

Structure and Computation of Straight Skeletons in 3-Space^{*}

Franz Aurenhammer and Gernot Walzl

Institute for Theoretical Computer Science,
Graz University of Technology, Graz, Austria
{auren, gernot.walzl}@igi.tugraz.at

Abstract. We characterize the self-parallel (mitered) offsets of a general nonconvex polytope \mathcal{Q} in 3-space and give a canonical algorithm that constructs a straight skeleton for \mathcal{Q} .

1 Introduction

The *straight skeleton* of a polygon P in the plane is a versatile skeletal structure, composed of angular bisectors of P and thus of piecewise linear components [1,8,12]. A well-known procedural definition exists, by a self-parallel shrinking process for P and the resulting ‘events’ that construct the skeleton nodes. Events are unique changes in the polygon boundary, yielding the *mitered offset* of P , in contradistinction to the Minkowski sum offset specified by the medial axis of P ; see e.g. [7,11]. Applications arise in diverse areas, including computer graphics, robotics, architecture, and geographical information systems.

To construct a straight skeleton for a given polytope \mathcal{Q} in 3-space, one tries to proceed in a way analogous to the planar case. The facet planes of \mathcal{Q} are offset simultaneously, self-parallel, and at unit speed. Thereby, the shrinking polytope undergoes changes of geometric, combinatorial, and topological nature. Geometrical changes, of course, take place continuously, whereas combinatorial changes (on the polytope boundary) and topological changes (like new tunnels, or breaking the polytope locally or globally apart) occur once in a while. Each type of change implies the former ones.

During the offsetting process, the edges and vertices of \mathcal{Q} trace out the facets and edges, respectively, of the spatial skeleton. Thus a concise specification of the offsetting process results in a concise definition of the skeleton. However, unlike parallel offsets of polygons, parallel offsets of polytopes in \mathbb{R}^3 are in general not unique. This might be among the reasons why not much literature exists on this topic. Barequet et al. [4] studied straight skeletons in 3-space, mainly for the special case of orthogonal (i.e., axes-aligned) polytopes, where the skeleton is the medial axis of the polytope for the L_∞ -metric (and thus can be defined via distances). They mention ambiguity problems for the case of general polytopes, referring to Demaine et al. [6]. Also, a nice offsetting possibility by means of

^{*} Supported by ESF Programme EuroGIGA - Voronoi, Austrian Science Foundation.

the planar weighted straight skeleton [2] is observed, but the offset treatment of general vertices remains unclear (as is explained in Subsection 2.2). In fact, it remains open whether a boundary-continuous and non-selfintersecting offset polytope always has to exist. For orthogonal polytopes, Martinez et al. [10] recently described an implementation of a straight skeleton construction algorithm that uses the space-sweep technique.

In this paper we give a systematic treatment of offsetting, in terms of generalized *vertex figures* of a given polytope \mathcal{Q} in \mathbb{R}^3 , and their spherical *bisector graphs* (Subsection 2.1). This includes a characterization of all possible offsets for \mathcal{Q} , and consequently, of all possible spatial straight skeletons for \mathcal{Q} .

As an algorithmic tool, the so-called *spherical skeleton* is introduced (Subsection 2.2), a generalization to the unit sphere of the classical unweighted straight skeleton in the plane. This structure can uniformly treat arbitrary polytope vertices, and all types of events that occur during the construction of a straight skeleton in \mathbb{R}^3 , including those which change the topology of the polytope. A canonical offsetting algorithm results (Section 3), which produces a small number of edges, among all offset possibilities, and comes with a certain optimality property concerning edge convexity. The algorithm also provides a means for converting boundary-triangulated (*simplicial*) nonconvex polytopes in \mathbb{R}^3 into 3-regular (*simple*) ones, via ε -thinning; see e.g. [11].

We define a *polytope* \mathcal{Q} in 3-space \mathbb{R}^3 as a bounded, closed, and interior-connected subset of \mathbb{R}^3 with piecewise linear boundary. The boundary components of \mathcal{Q} of dimensions 2, 1, and 0, respectively, are called *facets*, *edges*, and *vertices* of \mathcal{Q} . In the easiest case, \mathcal{Q} is homeomorphic to a ball in \mathbb{R}^3 .

A polytope is, in general, nonconvex and may contain tunnels, and even ‘voids’ that make its boundary disconnected. However, if the input polytope \mathcal{Q} is boundary-connected, no holes can be created in its offsetting process.

