On non-smooth convex distance functions

Ngọc-Minh Lê *

(Extended abstract)

Abstract

Under the Euclidean metric in 3-space, the bisectors of three points intersect in at most one connected component — namely, a line. In contrast to this, we show that, under non-smooth convex distance functions, there is no general upper bound to the number of connected components of the intersection of the bisectors of three points in 3-space. Our result is important for the further study of abstract Voronoi diagrams in 3-space.

1 Introduction

Voronoi diagrams are an important construct in computational geometry. In high-dimensional Euclidean spaces, there are few studies of Voronoi diagrams under convex distance functions other than the Euclidean one. In this paper, we continue the study of general convex distance functions initiated in [4, 6], where only smooth distance functions are considered. The reasons for this study are: 1. To generalize the abstract Voronoi diagram from 2-space [5] to 3-space, we need to know structural properties of bisectors, as well as of their intersections. The work of Boissonnat *et al* [2] indicates that a direct approach to computing Voronoi diagrams under general polyhedral distance functions might be very difficult. 2. Study of a restricted class of the multivariate generalization of the Davenport-Schinzel problem [10]. 3. Knowing properties of bisectors may help in solving problems from related areas, e.g. [8].

We prove some structural properties for the intersection of bisectors, and that there is no universal bound on the number of connected components of the intersection of bisectors under non-smooth distance functions.

2 Preliminaries

Let M be a convex body in d-space containing the origin O in its interior. For any $r \geq 0$ and $p \in \mathbb{R}^d$, we call the homothet $p + rM$ of M the *M-sphere centered at p with radius r*. Given any points x and $y \in \mathbb{R}^d$, we define the distance induced by M from x to y, denoted $d_M(x, y)$, as the radius of the smallest M-sphere centered at x that contains y. Equivalently, set $N = -M$, then the ray from y to x intersects the boundary of $N + y$ at a point x'. Clearly, we have $d_M(x, y) = |xy|/|x'y|$. We call $d_M(\cdot, \cdot)$ the *convex distance function induced by M*. Call *d_M smooth* if M is smooth.

We define the *bisector* of two distinct points p and $q \in \mathbb{R}^d$, denoted $B(p, q)$, to be the set of all points $x \in \mathbb{R}^d$ so that $d_M(x, p) = d_M(x, q)$ holds.

Let $\mathcal{C}$ be a convex body in $\mathbb{R}^d$. Let m be a point on the boundary of $\mathcal{C}$. A *supporting line* of $\mathcal{C}$ at m is a line contained in a supporting hyperplane of $\mathcal{C}$ at m.

Any two distinct points in $\mathbb{R}^d$ are called to be *in general position* with respect to the distance function d_M if any supporting line of M that is parallel to the line connecting them has only one point in common with M. From now on we always assume the general-position condition for any pair of sites we consider.

We now recall the facts needed to prove our results in the next section. Most of these facts can be found in [4], they hold also in non-smooth distance functions. Denote by $\Sigma_{pq}(\mathcal{C})$ the set of all points of the boundary of $\mathcal{C}$ at which there is a supporting line parallel to pq. In the planar case, the set $\Sigma_{pq}(\mathcal{C})$

*EsHag Research, Rönselstr. 17, 58135 Hagen, Germany. Email: minhle@ibm.net

consists of exactly two points, denoted $\sigma_{pq,1}(C)$ and $\sigma_{pq,2}(C)$. In the 3-dimension case, the set $\Sigma_{pq}(C)$ is a simple closed curve.

For any $p \in \mathbb{R}^d$, call the *placement of* N *at* p the set $N(p) = N + p$. For simplicity, we set $\Sigma_{pq}(p) = \Sigma_{pq}(N(p))$. The set $\Sigma_{pq}(p)$ divides $\partial N(p)$ into two parts, one of them is closer to q [1], denote it $F_{pq}(p)$. Let x be any point of the bisector of p and q. The ray from p to x intersects $\partial N(p)$ at a point u of $F_{pq}(p)$. Call u the *footmark* on $N(p)$ of x. The set of all footmarks on $N(p)$ of $B(p,q)$ is the set $F_{pq}(p)$.

