

Visibility maps and spherical algorithms

Tony C Woo

By the extraction of ideas from computer vision, geometrical design and complexity analysis, a structure called *visibility* emerges. The paper describes a way in which a 3D workpiece is mapped onto the unit sphere, and its visibility is determined. For applications, manufacturing machines are classified by their degrees of freedom into *point*, *line* and *surface* visible processes. Algorithms for optimal workpiece orientation are then formulated as simple intersections on the sphere.

Keywords: manufacturing, computational geometry, data structures, algorithms, surface normals, curvature, mappings

Convexity is a notion that facilitates geometrical computation. An entity is said to be *convex* if the line segment pq that joins two arbitrary points p and q from a set lies entirely in the set. However, entities such as curves and surfaces are not always convex. The notion of *visibility* promises to extend beyond convexity. Take the art-gallery formulation¹, for example. Given a nonconvex polygon that resembles the floor plan of an art gallery, one asks how many guards are needed such that the entire interior is visible to the guards. Intuition suggests that the lines of sight, or *point-to-point visibility*, partition a nonconvex polygon into convex (visible) regions, hence bringing convexity to an entity that was nonconvex. This paper takes analytically complex entities, and maps them onto a unit sphere into analytically simple entities called *visibility maps* that turn out to be convex. Problems that are combinatorially complex are then addressed by a set of simple *spherical algorithms* that invoke the intersection between the visibility map and a point, a great circle or a spherical rectangle; these are also convex.

The establishment of a line of sight between a point on a workpiece and a point on an effector can be quite

useful in manufacturing. In numerical-control (NC) machining, assembly with robots, or inspection by coordinate-measurement machines (CMM), the computation for a tool path often relies on an intersection-free line of sight. Though the geometries of the effector of the manufacturing processes differ, the question about visibility in its most basic form can be stated as follows.

Basic enquiry about visibility: Given the geometries of the workpiece and the effector, *find* a number of orientations of the workpiece such that all the points of interest on the workpiece are visible to the effector.

That a *minimum* of orientations is crucial to productivity in manufacturing is evident. In an automotive application of the laser inspection of cylinder heads, the scanning time is approximately 20 min at 70 point/line/s. The setup time for fixturing and calibrating the workpiece is about 6 h. Since a reorientation requires the dismounting and setting up of the workpiece, the total number of orientations (to achieve complete visibility) should be minimized.

To be sure, there are manufacturing operations in which only a *single* setup is affordable or permitted. The use of the stereolithography apparatus (SLA) in layered manufacturing is one such instance. (In layered manufacturing, a facsimile of a 3D computer representation is fabricated layer by layer, stacking along the z axis. The material, which is a powder or a liquid, is selectively solidified by sintering or photocuring in the xy plane. For sheet material, a contour in the xy plane is created, and unwanted portions are discarded. However, regardless of the process, there is only one setup, and hence single implied orientation is permitted.) The choice of an orientation in layered manufacturing can make a difference. In an experiment in which the housing of a mobile phone was fabricated, two orientations showed a difference in fabrication time. In the orientation in which the antenna pointed upwards, it took over 22 h to fabricate the housing on an SLA 250. In the orientation

Department of Industrial and Operations Engineering, University of Michigan, Ann Arbor, MI 48109-2117, USA
Paper received: 10 November 1992. Revised: 26 April 1993

in which the antenna pointed sideways, the fabrication time was less than 10 h.

An intellectual curiosity emerges. If layered manufacturing is viewed as a physical realization of a triple integral, why should the order of integration matter? A layer or a crosssection S can be thought of as a double integral:

$$S = \iint_c dx dy$$

where c is a closed curve in the xy plane. The integration of the layers S along the z direction yields a volume V :

$$V = \int S dz$$

The same volume should result if, for instance, the crosssections are taken in the yz plane or any other plane. However, an examination of the SLA manual reveals an explanation for the apparent paradox: the physical integration rates (the optical scanning rate in the xy plane and the mechanical elevation rate in the z direction) differ. Therefore, the physical constraint of a machine may lead to a preference about the order of integration, and hence the orientation of the workpiece.

Now, given that there exists *asymmetry* in time for a certain family of manufacturing processes, it is reasonable to expect that problem formulations for optimal workpiece orientations can be induced. However, what about manufacturing processes that do not exhibit asymmetry? In particular, time and ordering are important if the process is *sequential*. (NC machining, CMM probing, robotic assembly, and layered manufacturing with an SLA are all sequential, in the sense that physical realization is achieved point by point, layer by layer, or component by component.) Stamping, die sinking, metal casting and plastic mould injection, to name but a few, may be viewed as nonsequential or *parallel* in the sense that the entire workpiece is effected in one 'fell swoop'. For such manufacturing processes, does the same basic enquiry about optimal workpiece orientation still hold? The answer turns out to be yes. Consider the stamping operation as an example. The direction of travel for the die implies an orientation of the workpiece. The demoulding process is another: the parting plane in a 2-piece mould implies a direction for demoulding, and this is the same for decasting. This single degree of freedom dictates the visibility of the punch or the removability of the workpiece from the cast or the mould. Intuitively, there appears to be an inverse relation between the degrees of freedom of the mechanism of a machine and its visibility.

