Abstract Voronoi Diagram in 3-Space

Ngọc-Minh Lê *

Abstract

We propose a class of abstract Voronoi diagrams in 3-space that generalizes the planar abstract Voronoi diagram of Klein. Our class of abstract Voronoi diagrams includes the Voronoi diagram of point sites in general position under any convex distance function. To characterize the abstract Voronoi diagram in 3-space we introduce the notion of intersection characteristic. We determine the intersection characteristic for the simplex, the L_∞, and the L_p distance function. We find that the intersection characteristic in case of the simplex distance function is similar to that of the usual Euclidean distance. This enables us to give a randomized incremental algorithm for computing the Voronoi diagram under the simplex distance function in quadratic expected time.

1 Introduction

The Voronoi diagram is one of the most useful data structures in computational geometry, with applications in diverse fields of science and engineering. We refer to [2, 3, 22] for a history of the Voronoi diagram and its applications.

The Voronoi diagram of a set of sites in d-space partitions d-space into Voronoi regions, one for each site. The Voronoi region of a site s is the set of points in d-space for which s is the closest site among all the sites. The Voronoi diagram is then defined to be the union of the boundaries of the Voronoi regions of the sites. Thus, the type of the Voronoi diagram depends on the underlying distance measure, e.g. Euclidean or L_p norm or convex distance function, as well as on the kind of sites, e.g. points or lines.

In the plane, several techniques for constructing the Voronoi diagram under the Euclidean norm or convex distance functions have been found: divide-and-conquer [25], sweepline [10], geometric transform [5, 8], and randomized incremental construction [7]. A major contribution to the theory of planar Voronoi diagrams was made by Klein [13]. Klein defines a new object, the planar abstract Voronoi diagram, by a set of axioms that are general enough to include most of the known Voronoi diagrams as concrete models for the planar abstract Voronoi diagram. Efficient algorithms for constructing the planar abstract Voronoi diagram are provided in [13, 14].

A generalization of the planar Voronoi diagram to 3-space was proposed in [17]. The abstract Voronoi diagram in [17] is rather restricted: it covers only some of the Voronoi diagrams based on convex distance functions.

In this paper, we continue the research on abstract Voronoi diagram in 3-space initiated in [17]. Similar to [13], the abstract Voronoi diagram in 3-space is developed from the notion of

*This work was supported by Grant A0E97/98.EG04 and Grant DAG97/98.EG15 while the author was with the Department of Computer Science of the Hong Kong University of Science and Technology. Part of this work was supported by Deutsche Forschungsgemeinschaft while the author was with Praktische Informatik VI of the FernUniversität Hagen. Author's address: HCM City University of Technology, Department of Information Technology, 268 Ly Thuong Kiet St., Dist. 10 HCM City, Vietnam. Email: minhle@dit.hcmut.edu.vn

bisecting surface as the basic concept: for each pair of sites p and q we assume the existence of a bisector $B(p,q)$, which is homeomorphic to the plane and separates 3-space into a p-region and a q-region. The Voronoi region of site p is then the intersection of all p-regions for different q's. We postulate that Voronoi regions are path-connected and partition 3-space.

Because the Voronoi diagram under convex distance functions should serve as a model for the abstract Voronoi diagram in 3-space, we further require that $B(p,q)$ and $B(q,r)$ intersect properly in finitely-many biinfinite curves and that $B(p,q)$, $B(q,r)$, and $B(r,s)$ intersect properly at finitely many points. It is known that bisectors of point sites under convex distance functions have such properties, refer to [11, 15, 18]. The abstract Voronoi diagram is then defined to be the union of the boundaries of the Voronoi regions.

As shown in Sections 4 and 5, the set of axioms above implies several important structural properties that are known to hold for the Voronoi diagram under convex distance functions.

To characterize the system of bisectors of a set of sites S satisfying the axioms described above, we introduce the notion of intersection characteristic. The intersection characteristic of the system of bisectors of S is defined to be the pair of the ranges of $|B(p,q) \cap B(q,r)|$ and of $|B(p,q) \cap B(q,r) \cap B(r,s)|$, for any distinct $p,q,r,s \in S$, where $|B(p,q) \cap B(q,r)|$ and $|B(p,q) \cap B(q,r) \cap B(r,s)|$ denote the number of the connected components of $B(p,q) \cap B(q,r)$ and of $B(p,q) \cap B(q,r) \cap B(r,s)$, respectively.

For example, the intersection characteristic of any system of bisectors of point sites in general position under the Euclidean norm is the pair $\{1\}$ and $\{1\}$. We find that the intersection characteristic of any system of bisectors of point sites is a pair of subsets of

- $\{0,1\}$ and $\{0,1\}$, under the simplex distance function.
- $\{0,1\}$ and $\{0,1,2\}$, under the L_∞ distance function.

The method for proving these results is also of practical interest, because we may make use of it to implement a straightforward incremental construction of the Voronoi diagrams under the convex distance functions above.

We still don't have an efficient algorithm for computing the general abstract Voronoi diagram. However, at present only few efficient algorithms for computing the Voronoi diagram under convex distance functions different from the Euclidean distance are known, namely for the simplex distance function [4], Subsection 7.1 of this paper, and for the L_∞ distance function [4].

The paper is organized as follows. General definitions are given in Section 2. In Section 3 we introduce the abstract Voronoi diagram and state some properties. In Section 4 we describe an important property of Voronoi regions. Further properties of the abstract Voronoi diagram are explored in Section 5. We discuss the complexity of the abstract Voronoi diagram in Section 6. Examples of the abstract Voronoi diagram are given in Section 7. Finally, some concluding remarks are made in Section 8.

2 Preliminaries

Let P be a convex body in $\mathbb{R}^d$ containing the origin in its interior. We define the distance induced by P from a point x to a point p by $d_P(x,p) = \inf\{t \geq 0 \mid p \in (x + tP)\}$. We call d_P the *(convex) distance function induced by P*. In particular, the L_∞ distance function is obtained by taking P to be a hypercube, and the L_1 distance function is given by $d_{L_1}(x,p) = \sum_{i=1}^{d} |x_i - p_i|$. Let S be a set of n points, or *sites*, in $\mathbb{R}^d$. The *bisector for sites p and q*, denoted $B(p,q)$, is the set

$$B(p,q) \quad = \quad \{x \in \mathbb{R}^d \mid d_P(x,p) = d_P(x,q)\}.$$

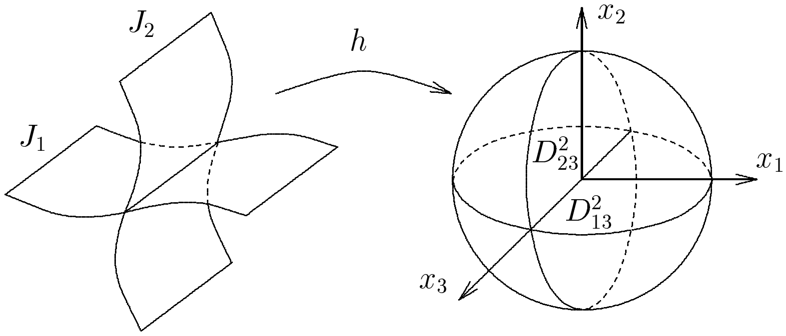

Figure 1: The surfaces J_1 and J_2 intersect properly.

The *Voronoi diagram of* S *under* d_P is defined as the partition of $\mathbb{R}^d$ into n *Voronoi regions* (or *cells*), one for each site. The region of a site p consists of all points that lie, with respect to d_P, closer to p than to any of the other sites.

We will use $\operatorname{bd} A, \operatorname{int} A$, and $\overline{A}$ to denote the boundary, the interior, and the closure of the set A. We set, for $d \geq 1$,

D^d	$=$	$\{x \in \mathbb{R}^d \mid \lvert x\rvert \leq 1\}$	the d-ball, d-disk
$\operatorname{int} D^d$	$=$	$\{x \in \mathbb{R}^d \mid \lvert x\rvert < 1\}$	the open d-ball, d-disk
S^{d-1}	$=$	$\{x \in \mathbb{R}^d \mid \lvert x\rvert = 1\}$	$(d-1)$-sphere
$\mathbb{R}^d_+$	$=$	$\{x = (x_1, x_2, \dots, x_d) \in \mathbb{R}^d \mid x_d \geq 0\}$	halfspace.

An *m-dimensional manifold M* in $\mathbb{R}^d$ is a subset of $\mathbb{R}^d$ each point of which has a neighbourhood that is homeomorphic to the m-ball. Usually we will say *m-manifold* for short. A connected 2-manifold is also called a *surface*. An m-manifold M is *locally flat* if for each point $x \in M$ there is an open set $U \subset \mathbb{R}^d$ containing x, an open set $U' \subset \mathbb{R}^d$ containing $\vec{0}$ and a homeomorphism $h : U \to U'$ so that $h(M \cap U) = U' \cap \mathbb{R}^{d-1}$, where we identify $\mathbb{R}^{d-1}$ with the hyperplane $x_d = 0$. A closed subset of $\mathbb{R}^d$ is called a *biinfinite curve* if it is homeomorphic to $\mathbb{R}$.

We denote by D^2_{ij}, with $i, j = 1, 2, 3$, the 2-disk in the $x_i x_j$-plane.

Definition 1 *Let J_1 and J_2 be two surfaces in 3-space so that $J_1 \cap J_2$ consists of disjoint curves. We say that J_1 and J_2* intersect properly *if for each point $x \in J_1 \cap J_2$ there is a neighbourhood U of x and a homeomorphism $h : U \to D^3$ so that*

1. *$h(J_1 \cap U) = D^2_{13}$,*
2. *$h(J_2 \cap U) = D^2_{23}$.*

See Figure 1.

Definition 2 *Let L be a curve and J a surface in 3-space so that $L \cap J$ consists of distinct points. We say that L and J* intersect properly *if for each point $x \in L \cap J$ there is a neighbourhood U of x and a homeomorphism $h : U \to D^3$ so that*

1. *$h(J \cap U) = D^2_{13}$,*
2. *$h(L \cap U) = I$, where I denotes the interval $[-1, 1]$ of the x_2-axis.*

See Figure 2.

Sometimes we abuse the language, and say, for instance, that a set A *is* a circle (an open disk, 2-sphere, ...) instead of saying that A *is homeomorphic* to the circle, ...

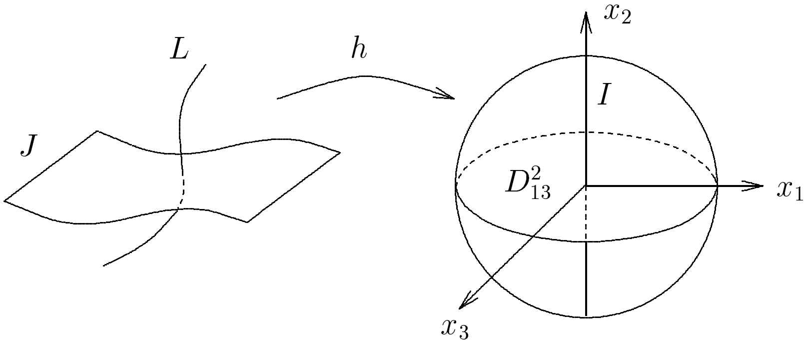

Figure 2: The curve L and the surface J intersect properly, I denotes the interval $[-1, 1]$ of the x_2-axis.

3 Abstract Voronoi diagrams

A closed subset J of $\mathbb{R}^3$ is called a *bisector* if J is homeomorphic to $\mathbb{R}^2$ or S^2. By the Jordan-Brouwer Separation Theorem, the set $\mathbb{R}^3 - J$ consists of exactly two disjoint domains of which J is the common boundary.

For $n \in \mathbb{N}$ let $S = \{1, \ldots, n\}$. We associate to each tuple $(p, q) \in S \times S$ with $p \neq q$ a bisector, denoted by $B(p, q)$, so that $B(p, q) = B(q, p)$. Furthermore, we associate to each pair $(p, q) \in S \times S$ with $p \neq q$ the sets $D(p, q)$ and $D(q, p)$ that are the two disjoint domains separated by $B(p, q)$. We call $B(p, q)$ the *bisector for sites* p and q and $D(p, q)$ the *region of dominance* of p over q. Let $\mathcal{J} = \{B(p, q) \mid p, q \in S, p \neq q\}$. The tuple $(S, \mathcal{J})$ is called a *system of bisectors.* Following Klein [13], we set

$$R(p,q) = \begin{cases} D(p,q) \cup B(p,q) & \text{if} \quad p < q, \\ D(p,q) & \text{if} \quad p > q, \end{cases} \tag{1}$$

$$VR(p,S) = \bigcap_{\substack{q \in S \\ q \neq p}} R(p,q). \tag{2}$$

The *abstract Voronoi diagram of S*, denoted by *V(S)*, is defined by

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

For any pairwise distinct sites $p, q, r, s, t \in S$ we set

$$\begin{aligned} B(p,q,r) &= B(r,p) \cap B(p,q), \\ B(p,q,r,s) &= B(s,p) \cap B(p,q,r), \\ B(p,q,r,s,t) &= B(t,p) \cap B(p,q,r,s). \end{aligned}$$

Definition 3 *The system of bisectors $(S, \mathcal{J})$ is called* admissible *if the following conditions are fulfilled:*

1. *For any distinct $p, q \in S$ the bisector $B(p, q)$ is locally flat. Moreover, after one-point compactifying $\mathbb{R}^3$ with a point ∞ the set $B(p, q) \cup \infty$ is locally flat in $\mathbb{R}^3 \cup \infty$.*

2. *The* proper-intersection condition*:*

(a) *For any four distinct sites $p, q, u, v \in S$ the set $B(p,q) \cap B(u,v)$ is a disjoint union of finitely many points, curves and open subsets of $B(p,q)$ and of $B(u,v)$.*

(b) *The set $B(p,q,r)$ has finitely many connected components, each of which is a biinfinite curve. Furthermore, $B(r,p)$ and $B(p,q)$ intersect properly.*

(c) *The set $B(p,q,r,s)$ consists of finitely many points, and $B(s,p)$ and $B(p,q,r)$ intersect properly.*

(d) *$B(p,q,r,s,t) = \emptyset$.*

3. *For each non-empty subset S' of S hold*

(a) *For any $p \in S'$, if the set $VR(p,S')$ is non-empty then it is path-connected.*

(b) *The Voronoi diagram partitions 3-space, i.e., $\mathbb{R}^3 = \bigcup_{p \in S'} VR(p,S')$.*

Remark 1 *The partition of 3-space by the Voronoi diagram given in 3(b) of the definition of admissible system of bisectors above is disjoint.*

Examples of admissible systems of bisectors are systems of bisectors for point sites in general position under semialgebraic convex distance functions such as the L_p, the L_1, and the L_∞ distance function.

Lemma 1 *For any distinct sites p, q, r in S holds*

$$\overline{R(p,q)} \cap \overline{R(q,r)} \subset \overline{R(p,r)}.$$

Proof: Let $T = \{p,q,r\}$. Using 3(b) of the definition of admissible system of bisectors (Definition 3) we obtain

$$\mathbb{R}^3 \quad = \quad VR(p,T) \cup VR(q,T) \cup VR(r,T).$$

Intersecting both sides of the above identity with $R(p,q) \cap R(q,r)$ and using the relations

$$\begin{aligned} R(p,q) \cap R(q,r) \cap VR(p,T) &= R(p,q) \cap R(q,r) \cap R(p,r), \\ R(p,q) \cap R(q,r) \cap VR(q,T) &= \emptyset, \text{ and} \\ R(p,q) \cap R(q,r) \cap VR(r,T) &= \emptyset, \end{aligned}$$

where the last two relations hold because, e.g., $R(p,q) \cap R(q,p) = \emptyset$, yields

$$R(p,q) \cap R(q,r) \quad = \quad R(p,q) \cap R(q,r) \cap R(p,r),$$

which implies

$$R(p,q) \cap R(q,r) \quad \subset \quad R(p,r). \tag{3}$$

If we could prove that

$$\overline{R(p,q)} \cap \overline{R(q,r)} \subset \overline{R(p,q) \cap R(q,r)},$$

then by the relation (3) we would prove the lemma. If $\overline{R(p,q)} \cap \overline{R(q,r)} \neq \emptyset$ then for each $x \in \overline{R(p,q)} \cap \overline{R(q,r)}$ we distinguish the following cases:
Case 1: $x \in D(p,q) \cap D(q,r)$ or $x \in B(p,q) \cap D(q,r)$ or $x \in D(p,q) \cap B(q,r)$. Then clearly $x \in \overline{R(p,q) \cap R(q,r)}$.
Case 2: $x \in B(p,q) \cap B(q,r)$. In this case, using the proper-intersection condition, we obtain $x \in \overline{R(p,q) \cap R(q,r)}$. □

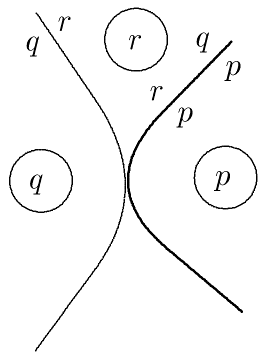

Figure 3: The bisectors $B(p,q)$, $B(p,r)$, and $B(q,r)$ don't intersect properly.

