The Reflex-Free Hull*

Hee-kap Ahn	Siu-Wing Cheng	Otfried Cheong	Jack Snoeyink
Utrecht University	Hong Kong UST	Utrecht University	UNC Chapel Hill

Abstract

We define a reflex-free hull in three dimensions as an intersection of reflex-free sets. The reflex-free hull allows a rich set of topological types, yet for polyhedral input with n edges, it remains a polyhedral set with $O(n)$ edges. This is in contrast to other possible hull definitions that give non-planar surfaces and higher combinatorial complexity. The reflex-free hull is related to filling cavities, as in the manufacturing process of casting, but we sketch examples to indicate that computing a reflex-free hull will be a challenging problem.

1 Introduction

Computational geometers have identified many classes of 2D polygons (convex, star-shaped, L-convex, externally visible, edge-visible, LR-visible, street, person… [5, 7]), but few classes of 3D polyhedra. Perhaps the fact that 3D polyhedra support rich classes of topological structure in the form of knots and links has overshadowed the identification of geometric structure. In applications such as manufacturing or molecular analysis, however, geometric structures such as mold cavities or docking sites are important.

In the plane, the difference between a simple polygon and its convex hull is a number of simple, polygonal bays, from which one can obtain a natural description of a polygon as a tree of unions and differences of convex pieces [6]. In space, it has been suggested that the same approach be used to define pockets in a search for casting directions [3], but in fact the difference between a polyhedron and its convex hull need not have a natural decomposition, and may have more complicated topology than the original polyhedron [1].

We propose a reflex-free hull, defined as an intersection of reflex-free sets. Section 2 defines the reflex-free sets and plane cavities, which may be seen as analogues to bays in polygons. Section 3 defines the reflex-free hull, *Rfh*, and Section 4 establishes some basic results about the *Rfh* of polyhedral sets, including the fact that the *Rfh* has linear complexity even though it allows a rich set of topological types. Section 5 shows that the reflex-free hull bounds the limit of a process of filling cavities, but that computing it in this manner would be challenging. Finally, Section 6 relates the reflex-free hull to other possible hull definitions that either have high complexities or limited topologies.

2 Definitions

We begin with basic geometric and topological definitions and notation [2, 4] for the sets that we consider in three-dimensional space, $\mathbb{R}^3$. We classify boundary points geometrically, and define reflex-free sets.

*We gratefully acknowledge support from various sources, including RGC CERG HKUST 6074/97E (Cheng) and NSF grants 9988742 and 0076984 (Snoeyink).

A *k-simplex* is a convex hull of $k+1$ affinely independent points. In $\mathbb{R}^3$, we have points, line segments, triangles, and tetrahedra as the 0-, 1-, 2-, and 3-simplices, respectively. The empty set is consider a (-1)-simplex. Notice that the boundary of a simplex is a collection of lower dimensional simplices. A *simplicial complex* is a collection of simplices with disjoint interiors that is closed under the operations of intersection and taking boundaries.

For our purposes, a *polyhedron* is the union of the simplices in a finite simplical complex. A *polyhedral set* is homeomorphic to a polyhedron. We restrict our discussion to polyhedral sets to avoid wild topological beasties like the Alexander horned sphere [4]. Section 4 further restricts the discussion to polyhedra when it investigates combinatorial properties of reflex-free hulls.

A set is *closed* if and only if it contains all of its limit points; the *closure* of a set, $\mathrm{cl}(S)$, is the union of S with its set of limit points. The *complement* of a set, $\overline{S}$, is $\mathbb{R}^3 \setminus S$. For any vector $v \in \mathbb{R}^3$, we define the *v-plane* $h_v(p) = \{q \mid (q-p)\cdot v = 0\}$, and the *closed v-halfspace* $h_v^-(p) = \{q \mid (q-p)\cdot v \le 0\}$. We may suppress subscripts or arguments and write h and h^- when they can be understood from context.

We use the Euclidean metric in $\mathbb{R}^3$, and denote the distance between two points by $d(p,q)$ and between two sets by $d(A,B) = \limsup\{d(a,b) \mid a \in A, b \in B\}$. For $\varepsilon > 0$, the *open* ε*-ball* is $B_\varepsilon(p) = \{q \mid d(p,q) < \varepsilon\}$. The *interior* of a set $\mathcal{Q}$, denoted $\mathrm{int}(\mathcal{Q})$, consists of the points $p \in \mathcal{Q}$ for which there exists an $\varepsilon > 0$ such that $B_\varepsilon(p) \subset \mathcal{Q}$. The *boundary* is defined $\mathrm{bd}(\mathcal{Q}) = \mathcal{Q} \setminus \mathrm{int}(\mathcal{Q})$.

If every boundary point has a neighborhood that is homeomorphic to a half-ball, then $\mathcal{Q}$ is called a *three-manifold with boundary*. We cannot restrict ourselves to manifolds, since non-manfold sets will be able to arise as reflex-free hulls in degenerate configurations. Examples of polyhedral sets that are not three-manifolds include any pair of tetrahedra joined at a vertex or along an edge, and any finite union of one and two dimensional simplices. Points that do not have ball or half-ball neighborhoods are called *singular*.

