

On the Construction of Abstract Voronoi Diagrams, II¹

by

R. Klein²

K. Mehlhorn³

S. Meiser³

A 03/89

Abstract: Abstract Voronoi Diagrams are defined by a system of bisecting curves in the plane, rather than by the concept of distance [K88a,b]. Mehlhorn, Meiser, Ó' Dúnlain [MMO] showed how to construct such diagrams in time $O(n \log n)$ by a randomized algorithm if the bisecting curves are in general position. In this paper we drop the general position assumption. Moreover, we show that the only geometric operation in the algorithm is the construction of a Voronoi diagram for five sites. Using this operation, abstract Voronoi diagrams can be constructed in a purely combinatorial manner. This has the following advantages: On the one hand, the construction of a five-site-diagram is the only operation depending on the particular type of bisecting curves and we can therefore apply the algorithm to all concrete diagrams by simply replacing this operation. On the other hand, this is the only operation computing intersection points; thus, problems arising from instable numerical computations can occur only there.

Key words: Voronoi diagrams, randomized algorithms

¹ This work was supported partially by the DFG, grant SPP Me 620/6, and partially by the ESPRIT II Basic Research Actions Program of the EC under contract No. 3075 (project ALCOM)

² FB Mathematik/Praktische softwareorientierte Informatik, Universität-GHS-Essen, D-4300 Essen 1, West Germany

³ FB 14 Informatik, Universität des Saarlandes, D-6600 Saarbrücken, West Germany

I. Introduction

The Voronoi diagram of a set of sites S in the plane partitions the plane into regions, called Voronoi regions, one for each site. The region of site p contains all points of the plane that are closer to p than to any other site in S .

Such partitions play an important role in different areas of science. In computer science, the Voronoi diagram belongs to the most useful data structures. Its structural properties depend on the underlying distance measure and the type of sites considered. Surveys on the variety of Voronoi diagrams that have been heretofore investigated in publications are given in [LS] and in [A].

In order to provide a unifying concept for the study of Voronoi diagrams, *abstract Voronoi diagrams* were introduced [K88a,b]. They are not based on the notion of distance but on systems of bisecting curves as primary objects. For any two sites, p and q , in S , let $J(p, q)$ denote a curve that is homeomorphic to the line and divides the plane into a p -region and a q -region. Then the Voronoi region of p with respect to S is the intersection of all p -regions, as q varies in $S - \{p\}$. The family of curves $J(p, q)$, $p \neq q$, $p, q \in S$, is called *admissible* if for each subset of sites $S' \subseteq S$ the Voronoi regions w.r.t. S' are path-connected and form a partition of the plane.

The algorithm of [MMO] is an instance of randomized incremental construction [CS], i.e., the Voronoi diagram is constructed incrementally by adding the sites one by one in random order. It constructs abstract Voronoi diagrams for a restricted class of bisecting curves, namely curves in general position. In particular, two different bisecting curves intersecting at a common point x must cross at x and no four curves may intersect in a common point. The first restriction forbids, for example, diagrams of point sites under the L_1 -metric (Manhattan-metric), as Figure 1 shows, cf. [L].




Figure 1.

In this paper, we drop the general position assumption. We only forbid some very special arrangements of bisecting curves. These forbidden arrangements do not occur for bisecting curves induced by *nice* metrics, including all norms, cf. [K88c]. The shapes of resulting Voronoi diagrams are described in section II. We show that Voronoi regions are homeomorphic to discs and that vertices and edges of the diagram can be described by only 3 and 4 sites, respectively. Next we show that when adding a new site to the diagram, all information necessary to compute the region of the new site can be obtained from the diagrams of the four characteristic sites of each intersected edge. This implies that the only non-combinatorial operation in the algorithm is the construction of a Voronoi diagram for 5 sites. All numerical operations take place in this particular operation. The algorithm of Guibas/Stolfi [GS] for Euclidean diagrams also has this property but neither the Plane-Sweep- nor the Divide & Conquer-algorithm do. The latter algorithms need to sort the sites

by x -coordinates. Moreover, the Plane-Sweep-algorithm sorts the computed events by x -coordinates; the Divide & Conquer-algorithm sorts the nodes of the diagram by y -coordinates in its merge step. In both cases, objects are compared to each other that are not at all related in their topology. Therefore, it may be difficult to make geometric decisions in a consistent manner.

The analysis of our algorithm makes a novel use of perturbation. Clarkson and Shor's analysis [CS] of randomized incremental construction relies on the input being in general position. Although Klein [K89] has shown how to perturb abstract Voronoi diagrams so as to generate bisecting curves in general position, this perturbation is much too costly to be used in an algorithm. We use it in the analysis by showing that the running time on the perturbed system dominates the running time on the original system and applying [CS] to bound the running time on the perturbed system. Thus perturbation is only used in the analysis; the algorithm itself does not know about it and works on the original system.

Throughout the paper, we use the following notation: For a subset $X \subseteq \mathbb{R}^2$ the closure, boundary and interior of X are denoted by \overline{X} , $bd\ X$ and $int\ X$, respectively.

II. Abstract Voronoi Diagrams

Let $n \in \mathbb{N}$, and for every pair of integers p, q such that $1 \leq p \neq q \leq n$ let $D(p, q)$ be either empty or an open unbounded subset of \mathbb{R}^2 and let $J(p, q)$ be the boundary of $D(p, q)$. We postulate:

- 1) $J(p, q) = J(q, p)$ and for each p, q such that $p \neq q$ the regions $D(p, q)$, $J(p, q)$ and $D(q, p)$ form a partition of \mathbb{R}^2 into three disjoint sets.
- 2) If $\emptyset \neq D(p, q) \neq \mathbb{R}^2$ then $J(p, q)$ is homeomorphic to the open interval $(0, 1)$.