Lemma 2 *Let p, q, r be distinct sites in S, then we have*

$$B(p,q) \cap B(q,r) \subset B(p,r).$$

Proof: By using the definition of $R(\cdot,\cdot)$ and applying Lemma 1 we have

$$\begin{aligned} B(p,q) \cap B(q,r) &= (\overline{R(p,q)} \cap \overline{R(q,p)}) \cap (\overline{R(q,r)} \cap \overline{R(r,q)}) \\ &= (\overline{R(p,q)} \cap \overline{R(q,r)}) \cap (\overline{R(r,q)} \cap \overline{R(q,p)}) \\ &\subset \overline{R(p,r)} \cap \overline{R(r,p)} \\ &\subset B(p,r). \end{aligned}$$

□

Lemma 2 may be false if the proper-intersection condition is violated; refer to Figure 3 for the two-dimensional case, which can be generalized to three-dimension easily.

Lemma 3 *For distinct $p, q, r \in S$ we have*

$$B(r,p) \cap B(p,q) = B(r,q) \cap B(p,q).$$

Proof: By applying Lemma 2 we obtain

$$B(r,p) \cap B(p,q) \quad \subset \quad B(r,q),$$

which clearly implies

$$B(r,p) \cap B(p,q) \quad \subset \quad B(r,q) \cap B(p,q); \tag{4}$$

analogously we have

$$B(r,q) \cap B(p,q) \quad \subset \quad B(r,p) \cap B(p,q). \tag{5}$$

Evidently the claim follows from the relations (4) and (5). □

For $T \subset S$ so that $|T| \geq 3$ we recursively define for distinct sites $p, q \in T$

$$B(T) \quad = \quad B(p,q) \cap B(T-p). \tag{6}$$

Lemma 4 *The set $B(T)$ does not depend on the special choice of distinct p and q in T.*

Proof: First we prove that q can be arbitrarily chosen from $T-p$. The proof will be by induction on $|T|$. By Lemma 3 the assertion is true for $|T|=3$. Assume that the lemma is true for any subset of T with cardinality $|T|-1$. Since $B(T-p)=B(q,r)\cap B(T-p-q)$ for any $r\in T-p-q$, we have $B(T-p)\subset B(q,r)$. So we may write

$$B(T-p)=B(q,r)\cap B(T-p).$$

Therefore

$$\begin{aligned} B(p,q)\cap B(T-p) &= B(p,q)\cap B(q,r)\cap B(T-p) \\ &\subset B(p,r)\cap B(T-p), \qquad \text{by Lemma 2.} \end{aligned}$$

Similarly, we have

$$B(p,r)\cap B(T-p) \quad\subset\quad B(p,q)\cap B(T-p).$$

Thus,

$$B(p,q)\cap B(T-p) \quad=\quad B(p,r)\cap B(T-p). \tag{7}$$

Now we prove that $B(T)$ is independent of the choice of $p\in T$, i.e.,

$$B(p,q)\cap B(T-p)=B(x,y)\cap B(T-x)$$

for any distinct $x,y\in T$.

If $x=p$ then clearly, by (7), the claim holds. Assume that $x\neq p$. We distinguish two cases
Case 1: $x\neq q$. By definition,

$$B(p,q)\cap B(T-p) \quad=\quad B(p,q)\cap B(x,q)\cap B(T-p-x). \tag{8}$$

On the other hand, by (7),

$$\begin{aligned} B(x,y)\cap B(T-x) &= B(x,q)\cap B(T-x) && (9)\\ &= B(x,q)\cap B(p,q)\cap B(T-x-p). && (10) \end{aligned}$$

Thus, the claim is true in this case.
Case 2: $x=q$. By (7), it suffices to prove that

$$B(p,q)\cap B(T-p) \quad=\quad B(q,p)\cap B(T-q). \tag{11}$$

If $|T|=3$, then, by Lemma 3 we are done. Otherwise $|T|>3$. Taking any $r\in T-\{p,q\}$, we obtain

$$\begin{aligned} B(p,q)\cap B(T-p) &= B(p,q)\cap B(q,r)\cap B(T-p-q), \qquad \text{and} && (12)\\ B(q,p)\cap B(T-q) &= B(q,p)\cap B(p,r)\cap B(T-q-p). && (13) \end{aligned}$$

Thus, by Lemma 3, (11) is true. □

Lemma 5 *Let $T=\{p_1,\ldots,p_m\}\subset S$. Then*

$$\begin{aligned} B(T) &= \bigcap_{i=1,\ldots,m-1} B(p_i,p_{i+1}), \qquad \textit{and}\\ B(T) &= \bigcap_{i=1,\ldots,m-1} B(p_m,p_i). \end{aligned}$$

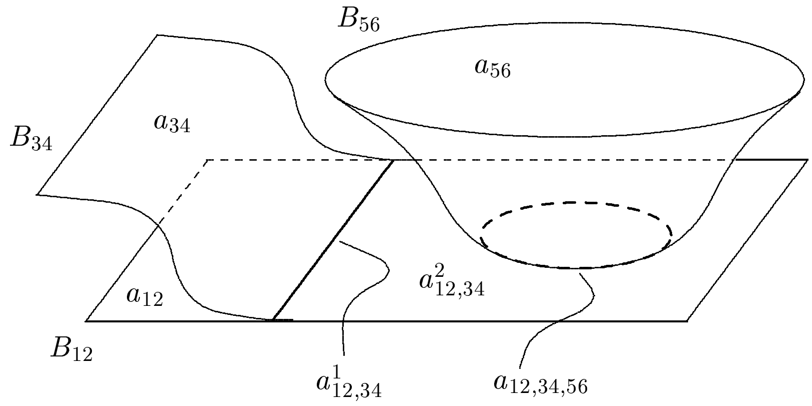

Figure 4: An example for a finite and disjoint decomposition of bisectors and its labelling. (Not all bisectors are shown.)

Proof: Starting from the definition of $B(T)$, see (6), and using Lemma 4, we carefully expand it to obtain the required relations. $\square$

Lemma 6 *Let $(S, \mathcal{J})$ be an admissible system of bisectors so that the so-called* Regularity Condition, *which is stated for convenience in course of the proof, is satisfied. Then there exists a surface $\Gamma \subset \mathbb{R}^3$ with the following properties:*

1. *The surface Γ is a locally flat 2-sphere.*
2. *For any pair of distinct sites $p, q \in S$, the intersection of $B(p,q)$ and Γ is homeomorphic to S^1; furthermore they intersect properly.*
3. *If $B(p,q,r) \neq \emptyset$, then each connected component of $B(p,q,r)$ intersects Γ at two distinct points; moreover, they intersect properly.*
4. *For any distinct $p, q \in S$, the intersection of $B(p,q)$ and the inner domain of Γ is homeomorphic to $int D^2$ and contains all points of the form $B(p,q,r,s)$ with pairwise distinct sites p, q, r, s in S.*

Proof: It is not hard to see that the Jordan curve theorem and the Schoenflies theorem hold in 2-spheres and in 2-manifolds that are closed and homeomorphic to $\mathbb{R}^2$. This fact will be used implicitly in the proof.

It is known that the 1-point compactification of $\mathbb{R}^3$ is homeomorphic to the 3-sphere $S^3 = \{x \in \mathbb{R}^4 \mid |x| = 1\}$. Let ∞ denote the point at infinity. The bisectors are regarded as being embedded in S^3. Clearly, $B(p,q) \cup \infty$ is homeomorphic to the 2-sphere S^2. By assumption, $B(p,q,r)$, if not empty, consists of biinfinite curves; thus, (after one-point compactifying $\mathbb{R}^3$) if $B(p,q,r) \neq \emptyset$ then each of its connected components together with the point ∞ is homeomorphic to the circle S^1.

The system of bisectors induces on each bisector a finite, disjoint decomposition of it into points, maximally-connected simple curves and relatively open faces each of which is labelled with the set of bisectors containing it, elements having the same label are distinguished by a counting index. See Figure 4.

Let $\{a_i^j\}$ denote the decomposition of the bisector $B(p,q)$, where i is the label and j is the counting index. (Note that since we have n sites, the total number of bisectors is $n(n-1)/2$.

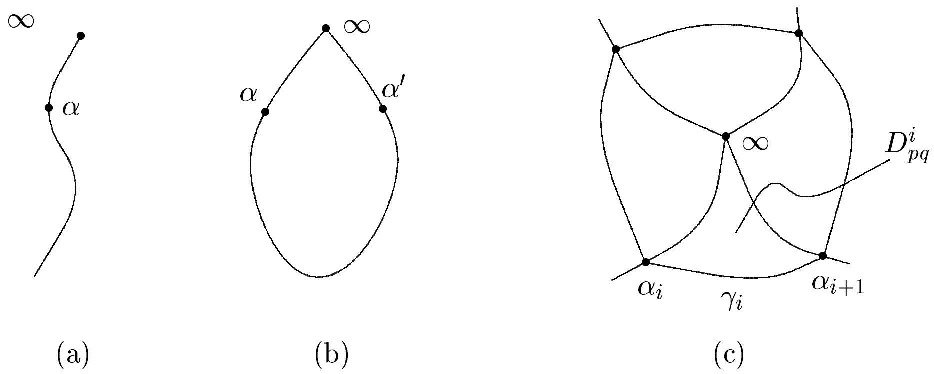

Figure 5: (a) Curve a_i^j is an infinite but not biinfinite curve. (b) Curve a_i^j is a biinfinite curve. (c) Definition of D_{pq} and D^i_{pq}.

So the index i contains at most $n(n-1)/2$ elements.) Thus, $B(p,q)$ is the union of all a_i^j with $pq \in i$ and j runs within an appropriate range.

Consider among the members of $\{a_i^j\}$ meeting ∞ those that are boundary curves of some member of $\{a_i^j\}$. For each such curve a_i^j, if any, we proceed as follows, refer to Figure 5(a) and (b).

If a_i^j is an infinite but not biinfinite curve, then we choose a point α on it so that the subsegment of $a_i^j \cup \infty$ connecting α and ∞ contains no endpoint of any boundary curve of $\{a_i^j\}$. Otherwise, if $a_i^j \cup \infty$ is a simple closed curve, then we choose two points α, α' on it so that the subsegment of $a_i^j \cup \infty$ connecting α, ∞ and α' contains, similarly to above, no endpoint of any curve in $\{a_i^j\}$. Finally, if there is no such curve on $B(p,q)$, then we simply choose a simple closed curve, denoted σ_{pq}, so that σ_{pq} separates $\{a_i^j\}$ and ∞.

Thus, for each bisector $B(p,q)$ we derive from its decomposition a set of arcs, if any, having only ∞ as their common endpoint. For each arc, we rename its endpoint other than ∞ determined as above by α_i, i being a counting index. Knowing that the union of two such arcs is locally flat in $B(p,q) \cup \infty$, for each pair of arcs we may connect their endpoints α_i and α_{i+1} with a simple curve denoted by γ_i so that the relative interior of γ_i does not intersect any curve a_i^j with $pq \in i$. Note that, such a γ_i may already have been constructed in a previous step, but for another bisector. Evidently, the union of all γ_i's with $pq \in i$ is a simple closed curve σ_{pq} in $B(p,q) \cup \infty$. Let D_{pq} be the 2-disk in $B(p,q) \cup \infty$ bounded by σ_{pq} and containing ∞; refer to Figure 5(c).

Furthermore, let D^i_{pq} denote the 2-disks in D_{pq} bounded by a γ_i and the arcs connecting the endpoints of γ_i with ∞. Let

$$H^* = \bigcup_{l,p',q',p'\neq q'} D^l_{p'q'}.$$

We denote by H_i the connected components of

$$H = H^* - \infty.$$

Consider first the case where H consists of only one connected component. The set H^* can be considered as union of 2-disks E_j's so that each E_j is union of $D^k_{p'q'}$'s (for k, p', q' in some index set depending on j) and has one side free from any other 2-disk; refer to Figure 6(a). Applying a result of Kirby [12] each 2-disk E_j is locally flat in S^3. By Flattening Theorem 3.4.1 and Corollary 3.4.1 in [24], the 2-disk E_j is flat in S^3. Thus, for each E_j there exists a locally

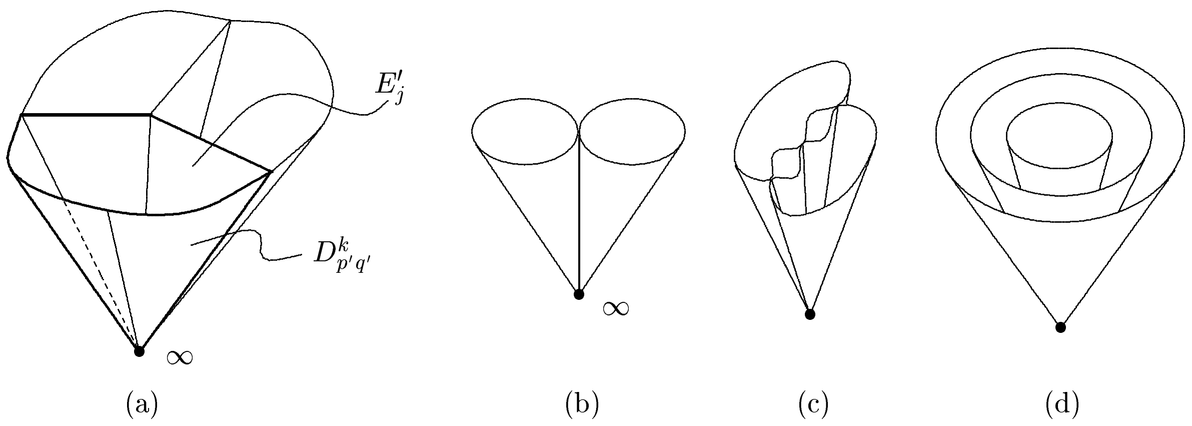

Figure 6: (a) E_j and E'_j. (b), (c) The set H consists of only one connected component. (d) Example: the set H consists of more than one connected components.

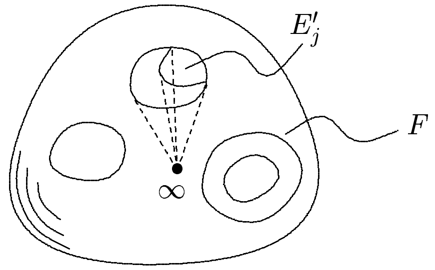

Figure 7: Constructing the sphere Γ.

flat 2-disk E'_j so that the relative interior of E'_j has no point in common with any bisector. If H consists of only one connected component, Figure 6(b) and (c), then the union of all E'_j's and (possibly) of an appropriate surface F, whose existence is ensured by a slight extension of Flattening Theorem 3.4.1 and Corollary 3.4.1 in [24], is a locally flat sphere [12]; refer to Figure 7. This sphere is the sphere Γ to be constructed, since it has all properties required.

However, if H consists of more than one connected component, Figure 6(d), then the situation is much more complicated: Although each H_i itself is not knotted, it is conceivable that different connected components may build links in a very complicated manner, see [23, 24].

Since we don't want to treat the problem in its full generality, let us make the following assumption in order to avoid such cases.

Regularity Condition: There exists a locally flat 2-sphere that intersects the bisectors $B(p,q)$'s exactly in the simple closed curves σ_{pq}'s.

