

Kinetic Stable Delaunay Graphs

Pankaj K. Agarwal
Duke University

Haim Kaplan
Tel Aviv University

Jie Gao
SUNY Stony Brook

Vladlen Koltun
Stanford University

Leonidas J. Guibas
Stanford University

Natan Rubin
Tel Aviv University

Micha Sharir
Tel Aviv University and
New York University

ABSTRACT

The best known upper bound on the number of topological changes in the Delaunay triangulation of a set of moving points in \mathbb{R}^2 is (nearly) cubic, even if each point is moving with a fixed velocity. We introduce the notion of a *stable Delaunay graph* (SDG in short), a dynamic subgraph of the Delaunay triangulation, that is less volatile in the sense that it undergoes fewer topological changes and yet retains many useful properties of the full Delaunay triangulation. SDG is defined in terms of a parameter $\alpha > 0$, and consists of Delaunay edges pq for which the (equal) angles at which p and q see the corresponding Voronoi edge e_{pq} are at least α . We prove several interesting properties of SDG and describe two kinetic data structures for maintaining it. Both structures use $O^*(n)$ storage. They process $O^*(n^2)$ events during the motion, each in $O^*(1)$ time, provided that the points of P move along algebraic trajectories of bounded degree; the $O^*(\cdot)$ notation hides multiplicative factors that are polynomial in $1/\alpha$ and polylogarithmic in n . The first structure is simpler but the dependency on $1/\alpha$ in its performance is higher.

Categories and Subject Descriptors

F.2.2 [Analysis of algorithms and problem complexity]: Non-numerical algorithms and problems—*Geometrical problems and computations*; G.2.1 [Discrete mathematics]: Combinatorics—*Combinatorial algorithms*

General Terms

Algorithms, Theory

Keywords

Delaunay triangulation, Voronoi diagram, kinetic data structures

1. INTRODUCTION

Delaunay triangulation and Voronoi diagram. Let P be a (finite) set of points in \mathbb{R}^2 . Let $\text{VD}(P)$ and $\text{DT}(P)$ denote the Voronoi di-

agram and Delaunay triangulation of P , respectively. For a point $p \in P$, let $\text{Vor}(p)$ denote the Voronoi cell of p . The Delaunay triangulation $\text{DT} = \text{DT}(P)$ consists of all triangles whose circumcircles do not contain points of P in their interior. Its edges form the *Delaunay graph*, which is the straight-edge dual graph of the Voronoi diagram of P . That is, pq is an edge of the Delaunay graph if and only if $\text{Vor}(p)$ and $\text{Vor}(q)$ share an edge, which we denote by e_{pq} . This is equivalent to the existence of a circle passing through p and q that does not contain any point of P in its interior—any circle centered at a point on e_{pq} and passing through p and q is such a circle. Delaunay triangulations and Voronoi diagrams are fundamental to much of computational geometry and its applications. See [6, 11] for a survey and a textbook on these structures.

In many applications of Delaunay/Voronoi methods (e.g., mesh generation and kinetic collision detection) the points are moving continuously, so these diagrams need to be efficiently updated as motion occurs. Even though the motion of the nodes is continuous, the combinatorial and topological structure of the Voronoi and Delaunay diagrams change only at discrete times when certain critical events occur. Their evolution under motion can be studied within the framework of *kinetic data structures* (KDS in short) of Basch *et al.* [7, 12, 13], a general methodology for designing efficient algorithms for maintaining such combinatorial attributes of mobile data.

For the purpose of kinetic maintenance, Delaunay triangulations are nice structures, because, as mentioned above, they admit local certifications associated with individual triangles. This makes it simple to maintain DT under point motion: an update is necessary only when one of these empty circumcircle conditions fails—this corresponds to co-circularities of certain subsets of four points. Whenever such an event happens, a single edge flip easily restores Delaunayhood. Estimating the number of such events, however, has been elusive—the problem of bounding the number of combinatorial changes in DT for points moving with constant velocities has been in the computational geometry lore for many years.

Let n be the number of moving points in P . We assume that each point moves along an algebraic trajectory of fixed degree (see Section 2 for a more formal definition). Guibas *et al.* [14] showed a roughly cubic upper bound of $O(n^2 \lambda_s(n))$ on the number of discrete (also known as *topological*) changes in DT, where $\lambda_s(n)$ is the maximum length of an (n, s) -Davenport-Schinzel sequence [21], and s is a constant depending on the motions of the points. A substantial gap exists between this upper bound and a quadratic lower bound [21].

It is thus desirable to find approaches for maintaining a substantial portion of DT that *provably* experiences only a nearly quadratic

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'10, June 13–16, 2010, Snowbird, Utah, USA.

Copyright 2010 ACM 978-1-4503-0016-2/10/06 ...\$10.00.

number of discrete changes, that is reasonably easy to define and maintain, and that retains useful properties for further applications.

Polygonal distance functions. If the “unit ball” of our underlying norm is *polygonal* then things improve considerably. In more detail, let Q be a convex polygon with a constant number, k , of edges. It induces a *convex distance function*

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

d_Q is a metric if Q is centrally symmetric with respect to the origin.

We can define the Q -Voronoi diagram of a set P of points in the plane in the usual way, as the partitioning of the plane into Voronoi cells, so that the cell $\text{Vor}^\circ(p)$ of a point p is $\{x \in \mathbb{R}^2 \mid d_Q(x, p) = \min_{p' \in P} d_Q(x, p')\}$. Assuming that the points of P are in general position with respect to Q , these cells are nonempty, have pairwise disjoint interiors, and cover the plane.

As in the Euclidean case, the Q -Voronoi diagram of P has its dual representation, which we refer to as the *Q -Delaunay triangulation* $\text{DT}^\circ(P) = \text{DT}^\circ$. A triple of points in P define a triangle in DT° if and only if they lie on the boundary of some homothetic copy of Q that does not contain any point of P in its interior. Assuming that P is in general position, these Q -Delaunay triangles form a triangulation of a certain simply-connected polygonal region that is contained in the convex hull of P . Unlike the Euclidean case, it does not always coincide with the convex hull (see Figures 3 and 11 for examples). See Chew and Drysdale [9] and Leven and Sharir [18] for analysis of Voronoi and Delaunay diagrams of this kind.

For kinetic maintenance, polygonal Delaunay triangulations are “better” than Euclidean Delaunay triangulations because, as shown by Chew [8], when the points of P move (in the algebraic sense assumed above), the number of topological changes in the Q -Delaunay triangulation is only nearly quadratic in n . One of the major observations in this paper is that the *stable portions* of the Euclidean Delaunay triangulation and the Q -Delaunay triangulation are closely related.

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 Δpqr^+ and Δpqr^- be the two Delaunay triangles incident to pq . Then pq is called α -stable if its opposite angles in these triangles satisfy $\angle pr^+q + \angle pr^-q < \pi - \alpha$. (The case where pq lies on the convex hull of P is treated as if one of r^+, r^- lies at infinity, so that the corresponding angle $\angle pr^+q$ or $\angle pr^-q$ 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 equal angles at which p and q see their common Voronoi edge e_{pq} are at least α . See Figure 1 (left).

A justification for calling such edges stable lies in the following observation: If a Delaunay edge pq is α -stable then it remains in DT during any continuous motion 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 at the time pq is α -stable we have $\angle pr^+q + \angle pr^-q < \pi - \alpha$ for any pair of points r^+, r^- lying on opposite sides of the line ℓ supporting pq , so, if each of these angles change by at most $\alpha/2$ we still have $\angle pr^+q + \angle pr^-q \leq \pi$, which is easily seen to imply that pq remains an edge of DT. (This argument also covers the cases when a point crosses ℓ from side to side.) 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 DT are those for p and q are almost co-circular with their two common Delaunay neighbors r^+, r^- , and hence more likely to get flipped “soon”.

Overview of our results. Let $\alpha > 0$ be a fixed parameter. In this paper we show how to maintain a subgraph of the full Delaunay triangulation DT, which we call a $(c\alpha, \alpha)$ -stable Delaunay graph (SDG in short), so that (i) every edge of SDG is α -stable, and (ii) every $c\alpha$ -stable edge of DT belongs to SDG, where $c > 1$ is some (small) absolute constant. Note that SDG is not uniquely defined, even when c is fixed.

In Section 2, we introduce several useful definitions and show that the number of discrete changes in the SDGs that we consider is nearly quadratic. What this analysis also implies is that if the true bound for kinetic changes in a Delaunay triangulation is close to cubic, then the overwhelming majority of these changes involve edges which never become stable and just flicker in and out of the diagram by co-circularity with their two Delaunay neighbors.

In Sections 3 and 4 we show that SDG can be maintained by a kinetic data structure that uses only near-linear storage (in the terminology of [7], it is *compact*), encounters only a nearly quadratic number of critical events (it is *efficient*), and processes each event in polylogarithmic time (it is *responsive*). For the second data structure described in Section 4, each point appears at any time in only polylogarithmically many places in the structure (it is *local*).

