

Casting a Polyhedron with Directional Uncertainty*

Hee-kap Ahn**

Otfried Cheong**

René van Oostrum**

Abstract

Casting is a manufacturing process in which molten material is poured into a cast (mould), which is opened after the material has solidified. As in all applications of robotics, we have to deal with imperfect control of the casting machinery. In this paper, we consider directional uncertainty: given a 3-dimensional polyhedral object, is there a polyhedral cast such that its two parts can be removed in opposite directions *with uncertainty* α without inflicting damage to the object or the cast parts? We give a necessary and sufficient condition for castability, and an algorithm that verifies castability and produces two polyhedral cast parts for a polyhedral object of arbitrary genus. Its running time is $O(n \log n)$. The resulting cast parts have $O(n)$ vertices in total. We also consider the case where the removal direction is not specified in advance, and give an algorithm that finds all feasible removal directions with uncertainty α in time $O(n^2 \log n / \alpha^2)$.

1 Introduction

Casting is a manufacturing process in which molten material is poured into a cavity inside a *mould* (cast). After the liquid material has hardened, the mould is opened, and we are left with an object [5, 11], which has the shape of the cavity.

An industrial CAD/CAM system can aid a part designer in verifying already *during the design* of an object whether the object in question can actually be manufactured using a casting process. At the basis of this verification is a geometric decision: is it possible to enclose the object in a mould that can be split into two parts, such that these two *cast parts* can be removed from the object without colliding with the object or each other. (We are not interested in casting processes where the mould has to be destroyed to remove the object.) Note that this is a preliminary decision meant to aid in part design—to physically create the mould for a part one needs to take into account other factors such as heat flow and how air can evade from the cavity.

This problem has been studied by Bose, Bremner, and van Kreveld [3], who considered the *sand casting model* relevant in iron casting, where the two cast parts have to be separated by a plane. Ahn et al. [2] gave, to our knowledge, the first complete algorithm to determine the castability of polyhedral parts for cast removal as we described above, under the assumption that the two cast parts have to be removed in opposite directions. This restriction is true for current casting machinery, and we will therefore assume it in this paper as well. Nevertheless, Ahn, Cheng, and Cheong [1] considered the castability of polyhedral parts in a relaxed model that may become relevant in the future.

The casting algorithms mentioned above assume perfect control of the casting machinery. When a cast part is removed, it is required that the part moves exactly in the specified direction. In practice, however, this will rarely be the case. As in all applications of robotics, we have to deal with imperfect control of the machinery, and a certain level of uncertainty in its movements. When a facet of the object

*This research was partially supported by the Hong Kong Research Grants Council.

**Institute of Information & Computing Sciences, Utrecht University, P.O.Box 80.089, 3508 TB Utrecht, The Netherlands.
Email: {heekap,otfried,rene}@cs.uu.nl

or of a cast part is almost parallel to the direction in which the cast parts are being moved, the two touching surfaces may damage each other when the mould is being opened. This can make the resulting object worthless, or it may wear away the surface of the mould so that it cannot be reused as often as desirable.

In Figure 1 (a), the mould can be opened by moving the two parts in direction \vec{d} and $-\vec{d}$. If, however, due to imperfect control, the upper part is translated in direction \vec{d}' , it will destroy the object. The cast parts in (b) are redesigned so that both cast parts can be translated without damage in the presence of some uncertainty.




Figure 1: (a) Insufficient angle: the upper part of the cast is stuck, (b) A new removal direction

In this paper, we consider directional uncertainty in the casting process: given a 3-dimensional polyhedral object, is there a polyhedral cast such that its two parts can be removed in opposite directions *with uncertainty* α without damage to the object or the cast parts? We call such an object *castable with uncertainty* α .

Directional uncertainty has been considered by researchers in motion planning, and robotics in general. A motion planning model with directional uncertainty was perhaps first proposed by Lozano-Pérez, Mason and Taylor [9]. An extensive treatment of motion planning with directional uncertainty is given in the book by Latombe [8].

We generalize the characterization of castable polyhedra by Ahn et al. [2] to incorporate uncertainty in the directions in which the cast parts are removed. A formal definition of our model is given in Section 2. It turns out that one of the main difficulties is to guarantee that the two cast parts are *polyhedral*—while this is trivial in the exact case, it requires approximation of a curved surface in our model with uncertainty. We give an algorithm that verifies whether a polyhedral object of arbitrary genus is castable for a given direction of cast part removal and given uncertainty $\alpha > 0$. The running time of the algorithm is $O(n \log n)$, where n is the number of vertices of the input polyhedron. If the object is castable, the algorithm also computes two polyhedral cast parts with $O(n)$ vertices in total.

