

“The Big Sweep”: On the Power of the Wavefront Approach to Voronoi Diagrams

F. Dehne¹ and R. Klein²

Abstract. We show that the wavefront approach to Voronoi diagrams (a deterministic line-sweep algorithm that does not use geometric transform) can be generalized to distance measures more general than the Euclidean metric. In fact, we provide the first worst-case optimal ($O(n \log n)$ time, $O(n)$ space) algorithm that is valid for the full class of what has been called *nice metrics* in the plane. This also solves the previously open problem of providing an $O(n \log n)$ -time plane-sweep algorithm for arbitrary L_k -metrics. Nice metrics include all convex distance functions but also distance measures like the Moscow metric, and composed metrics. The algorithm is conceptually simple, but it copes with all possible deformations of the diagram.

Key Words. Computational geometry, Delaunay triangulation, Voronoi diagram, Sweepline.

1. Introduction. Given a set S of n sites in the plane, their *Voronoi diagram* is a partition of the plane into regions, one to each site, such that the region of site p contains all points of the plane that are closer to p than to any other site in S . This structure has proven most useful in computational geometry as well as in other areas; see, e.g., [2] for a survey on variations and applications.

Four different algorithm schemes have been developed for computing the Voronoi diagram efficiently. First, a divide-and-conquer algorithm was presented in [18] that runs in optimal time $O(n \log n)$ in the worst case. Second, geometric transformations were discovered in [3] and [8] that reduce the problem to computing convex hulls in 3-space. A third worst-case optimal algorithm was proposed in [9]; after applying a transformation in the plane, a line-sweep algorithm is used. Finally, a randomized incremental construction was presented in [4] that allows the Voronoi diagram of n points to be computed in expected time $O(n \log n)$, the average being taken over the $n!$ many orders of insertion.

The concept of the Voronoi diagram and the algorithms for its construction have been generalized to different types of sites and distance measures [2], and to an abstract setting [11], [12].

In this paper we study the line-sweep approach. This paradigm can also be used for computing the Voronoi diagram of points on a cone [7]. Furthermore, it has been pointed out in [17] and [6] that the planar transform suggested in the original paper [9] is not necessary. Rather, the Voronoi diagram can be constructed directly by sweeping the plane

¹ School of Computer Science, Carleton University, Ottawa, Canada K1S 5B6. dehne@scs.carleton.ca. Research partially supported by the Natural Sciences and Engineering Research Council of Canada.

² Praktische Informatik VI, FernUniversität Hagen, Elberfelder Str. 95, 5800 Hagen, Germany. rolf.klein@fernuni-hagen.de. Research partially supported by the Deutsche Forschungsgemeinschaft, Grant No. Kl 655/2-1.

with a vertical line l from left to right, maintaining that part of the diagram to the left of l that cannot change any more as new sites are discovered during the sweep. Its right boundary is a sequence of parabola segments that looks like a *wavefront*. As a matter of fact, this is just the Voronoi diagram of the points to the left of l and of the site l itself, and the wavefront is the boundary of the region of l . A sweep algorithm without transform has been used in [19] for computing the Voronoi diagram of points with respect to the Manhattan distance L_1 and, analogously, L_∞ .

Since this wavefront algorithm is extremely natural and simple, we would like to generalize it to distance measures other than L_1 , L_2 , and L_∞ . For example, in order to plan collision-free motions of a convex robot, convex distance functions are required [5].

We show that the wavefront approach can even be generalized to the full class of *nice metrics* in the plane introduced in [14].

Roughly, a metric d is called nice if convergency of point sequences means the same in d as in the Euclidean metric, if for any two points there is a third one between them such that their mutual distances add up, and if bisectors are tractable; see Definition 1 below. This class contains not only all symmetric convex distance functions, but also distance measures like the Moscow metric, composed metrics, etc.; see [12].

So far, only for a proper subclass of nice metrics has a deterministic worst-case optimal algorithm been known [12]. Since it is of divide-and-conquer type, it works only if each point set can be partitioned into subsets of equal size whose bisector is acyclic. In practice, this condition is not easy to verify. Also, the divide-and-conquer algorithm is quite complicated in its general version. A randomized $O(n \log n)$ algorithm suitable for all nice metrics can be obtained from the results on abstract Voronoi diagrams presented in [16] and [13].

