

PhD THESIS

prepared at
INRIA Sophia Antipolis

and presented at the
University of Nice-Sophia Antipolis
Graduate School of Computer Science

*A dissertation submitted in partial fulfillment
of the requirements for the degree of*

DOCTOR OF COMPUTER SCIENCE
Specialized in Computational Geometry

Generalized Voronoi Diagrams and Applications

Camille WORMSER

Advisors	Pr. Jean-Daniel Boissonnat	INRIA Sophia Antipolis, France
	Pr. Mariette Yvinec	INRIA Sophia Antipolis, France
Reviewers	Pr. Aurenhammer	Technische Universitaet Graz
	Pr. Shewchuk	University of California at Berkeley
Examiners	Pr. Habib	Université Paris VII
	Pr. Hecht	Université Paris VI
	Pr. Lazard	Loria

UNIVERSITÉ NICE-SOPHIA ANTIPOLIS - UFR Sciences

École doctorale STIC

(Sciences et techniques de l'information et de la communication)

THÈSE

pour obtenir le titre de
DOCTEUR EN INFORMATIQUE
de l'UNIVERSITÉ de Nice-Sophia Antipolis
Discipline: Géométrie algorithmique

présentée et soutenue par

Camille WORMSER

Diagrammes de Voronoï généraux et applications

Thèse dirigée par Pr. Mariette Yvinec
et co-encadrée par Pr. Jean-Daniel Boissonnat

Soutenue le 1er décembre 2008

Composition du jury:

<i>Rapporteurs</i>	Pr. Aurenhammer	Technische Universitaet Graz
	Pr. Shewchuk	University of California at Berkeley
<i>Examineurs</i>	Pr. Habib	Université Paris VII
	Pr. Hecht	Université Paris VI
	Pr. Lazard	Loria

Abstract

Voronoi diagrams are fundamental data structures that have been extensively studied in Computational Geometry. A Voronoi diagram can be defined as the minimization diagram of a finite set of continuous functions. Usually, each of those functions is interpreted as the distance function to an object. The associated Voronoi diagram subdivides the embedding space into regions, each region consisting of the points that are closer to a given object than to the others. We may define many variants of Voronoi diagrams depending on the class of objects, the distance functions and the embedding space. Affine diagrams, i.e. diagrams whose cells are convex polytopes, are well understood. Their properties can be deduced from the properties of polytopes and they can be constructed efficiently.

The first part of this thesis is dedicated to the presentation and classification of Voronoi diagrams. We discuss the most studied varieties of Voronoi diagrams, before putting these diagrams in the context of abstract Voronoi diagrams, a notion inherited from Klein. This allows us to present in a general setting the question of recognizing classical Voronoi diagrams by looking at their bisectors, a point of view initiated by Aurenhammer.

In the second part, we focus on the study of anisotropic Voronoi diagrams, and the ways of computing their dual mesh, if it is well defined. If the dual mesh is not well defined, we study some ways of refining the diagram in order to obtain a well-defined dual. We first use the definitions of Labelle and Shewchuk and the linearization procedure, as presented in the previous part. This allows us to define an algorithm which is the natural consequence of Part I.

The third part is then devoted to a different approach to anisotropic meshing. By changing the definition of an anisotropic mesh into the one of a locally uniform anisotropic mesh, we allow the design of simple anisotropic mesh generation algorithms in 2D and 3D.

Finally, the fourth part of this thesis is devoted to the application of a different kind of Voronoi diagrams, namely power diagrams, to the question of greedy routing in ad hoc networks. There again, the local properties of triangulations play a crucial role. We prove how some local properties of regular triangulations, which are a generalization of Delaunay triangulations, imply global properties in terms of routing.

Acknowledgments

I worked for four years at INRIA Sophia-Antipolis, where I was advised by Mariette and Jean-Daniel. I would like to thank them for their advising and support, their enthusiasm, and their eagerness to transmit their knowledge. Maybe most importantly, I am very grateful for their making themselves extremely available for advising me.

This work moved forward thanks to their support, but also thanks to the help and support of the whole Geometrica team, students, postdocs and researchers alike, who provided me a rich environment, as experienced work colleagues, and very much as friends too.

I would also like to thank Pierre Pansu, who supported my orientation towards computer science, and Craig Gotsman, who provided me nice opportunities of extending my field of research.

My family and my friends have provided me so much help and support, I love them and thank them for that.

Contents

I	Voronoi Diagrams	21
1	Affine Diagrams	23
1.1	Lower Envelopes and Minimization Diagrams	24
1.2	Affine Voronoi Diagrams	25
1.2.1	Euclidean Voronoi Diagrams of Points	25
1.2.2	Delaunay Triangulation	27
1.2.3	Power Diagrams	30
2	Curved Voronoi Diagrams	33
2.1	Voronoi Diagrams with Algebraic Bisectors	34
2.1.1	Möbius Diagrams	34
2.1.2	Anisotropic Diagrams	37
2.1.3	Apollonius Diagrams	39
2.2	Linearization	42
2.2.1	Abstract Diagrams	42
2.2.2	Inverse Problem	46
II	Anisotropic Meshing by Linearization of Anisotropic Voronoi Diagrams	49
2.3	Introduction	51
2.4	Labelle and Shewchuk's Approach	51
2.5	Relation to Power Diagrams	54
2.6	Basic Operations and Primitives	55
2.7	Description of the Algorithm	57
2.8	Termination of the Algorithm	61
2.9	Conclusion and Future Work	67
III	Locally Uniform Anisotropic Meshes	69
3	Local Uniformity	71
3.1	Introduction	72
3.2	Preliminaries	72
3.2.1	Anisotropic Metric	72
3.2.2	Distortion	73
3.3	Stars and Refinement	74
3.3.1	Stars	74
3.3.2	Quasi-Cosphericity	74

3.3.3	Picking Region	76
3.3.4	Encroachment and Star Initialization	79
3.4	Algorithm	80
3.4.1	Algorithm Outline	80
3.4.2	Termination of the Algorithm and Quality of the Mesh	81
3.5	Conclusion	84
IV	Greedy Ad Hoc Routing	87
4	Greedy Power Routing	89
4.1	Sensor Networks	90
4.2	Previous Work	90
4.3	Greedy Power Routing	91
4.3.1	Power Routing	91
4.3.2	Contained Power Diagrams	91
4.3.3	Equivalence to Polyhedral Routing	94
5	Algorithmic Solutions	97
5.1	Circle Packing	98
5.1.1	Definitions	98
5.1.2	Practical Computations of Circle Packings	99
5.2	Local Termination Conditions	100
5.2.1	Triangulated Case	100
5.2.2	Generalized Papadimitriou and Ratajczak Result	104
5.2.3	Non Triangulated Case	105
5.2.4	Relation between Circle Packings, LPD and GLPD	108
5.3	Algorithms	108
5.3.1	Computing a Greedy Power Embedding	109
5.3.2	Termination	109
5.3.3	Locality	110
5.4	Discussion	110
5.4.1	Experimental Validation	110
5.4.2	Possible Improvements	111
	Bibliography	119

List of Figures

1	The Voronoi diagram of a set of 9 points.	13
2	The Delaunay triangulation of a set of 9 points.	14
3	The power diagram of a set of 9 weighted points.	15
4	The Apollonius diagram of a set of 9 circles.	15
5	Giraffe and its Voronoi fur.	16
6	A giant turtle and its Voronoi shell.	16
7	Two E.coli colonies, growing into an Apollonius diagram	17
8	Three Proteus mirabilis colonies, growing into a Moebius diagram	17
1.1	The lower envelope of a set of univariate functions. The minimization diagram is drawn on the horizontal line with the corresponding indices. The face of index $\{1\}$ consists of two components.	24
1.2	The Voronoi diagram of a set of 9 points.	26
1.3	The polyhedron, with one of its faces projected.	27
1.4	The Delaunay triangulation of a point set (in bold) and its dual Voronoi diagram (thin lines).	28
1.5	The polar hyperplane of a sphere.	29
1.6	A power diagram.	31
2.1	A Möbius diagram.	35
2.2	The Apollonius diagram of a set of circles. Compare with the power diagram of the same set of circles in Fig. 1.6.	40
2.3	A cell in an Apollonius diagram of hyperspheres projects vertically onto a Möbius diagram in σ_0	42
2.4	An anisotropic diagram (Courtesy of J. Shewchuk)	52
2.5	Two umbrellas (left) and one umbrella winding twice (right) around p	57
2.6	Edge e is a constraint segment, with the cell of a being completely included in the cell of b . Voronoi bisectors are represented by dashed curves.	59
2.7	Impossible case described in Lemma 2.8.4	63
3.1	Example of inconsistent stars in 2D: stars S_v and S_w are inconsistent because edge $[vw]$ belongs to S_v but not to S_w	74
3.2	Example of quasi-cospherical points in 2D: a , b , c and d are quasi-cospherical because d is outside of $\mathcal{C}_N(abc)$ but inside $\mathcal{C}_{N'}(abc)$	75
3.3	q , r , s define a forbidden region (black annulus) for p in the picking region (grey area)	79
3.4	Output of the algorithm with a zoom on the central part: the red lines delimit the zoomed region.	84
3.5	Output of the algorithm, implemented in 3D by Yuanmi Chen	85

4.1	A contained embedding obtained from the <i>kissing disks</i> theorem of Koebe and Andre'ev.	92
4.2	As the radius of the circle around w grows, $\text{Cell}(w)$ grows and $\text{Cell}(v)$ shrinks. The power diagram becomes uncontained when v is no longer in $\text{Cell}(v)$	93
4.3	Two circles in the plane and their images (P and Q) on the sphere through $\mathcal{C}^{-1} \circ \psi$	96
5.1	The local cell $\text{Cell}_G(v)$ (solid lines) contains the power diagram cell (dashed lines) and contains another vertex w	101
5.2	As w_4 moves away from v , $\text{LPD}(v, \phi)$ becomes unsatisfied (while the solid lines diagram becomes the dashed lines diagram), because $\text{Cell}_G(v)$ and $\text{Cell}_G(w_4)$ are not adjacent anymore, whereas edge $[vw_4]$ exists in G	102
5.3	Lifting two local cells that share an edge.	103
5.4	A face with 6 vertices embedded by ϕ with the regular triangulation of its vertices: G (solid lines) is triangulated into $G(\phi)$ (solid and dashed lines).	106
5.5	A face (solid edges) with 5 vertices, with $\text{GLPD}(v, \phi)$ not satisfied: the local cell of v (dashed lines) crosses the boundary of the face not only on $[w_1v]$ and $[vw_4]$ but also on $[w_2w_3]$, which is forbidden.	107
5.6	After 6 iterations, the colored circles are the ones that already satisfy LPD.	111
5.7	After 29 iterations, only 2 circles still do not satisfy LPD.	112
5.8	After 32 iterations, LPD is staisfied everywhere: the embedding is a greedy power embedding.	113
5.9	After 128 iterations, the embedding is a circle packing.	113

Figure 5 is a photograph by Miroslav Duchacek, licensed under the terms of the GNU Free Documentation License. Figure 6 is a photograph by Matthew Field, licensed under the terms of the GNU Free Documentation License. Figures 7 and 8 are pictures from University of Chicago magazine, The Complexity Complex, December 2002, Volume 95, Issue 2.

Introduction

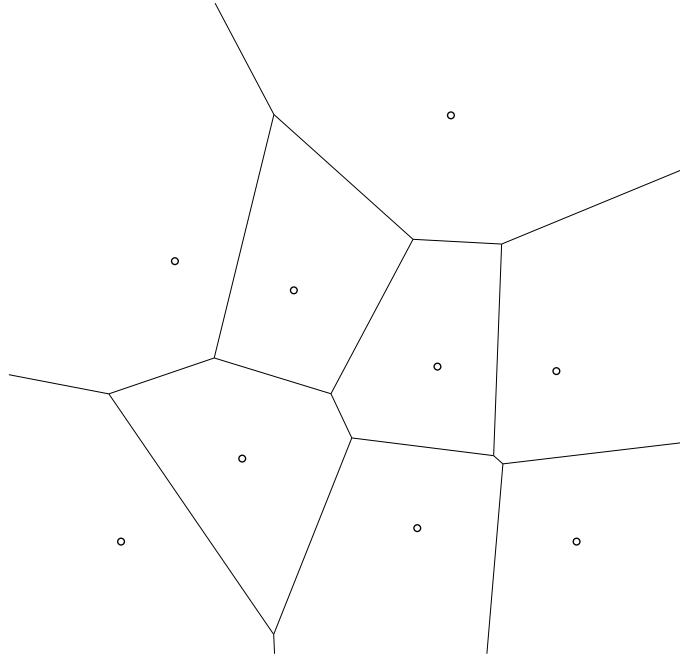


Figure 1: The Voronoi diagram of a set of 9 points.

History

Voronoi diagrams are fundamental data structures that have been extensively studied in Computational Geometry. Originally, these diagrams have been defined for a given set of points, called *sites* in Euclidean space, and for the Euclidean distance. The diagram consists of a decomposition of the space into regions called *cells*, each cell consisting of the points that are closer to a given site than to the others:

Let $\mathcal{P} = \{p_1, \dots, p_n\}$ be a set of points of \mathbb{R}^d . To each p_i , we associate its Voronoi region $V(p_i)$

$$V(p_i) = \{x \in \mathbb{R}^d : \|x - p_i\| \leq \|x - p_j\|, \forall j \leq n\}.$$

The region $V(p_i)$ is the intersection of $n-1$ half-spaces. Each such half-space contains p_i and is bounded by the bisector of p_i and some other point of \mathcal{P} . Since the bisectors are hyperplanes, $V(p_i)$ is a convex polyhedron, possibly unbounded. See Fig. 1 for an example in \mathbb{R}^2 .

This construction is natural for defining some kind of domain of influence around each site. Furthermore, the adjacency relations between the cells provide a natural construction of a graph of neighbors. This graph, called the *Delaunay triangulation*, has several interesting properties, which go beyond this intuition. First of all, it is a triangulation (see Fig. 2 for an example). And this triangulation has the unique property of being the triangulation with the best possible angles for its triangles, which are as close as possible to being equilateral. Several interesting mathematical and algorithmic properties derive from these central facts.

A Voronoi diagram can be defined more generally as the minimization diagram of a finite set of continuous functions. Usually, each of those functions is interpreted as the distance function to an object. The associated Voronoi diagram subdivides the embedding space into regions, each region consisting of the points that are closer to a given object than to the others, or, in other words, the points such that some func-

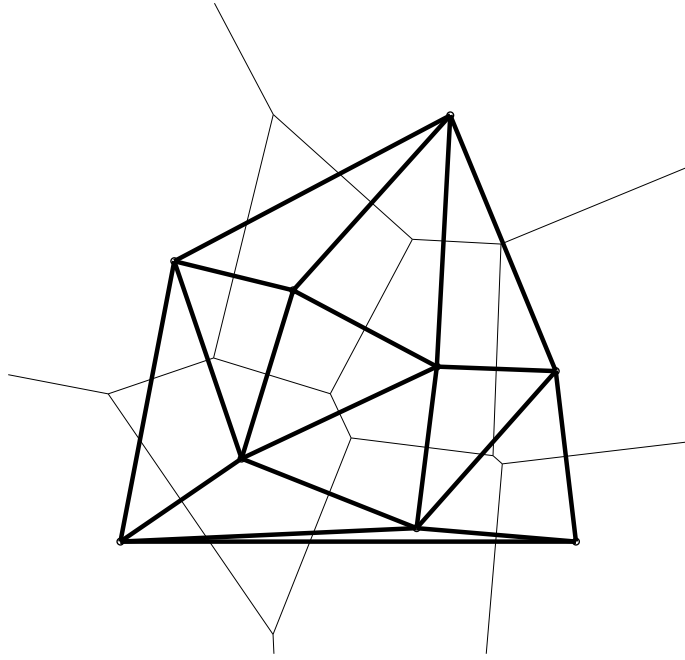


Figure 2: The Delaunay triangulation of a set of 9 points.

tion has a lower value than any other function. Studying how the unique properties of Voronoi diagrams and Delaunay triangulations are modified or extended in this broader context provides insights useful for both algorithmic improvements of their computation and better understanding of their appearance in natural contexts. See Fig. 3 and Fig. 4 for two examples of such diagrams.

Dual Contexts

Such diagrams occur in many scientific contexts, ranging from natural sciences, where they appear in nature, both in animals and inanimate objects, to modelization, where they appear as a tool for describing the closeness relationship between objects, which can be used for evaluating the interactions between these objects.

A reason for the appearance of Voronoi diagrams in nature is that they typically model growth phenomena. Euclidean Voronoi diagrams correspond to constant speed growth, while most other kind of Voronoi diagrams model more complex system, with non uniform speed, either constant for each site and depending on the site, or variable for each site, depending for example on its current size. These models explain that the giraffe fur (see Fig. 5), the turtle shell (see Fig. 6), bacteria colonies (see Fig. 7 and see Fig. 8), all feature the typical Euclidean Voronoi pattern or the more general Voronoi patterns. Giraffes and turtles present Euclidean patterns, showing that uniform speed growth processes with a common starting time can modelize the generation of the pattern. *Escherichia coli* colonies (see Fig. 7) present an Apollonius diagram pattern, showing that the growth processes have synchronized speeds, while the *Proteus mirabilis* colonies (see Fig. 8) present a Moebius diagram pattern, showing that the growth processes have independant speeds.

These two perspectives, natural sciences and experimental sciences, give an idea of two very typical uses that Voronoi diagrams and associated structures have in

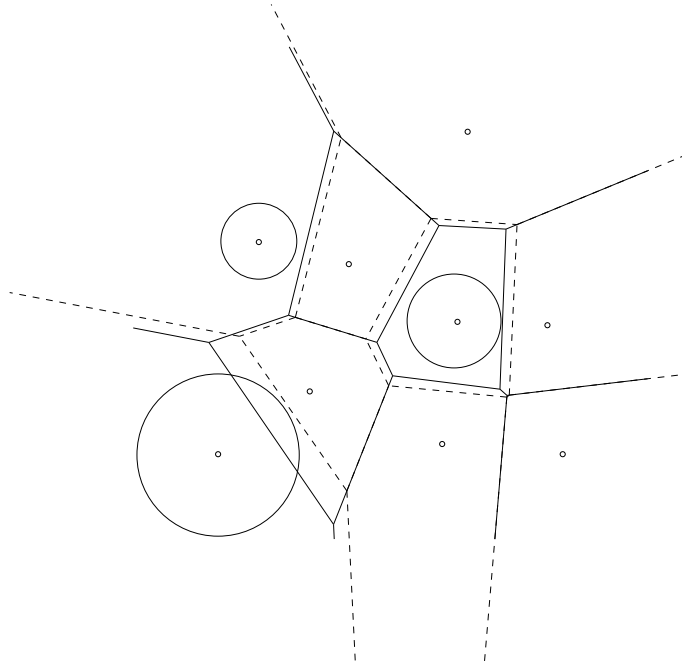


Figure 3: The power diagram of a set of 9 weighted points.

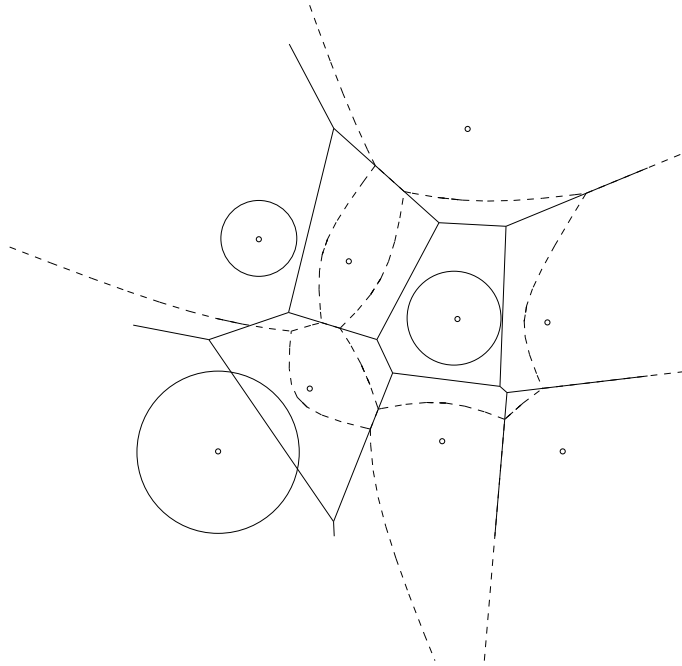


Figure 4: The Apollonius diagram of a set of 9 circles.

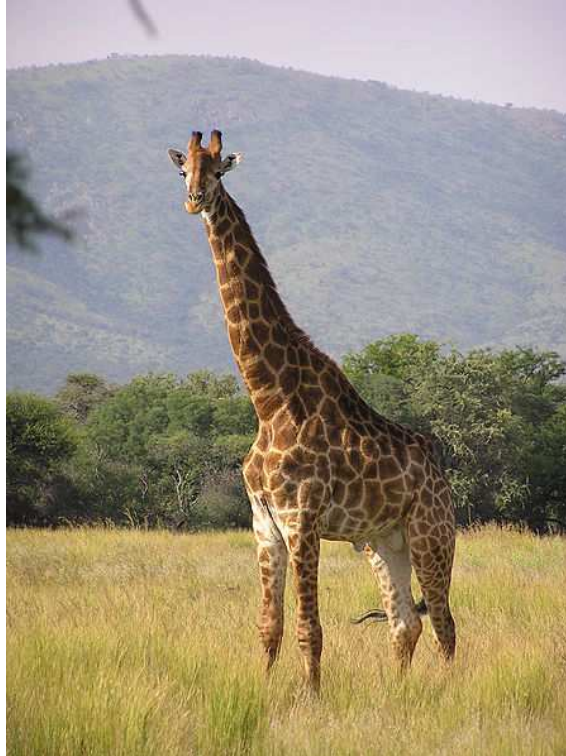


Figure 5: Giraffe and its Voronoi fur.



Figure 6: A giant turtle and its Voronoi shell.

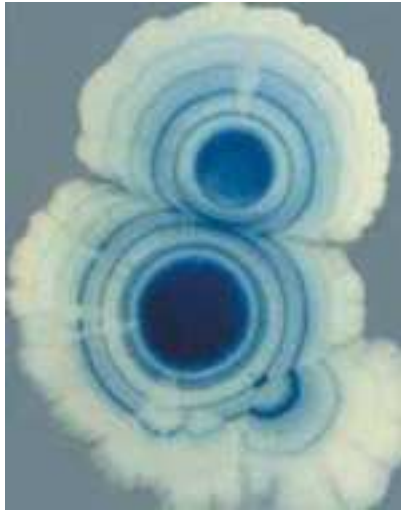


Figure 7: Two *E.coli* colonies, growing into an Apollonius diagram



Figure 8: Three *Proteus mirabilis* colonies, growing into a Moebius diagram

computer science. On the one hand, it is a natural object, whose existence should be interpreted and modeled (see Murray, *Mathematical Biology*, for examples). On the other hand, it is an underlying tool for computing and building models, thanks to its regularity properties, and the ease of computing it. In computer science, we can see the same structure, with the same two directions: Voronoi diagrams arise as a natural way of representing the closeness relations between objects, which can be studied for itself, and, for example, which makes them the tool of choice for the reconstruction of manifolds from pointsets. And they also appear as the classical way of building good underlying structures for numerical computations, as seen in meshing applications.

Approach

This thesis explores these two directions for general Voronoi diagrams. In a first part, Voronoi diagrams are studied as objects to be recognized, or computed. It could be presented as the natural study of Voronoi diagrams. The next part of this thesis focuses on the generation of anisotropic meshes, by using and adapting the tools presented in the first part. Anisotropic meshes are generated as duals of Voronoi diagrams, i.e. the generalization of the Delaunay triangulation, with specific properties. By generation of diagrams, we do not mean the mere computation of the diagram of a given set of objects, but the construction of suitable sets of objects, such that their Voronoi diagrams have the properties that are desired for the considered application, typically numerical computations.

The natural study of Voronoi diagrams has been pioneered by Aurenhammer and Klein. Aurenhammer gave the complete description of diagrams with linear bisectors, showing that any such diagram could be obtained as the power diagram of a set of weighted points, i.e. as a Voronoi diagram for a specific kind of objects and distances. Klein pushed this further into considering abstract Voronoi diagrams, where the primary definition of the diagram is given in terms of bisectors rather than objects and distances.

We may define many variants of Voronoi diagrams depending on the class of objects, the distance functions and the embedding space. Affine diagrams, i.e. diagrams whose cells are convex polytopes, are well understood. Their properties can be deduced from the properties of polytopes and they can be constructed efficiently. The first part of this thesis recalls the classical results on affine diagrams, and other more exotic varieties. This description is then unified into a description of abstract Voronoi diagrams, in a way derived and extended from Klein.

A recurring tool in this study of the classical types of Voronoi diagrams is the linearization of curved diagrams into higher dimensional ones. The anisotropic mesh generation method studied in the second part of this thesis builds upon this tool. But the important question here is to go from diagram generation to the dual mesh generation: the dual of a general Voronoi diagram is not always well defined. For this reason, we focused on desired local properties of the dual graph, which would be progressively satisfied everywhere in the diagram, thanks to the insertion of well chosen points into the diagram, until the dual would be well-defined. This first approach can be seen as the direct application of the linearization idea to the anisotropic mesh generation question. However, it did not allow us to solve all the challenges for defining a proven framework for anisotropic mesh generation in 3D.

Still, this first method underlined the usefulness of focusing on local properties of the dual graph itself (rather than properties of the diagram), while it was being

generated, even before being a valid mesh. Considering the seminal work of Labelle and Shewchuk, it also appeared that being able to define datastructures independently of the insertion order of the points would be crucial for proving the algorithm in a reasonably simple and robust way. As a result, we designed a new procedure, with these desired properties: rather than considering local properties of a global datastructure, we defined the intermediate datastructure as a union of local meshes, defined for a given set of points, independantly of the order of insertion. This union of local meshes would then be refined so that they could, in the end, be merged into one global mesh. This new framework proved simple and general enough to allow us to overcome the challenge of extending it to 3D in a provably good way. This is the subject of the third part of this thesis.

Finally, the last part of this text should be seen as an illustration of the mixed status of Voronoi diagrams, at the same time natural objects and engineered objects. In the context of greedy distributed routing, where Voronoi diagrams naturally appear to represent the possible connections between neighboring elements, we showed how specific local properties relating to routing could be designed, so that we could detect when the approximate evaluation of virtual coordinates for the communicating elements would be good enough for greedy routing to deliver. This idea of virtual coordinates was proposed by Papadimitriou and Ratajczak: the elements could be embedded in a virtual space where greedy routing would be guaranteed to work. Our contribution consisted in tailoring local properties of Voronoi diagrams so that detecting good enough embedding would be feasible locally, allowing distributed detection of this fact, but also reliance on approximate computations, particularly well suited to the distributed context.

Contribution

While the first part of this thesis relates the classical Voronoi diagrams taxonomy, it concludes with a general and unified presentation of the linearization method, and a new presentation of abstract Voronoi diagrams, in a simpler setting as the original one pioneered by Klein, and a more complete account of the relations between properties of abstract diagrams.

The second part focuses on anisotropic mesh generation. It first builds upon the definitions proposed by Labelle and Shewchuk, and shows how to obtain a practical meshing algorithm from the linearization scheme. The main contribution consists of obtaining a simple and generic computation scheme, by not computing the Voronoi diagram itself and its topological properties. Instead, we rely on dual computations only, and mesh validity checks.

Then, we consider a completely new approach to the anisotropic meshing question, relying on locally uniform anisotropy. This new approach is conceptually simple, and allows us to use proven and robust components, such as the Euclidean Delaunay triangulation computation. An important contribution is acknowledging the fact that procedures originally designed for sliver removal in 3D could be adapted for the removal of inconsistencies, in both 2D and 3D. This whole approach is generic enough, that we can envision further uses in difficult contexts, where classical uniform geometric datastructures are not versatile or flexible enough either to accomodate the variable parameters of the input, or to allow the computation of intrinsic datastructures, rather than global ones.

Finally, building upon the greedy routing framework developed by Papadimitriou

and Ratajczak, we present in the last part of this thesis an original method for the distributed computation of greedy embeddings. Most importantly, we show how local properties of the connectivity of power diagrams can be turned into global ones (which, in turn, have been shown to allow greedy routing). The Thurston circle packing algorithm was known to converge towards a solution to the problem we were considering. The design of such local properties allowed us to turn this approximation algorithm into an exact, proven algorithm for our purpose. Another important aspect of this work was the adaptation of the local properties to the case of non triangulated communication graphs, which amounted to dealing with degenerate triangulations. The difficulties arose from the fact that one cannot expect an approximation algorithm to reach a degenerate configuration in finite time. We showed how connectivity conditions can be adapted and made more flexible while retaining the same global consequences.

Applications

Mesh generation is a natural tool for numerical computations. Shewchuk detailed why anisotropic meshes should be preferred, in order to allow the most precise computations, for a fixed number of elements. However, the potential applications of the work presented in this thesis are broader.

Curved Voronoi diagrams are also used for modeling antenna placement, The work presented in this thesis is at the border between theory and applications. Anisotropic meshes are typically used for numerical computations (see Shewchuk for a detailed discussion of this fact), and can also be used during communication network design, for computing the regions of influence of antennas and optimizing the antennas' placement.

The diagrams and meshes that we study and compute appear in practical problems such as numerical computing, communication networks design, where curved diagrams allow to represent the antennas' zones of influence, sensor networks.

Part I

Voronoi Diagrams

AFFINE DIAGRAMS

OVERVIEW

An important class of Voronoi diagrams is the class of affine diagrams, whose bisectors are hyperplanes. The following chapter presents them and reviews some of their classical properties. Euclidean Voronoi diagrams of finite point sets are affine diagrams. Other examples of affine diagrams are the so-called power (or Laguerre) diagrams, where the objects are no longer points but hyperspheres and the Euclidean distance is replaced by the power of a point to a hypersphere. After presenting the general concept of minimization diagram in Sect. 1.1, we recall well-known facts about affine diagrams in Sect. 1.2. In particular, we characterize affine diagrams and establish a connection between affine diagrams and polytopes. As a consequence, we obtain tight combinatorial bounds and efficient algorithms. We also obtain a dual structure that is a triangulation under a general position assumption.

Contents

1.1 Lower Envelopes and Minimization Diagrams	24
1.2 Affine Voronoi Diagrams	25
1.2.1 Euclidean Voronoi Diagrams of Points	25
1.2.2 Delaunay Triangulation	27
1.2.3 Power Diagrams	30

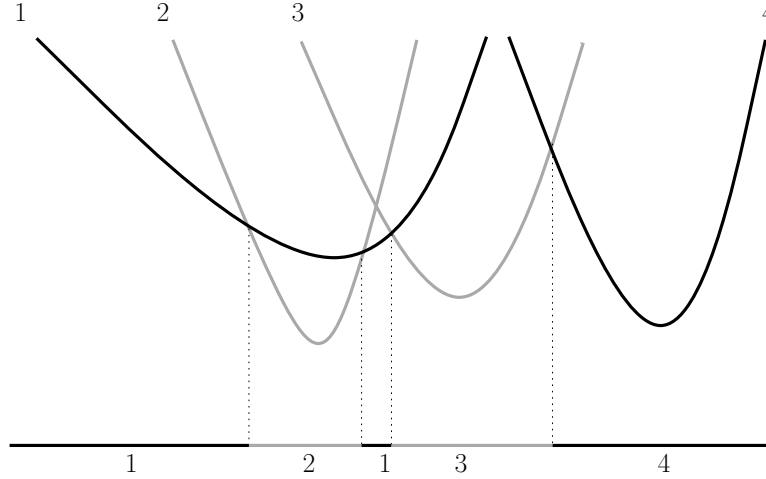


Figure 1.1: The lower envelope of a set of univariate functions. The minimization diagram is drawn on the horizontal line with the corresponding indices. The face of index $\{1\}$ consists of two components.

1.1 LOWER ENVELOPES AND MINIMIZATION DIAGRAMS

Let $\mathcal{F} = \{f_1, \dots, f_n\}$ be a set of d -variate continuous functions defined over \mathbb{R}^d . The *lower envelope* of \mathcal{F} is defined as

$$\mathcal{F}^- = \min_{1 \leq i \leq n} f_i.$$

From \mathcal{F} and \mathcal{F}^- , we define a natural partition of \mathbb{R}^d called the *minimization diagram* of \mathcal{F} . For a point $x \in \mathbb{R}^d$, we define the *index set* $I(x)$ of x as the set of all indices i such that $\mathcal{F}^-(x) = f_i(x)$. An equivalence relation noted \equiv can then be defined between two points of \mathbb{R}^d if they have the same index set:

$$x \equiv y \iff I(x) = I(y).$$

The equivalence classes \mathbb{R}^d / \equiv are relatively open sets that cover \mathbb{R}^d . Their closures are called the *faces* of the minimization diagram of \mathcal{F} (see Fig. 1.1). The index set of a face is defined as the largest subset of indices common to all the points of the face. Conversely, the face of index set I is the set of all points x such that $I \subset I(x)$.

Observe that the faces of this diagram are not necessarily contractible nor even connected. In particular, a 0-dimensional face may consist of several distinct points.

Lower envelopes and minimization diagrams have been well studied. We recall an important result due to Sharir [38] which provides an almost optimal result when the f_i are supposed to be multivariate polynomials of constant maximum degree.

Theorem 1.1.1 (Sharir). *The number of faces of the minimization diagram of a set \mathcal{F} of n multivariate polynomials of constant maximum degree η is $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$, where the constant of proportionality depends on ε , d and η . The vertices, edges and 2-faces of the diagram can be computed in randomized expected time $O(n^{d+\varepsilon})$ for any $\varepsilon > 0$.*

This general result is close to optimal in the worst-case: consider n ellipsoids inscribed in a $(d-1)$ -sphere S and intersecting S along great n $(d-2)$ -spheres $\sigma_1, \dots, \sigma_n$.

The arrangement of the σ_i has $\Theta(n^{d-1})$ faces. Since the non-bounded faces of the Euclidean Voronoi diagram of n objects are in 1-1 correspondence with the faces of their convex hull, we get a lower bound on the size of the Voronoi diagram of n ellipsoids of \mathbb{R}^d .

It has been improved in some special cases. For more information and other related results, one should consult the book by Sharir and Agarwal [39].

Voronoi diagrams, in their general setting, are just minimization diagrams of a finite set of continuous functions. This general definition encompasses the more traditional definition of Voronoi diagrams where the functions are defined as distance functions to a finite set of objects. Consider a set of objects $\mathcal{O} = \{o_1, \dots, o_n\}$. To each object o_i is attached a continuous function δ_i that measures the distance from a point x of \mathbb{R}^d to o_i . In the simplest case, \mathcal{O} is a finite set of points and $\delta_i(x)$ is the Euclidean distance from x to o_i . The Voronoi diagram of \mathcal{O} is defined as the minimization diagram of $\Delta = \{\delta_1, \dots, \delta_n\}$. The concept of Voronoi diagram has been generalized and various other diagrams have been defined by considering more general objects and other distance functions. Distance is then not to be taken literally. The function δ_i is only supposed to be continuous.

Theorem 1.1.1 provides very general bounds on the complexity of Voronoi diagrams. However, this result calls for improvement. First, in some special cases, much better bounds can be obtained by other approaches to be discussed later in this chapter. In particular, we will see that the most popular Euclidean Voronoi diagram of points has a much smaller combinatorial complexity than the one given in the theorem.

A second issue is the algorithmic complexity. The algorithm mentioned in the theorem fails to provide a complete description of the diagram since only faces of dimensions up to 2 are computed.

Moreover, the implementation of such an algorithm remains a critical issue. Computing lower envelopes of algebraic functions is a formidable task, even in the simplest cases, e.g. quadratic bi-variate functions.

1.2 AFFINE VORONOI DIAGRAMS

We first introduce Euclidean Voronoi diagrams of points and establish a correspondence between those diagrams and convex polyhedra in one dimension higher. Polarity allows to associate to a Voronoi diagram its dual cell complex, called a Delaunay triangulation.

Almost identical results can be obtained for power (or Laguerre) diagrams where points are replaced by hyperspheres and the Euclidean distance by the power of a point to a hypersphere. Power diagrams constitute a natural extension of Euclidean Voronoi diagrams and are still affine diagrams. In fact, we will see that any affine diagram is the power diagram of a finite set of hyperspheres.

1.2.1 Euclidean Voronoi Diagrams of Points

Let $\mathcal{P} = \{p_1, \dots, p_n\}$ be a set of points of \mathbb{R}^d . To each p_i , we associate its Voronoi region $V(p_i)$

$$V(p_i) = \{x \in \mathbb{R}^d : \|x - p_i\| \leq \|x - p_j\|, \forall j \leq n\}.$$

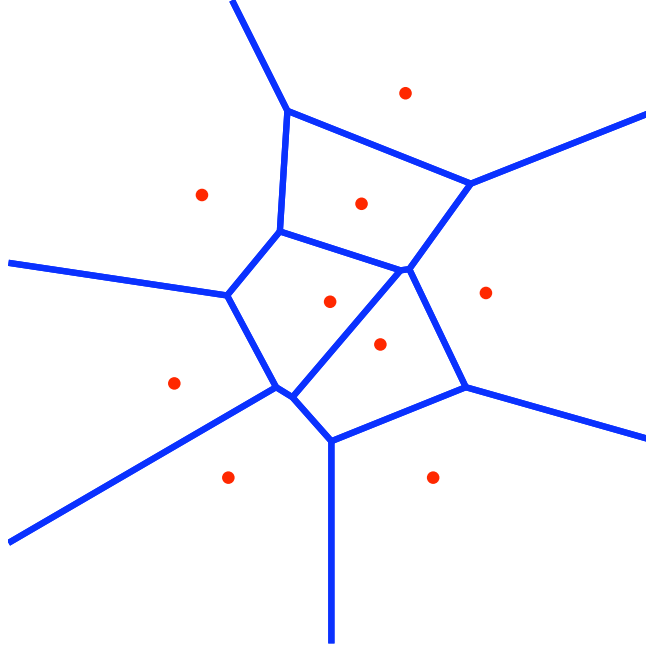


Figure 1.2: The Voronoi diagram of a set of 9 points.

The region $V(p_i)$ is the intersection of $n-1$ half-spaces. Each such half-space contains p_i and is bounded by the bisector of p_i and some other point of \mathcal{P} . Since the bisectors are hyperplanes, $V(p_i)$ is a convex polyhedron, possibly unbounded.

The *Euclidean Voronoi diagram* of \mathcal{P} , noted $\text{Vor}(\mathcal{P})$, is the cell complex whose cells are the Voronoi regions and their faces. Equivalently, the Euclidean Voronoi diagram of \mathcal{P} can be defined as the minimization diagram of the distance functions $\delta_i, \dots, \delta_n$, where

$$\delta_i(x) = \|x - p_i\|.$$

In other words, the Euclidean Voronoi diagram of \mathcal{P} is the minimization diagram of a set of functions whose graphs are vertical¹ cones of revolution of \mathbb{R}^{d+1} . Since minimizing $\|x - p_i\|$ over i is the same as minimizing $(x - p_i)^2$, the Euclidean Voronoi diagram of \mathcal{P} can alternatively be defined as the minimization diagram of the smooth functions $(x - p_i)^2$ whose graphs are translated copies of a vertical paraboloid of revolution of \mathbb{R}^{d+1} .

Observing further that, for any x , $\arg \min_i (x - p_i)^2 = \arg \min_i (-2p_i \cdot x + p_i^2)$, we obtain that the Euclidean Voronoi diagram of \mathcal{P} is the minimization diagram of a set of affine functions, namely the functions

$$d_i(x) = -2p_i \cdot x + p_i^2$$

whose graphs are hyperplanes of \mathbb{R}^{d+1} . Let us call h_{p_i} , $i = 1, \dots, n$, those hyperplanes and let $h_{p_i}^-$ denote the half-space lying below h_{p_i} . The minimization diagram of the d_i is obtained by projecting the polyhedron

$$\mathcal{V}(\mathcal{P}) = h_{p_1}^- \cap \dots \cap h_{p_n}^-.$$

vertically onto \mathbb{R}^d . See Fig. 1.3.