An edge e of \mathcal{Q} is called *reflex* if there exists some line segment $\ell \subset \mathcal{Q}$ whose interior intersects the interior of e in a single point. Otherwise, edge e is called *convex*. (Topological notion like interior, boundary, etc. is meant relative to the dimension of an object.)

2 Vertex Figure Resolution

2.1 Offset Characterization

In the offsetting process, the local facial structure of the input polytope has to be updated, at existing or arising vertices of degree 4 or higher (apart from certain degenerate cases). We describe these changes in terms of so-called vertex figures of a polytope [9].

Consider some vertex, v , of the current polytope \mathcal{Q} . Center a sphere at v , sufficiently small to intersect only faces of \mathcal{Q} incident to v . We can assume that, w.l.o.g., this sphere is the unit sphere, U . The *vertex figure*, $\mathcal{F}(v)$, of v is then defined as the intersection of \mathcal{Q} with the unit ball; see Figure 1 (left).

The boundary of \mathcal{Q} intersects U in a Jordan curve \mathcal{J} composed of great arc segments on U . (For ease of exposition, assume only one such curve for now.)

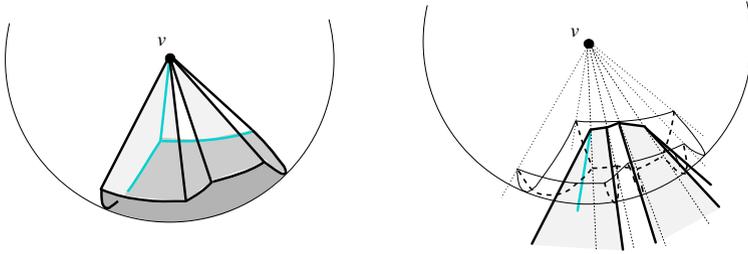


Fig. 1. Vertex figure and spherical polygon \mathcal{P} (left). Offsetting the vertex figure (right). The bisector graph G (dashed style) used for projecting is unique for \mathcal{P} in this example. The degree-6 vertex v splits into four vertices of degree 3 in the offset surface $\Gamma(G)$.

Then $\mathcal{Q} \cap U$ is a simply connected spherical polygon, \mathcal{P} . The curve \mathcal{J} consists of *nodes* and *segments*; the former are the intersections of U with the edges of \mathcal{Q} incident to v , and the latter belong to the boundary of the planar facets f_1, \dots, f_m of the vertex figure $\mathcal{F}(v)$. Here $m \geq 3$ is the degree of vertex v .

Let now H_i be the supporting plane of facet f_i . Denote with H_i^Δ the parallel offset of H_i by $\Delta > 0$, inward with respect to $\mathcal{F}(v)$. An (*inward*) *offset* of $\mathcal{F}(v)$ is defined as a radially monotone (w.r.t. v) and continuous surface over \mathcal{P} with facets from the planes H_i^Δ , including an unbounded facet for each plane.

For increasing Δ , the intersection line $\ell_{ij} = H_i^\Delta \cap H_j^\Delta$ of two offset planes sweeps along an angular bisector plane, B_{ij} , of H_i and H_j . The point of intersection of three offset planes moves along the trisector line of these planes. Note that these bisector planes and trisector lines all pass through the vertex v .

We next consider the picture on the unit sphere U . Define $b_{ij} = B_{ij} \cap U$, which is a great circle on U . Exactly three great circles b_{ij}, b_{ik}, b_{jk} meet in two common points, which are the intersections of U with the respective trisector line. (We only discuss the non-degenerate situation here.) For a system of circles with this property, we can consider any *bisector graph* G they define inside the spherical polygon \mathcal{P} .

G is defined as a crossing-free graph with labelled arcs $a_{ij} \subseteq b_{ij}$, where the ordering of the labels (ij) indicates the position of the facets f_i and f_j with respect to the bisector plane B_{ij} . G contains nodes of degree 1 (the nodes of \mathcal{P}), and of degree 3 whose incident arcs have labels of the form $(ij), (ik), (kj)$. G can be disconnected and can even contain cycles. We have the following characterization.

