ELSEVIER

Computational Geometry 21 (2002) 87–120

Computational Geometry
Theory and Applications

www.elsevier.com/locate/comgeo

Computing Voronoi skeletons of a 3-D polyhedron by space subdivision

Michal Etzion, Ari Rappoport *

Institute of Computer Science, The Hebrew University, Jerusalem 91904, Israel

Communicated by R. Seidel; received 6 July 1998; received in revised form 15 May 2000; accepted 14 August 2001

Abstract

We tackle the problem of computing the Voronoi diagram of a 3-D polyhedron whose faces are planar. The main difficulty with the computation is that the diagram's edges and vertices are of relatively high algebraic degrees. As a result, previous approaches to the problem have been non-robust, difficult to implement, or not provenly correct.

We introduce three new proximity skeletons related to the Voronoi diagram: (1) the *Voronoi graph* (*VG*), which contains the complete symbolic information of the Voronoi diagram without containing any geometry; (2) the *approximate Voronoi graph* (*AVG*), which deals with degenerate diagrams by collapsing sub-graphs of the VG into single nodes; and (3) the *proximity structure diagram* (*PSD*), which enhances the VG with a geometric approximation of Voronoi elements to any desired accuracy. The new skeletons are important for both theoretical and practical reasons. Many applications that extract the proximity information of the object from its Voronoi diagram can use the Voronoi graphs or the proximity structure diagram instead. In addition, the skeletons can be used as initial structures for a robust and efficient global or local computation of the Voronoi diagram.

We present a space subdivision algorithm to construct the new skeletons, having three main advantages. First, it solves at most uni-variate quartic polynomials. This stands in sharp contrast to previous approaches, which require the solution of a non-linear tri-variate system of equations. Second, the algorithm enables purely local computation of the skeletons in any limited region of interest. Third, the algorithm is simple to implement.

Keywords: Voronoi diagram; Voronoi graph; Voronoi diagram of a polyhedron; Medial axis of a polyhedron

1. Introduction

The Voronoi diagram is a fundamental geometric structure [2,7,12]. We are interested in Voronoi diagrams of 3-D linear polyhedra (i.e., polyhedra whose faces are planar), because they support many important applications in geometric computation [1,13,21]. The Voronoi diagram of an object is closely

* Corresponding author.
E-mail address: arir@cs.huji.ac.il (A. Rappoport).

0925-7721/02/$ – see front matter
PII: S0925-7721(01)00056-6

related to its medial axis. In the case of linear polyhedra, the Voronoi diagram of an object can be easily constructed from its medial axis, and vice versa.

The Voronoi diagram of a non-convex linear polyhedron contains non-linear algebraic entities. Its faces lie on quadratic surfaces, its edges are intersections of two quadratic surfaces, and its vertices are intersections of three quadratic surfaces. The combination of a complex connectivity structure and non-linear geometric elements makes the construction of the Voronoi diagram of a polyhedron a difficult problem. Computing the exact diagram requires solving systems of tri-variate non-linear equations [8, 14,15,18], resulting in algorithms that are not robust, difficult to implement, and difficult to prove correct.

Since construction of the exact geometry of the Voronoi diagram cannot avoid intersecting non-linear 3-D surfaces, several approximate structures have been suggested. Canny and Donald [4] define 'simplified Voronoi diagrams' based on a distance measure that is not a true metric. While this measure is appropriate for robot motion planning, it is not clear whether it can be used for other applications. Sudhalkar et al. [22] proposes the box-skeleton, which uses the maximum norm instead of the Euclidean norm, and therefore does not provide proximity information. Rezayat [16] builds a so-called 'midsurface' of an object, which is only implicitly defined by an algorithm to construct it. The algorithm is heuristic in nature, and user intervention is recommended. Reddy and Turkiyyah [14] construct approximate Voronoi diagrams in the sense that the geometry of the edges and surfaces of the Voronoi diagram is not computed exactly. However, the exact location of the vertices is computed, thus still requiring the computations of non-linear intersections. Milenkovic [11] uses a numeric predicate that identifies vertices without necessarily computing their exact locations, but its convergence is not guaranteed.

Another type of approximate Voronoi diagram of an object is the Voronoi diagram of a set of points on the object's boundary. Bertin and Chassery [3] prove that the Voronoi diagram of such points converges toward the Voronoi diagram of the polyhedron when the step of discretization tends to zero. Etzion [5] constructs a finite set of points on the boundary of a 2-D polygon, whose Voronoi diagram carries the complete symbolic information of the Voronoi diagram of the polygon. Several works [17,23,25] use a Delaunay triangulation of points on the polyhedron's boundary to build the medial axis of the polyhedron. However, the convergence of these algorithms has not been proven.

Lavender et al. [9] use an octree in order to provide an elegant 'black box' to answer proximity queries concerning specific points. For answering such queries, the method is general, easy to implement, and very practical. However, it does not provide any information regarding the symbolic structure of the Voronoi diagram, hence is not suitable for skeletal shape analysis. Vleugels and Overmars [24] also use a space subdivision to construct a geometric approximation of the Voronoi diagram of a set of disjoint convex sites. The symbolic information analyzed is limited to the connectivity of the Voronoi diagram; the different Voronoi elements are not identified.

Contribution

In this paper we introduce a new approach for dealing with non-linear Voronoi diagrams, based on computing their symbolic and geometric parts separately. We use the term *Voronoi Graph* (*VG*) to describe the symbolic part. We present a simple space subdivision algorithm for computing the Voronoi graph of a 3-D linear polyhedron. The algorithm constructs a *Proximity Structure Subdivision*, a subdivision whose cells are labeled according to relative proximities to polyhedron entities. The Voronoi graph is constructed from the subdivision in three stages: computing witnesses of Voronoi edges, using them to identify Voronoi vertices, and finally determining the connectivity structure. The algorithm utilizes only distance comparisons and 2-D geometric computations, the most complex of which is intersecting two conic sections. The algorithm has been implemented.

To tackle degeneracies, we define and compute the *Approximate Voronoi Graph* (*AVG*), in which degenerate and almost-degenerate parts of the Voronoi graph are identified and simplified. The space subdivision allows us to also compute a well-defined approximation to the geometric part of the Voronoi diagram to any desired accuracy. We refer to this type of approximate Voronoi diagram as a *Proximity Structure Diagram* (*PSD*). Computation of the PSD is very stable, since it does not involve symbolic decisions, and it utilizes the same simple geometric operations used in the computation of the Voronoi graph.

The algorithm has several important advantages over previous approaches. First, it utilizes only relatively simple 2-D geometric computations, thus avoiding complex and unstable intersections of 3-D surfaces. Second, all three proximity skeletons can be computed locally, in a given spatial region of interest. Third, the algorithm allows purely local computation of *partial* information contained in the skeletons, such as the identities and approximate locations of Voronoi vertices or edges, and it does so efficiently without requiring global curve tracing. Finally, its correctness has been formally proven.

The proximity skeletons we introduce are important by themselves for several reasons. First, they preserve proximity information, unlike approximations that use a different metric. Second, many applications that currently compute the Voronoi diagram or medial axis are actually only interested in partial proximity information present in the VG, AVG or PSD. Third, these skeletons can be used in order to efficiently identify regions of interest in which more detailed information is needed. Finally, the skeletons constitute initial structures for robust and efficient computation of the Voronoi diagram.

The paper is organized as follows. In Section 2 we formally define the Voronoi graph, and provide notations and basic definitions. In Section 3 we discuss properties of the Voronoi diagram and of the point sets used to define it. In Section 4 we define the proximity structure subdivision and give an algorithm for constructing it. In Section 5 we describe how the Voronoi graph is constructed from the subdivision. In Sections 6 and 7 we define the two other proximity skeletons and describe their construction. For clarity of exposition, in Sections 4 and 5 we assume that the Voronoi diagram of the polyhedron is not degenerate. Handling of degenerate Voronoi diagrams is done in Section 6. A detailed proof for the fact that Voronoi edges are 1-manifold curves is given in Appendix A. The discussion in Section 8 includes a description of a single minor configuration for which the proof of correctness of our algorithm has not been completed.

2. Definitions and notations

Let Q be a bounded 3-D linear polyhedron having a 2-manifold connected boundary composed of convex faces [10]. The requirement that Q has convex faces does not limit the range of polyhedra. For any polyhedron Q, we can decompose its faces into convex pieces, compute the Voronoi diagram (or Voronoi graph or proximity structure diagram) of the resulting polyhedron Q', and then easily obtain the Voronoi diagram of Q from the Voronoi diagram of Q' (see Section 8).

The *entities* of Q are the vertices, edges and faces of Q, and are denoted by lower-case letters a, b, c. The entities are closed sets, i.e., an edge contains its vertices, and a face contains its edges and vertices.

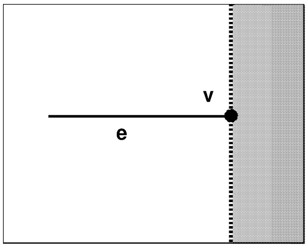

Fig. 1. A 2-D example: v is a vertex incident on edge e. If *CloserEq* is defined in the standard way, then $\mathit{CloserEq}(v,e) \cap \mathit{CloserEq}(e,v)$ is the 2-D gray region. If *CloserEq* is defined as in this paper, then $\mathit{CloserEq}(v,e) \cap \mathit{CloserEq}(e,v)$ is the dotted line, which is a 1-D region.

For two entities a and b, we say that $a \subset b$ (or $a \subseteq b$) if the point set of a is a proper subset (or subset) of the point set of b.

$d(x,y)$ denotes the distance of two points as well as the distance between a point and an entity. The distance between a point x and an entity a is defined as $\inf_{y\in a} d(x,y)$. For a point x, $B(x,r)$ denotes the locus of points y s.t. $d(x,y) < r$. For two points y,z, $[y,z]$ denotes the locus of points x s.t. $x = ty + (1-t)z$ for $0 \leqslant t \leqslant 1$, and (y,z) denotes the locus of points x s.t. $x = ty + (1-t)z$ for $0 < t < 1$. For a point set A, ∂A denotes the boundary of A, $\mathit{int}(A)$ denotes the interior of A, and $\mathit{cl}(A)$ denotes the closure of A. ∂A, $\mathit{int}(A)$ and $\mathit{cl}(A)$ are defined relative to the affine hull of A. $\mathit{dim}(A)$ denotes the dimension of the affine hull of A.

$\pi_a(x)$ denotes the *projection* of a point x on an entity a, i.e., the point on a nearest to x. $\pi_a(x)$ is a single point, since a is either a vertex or an edge or a convex face. A *footpoint* of a point x on a polyhedron Q is a point y s.t. $d(x,y) \leqslant d(x,z)$ for every point $z \in Q$. The *carrier* of an edge (face) is the infinite line (plane) containing the entity, i.e., it is the affine hull of the entity. The carrier of a vertex is the vertex itself. The carrier of an entity a is denoted by $\mathit{car}(a)$. Sets of entities are denoted by lower-case Greek letters α, β, γ. $\alpha\star$ denotes a set of entities containing α. $|\alpha|$ denotes the number of entities in α.

Let a and b be two entities. We would have liked to use the following standard definitions for the point sets $\mathit{Closer}(a,b)$ and $\mathit{CloserEq}(a,b)$: $\mathit{Closer}(a,b) = \{x \mid d(x,a) < d(x,b)\}$ and $\mathit{CloserEq}(a,b) = \{x \mid d(x,a) \leqslant d(x,b)\}$. However, if a and b intersect each other, then $\mathit{CloserEq}(a,b) \cap \mathit{CloserEq}(b,a)$ might be a 3-D region (a 2-D example is shown in Fig. 1).

In order to ensure that Voronoi faces are two-dimensional, we define $\mathit{Closer}(a,b)$ and $\mathit{CloserEq}(a,b)$ as follows. If $a \cap b = \emptyset$ or $a \subset b$, then $\mathit{CloserEq}(a,b) = \{x \mid d(x,a) \leqslant d(x,b)\}$ and $\mathit{Closer}(a,b) = \mathit{int}(\mathit{CloserEq}(a,b))$. Otherwise, $\mathit{Closer}(a,b) = \{x \mid d(x,a) < d(x,b)\}$ and $\mathit{CloserEq}(a,b) = \mathit{cl}(\mathit{Closer}(a,b))$. In addition we define $\mathit{Closer}(a,a) = \emptyset$ and $\mathit{CloserEq}(a,a) = \Re^3$. In Section 3 we study the properties of the $\mathit{Closer}(a,b)$ and $\mathit{CloserEq}(a,b)$ sets.

Let α be a set of entities. The *bisector* of α is $\mathit{bis}(\alpha) = \bigcap_{a,b\in\alpha} \mathit{CloserEq}(a,b)$. The bisector of the carriers of α is $\mathit{carbis}(\alpha) = \{x \mid \forall_{a,b\in\alpha} d(x,\mathit{car}(a)) = d(x,\mathit{car}(b))\}$. The *Voronoi region* of α is $R_\alpha = \bigcap_{a\in\alpha, b\in Q} \mathit{CloserEq}(a,b)$. If a point $x \in R_\alpha$, then we say that the entities in α are the *governors* of the point. Note that for every set of entities α, $R_\alpha \subseteq \mathit{bis}(\alpha)$.

The boundaries of the Voronoi regions R_α for $|\alpha| = 1$ comprise the *Voronoi diagram* of Q, $\mathit{VD}(Q)$. A point x on $\mathit{VD}(Q)$ satisfies that there exists a set of entities α whose size is greater than 1, s.t. $x \in R_\alpha$. For a specific set of entities α, consider a maximal connected region R in R_α s.t. $R \not\subset R_\beta$ for any $\beta \supset \alpha$.

If the region is a surface, then it is a *face* f_α of $VD(Q)$. If the region is a curve, then it is an *edge* e_α of $VD(Q)$. If the region is a point, then it is a *vertex* v_α of $VD(Q)$.

The *medial axis* of Q, $MA(Q)$, is the locus of points in $\Re^3$ having more than one footpoint on the boundary of Q.

The Voronoi graph

The Voronoi diagram of Q defines a labeled graph whose nodes are the elements (vertices, edges and faces) of the diagram, and whose arcs connect elements that are co-incident. Every node of the graph is labeled by the governors of the corresponding Voronoi element. We call this graph the Voronoi graph of Q, which is formally defined as follows.

Let G be an undirected graph such that every node is labeled by: (1) a set of entities of Q, (2) type: `face`, `edge` or `vertex`. G is a *Voronoi Graph* of Q if there exists a bijection F from the set of nodes of G to the set of elements of $VD(Q)$ such that: (1) For every node $n \in G$, if type of n is `face` then $F(n)$ is a Voronoi face. Similarly for types `edge` and `vertex`. (2) For every node $n \in G$, if the set of entities of n is α, then $F(n)$ is governed by α. (3) n_1 and n_2 share an arc in G iff there is an incidence relationship between $F(n_1)$ and $F(n_2)$ in $VD(Q)$.

We say that the Voronoi graph contains all the symbolic information present in the Voronoi diagram; it does not contain any geometry.

3. Properties of the Voronoi diagram

In this section we study the properties of the point sets and structures defined in the previous section. Lemmas 1–2 are auxiliary lemmas. Lemmas 3–9 give properties of the pointsets $Closer(a,b)$, $CloserEq(a,b)$, R_α, $bis(\alpha)$, $carbis(\alpha)$. Lemmas 10–14 give properties of $VD(Q)$. The proofs of Lemmas 1–4 are simple and therefore omitted.

Lemma 1 (The triangle inequality between two points and an entity). *Let a be an entity. Let x, y be two points.* (1) $d(x,a) \leqslant d(x,y) + d(y,a)$. (2) *If $d(x,a) = d(x,y) + d(y,a)$, then there exists a point z s.t. $z = \pi_a(y) = \pi_a(x)$ and $y \in [x,z]$.*

Lemma 2 (The conditions in which the interior of $\{x \mid d(x,a) = d(x,b)\}$ is empty). *Let a and b be two entities. Let x be a point s.t. $d(x,a) = d(x,b)$ and there does not exist a point z s.t. $z = \pi_b(x) = \pi_a(x)$. For every $\varepsilon > 0$ there exists a point $y \in B(x,\varepsilon)$ s.t. $d(y,a) > d(y,b)$.*

Throughout this section we will use the table of Fig. 2. The table is implied from the definitions of *Closer* and *CloserEq* together with Lemma 2.

