Power Diagram Depth Sorting

Paolo Cignoni
Istituto di Elaborazione dell'Informazione – Consiglio Nazionale delle Ricerche
Via S. Maria 46 - 56126 Pisa, ITALY Email: `cignoni@iei.pi.cnr.it`

Leila De Floriani
Dipartimento di Informatica e Scienze dell'Informazione – Università di Genova
Via Dodecaneso, 35 - 16146 Genova, ITALY Email: `deflo@disi.unige.it`,

Abstract

In this paper we propose a new approach to the depth sorting problem. Given a simplicial complex Γ in $\mathbb{E}^d$, our approach is based on the preliminary construction of a convex complex Γ^* in $\mathbb{E}^{d+1}$, whose projection on $\mathbb{E}^d$ corresponds to Γ, and on the representation of Γ^* as a power diagram. This approach exhibits a $O(m \log m)$ runtime complexity to sort a simplicial complex with m cells and requires only linear storage.

1 Introduction

The problem of sorting a set of objects in space according to their distance from a given viewpoint has been extensively studied in the literature mainly for its important applications in computer graphics. We were interested in this problem because its relevance in volume visualization; another area where the two-dimensional version of this problem arises is visibility computation on terrains [8].

Many techniques exist [3] for the direct visualization of an unstructured volumetric dataset orgainized as a tetrahaedral complex; projective techniques are extensivley used [9, 11, 14]. Such techniques are based on a two-phase process: first, cells are sorted in depth and then projected onto a view plane; the visual contribution of each cell is composed on the frame buffer. The depth order of the cells is necessary to guarantee the correctness of the process. Volume datasets can be composed of a considerably large number of cells (from 100k to 10M cells), thus efficiency and memory requirements of a depth sorting algorithm are relevant.

The main contribution of this paper is a proposal for a new technique for depth sorting a certain class of cell complexes that include those used in applications (e.g., Delaunay complexes). Given a cell complex Γ in $\mathbb{E}^d$, our approach is based on the preliminary construction of a convex complex Γ^* in $\mathbb{E}^{d+1}$ corresponding to Γ and on its representation as a power diagram. This approach exhibits a $O(m \log m)$ runtime complexity to sort a simplicial complex with m cells and requires only linear storage.

The remainder of the paper is organized as follows. In Section 2, we introduce some background notions, in Section 3, we shortly summarize some of the most important results in the literature concerning the depth-sorting problem. Section 4 describes our technique for sorting, while in Section 5 the problem of computing Γ^*, called the *lifting problem*, is discussed.

2 Background

2.1 Obstruction relation and cell complexes

The *obstruction relation* $\prec_p$ (usually called *in-front/behind* relation) for a pair of non self-intersecting objects A and B in the Euclidean d-dimensional space with respect to a viewpoint p, denoted as $\prec_p$, can be formally expressed as follows:

$$A \prec_p B \text{ iff } \exists \text{ a ray } r \text{ emanating from } p$$

and intersecting both A and B, such that any points in $r \cap A$ is closer to p than any point in $r \cap B$.

A *depth order* of a set of objects, with respect to a viewpoint p, is a sequence of such objects such that, if object A obstructs object B when seen from p, then A comes before B in the sequence.

Informally, a cell compex is a subdivision of the d-dimensional Euclidean space into cells. A k-cell in $\mathbb{E}^d$, with $0 \leq k \leq d$ is a closed connected subset of $\mathbb{E}^d$ such that its interior is homeomorfic to an open k-dimensional disk and whose boundary is not null.

Simplicial complexes can be viewed as a sbu-class of cell complexes. Consider a set V_σ of $k+1$ linearly indipendent points in the d dimensional space E^d, with $k \leq d$. The subset σ of E^d formed by all the points that can be expressed as the linear convex combination of the points of V is called a *k-simplex.* The points of V ar called *vertices* of σ, while k is the *order* of σ. Using this definition, a vertex is a 0-simplex, an edge is a 1-simplex, a triangle is a 2-simplex, and a tetrahedron is a 3-simplex.

Any s-simplex χ, $0 \leq s \leq d$, which is generated by a subset of $s+1$ vertices of σ, is called an *s-face* of σ or, shortly a face of σ. A collection Σ of simplices is called a *d-simplicial complex* when the following conditions hold:

- $\forall \sigma \in \Sigma$, all the faces of σ belong to Σ;
- $\forall \sigma, \tau \in \Sigma$ either $\sigma \cap \tau = \emptyset$ or $\sigma \cap \tau$ is a face of both σ and τ;
- d is the maximum order of simplices $\sigma \in \Sigma$

A cell complex Γ is called *acyclic* with respect to a given viewpoint p if and only if relation $\prec_p$ defines a partial order over the cells of Γ. In this case, it is possible to extract a depth order of the cells of Γ with respect to the viewpoint.