We classify each boundary point of a polyhedral set, $p \in \mathcal{Q}$, based intuitively on whether $\mathcal{Q}$ or $\overline{\mathcal{Q}}$ can be oriented to hold water at p. Non-manifold sets complicate these definitions. For example, a singular point p appears on the boundary more than once when $\overline{\mathcal{Q}}$ is not connected in the neighborhood of p. We call a connected component $C \subset B_\varepsilon(p) \cap \overline{\mathcal{Q}}$ an *appearance of p on* $\mathcal{Q}$ if $p \in \mathrm{cl}(C)$, and define the *neighborhood of an appearance* as $B_\varepsilon(p) \setminus C$.

For any $\varepsilon > 0$ and vector $v \in \mathbb{R}^3$, we define the *hemisphere* $H_{v,\varepsilon}(p) = B_\varepsilon(p) \cap h_v^-(p) \cap \overline{\{p\}}$. To simplify classification, $H_{v,\varepsilon}(p)$ does not contain p or the boundary points of $\overline{B_\varepsilon(p)}$. Again, we suppress subscripts or argument when they can be understood from context.

We classify an appearance of point $p \in \mathrm{bd}(\mathcal{Q})$ based on the relation of a hemisphere to the neighborhood of the appearance, which we denote $\mathcal{N}$. We say that p appears as a

- *reflex point* if there is a hemisphere at p inside the neighborhood $\mathcal{N}$. That is, if there exists a vector v and $\varepsilon > 0$ such that $H_{v,\varepsilon}(p) \subset \mathrm{int}(\mathcal{N})$.
- *convex point* if there is a hemisphere outside $\mathcal{N}$. That is, if there exists $H_{v,\varepsilon}(p) \subset \overline{\mathcal{N}}$.
- *flat point* if there exists an $\varepsilon > 0$ and $v \in \mathbb{R}^3$ such that $H_{v,\varepsilon}(p) \subset \mathcal{N}$ and $H_{-v,\varepsilon}(p) \subset \mathrm{cl}(\overline{\mathcal{N}})$.
- *nearly reflex point* if p is neither reflex nor flat and there exists $H_{v,\varepsilon}(p) \subset \mathcal{N}$.
- *nearly convex point* if p is neither convex nor flat and there exists $H_{v,\varepsilon}(p) \subset \mathrm{cl}(\overline{\mathcal{N}})$.
- *saddle point* otherwise. That is, for every $\varepsilon > 0$ and vector v, hemisphere $H_{v,\varepsilon}(p)$ intersects both the interior and the complement of $\mathcal{N}$.

For an example, we can classify points on the boundary of a three-manifold polyhedron. Points on faces are flat, points on edges are nearly reflex or nearly convex (or flat in the degenerate case of a dihedral angle of 180°), and points at vertices are convex, reflex, or saddle (except in degenerate cases of incident coplanar faces/edges). In a coffee mug, the reflex points are at the bottom of the bowl.

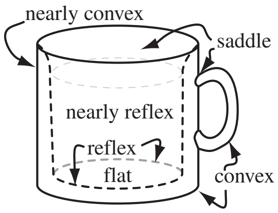

Figure 1: Classifying

Given a closed half-space h^-, we call any bounded connected component of $\overline{Q} \cap h^-$ a *plane-cavity*. A plane-cavity is maximal if no plane parallel to h defines a plane-cavity that contains it. It is not hard to see that you can fill a plane-cavity until it spills over at a saddle, in general, or at convex or nearly convex points in degenerate cases.

Lemma 1 *A polyhedral set Q has a plane-cavity $\overline{Q} \cap h^-$ for some half-space h^- if and only if there is a reflex point in* $\mathrm{bd}(Q)$.

Proof: If Q is not reflex free, then some point $p \in \mathrm{bd}(Q)$ appears as a reflex point, which means that there is a neighborhood $\mathcal{N}$ of this appearance of p and a hemisphere such that $H_{v,\varepsilon}(p) \subset \mathrm{int}(\mathcal{N})$. In the plane $h_v(p)$ that defines the hemisphere, choose a circle γ centered at p with radius less than ε. Because every point of γ is in the interior of $\mathcal{N}$, and $\mathcal{N}$ is compact, there is some $\delta > 0$ so that translating γ to $\gamma + \delta v$ remains strictly inside $\mathcal{N}$. The halfspace $h_v^-(p + \delta v)$ thus cuts off a bounded connected component from $\overline{Q}$.

For the inverse, let X be a plane-cavity of $h^- \cap \overline{Q}$ for some half-space h^-. Take a large sphere S that contains X strictly in its interior. Now, move S until S touches the boundary of X; by making S sufficiently large, the contact between S and X will be a single point $p \in \mathrm{bd}(Q)$ that is not in the plane h. Let v be a vector from p towards the center of the sphere S. We may choose $\varepsilon > 0$ and $\mathcal{N}$ to be the neighborhood of the appearance of p on X, and form a hemisphere $H_{v,\varepsilon}(p) \subset \mathrm{int}(\mathcal{N})$. Thus, p is a reflex point of Q. ◻

We say that Q is a *reflex-free set* iff Q is a polyhedral set that has no reflex points. By the previous lemma, a polyhedral set has no plane-cavities if and only if it is reflex free. The reflex-free sets are closed under intersection, provided they remain polyhedral.