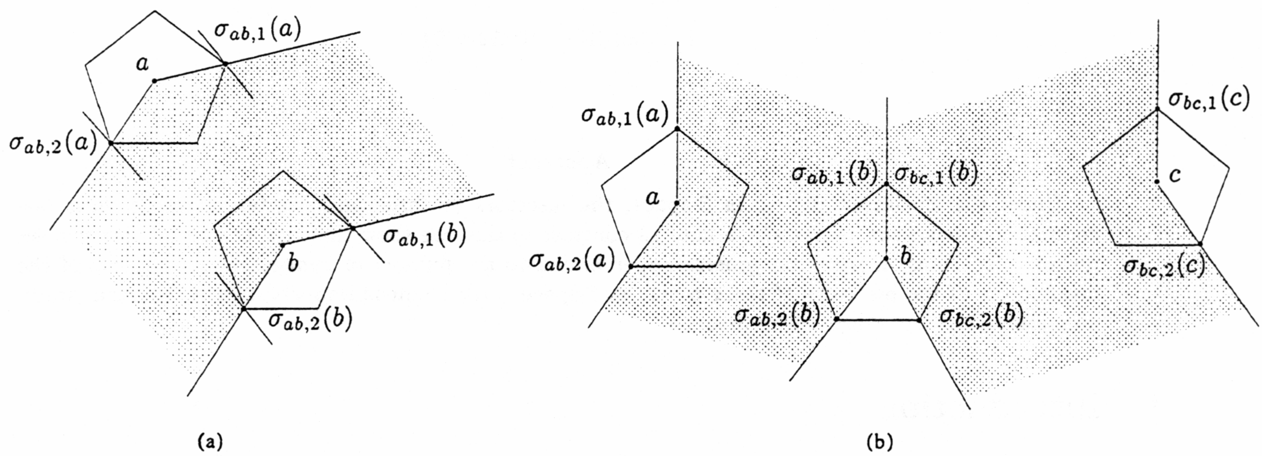

Figure 1: Strip bounding the bisector.

To avoid notational confusions in latter sections, we refer to sites in the plane often as a, b and c instead of p, q and r. The strip bounded by the rays from a to the points of $\Sigma_{ab}(a)$ and from b to the points of $\Sigma_{ab}(b)$ bounds the bisector of a and b, call it the *strip bounding the bisector of* a *and* b, see Figure 1(a).

Set $B(p,q,r) = B(p,q) \cap B(q,r)$. We now describe how to reduce the problem of constructing $B(p,q,r)$ to a 1-dimensional family of planar problems.

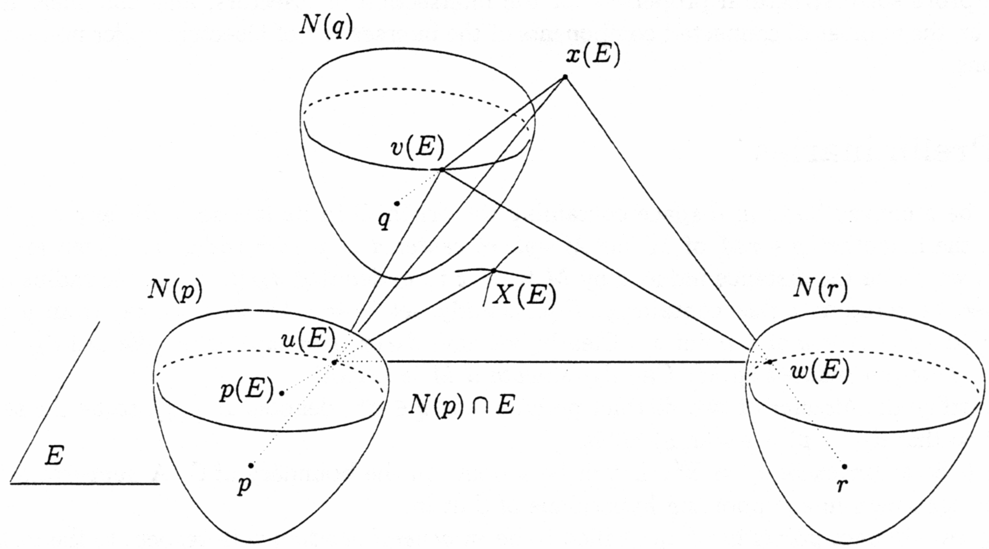

Figure 2: Constructing $B(p,q,r)$. We down-scale the spheres so that they are disjoint.

