Draining a Polygon –or– Rolling a Ball out of a Polygon

Greg Aloupis* Jean Cardinal* Sébastien Collette†* Ferran Hurtado‡ Stefan Langerman§* Joseph O'Rourke¶

Abstract

We introduce the problem of draining water (or balls representing water drops) out of a punctured polygon (or a polyhedron) by rotating the shape. For 2D polygons, we obtain combinatorial bounds on the number of holes needed, both for arbitrary polygons and for special classes of polygons. We detail an $O(n^2 \log n)$ algorithm that finds the minimum number of holes needed for a given polygon, and argue that the complexity remains polynomial for polyhedra in 3D. We make a start at characterizing the *1-drainable* shapes, those that only need one hole.

1 Introduction

Imagine a closed polyhedral container P partially filled with water. How many surface point-holes are needed to entirely drain it under the action of gentle rotations of P? It may seem that one hole suffices, but we will show that in fact sometimes $\Omega(n)$ holes are needed for a polyhedron of n vertices. Our focus is on variants of this problem in 2D, with a brief foray in Sec. 5 into 3D. We address the relationship between our problem and injection-filling of polyhedral molds [BvKT98] in Sec. 4.

A second physical model aids the intuition. Let P be a 2D polygon containing a single small ball. Again the question is: How many holes are needed to ensure that the ball, regardless of its initial placement, will escape to the exterior under gentle rotation of P? Here the ball is akin to a single drop of water. We will favor the ball analogy, without forgetting the water analogy.

Models. We consider two models, the (gentle) *Rotation* and the *Tilt* models. In the first, P lies in a vertical xy-plane, and gravity points in the $-y$ direction. The ball B sits initially at some convex vertex v_i; vertices are labeled counterclockwise (ccw). Let us assume that v_i is a local minimum with respect to y, i.e., both v_{i-1} and v_{i+1} are above v_i. Now we are permitted to rotate P in the vertical plane (or equivalently, alter the gravity vector). In the Rotation model, B does not move from v_i until one of the two adjacent edges, say $e_i = v_i v_{i+1}$, turns infinitesimally beyond the horizontal, at which time B rolls down e_i and falls under the influence of gravity until it settles at some other convex vertex v_j. For example, in Fig. 1, B at v_4 rolls ccw

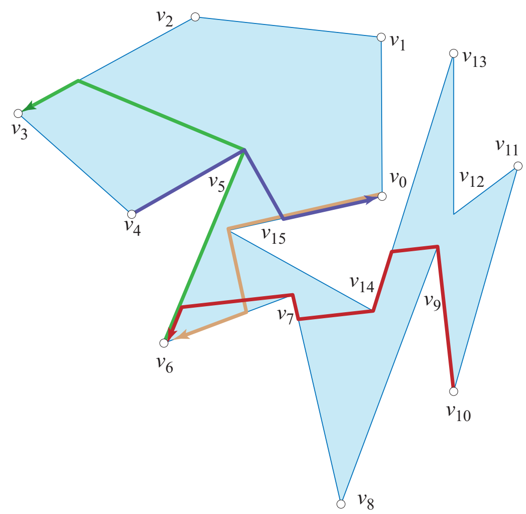

Figure 1: Polygon with several ball paths.

when v_4v_5 is horizontal, falls to edge $v_{15}v_0$, and comes to rest at v_0. Similarly, B at v_{10} rolls clockwise (cw) to v_6 after three falls. Note that all falls are parallel, and (arbitrarily close to) orthogonal to the initiating edge (in the Rotation model). After B falls to an edge, it rolls to the endpoint on the obtuse side of its fall path.

The only difference in the Tilt model is that any gravity vector may be selected. Only vectors between e_{i-1} and e_i will initiate a departure of B from v_i, i.e., the entire wedge is available rather than just the two incident edges. For example, in Fig. 1, B at v_4 rolls to $\{v_0, v_3\}$ in the Rotation model, but can roll to $\{v_0, v_1, v_2, v_3\}$ in the Tilt model. The Rotation model more accurately represents physical reality, for rain drops or for balls. The Tilt model mimics various ball-rolling games

* [greg.aloupis,jcardin,secollet,slanger]@ulb.ac.be Université Libre de Bruxelles (ULB), CP212, Bld. du Triomphe, 1050 Brussels, Belgium.

† Chargé de Recherches du FRS-FNRS.

‡ ferran.hurtado@upc.edu Universitat Politècnica de Catalunya, Jordi Girona 1–3, E-08034 Barcelona, Spain.

§ Chercheur Qualifié du FRS-FNRS.

¶ orourke@cs.smith.edu Smith College, Northampton, MA 01063, USA.