The sphere assumed above is the sphere Γ to be constructed. □

Results in [15, 18] show that systems of bisectors of point sites in general position under semi-algebraic convex distance functions satisfy the Regularity Condition. Moreover, the Regularity Condition is not strong as it appears since the following result holds.

Lemma 7 *If $B(p,q,r) \neq \emptyset$ for any pairwise distinct p,q,r in S then the Regularity Condition is satisfied.*

Proof: Consider any two bisectors $B(p,q)$ and $B(r,s)$. Take any $t \in S$ with $t \notin \{p,q,r,s\}$. By assumption, $B(t,p)$ and $B(t,r)$ intersect $B(p,q)$ and $B(r,s)$, respectively. Also by assumption,

$B(t,p)$ and $B(t,r)$ intersect. Thus the set H, defined in the proof of Lemma 6, consists of only one connected component. $\square$

Let $(R,\mathcal{I})$ be an admissible system of bisectors. A system of bisectors $(S,\mathcal{J})$ is called an *augmented* system of bisectors associated to $(R,\mathcal{I})$ if the following conditions hold

1. The set S is the union of R and a site ∞, and $\mathcal{J}$ consists of bisectors in $\mathcal{I}$ and a surface Γ so that $B(p,\infty) = B(\infty,p) = \Gamma$ for all $p \in R$.
2. The surface Γ has the properties listed in Lemma 6.
3. $p < \infty$ for all $p \in R$, and $D(p,\infty)$ is the inner domain of Γ.

Let $(S,\mathcal{J}) = (R \cup \infty, \mathcal{I} \cup \Gamma)$ be an augmented system of bisectors associated to $(R,\mathcal{I})$. Let $T \subset S$. If $\infty \in T$ then define $B(T) = B(p,\infty) \cap B(T-\infty)$ for any $p \in T-\infty$. Since $B(p,\infty) = \Gamma$ for any $p \in S$, and $B(T-\infty)$ is well-defined by Lemma 4, the definition of $B(T)$ does not depend on the special choice of $p \in T - \infty$.

Theorem 1 *Let $(R,\mathcal{I})$ be an admissible system of bisectors. Then $(S,\mathcal{J})$, the augmented system of bisectors associated to $(R,\mathcal{I})$, is admissible in the following sense:*

1. *The bisector $B(p,q)$ is locally flat for any distinct $p,q \in S$.*
2. *The* extended proper-intersection condition *is satisfied, i.e,*
   (a) *For any four distinct sites $p,q,u,v \in S$ the set $B(p,q) \cap B(u,v)$ is a disjoint union of finitely many points, curves and open subsets of $B(p,q)$ and $B(u,v)$.*
   (b) *The set $B(p,q,r)$ has at most finitely many connected components. If $\infty \notin \{p,q,r\}$ then each of which (if $B(p,q,r) \neq \emptyset$) is a biinfinite curve; otherwise $B(p,q,r)$ is a simple closed curve. Furthermore, for arbitrary $p \neq \infty$ the bisectors $B(r,p)$ and $B(p,q)$ intersect properly.*
   (c) *The set $B(p,q,r,s)$ consists of at most finitely many points. If $\infty \in \{p,q,r,s\}$ then $B(p,q,r,s)$ consists of exactly two distinct points. For arbitrary $p \neq \infty$, $B(s,p)$ and $B(p,q,r)$ intersect properly.*
   (d) *$B(p,q,r,s,t) = \emptyset$.*
3. *For each non-empty subset S' of S hold*
   (a) *For each $p \in S'$, if the set $VR(p,S')$ is non-empty then it is path-connected.*
   (b) *The Voronoi diagram partitions 3-space: $\mathbb{R}^3 = \bigcup_{p \in S'} VR(p,S')$.*

Proof: Claims 1 and 2 follow from the admissibility of $(R,\mathcal{I})$ (refer to Definition 3) and the definition of an augmented system of bisectors.
Claim 3 follows from 3(a) of Definition 3 and the construction of Γ. $\square$

Remark 2 *The partition defined by the Voronoi diagram in 3(b) of the theorem above is disjoint.*

4 Properties of Voronoi regions

From now on we consider only augmented systems of bisectors. Let $(S, \mathcal{J})$ be an augmented system of bisectors. Let T be any subset of S, and $p \in T$ with $p \neq \infty$. We set $G(p,q,T) = B(p,q) \cap \overline{VR(p,T)}$ for any $q \in T - p$. Define $T^* = \{q \in T \mid G(p,q,T) \text{ contains } x\}$, for any $x \in \operatorname{bd} VR(p,T)$. The set T^* depends on T, site p and point x.

Lemma 8 *1. $bd VR(p,T) = \bigcup_{q \in T-p} G(p,q,T)$.*

2. $G(p,q,T) = B(p,q) \cap (\bigcap_{q' \in T - \{p,q\}} \overline{R(p,q')})$.

Proof: 1.

$$\begin{aligned}
\operatorname{bd} VR(p,T) &= \overline{VR(p,T)} - \operatorname{int} VR(p,T) \\
\text{by (1),} \quad &= \overline{VR(p,T)} - \operatorname{int} \bigcap_{q \in T-p} R(p,q) \\
\text{by (2),} \quad &= \overline{VR(p,T)} - \bigcap_{q \in T-p} \operatorname{int} R(p,q) \\
&= \overline{VR(p,T)} - \bigcap_{q \in T-p} D(p,q) \\
&= \overline{VR(p,T)} \cap (\bigcup_{q \in T-p} (D(q,p) \cup B(p,q))) \\
&= \bigcup_{q \in T-p} \overline{VR(p,T)} \cap (D(q,p) \cup B(p,q)) \\
&= \bigcup_{q \in T-p} (\overline{VR(p,T)} \cap D(q,p)) \cup (\overline{VR(p,T)} \cap B(p,q)).
\end{aligned}$$

Since in the relation above $\overline{VR(p,T)} \cap D(q,p) = \emptyset$, Claim 1 follows.
2.

$$\begin{aligned}
\overline{VR(p,T)} \cap B(p,q) &= (\bigcap_{q' \in T-p} \overline{R(p,q')}) \cap B(p,q) \\
&= (\bigcap_{q' \in T-\{p,q\}} \overline{R(p,q')}) \cap \overline{R(p,q)} \cap B(p,q).
\end{aligned}$$

The claim follows from the identity above by use of $\overline{R(p,q)} \cap B(p,q) = B(p,q)$. $\square$

We will need the following simple result, which we state without proof.

Lemma 9 *Let X be an open set. Then, for any set Y and Z,*

$$X \cap Y = X \cap Z \Rightarrow X \cap \overline{Y} = X \cap \overline{Z}.$$

Lemma 10 *Let $x \in bd VR(p,T)$. Then for all sufficiently small open neighbourhoods W of x*

1. *$W \cap VR(p,T) = W \cap VR(p, T^* \cup p)$.*
2. *$W \cap \overline{VR(p,T)} = W \cap \overline{VR(p, T^* \cup p)}$.*

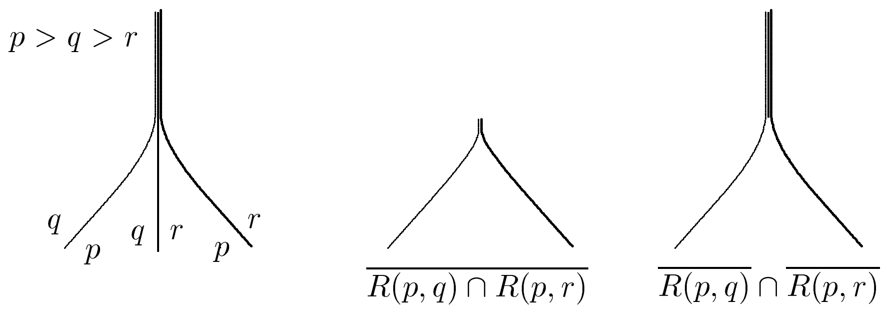

Figure 8: If the bisectors don't cross properly then we may have $\overline{R(p,q) \cap R(p,r)} \neq \overline{R(p,q)} \cap \overline{R(p,r)}$.

3. $W \cap \mathit{int}\, VR(p,T) = W \cap \mathit{int}\, VR(p, T^* \cup p)$.

4. $W \cap \mathit{bd}\, VR(p,T) = W \cap \mathit{bd}\, VR(p, T^* \cup p)$.

Proof:

1. $W \cap VR(p,T) = \overbrace{(W \cap \bigcap_{q \in T^*} R(p,q))}^{\alpha} \cap \overbrace{(W \cap \bigcap_{q \in T - T^* - p} R(p,q))}^{\beta}$. For every $q \in T - T^* - p$, the bisector $B(p,q)$ does not contain x. Thus, applying the Jordan-Brouwer Separation Theorem to the surface $B(p,q)$, for all $q \in T - T^* - p$, we can choose open set W sufficiently small so that $W \subset R(p,q)$ for any $q \in T - T^* - p$. Therefore $\beta = W$.
2. Claim 2 follows from Claim 1 by applying Lemma 9.
3. Claim 3 follows from Claim 1 using the identity $\mathrm{int}\,(\alpha \cap \beta) = \mathrm{int}\,\alpha \cap \mathrm{int}\,\beta$.
4. By Claim 1 of Lemma 8 we have

$$\begin{aligned} W \cap \mathrm{bd}\, VR(p,T) &= W \cap \bigcup_{q \in T-p} G(p,q,T) \\ &= \bigcup_{q \in T-p} W \cap B(p,q) \cap \overline{VR(p,T)}. \end{aligned}$$

For any $q \in T - T^* - p$, the bisector $B(p,q)$ does not pass through x. Thus, for every open set W small enough, we have $W \cap B(p,q) = \emptyset$. Therefore,

$$\begin{aligned} W \cap \mathrm{bd}\, VR(p,T) &= \bigcup_{q \in T^*} W \cap B(p,q) \cap \overline{VR(p,T)} \\ \text{by Claim 2,} \qquad &= \bigcup_{q \in T^*} W \cap B(p,q) \cap \overline{VR(p, T^* \cup p)} \\ &= W \cap \mathrm{bd}\, VR(p, T^* \cup p). \end{aligned}$$

$\square$

Lemma 11 *Let $T \subset S$. Then,* $\overline{\bigcap_{q \in T-p} R(p,q)} = \bigcap_{q \in T-p} \overline{R(p,q)}$.

Proof: We need only to show that $\bigcap_{q \in T-p} \overline{R(p,q)} \subset \overline{\bigcap_{q \in T-p} R(p,q)}$. Let $x \in \bigcap_{q \in T-p} \overline{R(p,q)}$. Using Claims 1 and 2 of Theorem 1 we can easily show that there exists a sequence (x_n) in $\bigcap_{q \in T-p} D(p,q)$

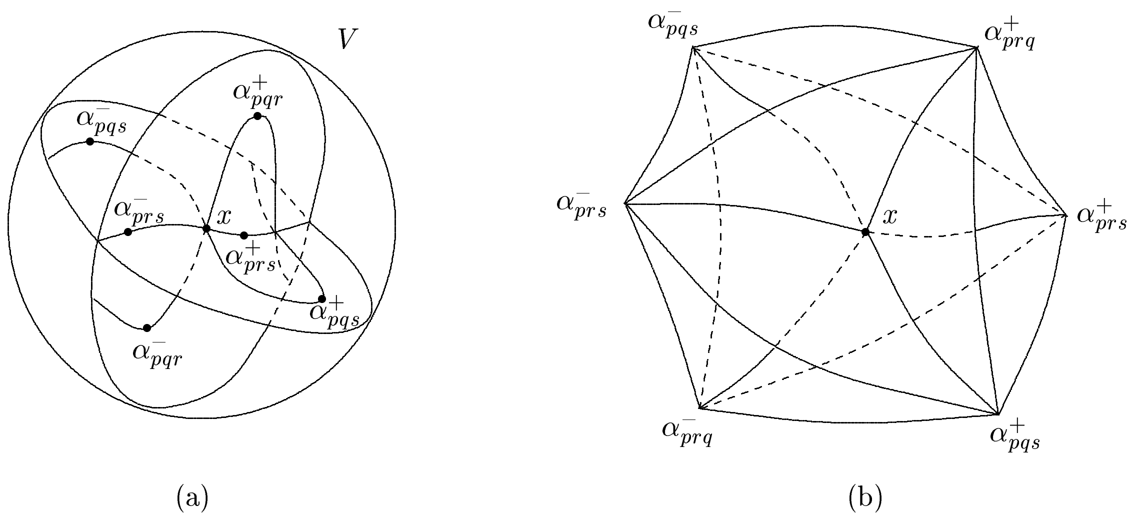

Figure 9: (a) The intersection of the bisectors in the neighbourhood V of x. (b) Construction of the required open set U.

converging to x. Thus $x \in \overline{\bigcap_{q \in T-p} R(p,q)}$. $\square$

Remark 3 *In the lemma above, the condition that the bisectors intersect properly is essential, as shown in Figure 8.*

From the classification theorem for 1-manifold we easily derive the following lemma.

Lemma 12 *Each compact, connected 1-manifold without boundary is homeomorphic to S^1.*

Let $T \subset S$ with $\infty \in T$. Let $p \in T$ with $p \neq \infty$, and $q \in S - T$.

Lemma 13 *We have*

1. *$bd\, VR(p,T) \cap B(p,q)$ is either empty or a compact 1-manifold without boundary.*
2. *$bd\, VR(p, T \cup q) = (bd\, VR(p,T) \cap \overline{R(p,q)}) \cup (B(p,q) \cap \overline{VR(p,T)})$.*
3. *$bd\, VR(p,T) \cap B(p,q) \subset bd\, VR(p, T \cup q)$.*

Proof: 1. Assume that $M = \mathrm{bd}\, VR(p,T) \cap B(p,q) \neq \emptyset$. Clearly M is closed. Since $\infty \in T$ we have

$$
\begin{aligned}
M &\subset \overline{VR(p,T)} \subset \overline{R(p,\infty)} && (14)\\
&\subset \text{interior domain of } \Gamma. && (15)
\end{aligned}
$$

Thus M is bounded. So M is compact. We will show that for any $x \in M$ there is an open set U containing x so that $U \cap M$ is homeomorphic to $I = (0,1)$.

By Claim 4 of Lemma 10, for all sufficiently small open neighbourhoods W of x we have

$$W \cap \mathrm{bd}\, VR(p,T) = W \cap \mathrm{bd}\, VR(p, T^* \cup p).$$

Intersecting both sides of the identity above with $B(p,q)$, we obtain

$$W \cap (\mathrm{bd}\, VR(p,T) \cap B(p,q)) = W \cap (\mathrm{bd}\, VR(p,T^* \cup p) \cap B(p,q)) \quad (16)$$
$$W \cap M = W \cap (\mathrm{bd}\, VR(p,T^* \cup p) \cap B(p,q)). \quad (17)$$

Equality 17 enables us to explore the local topology of the set M as follows. Using Claim 2(d) of Theorem 1, we conclude that the set T^* has at most two elements. If T^* has exactly one element, then the claim follows from Equality 17, by using Claim 2(b) of Theorem 1. Otherwise T^* has exactly two elements, say r and $s \in T - p$. Since $B(p,r)$ and $B(p,s)$ intersect properly, there exists a neighbourhood V of x with the properties given in Definition 1, see Figure 9(a). Since $B(p,q,r)$ intersects $B(p,s)$ properly, the connected component of $B(p,q,r)$ passing through x is partitioned by x into two connected components lying locally around x in disjoint complementary domains of $B(p,s)$.

Similarly, the connected components of $B(p,q,s)$ going through x also lie locally around x in disjoint complementary domains of $B(p,r)$. Let α^+_{pqr} and α^-_{pqr} be points of $B(p,q,r)$ sufficiently close to x so that $\alpha^-_{pqr} \subset D(s,p)$, $\alpha^+_{pqr} \subset D(p,s)$, and the arcs $x\alpha^-_{pqr}$ and $x\alpha^+_{pqr}$ do not contain any point of $B(p,q,r,s)$ other than x, see Figure 9(a). We analogously define the points α^+_{pqs} and α^-_{pqs} of $B(p,q,s)$.