Cell complexes in $\mathbb{E}^d$ that can be obtained as the orthogonal projection of the lower part of the boundary of a convex polytope in $\mathbb{E}^{d+1}$ are called *projective* [2]. Projective cell complexes have been shown to be acyclic with respect to any viewpoint [6].

Theorem 2.1. *The in-front/behind relation defined for the cells of any projective complex and for any fixed viewpoint in $\mathbb{E}^d$ is acyclic.*

It is an open problem to show the existence of an acyclic complex in $\mathbb{E}^d$ that cannot be obtained as the orthogonal projection of a convex complex in $\mathbb{E}^{d+1}$.

Delaunay simplicial complexes (as well as any proper subcomplex of a Delaunay one) [5] are projective and, thus, acyclic with respect to any viewpoint [6]. This is an important property for application purposes, because it assures that a volume dataset organized as a Delaunay complex can always be sorted and correctly visualized. In the two-dimensional case, this means that an algorithm which exploits the acyclicity property can be used to compute the visiblity map of a Delaunay-based triangulated terrain.

2.2 Power Diagrams and Convex Polyhedra

In this Subsection, we review some definitions and results regarding power diagrams; a complete introduction about power diagrams can be found in [2]. Power diagrams are considered one of the generalization of Voronoi diagrams that have the strongest similarities to the original diagrams.

The power of a point p with respect to a sphere $s \subset \mathbb{E}^d$ with center z and radius r, is defined as $pow(p,s) = d(p,z)^2 - r^2$. Thus $pow(p,s) < 0$ if p belongs to the ball bounded by s, $pow(p,s) = 0$ if p lies on the surface of s and is greater than zero otherwise; in this latter case it easy to show that $pow(p,s)$ is equal to the squared distance of p from the touching point of a line tangent to s through p.

Let s and t be two spheres in $\mathbb{E}^d$ with centers $z_s \neq z_t$ and radii r_s, r_t. The points x satisfying $pow(x,s) = pow(x,t)$ describe a hyperplane h perpendicular to the straight line joining z_s and z_t, known as the *chordale* of s and t, ond denoted $chor(s,t)$ for short. A nice property of chordales is that, if $s \cap t \neq \emptyset$, then $s \cap t \subset chor(s,t)$; moreover, s and t are contained into the same open halfspace bounded by $chor(s,t)$ if and only if s encloses or it is enclosed in t.

Let S denote a finite set of n spheres in $\mathbb{E}^d$, for $s \in S$ we call the set:

$$cell(s) = \{x \in \mathbb{E}^d | pow(x,s) > pow(x,t) \forall t \in S - \{s\}\}$$

as the *power cell* of s and the collection of all $cell(s)$, for $s \in S$, the *Power Diagram* of S, or PD(S) for short.

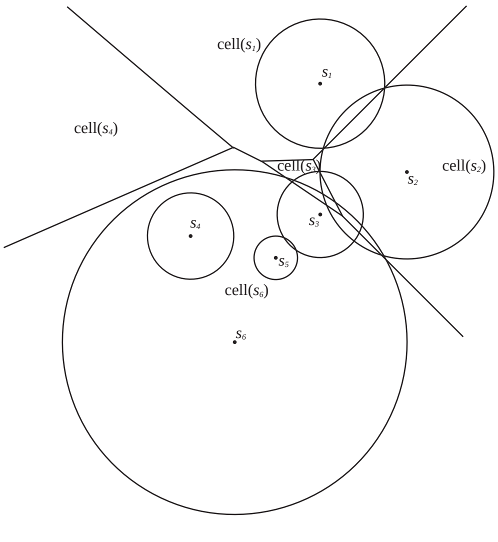

Figure 1: A power diagram in two dimensions.

By definition *cell*(s) is the intersection of $n-1$ halfspaces bounded by chordales, and therefore is a $d-$polyhedron with at most $n-1$ faces. Thus PD(S) is a cell complex in $\mathbb{E}^d$. Figure 1 shows a PD of six circles in $\mathbb{E}^2$; there can be empty cells (like *cell*(s_5)) and cells separated from their generating spheres, (*cell*(s_4) is distinct from s_4).

It can be shown that the PD of a set of congruent spheres is the Voronoi diagram of their centers.