We then consider the case where the direction of cast part removal is not specified in advance. We give an algorithm that finds all possible removal directions in which the polyhedral object is castable with uncertainty $\alpha > 0$ in time $O(n^2 \log n / \alpha^2)$.

2 Preliminaries

Throughout this paper, P denotes a polyhedron, that is, a not necessarily convex solid bounded by a piecewise linear surface. The union of vertices, edges, and facets on this surface forms the boundary of P , which we denote by $\text{bd}(P)$. We require $\text{bd}(P)$ to be a connected 2-manifold. Each facet of P is a

connected planar polygon, which is allowed to have polygonal holes. Two facets of P are called *adjacent* if they share an edge. We also assume that P is *simple*, which means that no two non-adjacent facets share a point. The polyhedron P may contain tunnels, and can indeed have arbitrary genus.

A polyhedron P is *monotone* in direction \vec{d} if every line with direction \vec{d} intersects the interior of P in at most one connected component. We say that P is α -monotone in direction \vec{d} for an angle α with $0 \leq \alpha < \pi/2$ if P is monotone in direction \vec{d}' for all directions \vec{d}' with $\angle(\vec{d}, \vec{d}') \leq \alpha$,

We say that a facet f of a polyhedron P is α -*steep* in direction \vec{d} if the angle β between a normal of f and \vec{d} lies in the range $\pi/2 - \alpha \leq \beta \leq \pi/2 + \alpha$. A polyhedron P is called α -*safe* in direction \vec{d} if none of its facets is α -steep for that direction.

A *terrain* is the graph of a (possibly partially defined) continuous, piecewise differentiable function with domain \mathbb{R}^2 and range \mathbb{R} . This means that a terrain is a surface with the property that every vertical line intersects it in at most one point. Hence, it is monotone in direction \vec{z} . We call a terrain α -*safe* if the normal vector of the surface makes an angle of at most $\pi/2 - \alpha$ with the vertical direction wherever it is defined. A terrain is *polyhedral* if the surface is piecewise linear.

A mould M with opening direction \vec{d} for a polyhedron P is a pair (C_r, C_b) of two polyhedra C_r and C_b , such that the interiors of C_r , C_b , and P are pairwise disjoint and the union $B := C_r \cup P \cup C_b$ is a rectangular box with an edge parallel to \vec{d} that completely contains P in its interior. We call C_r and C_b the *red cast part* and the *blue cast part* of M .

A mould M with opening direction \vec{d} is α -feasible, if for each pair of directions (\vec{d}_r, \vec{d}_b) with $\angle(\vec{d}, \vec{d}_r) \leq \alpha$ and $\angle(-\vec{d}, \vec{d}_b) \leq \alpha$, the red cast part C_r can be translated to infinity in direction \vec{d}_r without colliding with P or C_b , and the blue cast part C_b can be translated to infinity in direction \vec{d}_b without colliding with P . Note that the order of removing the cast parts is actually irrelevant.

A polyhedron P is α -*castable in direction* \vec{d} if an α -feasible mould with opening direction \vec{d} exists. For the special case $\alpha = 0$, we say that P is *castable in direction* \vec{d} .

The following simple lemma characterizes polyhedra castable in direction \vec{d} [2].

Lemma 1 *A polyhedron P is castable in direction \vec{d} if and only if it is monotone in direction \vec{d} .*

The main result of the present paper is a generalization of this result to α -castability. We state the result here—it will take us a few more pages to prove it.

Theorem 1 *A polyhedron P is α -castable in direction \vec{d} if and only if P is α -monotone and α -safe in direction \vec{d} .*

The following lemma proves the necessity of the condition.

Lemma 2 *If a polyhedron P is α -castable in direction \vec{d} , then P is α -monotone and α -safe in direction \vec{d} .*

Proof: Assume that P is not α -safe, so a facet f is α -steep with respect to \vec{d} . A point p in the interior of f can be neither on the boundary of C_r nor on the boundary of C_b , and so P is not α -castable in direction \vec{d} .

