Eurographics Symposium on Geometry Processing (2005)
M. Desbrun, H. Pottmann (Editors)

Surface Reconstruction from Noisy Point Clouds

Boris Mederos, Nina Amenta, Luiz Velho and Luiz Henrique de Figueiredo

University of California at Davis
IMPA - Instituto Nacional de Matematica Pura e Aplicada

Abstract

We show that a simple modification of the power crust algorithm for surface reconstruction produces correct outputs in presence of noise. This is proved using a fairly realistic noise model. Our theoretical results are related to the problem of computing a stable subset of the medial axis. We demostrate the effectiveness of our algorithm with a number of experimental results.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Surface Reconstruction, Medial Axis, Noisy samples

1. Introduction

Surface reconstruction is an important problem in geometric modeling. It has received a lot of attention in the computer graphics community in recent years because of the development of laser scanner technology and its wide applications in areas such as reverse engineering, product design, medical appliance design and archeology, among others.

Different approaches have been taken to the problem, including the work of Hoppe, DeRose et al which popularized laser range scanning as a graphics tool [HDD*92], the rolling ball technique of Bernardini et al [BMR*99], the volumetric approach of Curless et al [CL96] used in the Digital Michelangelo project [LPC*00], and the radial basis function method of Beatson et al. [CBC*01].

The algorithms [ABK98, ACDL00, BC02, ACK01a] uses the Voronoi diagram of the input set of point samples to produce a polyhedral output surface. A fair amount of theory was developed along with these algorithms, which was used to provide guarantees on the quality of the output under the assumption that the input sampling is everywhere sufficiently dense. The theory relates surface reconstruction to the problem of medial axis estimation in interesting ways, and shows that the Voronoi diagram and Delaunay triangulation of a point set sampled from a two-dimensional surface have various special properties. Some strengths of the sampling model used are that the required sampling density can vary over the surface with the local level of detail, and that over-sampling, in arbitrary ways, is allowed. One drawback is that it assumes that the sample is free of noise.

When noise is considered as well, the quality of the output is related to both the density and to the noise level of the sample. A small number of recent results have begun to explore the space of what it is possible to prove under various noisy sampling assumptions. Dey and Goswami [DG04] proposed an algorithm for which they could provide many of the usual theoretical guarantees, using a model in which both the sampling density and the noise level can vary with the local level of detail, but which gives up the arbitrary over-sampling property. A real noisy input, however, might well have arbitrary over-sampling but the sampling density and noise level usually varies unpredictably, independent of the local level of detail.

In this paper, we show that similar results can be achieved given bounds on the minimum sampling density and maximum noise level, but allowing arbitrary over-sampling.

Related Work

Most of the algorithms using the Voronoi diagram and Delaunay triangulation of the samples, for which a variety of theoretical guarantees can be provided, require the input to be noise-free [AB99, ACDL00, ACK01b, BC02]. In practice some of these algorithms are more sensitive to noise than others. The recent algorithm of Dey and Goswami [DG04] extends much of the theory developed in the noise-free case

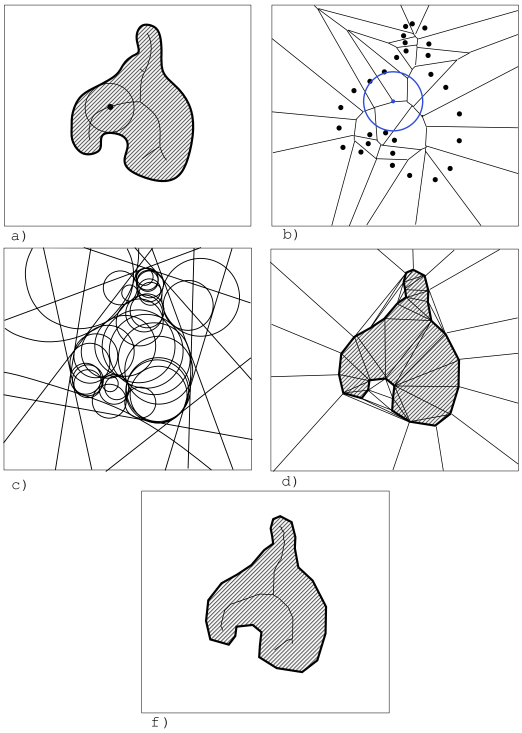

Figure 1: *A two dimensional example of the power crust algorithm. a) An object and its medial axis. b) The voronoi diagram and its poles, the blue points corresponding to poles and the circles corresponding to polar balls. c) The set of inner and outer polar balls. d) The power diagram of the set of polar balls. The algorithms labels the cells of this power diagram inner or outer. e) The set of faces in the power diagram which separate inner from outer cells.*

to inputs with noise. We do the same with a less restrictive sampling model, as described in more detail in Section 2.2.

Both our algorithm and that of Dey and Goswami are extensions of the *power crust* algorithm proposed by Amenta, Choi and Kolluri [ACK01b]. This algorithm is illustrated in Figure 1. Given an input sample P of points on a surface S, it selects from the Voronoi diagram of P a set V of Voronoi vertices, the *poles*, which approximate the medial axis transform of S. It then uses the *power diagram* (a kind of weighted Voronoi diagram) of the set of Delauanay balls centered at V (the *polar balls*) to recover a polyhedral surface representation.

Voronoi-based surface reconstruction techniques in general are closely related to Voronoi-based algorithms for medial axis estimation (in fact the power crust code is probably more often used for the latter problem). Yet another noisy sampling model was used by Chazal and Lieutier [CL05] in a recent paper on medial axis estimation: their sampling requirement is simply that the Hausdorff distance between the point sample and the surface itself is bounded by some constant r. Notice that this allows for arbitrary over-sampling, but does not allow the sampling density to vary over the surface according to the local level of detail. Chazal and Lieutier proved, drawing on more general results, that a subset of the Voronoi diagram of P approaches a subset of the medial axis of S as $r \rightarrow 0$, and that both converge to the entire medial axis. It is tempting to apply Chazal and Lieutier's result directly to the surface reconstruction problem, by using the power crust approach to produce a polyhedral surface from their approximate medial axis. But this is not as straightforward as it might seem: their medial axis estimation includes Voronoi edges and two-faces as well as vertices, while the analysis of the power crust relays on having an approximation of the medial axis by Voronoi vertices. Also, the subset of the medial axis approximated by Chazal and Lieuter is not guaranteed to be homotopy equivalent to the complete medial axis, or to the object, since the sampling is not required to be dense enough to capture the smallest topological feature.

Recently similar techniques have been used to analyze a particular smooth surface determined by a noisy sets of samples [Kol05], a variant of the MLS surface definition of Levin [Lev03]. In this case arbitrary over-sampling seems to be ruled out, since the surface locally averages the input samples and malicious over-sampling could influence the local averages. There is also a recent algorithm for curve reconstruction from a noisy sample [CFG*03] with theoretical guarantees, for which the sampling model has the interesting property that the quality of the output improves with increased sampling density, even when the noise level remains constant. The sampling model used is not particularly realistic, but the property seems quite relevant to practice.

2. Geometric Definitions and Sampling Assumptions

2.1. Defnitions and Notation

We will use the following notation. For any set $X \subset \mathbb{R}^3$, $\overset{o}{X}$, X^c and ∂X denote respectively the interior of X, the complement of X and the boundary of X. Given a point x and a set Y we denote by $d(x,Y) = \inf_{y \in Y} d(x,y)$. Given any two set X and Y we denote by $\tilde{d}_H(X,Y) = \sup_{x \in X} d(x,Y)$ the one-sided Hausdorff distance from X to Y and by $d_H(X,Y) = \max\{\tilde{d}_H(X,Y), \tilde{d}_H(Y,X)\}$ the Hausdorff distance between X and Y. We denote by $B_{c,\rho}$ a ball with center c and radius ρ.

We will consider two-dimensional, compact, and C^2 manifolds without boundary, and we will call such a manifold a *smooth surface*. Let S be a smooth surface. We will assume that S is contained in an open, bounded domain Ω (eg, a big open ball). The surface S divides Ω into two open solids, the inside (inner region) and the outside (outer region) of S, which are disconnected.

The *medial axis* M of a surface S is the closure of the set of points in Ω that have at least two distinct nearest points on S. Note that the set M is divided into two parts, the inner and outer medial axis, belonging to the inner or outer region of the surface S, respectively. The ball B_{m,ρ_m} centered at a medial axis point m with radius $\rho_m = d(m,S)$ will be called a *medial ball*. It is easy to see that a medial ball is maximal in the sense that there is no ball B with $\overset{o}{B} \cap S = \emptyset$ which contains B_{m,ρ_m}.

The medial axis M is a bounded set, since in our definition it is contained in the bounded domain Ω. So there exits an upper bound Δ_0 for the radius of the medial balls.

2.2. Sampling and Noise Models

There are at least two good approaches to defining sampling and noise models. First, we can begin with a model which we believe roughly describes the characteristics of reasonable input data sets, and then show that our algorithm works correctly on data that fits the model. The second approach would be to begin with the algorithm, and describe the data sets for which the algorithm is correct as broadly as possible, and then argue that this broad class of possible inputs includes reasonable input data sets (possibly among others). This is the approach taken in the analysis of many of the Voronoi-based surface reconstruction algorithms, as follows.

For a point $x \in S$, we define $\text{lfs}(x) = d(x,M)$. This lfs function is used to determine the required sampling density; it is small in regions of high curvature or where two patches of surface pass close together, and larger away from such regions of fine detail.

A finite set of points P is a *r-sample* of the surface S if $P \subset S$ and if for any $x \in S$ there is a point $p \in P$ with $d(x,p) \leq r\,\text{lfs}(x)$.

The points of a noisy sample P for S lie near but not on the surface. Let $\tilde{P}$ be the projection of the set P onto S, taking each point $p \in P$ to its closest point $\tilde{p} \in S$. Dey and Goswami in [DG04] introduced the definition of a *noisy* (k,r)*-sample*:

Definition 1 Noisy (k,r)-sample. A finite set of points P is a noisy (k,r)-sample if the following conditions hold:

1. $\tilde{P}$ is a r-sample of S.
2. For any $p \in P$; $d(p,\tilde{p}) \leq c_1\, r\, \text{lfs}(\tilde{p})$ for some constant c_1.
3. For any $p \in P$; $d(p,q) \geq c_2\, r\, \text{lfs}(\tilde{p})$, where q is the k^{th} nearest sample to p, for some constant c_2.

