



Computing global visibility maps for regions on the boundaries of polyhedra using Minkowski sums

Min Liu^{a,*}, Yu-shen Liu^b, Karthik Ramani^b

^a Institute of Manufacturing Engineering, Tsinghua University, Beijing, 100084, China

^b School of Mechanical Engineering, Purdue University, West Lafayette, IN, 47907, USA

ARTICLE INFO

Article history:

Received 30 May 2008

Accepted 26 March 2009

Keywords:

Global visibility map
Global occlusion map
Local visibility map
Minkowski sum

ABSTRACT

A global visibility map is a spherical image built to describe the complete set of global visible view directions for a surface. In this paper, we consider the computation of global visibility maps for regions on the boundary of a polyhedron. Both the self-occlusions introduced by a region and the global occlusions introduced by the rest of the surfaces on the boundary of the polyhedron are considered for computing a global visibility map. We show that the occluded view directions introduced between a pair of polyhedral surfaces can be computed from the spherical projection of the Minkowski sum of one surface and the reflection of the other. A suitable subset of the Minkowski sum, which shares the identical spherical projection with the complete Minkowski sum, is constructed to obtain the spherical images representing global occlusions. Our method has been successfully tested on many CAD models. It extends the previous methods for computing global visibility maps using convex decomposition, and it exhibits a better performance.

© 2009 Elsevier Ltd. All rights reserved.

1. Introduction

1.1. Background

A point \mathbf{p} in a 3D scene is visible along a view direction \mathbf{v} if the ray starting from \mathbf{p} goes to infinity, in the direction $-\mathbf{v}$, without intersecting any entities in the scene. The notion of point visibility can be extended to higher dimensions, such as 1D curves or 2D surfaces. A curve or a surface is visible along a direction \mathbf{v} if all the points on the curve or the surface are visible along the direction \mathbf{v} . Given a view direction, finding the set of entities that are visible along this direction is a well-known problem in computer graphics and is referred to as the *hidden surface/line removal* problem (under orthographic projection) [1,2]. The inverse problem is called a *visible set* problem, defined as: given an entity in the scene, find the set of view directions along which the given entity is completely visible. In this paper, we restrict our attention to the visible set calculation for polygonal surface meshes. For a sub-region on the boundary of a polyhedron, we calculate the complete set of view directions along which all points on the region can be seen from the

exterior. The mapping of the visible directions on the unit sphere is a spherical image called the global visibility map (GVM).

GVM calculation for polyhedral models has broad applications in process planning for many manufacturing processes such as 3-axis Numerical Controlled (NC) machining, inspection by Coordinate Measuring Machines (CMMs) [3–5], and mold design [6,7]. Given a region, its GVM associates the set of global interference free access directions by a tool or a probe, if one can ignore the small radii of the tool or the probe. In mold design, the GVM of a surface associates a set of candidate parting directions with the corresponding mold piece that forms the surface. In general, the GVM computation serves as a query base for further process planning applications in which the general objective is to minimize the number of setups or the number of mold pieces. In addition, GVM is inherently related to some machine vision-related tasks such as object recognition and object inspection. These tasks require the computation of the visible sets, with the viewpoints located at finite positions [8,9].

1.2. Related work

Earlier work on visible set calculation of 3D objects dealt with the visibility map (or visibility cone) calculation for surfaces, which was first introduced by Kim et al. [10,11] for computing visibility maps of Bézier surfaces and further explored by Elber and Cohen [12] to cover the visibility calculation for free-form surfaces. The general idea behind computing a visibility map is to compute

* Corresponding address: Institute of Manufacturing Engineering, Tsinghua University, Rm. 2402B, Building 9003, Beijing, 100084, China. Tel.: +86 10 6278 5548.

E-mail addresses: minliu@mail.tsinghua.edu.cn (M. Liu), liuyushen00@gmail.com (Y.-s. Liu), ramani@purdue.edu (K. Ramani).

the dual image [10] of the Gaussian map on the unit sphere. It ignores the fact that other surfaces on the object boundary may occlude the visibility of the target surface. Therefore, the visibility map applies only to applications such as the local tool path planning in machining or mold design, wherein a region under consideration is the whole pocket [13–15].

Approximation methods for computing the visible set for a 3D object were designed to build a visibility matrix based on the discretized approximation of the object and solution spaces [5,8]. The elements of the visibility matrix encode the visibility of discrete surface entities from each sampled point in a quantized viewing sphere, which can be calculated using hidden surface removal from given viewpoints. The computation can be done using either ray casting techniques [8] or z-buffer method through graphics hardware [5,16,17]. However, the drawback of this sample-and-test approach is its incompleteness. No matter how good the sampling of the solution space, it is still possible to falsely report some visible entities as invisible when the whole visible set contains no sampled viewpoints or directions.