Lemma 2 *Let $\{Q_\alpha\}$ be a family of reflex-free sets whose intersection is polyhedral. The intersection $\bigcap Q_\alpha$ is also reflex-free.*

Proof: Suppose that some point p is in a plane-cavity of $\bigcap Q_\alpha$ defined by half-space h^-. By Lemma 1, it is sufficient to show that for some α, point p is also in a plane-cavity of set Q_α.

Notice that the plane-cavities of the intersection can be written as a union of individual plane cavities:

$$\overline{\left(\bigcap Q_\alpha\right)} \cap h^- = \bigcup (\overline{Q_\alpha} \cap h^-).$$

Since p is in a bounded connected component of the union, point p must be in a connected component of $h^- \cap \overline{Q_\alpha}$, for some α. This component must be bounded, since the component in the union is bounded. ◻

We can use intersections to sculpt reflex-free sets. For example, drilling a hole through a reflex-free set keeps the set reflex free.

Lemma 3 *Let Q be a reflex-free set and X be a polyhedral set. If no reflex point of X is in $\mathrm{int}(Q)$, then the intersection $Q \cap X$ is reflex free.*

Proof: Consider a point p on the boundary of $Q \cap X$: either $p \in \mathrm{bd}(Q)$ or $p \in \mathrm{bd}(X) \cap \mathrm{int}(Q)$. In the first case, we know that no hemisphere $H(p) \subset \mathrm{int}(Q)$, and in the second we know that no hemisphere $H(p) \subset \mathrm{int}(X)$. Thus, there can be no hemisphere in their intersection. ◻

3 The reflex-free hull

Define $Rfh(Q)$, the *reflex-free hull* of a set Q, as the intersection of all reflex-free sets that contain Q. For example, the reflex-free hull of a torus is itself; reflex-free hull of a coffee cup would fill the cup but preserve the handle. The reflex-free hull of a set of discrete points would be these points, because any union of balls around the points is reflex free. The motivation for defining the reflex-free hull is to find a structure that surrounds a set, but fills in depressions or docking sites. We show that the reflex-free hull is reflex-free and idempotent.

Theorem 4 *For a closed set Q, the reflex-free hull $Rfh(Q)$ is reflex free, and satisfies $Rfh(Rfh(Q)) = Rfh(Q)$.*

Proof: Lemma 2 implies that $Rfh(Q)$ is reflex free.

Since $Q \subseteq Rfh(Q)$, the sets whose intersection defines $Rfh(Rfh(Q))$ are a subset of those that define $Rfh(Q)$. Thus, it is clear that $Rfh(Q) \subseteq Rfh(Rfh(Q))$. We prove the reverse inclusion. By the definition of $Rfh(Q)$, if a point p is not in $Rfh(Q)$, then there is a reflex-free set R_p that includes Q and not p. Since R_p participates in the intersection defining $Rfh(Q)$, we also know that $Rfh(Q) \subseteq R_p$. But then R_p also participates in the intersection defining $Rfh(Rfh(Q))$. Thus, p is not in $Rfh(Rfh(Q))$. ◻

We can use sculpting to prove the following technical lemma. Define the ε-*tube* for a line segment s as $C_\varepsilon(s) = \{x \mid d(x,s) < \varepsilon\}$.

Lemma 5 *Suppose that Y is a reflex-free set and that Y contains a point q in the convex hull of a finite set of points outside of Y, namely $\{p_1, p_2, \ldots, p_k\} \subset \overline{Y}$. For some $\varepsilon > 0$, the set $Y \setminus \bigcup_{1 \le i \le k} C_\varepsilon(p_i q)$ is reflex free.*

Proof: We may choose $\varepsilon > 0$ so that the balls $B_\varepsilon(p_i)$ do not intersect Y. The union of ε-tubes, $X = \bigcup_{1 \le i \le k} C_\varepsilon(p_i q)$ is an open set; its complement $\overline{X}$ is a closed set that has reflex points only on ball boundaries $\mathrm{bd}(B_\varepsilon(p_i))$. By Lemma 3, the intersection $Y \cap \overline{X}$ is reflex free. ◻

We use Lemma 5 to show that the reflex-free hull $Rfh(Q)$ inherits its convex, nearly convex, and saddle points from the underlying set Q.

Lemma 6 *For a closed set Q, every convex, nearly convex, or saddle point p of $Rfh(Q)$ is a point of Q. In particular, convex points of $Rfh(Q)$ are convex points of Q, nearly convex points of $Rfh(Q)$ are convex or nearly convex points of Q, and saddle points of $Rfh(Q)$ are convex, nearly convex or saddle points of Q.*

Proof: Because $\mathcal{Q} \subset Rfh(\mathcal{Q})$, the points inside $\mathcal{Q}$ are clearly interior points of $Rfh(\mathcal{Q})$.