¹By vertical, we mean that the axis of revolution is perpendicular to \mathbb{R}^d .

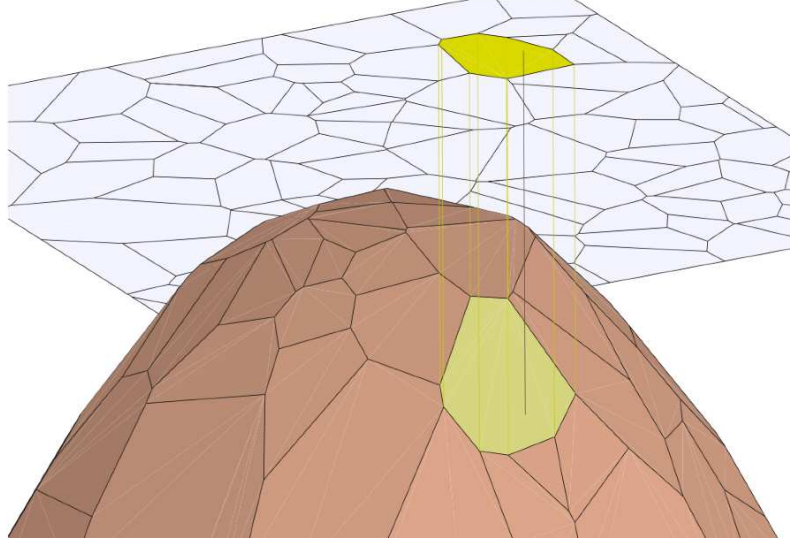


Figure 1.3: The polyhedron, with one of its faces projected.

We have therefore proved the following theorem:

Theorem 1.2.1. *The faces of the Euclidean Voronoi diagram $\text{Vor}(\mathcal{P})$ of a set of points \mathcal{P} are the vertical projections of the faces of the convex polyhedron $\mathcal{V}(\mathcal{P})$.*

1.2.2 Delaunay Triangulation

Two cell complexes V and D are said to be dual if there exists an involutive correspondence between the faces of V and the faces of D that reverses the inclusions, i.e. for any two faces f and g of V , their dual faces f^* and g^* satisfy: $f \subset g \Rightarrow g^* \subset f^*$. We introduce now a cell complex that is *dual* to the Voronoi diagram of a finite set of points \mathcal{P} .

We assume for now that the set of points \mathcal{P} is in *general position*, which means that no subset of $d + 2$ points of \mathcal{P} lie on a same hypersphere. Let f be a face of dimension k of the Voronoi diagram of \mathcal{P} . All points in the interior of f have the same subset \mathcal{P}_f of closest points in \mathcal{P} . The face dual to f is the convex hull of \mathcal{P}_f . The *Delaunay triangulation* of \mathcal{P} , noted $\text{Del}(\mathcal{P})$, is the cell complex consisting of all the dual faces. Because points of \mathcal{P} are assumed to be in general position, $|\mathcal{P}_f| = d - k + 1$, all the faces of $\text{Del}(\mathcal{P})$ are simplices and $\text{Del}(\mathcal{P})$ is a simplicial complex. The fact that $\text{Del}(\mathcal{P})$ is indeed a triangulation, i.e. a simplicial complex embedded in \mathbb{R}^d and covering the convex hull of \mathcal{P} , will be proved now using a duality between points and hyperplanes in the so-called space of spheres.

Polarity

Let σ be the hypersphere of \mathbb{R}^d of equation

$$\sigma(x) = (x - c)^2 - r^2 = x^2 - 2c \cdot x + s = 0,$$

where c is the center of σ , r its radius and $s = \sigma(0) = c^2 - r^2$.

We define the following bijective mapping

$$\phi : \sigma \in \mathbb{R}^d \longrightarrow \phi(\sigma) = (c, -s) \in \mathbb{R}^{d+1}$$

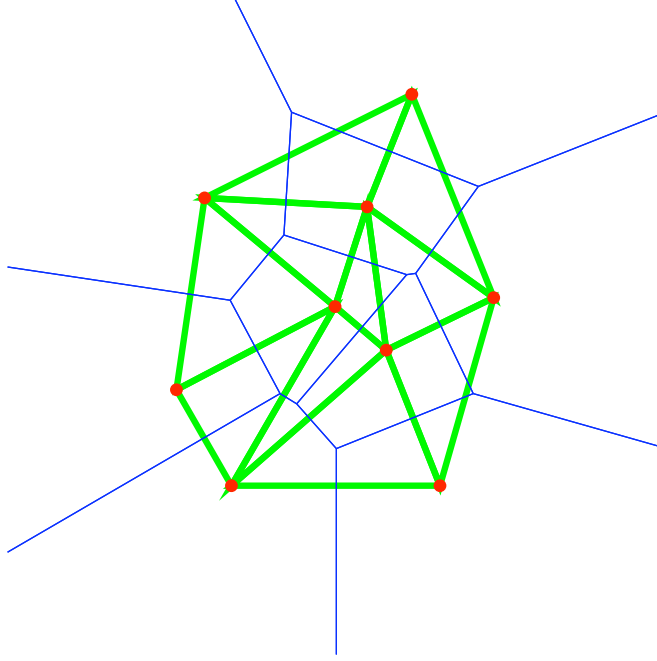


Figure 1.4: The Delaunay triangulation of a point set (in bold) and its dual Voronoi diagram (thin lines).

that maps a hypersphere of \mathbb{R}^d to a point of \mathbb{R}^{d+1} . We thus consider \mathbb{R}^{d+1} as the images by ϕ of the hyperspheres of \mathbb{R}^d and call \mathbb{R}^{d+1} *the space of spheres*. We note $\phi(p) = (p, -p^2)$ the image by ϕ of a point, considered as a hypersphere of radius 0. Observe that $\phi(p)$ is a point of the paraboloid \mathcal{Q} of \mathbb{R}^{d+1} of equation $x^2 + x_{d+1} = 0$. The points of \mathbb{R}^{d+1} that lie above \mathcal{Q} are images of imaginary hyperspheres whose squared radii are negative. The points below \mathcal{Q} are images of real hyperspheres.

We now introduced a mapping between points and hyperplanes of the space of spheres, known as *polarity*. Polarity associates to the point $\phi(\sigma)$ its *polar hyperplane* h_σ which is the hyperplane of \mathbb{R}^{d+1} of equation $2c \cdot x + x_{d+1} - s = 0$. Observe that the intersection of h_σ with \mathcal{Q} projects vertically onto σ , and that h_σ is the affine hull of the image by ϕ of the points of σ . If p is a point of \mathbb{R}^d , the polar hyperplane h_p of $\phi(p)$ is the hyperplane tangent to \mathcal{Q} at $\phi(p)$.

We deduce the remarkable following property: $x \in \sigma$ if and only if $\phi(x) = (x, -x^2) \in h_\sigma$ and σ encloses x if and only if $\phi(x) \in h_\sigma^+$, where h_σ^+ (resp. h_σ^-) denotes the closed half-space above (resp. below) h_σ . Indeed

$$\begin{aligned} \sigma(x) = 0 &\iff x^2 - 2c \cdot x + s = 0 \iff \phi(x) \in h_\sigma \\ \sigma(x) < 0 &\iff x^2 - 2c \cdot x + s < 0 \iff \phi(x) \in \text{int } h_\sigma^+, \end{aligned}$$

where $\text{int } h_\sigma^+$ denotes the open half-space above h_σ .

Polarity is an involution that preserves incidences and reverses inclusions. Indeed, if σ and σ' are two hyperspheres, we have

$$\begin{aligned} \phi(\sigma) \in h_{\sigma'} &\iff 2c' \cdot c - s - s' = 0 \iff \phi(\sigma') \in h_\sigma \\ \phi(\sigma) \in h_{\sigma'}^+ &\iff 2c' \cdot c - s - s' > 0 \iff \phi(\sigma') \in \text{int } h_\sigma^+. \end{aligned}$$

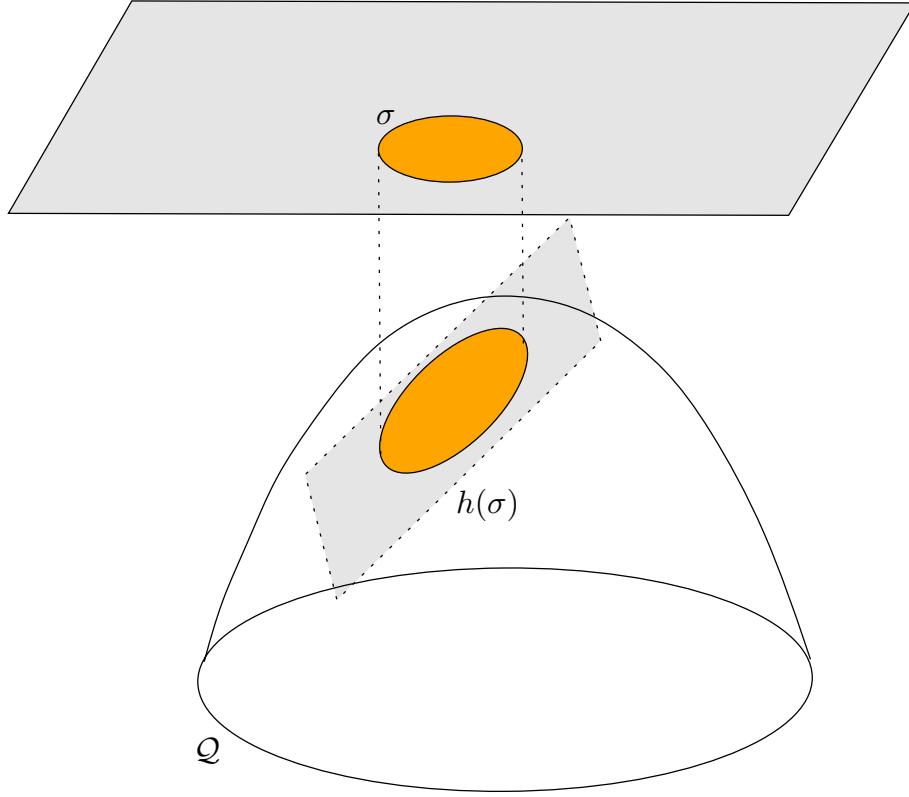


Figure 1.5: The polar hyperplane of a sphere.

Consider now a set $\mathcal{P} = \{p_1, \dots, p_n\}$ of n points and let $\mathcal{V}(\mathcal{P})$ denote, as in Sect. 1.2.1, the convex polyhedron defined as the intersection of the n half-spaces below the n polar hyperplanes h_{p_1}, \dots, h_{p_n} . Let f be a face of $\mathcal{V}(\mathcal{P})$ and assume that f is contained in $k + 1$ hyperplanes among the h_{p_i} . Without loss of generality, we denote those hyperplanes $h_{p_1}, \dots, h_{p_{k+1}}$. Let σ denote a hypersphere of \mathbb{R}^d such that $\phi(\sigma)$ belongs to the relative interior of f . From the above discussion, we have

$$\forall i, \quad 1 \leq i \leq k + 1, \quad \phi(\sigma) \in h_{p_i} \iff \phi(p_i) \in h_\sigma \quad (1.1)$$

$$\forall i, \quad k + 1 < i \leq n, \quad \phi(\sigma) \in \text{int } h_{p_i}^- \iff \phi(p_i) \in \text{int } h_\sigma^- \quad (1.2)$$

Given a convex polyhedron \mathcal{D} , we say that a hyperplane h *supports* \mathcal{D} if $\mathcal{D} \cap h$ is non-empty and \mathcal{D} is included in one of the two halfspaces, h^+ or h^- , bounded by h . If h is a supporting hyperplane of \mathcal{D} , $g = \mathcal{D} \cap h$ is a face of \mathcal{D} . If $\mathcal{D} \subset h^-$, g is called an upper face of \mathcal{D} . The collection of all upper faces of \mathcal{D} constitutes the *upper hull* of \mathcal{D} , which we denote by $\partial^+ \mathcal{D}$.

Let $\mathcal{D}(\mathcal{P}) = \text{Conv}(\phi(\mathcal{P}))$ be the convex hull of the set $\phi(\mathcal{P})$ and consider again the face f of $\mathcal{V}(\mathcal{P})$ defined above. Write $\mathcal{P}_f = \{p_1, \dots, p_{k+1}\}$. We deduce from (1.1) and (1.2) that, for any $\phi(\sigma)$ in the relative interior of f :

1. The hyperplane h_σ is a supporting hyperplane of $\mathcal{D}(\mathcal{P})$.
2. h_σ supports $\mathcal{D}(\mathcal{P})$ along the face $f^* = h_\sigma \cap \mathcal{D}(\mathcal{P}) = \text{Conv}(\phi(\mathcal{P}_f))$.
3. $\mathcal{D}(\mathcal{P}) \subset h_\sigma^-$ and f^* is a face of $\partial^+ \mathcal{D}(\mathcal{P})$.

To each face f of $\partial\mathcal{V}(\mathcal{P})$, we associate the face f^* of $\partial^+\mathcal{D}$ obtained as described above. This correspondence between the faces of $\partial\mathcal{V}(\mathcal{P})$ and the faces of $\partial^+\mathcal{D}(\mathcal{P})$ is bijective, preserves incidences and reverses inclusions, hence it is a duality.

The upper hull $\partial^+\mathcal{D}(\mathcal{P})$ projects vertically onto a cell complex of \mathbb{R}^d whose vertices are the points of \mathcal{P} . Because the projection is 1-1, this projected cell complex is properly embedded in \mathbb{R}^d and, since the projection preserves convexity, it covers the convex hull of \mathcal{P} . Under the general position assumption, the convex polyhedron $\mathcal{D}(\mathcal{P})$ is simplicial and the projected complex is a triangulation of \mathcal{P} . The duality between the faces of $\partial\mathcal{V}(\mathcal{P})$ and the faces of $\partial^+\mathcal{D}(\mathcal{P})$ implies that the projection of $\partial^+\mathcal{D}(\mathcal{P})$ is the Delaunay triangulation $\text{Del}(\mathcal{P})$ of \mathcal{P} introduced at the beginning of this section. This concludes the proof that, under the general position assumption, the Delaunay triangulation $\text{Del}(\mathcal{P})$ is a triangulation of \mathcal{P} . We have the following diagram:

$$\begin{array}{ccc} \partial\mathcal{V}(\mathcal{P}) = \partial(h_{p_1}^- \cap \dots \cap h_{p_n}^-) & \longleftrightarrow & \partial^+\mathcal{D}(\mathcal{P}) = \partial^+(\text{Conv}(\phi(\mathcal{P}))) \\ \downarrow & & \downarrow \\ \text{Voronoi Diagram } \text{Vor}(\mathcal{P}) & \longleftrightarrow & \text{Delaunay Triangulation } \text{Del}(\mathcal{P}) \end{array}$$

It follows from the above correspondence that the combinatorial complexity of the Delaunay triangulation of n points is the same as the combinatorial complexity of its dual Voronoi diagram. Moreover, the Delaunay triangulation of n points of \mathbb{R}^d can be deduced from the dual Voronoi diagram or vice versa in time proportional to its size. We also deduce from what precedes that computing the Delaunay triangulation of n points of \mathbb{R}^d reduces to constructing the convex hull of n points of \mathbb{R}^{d+1} . The following theorem is then a direct consequence of known results on convex hulls [11].

Theorem 1.2.2. *The combinatorial complexity of the Voronoi diagram of n points of \mathbb{R}^d and of their Delaunay triangulation is $\Theta\left(n^{\lfloor \frac{d+1}{2} \rfloor}\right)$. Both structures can be computed in optimal time $\Theta\left(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor}\right)$.*

The bounds in this theorem are tight. In particular, the Voronoi diagram of n points of \mathbb{R}^3 may be quadratic: if we take points on two non coplanar lines of \mathbb{R}^3 , say $n_1 + 1$ on one of the lines and $n_2 + 1$ on the other, their Voronoi diagram has $n_1 n_2$ vertices.

These bounds are worst-case bounds. Under some assumptions on the point distribution, better bounds can be obtained. For a set \mathcal{P} of n points uniformly distributed in a ball of \mathbb{R}^d , the combinatorial complexity of the Voronoi diagram of \mathcal{P} is $O(n)$ where the constant depends on the dimension d [21]. Other results are known for other point distributions [2, 3, 23].

In the discussion above, we have assumed that the points of \mathcal{P} were in general position. If this is not the case, some faces of $\mathcal{D}(\mathcal{P})$ are not simplices, and the complex $\partial^+\mathcal{D}(\mathcal{P})$ projects vertically onto a cell complex, dual to the Voronoi diagram and called the *Delaunay complex*. The faces of the *Delaunay complex* are convex and any triangulation obtained by triangulating those faces is called a Delaunay triangulation. Since there are several ways of triangulating the faces of the Delaunay complex, the Delaunay triangulation of \mathcal{P} is no longer unique.

1.2.3 Power Diagrams

A construction similar to what we did for the Euclidean Voronoi diagrams of points and their dual Delaunay triangulations can be done for the so-called power or Laguerre diagrams. Here we take as our finite set of objects a set of hyperspheres

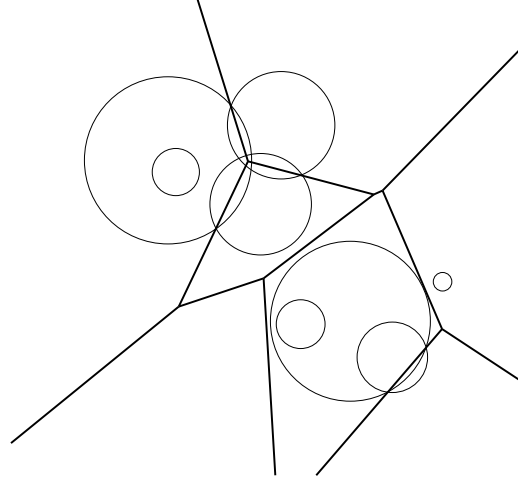


Figure 1.6: A power diagram.

(instead of points) and consider as distance function of a point x to a hypersphere σ the power of x to σ . As we will see, the class of power diagrams is identical to the class of affine diagrams, i.e. the diagrams whose bisectors are hyperplanes.

Definition of Power Diagrams

We call *power* of a point x to a hypersphere σ of center c and radius r the real number

$$\sigma(x) = (x - c)^2 - r^2.$$

Let $S = \{\sigma_1, \dots, \sigma_n\}$ be a set of hyperspheres of \mathbb{R}^d . We denote by c_i the center of σ_i , r_i its radius, $\sigma_i(x) = (x - c_i)^2 - r_i^2$ the power function to σ_i , and $s_i = c_i^2 - r_i^2$ the power of the origin. To each σ_i , we associate the region $\text{Pow}(\sigma_i)$ consisting of the points of \mathbb{R}^d whose power to σ_i is not larger than their power to the other hyperspheres of S :

$$\text{Pow}(\sigma_i) = \{x \in \mathbb{R}^d : \sigma_i(x) \leq \sigma_j(x), 1 \leq j \leq n\}.$$

The set of points that have equal power to two hyperspheres σ_i and σ_j is a hyperplane, noted π_{ij} , called *the radical hyperplane* of σ_i and σ_j . Hyperplane π_{ij} is orthogonal to the line joining the centers of σ_i and σ_j . We denote by π_{ij}^i the half-space bounded by π_{ij} consisting of the points whose power to σ_i is smaller than their power to σ_j . The region $\text{Pow}(\sigma_i)$ is the intersection of all half-spaces π_{ij}^i , $j \neq i$. If this intersection is not empty, it is a convex polyhedron, possibly not bounded. We call *power regions* the non empty regions $\text{Pow}(\sigma_i)$.

We define the *power diagram* of S , noted $\text{Pow}(S)$, as the cell complex whose cells are the power regions and their faces. When all hyperspheres have the same radius, their power diagram is identical to the Voronoi diagram of their centers.

Equivalently, the power diagram of S can be defined as the minimization diagram of the functions $\sigma_i, \dots, \sigma_n$. Observing that for any x

$$\arg \min_i \sigma_i(x) = \arg \min_i (-2c_i \cdot x + s_i),$$

we obtain that the power diagram of S is the minimization diagram of the set of affine functions

$$d_i(x) = -2p_i \cdot x + s_i$$

whose graphs are hyperplanes of \mathbb{R}^{d+1} . Let us call h_{σ_i} , $i = 1, \dots, n$, those hyperplanes and let $h_{\sigma_i}^-$ denote the half-space lying below h_{σ_i} . The minimization diagram of the δ_i is obtained by projecting vertically the convex polyhedron

$$\mathcal{L}(S) = h_{p_1}^- \cap \dots \cap h_{p_n}^-.$$

Theorem 1.2.3. *The faces of the power diagram $\text{Pow}(S)$ of S are the vertical projections of the faces of the convex polyhedron $\mathcal{L}(S)$.*

Power diagrams are very similar to Voronoi diagrams: the only difference is that the hyperplanes supporting the faces of $\mathcal{L}(S)$ are not necessarily tangent to the paraboloid \mathcal{Q} and that some hyperplane may not contribute a face. In other words, some hypersphere σ_i may have an empty power region (see the small circle in the upper left corner of Fig. 1.6).

By proceeding as in Sect. 1.2.2, we can define a convex polyhedron $\mathcal{R}(S)$ whose upper hull $\partial^+ \mathcal{R}(S)$ is dual to $\partial \mathcal{L}(S)$. The vertical projection of the faces of $\partial^+ \mathcal{R}(S)$ constitute the faces of a cell complex which, in general, is a simplicial complex. We call such a complex the *regular triangulation* of S and denote it by $\text{Reg}(S)$. We have the following diagram :

$$\begin{array}{ccc} \partial \mathcal{L}(S) = \partial (h_{\sigma_1}^- \cap \dots \cap h_{\sigma_n}^-) & \longleftrightarrow & \partial^+ \mathcal{R}(S) = \partial^+ \text{Conv}(\phi(S)) \\ \downarrow & & \downarrow \\ \text{Power diagram } \text{Pow}(S) & \longleftrightarrow & \text{Regular triangulation } \text{Reg}(S) \end{array}$$

We deduce the following theorem that states that computing the power diagram of n hyperspheres of \mathbb{R}^d (or equivalently its dual regular triangulation) has the same asymptotic complexity as computing the Euclidean Voronoi diagram or the Delaunay triangulation of n points of \mathbb{R}^d .

Theorem 1.2.4. *The combinatorial complexity of the power diagram of n hyperspheres of \mathbb{R}^d and of its dual regular triangulation are $\Theta\left(n \lfloor \frac{d+1}{2} \rfloor\right)$. Both structures can be computed in optimal time $\Theta\left(n \log n + n \lfloor \frac{d+1}{2} \rfloor\right)$.*

Affine Voronoi Diagrams

Euclidean Voronoi diagrams of points and power diagrams of hyperspheres are two examples of minimization diagrams whose bisectors are hyperplanes. It is interesting to classify Voronoi diagrams with respect to their bisectors. A first important class of Voronoi diagrams is the class of *affine diagrams* which consists of all Voronoi diagrams whose bisectors are hyperplanes.

In Sect. 2.2, we will prove that any affine Voronoi diagram of \mathbb{R}^d is identical to the power diagram of some set of hyperspheres of \mathbb{R}^d (Theorem 2.2.2), therefore showing that the class of affine Voronoi diagrams is identical to the class of power diagrams.

CURVED VORONOI DIAGRAMS

OVERVIEW

The previous chapter presented affine Voronoi diagrams. We now focus on more general Voronoi diagrams, with algebraic bisectors. We first present several classical variants of such diagrams, and some ways of computing them. While some of them have special properties which allow relatively easy computation, such as Möbius diagrams, all of them can be unified into a framework of abstract diagrams, simplified and adapted from Klein [28], and that we present in Section 2.2.1. On the one hand, we prove that all considered diagrams fit into this context of abstract diagrams, while on the other hand, we prove that the some necessary conditions for abstract diagrams to be one of the kind of diagrams that we have presented before are in fact sufficient conditions, hence proving the equivalence between these two approaches.

Contents

2.1 Voronoi Diagrams with Algebraic Bisectors	34
2.1.1 Möbius Diagrams	34
2.1.2 Anisotropic Diagrams	37
2.1.3 Apollonius Diagrams	39
2.2 Linearization	42
2.2.1 Abstract Diagrams	42
2.2.2 Inverse Problem	46
2.3 Introduction	51
2.4 Labelle and Shewchuk's Approach	51
2.5 Relation to Power Diagrams	54
2.6 Basic Operations and Primitives	55
2.7 Description of the Algorithm	57
2.8 Termination of the Algorithm	61
2.9 Conclusion and Future Work	67

2.1 VORONOI DIAGRAMS WITH ALGEBRAIC BISECTORS –

In this section, we introduce a first class of non-affine diagrams, namely the class of diagrams whose bisectors are algebraic hypersurfaces. We first consider the case of Möbius diagrams whose bisectors are hyperspheres and the case of *anisotropic* diagrams whose bisectors are quadratic hypersurfaces. These diagrams can be computed through linearization, a technique to be described in full generality in Sect. 2.2. Apollonius (or Johnson-Mehl) diagrams, although semi-algebraic and not algebraic, are also described in this section since they are closely related to Möbius diagrams and can also be linearized.

2.1.1 Möbius Diagrams

In this section, we introduce a class of non-affine Voronoi diagrams, the so-called Möbius diagrams, introduced by Boissonnat and Karavelas [6].

The class of Möbius diagrams includes affine diagrams. In fact, as we will see, the class of Möbius diagrams is identical to the class of diagrams whose bisectors are hyperspheres (or hyperplanes).

Definition of Möbius Diagrams

Let $\omega = \{\omega_1, \dots, \omega_n\}$ be a set of so-called *Möbius sites* of \mathbb{R}^d , where ω_i is a triple (p_i, λ_i, μ_i) formed of a point p_i of \mathbb{R}^d , and two real numbers λ_i and μ_i . For a point $x \in \mathbb{R}^d$, the distance $\delta_i(x)$ from x to the Möbius site ω_i is defined as

$$\delta_i(x) = \lambda_i(x - p_i)^2 - \mu_i.$$

Observe that the graph of δ_i is a paraboloid of revolution whose axis is vertical. The Möbius region of the Möbius site ω_i , $i = 1, \dots, n$, is

$$M(\omega_i) = \{x \in \mathbb{R}^d : \delta_i(x) \leq \delta_j(x), 1 \leq j \leq n\}.$$

Observe that a Möbius region may be non-contractible and even disconnected.

The minimization diagram of the δ_i is called the *Möbius diagram* of ω and noted $\text{Möb}(\omega)$ (see Fig. 2.1.1).

Möbius diagrams are generalizations of Euclidean Voronoi and power diagrams. In particular, if all λ_i are equal to some positive λ , the Möbius diagram coincides with the power diagram of a set of spheres $\{\sigma_i, i = 1, \dots, n\}$, where σ_i is the sphere centered at p_i of squared radius μ_i/λ . If all μ_i are equal and all λ_i are positive, then the Möbius diagram coincides with the so-called *multiplicatively weighted Voronoi diagram* of the weighted points $(p_i, \sqrt{\lambda_i})$.

The following lemma states that the bisector of two Möbius sites is a hypersphere (possibly degenerated in a point or in a hyperplane). Its proof is straightforward.

Lemma 2.1.1. *Let $\omega_i = \{p_i, \lambda_i, \mu_i\}$ and $\omega_j = \{p_j, \lambda_j, \mu_j\}$, $\omega_i \neq \omega_j$ be two Möbius sites. The bisector σ_{ij} of ω_i and ω_j is the empty set, a single point, a hypersphere or a hyperplane.*

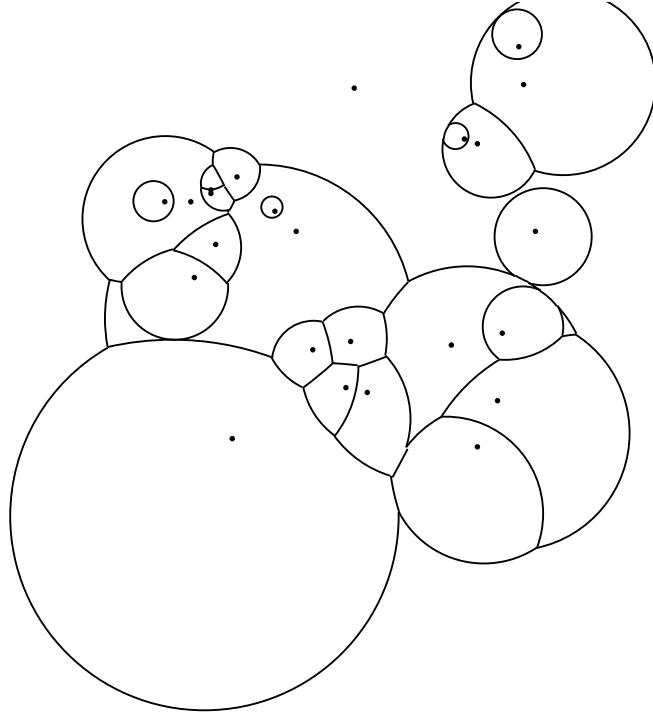


Figure 2.1: A Möbius diagram.

Möbius Diagrams and Power Diagrams

We now present an equivalence between Möbius diagrams in \mathbb{R}^d and power diagrams in \mathbb{R}^{d+1} . This result is a direct generalization of a similar result for multiplicatively weighted diagrams [4]. Given a cell complex \mathcal{C} covering a subspace X , we call *restriction of \mathcal{C} to X* the subdivision of X whose faces are the intersections of the faces of \mathcal{C} with X . The restriction of \mathcal{C} to X is denoted by \mathcal{C}_X . Note that the restriction \mathcal{C}_X is not, in general, a cell complex and that its faces may be non-contractible and even non-connected.

We associate to $\omega = \{\omega_1, \dots, \omega_n\}$ the set of hyperspheres $\Sigma = \{\Sigma_1, \dots, \Sigma_n\}$ of \mathbb{R}^{d+1} of equations

$$\Sigma_i(X) = X^2 - 2C_i \cdot X + s_i = 0,$$

where $C_i = (\lambda_i p_i, -\frac{\lambda_i}{2})$ and $s_i = \lambda_i p_i^2 - \mu_i$. We denote by \mathcal{Q} the paraboloid of \mathbb{R}^{d+1} of equation $x_{d+1} - x^2 = 0$.

Theorem 2.1.1 (Linearization). *The Möbius diagram $\text{Möb}(\omega)$ of ω is obtained by projecting vertically the faces of the restriction $\text{Pow}_{\mathcal{Q}}(\Sigma)$ of the power diagram of Σ to \mathcal{Q} .*

Proof. If $x \in \mathbb{R}^d$ is closer to ω_i than to ω_j with respect to ϑ , we have for all $j = 1, \dots, n$,

$$\begin{aligned} \lambda_i(x - p_i)^2 - \mu_i &\leq \lambda_j(x - p_j)^2 - \mu_j \\ \iff \lambda_i x^2 - 2\lambda_i p_i \cdot x + \lambda_i p_i^2 - \mu_i &\leq \lambda_j x^2 - 2\lambda_j p_j \cdot x + \lambda_j p_j^2 - \mu_j \\ \iff (x^2 + \frac{\lambda_i}{2})^2 + (x - \lambda_i p_i)^2 - \frac{\lambda_i^2}{4} - \lambda_i^2 p_i^2 + \lambda_i p_i^2 - \mu_i & \\ &\leq (x^2 + \frac{\lambda_j}{2})^2 + (x - \lambda_j p_j)^2 - \frac{\lambda_j^2}{4} - \lambda_j^2 p_j^2 + \lambda_j p_j^2 - \mu_j \\ \iff (X - C_i)^2 - r_i^2 &\leq (X - C_j)^2 - r_j^2 \end{aligned}$$

$$\iff \Sigma_i(X) \leq \Sigma_j(X)$$

where $X = (x, x^2) \in \mathcal{Q} \subset \mathbb{R}^{d+1}$, $C_i = (\lambda_i p_i, -\frac{\lambda_i}{2}) \in \mathbb{R}^{d+1}$ and $r_i^2 = \lambda_i^2 p_i^2 + \frac{\lambda_i^2}{4} - \lambda_i p_i^2 + \mu_i$. The above inequality shows that x is closer to ω_i than to ω_j if and only if X belongs to the power region of Σ_i in the power diagram of the hyperspheres Σ_j , $j = 1, \dots, n$. As X belongs to \mathcal{Q} and projects vertically onto x , we have proved the result. \square

Corollary 2.1.1.1. *Let Σ be a finite set of hyperspheres of \mathbb{R}^{d+1} , $\text{Pow}(\Sigma)$ its power diagram and $\text{Pow}_{\mathcal{Q}}(\Sigma)$ the restriction of $\text{Pow}(\Sigma)$ to \mathcal{Q} . The vertical projection of $\text{Pow}_{\mathcal{Q}}(\Sigma)$ is the Möbius diagram $\text{Möb}(\omega)$ of a set of Möbius sites of \mathbb{R}^d .*

Easy computations give ω .

Combinatorial and Algorithmic Properties

It follows from Theorem 2.1.1 that the combinatorial complexity of the Möbius diagram of n Möbius sites in \mathbb{R}^d is $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$. This bound is tight since Aurenhammer [4] has shown that it is tight for multiplicatively weighted Voronoi diagrams.

We easily deduce from the proof of the Linearization Theorem 2.1.1 an algorithm for constructing Möbius diagrams. First, we compute the power diagram of the hyperspheres Σ_i of \mathbb{R}^{d+1} , intersect each of the faces of this diagram with the paraboloid \mathcal{Q} and then project the result on \mathbb{R}^d .

Theorem 2.1.2. *Let ω be a set of n Möbius sites in \mathbb{R}^d , $d \geq 2$. The Möbius diagram $\text{Möb}(\omega)$ of ω can be constructed in worst-case optimal time $\Theta(n \log n + n^{\lfloor \frac{d}{2} \rfloor + 1})$.*

Another consequence of the linearization theorem is the fact that any Möbius diagram can be represented as a simplicial complex $T_{\mathcal{Q}}$ embedded in \mathbb{R}^{d+1} . $T_{\mathcal{Q}}$ is a sub-complex of the regular triangulation T dual to the power diagram $\text{Pow}(\Sigma_i)$ of the hyperspheres Σ_i . Since T is embedded in \mathbb{R}^{d+1} , $T_{\mathcal{Q}}$ is a simplicial complex of \mathbb{R}^{d+1} . More precisely, $T_{\mathcal{Q}}$ consists of the faces of T that are dual to the faces of $\text{Pow}_{\mathcal{Q}}(\Sigma)$, i.e. the faces of the power diagram that intersect \mathcal{Q} . We will call $T_{\mathcal{Q}}$ the *dual* of $\text{Pow}_{\mathcal{Q}}(\Sigma)$. Observe that since, in general, no vertex of $\text{Pow}(\Sigma)$ lies on \mathcal{Q} , $T_{\mathcal{Q}}$ is a d -dimensional simplicial complex (embedded in \mathbb{R}^{d+1}).

Moreover, if the faces of $\text{Pow}(\Sigma)$ intersect \mathcal{Q} transversally and along topological balls, then, by a result of Edelsbrunner and Shah [22], $T_{\mathcal{Q}}$ is homeomorphic to \mathcal{Q} and therefore to \mathbb{R}^d . It should be noted that this result states that the simplicial complex $T_{\mathcal{Q}}$ has the topology of \mathbb{R}^d . This result, however, is mainly combinatorial, and does not imply that the embedding of $T_{\mathcal{Q}}$ into \mathbb{R}^{d+1} as a sub-complex of the regular triangulation T may be projected in a 1-1 manner onto \mathbb{R}^d .

Spherical Voronoi Diagrams

Lemma 2.1.1 states that the bisectors of two Möbius sites is a hypersphere (possibly degenerated in a hyperplane). More generally, let us consider the Voronoi diagrams such that, for any two objects o_i and o_j of \mathcal{O} , the bisector $\sigma_{ij} = \{x \in \mathbb{R}^d, \delta_i(x) = \delta_j(x)\}$ is a hypersphere. Such a diagram is called a *spherical Voronoi diagram*.

In Sect. 2.2, we will prove that any spherical Voronoi diagram of \mathbb{R}^d is a Möbius diagram (Theorem 2.2.4).

Möbius transformations are the transformations that preserve hyperspheres. An example of a Möbius transformation is the inversion with respect to a hypersphere. If

the hypersphere is centered at c and has radius r , the inversion associates to a point $x \in \mathbb{R}^d$ its image

$$x' = c + \frac{r(x - c)}{(x - c)^2}.$$

Moreover, it is known that any Möbius transformation is the composition of up to four inversions [18]. An immediate consequence of Theorem 2.2.4 is that the set of Möbius diagrams in \mathbb{R}^d is stable under Möbius transformations, hence their name.

Möbius Diagrams on Spheres

Given a set ω of n Möbius sites of \mathbb{R}^{d+1} , the restriction of their Möbius diagram to a hypersphere \mathbb{S}^d is called a Möbius diagram on \mathbb{S}^d . Such a diagram is also the restriction of a power diagram of hyperspheres of \mathbb{R}^{d+1} to \mathbb{S}^d , and the converse is also true. This easily follows by considering the linearization through intersecting a sphere and applying the stereographic projection, instead of intersecting a paraboloid and applying the vertical projection.

We define spherical diagrams on \mathbb{S}^d as the diagrams on \mathbb{S}^d whose bisectors are hyperspheres of \mathbb{S}^d and that satisfy the two properties detailed in Sect. 2.2.1 and 2.2.2. It follows that the restriction of a Möbius diagram, i.e. a Möbius diagram on \mathbb{S}^d , is a spherical diagram.

Let us now prove the converse: any spherical diagram on \mathbb{S}^d is a Möbius diagram on \mathbb{S}^d . Let h be a hyperplane of \mathbb{R}^{d+1} . The stereographic projection that maps \mathbb{S}^d to h maps any spherical diagram \mathcal{D} on \mathbb{S}^d to some spherical diagram \mathcal{D}' on h . Theorem 2.2.4 implies that this \mathcal{D}' is in fact a Möbius diagram. Since the linearization theorem and its corollary still hold if one replaces the paraboloid \mathcal{Q} by any hypersphere of \mathbb{R}^{d+1} and the vertical projection by the corresponding stereographic projection, it follows that \mathcal{D} , which is the image of \mathcal{D}' by the inverse of the stereographic projection, is the restriction of some power diagram of \mathbb{R}^{d+1} to \mathbb{S}^d . The result follows.

2.1.2 Anisotropic Diagrams

The definition of anisotropic Voronoi diagrams presented in this section is a slight extension of the definition proposed by Labelle and Shewchuk [33]. This extended version of anisotropic Voronoi diagrams is alternatively called *flower power diagrams*, a term coined by Shewchuk, and a tribute to their seemingly hippy roots. The objects are points and the distance to a point is a quadratic form with an additive weight.

Anisotropic diagrams appear to be a natural generalization of Möbius diagrams and reduce to Möbius diagrams when the matrices are taken to be a scalar times the identity matrix. As will be shown, the class of anisotropic diagrams is identical to the class of diagrams whose bisectors are quadratic hypersurfaces.