We call $J(p, q)$ the bisecting curve for sites p and q . The abstract Voronoi diagram is now defined as follows:

Definition 1: Let $S = \{1, \dots, n-1\}$ and

$$R(p, q) := \begin{cases} D(p, q) \cup J(p, q) & \text{if } p < q \\ D(p, q) & \text{if } p > q \end{cases}$$

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

$$V(S) := \bigcup_{p \in S} bd\ VR(p, S)$$

$VR(p, S)$ is called the Voronoi region of p w.r.t. S and $V(S)$ is called the Voronoi diagram of S . \square

We postulate that the Voronoi regions and the bisecting curves satisfy the following three conditions:

- 1) Any two bisecting curves intersect in only a finite number of connected components.
- 2) For any non-empty subset S' of S
 - A) if $VR(p, S')$ is non-empty then $VR(p, S')$ is path-connected and has non-empty interior for each $p \in S'$,
 - B) $\mathbb{R}^2 = \bigcup_{p \in S'} VR(p, S')$ (disjoint)
- 3) For any two bisecting curves $J(p, q)$ and $J(p, r)$ and for all points $v \in J(p, q) \cap J(p, r)$, if for any neighborhood $U(v)$ of v
 - a) $J(p, r)|_{U(v)} \cap D(q, p)|_{U(v)} = \emptyset$ then $D(p, r)|_{U(v)} \subseteq D(p, q)|_{U(v)}$
 - or
 - b) $J(p, r)|_{U(v)} \cap D(p, q)|_{U(v)} = \emptyset$ then $D(r, p)|_{U(v)} \subseteq D(q, p)|_{U(v)}$.

Informally speaking, Condition 3 says that any two bisecting curves $J(p, q)$ and $J(p, r)$ intersecting in a point v without crossing, i.e., only touching at v , must have the p -region on the same side. Note that both bisecting curves belong to the same site p .

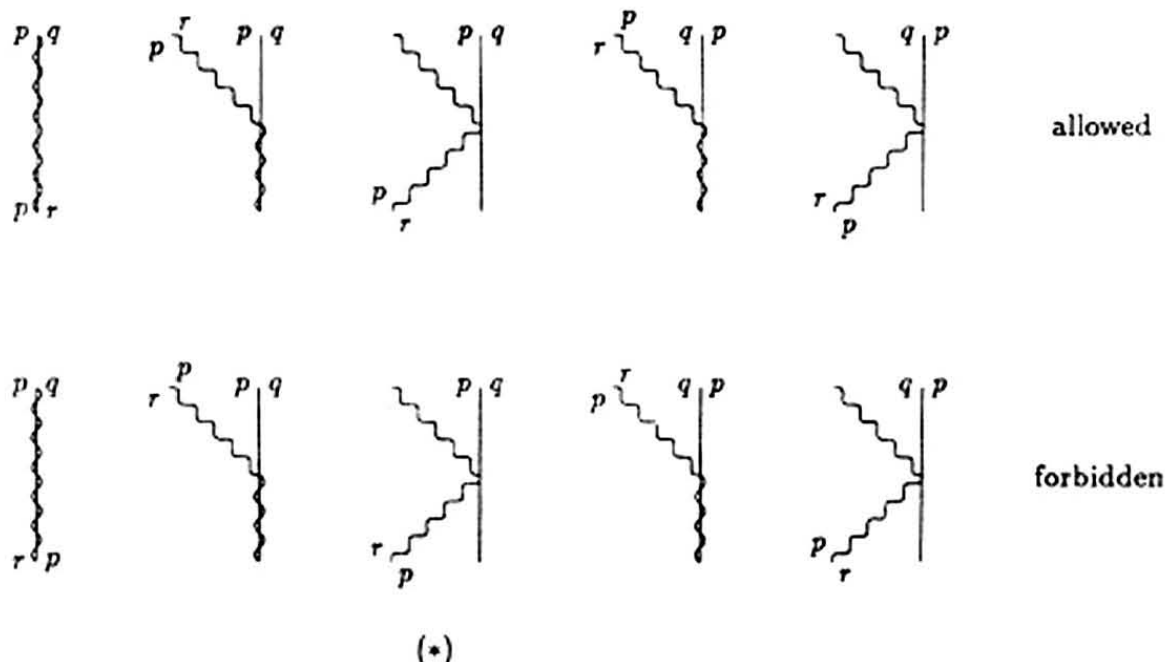


Figure 2. The upper row shows situations allowed by Condition 3, the lower row shows the corresponding forbidden situations

Condition 3 excludes the degenerate shapes of Voronoi regions shown in Figure 3. As a




Figure 3.

consequence of Theorem 2.3.5 [K88a] the point v in Figure 3 (i) must belong to the region of site p and both p -sectors are connected via v . This implies that these sectors cannot be separated by $J(p, q)$ and $J(p, r)$. Since both bisecting curves must pass through v , they are arranged as in Figure 2 (*). This is a contradiction to Condition 3. Situation (ii) in Figure 3 is impossible for the same reason; on both sides of the thin part of p -region there must be regions of different sites r, q (Theorem 2.3.5, [K88a]). In general we have:

Fact 1: The following holds for each point $v \in V(S)$: Arbitrarily small neighborhoods $U(v)$ of v exist having the following properties. Let $VR(p_1, S), VR(p_2, S), \dots, VR(p_k, S)$ be the sequence of Voronoi regions traversed on a counterclockwise march around the boundary of $U(v)$ and let I_1, I_2, \dots, I_k denote the corresponding intervals of $\partial U(v)$, where $I_j = (w_j, w_{j+1}) \subseteq VR(p_j, S)$ for $1 \leq j \leq k$ (indices must be read mod k). The intervals may be open, half-open or closed. We have $w_j \neq w_{j+1}$ for $1 \leq i \leq k$. The common boundary of $VR(p_{j-1}, S)$ and $VR(p_j, S)$ defines a curve segment $\beta_j \subseteq J(p_{j-1}, p_j)$ connecting v and w_j . $V(S) \cap U(v)$ is the union of the curve segments β_j together with the point v . Each β_j is contained in the Voronoi region of $\min\{p_{j-1}, p_j\}$. The open "piece of pie" bordered by β_j, β_{j+1} and I_j is non-empty and belongs to $VR(p_j, S)$. The point v belongs to the region of $\min\{p_1, \dots, p_k\}$. Finally, $p_i \neq p_j$ for $i \neq j$. For a proof see Theorem 2.3.5 [K88a] and the discussion above. \square

For the sequel, it is helpful to restrict attention to the "finite part" of $V(S)$. Let Γ be a simple closed curve such that all intersections between bisecting curves lie in the inner domain of Γ . We add a site ∞ to S , define $J(p, \infty) = J(\infty, p) = \Gamma$ for all p , $1 \leq p < n$, and $D(\infty, p)$ to be the outer domain of Γ for each p , $1 \leq p < n$.

Fact 2: Each Voronoi region is simply-connected and its boundary is a simple closed curve. Moreover, the closure of each non-empty Voronoi region $VR(p, S)$, $p \neq \infty$, is homeomorphic to a closed disc. A Voronoi diagram can be represented as a planar graph in a natural way. The vertices of the graph are the points of $V(S)$ which belong to the boundary of three or more Voronoi regions; the edges of the graph correspond to the maximal connected subsets of $V(S)$ belonging to the boundary of exactly two Voronoi regions. The faces of the graph correspond to the non-empty Voronoi regions. We use $V(S)$ to also denote this graph. For a proof see Lemma 2.2.4 and Theorem 2.3.5 [K88a]. \square

Condition 3 puts no restriction on the bisecting curves for *nice* metrics. Nice metrics are subject to four simple axioms, cf. [K88c], and include all norms in the plane. With nice metrics the connectivity of Voronoi regions is independent of the ordering of the sites, cf. Theorem 1.2.13, [K88a], and therefore the situation of Figure 4 cannot arise. In this situation,




Figure 4.

p -region would not be connected in $V(\{p, q, r\})$ if the sites are ordered $r < p < q$. Segment s would belong to r -region.

R. Klein showed recently [K89] that the following conditions are equivalent to our conditions (2A) and (2B). These conditions can be tested more easily.

Lemma 1: In any system of bisecting curves Conditions (2A) and (2B) are fulfilled iff the following assertions are fulfilled.

- 1) $R(p, q) \cap R(q, r) \subseteq R(p, r)$ holds for any three sites $p, q, r \in S$. (Transitivity)
- 2) Let $J(p, q)$ and $J(p, r)$ be such oriented that the p -region lies on their left hand side. Then the two curves cross at most twice and do not constitute a clockwise cycle in the plane. \square

We next take a closer look at the elements of $V(S) = \bigcup_{r \in S} \text{bd } VR(r, S)$. Let x be a point of $V(S)$. Since Voronoi regions partition the plane, x lies on the boundary of at least two Voronoi regions.

Case 1: x lies on the boundary of exactly two Voronoi regions, say p -region and q -region.

By Fact 1, Figure 5 shows $V(S)$ in a sufficient small neighborhood $U(x)$ of x . Thus x is a




Figure 5.

point on an edge of the Voronoi diagram. The bisecting curve defining the edge is determined by the sites p and q . Note that the endpoints of an edge do not belong to the edge.

Case 2: x lies on the boundary of $k \geq 3$ Voronoi regions.

In this case a neighborhood of x consists of k sectors belonging to different Voronoi regions, and x is a common endpoint of k edges of $V(S)$, i.e., a vertex of $V(S)$. Let $p, q, r_1, r_2, \dots, r_{k-2}$




Figure 6.

be the Voronoi regions involved, enumerated in counterclockwise order. Then, in particular, x is the endpoint of an edge e separating p -region from q -region.

The following lemma shows that every site r_1, \dots, r_{k-2} takes equal responsibility in the definition of x as an endpoint of e .

Lemma 2: Let $L \subseteq S$ and $\{p, q, r_i\} \subseteq L$ for some i , $1 \leq i \leq k-2$. Then x is also an endpoint of e in $V(L)$.

Proof: The regions of p , q , r_i cannot become smaller in $V(L)$ compared to $V(S)$. Thus in a neighborhood of x in $V(L)$, at least the regions of p , q and r_i are represented. Since each region can be represented only by one sector (Fact 1), the edge e must end in the point x in $V(L)$ as well. \square

Besides x the site r_i can define at most one more endpoint x' of an edge separating the p - and q -regions, because the Voronoi diagram of the three sites p , q and r_i can have at most two vertices in its "finite part", cf. Figure 7 (for a proof see Lemma 3.5.2.5, [K88a]). The points x and x' must necessarily be endpoints of the same edge. We can distinguish




Figure 7.

between x and x' by looking at the counterclockwise orders p, r_i, q resp. p, q, r_i of regions around the vertex.