In this paper we present the first deterministic algorithm for computing the Voronoi diagram of n points in an arbitrary nice metric in the plane within $O(n \log n)$ time and linear space. This also extends the results of [19] by providing an $O(n \log n)$ -time plane-sweep algorithm for arbitrary L_k .

Among the efficient deterministic algorithms, the wavefront approach might be the easiest to implement. It gracefully deals with all kinds of ugly phenomena like multiple vertices, two-dimensional bisector pieces, one-dimensional pieces of Voronoi regions.

The paper is organized as follows. After providing the basic definitions in Section 2, we briefly review, in Section 3, how the wavefront algorithm works in the Euclidean metric. In Section 4 we study the dynamic properties of the wavefront during the sweep. Then, in Section 5, we present the general wavefront algorithm. Four different kinds of events call for an update of the wavefront; how to handle them is described in respective subsections. In the last section we mention possible generalizations and propose some problems for further research.

2. Nice Metrics and Voronoi Diagrams. Most of the material contained in this section has been presented in [12]. We include it here for the convenience of the reader.

Let d be a metric in the plane, i.e., a function that assigns to each pair of points a, b in the plane a nonnegative distance $d(a, b)$, such that $d(a, b) = 0$ if and only if $a = b$, $d(a, b) = d(b, a)$, and $d(a, c) \leq d(a, b) + d(b, c)$ hold for all a, b, c .




Fig. 1. The bisector $B(q, l)$, where $q \in l$ and $q < l$, in the Euclidean metric (L_2) and in the L_∞ -metric defined by $L_\infty(a, b) = \max(|x(a) - x(b)|, |y(a) - y(b)|)$.

Let l denote a (vertical) line. By

$$d(p, l) := \min\{z \in l; d(p, z)\}$$

we denote the distance between point p and line l . We have $d(p, l) = 0$ if and only if $p \in l$. Now let p be a point, and let r denote a point or a vertical line. Let

$$\begin{aligned} B(p, r) &= \{z; d(p, z) = d(r, z)\}, \\ D(p, r) &= \{z; d(p, z) < d(r, z)\}, \end{aligned}$$

and put $C(p, r) = D(p, r) \cup B(p, r)$. The set $B(p, r)$ is called the *bisector* of p and r (see Figure 1). It need not be a curve; in L_1 it contains quarter-planes if the points are diagonal vertices of a square. For a point p and nonnegative distance α , the d -circle for p with distance α is defined as $\{z; d(p, z) = \alpha\}$.

We consider the following class of metrics in the plane.

DEFINITION 1. A metric d on \mathbb{R}^2 is called *nice* if:

1. Each d -circle contains a standard circle, and vice versa.
2. Each d -circle is contained in a standard circle.
3. For any two points a and c there exists a point $b \notin \{a, c\}$ such that $d(a, c) = d(a, b) + d(b, c)$ holds.
4. If p, r are two points, or a point and a line, then the boundary of $B(p, r)$ consists of two curves each of which is homeomorphic to a line. The intersection of two such curves consists of finitely many connected components.

The curves referred to in property 4 of Definition 1 will be the edges of the Voronoi diagram. However, we can choose between the left and the right boundary curve of $B(p, r)$, $C(p, r) \cap \overline{D(r, p)}$ or $C(r, p) \cap \overline{D(p, r)}$.

In order to make a consistent choice, let $S = \{p_1, \dots, p_m, l\}$ be a set of m points and one vertical line, l , and let $<$ be a total order on S . By ∂M we denote the boundary of a set M .

DEFINITION 2. For sites $p, r \in S$, $p \neq r$, let

$$R(p, r) := \begin{cases} D(p, r) \cup B(p, r) & \text{if } p \prec r, \\ D(p, r) & \text{if } r \prec p. \end{cases}$$

Then

$$VR(p, S) := \bigcap_{\substack{r \in S \\ r \neq p}} R(p, r)$$

is the *Voronoi region of p with respect to S* , and

$$V(S) := \bigcup_{p \in S} \partial R(p, S)$$

is the *Voronoi diagram of S* .

Clearly, $\partial R(p, r) = \partial R(r, p)$ holds. We denote this *separating curve* by $J(p, r)$.

