Approximate Medial Axis for CAD Models *

[Extended Abstract]

Tamal K. Dey
Dept. of CIS
The Ohio State University
Columbus, OH 43210, USA
tamaldey@cis.ohio-state.edu

Hyuckje Woo
Dept. of CIS
The Ohio State University
Columbus, OH 43210, USA
wooh@cis.ohio-state.edu

Wulue Zhao
Dept. of CIS
The Ohio State University
Columbus, OH 43210, USA
zhaow@cis.ohio-state.edu

ABSTRACT

Several research have pointed out the potential use of the medial axis in various geometric modeling applications. The computation of the medial axis for a three dimensional shape often becomes the major bottleneck in these applications. Towards this end, in a recent work, we suggested an efficient algorithm that approximates the medial axis of a shape from a point sample. The input to this algorithm is only the coordinates of the sample points. As a result the approximation quality is limited by the input sample density. However, in geometric applications involving CAD models, the surfaces from which samples need to be derived are known. In this paper we show how to take advantage of this a priori knowledge in our medial axis approximation algorithm. This is achieved by first sampling the CAD surfaces appropriately and then modifying the medial axis approximation algorithm to exploit the known 'features' of the input surfaces. The quality of the approximation achieved by the method is surprisingly high as our experimental results exhibit.

Categories and Subject Descriptors

I.3.5 [**Computer Graphics**]: Computational Geometry and Object Modeling; I.4 [**Image Processing and Computer Vision**]: Reconstruction,Image Representation

General Terms

Algorithms, Experimentation, Theory, Design.

Keywords

Medial axis, Point cloud, Voronoi diagram.

*This work is partially supported by NSF under grant DMS-0138456 and subcontract from Stanford University and also the grant EIA-0102280.

1. INTRODUCTION

The medial axis of a three dimensional shape embedded in three dimensions is the closure of all points that have more than one closest point on the shape boundary. As a skeleton shape representation, medial axis has been used in various applications including image processing [21, 22], computer vision [8, 15], solid modeling [17, 23, 25, 26], mesh generation [24] and many others [16, 19, 27]. The medial axis together with the distance to the boundary at each point is called the medial axis transform. Since it essentially captures the shape information, the medial axis transform has been proposed as an alternate representation to B-rep or CSG system for shape manipulation [12, 18].

Although many research have identified the potential application of the medial axes in various shape manipulations, the difficulty in their computations has stymied the progress. The difficulty in computing medial axes can be attributed to their high algebraic degree and their sensitivity to small changes in shape [11, 28]. Both contribute to numerical instability and inefficiency in their computations. As a result, few algorithms that have been proposed to compute exact medial axes are limited to special class of shapes [12, 18]. This seriously restricts the applications of the medial axis. Consequently efforts have been made to approximate the medial axis.

Voronoi diagrams have been shown to be useful for approximating the medial axes. In two dimensions Brandt and Algazi [10] show that that the Voronoi vertices derived from a point sample of the boundary curve of a shape approximate the medial axis. Amenta, Bern and Eppstein [2] observed that this property does not hold in three dimensions. The circumcenters of the so called 'slivers' can be arbitrarily close to the surface no matter how dense the sample is. However, Amenta and Bern [1] show that certain Voronoi vertices called *poles* lie far away from the surface. Amenta, Choi and Kolluri [4] and Boissonnat and Cazals [9] show that the poles indeed lie close to the medial axis and converge to it as the sample density approaches infinity. Based on this observation Amenta, Choi and Kolluri also give an algorithm that connects the poles with a cell complex approximating the medial axis [4].

In another work, Attali and Lachaud [6] and Attali and Montanvert [7] show how to prune the Voronoi diagram with an angle and thickness criterion to approximate the medial axis. They observe that the noisy vertices of the medial axis have small bisector angle or small thickness. Conse-

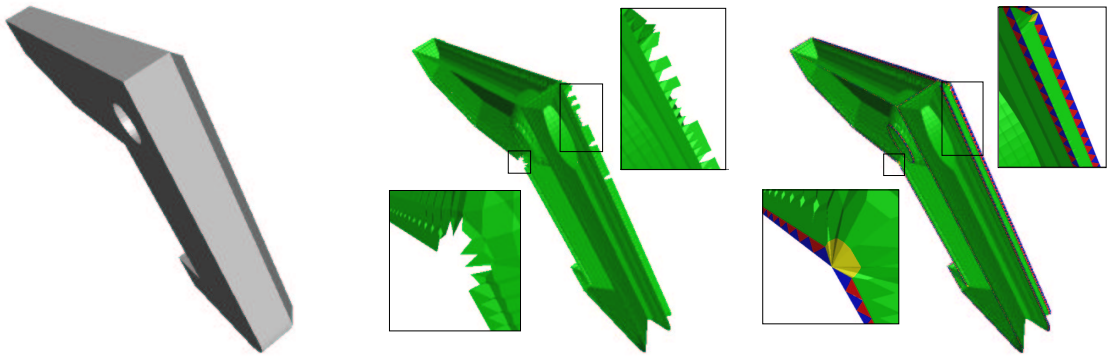

CLAW MEDIAL output New Algorithm

Figure 1: Computed medial axis of CLAW with MEDIAL and the new algorithm CAD_MEDIAL. The artifacts in the regions leading to the convex edges and non-concave corners are rectified with the new algorithm.

quently, these vertices are removed from the medial axis if they have bisector angle or thickness smaller than the user defined thresholds. Although it gives good results, the users have to fine tune the two parameters. The shock graph approach of Leymarie and Kimia [20] also relates to the Voronoi diagrams.

In a recent work [14], we also proposed a Voronoi based algorithm called MEDIAL that approximates the medial axis as a Voronoi subcomplex. This algorithm selects a set of Delaunay edges with two conditions called *Angle* and *Ratio* conditions and then output their dual Voronoi facets to approximate the medial axis. The algorithm works in scale and density independent manner so that no parameter is needed to be tuned. Also, a convergence guarantee between the output and the true medial axis is proved.