Lemma 3: In $V(S)$ at most one vertex can exist with the regions of p , q and r appearing in counterclockwise order in its neighborhood. Moreover, this vertex also exists in any diagram $V(L)$ where $\{p, q, r\} \subseteq L \subseteq S$.

Proof: Otherwise, consider two occurrences v_1, v_2 of such vertices in $V(\{p, q, r\})$, as in Figure 8. Since the p - and q -regions are each connected, two paths P and Q exist, connecting




Figure 8.

v_1 and v_2 but not containing them: P is contained in p -region and Q is contained in q -region. Now consider a tour on the cycle $v_1 \circ P \circ v_2 \circ Q$. In v_1 , we have r -land to the right; in v_2 , we have r -land to the left; i.e., r -land lies inside the circle as well as outside. But since every site can be represented only once at a vertex, r -land can cross the circle neither at v_1 nor at v_2 from the inside to the outside. \square

Corollary 1: Any edge e of $V(S)$ is determined uniquely by at most four sites and their counterclockwise order, say (p, r, q, s) . The same edge also exists in any diagram $V(L)$ where $\{p, q, r, s\} \subseteq L \subseteq S$.




Figure 9.

Proof: This follows from Lemma 3 and the fact that $VR(o, S) \subseteq VR(o, L)$ for $o \in \{p, q\}$. \square

As an important instance of the corollary the case $L = \{p, q, r, s\}$ tells us that the edge e exists even in $V(\{p, q, r, s\})$. This fact is important for the algorithm of [MMO]. This algorithm is incremental, i.e., at time t the Voronoi diagram $V(R)$ for a certain subset $R \subseteq S$ is already constructed and we add a new site not already in R to the Voronoi diagram. Some edges in $V(R)$ are thereby shortened and others are erased if they are covered by the new region. In order to construct the new region, we have to determine these intersections; this implies examining intersections between bisecting curves. Theorem 1 characterizes which bisecting curves we need to consider.

Theorem 1: Let $e = (p, r, q, s)$ be an edge of the Voronoi diagram $V(R)$ for sites $p \neq q \in R \subseteq S$, $r, s \in R - \{p, q\}$. Then for any set $L \subseteq R$ with $p, q, r, s \in L$, the edge e is also an edge of $V(L)$ and we have $e \cap \overline{VR(t, R \cup \{t\})} = e \cap \overline{VR(t, L \cup \{t\})}$ for any site $t \notin R$.

Proof: By corollary 1, edge e also exists in $V(L)$. Now consider any point $x \in e$. We show that a sufficiently small neighborhood $U(x)$ looks the same in $V(L \cup \{t\})$ as in $V(R \cup \{t\})$. Choose $U(x)$ such that $y \in VR(p, R) \cup VR(q, R)$ for all points $y \in U(x)$. Let $o \in \{p, q, t\}$ be such that $y \in VR(o, R \cup \{t\})$ for a point $y \in U(x)$. From $VR(o, R \cup \{t\}) \subseteq VR(o, L \cup \{t\})$ we conclude $y \in VR(o, L \cup \{t\})$. The theorem follows since the Voronoi regions form a partition of $U(x)$. \square

If we take $L = \{p, q, r, s\}$ then the intersection between the edge e and the region of a new site t can be constructed by looking at the diagram of only five sites p, q, r, s and t . We take this as the basic operation of our algorithm.

Basic Operation

Input: An edge $e = (p, r, q, s)$ and a site $t \notin \{p, q, r, s\}$

Output: The combinatorial structure of $e \cap \overline{VR(t, \{p, q, r, s, t\})}$, i.e., one of the following:

- 1) intersection is empty
- 2) intersection is non-empty and consists of a single component:
 - a) e itself
 - b) a segment of e adjacent to the (p, r, q) endpoint
 - c) a segment of e adjacent to the (p, q, s) endpoint
 - d) a segment not adjacent to any endpoint of e
- 3) intersection is non-empty and consists of exactly two components
- 4) neither of the above

Remark: We show in Lemma 6, section III, that case 4 never arises and that the two components in case 3 are adjacent to one endpoint of e each.

III. Incremental construction of abstract Voronoi diagrams — the general case

In this section, we describe the incremental construction algorithm. We start with three sites ∞, p, q where p and q are chosen at random and then add the remaining sites in random order. At the general step, we have to consider a set $R \subseteq S$ of sites with $\infty \in R$ and $|R| \geq 3$. We maintain the following data structures.

- 1) The Voronoi diagram $V(R)$: It is stored as a planar graph as described in the previous section.
- 2) The conflict graph $G(R)$: The conflict graph is a bipartite graph. The nodes on one side are the vertices and edges of the Voronoi diagram $V(R)$ and on the other side the sites in $S - R$. Correspondingly, there are two types of edges in the conflict graph. An edge of the first type connects edge e of $V(R)$ and site $s \in S - R$ iff $e \cap \overline{VR(s, R \cup \{s\})} \neq \emptyset$. An edge of the second type connects a vertex $v \in V(R)$ and a site $s \in S - R$ iff $v \in \overline{VR(s, R \cup \{s\})}$. We say: the site s is in conflict with the edge e (the vertex v).

In the sequel, we proceed as in the argumentation of [MMO]. Because we drop the general position assumption, the boundary of a new region can now contain vertices or even edges of the Voronoi region constructed so far. Therefore we will go through all proofs of [MMO] again and take care of the boundary of Voronoi regions.