The scheme described in Section 3 is based on a useful and interesting “equivalence” connection between the (Euclidean) SDG and a suitably defined “stable” version of the Delaunay triangulation of P under the “polygonal” norm whose unit ball Q is a regular k -gon, for $k = \Theta(1/\alpha)$. As noted above, Voronoi and Delaunay structures under polygonal norms are particularly favorable for kinetic maintenance because of Chew’s result [8], showing that the number of topological changes in $\text{DT}^\circ(P)$ is $O^*(n^2k^4)$; the $O^*(\cdot)$ notation hides a factor that depends sub-polynomially on both n and k . In other words, the scheme simply maintains the “polygonal” diagram $\text{DT}^\circ(P)$ in its entirety, and selects from it those edges that are also stable edges of the Euclidean diagram DT.

The major disadvantage of the solution in Section 3 is the rather high (proportional to $\Theta(1/\alpha^4)$) dependence on $1/\alpha$ ($\approx k$) of the bound on the number of topological changes. We do not know whether the upper bound $O^*(n^2k^4)$ on the number of topological changes in $\text{DT}^\circ(P)$ is nearly tight (in its dependence on k). To remedy this, we present in Section 4 an alternative scheme for maintaining stable (Euclidean) Delaunay edges. The scheme is reminiscent of the kinetic schemes used in [2] for maintaining closest pairs and nearest neighbors. It extracts $O^*(n)$ pairs of points of P that are candidates for being stable Delaunay edges. Each point $p \in P$ then runs $O(1/\alpha)$ kinetic and dynamic tournaments involving the other points in its candidate pairs. Roughly, these tournaments correspond to shooting $O(1/\alpha)$ rays from P in fixed directions and finding along each ray the nearest point equally distant from p and from some other candidate point q . We show that pq is a stable Delaunay edge if and only if q wins many consecutive tournaments of p (or p wins many consecutive tournaments of q). A careful analysis shows that the number of events that this scheme processes (and the overall processing time) is only $O^*(n^2/\alpha^2)$.

Section 5 reviews several useful properties of stable Delaunay graphs. In particular, we show that at any given time the stable subgraph contains at least $\left[1 - \frac{3}{2(\pi/\alpha - 2)}\right] n$ Delaunay edges, i.e., at least about one third of the maximum possible number of edges. In addition, we show that at any moment the SDG contains the closest pair, the so-called β -skeleton of P , for $\beta = 1 + \Omega(\alpha^2)$ (see [5, 17]), and the *crust* of a sufficiently densely sampled point set along a smooth curve (see [4, 5]). We also extend the connection in Section 3 to arbitrary distance functions d_Q whose unit ball Q is sufficiently close (in the Hausdorff sense) to the Euclidean one (i.e.,

the unit disk). Completing $DT^\diamond(P)$ into a triangulation of the entire hull is an interesting challenge that we are currently exploring in a companion paper [16].

2. PRELIMINARIES

Stable edges in Voronoi diagrams. Let $\{u_0, \dots, u_{k-1}\} \subset \mathbb{S}^1$ be a set of $k = \Theta(1/\alpha)$ equally spaced directions in \mathbb{R}^2 . For concreteness take $u_i = (\cos(2\pi i/k), -\sin(2\pi i/k))$, $0 \leq i < k$ (so our directions u_i go clockwise as i increases).¹ For a point $p \in P$ and a (unit) vector u let $u[p]$ denote the ray $\{p + \lambda u \mid \lambda \geq 0\}$ that emanates from p in direction u . For a pair of points $p, q \in P$ let b_{pq} denote the perpendicular bisector of p and q . If b_{pq} intersects $u_i[p]$, then the expression

$$\varphi_i[p, q] = \frac{\|q - p\|^2}{2\langle q - p, u_i \rangle} \quad (1)$$

is the distance between p and the intersection point of b_{pq} with $u_i[p]$. If b_{pq} does not intersect $u_i[p]$ we define $\varphi_i[p, q] = \infty$. The point q minimizes $\varphi_i[p, q]$, among all points q' for which $b_{pq'}$ intersects $u_i[p]$, 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 $N_i(p)$; see Figure 1 (right).

The (*angular extent*) of a Voronoi edge e_{pq} of two points $p, q \in P$ is the angle at which it is seen from either p or q (these two angles are equal). For a given angle $\alpha \leq \pi$, e_{pq} is called α -long (resp., α -short) if the extent of e_{pq} is at least (resp., smaller than) α . We also say that $pq \in DT(P)$ is α -long (resp., α -short) if e_{pq} is α -long (resp., α -short). As noted in the introduction, these notions can also be defined (equivalently) in terms of the angles in the Delaunay triangulation: A Delaunay edge pq is α -long if and only if $\angle pr^+q + \angle pr^-q \leq \pi - \alpha$, where $\triangle pr^+q$ and $\triangle pr^-q$ are the two Delaunay triangles incident to pq . See Figure 1(left).

Given parameters $\alpha' > \alpha > 0$, we seek to construct (and maintain under motion) an (α', α) -stable Delaunay graph (or *stable Delaunay graph*, for brevity, which we abbreviate as SDG) of P , which is any subgraph G of $DT(P)$ with the following properties:

- (S1) Every α' -long edge of $DT(P)$ is an edge of G .
- (S2) Every edge of G is an α -long edge of $DT(P)$.

An (α', α) -stable Delaunay graph need not be unique. In what follows, α' will always be some fixed (and reasonably small) multiple of α .

Kinetic tournaments. Kinetic tournaments were first studied by Basch *et al.* [7], for kinetically maintaining the lowest point in a set P of n points moving on some vertical line, say the y -axis, so that their trajectories are algebraic of bounded degree, as above. We also want the tournament to be dynamic, so that we can insert and delete points into/from it. This can be achieved by maintaining the tournament as a heap, stored as a *weight-balanced* ($BB(\alpha)$) tree [20] (and [19]). Omitting all further details, which can be found in [2] (and in the full version [1]), we state:

THEOREM 2.1 (AGARWAL *et al.* [2]). *A sequence of m insertions and deletions into a kinetic tournament, whose maximum size at any time is n (assuming $m \geq n$), when implemented as a weight-balanced tree, generates at most $O(m\beta_{r+2}(n) \log n)$ events, with a total processing cost of $O(m\beta_{r+2}(n) \log^2 n)$. Here r is the maximum number of times a pair of points intersect, and $\beta_{r+2}(n) = \lambda_{r+2}(n)/n$. Processing an update or a tournament event takes $O(\log^2 n)$ worst-case time. A dynamic kinetic tournament on n elements can be constructed in $O(n)$ time.*

¹The index arithmetic is modulo k , i.e., $u_i = u_{i+k}$.

Maintenance of an SDG. Let $P = \{p_1, \dots, p_n\}$ be a set of points moving in \mathbb{R}^2 . Let $p_i(t) = (x_i(t), y_i(t))$ denote the position of p_i at time t . We call the motion of P *algebraic* if each $x_i(\cdot), y_i(\cdot)$ is a polynomial, and the *degree* of motion of P is the maximum degree of these polynomials. Throughout this paper we assume that the motion of P is algebraic and that its degree is bounded by a constant. In this subsection we present a simple technique for maintaining a $(2\alpha, \alpha)$ -stable Delaunay graph. Unfortunately this algorithm requires quadratic space. It is based on the following easy observation (see Figure 1 (right)), where k is an integer, and the unit vectors (directions) u_0, \dots, u_{k-1} are as defined earlier.

LEMMA 2.2. *Let $\alpha = 2\pi/k$. (i) If the extent of e_{pq} is larger than 2α then there are two consecutive directions u_i, u_{i+1} , such that q is the neighbor of p in directions u_i and u_{i+1} . (ii) If there are two consecutive directions u_i, u_{i+1} , such that q is the neighbor of p in both directions u_i and u_{i+1} , then the extent of e_{pq} is at least α .*

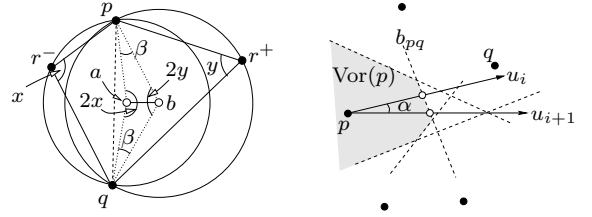


Figure 1. Left: The points p and q see their common Voronoi edge ab at (equal) angles β . This is equivalent to the angle condition $x + y = \pi - \beta$ for the two adjacent Delaunay triangles. Right: q is the neighbor of p in the directions u_i and u_{i+1} , so the Voronoi edge e_{pq} is α -long.

The algorithm maintains Delaunay edges pq such that there are two consecutive directions u_i and u_{i+1} along which q is the neighbor of p . For each point p and direction u_i we get a set of at most $n - 1$ piecewise continuous functions of time, $\varphi_i[p, q]$, one for each point $q \neq p$, as defined in (1). (Recall that $\varphi_i[p, q] = \infty$ when $u_i[p]$ does not intersect b_{pq} .) By assumption on the motion of P , for each p and q , the domain in which $\varphi_i[p, q](t)$ is defined consists of a constant number of intervals.