Let $\mathbb{E}^{d+1}$ be spanned by the coordinate axes $x_1, \cdots, x_{d+1}$, and let h_0 denote the hyperplane $x_{d+1} = 0$. The following result, that relates PD and convex polyhedra, is presented in [2]:

Theorem 2.2. *For any (d+1)-polyhedron P, which can be expressed as the intersection of upper halfspaces, there exists an affinely equivalent power diagram in h_0, and viceversa.*

It is assumed, without loss of generality, that all the upper halfspaces generating polyhedron P cross the unitary paraboloid $U : x_{d+1} = \sum_{i=1}^{d} x_i^2$. The above theorem is based on the following transform Π that maps a sphere $s \subset h_0$ with center z and radius r into the hyperplane:

$$\Pi(s) : x_{d+1} = 2x \cdot z - z \cdot z + r^2 \quad (1)$$

Hyperplane $\Pi(s)$ crosses the unitary paraboloid U in such a way that the projection of $U \cap \Pi(s)$ onto h_0 is equal to s. The reverse mapping from planes crossing U back to spheres in h_0, can be defined.

An interesting property of this transform is that, if s and t are two non-concentric spheres in h_0, then $chor(s,t)$ is the vertical projection of $\Pi(s) \cap \Pi(t)$ onto h_0. Therefore, given a polyhedron P, it is possibile to obtain a set of spheres S whose PD is the same as the projection of P onto h_0.

3 Related Work

The problem of depth sorting a set of objects in space has received considerable attention in the literature. Here, we review just some of the most important results concerning the three-dimensional depth-sorting problem. For a general and thorough overview of the problem see de Berg's PhD Thesis [4].

It should be noted that in the general case of a non-convex, possibly unconnected, arbitrary tetrahedral complex of m cells, the problem of calculating a depth ordering (if it exists), with respect to a given viewpoint, has a lower bound of $O(m \log m)$ [12], and none of the proposed algorithms still matches this bound.

An important result can be found in [4] where de Berg observes that the depth sorting problem is the same as computing a linear extension of the $\prec_p$ relation; de Berg describes an algorithm that solves this problem in a general way for a given relation $\prec$ and its transitive closure $\prec_*$ on a set S of m objects. The resulting algorithm can compute a depth order for a set of segments, or triangles in space, or decide that there is a cyclic overlap among them with a worst-case complexity of $O(m^{4/3+\epsilon})$.

Another solution to the depth sorting problem has been given in [1] by Agarwal et al.; their approach solves the linear extension problem, when the relation $\prec$ is not cyclic, and for triangles whose xy projections are *fat* enough in $O(m \log^6 m)$.

Because of the importance of depth sorting for applications in volume visualization, several algorithm has been developed and used for ordering a three-dimensional simplicial complex composed of m cells:

	Topological	Numerical	de Berg	NNS
Class of complexes	convex complexes	Delaunay (subset of)	any complex	any complex
Worst Case Complexity	$O(m)$	$O(m \log m)$	$O(m^{4/3+\epsilon})$	$O(m^2)$
Storage OverHead	$O(m)$	$O(m)$	$O(1)$	$O(1)$

Table 1: Results on computing depth orders for tetrahedral complexes of m cells.

Topological Sort The cells of an acyclic convex complex can be sorted by exploiting the face-adjacency relation between tetrahedra and face orientation. An algorithm, called Meshed Polyhedra Visibility Ordering (MPVO), based on this approach was proposed by Williams [14, 15]. In a first phase the obstruction relation is computed for all pairs of adjacent cells, and the adjacency graph into a directed acyclic graph (DAG). Then, a total order of the cells is obtained in linear time by a topological sorting the DAG.

Numerical Sort A technique, working only for Delaunay tetrahedral complexes, was proposed in [9]; it is based on the fact that the length of the tangent from the viewpoint to the sphere circumscribed to any tetrahedron reflects the depth ordering of the complex. The worst case complexity of this technique is $O(m \log m)$. This technique can be applied only to Delaunay complexes and is sensitive to degenerate situations.

Newell, Newell and Sancha's Sort The classical Newell, Newell and Sancha (NNS) sort algorithm [10] for triangles scattered in $\mathbb{E}^3$ has been extended to manage polyhedral cells instead of polygons by Steiner et al. [13]. In the same way as the original NNS sorting algorithm, the sorting process is organized into two phases. In the first phase an approximate depth sort of the polyhedra is calculated. The second step is a quadratic *fine tuning* of the sort where we check each cell against the ones that are nearest in the approximate order. In Table 1 we summarize the main characteristics of algorithms for sorting three-dimensional cell complexes; de Berg's algorithm has never been practically adopted for visualization purposes, but it is the best theoretical result for sorting a generic cell complex.