We assume that the Voronoi diagram $V(R)$ and the conflict graph $G(R)$ for $R \subseteq S$ are already constructed and discuss how to adjust $V(R)$ and $G(R)$ when adding a new site $s \in S - R$ to R . We first concentrate on $V(R)$. From now on, let $S := \overline{VR(s, R \cup \{s\})}$.

Lemma 4: $\bar{S} = \emptyset \iff \deg_{G(R)}(s) = 0$.

Proof: If $\bar{S} = \emptyset$ then clearly $\deg_{G(R)}(s) = 0$. Thus let $\deg_{G(R)}(s) = 0$. Then $\bar{S} \cap V(R) = \emptyset$ and $\bar{S} \subseteq \text{int } VR(p, R)$ for a site $p \in R$, since $V(R) = \bigcup_{r \in R} \text{bd } VR(r, R)$. Thus $VR(p, R \cup \{s\})$ is not simply-connected. \square

We are interested in the part of the diagram that is covered by S . Thus let $I := V(R) \cap \bar{S}$.

Lemma 5: I is a connected set and intersects $\text{bd } S$

Proof: Let I_1, I_2, \dots, I_k be the connected components of I for some k . No I_j can be contained entirely in the interior $\text{int } S$ of S . Otherwise, a simple closed curve $C \subseteq \text{int } S$ would exist, such that I_j is contained in the inner domain of C and C does not intersect $V(R)$. Thus $C \subseteq VR(r, R)$ for some $r \in R$. Since Voronoi regions are simply-connected, C and its interior must belong to $VR(r, R)$ and hence, C cannot contain I_j in its interior.

Thus let $k \geq 2$. Then a path $P \subseteq \bar{S} - I$ exists, connecting two points x and y on the boundary of S and separating I_1 from I_2 , cf. Figure 10. From $P \cap I = \emptyset$ we have $P \cap V(R) = \emptyset$




Figure 10.

and thus $P \subseteq \text{int } VR(r, R)$ for a site $r \in R$. Since $x, y \in P$ all sufficiently small neighborhoods $U(x)$ and $U(y)$ are entirely contained in $VR(r, R)$. The points in the intersection of these neighborhoods with the complement of S thus lie in $VR(r, R \cup \{s\})$ and can be connected by a path $Q \subseteq \text{int } VR(r, R \cup \{s\}) \subseteq \text{int } VR(r, R)$. The cycle $P \circ Q$ is therefore entirely contained in $\text{int } VR(r, R)$ and contains I_1 or I_2 in its interior. This is a contradiction. \square

Lemma 6: Let e be an edge of $V(R)$. If $e \cap \bar{S} \neq \emptyset$, then either $e \cap \bar{S} = V(R) \cap \bar{S}$ and $e \cap \bar{S}$ is a single component or $e - \bar{S}$ is a single component.

Proof: Assume first that $e \cap \bar{S} = V(R) \cap \bar{S}$. Because $V(R) \cap \bar{S}$ is connected by Lemma 5, we conclude that $e \cap \bar{S}$ is connected. Assume next that $e \cap \bar{S} \neq V(R) \cap \bar{S}$. Then with every point $x \in e \cap \bar{S}$ one of the subpaths of e connecting x to an endpoint of e must be contained in \bar{S} . Hence $e - \bar{S}$ is a single component. \square

We need to construct the boundary of S . From the discussion in section II we know that the boundary of S is a simple closed curve. Therefore $\text{bd } S \cap V(R)$ is a cyclic sequence of

- proper crossings between edges of $V(R)$ and $\text{bd } S$
- vertices of $V(R)$
- parts of edges of $V(R)$.

In the case of general position, $bd\ S \cap V(R)$ contains only intersection points of the first type. Also, the cyclic sequence of these intersection points and, thus, the new edges on $bd\ S$ are easy to find by simply walking along the boundary edges of the splitted regions, cf. Figure 11. The correctness of this procedure follows from the connectivity of $\bar{S} \cap V(R)$. We will now




Figure 11.

discuss how to extend this strategy to degenerate cases. We distinguish two major cases.

Case 1: There is a single edge e which is in conflict with s .

Assume that e separates the p -region from the q -region, i.e., $e \subseteq J(p, q)$ for two sites $p, q \in S$. We compute $e \cap \bar{S}$ using our basic operation. Note that $e \cap \bar{S}$ cannot be a single point. If $e \cap \bar{S}$ were a single point then the situation shown in Figure 12 would result, a contradiction to Fact 1. Thus $e \cap \bar{S}$ is a proper segment. Let v_1 and v_2 be the two endpoints of $e \cap \bar{S}$. We




Figure 12.

delete $e \cap \bar{S}$ from the diagram, add v_1 and v_2 as vertices and two edges with endpoints v_1 and v_2 . The two edges are part of $J(p, s)$ and $J(q, s)$, respectively.

Case 2: There are at least two edges which are in conflict with s .

Then $e - \bar{S}$ is a single component for each edge e of $V(R)$. The conflict graph gives us all edges in conflict with s . Let $e \subseteq J(p, q)$ be such an edge, i.e., $e \cap \bar{S} \neq \emptyset$. If $e \subseteq \bar{S}$ then e disappears completely. If $\emptyset \neq e - \bar{S} \neq e$ then some part of e survives. Let e_1 be one of the components of $e \cap \bar{S}$; e_1 can be computed by our basic operation. Then \bar{e}_1 contains exactly one endpoint of e . Let w be the other endpoint of \bar{e}_1 . Then w lies on $J(p, q)$ and also on the boundary of S . Thus in a small neighborhood $U(w)$ of w we have p -, q - and s -region, i.e., w is a vertex of $V(R \cup \{s\})$. Also, in $V(R \cup \{s\})$ exactly three edges are incident to w ; they are parts of $J(p, q)$, $J(q, s)$ and $J(p, s)$, respectively.