Definition and linearization

Consider a finite set of anisotropic sites $S = \{s_1, \dots, s_n\}$. Each site s_i , $i = 1, \dots, n$, is a triple (p_i, M_i, π_i) formed by a point $p_i \in \mathbb{R}^d$, a $d \times d$ symmetric positive definite matrix M_i and a scalar weight π_i . The distance $\delta_i(x)$ of point $x \in \mathbb{R}^d$ to site s_i is defined by

$$\delta_i(x) = (x - p_i)^t M_i (x - p_i) - \pi_i.$$

The anisotropic Voronoi region of site s is then defined as

$$AV(s_i) = \{x \in \mathbb{R}^d, \delta_i(x) \leq \delta_j(x), \forall 1 \leq j \leq n\},$$

The *anisotropic Voronoi diagram* is the minimization diagram of the functions $\delta_i(x)$.

Let $D = \frac{d(d+3)}{2}$. To each point $x = (x_1, \dots, x_d) \in \mathbb{R}^d$, we associate the two points

$$\begin{aligned}\tilde{\phi}(x) &= (x_r x_i, 1 \leq r \leq s \leq d) \in \mathbb{R}^{\frac{d(d+1)}{2}} \\ \hat{\phi}(x) &= (x, \tilde{\phi}(x)) \in \mathbb{R}^D,\end{aligned}$$

and we denote by \mathcal{Q} the d -manifold of \mathbb{R}^D defined as

$$\mathcal{Q} = \left\{ \hat{\phi}(x), x \in \mathbb{R}^d \right\}.$$

To each site $s_i = (p_i, M_i, \pi_i) \in S$, we associate:

1. the point $\tilde{m}_i \in \mathbb{R}^{\frac{d(d+1)}{2}}$ defined as

$$\begin{aligned}\tilde{m}_i^{u,u} &= -\frac{1}{2}M_i^{u,u}, \text{ for } 1 \leq u \leq d; \\ \tilde{m}_i^{u,v} &= -M_i^{u,v}, \text{ for } 1 \leq u < v \leq d,\end{aligned}$$

2. the point $\hat{p}_i = (M_i p_i, \tilde{m}_i)$,

3. the sphere σ_i of center \hat{p}_i and radius $\sqrt{\|\hat{p}_i\|^2 - p_i^t M_i p_i - \pi_i}$.

Let Π be the projection $\hat{y} = (y, \tilde{y}) \in \mathbb{R}^D \mapsto y \in \mathbb{R}^d$ and let Σ be the set of spheres σ_i , $i = 1, \dots, n$.

Theorem 2.1.3 (Linearization). *The anisotropic diagram of S is the image by Π of the restriction of the power diagram $\text{Pow}(\Sigma)$ to the d -manifold \mathcal{Q} .*

Proof. We have

$$\begin{aligned}\delta_i(x) &= (x - p_i)^t M (x - p_i) - \pi_i \\ &= x^t M_i x - 2p_i^t M_i x + p_i^t M_i p_i - \pi_i \\ &= -2\hat{p}_i^t \hat{\phi}(x) + p_i^t M_i p_i - \pi_i\end{aligned}$$

This implies that $\delta_i(x) < \delta_j(x)$ if and only if

$$(\hat{\phi}(x) - \hat{p}_i)^2 - (\hat{p}_i^2 - p_i^t M_i p_i - \pi_i) < (\hat{\phi}(x) - \hat{p}_j)^2 - (\hat{p}_j^2 - p_j^t M_j p_j - \pi_j)$$

Hence, x is closer to s_i than to s_j if and only if the power of $\hat{\phi}(x)$ to σ_i is smaller than its power to σ_j . Equivalently, a point $\hat{\phi}(x) \in \mathcal{Q}$ belongs to the power cell of $\sigma(s_i)$ if and only if its projection $x = \Pi(\hat{\phi}(x))$ belongs to the anisotropic Voronoi region $AV(s_i)$. \square

We easily deduce the following theorem.

Theorem 2.1.4. *The Voronoi diagram of n anisotropic sites of \mathbb{R}^d can be computed in time $O(n^{\lfloor \frac{D+1}{2} \rfloor})$ where $D = \frac{d(d+3)}{2}$.*

This result is to be compared to Theorem 1.1.1 which provides a better combinatorial bound. We let as an open question to fill the gap between those two bounds.

Quadratic Voronoi Diagrams

The bisectors of anisotropic diagrams, as defined in the previous section, are quadratic hypersurfaces. A minimization diagram whose bisectors are hyperquadrics is called a *quadratic Voronoi diagram*. In Sect. 2.2, we will prove that any quadratic Voronoi diagram is the anisotropic Voronoi diagram of a set of anisotropic sites (Theorem 2.2.5).

2.1.3 Apollonius Diagrams

In this section, we present diagrams that are closely related to Möbius diagrams: namely, the Euclidean Voronoi diagrams of hyperspheres, also called Apollonius or Johnson-Mehl diagrams. Contrary to Möbius and anisotropic diagrams, the bisectors of Apollonius diagrams are not algebraic hypersurfaces since the bisector between two hyperspheres is only one sheet of a hyperboloid. As a consequence, Apollonius diagrams cannot be linearized in the same way as Möbius and anisotropic diagrams. Nevertheless, another linearization scheme can be applied, leading to interesting combinatorial and algorithmic results.

Definition of Apollonius Diagrams

Let us consider a finite set of weighted points $S = \{\sigma_0, \sigma_1, \dots, \sigma_n\}$ where $\sigma_i = (p_i, r_i)$, $p_i \in \mathbb{R}^d$ and $r_i \in \mathbb{R}$. We define the distance from x to σ_i as

$$\delta_i(x) = \|x - p_i\| - r_i.$$

This distance is also called the *additively weighted* distance from x to the weighted point σ_i . The minimization diagram of the distance functions δ_i , $i = 1, \dots, n$, is called the additively weighted Voronoi diagram, or the Apollonius diagram of S . We denote it by $\text{Apo}(S)$ (see Fig. 2.2).

The Apollonius region $A(\sigma_i)$ of σ_i is defined as

$$A(\sigma_i) = \{x \in \mathbb{R}^d, \delta_i(x) \leq \delta_j(x)\}.$$

It is easy to see that $A(\sigma_i)$ is either empty or star-shaped from p_i . The boundary of $A(\sigma_i)$ may have a complicated structure. In fact, as we will see, the boundary of $A(\sigma_i)$ has the same combinatorial structure as a Möbius diagram in \mathbb{R}^{d-1} .

Since the diagram is not changed if we replace all r_i by $r_i + r$ for any $r \in \mathbb{R}$, we can assume, without loss of generality, that all r_i are non negative. The weighted points are then hyperspheres and the distance to a weighted point is the signed Euclidean distance to the corresponding hypersphere, counted positively outside the hypersphere and negatively inside the hypersphere.

Observe that, in the plane, a vertex of an Apollonius diagram is the center of a circle tangent to three circles of S (assuming all r_i non negative). Computing such a point is known as Apollonius' Tenth Problem, hence the name of the diagram.

Apollonius Diagrams and Power Diagrams

The graph of the distance function $\delta_i(x)$ is the half-cone of revolution \mathcal{C}_i of equation

$$\mathcal{C}_i : x_{d+1} = \|x - p_i\| - r_i, \quad x_{d+1} + r_i \geq 0$$

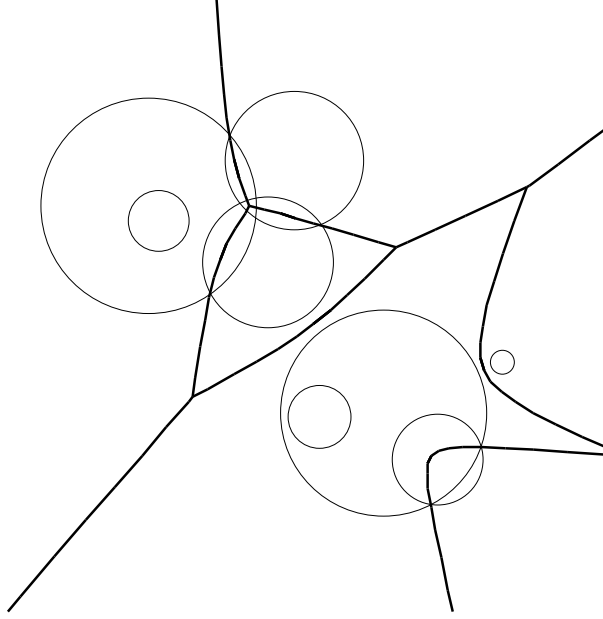


Figure 2.2: The Apollonius diagram of a set of circles. Compare with the power diagram of the same set of circles in Fig. 1.6.

The bisector of two hyperspheres of \mathcal{S} is thus the projection of the intersection of two half-cones. This intersection is a quadratic hypersurface (in fact, a sheet of a two sheet hyperboloid) contained in a hyperplane. Indeed, we have

$$\begin{aligned} \mathcal{C}_1 : (x_{d+1} + r_1)^2 &= (x - p_1)^2, \quad x_{d+1} + r_1 > 0, \\ \mathcal{C}_2 : (x_{d+1} + r_2)^2 &= (x - p_2)^2, \quad x_{d+1} + r_2 > 0. \end{aligned}$$

The intersection of the two half-cones is contained in the hyperplane h_{12} whose equation is obtained by subtracting the two sides of the above equations:

$$h_{12} : -2(p_1 - p_2) \cdot x - 2(r_1 - r_2)x_{d+1} + p_1^2 - r_1^2 - p_2^2 + r_2^2 = 0.$$

This shows that there exists a correspondence between the diagram $\text{Apo}(\mathcal{S})$ and the power diagram of the hyperspheres Σ_i in \mathbb{R}^{d+1} ($i = 1, \dots, n$), where Σ_i is centered at (p_i, r_i) and has radius $r_i\sqrt{2}$. More precisely, $A(\sigma_i)$ is the projection of the intersection of the half-cone \mathcal{C}_i with the power region $L(\Sigma_i)$. Indeed, x is in $A(\sigma_i)$ if and only if the projection X_i of x onto \mathcal{C}_i has a smaller x_{d+1} -coordinate than the projections of x onto the other half-cones \mathcal{C}_j , $j \neq i$. In other words, the coordinates (x, x_{d+1}) of X_i must obey

$$\begin{aligned} (x_{d+1} + r_i)^2 &= (x - p_i)^2 \\ (x_{d+1} + r_j)^2 &\leq (x - p_j)^2 \quad \text{for any } j \neq i, \end{aligned}$$

and by subtracting both sides, it follows that $\Sigma_i(X_i) \leq \Sigma_j(X_i)$ for all j .

The Apollonius diagram of \mathcal{S} can be computed using the following algorithm:

The power diagram of the Σ_i can be computed in time $O(n^{\lfloor \frac{d}{2} \rfloor + 1} \log n)$. The intersection involved in Step 3 can be computed in time proportional to the number of faces of the power diagram of the Σ_i 's, which is $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$. We have thus proved the following theorem due to Aurenhammer [4]:

Algorithm 1 Construction of Apollonius diagrams

INPUT: a set of hyperspheres \mathcal{S}

1. Compute Σ_i , for $i = 1, \dots, n$;
2. Compute the power diagram of the Σ_i 's;
3. For all $i = 1, \dots, n$, project vertically the intersection of the power region $L(\Sigma_i)$ with the half-cone \mathcal{C}_i .

OUTPUT: the Apollonius diagram of \mathcal{S} .

Theorem 2.1.5. *The Apollonius diagram of a set of n hyperspheres in \mathbb{R}^d has complexity $O(n^{\lfloor \frac{d}{2} \rfloor + 1})$ and can be computed in time $O(n^{\lfloor \frac{d}{2} \rfloor + 1} \log n)$.*

This result is optimal in odd dimensions, since the bounds above coincide with the corresponding bounds for the Voronoi diagram of points under the Euclidean distance. It is not optimal in dimension 2, where the combinatorial complexity of the Apollonius diagram of n circles has linear size. We also conjecture that it is not optimal in any even dimension.

Computing a Single Apollonius Region

We now establish a correspondence, due to Boissonnat and Karavelas [6], between a single Apollonius region and a Möbius diagram on a hypersphere.

To give the intuition behind the result, we consider first the case where one of the hyperspheres, say σ_0 , is a hyperplane, i.e. a hypersphere of infinite radius. We take for σ_0 the hyperplane $x_d = 0$, and assume that all the other hyperspheres lie the half-space $x_d > 0$. The distance $\delta_0(x)$ from a point $x \in \mathbb{R}^d$ to σ_0 is defined as the Euclidean distance.

The points that are at equal distance from σ_0 and σ_i , $i > 0$, belong to a paraboloid of revolution with vertical axis. Consider such a paraboloid as the graph of a $(d-1)$ -variate function ϑ_i defined over \mathbb{R}^{d-1} . It follows from Sect. 2.1.1 that the minimization diagram of the ϑ_i , $i = 1, \dots, n$, is a Möbius diagram (see Fig. 2.3).

Easy computations give the associated weighted points. Write $p_i = (p'_i, p''_i)$, $p'_i \in \mathbb{R}^{d-1}$, $p''_i \in \mathbb{R}$, $i > 0$ and let $\omega = \{\omega_1, \dots, \omega_n\}$ be the set of Möbius sites of \mathbb{R}^d where $\omega_i = \{p'_i, \lambda_i, \mu_i\}$, and

$$\lambda_i = \frac{1}{r_i + p''_i}, \quad \mu_i = r_i - p''_i, \quad i > 0.$$

One finally obtains that the vertical projection of the boundary of the Apollonius region $A(\sigma_0)$ of σ_0 onto σ_0 is the Möbius diagram of ω .

We have assumed that one of the hyperspheres was a hyperplane. We now consider the case of hyperspheres of finite radii. The crucial observation is that the radial projection of $A(\sigma_0) \cap A(\sigma_i) \cap A(\sigma_j)$ onto σ_0 , if not empty, is a hypersphere. It follows that the radial projection of the boundary of $A(\sigma_0)$ onto σ_0 is a Möbius diagram on σ_0 .

Such a Möbius diagram on σ_0 can be computed by constructing the restriction of the power diagram of n hyperspheres of \mathbb{R}^d with the hypersphere σ_0 , as we have seen in the previous section.

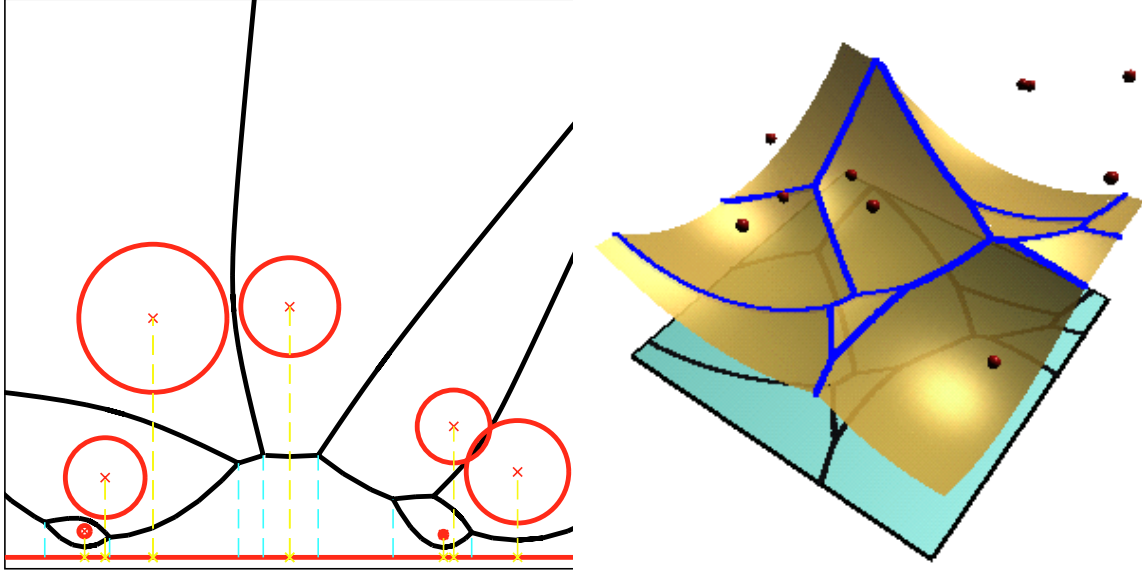


Figure 2.3: A cell in an Apollonius diagram of hyperspheres projects vertically onto a Möbius diagram in σ_0 .

Theorem 2.1.6. *Let S be a set of n hyperspheres in \mathbb{R}^d . The worst-case complexity of a single Apollonius region in the diagram of n hyperspheres of \mathbb{R}^d is $\Theta(n^{\lfloor \frac{d+1}{2} \rfloor})$. Such a region can be computed in optimal time $\Theta(n \log n + n^{\lfloor \frac{d+1}{2} \rfloor})$.*

2.2 LINEARIZATION

In this section, we introduce abstract diagrams, which are diagrams defined in terms of their bisectors. We impose suitable conditions on these bisectors so that any abstract diagram can be built as the minimization diagram of some distance functions, thus showing that the class of abstract diagrams is the same as the class of Voronoi diagrams.

Given a class of bisectors, such as affine or spherical bisectors, we then consider the inverse problem of determining a small class of distance functions that allows to build any diagram having such bisectors. We use a linearization technique to study this question.

2.2.1 Abstract Diagrams

Voronoi diagrams have been defined (see Sect. 1.1) as the minimization diagram of a finite set of continuous functions $\{\delta_1, \dots, \delta_n\}$. It is convenient to interpret each δ_i as the distance function to an abstract object o_i , $i = 1, \dots, n$. We define the bisector of two objects o_i and o_j of $\mathcal{O} = \{o_1, \dots, o_n\}$ as

$$b_{ij} = \{x \in \mathbb{R}^d, \delta_i(x) = \delta_j(x)\}.$$

The bisector b_{ij} subdivides \mathbb{R}^d into two open regions: one, b_{ij}^i , consisting of the points of \mathbb{R}^d that are closer to o_i than to o_j , and the other one, b_{ij}^j , consisting of the points of \mathbb{R}^d that are closer to o_j than to o_i . We can then define the Voronoi region of o_i as the intersection of the regions b_{ij}^i for all $j \neq i$. The union of the closures of these Voronoi

regions covers \mathbb{R}^d . Furthermore, if we assume that the bisectors are $(d-1)$ -manifolds, the Voronoi regions then have disjoint interiors and we can define the closed region associated to b_{ij}^i as $\bar{b}_{ij}^i = b_{ij}^i \cup b_{ij}$.

In a way similar to Klein [29], we now define diagrams in terms of bisectors instead of distance functions. Let $B = \{b_{ij}, i \neq j\}$ be a set of closed $(d-1)$ -manifolds without boundary. We always assume in the following that $b_{ij} = b_{ji}$ for all $i \neq j$. We assume further that, for all distinct i, j, k , the following incidence condition (I.C.) holds:

$$b_{ij} \cap b_{jk} = b_{jk} \cap b_{ki} \quad (I.C.)$$

This incidence condition is obviously needed for B to be the set of bisectors of some distance functions.

By Jordan's theorem, each element of B subdivides \mathbb{R}^d into at least two connected components and crossing a bisector b_{ij} implies moving into another connected component of $\mathbb{R}^d \setminus b_{ij}$. Hence, once a connected component of $\mathbb{R}^d \setminus b_{ij}$ is declared to belong to i , the assignments of all the other connected components of $\mathbb{R}^d \setminus b_{ij}$ to i or j are determined.

Given a set of bisectors $B = \{b_{ij}, i \neq j\}$, an *assignment* on B associates to each connected component of $\mathbb{R}^d \setminus b_{ij}$ a label i or j so that two adjacent connected components have different labels.

Once an assignment on B is defined, the elements of B are called *oriented bisectors*.

Given B , let us now consider such an assignment and study whether it may derive from some distance functions. In other words, we want to know whether there exists a set $\Delta = \{\delta_1, \dots, \delta_n\}$ of distance functions such that

1. the set of bisectors of Δ is B ;
2. for all $i \neq j$, a connected component C of $\mathbb{R}^d \setminus b_{ij}$ is labeled by i if and only if

$$\forall x \in C, \delta_i(x) \leq \delta_j(x).$$

We define the *region* of object o_i as $\cap_{j \neq i} \bar{b}_{ij}^i$.

A necessary condition for the considered assignment to derive from some distance functions is that the regions of any subdiagram cover \mathbb{R}^d . We call this condition the assignment condition (A.C.):

$$\forall I \subset \{1, \dots, n\}, \cup_{i \in I} \cap_{j \in I \setminus \{i\}} \bar{b}_{ij}^i = \mathbb{R}^d \quad (A.C.)$$

Given a set of bisectors $B = \{b_{ij}, i \neq j\}$ and an assignment satisfying I.C. and A.C., the *abstract diagram* of \mathcal{O} is the subdivision of \mathbb{R}^d consisting of the regions of the objects of \mathcal{O} and of their faces. The name *abstract Voronoi diagram* was coined by Klein [29], referring to similar objects in the plane.

For any set of distance functions δ_i , we can define the corresponding set of oriented bisectors. Obviously, I.C. and A.C. are satisfied and the abstract diagram defined by this set is exactly the minimization diagram for the distance functions δ_i . Hence any Voronoi diagram allows us to define a corresponding abstract diagram. Let us now prove the converse: any abstract diagram can be constructed as a Voronoi diagram.

We now prove that I.C. and A.C. are sufficient conditions for an abstract diagram to be the minimization diagram of some distance functions, thus proving the equivalence between abstract diagrams and Voronoi diagrams. We need the following technical lemmas.

Lemma 2.2.1. *The assignment condition implies that for any distinct i, j, k , we have*

$$b_{ij}^j \cap b_{jk}^k \cap b_{ki}^i = \emptyset.$$

Proof. A.C. implies that $\mathbb{R}^d = \cup_{1 \leq i \leq n} \cap_{j \neq i} \bar{b}_{ij}^i \subset \bar{b}_{ij}^i \cup \bar{b}_{jk}^j \cup \bar{b}_{ki}^k$. Hence, $\bar{b}_{ij}^i \cup \bar{b}_{jk}^j \cup \bar{b}_{ki}^k = \mathbb{R}^d$. Taking the complementary sets, we obtain $b_{ij}^j \cap b_{jk}^k \cap b_{ki}^i = \emptyset$. \square

Lemma 2.2.2. *For any distinct i, j, k , we have*

$$b_{ij} \cap b_{jk}^k \subset b_{ik}^k \quad \text{and} \quad b_{ij} \cap \bar{b}_{jk}^k \subset \bar{b}_{ik}^k \quad (2.1)$$

$$b_{ij} \cap b_{jk}^j \subset b_{ik}^i \quad \text{and} \quad b_{ij} \cap \bar{b}_{jk}^j \subset \bar{b}_{ik}^i \quad (2.2)$$

Proof. Let us first prove that $b_{ij} \cap b_{jk}^k \subset b_{ik}^k$:

Consider $x \in b_{ij} \cap b_{jk}^k$. Assume, for a contradiction, that $x \notin b_{ik}^k$. It follows that $x \in \bar{b}_{ik}^k$, but x cannot lie on b_{ik}^k , because this would imply that $x \in b_{ik} \cap b_{ij}$, which does not intersect b_{jk}^k . Hence, $x \in \bar{b}_{ik}^k$ and therefore, $x \in b_{ij} \cap b_{jk}^k \cap \bar{b}_{ik}^k$. We can then find x' in the neighborhood of x such that $x' \in b_{ij}^j \cap b_{jk}^k \cap b_{ki}^i$, contradicting Lemma 2.2.1.

Let us now prove that $b_{ij} \cap \bar{b}_{jk}^k \subset \bar{b}_{ik}^k$. We have proved the inclusion for $b_{ij} \cap b_{jk}^k$. It remains to prove that $b_{ij} \cap b_{jk} \subset \bar{b}_{ik}^k$ which is trivially true, by I.C.

The two other inclusions are proved in a similar way. \square

We can now prove a lemma stating a transitivity relation:

Lemma 2.2.3. *For any distinct i, j, k , we have $b_{ij}^i \cap b_{jk}^j \subset b_{ik}^i$.*

Proof. Let $x \in b_{ij}^i \cap b_{jk}^j$. Assume, for a contradiction, that $x \notin b_{ik}^i$. If $x \in b_{ik}^k$, we have $b_{ij}^i \cap b_{jk}^j \cap b_{ik}^k \neq \emptyset$, contradicting Lemma 2.2.1. Therefore, x has to belong to b_{ik} , which implies that $x \in b_{ij}^i \cap b_{ik} \subset b_{jk}^j$ by Lemma 2.2.2. This contradicts $x \in b_{jk}^j$. We deduce that $x \in b_{ik}^i$, as needed. \square

The following lemma states that at most two assignments are likely to derive from some Voronoi diagram.

Lemma 2.2.4. *For a given set B satisfying I.C. and assuming that we never have $b_{ij} \subset b_{ik}$ for $j \neq k$, there are at most two ways of labeling the connected components of each $\mathbb{R}^d \setminus b_{ij}$ as b_{ij}^i and b_{ij}^j such that A.C. is verified.*

Proof. First assume that the sides b_{12}^1 and b_{12}^2 have been assigned. Consider now the labeling of the sides of b_{1i} , for some $i > 2$: let x be a point in the non empty set $b_{2i} \setminus b_{12}$. First assume that $x \in b_{12}^1$. Lemma 2.2.2 then implies that $x \in b_{1i}^1$. Conversely, if $x \in b_{12}^2$, $x \in b_{1i}^2$. In both cases, the assignment of the sides of b_{1i} is determined.

All other assignments are determined in a similar way. One can easily see that reversing the sides of b_{12} reverses all the assignments. Thus, we have at most two possible global assignments. \square

Theorem 2.2.1. *Given a set of bisectors $B = \{b_{ij}, 1 \leq i \neq j \leq n\}$ that satisfies the incidence condition (I.C.) and an assignment that satisfies the assignment condition (A.C.), there exists a set of distance functions $\{\delta_i, 1 \leq i \leq n\}$ defining the same bisectors and assignments.*

Proof. Let δ_1 be any real continuous function over \mathbb{R}^d . Let $j > 1$ and assume the following induction property: for all $i < j$, the functions δ_i have already been constructed so that

$$\forall i, i' < j, \quad \delta_i(x) \leq \delta_{i'}(x) \Leftrightarrow x \in \bar{b}_{ii'}^i.$$

Let us build δ_j . We consider the arrangement of all bisectors b_{ij} , for $i < j$: for each $I \subset J = \{1, \dots, j-1\}$, we define $V_I = (\cap_{i \in I} \bar{b}_{ij}^i) \cap (\cap_{k \in J \setminus I} \bar{b}_{jk}^j)$. The set V_I is a non necessarily connected region of the arrangement where we need $\delta_j > \delta_i$ if $i \in I$ and $\delta_j < \delta_i$ if $i \in J \setminus I$. This leads us to the following construction.

The interior of V_I is $\text{int } V_I = (\cap_{i \in I} b_{ij}^i) \cap (\cap_{k \in J \setminus I} b_{jk}^j)$. Lemma 2.2.3 and the induction hypothesis imply that

$$\forall i \in I, \forall k \in J \setminus I, \forall x \in \text{int } V_I, \delta_i(x) < \delta_k(x).$$

In particular, if we define $\nu_I = \min_{k \in J \setminus I} \delta_k$ and $\mu_I = \max_{i \in I} \delta_i$ on V_I , we have $\mu_I < \nu_I$ on $\text{int } V_I$.

Let us now consider some point x on the boundary of V_I . We distinguish two cases. We can first assume that $x \in b_{ij}$ for some $i \in I$. Then, by Lemma 2.2.2, for any $i' \in I \setminus \{i\}$, $x \in b_{ij} \cap \bar{b}_{i'j}^{i'} \subset \bar{b}_{i'i}^{i'}$ so that $\delta_{i'}(x) \leq \delta_i(x)$. It follows that $\mu_I(x) = \delta_i(x)$.

Consider now the case when $x \in \partial V_I \cap b_{jk}$ with $k \in J \setminus I$, we have $\nu_I(x) = \delta_k(x)$. Finally, if $x \in \partial V_I \cap b_{ij} \cap b_{jk}$ with $i \in I$ and $k \in J \setminus I$, we have $\mu_I(x) = \delta_i(x)$ and $\nu_I(x) = \delta_k(x)$. By the induction hypothesis, $\delta_i(x) = \delta_k(x)$, which implies that $\mu_I(x) = \nu_I(x)$.

It follows that we can define a continuous function ρ on ∂V_I in the following way:

$$\begin{aligned} \rho_I(x) &= \mu_I(x) \text{ if } \exists i \in I, x \in b_{ij} \\ &= \nu_I(x) \text{ if } \exists k \in J \setminus I, x \in b_{jk} \end{aligned}$$

Furthermore, on $\partial V_I \cap b_{ij} = \partial V_{I \setminus \{i\}} \cap b_{ij}$, if $i \in I$, we have

$$\rho_I(x) = \mu_I(x) = \nu_{I \setminus \{i\}}(x) = \rho_{I \setminus \{i\}}(x). \quad (2.3)$$

The definitions of the ρ_I are therefore consistent, and we can now use these functions to prove that the following definition of δ_j satisfies the induction property.

Finally, we require δ_j to be any continuous function verifying

$$\mu_I < \delta_j < \nu_I$$

on each $\text{int } V_I$. By continuity of δ_j , we deduce from 2.3 that if $x \in \partial V_I \cap b_{jk} = \partial V_{I \setminus \{i\}} \cap b_{ij}$ with $k \in J \setminus I$, we have $\rho_I(x) = \mu_I(x) = \nu_{I \setminus \{i\}}(x) = \rho_{I \setminus \{i\}}(x) = \delta_j(x)$.

It follows that on each V_I , for all $i < j$, $\delta_i(x) < \delta_j(x)$ iff $x \in b_{ij}^i$ and $\delta_i(x) = \delta_j(x)$ iff $x \in b_{ij}$. The induction follows. \square

One can prove that, in the proof of Lemma 2.2.4, the assignment we build satisfies the consequences of A.C. stated in Lemmas 2.2.1, 2.2.2 and 2.2.3. The proof of Theorem 2.2.1 does not need A.C. but only the consequences of A.C. stated in those three lemmas. It follows that any of the two possible assignments determined in the proof of Lemma 2.2.4 allows the construction of distance functions, as in Theorem 2.2.1, which implies that A.C. is indeed verified. We thus obtain a stronger version of Lemma 2.2.4.

Lemma 2.2.5. *For a given set B satisfying I.C. and assuming that we never have $b_{ij} \subset b_{ik}$ for $j \neq k$, there are exactly two ways of labeling the connected components of each $\mathbb{R}^d \setminus b_{ij}$ as b_{ij}^i and b_{ij}^j such that A.C. is verified.*

Theorem 2.2.1 proves the equivalence between Voronoi diagrams and abstract diagrams by constructing a suitable set of distance functions. In the case of affine bisectors, the following result of Aurenhammer [4] allows us to choose the distance functions in a smaller class than the class of continuous functions.

Theorem 2.2.2. *Any abstract diagram of \mathbb{R}^d with affine bisectors is identical to the power diagram of some set of spheres of \mathbb{R}^d .*

Proof. In this proof, we first assume that the affine bisectors are in general position, i.e. four of them cannot have a common subspace of co-dimension 2: the general result easily follows by passing to the limit.

Let $B = \{b_{ij}, 1 \leq i \neq j \leq n\}$ be such a set. We identify \mathbb{R}^d with the hyperplane $x_{d+1} = 0$ of \mathbb{R}^{d+1} . Assume that we can find a set of hyperplanes $\{H_i, 1 \leq i \leq n\}$ of \mathbb{R}^{d+1} such that the intersection $H_i \cap H_j$ projects onto b_{ij} . Sect. 1.2 then shows that the power diagram of the set of spheres $\{\sigma_i, 1 \leq i \leq n\}$ obtained by projecting the intersection of paraboloid \mathcal{Q} with each H_i onto \mathbb{R}^d admits B as its set of bisectors¹ (see Fig. 1.5).

Let us now construct such a set of hyperplanes, before considering the question of the assignment condition.

Let H_1 and H_2 be two non-vertical hyperplanes of \mathbb{R}^d such that $H_1 \cap H_2$ projects vertically onto b_{12} . We now define the H_i for $i > 2$: let Δ_i^1 be the maximal subspace of H_1 that projects onto b_{1i} and let Δ_i^2 be the maximal subspace of H_2 that projects onto b_{2i} . Both Δ_i^1 and Δ_i^2 have dimension $d - 1$. I.C. implies that $b_{12} \cap b_{2i} \cap b_{i1}$ has co-dimension 2 in \mathbb{R}^d . Thus $\Delta_i^1 \cap \Delta_i^2$, its preimage on H_1 (or H_2) by the vertical projection, has the same dimension $d - 2$. This proves that Δ_i^1 and Δ_i^2 span a hyperplane H_i of \mathbb{R}^{d+1} . The fact that $H_i \neq H_1$ and $H_i \neq H_2$ easily follows from the general position assumption.

We still have to prove that $H_i \cap H_j$ projects onto b_{ij} for $i \neq j > 2$. I.C. ensures that the projection of $H_i \cap H_j$ contains the projection of $H_i \cap H_j \cap H_1$ and the projection of $H_i \cap H_j \cap H_2$, which are known to be $b_{ij} \cap b_{1i}$ and $b_{ij} \cap b_{2i}$, by construction. The general position assumption implies that there is only one hyperplane of \mathbb{R}^d , namely b_{ij} , containing both $b_{ij} \cap b_{1i}$ and $b_{ij} \cap b_{2i}$. This is the projection of $H_i \cap H_j$.

As we have seen, building this set of hyperplanes of \mathbb{R}^{d+1} amounts to building a family of spheres whose power diagram admits B as its set of bisectors. At the beginning of the construction, while choosing H_1 and H_2 , we may obtain any of the two possible labellings of the sides of b_{12} . Since there is no other degree of freedom, this choice determines all the assignments. Lemma 2.2.4 shows that there are at most two possible assignments satisfying A.C., which proves we can build a set of spheres satisfying any of the possible assignments. The result follows. \square

2.2.2 Inverse Problem

We now assume that each bisector is defined as the zero-set of some real-valued function over \mathbb{R}^d , called a *bisector-function* in the following. Let us denote by B the set of bisector-functions. By convention, for any bisector-function β_{ij} , we assume that

$$b_{ij}^i = \{x \in \mathbb{R}^d : \beta_{ij}(x) < 0\} \text{ and } b_{ij}^j = \{x \in \mathbb{R}^d : \beta_{ij}(x) > 0\}.$$

¹We may translate the hyperplanes vertically in order to have a non-empty intersection, or we may consider imaginary spheres with negative squared radii.

We now define an algebraic equivalent of the incidence relation in terms of pencil of functions: we say that B satisfies the linear combination condition (L.C.C.) if, for any distinct i, j, k , β_{ki} belongs to the pencil defined by β_{ij} and β_{jk} , i.e.

$$\exists(\lambda, \mu) \in \mathbb{R}^2 \quad \beta_{ki} = \lambda\beta_{ij} + \mu\beta_{jk} \quad (L.C.C.)$$

Note that L.C.C. implies I.C. and that in the case of affine bisectors L.C.C. is equivalent to I.C. Furthermore, it should be noted that, in the case of Voronoi diagrams, the bisector-functions defined as $\beta_{ij} = \delta_i - \delta_j$ obviously satisfy L.C.C.

We now prove that we can view diagrams satisfying L.C.C. as diagrams that can be linearized.

Definition 2.2.1. A diagram \mathcal{D} of n objects in some space E is said to be a pullback of a diagram \mathcal{D}' of m objects in space F by a function $\phi : E \rightarrow F$ if $m = n$ and if, for any distinct i, j , we have

$$b_{ij}^i = \phi^{-1}(c_{ij}^i)$$

where b_{ij}^i denotes the set of points closer to i than to j in \mathcal{D} and c_{ij}^i denotes the set of points closer to i than to j in \mathcal{D}' .

Theorem 2.2.3. Let $B = \{\beta_{ij}\}$ be a set of real-valued bisector-functions over \mathbb{R}^d satisfying L.C.C. and A.C. Let V be any vector space of real functions over \mathbb{R}^d that contains B and constant functions.

If N is the dimension of V , the diagram defined by B is the pullback by some continuous function of an affine diagram in dimension $N - 1$.

More explicitly, there exist a set $C = \{\psi_{ij} \cdot X + c_{ij}\}$ of oriented affine hyperplanes of \mathbb{R}^{N-1} satisfying I.C. and A.C. and a continuous function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^{N-1}$ such that for all $i \neq j$,

$$\bar{b}_{ij}^i = \{x \in \mathbb{R}^d, \beta_{ij}(x) \leq 0\} = \phi^{-1}\{y \in \mathbb{R}^{N-1}, \psi_{ij}(y) \leq c_{ij}\}.$$

Proof. Let $(\gamma_0, \dots, \gamma_{N-1})$ be a basis of V such that γ_0 is the constant function equal to 1.

Consider the evaluation application,

$$\phi : x \in \mathbb{R}^d \mapsto (\gamma_1(x), \dots, \gamma_{N-1}(x)) \in \mathbb{R}^{N-1}.$$

If point x belongs to some b_{ij}^i , we have $\beta_{ij}(x) < 0$. Furthermore, there exists real coefficients $\lambda_{ij}^0, \dots, \lambda_{ij}^{N-1}$ such that $\beta_{ij} = \sum_{k=0}^{N-1} \lambda_{ij}^k \gamma_k$. The image $\phi(x)$ of x thus belongs to the affine half-space B_{ij}^i of \mathbb{R}^{N-1} of equation

$$\sum_{k=1}^{N-1} \lambda_{ij}^k X_k < -\lambda_{ij}^0.$$

In this way, we can define all the affine half-spaces B_{ij}^i of \mathbb{R}^{N-1} for $i \neq j$: B_{ij}^i is an oriented affine hyperplane with normal vector $(\lambda_{ij}^1, \dots, \lambda_{ij}^{N-1})$ and constant term λ_{ij}^0 . Plainly, L.C.C. on the β_{ij} translates into I.C. on the B_{ij}^i , and we have

$$\bar{b}_{ij}^i = \{x \in \mathbb{R}^d, \beta_{ij}(x) \leq 0\} = \phi^{-1}\{y \in \mathbb{R}^{N-1}, B_{ij}^i(y) \leq -\lambda_{ij}^0\} \quad (2.4)$$

Finally, let us prove that A.C. is also satisfied. Lemma 2.2.5 states that the B_{ij}^i have exactly two inverse assignments satisfying A.C. Furthermore, Equation 2.4 implies that any of these two assignments defines an assignment for the b_{ij} that also

satisfies A.C. It follows that if the current assignment did not satisfy A.C., there would be more than two assignments for the b_{ij} that satisfy A.C. This proves that A.C. is also satisfied by the B_{ij} and concludes the proof. \square

We can now use Theorem 2.2.2 and specialize Theorem 2.2.3 to the specific case of diagrams whose bisectors are hyperspheres or hyperquadrics, or, more generally, to the case of diagrams whose class of bisectors spans a finite dimensional vector space.

Theorem 2.2.4. *Any abstract diagram of \mathbb{R}^d with spherical bisectors such that the corresponding degree 2 polynomials satisfy L.C.C. is a Möbius diagram.*

Proof. Since the spherical bisectors satisfy L.C.C., we can apply Theorem 2.2.3 and Theorem 2.2.2. Function ϕ of Theorem 2.2.3 is simply the lifting mapping $x \mapsto (x, x^2)$, and we know from Theorem 2.2.2 that our diagram can be obtained as a power diagram pulled-back by ϕ . That is to say $\delta_i(x) = \Sigma_i(\phi(x))$, where Σ_i is a hypersphere in \mathbb{R}^{d+1} .

Another way to state this transformation is to consider the diagram with spherical bisectors in \mathbb{R}^d as the projection by ϕ^{-1} of the restriction of the power diagram of the hyperspheres Σ_i to the paraboloid $\phi(\mathbb{R}^d) \subset \mathbb{R}^{d+1}$ of equation $x_{d+1} = x^2$.