Let x be any point of $B(p,q,r)$. Let u, v and w be the footmarks of x on $N(p), N(q)$ and $N(r)$, respectively. The footmarks span a plane $\langle u, v, w\rangle$ parallel to $\langle p, q, r\rangle$, the plane spanned by the sites p, q and r. In each plane spanned by the footmarks there are no other footmarks of any point of $B(p,q,r)$. Set

[1] We down-scale $N(p)$ and $N(q)$ so that they are disjoint.

$E = \langle u, v, w \rangle$. It has been shown that u, v and w are exactly the common points of a unique convex curve homothetic to the reflection of $N(p) \cap E$ (with respect to a point in its interior) with the cross section of $N(p), N(q)$ and $N(r)$ with E, respectively. Conversely, if E is a plane parallel to $\langle p, q, r \rangle$ that properly intersects $N(p)$ then the triple of points that are determined that way are the footmarks of a unique point of $B(p, q, r)$. Using this property, we construct the set $B(p, q, r)$ as follows. See Figure 2.

We sweep the body $N(p)$ from one end to the other end with a plane E parallel to the plane $\langle p, q, r \rangle$. At each such position of E, determine the footmarks $u(E), v(E)$ and $w(E)$ on $N(p) \cap E, N(q) \cap E$ and $N(r) \cap E$, respectively.

How to determine the footmarks? We fix *any* interior point of $N(p) \cap E$, denoted $p(E)$, as a center for defining a (planar) convex distance function $\delta(E)$ having the reflection of $N(p) \cap E$ as its unit circle. Set $q(E) = p(E) + \overrightarrow{pq}$, and $r(E) = p(E) + \overrightarrow{pr}$. Clearly, $N(q) \cap E$ and $N(r) \cap E$ are the placements of $N(p) \cap E$ at $q(E)$ and $r(E)$, respectively.

Consider, under the distance function $\delta(E)$, the bisectors $B(p(E), q(E))$ and $B(q(E), r(E))$. If they intersect then denote the intersection by $X(E)$. The footmarks $u(E), v(E)$ and $w(E)$ are exactly the intersection of the segments $X(E)p(E), X(E)q(E)$ and $X(E)r(E)$ with the boundary of $N(p) \cap E, N(q) \cap E$ and $N(r) \cap E$, respectively. (The intersections do not not depend on the choice of $p(E)$.)

If M is smooth then each slice $N(p) \cap E$ is smooth too. This implies that the strips bounding the bisectors $B(p(E), q(E))$ and $B(q(E), r(E))$ cross. Therefore, $B(p(E), q(E))$ and $B(q(E), r(E))$ intersect for each E. We conclude that $B(p, q, r)$ is a single curve.

However, if M is not smooth — e.g., polyhedral bodies — then the slices $N(p) \cap E$ need not be smooth. Consequently, $B(p(E), q(E))$ and $B(q(E), r(E))$ need not intersect, i.e., point $X(E)$ need not exist. In fact, as shown in Figure 1(b), the strips bounding the bisectors $B(a, b)$ and $B(b, c)$ do not cross because the points $\sigma_{ab,1}(b)$ and $\sigma_{bc,1}(b)$ coincide. Therefore $B(a, b)$ and $B(b, c)$ do not intersect.

We need the following in the next section.

Lemma 1 *Let E_0 be a position of E so that the point $X(E_0)$ exists. Then, for any position of E sufficiently close to E_0, the point $X(E)$ exists.*

Proof: The point $X(E_0)$ exists, i.e., the planar bisectors $B(p(E_0), q(E_0))$ and $B(q(E_0), r(E_0))$ intersect. Thus, there are two crossing strips bounding these bisectors. Moreover, because each two of the points p, q and r are in general position the intersections of the curves $\Sigma_{yz}(x)$ with E, where $x, y, z \in \{p, q, r\}$, change continuously with E (in the obvious sense). Therefore, for any E sufficiently close to E_0, there are two crossing strips bounding some two (planar) bisectors. (Actually, these strips have the same "names" as the strips above.) Clearly, the point $X(E)$ exists. □

3 Main results

Let F be a 2-dimensional closed surface homeomorphic to a plane. Any three disjoint bi-infinite curves on F divide F into disjoint regions. Call any three bi-infinite curves on F to be *parallel* if the resulting decomposition of F is homeomorphic to the decomposition of the plane that results from an arrangement of 3 parallel distinct lines on it. Figure 3(a) shows non-parallel bi-infinite curves. Call a set of disjoint bi-infinite curves on F to be *star-arranged* if no three curves among them are parallel. It is not hard to see that the decompositions of F that result from any two star-arranged sets of bi-infinite curves having the same cardinality are homeomorphic.