Consider a point $p \in Rfh(\mathcal{Q})$ that is outside of $\mathcal{Q}$. We may choose $\varepsilon > 0$ such that the ball $B_\varepsilon(p)$ does not intersect $\mathcal{Q}$. The difference $Y = B_\varepsilon(p) \setminus \mathcal{Q}$ must be convex, since otherwise we could find two, three, or four points in Y whose convex hull contains a point $q \in Rfh(\mathcal{Q})$ and then apply Lemma 5 to obtain a smaller reflex-free set that still contains $\mathcal{Q}$. It is readily checked that for a convex, nearly convex, or saddle point p, the difference Y is not a convex set.

Therefore, a convex, nearly convex, or saddle point $p \in Rfh(\mathcal{Q})$ must come from $\mathrm{bd}(\mathcal{Q})$. By checking the hemispheres of p with respect to $Rfh(\mathcal{Q})$ and $\mathcal{Q}$, we observe that a convex point of $Rfh(\mathcal{Q})$ is a convex point of $\mathcal{Q}$, a nearly convex point of $Rfh(\mathcal{Q})$ is a convex or nearly convex point of $\mathcal{Q}$, and a saddle point p of $Rfh(\mathcal{Q})$ is a convex, nearly convex, or saddle point of $\mathcal{Q}$. ◻

4 The reflex-free hull of a polyhedron

In this section, we consider the reflex-free hull of a polyhedron. We define the *size* of a polyhedron to be the number of vertices, edges, and faces on its boundary. We show that the reflex-free hull of a polyhedron is a polyhedron of the same asymptotic size. We find this surprising in light of the high complexity of other definitions of hulls that we sketch in the next section.

Theorem 7 *The reflex-free hull of a polyhedron of size n is a polyhedron whose size is $O(n)$.*

We use $\mathcal{P}$ as our polyhedron and establish a sequence of lemmas before formally proving this theorem. We first wish to show that the boundary of $Rfh(\mathcal{P})$ consists of flat faces, straight segment edges, and point vertices. Since the convex, nearly convex, and saddle points of $Rfh(\mathcal{P})$ come from $\mathcal{P}$ by Lemma 6, our main task is to show in Lemma 8 that the nearly reflex points form line segments. Next we observe in Lemma 9 that each convex edge of $\mathcal{P}$ contributes at most two vertices to $Rfh(\mathcal{P})$. This allows us to establish that $Rfh(\mathcal{P})$ is a polyhedron in Lemma 10. Finally, we establish the theorem by relating the size to genus, and bounding the genus of $Rfh(\mathcal{P})$.

Lemma 8 *For a polyhedral set $\mathcal{P}$, let R be the set of nearly reflex points of $Rfh(\mathcal{P})$ that do not lie at vertices or on edges of $\mathcal{P}$. The connected subsets of R are line segments.*

Proof: Let $p \in R$ be a nearly reflex point of $Rfh(\mathcal{P})$ that is not a vertex of $\mathcal{P}$. Choose $\varepsilon > 0$ sufficiently small that the only facets of $\mathcal{P}$ that intersect $B_\varepsilon(p)$ are those incident on p. Let h^- be the halfspace that contains the hemisphere that shows that (this appearance of) p is nearly reflex.

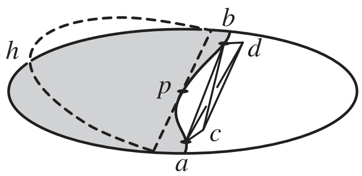

Figure 2: Sculpt with tetrahedron $abcd$ or rotate h.

Figure 2 illustrates the disk $D = h \cap B_\varepsilon(p)$ that is contained in $Rfh(\mathcal{P})$, drawn with shading where D intersects the interior $\mathrm{int}(Rfh(\mathcal{P}))$. Notice that the points of $R \cap D$ serve as the boundary between shaded and unshaded, that is, between $D \cap \mathrm{int}(Rfh(\mathcal{P}))$ and $D \cap \mathrm{bd}(Rfh(\mathcal{P}))$.

Choose a and b as points of $R \cap D$, as in Figure 2. We observe that the segment ab cannot intersect the interior $\mathrm{int}(Rfh(\mathcal{P}))$: Construct an open tetrahedron τ as the interior of the convex hull of a, b, and two other points, $c, d \in B_\varepsilon(p) \setminus Rfh(\mathcal{P})$ so that a, b, c, and d are not coplanar. The reflex vertices of $\overline{\tau}$ are a, b, c, and d, which are not in the interior of $Rfh(\mathcal{P})$. The open tetrahedron τ does not intersect $\mathcal{P}$, since c and d lie outside and a and b lie on or outside the facets of $\mathcal{P}$ in $B_\varepsilon(p)$. Thus, $Rfh(\mathcal{P}) \cap \overline{\tau}$ is reflex free by Lemma 3. But this sculpting operation cannot remove points from $Rfh(\mathcal{P})$, so the segment ab does not intersect the interior $\mathrm{int}(Rfh(\mathcal{P}))$.