Here the first condition requires the sample to be dense enough, the second condition bounds the noise level, and the third condition requires that the sample is nowhere too dense (by requiring the k^{th} nearest sample to be far enough away). The third condition does not seem strictly necessary, and one of the contributions of this paper is to show that indeed it is not, at least for many of the geometric results used in the analysis. We will adopt a definition which we call a *noisy r-sample*, essentially only using conditions i) and ii):

Definition 2 Noisy r-sample. A finite set of points P is a noisy r-sample if the following two conditions hold:

1. $\tilde{P}$ is a r-sample of S.
2. For any $p \in P$, $d(p,\tilde{p}) \leq k_1\, r\, \text{lfs}(\tilde{p})$, for some constant k_1.

We define $\text{lfs}(S) = \min_{x\in S} \text{lfs}(x)$ for the surface as a whole. Assuming S is C^2 we have $\text{lfs}(S) > 0$ [APR02]. We also define the maximum local feature size $\Delta_1 = \max_{x\in S} \text{lfs}(x)$ and we have $\Delta_1 \leq \Delta_0$ (recall that Δ_0 is the radius of the largest medial ball).

3. Geometric constructions and the algorithm

To avoid to dealing with infinite Voronoi cells, we add to the sample set P a set Z of eight points, the vertices of a large box containing Ω.

The concept of *poles* was defined by Amenta and Bern [ACK01b] as follows:

Definition 3 The poles p_i, p_o of a sample $p \in P$, are the two vertices of its Voronoi cell farthest from p, one on either side of the surface. The Voronoi balls $B_{p_i,\rho_{p_i}}$, $B_{p_o,\rho_{p_o}}$ are the polar balls with radii $\rho_{p_i} = d(p_i,p)$ and $\rho_{p_o} = d(p_o,p)$ respectively.

Notice that given a noisy sample set not all Voronoi cells are long and skinny, as they are in the noise-free case.

A polar ball B_{v,ρ_v} is classified as an inner (outer) polar ball if its center is inside the inner (outer) region of $\mathbb{R}^3 \setminus S$. We denote by $\mathbb{P}_I$ and $\mathbb{P}_O$ the set of all inner and outer polar balls, respectively.

Algorithm

Our algorithm consists of a very simple modification to the power crust algorithm: we discard any poles such that the radius of the associated polar ball is smaller than $\frac{\text{lfs}(S)}{c}$ where $c > 1$ is a constant.

This can be summarized as follows.

Algorithm 3.1 Power Crust

1. Compute the Delaunay Diagram of $P \cup Z$.
2. Determine the set $\mathbb{P}$ of polar balls.
3. Delete from $\mathbb{P}$ any ball of radius $< \frac{\text{lfs}(S)}{c}$, producing $\mathbb{P}'$.
4. Compute the power diagram of $\mathbb{P}'$
5. Label the balls in $\mathbb{P}'$ as outer balls or inner balls, resulting in the sets $\mathbb{B}_O$ and $\mathbb{B}_I$.
6. Determine the faces in $\text{Pow}(\mathbb{B}_O \cup \mathbb{B}_I)$ separating inner from outer cells.

We discuss the labeling in step five in the Appendix A. It is done using exactly the same method as in the original power crust algorithms, but to show that it remains correct in the noisy case we need to prove a few more lemmas.

Analysis Overview

Most of our paper is concerned with the proof that this simple modification produces an output polyhedral surface which is correct, topologically and geometrically, given a noisy r-sample. Some of the lemmas are true for constant r independent of S, The lemmas 6-9 and Theorems 1 and 2 requires $r = O(\frac{\text{lfs}(S)}{\Delta_1})$.

We prove that a subset of the medial axis can be well approximated by the set of poles, this is stated in Lemma 6. As a consequence of this fact we prove in Lemma 8 that the boundary of the union of the set of big inner (outer) polar balls (see Equations 1 and 2) is close to the sampled surface, in the sense of Hausdorff distance. We use this fact in turn to show that the Hausdorff distance between the power crust and the sampled surface is $O(\sqrt[4]{r})$ (Theorem 1) and that the power crust is homeomorphic to the original surface S (Theorem 2).

4. Union of polar balls

Given a constant $c > 1$ we define the following two polar ball subsets:

$$\mathbb{B}_I = \{ B_{c,\rho_c} \in \mathbb{P}_I \,:\, \rho_c \geq \frac{\text{lfs}(S)}{c} \} \quad (1)$$

$$\mathbb{B}_O = \{ B_{c,\rho_c} \in \mathbb{P}_O \,:\, \rho_c \geq \frac{\text{lfs}(S)}{c} \} \quad (2)$$

The sets B_I and B_O are the sets of balls retained in our modified power crust algorithm. Their respective boundary sets are: $S_I = \partial(\bigcup_{B\in\mathbb{B}_I} B)$ and $S_O = \partial(\bigcup_{B\in\mathbb{B}_O} B)$. Our goal will be to prove that the boundary sets S_I and S_O are close to the surface S. Moreover, we will prove that a subset of the two-dimensional faces of the power diagram of $\mathbb{B}_I \cup \mathbb{B}_O$ is homeomorphic to the surface S.

Our proofs will also use another pair of subsets of the polar balls. We denote by $\mathbb{B}'_I$ and $\mathbb{B}'_O$ the set of inner and outer polar balls where each ball contains a medial axis point. That is,

$$\mathbb{B}'_I = \{ B_{c,\rho_c} \in \mathbb{P}_I \,:\, B_{c,\rho_c} \cap M \neq \emptyset \} \quad (3)$$

$$\mathbb{B}'_O = \{ B_{c,\rho_c} \in \mathbb{P}_O \,:\, B_{c,\rho_c} \cap M \neq \emptyset \} \quad (4)$$

The following lemma proves that $\mathbb{B}'_I \subset \mathbb{B}_I$ and $\mathbb{B}'_O \subset \mathbb{B}_O$ respectively.

Lemma 1 $\mathbb{B}'_I \subset \mathbb{B}_I$ and $\mathbb{B}'_O \subset \mathbb{B}_O$, for $c > 2$ and $r < \frac{c-2}{k_1 c}$.

Proof Take a ball $B_{x,\rho} \in \mathbb{B}'_I$ ($B_{x,\rho} \in \mathbb{B}_O$). There exists a sample p on $\partial B_{x,\rho}$ and there exists an inner (outer) medial axis point m inside $B_{x,\rho}$. Then we have that $d(\tilde{p},p) + 2\rho \geq d(\tilde{p},p) + d(p,m) \geq d(\tilde{p},m) \geq \text{lfs}(\tilde{p})$, and consequently $\rho \geq \frac{\text{lfs}(\tilde{p}) - d(\tilde{p},p)}{2} \geq \frac{1-k_1\cdot r}{2}\cdot\text{lfs}(\tilde{p})$. Taking $r \leq \frac{c-2}{k_1 c}$ we get that $\rho \geq \text{lfs}(\tilde{p})/c \geq \text{lfs}(S)/c$. □

The next lemma is a consequence of the sampling requirements and will be used for later proofs.

Lemma 2 Given P a noisy r-sample of S, let D be a ball with $\overset{o}{D} \cap P = \emptyset$ and $D \cap S \neq \emptyset$, let x be a point in $D \cap S$. If $B(x,\rho_x) \subset D$ then $\rho_x \leq r(1+2k_1)\text{lfs}(x)$.

Proof By sampling condition 1 of Definition 2. there exists a sample q such that $d(x,\tilde{q}) \leq r\,\text{lfs}(x)$. Using the fact that lfs is a one-Lipschitz function we have that $\text{lfs}(\tilde{q}) \leq d(\tilde{q},x) + \text{lfs}(x) \leq r\,\text{lfs}(x) + \text{lfs}(x) = (1+r)\text{lfs}(x)$.
By the sampling condition 2 and the previous equation we get $d(x,q) \leq d(x,\tilde{q}) + d(\tilde{q},q) \leq r\,\text{lfs}(x) + k_1\, r\,\text{lfs}(\tilde{q}) \leq (r + 2k_1 r)\text{lfs}(x)$. Since $\overset{o}{D} \cap P = \emptyset$ one deduces that $\overset{o}{B}(x,\rho_x) \cap P = \emptyset$, hence $\rho_x \leq d(x,q) \leq r(1+2k_1)\text{lfs}(x)$. □

Also we have the following lemma from Amenta and Bern [AB99] which estimates the angle between the normals to the surface at two close points.

Lemma 3 For any two points p and q on S with $d(p,q) \leq r\min\{\text{lfs}(p),\text{lfs}(q)\}$, for any $r \leq \frac{1}{3}$, the angle between the normals to S at p and q is at most $\frac{r}{1-3r}$.

A central idea in Voronoi-based surface reconstruction is that the Voronoi cells of a dense enough noise-free sample are long, skinny and perpendicular to the surface. This is not true for all Voronoi cells when there is noise, but the following lemma shows that it is true for large enough Voronoi cells. Specifically, given a sample point p and a point $x \in Vor(p)$ we bound the angle between the vector $\vec{xp}$ and the surface normal $\vec{n}_{\tilde{p}}$ at the projection of the sample p onto S. The lemma states that when x is far away from p, then this angle has to be small. In the noise-free case, "small" meant $O(r)$; here we achieve a bound of only $O(\sqrt{r})$.

Lemma 4 Let $p \in P$ be a sample such that there exists a point x on the inner (outer) region of the Voronoi cell of p with distance ρ_x between x and p satisfying the inequality $\rho_x \geq \frac{\text{lfs}(\tilde{p})}{c_1}$ for some constant c_1. Then the angle between the vector $\vec{xp}$ and the oriented outward (inward) surface normal $\vec{n}_{\tilde{p}}$ is $O(\sqrt{r})$.