On the other hand, if P is α -castable in direction \vec{d} , it is castable in any direction \vec{d}' with $\angle(\vec{d}, \vec{d}') \leq \alpha$. By Lemma 1, it follows that P is monotone in direction \vec{d}' . It follows that P is α -monotone. \square

3 Finding a mould

It remains to prove the sufficiency of the condition in Theorem 1. We do so by showing how to construct an α -feasible mould for any α -monotone and α -safe polyhedron. To simplify the presentation, we will assume, without loss of generality, that \vec{d} is the upward vertical direction (the positive z -direction). We say that P is α -castable if it is α -castable in the vertical direction.

A facet of P is called an *up-facet* if its outward normal points upwards, and a *down-facet* if its outward normal points downwards. Assuming P is α -safe, there are no vertical facets, and so each facet is either an up-facet or a down-facet. Clearly an up-facet of P must be a facet of the red cast part C_r , while a down-facet of P must be a facet of the blue cast part C_b . The difficulty is finding the separating surface between C_r and C_b “elsewhere.”

Assume that P is α -castable and that (C_r, C_b) is an α -feasible mould for P . Again we denote by B the axis-parallel box that forms the outside of the mould. We define the *blue parting surface* S_b as the common boundary of C_b and $C_r \cup P$, and the *red parting surface* S_r as the common boundary of C_r and $C_b \cup P$. Any upwards directed vertical line ℓ must intersect C_b , P , and C_r in this order, each in a single connected component that can be empty. It follows that both S_b and S_r are polyhedral terrains. The two terrains coincide except where they bound the polyhedron P . If we let $S := S_b \cap S_r$, define S_u to be the union of all up-facets, and S_d to be the union of all down-facets, we can write $S_r = S \cup S_u$ and $S_b = S \cup S_d$. The boundary of S is the set of silhouette edges of P (an edge is a silhouette edge if it separates an up-facet from a down-facet).

Constructing a mould therefore reduces to the construction of the terrain S . For the special case $\alpha = 0$, Ahn et al. [2] gave a triangulation method for constructing S as follows: Let h be a horizontal plane cutting the box B in two roughly equal halves. Let R be the rectangle $h \cap B$. We project P onto h and obtain a polygon \bar{P} , possibly with holes. Let T be a triangulation of $R \setminus \bar{P}$. Every triangle in T is “lifted” into 3-dimensional space by replacing each vertex \bar{v} of \bar{P} by its original vertex v of P . The resulting 3-dimensional surface is the desired terrain S separating the red and blue cast parts. (The description by Ahn et al. [2] is more complicated as it handles vertical facets.)

Unfortunately, this construction does not necessarily produce an α -feasible mould, even when the polyhedron is α -castable. Figure 2 illustrates this possibility. \bar{P} is the projection of a polyhedron P that is α -monotone and α -safe. The z -coordinates of vertices a and b are identical (and so the segment ab is horizontal). The z -coordinate of c is chosen such that both ac and bc make an angle of α with the vertical direction. Any triangulation of $R \setminus \bar{P}$ contains the triangle abc . This implies that the midpoint p of ab lies on S , and therefore on the boundary of the red cast part. However, translating p upwards with uncertainty α may cause it to collide with the polyhedron at c , and so the mould is not α -feasible.

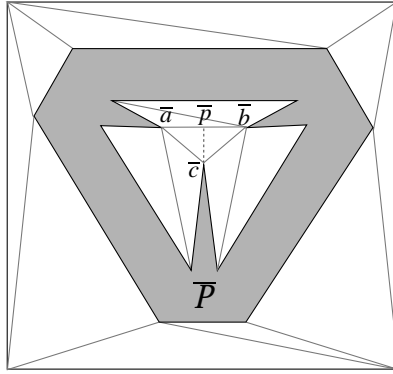


Figure 2: The triangulation method fails: the line segment pc is too steep.

The problem with this approach is that even if the polyhedron is α -monotone and α -safe, the constructed terrain S is not: the triangle abc is in fact α -steep. We now prove that it suffices to make sure this does not happen.

Lemma 3 *Let B be an axis-parallel box, and let S be an α -safe terrain separating the top and bottom facets of B . Let C be the part of B above S , and let $C' := B \setminus C$. Let \vec{d} be the upward vertical direction,*

and let \vec{d}' be such that $\angle(\vec{d}, \vec{d}') \leq \alpha$. Then C can be translated to infinity in direction \vec{d}' without colliding with C' .