For each point p , and ray $u_i[p]$, consider each function $\varphi_i[p, q]$ as the trajectory of a point moving along the ray and corresponding to q . The algorithm maintains these points in a dynamic and kinetic tournament $K_i(p)$ (see Theorem 2.1) that keeps track of the minimum of $\{\varphi_i[p, q](t)\}_{q \neq p}$ over time. For each pair of points p and q such that q wins in two consecutive tournaments, $K_i(p)$ and $K_{i+1}(p)$, of p , it keeps the edge pq in the stable Delaunay graph. It is trivial to update this graph as a by-product of the updates of the various tournaments. The analysis of this data structure is straightforward using Theorem 2.1, and yields the following result.

THEOREM 2.3. *Let P be a set of n moving points in \mathbb{R}^2 under algebraic motion of bounded degree, let k be an integer, and let $\alpha = 2\pi/k$. A $(2\alpha, \alpha)$ -stable Delaunay graph of P can be maintained using $O(kn^2)$ storage and processing $O(kn^2\beta_{r+2}(n) \log n)$ events, for a total cost of $O(kn^2\beta_{r+2}(n) \log^2 n)$ time. The processing of each event takes $O(\log^2 n)$ worst-case time. Here r is a constant that depends on the degree of motion of P .*

3. AN SDG BASED ON POLYGONAL VD

Let $Q = Q_k$ be a regular k -gon for some even $k = 2s$, circumscribed by the unit disk, and let $\alpha = \pi/s$ (this is the angle at which

the center of Q sees an edge). Let $\text{VD}^\diamond(P)$ and $\text{DT}^\diamond(P)$ denote the Q -Voronoi diagram and the dual Q -Delaunay triangulation of P , respectively. In this section we show that the set of edges of $\text{VD}^\diamond(P)$ with sufficiently many *breakpoints* form a (β, β') -stable (Euclidean) Delaunay graph for appropriate multiples β, β' of α . Thus, by kinetically maintaining $\text{VD}^\diamond(P)$ (in its entirety), we shall get “for free” a KDS for keeping track of a stable portion of the Euclidean DT.

3.1 Properties of $\text{VD}^\diamond(P)$

We first review the properties of the (stationary) $\text{VD}^\diamond(P)$ and $\text{DT}^\diamond(P)$. Then we consider the kinetic version of these diagrams, as the points of P move, and review Chew’s proof [8] that the number of topological changes in these diagrams, over time, is only nearly quadratic in n . Finally, we present a straightforward kinetic data structure for maintaining $\text{DT}^\diamond(P)$ under motion that uses linear storage, and that processes a nearly quadratic number of events, each in $O(\log n)$ time. Although later on we will take Q to be a regular k -gon, the analysis in this subsection is more general, and we only assume here that Q is an arbitrary convex k -gon.

Stationary Q -diagrams. The bisector b_{pq}^\diamond between two points p and q , with respect to $d_Q(\cdot, \cdot)$, is the locus of all placements of the center of any homothetic copy Q' of Q that touches p and q . Q' can be classified according to the pair of its edges, e_1 and e_2 , that touch p and q , respectively. If we slide Q' so that its center moves along b_{pq}^\diamond (and its size expands or shrinks to keep it touching p and q), and the contact edges, e_1 and e_2 , remain fixed, the center traces a straight segment. The bisector is a concatenation of $O(k)$ such segments. They meet at *breakpoints*, which are placements of the center of a copy Q' that touches p and q and one of the contact points is a vertex of Q ; see Figure 2. We call such a placement a *corner contact* at the appropriate point. Note that a corner contact where some vertex w of (a copy Q' of) Q touches p has the property that the center of Q' lies on the fixed ray emanating from p and parallel to the directed segment from w to the center of Q .

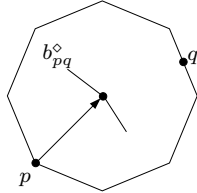


Figure 2. Each breakpoint on b_{pq}^\diamond corresponds to a corner contact of Q at one of the points p, q , so that ∂Q also touches the other point.

A well known property of Q -bisectors and Voronoi edges is that two bisectors $b_{pq_1}^\diamond, b_{pq_2}^\diamond$, can intersect at most once (again, assuming general position) [15], so every Q -Voronoi edge e_{pq}^\diamond is connected. Equivalently, this asserts that there exists at most one homothetic placement of Q at which it touches p, q_1 , and q_2 .

Another useful property of bisectors and Delaunay edges, in the special case where Q is a regular k -gon, which will be used in the next subsection, is that the breakpoints along a bisector b_{pq}^\diamond alternate between corner contacts at p and corner contacts at q . The proof is omitted here.

Consider next an edge pq of $\text{DT}^\diamond(P)$. Its dual Voronoi edge e_{pq}^\diamond is a portion of the bisector b_{pq}^\diamond , and consists of those center placements along b_{pq}^\diamond for which the corresponding copy Q' has an *empty interior* (i.e., its interior is disjoint from P). Following the notation of Chew [8], we call pq a *corner edge* if e_{pq}^\diamond contains a

breakpoint (i.e., a placement with a corner contact); otherwise it is a *non-corner edge*, and is therefore a straight segment.

Kinetic Q -diagrams. Consider next what happens to $\text{VD}^\diamond(P)$ and $\text{DT}^\diamond(P)$ as the points of P move continuously with time. In this case $\text{VD}^\diamond(P)$ changes continuously, but undergoes topological changes at certain critical times, called *events*. There are two kinds of events:

- (i) [FLIP EVENT.] A Voronoi edge e_{pq}^\diamond shrinks to a point, disappears, and is “flipped” into a newly emerging Voronoi edge $e_{p'q'}^\diamond$.
- (ii) [CORNER EVENT.] An endpoint of some Voronoi edge e_{pq}^\diamond becomes a breakpoint (a corner placement). Immediately after this time e_{pq}^\diamond either gains a new straight segment, or loses such a segment.

Some comments are in order:

- (a) A flip event occurs when the four points p, q, p', q' become “co-circular”: there is an empty homothetic copy Q' of Q that touches all four points.
- (b) Only non-corner edges can participate in a flip event, as both the vanishing edge e_{pq}^\diamond and the newly emerging edge $e_{p'q'}^\diamond$ do not have breakpoints near the event.
- (c) A flip event entails a discrete change in the Delaunay triangulation, whereas a corner event does not. Still we will keep track of both kinds of events.

We first bound the number of corner events.

LEMMA 3.1. *Let P be a set of n points in \mathbb{R}^2 under algebraic motion of bounded degree, and let Q be a convex k -gon. The number of corner events in $\text{DT}^\diamond(P)$ is $O(k^2 n \lambda_r(n))$, where r is a constant that depends on the degree of motion of P .*

PROOF. Fix a point p and a vertex w of Q , and consider all the corner events in which w touches p . As noted above, at any such event the center c of Q lies on a ray γ emanating from p at a fixed direction. For each other point $q \in P \setminus \{p\}$, let $\varphi_\gamma^\diamond[p, q]$ denote the distance, at time t , from p along γ to the center of a copy of Q that touches p (at w) and q . The value $\min_q \varphi_\gamma^\diamond[p, q](t)$ represents the intersection of $\partial \text{Vor}^\diamond(p)$ with γ at time t , where $\text{Vor}^\diamond(p)$ is the Voronoi cell of p in $\text{VD}^\diamond(P)$. The point q that attains the minimum defines the Voronoi edge e_{pq}^\diamond (or vertex if the minimum is attained by more than one point q) that intersects γ .

In other words, we have a collection of $n - 1$ partially defined functions $\varphi_\gamma^\diamond[p, q]$, and the breakpoints of their lower envelope represent the corner events that involve the contact of w with p . By our assumption on the motion of P , each function $\varphi_\gamma^\diamond[p, q]$ is piecewise algebraic, with $O(k)$ pieces. Each piece encodes a continuous contact of q with a specific edge of Q' , and has constant description complexity. Hence (see, e.g., [21, Corollary 1.6]) the complexity of the envelope is at most $O(k \lambda_r(n))$, for an appropriate constant r . Repeating the analysis for each point p and each vertex w of Q , the lemma follows. \square

Consider next flip events. As noted, each 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 $e_{p_1 p_3}^\diamond$, which is a non-corner edge before the event, shrinks to a point and is replaced by the non-corner edge $e_{p_2 p_4}^\diamond$. Let e_i denote the edge of Q' that touches p_i , for $i = 1, 2, 3, 4$.

We fix the quadruple of edges e_1, e_2, e_3, e_4 , bound 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 Q_0 of these edges, and note that any flip event involving these edges is also a flip event for Q_0 . We therefore restrict our attention to Q_0 , which is a convex k_0 -gon, for some $k_0 \leq 8$.

