[16] D. Kirkpatrick and J. D. Radke, A framework for computational morphology, in G. Toussaint, ed., *Computational Geometry*, North-Holland, 1985, 217–248.

[17] 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.

[18] Z. Rahmati, M. Ali Abam, V. King, S. Whitesides, and A. Zarei, A dimple, faster method for kinetic proximity problems, CoRR abs/1311.2032 (2013).

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

[20] N. Rubin, On kinetic Delaunay triangulations: a near quadratic bound for unit speed motions, *Proc. 54th Annu. IEEE Sympos. Found. Comp. Sci.*, 2013, 519–528.

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

Acknowledgements. P.A. and M.S. were supported by Grant 2012/229 from the U.S.-Israel Binational Science Foundation. P.A. was also supported by NSF under grants CCF-09-40671, CCF-10-12254, and CCF-11-61359, and by an ERDC contract W9132V-11-C-0003. L.G. was supported by NSF grants CCF-10-11228 and CCF-11-61480. H.K. was 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). N.R. was 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). M.S. was 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.