This paper introduces a hierarchy that classifies the various manufacturing processes by their visibilities. The geometries are then mapped onto a unit sphere offering

a data structure called a *visibility map*. Algorithms for determining the optimal partitioning on the sphere are then developed, leading to answers to the basic enquiry about visibility.

VISIBILITY HIERARCHY OF MANUFACTURING PROCESSES

Visibility is meaningful when a pair consisting of a *source* and a *target* is considered. Geometrically, the target (or the workpiece) can be a point, a line segment, a curve segment, a planar polygon, or a surface patch. Similarly, the source (or the effector) can be modelled hierarchically by its topological dimensionalities: 0D points, 1D lines and curves, and 2D planes and surfaces, giving rise to a hierarchy, a summary of which is given in *Figure 1*. The geometries of the source are discussed in turn below.




Figure 1 Visibility hierarchy; (a) point visibility, (b) line visibility, (c) curve visibility, (d) plane visibility, (e) surface visibility [(a) numerically controlled machining, (b) wire-cutting electric-discharge machining, (d) stereolithography apparatus, (e) mould and die design.]

The probe of a CMM is spherical; the contact that it makes with a workpiece under inspection is often assumed to be a point on the sphere. In NC machining, because of the small stepping distance now made available by the machine-tool makers (of 10^{-5} – 10^{-6} m), the material-removal process can also be considered as a point process. NC and CMM are thus said to be processes that require *point visibility* (which is short for point-to-point visibility).

For wire-cutting electric-discharge machining (EDM), the effector is a line. Because of the small radius of the cutting wire, fine features can be created on the workpiece. Also, because of the process of discharging, there must be *line visibility* between the effector and the workpiece such that all the points in a segment of the cutting wire are visible to the corresponding points in a line segment on the workpiece. A noncontact variation of the line-visibility process is laser scanning for inspection. Diffracting a point source by a cylindrical lens into a line source, the laser casts stripes of lines and curves onto the surface of a workpiece, giving rise to *curve visibility*.

In layered manufacturing, there must be *plane visibility* for an entire layer. The energy source of a layered-manufacturing device can be conceptualized as a planar source, giving rise to a point-by-point correspondence between the source and the target. In the medical application of computer tomography, another noncontact variation, the slicing of the target can also be said to exhibit planar visibility.

The simplest example of *surface visibility* is probably sheet-metal stamping, in which the points on the sheet metal to be formed (the target) must maintain visibility with the corresponding points on the die (the source). Die-sinking EDM exhibits a similar characteristic of surface visibility. Similarly, die casting and plastics mould injection also demand surface visibility.

Degrees of freedom and number of setups

Recall that the amount of sequentialism varies from process to process. From *Figure 1*, it can be noted that, as the hierarchy traverses from surface to point visibility, the *degrees of freedom* (DOF) of the corresponding machines increases. While a sheet-metal stamping machine (possessing surface visibility) has only one DOF, an SLA (with planar visibility) has three DOF. Moving along the hierarchy, an NC milling machine (with point visibility) has up to five DOF.

Because of the way in which the hierarchy is established, the greater the number of degrees of freedom that a machine has, the simpler its *unit work* is. The unit work for an NC milling machine is to remove a small amount of material modelled as a point. In contrast, the unit work for a 1-DOF machine is the creation of a multiple number of surfaces in one motion. To enable the unit work to be done, the workpiece has to be set up

properly. It also may be interesting to examine the tradeoff between the number of setups and the degrees of freedom afforded by the machine.

Consider a 3-axis NC milling machine and a 5-axis NC milling machine. Both are point-visible processes, with different degrees of freedom. Suppose, to avoid excessive gouging, a ball endmill is used in both NC machines. As shown in *Figure 2*, because of the fourth and the fifth axes which pivot the tool, a workpiece can be machined on a 5-axis machine with a smaller number of dismountings, recalibrations and startups. Given that there is a choice of the machine to be used, it is reasonable to choose the one that offers fewer setups. This is a




Figure 2 3-axis machining compared with 5-axis machining

machine-selection problem that involves scheduling. However, given a fixed type of machine, say a 3-axis NC milling machine, it is also reasonable to ask whether the number of setups can be minimized. This leads to the following instance of the basic enquiry about visibility.

Enquiry about visibility with a high number of DOF: Given the geometries of the workpiece and the kinematics of the machine, *find* the minimum number of setups which require a reorientation of the workpiece.

Partial visibility and auxiliaries

A variation of the basic enquiry arises in manufacturing processes with a low number of degrees of freedom. Consider a surface-visible process such as moulding or casting. Suppose that the workpiece shown in *Figure 3a* had been moulded. To remove the workpiece, the choice of the *parting plane* is crucial. The direction that is orthogonal to the parting plane yields the direction for the lines of sight, giving rise to the single degree of freedom in the demoulding process. Now, the limitations of a single parting plane (a 1-DOF surface-visible process) should be noted here. Regardless of the direction of 'illumination', there can always be 'shadows' and hence *partial visibility* in a cavity, as shown in *Figure 3b*.