The number of flip events for a fixed Q_0 is bounded by extending and adapting the analysis of Chew [8]. Basically, we surround a non-corner edge by $O(1)$ nearby corner edges, and argue that, as long as none of them undergoes a corner event, only $O(1)$ flip events can occur in the “middle” edge. This allows us to charge a flip event to a corner event so that each corner event is charged only $O(1)$ times. See [8] and the full version [1] for details. Repeating this to each quadruple of edges of Q , we thus obtain:

THEOREM 3.2. *Let P be a set of n moving points in \mathbb{R}^2 under algebraic motion of bounded degree, and let Q be a convex k -gon. The number of topological changes in $\text{VD}^\circ(P)$ with respect to Q is $O(k^4 n \lambda_r(n))$, where r is a constant that depends on the degree of motion of P .*

Kinetic maintenance of $\text{VD}^\circ(P)$ and $\text{DT}^\circ(P)$. As already mentioned, it is a fairly trivial task to maintain $\text{DT}^\circ(P)$ and $\text{VD}^\circ(P)$ kinetically, as the points of P move. All we need to do is to assert the correctness of the present triangulation by a collection of local certificates, one for each edge of the diagram, where the certificate of an edge asserts that the two homothetic placements Q^-, Q^+ of Q that circumscribe the two respective adjacent Q -Delaunay triangles $\triangle pqr^-, \triangle pqr^+$, are such that Q^- does not contain r^+ and Q^+ does not contain r^- . The failure time of this certificate is the first time (if one exists) at which p, q, r^- , and r^+ become Q -cocircular—they all lie on the boundary of a common homothetic copy of Q . If pq is an edge of the periphery of $\text{DT}^\circ(P)$, so that $\triangle pqr^+$ exists but $\triangle pqr^-$ does not, then Q^- is a limiting wedge bounded by rays supporting two consecutive edges of (a copy of) Q , one passing through p and one through q (see Figure 3). The failure time of the corresponding certificate is the first time (if any) at which r^+ also lies on the boundary of that wedge. This corresponds to a flip event.

We maintain the breakpoints using “sub-certificates”, each of which asserts that Q^- , say, touches each of p, q, r^- at a respective specific edge (and similarly for Q^+). The failure time of this sub-certificate is the first failure time when one of p, q or r^- touches Q^- at a vertex. In this case we have a corner event—two of the adjacent Voronoi edges terminate at a corner placement. Note that the failure time of each sub-certificate can be computed in $O(1)$ time. Moreover, for a fixed collection of valid sub-certificates, the failure time of an initial certificate (asserting non-co-circularity) can also be computed in $O(1)$ time (provided that it fails before the failures of the corresponding sub-certificates), because we know the four edges of Q^- involved in the contacts.

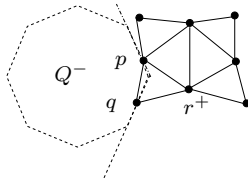


Figure 3. If r^- does not exist then Q^- is a limiting wedge bounded by rays supporting two consecutive edges of (a copy of) Q .

We therefore maintain an event queue that stores and updates all the active failure times (there are only $O(n)$ of them—the bound is independent of k). When a sub-certificate fails we do not change

$\text{DT}^\circ(P)$, but only update the corresponding Voronoi edge, by adding or removing a segment and a breakpoint, and by replacing the sub-certificate by a new one; we also update the real certificate associated with the edge, because one of the contact edges has changed. When a real certificate fails we update $\text{DT}^\circ(P)$ and construct $O(1)$ new sub-certificates and certificates. Altogether, each update of the diagram takes $O(\log n)$ time. We thus have

THEOREM 3.3. *Let P be a set of n moving points in \mathbb{R}^2 under algebraic motion of bounded degree, and let Q be a convex k -gon. $\text{DT}^\circ(P)$ and $\text{VD}^\circ(P)$ can be maintained using $O(n)$ storage and $O(\log n)$ update time, so that $O(k^4 n \lambda_r(n))$ events are processed, where r is a constant that depends on the degree of motion of P .*

3.2 Stable Delaunay edges in $\text{DT}^\circ(P)$

We now restrict Q to be a regular k -gon. Let v_0, \dots, v_{k-1} be the vertices of Q arranged in a clockwise direction, with v_0 the leftmost. We call a homothetic copy of Q whose vertex v_j touches a point p , a v_j -placement of Q at p . Let u_j be the direction of the vector that connects v_j with the center of Q , for each $0 \leq j < k$ (as in Section 2). See Figure 4 (left).

We follow the machinery in the proof of Lemma 3.1. That is, for any pair $p, q \in P$ let $\varphi_j^\circ[p, q]$ denote the distance from p to the point $u_j[p] \cap b_{pq}^\circ$; we put $\varphi_j^\circ[p, q] = \infty$ if $u_j[p]$ does not intersect b_{pq}° . If $\varphi_j^\circ[p, q] < \infty$ then the point $b_{pq}^\circ \cap u_j[p]$ is the center of the v_j -placement Q' of Q at p that also touches q . The value $\varphi_j^\circ[p, q]$ is equal to the circumradius of Q' . See Figure 4 (middle).

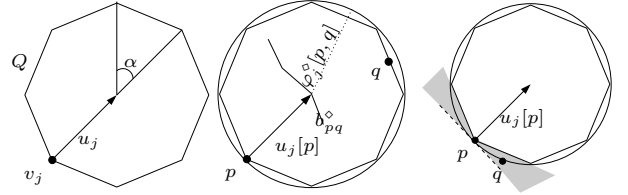


Figure 4. Left: u_j is the direction of the vector connecting vertex v_j to the center of Q . Middle: The function $\varphi_j^\circ[p, q]$ is equal to the radius of the circle that circumscribes the v_j -placement of Q at p that also touches q . Right: The case when $\varphi_j^\circ[p, q] = \infty$ while $\varphi_j[p, q] < \infty$. In this case q must lie in one of the shaded wedges.

The neighbor $N_j^\circ[p]$ of p in direction u_j is defined to be the point $q \in P \setminus \{p\}$ that minimizes $\varphi_j^\circ[p, q]$. Clearly, for any $p, q \in P$ we have $N_j^\circ[p] = q$ if and only if there is an empty v_j -placement Q' of Q at p so that q touches one of its edges.

Remark: Note that, in the Euclidean case, we have $\varphi_j[p, q] < \infty$ if and only if the angle between \overline{pq} and $u_j[p]$ is at most $\pi/2$. In contrast, $\varphi_j^\circ[p, q] < \infty$ if and only if the angle between \overline{pq} and $u_j[p]$ is at most $\pi/2 - \pi/k = \pi/2 - \alpha/2$. Moreover, we have $\varphi_j[p, q] \leq \varphi_j^\circ[p, q]$. Therefore, $\varphi_j^\circ[p, q] < \infty$ always implies $\varphi_j[p, q] < \infty$, but not vice versa; see Figure 4 (right).

LEMMA 3.4. *Let $p, q \in P$ be a pair of points such that $N_j(p) = q$ for $h \geq 3$ consecutive indices, say $0 \leq j < h$. Then for each of these indices, except possibly for the first and the last one, we also have $N_j^\circ[p] = q$.*

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 $\text{VD}(P)$. (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

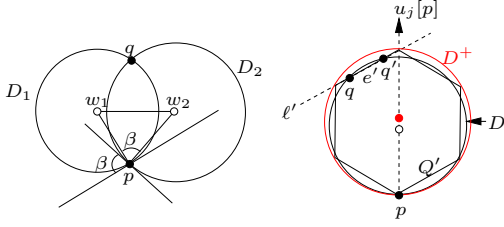


Figure 5. Left: The angle between the tangents to D_1 and D_2 is equal to $\angle w_1 p w_2 = \beta = (h-1)\alpha$. Right: The line ℓ' crosses D in a chord qq' which is fully contained in e' (right).

tangents to D_1 and D_2 at p or at q (these angles are equal) is $\beta = (h-1)\alpha$; see Figure 5 (left).

Fix an arbitrary index $1 \leq j \leq h-2$, so $u_j[p]$ intersects e_{pq} and forms an angle of at least α with each of pw_1, pw_2 . Let Q' be the v_j -placement of Q at p that touches q . The existence of such a placement follows from the preceding remark. We claim that $Q' \subset D_1 \cup D_2$. Establishing this property for every $1 \leq j \leq h-2$ will complete the proof of the lemma. Let e' be the edge of Q' passing through q . 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' . Since $q \in \partial D$ and is interior to D^+ , and since D and D^+ are centered on the same ray $u_j[q]$ and pass through p , it follows that $D \subset D^+$. See Figure 5 (right). The line ℓ' containing e' crosses D in a chord qq' that is fully contained in e' . The angle between the tangent to D at q and the chord qq' is equal to the angle at which p sees qq' . This is smaller than the angle at which p sees e' , which in turn is equal to $\alpha/2$.

