CONSTRUCTING THE VISIBILITY COMPLEX OF POLYTOPES IN 3-SPACE

SAMUEL HORNUS AND MAXIME WOLFF

Abstract. We prove the connexity of the boundary of each 4-dimensional "cell "of the 3D visibility complex of disjoint convex polyhedra. Using this property, we outline a simple algorithm for the construction of the visibility complex of k disjoint, convex polyhedra in time $O(n^2k^2\log n)$ where n is the total complexity of the polyhedra.

1. Introduction

1.1. **The 3D visibility complex.** Let $\mathcal{O}$ be a set of k pairwise disjoint convex polyhedra, totalling up n edges. We call these convex polyhedra the *occluders* or the *objects*. Together, they define the *free space* $\mathcal{F}$ as the complement of the union of their interior:

$$\mathcal{F} = \mathbb{R}^3 \setminus \bigcup_{O\in\mathcal{O}} \mathring{O}.$$

A *free segment* is a closed oriented line segment inside the free space. It is *maximal* if no other free segment contains it properly. We let $\mathcal{S}$ denote the set of maximal free segments. An endpoint of a maximal free segment either lies on the boundary of an occluder, or extends to infinity. The starting point of a maximal free segment s is called s^0 and its ending point s^1. Letting $\widetilde{\mathcal{O}} = \mathcal{O}\cup\{\infty\}$, each maximal free segment s can be associated with a pair $\mathfrak{V}(s) = (B, F) \in \widetilde{\mathcal{O}}^2$, such that s^0 is on the boundary of occluder B and s^1 is on the boundary of F. if s^0 (resp. s^1) is at infinity, then B (resp. F) is set to ∞. B is called the *back blocker* of the free segment s while F is called its *front blocker*. See figure 1. We refer to $\mathfrak{V} : \mathcal{S} \to \widetilde{\mathcal{O}}^2$ as the *visibility function*. We also extend the visibility function to any subset of $\mathcal{S}$ on which $\mathfrak{V}$ is constant (e.g., to the "cells" of the visibility complex that we define below). Thanks to the requirement that occluders be disjoint, all maximal free segments have positive or infinite length, and thus have a well defined orientation.

In order to endow $\mathcal{S}$ with a topology, we define $\overline{\mathcal{S}} = \mathcal{F} \times \mathbb{S}^2$, where $\mathbb{S}^2$ is the unit sphere in $\mathbb{R}^3$ interpreted as the space of orientations. Each element $(p, u) \in \overline{\mathcal{S}}$ defines a unique oriented maximal free segment in $\mathcal{S}$; however this function is not one-to-one. We then obtain the topological space $\mathcal{S}$ as the quotient of $\overline{\mathcal{S}}$ by the equivalence relation $\sim$ on $\overline{\mathcal{S}}$ defined as

$$(p, u) \sim (q, v) \iff \begin{cases} & u = v = \omega_{p,q} \\ \text{and} & [p, q] \text{ is a free segment,} \end{cases}$$

where $\omega_{p,q} \in \mathbb{S}^2$ is the orientation of segment $[p, q]$. As such, defining $\mathcal{S} = \overline{\mathcal{S}}/\sim$ endows the set of maximal free segments with the quotient topology [Hat01].

The *visibility complex* $\mathcal{VC}$ of $\mathcal{O}$ is a partition of $\mathcal{S}$ into maximally connected components of dimensions 0 to 4. Although a component of this partition is not necessarily homeomorphic to a ball, we *improperly* call it a *cell* in this paper. The 4-dimensional cells (or 4-cells) of $\mathcal{VC}$ are the interior of the connected components of $\mathfrak{V}^{-1}((B, F))$ for all $(B, F) \in \widetilde{\mathcal{O}}^2$. A 4-cell in $\mathfrak{V}^{-1}((B, F))$ is said to be *of type* (B, F) or to have *visibility* (B, F). For $i \in \{0, 1, 2, 3\}$, the i-cells of $\mathcal{VC}$ are the maximally connected set of free segments in $\mathfrak{V}^{-1}((B, F))$ tangent to the same $4 - i$ occluders, for all $(B, F) \in \widetilde{\mathcal{O}}^2$.

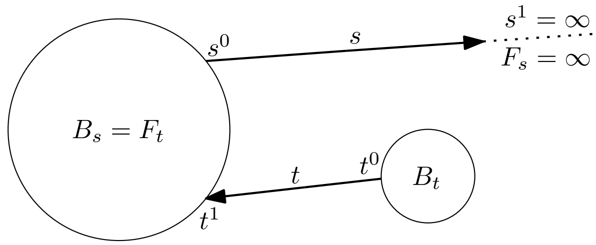

FIGURE 1. Definitions.

For each object $P \in \mathcal{O}$, define $\mathcal{T}_P \subset \mathcal{S}$ as the set of the maximal free segments tangent to P. Then, one can also define the visibility complex $\mathcal{VC}$ of $\mathcal{O}$ as the arrangement in $\mathcal{S}$ induced by the hypersurfaces $\mathcal{T}_P$ for all $P \in \mathcal{O}$.

We define $\mathcal{S}^3 = \cup_{P \in \mathcal{O}} \mathcal{T}_P$ as the set of maximal free segments that are tangent to at least one occluder. $\mathcal{S}^3$ is the union of all i-cells for $i \leq 3$. The i-cells for $i \leq 3$ also induce a decomposition of $\mathcal{S}^3$ into a "quasi-cell" complex $\mathcal{VC}^3$, which we call the 3-skeleton of $\mathcal{VC}$. The 1-skeleton of $\mathcal{VC}$ is defined similarly.