Now, consider the convex hull of the nearly reflex points R within disk D. This hull cannot contain a point of the interior, $D \cap \mathrm{int}(Rfh(\mathcal{P}))$, by the previous paragraph. Since p is not flat, there are interior points in any neighborhood of p, so p must be on the boundary of the hull. If we rotate the plane h around a tangent to the hull at p and define a new hemisphere $H(p)$ as indicated by the dashed line in Figure 2, then we can see that p must lie on a line segment on this hull, or we would observe a reflex point at p. This concludes the proof of the lemma. ◻

We next observe that an edge from $\mathcal{P}$ appears at most once as a nearly convex edge on $Rfh(\mathcal{P})$.

Lemma 9 *If p and q are nearly convex points of $Rfh(\mathcal{P})$ that come from the same edge e of $\mathcal{P}$, then the segment pq consists entirely of nearly convex points of $Rfh(\mathcal{P})$.*

Proof: Let p and q be points satisfying the hypothesis of the lemma. Choose $\varepsilon > 0$ sufficiently small that the ε-tube $\mathcal{C}_\varepsilon$(pq) does not intersect any facets of $\mathcal{P}$ except those incident on segment pq. At the left side of Figure 3, we draw two ε disks around p and q, and shade their intersections with $\mathcal{P}$ (dark) and $Rfh(\mathcal{P})$ (light).

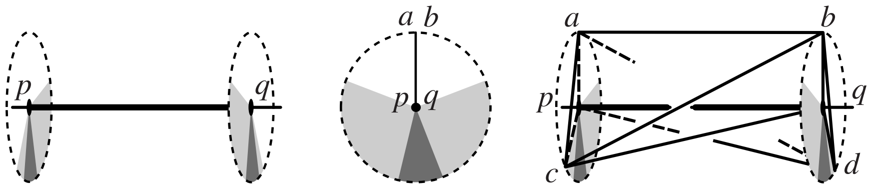

Figure 3: Sculpting $Rfh(\mathcal{P})$ between two nearly convex points p and q of an edge e of $\mathcal{P}$.

We may choose points a and b outside of $Rfh(\mathcal{P})$ so that the segment ab is parallel to pq: simply choose a and b in the antipodes of the regions intersecting $\mathcal{P}$ in the two disks, as in the middle of Figure 3. We may then choose c and d outside of $Rfh(\mathcal{P})$ so that $\angle apc$ is obtuse in the disk containing p and $\angle bqd$ is obtuse in the disk containing q, and c and d lie on opposite sides of the plane through $abpq$.

Now, form the sculpting region X as the closure of the complement of the two tetrahedra $acpq$ and $bdpq$, as at the right of Figure 3. The reflex points of X, which are a, b, c, and d, are all outside of $Rfh(\mathcal{P})$, so by Lemma 3 we know that $Rfh(\mathcal{P}) \subset X$. Since all points of pq are on $e \subset Rfh(\mathcal{P})$, we can find hemispheres to show that they are all nearly convex points of $Rfh(\mathcal{P})$. ◻

We summarize what we know about $Rfh(\mathcal{P})$ thus far.

Lemma 10 *The reflex-free hull of a polyhedron of size n is bounded by a polyhedron with $O(n)$ vertices.*

Proof: Let $\mathcal{P}$ be a polyhedron of size n, and classify the points of $\mathrm{bd}(Rfh(\mathcal{P}))$. By Lemmas 6 and 8, the nearly convex and nearly reflex points may be organized into line segments. These line segments must end at vertices that are convex or saddle points, since there are no reflex points. But each such vertex lies on a nearly convex edge of $\mathcal{P}$, and each edge can contribute at most two vertices by Lemma 9. The remaining points are flat points, which may therefore be grouped into polygons. A finite number of polygons may be formed on n vertices, so $Rfh(\mathcal{P})$ is a polyhedron. $\square$

We can now complete the proof of Theorem 7, and show that, for a polyhedron $\mathcal{P}$ the reflex-free hull $Rfh(\mathcal{P})$ is a polyhedron of the same asymptotic size.

Proof: From Lemma 10 we know that $Rfh(\mathcal{P})$ is a polyhedron with $O(n)$ vertices, where n is the size of $\mathcal{P}$. We have only to bound the numbers of faces and edges.

From Euler-Poincare formula, we know that $V-E+F=2-2g$, where V, E, F, and g are the number of vertices, edges, faces, and genus of $Rfh(\mathcal{P})$. Since $3F \leq 2E$, we deduce that $V-E/3 \geq 2-2g$ and so $E \leq 3V+6g-6$. Since $V \in O(n)$, we have only to bound the genus $g \in O(n)$ to complete the proof.

