Decimating Samples for Mesh Simplification

Tamal K. Dey   Joachim Giesen   James Hudson *

Ohio State University Columbus, OH 43210

Abstract

Modern scanning devices allow to obtain a dense sample of discrete points from the surface of a physical object. A piecewise linear surface interpolating these points is computed to *reconstruct* the sampled surface. Typically such surfaces have a large combinatorial description since the input is usually too dense. In this paper we present an algorithm to decimate the samples to eliminate oversampling. The algorithm decimates the sample with the guarantee that the remaining points are sufficient to reconstruct the surface and has a density controlled by an user input.

1 Introduction

A wide range of application domains and the modern laser technology have sparked new developments in modeling surfaces from samples. In a typical application the surface of a physical object is scanned with a scanning device such as laser range scanner that generates a dense set of discrete samples. Surface reconstruction is the problem of interpolating these samples with a piecewise linear surface that approximates the original surface. A number of algorithms with various guarantees and capabilities have been designed in recent years for the problem [1, 2, 3, 4, 5, 6, 7, 11, 14, 15, 17, 19].

Reconstructed surfaces are used for various purposes such as display in computer graphics, finite element methods in physical simulations or in computer aided manufacturing. Often the surface becomes unwieldy for these postprocessings because of their large combinatorial description. Naturally, a demand arose to simplify these surfaces in a meaningful way. A popular approach contracts the edges of a surface mesh for simplification, see [10, 16, 18, 20] for example.

Extra care is needed to ensure that edge contraction does not cause any major change in shape. This means topological and geometric features of the shape should be preserved during edge contractions. To this end, the topology preserving edge contractions of [12] and the quadric error method of [16] can be combined. However, these edge contractions can introduce self-intersection warranting a geometric validity check, which adds more complications.

In this paper we develop an alternative approach of decimating the samples before reconstruction as suggested first in [2]. Recently, in a simultaneous work Boissonnat and Cazals also consider decimating samples for simplification [9]. The advantage of this approach is that one does not need to worry about invalid geometry since the reconstructed surface is a subcomplex of the Delaunay triangulation of the decimated sample. In addition, if one makes sure that the samples are decimated up to a level so that the remaining samples are still sufficient for a meaningful reconstruction, topological and geometric features are automatically preserved. We demonstrate such an algorithm DECIMATE in this paper. It exploits the Voronoi cell structure to determine the sample points to be deleted and an input parameter ρ to determine the level of decimation. This algorithm is also useful in multiresolution modeling where a hierarchy of surface meshes can be computed out of a hierarchy of sample points.

Our algorithm has theoretical guarantees that we prove in an asymptotic framework. This asymptotic analysis assumes that the input sample is extremely dense and the target density to which the point set is to be decimated is extremely high. It is important to note that the constants hidden in the asymptotic notations depend on the sampled surface. However, they are independent of the input and target densities. We adopted the asymptotic analysis to ease the onus of calculating the constants rigorously. These constants are irrelevant for most practical data anyway. Nevertheless, the proofs embody the essential ingredients needed for the correctness of the algorithm and, in fact, guide its design. The practical validity of the algorithm is exhibited by our empirical results.

2 Voronoi cell geometry

The main tool we use to detect oversampling is the geometry of the Voronoi cell structure. Our reasoning is based on a density condition called ε-sampling as introduced in [2]. This definition builds on the *medial axis* of a surface and the related term of *local feature size*. The medial axis of a surface $S \subset \mathbb{R}^3$ is defined as the locus of all points that have more than one closest point on S. The local feature size at a point $x \in S$, denoted as $f(x)$, is the least distance of x to the medial axis. Let $P \subset \mathbb{R}^3$ be a sample of a surface S. We say P is an ε-sample of S if each point $x \in S$ has a sample within $\varepsilon f(x)$ distance. It has been observed that typically $\varepsilon < 0.4$

* This research is partially supported by NSF grant CCR-9988216.
e-mail:{tamaldey, giesen, jhudson}@cis.ohio-state.edu

gives a dense enough sampling for surface reconstruction. For a dense sample each Voronoi cell is long and thin as shown in Figure 1. We introduce a measure of this structural property that is easily computable. We need the following definitions for further expositions.

Definitions

Let V_P and D_P denote the Voronoi diagram and Delaunay triangulation of the sample $P \subset \mathbb{R}^3$ respectively. For a sample $p \in P$, let V_p be its Voronoi cell. For convenience we use the notation $\angle(\mathbf{v}, \mathbf{w})$ to denote the *acute* angle between the lines supporting two vectors $\mathbf{v}$ and $\mathbf{w}$.