A few complete visibility methods have treated the solution space as continuous and have sought an exact solution [18–23]. The *aspect graph* approaches partition viewpoint space according to the qualitative aspect of the view [24], while the *visibility complex* deals with the creation of a data structure that encodes all visibility relations between objects in the collection [25]. Dhaliwal et al. [22] presented an algorithm for computing the global visibility for triangle meshes, based on the occlusion calculation between a pair of triangle facets. Our previous work [18] extended Dhaliwal's work to the occlusion calculation between a pair of convex facets. Besides, we considered the completeness of a GVM and treated a GVM as a spherical image which might degenerate to lower dimensions. A similar idea for calculating the occlusions between a pair of convex facets was explored from a different perspective in [23]. The bottleneck of the above *convex decomposition based method* [18,22,23] exists in the complex intermediate spherical arrangement calculation for obtaining a final GVM. The method needs to perform union operations among all spherical polygons representing pairwise occlusions among all convex facets.

For computing the GVMs of general polyhedral surfaces, we use the Minkowski sum operations. The Minkowski sum of two sets is the set of all possible vector sums of a point from one set and a point from the other. The Minkowski sum is a useful geometric tool that has been used in a number of applications such as robotics [26,27], morphing [28], assembly planning [29], and CAD/CAM [30,31]. The computation of the exact Minkowski sum of two convex objects, such as planar polygons or polytopes, is known [32]. For computing the Minkowski sums of general polyhedra, a prevailing approach is to decompose each polyhedron into convex pieces, then compute the pairwise Minkowski sums of pieces of the two, and finally compute the union of pairwise sums [33]. The limitations are that the optimal convex decomposition is known to be NP-hard itself and there is also a bottleneck in computing the union of pairwise Minkowski sums. Until recently, Hachenberger [34] presented the first robust implementation of the 3D Minkowski sum of two non-convex polyhedra, based on convex decomposition. Another approach for computing the Minkowski sums of general polyhedra is based on the convolution of the boundary of two polyhedra. Convolution was proposed originally on 2D planar tracings [35] and later extended to 3D on polyhedra tracings [36]. It can be considered as an *implicit representation* of the Minkowski sum, while being smaller and easier to compute. Kaul and Rossignac [28] used weighted Minkowski sums to construct a smooth interpolation between two polyhedra for animation. Ramkumar [37] and Wein [31] used convolution to compute the outer boundary of the Minkowski sum of two polygons.

In this paper, we extend the previous works [18,22,23] on the global visibility calculation. The terms, *Global Visibility Map*, *Local*

Visibility Map, and *Global Occlusion Map* are formally defined to differentiate the set of self-occlusion free view directions, the set of total occlusion free view directions, and the set of occluded view directions introduced by other obstacles in a global environment. A general method for computing the global visibility maps for regions on the boundary of a polyhedron is proposed. We borrowed the M-region idea introduced by Halperin et al. [29] for global occlusion map computation using Minkowski sums. We show that the global occlusion maps of a polyhedral surface can be computed from the spherical projection of the Minkowski sum between the reflection of the given surface through the origin, and the rest of the surfaces on the given polyhedron. We differentiate our method from the M-region [29] by the way the spherical projections of the Minkowski sum are computed. Specifically, we construct a suitable subset of the Minkowski sum which shares the identical spherical projection with the complete Minkowski sum; thus, our method avoids the explicit Minkowski sum boundary calculation which is computationally expensive. Compared to convex decomposition-based methods, our method reduces the size of the intermediate spherical arrangement generated for computing a GVM in general cases.

2. Preliminaries

Throughout this paper, \mathcal{P} denotes a polyhedron, which is a watertight solid bounded by a piecewise linear surface. We denote the boundary surface of \mathcal{P} by $\partial\mathcal{P}$ and the interior of \mathcal{P} by $Int(\mathcal{P})$. The symbols $\partial(\cdot)$ and $Int(\cdot)$ are also used for denoting boundaries and interiors of general sets. $\partial\mathcal{P}$ is a closed polyhedral surface and is required to be an orientable 2-manifolds. S denotes a polyhedral surface with border curves and S also represents a sub-region on $\partial\mathcal{P}$ in this paper.

The symbols V_S , E_S , and F_S refer to a vertex, an edge, or a face of a polyhedral surface S , while $S.vertices$, $S.edges$, and $S.faces$ refer, respectively, to the set of all vertices, edges, and faces of S . Corresponding notations are used for any other polyhedral surfaces referred to in this paper. Each face F is associated with an *outward normal*, denoted \mathbf{n}_F . The corresponding symbol of an outward normal is used for any other face referred to in this paper.

A face F of $\partial\mathcal{P}$ is a *hull facet*, if F lies on the convex hull of \mathcal{P} ; otherwise, F is a *non-hull facet*. A maximal connected set of non-hull facets on \mathcal{P} forms a *concave region* of \mathcal{P} .