The above algorithms can generate a correct depth order only if the starting complex is acyclic with respect to the specified viewpoint. However, practical algorithms for testing if a cell complex is acyclic for any arbitrary viewpoint are not known. A brute-force algorithm, based on the idea of sorting the complex from all significant possible viewpoints, was sketched as a personal communication between H. Edelsbrunner and P. Williams [15]. The idea is to place all significant viewpoints in all the cells γ of the plane arrangement $\mathcal{H}$ generated by the partition of the space with planes passing through all the faces of the cells. All the points p in the same cell γ of $\mathcal{H}$ certainly share the same depth ordering of the original complex Σ w.r.t. p, because the occlusion relation between two cells can change only when crossing the plane passing through one of the faces of the two cells. The main drawback of this approach is its complexity, since the arrangement of the planes defined by the faces of a 3D simplicial complex with n vertices can contain $O(n^3)$ cells, therefore if we sort the complex for each reasonable direction using a topological approach we obtain an $O(n^4)$ complexity, that imposes serious restictions to the usability of the algorithm.

4 Depth Sorting a Power Diagram

Our motivation for using Power Diagrams lies in the possibility of easily calculating a depth ordering of the cells of such a structure. For this purpose, first we extend a result by Aurenhammer, then we show that the $pow()$ function "agrees" with the obstruction relation with respect to a given viewpoint.

In Theorem 2.2 Aurenhammer considers polyhedra formed by the intersection of non-vertical upper halfspaces, therefore unbounded infinite polyhedra whose projection onto h_0 is a partition of the whole plane. With the following theorem we extend this relation in order to be able to manage bounded polyhedra such as the lower part of a convex hull.

Theorem 4.1. *Let P be the lower part of a convex $(d+1)$-polyedron in $\mathbb{E}^{d+1}$ bounded by simplicial faces; there exists a power diagram $PD(S)$ of a set of spheres S $\mathbb{E}^d$ whose cells are a superset of the projection of the faces of P onto h_0.*

Proof As in the previous case we assume, without loss of generality, that the hyperplanes defined by all the d-faces of P cross the unitary paraboloid U.

Consider the halfspaces passing through the d-faces of P and containing P; let P' be the unbounded $(d+1)$-polyhedron which is obtained by the intersection of these upper halfspaces. The internal faces of P (the ones not having a $(d-1)$-face on the boundary of P) have a direct correspondence in P'. The d-faces having one or more $(d-1)$-faces on the boundary of P become larger, possibly unbounded, d-faces.

Let P_0 be the projection of P on h_0. As a consequence of Theorem 2.2 we can build the PD that is affine to the projection P_0' of P' on h_0. It is easy to see that each d-face σ of P_0 corresponds to a d-face σ' of P_0' such that $\sigma = \sigma'$ if σ is an internal d-face, or $\sigma \subset \sigma'$ if σ is a boundary d-face. □

It should be noted that in this case the direct corrispondence between faces of P_0 and P_0' holds only for d-faces. In Figure 2, we can see an example of this situation: the cells σ_i with a $(d-1)$-face on the boundary correspond to unbounded σ_i' cells of the PD that have lost their boundary $(d-1)$-faces.

In order to calculate the depth sorting of a PD we introduce a numerical function Φ that, given a viewpoint p, *agrees* with the infront/behind relation that is: $\Phi_p(\sigma)$ such that $\Phi_p(\sigma) < \Phi_p(\tau)$ if $\tau \prec_p \sigma$. Clearly if such a numerical function exists for each given viewpoint, then $\prec_p$ is acyclic since it is impossible to have a set of cells $\sigma_1, \dots, \sigma_k$ that forms an occluding cycle.

Let Γ^* be the convex polyhedron in $\mathbb{E}^{d+1}$ whose projection onto h_0 corresponds to the simplicial complex Γ; for Theorem 2.2 using transform Π we can build the set of spheres S_Γ such that the power diagram $PD(S_\Gamma)$ is affine to Γ in the sense of Theorem 4.1, that is for each d-cell σ of Γ correspond a possibly larger cell σ' of $PD(S_\Gamma)$.

With the following theorem we show the relation between the $pow()$ function and the occlusion relation.

Theorem 4.2. *Given a viewpoint p, the infront/behind relation among two simplexes σ, τ agrees with the power of the viewpoint with respect to the corresponding spheres s_σ and s_τ of S_Γ; that is: $\sigma \prec_p \tau \rightarrow pow(p, s_\sigma) < pow(p, s_\tau)$,*

Proof We can see that $\sigma \prec_p \tau \rightarrow \sigma' \prec_p \tau'$ so we has to show that $\sigma' \prec_p \tau' \rightarrow pow(p, s_\sigma) < pow(p, s_\tau)$, It easy to show that, if $\sigma' \prec_p \tau'$ and σ' and τ' are not adjacent, we can find a chain $\sigma' \prec_p \sigma_1' \prec_p \cdots \prec_p \sigma_k' \prec_p \tau'$ such that all these cells are consecutively adjacent; for example such a chain can be built by choosing the d-cells crossed by a ray starting from p and crossing both σ' and τ'. For this reason we can simplify our proof and reduce it to the case of two (d-1)-adjacent cells.

