Geometric and Computational Aspects of Gravity Casting*

Prosenjit Bose and Godfried Toussaint

School of Computer Science
McGill University
3480 University Street
Montréal, Québec, Canada
H3A 2A7

Abstract

In the manufacturing industry, finding an orientation for a mold that eliminates surface defects and insures a complete fill after the termination of the gravity casting process is an important and difficult problem which has not previously been investigated formally. We initiate the study of the gravity casting process from a geometric perspective and present an optimal $\theta(n \log n)$ time algorithm that solves this problem in two dimensions given an object of size n. We also characterize the object shapes (modeled as simple polygons) that can be 1-filled and relate fillability to well known classes of polygons. For certain classes of objects, an optimal direction of fillability can be determined in linear time.

1 Introduction

'Mold orientation during fill is a cut-and-try process to find the most favorable position.' [13]

Our motive in studying gravity casting is to understand the geometric and computational aspects that will allow us to find a favorable mold position.

A mold, as defined in [2], refers to the whole assembly of parts that make up a cavity into which liquid is poured to give the shape of the required component when the liquid hardens. Given a three dimensional object as a mold, 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 important 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. However, when viewed from a purely geometric perspective, optimal solutions exist. We study these problems in the two-dimensional case in order to shed light on the difficulties involved in three dimensions. To this end, we solve the following problems of gravity casting in the plane:

1. Given an n-sided polygon with a point representing the pin gate and given the direction of gravity, we determine in $O(n)$ time whether or not the polygon can be filled without forming air pockets. We extend this result to polygons with holes.

*Research supported in part by NSERC-OGP0009293, FCAR-93ER0291, and an NSERC postgraduate scholarship.

2. Given a simple polygon with or without holes, we determine in $O(n \log n)$ time all the orientations that allow a filling of the polygon from one pin gate without forming any air pockets. For example, a convex polygon allows all orientations.

3. If we are only allowed to fill a polygon from one fixed orientation, we find in $O(n \log n)$ time the orientation that minimizes the number of venting holes needed to fill the polygon with no air pockets forming.

4. We show that some well known classes of polygons can be filled with no air pockets using one pin gate, namely monotone polygons and open-edge visible polygons.

5. 'Molds ... are most commonly of the clam-shell type to facilitate easy removal of the part.' [13] The corresponding two-dimensional object is a clam-shell polygon. A polygon is clam-shell if the 2 pieces of the mold can be removed by translation (these polygons are also referred to as 2-Moldable polygons in [14] where the problem of castability is investigated. The castability problem occurs after the gravity casting process is completed and it is thus closely related to the problem considered here). We show that all clam-shell molds have an orientation that allows them to be filled with one pin gate, and we find all such orientations in linear time.

6. We show that star-shaped polygons are not necessarily 1-fillable, but they are always 2-fillable. The orientation and placement of the pin gate that allows the 2-fillability of a star-shaped polygon are determined in $O(n)$ time.

7. We show that L-convex polygons are not necessarily 1-fillable, but they are always 2-fillable. The orientation and placement of the pin gate that allows the 2-fillability of an L-convex polygon are determined in $O(n \log n)$ time.

8. The class of polygons known as weakly-externally visible includes the class of L-convex polygons (which are 2-fillable). We show however that weakly-externally visible polygons are not 2-fillable.

9. We prove that any polygon that is weakly-internally visible from a chord is 2-fillable.

2 Notation and Preliminaries

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

A *simple polygon* P is a simply connected subset of the plane whose boundary is a closed chain of line segments. We will denote a polygon P by a set of vertices $v_1, v_2, \ldots, v_{n-1}, v_n$ 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 interior of the polygon lies to the right as the boundary of the polygon is traversed. We will also refer to a simple polygon as a polygonal mold or mold since we are modeling the molds by simple polygons.

We will denote the open interior of the polygon P by $int(P)$, the boundary by $bd(P)$, and the open exterior by $ext(P)$. The boundary is considered part of the polygon; that is, $P = int(P) \cup bd(P)$. A *clockwise chain* $C_{ij}(P) = (e_i, e_{i+1}, \ldots, e_j)$ is the clockwise sequence of edges on the boundary of P from i to j. A *counter-clockwise chain* $CC_{ij}(P) = (e_i, e_{i+1}, \ldots, e_j)$ is the counter-clockwise sequence of edges on the boundary of P from i to j. We assume that the vertices of the polygon are in general position, i.e., no three vertices are collinear. This assumption simplifies our discussion but is not an inherent limitation of our methods.

Given a line segment e, we denote the line containing e as $L(e)$. A *convex edge* of a simple polygon refers to an edge e where both endpoints of e are convex vertices. Similarly, a *reflex edge* of a simple polygon refers to an edge e where both endpoints of e are reflex vertices.

Given two points a and b in 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 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.

By a *direction* we mean an equivalence class of *oriented* parallel lines. A given direction Θ will be specified by a point on a unit circle in the following way. Let C be a unit circle with center o. Let x be a point on the boundary of the circle such that *ray*(ox) is parallel to and with the same orientation as Θ. Then direction Θ is represented by the point x. (Refer to Figure 1). A point that is diametrically opposite to x on the unit circle represents the *inverse or opposite direction* to Θ and is denoted by *opp*(Θ). A *right (left) normal* to a given direction Θ is an equivalence class N of *oriented* parallel lines with the property that every member of N is orthogonal to Θ and oriented to the right (left) of Θ. The right normal of Θ will be denoted by $N^+(\Theta)$ and the left normal will be denoted by $N^-(\Theta)$.

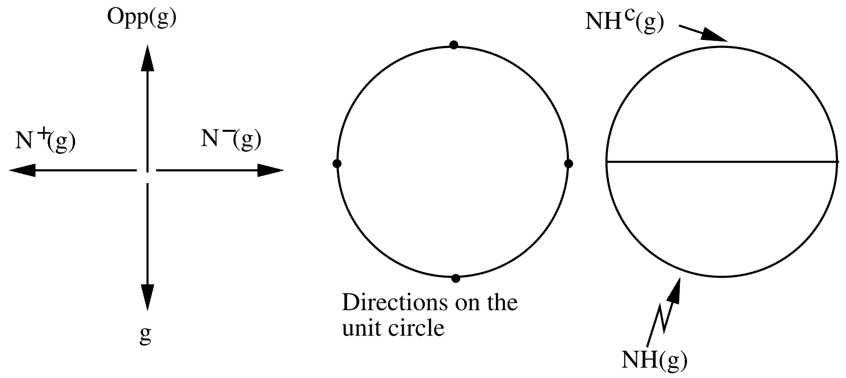

Figure 1: Illustrating the representation of directions.

An equivalence class of parallel lines H will be specified by a pair of points p_1 and p_2 that are diametrically opposite on the unit circle, such that the line determined by the two points is parallel to a line in H. A *normal* to a line in H is an equivalence class N of parallel lines with the property that every member of N is orthogonal to H. Given an oriented direction Θ, we define the *normal* to Θ, denoted by $N(\Theta)$, as the equivalence class of parallel lines that are orthogonal to Θ. Notice that the two points representing the normal divide the boundary of the unit circle into two semi-circles. We refer to the open semi-circle containing the point representing the direction Θ as the open *normal semi-circle* or the open *normal half-plane* of Θ and denote it as $NH(\Theta)$. The closed semi-circle is denoted by $NH[\Theta]$. The open semi-circle not containing the point representing the direction of Θ will be denoted by $NH^c(\Theta)$. Similarly, the closed semi-circle not containing the point is $NH^c[\Theta]$.