Assume that the center of Σ_j is $(u_1^j, \dots, u_{d+1}^j)$, and that the radius of Σ_j is ρ^j . We denote by Σ_j the power to Σ_j . Distance δ_j can be expressed in terms of these parameters:

$$\delta_j(x) = \Sigma_j(\phi(x)) = \sum_{1 \leq i \leq d} (x_i - u_i^j)^2 + ((\sum_{1 \leq i \leq d} x_i^2) - u_{d+1}^j)^2 - \rho^{j^2}.$$

Subtracting from each δ_j the same term $(\sum_{1 \leq i \leq n} x_i^2)^2$ leads to a new set of distance functions that define the same minimization diagram as the δ_j . In this way, we obtain new distance functions which are exactly the ones defining Möbius diagrams.

This proves that any diagram whose bisectors are hyperspheres can be constructed as a Möbius diagram. \square

The proof of the following theorem is similar to the previous one:

Theorem 2.2.5. *Any abstract diagram of \mathbb{R}^d with quadratic bisectors such that the corresponding degree 2 polynomials satisfy L.C.C. is an anisotropic Voronoi diagram.*

Part II

Anisotropic Meshing by Linearization of Anisotropic Voronoi Diagrams

2.3 INTRODUCTION

Anisotropic meshes are triangulations of a given domain in the plane or in higher dimension, with elements elongated along prescribed directions. Anisotropic triangulations have been shown [42] to be particularly well suited for interpolation of functions or numerical modeling. They allow to minimize the number of triangles in the mesh while retaining a good accuracy in computations. For such applications, the elongation directions are usually given as quadratic forms at each point. These directions may be related to the curvature of the function to be interpolated, or to some specific directions taken into account in the equations to be solved.

Various heuristic solutions for generating anisotropic meshes have been proposed. Li et al. [35] and Shimada et al. [48] use packing methods in 2D and 3D respectively. Bossen and Heckbert [10] use a 2D method consisting in centroidal smoothing, retriangulating and inserting or removing sites. Borouchaki et al. [8] adapt the classical 2D Delaunay refinement algorithm to the case of an anisotropic metric. In terms of applications, the question of tailoring anisotropic meshes to the specific needs of partial differential equations solvers has been studied by Simpson [44]. An example of strategy used to adapt anisotropic meshes thanks to a posteriori computations of the error in finite elements computations has been presented by Apel et al. [1], and typical examples of applications to fluid dynamics computations have been investigated by Frey and Alauzet [24], adapting [8] in 3D. Recently, the practicality of such methods for numerical solving of fluid dynamics has been demonstrated further, by Dobrzynski et al. [20], who deal in particular with adaptive anisotropic meshing.

In a different context, the one of approximation of functions, Mirebeau and Cohen [36] have developed a greedy bisection method which is asymptotically optimal in terms of interpolation error.

Labelle and Shewchuk [33] have settled the foundations for a rigorous approach in the non asymptotic case, based on the so-called anisotropic Voronoi diagrams. These diagrams are computed and refined until their dual is a well-defined triangulation, with well-shaped triangles. An extension of Labelle and Shewchuk results to the 2-manifold case was proposed by Cheng et al. [13], where a 3D anisotropic Voronoi diagram is considered to build an anisotropic mesh of the closed 2-manifold embedded in 3D. This approach, however, does not solve the question of obtaining a dual anisotropic triangulation in 3D.

We present the ideas of Labelle and Shewchuk in the first two sections, and we propose an alternative view of the construction of these diagrams in Section 2.5. After detailing in Section 2.6 the computations that we need, we expose a variant of the meshing algorithm of Labelle and Shewchuk in Section 2.7. This variant computes the Voronoi vertices using a higher dimensional power diagram and refines the diagram as long as dual triangles overlap. The last sections prove the correctness of this approach.

2.4 LABELLE AND SHEWCHUK'S APPROACH

Labelle and Shewchuk [33] have proposed a discrete definition of anisotropic Voronoi diagrams. This section presents the basis of their work. The diagram is defined over a domain $\Omega \subset \mathbb{R}^d$, and each point $p \in \Omega$ has an associated metric. More specifically, a point p is given a symmetric positive definite quadratic form represented by a $d \times d$ matrix M_p . The *distance* between two points x and y as viewed by p is

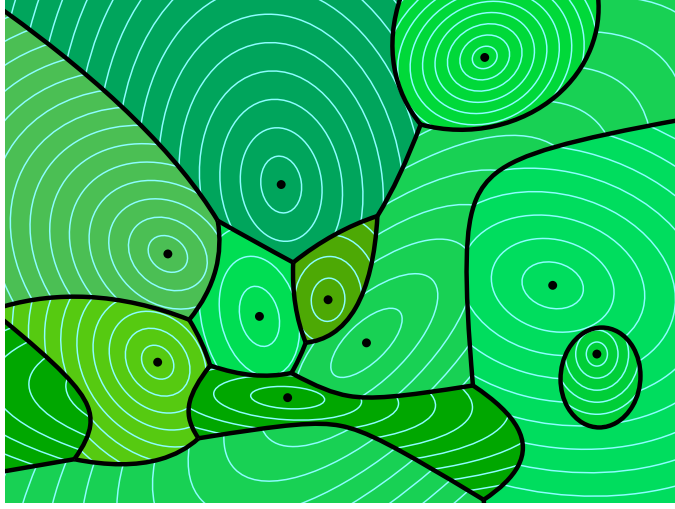


Figure 2.4: An anisotropic diagram (Courtesy of J. Shewchuk)

defined as $d_p(x, y) = \sqrt{(x - y)^t M_p (x - y)}$, and we write $d(p, q) = \min(d_p(p, q), d_q(p, q))$. Note that d_p is a distance, whereas d is not, since it does not necessarily verifies the triangular inequality.

In a similar way, $\angle_p xqy$ is defined as $\arccos \frac{(x-q)^t M_p (y-q)}{d_p(x, q) d_p(y, q)}$.

In order to compare the metric at points p and q , a transfer application is needed. Given the quadratic form M_p of a point p , we denote by F_p a matrix such that $\det(F_p) > 0$ and $F_p^t F_p = M_p$. Then $d_p(x, y) = \|F_p(x - y)\|_2$ and the transfer application from p to q is $F_{p,q} = F_q F_p^{-1}$. This application $F_{p,q}$ is in fact an isometry between the metric spaces (\mathbb{R}^d, M_p) and (\mathbb{R}^d, M_q) . The *distortion* between p and q is then defined as $\gamma(p, q) = \gamma(q, p) = \max\{\|F_{p,q}\|_2, \|F_{q,p}\|_2\}$. For any points x, y , we have $1/\gamma(p, q) d_q(x, y) \leq d_p(x, y) \leq \gamma(p, q) d_q(x, y)$.

Labelle and Shewchuk [33] define the anisotropic Voronoi diagram in the following way (and provide some examples):

Definition 2.4.1. *Let S be a set of points, called sites hereafter. The Voronoi cell of a site p in S is $\text{Vor}(p) = \{x \in \mathbb{R}^d : d_p(p, x) \leq d_q(q, x) \text{ for all } q \in S\}$. Any subset $R \subset S$ induces a Voronoi face $\text{Vor}(R) = \bigcap_{q \in R} \text{Vor}(q)$ which is the locus of points equally close to the sites in R and no closer to any other site. If not empty, such a face has dimensionality $\dim(\text{Vor}(R)) \geq d + 1 - |R|$, achieving equality if the sites are in general position. The anisotropic Voronoi diagram of S is the arrangement of the Voronoi faces $\{\text{Vor}(R) : R \subset S, \text{Vor}(R) \neq \emptyset\}$.*

It should be noted that

- each site is in the topological interior of its cell, which has dimensionality d ;
- the bisectors are quadric surfaces (conic curves in dimension 2);
- the Voronoi faces are not always connected.

See Figure 2.4 for an example of anisotropic diagram.

For brevity, we use in the sequel the term k -Vface to name Voronoi faces that have dimensionality k . The *label* of a Vface $\text{Vor}(R)$ is the set R . As noted, faces are not

necessarily connected. In particular, a 0-Vface is not necessarily a unique point, but may consist of several ones. We call each of these points a *Voronoi vertex*.

For any diagram D , and any domain Ω , we denote by D_Ω the diagram D restricted to Ω , i.e. the diagram obtained by intersecting the cells of D with Ω .

Definition 2.4.2. *The dual complex of the anisotropic Voronoi diagram of S is the simplicial complex whose set of vertices is the set S , with a simplex associated to each subset $R \subset S$ such that $\text{Vor}(R) \neq \emptyset$.*

In two dimensions and with points in general position, the dual complex includes, for each Voronoi vertex v , a dual triangle whose vertices are the three sites that compose the label of v . There is no reason why these triangles should form a triangulation. The two issues to be considered are the combinatorial planarity of the graph, which depends on the connectivity of the cells, and the ability to straighten its edges without crossing, which depends on the curvature of the bisectors.

The goal of the meshing algorithm is to refine the anisotropic Voronoi diagram by inserting new sites, so that its geometric dual becomes a triangulation, with well-shaped triangles.

In order to prove the correctness of their algorithm, Labelle and Shewchuk [33] have defined the wedge property and have proved the following results to ensure that their algorithm converges to a triangulation.

Definition 2.4.3. *The wedge between two sites p and q is the locus of points x such that the angle $\angle_p x p q$ and the angle $\angle_q x q p$ are less than $\pi/2$, or equivalently $d_p(x, q)^2 \leq d_p(p, x)^2 + d_p(p, q)^2$ and $d_q(x, p)^2 \leq d_q(q, x)^2 + d_q(p, q)^2$.*

A k -Vface f , with $k < d$, is said to be wedged if, for any pair p, q of distinct sites such that $f \subset \text{Vor}(p) \cap \text{Vor}(q)$, we have $f \subset \text{wedge}(p, q)$.

Theorem 2.4.1. *If every subface of a d -Vface $\text{Vor}(p)$ is wedged, then the d -Vface is star-shaped around p .*

The following lemma is only valid in the two-dimensional case, i.e. $d = 2$.

Lemma 2.4.1. *Let v be a Voronoi vertex labeled by the sites p, q and r . If v is wedged, then the orientation of the triangle pqr matches the ordering of the cells $\text{Vor}(p), \text{Vor}(q), \text{Vor}(r)$ locally around v .*

Let Ω be a polygonal domain of the plane and S be a set of sites in Ω that includes every vertex of Ω . We denote by D the anisotropic Voronoi diagram of S and D_Ω its restriction to Ω . The following result is central to the proof of correctness of Labelle and Shewchuk's algorithm.

Theorem 2.4.2. *Suppose that each 1-Vface of D that intersects the boundary $\partial\Omega$ intersects a single edge of $\partial\Omega$ and that each edge of $\partial\Omega$ is intersected exactly once. If all the 1-Vfaces and vertices of D_Ω are wedged, then the dual complex of D_Ω is a triangulation of Ω if S is in general position, i.e. if all Voronoi vertices have degree 3.*

If S is not in general position, the geometric dual is a polygonalization of Ω with strictly convex polygons. Labelle and Shewchuk represent the Voronoi diagram as the lower envelop of a set of paraboloids. When inserting a new site, this lower envelop is updated in a lazy way, which amounts to computing only the connected component

of the cell that contains the new site. Theorem 2.4.2 validates their lazy computation of the diagram².

Labelle and Shewchuk's algorithm consists in incrementally inserting points

- on edges of $\partial\Omega$ until these segments appear in D_Ω ;
- on non-wedged Voronoi edges;
- at the center of triangles that are badly shaped, or are too large, or do not have the same orientation as the three Voronoi cells around their dual Voronoi vertices.

2.5 RELATION TO POWER DIAGRAMS

In this section, we reduce the construction of an anisotropic Voronoi diagram in \mathbb{R}^d to the computation of a power diagram in \mathbb{R}^D where $D = d(d+3)/2$ and its restriction to a d -manifold. In the following, $\|\cdot\|$ denotes the Euclidean distance.

Definition 2.5.1. *A power diagram is defined for a set of spheres. Given a sphere σ centered at y and of radius r , the power distance of a point x with respect to σ is defined as $\pi_\sigma(x) = \|x - y\|^2 - r^2$.*

The power diagram of a set of hyperspheres Σ of \mathbb{R}^D is the subdivision induced by the power cells of the spheres in Σ , where the power cell $Pow(\sigma)$ of a sphere σ is the locus of points with a smaller power distance with respect to σ than to any other sphere in Σ : $Pow(\sigma) = \{x \in \mathbb{R}^D, \pi_\sigma(x) \leq \pi_\tau(x), \forall \tau \in \Sigma\}$.

We define the power cell of a set of spheres $\{\sigma_i\}_i$ as $Pow(\{\sigma_i\}_i) = \cap_i Pow(\sigma_i)$. The dual of the power diagram of Σ is called the regular complex of Σ .

Let $D = \frac{d(d+3)}{2}$. Associate to each point $x = (x_1, \dots, x_d) \in \mathbb{R}^d$

- the point $\tilde{x} \in \mathbb{R}^{\frac{d(d+1)}{2}}$, whose coordinates are $x_r x_s$ in increasing lexicographic ordering of (r, s) , with $1 \leq r \leq s \leq d$;
- the point $\dot{x} = (x, \tilde{x}) \in \mathcal{P} \subset \mathbb{R}^D$.

where \mathcal{P} denotes the d -manifold of \mathbb{R}^D $\{\dot{x} \in \mathbb{R}^D : x \in \mathbb{R}^d\}$.

Let $S = \{p_1, \dots, p_n\}$ be a finite set of sites in \mathbb{R}^d . To each point p_i of S , we attach a symmetric positive definite matrix M_{p_i} , whose elements are denoted by $(M_{p_i}^{r,s})_{1 \leq r, s \leq d}$, and we define

- the point $q_i = (q_i^{r,s}, 1 \leq r \leq s \leq d) \in \mathbb{R}^{\frac{d(d+1)}{2}}$ defined as
 - $q_i^{r,r} = -\frac{1}{2} M_{p_i}^{r,r}$, for $1 \leq r \leq d$;
 - $q_i^{r,s} = -M_{p_i}^{r,s}$, for $1 \leq r < s \leq d$.
- the point $\hat{p}_i = (M_{p_i} p_i, q_i) \in \mathbb{R}^D$;
- the sphere $\sigma(p_i) \subset \mathbb{R}^D$ of center \hat{p}_i and radius $\sqrt{\|\hat{p}_i\|^2 - p_i^t M_{p_i} p_i}$.

²In fact, there is a slight imprecision in their claim about the triangulation output by their algorithm: since the algorithm cannot check the wedge property for Voronoi edges that have not been computed, it does not ensure that no disconnected cell remains in the complete diagram.

Let Π be the projection $(x, \tilde{x}) \in \mathbb{R}^D \mapsto x \in \mathbb{R}^d$. Let Σ be the set of spheres $\{\sigma(p), p \in S\}$.

Lemma 2.5.1. *The anisotropic Voronoi diagram of $S \subset \mathbb{R}^d$ is the image by Π of the restriction of the D -power diagram of Σ to the d -manifold \mathcal{P} .*

Proof. We have the following equalities:

$$\begin{aligned} d_{p_i}(x, p_i)^2 &= x^t M_{p_i} x - 2p_i^t M_{p_i} x + p_i^t M_{p_i} p_i = -2q_i^t \tilde{x} - 2p_i^t M_{p_i} x + p_i^t M_{p_i} p_i \\ &= -2\hat{p}_i^t \dot{x} + p_i^t M_{p_i} p_i. \end{aligned}$$

This implies that $d_{p_i}(x, p_i) < d_{p_j}(x, p_j)$ if and only if

$$\|\dot{x} - \hat{p}_i\|^2 - (\|\hat{p}_i\|^2 - p_i^t M_{p_i} p_i) < \|\dot{x} - \hat{p}_j\|^2 - (\|\hat{p}_j\|^2 - p_j^t M_{p_j} p_j).$$

It follows that x is closer to p_i than to p_j if and only if the power of \dot{x} with respect to σ_i is smaller than its power with respect to σ_j . This proves that, for a point $z \in \mathcal{P}$, being in the power cell of σ_i is equivalent to $\Pi(z)$ being in the cell of p_i in the anisotropic diagram of S . \square

The previous lemma gives a construction of the anisotropic Voronoi diagram. As is well-known, computing a power diagram in \mathbb{R}^D reduces to computing a lower convex hull in \mathbb{R}^{D+1} . Hence, in the two-dimensional case, the computation of a six-dimensional convex hull is needed. To get the anisotropic Voronoi diagram, it remains to compute the intersection of the power diagram with the manifold \mathcal{P} . We detail the computations required by our algorithm in the following section.

2.6 BASIC OPERATIONS AND PRIMITIVES

Computing the complete anisotropic Voronoi diagram explicitly is not easy. However, our meshing algorithm only requires computing Voronoi vertices. We now explain how to compute these vertices, in the two-dimensional case. Recall that a 0-Vface of \mathbb{R}^2 may be seen as the projection of a finite subset of \mathbb{R}^5 . This set is obtained as the intersection of a linear subspace of codimension 2 (obtained as the intersection of three cells of the power diagram of Σ) with the 2-dimensional manifold \mathcal{P} (see Lemma 2.5.1).

The computation of the Voronoi vertices whose label is $\{a, b, c\}$ consists of the following steps:

- (1) Compute the power diagram of Σ and consider three sites a , b and c such that $(\sigma(a), \sigma(b), \sigma(c))$ corresponds to a triangle in the regular complex of Σ (see Definition 2.5.1), which means that their cells have a common non-empty intersection.
- (2) Compute the hyperplane H_{ab} , which is the bisector of $\sigma(a)$ and $\sigma(b)$, and the hyperplane H_{bc} , which is the bisector of $\sigma(b)$ and $\sigma(c)$, and then their intersections D_{ab} and D_{bc} with \mathcal{P} . Practically, D_{ab} and D_{bc} are represented by their projections by Π , named respectively C_{ab} and C_{bc} . The curves C_{ab} and C_{bc} are conics of \mathbb{R}^2 , and the equation of C_{ab} in \mathbb{R}^2 is:

$$(x^t M_a x - 2a^t M_a x + a^t M_a a) - (x^t M_b x - 2b^t M_b x + b^t M_b b) = 0$$

We denote this equation by $C_{ab}(x) = 0$. The equation of C_{bc} is obtained similarly.

- (3) Compute the intersection points of C_{ab} and C_{bc} . This intersection is the set of Voronoi vertices whose label is $\{a, b, c\}$ in the Voronoi diagram of $\{a, b, c\}$.
- (4) In the previous steps, we have only considered the bisectors of the spheres $\sigma(a), \sigma(b), \sigma(c)$ corresponding to the three sites involved, or, equivalently, the Voronoi diagram of $\{a, b, c\}$ alone. We now consider the bisectors of a, b, c in the Voronoi diagram of S , or, equivalently, the bisectors of the spheres $\sigma(a), \sigma(b), \sigma(c)$ in the power diagram of Σ . In the Voronoi diagram of S , some of the elements of $C_{ab} \cap C_{bc}$ are not Voronoi vertices because they belong to the cell of a closer site. Accordingly, in \mathbb{R}^D , the linear subspace $H_{ab} \cap H_{bc}$ may intersect the power cells of some other sphere $\sigma(x)$ for $x \in S \setminus \{a, b, c\}$. The pre-image by Π of a point z of $C_{ab} \cap C_{bc}$ lies on $H_{ab} \cap H_{bc}$. It belongs to the power cell $Pow(\{\sigma(a), \sigma(b), \sigma(c)\})$ if and only if its power to $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$ is smaller than its power to any other $\sigma(c')$ in Σ . We do not have to check this fact for all the other spheres $\sigma(c')$ with $c' \in S$, but only for the spheres whose cells are incident to $Pow(\{\sigma(a), \sigma(b), \sigma(c)\})$, since the cells of a power diagram are always connected. We realize this computation after projecting onto the plane.

Among the points z of $C_{ab} \cap C_{bc}$, we keep the ones such that for each tetrahedron of the regular complex defined by $\sigma(a)\sigma(b)\sigma(c)\sigma(f)$, the inequalities $C_{af}(z) < 0$, $C_{bf}(z) < 0$ and $C_{cf}(z) < 0$ are verified. Note that those three inequalities are equivalent, since z has the same power with respect to the three spheres $\sigma(a)$, $\sigma(b)$ and $\sigma(c)$. The points kept are in fact the Voronoi vertices labeled by $\{a, b, c\}$.

Our algorithm takes as input a set of segments which are required to appear in the final triangulation. These segments are called *constraint segments*. They may be refined during the algorithm, by the insertion of sites located on them. In such a case, the different pieces delimited by the sites inserted on the constraint segment are called *constraint subsegments*.

Most notably, among them are the boundaries of the domain we want to triangulate. We now present how to compute the classical property of *encroachment* of a constraint subsegment.

Definition 2.6.1. A constraint subsegment $e = (a, b)$ is encroached by a point $p \notin \{a, b\}$ if $\text{Vor}(p) \cap [a, b] \neq \emptyset$ in the Voronoi diagram of $\{a, b, p\}$.

During the algorithm, we need to compute whether a constraint subsegment $e = (a, b)$, that was not previously encroached, is encroached by a point p to be inserted.

First note that, when inserting a site p , we have a small set of potentially encroached edges: among the constraint subsegments, it is sufficient to consider the ones that would have at least one of their endpoints joined to p in the dual complex, if p were inserted in the diagram. Indeed, if p encroaches e , the cell $\text{Vor}(p)$ is adjacent to the cells of at least one of the endpoints of e : before the insertion of p , $e = [a, b]$ was not encroached and was covered by $\text{Vor}(a)$ and $\text{Vor}(b)$. After the insertion of p , $\text{Vor}(p)$ covers a part of e , while $\text{Vor}(a)$ and $\text{Vor}(b)$ cover the rest of it.

Practically, let $e = [a, b]$ be such a constraint subsegment. Then, let E be the intersection $C_{pa} \cap [a, b]$ of the bisector of p and a with $[a, b]$. If some $z \in E$ verifies $C_{pb}(z) < 0$, we have in fact $z \in \text{Vor}(\{a, p\}) \cap [a, b]$ and $\text{Vor}(p)$ intersects $[a, b]$. A constraint subsegment may also completely disappear from the dual when a site p is inserted. Such a segment is obviously encroached by p .

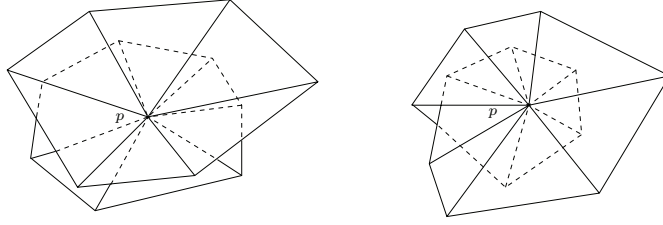


Figure 2.5: Two umbrellas (left) and one umbrella winding twice (right) around p

2.7 DESCRIPTION OF THE ALGORITHM

As above, let Ω be a polygonal domain of the plane, whose boundary is denoted by $\partial\Omega$. We denote by C the set of constraint subsegments and by S a finite set of sites in Ω . The set C is updated during the course of the algorithm to reflect the fact that some constraint segments have been refined into constraint subsegments. At the beginning, we assume that the edges of $\partial\Omega$ belong to C and that the vertices of $\partial\Omega$ belong to S . Refining the Voronoi diagram consists in adding sites to the set S . We assume that the quadratic form associated to any point of Ω can be obtained.

We have seen in the previous section how to compute the Voronoi vertices. If the label of a vertex v is $\{a, b, c\}$, the triangle abc is called the *dual triangle* of v . We now introduce some properties that will ensure that the dual triangles define a triangulation of the domain they cover.

We consider a set of non-degenerate triangles T (that is, triangles with non collinear vertices) such that

- (i) the set of vertices of the triangles in T is exactly S ;
- (ii) each edge on $\partial\Omega$ is the edge of exactly one triangle in T ;
- (iii) if e is the edge of some triangle in T and is not an edge on $\partial\Omega$, e belongs to exactly two triangles in T , which do not overlap³.

We prove that under those assumptions, T is a triangulation of Ω .

Definition 2.7.1. Let $p \in S$ be one of the sites and T_p be the set of triangles incident to p . Two triangles are said to be adjacent if they share an edge. The equivalence classes for the transitive closure of the adjacency relation in T_p are called the *umbrellas* of p .

The link $link(p)$ of a site p is the set of edges opposite to p in all the triangles of T_p .

Lemma 2.7.1. If the finite set of triangles T verifies Rules (i), (ii) and (iii), we claim that:

- (a) all the triangles in T are inside Ω ;
- (b) if p is an internal site, its umbrellas are combinatorial disks and p is inside each of its embedded umbrellas;
- (c) if p is a vertex of $\partial\Omega$, p has a unique umbrella, and p is on the boundary of this umbrella. Furthermore, the triangles of the umbrella do not overlap.

³Since all triangles are non-degenerate, the overlapping is well-defined.

Proof. For the sake of simplicity, we prove the result under the hypothesis that Ω is simply connected. The result is still true without this hypothesis. However the proof would be more complicated.

- (a) We consider an edge e of the boundary of the union U of all the triangles. From Rules (ii) and (iii), e has to be an edge of $\partial\Omega$. Thus the boundary of U is included in the boundary $\partial\Omega$. Since Ω is a simply connected polygon, $\partial\Omega$ is a topological circle embedded in the plane. The set U is closed, and so is its boundary ∂U . It follows that ∂U is a closed non empty subset of the topological circle $\partial\Omega$, which implies that $\partial\Omega = \partial U$. Finally, both U and Ω are bounded domains, with the same circle as boundary, hence $U = \Omega$.
- (b) If p is an internal site, Rule (iii) implies that $\text{link}(p)$ is a union of closed polygonal curves (not necessarily simple curves), since Rule (iii) prevents any vertex of degree different from 2 to appear on the link. An umbrella is then obtained by choosing one of those closed curves, and linking p to every vertex of it. This proves that an umbrella is a combinatorial disk, since it has a combinatorial circle as boundary.

Consider an embedded umbrella, i.e. the union U of the triangles of an umbrella. Assume for a contradiction that p is not in the interior of this union. Then p is on the boundary ∂U and both edges of this boundary that are incident to p belong to two triangles of the umbrella which have to overlap. This contradicts Rule (iii). In other words, we have proved that if there is a closed curve in the link of p , p is enclosed by it.

- (c) If p is a vertex of $\partial\Omega$, $\text{link}(p)$ may a priori contain some closed curves and some curves joining the two neighbors of p on $\partial\Omega$. As seen in the proof of (b), the closed curves have to enclose p . Thanks to (a) and to the fact that p is on $\partial\Omega$, this is not possible. Therefore, the link of vertex p cannot include a closed curve. Rule (ii) then implies that all curves in $\text{link}(p)$ have the same first and last segment and because Rule (iii) prevents any branching vertex in $\text{link}(p)$, the link $\text{link}(p)$ is a single curve. The fact that the triangles of the unique umbrella do not overlap follows from (iii) too.

□

Theorem 2.7.1. *Under assumptions (i), (ii), (iii), T is a triangulation of Ω .*

Proof. A priori, an internal site may have multiple umbrellas and each of those umbrellas may wind more than once around p . To prove that T is a triangulation, we now glue the triangles of T along their common edges and vertices to build a surface: we denote by $\mathcal{T} = \{(x, t) \in \Omega \times T \mid x \in t\}$ the set of points associated to the triangles they belong to, and we define on \mathcal{T} the equivalence relation \sim by setting $(x, t) \sim (x', t')$ if $x = x'$, $x \in \partial t$ and $x' \in \partial t'$, so that taking the quotient of the set \mathcal{T} by the equivalence relation \sim amounts to gluing the common edges and vertices. The final glued space is denoted by $\mathcal{G} = \mathcal{T} / \sim$.

Let $h : (x, t) \in \mathcal{G} \mapsto x$ be the first projection, mapping \mathcal{G} to Ω . The correctness of the triangulation is equivalent to h being a homeomorphism. Let Ω_p be the punctured space obtained by removing from Ω the vertices of the triangles of T , and let \mathcal{G}_p be $h^{-1}(\Omega_p)$.

From assumption (iii), the restriction h_p of h to \mathcal{G}_p is a local homeomorphism. Using the fact that \mathcal{G}_p is a separated space, that h_p is a proper map, and that Ω_p is

connected, it follows that h_p is a covering of Ω_p . As the points close to $\partial\Omega$ have only one pre-image, from assumption (ii), $h_p : \mathcal{G}_p \rightarrow \Omega_p$ has only one sheet and is in fact a homeomorphism.

This shows that each site p has a unique umbrella, which is well embedded and that h_p may be extended to \mathcal{G} as a homeomorphism. Thus, Ω is triangulated by T . \square

In order to present the refinement algorithm, we need to define a shape criterion. Let v be a Voronoi vertex of an anisotropic Voronoi diagram. The label of v consists of three sites that form a dual triangle $t_v = abc$. The *radius* of v is $r(v) = d_a(a, v) = d_b(b, v) = d_c(c, v)$ (we define the radius of the center instead of the radius of the triangle, because the triangle may have multiple centers). The length of an edge (a, b) is $d(a, b) = \min(d_a(a, b), d_b(a, b))$. We denote the shortest edge of t_v by $\delta(t_v)$. The *radius-edge ratio* of v is $\beta(v) = r(v)/\delta(t_v)$.

For a given *shape bound* B , a vertex v or the associated triangle are said to be *badly-shaped* if $\beta(v) > B$. Otherwise, they are said to be *well-shaped*.

Let us now present the algorithm, which refines an anisotropic Voronoi diagram V until the triangles dual to the Voronoi vertices of V_Ω , the restriction of V to Ω , have a good shape and satisfy conditions (i), (ii) and (iii) (stated in Section 2.7), and therefore form a triangulation of Ω , by Theorem 2.7.1.

First recall that, thanks to the monotonicity of the distance function associated to each point, there is always a unique point on a line segment that is equidistant from both of its endpoints.

Definition 2.7.2. Assume that a constraint subsegment $e = (p, q)$ is encroached. The breakpoint of the edge (p, q) is defined as the point of $[p, q] \setminus (\text{Vor}(p) \cup \text{Vor}(q))$ closest to the midpoint of $[p, q]$ (this point is independent of the considered metric). By midpoint, we mean the intersection of $[p, q]$ with the bisector of p and q , i.e. the point z of $[p, q]$ such that $d_p(p, z) = d_q(q, z)$.

We now present our refinement algorithm. We are given a shape bound B . At each step of the algorithm, we maintain the set T of dual triangles, obtained as the labels of the computed Voronoi vertices that are inside Ω (see Section 2.6). We define a procedure of conditional insertion, needed for the presentation of the algorithm:

CONDITIONALLY INSERT(x): if x encroaches some constraint subsegment e , insert a site at the breakpoint of e . Otherwise, insert x .

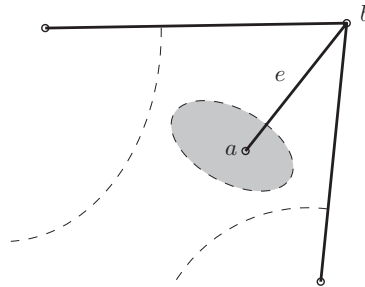


Figure 2.6: Edge e is a constraint segment, with the cell of a being completely included in the cell of b . Voronoi bisectors are represented by dashed curves.

The algorithm inserts points iteratively, applying the following rules. Rule i is applied only if no Rule j with $j < i$ applies:

- Rule (1) if some constraint subsegment $e \in C$ does not appear as the edge of a dual triangle because it is encroached, insert a site at the breakpoint of edge e ;
- Rule (2) if some constraint subsegment $e \in C$ does not appear as the edge of a dual triangle, because its dual Vface is a complete ellipse (it can happen if the constraint subsegment has a free endpoint, i.e. an endpoint which is not incident to any other constraint subsegment, see Figure 2.6 for an example), denote by Δ the support line of e . Then conditionally insert a site located at the intersection of $\Delta \setminus e$ with the ellipse;
- Rule (3) if a Voronoi vertex v is badly shaped (see Section 2.7), conditionally insert a site located at that vertex;
- Rule (4) if a triangle abc is the dual of several Voronoi vertices, conditionally insert a site located at the vertex that is the furthest from a , b and c ;
- Rule (5) if two triangles share an edge and overlap, conditionally insert a site at the dual Voronoi vertex of one of them: choose the triangle which contains the edge (x, y) such that $\gamma(x, y)$ is maximal ($\gamma(x, y)$ is the distortion between x and y defined in Section 2.4).

Note that this set of rules is designed for segment constraints. It can still be adapted for dealing with point constraints too, that is to say, isolated points which are required to appear in the final mesh: in order to prevent an isolated vertex from staying undetected, we virtually adjoin it a mirror vertex, infinitesimally close to it. This doubled vertex is a tiny constraint segment which can be dealt with symbolically, in the framework we have presented.

We will now prove that if the algorithm terminates, Conditions (i), (ii) and (iii) of Section 2.7 are verified. By Theorem 2.7.1, the dual complex is therefore a triangulation, without any badly-shaped vertex.

Lemma 2.7.2. *Upon termination of the algorithm, the dual triangles in T form a triangulation of the domain Ω and all the constraint subsegments appear in this triangulation.*

Proof. First, let us prove that each constraint subsegment is incident to at least one triangle in T . Consider some constraint subsegment s with endpoints a and b .

- Thanks to Rule (1), s is not encroached and therefore lies in the union of the cells of its endpoints.
- Since each site lies in its own cell, s cannot be included in one cell only. This proves that the dual edge $\text{Vor}(\{a, b\})$ is not empty and intersects s and the domain Ω .
- If the bisector of a and b is an ellipse, Rule (2) implies that the Voronoi edge $\text{Vor}(\{a, b\})$ has endpoints within Ω . In all cases, observe that $\text{Vor}(\{a, b\})$ is a union of curved segments, with an even number of endpoints. Furthermore, owing to the monotonicity of the distance $d_a(a, x)$ along ab , $\text{Vor}(\{a, b\})$ intersects s in at most one point (and at least once, thanks to Rule (1)). Consider the curved segment ℓ of $\text{Vor}(\{a, b\})$ which intersects s . One of the two endpoints of ℓ has to be inside Ω because $\text{Vor}(\{a, b\})$ cannot intersect any other constraint subsegment, since the other constraint subsegments are not encroached either. It follows that $\text{Vor}(\{a, b\})$ has at least one endpoint in Ω .

Therefore in any case the dual edge $\text{Vor}(\{a, b\})$ has endpoints in Ω , and the dual triangles of those endpoints are incident to s . We still have to ensure that the three hypotheses (i), (ii) and (iii) of Theorem 2.7.1 are verified. (i) is obviously verified and (iii) is implied by Rule (5). Let us now prove (ii): consider a constraint subsegment s of $\partial\Omega$. From the first part of the proof, we know that the dual Voronoi edge e of s intersect $\partial\Omega$ in one point and therefore has an odd number of endpoints within Ω . If e had more than one endpoint, i.e. if s had more than one incident triangle, it would in fact have at least three, and s would have at least three incident triangles, contradicting Rule (5). This proves that s has exactly one incident triangle, as required by hypothesis (ii). All three hypothesis are verified. In case of termination, Theorem 2.7.1 shows that the set T is a triangulation of Ω . \square

Note that Rule 2 can be omitted if we assume that the graph consisting of all constraint segments of C has no vertex of degree 1. Indeed, in such a case, if no constraint subsegment is encroached, none of them can have an ellipse as a dual Vface.

2.8 TERMINATION OF THE ALGORITHM

We now consider the conditions needed to ensure the termination of the algorithm. These conditions depend on the shape bound K and on the geometry of the initial set of constraint segments C .

Let us prove that two well-shaped dual triangles (as defined in Section 2.7) cannot overlap if the relative distortion between adjacent sites is small enough. In the following, abc and abc' are two adjacent triangles that are respectively dual to Voronoi vertices q_c and $q_{c'}$. The points q_c and $q_{c'}$ lie inside Ω , otherwise, their dual triangles would not be considered. We define γ as the maximum of the distortion $\gamma(x, y)$ (see Section 2.4) where the maximum is taken over all edges (x, y) of the two triangles, and $\delta = \max(\delta(abc), \delta(abc'))$ (as defined in Section 2.7).

If q_c and $q_{c'}$ are well-shaped, i.e. $\beta(q_c) \leq K$ and $\beta(q_{c'}) \leq K$, we have the following inequalities:

$$\begin{aligned} d_c(c, q_{c'}) &\leq d_c(c, q_c) + d_c(q_c, a) + d_c(a, q_{c'}) && \text{(triangular ineq.)} \\ &\leq d_c(c, q_c) + \gamma(a, c)d_a(q_c, a) + \gamma(a, c)d_a(a, q_{c'}) && \text{(distortion)} \\ &\leq (1 + \gamma(a, c))K\delta(abc) + \gamma(a, c)K\delta(abc') \\ &\leq (1 + 2\gamma)K\delta \end{aligned}$$

The same inequality holds when c and c' are exchanged. In the same way,

$$\begin{aligned} d_c(c, a) &\leq d_c(c, q_c) + d_c(q_c, a) \leq d_c(c, q_c) + \gamma(a, c)d_a(a, q_c) \\ &\leq (1 + \gamma)K\delta(abc) \leq (1 + \gamma)K\delta \\ &\text{and} \\ d_b(b, a) &\leq (1 + \gamma)K\delta(abc) && (*) \\ &\text{and} \\ d_c(c, c') &\leq d_c(c, q_{c'}) + d_c(q_{c'}, c') \leq (1 + 2\gamma)K\delta + \gamma(c, c')d_{c'}(c', q_{c'}) \\ &\leq (1 + \gamma)^2K\delta. \end{aligned}$$

Let $r = (1 + \gamma)^2K\delta$. We consider the zones $Z_3 = B(a, r) \cap B(b, r) \cap B(c, r)$, $Z'_3 = B(a, r) \cap B(b, r) \cap B(c', r)$ and $Z_4 = B(a, r) \cap B(b, r) \cap B(c, r) \cap B(c', r)$, where $B(p, r) =$

$\{x \in \mathbb{R}^2, d_p(p, x) \leq r\}$. As shown by the previous inequalities, the four sites a, b, c and c' are in Z_4 , as well as the two centers q_c and $q_{c'}$.

Lemma 2.8.1. *If a triangle abc is well-shaped, any point $q \notin Z_3$, is far from each of the three sites a, b and c . More precisely, for any $x \in \{a, b, c\}$, we have $d_x(x, q) > 2K\delta(abc)$.*

Proof. Assume that $q \notin B(b, r)$ for example. We then have

$$\begin{aligned} d_a(a, q) &\geq d_b(a, q)/\gamma \geq (d_b(b, q) - d_b(a, b))/\gamma \\ &\geq (r - (1 + \gamma)K\delta(abc))/\gamma && \text{(by (*))} \\ &\geq ((1 + \gamma)^2 K\delta(abc) - (1 + \gamma)K\delta(abc))/\gamma > 2K\delta(abc) \end{aligned}$$

□

Let V_{Z_4} be the restriction of the Voronoi diagram to Z_4 . We now establish a sufficient condition on the bound K and on the distortion bound γ so that the vertices and the edges of the Voronoi diagram $V_{Z_4}(\{a, b, c, c'\})$ are wedged.

