

Filling Polyhedral Molds

Prosenjit Bose *
Carleton University

Marc van Kreveld †
Utrecht University

Godfried Toussaint ‡
McGill University

Abstract

In the manufacturing industry, finding an orientation for a mold that eliminates surface defects and ensures a complete fill after termination of the gravity casting process, is an important and difficult problem. We study the problem of determining a favorable position of a mold (modeled as a polyhedron), such that when it is filled, no air pockets and ensuing surface defects arise. Given a polyhedron in a fixed orientation, we present a linear time algorithm that determines whether the mold can be filled from that orientation without forming air pockets. We also present an algorithm that determines the most favorable orientation for a polyhedral mold in $O(n^2)$ time. A reduction from a well-known problem indicates that improving the $O(n^2)$ bound is unlikely for general polyhedral molds. We relate fillability to some well known classes of polyhedra. For some of these classes of objects, an optimal direction of fillability can be determined in linear time. Finally, for molds that satisfy a local regularity condition, we give an improved algorithm that runs in time $O(nk \log^2 n \log \log(n/k))$, where k is the number of venting holes needed to avoid air pockets in an optimal orientation.

1 Introduction

A well-known technique used in the manufacturing of goods is gravity casting. A mold, as defined in [4], refers to the whole assembly of parts that make up a cavity into which liquid is poured to give the shape of the desired component when the liquid hardens. Given a mold (modeled as a polyhedron representing the surface of the cavity), establishing whether there exists an orientation that allows the filling of the mold using only one pin gate (the pin gate is the point from which the liquid is poured into the mold) as well as determining an orientation that allows the most complete fill are two major problems in the field of gravity casting. These problems are difficult when the focus is on the fluid dynamics and physics of the whole molding process. To date only heuristics have been proposed as solutions to these two problems [4], [12], [24]. In fact, until now, determining the favorable position for filling a mold was considered a cut-and-try process [24].

*Research supported by NSERC Grant OGP0183877 and a FIR Grant. School of Computer Science, Carleton University, 1125 Colonel By Drive Ottawa, Ontario, Canada, K1S 5B6. E-mail: jit@scs.carleton.ca.

†Research supported by an NSERC postdoctoral fellowship, and the ESPRIT Basic Research Action 7141 (project ALCOM II: *Algorithms and Complexity*). Department of Computer Science, Utrecht University, POB 80.089, 3508 TB, Utrecht, The Netherlands. E-mail: marc@cs.ruu.nl

‡Research supported in part by grants NSERC-OGP0009293 and FCAR-93ER0291. School of Computer Science, McGill University, 3480 University Street, Montréal, Québec, Canada, H3A 2A7. E-mail: godfried@cs.mcgill.ca

However, when viewed from a purely geometric perspective, efficient solutions do exist. An initial study of the geometric and computational aspects of mold filling was carried out by Bose and Toussaint [2] who showed that for any mold modeled by a simple polygon P with n vertices, one can decide in $O(n)$ time whether a given orientation allows for a complete fill (the point from which a polygon is filled is always the highest point with respect to the direction of gravity). They also presented an optimal $O(n \log n)$ time algorithm which determines all the orientations that minimize the number of venting holes needed to avoid air pockets (a venting hole is a point from which air, but no liquid, is allowed to escape). This problem is equivalent to finding the orientation that minimizes the number of local maxima of P . Finally, they related fillability to certain known classes of polygons. For a survey on applications of computational geometric techniques to certain manufacturing processes the reader is referred to [3].




Figure 1: Gravity casting of a star-shaped object using one filling hole and two additional venting holes.

In this paper, we study the three-dimensional aspects of the mold filling problem (a preliminary version of this paper appeared in [1]). We show that given a mold, represented by a polyhedron with n vertices in a fixed orientation, we can determine in $O(n)$ time whether or not the mold can be filled without forming air pockets. Then, we show that in $O(n^2)$ time, we can find all the orientations of a polyhedron that allow the mold to be filled without forming air pockets. A polyhedron with this property is referred to as *1-fillable*. This algorithm also finds the orientation that minimizes the number of venting holes needed to ensure a complete fill when the polyhedron is not 1-fillable. The above problem is equivalent to finding the set of orientations of the polyhedron that minimizes the number of local maxima in the positive z -direction assuming that gravity points in the negative z -direction. The pin gate is placed at the global z -maximum of the polyhedron with gravity pointing in the negative z -direction, and the venting holes are placed at the local z -maxima (see Figure 1).

We give a *pseudo-lower bound* on the complexity of this problem by a reduction from the problem

‘ $A+B=C$?’ to mold filling. The problem ‘ $A+B=C$?’ is defined as follows: Given three sets A , B and C of n real numbers each, decide if there exists $a \in A$, $b \in B$ and $c \in C$ such that $a + b = c$. The best known algorithm for ‘ $A+B=C$?’ uses $O(n^2)$ time. Gajentaan and Overmars [11] have shown that there exist many problems in geometry that also reduce to ‘ $A+B=C$ ’, such as: ‘Given a set of n points in the plane, are there three collinear points?’ and ‘Given a set of n rectangles in the plane, do they cover a given rectangle *RECT* completely?’. Since the best known algorithms take $O(n^2)$ time to solve any one of these problems, a problem which can be reduced to one of these is referred to as a *3Sum – Hard* problem. We reduce the rectangle covering problem to the filling problem. Our reduction takes $O(n \log n)$ time. Therefore, the mold filling problem is *3Sum – Hard* making the quadratic bound difficult to improve even to determine whether there is an orientation with only one local maximum.

The interesting question that arises is whether one can improve the $O(n^2)$ time bound for some restricted classes of polyhedra. We relate fillability to certain known classes of polyhedra, namely, star-shaped, monotone and facet-visible polyhedra. In the case of star-shaped polyhedra, this reduces the time bound for finding an optimal orientation to $O(n)$ as opposed to $O(n^2)$. Finally, we show that if a polyhedron is in some sense not too irregular, then we can determine in $O(nk \log^2 n \log \log(n/k))$ time the orientation that minimizes the number of venting holes needed to ensure a complete fill. Here k is the number of local maxima in that orientation. The main idea is that the restrictions imposed on the polyhedron, under a suitable transformation, lead to a set of *fat convex polygons* in the plane, and we use the fact that a set of fat convex polygons has small union size (i.e. the complexity of the boundary of the union is *not* quadratic in the number of vertices of the convex polygons). Both the general algorithm and the improved but restricted algorithm are fairly simple and should perform well in practice. The improved but restricted algorithm provides the correct result for any polyhedron, but a guarantee on the asymptotic running time can only be given for polyhedra with the restriction.