Lemma 3 (Basic properties of *Closer* and *CloserEq*). *Let a, b be two entities.*
1. *$Closer(a,b) \subseteq CloserEq(a,b)$.*
2. *$Closer(a,b)$ is an open set.*
3. *$CloserEq(a,b)$ is a closed set.*
4. *$\Re^3 \setminus Closer(a,b)$ is connected and unbounded.*

	$Closer(a,b)$	$CloserEq(a,b)$
$a=b$	$\emptyset$	$\Re^3$
$a\cap b=\emptyset$	$d(x,a)<d(x,b)$	$d(x,a)\leqslant d(x,b)$
$a\subset b$	$int(d(x,a)=d(x,b))$	$d(x,a)=d(x,b)$
$b\subset a$	$d(x,a)<d(x,b)$	$cl(d(x,a)<d(x,b))$
$a\cap b=c\neq a,b$	$d(x,a)<d(x,b)$	$cl(d(x,a)<d(x,b))$

Fig. 2. The point sets $Closer(a,b)$ and $CloserEq(a,b)$.

Lemma 4 (The relationship between $Closer(a,b)$ and $CloserEq(b,a)$). *Let a,b be two entities.*
1. *If $a=b$ or $a\subset b$ or $b\subset a$ or $a\cap b=\emptyset$, then $\Re^3 = Closer(a,b)\cup CloserEq(b,a)$.*
2. *$Closer(a,b)\cap CloserEq(b,a)=\emptyset$.*

Lemma 5 (*Closer* and *CloserEq* of co-incident entities). *Let a,b be two entities s.t. $b\subseteq a$. $d(x,a)=d(x,b)=d(x,car(b))$ iff $x\in CloserEq(b,a)\setminus\bigcup_{c\subset b} Closer(c,a)$.*

Proof. Consider the three cases:
1. a is a vertex. Then $b=a$, and it is clear.
2. a is an edge. If b is a vertex then $d(x,a)=d(x,car(b))\Leftrightarrow d(x,a)=d(x,b)\Leftrightarrow x\in CloserEq(b,a)$. If $b=a$, then $d(x,a)=d(x,car(a))\Leftrightarrow \pi_{car(a)}(x)\in a\Leftrightarrow$ for every $c\subset a$ and for every $\varepsilon>0$ there exists a point y s.t. $d(x,y)<\varepsilon$ and $d(y,a)<d(y,c)\Leftrightarrow x\notin Closer(c,a)$ for every $c\subset a$.
3. a is a face. If b is a vertex then $d(x,a)=d(x,car(b))\Leftrightarrow d(x,a)=d(x,b)\Leftrightarrow x\in CloserEq(b,a)$. If b is an edge then $d(x,a)=d(x,b)=d(x,car(b))\Leftrightarrow x\notin Closer(c,b)$ for every $c\subset b$ and $x\in CloserEq(b,a)$. If $b=a$ then $d(x,a)=d(x,car(a))\Leftrightarrow \pi_{car(a)}(x)\in a\Leftrightarrow$ for every ε there exists a point y s.t. $d(x,y)<\varepsilon$ and $d(y,a)<d(y,c)$ for every $c\subset a\Leftrightarrow x\notin Closer(c,a)$ for every $c\subset a$. □

Lemma 6 (Properties of $bis(a,b)$). *Let a,b,c be three entities.*
1. *$dim(bis(a,b))\leqslant 2$.*
2. *Let a and b be two entities s.t. $a\cap b=c\neq a,b$. Let x be a point s.t. $\pi_{car(a)}(x)\in a$ and $\pi_{car(b)}(x)\in b$. If $x\in bis(a,c)\cap bis(b,c)$ then $x\in bis(a,b)$.*
3. *If $x\in carbis(a,b)$, $\pi_{car(a)}(x)\in a$ and $\pi_{car(b)}(x)\in b$, then $x\in bis(a,b)$.*

Proof.
1. If $x\in bis(a,b)$ then $x\in CloserEq(a,b)\cap CloserEq(b,a)$. Lemma 4.2 implies that $x\in CloserEq(a,b)\setminus Closer(a,b)$. The definitions of *CloserEq* and *Closer* imply that the dimension of the locus of points $\{x|x\in CloserEq(a,b)\setminus Closer(a,b)\}$ is not greater than 2.
2. We show in the following that for every $\varepsilon>0$ there exist points y_1,y_2 s.t. $d(x,y_1)<\varepsilon$, $d(x,y_2)<\varepsilon$, $d(y_1,a)<d(y_1,b)$ and $d(y_2,b)<d(y_2,a)$. This implies that $x\in bis(a,b)$. Consider the following cases:
   (a) a and b are edges, and c is a vertex. Let P be the plane of a and b. $x\in bis(a,c)$, and therefore x is on the plane orthogonal to a at c. $x\in bis(b,c)$, and therefore x is on the plane orthogonal to b at c. If a and b are not colinear, then these planes intersect in a line l orthogonal to P at c.

$x \in l$, and therefore for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $\pi_P(y_1) \in int(a)$ and $\pi_P(y_2) \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$. If a and b are colinear on the line l', then x is on the plane orthogonal to l' at c. Therefore for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $\pi_{l'}(y_1) \in int(a)$ and $\pi_{l'}(y_2) \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$.

(b) a and b are faces, and c is a vertex. $x \in bis(a, c)$ and satisfies that $\pi_{car(a)}(x) \in a$. Therefore x is on the line orthogonal to $car(a)$ at c. Similarly x is on the line orthogonal to $car(b)$ at c. If $car(a) \neq car(b)$, then these lines intersect in c. Therefore $x = c$. In this case for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $y_1 \in int(a)$ and $y_2 \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$. If $car(a) = car(b) = P$, then x is on the line orthogonal to P at c. In this case for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $\pi_P(y_1) \in int(a)$ and $\pi_P(y_2) \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$.

(c) a and b are faces, and c is an edge. $x \in bis(a, c)$ and satisfies that $\pi_{car(a)}(x) \in a$. Therefore $\pi_{car(a)}(x) \in c$. Similarly $\pi_{car(b)}(x) \in c$. If $car(a) \neq car(b)$, then $x \in c$. In this case for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $y_1 \in int(a)$ and $y_2 \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$. If $car(a) = car(b) = P$, then $\pi_P(x) \in c$. In this case for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $\pi_P(y_1) \in int(a)$ and $\pi_P(y_2) \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$.

(d) a is a face, b is an edge, and c is a vertex. $x \in bis(a, c)$ and satisfies that $\pi_{car(a)}(x) \in a$. Therefore x is on the line l orthogonal to $car(a)$ at c. $x \in bis(b, c)$ and therefore is on the plane P orthogonal to b at c. If $l \not\subset P$ then $l \cap P = c$. In this case there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $y_1 \in int(a)$ and $y_2 \in int(b)$. Therefore $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$. If $l \subset P$ then a and b share a plane Q. In this case for every $\varepsilon > 0$ there exist points y_1, y_2 s.t. $d(x, y_i) < \varepsilon$, $\pi_Q(y_1) \in int(a)$ and $\pi_Q(y_2) \in int(b)$. $d(y_1, a) < d(y_1, b)$ and $d(y_2, b) < d(y_2, a)$.

3. We show in the following that $x \in CloserEq(a, b)$. $\pi_{car(a)}(x) \in a$ therefore $d(x, car(a)) = d(x, a)$. Similarly $d(x, car(b)) = d(x, b)$. Therefore $d(x, a) = d(x, b)$. Suppose on the contrary $x \notin CloserEq(a, b)$. Then $b \subset a$ or $b \cap a = d \neq a, b$, and there exists an $\varepsilon > 0$ s.t. if $y \in B(x, \varepsilon)$, then $d(y, a) \geqslant d(y, b)$. Consider the two cases:
   (a) $b \subset a$. $\pi_{car(a)}(x) \in a$. Therefore (Lemma 5) $x \notin Closer(b, a)$. Contradiction (Lemma 4.1).
   (b) $b \cap a = d \neq a, b$. $\pi_{car(a)}(x) \in a$. Therefore (Lemma 5) $x \notin Closer(d, a)$. Therefore $x \in CloserEq(a, d)$ (Lemma 4.1). Similarly $x \in CloserEq(b, d)$. Lemma 2 implies that $\pi_a(x) = \pi_b(x)$, and therefore $d(x, a) = d(x, b) = d(x, d)$. Therefore $x \in CloserEq(d, a) \cap CloserEq(d, b)$. Therefore $x \in bis(a, d) \cap bis(b, d)$. Lemma 6.2 implies that $x \in bis(a, b)$. Contradiction. □

Lemma 7 (Transitivity of *Closer* and *CloserEq*). *Let a, b, c be three entities.*
1. $Closer(a, b) \cap Closer(b, c) \subseteq Closer(a, c)$.
2. $CloserEq(a, b) \cap Closer(b, c) \subseteq CloserEq(a, c)$.
3. *Let x be a point s.t. $\pi_{car(a)}(x) \in a$. If $x \in CloserEq(a, b) \cap CloserEq(b, c)$ then $x \in CloserEq(a, c)$.*

Proof.
1. If $a = c$ then Lemma 4.2 implies that $Closer(a, b) \cap Closer(b, c) = \emptyset$. If $a \neq c$ let $x \in Closer(a, b) \cap Closer(b, c)$. $d(x, a) \leqslant d(x, b)$ and $d(x, b) \leqslant d(x, c)$. If $d(x, a) < d(x, b)$ or $d(x, b) < d(x, c)$ then we are done. Otherwise $d(x, a) = d(x, b)$ and $d(x, b) = d(x, c)$. $x \in Closer(a, b)$, therefore $a \subset b$, and there exists an $\varepsilon > 0$ s.t. every $y \in B(x, \varepsilon)$ satisfies that $d(y, a) = d(y, b)$. $x \in Closer(b, c)$, therefore $b \subset c$, and there exists an $\varepsilon > 0$ s.t. every $y \in B(x, \varepsilon)$ satisfies that $d(y, b) = d(y, c)$.

Therefore $a \subset c$, and there exists an $\varepsilon > 0$ s.t. every $y \in B(x,\varepsilon)$ satisfies that $d(y,a) = d(y,c)$. Therefore $x \in \mathit{Closer}(a,c)$.

2. If $a = b$, then it is implied from Lemma 3.1. Let $x \in \mathit{CloserEq}(a,b) \cap \mathit{Closer}(b,c)$. $d(x,a) \leqslant d(x,b)$ and $d(x,b) \leqslant d(x,c)$. Suppose on the contrary $x \notin \mathit{CloserEq}(a,c)$. Then (1) $d(x,a) = d(x,b) = d(x,c)$, (2) $c \subset a$, or $c \cap a = d \neq a, c$, (3) $b \subset c$, and (4) there exists an $\varepsilon > 0$ s.t. every $y \in B(x,\varepsilon)$ satisfies that $d(y,a) \geqslant d(y,c)$. $x \in \mathit{Closer}(b,c)$, therefore there exists an $\varepsilon > 0$ s.t. every $y \in B(x,\varepsilon)$ satisfies that $d(y,b) \leqslant d(y,c)$. Therefore there exists an $\varepsilon > 0$ s.t. every $y \in B(x,\varepsilon)$ satisfies that $d(y,a) \geqslant d(y,b)$. (2) and (3) imply that $a \not\subset b$, and therefore if $x \in \mathit{CloserEq}(a,b)$ then for every $\varepsilon > 0$ there is a point y s.t. $d(x,y) < \varepsilon$, and $d(y,a) < d(y,b)$. Contradiction.
3. $x \in \mathit{CloserEq}(a,b)$ therefore $d(x,a) \leqslant d(x,b)$. $x \in \mathit{CloserEq}(b,c)$ therefore $d(x,b) \leqslant d(x,c)$. Therefore $d(x,a) \leqslant d(x,c)$. If $d(x,a) < d(x,c)$ we are done. Otherwise $d(x,a) = d(x,b) = d(x,c)$. Suppose on the contrary $x \notin \mathit{CloserEq}(a,c)$. Then (1) $c \subset a$ or $a \cap c = d \neq a, c$ and (2) there is an $\varepsilon > 0$ s.t. if $y \in B(x,\varepsilon)$ then $d(y,a) \geqslant d(y,c)$. Consider the two cases:
   (a) $c \subset a$. Then $x \in \mathit{Closer}(c,a)$ (Lemma 4.1). Then $\pi_{\mathit{car}(a)}(x) \notin a$ (Lemma 5). Contradiction.
   (b) $a \cap c = d \neq a, c$. The existence of $B(x,\varepsilon)$ implies that $d(x,a) = d(x,c) = d(x,d)$ (Lemma 2). Therefore $x \in \mathit{CloserEq}(d,a) \cap \mathit{CloserEq}(d,c)$. $\pi_{\mathit{car}(a)}(x) \in a$, therefore $x \in \mathit{CloserEq}(a,d)$ (Lemma 5), therefore $x \in \mathit{bis}(a,d)$. If $\pi_{\mathit{car}(c)}(x) \in c$, then $x \in \mathit{CloserEq}(c,d)$ and $x \in \mathit{bis}(c,d)$. In this case Lemma 6.2 implies that $x \in \mathit{CloserEq}(a,c)$, and contradiction. If $\pi_{\mathit{car}(c)}(x) \notin c$, then $x \in \mathit{Closer}(d,c)$. In this case Lemma 7.2 implies that $x \in \mathit{CloserEq}(a,c)$, and contradiction. □

Lemma 8 (Properties of R_α). *Let α be a set of entities.*
1. *R_α is a closed set.*
2. *$R_\alpha \subseteq \mathit{carbis}(\alpha)$.*
3. *If $x \in R_\alpha$ and $b \notin \alpha$, then there exists an entity $a \in \alpha$ s.t. $x \in \mathit{Closer}(a,b)$.*
4. *If $x \in \partial R_\alpha$ in the relative topology of $\mathit{carbis}(\alpha)$, and $\dim(\mathit{carbis}(\alpha)) > 0$, then $x \in R_\beta$ for $\beta \supset \alpha$.*
5. *$\dim(R_\alpha) = \dim(\mathit{carbis}(\alpha))$.*

Proof.
1. Finite intersection of closed sets is a closed set.
2. Let $x \in R_\alpha$. Let a, b be two entities in α. $x \in \mathit{bis}(a,b)$. Therefore $d(x,a) = d(x,b)$. If $d(x,a) \neq d(x,\mathit{car}(a))$, then there exits $a' \subset a$ s.t. $x \in \mathit{Closer}(a',a)$ (Lemma 5). Then $x \notin \mathit{CloserEq}(a,a')$ (Lemma 4.1) in contradiction to being x in R_α. Therefore $d(x,\mathit{car}(a)) = d(x,a) = d(x,b) = d(x,\mathit{car}(b))$.
3. We first show that if $b \notin \alpha$, then there exists an entity e s.t. $x \in \mathit{Closer}(e,b)$. Then we show that this implies that exists an entity $a \in \alpha$ s.t. $x \in \mathit{Closer}(a,b)$.
   Suppose on the contrary that $x \notin \mathit{Closer}(e,b)$ for any entity e. If $\pi_x(\mathit{car}(b)) \notin b$, then there exists an entity $e \subset b$ s.t. $x \in \mathit{Closer}(e,b)$ (Lemma 5), and contradiction. Therefore $\pi_x(\mathit{car}(b)) \in b$. $b \notin \alpha$, therefore there exists an entity e s.t. $x \notin \mathit{CloserEq}(b,e)$. $b \cap e = c \neq b, e$ (Lemma 4.1). $x \notin \mathit{Closer}(e,b)$ therefore $d(x,e) \geqslant d(x,b)$. $x \notin \mathit{CloserEq}(b,e)$ therefore there exists an $\varepsilon > 0$ s.t. every $y \in B(x,\varepsilon)$ satisfies that $d(y,e) \leqslant d(y,b)$. Therefore (Lemma 2) $d(x,e) = d(x,b) = d(x,c)$, and $x \in \mathit{CloserEq}(c,e)$. If $x \in \mathit{CloserEq}(b,c)$ then $x \in \mathit{CloserEq}(b,e)$ (Lemma 7.3) and contradiction. Therefore $x \in \mathit{Closer}(c,b)$ (Lemma 4.1). Contradiction.
   Suppose on the contrary that there does not exist an entity $a \in \alpha$ s.t. $x \in \mathit{Closer}(a,b)$. We have shown that there exists an entity e_1 s.t. $x \in \mathit{Closer}(e_1,b)$. $e_1 \notin \alpha$, therefore there exists an entity

e_2 s.t. $x \in Closer(e_2, e_1)$. $x \in Closer(e_2, b)$ (Lemma 7.1). Therefore there exists an infinite sequence of entities $\{e_i\}$ s.t. $x \in Closer(e_j, e_i)$ for any $j > i$. Contradiction.