The common drawback of the Voronoi-based algorithms applied to CAD models is that the structure of the medial axis is not complete. The Voronoi-based algorithms work on the assumption that the sample is dense. However, with finite sampling, this density condition can never be satisfied at non-smooth edges and corners which are common to CAD models. The approximate medial axis with Voronoi facets stops before touching the convex edges and non-concave corners while the true medial axis does pass through them. For example, see Figure 1. Further, due to inherent undersampling the medial axis is approximated poorly near the sharp edges and corners. In some applications, especially in CAD/CAM, it is sometime needed that the medial axis is properly approximated even in the vicinity of sharp edges and corners. For example, in dimensional reduction [5], a medial vertex and two medial lines which connect to two vertices of a beam are used to identify a beam end. Applications such as decomposition of general solids into sub regions, and recognition of small features for suppression also need good approximation of the medial axis near edges and corners.

The work in this paper which aims to solve the above problems grows out of our previous work on medial axis approximation. The new algorithm constructs a complete structure of the medial axis and connect it to convex edges and non-concave corners. We achieve this by sampling the known CAD surface in a specific way and then using the information about the surface in approximating the medial axis from the Voronoi diagram. The new algorithm CAD_MEDIAL uses MEDIAL as basis and makes nontrivial extensions and modifications to complete the structure of the medial axis. Although we focus on the CAD objects, our algorithm works for other objects where its boundary is known a priori.

The rest of the paper is structured as follows. Section 2 provides some necessary definitions. Section 3 gives an overview of our original algorithm MEDIAL. Section 4 describes the difficulties we face for the CAD models and how we solve them with the new algorithm CAD_MEDIAL. Experimental results are shown in section 5. We conclude the paper in section 6.

2. PRELIMINARIES

Let P be a point sample from a compact surface $S \subset \mathbb{R}^3$ without boundary. A ball is called *medial* if it meets S only tangentially in at least two points. The medial axis M of S is the closure of the set of centers of all medial balls. Each point on S has two medial balls, one touching it from outside and the other touching it from inside. It follows that the line going through a point $p \in S$ and the centers of its medial balls is normal to S at p.

The surface S separates $\mathbb{R}^3$ into two open subsets S^+ and S^- where S^+ is bounded, S^- is unbounded and $\mathbb{R}^3 = S^+ \cup S^{-1} \cup S$. The medial axis M of S has two components, the *inner medial axis* which is the closure of $M \cap S^+$ and the *outer medial axis* which is the closure of $M \cap S^-$. We are interested in computing the inner medial axis which resides inside the solid bounded by S. The algorithm initially approximates the entire medial axis M and later separates out the inner medial axis from the outer one.

Obviously, a sample P needs to be dense enough to contain information about the features of S. Following Amenta et al. [2] we define the local feature size $f()$ as a function $f : S \to \mathbb{R}$ where $f(x)$ is the distance of the point $x \in S$ to the medial axis M. Intuitively, $f()$ measures how complicated S is locally. A sample is an ε-sample if each point $x \in S$ has a sample point within $\varepsilon f(x)$ distance. We have observed that $\varepsilon \leq 0.4$ is sufficient for approximating S with a piecewise linear surface in practice [13].

The Voronoi diagram and its dual, the Delaunay triangu-

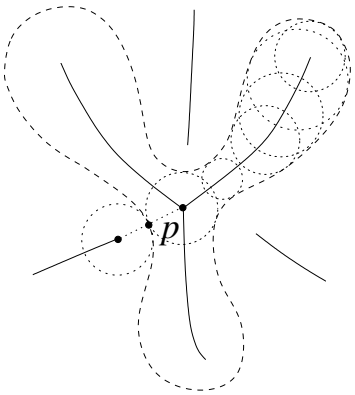

Figure 2: A curve and its medial axis in 2D.

lation are the main data structure used in our algorithms. The Voronoi diagram V_P for P is a cell complex consisting of Voronoi cells $\{V_p\}_{p \in P}$ and their facets, edges and vertices, where $V_p = \{x \in \mathbb{R}^3 \mid \|p - x\| \leq \|q - x\|, \forall q \in P\}$. The dual complex D_P of V_P, called the Delaunay triangulation of P, consists of Delaunay tetrahedra and their incident triangles, edges and vertices.

Amenta and Bern [1] proved that the Voronoi cells for a dense sample P of a surface are elongated along the normal directions to the surface. This motivates the following definitions of *poles* and *pole vectors*.

DEFINITION 1. *The* pole p^+ *for a sample p is the farthest Voronoi vertex from p in the Voronoi cell V_p. The vector $\mathbf{v}_p = p^+ - p$ is called the* pole vector. *In case V_p is unbounded, p^+ is taken at infinity and the pole vector is taken in the average direction of all unbounded Voronoi edges.*

It is proved in [1] that, up to orientation, the pole vector $\mathbf{v}_p$ approximates the normal $\mathbf{n}_p$ to the surface at p. In what follows we will use $\angle \mathbf{u}, \mathbf{v}$ to denote the acute angle between the supporting lines of two vectors $\mathbf{u}$ and $\mathbf{v}$.

3. ALGORITHM MEDIAL

Since we use the algorithm MEDIAL as the basis for the new algorithm, we describe its main concepts briefly. The algorithm MEDIAL was mainly designed for approximating the medial axes of smooth surfaces. It assumes that the input is a dense sample for sufficiently small value of $\varepsilon > 0$. Unfortunately, this sampling condition can never be achieved for CAD models which have non-smooth edges and corners. This warrants special action. Fortunately, we can leverage the knowledge of the input CAD surface into our new algorithm to handle the problem of non-smoothness.

The important observations used in MEDIAL are that the Voronoi facets lie close to the medial axis if their dual Delaunay edges tilt away from the surface or are very long. To filter the Delaunay edges that tilt away from the surface, we use an *Angle condition* while for the long Delaunay edges we use a *Ratio condition*. In order to measure the angle between the Delaunay edges and the surface, approximations to the tangent planes at the sample points are required. We use the pole vectors for this approximation.

