Extraction and Visualization of Dimensions from a Geometric Model

Terence M. Bahlen[1], Willem F. Bronsvoort[2] and Allan D. Spence[3]

[1]Delft University of Technology, t.m.bahlen@student.tudelft.nl
[2] Delft University of Technology, w.f.bronsvoort@tudelft.nl
[3]McMaster University, adspence@mcmaster.ca

ABSTRACT

Many applications require information on dimensions of a geometric model, but these are usually not all explicitly present in the model. A method is introduced that extracts and visualizes information on dimensions from a geometric model. It first computes the medial axis of the model, and then uses this representation to determine certain dimensions in the model, in particular thicknesses and angles, and to visualize these on the boundary of the model. Presented results show that the method can visualize important dimensional information in a geometric model.

Keywords: geometric modeling, dimensions, medial axis transform, visualization.
DOI: 10.3722/cadaps.2010.579-589

1 INTRODUCTION

Dimensions, such as thicknesses and angles, play an important role in many CAD/CAM applications. An example is in the design of products to be manufactured by injection molding. There are many design guidelines related to thicknesses and angles for a proper CAD model, to avoid potential problem areas that cause warpage and surface sink marks during the injection molding. This research focuses on thickness and angle dimensions, because these two dimensions are sufficient to get the required insight into the dimensions of the geometric model for injection molding and many other applications.

A major problem is that usually not all dimensions are explicitly present in a model, and thus have to be made explicit when needed. Three situations where this occurs will be mentioned here. First, in a CAD system only a limited number of dimensions are explicitly input by the designer, and all other dimensions are implicitly defined. Second, when a CAD model is transferred from one system to another system, often the parametric design dimensions do not transfer, and only a geometric model remains. If dimensions are needed in the other system, these have to be recovered from the geometric model. Third, when an object model has been created with reverse engineering, e.g. with laser scanning, again no dimensions are present in the geometric model, and so, if needed, have to be computed.

In this paper, a method is introduced for extracting dimensions from a geometric model in a consistent way. It computes the medial axis, or skeleton, of the model, determines dimensions on the basis of this representation, and visualizes these in a straightforward way on the boundary of the model.

With the method, dimensions in a model can be verified, e.g. in designs for injection molding as mentioned above, but also certain properties, such as symmetry or the lack thereof, can be detected. For example, a model created with reverse engineering can contain multiple instances of the same feature and visualizing the dimensions of each of these features, allows the user to see whether they are indeed the same. One of the features may differ from the others, because it endured larger forces and has therefore a larger amount of wear.

Section 2 gives background information on the definition and possible computation of the medial axis transform of a model. Section 3 explains how we define dimensions on the basis of the medial axis and how they can be calculated. Section 4 describes how the medial axis is computed in our method and how the dimensions are visualized on the boundary of the model. Section 5 shows some results. Section 6 enumerates our conclusions on the method.

2 MEDIAL AXIS TRANSFORM

The medial axis transform was introduced in [1] to describe biological shapes and is sometimes referred to as the skeleton of an object. Since then the medial axis transform has been extensively researched and developed in many areas involving shape analysis. The medial axis is used in shape simplification [9], shape matching [8], routing in sensor networks [2], etc.

The definition of the medial axis transform given in [7] is as follows. *Let D be a subset of* R^n. *The medial axis is the locus of points which lie at the centers of all closed spheres which are maximal within D, together with the limit points of this locus. A closed sphere is said to be maximal in D if it is contained in D, but not a proper subset of any other sphere contained in D. The radius function of the medial axis of D is a continuous, real-valued function defined on the medial axis, whose value at each point on the medial axis is equal to the radius of the associated maximal sphere. The medial axis transform is the medial axis together with its associated radius function.* Fig. 1 illustrates an example of a medial axis in 3D, and Fig. 2 examples of 2D medial axes.

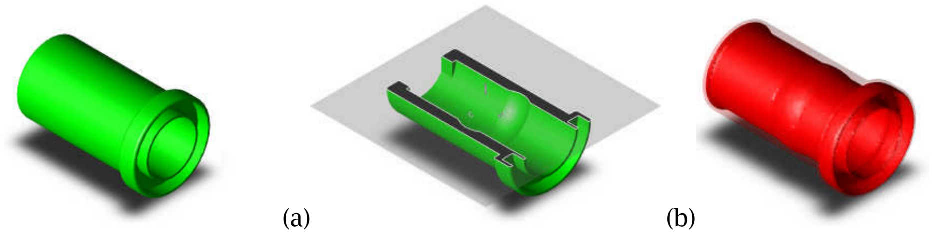

(a) (b) (c)

Fig. 1: An example of a 3D medial axis: (a) original model, (b) cut-away view, (c) the medial axis (with the original object boundary transparent), which shows information about relationships between faces, edges and vertices not apparent in the original model (a), without further inspection (b).