4. Let b be a governor of a neighborhood of x in $carbis(\alpha) \setminus R_\alpha$. $x \in R_b$ (Lemma 8.1). If $b \notin \alpha$, then we are done. Otherwise $b \in \alpha$. Let y be a point in this neighborhood. We show in the following that there exists an entity $a \in \alpha$ s.t. $\pi_{car(a)}(y) \notin a$.
   Suppose on the contrary that for every $a \in \alpha$ $\pi_{car(a)}(y) \in a$. Then for every $a \in \alpha$, $y \in bis(a, b)$ (Lemma 6.3). Then $y \in R_\alpha$ (Lemma 7.3), and contradiction.
   $\pi_{car(a)}(y) \notin a$, therefore there exists an entity $a' \subset a$ s.t. $y \in Closer(a', a)$ and $\pi_{car(a')}(y) \in a'$ (Lemma 5). $x \in CloserEq(a', a)$ (Lemma 8.1). Therefore $x \in R_{a'}$ (Lemma 7.3). If $a' \notin \alpha$, then we are done. Otherwise $a' \in \alpha$. Then $d(y, a) > d(y, car(a)) = d(y, car(a')) = d(y, a')$. Contradiction, since $a' \subset a$.
5. Implied from Lemma 8.2 and Lemma 8.4. □

Lemma 9 (Starness of R_a). *If $x \in R_a$ then $[x, \pi_a(x)] \subseteq R_a$.*

Proof. $x \in R_a$ therefore $\pi_{car(a)}(x) = \pi_a(x)$ (Lemma 5), and for every $e \in Q$ $x \in CloserEq(a, e)$. Let y be a point in $[x, \pi_a(x)]$. We have to show that $y \in CloserEq(a, e)$. $d(x, a) - d(y, a) = d(x, y)$. By Lemma 1 $d(x, y) \geqslant d(x, e) - d(y, e)$. These two equations imply that $d(x, a) - d(y, a) \geqslant d(x, e) - d(y, e)$. $x \in CloserEq(a, e)$ and therefore $d(x, a) \leqslant d(x, e)$. The last two equations imply that $d(y, a) \leqslant d(y, e)$. Consider the following cases:
1. $a \cap e = \emptyset$ or $a \subset e$. The fact that $d(y, a) \leqslant d(y, e)$ implies that $y \in CloserEq(a, e)$.
2. $a \supset e$. The fact that $x \in CloserEq(a, e)$ implies that $y \in CloserEq(a, e)$.
3. $a \cap e = b \neq a, e$. If $y \notin CloserEq(a, e)$, the fact that $d(y, a) \leqslant d(y, e)$ implies that $d(y, a) = d(y, e) = d(y, b)$ (Lemma 2). Therefore $y \in CloserEq(b, e)$. The fact that $x \in CloserEq(a, b)$ implies that $y \in CloserEq(a, b)$. Therefore $y \in CloserEq(a, e)$ (Lemma 7.3). □

Lemma 10 (The endpoint of a Voronoi edge (face) is a Voronoi vertex (edge)). *Let α be a set of entities of the polyhedron Q.*
1. *Let e_α be an edge of $VD(Q)$. If x is a point on ∂e_α in the relative topology of $carbis(\alpha)$, then x is a vertex v_β of $VD(Q)$ s.t. $\alpha \subset \beta$.*
2. *Let f_α be a face of $VD(Q)$. If x is a point on ∂f_α in the relative topology of $carbis(\alpha)$, then x is on an edge e_β of $VD(Q)$ s.t. $\alpha \subset \beta$.*

Proof. Implied from Lemma 8.4. □

Lemma 11 (A lower bound to the number of governors of a Voronoi element).
1. *If f_α is a Voronoi face, then $|\alpha| \geqslant 2$.*
2. *If e_α is a Voronoi edge, then $|\alpha| \geqslant 3$.*
3. *If v_α is a Voronoi vertex, then $|\alpha| \geqslant 4$.*

Proof. If α contains one entity, then $carbis(\alpha) = \Re^3$. Therefore if f_α is a Voronoi face, then $|\alpha| \geqslant 2$ (Lemma 8.5). Item 2 and item 3 are implied from item 1 by Lemma 10. □

Lemma 12 (The relationship between the Voronoi diagram and the medial axis). *For a set of entities α define $E(\alpha) = \alpha \setminus \{a\colon\ a \supset b, b \in \alpha\}$. $MA(Q) = VD(Q) \setminus \bigcup\{R_\alpha\colon\ |E(\alpha)| = 1\}$.*

Proof.

1. $MA(Q) \subset VD(Q) \setminus \bigcup\{R_\alpha\colon\ |E(\alpha)| = 1\}$. Let $x \in MA(Q)$. First we show that $x \in VD(Q)$. Let $p_1, \ldots, p_n$ be the footpoints of x on ∂Q. $x \in MA(Q)$, therefore $n \geqslant 2$. Let a_i be the entity p_i is incident on. If a point p_i is incident on more than one entity, then we take the lowest dimensional among these entities. Let $\alpha = \{a_1, \ldots, a_n\}$. In order to show that $x \in VD(Q)$, it is enough to show that $x \in R_\alpha$ since $|\alpha| \geqslant 2$. We have to show that $x \in \mathit{CloserEq}(a_i, b)$ for every $a_i \in \alpha$ and $b \in Q$. $d(x, a_i) \leqslant d(x, b)$ since $d(x, p_i) \leqslant d(x, q)$ for every $q \in \partial Q$. Consider the following cases:
   (a) $a_i \cap b = \emptyset$ or $a_i \subset b$. Then $d(x, a_i) \leqslant d(x, b)$ implies that $x \in \mathit{CloserEq}(a_i, b)$.
   (b) $b \subset a_i$. If $x \notin \mathit{CloserEq}(a_i, b)$ then $d(x, a_i) = d(x, b)$. In this case $p_i \in b$, and $b \in \alpha$. Therefore $a_i \notin \alpha$. Contradiction.
   (c) $a_i \cap b = c \neq a_i, b$. If $x \notin \mathit{CloserEq}(a_i, b)$ then $d(x, a_i) = d(x, b)$ and there exists an $\varepsilon > 0$ s.t. every $y \in B(x, \varepsilon)$ satisfies that $d(y, a_i) \geqslant d(y, b)$. Therefore $\pi_{a_i}(x) = \pi_b(x)$ (Lemma 2), and $d(x, a_i) = d(x, b) = d(x, c)$. Therefore $x \in \mathit{CloserEq}(c, b)$. The previous item implies that $x \in \mathit{CloserEq}(a_i, c)$. $\pi_{car(a_i)}(x) \in a_i$, since otherwise $x \in \mathit{Closer}(d, a_i)$ for some $d \subset a_i$ (Lemma 8), in contradiction to previous item. Therefore $x \in \mathit{CloserEq}(a_i, b)$ (Lemma 7.3).

   Now we show that $|E(\alpha)| \geqslant 2$. It is enough to show that $E(\alpha) = \alpha$, since $|\alpha| \geqslant 2$. Suppose on the contrary there is an entity $a_i \in \alpha \setminus E(\alpha)$. Then there exists an entity $b \in \alpha$ s.t. $b \subset a_i$. $p_i \in b$ therefore $a_i \notin \alpha$, contradiction.
2. $MA(Q) \supset VD(Q) \setminus \bigcup\{R_\alpha\colon\ |E(\alpha)| = 1\}$. Let $x \in VD(Q) \setminus \bigcup\{R_\alpha\colon\ |E(\alpha)| = 1\}$. Let α be a set of entities s.t. $x \in R_\alpha$ and $|E(\alpha)| = 1$. Let $a_1, \ldots, a_n$ be the entities of $E(\alpha)$. $n \geqslant 2$. $x \in \mathit{CloserEq}(a_i, b)$ for every $a_i \in \alpha$ and $b \in Q$. Therefore $d(x, a_i) \leqslant d(x, b)$ for every $a_i \in \alpha$ and $b \in Q$. Let $p_i = \pi_{a_i}(x)$. $d(x, p_i) \leqslant d(x, q)$ for every $q \in \partial Q$. In order to prove that $x \in MA(Q)$, it is enough to show that $p_i \neq p_j$ for every $i \neq j$. If $p_i = p_j$, then $a_i \cap a_j \neq \emptyset$. Let $b = a_i \cap a_j$. $b \subset a_i$ or $b \subset a_j$ or both. Therefore $a_i \notin E(\alpha)$, or $a_j \notin E(\alpha)$ or both. Contradiction. □

Lemma 13 (Voronoi faces are simply connected). *If the boundary of Q is connected, and the faces of Q are simply connected, then the faces of $VD(Q)$ are also simply connected.*

Proof. Sherbrooke [19] proves this claim for the faces of $MA(Q)$. In order to complete the proof of the present lemma, we have to show that a face $f_\alpha \in VD(Q) \setminus MA(Q)$ is simply connected. Lemma 12 implies that such a face f_α satisfies that $|E(\alpha)| = 1$. Therefore α contains an entity b s.t. every entity $a \in \alpha$ satisfies that $b \subseteq a$. Let $x \in R_\alpha$. $\pi_b(x) \in a$ for every $a \in \alpha$, therefore $[x, \pi_b(x)] \subseteq R_\alpha$ (Lemma 9). Therefore R_α is connected.

Suppose on the contrary that R_α is not simply connected. Then $\mathit{carbis}(\alpha)$ is a plane, and there exists a point $x \in \mathit{carbis}(\alpha) \setminus R_\alpha$ which is enclosed by a loop $L \subseteq R_\alpha$. Consider the line M through x and $\pi_b(x)$. $M \subseteq \mathit{carbis}(\alpha)$. Let y be the intersection point of L and M which is farthest from $\pi_b(x)$. $y \in R_\alpha$. Therefore $[y, \pi_b(x)] \subseteq R_\alpha$ (Lemma 9). Contradiction, since $x \in [y, \pi_b(x)]$. □

Lemma 14 ($VD(Q)$ does not contain a loop of edges $e_{abc\star}$). *Let Q be a polyhedron whose boundary is connected, and whose faces are simply connected. Let f_α be a bounded Voronoi face of $VD(Q)$. There does not exist a set of entities $\beta \supset \alpha$ s.t. all the edges of f_α are governed by $e_{\beta\star}$.*

Proof. Suppose on the contrary that there exists such a set of entities β. We first show that there do not exist two entities $a, b \in \beta$ s.t. $a \supset b$. Suppose there are. Let $c \in \beta \setminus \{a, b\}$. $\partial f_\alpha \subseteq \mathit{carbis}(a, b)$

(Lemma 8.2). $carbis(a, b)$ is either a line or a plane. Since f_α is a bounded face, $carbis(a, b)$ cannot be a line, so it is a plane. Let x be a point in f_α. The line through x and $\pi_b(x)$ intersects ∂f_α in two points x_1 and x_2. $d(x_1, b) = d(x_1, c)$ and also $d(x_2, b) = d(x_2, c)$. Therefore $\pi_b(x) = \pi_c(x)$ (Lemma 1). Therefore $E(\beta) = 1$. Therefore ∂f_α is a line. Contradiction.

Let $a, b \in \alpha$, and $c \in \beta \setminus \{a, b\}$. Define S_a to be the solid composed of the projection segments of f_α on a. Define S_b similarly. Let C_c be the projection of ∂f_α on c. Since c is simply connected, the region bounded by C_c is in c. Define T_c to be the surface composed of the projection segments of ∂f_α on c together with the part of c enclosed by C_c. $S_a \subseteq R_a$, $S_b \subseteq R_b$, $T_c \subseteq R_c$ (Lemma 9). Therefore $int(S_a)$ does not intersect S_b and T_c, and $int(S_b)$ does not intersect S_a and T_c. Therefore S_a (or S_b) is in the interior of the solid defined by T_c. Therefore a is in the interior of the solid defined by T_c. We show in the following that this implies that a and c are not in the same connected component of the boundary of Q, in contradiction to the assumption of the lemma.

Entities a and c are not incident one on the other, therefore if they are connected, there is an entity d that intersects T_c. Since $T_c \subset R_c$, d must intersect T_c in a point incident on c and d. Therefore d is wholly in the interior of the solid defined by T_c, and d either contains c or is adjacent to c. In this case there exists a point $x \in \partial f_\alpha$ s.t. $x \in Closer(d, c)$ in contradiction to $\partial f_\alpha \subseteq R_c$. □

4. The space subdivision algorithm

In this section we define the proximity structure subdivision and give an algorithm for constructing it. We prove that the algorithm halts, and show that when utilizing cells with linear boundaries, the geometric operations involved amount to solving a quadratic equation in a single variable.

Intuitively, the general idea is to recursively subdivide space according to the distances of the cells from the entities of the polyhedron, such that all the points in a cell share the same nearest entities. We would like the cells to separate Voronoi vertices, i.e., that each cell will contain no more than one Voronoi vertex. Therefore we stop the subdivision process when the number of entities attached to a cell is smaller than or equal to four. This subdivision process might not halt, since it is possible that a point has more than four governors. For example, every vertex of Q has a set of governors that includes all the entities of Q containing that vertex. Note that this situation is not degenerate, since a small perturbation of the polyhedron does not necessarily modify the symbolic structure of the Voronoi diagram.[1] Lemma 18 states the situations in which a point has more than four governors in a non-degenerate diagram. These situations are added to the halting criteria of the recursion.

4.1. Definition and algorithm

Definition 1. A *proximity structure subdivision* (*PSS*) is a space subdivision[2] in which each cell C is labeled by a set α of polyhedron entities, such that two conditions hold. Let C_α be a cell that is labeled by a set α of polyhedron entities. The two conditions are the following:
1. $b \notin \alpha$ iff there exists an entity a of Q such that $C_\alpha \subseteq Closer(a, b)$.

[1] As a result, it is inaccurate to define 'degeneracy of a Voronoi diagram of a polyhedron' by saying that there exists a point with more than four nearest sites.

[2] We treat all subdivision cells as closed sets, hence they include their boundaries.

2. At least one of the following holds:
   (a) $|\alpha| \leqslant 4$.
   (b) $|\alpha| = 5$, and α includes an edge and two coplanar faces containing that edge.
   (c) $|\alpha| = 5$, and α includes a vertex and two colinear edges containing that vertex.
   (d) $|\alpha| = 6$, and α is composed of two disjoint sets, each consists of an edge and two coplanar faces containing that edge.
   (e) $|\alpha| = 6$, and α is composed of two disjoint sets, each consists of a vertex and two colinear edges containing that vertex.
   (f) $|\alpha| = 6$, and α is composed of two disjoint sets, one consists of an edge and two coplanar faces containing that edge, and the other consists of a vertex and two colinear edges containing that vertex.
   (g) All the entities in α share a vertex.
   (h) All the entities in α except one share a vertex and a plane.

The first condition serves for reducing the number of polyhedron entities relevant to proximity information of a cell, and is thus similar in purpose to the condition used in [9]. The second condition refines the subdivision to enable extraction of the structure of the Voronoi graph. The following lemmas give basic properties of the subdivision.

Lemma 15. *Let C_α be a cell in a PSS. Let b be an entity. If $b \notin \alpha$, then $C_\alpha \cap R_b = \emptyset$.*

Proof. If $b \notin \alpha$, then there exists an entity a s.t. $C_\alpha \subseteq \mathit{Closer}(a, b)$. Therefore, by Lemma 4.2,

$$C_\alpha \cap \mathit{CloserEq}(b, a) = \emptyset. \qquad \square$$

Lemma 16. *Let C_α be a cell in a PSS. Let b be an entity. If $b \notin \alpha$, then there exists an entity $a \in \alpha$ s.t. $C_\alpha \subseteq \mathit{Closer}(a, b)$.*

Proof. We show in the following that if $b \notin \alpha$ and there does not exist an entity $a \in \alpha$ s.t. $C_\alpha \subseteq \mathit{Closer}(a, b)$, then there is an infinite number of entities in Q. Let $a_1 = b$. $a_1 \notin \alpha$, therefore there exists an entity a_2 of Q such that $C_\alpha \subseteq \mathit{Closer}(a_2, a_1)$. $a_2 \notin \alpha$, therefore there exists an entity a_3 of Q such that $C_\alpha \subseteq \mathit{Closer}(a_3, a_2)$. Lemma 7.1 implies that $C_\alpha \subseteq \mathit{Closer}(a_3, a_1)$ and therefore $a_3 \notin \alpha$. Thus there exists an infinite sequence of entities $\{a_i\}$ s.t. $C_\alpha \subseteq \mathit{Closer}(a_j, a_i)$ for any $i < j$. Therefore for any $i \neq j$ $a_i \neq a_j$. $\square$

Subdivision process

A proximity structure subdivision is easily computed recursively. We start with a cell that bounds the world of interest. For each cell, the set α is computed according to the first condition. Cells for which the second condition does not hold are subdivided, and the algorithm is invoked recursively on the sub-cells. Obviously, if $C_\alpha \subseteq C_\beta$ then $\alpha \subseteq \beta$, and the computation of α for sub-cells can be done efficiently by considering only the entities attached to the parent cell. In practice, the simplest way to implement the algorithm is by using an octree to represent the subdivision.

4.2. Halting of the subdivision process

In this section we prove that the subdivision process halts if $VD(Q)$ is not degenerate. If $VD(Q)$ is degenerate then an additional halting condition is needed (Section 6).

