Voronoi Diagrams for Direction-Sensitive Distances*

Oswin Aichholzer†
Franz Aurenhammer
Institute for Theoretical Computer Science
Graz University of Technology
Graz, Austria
e-mail: {oaich,auren}@igi.tu-graz.ac.at

Danny Z. Chen‡
Department of Computer Science & Engineering
University of Notre Dame
Notre Dame, IN 46556, U.S.A.
e-mail: chen@cse.nd.edu

D. T. Lee§
Department of Electrical and Computer Engineering
Northwestern University
Evanston, IL 60208, U.S.A.
e-mail: dtlee@ece.nwu.edu

Asish Mukhopadhyay
Department of Computer Science
Indian Institute of Technology, Kanpur
UP-208016, India
e-mail: am@iitk.ernet.in

Evanthia Papadopoulou
IBM Watson Research Center
Yorktown Heights, NY 10598, U.S.A.
e-mail: evanthia@watson.ibm.com

1 Introduction

Most computational geometry research on planar problems assumes that the underlying plane is perfectly 'flat', in the sense that movement between any two points on the plane always takes the same cost as long as the Euclidean distance between the two points is the same. In real environments, distances may depend on the direction one moves along [10], or even may be influenced by local properties of the plane [8]. These situations sometimes can be modeled by considering a piecewise-linear surface as the underlying 'plane', and measuring distances therein; see e.g. [7]. In fact, many distance problems on non-flat planes are hard to deal with from the computational geometry point of view.

We study distance problems for the basic case of a 'tilted' plane in three-space. In this model, the cost of moving depends not only on the Euclidean distance but also on how much upwards or downwards the movement has to travel, simulating the situation when driving a vehicle on the tilted plane. Direction-sensitive distances and, in particular, their induced Voronoi diagrams have not received much attention in the computational geometry literature. Exceptions are convex distance functions [2, 5] and the so-called boat-sail distance [10] which comes closest to our concept.

*Part of this work was carried out when the authors were visiting the Center for Applied Science and Engineering and the Institute of Information Science, Academia Sinica, Nankang, Taiwan, June – July 1996.

†The research of this author was supported by the Austrian Spezialforschungsbereich F003, Optimierung und Kontrolle.

‡The research of this author was supported in part by the National Science Foundation under Grant CCR-9623585.

§The research of this author was supported in part by the National Science Foundation under the Grant CCR-9309743, and by the Office of Naval Research under the Grants No. N00014-93-1-0272 and No. N00014-95-1-1007.

2 Skew distances

Let T be a tilted plane in $\mathcal{R}^3$ such that the angle between T and the xy-plane is α, with $0 \leq \alpha \leq \pi/2$. By convention, T is obtained by rotation along the x-axis. We define a coordinate system on T by taking the accordingly rotated coordinate system of the xy-plane. Based on this system, a point p on T is described by its coordinates $x(p)$ and $y(p)$. The Euclidean distance in $\mathcal{R}^3$ between $p, q \in T$ then can be expressed as $d_2(p,q) = ((x(p)-x(q))^2 + (y(p)-y(q))^2)^{1/2}$.

A simple distance function on T is obtained by taking, for points $p, q \in T$, their Euclidean distance plus their signed difference in height, that is, $d_2(p,q)+(h(q)-h(p))$. The latter term is equal to $(y(q)-y(p)) \cdot \sin\alpha$, which means that the slope of T affects the y-difference of p and q by a factor in the range $[0,1]$. To obtain a richer, but still realistic class of distance functions on T, we define the *skew distance* from p to q as

$$d(p,q) = d_2(p,q) + k \cdot (y(q) - y(p)),$$

where $k \geq 0$ is a constant. The parameter k has a nice physical interpretation. Imagine a ball moving on T, and let $\frac{1}{k}$ be the frictional coefficient on T. For $k < 1$, friction dominates gravity and the ball sticks on the skew plane. Friction and gravity balance out if $k = 1$. For $k > 1$, gravity dominates friction and the ball rolls downhill. The last case also models the assumption that energy is gained when a (electric-driven) vehicle moves downhill. Note that d is non-symmetric as $d(p,q) \neq d(q,p)$ unless $y(p) = y(q)$ or $k = 0$, but d still obeys the triangle inequality.