Another way to look at the medial axis is the grassfire analogy, which is a more dynamic interpretation of the medial axis. In this analogy, each point on the boundary is considered to be a point of fire, all points burning with the same intensity. The fire spreads perpendicular from the boundary point, at which it starts, towards the inside of the object, and burns with a constant rate of one unit distance per unit time. At time t the outer extent of the burned area is the curve parallel to the boundary offset by distance t. The medial axis consists of the closure of the quench points of the fires, i.e. the points where the fires meet and douse one another.

There is a one-to-one correspondence between the geometric model and its medial axis transform. To each geometric model belongs a unique medial axis transform, and vice versa. The geometric model can be determined from a medial axis transform by taking the union of all points on the medial axis and the associated maximal spheres. This reconstructability relies on the fact that, although the medial axes of two geometric models can be the same, the medial axis transforms, which include the radius functions, are different if the two models are different.

Many algorithms that compute the medial axis transform from a geometric model make use of some basic concepts related to the medial axis transform. These basic concepts involve a classification of the points on the medial axis, which is determined by so-called foot points. A foot point is a point of contact with the object boundary of the maximal circle, or in 3D the maximal sphere, of a point on the medial axis. In the footpoint, the circle is tangent to the object boundary. Depending on the shape of the object boundary, there can be either a discrete point contact or an area contact. In the latter case, the maximal circle coincides over an area with the boundary, i.e. the radii of curvature of the boundary and the circle are equal. Since the maximal circle is tangent to the object boundary, the lines from the medial axis point to its footpoints are perpendicular to the object boundary. Together with the notion of a foot point, comes the notion of a governor, which is the face, edge or vertex in which the foot point lies, as illustrated in Fig. 2.

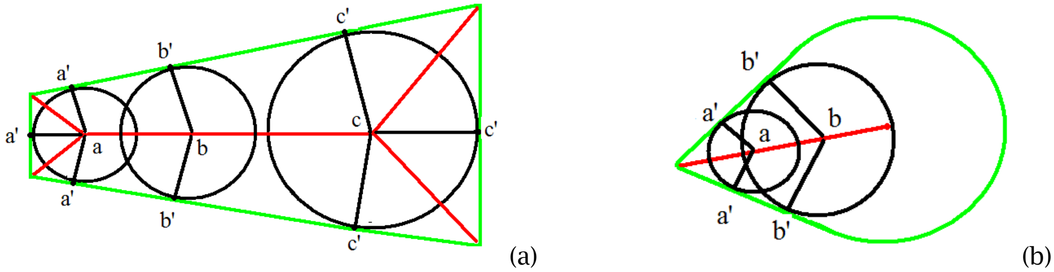

Fig. 2: Examples of 2D medial axis transforms; the red lines indicating the medial axis, the green lines the boundary of the object, and the black circles maximal circles, with their centers at *a*, *b* and *c* and their footpoints at *a'*, *b'* and *c'*; (a) having four edges and four vertices as governors, and (b) having one curved edge and one vertex as governors.

Based on the number of governors of a point on a 3D medial axis, the medial axis can be subdivided into several types of points, see [7]:

- a seam point: a point that has three or more governors (a seam is a connected curve of seam points)
- a seam-end-point: a point where a seam runs into the boundary
- a junction point: a point where three or more seams intersect
- a sheet point: a point with exactly two governors (a sheet is a connected surface of sheet points).

The types of points are illustrated in Fig. 3, where (a) is the original shape and (b) the corresponding medial axis.

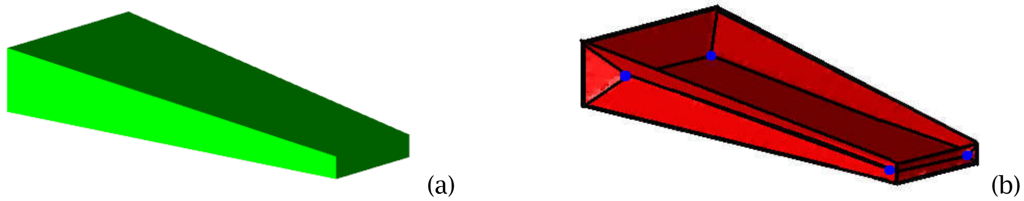

Fig. 3: Classification of points on a medial axis: (a) the original shape and (b) the medial axis; black lines indicate seams and edges, blue dots junction points, and red areas sheets. This medial axis contains 13 sheets, 4 junction points, 8 seam-end points and 12 seams.