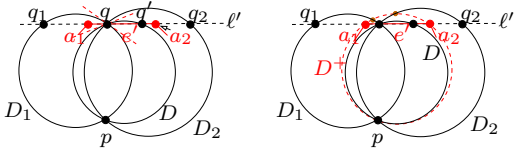


Figure 6. Left: The line ℓ' forms an angle of at least $\alpha/2$ with each of the tangents to D_1, D_2 at q . Right: The edge $e' = a_1 a_2$ of Q' is fully contained in $D_1 \cup D_2$.

Arguing as in the analysis of D_1 and D_2 , the tangent to D at q forms an angle of at least α with each of the tangents to D_1, D_2 at q , and hence e' forms an angle of at least $\alpha/2$ with each of these tangents; see Figure 6 (left). The line ℓ' marks two chords $q_1 q, q_2 q$ within the respective disks D_1, D_2 . We claim that e' is fully contained in their union $q_1 q_2$. Indeed, the angle $q_1 p q$ is equal to the angle between ℓ' and the tangent to D_1 at q , so $\angle q_1 p q \geq \alpha/2$. On the other hand, the angle at which p sees e' is $\alpha/2$, which is smaller. This, and the symmetric argument involving D_2 , are easily seen to imply the claim.

Now consider the circumscribing disk D^+ of Q' . Denote the endpoints of e' as a_1 and a_2 , where a_1 lies in $q_1 q$ and a_2 lies in $q_2 q$. Since the ray $\overline{p a_1}$ hits ∂D^+ before hitting D_1 , and the ray $\overline{p q}$ 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 $\overline{p a_1}, \overline{p q}$ and thus crosses e' . See Figure 6 (right). 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' (which lie on this arc) lie in $D_1 \cup D_2$. This, combined with the argument in the preceding paragraphs, is easily seen to imply that $Q' \subseteq D_1 \cup D_2$, so its interior does not contain points of P , which in turn implies that $N_j^\circ[p] = q$. As noted, this completes the proof of the lemma. \square

Since Q -Voronoi edges are connected, Lemma 3.4 implies that e_{pq}° is “long”, in the sense that it contains at least $h-2$ breakpoints that represent corner placements at p , interleaved (as promised in Section 3.1) with at least $h-3$ corner placements at q . This property is easily seen to hold also under the weaker assumptions that: (i) for the first and the last indices $j = 0, h-1$, the point $N_j[p]$ either is equal to q or is undefined, and (ii) for the rest of the indices j we have $N_j[p] = q$ and $\varphi_j^\circ[p, q] < \infty$ (i.e., the v_j -placement of Q at p that touches q exists). In this relaxed setting, it is now possible that any of the two points w_1, w_2 lies at infinity, in which case the corresponding disk D_1 or D_2 degenerates into a halfplane. This stronger version of Lemma 3.4 is used in the proof of the converse Lemma 3.5, asserting that every edge e_{pq}° in $\text{VD}^\circ(P)$ with sufficiently many breakpoints has a stable counterpart e_{pq} in $\text{VD}(P)$.

LEMMA 3.5. *Let $p, q \in P$ be a pair of points such that $N_j^\circ[p] = q$ for at least three consecutive indices $j \in \{0, \dots, k-1\}$. Then for each of these indices, except possibly for the first and the last one, we have $N_j[p] = q$.*

PROOF. Again, assume with no loss of generality that $N_j^\circ[p] = q$ for $0 \leq j \leq h-1$, with $h \geq 3$. Suppose to the contrary that, for some $1 \leq j \leq h-2$, we have $N_j[p] \neq q$. Since $N_{j-1}^\circ[p] = q$ by assumption, we have $\varphi_j[p, q] \leq \varphi_{j-1}^\circ[p, q] < \infty$, 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 , so that $\varphi_{j-1}[p, r] < \varphi_{j-1}[p, q] < \infty$, which follows because (i) $N_{j-1}^\circ[p] = q$ by assumption, so $\varphi_{j-1}^\circ[p, q] < \infty$, and (ii) $\varphi_{j-1}[p, q] \leq \varphi_{j-1}^\circ[p, q]$. See Figure 7 (left). Similarly, we get that either $\varphi_{j-2}[p, r] < \varphi_{j-2}[p, q] < \infty$ or $\varphi_{j-2}[p, r] \leq \varphi_{j-2}[p, q] = \infty$ (where the latter can occur only for $j = 1$). Now applying (the extended version of) Lemma 3.4 to the point set $\{p, q, r\}$ and to the index set $\{j-2, j-1, j\}$, we get that $\varphi_{j-1}^\circ[p, r] < \varphi_{j-1}^\circ[p, q]$. But this contradicts the fact that $N_{j-1}^\circ[p] = q$. \square

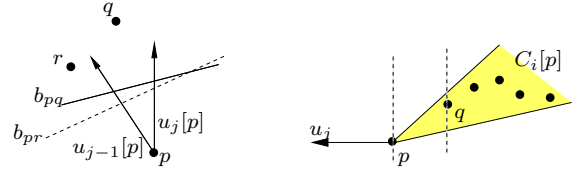


Figure 7. Left: Proof of Lemma 3.5. If $N_j[p] \neq q$ because some r , lying to the left of the line from p to r , satisfies $\varphi_j[p, r] < \varphi_j[p, q]$. Since $\varphi_{j-1}[p, q] < \varphi_{j-1}^\circ[p, q] < \infty$, we have $\varphi_{j-1}[p, r] < \varphi_{j-1}[p, q]$. Right: q is j -extremal for p .

Maintaining an SDG using $\text{VD}^\circ(P)$. Lemmas 3.4 and 3.5 together imply that SDG can be maintained using the fairly straightforward kinetic algorithm for maintaining the whole $\text{VD}^\circ(P)$, provided by Theorem 3.3. We use $\text{VD}^\circ(P)$ to maintain the graph G on P , whose edges are all the pairs $(p, q) \in P \times P$ such that p and q define an edge e_{pq}° in $\text{VD}^\circ(P)$ that contains at least seven breakpoints. As shown in Theorem 3.3, this can be done with $O(n)$ storage, $O(\log n)$ update time, and $O(k^4 n \lambda_r(n))$ updates (for an appropriate r). We claim that G is a $(6\alpha, \alpha)$ -SDG in the Euclidean norm.

Indeed, if two points $p, q \in P$ define a 6α -long edge e_{pq} in $\text{VD}(P)$ then this edge stabs at least six rays $u_j[p]$ emanating from p , and at least six rays $u_j[q]$ emanating from q . Thus, according to Lemma 3.4, $\text{VD}^\circ(P)$ contains the edge e_{pq}° with at least four breakpoints corresponding to corner placements of Q at p that touch q , and at least four breakpoints corresponding to corner placements of Q at q that touch p . Therefore, e_{pq}° contains at least 8 breakpoints, so $(p, q) \in G$.

For the second part, if $p, q \in P$ define an edge e_{pq}° in $\text{VD}^\circ(P)$ with at least 7 breakpoints then, by the interleaving property of breakpoints, we may assume, without loss of generality, that at least four of these breakpoints correspond to P -empty corner placements of Q at p that touch q . Thus, Lemma 3.5 implies that $\text{VD}(P)$ contains the edge e_{pq} , and that this edge is hit by at least two consecutive rays $u_j[p]$. But then, as observed in Lemma 2.2, the edge e_{pq} is α -long in $\text{VD}(P)$. We thus obtain the main result of this section.

THEOREM 3.6. *Let P be a set of n moving points in \mathbb{R}^2 under algebraic motion of bounded degree, and let $\alpha \geq 0$ be a parameter. A $(6\alpha, \alpha)$ -stable Delaunay graph of P can be maintained by a KDS of linear size that processes $O(n\lambda_r(n)/\alpha^4)$ events, where r is a constant that depends on the degree of motion of P , and that updates the SDG at each event in $O(\log n)$ time.*

4. A FASTER DATA STRUCTURE

The data structure of Theorem 3.6 requires $O(n)$ storage but the best bound we have on the number of events it may encounter is $O^*(n^2/\alpha^4)$, which is much larger than the number of events encountered by the data structure of Theorem 2.3. In this section we present an alternative data structure that requires $O^*(n/\alpha^2)$ space and $O^*(n^2/\alpha^2)$ overall processing time. The algorithm is *local*, and it processes each event in $O^*(1/\alpha)$ time.