Consider the vertices of $V(R)$ next. If no edge e of $V(R)$ has v as an endpoint of $e - \bar{S}$ then v disappears from the diagram. So let us assume that at least one edge e of $V(R)$ has v as an endpoint of $e - \bar{S}$, i.e., v does not lie in the interior of S . The vertex v has degree

at least three in $V(R)$. Also, either v lies outside \bar{S} or at least one edge incident to v must intersect \bar{S} in all neighborhoods $U(v)$ because otherwise the situation shown in Figure 13 would result, contradicting Fact 1. If v lies outside \bar{S} then nothing changes in the vicinity




Figure 13.

of v . So let us assume that v lies on the boundary of S . Since $bd S$ is a simple closed curve which passes through v the edges incident to v split into two groups, those which intersect \bar{S} in all neighborhoods $U(v)$ of v and those which do not. Both groups form a contiguous subsequence of the cyclic adjacency list of v , cf. Figure 14. Let f' and f'' be the first and




Figure 14.

the last edge of the subsequence of edges which run outside $\bar{S} \cap U(v)$ for a sufficiently small neighborhood $U(v)$ of v ; $f' = f''$ is possible. Let e'' and e' be the first and last edge in counterclockwise order of the subsequence of edges which run inside $\bar{S} \cap U(v)$ for all small neighborhoods $U(v)$ of v . Let f' and e' border p -region and let f'' and e'' border q -region. Then we delete the edges between e'' and e' and add two new edges incident to v . These edges are part of $J(q, s)$ and $J(p, s)$, respectively.

At this point we have constructed all vertices of $V(R \cup \{s\})$ and their cyclic adjacency lists. We have not yet linked the two occurrences of each edge. Let $I := V(R) \cap \bar{S}$. We know an embedding of I into the plane. The boundary of the outer face is a closed curve since I is connected. Also, all the vertices on $bd S$ lie on I and hence a clockwise traversal of the boundary of the outer face yields the cyclic ordering of the vertices of $bd S$. This allows us to link the two occurrences of each edge. We conclude in

Lemma 7: $V(R \cup \{s\})$ can be constructed from $V(R)$ and $G(R)$ in $O(\deg_{G(R)}(s) + 1)$ steps. Here our basic operation is considered as one step. \square

Let us now turn to the computation of $G(R \cup \{s\})$. We distinguish three cases: edges of $V(R \cup \{s\})$ which already were edges of $V(R)$, edges which are part of edges of $V(R)$, and edges which are completely new.

For any edge e of $V(R)$ not in conflict with s , the conflict information remains the same. Since $e \cap \overline{VR(s, R \cup \{s, t\})} \subseteq e \cap \overline{VR(s, R \cup \{s\})} = \emptyset$ the intersection $e \cap \overline{VR(t, R \cup \{t\})}$ remains unchanged when we add s to R for any $t \in S - R - \{s\}$.

If e is in conflict with s but $e \not\subseteq \bar{S}$ then $e - \bar{S}$ consists of at most two subsegments by Lemma 6. Let e' be one of those subsegments. For e' the same argumentation applies as above since $e' \cap \overline{VR(s, R \cup \{s\})} = \emptyset$. Hence any site in conflict with e' must be in conflict with e .

For a newly constructed edge, we show that it is sufficient to gather its conflict information from edges on that part of the boundary of the splitted region that is covered by S . Let e_{12} be such a new edge with endpoints x_1 and x_2 . The new edge e_{12} either intersects the interior of exactly one Voronoi region $VR(p, R)$ or it runs on the boundary of a region $VR(p, R)$ such that $VR(p, R)$ and S are on different sides of e_{12} . This follows from the discussion in Case 1 and Case 2 above.

Let P be the part of the boundary of $VR(p, R)$ between x_1 and x_2 such that in all sufficiently small neighborhoods of x_1 and x_2 the path P is contained in \bar{S} .




Figure 15.

Claim: $P \subseteq \bar{S}$

Proof: If P crosses the boundary $bd S$ of S then $V(R) \cap \bar{S}$ becomes unconnected, contradicting Lemma 5. \square

Lemma 8: Let $t \in S - R - \{s\}$ be in conflict with e_{12} in $V(R \cup \{s\})$. Then t is in conflict with P in $V(R)$.

Proof: Consider $VR(p, R)$. By the definition of conflict, a point $x \in e_{12}$ exists such that $x \in$




Figure 16.

$\overline{VR(t, R \cup \{s, t\})} \subseteq \overline{VR(t, R \cup \{t\})}$. Since we claim a contradiction we assume that t is not in conflict with P in $V(R)$, i.e., $\overline{VR(t, R \cup \{t\})} \cap P = \emptyset$. If x is on P the contradiction is evident. Thus, assume $x \notin P$. By assumption, $VR(t, R \cup \{s, t\}) \cap U(x_1) \subseteq VR(t, R \cup \{t\}) \cap U(x_1) = \emptyset$ for any sufficiently small neighborhood $U(x_1)$ of x_1 . Now consider in any such neighborhood of x_1 the wedge spanned by e_{12} and the part of $bd VR(p, R)$ outside S . The points in this wedge all belong to $VR(p, R \cup \{s, t\})$. The same is true for any sufficiently small neighborhood of x_2 . Since $VR(p, R \cup \{s, t\})$ is connected, there is a path Q from x_1 to x_2 running completely inside $VR(p, R \cup \{s, t\}) \subseteq VR(p, R \cup \{t\})$ except at the endpoints.