Definition 2.8.1. *The three following conditions are called condition (H):*

- (i) $K > 1$ and $K^4(\gamma^2 - 1)(1 + \gamma)^6 \leq 1$
- (ii) *the triangles are well-shaped (for the bound K);*
- (iii) γ *is an upper bound on the distortion between the considered sites.*

Lemma 2.8.2. *Under condition (H), all the 0 and 1-Vfaces of the Voronoi diagram of $V_{Z_3}(\{a, b, c\})$ and all the 0 and 1-Vfaces of the Voronoi diagram of $V_{Z'_3}(\{a, b, c'\})$ are wedged.*

Proof. Let $x, y \in \{a, b, c\}$ with $x \neq y$. Let z be a point of Z_4 on the bisector of x and y . We want to ensure that $d_x(z, y)^2 \leq d_x(x, z)^2 + d_x(x, y)^2$. We have $d_x(z, y)^2 \leq \gamma^2 d_y(y, z)^2$ and since z is on the bisector between x and y , $d_x(x, z) = d_y(y, z)$. †

Now, if $[x, y]$ is the common edge of the two triangles, we have $d_x(x, y) \geq \delta$. Otherwise, we have by (*) $d_x(x, y) \geq \delta(abc) \geq d(a, b)/(K(1 + \gamma)) \geq \delta/(K(1 + \gamma))$. ‡

Finally, by inequalities † and ‡, $d_x(x, z)^2 + d_x(x, y)^2 \geq d_y(y, z)^2 + \frac{\delta^2}{K^2(1 + \gamma)^2}$. And if $\gamma^2 d_y(y, z)^2 \leq d_y(y, z)^2 + \frac{\delta^2}{K^2(1 + \gamma)^2}$, we have $d_x(z, y)^2 \leq d_x(x, z)^2 + d_x(x, y)^2$.

Thus, a sufficient condition for z to be wedged is $\gamma^2 d_y(y, z)^2 \leq d_y(y, z)^2 + \frac{\delta^2}{K^2(1 + \gamma)^2}$ (and the condition obtained by swapping x and y). The domain Z_4 was chosen so that $d_y(y, z) < r = (1 + \gamma)^2 K\delta$. Hence, a sufficient condition for the point z to be wedged is $(\gamma^2 - 1)(1 + \gamma)^2(1 + \gamma)^4 K^4 \leq 1$, i.e. $(\gamma^2 - 1)(1 + \gamma)^6 K^4 \leq 1$.

□

Lemma 2.8.3. *Under condition (H), the cells of a, b and c in $V_{Z_3}(\{a, b, c\})$ are connected.*

Proof. Under condition (H), the proof of Theorem 2.4.1 (Theorem 4 in [33]) can easily be adapted to show that every cell is connected in Z_3 , by showing that it is star-shaped around its site: let y be some point of the cell of a in $V_{Z_3}(\{a, b, c\})$. The segment $[ay]$ is entirely included in Z_3 because Z_3 is convex, as an intersection of ellipses. In order to show that y is visible from site a , we only need to consider the Voronoi edges that are intersected by the segment $[ay]$. Those intersection points lie inside Z_3 . From Lemma 2.8.2, the intersection points are wedged, and the proof of Theorem 4 [33] shows that the cell of a in $V_{Z_3}(\{a, b, c\})$ is star-shaped.

□

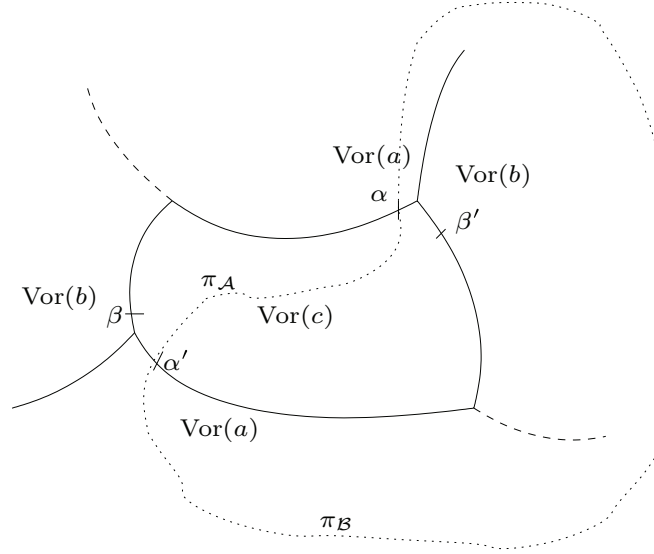


Figure 2.7: Impossible case described in Lemma 2.8.4

Lemma 2.8.4. *Consider three connected components of distinct 2-Vfaces, whose topological interiors are denoted by \mathcal{A} , \mathcal{B} and \mathcal{C} . On the boundary of \mathcal{C} , we cannot find four points $\alpha, \beta, \alpha', \beta'$ in this order such that $\alpha, \alpha' \in \partial\mathcal{A}$ and $\beta, \beta' \in \partial\mathcal{B}$ (see Figure 2.7).*

Proof. Assume to the contrary that $\alpha, \beta, \alpha', \beta'$ exist. Then, since \mathcal{C} and \mathcal{A} are connected, there exists a simple path $\pi_{\mathcal{C}}$ in \mathcal{C} and a simple path $\pi_{\mathcal{A}}$ in \mathcal{A} joining α and α' . The union of those two paths forms a closed curve π . As \mathcal{B} is connected, there is also a path $\pi_{\mathcal{B}}$ in \mathcal{B} joining β and β' . By Jordan theorem, β and β' should therefore be in the same connected component delimited by π . However, if we follow the boundary of \mathcal{C} from β to β' , we cross π exactly once. So β and β' do not belong to same connected component, which contradicts our hypothesis. \square

Lemma 2.8.5. *If (H) is verified and if all the Voronoi vertices in Z_3 are well-shaped, there is a unique Voronoi vertex with label $\{a, b, c\}$ in $V_{Z_3}(\{a, b, c\})$.*

Proof. Assume for a contradiction that two Voronoi vertices v and v' of $V_{Z_3}(\{a, b, c\})$ have the same label $\{a, b, c\}$. By Lemma 2.4.1, the cells around v and v' have the same cyclic order.

By Lemma 2.8.3, the three cells $\text{Vor}(a) \cap Z_3$, $\text{Vor}(b) \cap Z_3$ and $\text{Vor}(c) \cap Z_3$ are connected. By considering the neighborhoods of v and v' , we can find four points $\alpha, \beta, \alpha', \beta'$ in this order on the boundary of $\text{Vor}(c) \cap Z_3$ such that α, α' belong to the boundary of the cell of a and β, β' belong to the boundary of the cell of b . This contradicts Lemma 2.8.4 (see Figure 2.7). \square

Lemma 2.8.6. *If (H) is verified and if all the Voronoi vertices labeled by $\{a, b, c\}$ and $\{a, b, c'\}$ in Z_4 are well-shaped, there is exactly one Voronoi vertex labeled by $\{a, b, c\}$ and one Voronoi vertex labeled by $\{a, b, c'\}$ in $V_{Z_4}(\{a, b, c, c'\})$.*

Proof. The result follows from Lemma 2.8.5, because $Z_4 \subset Z_3$ and any vertex labeled by $\{a, b, c\}$ in Z_4 is also a vertex in $V_{Z_3}(\{a, b, c\})$. \square

The following lemma states that under low distortion of the metric, the cells are arranged along the border of Z_4 in the same order as the vertices of the convex hull of $\{a, b, c, c'\}$. This topological property will help us prove that we have a triangulation in Z_4 .

Lemma 2.8.7. *Let x be a , b or c . If (H) holds, the cell $\text{Vor}(x)$ in $V_{Z_3}(\{a, b, c\})$ contains a segment that joins x to a point on the boundary of Z_3 and does not intersect the convex hull of the sites.*

Proof. Let us assume that $x = a$ in the following. As proved in Lemma 2.8.2, under condition (H), any point in Z_3 equidistant to b and a is in the wedge defined by b and a . Therefore the cell of a in $V_{Z_3}(a, b)$ contains the intersection of Z_3 with a half-plane H_b^+ defined as follows. H_b^+ is the half-plane not containing b and bounded by the hyperplane H_b that goes through a and is normal to $[ab]$, from the point of view of a . Since a is on the boundary of the convex hull of a, b, c , the domain $H_b^+ \cap H_c^+$ contains at least one half-line r with origin a : this half-line is any half-line contained in the cone orthogonal (in the sens of the metric of a) to the cone delimited by the tangents to the convex hull at point a . This ray r does not intersect the convex hull of the three sites, and it is inside the cell of site a in the three-sites-diagram. \square

Lemma 2.8.8. *If (H) is verified and if all the Voronoi vertices labeled by $\{a, b, c\}$ and $\{a, b, c'\}$ in Z_3 and Z'_3 respectively are well-shaped, the 1-Vface of the restricted diagram $V_{Z_4}(\{a, b, c, c'\})$ labeled $\{a, b\}$ is connected.*

Proof. Let e be the dual 1-Vface of (a, b) in $V_{Z_4}(\{a, b, c, c'\})$. If e does not intersect the boundary of Z_4 or intersects it once, e has to be connected. Indeed, thanks to Lemma 2.8.6, e has at most two endpoints, labeled $\{a, b, c\}$ and $\{a, b, c'\}$, within Z_4 .

We now prove that e does not touch the boundary of Z_4 . From Lemma 2.8.3 and 2.8.5, the 1-Vface A labeled by $\{a, b\}$ in $V_{Z_3}(\{a, b, c\})$ is connected. Since a vertex labeled by $\{a, b, c'\}$ exists in $V_{Z_4}(\{a, b, c, c'\})$, it has to belong to A . Consider the arc $\ell \subset A$ of the bisector of $\{a, b\}$ which links the vertex q_c labeled by $\{a, b, c\}$ and the vertex $q_{c'}$ labeled $\{a, b, c'\}$. Let us prove that ℓ is entirely included in Z_4 . Assume to the contrary that ℓ is not entirely included in Z_4 . The boundary of Z'_3 has to intersect it twice, because $Z_4 = Z_3 \cap Z'_3$. It follows that the 1-Vface A' labeled by $\{a, b\}$ in $V_{Z_3}(\{a, b, c'\})$, which contains $\ell \cap Z_4$, intersects the boundary of Z'_3 twice. Since there is only one vertex labeled $\{a, b, c'\}$ in Z'_3 , there is a sub-arc ℓ' of A' without any vertex on it. ℓ' cuts Z_3 into two parts (called the two sides of ℓ' in the following). The cell of c' is connected, and is on one side of ℓ' . The other side of ℓ' belongs to another cell. Let us assume, without loss of generality, that it belongs to the cell of a . This part of the cell of a is not adjacent to the cell of c' , which implies that the cell of a is disconnected. This contradicts the fact that the cells of $V_{Z_3}(\{a, b, c'\})$ are connected. We have proved that ℓ is entirely included in Z_4 . Then, since there is only one vertex labeled by $\{a, b, c'\}$ in Z_4 , ℓ is exactly the 1-VFace labeled by $\{a, b\}$ in $V_{Z_4}(\{a, b, c, c'\})$. This concludes the proof. \square

Lemma 2.8.9. *If (H) is verified and if all the Voronoi vertices in Z_4 are well-shaped, the two triangles abc and abc' do not overlap each other.*

Proof. From Lemma 2.8.8, the 1-Vface labeled $\{a, b\}$ in the restricted diagram $V_{Z_4}(\{a, b, c, c'\})$ is connected. The two endpoints of the 1-Vface labeled $\{a, b\}$ are the Voronoi vertices q_c and $q_{c'}$. It follows that the cells of a , b and c around q_c and the

cells of a , b and c' around $q_{c'}$ have opposite cyclic orders. Lemma 2.4.1 applied to $V_{Z_3}(\{a, b, c\})$ and $V_{Z'_3}(\{a, b, c'\})$ then implies that the triangles abc and abc' do not overlap each other. \square

We now consider the algorithm at some point during its execution. The proof makes use of an arbitrary shape bound K and a distortion coefficient G , chosen so that the following condition (C) is satisfied: any pair of adjacent segments of C forms an angle of at least $2 \arcsin(G^2/2)$ and $(G^2 - 1)(1 + G)^6 K^4 \leq 1$. (C) This section aims at proving a lower bound on the insertion radius d_{\min}^w of the next inserted site w . By insertion radius, we mean the shortest Euclidean distance between the new site and all the previously inserted sites. It may depend on the current shortest anisotropic distance d_{\min} between the sites, on the shape bound K , on the geometry of the constraint segments and on the metric field on Ω . The distortion coefficient G is used as a way to discriminate different configurations inside the proof. As we have seen, no such coefficient intervenes in the algorithm itself. The following definitions are taken from [33]:

Definition 2.8.2. *The bounded distortion radius $\text{bdr}(p, \gamma)$ is defined as $\sup\{\ell : d_p(p, q) \leq \ell \Rightarrow \gamma(p, q) \leq \gamma\}$ and $\text{bdr}_{\min}(\gamma)$ is the lower bound of these radii: $\text{bdr}_{\min}(\gamma) = \inf\{\text{bdr}(p, \gamma) : p \in \Omega\}$.*

Definition 2.8.3. *Given some bound $G > 0$, two points q and q' that belong to constraint segments in C are said to be G -intertwined if they lie on a common segment of C or if they lie on two edges e and e' of C that share an endpoint b and are such that $\gamma(q, b) < G$ and $\gamma(q', b) < G$. For a set of constraint segments C , the local feature size $\text{lfs}_{\min}^G(C)$ is the upper bound on the distances r such that $x < r$ implies that for all $p \in \Omega$, $B(p, x)$ does not contain two non- G -intertwined points of $\cup C$.*

The following four lemmas are Lemma 5, 14, 16 and 17 of [33]:

Lemma 2.8.10. *Let w be a point on the bisector of a and b that lies outside the wedge of a and b , on the side of b . Let $G \geq 1$ be a constant for which $\gamma(a, b) \leq G$. Then the proximity of w to a and b is bounded by $d_a(a, w) = d_b(b, w) \geq d_b(b, a)/\sqrt{G^2 - 1}$.*

Lemma 2.8.11. *Let a and b be two sites of a Voronoi diagram D , and w a point on the bisector $\text{Vor}(\{a, b\})$ in D . Assume that there exists some $G > 1$ such that $d_a(a, b) \geq \text{bdr}(a, G)$. Then for any site x of D , $d(x, w) \geq \text{bdr}_{\min}(G)/(G^3 + G)$.*

Lemma 2.8.12. *Let p be a point in Ω . For any $G > 1$ and for every site x , $d_p(p, x) \geq \min\left(\frac{d_x(x, p)}{G}, \text{bdr}(p, G)\right)$.*

Lemma 2.8.13. *Assume that any pair of adjacent segments of C forms an angle of at least $2 \arcsin(G^2/2)$, as measured by the common endpoint. Let $e = (a, b)$ be a subsegment of C . Let s be a site that encroaches e . Let w be a point in $\text{Vor}(s) \cap e$. Let $m = \min\{d_a(a, s), d_b(b, s)\}$. Then for any site x of the diagram, $d(x, w) \geq \min(m, \text{lfs}_{\min}^G(C)/G, \text{bdr}_{\min}(G))$.*

We study now the inter-site distances created while inserting a new site w along the five rules of the algorithm, as presented in Section 2.7. Recall that G is assumed to satisfy Condition (C).

Rule 1: If Rule 1 applies, the inter-site distances created by the insertion of the breakpoint of the encroached subsegment are bounded by Lemma 2.8.13: for any site x of the diagram, $d(x, w) \geq \min(m, \text{lfs}_{\min}^G(C)/G, \text{bdr}_{\min}(G))$.

We call *original refinement point* the point passed as argument to the conditional insertion procedure. We now consider the cases of Rules 2, 3, 4 and 5 when the inserted point w is the original refinement point and not a point lying on an encroached edge.

Rule 2: If Rule 2 applies, the inserted site w lies on an edge $\text{Vor}(a) \cap \text{Vor}(b)$ but outside $\text{wedge}(a, b)$. We have two cases to consider. If $\gamma(a, b) \leq G$, we can apply Lemma 2.8.10 so that for every site x , $d_x(x, w) \geq d_a(a, w) = d_b(b, w) \geq \frac{d_{\min}}{\sqrt{G^2-1}}$, and Lemma 2.8.12 then implies $d(x, w) \geq \min\left(\frac{d_{\min}}{G\sqrt{G^2-1}}, \text{bdr}_{\min}(G)\right)$. If $\gamma(a, b) > G$, Lemma 2.8.11 implies $d(x, w) \geq \frac{\text{bdr}_{\min}(G)}{G^3+G}$.

Rule 3: If Rule 3 applies, w is located at a Voronoi vertex dual to the triangle abc and at distance $d_x(x, w) \geq r = d_a(a, w) > K\delta(abc)$ from any site x . Lemma 2.8.12 implies that for every site x and any coefficient G , $d(x, w) \geq \min\left(\frac{K}{G}d_{\min}, \text{bdr}_{\min}(G)\right)$.

Rule 4: If Rule 4 applies, no vertex is badly shaped, and w is one of the vertices dual to triangle abc . Because w is located at the furthest vertex from a, b and c , Lemma 2.8.5, implies that either the distortion between the sites a, b and c is greater than G , or w does not belong to the zone Z_3 . If the distortion is greater than G , we can use Lemma 2.8.11. If w is not in Z_3 , thanks to Lemma 2.8.1, for every site x , $d_x(x, w) \geq K\delta(abc) \geq Kd_{\min}$, so that, using Lemma 2.8.12, $d(x, w) \geq \min\left(\frac{K}{G}\delta(abc), \text{bdr}_{\min}(G)\right)$. In summary, if w is inserted by Rule 4, $d(x, w) \geq \min\left(\frac{K}{G}d_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right)$.

Rule 5: Finally, if Rule 5 applies, w is located at the Voronoi vertex of a triangle abc overlapping another triangle abc' . Rule 4 implies that abc has a unique dual vertex. Lemma 2.8.9 proves that this is only possible if γ , the maximum of the distortion $\gamma(x, y)$ where the maximum is taken over all edges (x, y) of the two triangles abc and abc' is greater than G , since both abc and abc' are well shaped. Then we have the bound given by Lemma 2.8.11 for every site x : $d(x, w) \geq \frac{\text{bdr}_{\min}(G)}{G^3+G}$.

Summary for Rules 2,3,4,5 without encroachment:

We have proved that, if the original refinement point is inserted, the minimal distance d_{\min}^w after insertion of w verifies $d_{\min}^w \geq \min\left(\frac{d_{\min}}{G\sqrt{G^2-1}}, \frac{K}{G}d_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right)$ where K is the shape bound and G is any value satisfying Condition (C), as stated at the beginning of Section 2.8.

Rules 2,3,4,5 with encroachment: Denote by $e = (a, b)$ the constraint sub-segment encroached by s . Since s encroaches e , we insert the corresponding break-point w on e . First recall the following fact, extracted from the proof of Lemma 23 in [33]: if w belongs to $\text{Vor}(a)$ and if for some $G > 1$, we have $d_s(s, a) \geq \text{bdr}(s, G)$, then for any site x , $d(x, w) \geq \text{bdr}_{\min}(G)/(G^3 + G)$. Otherwise, we have $d_x(x, w) \geq d_a(a, s)/(G\sqrt{G^2 + 1})$. We use now the bounds established for Rules 2, 3 and 4 (with w replaced by s):

$$d_x(x, s) \geq \min\left(\frac{d_{\min}}{\sqrt{G^2-1}}, Kd_{\min}, \frac{\text{bdr}_{\min}(G)}{G^3+G}\right), \text{ and } d_x(x, w) \geq \min\left(\frac{d_{\min}}{G\sqrt{G^4-1}}, \frac{Kd_{\min}}{G\sqrt{G^2+1}}, \frac{\text{bdr}_{\min}(G)}{(G^4+G^2)\sqrt{G^2+1}}\right). \text{ Lemma 2.8.12 then implies that } d(x, w) \geq \min\left(\frac{d_{\min}}{G^2\sqrt{G^4-1}}, \frac{Kd_{\min}}{G^2\sqrt{G^2+1}}, \frac{\text{bdr}_{\min}(G)}{(G^5+G^3)\sqrt{G^2+1}}\right).$$

Termination

In order to handle the first two terms in the previous equation and to respect the condition of Lemma 2.8.2, let assume that $K > 1$ and $G > 1$ satisfy (C) and the two additional conditions $G^2\sqrt{G^4-1} \leq 1$ and $G^2\sqrt{G^2+1} \leq K$.

Note that for any $K > \sqrt{2}$, a suitable $G > 1$ may be found, since all conditions are verified when $G \rightarrow 1^+$. We also demand that any pair of incident edges of C

forms an angle of at least $2 \arcsin(G^2/2)$, so that it complies to the requirements of Lemma 2.8.13. Under those conditions, the minimal inter-distance d'_{\min} after the insertion of a new site is bounded from below: $d_{\min} \geq \min \left(\frac{\text{bdr}_{\min}(G)}{(G^5+G^3)\sqrt{G^2+1}}, \frac{\text{lfs}_{\min}^G(C)}{G} \right)$. Finally, if we can find G satisfying the conditions and such that $\text{bdr}_{\min}(G) > 0$, the above bound is not trivial, and an easy induction shows that we indeed have a lower bound on the minimal inter-distance. This proves that the algorithm will not insert sites indefinitely, by a classical volume argument. Moreover, because $(G^2 - 1)K^2 < 1$, the shape condition parametrized by K may be translated into a condition in terms of a lower bound on the angles of the triangles, as measured by any point inside the triangle (see Corollary 10 of [33]).

Theorem 2.8.1. *Let $K > \sqrt{2}$ be a constant, and let C be a set of constraint segments which bounds a polygonal domain of the plane such that incident segments always form an angle greater than 60° . Under these assumptions, the algorithm presented in Section 2.7 terminates and provides a triangulation whose dual Voronoi vertices respect the shape bound K .*

Proof. Let $G > 1$ be such that $(G^2 - 1)(1 + G)^6 K^4 \leq 1$ and $G^2 \sqrt{G^4 - 1} \leq 1$ and $G^2 \sqrt{G^2 + 1} \leq K$. We can also assume that G is close enough to 1, so that incident segments of C always form an angle greater than $2 \arcsin(G^2/2)$. We have seen that such a G can always be found. And if $\text{bdr}_{\min}(G) > 0$, which is always the case if the field of metrics is continuous, the algorithm terminates. \square

2.9 CONCLUSION AND FUTURE WORK

The approach that we have presented is built upon the work of Labelle and Shewchuk. Instead of using a lower envelop of paraboloids, computed in a greedy way, we rely on a power diagram in higher dimension. As we have shown, we do not need all the combinatorial informations given by such a diagram, but only the zero-dimensional intersections of it with a 2-manifold. Indeed, we present the algorithm by focusing on the overlapping condition on dual triangles, thus minimizing the dependence over the Voronoi diagram itself, apart from the computation of the Voronoi vertices. As an aside, we also rely only on the Voronoi vertices that are inside the domain Ω , while Labelle and Shewchuk compute the whole diagram.

The simplicity of the structure of our algorithm makes it a good candidate for an extension to the 3-dimensional case, especially because of the absence of topological considerations. However, we currently cannot prove that this meshing algorithm terminates in three dimensions because flat tetrahedra may overlap their neighbors, without inducing a large insertion distance for the new refining point. This may happen even in the case of low distortion of the metric field. The extension to the 3-dimensional case, while relying on a simple framework, raises interesting issues in terms of complexity of the computation of the restriction of a high dimensional power diagram, and in terms of termination conditions and proper embedding of a three-dimensional triangulation.

Part III

Locally Uniform Anisotropic Meshes

LOCAL UNIFORMITY

OVERVIEW

In the previous part, we have presented an approach which is mostly theoretical, in the sense that it is computationally expensive, and cannot yet be easily applied to the 3D case, for reasons related both to the computational cost and the geometrical problem encountered because of sliver tetrahedra.

The following chapter takes a very different view at the same problem. Changing the definition of the object to be built, we are then able to provide a simple algorithm that works in both 2D and 3D. The considered object is first an unorganized soup of stars which is refined until they are consistent. It is not anymore defined as the dual of a Voronoi diagram. However, once the algorithm terminates, the same kind of shape guarantees will apply.

Contents

3.1 Introduction	72
3.2 Preliminaries	72
3.2.1 Anisotropic Metric	72
3.2.2 Distortion	73
3.3 Stars and Refinement	74
3.3.1 Stars	74
3.3.2 Quasi-Cosphericity	74
3.3.3 Picking Region	76
3.3.4 Encroachment and Star Initialization	79
3.4 Algorithm	80
3.4.1 Algorithm Outline	80
3.4.2 Termination of the Algorithm and Quality of the Mesh	81
3.5 Conclusion	84

3.1 INTRODUCTION

Neither the method of Labelle and Shewchuk [33] nor the method presented in the previous chapter can be extended to 3D in a straightforward manner because the occurrence of sliver tetrahedra results in endless loops.

We propose a new approach for the generation of anisotropic meshes. Given a set of sites V , for each site $v \in V$, computing the Delaunay triangulation $\text{Del}_v(V)$ for the metric M_v is simple, since it is just the image of a Euclidean Delaunay triangulation under a stretching transformation. We define the star S_v of a site v as the set of simplices incident to v in $\text{Del}_v(V)$. With this notation, we can define a *locally uniform anisotropic mesh* as a mesh such that for each site v , the set of elements incident to v in the mesh is exactly its star S_v . Our algorithm allows to build such a *locally uniform anisotropic mesh*.

Initially, there are *inconsistencies* among the stars of the sites, in the sense that it is impossible to merge these stars into a mesh. Then, by adding new points in V at carefully chosen locations, we show how to remove all the inconsistencies. The data structure involved is similar to the one presented by Shewchuk[43], in the context of maintaining triangulations of moving points. Furthermore, the method for guaranteeing termination is inspired by the method for sliver removal of Chew[14], also used later by Li and Teng[34].

Some notable advantages of this new method are:

- programming this algorithm is simple and straightforward, since it relies on the usual Delaunay predicates (applied to some stretched spaces);
- it is valid in 3D;
- in 3D, the termination of the algorithm relies on the sliver removal method of Chew[14], adapted to avoid configurations unsuitable for the algorithm. Consequently, slivers tetrahedra, which are a typical problem for numerical computations, are also avoided without further expense.

We present the results in dimension 3. However most of these results are still true in higher dimensions (and dimension 2), with few or without modifications.

3.2 PRELIMINARIES

3.2.1 Anisotropic Metric

We consider a domain $\Omega \subset \mathbb{R}^d$ and assume that each point $p \in \Omega$ is given a symmetric positive definite quadratic form represented by a $d \times d$ matrix M_p , called the metric at p . The distance between two points a and b , as measured by a metric M is defined as

$$d_M(a, b) = \sqrt{(a - b)^t M (a - b)}$$

and we use the notations $d_p = d_{M_p}$, $d_p(a) = d_p(p, a)$ and $d(a, b) = \min(d_a(b), d_b(a))$.

Given the positive definite quadratic form M_p of a point p , we denote by F_p any matrix such that $\det(F_p) > 0$ and $F_p^t F_p = M_p$. The Cholesky decomposition provides such a square root matrix F_p . Note however that F_p is not unique. The Cholesky decomposition provides an upper triangular F_p , while a symmetric F_p can be obtained

by diagonalizing the quadratic form M_p and computing the quadratic form with the same eigenvectors and the square root of each eigenvalue.

The Delaunay triangulation $\text{Del}_p(V)$ of a finite set of points V with metric M_p is simply obtained by computing the Euclidean Delaunay triangulation of the stretched image $F_p(V)$, and stretching the triangulation back with F_p^{-1} . In the sequel, the points of V , associated with their metrics, are called *sites*, and we refer to the elements of maximal dimension in the triangulation (tetrahedra in 3D) as *simplices*.

Definition 3.2.1. *Given some metric M , a sphere or a ball computed for M are called M -sphere and M -ball. In the same way, we define the M -circumsphere $\mathcal{C}_M(\tau)$, the M -circumball $\mathcal{B}_M(\tau)$ and the M -circumradius $R_M(\tau)$ of a simplex τ , and the M -volume of a domain.*

Given some metric M , the M -radius-edge ratio $\rho_M(\tau)$ of a simplex τ is the ratio $R_M(\tau)/d_M(\tau)$, where $d_M(\tau)$ denotes its shortest edge, as measured by M .

Note that if M and N are two metrics, an M -sphere is in general an ellipsoid for N . In particular, an M -sphere is an empty Euclidean ellipsoid, with axes aligned along the eigenvectors of M_p . $\text{Del}_p(V)$ is the triangulation of V such that each simplex has an empty M_p -circumsphere. By empty, we mean that the circumsphere contains no site of the triangulation.

3.2.2 Distortion

The definitions in this section are mostly the ones proposed by Labelle and Shewchuk[33] that we presented in Section 2.4. Still, we recall them because we slightly modify them so that they suit our context.

Given two metrics M and N , and their square-roots F_M and F_N , the relative *distortion* between M and N is then defined as $\gamma(M, N) = \max\{\|F_M F_N^{-1}\|_2, \|F_N F_M^{-1}\|_2\}$, where $\|\cdot\|_2$ denotes the operator norm associated to the Euclidean metric. Similarly, given two points p and q , the relative *distortion* between p and q is then defined as $\gamma(p, q) = \gamma(M_p, M_q)$.

A fundamental property of $\gamma(p, q)$ is that it bounds the difference between d_p and d_q : for any points x, y , we have $1/\gamma(p, q) d_q(x, y) \leq d_p(x, y) \leq \gamma(p, q) d_q(x, y)$. The *bounded distortion radius* $\text{bdr}(p, \gamma)$ is the upper bound of numbers ℓ such that for all q and r in Ω , $\max(d_p(q), d_p(r)) \leq \ell \Rightarrow \gamma(q, r) \leq \gamma$. Furthermore, the *minimal bounded distortion radius* associated to γ is $\text{bdr}_{\min}(\gamma) = \inf \text{bdr}(p, \gamma)$, with the minimum taken over all points $p \in \Omega$. Note that this is not exactly the same definition as the one proposed by Labelle and Shewchuk (denoted bdr_{LS} here), but we have

Lemma 3.2.1. *The two notions of bounded distortion radius are related by the following inequalities: $\text{bdr}_{LS}(p, \sqrt{\gamma}) < \text{bdr}(p, \gamma) < \text{bdr}_{LS}(p, \gamma)$.*

Proof. $\text{bdr}_{LS}(p, \gamma)$ is the upper bound of numbers ℓ such that for all q in Ω , $d_p(q) \leq \ell \Rightarrow \gamma(p, q) \leq \gamma$. In particular, if $\ell < \text{bdr}_{LS}(p, \sqrt{\gamma})$, we have that $\max(d_p(q), d_p(r)) \leq \ell$ implies $\gamma(q, r) \leq \gamma(q, p)\gamma(p, r) \leq \gamma^2$. The lower bound follows: $\text{bdr}_{LS}(p, \sqrt{\gamma}) < \text{bdr}(p, \gamma)$. The other inequality is a direct consequence of the definition. \square

In dimension 3, each simplex $\tau = abcd$ has four circumspheres $\mathcal{C}_a(\tau)$, $\mathcal{C}_b(\tau)$, $\mathcal{C}_c(\tau)$ and $\mathcal{C}_d(\tau)$. We define the *total distortion* over τ as the maximal distortion between any pairs of points of Ω which are both inside $\mathcal{C}_a(\tau)$ or both inside $\mathcal{C}_b(\tau)$, or $\mathcal{C}_c(\tau)$ or $\mathcal{C}_d(\tau)$. This total distortion is denoted by $\gamma(\tau)$.

In the following, we assume that the domain Ω to be meshed is compact, and that the metric field is continuous over Ω . It follows that $\Gamma = \max_{x,y \in \Omega} \gamma(x,y)$ is finite.

3.3 STARS AND REFINEMENT

We now define the local structures that are built and refined by our algorithm. These definitions rely on the notion of restricted Delaunay triangulation.

Let Ω be a domain of \mathbb{R}^3 , and let V be a finite set of points of Ω .

Definition 3.3.1. *The restriction to Ω of the Delaunay triangulation $\text{Del}(V)$ of V is the sub-complex of $\text{Del}(V)$ consisting of the simplices whose Voronoi dual belongs to Ω .*

3.3.1 Stars

Definition 3.3.2. *We define the star S_v of a site v as the set of simplices incident to v in $\text{Del}_v(V)$ restricted to Ω .*

Definition 3.3.3. *Two stars S_v and S_w are said to be inconsistent if edge $[vw]$ appears in only one of the two stars S_v and S_w . Any simplex containing $[vw]$ is also said to be inconsistent (see Figure 3.1).*

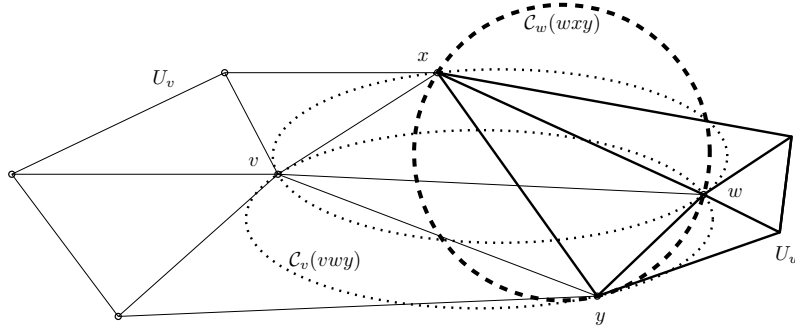


Figure 3.1: Example of inconsistent stars in 2D: stars S_v and S_w are inconsistent because edge $[vw]$ belongs to S_v but not to S_w .

Definition 3.3.4. *The conflict zone of a star S_v is the union of the balls $B_{M_v}(\tau)$ circumscribing the simplices τ that compose S_v . We denote it by Z_v .*

The following result is a simple property of the Delaunay triangulation:

Lemma 3.3.1. *The conflict zone of a star S_v is non-increasing upon insertion of new sites.*

It follows that the star of a site v can be maintained by maintaining a local triangulation around v : to each site v is attached a triangulation \mathcal{T}_v , computed as the Delaunay triangulation for metric M_v , and a new site s is inserted into \mathcal{T}_v only if s belongs to the conflict zone of S_v .

3.3.2 Quasi-Cosphericity

Let $\gamma_0 > 1$ be a bound on the distortion. We introduce now the notion of γ_0 -cosphericity and show its link with inconsistent simplices.

Definition 3.3.5. Five sites a, b, c, d, e are said to be γ_0 -cospherical for metric M if there exist two metrics N, N' such that

- $\gamma(M, N) \leq \gamma_0, \gamma(M, N') \leq \gamma_0, \gamma(N, N') \leq \gamma_0$;
- the triangulations $\text{Del}_N(\{a, b, c, d, e\})$ and $\text{Del}_{N'}(\{a, b, c, d, e\})$ are different.

If γ_0 is implicit, we say that a, b, c, d, e are quasi-cospherical.

See Figure 3.2 for an illustration in 2D. Note that the five points a, b, c, d, e play symmetric roles in the definition of γ_0 -cosphericity. We have the following simple fact:

Lemma 3.3.2. Five points a, b, c, d, e are γ_0 -cospherical for metric M if there exist two metrics N, N' such that

- $\gamma(M, N) \leq \gamma_0, \gamma(M, N') \leq \gamma_0, \gamma(N, N') \leq \gamma_0$;
- e is outside $\mathcal{C}_N(abc)$;
- e is inside $\mathcal{C}_{N'}(abc)$.

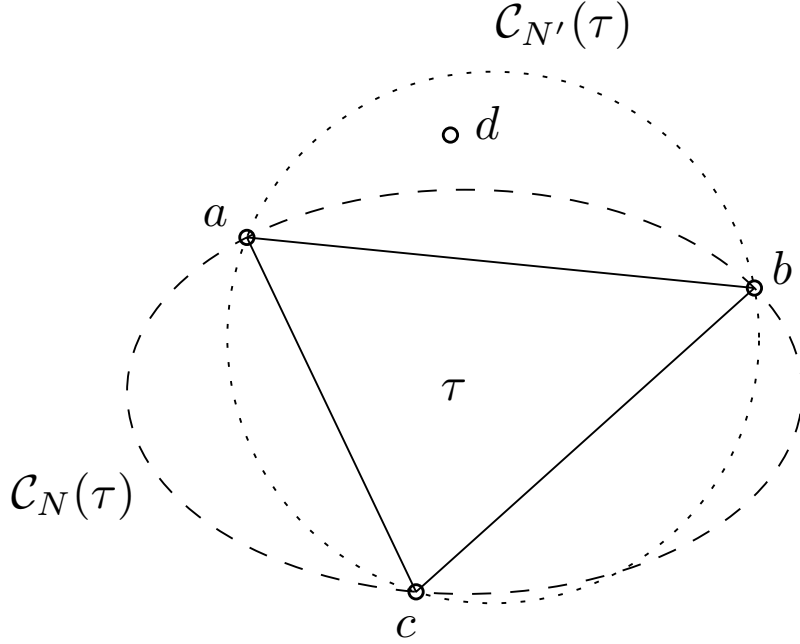


Figure 3.2: Example of quasi-cospherical points in 2D: a, b, c and d are quasi-cospherical because d is outside of $\mathcal{C}_N(abc)$ but inside $\mathcal{C}_{N'}(abc)$

Let us now show how the notion of γ_0 -cosphericity is related to inconsistencies:

Lemma 3.3.3. Let $\tau = (v, w, x, y)$ be some inconsistent simplex with distortion $\gamma(\tau) < \gamma_0$, which appears in star S_v but not in star S_w . Then there exists a vertex p of S_w such that $\{v, w, x, y, p\}$ are γ_0 -cospherical for metric M_v .

Proof. Since $\tau \in S_v$, $\mathcal{C}_v(vwxy)$ is empty. But since $\tau \notin S_w$, there exists some site p of S_w which is inside $\mathcal{C}_w(vwxy)$. It follows that v, w, x, y, p are γ_0 -cospherical for metric M_v . \square

Definition 3.3.6. Given some metric M and five points x_1, \dots, x_5 γ_0 -cospherical for the metric M , the M -radius r of the quasi-cospherical configuration is the minimum of the M -circumradii of the simplices $x_i x_j x_k x_l$, for i, j, k, l distinct integers in $\{1, \dots, 5\}$.

The M -radius-edge ratio of the quasi-cospherical configuration is the ratio r/d_{\min} , where $d_{\min} = \min_{1 \leq i \neq j \leq 5} d(x_i, x_j)$.

3.3.3 Picking Region

The refinement algorithm consists of refining the simplices which do not satisfy the required conditions in terms of size, shape, distortion radius or consistency by inserting a point in the empty circumscribing ball of each bad simplex (the circumscribing ball being computed for the metric of the star currently considered). In the usual Delaunay refinement, this point is simply the circumcenter of the simplex.

However, we cannot guarantee that the consistency problems will disappear if new sites are inserted exactly at the circumcenter of the simplices. As we have seen in the previous section, once the distortion radii of all elements are small, remaining inconsistencies are related to the occurrence of quasi-cospherical configurations. At this point, if the exact circumcenter is inserted, cascading configurations are possible: the refinement could create smaller and smaller inconsistent quasi-cospherical simplices. This is easily seen from the fact that the classical Delaunay refinement cannot get rid of almost flat and cocyclic tetrahedra, called *slivers*. We quantify this by measuring the shortest distance between sites:

Definition 3.3.7. The shortest interdistance $\ell(V)$ of the set of sites V is the shortest distance between pairs of sites of V :

$$\ell(V) = \min_{a, b \in V} d(a, b)$$

In order to prove the termination of the refinement procedure, we need to provide a positive lower bound on $\ell(V)$. In the same way as Chew[14] and Li and Teng[34] did for avoiding slivers in 3D Delaunay refinement, we define for each simplex, face and edge (generically called *face* in the sequel) a picking region. Let $\delta < 1$ be a constant to be specified later. If c_τ and r_τ are the M -circumcenter and M -circumradius of a face τ , where M is the metric of some site, we define the M -picking region of τ as the intersection of the M -ball $\mathcal{D}_M(c_\tau, \delta r_\tau)$ with the affine subspace generated by τ . For this reason, δ is called the *picking ratio*.

To avoid cascading constructions, we need to insert a point which is not γ -cospherical with any of the existing simplices. Writing $W(\tau)$ for the set of points that are γ -cospherical with a given simplex τ , we therefore need to bound the M -volume of $W(\tau)$.