Each face F on $\partial\mathcal{P}$ divides the three-space into two open half spaces separated by the carrying plane of F : the *positive half space* of F , or $H^+(F)$, is the open space on the side of F 's outward normal points, and the *negative half space* of F , or $H^-(F)$, is the half space on the other side.

Face F is *front-facing* to a point \mathbf{p} if \mathbf{p} lies in $H^+(F)$, and F is *back-facing* to \mathbf{p} if \mathbf{p} lies in $H^-(F)$. F is *tangent-facing* to \mathbf{p} if \mathbf{p} is on the carrying plane of F . If all the faces of a polyhedral surface S have the same facing type to point \mathbf{p} , S is called a *pure facing* surface to \mathbf{p} . Corresponding notations of *pure front-facing*, *pure back-facing* are defined for a surface S with respect to point \mathbf{p} if $S.faces$ are all front-facing or back-facing to \mathbf{p} .

The *viewing sphere* (or *unit sphere*), \mathbb{S}^2 , represents the manifold of all unit vectors. The spherical geometric terms *spherical point*, *spherical segment*, and *spherical polygon* are directly analogous to the point, line segment, and polygon in planar geometry. Note that a spherical segment is always a geodesic arc on \mathbb{S}^2 . A *spherical image* is a point set on \mathbb{S}^2 which may present itself as a 2-dimensional region (2-cell), 1-dimensional curve (1-cell), 0-dimensional point (0-cell), or their combinations. Two spherical points $\mathbf{u} \in \mathbb{S}^2$ and $\mathbf{v} \in \mathbb{S}^2$ are antipodal if $\mathbf{u} = -\mathbf{v}$, where \mathbf{u} is the antipode of \mathbf{v} and vice versa. Two spherical images are *opposite* to each other if and only if, for each spherical point in one image, there exists its antipode in the other image. A spherical image is *hemispherical* if it




Fig. 1. Dualities on the unit sphere.

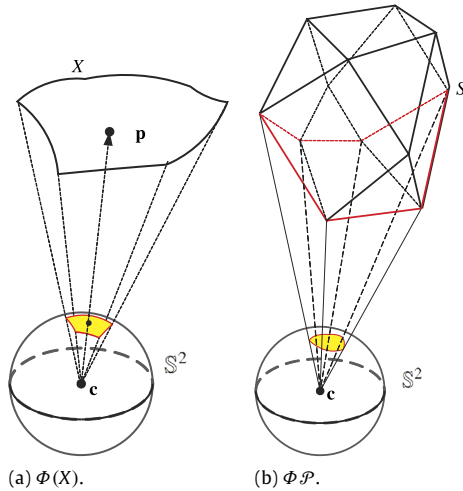


Fig. 2. Spherical projections. (a) The spherical projection of a general point set X . (b) When \mathbf{c} is exterior to a closed polyhedron \mathcal{P} , $\Phi(\mathcal{P})$ is equal to the spherical projections of all the back-facing faces on $\partial\mathcal{P}$, and the red curves in (b) show the border curves of those pure back-facing surfaces. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

is contained in any closed hemisphere. A spherical image is *convex* if it is hemispherical and, for any pair of points \mathbf{u} and \mathbf{v} in the image, the spherical segment $\overline{\mathbf{u}\mathbf{v}}$ is also in the image. The *spherical convex hull* (SCH) for a spherical image $X \in \mathbb{S}^2$ is the minimal convex spherical image containing X . Note that we have to consider special geometric properties of the spherical geometry. For example, there are special cases of spherical polygons, such as a polygon bounded by two semi great circles (lune), a polygon bounded by a full great circle (hemisphere), or an unbounded polygon (full sphere).

The *dual* image of a spherical point $\mathbf{v} \in \mathbb{S}^2$ is a closed hemisphere H , defined by $H = \{\mathbf{x} \mid (\mathbf{v} \cdot \mathbf{x}) \geq 0, \mathbf{x} \in \mathbb{S}^2, \mathbf{v} \in \mathbb{S}^2\}$ and vice versa (see Fig. 1(a) and (b)). The dual of a spherical convex polygon P with vertices $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$ is also a spherical convex polygon P^* . Let H_i be the dual hemisphere of \mathbf{v}_i , P^* is defined by $H_1 \cap H_2 \cap \dots \cap H_n$. Each vertex of P corresponds to an edge of P^* and vice versa (see Fig. 1(c)–(e)). More discussion on the details of dualities on the unit sphere can be found in [10].