Figure 17.

We can assume that Q does not touch the boundary $bd S$ of S (and therefore not includes x). By definition of P and Q the Voronoi region $VR(t, R \cup \{t\})$ cannot intersect these two paths. Moreover, x lies in the interior of the cycle $x_1 \circ P \circ x_2 \circ Q$; otherwise $VR(p, R)$ would not be simply connected. From $P \cap \overline{VR(t, R \cup \{t\})} = \emptyset$ and $x \in \overline{VR(t, R \cup \{t\})}$ we conclude that $VR(t, R \cup \{t\})$ lies in the interior of the cycle. Since $bd VR(t, R \cup \{t\}) \cap x_1 \circ P \circ x_2 = \emptyset$ this is a contradiction to the fact that $VR(p, R \cup \{t\})$ is simply connected. \square

Now we know all candidate-sites possibly in conflict with e_{12} , namely the sites in conflict with an edge or a vertex on P . Using our basic operation, we can test whether a conflict exists with these candidate sites or not. The conflict information of a newly constructed vertex $v = e \cap bd S$ is obtained from the conflict information of edge e in the same way. Note that we visited only conflicting edges and that each such edge can belong at most 2 times to a path P by planarity. We conclude in

Lemma 9: $G(R \cup \{s\})$ can be constructed from $V(R)$ and $G(R)$ in time

$$O\left(\sum_{\{f,s\} \in G(R)} \deg_{G(R)}(f)\right).$$

Theorem 2: a) Let $s \in S - R$. Then the data structures $G(R \cup \{s\})$ and $V(R \cup \{s\})$ can be obtained from $G(R)$ and $V(R)$ in time

$$O\left(\sum_{\{f,s\} \in G(R)} \deg_{G(R)}(f)\right).$$

b) For $R \subseteq S$, $|R| = 3$ and $\infty \in R$, the data structures $V(R)$ and $G(R)$ can be set up in time $O(n)$ where $n = |S|$.

Proof: a) This point summarizes Lemmas 7 and 9.

b) The Voronoi diagram $V(R)$ for three sites ∞ , p and q has the structure shown in Figure 18 and can certainly be set up in time $O(1)$. Also for each of the edges e of $V(R)$ and each of the $n - 3$ sites in $S - R$ one can test $e \cap VR(t, R \cup \{t\}) \neq \emptyset$ by our basic operation in time $O(1)$. This proves (b). \square




Figure 18.

We next turn to the analysis of the algorithm. Our goal is to analyze the running time of our algorithm by an appeal to Clarkson and Shor's results about the running time of randomized incremental construction. However, Clarkson and Shor's analysis relies heavily on their general position assumption. It states that each region (here edges and vertices) is *uniquely* determined by a bounded number of objects (here sites). Now observe that any vertex that has r_1, \dots, r_l -regions, $l > 3$, incident to it is defined by any 3-tuple (r_i, r_j, r_k) , $1 \leq i < j < k \leq l$, by Lemma 3. A similar observation holds for any edge with endpoints of degree greater than 3. This shows that the general position assumption of [CS] is generally not satisfied, when the bisecting curves are not in general position. However, the general position assumption of [CS] is satisfied if the bisecting curves are in general position as defined in [MMO], i.e., any intersection is a proper crossing and no four curves have a common point.

The idea behind our proof is now as follows. We take our system of bisecting curves and conceptually perturb it so as to obtain a system in general position. The perturbation is carried out in such a way that no conflicts are lost. This allows us to show that the cost of one iteration step of the algorithm on the original system is bounded by a constant times the cost of the corresponding step on the perturbed system; the asymptotic time bound thus holds for both. The perturbed system is in general position and hence Clarkson and Shor's analysis applies.

In [K89] Klein proves the following

Theorem 3: Let $\mathcal{J} = \{J(p, q); p, q \in S, p \neq q\}$ be a system of bisecting curves as defined in section II. Then \mathcal{J} can be deformed in such a way that the resulting system $\tilde{\mathcal{J}}$ has the following properties:

- 1) No other intersections than cross-points exist between the curves in $\tilde{\mathcal{J}}$. At most 3 curves cross at a point (general position)
- 2) For each subset $R \subseteq S$: After contracting some edges, the perturbed Voronoi diagram $\tilde{V}(R)$ differs from the original diagram $V(R)$ only by a deformation.

A perturbation is necessary for those curves that touch each other or even share a common segment and for more than 3 curves intersecting in a common point. Curves sharing a common segment are spreaded out like wires of a cable. A point common to more than 3

curves or a touching point of two curves is split into a graph of newly constructed intersection points and curves segments connecting them. These segments give rise to the edges in $\tilde{V}(R)$ that are contracted in Theorem 3.

We will now show that no conflict information is lost by the perturbation. First consider an edge e in $V(R)$ that is in conflict with some site t in $S - R$. Edge e also exists in $\tilde{V}(R)$. Moreover, Theorem 3 implies that e exists in $\tilde{V}(R \cup \{t\})$ iff e exists in $V(R \cup \{t\})$. Thus, each edge of $V(R)$ is also represented in $\tilde{V}(R)$ and their conflict information is the same.