Lemma 1. *For every offset surface for $\mathcal{F}(v)$ its edge graph maps, by central projection with respect to v , to some bisector graph for the system $(b_{ij})_{1 \leq i < j \leq m}$. Conversely, each bisector graph G for this system lifts, by this projection, to the edge graph of an offset surface $\Gamma^\Delta(G)$ for $\mathcal{F}(v)$, which is unique for fixed $\Delta > 0$.*

Proof. From the way how offset planes intersect, it is obvious that each valid offset surface radially projects to a crossing-free degree-3 graph with the required labelling. To prove the converse, we lift G 's arcs a_{ij} to the lines $\ell_{ij} = H_i^\Delta \cap H_j^\Delta$ by central projection with respect to v . This can be done because a_{ij}, ℓ_{ij} , and v

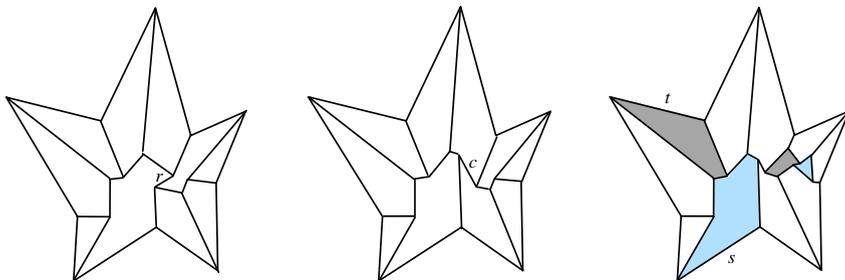


Fig. 2. Different offset surfaces for a vertex figure of degree 10. The figure is based on a pentagonal star \mathcal{P} (of small area, so that \mathcal{P} is almost flat.) Our algorithm will produce the offset with the convex edge c (middle), avoiding the reflex edge r (left). The bisector graph is highly ambiguous and even need not be a tree. Additional small facets for two segments s and t bounding \mathcal{P} might show up in the surface, either together (right), or separately (not shown).

are contained in the same plane, B_{ij} . Each connected face of G thus lifts to a polygon in 3-space, which is planar because its edges e_{ij} are labelled with the same offset plane index j on the ‘inside’, by the labelling of G . That is, G lifts to a unique piecewise linear surface, whose facets fit continuously because their edges e_{ij} are part of ℓ_{ij} and thus lie in both offset planes H_i^Δ and H_j^Δ . The surface is radially monotone because G is crossing-free.

See Figure 1 (right) for an illustration. The combinatorial structure of $\Gamma^\Delta(G)$ does not change for $\Delta > 0$, so we may just write $\Gamma(G)$.

In general, there is more than one valid bisector graph for a given system (b_{ij}) . In fact, there can be exponentially many, in the degree m of v (which can be obtained, for example, by combining the graphs shown in Figure 2 into bigger ones). Consequently, there may be more than one valid offset surface for a given vertex figure. That is, parallel offsets of polytopes in \mathbb{R}^3 , and with it, their straight skeletons, are not unique.

2.2 Spherical Skeleton

By Lemma 1, offsets of vertex figures – and in fact, of the entire polytope Q – can be computed by generating bisector graphs. In principle, any valid bisector graph for a spherical polygon \mathcal{P} could be extracted from the arrangement of the great circles b_{ij} defined on the unit sphere U . A more direct and also unique way is to construct the *spherical skeleton* of \mathcal{P} , denoted by $\text{Sph}(\mathcal{P})$. We now define this skeleton, by describing its events (and peculiarities).

We simulate the offsetting process, by considering the intersection of a vertex figure’s facets with U . That is, we offset the Jordan curve \mathcal{J} that bounds \mathcal{P} , inwards with respect to \mathcal{P} . The moving curve is denoted by \mathcal{J}^Δ , where Δ means offset distance, hence time. (The case ‘outwards’ is analogous, and the case ‘more than one Jordan curve’ is similar.)

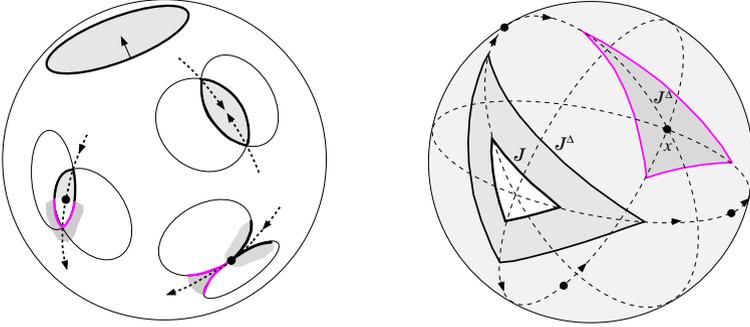


Fig. 3. Void event and arc merge event (left upper). Two stopping events; the interior of the offsetting region is shaded (left lower). A triangular Jordan curve \mathcal{J} and two offsets \mathcal{J}^Δ , for $\Delta < 1$ and $\Delta > 1$ (right). The three extending skeleton arcs (dashed) stop once each, at the marked positions. They meet at point x where \mathcal{J}^Δ vanishes in a triangle collapse.