Proof Denote by B_{m,ρ_m} the outer (inner) medial ball tangent to the surface S at $\tilde{p}$. Let B_{x,ρ_x} be the ball centered at x with radius $\rho_x = d(x,p)$. Since x is in the Voronoi cell of p we have $\overset{o}{B}_{x,\rho_x} \cap P = \emptyset$.
The angle between the vectors $\vec{xp}$ and $\vec{n}_{\tilde{p}}$ is the sum $\angle(t,x,p) + \angle(t,m,p)$, where the segment pt is perpendicular to xm, see figure 2. Our aim will be to find upper bounds for the angles $\angle(t,x,p)$ and $\angle(t,m,p)$, respectively. Since $d(x,t) < d(x,p) = \rho_x$, we have that $t \in B_{x,\rho_x}$, and the following two situations are possible: either $t \in B_{m,\rho_m} \cap B_{x,\rho_x}$ or $t \in B^c_{m,\rho_m} \cap B_{x,\rho_x}$.
First case: $t \in B_{m,\rho_m} \cap B_{x,\rho_x}$, see figure 2 left. Since $t \in B_{m,\rho_m}$ we have that t is on the outer (inner) region of $\Omega \setminus S$ and the ray l_x containing x and t intersects the surface at the point t_s lying between the points x and t, therefore $t_s \in B_{x,\rho_x}$ since the

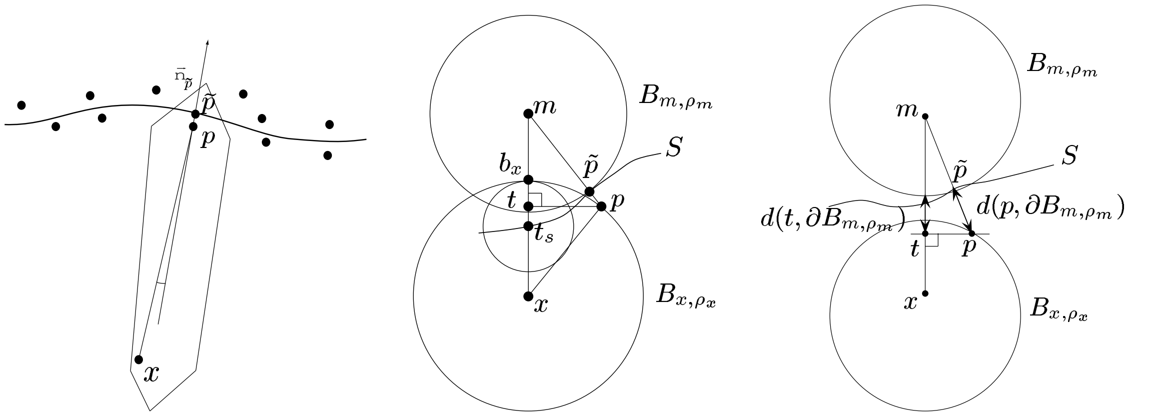

Figure 2: *Left: Illustration of Lemma 4, a fundamental result describing the shape of the Voronoi cells. When there exists a point x in $Vor(p)$ such that $d(x,p) \geq \frac{lfs(\tilde{p})}{c}$ then the angle between the segment $\vec{xp}$ and the normal $\vec{n}_{\tilde{p}}$ is a $O(\sqrt{r})$. Center: Figures used in the proof of Lemma 4. $t \in B_{m,\rho_m} \cap B_{x,\rho_x}$. Right $t \in B^c_{m,\rho_m} \cap B_{x,\rho_x}$*

segment $[x,t] \subset B_{x,\rho_x}$. Moreover, the ray l_x intersects $\partial B_{x,\rho_x}$ at the point b_x, see figure 2. Using that $t_s \in B_{x,\rho_x}$ we have, for small enough r, the following inequality that will be useful later:

$$\begin{aligned} \text{lfs}(t_s) &\leq d(t_s,\tilde{p}) + \text{lfs}(\tilde{p}) \leq \rho_x + d(x,p) + d(p,\tilde{p}) + \text{lfs}(\tilde{p}) \\ &\leq 2\rho_x + (1 + r\cdot k_1)\text{lfs}(\tilde{p}) \leq (2+2c_1)\rho_x = k_c\,\rho_x \quad (5) \end{aligned}$$

Because the points t_s and b_x are inside the ball B_{x,ρ_x} we have that $B_{t_s,d(t_s,b_x)} \subset B_{x,\rho_x}$. Since B_{x,ρ_x} is empty of samples (because ρ_x is the distance of x to its closest point in P), we have that $B_{t_s,d(t_s,b_x)}$ is also empty of samples. Consequently, by Lemma 2, we obtain $d(t_s,b_x) \leq O(r)\,\text{lfs}(t_s)$. From this last equation together with equation 5 and the fact that $t \in [b_x,t_s]$ we obtain the following two inequalities:

$$d(t,b_x) \leq d(t_s,b_x) \leq O(r)\text{lfs}(t_s) \leq O(r)\rho_x \quad (6)$$

$$d(t,t_s) \leq d(t_s,b_x) \leq O(r)\text{lfs}(t_s) \leq O(r)\rho_x \quad (7)$$

Consequently, by 6 we have $d(t,x) = \rho_x - d(t,b_x) \geq (1 - O(r))\rho_x$, hence

$$d(p,t) = \sqrt{d(p,x)^2 - d(x,t)^2} = O(\sqrt{r})\rho_x \quad (8)$$

so, the angle $\angle(t,x,p)$ is bounded by

$$\angle(t,x,p) = \arcsin\left(\frac{d(p,t)}{\rho_x}\right) = O(\sqrt{r}) \quad (9)$$

On the other hand, since $t \in B_{m,\rho_m}$, we have that $\text{lfs}(t_s) < d(t_s,t) + d(t,m) \leq O(r)\text{lfs}(t_s) + \rho_m$, thus obtaining $\text{lfs}(t_s) < \frac{\rho_m}{1-O(r)}$. Because the points m, t, t_s are collinear, $t \in B_{m,\rho_m}$ and $t_s \notin \overset{o}{B}_{m,\rho_m}$. So we obtain the following lower bound for the distance between t and m:

$$d(t,m) \geq \rho_m - d(t,t_s) \geq \rho_m - O(r)\text{lfs}(t_s) > (1 - O(r))\rho_m$$

.

Since $\text{lfs}(\tilde{p}) < \rho_m$, and using the sampling conditions, we get that $d(p,m) < d(m,\tilde{p}) + d(\tilde{p},p) \leq (1+O(r))\rho_m$, consequently

$$d(p,t) = \sqrt{d(p,m)^2 - d(t,m)^2} = O(\sqrt{r})\rho_m \quad (10)$$

We have that $\rho_m = d(m,\tilde{p}) \leq d(m,p) + d(p,\tilde{p})$, so using that $\text{lfs}(\tilde{p}) < \rho_m$ we have $d(m,p) \geq \rho_m - d(\tilde{p},p) \geq (1 - O(r))\rho_m$. From this equation and Equation 10 we can bound the angle $\angle(t,m,p)$ as follows:

$$\angle(t,m,p) = \arcsin\left(\frac{d(p,t)}{d(m,p)}\right) = O(\sqrt{r}) \quad (11)$$

Therefore, from 9 and 11 we have that our target angle $\angle(t,x,p) + \angle(t,m,p)$ is $O(\sqrt{r})$.
Second case: $t \in B^c_m \cap B_{x,\rho_x}$ (Note that this case implies that $B_{m,\rho_m} \cap B_{x,\rho_x} = \emptyset$). Since $t \notin B_{m,\rho_m}$ and $d(p,m) \geq d(t,m)$ we obtain that $p \notin B_{m,\rho_m}$, consequently we have that $d(t,\partial B_{m,\rho_m}) \leq d(p,\partial B_{m,\rho_m}) = d(p,\tilde{p}) \leq O(r)\text{lfs}(\tilde{p})$, see Figure 2 right. From this inequality and using the fact that $t \in B_{x,\rho_x}$ we get

$$d(t,x) \geq \rho_x - d(t,\partial B_{m,\rho_m}) \geq \rho_x - O(r)\text{lfs}(\tilde{p}) \quad (12)$$

Since $\text{lfs}(\tilde{p}) \leq d(\tilde{p},p) + d(p,x) \leq O(r)\text{lfs}(\tilde{p}) + \rho_x$ we get $\text{lfs}(\tilde{p}) \leq \frac{\rho_x}{1-O(r)}$. Consequently Equation 12 can be rewritten in terms of ρ_x, that is $d(t,x) \geq \rho_x - O(r)\text{lfs}(\tilde{p}) \geq (1 -$

$O(r))\rho_x$. We deduce the following upper bound for the distance between p and t

$$d(p,t) = \sqrt{d(p,x)^2 - d(t,x)^2} = O(\sqrt{r})\rho_x$$

Therefore, we have $\angle(t,x,p) = \arcsin\left(\frac{d(p,t)}{d(x,p)}\right) \leq \arcsin\left(\frac{O(\sqrt{r})\rho_x}{\rho_x}\right) = O(\sqrt{r})$.
On the other hand, since $t \notin B_{m,\rho_m}$ we get $d(t,m) > \rho_m$. $d(p,m) \leq d(p,\tilde{p}) + d(\tilde{p},m) \leq O(r)\text{lfs}(\tilde{p}) + \rho_m = (1 + O(r))\rho_m$, and hence

$$d(p,t) = \sqrt{d(p,m)^2 - d(t,m)^2} = O(\sqrt{r})\rho_m$$

and the angle $\angle(t,m,p) = \arcsin\left(\frac{d(p,t)}{d(p,m)}\right) \leq \arcsin\left(\frac{O(\sqrt{r})\rho_m}{\rho_m}\right) = O(\sqrt{r})$. Thus we conclude that the angle $\angle(t,x,p) + \angle(t,m,p)$ is $O(\sqrt{r})$. □

As a consequence of this lemma, we have that the inner and outer parts of the medial axis M are inside the sets $\bigcup_{B \in \mathbb{B}_I} \overset{o}{B}$ and $\bigcup_{B \in \mathbb{B}_O} \overset{o}{B}$ respectively, this is stated in the next lemma.

