

Homotopy-Preserving Medial Axis Simplification

Avneesh Sud*

Mark Foskey†

Dinesh Manocha‡

Department of Computer Science

Department of Radiation Oncology

Department of Computer Science

University of North Carolina at Chapel Hill

Abstract

We present a novel algorithm to compute a simplified medial axis of a polyhedron. Our simplification algorithm tends to remove unstable features of Blum’s medial axis. Moreover, our algorithm preserves the topological structure of the original medial axis and ensures that the simplified medial axis has the same homotopy type as Blum’s medial axis. We use the separation angle formed by connecting a point on the medial axis to closest points on the boundary as a measure of the stability of the medial axis at the point. The medial axis is decomposed into its parts that are the sheets, seams and junctions. We present a stability measure of each part of the medial axis based on separation angles and examine the relation between the stability measures of adjacent parts. Our simplification algorithm uses iterative pruning of the parts based on efficient local tests. We have applied the algorithm to compute a simplified medial axis of complex models with tens of thousands of triangles and complex topologies.

Keywords: Medial Axis, Voronoi diagram, homotopy, simplification, separation angle

1 Introduction

The medial axis of a geometric object is the set of interior points that have at least two closest points on the boundary of the object. The medial axis can also be defined as the set of centers of at least twice tangent maximal balls contained inside the object. This formulation was originally proposed by Blum [1967] and many authors have proposed extensions to this formulation. Given a 3D solid, its medial axis consists of a union of surfaces that provide useful information about its shape and topology. The medial axis transform (MAT) consists of all the medial axis points and the distance to the boundary from each medial point. MAT has applications in image analysis and computer vision [Pizer et al. 2003], solid modeling [Blanding et al. 1999], mesh generation and finite element analysis [Sheffer et al. 1998; Suresh 2003], shape simplification [Tam and Heidrich 2003], motion planning [Foskey et al. 2001], etc.

Given a 3D solid, there are two main issues in the computation of MAT: algebraic complexity and instability. The

algebraic complexity arises from the fact that the sheets, seams and junctions of a medial axis are high degree algebraic primitives and it is hard to compute the exact MAT reliably for complex models. Many techniques have been proposed in the literature to approximate the medial axis. At a broad level these techniques either compute the Voronoi diagram of a point sample on the boundary of the solid or evaluate the distance field of the primitives on a spatial grid followed by isosurface extraction.

The instability refers to the property that small modifications to the boundary of the solid can induce large modifications in its medial axis. Different algorithms have been proposed to compute a stable subset of the medial axis based on different geometric criteria. However, current algorithms are either limited to point datasets or may not preserve the topology of the medial axis.

Main Results: We present a novel algorithm to compute a homotopy preserving simplification of the medial axis of polyhedral models. Our algorithm computes a polygonal approximation that has the same homotopy type as Blum medial axis and thereby preserves its topological structure. Our simplification algorithm is based on the *separation angle* [Amenta et al. 2001; Dey and Zhao 2002; Dimitrov et al. 2003] and simplified medial axis [Foskey et al. 2003] and attempts to remove features on the medial axis for which the separation angle is below a certain threshold.

We initially compute a connectivity graph that represents the connectivity of the original medial axis. We present an iterative algorithm that removes sheets from the medial axis based on the separation angle criterion without changing the homotopy type of the medial structure. We maintain a priority queue of the sheets ordered by the separation angle and compute a simplification of the medial axis by removing sheets from the queue that correspond to the unstable parts. We give rigorous guarantees on the homotopy of the simplified medial axis. Our simplification algorithm can also be applied to other medial axis formulations as long as we are given an approximation with correct homotopy.

Some of the novel results of our work include:

- A homotopy preserving medial axis simplification algorithm applicable to continuous solid representations.
- Relationship between the stability of medial axis junctions and seams to stability of incident sheets using separation angles.
- Algorithm to compute the stability of a medial axis sheet using discrete sampling.

We have implemented the algorithm on a PC with 2.4GHz Pentium IV processor and applied it to complex CAD models consisting of tens of thousands of triangles. Our algorithm is able to simplify the medial axis of these models while preserving the homotopy in a few seconds.

As compared to prior algorithms for medial axis simplification, our algorithm offers the following *advantages*:

*e-mail: sud@cs.unc.edu

†e-mail: mark_foskey@unc.edu

‡e-mail: dm@cs.unc.edu

- **Polyhedral Models:** Our simplification algorithm is applicable to models with continuous piecewise linear boundaries, possibly with internal voids.
- **Complex Models:** Our algorithm is able to handle complex models with a high number of boundary primitives and complex topologies.
- **Homotopy Preservation:** The simplified medial axis has the same topological structure as the original medial axis. These properties are important for shape analysis, motion planning and mesh generation.
- **Efficiency:** Our simplification is based on iterative pruning of the connectivity graph and is computed using local operations.

Organization: The rest of the paper is organized in the following manner. We give a brief overview of previous work in Section 2. We introduce our notation and present background material in Section 3. We present the formulation of a topology preserving simplified medial axis in Section 4 and describe our algorithm in Section 5. We prove the correctness of algorithm in Section 6. We describe its implementation in Section 7 and highlight its performance on different benchmarks. We analyze its performance and compare it with other approaches in Section 8.

2 Related Work

In this section we give a brief overview of prior work on medial axis computation as well as medial axis simplification. We make this separation for convenience, but it is important to realize that the two are often integrated in practice.

2.1 Medial Axis Computation

In this section we focus on methods to generate a (possibly approximate) medial axis of a figure or object. The algorithms are categorized based on different model representations.

Image datasets. The problem of MAT computation of a point dataset has been extensively studied in computer vision and image processing. In two and three dimensions, approximations to the medial axis have been computed using *thinning algorithms* [Lam et al. 1992; Zhang and Wang 1993]. Many algorithms based on partial differential equations of front propagation have also been proposed [Kimmel et al. 1995; Siddiqi et al. 1997]. Pizer et al. [2003] have generated structures related to the medial axis using filters which yield high values for points near the medial axis of an object.

Boundary point samplings. Algorithms for computing the medial axis of an object from a sample of points on the boundary typically begin by constructing a Voronoi diagram of the point set, after which they use various criteria to prune faces [Amenta et al. 2001; Dey and Zhao 2002; Naf et al. 1996; Ogniewicz and Kübler 1995; Sheehy et al. 1995; Shaham et al. 2004].

Polyhedral Models. Many MAT computation algorithms have been proposed for polyhedral models based on 3D tracing of the seam curves [Milenkovic 1993; Reddy and Turkiyyah 1995; Sherbrooke et al. 1996]. These algorithms

solve a system of algebraic equations to compute the junction points and the seam curves. Culver et al. [1999] used exact arithmetic to compute the MAT accurately. Etzion and Rappoport [2002] used spatial subdivision techniques to determine connectivity of the MAT and only can guarantee correctness up to a certain resolution. All these algorithms have been applied to polyhedral models with a few hundred triangles. Foskey et al. [2003] used graphics hardware to generate an image-space representation of the gradient of the distance field to the boundary, which can be analyzed to find the medial axis. The gradient field in their method is actually the same as the velocity field of the propagating front in the methods of Siddiqi et al. [2002] mentioned above. Du and Qin [2004] also computed an approximation of the medial axis using diffusion partial differential equations solved at a discrete sample of boundary points. Yang et al. [2004] generated sample points on the boundaries of maximal spheres, and apply a separation angle criterion (see below) to select the points approximately on the medial axis.

2.2 Medial Axis Simplification

In this section we give a brief overview of medial axis simplification algorithms. The instability of the medial axis, and its resulting complexity for objects with boundaries exhibiting fine detail, has been known for some time (see for instance, Blum and Nagel [1978]). A number of methods for simplifying the medial axis have been proposed. Pizer et al. [2003] have presented an extensive survey of methods for approximating and simplifying the medial axis.

A well known criterion for medial axis simplification is based on the *object angle* [Dimitrov et al. 2003]. The separation angle is twice the object angle at any point on the medial axis. The underlying methods involve computing subsets for which the object angle is above a certain threshold. Malandain and Fernández-Vidal [1998] traced the idea, in varying forms, back to Meyer [1979] and Kruse [1991]. Our simplification algorithm also uses this criterion.

Siddiqi et al. [2002] formulated the detection of gradient discontinuities in terms of the average gradient flux into a neighborhood, which has been shown to be closely related to the object angle [Dimitrov et al. 2003]. Malandain and Fernández-Vidal [1998] used a criterion combining the object angle and the distance between the two points nearest to the medial axis point. Foskey et al. [2003] detected gradient discontinuities across adjacent voxels by comparing the directions of neighboring vectors.

Another class of approaches are based on using a point sampling of the boundary. These algorithms approximate the medial axis by computing the Voronoi diagram of the set of points and eliminating some of the Voronoi faces using different criteria. Amenta et al. [2001] used the distance between the two points nearest to the medial axis point as a criterion for medial axis simplification. Dey and Zhao [2002] combined a similar distance criterion with an object angle criterion and observed that the two criteria together tend to eliminate spurious holes. Tam and Heidrich [Tam and Heidrich 2003] used a volume criterion to remove parts of the medial axis while preserving the topology. Leymarie and Kimia [2001] also began with surface point samples, but their algorithms are based on the differential equations of front propagation.

