Discrete Comput Geom 36:573–592 (2006)
DOI: 10.1007/s00454-006-1266-7

Discrete & Computational
Geometry

A Two-Dimensional Kinetic Triangulation with Near-Quadratic Topological Changes*

Pankaj K. Agarwal,[1] Yusu Wang,[2] and Hai Yu[1]

[1]Department of Computer Science, Duke University,
Box 90129, Durham, NC 27708, USA
{pankaj,fishhai}@cs.duke.edu

[2]Department of Computer Science and Engineering, Ohio State University,
2015 Neil Avenue, Columbus, OH 43210, USA
yusu@cse.ohio-state.edu

Abstract. A triangulation of a set S of points in the plane is a subdivision of the convex hull of S into triangles whose vertices are points of S. Given a set S of n points in $\mathbb{R}^2$, each moving independently, we wish to maintain a triangulation of S. The triangulation needs to be updated periodically as the points in S move, so the goal is to maintain a triangulation with a small number of topological events, each being the insertion or deletion of an edge. We propose a kinetic data structure (KDS) that processes $n^2 2^{O(\sqrt{\log n \cdot \log\log n})}$ topological events with high probability if the trajectories of input points are algebraic curves of fixed degree. Each topological event can be processed in $O(\log n)$ time. This is the first known KDS for maintaining a triangulation that processes a near-quadratic number of topological events, and almost matches the $\Omega(n^2)$ lower bound [1]. The number of topological events can be reduced to $nk \cdot 2^{O(\sqrt{\log k \cdot \log\log n})}$ if only k of the points are moving.

1. Introduction

A *triangulation* of a set S of points in $\mathbb{R}^2$ is a subdivision of the convex hull of S into triangles whose vertices are points of S. Motivated by applications in several areas, including computer graphics, physical simulation, collision detection, and geographic information systems, triangulations have been widely studied [9], [12]. With the advancement

* This research was supported by NSF under Grants CCR-00-86013, EIA-98-70724, EIA-01-31905, CCR-02-04118, and DEB-04-25465, by ARO Grants W911NF-04-1-0278 and DAAD19-03-1-0352, by a grant from the U.S.–Israel Binational Science Foundation. This work was done while the second author was a graduate student at Duke University.

in technology, many applications, for instance, video games, virtual reality, dynamic simulations, and robotics, call for maintaining a triangulation as the points move. For example, the arbitrary Eulerian–Lagrangian method [11] provides a way to integrate the motion of fluids and solids within a moving finite-element mesh. The time axis is discretized and the mesh vertices are moved between each time step so as to respect the interfaces between the different media. However, numerical problems arise when the mesh becomes too distorted, and the mesh generated depends on the discretization of time. Another approach to build more adaptive triangulations for such problems is to work in the space–time domain [13].

Given a set S of points in $\mathbb{R}^2$, each moving independently, we wish to maintain a triangulation of S. As the points in S move, any fixed triangulation of S also deforms continuously. However, a triangulation computed for the initial configuration cannot be guaranteed to remain valid all the time. Therefore, it becomes necessary to *update* a triangulation over time by deleting some of the existing edges and inserting some other edges. We refer to each such insertion or deletion in the triangulation as a *topological event*. The topological events occur only at discrete time instances. In this paper we study how to maintain a triangulation so that the number of topological events is near-quadratic.

Related Work. Since Atallah's seminal paper [7] on kinetic geometry, much work has been devoted to this area due to its importance in both theory and applications of computational geometry; see [3], [5], and [14] for reviews on kinetic geometric algorithms and data structures. The early work on kinetic geometry mostly focused on bounding the number of combinatorial or topological changes in various geometric structures as the input objects move. Basch et al. [8] introduced a general framework, the so-called *kinetic data structure* (KDS), for maintaining a discrete attribute of objects in predictable motion. Their approach to maintain a given attribute $A(t)$ for a continuously changing scene $S(t)$ is as follows: at a given time t, we create a proof of correctness of the attribute based on elementary tests called *certificates*. For each certificate, we compute the time at which it fails and put it in a global event queue. As the attribute cannot change while all tests remain valid, it is unnecessary to perform any computation until the first certificate fails. When a certificate fails, the discrete attribute is updated if it needs to be, and a new proof of correctness is constructed by making certain modifications to the previous proof of correctness. Their approach led to efficient algorithms for several kinetic problems [14].

In the context of triangulation, a longstanding open problem is to bound the number of topological events in the Delaunay triangulation of a set of moving points in $\mathbb{R}^2$. The best known upper bound is near-cubic if trajectories of input points are algebraic curves of fixed degree; the bound is cubic if each point moves with unit speed in a fixed direction [6]. Although it is conjectured that the number of topological changes is $O(n^2)$, no such bound is known even for maintaining an arbitrary triangulation of a set of moving points in $\mathbb{R}^2$. Agarwal et al. [2] described a scheme for maintaining a triangulation of a set of points that incurs roughly $n^{7/3}$ topological changes if the points are moving linearly. Chew [10] proved that the Delaunay triangulation of S under L_1-metric changes $O(n^2\alpha(n))$ times, where $\alpha(n)$ is the inverse Ackermann function; however, the Delaunay triangulation under L_1-metric is not necessarily a triangulation of the convex hull of the point set.

Agarwal et al. [1] showed that the convex hull of a moving point set may change $\Theta(n^2)$ times if the points are moving linearly; this result immediately implies a lower bound $\Omega(n^2)$ on the number of topological changes to any triangulation. This lower bound on the number of topological events holds even if a linear number of Steiner points are allowed [4].

Our Results. Let $S = \{p_1, \ldots, p_n\}$ 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, and let $S(t)$ denote the configuration of S at time t. We assume that $x_i(\cdot)$, $y_i(\cdot)$ are polynomials of fixed degree. We describe a KDS for maintaining a triangulation of $S(t)$ that processes $n^2 2^{O(\sqrt{\log n \cdot \log\log n})}$ topological events with high probability, each of which is insertion or deletion of an edge. Each topological event can be processed in $O(\log n)$ time. This is the first KDS for maintaining a triangulation that processes near-quadratic topological events. The number of events can be reduced to $nk \cdot 2^{O(\sqrt{\log k \cdot \log\log n})}$ if only k of the points are moving.

Our algorithm relies on a randomized hierarchical scheme. We first describe the so-called fan triangulation (Section 2), and then introduce the notion of constrained fan triangulation with respect to a planar subdivision (Section 3). We choose a random sample $R \subseteq S$, compute a triangulation of R recursively, and then compute the constrained fan triangulation of S with respect to the triangulation computed for R (Section 4). We analyze the events in constrained fan triangulation and show that if R is a random subset of appropriate size, the total number of events is near-quadratic (Section 5).

2. Fan Triangulation

Let $S = \{p_1, \ldots, p_n\}$ be a set of n (stationary) points in $\mathbb{R}^2$, sorted in non-increasing order of their y-coordinates, i.e., $y(p_1) \geq y(p_2) \geq \cdots \geq y(p_n)$. For a point $q \in \mathbb{R}^2$, let $S_q = \{p_i \in S \mid y(p_i) > y(q)\}$. Denote by $\mathcal{V}(q) \subseteq S$ the set of points on $\partial\,\mathrm{conv}(S_q)$ that are visible from q, i.e., $p_i \in \mathcal{V}(q)$ if the relative interior of the segment qp_i does not intersect $\mathrm{conv}(S_q)$. Furthermore, let $\rho(q)$ denote the point from $\mathcal{V}(q)$ such that the oriented line $\overrightarrow{q\rho(q)}$ is the left tangent of $\mathrm{conv}(S_q)$, and let $\gamma(q)$ denote the point from $\mathcal{V}(q)$ such that the oriented line $\overrightarrow{q\gamma(q)}$ is the right tangent of $\mathrm{conv}(S_q)$. Obviously, $\mathcal{V}(q)$ is the subset of vertices on $\partial\,\mathrm{conv}(S_q)$ lying between $\rho(q)$ and $\gamma(q)$, assuming that the vertices are ordered in counterclockwise direction.

The *fan triangulation* of S is constructed by sweeping a horizontal line h from $y = +\infty$ to $y = -\infty$. At any time the algorithm maintains the fan triangulation of points from S that lie above h. It updates the triangulation when the sweep line crosses a point $p_i \in S$ by adding the edges $p_i p_j$ for all $p_j \in \mathcal{V}(p_i)$; see Fig. 1. The triangulation at the end of the sweep is the fan triangulation of S, which we denote as $\mathbb{F}(S)$.

We classify the edges of $\mathbb{F}(S)$ incident upon a point $p_i \in S$ into two classes:

(i) *Up edges*: edges $p_i p_j$ so that $j < i$; p_j is also referred to as an *up neighbor* of p_i.
(ii) *Down edges*: edges $p_i p_j$ so that $j > i$; p_j is also called a *down neighbor* of p_i. Furthermore, if $p_i = \rho(p_j)$ or $p_i = \gamma(p_j)$, then edge $p_i p_j$ is referred to as a *convex edge*; otherwise, it is a *reflex edge*.