Proof: Assume the claim was false, and consider a point $p \in C$ that when translated in direction \vec{d}' collides with a point $q \in C'$. The line segment pq lies completely inside B , and so its vertical projection onto S is a path π . Since p lies above one end-point of π , q lies below the other end-point, and the slope of pq is at least $\pi/2 - \alpha$, there must be a segment on π where the slope is at least $\pi/2 - \alpha$. This is a contradiction to the assumption that S is α -safe. \square

Lemma 4 *Let P be an α -safe polyhedron, B an axis-parallel box enclosing P , and let S be an α -safe polyhedral terrain bounded by the silhouette edges of P . Then the mould defined by the parting surfaces $S_r := S \cup S_u$ and $S_b := S \cup S_d$ is α -feasible.*

Proof: Since P is α -safe, both S_u and S_d are α -safe terrains. Since S is α -safe, both S_r and S_b are therefore α -safe. Lemma 3 now implies that the mould is α -feasible. \square

We will now show how to construct a terrain S as in Lemma 4 by forming the lower envelope of a set of cones. Given a point p on an up-facet of P , the α -cone $D(p)$ of p is the solid vertical upwards oriented cone of angle α with apex p . Formally, if p' is a point vertically above p , then $D(p) := \{x \mid \angle(xpp') \leq \alpha\}$. Let now D_1 be the union of $D(p)$ over all points $p \in S_u$, and let E_1 be the lower envelope of D_1 . Clearly, E_1 contains S_u , and so $S := E_1 \setminus S_u$ is bounded by the silhouette edges of P . Since E_1 consists of patches of α -cones, it is clearly α -safe. It follows that S fulfills the requirements of Lemma 4, except that it is not a polyhedral terrain.

We will see below that we can easily “approximate” S by a polyhedral, α -safe terrain S' that contains all the linear edges of S and lies below (or coincides with) S everywhere. (The reader might also rightfully ask why a mould has to be polyhedral—perhaps a mould bounded by the conic patches resulting from our construction might work better in practice than the polyhedral version we will construct below.)

The construction of S above appears to require taking the union of an infinite family of cones. We now give an alternative definition of S as the lower envelope of h constant-complexity objects, where h is the number of silhouette edges of P .

In fact, let pq be a silhouette edge of P . The α -region $D(pq)$ of pq is the convex hull of $D(p) \cup D(q)$. The lower envelope of $D(pq)$ consists of three components: two conic surfaces supported by the α -cones $D(p)$ and $D(q)$, and a connecting area consisting of two planar facets.

Let now D_2 be the union of $D(pq)$, over all silhouette edges pq , and let $E = E_2$ be the lower envelope of D_2 . It is easy to see that E_1 is in fact the lower envelope of S_u and E_2 , and so E_1 and E_2 coincide “outside” of P . Thus, if we define S to be the part of $E = E_2$ not lying above S_u , we define the same terrain S as above.

The lower envelope E consists of $O(h)$ faces, which are either planar, or supported by a single α -cone $D(x)$ for a vertex x of P . An edge of E is either a silhouette edge of P , a straight edge separating a conic patch supported by an α -cone $D(x)$ from an adjacent planar patch supported by an α -region $D(xy)$, or is an arc supported by the intersection curve of two α -cones, an α -cone and a plane, or two planes. Such arcs are either straight segments, arcs of parabolas, or arcs of hyperbolas. In all cases, they are contained in a plane. Figure 3 shows the two types of conic sections arising.

We can represent E by its projection on the xy -plane. The projection is in fact a planar subdivision, whose faces are supported by a single plane or α -cone. If we annotate each face with the vertex or silhouette edge of P whose α -cone or α -region supports it, the resulting map is a complete representation of E .

In general, the lower envelope of m well-behaved, constant-complexity objects can have complexity $\Theta(m^2)$ [10]. We will show in the following that our planar subdivision has in fact *linear* complexity.

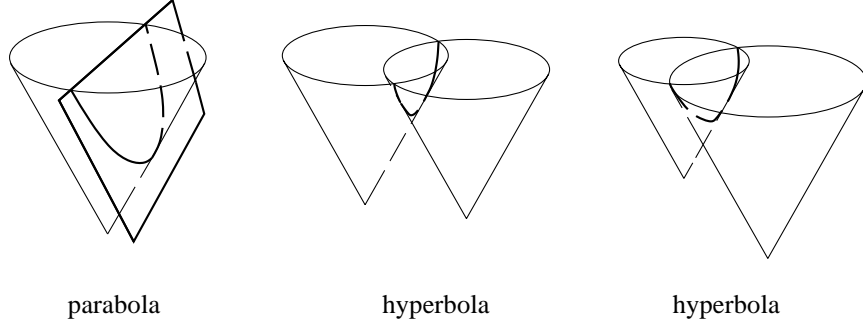


Figure 3: Types of conic sections: parabola and hyperbola

Roughly speaking, we interpret the planar map as a kind of Voronoi diagram. Our sites are the projections of silhouette edges onto the xy -plane, additively weighed by the “height” of the edge above the xy -plane. (This is, indeed, a strange notion of “weight,” as it is not constant for a given site. The concerned reader is asked to wait for the formal definition below.) This diagram does not appear to have been studied before, but it does fit into Klein’s framework of *abstract Voronoi diagrams* [6], and his results on complexity and computation apply.