2.3 Topological and Smoothness Properties of MAT

Attali, Boissonat, and Edelsbrunner [Attali et al. 2004] have surveyed different techniques that generate a stable and homotopy preserving medial structure. The homotopy relationship between an object and its medial axis has been proven in a particularly general form by Lieutier [2003], who showed that homotopy preservation holds for any bounded open subset of \mathbb{R}^n . Chazal and Soufflet [2004] presented smoothness constraints on the boundary of a solid, which need not be polyhedral, under which the medial axis obeys certain stability and finiteness conditions. Chazal and Lieutier [2004] have also proven results about stability, and presented a homotopy preserving medial axis simplification, however the approach has not been demonstrated on complex models. We discuss some of these simplification methods in relation to our work in section 8.

3 Notation and Background

In this section, we introduce some of the terminology used in this paper. We also give a brief overview of the θ -simplified medial axis (θ -SMA).

3.1 Basic Terminology

The notations are summarized in table 1, and are formally defined below. In this paper we will consider only polyhedral solids, which we refer to as an object \mathcal{O} . The solid \mathcal{O} can have internal voids. For a point $\mathbf{x} \in \mathcal{O}$, any point on the boundary of \mathcal{O} that is at least as close to \mathbf{x} as any other will be called a *nearest neighbor* of \mathbf{x} , and the set of nearest neighbors will be called the *neighbor set* of \mathbf{x} and denoted $NS(\mathbf{x})$. With a distance function $d()$,

$$NS(\mathbf{x}) = \{\mathbf{y} \in \partial\mathcal{O} \mid d(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \partial\mathcal{O})\}.$$

Then the *medial axis* of \mathcal{O} , denoted \mathcal{M} , is defined as the set of points inside \mathcal{O} with at least two nearest neighbors.

The boundary of \mathcal{O} can be decomposed disjointly into vertices, open edges, and open faces, which we refer to collectively as *sites*. Each nearest neighbor of \mathbf{x} will be in exactly

Notation	Meaning
\mathcal{X}^c	Closure of a set \mathcal{X}
$\overline{\mathcal{X}}$	Compliment of \mathcal{X}
\mathcal{X}°	Interior of \mathcal{X}
$\partial\mathcal{X}$	Boundary of \mathcal{X}
$ \mathcal{X} $	Cardinality of \mathcal{X}
\mathcal{O}	Polyhedral solid in \mathbb{R}^3
p_i	A face, edge or vertex site on $\partial\mathcal{O}$
$\mathbf{n}_i(\mathbf{x})$	Normal to a site p_i from a point \mathbf{x}
$NS(\mathbf{x})$	Set of boundary points closest to $\mathbf{x} \in \mathcal{O}$
$Gov(\mathbf{x})$	Set of boundary sites closest to $\mathbf{x} \in \mathcal{O}$
\mathcal{M}	Medial axis of \mathcal{O}
\mathcal{F}, f_i	Set of sheets of \mathcal{M} , one sheet of \mathcal{M}
\mathcal{E}, e_i	Set of seams of \mathcal{M} , one seam of \mathcal{M}
\mathcal{V}, v_i	Set of junctions of \mathcal{M} , one junction of \mathcal{M}
$\mathcal{R}(f_i)$	Set of <i>rim curves</i> of a sheet f_i
$\mathcal{S}(f_i)$	Set of <i>seam curves</i> of a sheet f_i

Table 1: Notation used in the paper

one site, and we define the set of neighboring sites $Gov(\mathbf{x})$ to be the set of sites containing a nearest neighbor of \mathbf{x} :

$$Gov(\mathbf{x}) = \{p_i \mid \mathbf{y} \in p_i \text{ for some } \mathbf{y} \in NS(\mathbf{x})\}.$$

For each point $\mathbf{x} \in \mathcal{M}$, the sites in the set $Gov(\mathbf{x})$ are called the *governors* of the point \mathbf{x} . Clearly, $|Gov(\mathbf{x})| \geq 2$ for any point \mathbf{x} on the medial axis.

3.2 Medial Axis Point Classification

We define a *sheet set* to be the set of all medial axis points governed by a specified pair of sites (or at least having that pair among their governors), and we define a *sheet* to be a connected component of a sheet set. The interior of a sheet is a smooth surface. A sheet may contain holes, because we do not require that the boundary of \mathcal{O} be connected or have only simply connected faces. A *seam curve*, or *seam*, is a connected component of the intersection of two or more sheets. The intersection of three or more seams is a *junction*. This definition corresponds approximately to those given in [Culver et al. 1999] and [Sherbrooke et al. 1996]. Finally, for any subset \mathcal{M}' of \mathcal{M} , the intersection of a seam with the boundary will be a *seam end*. The intersection of a sheet with seam ends removed, and the boundary of \mathcal{M}' will be a *rim set*. An example of seam points, junction points, and rim points is given in figure 1. A similar classification of medial axis points for any bounded set in \mathbb{R}^3 is given in [Giblin and Kimia 2000].

We make one special proviso about rims and seam ends. In general, including the case when $\mathcal{M}' = \mathcal{M}$, the boundary of \mathcal{M}' will not be contained in \mathcal{M}' . In this case, it is possible that two sheets that do not intersect will have boundaries that do intersect. If this occurs, their rim curves and seam ends will be treated as distinct combinatorial entities, since the goal is to reflect the connectivity properties of \mathcal{M} , not its closure.




Figure 1: Medial axis point classification: (a) Classification of the points on the medial axis (thin lines) of a simple polyhedron (thick lines) (b) A subset $\mathcal{M}' \subset \mathcal{M}$ is shaded in gray. A rim point and a seam point on the boundary of the central sheet are shown.

3.3 Homotopy Equivalence

One of the major goals of our work is to compute a simplification of the MAT that is homotopy equivalent to the exact MAT. The notion of homotopy equivalence between topological sets enforces a one-to-one correspondence between connected components, holes, tunnels or cavities and also the

way in which they are related. It has been shown by Lieutier [Lieutier 2003] that any bounded open subset $\mathcal{X} \subseteq \mathbb{R}^n$ is homotopy equivalent to its medial axis. Intuitively this implies that the medial axis and the shape are connected in the same way.

Formally, two maps $f: \mathcal{X} \rightarrow \mathcal{Y}$ and $g: \mathcal{X} \rightarrow \mathcal{Y}$ are *homotopic* if there exists a continuous family of maps $h_t: \mathcal{X} \rightarrow \mathcal{Y}$, for $t \in [0, 1]$, such that $h_0 = f$ and $h_1 = g$. Thus, a homotopy is a deformation of one map to another. Two spaces \mathcal{X} and \mathcal{Y} are *homotopy equivalent* if there exist continuous maps $f: \mathcal{X} \rightarrow \mathcal{Y}$ and $g: \mathcal{Y} \rightarrow \mathcal{X}$ such that $g \circ f$ and $f \circ g$ are homotopic to the identity maps on their respective spaces. As an example, f could be the inclusion of a circle into an annulus, and g could be radial projection of the annulus onto the circle.

In situations such as this one, where f is an inclusion and $f \circ g$ is actually equal to the identity map, the homotopy equivalence is called a *deformation retraction*. See Spanier [1989] for details of these definitions. Our simplification algorithm also performs a sequence of deformation retractions on the original medial axis to generate a simplified medial axis with the same homotopy type as the original.

3.4 θ -Simplified Medial Axis

Given a polyhedral model \mathcal{O} and a medial axis \mathcal{M} , the *separation angle* $\Theta(\mathbf{x})$ at each point \mathbf{x} on \mathcal{M} is the largest angle subtended by a pair of nearest neighbor points on $\partial\mathcal{O}$, and is given by

$$\Theta(\mathbf{x}) = \max_{\mathbf{y}_i, \mathbf{y}_j \in \mathcal{NS}(\mathbf{x})} (\angle \mathbf{y}_i \mathbf{x} \mathbf{y}_j)$$

Given an angle θ , the *θ -simplified medial axis* (θ -SMA) of \mathcal{O} , denoted by \mathcal{M}_θ , is the set of points of \mathcal{M} with separation angle greater than θ [Foskey et al. 2003] (see figure 2). Foskey et al. [2003] discuss the convergence and stability




Figure 2: θ -Simplified Medial Axis, \mathcal{M}_θ : (a) The medial axis (black) of a part of a polyhedron (blue) (b) \mathcal{M}_θ for $\theta = \pi/2$.

of \mathcal{M}_θ and provide error bounds on the boundary reconstructed from \mathcal{M}_θ . The speed of medial axis formation at point \mathbf{x} is proportional to $\frac{1}{\sin \Theta(\mathbf{x})}$ [Pizer et al. 2003]. Parts of the medial axis with a higher speed of formation are regarded as more important [Blum 1967], and the separation angle $\Theta(\mathbf{x})$ has been used as a measure of the stability of the medial axis at the point \mathbf{x} .