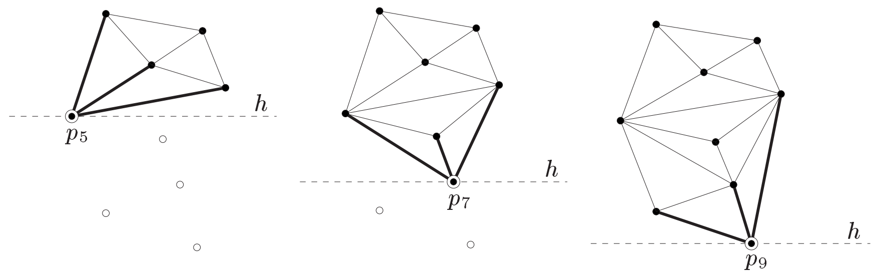

Fig. 1. Construction of fan triangulation at various stages—the point denoted by double circle is being inserted, and the thick edges are added. The points are ordered from top to bottom by their indices.

The following properties of $\mathbb{F}(S)$ are straightforward to prove:

(F1) For $1 \le i < n$, $p_i p_{i+1}$ is an edge of $\mathbb{F}(S)$.

(F2) For each p_i, at most one of its down edges is a reflex edge. Indeed if $p_i p_j$ is a reflex edge of p_i, then p_i lies in the interior of $\mathrm{conv}(S_{p_{j+1}})$, and therefore there is no down edge $p_i p_k$, for $k > j$. Let the edges $p_i p_{j_1}, p_i p_{j_2}, \ldots, p_i p_{j_k}$ be the sequence of the down edges incident upon p_i sorted in the clockwise direction. Then either $j_1 < \cdots < j_k$ and $p_i = \rho(p_{j_l})$ for $l < k$ ($l \le k$ if $p_i p_{j_k}$ is not a reflex edge, e.g., p_5 in Fig. 1(c)), or $j_1 > \cdots > j_k$ and $p_i = \gamma(p_{j_l})$ for $l < k$ ($l \le k$ if $p_i p_{j_k}$ is not a reflex edge, e.g., p_4 in Fig. 1(c)).

Now suppose we are given a set of moving points $S(t) = \{p_1(t), p_2(t), \ldots, p_n(t)\}$ in the plane, and we wish to maintain $\mathbb{F}(S(t))$ for any $t \in \mathbb{R}$. As the points from S move, $\mathbb{F}(S(t))$ deforms continuously. However, the topology of $\mathbb{F}(S(t))$ changes only at discrete times, which we refer to as *events*. Note the difference between events and topological events: an event may cause multiple topological events to the triangulation. As observed in [2], there are two types of events.

Ordering Event. The y-coordinates of two points p_i and p_j become equal at time t_0. Assume that $y(p_i(t_0^-)) > y(p_j(t_0^-))$ and $y(p_i(t_0^+)) < y(p_j(t_0^+))$, where t_0^- is the time immediately before t_0 and t_0^+ is the time immediately after t_0. By property (F1), $p_i(t_0^-)p_j(t_0^-)$ and $p_j(t_0^+)p_i(t_0^+)$ are present in $\mathbb{F}(S(t_0^-))$ and $\mathbb{F}(S(t_0^+))$, respectively. In fact, both of them are necessarily convex edges; assume $p_i(t_0^-) = \rho(p_j(t_0^-))$ and $p_j(t_0^+) = \gamma(p_i(t_0^+))$. Let $p_{k_1}, p_{k_2}, \ldots, p_{k_u}$ be the sequence of vertices in $\mathcal{V}_{p_i p_j}(t_0) = \mathcal{V}(p_i(t_0^-)) \cap \mathcal{V}(p_j(t_0^+))$ ordered in the counterclockwise direction along $\partial\,\mathrm{conv}(S_{p_i}(t_0))$. Obviously, at time t_0^-, $p_{k_l} p_i$, for all $1 \le l \le u$, and $p_{k_u} p_j$ are edges in $\mathbb{F}(S(t_0^-))$, while at time t_0^+, $p_{k_l} p_j$, for all $1 \le l \le u$, and $p_{k_1} p_i$ are edges in $\mathbb{F}(S(t_0^+))$. So to update the triangulation, we delete the edges $p_{k_2} p_i, \ldots, p_{k_u} p_i$ and insert the edges $p_{k_1} p_j, \ldots, p_{k_{u-1}} p_j$ to obtain $\mathbb{F}(S(t_0^+))$. See Fig. 2(a). Hence, an ordering event induces $O(|\mathcal{V}_{p_i p_j}(t_0)|)$ topological events.

Visibility Event. For a point $p_j \in S$, either $\rho(p_j)$ or $\gamma(p_j)$ changes at time t_0. Suppose that $p_i(t_0^-) = \rho(p_j(t_0^-))$ and $p_k(t_0^+) = \rho(p_j(t_0^+))$ with $y(p_k(t_0)) > y(p_i(t_0))$; the other cases are symmetric. Then $p_i(t_0)$, $p_j(t_0)$, and $p_k(t_0)$ are collinear. Furthermore, among

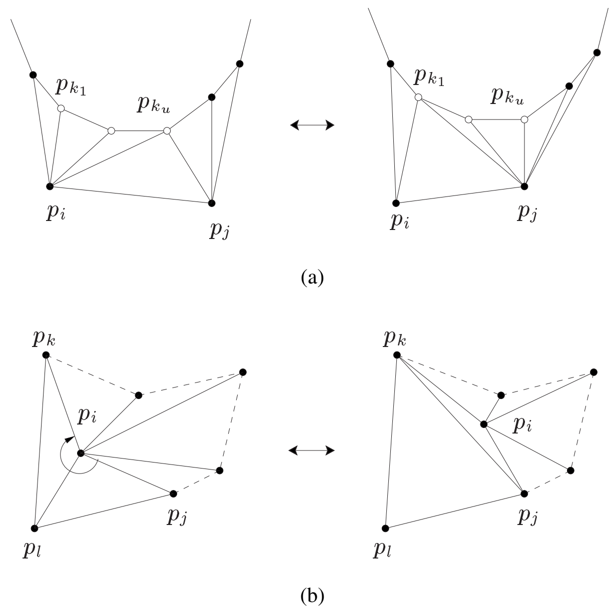

Fig. 2. (a) Ordering event. Points denoted by hollow circles are from $\mathcal{V}_{p_i p_j}(t_0)$. (b) Visibility event.

all the convex edges incident upon p_i at t_0^-, $p_i p_j$ is the leftmost edge. It then follows from property (F2) of the fan triangulation that there is at most one edge (the reflex edge), say $p_i p_l$, between $p_i p_j$ and $p_i p_k$ in clockwise order around p_i. If $p_i p_l$ does not exist, then p_i, p_j, and p_k are collinear on $\partial\, \mathrm{conv}(S(t_0))$. To update the fan triangulation, we delete the edge $p_i p_l$ (if it exists) and insert the edge $p_j p_k$ (Fig. 2(b)). Each visibility event induces $O(1)$ topological events.

In order to detect the above events, we maintain three families of certificates in a global priority queue:

(i) For each edge $p_i p_j \in \mathbb{F}(S(t))$, the next time t_0 at which $y(p_i(t_0)) = y(p_j(t_0))$.
(ii) For each triangle $p_i p_j p_k \in \mathbb{F}(S(t))$, with $y(p_j) < y(p_i) < y(p_k)$, the next time t_0 at which $p_i(t_0) \in p_j(t_0) p_k(t_0)$.
(iii) For each point p_i, let p_j be the point with the minimum y-coordinate so that $p_i = \rho(p_j)$ (resp. $p_i = \gamma(p_j)$) if it exists. We add the time t_0 at which $p_j(t_0)$, $p_i(t_0)$, and $\rho(p_i(t_0))$ (resp. $\gamma(p_i(t_0))$) become collinear.

By our above discussion, it is easy to verify that these certificates detect all events. Moreover, each topological event can be processed in $O(\log n)$ time, including the time spent in updating the global event queue.

3. Constrained Fan Triangulation

In this section we introduce the notion of *constrained fan triangulation* and show how to maintain it under motion. As earlier, let $S = \{p_1, \ldots, p_n\}$ be a set of n (stationary)

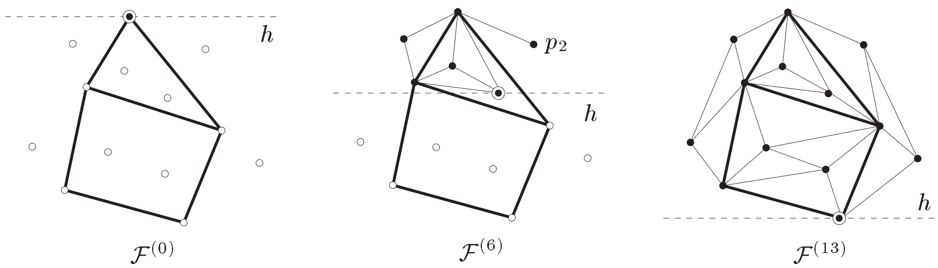

Fig. 3. Constructing constrained fan triangulation with respect to Π (thick edges) at various stages.

points in $\mathbb{R}^2$ sorted in non-increasing order of their y-coordinates, and let Π be a set of segments with pairwise-disjoint interiors whose endpoints lie in S. We are mostly interested in the case in which Π is a triangulation of a subset $R \subseteq S$, but we give the definition for the general setting.