Lemma 5 Given an inner (outer) medial axis point m, then there exists an inner (outer) polar ball $B \in \mathbb{B}_I$ ($B \in \mathbb{B}_O$) such that $m \in \overset{o}{B}$.

Proof There exists a sample p such that m is inside its Voronoi cell. Denote by q the inner (outer) pole of p. Then by the definition of local feature size we have $d(m,\tilde{p}) \geq \text{lfs}(\tilde{p})$. By the triangle inequality we have $d(m,p) + d(p,\tilde{p}) \geq d(m,\tilde{p})$, so we have $d(m,p) \geq d(m,\tilde{p}) - d(p,\tilde{p}) \geq \text{lfs}(\tilde{p}) - rk_1\,\text{lfs}(\tilde{p})$. Taking $r \leq \frac{1}{2k_1}$ we get $d(m,p) \geq \frac{\text{lfs}(\tilde{p})}{2}$. This fact along with Lemma 4 implies that the angle $\angle(\vec{mp},\vec{n}_{\tilde{p}}) = O(\sqrt{r})$, using the same argument. Since $d(q,p) \geq d(m,p)$ we obtain $\angle(\vec{qp},\vec{n}_{\tilde{p}}) = O(\sqrt{r})$. Hence we obtain $\angle(\vec{qp},\vec{mp}) = O(\sqrt{r})$.
We take r small enough such that $\angle(\vec{qp},\vec{mp}) \leq \frac{\pi}{4}$. Since $d(m,p) \leq d(q,p)$ we find that m is inside the interior of the inner (outer) polar ball $B_{q,d(q,p)}$. Hence, we have that $B_{q,d(q,p)} \in \mathbb{B}'_I$ ($B_{q,d(q,p)} \in \mathbb{B}'_O$). By Lemma 1, $\mathbb{B}'_I \subset \mathbb{B}_I$ ($\mathbb{B}'_O \subset \mathbb{B}_O$), completing the proof. □

From now on assure that $r = O(\text{lfs}(S)/\Delta_1)$. We will show that the medial axis points m with angle $\angle qmx_1$ sufficiently large are well approximated by poles. The point q is the closest sample to m and x_1 is the closest sample of the closest surface point to m.

Lemma 6 Let m be a inner (outer) medial axis point such that $m \in \text{Vor}(q)$ for some sample q and let p be the inner (outer) pole of $\text{Vor}(q)$. Let $\tilde{x} \in S$ be the closest point to m on S and x_1 the closest sample to $\tilde{x}$. Then we have $|d(m,x_1) - d(m,q)| \leq O(r)$ and if the angle $\angle x_1mq > \sqrt{r}$, then $d(m,p) = O(\sqrt{r})$ and $|d(m,\tilde{x}) - d(p,q)| \leq O(\sqrt{r})$

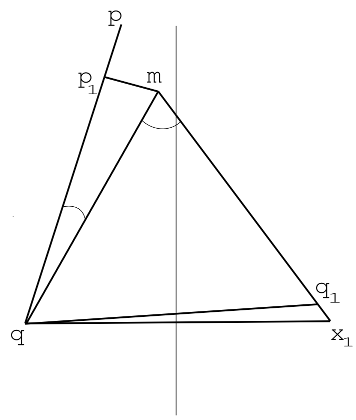

Figure 3: *The medial axis point $m \in Vor(q)$ is close to the pole p of $Vor(q)$ when angle $\angle qmx_1 > \sqrt{r}$*

Proof First we proof that the distances $d(q,m)$ and $d(m,x_1)$ are close. Let $\tilde{q}$ be the projection of the sample q onto S, since $\tilde{x}$ is the closest sample to m then we have that $d(m,\tilde{q}) \geq d(m,\tilde{x})$. Using that $d(m,\tilde{q}) \leq d(m,q) + d(q,\tilde{q})$ we have that $d(m,q) \geq d(m,\tilde{x}) - O(r)\text{lfs}(\tilde{q}) \geq d(m,\tilde{x}) - O(r)\Delta_1$. Reasoning in a similar way we have that $d(m,p) \leq d(m,x_1)$ where x_1 is the closest sample of $\tilde{x}$, we have using the triangular inequality $d(m,p) \leq d(m,x_1) \leq d(m,\tilde{x}) + d(\tilde{x},x_1) = d(m,\tilde{x}) + O(r)\text{lfs}(\tilde{x}) \geq d(m,\tilde{x}) + O(r)\Delta_1$ we get that:

$$|d(m,q) - d(m,\tilde{x})| \leq O(r\Delta_1) \tag{13}$$

Using that $|d(m,\tilde{x}) - d(m,x_1)| \leq O(r\Delta_1)$ together with equation 13 we get that

$$|d(m,q) - d(m,x_1)| \leq O(r\Delta_1) \tag{14}$$

Since $d(m,q) \leq d(m,x_1)$ then there exists a point q_1 in the segment mx_1 such that $d(m,q_1) = d(m,q)$, see figure 3. Using 14 we have that $d(q_1,x_1) \leq O(r\Delta_1)$.
Our goal is to bound the distance between p and m, in that direction we will prove that the distance $d(p,p_1) = d(p,q) - d(m,q)$ and $d(m,q_1)$ are small, so we will start by bounding $d(p,q)$.
The triangle qmq_1 is isosceles we have that $d(q,q_1) = 2\sin(\angle qmq_1/2)d(m,q)$ from this we get that $d(q,x_1) \leq d(q,q_1) + d(q_1,x_1) \leq 2\sin(\angle qmq_1/2)d(q,m) + O(r\Delta_1)$. Using this last equation we can bound the angle $\angle q_1qx_1 \leq \arcsin(\frac{d(q_1,x_1)}{d(q,q_1)}) \leq \arcsin(\frac{O(r\Delta_1)}{2\sin(O(\sqrt{r})d(m,q)}) \leq \arcsin(\frac{O(r\Delta_1)}{2\sin(O(\sqrt{r})\text{lfs}(S)}) = O(\sqrt{r}\Delta_1/\text{lfs}(S))$. The angle $\angle mqq_1$ satisfies that $\angle mqq_1 = \pi/2 - \angle qmq_1/2$. By the lemma 4 we can deduce that angle $\angle pqm = O(\sqrt{r})$. Combining all this angle we derive that $\angle pqx_1 = \pi/2 - \angle qmq_1/2 + \angle pqm + \angle q_1qx_1$. Using that $d(p,q) \leq d(p,x_1)$ we derive an upper bound for $d(p,q)$

$$\begin{aligned} d(p,q) &\leq \frac{d(q,x_1)/2}{\sin(\pi/2-\angle pqx_1)} \\ &= \frac{d(q,x_1)/2}{\sin(\angle qmq_1/2-\angle pqm-\angle q_1qx)} \\ &\leq \frac{d(m,q)}{\cos(\angle pqm+\angle q_1qx)} \\ &+ \frac{O(r\Delta_1)}{\sin(\angle qmq_1/2)\cos(\angle pqm+\angle q_1qx)} \\ &= \frac{d(m,q)}{\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S)))} \\ &+ \frac{O(\sqrt{r}\Delta_1)}{\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S)))} \end{aligned} \quad (15)$$

Recall that $\max_{s\in S}\mathrm{lhf}(s)=\Delta_1<\Delta_0$, where Δ_0 is the maximum of the radius of the medial balls. From equation 15 and taking the constant r such that $\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S)))\geq 1/2$ which implies that $r=O(\mathrm{lfs}(S)/\Delta_1)$ we have that

$$\begin{aligned} d(p,p_1) &= d(p,q)-d(m,q) \\ &\leq \frac{d(m,q)}{\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S)))} \\ &+ \frac{O(\sqrt{r}\Delta_1)}{\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S)))}-d(m,q) \\ &= \frac{1-\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S))}{\cos(O(\sqrt{r}\Delta_1/\mathrm{lfs}(S)))}d(m,q)+O(\sqrt{r}\Delta_1) \\ &= O(\sqrt{r}\Delta_1(1+\Delta_0/\mathrm{lfs}(S))) \end{aligned} \quad (16)$$

On the other hand using 13 we have that $d(m,q)\leq d(m,\tilde{x})+O(r\Delta_1)=\Delta_1+O(r\Delta_1)$, therefore we have that

$$\begin{aligned} d(m,p_1) &= 2\sin(\angle pqm/2)d(m,q) \\ &\leq 2\sin(O(\sqrt{r}))d(m,q)=O(\sqrt{r}\Delta_0) \end{aligned} \quad (17)$$

Combining equations 17, 16 we get that

$$\begin{aligned} d(m,p) &\leq d(m,p_1)+d(p_1,p) \\ &\leq O(\sqrt{r}\Delta_0)+O(\sqrt{r}\Delta_1(1+\Delta_0/\mathrm{lfs}(S))) \\ &= O(\sqrt{r}\Delta_0(2+\Delta_0/\mathrm{lfs}(S))). \end{aligned} \quad (18)$$

Finally, using that $|d(m,q)-d(q,p)|\leq d(p,p_1)$ and the equation 13 we have that $|d(m,\tilde{x})-d(p,q)|\leq O(\sqrt{r}\Delta_1)$ □

Using this fact, lemma 6, one can derives that the boundaries S_I and S_O of the union of balls $\bigcup_{B\in\mathbb{B}_I}B$ and $\bigcup_{B\in\mathbb{B}_O}B$ are close to the surface S. This is stated in lemma 8, to prove this lemma a technical lemma 7 is first introduced.

Lemma 7 Let B_{c_1,ρ_1} and B_{c_2,ρ_2} be two balls with $B_{c_1,\rho_1}\cap B_{c_2,\rho_2}\neq\emptyset$ and $B_{c_2,\rho_2}\not\subset B_{c_1,\rho_1}$. Let $\varepsilon<\rho_i, i=1,2$ be such that $d(c_1,c_2)\leq\varepsilon$ and $|\rho_1-\rho_2|\leq\varepsilon$. Let x_2 be a point on $\partial B_{c_2,\rho_2}\setminus B_{c_1,\rho_1}$ and $\{x_1\}=[c_2,x_2]\cap\partial B_{c_1,\rho_1}$. Then $d(x_1,x_2)\leq 2\varepsilon$.