The Voronoi regions form a disjoint decomposition of the plane. It can be derived from property 3 of Definition 1 that for any two points, p and q , there exists a *d -straight path* from p to q , satisfying $d(a, c) = d(a, b) + d(b, c)$ for any three consecutive points a, b , and c on the path. Since the Voronoi regions $VR(p, S)$ are *d -star shaped*—each d -straight path from p to a point in $VR(p, S)$ is contained in $VR(p, S)$ —we obtain the following consequence:

LEMMA 3. For each point $p \in S$, $VR(p, S)$ is connected. The Voronoi region of line l is connected if no point in S lies on l .

The Voronoi diagram is a planar graph of linear complexity whose edges consist of pieces of bisecting curves $J(p, r)$, and whose faces are the Voronoi regions. However, the regions may contain one-dimensional pieces (cut-points whose removal leaves the region disconnected). Examples are shown in Figure 2 for the L_∞ -metric and for the




Fig. 2. Let $S = \{p, q, r\}$ and $p \prec q, r$. Then the region $VR(p, S)$ contains one-dimensional pieces. The thick half-line in the left picture and the shaded quarter-plane in the right picture consist of points equally far from p, q , and r ; they belong to the region of p .

Moscow metric, where distance is defined by minimum length paths that consist only of segments radial to the center c , and of segments of circles around c ; see [12]. If the region of site p has a cut point, v , then p must, with respect to \prec , be the minimum of all other sites whose regions are adjacent to v .

3. Review of the Wavefront Algorithm for the Euclidean Metric. Let $S = \{p_1, \dots, p_n\}$ be a set of n point sites in the plane. We want to construct the Euclidean Voronoi diagram $V(S)$. To this end we compute, for each value of t from $-\infty$ to ∞ , the Voronoi diagram $V(S_t)$, where

$$S_t = \{p \in S; x(p) < t\} \cup \{l_t\}.$$

Here l_t denotes the vertical line whose x -coordinate equals t . Though l_t works as the sweepline it is most useful to add it to the set of sites.

First, we sort the points p_i by their x -coordinates. We may without loss of generality assume that the points p_i have pairwise different coordinates $x(p_i)$.

In Figure 3 two Voronoi diagrams, $V(S_t)$ and $V(S_{t'})$, are depicted. We first discuss the situation at time t . Since none of p_1, \dots, p_5 lies on line l_t , the bisecting curves $J(p_i, l_t) = B(p_i, l_t)$ are parabolae. In this example, all of them contribute to the wavefront W_t , i.e., to the boundary of the Voronoi region $VR(l_t, S_t)$. The points on the p -side of $J(p, l_t)$ are closer to p than to l_t , so they are *a priori* closer to p than to any line $l_{t'}$, where $t < t'$,




Fig. 3. The Euclidean Voronoi diagrams $V(S_t)$ and $V(S_{t'})$.

and to any point site to the right of l_t . Consequently, as l_t moves on, the waves move on as well, whereas the Voronoi regions of $V(S_t)$ that do not contribute to the wavefront do not change anymore.

There are two possible events that call for an update of the wavefront, namely when a new wave appears in W_t , or when an old wave disappears. The first type of event is called a *site event*. In Figure 3 it occurs when the sweepline hits the point site p_6 . Shortly after, at time t' , there is a new wave formed by a segment of $J(p_6, l_{t'})$ glued onto the wave of p_4 (which now contributes two segments to the wavefront $W_{t'}$). When the sweepline hits p_6 the new wave starts out as a left half-line; see Figure 1.

Let p, q denote two point sites whose waves are adjacent in W_t . The bisector of p and q gives rise to an edge of $V(S_t)$ to the left of W_t . Its prolongation into the region of l_t is called a *spike*. In Figure 3 spikes are depicted by dashed lines. The spikes can be thought of as tracks the waves move along. A wave disappears from the wavefront once it has reached the point where its two neighboring spikes cross. This is called a *spike event*. At point v in Figure 3 a spike event could occur. Without site p_6 , the wave of p_3 would disappear, after reaching v , and the neighboring waves of p_4 and p_5 would become adjacent. However, after detecting site p_6 point v' gives rise to an earlier spike event that occurs when the wave of p_4 (together with the wave of p_3) arrives at v' .

If, at time t_f , all point sites have been detected and all pending spike events have been processed, the diagram $V(S)$ can be obtained from $V(S_{t_f})$ simply by removing the wavefront.