To construct the required neighbourhood U of x, we construct its boundary $\mathrm{bd}\, U$ so that it is homeomorphic to the 2-sphere. Indeed, $\mathrm{bd}\, U$ is composed of the outer curved triangles as shown in Figure 9(b). To achieve this we may proceed exactly like in the proof of Lemma 6; the main tool for the proof is the Flattening Theorem. We omit the details. Clearly, $U \cap M$ is the interior of the arc $x\alpha^+_{pqr} \cup x\alpha^+_{pqs}$ which is homeomorphic to an open interval.
2. By Lemma 8,

$$\begin{aligned}
\mathrm{bd}\, VR(p,T\cup q) &= \bigcup_{q' \in T \cup q - p} (B(p,q') \cap \overline{VR(p,T\cup q)}) \\
&= \bigcup_{q' \in T \cup q - p} (B(p,q') \cap \overline{VR(p,T)} \cap \overline{R(p,q)}), \quad \text{by Lemma 11} \\
&= \bigcup_{q' \in T - p} (B(p,q') \cap \overline{VR(p,T)} \cap \overline{R(p,q)}) \\
&\quad \bigcup (B(p,q) \cap \overline{VR(p,T)} \cap \overline{R(p,q)}) \\
&= (\mathrm{bd}\, VR(p,T) \cap \overline{R(p,q)}) \cup (B(p,q) \cap \overline{VR(p,T)}), \quad \text{since } B(p,q) \subset \overline{R(p,q)}.
\end{aligned}$$

3.

$$\begin{aligned}
\mathrm{bd}\, VR(p,T) \cap B(p,q) &\subset \mathrm{bd}\, VR(p,T) \cap \overline{R(p,q)}, \quad \text{since } B(p,q) \subset \overline{R(p,q)} \\
&\subset \mathrm{bd}\, VR(p,T\cup q), \quad \text{from Claim 2 of this lemma.}
\end{aligned}$$

□

Lemma 14 *The set $\mathrm{bd}\, VR(p,T) \cap B(p,q)$ is either empty or consists of finitely many disjoint simple closed curves.*

Proof: From Lemma 12 and 13 it follows that each connected component of the set $M = \mathrm{bd}\, VR(p,T) \cap B(p,q)$ is a simple closed curve; thus, M consists of disjoint simple closed curves. Let N_i, for $1 \leq i \leq k$, be the disjoint simple closed curve whose union yields M. Lemma 8 implies that for each i the simple closed curve N_i is the union of parts of $B(p,q) \cap B(p,q')$, where $q' \in T$. If some N_i is exactly $B(p,q) \cap B(p,q')$ for some q' in T then, by Claim 2(b) of Theorem 1, the site q' must be ∞. Thus, at most one of the simple closed curves N_i's is the set $B(p,q) \cap \Gamma$; and each of the remaining simple closed curves is the union of parts of $B(p,q) \cap B(p,q')$ for at least two distinct sites $q' \in T$. Clearly the endpoints of such parts are contained in $B(p,q) \cap B(p,q') \cap B(p,q'')$ which is by Lemma 5 the set $B(p,q,q',q'')$. By Claim 2(c) of Theorem 1, the number of points in $B(p,q,q',q'')$ is finite. Consequently, the number of the simple closed curves N_i's is also finite. □

Let N_i with $1 \leq i \leq k$ be defined as in the proof of Lemma 14. For each i let K_i denote the 2-disks in $B(p,q)$ with boundary N_i; that is, K_i is the union of N_i and the interior domain of N_i.

Lemma 15 *For all distinct i,j with $1 \leq i,j \leq k$ the 2-disks K_i and K_j are disjoint.*

Proof: If $k = 1$ then we are finished. Assume that $k \geq 2$. For contradiction, suppose without loss of generality that $K_1 \cap K_2 \neq \emptyset$. By Lemma 14, in $B(p,q)$, either N_1 is a subset of the interior domain of N_2 or N_2 is a subset of the interior domain of N_1; without loss of generality suppose the first case. In the proof of Lemma 14 we have shown that at most one of the simple closed curves is the set $B(p,q) \cap \Gamma$. Thus only two cases are possible:
Case 1: $N_1 \neq B(p,q) \cap \Gamma$. Let $x \in N_1 - \Gamma$. Let $q' \in T$ so that $G(p,q')$ contains x; clearly $q' \neq \infty$. Since $B(p,q,q')$ is a biinfinite curve containing the point x of the bounded interior domain of N_2, it must intersect N_2. Let $y \in B(p,q,q') \cap N_2$. Since $B(p,q)$ and $B(p,q')$ intersect properly (Claim 2 of Theorem 1), the simple closed curve N_2 would have y as a boundary point. This is a contradiction.
Case 2: $N_1 = B(p,q) \cap \Gamma$. We will prove that this case is also impossible. By the same argument as above, N_2 is the union of parts of $B(p,q) \cap B(p,q')$ for at least two distinct sites $q' \in T$. Let q'_1 and q'_2 denote two such sites so that $B(p,q) \cap B(p,q'_1)$ and $B(p,q) \cap B(p,q'_1)$ have a point in common that lies in N_2. Clearly the intersection point lies in $B(p,q,q'_1,q'_2)$. Since N_2 lies in the exterior domain of N_1, it follows that $B(p,q,q'_1,q'_2)$ is contained in the exterior domain of N_1; consequently $B(p,q,q'_1,q'_2)$ is contained in the exterior domain of Γ. This contradicts Claim 4 of Lemma 6. □

Note that we are still assuming that $p \in S - \infty$.

Theorem 2 *The set bd$VR(p,S)$ is homeomorphic to the sphere S^2.*

Proof: We prove the assertion using an induction argument on $|T|$ with $T \subset S$ and $p, \infty \in T$. Obviously, $\operatorname{bd} VR(p,T)$ is homeomorphic to the 2-sphere S^2 if $T = \{p, \infty\}$. Assume that the assertion is true for $|T| = m$; that is, $\operatorname{bd} VR(p,T)$ is homeomorphic to S^2. We will prove the assertion for $|T| = m+1$. Let $q \in S - T$. By Lemma 13 we have

$$\operatorname{bd} VR(p, T \cup q) = \overbrace{(B(p,q) \cap \overline{VR(p,T)})}^{\alpha} \cup \overbrace{((\operatorname{bd} VR(p,T)) \cap \overline{R(p,q)})}^{\beta}.$$

We will prove that the set α consists of k pairwise disjoint 2-disks, and the set β is homeomorphic to a sphere with k holes so that each 2-disk fits precisely to one hole along their intersection which is exactly their boundary. Using the concept of connected sum from the algebraic topology [20] the claims follows.

We may write α as

$$\alpha = \overbrace{(\operatorname{bd} VR(p,T) \cap B(p,q))}^{\alpha_1} \cup \overbrace{(B(p,q) \cap \operatorname{int} VR(p,T))}^{\alpha_2}.$$

By Lemma 14 the set α_1 consists of simple closed curves $N_i \subset B(p,q)$ for $1 \leq i \leq k$. Lemma 15 and the proper-intersection condition imply that the set α_2 either consists of open disks, denoted by C_i, whose boundaries are N_i for $1 \leq i \leq k$, or is the unbounded subset of $B(p,q)$ having the simple closed curves N_i for $1 \leq i \leq k$ as its boundary. The latter case is impossible, since $VR(p,T)$ is bounded. Thus, α is the union of pairwise disjoint disks $C_i \cup N_i$ with $1 \leq i \leq k$.

We consider now the set β; it may be written as

$$\beta = \overbrace{(\operatorname{bd} VR(p,T)) \cap (B(p,q))}^{\beta_1} \cup \overbrace{(\operatorname{bd} VR(p,T) \cap D(p,q))}^{\beta_2}.$$

By Lemma 14 the set β_1 consists of simple closed curves $N_i \subset \operatorname{bd} VR(p,T)$, for $1 \leq i \leq k$. By Lemma 15, the simple closed curves N_i's bound a connected, open subset J of $\operatorname{bd} VR(p,T)$ (thus, J is a sphere with k holes). Let D_i denote the open disks that are the connected components of $\operatorname{bd} VR(p,T)$ minus the closure of J. The proper-intersection condition implies that the set β_2 is either (1) the disks D_i with $i = 1, \ldots, k$ or (2) the set J; but (1) is impossible by the following argument.

Case 1: The set β consists of pairwise disjoint disks $D_i \cup N_i$, for $1 \leq i \leq k$. Thus, the set $\operatorname{bd} VR(p, T \cup q)$ are the connected sums $(C_i \cup N_i)\#(D_i \cup N_i)$ that arise by pasting the disjoint, 2-disks $C_i \cup N_i$ and $D_i \cup N_i$ along their boundaries N_i, for $1 \leq i \leq k$. It follows that the sets $(C_i \cup N_i)\#(D_i \cup N_i)$ are pairwise disjoint spheres. Since $VR(p, T \cup q)$ is bounded, its interior must be the inner domains of these spheres. But this is impossible: Claim 3(a) of Theorem 1 says that the set $VR(p, T \cup q)$ is path-connected.

Case 2: The set β is a sphere with k holes. Clearly the connected sum $\beta \# \bigcup (D_i \cup N_i)$ of β and the 2-disks $D_i \cup N_i$ with respect to their boundaries N_i, for $1 \leq i \leq k$, is a sphere. □

Let Σ be a $(d-1)$-sphere in $\mathbb{R}^d$. Recall that, by the Jordan-Brouwer Seperation Theorem, the set Σ seperates $\mathbb{R}^d$ into exactly two disjoint domains of which Σ is the common boundary. The following theorem may be found in [24] (Theorem 1.8.2 in that book).

Theorem 3 *(Generalized Schoenflies Theorem) Let Σ be a locally flat $(d-1)$-sphere in $\mathbb{R}^d$. Then the closure of the bounded, complementary domain of Σ in $\mathbb{R}^d$ is homeomorphic to the closed d-ball D^d.*

Note that in the generalized Schoenflies theorem for $d = 2$ the condition "locally flat" can be dropped (The Schoenflies theorem). In $\mathbb{R}^3$ there exists a 2-sphere (*the Alexander horned sphere*, [1]) so that one of the complementary domains is not simply connected.

Theorem 4 *The set $\overline{VR(p,S)}$ is homeomorphic to the 3-ball D^3.*

Proof: The definition of admissible system of bisectors requires that each bisector is locally flat. Applying a result of Kirby (Theorem on page 614 in [12]), it follows that $\operatorname{bd} VR(p,S)$ is also locally flat. By Theorem 2, the set $\operatorname{bd} VR(p,S)$ is a 2-sphere. Thus, by Theorem 3, the set $\overline{VR(p,S)}$ is homeomorphic to D^3. □

5 Properties of abstract Voronoi diagrams

Lemma 16 *Let $T \subset S$ with $|T| \geq 2$. Then*

$$\bigcap_{a \in T} \overline{VR(a,S)} \subset B(T).$$

Proof: Take $w \in T - \infty$. Clearly,

$$\bigcap_{a \in T} \overline{VR(a,S)} = (\bigcap_{a \in T-w} \overline{VR(a,S)}) \cap \overline{VR(w,S)}.$$

Using the relations $VR(a,S) \subset R(a,w)$ and $VR(w,S) \subset VR(w,T)$ we obtain

$$\begin{aligned}
\bigcap_{a \in T} \overline{VR(a,S)} &\subset (\bigcap_{a \in T-w} \overline{R(a,w)}) \cap \overline{VR(w,T)} \\
&\subset \bigcap_{a \in T-w} \overline{R(a,w)} \cap \bigcap_{a \in T-w} \overline{R(w,a)}, \qquad \text{since } VR(w,T) \subset R(w,a) \\
&\subset \bigcap_{a \in T-w} B(a,w) \\
&\subset B(T), \qquad \text{by Lemma 5.}
\end{aligned}$$

□

For each $x \in \mathbb{R}^3$ let $N(x)$ denote the cardinality of the set $\{p \in S \mid x \in \overline{VR(p,S)}\}$.

Lemma 17 *$N(x) \leq 4$, and $N(x) = 1$ if and only if $x \in int VR(p,S)$ for exactly one $p \in S$.*

Proof: Let $T = \{p \in S \mid x \in \overline{VR(p,S)}\}$. Applying Lemma 16, we see that $x \in B(T)$. By Claim 2(d) of Theorem 1, we have

$$B(T) = \emptyset$$

for any $T \subset S$ with $|T| > 4$. Thus $N(x) \leq 4$.

Let $x \in \mathbb{R}^3$ such that $N(x) = 1$. By assumption, $x \in \overline{VR(p,S)}$ for exactly one p in S. If $x \in \operatorname{int} VR(p,S)$ then we are done. Suppose for contradiction that $x \in \operatorname{bd} \overline{VR(p,S)} = \operatorname{bd} VR(p,S)$. By Claim 3(b) of Theorem 1 and since $|S|$ is finite, there exists a site q in S such that each neighbourhood of x contains a point x' of $VR(q,S)$. Thus, $x \in \overline{VR(q,S)}$ contradicting $N(x) = 1$.

Conversely, let $x \in \operatorname{int} VR(p,S)$. Choose a sufficiently small neighbourhood $U \subset VR(p,S)$ of x such that U does not intersect $\operatorname{bd} VR(p,S)$. Since the decomposition of $\mathbb{R}^3$ into Voronoi regions is disjoint, U contains only points of $VR(p,S)$. Thus, $N(x) = 1$. □

We set

$$\begin{aligned}
F &= \{x \in \mathbb{R}^3 \mid N(x) = 2\}, \\
E &= \{x \in \mathbb{R}^3 \mid N(x) = 3\}, \\
V &= \{x \in \mathbb{R}^3 \mid N(x) = 4\}.
\end{aligned}$$

Lemma 18 *$V(S) = F \cup E \cup V$ (disjoint).*

Proof: The claim follows immediately from the definition of $V(S)$ and Lemma 17. □

We call a maximal connected subset of F, E, V of $V(S)$ a *face, edge, vertex*, respectively.

Lemma 19
1. *Let f be a face of $V(S)$. Then there are exactly two distinct sites $p, q \in S$ so that for each $x \in f$ we have $x \in \overline{VR(p,S)} \cap \overline{VR(q,S)}$.*
2. *Let e be any edge of $V(S)$. Then there are exactly three distinct sites p, q and r in S so that for each $x \in e$ we have $x \in \overline{VR(p,S)} \cap \overline{VR(q,S)} \cap \overline{VR(r,S)}$.*

3. *Each vertex v of $V(S)$ is a point. There exist exactly four distinct sites p, q, r and s in S such that with $T = \{p, q, r, s\}$*

$$v \in \bigcap_{w \in T} \overline{VR(w, S)}.$$

Proof: 1. Let $x \in f$. By definition, there are exactly two sites p and q in S so that $x \in \overline{VR(p, S)} \cap \overline{VR(q, S)}$. Thus, we can find a neighbourhood N of x that does not intersect $\overline{VR(r, S)}$ for any $r \in S$ with $r \neq p, q$. Clearly, for every $y \in f \cap N$, we have $y \in \overline{VR(p, S)} \cap \overline{VR(q, S)}$. (Note that by Lemma 16 we obtain $f \cap N = B(p, q) \cap N$).
2. The proof is similar to that of Claim 1.
3. By arguments similar to those for proving Claim 1, and by Lemma 16, there are exactly four distinct sites p, q, r, s of S so that we have $v \in \overline{VR(p, S)} \cap \overline{VR(q, S)} \cap \overline{VR(r, S)} \cap \overline{VR(q, S)} \subset B(p, q, r, s)$. Now observe that $B(p, q, r, s)$ is a finite set of points. $\square$

Lemma 20 *Let f be a face of $V(S)$ incident to $\overline{VR(p, S)}$ and $\overline{VR(q, S)}$. Then,*

1. *If $p, q \neq \infty$, then f is homeomorphic to the open disk.*

2. *Each connected component of the boundary of f is homeomorphic to S^1.*

Proof: (Sketch) The union of the faces of $V(S)$ incident to $\overline{VR(p, S)}$ and $\overline{VR(q, S)}$, denoted $F(p, q)$, can be written as

$$F(p, q) = \bigcap_{q' \in S, q' \neq p, q} (B(p, q) \cap D(p, q')).$$

We prove the claim by induction, using the expression above. We start with three sites $\{p, q, \infty\}$; the claim is true obviously. The induction step can be proved essentially using the fact that the bisectors of any three sites intersect properly and the connected components of $B(p, q, r)$ and $B(p, q, s)$ are biinfinite curves on $B(p, q)$ that intersect each other properly, for any $r, s \in S$ with $r, s \neq p, q$. $\square$

Lemma 21 *Locally the abstract and the ordinary Euclidian Voronoi diagram for point sites in general position look identical at points on their edges. To be precise, let e be any edge of $V(S)$. Assume that e is incident to the Voronoi regions of p, q, r. Then for any point $x \in e$ there exists a neighbourhood U of x with the following properties, refer to Figure 11(b):*

1. *$U \cong D^3$.*

2. *$U \cap \overline{VR(w, S)} \cong D^3$ for each $w \in \{p, q, r\}$.*

3. *$U \cap \overline{VR(w, S)} \cap \overline{VR(w', S)}$ is homeomorphic to the 2-disk for distinct $w, w' \in \{p, q, r\}$.*

4. *$U \cap \overline{VR(p, S)} \cap \overline{VR(q, S)} \cap \overline{VR(r, S)}$ is homeomorphic to the closed interval $[0, 1]$.*

5. *$U \cap bd VR(w, S)$ and $bd U \cap \overline{VR(w, S)}$ are homeomorphic to the 2-disk; and $bd U \cap bd VR(w, S)$ is a simple closed curve for every $w \in \{p, q, r\}$.*

Globally, the set e is homeomorphic either to the open interval $(0, 1)$ or to S^1. In the latter case, the boundary of e consists of two distinct points, refer to Figure 10.