The *spherical projection* of a point set $X \in \mathbb{R}^3$ is defined as $\Phi : X \rightarrow \mathbb{S}^2$ such that $\Phi(p) = \left\{ \frac{\mathbf{c}\mathbf{p}}{|\mathbf{c}\mathbf{p}|} \mid p \in X \right\}$, where \mathbf{c} is the center of \mathbb{S}^2 . In this work, \mathbf{c} is fixed at the origin and it is also used to denote the origin. The spherical projection of a closed point set X is equal to the spherical projection of its boundary, i.e. $\Phi(X) = \Phi(\partial(X))$. If a polyhedral surface S is *pure facing* to \mathbf{c} , the boundary of $\Phi(S)$ can be computed from the spherical projection of the border curves of S . For a polyhedron \mathcal{P} , $\Phi(\mathcal{P})$ is equal to the union of the spherical projections of all the pure back-facing surfaces on $\partial\mathcal{P}$ with respect to \mathbf{c} . See Fig. 2 for an illustration of spherical projections.

3. Local visibility vs. global visibility

Consider a surface S in a scene. When S is visible along direction \mathbf{v} , two conditions must be held. First, any line of sight along

direction \mathbf{v} intersects S no more than once. Second, any line of sight along \mathbf{v} intersects no other surfaces in the scene before it intersects with S . If the first condition does not hold, we say S is self-occluded in direction \mathbf{v} ; and if the second condition does not hold, we say S is globally occluded in direction \mathbf{v} . Using the concepts of self-occlusion and global occlusion, the terms, local visibility map and global visibility map, are defined for a surface. We let the direction that tails at a spherical point and heads into the center of the unit sphere be the corresponding view direction.

3.1. Local visibility map

Definition 1 (Local Visibility Map (LVM)). The local visibility map of a surface S , denoted by $LVM(S)$, defines the set of spherical points on \mathbb{S}^2 corresponding to all the view directions which are not self-occluded by S itself.

$LVM(S)$ can be calculated as the set of spherical points $\{\mathbf{v} \mid \mathbf{v} \cdot \mathbf{n}_p \geq 0, \mathbf{p} \in S, \mathbf{v} \in \mathbb{S}^2\}$, where \mathbf{n}_p is the outward normal of surface S at point \mathbf{p} [10].

For a single polygonal face F , its LVM is a closed hemisphere on \mathbb{S}^2 dual to the outward normal or F . The LVM of a polyhedral surface S is the intersection of all the hemispheres which are the LVMs of S 's faces [15]. That is,

$$LVM(S) = LVM(F_1) \cap LVM(F_2) \cap \dots \cap LVM(F_n) \\ = H_1 \cap H_2 \cap \dots \cap H_n, \tag{1}$$

where $\{F_1, F_2, \dots, F_n\} = S.faces$, and $\{H_1, H_2, \dots, H_n\}$ are the set of hemispheres dual to $\{\mathbf{n}_{F_1}, \mathbf{n}_{F_2}, \dots, \mathbf{n}_{F_n}\}$.

The calculation of LVM uses the notion of a Gaussian map. Given a surface S lying in \mathbb{R}^3 , the Gaussian map is a continuous map $\mathbf{GM} : S \rightarrow \mathbb{S}^2$ such that $\mathbf{GM}(p)$ is the normal vector to S at p . We adapt the original definition of Gaussian map and define a discrete Gaussian map for a polyhedral surface S or $DGM(S)$ by considering S as a bucket of facets and ignoring all the vertices and edges of S . Each facet in S is then mapped to a point on \mathbb{S}^2 , and $DGM(S) = \{\mathbf{n}_F \mid F \in S.faces\}$.

An important relation established in [10,15] between the discrete Gaussian map and the local visibility map for a polyhedral surface is as follows.

Proposition 1. For a polyhedral surface S , if the discrete Gaussian map $DGM(S)$ is not contained in a hemisphere, then $LVM(S)$ is empty. Otherwise, let $SCH(DGM(S))$ be the spherical convex hull (SCH) of the $DGM(S)$. $LVM(S)$ is equal to the dual image of the $SCH(DGM(S))$.

The procedure for computing $LVM(S)$ based on Proposition 1 is illustrated in Fig. 3. More discussion on the local visibility map and its properties can be found in [10,15].

3.2. Global visibility map

Definition 2 (Global Visibility Map (GVM)). The global visibility map of a surface S , denoted by $GVM(S)$, is the set of spherical points on \mathbb{S}^2 corresponding to all the view directions free of both self-occlusion and global occlusion.

In this work, if a ray emanating from S in direction \mathbf{v} touches another point on $\partial\mathcal{P}$ without intersecting $Int(\mathcal{P})$, $-\mathbf{v}$ is assumed




Fig. 3. Illustration of LVM calculation. (a) the surface S , (b) the $DGM(S)$. The $LVM(S)$ will be the intersection of all hemispheres dual to each point in $DGM(S)$. (c) the boundaries of all the hemispheres in this set. (d) the $SCH(DGM(S))$, and (e) the dual image of the spherical convex hull in (d). According to Proposition 1, the spherical polygon in (e) is the $LVM(S)$, which is equal to the intersection region in (c).

to be a global visible direction. In manufacturing applications, $-\mathbf{v}$ is considered as “just touching” and is generally treated as accessible direction. Since the boundaries of a GVM or a LVM correspond to the extreme directions, GVM and LVM are defined as closed spherical images.