Indeed, analogously to the number of setups for a process with a high number of DOF, there can be multiple pieces in the mould or die for a process with a low number of DOF. In place of the multiple pieces, *auxiliaries* (called

'cores' or 'pins') are used in practice. *Figure 4* shows three kinds of auxiliaries in a 2-piece mould. To reduce the number of auxiliaries which incur expense, it is desirable to minimize them, leading to the second variation of the basic enquiry.

Enquiry about visibility with a low number of DOF: Given the geometries of the workpiece and the kinematics of the machine, *find* the minimum number of auxiliaries that compensate for partial visibility.

In the fourth section, the 'low' and 'high' numbers of DOF in these two enquiries are instantiated, and the two variations of the basic enquiry are once again generalized as intersections of spherical maps by a point, a great circle and a spherical rectangle. Before this is done, a discussion on the data structure of spherical maps is required.

VISIBILITY MAP: A DATA STRUCTURE

Complete visibility

Analogously to a hodograph², which is a map of tangents, a Gauss map (G map) is a map of normals on the unit sphere^{3,4}. The G map has many applications. It provides cutting-tool orientations⁵ and surface offsets^{6,7} for NC machining, and it assists in computing the aspect graph⁸, and in shape reconstruction from imaging⁹ in computer vision. In the following discussion, given a surface *S* in




Figure 3 Demoulding; (a) visibility, (b) partial visibility




Figure 4 Auxiliaries in mould; (a) side core, (b) side cavity, (c) form pin

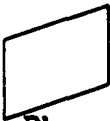



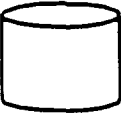



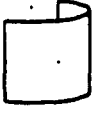







Surfaces	Gauss Maps	Surfaces	Gauss Maps
 Plane	 A Point	 Sphere	 Sphere
 Cylinder	 Great Circle	 Hemisphere	 Hemisphere
 1/2 Cylinder	 1/2 Great Circle	 Cone	 Small Circle
 1/4 Cylinder	 1/4 Great Circle	 Truncated Cone	 Small Circle

Figure 5 Examples of Gaussian maps

a parametric form, it is assumed that its G map exists. Figure 5 gives examples of G maps that are regions on the unit sphere. Note that each point in a G map denotes a direction from which a point in the surface is visible.

A visibility map (or V map) is an enhanced G map in the following sense. Like a G map, a V map is a spherical region. However, unlike a G map, any point in a V map denotes a direction such that the entire surface S is visible to its exterior. The idea of this powerful notion is simple. The normal at a point on a surface gives the direction in which the point is visible from infinity if the ray admits no intersection. However, the same point can be visible from many other directions, up to a hemisphere of directions bounded by the tangent plane at that point. This observation leads to a (rudimentary) procedure for computing a V map from a given G map. (The variable names are now joined by an underscore, to distinguish them from the procedure names.)

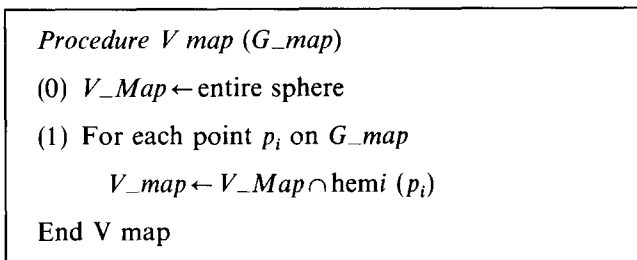


Figure 6 shows the V-map procedure in graphical form.

In Step 1, two calls are made: \cap for computing the intersection of two sets on a sphere, and $\text{hemi}(p_i)$ for providing a hemisphere whose 'pole' is at point p_i . The following observations can be made about these two calls in relation to the efficiency of computing a V_map . First, an intersection can be computed quickly if the two sets are convex. A hemisphere is a convex set. (This can be verified easily through central projection. The great circle bounding a hemisphere maps into a line in the plane of projection, and the hemisphere maps into a half plane.) Since the intersection of convex sets is always convex, the V_map generated by successive intersections of hemispheres at p_i must also be convex. Second, revisiting Step 1 of the V-map procedure, it is noted that the number of iterations is not specified. In particular, if a G_map is




Figure 6 Construction of visibility maps

to be sampled at p_i , should its entire interior be included? (Recall that a G_map is not necessarily convex, unlike a V_map . To ensure that a sample p_i falls into a G_map , it has to be tested by counting the even/odd number of intersections between a ray emanating from p_i and the boundary of the G_map , all adding to the time complexity.) Two steps are crucial to achieving savings in computation time. First, if only the boundary of a G_map is sampled, then there is no need for ray testing. (This sufficiency condition is shown below.) Second, the boundary of a G_map can be quite arbitrary, as it is, in general, nonconvex. If only the convex portion of the boundary of a G_map is sampled, there can be further savings. (The necessity is also established below.)