Given two points a and b on the unit circle, let $arc[a,b]$ and $arc(a,b)$ denote respectively the closed and open arcs of the unit circle from a to b in the clockwise direction.

2.1 Geometric Model

We are now in a position to define the geometric model of the gravity casting process referred to as the *gravity model.*

The point on the polygon boundary from which the liquid is injected into the polygon is called the *pin gate*. A *venting hole* is a point from which only air and not any liquid is allowed to escape. The pin gate is considered to be a venting hole. 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.

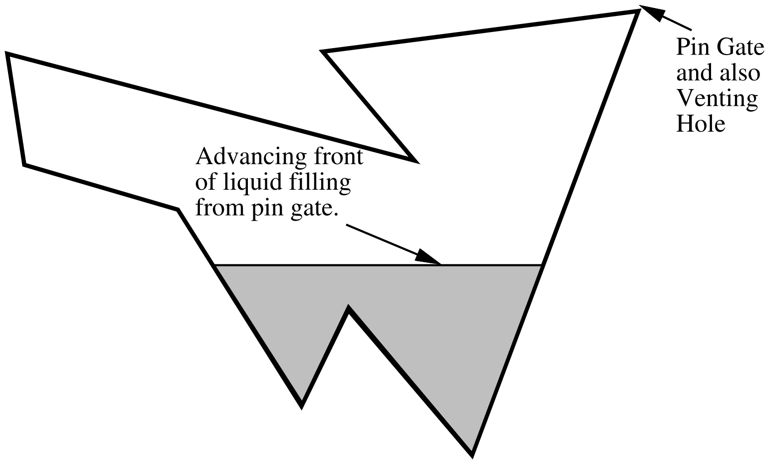

Figure 2: Illustration of the gravity model.

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

When liquid is poured into a polygon, 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 line. The lowest horizontal line such that all the liquid in the polygon is contained below it, is defined as the *level line*.

It is possible for the level line to be higher than the level of the liquid in some section of the polygonal mold. For example, the situation depicted in Figure 3 can occur while the mold is being filled with liquid. Thus we define a *level chord* to be the horizontal chord representing the level of liquid in the subpolygon lying below the chord. The region inside the polygon and above the level line contains air. Similarly, the subpolygon containing the level chord, below the level line inside the polygon, contains air above the level chord.

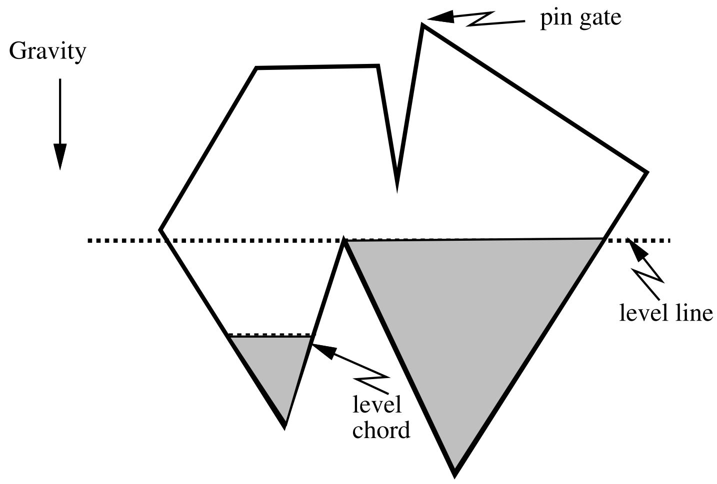

Figure 3: Level line and level chord.

When the level line contains the pin gate, we say the polygon is *maximally filled.* A region containing air in a maximally filled polygon is called an *air pocket.* The highest point (there may be more than one) of an air pocket in a maximally filled mold is the *peak* of the air pocket. A polygon is said to be *k-fillable* if there exists a fixed orientation of the polygon, a placement of the pin gate and a placement of $k-1$ venting holes such that when liquid is poured into the polygon through the pin gate, there are no air pockets when the polygon is maximally filled. A polygon is said to be *k-fillable with re-orientation* provided that the polygon 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 is poured into the polygon hardens. The number k in this case refers to the number of times that the polygon 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.

3 The Decision Problem

We begin by showing that given a simple polygon and the direction of gravity, in linear time, we can decide whether or not the polygon is 1-fillable with respect to the direction of gravity.

Let g be the point on the unit circle representing the direction of gravity. We make the following key observation (refer to Figure 4).

Observation 1 *The peak of an air pocket is a local maximum of polygon P with respect to the direction of gravity. It is either a convex vertex v_i of the polygon P such that* $\mathrm{ray}(v_iv_{i-1}) \in NH(g)$ *and* $\mathrm{ray}(v_iv_{i+1}) \in NH(g)$ *or a convex edge e_i with endpoints v_i and v_{i+1} such that* $\mathrm{ray}(v_iv_{i-1}) \in NH(g)$ *and* $\mathrm{ray}(v_{i+1}v_{i+2}) \in NH(g)$*. Such a vertex or edge will be referred to as a local maximum vertex or local maximum edge.*

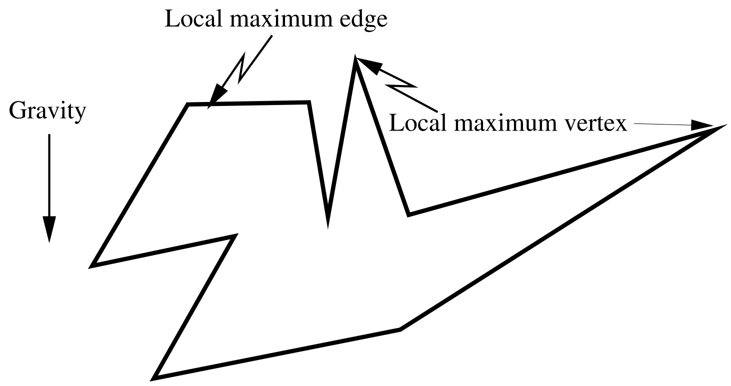

Figure 4: Illustration of local maximum vertex and edge.

The above observation forms the basis of the following theorem characterizing 1-fillable polygons.

Theorem 1 *A polygon is 1-fillable if and only if it contains one local maximum vertex or one local maximum edge with respect to the direction of gravity.*

Corollary 1 *A polygon P is 1-fillable if and only if $\forall p \in P$, the shortest path from p to the pin gate is monotonic with respect to the direction of gravity.*

A chain $C_{ij}(P)$ is *monotonic* with respect to direction Θ if the projections of the vertices $p_i, p_{i+1}, \ldots, p_j$ onto a line $L(\Theta)$ are ordered as the vertices in $C_{ij}(P)$. P is a *monotone polygon* if there exists a line $L(\Theta)$ such that the boundary of P can be partitioned into two chains $C_{ij}(P)$ and $C_{ji}(P)$ that are monotonic with respect to Θ. (Refer to Figure 8).