Notation. We use the directions u_i and the associated quantities $N_i[p]$ and $\varphi_i[p, q]$ defined in Section 2. We assume that k is even, and write, as in Section 2, $k = 2s$. We denote by C_i the cone (or wedge) with apex at the origin that is bounded by u_i and u_{i+1} . Note that C_i and $C_{i\pm s}$ are antipodal. As before, for a vector u , we denote by $u[x]$ the ray emanating from x in direction u . Similarly, for a cone C we denote by $C[x]$ the translation of C that places its apex at x . Let $0 \leq \beta \leq \pi/2$ be an angle. For a direction $u \in \mathbb{S}^1$ and for two points $p, q \in P$, we say that the edge $e_{pq} \in \text{VD}(P)$ is β -long around the ray $u[q]$ if p is the Voronoi neighbor of q in all directions in the range $[u-\beta, u+\beta]$, i.e., for all $v \in [u-\beta, u+\beta]$, the ray $v[q]$ intersects e_{pq} . The β -cone around $u[q]$ is the cone whose apex is q and each of its bounding rays makes an angle of β with $u[q]$.

j -extremal points. Let $p, q \in P$, let i be the index such that $q \in C_i[p]$, and let u_j be a direction such that $\langle u_j, x \rangle \leq 0$ for all $x \in C_i$. We say that q is j -extremal for p if $q = \arg \max\{\langle p', u_j \rangle \mid p' \in C_i[p] \cap P \setminus \{p\}\}$. That is, q is the nearest point to p in this cone, in the $(-u_j)$ -direction. Clearly, a point p has at most s j -extremal points, one for every admissible cone $C_i[p]$, for any fixed j . See Figure 7 (right).

For $0 \leq i < k$, let C'_i denote the extended cone that is the union of the seven consecutive cones C_{i-3}, \dots, C_{i+3} . Let $p, q \in P$, let i be the index such that $q \in C_i[p]$, and let u_j be a direction such that $\langle u_j, x \rangle \leq 0$ for all $x \in C'_i$. We say that the point $q \in P$ is *strongly j -extremal* for p if $q = \arg \max\{\langle p', u_j \rangle \mid p' \in C'_i[p] \cap P \setminus \{p\}\}$.

We say that a pair $(p, q) \in P \times P$ is (strongly) (j, ℓ) -extremal, for some $0 \leq j, \ell \leq k-1$, if p is (strongly) ℓ -extremal for q and q is (strongly) j -extremal for p .

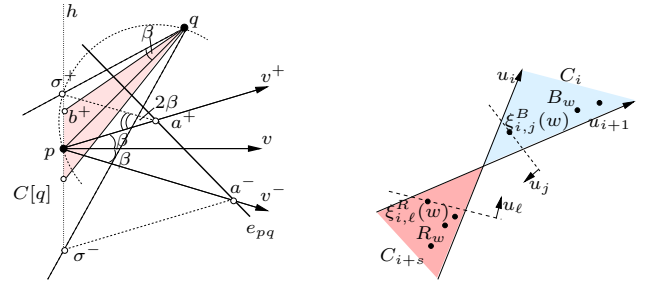


Figure 8. Left: Illustration of the setup in Lemma 4.1: the edge e_{pq} is β -long around $v[p]$, and the “tip” $\Delta\sigma^+q\sigma^-$ of the cone $C[q]$ is empty. Right: The points $\xi_{i,\ell}^R(w), \xi_{i,j}^B(w)$ for an appropriate node w .

LEMMA 4.1. *Let $p, q \in P$, and let v be a direction such that the edge e_{pq} appears in $\text{VD}(P)$ and is β -long around the ray $v[p]$. Let $C[q]$ be the β -cone around the ray from q through p . Then $\langle v, v \rangle \geq \langle p', v \rangle$ for all $p' \in P \cap C[q] \setminus \{q\}$.*

PROOF. Refer to Figure 8 (left). Without loss of generality, we assume that v is the $(+x)$ -direction and that q lies above the x -axis and to the right of p . (In this case the slope of the bisector b_{pq} is negative.) Let v^+ (resp., v^-) be the direction that makes a counterclockwise (resp., clockwise) angle of β with v . Let a^+ (resp., a^-) be the intersection of e_{pq} with $v^+[p]$ (resp., with $v^-[p]$); by assumption, both points exist. Let h be the vertical line passing through p . Let σ^+ (resp., σ^-) be the intersection point of h with the ray emanating from a^+ (resp., a^-) in the direction opposite to v^- (resp., v^+); see Figure 8 (left).

Note that $\angle pa^+\sigma^+ = 2\beta$, and that $\|a^+\sigma^+\| = \|pa^+\| = \|qa^+\|$, i.e., a^+ is the circumcenter of $\Delta p\sigma^+q$. Therefore $\angle\sigma^+qp = \frac{1}{2}\angle\sigma^+a^+p = \beta$. That is, σ^+ is the intersection of the upper ray of $C[q]$ with h . Similarly, σ^- is the intersection of the lower ray of $C[q]$ with h . Moreover, if there exists a point $x \in P$ properly inside the triangle $\Delta pq\sigma^+$ then $\|a^+x\| < \|a^+p\|$, contradicting the fact that a^+ is on e_{pq} . So the interior of $\Delta pq\sigma^+$ (including the edges pq, σ^+q) is disjoint from P . Similarly, by a symmetric argument, no points of P lie inside $\Delta pq\sigma^-$ or on its edges pq, σ^-q . Hence, the portion of $C[q]$ to the right of p is openly disjoint from P , and therefore p is a rightmost point of P (extreme in the v direction) inside $C[q]$. \square

COROLLARY 4.2. *Let $p, q \in P$. (i) If the edge e_{pq} is 3α -long in $\text{VD}(P)$ then there are $0 \leq j, \ell < k$ for which (p, q) is a (j, ℓ) -extremal pair. (ii) If the edge e_{pq} is 9α -long in $\text{VD}(P)$ then there are $0 \leq j, \ell < k$ for which (p, q) is a strongly (j, ℓ) -extremal pair.*

PROOF. To prove part (i), choose $0 \leq j, \ell < k$, such that e_{pq} is α -long around each of $u_\ell[p]$ and $u_j[q]$. By Lemma 4.1, p is u_ℓ -extremal in the α -cone $C[q]$ around the ray from q through p . Let i be the index such that $p \in C_i[q]$. Since the opening angle of $C[q]$ is 2α , it follows that $C_i[q] \subseteq C[q]$, so p is ℓ -extremal with respect to q , and, symmetrically, q is j -extremal with respect to p . To prove part (ii) choose $0 \leq j, \ell < k$, such that e_{pq} is 4α -long around each of $u_\ell[p]$ and $u_j[q]$ and apply Lemma 4.1 as in the proof of part (i). \square

The stable Delaunay graph. We kinetically maintain a $(9\alpha, \alpha)$ -stable Delaunay graph, whose precise definition is given below, using a data-structure which is based on a collection of 2-dimensional orthogonal range trees similar to the ones used in [2].

Fix $0 \leq i < s$, and choose a “sheared” coordinate frame in which the rays u_i and u_{i+1} form the x - and y -axes, respectively.

That is, in this coordinate frame, $q \in C_i[p]$ if and only if q lies in the upper-right quadrant anchored at p .

We define a 2-dimensional range tree \mathcal{T}_i consisting of a *primary* balanced binary search tree with the points of P stored at its leaves ordered by their x -coordinates, and of secondary trees, introduced below. Each internal node v of the primary tree of \mathcal{T}_i is associated with the *canonical subset* P_v of all points that are stored at the leaves of the subtree rooted at v . A point $p \in P_v$ is said to be *red* (resp., *blue*) in P_v if it is stored at the subtree rooted at the left (resp., right) child of v in \mathcal{T}_i . For each primary node v we maintain a secondary balanced binary search tree \mathcal{T}_i^v , whose leaves store the points of P_v ordered by their y -coordinates. We refer to a node w in a secondary tree \mathcal{T}_i^v as a *secondary node* w of \mathcal{T}_i .

Each node w of a secondary tree \mathcal{T}_i^v is associated with a canonical subset $P_w \subseteq P_v$ of points stored at the leaves of the subtree of \mathcal{T}_i^v rooted at w . We also associate with w the sets $R_w \subset P_w$ and $B_w \subset P_w$ of points residing in the *left* (resp., *right*) subtree of w and are red (resp., blue) in P_v . It is easy to verify that the sum of the sizes of the sets R_w and B_w over all secondary nodes of \mathcal{T}_i is $O(n \log^2 n)$.

For each secondary node $w \in \mathcal{T}_i$ and each $0 \leq j < k$ we maintain the points

$$\xi_{i,j}^R(w) = \arg \max_{p \in R_w} \langle p, u_j \rangle, \quad \xi_{i,j}^B(w) = \arg \max_{p \in B_w} \langle p, u_j \rangle,$$

provided that both R_w, B_w are not empty. See Figure 8 (right). It is straightforward to show that if (p, q) is a (j, ℓ) -extremal pair, so that $q \in C_i[p]$, then there is a secondary node $w \in \mathcal{T}_i$ for which $q = \xi_{i,j}^R(w)$ and $p = \xi_{i,\ell}^B(w)$.