Definition 2. For a point x, let $f_1(x), \ldots, f_k(x)$ be the footpoints of x on Q, and let $\alpha_i(x)$ be the set of entities governing x and containing $f_i(x)$. We say that $VD(Q)$ is *non-degenerate* iff for every point x the two following conditions are satisfied:
1. For any permutation on $\{\alpha_i\}$: Let $\alpha(x) = \alpha_1(x) \cup \cdots \cup \alpha_i(x)$ for $1 \leqslant i \leqslant k-1$. $dim(carbis(\alpha(x) \cup \alpha_{i+1}(x))) < dim(carbis(\alpha(x)))$.
2. For every $1 \leqslant i \leqslant k$ and $1 \leqslant j \neq i \leqslant k$, if $|\alpha_j(x)| > 1$, then $dim(carbis(\alpha_i(x) \cup \alpha_j(x))) < dim(carbis(\alpha_i(x))) - 1$.

The first item of the above definition is closely related to the definition usually used for degeneracy of the medial axis or of the Voronoi diagram of disjoint sites. This item states that if the diagram is not degenerate, then the dimension of the locus of points equidistant from a partial set of the footpoints of a point decreases as additional footpoints are added to the set.

The second item of the above definition handles the case of non-disjoint sites. Consider a point with two footpoints f_1 and f_2 incident on α_1 and α_2, respectively. The locus of points equidistant from the entities of $\alpha_1 \cup \alpha_2$ is the intersection of three sets: (1) the set of points equidistant from α_1, (2) the set of points equidistant from α_2, and (3) the set of points equidistant from an entity $a_1 \in \alpha_1$ and $a_2 \in \alpha_2$. If it is not a degenerate case, then the dimension of the intersection set decreases as each of the three sets is added.

In Lemmas 17–19 we assume that $VD(Q)$ is not degenerate. Lemma 18 states the conditions in which a point has more than four governors. Lemma 17 is an auxiliary lemma of Lemma 18.

Lemma 17 (The *carbis* of entities sharing a vertex). *Let v be a vertex of Q. Let $e_1, \ldots, e_n$ be edges of Q containing v. Let $f_1, \ldots, f_k$ be faces of Q containing v. Let $\alpha = \{v, e_1, \ldots, e_n, f_1, \ldots, f_k\}$. Suppose $n > 1$ or $k > 0$ (or both).*
1. *If there exists a line L s.t. $a \subset L$ for every $a \in \alpha$, then $carbis(\alpha)$ is a plane orthogonal to L at v.*
2. *If all the entities of α share a plane P, and do not share a line, then $carbis(\alpha)$ is a line orthogonal to P at v.*
3. *If the entities of α do not share a plane, then $carbis(\alpha) = v$.*

Proof.
1. The bisector of a line and a point incident on the line is a plane orthogonal to the line at the point.
2. Let L be the line orthogonal to P at v. First we prove that $L \subseteq carbis(\alpha)$. Let $x \in L$. $d(x, P) = d(x, v)$. Therefore $d(x, car(f_i)) = d(x, v)$ for every $1 \leqslant i \leqslant k$, since $car(f_i) = P$ for every $1 \leqslant i \leqslant k$. Similarly, $d(x, car(e_i)) = d(x, v)$ for every $1 \leqslant i \leqslant n$, since $v \in car(e_i)$, and $car(e_i) \subset P$ for every $1 \leqslant i \leqslant n$.
   Now we prove that $carbis(\alpha) \subseteq L$. Let $x \in carbis(\alpha)$. If $k > 0$ then $d(x, v) = d(x, P)$, and therefore $x \in L$. If $k = 0$ then $n > 1$. Let e_1 and e_2 be two edges in α s.t. $car(e_1) \neq car(e_2)$. $carbis(v, e_1)$ and $carbis(v, e_1)$ are two different planes, and their intersection is a line.
3. It is clear that $v \subseteq carbis(\alpha)$, since v is incident on all the entities of α. We prove in the following that $carbis(\alpha) \subseteq v$. Let β be a maximal subset of α s.t. all the entities in β share a plane P. Lemma 17.2

implies that $carbis(\beta)$ is a line L orthogonal to P at v. Let $a \in \alpha \setminus \beta$. If a is a face, then L and $carbis(v,a)$ are two different lines, and their intersection is a point. Otherwise a is an edge. Let $R = carbis(v,a)$. R is a plane orthogonal to $car(a)$ at v. Suppose on the contrary that $carbis(\alpha) \nsubseteq v$, then $L \subseteq R$. Therefore $car(a)$ is orthogonal to L at v, and $car(a) \subseteq P$. Contradiction. □

Lemma 18 (The number of governors of a point). *Let Q be a polyhedron s.t. $VD(Q)$ is not degenerate. Let α be a set of entities of Q s.t. $R_\alpha \neq \emptyset$. One of the conditions* 2a–2h *of the definition of the PSS* (*Definition 1*) *holds.*

Proof. Suppose $|\alpha| > 4$. Let x be a point in R_α. Lemma 8.2 implies that $x \in carbis(\alpha)$. Let k be the number of footpoints of x on Q. Definition 2.1 implies that $k \leqslant 4$. Let $\alpha_1, \ldots, \alpha_k$ be the subsets of α, s.t. α_i is the set of entities sharing the footpoint f_i. Let $l = |\alpha|$, and $l_i = |\alpha_i|$. The sets $\alpha_1, \ldots, \alpha_k$ are disjoint, since otherwise if $a \in \alpha_i \cap \alpha_j$ for $i \neq j$, then a includes two different footpoints of x, in contradiction to the linearity and convexity of a. Therefore the sets $\alpha_1, \ldots, \alpha_k$ are disjoint, and $\sum_{1 \leqslant i \leqslant k} l_i = l$. Claim: there exists $1 \leqslant i \leqslant k$ s.t. $dim(carbis(\alpha_i)) > 4 - l_i$.

Suppose on the contrary that for every $1 \leqslant i \leqslant k$ $dim(carbis(\alpha_i)) \leqslant 4 - l_i$. Consider the two cases:

1. There exist two sets α_i and α_j s.t. $l_i > 1$ and $l_j > 1$. Then Definition 2.2 implies that $dim(carbis(\alpha_i \cup \alpha_j)) < \min(4 - l_i, 4 - l_j) - 1$. Consider the two cases:
   (a) $l_i > 2$ or $l_j > 2$. Then $dim(carbis(\alpha_i \cup \alpha_j)) < 0$, in contradiction to the existence of x.
   (b) $l_i = 2$ and $l_j = 2$. Then $dim(carbis(\alpha_i \cup \alpha_j)) = 0$. $l_i + l_j = 4 < l$, therefore there exists a third footpoint f_m. Definition 2.1 implies that $dim(carbis(\alpha_i \cup \alpha_j \cup \alpha_m)) < 0$, in contradiction to the existence of x.
2. Only one set α_i satisfies that $l_i > 1$. $l_i = l - (k-1)$. Definition 2.1 implies that $dim(carbis(\alpha_i)) \geqslant k - 1$. These two equations imply that $dim(carbis(\alpha_i)) \geqslant l - l_i > 4 - l_i$. Contradiction.
3. There does not exist a set α_i s.t. $l_i > 1$. Then $l \leqslant 4$, and contradiction.

This completes the proof of the claim, i.e., there exists $1 \leqslant i \leqslant k$ s.t. $dim(carbis(\alpha_i)) > 4 - l_i$.

Let f_i be a footpoint s.t. $dim(carbis(\alpha_i)) > 4 - l_i$. f_i is either a vertex v of Q, or incident on an edge e of Q. Consider the two cases:

1. f_i is a vertex of Q. Lemma 17 implies that:
   (a) If the entities of α_i do not share a plane, then $dim(carbis(\alpha_i)) = 0$. Definition 2.1 implies that $l = l_i$, i.e., Definition 1.2g is satisfied.
   (b) If all the entities of α_i share a plane, and do not share a line, then $dim(carbis(\alpha_i)) = 1$. Definition 2.1 implies that $k \leqslant 2$. If $k = 1$, then Definition 1.2g is satisfied. If $k = 2$, let f_j be the other footpoint. Definition 2.1 implies that $|l_j| \leqslant 1$, and therefore Definition 1.2h is satisfied.
   (c) If all the entities of α_i share a line, i.e., α_i consists of the vertex f_i and two colinear edges containing that vertex, then $dim(carbis(\alpha_i)) = 2$. Consider the two cases:
      i. $k > 2$. Definition 2 implies that there are two additional footpoints f_j and f_m s.t. $l_j = l_m = 1$. Therefore Definition 1.2c is satisfied.
      ii. $k = 2$. Let f_j be the other footpoint. If $l_j = 1$, then Definition 1.2a is satisfied. If $l_j = 2$, then Definition 1.2c is satisfied. Suppose $l_j > 2$. $l_i = 3$, therefore $dim(carbis(\alpha_j)) \geqslant 2$ (Definition 2.2), and because $l_j > 2$, $dim(carbis(\alpha_j)) > 4 - l_j$. f_j is a footpoint satisfying that $dim(carbis(\alpha_j)) > 4 - l_j$, and therefore the discussion in the previous items (item 1a and item 1b) applies also to f_j as well. Therefore if f_j is a vertex, then it is a vertex incident on two colinear edges. Recall that α_i and α_j are disjoint. Therefore Definition 1.2c is satisfied. If

f_j is on an edge, then α_j consists of the edge and two coplanar faces containing that edge, and Definition 1.2f is satisfied.

2. f_i is on an edge of Q. In this case α_i consists of the edge and two coplanar faces containing that edge. Therefore $|\alpha_i| = 3$ and $dim(carbis(\alpha_i)) = 2$. This case is analogous to item 1c. Therefore in this case one of items 2a, 2b, 2d and 2f of Definition 1 is satisfied. □

Lemma 19. *If $VD(Q)$ is not degenerate the subdivision process halts.*

Proof. Suppose the subdivision process does not halt. Then there exists an infinite sequence of cells C_{α_i} s.t. (1) $size(C_{\alpha_i}) \to 0$, (2) for every i, C_{α_i} is not a leaf, and (3) $C_{\alpha_{i+1}} \subseteq C_{\alpha_i}$. The sequence converges. Let x be $\bigcap_{\forall i} C_{\alpha_i}$. Let $\alpha(x)$ be the set of governors of x. For every entity $b \notin \alpha(x)$ there exists an entity a s.t. $x \in Closer(a, b)$ (Lemma 8.3). $Closer(a, b)$ is an open set (Lemma 3.2). Therefore for every entity $b \notin \alpha(x)$ there exists an entity a, and $\varepsilon(b) > 0$ s.t. if point $y \in B(x, \varepsilon)$, then $y \in Closer(a, b)$. Let D_x be the minimum of the ε_b for all $b \notin \alpha(x)$. There exists an integer N s.t. for every $i > N$, $C_{\alpha_i} \in B(x, D_x)$. Let $i > N$. If $c \in \alpha_i$ then there does not exist an entity d s.t. $C_{\alpha_i} \subseteq Closer(d, c)$ (definition of PSS), and therefore $c \in \alpha(x)$. Therefore for $i > N$, $\alpha_i \subseteq \alpha(x)$, and C_{α_i} is a leaf (Lemma 18). Contradiction. □

4.3. Geometric operations of the subdivision process

In order to compute the set of entities attached to a cell, we have to answer the query: Given a cell C, and entities a, b, is $C \subseteq Closer(a, b)$? Lemma 3.4 implies that testing whether $C \subseteq Closer(a, b)$ is equivalent to testing whether $\partial C \subseteq Closer(a, b)$.

Using linear cell boundaries, the algorithm in Fig. 3 tests whether $\partial C \subseteq Closer(a, b)$. In order to test whether a face F of C is in $Closer(a, b)$, it is not enough to test the vertices of F. Even if all vertices of F are in $Closer(a, b)$, there might still be a point $x \in F$ s.t. $x \notin Closer(a, b)$. Therefore we have to test whether F intersects the bisector $bis(a, b)$.

a and b are linear entities, therefore $bis(a, b)$ is a piecewise quadratic surface. The bisector is a *piecewise* quadratic surface, and not a quadratic surface, because a and b are polyhedron entities, not infinite lines or planes. Each section of $bis(a, b)$ is a part of $carbis(a', b')$ s.t. $a' \subseteq a$ and $b' \subseteq b$. $carbis(a', b')$ is a quadratic surface for any two entities a' and b'.

In order to work with quadratic surfaces, and not piecewise quadratic surface, we first decompose each face of C into polygons $P_{a'b'}$ s.t. (1) $a' \subseteq a$, (2) $b' \subseteq b$, and (3) a point $x \in P_{a'b'}$ iff $d(x, a) = d(x, car(a'))$ and $d(x, b) = d(x, car(b'))$ (line 2). The part of $bis(a, b)$ in $P_{a'b'}$ is equal to $carbis(a', b')$, and therefore the location of $P_{a'b'}$ with respect to $bis(a, b)$ can easily be tested (lines 4–23).

If $a' = b'$ then $P_{a'b'} \notin Closer(a, b)$ iff $a \not\subset b$ or there exists a vertex of $P_{a'b'}$ on $bis(a, b)$ (lines 5–11). Note that in this case $(a \subset b)$ $bis(a, b)$ is a piecewise *linear* surface which can be easily computed. If $a' \neq b'$ then $P_{a'b'} \in Closer(a, b)$ iff $d(x, car(a')) < d(x, car(b'))$ for all points $x \in P_{a'b'}$ (lines 12–23). This condition is tested by comparing the distances from an arbitrary point x to $car(a')$ and $car(b')$. If $d(x, car(a')) \geqslant d(x, car(b'))$, then $P_{a'b'} \nsubseteq Closer(a, b)$ (lines 12–14). Otherwise, $P_{a'b'} \nsubseteq Closer(a, b)$ iff $carbis(a', b')$ intersects $P_{a'b'}$ (lines 15–23). This is tested by testing whether $carbis(a', b')$ intersects the plane containing $P_{a'b'}$ (lines 16–17), the boundary of $P_{a'b'}$ (lines 18–20), or the interior of $P_{a'b'}$ (lines 21–23).

The algorithm of Fig. 3 uses three auxiliary functions. The function *PointOnPolygon*(P) picks any point on the polygon P, and the function *PointOnConicSection*(B) picks any point on the conic section B.

```
CellIsCloser (Cell C, Entity a, Entity b)
for every face F of C
    PL = DecomposeCellFace(F, a, b);
    for every polygon P_{a'b'} in PL
        if a' = b'
            if a ⊂ b
                for every vertex v of P_{a'b'}
                    if v ∈ bis(a, b)
                        return NO;
                continue;
            else
                return NO;
        v = PointOnPolygon(P_{a'b'});
        if d(v, car(a') ⩾ d(v, car(b'))
            return NO;
        B = carbis(a', b') ∩ plane(P_{a'b'});
        if B = ∅
            continue;
        for every edge E of P_{a'b'}
            if B ∩ E ≠ ∅
                return NO;
        x = PointOnConicSection(B);
        if x ∈ P_{a'b'}
            return NO;
return YES;
```

Fig. 3. $\mathit{CellIsCloser}(C, a, b)$ returns YES iff $C \subseteq \mathit{Closer}(a, b)$. The function solves at most a quadratic equation.

The function $\mathit{DecomposeCellFace}(F, a, b)$ decomposes a face F of a cell C into polygons $P_{a'b'}$ s.t. (1) $a' \subseteq a$, (2) $b' \subseteq b$, and (3) $x \in P_{a'b'}$ iff $d(x, a) = d(x, a') = d(x, \mathit{car}(a'))$ and $d(x, b) = d(x, b') = d(x, \mathit{car}(b'))$.

Each polygon $P_{a'b'}$ is the intersection of two polygons $P_{a'}$ and $P_{b'}$. $P_{a'} = F \cap H(a', a)$ where $H(a', a) = \{x \mid d(x, a) = d(x, a') = d(x, \mathit{car}(a'))\}$. $P_{b'}$ is defined similarly. $H(a', a)$ is an intersection of a finite number of half-spaces each defined by a single plane. Consider the three cases:

1. a is a vertex v. Then $a' = v$ and $H(v, v)$ is the whole space.
2. a is an edge e. If v is a vertex of e then $H(v, e)$ is the half-space defined by the plane orthogonal to e at v, and which does not contain e. $H(e, e)$ is the intersection of two half-spaces defined by the two planes orthogonal to e at its vertices, and which contain e.
3. a is a face f. If v is a vertex of f, then $H(v, f)$ is the intersection of $H(v, e_1)$ and $H(v, e_2)$ where e_1 and e_2 are the two edges containing v in f. If e is an edge of f, then $H(e, f)$ is the intersection of $H(e, e)$ and the half-space defined by the plane orthogonal to f at e and which does not contain f. $H(f, f)$ is the intersection of half spaces each defined by the plane orthogonal to f at one of its edges, and which contains f.