Theorem 1 *Let p, q and r be three points in 3-space, each two of them are assumed to be in general position with respect to d_M. Then we have:*

1. *Each connected component of the intersection of $B(p, q)$ and $B(q, r)$ is a bi-infinite curve.*
2. *The bisectors $B(p, q)$ and $B(q, r)$ intersect transversely whenever they intersect at all.*
3. *On each bisector, the connected components of $B(p, q) \cap B(q, r)$ form a star-arranged set of bi-infinite curves.*

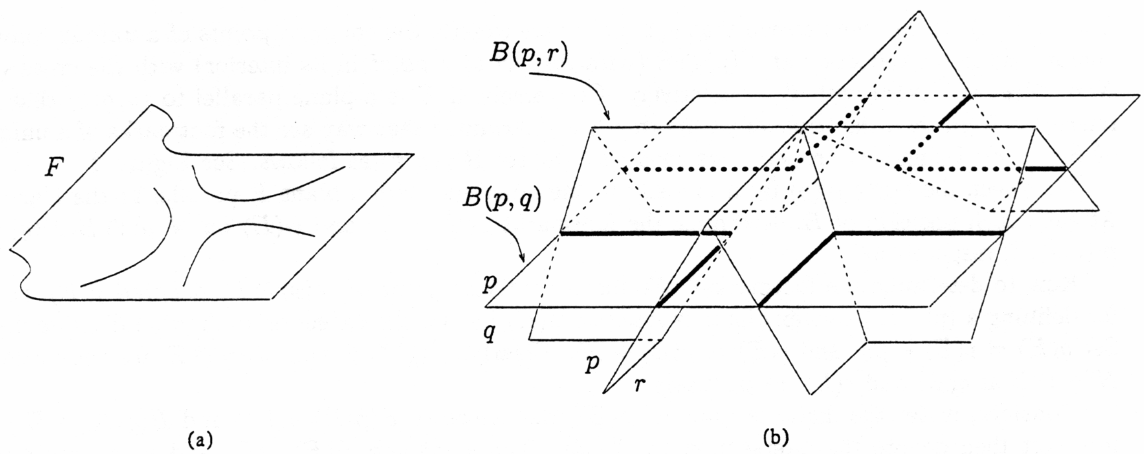

Figure 3: Star-arranged bi-infinite curves.

Proof: (Sketch) 1. Each component of $B(p,q) \cap B(q,r)$ is closed and, by Lemma 1, does not have a boundary. So it is either a bi-infinite curve or a simple closed curve. The latter case cannot occur because the footmarks of $B(p,q) \cap B(q,r)$ are unique in each sweeping plane E during the monotone sweep of the body $N(p)$.

2. This is a consequence of 1, the starshaped-ness of the Voronoi regions, and the fact that bisectors are surfaces.

3. If there are three connected component that are parallel curves then some Voronoi region is disconnected, which is impossible. (This is related to the planar non-degenerate abstract Voronoi diagram: any two bisectors of three sites intersect in at most 2 points, see [5, Lemma 3.5.2.5]). ☐

Theorem 2 *For every $n \in N$, there is a distance function and three distinct points p, q, and r so that the bisectors $B(p,q)$ and $B(q,r)$ intersect in at least n connected components.*

Proof: We will construct a (symmetric) body H_n so that the set of footmarks on $H_n(p)$ of points of $B(p,q) \cap B(q,r)$ consists of at least n connected components. Informally speaking, this can be achieved by constructing the body as a stack of convex bodies so that a sweep through the arising non-smooth body, using a plane parallel to some fixed appropriate one, alternately yields layers of smooth and non-smooth convex cross-sections. Each layer of smooth cross-sections will ensure the existence of a connected component of $B(p,q) \cap B(q,r)$, and each layer of non-smooth cross-sections ensures that its neighbor layers really yield different connected components.