4 θ -Homotopy Medial Axis

In this section, we analyze the topological characterization of θ -SMA and present a formulation for computing a

homotopy-preserving simplified medial axis, the *θ -homotopy medial axis*. The problem with the θ -SMA is that it does not in general preserve the homotopy type of the medial axis. The θ -SMA can be disconnected when the medial axis is connected, or have holes when the medial axis does not, and lack holes when the medial axis has them. An illustration of the failure of connectivity is shown in Figure 3. The other kinds of connectivity problems also arise because




Figure 3: The θ -Simplified Medial Axis, $\mathcal{M}_{\pi/3}$ is disconnected even though the original object \mathcal{O} is connected. Note that the separation angle at \mathbf{x} is less than $\pi/3$, while it exceeds $\pi/3$ for the portions of the medial axis shown.

the angle criterion may discard topologically significant portions. The fundamental issue here is that homotopy type is a global property, whereas the separation angle is a local measure.

Decreasing the θ threshold does not provide a guaranteed solution to fix the problems. As illustrated in Figure 3, the problem is associated with local minima of the separation angles, and such a local minimum can occur for any value of θ . In any event, decreasing θ only increases the number of unstable features of the θ -SMA.

Our goal is to compute a simplified medial axis that would allow significant simplification corresponding to large values of θ , while preserving the homotopy type of \mathcal{M} . Clearly such a simplified medial axis has to be a superset of \mathcal{M}_θ . However, we would like such an axis to be minimal in some regard in order to minimize the unstable parts. We now formally present the desired subset of the medial axis. Let \mathcal{H}_θ denote the class of subsets of \mathcal{M} which are supersets of \mathcal{M}_θ and are homotopy equivalent to \mathcal{M} .

$$\mathcal{H}_\theta = \{\mathcal{X} | \mathcal{X} \subseteq \mathcal{M}, \mathcal{X} \supseteq \mathcal{M}_\theta, \mathcal{X} \simeq \mathcal{M}\}$$

Define a set $\mathcal{X} \in \mathcal{H}_\theta$ to be *irreducible* if the removal of any sheet yields a set that either has a different homotopy type, or is no longer a superset of \mathcal{M}_θ . That is,

$$\mathcal{M}_\theta^* = \{\mathcal{X} | \mathcal{X} \in \mathcal{H}_\theta, \text{ for all } f_i \in \mathcal{X}, (\mathcal{X} \setminus \{f_i\}) \notin \mathcal{H}_\theta\}$$

We will refer to any irreducible set in \mathcal{H}_θ as a *θ -homotopy medial axis*, or θ -HMA. We will typically denote a θ -HMA by \mathcal{M}_θ^* . The set \mathcal{M}_θ^* is not unique. A discussion about lack of uniqueness is presented in section 8.

5 θ -Homotopy Medial Axis Computation

In this section we present an algorithm for computing a θ -HMA \mathcal{M}_θ^* . We begin by computing a Voronoi diagram of the polyhedron using spatial subdivision techniques presented

in [Sud and Manocha 2005]. Our approach is similar to the subdivision algorithm of [Etzion and Rappoport 2002]. A Voronoi graph is computed that represents the connectivity of the Voronoi diagram of the polyhedron. A node in the graph corresponds to a Voronoi face, edge or vertex and an edge indicates an incidence relationship between two nodes. The algorithm computes an exact Voronoi graph of the polyhedron if the Voronoi diagram is not degenerate. In presence of degeneracies, it computes an ϵ -Voronoi graph, similar to [Etzion and Rappoport 2002]. However, in contrast to this work, our algorithm subdivides a cell till the arrangement of Voronoi faces can be unambiguously determined from a labeling of the governors at the corners of the cell. We use a simple subdivision criteria, and reduce computation of intersections of conic sections on faces of a cell. Furthermore, the algorithm can handle internal voids in the polyhedron. Given the Voronoi graph, a sub-graph corresponding to the medial axis \mathcal{M} of the polyhedron is computed using the property of Lemma 12 in [Etzion and Rappoport 2002]. The Voronoi faces, edges and vertices correspond to the medial axis sheets, seams and junctions respectively. We refer to this sub-graph, which captures the connectivity of different elements of the medial axis \mathcal{M} , as the *connectivity graph* of \mathcal{M} .

The diameter of a cell after the spatial subdivision gives a polygonal approximation to the geometric part of the Voronoi diagram. The approximation has bounded Hausdorff error to the exact Voronoi diagram, like the *Proximity Structure Diagram* [Etzion and Rappoport 2002]. This geometric approximation is used to construct a polygonal mesh approximation of the θ -HMA consisting of axis aligned faces.

Given the exact medial axis \mathcal{M} , our simplification algorithm is presented in Section 5.2 and it simplifies the medial axis by pruning sheets of the medial axis. We first define the separation angle of a sheet f_i to be the supremum of the separation angles for all points interior to the sheet:

$$\Theta(f_i) = \max_{\mathbf{x} \in f_i^\circ} (\Theta(\mathbf{x})).$$

$\Theta(f_i)$ gives a measure of the stability of the sheet f_i . We use a conservative definition for the separation angle of the sheet to ensure that the simplified medial \mathcal{M}_θ^* is a superset of \mathcal{M}_θ . Similarly we define the separation angle of a seam e_i as:

$$\Theta(e_i) = \max_{\mathbf{x} \in e_i^\circ} (\Theta(\mathbf{x})).$$

5.1 Sheet Separation Angle Computation

The Voronoi diagram computation algorithm computes a piecewise linear approximation of each sheet based on a discrete sampling introduced by spatial subdivision. In this section, we address the problem of computing a bounded approximation of the sheet separation angle $\Theta(f_i)$. Each sheet of the medial axis of a polyhedron is trimmed quadric surface [Culver 2000]. Exact computation of the sheet separation angle involves computing the extreme value of a non-linear function on a quadric surface. Instead we present an efficient approach to compute a conservative upper bound on the separation angle using spatial subdivision. The tightness of the bound depends on the degree of subdivision.

Given a cell C and a sheet f_i intersecting the cell, our goal is to compute the maximum separation angle for all points on the sheet inside the cell. Let $\{p_1, p_2\}$ be the two governors

of the sheet f_i and \mathbf{c} be the center of the cell C . We classify the inputs into 2 cases:

1. The governors do not intersect the cell C , i.e. $C \cap \{p_1, p_2\} = \emptyset$
2. At least one of the governors intersects the cell C , i.e. $C \cap \{p_1, p_2\} \neq \emptyset$.

Case 1: $C \cap \{p_1, p_2\} = \emptyset$. We simplify the problem to computing the maximum of the separation angles for all points inside the cell to the two governors. We compute the separation angle from the center \mathbf{c} of the cell to each of the two governors and add conservative error bounds to get the maximum separation angle. Let \mathbf{x} be any point inside cell C . Let $\mathbf{n}_i(\mathbf{x})$ denote the normal vector from a point \mathbf{x} to the sites p_i , ($i = 1, 2$), and $\alpha_i(\mathbf{x})$ represent the angle between $\mathbf{n}_i(\mathbf{c})$ and $\mathbf{n}_i(\mathbf{x})$. If $\Delta\theta_i$ is an upper bound on $\alpha_i(\mathbf{x})$ for all $\mathbf{x} \in C$, then the maximum separation angle for sheet f_i inside cell C is given by:

$$\Theta(f_i \cap C) \leq \cos^{-1} \left(\frac{\mathbf{n}_1(\mathbf{c}) \cdot \mathbf{n}_2(\mathbf{c})}{\|\mathbf{n}_1(\mathbf{c})\| \|\mathbf{n}_2(\mathbf{c})\|} \right) + \Delta\theta_1 + \Delta\theta_2$$

The computation of the error bounds $\Delta\theta_i$ for each of the three types of governors (point site, line site and triangle site) is presented below:




Figure 4: Normal Cone to compute $\Delta\theta$ for a point site

Point Site p_i : \mathbf{p} . The range of angles subtended by a point site to all points in the cell is given by a normal cone. The normal cone is the smallest cone enclosing the cell C with the apex at \mathbf{p} and axis along $\mathbf{n}_i(\mathbf{c})$ (see figure 4). Let $\Delta\theta_i$ be the half opening angle of the cone. The angle $\alpha_i(\mathbf{x})$ is maximized when point \mathbf{x} is one of the corner vertices \mathbf{v}_j ($1 \leq j \leq 8$) of the cell C . Thus, for the smallest cone enclosing the cell C ,

$$\Delta\theta_i = \max_{1 \leq j \leq 8} \left[\cos^{-1} \left(\frac{\mathbf{n}_i(\mathbf{v}_j) \cdot \mathbf{n}_i(\mathbf{c})}{\|\mathbf{n}_i(\mathbf{v}_j)\| \|\mathbf{n}_i(\mathbf{c})\|} \right) \right], \text{ where } \mathbf{n}_i(\mathbf{x}) = \mathbf{p} - \mathbf{x}.$$

Line Site p_i : $\mathbf{p} + \lambda(\mathbf{q} - \mathbf{p})$. The range of angles subtended




Figure 5: Wedge to determine $\Delta\theta$ for a line site

by a line to all the points in the cell is given by the smallest wedge enclosing the cell, with the top edge of the wedge being the line site (see figure 5). Let $\Delta\theta_i$ be the half angle of the wedge. As in the point site case, the angle $\alpha_i(\mathbf{x})$ is maximized when point \mathbf{x} is one of the corner vertices \mathbf{v}_j ($1 \leq j \leq 8$) of the cell C . Thus, for the smallest wedge enclosing the cell C ,

$$\Delta\theta_i = \max_{1 \leq j \leq 8} \left[\cos^{-1} \left(\frac{\mathbf{n}_i(\mathbf{v}_j) \cdot \mathbf{n}_i(\mathbf{c})}{|\mathbf{n}_i(\mathbf{v}_j)| |\mathbf{n}_i(\mathbf{c})|} \right) \right],$$

where $\mathbf{n}_i(\mathbf{x}) = \mathbf{p} + \lambda(\mathbf{q} - \mathbf{p}) - \mathbf{x}$, $\lambda = \frac{(\mathbf{x} - \mathbf{p}) \cdot (\mathbf{q} - \mathbf{p})}{(\mathbf{q} - \mathbf{p})^2}$.

