The 3D visibility complex : a new approach to the problems of accurate visibility.

Frédo Durand, George Drettakis and Claude Puech

iMAGIS *

Abstract: Visibility computations are central in any computer graphics application. The most common way to reduce this expense is the use of approximate approaches using spatial subdivision. More recently analytic approaches efficiently encoding visibility have appeared for 2D (the visibility complex) and for certain limited cases in 3D (aspect graph, discontinuity meshes). In this paper we propose a new way of describing and studying the visibility of 3D space by a dual space of the 3D lines, such that all the visibility events are described. A new data-structure is defined, called the *3D visibility complex*, which encapsulates all visibility events. This structure is global and complete since it encodes all visibility relations in 3D, and is spatially coherent allowing efficient visibility queries such as view extraction, aspect graph, discontinuity mesh, or form factor computation. A construction algorithm and suitable data structures are sketched.

Keywords: visibility, visibility complex, spatial coherence, discontinuity meshing, form factor

1 Introduction

Visibility calculations are central to any computer graphics application. To date, no approach has been presented to encode all visibility information in a 3D scene.

In this paper we will present a new approach, which we call the *3D visibility complex*, which encodes all visibility information contained in a three dimensional scene. This research is in a preliminary phase, since an implementation has not yet been undertaken, but we believe that the importance and potential use of such a structure justify its presentation even at the stage of conception.

Related works The first attempts to cope with the cost of visibility computations involved space partitioning structures but they provided only local visibility information. Arvo and Kirk [1] subdivide the 5D ray-space for ray-tracing. Teller [13] uses the 5D Plücker duality to compute the antipenumbra cast by an area light source. He also developed algorithms for scenes naturally divided into cells [15] where the visibility is propagated through portals. In computer vision the *aspect graph* [7, 6] has been developed to group all the viewpoints for which an object has the same "aspect". An aspect changes along visibility events which are the same as for the discontinuity meshing techniques [8]. These techniques have thus been extended with *backprojections* [3, 12] to provide the aspect of the source. Recently, efficient data structures have been developed for the 2D case [10, 5] and have inspired our research, although the new approach has been developed from scratch with the specifically three-dimensional problem in mind.

* Laboratoire GRAVIR / IMAG. iMAGIS is a joint research project of CNRS/INRIA/INPG/UJF. Postal address: B.P. 53, F-38041 Grenoble Cedex 9, France. Contact E-mail: `Frederic.Durand@imag.fr`.

2 Description of the 3D Visibility Complex

In this discussion we will consider scenes of general convex objects, but the concepts will also be given for the polygonal scenes where appropriate. Visibility will be defined in terms of ray-objects intersections. If we consider the objects to be transparent, a ray is not blocked and all the objects a line intersects must be considered. If however we want to take occlusions into account, we will consider maximal free segments which are segments having no intersection with the inside of the objects and whose length is maximal (their two extremities lie on the boundary of two objects or are at infinity). In what follows we will often refer to them simply as *segments*. Segments can be interpreted as rays which can *see* the two objects on their extremities. A 3D line can be collinear to many segments, separated by the objects the line intersects. In this paper, we will introduce concepts first in terms of line visibility (where all the objects intersected by a line are considered) and then in terms of segment visibility (where the occlusions are taken into account).

We wish to group the segments (or the lines) which see the same objects. A partition of the set of segments into connected components according to their visibility is thus required. Since sets of segments are not intuitive objects, we will try to represent them in a dual space which will afford a better understanding of intricate visibility relationships. A suitable duality will thus be used for the purposes of illustration and presentation.

2.1 Duality

We have chosen to decompose the 4 dimensions of line space into two dimension of direction (the spherical coordinates (θ, φ) of the director vector of the lines) and a projection (u, v) onto the plane perpendicular to the line and going through the origin. The axes of the planes are chosen such as u is along $\boldsymbol{t} \wedge \boldsymbol{y}$ [2]. The intersections of a line with two parallel planes could also be used. Nonetheless, we believe that such an approach makes the interpretation of lines sharing one coordinate harder.

Visualizing 4D space is very hard. It can be seen as a moving 3D world with the 4th dimension being time. One approach is to use slices (in this paper we will fix $\varphi = ct$) which can be seen as frames in time. Such a slice will be called a φ-slice. Since each slice will be a 3D space (θ, u, v), it will sometimes be useful to cut one more time and consider φ and θ constant. We will obtain a 2D slice where only u and v vary, composed of all the lines which are parallel and have the direction (θ, φ). Such a slice will be called a $\theta\varphi$-slice. These 2D $\theta\varphi$-slices are easier to handle and visualize. They justify in part the choice of the duality because they can be interpreted as orthographic projections of the scene.