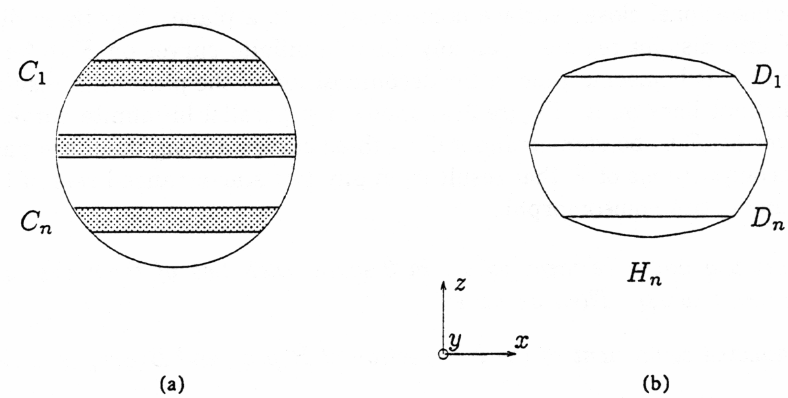

Figure 4: Constructing the body H_n. The y-axis is normal to the drawing plane and points into it.

Start with a sphere S. We cut from it n slices (the slices $C_1, \ldots, C_n$ in Figure 4(a)) all of which are bounded by discs parallel to the xy-plane and have the same height.

To construct the body H_n from the remaining pieces, we glue the two innermost ones together, and then successively glue the other, appropriately scaled up, outer pieces with the previously resulting inner body. Figure 4(b) shows the resulting body H_n.

The body H_n is not smooth exactly at the points of the circles bounding the discs along which we have glued the spherical pieces together. Denote these circles by D_i, for $i = 1, \dots, n$. To choose the points p, q, and r claimed in the theorem, we proceed as follows.

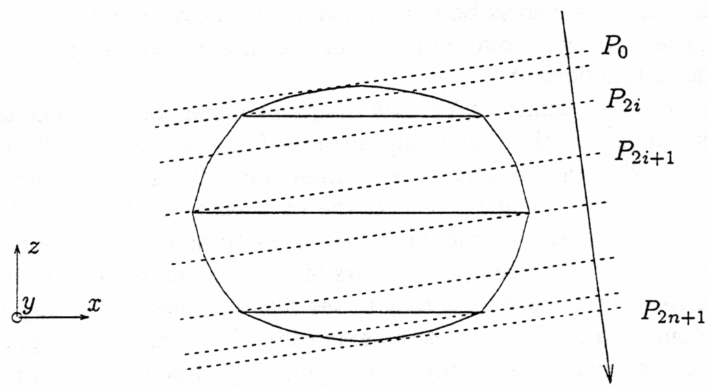

Figure 5: Critical positions of the sweep plane.

We sweep H_n with a plane that is slightly inclined to the planes of the circles D_i. See Figure 5. During the sweep, the cross sections of H_n change depending on the position of the sweep plane. Call $P_0, \dots, P_{2n+1}$ the critical positions of the sweep plane, i.e., the positions where the (names of the) spherical pieces contributing to the cross section change. Between any two such critical positions, the cross section of H_n is either a circle or the intersection of two distincts circles, refer to Figure 6(a) and (b). In the latter case, the cross section is not smooth.

We are interested in the positions of the sweep plane where the discontinuity of the cross section occurs exactly at its extreme point with respect to some fixed direction, say the y-direction, see Figure 6(b).

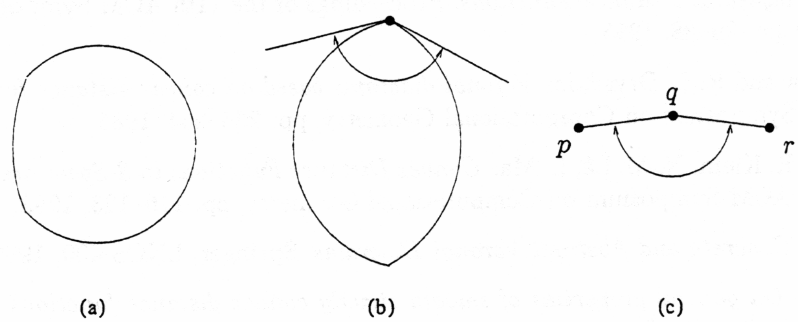

Figure 6: Cross sections determine how to choose p, q, r.

At any such position, the one-sided tangents of the cross section are not parallel to the x-axis. (As we will soon see, this is needed for choosing *non-collinear* points p, q and r). This can be seen by observing that, at that position, the centers of the intersecting circles do not lie on the segment connecting the intersection points. Denote such positions of the sweeping plane by $Q_1, \dots, Q_n$. Then we can choose, in a plane parallel to the sweeping plane, non-collinear points p, q and r so that, at any position Q_i of the sweeping plane, the supporting lines of the cross section parallel to the lines pq and qr touch it at its y-extreme point. See Figure 6(c).