Definition 3 (Global Occlusion Map (GOM)). The global occlusion map introduced by a surface O to another surface S , or $GOM(S, O)$, is the set of spherical points describing all the globally occluded view directions of S due to O .

Let the set of spherical points describing all the stabbing lines from S to O be \mathcal{Q} ,

$$\mathcal{Q} = \left\{ \frac{\mathbf{p}\mathbf{q}}{|\mathbf{p}\mathbf{q}|} \mid \mathbf{p} \in S, \mathbf{q} \in O \right\}. \quad (2)$$

For two disjointed surfaces S and O , $\partial\mathcal{Q}$ represents the set of “just touching” stabbing lines. The set of directions on $\partial\mathcal{Q}$ are not considered as globally occluded directions of S due to O . Thus, for two disjointed surfaces S and O , we have

$$GOM(S, O) = Int(\mathcal{Q}). \quad (3)$$

Since we only consider one polyhedron \mathcal{P} in our scene, for a sub region surface $S \in \partial\mathcal{P}$, only the surfaces in $(\partial\mathcal{P} - S)$ might introduce global occlusions to surface S . The set of points in $LVM(S)$ excludes all possible self-occlusions by S itself while $GOM(S, \partial\mathcal{P} - S)$ includes all possible global occlusions. Therefore, the difference between $LVM(S)$ and $GOM(S, \partial\mathcal{P} - S)$ represents the set of global visible directions of S ; i.e.,

$$GVM(S) = LVM(S) - GOM(S, \partial\mathcal{P} - S). \quad (4)$$

Practically, the global occlusion maps due to the faces in $(\partial\mathcal{P} - S)$ are possibly overlaid with each other. We wish to define the subset faces of $(\partial\mathcal{P} - S)$ which can minimize the overlay. Assume a ray emanating from a point on $\partial\mathcal{P}$ can only go into the positive half space defined by the oriented tangent plane of the point. (Otherwise, the ray is always blocked by $\partial\mathcal{P}$ itself.)

Property 1. The rays starting from a hull facet of \mathcal{P} can always go to infinity. The rays emanating from a non-hull facet F of \mathcal{P} will either go to infinity or be blocked first by another non-hull facet that “potentially faces” F [18,22].

Note that a pair of non-hull faces $F \in \partial\mathcal{P}$ and $F' \in \partial\mathcal{P}$ “potentially face” each other if all the following conditions are satisfied: (1) F and F' are presented in the same concave region of \mathcal{P} ; (2) F and F' are not coplanar; (3) F lies fully or partially in $H^+(F')$, and F' lies fully or partially in $H^+(F)$ [18].

Consider the set of parallel rays starting from the points in S along direction \mathbf{v} . If \mathbf{v} does not belong to $LVM(S)$, the direction \mathbf{v} will be self-occluded. If $\mathbf{v} \in LVM(S)$, for a specific face $F' \in (\partial\mathcal{P} - S)$, if $\mathbf{v} \cdot \mathbf{n}_{F'} > 0$, none of the rays in direction \mathbf{v} can reach F' from its front side. Thus, we have the following property.




Fig. 4. Occlusion maps between two neighboring faces which share a concave edge.

Property 2. For a face $F' \in (\partial\mathcal{P} - S)$, if the hemisphere dual to $\mathbf{n}_{F'}$ is a superset of the $LVM(S)$, then without self-occlusion, no set of parallel rays emanating from points of S can reach F' .

To check whether the hemisphere dual to $\mathbf{n}_{F'}$ is superset to $LVM(S)$, it is sufficient to check if $\mathbf{n}_{F'} \in SCH(DGM(S))$. See the proof in [15].

We have discussed the GOMs between two disjointed surfaces. See Eq. (3). Now we consider the set of faces adjacent to S . When a face F' in $(\partial\mathcal{P} - S)$ is adjacent to $F_S \in S$, depending on the convexity of the common edge, the global occlusions between F_S and F' have different properties.

Property 3. If a face $F' \in (\partial\mathcal{P} - S)$ shares a convex edge with $F_S \in S$, no set of parallel rays emanating from S can reach F' without self-occlusion.

This is obvious since the unblocked rays starting from F_S go only to the $H^+(F_S)$.