We construct the constrained fan triangulation of S by sweeping a horizontal line h from $y = +\infty$ to $y = -\infty$. Let $\mathbb{F}^{(0)} = \Pi$, and let $\mathbb{F}^{(i-1)}$ be the *partial* triangulation computed after the sweep line has processed p_{i-1}. If the interior of segment $p_i p_j$, for $j < i$, does not intersect $\mathbb{F}^{(i-1)}$, then $p_i p_j$ is called *visible* with respect to $\mathbb{F}^{(i-1)}$. Define $\mathcal{V}(p_i)$, $\rho(p_i)$, and $\gamma(p_i)$ similarly as before, under the modified concept of visibility (with respect to $\mathbb{F}^{(i-1)}$); note that $\rho(p_i)$ and $\gamma(p_i)$ depend on Π. When the sweep line crosses p_i, we compute $\mathbb{F}^{(i)}$ by adding the edges $p_i p_j$ for all vertices p_j in $\mathcal{V}(p_i)$. See Fig. 3. Note that unlike the sweeping process to construct the fan triangulation, $\mathbb{F}^{(i)}$ is not necessarily a constrained fan triangulation of the already swept points. In fact, $\mathbb{F}^{(i)}$ might not even be a triangulation; it is possible that only one point from $p_1, \ldots, p_{i-1}$ is visible from p_i due to the constraint Π, in which case we add only one edge to $\mathbb{F}^{(i)}$, and this "dangling" edge of $\mathbb{F}^{(i)}$ will become part of a triangle at a later stage; see the edge incident upon p_2 in $\mathbb{F}^{(6)}$ of Fig. 3. The final triangulation $\mathbb{F}^{(n)}$ is the constrained fan triangulation of S (with respect to Π), denoted by $\mathbb{F}(S, \Pi)$.

Observe that if $\Pi = \partial\,\mathrm{conv}(S)$, then $\mathbb{F}(S, \Pi) = \mathbb{F}(S)$. If Π is a triangulation of a subset $R \subseteq S$, then Π partitions S into various subsets: $S_\Delta = S \cap \Delta$ (including vertices of Δ), where $\Delta \in \Pi$ is a triangle or $\Delta = \Sigma$ is the exterior of the convex hull of R. By the above observation, the constrained fan triangulation $\mathbb{F}(S, \Pi)$ restricted to S_Δ is the same as $\mathbb{F}(S_\Delta)$ for $\Delta \in \Pi$. Hence, $\mathbb{F}(S, \Pi)$ can be computed by constructing independently $\mathbb{F}(S_\Delta)$ for each triangle $\Delta \in \Pi$ and constructing $\mathbb{F}(S_\Sigma, \partial\,\mathrm{conv}(R))$ within the exterior of $\mathrm{conv}(R)$, i.e., the constrained fan triangulation of S_Σ with respect to the boundary of $\mathrm{conv}(R)$.

The following properties of constrained fan triangulation are generalizations of (F1) and (F2):

(C1) For $j < i$, if p_j is the lowest vertex visible from p_i, then $p_j p_i$ is an edge in $\mathbb{F}(S, \Pi)$.

(C2) If $p_i p_{j_1}, \ldots, p_i p_{j_k}$ are the down edges incident upon p_i sorted in clockwise direction so that no edges of Π lie between them, then either $j_1 < \cdots < j_k$ and $p_i = \rho(p_{j_l})$ for $l < k$ ($l \le k$ if $p_i p_{j_k}$ is not a reflex edge), or $j_1 > \cdots > j_k$ and $p_i = \gamma(p_{j_l})$ for $l < k$ ($l \le k$ if $p_i p_{j_k}$ is not a reflex edge).

Next, we describe how to maintain $\mathbb{F}(S, \Pi)$ as the points in S move. For the time being, we assume that Π is a triangulation of a subset $R \subseteq S$ and that motion is such that the topology of Π does not change and remains a valid triangulation of R throughout the motion. In addition to ordering and visibility events, a new type of event, called the *crossing event*, arises when a point of $S \backslash R$ crosses an edge of Π. In the following we discuss each of them.

Ordering Event. There are two points $p_i, p_j \in S$ so that (i) $p_i(t_0)$ is visible from $p_j(t_0)$ with respect to Π, and (ii) $y(p_i(t_0)) = y(p_j(t_0))$. An ordering event is processed in the same way as in Section 2. If $p_i(t_0)$ is not visible from $p_j(t_0)$, then $\mathbb{F}(S, \Pi)$ does not change at t_0. Note that the visibility relation between two points can change only by a crossing or a visibility event described below.

Visibility Event. For a point $p_i \in S$, either $\rho(p_i)$ or $\gamma(p_i)$ changes at time t_0, and p_i does not cross any edge of Π during $[t_0^-, t_0^+]$. We process this event in the same way as in Section 2. Note that $\rho(p_i)$ or $\gamma(p_i)$ could also change as p_i crosses some edge from Π, which will be covered by the crossing event below.

Crossing Event. A point p_i crosses an edge $p_j p_k$ of Π at time t_0; assume $y(p_k(t_0)) > y(p_j(t_0))$. Recall that Π is a triangulation of a subset of S. Π partitions $\mathbb{R}^2$ into several connected components, each region being either a triangle of Π or the exterior of $\text{conv}(R)$. Suppose p_i moves from the region Δ^- to Δ^+. Let S_{Δ^-} (resp. S_{Δ^+}) be the subset of points of S in Δ^- (resp. Δ^+) at t_0^-. Then the crossing event corresponds to deleting p_i from S_{Δ^-}, inserting it into S_{Δ^+}, and updating the triangulations in Δ^- and Δ^+. First, we consider Δ^-.

Given an arbitrary point $q \in S$, the *star* of q, denoted by $\text{St}(q)$, is the union of triangles adjacent to q. $\text{St}(q)$ is a star-shaped polygon with q in its kernel—every point in $\text{St}(q)$ is visible from q. The *link* of q, denoted by $\text{Lk}(q)$, is defined as $\partial\,\text{St}(q)$. $\text{Lk}(q)$ is a closed polygonal curve. If $q \in \partial\,\text{conv}(S)$, then $q \in \text{Lk}(q)$; otherwise q lies in the interior of $\text{St}(q)$. Given the fan triangulation of S, $\text{Lk}(q)$ consists of a convex chain, corresponding to the up neighbors of q, a y-monotone polygonal chain, corresponding to the down neighbors of q, and one more edge that connects the lowest neighbor of q to an up neighbor of q (or to q if $q \in \partial\,\text{conv}(S)$); see Fig. 4(a).

If $p_i \in \partial\,\text{conv}(S)$ at time t_0^- (in which case Δ^- represents the exterior of $\text{conv}(R)$), we can simply remove edges $p_i p_k$ and $p_i p_j$ from $\mathbb{F}(S(t_0^-), \Pi)$. We now assume that $p_i(t_0^-) \notin \partial\,\text{conv}(S(t_0^-))$, implying that $p_i(t_0^-) \notin \text{Lk}(p_i(t_0^-))$.

Lemma 3.1. *Within Δ^-, any edge from $\mathbb{F}(S(t_0^-), \Pi)$ not incident upon p_i at time t_0^- is present in $\mathbb{F}(S(t), \Pi)$ for all $t \in [t_0^-, t_0^+]$.*

The above lemma is straightforward, as any two points that are previously visible to each other within Δ^- will remain visible after p_i moves out of Δ^-. In view of Lemma 3.1, we delete the edges incident upon p_i at time t_0^- and re-triangulate within $\text{Lk}(p_i(t_0^-))$. The portion of the triangulation on and outside $\text{Lk}(p_i(t_0^-))$ remains unchanged. We re-triangulate within $\text{Lk}(p_i(t_0^-))$ as follows (Fig. 4). Let $Q = \langle p_k = q_0, q_1, \ldots, q_w, q_{w+1}, \ldots, q_u = p_j \rangle$ be the sequence of vertices on $\text{Lk}(p_i)$, where

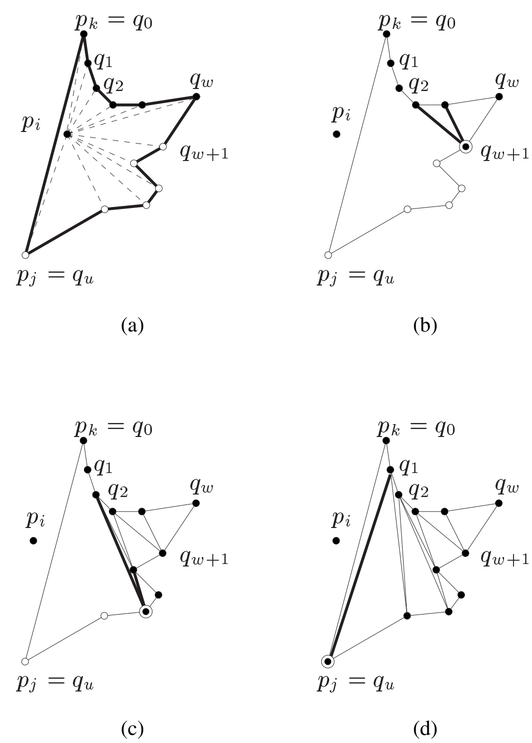

Fig. 4. Re-triangulation of $\text{St}(p_i(t_0^-))$ after deleting p_i. The thick polygonal chain in (a) is $\text{Lk}(p_i(t_0^-))$.