Consider a silhouette edge e of P . Let \bar{e} be the projection of e onto the xy -plane. For a point $\bar{p} \in \bar{e}$, let p_z be the z -coordinate of the point $p = (p_x, p_y, p_z) \in e$ whose projection on the xy -plane is \bar{p} , and let $w(\bar{p})$ be $p_z \tan \alpha$. We can now define a distance measure in the plane as follows: For $x \in \mathbb{R}^2$ and $\bar{p} \in \bar{e}$, we define

$$d(x, \bar{p}) := |x\bar{p}| + w(\bar{p}) = |x\bar{p}| + p_z \tan \alpha.$$

The distance of a point x to a segment \bar{e} is then

$$d(x, \bar{e}) := \min_{\bar{p} \in \bar{e}} d(x, \bar{p}).$$

Lemma 5 *The vertical projection of the lower envelope E coincides with the Voronoi diagram of the projected silhouette edges under the distance function defined above.*

Proof: Let x be a point in the plane, and let e be a silhouette edge of P . Let x^* be the point where the vertical line through x intersects the boundary of the α -region $D(e)$. We observe that $d(x, \bar{e}) = |xx^*| \tan \alpha$. The lemma follows. \square

In the following lemma, we show some properties of this Voronoi diagram.

Lemma 6 *Let P be an α -safe and α -monotone polyhedron. Consider the Voronoi diagram defined by the projections of a subset G' of silhouette edges of P with the distance function above. It has the following properties:*

- A projected silhouette edge \bar{e} lies in its own Voronoi cell.
- Given a point x in the Voronoi cell of \bar{e} . Let $y \in \bar{e}$ be the point on \bar{e} minimizing the distance from x . Then the segment xy is contained in the Voronoi cell of \bar{e} .
- Each Voronoi cell is simply connected.
- The Voronoi diagram is an abstract Voronoi diagram as defined by Klein et al. [7].

Proof: Let G' be a non-empty subset of silhouette edges and vertices, and let E' be the lower envelope of the α -regions of the silhouette edges in G' .

(i) The claim is identical to stating that the silhouette edge e appears on the lower envelope E' . If it didn't, a point $p \in e$ would have to lie inside the α -region $D(e')$ of some other silhouette edge e' , in contradiction to the assumption that P is α -monotone.

(ii) Assume there is a point $z \in xy$ such that the nearest site point to z is $t \neq y$. Then

$$\begin{aligned} d(x,t) &= |xt| + w(t) \leq |xz| + |zt| + w(t) = |xz| + d(z,t) \\ &< |xz| + d(z,y) = |xz| + |zy| + w(y) = |xy| + w(y) = d(x,y), \end{aligned}$$

in contradiction to the definition of y . So the nearest point on a site is y , for all points on xy , and the segment xy is contained in the Voronoi cell of \bar{e} .

(iii) Follows from (i) and (ii).

(iv) The abstract Voronoi diagram framework by Klein et al. assumes a set of (abstract) objects, each pair of which defines a bisector partitioning the plane into two unbounded regions. The system of bisectors has to adhere to a set of four axioms. It is straightforward to verify that the bisectors defined by pairs of silhouette edges do fulfill these axioms, using (i)-(iii) and elementary calculations. \square

Figure 4 shows the bisector of two projected silhouette edges \bar{e} and \bar{e}' . Note the drop-shaped curves surrounding each edge: these are curves of equal distance from the segment.




Figure 4: $d(\bar{e}, \bar{x}) = d(\bar{e}', \bar{x})$

Lemma 7 Let P be an α -monotone and α -safe polyhedron with n vertices, and let E be the lower envelope of the α -regions of its silhouette edges. Then E has complexity $O(n)$ and can be computed in time $O(n \log n)$.

Proof: From Lemma 5 and Lemma 6 (i)-(iii), we can conclude that E has linear complexity.