DEFINITION 2. *The* tangent polygon *for a sample point is defined as the polygon in which the plane through p with normal $\mathbf{v}_p$ intersects V_p.*

Since $\mathbf{v}_p$ approximates the normal $\mathbf{n}_p$, the tangent polygon approximates the tangent plane at p restricted within V_p. It turns out that a dual structure to the tangent polygon from the Delaunay triangulation D_P is useful for checking both the Angle and Ratio conditions.

DEFINITION 3. *The* umbrella U_p *for a sample point p is defined as the topological disc made by the Delaunay triangles incident to p that are dual to the Voronoi edges intersected by the tangent plane.*

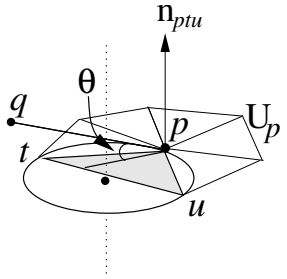

Figure 3: The angle θ for the Angle condition $[\theta]$. The Voronoi edge shown with the dotted line is parallel to the normal of the shaded triangle.

We take umbrella triangles in U_p for approximating the tangent plane at the sample point p and determine all Delaunay edges pq that make relatively large angle with the planes of the triangles. To compute this we measure the angle $\angle \mathbf{n}_\sigma, \mathbf{t}_{pq}$ between the vector $\mathbf{t}_{pq}$ from p to q and the normal $\mathbf{n}_{ptu}$ to a triangle ptu in U_p. We select those edges that make this angle less than a threshold angle $\frac{\pi}{2} - \theta$ for all umbrella triangles in U_p for $0 < \theta < \frac{\pi}{2}$, see Figure 3. This means a selected edge pq satisfies:

Angle condition $[\theta]$: $\max_{ptu \in \mathrm{U}_p} \angle \mathbf{n}_{ptu}, \mathbf{t}_{pq} < \frac{\pi}{2} - \theta$.

To filter the Delaunay edges which are relatively long, we compare the length of the Delaunay edges with the circumradii of the umbrella triangles. The rationale behind this choice is explained in our previous paper [14]. Let R_{ptu} denote the circumradius of a triangle ptu, refer to Figure 4. We measure the ratio of the length of a Delaunay edge to the circumradii of the umbrella triangles and consider those that satisfy:

Ratio condition $[\rho]$: $\min_{ptu \in \mathrm{U}_p} \frac{\|p-q\|}{R_{ptu}} > \rho$.

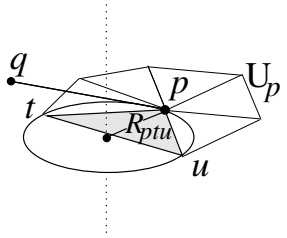

Figure 4: Radius of interest for the Ratio condition $[\rho]$.

Intuitively, the Angle condition captures the Delaunay edges that lie away from the surface, and the rest of the

Delaunay edges whose duals lie close to the medial axis but make small angles with the surface are captured by the Ratio condition. As shown earlier [14], the Angle and Ratio condition individually is not sufficient for a good medial axis approximation. Combining these two conditions we get the algorithm MEDIAL.

```
MEDIAL(P)
 1 Compute V_P and D_P;
 2 F = ∅;
 3 for each p ∈ P
 4    Compute U_p;
 5    for each Delaunay edge pq ∈ U_p
 6       if pq satisfies Angle Condition [θ]
             or Ratio condition [ρ]
 7          F := F ∪ Dual pq
 8       endif
 9    endfor
10 endfor
11 output closure(F)
```

The values of θ and ρ remained fixed for all tested data sets. Our experiments suggested that $\theta = \frac{\pi}{8}$ and $\rho = 8$ are appropriate for all tested data sets. Using the properties of the Voronoi diagram and ε-sampling we proved the following result [14].

THEOREM 1. *Let F_ε be the subcomplex computed by* MEDIAL *for an ε-sample of a surface $S \subseteq \mathbb{R}^3$ without boundary. We have $\lim_{\varepsilon \to 0} F_\varepsilon = M$, where M is the medial axis of S.*

4. ALGORITHM CAD_MEDIAL

The algorithm MEDIAL produces nice results if the input point set is sufficiently dense. Usually the surfaces of the CAD objects are non-smooth and it is impossible to sample non-smooth regions satisfying the ε-sampling condition for any $\varepsilon > 0$. As a result, the approximation becomes poor near these regions. We use the knowledge of the input surface to solve this problem.

4.1 Difficulties

4.1.1 Completion

It follows from the definition that the medial axis of a non-smooth surface passes through the points where the surface is non-smooth. We would be interested in the inner medial axis, i.e., the subset of the medial axis enclosed by the surface. This inner medial axis passes through sharp edges that are convex and corners that are non-concave. The algorithm MEDIAL approximates the medial axis with the Voronoi facets which cannot meet sample points. Consequently, the approximated medial axis is not *complete*. A 2D example is shown in Figure 5. The medial axis is approximated by the Voronoi edges near the corner cannot meet the corner point though it should.

4.1.2 Normal estimation

The next difficulty we face is that the normal estimation at the sample points near non-smooth regions is poor. Recall that the pole vectors estimate true normals only if the sampling density is sufficiently high. But, since this density cannot be achieved by finite sampling in the neighborhoods of non-smooth regions, we face a difficulty. Both Angle and Ratio conditions are affected by poor normal estimation.

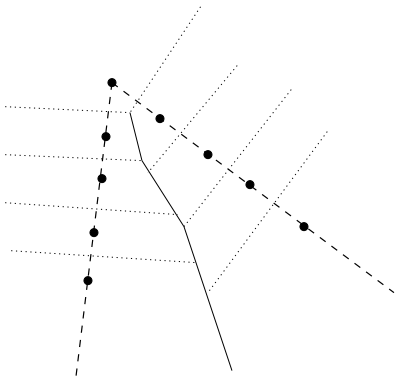

Figure 5: Approximated medial axis shown with solid Voronoi edges does not pass through the corner point.