\mathcal{J}^Δ 's nodes move on the great circles b_{ij} and draw out the skeleton arcs. The moving nodes bound shifted segments s_i from \mathcal{J} on circles $H_i^\Delta \cap U$ whose radii all shrink identically. The speed of an expanding skeleton arc a_{ij} is given by the (fixed) angle between the two planes H_i^Δ and H_j^Δ and by the time Δ . As one possible event, three such arcs a_{ij} , a_{jk} , and a_{ik} meet at the same time at a point of intersection of U with the respective trisector line. This means a *segment collapse* for s_j , making \mathcal{J}^Δ lose this segment. Or, one of the arcs, say a_{ij} , bumps into some shifting segment s_k and splits it (*segment split* event). This breaks \mathcal{J}^Δ into two closed curves, and lets two arcs a_{ik} and a_{jk} start at the breaking point. In both cases, a new node of the skeleton $\text{Sph}(\mathcal{P})$ is created, quite similar to the planar skeleton case [1]. Also similar is the simultaneous collapse of three segments forming a spherical 3-gon (*triangle collapse*). On the other hand there are new events, for example the collapse of a spherical 2-gon (*arc merge*) or of a spherical 1-gon, i.e., a full circle (*void event*); see Figure 3 (left). These events do not create nodes of $\text{Sph}(\mathcal{P})$.

If the vertex figure $\mathcal{F}(v)$ based on \mathcal{P} is *pointed*, that is, if there exists a plane through v that has \mathcal{P} on a fixed side, then the construction of $\text{Sph}(\mathcal{P})$ will be completed at time $\Delta < 1$, when all the planes H_i^Δ are still intersecting the sphere U . Otherwise, more substantial differences to the planar case occur: At time $\Delta = 1$, the planes H_i^Δ will start avoiding U , although $\text{Sph}(\mathcal{P})$ is still incomplete. To continue the construction, we let the planes return and intersect U in circles with negative radii, which are expanding now. (That is, the *complements* of the disks they bound are shrinking now, rather than the disks themselves.) Arc extension along b_{ij} , which stopped at time $1 - \Delta_{ij}$ and at point p_{ij} when and where the line $\ell_{ij} = H_i^\Delta \cap H_j^\Delta$ left U , continues after this *stopping event* at p_{ij} at time $1 + \Delta_{ij}$, with the arc endpoint carrying two segments $\overline{s_i}$ and $\overline{s_j}$ of mirrored shape which stem from negative circles; see Figure 3 (left lower). The convexity status of their common endpoint alters thereby.

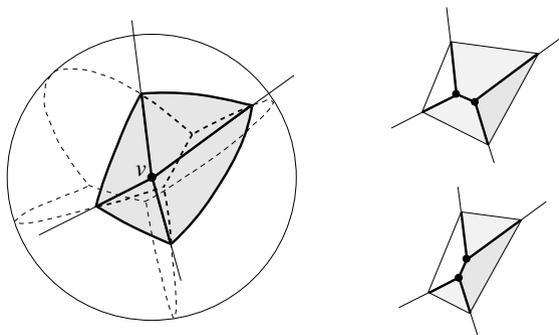


Fig. 4. Vertex touch; the convex case. A polytope edge shrinks to length zero, creating a vertex figure $\mathcal{F}(v)$ with four convex edges (left). Its base is a spherical rectangle \mathcal{P} . The two longest segments of \mathcal{P} 's boundary curve \mathcal{J} yield an arc in the spherical skeleton (dashed style, inside \mathcal{P}) that connects two nodes. Accordingly, vertex v splits into two offset vertices after the event (right lower).

For the inverse event, the outward offset is relevant. The spherical skeleton now lives in the outer hemisphere $U \setminus \mathcal{P}$, which corresponds to the complement of $\mathcal{F}(v)$. Four skeleton arcs extend from \mathcal{J} around U , but stop in between this time. The two shortest segments of \mathcal{J} yield a skeleton arc now. Vertex v splits into two vertices again, yielding different facet adjacencies that describe the situation before the event (right upper).