Poles: The farthest Voronoi vertex p^+ in V_p is called the positive *pole* of p. The negative pole of p is the farthest point $p^- \in V_p$ from p such that two vectors from p to p^+ and p^- make an angle more than $\frac{\pi}{2}$. We call $\mathbf{v}_p = p^+ - p$, the *pole vector* for p. If V_p is unbounded, p^+ is taken at infinity, and the direction of $\mathbf{v}_p$ is taken as the average of all directions given by unbounded Voronoi edges.

Cocone: The set $C_p(\theta, \mathbf{v}) = \{y \in V_p \ : \ \angle((y - p), \mathbf{v}) \geq \frac{\pi}{2} - \theta\}$ is called the cocone of p with axis $\mathbf{v}$ and angle θ. In words, $C_p(\theta, \mathbf{v})$ is the complement of a double cone (clipped within V_p) centered at p with an opening angle $\frac{\pi}{2} - \theta$ around the axis aligned with $\mathbf{v}$. Usually, we take $\mathbf{v} = \mathbf{v}_p$, the *pole* vector and $\theta < \frac{\pi}{8}$ small. See Figure 1 for an example of a cocone.

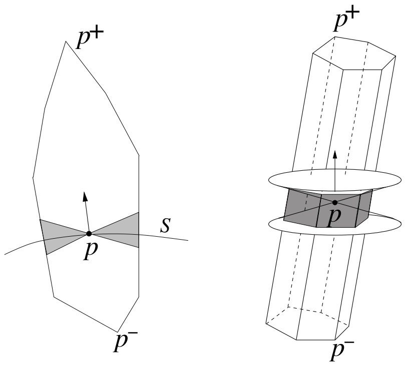

Figure 1: A Voronoi cell together with the normalized pole vector and the cocone (shaded) in two dimensions (left) and three dimensions (right).

Radius and height: The *radius* for a point p, r_p, is defined as the distance $\max\{||y - p|| \ : \ y \in C_p(\theta, \mathbf{v}_p)\}$. Define the *height* h_p as $||p^- - p||$. Obviously, $h_p \leq ||p^+ - p||$.

3 Algorithm DECIMATE

Our decimation scheme detects sample points that oversample the surface and then deletes them iteratively. Let d_p denote the maximum distance of any point from p on S that has p as its nearest sample. A point p lies in an oversampled region, if d_p is small compared to the local feature size $f(p)$. However, we cannot perform these comparisons exactly since S is unknown. Instead, we use the approximations to d_p and $f(p)$.

Recent results in [4, 8] show that poles approximate the medial axis. We can take the minimum among all distances of p to the poles as an approximation to $f(p)$. However, if one examines carefully, it can be deduced that the result of [4, 8] can be stated in asymptotic notation as in the next lemma. It implies that h_p can be taken as an approximation to $f(p)$ saving the computation for minimum distance to the poles. We emphasize that the big-O notation hides some surface dependent constants.

Lemma 1 *Let P be an ε-sample of a smooth surface S. Then, for any $p \in P$, $h_p \leq \lambda f(p)$ where $\lambda = 1 + O(\varepsilon)$ can be chosen as a constant.*

Let $V_{p,S} = V_p \cap S$ denote the restriction of the Voronoi cell V_p on S. The distance d_p of the farthest point on $V_{p,S}$ from p is approximated by the cocone radius r_p since cocone is a thin neighborhood around $V_{p,S}$ as proved in [3]. The following results from [2, 3] are used in our proofs in an asymptotic framework. The assumption is that there exist small enough ε, ρ, with $\varepsilon < \rho < 1$ for which proofs hold. We denote the normal vector at p on S as $\mathbf{n}_p$.