The space $\mathcal{S}$ of maximal free segments is Hausdorff. Furthermore, although the visibility complex itself is not a cell-complex structuring of $\mathcal{S}$, the latter is a cell-complex (or CW-complex). Hausdorff cell-complexes have nice properties that we use in the proof of theorem 1.

1.2. **Contributions.** Our main result is

Theorem 1. *The boundary of a 4-cell of the visibility complex of a finite set of pairwise disjoint convex polyhedra is connected.*

Example: Let O_1 and O_2 be two disjoint balls in $\mathbb{R}^3$ with same radii. Let O_3 be a tiny ball in between O_1 and O_2. Let C be the only 4-cell whose visibility is (O_1, O_2). We invite the reader to check that the boundary of C is indeed connected, and that C is not simply connected, implying that $\mathcal{VC}$ is not a cell-complex.

Our second contribution is the application of this result to a simple algorithm for the construction of the visiblity complex of a set of disjoint convex polyhedra.

1.3. **Related work.** The 3D visibility complex of a set of pairwise disjoint convex objects in 3-space was introduced by Durand *et al.* [DDP02] as an generalization of the planar visibility complex of Pocchiola and Vegter [PV96b]. Contrary to the 3D case, the 2D visibility complex has good topological properties (for example, each cell is a ball) that afford efficient algorithms for its construction. In particular, the visibility complex of n disjoint convex objects [1] in the plane can be constructed in optimal time $O(n \log n + k)$ and optimal space, where k is the size of the complex, with the *Greedy Flip Algorithm* [PV96a]. An implementation of the latter algorithm [Ang02] in the CGAL library [CGA] allowed to study the size of the complex for particular distributions of objects [ELPZ05]. This experimental study shows that the complex has linear size for uniformly distributed unit discs in the plane, thus encouraging an experimental study of the practical size of the complex in 3D. Other encouragements come from the result of Devillers *et al.* [DDE$^+$03] which shows that the average size of the visibility complex of uniformly distributed unit 3-balls in space has linear size. For the general case of k polytopes totalling up n edges, the complex was shown to have complexity $O(n^2k^2)$ in [BDD$^+$06]. A particular subset of the visibility complex, namely the set of *free lines* (lines avoiding all objects) was shown to have simpler $O(n^{3+\varepsilon})$ complexity when the objects are unit 3-balls, by Agarwal *et al.* [AAKS05].

[1] The bitangents to two objects must be computable in constant time.

1.3.1. *Existing algorithms for constructing the 3D visibility complex.* Durand [DDP02] gives an algorithm that maintains the 2D visibility complex in a planar slice of 3-space, as the slice sweeps the whole line space. The algorithm is succintly presented and affords the construction of the visibility complex of polyhedral scenes in time $O((q + n^3) \log n)$ where q is the number of 0-cells in the complex. Goaoc [Goa04] proposes an algorithm for the construction of the visibility complex of semi-algebraic sets in space. It uses a "black box" that performs algebraic computations on the sets of lines piercing an object. The algorithm is interesting due to its simplicity. Howevever, some algebraic operations used in the algorithm are not currently implemented, and temporarily prevents its practical implementation. Brönnimann *et al.* [BDD+06] also show how the 1-skeleton of the visibility complex of convex polyhedra can be constructed in time $O(n^2 k^2 \log n)$. In this paper we use the connexity property proved in section 2 to extend the construction of the 1-skeleton of $\mathcal{VC}$ in [BDD+06, Goa04] to the construction of the whole visibility complex, within the same time bound.

In the next section we give a proof of theorem 1. In section 3 we outline a simple algorithm for the construction of the visibility complex. We conclude the paper in section 4.

2. Proof of theorem 1

Let C be a 4-cell of the visibility complex and denote by ∂C and $C^c = \mathcal{S} \setminus C$ respectively the boundary of C and its complement in $\mathcal{S}$. By definition, C is connected, and our goal is to show that ∂C is connected. Observe that theorem 1 holds in the case $k = 1$ of only one occluder. In the sequel we assume $k > 1$. We begin with three lemmas setting up connectedness properties of the complement C^c of a 4-cell C, the free space $\mathcal{F}$ and the space of maximal free segments $\mathcal{S}$.

Lemma 2. *C^c is path-connected.*

Proof. Let s and $t \in C^c$. Let $(B_s, F_s) = \mathfrak{V}(s)$ and $(B_t, F_t) = \mathfrak{V}(t)$ be the visible sets of s and t respectively, and $(B_C, F_C) = \mathfrak{V}(C)$ the visible set of the 4-cell C. $\mathfrak{V}(C)$ is to be avoided while moving s continuously to t.

Case 1: $B_C = F_C = \infty$. If $B_s = F_s = \infty$ (that is, s has the same visibility as C but lies in a different connected component of $\mathfrak{V}^{-1}(\infty, \infty)$) then we move s randomly until B_s becomes an object in $\mathcal{O}$. We move t similarly so that B_t becomes an object in $\mathcal{O}$.

We number the occluders in $\mathcal{O}$ so that the sequence $O_1 = B_s, O_2, O_3 \ldots, O_j = B_t$ is the ordered sequence of occluders intersected by the line segment $S = [s^0 t^0]$.[2] Assuming that s^0 lies on the boundary of O_i, we define a *jump* motion which continuously takes s^0 from O_i to O_{i+1}.