2 Notation and Preliminaries

Let us first introduce some of the terminology we will be using in this paper.

A *simple polygon* is a simply connected subset of the plane whose boundary is a closed chain of line segments. We will denote a polygon by a set of vertices $v_1, v_2, \dots, v_{n-1}, v_n$ such that each pair of consecutive vertices is joined by an edge, including the pair $\{v_n, v_1\}$. We assume that the vertices are in clockwise order, so that the local interior of the polygon lies to the right as the boundary of the polygon is traversed.

Given two points a and b on the plane, let $[ab]$ and (ab) denote respectively the closed and open line segments between the two points. A *chord* of a polygon is a line segment between two points, a and b , on the polygon boundary such that the open line segment is contained in the interior of the polygon. A chord divides a polygon into two subpolygons. Given a line segment e , we denote the line containing e as $L(e)$.

Intuitively, one can think of a *simple polyhedron* as an object which can be continuously deformed into a sphere. More precisely, we define a *simple polyhedron* P as in O’Rourke [21]. The boundary of P is a finite collection of planar, bounded convex polygonal faces such that

1. The faces are disjoint or intersect *properly*. (A pair of faces intersect properly if either they have a single vertex in common or have two vertices, and the edge joining them, in common.)
2. The *link* of every vertex is a simple polygon. (Triangulate the faces that have vertex v on their boundary by adding diagonals adjacent to v . The link of v is the collection of edges opposite v in all the triangles incident to v .)
3. The *one-skeleton* is connected. (The *one-skeleton* is the graph of edges and vertices of the polyhedron.)

The boundary of P is closed and is denoted as ∂P . The boundary encloses a bounded interior region of space, denoted as $int(P)$. The polyhedron consists of the boundary and its interior, (i.e. $P = int(P) \cup \partial P$). The (unbounded) *exterior* of P is denoted as $ext(P)$. Since this paper only deals with simple polyhedra, we will refer to them as polyhedra in the remainder of the paper. The vertices and the edges of the faces are the *vertices* and the *edges* of the polyhedron. The open interiors of the faces are called the *facets* of the polyhedron.

If we intersect a polyhedron with an arbitrary plane, the result is a collection (possibly empty) of polygons (or line segments or points) lying on the plane. A polygon in this collection will be referred to as a *sectional* polygon. Notice that a sectional polygon divides the polyhedron into two simple polyhedra. Thus a sectional polygon is the three dimensional equivalent to a chord in a polygon.

It will be convenient to have the set of all directions in space represented by two planes. Although this is not standard, it will help simplify the exposition. Let the plane $z = -1$, denoted by $DP^{(-)}$, represent all directions with a negative z -component. Let the plane $z = 1$, denoted by $DP^{(+)}$, represent all directions with a positive z -component. We do not consider the horizontal directions. This assumption simplifies our discussion but is not an inherent limitation of our methods. A point q in $DP^{(-)}$ or $DP^{(+)}$ represents the direction \vec{oq} , where o represents the origin of E^3 . Given a direction d , represented by \vec{oq} , we define $opp(d)$ to be the opposite direction. Thus, $opp(d)$ is pointing in the direction of the vector \vec{qo} .

A polygonal chain $C = p_0, p_1, \dots, p_n$ is *monotonic* with respect to direction Θ if the projections of the vertices p_0, p_1, \dots, p_n onto a line in direction Θ are ordered as the vertices in C .

3 Geometric Model

We now define the geometric model of the gravity casting process, referred to as the *gravity model*.

A mold is modeled by a simple polyhedron. The point on the boundary of a mold through which the liquid is poured into the polyhedron is called the *pin gate*. We assume that the pin gate is the only point from which air is allowed to escape unless stated otherwise. A *venting hole* is a point from which only air and no liquid is allowed to escape. We assume that neither the liquid being poured into the mold, nor the air in the mold are compressible. Finally, we assume that air cannot bubble out through the liquid; that is a bubble of air can never be completely surrounded by liquid.

The sole force acting on the liquid is gravity. When a direction of gravity is not specified, we assume, for simplicity of exposition, that gravity points in the negative z -direction. Thus, if only one pin gate is used, we assume it to be a point on the boundary with the highest z -coordinate, since otherwise, the polyhedron could not be completely filled.

When liquid is poured into a polyhedron, the level of the liquid rises in the direction opposite that of gravity. We assume that the advancing *front* of the rising liquid is a plane. The lowest horizontal plane such that all the liquid in the polyhedron is contained below it, is defined as the *level plane*.

When the level plane contains the pin gate, we say the polyhedron is *maximally filled*. A region containing air in a maximally filled polyhedron is called an *air pocket*. A polyhedron is said to be *1-fillable* if there exists a pin gate and direction of gravity such that when the liquid is poured into the polyhedron through the pin gate, there are no air pockets when the polyhedron is maximally filled. We call the highest point (there may be more than one) of an air pocket in a maximally filled mold, the *peak* of the air pocket. This leads to the following characterization.

Observation 3.1 *A polyhedron P in 3-space is 1-fillable in direction $-z$ provided that for every point of P there is a $+z$ -monotone path from it to the z -maximum of P . Thus, a polyhedron is 1-fillable if there is an orientation of P from which it is 1-fillable.*

We extend the notion of fillability in the following two ways. A polyhedron is said to be *k -fillable* if there exists a fixed orientation of the polyhedron, a placement of the pin gate and a placement of $k - 1$ venting holes such that when liquid is poured into the polyhedron through the pin gate, there are no air pockets when the polyhedron is maximally filled. A polyhedron is said to be *k -fillable with re-orientation* provided that the polyhedron can be re-oriented and filled from a new pin-gate after partial filling from an initial orientation and pin gate. We assume that after the completion of a partial filling, the liquid that has been poured into the polyhedron hardens. The number k in this case refers to the number of times that the polyhedron needs to be re-oriented before it is completely filled. Notice that both definitions are identical when $k = 1$. Unless stated otherwise, we will always refer to *k -fillable* as filling from a fixed orientation.

4 The Decision Problem

In this section we present an $O(n)$ time algorithm to decide whether a polyhedron P is 1-fillable given an orientation of the polyhedron.

Let P be a simple polyhedron of which all faces are triangulated, and let v be an arbitrary vertex of P . We define P_v to be the union of the faces incident to v . Let f_1, \dots, f_m be the sequence of faces of P_v such that f_i and f_{i+1} are incident to an edge denoted e_i , and f_m and f_1 are incident to an edge e_m . If v is incident to m faces, and if P has a triangulated boundary, then P_v has $2m$ edges and m vertices besides v . Let S_v be a sphere centered at v , such that S_v intersects only the m edges and faces incident to v , and no other facets, edges or vertices of P .

