

The 3D Visibility Complex

FRÉDO DURAND

iMAGIS-GRAVIR/IMAG-INRIA and MIT LCS

and

GEORGE DRETTAKIS

iMAGIS-GRAVIR/IMAG-INRIA and REVES/INRIA SOPHIA

and

CLAUDE PUECH

iMAGIS-GRAVIR/IMAG-INRIA

Visibility problems are central to many computer graphics applications. The most common examples include hidden-part removal for view computation, shadow boundaries, mutual visibility of objects for lighting simulation. In this paper, we present a theoretical study of 3D visibility properties for scenes of smooth convex objects. We work in the space of light rays, or more precisely, of *maximal free segments*. We group segments that “see” the same object; this defines the *3D visibility complex*. The boundaries of these groups of segments correspond to the *visual events* of the scene (limits of shadows, disappearance of an object when the viewpoint is moved, etc.). We provide a worst case analysis of the complexity of the visibility complex of 3D scenes, as well as a probabilistic study under a simple assumption for “normal” scenes. We extend the visibility complex to handle temporal visibility. We give an output-sensitive construction algorithm and present applications of our approach.

Categories and Subject Descriptors: I.3.7 [Computing Methodologies]: Computer Graphics—*Three-Dimensional Graphics and Realism*

General Terms: Global Visibility

Additional Key Words and Phrases: line-space, visibility classification, Visibility Complex, Visibility Skeleton

1. INTRODUCTION

Visibility computations are central in many computer graphics algorithms. Some of the most common include the determination of the objects visible from a viewpoint, and the computation of umbra and penumbra. Recent techniques such as global illumination simulations require a more global approach, since the mutual visibility of all pairs of points of a scene must be determined. In many cases, visibility is the computation bottleneck, requiring the development of efficient solutions [Holzschuch et al. 1994].

Part of this article builds upon the material from a Eurographics rendering workshop paper [Durand et al. 1996] and a paper at CCCG’97 [Durand et al. 1997a]. A more detailed discussion and comparison with related work has been added (section 6). The catalogue of adjacencies has been added in appendix A. In section 3.2 we have included a probabilistic study of the complexity of our structure in scenes with reasonable assumptions, and in section 4 we discuss extensions of the 3D visibility complex to handle moving objects.

iMAGIS is a joint research project of CNRS/INRIA/UJF/INPG.

Address: iMAGIS/GRAVIR, INRIA Rhône-Alpes, ZIRST, 655 avenue de l’Europe, Montbonnot, F-38334 Saint Ismier Cedex, France

{Fredo.Durand | George.Drettakis | Claude.Puech}@imag.fr

<http://www-imagis.imag.fr/>

Unfortunately, visibility is intricate and hard to comprehend. It is by nature global; a far away object can be visible, and spatially distant objects can have complex interactions. Previous approaches have focused on answering precise queries, and have not really sought to understand the coherence and the globality of visibility. This has led to efficient solutions, but the specificity of visibility issues has often been occluded. We believe that it is useful to better understand visibility, leading to more efficient algorithms.

The computational geometry research community has shown great interest in planar visibility. Work such as the visibility complex by Pocchiola and Vegter [1996b] provides a powerful and elegant framework to comprehend visibility properties in the plane, where visibility problems and coherence are expressed naturally.

Unfortunately, theoretical literature is poorer when 3D visibility is concerned. Researchers typically study the theoretical complexity of problems such as view computation or ray-shooting. The results available are often reductions of algorithms. Even in theory, the approach is mostly constructive and not analytical.

Our aim is first to develop a framework which describes the visibility properties of a three dimensional scene. Any query or visibility property must be easily and naturally expressed. This should be true both for simple (point seen by a ray) or more global queries or properties (limits of shadows, parts of the scene visible from a volume, set of possible views of an object). The notion of coherence must also be explicit, clearly explaining where and how a change occurs. Our approach builds upon the 2D visibility structure and framework developed by Pocchiola and Vegter [Pocchiola and Vegter 1996b; Pocchiola and Vegter 1996a; Durand and Puech 1995].

We borrow the metaphor of function analysis from Koenderink and vanDoorn [1976]. Most visibility methods permit the calculation of the values of a function: What is the point visible along this ray? Are these two points mutually visible? Nonetheless, these queries at a point do not provide more global or structural information which would permit a better description of the function: is it monotonic, continuous, are there inflexions, etc.?

When studying a question or a system, information is not encoded in “uniform”, continuous regions where by definition nothing occurs. The organization or structure are defined by discontinuities, boundaries or *catastrophes* (in the sense of Thom [1972]). We will attempt to determine the sets of rays for which visibility properties remain invariant, and those for which these properties change. In practice, we will determine which rays see the same object, as well as the boundaries of these sets.

This paper proposes a mainly morphological and phenomenological study of visibility. Our approach is *topological*. Topology is based on the notion of neighbourhood. It is particularly well suited to formal study of the notion of coherence. We will however not use a strict mathematical formalism, even though the underlying notions will be present; We have chosen a more descriptive and concrete language.

1.1 Overview

In the next section, we express visibility in line and segment space, and introduce a new data structure, the 3D visibility complex. In Section 3, we analyse the complexity of this data structure: We demonstrate a worst-case bound of $O(n^4)$, and a probabilistic bound of $O(n^{\frac{8}{3}})$ based on a reasonable model of “normal” scenes. Section 4 informally discusses the extension of these concepts to the temporal domain, in order to study visibility among moving objects. We then describe a output-sensitive construction algorithm in Section 5. The review of previous work is postponed until Section 6, which allows us to discuss and

interpret previous within the framework we introduce. We conclude with a discussion of future work issues.

2. INTRODUCTION TO THE 3D VISIBILITY COMPLEX

We will be dealing with scenes composed of convex smooth algebraic surfaces (we do not address their precise representation, which is irrelevant for our purpose). We consider algebraic surfaces of bounded degree. The approach can however be extended to larger class of surfaces. In particular, we have extended it to polygons [Durand et al. 1997a], and piecewise concave algebraic surfaces could be handled using results from singularity theories. See, e.g., the Ph.D. thesis of the first author [Durand 1999], the work by Kergosien [1981], Rieger [1992, 1993] and Petitjean et al. [1992, 1992]. We consider scenes in general configurations. In our context, this means that we exclude incidences (tangency) between lines and objects that are not stable under small perturbations of the scene.

2.1 Line vs. ray visibility

Visibility will be defined in terms of ray-object intersections. If we consider the objects to be transparent, a ray passing through them is not blocked and all the objects a line intersects must be considered. However, if we want to take occlusions into account, only the first object intersected by a ray is relevant. Nevertheless, considering rays also induces redundant information since many collinear rays “see” the same object: Consider a ray r with origin A which intersects an object at point B (Fig. 1(a)). All the rays collinear to r with an origin A' between A and B “see” the same point B . We can thus group these rays into a segment S (Fig. 1(b)).

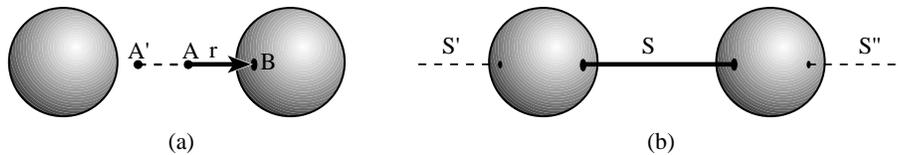


Fig. 1. Maximal free segment. (a) All the rays collinear to r whose origin is between the two spheres “see” point B . (b) These rays are grouped into a *maximal free segment* S . Two other maximal free segments S' and S'' are collinear to S .

We will consider *maximal free segments* which are segments having no intersection with the inside of the objects and whose length is maximal, i.e. 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*. Examples of maximal free segments are given in Fig. 1(b). Segments can also be interpreted as rays which can *see* the two objects at their extremities. A 3D line can be collinear to many segments, separated by the objects the line intersects (see Fig. 1(b)). In this article, we will first introduce concepts 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.

Since sets of segments are not intuitive objects, we will represent them in an appropriate space which will afford a better understanding of intricate visibility relationships. A suitable parameterization will be used for the purposes of illustration and presentation, and for the construction algorithm.

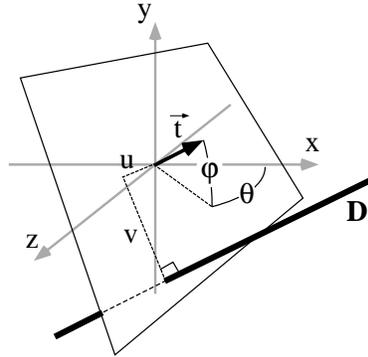


Fig. 2. Line parameterization using two spherical angles for the direction and the Cartesian coordinates of the intersection with an orthogonal plane.

2.2 Parameterization for visualization

We use a parameterization which maps lines in 3D space onto points in \mathbb{R}^4 . We have chosen to decompose the 4 dimensions of line space into two dimensions of direction (the spherical coordinates (θ, φ) of the direction vector of the lines) and a projection (u, v) onto the plane perpendicular to the line and going through the origin (Fig. 2). The axes of the planes are chosen such that u is along $\vec{t} \times \vec{y}$. Our mapping is thus into $[0, 2\pi] \times [-\pi/2, \pi/2] \times \mathbb{R}^2$. Singularities occur at $\varphi = \pm \frac{\pi}{2}$, but since we use this parameterization for the purpose of presentation and visualization we can ignore them without loss of generality.

Visualizing 4D space is very hard. One approach is to use slices or cross sections. In this paper we will fix $\varphi = ct$. 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 two-dimensional $\theta\varphi$ -slices are easier to visualize. They justify in part the choice of the parameterization because they can be interpreted as orthographic projections of the scene.

Note that this parameterization will be used mainly for illustrative and presentation purposes. The concepts presented in Section 2 through 4 are not inherent to any parameterization of lines (except when slices are involved). They could be expressed using only topological notions, although we believe that visualizing in \mathbb{R}^4 greatly helps their understanding. However, the construction sweep introduced in Section 5 will explicitly use the parameterization.

2.3 Tangency volumes

Line visibility. Visibility changes whenever a line becomes tangent to an object: A line which is tangent to an object is the limit between lines which intersect the object and lines which do not intersect it.

The set of lines tangent to one object is a 3-D set in the four-dimensional line space. This means, more intuitively, that a line has 3 degrees of freedom while staying tangent to one object (see Fig. 3). We will call the set of lines tangent to an object the *tangency*

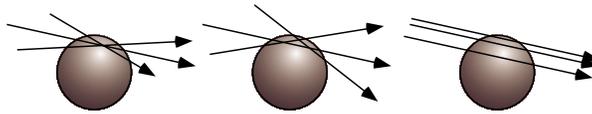


Fig. 3. To remain tangent to an object, a line has three degrees of freedom: (Rotation along the contact point, rotation backward or forward and side translation.

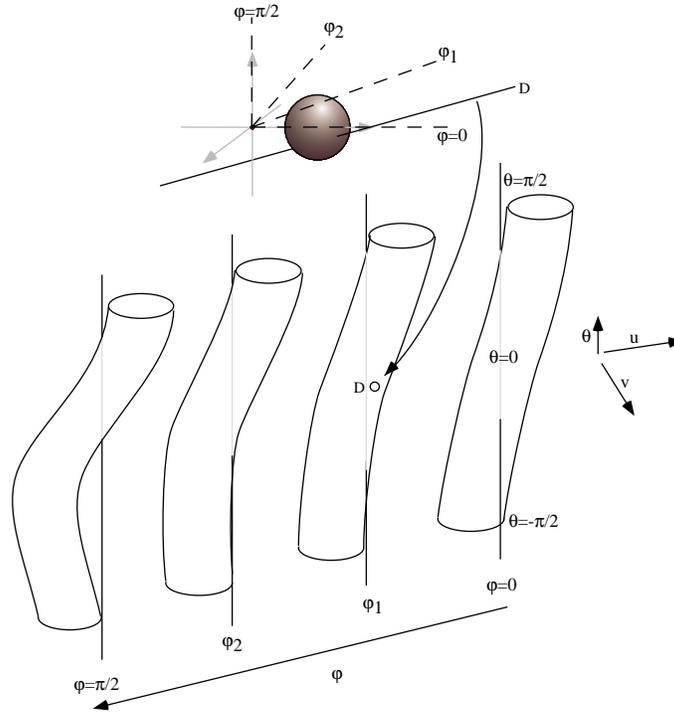


Fig. 4. Tangency Volume of a sphere visualized in \mathbb{R}^4 . The θ axis ($u = 0, v = 0$) is shown for each ϕ -slice providing a better 3D visualization. In the left-hand ϕ -slice, which corresponds to the discontinuity in the parameterization for $\phi = \frac{\pi}{2}$, the “cylinder” just turns around the θ axis. The line D intersects the object and is inside the tangency volume.

volume of this object.