Before proving Theorem 1, we establish a few lemmas. Given a point p (or horizontal edge e) in the plane, let $h(p)$ $(h(e))$ denote the horizontal line containing p (e). Let P be a polygon with two convex vertices c_1 and c_2 having the following properties. The y-coordinate of c_1 is greater than the y-coordinate of c_2; that is, c_1 is higher than c_2. The edges incident with c_1 lie below $h(c_1)$ and the edges incident with c_2 lie below $h(c_2)$.

Let a be the first point of intersection between $h(c_2)$ and chain $C_{c_2 c_1}(P)$ as the chain is traversed from c_2 to c_1 and let b be the first point of intersection between $h(c_2)$ and chain $CC_{c_2 c_1}(P)$ as the chain is traversed from c_2 to c_1. Let x be the intersection point closest to c_2. Let C be the chain from c_2 to x. All of the vertices of this chain lie below $h(c_2)$.

Lemma 1 *Chain C contains a reflex vertex v or reflex edge e such that the vertices adjacent to v lie above $h(v)$ or the vertices adjacent to e lie above $h(e)$.*

Proof: Both c_2 and x have the same y-coordinate, by construction. We now prove the lemma by induction on the number of vertices in chain C.

Basis: $n = 3$. If the chain has only three vertices including c_2 and x, then the lemma holds trivially.

Inductive Hypothesis: Assume that for a chain with $n \leq k$, $k \geq 3$ vertices, the lemma holds.

Inductive Step: Let $n = k+1$. We have a simple polygonal chain C with $k+1$ vertices. Define P to be $C \cup [c_2 x]$. By definition of x, P must be a simple polygon. Remove an *ear* of P that does not contain c_2 or x. An ear of P is a triangle such that one of its edges is a diagonal of P and the remaining two edges are edges of P. Such an ear must exist since P contains two non-overlapping ears [11].

The new chain C' from c_2 to x contains one fewer vertex. Therefore, by the induction hypothesis, it must contain a reflex vertex v_j with both adjacent vertices lying above $h(v_j)$ or an edge e with both adjacent vertices lying above $h(e)$.

We first consider the case where the chain contains a vertex v_j. The case where the chain contains an edge is similar. Suppose the vertex that was removed from the chain C is not adjacent to v_j, then v_j still satisfies the requirements in C.

Suppose the vertex that was removed is adjacent to v_j, then we must show that C contains a vertex or edge satisfying the required properties. Without loss of generality, let v_{j-1} represent the vertex that was removed.

Consider the horizontal line $h(v_j)$. If v_{j-1} is above the horizontal line, then both vertices adjacent to v_j lie above $h(v_j)$ in C. If v_{j-1} lies on the horizontal line, then both vertices adjacent to $[v_{j-1}v_j]$ lie above $h(v_j)$ in C. Finally, if v_{j-1} lies below $h(v_j)$ then both vertices adjacent to v_{j-1} lie above $h(v_{j-1})$ in C. ∎

Lemma 2 *A polygon with more than one local maximum vertex (or edge) will contain at least one air pocket when maximally filled.*

Proof: We first show that a polygon with more than one local maximum vertex will contain at least one air pocket when filled. We proceed by contradiction. Suppose a polygon P containing at least two local maximum vertices can be maximally filled with no air pockets. Let c_1 and c_2 be the two local maxima, with c_1 having the larger y-coordinate.

By Lemma 1, the polygonal chain between c_1 and c_2 contains a reflex vertex v or reflex edge e such that the vertices adjacent to v lie above $h(v)$ or the vertices adjacent to e lie above $h(e)$. Let us assume it is a vertex v. A similar argument holds for edge e.

Since the polygon can be maximally filled with no air pockets, at some point in time while liquid is being poured in the polygon, there will exist a level chord $[bc]$ containing v. Let b, v, c be the sequence of these three points when viewed in clockwise order starting at b. Since v is a reflex vertex, both $[vb]$ and $[cv]$ are chords. Let P_1 be the subpolygon consisting of the clockwise chain from b to v and the edge $[vb]$ and let P_2 be the subpolygon consisting of the clockwise chain from v to c and the edge $[vc]$. Without loss of generality, let the pin gate be contained in P_2. Polygon P_1 must contain some air since $[bc]$ is a level chord.

Now, every path from a point in P_1 to the pin gate must intersect $[vb]$. But this implies that the air in P_1 is trapped since we assumed that air cannot bubble through liquid. Thus, P_1 contains an air pocket, contradicting the fact that P can be filled with no air pockets.

A similar argument shows that the lemma holds for local maximum edges. ∎

We now establish the converse of the above lemma, thereby proving Theorem 1.

Lemma 3 *A maximally filled polygon containing an air pocket must have more than one local maximum vertex or edge.*

Proof: Let P be a filled polygonal mold with an air pocket. Let p be the peak of the air pocket, and pg be the pin gate. By Observation 1, either p is a local maximum vertex or edge. Also, by assumption, pg is the highest point on the polygon with respect to the direction of gravity. Therefore, it is also a local maximum vertex or contained in a local maximum edge. Since the pin gate cannot be the peak of an air pocket, the polygon contains at least two local maxima. ■

A simple linear time algorithm for the decision problem is implied by Theorem 1. By testing locally, with respect to the direction of gravity, every convex vertex to determine whether or not it is a local maximum vertex, and testing every edge to determine whether or not it is a local maximum edge, we can determine if a polygon is 1-fillable with respect to the direction of gravity. In fact, the number of local maximum vertices and edges determines the number of venting holes that need to be placed in order to fill the polygon in the given orientation.

Lemma 4 *The number of venting holes needed is equal to the number of local maximum vertices and edges.*

Proof: ($\Rightarrow$) We proceed by induction on the number of local maximum vertices and edges. Recall that gravity is assumed to point in the negative y direction.

Basis: The number of local maximum vertices or edges, max, is 1. This implies that the polygon is 1-fillable by Theorem 1. Thus, one venting hole is necessary.

Inductive Hypothesis: Assume that the number of venting holes needed = the number of local maximum vertices and edges, when $max \leq k$, $k \geq 1$.

Inductive Step: Let $max = k + 1$. Suppose that polygon P has at least two local maximum vertices. The argument is similar for local maximum edges.

By Lemma 1 either the clockwise chain or counter-clockwise chain between the two local maxima contains a reflex vertex v or reflex edge e such that the vertices adjacent to v lie above $h(v)$ or the vertices adjacent to e lie above $h(e)$. Without loss of generality, let us assume that there is a reflex vertex v.

Extend a horizontal ray from v to the right until it intersects the polygon boundary. Let i be the intersection point. The chord $[vi]$ partitions the polygon into two subpolygons, P_1 and P_2. Each has less than $k+1$ local maximum vertices or edges by construction. Suppose P_1 has $w \geq 1$ local maximum vertices, then P_2 has $k+1-w$ local maximum vertices. Thus by the induction hypothesis, P_1 needs w venting holes and P_2 needs $k+1-w$ venting holes. This totals to $k+1$ venting holes in P, as required.

