


Forschungsschwerpunkt S92

Industrial Geometry

<http://www.ig.jku.at>



FSP Report No. 50


Voronoi Diagrams for Oriented Spheres

F. Aurenhammer, M. Peternell,
H. Pottmann and J. Wallner

December 2006

FWF

Der Wissenschaftsfonds.



Voronoi Diagrams for Oriented Spheres*

F. Aurenhammer
J. Wallner

University of Technology Graz, Austria
auren@igi.tugraz.at
j.wallner@tugraz.at

M. Peternell
H. Pottmann

Vienna University of Technology, Austria
peternell@geometrie.tuwien.ac.at
pottmann@geometrie.tuwien.ac.at

Abstract

We consider finite sets of oriented spheres in \mathbb{R}^{k-1} and, by interpreting such spheres as points in \mathbb{R}^k , study the Voronoi diagrams they induce for several variants of distance between spheres. Our results are motivated by applications to special relativity theory.

1 INTRODUCTION

Consider the general quadratic-form distance Q for two points $p, q \in \mathbb{R}^k$, given by

$$Q(p, q) = (q - p)^T \cdot M \cdot (q - p)$$

where M is a nonsingular $k \times k$ matrix. We may assume that, without loss of generality, M is symmetric.¹ Voronoi diagrams for this kind of distance found some attention in computational geometry [8, 5, 6, 2], in particular, for special Jordan matrices $M = \text{diag}(a_i)_{a_i \in \{1, -1\}}$. For each pair of points, their separator (locus of equal distance) is a hyperplane; such diagrams are sometimes called affine diagrams [2, 4]. Clearly, the cases $M = I$ and $M = -I$ give the classical Voronoi diagram and its farthest-point variant, respectively. For $k = 2$ and $M = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$, the distance $Q(p, q)$ describes the area of the axis-parallel rectangle with diagonal pq . This diagram proved useful for finding the largest empty axis-parallel rectangle among a finite set of points [5]. A general result in [2] shows that all diagrams obtainable by Q are power diagrams [1] for a suitable set of spheres.

*Supported by the Austrian FWF Joint Research Project 'Industrial Geometry', S9200.

¹We have $Q(p, q) = \frac{1}{2}(q - p) \cdot (M + M^T) \cdot (q - p)$.

In this note, we are interested in the quasi-euclidean distance, d , generated by the matrix $M = \text{diag}(1, \dots, 1, -1)$. Our main motivation comes from special relativity theory, see e.g. [3, 10], where one describes events as points in the quasi-euclidean space whose metric is governed by the matrix M and whose isometric mappings are the Lorentz transformations. The line spanned by two events p and q can be time-like, or space-like, or light-like – these cases correspond to $d(p, q) < 0$, $d(p, q) > 0$, and $d(p, q) = 0$, respectively. The value of $d(p, q)$ is a Lorentz invariant: If positive, it has an interpretation as square of distance, and if negative, then the square root of its absolute value is the proper time experienced by a particle whose life line is the line segment pq .

For any finite collection of events we can now ask questions like the following: Can we decompose space into natural neighbourhood regions, and can we classify the single events as belonging to certain neighbourhood clusters? Does it make sense to distinguish strictly between time-like and space-like difference vectors when doing distance computations? The present paper studies several types of cell decomposition of space, together with illustrations in 2D, which are related to these questions. We consider decomposition into cells defined by the nearest neighbour property, but due to the nonpositivity of the function d , the Voronoi diagrams generated by $|d|$, $\max(d, 0)$, and related functions differ from the Voronoi diagram generated by d itself.

Let us mention that the rectangle area related distance for the matrix $M = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ considered above transforms to the two-dimensional instance of d , $M = \text{diag}(1, -1)$, by a rotation of $\frac{\pi}{4}$.




Figure 1: Separator of two oriented circles

2 STANDARD CASE

We first discuss the Voronoi diagram induced by the quasi-euclidean distance d , without any restriction on its sign. To this end, we use an interpretation of d based on the following simple but useful property: If the k^{th} coordinates of two points $p, q \in \mathbb{R}^k$ are interpreted as 'radii', and their first $k - 1$ coordinates specify 'centers', then $d(p, q)$ expresses the squared tangent length between the respective two spheres in \mathbb{R}^{k-1} . That is, the space of (oriented) spheres in \mathbb{R}^{k-1} is accordingly divided by a power diagram in \mathbb{R}^k .

To illustrate this fact, Figure 1 displays two circles σ_1 and σ_2 (shown as shaded disks) in the xy -plane, lying below their corresponding two points in \mathbb{R}^3 . Their separator with respect to the quasi-euclidean distance d is a plane, E , that intersects the xy -plane in the power line of the circles σ_1 and σ_2 . In fact, E is the power plane of the two 3-dimensional spheres S_1 and S_2 , where S_i intersects the xy -plane in σ_i at an angle of $\frac{\pi}{4}$, for $i = 1, 2$. This can be shown by simple analytic calculations. Let us call S_i the *principal sphere* of σ_i . Note that the center of S_i is the reflection through the xy -plane of the point in \mathbb{R}^3 that corresponds to σ_i . Each point $p \in E$ corresponds to a circle in the xy -plane with tangents of equal length to σ_1 and σ_2 .