Lemma 2 *Let P be an ε-sample of a smooth surface S. Then $\angle(\mathbf{v}_p, \mathbf{n}_p) = O(\varepsilon)$.*

Lemma 3 *Let $\mathbf{v}$ be a vector with origin at p so that $\angle(\mathbf{v}, \mathbf{n}_p) = O(\varepsilon)$. Then there exists a constant c so that if P is an $O(\rho)$-sample and $\theta = c\rho$, a small angle, then (a) any point $x \in V_{p,S}$ satisfies $\angle((x - p), \mathbf{n}_p) \geq \frac{\pi}{2} - O(\rho)$, (b) $V_{p,S}$ lies inside $C_p(\theta, \mathbf{v})$, and (c) $r_p = O(\rho) f(p)$.*

We determine the samples responsible for oversampling by computing the ratio $\frac{r_p}{h_p}$ for each point p as follows. The height h_p is computed from the Voronoi diagram V_P. As P changes with deletion of points, so does V_P. We do not recompute h_p. The reason is that poles approximate the medial axis better with denser sample than the sparser ones; and we use h_p to approximate the distance of p to the medial axis. This also means that the cocone axis remains same for a point p throughout the decimation. However, $C_p(\theta, \mathbf{v})$ and hence its radius r_p changes as Voronoi diagram changes. If $\frac{r_p}{h_p}$ is less than the given parameter ρ, we delete p and update the Voronoi diagram. This process continues till we find no such p.

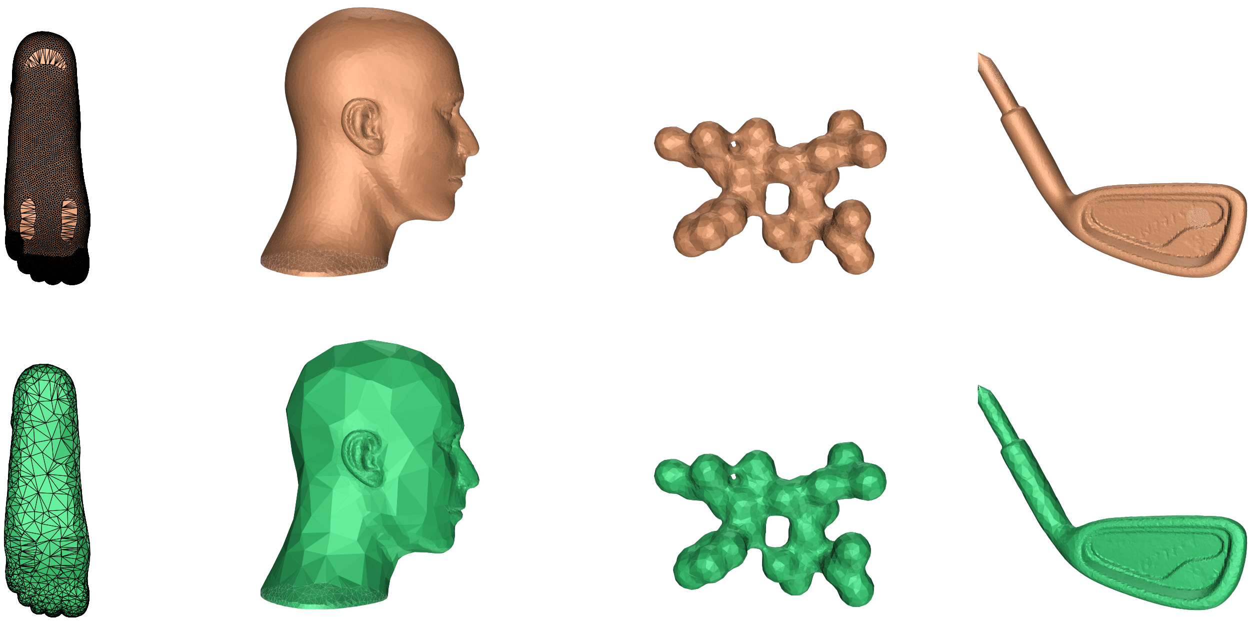

Figure 2: Data sets FOOT, MANNEQUIN, MOLECULE and CLUB. The yellow models (upper row) show reconstructions without any sample decimation. The green models (lower row) show reconstruction after DECIMATE.

DECIMATE(P, ρ)
1 **for** each $p \in P$ compute h_p
2 $Q := P$
3 **while** $\exists p \in Q$ so that $\frac{r_p}{h_p} < \rho$
4 delete p from Q.
5 update V_Q.
6 **if** V_q is changed for any $q \in Q$ update r_q
7 **endwhile**
8 **return** Q

We prove that, for an appropriate choice of ρ, DECIMATE produces a subset $Q \subset P$ so that Q is a dense enough sample for reconstruction. Additionally, for each remaining sample $p \in Q$ there is a point, say $x \in S$, so that x has p as its nearest sample and the distance $||x - p||$ is more than a constant factor of $\rho f(x)$. It means that p is necessary for x to maintain a sample density determined by the user parameter ρ. Although in theory we need an extremely small value of ρ, we obtain good results for ρ in the range of $0.25 - 0.4$.