To implement this algorithm the segments of the wavefront can be stored in a balanced binary tree and a priority queue maintained for the site and spike events. Together with the initial sorting step, all this can be done in time $O(n \log n)$ and space $O(n)$, in the worst case.

4. Proofs of Wavefront Properties for the General Case. In the general case there are two additional types of events. Two nonintersecting waves may touch, and then intersect (*touch event*), and of two intersecting waves one may outrun the other (*pass event*). These event types are illustrated in Figure 6. They do not occur in the Euclidean metric because any two parabolae $B(p, l)$ and $B(q, l)$ intersect, if p and q are not on l . Also, we have to replace the intuitive arguments given in Section 3 by formal proofs based on the properties of nice metrics, as stated in Definition 1.

Let p_1, \dots, p_n denote the given point sites. As before, let

$$S_t = \{p \in S; x(p) < t\} \cup \{l_t\}.$$

As tie break order $<$ in Definition 2 of the Voronoi diagram we choose the order induced by the (unique) x -coordinates; thus, the line l_t is always the maximal element of all sites currently considered. It is this choice of $<$ that helps us cope with deformations.

First, we study the behavior of a single wave $J(p, l_t)$, as t grows bigger.

LEMMA 4. *For every point p , the function $f_p(t) = d(p, l_t)$, $t \geq x(p)$, is strictly increasing and continuous. The function $f_p(t)$ is unbounded, that is, $f_p(t) \rightarrow \infty$ for $t \rightarrow \infty$.*




Fig. 4. The point-to-sweepline distance increases continuously.

PROOF. Let $x(p) \leq t < t'$, and let π denote a d -straight path from p to a point $q' \in l_{t'}$ satisfying $d(p, q') = d(p, l_{t'})$; see Figure 4. The path π intersects l_t at some point q , and we obtain $d(p, l_t) \leq d(p, q) < d(p, q') = d(p, l_{t'})$. Hence, $f_p(t)$ is strictly increasing. In order to show the continuity of $f_p(t)$ we consider a d -straight path δ from p to l_t ending in r , where $d(p, r) = d(p, l_t)$. Let r' be the point on $l_{t'}$ with $y(r) = y(r')$. Then $d(p, l_t) < d(p, l_{t'}) \leq d(p, l_t) + d(r, r')$. If $|t' - t| \rightarrow 0$, then $d(r, r') \rightarrow 0$ and, hence, $|f_p(t') - f_p(t)| \rightarrow 0$. The unboundedness of $f_p(t)$ follows from the assumption that the d -circles are bounded. \square

Now we show that the waves keep moving, as the sweepline proceeds.

LEMMA 5. For any $t < t'$, and for any point p with $x(p) < t$, the bisecting curve $J(p, l_t)$ is contained in the domain $D(p, l_{t'})$. In particular, $J(p, l_t) \cap J(p, l_{t'}) = \emptyset$.

PROOF. Let $w \in J(p, l_t) \subset B(p, l_t)$, as shown in Figure 5. From Lemma 4 it follows that $d(w, l_{t'}) > d(w, l_t) = d(w, p)$. Hence, $w \in D(p, l_{t'})$. \square

As a consequence, the Voronoi regions of $V(S_t)$ can only grow bigger, as the sweepline proceeds.




Fig. 5. The waves keep moving, as the sweepline proceeds.

LEMMA 6. *For any $t < t'$, and for any point site $p \in S_t$, we have $VR(p, S_t) \subseteq VR(p, S_{t'})$; equality holds if the Voronoi region of p in $V(S_t)$ does not share an edge with the wavefront $\partial VR(l_t, S_t)$.*

PROOF. Lemma 5 implies $R(p, l_t) \subset R(p, l_{t'})$ for each point site $p \in S_t$. For each point site $q \in S_{t'} - S_t$ we have $VR(p, S_t) \subset R(p, q)$. Namely, if z belongs to the region of p in $V(S_t)$, then, in particular, $z \in R(p, l_t)$, hence $d(p, z) \leq d(z, l_t) \leq d(z, q)$; the latter could be an equality if $q \in l_t$. This implies $z \in C(p, q) = R(p, q)$, due to $p \prec q$. \square

Next, we show that there is no bound to the expansion of a wave.