For each occluder O_i define the *center point* of O_i as the midpoint of segment $O_i \cap [s^0 t^0]$. The first step of the jump is to align the supporting line of s with then center point of O_i. This is achieved by rotating s around its endpoint s^0 (this step is only necessary at the begining, when s^0 lies on ∂O_1). The second step of the jump rotates s around the center of O_i until it touches O_{i+1} and its supporting lines goes through the center of O_{i+1}. At the end of the jump, s lies on the line segment S. For the next jump, s will rotate around the center of O_{i+1} with its ending point s^1 on ∂O_{i+1} instead of its starting point s^0. A jump sequence is illustrated in figure 2-left.

After $j - 1$ jumps, one end of s lies on $\partial O_j = \partial B_t$ on which also lies t^0. If necessary (i.e., s^1 lies on ∂B_t) we can change the orientation of s as illustrated in figure 2-right so that s^0 gets on ∂B_t. It is then easy matter to move s to t while constraining s^0 to stay on ∂B_t.

Case 2: $B_C \neq \infty$. Again, if necessary, we perturb s and t so that their back blocker gets different from B_C. We need to find a sequence of jumps from B_s onto B_t, avoiding occluder B_C. This is straightforward since occluders are convex. Note however that we may need to get s jump on

[2] Some occluder may be tangent to segment $S = [s^0 t^0]$. We could simply ignore these occluders, but this is not necessary and would complicate the construction of the motion from s to t.

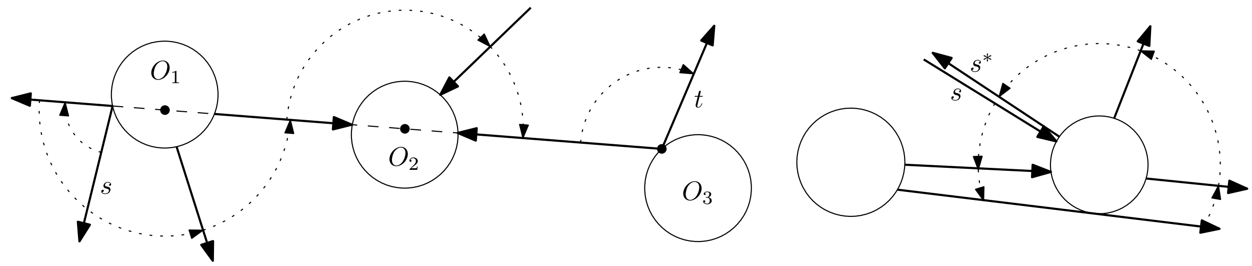

FIGURE 2. **Left**: a sequence of jumps takes s to t continuously, while avoiding cells of visibility (∞, ∞). **Right**: providing $k > 1$, the maximal free segment s can switch its orientation (thus becoming s^*) without crossing a 4-cell of type (∞, ∞).

additionnal *virtual* occluders if B_C is a large occluder in between B_s and B_t, and no other occluder can be jumped on from B_s or B_t. □

Lemma 3. *The free space $\mathcal{F}$ is path connected and simply connected.*

Proof. We let $\mathcal{I}$ denote the unit interval $[0, 1]$ of $\mathbb{R}$. We prove only the simple-connectness of $\mathcal{F}$, as proving its connectedness is similar. We use induction over the number k of occluders in $\mathcal{O}$. The case $k = 0$ is trivially handled. Now assume that $\mathcal{O}$ contains k occluders, $\{O_1, \ldots, O_k\}$, $k > 0$. Define $\mathcal{F}' = \mathbb{R}^3 \setminus \bigcup_{i=1}^{k-1} \mathring{O}_i$. By the induction hypothesis, $\mathcal{F}'$ is connected. $\mathcal{F}$ is obtained from $\mathcal{F}'$ by substracting the occluder O_n: $\mathcal{F} = \mathcal{F}' \setminus \mathring{O}_n$.

Let $x_0 \in \mathcal{F}$ and f a loop in $\mathcal{F}$ whose basepoint if x_0. f is a continous map $f : \mathcal{I} \mapsto \mathcal{F}$ so that $f(0) = f(1) = x_0$. Since $\mathcal{F} \subset \mathcal{F}'$, f is also a loop in $\mathcal{F}'$. Hence, there exists a homotopy F in $\mathcal{F}'$, deforming f to the constant loop x_0: $F : \mathcal{I} \times \mathcal{I} \mapsto \mathcal{F}'$ is continuous and satisfies $F(0, t) = f(t)$ and $F(1, t) = x_0$ for all t in $\mathcal{I}$. We now build a homotopy $G : \mathcal{I} \times \mathcal{I} \mapsto \mathcal{F}$ that satisfies $G(0, t) = f(t)$ and $G(1, t) = x_0$ for all t in $\mathcal{I}$. Let $\operatorname{Im} F$ be the image of F. As $\operatorname{Im} F$ is the image of the 2-dimensional square, there exists a point $O_k^* \in \mathring{O}_k$ which is not in $\operatorname{Im} F$.