Triangle Site p_i with face normal $\hat{\mathbf{n}}$. The shortest path from any point to the triangle is perpendicular to the face. Thus $\mathbf{n}_i(\mathbf{x}) = \hat{\mathbf{n}}$ for all \mathbf{x} , and $\Delta\theta_i = 0$.

Case 2 $C \cap \{p_1, p_2\} \neq \emptyset$. If the two sites do not intersect ($p_1 \cap p_2 = \emptyset$), then the bisector surface (and sheet f_i) also do not intersect either site. In such a case we can subdivide the cell C into sub-cells $\{C_k\}$ such that $C_k \cap \{p_1, p_2\} = \emptyset$ if $C \cap f_i \neq \emptyset$. The computation of the sheet separation angle is then reduced to Case 1.

If the two sites intersect ($p_1 \cap p_2 \neq \emptyset$), then the sheet corresponds to one of the non-generic cases of a bisector surface [Culver 2000], and the separation angle $\Theta(f_i)$ can be determined exactly from the pairs of governors. The case of two point governors case never occurs, we examine each of the other 5 pairs of governors individually.

Point-Triangle The bisector surface is a redundant line, and never occurs on the medial axis [Culver 2000].

Point-Line The bisector surface is a plane through the point and perpendicular to the line, $\Theta(f_i) = 0$

Line-Line The bisector surface is an orthogonal plane pair, $\Theta(f_i) = \text{angle between the two lines}$.

Line-Triangle The bisector surface is a right circular cone, or a plane if the line is incident on the triangle. In first case, the separation angle $\Theta(f_i) = \pi/2 - \cos^{-1}(\hat{\mathbf{l}} \cdot \hat{\mathbf{n}})$, where $\hat{\mathbf{l}}$ and $\hat{\mathbf{n}}$ are unit normals along the line and to the triangle respectively. In the second case, $\Theta(f_i) = 0$.

Triangle-Triangle The bisector surface is an orthogonal plane pair, and the separation angle is given by $\Theta(f_i) = \cos^{-1}(\hat{\mathbf{n}}_1 \cdot \hat{\mathbf{n}}_2)$, where $\hat{\mathbf{n}}_1, \hat{\mathbf{n}}_2$ are unit normals to the two triangles.

5.2 Simplification Algorithm

We now present our medial axis simplification algorithm. We treat the medial axis \mathcal{M} as an abstract 2-dimensional complex consisting of faces, edges, and vertices. Initially, edges correspond either to the seam curves, which lie between sheets, or rim curves, which lie on the boundary.

The key idea in our algorithm is a simple criterion for determining whether a sheet can be removed without changing the homotopy type of the medial structure. We call such sheets *frontier sheets* (Figure 6). We will describe this criterion below, but first give an overview of how it is used in the algorithm. We maintain a set \mathcal{Q} of all frontier sheets. We successively remove sheets from this set until it is empty. As each sheet is removed from \mathcal{Q} , it is also removed from the




Figure 6: Classification of sheets for iterative pruning: The sheets colored gray are frontier sheets, and can be removed without changing the homotopy type. For the ‘loop’ sheets the rim set is not connected and they will never become frontier sheets. The ‘interior’ sheet has an empty rim set, however it may become a frontier sheet after removal of one of its adjacent sheets.

medial structure if its separation angle is no greater than θ . Removal of a sheet from the structure can affect whether its neighbors are frontier sheets, and so each time we remove a sheet we check each neighbor of that sheet to see if it needs either to be added or removed.

A sheet f_i is defined to be a frontier sheet provided that its set of rim points $\mathcal{R}(f_i)$ and its set of seam points $\mathcal{S}(f_i)$ are both connected and nonempty. The set \mathcal{Q} is defined as:

$$\mathcal{Q} = \{f_i \mid \mathcal{R}(f_i) \neq \emptyset, \mathcal{R}(f_i) \text{ is connected}, \mathcal{S}(f_i) \neq \emptyset, \mathcal{S}(f_i) \text{ is connected}\}. \quad (1)$$

In Section 6.2 we will prove that the frontier sheets are precisely those sheets which may be removed without changing the homotopy type. We present an intuitive justification for that claim here. If the rim set and seam set are both connected then each set is a single curve, and removing the sheet is equivalent to retracting the sheet onto its seam set via a homotopy (see Figure 7(a)). On the other hand, if




Figure 7: Sheet pruning: (a) The cyan sheet is a valid frontier sheet, and has a deformation retract to its seam set. (b) The cyan sheet is not a frontier sheet. (c) Removing the sheet makes the two adjacent sheets disconnected.

the rim set is disconnected or empty, removing the sheet removes a path between two points on different seam components and hence does not preserve the homotopy type (see Figures 7(b),(c)). Note that, when we remove a sheet, we remove its interior and its rim set, but not its seam set.

We noted earlier that removing a sheet can cause other sheets either to lose or gain frontier status, and we can now explain why this is true. A sheet with an empty rim set can gain a rim edge if one of its neighboring sheets is removed, and thereby become a frontier sheet. Conversely, a sheet with a single seam component can find that its seam set is broken into two components if an adjacent sheet is removed.

Algorithm 1 simplifies \mathcal{M} based on removal of frontier sheets. Let the resulting medial subset after the j th iteration be \mathcal{M}_j , and let the corresponding frontier set be \mathcal{Q}_j . The frontier set is maintained as a priority queue, the priority deter-

Input: Initial (Blum) medial axis \mathcal{M}_0 , angle θ
Output: Final medial subset \mathcal{M}_f

```

1 Label all sheets in  $\mathcal{M}$  as unmarked
2 Initialize  $\mathcal{Q}_0, j \leftarrow 0$ 
3 repeat
4    $f_i \leftarrow \text{ExtractSheet}(\mathcal{Q}_j)$ 
5    $(\mathcal{Q}_{j+1}, \mathcal{M}_{j+1}) \leftarrow \text{RemoveSheet}(f_i, \mathcal{Q}_j, \mathcal{M}_j, \theta)$ 
6    $j \leftarrow j + 1$ 
7 until  $(\mathcal{Q}_j = \emptyset)$ 
8  $\mathcal{M}_f \leftarrow \mathcal{M}_{j+1}$ 

```

Algorithm 1: SimplifyMAT(\mathcal{M}_0, θ).

mined by the sheet separation angle. Initially, $\mathcal{M}_0 = \mathcal{M}$. \mathcal{Q}_0 is computed using \mathcal{M}_0 in equation (1). The function $\text{ExtractSheet}(\mathcal{Q}_j)$ in line 4 returns a sheet with minimum separation angle from the set \mathcal{Q}_j (but does not remove it from \mathcal{Q}_j). The key step in the algorithm is the removal of a frontier sheet in line 5, which is described in Algorithm 2.

Input: A frontier sheet f_i , frontier set \mathcal{Q}_j , medial subset \mathcal{M}_j , angle θ
Output: Frontier set \mathcal{Q}_{j+1} , medial subset \mathcal{M}_{j+1}

```

1 if  $\Theta(f_i) \geq \theta$  then
2   Label  $f_i$  as fixed
3    $\mathcal{Q}_{j+1} \leftarrow \mathcal{Q}_j \setminus \{f_i\}$ 
4 else
5    $\mathcal{M}_{j+1} \leftarrow \mathcal{M}_j \setminus \{f_i\}$ 
6    $\mathcal{Q}_j \leftarrow \mathcal{Q}_j \setminus \{f_i\}$ 
7    $\mathcal{Q}_{j+1} \leftarrow \text{UpdateFrontierNbrs}(f_i, \mathcal{Q}_j)$ 
8 end

```

Algorithm 2: RemoveSheet($f_i, \mathcal{Q}_j, \mathcal{M}_j, \theta$).

Algorithm 2 removes a frontier sheet from the medial subset \mathcal{M}_j only if the separation angle of the sheet lies below the angle threshold θ . (line 1). The removal of a frontier sheet does not change the homotopy type of \mathcal{M}_j . As we noted earlier, removal of the sheet from \mathcal{M}_j may change the frontier status of its neighboring sheets. Neighboring sheets are checked for such changes and the frontier set is updated in (line 7), which is described in detail as Algorithm 3.