$q_0, q_1, \ldots, q_w$ is the convex chain formed by the up neighbors of p_i, and $y(p_i) > y(q_{w+1}) > y(q_{w+2}) > \cdots > y(q_u)$. Since $q_0, \cdots, q_w$ are already in convex position, we only visit $q_{w+1}, \ldots, q_u$ in order. Without loss of generality, assume that p_i lies to the right of the edge $p_k p_j$ at time t_0^-, i.e., p_i crosses $p_k p_j$ from right to left; the other case can be handled symmetrically. Suppose we have processed $q_{w+1}, \ldots, q_{z-1}$, i.e., added the new edges incident upon them. Letting $\rho^{(i)}(\cdot) = \rho(\rho^{(i-1)}(\cdot))$, we maintain a subsequence of Q,

$$\Phi^{(z-1)} = \langle q_{z-1} = \rho^{(0)}(q_{z-1}), \rho(q_{z-1}), \rho^{(2)}(q_{z-1}), \rho^{(3)}(q_{z-1}), \ldots \rangle,$$

i.e., the vertices that appear on the left boundary of the convex hull of $q_0, \ldots, q_{z-1}$. If the vertices $q_{z-1}, \rho(q_{z-1}), \ldots, \rho^{(l)}(q_{z-1})$ are visible from q_z, we then add the edges $q_z q_{z-1}, q_z \rho(q_{z-1}), \ldots, q_z \rho^{(l)}(q_{z-1})$, delete $q_{z-1}, \ldots, \rho^{(l-1)}(q_{z-1})$ from the sequence, set

$$\Phi^{(z)} = \langle q_z, \rho(q_z) = \rho^{(l)}(q_{z-1}), \rho^{(l+1)}(q_{z-1}), \ldots \rangle,$$

and repeat this process for q_{z+1} unless $z = u$.

Next, we describe how to insert p_i into Δ^+ and construct $\mathbb{F}(S_{\Delta^+} \cup \{p_i\})$ at time t_0^+ from $\mathbb{F}(S_{\Delta^+})$. Roughly speaking, we need to do the opposite of what we did in Δ^-. That is, we identify $\text{Lk}(p_i(t_0^+))$ in $\mathbb{F}(S_{\Delta^+} \cup \{p_i\})$, delete the edges of $\mathbb{F}(S_{\Delta^+})$ that lie within

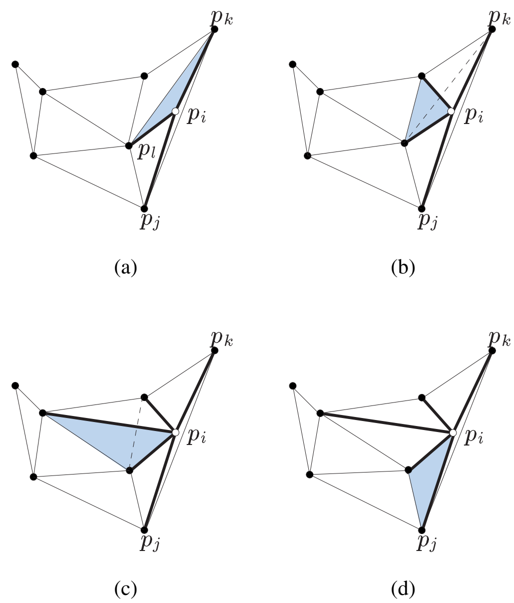

Fig. 5. Updating $\mathbb{F}(S_{\Delta^+})$ after inserting p_i—the shaded triangle is being processed, and the dashed edge was flipped.

the polygon formed by $\mathrm{Lk}(p_i(t_0^+))$, and connect p_i to all the vertices on $\mathrm{Lk}(p_i(t_0^+))$ to form $\mathrm{St}(p_i(t_0^+))$. We assume that Δ^+ lies to the left of $p_k p_j$. The following procedure performs these steps simultaneously.

Let p_l be the vertex in Δ^+ adjacent to the edge $p_j p_k$ at time t_0^-. If p_l does not exist, then $p_j p_k$ is an edge of $\mathrm{conv}(S)$, and p_i becomes a vertex of $\mathrm{conv}(S)$, in which case we simply add the triangle $p_i p_j p_k$. We thus assume that p_l exists. We add the edges $p_i p_j$, $p_i p_k$, and $p_i p_l$. We maintain a stack $\mathcal{S}$ of triangles. Initially, we push $p_i p_l p_j$ and $p_i p_k p_l$ to $\mathcal{S}$ (with the latter being on the top of $\mathcal{S}$). We perform the following procedure until $\mathcal{S}$ becomes empty. An example is illustrated in Fig. 5. For a triangle $p_i p_w p_z$ with $y(p_w) > y(p_z)$, we define the region $\tau(p_i p_w p_z)$ to be the intersection of the following three halfplanes: (i) $y > y(p_z)$; (ii) the halfplane lying below the line $p_i p_w$; and (iii) the halfplane bounded by the line $p_w p_z$ that does not contain p_i; see Fig. 6. Intuitively, this

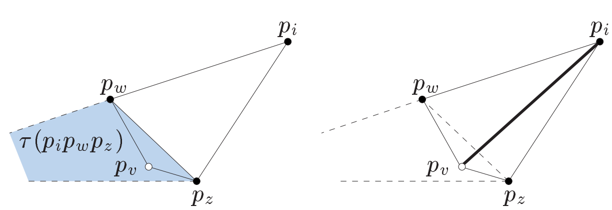

Fig. 6. If p_v lies inside the shaded region $\tau(p_i p_w p_z)$, then an edge flip is performed.

region $\tau(p_i p_w p_z)$ contains points that potentially should be visible to p_i if edge $p_w p_z$ were not present.

(1) Remove the top triangle $p_i p_w p_z$ from $\mathcal{S}$. Assume that $y(p_w) > y(p_z)$.
(2) If $p_w p_z$ is an edge of the convex hull or $y(p_z) > y(p_i)$, go to step (1).
(3) Let p_v be the other vertex adjacent to edge $p_w p_z$. If $p_v \notin \tau(p_i p_w p_z)$, go to step (1).
(4) Delete the edge $p_w p_z$, and insert the edge $p_i p_v$ (an edge flip).
(5) Push the triangle $p_i p_v p_z$ to $\mathcal{S}$.
(6) Push the triangle $p_i p_w p_v$ to $\mathcal{S}$.

By induction, one easily observes that the above algorithm constructs a valid triangulation F within Δ^+. Let F^- and F^+ denote the fan triangulation within Δ^+ at time t_0^- and t_0^+, respectively.

Lemma 3.2. *F as constructed by the above procedure is the same as F^+.*

Proof. Observe that if an edge from F or F^+ does not have p_i as an endpoint, then the edge is present in F^- as well, because the newly added edges in F or F^+ are all adjacent to p_i. Therefore, we only have to prove that the above procedure correctly identifies $\mathrm{Lk}(p_i)$ in F^+. Let $Q = \langle p_k = q_0, q_1, \ldots, q_u = p_j \rangle$ be the sequence of neighbors of p_i in the resulting triangulation F. To prove the lemma, we need to show that (i) any $p_i q_z$ for $0 \le z \le u$ exists in F^+, and (ii) for any $p_i p_z \in F^+$, $p_z \in Q$. The first claim is easily shown by induction on the order in which the edges were added by the above procedure. At any time, the newly added edge is guaranteed to be present in F^+. We sketch the proof for the second claim below.