Suppose that p_i is any point in the interior of a G_map . Let C be a circle passing through it. C must make at least two intersections q_j and q_k with the boundary of the G_map that lie on either side of p_i . Construct three hemispheres whose poles are at the three points p_i , q_j and q_k on C . If $\text{hemi}(q_j)$ intersects $\text{hemi}(q_k)$, and they must, unless q_j and q_k are exactly 180° apart, it is seen that $\text{hemi}(p_i)$ is redundant in its contribution to the V_map . This establishes the sufficiency of sampling the boundary of a G_map . Now, the necessity of sampling the convex portion of the boundary follows a similar reasoning. Let p_i be any point on the nonconvex portion of the boundary, and the other two points q_j and q_k be on the convex portion. (The three points must coexist on any great circle c passing through p_i .) It is seen that, again, $\text{hemi}(p_i)$ is redundant, hence establishing the necessity criterion.

Enjoying convexity, the V-map procedure can be implemented in time $O(n \log n)$, where n is the number of samples on the spherical convex hull of a G_map . (For further details on the intersection of hemispheres, the construction of a spherical convex hull, and the determination of whether a given set of points lies in a hemisphere, the reader is referred to Reference 10.) A V_map can now be computed quickly. Further, its representation can be compacted. An incircle, a spherical circle inscribed in a V_map , gives a conservative approximation, and is much more compact in its representation. However, unlike circumscription, inscription is very difficult, because of the lack of 'extreme points'. In Reference 11, a 2-step procedure for finding the incircle of a V_map is described. The first step is to take the farthest point Voronoi diagram of a G_map to obtain the circumcircle for a G_map . Next, it is shown that its 180° complement yields the desired incircle of a V_map . In Reference 11, implementations of the G-map procedure and the V-map procedure are also given.

Partial visibility

A V_map obtained from a G_map which in turn is computed from a given surface S has the following property. Any point in a V_map yields a direction such

that the entire S is potentially visible to the exterior. For n surfaces to be visible, their corresponding V_maps must overlap. Sometimes, the overlap is empty. As shown in *Figure 7*, the planar surfaces yield single points as their G_maps ; their V_maps are hemispheres, the intersections of which are empty. Indeed, there is not a single direction from which the interior of the entire cavity is visible from the exterior.

Given that the visibility is *partial*, there will be 'shaded' regions for a given direction of illumination. (In the context of manufacturing, the shaded regions correspond to the auxiliaries, and the direction of illumination to the single degree of freedom, as in demoulding.) *Figure 8* shows two configurations. It can be seen that the geometry and the topology of the auxiliaries differ in these two instances. In the viewing direction of *Figure 8a*, the geometry of the auxiliary changes, but its topology remains invariant, as the viewing direction is changed slightly. The transition from *Figure 8a* to *Figure 8b* is termed the *critical viewing direction (CVD)*. Once a CVD has been determined, the augmented visibility maps (A_V_maps) can be constructed, in which hemispheres corresponding to the (planar) surfaces are intersected. (Note that, unlike a V_map , an A_V_map is valid only with respect to a CVD.) The (shaded) auxiliaries can be computed by taking the regularized set difference between the cavity in question and the (unshaded) *visibility polyhedron*¹². Because the visibility polyhedron is a proper subset of the cavity, an $O(n \log n)$ time algorithm¹³ can be used in computing for their intersection.