LEMMA 7. *Let $p \in S_t$, and let $z \in D(l_t, p)$. Then there is a real number $t' > t$ such that z lies on or to the left of $J(p, l_{t'})$.*

PROOF. At time t we have $d(z, l_t) < d(z, p)$. Due to Lemma 4 the value of $d(z, l_t)$ is continuously increasing, as t tends to ∞ (it may be decreasing first, if z lies to the right of l_t). Thus, there is a unique t' such that $d(z, l_{t'}) = d(z, p)$, which means $z \in B(p, l_{t'}) \subset R(p, l_{t'})$. \square

DEFINITION 8. For each point z to the right of point p let $t_{\text{reach}}(p, z) = \inf\{t; z \in R(p, l_t)\}$.

To simplify the discussion we assume that the bisector $B(p, l_t)$ is a curve, i.e., that $B(p, l_t) = J(p, l_t)$ holds if $p \notin l$. This can be shown to be true for all symmetric convex distance functions. The case where $B(p, l_t)$ contains two-dimensional pieces does not cause any problems. Under this assumption, $t_{\text{reach}}(p, z)$ marks the unique time when $J(p, l_t)$ hits z .

Now we look at the possible interaction of two waves.

DEFINITION 9. Two bisecting curves, $J(p, q)$ and $J(q, r)$, are said to *cross* at point v if, in a neighborhood of v , one piece of $J(p, q)$ is a Voronoi edge that separates the regions of p and q in the Voronoi diagram $V(\{p, q, r\})$, and the other piece of $J(p, q)$ is not.

This definition is symmetric in $J(p, q)$ and $J(q, r)$.

Two bisectors $J(p, l_t)$ and $J(q, l_t)$ can cross at most twice, or the Voronoi diagram of $\{p, q, l_t\}$ would have a disconnected Voronoi region, contradicting Lemma 3. It is easy to distinguish the two vertices that two bisectors represented in the wavefront may have in common. Namely, the cyclic sequences of Voronoi regions in counterclockwise order around them are different; see, for example, the waves of p_4 and p_6 in Figure 3.

DEFINITION 10. For $p, q \in S_t$ let

$$\begin{aligned} t_{\text{start}}(t', p, q) &= \inf\{t \geq t'; J(p, l_t) \text{ crosses } J(q, l_t) \text{ with region order } (p, q, l_t)\}, \\ t_{\text{stop}}(t', p, q) &= \sup\{t \geq t'; J(p, l_t) \text{ crosses } J(q, l_t) \text{ with region order } (p, q, l_t)\}. \end{aligned}$$




Fig. 6. At time t_1 a touch event occurs. At time t_2 , the q -wave outruns the p -wave, giving rise to a pass event.

If $J(p, l_t)$ is strictly above $J(q, l_t)$ then $t_{\text{start}}(t', p, q)$ marks the time when the two curves touch, as depicted in Figure 6, if they ever do. Otherwise, $t_{\text{start}}(t', p, q) = \infty$. Once two such bisectors have started to intersect, they can only get disentangled if one of them passes the other, because they never recede, due to Lemma 5. In Figure 6 this happens at time $t_2 = t_{\text{stop}}(t_1, p, q)$.

Next, we look at the wavefront $W_t = \partial VR(l_t, S_t)$ as a whole. Since l_t is maximal with respect to \prec , its Voronoi region does not contain cutpoints, according to Definition 2. The wavefront can consist of finitely many disconnected pieces that are separated by parts of $VR(l_t, S_t)$ extending to infinity. Each wavefront segment consists of finitely many waves, some of which may have degenerated into points. Conceptually, we assume that the “essential” part of the diagram is encircled by a closed curve Γ consisting of a segment of l_t and a \subset -shaped segment to the left, so large that only semi-infinite bisectors are outside of Γ , which either coincide or stay disjoint. Each of the wavefront segments hits Γ at two points, thereby introducing a top-down order among these segments, just as if they were connected.

The right drawing of Figure 3 shows that the same site may contribute more than one wave to the wavefront.

LEMMA 11. *At each time t , the number of waves in W_t is $O(n)$.*

PROOF. Since any two bisecting curves can cross at most twice, the assertion follows from the fact that $\lambda_2(n) = O(n)$; see [1]. □

As in Section 3 we denote the part in $VR(l_t, S_t)$ of the curve bisecting the sites of two neighboring waves of W_t a *spike*. It is easy to see that the two spikes of a p -wave in W_t can cross at most once, and that they do not intersect at all if they belong to the same bisecting curve.