Definition 4.1 *A vertex v is a locally extreme vertex of P provided that there exists a supporting plane h_v , with $v \in h_v$, such that all adjacent vertices of v lie to the same side of h_v .*

For a locally extreme vertex v and supporting plane h_v , let h_v^+ and h_v^- denote the closed half-spaces above and below the plane h_v . Let h_v^\diamond be the closed half-space bounded by the plane h_v with normal d , containing the vertex v and where $\diamond \in \{-, +\}$ is the opposite of the sign of the z -component in d . Recall that we assume, for simplicity, that d is not a horizontal direction.

Definition 4.2 *A vertex v is a local maximum of P in direction d provided that P_v lies in the closed half-space h_v^\diamond .*

We now prove the theorem used to establish the linear time decision algorithm.

Theorem 4.1 *A polyhedron P is 1-fillable in direction $-z$ if and only if P has precisely one local maximum in direction $+z$.*

Proof: (\Rightarrow) We assume that gravity is in the $-z$ direction. Suppose that P is 1-fillable, and suppose that P has more than one local z -maximum. Let q be a local z -maximum of P which is not the global z -maximum M of P . Let Π be any path from q to M . Since q is a local z -maximum, Π has negative value in its z -component when it leaves q , contradicting observation 3.1

(\Leftarrow) On the other hand, suppose that P has only one local z -maximum M , which must also be the global z -maximum of P . Let p be any point inside P , and let f be the face of P hit by a ray emanating from p vertically upward. Let q be the vertex incident to this face with maximum z -coordinate. Clearly, there is a $+z$ -monotone path from p to q consisting of two segments. If $q = M$ we are done, otherwise q is not a local z -maximum, and it must be incident to an edge with endpoints q and q' such that q' has greater z -coordinate. We repeat the argument with q' for q until the path reaches M . ■

From this theorem, we see that given a polyhedron P and a direction of gravity g , to test 1-fillability of P with respect to g , we need only determine the number of local maxima with respect to gravity. We can determine if a vertex is a local maximum in time linear in the degree of the vertex. Since the sum of the degrees of the vertices of P is linear in n , this immediately gives us a linear time algorithm to determine whether or not a polyhedron is 1-fillable from a fixed orientation.

Theorem 4.2 *Given a polyhedron P , we can determine in $O(n)$ time whether or not the polyhedron is 1-fillable with respect to gravity.*

5 Determining all Directions of Fillability

In this section we present an $O(n^2)$ time algorithm to find the orientation of a given polyhedron P that minimizes the number of venting holes needed in order to ensure a complete fill from a fixed orientation. This orientation is equivalent to the orientation that minimizes the number of local maxima. The algorithm has two stages. In the first stage, the fillability problem is transformed to a planar problem for a set of convex (possibly unbounded) polygons that cover the plane. In the second stage, the following problem is solved: Given a set of n convex polygons in the plane, find the point that is covered by a minimum number of them.

5.1 Transforming Fillability to Covering

Let P be a bounded polyhedron with n vertices, and assume that P is given by its incidence graph (see e.g. [8]). First, we triangulate every face of P (note that each face of P is convex). We choose an initial orientation of P such that no edge of P is vertical. This can be done in $O(n)$ time. Let v be any vertex of P . We extract the description of P_v from the description of P in time proportional to the size of P_v . Let f_1, \dots, f_m be the sequence of disjoint facets incident to v , such that f_i and f_{i+1} are incident to an edge e_i of P_v (and f_m and f_1 are incident to an edge e_m). Let w_1, \dots, w_m be the sequence of endpoints corresponding to e_1, \dots, e_m , see Figure 2.

Suppose that v is a locally extreme vertex. We define the *cone* C_v of v to be the unbounded polyhedron consisting of v as its only vertex, m half-lines E_1, \dots, E_m starting at v , which contain the edges e_1, \dots, e_m , respectively, and m unbounded facets bounded by E_i and E_{i+1} ($1 \leq i \leq m-1$), or E_m and E_1 . Note that C_v need not be a convex polyhedron. Let CC_v be the convex hull of C_v , which is a *convex cone*. The half-lines that are the edges of CC_v are a subset of the edges of C_v ; we denote them by E_{i_1}, \dots, E_{i_j} , where $1 \leq i_1 < \dots < i_j \leq m$. Finally, we define the *normal cone* NC_v of the convex cone CC_v as follows. Let h_{i_1}, \dots, h_{i_j} be the set of planes that pass through v and are perpendicular to E_{i_1}, \dots, E_{i_j} . Let H_{i_1}, \dots, H_{i_j} be the closed half-spaces bounded by h_{i_1}, \dots, h_{i_j} such that they contain E_{i_1}, \dots, E_{i_j} , respectively. Then NC_v is the convex region that is bounded by $H_{i_1} \cap \dots \cap H_{i_j}$. Notice that if CC_v is a “sharp” cone then NC_v is a “blunt” cone, and vice versa.




Figure 2: Left: P_v . Middle: the convex hull CC_v of C_v . Right: the convex cone CC_v and the normal cone NC_v .

Each locally extreme vertex of the polyhedron P defines a convex region in $DP^{(-)}$ and/or $DP^{(+)}$, which corresponds to the directions with respect to which it is a local maximum. Recall that a point q in $DP^{(-)}$ or $DP^{(+)}$ represents the direction \vec{oq} , where o represents the origin of E^3 . Hence, P gives rise to $O(n)$ convex regions in these planes. It follows that a direction for which P has the smallest number of local maxima corresponds to some point in the plane that is covered by the smallest number of convex regions. The following lemma relates the normal convex cones to the direction planes, $DP^{(-)}$ and $DP^{(+)}$.

Lemma 5.1 *For every locally extreme vertex v of a polyhedron P such that v coincides with the*

origin o and direction $d = \vec{oq}$ where q is a point on one of the direction planes, it holds that v is a local maximum in (non-horizontal) direction $-d$ if and only if $q \in NC_v \cap DP^{(-)}$ or $q \in NC_v \cap DP^{(+)}$.