Proof We have $d(x_1,x_2)=\rho_2-d(c_2,x_1)$. By the triangular inequality we have $d(c_2,x_1)\geq d(c_1,x_1)-d(c_2,c_1)=\rho_1-d(c_2,c_1)$. From these two inequalities we obtain $d(x_1,x_2)\leq\rho_2-\rho_1+d(c_2,c_1)\leq|\rho_2-\rho_1|+d(c_2,c_1)\leq 2\varepsilon$ □

Lemma 8 $d_H(S_I,S)\leq O(\sqrt[4]{r})$ and $d_H(S_O,S)\leq O(\sqrt[4]{r})$.

Proof We show that $d_H(S_I,S)\leq O(\sqrt[4]{r})$; the argument for S_O is identical. We begin by showing that $\tilde{d}_H(S_I,S)\leq O(\sqrt[4]{r})$.
Consider any point $x\in S_I$. First assume that x is on the outside of S. Let B_{c,ρ_c} be a polar ball in B_I such that $x\in\partial B_{c,\rho_c}$. Then the segment $[c,x]$ from the center of the polar ball to x intersects S in a point s. Since the ball $B_{s,d(s,x)}$ is inside the polar ball $B_{c,\rho}$, Lemma 2 implies that $d(x,S)\leq d(x,s)=O(r)$ and we are done.
So let us assume that $x\in S_I$ is in the inner region of S. Let $\tilde{x}\in S$ be the closest point to x on S, and let $m_{\tilde{x}}$ be the center of the inner medial axis ball $B_{m_{\tilde{x}},\rho_{\tilde{x}}}$ tangent to S at $\tilde{x}$. Then we have that x is inside the segment $[\tilde{x},m_{\tilde{x}}]$; otherwise the ball $B_{x,d(x,\tilde{x})}$ with $\overset{o}{B}_{x,d(x,\tilde{x})}\cap S=\emptyset$ contains $B_{m_{\tilde{x}},\rho_{\tilde{x}}}$ which is a contradiction due to the ball $B_{m_{\tilde{x}},\rho_{\tilde{x}}}$ is maximal.
The medial axis point $m_{\tilde{x}}$ belongs to the Voronoi cell of some sample point q, let p and B_{p,ρ_p} be the inner pole of q and its polar ball respectively. The first part of lemma 6 states that the distances $d(m_{\tilde{x}},x_1)$ and $d(m_{\tilde{x}},q)$ where x_1 is the closest sample to $\tilde{x}$ are very close, that is $|d(m_{\tilde{x}},x_1)-d(m_{\tilde{x}},q)|\leq O(r)$. Suppose that the angle $\angle qmx_1\leq\sqrt{r}$, this implies that $d(x_1,q)$ is small. Since $d(m_{\tilde{x}},q)\leq d(m_{\tilde{x}},x_1)$, then there exists a point q_1 in the segment $m_{\tilde{x}}x_1$ such that $d(m_{\tilde{x}},q_1)=d(m_{\tilde{x}},q)$ and $d(x_1,q_1)=|d(m_{\tilde{x}},x_1)-d(m_{\tilde{x}},q)|\leq O(r)$. Therefore, using that $\angle x_1mq\leq\sqrt{r}$ we have $d(q,x_1)\leq d(q,q_1)+d(q_1,x_1)\leq 2\sin(\angle x_1mq/2)d(m,q)+O(r)\leq O(\sqrt{r})$ and consequently $d(\tilde{x},q)\leq d(\tilde{x},x_1)+d(x_1,q)\leq O(r)\mathrm{lfs}(\tilde{x})+O(\sqrt{r})=O(\sqrt{r})$.
On the other hand, the lemma 4 implies that $\angle pqm=O(\sqrt{r})$, from this fact and using that $\mathrm{lfs}(S)/2\leq\mathrm{lfs}(\tilde{q})/2\leq\mathrm{lfs}(\tilde{q})-d(q,\tilde{q})\leq d(m,q)\leq d(p,q)$ we have for r small enough that $m_{\tilde{x}}\in B_{p,\rho_p}$ and consequently $B_{p,\rho_p}\in\mathbb{B}_I'\subset\mathbb{B}_I$. The point $\tilde{x}\notin B_{p,\rho_p}$ because, otherwise $x\in B_{p,\rho_p}$ which is a contradiction with $x\in S_I$. Therefore, there exists $x_2=[\tilde{x},m_{\tilde{x}}]\cap\partial B_{p,\rho_p}$ and $x\in[\tilde{x},x_2]$. From the fact that $d(\tilde{x},q)\leq O(\sqrt{r})$ and $q\in\partial B_{p,\rho_p}$ we get that $d(\tilde{x},x)\leq d(\tilde{x},x_2)\leq\sqrt{d(\tilde{x},q)^2+2\rho_p d(\tilde{x},q)}=O(\sqrt[4]{r})$ and we are done.
Now we consider the case $\angle qmx_1>\sqrt{r}$. The lemma 6 implies $d(m,p)\leq O(\sqrt{r})$ and $|\rho_{\tilde{x}}-\rho_p|=O(\sqrt{r})$. Since $d(m,p)\leq O(\sqrt{r})$, then for r small enough $m_{\tilde{x}}$ belongs to B_{p,ρ_p} and consequently $B_{p,\rho_p}\in\mathbb{B}_I'\subset\mathbb{B}_I$. Recall that $\tilde{x}\in S$ and that $x\in[\tilde{x},m_{\tilde{x}}]$. If $\tilde{x}$ is inside $\overset{o}{B}_{p,\rho_p}$, then $[\tilde{x},m_{\tilde{x}}]\subset\overset{o}{B}_{p,\rho_p}$, so that $x\in\overset{o}{B}_{p,\rho_p}$. But this contradicts the fact that $x\in S_I$. Hence it must be the case that $\tilde{x}$ is on $\partial B_{m_{\tilde{x}},\rho_{\tilde{x}}}\setminus\overset{o}{B}_{p,\rho_p}$.
Let $x_2=[m_{\tilde{x}},\tilde{x}]\cap\partial B_{p,\rho_p}$ (this intersection point is unique). We have that $x\in[x_2,\tilde{x}]$; otherwise, $x\in(m_{\tilde{x}},x_2)$, the portion of the segment inside $\overset{o}{B}_{p,\rho_p}$, which again is a contradiction with the fact that $x\in S_I$ and $B_{p,\rho_p}\in\mathbb{B}_I$. Now applying

Lemma 7, we have that $d(x_2,\tilde{x}) \leq 2O(\sqrt{r})$ and $d(x,\tilde{x}) \leq d(x_2,\tilde{x}) \leq O(\sqrt{r})$, so we have proved that $\tilde{d}_H(S_I,S) \leq O(\sqrt{r})$.
Now we will prove that $\tilde{d}_H(S,S_I) \leq O(\sqrt{r})$. Let x be an arbitrary point on S and let B_m and $B_{m'}$ be the inner and outer medial balls tangent to S at x respectively. The segment $[m,m']$ is orthogonal to S at x.
Now we will establish that there exists a point x_1 on $S_I \cap (m,m')$. Suppose not; then $S_I \cap (m,m') = \emptyset$, and there exists a ball $B_{c,\rho} \in \mathbb{B}_I$ such that $m' \in B_{c,\rho}$. Since c and m' are on opposite sides of S, then the segment $[c,m']$ intersects S at a point s, so we have that $m' \in B_{s,d(s,\partial B_{c,\rho})} \subset B_{c,\rho}$ with $B_{s,d(s,\partial B_{c,\rho})}$ empty of samples. From Lemma 2 we have $d(s,\partial B_{c,\rho}) = O(r)\text{lfs}(s) < d(s,m')$, which implies that $m' \notin B_{s,d(s,\partial B_{c,\rho})}$, obtaining a contradiction we the fact that the segment $[s,m']$ is contained in $B_{s,d(s,\partial B_{c,\rho})}$.
We can conclude there exists a point x_1 on $S_I \cap (m,m')$. Since the closest point to x_1 on S is the point x (the segment $[x_1,x]$ is orthogonal to the surface at x), we have $d(x,x_1) = d(x_1,S) \leq \tilde{d}_H(S_I,S) \leq O(\sqrt[4]{r})$. Hence $d(x,S_I) \leq O(\sqrt[4]{r})$ and consequently $\tilde{d}_H(S,S_I) \leq O(\sqrt[4]{r})$. □

5. Power Crust

The *power diagram* of a set of balls $\mathbb{B}$ is the weighted Voronoi diagram which assigns an unweighted point x to the cell of the ball $B \in \mathbb{B}$ which minimizes the power distance $d_{pow}(x,B)$. The power distance between a point and a ball $d_{pow}(x,B_{c,\rho}) = d(x,c)^2 - \rho^2$. We denote it by $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O)$. In the next two theorem we will prove that $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O)$ is a polyhedral surface homeomorphic and close to the original surface S.

Taking $\varepsilon < \text{lfs}(S)$ we denote by $N_\varepsilon = \{x \in \mathbb{R}^3 \,:\, d(x,\tilde{x}) \leq \varepsilon\}$ a tubular neighborhood around S. The boundary of N_ε is $S_\varepsilon \bigcup S_{-\varepsilon}$ where $S_{\pm\varepsilon} = \{x \in \mathbb{R}^3 \,:\, x = \tilde{x} \pm \varepsilon n_{\tilde{x}}\}$ are two offset surfaces. When $d_H(S_I,S) < \varepsilon$ and $d_H(S_O,S) < \varepsilon$ (Lemma 8), the boundary S_I (S_O) of the sets $\bigcup_{B\in\mathbb{B}_I} B$ ($\bigcup_{B\in\mathbb{B}_O} B$) is inside the set N_ε and consequently the sets S_ε and $S_{-\varepsilon}$ are inside the interior of the sets $\bigcup_{B\in\mathbb{B}_O} B$ and $\bigcup_{B\in\mathbb{B}_I} B$ respectively.