Lemmas 20 and 21 prove that the algorithm of Fig. 3 is correct.

Lemma 20. *Let $a, b, a' \subseteq a, b' \subseteq b$ be entities. Let $P_{a'b'}$ be a planar polygon s.t. $x \in P_{a'b'}$ iff $d(x,a) = d(x, car(a'))$ and $d(x,b) = d(x, car(b'))$.*
1. *If $a' = b'$ and $a \not\subset b$, then $P_{a'b'} \nsubseteq Closer(a,b)$.*
2. *If $a' = b'$ and $a \subset b$, then $P_{a'b'} \subseteq Closer(a,b)$ iff every vertex v of $P_{a'b'}$ satisfies that $v \notin bis(a,b)$.*
3. *If $a' \neq b'$ then $P_{a'b'} \subseteq Closer(a,b)$ iff $\forall_{x \in P_{a'b'}} d(x, car(a')) < d(x, car(b'))$.*

Proof.
1. For every $x \in P_{a'b'}$ $d(x, car(a')) = d(x, car(b'))$, and therefore $d(x,a) = d(x,b)$. Therefore $P_{a'b'} \nsubseteq Closer(a,b)$, since $a \not\subset b$.
2. Suppose there exists a vertex v of $P_{a'b'}$ s.t. $v \in bis(a,b)$. Then $v \in CloserEq(b,a)$, and by Lemma 4.2 $v \notin Closer(a,b)$.
   Suppose every vertex v of $P_{a'b'}$ satisfies that $v \notin bis(a,b)$. For every $x \in P_{a'b'}$ $d(x, car(a')) = d(x, car(b'))$, and therefore $d(x,a) = d(x,b)$. Therefore $P_{a'b'} \subseteq CloserEq(a,b)$. Suppose on the contrary that there is a point $x \in P_{a'b'}$ s.t. $x \notin Closer(a,b)$. Therefore for every $\varepsilon > 0$ there exists a point y s.t. $d(x,y) < \varepsilon$ and $d(y,b) < d(y,a)$. $y \notin P_{a'b'}$. Therefore $x \in \partial P_{a'b'}$. We show in the following that if there exists a point $x \in bis(a,b) \cap \partial P_{a'b'}$, then one at least of the vertices of $P_{a'b'}$ satisfies that $v \in bis(a,b)$.
   Suppose on the contrary that there exists such a point x, and no vertex v satisfies that $v \in bis(a,b)$. Let v_1 and v_2 be the vertices of the edge of $P_{a'b'}$ containing x. $v_1, v_2 \in Closer(a,b)$. If a is a vertex then $Closer(a,b)$ is convex, and contradiction. Otherwise a is an edge, and b is a face. Let u_1 and u_2 be the two vertices of a. $Closer(a,b)$ is composed of three regions: (1) $Closer(a,b) \cap Closer(u_1,a)$, (2) $Closer(a,b) \cap Closer(u_2,a)$ and (3) $Closer(a,b) \cap Closer(a,u_1) \cap Closer(a,u_2)$. Each of the three regions is convex. Therefore v_1 and v_2 are in two different regions. If v_i is in the first region, then $a' = u_1$. If v_i is in the second region, then $a' = u_2$. If v_i is in the third region, then $a' = a$. Contradiction.
3. Suppose $P_{a'b'} \subseteq Closer(a,b)$. Let $x \in P_{a'b'}$. $a \not\subset b$ since if $a \subset b$ then $\pi_a(x) = \pi_b(x)$, and therefore $a' = b'$ (since a' and b' is the lowest dimensional entity of Q containing $\pi_a(x) = \pi_b(x)$). The facts that $a \not\subset b$ and $P_{a'b'} \subseteq Closer(a,b)$, imply that for every x in $P_{a'b'}$ $d(x,a) < d(x,b)$, and therefore $d(x, car(a')) < d(x, car(b'))$.
   Suppose $\forall_{x \in P_{a'b'}} d(x, car(a')) < d(x, car(b'))$. Then $\forall_{x \in P_{a'b'}} d(x,a) < d(x,b)$. Then $\forall_{x \in P_{a'b'}} x \in Closer(a,b)$. □

The following lemma justifies lines 21–23 of the algorithm. $B = carbis(a', b') \cap plane(P_{a'b'})$, and therefore a conic section. We show in this lemma that if B does not intersect any edge of a polygon (lines 18–20), then it is enough to test one point of B in order to determine whether B intersects the polygon.

Lemma 21. *Let B be a conic section, and P a polygon. If $B \cap P \neq \emptyset$, and $B \cap \partial P = \emptyset$, then B is wholly in the interior of P.*

Proof. It is clear that $B \cap P$ is wholly in the interior of P. B is not wholly in the interior of P, if B has more than one connected component, and one of the connected components is bounded. This is not the case, since B is a conic section. □

The only geometric operations used in the algorithm are the ones used in order (1) to decompose a planar polygon by planes, (2) to decide whether a point x is closer to the carrier of entity a than to the carrier of entity b, (3) to decide whether an edge of a polygon intersects a conic section, and (4) to pick a point on a conic section. The first two queries are answered by linear operations. The last two queries are answered by solving a uni-variate quadratic equation.

5. Extraction of Voronoi elements

In this section we show how to construct the Voronoi graph from a proximity structure subdivision.

5.1. Computing Voronoi edge witnesses

As a first step, we find which Voronoi edges intersect the boundaries of the cells. Only cells labeled by three or more entities should be considered, since the other cells do not intersect Voronoi edges. The computation is done separately for each cell face F (Fig. 4). For a given cell C_α, a face F of the cell, and three entities $a, b, c \in \alpha$, CellFaceVoronoiEdgeIntersection computes the intersection points of F and Voronoi edges e_β s.t. $a, b, c \in \beta$.

A point is on a Voronoi edge e_β iff it lies on $bis(a, b)$ for any $a, b \in \beta$, and is not closer to any other polyhedron entity than to the entities of β. The algorithm intersects the bisectors of the carriers of $a, b, c \in \beta$ with the plane of the face F (lines 1–2), resulting in two conic sections, which are then intersected (line 3). Intersection points that are outside of the face (lines 5–6) or that do not obey the above

CellFaceVoronoiEdgeIntersection
(CellEntities α, CellFace F, Entity a, Entity b, Entity c)
1 $W_{ab} = carbis(a, b) \cap plane(F)$;
2 $W_{ac} = carbis(a, c) \cap plane(F)$;
3 $W = W_{ab} \cap W_{ac}$;
4 **for** every point $x \in W$
5   **if** $x \notin F$
6   **goto** 4;
7   **if** $(\pi_{car(a)}(x) \notin a)$ **or** $(\pi_{car(b)}(x) \notin b)$ **or** $(\pi_{car(c)}(x) \notin c)$
8   **goto** 4;
9   $\beta(x) = \{a, b, c\}$;
10   **for** every entity $e \in \alpha \setminus \{a, b, c\}$
11   **if** $(x \in carbis(a, e))$ **and** $(\pi_{car(e)}(x) \in e)$
12       $\beta(x) = \beta(x) \cup \{e\}$;
13   **goto** 10;
14   **if** $x \in Closer(e, a)$
15   **goto** 4;
16   **output** $(x, \beta(x))$;
17 **return**;

Fig. 4. Computing the intersection points of a face F of $C = C_\alpha$ and Voronoi edges e_β s.t. $a, b, c \in \beta$. The function computes the intersection of two conic sections, i.e., the roots of at most a quartic uni-variate polynomial.

criterion (lines 7–8 and lines 14–15) are removed. Voronoi edges having more than three governors are detected in lines 11–12.

If W includes an infinite number of points, then W is part of a conic section contained in F. In this case W is modified to contain only the intersection points between W and ∂F.

Lemma 22 (The algorithm of Fig. 4 is correct). *Let the set of pairs* $\{(x_i, \beta_i)\}$ *for* $1 \leqslant i \leqslant n$ *be the output of the algorithm of Fig. 4. Let* $X = \{x_i\}$ *for* $1 \leqslant i \leqslant n$.
1. *For every* $1 \leqslant i \leqslant n$: $x_i \in e_{\beta_i}$.
2. $X = e_{abc\star} \cap F$.

Proof.
1. (a) $x_i \in bis(a, b)$ *for every* $a, b \in \beta_i$. Implied by Lemma 6.3.
   (b) *For every pair of entities* $e \in \alpha$ *and* $b \in \beta_i$ $x_i \notin Closer(e, b)$. Suppose on the contrary that $x_i \in Closer(e, b)$. $\pi_{car(b)}(x_i) \in b$, and therefore $e \not\subseteq b$ (Lemma 5). Therefore $x_i \in Closer(e, b)$ implies that $d(x_i, e) < d(x_i, b)$. $x_i \notin Closer(e, a)$ (lines 14–15), therefore $d(x_i, e) \geqslant d(x_i, a)$. Therefore $d(x_i, b) > d(x_i, a)$, in contradiction to item 1a.
   (c) *For every pair of entities* $e \in Q$ *and* $b \in \beta_i$ $x_i \notin Closer(e, b)$. Suppose on the contrary $x_i \in Closer(e, b)$. $e \notin \alpha$ (item 1b), therefore there exists an entity $f \in \alpha$ s.t. $x_i \in Closer(f, e)$ (Lemma 16). Therefore $x_i \in Closer(f, b)$ (Lemma 7.1). Contradiction to item 1b.
   (d) *For every pair of entities* $e \in Q$ *and* $b \in \beta_i$ $x_i \in CloserEq(b, e)$, *i.e.,* $x_i \in e_{\beta_i}$. If $b \subset e$ or $e \subset b$ or $b \cap e = \emptyset$, then it is implied from item 1c (Lemma 4.1). Suppose $e \cap b = d \neq b, e$. $x_i \notin Closer(e, b)$, therefore $d(x_i, b) \leqslant d(x_i, e)$. Suppose on the contrary $x_i \notin CloserEq(b, e)$. Then there exists an $\varepsilon > 0$ s.t. if $y \in B(x_i, \varepsilon)$ then $d(y, b) \geqslant d(y, e)$. Therefore $d(x_i, b) = d(x_i, e) = d(x_i, d)$ (Lemma 2). Therefore $x_i \in CloserEq(d, b) \cap CloserEq(d, e)$. By the result of the present item for d and b, $x_i \in CloserEq(b, d)$. Therefore $x_i \in CloserEq(b, e)$ (Lemma 7.3). Contradiction.
2. (a) $X \subseteq e_{abc\star} \cap F$. Let $x_i \in X$. $x_i \in F$ (lines 5–6). $x_i \in e_{\beta_i}$ (item 1). $a, b, c \in \beta_i$ (lines 9, 12). Therefore $x_i \in e_{abc\star} \cap F$.
   (b) $X \supseteq e_{abc\star} \cap F$. Let $x \in e_{abc\star} \cap F$. We show in the following that there exists $1 \leqslant i \leqslant n$ s.t. $x = x_i$. Let β be a set of entities s.t. $x \in e_\beta \cap F$ and $\beta = abc\star$. Let b be an entity in β. $x \in e_\beta$, therefore $x \in CloserEq(b, e)$, for any entity e of Q, and in particular for $e \subset b$. Therefore $d(x, b) = d(x, car(b))$ (Lemma 5). Therefore $\pi_{car(b)}(x) \in b$. If a is also an entity in β, then $x \in CloserEq(a, b) \cap CloserEq(b, a)$, and therefore $d(x, a) = d(x, b)$. Therefore $d(x, car(a)) = d(x, car(b))$. It is clear that $x \notin Closer(e, b)$ since $Closer(e, b) \cap CloserEq(b, e) = \emptyset$ (Lemma 4.2). □

In lines 7–8 and 11 we test whether $\pi_{car(a)}(x) \in a$ for every entity $a \in \beta$. If $\pi_{car(a)}(x) \notin a$, then the facts that $x \in carbis(a, b)$ and $x \notin Closer(e, b)$ for any entity $e \in Q$ do not imply that $x \notin Closer(e, a)$. This case is demonstrated in Fig. 5. In this figure $x \in carbis(a, b) \cap R_b$. However $x \in Closer(e, a)$.

The highest degree operation performed in the algorithm of Fig. 4 is the intersection of two conic sections in line 3. Therefore the geometric operations performed in the algorithm of Fig. 4 amount to solving a uni-variate polynomial whose degree is (1) 1, if all three entities a, b, c are faces or all are vertices, (2) not more than 2, if two of the entities are faces or two are vertices, or (3) not more than 4, in

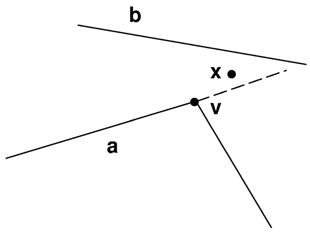

Fig. 5. $x \in carbis(a, b) \cap R_b$. In spite of that, $x \in Closer(v, a)$.

all other cases. In the last two cases, the degree is lower than 2 or 4 when (1) the entities are incident on each other, or (2) two of the entities are edges sharing a plane. In summary so far, we have

Lemma 23. *All intersection points between Voronoi edges and subdivision cell boundaries can be computed using linear operations, distance comparisons, and computing roots of at most quartic univariate polynomials.*

5.2. Extraction of Voronoi vertices

After computing edge witnesses, we identify Voronoi vertices. In the following we prove that a cell C_α does not contain a vertex of $VD(Q)$ not labeled by α. Assuming that a cell does not contain two different vertices with the same governors, we provide a simple criterion to determine whether a cell contains a vertex or not, using the set of Voronoi edge witnesses computed earlier. The implications of the assumption are discussed in Section 8.

Lemma 24. *Let C_α be a cell in a PSS. If it contains a vertex of $VD(Q)$, it is v_α.*

Proof. Suppose on the contrary that there exists a vertex v_β in C_α, s.t. $\alpha \neq \beta$. Lemma 15 implies that $\beta \subseteq \alpha$, and therefore $\beta \subset \alpha$. In the following we show that $dim(carbis(\beta)) > 0$, in contradiction to Lemma 8.5. α satisfies one of the conditions 2a–2h of Definition 1. Consider the following cases:

1. Condition 2a of Definition 1 holds. $|\alpha| \leqslant 4$. Then $|\beta| < 4$ in contradiction to Lemma 11.
2. One of the conditions 2b–2f of Definition 1 holds. The proof is identical for all these cases. Consider for example that condition 2b holds. $|\alpha| = 5$, and α includes an edge e and two coplanar faces f_1 and f_2 containing e. Let P be the plane carrying e, f_1, f_2. If two of e, f_1, f_2 are in β, then $\pi_P(v_\beta) \in e$, and therefore the third is also in β. Therefore $\beta = \{e, f_1, f_2, a\}$.

$$dim\big(carbis(\beta)\big) = dim\big(carbis(e, f_1, f_2) \cap carbis(a, e)\big) \geqslant dim\big(carbis(e, f_1, f_2)\big) - 1 = 1.$$

3. Condition 2g of Definition 1 holds. All the entities of α share a vertex v. $v = v_\alpha \neq v_\beta$. Let R be the ray from v through v_β. Let $S = R \cap C_\alpha$. We show in the following that $S \subset carbis(\beta)$. Therefore $dim(carbis(\beta)) > 0$.

Let b be an entity in β. Let x be a point in S. There exists a real number $t \geqslant 0$ s.t. $x = tv_\beta + (1-t)v$. We show in the following that $d(x, car(b)) = td(v_\beta, car(b))$. This implies that $x \in carbis(\beta)$, i.e., $S \subset carbis(\beta)$.
If $b = v$ then it is clear that $d(x, car(b)) = td(v_\beta, car(b))$. Otherwise $car(b)$ is a line or a plane passing through v. Consider the two triangles: $\triangle vx\pi_{car(b)}(x)$ and $\triangle vv_\beta\pi_{car(b)}(v_\beta)$. They are similar triangles, and therefore

$$\frac{d(v_\beta, car(b))}{d(v_\beta, v)} = \frac{d(x, car(b))}{d(x, v)}.$$

4. Condition 2h of Definition 1 holds. All the entities of α except one (a) share a vertex v and a plane P. $|\beta| \geqslant 4$. Therefore β contains at least three entities incident on P and containing v. The bisector of the carriers of three such entities is the line L orthogonal to P at v. Therefore $v_\beta \in L$. Every point on L is equidistant from *all* the carriers of entities incident on P and containing v, and therefore if $v_\beta \in L$, then $\alpha \setminus \{a\} \subseteq \beta$. If $a \in \beta$, then $\alpha = \beta$, and contradiction. If $a \notin \beta$, then $carbis(\beta)$ is L and therefore $dim(carbis(\beta)) > 0$. □

Lemma 25. *Let C be a cell in a PSS. Let $k > 0$ be the number of intersection points of a Voronoi edge e_β and ∂C. There exists a vertex of $VD(Q)$ in C iff k is odd.*

Proof. $carbis(\beta)$ is a 1-manifold curve (Lemma A.10). Therefore if $carbis(\beta) \cap C \neq \emptyset$, then $carbis(\beta) \cap C$ is composed of disjoint portions of $carbis(\beta)$, each homeomorphic to a linear segment. [3] Suppose there does not exist a vertex of $VD(Q)$ in C. Hence, if $carbis(\beta)$ enters C in a point in e_β, it exits C in a point in e_β (Lemma 10). Suppose there exists a vertex of $VD(Q)$ in C. This vertex is $v_{\beta\star}$. Assuming that the cell does not contain two vertices with the same governors, Lemma 24 states that there exists a single vertex in C. Therefore there is exactly one connected portion of $carbis(\beta)$ in which it enters into C in a point in e_β, and exits in a point outside of e_β (Lemma 10). □

Lemma 25 provides a criterion to decide whether a cell contains a Voronoi vertex. If no Voronoi edge intersects the cell, then the cell does not contain a Voronoi vertex, otherwise either there exists more than one Voronoi vertex in the cell, or the edge is a closed loop, in contradiction to Lemma 14. Voronoi vertices that are on the boundary of a cell are detected when computing Voronoi edge witnesses. There is one type of vertices that the criterion of Lemma 25 might not detect. The criterion will not detect a Voronoi vertex v_α s.t. for every edge e_β emanating from v_α, there exists another edge e_β emanating from v_α. Such a vertex v_α cannot be detected without computing its exact location. Such vertices can be thought of as vertices lying in the interior of edges; their presence results from a degenerate configuration.