One of the drawbacks of the medial axis transform is its sensitivity to small changes in the boundary, as illustrated in Fig. 4, where it can be seen that a small change in the object boundary can lead to a significant change in the medial axis. Some algorithms try to alleviate this problem by classifying offshoots of the medial axis as less significant, based on a substance measure, see [5]. These extraneous axes are then removed and only the major portions of the medial axis are maintained.

Some shapes exhibit special cases. The most common ones are a junction point with more than four governors and a seam with more than three governors. These special cases are incompatible with the generic case algorithms and thus separate, special treatment is needed for them.

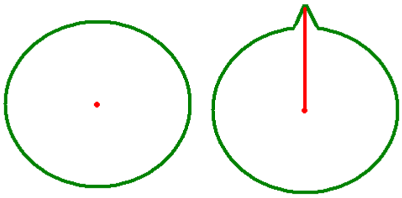

Fig. 4: The problem of sensitivity of the medial axis transform, illustrating that a small change in the object's boundary can lead to a significant change in the medial axis transform.

The computation of the medial axis transform is hard in general, because of numerical instabilities and because the medial axis transform is sensitive to small changes in the boundary. Therefore there are many algorithms that compute the medial axis transform for various types of shapes. These methods can be subdivided into continuous and discrete methods.

Continuous methods rely on the fact that the algebraic form is explicitly known for each sheet. The continuous method [7] calculates the medial axis transform for 3D polyhedra based on the relationships that exists between the various types of points. For example, if G1 is a governor of one sheet and G2 of another, then the seam connecting the sheets has as governors G1 and G2.

Discrete methods usually employ a surface sampling approach [4] or some spatial subdivision. Surface sampling algorithms represent the initial object as a dense cloud of sample points, presumed to be on the boundary. The medial axis is then approximated with a subset of the Voronoi diagram of the point cloud. One of the main issues when applying these discrete methods is the generation of an appropriate set of point samples on the boundary, to ensure a close approximation of the medial axis transform. Under-sampling may result in an approximation of the medial axis that is not accurate enough and thus in jagged results. One method that achieves a good approximation of the medial axis, using surface sampling, keeps a link to the original CAD model [3]. This allows for a better approximation using certain heuristics.

3 THE MEDIAL AXIS AND DIMENSIONS

In this section it will be shown how the medial axis can be used to extract dimensional information, in particular thicknesses and angles, from a geometric model.

The dimension thickness is, in general, not properly defined. For instance, when people are asked what the thickness of a 5 by 5 by 5 cube is, most would answer 5. However, for a cube-like model, which does not have all angles at 90°, there is ambiguity, since the thickness can now be measured in multiple directions.

The medial axis is a dimensionally-reduced structure that captures, in a consistent way, a notion of thicknesses and also of angles. As mentioned in Section 2, there is a one-to-one correspondence between the medial axis transform and the original model, which reduces to a relation between each medial axis point and its corresponding foot points and governors. We make use of this by extracting a dimensional value for each medial axis point, by looking at its foot points and governors, and assigning this value to these foot points.

To determine both types of dimensions, thickness and angle, a further classification of points on the medial axis is needed. Sheet points of which the governors are not directly connected give a good indication of thickness. Since angles are defined between edges or faces sharing a vertex or edge, sheet points of which the governors are directly connected, i.e. share a vertex or an edge, correspond to angle dimensions.

This distinction between thickness and angle sheet points subdivides the medial axis into areas capturing thicknesses and areas capturing angles, as illustrated for a 2D medial axis in Fig. 5(a), where the area *a-b* captures thicknesses and all other areas capture angles. The subdivision is based on the junction points of the medial axis. The foot points of these junction points (*a'* and *b'*), in their turn, subdivide the governors, and hence the boundary, into areas capturing thicknesses (blue areas) and areas capturing angles (grey areas). These areas thus correspond to specific areas of the medial axis. The sheets on a 3D medial axis are bounded by junction points, seam-end points and seams, as can be seen in Fig. 5(b). This figure illustrates the two types of sheets capturing thicknesses (blue) and angles (grey), respectively. In a similar way as in 2D, areas on the governors capturing thicknesses and angles are determined.

For the case in Fig. 5(c), there are no parts of the medial axis that can be directly classified as being a thickness or an angle area, since there is no junction point that divides the medial axis. Considering the area at *c*, which clearly is an angle, the question arises where the angle area should stop and the thickness area should begin. In this case, the medial axis at point *a*, and hence the foot points *a*' on the boundary, give a good subdivision, because the angle of the normals at the foot points *a'* changes compared to the angle at *c*. If the object boundary from *c* to the footpoints *a'* would consists of curves, the angle of the normals would change immediately and only point *c* would remain as the angle area. If this is not desirable, a value can be specified that defines how far along the curved edges the angle area is extended.