In the following, we will rely on the slivers having been removed first, so that we can use this lemma:

Lemma 3.3.4. For any tetrahedron $vwx y$, with bounded radius-edge ratio $\rho \leq \rho_0$, and which is not a sliver, i.e. which has a sliverity ratio (the ratio between its volume and the cube of its shortest edge length) $\sigma \geq \sigma_0$, and for any direction n ,

$$\max(|\cos(n, vx)|, |\cos(n, vy)|, |\cos(n, vw)|) \geq \sin \alpha,$$

where $\alpha = \frac{1}{2} \sin^{-1} \frac{4\sigma_0}{\sqrt{3}\rho_0^2}$.

Lemma 3.3.5. *Let M be a metric, let $\tau = vwxy$ be some simplex with M -circumradius R , radius-edge ratio smaller than ρ_0 , and sliverity ratio (the ratio between its volume and the cube of its shortest edge length, as measured by M) greater than σ_0 , and let $\gamma_0 > 1$ be a distortion bound. The set $W(\tau)$ of points z such that v, w, x, y, z are γ_0 -cospherical is included in a region of M -volume $V_M < R^3 f(\gamma_0)$, where f is such that $f(x)$ tends to 0 when x tends to 1.*

Proof. Denote by c_M the center of $\mathcal{C}_M(vwxy)$. Denote by N and N' the two metrics involved in the definition of γ_0 -cosphericity. Assume that $vwxy$ is a Delaunay simplex for metric N . By definition, z is outside $\mathcal{C}_N(vwxy)$ but inside $\mathcal{C}_{N'}(vwxy)$. Denote by c_N and $c_{N'}$ the centers of these circumscribed spheres, for metrics N and N' respectively.

We can assume, without loss of generality, that N is the Euclidean distance. Recall that the Euclidean circumcenter of $vwxy$ can be expressed as

$$c_N = f(v, w, x, y) = v + \frac{c(1, 1, 1)}{\det(w - v, x - v, y - v)}, \text{ with}$$

$$\begin{aligned} c(1, 1, 1) &= (y - v)^2(w - v) \times (x - v) \\ &\quad + (w - v)^2(x - v) \times (y - v) \\ &\quad + (x - v)^2(y - v) \times (w - v) \end{aligned}$$

Denote now by A a square root of N' (see Section 3.2.1 for a definition of F_p , the square root of M_p). We can assume that $A = \text{Diag}(\lambda, \mu, \nu)$ with $0 < \lambda \leq \mu \leq \nu \leq \gamma_0$ and $\nu \geq 1/\lambda$ (by changing the frame of coordinates and exchanging N and N' if needed). We then have

$$\begin{aligned} c_{N'} &= A^{-1} f(Av, Aw, Ax, Ay) \\ &= v + \frac{A^{-1} \text{Com}(A)}{\det(A)} \frac{c(\mu_1, \mu_2, \mu_3)}{\det(w - v, x - v, y - v)} \\ &= v + A^{-2} \frac{c(\mu_1, \mu_2, \mu_3)}{\det(w - v, x - v, y - v)}, \text{ with} \end{aligned}$$

$$\begin{aligned} c(\mu_1, \mu_2, \mu_3) &= \mu_1(y - v)^2(w - v) \times (x - v) \\ &\quad + \mu_2(w - v)^2(x - v) \times (y - v) \\ &\quad + \mu_3(x - v)^2(y - v) \times (w - v) \end{aligned}$$

with $\lambda \leq \mu_1, \mu_2, \mu_3 \leq \nu$ and $\text{Com}(A) = \text{Diag}(\mu\nu, \nu\lambda, \lambda\mu)$. Furthermore, we have

$$c(\mu_1, \mu_2, \mu_3) \cdot (y - v) = \mu_1(y - v)^2 \det(w - v, x - v, y - v),$$

and the same formulas with cyclic permutations of w, x, y . It follows from Lemma 3.3.4 that

$$\left\| \frac{c(\mu_1, \mu_2, \mu_3) - c(1, 1, 1)}{\det(w - v, x - v, y - v)} \right\| \leq 3(\gamma_0 - 1) \frac{2R_N}{\sin(\alpha)},$$

where α is the one defined in Lemma 3.3.4.

We define

$$\tilde{c}_N = v + \frac{1}{\det(w - v, x - v, y - v)} c(\mu_1, \mu_2, \mu_3).$$

The triangular inequality then shows that $d_N(c_N, c_{N'}) \leq d_N(c_N, \tilde{c}_N) + d_N(\tilde{c}_N, c_{N'}) \leq 3(\gamma_0 - 1)2R_N / \sin(\alpha) + \|A^{-2} - I\|R_N \leq 7(\gamma_0^2 - 1)R_N / \sin(\alpha)$.

Finally, $d_N(c_N, c_{N'}) < C\gamma_0(\gamma_0^2 - 1)R$. Note that this inequality is valid for any metrics N, N' such that the distortions $\gamma(N, N'), \gamma(M, N), \gamma(M, N')$ are smaller than γ_0 . In particular, we may have $M = N$ or $M = N'$.

For metric M , $\mathcal{C}_N(vwxy)$ is an ellipsoid whose minor half-axis is bigger than R/γ_0 . It follows from the upper bound of the distance between c_N and $c_{N'}$ that $\mathcal{C}_N(vwxy)$ contains the Euclidean sphere centered at c_M with radius $(1/\gamma_0 - C\gamma_0(\gamma_0^2 - 1))R > (2 - \gamma_0 - C\gamma_0(\gamma_0^2 - 1))R$.

Similarly, for metric M , $\mathcal{C}_{N'}(vwxy)$ is an ellipse whose major half-axis is smaller than $\gamma_0 R$. It follows from the upper bound of the distance between c_N and $c_{N'}$ that $\mathcal{C}_{N'}(vwxy)$ is contained in the Euclidean sphere centered at c_M with radius $(\gamma_0 + C\gamma_0(\gamma_0^2 - 1))R$.

Finally, the volume V_M is bounded by $4/3\pi R^3 ((\gamma_0 + C\gamma_0(\gamma_0^2 - 1))^3 - (2 - \gamma_0 - C\gamma_0(\gamma_0^2 - 1))^3) = R^3 f(\gamma_0)$. \square

Similarly, we need to bound the M -area of the intersection of $W(\tau)$ with a plane and the M -length of the intersection of $W(\tau)$ with a line: in order to conform the mesh to the prescribed boundary, the algorithm may need to restrict the insertion of a point to a given triangle or segment.

Lemma 3.3.6 (Plane restriction). *Given a metric M , and a simplex $\tau = vwxy$ with M -circumradius R and radius-edge ratio smaller than ρ_0 , and a bound $\gamma_0 > 1$, the set $W(\tau)$ of points z such that v, w, x, y, z are γ_0 -cospherical, intersected with a plane π , is included in a region of M -area $V_M < R^2 g(\gamma_0)$, where g is such that $g(x)$ tends to 0 when x tends to 1.*

Lemma 3.3.7 (Line restriction). *Given a metric M , and a simplex $\tau = vwxy$ with M -circumradius R and radius-edge ratio smaller than ρ_0 , and a bound $\gamma_0 > 1$, the set $W(\tau)$ of points z such that v, w, x, y, z are γ_0 -cospherical, intersected with a line ℓ , is included in a region of M -length $V_M < Rh(\gamma_0)$, where h is such that $h(x)$ tends to 0 when x tends to 1.*

In Lemmas 3.3.6 and 3.3.7, functions g and h follow from the volume formulas of the intersection of a torus with a plane or a line.

Lemma 3.3.8. *Let ρ_0 be a positive bound, and let $\epsilon > 0$ be the shortest interdistance. There is at most a constant number $K(\rho_0)$ of possible new γ_0 -cospherical configurations p, q, r, s, t if a point p is inserted in the picking region $\mathcal{D}(c_\tau, \delta r_\tau)$ of a face τ with radius-edge ratio smaller than τ (see Figure 3.3).*

Proof. Let q, r, s, t be four points such that p, q, r, s, t are γ_0 -cospherical for metric M_p . Since q, r, s, t are within a bounded distance from p , a volume argument follows from the fact that all sites have an interdistance greater than ϵ . \square

Lemma 3.3.9. *If γ_0 is such that $K(\rho_0) \max(f(\gamma_0), g(\gamma_0), h(\gamma_0))\beta^3 < 4/3\pi\delta^3$, the set of points p that would create new γ_0 -cospherical configurations, with radius smaller than βr_τ and radius-edge ratio smaller than ρ_0 , does not cover the entire picking region.*

Proof. The total area of the set of points that may create such γ_0 -cospherical configurations upon insertion of p is smaller than $K(\rho_0)f(\gamma_0)(\beta r_\tau)^3$. If γ_0 is chosen so that

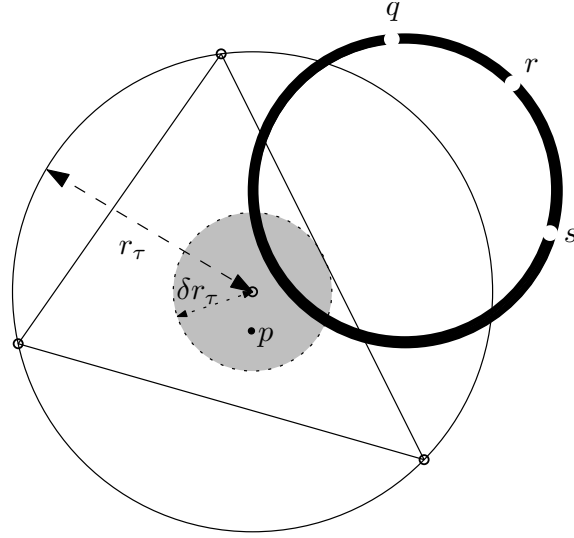


Figure 3.3: q, r, s define a forbidden region (black annulus) for p in the picking region (grey area)

this volume is smaller than the volume $4/3\pi\delta^3r_\tau^3$ of the picking region, the picking region is not entirely covered.

The same proof remains valid in the case of restricted picking if one replaces f by g and h . \square

3.3.4 Encroachment and Star Initialization

Let us now present how the boundary of the domain is preserved during the refinement process. We assume that the domain Ω to be meshed is a polyhedral domain in dimension 3. By preserving the boundary $\partial\Omega$ of the domain, we mean that the vertices, edges and faces of $\partial\Omega$ appear as elements of the final mesh.

As in the usual Delaunay refinement algorithms, this goal is reached by protecting the boundary $\partial\Omega$ from encroachment by inserted points. Let us recall these notions precisely, in the Euclidean context. See [41] for the original and detailed presentation of this method.

Definition 3.3.8. *A point p is said to encroach a boundary edge or facet f if p is inside the smallest circumscribing sphere of f . This sphere is called the diametral sphere of an edge, and the equatorial sphere of a facet. This sphere being empty is called the Gabriel property for f .*

Maintaining the Gabriel property for each boundary edge and facet provides the protection needed for the boundary. Recall that maintaining the Gabriel property of boundary edges and facets upon insertion of a new site v means applying the insertion function `Insert_or_snap_e(v)` defined as follow:

- `GInsert_or_snap_e(c)`:
if c encroaches some boundary edge e , insert the circumcenter of e . Otherwise, `Insert_or_snap_f(c)`.
- `GInsert_or_snap_f(c)`:
if c encroaches some boundary triangle f , insert the circumcenter of f . Otherwise, insert (c) .

In this manner, all protected edges and facets do appear in the final mesh and no circumcenter is ever inserted outside the domain.

In our context, we do the same for each of the stars: all constraints are inserted in all stars, and the Gabriel property is maintained in each star for the corresponding metric.

Note that in practice, as soon as the conflict zone Z_v of S_v has an empty intersection with the union of the diametral balls of the constraints, updating S_v is done without taking the constraints into account anymore. This immediately follows from the fact that Z_v is a non increasing set.

This procedure guarantees that boundary facets and edges will be kept all along the algorithm.

3.4 ALGORITHM

3.4.1 Algorithm Outline

The refinement algorithm that we consider constructs the set of sites V in a greedy way while maintaining the set of stars $\{S_v\}_{v \in V}$ and the corresponding sets of constraints E_v whose diametral balls intersect Z_v .

The algorithm refines the simplices of inconsistent stars, until inconsistent stars disappear. Once all stars are consistent, they can be merged together to form a triangulation \mathcal{T} of the domain, with the property that the 1-neighborhood of any vertex v in \mathcal{T} is Delaunay for metric M_v . For this reason, we call the resulting triangulation a *locally uniform anisotropic mesh*.

As we have seen in Section 3.3.3, simply refining inconsistent simplices by inserting their circumcenter does not allow to maintain a lower bounded insertion radius, which is the condition for the algorithm to terminate. In order to avoid this problem, we manage not to create a forbidden quasi-cospherical configuration by selecting a suitable new site in the picking region around the circumcenter of the simplex to be refined: a point is picked randomly. If it creates any γ -cospherical configuration with γ too small, it is discarded, and a new point is picked in the picking region.

Let $\gamma_0 > 1$, $\delta > 0$, $\rho_0 > 0$ and $\beta > 0$ be constants to be specified in Section 3.4.2. In order to describe precisely the algorithm, we define the insertion procedures to be used. Face τ is either a simplex, a triangle or an edge:

- `Pick_valid(τ, M)`:
denote by c and r the center and radius of $\mathcal{C}_M(\tau)$. Pick randomly a point x in the picking region $B_M(c, \delta r) \cap H$, where H is the affine subspace spanned by τ . If there exists points p, q, r, s such that $xpqr$ is a new simplex with $\gamma(xpqr) < \gamma_0$ and x, p, q, r, s are γ_0 -cospherical with radius smaller than βr_τ and radius-edge ratio smaller than ρ_0 , discard x and pick another random point x , until no such points p, q, r, s exist. Return x .
- `Pick_valid_for_sliver(τ, M)`:
denote by c and r the center and radius of $\mathcal{C}_M(\tau)$. Pick randomly a point x in the picking region $B_M(c, \delta r) \cap H$, where H is the affine subspace spanned by τ . If inserting x creates any sliver tetrahedron (i.e. a tetrahedron with $\sigma < \sigma_0$), discard x and pick another random point x , until no such sliver tetrahedron appears. Return x .

- `Refine(τ): Insert_or_snap_e(Pick_valid(τ, M))`,
where M is the metric of the star that is being refined.
- `Refine_sliver(τ): Insert_or_snap_e(Pick_valid_for_sliver(τ, M))`,
where M is the metric of the star that is being refined.
- `Insert_or_snap_e(c):`
if c encroaches some boundary edge e , `Refine(e)`. Otherwise,
`Insert_or_snap_f(c)`.
- `Insert_or_snap_f(c):`
if c encroaches some boundary triangle f , `Refine(f)`. Otherwise, insert c .

The algorithm consists of applying the following rules. Rule (i) is applied only if Rule (j) with $j < i$ cannot be applied:

Rule (1) Encroachment: Refine encroached elements (edges and then faces) e by calling `Refine(e)`.

Rule (2) Distortion: If a simplex τ is such that $\gamma(\tau) \geq \gamma_0$, `Refine(τ)`;

Rule (3) Radius-edge ratio: If a simplex τ of S_v is such that $\rho_{M_v}(\tau) > \rho_0$, `Refine(τ)`;

Rule (4) Sliverity: If a simplex $\tau = vxyz$ of star S_v has a sliverity ratio (the ratio between its volume and the cube of its shortest edge length, as measured by M_v) smaller than σ_0 for M_v , `Refine_sliver(τ)`.

Rule (5) Cosphericity: If a simplex $\tau = vxyz$ of star S_v is such that there exists a site p such that v, x, y, z, p are γ_0 -cospherical for M_v , `Refine(τ)`.

Once the algorithm terminates, a simple sweep allows to merge all the stars into the final *locally uniform anisotropic mesh*.

3.4.2 Termination of the Algorithm and Quality of the Mesh

Let us now prove that the algorithm presented in the previous section does terminate, for suitable choices of distortion bound γ_0 , picking ratio δ , radius-edge ratio ρ_0 and size ratio β . Let us consider the refinement rules, in their order of priority.

Lemma 3.4.1. *Assume that for any boundary edge e , the angle between the two boundary facets incident to e , computed for the metric of any point belonging to e , is greater than 90° . Then Rule (1) is applied only a finite number of times during the algorithm.*

Proof. Once the boundary is sufficiently refined, the diedral angle at any boundary edge, as computed for the metric at any point in the star of its vertices, is greater than 90° , thanks to the continuity of the metric field. At this point, the usual proofs apply. \square

Denote by ϵ_1 the shortest interdistance between sites once Rule (1) cannot be applied anymore. Recall the definition $\Gamma = \max_{x,y \in \Omega} \gamma(x,y)$. Let us now consider the shortest interdistance created by Rule (2):

Lemma 3.4.2. *Let $\gamma_0 > 0$ be a distortion bound. Denote by r_0 the minimal bounded distortion radius associated to γ_0 . Any simplex τ such that $\gamma(\tau) > \gamma_0$ can be refined, while creating no interdistance shorter than $(1 - \delta)^3 r_0 / (4\Gamma^3)$.*

Proof. If a M_x -sphere $\mathcal{C}(x, r)$ has a radius r less than $r_0/2$, then $\gamma(p, q) < \gamma_0$ for any $p, q \in \mathcal{C}$. Let τ be a simplex such that $\gamma(\tau) > \gamma_0$, and denote by a a vertex of τ such that $\gamma(x, y) > \gamma_0$ for two points x, y which are inside $\mathcal{C}_a(\tau)$. It follows that $R_{M_a}(\tau) > r_0/2$. Denote by c_a the center of $\mathcal{C}_a(\tau)$. For any site $w \neq a$, and any point x in the picking region around c_a , we have $d_w(x) \geq d_a(w, x)/\Gamma \geq (d_a(w, c_a) - \delta R_{M_a})/\Gamma$. The Delaunay empty ball property then implies $(d_a(w, c_a) - \delta R_{M_a})/\Gamma \geq (1 - \delta)R_{M_a}/\Gamma$ and by the high distortion condition $R_{M_a}(\tau) > r_0/2$, we finally have $(1 - \delta)r_a/\Gamma \geq (1 - \delta)r_0/(2\Gamma)$.

To summarize, we have proved that $d_w(x) \geq (1 - \delta)r_0/(2\Gamma)$. The same lower bound is obviously also valid for $d_x(w)$.

In case boundary elements are encroached, the same proof can be applied to the boundary elements instead of τ , with a penalty of a factor at most $(1 - \delta)^2/(2\Gamma^2)$: if a point x , chosen in the picking region of a simplex τ , encroaches a boundary facet f (for a metric M), the distance r_x from x to any site is at most $\sqrt{2}R_M(f)$. Furthermore, as we have seen in the first part of the proof, the point y picked in the picking region of f has a distance r_y to any site of at least $(1 - \delta)R_M(f)/\Gamma$. It follows that $r_y \geq (1 - \delta)R_M(f)/\Gamma \geq (1 - \delta)r_x/(\sqrt{2}\Gamma)$.

Hence, the penalty for one encroachment is a factor of $(1 - \delta)/(\sqrt{2}\Gamma)$. It follows that the penalty for two consecutive encroachments (of a face and then of an edge) is a factor of $(1 - \delta)^2/(2\Gamma^2)$. This concludes the proof. \square

Denote by ϵ_2 the shortest interdistance obtained after Rule (1) and Rule (2) have been applied: $\epsilon_2 = \min(\epsilon_1, (1 - \delta)^3 r_0/(4\Gamma^3))$. In the following, we can assume that all simplices have a distortion less than γ_0 , and that the interdistance is greater than $\epsilon_2 > 0$. In case simplices with high distortion were to appear again later in the process, the previous lemma shows that we could again refine them and maintain the same bound ϵ_2 . Let us now consider the case of simplices with high radius-edge ratio.

Lemma 3.4.3. *If $(1 - \delta)^3 \rho_0 > 2\gamma_0^3$, refining the simplices with a radius-edge ratio larger than ρ_0 does not decrease the shortest interdistance.*

Proof. Denote by ϵ the shortest interdistance before the refinement of a simplex with a radius-edge ratio larger than ρ_0 . In a way similar to the proof of Lemma 3.4.2, one computes easily that after the refinement, the shortest interdistance is still greater than $(1 - \delta)^3 \rho_0 \epsilon / (2\gamma_0^3)$. The result follows. \square

Hence, ϵ_2 remains lower bound of the interdistance.

Lemma 3.4.4. *For a sufficiently small σ_0 , Rule (4) can be applied, and does not decrease the shortest interdistance by more than $(1 - \delta)/4$.*

Proof. In the Euclidean case, Li and Teng[34] provide in Section 4.1 a bound on σ_0 which allows the picking to occur by guaranteeing a positive volume for the set of points which can be accepted by the picking method. This bound relies on a volume argument, and on the fact that all tetrahedra have a good radius-edge ratio $\rho < \rho_0$.

In our case, thanks to Rule (3), we have the same condition on ρ . However, the volumes of the forbidden regions have to be computed in different metrics. Using very conservative bounds, we obtain that our bound σ_0 should be at least Γ^4 times the Euclidean bound (Γ^3 would be enough for the 3D case, but the 2D and 1D cases needed for dealing with encroachment require Γ^4). With this bound σ_0 , Rule (4) can be applied because there exists a valid point to be picked, and Lemma 4.2 of [34] show that the minimal interdistance does not decrease by more than a factor $(1 - \delta)/4$. Since this result is true for each local (Euclidean) triangulation, this is true globally. \square

This proves that applying Rule (4) keeps the interdistance greater than $\epsilon_3 = \frac{1-\delta}{4}\epsilon_2$.

Finally, we can compute how much the interdistance is decreased when Rule (5) is applied.

Lemma 3.4.5. *Let $\tau = vxyz$ be a simplex of star S_v with a site p such that v, x, y, z, p are γ_0 -cospherical for M_v . Refining all such configurations does not create any interdistance shorter than $(1 - \delta)^3 \epsilon_3 / 2$ if $(1 - \delta)^3 \beta > 2\gamma_0^3$.*

Proof. Denote by ϵ the current shortest interdistance. In a way similar to the proof of Lemma 3.4.2, one computes easily that the shortest interdistance after the refinement of such a γ_0 -cospherical configuration stays bigger than $(1 - \delta)^3 \epsilon / (2\gamma_0^3)$.

Recall that, thanks to the definition of `Pick_valid`, no γ_0 -cospherical configuration is ever created by the refinement of any simplex τ , except γ_0 -cospherical configurations with radius bigger than βr_τ or radius-edge ratio bigger than ρ_0 . If the radius-edge ratio is bigger than ρ_0 , the configuration is to be refined by Rule (3). As we have just seen, if the radius is bigger than βr_τ , the shortest interdistance created to refine this new γ_0 -cospherical configuration is at least $(1 - \delta)^3 \beta r_\tau / (2\gamma_0^3)$. Hence, if we choose β large enough, so that $(1 - \delta)^3 \beta > (2\gamma_0^3)$, refining this kind of new γ_0 -cospherical configuration does not reduce the shortest interdistance.

It follows that $(1 - \delta)^3 \epsilon_2 / 2$ is a lower bound on the interdistance after applying Rule (4), under the condition that $(1 - \delta)^3 \beta > (2\gamma_0^3)$. \square

Lemma 3.4.1, 3.4.2, 3.4.3, 3.4.4 and 3.4.5 show that the insertion radius admits a positive lower bound. This concludes the proof of the termination of the algorithm. Let us summarize this result in the following theorem, which also relies on Lemma 3.3.8:

Theorem 3.4.1. *Given a polyhedral domain Ω and a continuous metric field over Ω , and the following properties for the parameters of the algorithm,*

- *the dihedral angle at each boundary edge e , computed for the metric of any point of e , is greater than 90° ;*
- *ρ_0 is larger than 2;*
- *δ is small enough, so that $(1 - \delta)^3 \rho_0 > 2$;*
- *β is large enough, so that $(1 - \delta)^3 \beta > 2$;*
- *γ_0 is close enough to 1, so that $K(\rho_0, \beta) \max(f(\gamma_0), g(\gamma_0), h(\gamma_0)) \beta^2 < 4/3\pi\delta^2$ and $(1 - \delta)^3 \beta > 2\gamma_0^3$ and $(1 - \delta)^3 \rho_0 > 2\gamma_0^3$.*

the refinement algorithm terminates, with a lower bound ρ_0 on the radius-edge ratio of the elements and an upper bound γ_0 on the distortion of the simplices. \square

Note that these bounds ρ_0 and γ_0 ensure that eventually all simplices are well-shaped for the metrics of their vertices. This guarantees the quality of the final mesh.

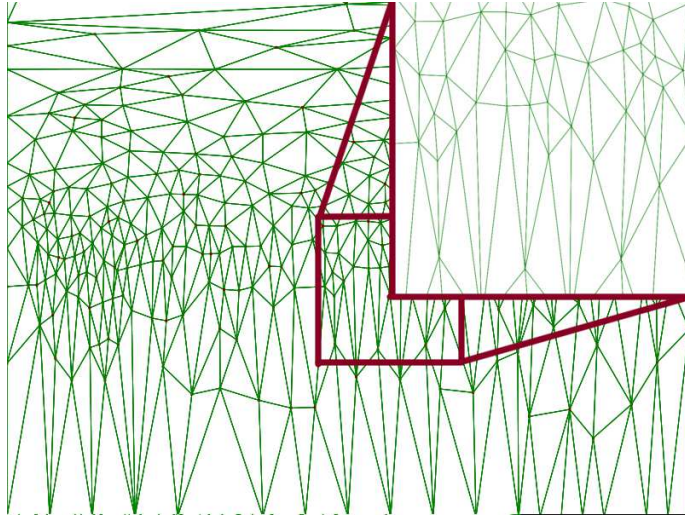


Figure 3.4: Output of the algorithm with a zoom on the central part: the red lines delimit the zoomed region.

3.5 CONCLUSION

We have proposed a new definition for an anisotropic mesh and an algorithm to generate such a mesh. The algorithm is the first to offer guarantees in 3-space. Moreover, the algorithm is simple and has been implemented in the plane in C++ using CGAL.

Although the implementation has not been optimized, we had still a much more scalable algorithm than the one we proposed in [7]: our datastructure has asymptotically the same space complexity as a triangulation of the same pointset. Interestingly, the assumption that the metric field was continuous appeared crucial not only in theory, but also in practical tests: discontinuities typically prevent the algorithm from terminating, because the algorithm refines the locus of the discontinuity (usually a curve) indefinitely.

Figure 3.4 shows the output of the algorithm on a domain where the metric is stretched horizontally in the upper part and vertically in the lower part. In this example, we did not enforce any size bound, so that the variable density of the result clearly shows where more refinement was needed for removing inconsistencies. As expected, the higher densities are located along the line of high distortion, where the eigenvectors exchange their eigenvalues.

Yuanmi Chen implemented the algorithm in 3D, an example of the output is presented in Figure 3.5.

Future directions of work include

- allowing more general constraints, in particular constraints with sharp edges, and using a protection scheme to avoid cascading insertions in the neighborhood of these edges;
- dealing with discontinuities by protecting points of discontinuity and by considering the curves of discontinuity as constraints of the triangulation;
- extending the results in dimension $d > 3$.

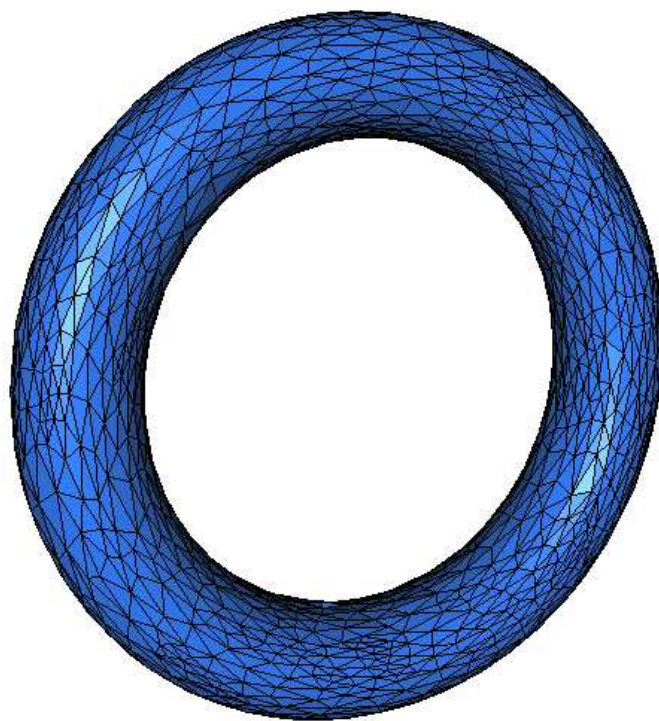


Figure 3.5: Output of the algorithm, implemented in 3D by Yuanmi Chen

Part IV

Greedy Ad Hoc Routing

GREEDY POWER ROUTING

OVERVIEW

In this chapter, we present an application of power diagrams to the field of adhoc networks routing. We consider a greedy routing scheme for planar 3-connected graphs, first introduced by Ben Chen et al. [12]. The embedding is in \mathbb{R}^2 , but the proximity measure used is not Euclidean. We show the relationship between this embedding and classical circle packings, and show how to modify Thurston's iterative algorithm for computing circle packings to compute our embeddings in a distributed manner.

The greedy routing scheme is described in Section 4.3. We prove that it is in fact equivalent to greedy polyhedral routing in Section 4.3.3.

Contents

4.1 Sensor Networks	90
4.2 Previous Work	90
4.3 Greedy Power Routing	91
4.3.1 Power Routing	91
4.3.2 Contained Power Diagrams	91
4.3.3 Equivalence to Polyhedral Routing	94

4.1 SENSOR NETWORKS

Sensor networks are a collection of (usually miniature) devices, each with limited computing and wireless communication capabilities, distributed over a physical area. The sensor network collects data from its environment and should be able to integrate it and answer queries related to this data. Sensor networks are becoming more and more attractive in many application domains.

The advent of sensor networks has posed a number of research challenges to the networking and distributed computation communities. Since each sensor can typically communicate only with a small number of other sensors within a short range, information generated at one sensor can reach another sensor only by routing it through the network. Traditional routing algorithms rely only on the combinatorial connectivity graph of the network, but the introduction of so-called *location-aware* sensors, namely, those that also *know* what their physical location is (e.g. by using a GPS receiver), permit more efficient *geographic* or *geometric* routing.

In geometric routing we consider the following problem: A packet is to be routed across the network from a source sensor to a destination sensor. The physical locations – the *coordinates* – of the source and destination sensors are known. When a sensor receives a packet, it must decide to which of its *neighbors* it should forward the packet based on a *local* decision. By local decision, we mean that the decision is made based *only* on local information - the coordinates of the current sensor, the destination, and the sensor's neighbors. Despite this restrictive locality, the routing algorithm should guarantee that the packet will indeed arrive at the destination.

4.2 PREVIOUS WORK

One simple geometric routing scheme is *greedy routing*. In greedy routing, when a sensor receives a packet, it forwards the packet to the neighbor that is *closest* in some sense to the destination sensor. The main problem with greedy routing is that it may encounter local minima, also known as *routing voids* or *holes*, when the current sensor has no neighbor closer to the destination than itself. When such a local minimum is encountered, the packet is "stuck", greedy routing cannot continue, and the delivery fails. Examples of greedy routing are greedy Euclidean routing, which is based on Euclidean distance to the destination, or compass routing, based on angular distance to the destination [32]. An important question is the design of proximity (i.e. closeness) measures that guarantee the delivery of all packets, irrespective of the source or destination node. Since this measure is usually a distance in some space where the nodes have been embedded, the problem of positioning the nodes in such a space is referred to as the problem of computing a greedy embedding of a given network.

The most natural example of greedy routing is greedy Euclidean routing, where the proximity of nodes is measured simply by the Euclidean distance. This scenario has been studied in detail by Papadimitriou and Ratajczak [37], who conjectured that any 3-connected planar graph admits a greedy Euclidean embedding, namely, a greedy embedding for the Euclidean distance. An easy example is the subset of Delaunay-realizable triangle graphs, since it is easy to check that Delaunay triangulations are greedy Euclidean embeddings of their underlying graph.

While not able to prove their conjecture, Papadimitriou and Ratajczak propose other greedy routing schemes, most notably, 3D polyhedral routing. This consists of embedding the 3-connected planar graph as a polyhedron edge-tangent to the unit

sphere in \mathbb{R}^3 . A packet is routed by forwarding it to the neighbor vertex that maximizes the dot-product with the destination vertex. Such an embedding always exists, and they prove that the routing scheme always delivers.

Recently, Dhandapani [19] proved the conjecture of Papadimitriou and Ratajczak for the special case of a triangle graph. Using Schnyder embeddings of triangulations of the sphere, Dhandapani showed the existence of a greedy Euclidean embedding for any such triangulation. Unfortunately, the proof is not constructive.

Other spaces have been considered as embedding spaces for the greedy routing problem. Kleinberg [30] studied the question of embedding the network in the hyperbolic plane, and was able to construct a greedy embedding in the hyperbolic plane for every connected finite graph. Note that the graph is neither assumed to be planar, nor to have any particular connectivity property, beyond connectedness. For these reasons, Kleinberg's results are particularly valuable in practical implementations. Furthermore, a distributed algorithm is presented, allowing the network to compute its own embedding, at the expense of a few broadcasting operations.

4.3 GREEDY POWER ROUTING

4.3.1 Power Routing

The routing algorithm is greedy routing where the nodes are embedded as circles in the plane, and the circle power functions are used as distance functions. Namely, to route to destination t when at vertex v , forward to the neighboring vertex u such that $u = \operatorname{argmin}_{w \in N(v)} \operatorname{Pow}(w, t)$, where $N(v)$ is the set of neighbors of v . This routing scheme was first introduced by Ben Chen et al. [12].

In general, this greedy routing algorithm is not guaranteed to deliver. However, the freedom to choose the radius of each circle gives us some flexibility beyond the usual Euclidean distance so that the embedding can be made greedy.

4.3.2 Contained Power Diagrams

An orthogonal dual of a convex tiling is a planar embedding of the graph dual to the tiling, such that primal-dual edge pairs lie on orthogonal lines. We consider the setting in which the faces dual to boundary vertices are unbounded, and the vertex dual to the outer face is not embedded. For a 3-connected planar graph, there may exist many orthogonal primal/dual embedding pairs. Here we will be interested in pairs with a special property.

Definition 4.3.1. *A contained embedding of a 3-connected planar graph is an orthogonal primal/dual embedding pair, such that each primal vertex is strictly contained in its dual face.*

Lemma 4.3.1. *Any 3-connected planar graph and its dual have a contained embedding.*

Proof. The celebrated *kissing disks* theorem of Koebe and Andre'ev [31] states that any 3-connected planar graph and its dual can be simultaneously embedded in the plane such that each face is a convex polygon with an inscribed circle whose center coincides with the vertex of the dual corresponding to the face, and such that edges

are perpendicular to their dual edges. Such an embedding is by definition a contained embedding (see Figure 4.1). \square

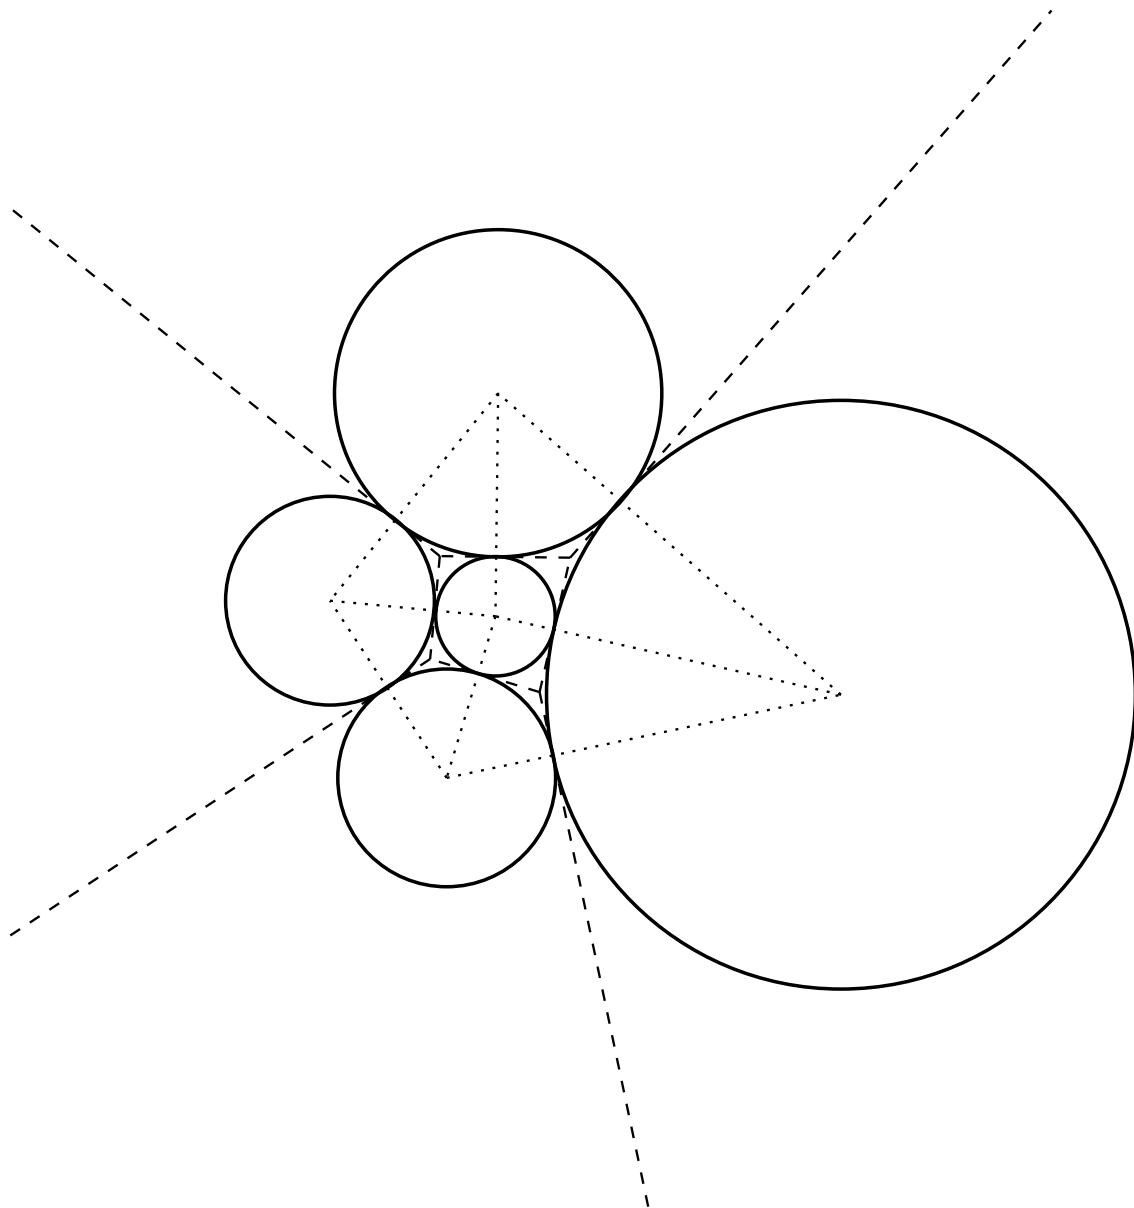


Figure 4.1: A contained embedding obtained from the *kissing disks* theorem of Koebe and Andre'ev.

Note that a contained embedding of a graph is not necessarily unique. For example, if the graph happens to be a Delaunay-realizable triangulation, then any Delaunay realization and its dual Voronoi diagram are also a contained embedding for that graph.

As we have seen in the previous section, such contained embeddings are contained power diagrams. In terms of power diagrams, we have the following definition:

Definition 4.3.2. A power diagram is said to be contained if each site is contained in its cell (see Figure 4.2).