2.2 Tangency curves

Line Visibility Visibility changes when a line becomes tangent to an object. The set of lines tangent to one object is a 3-D set in the 4D dual space. This means, more intuitively, that a line has 3 degrees of freedom to stay tangent to one object. We will call the dual of the set of lines tangent to an object the *tangency volume* of this object.

Figure 1b shows a representation of the tangency volume of a sphere. For each φ-slice, the set of tangents is a sort of 2D "cylinder", forming a 3D structure in the 4D dual space. If we consider a 2D $\theta\varphi$-slice (horizontal in figure 1b) the set of tangents sharing

[2] Discontinuities occur at $\varphi = \pm\frac{\pi}{2}$, but since we use this duality for the purpose of presentation and visualization we can ignore them without loss of generality.

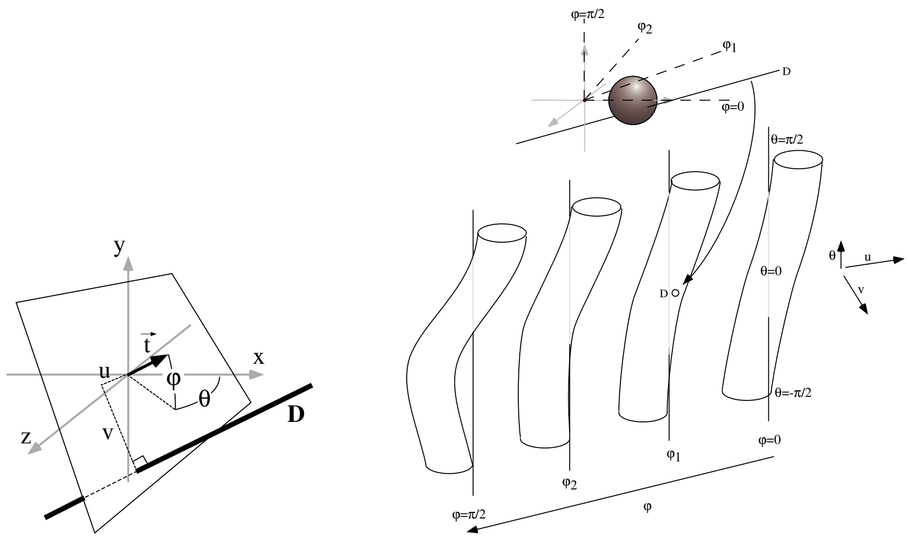

(a) (b)

Fig. 1. (a) Duality (b) Tangency Volume of a sphere. The θ axis ($u = 0$, $v = 0$) is shown for each φ-slice providing a better 3D visualization. In the left-hand φ-slice, which corresponds to the discontinuity in the duality for $\varphi = \frac{\pi}{2}$, the “cylinder” just turns around the θ axis. The line D intersects the object and has its dual inside the tangency volume.

that direction is a circle in the dual space. This is general: because of the definition of u and v, the set of tangents to one object in one direction is the outline of the object in this direction.

If a line has its dual on the tangency volume, it is tangent to the object. If the dual is inside the 4D set bounded by the tangency volume, it intersects the object, similarly to line D on figure 1b.

Segment Visibility Let us now consider visibility with occlusion. A line which intersects the object is collinear to at least two segments, one before and one after the object.

Consider a $\theta\varphi$-slice such as that on the lower left of figure 2. The sets of lines that intersect and that do not intersect the object are bounded by the outline of the object. For segment visibility we have to consider the segments that see the front of the object and those that see its back. Since such segments are collinear to the same line, they are projected on the same point in the 4D line dual space. Consequently the set of segments that see the front and the set of segments that see the back of the object are projected onto the same position of the 4D dual space as shown in the right of figure 2. The outline, which is the set of tangents to the object for the chosen θ and φ, is incident to the three sets (front, back and no intersection). This means that a segment tangent to the object has topological neighbours that do not intersect the objects, some that see the front, and some that see the back.

To differentiate the segments, we add a pseudo-dimension. It is not a continuous dimension since we just have to sort all the collinear segments. If we impose $\theta = ct$,

$\varphi = ct$ and $v = ct$, the sets of segments can be represented by a graph [3] shown on the lower right. Each tangent corresponds to a vertex of the graph. This graph is a 1D structure embedded in 2D. Similarly, for a $\theta\varphi$-slice, the sets of segments are represented by a 2D structure embedded into 3D. We call the partition of the segments of direction (θ, φ) according to their visibility the *auxiliary complex* for (θ, φ) (see also figure 4).