Let $p \in \mathbb{R}^3 \setminus \{O_k^*\}$. We define p^* as the unique intersection of the boundary ∂O_k of O_k with the half-line having origin O_k^* and direction $p - O_k^*$. One can easily check that $p^* \in \mathcal{F}$. Now, define the map

$$\mathsf{push} : \mathcal{F}' \setminus \{O_k^*\} \to \mathcal{F}, \mathsf{push}(p) = \begin{cases} p, & \text{if } p \notin O_k \\ p^*, & \text{if } p \in O_k. \end{cases}$$

Since $\operatorname{Im} F$ does not contain O_k^*, the composition map $G = \mathsf{push} \circ F$ is continuous, and therefore, is a homotopy taking f to the constant loop x_0, while staying inside $\mathcal{F}$. This concludes the proof. □

Lemma 4. *The set $\mathcal{S}$ of maximal free segments is path connected and simply connected.*

Proof. Path-connectedness follows easily from lemma 3. We focus on simple connectedness.

Let $c : \mathbb{S}^1 \mapsto \mathcal{S}$ be a closed loop in the visibility complex. Equivalently, c is a continous function $\mathcal{I} \mapsto \mathcal{S}$, satisfying $c(0) = c(1)$. In order to prove lemma 4, we need to show that c is continously contractible to a point (a maximal free segment) in $\mathcal{S}$.

First, we extract a closed curve $p : \mathbb{S}^1 \mapsto \mathcal{F}$ such that $p(t)$ is a point of segment $c(t)$. Because $\mathcal{F}$ is simply connected (lemma 3), we can contract p to a point. This yields a possible contraction of c to a single maximal free segment in $\mathcal{S}$ thereby proving lemma 4.

We make the hypothesis that c goes from a 4-cell of $\mathcal{VC}$ to another in a transversal way, that is, if $c(t) \in \mathcal{S}^3$ ($c(t)$ is tangent to some occluder), then there exists $\varepsilon > 0$ such that $0 < |t' - t| < \varepsilon$ implies $c(t') \notin \mathcal{S}^3$. We justify our hypothesis by arguing that the path c' obtained by applying any sufficiently random perturbation to c does satisfy this hypothesis.

The length of segments $c(t)$ varies discontinuously when c crosses the boundary of a 4-cell of $\mathcal{VC}$. Let L be the set of point in $\mathcal{I}$ where the length of $c(t)$ changes discontinuously. L contains a finite number of points. Indeed, assume otherwise; then L admits an accumulation point $a \in \mathcal{I}$, which contradicts our transversality hypothesis. We write $L = \{t_i \in [0,1], i = 0..N_c\}$. The finite type of L [3] allows us to extract loop of finite length in $\mathcal{F}$ that intersects every free segment in c, as follows.

Let $\mathfrak{B}$ be a ball large enough to enclose all the occluders and intersect every maximal free segment of c. The midpoint of a segment $s = c(t)$ is defined as the midpoint of $s \cap \mathfrak{B}$. Let $t \in [0,1]$. If $t \notin L$, we define $p(t)$ as the midpoint of segment $c(t)$. p is not yet defined at point $t_i, i = 0..N_c$. However, $c(t_i)$ is tangent to at least one occluder and contains (considered as a set of points in $\mathbb{R}^3$) two limit points:

$$T_i = \lim_{t \to t_i^-} p(t) \quad \text{and} \quad T_i' = \lim_{t \to t_i^+} p(t)$$

Instead of defining p at point t_i, we insert the line segment $[T_i, T_i'] \subset c(t_i)$ to "glue" the parts of p already defined. This gluing process is achieved by enlarging, in $[0,1]$, every singletons $\{t_i\}$ to a small interval $[a_i, b_i]$ with positive length (we effectively reparameterize c to make room for a finite number of intervals). Then, we map the interval $[a_i, b_i]$ to the segment $[T_i, T_i']$. Now, p is a continuous loop in $\mathcal{F}$. Likewise, we build the continuous application $\omega : [0,1] \mapsto \mathbb{S}^2$ such that

$$(p(t), \omega(t)) = c \circ r(t), t \in [0,1]$$

where r is a continuous reparameterization of $[0,1]$ (r "smashes" each interval $[a_i, b_i]$ whose image is $[T_i, T_i']$ on $\{t_i\}$). Since $\mathcal{F}$ (lemma 3) and $\mathbb{S}^2$ are simply connected, one can contract $c \approx (p, \omega)$ to a point in $\mathcal{S}$. □

We can now prove theorem 1.

Proof of theorem 1. Lemma 2 yields the following idea to prove theorem 1: take s_1 and s_2 in the boundary of C. We consider a loop $\ell : [0,1] \mapsto \mathcal{S}$ such that $\ell(0) = \ell(1) = s_1$, $\ell(0.5) = s_2$, $\ell((0, 0.5)) \subset C$ and $\ell((0.5, 1)) \subset C^c$. Such a loop exists because both C and C^c are path-connected. As we have shown in lemma 4, the set $\mathcal{S}$ of maximal free segments is simply connected, hence we can contract ℓ to a single point inside C. Using this homotopy, we prove the existence of a path from s_1 to s_2 inside ∂C, thus proving our main theorem.

Let $H\colon [0,1] \times [0,1] \to \mathcal{F}$ be a homotopy between ℓ and a constant loop in C (that is, for all $t \in [0,1]$, $H(0,t) = \ell(t)$, $H(1,t)$ is a fixed point in C, and for all $x \in [0,1]$, $H(x,0) = H(x,1)$). We may suppose that H is transverse to C, or equivalently:

- for all but finitely many $x \in [0,1]$, the loop $H(x,\cdot)$ crosses ∂C transversely,
- for finitely many $x_1, \dots, x_n \in (0,1)$, the loop $H(x_i,\cdot)$ crosses ∂C transversely, except at some $t \in [0,1]$, and the evolution of this intersection is as depicted in figure 3.