Input: A frontier sheet f_i , frontier set \mathcal{Q}_j
Output: Frontier set \mathcal{Q}_{j+1}

```

1 Initialize  $\mathcal{Q}_{j+1} \leftarrow \mathcal{Q}_j$ 
2 foreach sheet  $f_k$  sharing a seam point with  $f_i$  do
3   if  $(\text{Label}(f_k) \neq \text{fixed})$  then
4     if  $(f_k \text{ is a frontier sheet})$  then
5        $\mathcal{Q}_{j+1} \leftarrow \mathcal{Q}_{j+1} \cup \{f_k\}$ 
6     else
7        $\mathcal{Q}_{j+1} \leftarrow \mathcal{Q}_{j+1} \setminus \{f_k\}$ 
8 end

```

Algorithm 3: UpdateFrontierNbrs(f_i, \mathcal{Q}_j)

6 Correctness

In this section we demonstrate that Algorithm 1 is correct, i.e. the final medial subset is a valid θ -HMA. Let \mathcal{M}_f be

the subset of \mathcal{M} obtained as the final out of algorithm 1. To prove correctness, we must show that \mathcal{M}_f contains \mathcal{M}_θ , \mathcal{M}_f has the homotopy type of \mathcal{M} , and \mathcal{M}_f is irreducible. We will first show that $\mathcal{M}_\theta \subset \mathcal{M}_f$.

6.1 Separation Angles of Medial Axis Parts

It is clear from the definition of the separation angle for a sheet that every sheet interior point that is removed will have a separation angle no greater than the threshold θ . So it remains to show that no seam or junction point is removed if its separation angle is greater than θ .

The set of governors for all points in the interior of the sheet, and on the boundary curves, remains the same and each governor is linear. Thus the separation angle $\Theta(x)$ is a continuous function of all points in the interior of a sheet, and on the rim points on the boundary of the sheet. However, the set of governors changes at a seam or a junction, causing the separation angle to be discontinuous on the boundary of the sheet (figure 8). Lemma 1 bounds the discontinuity in the separation angle at the seam and junction boundaries of a sheet.




Figure 8: Separation angle of a seam e_i : Three sheets f_1, f_2 and f_3 meet at a seam e_i . For any point \mathbf{y} on e_i , $\Theta(\mathbf{y}) \leq \Theta(f_1)$.

Lemma 1.

- (i) Let e_i be a (non-degenerate) seam of a medial axis, formed by intersection of three sheets f_1, f_2 and f_3 . Then, $\Theta(e_i) \leq \max_{1 \leq j \leq 3}(\Theta(f_j))$
- (ii) Let v_i be a (non-degenerate) junction of a medial axis, formed by intersection of four sheets f_1, f_2, f_3 and f_4 . Then, $\Theta(v_i) \leq \max_{1 \leq j \leq 4}(\Theta(f_j))$

Proof. (i) Let the set of governors of sheet f_1 be $\text{Gov}(f_1) = \{a, b\}$. Since f_1 and f_2 intersect, $\text{Gov}(f_1) \cap \text{Gov}(f_2) \neq \emptyset$. Also $\text{Gov}(f_1) \neq \text{Gov}(f_2)$ as two intersecting sheets cannot have same set of governors. Thus $|\text{Gov}(f_1) \cap \text{Gov}(f_2)| = 1$, and $\text{Gov}(f_2) = \{b, c\}$. Similarly $\text{Gov}(f_3) = \{a, c\}$, and the set of governors of e_i is $\text{Gov}(e_i) = \{a, b, c\}$. For any point $\mathbf{x} \in e_i$, the closest points on a, b and c be $\mathbf{y}_a, \mathbf{y}_b, \mathbf{y}_c$. Then $NS(\mathbf{x}) = \{\mathbf{y}_a, \mathbf{y}_b, \mathbf{y}_c\}$, and by definition of $\Theta(\mathbf{x})$,

$$\begin{aligned} \Theta(\mathbf{x}) &= \max(\angle \mathbf{y}_a \mathbf{x} \mathbf{y}_b, \angle \mathbf{y}_b \mathbf{x} \mathbf{y}_c, \angle \mathbf{y}_a \mathbf{x} \mathbf{y}_c) \\ &= \angle \mathbf{y}_a \mathbf{x} \mathbf{y}_b \text{ (assume WLOG)} \end{aligned}$$

Let \mathbf{y} be a point on sheet f_1 inside a δ -neighborhood of \mathbf{x} . Since $\Theta(f_i)$ is continuous and $\text{Gov}(f_1) = \{a, b\}$, $\lim_{\mathbf{y} \rightarrow \mathbf{x}} \Theta(\mathbf{y}) = \angle \mathbf{y}_a \mathbf{x} \mathbf{y}_b = \Theta(\mathbf{x})$. By definition of




Figure 9: Flange Plate Model (990 polygons): Medial axis sheets through a cut-out of the model, $\theta = 150^\circ$ (b) The sheets become disconnected, holes disappear for the θ -SMA. (c) In the θ -HMA the holes are preserved, and the entire medial axis remains connected.

$\Theta(f_1), \Theta(e_1) \geq \lim_{\mathbf{y} \rightarrow \mathbf{x}} \Theta(\mathbf{y})$. Hence, $\Theta(x) \leq \Theta(f_1)$. Since choice of point \mathbf{x} on e_i was arbitrary,

$$\Theta(e_i) = \Theta(\mathbf{x}) \leq \Theta(f_1) \leq \max_{1 \leq j \leq 3} (\Theta(f_j))$$

- (ii) Proof follows as above, using 4 governors of the junction, instead of 3 governors of the seam. \square

The implication of Lemma 1 is that we can get an upper bound on the separation angle of a non-degenerate seam (junction) from the separation angles of the incident sheets. This ensures that during simplification, if a seam (junction) belongs to \mathcal{M}_θ^* , then at least one of the incident sheets will belong to \mathcal{M}_θ^* . Conversely, if all incident sheets do not belong to \mathcal{M}_θ^* , then the seam (junction) will not belong to \mathcal{M}_θ^* . Hence, it suffices to compute separation angles and test the sheets for pruning during simplification.

Lemma 2. For a non-degenerate \mathcal{M} , $\mathcal{M}_\theta \subseteq \mathcal{M}_f$.

Proof. Let $\mathbf{x} \in \mathcal{M}_\theta$, i.e. $\Theta(x) \geq \theta$. If \mathbf{x} is in the interior of a sheet f_i , then $\Theta(f_i) \geq \theta$. If \mathbf{x} is in the interior of a seam e_j , then by Lemma 1, $\Theta(f_i) \geq \theta$ for some sheet f_i incident on that seam. Thus, f_i will never be removed from the medial subset, and so e_j , being incident on f_i will be in \mathcal{M}_f . Therefore, $\mathbf{x} \in e_j$ will also be in \mathcal{M}_f . In the same way, Lemma 1 also implies that $\mathbf{x} \in \mathcal{M}_f$ if \mathbf{x} is a junction point. \square

6.2 Homotopy Preservation

Lemma 3. \mathcal{M}_f is homotopy equivalent to \mathcal{M} .

Proof. We perform induction on j . Our proof is complete if we show that \mathcal{M}_j is homotopy equivalent to \mathcal{M}_{j+1} , or, equivalently, that removing a frontier sheet f_i does not change the homotopy type. If both the seam set $\mathcal{S}(f_i)$ and the rim set $\mathcal{R}(f_i)$ are non-empty and connected, then the boundary of the sheet can have at most two components. If the boundary has one component, then the sheet is a topological disk, with a boundary consisting of two curves, the seam set and the rim set. If the boundary has two components, then one component must be the seam set, and the other the rim set. In that case, the sheet is an annulus, which can also be retracted onto the seam set.

The existence of a retraction means that there is a map $h : f_i \rightarrow \mathcal{S}(f_i)$ such that (a) the restriction of h to $\mathcal{S}(f_i)$ is equal to the identity on $\mathcal{S}(f_i)$, and (b) $g \circ h$ is homotopic to the identity on f_i , where g is the inclusion $\mathcal{S}(f_i) \rightarrow f_i$. We can then define $\hat{h} : \mathcal{M}_j \rightarrow \mathcal{M}_{j+1}$ to be equal to the identity on $\mathcal{M}_{j+1} \subset \mathcal{M}_j$, and equal to h on f_i . Then, if \hat{g} is the inclusion $\mathcal{M}_{j+1} \rightarrow \mathcal{M}_j$, it is clear that $\hat{h} \circ \hat{g}$ is equal to the identity on \mathcal{M}_{j+1} , and $\hat{g} \circ \hat{h}$ is homotopic to the identity on \mathcal{M}_j . Thus, the two spaces are homotopy equivalent to one another. \square

Lemma 4. \mathcal{M}_f is irreducible.

Proof. Let f_i be any frontier sheet in the final subset \mathcal{M}_f . Then f_i is labeled fixed, and either $\Theta(f_i) \geq \theta$ or f_i is an isolated component. Thus $\mathcal{M}_f \setminus \{f_i\}$ is not a subset of \mathcal{M}_θ , or does not have the same number of components as \mathcal{M} .

Let f_i be any non-frontier node in the final connectivity graph \mathcal{M}_f . If $\Theta(f_i) \geq \theta$, $\mathcal{M}_f \setminus \{f_i\}$ is not a subset of \mathcal{M}_θ . If $\Theta(f_i) < \theta$, then $\mathcal{M}_f \setminus \{f_i\}$ is not homotopy equivalent to \mathcal{M}_f .