Theorem 1 If $d_H(S_I,S) \leq \varepsilon$ and $d_H(S_O,S) \leq \varepsilon$ then the Hausdorff distance between $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O)$ and S is smaller than 2ε.

Proof Let $\text{I}(S_{-2\varepsilon})$ be the part of $\Omega \setminus S_{-2\varepsilon}$ inside the interior part of S and let $\text{O}(S_{2\varepsilon})$ be the part of $\Omega \setminus S_{2\varepsilon}$ inside the exterior of S. Hence, we have $\Omega \setminus N_{2\varepsilon} = \text{I}(S_{-2\varepsilon}) \cup \text{O}(S_{2\varepsilon})$ with $\text{I}(S_{-2\varepsilon}) \cap \text{O}(S_{2\varepsilon}) = \emptyset$. From the conditions $d_H(S_I,S) \leq \varepsilon$ and $d_H(S_O,S) \leq \varepsilon$ we can deduce that $\text{O}(S_{2\varepsilon}) \subset (\bigcup_{B\in\mathbb{B}_O} B)$ and $\text{I}(S_{-2\varepsilon}) \subset (\bigcup_{B\in\mathbb{B}_I} B)$. Also one has $(\bigcup_{B\in\mathbb{B}_I} B) \cap \text{O}(S_{2\varepsilon}) = \emptyset$ and $(\bigcup_{B\in\mathbb{B}_O} B) \cap \text{I}(S_{-2\varepsilon}) = \emptyset$.
First we will prove that $\tilde{d}_H(\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O),S) \leq 2\varepsilon$. This is equivalent to proving that $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O) \subset N_{2\varepsilon}$. Let f be a face of $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O)$ separating the cell of the ball $B_1 \in \mathbb{B}_I$ from the cell of the ball $B_2 \in \mathbb{B}_O$ and let x be a point on f. Because $d_{pow}(x,B_2) = d_{pow}(x,B_1)$ we know that $d_{pow}(x,B_2)$ and $d_{pow}(x,B_1)$ have the same sign, implying that when it is negative then $x \in B_1 \cap B_2$ and otherwise $x \notin \bigcup_{B\in\mathbb{B}_I \bigcup \mathbb{B}_O} B$. In the first case because x is simultaneously in $(\bigcup_{B\in\mathbb{B}_I} B)$ and $(\bigcup_{B\in\mathbb{B}_O} B)$ then from the previous observation at the beginning of the lemma one deduces that $x \in N_{2\varepsilon}$.
The second cases we have $x \notin \bigcup_{B\in\mathbb{B}_I \bigcup \mathbb{B}_O} B$, but due to $\text{O}(S_{2\varepsilon}) \subset (\bigcup_{B\in\mathbb{B}_O} B)$ and $\text{I}(S_{-2\varepsilon}) \subset (\bigcup_{B\in\mathbb{B}_I} B)$ then we have that $x \in N_{2\varepsilon}$.
Now we will prove that $\tilde{d}_H(S,Pow(\mathbb{B}_I \cup \mathbb{B}_O)) \leq 2\varepsilon$. Given a point $x \in S$ the interval $[x + 2\varepsilon n_{\tilde{x}}, x - 2\varepsilon n_{\tilde{x}}]$ has boundary points $x + 2\varepsilon n_{\tilde{x}}$ and $\tilde{x} - 2\varepsilon n_{-\varepsilon}$ in the interior of the set $\bigcup_{B\in\mathbb{B}_I} B$ and $\bigcup_{B\in\mathbb{B}_O} B$ respectively, hence we have that $\tilde{x} + 2\varepsilon n_{\tilde{x}}$ is in the power cell of some ball in $\mathbb{B}_O$ and $\tilde{x} - 2\varepsilon n_{\tilde{x}}$ is in the power cell of some ball in $\mathbb{B}_I$, therefore moving a point along the interval $[\tilde{x} + 2\varepsilon n_{\tilde{x}}, \tilde{x} - 2\varepsilon n_{\tilde{x}}]$ it will meet at a face of the power crust at some point, otherwise it will stay forever in outer power cells which is a contradiction with the fact that $\tilde{x} - 2\varepsilon n_{\tilde{x}}$ belongs to some inner power cell. □

From the above theorem and the fact that $d_H(S_I,S) = O(\sqrt[4]{r})$ and $d_H(S_O,S) = O(\sqrt[4]{r})$ we can deduce that $d_H(\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O),S) = O(\sqrt[4]{r})$.

Now we extend the lemma [23] of Amenta, Choi and Kolluri [ACK01b] to a more general setting in which the point u does not need to be on the surface.

Lemma 9 Given a point u and a ball $B_{c,\rho} \in \mathbb{B}_I$ ($B_{c,\rho} \in \mathbb{B}_O$) such that $d(u,\partial B_{c,\rho}) \leq O(\varepsilon)$ and $u \in N_\varepsilon$, then the angle between the vector $\vec{cu}$ and the outward (inward) normal $\vec{n}_{\tilde{u}}$ is $O(\sqrt{\varepsilon})$.

Proof See lemma 9 in Appendix A □

Define by $f_I(x) = \min_{B\in\mathbb{B}_I} d_{pow}(x,B)$ and $f_O(x) = \min_{B\in\mathbb{B}_O} d_{pow}(x,B)$ the functions which return the minimum *power distance* from x to the sets $\mathbb{B}_I$ and $\mathbb{B}_O$ respectively. Based in this two function the following lemma 2 from Amenta, Choi and Kolluri [ACK01b] is also valid under our sampling assumption and for our particular polar ball sets $\mathbb{B}_I$ and $\mathbb{B}_O$. We show functions f_I and f_O are strictly monotonic and have a single intersection point along the segment $[\tilde{x} + 2\varepsilon n_{\tilde{x}}, \tilde{x} - 2\varepsilon n_{\tilde{x}}]$ since $f_I(\tilde{x} + 2\varepsilon n_{\tilde{x}}) f_O(\tilde{x} + 2\varepsilon n_{\tilde{x}}) < 0$ and $f_I(\tilde{x} - 2\varepsilon n_{\tilde{x}}) f_O(\tilde{x} - 2\varepsilon n_{\tilde{x}}) < 0$.

Theorem 2 The power crust of $\mathbb{B}_I \bigcup \mathbb{B}_O$ is a polyhedral surface homeomorphic to S.

Proof From the Lemma 8 we have that $d_H(S_I,S) = O(\sqrt[4]{r})$ and $d_H(S_I,S) = O(\sqrt[4]{r})$ and from theorem 1 we have $d_H(\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O),S) = O(\sqrt[4]{r})$. We will take $\varepsilon = 2d_H(\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O),S)$ which is smaller than $\text{lfs}(S)$ for small r. Given a point $\tilde{x} \in S$ we have $[\tilde{x} - \varepsilon n_{\tilde{x}}, \tilde{x} + \varepsilon n_{\tilde{x}}] \subset N_\varepsilon$. Let $d : \text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O) \to S$ the function that given a point $x \in \text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O)$ assigns the closest point $d(x) \in S$. Due to the previous lemma we have $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O) \subsetneq N_\varepsilon$ and since the set of points where the distance function is undefined is

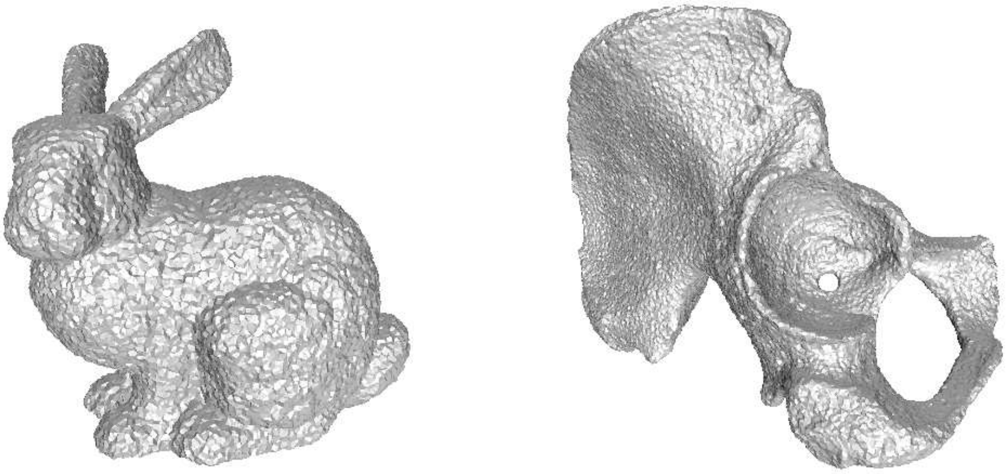

Figure 4: *Bunny and hip-bone models. The vertices of the hip-bone model were randomly perturbed using Gaussian noise, while noisy points were added to the vertex set of the bunny model to increase the density. The bumpy but topologically correct outputs shown here were produced by applying our modified power crust algorithm to the noisy point clouds.*

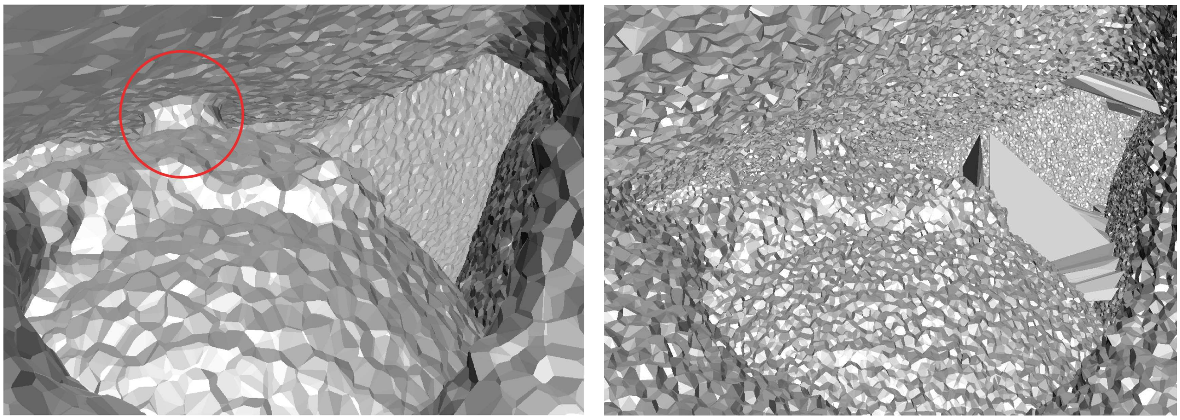

Figure 5: *View from inside of the hip model. On the left, our proposed method. The feature inside the red circle is the inside view of the small hole in the middle of the hip which can be seen in Figure 4. On the right, the original power crust algorithm, which has some artifacts on the interior.*

the medial axis then the distance function is well defined on the power crust.