Having determined the points p, q and r as above, we now prove that the intersection of the bisectors $B(p, q)$ and $B(q, r)$ satisfies the claimed properties. Recall the construction of the footmarks of the points of $B(p, q) \cap B(q, r)$ in the previous section. For any position of the sweep plane within the range bounded by the critical positions P_{2i} and P_{2i+1}, with $i = 0, \dots, n$, the cross section of $H_n(p)$ is smooth. Therefore, in each such plane there is exactly one footmark on $H_n(p)$ of a unique point of $B(p, q) \cap B(q, r)$. Thus, each layer of planes between P_{2i} and P_{2i+1} corresponds to a connected component of $B(p,q) \cap B(q,r)$. Moreover, these components are distinct because at the positions $Q_1, \ldots, Q_n$ of the sweeping plane, which separate the layers from each other, there is no footmark on $H_n(p)$ of any point of $B(p,q) \cap B(q,r)$. □

4 Discussion

We have shown that if the convex body that defines the distance function is not smooth then the bisectors of 3 points in general position may intersect in disconnected bi-infinite curves — in contrast to the case of smooth distance functions [4].

In the plane, the complexity of Voronoi diagrams under *any* distance function is linear, the constant factor does not depend on the underlying distance function. Essentially (and loosely speaking), this is due to the property that the bisectors of any three points in general position intersect in at most 1 point.

In 3-space, Boissonnat *et al* prove that the complexity of Voronoi diagrams of n points under the simplicial, L_1, and L_∞ distance function is $O(n^2)$. (Actually, they solve the problem in any dimension.)

Under the simplicial distance, the bisectors of 4 points intersect in at most 1 point, and following the argument patterns in this paper, it is easy to see that the bisectors of any 3 points intersect in at most 2 bi-infinite polygonal chains. We suggest an extension of the "abstract" approach in [9] to prove the bound $O(n^2)$ without using arguments specific to the geometry of homothetic simplexes as in [2].

More generally, we believe that the complexity bound $O(n^2)$ for Voronoi diagrams under the L_1 and L_∞ distance can be proved using only structural properties of the diagram. As we have seen, some of these structural properties can be expressed by the number of connected components of bisector intersections for 3 and 4 sites.

References

[1] F. Aurenhammer: *Voronoi Diagrams—A Survey of a Fundamental Data Structure.* ACM Computer Surveys 23(3), 1991.

[2] J.-D. Boissonnat, M. Sharir, B. Tagansky, M. Yvinec: *Voronoi Diagrams in Higher Dimensions under Certain Polyhedral Distance Functions.* Proceedings of the 11th ACM Symposium on Computational Geometry, pp. 79–88, 1995.

[3] L. P. Chew and R. L. Drysdale: *Voronoi diagrams based on convex distance functions.* Proceedings of the ACM Symposium on Computational Geometry, pp. 235–244, 1985.

[4] C. Icking, R. Klein, N.-M. Lê, L. Ma: *Convex Distance Functions in 3-Space are Different.* Proceedings of the 9th ACM Symposium on Computational Geometry, pp. 116–123, 1993.

[5] R. Klein: *Concrete and Abstract Voronoi Diagrams.* Springer, LNCS 400, 1989.

[6] N.-M. Lê: *On general properties of smooth strictly convex distance functions in $\mathbb{R}^D$.* Proceedings of the 5th Canadian Conference on Computational Geometry, pp. 375–380, 1993.

[7] N.-M. Lê: *On Voronoi diagrams in the L_p-metric in higher dimensions.* Proceedings of the 11th Symposium on Theoretical Aspects of Computer Science, P. Enjalbert *et al.* (Eds.), Springer, LNCS 775, pp. 711–722, 1994.

[8] N.-M. Lê: *On determining optimal strategies for pursuit games in the plane.* Proceedings of the 22nd International Colloquium on Automata, Languages, and Programming, Springer, LNCS 944, pp. 499–510, 1995.

[9] N.-M. Lê: *Randomized incremental construction of simple abstract Voronoi diagrams in 3-space.* Proceedings of the 10th Conference on Fundamentals of Computation Theory, Springer, LNCS 965, pp. 333–342, 1995.

[10] J. T. Schwartz and M. Sharir: *On the Two-Dimensional Davenport-Schinzel Problem.* J. Symbolic Computation (1990) 10, pp. 371–393.