We can identify the h silhouette edges of P in $O(n)$ time by inspecting the normals of all facets. By Lemma 6 (iv), the projection of E onto the xy -plane can be computed in time $O(h \log h)$ by the randomized incremental algorithm of Klein et al. [7]. Each face of the Voronoi diagram carries information about the site creating it, and so we can construct the envelope E in linear time $O(h)$. \square

We have now seen how to compute an α -safe terrain E bounded by the silhouette edges of P in time $O(n \log n)$. All that remains to be done to fulfill the assumptions of Lemma 4 is to turn E into a *polyhedral* terrain. We proceed as follows.

The edges of E consist a constant number of segments of two types: straight line segments and conic arcs. Let $\delta = v_1v_2$ be such a conic arc, with endpoints v_1 and v_2 . Its projection $\bar{\delta}$ separates two cells of the Voronoi diagram, say of \bar{e} and \bar{e}' .

We conceptually add four straight line segments to the graph of the Voronoi diagram by connecting both \bar{v}_1 and \bar{v}_2 to the nearest point on both \bar{e} and \bar{e}' . We do this for all conic arcs of E , adding a linear number of “spokes” to the Voronoi diagram graph. The spokes do not intersect, and so we have increased the complexity of the diagram by a constant factor only. As a result, any conic arc $\bar{\delta}$, is now incident to two constant-complexity faces in the diagram. There are two cases, depicted in Figure 5 (a), depending on whether the spokes meet on one or two sides. Without loss of generality, we can assume that the spokes always meet on \bar{e}' .




Figure 5: Approximation of curved surfaces. (a) adding spokes to the diagram: bisector of a vertex and an edge (left) and bisector of two vertices (right). (b) in the new terrain E' , conical facets have been replaced by triangles.

As we have seen before, the conic arc δ is contained in a plane Γ . We now choose a line ℓ in Γ tangent to δ on its convex side, such that its projection $\bar{\ell}$ separates $\bar{\delta}$ from \bar{e} . (If Γ is a vertical plane, then $\bar{\delta}$ is a straight segment, and $\bar{\ell}$ contains $\bar{\delta}$.) Let furthermore ℓ_1 and ℓ_2 be the lines in Γ tangent to δ in v_1 and v_2 . Let $x := \ell \cap \ell_1$ and $y := \ell \cap \ell_2$.

We now construct a new terrain E' by replacing the conic arc δ with the polygonal chain v_1xyv_2 , and replacing the conic surface patches supported by $D(e')$ and $D(e)$ each by three triangles $e'v_1x$, $e'xy$, $e'yv_2$ (and analogously for e if the spokes meet on both sides). Figure 5 (b) shows the projection of the new terrain E' .

We can perform this operation for all conic arcs of E simultaneously, resulting in a polyhedral terrain E' . Note that the triangles lie on planes that are tangent to α -cones $D(e)$ or $D(e')$, and so they are not α -steep. This implies that E' is α -safe. By Lemma 4, the terrain E' defines an α -feasible mould, and

we have the following result.

Lemma 8 *If a polyhedron P is α -monotone and α -safe in direction \vec{d} , then P is α -castable in direction \vec{d} .*

This concludes the proof of Theorem 1; the theorem follows immediately from Lemmas 2 and 8.

Our proof of Theorem 1 is constructive: Given an α -safe and α -monotone polyhedron, we can compute a feasible cast with uncertainty α in time $O(n \log n)$. The construction uses the randomized incremental algorithm by Klein et al. [7], as in Lemma 7.

This procedure does not yet allow us to *decide* whether a polyhedron is indeed α -castable, as we can only guarantee the correctness of the envelope construction if the envelope is indeed an abstract Voronoi diagram. This is not necessarily the case if the polyhedron is not α -monotone. Fortunately, it is not difficult to add a test to each stage of the algorithm by Klein et al. that will detect if P is not α -monotone. This is based on the following lemma.

Lemma 9 *Let P be α -safe and monotone, and let G' be a non-empty subset of the silhouette edges of P . G' is α -monotone if and only if each edge $e \in G'$ appears completely on the lower envelope of the α -regions of G' .*

Proof: The necessity of the condition was already proven in Lemma 6.

Assume that G' is not α -monotone. Then there are two silhouette edges e, e' and two points $p \in e$ and $q \in e'$ such that the slope of pq is greater than $2\pi - \alpha$. The point p , therefore, lies inside the α -region $D(e')$, and so e does not appear completely on the lower envelope. \square

We can now augment the algorithm by Klein et al. to achieve the following result.