Fig. 4 shows a representation of the tangency volume of a sphere. For each ϕ -slice, the set of tangents is a sort of 2D “tube”, forming a 3D structure in \mathbb{R}^4 . If we consider a 2D $\theta\phi$ -slice (horizontal in Fig. 4), the set of tangents sharing that direction is mapped to a circle in \mathbb{R}^4 . This is also true in the general case: Because of the definition of u and v , the visualization of the set of tangents in one direction has the shape of the silhouette of the object along the corresponding orthographic projection. More precisely, the visualization of the set of tangents in one direction correspond to the *fold* along that direction. The fold is the “silhouette” of a semi-transparent version of the object [Petitjean 1992], which is logical since we deal with line visibility and do not take occlusion into account at this point.

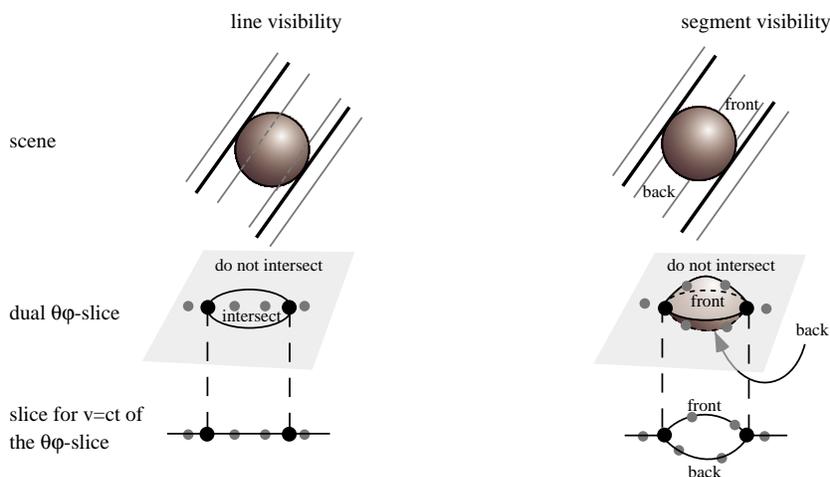


Fig. 5. Visibility for $\theta = ct$ and $\phi = ct$. If we consider lines (on the left), visibility can be described by a planar structure (below). But if we consider segments (on the right) we have different levels on this plane depending on the side of the object. The set of segments which do not 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 silhouette.

If a line lies *on* the tangency volume, it means that it is tangent to the object. If the line lies “inside” the 4D set bounded by the tangency volume, the line intersects the object. An example is line *D* in Fig. 4.

Segment visibility and auxiliary complex. Let us now consider visibility with occlusion, that is, segment visibility rather than line visibility. A line which intersects the object is collinear to two segments, one before and one after the object.

Consider a $\theta\phi$ -slice such as that on the lower left of Fig. 5. Viewed in that slice, the sets of lines that intersect, and those that do not intersect the object are bounded by the fold 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 both visualized with the same point in \mathbb{R}^4 , as shown in the left of Fig. 5. 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 neighbourhood composed of segments that do not intersect the objects, some that see the front, and some that see the back.

To differentiate the segments in our visualization, we need to add a pseudo-dimension. This is not a continuous dimension since it is only used to sort all the collinear segments. If we impose $\theta = ct$, $\phi = ct$ and $v = ct$, the sets of segments can be represented by the embedding of a graph shown on the lower right. Each tangent corresponds to a vertex of the graph. This graph is a 1D structure embedded in the uv real plane.

We recall that a n -manifold is a set for which the neighbourhood of each point is homeomorphic to \mathbb{R}^n . The embedding of a graph is manifold almost everywhere (homeomorphic to \mathbb{R}), except at nodes of valence greater than 2, where the branching needs to be embedded in a higher dimensional space.

Similarly, for a $\theta\phi$ -slice, the sets of segments is a 2-manifold almost everywhere, except

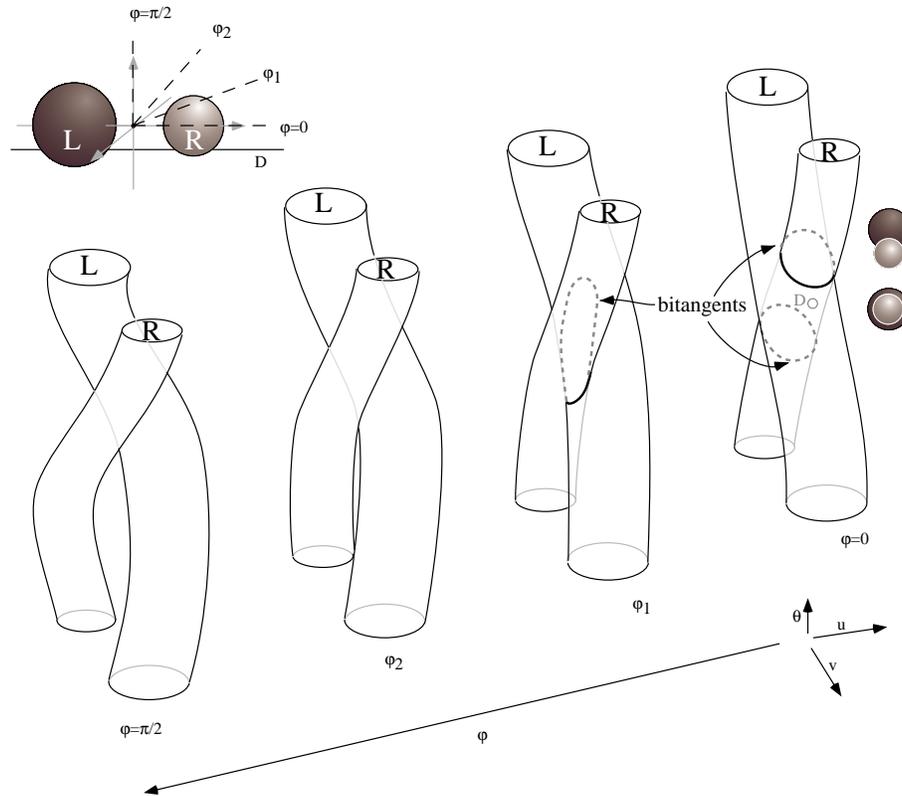


Fig. 6. Dual arrangement for two spheres.

at the tangents. We call the partition of the segments of direction (θ, φ) according to their visibility the *auxiliary complex* for (θ, φ) (see also Fig. 7). It can be seen as a generalized orthographic view where all the objects, hidden or visible, are organized into layers. It is a 2-dimensional cellular decomposition, but it is not formally a cell-complex (a.k.a. CW-complex) because its faces are not necessarily homeomorphic to disks [Hatcher 2001] (e.g., they can have holes if one object is in front of another object).

In a similar manner, the space of maximal free segments of a scene is a 4-manifold almost everywhere, except at the branchings on the tangency volumes.

2.4 Bitangents dual arrangement and the Visibility Complex

Line visibility and the dual arrangement. We now consider two objects. If a line is inside the tangency volumes of both objects, it intersects them both. The tangency volumes give us a partition of the space of lines according to the objects they intersect.

Definition: *The Dual Arrangement is the partition of the lines of 3D space into connected components according to the objects they intersect.*

Its 4-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 the 2D set of lines tangent to the two objects (bitangents). A bitangent corresponds to a

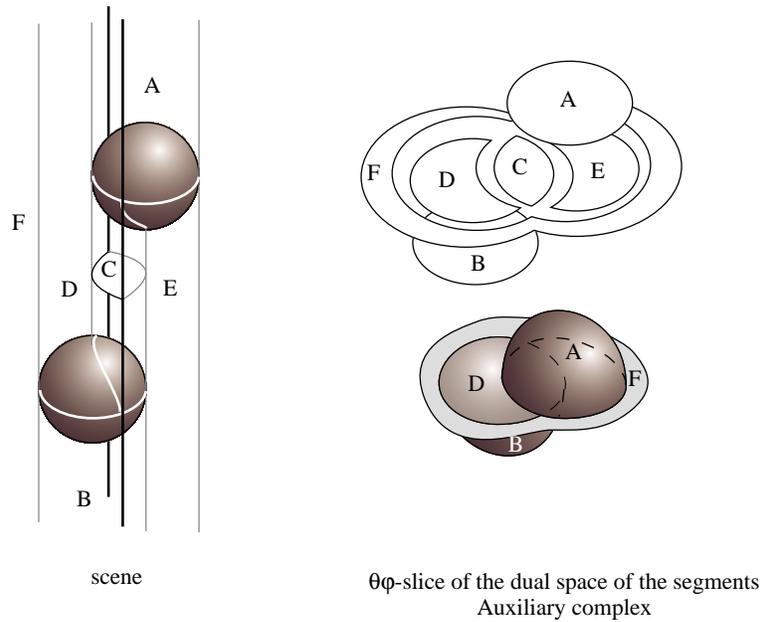


Fig. 7. Auxiliary Complex for two spheres. In the lower representation, only the points on the surfaces represented are associated with segments. In the upper view, the faces of the auxiliary complex have been moved outwards to make their incidences easier to understand. Face F is infinite.

t -vertex in an image, that is, the visual intersection of two object folds.

For a ϕ -slice, the set of bitangents is mapped onto a space curve. This is shown as a dashed line in Fig. 6 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 visualization slice of a 4D face is a volume corresponding to the intersection of the interior of the two cylinders.

Segment visibility and the Visibility Complex. An auxiliary complex for two objects is visualized on Fig. 7 for a given direction. Its visualization is still delimited by the “silhouette” (corresponding to tangents) of the objects, but for example the “silhouette” 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.

Definition: *The 3D visibility complex is the partition of the maximal free segments of 3D space into connected components according to the objects they touch.*

The 4-faces of the complex are 4D sets of segments with the same objects at their extremities. They are bounded by portions of the tangency volumes which are 3D. 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 not correspond to a bitangent segment because of occlusions.

The 3D visibility complex is a 4-dimensional cellular decomposition, but like the auxiliary complex, it is not formally a cell-complex [Hatcher 2001] because of the complex topology of its faces. In contrast, the 2D visibility complex [Pocchiola and Vegter 1996b] is actually a cell-complex.

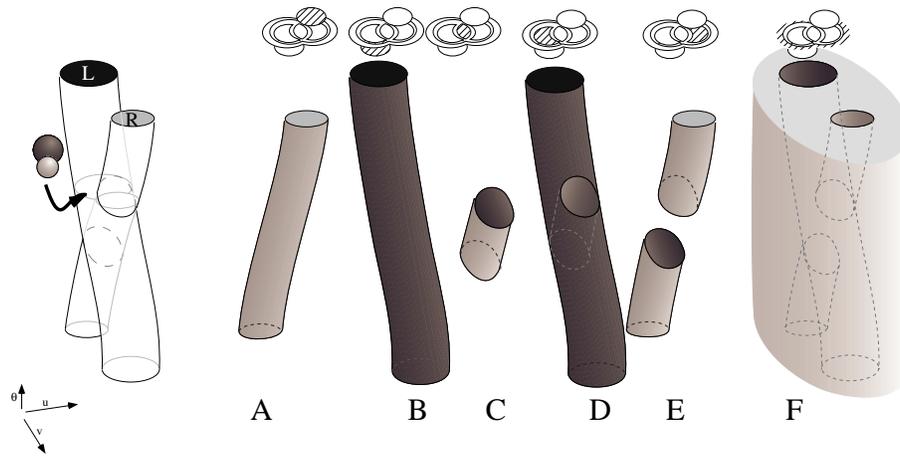


Fig. 8. φ -slice for $\varphi = 0$ of the faces of the visibility complex of the previous scene (Fig. 6 and 7). A is the set of segments that see the front of R , B is the set of segments that see the back of L . C is the set of segments between L and R . It can be interpreted as the intersection of set of lines that see L and the set of lines that see R , and in line space it has the shape of $A \cap B$. D is the set of segments that see the front of L . Since 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 neither of the two spheres. It is the complement of $A \cup B$.

Fig. 8 is a φ -slice for $\varphi = 0$ of all the faces of the 3D visibility complex for the scene composed of two spheres of Fig. 6. 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 in our visualization, some of these slices overlap in \mathbb{R}^3 , since many collinear segments are mapped onto the same point. A φ -slice of the complex is not manifold at the branchings of the tangency volumes.

2.5 Tritangents

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

Fig. 9 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_{θ_1} . It is the intersection of the tangency volume of R and B minus the tangency volume of G . The tritangents are the points in white. Note that because of occlusion by the sphere G , lines that are bitangents of the R and B do not correspond to bitangent segments. This is shown in Fig. 10 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 do not correspond to bitangent segments. 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 .