5. The General Wavefront Algorithm. During the sweep, we maintain the combinatorial structure of the wavefront $W_t = \partial VR(l_t, S_t)$, i.e., the sequence of boundary edges of $VR(l_t, S_t)$ in top-down order, and the event queue Q_t . In the latter, future events of four types are stored, together with the time when they will occur:

- **Site events.** For each point site p to the right of l_t the time $x(p)$.
- **Spike events.** For each pair of spikes of a p -wave in W_t that cross at point v the time $t_{\text{reach}}(p, v)$; see Definition 8.
- **Touch events.** For each pair of disjoint segments of W_t the time $t_{\text{start}}(t, p, q)$, if less than ∞ , if the lowest wave of the upper segment is a p -wave and the uppermost wave of the lower segment is a q -wave; see Definition 10.
- **Pass events.** For each end of a segment of W_t the time $t_{\text{stop}}(t, p, q)$, if less than ∞ , if the last wave in the segment is a p -wave, its predecessor a q -wave, and if the p -wave is above the q -wave.

We assume that events scheduled for the same time are sorted in such a way that spike events come first, next pass events, then touch events, and finally site events.

The correctness of the wavefront approach is due to the following:

LEMMA 12. *The wavefront can only change its structure when one of the above events occurs.*

PROOF. Suppose that no event occurs within the time interval (t', t'') . Then disjoint wavefront segments remain disjoint, because there is no touch event. If a wave outruns its neighbor, the latter must be situated at the end of a wavefront segment (otherwise there would be a spike event before), but such pass events do not occur in (t', t'') , by assumption.

Therefore, the waves run along their spikes. Since the spikes do not cross it follows that no wave can disappear from $W_{t'}$.

Suppose that at time $t \in (t', t'')$ a new wave of site p appears. Then p belongs to $S_{t'}$, and has, due to Lemma 6, already contributed one or several waves to $W_{t'}$. None of them has yet disappeared. For each p -wave in W_t we consider a d -straight arc to p ; since it is contained in the region of p , it must pass through a p -wave of $W_{t'}$. Since the latter contains fewer p -waves than W_t , there must be two paths leading through the same p -wave of $W_{t'}$; see Figure 7. Each wave of W_t between the p -waves these paths start from is separated from its site—a contradiction. Therefore, the sequence of waves in the wavefront does not change before time t'' . \square

If the two spikes of a p -wave cross at point v , then the p -wave reaches v at time $t_{\text{reach}}(p, v)$ and not before, by definition. However, some other part of the wavefront could reach v at an earlier time.

LEMMA 13. *Assume that the first event in Q_t is a spike event, and let v be the cross-point associated with it. Then v lies in front of W_t , i.e., in $VR(l_t, S_t)$.*

PROOF. Suppose the spike event is scheduled for time $t' = t_{\text{reach}}(p, v)$. If some piece of the wavefront reaches v before time t' , then it is bound to hit the p -wave head-on




Fig. 7. The wave of q cannot be connected to q by a path contained in the region of q .

before the latter arrives at v . However, then there must be a spike event before time t' —a contradiction. \square

Next, we describe how to update W_t and Q_t on processing an event.

5.1. *Spike Events.* When a spike event occurs we delete the corresponding wave from the wavefront. If any of the two spikes involved has had a later cross-point with its other neighbor, this spike event is deleted from Q_t . For example, in Figure 3 we would at time $t_{p_4}(v')$ delete the event associated with v . Finally, we form the spike of the two newly adjacent waves and compute the cross-points with its neighbors. The corresponding spike events are inserted into Q_t . They could occur at time t , too, but they would be processed before the sweepline moves on.

Multiple spike events (leading to Voronoi vertices of degree larger than three) are dealt with like simple ones. If we have a sequence of spikes crossing at the same point, v , then all the associated waves arrive at v at the same time. Within this sequence, neighboring pairs of spikes can be processed in any order.

5.2. *Touch and Pass Events.* When two formerly disjoint segments of W_t become united we have to update the sequence of waves, because the piece of the encircling curve Γ that has separated the two segments disappears. A new spike appears between the newly touching waves. We compute the cross-points with its neighbors, and insert any resulting spike event into the queue Q_t .