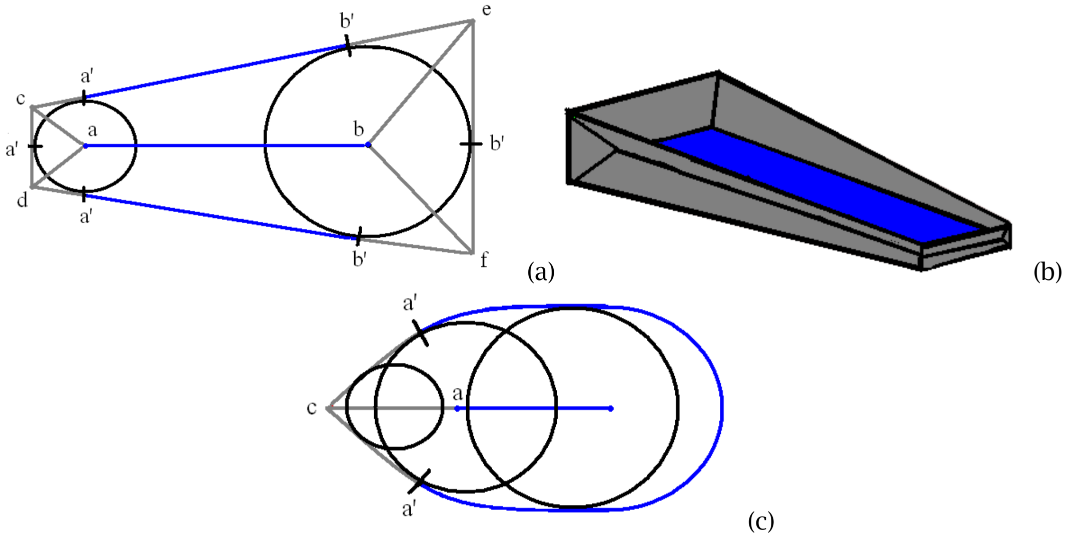

Fig. 5: The division between thickness and angle areas, (a) the 2D example of the medial axis, *a'* and *b'* indicating foot points of the junction points, dividing both the medial axis and the boundary into areas capturing thicknesses (blue) and areas capturing angles (grey), (b) an example in 3D of the two types of sheets on the medial axis, blue again capturing thicknesses and grey angles, and (c) an example of a medial axis without junction points, but still with a proper identification of areas capturing thicknesses (blue) and angles (grey).

The calculation of the actual thickness value if two foot points *m'* correspond to a point *m* on the medial axis, is illustrated in Fig. 6(a). We take the distance *d* between the two foot points. In the case of Fig. 5(c), there is ambiguity when calculating the thickness at the rightmost limit point of the medial axis, because the maximal circle and the object boundary coincide. The thickness can here be calculated between different combinations of two footpoints, which can result in different thickness values. In such cases, the radius value is taken as the thickness value, since this gives a good indication of the thickness. In general, this occurs at rounded regions of an object.

The calculation of the value for an angle *a* corresponding to a point on the medial axis, is illustrated in Fig. 6(b). We take the angle *b,* at medial axis point *m,* between the surface normals at the corresponding foot points *m'* and calculate the value at the actual corner as $a = 180^{\circ} - b$.

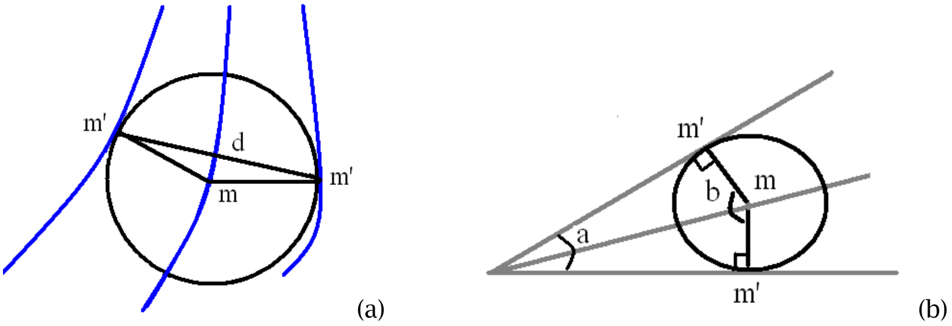

Fig. 6: The actual calculation of the thickness (a) and angle (b) values; for the thickness we take the distance *d* between the foot points *m′*, and for the angle we take 180°– *b*.