Theorem 2 *Given a polyhedron P with n vertices and a direction \vec{d} , we can test the α -castability of P in \vec{d} in time $O(n \log n)$. If it is α -castable, then we can construct an α -feasible mould in $O(n \log n)$ time. The resulting mould has $O(n)$ vertices.*

Proof: We first examine every facet of P and decide whether P is α -safe. If so, we test whether P is monotone in direction \vec{d} , for instance using the algorithm of Ahn et al. [2]. If either step fails, we report that P is not α -castable in direction \vec{d} .

Otherwise, we now use the algorithm by Klein et al. [7] to compute the Voronoi diagram of the projected silhouette edges G . The algorithm incrementally constructs the diagram, while adding the projecting silhouette edges one by one in random order. At each step, it maintains the Voronoi diagram $V(\overline{G'})$ and a so-called history graph $H(\overline{G'})$ of the subset G' of edges inserted so far. When inserting a new silhouette edge $s \in \{G \setminus G'\}$, the algorithm first computes the set E_s of Voronoi edges that are intersected by the Voronoi region $V(s)$ in $V(\overline{G' \cup \{s\}})$. Then it constructs the updated diagram $V(\overline{G' \cup \{s\}})$ and the updated history graph $H(\overline{G' \cup \{s\}})$ by using E_s . This can be done in $O(E_s)$ [7] time.

We know that this procedure works correctly as long as the subset G' is α -monotone. We augment the algorithm such that it recognizes, as soon as a new silhouette edge $s \in G \setminus G'$ is added, whether $G' \cup \{s\}$ is no longer α -monotone. The test relies on Lemma 9. $G' \cup \{s\}$ is not α -monotone if and only if there is $s' \in G'$ such that either $s \cap D(s') \neq \emptyset$ or $s' \cap D(s) \neq \emptyset$. The silhouette edge s' must participate in the definition of the Voronoi edges of E_s , and so we can test this in $O(E_s)$ time. It now suffices to verify that E_s is indeed correctly computed by the algorithm even if $G' \cup \{s\}$ is not α -monotone. \square

4 Computing feasible directions

We now describe an algorithm to solve the following problem: Given a polyhedron P and an angle α , decide whether there is a direction \vec{d} such that P is α -castable in direction \vec{d} . In fact, we will solve the more general problem of finding all directions \vec{d} for which P is α -castable.

We identify the set of directions with the set of points on the unit sphere S^2 centered at the origin. A point p on S^2 corresponds to the direction \vec{d}_p from the origin o to p . Our goal is to identify the region of S^2 corresponding to directions in which P is α -castable.

If we imagine the direction \vec{d} changing continuously, there are directions where an up-facet may become a down-facet, or vice versa. The set of these directions forms a collection M of $O(n)$ great circles on S^2 . We note that P is α -safe in a direction \vec{d}_p if and only if p has distance at least α to all great circles in M .

Let C be a cell of the great circle arrangement of M . If \vec{d} varies inside C , the silhouette edges of P remain the same, but at certain directions the monotonicity of P changes. In fact, this happens when a line parallel to \vec{d} through a silhouette vertex crosses a silhouette edge. The set of directions for which this occurs forms a collection N of $O(n^2)$ arcs of great circles. We note that P is α -monotone in direction \vec{d}_p if and only if P is monotone in direction \vec{d}_p and p has distance at least α to all the arcs in N .

Instead of computing the complete arrangement of $M \cup N$, we can work with a set S of $O(1/\alpha^2)$ sampling points on S^2 . The sampling points S are chosen such that any spherical disc of radius α on S^2 contains a point of S .

For each $s \in S$, we first test whether P is monotone in direction \vec{d}_s in time $O(n \log n)$, using the algorithm by Ahn et al. [2]. If it is, we then construct the cell of the arrangement of M containing s by computing the intersection C_1 of n hemispheres in time $O(n \log n)$. We then compute the $O(n^2)$ arcs of great circles where the monotonicity of P changes within C_1 , and compute the single cell C_2 containing s in their arrangement in time $O(n^2 \log n)$ using the randomized incremental construction algorithm by de Berg et al. [4]. By the observations above, if $p \in C_2$ then P is α -monotone and α -castable in direction \vec{d} if and only if p has distance at least α to the boundary of C_2 . We can compute this set of directions by taking the Minkowski-difference of C_2 and a disc of radius α .