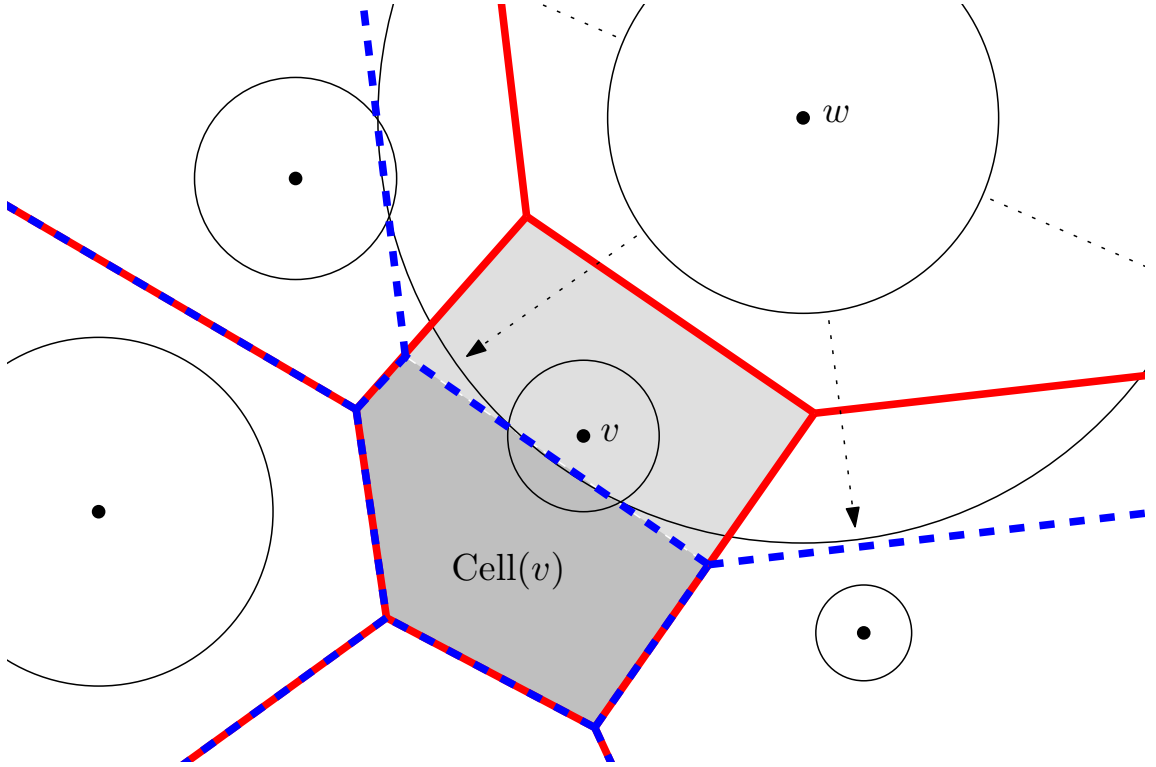


Figure 4.2: As the radius of the circle around w grows, $\text{Cell}(w)$ grows and $\text{Cell}(v)$ shrinks. The power diagram becomes uncontained when v is no longer in $\text{Cell}(v)$.

This key containment property is a sufficient condition for the greedy power routing to deliver. To state this result, we adopt the following notations: let $G(V, E)$ be a combinatorial triangulation. Assume that G is planar and denote by B its boundary, which is a cycle. In the following, we study a map $\phi : V \rightarrow D^2 \times \mathbb{R}$, which associates to each vertex v a point $p(v)$ in the unit disk and a scalar weight $\sigma(v)$. We denote by $\text{Conv}(p(V))$ the convex hull of the associated points.

Theorem 4.3.1. *If the restriction of the power diagram of $\phi(V)$ to $\text{Conv}(p(V))$ is contained and if its adjacency graph (i.e. the combinatorial dual) is a subgraph of G , then greedy power routing delivers on ϕ .*

Proof. First note that in the special case that the embedding is a Delaunay triangulation, then all the radii are equal and greedy power routing is the same as greedy Euclidean routing.

In the general case, we must show that given a destination vertex t , each vertex v has a neighbor u in G such that $\text{Pow}(u, t) < \text{Pow}(v, t)$. This may be shown using an argument similar to that of Bose et al. [9] that the Delaunay triangulation is greedy. Specifically, consider the power diagram of the primal vertices with the given radii. Let e be the first edge of the power diagram which the line $v \rightarrow t$ intersects. There must exist such an edge, because in a contained embedding each vertex is strictly contained in its dual face, so v and t must lie in different cells of the power diagram. Let u be the vertex whose cell is adjacent to v 's cell through e , and l be the line supporting e . Edge e is part of the restricted power diagram. Since the adjacency graph of the restricted power diagram is a subgraph of G , u is a neighbor of v in G .

Every point x on l is equidistant from u and v : $\text{Pow}(x, v) = \text{Pow}(x, u)$. Every point y in the half plane created by l that contains u is closer to u than to v : $\text{Pow}(y, u) < \text{Pow}(y, v)$. By the definition of u , t lies in the half plane which is closer to u , hence $\text{Pow}(t, u) < \text{Pow}(t, v)$. It remains to show that the routing terminates at the destination vertex t . But, by construction, every vertex is strictly contained in its dual cell, hence all vertices $v \neq t$ in the embedding satisfy $\text{Pow}(t, t) < \text{Pow}(t, v)$. Thus, $\text{Pow}(\cdot, t)$ has a global minimum at t . This concludes the proof. \square

4.3.3 Equivalence to Polyhedral Routing

Before going deeper into the study of greedy power routing, we first show the equivalence between greedy power routing and greedy polyhedral routing, as described by Papadimitriou and Ratajczak [37].

Definition 4.3.3. Greedy polyhedral routing is greedy routing among the vertices of a convex polyhedron in \mathbb{R}^3 containing the origin O , by greedily maximizing $\langle Ov, Ot \rangle$ where v is the current vertex, and t the destination vertex.

We use elementary geometric arguments, but rely on what is simply the projective equivalence of polarities with respect to the paraboloid and with respect to the sphere.

Polarity

Denote by \mathbb{S}^2 the unit sphere of \mathbb{R}^3 , and by O its center.

Definition 4.3.4. The polar hyperplane of a point P different from O , denoted $\pi(P)$, is the plane defined by the equation $\langle \vec{OP}, x \rangle = 1$. The point P is called its polar point. We denote by $\mathcal{C}(P)$ the intersection $\pi(P) \cap \mathbb{S}^2$, and by $\text{pr}_P(Q)$ the oriented distance between O and the projection of Q on (OP) : $\text{pr}_P(Q) = \langle OP, OQ \rangle / \sqrt{\langle OP, OP \rangle}$.

In other words, if P is outside \mathbb{S}^2 , the circle $\mathcal{C}(P)$ is the locus of points x such that (Px) is tangent to \mathbb{S}^2 , and $\pi(P)$ is the plane containing $\mathcal{C}(P)$. Note that, by definition, $\pi(P)$ is orthogonal to (OP) . Let us now recall the following lemma.

Lemma 4.3.2. For any two points P and Q outside \mathbb{S}^2 ,

$$P \in \pi(Q) \Leftrightarrow Q \in \pi(P) \Leftrightarrow \mathcal{C}(P) \perp \mathcal{C}(Q).$$

Proof. In the following, we use the notations $p = \vec{OP}$ and $q = \vec{OQ}$. We have the following equivalences:

$$P \in \pi(Q) \Leftrightarrow \langle q, p \rangle = 1 \Leftrightarrow Q \in \pi(P).$$

Furthermore, the tangent vectors to $\mathcal{C}(P)$ and $\mathcal{C}(Q)$ at an intersection point x are collinear to $p \times x$ and $q \times x$. The scalar product of these vectors is

$$\langle p \times x, q \times x \rangle = \langle (q \times x) \times p, x \rangle = \langle q, p \rangle \langle x, x \rangle - \langle x, p \rangle \langle x, q \rangle = 1 \cdot 1 - 1 \cdot 1 = 0.$$

This concludes the proof. \square

It easily follows from this lemma that the oriented angle of intersection $\alpha(P, Q)$ of two circles $\mathcal{C}(P)$ and $\mathcal{C}(Q)$ (0° in the case of tangency) depends only on the distance $\text{pr}_P(Q)$ between O and the projection of Q on (OP) , and is a locally increasing function of this parameter. If we restrict Q so that $\mathcal{C}(Q)$ does not contain P we obtain an increasing function.

Stereographic Projection

Recall that the stereographic projection and its inverse $\psi : \mathbb{R}^2 \rightarrow \mathbb{S}^2$ map circles to circles and preserve the angles of intersection between circles.

As in the previous section, we denote by $\text{pr}_P(Q)$ the distance between O and the projection of Q on (OP) . Let D , C_1 and C_2 be three circles in the plane, such that $\text{Pow}(D, C_1) < \text{Pow}(D, C_2)$. This property does not depend on the radius of D . Thus, by adapting this radius, we may assume that C_1 and D intersect. Let us first assume that C_2 intersects D too. In this case, the angles of intersection satisfy $\alpha(D, C_1) > \alpha(D, C_2)$. Denote by P , Q_1 and Q_2 the points such that $\psi(D) = \mathcal{C}(P)$, $\psi(C_1) = \mathcal{C}(Q_1)$ and $\psi(C_2) = \mathcal{C}(Q_2)$. Since ψ preserves the angles of intersection, we have $\alpha(P, Q_1) > \alpha(P, Q_2)$. The previous section then implies that $\text{pr}_P(Q_1) > \text{pr}_P(Q_2)$. If C_2 does not intersect D , considering a second larger circle with the same center as D and orthogonal to C_2 provides the same conclusion.

This fact can be summarized as follows:

Lemma 4.3.3. *If \mathcal{X} is a set of circles, and Y is another circle such that no circle of \mathcal{X} contains the center of Y , then for any point P the extrema*

$$\min_{X \in \mathcal{X}} \text{Pow}(P, X) \text{ and } \max_{X \in \mathcal{X}} \text{pr}_P(\mathcal{C}^{-1}(\psi(X)))$$

are obtained at the same $X_0 \in \mathcal{X}$.

Note that these two quantities are mapped to each other by a homography. This explains why a restriction is needed in order to have a monotonic function.

Routing Equivalence

Given a set of circles \mathcal{X} such that no circle contains the center of another circle, Lemma 4.3.3 shows that greedy polyhedral routing on $\mathcal{C}^{-1}(\psi(X))$ (see Definition 4.3.3) generates exactly the same paths as greedy power routing on \mathcal{X} .

It follows that any set of circles on which greedy power routing delivers, composed with the mapping $\mathcal{C}^{-1} \circ \psi$, provides a polyhedron on which greedy polyhedral routing delivers. Furthermore, the following lemma relates the equivalent special cases which interest us:

Lemma 4.3.4. *The transformation $\mathcal{C}^{-1} \circ \psi$ maps a circle packing to a polyhedron edge-tangent to \mathbb{S}^2 (see Figure 4.3).*

Proof. ψ maps tangent circles in the plane to tangent circles on the sphere. Denote by $\mathcal{C}(P)$ and $\mathcal{C}(Q)$ two such tangent circles on \mathbb{S}^2 , with P and Q being their polar points, and denote by T their tangency point. By construction, (PT) is tangent to \mathbb{S}^2 at T and orthogonal to $\mathcal{C}(P)$ at T . Similarly, (QT) is tangent to the sphere at T and orthogonal to $\mathcal{C}(Q)$ at T . Hence, both lines (PT) and (QT) belong to the tangent plane of \mathbb{S}^2 at point T , and both are orthogonal to the common tangent of $\mathcal{C}(P)$ and $\mathcal{C}(Q)$ at T . It follows that P , T and Q are co-linear. This proves that the segment linking the images of two tangent circles by $\mathcal{C}^{-1} \circ \psi$ is tangent to \mathbb{S}^2 . The result follows. \square

One can also show the connection between the containment property of power diagrams, and the property that Papadimitriou and Ratajczak proved to be a sufficient condition for greedy polyhedral routing to deliver:

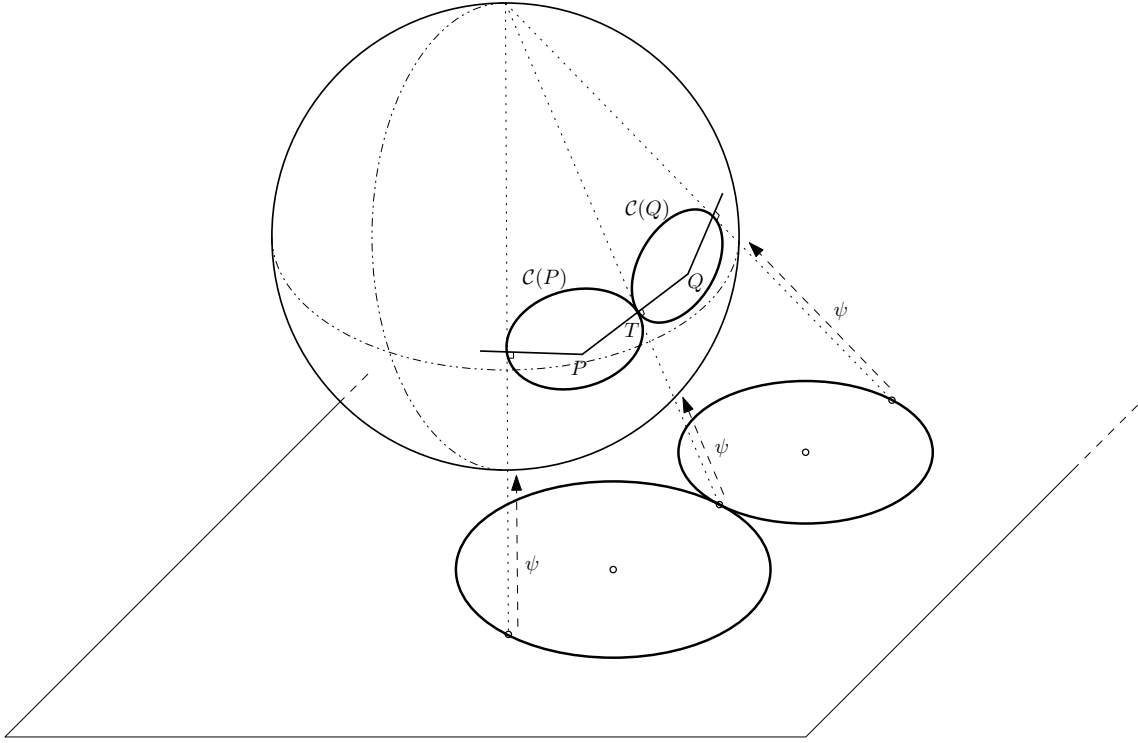


Figure 4.3: Two circles in the plane and their images (P and Q) on the sphere through $\mathcal{C}^{-1} \circ \psi$

Definition 4.3.5. Let P be a convex polyhedron in \mathbb{R}^3 containing the origin O . A supporting hyperplane of P at vertex v of P is a hyperplane that contains v but does not intersect P otherwise.

A polyhedron P is said to have orthogonal support if for each vertex v of P , the plane orthogonal to (Ov) and containing v is a supporting hyperplane.

Papadimitriou and Ratajczak proved that having orthogonal support is a sufficient condition for a polyhedron to provide greedy routing that delivers.

Lemma 4.3.5. A set of circles defines a contained power diagram if and only if its image by $\mathcal{C}^{-1} \circ \psi$ is a polyhedron P with orthogonal support.

Proof. Denote by v and w two vertices of a convex polyhedron P . Circles $C_v = C(c_v, r_v) = \psi^{-1} \circ \mathcal{C}(v)$ and $C_w = C(c_w, r_w) = \psi^{-1} \circ \mathcal{C}(w)$ are the corresponding circles in the plane. Denote by h_v the plane orthogonal to (Ov) and containing v . Then w belongs to h_v if and only if the radical axis of C_v and C_w (i.e. the power diagram bisector, which is, in this case, the line passing through the intersection points of the two circles) passes through c_v . In other words, w belongs to h_v if and only if c_v belongs to the boundary of the power region of C_w in the power diagram of $\{C_v, C_w\}$.

The result then follows from Lemma 4.3.3. \square

This completes the parallel between the two routing schemes. This parallel implies that the algorithm we design in the following sections allows the computation of more general greedy polyhedral embeddings than those edge-tangent to the sphere, as proposed by Papadimitriou and Ratajczak.

ALGORITHMIC SOLUTIONS

OVERVIEW

In this chapter, we present a novel way of using Thurston algorithm for the distributed computation of greedy embeddings, i.e. embeddings for which the greedy routing algorithm works.

Section 5.1 reviews the notion of circle packings and Thurston's algorithm to compute them, while the design of suitable termination conditions for the Thurston algorithm is studied in Section 5.2. Section 5.3 is devoted to the description of the algorithm, before a final discussion on our validation experiments and future work in Section 5.4.

Contents

5.1 Circle Packing	98
5.1.1 Definitions	98
5.1.2 Practical Computations of Circle Packings	99
5.2 Local Termination Conditions	100
5.2.1 Triangulated Case	100
5.2.2 Generalized Papadimitriou and Ratajczak Result	104
5.2.3 Non Triangulated Case	105
5.2.4 Relation between Circle Packings, LPD and GLPD	108
5.3 Algorithms	108
5.3.1 Computing a Greedy Power Embedding	109
5.3.2 Termination	109
5.3.3 Locality	110
5.4 Discussion	110
5.4.1 Experimental Validation	110
5.4.2 Possible Improvements	111

5.1 CIRCLE PACKING

5.1.1 Definitions

As we have seen in Lemma 4.3.1, kissing disks, also called circle packings, are an example of contained power diagrams, and, as such, are a greedy power embedding of their tangency graph. More formally:

Definition 5.1.1. *Given a planar triangulation $G(V, E)$, a G -circle packing is a set \mathcal{C} of circles in the plane with a bijection $\gamma : V \rightarrow \mathcal{C}$ such that $\gamma(v)$ and $\gamma(w)$ are externally tangent if and only if $\{v, w\}$ is an edge of G .*

Definition 5.1.2. *A G -circle packing is said to be locally univalent if for any vertex $v \in V$, the circles corresponding to v and to its neighbors in G have mutually disjoint interiors.*

We now state a few important results about circle packings. A detailed presentation of the subject can be found in Stephenson [45].

Theorem 5.1.1 ([45], p. 18). *Given a planar triangulation $G(V, E)$, and any assignment of positive radii to the boundary vertices of G , there exists (in the Euclidean and in the hyperbolic plane) an essentially unique locally univalent circle packing for G whose boundary circles have these values as their radii.*

Essentially unique is to be understood as up to isometry.

Definition 5.1.3. *A G -circle packing is said to be univalent if its circles have mutually disjoint interiors.*

In the sequel, we will need circle packings that are univalent. Thus, we will use the following result:

Theorem 5.1.2 ([45], p. 62). *Let G be a combinatorial closed disc (that is, simply connected, finite, triangulation). Then there exists an essentially unique univalent circle packing \mathcal{P}_G contained in the unit disc such that any boundary circle is internally tangent to the unit disc.*

We will refer to this kind of packing as a G -circle packing of the unit disc.

Note that the previous results are stated for a triangulated graph. However, these two theorems are still true for 3-connected planar graphs, if a rigidity condition is added to the definition of circle packing:

Definition 5.1.4. *Given a 3-connected planar graph $G(V, E)$, a G -circle packing is a set \mathcal{C} of circles in the plane with a bijection $\gamma : V \rightarrow \mathcal{C}$ such that $\gamma(v)$ and $\gamma(w)$ are externally tangent if and only if $\{v, w\}$ is an edge of G , and such that for each face $f = (w_1, \dots, w_n)$ of G , there exists a circle $c(f)$ which is orthogonal to all circles $\gamma(w_i)$, $1 \leq i \leq n$.*

5.1.2 Practical Computations of Circle Packings

Various methods exist for the computation of circle packings. The oldest and simplest one, which we will study in detail and build upon, is the Thurston algorithm [46]. It is an iterative algorithm which greedily updates the radii of the circles until they converge to values compatible with circle packing. Various other algorithms have surfaced since the inception of the original Thurston algorithm. Before presenting the details of the Thurston algorithm, we briefly describe two other algorithms relevant to our study.

The Springborn-Bobenko Algorithm

Springborn and Bobenko [5] have proposed a general framework for dealing with so-called *circle patterns*, which are sets of circles with non-zero intersection angles instead of the simpler tangency condition of circle packings. They characterize the intersection angles for which such circle patterns exist, and then define convex functionals on circle patterns which are minimized when the required conditions on these intersection angles are satisfied.

These ideas have been applied by Kharevych, Schroeder and Springborn [27] to the conformal parametrization of discrete 3D surfaces. They show how to apply the variational characterization of circle patterns of Springborn and Bobenko [5] to the practical computation of circle patterns with prescribed intersection angles.

Applying these methods to the special case of circle packings is easy. However, the minimization procedure is not directly amenable to distribution among network nodes.

Discrete Ricci Flow

Chow and Luo [15] have considered a completely different approach to the question of circle packing. They describe a discretization of Hamilton's Ricci flow and prove that it converges to a circle packing with prescribed adjacency relations. This implies an algorithm for computing circle packings, which is proved to converge exponentially fast.

While very efficient, this algorithm requires a periodic global rescaling of the circle radii, which prevents distribution of the computation among network nodes.

The Thurston Algorithm

In this section we present the algorithm that Thurston [46] designed for the numerical computation of circle packings.

The algorithm consists of setting the value of the boundary radii and updating all internal radii in order to satisfy local univalence. This step is repeated until some error bound on the local univalence error (measured as an angular error) is reached. At this point, a layout process is required to translate the radii values into planar coordinates of the centers. The convergence of this process to a locally univalent circle packing, in the Euclidean and hyperbolic case, is proved in [16]. See Collins and Stephenson [17] for a practical and efficient implementation of this algorithm.

Note that this algorithm works for triangulations only. However, it can be generalized to more general 3-connected planar graphs, with the additional constraint specified in Definition 5.1.4.

In the following, we represent the Thurston algorithm by a sequence of so-called *circle mapping functions* $(\phi_n)_{n \in \mathbb{N}}$ that map vertices of V to circles in the plane. The distance between two such functions is measured as the Euclidean distance d on $\mathbb{R}^{3|V|}$. We denote by Φ_G the function that maps the vertices to the limit circle packing Φ_G , which is unique up to some isometry of the hyperbolic plane, namely, some Möbius transformation.

There are two reasons why we focus on Thurston's algorithm: it is an extremely simple algorithm, and it can be distributed in a straightforward manner. However, there is one drawback in this algorithm, beyond its relative slowness. It provides only an *approximation* of the desired circle packing. Computing the exact one would require an infinite number of steps.

In the sequel, we show how to overcome this such that only a finite number of steps are required.

5.2 LOCAL TERMINATION CONDITIONS

In order to stop the iterations of Thurston algorithm, we need a termination condition that would guarantee that the result is at least a contained power diagram, with the correct adjacency relations. This is sufficient to enable greedy power routing. We need to ensure, however, that the algorithm may be distributed, including checking the termination condition.

5.2.1 Triangulated Case

Recall that we study a map $\phi : V \rightarrow D^2 \times \mathbb{R}$, which associates to each vertex v a point $p(v)$ in the unit disk and a radius $\sigma(v)$:

$$\phi = (p, \sigma).$$

The boundary of G is denoted by B .

Definition 5.2.1. *If w_1, \dots, w_n are the neighbors of v in G , the local cell of v in G , denoted by $\text{Cell}_G(v)$, is the cell of v in the power diagram of $\{\phi(v), \phi(w_1), \dots, \phi(w_n)\}$ (see Figure 5.1).*

In the following definition, when we refer to the order of vertices around another vertex, we mean the cyclic order of vertices, which is independent of the embedding in the case of a triangulation (except that we can reverse all orientations).

Definition 5.2.2. *For any vertex $v \in V$, we say that property $\text{LPD}(v, \phi)$ (Local Power Diagram) is satisfied if and only if*

- *if w_1, \dots, w_n are the neighbors of v in G (in this order), then the cell $\text{Cell}_G(v)$ contains $p(v)$ and the cells adjacent to it are exactly those of w_1, \dots, w_n (in this order, see Figure 5.2);*
- *Let $v \in B$. Denote by w_1 and w_n the two neighbors of v that belong to B and that are linked to v by boundary edges. Then in the power diagram of $\{\phi(v), \phi(w_1), \dots, \phi(w_n)\}$, $\text{Cell}(v) \cap \text{Cell}(w_1) \cap \text{Cell}(w_n)$ is either empty (which means that $\text{Cell}(v)$ is unbounded) or it is a point outside the unit disk \mathbb{D}^2 .*

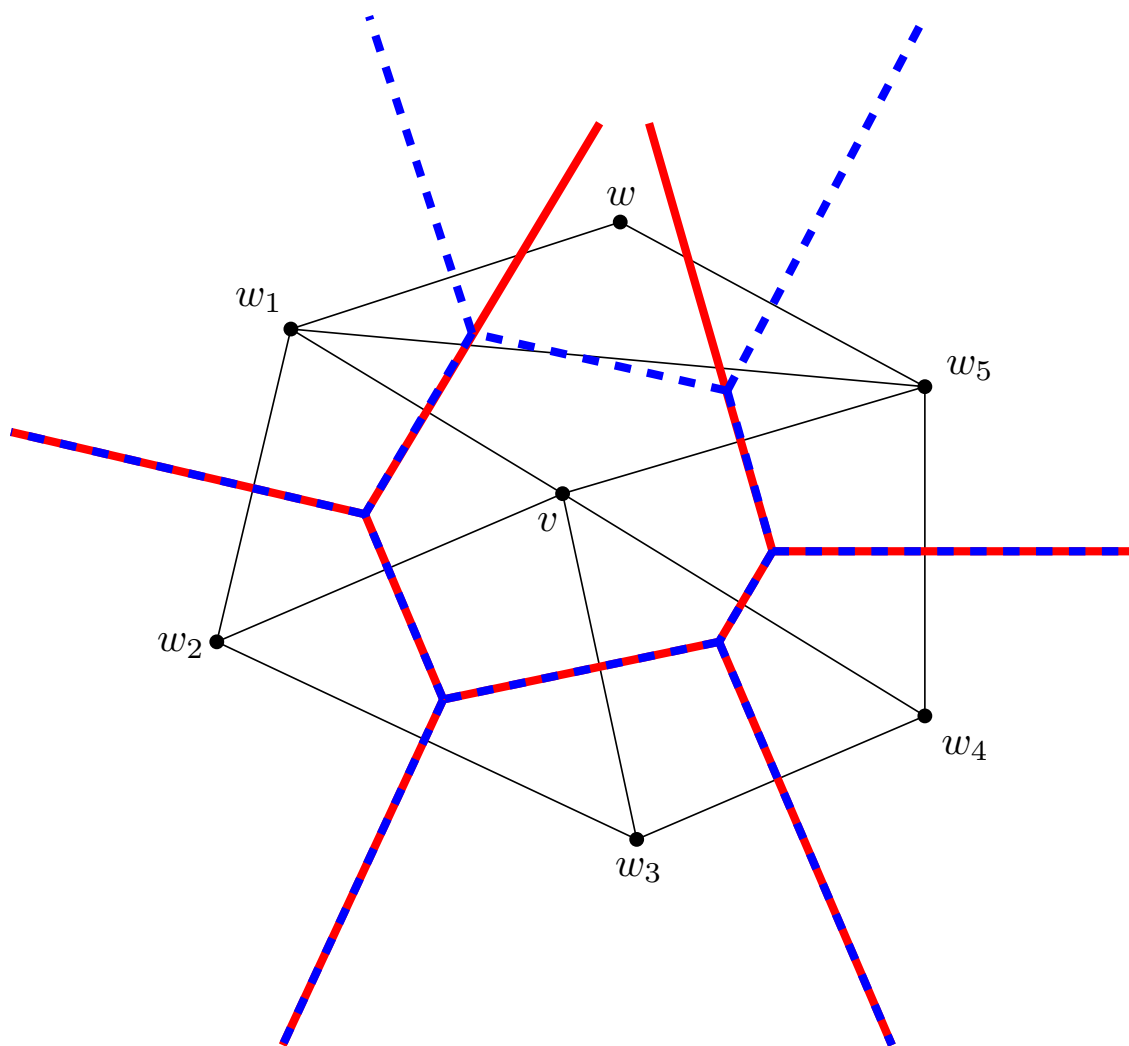


Figure 5.1: The local cell $\text{Cell}_G(v)$ (solid lines) contains the power diagram cell (dashed lines) and contains another vertex w .

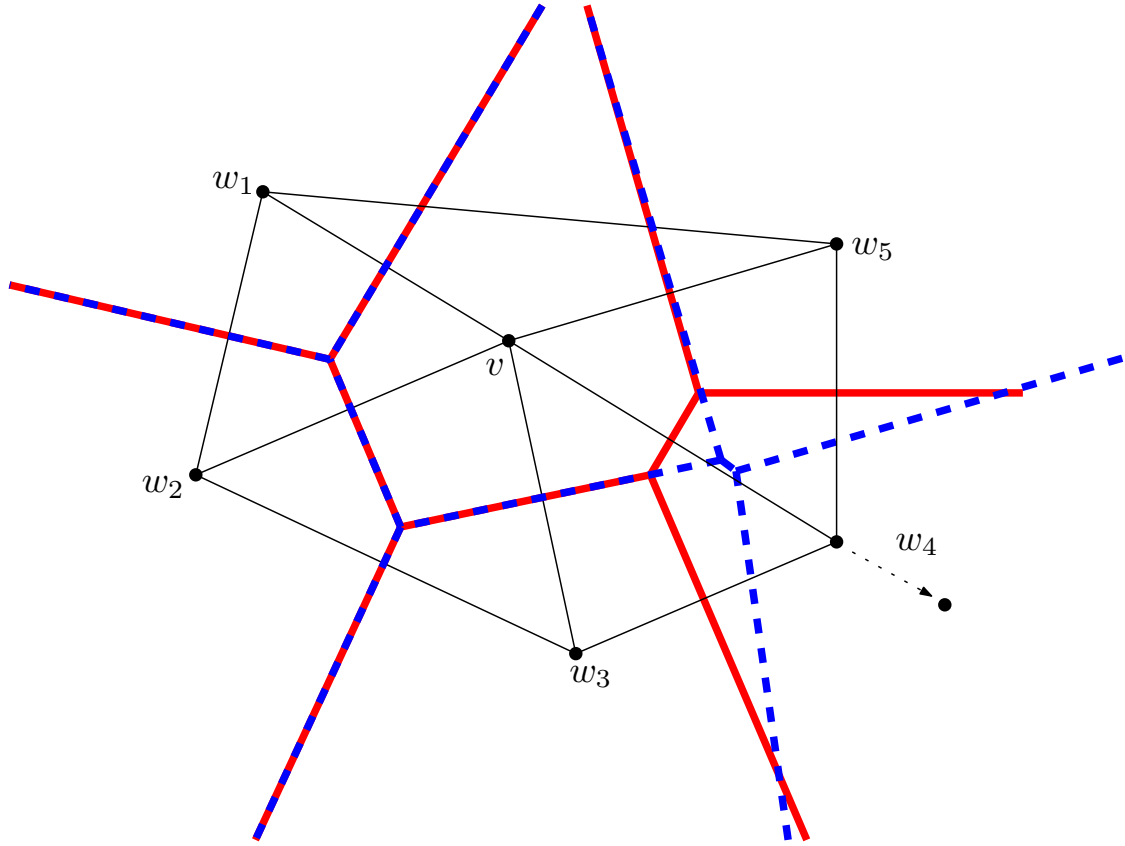


Figure 5.2: As w_4 moves away from v , $\text{LPD}(v, \phi)$ becomes unsatisfied (while the solid lines diagram becomes the dashed lines diagram), because $\text{Cell}_G(v)$ and $\text{Cell}_G(w_4)$ are not adjacent anymore, whereas edge $[vw_4]$ exists in G .

Note that the condition about the order of neighbor cells around a given cell is equivalent to requiring that the graph is properly embedded (this follows from the convexity of the power diagram cells). Thus, if G is known to be embedded, specifying the order of neighbor cells is not necessary.

We are now ready to state the central theorem of this section:

Theorem 5.2.1. *If*

$$\forall v \in V, \quad \text{LPD}(v, \phi),$$

then the restriction of the power diagram of $\phi(V)$ to the convex hull $\text{Conv}(p(V))$ is contained and its adjacency graph is G .

Proof. From now on, we denote by $\text{Cell}(w)$ the cell of $\phi(w)$ in the power diagram of $\phi(V)$, and by $\text{Cell}_G^v(w)$ the cell of w in the power diagram of $\{\phi(v), \phi(w_1), \dots, \phi(w_n)\}$, where w_1, \dots, w_n are the neighbors of v in G . Let ρ be the restriction to $\text{Conv}(p(V))$.

We now prove that $\rho(\text{Cell}_G^v(v)) = \rho(\text{Cell}(v))$ for all $v \in V$. First note that $\text{Cell}(v) \subset \text{Cell}_G^v(v)$ for all $v \in V$ implies that $\cup_{v \in V} \rho(\text{Cell}_G^v(v)) = \text{Conv}(p(V))$.

For each vertex $v \in V$, we consider the usual lifting to the polar hyperplane $\ell_v : x \mapsto (x, 2\langle x | \phi(v) \rangle - \|\phi(v)\|^2 + r(v)^2)$ in dimension 3. The power diagram of $\phi(V)$ is the projection of the upper envelope of the hyperplanes $\ell_v(\mathbb{R}^2)$. We now show that the $\ell_v(\rho(\text{Cell}_G^v(v)))$ can be glued into a convex terrain over the convex domain $\text{Conv}(p(V))$ (see Figure 5.3).

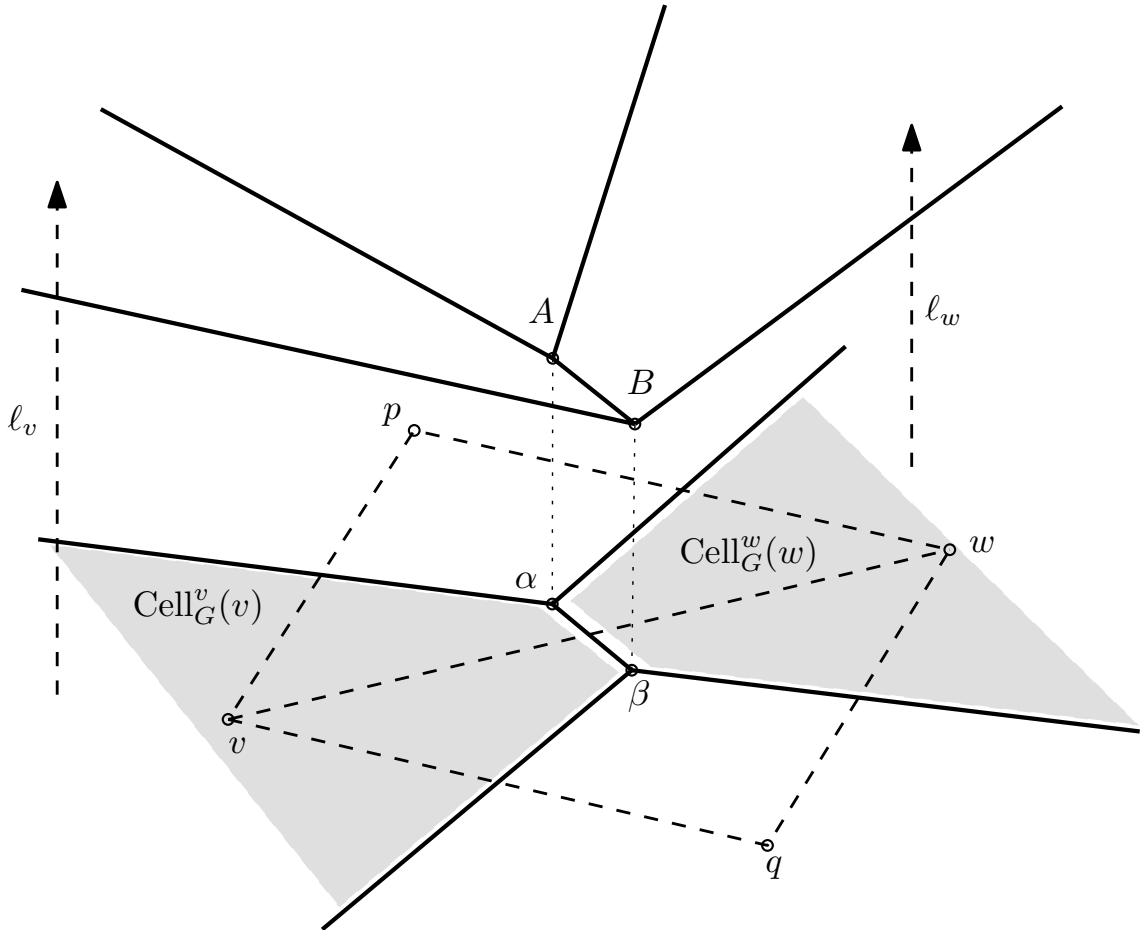


Figure 5.3: Lifting two local cells that share an edge.

If v and w are neighbors in G and $v \notin B$, let p and q be the two vertices opposite the edge (v, w) . Let α be the power diagram vertex defined by v, w and p , and let β be the power diagram vertex defined by v, w and q . The hypotheses $\text{LPD}(v, \phi)$ and $\text{LPD}(w, \phi)$ imply that the segment $[\alpha\beta]$ is an edge common to $\text{Cell}_G^v(v)$ and $\text{Cell}_G^w(w)$ because the four vertices v, w, p and q will all appear in the computations of the border of both cells.

This implies that $\ell_v(\text{Cell}_G^v(v))$ and $\ell_w(\text{Cell}_G^w(w))$ can be glued together along their common edge which is $[AB] = \ell_v([\alpha\beta]) = \ell_w([\alpha\beta])$. Furthermore, one can see that the angle between $\ell_v(\text{Cell}_G^v(v))$ and $\ell_w(\text{Cell}_G^w(w))$ along $[AB]$ is convex, because it is true for the local diagram of v and its neighbors,

Now consider the case where both v and w are boundary vertices. Let p be the vertex opposite (v, w) in G and consider the edge $e(v, w) = \text{Cell}_G^v(v) \cap \text{Cell}_G^w(w)$. Hypothesis $\text{LPD}(v, \phi)$ implies that this edge $e(v, w)$, whether infinite or not, has only one vertex inside the unit disk \mathbb{D}^2 , which is the power diagram vertex defined by v, w and p . $e(v, w)$ is also perpendicular to the line $(p(v)p(w))$ and reaches the boundary of \mathbb{D}^2 . By symmetry, $e(w, v)$ has the same properties. It follows that $\rho(e(v, w)) = \rho(e(w, v))$. This proves again that $\ell_v(\text{Cell}_G^v(v))$ and $\ell_w(\text{Cell}_G^w(w))$ can be glued together along this convex edge.

Finally, we obtain that the $\ell_v(\text{Cell}_G^v(v))$ can be glued together into a locally convex polyhedral terrain \mathcal{P} over the convex domain $\text{Conv}(p(V))$. It follows that \mathcal{P} is globally convex and is in fact the restriction of a convex polytope and that the projection of its edges onto $\text{Conv}(p(V))$ is a restricted power diagram, whose sites happen to be the elements of $\phi(V)$, by construction. The way the patches have been glued together shows that the adjacency graph of this restricted power diagram is exactly G .