As a final step, the thickness and angle values are assigned to the footpoints, and the values are mapped to different color maps, to illustrate the difference between the two areas. Since the interval in which the thickness values can lie is not bounded, the values are taken relative to the maximum thickness within the object, when mapped to the thickness color map. The thickness values are interpolated across the surface, because the thickness changes from point to point on the surface. The angle values, having a limited range, are mapped to fixed color values, and an angle area gets a single color, since the angle value does not change within the area. How these colors are then used to create a complete visualization is discussed in Sections 4 and 5.

A previous approach for extracting and visualizing thickness information [6], visualizes the thickness according to the medial axis transform by taking the value of the radius function at each point of the medial axis, and using this value as an indication for thickness at the corresponding foot points. This approach therefore always indicates a thickness of value zero in the corners of the model, which is not useful at all. We, on the other hand, display the size of the angle at a corner and also have a better measure for thickness (the distance between foot points, instead of the radius of the maximal sphere).

4 COMPUTATION OF THE MEDIAL AXIS AND VISUALIZATION OF DIMENSIONS

This section will explain how the medial axis is computed, and how this representation is used for the visualization of the thickness and angle dimensions.

As mentioned in Section 2, the computation of the medial axis is hard in general. Because of that, and because the extraction of dimensions needs to work for any kind of geometric model, we decided to use a mesh as input. This is not a severe requirement, because most environments have meshing capabilities, either from a CAD model or from a point cloud. The algorithm expects a fairly uniform and dense (triangular) mesh, not only to get a good approximation of the medial axis, but also to get a good visualization of the dimensions. This means that in some situations remeshing may be necessary.

Governors, as mentioned in Section 3, are faces, edges or vertices in which a foot point of the medial axis lies. These governors are properly defined for a CAD model, since it is the entire face, edge or vertex in which a foot point lies. For a mesh, however, the governors are not properly defined, since faces and edges are subdivided into mesh elements. If there is a link from the mesh elements to an original CAD model, the proper governors can still be constructed, but if only a mesh is given, a different approach has to be used. Our approach depends on a so-called significant angle (usually between 0 and 15°) between the normals of two adjacent mesh elements, which specifies whether neighboring mesh elements are considered to be smoothly connected and thus part of the same governor. This subdivides the boundary into sets of mesh elements that we consider governors of the medial axis. During the construction of the governors, a link is created between each mesh element and the corresponding governor(s). This is needed for the correct classification of the medial axis points and their footpoints.

For the computation of the medial axis, all vertices of the mesh are considered to be footpoints. The algorithm traverses all footpoints, and finds the corresponding medial axis point for each footpoint by searching another foot point of that medial axis point. It is explained here for the 2D case; the algorithm for the 3D case is similar, but instead of circles, spheres have to be considered. The criterion for identifying the corresponding medial axis point is illustrated in Fig. 7(a). For a foot point f_a, the corresponding medial axis point m is always along the normal to the object boundary at f_a. Any other footpoint f_b of m has an equal distance to m, because of the nature of the medial axis. Using this criterion, the triangle $f_a m f_b$ can be constructed, in which the edges $f_a m$ and $f_b m$ have equal length and the angles a and b are equal. The maximal circle of m has the length of $f_a m$ as its radius, and two of its foot points are f_a and f_b.

It is worth mentioning that the algorithm in [4], mentioned in Section 2, uses a similar criterion based on the properties of the medial axis, but the focus there is on a convergence guarantee for the medial axis approximation from the Voronoi diagram, using a subset of Voronoi vertices called poles. Our approach is directed towards the identification of maximal spheres.

In the algorithm, given a footpoint f_1, a triangle is constructed for every other footpoint f, i.e. every other vertex of the mesh, by taking the same angle a at f as at f_1. If no triangle can be constructed, because $f_1 m$ and fm do not intersect, f is discarded. If a triangle can be constructed, the corresponding circle could be a maximal circle. Fig. 7(b) illustrates that this can result in multiple circles. The circle with the smallest radius value is taken, since all other circles contain at least one mesh point and are therefore not a maximal circle. In this way, the algorithm finds a medial axis point for each footpoint, and so we get a discrete approximation of the medial axis.

If a footpoint f_1 lies at a concave edge or vertex, there is a possibility of identifying multiple footpoints f with different medial axis points m and hence with different radius values. For such footpoints f_1, therefore no corresponding footpoint f is searched, i.e. they are skipped within the iteration over all footpoints. However, they can be a footpoint of another footpoint that is not at a concave edge or vertex, and will be found when for the latter the corresponding footpoint is searched. As will be discussed later, these concave footpoints are visualized in a different way.