Proof: Let ℓ be the half-line rooted at o with direction d . By construction, the following equivalence holds for any locally extreme vertex v located at o and $\diamond \in \{-, +\}$: There exists a plane h through v with normal d such that $CC_v \subseteq h^\diamond$ if and only if $\ell \subseteq \text{interior}(NC_v) \cup NC_v$. Since the direction d is represented by the point $q = \ell \cap DP^{(\diamond)}$, the lemma follows immediately. ■

Therefore we first determine if v is a locally extreme vertex. This is the case if and only if v is an extremal point in the set $\{v, w_1, \dots, w_m\}$. This is equivalent to the problem of determining if v can be separated from $\{w_1, \dots, w_m\}$ by a plane, which in turn is equivalent to linear programming [7]. Therefore we can determine if v is convex by linear programming in linear time (see e.g. [8, 17, 25]). If v is not a locally extreme vertex, then v is not a local maximum for any direction, and we stop considering v . Otherwise, let h_v be a plane that contains v and has w_1, \dots, w_m to one side of it. Such a plane is returned by the linear programming test. Let h'_v be a plane parallel to h_v which intersects all edges e_1, \dots, e_m . The intersection of h'_v with P_v is a simple polygon \bar{P}_v with m vertices (corresponding to e_1, \dots, e_m) and m edges (corresponding to f_1, \dots, f_m). We compute the convex hull of \bar{P}_v in linear time [16], [18]. Let us denote the convex hull by $CH(\bar{P}_v)$. Let $\bar{e}_{i_1}, \dots, \bar{e}_{i_j}$ be the sequence of vertices of $CH(\bar{P}_v)$, where $1 \leq i_1 < \dots < i_j \leq m$. Clearly, these vertices correspond to the edges e_{i_1}, \dots, e_{i_j} of P_v . We have in fact computed the edges adjacent to v on the convex hull of P_v . This information gives us the description of the convex cone CC_v of v in linear time. Furthermore, the normal cone NC_v can also be computed in additional linear time.

Translate NC_v such that v coincides with the origin o . Let $Q_v^{(-)}$ be the convex polygon $NC_v \cap DP^{(-)}$ and let $Q_v^{(+)}$ be $NC_v \cap DP^{(+)}$. Either $Q_v^{(-)}$ is a bounded convex polygon and $Q_v^{(+)}$ is empty, or vice versa, or both $Q_v^{(-)}$ and $Q_v^{(+)}$ are unbounded convex polygons. The convex polygons have the following meaning: v is a local maximum in a non-horizontal direction $-d$ if and only if the half-line starting at the origin o in direction d intersects the interior of one of the polygons $Q_v^{(-)}$ or $Q_v^{(+)}$ by Lemma 5.1. We compute the convex polygons $Q^{(-)}$ and $Q^{(+)}$ for all vertices of P , giving sets $\mathcal{Q}^{(-)}$ and $\mathcal{Q}^{(+)}$ of at most n convex polygons in the planes $DP^{(-)}$ and $DP^{(+)}$, respectively. The total complexity of the polygons in $\mathcal{Q}^{(-)}$ and $\mathcal{Q}^{(+)}$ is $O(n)$. The question: ‘Is P 1-fillable?’ or ‘Is there an orientation of P such that it has only 1 maximum?’ translates to the question: ‘Is there a point in $DP^{(-)}$ or $DP^{(+)}$ that is covered by only one convex polygon?’ Similarly, the question of k -fillability translates to deciding whether there exists a point that is covered by only k convex polygons. We therefore have established the following result:

Lemma 5.2 *In $O(n)$ time, one can transform the problem of k -fillability to the problem of finding a point in the plane covered by only k convex polygons.*

5.2 Solving the Covering Problem

The next step in the algorithm involves solving the following problem: ‘Given a set \mathcal{Q} of n convex, but not necessarily bounded, polygons in the plane, with total complexity $O(n)$, find a point that

is covered by the minimum number of polygons of \mathcal{Q} .' Our algorithm constructs the subdivision induced by \mathcal{Q} , and associates to each cell the number of polygons that contain it.

The subdivision induced by \mathcal{Q} without the numbering can be constructed deterministically in $O(n \log n + A)$ time by the algorithm of Chazelle and Edelsbrunner[5], where A is the total number of intersection points of all polygons in \mathcal{Q} . Alternatively, a simpler randomized algorithm performs the task with the same time bound, see Clarkson[6] or Mulmuley[19]. Clearly $A = O(n^2)$, and we obtain a planar subdivision \mathcal{S} with $O(n^2)$ vertices, edges and cells. Consider the graph G which has a node for every cell of \mathcal{S} , and an edge between two nodes if the corresponding cells are incident to the same edge of \mathcal{S} . The graph G has $O(n^2)$ nodes and edges. Start at any node a_1 , and compute in $O(n)$ time how many polygons of \mathcal{Q} cover it. Store this number with a_1 . Start from a_1 with a depth first search. Every edge (a_i, a_j) of G we traverse corresponds to going inside or outside a polygon of \mathcal{Q} , in which case we take the number of a_i , add or subtract one from it, and assign this number to a_j . Thus the whole process of assigning values to cells of \mathcal{S} requires only $O(n^2)$ time. The cell with the minimum number assigned to it is covered by the minimum number of polygons.

Returning to the k -fillability problem, the above algorithm finds the direction d such that the polyhedron has the minimum number of local maxima, if we apply it to both the set $\mathcal{Q}^{(-)}$ of convex polygons in the plane $DP^{(-)}$ and $\mathcal{Q}^{(+)}$ in the plane $DP^{(+)}$. We summarize the algorithm below.

Algorithm 1: *Find all orientations such that P is fillable with the minimum number of venting holes.*

1. Select all locally extreme vertices of polyhedron P . Compute the convex cone of each locally extreme vertex. Compute the normal cone of each convex cone. Call this set NC .
2. Intersect each normal cone in NC with $DP^{(+)}$ and $DP^{(-)}$. Call this set of (possibly unbounded) convex polygons R . Compute the arrangement $\mathcal{Q}^{(+)}$ induced by R on $DP^{(+)}$ and $\mathcal{Q}^{(-)}$ induced by R on $DP^{(-)}$.
3. Find all regions on $\mathcal{Q}^{(+)}$ and $\mathcal{Q}^{(-)}$ covered by the least number of convex polygons of the set R . These regions represent the orientations minimizing the number of venting holes need to fill P .

We conclude with the following theorem:

Theorem 5.1 *Given a simple bounded polyhedron P in 3-space, one can find in $O(n^2)$ time an orientation for P such that P is fillable with the minimum number of venting holes.*

Remark: Observe that the solution presented above makes no use of the fact that P is topologically equivalent to a sphere. The algorithm works equally well for a polyhedron topologically equivalent to a torus or polyhedron of yet higher genus. In practice, this is relevant because a plastic cup with an ear is modeled by a polyhedron that is a torus.