4.1.3 Separation

The algorithm MEDIAL approximates the entire medial axis including both inner and outer ones. We separate the inner medial axis from the outer one by first reconstructing a water-tight surface from the input point sample and then retaining only that part of the approximated medial axis which resides inside the water-tight surface. This strategy works well with smooth surfaces. If we apply the same strategy in case of CAD models, we face the difficulty that the reconstructed water-tight surface may have small artifacts near non-smooth regions. These artifacts also propagate to the approximated medial axis. Since we aim for a complete medial axis as accurate as possible, we seek a better solution to this problem.

4.2 Remedies

The undersampling caused by non-smooth edges and corners is the source of all problems in approximating the medial axis for CAD surfaces. Therefore, we pay special attention while sampling the vicinity of these regions. First of all, we can sample non-smooth edges and vertices since the CAD surface is known. Secondly, we can choose the sample placements so that the medial axis can be completed by connecting it to proper edges and corners.

4.2.1 Sampling

First, we make sure that the non-smooth edges and vertices can be recovered from the Delaunay triangulation, i.e. if p and q are two consecutive sample points on a non-smooth edge, then the edge pq should be present in the Delaunay triangulation. This needs that the sampling is dense enough in the following sense. Let $r_1 > 0$ be such that any ball with radius smaller than r_1 and with a center on a non-smooth edge e intersects the surface in a topological disk and does not intersect any vertex or facet that is not incident on e. Similarly, let $r_2 > 0$ be such that any ball with radius smaller than r_2 and centering a non-smooth vertex v intersects the surface in a topological disk and does not intersect any vertex, edge, or facet that is not incident to v. We call $r = \min\{r_1, r_2\}$ the *minimum size* of the input surface.

For each non-smooth vertex v we put a ball around it with radius less than r. The intersection point, say p, of this ball with each edge incident to v is introduced in the

sample. No other sample point is introduced inside this ball making it empty. Thus, all edges between v and intersection points such as p appear in the Delaunay triangulation of the sample. The rest of the non-smooth edges are sampled at regular interval of distance less than $\frac{r}{2}$. This guarantees that each edge pq between two consecutive sample points on non-smooth edges has a diametric ball which is empty of any other sample point. Therefore, pq appears in the Delaunay triangulation. In our implementation we avoided computing the minimum size by sampling the sharp edges and corners densely with conservative estimates.

The structure of the medial axis will be complete if we can identify the Voronoi facets that should be connected to the convex edges and non-concave corners. Keeping this in mind we sample the convex edges and its vicinity in the following way. Let f_l and f_r be two facets incident to a convex edge e. First we place sample points sufficiently densely on e as described before. For each such sample s, we compute two points s_l and s_r on f_l and f_r respectively in such way that the segments $s_l s$ and $s_r s$ are perpendicular to e at s and their Euclidean lengths are equal. Call the three sample points s, s_l and s_r a *triplet*, and s_l and s_r the neighbors of s. See Figure 6.

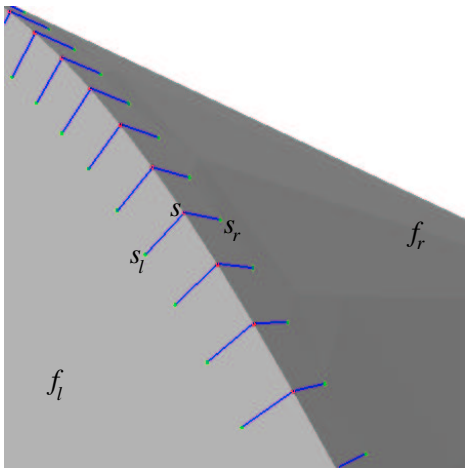

Figure 6: Edge sampling. The red points are samples on the edge and the blue line segments connect the corresponding neighbors to samples.

4.2.2 *Convex edges*

Since the sampling respects the minimum size, the triangle $ss_r s_l$ with a triplet has an empty circumcircle. This means that it is a Delaunay triangle. If the dihedral angle at the convex edge e is not too shallow, the Delaunay edge $s_r s_l$ is selected by the Angle condition meaning that its dual Voronoi facet participates in the approximate medial axis. In fact, if the dihedral angle at the edge e is more than $\pi - 2\theta$, the Angle condition with the threshold angle θ cannot capture the Delaunay edge $s_r s_l$. We have observed that some of the convex edges of CAD surfaces can have dihedral angles as large as 170°. As a result we have lowered down the threshold of θ to 5° for the new algorithm.

The dual Voronoi edge of the Delaunay triangle $ss_r s_l$ is an edge of the Voronoi facet dual to $s_r s_l$. Let v_1 and v_2 be the two end points of this dual Voronoi edge. The triangle sv_1v_2 is placed to connect the medial axis to the sample point s on the convex edge e. All such triangles form a zigzag shape (marked as red triangles in Figure 7). We fill the triangular holes left (marked as blue triangles) completing the connection of the medial axis to the convex edge e except possibly at the ends.

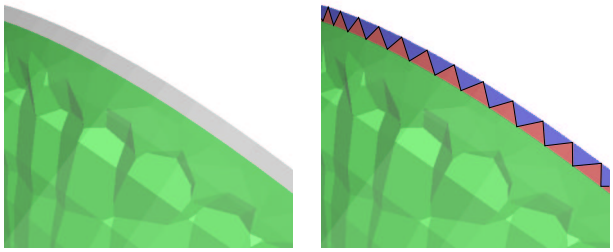

Figure 7: The picture in left shows the medial axis before connecting the Voronoi facets to the edge. The right picture shows how they are connected.

Let F be the collection of Voronoi facets chosen by Angle and Ratio conditions. Let s' and s'' be the two adjacent sample points of s on the convex edge. We summarize the above steps as follows.

ProcConvexEdges(P,F)
1 for each triplet $ss_l s_r$
2 find the dual Voronoi edge v_1v_2 of $ss_l s_r \in D_P$;
3 $F := F \cup sv_1v_2$
4 if $v_1 \in V_{s'}$ and $v_2 \in V_{s''}$
$F := F \cup ss'v_1$ and $F := F \cup ss''v_2$;
5 end for

Notice that the true medial axis bisects the dihedral angle formed by the two facets f_r and f_l at e. Since we choose the lengths of ss_r and ss_l equal, its dual Voronoi facet nicely approximates the medial axis. This Voronoi facet lies on the bisector plane of $s_r s_l$ and passes through the point s. So the triangle sv_1v_2 joins the dual Voronoi facet of $s_l s_r$ smoothly.