For $k > 1$, $d(p, q)$ may be negative. Moreover, $d(p, q)$ may *decrease* when a point q below p is moved downwards T. Let

$$L_0(p) = \{a \mid d(p, a) = 0\}.$$

Clearly, for $k < 1$, $L_0(p) = \{p\}$. For $k = 1$, $L_0(p)$ is the vertical ray emanating from p and extending below p. For $k > 1$, $L_0(p)$ is composed of two rays of slopes $\pm 1/(\sqrt{k^2 - 1})$, emanating from and extending below p. We refer to these two rays as the *0-rays* of p.

As d is non-symmetric, two different 'unit circles' with p as the fixed center can be defined, $\sigma(p) = \{a \mid d(p, a) = 1\}$ and $\sigma'(p) = \{a \mid d(a, p) = 1\}$. However, $\sigma'(p)$ is just the reflection of $\sigma(p)$ through the horizontal line passing through p. We adopt the convention of considering only the 'outgoing' skew distance in the definitions of geometric structures like $\sigma(p)$, $L_0(p)$, etc., keeping in mind that their 'incoming' versions can be obtained by reflection.

Lemma 1 *For $k > 0$, $\sigma(p)$ is a conic with focus p, directrix the horizontal line at y-distance $1/k$ above p, and eccentricity k. Thus $\sigma(p)$ is an ellipse for $0 < k < 1$, a parabola for $k = 1$, and a hyperbola for $k > 1$.*

3 Skew Voronoi diagrams and related structures

Let S be a set of n point sites on the tilted plane T. The *skew Voronoi diagram* of S, for short $SV(S)$, associates each site $p \in S$ with the region

$$reg(p) = \{a \mid d(p, a) < d(q, a), \forall q \in S\}.$$

Figure 1 depicts a skew Voronoi diagram for six sites. The dashed line segments indicate the upper envelope of their 0-rays. Note the emptiness of $reg(q)$ and $reg(t)$.

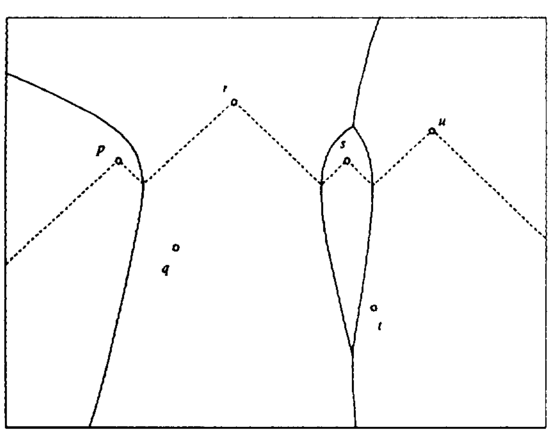

Figure 1: Skew Voronoi diagram for $k = 1.5$

Of course, $SV(S)$ depends on the 'slope parameter' k. The case $k < 1$ leads to known structures: According to Lemma 1, the skew distance d then is a *convex distance function* as defined in [2], and Voronoi diagrams for convex distance functions are well-studied objects; see e.g. [5].