In a similar manner, a φ-slice is in fact a 3D structure embedded into 4D, and the sets of segments is a 4D space embedded into 5D.

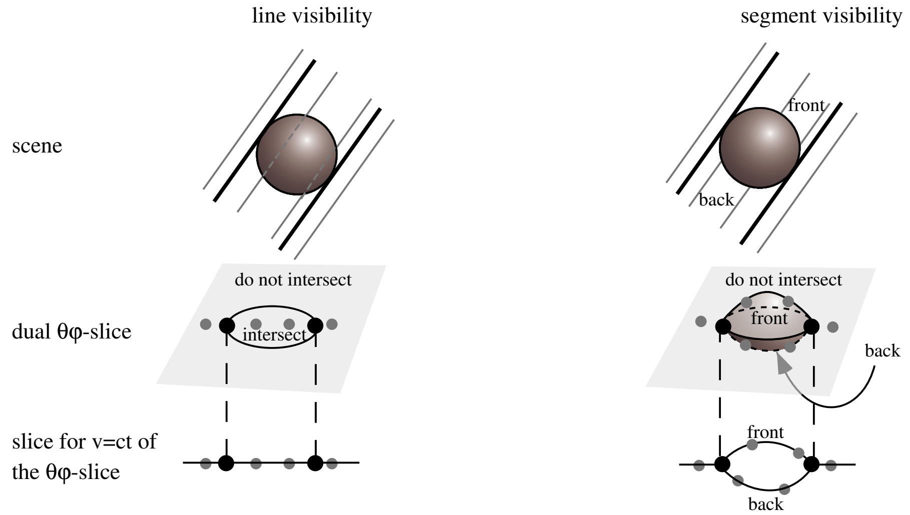

Fig. 2. Visibility for $\theta = ct$ and $\varphi = ct$. If we consider lines (on the left), visibility can be described by a planar structure (below). But if we consider segments (on the right) we have different levels on this plane depending on the side of the object. The set of segments which do not intersect and the sets of those that intersect the front or the back of the object share the same boundary, the tangents to the object which correspond to its outline. Recall that the Auxiliary Complex shown on the lower right is a 2D structure embedded into 3D, i.e. it is "empty ", since the points outside the surfaces have no meaning.

2.3 Bitangents

Line Visibility Now consider two objects. If a line has its associated dual point inside the tangency volumes of both objects, it intersects them both. The tangency volumes give us a partition of the dual space of the 3D lines according to the objects they intersect. We call this partition the *dual arrangement*. Its faces are 4D sets of lines which intersect the same objects. They are bounded by portions of the tangency volumes which are 3D. The intersection of two tangency volumes is a 2D set corresponding to the lines tangent to the two objects (bitangents).

For a φ-slice the set of bitangents is a space curve (shown as dashed line in figure 3 on the two φ-slices on the right). It corresponds to the intersection of the two "cylinders" which are the φ-slices of the tangency volumes. The slice of a 4D face is a volume corresponding to the intersection of the inside of the two cylinders.

[3] It is in fact an embedding of a graph since the points on the edges also have a meaning

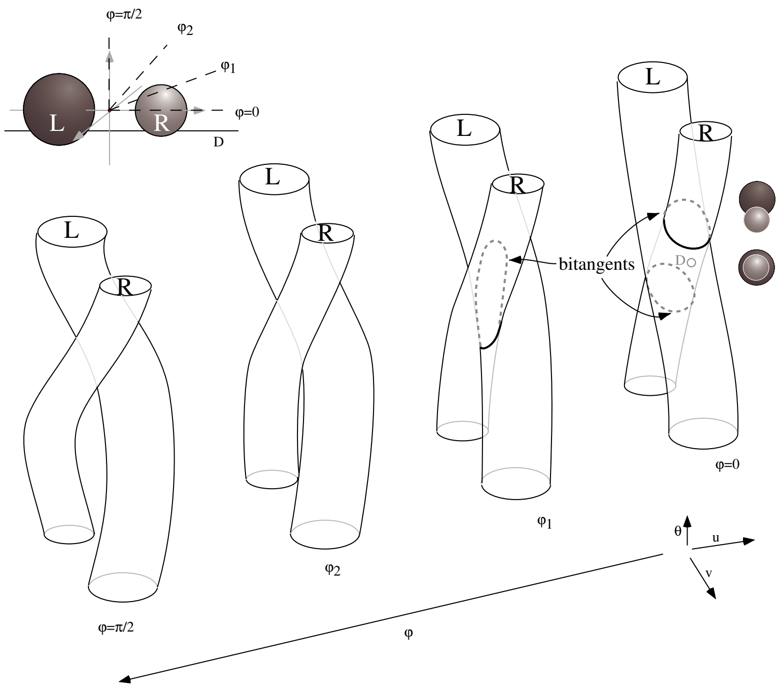

Fig. 3. Dual arrangement for two spheres.