6 A Reduction from Covering to 1-Fillability

In this section, we present an $O(n \log n)$ time reduction from the rectangle covering problem to the problem of 1-fillability of polyhedra. Since a reduction from the ‘ $A+B=C$?’ problem to rectangle covering is given by Gajentaan & Overmars[11], it follows that 1-fillability is at least as hard as ‘ $A+B=C$?’.

Theorem 6.1 *The rectangle covering problem can be reduced to the 1-fillability problem in $O(n \log n)$ time.*

Proof: Let I be an instance of the rectangle covering problem, i.e., given a set \mathcal{R} of n rectangles in the plane, and also a rectangle $RECT$, decide if the union of the rectangles in \mathcal{R} cover $RECT$. We now describe the construction of a polyhedron P such that it is 1-fillable if and only if the rectangle $RECT$ is not covered by \mathcal{R} .




Figure 3: Left: an instance of the rectangle covering problem. Middle: a rectangle r_i and its convex cone $CC(r_i)$. Right: the normal convex cone $NC(r_i)$ and the spike s_i .

We associate the plane in which \mathcal{R} and $RECT$ lie with the plane $z = -1$, such that the center of $RECT$ is the point $(0, 0, -1)$. For every $r_i \in \mathcal{R}$, we associate the convex cone $CC(r_i)$ to be the cone with apex the origin o of 3-space, and whose intersection with the plane $z = -1$ is the rectangle r_i . Then we invert $CC(r_i)$ to obtain a convex cone $NC(r_i)$, and we intersect $NC(r_i)$ with the plane $z = -1$ to obtain a possibly unbounded convex polygon Q_i . For each Q_i , we choose a point q_i in its interior such that all of the q_i are distinct. (The convex hull of the q_i should contain the point $(0, 0, -1)$; if not, we add suitably chosen dummy rectangles to \mathcal{R} outside of $RECT$ to enforce this.) Let h_i be the plane through o with normal $o\vec{q}_i$. Translate h_i in direction $o\vec{q}_i$ by an amount such that the interior of $h_i \cap NC(r_i)$ has positive area, but is contained in a disk with diameter 1. Define the *spike* s_i to be the polyhedron $h_i^+ \cap NC(r_i)$. Translate h_i and the spike s_i simultaneously back in direction $q_i\vec{o}$, such that h_i passes through o again.

Let γ be the minimum distance between any two of the distinct points q_i . Let δ be the maximum distance of any q_i to the origin o . Let S be a sphere centered at o with radius at least $2\delta/\gamma + 1$. Translate every pair h_i and s_i in direction $q_i\vec{o}$ such that h_i is tangent to S ($S \subseteq h_i^-$). By the choice of the radius of S and the area of $h_i \cap NC(r_i)$ (the ‘base’ of the spike), no two spikes s_i and s_j intersect. Compute the convex polytope $P = (z \geq -1) \cap \bigcap_{1 \leq i \leq n} h_i^-$. By construction (the addition of dummy rectangles), P is a bounded convex polyhedron. To P , we add each spike s_i on the facet

of P that lies in h_i . To finish the construction, we add one more gadget to the facet contained in the plane $z = -1$. This is the new spike s_{RECT} for $RECT$, which is translated in the $-z$ -direction over a distance so that its topmost point penetrates the lower facet of P .




Figure 4: An example of the polyhedron constructed for theorem 6.1.

Without all the spikes, P is a convex polyhedron, and thus has exactly one maximum for every direction. The spike s_{RECT} gives additional local maxima for every direction corresponding to a point in $z = -1$ outside of $RECT$. The other spikes give a local maximum for every direction that corresponds to a point inside the corresponding rectangles of \mathcal{R} . Hence, P is 1-fillable if and only if $RECT$ is not covered by the union of the rectangles in \mathcal{R} . The construction can be performed in $O(n \log n)$ time using the half-space intersection algorithm of Preparata and Muller [22]. ■

7 Fillability of Certain Classes of Polyhedra

In this section, we investigate the relationship between the notion of fillability and certain known classes of restricted polyhedra. These results are relevant to the manufacturing industry because in practice many objects are not modeled by polyhedra of arbitrary shape complexity.

7.1 Monotone Polyhedra

A polygon P is monotone in direction l if for every line L orthogonal to l that intersects P , the intersection $L \cap P$ is a line segment (or point). We generalize this notion to 3-dimensions to obtain a large family of monotone polyhedra. We define the class as follows.

Definition 7.1 *A polyhedron P is weakly monotone in direction l if there exists a direction l such that the intersection, of each plane orthogonal to l that intersects P , is a simple polygon (or a line segment or point). The direction l is referred to as the direction of monotonicity.*

Note that there exist many different classes of simple polygons [20], [23], [26]. By substituting one of these classes for the word *simple* in the above definition, we obtain a score of families

of *weakly monotonic* polyhedra. Thus we say that if all the intersections are *convex* polygons, we have a weakly monotonic polyhedron in the *convex sense*. If the intersections are *monotone* polygons, then we have a weakly monotonic polyhedron in the *monotone sense*, and so on. Refer to figure 5. Weakly monotone polyhedra have been previously investigated in the context of movable separability of polyhedra [26].




Figure 5: Weakly Monotonic Polyhedron

We observe that a weakly monotonic polyhedron P is 1-fillable if it is oriented such that gravity points in the direction of monotonicity. We omit the easy proof of this fact.

7.2 Facet Visible Polyhedra and Star-Shaped Polyhedra

Two points inside a polyhedron are said to be *visible* if the line segment between them does not intersect the exterior of the polyhedron. A point p is *weakly visible* from a facet f if there is a point x on f such that p is visible from x .

A polyhedron P is *facet visible* if there is a facet of the polyhedron from which all the points in the polyhedron are weakly visible. (Recall that a facet is the open interior of a face.)

Let P be a facet visible polyhedron. Without loss of generality, let f_1 be the open facet from which the polyhedron is weakly visible. Let d^* denote the direction of the interior normal to the facet.

Theorem 7.1 *A facet visible polyhedron P is 1-fillable if it is oriented such that d^* points in the direction of gravity.*

Proof: For ease of exposition, let us assume that gravity is in the negative z-direction.

Let p_1 , an arbitrary point of the facet, be the pin gate. Let a be an arbitrary point in P . Since P is facet visible, there must be a point b on f_1 that sees point a , i.e. $[ab] \in P$.

Let Π be the path $= (a, b, p_1)$ in P . Since Π is monotone with respect to d^* , the theorem follows.