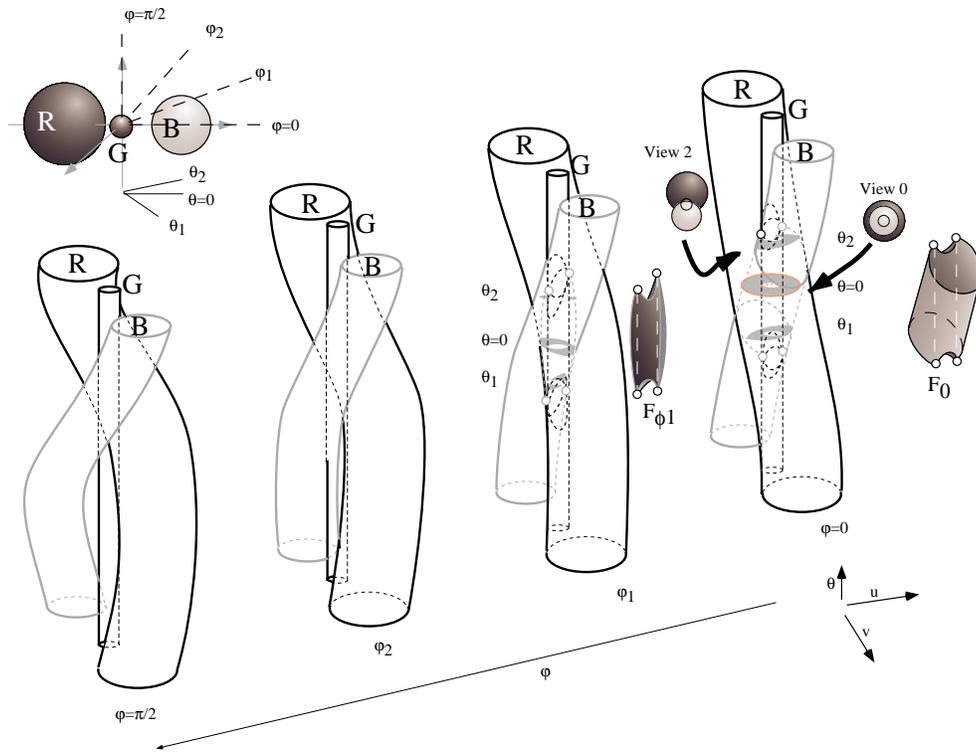


Fig. 9. Visibility Complex of a scene of three spheres.

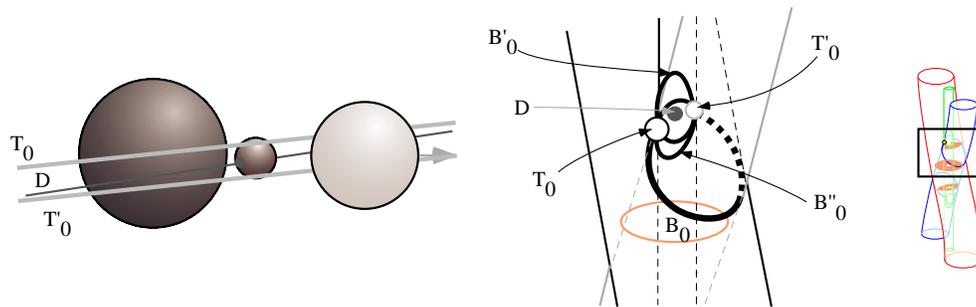


Fig. 10. Zoomed view of the ϕ -slice $\phi = 0$.

Tritangents induce *visual events* or *visibility events*, e.g. [Rieger 1987; Rieger 1990; Petitjean 1992; Petitjean et al. 1992]. Visual events describe the *qualitative* (or topological) changes in visibility. Consider the example in Fig. 11. As the viewpoint moves downwards, sphere *A* becomes hidden by the conjunction of *B* and *C*. This occurs when the viewpoint lies on a tritangent. In what follows, we will interchangeably refer to a visual events or to the set of lines inducing it.

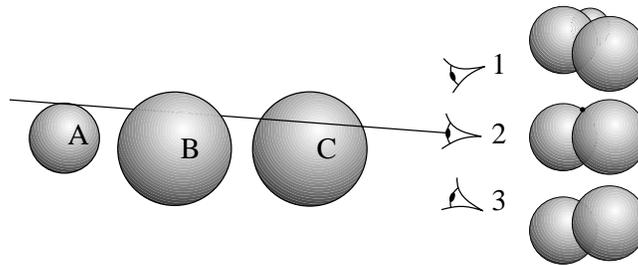


Fig. 11. Tritangency visual event. At viewpoint 2, sphere *A* becomes hidden by sphere *B* and *C*. This occurs as the viewpoint lies on a line tangent to the three spheres.

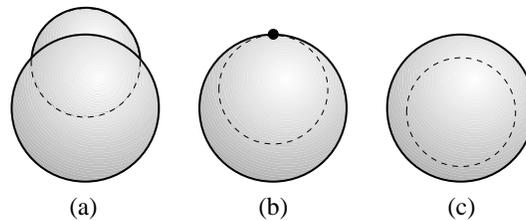


Fig. 12. Tangent crossing singularity. As the viewpoint moves downwards, the back sphere becomes hidden by the front-most one. At viewpoint (b) a singularity occurs (highlighted with a point): the two spheres are visually tangent.

Note that a scene does not necessarily contain tritangents (this is for example the case in Fig. 13(a) page 13).

2.6 Tangent crossing

There is another type of visual event called a *tangent crossing*. It is induced by planes tangent to two objects, and to the lines going through the corresponding two points of tangency. Consider the case of a sphere hidden behind another sphere. If the viewpoint is moved, the sphere will become visible when the two spheres are visually tangent (Fig. 12). The line going through the viewpoint and the point of visual tangency is contained in a plane tangent to the two spheres.

How is this interpreted in our parameterization? These tangent crossing critical lines are subsets of bitangents. They correspond to minima or maxima with respect to θ in the φ -slices. To understand why, remember that a (φ, θ) -slice of the complex corresponds to an orthographic view. Consider a rotation at $\varphi = ct$ starting at a direction for which the two objects are distinct in the view. Bitangents (*t*-vertices) will appear in the orthographic view at a tangent crossing. The tangent crossing thus corresponds to the first θ at which a set of bitangents starts.

The same reasoning can be applied to show that tangent crossings are also minima or maxima in θ -slices. If occlusion occurs because of a third object, the tangent crossing can be discarded and the minimum or maximum of a bitangent set can be another event, as is the case in Fig. 10 where the tritangent T'_0 is the maximum of the set B_0 .

Tangent crossings, as all visual events, are induced by 1D sets of lines. These events

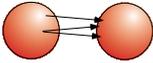
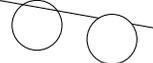
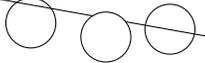
Dimension	Scene configuration	ϕ -slice in line space	Name	Notation
4			4-face	
3			tangency volume	T
2			bitangency face	T+T
1			tritangency arc	T+T+T
1			tangent crossing	T++T
0			vertex	T+T+T+T
0			vertex	T++T+T

Table 1. Elements of the visibility complex

are crucial for dynamic maintenance of views, aspect graphs and discontinuity meshes. Triple-tangency and tangent-crossing are the only visual events for convex smooth objects (see e.g. the catalogue of Kergosien [1981] or Petitjean [1992] and his co-authors [1992]).

2.7 Notation and vertices

Before describing the 0-dimensional faces of the 3D visibility complex, we introduce a notation for the different faces and the corresponding events, which we borrow from Petitjean et al. [1992, 1992]. The tangency between an object and a line is noted T , and multiple tangencies are noted using $+$. A bitangent is noted $T + T$, and a tritangent is $T + T + T$. The higher-order incidence at a tangent crossing is denoted by $++$. A tangent crossing is thus noted $T + +T$.

Using this notation, the first kind of 0-faces of the complex are $T + T + T + T$ segments (tangent to four objects). The second sort of vertices are $T + +T + T$ segments, which are tangent to 3 three objects, two of them sharing a tangent plane.

The dimensions and incidences of the boundaries of the faces of the 3D Visibility Complex are summarized in table 2.7. Note that a scene can have a visibility complex with neither vertex nor tritangency arc.

3. COMPLEXITY ANALYSIS

In this section, we analyze the complexity of the 3D visibility complex. The algebraic objects are assumed to have bounded degree. In what follows, n represents the number of smooth convex objects.

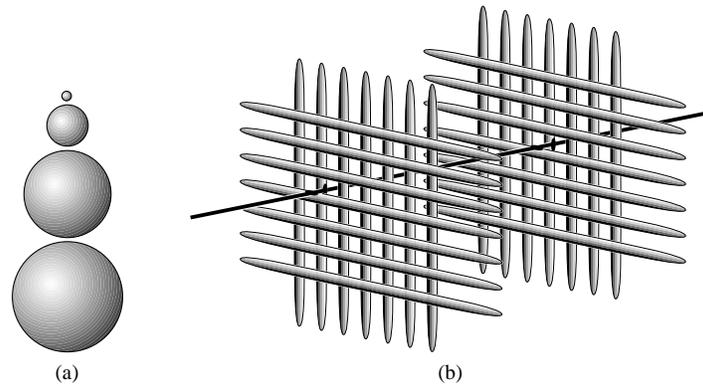


Fig. 13. Lower and upper bound scenes for the visibility complex (a) Scene with an $O(n)$ Visibility Complex. (b) Scene with a $O(n^4)$ Visibility Complex (an example of $T+T+T+T$ critical segment is shown).

3.1 Lower and upper bound

THEOREM 1. *The size of the 3D visibility complex is $\Omega(n)$ and $O(n^4)$ where n is the number of objects of the scene.*

The $\Omega(n)$ bound is trivial since there exists at least one tangency volume for each 3D object. The $O(n^4)$ bound comes from the number of $T+T+T+T$ events. However, we have to prove that the size of the complex is dominated by their number¹. For this purpose, we need to prove that the number of faces of dimension 4, 3 and 2 is bounded by a constant factor of the number of faces of dimension 1. A sketch of the proof can be found in the appendix.

The size of the complex is thus bounded by the number of 1-faces which are not adjacent to a 0-face plus the number of 0-faces. For each event type, the number of possible systems of algebraic equations depends on the number of objects implicated, the $T+T+T+T$ critical line sets are thus the most numerous with their number being bounded by $O(n^4)$.

We show in figure 13(a) an example of a scene with a visibility complex of size $O(n)$: there is one $T++T$ face for each pair of neighbours spheres. Note there is no $T+T+T$ -face in this case because of the decreasing radius of the spheres. There is no 0-face because there can be no $T+++T$ nor $T+T+T+T$ segment. The scene in figure 13(b) is similar to the one given by Plantinga and Dyer [1990] and has a $O(n^4)$ visibility complex. There are two “grids”, each one composed of two very slightly perturbed (to avoid degeneracies) orthogonal sets of $\frac{n}{4}$ parallel thin ellipsoids. Consider an ellipsoid in each of the four sets: there is always a $T+T+T+T$ critical segment.

3.2 Probabilistic complexity

Unfortunately, these theoretical bounds do not provide much information about practical complexity in the case of “normal” scenes. The scenes shown in figure 13 are very different from scenes typically used in computer graphics. Some probabilistic approaches using a model of “reasonable” scenes in the spirit of de Berg et al. [1997] should be made.

¹for example, the size of a tetrahedralisation is not necessarily dominated by the number of vertices, it can be quadratic. See, e.g. the book by Boissonnat and Yvinec [1998], chapter 13

As a first step, we study the number of triples of objects generating tritangents in “reasonable” scenes. Potentially, any triple of object can interact and generate tritangents. However, in order to obtain scenes exhibiting $O(n^3)$ triple interactions, one has to consider asymptotically infinitely thin objects, such as in Fig. 13(b). In contrast, implementations by e.g. Drettakis and Fiume [1994] or Durand et al. [1997b] have shown that the practical scenes they tried do not contain $O(n^3)$ triple-tangent interactions. This discrepancy can be explained because, in typical computer graphics scenes, the size of the objects is bounded, and the geometric complexity is increased by adding smaller details or by placing many objects next to each other, not by adding many interleaved infinite lines. We propose here a simplified probabilistic model for these “reasonable” scenes, and show that under our assumptions the mean number of triples of objects generating tritangents is $O\left(n^{\frac{7}{3}}\right)$.

We consider a scene model where objects of bounded extent are uniformly distributed inside a finite sphere. R is the radius of the scene, and r is the maximum radius of an object. The density of objects is constant as their number varies: If n is the number of objects we have $\frac{R}{r} = \Theta(\sqrt[3]{n})$. We first study the probability that, given two objects, a third object generates a tritangent with them. We will show that this can be expressed as a volume ratio. We will then integrate over all objects to get a bound on the mean total number of interacting triples. To simplify our calculation, we use bounding spheres enclosing the objects.

Probability given two objects. Consider two objects A and B separated by a distance x . We want to obtain a bound on the probability $P_{AB,x}$ that a third object C generates a tritangent with A and B . This will happen only if the sphere bounding C intersects the hourglass volume formed by two cones and a cylinder tangent to the two spheres enclosing A and B (Fig. 14). The probability that a sphere of radius r intersects this volume is equal to the probability that the center of the sphere intersects an hourglass volume dilated by r (i.e. the Minkowski sum of the hourglass and the sphere). Since we assume that the object centers are uniformly distributed, this probability is equal to the ratio between the volume of the dilated hourglass and the volume of the sphere bounding the scene.