Segment visibility An auxiliary complex for two objects is shown on figure 4 for a given direction. It is still delimited by the outline of the objects, but for example the outline of the upper sphere has no influence on the set B of segments that see the back of the lower sphere. Note that the two bitangents (shown in fat black lines) are incident to all faces.

Figure 5 is a φ-slice for $\varphi = 0$ of all the faces of the scene composed of two spheres of figure 3. The view in a given direction is shown on the left of the cylinders, and we consider the associated auxiliary complex shown six times on the top of the schema. Each time, a face is hatched and a volume is drawn below which corresponds to the φ-slice of the face of the visibility complex at $\varphi = 0$. Note that the union of these volumes is more than the entire 3D space, since a φ-slice of the complex is a 3D structure embedded into 4D.

2.4 Tritangents

Consider now a scene of three objects. A line tangent to the three objects has its dual at the intersection of the three tangency volumes. A set of connected tritangents is a 1D set in the 4D dual space. Its projection on a φ -slice is a point. The set of tritangents can be also interpreted as the intersection of the three sets of bitangents.

Figure 6 shows part of the visibility complex of a scene of three spheres. On the φ-slice $\varphi = 0$ two orthographic views of the scene for $\theta = 0$ (View 0) and for $\theta = \theta_2$ (View 2) are drawn next to the corresponding θ in the φ-slice. The set F of segments that see the spheres R and B is shown by its two slices F_0 and $F_{\varphi 1}$. Note that it is the intersection of the tangency volume of R and B minus the tangency volume of G. The

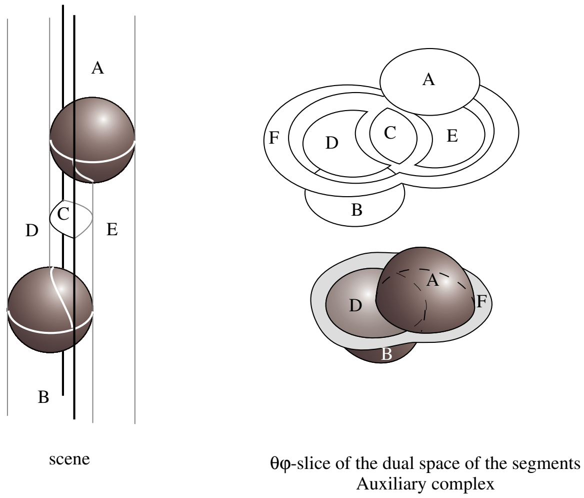

Fig. 4. Auxiliary Complex for two spheres. Recall that the auxiliary complex is a 2D structure embedded in 3D. In the lower representation, only the points on the surfaces represented are associated with segments. In the upper view, the faces of the auxiliary complex have been moved out to make their incidences easier to understand.

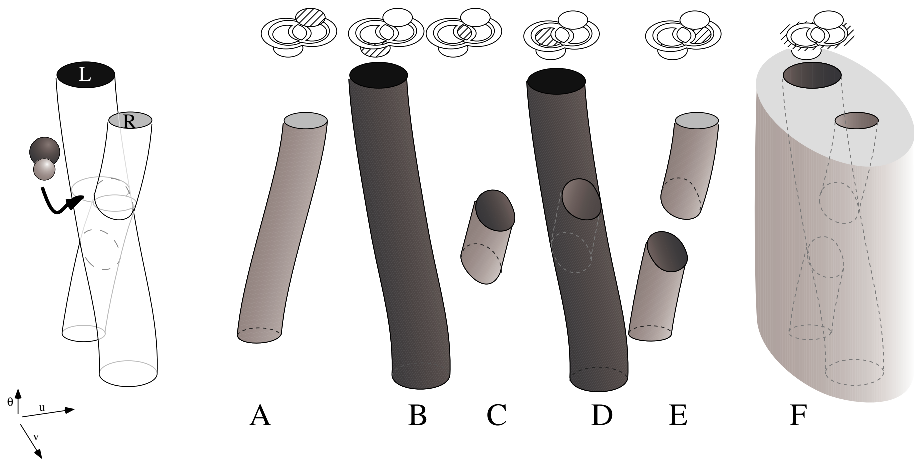

Fig. 5. φ-slice for $\varphi = 0$ of the faces of the visibility complex of the previous scene. A is the set of segments that see the front of R, B is the set of segments that see the back of L. C is the set of segments between L and R. It can be interpreted as the intersection of set of lines that see L and the set of lines that see R, and in the dual space it has the shape of $A \cap B$. D is the set of segments that see the front of L. Since the visibility is occluded by R in this direction, D has the shape of $B - A$. Similarly, E is the set of segments that see the back of R. Finally, F is the set of segments that see none of the two spheres. It is the complement of $A \cup B$.