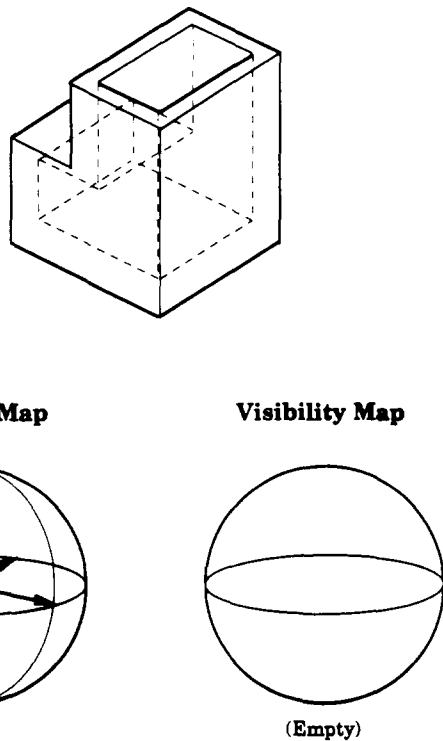


Figure 7 Empty visibility map

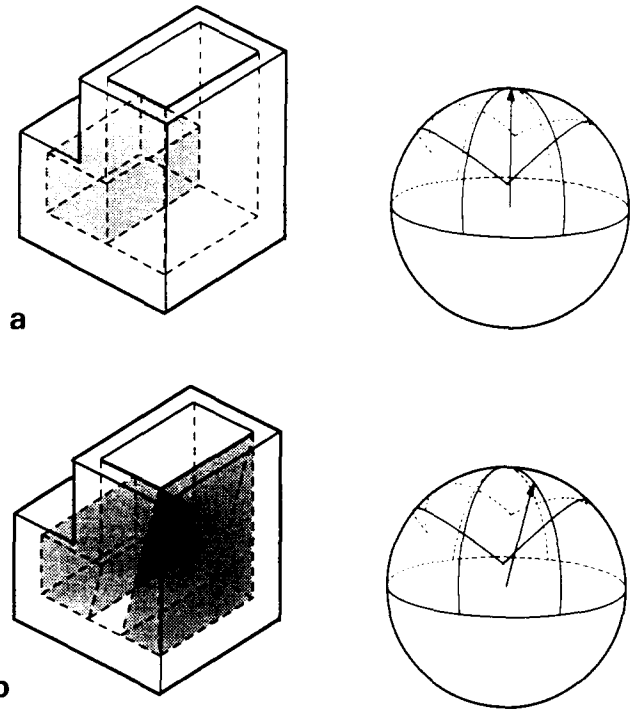


Figure 8 Augmented visibility maps

SPHERICAL ALGORITHM

The workpiece geometry having been transformed onto the unit sphere as V_maps , it is time to see how the mechanisms of the machines are related to the geometry of the workpiece. The basic idea is as follows. Suppose that the degrees of freedom of the mechanism of a machine can be represented, in terms of visibility, on a unit sphere. Then, given the V_maps which are also on a unit sphere, the problem of setting up a workpiece on a machine becomes that of rotating the sphere with the V_maps against that with the machine, according to some criteria. A criterion is *maximal intersection*. Rotate one of the two nested unit spheres such that their maps intersect maximally. (This corresponds to a setup in which the greatest number of workpiece surfaces are visible to the machine.) Delete those intersected V_maps , and continue until all the V_maps of interest are deleted. Relating the visibility hierarchy to the degrees of freedom is now appropriate. A summary is given in *Figure 9*. The entities on the sphere, a point, a segment of a great circle, and a spherical rectangle, denote the varying DOF of the machines. For clarity, the V_maps representing the workpiece geometry are omitted from the figure.

First, consider a 3-axis NC milling machine. Let the three DOF of the machine be denoted by x, y, z , meaning that they are the translations in the x, y and z directions. The tool axis of such a machine can be represented by a *point* p on the unit sphere. In other words, the visibility of the 3-DOF NC machine is in the direction that is opposite to the vector p represented as a point on the unit sphere. To machine a workpiece, the V_maps of the




Figure 9 Spherical algorithms

workpiece must be oriented by rotating the maps and keeping the point p fixed in such a way that the V_maps contain p . The rotation of the V_maps therefore corresponds to a setup. It is obvious that the point p may not lie in all the maps. Dismounting the workpiece, reorienting, and reclamping the workpiece (another setup) permits other maps to contain the point p , and hence the visibility and machining of other portions of the workpiece. This corresponds to three DOF and point visibility, as shown in Figure 9.

Next, consider a 4-axis NC milling machine with a rotary table. Let the fourth axis be denoted by θ . Figure 10 shows an example of such a 4-axis NC machine. Intuitively, the additional axis offers greater visibility than a 3-axis machine. Indeed, the additional degree of freedom from the rotary table can be represented by a great circle C on the unit sphere, if the table rotates through 360° . (In Figure 9, a segment of a great circle is shown.) To ensure the visibility of the workpiece on a four-DOF point-visible machine, the V_maps must be rotated in such a way that they are maximally intersected by C .

A 5-axis NC machine has an additional degree of freedom, denoted by ϕ . This rotary DOF is again a segment of another great circle C' on the unit sphere. (It is customary that the fifth axis does not rotate through

360° to prevent the fixtures and the table from tool collision and damage.) The combination of the fourth and the fifth axes forms a band, or just a spherical rectangle R on the sphere. To ensure visibility, the maps must again be rotated such that they are intersected by R .

The enquiries about visibility for machines with low and high numbers of degrees of freedom posed in the second section can now be succinctly stated in terms of the relation between the maps (which represent the workpiece geometry and the effector geometry) and three entities: a point, a great circle, and a spherical rectangle (which represent the increasing number of degrees of freedom in the mechanism of the machine), all on a unit sphere.

Visibility problem: Given the geometries of the workpiece and the effector, and the kinematics of the machine, find an orientation of the maps such that they are intersected by a point, a great circle or a spherical rectangle.

Returning to Figure 9, it can be seen that the visibility problem covers all the manufacturing processes that possess point visibility, line visibility, or surface visibility. Take the 1-DOF surface-visible process as an example, one instance of which is stamping. The relative orientation of the punch and the sheet metal on the die must be such that the direction of travel offers visibility to all the points on the workpiece, so as to result in the desired shape in the sheet metal bound on the die. The removal of a casting from a 2-piece cast is another example. The parting plane must be such that the casting can be removed in the direction normal to the parting plane. The direction of the punch and the direction of the casting can both be represented by a point on the unit sphere; the single DOF is denoted by z , without loss of generality.