5.3. Extraction of Voronoi edges

After computing edge witnesses and identifying Voronoi vertices, we identify Voronoi edges. We describe how to determine the edges of $VD(Q)$ and the incidence relationships between the edges and the vertices of $VD(Q)$. We first prove that the algorithm of Fig. 4 computes witnesses for *every* edge of $VD(Q)$ (Lemma 26).

[3] Unless $carbis(\beta)$ is tangent to C. This situation is avoided as explained in Section 5.1.

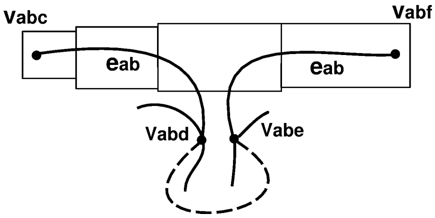

Fig. 6. v_{abc} and v_{abf} share a path of cells intersecting e_{ab}, but do not share an edge e_{ab}. Edges of the Voronoi diagram are shown by solid curves. The dashed curve shows a part of $carbis(a, b)$ that is not a Voronoi edge.

Lemma 26. *Let e be an edge of $VD(Q)$. There exists a cell C in a PSS s.t. e intersects the boundary of C.*

Proof. Suppose on the contrary that there exists a cell C and an edge e s.t. $e \subseteq int(C)$. Assuming that the cell does not contain two vertices with the same governors, there is at most one vertex of $VD(Q)$ in C (Lemma 24). Therefore e is a closed loop, in contradiction to Lemma 14. □

Lemma 26 implies that all edges of $VD(Q)$ are witnessed. In order to complete the identification of Voronoi edges, we have to determine which witnesses share the same Voronoi edge. Note that there may be several Voronoi edges having identical labels. We would like to say that two points $x, y \in R_\alpha$ share the same edge e_α if there exists a path of cells connecting them s.t. every pair of consecutive cells in the path shares a witness of e_α. This might be incorrect, as shown for the 2-D case in Fig. 6. Therefore we subdivide leaf cells with more than two witnesses of $carbis(\alpha)$. Lemma 27 proves that this refinement process halts. We call the resulting structure a *refined proximity structure subdivision*. Note that the new generated sub-cells also satisfy the halting conditions of the PSS process.

Lemma 27. *The refinement process defined above halts.*

Proof. Let C be a cell in a PSS. $carbis(\alpha)$ is an intersection of two quadratic surfaces, and therefore intersects a plane in a finite number of points ($\leqslant 4$).[4] Therefore it intersects C in a finite number of points. Therefore there is a finite number of portions of $carbis(\alpha)$ in C. Since $carbis(\alpha)$ is a 1-manifold curve (Lemma A.10), these intervals are disjoint, and each of them is homeomorphic to a linear segment. Let $m(C)$ be the minimal distance between two of these intervals. Since these intervals are disjoint $m(C) > 0$. A cell of size smaller than $m(C)$ contains only one interval of $carbis(\alpha)$, and therefore intersects $carbis(\alpha)$ in no more than two points. □

Lemma 28 (A criterion to determine whether two points share a Voronoi edge). *Let S be a refined PSS. Let α be a set of entities s.t. $dim(carbis(\alpha)) = 1$. Let x and y be points in e_α. Let C_x be a cell of S containing x, and let C_y be a cell of S containing y. x and y are incident on the same Voronoi edge e_α iff there exists a sequence of cells $C_1, \ldots, C_n$ s.t. $C_1 = C_x$, $C_n = C_y$, and C_i and C_{i+1} share a witness of e_α.*

[4] Unless $carbis(\alpha)$ is incident on the plane. This situation is avoided as explained in Section 5.1.

Proof. If x and y are incident on the same edge e_α, then it is clear that the condition is satisfied. Suppose now that the condition is satisfied. First we show that there exists a connected part P of $carbis(\alpha)$ which connects x and y and which is contained in the cells $C_1, \ldots, C_n$. Then we prove that P is wholly in e_α.

If $C_1, \ldots, C_n$ do not include a connected part of $carbis(\alpha)$, then the boundary of one of these cells intersects $carbis(\alpha)$ in more than two points, contradicting the fact that S is a refined PSS. Suppose on the contrary that P contains a point $x \in C_i$ s.t. $x \notin e_\alpha$. Then C_i contains two Voronoi vertices (Lemma 10), in contradiction to Lemma 24 (assuming that the cell does not contain two vertices with the same governors). □

Lemma 28 determines which witnesses share the same Voronoi edge. It also determines which Voronoi vertices share the same Voronoi edge. Thus determining the edges of $VD(Q)$ and the incidence relationships between the edges and the vertices of $VD(Q)$.

5.4. Extraction of Voronoi faces

Lemma 29. *A set $E = \{e_1, \ldots, e_n\}$ of Voronoi edges defines a Voronoi face f_α iff the following conditions are satisfied*:

1. $dim(carbis(\alpha)) = 2$.
2. *Every edge $e \in E$ is governed by $\alpha\star$.*
3. *There does not exist a set of entities $\beta \supset \alpha$ s.t. every edge $e \in E$ is governed by $\beta\star$.*
4. *E is connected, i.e., every two edges e_i and e_{i+1} share a vertex of $VD(Q)$.*

Proof. Suppose there exists a set of edges E as defined above. The set of edges E establish a connected region in R_α. $dim(carbis(\alpha)) = 2$, therefore this region is a Voronoi face f_α iff there does not exist $\beta \supset \alpha$ s.t. the region is contained in R_β.

Suppose there exists a face f_α. Then $dim(carbis(\alpha)) = 2$ (Lemma 8.5). f_α is simply connected (Lemma 13). Lemma 10 implies that f_α is bounded by a set of edges $e_{\alpha\star}$. Lemma 14 implies that it cannot be that all the edges of f_α are governed by β for $\beta \supset \alpha$. □

6. Dealing with degenerate diagrams

In Section 4 we assumed that $VD(Q)$ is not degenerate. If $VD(Q)$ is degenerate, then the subdivision process might not halt. In the following we describe the modifications that should be applied to the algorithm in order to handle degenerate diagrams as well.

The modifications are the following:

1. Subdivision process: An additional halting condition is added. The subdivision process is stopped also when the diameter of a cell is smaller than a given tolerance parameter ε. In the following we will refer to such cells as ε cells.
2. Extraction of the Voronoi graph from the subdivision:
   (a) ε cells are ignored in the extraction of Voronoi vertices.
   (b) The condition of Lemma 28 used in the extraction of Voronoi edges is modified as follows. Two points are incident on the same Voronoi edge iff there exists a sequence of cells $C_1, \ldots, C_n$ as defined in Lemma 28, *and* the intermediate cells are not ε cells.

In Section 5 we did not assume that the diagram is not degenerate, but we handled only cells that satisfy the conditions 2a–2e of a PSS cell (Section 4.1). Therefore applying the algorithm (with the above modifications) on a degenerate diagram, yields a correct Voronoi graph in the cells that are not ε cells. In the ε cells we know the governing entities, but we do not know how these governors share the cell. An ε cell is a small area where a degeneracy or an almost-degeneracy occurs. We do not want to further investigate these small areas, therefore we regard each ε cell as a single node in the Voronoi graph. Note that the extraction of the Voronoi edges emanating from the ε cells is correct.

The graph extracted by applying the above algorithm on a degenerate diagram is called an *Approximate Voronoi Graph* (*AVG*). An AVG approximates the Voronoi graph of Q to a tolerance of ε in the sense that a connected subgraph of the Voronoi graph that lies in a region of space of size smaller than ε is replaced by a single graph node.

Formally we define an approximate Voronoi graph as follows. Let G be an undirected graph s.t. every node is labeled by: (1) a set of entities of Q, (2) type: 'subgraph', 'face', 'edge' or 'vertex'. G is an *ε-approximation of the Voronoi graph of Q* if for every node n of type 'subgraph' there exists a subgraph G_n of the Voronoi graph of Q s.t. (1) G_n is governed only by the entities attached to n, (2) the part of $VD(Q)$ corresponding to G_n is bounded by a sphere of radius ε, and substitution of all such nodes n by their corresponding subgraphs G_n results in the Voronoi graph of Q.

7. The proximity structure diagram

The main contribution of this paper is the introduction and computation of the Voronoi graph, containing the *structure* of the Voronoi diagram of a polyhedron. In addition, the specific space subdivision algorithm that we use enables us to easily compute a quantifiable approximation to the *geometry* of the diagram as well.

We define a *Proximity Structure Diagram* (*PSD*) *of Q with a parameter* δ to be a Voronoi graph of Q s.t. every node of the Voronoi graph carries also a geometric approximation (of the appropriate type) to the corresponding element in $VD(Q)$, to an accuracy of δ. Formally, if h is a Voronoi element and h_a its geometric approximation, then $\forall x \in h, \exists y \in h_a$ s.t. $d(x, y) < \delta$ and $\forall y \in h_a, \exists x \in h$ s.t. $d(x, y) < \delta$.

We use the term 'proximity structure diagram' for what many readers would informally call an 'approximate Voronoi diagram'. We feel that the latter term is misleading, because it does not specify whether the approximation is of the connectivity of the Voronoi diagram, its geometry, or both. In our terminology, an AVG has approximate connectivity, and a PSD has exact connectivity and approximate geometry. The parameter controlling the connectivity approximation is ε, and the one controlling the geometry approximation is δ.

An easy way to construct a PSD is to first construct the Voronoi graph using the proximity structure subdivision algorithm, and then subdivide each cell that intersects a Voronoi edge until its diameter is smaller than δ. To obtain the desired approximation, we can either approximate directly in 3-D or work in the parameter space of the carrier surfaces of the entity bisectors. Direct 3-D approximation works best for vertices and edges, since centers of cells that contain Voronoi vertices, and piecewise linear curves connecting Voronoi edge witnesses, obviously provide δ approximations to the vertices and edges of $VD(Q)$. Faces are most efficiently approximated by representing them as trimmed surfaces in parameter space. Note that in this case if it is desired that the vertex, edge and face approximations be self-consistent then they must all be represented by mappings from parameter space.

8. Discussion

In this paper we introduced the Voronoi graph, the approximate Voronoi graph, and the proximity structure diagram of a polyhedron, and presented a simple approach to construct them for 3-D linear polyhedra. The Voronoi graph contains the complete symbolic information of the Voronoi diagram. The AVG and PSD complement each other in the sense that the first approximates the symbolic part of the Voronoi diagram and the second approximates the geometric part of the Voronoi diagram.

The skeletons are important for both theoretical and practical reasons. The main advantages of our computational approach are that it uses relatively low-degree algebraic operations in a single variable and that it enables local computation of the skeletons. Our results thus constitute a substantial improvement over the many previous approaches for computing Voronoi diagrams of 3-D polyhedra and for defining related approximations.

The algorithm has been implemented. Examples of its output are given in Figs. 7–9. Each of these figures includes a polyhedron and part of its Voronoi graph. The polyhedron edges are shown in black. The Voronoi graph does not contain any geometry; in order to visualize it, spheres denoting Voronoi vertices are displayed in the centers of the subdivision cells containing them, and gray polylines denoting Voronoi edges connect their Voronoi vertices while passing through the edge witnesses. Note that these edge polylines are not geometric approximations to the edges and are given only for visualization purposes. A geometric approximation could easily be made much more accurate.

In order to make the figures less cluttered, only part of the graph is displayed. The displayed part is the ‘central’ part of the graph: only its portion inside the polyhedron, and without Voronoi elements that are

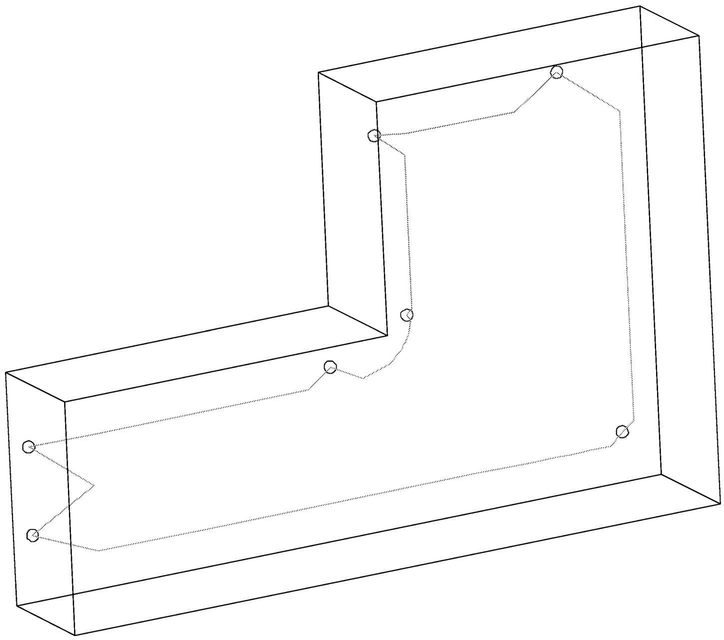

Fig. 7. Visualization of the central part of the Voronoi graph of the polyhedron. Polyhedron edges are shown as black lines, Voronoi edges as gray lines, and Voronoi vertices as spheres.

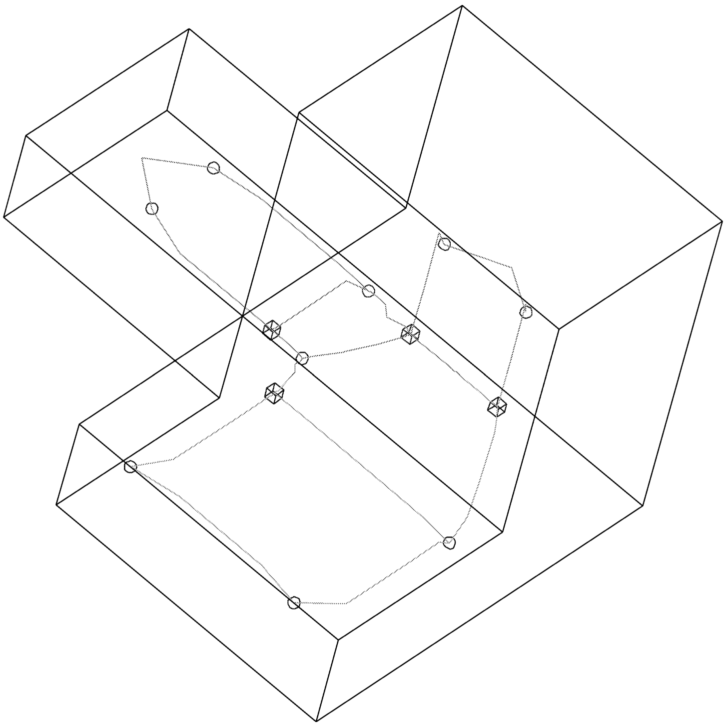

Fig. 8. ε cells are shown as cubes.