We prove the second claim by contradiction. By the first claim, Q is a subset of $\mathrm{Lk}(p_i)$ at time t_0^+ in F^+. Therefore, similar to $\mathrm{Lk}(p_i)$, Q consists of a convex chain $q_0, q_1, \ldots, q_w$, corresponding to the up neighbors of p_i in Q, a y-monotone polygonal chain $q_{w+1}, \ldots, q_u$, corresponding to the down neighbors of p_i in Q, and edge $q_0 q_u$ (i.e., $p_k p_j$). See Fig. 7. Moreover, edges $p_i q_0, \ldots, p_i q_u$ are ordered in clockwise or counterclockwise order around p_i. Assume that q_s and q_{s+1} are two consecutive vertices of Q such that there exists a point $q \in \mathrm{Lk}(p_i)$ lying between points q_s and q_{s+1} along $\mathrm{Lk}(p_i)$. Obviously, q lies inside the wedge formed by $\overrightarrow{p_i q_s}$ and $\overrightarrow{p_i q_{s+1}}$, as $\mathrm{Lk}(p_i)$ is star-shaped. We now distinguish two cases, each of which leads to a contradiction:

1. $s + 1 \le w$. As $q_0, \ldots, q_w$ form a subset of a convex chain, triangle $p_i q_s q_{s+1}$ contains point q (Fig. 7(a)). By construction, $p_i q_s q_{s+1}$ is a triangle in F. However, F is a valid triangulation, so triangle $p_i q_s q_{s+1}$ cannot contain any of the input points, a contradiction.
2. $s + 1 > w$. Recall that q lies in the wedge formed by the rays $\overrightarrow{p_i q_s}$ and $\overrightarrow{p_i q_{s+1}}$. If $q \in \Delta p_i q_s q_{s+1}$, then we arrive at the same contradiction as in case 1 (Fig. 7(b)). So q lies in the halfplane bounded by $q_s q_{s+1}$ that does not contain p_i. Moreover, since q_{s+1} lies on the y-monotone part of $\mathrm{Lk}(p_i)$ and q lies between q_s and q_{s+1} along $\mathrm{Lk}(p_i)$, we have $y(q) > y(q_{s+1})$, implying that $q \in \tau(p_i q_s q_{s+1})$. As q_s and q_{s+1} are two consecutive neighbors of p_i in F, $p_i q_s q_{s+1}$ must have been

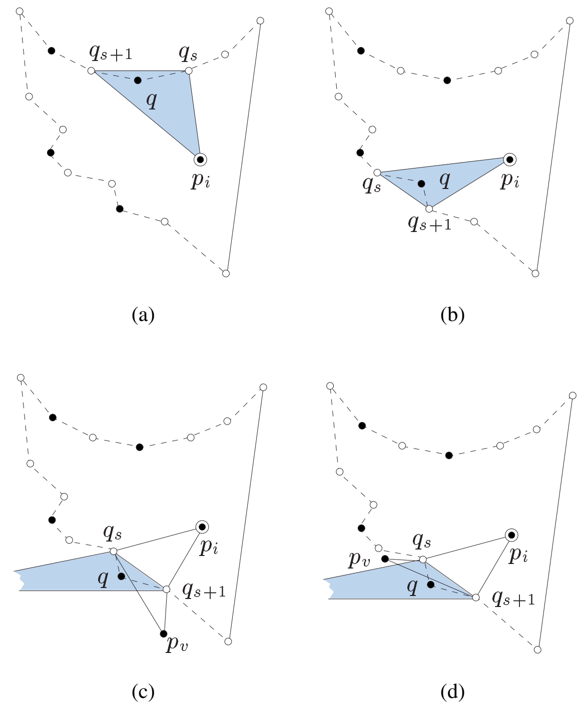

Fig. 7. Dashed polygon is $\mathrm{Lk}(p_i)$, with points denoted by hollow circles being vertices of Q. (a),(b) q lies inside triangle $p_i q_s q_{s+1}$. (c),(d) q lies in the region $\tau(p_i q_s q_{s+1})$ (shaded), but triangle $p_v q_s q_{s+1}$ prevents vertex q from being connected to p_i. All cases lead to a contradiction.

processed at some moment. The fact that $p_i q$ was not added at that time implies that there was another triangle, say, $p_v q_s q_{s+1} \in F^-$, incident upon edge $q_s q_{s+1}$, and $p_v \notin \tau(p_i q_s q_{s+1})$. Notice that q is not contained in either triangle $p_v q_s q_{s+1}$ or $p_i q_s q_{s+1}$. Now assume p_v lies below q_{s+1} (Fig. 7(c)). However, in this case, edge $p_v q_s$ could not have been present in F^- as it crosses edge $q q_{s+1}$ of F^-, but both points q and q_{s+1} are above p_v, a contradiction. The other case in which p_v lies above line $p_i q_s$ can be handled similarly (Fig. 7(d)).

This proves the second claim, and the lemma follows. □

The following lemma follows immediately from the above algorithm.

Lemma 3.3. *The number of topological events induced by a crossing event of p_i is proportional to the old degree plus the new degree of p_i.*

Each topological event can be handled in $O(\log n)$ time, including the time spent in updating the global event queue. To detect all the three types of events, we can maintain the same three families of certificates as for a fan triangulation.

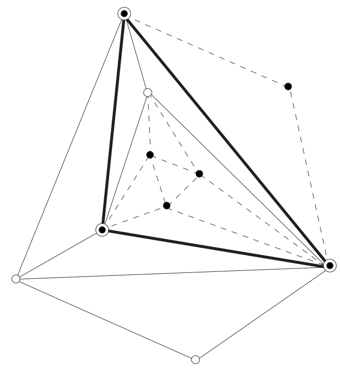

Fig. 8. A hierarchical fan triangulation with three levels—points in R_1 are denoted by double circles, in $R_2 \backslash R_1$ by hollow circles, and in $R_3 \backslash R_2$ by black circles.

4. Hierarchical Fan Triangulation

We now use the constrained fan triangulation to define a hierarchical fan triangulation $\mathbb{F}$. Let $\emptyset = R_0 \subseteq R_1 \subseteq \cdots \subseteq R_w = S$. Set $\mathbb{F}_0 = \emptyset$, and for $i \geq 1$, define $\mathbb{F}_i = \mathbb{F}(R_i, \mathbb{F}_{i-1})$, i.e., $\mathbb{F}_i$ is the constrained fan triangulation of R_i with respect to $\mathbb{F}_{i-1}$. By construction, $\mathbb{F}_1$ is the fan triangulation of R_1 and $\mathbb{F}_{i-1} \subseteq \mathbb{F}_i$. Set $\mathbb{F} = \mathbb{F}_w$. See Fig. 8 for an example.

Hereditary Event. As the points move, each $\mathbb{F}_i$ deforms continuously. The topology of $\mathbb{F}$ changes when there is an ordering, a visibility, or a crossing event in one of the $\mathbb{F}_i$'s, and it can be processed as described in Sections 2 and 3. However, a topological change in $\mathbb{F}_i$ also propagates changes in $\mathbb{F}_j$ for $j > i$, as the insertion or deletion of an edge in $\mathbb{F}_i$ affects the visibility of points in R_j. We refer to an event in $\mathbb{F}_{i+1}$ caused by another event in $\mathbb{F}_i$, as a *hereditary event* of $\mathbb{F}_{i+1}$. We process hereditary events as follows. If $\operatorname{conv}(R_i)$ changes, i.e., a point $q \in R_i$ appears or disappears on $\operatorname{conv}(R_i)$, we update $\mathbb{F}_{i+1}$ within the exterior of $\operatorname{conv}(R_i)$ by inserting or deleting q as described in the crossing event of Section 3. The number of topological events in this case is proportional to the new degree (if q is inserted) or old degree (if q is deleted) of q within the exterior of $\operatorname{conv}(R_i)$. Moreover, if a triangle Δ in $\mathbb{F}_i$ is destroyed due to an edge insertion or deletion, we simply delete all the edges of $\mathbb{F}_{i+1}$ lying inside Δ. As destroying a triangle unavoidably leads to creation of other triangle(s), we only describe how we perform reconstruction after creating a new triangle. Suppose a triangle Δ is created in $\mathbb{F}_i$, we reconstruct the affected portion of $\mathbb{F}_{i+1}$. More specifically, let $R_{i+1}^{\Delta} = R_{i+1} \cap \Delta$. We construct the fan triangulation $\mathbb{F}(R_{i+1}^{\Delta})$ inside Δ. The number of topological events in each newly created triangle Δ can be bounded by $O(|R_{i+1}^{\Delta}|)$. Each topological event induced by hereditary events can be handled in $O(\log n)$ time, including the time spent in updating the global event queue.

In the next section we analyze the performance of the hierarchical fan triangulation, assuming that R_i is a random subset of R_{i+1} of an appropriate size.

5. Analysis

In the hierarchical fan triangulation $\mathbb{F}$ defined in Section 4, we choose R_{i-1} to be a random subset of R_i of size

$$\min\{|R_i|, 5|R_i|^{1-1/i} \log|R_i|\}, \tag{1}$$

for $1 < i \le w$. Let $w = \lceil\sqrt{\log n/\log\log n}\rceil$. In this section we show that $\mathbb{F}$ has a near-quadratic number of topological changes as promised. To do this, we first focus on a specific level of the construction of $\mathbb{F}$. For notational convenience, let $P = R_i$, $R = R_{i-1}$, $\mathcal{T}(P) = \mathbb{F}_i$, $\mathcal{T}(R) = \mathbb{F}_{i-1}$, $n = |P|$, and $r = n^{1-1/i}$, for some $1 < i \le w$. It follows that $\mathcal{T}(P) = \mathbb{F}(P, \mathcal{T}(R))$, and $|R| = \min\{n, 5r \log n\}$.