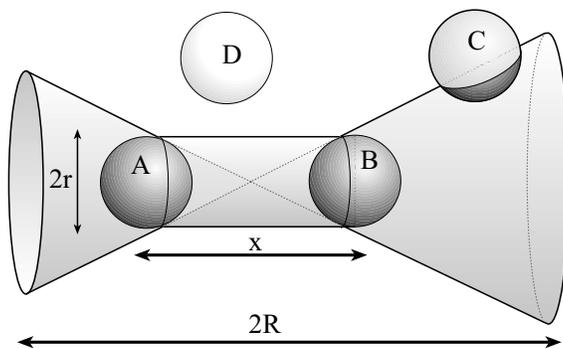


Fig. 14. An object C can generate a tritangent with A and B only if it intersects the hourglass volume. D does not generate tritangents with A and B .

We now compute the volume of the dilated hourglass. Recall that x is the distance between the centers of the bounding spheres of A and B . We distinguish the case where

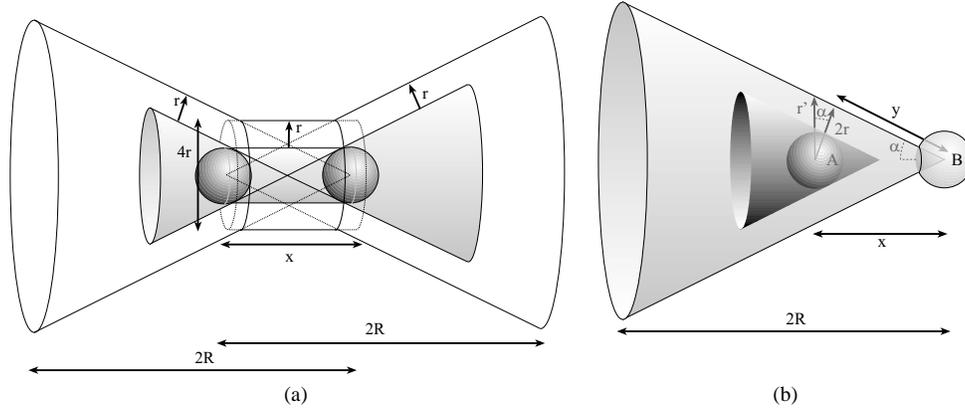


Fig. 15. (a) Dilated hourglass volume. (b) Construction of one cone of the dilated hourglass.

$x \leq 4r$ and the case where $x > 4r$. In the first case we trivially bound the volume by the volume of the scene sphere, *i.e.*, $P_{AB, x \leq 4r} \leq 1$. We will see that this case is asymptotically negligible.

In the latter case, we bound the volume of the dilated hourglass by the sum of a cylinder of length x and diameter $4r$ and two cones with apex at the centers of A and B and height $2R$, as shown in Fig. 15. The volume of the cylinder is $vol_{cylinder} = 4\pi r^2 x$.

Let us now bound the volume of one cone. We compute its radius r' at height x , using similar right-angle triangles:

$$\begin{aligned} \frac{r'}{2r} &= \frac{x}{y} \\ r' &= 2r \frac{x}{\sqrt{x^2 - 4r^2}} \end{aligned}$$

To show that r' is a decreasing function of x , we compute the derivative:

$$\frac{dr'}{dx} = \frac{-2r}{(x^2 - 4r^2) \sqrt{\frac{x^2 - 4r^2}{x^2}}}$$

And since $x > 4r$, this derivative is always negative and r' is a decreasing function of x . We can thus bound it by its value for $x = 4r$, that is less than $r' \leq \frac{4}{\sqrt{3}}r$. The volume of one cone is then:

$$vol_{cone} = \pi \left(r' \frac{2R}{x} \right)^2 \frac{2R}{3} \leq \pi \left(\frac{4}{\sqrt{3}} r \frac{2R}{x} \right)^2 \frac{2R}{3} \leq \frac{128 \pi r^2 R^3}{9x^2}$$

The total volume of the dilated hourglass is then bounded by:

$$\begin{aligned} vol_{hourglass} &\leq 2vol_{cone} + vol_{cylinder} \\ &\leq 4\pi r^2 + 2 \frac{128\pi r^2 R^3}{9x^2} \end{aligned}$$

We divide by the volume of the scene sphere to obtain a bound on the probability:

$$\begin{aligned} P_{AB,x} &\leq \frac{\text{vol}_{\text{hourglass}}}{\text{vol}_{\text{sphere}}} \\ &\leq \frac{3r^2x}{R^3} + \frac{64r^2}{3x^2} \end{aligned}$$

Probability for all the objects. We now have a bound on the probability, given a pair of objects at distance x that a third object C generates a tritangent with them. We have to compute the total mean number of such triples, that is consider every possible C , and every possible pair A, B at every possible distance x .

Until now, we have considered that the distance x between A and B is fixed. Consider now a random pair A, B . We have to take into account the probability that they lie at distance x . We use a result of integral geometry by Santaló ([1976], page 212, or see also the book by Mathai [1999], page 296) on the probability distribution of distances of pairs of points inside a sphere of radius R :

$$\mu(x) = \frac{x^2}{R^3} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right), \quad 0 \leq x \leq 2R,$$

that is, the probability that two random points in a sphere of radius R lie at a distance between x and $x + dx$ is $\mu(x) dx$.

To obtain the total mean number of triples M we integrate over x and multiply by the number of triples of objects, $\frac{n(n-1)(n-2)}{6}$. Recall that we have distinguished the case $x \leq 4r$ where the probability is bounded by 1 and the case $x > 4r$ where the probability is given by a ratio of volumes.

$$\begin{aligned} M &< \frac{n(n-1)(n-2)}{6} \int_0^{2R} P_{AB,x} \mu(x) dx \\ &< \frac{n(n-1)(n-2)}{6} \left(\int_0^{4r} 1 * \mu(x) dx + \int_{x=4r}^{2R} \left(\frac{3r^2x}{R^3} + \frac{64r^2}{3x^2} \right) \mu(x) dx \right) \end{aligned} \quad (1)$$

We now derive bounds on the two integrals in inequality (1). The first integral is:

$$\begin{aligned} \int_{x=0}^{4r} \frac{x^2}{R^3} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx &= \frac{(4r)^3}{R^3} - \frac{9(4r)^4}{16R^4} + \frac{3(4r)^6}{96R^6} \\ &= 128 \frac{r^6}{R^6} - 144 \frac{r^4}{R^4} + 64 \frac{r^3}{R^3} \end{aligned} \quad (2)$$

The second integral in inequality 1 is:

$$\int_{x=4r}^{2R} \left(\frac{3r^2x}{R^3} + \frac{64r^2}{3x^2} \right) \mu(x) dx = \int_{x=4r}^{2R} \frac{3r^2x}{R^3} \frac{x^2}{R^3} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx + \int_{x=4r}^{2R} \frac{64r^2}{3x^2} \frac{x^2}{R^3} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx \quad (3)$$

Since the integrand is positive, we can bound the first term by extending the lower limit to

0:

$$\begin{aligned}
\int_{x=4r}^{2R} \frac{3r^2 x^3}{R^6} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx &\leq \int_{x=0}^{2R} \frac{3r^2 x^3}{R^6} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx \\
&\leq \frac{3r^2}{R^6} \left(\frac{3(2R)^4}{4} - \frac{9(2R)^5}{20R} + \frac{3(2R)^7}{112R^3} \right) \\
&\leq \frac{108r^2}{35R^2}
\end{aligned} \tag{4}$$

The second term in Eq. 3 gives:

$$\begin{aligned}
\int_{x=4r}^{2R} \frac{64r^2}{3R^3} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx &\leq \int_{x=0}^{2R} \frac{64r^2}{3R^3} \left(3 - \frac{9x}{4R} + \frac{3x^3}{16R^3} \right) dx \\
&\leq \frac{64r^2}{3R^3} \left(3(2R) - \frac{9(2R)^2}{8R} + \frac{3(2R)^4}{64R^3} \right) \\
&\leq \frac{48r^2}{R^2}
\end{aligned} \tag{5}$$

Combining the results of equations (1), (2), (4) and (5) we obtain:

$$\begin{aligned}
M &\leq \frac{n(n-1)(n-2)}{6} \left(128 \frac{r^6}{R^6} - 144 \frac{r^4}{R^4} + 64 \frac{r^3}{R^3} + \frac{108r^2}{35R^2} + \frac{48r^2}{R^2} \right) \\
M &= O\left(n^3 \frac{r^2}{R^2}\right)
\end{aligned}$$

And since we assumed that $\frac{R}{r} = \Theta(\sqrt[3]{n})$, we obtain

$$\begin{aligned}
M &= O\left(n^3 \left(\frac{1}{\sqrt[3]{n}}\right)^2\right) \\
&= O\left(n^{\frac{7}{3}}\right)
\end{aligned}$$

The mean number of triples of objects generating tritangents is $O\left(n^{\frac{7}{3}}\right)$ (that is, $O(n^{2.33})$) in our model instead of $O(n^3)$ when objects of unbounded extent are considered. In the statistics which we have obtained for a set of test scenes [Durand et al. 1997b], the number of $T + T + T$ is closer to quadratic than to cubic, confirming our probabilistic bound.

Similar arguments can be used to compute the mean number of $T + T + T + T$ 0-faces. Since the presence of a second sphere inside the hourglass is independent of the presence of the first sphere, we can show that the bound is then $O\left(n^{\frac{8}{3}}\right)$ (i.e., $O(n^{2.67})$) instead of $O(n^4)$.

4. EXTENSION TO TEMPORAL VISIBILITY

In this section we describe an informal extension of the Visibility Complex to dynamic scenes. The formal analysis of temporal visibility requires the handling of higher-order singularities, which are to our knowledge still poorly understood.

We however show that a framework similar to the static case provides an intuition about the changes in visibility as objects move smoothly. We consider any smooth movement of the objects (including deformations), and that all objects can be dynamic. The temporal

dimension is considered continuous, even though a representation at certain time steps will be necessary for clarity.

4.1 Temporal visual event

We adopt the same approach as for the static case: we try to determine when visibility changes in a dynamic 3D scene. We are interested in the topological changes of the visibility complex as the objects move. Consider the situation depicted in Fig. 16(a). Sphere A moves upward. At time t_1 , there is no segment with extremities on A and C because of the occlusion by B . At time t_3 , some segments see A and C because A has moved. The limit configuration occurs at time t_2 where a plane is tangent to the three spheres. Segment $A++B++C$ corresponds to this tritangent plane. It is a *temporal visual event*. It is similar to the 0-faces of the static complex, with one additional incidence.

In Fig. 16(b) we represent the 0 and 1-dimensional faces of the static visibility complex for some timestep. At time t_1 the complex only has two 1-faces and no 0-face. At time t_3 six $T++T+T$ 0-faces have appeared, as well as some new $T++T$ and $T+T+T$ 1-faces. A topological change in the structure of the 3D visibility complex has occurred.

4.2 The Temporal Visibility Complex

We add the temporal dimension to the space of maximal free segments. A temporal segment is defined as a maximal free segment in space together with a time value. Temporal segment space for a dynamic scene is thus a 5-manifold almost everywhere, except at the tangents.

We define the *Temporal Visibility Complex* as the partition of the temporal maximal free segments according to the object that they touch.

The k -faces of the Temporal Visibility Complex correspond to the $(k-1)$ -faces of the static Visibility Complex. A time-slice of the temporal Visibility complex is an instance of static Visibility Complex. The 0-faces of the Temporal Visibility Complex correspond to the *temporal visual events*. These correspond to changes in the topological structure of the visibility complex. Maintaining the Visibility Complex as objects move is equivalent to a time-sweep of the Temporal Visibility Complex.

Fig. 16(c) shows the 0, 1 and 2-faces of the temporal complex for the scene in Fig. 16(a). At the temporal visual event, the slice of the temporal complex is modified, new $T+T++T$ faces appear as well as some 1-faces of the static complex. Temporal visual events are summarized in table 2 for scenes of smooth objects. They are very similar to the 0-faces of the static visibility complex. As aforementioned, the study of their precise catalogue and adjacencies may require the treatment of complex singularities.

5. OUTPUT-SENSITIVE SWEEP

We now present an output-sensitive construction algorithm for the static Visibility Complex of convex algebraic surfaces. This algorithm has not been implemented, but gives a bound on the complexity of the construction of the complex. The algorithm can easily be extended to handle polyhedra as well [Durand et al. 1997a].