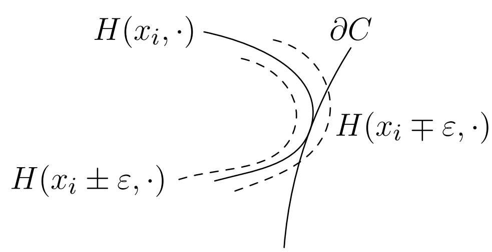

FIGURE 3. Non-transversal crossing of ∂C.

Now fix a connected component C_0 of ∂C. The number $n(x)$ of intersection points between the loop $H(x,\cdot)$ and the set C_0 changes only at $x_1, \dots, x_n$, and is constant modulo 2. Now let C_0 be the connected component of ∂C containing s_1. Then $n(1) = 0$ hence $n(0)$ has to be even. It follows

[3] A countable set L would have been sufficient for our purpose.

that $s_2 \in C_0$. Since connectedness and path-connectedness coincide in CW-complexes, there exists a path joining s_1 and s_2 in ∂C. □

2.1. **Discussion.** All the results presented above stay valid when one consider general convex objects for which the set of maximal free segments is a cell-complex. This is true for convex polyhedra and for semi-algebraic convex sets. We also believe it to be true for any kind of convex objects.

A natural question would be to extend theorem 1 to non necessarily convex objects, e.g., to a finite set of disjoint bounded 3-manifolds, each homeomorphic to the 3-ball. In this setting, lemmas 3 and 4 stay valid. The proof of lemma 2 however makes use of the convexity of the occluders, but we conjecture the statement of lemma 2 to be true in this more general setting.

3. Constructing the 3D visibility complex

3.1. **Tesselation and adjacencies in the visibility complex.** Throughout the paper, we assume the generic position of the occluders, so that each i-cell is of dimension i and no maximal free segment is tangent to more than four occluders. However, the definitions given above do not lend themselves to a simple computationally tractable data structure. Consider for example, a single convex polyhedron A in space. Using the definitions given above, the visibility complex consists of three 4-cells with visibility (∞, ∞), (A, ∞) and (∞, A). Faces (A, ∞) and (∞, A) have exactly the same geometry and differ only in the orientation of their maximal free segments. Additionally, the three 4-cells have the same boundary which is the 3-cell consisting of all maximal free segments tangent to A. The visibility complex does not comprise any 0-, 1- or 2-cell. Note also that the only 3-cell has no simple description in terms of data-structures. Let $\mathcal{T}_A$ be the set of free segments tangent to polyhedron A. In order to simplify the data structures, we want to subdivide $\mathcal{T}AP$ in such a way that every i-cell has a simple description and each $i+1$-cell is adjacent to some i-cell.

A maximal free segment s with orientation ω_s is said to be *tangent to edge* e of polytope A if it is tangent to A at a point on e and edge e appears as a silhouette edge with positive length on the orthogonal view of A along direction ω_s. When s intersects the interior of a face f along a line segment, all edges of f intersected by s are considered as silhouette edges (which amounts to two, three or four edges).

Using this definition an i-cell is simply defined as a maximally connected set of free segments tangent to the same $4-i$ edges. In particular, every free segment containing two vertices of a polytope forms a 0-cell (since it is tangent to exactly four edges), whose type we call VV since the free segment passes through two vertices. 0-cells can have type VV (with one or two polytopes involved), VEE (with two or three polytopes) or EEEE (when four polytopes are involved). A free segment passing through a vertex v of face f and an edge of f not containing v is tangent to three edges of the polyhedron and is therefore part of a 1-cell, bounded by two 0-cells of type VV or VEE. A free segment in the interior of a 2-cell is either tangent to two edges in their interior (the 2-cell is of type EE), or tangent to a vertex of a polyhedron (the 2-cell is of type V). Free segments inside a same 3-cell are always tangent a same edge. This is a key property of 3-cells. Indeed, in order to build the 3-skeleton of the visibility complex, we consider each edge e in turn and construct all 3-cells whose edge of tangency is e at once with a sweep. Thus, at the end of each sweep along an edge, we have the garantee that all swept 3-cells and their boundaries have been completely constructed. So we never have to "leave" a cell in a partially constructed state. When a 3-cell is swept and one of its boundary 2-cell, b, already exists, one just has to update the adjacency relations for cell b.

In essence, the new definition of i-cells gives us a fine tesselation of the set $\mathcal{T}_A$ of free segments tangent to each polyhedron A, wherein each i-cell has a computationally tractable geometry; in particular, 2-cells are represented by either spherical polygons (for 2-cells of type V) or 2-dimensional regions in the 2-dimensional space of maximal free segments tangent to two edges (for 2-cells of type

EE). A 3-cell c forms a 3-dimensional volume that can be swept along the unique edge of tangency common to all free segments in c.