Let σ' and τ' be two (d-1)-adjacent cells of $PD(S_\Gamma)$, s_σ and s_τ be the two spheres of S_Γ, such that $cell(s_\sigma) = \sigma'$ and $cell(s_\tau) = \tau'$; the hyperplanes $\Pi(s_\sigma)$ and $\Pi(s_\tau)$ pass through σ^* and τ^*, respectively.

The (d-1)-face f common to σ' and τ' lies on the chordale $chor(s_\sigma, s_\tau)$, since $chor(s_\sigma, s_\tau)$ is the vertical projection of $\Pi(s_\sigma) \cap \Pi(s_\tau)$ onto h_0.

Now consider the obstruction relation between the two cells. For the convexity of σ' and τ', $\sigma' \prec_p \tau'$ if and only if the viewpoint p belongs to the halfspace bounded by the hyperplane passing through the (d-1)-face f, that is $chor(s_\sigma, s_\tau)$; but $chor(s_\sigma, s_\tau)$ partition $\mathbb{E}^d$ in the region where $pow(x, s_\sigma) < pow(x, s_\tau)$ and viceversa. Therefore, $\sigma' \prec_p \tau' \rightarrow pow(p, s_\sigma) < pow(p, s_\tau)$, so the $pow()$ function agrees with the obstruction relation. □

This results suggests a technique for depth sorting a simplicial complex in $\mathbb{E}^d$ that is the projection of the lower part of a convex polyhedron in $\mathbb{E}^{d+1}$. Given the viewpoint p it is sufficient to sort the d-cells σ_i of the complex according to $pow(p, s_{\sigma_i})$, where s_{σ_i} is the sphere obtained by the transformation Π from the plane h_i passing through the d-face σ_i^* of the $\mathbb{E}^{d+1}$ complex Γ^*. Hereafter we will refer this approach to depth sorting as *Power Diagram Depth Sorting* or PDD sorting.

The following results extend the applicability of this sorting technique to non-convex domains.

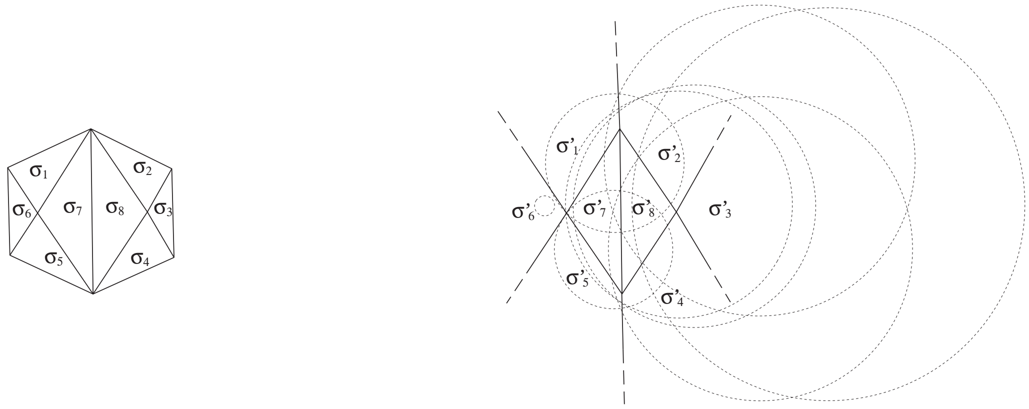

Figure 2: A simplicial complex Γ and the corresponding power diagram in two dimensions.

Corollary 4.1. *The agreement between* $pow()$ *and* $\prec_p$ *holds also for a complex* Γ' *that is a subset of a complex* Γ *that is obtainable as a projection of a convex polyhedron* Γ^*.

The sort is based only on the numerical value of the Φ function; once this function has been calculated we can discard part of the complex Γ without any risk. Thus we can use the PD sorting also for non-convex complexes which are subsets (maybe *carved out*) of large convex complexes. An important class of such complexes is given by alpha shapes [7].

Corollary 4.2. *Let* Γ *be a complex obtainable as a projection of a convex polyhedron* Γ^**, the agreement between* Φ_p *and* $\prec_p$ *holds also for any complex* Γ' *satisfying the following property: for each* d*-cell* $\sigma' \in \Gamma'$ *it is possible to find a* d*-cell* $\sigma \in \Gamma$ *such that* $\sigma' \subseteq \sigma$ *and if* $\sigma', \tau' \in \Gamma'$, $\sigma, \tau \in \Gamma$, $\sigma' \subseteq \sigma$, $\tau' \subseteq \tau$ *then* $\sigma' \neq \tau' \rightarrow \sigma \neq \tau$.