4 Guarantees

The guarantees of DECIMATE are summarized in the following main theorem.

Theorem 1 *There exist $\varepsilon, \rho > 0$ so that if P is an ε-sample of a surface S with $0 < \varepsilon < \rho < 1$, P can be decimated to a new sample $Q \subset P$ with the guarantees that:*
(i) Q is a $c_1\rho$-sample of S,
(ii) each sample $q \in Q$ has a point $x \in S$ so that x has q as nearest sample and $||x - q|| > c_2\rho f(q)$ where c_1 and c_2 are constants independent of ε, ρ.

The guarantees (i) and (ii) are proved in Lemma 4 and 5 respectively. Let Q_i denote the point set after the *ith* deletion in DECIMATE.

Lemma 4 *Q_i is an ε_n-sample where $\varepsilon_n = O(\rho)$.*

PROOF. Assume inductively that Q_{i-1} satisfies the lemma for $\varepsilon_n = \frac{2\lambda\rho}{1-2\lambda\rho}$. Certainly, $Q_0 = P$ satisfies the condition for small ε. Let p be the point that is deleted in the *ith* deletion. Since we keep the cocone axis as $\mathbf{v}_p$, Lemma 2 and 3 assure that $V_{p,S}$ lies inside $C_p(\theta, \mathbf{v})$. Condition of the *while* loop requires that any point in $V_{p,S}$ is at a distance of $r_p \leq \rho h_p \leq \lambda\rho f(p)$ from p (Lemma 1). Therefore, any point $x \in V_{p,S}$ is at most $2\lambda\rho f(p)$ distance away from its nearest sample, say q, when p is deleted. We have $||x - q|| \leq 2\lambda\rho f(p) \leq \frac{2\lambda\rho}{1-2\lambda\rho} f(x)$. Thus, Q_i is an ε_n-sample where $\varepsilon_n = \frac{2\lambda\rho}{1-2\lambda\rho} = O(\rho)$.

Lemma 5 *Let p be any point in Q. There is a point $x \in V_{p,S}$ so that $||p - x|| > c\rho f(p)$ where c is a constant.*

PROOF. Since p does not satisfy the decimation condition in step 3, we must have $\frac{r_p}{h_p} > \rho$, or $r_p > \lambda\rho f(p)$. This implies that there exists a point $y \in C_p(\theta, \mathbf{v})$ so that $||p - y|| > \lambda\rho f(p)$. Let $x \in V_{p,S}$ be a point closest to y. It can be shown that $||x - y|| \leq c_1\rho^2 f(p)$ where Q is an $O(\rho)$-sample and c_1 is a constant. The fact that Q is an $O(\rho)$-sample follow from Lemma 4. Therefore, $||p - x|| > ||p - y|| - ||x - y|| = \lambda\rho f(p) - c_1\rho^2 f(p) > c\rho f(p)$ where c is a suitable constant.

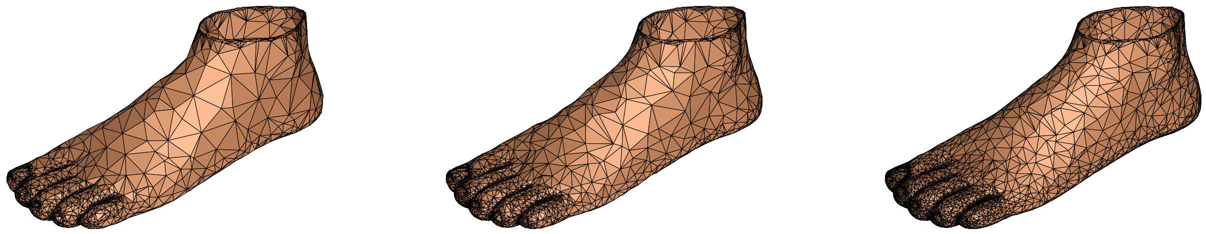

Figure 3: Decimation of the data set FOOT with DECIMATE with $\rho = 0.4$, 2046 points (left) $\rho = 0.33$, 2714 points (middle) and $\rho = 0.25$, 4116 points (right).