Let $p, p_1, p_2, p_3 \in P$, $t, t_1, t_2 \in \mathbb{R}$, and $m \in \mathbb{N}$. Let $h(p_1, p_2)$ be the open halfplane to the left of the oriented line $\overrightarrow{p_1 p_2}$. We define (see Fig. 9)

$$\begin{aligned}
\langle p_1, p_2, p_3; t\rangle &= \{p \in P \mid p(t) \in \Delta p_1(t)p_2(t)p_3(t)\},\\
\langle p_1, p_2; t_1, t_2\rangle_s &= \{p \in P \mid \exists t \in [t_1, t_2] \text{ s.t., } p(t) \in p_1(t)p_2(t)\},\\
\langle p_1, p_2; t_1, t_2\rangle_h &= \{p \in P \mid \exists t \in [t_1, t_2] \text{ s.t., } p(t) \in h(p_1(t), p_2(t))\},\\
\langle p; m; t\rangle_\rho &= \{\rho^{(k)}(p(t)) \mid 1 \le k \le m\},\\
\langle p; m; t\rangle_\gamma &= \{\gamma^{(k)}(p(t)) \mid 1 \le k \le m\}.
\end{aligned}$$

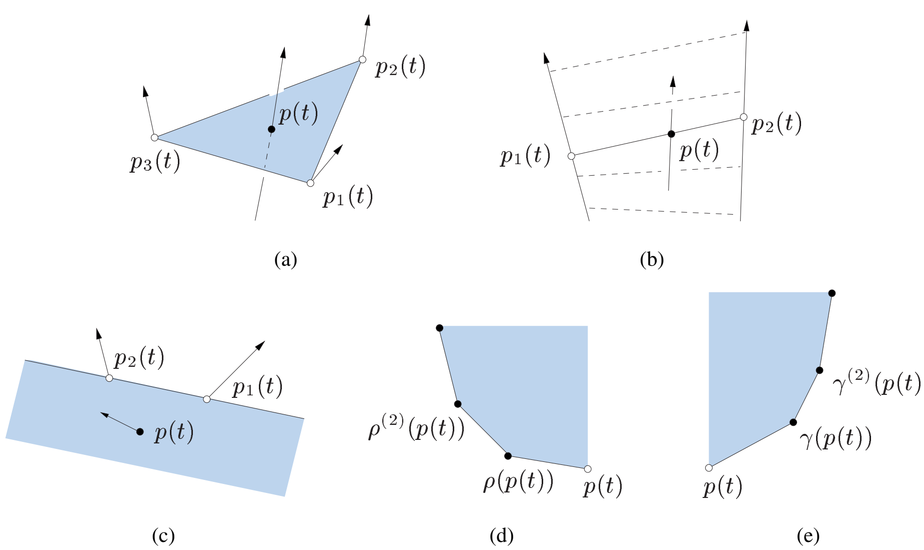

Fig. 9. Examples of different ranges: (a) $\langle p_1, p_2, p_3; t\rangle$, (b) $\langle p_1, p_2; t_1, t_2\rangle_s$, (c) $\langle p_1, p_2; t_1, t_2\rangle_h$, (d) $\langle p; m; t\rangle_\rho$, (e) $\langle p; m; t\rangle_\gamma$.

We then let

$$\mathcal{R}_1 = \{\langle p_1, p_2, p_3; t\rangle \mid p_1, p_2, p_3 \in P, t \in \mathbb{R}\},$$
$$\mathcal{R}_2 = \{\langle p_1, p_2; t_1, t_2\rangle_s \mid p_1, p_2 \in P, t_1, t_2 \in \mathbb{R}\},$$
$$\mathcal{R}_3 = \{\langle p_1, p_2; t_1, t_2\rangle_h \mid p_1, p_2 \in P, t_1, t_2 \in \mathbb{R}\},$$
$$\mathcal{R}_4 = \{\langle p; m; t\rangle_\rho \mid p \in P, m \in \mathbb{N}, t \in \mathbb{R}\},$$
$$\mathcal{R}_5 = \{\langle p; m; t\rangle_\gamma \mid p \in P, m \in \mathbb{N}, t \in \mathbb{R}\}.$$

Finally, let $\mathbb{X} = (P, \mathcal{R})$ be the range space with $\mathcal{R} = \bigcup_{1\le j\le 5} \mathcal{R}_j$. We have the following lemma.

Lemma 5.1. $|\mathcal{R}| = O(n^5)$.

Proof. Consider three points p_1, p_2, and p_3 of P. As t increases, $\langle p_1, p_2, p_3; t\rangle$ changes only when some point $q \in P$ moves in or out of the triangle $p_1 p_2 p_3$. For any point q, this can only happen a constant number of times. Thus $\langle p_1, p_2, p_3; t\rangle$ can change $O(n)$ times as t goes from $-\infty$ to $+\infty$ for fixed p_1, p_2, and p_3. There are $O(n^3)$ different choices of p_1, p_2, p_3, implying that $|\mathcal{R}_1| = O(n^4)$. Similarly we can prove $|\mathcal{R}_2| = O(n^4)$ and $|\mathcal{R}_3| = O(n^4)$.

For any $p, q \in P$, if $y(p(t)) < y(q(t))$, let $\theta(p(t), q(t))$ denote the angle formed by $\overrightarrow{pq}$ and the $+x$-direction at time t; $\theta(p(t), q(t))$ is undefined at t otherwise. The number of changes to $\rho(p)$ is the same as the complexity of the upper envelope of $\{\theta(p(t), q(t)) \mid q \neq p, q \in P\}$, which can be bounded by $\lambda_s(n)$ [16], where $\lambda_s(\cdot)$ is the maximum length of a Davenport–Schinzel sequence of order s. The value of s depends on the maximum degree of polynomial of the trajectories of points in P; for example, $s = 4$ if points move linearly. For fixed p and m, as t increases, $\langle p; m; t\rangle_\rho$ changes only when $\rho(q)$ changes for some $q \in P$. There are at most n different choices for p and m each, thus we have $|\mathcal{R}_4| = O(n^3\lambda_s(n))$, which is slightly larger than $O(n^4)$. Similarly $|\mathcal{R}_5| = O(n^3\lambda_s(n))$. Putting everything together, $|\mathcal{R}| = O(n^3\lambda_s(n)) = O(n^5)$. □

As time t varies, the topological structure of $\mathcal{T}(P)$ may change: features such as edges or triangles may appear or disappear. The *lifetime* of a feature is the period between two of its consecutive appearances and disappearances. If an edge or a triangle occurs more than once, we count each occurrence as a different feature.

By (1), Lemma 5.1, and standard results from random sampling theory [15], R is a $(1/r)$-net of the range space $\mathbb{X}$ with high probability. Thus we obtain the following lemmas.

Lemma 5.2. *With high probability*:

(a) *At any time, there are less than n/r points of P inside any triangle of $\mathcal{T}(R)$.*
(b) *Less than n/r points of P can cross any edge ever appearing in $\mathcal{T}(R)$ during its lifetime.*
(c) *Less than n/r points of P can ever appear in $h(p, q)$ during the period when $p, q \in R$ are two neighboring points on the convex hull of R in clockwise order.*

Lemma 5.3. *With high probability*:

(a) *For any* $p \in P, m \in \mathbb{N}, t \in \mathbb{R}$,

$$|\langle p; m; t\rangle_\rho| \leq (|\langle p; m; t\rangle_\rho \cap R| + 1) \cdot n/r.$$

(b) *For any* $p \in P, m \in \mathbb{N}, t \in \mathbb{R}$,

$$|\langle p; m; t\rangle_\gamma| \leq (|\langle p; m; t\rangle_\gamma \cap R| + 1) \cdot n/r.$$

Proof. Suppose that

$$\langle p; m; t\rangle_\rho \cap R = \{\rho^{(i_1)}(p(t)), \rho^{(i_2)}(p(t)), \ldots, \rho^{(i_v)}(p(t))\},$$

where $1 \leq i_1 < i_2 < \cdots < i_v \leq m$. Let $i_0 = 0$ and $i_{v+1} = m + 1$. We have

$$\langle p; m; t\rangle_\rho = \left(\bigcup_{j=1}^{v+1} \langle \rho^{(i_{j-1})}(p(t)); i_j - i_{j-1} - 1; t\rangle_\rho\right) \cup \{\rho^{(i_1)}(p(t)), \rho^{(i_2)}(p(t)), \ldots, \rho^{(i_v)}(p(t))\}.$$

Moreover, by the properties of a $(1/r)$-net of $\mathbb{X}$, we know that with high probability,

$$|\langle \rho^{(i_{j-1})}(p(t)); i_j - i_{j-1} - 1; t\rangle_\rho| \leq n/r - 1,$$

for $1 \leq j \leq v + 1$. We thus obtain

$$|\langle p; m; t\rangle_\rho| \leq (v + 1) \cdot (n/r - 1) + v \leq (v + 1) \cdot n/r.$$

We can prove (b) in a similar manner. □

Let $\deg(p(t))$ denote the degree of point $p \in P$ in $\mathcal{T}(P)$ at time t.