We prove this by treating \mathcal{M}_j as a cell complex and considering its Euler characteristic. A 2-dimensional cell complex in \mathbb{R}^3 is a space that can be decomposed into open topological disks (faces), open curves (edges) and points (vertices) in such a way that the boundary of each face is a union of edges and vertices from the decomposition, and the boundary (that is, the endpoints) of each edge are vertices from the decomposition. Strictly speaking, the medial axis is not a cell complex, because the curves bounding frontier sheets are not in general part of the medial axis. However, we may add abstract edges without changing the homotopy type to construct a cell complex. The Euler characteristic, given by $\chi = F - E + V$ where F , E , and V are the numbers of faces, edges and vertices respectively, is a well-known homotopy invariant (see, e.g., [Spanier 1989]).

When we remove a sheet f_i from \mathcal{M}_j to get \mathcal{M}_{j+1} , we remove all of the faces, edges, and vertices of f_i except for the edges and vertices that are part of the seam set $\mathcal{S}(f_i)$. Thus, the change in Euler characteristic resulting from removing the sheet is given by

$$\chi(\mathcal{M}_j) - \chi(\mathcal{M}_{j+1}) = \chi(f_i) - \chi(\mathcal{S}(f_i)).$$

We wish to show that $\chi(f_i) - \chi(\mathcal{S}(f_i))$ is nonzero unless f_i is a frontier sheet.

The sheet f_i (which is connected by definition) is homotopy equivalent to a disk with n holes removed, for some n . The Euler characteristic of such a complex is given by $\chi = 1 - n$. The seam set consists of components of two types. There are loops, for which the number of vertices equals the number of edges, and $\chi = 0$. There are also unclosed chains of edges, for which there is one more vertex than edges, and $\chi = 1$. Therefore, $\chi(\mathcal{S}(f_i))$ cannot be negative, so that there are only two ways $\chi(f_i) - \chi(\mathcal{S}(f_i))$ can be zero. First we may have $\chi(f_i) = \chi(\mathcal{S}(f_i)) = 0$, in which case f_i is an annulus with connected, non-empty seam and rim sets. Second, we may have $\chi(f_i) = \chi(\mathcal{S}(f_i)) = 1$, in which case f_i is a (topological) disk, also with connected seam and rim sets. These cases are precisely the two kinds of frontier sets. \square

Together, the foregoing results show that $\mathcal{M}_f = \mathcal{M}_\theta^*$, as desired.




Figure 10: Brake Rotor Model (4.7k polygons): Rim curves are shown in green, $\theta = 150^\circ$ (b) The small holes in the center disappear in the θ -SMA, and the outer boundary becomes disconnected. (c) the θ -HMA the holes in the center are preserved, and the entire medial axis remains connected.

7 Implementation and Results

In this section we describe the implementation of our algorithm and highlight its performance on a number of complex benchmarks.

7.1 Implementation

We implemented the system in Microsoft Visual C++ and use OpenGL as the graphics API. All the timings reported in this paper were generated on a Pentium IV 2.4GHz PC with 2GB RAM running Windows XP. Our implementation for computing the Voronoi diagram is based on the techniques described in [Sud and Manocha 2005], from which the connectivity graph is extracted.

To test if a sheet f_i is a frontier sheet, we first extract a sub-graph of the connectivity graph. The sub-graph corresponds to f_i and its incident set of seam curves $\mathcal{S}(f_i)$. We then perform a depth-first-search on the sub-graph to determine the number of components in $\mathcal{S}(f_i)$ and in $\mathcal{R}(f_i)$. If sheet f_i is a frontier sheet, then number of components in $\mathcal{R}(f_i)$ and $\mathcal{S}(f_i)$ is 1. Iterative pruning during the medial axis simplification algorithm involves removal of nodes corresponding to the sheets and incident seam curves. The final graph captures the connectivity of the θ -HMA. The priority queue \mathcal{Q}_j is implemented as a heap.

7.2 Models and Results

We have applied our algorithm to compute the θ -HMA of polyhedral models of various sizes, ranging from 1000 triangles to 60k triangles. The complexity of the Blum medial axis ranged from 1.3k sheets to 89k sheets. Our benchmark models include CAD models with many sharp edges and high-aspect-ratio triangles. Such models can be relatively hard for medial axis algorithm that compute a point sampling on the boundary of the objects and a Voronoi diagram of the point samples. Some of the benchmark models have a high genus and holes that are preserved during medial axis simplification. We also tested our algorithm on synthetic benchmark models obtained by performing boolean operations with various solids.

For simplicity, in the figures we only show seam curves that are the intersections of three or more sheets. Also, maximally connected 2-manifolds have been grouped into one




Figure 11: Primer Anvil Model (4.3k polygons): Model boundary is shown in wireframe. The medial axis sheets are in blue, rim curves in green, seam curves in magenta, $\theta = 150^\circ$ (a) In the θ -SMA the sheets become disconnected, a thin sheet remains at the bottom (b) In the θ -HMA the sheets remain connected.

sheet. The models and their corresponding medial axes are shown in Figures 9 - 14. The polygonal meshes corresponding to the θ -SMA and θ -HMA have been smoothed using Taubin's algorithm [1995]. Table 2 lists the complexity of the polyhedral models, and of the original medial axis and corresponding simplifications θ -HMA and θ -SMA. The time to simplify the Blum medial axis to θ -HMA is also listed.

Model	Polys	θ ($^\circ$)	Num Sheets			Time (s)
			BMA	θ -HMA	θ -SMA	
Plate	990	150	1896	21	22	1.29
Rotor	4736	150	1365	41	17	1.23
Mount	2442	45	7455	536	283	2.43
Ridge-Rod	5012	120	30676	74	36	18.10
Anvil	4340	150	32102	4	4	17.51
Drivewheel	60712	150	89885	4	3	26.05

Table 2: Medial Axis Complexity: Polygon and sheet count of various models. θ is the separation angle (in degrees) used for computing θ -HMA and θ -SMA. Num Sheets refers to number of sheets in the exact Blum medial axis (BMA), and the simplified θ -HMA and θ -SMA. Time is the time in seconds used by Algorithm 1 to compute the θ -HMA from the Blum medial axis.

8 Analysis and Limitations

In this section we analyze the performance of our simplification algorithm. We highlight its computational complexity, give topological guarantees on the output and perform comparisons with some related algorithms for medial axis simplification.

8.1 Time Complexity

We provide the complexity of the algorithm as a function of the combinatorial complexity of the Blum MAT. A key step in our simplification algorithm is the operation to check if a sheet f_i is a frontier sheet. Let $|\mathcal{S}(f_i)|$ denote the number of seam curves incident on f_i , given by the number of sheets adjacent to f_i , and $\langle |\mathcal{S}(f)| \rangle$ be the average number of seam




Figure 12: CAD Mount (2.4k polygons), $\theta = 45^\circ$: The sheets emerging from the center of the vertical rod have low separation angle and have been removed. Note that the removal does not change the homotopy type.

curves of a sheet. Then the cost of checking if a sheet f_i is frontier is $O(|\mathcal{S}(f_i)|)$. We first present the cost of Algorithm 3. In the worst case, the frontier sheet check is performed on each sheet f_k adjacent to a sheet f_i , i.e. $|\mathcal{S}(f_i)|$ times. The cost of each frontier check is $O(|\mathcal{S}(f_k)|)$. The cost of adding or deleting a sheet from the priority queue \mathcal{Q}_j is $O(\log |\mathcal{Q}_j|)$. Hence the cost of a single instance Algorithm 3 is $\sum_{k=1}^{|\mathcal{S}(f_i)|} [O(|\mathcal{S}(f_k)|) + O(\log |\mathcal{Q}_j|)]$. Therefore, the cost of a single instance of Algorithm 2 is $O(\log |\mathcal{Q}_j|) + \sum_{k=1}^{|\mathcal{S}(f_i)|} [O(|\mathcal{S}(f_k)|) + O(\log |\mathcal{Q}_j|)] = O(|\mathcal{S}(f)|^2 + \log |\mathcal{Q}_j|)$. A sheet f_i can get added to the frontier set \mathcal{Q}_j at most $|\mathcal{S}(f_i)|$ times. Hence, the number of iterations in Algorithm 1 is at most $\sum_{i=1}^{|\mathcal{F}|} |\mathcal{S}(f_i)| = O(|\mathcal{F}| \langle |\mathcal{S}(f)| \rangle)$. Moreover, the size of the frontier set is bounded by the number of sheets, $|\mathcal{Q}_j| \leq |\mathcal{F}|$. Thus the total cost of the Algorithm 1 is $O(|\mathcal{F}| \langle |\mathcal{S}(f)| \rangle^3 + |\mathcal{F}| \log |\mathcal{F}| \langle |\mathcal{S}(f)| \rangle)$. Typically, $\langle |\mathcal{S}(f)| \rangle$ is a constant, the size of the frontier set is much smaller than $|\mathcal{F}|$, and the simplification cost is usually linear (or better) in $|\mathcal{F}|$.