($\Leftarrow$) Suppose one of the local maxima in P did not have a venting hole or pin gate. Let v_i represent that local maximum in P. Let x be the farthest point from v_i on $[v_i v_{i+1}]$ visible from v_{i-1}. Triangle(v_{i-1}, v_i, x) contains neither a venting hole nor a pin gate. Therefore it is an air pocket. ■

Theorem 2 *Given a simple polygon and the direction of gravity, in $O(n)$ time one can determine the minimum k for which the polygon is k-fillable.*

Let us turn our attention to polygons with holes. A polygon with holes is defined as a polygon P enclosing several other polygons $H_1, \ldots, H_k$, the holes, such that none of the boundaries of $P, H_1, \ldots, H_k$ intersect and each of the holes is a simply connected region. A polygon with holes is an object that can be constructed with cores and inserts [9, 4]. In a polygon with holes, the peak of an air pocket may not involve a vertex or edge of the polygon, but a vertex or edge of a hole (see Figure 5). Therefore, there are two types of peaks of air pockets that may exist in a polygon with holes. The first type was described in Observation 1. The other type is described in the following observation.

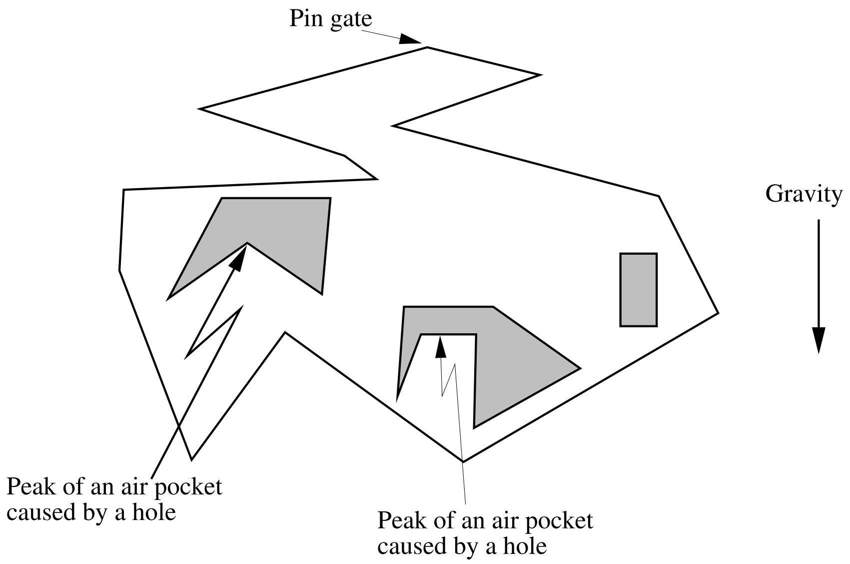

Figure 5: 1-fillable polygon that is no longer 1-fillable because of holes.

Observation 2 *A reflex vertex v or a reflex edge e of a hole H_j in polygon P is a peak of an air pocket provided that it is a local maximum of $bd(H_j)$ with respect to the direction of gravity (see Figure 5). Such a vertex or edge will be referred to as a* reflex maximum vertex *or* reflex maximum edge.

This observation provides a characterization of the peaks of air pockets caused by the presence of holes in a simple polygon. The characterization is similar to that of the peaks of air pockets in simple polygons without holes. Therefore, we have the following theorem.

Theorem 3 *A polygon with holes is 1-fillable if and only if the polygon has only one local maximum vertex or edge with respect to the direction of gravity and none of the holes have a reflex maximum vertex or reflex maximum edge.*

Proof: The argument in one direction is similar to the argument in Lemma 2 and the argument in the other direction is similar to the argument in Lemma 3. ∎

Corollary 2 *A polygon P with holes is 1-fillable if and only if $\forall p \in P$, the shortest path from p to the pin gate is monotonic with respect to the direction of gravity.*

Similar to the case of simple polygons without holes, to determine if a simple polygon with holes is k-fillable, simply test with respect to the direction of gravity, every convex vertex and convex edge of the polygon to determine whether or not it is a local maximum and test every reflex vertex and reflex edge of the holes to determine whether or not it is a reflex maximum. Since testing a vertex or edge can be done in constant time, we have the following theorem.

Theorem 4 *Given a simple polygon with holes and the direction of gravity, one can determine the minimum k for which the polygon is k-fillable in $O(n)$ time where n is the number of vertices of the polygon and the holes.*

4 Determining all Directions of Fillability

In the previous section, we showed that given a simple polygon without holes and the direction of gravity, we can determine in linear time the minimum k for which the polygon is k-fillable with respect to gravity. The extension to polygons with holes was immediate from Observation 2, so for simplicity of exposition, we continue the discussion with simple polygons without holes.

Suppose that we are given a polygon and asked whether there exists an orientation of the polygon such that the polygon is 1-fillable with respect to the direction of gravity. For example, the polygon in Figure 6 is 1-fillable in one orientation but not in another. In this section, we show that in $O(n \log n)$ time the complete range of directions of gravity that allow the 1-fillability of a polygon can be determined. In fact, in optimal $\Theta(n \log n)$ time we can determine all the orientations of the polygon that allow it to be k-filled where k is minimum over all orientations.

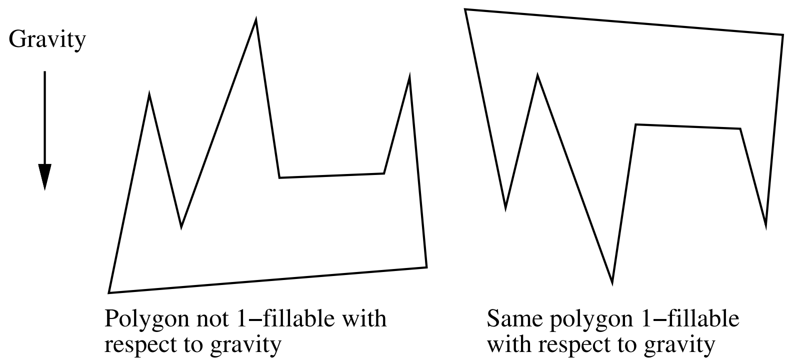

Figure 6: A polygon that is 1-fillable from one orientation but not another.

Let us examine the set of directions that cause a convex vertex to be a local maximum. Given a convex vertex v_i of polygon P, the set of directions for which v_i is a local maximum is defined as follows and will be denoted by $M(v_i)$ (refer to Figure 7).

$$M(v_i) = \{\forall \text{ directions } d \parallel ray(v_iv_{i+1}) \in NH(d) \text{ and } ray(v_iv_{i-1}) \in NH(d)\}$$

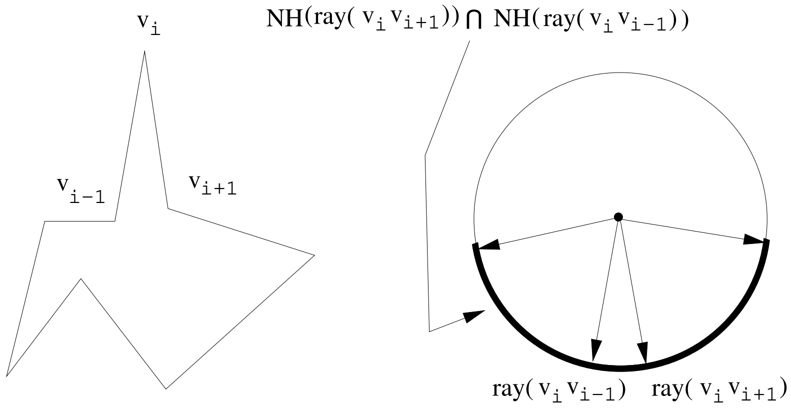

Figure 7: When a convex vertex is a local maximum.