When two neighboring faces F_S and F' share a concave edge (Fig. 4), all the stabbing lines from F_S to F' have the direction \mathbf{v} satisfying $\{\mathbf{v} \cdot \mathbf{n}_{F_S} \geq 0, \mathbf{v} \cdot \mathbf{n}_{F'} \leq 0\}$. The set of stabbing lines prescribes a lunar region on S^2 representing the $GOM(F_S, F')$, see Fig. 4(b). Assume the common edge of F_S and F' be E . The stabbing lines from F_S to E form a semi-circle with a pole at \mathbf{n}_{F_S} , bounded by two spherical points, $(\mathbf{n}_{F_S} \times \mathbf{n}_{F'})$ and its antipode. This semi-circle corresponds to the $GOM(F_S, E)$ (Fig. 4(b)). An observation is that any point in the interior of this semi-circle is a globally occluded direction. The stabbing lines from E to F' will be another semi-circle with a pole at $-\mathbf{n}_{F'}$, bounded by the same two spherical points. The set of stabbing lines forming this semi-circle touches the face of F' without going into $Int(\mathcal{P})$ there in F' . Thus, the points on the later semi-circle are not globally occluded by F' (see the dashed circles in Fig. 4(b)). In summary, we have the following property.

Property 4. If two faces F_S and F' share a concave edge, $GOM(F_S, F')$ is equal to a lune dual to the spherical image defined by $\{\mathbf{n}_{F_S}, -\mathbf{n}_{F'}\}$. The semi-circle with the pole at \mathbf{n}_{F_S} is in $GOM(F_S, F')$, while the semi-circle with the pole at $-\mathbf{n}_{F'}$ is not in $GOM(F_S, F')$. The two bounding points of the lune are not in $GOM(F_S, F')$.




Fig. 5. 2D Illustration of Eq. (6) for a polyline S on a polygon P .

Based on Properties 1–4, we classify the faces in $(\partial\mathcal{P} - S)$ into different categories.

Definition 4 (Neighboring Occluder). The neighboring occluder \mathcal{N}_S for a surface S is the set of faces in $(\partial\mathcal{P} - S)$ which shares a concave edge with S .

Definition 5 (Potential Occluder). The potential occluder \mathcal{O}_S for a surface S is the set of faces in $(\partial\mathcal{P} - S)$ which “potentially face” S . A face F' in $(\partial\mathcal{P} - S)$ “potentially faces” S if the following conditions are satisfied: (1) F' is disjoint with S ; (2) F' “potentially faces” at least one face of S ; and (3) $\mathbf{n}_{F'} \notin \text{SCH}(\text{DGM}(S))$.

Besides the faces in \mathcal{N}_S and \mathcal{O}_S , all the other faces of $(\partial\mathcal{P} - S)$ can be ignored for global occlusion calculation because a ray emanating from a point in S either touches another point in S or touches a point in \mathcal{N}_S or \mathcal{O}_S before it can reach the points in $(\partial\mathcal{P} - S - \mathcal{N}_S - \mathcal{O}_S)$.

For each face $F' \in \mathcal{N}_S$, F' shares a concave edge with one face $F_S \in S$. Using Property 4, $\text{GOM}(F_S, F')$ is a lune (see Fig. 4). If one considers only the set view directions in $\text{LVM}(S)$, the global occlusions introduced to all the other faces in S , due to F , are always subset to the lune $\text{GOM}(F_S, F')$. Assume $\mathcal{N}_S = \{F'_1, \dots, F'_n\}$, and the hemispheres dual to $\mathbf{n}_{F'_1}, \dots, \mathbf{n}_{F'_n}$ are H_1, \dots, H_n , $\text{LVM}(S) - \text{GOM}(S, \mathcal{N}_S)$ is equal to $\text{LVM}(S) \cap H_1 \cdots \cap H_n$, which is meanwhile equal to $\text{LVM}(S + \mathcal{N}_S)$. Therefore, the global occlusion effect due to \mathcal{N}_S can be virtually merged to the self-occlusion for an augmented surface $(S + \mathcal{N}_S)$; i.e.,

$$\text{LVM}(S) - \text{GOM}(S, \mathcal{N}_S) = \text{LVM}(S + \mathcal{N}_S). \quad (5)$$

In summary, the GVM of a polyhedral surface S can be calculated by,

$$\text{GVM}(S) = \text{LVM}(S + \mathcal{N}_S) - \text{GOM}(S, \mathcal{O}_S). \quad (6)$$

It is obvious that the GVM of a surface containing only hull facets, or a surface forming a whole concave region, is equal to the LVM of the surface. In general, the GVM for a surface $S \in \partial\mathcal{P}$ is subset to $\text{LVM}(S)$. Fig. 5 gives a 2D example for illustrating the spherical maps defined in this section and the relation established in Eq. (6). For a polyline $S = \{l_8, l_9\}$ on a polygon P bounded by segment $\{l_1, \dots, l_{12}\}$, consider all the line segments for occlusion testing. l_1 to l_4 are not potentially facing S since they are on the convex hull of P . S is contained in the negative half space defined by l_5 . Thus, l_5 does not potentially face S . l_6 belongs to $\text{SCH}(\text{DGM}(S))$; it is not a potential occluder of S either. l_7 is convex connected with S and is not a neighboring occluder. Therefore, $\mathcal{O}_S = \{l_{11}, l_{12}\}$ and $\mathcal{N}_S = \{l_{10}\}$, as shown in Fig. 5(a). The discrete Gaussian map of S contains two points describing the line normals of S (Fig. 5(b)). To get the $\text{LVM}(S + \mathcal{N}_S)$, one needs to construct $\text{DGM}(S + \mathcal{N}_S)$ (Fig. 5(c)) and compute its dual image (Fig. 5(d)). $\text{LVM}(S + \mathcal{N}_S)$ describes the set of view directions free of self-occlusion among S and the global occlusion introduced by \mathcal{N}_S . $\text{GOM}(S, \mathcal{O}_S)$ represents the global occlusions introduced by \mathcal{O}_S , which is the spherical