Skew Voronoi diagrams have another nice and helpful interpretation for general k. Consider the *bisector* of two sites $p, q \in S$, $b(p, q) = \{a \mid d(p, a) = d(q, a)\}$. It can be described by the equation

$$d_2(p, a) - k \cdot y(p) = d_2(q, a) - k \cdot y(q).$$

Without loss of generality, let S be contained in the y-positive halfplane of T. For each site $p \in S$, define $C(p)$ as the (Euclidean) circle with center p and radius $k \cdot y(p)$. Then $b(p, q)$ is just the set of points equidistant under the Euclidean distance from the two circles $C(p)$ and $C(q)$. We hence conclude:

Theorem 3.1 *Let S be a set of point sites in the y-positive halfplane of T, and let $\mathcal{C}(S) = \{C(p) \mid p \in S\}$. Then $SV(S)$ is the Euclidean closest-point Voronoi diagram of $\mathcal{C}(S)$.*

Various properties are known for the Voronoi diagram of circles in the plane; see [3, 6, 9]. Its edges are hyperbolic arcs, its regions are star-shaped, and its size is linear. Among the available techniques for its construction, the $O(n \log n)$ time plane-sweep algorithm in [3] is reasonably simple to implement. Clearly, $\mathcal{C}(S)$ can be computed from S in $O(n)$ time. So, for any fixed $k \geq 0$, the skew Voronoi diagram can be computed in time $O(n \log n)$ and space $O(n)$.

Another kind of direction-sensitive diagram, the so-called *Voronoi diagram in a river*, has been investigated in [10]. The underlying distance function is the time required to reach a point on the river by a boat starting from a site (and fighting against a constant river flow). Despite apparent similarity, the 'boat-sail' distance is different from the skew distance. For example, if the river is slower than the boat then the river Voronoi diagram and the corresponding Euclidean Voronoi diagram are combinatorially the same. If the river is faster than the boat then all regions start at their respective sites.

Let us now give a brief catalogue of properties of $SV(S)$ in dependence of k. In fact, for $k < 1$, $SV(S)$ has a combinatorial structure identical to the Euclidean Voronoi diagram of some set S' of point sites. This can be seen by applying an affine transformation that takes S into S', and the elliptic circles defined by the skew distance into Euclidean circles[1]. Observe that the convex hulls of S and S' are combinatorially the same, so exactly the sites lying on the convex hull S have unbounded regions in $SV(S)$. Still, $SV(S)$ and the Euclidean Voronoi diagram of S are combinatorially different, in general. For example, two sites which are

[1]Thanks go to Olivier Devillers for pointing out this property.

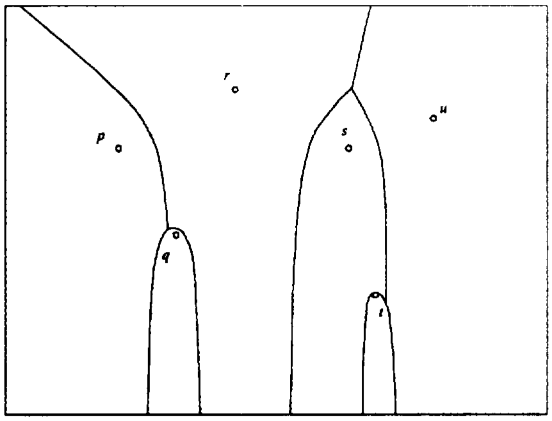

Figure 2: Skew Voronoi diagram for $k = 1$.

neighbored in the x-sorting of S but have non-adjacent Euclidean regions will have adjacent skew regions, provided k is sufficiently close to 1; see the case $k = 1$.

The case $k > 1$ is the most interesting one from the geometric point of view. Let the *0-envelope*, $E_0(S)$, be defined as the upper envelope of the 0-rays for all $p \in S$. From the triangle inequality we get $reg(p) \neq \emptyset$ if and only if $p \in E_0(S)$. As $p \in reg(p)$ for each non-empty region, $E_0(S)$ cuts each such region into an upper and a lower part. This implies that the edges of $SV(S)$ lying above and below $E_0(S)$, respectively, have the structure of a forest[2]. Using the interpretation of $SV(S)$ as a Voronoi diagram for circles, sites with unbounded regions can be characterized: A region $reg(p)$ is unbounded if and only if p lies on both $E_0(S)$ and the upper convex hull of S.