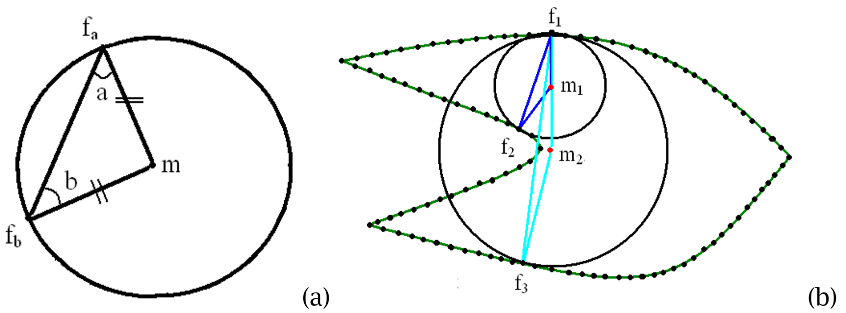

Fig. 7: The identification of a medial axis point, (a) the construction of the medial axis point m, given the surface normal at f_a, and the line $f_a f_b$, (b) an actual situation as it can occur within a mesh; the triangle $f_1 m_1 f_2$ is chosen, and not the triangle $f_1 m_2 f_3$, because the circle corresponding to $f_1 m_1 f_2$ has a smaller radius.

To make this approach plausible, we consider a continuous approach, in which not a discrete model (a mesh) but a continuous model is taken. Such an approach can identify a medial axis point for a footpoint f_a by taking a continuous sequence of circles tangent to the boundary at f_a, which all have their midpoint on the surface normal at f_a, with an increasing radius value. The smallest circle c that touches another part of the boundary, as illustrated in Fig. 8, is the circle of which the center is the corresponding medial axis point m, and the point where the circle touches another part of the boundary is another footpoint f_b. Our approach identifies all possible circles for the discrete model with their center along the surface normal at f_a and touching another part of the boundary. The smallest circle we find corresponds to the circle c of the continuous approach; compare Fig. 8 and Fig. 7(b).

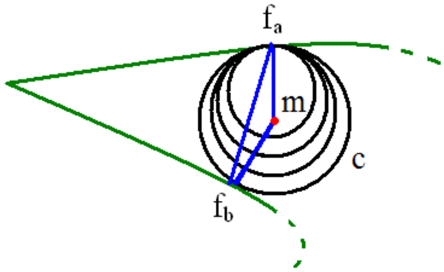

Fig. 8: A possible continuous approach in which the maximal circle c for footpoint f_a is the smallest circle tangent to the boundary at f_a that touches another part of the boundary.

For each medial axis point, the algorithm finds the corresponding footpoints and thus the corresponding governors; see begin of this section. All sheet points are classified as a thickness or an angle sheet point, according to the connectedness of the governors as mentioned in Section 3. Depending on the classification of a sheet point, either a thickness or an angle value is calculated, which is then assigned to its footpoints. These values are mapped to their corresponding color maps; for more details see Sections 3 and 5. If a footpoint lies at a concave edge or vertex, as discussed earlier, the value for that point is based on the dimensional values at its neighboring mesh vertices.

The triangular mesh is used to visualize the dimensions, i.e. the color values at the mesh vertices are interpolated across the triangles, resulting in a visualization of the dimensions across the whole boundary. Some mesh elements may have vertices in two different angle areas or in both a thickness and an angle area. The visualization between these and the neighboring mesh elements can give a jagged result, but as long as the density of the mesh is high enough, this is not really distracting.

5 RESULTS

This section will illustrate the general principle, show two examples with more details, and finally discuss some results of a parallel implementation.

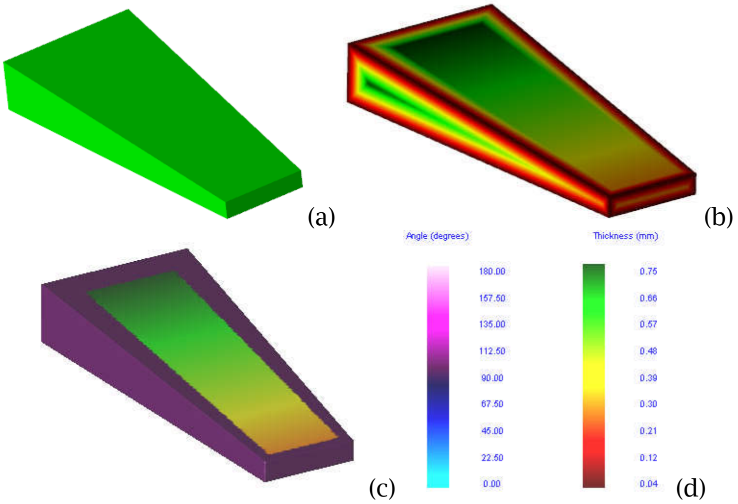

Fig. 9: Model (a), visualization of only thickness according to [6] (b), and visualization of both thickness and angle areas (c) according to different color maps (d).