Note that the sites on the boundary of G may not be in convex position. In particular, if the power diagram were not restricted to $\text{Conv}(p(V))$ (as we have seen in the proof, restricting to \mathbb{D}^2 is in fact sufficient), the cells of vertices which are not connected in G may be adjacent, creating an adjacency graph bigger than G . \square

We can now state the following corollary of Theorems 4.3.1 and 5.2.1:

Corollary 5.2.1.1. *If*

$$\forall v \in V, \quad \text{LPD}(v, \phi),$$

then greedy power routing delivers on ϕ . \square

5.2.2 Generalized Papadimitriou and Ratajczak Result

Papadimitriou and Ratajczak [37] provided geometric conditions on embeddings of 3-connected planar graphs which characterize greedy Euclidean embeddings. We now present this result in the more general context of arbitrary distance functions, and explain how it relates to Section 5.2.1. We will need this for the extension of the results of Section 5.2.1 to more general planar graphs.

Given a field d of distance functions $\{d_x : \mathbb{R}^2 \rightarrow \mathbb{R}, x \in \mathbb{R}^2\}$ (these functions are arbitrary real functions) and a set of sites $V \subset \mathbb{R}^2$, we can define two kinds of distance diagrams:

- the usual one, where the cell of a site v is defined as

$$\text{Cell}(v) = \{x \in \mathbb{R}^2, d_v(x) \leq d_w(x), \forall w \in V\}$$

- the reciprocal one, where the cell of a site v , called the *reciprocal cell* is defined as

$$\text{Cell}^\circ(v) = \{x \in \mathbb{R}^2, d_x(v) \leq d_x(w), \forall w \in V\}$$

Note that in the first case, the computation of a cell depends only on the distance functions of the sites. In contrast, in the second case, it depends on the distance functions at each point in the plane. Thus, the reciprocal diagram is usually impossible to compute (locally) if the distance functions are too general.

Just as we defined the local cell $\text{Cell}_G(v)$ of a vertex v of an embedded graph G , we can define the local *reciprocal cell* $\text{Cell}_G^\circ(v)$ and state a generalized version of the characterization of Papadimitriou and Ratajczak.

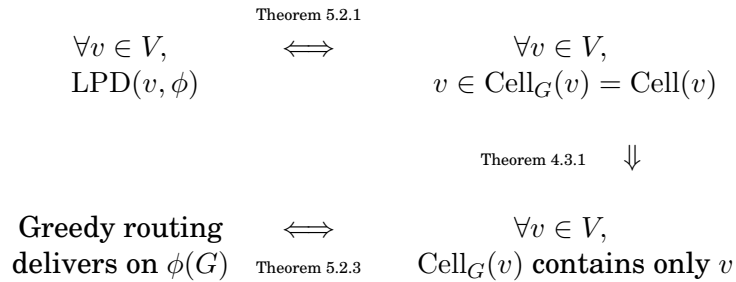
Theorem 5.2.2. *Given a field d of distance functions $\{d_x : \mathbb{R}^2 \rightarrow \mathbb{R}, x \in \mathbb{R}^2\}$, greedy routing on a graph $G(V, E)$ with respect to d delivers if and only if for each vertex $v \in V$, the local reciprocal cell $\text{Cell}_G^\circ(v)$ contains no vertex other than v .*

Proof. The proof is exactly the same as in Theorem 1 of Papadimitriou and Ratajczak [37]. \square

This is not a practical result. However, in the case of symmetrical distance functions, i.e. distance functions such that $\forall x, y \in \mathbb{R}^2, d_x(y) = d_y(x)$, the usual cell and the reciprocal cell are identical, namely $\text{Cell} = \text{Cell}^\circ$ and $\text{Cell}_G = \text{Cell}_G^\circ$. This is the case not only for the Euclidean distance, but also for the power distance: each point x in the plane is endowed with an arbitrary radius r_x , and the distance between two points x and y is defined as $d_x(y) = d_y(x) = \|x - y\|^2 - r_x^2 - r_y^2$ (if x is not a site, we may choose $r_x = 0$ or any arbitrary real value). Thus, we can now generalize Theorem 5.2.2:

Theorem 5.2.3. *Greedy power routing delivers if and only if for each vertex $v \in V$, the local cell $\text{Cell}_G(v)$ for the power distance contains no vertex other than $p(v)$ (see Figure 5.1). \square*

We summarize our results so far in the following diagram, which details the links between the various conditions. These hold for both Euclidean and power distances:



Note that the upper right condition may also be stated as “ G is the dual graph of the contained distance (power or Voronoi) diagram of $\phi(V)$.” Theorem 5.2.1 proves the left-to-right implication, and the right-to-left one is easy to check.

5.2.3 Non Triangulated Case

Let us now consider the more general case of a 3-connected planar graph. As in Section 5.2.1 for triangulated graphs, we present local sufficient conditions for greedy

power routing to deliver on general 3-connected planar graphs. The locality of the conditions is discussed in section 5.3.3.

In the previous section, we proved that satisfying LPD at every vertex implied that G is the adjacency graph of the power diagram of $\phi(V)$. This cannot be the case if G is not a triangulation: such a graph can only be the dual graph of a degenerate power diagram, which would be unstable under perturbation of the vertices, whereas LPD is stable.

In order to state the next definition, we need the following result:

Lemma 5.2.1. *If a set of points $\{p_1, \dots, p_n\}$ is in convex position, for any radii $(\sigma_i)_{1 \leq i \leq n}$, the adjacency graph of the power diagram of the circles $\mathcal{C}(p_i, \sigma_i)$ is a triangulation of $\text{Conv}(\{p_1, \dots, p_n\})$.*

Proof. The dual of a power diagram is known to be a (regular) triangulation. However, in order to have a triangulation of the convex hull $\text{Conv}(\{p_1, \dots, p_n\})$, each point p_i must be a vertex of this triangulation. In other words, it has to have a non-empty cell, which is guaranteed by the convexity assumption. \square

Definition 5.2.3. *If p is a convex embedding of G , the ϕ -triangulation of G is defined in the following way: if f is a non-triangle face, $p(f)$ is convex and we glue along f the dual graph of the power diagram of the vertices of f (which is indeed a triangulation of f , thanks to Lemma 5.2.1). The resulting triangulation of G is called the ϕ -triangulation of G and is denoted by $G(\phi)$ (see Figure 5.4).*

In case we are in a degenerate configuration, we choose a triangulation obtained after some infinitesimal perturbation.

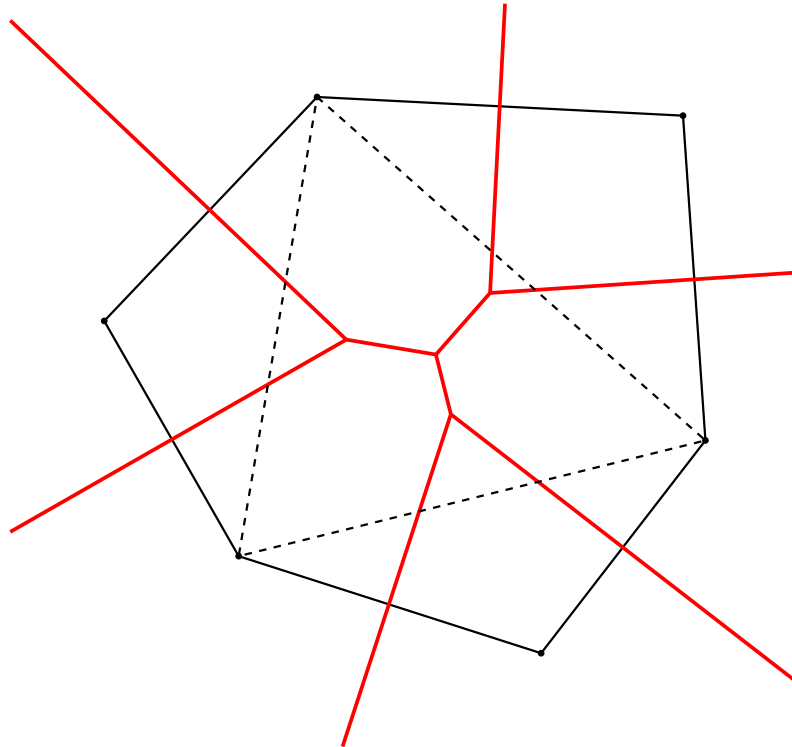


Figure 5.4: A face with 6 vertices embedded by ϕ with the regular triangulation of its vertices: G (solid lines) is triangulated into $G(\phi)$ (solid and dashed lines).

We are now able to present the generalized version of the condition that we proved sufficient in the triangulated case:

Definition 5.2.4. For any vertex $v \in V$, we say that property $\text{GLPD}(v, \phi)$ (*Generalized Local Power Diagram*) is satisfied if and only if the faces incident to v are convex, property $\text{LPD}(v, \phi)$ is satisfied in $G(\phi)$ and for each non-triangle face $f = (v, w_1, \dots, w_n)$ incident to v , the local cell $\text{Cell}_G(v)$ of v in G intersects f only along segments $[w_nv]$ and $[vw_1]$ (see Figure 5.5).

Note that, in the last condition, the local cell is computed in G , and not in $G(\phi)$ (see Definition 5.2.3): otherwise, the condition is trivially satisfied.

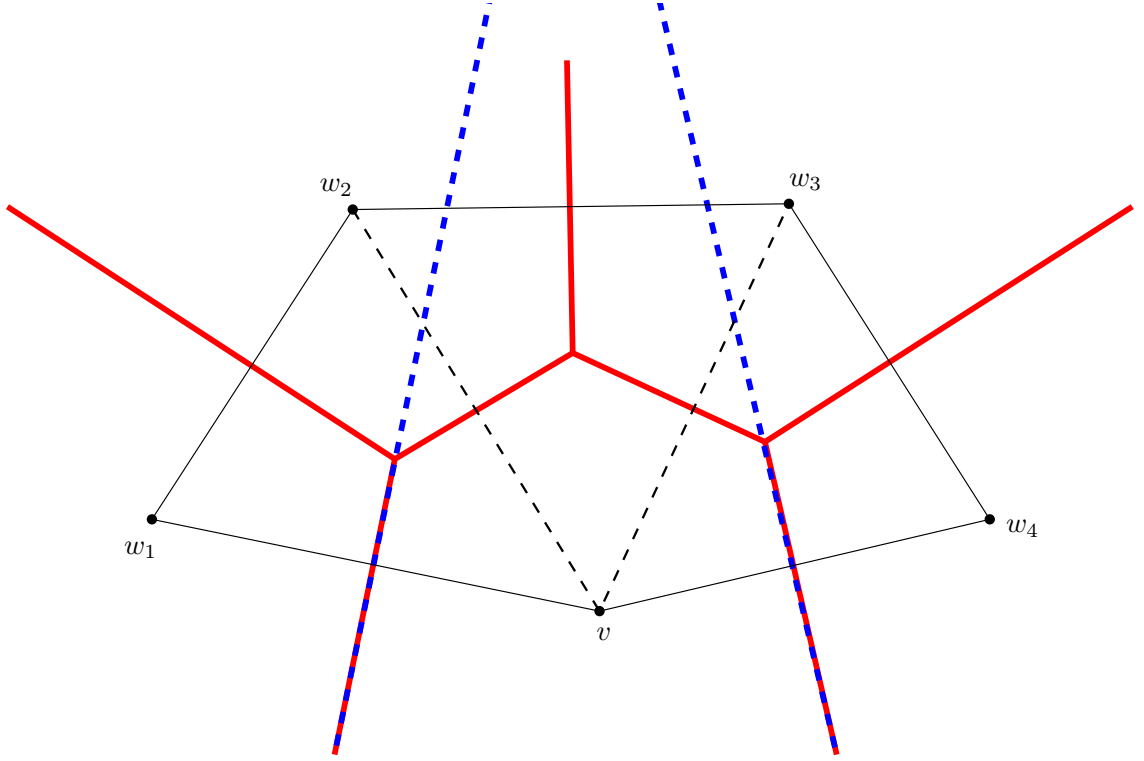


Figure 5.5: A face (solid edges) with 5 vertices, with $\text{GLPD}(v, \phi)$ not satisfied: the local cell of v (dashed lines) crosses the boundary of the face not only on $[w_1v]$ and $[vw_4]$ but also on $[w_2w_3]$, which is forbidden.

Theorem 5.2.4. If p is a convex embedding and

$$\forall v \in V, \quad \text{GLPD}(v, \phi),$$

then each local cell $\text{Cell}_G(v)$ contains only its site $p(v)$.

Proof. From the proof of Theorem 5.2.1, we know that $\text{LPD}(v, \phi)$ being satisfied for every vertex v implies that the local cell $\text{Cell}_{G(\phi)}(v)$ computed in $G(\phi)$ is exactly the cell of the power diagram of $\phi(V)$, and that this diagram is a contained embedding of $G(\phi)$.

We need the local cell $\text{Cell}_G(v)$ computed in G to be empty of other vertices. We know that $\text{Cell}_{G(\phi)}(v) \subset \text{Cell}_G(v)$. We now prove that the difference $\text{Cell}_G(v) \setminus$

$\text{Cell}_{G(\phi)}(v)$ is contained in the union of the faces incident to v . Note that $\text{Cell}_{G(\phi)}(v)$ is not itself contained in this union.

Let us consider now a non-triangle face $f = (v, w_1, \dots, w_n)$ incident to v . We denote by $W_f = \{w_{i_1}, \dots, w_{i_k}\}$ the set of vertices of f that belong to $W = N_{G(\phi)}(v) \setminus N_G(v)$. Denote by $\text{Cell}_f(v)$ the cell of v in the power diagram of $\{v\} \cup N_G(v) \cup W_f$.

By convexity of f , and using the fact the the local cells of the w_i are not allowed to cross f along the segments $[w_nv]$ and $[vw_1]$, one can easily see that $\text{Cell}_G(v) \setminus \text{Cell}_f(v)$ is contained in f . Since $\text{Cell}_{G(\phi)}(v) = \cap_f \text{Cell}_f(v)$, where the intersection is taken over all non-triangle faces f incident to v , the result follows. \square

One could wonder why we do not impose the stronger condition that triangle faces should satisfy the same property as non-triangle faces. The reason is that this condition is not equivalent to LPD in the triangulated case, whereas GLPD is. Since we want a condition as weak as possible, we avoid this.

The following corollary is a consequence of Theorems 5.2.3 and 5.2.4:

Corollary 5.2.3.1. *If the first component p of ϕ is a convex embedding and if*

$$\forall v \in V, \quad \text{GLPD}(v, \phi),$$

then greedy power routing delivers on ϕ . \square

5.2.4 Relation between Circle Packings, LPD and GLPD

The following theorems show that the conditions that we have described are indeed satisfied by the limits of the Thurston algorithm, namely circle packings.

Theorem 5.2.5. *If G is a planar triangulation and if $\phi(G)$ is a G -circle packing of the unit disc, then*

$$\forall v \in V, \quad \text{LPD}(v, \phi)$$

Proof. Since the bisector between two tangent circles is their common tangent line, the local cell of a circle is the intersection of the halfspaces delimited by some tangent lines. \square

Theorem 5.2.6. *If G is a 3-connected planar graph and if $\phi(G)$ is a G -circle packing of the unit disc, then*

$$\forall v \in V, \quad \text{GLPD}(v, \phi)$$

Proof. Let f be a face of G . By definition of the G -circle packing, there exists a circle $c(f)$ which is orthogonal to the circles of the vertices of f . It follows that c_f is inscribed in f , thus p is a convex embedding. We are in fact in the most degenerate case, and the faces can be triangulated arbitrarily to obtain a ϕ -triangulation of G . However, whichever triangulation we choose, the power diagram face of v is the polygon whose vertices are the centers of circles c_f , for the faces f incident to v . \square

5.3 ALGORITHMS

5.3.1 Computing a Greedy Power Embedding

We now derive from Sections 5.2.1 and 5.2.3 a distributed algorithm for the computation of a contained power diagram.

The algorithm consists simply of augmenting Thurston's iterative circle packing algorithm (see Section 5.1.2) with the conditions LPD (or GLPD) as termination conditions. Note that the Thurston algorithm itself has no concrete termination condition: it is an iterative process which is guaranteed to converge, and that in practice is run as many times as needed until some condition measuring convergence is met. Typically, some threshold on the angular error is used as a termination condition. However, it is not obvious that any such threshold on the angular error can guarantee that a contained power diagram is achieved.

The correctness of the algorithm follows from Section 5.2.4, since, in the worst case, the conditions LPD (or GLPD) will be satisfied when the algorithm converges to a circle packing, which is guaranteed. We now describe the algorithm and discuss its correctness.

5.3.2 Termination

Our algorithm consists of running the Thurston algorithm to compute a circle packing in the Poincaré model of the hyperbolic plane, initialized with infinite radii for all boundary circles. This amounts to requiring that the boundary circles are internally tangent to the unit circle. Theorem 5.1.1 implies that the locally univalent circle packing that we would obtain upon convergence is essentially unique. Since Theorem 5.1.2 states that there exists a univalent circle packing satisfying such boundary conditions, we know that the circle packing the algorithm is converging to is not only locally univalent, but also globally univalent.

We stop the Thurston algorithm as soon as the LPD condition is satisfied (or the GLPD condition, in case the graph is not a triangulation but a general 3-connected planar graph).

More precisely, the steps of the algorithm are as follows (with some integer parameter $N > 0$):

- 1) set all boundary radii to infinity and all internal radii to 1;
- 2) update all internal radii by applying N steps of Thurston's algorithm in the hyperbolic plane;
- 3) fix the positions of two neighbor disks and sweep the network to compute the Euclidean layout ϕ of the circles in the Poincaré unit disk representation of the hyperbolic plane;
- 4) if $\text{LPD}(v, \phi)$ (or $\text{GLPD}(v, \phi)$ in the non triangulated case) is not satisfied for some v , go to step 2. Otherwise, return the current layout.

Note that in the non-triangulated case, steps 2, 3 and 4 will require the network to emulate a triangulation of the graph. Additionally, the network has to be able to detect the state at which $\text{LPD}(v, \phi)$ is satisfied at *all* nodes (Step 4), at which point the algorithm terminates. This is complicated by the fact that $\text{LPD}(v, \phi)$ being satisfied does not imply that it will continue to be satisfied at subsequent iterations (because of the activity at neighboring nodes). However, the following lemma proves that ultimately the algorithm will converge, namely, reach a state in which $\text{LPD}(v, \phi)$

is satisfied for all v . This state may then be detected by standard distributed algorithmic techniques.

Lemma 5.3.1. *Conditions LPD and GLPD are open conditions in the neighborhood of circle packings in the sense that for all G and limit circle packing Φ_G , there exists a distance $\epsilon > 0$ such that for all circle mapping functions ϕ , we have $d(\phi, \Phi_G) < \epsilon \Rightarrow \forall v \in V, \text{LPD}(v, \phi)$ if G is a triangulation, and $d(\phi, \Phi_G) < \epsilon \Rightarrow \forall v \in V, \text{GLPD}(v, \phi)$ if G is a 3-connected planar graph.*

Proof. Using Theorems 5.2.5 and 5.2.6, it suffices to observe that, in the case of circle packings, two neighboring circles have a common power diagram edge of positive length, and that the corresponding embedding of the centers is always strictly convex, i.e. all the faces of the embedding are strictly convex. \square

5.3.3 Locality

Let us now examine the locality of the computations involved in the algorithm. In the triangulated case, each node of the triangulation needs to know the radii associated with its neighbors in order to update its own radius. This is the most local level of communication possible. We call it *G-locality*. In the case of a non-triangle 3-connected planar graph, each vertex needs to know the radii of the vertices it shares a *face* with. This level of communication, which is less local, is called *G-face-locality*.

The algorithm generates a set of radii, but in order to check the LPD or GLPD conditions, we need an actual embedding of the node and its neighbors. Such a layout of circles may be obtained by positioning the circles in a breadth-first order: once two neighbor vertices have their positions set, all other positions can be computed in this order. As for the computation of radii, this step is *G-local* in the case of a triangular graph, but *G-face-local* in the case of 3-connected planar graphs. Similarly, one can see that checking LPD is *G-local*, whereas checking GLPD is *G-face-local*.

5.4 DISCUSSION

5.4.1 Experimental Validation

We have implemented a simulation of the algorithm of Section 5.3 in MATLAB and tested it on random triangular graphs and 3-connected planar graphs containing around 50 vertices each, generated by E. Fusy's software [26]. We obtained greedy power embeddings after a few hundred iterations (in general, less than 100 for triangulations, and between 100 and 500 for general 3-connected graphs). If we define an exact packing as a circle packing such that circles which should be tangent are indeed tangent, with an error on the distance between their centers within 1% of the smallest of the two radii, we can compare the number of iterations required to obtain a greedy power embedding with the number of iterations needed to obtain an exact packing: in the case of triangle graphs, we needed, on the average, a factor of 3.8 less iterations. In the case of general 3-connected planar graphs, we needed, on the average, a factor of 1.8 less iterations. Figures 5.6, 5.7, 5.8, and 5.9 show two intermediate steps, the greedy power embedding and the exact packing generated for the same input graph.

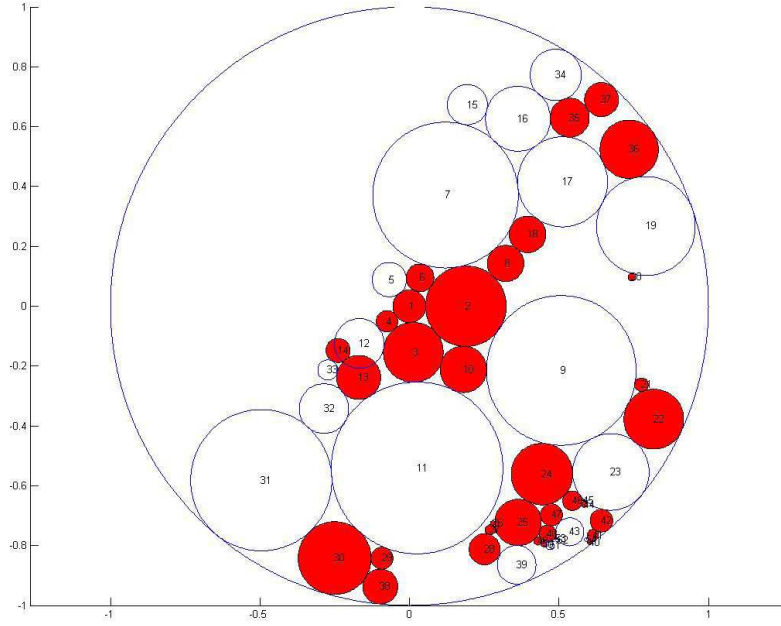


Figure 5.6: After 6 iterations, the colored circles are the ones that already satisfy LPD.

Note that the high non-uniformity of these random graphs, i.e. a short loop of edges may bound a region containing a large number of vertices (i.e. the graph contains small cuts), is a reason for the relatively low efficiency of the algorithm. This kind of setting is not realistic in the case of sensor networks, where one would expect the planar graph to be a subgraph of a realistic communication graph such as a unit disk graph.

We did not implement the heuristic acceleration schemes proposed by Collins and Stephenson [17] because these heuristics rely on the global evaluation of the so-called *error reduction factor*. It would however be interesting to check whether a much more local evaluation of this parameter could still speed up the process significantly.

5.4.2 Possible Improvements

We have described a modification of the Thurston algorithm originally designed for generating circle packings, so that it is able to generate the embeddings required to support greedy power routing on a sensor network. The algorithm is simple and *Gface-local*, thus may easily be implemented in a distributed manner on the sensor network. However, our algorithm is not practical in case the domain contains big holes, which would function as large non-triangulated faces. A natural way of dealing with this problem would be to analyze the topology of the underlying domain and split it into simply connected parts which could be treated separately (see [25]).

Our current implementation uses a breadth-first traversal to locally compute the position of a vertex at each iteration once the radii have been adjusted. This involves simple and local computations, but may accumulate error in large networks. An optimized layout process that would spread the error evenly among the vertices could improve our results by triggering the termination conditions earlier. One way to do

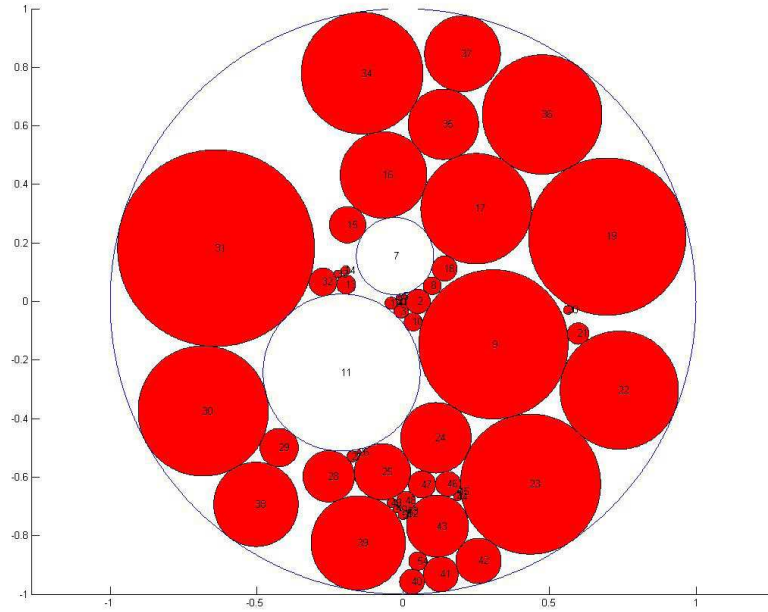


Figure 5.7: After 29 iterations, only 2 circles still do not satisfy LPD.

this is using the triangle layout method of ABF++ (Angle Based Flattening) [40], which involves solving a linear system for the vertex coordinates. Since this type of computation may be distributed among the vertices, it is a promising direction for future research. Alternatively, it might be possible to devise a way of checking LPD or GLPD from the radii only, without explicitly computing the vertex positions.

Most algorithms for greedy routing rely on the input being a planar 3-connected graph, which is not very realistic. The simplest remedy is to extract a spanning subgraph of this type from a more general input and embed this. It is easy to see that adding back the non-planar edges after the embedding process does not harm the greediness of the embedding. However, extracting such a subgraph is in itself a difficult problem. Thus an important problem is to devise a greedy embedding algorithm for general graphs.

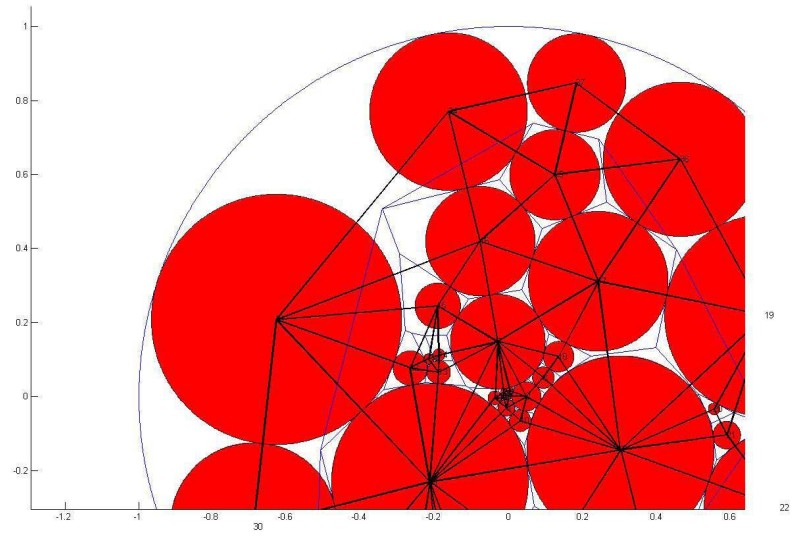


Figure 5.8: After 32 iterations, LPD is satisfied everywhere: the embedding is a greedy power embedding.

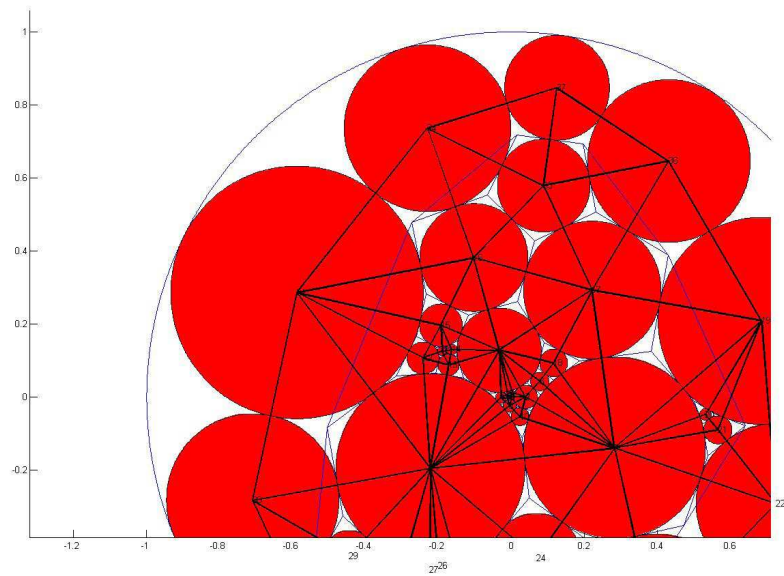


Figure 5.9: After 128 iterations, the embedding is a circle packing.

Conclusion and perspectives

This thesis presents both theoretical studies of Voronoi diagrams and some of their applications. These applications have been developed by extending existing methods, such as the linearization methods, into algorithms, but also by developing new frameworks, such as the locally uniform mesh framework, and the local power diagram property for greedy routing.

The importance of local geometric properties appears strongly in this work. In each case, we made use of local properties of diagrams or triangulations, and showed that satisfying these properties would solve our problem. Then, the algorithmic part of the work consisted in developing algorithms that would allow computing geometric structures with these properties. For mesh generation applications, the locality of the properties was more of a tool than a goal. In the case of locally uniform anisotropic meshes, locality helped us to design a meshing algorithm using refinement, and with the property that the current state was independent from the order of insertion of the points, a property which is not easily obtained when dealing with anisotropic meshes, and which proved useful for proving the correctness of the algorithm. Interestingly, in the greedy routing case, locality was a strong goal, rather than a tool, because our aim was to provide a distributed method for computing an embedding, which implied that the quality of such an embedding had to be checked locally.

The proofs of correctness of the use of the LPD property are important in that they show an unintuitive link between local and global containment in affine diagrams. This link can be compared to the link between local and global convexity of polyhedra. In our view, studying further local geometric properties of geometric structures such as diagrams and triangulations is an interesting and promising area of research. A natural extension of this work consists in using it in the context of search datastructure, and more specifically peer-to-peer networks.

The interest of the locally uniform meshing is different: it consists in developing a simple framework, apparently too simple for the task, and showing that this framework is still sufficient while being quite generic. This simplicity and genericity leaves a lot of space for improvements, which would be much more complex to develop in more classical ad hoc frameworks. For example, we considered locally uniform meshes, but one can imagine doing the same kind of constructions with a higher order of approximation of the metric. This local framework makes the issues to be solved in order to design such extensions rather independent and well defined. Another direction for extension consists in using this local stars framework for cases where computing in an embedding space would be too costly: maintaining the local stars means working mostly on the intrinsic geometry of the considered object.

Furthermore, it is important to note that most of the proved algorithms for meshing have much better practical behaviors than what can be expected from the bounds proved theoretically. Understanding better these behaviors and modeling the real quality achievable by the known algorithms are important questions that are mostly open, and that would allow better designs in the future. Developing methods for such analyses should not rely solely on statistical methods. Beyond statistics, there is much room for improvements in terms of geometric understanding. In our view, the design of the off-centers methods [47] is a typical example of the kind of improvements in the geometric understanding of meshing tools that are needed for advancing further the study of practical complexity, from a theoretical viewpoint.

Bibliography

- [1] T. Apel, S. Grosman, P.K. Jimack, and A. Meyer. A new methodology for anisotropic mesh refinement based upon error gradients. *Applied Numerical Mathematics*, 50(3-4):329–341, 2004.
- [2] D. Attali, J.-D. Boissonnat, and A. Lieutier. Complexity of the Delaunay triangulation of points on surfaces the smooth case. In *Proc. 19th Ann. Symposium on Computational Geometry*, pages 201–210, San Diego, 2003. ACM Press.
- [3] Dominique Attali and Jean-Daniel Boissonnat. A linear bound on the complexity of the delaunay triangulation of points on polyhedral surfaces. *Discrete and Comp. Geometry*, 31:369–384, 2004.
- [4] F. Aurenhammer. Power diagrams: properties, algorithms and applications. *SIAM J. Comput.*, 16:78–96, 1987.
- [5] A. Bobenko and B. Springborn. Variational principles for circle patterns and Koebe’s theorem. *Trans. Amer. Math. Soc.*, 356:659–689, 2004.
- [6] Jean-Daniel Boissonnat and Menelaos Karavelas. On the combinatorial complexity of Euclidean Voronoi cells and convex hulls of d -dimensional spheres. In *Proc. 14th ACM-SIAM Sympos. Discrete Algorithms (SODA)*, pages 305–312, 2003.
- [7] Jean-Daniel Boissonnat, Camille Wormser, and Mariette Yvinec. Anisotropic diagrams: Labelle shewchuk approach revisited. *Theoretical Computer Science*, to appear.
- [8] Houman Borouchaki, Paul Louis George, Frédéric Hecht, Patrick Laug, and Eric Saltel. Delaunay mesh generation governed by metric specifications. part i algorithms. *Finite Elem. Anal. Des.*, 25(1-2):61–83, 1997.
- [9] P. Bose, P. Morin, I. Stojmenovic, and J. Urrutia. Routing with guaranteed delivery in ad hoc wireless networks. *Wireless Networks*, 7(6):609–616, 2001.
- [10] Frank Bossen and Paul Heckbert. A pliant method for anisotropic mesh generation. In *5th International Meshing Roundtable*, October 1996.
- [11] Bernard Chazelle. An optimal convex hull algorithm in any fixed dimension. *Discrete Comput. Geom.*, 10:377–409, 1993.
- [12] M. Ben Chen, C. Gotsman, and S. J. Gortler. Routing with guaranteed delivery on virtual coordinates. In *Proc. Canadian Conf. on Comp. Geom.*, 2006.

- [13] S.W. Cheng, T.K. Dey, E.A. Ramos, and R. Wenger. Anisotropic surface meshing. *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithms*, pages 202–211, 2006.
- [14] L. Paul Chew. Guaranteed-Quality Delaunay Meshing in 3D. In *Proceedings of the Thirteenth Annual Symposium on Computational Geometry*, pages 391–393, Nice, France, June 1997. Association for Computing Machinery.
- [15] B. Chow and F. Luo. Combinatorial Ricci flows on surfaces.
- [16] Y. Colin de Verdière. Empilements de cercles: Convergence d’une méthode de point fixe. In *Forum Mathematicum*, volume 1, pages 395–402, 1989.
- [17] C. Collins and K. Stephenson. A circle packing algorithm. *Computational Geometry: Theory and Applications*, 25:233–256, 2003.
- [18] H. S. M. Coxeter. *Introduction to Geometry*. John Wiley & Sons, New York, 2nd edition, 1969.
- [19] R. Dhandapani. Greedy drawings of triangulations. In *Accepted to SODA*, 2008.
- [20] C. Dobrzynski, P.J. Frey, B. Mohammadi, and O. Pironneau. Fast and accurate simulations of air-cooled structures. *Computer Methods in Applied Mechanics and Engineering*, 195(23-24):3168–3180, 2006.
- [21] R. A. Dwyer. Higher-dimensional Voronoi diagrams in linear expected time. *Discrete Comput. Geom.*, 6:343–367, 1991.
- [22] H. Edelsbrunner and N. R. Shah. Triangulating topological spaces. *Int. J. on Comp. Geom.*, 7:365–378, 1997.
- [23] Jeff Erickson. Dense point sets have sparse delaunay triangulations. *Discrete Comput. Geom.*, 33:85–115, 2005.
- [24] PJ Frey and F. Alauzet. Anisotropic mesh adaptation for CFD computations. *Computer Methods in Applied Mechanics and Engineering*, 194(48):5068–5082, 2005.
- [25] S. Funke and N. Milosavljevic. Network sketching or: “How much geometry hides in connectivity? – part II”. In *Proc. ACM-SIAM Symp. on Discrete Algorithms*, 2007.
- [26] É. Fusy. Quadratic exact-size and linear approximate-size random sampling of planar graphs. In *Proc. Analysis of Algorithms*, 2005.
- [27] L. Kharevych, P. Schroeder, and B. Springborn. Discrete conformal mappings via circle patterns. *ACM Transactions on Graphics*, 2006.
- [28] R. Klein. Abstract Voronoi Diagrams and their Applications. *Proceedings of the International Workshop on Computational Geometry on Computational Geometry and its Applications*, pages 148–157, 1988.
- [29] Rolf Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes Comput. Sci.* Springer-Verlag, 1989.

- [30] R. Kleinberg. Geographic routing using hyperbolic space. In *Proceedings of the 26th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, 2007.
- [31] P. Koebe. Kontaktprobleme der konformen Abbildung. *Ber. Sächs. Akad. Wiss. Leipzig, Math.-Phys. Kl.*, 88:141–164, 1936.
- [32] E. Kranakis, H. Singh, and J. Urrutia. Compass routing on geometric networks. In *Proc. Canadian Conf. on Comp. Geom.*, pages 51–54, 1999.
- [33] F. Labelle and J. Shewchuk. Anisotropic voronoi diagrams and guaranteed-quality anisotropic mesh generation. In *Proc. 19th Ann. Symposium on Computational Geometry*, pages 191–200. ACM Press, 2003.
- [34] Xiang-Yang Li and Shang-Hua Teng. Generating well-shaped delaunay meshed in 3d. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 28–37, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics.
- [35] Xiang-Yang Li, Shang-Hua Teng, and Alper Üngör. Biting ellipses to generate anisotropic meshes. In *8th International Meshing Roundtable*, October 1999.
- [36] J.M. Mirebeau and A. Cohen. Greedy bisection generates optimally adapted triangulations. 2008.
- [37] C. Papadimitriou and D. Ratajczak. On a conjecture related to geometric routing. In *Proc. ALGOSENSORS*, pages 9–17, 2004.
- [38] Micha Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete Comput. Geom.*, 12:327–345, 1994.
- [39] Micha Sharir and P. K. Agarwal. *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, New York, 1995.
- [40] A. Sheffer, B. Lévy, M. Mogilnitsky, and A. Bogomyakov. Abf++ : Fast and robust angle based flattening. *ACM Transactions on Graphics*, 24(2):311–300, 2005.
- [41] Jonathan Richard Shewchuk. Tetrahedral mesh generation by delaunay refinement. In *SCG '98: Proceedings of the fourteenth annual symposium on Computational geometry*, pages 86–95, New York, NY, USA, 1998. ACM.
- [42] Jonathan Richard Shewchuk. What is a good linear finite elements interpolation, conditioning, anisotropy, and quality measures. Manuscript 2002.
- [43] Richard Shewchuk. Star splaying: an algorithm for repairing delaunay triangulations and convex hulls. In *SCG '05: Proceedings of the twenty-first annual symposium on Computational geometry*, pages 237–246, New York, NY, USA, 2005. ACM.
- [44] RB Simpson. Anisotropic mesh transformations and optimal error control. *Proceedings of the third ARO workshop on Adaptive methods for partial differential equations*, pages 183–198, 1994.
- [45] K. Stephenson. *Introduction to Circle Packing. The Theory of Discrete Analytic Functions*. Cambridge University Press, April 2005.

- [46] W. Thurston. The finite riemann mapping theorem. In *Invited talk, An International Symposium at Purdue University on the occasion of the proof of the Bieberbach conjecture*, March 1985.
- [47] A. Ungor. Off-Centers: A New Type of Steiner Points for Computing Size-Optimal Quality-Guaranteed Delaunay Triangulations. In *Latin 2004: Theoretical Informatics: 6th Latin American Symposium, Buenos Aires, Argentina, April 5-8, 2004: Proceedings*. Springer, 2004.
- [48] S. Yamakawa and K. Shimada. High Quality Anisotropic Tetrahedral Mesh Generation via Ellipsoidal Bubble Packing. *Proceedings of the 9th International Meshing Roundtable*, pages 263–273, 2000.