Lemma 5.4. *If a point $p \in P$ lies on an edge of $\mathcal{T}(R)$ at time t_0, then both* $\deg(p(t_0^-))$ *and* $\deg(p(t_0^+))$ *are bounded by n/r with high probability.*

Proof. If $p(t_0^-)$ or $p(t_0^+)$ lies in the interior of $\mathrm{conv}(R)$, then we have $\deg(p(t_0^-)) \leq n/r$ or $\deg(p(t_0^+)) \leq n/r$ with high probability by Lemma 5.2(a). If p crosses an edge p_1p_2 of $\partial\,\mathrm{conv}(R)$ at t_0, say, from inside to outside, then at t_0^+, all points adjacent to p lie in the open halfplane bounded by p_1p_2 that is disjoint from $\mathrm{conv}(R)$. It then follows from Lemma 5.2(c) that $\deg(p(t_0^+)) \leq n/r$. A symmetric argument shows $\deg(p(t_0^-)) \leq n/r$ if p moves from the exterior to the interior of $\mathrm{conv}(R)$ at t_0. □

For a set $A \subseteq P$ of points, let $\psi(A)$ denote the total number of topological changes to $\mathcal{T}(A)$ over time. We will bound $\psi(P)$ in terms of $\psi(R)$, where R is a $(1/r)$-net with respect to the range space $\mathbb{X}$. For simplicity, we omit the phrase "with high probability." We now bound the number of topological changes induced by each type of event as discussed in Section 3 and 4.

Hereditary Events. Each hereditary event either causes us to re-triangulate $\mathcal{T}(P)$ inside $O(1)$ triangles of $\mathcal{T}(R)$, or insert or delete a point of R that appears or disappears on $\partial\,\mathrm{conv}(R)$. By Lemma 5.2(a) and 5.4, each such event induces $O(n/r)$ topological changes. Since each hereditary event is caused by a topological event of $\mathcal{T}(R)$, the total number of topological changes induced by the hereditary events is

$$O(n/r)\cdot\psi(R). \tag{2}$$

Crossing Events. By Lemma 5.2(b), at most n/r points cross an edge of $\mathcal{T}(R)$ during its lifetime. By Lemmas 3.3 and 5.4, each such event induces $O(n/r)$ topological changes. Since there are $O(\psi(R)+|R|)$ distinct edges in $\mathcal{T}(R)$ over the entire history, the total number of topological changes induced by crossing events is

$$O((n/r)\cdot(n/r)\cdot(\psi(R)+|R|)) = O(n^2/r^2)\cdot\psi(R). \tag{3}$$

Visibility Events. Recall that a visibility event occurs when there is a change in $\rho(p)$ or $\gamma(p)$ with respect to $\Delta\in\mathcal{T}(R)$ for some point $p\in P_\Delta$ (recall P_Δ is the set of points of P inside Δ). As discussed earlier, each such event induces a constant number of topological changes. Therefore, we only need to bound the number of visibility events.

A visibility event occurs due to the collinearity of p and two other points q and z, where p, q, and z lie within the same face, say Δ, of $\mathcal{T}(R)$. If Δ is a triangle in $\mathcal{T}(R)$ (i.e., $p\in\mathrm{conv}(R)$), then by Lemma 5.2, only $O(n/r)$ points can ever appear in Δ during its lifetime. Using a lower-envelope argument as in the proof of Lemma 5.1, it can be shown that the number of visibility events happening to p while $p\in\Delta$ is $O(\lambda_s(n/r))$. Since there are $O(\psi(R)+|R|)$ triangles in $\mathcal{T}(R)$ over the entire history, the total number of visibility events inside $\mathrm{conv}(R)$ is

$$O((n/r)\cdot\lambda_s(n/r)\cdot(\psi(R)+|R|)) = O(n^2\log(n)/r^2)\cdot\psi(R).$$

Next, suppose a point $p\notin\mathrm{conv}(R)$ causes a visibility event. Let l_{t} (resp. l_{b}) be the horizontal line passing through the highest (resp. lowest) point of R. We distinguish three cases: (a) p lies between l_{t} and l_{b} and to the left (resp. right) of $\mathrm{conv}(R)$, and $\gamma(p)$ (resp. $\rho(p)$) changes; (b) p lies between l_{t} and l_{b} and to the left (resp. right) of $\mathrm{conv}(R)$, and $\rho(p)$ (resp. $\gamma(p)$) changes; and (c) p lies above l_{t} or below l_{b}.

It is easy to observe that cases (b) and (c) at $t=t_0$ correspond to a change in

$$\arg\max_{q\in P}\theta(p(t),q(t)) \quad\text{or}\quad \arg\min_{q\in P}\theta(p(t),q(t))$$

at time t_0. Therefore, there are an $O(n\lambda_s(n)) = O(n^2\log n)$ number of such events, by a lower-envelope argument similar to the one in the proof of Lemma 5.1.

Now consider case (a) (depicted in Fig. 10). Assume without loss of generality that p lies to the left of $\mathrm{conv}(R)$. Let p_1p_2 be the edge on $\mathrm{conv}(R)$ that is hit by the ray emanating from p towards its right (assuming p_1, p_2 are in clockwise order). Observe that q, z, the two points that are involved in the visibility event with p, must lie in the triangle pp_1p_2, and therefore $p,q,z\in h(p_1,p_2)$. Thus, during the period when p_1, p_2 are two neighboring points on $\mathrm{conv}(R)$ in clockwise order, the number of events that three points of P in $h(p_1,p_2)$ become collinear is $O(n^3/r^3)$, by Lemma 5.2(c). If

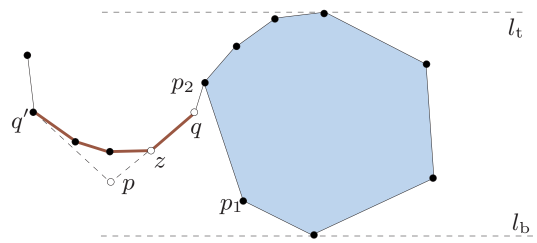

Fig. 10. The shaded region is conv(R), and the thick chain is $\mathcal{V}(p)$.

points move along algebraic trajectories, the combinatorial structure of conv(R) changes $O(|R|\lambda_s(|R|)) = O(|R|^2 \log n)$ times [16]. Therefore there are $O(n^3|R|^2 \log(n)/r^3)$ such events.

In total, the number of topological changes to the triangulation of P caused by visibility events is bounded by

$$O(n^3|R|^2 \log(n)/r^3 + n^2 \log n) + O(n^2 \log(n)/r^2) \cdot \psi(R). \tag{4}$$

Ordering Events. The number of topological changes induced by an ordering event at time t_0 when the y-coordinates of two mutually visible points p and q become equal is proportional to $|\mathcal{V}_{pq}(t_0)|$, i.e., the number of points that are visible from both p and q at time t_0. If p and q lie inside a triangle face of $\mathcal{T}(R)$, then $|\mathcal{V}_{pq}(t_0)| = O(n/r)$. Since there are $O(n^2)$ ordering events, the total number of topological changes induced by such events is $O(n^3/r)$. Now assume that p and q lie in the exterior of conv(R), and let I be the set of such events. Obviously, $|I| \le n^2$. Denote by m_i the number of points of R on $\mathcal{V}_{pq}(t_0)$ for the ith such event from I. By Lemma 5.3, the total number of topological changes caused by this type of event is at most

$$\sum_i (m_i + 2)n/r = O(n^3/r) + (n/r) \cdot \sum_i m_i.$$

Lemma 5.5. $\sum_{i \le |I|} m_i = O(n^2|R|^2 \log(n)/r^2)$.