3.2. **Construction of $\mathcal{VC}$.** A three-dimensional object is often represented by expliciting its external boundary together with the boudaries of its cavities. In the same way, we want to build the boundary of each 4-cell of the visibility complex. We now give a simple algorithm to do so for a set of disjoint convex polyhedra. More precisely, we first explicitely build the 3-skeleton of $\mathcal{VC}$, and then use the connexity property of theorem 1 to extract the boundary of each 4-cell. Since the boundary of a 4-cell is connected, extracting it amounts to finding some connected component in the adjacency graph given by the 3-skeleton. Let $\mathcal{G} = (\mathrm{Sk}_3, E_3)$ the graph wherein Sk_3 is the set of 3-cells of $\mathcal{VC}$. and $E_3 \subset \mathrm{Sk}_3 \times \mathrm{Sk}_3$ is the adjacency relation between 3-cells. During the construction of $\mathcal{G}$, each 3-cell b is tagged with a set $\mathrm{adj}(b)$ of three pairs $(A, B) \in (\mathcal{O} \cup \{\infty\})^2$ corresponding to the visibility of the three 4-cells adjacent to b. Then, given a 4-cell C whose visibility is (A, B) and a 3-cell $b \in \partial C$, we can retrieve the whole boundary of C as the connected component $K_b \subset \mathrm{Sk}_3$ which is maximaly connected in $\mathcal{G}$, and such that for all $b' \in K_b$, we have $(A, B) \in \mathrm{adj}(b')$. We now examine how to construct the 3-skeleton $\mathcal{VC}^3$ of $\mathcal{VC}$ from which the graph (Sk_3, E_3) can be (implicitly) extracted.

$\mathcal{O}$ is a set of polytopes. Hence, a 3-cell b consists of a maximaly connected set of free segments tangent to a given edge e of one polytope, and having the same visibility. When the visible set of b is $(A.B)$, we say that b is of the form (e, A, B). For each edge e of the scene, we construct the 3-cells of the form (e, o_1, o_2), where $o_1, o_2 \in \mathcal{O} \cup \{\infty\}$.

We assume that the 1-skeleton of $\mathcal{VC}$ has already been constructed. This can be achieved in time $O(n^2 k^2 \log n)$ with the algorithm of Goaoc *et al.* [Goa04, BDD+06]. k is the number of polytopes and n the total number of edges. The algorithm sweeps the set of maximal free segments tangent to e with a plane, containing e and rotating around e. Although it may be possible to construct the whole set of 3-cells of the form $(e, *, *)$ using the same plane-sweep, we prefer using a different kind of sweep, that is much easier to visualize.

3.2.1. *2-dimensional radial cuts in $\mathcal{S}$.* Let P be a point in the interior of the free space. The visibility graph $V(P)$ of P is the projection of the silhouette edges of $\mathcal{O}$ visible from P, onto the unit sphere centered on P. $V(P)$ is a "spherical" graph. A vertex of $V(P)$ is either a (silhouette) vertex of some occluder $o \in \mathcal{O}$ or a *t-vertex*, that is, the visual intersection of two visible silhouette edges.

Now let us consider the set $\mathcal{S}_P$ of maximal free segments passing through P. It is the intersection of a bundle of oriented lines with the visibility complex, and is homeomorphic to $\mathbb{S}^2$. Again, we can partition $\mathcal{S}_P$ into "connected" equivalence classes with respect to the visible set of the maximal free segments. This is equivalent to considering the intersection of the partition $\mathcal{VC}$ with $\mathcal{S}_P$. Such a decomposition of $\mathcal{S}_P$ can be visualized as the arrangement of $V(P)$ with the central symetrical of $V(P)$ itself; We thus view $\mathcal{S}_P$ as a central symetrical planar graph on $\mathbb{S}^2$. When P is in the interior of the free space, vertices of $\mathcal{S}_P$ are points in some 2-cells of $\mathcal{VC}$, that is, bitangents or tangents through a vertex of $\mathcal{O}$. Edges of $\mathcal{S}_P$ are 1-dimensional curves in some 3-cells of $\mathcal{VC}$, that is, segments tangent to a same edge. Finally the faces of $\mathcal{S}_P$ are 2-dimensional sheets in some 4-cells of $\mathcal{VC}$, that is, the intersection of $\mathcal{S}_P$ with the interior of one 4-cell.

We are interested in the free segments that are tangent to edges of the polytopes in $\mathcal{O}$. By placing point P on edge e, $\mathcal{S}_P$ becomes the set of free segments that are tangent to e at point P (figure 4).[4] It is clear that when P moves along an edge e, $\mathcal{S}_P$ sweeps all the 3-cells of the form $(e, *, *)$. If we can maintain the decomposition of $\mathcal{S}_P$ as P moves from one end of e to the other, we'll be able to construct all 3-cell whose edge of tangency is e. The spherical graph $\mathcal{S}_P$ retains the same combinatorics as P moves, except at some positions of P when three edges of the graph

[4] One ignores the free segments that touch e but are not tangent to e at P.

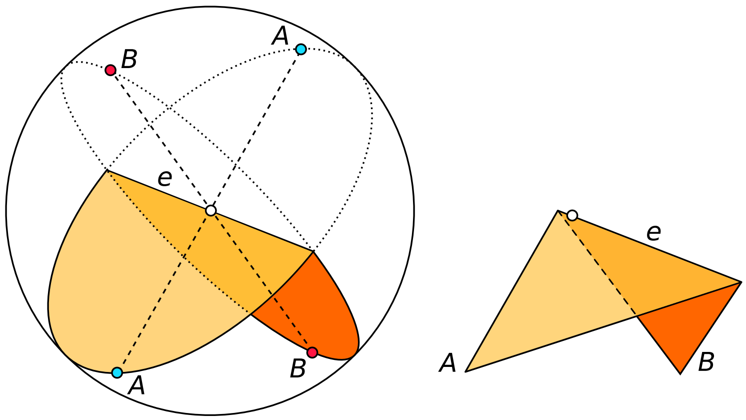

FIGURE 4. Sweeping a 3-cell of the form $(e, *, *)$. The white dot indicates the position of point P at some time of the sweep. **On the right**, one observes the edge e being swept, and both its adjacent polygons. **On the left**, one observes how these two polygons partition the sphere of directions around P.