■

Corollary 7.1 *Every polyhedron that is weakly visible from a sectional polygon is 2-fillable with re-orientation.*




Figure 6: A star-shaped polyhedron that is not 1-fillable

A *star-shaped polyhedron* is a polyhedron that contains at least one point x from which all points of the polyhedron are visible (see figures 1 and 6 for a star-shaped polyhedron). The set of points from which all points are visible is known as the *kernel* of the star-shaped polyhedron. A point in the kernel of a star-shaped polyhedron can be computed in $O(n)$ time using Megiddo's linear programming technique [17]. This implies that in $O(n)$ time, a sectional polygon can be found from which the star-shaped polyhedron is weakly visible. However, a star-shaped polyhedron may not necessarily be 1-fillable (see figure 6). In fact, if a star-shaped polyhedron is filled from one fixed orientation, it may need $\Omega(n)$ venting holes.

Theorem 7.2 *A star-shaped polyhedron is not necessarily 1-fillable but can always be 2-filled with re-orientation in $O(n)$ time.*

8 An Improved Algorithm for Restricted Polyhedra

We next consider the issue of improving our $O(n^2)$ time algorithm for polyhedra that satisfy certain regularity conditions. We place a local condition on each vertex to ensure that the covering problem we obtain can be solved more efficiently. Our new algorithm runs in $O(nk \log^{O(1)} n)$ time, and thus is dependent upon the number of local maxima of P . If this number is small compared to n , the new algorithm improves considerably upon the previous algorithm. The local restriction is such that the convex polygons that are obtained for the covering problem are *fat*, that is, the ratio of the diameter to the width of each polygon is bounded from above by a constant (a convex polygon is fat provided that the ratio of the diameter to the width is bounded by a constant). We have identified two different restrictions that each lead to fat polygons. We note that the algorithm

presented in this section provides the correct result for any polyhedron, but a guarantee on the asymptotic running time can only be given for polyhedra satisfying the local conditions.

Definition 8.1 *A locally extreme vertex v of a polyhedron P in 3-space is non-flat provided that there exists a positive constant β and plane h through v such that $P_v \subset h^-$ or $P_v \subset h^+$, and every edge e incident to v makes an angle at least β with h .*

For a point p , a normalized vector v in 3-space, and a non-negative real λ , let ℓ be the half-line $p + \lambda \cdot \vec{v}$. A *cone annulus* CA with radii r_1 and r_2 is the geometric object such that for any plane h that is perpendicular to ℓ and intersects ℓ at a point represented by some value of λ in the equation of the line ℓ , the intersection $h \cap CA$ is an annulus defined by circles with radii λr_1 and λr_2 . The *ratio of the cone annulus* CA is defined to be r_2/r_1 , assuming that $r_2 \geq r_1$.

Definition 8.2 *A locally extreme vertex v of a polyhedron P in 3-space is annulus-bounded provided that there exists a constant ρ and a cone annulus CA_v with apex v and ratio at most ρ , such that $P_v \subset CA_v$.*

Observe that a locally extreme vertex v of P can be non-flat but not annulus-bounded, or annulus-bounded but not non-flat. Non-flatness will usually be the case for any vertex v that is not used to approximate a convex surface in 3-space. If v is used to approximate a convex surface in 3-space, then it may be annulus-bounded or not. Roughly speaking, v will be annulus-bounded if in the neighborhood of v on the surface, the number of approximating points is chosen linearly dependent upon the change in derivative (and not dependent upon the distance).

Definition 8.3 *A bounded polyhedron P in 3-space is restricted provided that each locally extreme vertex v of P is non-flat or annulus-bounded.*

We will show that any restricted polyhedron P yields—using an adapted transformation to the covering problem—a set of fat convex polygons. The second problem is that of computing regions that are covered not too often without having to compute the full subdivision of the convex polygons (which may have quadratic complexity even for fat objects). We will show that one can compute the regions that are at most k -covered in $O(nk \log^2 n \log \log(n/k))$ time, if the convex polygons are fat.

8.1 Transforming a Restricted Polyhedron to Fat Polygons

Let P be a restricted polyhedron with n vertices. As before, we triangulate all facets of P , and for each vertex $v \in P$, we test whether v is locally extreme. If so, we compute the convex cone CC_v and the normal convex cone NC_v . However, instead of letting all non-horizontal directions be represented by the planes $z = -1$ and $z = 1$, we let all directions be represented by an axis-parallel cube. Let DC be the axis-parallel cube centered at the origin o and with edge length 2. The normal cone may intersect all six facets of DC , or its interior may contain a whole facet. For any facet F of DC , consider the set Q of clipped convex polygons $NC_v \cap F$, see Figure 7.



Figure 7: Left: the intersection of the cube with two convex cones. Right: one facet of the cube intersected by several convex cones.

Lemma 8.1 *Any clipped convex polygon $Q \in \mathcal{Q}$ can be extended to a fat convex polygon that coincides with Q inside the facet F .*

Proof: Let h_F be the plane that contains facet F , and let $F' \subset h_F$ be the square with edge length 4 that contains F such that their centers coincide and their edges are parallel. Suppose that $Q = NC_v \cap F$ for an normal cone NC_v . Define $Q' = NC_v \cap F'$ and $Q'' = NC_v \cap h_F$, see Figure 8.




Figure 8: The facet F , the square F' and the polygons Q , Q' and Q'' in h_F .

First, suppose that v is a non-flat vertex. Let h be the plane and β the constant as in Definition 8.1. It follows that CC_v is contained in a circular cone with radius r bounded from above by the constant $1/\tan(\beta)$. Therefore, the normal cone NC_v contains a circular cone with radius $r' \geq \tan(\beta)$. There are two cases: (i) $Q'' \subseteq F'$ (or, equivalently, $Q' = Q''$), and (ii) $Q'' \not\subseteq F'$ (or $Q' \subset Q''$). In the former case, Q' contains a circle with radius at least r' , which provides a lower bound on the width of Q' . The diameter of Q' is at most $4\sqrt{2}$, the diameter of F' . It follows that Q' is fat. In the latter case, consider the convex cone NC_v . It contains a half-line that intersects F , and it contains a circular cone with radius r' . Since the cone is convex, it must contain the closure of the half-line and the circular cone. This object intersects F' in a region with width at least $\min\{\tan(\beta/\sqrt{2}), 1/2\}$. Since Q' contains this region, Q' is fat.

Second, suppose that v is annulus bounded, and let CC_v be contained in a cone annulus with radii r'_1 and r'_2 . It follows easily that NC_v is contained in a cone annulus with radii $r_1 = 1/r'_2$ and

$r_2 = 1/r'_1$. We will show that if $r_2/r_1 = r'_2/r'_1$ is bounded by a constant, then Q' is fat (this only holds if Q' intersects F).