For $k = 1$, all regions of $SV(S)$ are unbounded and extend to minus infinity in y-direction. Their left-to-right order reflects the x-sorted order of S.

See Figure 1 for an example of a 0-envelope and the corresponding skew Voronoi diagram. For the same set of sites, the diagram that results from setting $k = 1$ is shown in Figure 2. Some pictures of skew Voronoi diagrams for different values of k and some animated sequences for steadily increasing k can be obtained from `http://www.cis.tu-graz.ac.at/igi/oaich/skewvd`.

To get an output-sensitive algorithm for $k > 1$, sites with empty regions can be pruned away by first constructing $E_0(S)$. By a coordinate transformation, the latter task is essentially equivalent to finding all the m maximal dominating elements among S, which can be done in $O(n \log m)$ time and $O(n)$ space [4]. This gives an algorithm with the same complexity for constructing $SV(S)$ if m regions are non-empty, which is optimal by reduction to the planar convex hull problem.

A generalization of the skew Voronoi diagram problem, still solvable in time $O(n \log n)$, is the following: Given a set S of n sites, let $S' = S$, compute $SV(S')$, and call it the *layer*-1 *diagram*, $SV_1(S)$. Next, remove from S' all sites with non-empty regions in $SV_1(S)$ (that is, all $p \in E_0(S')$), and repeat the computation on S', which next gives $SV_2(S)$, until $S' = \emptyset$. Note that the collection of 0-envelopes for the shrinking set S' can be computed in overall time $O(n \log n)$ using the max-dominance algorithm by [1].

$SV_1(S)$ and $SV_2(S)$ can be used to solve, in time $O(n \log n)$, the all nearest neighbors problem under the skew distance: For each $p \in S$, find a site $q \neq p$ such that $d(q,p)$ is minimized. Note that q may *not* be any of the Voronoi neighbors of p in $SV(S)$. For example, when $E_0(S)$ contains only one site p of S, then p has no Voronoi neighbor in $SV(S)$ but still has a closest $q \neq p$. However, it can be shown that q has a non-empty region in $SV_2(S)$ then.

[2]Thanks go to Günter Rote for pointing out the forest property.

References

[1] Atallah, M.J., Kosaraju, S.R., *An efficient algorithm for maxdominance, with applications*, Algorithmica 4 (1989), 221–236.

[2] Chew, L.P., Drysdale, R.L., *Voronoi diagrams based on convex distance functions*, Proc. 1st Ann. ACM Symp. Computational Geometry, 1985, 235–244.

[3] Fortune, S., *A sweepline algorithm for Voronoi diagrams*, Algorithmica 2 (1987), 153–174.

[4] Kirkpatrick, D.G., Seidel, R., *Output-size sensitive algorithms for finding maximal vectors*, Proc. 1st Annu. ACM Sympos. Computational Geometry, 1985, 89–96.

[5] Klein, R., *Concrete and Abstract Voronoi Diagrams*, Lecture Notes in Computer Science 400, Springer, 1989.

[6] Lee, D.T., Drysdale, R.L., *Generalization of Voronoi diagrams in the plane*, SIAM J. Computing 10 (1981), 73–87.

[7] Mitchell, J.S.B., Mount, D.M., Papadimitriou, C.H., *The discrete geodesic problem*, SIAM J. Computing 16 (1987), 647–668.

[8] Mitchell, J.S.B., Papadimitriou, C.H., *The weighted region problem: Finding shortest paths through a weighted planar subdivision*, Journal of the ACM 38 (1991), 18–73.

[9] Sharir, M., *Intersection and closest-pair problems for a set of planar discs*, SIAM J. Computing 14 (1985), 448–468.

[10] Sugihara, K., *Voronoi diagrams in a river*, Int'l. J. Computational Geometry & Applications 2 (1992), 29–48.