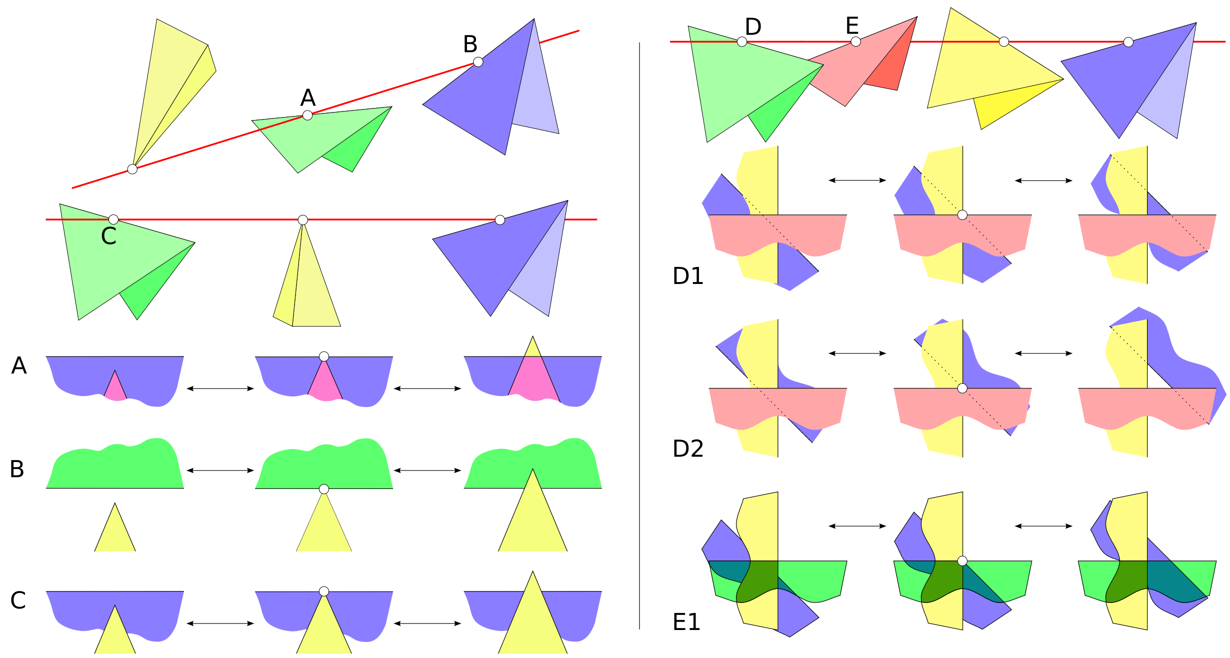

FIGURE 5. Updating $\mathcal{S}_P$ when P crosses a 0-cell of type VEE (to the left) or EEEE (to the right). In the bottom illustrations, each coloured region is a 2D slice of a 3-cell. To the right, cases D1 and D2 are symetrical. Case E2, symetrical to E1, is not illustrated here.

cross. Each such crossing corresponds to a 0-cell of the visibility complex, that is, to a maximal free segment tangent to four edges. If the 1-skeleton of the visibility has already been constructed, we can thus directly know the combinatorial changes that happen as P moves on edge e. We have informally discussed how the visibility complex can be constructed from its 1-skeleton.

3.2.2. *Constructing 3-cells attached to segment e.* Let $e = \overline{v_0 v_1}$ be an edge of polytope A. e is adjacent to two polygons $t_1(e)$ and $t_2(e)$. We denote by $\mathcal{D}(e)$ the subset of $\mathbb{S}^2$ equal to the set of directions of the tangents to A in the relative interior of e. $\mathcal{D}(e)$ is a double-moon on the sphere of direction, in between the two planes supporting $t_1(e)$ and $t_2(e)$ (figure 4). Let $\mathcal{VC}(e)$ the subset of $\mathcal{VC}$ of maximal free segment tangent to e, and call it the *column* of e. The 3-cells in $\mathcal{VC}(e)$ are constructes, together, as a partition of $[0,1] \times \mathcal{D}(e)$. First, the point P is placed on vertex v_0 and the 2-cells attached to vertex v_0 are constructed from the 1-skeleton. Those 2-cells form parts of the boundary of some 3-cells in $\mathcal{VC}(e)$. Then, we move P along e toward v_1.

3.2.3. *Construction of $\mathcal{S}_{v_0}$ and its maintenance along e.* Since the 1-skeleton of $\mathcal{VC}$ has already been computed, the 0- and 1-cells in $\mathcal{S}_{v_0}$ are readily available, and one computes the 2-cells of $\mathcal{S}_{v_0}$ by sweeping the sphere of direction centered on v_0. This sweep process is analogous to the sweep of a set of polygons in the plane. In particular, no intersection can be found during the sweep; the sweep process is only used to bind each 2-cells to its boundaries. The construction of $\mathcal{S}_v$ for all vertices of the polytopes takes time $O(n^2k^2 \log(n^2k^2)) = O(n^2k^2 \log n)$.