There are two cases: (i) $Q'' \subseteq F'$ (or, equivalently, $Q' = Q''$), and (ii) $Q'' \not\subseteq F'$ (or $Q' \subset Q''$).

In the former case, Q' contains a circle with radius at least r_1 , and Q' is contained in a circle with radius at most $9r_2$. Therefore, the width of Q' is at least r_1 and the diameter is at most $9r_2$, and since r_2/r_1 is bounded, Q' is fat.

In the latter case, let ω be the width of Q' . We will show that $\omega \geq r_1/(3r_2)$. If r_2/r_1 is bounded from above by a positive constant, then ω is bounded from below by a positive constant. Furthermore, the diameter δ of Q' is at least 1 (since Q' intersects the boundaries of both F and F'), and δ is at most $4\sqrt{2}$, the diameter of F' . It follows that the ratio of diameter and width of Q' is bounded by a constant, and therefore, Q' is fat.

Let p_1 and p_2 be two points on Q' that realize the width, and assume without loss of generality that $\omega \leq 1/2$ (otherwise, it follows immediately that Q' is fat). Let ℓ_1 and ℓ_2 be the lines through p_1 and p_2 , and tangent to Q'' . Therefore, they are also tangent to Q' , and Q'' lies between the lines. Let h_1 and h_2 be the planes through o and containing ℓ_1 and ℓ_2 , respectively. Observe that NC_v lies between the planes h_1 and h_2 . Furthermore, $p_1 \in h_1$ and $p_2 \in h_2$ have distance ω in 3-space, and the segment $\overline{p_1 p_2}$ has distance at least 1 to the line $h_1 \cap h_2$. It follows that $r_1 \leq \omega/2$.

Also, since Q' intersects both F and F' , it follows that $r_2 \geq 1/6$. Hence, $\omega \geq r_1/(3r_2)$ as required. ■

8.2 Computing the k -covered Regions for Fat Polygons

A set of n fat convex polygons with total complexity $O(n)$ has the property that the boundary of the union has close to linear complexity, whereas it may be quadratic for non-fat convex polygons. A more general result, which we require for our algorithm, is:

Theorem 8.1 [13, 15] *Given a set \mathcal{Q} of n fat convex polygons with total complexity $O(n)$, the total complexity of all cells in the subdivision induced by \mathcal{Q} which are covered by at most k polygons of \mathcal{Q} is $O(nk \log \log(n/k))$.*

If we clip each polygon in the set \mathcal{Q} with a square F , then the theorem still holds. For some value of k , we will test whether every point in F is at least k -covered. If not, then we know that the polyhedron that gave rise to \mathcal{Q} has less than k local maxima for some orientation. If for each of the six facets of DC every point is at least k -covered, then we know that the polyhedron has at least k local maxima for every orientation.

For a facet F , a set \mathcal{Q} of polygons with total complexity $O(n)$, and an integer $k \geq 1$, we test whether there exists a point inside F that is covered by at most k polygons of \mathcal{Q} . This is done by computing the subdivision defined by \mathcal{Q} using divide-and-conquer. However, we only compute edges and vertices of regions that are covered by at most k polygons of \mathcal{Q} . With every region, an integer is associated that represents the number of polygons that cover the region. The regions for which the integer is greater than k are associated with the special integer ∞ .

We partition \mathcal{Q} into two subsets \mathcal{Q}_1 and \mathcal{Q}_2 such that the total complexity of the polygons in each of the subsets is approximately the same. Recursively, we compute the subdivisions \mathcal{S}_1 and \mathcal{S}_2 of the regions that are covered by at most k polygons of \mathcal{Q}_1 and \mathcal{Q}_2 , respectively. Then we merge these two subdivisions \mathcal{S}_1 and \mathcal{S}_2 as follows.

Perform a plane sweep over $\mathcal{S}_1 \cup \mathcal{S}_2$ to compute the union $\bar{\mathcal{S}}$ of these subdivisions ($\bar{\mathcal{S}}$ is a refinement of both \mathcal{S}_1 and \mathcal{S}_2). The plane sweep stops at every vertex of \mathcal{S}_1 and of \mathcal{S}_2 , and at every intersection point of an edge in \mathcal{S}_1 and an edge in \mathcal{S}_2 . Any region R of $\bar{\mathcal{S}}$ is assigned an integer as follows. Assume $R \subseteq R_1$ and $R \subseteq R_2$ for regions R_1 in \mathcal{S}_1 and R_2 in \mathcal{S}_2 . Let R_1 be associated to i_1 , and R_2 to i_2 . Then the integer i for R is $i_1 + i_2$, unless $i_1 = \infty$, $i_2 = \infty$, or $i_1 + i_2 > k$, in which case i is assigned ∞ . This gives a correct assignment of integers to the subdivision $\bar{\mathcal{S}}$. However, $\bar{\mathcal{S}}$ contains edges between two regions with integer ∞ . To complete the merge step, we remove all the edges from $\bar{\mathcal{S}}$ to obtain the desired subdivision \mathcal{S} . For a detailed account of this plane sweep with merge refer to Mairson and Stolfi [14].

If the polygons in \mathcal{Q}_1 and \mathcal{Q}_2 have complexity $O(n)$, then the subdivisions \mathcal{S}_1 and \mathcal{S}_2 have complexity $O(nk \log \log(n/k))$ by Theorem 8.1. The number of intersection points of some edge of \mathcal{S}_1 and some edge of \mathcal{S}_2 is $O(nk \log \log(n/k))$, since such an intersection point is obtained from two edges that bound a region that is covered at most k times. Hence, any intersection point lies on the closure of a cell that is covered at most $2k$ times. By the complexity result of Theorem 8.1, the plane sweep takes $O(nk \log n \log \log(n/k))$ time. We conclude that the divide-and-conquer algorithm requires $O(nk \log^2 n \log \log(n/k))$ time in total.

We show how to use the above algorithm for k -fillability of a polyhedron P . Let F_1, \dots, F_6 be the six facets of the cube DC , and let $\mathcal{Q}_1, \dots, \mathcal{Q}_6$ be the six sets of polygons on these facets. Let $j = 1$ and determine for each F_i and \mathcal{Q}_i whether there exists a point on F_i that is covered by only j polygons of \mathcal{Q}_i . If the answer is no for all F_i , then we double j and try again. If the answer is yes for some facet F_i , then traverse all subdivisions of F_1, \dots, F_6 that we computed, and find the region with smallest associated integer $k \leq j$. Any point in this region corresponds to an orientation of the polyhedron P for which it has k local maxima. Any other orientation gives at least as many maxima. The time taken by the algorithm is $O(\sum_{b=0}^{\lceil \log_2 k \rceil} n \cdot 2^b \log^2 n \log \log(n/2^b)) = O(nk \log^2 n \log \log(n/k))$ time.