We will prove it is a homeomorphism. Because the power crust is a compact set (it is a finite union of compact sets in this case faces) then we only need to prove that $d(\cdot)$ is a continuous, one-to-one and onto mapping. The continuity follows because the distance function to any set is an one-Lipschitz function. The onto condition follows from $d_H(\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O), S) \leq \varepsilon$, that is for any point $\tilde{x} \in S$ there exists at least a power crust point in $[\tilde{x} - \varepsilon n_{\tilde{x}}, \tilde{x} + \varepsilon n_{\tilde{x}}]$ and given a point in $[\tilde{x} - \varepsilon n_{\tilde{x}}, \tilde{x} + \varepsilon n_{\tilde{x}}]$ with $\varepsilon \leq \text{lfs(S)}$ its closest point on S is $\tilde{x}$.

The one-to-one condition. Suppose that it is false, it implies that there are two points x_1 and x_2 on $\text{Pow}(\mathbb{B}_I \cup \mathbb{B}_O)$ such that $d(x_1) = d(x_2)$ or equivalent $\tilde{x}_1 = \tilde{x}_2$ where x_1 and x_2 belong to $[\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}, \tilde{x}_1 + \varepsilon n_{\tilde{x}_1}]$. Given a point $x \in [\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}, \tilde{x}_1 + \varepsilon n_{\tilde{x}_1}]$ let $B_{c_x,\rho_x} \in \mathbb{B}_I$ be a ball which satisfies $d_{pow}(x, B_{c_x,\rho_x}) = f_I(x)$. Let $B_{\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}} \in \mathbb{B}_I$ be a ball which contains the point $\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}$ then we have $d_{pow}(x, B_{c_x,\rho_x}) \leq d_{pow}(x, B_{\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}}) \leq (\rho + d(x, \partial B_{\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}}))^2 - \rho^2 = O(\varepsilon^2)$ where ρ is the radius of the ball $B_{\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}}$. From this fact $d_{pow}(x, B_{c_x,\rho_x}) < O(\varepsilon^2)$ we obtain that $d(x, \partial B_{c_x,\rho_x}) \leq O(\varepsilon)$, so applying the lemma 9 to the point x we obtain that the angle between the outward normal $n_{\tilde{x}}$ and the vector $\vec{c_x x}$ is $O(\sqrt{\varepsilon})$ and consequently for small enough r we obtain that this angle is smaller than $\pi/2$. This means that when

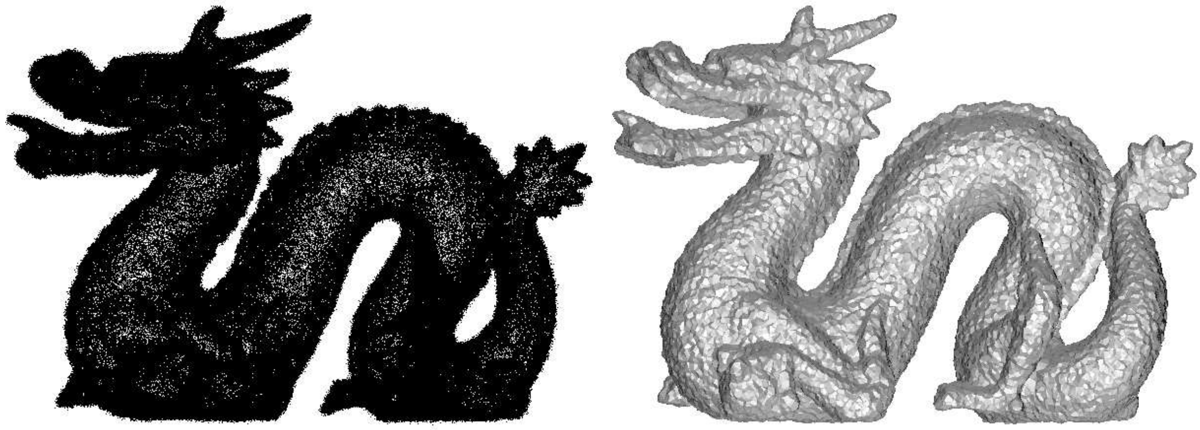

Figure 6: *Reconstruction of the dragon model perturbed with Gaussian noise. The perturbed point cloud is shown on the left.*

we move the point x from $\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}$ to $\tilde{x}_1 + \varepsilon n_{\tilde{x}_1}$ along the segment $[\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}, \tilde{x}_1 + \varepsilon n_{\tilde{x}_1}]$ we have that the function f_I is strictly decreasing. The same argument shows that the function f_O is strictly increasing.

A power crust points x is characterized by the following equality $f_I(x) = f_O(x)$. Using that $f_I(\tilde{x}_1 \pm \varepsilon n_{\tilde{x}_1}) \cdot f_O(\tilde{x}_1 \mp \varepsilon n_{\tilde{x}_1}) < 0$ and the functions f_I and f_O are strictly decreasing and increasing respectively along the interval $[\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}, \tilde{x}_1 + \varepsilon n_{\tilde{x}_1}]$ then there exist an unique point x_3 on $[\tilde{x}_1 - \varepsilon n_{\tilde{x}_1}, \tilde{x}_1 + \varepsilon n_{\tilde{x}_1}]$ such that $f_I(x_3) = f_O(x_3)$. From this we conclude that $x_1 = x_2 = x_3$ and the function $d(\cdot)$ is one-to-one. $\square$

6. Implementation and Experiments

Since we do not know $\mathrm{lfs}(S)$ for a given input surface, we choose the size of the balls to eliminate by trial and error in each case.

Our experiments were done using an in-house implementation of the power crust algorithm, due to Ravi Kolluri. This code uses Jonathan Shewchuk's currently unreleased `pyramid` code for Delaunay triangulation. Filtering the polar balls required adding exactly eleven lines of code to the power crust implementation.

We tested the algorithm with several data sets, produced by taking polyhedral models and adding noise. The results are shown in Figures 4, 5 and 6. The bunny and the dragon were taken from the Stanford 3D scanning repository, and the hip-bone is from the Cyberware Web site. For the Stanford bunny we added four new samples per vertex respectively, each perturbed with Gaussian noise. For the hip-bone and the dragon models, which are already fairly large, we just perturbed the input samples. The bunny point set consisted of 179,736 points and the reconstruction was computed in less than a minute. The hip-bone set contained 397,625 points and the reconstruction required about 3 minutes, while the dragon point set contained 875,290 and required about 10 minutes. Experiments were done on a Pentium 4, 2.4GHz, with 1Gb of memory.

In each reconstruction we chose the constant δ used to filter the polar balls based on the noise level, with δ being four times the variance of the Gaussian. The noise level in turn was chosen to be less than the smallest feature of the input model, for instance to avoid filling in the hole in the hip-bone or connecting the neck of the dragon to its back.

References

[AB99] AMENTA N., BERN M.: Surface reconstruction by voronoi filtering. *Disc. Comput. Geom 22* (1999), 481–502. 1, 4

[ABK98] AMENTA N., BERN M., KAMVYSSELIS M.: A new voronoi-based surface reconstruction algorithm. In *Proceedings of SIGGRAPH* (1998), pp. 415–421. 1

[ACDL00] AMENTA N., CHOI S., DEY T. K., LEEKHA N.: A simple algorithm for homeomorphic surface reconstruction. *Internat J. Comput. Geom & Applications* (2000), 125–141. 1

[ACK01a] AMENTA N., CHOI S., KOLLURI R.: The power crust. In *Proceedings of 6th ACM Symposium on Solid Modeling* (2001), pp. 249–260. 1

[ACK01b] AMENTA N., CHOI S., KOLLURI R.: The power crust, union of balls and medial axis transform. *Comput. Geom: Theory & Applications 19* (2001), 127–153. 1, 2, 3, 8, 11

[APR02] AMENTA N., PETERS T. J., RUSSELL. A.: Computational topology: Ambient isotopic approximation of 2-manifolds. *Theoretical Computer Science* (2002). 3

[BC02] BOISSONNAT J. D., CAZALS F.: Smooth surface reconstruction via natural neighbor interpolation of distance function. *Comp. Geometry Theory and Applications* (2002), 185–203. 1

[BMR*99] BERNARDINI F., MITTLEMAN J., RUSHMEIR H., SILVA C., TAUBIN G.: The ball-pivoting algorithm for surface reconstruction. *IEEE Transactions on Vision and Computer Graphics 5*, 4 (1999). 1

[CBC*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3d objects with radial basis functions. In *Proceedings of SIGGRAPH* (2001), pp. 67–76. 1

[CFG*03] CHEN W. S., FUNKE F., GOLIN M., KUMAR P., POON H. S., RAMOS E.: Curve reconstruction from noisy samples. *Proc. 19th Annu. Sympos. Comput. Geom.* (2003), 302–311. 2

[CL96] CURLESS B., LEVOY M.: A volumetric method for building complex models from range images. In *Proceedings of SIGGRAPH* (1996), pp. 303–312. 1

[CL05] CHAZAL F., LIEUTIER A.: The λ-medial axis. *Graphical Models 67*, 4 (2005), 304–331. 2

[DG04] DEY T. K., GOSWAMI S.: Provable surface reconstruction from noisy samples . *Discrete and Computational Geometry 29*, 3 (2004), 15–30. 1, 3

[HDD*92] HOPPE H., DEROSE T., DUCHAMP T., MCDONALD J., STUETZLE W.: Surface reconstruction from unorganized points. In *Proceedings of SIGGRAPH* (1992), pp. 71–78. 1

[Kol05] KOLLURI R.: Provably good moving least squares. In *Proceedings of ACM Symposium on Discrete Algorithms* (2005). 2

[Lev03] LEVIN D.: Mesh-independent surface interpolation. In *Geometric Modeling for Scientific Visualization*, Brunnett G., Hamann B., Mueller K.„ Linsen L., (Eds.). Springer-Verlag, 2003. 2

[LPC*00] LEVOY M., PULLI K., CURLESS B., RUSINKIEWICZ S., KOLLER D., PEREIRA L., GINZTON M., ANDERSON S., DAVIS J., GINSBERG J., SHADE J., FULK D.: The digital michelangelo project: 3D scanning of large statues. In *Proceedings of SIGGRAH* (2000), pp. 131–144. 1

Appendix A: Labeling Algorithm

Once we have determined the set $\mathbb{P}'$ of polar balls to be retained in the noisy version of the power crust algorithm, and computed their power diagram, the next step of the algorithm is to label each of the balls in $\mathbb{P}'$ as an outer or inner ball, thus determining the sets $\mathbb{B}_I$ and $\mathbb{B}_O$. We use exactly the same labeling algorithm as in the original power crust implementation [ACK01b], but we explain it here for completeness. Then we prove a couple of lemmas which guarantee that the labeling algorithm is correct. These proofs a similar to analogous proofs in the noise-free case, but again we include them for completeness.