Let $e_1, e_2, \ldots, e_j$ be the edges of A adjacent to v_0. The support of $\mathcal{S}_{v_0}$ in $\mathbb{S}^2$ is $\cup_{i=1..j}\mathcal{D}(e_i)$. Also, a free segment s in $\mathcal{S}_{v_0}$ is adjacent to two columns: the columns $\mathcal{VC}(e)$ and $\mathcal{VC}(e')$ where e and e' are the two silhouette edges adjacent to v_0 in the orthographic view parallel to s. Prior to sweeping $\mathcal{VC}(e)$, we extract the 2-cells of $\mathcal{S}_{v_0}$ included in $\mathcal{D}(e)$; this gives us $\mathcal{S}_P$ for a point P on e infinitely close to v_0. Each face of $\mathcal{S}_P$ starts a 3-cell in the column of e, partially bounded by 2-cells in $\mathcal{S}_{v_0}$.

3.2.4. *Handling the sweep events.* Let P_0 be a point of e where a combinatorial change of $\mathcal{S}_P$ happens. Then, in $\mathcal{S}_{P_0}$, 3 edges cross at point $u_0 \in \mathbb{S}^2$. This is an event of type VE or EEE, which indicates a maximal free segment $s = (P_0, u_0)$ tangent to e and to three other polytope edges. Thus, s is tangent to four edges and is a 0-cell of the visibility complex: s has already been constructed as an element of the 1-skeleton of $\mathcal{VC}$. For each column of and edge e, we thus extract all 0-cells tangent to e and sort them according to their intersection point with e. This gives all the events needed for sweeping the column of e. Again, sorting the sweep events for all columns takes time $(On^2k^2 \log n)$. Consequently, all the 3-cells of the visibility complex can be constructed in the same time bound of $O(n^2k^2 \log n)$. Finding the connected components forming the boundary of each 4-cell in the graph (Sk_3, E_3) takes linear time.

4. Conclusions

We have proved that the boundary of every 4-cell in the visibility complex of disjoint convex objects in space is path connected. It may be interesting to look for other topological or conbinatorial properties of the 3D visibility complex. Other problems arise regarding the construction of the visibility complex: for example, given one maximal free segment s, how can we build the only cell of $\mathcal{VC}$ that contains s, without constructing other 4-cells? For problems related to the generation of photorealistic pictures, light rays are often simulated and bounced several times in the 3D scene, thus visiting a set of 4-cells of the visibility complex. Computing this efficiently may help in driving the accuracy of these light simulations up.

References

[AAKS05] Pankaj K. Agarwal, Boris Aronov, Vladlen Koltun, and Micha Sharir. On lines avoiding unit balls in three dimensions. *Discrete and Computational Geometry*, 34(2):231–250, August 2005.

[Ang02] Pierre Angelier. *Algorithmique des graphes de visibilité*. PhD thesis, École Normale Supérieure (Paris), February 2002.

[BDD+06] Hervé Brönnimann, Olivier Devillers, Vida Dujmović, Hazel Everett, Marc Glisse, Xavier Goaoc, Sylvain Lazard, Hyeon-Suk Na, and Sue Whitesides. Lines and free line segments tangent to arbitrary three-dimensional convex polyhedra. *SIAM Journal on Computing*, 2006. To appear. A preliminary version of this article appeared in the proceedings of the 20th SoCG, Brooklyn, New York, USA, 2004.

[CGA] CGAL. *Computational Geometry Algorithm Library.* `http//www.cgal.org/`.

[DDE+03] Olivier Devillers, Vida Dujmović, Hazel Everett, Xavier Goaoc, Sylvain Lazard, Hyeon-Suk Na, and Sylvain Petitjean. The expected number of 3d visibility events is linear. *SIAM Journal on Computing*, 32(6):1586–1620, 2003.

[DDP02] Frédo Durand, George Drettakis, and Claude Puech. The 3D visibility complex. *ACM Transactions on Graphics*, 2002.

[ELPZ05] Hazel Everett, Sylvain Lazard, Sylvain Petitjean, and Linqiao Zhang. An experimental assessment of the 2d visibility complex. In *Proceedings of 17th Canadian Conference on Computational Geometry (CCCG)*, 2005.

[Goa04] Xavier Goaoc. *Structures de visibilité globales: tailles, calculs et dégénérescences.* PhD thesis, Université de Nancy, 2004. in french.

[Hat01] Allen Hatcher. *Algebraic Topology.* Cambridge University Press, 2001.

[PV96a] Michel Pocchiola and Gert Vegter. Topologically sweeping visibility complexes via pseudo-triangulations. *Discrete Computational Geometry*, 16(4):419–453, December 1996. Special issue devoted to the proceedings of the 11th Annual ACM Symposium on Computational Geometry (SoCG'95).

[PV96b] Michel Pocchiola and Gert Vegter. The visibility complex. *International Journal on Computational Geometry and Applications*, 6(3):279–308, 1996. Special issue devoted to the proceedings of the 9th Annual ACM Symposium on Computational Geometry (SoCG'93).

Samuel Hornus, Division of computer science, Korea advanced institute of science and technology, Daejeon, Korea
E-mail address: `hornus@tclab.kaist.ac.kr`

Maxime Wolff, Institut Fourier, Université Joseph Fourier, Grenoble, France
E-mail address: `Maxime.Wolff@ujf-grenoble.fr`