(Intuitively speaking, the returning planes collect up their stopped arc endpoints in reverse order, letting them continue at the right time.)

This ‘positive/negative switch’ at time $\Delta = 1$ ensures that, for each trisector line, both intersection points with U are taken into account as possible skeleton nodes; see Figure 3 (right). For each stopped arc endpoint, its two incident segments are temporarily inactive for split events.

$\text{Sph}(\mathcal{P})$ is a bisector graph which is outerplanar. Each segment s_i of \mathcal{J} sweeps out a single region, which is incident to s_i . (Segments on the same great circle might give the same region.) For, once all segments for an offsetting circle $H_i^\Delta \cap U$ have collapsed (unless in a stopping event), the plane H_i^Δ cannot define any further piece of $\text{Sph}(\mathcal{P})$. Thus $\text{Sph}(\mathcal{P})$ contains $\leq m$ regions and $O(m)$ arcs and nodes, where m is the number of segments of \mathcal{J} . Figures 4 to 9 give illustrations.

Using *weighted* planar straight skeletons [2] to resolve high-degree vertices has been considered in Barequet et al. [4]. In their approach, partial skeletons in two osculating planes need to be merged into one skeleton. Thereby, components from one plane might destroy the skeleton structure in the other, however. For example, skeleton faces can expand, rather than shrink as they do for Voronoi-like partitions. It remains unclear how to deal with this problem, which is equivalent to the (unsolved) problem of computing a planar straight skeleton by divide & conquer, or incremental insertion.

To base the offsetting process on $\text{Sph}(\mathcal{P})$ is natural, as it is consistent with the well-known process that constructs the straight skeleton in \mathbb{R}^2 . This choice is also preferable, not least because of the linear size of $\text{Sph}(\mathcal{P})$. For general bisector graphs, examples with $\Omega(m^2)$ interior faces can be constructed. Also,

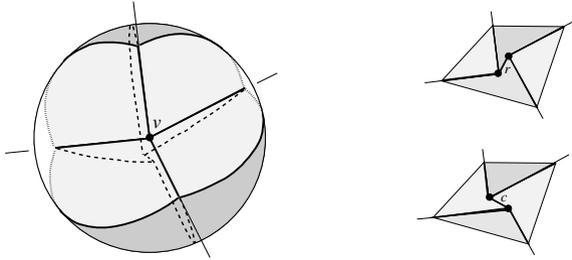


Fig. 5. Vertex touch; saddle point. $\mathcal{F}(v)$ has four edges, of alternating convexity type (left). Unlike the convex case in Fig. 4, two offsets exist (right). The spherical skeleton produces the lower solution. The short skeleton arc corresponds to a convex edge c in the offset polytope. It connects the two vertices which v is split into. In the upper solution, the offset edge r is reflex.

$\text{Sph}(\mathcal{P})$ is unweighted, and thus behaves like the classical planar straight skeleton in certain respects. For example, if $\mathcal{F}(v)$ is pointed then one can show that no new *reflex* edges are generated in the offset surface $\Gamma(\text{Sph}(\mathcal{P}))$, only the ‘forced’ ones that are also present in $\mathcal{F}(v)$. A similar property holds for the ‘roof’ of the planar straight skeleton [1], but not for weighted straight skeletons, in general.

$\text{Sph}(\mathcal{P})$ can be computed in $O(m^2 \log m)$ time with a trivial implementation, like in the planar case [1]. Note that m (the degree of the vertex figure) will be usually quite small in practice, for two reasons: Most solids can be accurately approximated by polytopes having vertices of small constant degree. Also, in the non-degenerate case, the straight skeleton events in \mathbb{R}^3 lead to vertex figures of degree at most 8, as we shall see in the following section.

3 Straight Skeleton Algorithm in \mathbb{R}^3

When attempting to construct a straight skeleton in \mathbb{R}^3 , the vertex figure resolution problem arises immediately, when vertices of degree $m \geq 4$ of the input polytope \mathcal{Q} have to be split. In the non-degenerate case, such a vertex splits into $m - 2$ vertices of degree 3 in the shrunken polytope. Interestingly, by nature of the straight skeleton events in \mathbb{R}^3 , the vertex resolution problem is encountered again each time a new event is to be handled.