For each $p \in P$ we define a set $\mathcal{N}[p]$ containing all points $q \in P$ for which (p, q) is a (j, ℓ) -extremal pair, for some pair of indices $0 \leq j, \ell < k$. Specifically, for each $0 \leq i < s$, and each secondary node $w \in \mathcal{T}_i$ such that $p = \xi_{i,\ell}^R(w)$ for some $0 \leq \ell < k$, we include in $\mathcal{N}[p]$ all the points q such that $q = \xi_{i,j}^B(w)$ for some $0 \leq j < k$. Similarly, for each $0 \leq i < s$, and each secondary node $w \in \mathcal{T}_i$ such that $p = \xi_{i,\ell}^B(w)$ for some $0 \leq \ell < k$ we include in $\mathcal{N}[p]$ all the points q such that $q = \xi_{i,j}^R(w)$ for some $0 \leq j < k$.

For each $0 \leq i < s$, any point $p \in P$ belongs to $O(\log^2 n)$ sets R_w and B_w , so the size of $\mathcal{N}[p]$ is bounded by $O(s^2 \log^2 n)$. Indeed, p may be coupled with up to $k = 2s$ neighbors at each of the $O(s \log^2 n)$ nodes containing it.

For each point $p \in P$ and $0 \leq \ell < k$ we maintain all points in $\mathcal{N}[p]$ in a kinetic and dynamic tournament $\mathcal{D}_\ell[p]$ whose winner q minimizes the directional distance $\varphi_\ell[p, q]$, as given in (1). That is, the winner in $\mathcal{D}_\ell[p]$ is $N_\ell[p]$ in the Voronoi diagram of $\{p\} \cup \mathcal{N}[p]$.

We are now ready to define the stable Delaunay graph G that we maintain. For each pair of points $p, q \in P$ we add the edge (p, q) to G if the following hold.

- (G1) There is an index $0 \leq \ell \leq k - 1$ such that q wins the 8 consecutive tournaments $\mathcal{D}_\ell[p], \dots, \mathcal{D}_{\ell+7}[p]$.
- (G2) The point p is strongly $(\ell+3)$ -extremal and strongly $(\ell+4)$ -extremal for q .

In the full version of the paper [1] we prove that G is $(9\alpha, \alpha)$ -stable. We also show how to kinetically maintain (a refined version of) G efficiently, and establish the following theorem.

THEOREM 4.3. *Let P be a set of n moving points in \mathbb{R}^2 under algebraic motion of bounded degree, and let $\alpha > 0$ be a parameter. A $(9\alpha, \alpha)$ -SDG of P can be maintained using a data structure that requires $O((n/\alpha^2) \log n)$ space and encounters two types of events: swap events and tournament events. There are $O(n^2/\alpha)$*

swap events, each processed in $O(\log^4(n)/\alpha)$ time. There are

$$O((n/\alpha)^2 \beta_{r+2} (\log(n)/\alpha) \log^2 n \log(\log(n)/\alpha))$$

tournament events, which are handled in a total of

$$O((n/\alpha)^2 \beta_{r+2} (\log(n)/\alpha) \log^2 n \log^2(\log(n)/\alpha))$$

processing time. The worst-case processing time of a tournament event is $O(\log^2(\log(n)/\alpha))$.

Remark: Comparing this algorithm with the space-inefficient one of Section 2, we note that they both use the same kind of tournaments, but here much fewer pairs of points ($O^*(n/\alpha^2)$ instead of $O(n^2/\alpha)$) participate in the tournaments. The price we have to pay is that the test for an edge pq to be stable is more involved. Moreover, keeping track of the subset of pairs that participate in the tournaments requires additional work, which is facilitated by the range trees \mathcal{T}_i .

5. PROPERTIES OF SDG

We conclude the paper by establishing some of the properties of stable Delaunay graphs.

Near co-circularities do not show up in an SDG. Consider a critical event during the kinetic maintenance of the full Delaunay triangulation, in which four points a, b, c, d become co-circular, in this order along their circumcircle, with this circle being empty. Just before the critical event, the Delaunay triangulation involved two triangles, say, abc, acd . The Voronoi edge e_{ac} shrinks to a point (namely, to the circumcenter of $abcd$ at the critical event), and, after the critical co-circularity, 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 e_{ac} becomes α -short (or even before this happens). It will then remove this edge from the stable subgraph, so the actual co-circularity will not be recorded. The new edge e_{bd} will then be detected by the algorithm only when it becomes sufficiently long (if at all), and will then enter the stable Delaunay graph. In short, critical co-circularities do not arise *at all* in our scheme.

As noted in the introduction, a Delaunay edge ab is just about to become α -short or α -long when the sum of the opposite angles in its two adjacent Delaunay triangles is $\pi - \alpha$ (see Figure 1 (left)). This shows that changes in the stable Delaunay graph occur when the ‘‘co-circularity defect’’ of a nearly co-circular quadruple (i.e., the difference between π and the sum of opposite angles in the quadrilateral spanned by the quadruple) is between α and $c\alpha$, where c is the constant used in our definitions in Section 3 or Section 4. Note that a degenerate case of co-circularity 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 Voronoi edge becomes α -short, in which case the angle $\angle acb$, where c is the middle point the (near-)collinearity becomes $\pi - \alpha$ (see Figure 9 (left)). Therefore a hull edge is removed if the Delaunay triangle is almost collinear. The edge re-appears when its corresponding Voronoi edge is long enough, as before.

SDGs are not too sparse. Consider the Voronoi cell $\text{Vor}(p)$ of a point p , and suppose that p has only one α -long edge e_{pq} . Since the angle at which p sees e_{pq} is at most π , the sum of the angles at which p sees the other edges is at least π , so $\text{Vor}(p)$ has at least π/α α -short edges. Let m_1 denote the number of points p with this property. Then the sum of their degrees in $\text{DT}(P)$ is at least $m_1(\pi/\alpha + 1)$. Similarly, if m_0 points do not have any α -long

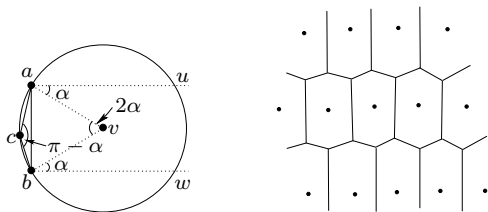


Figure 9. Left: The near collinearity that corresponds to a Voronoi edge becoming α -short. Right: If the points of P lie on a sufficiently spaced shifted grid then the number of α -long edges in $\text{VD}(P)$ (the vertical ones) is close to n .

Voronoi edge, then the sum of their degrees is at least $2\pi m_0/\alpha$. Any other point has degree at least 3 in $\text{DT}(P)$ and has at least two α -long Voronoi edges. So the number of α -long edges is at least (recall that each α -long edge is counted twice)

$$n - m_1 - m_0 + m_1/2 = n - (m_1 + 2m_0)/2. \quad (2)$$

Let h denote the number of hull vertices. Since the sum of the degrees is $6n - 2h - 6$, we get

$$3(n - h - m_1 - m_0) + 2h + m_1 \left(\frac{\pi}{\alpha} + 1 \right) + 2m_0 \frac{\pi}{\alpha} \leq 6n - 2h - 6,$$

implying that

$$m_1 + 2m_0 \leq \frac{3n}{\pi/\alpha - 2}.$$

Plugging this inequality in (2), we conclude that the number of α -long edges is at least

$$n \left[1 - \frac{3}{2(\pi/\alpha - 2)} \right].$$

As α decreases, the number of edges in the SDG is always at least a quantity that gets closer to n . This is nearly tight, since there exist n -point sets for which the number of stable edges is only roughly n , see Figure 9 (right).

Closest pairs, crusts, β -skeleta, and the SDG. Let $\beta \geq 1$, and P 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., [5, 17] for properties of the β -skeleton, and for its applications in surface reconstruction. We show that the edges of the β -skeleton are α -stable in $\text{DT}(P)$, provided $\beta \geq 1 + \Omega(\alpha^2)$. 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 [4, 5] for the definition of crusts and their applications in surface reconstruction). See the full version of the paper [1] for the proofs of these claims.

In contrast, 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 is given in Figure 10 (left). As a matter of fact, the stable Delaunay graph need not even be connected, as is illustrated in Figure 10 (right).

Completing SDG into a triangulation. As argued above, the Delaunay edges that are missing in the stable subgraph correspond to nearly co-circular quadruples of points, or to nearly collinear triples of points near the boundary of the convex hull. Arguably, these

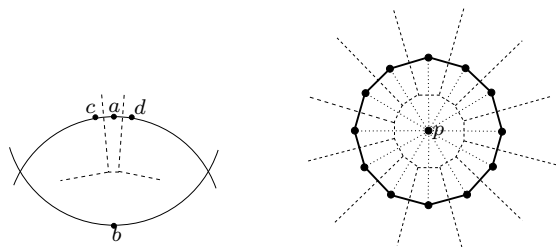


Figure 10. Two constructions with an appropriate choice of α . Left: ab is an edge of the relative neighborhood graph but not of SDG. Right: A wheel-like configuration that disconnects p in the stable Delaunay graph. The Voronoi diagram is drawn with dashed-dotted lines, the stable Delaunay edges are drawn as solid, and the remaining Delaunay edges as dotted edges. The points of the “wheel” need not be co-circular.