Curvature is defined at all points on a surface, and the sum of curvature over $Rfh(\mathcal{P})$ equals $-4\pi(g-1)$. For a polyhedron, points on faces or edges have curvature zero, and the curvature of a vertex v equals 2π minus the sum of the angles of faces incident on v. Thus, $4\pi g$ equals 2π plus the sum of all face angles at all vertices of $Rfh(\mathcal{P})$. Since the genus of $\mathcal{P}$ is less than n by the Euler-Poincare formula, we bound the increase of genus by bounding the increase in the sum of face angles when we go from $\mathcal{P}$ to $Rfh(\mathcal{P})$.

Three types of changes to vertices occur when we go from $\mathcal{P}$ to $Rfh(\mathcal{P})$:

(i) a vertex v of $\mathcal{P}$ may disappear into the interior of $Rfh(\mathcal{P})$,

(ii) a new vertex v may be created on $Rfh(\mathcal{P})$, or

(iii) a vertex v of $\mathcal{P}$ may become incident on new faces.

We can bound how each type of change increases the sum of face angles.

(i) When a vertex v of $\mathcal{P}$ disappears, it no longer contributes to sum of face angles.

(ii) A new vertex v is incident on exactly one convex edge, since v is not reflex and is not a vertex in $\mathcal{P}$. The sum of angles of the two faces incident to the convex edge is less than 2π, and the sum of angles of the remaining faces incident to v is less than 2π. So a new vertex increases the sum of face angles by less than 4π.

(iii) Since $Rfh(\mathcal{P})$ is a polyhedron, we may organize the faces incident to v into one or more topological disks. At most one of these disks may be flat, contributing an angle of 2π. Any other must have at least one convex edge.

Lemma 6 implies that no new convex edge can be created incident to v. Thus, this change replaces the faces between two convex edges (possibly identical) by a new set of faces that are joined by reflex edges. This decreases the sum of face angles, except possibly where a new face angle of greater than π is created incident to v. This new face adds less than 2π, but also consumes one quarter of the neighborhood of v. The increase in face angles at v will be less than 8π.

Hence, the maximum increase in face angles is $8\pi n$, and the increase in genus of $Rfh(\mathcal{P})$ is thus $O(n)$. This completes the proof of Theorem 7. $\square$

5 The reflex-free hull and cavities

Recall that for a closed polyhedral set $\mathcal{Q}$ and halfspace h^-, we defined a *plane-cavity* as a connected component of $\overline{\mathcal{Q}} \cap h^-$. We can enlarge a plane-cavity by translating its plane h unless h contains a saddle point or nearly convex points. We say that a plane-cavity is *limited* if its plane contains three saddle points or a closed curve of convex and nearly convex points.

Lemma 11 *Any plane-cavity is contained in the union of four limited plane-cavities.*

Proof: Consider a plane-cavity $C \subset \overline{\mathcal{Q}} \cap h^-$. We may translate h to enlarge C unless doing so would cause C to be connected to the unbounded component of $\overline{\mathcal{Q}}$. This may happen if h contains a closed chain of (nearly) convex points satisfying the lemma.

Otherwise h contains one or more saddle points. If h contains two saddle points, a and b, then we may duplicate h and rotate the two copies in opposite directions around the line ab. We stop each rotation when a third saddle point or chain of nearly convex points is reached. If there is a single saddle, we again duplicate and rotate h around some line through the saddle until we hit a second saddle and reduce to two instances of the previous case. $\square$

If we iteratively fill up plane cavities for a polyhedron $\mathcal{P}$, then we obtain a sequence of interesting sets. We describe this process precisely as follows. Let $\mathcal{P}_0 = \mathcal{P}$. Given some plane-cavity C_k of $\mathcal{P}_k$, we form the union $\mathcal{P}_k \cup C_k$ to obtain a new polyhedron $\mathcal{P}_{k+1}$. We may choose our plane-cavities by always choosing the one with largest volume, or by always choosing four limited plane-cavities that enclose the largest volume.

We call a connected component of $\mathcal{P}_k \setminus \mathcal{P}$ a *cavity*. We believe, but have not been able to formally prove, that in the limit we can obtain the reflex-free hull by filling cavities. Equivalently, the cavities of a closed polyhedron $\mathcal{P}$ are the connected components of $Rfh(\mathcal{P}) \setminus \mathcal{P}$.

Theorem 12 *For a closed polyhedron $\mathcal{P}$, the limit of the process of filling cavities is a subset of the reflex-free hull, $Rfh(\mathcal{P})$.*

Proof: We show by induction that the cavities identified by the filling process are inside all reflex-free sets that contain $\mathcal{P}$. Specifically, we prove that any polyhedral set $\mathcal{Q}$ that contains $\mathcal{P}$, but does not contain $\mathcal{P}_k$, has a reflex point.

The base case, $k = 0$, is trivial; no set containing $\mathcal{P}$ can omit a point of $\mathcal{P}_0 = \mathcal{P}$.