Finally, the line-visible processes in the middle column of Figure 9 show the formulation of the visibility problem. Consider a 2-DOF line-visible process such as the wire-cutting EDM. Let the direction of the cutting wire be in the z axis, represented by a point p on the unit sphere. There are two degrees of freedom x and y for the cutting wire to travel. As shown in Figure 11, such a configuration of a wire-cutting EDM can create planar surfaces whose normals are in the x or y direction. Consider the surface facing in the x direction. Its G_map is a point g_1 on the unit sphere. Now, the same surface can be cut in many directions, some of which are shown by the crosshatching on the surface in Figure 11. In fact, the cutting wire can assume any of the directions for all the lines lying in the plane facing in the x direction; they form a V_map which is a great circle V_1 on the unit sphere. Similarly, all the possible cutting directions for the surface facing in the y direction yield a V_map of another great circle V_2 . To cut both surfaces, their V_maps must intersect, and they do, at the 'north pole'. Hence, the cutting wire must be oriented along the z axis, coinciding with the pole. No other orientation would

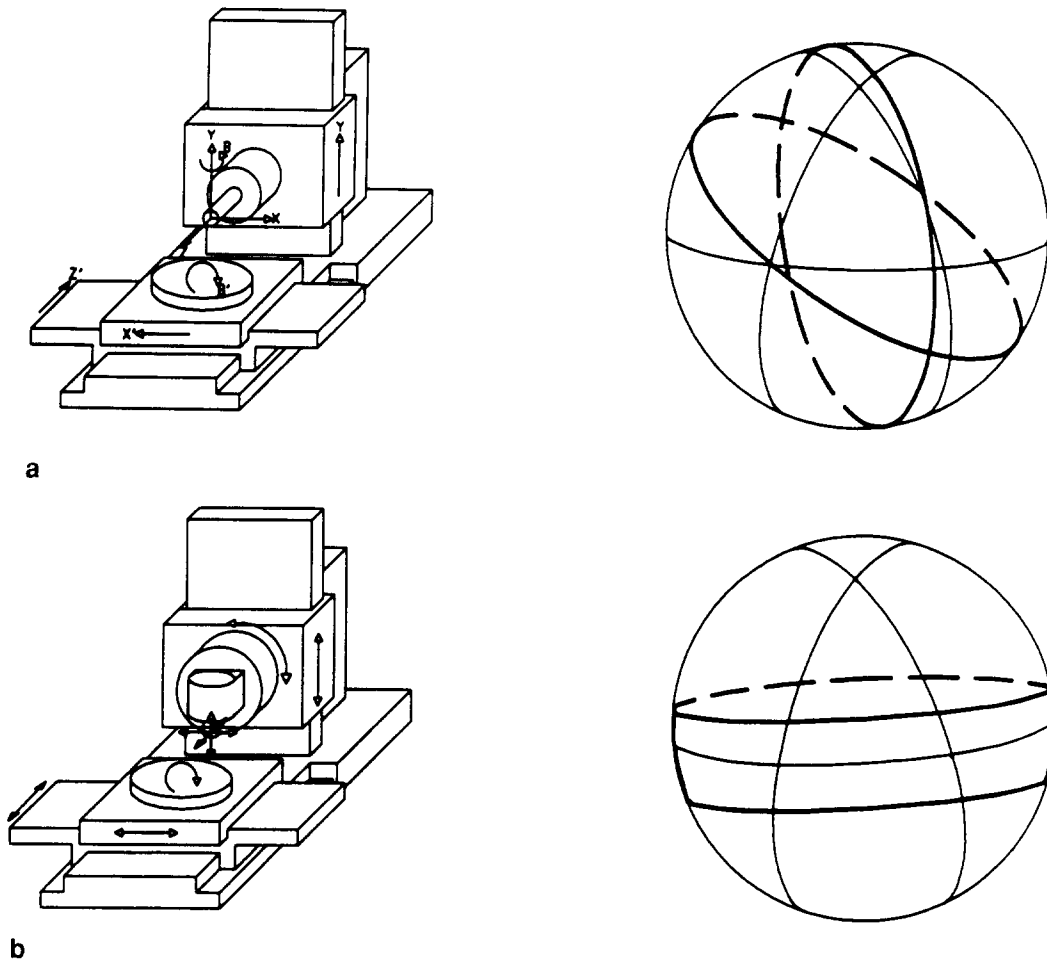


Figure 10 Machines; (a) 4-axis machine and great circle, (b) 5-axis machine and spherical band




Figure 11 Wire-cutting EDM and line visibility

suffice, as the intersection of the V_maps is a single point in this case. To accommodate more complex geometries of the workpiece, in wire-cutting EDM, the cutting wire needs additional degrees of freedom in rotation, giving rise to the 3-DOF and 4-DOF configurations in the middle of Figure 9, all of which have the generic form of the visibility problem.

The basic algorithms that enable the visibility problem to be solved are now easy to describe.

Intersection for constructing V_maps : For point-visible processes, the V_maps are constructed by intersecting hemispheres whose poles are points in the convex hull

of the G_maps of the workpiece surfaces. For line-visible processes, the V_maps are obtained by intersecting great circles whose duals are points in the G_maps . For surface-visible processes, the V_maps are the G_maps .