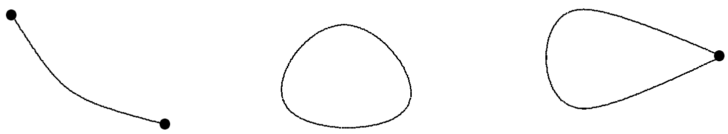

(a) (b) (c)

Figure 10: (a), (b) Possible Voronoi edges. (c) Impossible.

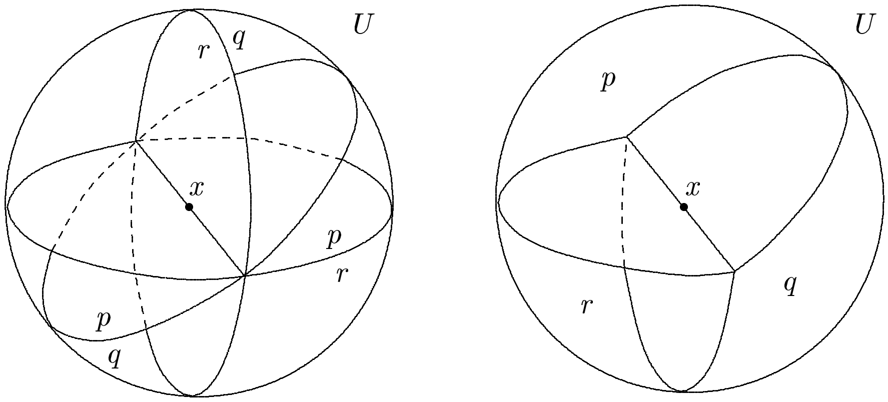

(a) (b)

Figure 11: The shape of the abstract Voronoi diagram in a small neighbourhood of points on Voronoi edges.

Proof: (Sketch) Let $x \in e$. By Lemma 10 we can restrict our local analysis of $V(S)$ around x to that of the Voronoi diagram of the sites p, q, r. Since $B(p,q)$ and $B(p,r)$ intersect properly, by Definition 1, there exists a neighbourhood V of x having the properties described there, see Figure 11(a). Let $J = V \cap B(p,q) \cap B(p,r)$. By Lemma 3 we have

$$B(q,r) \cap B(p,q) = B(q,r) \cap B(p,r) = B(p,q) \cap B(p,r).$$

By the above identities we see that in V the bisector $B(q,r)$ intersects $B(p,q)$ and $B(p,r)$ exactly in J. The bisectors $B(p,q)$ and $B(p,r)$ partition the 3-ball V into four parts, call them V_{pp}, V_{pq}, V_{pr} and V_{qr}.

Moreover, since $B(q,r)$ crosses $B(p,q)$ and $B(p,r)$ properly, it follows that $B(q,r)$ either lie in $V_{pp} \cup V_{qr}$ or in $V_{pq} \cup V_{pr}$. Proceeding like in the proof of Lemma 6 we may choose a neighbourhood U of x so that the properties as shown in Figure 11(b) hold (we omit a formal description). By an easy case analysis, we see that only the configuration shown in Figure 11(b) is possible.

The global behaviour of e follows from Claim 4 and by additionally considering the next lemma (Lemma 22). (We put the claim here, in this lemma, for convenience.) □

Lemma 22 *The abstract and the ordinary Euclidian Voronoi diagram of point sites in general position look like the same locally at their vertices. To be precise, assume that the vertices are incident to the closures of the Voronoi regions of p, q, r, s. Let v be any point of these vertices. Then there exists a neighbourhood U of v with the following properties, see Figure 12:*

1. $U \cong D^3$.

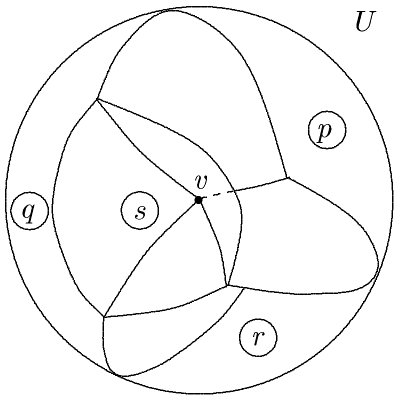

Figure 12: The shape of the abstract Voronoi diagram in a small neighbourhood of a Voronoi vertex.

2. *$U \cap \overline{VR(w,S)} \cong D^3$ for every $w \in T$, where $T = \{p,q,r,s\}$.*

3. *$U \cap \overline{VR(w,S)} \cap \overline{VR(w',S)}$ is homeomorphic to the 2-disk for distinct $w, w' \in T$.*

4. *$U \cap \overline{VR(w,S)} \cap \overline{VR(w',S)} \cap \overline{VR(w'',S)}$ is homeomorphic to the interval $[0,1]$ for distinct $w, w', w'' \in T$.*

5. *$U \cap \bigcap_{w \in T} \overline{VR(w,S)} = v$.*

6. *$U \cap bd\, VR(w,S)$ and $bd\, U \cap \overline{VR(w,S)}$ are homeomorphic to the 2-disk; and $bd\, U \cap bd\, VR(w,S)$ is homeomorphic to S^1 for all $w \in T$.*

Proof: (Sketch) By Lemma 10 we know that locally at v the sets $V(S)$ and $V(\{p,q,r,s\})$ look identical. Thus we need only to investigate how $V(\{p,q,r,s\})$ locally at v looks like.

We regard $V(\{p,q,r,s\})$ as araising from $V(\{p,q,r\})$ after inserting site s to $\{p,q,r\}$. The proof of Claim 2 of Lemma 21 shows how an abstract Voronoi diagram locally looks like at points of its edges. Let W be a neighbourhood of point v now lying on an edge of $V(\{p,q,r\})$, W having the properties listed in that claim.

First we want to see how $B(p,s)$ intersects $W \cap V(\{p,q,r\})$. Since $B(p,q,r)$ intersects $B(p,s)$ properly at v, in W the set $B(p,q,r)$ is divided by v into two parts lying in disjoint complementary domains of $B(p,s)$. Let α^+_{pqr} and α^-_{pqr} be the intersections of $W \cap B(p,q,r)$ with bd W, where the naming is such that $\alpha^+_{pqr} \in D(p,s)$ and $\alpha^-_{pqr} \in D(s,p)$. Then, after inserting site s, the "local edge" $\alpha^-_{pqr}\alpha^+_{pqr} \subset B(p,q,r)$ is "shortened" by $\overline{VR(s,\{p,q,r,s\})}$ to yield the "local edge" $v\alpha^+_{pqr}$ of $V(\{p,q,r,s\})$.

Arguing in a similar way, eventually after making the neighbourhood W smaller, (the points above remain defined), we obtain the new local edges of $V(\{p,q,r,s\})$, namely $v\alpha^+_{pqs} \subset B(p,q,s)$, $v\alpha^+_{prs} \subset B(p,r,s)$, and $v\alpha^+_{qrs} \subset B(q,r,s)$ after inserting site s.

Now as in the proof of Lemma 6 we can choose a neighbourhood U of v so that the configuration described in Figure 12 holds, as required. □

6 Complexity of the abstract Voronoi diagram

Let $(S, \mathcal{J})$ be an admissible system of bisectors. Denote the number of the connected components of $B(p,q,r)$ by $|B(p,q,r)|$, and the number of the connected components of $B(p,q,r,s)$ by

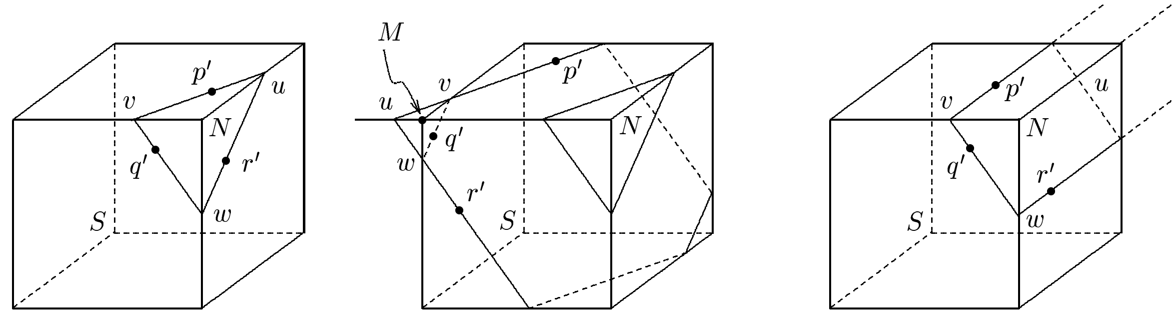

(a) (b) (c)

Figure 13: With respect to the scan plane: (a) Vertices N and S are extremal vertices. (b) Vertex M is a nonextremal vertex. (c) Vertex N is an extremal vertex.

$|B(p,q,r,s)|$ for any distinct $p,q,r,s \in S$. Denote by $\mathcal{C}_1(\mathcal{J})$, or $\mathcal{C}_1$ for short, the range of the values that $|B(p,q,r)|$ takes for any distinct $p,q,r \in S$. Similarly, denote by $\mathcal{C}_2(\mathcal{J})$, or simply $\mathcal{C}_2$, the range of the values that $|B(p,q,r,s)|$ takes for any distinct $p,q,r,s \in S$. We call the pair $\mathcal{C}_1$ and $\mathcal{C}_2$ the *intersection characteristic* of $(S,\mathcal{J})$.

Denote by k_1 and k_2 the maxima of $\mathcal{C}_1$ and $\mathcal{C}_2$, respectively. Let $n = |S|$. A trivial upper bound on the combinatorial complexity of $V(S)$ is $O(k_1k_2n^4)$. This can be seen easily as follows. Consider the Voronoi faces separating the Voronoi regions of any sites p and p'. These faces are the maximal connected subsets of the intersection of all the sets $B(p,p') \cap D(p,p'')$, where $p'' \neq p,p'$. Observe that $B(p,p') \cap D(p,p'')$ is bounded by $B(p,p',p'')$, which consists of at most k_1 biinfinite curves. Furthermore, $B(p,p',p'')$ and $B(p,p',p''')$ intersect in at most k_2 points. Thus, the number of vertices of the faces separating the Voronoi regions of p and p' is $O(k_1k_2n^2)$.

By the Euler's formula, the number of edges and faces incident to $VR(p,S)$ and $VR(p',S)$ is still $O(k_1k_2n^2)$. Since the boundary of any Voronoi region is built up by at most $n-1$ bisectors, the complexity of any Voronoi region is $O(k_1k_2n^3)$. Moreover $V(S)$ has at most n Voronoi regions. It follows that the complexity of $V(S)$ is $O(k_1k_2n^4)$.

A simple way to construct the abstract Voronoi diagram is to compute it incrementally. That is, starting with the Voronoi diagram of some 4 sites $\{p,q,r,\infty\}$, we insert the remaining sites to the current set of sites one by one, and update the current Voronoi diagram correspondingly whenever a new site is inserted.

Assume a model of computation in which infinite precision real arithmetic is possible and standard operations—e.g., computing the intersections of two curves on a surface, or testing if a point lies in a halfspace bounded by a surface homeomorphic to the plane—take $O(1)$ time. Then, clearly, the described algorithm has polynomial time complexity.

7 Examples

7.1 The simplicial Voronoi diagram

Let T denote a simplex in 3-space, i.e. a tetrahedron, that contains the origin of the coordinate system in its interior domain. We consider the convex distance function admitting T as the unit "sphere": the simplex or simplicial distance function.

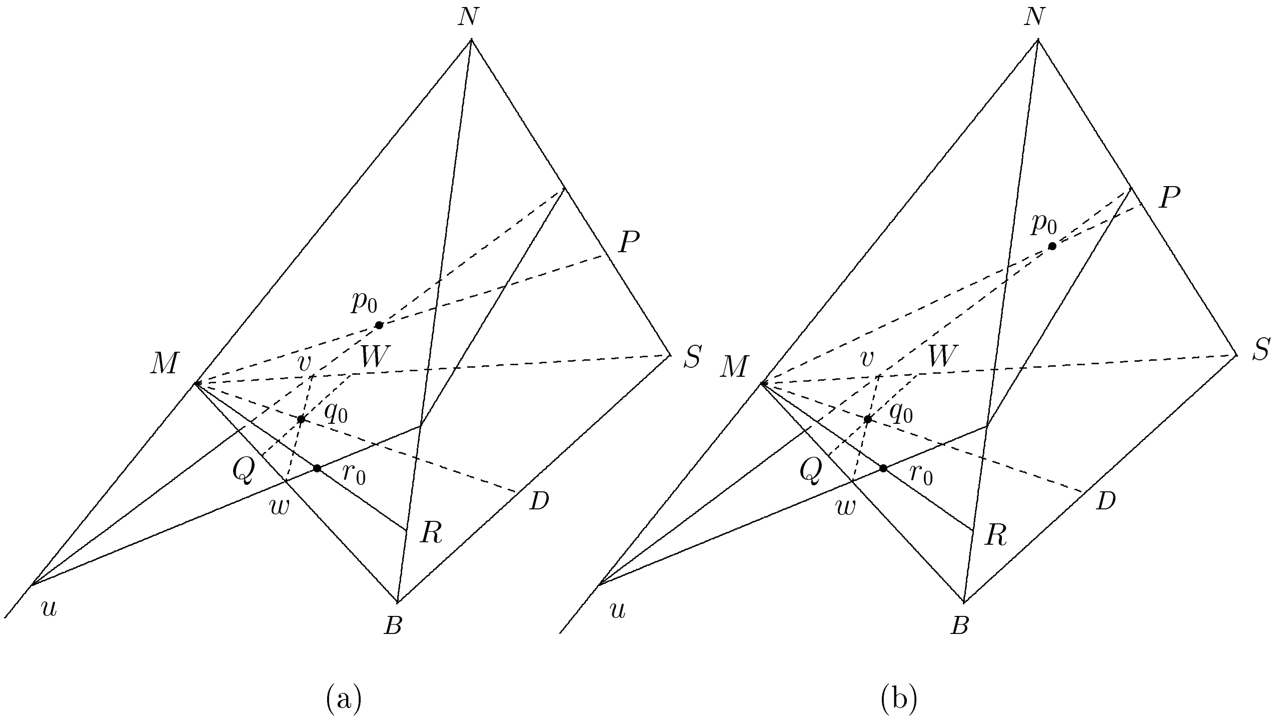

Figure 14: (a) $MR/Mr_0 < MP/Mp_0$. (b) $MR/Mr_0 > MP/Mp_0$.

7.1.1 The simplicial distance function

Theorem 5 *Let S be a a set of points in $\mathbb{R}^3$ in general position with respect to the simplex distance function. Then the system of bisectors of S under the simplex distance function is simple.*

The theorem follows from results in [18] and the following two lemmas.