Let now $\{\sigma_1, \dots, \sigma_n\}$ be a given set of oriented spheres in \mathbb{R}^{k-1} , treated as points in \mathbb{R}^k in the following. Separators are, of course, hyperplanes in \mathbb{R}^k , and each separator indeed separates its defining spheres σ_i, σ_j . However, a sphere need not be contained in its halfspace of smaller distance. For each sphere σ_i , its region $reg(\sigma_i)$ represents the set of spheres showing smallest squared tangent length to σ_i (which might be negative). The set of all spheres at distance zero from σ_i is a vertical double-cone, $cone(\sigma_i)$, with apex σ_i and aperture angle of $\frac{\pi}{2}$. By slight abuse of notation, we will call the *interior* of $cone(\sigma_i)$ the part of \mathbb{R}^k where the distance to σ_i is negative. The regions of the diagram are connected (as being convex polyhedra), but $\sigma_i \in reg(\sigma_i)$ does not hold, in general.




Figure 2: Standard distance

Figure 2 gives an example. Encirculated numbers indicate affiliation of regions to spheres.

It can be shown that $reg(\sigma_i)$ is unbounded if and only if σ_i lies on the boundary of the convex hull of (the points) $\sigma_1, \dots, \sigma_n$: Let c_i denote the center of the principal sphere S_i for σ_i ; cf. Figure 1. As c_i is the reflection of σ_i through the hyperplane $x_k = 0$ of \mathbb{R}^k , center c_i is extreme in c_1, \dots, c_n if and only if σ_i is extreme in $\sigma_1, \dots, \sigma_n$. Extremality of centers, however, is well known to characterize unboundedness of regions in the respective power diagram. Note that $reg(\sigma_i) = \emptyset$ is possible for non-extreme σ_i . The diagram, as being the power di-

agram of n principal spheres S_i , can be computed in $O(n \log n)$ time and $O(n)$ space in \mathbb{R}^2 , and in $O(n^2)$ time and space in \mathbb{R}^3 ; see [1, 8]. Both results are asymptotically optimal in the worst case.

3 VARIANT 1

A different behaviour is exhibited by a variant of the quasi-euclidean distance. We let

$$D(p, q) = |d(p, q)|.$$

Now each sphere σ_i gets assigned the region of all spheres that show squared tangent lengths of smallest absolute value to σ_i .




Figure 3: Distance D

Two types of separators arise, namely, for the equation $d(x, p) = d(x, q)$ (the hyperplanes from before), as well as for the equation $d(x, p) = -d(x, q)$ (which are vertical hyperboloids). As $D(p, q) \geq 0$, the cones of distance zero get completely swallowed by regions; more precisely, $\text{cone}(\sigma_i) \subset \text{reg}(\sigma_i)$ holds for all input spheres σ_i . In particular, we have $\sigma_i \in \text{reg}(\sigma_i)$. Figure 3 gives an illustration.

The resulting Voronoi diagram has a size of $\Theta(n^2)$ even in \mathbb{R}^2 , because the edges of the line arrangement induced by the n cones are in bijection with the interior-connected parts of the dia-

gram regions. (An exception are the $O(n)$ line arrangement edges incident to some σ_i .) Thus the combinatorial structure of the diagram is essentially determined by the line arrangement above. This leads to an $\Theta(n^2)$ construction algorithm in \mathbb{R}^2 . Note that, as the diagram is a planar graph in this case, each single region is bounded by $\Theta(n)$ edges.

4 VARIANT 2

We now restrict attention to positive tangent lengths. As one possibility, this leads us to define

$$\Delta_\infty(p, q) = \begin{cases} d(p, q) & \text{if } \geq 0 \\ \infty & \text{otherwise.} \end{cases}$$

In the resulting Voronoi diagram, spheres to which squared tangent lengths are negative are, sloppily speaking, out of the game. More precisely, for each sphere σ_i , its region $\text{reg}(\sigma_i)$ is contained in the exterior of $\text{cone}(\sigma_i)$. On the other hand, $\text{cone}(\sigma_i)$ has to be contained in $\text{reg}(\sigma_i)$, as 0 is the smallest value achieved by Δ_∞ . This implies that the double-cones show up completely in the diagram; in fact, the diagram is a refinement of the arrangement induced in \mathbb{R}^k by these n double-cones. Thus its size is $\Omega(n^2)$ already in \mathbb{R}^2 . An example is given in Figure 4.




Figure 4: Distance Δ_∞

For each sphere, its index is placed in its region. The hatched area is at distance ∞ from all three

spheres. It is the intersection of the interiors of all double-cones.