An *event* in this sense is an increase in the number of facets for some vertex of the shrinking polytope (for example, by touch with another degree-3 vertex). We are then left with the problem of resolving that vertex, to continue the skeleton construction for \mathcal{Q} . As we have seen in Subsection 2.1, there are several possibilities to proceed. This gives a tree of possible offset polytopes, with root \mathcal{Q} . Each path from \mathcal{Q} to a tree leaf corresponds to a different straight skeleton for \mathcal{Q} . (If \mathcal{Q} is *convex* then there is only one path, leading to the *medial axis* of \mathcal{Q} .)

When the spherical skeleton approach from Subsection 2.2 is used, we always obtain a unique path. This path is short because offsets with a small number of edges are created, especially in the beginning where high-degree vertices of \mathcal{Q}

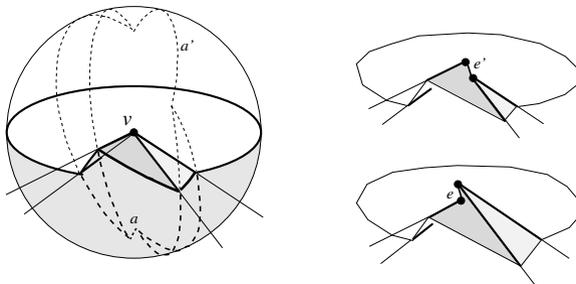


Fig. 6. Vertex touch; pointed saddle point. Again, $\mathcal{F}(v)$ has two reflex and two convex edges, but they do not positively span \mathbb{R}^3 as in Fig. 5. The skeleton lives in the lower hemisphere. It is generated by a segment collapse and a triangle collapse, followed by a void event at the south pole. The short skeleton arc, a , gives a convex edge e in the offset polytope (right lower), where the horizontal facet and the dark-shaded facet are adjacent. The same adjacency existed already as edge e' in the polytope before the event (right upper), indicated by arc a' in the skeleton in the upper hemisphere, which specifies the inverse event. Observe that a and a' lie on the same great circle.

may be present. The skeleton $\text{Sph}(\mathcal{P})$ does not only encode the local facial structure at a vertex v after the event, but can also restore this structure *before* the event (except for certain collapse events, see later). To this end, the *outward* offset of the vertex figure $\mathcal{F}(v)$ based on \mathcal{P} is considered. This is just the (inward) offset of the complement of $\mathcal{F}(v)$. We obtain it by using the spherical skeleton $\text{Sph}(U \setminus \mathcal{P})$, as is illustrated in Figures 4, 6, 7, 8, and 9. In this way, an event can be associated with a unique *inverse event*.

As a consequence, the resulting unique straight skeleton, $\text{SKEL}(Q)$, for the polytope Q is given locally at each of its vertices v by

$$\text{SKEL}(Q) \cap U = \text{Sph}(\mathcal{P}) \cup \text{Sph}(U \setminus \mathcal{P}) .$$

Our vertex resolution procedure works, without conceptual changes, for all occurring degenerate cases, be they by the special structure of Q (for example, regular pyramid vertices that offset into vertices of degree ≥ 4) or by construction (during event handling, where coplanarities may arise naturally).

3.1 Event Categorization

Categorizing the (non-degenerate) events that construct $\text{SKEL}(Q)$ is relatively easy. A general distinction is between *surface events* which change the combinatorial structure of Q 's edge graph without topological effect, and *solid events* which also change the topology of Q . For the former, the vertex figure yields a single Jordan curve on U , whereas it gives two disjoint curves for the latter.

Note that a categorization of all possible events is *not* needed for handling them; our algorithm automatically computes a unique local offset structure.

In the generic case, all vertices of the offsetting polytope are of degree 3 before each event (except initially), and faces of dimension ≥ 1 touch only in their interiors

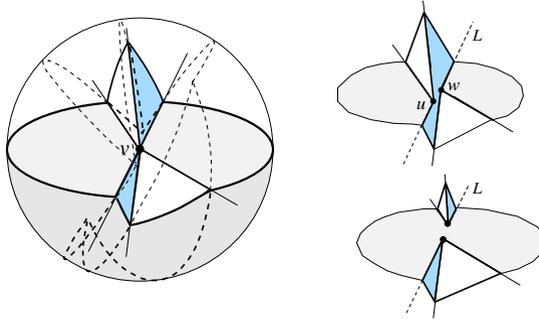


Fig. 7. Vertex touch; degree 6. In $\mathcal{F}(v)$, the two dark-shaded facets and the two horizontal facets are coplanar, respectively (left). Thus the coincidence of the two approaching polytope vertices u and w (right upper) that touch in v is not by degeneracy but by construction; they move on the same straight line L . Vertex v splits into two degree-3 vertices after the event (right lower), as is witnessed by two nodes of the same degree in the skeleton in the lower hemisphere. This skeleton is disconnected, as is the one for the inverse event in the complementary hemisphere.