In the following lemma, we characterize the directions in the set $M(v_i)$.

Lemma 5 $M(v_i) = NH(ray(v_iv_{i-1})) \cap NH(ray(v_iv_{i+1}))$

Proof: Let $r_{i-1} = ray(v_iv_{i-1})$ and $r_{i+1} = ray(v_iv_{i+1})$. For v_i to be a local maximum with respect to a given direction d, both r_{i-1} and r_{i+1} must be in $NH(d)$. Every direction $\theta \in NH(r_{i-1}) \cap NH(r_{i+1})$ is contained in $M(v_i)$ since both r_{i-1} and r_{i+1} are contained in $NH(\theta)$. Therefore let us consider an arbitrary direction $y \in M(v_i)$ that is not in $NH(r_{i-1}) \cap NH(r_{i+1})$. Then, one of r_{i-1} or r_{i+1} is not contained in $NH(y)$, contradicting the fact that v_i is a local maximum with respect to y. ■

A convex edge can be a local maximum edge only when the direction of gravity is orthogonal to the line containing the edge. Therefore, the proofs of Lemma 4 and Lemma 5 suggest the following algorithm to find the minimum number of venting holes needed to fill a polygon given that it must be filled in only one orientation.

Algorithm 1

1. Find the direction that minimizes the number of local maximum vertices and edges.
2. Place the pin gate at the global maximum. [This can be done in $O(n)$ time].
3. Place a venting hole at every local maximum that is not a global maximum. [This can be done in $O(n)$ time by just scanning the boundary of the polygon].

Let us elaborate on the first step. Each convex vertex has an open arc representing the set of directions that cause that vertex to be a local maximum. If vertex v_i is convex, we denote the arc by $arc(a_i, b_i)$. Let A be the set of endpoints of all the arcs. Similarly, each convex edge has a point representing the direction that causes it to be a local maximum. If edge e_i is convex, let p_i represent this point. Let E be the set of all of these points. Pick an arbitrary direction $d \notin E$, and radially sort $A \cup E$ in a clockwise manner with respect to d. Let $S = s_1, s_2 \ldots s_m$ represent this sorted order.

We now perform a rotational sweep to determine the set of directions that minimizes the number of local maxima. Let c be the number of vertices that are local maxima with respect to d. Consider the first element of the sequence S. If s_1 is the end of an arc, then we know that the directions represented by $arc[d, s_1)$ have c local maxima and at direction s_1 have $c-1$ maxima. If s_1 is the beginning of an arc, then we know that the directions represented by $arc[d, s_1]$ have c local maxima and the directions after s_1 have $c+1$ local maxima. Finally, if s_1 is a point of E, then we know that the directions represented by $arc[d, s_1)$ have c local maxima, and at direction s_1 there are $c+1$ local maxima. By proceeding in this manner, the intervals on the unit circle induced by the set S are labeled with the number of local maxima present for each interval of directions. By choosing all the intervals with the smallest number, we have a complete description of all the directions from which the given polygon can be filled with a minimum number of venting holes. Some care must be taken when the endpoints of arcs and points coincide. The details of this technique may be found in [8]. The time complexity of this step is $O(n \log n)$ due to sorting.

The correctness of the algorithm follows from Lemma 4 and Lemma 5. The time complexity of the algorithm is dominated by step 1. Therefore, the total time complexity of the algorithm is $O(n \log n)$.

Theorem 5 *The minimum number of venting holes needed to fill a simple polygon from one fixed direction can be computed in $O(n \log n)$ time.*

In [5], it was shown that an $\Omega(n \log n)$ lower bound exists for the problem of determining the minimum number of venting holes to fill a simple polygon from one fixed direction by a reduction from *Element Uniqueness*.

Notice that the technique used in Algorithm 1 is not restricted to finding an orientation of a polygon that minimizes the number of local maxima. A local minimum can only be a convex vertex or convex edge, and the directions that cause such a vertex or edge are defined similarly to the directions that cause them to be local maxima. Therefore, this technique can be used to maximize the number of local maxima, minimize the number of local minima, maximize the number of local minima, minimize the combined number of local minima and maxima and maximize the combined number of local minima and maxima.

Theorem 6 *Given a simple polygon, in $O(n \log n)$ time, we can find the set of directions that minimize or maximize the number of local minima, the number of local maxima or the combined number of local minima and maxima.*

5 Fillability of Certain Classes of Polygons

As there is an $\Omega(n \log n)$ lower bound for determining the orientation minimizing the number of venting holes needed to fill a simple polygon, we study the relationship between certain known classes of polygons and fillability. We show that for some restricted classes of polygons, the optimal orientation for filling can be determined in linear time.

5.1 Monotone Polygons

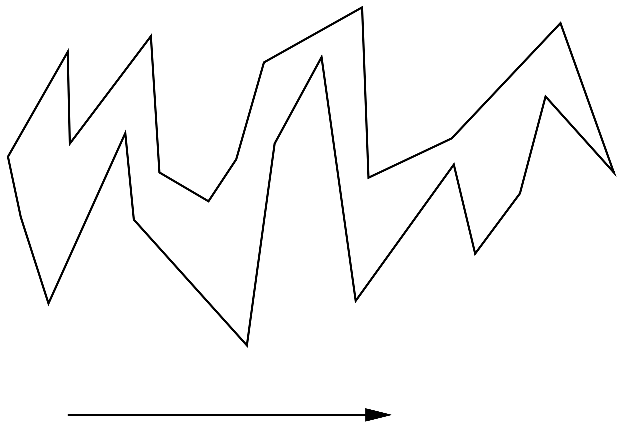

Figure 8: Monotone polygon.

Theorem 7 *A monotone polygon is 1-fillable if it is oriented such that gravity is a direction of monotonicity.*

Proof: Let P be a monotone polygon oriented such that g represents a direction of monotonicity and the direction of gravity. Without loss of generality, let this direction be the negative y direction.

If P is not 1-fillable, then by Theorem 1, it must contain at least two local maxima with respect to the direction of gravity. By Lemma 1, there exists a reflex vertex v or reflex edge e such that the vertices adjacent to v lie above $h(v)$ or the vertices adjacent to e lie above $h(e)$. The existence of such a vertex or edge violates monotonicity. Therefore, a monotone polygon is 1-fillable. ∎

Since monotone polygons can be recognized in linear time [12] and the direction of monotonicity delivered as a witness, Theorem 7 provides a linear time algorithm for determining the optimal orientation of simple polygons if they are monotone.

5.2 Weakly-Edge Visible Polygons and Star-Shaped Polygons

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

A polygon P is *edge visible* if there is an edge in the polygon from which all the points in the polygon are weakly visible. A polygon P is *open-edge visible* if there is an edge e in P such that all points p are visible from some point x on e other than the endpoints of the edge.

Let P be an open-edge visible polygon. Without loss of generality, let (v_1v_2) be the open edge from which the polygon is weakly visible. Let the polygon be oriented such that gravity g^* is the clockwise normal to $ray(v_1v_2)$. (Refer to Figure 9).