For each sample in the special set Z of vertices of the bounding box, its polar ball is inserted in a queue and labeled as outer. Then we iteratively propagate the labeling. While the queue is not empty, we remove a ball B_p from the queue. We examine each of the balls B_q whose cells neighbor that of B_p in the power diagram of $\mathbb{P}'$. If the intersection between B_p and B_q is at a angle bigger than $\pi/4$, we assign to B_p the same label as B_q. Also we assign the opposite label to the ball of the other pole of p, if there is one. This process is repeated until there is not a new ball that can be classified as outer or inner. Once we have finished the labeling we determine the faces in $\mathrm{Pow}(\mathbb{P}')$ separating inner balls from outer ones.

Lemma 10 The angle of intersection between a polar ball $B_{c_1,\rho_1} \in \mathbb{B}_I$ and a polar ball $B_{c_2,\rho_2} \in \mathbb{B}_O$ with $B_{c_2,\rho_2} \cap B_{c_1,\rho_1} \neq \emptyset$ is $O(r)$.

Proof We have that the center c_1 and c_2 of B_{c_1,ρ_1} and B_{c_2,ρ_2} are in different sides of S, thus the segment $[c_1,c_2]$ intersects the surface at a point $x \in S$. Let $\{b_i\} = [c_1,c_2] \cap \partial B_{c_i,\rho_i}$ with $i = 1,2$.
The ball $B_{x,d(x,b_i)}$ is inside B_{c_i,ρ_i} which is empty of samples. Then by lemma 2 we obtain that $d(x,b_i) = O(r)\mathrm{lfs}(x)$. Hence $d(b_1,b_2) \leq d(x,b_1) + d(x,b_2) \leq O(r)\mathrm{lfs}(x)$, consequently $d(b_1,b_2) \leq O(r)\max_{x\in S}\mathrm{lfs}(x) = O(r)\Delta_1$.
Let Π a plane containing the intersection circle between the balls B_{c_1,ρ_1} and B_{c_2,ρ_2} and $\{z\} = \Pi \cap [c_1,c_2]$. Let us bound the distance c_i to z we have that $d(c_i,z) \geq \rho_i - d(b_1,b_2) \geq \rho_i - O(r)\Delta_1$. Since that $\rho_i \geq \mathrm{lfs}(S)/c$, then for small enough r we have:

$$\cos(\alpha_i) = \frac{d(c_i,z)}{\rho_i} \geq \frac{\rho_i - O(r)\Delta_1}{\rho_i} \geq 1 - O(r)\frac{c\,\Delta_1}{\mathrm{lfs}(S)} = 1 - O(r)$$

Hence we have that $\alpha_i = O(r)$ for $i = 1,2$ and the angle between the two balls is $\alpha_1 + \alpha_2 = O(r)$ □

Lemma 11 Let ε be smaller than $\mathrm{lfs}(S)$. Given a point u and a ball $B_{c,\rho} \in \mathbb{B}_I$ ($B_{c,\rho} \in \mathbb{B}_O$) such that $d(u,\partial B_{c,\rho}) \leq O(\varepsilon)$ and $u \in N_\varepsilon$, then the angle between the vector $\vec{cu}$ and the outward (inward) normal $\vec{n}_{\tilde{u}}$ is $O(\sqrt{\varepsilon})$.

Proof Let B_{m,ρ_m} be the outer medial axis ball tangent to S at $\tilde{u}$ and let $B_{c,\rho} \in \mathbb{B}_I$ be a ball such that $d(u,\partial B_{c,\rho}) = O(\varepsilon)$.
Let N_ε be a tubular neighborhood of S. It is easy to see that the ball $B_{m,\rho_m-\varepsilon}$ is inside the outer solid which is delimited by S_ε and Ω, therefore $B_{c,\rho} \cap B_{m,\rho_m-\varepsilon} = \emptyset$.
The points c, m and u form a triangle and the point t is

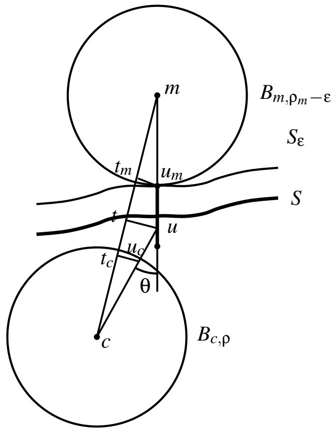

Figure 7: *The angle* θ *between the vector* $\vec{cu}$ *and the normal* $\vec{n}_{\tilde{u}}$ *at* $\tilde{u}$

.

the projection of u onto the segment cm. Our aim is to find an upper bound for angle between the vectors $\vec{cu}$ and $\vec{n}_{\tilde{u}}$ which is $\theta = \angle(m,c,u) + \angle(u,m,u)$ see figure 7. We have the following identities $\angle(c,m,u) = \arcsin\left(\frac{d(u,t)}{d(u,m)}\right)$ and $\angle(m,c,u) = \arcsin\left(\frac{d(u,t)}{d(u,c)}\right)$.

There are three possibilities: $t \in B_{m,\rho_m-\varepsilon}$, $t \in B_{c,\rho}$ and $t \notin B_{m,\rho_m-\varepsilon} \bigcup B_{c,\rho}$. When $t \in B_{m,\rho_m-\varepsilon}$ ($t \in B_{c,\rho}$) using the equations $d(u,t) = \sqrt{d(u,m)^2 - d(t,m)^2}$ and $d(u,t) = \sqrt{d(u,c)^2 - d(t,c)^2}$ one deduce that $d(u,t) \leq \sqrt{(\rho + d(u,\partial B_{c,\rho}))^2 - \rho^2} = O(\sqrt{\varepsilon})$ when $t \in B_{m,\rho_m-\varepsilon}$ and in the other case $t \in B_{c,\rho}$ we get $d(u,t) \leq \sqrt{(\rho_m + d(\tilde{u},\partial B_{m,\rho_m-\varepsilon}))^2 - \rho_m^2} = O(\sqrt{\varepsilon})$. From this two bounds of the distance $d(u,t)$ we obtain that

$$\angle(c,m,u) \leq \arcsin\left(\frac{d(u,t)}{\rho_m - \sqrt{\rho_m^2 - d(u,t)^2}}\right) = O(\sqrt{\varepsilon})$$

$$\angle(m,c,u) \leq \arcsin\left(\frac{d(u,t)}{\rho - \sqrt{\rho^2 - d(u,t)^2}}\right) = O(\sqrt{\varepsilon})$$

Now consider $t \notin B_{m,\rho_m-\varepsilon} \bigcup B_{c,\rho}$. Let $t_m u_m$ with $u_m \in \partial B_{m,\rho_m}$ be a segment parallel to ut and intersecting the segments cm also let $t_c u_c$ be parallel segment to tu with $u_c \in \partial B_{c,\rho}$ and $t_c \in uc$. From this we get that

$$\angle(c,m,u) = \arcsin\left(\frac{d(u_m,t_m)}{\rho_m}\right) \quad (19)$$

$$\angle(m,c,u) = \arcsin\left(\frac{d(u_c,t_c)}{\rho}\right) \quad (20)$$

Due to $d(u_m,t_m) \leq \sqrt{(\rho_m - d(u,\partial B_{m,\rho_m-\varepsilon}))^2 - {\rho_m}^2} = O(\sqrt{\varepsilon})$ and $d(u_c,t_c) \leq \sqrt{(\rho - d(u,\partial B_{c,\rho}))^2 - \rho^2} = O(\sqrt{\varepsilon})$ we obtain that $\angle(c,m,u) = O(\sqrt{\varepsilon})$ and $\angle(m,c,u) = O(\sqrt{\varepsilon})$ □

Corollary 1 The angle of intersection between two balls $B_{c_1,\rho_1} \in \mathbb{B}_I$ and $B_{c_2,\rho_2} \in \mathbb{B}_I$ such that $B_{c_1,\rho_1} \cap B_{c_2,\rho_2} \cap N_\varepsilon \neq \emptyset$, is $\pi - O(\sqrt{\varepsilon})$.

Proof Take $x \in N_\varepsilon \cap B_{c_1,\rho_1} \cap B_{c_2,\rho_2}$. We have that $d(x,\partial B_{c_i,\rho_i}) = 0$ for $i = 1,2$.
Therefore, applying the lemma 11 we have the angle between the surface normal $\vec{n}_{\tilde{x}}$ and the vector $\vec{c_i x}$, is $O(\sqrt{\varepsilon})$ for $i = 1,2$, consequently the angle between the vectors $\vec{c_1 x}$ and $\vec{c_2 x}$ is $O(\sqrt{\varepsilon})$ and the angle between the tangent planes at x is $\pi - O(\sqrt{\varepsilon})$. □