4.2.3 *Non-concave corners*

The non-concave corners usually connect to several branches of the medial axis. We have a sample point on every corner. Consider the Voronoi cell V_p of a non-concave corner p. Observe that some Voronoi edges in V_p are incident to the Voronoi facets that are chosen to approximate the medial axis as shown in the left picture of Figure 8. We need to identify these Voronoi edges and connect them to p to complete the medial axis as shown in the right picture of Figure 8.

However not all the Voronoi edges of V_p appearing on the approximated medial axis should be connected to p. We find that if a Voronoi facet of V_p is chosen in the medial axis, then none of its edges should be connected to p. In Figure 9 we show an example. The left picture shows that all the Voronoi edges of the chosen Voronoi facets are connected to the corner (yellow triangles). We can see that this is obviously not correct because there is a branch which should not meet the corner though it shares a Voronoi facet with the Voronoi cell of the corner. The right picture shows the correct result.

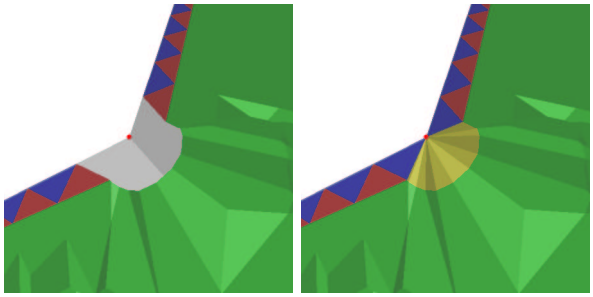

Figure 8: Nonconcave corners handling

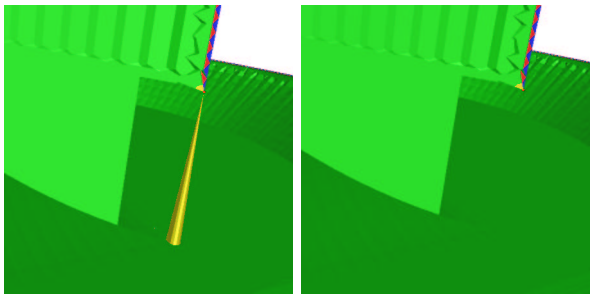

Figure 9: Non-concave corners special case

PROCNONCONCAVECORNERS(P, F)
1 for each non-concave corner p ;
2 for each Voronoi edge v_1v_2 on the medial axis;
3 if no Voronoi facet of V_p incident to v_1v_2 is in F
4 $F := F \cup pv_1v_2$;
5 end if
6 end for
7 end for

4.2.4 *Umbrella*

As we pointed out earlier, the umbrella for sample points in the vicinity of the non-smooth regions may deviate considerably from the tangent plane. This is mainly because the normal estimation suffers in these regions. As a remedy we redefine the umbrellas as follows.

Consider the sample points on a facet of the model. We can reconstruct the facet from its sample points using our COCONE algorithm [13]. This algorithm can reconstruct surfaces with boundaries from a dense sample. Since each individual facet can be thought of as a smooth surface with boundaries, the sample points on each facet constitute a sufficiently dense sample. Therefore, each facet is reconstructed accurately with our COCONE algorithm. For details of this algorithm see [13].

Thus, we reconstruct the entire input CAD surface facet by facet and stitching them together along the non-smooth edges and corners. In our actual algorithm we do not need this explicit stitching. The set of triangles incident to a sample point over all facet reconstructions constitute its umbrella. Notice that, since facet reconstructions are accurate, the umbrellas as defined above approximate the surface and hence the tangent planes quite accurately.

The left and right pictures in Figure 10 shows the medial axis approximation by using the old and new definitions of

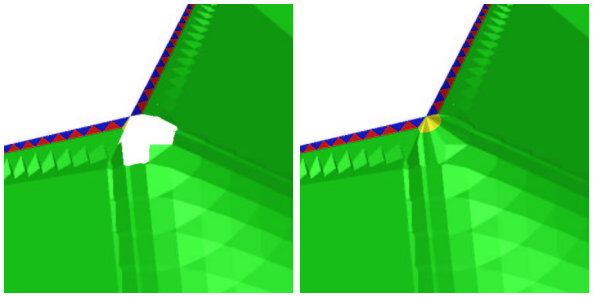

Figure 10: The effect of different umbrellas on the medial axis

the umbrella respectively. We can see that the hole near the corner point gets shrunk when the new definition of umbrella is used.

4.2.5 *Inner medial axis*

We are interested in the inner medial axis. We know that for a CAD model with a single boundary, each of the component of the outer medial axis must extend to infinity. This means each component in the approximation of the outer medial axis must have an unbounded Voronoi facet. Further, the inner and outer medial axis can only be connected through non-concave corner points. We use these facts to separate the inner medial axis from the outer one.

We start a walk from an infinite Voronoi facet that is on the medial axis and then continue collecting all other facets in the outer medial axis without ever going through the surface triangles around a non-concave corner point. If some infinite Voronoi facets chosen for the medial axis are left after this walk, we start a new walk from one of these infinite facets. This continues until we are not left with any infinite Voronoi facet that has been chosen for the medial axis approximation. At the end of this process we only collect the Voronoi facets approximating the outer medial axis. The rest of the Voronoi facets and the forced triangles are output as the inner medial axis.

4.2.6 *New algorithm*

With all the modifications described above the steps of the new algorithm CAD_MEDIAL can be enumerated as follows.