Lemma 23 *Let p, q, r, s be any points in $\mathbb{R}^3$ in general position with respect to the simplex distance function. Then there is at most one homothetic copy of T passing through p, q, r, s.*

Proof: The lemma follows from the general position condition and from the observation that the number of faces of the unit simplex equals 4, the cardinality of $\{p, q, r, s\}$. $\square$

We remark that the lemma above can be generalized to higher dimensions.

Lemma 24 *Let p, q, r be any points in $\mathbb{R}^3$ in general position with respect to the simplex distance function. Then the set $B(p, q, r)$ ist either empty or a single biinfinite polygonal chain.*

We have to prove that, given any two homothetic copies T_1 and T_2 of T that pass through p, q, r, we can homothetically and continuously deform T_2 to T_1 under the constraint that, during the deformation, T_2 always passes through p, q, r. Equivalently, by a homothety argument, we have to prove that, given any homothetic copy $\triangle p'q'r'$ of the triangle $\triangle pqr$ whose vertices lie on T_1, we can homothetically and continuously deform $\triangle p'q'r'$ to $\triangle pqr$ in such a way that during the deformation the vertices of $\triangle p'q'r'$ always lie on T_1 — the trick is that, instead of keeping the three sites fixed and working with *arbitrary large* homothetic simplices passing through the sites,

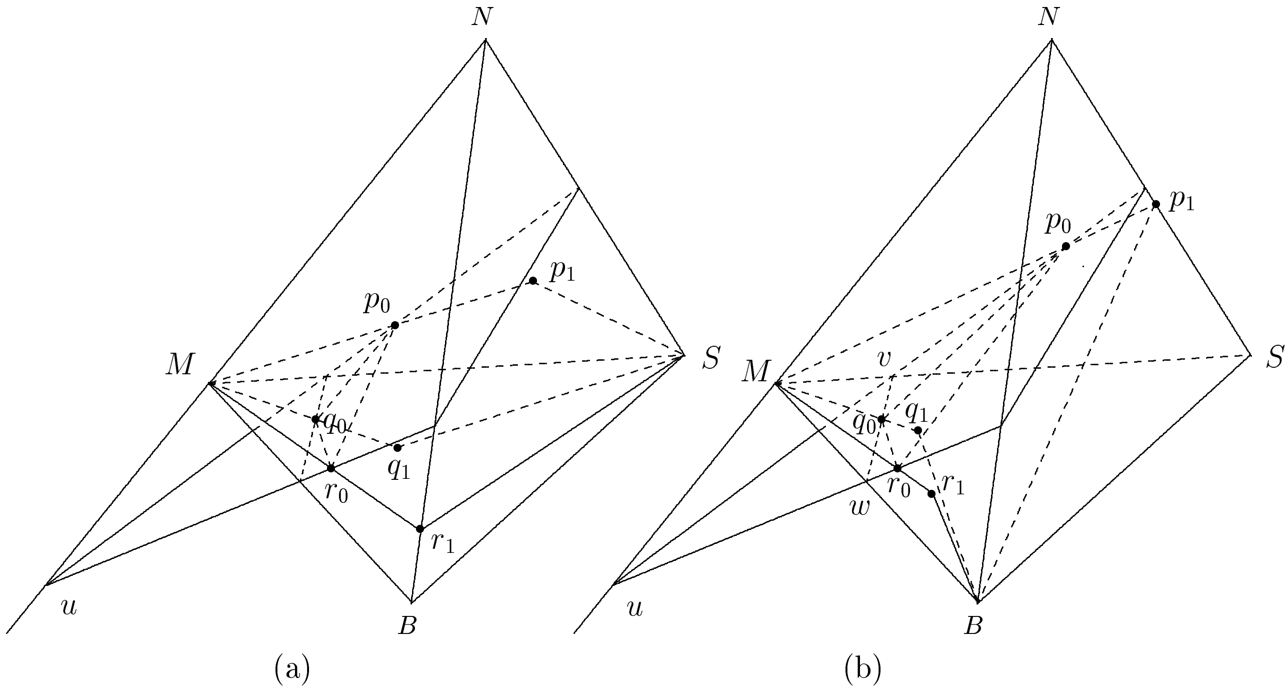

Figure 15: (a) The inscribed homothetic triangle $\triangle p_1q_1r_1$ shrinks to the extremal vertex S. (b) The inscribed homothetic triangle $\triangle p_1q_1r_1$ shrinks to the nonextremal vertex B.

we keep the simplex fixed and work with homothetic triangles having vertices on the *compact* simplex.

In general, we are confronted with the problem of determining the connected components of the set of homothetic triangles inscribed in a convex body. To detect all connected components of the set of all such homothetic triangles, we scan the convex body, denoted P, with a plane parallel to $\triangle pqr$. In any position of the scan plane, there is at most one triangle $\triangle p'q'r'$ homothetic to $\triangle pqr$ whose vertices lie on P (because, by a homothety argument, this is equivalent to the fact that there is at most one homothetic convex curve passing through three distinct points [11]). We call the set of all vertices of such $\triangle p'q'r'$'s that correspond to p the *p-trace* of $\triangle pqr$. For example, under the L_2 norm — P is a sphere — the p-trace of any $\triangle pqr$ is a relatively open half-bigcircle connecting the two points at which the scan plane just touches the sphere.

To determine the connected components of the set of triangles homothetic to $\triangle pqr$ and inscribed in P, we equivalently determine the connected components of the p-trace (or any of the two other traces, i.e. the q-trace or the r-trace) of $\triangle pqr$.

The case when P is a polyhedron naturally deserves our attention. Because the faces of P are planar, the deformations acting on $\triangle p'q'r'$ are of the following types.

- Shrinking or expanding. If p', q', r' lie on faces of P whose planes have a point in common, we can shrink or expanding $\triangle p'q'r'$ toward or away from this common point.

- Sliding. If p', q', r' lie on faces of P whose planes don't have a point in common, we can slide $\triangle p'q'r'$ along some fixed direction.

Observe that, relatively to the scan direction, each connected component of the p-trace arises by expanding $\triangle p'q'r'$ from some vertex of P, and ends by shrinking $\triangle p'q'r'$ to some other vertex

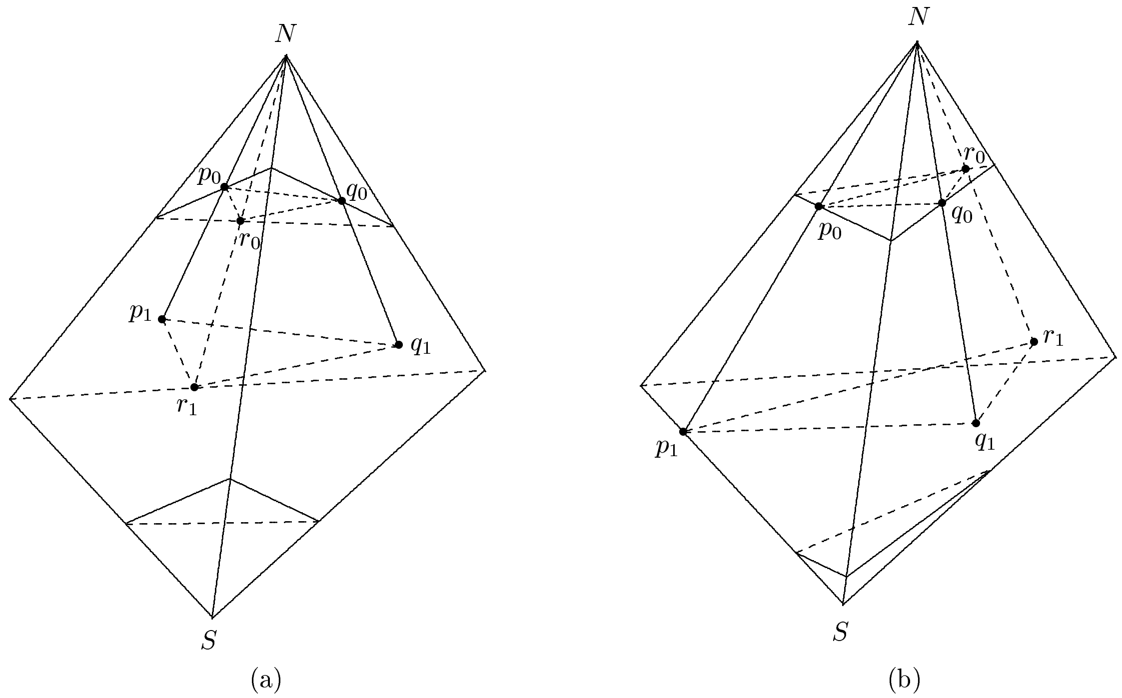

Figure 16: (a) The edge that is reached first by a vertex of the expanding homothetic triangle (here r_1) is opposite to S. (b) The edge that is reached first by a vertex of the expanding homothetic triangle (here p_1) is incident to S.

of P. There are two kinds of such vertices:

- Vertex with the property that P lies in one of the halfspace bounded by the scan plane going through that vertex. Call any such vertex an *extremal vertex (with respect to the scan direction)*.
- Vertex with the property that P is crossed by the scan plane going through that vertex. Call any such vertex a *nonextremal vertex (with respect to the scan direction)*.

Note that any scan plane that is sufficiently close to an extremal vertex and that intersects P can be parameterized by three points, denoted by u, v, w, on the edges incident to the vertex as depicted in Figure 13(a). It may happen that u lies "at infinity" as shown in Figure 13(c).

Figure 13(b) shows how to parameterize a scan plane that is sufficiently close to a nonextremal vertex and that intersects all the three faces incident to the vertex; observe that point u lies outside the polyhedron P. It may occur that u lies "at infinity", see Figure 13(c). We now turn to the proof of the lemma.

Proof: We prove the lemma by proving the following:

(1) Whenever the p-trace has a connected component that starts from a nonextremal vertex, then it consists solely of this connected component.

(2) The p-trace has at most one connected component that starts from an extremal vertex.

(1) Assume that the scan plane is not parallel to any edge of T_1. Let M be the nonextremal vertex of T_1, from which the triangle expands homothetically. Consider an intermediate position $\triangle p_0q_0r_0$ of the expanding triangle that is sufficiently close to M. The plane passing through

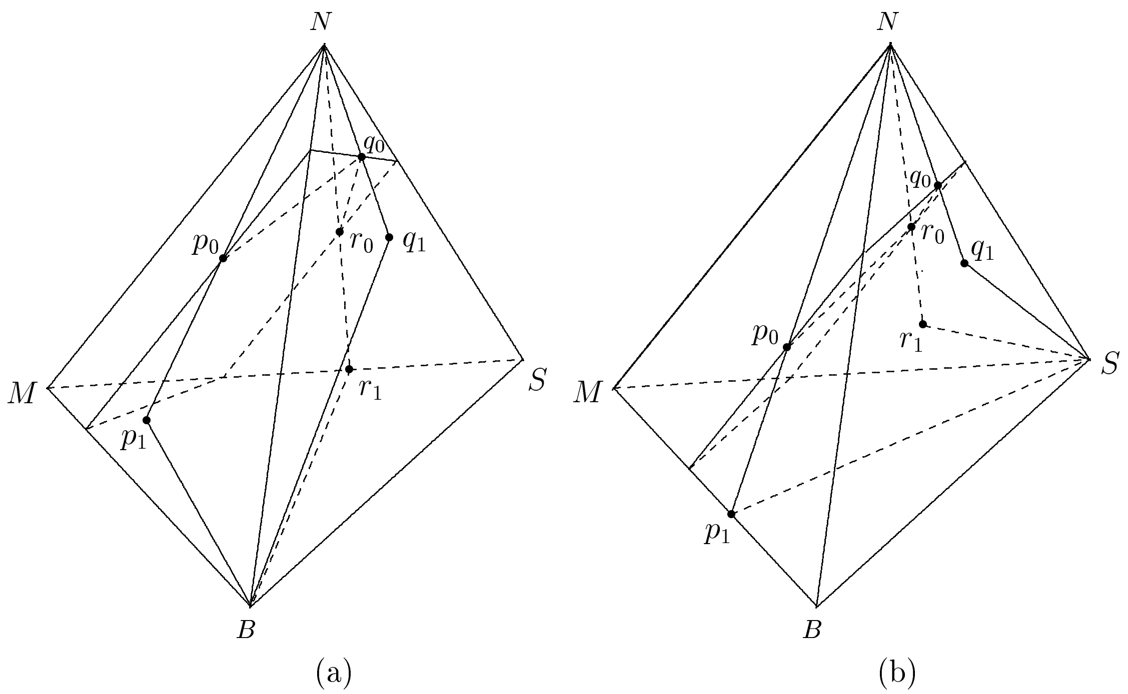

Figure 17: The scan plane through the extremal vertex N has the entire edge MN in common with T_1.

M and parallel to $\triangle p_0q_0r_0$ divides T_1 in two parts, each of which contains an extremal vertex of T_1. Therefore the extremal vertices N and S of T_1 lie in distinct halfspaces bounded by the plane spanned by p_0, q_0, w_0 as shown wlog in Figure 14.

First observe that in any scan plane between the extremal vertex N and the nonextremal vertex M there can't be any inscribed homothetical triangle: if such a triangle would exist, we would homothetically expand it to one that lies in the scan plane passing through the nonextremal vertex M, which is impossible.

The line through q_0 parallel to BS intersects MB and MS at Q and W. Because S is an extremal vertex, we have $MQ < Mw$. Thus $MD/Mq_0 > MB/Mw > MR/Mr_0$. We distinguish two cases:

1. $MR/Mr_0 < MP/Mp_0$, refer to Figure 14(a).

2. $MR/Mr_0 > MP/Mp_0$, see Figure 14(b).

(The case $MR/Mr_0 = MP/Mp_0$ is excluded because of the general position assumption.) We continue expanding the triangle homothetically from the position $\triangle p_0q_0r_0$ until one vertex of the expanding triangle reaches an edge of T_1. The vertex of the expanding triangle that reaches an edge first is the one that comes from r_0 in the first case, and the one that comes from p_0 in the second case, as shown in Figure 15(a) and (b), resp. The triangle is now at the position $\triangle p_1q_1r_1$. In the first case, the triangle shrinks to the extremal vertex S, while in the second case it shrinks to the nonextremal vertex B. Moreover, in the second case, there can't be any inscribed homothetic triangle lying in any scan plane between B and S. This can be seen using arguments similar to that for the case above where the scan plane is between N and M.

(2) Assume first that the scan plane through any extremal vertex of T_1 has only this point in common with T_1. As above, we expand the triangle homothetically from the extremal vertex N.

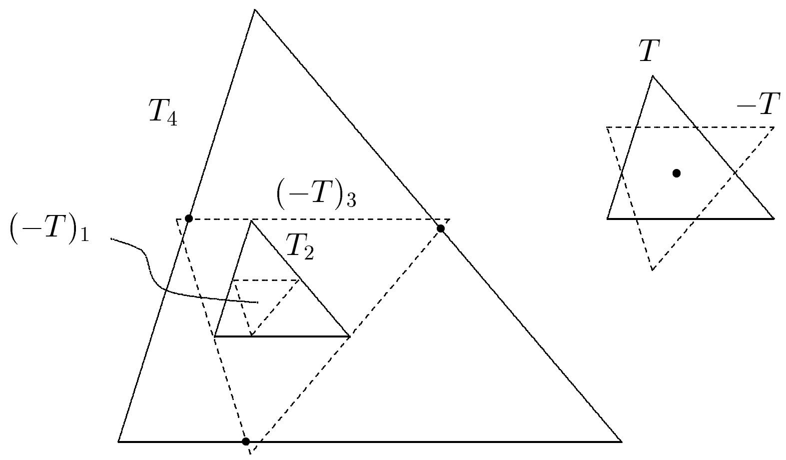

Figure 18: The sentinel sites are chosen randomly from small neighbourhoods of the vertices of the simplex $(-T)_3$.