Intersection for orienting the workpiece: As in Figure 9, visibility is computed by intersecting a point, a great circle or a spherical rectangle with the V_maps , depending on the degrees of freedom that are available in the mechanism of the manufacturing machine.

Since there can be many surfaces, and hence many V_maps , it is desirable to intersect them optimally. A globally optimal solution to the problem of finding the minimum number of points (great circles or spherical rectangles) that intersect all the maps of interest is difficult. This has been proven¹² to be equivalent to the set-covering problem, which is strongly NP-hard¹⁴. However, a locally optimal solution to the problem of finding one point (or a great circle or a spherical rectangle) that intersects the maximal number of maps is in hand. In particular, intersecting a point with the greatest number of maps is equivalent to finding the densest hemisphere. An $O(nv \log n)$ time algorithm, where n is the number of maps and v is the total number of vertices in them, has recently emerged¹⁵. Now,

intersecting a great circle with the greatest number of maps on the sphere is equivalent to *sandwich cutting*. It can also be achieved in $O(nv \log n)$ time¹⁵ by exploiting a novel partitioning scheme which reduces infinitely many possible solutions to a finite number of equivalent ones in $O(nv)$ time¹⁶.

EXAMPLE

The machining of an object with 12 surfaces (represented as planar, ruled, Coons and Bézier surfaces (the mixture was intentional)) on a 4-axis milling machine is shown¹⁷. The geometry of these surfaces is shown in *Figure 12*.

The *G_map* for the top surface (there are ten on the side and one at the bottom) is shown as the dark region on a unit sphere, which has been ‘opened up’ for clarity (see *Figure 13*). This surface will be referred to as ‘surface number 12’. The *V_maps* for all 12 surfaces are given in *Figure 14*, as incircle approximations to the convex *V_maps*.

These 12 (circular) maps are seen to overlap each other. Map 12, for example, overlaps map 2, which, in turn, overlaps maps 6 and 7. However, there is no overlap between maps 12 and 6. A way of sorting out this intransitive relation is to use *clustering*. *Table 1* shows, for example, that map M_2 is in cluster x_2 with map M_{12} . It is also in cluster x_1 , with maps M_6, M_7 and M_3 (whose intersection is degenerated by the limiting view of *Figure*




Figure 12 Object with 12 surfaces




Figure 13 *G_map* for one surface




Figure 14 *V_map* for all 12 surfaces




Figure 15 Intersecting the maps

Table 1 Clustering of maps

Map	Cluster					
	x_1	x_2	x_3	x_4	x_5	x_6
M_1				1	1	
M_2	1	1				1
M_3	1				1	1
M_4			1			
M_5			1	1		
M_6	1					
M_7	1					
M_8					1	1
M_9						1
M_{10}				1	1	
M_{11}			1	1		1
M_{12}		1				

14. *Table 1* reveals that map M_2 is also in cluster x_6 , which contains a collection of maps, as shown on the right-hand side of *Figure 14*.

If a 3-axis machine is to be used, then five clusters x_1, x_2, x_3, x_4 and x_6 are needed, corresponding to five setups. If a 4-axis machine is to be used, only two setups are necessary. This is shown by the great circle shown by the broken line in *Figure 15*. This corresponds to the fourth axis, the rotating table. The broken line intersects all the maps, except for map M_{12} . Thus, the first setup on the 4-axis machine entails lining the cutter up (using three axes) with the great circle, and rotating the table (using the fourth axis). The second setup cleans up the remaining M_{12} .

CONCLUDING REMARKS

The purpose of presenting visibility maps as data structures that support spherical algorithms is to enable toolpaths to be automatically generated, in a context that lies beyond the traditional NC machining of sculptured surfaces. Elsewhere, researchers have begun to analyse 'accessibility'¹⁸ for CMMs on the basis of the notion of a visibility 'cone'. Recently, 'global accessibility'¹⁹ has been considered to account for the partial visibility discussed in the third section.

In this paper, a wide range of manufacturing processes are categorized by their point, line and surface visibilities. Further, the various degrees of freedom in the mechanisms of the machines gives a computational structure in terms of the maximal intersection of a point, a great circle and a spherical rectangle with the maps.

The above formulation has been made possible by taking into account the scientific foundations laid down by researchers in three seemingly disparate areas: computer vision²⁰, geometric design⁵, and computational complexity²¹. Curiously, these three areas were all given the same name, *computational geometry*, a term coined independently by Minsky and Papert, Robin Forrest²² and Michael Shamos²³. Perhaps this paper serves to unify these three connotations of computational geometry by offering a structure called *visibility*, from which efficient solutions to many classes of problem in design and manufacturing can now be systematically derived.

ACKNOWLEDGEMENTS

The author gratefully acknowledges the insight of the referees of this paper, and the financial support of the Fulbright Foundation, the US National Science Foundation (grants DDM-8903029 and DDM-9113791), and the University Research Fund of the Scientific Research Laboratory, Ford Motor Company, USA.