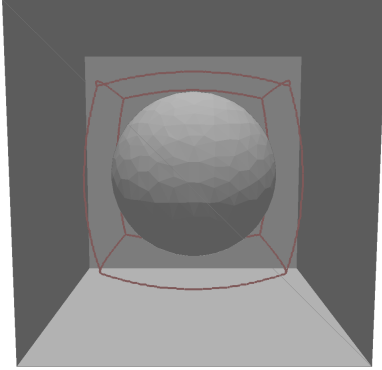


Figure 13: Cube with spherical void (1.5k polygons: A cut-out showing the cube and the spherical void in the center. The θ -HMA curves are drawn in magenta, $\theta = 180^\circ$. The θ -HMA remains connected, and preserves the void.

8.2 Comparisons with Other MAT Simplification Algorithms

In this section, we compare some features of our MAT simplification algorithm with prior techniques. There are many known approaches for computing and simplifying the medial axis. It is hard to make direct comparisons between

all these algorithms, as different algorithms make varying assumptions about the input and generate different kind of approximations.

The main feature of our approach is that we preserves the homotopy type of the medial axis while allowing for significant simplification of the medial axis. Our algorithm has been applied to polyhedral models as input, and faithfully captures the medial axis near sharp edges and corners in the input. Further, the algorithm preserves cavities corresponding to internal voids in the medial axis.

Some of the earlier analytic algorithms for MAT computation are based on tracing the seam curves [Culver et al. 1999; Reddy and Turkiyyah 1995; Sherbrooke et al. 1996]. These algorithms are relatively expensive and the worst case complexity is $O(n^3)$, where n is the number of features in the input solid. In practice, they have been applied to polyhedral models with few thousand triangles and compute the Blum medial axis and not a simplification of the medial axis. The adaptive subdivision algorithms [Vleugels and Overmars 1995; Etzion and Rappoport 2002] compute the generalized Voronoi Diagram, rather than a simplified medial axis. Further, these approaches may not be able to handle polyhedral models with internal voids.

The surface sampling approaches, such as [Amenta et al. 2001; Dey and Zhao 2002], take a point sampling on the surface as input and approximate the medial axis using the Voronoi diagram. Robust and efficient methods for computing the Voronoi diagram for point samples are well known. It is hard to make a direct comparison, as the output generated by these algorithms is different than our approaches which compute a distance field on a spatial grid. Many times the algorithms based on a point samples of the boundary may not be able to generate a good quality of approximation of the medial axis near the sharp features of the polyhedral model. The convergence of the Voronoi diagram to the medial axis with a finite discrete sampling has been proven, and extended algorithms have been proposed to generate good quality approximations for CAD models [Dey and Zhao 2002]. However, these methods guarantee a convergence to the medial axis in the limit, and may not provide topological guarantees on the computed medial axis approximation. Tam and Heidrich [2003] describe an iterative algorithm to simplify the medial axis of polyhedral models while avoiding some topological artifacts during the construction. Their work builds upon point sampling approaches, and has been applied to scanned models without many sharp features. There are no guarantees on the homotopy equivalence of the medial axis. Furthermore, the pruning algorithm needs to perform expensive global operations for topology preservation.

The λ -medial axis [Chazal and Lieutier 2004] provides a simplification of the medial axis for any open bounded shape in \mathbb{R}^n with homotopy equivalence to the original medial axis. The constraints on λ depend on the critical points in the gradient field of the distance function. An ϵ -sampling of the boundary of the shape is required, the choice of ϵ depends on different heuristics. Also, a single value of λ may not be appropriate to provide significant simplification for the entire shape. We are not aware of a practical implementation of this method. Attali et al. [2004] acknowledge these open issues and suggest a nested sequence of λ -Voronoi graphs with different values of λ for portions of the shape. In fact, a λ -medial axis with a small value of λ can be used as the original medial axis for our simplification algorithm, which

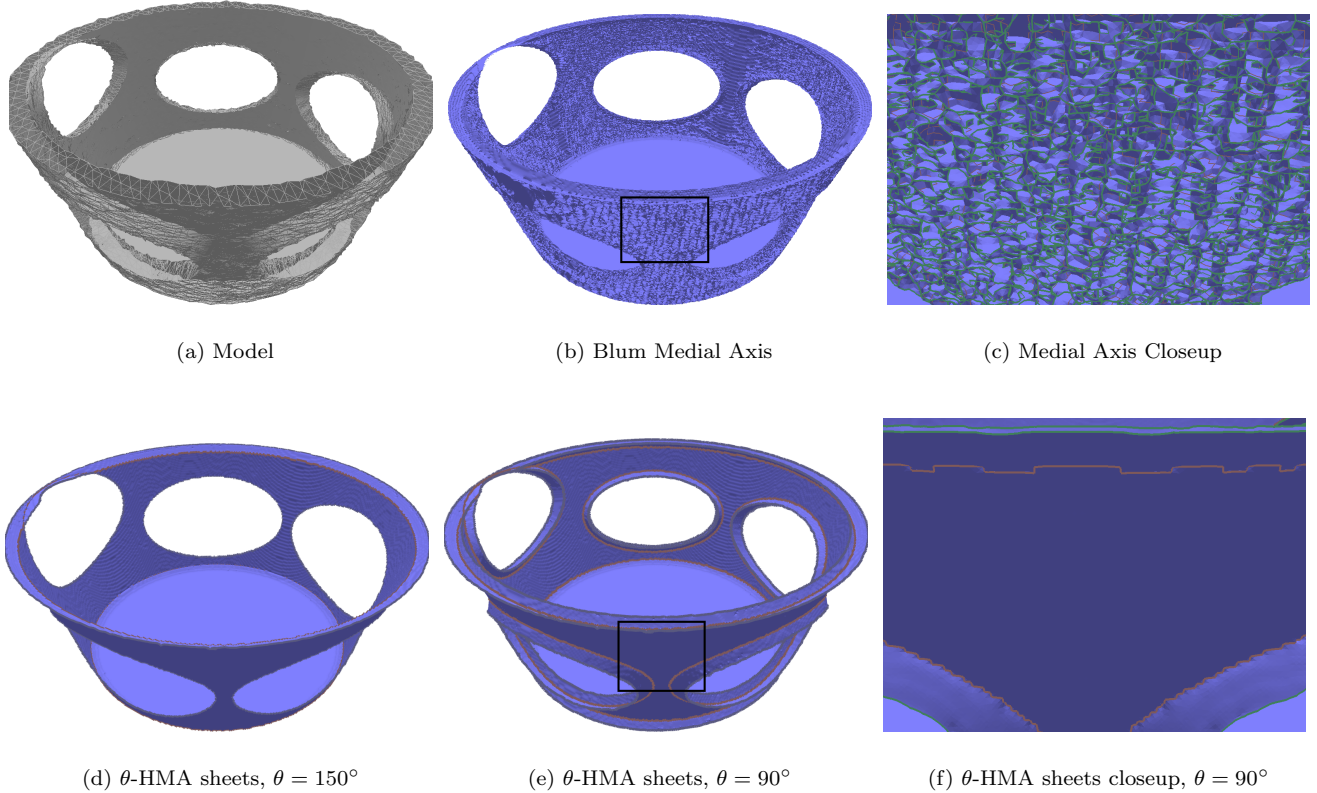


Figure 14: DriveWheel model (60k Polygons) and medial axis at different resolutions: Artificial noise was added to the model. Rim curves are shown in green, seam curves are shown in magenta. (a) The Model, with the front faces shown in wireframe (b) Blum medial axis, black box highlights the zoomed in region (c) A closeup highlighting the tiny sheets corresponding to the unstable parts. (d) Sheets of the θ -HMA, $\theta = 150^\circ$. Connectivity of the model and all holes are preserved. (e) Sheets of the θ -HMA, $\theta = 90^\circ$, black box highlights the zoomed in region (f) A closeup of the θ -HMA, $\theta = 90^\circ$, showing the stable subset of the medial axis.

subsequently allows significant simplification while preserving homotopy equivalence.

8.3 Limitations

Our approach has a few limitations. Our simplification algorithm depends on a spatial subdivision scheme to compute the Voronoi graph of the polyhedron. Similar to [Etzion and Rappoport 2002], the subdivision scheme generates a topologically accurate Voronoi diagram in absence of degeneracies. In degenerate configurations, the algorithm computes an approximate Voronoi graph which may not preserve homotopy equivalence to the original polyhedral model. The measure of stability that depends on separation angles, provides scale invariance but may retain noisy features if they exhibit high separation angles. Our simplification algorithm uses a greedy approach for pruning the unstable parts of the medial axis and a global minimum of the stability measure is not guaranteed. The elementary primitive in our pruning algorithm is a sheet, and the amount of simplification is influenced by the size of sheets. Finally, the simplified medial axis is not unique for a fixed value of θ , but depends on the pruning order. Actually, determining a unique order for iterative pruning for 3D models using topological constraints is still an open problem [Pizer et al. 2003].

9 Conclusions and Future Work

We have presented a medial axis approximation, the θ -HMA, that computes a stable subset of Blum’s medial axis, and preserves the homotopy type of Blum’s medial axis. The stability of the medial axis is guided by a separation angle criterion, which has been well studied. For polyhedral models, we present a formal characterization of the relationship between the stability of medial axis junctions and seams to the stability of incident sheets, based on separation angles. Our algorithm computes a bounded measure of stability of a medial axis sheet using discrete sampling. The construction of the θ -HMA is based on an iterative pruning algorithm which uses efficient local tests.