For the inductive step, we assume the induction hypothesis for some $k \geq 0$, and prove it for $k+1$. Thus, assume that $\mathcal{Q}$ is a polyhedral set that contains $\mathcal{P}$ but does not contain $\mathcal{P}_{k+1}$. If $\mathcal{Q}$ does not contain $\mathcal{P}_k$, then the induction hypothesis applies immediately, so we assume that $\mathcal{Q}$ contains $\mathcal{P}_k$. The boundary of $\mathcal{Q}$ therefore intersects the plane-cavity $C_k = \mathcal{P}_{k+1} \setminus \mathcal{P}_k$. But since $\mathrm{bd}(\mathcal{Q})$ can only escape through the plane defining C_k, if at all, $\mathcal{Q}$ has a plane cavity in C_k. Lemma 1 says that $\mathcal{Q}$ has a reflex point. $\square$

Unfortunately, we do not know how to turn this definition into an efficient procedure to compute the reflex-free hull of a polyhedron. The process of filling in one reflex vertex can create others at reflex edges. Figure 4 illustrates one example in which filling cavities must be taken to the limit to attain the reflex-free hull.

Start with the cube $[-5,5]^3$ and subtract the following sets: $\{(x,y,z) \mid x \in [-1,1], z > |y|\}$, $\{(x,y,z) \mid |x| \in [1,3], z > |y| - |x| + 1\}$, $\{(x,y,z) \mid x \in [3,5], z > y/2+1\}$, and $\{(x,y,z) \mid x \in [-3,-5], z > -y/2+1\}$, to obtain an object illustrated in Figure 4. There are four labeled lines that are relevant in this example. We parameterize them by z. Two lines are *pivot lines*, $\alpha(z) = (-3, 2z-2, z)$ and $\delta(z) = (3, 2-2z, z)$, and two lines end at saddle points, $\beta(z) = (-1, z, z)$, and $\gamma(z) = (1, -z, z)$. With a little algebra, we observe that a plane that contains the entire line δ and point $\beta(t)$, for any chosen value of $t \geq 0$, must intersect the line $\gamma(z)$ at $z = (t+2)/6$. By symmetry, the plane containing the line α and point $\gamma(t)$ intersects the line $\beta(z)$ at $z = (t+2)/6$.

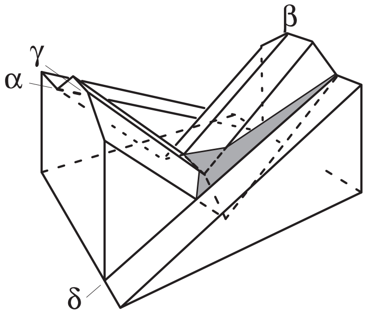

Figure 4: Each filling step creates reflex and saddle points

Initially, there are two reflex vertices with coordinates $(\pm 3, 0, -3)$. We can eliminate the first by filling the cavity defined by the plane through pivot line δ and the saddle point $\beta(0)$; this plane intersects γ at $z_1 = 2/6$. We eliminate the second by filling to the plane through α and the newly created saddle point $\gamma(z_1)$; this plane intersects β at $z_2 = 7/18$, and creates a new reflex vertex where the two filling planes meet. From now on, we fill from a pivot line to a saddle at $z_i = (z_{i-1}+2)/6$. The reflex-free hull for this example has a reflex edge along the line through $(-1, 2/5, 2/5)$ and $(1, -2/5, 2/5)$, which happens to be the unique line incident to α, β, γ, and δ. Thus, we approach, but never reach the reflex-free hull.

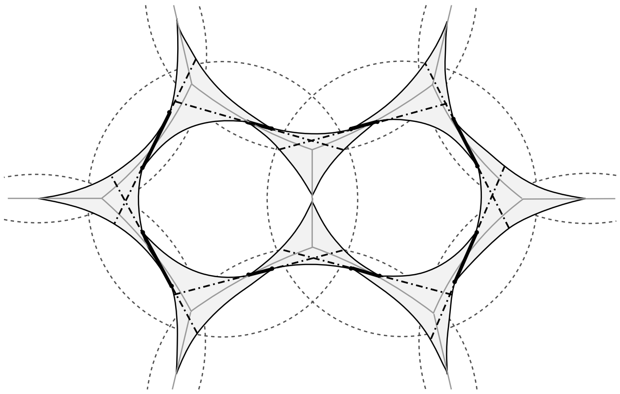

Figure 5: Part of the reflex hull of an appropriately placed set of eight spheres. The planes determining the new boundaries of this hull are defined by a sequence of saddles in such a way that if any saddle is moved then all plane equations change.

If reflex edges incident on four polyhedron edges were the worst that could occur, we could still hope for a polynomial-time algorithm for the reflex-free hull by inspecting all 4-tuples of edges to see if they support a common line. Unfortunately, however, reflex edges may be defined by lines that hit only two polyhedron edges. Figure 5 shows such an example made of eight spheres, which could be approximated by polyhedra. In it, we have a sequence of eight geodesic triangles that share the thick segments, which are reflex edges of the hull. The extensions of reflex edges (dash-dotted) intersect within the incident geodesic triangles. The reader who would like to find an algorithm to compute reflex-free hulls is advised to build similar examples from modeling clay.

6 Other hulls

The fact that the reflex-free hull has linear complexity may not at first seem surprising. In this section, we consider two other natural definitions for hulls that have far worse complexities.