in an event. Then a surface event is either a *vertex touch* (of one of various types displayed in Figures 4 to 7) which give a vertex figure $\mathcal{F}(v)$ of degree at most 6, a *vertex/edge touch* (like in Figure 8) where $\mathcal{F}(v)$ is of degree 5, or an *edge/edge touch* where $\mathcal{F}(v)$ is of degree 8. Solid events include the *piercing event* (vertex/facet touch, Figure 9), the *kissing event* (edge/edge touch), and the *splitting event*. For each of them, $\mathcal{F}(v)$ has 4 facets. *Tetrahedron collapse* is the last one (and also the chronologically final event), with 4 facets shrunk to a common point.

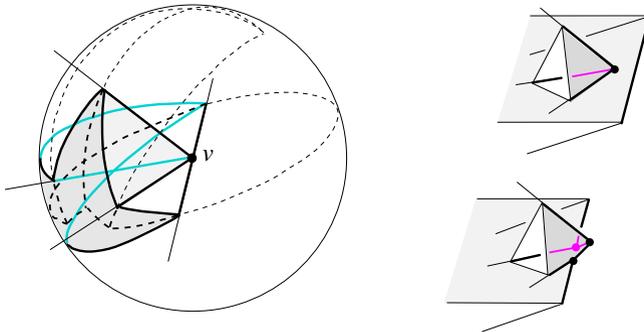


Fig. 8. Vertex/edge touch. The flat wedge retracts to the left, faster than does the triangular pyramid on its top (right upper). At v , the pyramid splits the rim of the wedge, whose lower facet then expands to above the rim in the offset polytope (right lower). Correspondingly, the degree-5 vertex v in $\mathcal{F}(v)$ splits into three vertices, as is witnessed by three nodes in the skeleton (left). This bisector graph for $\mathcal{F}(v)$ is unique. Therefore no other offsets are possible. (The case where the upper (horizontal) wedge facet expands at v leads to a wrong orientation of this plane. The corresponding graph leaves \mathcal{P} .) The situation before the event is encoded in the complementary skeleton, which is disconnected. Its single arc represents the unsplit rim.

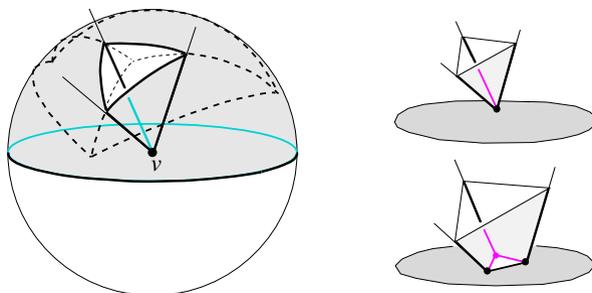


Fig. 9. Piercing event. $\mathcal{F}(v)$ is based on a ‘spherical polygon’ bounded by *two* Jordan curves, a spherical triangle and a full circle (left). The skeleton contains a triangular cycle that corresponds to the hole pierced into the horizontal facet by the offsetting pyramidal pit (right lower). Thus vertex v splits into three vertices.

For the inverse of the piercing event, the skeleton is just a degree-3 star inside the triangular Jordan curve, corresponding to the pyramid whose apex now rises above the horizontal polytope facet (right upper). The other curve, the full circle, contracts at the south pole without leaving a trace – a void event.

The anatomy of events can be quite complex. We therefore provided examples for most of them (but not all, due to space constraints), with detailed explanations in the figures’ legends. The meaning of the events for the skeleton $\text{SKEL}(Q)$ in \mathbb{R}^3 being constructed is briefly commented below.

Each event constructs a particular vertex v of $\text{SKEL}(Q)$. The surface events in Figures 4, 5, and 6 complete a facet of $\text{SKEL}(Q)$ at v , and start a new one. Adjacency of skeleton cells ‘flips’ in the first two cases but, interestingly, remains in the last case. For Figure 7 even the same facet, swept over by line L , continues at v , with two parts. Two new facets, traced out by small pyramid edges, start for Figure 8. Three new facets are created at v in a piercing event (Figure 9). This event yields a tunnel in the shrinking polytope.