Similarly, if a wave at the end of a wavefront segment is outrun by its neighbor, we delete it from W_t , and remove from Q_t the spike event possibly caused by the spike between these two waves.

5.3. *Site Events.* When the sweepline hits a new site, q , at time t , we insert a new wave into the wavefront W_t . Before that, we have processed all other events of time t that were stored in the queue.

From the examples depicted in Figure 1 we know that the new wave $B(q, l_t)$ can be a curve through q that is still folded, like the left half-line in L_2 , or one that has already begun to open up, like the contour of the left quarter-plane in L_∞ . We treat the first situation as a special case of the second. Thus, for each of the two arcs of $B(q, l_t)$ we

have to find the first point where it crosses the wavefront. This is greatly facilitated by the following observation.

LEMMA 14. *Let A be an arc of $B(q, l_t)$, where $q \in l_t$. Then there is at most one cross-point of A with W_t , namely, the first point on A that belongs to W_t .*

PROOF. The first point w of W_t on A must be a cross-point. Namely, if w belongs to the wave of site $p \in S_t$, then $p \prec q \prec l_t$ holds, by definition of \prec . However, then A stops being a (q, l_t) -Voronoi edge at point w , even if it only touches the wavefront. In fact, we have $w \in B(p, l_t) \cap B(q, l_t) \subset B(p, q) \subset R(p, q)$, so w belongs to the region of p , and not q , in $V(\{p, q, l_t\})$. Moreover, once A has touched the wavefront it cannot return into the region of l_t ; either the region of q or the region of l_t would not be connected. This shows that there can be only one cross-point. \square

Thanks to Lemma 14 we can locate each of the two cross-points of $B(q, l_t)$ with W_t by a *binary search* on the ordered sequence of waves in W_t . We start with the wave s in the middle of W_t and test if arc A has a cross-point with s . If not, we check whether s lies above or below A , to direct the further search. Note that this search works correctly even if W_t is not y -monotone.

Once both cross-points have been found, the new wave is inserted into the wavefront.

LEMMA 15. *The waves of W_t that are covered by the new wave $B(q, l_t)$ now become Voronoi edges separating the regions of their point site from q ; see Figure 8.*

PROOF. For the two points marked x and z we have $d(r, x) \leq d(x, l_t) \leq d(x, q)$, hence $x \in R(r, q)$ because of $r \prec q$. Also, we have $d(z, q) = d(z, l_t) < d(r, z)$, which implies $z \in R(q, r)$ for each $r \in S_t$. \square

After inserting the new wave we check its two spikes for cross-points with their neighbors, and insert the resulting spike events into the event queue. Before that, we remove from Q_t all spike events involving spikes that are covered by the new wave.




Fig. 8. The part of W_t covered by the new q -wave belongs to $V(S)$.

Note that, for L_2 we need to consider only spike events and site events, and Lemmas 14 and 15 are obvious for this special case.

THEOREM 16. *The Voronoi diagram of n points based on a nice metric in the plane can be computed by the wavefront algorithm in optimal time $O(n \log n)$, using linear space.*

PROOF. Only the performance bounds need proof. Due to Lemma 11, the update operations on W_t and Q_t can be carried out in time $O(\log n)$ per event, and linear space is sufficient to hold these structures. Clearly, we have n site events and $O(n)$ spike events, each giving rise to a Voronoi vertex. Each touch or pass event results in an unbounded Voronoi edge. Hence there are $O(n)$ events altogether. \square

Here we assume that $O(1)$ implementations of the following *elementary operations* are available. To find out if and where two neighboring p -spikes cross, and to test if and where a segment c of a bisector $B(p, l)$ is crossed by an arc A of $B(q, l)$ starting from $q \in l$, or whether c lies above or below A . Finally, to compute the functions $t_{\text{reach}}(p, z)$, $t_{\text{start}}(t', p, q)$, and $t_{\text{stop}}(t', p, q)$.

6. Conclusion. We have shown that the wavefront approach to computing the Voronoi diagram is very natural, that it applies to a variety of interesting metrics, and that it can easily cope with all kinds of degeneracies. These properties should make it a tool well suited for practical applications.

An obvious question is if the wavefront algorithm can handle even more general situations than point sites in nice metrics. For example, as long as there is a substitute for d -straight paths that connect each point to its site, a further generalization seems possible. Another open problem is whether the approach can also be applied to general (not necessarily symmetric) convex distance functions. Second, sites other than points should be considered. We expect that without major modifications the Voronoi diagram of n line segments can be computed, as is the case for Fortune’s approach [9] that uses a geometric transform.