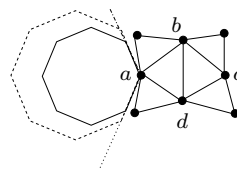


Figure 11. The triangulation $\text{DT}^\circ(P)$ of an 8-point set P . The points a, b, c, d , which do not lie on the convex hull of P , still lie on the boundary of the union of the triangles of $\text{DT}^\circ(P)$ because, for each of these points we can place an arbitrary large homothetic interior-empty copy of Q which touches that point.

missing edges carry little information, because they may “flicker” in and out of the Delaunay triangulation even when the points move just slightly (so that all angles determined by the triples of points change only slightly). Nevertheless, in many applications it is desirable (or essential) to complete the stable subgraph into *some* triangulation, preferably one that is also stable in the combinatorial sense—it undergoes only nearly quadratically many topological changes.

By the analysis in Section 3 we can achieve part of this goal by maintaining the full Delaunay triangulation $\text{DT}^\circ(P)$ under the polygonal norm induced by the regular k -gon Q_k . This diagram experiences only a nearly quadratic number of topological changes, is easy to maintain, and contains all the stable Euclidean Delaunay edges, for an appropriate choice of $k \approx 1/\alpha$. Moreover, the union of its triangles is simply connected — it has no holes. Unfortunately, in general it is not a triangulation of the entire convex hull of P , as illustrated in Figure 11.

For the time being, we leave it as an open problem to come up with a simple and “stable” scheme for filling the gaps between the triangles of $\text{DT}^\circ(P)$ and the edges of the convex hull. It might be possible to extend the kinetic triangulation scheme developed in [16] to kinetically maintain a triangulation of “fringes” between $\text{DT}^\circ(P)$ and the convex hull of P .

Of course, if we only want to maintain a triangulation of P that experiences only a nearly quadratically many topological changes, then we can use the scheme in [16], or the earlier, somewhat more involved scheme in [3]. However, if we want to keep the triangulation “as Delaunay as possible”, we should include in it the stable portion of DT , and then the efficient completion of it, as mentioned above, becomes an issue.

Nearly Euclidean norms and some of their properties. One way

of interpreting the results of Section 3 is that the stability of Delaunay edges is preserved, in an appropriately defined sense, if we replace the Euclidean norm by the polygonal norm induced by the regular k -gon Q_k (for $k \approx 1/\alpha$). That is, stable edges in one Delaunay triangulation are also edges of the other triangulation, and are stable there too. Here we note that there is nothing special about Q_k : The same property holds if we replace the Euclidean norm by any sufficiently close norm (or convex distance function [9]).

Specifically, let Q be a closed convex set in the plane that is contained in the unit disk D_0 and contains the disk $D'_0 = (\cos \alpha)D_0$ that is concentric with D_0 and scaled by the factor $\cos \alpha$. This is equivalent to requiring that the Hausdorff distance $H(Q, D_0)$ between Q and D_0 be at most $1 - \cos \alpha$. We define the center of Q to coincide with the common center of D_0 and D'_0 .

Q induces a convex distance function d_Q , defined by $d_Q(x, y) = \min\{\lambda \mid y \in x + \lambda Q\}$. Consider the Voronoi diagram $\text{Vor}^Q(P)$ of P induced by d_Q , and the corresponding Delaunay triangulation $\text{DT}^Q(P)$. We omit here the detailed analysis of the structure of these diagrams, which is similar to that for the norm induced by Q_k , as presented in Section 3. See also [8, 9] for more details. Call an edge e_{pq} of $\text{Vor}^Q(P)$ α -stable if the following property holds: Let u and v be the endpoints of e_{pq} , and let Q_u, Q_v be the two homothetic copies of Q that are centered at u, v , respectively, and touch p and q . Then we require that the angle between the supporting lines at q (for simplicity, assume that Q is smooth; otherwise, the condition should hold for any pair of supporting lines at p or at q) to Q_u and Q_v at p is at least α , and that the same holds at q . In this case we refer to the edge pq of $\text{DT}^Q(P)$ as α -stable.

Note that Q_k -stability was (implicitly) defined in a different manner in Section 3, based on the number of breakpoints of the corresponding Voronoi edges. Nevertheless, it is easy to verify that the two definitions are essentially identical. In the full version of the paper [1] we show:

THEOREM 5.1. *Let P, α , and Q be as above. Then (i) every 4α -stable edge of the Euclidean Delaunay triangulation is an α -stable edge of $\text{DT}^Q(P)$. (ii) Conversely, every 4α -stable edge of $\text{DT}^Q(P)$ is also an α -stable edge in the Euclidean norm.*

The proof follows the analysis in Section 3, and adapts it to handle the more general situation that arises here, requiring the derivation of several additional geometric properties of Q -stable edges.

There are many interesting open problems that arise here. One of the main problems is to bound the number of topological changes in $\text{DT}^Q(P)$ under algebraic motion of bounded degree of the points of P , and to extend the class of sets Q for which a near quadratic bound on the number of these changes can be established.

References

- [1] P. K. Agarwal, J. Gao, L. Guibas, H. Kaplan, V. Koltun, N. Rubin and M. Sharir, Kinetic stable Delaunay graphs, <http://www.cs.tau.ac.il/~rubinnat/StableF.pdf>.
- [2] P. K. Agarwal, H. Kaplan and M. Sharir, Kinetic and dynamic data structures for closest pair and all nearest neighbors, *ACM Trans. Algorithms* 5 (1) (2008), Art. 4.
- [3] P. K. Agarwal, Y. Wang and H. Yu, A 2D kinetic triangulation with near-quadratic topological changes, *Discrete Comput. Geom.* 36 (2006), 573–592.
- [4] N. Amenta and M. Bern, Surface reconstruction by Voronoi filtering, *Discrete Comput. Geom.*, 22 (1999), 481–504.
- [5] 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.
- [6] F. Aurenhammer and R. Klein, Voronoi diagrams, in *Handbook of Computational Geometry*, J.-R. Sack and J. Urrutia, Eds., Elsevier, Amsterdam, 2000, pages 201–290.
- [7] J. Basch, L. J. Guibas and J. Hershberger, Data structures for mobile data, *J. Algorithms* 31 (1) (1999), 1–28.
- [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. First Annu. ACM Sympos. Comput. Geom.*, 1985, pp. 235–244.
- [10] B. Delaunay, Sur la sphère vide. A la memoire de Georges Voronoi, *Izv. Akad. Nauk SSSR, Otdelenie Matematicheskikh i Estestvennyh Nauk* 7 (1934), 793–800.
- [11] H. Edelsbrunner, *Geometry and Topology for Mesh Generation*, Cambridge University Press, Cambridge, 2001.
- [12] L. J. Guibas, Modeling motion, In J. E. Goodman and J. O'Rourke, editors, *Handbook of Discrete and Computational Geometry*. CRC Press, Inc., Boca Raton, FL, USA, second edition, 2004, pages 1117–1134.
- [13] L. J. Guibas, Kinetic data structures — a state of the art report, In P. K. Agarwal, L. E. Kavradi and M. Mason, editors, *Proc. Workshop Algorithmic Found. Robot.*, pages 191–209. A. K. Peters, Wellesley, MA, 1998.
- [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.*, pages 113–125. Springer-Verlag, 1992.
- [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] H. Kaplan, N. Rubin and M. Sharir, A kinetic triangulation scheme for moving points in the plane, this proceedings.
- [17] D. Kirkpatrick and J. D. Radke, A framework for computational morphology, *Computational Geometry* (G. Toussaint, ed.), North-Holland (1985), 217–248.
- [18] 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.
- [19] K. Mehlhorn, *Data Structures and Algorithms 1: Sorting and Searching*, Springer Verlag, Berlin 1984.
- [20] J. Nievergelt and E. M. Reingold, Binary search trees of bounded balance, *SIAM J. Comput.* 2 (1973), 33–43.
- [21] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

Acknowledgements. Pankaj Agarwal was supported by NSF under grants CNS-05-40347, CCF-06 -35000, CCF-09-40671 and DEB-04-25465, by ARO grants W911NF-07-1-0376 and W911NF-08-1-0452, by an NIH grant IP50-GM-08183-01, by a DOE grant OEG-P200A070505, and by a grant from the U.S.–Israel Binational Science Foundation. Leo Guibas was supported by NSF grants CCR-0204486, ITR-0086013, ITR-0205671, ARO grant DAAD19-03-1-0331, as well as by the Bio-X consortium at Stanford. Haim Kaplan and Natan Rubin were partially supported by Grant 2006/204 from the U.S. - Israel Binational Science Foundation, project number 2006204, and by the Israel Science Foundation grant no. 975-06. Micha Sharir and Natan Rubin were supported by NSF Grants CCF-05-14079 and CCF-08-30272, by Grant 338/09 from Israel Science Fund, and by the Hermann Minkowski–MINERVA Center for Geometry at Tel Aviv University.