A vertex of the expanding inscribed homothetic triangle will reach an edge of the face opposite to N first. Only two cases can occur:

1. The edge is opposite to S, Figure 16(a).

2. The edge is incident to S, Figure 16(b).

In the first case the triangle shrinks to the extremal vertex S. In the second case it shrinks to a nonextremal vertex; furthermore by the same arguments as in (1), there can't be any inscribed homothetic triangle lying in any scan plane between this nonextremal vertex and S.

Finally, the case where the scan plane through some extremal vertex of T_1 has an entire edge in common with T_1 can be handled similarly as above, refer to Figure 17. □

7.1.2 Computing the simplicial Voronoi diagram

The strictly simple abstract Voronoi diagram as described in [17] cannot be applied to our problem directly. This is because systems of bisectors under the simplicial distance function are not strictly simple; indeed, for sites $p, q, r, s \in \mathbb{R}^3$, the intersection $B(p,q) \cap B(q,r) \cap B(r,s)$ as well as the intersection $B(p,q) \cap B(q,r)$ are often empty. However, a closer inspection of [17] shows that we can still resort to this approach after appropriately adapting the initialization step of the algorithm in [17] to our problem.

Theorem 6 *In 3-space the Voronoi diagram under the simplicial distance function of n points in general position can be computed in expected time $O(n^2)$, using $O(n^2)$ expected space.*

Proof: Let S be the set of the given point sites. We first observe that properly-bounded Voronoi faces, i.e. Voronoi faces on the original bisectors, are always simply-connected; we omit the simple proof. The algorithm in [17] assumes that the system of bisectors is strictly simple to achieve that even Voronoi faces on Γ (the symbolic huge sphere defined in Lemma 6) are simply-connected. Note, however, that simplicity already suffices to ensure that, when a new site is inserted, no old properly-bounded Voronoi faces become disconnected. So the algorithm still works in our case if we can arrange that after the initialization step the Voronoi regions of the sites to be added won't meet Γ; in other words, if they are properly bounded. This can

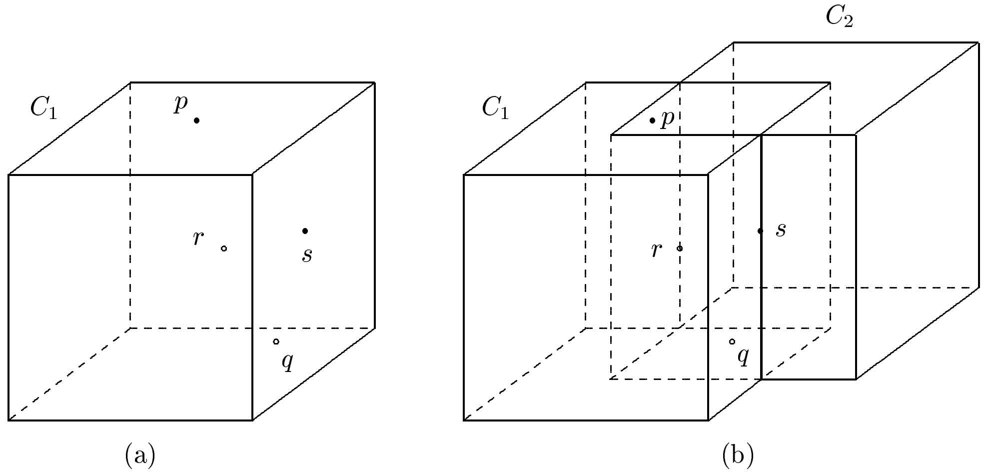

Figure 19: Points p, q, r, s being in general position. (a) The configuration of faces of any cube passing through p, q, r, s. (b) There are at most 2 homothetic cubes passing through p, q, r, s.

be achieved, borrowing an idea of [4], by augmenting the original sites with a set S_0 of certain artificial *sentinel sites* that surround the original sites in an appropriate way.

The initialization step for our case consists then in constructing the data structures (i.e., the incidence graph representing the Voronoi diagram and the history graph) corresponding to $\{\infty\} \cup S_0$, instead of those corresponding to the set $\{\infty, p, q, r\}$ for some sites p, q, r chosen randomly from S as in [17]. We then incrementally compute the diagram of the augmented set of sites $S' = S \cup \{\infty\} \cup S_0$ by inserting the sites of S in random order, as described in [17], and finally, we discard the sentinels to recover $V(S \cup \{\infty\})$ from $V(S')$.

In particular, let $(-T)_1$ be a homothetic copy of $-T$ so that $(-T)_1$ contains the original sites. Further, let $(-T)_3$ be a homothetic copy of $-T$ with a sufficiently large scaling factor containing $(-T)_1$ so that the clearance between $(-T)_1$ and $(-T)_3$ are big enough, i.e. as specified in Figure 18. We then choose 4 artificial sites sufficient close to the vertices of $(-T)_3$ and lie on distinct faces of $(-T)_3$. We finally obtain the sentinels by perturbing the artificial sites as a whole by a small random move, which ensures the general position assumption for the augmented set of sites.

Each original site has a properly-bounded Voronoi cell in the Vonoroi diagram of the augmented set of sites because any sufficiently large homothetic simplex passing this site contains a sentinel in it. Moreover, any homothetic simplex passing through 4 sites of S is entirely contained in T_4; therefore, the vertices of $V(S)$ are also vertices of $V(S')$. Thus, when we discard the sentinels, no new properly-bounded Voronoi vertex arises. From [17] we obtain that the complexity of any Voronoi region is $O(n)$, so discarding the sentinels does not increase the overall complexities. $\square$

7.2 The L_∞ distance function

Let C denote the unit L_∞ sphere, i.e., a cube containing the origin of the coordinate system in its interior domain.

Theorem 7 *Let S be a set of points in $\mathbb{R}^3$ in general position with respect to the L_∞ distance function. Then the intersection characteristic of the system of bisectors of S satisfies $\mathcal{C}_1 \subset \{0,1\}$ and $\mathcal{C}_2 \subset \{0,1,2\}$.*

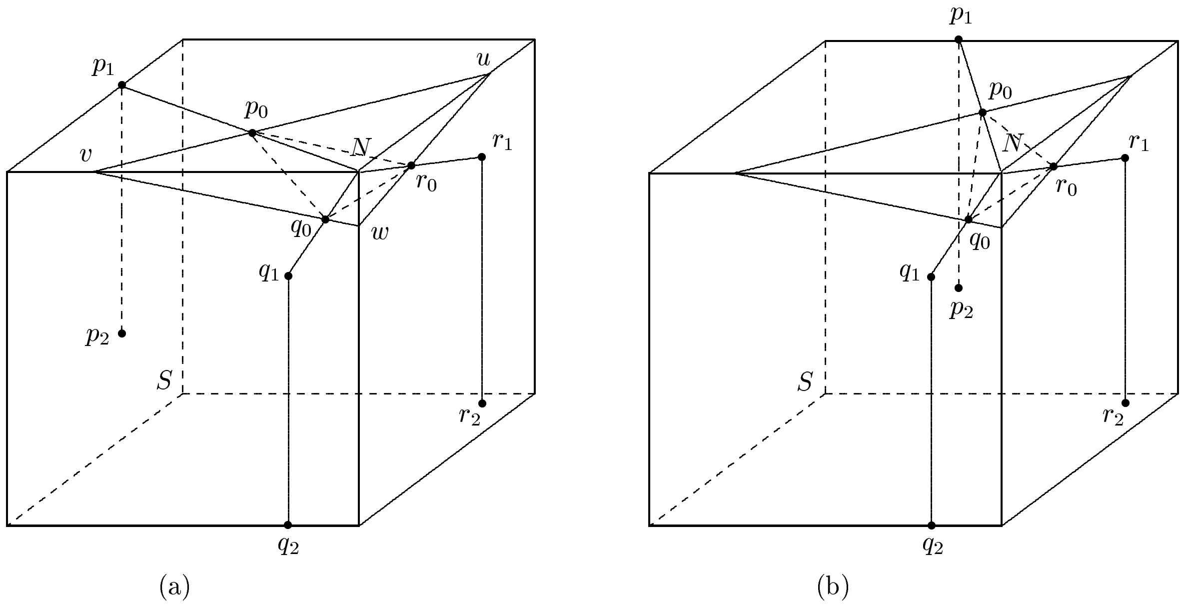

Figure 20: Expanding homothetic triangles from an extremal vertex N.

The theorem follows from results in [18] and the following two lemmas.

Lemma 25 *Let p, q, r, s be any 4 points in $\mathbb{R}^3$ in general position with respect to the L_∞ distance function. Then there are at most 2 homothetic copies of C passing through p, q, r, s.*

Proof: Let C_1 be any homothetic copy of the unit cube C passing through p, q, r, s. (If there is no such cube then we are done.) By the general position assumption, the supporting planes of the faces of C_1 passing through p, q, r, s must have empty intersection, even "at infinity". (For example, after removing the bottom and top faces of C, any four points lying on the remaining faces of C are not in general position, because the supporting planes of these faces intersect "at infinity".)

Consequently, we can show that whenever p, q, r, s are in general position then the faces (of any homothetic cube passing through them) will consist of three faces having a single vertex in common and one of the three remaining faces. Thus, two of the points p, q, r, s, say p and q, lie on two parallel faces of C_1, refer to Figure 19(a).

Let C_2 be any homothetic copy of the unit cube C, different from C_1, passing through p, q, r, s. (We are done if there is no such cube.) Then, by arguments similar to above, there are two parallel faces of C_2 passing through two distinct points among p, q, r, s. Therefore the size of C_2 must be smaller than or equal to the size of C_1. This follows by observing that any pair of points among p, q, r, s other than the pair p and q define, by means of parallel faces, a homothetic cube whose size is smaller than C_1.

By symmetry, the size of C_1 is smaller than or equal to that of C_2. So C_1 and C_2 must have the same size, which implies further that the points among p, q, r, s lying on parallel faces of C_2 are p and q. We infer that the intersection of C_1 and C_2 consists of two rectangles passing through p and q and two line segments passing through r and s as shown in Figure 19(b).

Call $\Pi(C_1)$ the projection of the shape of C_1 on P, the supporting plane of the face of C_1 passing through q. Call $\Pi(r)$ the projection of r on P. Suppose for contradiction that there is a third homothetic cube C_3 passing through p, q, r, s. Then the size of C_3 would equal the size

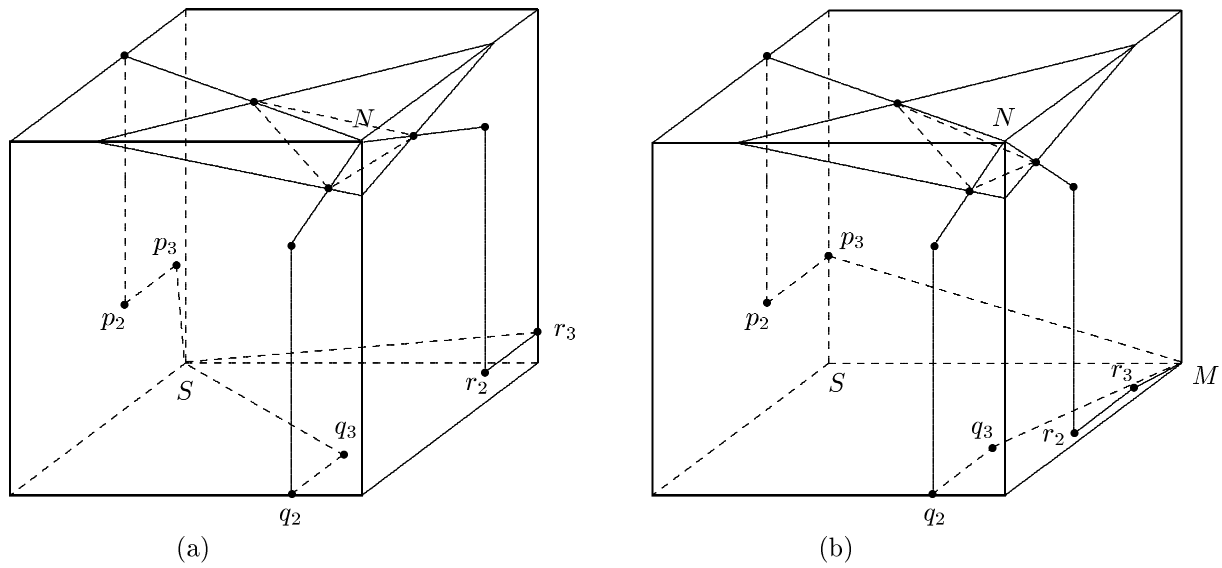

Figure 21: Shrinking homoth. triangles: (a) to extremal vertex S. (b) to nonextremal vertex M.

of C_1. Thus $\Pi(C_3)$ would pass through $\Pi(r)$ and $\Pi(s)$, and cross $\Pi(C_1)$ and $\Pi(C_2)$, which is impossible. $\square$

Lemma 26 *Let p, q, r be any 3 points in $\mathbb{R}^3$ in general position with respect to the L_∞ distance function. Then the set $B(p, q, r)$ ist either empty or a single biinfinite polygonal chain.*

Proof: We have to prove that the p-trace (or any of the two other traces) of $\triangle pqr$ is connected. We prove the lemma by proving the following.

(1) The p-trace has at most one connected component that starts from an extremal vertex, and

(2) The p-trace has no connected component that starts from a nonextremal vertex and ends at a nonextremal vertex.

We proceed with the proof by finding all possible configurations of the traces on a homothetic cube C_1, relying on a thorough case analysis. We assume that the scan plane is not parallel to any edge of C_1. A proof sketch for the remaining case is given in Appendix C.

(1) Let N be the extremal vertex from which $\triangle p'q'r'$ expands, see Figure 20. Let $\triangle p_0q_0r_0$ be a sufficiently small intermediate position of $\triangle p'q'r'$. To construct the p, q, r-traces we concurrently extend the rays from N through p_0, q_0, r_0 proportionally to the lengths of the segments Np_0, Nq_0, Nr_0 until one of them hits a face different than the one on which it lies. The expanding triangle is now at some position $\triangle p_1q_1r_1$. We then slide $\triangle p_1q_1r_1$ along the faces of the cube until one of its vertices reaches a new face. Due to the geometry of the cube, this vertex does not correspond to the vertex that has changed face before. Denote the new position of the sliding triangle by $\triangle p_2q_2r_2$. By exploiting symmetry operations on the cube such as rotation or reflection — they preserve homothety — and after renaming (if nesessary) the sites, we need to consider only two cases shown in Figure 20.

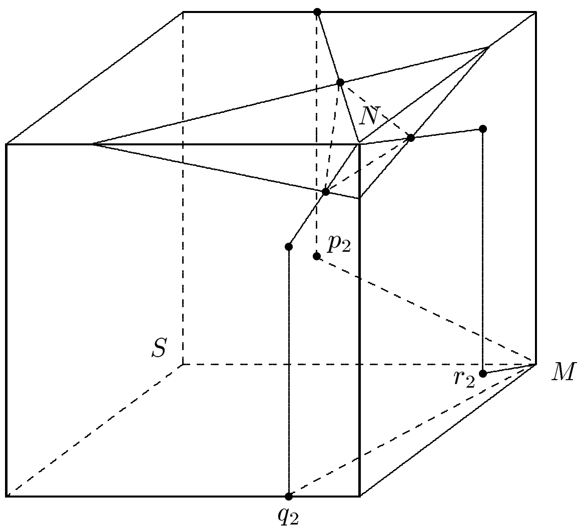

Figure 22: Shrinking homothetic triangles to a nonextremal vertex M.

Case (a) of Figure 20. We slide $\triangle p_2q_2r_2$ until either p_2 or r_2 reaches another face of the cube; note that q_2 lies relatively to the new slide direction behind p_2 and r_2. Depending on which one of the points p_2 and r_2 reaches a new face first, we distinguish two cases as depicted in Figure 21. In both cases, the triangle is at the position $\triangle p_3q_3r_3$.

– Subcase (a) of Figure 21. We shrink $\triangle p_3q_3r_3$ to the extremal vertex S. Clearly, in this subcase each trace has a single connected component.

– Subcase (b) of Figure 21. We shrink $\triangle p_3q_3r_3$ to the nonextremal vertex M. This subcase will be considered further at $(*)$.

Case (b) of Figure 20. Triangle $\triangle p_2q_2r_2$ can be shrunk to the nonextremal vertex M, as shown in Figure 22. We continue considering this case at $(*)$.