CAD_MEDIAL(S)
1 Get P by sampling S appropriately;
2 Compute D_P and V_P ;
3 Reconstruct S facet by facet;
4 Let U_p be the set of triangles incident to p in this reconstruction;
5 for each Delaunay edge $pq \in U_p$
6 if either p or q satisfies Angle or Ratio condition
7 $F := F \cup \text{Dual}\, pq$
8 endif
9 endfor
10 ProcConvexEdges(P,F);
11 ProcNonconcaveCorners(P,F);
12 Extract inner medial axis into F;
13 output closure(F);

5. EXPERIMENTAL RESULTS

Figure 11 shows our results on some example CAD objects. For all the data sets, CAD_MEDIAL constructed a complete medial axis in terms of the structure. We used $\rho = 8$ and $\theta = 5^\circ$ for all models.

In general, the medial axis of a surface is a CW-complex where each cell can be zero, one or two dimensional. In CAD_MEDIAL, we approximate all cells with two dimensional cells, namely the Voronoi facets. The medial axis for the data sets CHUCK, MAILBOX and TAPBOLT contains one dimensional components along the cylinders which are approximated by thin Voronoi facets and are too thin to be visible in the picture.

CAD_MEDIAL works for complicated shapes as the examples in Figure 11 exhibit. The QUIDE example shows the result for objects with smooth parametric facets. CAP shows objects without corners. ANCHOR shows a complex shape.

Figure 12 shows the medial axis with different sampling density. We approximate the medial axis with piecewise linear polygon, i.e. Voronoi facets. The geometric accuracy of the approximation increases as the density increases. With more sample points, we can get smoother approximation of the high curvature regions by smaller Voronoi facets, see the middle picture. However, this high geometric accuracy is not necessary for some applications. In these cases low sample density can be used. The right picture in the figure shows a low resolution version of the medial axis with smaller sample. The number of Voronoi facets decreases considerably, while the size of the Voronoi facets increases which can be seen from the magnified pictures. But if the density is too low, the artifacts such as holes will appear in the medial axis.

All of our softwares are written in C++. We used the CGAL library [29] for the Voronoi and Delaunay computation and used filtered floating point arithmetic for robust geometric computations. Experiments were conducted on a PC with 933 MHz CPU and 512MB memory. The code was compiled with CGAL2.3 library and g++ compiler with `O1` level of optimization. The time for the Delaunay triangulation, reconstruction by facets and Delaunay edge selection are listed in Table 1.

object	# points	Delaunay time(sec.)	Recon. time(sec.)	Filtering time(sec.)
ANCHOR	27,619	129.66	75.6	2.62
CAP	17,805	72.95	56.55	1.43
CHUCK	46,906	301.63	148.08	3.68
MAILBOX	6,980	29.51	27.61	0.56
QUIDE	14,951	52.52	44.16	1.13
TAPBOLT	47,870	241.84	263.3	4.04
WIDGET	66,686	440.53	300.43	5.83

Table 1: Time data.

6. CONCLUSIONS

In this paper, we extended a general purpose medial axis algorithm to three dimensional CAD objects. Several difficulties that one encounters with non-smooth boundaries in CAD objects are addressed successfully with the new algorithm. The original algorithm does not connect the medial axis to the convex edges and non-concave corners which it should to complete the structure. We solved this problem by extracting the sharp features information from the CAD objects, then forcing the connection of the medial axis to the convex edges and non-concave corners. The second problem due to poor normal estimations near the non-smooth regions is solved by reconstructing the surface 'perfectly' through the facet by facet reconstruction. The third problem of separating the inner medial axis from the outer one is solved by a simple walk over the approximated medial axis.

We have indicated that the sampling density near the non-smooth regions should be sufficiently high so that all required edges appear in the Delaunay triangulation. We have also suggested how one can guarantee this property. However, this guarantee requires some extra computation that we have avoided by simply being conservative during sampling. But as more samples increase the computation time and may create unnecessary burdon for further applications, it would be worthwhile to sample as much as needed for the good approximation. Can this be determined a priori efficiently? We plan to look into this problem in future.

7. REFERENCES

[1] N. Amenta and M. Bern. Surface reconstruction by Voronoi filtering. *Discr. Comput. Geom.* **22** (1999), 481–504.

[2] N. Amenta, M. Bern and D. Eppstein. The crust and the β-skeleton: combinatorial curve reconstruction. *Graphical Models and Image Processing*, **60** (1998), 125-135.

[3] N. Amenta, S. Choi, T. K. Dey and N. Leekha. A simple algorithm for homeomorphic surface reconstruction. *Internat. J. Comput. Geom. Applications*, **12** (2002), 125–141.

[4] N. Amenta, S. Choi and R. K. Kolluri. The power crust. *Proc. Solid Modeling '01*, (2001), 249–260.

[5] C. G. Armstrong, D.J. Robinson, R. M. McKeag, T. S. Li, S. J. Bridgett, R. J. Donaghy and C. A. McGleenan. Medials for Meshing and More. *Proc. 4th Internat. Meshing Roundtable*, (1995), 277–288.

[6] D. Attali and J.-O. Lachaud. Delaunay conforming iso-surface, skeleton extraction and noise removal. *Comput. Geom. : Theory Appl.*, 2001, to appear.

[7] D. Attali and A. Montanvert. Computing and simplifying 2D and 3D continuous skeletons. *Computer Vision and Image Understanding* **67** (1997), 261–273.

[8] S. Bouix and K. Siddiqi. Divergence-based medial surfaces. *Proc. European Conference on Computer Vision*, 2000.

[9] J. D. Boissonnat and F. Cazals. Smooth surface reconstruction via natural neighbor interpolation of distance functions. *Proc. 16th. ACM Sympos. Comput. Geom.*, (2000), 223–232.

[10] J. W. Brandt and V. R. Algazi. Continuous skeleton computation by Voronoi diagram. *Comput. Vision, Graphics, Image Process.* **55** (1992), 329–338.

[11] S. W. Choi and H.-P. Seidel. Linear one-sided stability of MAT for weakly injective 3D domain. *Proc. 7th ACM Sympos. Solid Modeling Applications*, (2002), 344–353.

[12] T. Culver, J. Keyser and D. Manocha. Accurate computation of the medial axis of a polyhedron. *Solid Modeling '99*, (1999), 179–190.