(e.g., Labyrinth) that permit quickly "tilting" the polygon/maze from the horizontal so that any departure vector from v_i can be achieved. We emphasize that, aside from this departure difference, the models are identical. In particular, inertia is ignored, and rotation while the ball is "in-flight" is forbidden (otherwise we could direct B along any path).

There are two "degenerate" situations that can occur. If B falls exactly orthogonal to an edge e, we arbitrarily say it rolls to the cw endpoint of e. If B falls directly on a vertex, both of whose edges angle down with respect to gravity, we stipulate that it rolls to the cw side.

Questions. Given P, what is the minimum number of point-holes needed to guarantee that any ball, regardless of starting position, may eventually escape from P under some sequence of rotations/tilts? Our main result is that this number can be determined in $O(n^2 \log n)$ time. In terms of combinatorial bounds, we show that some polygons require $\lfloor n/6 \rfloor$ and $\lfloor n/7 \rfloor$ holes (in the Rotation/Tilt models respectively), but $\lceil n/4 \rceil$ holes always suffice. We make a start at characterizing the *1-drainable* polygons, those that only need one hole Finally we argue that the minimum number of holes can be computed for a 3D polyhedron in polynomial time. (Omitted proofs may be found in the full version.)

2 Traps

We start by exhibiting polygons that need $\Omega(n)$ holes to drain. The basic idea is shown in Fig. 2(a) for the Rotation model. We create *traps* with 5 vertices forming

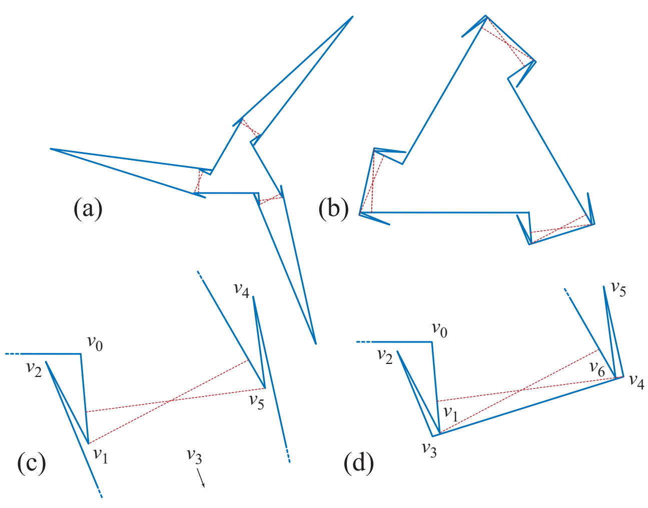

Figure 2: (a) Trap for Rotation model. (b) Trap for Tilt model. (c,d) Details of traps.

an "arrow" shape, connected together around a convex polygonal core so that 6 vertices are needed per trap. A ball in v_4 rolls to fall on edge v_0v_1, but because of the slightly obtuse angle of incidence, rolls to v_2; and symmetrically, v_2 leads to v_4. So there is a cycle (defined precisely in Sec. 3) that "traps" ball between $\{v_2, v_3, v_4\}$ and isolates it from the other two traps. Therefore three holes are required to drain this polygon. In the Tilt model, Fig. 2(a) only needs one hole, because B could roll directly from v_3 through the $v_1 - v_5$ "gap." However, the polygon in Fig. 2(b) requires 3 holes. Here the range of effective gravity tilt vectors from v_4 is so narrow that the previous analysis holds. These examples establish the necessity half of this theorem:

Theorem 1 (Combinatorial Bounds) *In the Rotation (resp. Tilt) model, $\lfloor n/6 \rfloor$ (resp. $\lfloor n/7 \rfloor$) holes are sometimes necessary to drain an n-vertex polygon. $\lceil n/4 \rceil$ holes suffice to drain any polygon.*

Although we believe that at least two reflex vertices are needed in every cycle, we were unable to show that they could not be shared between traps. We nevertheless conjecture that $\lceil n/5 \rceil$ holes suffice. (A variation on Fig. 2(a) permits the formation of two traps with $n = 11$.)

Proposition 2 *$\lfloor n/28 \rfloor$ holes are sometimes necessary to drain an n-vertex orthogonal polygon, and $\lceil n/8 \rceil$ holes suffice.*

3 The Pin-Ball Graph

Let G be a directed graph whose nodes are the convex vertices of P, with v_i connected to v_j if B can roll in one "move" from v_i to v_j. Here, a move is a complete path to the local y-minimum v_j, for some fixed orientation of P. We conceptually label the arcs of G with the sequence of vertices and edges along the path $\rho(v_i, v_j)$. Thus, the (v_{10}, v_6) arc in Fig. 1 is labeled $(v_9, e_{13}, v_{14}, e_7, v_7, e_5)$. We use G_R and G_T to distinguish the graphs for the Rotation and Tilt models respectively, and G when the distinction is irrelevant.