The agreement of Φ_p function with $\prec_p$ relation depends on the intepretation of chordales as separating planes. This separation works also if the cells of the complex are strictly included in the larger power cells of the power diagram. The latter condition of this corollary requires that for each cell of the PD there is at most one smaller cell.

Thus, we can use the PDD sorting for complexes whose d-cells can be seen as *shrinking* of larger cells. The requirements for the $(d+1)$ convex polyhedron that we search are that its orthogonal projection is a complex whose d-cells cover all the cells of our d-complex. An example of this situation can be seen in Figure 3: the grey triangles are strictly included in the cells of the Power Diagram, so we can use the PDD sorting for them. The problem of how effectively find the spheres to build this corresponding PD will be faced in Section 5.

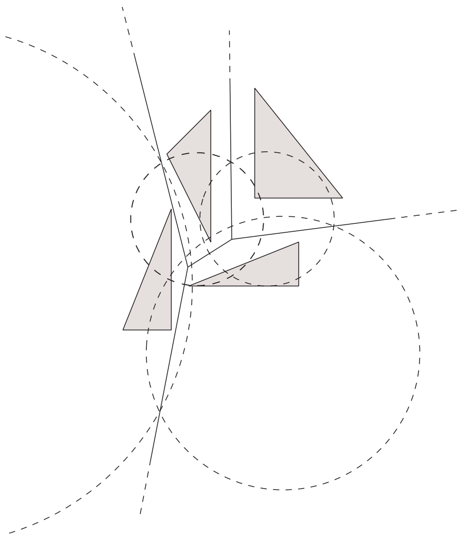

Figure 3: A complex formed by a sparse set of triangles and the corresponding circles and power diagram.

Another consequence of the above corollaries is that this approach can also be used to sort scattered triangles in space: in this case, for each triangle f, it is necessary to find a supporting tetrahedron σ such that f belongs to the faces of σ and σ does not intersect any other triangle. This can be done by adding a vertex sufficiently close to face f. Using this approach we could sort a set of m triangles with a complexity of $O(m \log m)$, with just a linear overhead in storage (the center and radius for each sphere).

An interesting aspect of our approach lies in the *simplicity* of the data structure needed for sorting, once the corresponding polyhedron and power diagram has been found: for each simplex it is only necessary to store its power circle. This property can exploited, for example, in the creation of data structures for secondary memory: applications that need to depth sort just a portion of the whole complex at a time can easily store the corresponding power circle together with each simplex so that each subset of the original complex can be extracted and sorted.

5 Lifting a simplicial complex

The main problem with PDD sorting is that generally we have just a simplicial complex Γ in $\mathbb{E}^d$, and not Γ^*. In some special cases, it is simple to find the convex polyhedron: for example, if Γ is a Delaunay simplicial complex we can exploit the well known correspondence with convex hull in $\mathbb{E}^{d+1}$ to find Γ^*. In this Section, given a generic complex Γ we address the following problem, hereafter denoted as the *lifting problem*: find a convex polyhedron Γ^*, if there exists one, such that Γ is the vertical projection of Γ^*.

Finding an efficient and usable solution to this problem means finding a new approach to the depth sorting problem. The *lifting problem* consists in lifting each vertex v of a complex Γ in $\mathbb{E}^d$ along the $d+1$ axis in order to obtain a convex polyhedron in $\mathbb{E}^{d+1}$. The first remark is that such a problem does not have just one single solution: there exists an infinite number of different convex polyhedra having the same projection onto h_0, and therefore there are infinite sets of spheres that can be used to sort our complex.

We will formulate the lifting problem as a linear programming problem. The lifted complex is convex if and only if the dihedral angle between any two lifted $(d-1)$-adjacent simplex is convex. This condition can be expressed in the following way: let σ and τ be two simplexes that are (d-1)-adjacent through the common (d-1)-face f; let v_σ and v_τ be the two vertices of σ and τ, which do not belong to f; let

$$h_\sigma : x_{d+1} = \mu_{\sigma,1}x_1 + \mu_{\sigma,2}x_2 + \cdots + \mu_{\sigma,d}x_d + k_\sigma$$

be the equation of the non vertical hyperplane in $\mathbb{E}^{d+1}$ passing through σ.

The convexity of the dihedral angle between σ and τ is guaranteed if the projection of v_τ onto h_σ is strictly higher than the $d+1$ coordinate of v_τ:

$$v_{\tau,d+1} > \mu_{\sigma,1}v_{\tau,1} + \mu_{\sigma,2}v_{\tau,2} + \cdots + \mu_{\sigma,d}v_{\tau,d-1} + k_\sigma$$

It can be observed that the coefficients $\mu_{\sigma,i}$ linearly depend only on the vertices of σ thus we can express such linear inequalities with the $d+1$ coordinates of vertices of σ and v_τ as unknowns.