It remains to argue that all feasible casting directions are found this way. Let \vec{d}_p be a direction in which P is α -castable. The spherical disc with center p and radius α contains a point $s \in S$, and does not intersect any great circle arc in M or N . This implies that p and s are contained in the same cell of the arrangement of $M \cup N$. Furthermore, P must be monotone in direction \vec{d}_s . It follows that p will be found by our algorithm.

Finding the direction of maximum uncertainty. It is desirable that the parting terrain of a cast is as “flat” as possible. So while a relatively small uncertainty α may be given as a minimum requirement for manufacturing, we actually prefer to generate casts with uncertainty as large as possible.

We can easily extend the algorithm described above to solve this problem. Again we are given an angle $\alpha > 0$ and wish to test whether P is α -castable. If the answer is positive, we now also want to determine the largest $\alpha^* > \alpha$ for which a direction \vec{d} exists such that P is α^* -castable in direction \vec{d} .

We proceed as above: We generate a sampling set S such that any spherical disc of radius α contains a point of S . We then compute, for each $s \in S$, the cell C_2 containing s . The direction of largest uncertainty within C_2 is the center of the maximum inscribed (spherical) disc for C_2 , which we compute in $O(n^2 \log n)$ time. The largest inscribed disc, over all cells computed, determines the largest uncertainty for which the object is still castable.

Theorem 3 *Let P be a polyhedron with n vertices, and $\alpha > 0$. All directions in which P is castable with uncertainty α can be computed in $O(n^2 \log n / \alpha^2)$ time. If such a direction exists, the largest $\alpha^* > \alpha$ for*

which P is castable with uncertainty α^* can be computed within the same time bound.

A heuristic. If an approximative solution is sufficient, the following heuristic can be applied. It runs in time $O(n \log n)$ for constant α .

Let $\alpha' := (1 - \epsilon)\alpha$, for some approximation parameter $\epsilon > 0$. We choose a set S of $O(1)$ sampling directions on \mathcal{S}^2 , sufficiently dense such that for any spherical disc D of radius α there is a point $s \in S$ such that the disc of radius α' with center s is contained in D .

For each $s \in S$ we test whether P is α -castable using the algorithm of Section 3. If we are successful, we report P to be α -castable. If not, we test each direction $s \in S$ again, this time with uncertainty α' . If no feasible casting direction with uncertainty α' is found, we report that P is not castable with uncertainty α . This is true by the choice of S . If a feasible direction for uncertainty α' is found, we report a “maybe” answer: P is castable with uncertainty $(1 - \epsilon)\alpha$, and may or may not be castable with uncertainty α .

The same idea can be used to approximate the largest feasible uncertainty. We can, for instance, set $\alpha' := \alpha/2$, and keep doubling α until P is no longer α -castable.

References

- [1] H.-K. Ahn, S.-W. Cheng, and O. Cheong. Casting with skewed ejection direction. In *Proc. 9th Annu. Internat. Sympos. Algorithms Comput.*, volume 1533 of *Lecture Notes Comput. Sci.*, pages 139–148. Springer-Verlag, 1998.
- [2] 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.
- [3] P. Bose, D. Bremner, and M. van Kreveld. Determining the castability of simple polyhedra. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 123–131, 1994.
- [4] M. de Berg, K. Dobrindt, and O. Schwarzkopf. On lazy randomized incremental construction. *Discrete Comput. Geom.*, 14:261–286, 1995.
- [5] R. Elliott. *Cast iron technology*. Butterworths, London, UK, 1988.
- [6] Rolf Klein. *Concrete and Abstract Voronoi Diagrams*, volume 400 of *Lecture Notes Comput. Sci.* Springer-Verlag, 1989.
- [7] Rolf Klein, Kurt Mehlhorn, and Stefan Meiser. Randomized incremental construction of abstract Voronoi diagrams. *Comput. Geom. Theory Appl.*, 3(3):157–184, 1993.
- [8] J.-C. Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.
- [9] T. Lozano-Pérez, M. T. Mason, and R. Taylor. Automatic synthesis of fine-motion strategies for robots. *Internat. J. Robot. Res.*, 3(1), 1984.
- [10] Micha Sharir. Almost tight upper bounds for lower envelopes in higher dimensions. *Discrete Comput. Geom.*, 12:327–345, 1994.
- [11] C. F. Walton and T. J. Opar, editors. *Iron castings handbook*. Iron Casting Society, Inc., 1981.