image describing the stabbing lines from S to \mathcal{O}_S , as shown in Fig. 5(e). The final GVM(S) is calculated as the non-regularized Boolean difference between $\text{LVM}(S + \mathcal{N}_S)$ and $\text{GOM}(S, \mathcal{O}_S)$. The final result is shown in Fig. 5(f). Note that the two bounding points in GVM(S) are global visible directions.

The Boolean difference and union operators in Eqs. (4)–(6) are non-regularized Boolean operators. A regularized Boolean operation between two bodies of the same dimension deletes the results with lower dimensions. However, the non-regularized Boolean operation maintains those degenerated outputs. Therefore, the final GVM for a surfaces, calculated through Eq. (6) using non-regularized Boolean operators, may contain lower dimensional elements. The GVM for a polyhedral surface may contain 0-cell(s), 1-cell(s), 2-cell(s), or a combination. Low-dimensional GVMs play an important role in many CAM applications since for the common features such as slots, steps, bosses, etc., 0-cell and 1-cell GVMs occur frequently and they represent *tight passages* of the tool’s accessibility [18].

4. Calculating global occlusion maps using Minkowski sums

The Minkowski sum of two sets, $S_1 \in \mathbb{R}^3$ and $S_2 \in \mathbb{R}^3$, denoted $S_1 \oplus S_2$, is defined as: $S_1 \oplus S_2 = \{\mathbf{p} + \mathbf{q} \mid \mathbf{p} \in S_1, \mathbf{q} \in S_2\}$. We now show that the global occlusion map of a surface S due to another surface O can be derived from the spherical projection of the Minkowski sum between $-S$ and O , where $-S = \{-\mathbf{p} \mid \mathbf{p} \in S\}$.

Proposition 2. For two disjoint surfaces S and O ,

$$\text{GOM}(S, O) = \text{Int}(\Phi(O \oplus -S)).$$

Proof. According to Definition 3, $\text{GOM}(S, O)$ is the open spherical image describing the vectors in the set $\left\{ \frac{\mathbf{pq}}{|\mathbf{pq}|} \mid \mathbf{p} \in S, \mathbf{q} \in O \right\}$. The spherical point describing a normalized vector $\frac{\mathbf{pq}}{|\mathbf{pq}|}$ is equal to the spherical projection of the point $(\mathbf{q} - \mathbf{p})$. Therefore, $\text{GOM}(S, O)$ is equal to the spherical projections for all the points in $\{\mathbf{q} - \mathbf{p} \mid \mathbf{p} \in S, \mathbf{q} \in O\}$, which is exactly equal to $O \oplus -S$. Since GOM is an open spherical image for two disjoint surfaces, we have $\text{GOM}(S, O) = \text{Int}(\Phi(O \oplus -S))$. We illustrate the proof in Fig. 6 with a 2D example. \square

The Minkowski sum of two polyhedral surfaces has the following property.

Property 5. Let O and S be two polyhedral surfaces, and $M = \partial(O \oplus S)$. Any face of M can be written as $V_0 \oplus F_S, F_0 \oplus V_S$ or $E_0 \oplus E_S$ [28].

We define all faces that can be written as $V_0 \oplus F_S, F_0 \oplus V_S$ or $E_0 \oplus E_S$ as *tentative faces*. The pair of V_0 and F_S, F_0 and V_S , or E_0 and E_S is the *parent* of the corresponding tentative face. A tentative face is said to be of type *VF, FV, or EE*. Fig. 7 shows an example for the different types of tentative faces on the Minkowski sum boundary $\partial(O \oplus S)$.




Fig. 6. 2D Illustration for the proof of Proposition 2. Note that the blue curve in (b) showing the GOM(S, O) is equal to the blue curve in (e), which is the $\Phi(O \oplus -S)$.

Fig. 7. Different types of tentative faces on $\partial(O \oplus S)$.