Because our surfaces have bounded degree, the computation of tangency and other differential properties can be performed in constant time. Our algorithm is a double sweep with a preprocessing phase. First the scene is swept by a horizontal plane and a 2D Visibility Complex [Pocchiola and Vegter 1996b] of the ϕv -slice is maintained (figure 17(b)). We then sweep ϕ (figure 17(c)), but some 0-faces cannot be detected during this sweep and

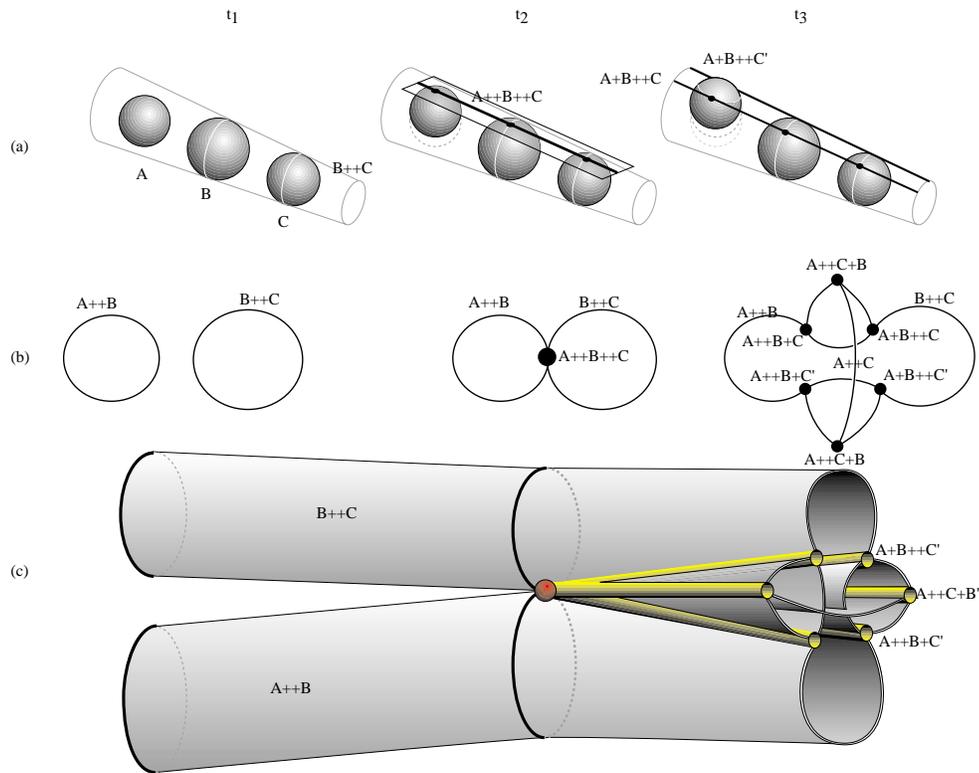


Fig. 16. Temporal visual events and the Temporal Visibility Complex. (a) Situation in 3D space. As A moves upward, it becomes visible from C at time t_2 where a segment $(A++B++C)$ lies in a plane tangent to the three spheres. The envelope of $B++C$ is represented as a cone. Only two $T++T++T$ 0-faces are represented, but there are actually 6. (b) 1 and 0-dimensional faces of the visibility complex for each timestep. (Note that this is a projection, not a slice of line space. 1D sets are actually represented as 1D curves). At time t_3 we have not labeled all the 1-faces. Note that the graph of 1-faces is not planar (it is 4-dimensional), to draw it the figure must contain the crossings of $A++C$. (c) Structure of a part of the Temporal Visibility Complex in temporal segment space. Only 0, 1, and 2-dimensional faces are represented. The faces of the static complex represented in (b) are slices of the temporal complex. The 1-faces of the temporal complex are represented as thin cylinders. Note that they correspond to the 0-faces of the static complex and that they appear only at t_3 . Note also that due to the non-planarity of the slice, we cannot represent the 2-faces of the temporal complex without intersections in the middle.

have to be preprocessed.

5.1 Sweeping the initial slice

To build the initial φ -slice, we first maintain a $\varphi\nu$ -slice of the 3D visibility complex which corresponds to the 2D visibility complex of the sweeping plane.

The 2D visibility complex. We first briefly review the 2D visibility complex (e.g. Pocchiola and Vegter [1996b, 1996a], Rivière [1997a] or Durand and Puech [1995]) and its relation with the 3D visibility complex. The 2D visibility complex is the partition of the maximal free segments of the plane according to the objects they touch. Its 2D faces are

Type	Configuration
T+T+T+T+T	
T++T+T+T	
T++T++T	

Table 2. Vertices of the temporal visibility complex. These configurations can occur because of the movement of any (or all) of the involved objects.

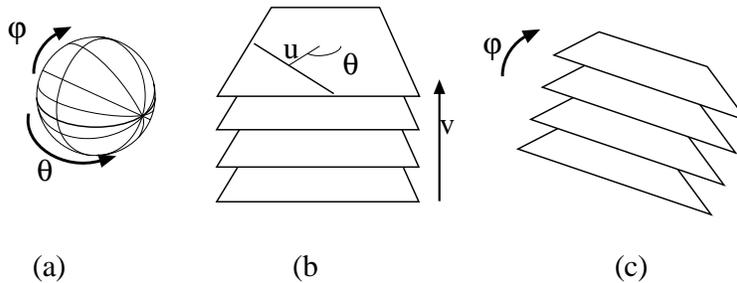


Fig. 17. (a) Parameterization of the directions. (b) Initial v sweep. (c) ϕ sweep.

connected components of segments touching the same objects (they are ϕv -slices of the 4-faces of the 3D visibility complex). They are bounded by edges which correspond to segments tangent to one object (ϕv -slices of the 3-faces T) and vertices which are free bitangents of the 2D scene (ϕv -slices of 2-faces $T + T$).

Since a view around a point corresponds to the extremities of the segments going through this point, it corresponds to the traversal of the 2D visibility complex along the 1D path of these segments. The object seen along a ray changes when the path traverses a new face, which occurs at an edge of the 2D complex.

Sweep. The 2D visibility complex has to be updated when the sweeping plane is tangent to an object and when three 2D slices of objects share a tangent.

When the sweeping plane starts intersecting an object, we have to “insert” this object in our 2D complex. This is done by computing a 2D view around the point of tangency or around the vertex using the current 2D visibility complex. This can be done in $O(v \log n)$ where v is the size of the view using the techniques described by Rivière [1997b]. When the path of this view crosses an edge of the 2D complex it corresponds to a new $T + T$ or $V + T$ face of our 3D complex.

Symmetrically, when an object stops intersecting the sweeping plane, the corresponding faces of the 2D visibility complex are collapsed. These faces are those along the chains of edges in line space corresponding to segments tangent to this object. Their removal can be done in $O(v)$ where v is again the size of the view.

As the plane moves, three slices of objects can share a tangent (corresponding to a $T + T + T$ face of the 3D complex), in which case the 2D visibility complex is updated using the technique of Rivière [1997b]. Basically, for each bitangent we compute the value of v where it will become tangent to a third object and store these sweep-events in a queue, which requires time $O(\log n)$ whenever a bitangent is created.

Finally, a bitangent of the 2D complex can correspond to a common tangent plane. For each bitangent, we compute the value of v for which it will lie on a bitangent plane and insert this sweep-event in the queue. Of course, these sweep-events have to be discarded if the bitangent is collapsed before.

5.2 Principle of the φ sweep

We now have computed a φ -slice of the 3D visibility complex. It is the partition of the segments contained in the set of horizontal planes. In this φ -slice, 1-faces of the complex have dimension 0, 2-faces have dimension 1, and so on.

During the φ -sweep (Fig. 17(c)) we maintain this φ -slice as well as a priority queue of sweep-events. In what follows, we will only describe the update of the 1-faces of the visibility complex, the update of the higher dimensions is done at each sweep-event using a catalogue of adjacencies of the 1-faces which can be found in appendix A. As stated before, the number of adjacent upper-dimensional faces is bounded; their update does not affect the complexity.

This φ sweep can also be understood by projecting the 1-faces and the vertices of the visibility complex onto the direction sphere S^2 represented in Fig. 17. The sweep then corresponds to the rotation of a great circle. Note that some intersections on S^2 are not relevant for the complex (this is the same difference as a t-vertex in a view and an actual vertex of the scene).

We first prove that some sweep-events are *regular*. A sweep event is regular if it corresponds to the collapsing of a 1D element of the sweep front as its two extremities merge [Mulmuley 1991]. Intuitively, this means that the event can be easily predicted from the current sweep front. In our case, such sweep-events can be detected by computing the value of φ for which it will collapse for each 1D component of the φ -slice. We will next study the case of irregular sweep-events, which are harder to treat.

5.3 Regular 0-faces

Consider a $T_1 + T_2 + T_3 + T_4$ segments with extremities O_0 and O_5 and elevation angle φ_0 (Fig. 18). Consider the 1D critical line set $T_1 + T_2 + T_3$. We locally parameterize it by φ and call it $l(\varphi)$. The ruled surface described by $l(\varphi)$ cuts O_4 at φ_0 . Two 1-faces of the complex are associated with $l(\varphi)$, one for $\varphi < \varphi_0$ and one for $\varphi > \varphi_0$; one has O_5 at its extremity, the other O_4 . Similarly for $T_2 + T_3 + T_4$. Moreover the two 1-faces before φ_0 are adjacent to a 2-face $T_2 + T_3$. In the φ -slice, this 2-face is a 1D set bounded by the slices of $T_1 + T_2 + T_3$ and $T_2 + T_3 + T_4$. This 1D set collapses at φ_0 , it is thus a regular sweep-event. It can be detected by considering the adjacent $T + T + T$ faces in the φ -slice and maintaining a priority queue.

The $T + T + T$ faces can be handled in the same way because they are adjacent to a pair of $T + T$ and a pair of $T + T + T$ 1-faces, and the faces of a pair are associated with the same line set.

The projection of regular events onto S^2 is similar to that illustrated in Fig. 19(a).

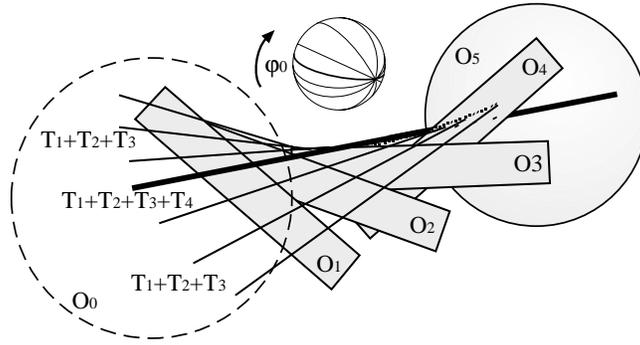


Fig. 18. $T + T + T$ critical line sets adjacent to a $T + T + T + T$ critical line.

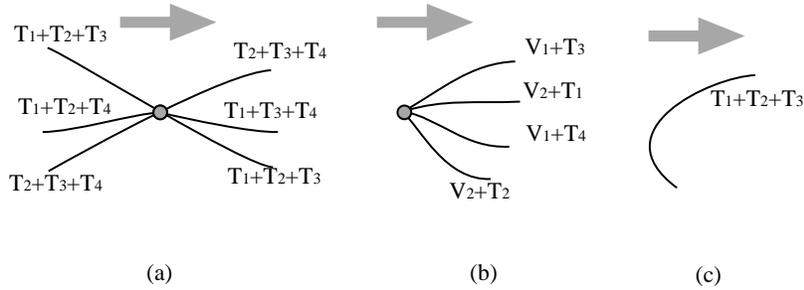


Fig. 19. Different sweep-events represented in line space. The $T + T + T + T$ event (a) is regular, but the $V + V$ event (b) has to be preprocessed as well as the null derivative with respect to φ of the $T + T + T$ events (c).

5.4 Irregular 0-faces and non monotonic 1-faces

Unfortunately, not all the 0-faces are regular sweep-events. A 1-face of the complex can appear during the sweep without a 0-face event. This is obviously the case for $T + T$ events since they can be adjacent to no 0-face, but this can also be the case for $T + T + T$ events. Consider the associated line set; It is not necessarily monotonic with respect to φ (see Fig. 19(c)).

These events thus have to be preprocessed by considering all the VV pairs and all the Object-Object- V triples. This determines a *potential* 0-face, which is a line with coordinates $(\theta_0, \varphi_0, u_0, v_0)$. φ_0 is used to insert it in a sweep event queue. Note that this is just a *potential* 0-face which then has to be tested for occlusion when the corresponding φ_0 is swept. If an object lies between the generators, the 0-face has to be discarded.

For an efficient search, we order the slices of $T + T$ faces using the v coordinate. For a given v there are at most four bitangent segments in a φ -slice. This is true because a φv -slice corresponds to a 2D visibility complex, and because two convex objects in the plane have four common tangent lines. Moreover, since a 2D object lies entirely on the positive or negative half-plane defined by such a line, we can classify these lines according to the side of each of the objects. This defines four categories of bitangents. For each category, there is at most one bitangent for each value of v in a φ -slice.

We thus use a search tree on the $T + T$ faces of the current φ -slice sorted by their gener-

ators, their category among the 4 possible categories and their v values.