Also, it would be possible to use curves different from a vertical line for the sweep. For example, an expanding circle would allow us to compute the Voronoi diagram of a large set of points *locally*, if the sites are given in increasing distance from the query point.

The existing general Voronoi diagram algorithms make use of the fact that the bisector of two sites is homeomorphic to a curve, and not to a circle. However, this condition is violated, e.g., if the sites are disjoint convex curve segments, or for point sites on the surface of a cone [7]. We think it is one of the major open problems to invent a general algorithm that can deal with this case, too.

References

- [1] M. Atallah. Dynamic computational geometry. *Comput. Math. Appl.* **11**, 1171–1181, 1985.
- [2] F. Aurenhammer. Voronoi diagrams—a survey of a fundamental data structure. *ACM Comput. Surveys* **23**(3), 345–405, 1991.

- [3] K. Q. Brown. Voronoi diagrams from convex hulls. *Inform. Process. Lett.* **9**(5), 223–228, 1979.
- [4] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete. Comput. Geom.* **4**, 387–421, 1989.
- [5] L. P. Chew and R. L. Drysdale III. Voronoi diagrams based on convex distance functions. *Proceedings of the 1st ACM Symposium on Computational Geometry*, 1985, 235–244.
- [6] R. Cole. Reported by C. Ó'Dúnlaing, 1989.
- [7] F. Dehne and R. Klein. A sweepcircle algorithm for Voronoi diagrams. In H. Güttler and H. J. Schneider, editors, *Graphtheoretic Concepts in Computer Science (WG '87)*, pp. 59–70, Staffelstein. LNCS 314, Springer-Verlag, Berlin, 1988.
- [8] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.* **1**, 25–44, 1986.
- [9] S. Fortune. A sweepline algorithm for Voronoi diagrams. *Algorithmica* **2**(2), 153–174, 1987.
- [10] Ch. Icking, R. Klein, N.-M. Le, and L. Ma. Convex distance functions in 3D are different. *Proceedings of the 9th ACM Symposium on Computational Geometry*, 1993, pp. 116–123.
- [11] R. Klein. Abstract Voronoi diagrams and their applications. In H. Noltemeier, editor, *Computational Geometry and Its Applications (CG '88)*, pp. 148–157, Würzburg. LNCS 333, Springer-Verlag, Berlin, 1988.
- [12] R. Klein. *Concrete and Abstract Voronoi diagrams*. LNCS 400, Springer-Verlag, Berlin, 1989.
- [13] R. Klein, K. Mehlhorn, and St. Meiser. On the construction of abstract Voronoi diagrams, II. In T. Asano, T. Ibaraki, H. Imai, and T. Nishizeki, editors, *Algorithms (SIGAL '90)*, pp. 138–154, Tokyo. LNCS 450, Springer-Verlag, Berlin, 1990.
- [14] R. Klein and D. Wood. Voronoi diagrams based on general metrics in the plane. *Proceedings of the 5th Annual Symposium on Theoretical Aspects of Computer Science (STACS '88)*, pp. 281–291, Bordeaux. LNCS 294, Springer-Verlag, Berlin, 1988.
- [15] M. L. Mazón and T. Recio. Voronoi diagrams based on strictly convex distances on the plane. Manuscript, Departamento De Matemáticas, Universidad de Cantabria, Santander, 1991.
- [16] K. Mehlhorn, St. Meiser, and C. Ó'Dúnlaing. On the construction of abstract Voronoi diagrams. *Discrete Comput. Geom.* **6**, 211–224, 1991.
- [17] R. Seidel. Constrained Delaunay Triangulations and Voronoi Diagrams with Obstacles. Technical Report 260, IIG-TU Graz, pages 178–191, 1988.
- [18] M. I. Shamos and D. Hoey. Closest-point problems. *Proceedings of the 16th IEEE Symposium on Foundations of Computer Science*, 1975, pp. 151–162.
- [19] G. M. Shute, L. L. Deneen, and C. D. Thomborson. An $O(n \log n)$ plane-sweep algorithm for L_1 and L_∞ Delaunay triangulations, *Algorithmica* **6**, 207–221, 1991.