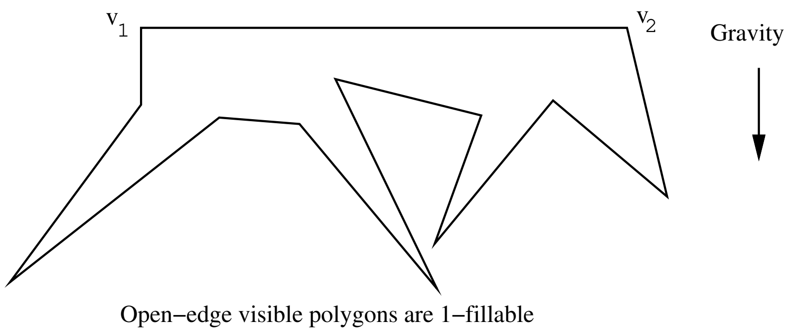

Figure 9: Open-edge visible polygon.

Theorem 8 *An open-edge visible polygon P with gravity g^* is 1-fillable.*

Proof: For ease of exposition, let us assume that gravity is in the negative y-direction. Without loss of generality, let v_1 be the pin gate. Let p be an arbitrary point in P. Since P is open-edge visible, there must be a point z on (v_1v_2) that sees point p, i.e. $[pz] \in P$.

Let Π be the shortest path from p to v_1 in polygon P. Since $[v_1p]$ is monotone with respect to gravity, and Π is a convex chain from p to v_1 contained in the triangle(z, p, v_1), Π is monotone with respect to gravity. Therefore, the lemma follows from Corollary 1. ■

Since open-edge visible polygons can be recognized in linear time [1, 15] and the required edge delivered as a witness, Theorem 8 provides a linear time algorithm for determining the optimal orientation of simple polygons if they are edge-visible.

Corollary 3 *Any polygon that is weakly visible from a chord is 2-fillable with re-orientation.*

A *star-shaped polygon* is a polygon that contains at least one point x from which all points of the polygon are visible. The set of points from which all points are visible is known as the *kernel* of the star-shaped polygon. The kernel of a star-shaped polygon can be computed in $O(n)$ time using the algorithm of Lee and Preparata [7] or a point in the kernel of a star-shaped polygon can be computed in $O(n)$ time using Megiddo's linear programming technique [10]. This implies that in $O(n)$ time, a chord can be found from which the star-shaped polygon is weakly visible. However, a star-shaped polygon may not necessarily be 1-fillable. The star-shaped polygon in Figure 10 is not 1-fillable since there are always two local maxima with respect to every direction. Therefore, we have the following.

Theorem 9 *A star-shaped polygon is not necessarily 1-fillable but is always 2-fillable with re-orientation in $O(n)$ time.*

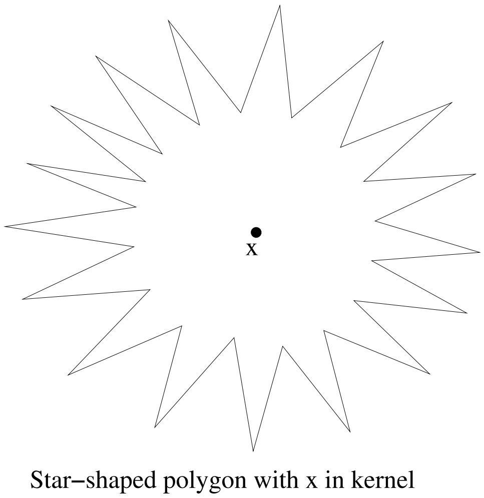

Figure 10: Star-shaped polygon.

5.3 Clam-Shell Polygons

A polygon is *clam-shell* if it can be partitioned into two chains $C_{ij}(P)$ and $C_{ji}(P)$ such that each chain can be removed from the mold by a single translation (not necessarily in a common direction). In Figure 11, we have a polygon that is not clam-shell. Clam-shell polygons were studied in [14] where the following result was proved. (Refer to Figure 12).

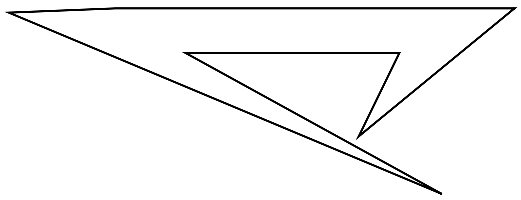

Figure 11: Simple polygon that is not clam-shell.

Theorem 10 *[14] A polygon is clam-shell if and only if the boundary can be decomposed into two chains, each monotonic to an arbitrary direction. Clam-shells can be recognized in linear time.*

Thus we see that this class is a generalization of monotonic polygons. Before showing that all clam-shell polygons are 1-fillable, we establish some key properties of clam-shell polygons. Let $C_{ij}(P)$ be monotonic with respect to some direction Θ. The monotonicity of $C_{ij}(P)$ implies that the chain has only one local maximum for all directions in $arc[\Theta, opp(\Theta)]$. Therefore, we have the following theorem.

Theorem 11 *A clam-shell polygon is 1-fillable.*

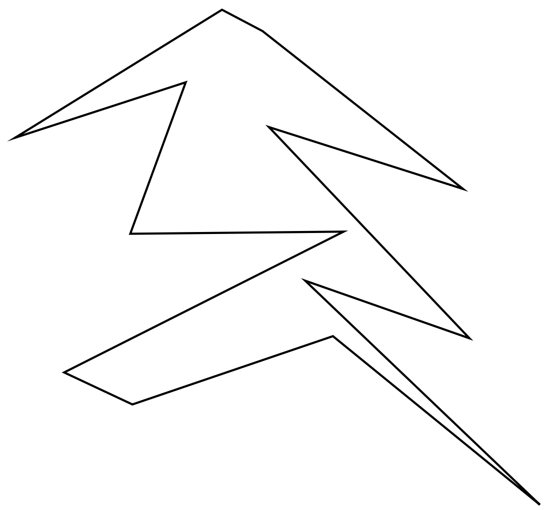

Figure 12: Clam-shell polygon.

Proof: Given a clam-shell polygon P, let $C_{ij}(P)$ be monotone with respect to direction Θ_1 and $C_{ji}(P)$ be monotone with respect to direction Θ_2. The intersection $arc[\Theta_1, opp(\Theta_1)] \cap arc[\Theta_2, opp(\Theta_2)]$ must be non-empty since both arcs are closed semi-circles. ■

Since clam-shell polygons can be recognized in linear time [14] and the required partition of the boundary delivered as a witness, Theorem 11 provides a linear time algorithm for determining the optimal orientation of simple polygons if they are clam-shell.

5.4 L-Convex Polygons

A polygon P is *L-convex* if $\forall x, y \in$P$, \exists z \in P$ such that $[xz] \in P$ and $[yz] \in P$. (Refer to Figure 13). From the definition of L-convexity, we see that a star-shaped polygon is L-convex since all points are seen by a kernel point.

Lemma 6 *[6] If P is an L-convex polygon, it has the property that for every point $x \in P$, there exists a chord of the polygon containing x from which P is weakly visible.*

The lemma proved in [6] is more general and applies to L-convex sets. From Corollary 3 together with the example of the star-shaped polygon that is not 1-fillable (Figure 10), we have the following.

Theorem 12 *An L-convex polygon is not necessarily 1-fillable but always 2-fillable with re-orientation.*

Lemma 6 only proves the existence of the chord, but does not offer a method of computing such a chord given an L-convex polygon. We now present an $O(n \log n)$ time algorithm to find such a chord. We first prove a few key lemmas.