5.5 Complexity of the algorithm

THEOREM 2. *The 3D visibility complex can be built in time $O((k + n^3) \log n)$ where n is the complexity of the scene, and k the number of 0-faces of the complex.*

During the initial v sweep, each view computation requires time $O(v \log n)$ where v is the size of the view. A view corresponds to the number of 3-faces of the 3d visibility complex adjacent to the appearing/disappearing 2 faces. The total cost is thus bounded by $O(k \log n)$. Each tritangent event requires time $O(\log n)$. Here again the cost is bounded by $O(k \log n)$.

During the ϕ sweep, each regular event requires $O(\log n)$ to maintain the priority queue.

The preprocessing of the non-monotonic 1-faces requires the enumeration of all the triplets of objects and the insertion of the computed faces in the priority queue, it is therefore $O(n^3 \log n)$.

6. INTERPRETATION OF VISIBILITY PROPERTIES AND RELATED WORK

In this section, we interpret many classical visibility properties and problems in the framework of the visibility complex. We also compare our approach to related work. Unfortunately, interpreting a visibility query does not always translate trivially to an efficient algorithm. Because the Visibility Complex provides a natural framework for visibility, its computational expense can match the difficulty of the visibility problem itself.

6.1 View computation

View computation is a good example of such difficulties. In what follows, we show that views can be easily interpreted in segment space, but no trivial extraction algorithm can be proposed.

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 2-dimensional subset of segment space. In our parameterization, u and v can be expressed with $\sin(\theta)$ and $\sin(\phi)$. The view can be expressed as the intersection of the visibility complex with the 2-manifold of segments going through the viewpoint. Each face intersected corresponds to an object seen. An intersection with a tangency volume corresponds to a silhouette in the image. The ray-tracing algorithm is equivalent to sampling such a manifold.

In Fig. 20, the set of 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 silhouette of the objects, such as D_1, D_2, D_3, D_4 and D_5 . However, all the intersections do not necessarily correspond to a silhouette since the objects are not transparent, and points such as D' must not be taken into account. Consider the ϕ -slice $\phi = 0$ and the slice V_0 of the lines going through V with $\phi = 0$. Fig. 21 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 (which is the set of segments which have both extremities at infinity). At D_1 , we leave face F and have to choose 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.

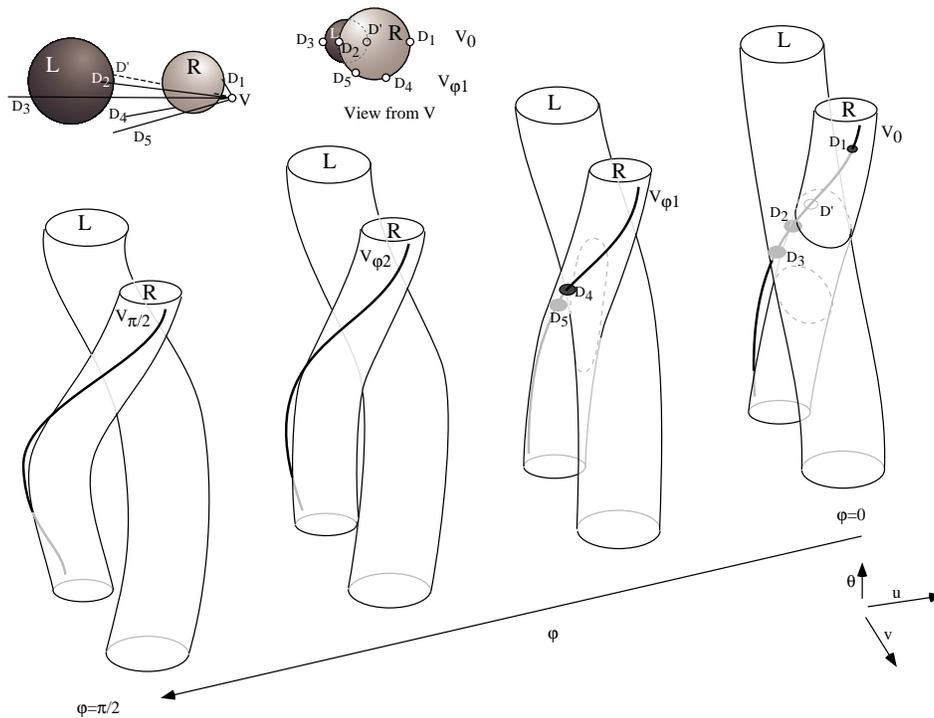


Fig. 20. The View around a point is the intersection of the visibility complex and the surface described by the set of segments going through this point.

The recent technique of *multiple center of projection images* by Rademacher and Bishop [1998] in fact corresponds to the intersection of the complex with other 2-manifolds.

For a walkthrough, 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 by Coorg and Teller [1996, 1997] where conservative visibility events are lazily computed.

6.2 Umbra and penumbra

The 1-faces inside the tangency volume of a light source correspond to the *discontinuity surfaces* of the discontinuity meshing methods, that is, the limits of umbra and penumbra. The visibility complex encodes all the events to compute a discontinuity mesh where all the objects are considered as sources.

In section 7.3 we discuss an efficient method based on the visibility complex concepts which uses visibility to guide the meshing and the refinement of a hierarchical lighting simulation, resulting in high quality images [Durand et al. 1999].

6.3 Form-factors

The form factor F_{ij} used in radiosity algorithms [Sillion and Puech 1994] 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 [Santaló 1976;

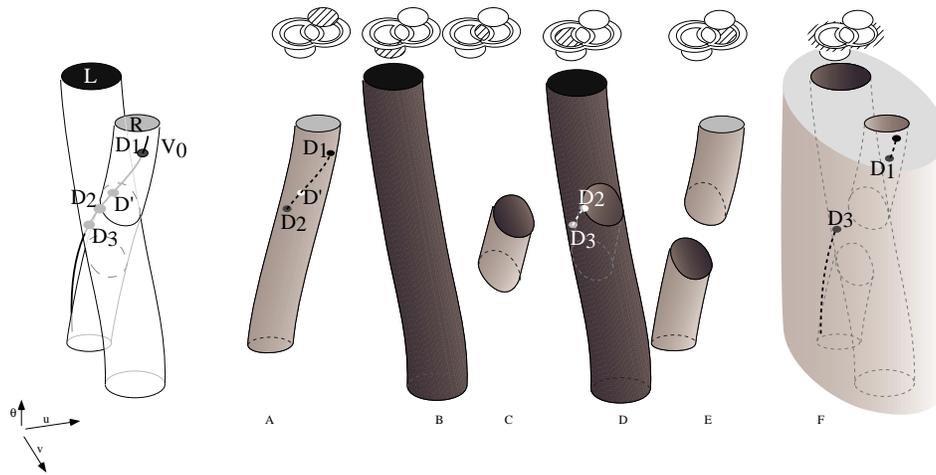


Fig. 21. Traversal of the ϕ -slice $\phi = 0$ of the complex to compute the view around point V .

Sbert 1993]). It correspond to the measure of the 4-face of the Visibility Complex grouping segments that touch i and j , divided by the measure of the interior of the tangency volume of i . An equivalent interpretation has been proposed and used for radiosity computation in 2D [Orti et al. 1996; Durand et al. 1996].

6.4 Computer vision and the aspect graph

The aspect graph is a powerful *viewer-centered* data-structure which encodes all possible views of an object [Koenderink and vanDoorn 1976; Koenderink and vanDoorn 1979; Koenderink 1987; Plantinga and Dyer 1990; Gigus and Malik 1990; Petitjean et al. 1992; Eggert et al. 1992]. It is based on the notion of visual events, which partition the space of all possible viewpoints into cells where views are qualitatively equivalent.

The 1-faces of the visibility complex correspond to the visual events of the aspect graph. 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.

This explains its larger size with respect to the complex. Consider the polygonal case. There can be $O(n^3)$ triples of objects generating tritangents. In the visibility complex, this induces a size $O(n^4)$ because these events are “split” only by $T + T + T + T$ and $T + +T + T$ vertices. In the case of the aspect graph for orthographic projection, the construction of the 2D arrangement on the sphere S^2 of directions can lead to as many as $O(n^6)$ intersections between pairs of events. In the perspective case where the viewing space is \mathbb{R}^3 , the size is $O(n^9)$.

The visibility complex in fact implicitly contains the information encoded in the aspect graph. The aspect graph can be built by projecting the 1-faces of the Visibility Complex into the relevant viewpoint space and then computing their arrangement.

Visual events are also involved in the computation of the *visual hull* [Laurentini 1994], or in the planning of visibility based pursuit motions in robotics [Guibas et al. 1999].

6.5 The asp

The *asp* data structure was developed as an intermediate data-structure to build the aspect graph [Plantinga and Dyer 1990]. Similar to our approach, it partitions line or ray space. Two different definitions are considered depending on whether orthographic or perspective projection is considered. In the orthographic case oriented lines are considered and the *asp* is a 4D cellular decomposition of line-space according to the first object intersected (starting from infinity). Information about visibility from inside the convex-hull of the object is not encoded.

If perspective projection is used, the *asp* is a 5D cellular decomposition of ray-space according to the first object a ray intersects. The same information is encoded as in the visibility complex, but the information is redundant because many collinear rays “see” the same object.

The visibility complex can be seen as offering the advantages of the two versions of the *asp*. It encodes the optimal amount of information, by considering maximal free segments.

6.6 Plenoptic function, lumigraph and light-field

The *plenoptic function* is a function in 5D ray-space which describes the light traveling at a point in a given direction [Adelson and Bergen 1991]. If light traveling in free-space is assumed invariant and if the observer is constrained outside the scene, it can be simplified into a function in 4D line space called *lumigraph* [Gortler et al. 1996] or *light field* [Levoy and Hanrahan 1996]. They are in a sense very similar to the *asp* concept, except that color information is encoded and that a discrete sampling is used instead of an analytic subdivision. However, the underlying multidimensional ray-spaces are the same.

Langer and Zucker [1997] use similar topological concepts as ours to describe the light field of a scene in a shape-from-shading context. As for the visibility complex, they note that the manifold of rays is 4 dimensional with some branchings. They do not however propose the analysis of the changes in visibility, nor the adjacencies between classes of rays.

6.7 Plücker space

Plücker parameterization represents lines in a five dimensional projective space where hyperplanes are used to characterize line-line incidence. It has been extensively used to compute visibility in polygonal scenes.

Note that the fifth dimension in Plücker space does not correspond to the pseudo-dimension which we have introduced in section 2.3 for visualization purposes. Indeed, not all points in Plücker space correspond to real lines: only those lying on the so-called *Plücker hypersurface*. The fifth dimension alleviates the problem of singularities which are always present in a parameterization of lines in \mathbb{R}^4 (as is also the case for any 2D parameterization of the S^2 sphere).

Arrangements in Plücker space [Pellegrini 1993; Teller 1992b; Teller 1992a] are very similar to the dual arrangement that we have presented. An arrangement of hyperplanes is computed in 5D and intersected with the Plücker hypersurface to obtain a 4D structure. This latter structure is in fact exactly the dual arrangement expressed in a different parameterization. This illustrates our claim in the introduction: the concepts described in this paper do not rely on a particular parameterization of lines.

However, approaches in Plücker space have considered only equivalents of the dual arrangement. They consider only line visibility, occlusion has not really been taken into

account.

6.8 The 2D visibility complex

Our definition of the 3D visibility complex as the partition of maximal free segment is the same as the 2D version introduced by Pocchiola and Vegter [1996b, 1996a]. We now compare the 2D and the 3D data-structures, which illustrates the large differences between 2D and 3D visibility.

First of all, the space of lines in 3D is 4-dimensional. The increase in the dimensionality of the problems is 2 and not only 1. Similarly, the theoretical complexity is $O(n^4)$ instead of $O(n^2)$ in 2D. In 2D, only tangent and bitangent segments are considered, while in 3D the vertices of the complex are segments tangent to four object or T++T+T segments.

The property which most clearly explains why 3D visibility is much harder than in 2D, is *separability*. In 2D, a line separates the plane into two half-planes. No such property holds in 3D because lines are no longer hyperplanes.

The consequence is that some convexity or monotonicity properties which hold in 2D do not hold in 3D. This is especially the case of the faces for the visibility complex. In 2D, depending on the parameterization of lines, the faces are at least monotonic with respect to direction. This has a number of useful consequences, including the possibility to perform efficient optimal sweeps or walks along the complex.

As aforementioned, the 3D visibility complex is not formally a cell-complex [Hatcher 2001], while the 2D visibility complex is a 2-dimensional cell-complex. This means that the faces of the 3D visibility complex can have a non 0-genus (they can have holes, like a torus).

In addition to their complex topology, we have seen that some 0-faces are irregular for the sweep of the 3D complex. The faces cannot be made monotonic with respect to one parameter. This makes our algorithm far from optimal. This also explains the difficulty in designing an efficient view extraction algorithm.

7. CONCLUSIONS

7.1 Summary

We have presented a new approach for visibility computation and described a powerful data-structure, the *3D visibility complex*, which encapsulates all the visibility information in a 3D scene.