Note that diagram vertices do arise which are not vertices of the double-cone arrangement. Let Z be a cell of this arrangement. Then, in the Voronoi diagram for Δ_∞ , cell Z is divided among the $m \leq n$ spheres σ_j where Z is exterior to $\text{cone}(\sigma_j)$. Division takes place with respect to the quasi-euclidean distance d (which is non-negative in the entire cell). That is, Z houses a part of the power diagram for m spheres. In \mathbb{R}^2 , this implies an upper bound of $O(n^3)$ on the size of the diagram. Observe that a sphere may occupy part of a cell Z without contributing to the boundary of Z with its double-cone. In Figure 4, this happens for sphere σ_2 and the cell in the upper left corner. Notice further that the induced complex is not face-to-face.

5 VARIANT 3

In order to return to a more 'harmless' diagram, we consider the distance

$$\Delta_0(p, q) = \begin{cases} d(p, q) & \text{if } > 0 \\ 0 & \text{otherwise.} \end{cases}$$

Now negative squared tangent length means being as close as possible. Consequently, the diagram lives in the complement, C , of the union of the interiors of the respective double-cones. Note that C consists of two or more maximal interior-connected components. For each point x in the interior of C , the tangent length to all spheres σ_i is positive. Within C , the diagram coincides with the standard case, i.e., the power diagram from Section 2. Notice that the induced cell complex is face-to-face.

Incident to the boundary of C but exterior, there are cells of the double-cone arrangement where the distance to exactly one sphere σ_i is zero. These cells are, therefore, part of $\text{reg}(\sigma_i)$. We exclude such parts of regions from consideration, as they do not change characteristic diagram properties.

Figure 5 depicts the diagram in \mathbb{R}^2 for five spheres. Again, encircled indices attach regions to their defining spheres. Parts of regions not contained in C are lightly shaded. Their index is iden-




Figure 5: Distance Δ_0

tical to that of the adjacent part. Note that the region of sphere σ_2 is empty.

In \mathbb{R}^2 , the boundary of C consists of $O(n)$ line segments. Diagram vertices are either vertices of C (distance zero), or vertices of the standard power diagram. Hence their number is $O(n)$. The complex induced in C is polygonal. After having computed the boundary of C , as well as the power diagram, in $O(n \log n)$ time, the parts within C can be singled out in $O(n)$ time using the edge-to-edge property of the complex.

The situation remains relatively simple in \mathbb{R}^3 . To compute the boundary of C , we first treat the lower parts and the upper parts of the double-cones separately and calculate their respective unions. As all cones are vertical and have the same angle of aperture, this task is equivalent to (twice) computing an additively weighted Voronoi diagram in \mathbb{R}^2 ; see e.g. [1, 2]. Such a diagram can be computed in $O(n \log n)$ time [7]. In a next step, the standard power diagram in \mathbb{R}^3 is constructed in time $O(n^2)$. Exploiting, again, the face-to-face property, we can find the parts of the diagram lying on the appropriate sides of the cone unions' boundaries. Observe that the combinatorial complexity of the obtained Voronoi diagram for distance Δ_0 in \mathbb{R}^3 is $O(n^2)$, as each boundary component of C injectively corresponds to a face of the power diagram.

6 DISCUSSION

We have considered several variants of decomposing space based on closeness with respect to the quasi-euclidean distance d . The resulting Voronoi diagrams exhibit quite different behavior concerning combinatorial size as well as geometric shape. Variants 1 and 2 behave quadratically in size already in \mathbb{R}^2 , and curved region boundaries add to the geometric complexity of these diagrams in \mathbb{R}^3 . Variant 3, on the other hand, is (essentially) composed of linear components well treatable even in \mathbb{R}^3 . We raise the questions of whether there are other meaningful variants based on the distance d .

Our results also substantiate the role power diagrams play in physical contexts. Another recent example are the Bregman Voronoi diagrams considered in [9].

References

- [1] F. Aurenhammer. *Power diagrams: properties, algorithms, and applications*. SIAM J. Computing 16 (1987), 78-96.
- [2] F. Aurenhammer. *Geometric relations among Voronoi diagrams*. Geometriae Dedicata 27 (1988), 65-75.
- [3] W. Benz. *Classical Geometries in Modern Contexts; Geometry of Real Inner Product Spaces*. Birkhauser, Basel, Boston, Berlin, 2005.
- [4] J.-D. Boissonnat, M. Teillaud. *Effective Computational Geometry for Curves and Surfaces*. Springer Berlin, Heidelberg, 2007.
- [5] B. Chazelle, R.L.S. Drysdale, D.T. Lee. *Computing the largest empty rectangle*. SIAM J. Computing 15 (1986), 300-315.
- [6] H. Edelsbrunner, R. Seidel. *Voronoi diagrams and arrangements*. Discrete & Computational Geometry 1 (1986), 25-44.
- [7] S. Fortune. *A sweepline algorithm for Voronoi diagrams*. Algorithmica 2 (1987), 153-174.
- [8] H. Imai, M. Iri, K. Murota. *Voronoi diagram in the Laguerre geometry and its applications*. SIAM J. Computing 14 (1985), 93-105.
- [9] F. Nielsen, J.-D. Boissonnat, R. Nock. *On Bregman Voronoi diagrams*. ACM-SIAM Symp. on Discrete Algorithms, SODA 2007.
- [10] N.M.J. Woodhouse. *Special Relativity*. Springer Undergraduate Mathematics Series, 2003.