The lifting problem can be expressed as a set of m linear inequalities, with n unknowns, where n is the number of vertices, and m is the number of internal (d-1)-faces of Γ.

A solution of this system can be easily found by using the simplex algorithm; this requires to transform all the strict inequalities by adding a slack constant δ in order to fix the minimum distance between v_τ and h_σ:

$$v_{\tau,d+1} - \mu_{\sigma,1}v_{\tau,1} - \mu_{\sigma,2}v_{\tau,2} - \cdots - \mu_{\sigma,d}v_{\tau,d-1} - k_\sigma \geq \delta$$

and to set as objective function the minimization of the complessive sum of the new (d+1) coordinates of vertices. It should be noted that the coefficient matrix of this LP problem is very large, but fortunately it is very sparse: each line of the system has at most d+2 non zero elements.

If the simplex algorithm does not find a solution then Γ is not a projective complex. It is still an open problem whether this fact implies the existence of a viewpoint p such that the $\prec_p$ relation for Γ is cyclic with respect to p.

Non Convex Simplicial Complexes The technique presented works only if the complex to be lifted is convex; for corollaries 4.1 and 4.2 we know that if we consider a non-convex complex Γ' as a subset of a larger convex polyhedron Γ, that is projection of the lower part of a convex polyhedron Γ^* in $\mathbb{E}^{d+1}$, our sorting method is still applicable, but the lifting technique previously proposed does not work.

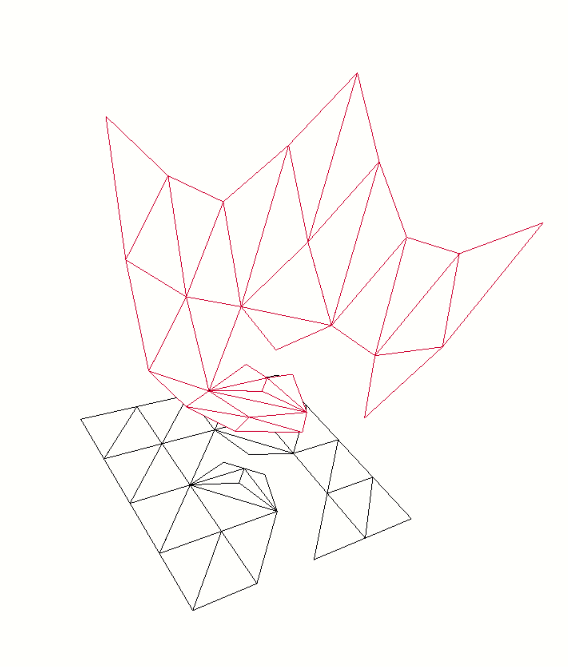

Figure 4: A the lifting of a 2D non-convex 2D complex.

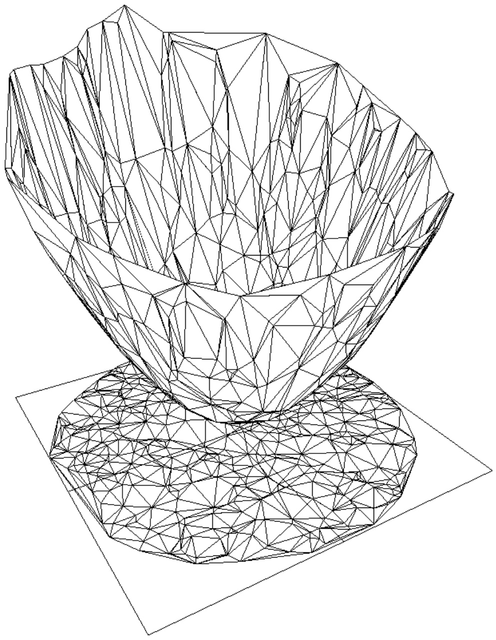

Figure 5: A 2D simplicial complex Γ composed of 1000 cells and the corresponding lifted complex Γ^*.

We can solve this problem including for each (d-1)-face f on the boundary, with f belonging to a simplex σ, the constraint that the halfspace h_σ in $\mathbb{E}^{d+1}$ passing through the lifting of σ, does not contain any other lifted vertex. It should be noted that this approach can increase the number of constraints of the linear programming problem from $O(m)$ to $O(m^2)$. An application of this techinque to the lifting of a non convex two dimensional complex is shown in Figure 4. This technique can also used for building a power diagram for a set of sparse triangles like those shown in Figure 3.