$(*)$ Consider the other extremal vertex S. If the whole trace would have another connected component starting from S, then due to symmetries of the cube one of the cases found above would hold for this connected component too. We then would have that the connected components starting from N and from S would meet each other, which is impossible.

(2) Let M be the nonextremal vertex from which the triangle expands homothetically, see Figure 23. Choose a sufficiently small intermediate position of the expanding triangle, denoted $\triangle p_0q_0r_0$. We proceed similarly to Case (1). After expanding and sliding the triangle, we obtain — refer to Appendix A and B, and up to symmetries of the cube — two cases as shown in Figure 23, where the triangle is now at the position $\triangle p_2q_2r_2$.

Case (a) of Figure 23. We claim that p_2 is closer to the left face of the cube than q_2; the simple proof can be found in Appendix B. We slide $\triangle p_2q_2r_2$ along the cube. Due to the claim, the vertex of the triangle which was at p_2 will reach a new face first. So the triangle will be at $\triangle p_3q_3r_3$, see Figure 24(a). Finally we shrink the triangle homothetically to the extremal vertex N. (Note that if q_2 would be closer to the left face of the cube than p_2, then after the sliding phase we would be able to shrink the triangle to a nonextremal vertex. So Claim (2) would be false!)

Case (b) of Figure 23. We shrink the triangle homothetically to the extremal vertex N, the result is shown in Figure 24(b). This completes the proof. $\square$

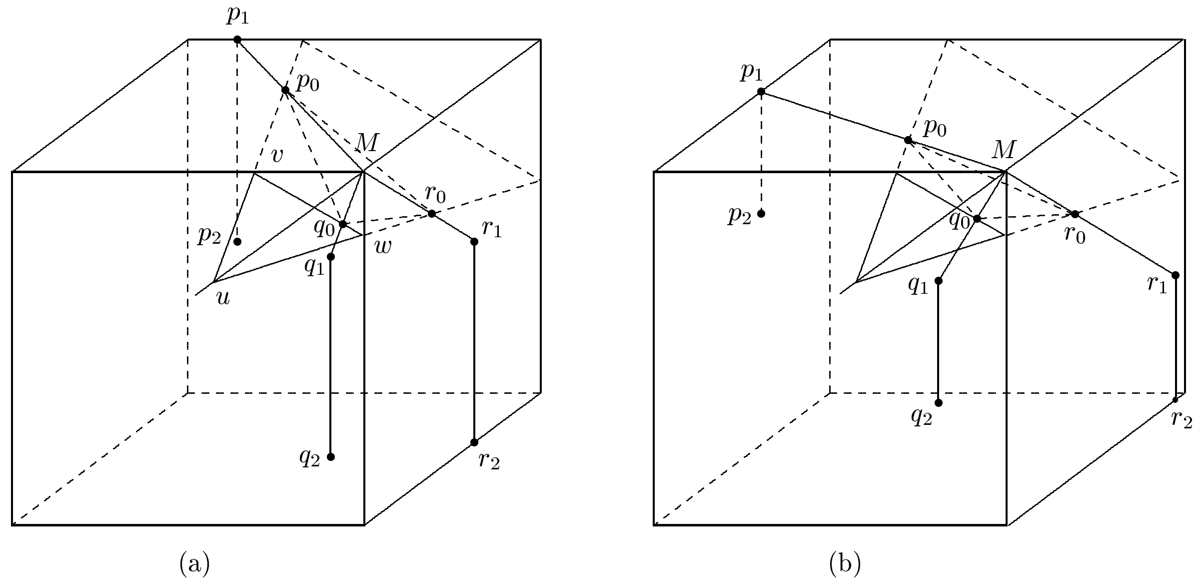

Figure 23: Expanding homothetic triangles from a nonextremal vertex M.

7.3 The L_p norm

The L_p norm, $1 \le p < \infty$, is defined by the equation

$$L_p(u,v) = \sqrt[p]{|u_1 - v_1|^p + |u_2 - v_2|^p + |u_3 - v_3|^p}$$

for any $u = (u_1, u_2, u_3)$ and $v = (v_1, v_2, v_3) \in \mathbb{R}^3$. Using the notion of intersection characteristic we can adapt the results of [16] to our concept of abstract Voronoi diagram as follows.

Theorem 8 *Let S be a set of at least four points in 3-space in general position with respect to the L_p distance function, p being even. Then for the intersection characteristic of the system of bisectors of S holds*

1. *$\mathcal{C}_1$ equals $\{1\}$.*
2. *$\mathcal{C}_2$ contains 1 and is a subset of $\{1, 3, \ldots, K\}$, where K is a positive odd integer independent of p.*

Extensive computations done by Maria Jose Gonzalez Lopez of the Computational Algebra and Geometry group of Tomas Recio strongly support the conjecture that K is rather small, namely $K = 3$.

8 Concluding remarks

We have proposed a class of abstract Voronoi diagrams in 3-space. To characterize them we have introduced the notion of intersection characteristic.

Systems of bisectors of point sites in general position under convex distance functions are examples of the abstract Voronoi diagram. We have given methods to determine the intersection characteristic for the simplex and the L_∞ distance functions. We note that these methods can be generalized to other polyhedral convex distance functions as well. Moreover, they might be

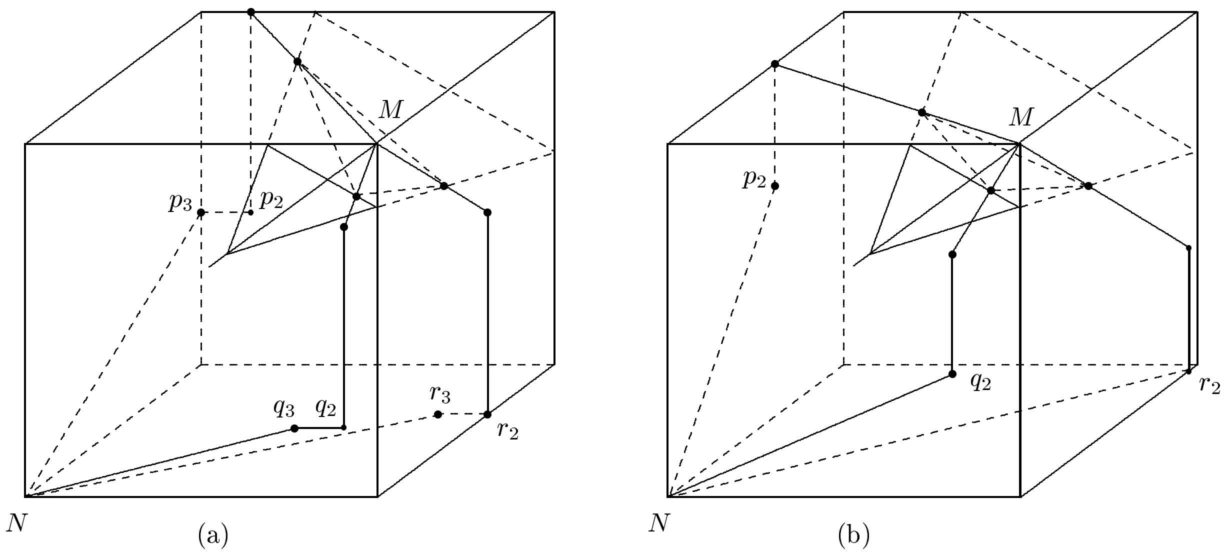

Figure 24: Shrinking homothetic triangles to an extremal vertex N.

useful in implementations of incremental algorithms to construct the Voronoi diagram under polyhedral convex distance functions.

The fact that the intersection characteristic for the simplex distance function satisfies $\mathcal{C}_1 \subset \{0,1\}$ and $\mathcal{C}_2 \subset \{0,1\}$ has been a hint for us to apply the algorithm originally developed for systems of bisectors with the intersection characteristic $\{1\}$ and $\{1\}$ in [17] to the simplex distance function.

To conclude the paper, we mention two major open problems. The first one is to give a tight bound on the complexity of the abstract Voronoi diagram. The second one is to find an efficient algorithm to compute the abstract Voronoi diagram.

Acknowledgements

I wish to thank Rolf Klein, Li-Hong Ma, and Kurt Mehlhorn for (directly or indirectly) contributing some initial ideas. Through Italiano I am informed of a counterexample showing that a previous version of Lemma 25 is false.

Thanks also go to Siw-Wing Cheng and Otfried Cheong for giving me the possibility for finishing this work.

References

[1] J. W. Alexander. An example of a simply connected surface bounding a region which is not simply connected. In *Proc. Nat. Acad. Sci. U.S.A.*, 10 (1924), pp. 6–8.

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

[3] F. Aurenhammer and R. Klein. Voronoi Diagrams. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry.* Elsevier Science Publishers B.V., Amsterdam, 1999.

[4] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec. Voronoi Diagrams in Higher Dimensions under Certain Polyhedral Distance Functions. *Discrete Comput. Geom.*, 19:485–519, 1998.

[5] K. Q. Brown. Voronoi diagrams from convex hulls. *Information Processing Letters*, 9:223–228 (1979).

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

[7] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.

[8] H. Edelsbrunner and R. Seidel. Voronoi diagrams and arrangements. *Discrete Comput. Geom.*, 1:25–44, 1986.

[9] P. E. Ehrlich and H.-C. Im Hof. Dirichlet regions in manifolds without conjugate points. *Comment. Math. Helvetici*, 54 (1979), pp. 642–658.

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

[11] C. Icking, R. Klein, N.-M. Lê, L. Ma. Convex Distance Functions in 3-Space are Different. *Fundam. Inform.*, 22:331–352, 1995.

[12] R. C. Kirby. The union of flat $(n-1)$-balls is flat in $\mathbb{R}^n$. *Bull. Amer. Math. Soc.*, 74:614–617 (1968).

[13] R. Klein. *Concrete and Abstract Voronoi Diagrams.* LNCS 400, Springer-Verlag, New York 1989.

[14] R. Klein, K. Mehlhorn, and S. Meiser. Randomized Incremental Construction of Abstract Voronoi Diagrams. In *Comput. Geometry: Theory and Applications*, 3 (1993), pp. 157–184.

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

[16] N. M. Lê. On Voronoi diagrams in the L_p-metric in $\mathbb{R}^D$. *Discrete Comput. Geom.*, 16:177–196, 1996.

[17] N.-M. Lê. Randomized incremental construction of simple abstract Voronoi diagrams in 3-space. *Computational Geometry*, 8 (1997), 279–298.

[18] N.-M. Lê. On non-smooth convex distance functions. *Information Processing Letters*, 63(6):323–329 (1997).

[19] N.-M. Lê. On Determining Optimal Strategies in Pursuit Games in the Plane. *Theoretical Computer Science*, 197(1–2):203–234, 1998.

[20] W. S. Massey. *Algebraic Topology: An Introduction.* GTM 56, Springer-Verlag, New York 1977.

[21] K. Mehlhorn, S. Meiser, and C. Ó'Dúnlaing. On the Construction of Abstract Voronoi Diagrams. *Discrete Comput. Geom.*, 6:211–224, 1991.

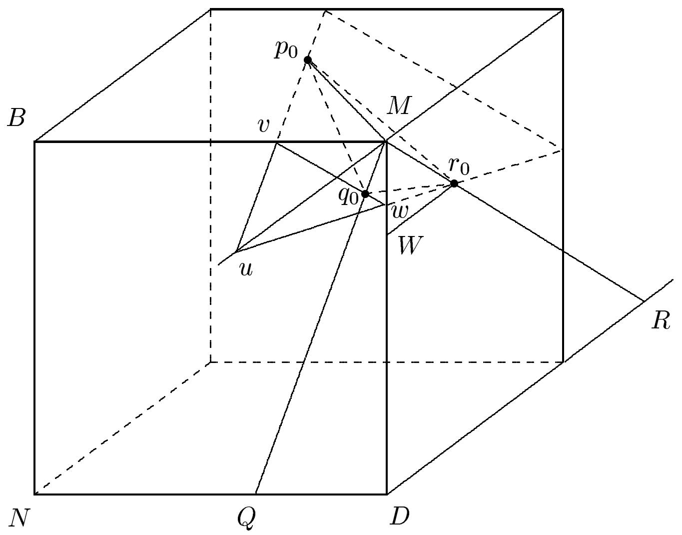

Figure 25: Illustrating the proof that $MR/Mr_0 < MQ/Mq_0$.

[22] A. Okabe, B. Boots, and K. Sugihara. *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams.* John Wiley & Sons, Chichester, UK, 1992.

[23] D. Rolfsen. *Knots and Links.* Mathematics Lecture Series, Publish or Perish, Inc., Texas 1990.

[24] T. B. Rushing. *Topological Embeddings.* Academic Press, New York and London 1973.

[25] S. I. Shamos and D. Hoey. Closest point problems. in *Proc. of the 16th Annual IEEE Symposium on Foundations of Computer Science*, pp. 151–162, 1975.

[26] S. Stifter. An Axiomatic Approach to Voronoi-Diagrams in 3D. *Journal of Computer and System Sciences*, 43(2):361–379, 1991.

A

We claim that, when $\triangle p'q'r'$ expands from the nonextremal vertex M homothetically, always keeping contact with the cube, then either p' or r' will reach an opposite edge of the cube first. To see this, imagine p', q', r' are moving concurrently with constant velocities $\overrightarrow{Mp_0}, \overrightarrow{Mq_0}, \overrightarrow{Mr_0}$, respectively, from M towards the boundaries of the faces containing them, see Figure 25. Let Q be the goal point for q'. Then q' needs MQ/Mq_0 time units to reach the opposite edge of the cube. If Q lies in the edge NB, then we have $MQ/Mq_0 > MB/Mv$; otherwise, Q lies in the edge ND, then we have $MQ/Mq_0 > MD/Mw$.

Assume the latter case, the first case is similar. Point r' needs at most MR/Mr_0 time units to reach the boundary of the face containing it. Let W be a point in MD so that Wr_0 and DR are parallel. Then, $MR/Mr_0 = MD/MW$. Moreover, due to the direction of ur_0, we have $MW > Mw$. Thus, $MR/Mr_0 < MD/Mw$. We conclude that $MR/Mr_0 < MQ/Mq_0$. A similar inequality holds for point p' if Q lies in the edge NB. Therefore, either p' or r' will reach the boundary of the face containing it before q'.

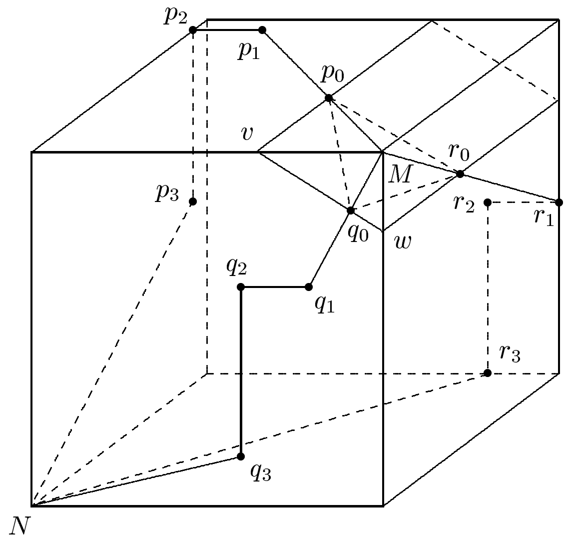

Figure 26: The scan plane is parallel to an edge of the cube.

B

As above, imagine that p', q', r' are moving concurrently with constant velocities $\overrightarrow{Mp_0}, \overrightarrow{Mq_0}, \overrightarrow{Mr_0}$, respectively, from M towards the boundaries of the faces containing them, see Figure 25. Due to the directions of uv and vw, the length of the projection of $\overrightarrow{Mp_0}$ on BM (which is orthogonal to the left face of the cube) is greater than the length of the projection of $\overrightarrow{Mq_0}$ on BM. Therefore, at the end of this expansion phase, p' being at p_1 is closer to the left face of the cube than q' being at q_1. Similarly, at the end of this expansion phase, r' being at r_1 is closer to the bottom face of the cube than q_1.

C

Consider the case when the scan plane is parallel to an edge of the cube. In this case, the trace has at most one connected component, as depicted in Figure 26. The proof is very similar to the previous ones, and is omitted.