The general principle of our approach is illustrated in Fig. 9. The model, only thickness visualization according to [6], and our visualization, including both thicknesses and angles, are displayed. The thickness values are mapped to colors ranging from red to green, where red was chosen for the thinner regions and green for the thicker regions. The angular values are mapped to colors ranging from blue to pink, to clearly illustrate the difference with the thickness areas. This simple example illustrates that the distinction between the two types of areas is clear.

A second example, see Fig. 10, illustrates that the differences between two parts can more easily be seen with our visualization. The part in Fig. 10(a) has rounded regions in two locations, whereas the part in Fig. 10(c) has straight edges. Our visualizations of the two parts, given in Fig. 10(b) and Fig. 10(d), are clearly different.

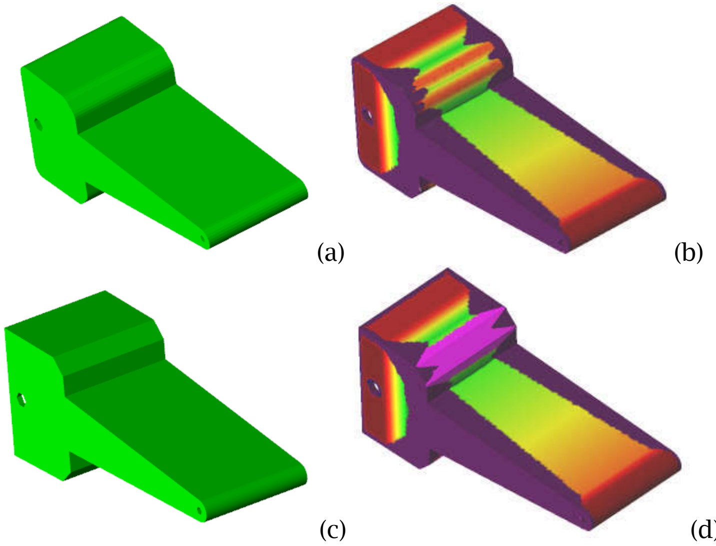

Fig. 10: Two parts: (a) has rounded regions in two locations, whereas (c) has straight edges. Their visualizations show thickness areas (b) and angle areas (d), respectively.

In the final example, see Fig. 11, a part contains four similar instances of the same feature (extruded cylinders). However, one of them is smaller than the others, which cannot be seen in a standard display of the part, see Fig. 11(a). The visualization in Fig. 11(b) shows a somewhat different color for one of the features, more clearly visible in the two close-ups. The features only differ 0.5 mm, hence the difference in visualization is also small, but large enough to be noticed.

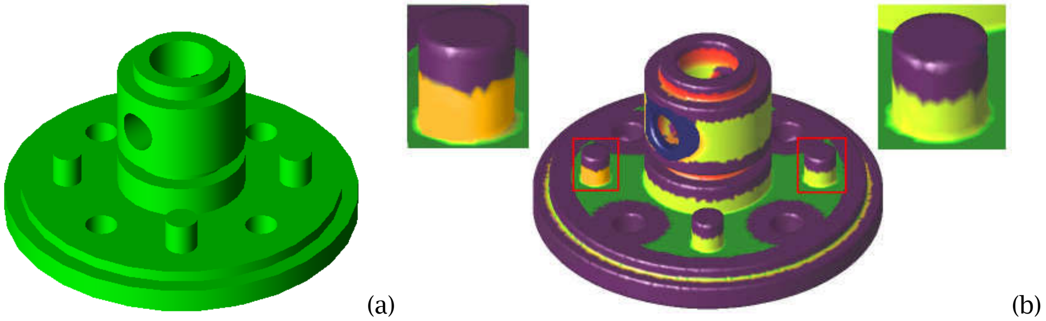

Fig. 11: A part that has several instances of the same feature (a), the extruded small cylinders that appear similar. However, one of them deviates by only 0.5 mm, and this is visible since the smaller feature has a somewhat different color (b).

Because the medial axis transform is computation intensive, especially if one wants a high degree of accuracy, the serial implementation was turned into a parallel one, with the help of openMP [10]. As can be seen in Tab. 1, the achieved speedup is very high, because the iteration over the footpoints can easily be divided among various threads. Implementation of a spatial subdivision scheme could further improve overall speed.