Proof. Suppose two points p and q lying outside conv(R) induce an ordering event at time t_0, i.e., $y(p(t_0)) = y(q(t_0))$ and p is visible from q at t_0. Let $z_1, z, z_2 \in R$ be three consecutive points on $\partial\,$conv(R) in clockwise order at t_0 and $z \in \mathcal{V}_{pq}(t_0)$. Recall that $h(z_1, z)$ is the open halfplane to the left of $\overrightarrow{z_1 z}$. We have that $p, q \in h(z_1, z) \cup h(z, z_2)$ as z is visible to both p and q. For fixed edges $z_1 z$ and $z z_2$ of $\partial\,$conv(R), z can only be involved in $O(n^2/r^2)$ such events because, by Lemma 5.2(c), only $O(n/r)$ points ever appear in $h(z_1, z)$ (resp. $h(z, z_2)$) during the lifetime of edge $z_1 z$ (resp. $z z_2$). Let u_z be the total number of times that z is involved in such an event. As points in R move, conv(R) may change $O(|R|\lambda_s(|R|)) = O(|R|^2 \log n)$ times [16], therefore the number of distinct triples of consecutive vertices on $\partial\,$conv(R) is $O(|R|^2 \log n)$. We then have

$$\sum_{z \in R} u_z = O(n^2|R|^2 \log(n)/r^2).$$

Using a simple double counting argument, we obtain that $\sum_{i\leq|I|} m_i = \sum_{z\in R} u_z$. This proves the lemma. □

It follows from the above lemma that the total number of topological changes induced by ordering events is bounded by

$$O(n^3/r + n^3|R|^2\log(n)/r^3) = O(n^3|R|^2\log(n)/r^3). \tag{5}$$

Summing (2)–(5) and substituting the values of $|R|$ and r, we obtain the following recurrence for $\psi(P)$:

$$\psi(P) \leq O(n^{2+1/i}\log^3 n) + O(n^{2/i}\log n)\cdot\psi(R).$$

Returning to the hierarchical fan triangulation, let n_i be the size of R_i, then

$$\psi(n_i) \leq c_1 n_i^{2/i}\log(n_i)\psi(n_{i-1}) + c_2 n_i^{2+1/i}\log^3 n_i,$$

where $c_1, c_2 > 0$ are constants, and

$$n_{i-1} = |R_{i-1}| = \min\{n_i, 5n_i^{1-1/i}\log n_i\}.$$

It can be verified by induction that the solution to the above recurrence is

$$\psi(n_i) \leq n_i^{2+1/i}\log^{3i}(n_i)2^{ci},$$

where c is a sufficiently large constant. In particular, for $i = w = \lceil\sqrt{\log n/\log\log n}\rceil$, i.e., the hierarchical fan triangulation, we have

$$\psi(n) = n^2 2^{O(\sqrt{\log n\cdot\log\log n})}.$$

We conclude with the following main result.

Theorem 5.6. *Let S be a set of n points moving in $\mathbb{R}^2$. If the motion of S is algebraic, a triangulation of S can be maintained by a randomized algorithm so that the number of topological events processed by the algorithm is $n^2 2^{O(\sqrt{\log n\cdot\log\log n})}$ with high probability, and each topological event requires $O(\log n)$ time.*

As a special case of our problem, we consider a scenario in which only k out of n points are moving. By extending our previous technique, we can show that there exists a triangulation whose number of topological changes is roughly $O(nk)$. We describe this triangulation briefly.

The overall framework follows that for the case with n moving points. The major difference is that, at each level of the hierarchical fan triangulation, instead of sampling the points all at once, we sample the static and moving points separately. For any point set P, let $P^s \subseteq P$ be the set of static points in P, and let $P^t \subseteq P$ be the set of moving points in P.

Set $w = \lceil\sqrt{\log k/\log\log n}\rceil$. Let $\emptyset = R_0 \subseteq R_1 \subseteq \cdots \subseteq R_w = S$, such that for $1 < i \leq w$, R_{i-1}^s is a random subset of R_i^s of size

$$\min\{|R_i^s|, O(|R_i|\cdot|R_i^t|^{-1/i}\log|R_i^s|)\},$$

and R^t_{i-1} is a random subset of R^t_i of size

$$\min\{|R^t_i|, O(|R^t_i|^{1-1/i}\log|R^t_i|)\}.$$

We construct the hierarchical fan triangulation $\mathbb{F}$ of S as described in Section 4.

Let $n_i = |R_i|$ and $k_i = |R^t_i|$, for $1 \le i \le w$. By the same method as in [2], one can show that the fan triangulation $\mathbb{F}_1$ changes $O(k_1^{4/3}\lambda_d(n_1)) = O(n_1 k_1^{4/3}\log n_1)$ times. Using this fact, and by similar analysis as above, it can be shown that the number of topological changes to $\mathbb{F}_i$ is bounded by $n_i k_i^{1+1/i}\log^{3i}(n_i)2^{ci}$ for some constant c, for $1 \le i \le w$. In particular, the number of topological changes to the hierarchical fan triangulation is $nk \cdot 2^{O(\sqrt{\log k \cdot \log\log n})}$. Without going into the detail, we conclude with the following theorem.

Theorem 5.7. *Let S be a set of n points in $\mathbb{R}^2$, with k of them moving with algebraic trajectories of bounded degree. A triangulation of S can be maintained by a randomized algorithm so that the number of topological events processed by the algorithm is $nk \cdot 2^{O(\sqrt{\log k \cdot \log\log n})}$ with high probability, and each topological event requires $O(\log n)$ time.*

6. Conclusions

In this paper we have described a randomized algorithm for maintaining a triangulation of a set of moving points in $\mathbb{R}^2$. If the motion is algebraic, then the expected number of topological events is $n^2 2^{O(\sqrt{\log n \cdot \log\log n})}$. Our result almost matches the $\Omega(n^2)$ lower bound, and improves over the previously best known result [2] by nearly a factor of $n^{1/3}$. If only k points of the point set are moving, the expected number of topological changes reduces to $nk \cdot 2^{O(\sqrt{\log k \cdot \log\log n})}$.

Although the triangulation that we maintain is conceptually simple, it is not clear how to derandomize the algorithm efficiently. It will be interesting to find a simple deterministic triangulation, easy to implement, and with a near-quadratic number of topological changes. A major open problem in this area is, of course, bounding the number of topological events in the Delaunay triangulation of a set of moving points. The best known upper bound is near cubic while the best known lower bound is quadratic [6].

Another interesting problem is whether Steiner points can help to reduce the complexity of kinetic triangulations. Agarwal et al. [1] showed that the $\Omega(n^2)$ lower bound on the total number of topological changes to a triangulation still holds even if $O(n)$ Steiner points are allowed and can move along any continuous trajectories. However, we expect that with the introduction of Steiner points, a much simpler triangulation with a near-quadratic number of topological changes may be found.

Acknowledgments

We thank Herbert Edelsbrunner, Leo Guibas, Alper Üngör, and the two anonymous referees for their helpful comments.

References

1. P. K. Agarwal, J. Basch, M. de Berg, L. J. Guibas, and J. Hershberger, Lower bounds for kinetic planar subdivisions, *Discrete Comput. Geom.*, **24** (2000), 721–733.
2. P. K. Agarwal, J. Basch, L. J. Guibas, J. Hershberger, and L. Zhang, Deformable free space tiling for kinetic collision detection, *Internat. J. Robotics Res.*, **21** (2002), 179–197.
3. P. K. Agarwal et al., Algorithmic issues in modeling motion, *ACM Comput. Surveys*, **24** (2002), 550–572.
4. P. K. Agarwal, L. J. Guibas, J. Hershberg, and E. Veach, Maintaining the extent of a moving point set, *Discrete Comput. Geom.*, **26** (2001), 253–274.
5. P. K. Agarwal and C. M. Procopiuc, Advances in indexing mobile objects, *IEEE Bull. Data Engr.*, **25** (2002), 25–34.
6. G. Albers, L. J. Guibas, J. Mitchell, and T. Roos, Voronoi diagrams of moving points, *Internat. J. Comput. Geom. Appl.*, **8** (1998), 365–380.
7. M. Atallah, Some dynamic computational geometry problems, *Comput. Math. Appl.*, **11** (1985), 1171–1181.
8. J. Basch, L. J. Guibas, and J. Hershberger, Data structures for mobile data, *J. Algorithms*, **31** (1999), 1–28.
9. M. W. Bern and D. Eppstein, Mesh generation and optimal triangulation, in: *Computing in Euclidean Geometry* (D.-Z. Du and F. K.-M. Hwang, eds.), 2nd edition, number 4 in Lecture Notes Series on Computing, World Scientific, Singapore, 1995, pp. 47–123.
10. L. P. Chew, Near-quadratic bounds for the L_1 Voronoi diagram of moving points, *Comput. Geom. Theory Appl.*, **7** (1997), 73–80.
11. J. Donea, Arbitrary Lagrangian-Eulerian finite element methods, in: *Computational Methods for Transient Analysis* (T. B. Belytschk and T. J. R. Hughes, eds.), North-Holland, Elsevier, Amsterdam, 1983, pp. 474–516.
12. H. Edelsbrunner, Triangulations and meshes in computational geometry, *Acta Numer.*, **9** (2000), 133–213.
13. J. Erickson, D. Guoy, J. Sullivan, and A. Üngör, Building space–time meshes over arbitrary spatial domains, *Proc. 11th Internat. Meshing Roundtable*, 2002, pp. 391–402.
14. L. J. Guibas, Kinetic data structures: a state of the art report, in: *Robotics: The Algorithmic Perspective* (P. K. Agarwal, L. Kavraki, and M. Mason, eds.), A. K. Peters, Natick, MA, 1998, pp. 191–209.
15. D. Haussler and E. Welzl, ε-Nets and simplex range queries, *Discrete Comput. Geom.*, **2** (1987), 127–151.
16. M. Sharir and P. K. Agarwal, *Davenport–Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

Received July 27, 2004, *and in revised form June* 30, 2005. *Online publication September* 11, 2006.