tritangents are the points in white. Note also that because of the occlusion by the sphere G, lines that are bitangents of the R and B do not correspond to bitangent segments. This is shown in figure 7 which is a zoomed view of the φ-slice $\varphi = 0$. The set of bitangents B_0 is cut because bitangent lines such as D intersect G and correspond to no bitangent segment. We can thus see that the tritangent T_0 and T_0' are the intersection of the φ-slices of the three tangency volumes, and are also incident to the three sets of bitangents B_0, B_0' and B_0''.

Note that a scene does not necessarily contain tritangents in the general case.

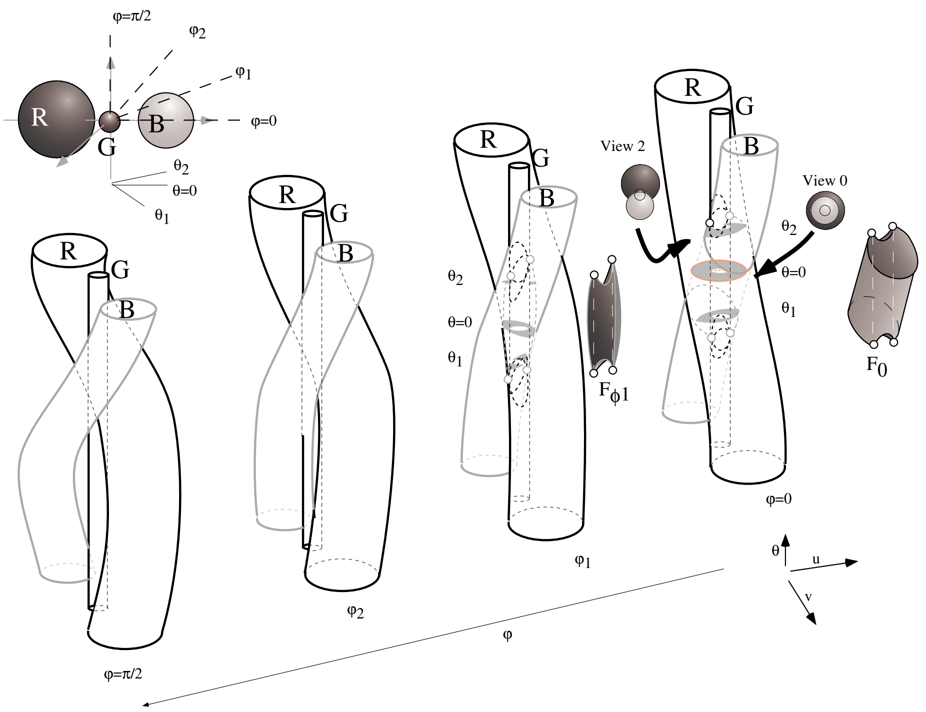

Fig. 6. Visibility Complex of a scene of three spheres.

3 Data Structure and Storage Complexity

3.1 Overview of the Data Structure

We have defined the dual arrangement which is the partition of the lines of the 3D space into connected components according to the objects they intersect. It is a 4D structure.

Similarly, the 3D visibility complex is the partition of the maximal free segments of 3D space into connected components according to the objects they touch. It is a 4D structure embedded into 5D. The dimensions and incidences of the boundaries of the faces are summarised in table 3.1.

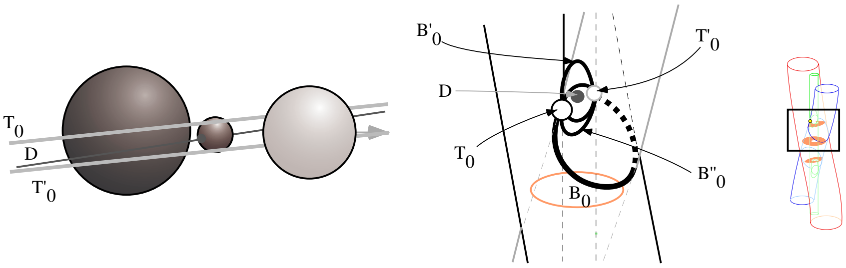

Fig. 7. Zoomed view of the φ-slice $\varphi = 0$.