5 Experiments

We implemented DECIMATE in C++ using the CGAL library [21]. The surfaces are reconstructed by our COCONE software built on the work of [13]. The running times we report here are due to experiments performed on a PC with 933 Mhz Pentium III processor and 512 MB main memory.

Figure 2 shows the results on the four data sets FOOT, MOLECULE, MANNEQUIN and CLUB. We achieve almost 65% reduction in data points when $\rho = 0.33$ is used for these experiments; see Table 1. Increasing ρ to 0.4 achieves almost 80% reduction as shown in table 2. Different levels of decimations for FOOT with three different values of ρ are shown in Figure 3.

Data	#points before	#points after	sec.
MOLECULE	8862	2742	118
FOOT	20021	2714	417
CLUB	16864	5967	345
MANNEQUIN	12772	5038	251

Table 1: Results of DECIMATE with $\rho = 0.33$.

Data	#points before	#points after	sec.
MOLECULE	8862	1921	117
FOOT	20021	2046	408
CLUB	16864	4542	333
MANNEQUIN	12772	4272	242

Table 2: Results of DECIMATE with $\rho = 0.4$.

References

[1] U. Adamy, J. Giesen, M. John. New Techniques for Topologically Correct Surface Reconstruction. *Proc. IEEE Visualization*, (2000), 373–380.

[2] N. Amenta, M. Bern and M. Kamvysselis. A new Voronoi-based surface reconstruction algorithm. *SIGGRAPH 98*, (1998), 415-421.

[3] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Proc. 16th. ACM Sympos. Comput. Geom.*, (2000), 213–222.

[4] N. Amenta, S. Choi and R. K. Kolluri. The power crust, unions of balls, and the medial axis transform. *Manuscript*, 2000.

[5] D. Attali. r-regular shape reconstruction from unorganized points. *Proc. 13th ACM Sympos. Comput. Geom.*, (1997), 248–253.

[6] C. Bajaj, F. Bernardini and G. Xu. Automatic reconstruction of surfaces and scalar fields from 3D scans. *SIGGRAPH 95*, (1995), 109–118.

[7] F. Bernardini, J. Mittelman, H. Rushmeir, C. Silva and G. Taubin. The ball-pivoting algorithm for surface reconstruction. *IEEE Trans. Vis. Comput. Graphics*, **5**, no. 4, 349–359.

[8] J.-D. Boissonat and F. Cazals. Natural coordinates of points on a surface. *Intl. J. Comput. Geom. Appl.*, to appear.

[9] J.-D. Boissonnat and F. Cazals. Coarse-to-fine surface simplification with geometric guarantees. *Manuscript*, 2001.

[10] J. Cohen, A. Varshney, D. Manocha, G. Turk, H. Weber, P. Agarwal, F. Brooks and W. Wright. Simplification envelopes. *SIGGRAPH 96*, (1996), 119–128.

[11] B. Curless and M. Levoy. A volumetric method for building complex models from range images. *SIGGRAPH 96*, (1996), 303-312.

[12] T. K. Dey, H. Edelsbrunner, S. Guha and D. Nekhayev. Topology preserving edge contractions. *Publ. Inst. Math. (Beograd) (N.S.)*, **66**, (1999), 23–45.

[13] T. K. Dey and J. Giesen. Detecting undersampling in surface reconstruction. *Proc. 17th ACM Sympos. Comput. Geom.*, (2001), 257–263.

[14] T. K. Dey, J. Giesen and J. Hudson. A Delaunay based shape reconstruction from large data. *Proc. IEEE Sympos. in Large Data Visualization and Graphics*, (2001), to appear.

[15] H. Edelsbrunner and E. P. Mücke. Three-dimensional alpha shapes. *ACM Trans. Graphics*, **13**, (1994), 43–72.

[16] M. Garland and P. S. Heckbert. Surface simplification using quadric error metrics. *Proc. SIGGRAPH 97*, (1997), 209–216.

[17] M. Gopi, S. Krishnan and C. Silva. Surface reconstruction based on lower dimensional localized Delaunay triangulation. *Eurographics 2000*.

[18] H. Hoppe. Progressive meshes. *Proc. SIGGRAPH 96*, (1996), 99-108.

[19] H. Hoppe, T. DeRose, T. Duchamp, J. McDonald and W. Stuetzle. Surface reconstruction from unorganized points. *SIGGRAPH 92*, (1992), 71-78.

[20] G. Turk. Re-tiling polygonal surfaces. *Siggraph 92*, (1992), 55–64.

[21] `http://www.cgal.org`