Theorem 8.2 *Let P be a polyhedron with n vertices, such that P is locally annulus bounded. Then an orientation for P can be found such that P can be filled using k venting holes, which is the minimum possible for P , and the algorithm takes $O(nk \log^2 n \log \log(n/k))$ time.*

9 Conclusions and Open Problems

This paper presented the first algorithm to compute an orientation of a polyhedron in 3-dimensional space which is good for a mold used for gravity casting. The orientation found minimizes the number of venting holes needed to guarantee a complete fill without forming air pockets. It was shown that the required orientation is that which minimizes the number of local z -maxima. For a polyhedron with n vertices, our algorithm runs in $O(n^2)$ time and should be easy to implement.

It was also shown that it will most likely be difficult to reduce the $O(n^2)$ bound if no restrictions

on the polyhedron are imposed. However, we gave a second algorithm that should improve the first one in most practical situations and, under some local conditions of the polyhedron, we showed that the second algorithm performs asymptotically better than the first one. If the minimum number of venting holes required to fill the polyhedron is k , our algorithm runs in $O(nk \log^2 n \log \log(n/k))$ time. In practice, k will most probably be small compared to n .

There are several directions for further research. First of all, the algorithms presented in this paper have not yet been implemented, and it would be interesting to test how they behave in practice. Secondly, for the second algorithm, we imposed restrictions on the polyhedron to be able to prove the stated running time. It may, however, be possible to weaken the restrictions that we imposed and still prove the same time bound. Also, a different approach may yield other algorithms that improve upon the $O(n^2)$ time bound for other types of restricted polyhedra. Thirdly, our algorithms require $O(n^2)$ and $O(nk \log n \log \log(n/k))$ storage, respectively. It may be possible to improve this, perhaps by using the topological line sweep algorithm of Edelsbrunner and Guibas[9].

10 Acknowledgements

The authors would like to thank David Thompson for providing a stimulating research environment.

References

- [1] Bose, P., M. van Kreveld, and G. Toussaint Filling polyhedral molds. *Proc. 3rd Workshop Algorithms and Data Struc.*, LNCS 709, Pp. 210–221, 1993.
- [2] Bose, P., and G. Toussaint, Geometric and computational aspects of gravity casting, *Computer-Aided Design*, **27**(6), pp. 455–464, 1995.
- [3] Bose, P., and G. Toussaint, Geometric and computational aspects of manufacturing processes, *Computers and Graphics*, 18:4, pp. 487–497, 1994.
- [4] Bown, J., *Injection moulding of plastic components*. McGraw-Hill, England, 1979.
- [5] Chazelle, B., and H. Edelsbrunner, An optimal algorithm for intersecting line segments in the plane. *J. ACM*, 39, pp. 1–54, 1992.
- [6] Clarkson, K.L., New applications of random sampling in computational geometry. *Discr. & Comp. Geom.*, 2, pp. 195–222, 1987.
- [7] Dobkin, D., and S. Reiss, The Complexity of linear programming. *Theoretical Computer Science*, 11, pp. 1–18, 1980.
- [8] Edelsbrunner, H., *Algorithms in combinatorial geometry*. Springer-Verlag, Berlin, 1987.
- [9] Edelsbrunner, H., and L. Guibas, Topologically sweeping an arrangement. *J. Comp. Sys. Sciences*, 38, pp. 165–194, 1989.

- [10] Fekete, S. and J. Mitchell, Geometric aspects of injection molding, manuscript, State University of New York at Stony Brook, 1992.
- [11] Gajentaan, A., and M.H. Overmars, $O(n^2)$ difficult problems in computational geometry. Tech. Report, Dept. of Computer Science, Utrecht University, 1993.
- [12] Hui, K.C., and S.T. Tan, Mould design with sweep operations - a heuristic search approach. *C.A.D.*, 24:2, pp. 81-90, 1992.
- [13] van Kreveld, M., *Fat dissection and covering, and the union size of polygons. Proc. 3rd Workshop Algorithms and Data Struct.*, LNCS 709, Pp. 452-463, 1993.
- [14] Mairson, H., and J. Stolfi, Reporting and counting intersections between two sets of line segments, in R.A. Earnshaw (Ed.), *Theoretical foundations of computer graphics and CAD*, Vol. F40, Springer-Verlag, Berlin, pp. 307-325, 1988.
- [15] Matoušek, J., N. Miller, J. Pach, M. Sharir, S. Sifrony, and E. Welzl, Fat triangles determine linearly many holes. *Proc. 32nd IEEE Symp. Found. Comp. Science*, pp. 49-58, 1991.
- [16] McCallum, D. and D. Avis, A linear time algorithm for finding the convex hull of a simple polygon. *Inform. Process. Lett.*, 8, pp. 201-205, 1979.
- [17] Megiddo, N., Linear-time algorithms for linear programming in R^3 and related problems. *SIAM J. Comp.*, 12, pp. 759-776, 1983.
- [18] Melkman, A., On-line construction of the convex hull of a simple polyline. *Inform. Process. Lett.* 25, pp. 11-12, 1987.
- [19] Mulmuley, K., A fast planar partition algorithm, I. *Proc. 29th IEEE Symp. Found. Comp. Science*, pp. 580-589, 1988.
- [20] O'Rourke, J., *Art Gallery Theorems and Algorithms*. Oxford University Press, New York, New York, 1987.
- [21] O'Rourke, J., *Computational Geometry in C*, Cambridge University Press, New York, 1994.
- [22] Preparata, F., and D. Muller, Finding the intersection of n half-spaces in time $O(n \log n)$. *Theoretical Computer Science*, 8, pp. 45-55, 1979.
- [23] Preparata, F.P., and M.I. Shamos, *Computational geometry - an introduction*. Springer-Verlag, New York, 1985.
- [24] Pribble, W.I., Molds for reaction injection, structural foam and expandable styrene molding. In: *Plastics mold engineering handbook*, J.H. DuBois and W.I. Pribble (Eds.), Van Nostrand Reinhold Company Inc., New York, 1987.
- [25] Seidel, R., Linear programming and convex hulls made easy, *Proc. 6th ACM Symp. Comp. Geom.*, pp. 211-215, 1990.
- [26] Toussaint, G., Movable Separability of Sets, in *Computational Geometry*, North-Holland, New York, pp. 335-375, 1990.