[13] T. K. Dey and J. Giesen. Detecting undersampling in

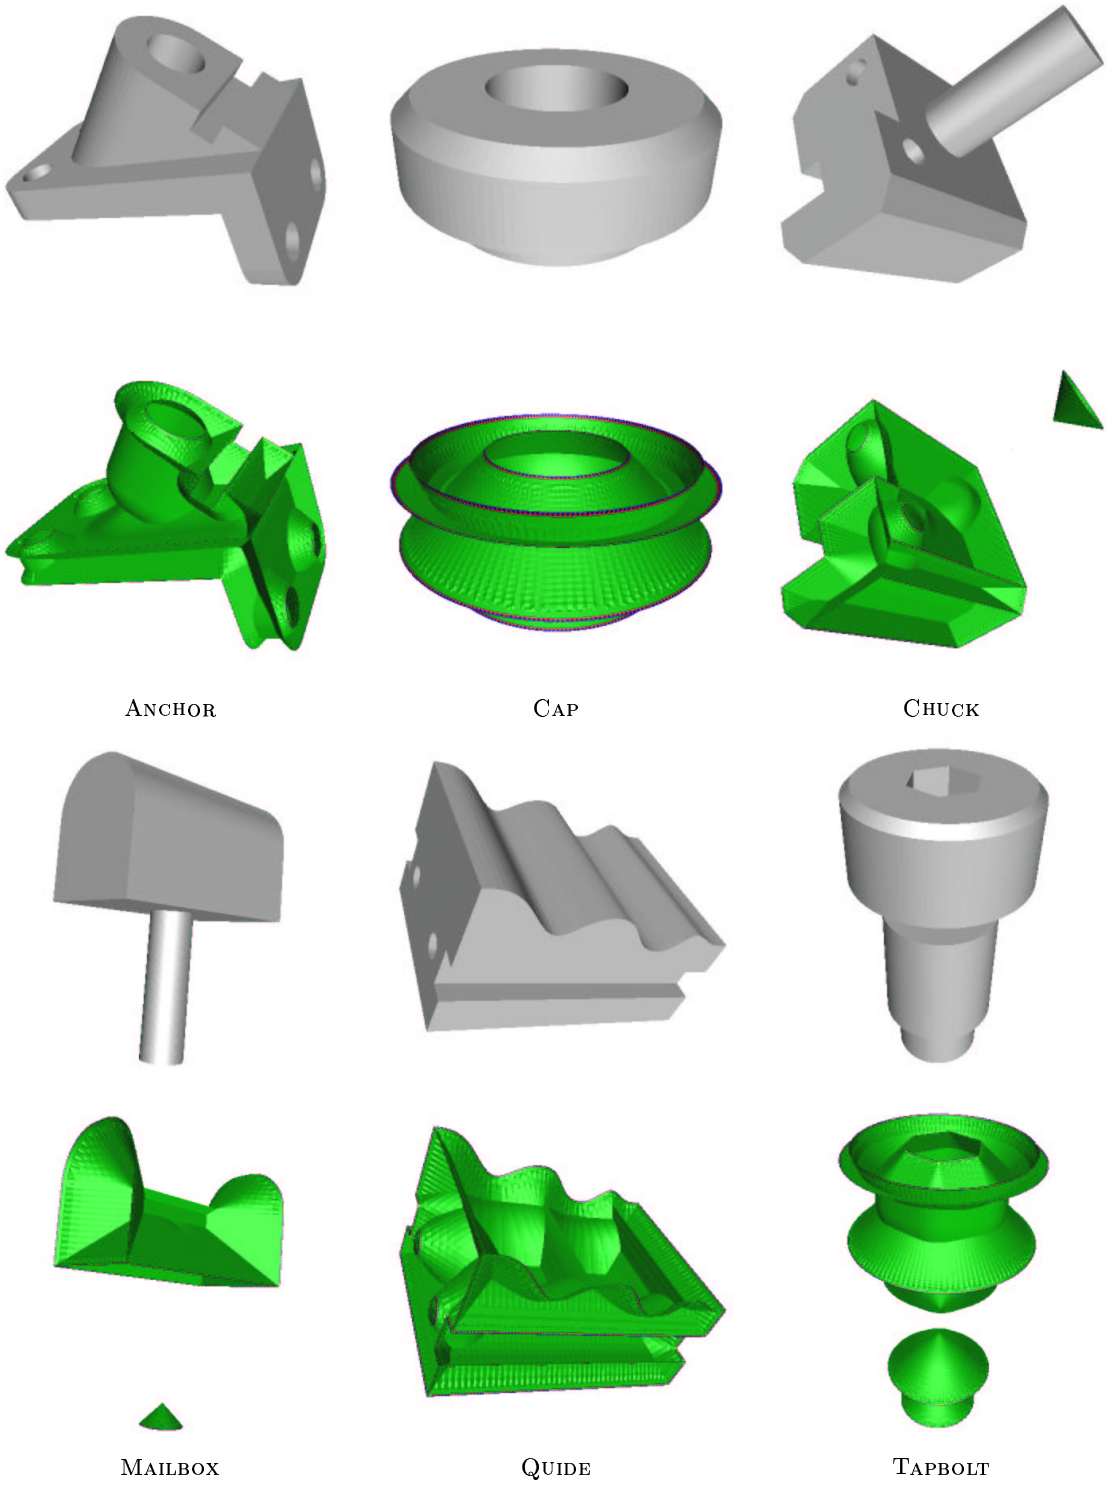

Figure 11: CAD_Medial **output: surfaces shown in the upper row, corresponding medial axes shown in the lower row.**