Next consider a vertex $v \in V(R)$ that is replaced by a graph $\mathcal{G}(v)$ in $\tilde{V}(R)$. Analogously, Theorem 3 implies that v exists in $V(R \cup \{t\})$ iff $\mathcal{G}(v)$ exists in $\tilde{V}(R \cup \{t\})$. Hence the vertices of $\mathcal{G}(v)$ all have the same conflict information as v .

Let us now compare the running time of our algorithm on both systems of bisecting curves \mathcal{J} and $\tilde{\mathcal{J}}$, respectively. Addition of a new site $s \in S - R$ to $V(R)$ and $\tilde{V}(R)$ takes time proportional to

$$\sum_{\{f,s\} \in G(R)} \deg_{G(R)}(f) \quad \text{and} \quad \sum_{\{f,s\} \in \tilde{G}(R)} \deg_{\tilde{G}(R)}(f), \quad \text{respectively.}$$

As discussed above we have $\deg_{G(R)}(e) = \deg_{\tilde{G}(R)}(e)$ for any edge $e \in V(R)$ and $\deg_{G(R)}(v) = \deg_{\tilde{G}(R)}(v')$ for any vertices $v \in V(R), v' \in \mathcal{G}(v) \subseteq \tilde{V}(R)$; thus $\deg_{G(R)}(s) \leq \deg_{\tilde{G}(R)}(s)$. Hence, any iteration step of the algorithm on the original curve system costs at most a constant factor more than the corresponding step for the perturbed system.

Theorem 4: The abstract Voronoi diagram $V(S)$ of n sites can be constructed by a randomized algorithm in time $O(n \log n)$.

Proof: In [CS] Clarkson and Shor proved that randomized incremental construction has expected running time

$$O(m(n) + n \cdot \sum_{1 \leq r \leq n/2} m(r)/r^2 + n)$$

provided that their general position assumption holds, initialization takes time $O(n)$ and addition of an object (here site) s to the set R takes time proportional to

$$\sum_{\{f,s\} \in G(R)} \deg_{G(R)}(f);$$

the summation is over all regions (here edges and vertices) of the current structure (here Voronoi diagram $V(R)$) which conflict with site s . Also, $m(r)$ is the expected size of the structure for a random subset $R \subseteq S$ of r elements. In our case we have $m(r) = O(r)$ since the Voronoi diagram is a planar graph with r regions and therefore has less than $3r - 6$ edges and $2r - 4$ vertices by Euler's relation, initialization takes time $O(n)$ by Theorem 2b and the time bound for the update step holds by Theorem 2a. Thus the running time of our algorithm on the perturbed system is $O(n \log n)$. From the considerations above we conclude that the same bound holds for the original system. \square

Remark: In our algorithm ∞ is always a member of R . An inspection of Clarkson's argument shows that this minor deviation from randomness does not change the time bound.

IV. Conclusions and Open Problems

We showed under very general assumptions on the bisecting curves that the construction of abstract Voronoi diagrams can be transformed efficiently and purely combinatorially into the construction of abstract Voronoi diagrams for five sites. Further research will be done on the question of how the construction of the 5-site-diagram can be carried out efficiently dependent on the type of bisecting curves. Our goal is a "universal" algorithm for Voronoi diagrams. Nevertheless, we still have open problems to solve: Can the concept of abstract Voronoi diagrams be generalized to higher dimensions? What can be done in two dimensions without the assumption that bisectors are non-closed curves?

References:

- [A] F. Aurenhammer (1988): *Voronoi Diagrams — A survey*, tech. report 263, Institutes for Information Processing, Graz Technical University, Austria
- [CS] K. L. Clarkson, P. W. Shor (1988): *Algorithms for Diametral Pairs and Convex Hulls that are Optimal, Randomized and Incremental*, Proc. 4th ACM Symposium on Computational Geometry, pp. 12–17
- [GS] L. Guibas, J. Stolfi (1985): *Primitives for the Manipulation of General Subdivisions and the Computation of Voronoi Diagrams*, ACM Transactions on Graphics 4, pp. 74–123
- [K88a] R. Klein (1988): *Abstract Voronoi Diagrams*, Habilitationsschrift, Mathematics Faculty of University of Freiburg i. Br., LNCS 400
- [K88b] R. Klein (1988): *Abstract Voronoi Diagrams and their Applications (extended abstract)*, in: H. Noltemeier (Ed.), *Computational Geometry and its Applications (CG '88)*, LNCS 333, pp. 148–157
- [K88c] R. Klein (1988): *Voronoi Diagrams in the Moscow Metric (extended abstract)*, in: J. van Leeuwen, (Ed.), *Graphtheoretic Concepts in Computer Science (WG '88)*, LNCS 344, pp. 434–441
- [K89] R. Klein (1989): *Combinatorial Properties of Abstract Voronoi Diagrams*, WG 89, Rolduc, to appear in LNCS
- [L] D. T. Lee (1980): *Two-Dimensional Voronoi Diagrams in the L_p -Metric*, J. ACM 27, pp. 604–618
- [LS] D. Leven, M. Sharir (1986): *Intersection and Proximity Problems and Voronoi Diagrams*, in: J. Schwartz and C. K. Yap (eds.), *Advances in Robotics*, Vol. 1, Lawrence Erlbaum, pp. 187–228
- [MMO] K. Mehlhorn, S. Meiser, C. Ó' Dúnlaing (1989): *On the Construction of Abstract Voronoi Diagrams*, Journal of Discrete & Computational Geometry, to appear