The set of tentative faces defines a superset of $\partial(O \oplus S)$; it is, meanwhile, a subset of $O \oplus S$. For a closed set, the spherical projection of the set is equal to the spherical projection of its boundary (see Section 2). Therefore, $\Phi(O \oplus S)$ is equal to $\Phi(\partial(O \oplus S))$; and it further equals the spherical projection of all the tentative faces. The set of tentative faces has a size complexity of $O(mn)$, given that the two surfaces O and S have sizes m and n respectively. But in fact, many tentative faces do not really show up on $\partial(O \oplus S)$. We wish to filter out the tentative faces which are not on $\partial(O \oplus S)$ thus, to save computation.

Each edge E in surface O or S associates with a *tangent vector*, denoted \mathbf{t}_E , which is the unit vector parallel to E . Given a face F in O or S , and an edge E incident to F , the *inface normal* $\mathbf{n}_{E,F}$ is defined as the unit vector with direction $\mathbf{n}_F \times \mathbf{t}_E$, oriented such that it points towards the interior of F .

A *VF* or an *FV* type of tentative face can be computed by translating the face parent along the vector corresponding to the vertex parent, and the normal direction of the tentative face is parallel to the normals of the face parent. For an *EE* type tentative face, it is a parallelogram with the normal parallel to $(\mathbf{t}_{E_0} \times \mathbf{t}_{E_S})$, where \mathbf{t}_{E_0} and \mathbf{t}_{E_S} are the tangent vectors of the two edge parents. Now we define two terms, the *tentative normal* \mathbf{n}_{ten} and the *critical set* C_{cri} , for a tentative face as follows.

For a *VF* or *FV* type of tentative face with the parent pair $\{V, F\}$, ($V \in S.vertices, F \in O.faces$, or $V \in O.vertices, F \in S.faces$), \mathbf{n}_{ten} is the normal of the face parent, i.e. $\mathbf{n}_{ten} = \mathbf{n}_F$. Let the set of tangent vectors for all the edges incident to V be $T = \{\mathbf{t}_{E_i}\}$ (each $\mathbf{t}_{E_i} \in T$ is defined as pointing away from V). Define the set of dot products in $\mathbf{n}_{ten} \cdot T$ as the *critical set* C_{cri} of the corresponding *VF* or *FV* type of tentative face.

For an *EE* type of tentative face with the parent pair $\{E_0, E_S\}$, define $\mathbf{n}_{ten} = \mathbf{t}_{E_0} \times \mathbf{t}_{E_S}$. Let N be the set of inface normals of the edges E_0 and E_S in their neighboring faces. The set of products in $\mathbf{n}_{ten} \cdot N$ is defined as the critical set C_{cri} for an *EE* type of tentative face.

There are four types of tentative faces classified by their critical sets C_{cri} .

- * $\{0, 0\}$: a tentative face with all zeros in its C_{cri} .

- * $\{+, -\}$: a tentative face with both negative and positive numbers in its C_{cri} .
- * $\{+, 0\}$: a tentative face with some positive numbers but without negative numbers in its C_{cri} .
- * $\{-, 0\}$: a tentative face with some negative numbers but without positive numbers in its C_{cri} .

Lemma 1. A tentative face of $O \oplus S$ is not on $\partial(O \oplus S)$ if its C_{cri} has the type $\{+, -\}$.

Lemma 2. If a tentative face of $O \oplus S$ is on $\partial(O \oplus S)$, and its C_{cri} has the type $\{-, 0\}$, the outward normal of the tentative face is \mathbf{n}_{ten} ; otherwise, if C_{cri} has the type $\{+, 0\}$, then the outward normal of the tentative face is $-\mathbf{n}_{ten}$.

The proofs for the above two lemmas are given in the [Appendix](#).

Consider the case when the C_{cri} of a tentative face are all zeros. This happens only when the two parents of the tentative face are in non-general positions. In a non-general position, the normal for a tentative face is unstable since the tentative face might belong to a degenerated region of $O \oplus S$. In this work, we consider all non-general positions and record the normal for tentative faces with type $\{0, 0\}$ as parallel to \mathbf{n}_{ten} and pointing to the direction such that the tentative face will be back-facing to the origin \mathbf{c} .

Recall that, for a closed polyhedron, its spherical projection is equal to the spherical projections of its pure back-facing surfaces. In a general position, $\partial(O \oplus S)$ is a closed polyhedron, while in a non-general position, $\partial(O \oplus S)$ may contain open surfaces corresponding to degenerated Minkowski sums. As mentioned, we set all the tentative faces resulting from non-general positions as back-facing to \mathbf{c} . Therefore, to obtain $\Phi(O \oplus S)$, it is sufficient to compute the spherical projection for the set of back-facing tentative faces to \mathbf{c} .

Definition 6 (Effective Minkowski Envelope (EME)). Given two disjointed surfaces S and O , the effective Minkowski Envelope for $O \oplus S$, denoted as $EME(O \oplus S)$, is defined as the set of backing-facing tentative faces of $O \oplus S$ (with respect to origin \mathbf{c}), whose critical sets are not in the type $\{+, -\}$.