There are several avenues of future work. We would like to study other medial axis simplification criteria in conjunction to separation angles. We would also like to compute a θ -HMA with better guarantees on the global minimum of the stability measures, possibly leading to a unique pruning order. We would like to extend our algorithm to handle degenerate configurations of the Voronoi diagram, both in terms of topology of the θ -HMA and the stability relationship between incident parts of the medial axis. We are interested in applying the simplification algorithm to other medial axis approximations. Finally, we would like to explore applications of the θ -HMA such as mesh generation and shape analysis.

Acknowledgments

This research is supported in part by ARO Contract DAAD 19-02-1-0390, and W911NF-04-1-0088, NSF Awards 0400134, 0118743, DARPA and RDECOM Contract N61339-04-C-0043, DOD Prostate Cancer Research Program DAMD17-03-1-0134 and Intel Corporation. We thank Ajith Mascarenhas and the UNC GAMMA group for many useful discussions and support. We are also grateful to the reviewers for their feedback.

References

- AMENTA, N., CHOI, S., AND KOLLURI, R. K. 2001. The power crust. In *Proc. ACM Symposium on Solid Modeling and Applications*, 249–260.
- ATTALI, D., BOISSONAT, J.-D., AND EDELSBRUNNER, H. 2004. Stability and computation of the medial axis. In *Mathematical Foundations of Scientific Visualization, Computer Graphics, and Massive Data Exploration*. Springer-Verlag.
- BLANDING, R., BROOKING, C., GANTER, M., AND STORTI, D. 1999. A skeletal-based solid editor. In *Proc. ACM Symposium on Solid Modeling and Applications*, 141–150.
- BLUM, H., AND NAGEL, R. 1978. Shape description using weighted symmetric axis features. *Pattern Recognition* 10, 167–180.
- BLUM, H. 1967. A transformation for extracting new descriptors of shape. In *Models for the Perception of Speech and Visual Form*, W. Wathen-Dunn, Ed. MIT Press, 362–380.
- CHAZAL, F., AND LIEUTIER, A. 2004. Stability and homotopy of a subset of the medial axis. In *Proc. ACM Symposium on Solid Modeling and Applications*.
- CHAZAL, F., AND SOUFFLET, R. 2004. Stability and finiteness properties of medial axis and skeleton. *Journal of Control and Dynamical Systems* 10, 2, 149–170.
- CULVER, T., KEYSER, J., AND MANOCHA, D. 1999. Accurate computation of the medial axis of a polyhedron. In *Proc. ACM Symposium on Solid Modeling and Applications*, 179–190.
- CULVER, T. 2000. *Accurate Computation of the Medial Axis of a Polyhedron*. PhD thesis, Department of Computer Science, University of North Carolina at Chapel Hill.
- DEY, T. K., AND ZHAO, W. 2002. Approximate medial axis as a Voronoi subcomplex. In *Proc. ACM Symposium on Solid Modeling and Applications*, 356–366.
- DIMITROV, P., DAMON, J. N., AND SIDDIQI, K. 2003. Flux invariants for shape. In *International Conference on Computer Vision and Pattern Recognition*.
- DU, H., AND QIN, H. 2004. Medial axis extraction and shape manipulation of solid objects using parabolic PDEs. In *Proc. ACM Symposium on Solid Modeling and Applications*.
- ETZION, M., AND RAPPOPORT, A. 2002. Computing Voronoi skeletons of a 3-d polyhedron by space subdivision. *Computational Geometry: Theory and Applications* 21, 3 (March), 87–120.
- FOSKEY, M., GARBER, M., LIN, M., AND MANOCHA, D. 2001. A voronoi-based hybrid planner. *Proc. of IEEE/RSJ Int. Conf. on Intelligent Robots and Systems*.
- FOSKEY, M., LIN, M., AND MANOCHA, D. 2003. Efficient computation of a simplified medial axis. *Proc. of ACM Solid Modeling*, 96–107.
- GIBLIN, P., AND KIMIA, B. 2000. A formal classification of 3D medial axis points and their local geometry. In *Proc. IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 566–573.
- KIMMEL, R., SHAKED, D., KIRYATI, N., AND BRUCKSTEIN, A. M. 1995. Skeletonization via distance maps and level sets. *Computer Vision and Image Understanding* 62, 3, 382–391.
- KRUSE, B. 1991. An exact sequential Euclidean distance algorithm with application to skeletonizing. In *7th Scandinavian Conference on Image Analysis (SCIA '91)*, 517–524.
- LAM, L., LEE, S.-W., AND CHEN, C. Y. 1992. Thinning methodologies—A comprehensive survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 14, 9, 869–885.
- LEYMARIE, F. F., AND KIMIA, B. B. 2001. The shock scaffold for representing 3D shape. In *Visual Form 2001*, Springer-Verlag, 216–229. Lecture Notes in Computer Science, no. LNCS 2059.
- LIEUTIER, A. 2003. Any open bounded subset of R^n has the same homotopy type than its medial axis. In *Proc. ACM Symposium on Solid Modeling and Applications*, 65–75.
- MALANDAIN, G., AND FERNÁNDEZ-VIDAL, S. 1998. Euclidean skeletons. *Image and Vision Computing* 16, 317–327.
- MEYER, F. 1979. *Cytologie quantitative et morphologie Mathématique*. PhD thesis, École des Mines.
- MILENKOVIC, V. 1993. Robust construction of the Voronoi diagram of a polyhedron. In *Proc. 5th Canad. Conf. Comput. Geom.*, 473–478.
- NAF, M., KUBLER, O., KIKINIS, R., SHENTON, M., AND SZEKELY, G. 1996. Characterization and recognition of 3D organ shape in medical image analysis using skeletonization. In *MMBIA96, MEDIAL AXES*.
- OGNIEWICZ, R. L., AND KÜBLER, O. 1995. Hierarchic Voronoi skeletons. *Pattern Recognition* 28, 3, 343–359.
- PIZER, S. M., SIDDIQI, K., SZEKELY, G., DAMON, J. M., AND ZUCKER, S. W. 2003. Multiscale medial loci and their properties. *International Journal of Computer Vision* 55, 155–179.
- REDDY, J., AND TURKIYYAH, G. 1995. Computation of 3D skeletons using a generalized Delaunay triangulation technique. *Comput. Aided Design* 27, 9, 677–694.
- SHAHAM, A., SHAMIR, A., AND COHEN-OR, D. 2004. Medial axis based solid representation. In *Proc. ACM Symposium on Solid Modeling and Applications*.
- SHEEHY, D. J., ARMSTRONG, C. G., AND ROBINSON, D. J. 1995. Computing the medial surface of a solid from a domain Delaunay triangulation. In *Proc. Symposium on Solid Modeling and Applications*.
- SHEFFER, A., ETZION, M., RAPPOPORT, A., AND BERCOVIER, M. 1998. Hexahedral mesh generation using the embedded voronoi graph. *7th International Meshing Roundtable*, 347–364.
- SHERBROOKE, E. C., PATRIKALAKIS, N. M., AND BRISSON, E. 1996. An algorithm for the medial axis transform of 3d polyhedral solids. *IEEE Trans. Visualizat. Comput. Graph.* 2, 1 (Mar.), 45–61.
- SIDDIQI, K., B.B., K., AND SHU, C.-W. 1997. Geometric shock-capturing eno schemes for subpixel interpolation, computation and curve evolution. *Graphical Models and Image Processing* 59, 5, 278–301.
- SIDDIQI, K., BOUIX, S., TANNENBAUM, A., AND ZUCKER, S. W. 2002. Hamilton-Jacobi skeletons. *International Journal of Computer Vision* 48, 215–231.
- SPANIER, E. H. 1989. *Algebraic Topology*. Springer.
- SUD, A., AND MANOCHA, D. 2005. A simple algorithm for bounded-error approximation of voronoi diagrams of 3D polygonal models. Tech. rep., University of North Carolina-Chapel Hill.
- SURESH, K. 2003. Automating the CAD/CAE dimensional reduction process. In *Proc. ACM Symposium on Solid Modeling and Applications*, 76–85.
- TAM, R., AND HEIDRICH, W. 2003. Shape simplification based on the medial axis transform. *IEEE Visualization*.
- TAUBIN, G. 1995. A signal processing approach to fair surface design. In *Proc. of ACM SIGGRAPH*, 351–358.
- VLEUGELS, J., AND OVERMARS, M. 1995. Approximating generalized Voronoi diagrams in any dimension. Tech. Rep. UU-CS-1995-14, Department of Computer Science, Utrecht University.
- YANG, Y., BROCK, O., AND MOLL, R. N. 2004. Efficient and robust computation of an approximated medial axis. In *Proc. ACM Symposium on Solid Modeling and Applications*, 15–24.
- ZHANG, Y. Y., AND WANG, P. S. P. 1993. Analytical comparison of thinning algorithms. *Int. J. Pattern Recognit. Artif. Intell.* 7, 1227–1246.