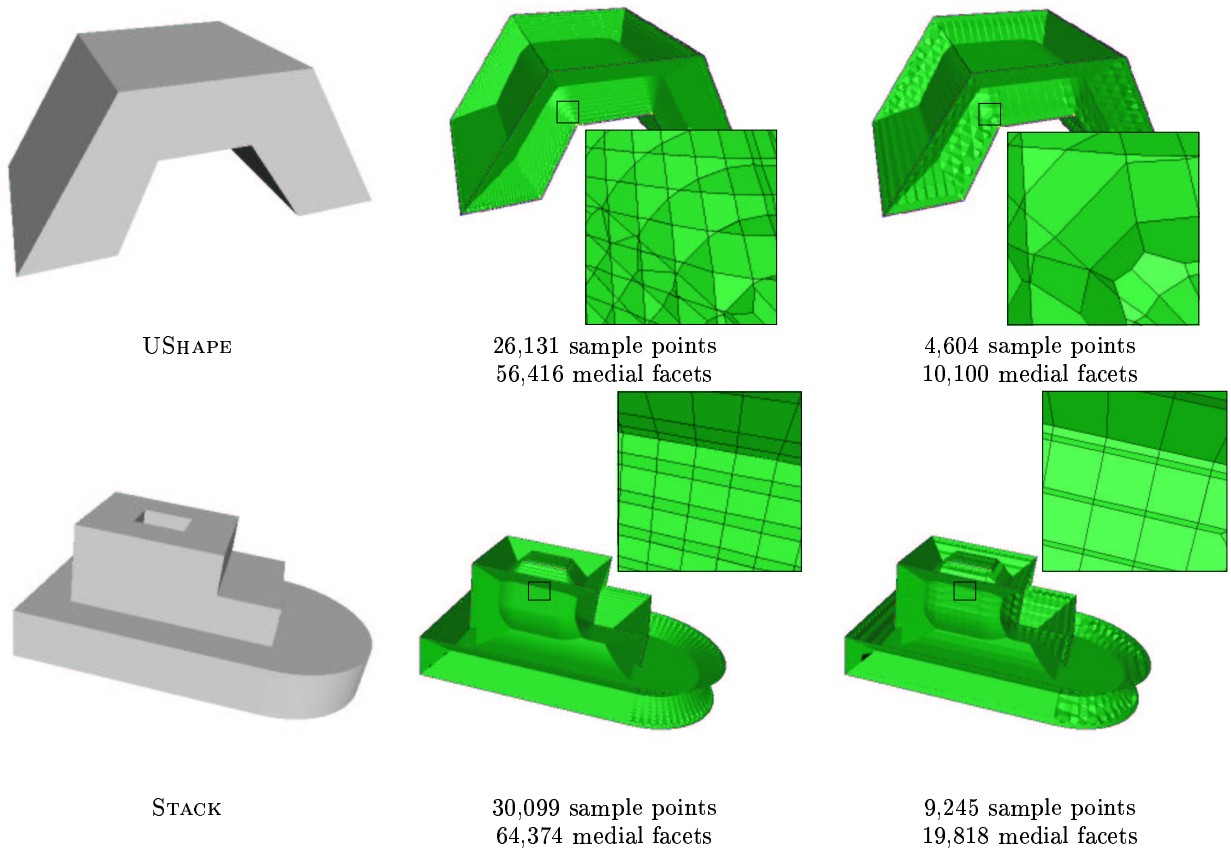

Figure 12: Computed medial axis at different sampling density.

surface reconstruction. *Proc. 17th Ann. Sympos. Comput. Geom.* (2001), 257–263.

[14] T. K. Dey and W. Zhao. Approximating medial axis from the Voronoi diagram with a convergence guarantee. *Algorithmica*, (2002), to appear. http://www.cis.ohio-state.edu/~tamaldey/paper/medial.pdf.

[15] P. J. Giblin and B. B. Kimia. A formal classification of 3D medial axis points and their local geometry. *Proc. Computer Vision and Pattern Recognition (CVPR)*, 2000.

[16] L. Guibas, R. Holleman and L. E. Kavraki. A probabilistic roadmap planner for flexible objects with a workspace medial axis based sampling approach. *Proc. IEEE/RSJ Intl. Conf. Intelligent Robots and Systems*, 1999.

[17] H. N. Gursoy and N. M. Patrikalakis. Automated interrogation and adaptive subdivision of shape using medial axis transform. *Advances in Engineering Software* **13** (1991), 287–302.

[18] C. Hoffman. How to construct the skeleton of CSG objects. *The Mathematics of Surfaces, IVA*, Bowyer and J. Davenport Eds., Oxford Univ. Press, 1990.

[19] P. Hubbard. Approximating polyhedra with spheres for time critical collision detection. *ACM Trans. Graphics* **15** (1996), 179–210.

[20] F. Leymarie and B. Kimia. The shock scaffold for representing 3D shape. *Proc. 4th Internat. Workshop Visual Form.*, (2001), LNCS 2059, Springer-Verlag, 216–229.

[21] C. Niblak, P. Gibbons and D. Capson. Generating skeletons and centerlines from the distance transform. *CVGIP: Graphical Models and Image Processing* **54** (1992), 420–437.

[22] R. L. Ogniewicz. Skeleton-space: A multiscale shape description combining region and boundary information. *Proc. Computer Vision and Pattern Recognition*, (1994), 746–751.

[23] D. Sheehy, C. Armstrong and D. Robinson. Shape description by medial axis construction. *IEEE Trans. Visualization and Computer Graphics* **2** (1996), 62–72.

[24] A. Sheffer, M. Etzion, A. Rappoport and M. Bercovier. Hexahedral mesh generation using the embedded Voronoi graph. *Engineering Comput.* **15** (1999), 248–262.

[25] E. C. Sherbrooke, N. M. Patrikalakis and E. Brisson. An algorithm for the medial axis transform of 3D polyhedral solids. *IEEE Trans. Vis. Comput. Graphics* **2** (1996), 44–61.

[26] D. Storti, G. Turkiyyah, M. Ganter, C. Lim and D. Stal. Skeleton-based modeling operations on solids. *Solid Modeling '97* (1997), 141–154.

[27] M. Teichman and S. Teller. Assisted articulation of closed polygonal models. *Proc. 9th Eurographics Workshop on Animation and Simulation*, 1998.

[28] F.-E. Wolter. Cut locus & medial axis in global shape interrogation & representation. *MIT Design Laboratory Memorandum 92-2*, 1992.

[29] www.cgal.org.

[30] www.cis.ohio-state.edu/~tamaldey/cocone.html.