Its size k is $\Omega(n)$ and $O(n^4)$. Using a probabilistic approach, we have shown that in scenes with reasonable assumptions, the number of tri-tangency events is $O(n^{2.33})$ instead of $O(n^3)$.

We have also sketched the definition of the *Temporal Visibility Complex* which encodes all the visibility of a time-varying 3D scene.

We have presented an output-sensitive algorithm to build the structure in time $O((n^3 + k) \log n)$.

Finally, we have interpreted many visibility queries and data structures and shown that visibility properties can be naturally expressed in this new framework.

7.2 Discussion

The Visibility complex is a useful framework because it permits an efficient encoding of visibility information. We believe it is a valuable tool for the interpretation of visibility

problems. By providing different insights to these questions, we hope it will permit a better understanding and the development of new methods, as we discuss below.

Its direct practical interest is however more questionable. The implementation of a 4D cellular subdivision is an intricate task, and traversing the adjacencies of such a complex is not straightforward.

Moreover, sweep constructions, though efficient, are prone to numerical precision and robustness problems. If an event of the sweep is not correctly handled, lost coherence will make the entire algorithm fail.

7.3 Applications of the study

This theoretical work has led to practical implementations, described in companion papers.

The Visibility Skeleton. To avoid the treatment of a four dimensional cell-complex, we have simplified the visibility complex of polygonal objects into a graph in line-space, the *visibility skeleton* [Durand et al. 1997b]. Only faces of dimension 0 and 1 are kept, defining the nodes and arcs of the graph.

The construction algorithm we implemented for the visibility skeleton is different from the sweep described in Section 5. It is topological and local: compute the nodes of the graph (corresponding to the vertices of the complex), and deduce the arcs using a catalogue of adjacencies similar to the one given in appendix. Thanks to this algorithm, we have been able to show that it is possible to implement a global visibility structure which is generic. Once the structure is built, very efficient queries are possible: *e.g.*, milliseconds for the retrieval of exact umbra and penumbra boundaries (see Fig. 22).

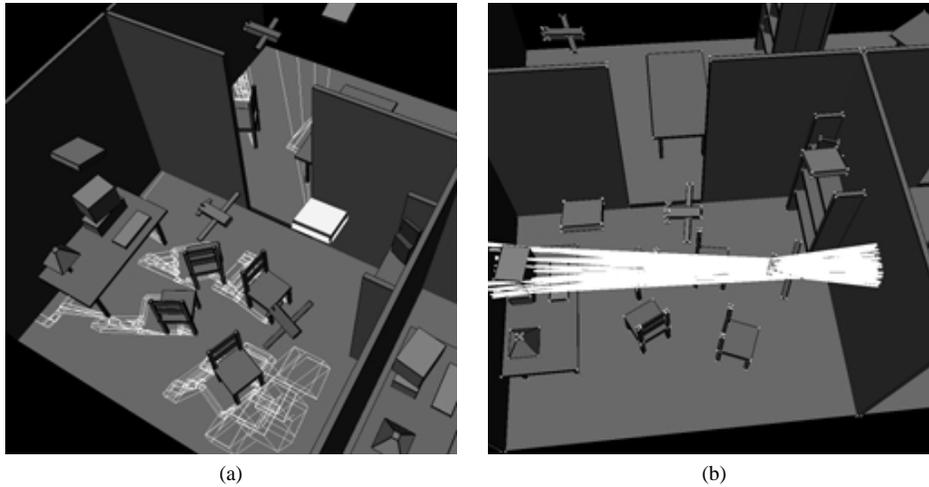


Fig. 22. (a) The discontinuity mesh with respect to the right source. (b) Limits of the occlusions caused by a part of the plane between the computer screen and the right wall.

Visibility driven hierarchical radiosity. The application of the visibility skeleton to lighting simulation [Durand et al. 1999] permits higher quality images together with improved computation time compared to previous algorithms. Previous methods compute visibility

by sampling, which is costly and inaccurate, while the visibility skeleton affords an efficient exact computation of the amount of light leaving one polygon arriving at a point (area-to-point form factor).

It also allows us to subdivide the mesh used to represent the lighting function along the limits of umbra and penumbra, with respect to any polygon, while previous methods were limited to a small number of primary light sources. In particular, we handle the discontinuities of the lighting function caused by the limits of visibility of indirect lighting.

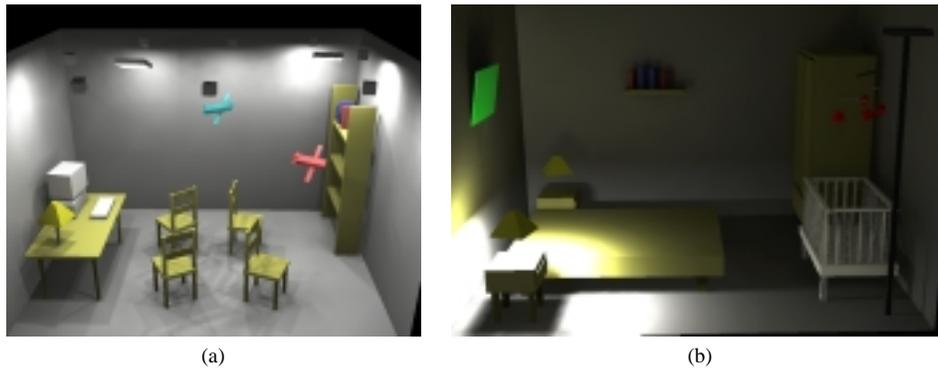


Fig. 23. Images computed using our new hierarchical radiosity algorithm based on the Visibility Skeleton and hierarchical triangulations. (a) A scene with multiple sources. The skeleton construction took 2min 23s and the lighting simulation 8min. (b) A scene mainly lit by indirect light. The skeleton construction took 4min 12s and the lighting simulation 6min 58s. Note the shadows caused by indirect illumination, cast by the books on the back wall.

7.4 Future work

The faces of the Visibility Complex directly translate the intricate nature of 3D visibility. They have no convexity nor monotonicity property, as illustrated by the irregular events of our sweep algorithm. Enforcing such properties could be achieved by appropriately subdividing the faces of the complex. The cylindrical partitions of Mulmuley [1991] or the cylindrical algebraic decomposition of Collins [1975] are two examples of such approaches. Unfortunately, this would increase the space complexity of the structure.

The recent breakthrough by Koltun on vertical decomposition of arrangements of algebraic surfaces in 4D [Koltun 2001] could be of great help in this context. He showed that the complexity of such a vertical decomposition is $O(n^{4+\epsilon})$ for any $\epsilon > 0$. A vertical decomposition then permits logarithmic point-location queries, which translates into efficient ray-shooting queries. The application to other visibility problems should be studied.

In particular, the definition of an efficient view computation retrieval using the complex should be explored. We have not obtained such results because of the aforementioned problems of non-monotonicity. View maintenance raises the same issues.

The design of an efficient view-computation algorithm is highly related to the problem of an efficient construction of the complex. For polygonal scenes, another approach would sweep the scene by planes rotating around the edges of the polygons while maintaining a

k	k-face	(k+1)-faces	(k+2)-faces	(k+3)-faces	(k+4)-faces
3	T	3 4-faces			
2	$T_1 + T_2$	2 T_1 2 T_2	6 4-faces		
1	$T_1 + T_2 + T_3$	2 $T_1 + T_2$ 1 $T_1 + T_3$ 2 $T_2 + T_3$	3 T_1 4 T_2 3 T_3	10-faces	
	$T_1 + T_2$	1 $T_1 + T_2$	2 T_1 2 T_2	6 4-faces	
0	$T_1 + T_2 + T_3 + T_4$	2 $T_1 + T_2 + T_3$ 1 $T_1 + T_2 + T_4$ 1 $T_1 + T_3 + T_4$ 2 $T_2 + T_3 + T_4$ 2 $T_2 + T_3 + T_4$ 2 $T_2 + T_3 + T_4$	3 $T_1 + T_2$ 4 $T_2 + T_3$ 3 $T_3 + T_4$ 2 $T_1 + T_3$ 2 $T_2 + T_4$ 1 $T_1 + T_4$	4 T_1 6 T_2 6 T_3 4 T_4	15 4-faces
	$T_1 + T_2 + T_3$	2 $T_1 + T_2$ 1 $T_1 + T_2 + T_3$	2 $T_1 + T_2$ 2 $T_2 + T_3$	3 T_1 4 T_2 3 T_3	10 4-faces

Table 3. Adjacencies of the faces of the visibility complex of polygons and smooth objects.

2D visibility complex. The sweeps have to be synchronized, and only the portion of the scene visible from each vertex has to be maintained in the corresponding 2D complex. The definition of a topological order to sweep the 0-faces of the complex is yet another issue required to reach an optimal output-sensitive algorithm.

In the case of algebraic scenes, the work of Petitjean [1995, 1996] on enumerative geometry could be used to obtain tight bounds on the size of the complex depending on the degree of the objects.

ACKNOWLEDGMENTS

We are indebted to Sylvain Petitjean and Sylvain Lazard who pointed out initial errors in the derivation of the probabilistic bound. Frédéric Cazals's probabilistic skills were also of great help, in particular for indicating the relevant formula in the Santaló book. Finally, thanks to the anonymous reviewers, whose comments greatly helped us improve this paper.

APPENDIX

A. COMPLEXITY PROOF AND CATALOGUE OF ADJACENCIES

The number of $(k+1)$ -faces adjacent to a k -face is bounded. For example a 1-face $T_1 + T_2 + T_3$ is adjacent to five 2-faces: two faces $T_1 + T_2$, $T_1 + T_3$ and two $T_2 + T_3$. There are two different $T_1 + T_2$ faces because one extremity of the segments can either lie on the object tangent at T_3 or not. See Fig. 24). The same applies to $T_2 + T_3$. The other adjacencies are summarized in table 3.

Each 4-face is adjacent to at least one 3-face, a 3-face to at least one 2-face, and a 2-face to at least one 1-face. However, a 1-face may not be adjacent to a 0-face. We sketch the demonstration. For a given k -face F of the complex (composed of segments), we consider the associated set of lines S with the same tangency properties. For example, a $T_1 + T_2$ face is associated with the set of all lines tangent to T_1 and T_2 . This set of lines contains a subset S' that has one more incidence: it is tangent to an additional object or some of the lines tangent to two objects belongs to their common tangent plane. Consider now

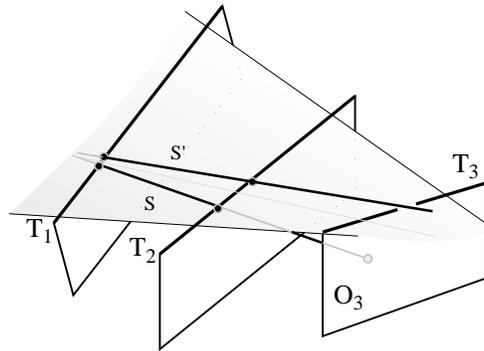


Fig. 24. Adjacencies at a $T_1 + T_2 + T_3$ face. Two different $T_1 + T_2$ sets of bitangents are adjacent: one with extremity O_3 (such as S) and one lying above T_3 (such as S').

a continuous path from the line associated with a segment s of F to one of S' , and the corresponding continuous path over the segments². If all the segments of this path have the same extremities, F is adjacent to the face with one additional incidence associated with S' , otherwise when the extremity changes there is a tangency local event and thus a lower-dimensional face.

We illustrate this by the example of a T 3-face and a maximal free segment s of this face (Fig. 25). We consider the case where s does not have an extremity lying at infinity. Consider a plane P containing s . We define l to be the first line in a polar order tangent to both the object at T and the object O_1 at one of the extremity of s (see Fig. 25). Consider a continuous path $l(t)$ in the line space from the line defined by s to l . A corresponding continuous path can be defined over the segments by considering for each $l(t)$ the segment $s(t)$ tangent to the object at T . If $s(t)$ always has the two same objects at its extremities, it belongs to the same T 3-face as s and this face is thus adjacent to a $T + T$ 2-face at segment s' (Fig. 25(a)), otherwise when the object at one extremity changes the segment is tangent to two objects and there is also a $T + T$ 2-face (Fig. 25(b)). The other cases can be demonstrated in a similar manner.

Note that a 1-face is not necessarily adjacent to a 0-face (the scene in Figure. 13(a) does not have a 0-face). This is why the 3D Visibility Complex is not a CW complex [Hatcher 2001].

REFERENCES

- ADELSON, E. H. AND BERGEN, J. R. 1991. The plenoptic function and the elements of early vision. *M. Landy and J. A. Movshon, (eds) Computational Models of Visual Processing*, 3–20.
- BOISSONNAT, J.-D. AND YVINEC, M. 1998. *Algorithmic geometry*. Cambridge University Press.
- CLARKSON, K. L. 1987. New applications of random sampling to computational geometry. *Discrete and Computational Geometry 2*, 195–222.
- COLLINS, G. 1975. Quantifier elimination for real closed fields by cylindrical algebraic decomposition. In *Proceedings Second GI Conference on Automata Theory and Formal Languages*, Volume 33 of *Lecture Notes in Computer Science* (Berlin, 1975), pp. 134–183. Springer-Verlag.