Note that the elements of the visibility complex and those of the dual arrangement are not the same. A line can be tangent to two objects and correspond to no bitangent segment because of occlusions.

In the general case, a scene can have a degenerate visibility complex with no vertex and no tritangency edge.

Dim	Scene configuration	φ-slice in the dual space	Name
4	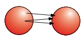	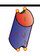	face
3	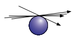	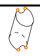	tangency face
2	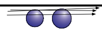		bitangency face
1	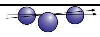		tritangency edge
0	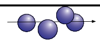		vertex

Table 1. Elements of the visibility complex

3.2 Polygonal case

In the case of polygonal scenes, the outlines of the objects can be decomposed into edges and vertices. Consequently the tangency volumes of a polygon can be divided into sets of lines going through the edges which are 3D sets, and sets of lines going through the vertices which are 2D sets. A 2D component of the complex corresponds to a segment touching two edges, or to a segment touching one vertex of a polygon. In the same manner, the 1-faces of the complex correspond to segments going through three edges (the EEE events of the aspect graphs or of the discontinuity mesh) or to segments going through an edge and a vertex (the EV events). Vertices of the complex can be $EEEE$ or EEV or VV events. In particular a line (or a segment) going through the vertex of a polygon can be interpreted as being tangent to the two edges incident to this vertex.

In the polygonal case the visibility complex is always non-degenerate since there are always VV vertices and EV 1-faces.

3.3 Complexity

In the general case, there exist convex objects for which the number of faces of the complex is unbounded. However, in the polygonal case, the storage complexity of the visibility complex is $O(n^4)$, where n is the number of edges of polygons. This complexity depends strongly on the configuration of the scene. We show below that the proposed construction algorithm is $O(n \log n$.

As mentioned in the introduction, practical experience with discontinuity meshing has shown that the scenes studied in computer graphics tend to have more optimistic visibility complexity than that predicted by the theoretical worst case [3].

4 Applications of the approach

4.1 View computation

A view around a point is defined by the extremities of the set of segments going through this point. The set of segments going through a point is a 2D surface in the dual space (u and v can be expressed with $\sin(\theta)$ and $\sin(\varphi)$). The view can be expressed as the intersection of the visibility complex with this surface. Each face intersected corresponds to an object seen. An intersection with a tangency volume corresponds to an outline in the image. The ray-tracing algorithm is equivalent to a sampling of such a surface.

In figure 8, the surface described by the lines going through viewpoint V is represented by its φ-slices which are curves. The intersections of these curves with the tangency volumes are the points of the view on the outline of the objects, such as D_1, D_2, D_3, D_4 and D_5. However,all the intersections do not necessarily correspond with an outline since the objects are not transparent, and points such as D' must not be taken into account. Consider the φ-slice $\varphi = 0$ and the slice V_0 of the lines going through V with $\varphi = 0$. Figure 9 shows the φ-slices of the faces of the visibility complex and their traversal. We traverse the visibility complex up and down along V_0. Initially, the segments see nothing, since we are in the face F. At D_1, we leave face F and have to chose between face A and E. Since V lies in the front of the sphere R, we now traverse A from D_1 to D_2. D' lies on no boundary of face A and is thus not considered. We then traverse face D and finally face F again. Once the φ-slice has been traversed, the intersections with the boundaries of the faces are maintained while φ is swept. Visibility changes will appear when V_φ meets a bitangency edge or a new tangency volume.

For a walkthrough, the view can be maintained since the events where the visibility changes correspond to intersections of the surface described by V with the 1-faces of the visibility complex. This approach is similar to the one described in [2] where conservative visibility events are lazily computed.

4.2 Form-Factors

The form factor F_{ij} [4] involved in radiosity computation is the proportion of light that leaves patch i which arrives at patch j. It can be expressed as the measure of lines which intersect i and j divided by the measure of lines which intersect i. In the dual space, it is the measure of the face F_{ij} divided by the measure of the inside of the tangency volume of i. See [9] [4] for the equivalent interpretation of the form factors with the 2D visibility complex.

[4] The same notation is used for the form factor and for the face between i and j though the form factor is a scalar and the face is a set of segments.