Let x be a point inside polygon P. The *visibility polygon* from x, denoted by $VP(x, P)$, is the set of points in P visible from x. It is formally defined as $VP(x, P) = \{z \mid z \in P \text{ and } [xz] \cap P = [xz]\}$

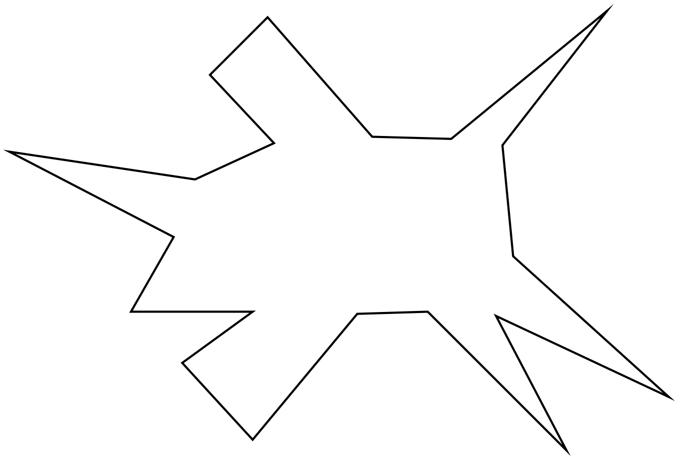

Figure 13: L-convex polygon.

Let v_i be an arbitrary vertex of an L-convex polygon P. By Lemma 6, there must be a chord containing v_i from which P is weakly visible. Let us denote this chord by $C(v_i)$. Thus, $C(v_i)$ must be contained in $VP(v_i, P)$. We will now show the relationship between a diagonal in $VP(v_i, P)$ and $C(v_i)$. Let b be a vertex of $VP(v_i, P)$. The chord $[v_i b]$ is a chord in P and divides the polygon into two subpolygons, P_1 and P_2. If both P_1 and P_2 are weakly visible from $[v_i b]$, then we have found $C(v_i)$. Otherwise, we will show that either P_1 or P_2 has to be weakly visible from $[v_i b]$.

Lemma 7 *Either P_1 or P_2 or both are weakly visible from $[v_i b]$.*

Proof: If both are weakly visible from $[v_i b]$ then the lemma holds. Suppose that one of P_1 or P_2 is not weakly visible from $[v_i b]$.

Case 1: $C(v_i) \in P_2$. By Jordan's Curve Theorem, every line segment with one endpoint in P_1 and one endpoint in P_2 must intersect $[v_i b]$. P_1 is weakly visible from $C(v_i)$. This means $\forall x \in P_1, \exists y \in C(v_i)$ such that $[xy] \in P$. But then $[xy]$ must intersect $[v_i b]$. Thus, P_1 is weakly visible from $[v_i b]$.

Case 2: $C(v_i) \in P_1$. Symmetric to case 1. ∎

Notice that if P_1 is not weakly visible from $[v_i b]$ then $C(v_i)$ is contained in P_1 and vice versa. Thus, $C(v_i)$ is either a diagonal of the polygon P or contained between two consecutive diagonals. The question now becomes, how do we compute $C(v_i)$ if it is contained between two consecutive diagonals?

Consider the following situation. Let r and l be two consecutive vertices in $VP(v_i, P)$ such that P_1 is weakly visible from $[v_i l]$ but not $[v_i r]$ and P_2 is weakly visible from $[v_i r]$ but not $[v_i l]$.

Let $a_1, a_2, \ldots, a_m$ be the vertices in P_1 that are not weakly visible from $[v_i r]$. They must be visible from $[lr]$ since $C(v_i)$ is contained in the triangle(v_i, l, r). Similarly, let $b_1, b_2, \ldots, b_t$ be the vertices of P_2 that are not weakly visible from $[v_i l]$. They must also be visible from $[lr]$. Let a'_i be the point on $[lr]$ farthest away from l, from which a_i is visible. Let $\overline{a} = \min_{1 \le i \le m} \|[la'_i]\|$. Similarly, let b'_i be the point on $[lr]$ farthest away from r, from which b_i is visible. Let $\overline{b} = \min_{1 \leq i \leq t} \|[lb'_i]\|$. Let $S = \overline{a} \cap \overline{b}$.

Lemma 8 *S is not empty.*

Proof: Suppose S was empty. We know that $C(v_i) \in \text{triangle}(l, r, v_i)$. Now, if $C(v_i) \cap [lr] \notin \overline{b}$, then there would be some b_i that was not weakly visible from $C(v_i)$ which is a contradiction. Similarly, if $C(v_i) \cap [lr] \notin \overline{a}$, then there would be some a_i that was not weakly visible from $C(v_i)$ which is a contradiction. ∎

Lemma 9 *For every $s \in S$, the polygon P is weakly visible from $[v_i s]$.*

Proof: If a point $p \in P$ is weakly visible from both $[v_i l]$ and $[v_i r]$ then it must also be weakly visible from $[v_i s]$ since any line segment that intersects both $[v_i l]$ and $[v_i r]$ also intersects $[v_i s]$.

If p is not weakly visible from $[v_i r]$ then let p' be the farthest point from l from which p is weakly visible. $\overline{a}$ must be contained in $[lp']$. Thus, p must be weakly visible from $[v_i s]$. A similar argument shows that if p is not weakly visible from $[v_i l]$, it is still weakly visible from $[v_i s]$. ∎

Lemma 7 and Lemma 9 suggest the following algorithm.

Algorithm 2

1. Choose a vertex $v_1 \in P$.
2. Compute $VP(v_1, P)$ in $O(n)$ time using the algorithm of Avis and Elgindy [3]. [Let $\{v_1, \ldots, v_k\}$ denote the vertices of $VP(v, P)$]
3. Triangulate $VP(v, P)$ by inserting all diagonals $d_i = [v_1 v_i]$, $2 \leq i \leq k$.
4. Let $l = 2$, $r = k$.
5. Let $s = \lfloor r/2 \rfloor$
6. Let P_1 be the polygon with vertices $\{v_1, v_l, v_{l+1}, \ldots, v_s\}$
7. Let P_2 be the polygon with vertices $\{v_1, v_s, v_{s+1}, \ldots, v_r\}$
8. If both P_1 and P_2 are weakly visible from d_s, exit with d_s. [Use the algorithm of Avis and Toussaint [1] to verify weak visibility from d_s in $O(n)$ time].
9. If $r - l = 1$ go to step 12. [two consecutive diagonals at this point].
10. If P_1 is not visible from d_s, then $r = s$, go to step 5.
11. If P_1 is not visible from d_s, then $l = s$, go to step 5.
12. Compute $\overline{a} \cap \overline{b}$ in $O(n)$ time using the algorithm of Avis and Toussaint [1]. Pick any point s in $\overline{a} \cap \overline{b}$ and exit with $[v_1 s]$.

The correctness of the algorithm follows from the discussion. Each step takes at most $O(n)$ time and we loop through steps 5 to 11 at most $O(\log n)$ times. Thus the total time complexity of the algorithm is $O(n \log n)$.