²the notion of corresponding path in segment space is not always properly defined because of the branchings due to tangency. In our case however, if tangency occurs, one incidence is added and we have our adjacent face

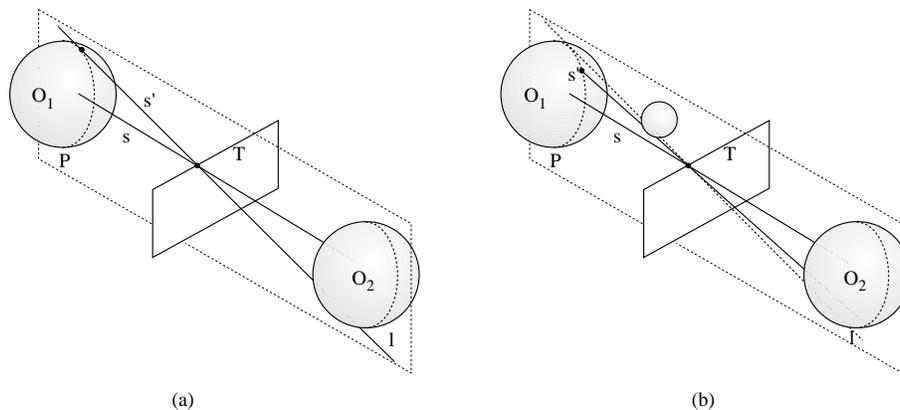


Fig. 25. Construction of a segment of a $T+T$ 2-face adjacent to a T 3-face.

- COORG, S. AND TELLER, S. J. 1996. Temporally coherent conservative visibility. In *Proceedings of the Twelfth Annual Symposium On Computational Geometry (ISG '96)* (New York, May 1996), pp. 78–87. ACM Press. <http://graphics.lcs.mit.edu/~satyan/pubs.html>.
- COORG, S. AND TELLER, S. J. 1997. Real-time occlusion culling for models with large occluders (color plate S. 189). In *Proceedings of the Symposium on Interactive 3D Graphics* (New York, April 27–30 1997), pp. 83–90. ACM Press. <http://graphics.lcs.mit.edu/~satyan/pubs.html>.
- DE BERG, M., KATZ, M. J., VAN DER STAPPEN, A. F., AND VLEUGELS, J. 1997. Realistic input models for geometric algorithms. In *Proceedings of the 13th International Annual Symposium on Computational Geometry (SCG-97)* (New York, June 4–6 1997), pp. 294–303. ACM Press.
- DRETTAKIS, G. AND FIUME, E. 1994. A fast shadow algorithm for area light sources using backprojection. *Computer Graphics* 28, Annual Conference Series (July), 223–230. <http://www-imagis.imag.fr/~George.Drettakis/pub.html>.
- DURAND, F. 1999. *3D Visibility, analysis and applications*. Ph. D. thesis, U. Joseph Fourier, Grenoble. <http://www-imagis.imag.fr>.
- DURAND, F., DRETTAKIS, G., AND PUECH, C. 1996. The 3D visibility complex: A new approach to the problems of accurate visibility. In X. PUEYO AND P. SCHRÖDER Eds., *Eurographics Rendering Workshop 1996* (New York City, NY, June 1996), pp. 245–256. Eurographics: Springer Wein. ISBN 3-211-82883-4, <http://www-imagis.imag.fr/Publications/>.
- DURAND, F., DRETTAKIS, G., AND PUECH, C. 1997a. The 3d visibility complex: a unified data-structure for global visibility of scenes of polygons and smooth objects. In *Canadian Conference on Computational Geometry* (Aug. 1997). <http://www-imagis.imag.fr/~Fredo.Durand>.
- DURAND, F., DRETTAKIS, G., AND PUECH, C. 1997b. The visibility skeleton: a powerful and efficient multi-purpose global visibility tool. In *SIGGRAPH 97 Conference Proceedings*, Annual Conference Series (Aug. 1997), pp. 89–100. ACM SIGGRAPH: Addison Wesley.
- DURAND, F., DRETTAKIS, G., AND PUECH, C. 1999. Fast and accurate hierarchical radiosity using global visibility. *ACM Transactions on Graphics* 18, 2 (April), 128–170.
- DURAND, F., ORTI, R., RIVIÈRE, S., AND PUECH, C. 1996. Radiosity in flatland made visibly simple. In *video 12th Annu. ACM Sympos. Comput. Geom.* (1996).
- DURAND, F. AND PUECH, C. 1995. The visibility complex made visibly simple. In *video 11th Annu. ACM Sympos. Comput. Geom.* (1995).
- EGGERT, D., BOWYER, K., AND DYER, C. R. 1992. Aspect graphs: State-of-the-art and applications in digital photogrammetry. In *Proceedings of the 17th Congress of the International Society for Photogrammetry and Remote Sensing, Part B5* (1992), pp. 633–645.
- GIGUS, Z. AND MALIK, J. 1990. Computing the aspect graph for the line drawings of polyhedral objects. *IEEE Trans. on Pat. Matching & Mach. Intelligence* 12, 2 (Feb.), 113–122.

- GORTLER, S. J., GRZESZCZUK, R., SZELISKI, R., AND COHEN, M. F. 1996. The lumigraph. In H. RUSHMEIER Ed., *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series (Aug. 1996), pp. 43–54. ACM SIGGRAPH: Addison Wesley. held in New Orleans, Louisiana, 04-09 August 1996, <http://hillbilly.deas.harvard.edu/~sjg/>.
- GUIBAS, L. J., LATOMBE, J.-C., LAVALLE, S. M., LIN, D., AND MOTWANI, R. 1999. Visibility-based pursuit-evasion in a polygonal environment. *International Journal of Computational Geometry and Applications* 9, 5 (Oct.), 471–494.
- HATCHER, A. 2001. *Algebraic Topology*. Cambridge University Press.
- HOLZSCHUCH, N., SILLION, F., AND DRETTAKIS, G. 1994. An Efficient Progressive Refinement Strategy for Hierarchical Radiosity. In *Fifth Eurographics Workshop on Rendering* (Darmstadt, Germany, June 1994), pp. 343–357. <http://www.loria.fr/~holzschu>.
- KERGOSIEN, Y. L. 1981. La famille des projections orthogonales d'une surface et ses singularités. *C.R. Acad. Sc. Paris* 292, 929–932.
- KOENDERINK, J. J. 1987. An internal representation for solid shape based on the topological properties of the apparent contour. In W. RICHARDS AND S. ULLMAN Eds., *Image Understanding 1985–86* (Norwood, NJ, 1987), pp. 257–285. Ablex.
- KOENDERINK, J. J. AND VANDOORN, A. J. 1976. The singularities of the visual mapping. *BioCyber* 24, 1, 51–59.
- KOENDERINK, J. J. AND VANDOORN, A. J. 1979. The internal representation of solid shape with respect to vision. *BioCyber* 32, 211–216.
- KOLTUN, V. 2001. Almost tight upper bounds for vertical decompositions in four dimensions. In *Proceedings 42nd Symposium on Foundations of Computer Science* (2001). <http://www.math.tau.ac.il/~vladlen/>.
- LANGER, M. S. AND ZUCKER, S. W. 1997. Casting light on illumination: A computational model and dimensional analysis of sources. *Computer Vision and Image Understanding: CVIU* 65, 2 (Feb.), 322–335.
- LAURENTINI, A. 1994. The visual hull concept for silhouette-based image understanding. *T-PAMI* 16, 150–162.
- LEVOY, M. AND HANRAHAN, P. 1996. Light field rendering. In H. RUSHMEIER Ed., *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series (Aug. 1996), pp. 31–42. ACM SIGGRAPH: Addison Wesley. held in New Orleans, Louisiana, 04-09 August 1996, <http://www-graphics.stanford.edu/papers/light/>.
- MATHAI, A. M. 1999. *An Introduction to Geometrical Probability*. Gordon and Breach Science Publishers.
- MULMULEY, K. 1991. Hidden surface removal with respect to a moving view point (extended abstract). In *STOC: ACM Symposium on Theory of Computing (STOC)* (1991), pp. 512–522.
- ORTI, R., RIVIÈRE, S., DURAND, F., AND PUECH, C. 1996. Radiosity for dynamic scenes in flatland with the visibility complex. In J. ROSSIGNAC AND F. SILLION Eds., *Computer Graphics Forum (Proc. of Eurographics '96)*, Volume 16(3) (Poitiers, France, Sept. 1996), pp. 237–249.
- PELLEGRINI, M. 1990. Stabbing and ray shooting in 3-dimensional space. In *Proc. 6th Annu. ACM Sympos. Comput. Geom.* (1990), pp. 177–186.
- PELLEGRINI, M. 1993. Ray shooting on triangles in 3-space. *Algorithmica* 9, 471–494.
- PELLEGRINI, M. 1994. On lines missing polyhedral sets in 3-space. *Discrete & Computational Geometry* 12, 203–221.
- PELLEGRINI, M. AND SHOR, P. W. 1992. Finding stabbing lines in 3-space. *GEOMETRY: Discrete & Computational Geometry* 8.
- PETITJEAN, S. 1992. Computing exact aspect graphs of smooth objects bounded by smooth algebraic surfaces. Master's thesis, University of Illinois, Urbana-Champaign, IL. available as technical report UIUC-BI-AI-RCV-92-04.
- PETITJEAN, S. 1995. The number of views of piecewise-smooth algebraic objects. *Proceedings of Proceedings of the Symposium on Theoretical Aspects of Computer Science, Lecture Notes in Computer Science* 900, 571–582.
- PETITJEAN, S. 1996. The enumerative geometry of projective algebraic surfaces and the complexity of aspect graphs. *International Journal of Computer Vision* 19, 3, 1–27.

- PETITJEAN, S., PONCE, J., AND KRIEGMAN, D. J. 1992. Computing exact aspect graphs of curved objects: Algebraic surfaces. *International of Computer Vision* 9, 3, 231–255.
- PLANTINGA, H. AND DYER, C. R. 1990. Visibility, occlusion, and the aspect graph. *International Journal of Computer Vision* 5, 2, 137–160.
- POCCHIOLA, M. AND VEGTER, G. 1996a. Topologically sweeping visibility complexes via pseudotriangulations. *Discrete & Computational Geometry* 16, 419–453.
- POCCHIOLA, M. AND VEGTER, G. 1996b. The visibility complex. *International Journal of Computational Geometry and Applications*, 279–308. special issue devoted to ACM-SoCG'93.
- RADEMACHER, P. AND BISHOP, G. 1998. Multiple-center-of-projection images. In M. COHEN Ed., *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series (July 1998), pp. 199–206. ACM SIGGRAPH: Addison Wesley. ISBN 0-89791-999-8.
- RIEGER, J. H. 1987. On the classification of views of piecewise smooth objects. *Image and Vision Computing* 5, 91–97.
- RIEGER, J. H. 1990. The geometry of view space of opaque objects bounded by smooth surfaces. *Artificial Intelligence(1-2)* 44, 1–40.
- RIEGER, J. H. 1992. Global bifurcation sets and stable projections of non-singular algebraic surface. *Internal Journal of Computer Vision* 7, 3, 171–194.
- RIEGER, J. H. 1993. Computing view graphs of algebraic surfaces. *Journal of Symbolic Computation* 16, 3 (Sept.), 259–272.
- RIVIÈRE, S. 1997a. *Calculs de visibilité dans un environnement polygonal 2D*. Ph. D. thesis, Université Joseph Fourier (Grenoble). PhD Thesis.
- RIVIÈRE, S. 1997b. Dynamic visibility in polygonal scenes with the visibility complex. In *Communication 13th Annual ACM Symposium on Computational Geometry* (1997), pp. 421–423.
- SANTALÓ, L. A. 1976. *Integral Geometry and Geometric Probability*. Addison-Wesley, Reading, MA.
- SBERT, M. 1993. An integral geometry based method for fast form-factor computation. *Computer Graphics Forum* 12, 3, C409–C420.
- SILLION, F. AND PUECH, C. 1994. *Radiosity and Global Illumination*. Morgan Kaufmann, San Francisco, CA.
- TELLER, S. J. 1992a. Computing the antipenumbra of an area light source. In E. E. CATMULL Ed., *Computer Graphics (SIGGRAPH '92 Proceedings)*, Volume 26-2 (July 1992), pp. 139–148. <http://graphics.lcs.mit.edu/~seth/pubs/pubs.html>.
- TELLER, S. J. 1992b. *Visibility Computations in Densely Occluded Polyhedral Environments*. Ph. D. thesis, CS Division, UC Berkeley. Tech. Report UCB/CSD-92-708, <http://graphics.lcs.mit.edu/~seth/pubs/pubs.html>.
- THOM, R. 1972. *Structural Stability and Morphogenesis*. Benjamin, New-York.