For a closed set S, we may obtain the convex hull $CH(S)$ by removing halfspaces do not intersect S or as the intersection of half-spaces that contain S. We may obtain the reflex-free hull $Rfh(S)$ by sculpting according to Lemma **??**, or as the intersection of reflex-free sets that contain S. In a similar way, we may define the *line hull*, $LH(S)$, by removing lines that do not intersect S, or, more formally, as the intersection of sets containing S that are the complements of lines.

We can order the hulls of a set S by inclusion: the convex hull contains the line hull, which contains the reflex-free hull.

Lemma 13 *For a polyhedral set S, we have $S \subseteq Rfh(S) \subseteq LH(S) \subseteq CH(S)$. In general, all inclusions are strict.*

Proof: Follows immediately from the definitions, once we notice that the complement of a line is reflex-free, and that a halfplane can be represented as the intersection of line-complements. $\square$

In the plane, the line hull of a connected set is the same as its convex hull, but the line hull of a disconnected polyhedral set of size n may can have $\Theta(n^4)$ complexity, as it is related to an arrangement of the $\Theta(n^2)$ lines tangent to pairs of vertices of S. In $\mathbb{R}^3$ the line hull is bounded by pieces of ruled surfaces, including hyperboloids. Sergei Bespamyatnikh, in private communication, described an example of a connected polyhedral set S of size n whose line hull has $\Theta(n^9)$ complexity. Begin with the six faces of a large, axis-aligned cube, and cut a small square hole in the center of each. Near the center of this cube we have three families of lines, each roughly parallel to one of the three axes. Block the lines parallel to the x axis with three squares parallel to the yz plane, and cut n parallel slits in different directions each square so that the lines that do pass through these slits form $\Omega(n^3)$ hyperboloids near the center of the cube. Repeat this for the y and z axes, so that these hyperboloids have $\Omega(n^9)$ intersections.

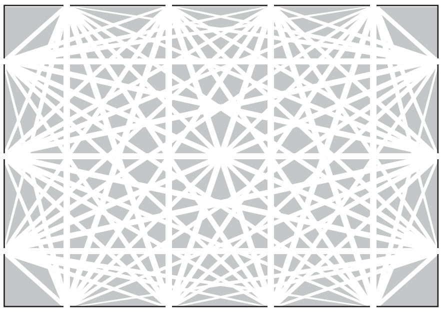

Figure 6: Line hull of 18 black segments

Some other natural hull definitions suffer from similar complexities. One could define the *star hull*, $SH(S)$, as the intersection of all star-shaped polyhedra that contain S. Each point p that is not in the star hull is excluded from some star-shaped polyhedron, which says that p has a ray to infinity that does not intersect the interior of S. Thus, one could define the *ray hull* as the intersection of sets containing S that are the complements of open rays. We say that a point p of a closed set X is *externally visible* if there is a ray from p that does not intersect the interior of X. A set X is *externally visible* if every point on its boundary is externally visible. Thus, one could define the *externally visible hull* of a closed set S to be the intersection of all externally visible sets that contain S.

It is not difficult to see that the star hull, ray hull, and externally visible hull are identical, and that $S \subseteq SH(S) \subseteq LH(S)$, with strict inclusion for many sets S. The reflex free hull and star hull cannot always be ordered by inclusion: The boundary of the star hull contains any reflex vertices from S that are externally visible. The boundary of the reflex free hull may contain points that are not externally visible, as can be seen in an example of nested tori rotated about a common axis so they form a spherical shell.

A minor modification of Bespamyatnikh's construction shows that the star hull can again be bounded by hyperboloids, and may have $\Omega(n^9)$ complexity.

Acknowledgments

We thank Herbert Edelsbrunner and Sergei Bespamyatnikh for discussions on the relation of this problem to casting and other hulls.

References

[1] H.-K. Ahn, M. de Berg, P. Bose, S.-W. Cheng, D. Halperin, J. Matoušek, and O. Schwarzkopf. Separating an object from its cast. In *Proc. 13th Annu. ACM Sympos. Comput. Geom.*, pages 221–230, 1997.

[2] M. A. Armstrong. *Basic Topology*. McGraw-Hill, London, UK, 1979.

[3] L. Chen, S. Chou, and T. Woo. Parting directions for mould and die design. *Comput. Aided Design*, 25:762–768, 1993.

[4] J. G. Hocking and G. S. Young. *Topology*. Addison-Wesley, Reading, MA, 1961.

[5] M. Soss and G. Toussaint. A hierarchy of polygons: a study of polygon properties. presentation at Canada-Cuba Workshop on Algorithms, May 2000.

[6] S. B. Tor and A. E. Middleditch. Convex decomposition of simple polygons. *ACM Trans. Graph.*, 3:244–265, 1984.

[7] Marc van Kreveld. Person polygons: Why not? In D. Avis and J. Bose, editors, *Snapshots of Computational and Discrete Geometry, Vol. 3, Tech. Rep. SOCS–94.50*. McGill School of Comp. Sci., 1995. `http://www.cs.uu.nl/people/marc/personpoly.ps.gz`.