4 Concluding Remarks

Concerning the size of $\text{SKEL}(Q)$, each event implies that four offset planes meet at the same point, which happens only once during the entire offsetting process. This bounds the number of events by $\binom{n}{4}$, if Q has n vertices. The size of $\text{SKEL}(Q)$ thus is $O(n^4)$. A lower bound of $\Omega(n^2\alpha(n)^2)$ has been shown in [4].

The construction time depends on the number of occurring events, which tends to stay linear for many inputs, as observed in our implementation. The detection of solid events is costly, though, due to the absence of a three-dimensional analogue of motorcycle graphs [8,12]. The direct method checks self-intersection of the offset polytope after each potential surface event, by intersecting $O(n)$ boundary triangles in $O(n^2)$ time. A theoretical speed-up to $O(n^{4/5+\epsilon}k^{4/5+\epsilon})$ is possible, provided the number k of triangle intersections is small [5].

The cells of $\text{SKEL}(Q)$ are *monotone* in direction normal to their defining polytope facets. This can be shown by generalizing the ‘roof’ argument in [1]

to one dimension higher. The monotonicity of skeleton cells may be useful in applications arising, for instance, in automatic meshing.

When not exclusively using $\text{Sph}(\mathcal{P})$ but also other bisector graphs, offsets with tailor-made features can be generated. We have investigated this issue in a companion paper [3]. In particular, offset polytopes with a minimum or a maximum number of reflex edges are generated. There can be super-exponentially many (in n) different straight skeletons for a given polytope with n vertices. Finding skeletons efficiently that require a minimum number of *constructing events* (or even characterizing them in terms of vertex figure resolution) is of vital interest, as such skeletons have a smallest number of vertices.

The skeleton construction approach in Sections 2 and 3 works in principle in any dimension, though details get involved rapidly. As a possibly still tractable instance, a spherical skeleton in \mathbb{R}^4 can be defined on the basis of $\text{SKEL}(Q)$, leading to a canonical straight skeleton for 4-dimensional polytopes.

Acknowledgements. We thank Günter Rote for interesting discussions.

References

1. Aichholzer, O., Alberts, D., Aurenhammer, F., Gärtner, B.: A novel type of skeleton for polygons. *Journal of Universal Computer Science* 1, 752–761 (1995)
2. Aurenhammer, F.: Weighted skeletons and fixed-share decomposition. *Computational Geometry: Theory and Applications* 40, 93–101 (2007)
3. Aurenhammer, F., Walzl, G.: Three-dimensional straight skeletons from bisector graphs. In: *Proc. 5th International Conference on Analytic Number Theory and Spatial Tessellations*, Kiev, Ukraine (to appear, 2013)
4. Barequet, G., Eppstein, D., Goodrich, M.T., Vaxman, A.: Straight skeletons of three-dimensional polyhedra. In: Halperin, D., Mehlhorn, K. (eds.) *ESA 2008*. LNCS, vol. 5193, pp. 148–160. Springer, Heidelberg (2008)
5. De Berg, M., Guibas, L.J., Halperin, D.: Vertical decompositions for triangles in 3-space. *Discrete & Computational Geometry* 15, 35–61 (1996)
6. Demaine, E.D., Demaine, M.L., Lindy, J.F., Souvaine, D.L.: Hinged dissection of polypolyhedra. In: Dehne, F., López-Ortiz, A., Sack, J.-R. (eds.) *WADS 2005*. LNCS, vol. 3608, pp. 205–217. Springer, Heidelberg (2005)
7. Devadoss, S.L., O’Rourke, J.: *Discrete and Computational Geometry*. Princeton University Press (2011)
8. Eppstein, D., Erickson, J.: Raising roofs, crashing cycles, and playing pool: Applications of a data structure for finding pairwise interactions. *Discrete & Computational Geometry* 22, 569–592 (1999)
9. Grünbaum, B.: *Convex Polytopes*. Interscience, New York (1967)
10. Martínez, J., Vigo, M., Pla-García, N.: Skeleton computation of orthogonal polyhedra. *Computer Graphics Forum* 30, 1573–1582 (2011)
11. Siddiqi, K., Pizer, S.M.: *Medial Representations. Mathematics, Algorithms, and Applications*. Springer Series on Computational Imaging and Vision 37 (2008)
12. Vigneron, A., Yan, L.: A faster algorithm for computing motorcycle graphs. In: *Proc. 29th Ann. ACM Symposium on Computational Geometry*, pp. 17–26 (2013)