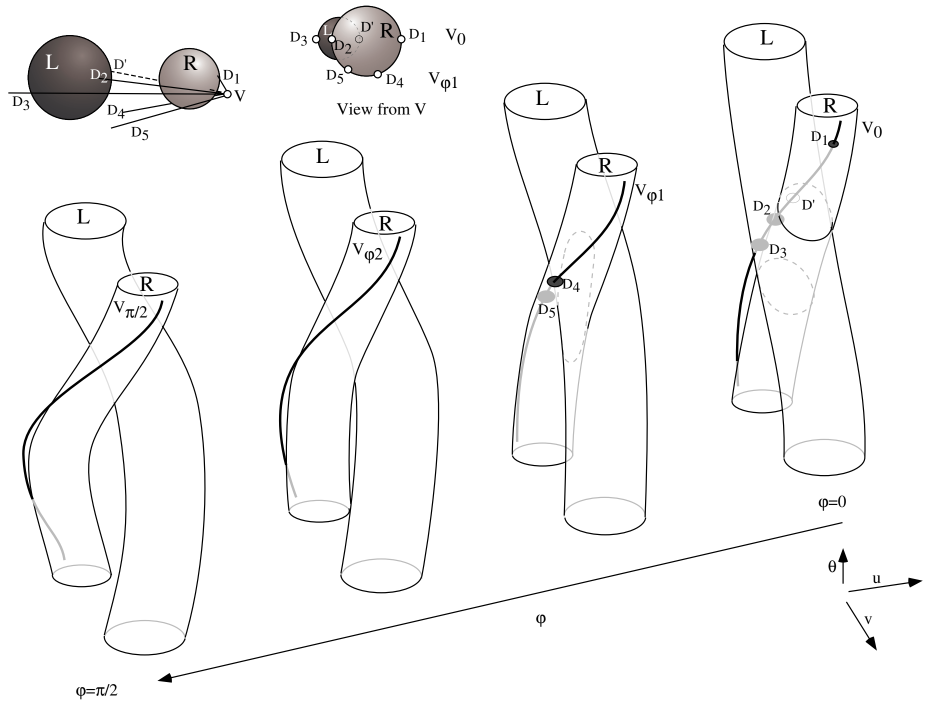

Fig. 8. The View around a point is the intersection of the visibility complex and the surface described by the set of segments going through this point.

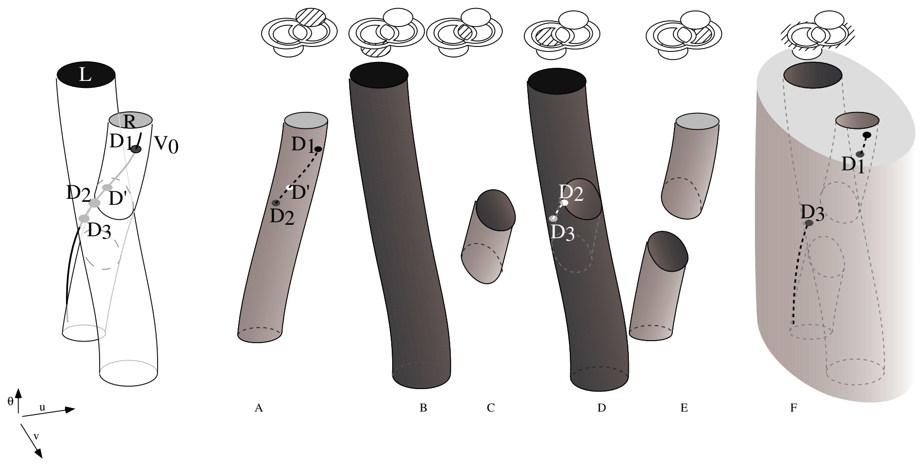

Fig. 9. Traversal of the φ-slice $\varphi = 0$ of the complex to compute the view around point V.

4.3 Other applications

The 1-faces of the visibility complex correspond to the visibility events of the aspect graph. The complex can thus help in its construction. The complexity of the aspect graph is $O(n^6)$ though the visibility complex is "only"$O(n^4)$ because the aspect graph is an arrangement of the $O(n^3)$ 1-faces of the complex.

In the same way, the 1-faces inside the tangency volume of the scene correspond to the discontinuity surfaces of the discontinuity meshing methods. The visibility complex gives all the events to compute a discontinuity mesh where all the objects are considered as sources.

In the context of hierarchical radiosity, whenever a link between two objects i and j is to be refined the boundary of the face F_{ij} of the visibility complex provides all the visibility information pertinent to this energy exchange. This information can be used to effect progressive discontinuity meshing and to improve the quality of the form-factor calculation.