Model	*# Vertices*	*# Triangles*	*1 Core*	*2 Core*	*Speedup*	*4 Core*	*Speedup*
Fig. 9	38866	77728	141.5 s	70.6 s	2.0	**36.1 s**	**3.9**
Fig. 10 (a)	42247	84502	168.3 s	84.3 s	1.9	**43.2 s**	**3.8**
Fig. 11	52028	104080	261.3 s	131.6 s	1.9	**66.2 s**	**3.9**

Tab. 1: Results from the parallel implementation using openMP [10], illustrating that the scalability of the solution is very high, since the achieved speedups are close to what can maximally be expected. The results were tested on an Intel Core 2 Quad @ 2.67 GHz, 6GB.

6 CONCLUSIONS

This paper discussed a method that extracts and visualizes dimensions, specifically thicknesses and angles. The method computes the medial axis and determines the dimensions based on the information provided by this representation. The medial axis gives a consistent way of measuring the dimensions in question, and makes a distinction which dimension gets preference over the other to give a useful visualization. Since the method needs to work for any input and format, the calculation and visualization of the dimensions is based on a mesh. By making a parallel implementation of the method, the dimensions can be extracted and visualized fast, making sure that this does not become a bottleneck of the design process.

The visualization gives intuitive feedback, making a clear distinction between thickness and angle dimensions. It is capable of illustrating symmetry, whether or not surfaces are parallel, what is the inclination of a face compared to the rest of the part, what is the relationship between different parts of an object, e.g. does this part of the object have a thickness of around 40 % of the thickness encountered elsewhere in the object, etc. So, it can be used to verify several properties of the geometry of an object.

Future work involves the removal of the jagged results between the thickness and angle areas and handling special cases, mentioned in Section 2. An example of a special case is a cube-like object in which there are no sheet points capturing thickness, but only one seam that contains this information. Since seams are not mapped to the boundary, the visualization only offers information on the angle dimensions and not the thickness dimensions. In such situations, the visualization needs to be adjusted to incorporate some thickness information. In realistic models, however, these special cases rarely occur, and the visualization is already very useful.

ACKNOWLEDGEMENTS

This work was supported by the Canadian Natural Sciences and Engineering Research Council of Canada (Discovery Grant for A. Spence). Academic access to the Spatial Corp. ACIS 3D Mesh generation software, and the Tech Soft 3D HOOPS graphics software, is gratefully acknowledged too.

REFERENCES

[1] Blum, H.: A transformation for extracting new descriptors of shape, ACM Computing Surveys, 23(3), 1991, 345-405.
[2] Bruck, J.; Gao, J.; Jiang, A.: MAP: medial axis based geometric routing in sensor networks, Wireless Networks, 13(6), 2007, 835-853.

[3] Dey, T. K.; Hyuckje, W.; Zhao, W.: Approximate medial axis for CAD models, Proceedings of the Eighth ACM Symposium on Solid Modeling and Applications, June 16-20, Seattle, WA, USA, ACM, New York, NY, USA, 2003, 280-285.
[4] Dey, T. K.; Zhao, W.: Approximate medial axis as a Voronoi subcomplex, Computer-Aided Design, 36(2), 2004, 195–202.
[5] Katz, R. A.; Pizer, S. M.: Untangling the Blum medial axis transform, International Journal of Computer Vision, 55(2/3), 2003, 139-153.
[6] Lambourne, J. G.; Brujic, D.; Djuric, Z.; Ristic, M.: Calculation and visualization of the thickness of 3D CAD models, Proceedings of the International Conference on Shape Modeling and Applications, June 15-17, Cambridge, MA, USA, IEEE Computer Society, Los Alamitos, CA, USA, 2005, 340-344.
[7] Sherbrooke, E. C.; Patrikalakis, N. M.; Brisson, E.: An algorithm for the medial axis transform of 3D polyhedral solids, IEEE Transactions on Visualization and Computer Graphics, 2(1), 1996, 44-61.
[8] Siddiqi, K.; Shokoufandeh, A.; Dickinson, S.; Zucker, S.: Shock graphs and shape matching, International Journal of Computer Vision, 35(1), 1999, 13-32.
[9] Tam, R.; Heidrich, W.: Shape simplification based on the medial axis transform, Proceedings of the 14th IEEE Visualization Conference, October 22-24, Seattle, WA, USA, IEEE Computer Society, Los Alamitos, CA, USA, 2003, 481-488.
[10] OpenMP, www.openmp.org