5.5 Weakly-Externally Visible Polygons

We have seen that the class of clam-shell polygons are 1-fillable and the class of star-shaped polygons are 2-fillable. We have also seen that the class of L-convex polygons, which contains the class of star-shaped polygons is also 2-fillable. A natural question is whether the class of weakly-externally visible polygons is 2-fillable, since the class of weakly-externally visible polygons contains the class of L-convex polygons (this is shown in [6]) and the class of clam-shell polygons. A polygon is weakly-externally visible if for every point x on its boundary there is an infinite ray emanating from that point in some direction that intersects the boundary only at x. In fact, the class of weakly-externally visible polygons is in some sense the largest class of interest since it is the largest class of polygons that allows the pin gate to be placed anywhere on its boundary, if we require the pin gate to be reachable by a line probe from infinity. Consider the weakly-externally visible

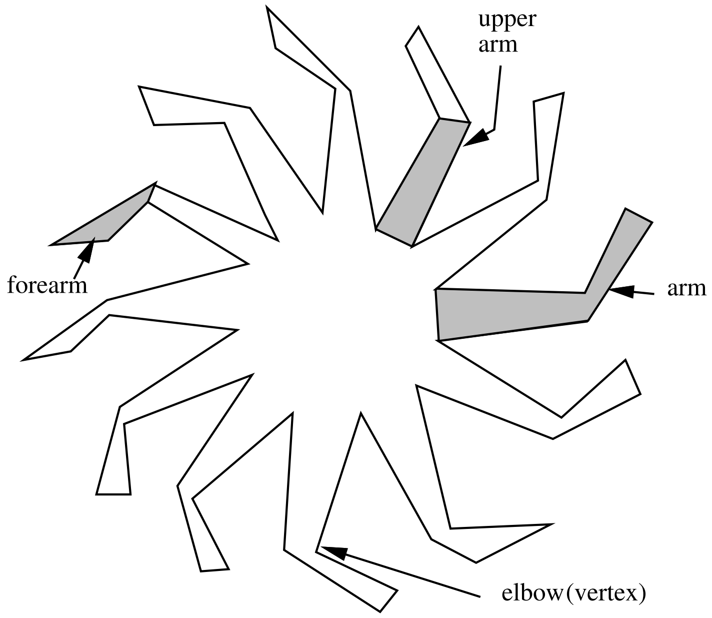

Figure 14: Weakly-externally visible polygon.

polygon in Figure 14. No matter what direction we inject it from, at most only 1/4 of the arms will be completely filled. But, 3/4 of the arms cannot be filled from a second pin gate and orientation. Therefore, we have the following theorem.

Theorem 13 *A weakly-externally visible polygon is not necessarily 2-fillable with re-orientation.*

6 Conclusion

In this paper, we have presented a linear time algorithm to decide whether an arbitrary simple polygon is 1-fillable in a given direction. We have also presented a $\Theta(n \log n)$ time algorithm to find all the directions from which a simple polygon is k-fillable, for minimum k. Finally, we have related the notion of fillability to some known classes of polygons and determined the complexity of filling them. We summarize those relationships by the chart in Figure 15. The fundamental ideas presented here generalize to the more important problem in three dimensions although the technical details are more complex. We are presently developing algorithms for this problem which will be the topic of a forthcoming paper.

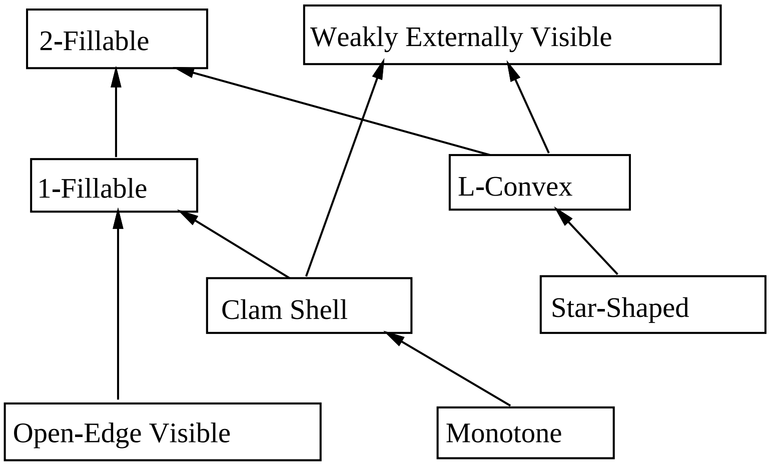

Figure 15: Summary of relation between fillability and classes of polygons.

7 Acknowledgements

The authors would like to thank Francois Trochu of the Mechanical Engineering Department of the Ecole Polytechnique for suggesting the problem to us. We thank an anonymous referee for many valuable suggestions. We also thank Paco Gomez, Marc Van Kreveld, Joe Mitchell, and Sandor Fekete for interesting discussions and David Thompson for providing a stimulating research environment.

References

[1] D. Avis and G. Toussaint. An Optimal Algorithm for Determining the visibility of a Polygon from an Edge. *IEEE Transactions on Computers,* C-30-12: 910-914, 1981.

[2] John Bown. *Injection Moulding of Plastic Components.* McGraw-Hill, England, 1979.

[3] H. ElGindy and D. Avis. A Linear Algorithm for Computing the visiblity Polygon from a point. *Journal of Algorithms,* 2:186-197, 1981.

[4] S.D. ElWakil. *Processes and design for manufacturing,* Prentice Hall, New Jersey, 1989.

[5] S.P. Fekete and J.S.B. Mitchell. Geometric Aspects of Injection Molding, State University of New York, Stony Brook, manuscript 1992.

[6] A. Horn and F.A. Valentine. Some Properties of L-Sets in the plane, *Duke Mathematics Journal*, 16:131-140, 1949

[7] D.T. Lee and F.P. Preparata. An optimal algorithm for finding the kernel of a polygon. *SIAM Journal of Computing*, 6-3: 594-606, 1979.

[8] C.C. Lee and D.T. Lee. On a Circle-Cover Minimization Problem. *Information Processing Letters*, 18:109-115, 1984.

[9] T. Lyman (editor). *Casting Design Handbook*, American Society for Metals, Ohio, 1962.

[10] N. Megiddo. Linear time algorithms for linear programming in R^3 and related problems. *SIAM Journal of Computing*, 12-4: 759-776, 1983.

[11] G.H. Meisters. Polygons have ears. *Amer. Math. Monthly*, 82: 648-651, 1975.

[12] F.P. Preparata and K.J. Supowit. Testing a simple polygon for monotonicity. *Information Processing Letters*, 12-4: 161-164, 1981.

[13] W. I. Pribble. Molds for Reaction Injection, Structural Foam and Expandable Styrene Molding. *Plastics Mold Engineering Handbook, Fourth Edition.* J. Harry DuBois and Wayne I. Pribble (eds.), Van Nostrand Reinhold Company Inc., New York, 1987.

[14] A. Rosenbloom and D. Rappaport. Moldable and Castable Polygons. *Proc. of the Fourth Canadian Conference on Computational Geometry*, St. John's, Newfoundland, 1992.

[15] J.R. Sack and S. Suri. An optimal algorithm for detecting weak visibility of a polygon. *IEEE Transactions on Computers*, 39, 1990.

[16] B.L. Simpson. *History of the metal casting industry*, American Foundrymen's Society Inc., Il., 1969.