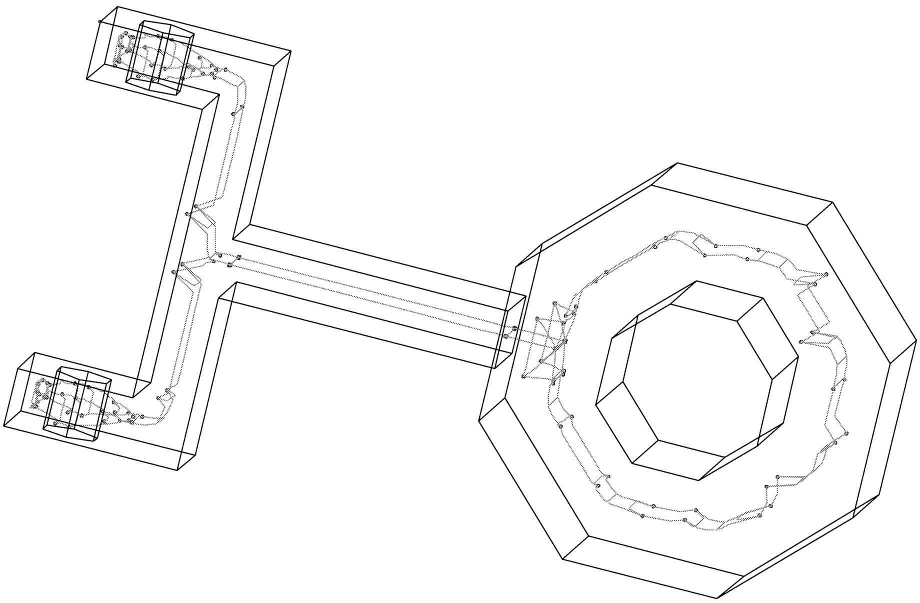

Fig. 9. A more complex example.

incident on polyhedron vertices (equivalent to the medial axis, without elements touching convex vertices and edges).

Fig. 7 shows a simple example. The polyhedron of Fig. 8 is degenerate, and therefore its PSS contains ε cells, denoted by small cubes. The geometry of the cubes is not identical to the geometry of the ε cells—a connected set of ε cells is displayed by a constant size cube. Fig. 9 shows a more complex part with three holes.

We assumed in this paper that the polyhedron's boundary is connected, and composed of convex faces. When the boundary is disconnected, the polyhedron contains cavities. In this case there might be (1) Voronoi edges that are loops, and (2) Voronoi faces that are multiply connected. A Voronoi edge that is a loop might be wholly in the interior of a cell (we have no example for such an occurrence). Such an edge will not be detected by the algorithm. The criterion to extract Voronoi faces should be extended if multiply connected Voronoi faces exist. If two Voronoi edges share the same loop in a Voronoi face f_{ab}, then there exists a sequence of Voronoi edges $e_{ab\star}$ connecting them. If two Voronoi edges share the same face f_{ab}, but not the same loop of f_{ab}, then there is a path in $carbis(a, b)$ connecting points in the two edges s.t. the interior of the path does not intersect an edge $e_{ab\star}$, and the path includes a point in R_a. While the first criterion can be implemented by finding paths in the already computed edge graph, the second criterion requires a search in the PSS and additional numerical computations similar to those executed when computing Voronoi edge witnesses.

Requiring that the faces of the polyhedron are convex makes both the proofs and the implementation simpler. This requirement does not limit the range of polyhedra handled by the algorithm. For any polyhedron Q, we can decompose its faces into convex pieces, compute the Voronoi diagram (or Voronoi graph or proximity structure diagram) of the resulting polyhedron Q', and then easily obtain the Voronoi diagram of Q from the Voronoi diagram of Q' in the following manner. For every element of $VD(Q')$ we know its set of governors in Q', and therefore its set of governors in Q. $VD(Q)$ is obtained from $VD(Q')$ by removing Voronoi elements whose set of governors in Q consist of a single entity, and by merging Voronoi edges (faces) whose connecting vertices (edges) were removed. This is how the part in Fig. 9 was handled.

The proofs in this paper are correct when assuming that there does not exist a cell with a multiplicity of Voronoi vertices all possessing the same set of governors (Section 5). If there exists a cell containing a multiplicity of Voronoi vertices, and all of these vertices are labeled by the same set of governors, then our algorithm might miss these vertices and identify the edges connecting them as the same edge. In all other cases the algorithm computes the correct result. Even in the former case, the inaccuracy in the Voronoi graph is limited to this specific cell, and the construction of the rest of the Voronoi graph is correct.

The skeletons introduced in this paper have many applications in geometric computing. For example, [20] presents a hexahedral mesh generation algorithm that uses the Voronoi graph to decompose the polyhedron into simple sub-volumes that are easy to mesh by basic methods. The medial axis of an object provides a natural subdivision of the object into simple parts. This application demonstrates that the exact location of the Voronoi elements is not always needed. The Voronoi graph contains enough information needed in order to determine where to decompose the polyhedron. If the polyhedron should be decomposed with respect to a specific Voronoi element, then a geometric approximation of this specific Voronoi element is computed. Fig. 10 shows the mesh generated using the algorithm of [20].

The focus in this paper has been on the new concepts and the correctness of the algorithm. The computational aspects, including implementational issues and timing are discussed in another paper [6].

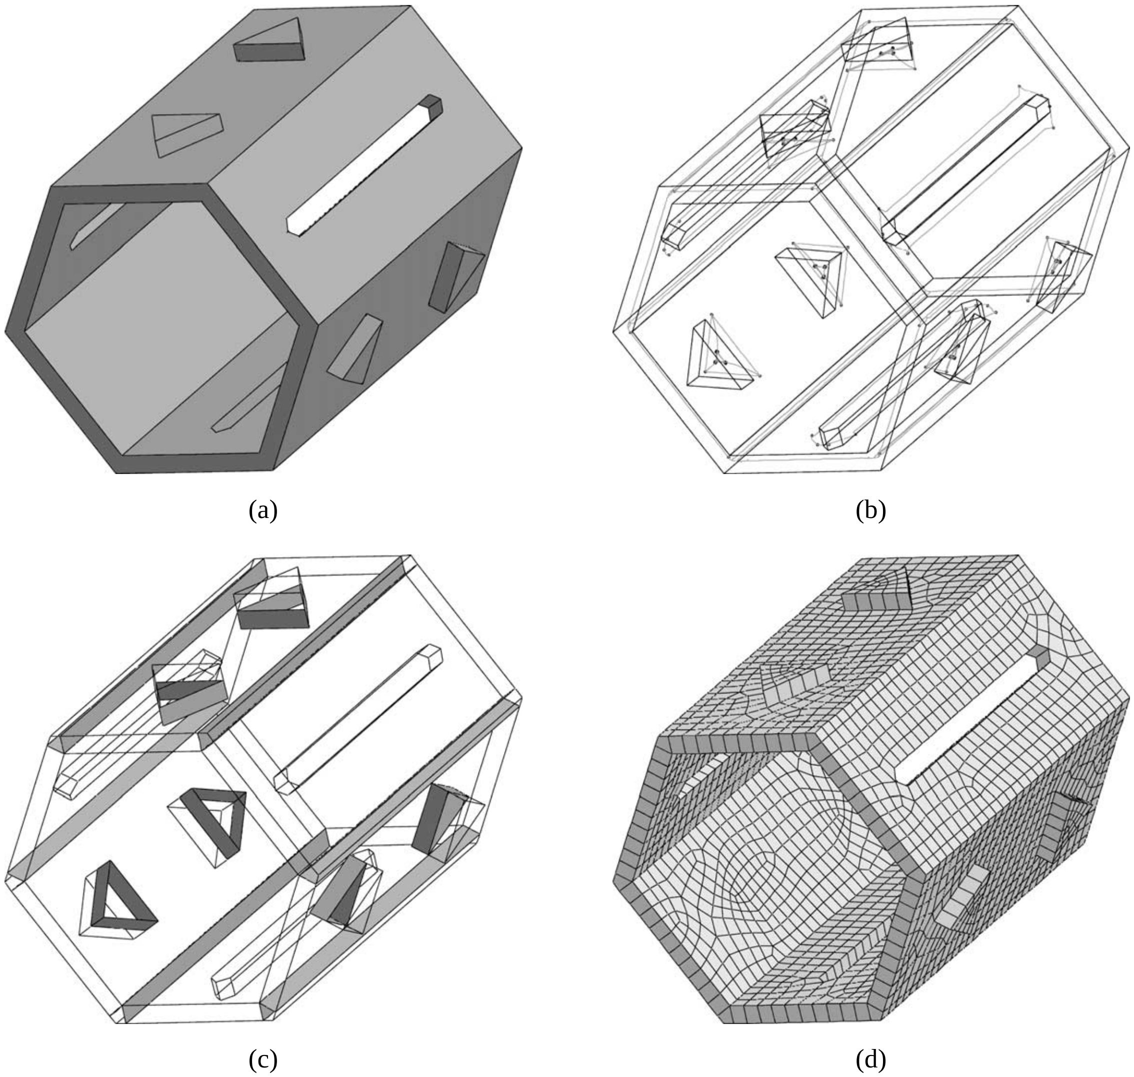

Fig. 10. Hexahedral mesh generation using the Voronoi graph: (a) the initial polyhedron; (b) the Voronoi graph of the polyhedron; (c) the decomposition faces generated based on the Voronoi graph; (d) the final mesh.

Additional topics for future work include enhancing the domain to curved polyhedra, and demonstrating further applications of the new skeletons.

Acknowledgements

We thank John Woodwark for a useful correspondence. Fluent provided the Gambit software used for generating the figures showing our algorithm's output. This research was supported by the Israeli Ministry of Science, the program for development of scientific and technological infrastructures and the Israel Science Foundation.

Appendix A. $carbis(a, b, c)$ is 1-manifold

Let α be a set of entities of Q. In this appendix we show that if $dim(carbis(\alpha)) = 1$, then $carbis(\alpha)$ is a 1-manifold curve. If $dim(carbis(\alpha)) = 1$ then α contains three entities a, b, c s.t. $carbis(\alpha) = carbis(a, b, c)$. Therefore it is sufficient to show that $carbis(a, b, c)$ is a 1-manifold curve for any three entities of a, b, c of Q.

This section is composed of two parts. In Appendix A.1, cases in which $carbis(a, b, c)$ might not be 1-manifold are identified, and the definition of $carbis(a, b, c)$ is slightly modified accordingly. In Appendix A.2 we prove that $carbis(a, b, c)$ is a 1-manifold curve, when using the new definition.

A.1. Splitting the bisectors

$carbis(a, b, c)$ might not be 1-manifold when a, b, c includes a plane, or two edges sharing a plane. In these cases $carbis(a, b, c)$ is composed of few 1-manifold parts. In order to split $carbis(a, b, c)$ into its 1-manifold components we use the notion of signed distance. The signed distance $d^\star$ between a point x and a plane P is defined as follows. If $x \in In(P)$, then $d^\star(x, P) = d(x, P)$, otherwise $d^\star(x, P) = -d(x, P)$. The signed distance between a point x and an oriented line L with respect to a plane P containing L, is defined as follows. If $\pi_P(x) \in In(L, P)$, then $d^\star(x, L) = d(x, L)$, otherwise $d^\star(x, L) = -d(x, L)$.

Lemma A.1. *Let a and b be two faces of Q. Suppose a and b are not parallel, and are not coplanar. $carbis(a, b)$ is composed of two planes P_1 and P_2 s.t. $x \in P_1$ iff $d^\star(x, car(a)) = d^\star(x, car(b))$, and $x \in P_2$ iff $d^\star(x, car(a)) = -d^\star(x, car(b))$.*

Lemma A.2. *Let a be a face of Q, and b be a vertex of Q. Suppose $a \not\supset b$. $carbis(a, b)$ is a paraboloid s.t. $x \in carbis(a, b)$ iff $d^\star(x, car(a)) = sign(d^\star(b, car(a))) * d(x, b)$.*

In the following when we say “half a cone”, we mean one part of the two parts of a cone obtained by intersecting the cone with a plane that intersects it only in its apex.

Lemma A.3. *Let a be a face of Q, and b be an edge of Q. Suppose $a \not\supset b$, and a and b are not parallel. $carbis(a, b)$ is a cone composed of two halves of a cone H_1 and H_2 s.t. $x \in H_1$ iff $d^\star(x, car(a)) = d(x, car(b))$, and $x \in H_2$ iff $d^\star(x, car(a)) = -d(x, car(b))$.*

Lemma A.4. *Let a and b be two edges of Q sharing a plane P. Suppose a and b are not parallel, and are not colinear. $carbis(a, b)$ is composed of two planes P_1 and P_2 s.t. $x \in P_1$ iff $d^\star(x, car(a)) = d^\star(x, car(b))$, and $x \in P_2$ iff $d^\star(x, car(a)) = -d^\star(x, car(b))$, where $d^\star$ is w.r.t. P.*

Let a and b be two entities that satisfy one of the following:
1. a and b are faces that are not parallel and are not coplanar.
2. a and b are two edges sharing a plane. a and b are not parallel and are not colinear.
3. a is a face and b is an edge. $a \not\supset b$, and a and b are not parallel.

Lemmas A.1–A.4 imply that $carbis(a, b)$ is composed of two parts, either two planes, or two halves of a cone. In the rest of Appendix A when we say $carbis(a, b, c)$, and a and b are of the types mentioned above, we mean the part of $carbis(a, b, c)$ that is incident on a specific half of $carbis(a, b)$. Lemmas A.5–A.6 prove that a Voronoi edge e_{abc} cannot be incident on two different halves of $carbis(a, b)$,

Lemma A.5. *Let e_α be a Voronoi edge, s.t. $|H(\alpha)| > 1$.*[5] *Let a be an entity in α that is a face of Q. Let x_1, x_2 be two points in e_α. $d^\star(x_1, car(a)) * d^\star(x_2, car(a)) > 0$.*

Proof. Consider the two cases:
1. $d^\star(x_i, car(a)) = 0$. Then $x_i \in car(a)$. Since $x_i \in R_a$, $x_i \in a$. $x_i \in e_\alpha$, therefore $x_i \in b$ for every $b \in \alpha$. Therefore $|H(\alpha)| = 1$. Contradiction.
2. $d^\star(x_1, car(a)) > 0$ and $d^\star(x_2, car(a)) < 0$. Then there exists a point $y \in e_\alpha$ s.t. $y \in car(a)$. $y \in e_\alpha$, and therefore $\pi_{car(a)}(y) \in a$. Therefore $y \in a$. $y \in e_\alpha$, therefore $y \in b$ for every $b \in \alpha$. Therefore $|H(\alpha)| = 1$. Contradiction. □

Lemma A.6. *Let e_α be a Voronoi edge, s.t. $|H(\alpha)| > 1$. Let a and b be two entities of α that are edges of Q, and share a plane. Let x_1, x_2 be two points in e_α. $d^\star(x_1, car(a)) * d^\star(x_2, car(a)) > 0$ and $d^\star(x_1, car(b)) * d^\star(x_2, car(b)) > 0$, where $d^\star$ is w.r.t. P.*

Proof. Consider the two cases:
1. $d^\star(x_i, car(a)) = 0$. Then $x_i \in car(a)$. Since $x_i \in R_a$, $x_i \in a$. $x_i \in e_\alpha$, therefore $x_i \in c$ for every $c \in \alpha$. Therefore $|H(\alpha)| = 1$. Contradiction.
2. $d^\star(x_1, car(a)) > 0$ and $d^\star(x_2, car(a)) < 0$. Let R be the plane orthogonal to P at a. There exists a point $y \in e_\alpha$ s.t. $\pi_P(y) \in car(a) \cap car(b)$. $y \in e_\alpha$, therefore $\pi_{car(a)}(y) \in a$. Therefore $\pi_P(y) \in a$. Similarly $\pi_P(y) \in b$. Therefore $\pi_P(y)$ is a vertex of a, b. Therefore $|H(\alpha)| = 1$. Contradiction. □

A.2. carbis(a, b, c) is 1-manifold

Lemma A.10 proves that $carbis(a, b, c)$ is 1-manifold. Lemmas A.7–A.9 are auxiliary lemmas of Lemma A.10.