6 Conclusions

We have presented a new approach to the depth sorting problem called Power Diagram Depth Sorting. Given a simplicial complex Γ in $\mathbb{E}^d$ we build the convex polyhedron Γ^* in $\mathbb{E}^{d+1}$ whose projection on the hyperplane h_0 gives back Γ. Then we use Γ^* to find the set of spheres S_Γ such that the power diagram $PD(S_\Gamma)$ is affine, in the sense of Theorem 4.1, to the starting complex Γ. Given a viewpoint, this set of spheres S_Γ can be used to depth sort the original complex Γ in $O(m \log m)$.

We have also shown a technique to effectively build S_Γ starting from Γ based on the transformation of the lifting problem in a LP problem. Some experiments in two dimensions showed that the solution of the linear programming problem generated by a complex of one thousand triangles (see Figure 5) can be found in less than a minute on a small personal computer using a public domain LP solver. This preliminary result appears to be a reasonable preprocessing step, and it allows us to claim the practical relevance of the proposed solution.

We can skectch some open problems whose solution could enhance the effective usability of this technique and the understanding of the relation between convexity and acyclicity:

- What is the lower bound of the lifting problem for convex and non-convex complexes?
- Given a complex Γ there exist many possible different lifted complexes Γ^*, what is the best one according to some measure of tightness for

the spheres to the corresponding PD cells? In other words, if Γ is a Delaunay complex, can we find some lifting technique that returns the set of spheres formed by the circumcircles of the Γ simplices?

- We know that convexity in $\mathbb{E}^{d+1}$ implies the acyclicity of the projected complex; what about the reverse relation?

Acknowledgements This work was partially financed by the Progetto Coordinato "Modelli multirisoluzione per la visualizzazione di campi scalari multidimensionali" and the Progetto Finalizzato "Beni Culturali" of the Italian National Research Council (CNR).

References

[1] P. K. Agarwal, M. J. Katz, and M. Sharir. Computing depth orders and related problems. *Comp. Geom. Theory and Appl.*, 5:187–206, 1995.

[2] F. Aurenhammer. Power diagrams: Properties, algorithms and applications. *Siam J. Comput.*, 16(1):78–96, February 1987.

[3] P. Cignoni, C. Montani, and R. Scopigno. Tetrahedra based volume visualization. In *LNCS (Proc. VisMath '97, Berlin)*, page (in press). Springer Verlag, 1998.

[4] M. de Berg. *Ray Shooting, Depth Orders and Hidden Surface Removal*. Number 703 in Lecture Notes in Computer Science. Springer-Verlag, 1993.

[5] H. Edelsbrunner. *Algorithms in Combinatorial Geometry*. Springer-Verlag, Berlin, 1987.

[6] H. Edelsbrunner. An acyclicity theorem for cell complexes in d dimensions. *Combinatorica*, 10:251–260, 1990.

[7] H. Edelsbrunner and E. P. Mücke. Three-Dimensional alpha shapes. *ACM Transactions on Graphics*, 13(1):43–72, January 1994. ISSN 0730-0301.

[8] L. De Floriani, P. Magillo, and E. Puppo. Geometric structures and algorithms for geographical information systems. In J.R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*. Elsevier Science, 1998.

[9] N. Max, P. Hanrahan, and R. Crawfis. Area and volume coherence for efficient visualization of 3D scalar functions. *Computer Graphics (San Diego Workshop on Volume Visualization)*, 24(5):27–33, November 1990.

[10] M. E. Newell, R. G. Newell, and T. L. Sancha. A new approach to the shaded picture problem. In *Proc. ACM Natl. Conf.*, pages 443–450, 1972.

[11] P. Shirley and A. Tuchman. A polygonal approximation to direct scalar volume rendering. *Computer Graphics (San Diego Workshop on Volume Visualization)*, 24(5):63–70, November 1990.

[12] C. T. Silva, J. Mitchell, and A. Kaufman. Fast rendering of irregular grids. In R. Crawfis and C. Hansen, editors, *Proceedings 1996 Symp. on Volume Visualization (Oct. 28-29)*, pages 15–22, 1996.

[13] C. Stein, B. Becker, and N. Max. Sorting and Hardware Assisted Rendering for Volume Visualization. In *Proceedings of 1994 Symposium on Volume Visualization*, pages 83–90. ACM Press, October 17-18 1994.

[14] P. L. Williams. Visibility ordering of meshed polyhedra. *ACM Transaction on Graphics*, 11(2):103–126, April 1992.

[15] P. L. Williams. *Interactive Direct Volume Rendering of Curvilinear and Unstructured Data*. PhD thesis, University of Illinois at Urbana–Champaign, 1993.