5 Implementation

We give here a general outline for the implementation of this data structure for scenes of polygons. The development of the actual implementation will present technical difficulties which we have not yet addressed. We simply sketch an outline of the form the data structure will have and give a general idea of how the construction will proceed.

5.1 Data Structure

To represent the 3D visibility complex, we can use a polytope structure. Each k-face has pointers to its boundaries (faces of a lower dimension) and to the faces of a larger dimension it is adjacent to. A tangency face has for example a list of the bitangency face of its boundary, and three pointers to its adjacent faces. For each k-face we also store the two objects it can see and the objects to which its segments are tangent.

5.2 Algorithm

We present here the outline of an algorithm to build the visibility complex. It consists of a direct enumeration of the vertices of the complex inspired by [6], and then a sweep of these vertices.

All the potential $O(n^3)$ EV and EEE events are first enumerated, and we then compute the intersection of the corresponding discontinuity surfaces with the n objects of the scene. This gives us all the vertices of the visibility complex which are then sorted in φ and stored in a priority queue.

We then maintain a φ-slice of the complex during the sweep of the vertices. For each vertex swept we link all the k-faces incident to this vertex.

The algorithm presented is $O(n^4 \log n)$, but experience in the field of discontinuity meshing and backprojections has shown that the cost can be much reduced thanks to accelerations techniques [3]; the number of EEE actually considered is usually far less than n^3.

6 Conclusions

We have presented a new approach for visibility computation and described a powerful data-structure which encapsulates all the visibility information in a 3D scene. The dual space used affords a better understanding of the visibility events, which have been presented in detail. Moreover, this representation gives all the relations of adjacency between these events.

The 3D visibility complex is a very promising data structure for numerous computer graphics applications: we have briefly outlined its potential use for the visibility computation of a view, its use in form-factor computations and discontinuity meshing as well as the computation of aspects or backprojections.

We have presented a first outline of the data structure and a construction algorithm. Current work focuses on the completion of the algorithm and the data structure and its subsequent implementation for polygonal scenes.

It is nonetheless evident that the when applied to large scenes, the 3D visibility complex will suffer from combinatorial growth in storage. To cope with this combinatorial complexity, two strategies will be explored. Lazy construction can allow the computation of only the most important visibility events and faces of the visibility complex when they are actually needed by the application. A hierarchical extension of the 3D visibility complex will be studied.

Finally the visibility complex, like its 2D equivalent, seems very promising for dynamic environments due to its inherently coherent construction.

References

1. Arvo J, Kirk D. Fast ray-tracing by ray classification. SIGGRAPH 87
2. Coorg S, Teller S. Temporally coherent conservative visibility. Proc. of the 12th Ann. ACM Symp. on Computational Geometry, May 1996.
3. Drettakis G, Fiume E. A fast shadow algorithm for area light sources using backprojections. SIGGRAPH 1994.
4. Durand F, Orti R, Rivière S, Puechc C. Radiosity in flatland made visibly simple. video of the 12th Ann. ACM Symp. on Computational Geometry, May 1996.
5. Durand F, Puech C. The Visibility Complex made visibly simple, video of the 11th Ann. ACM Symp. on Computational Geometry, June 1995.
6. Gigus Z, Canny J, Seidel R. Efficiently computing and representing aspect graphs of polyhedral objects. IEEE Trans. on Pattern Analysis and Machine Intelligence, June 1991.
7. Gigus Z, Malik J. Computing the Aspect graph for line drawings of polyhedral objects. IEEE Trans on Pattern Analysis and Machine Intelligence, vol. 12, February 1990.
8. Lischinski D, Tampieri F, Greenberg D. Combining hierarchical radiosity and discontinuity meshing. SIGGRAPH 1993.
9. Orti R, Rivière S, Durand F, Puech C. Radiosity for dynamic scenes in flatland with the visibility complex. Eurographics 1996.
10. Pocchiola M, Vegter G. The visibility Complex, To appear in the special issue of Internat. J. Comput. Geom. Appl. devoted to the 9th Annu. ACM Sympos. on Comput. Geometry, held in San Diego, June 1993.
11. Sbert M. An integral geometry based method for fast form-factor computation. Eurographics 1993.
12. Stewart J, Ghali S. Fast computation of shadow boundaries using spatial coherence and backprojections. SIGGRAPH 1994.
13. Teller S. Computing the antipenumbra cast by an area l light source. SIGGRAPH 1992.
14. Teller S. Visibility computation in densely occluded polyhedral environments. PhD thesis UC Berkeley, 1992
15. Teller S, Hanrahan P. Global visibility algorithm for illumination computations. SIGGRAPH 1993.