Lemma A.7. *Let q be a point. Let L a line or a plane s.t. $q \notin L$. Let p be a point on $bis(q, L)$. If a plane T is tangent to $bis(q, L)$ at p, then $T = bis(q, \pi_L(p))$.*

Proof. In order to prove that $bis(q, \pi_L(p))$ is tangent to $bis(q, L)$ at p, it is sufficient to show that (1) every point $x \in bis(q, \pi_L(p))$ satisfies $d(x, L) \leqslant d(x, q)$ and (2) $p \in bis(q, \pi_L(p))$. (1) is correct since if $x \in bis(q, \pi_L(p))$ then $d(x, L) \leqslant d(x, \pi_L(p)) = d(x, q)$. (2) is correct since $d(p, q) = d(p, L) = d(p, \pi_L(p))$. □

Lemma A.8. *Let L_1 and L_2 be two lines that do not share a plane. Let p be a point on $bis(L_1, L_2)$. If a plane T is tangent to $bis(L_1, L_2)$ at p, then $T = bis(\pi_{L_1}(p), \pi_{L_2}(p))$.*

Proof. Let $p_1 = \pi_{L_1}(p)$. Let $p_2 = \pi_{L_2}(p)$. Let R_1 be the plane orthogonal to L_1 at p_1. Let $C_1 = R_1 \cap bis(L_1, L_2)$. We show in the following that $C_1 = R_1 \cap bis(p_1, L_2)$. Let x be a point in C_1. $d(x, p_1) = d(x, L_1) = d(x, L_2)$. Therefore $x \in R_1 \cap bis(p_1, L_2)$. Let x be a point $R_1 \cap bis(p_1, L_2)$. $d(x, L_1) = d(x, p_1) = d(x, L_2)$. Therefore $x \in C_1$. Therefore $C_1 = R_1 \cap bis(p_1, L_2)$. Therefore C_1 is intersection of a plane and a swept parabola, and therefore 1-manifold. $p \in C_1$. Let t_1 be the line tangent to C_1 at p. Since $C_1 \subseteq bis(p_1, L_2)$, t_1 is incident on the plane tangent to $bis(p_1, L_2)$ at p. Lemma A.7

[5] Recall that $H(\alpha) = \alpha \setminus \{a\colon\ a \supset b, b \in \alpha\}$. If $|H(\alpha)| = 1$, then no splitting of $carbis(\alpha)$ is done.

implies that this plane is $bis(p_1, p_2)$. Similarly we define R_2, C_2 and t_2. t_1 and t_2 are both incident on T and on $bis(p_1, p_2)$. We show in the following that $t_1 \neq t_2$. This implies that $T = bis(p_1, p_2)$.

Suppose on the contrary that $t_1 = t_2$. Let t be $t_1 = t_2$. $C_1 \subseteq R_1$, therefore $t = t_1 \subseteq R_1$. Similarly $t = t_2 \subseteq R_2$. Also $t \subseteq bis(p_1, p_2)$. Therefore every point $x \in t$ satisfies that $d(x, L_1) = d(x, p_1) = d(x, p_2) = d(x, L_2)$. Therefore t is a line incident on the swept parabola $bis(p_1, L_2)$. Therefore t is orthogonal to the plane of p_1 and L_2. Similarly t is orthogonal to the plane of p_2 and L_1. Therefore L_1 and L_2 share a plane. Contradiction. □

Lemma A.9. *Let R be a plane. Let L be a line s.t. $L \nsubseteq R$. Let p be the point on $bis(R, L)$. Let P be the plane passing through $\pi_L(p)$ and whose normal is $[p, \pi_L(p)]$.*
1. *$L \subset P$*
2. *If a plane T is tangent to $bis(R, L)$ at p, then $T = bis(R, P)$.*

Proof.
1. Let $x \in L$. $[x, \pi_L(p)]$ is orthogonal to $[p, \pi_L(p)]$. Therefore $x \in P$. Therefore $L \subset P$.
2. $p \in bis(R, P)$ since $d(p, R) = d(p, L) = d(p, \pi_L(p)) = d(p, P)$. Every point $x \in bis(R, P)$ satisfies that $d(x, R) = d(x, P) \leqslant d(p, L)$, since $L \subset P$. □

Lemma A.10. *If $dim(carbis(a, b, c)) = 1$, then $carbis(a, b, c)$ is a 1-manifold curve.*

Proof. Consider the following cases:
1. a, b, c are vertices. Then $car(a)$, $car(b)$ and $car(c)$ are points, and $carbis(a, b, c)$ is a line.
2. a, b, c are faces. Then $car(a)$, $car(b)$ and $car(c)$ are planes, and $carbis(a, b, c)$ is a line.
3. a and b are vertices, and c is a face. $car(a)$ and $car(b)$ are points, and $car(c)$ is a plane. $carbis(a, b)$ is a plane, and $carbis(a, c)$ either is a line or a paraboloid. Therefore $carbis(a, b, c)$ is either a line an intersection of a plane and a paraboloid. Therefore $carbis(a, b, c)$ is 1-manifold.
4. a and b are vertices, and c is an edge. $car(a)$ and $car(b)$ are points, and $car(c)$ is a line. $carbis(a, b)$ is a plane, and $carbis(a, c)$ is either a linear swept parabola or a plane. Therefore $carbis(a, b, c)$ is either the intersection of two planes or the intersection of a plane and a linear swept parabola, and therefore 1-manifold.
5. a and b are faces, and c is a vertex. $car(a)$ and $car(b)$ are planes, and $car(c)$ is a point. $carbis(a, b)$ is a plane, and $carbis(a, c)$ is either a line or a paraboloid. Therefore $carbis(a, b, c)$ is either a line or the intersection of a plane and a paraboloid. Therefore $carbis(a, b, c)$ is 1-manifold.
6. a and b are faces, and c is an edge. $car(a)$ and $car(b)$ are planes, and $car(c)$ is a line. $carbis(a, b)$ is a plane, and $carbis(a, c)$ is either a plane, or half a cone, or a swept parabola. The intersection of a plane with a plane or half a cone is a 1-manifold curve. The intersection of a plane with half a cone is not 1-manifold curve only if the plane is tangent to the cone. In this case, Lemma A.9 implies that $car(c) \subseteq car(b)$. Therefore $carbis(b, c)$ is a plane, and $carbis(a, b, c)$ is a line, i.e., a 1-manifold curve.
7. a is a vertex and b and c are edges. $car(a)$ is a point, and $car(b)$ and $car(c)$ are lines. Consider the two cases:
   (a) $a \in car(b)$ or $a \in car(c)$. Then $carbis(a, b, c)$ is the intersection of a plane and a swept parabola, and therefore it is a 1-manifold curve.
   (b) $a \notin car(b)$ and $a \notin car(c)$. Suppose on the contrary that $carbis(a, b, c)$ is not 1-manifold. Then there exists a point $p \in carbis(a, b, c)$ s.t. the tangent planes of $carbis(a, b)$ and $carbis(a, c)$ at

p are the same plane. Therefore $bis(a, \pi_{car(b)}(p)) = bis(a, \pi_{car(c)}(p))$ (Lemma A.7). Therefore $\pi_{car(b)}(p) = \pi_{car(c)}(p)$. Therefore $car(b)$ and $car(c)$ intersect, and therefore share a plane. In this case $carbis(b, c)$ is a plane, and $carbis(a, c)$ is either a linear swept parabola or a plane. Therefore $carbis(a, b, c)$ is either the intersection of a plane and a linear swept parabola, or the intersection of two planes, and therefore a 1-manifold curve.

8. a is a vertex, b is an edge, and c is a face. $car(a)$ is a point, $car(b)$ is a line, and $car(c)$ is a plane. Consider the three cases:
   (a) $a \in car(b)$. Then $carbis(a, b)$ is a plane, and $carbis(a, c)$ is a paraboloid. $carbis(a, b, c)$ is the intersection of a line and a paraboloid, i.e., a 1-manifold curve.
   (b) $a \in car(c)$. Then $carbis(a, c)$ is a line. Since $dim(carbis(a, b, c)) = 1$, $carbis(a, b, c)$ is a line.
   (c) $a \notin car(b)$ and $a \notin car(c)$. Suppose on the contrary that $carbis(a, b, c)$ is not 1-manifold. Then there exists a point $p \in carbis(a, b, c)$ s.t. the tangent planes of $carbis(a, b)$ and $carbis(a, c)$ at p are the same plane. Therefore $bis(a, \pi_{car(b)}(p)) = bis(a, \pi_{car(c)}(p))$ (Lemma A.7). Therefore $\pi_{car(b)}(p) = \pi_{car(c)}(p)$. Consider the two cases:
      i. $car(b) \subset car(c)$. Then $carbis(b, c)$ is a plane, and $carbis(a, b, c)$ is the intersection of a plane and a paraboloid, and therefore 1-manifold.
      ii. $car(b) \not\subset car(c)$. Then $car(b)$ and $car(c)$ intersect in a point q. $q = \pi_{car(b)}(p) = \pi_{car(c)}(p)$. If $q \neq p$, then $[p, q]$ is orthogonal to $car(b)$, and also $[p, q]$ is orthogonal to $car(c)$, and therefore $car(b) \subset car(c)$. Therefore $p = q$, and $q = \pi_{car(a)}(p) = a$. Therefore $a \in car(b) \cap car(c)$.
9. a is a face and b and c are edges. $car(a)$ is a plane, and $car(b)$ and $car(c)$ are lines. Consider the three cases:
   (a) $car(b) \subset car(a)$ or $car(c) \subset car(a)$. Suppose w.l.g. $car(b) \subset car(a)$. Then $carbis(a, b)$ is a plane, and $carbis(a, c)$ is either a plane, or half a cone, or a swept parabola. The intersection of two planes is a 1-manifold curve. The intersection of a plane and a swept parabola is a 1-manifold curve. The intersection of a plane and half a cone is not 1-manifold only if the plane is tangent to the cone. If $carbis(a, b)$ is tangent to $carbis(a, c)$, then Lemma A.9 implies that $car(c) \subseteq car(a)$. Therefore $carbis(a, c)$ is a plane, and $carbis(a, b, c)$ is a line.
   (b) b and c share a plane. Then $carbis(b, c)$ is a plane, and $carbis(a, c)$ is either a plane, or half a cone, or a swept parabola. The intersection of two planes is a 1-manifold curve. The intersection of a plane and a swept parabola is a 1-manifold curve. The intersection of a plane and half a cone is not 1-manifold only if the plane is tangent to the cone. If $carbis(b, c)$ is tangent to $carbis(a, c)$, then Lemma A.9 implies that $car(b) \subseteq car(a)$. Therefore $carbis(a, b)$ is a plane, and $carbis(a, b, c)$ is a line.
   (c) $car(b) \not\subset car(a)$, $car(c) \not\subset car(a)$ and b and c do not share a plane. Suppose on the contrary that $carbis(a, b, c)$ is not 1-manifold. Then there exists a point $p \in carbis(a, b, c)$ s.t. the tangent planes of $carbis(a, b)$ and $carbis(a, c)$ at p are the same plane. [6] Therefore b and c share a plane (Lemma A.9).
10. a, b, c are edges. $car(a)$, $car(b)$ and $car(c)$ are lines. Let k be the number of pairs of edges in $\{a, b, c\}$, s.t. a pair consists of two edges sharing a plane. Consider the following cases:
   (a) $k \geqslant 2$. Then $carbis(a, b, c)$ is the intersection of two planes, and therefore 1-manifold.

[6] If there does not exist a tangent plane to a cone at a point q, then q is the apex of the cone. If the apex q of the cone $carbis(a, b)$ is on $carbis(a, b, c)$, then $car(b)$ and $car(c)$ share a point (q), and therefore b and c share a plane.

(b) $k = 1$. Suppose w.l.g. a and b share a plane. Suppose on the contrary that $carbis(a, b, c)$ is not 1-manifold. Then there exists a point $p \in carbis(a, b, c)$ s.t. the tangent planes of $carbis(a, b)$, $carbis(a, c)$ and $carbis(b, c)$ at p are the same plane T. Since $carbis(a, b)$ is a plane, $T = carbis(a, b)$. Lemma A.8 implies that $T = bis(\pi_{car(a)}(p), \pi_{car(c)}(p)) = bis(\pi_{car(b)}(p), \pi_{car(c)}(p))$. Therefore $\pi_{car(a)}(p) = \pi_{car(b)}(p)$. Therefore $\pi_{car(a)}(p)$ is the intersection point of $car(a)$ and $car(b)$, and therefore $\pi_{car(a)}(p) \in carbis(a, b) = T$. Contradiction to $T = bis(\pi_{car(a)}(p), \pi_{car(c)}(p))$.

(c) $k = 0$. Suppose on the contrary that $carbis(a, b, c)$ is not 1-manifold. Then there exists a point $p \in carbis(a, b, c)$ s.t. the tangent planes of $carbis(a, b)$ and $carbis(a, c)$ at p are the same plane T. Lemma A.8 implies that $T = bis(\pi_{car(a)}(p), \pi_{car(b)}(p)) = bis(\pi_{car(a)}(p), \pi_{car(c)}(p))$. Therefore $\pi_{car(b)}(p) = \pi_{car(c)}(p)$. Therefore $car(b)$ and $car(c)$ intersect. Therefore b and c share a plane. Contradiction. □

References

[1] C.G. Armstrong, Modeling requirements for finite-element analysis, Computer-Aided Design 26 (7) (1994) 573–578.

[2] F. Aurenhammer, Voronoi diagrams: a survey of a fundamental geometric data structure, ACM Computing Surveys (1991) 345–405.

[3] E. Bertin, J.M. Chassery, A 3D generalized Voronoi diagram for a set of polyhedra, in: Curves and Surfaces in Geometric Design, Peters, Wellesley, MA, 1994, pp. 43–50.

[4] J. Canny, B. Donald, Simplified Voronoi diagrams, Discrete and Computational Geometry 3 (1988) 219–236.

[5] M. Etzion, Computing the Voronoi graph and the Voronoi diagram of a 3-D solid, PhD Thesis, The Hebrew University, Jerusalem, 1999.

[6] M. Etzion, A. Rappoport, Computing the Voronoi diagram of a 3-D polyhedron by seperate computation of its symbolic and geometric parts, in: ACM Symposium on Solid Modeling and Applications, MI, 1999, pp. 167–178.

[7] S. Fortune, Voronoi diagrams and Delaunay triangulations, in: D.-Z. Du, F. Hwang (Eds.), Computing in Euclidean Geometry, World Scientific, Singapore, 1992, pp. 193–234.

[8] C.M. Hoffmann, How to construct the skeleton of CSG objects, in: A. Bowyer (Ed.), Computer-Aided Surface Geometry and Design, Oxford University Press, 1994, pp. 421–437 (Proceedings, Mathematics of Surfaces IV).

[9] D. Lavender, A. Bowyer, J. Davenport, A. Wallis, J. Woodwark, Voronoi diagrams of set-theoretic solid models, IEEE Computer Graphics and Applications 12 (5) (1992) 69–77.

[10] M. Mantyla, An Introduction to Solid Modeling, Computer Science Press, MD, 1988.

[11] V.J. Milenkovic, Robust construction of the Voronoi diagram of a polyhedron, in: Proc. Fifth Canadian Conference on Computational Geometry, 1993, pp. 473–478.

[12] A. Okabe, B. Boots, K. Sugihara, Spatial Tessellations: Concepts and Applications of Voronoi Diagrams, Wiley, 1992.

[13] N.M. Patrikalakis, H.N. Gursoy, Shape interrogation by medial axis transform, in: B. Ravani (Ed.), Advances in Design Automation, Vol. 1: Computer Aided and Computational Design, ASME, 1990.

[14] J.M. Reddy, G.M. Turkiyyah, Computation of 3D skeletons using a generalized Delaunay triangulation technique, Computer-Aided Design 27 (9) (1995) 677–694.

[15] G. Renner, I.A. Stroud, Medial surface generation and refinement, Proc. IFIP Workshop on Geometric Modeling in CAD, Airlie, VA, 1996.

[16] M. Rezayat, Midsurface abstraction from 3D solid models: general theory and applications, Computer-Aided Design 28 (11) (1995) 905–915.
[17] D.J. Sheehy, C.G. Armstrong, D.J. Robinson, Computing the medial surface of a solid from a domain Delaunay triangulation, IEEE Transactions on Visualization and Computer Graphics 2 (1) (1996) 62–72.
[18] E.C. Sherbrooke, N.M. Patrikalakis, E. Brisson, Computation of the medial axis transform of 3D polyhedra, IEEE Transactions on Visualization and Computer Graphics 2 (1) (1996) 44–61.
[19] E.C. Sherbrooke, 3-D shape interrogation by medial axis transform, PhD Thesis, MIT, 1995.
[20] A. Sheffer, M. Etzion, A. Rappoport, M. Bercovier, Hexahedral mesh generation using Voronoi skeletons, in: Proc. Seventh International Meshing Roundtable ’98, 1998, pp. 347–364.
[21] D.W. Storti, G.M. Turkiyyah, M.A. Ganter, C.T. Lim, D.M. Stal, Skeleton-based modeling operations on solids, in: ACM Symposium on Solid Modeling Foundations and Applications, Atlanta, May 1997, pp. 141–154.
[22] A. Sudhalkar, L. Gursoz, F. Prinz, Continuous skeletons of discrete objects, in: ACM Symposium on Solid Modeling Foundations and Applications, UT, 1993, pp. 85–94.
[23] G.M. Turkiyyah, D.W. Storti, M. Ganter, H. Chen, M. Vimawala, An accelerated triangulation method for computing the skeletons of free-form solid models, Computer-Aided Design 29 (1) (1997) 5–19.
[24] J. Vleugels, M. Overmars, Approximating generalized Voronoi diagrams in any dimension, Int. J. Comp. Geom. Appl. 8 (1998) 201–221.
[25] X. Yu, J.A. Goldak, L. Dong, Constructing 3D discrete medial axes, in: ACM Symposium on Solid Modeling, Austin, 1991, pp. 481–489.