The work reported here would not have been possible if it were not for the diligence of the authors' colleagues at large: L L Chen (National Taiwan Institute of Technology, Taiwan), S Y Chou (National Taiwan Institute of Technology, Taiwan), J Gan (Nanyang Technological University, Singapore), S Y Shin (Korea Advanced Institute of Science and Technology, Korea), and K Tang (Schlumberger Technologies, USA) (listed in alphabetical order).

A preliminary version of this paper was given as an invited talk at the *SIAM Conference on Geometric Design* Tempe, AZ, USA (4–8 Nov 1991).

REFERENCES

- 1 O'Rourke, J *Art Gallery Theorems and Algorithms* Oxford University Press (1987)
- 2 Bézier, P *Numerical Control: Mathematics and Applications* John Wiley, UK (1972)

- 3 do Carmo, M P *Differential Geometry of Curves and Surfaces* Prentice-Hall (1976)
- 4 Hilbert, D and Cohn-Vossen, S *Geometry and the Imagination* Chelsea Publishing, USA (1983)
- 5 Faux, I D and Pratt, M J *Computational Geometry for Design and Manufacture* Halsted Press, USA (1979)
- 6 Farouki, R T 'The approximation of non-degenerate surfaces' *Comput. Aided Geom. Des.* Vol 2 No 1 (1985) pp 257–279
- 7 Chen, Y J and Ravani, B 'Offset surface generation and contouring in computer aided design' *ASME Trans. J. Mechanisms, Transmission & Automat. Des.* Vol 109 (Mar 1987) pp 133–142
- 8 Sripradisvarakul, T and Jain, R 'Generating aspect graph for curved objects' *Proc. IEEE Wkshp. Interpretation of 3D Scenes* Austin, TX, USA (27–29 Nov 1989) pp 109–115
- 9 Horn, B K 'Extended Gaussian image' *Proc. IEEE* Vol 72 No 12 (1984) pp 1671–1686
- 10 Chen, L L and Woo, T C 'Computational geometry on the sphere with application to automated machining' *ASME Trans. J. Mech. Des.* Vol 114 (Jun 1992) pp 288–295
- 11 Gan, J G, Woo, T C and Tang, K 'Spherical maps: their construction, properties, and approximation' *ASME Trans. J. Mech. Des.* (to appear) (also available as *Technical Report 91-3* Dep. Industrial & Operations Engineering, University of Michigan, USA (Feb 1991))
- 12 Chen, L L 'Visibility algorithm for mold and die design' *PhD Dissertation* Dep. Industrial & Operations Engineering, University of Michigan, USA (1991)
- 13 Tang, K and Woo, T C 'Algorithmic aspects of alternating sum of volumes. Part I: Data structure and difference operation' *Comput.-Aided Des.* Vol 23 No 5 (1991) pp 357–366
- 14 Megiddo, N and Tamir, A 'On the complexity of locating linear facilities in the plane' *Oper. Res. Lett.* Vol 1 No 5 (1982) pp 194–197
- 15 Chen, L L, Chou, S Y and Woo, T C 'Separating and intersecting spherical polygons for computing visibility on 3-, 4-, and 5-axis machines' *ACM Trans. Graph.* (to appear) (also available as *Technical Report 91-9* Dep. Industrial & Operations Engineering, University of Michigan, USA (Mar 1991))
- 16 Tang, K, Woo, T C and Gan, J 'Maximum intersection of spherical polygons and workpiece orientation for 4- and 5-axis machining' *ASME Trans. J. Mech. Des.* Vol 114 (Aug 1992) pp 271–276
- 17 Gan, J G 'Spherical algorithms for setup orientations for workpieces with sculptured surfaces' *PhD Dissertation* Dep. Industrial & Operations Engineering, University of Michigan, USA (1990)
- 18 Spyridi, A J and Requicha, A A G 'Accessibility analysis for the automatic inspection of parts' *Proc. IEEE Int. Conf. Robotics & Automation* Cincinnati, OH, USA (13–18 May 1990) pp 1284–1289
- 19 Requicha, A A G 'Computing global accessibility directions for the faces of a solid' *SIAM Conf. Geometric Design* Tempe, AZ, USA (4–8 Nov 1991)
- 20 Minsky, M and Papert, S *Perceptions: An Introduction to Computational Geometry* MIT Press (1972)
- 21 Preparata, F P and Shamos, M I *Computational Geometry – An Introduction* Springer-Verlag (1985)
- 22 Forrest, A R 'Computational geometry' *Proc. Roy. Soc.* Vol 321 Ser. 4 (1971) pp 187–195
- 23 Shamos, M I 'Computational geometry' *PhD Dissertation* Dep. Computer Science, Yale University, USA (1978)



Tony C Woo is a professor of industrial and operations engineering and a professor of mechanical engineering and applied mechanics. His research interest is in the development of efficient data structures and algorithms for design and manufacturing. He has been a Fulbright Scholar at the NTH, Trondheim, Norway, and has served at the US National Science Foundation, Washington DC, USA.