We gather together a number of basic properties of G in the following lemma.

Lemma 3 (G Properties)

1. *Every node of G_R has out-degree 2; a node of G_T has out-degree at least 2 and at most $O(n)$.*
2. *Both G_R and G_T have $O(n)$ nodes (one per convex vertex). G_R has at most $2n$ arcs, while G_T has $O(n^2)$ arcs, and sometimes $\Omega(n^2)$ arcs.*
3. *Each path label has length $O(n)$ (in either model).*
4. *The total number of path labels on the arcs of G_R is $O(n^2)$, and sometimes $\Omega(n^2)$.*
5. *The total number of labels in G_T is $O(n^3)$, and sometimes $\Omega(n^3)$.*

We will see below that G_T can be constructed more efficiently than what the cubic total label size in Lemma 3(5) might indicate.

Noncrossing Paths. A ball path corresponding to one arc of G is a polygonal curve, monotone with respect to gravity $\overrightarrow{g}$. The path is composed of subsegments of polygon edges, as well as *fall* segments, each of which is parallel to $\overrightarrow{g}$ and incident to a reflex vertex. A directed path ρ naturally divides P into a "left half" $L = L(\rho)$ of points left of the traveling direction, and a "right half" $R = R(\rho)$, where L and R are disjoint, and $L \cup R \cup \rho = P$. Two ball paths ρ_1 and ρ_2 (properly) *cross* if ρ_2 contains points in both $L(\rho_1)$ and $R(\rho_1)$. For example, in Fig. 1, $\rho(v_0, v_6)$ crosses $\rho(v_{10}, v_6)$. Let $\overline{L} = L(\rho) \cup \rho$ be the closure of $L(\rho)$, and similarly define $\overline{R}$.

Two paths can only cross at a reflex vertex (as do $\rho(v_4, v_0)$ and $\rho(v_6, v_3)$ in Fig. 1) or on fall segments of each (as do $\rho(v_0, v_6)$ and $\rho(v_{10}, v_6)$).

Lemma 4 (Noncrossing) *Two paths ρ_1 and ρ_2 from the same source vertex v_0 never properly cross (in either model).* See Fig. 3.

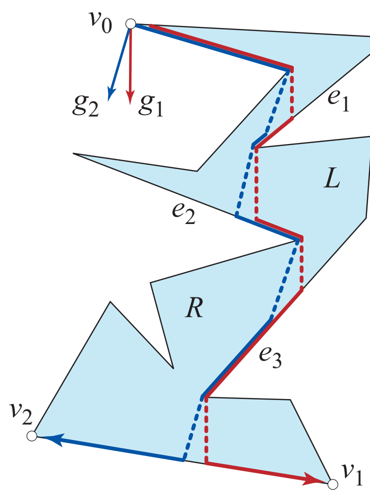

Figure 3: Paths from the same source v_0 do not cross. Fall segments are dashed. L and R indicate polygon "halves" left and right of the directed paths.

Lemma 5 (Label Intervals) *In the Tilt model, a particular label λ appears on the arcs of G_T originating at one particular v_i within an interval $[\overrightarrow{g_1}, \overrightarrow{g_2}]$ of gravity directions.*

Cycles and Strongly Connected Components. The directed path from any vertex v_i of G leads to a cycle in G, because every node has at least two outgoing edges by Lemma 3(1). Any maximal cycle in G has length at least 3. Anything less would involve a pair (v_i, v_j) connecting only to each other, which would contradict Lemma 3(1). Note that any pair of convex vertices adjacent on ∂P form a non-maximal cycle of length 2.

A cycle is a particular instance of a *strongly connected component* (SCC) of G, a maximal subset $C \subset G$ in which each node has a directed path to all others.

Define a graph G^* as follows. Let $C_1, C_2, \ldots$ be the SCC's of G. Contract each C_k to a node c_k of G^*, while otherwise maintaining the connectivity of G. Then G^* is a DAG (because all cycles have been contracted).

Lemma 6 (Sinks) *The minimum number m of holes needed to drain P is the number of sinks of G^*.*

Lemma 7 *The locations of the minimum number m of holes needed to drain P can be found in linear time in the size $|G|$ of G, once G has been constructed.*

Construction of G. Our goal is to construct the unlabeled G. Labels merely represent the paths that realize each arc of G. An example given in the long version seems to require $\Omega(n^2)$ ray-shooting queries in the Rotation model, and as we do not know how to avoid this, our goal becomes an $O(n^2 \log n)$ algorithm. This is straightforward for G_R, so we focus on G_T, which by Lemma 3(5) is potentially cubic.

We first preprocess P for efficient ray-shooting queries, using fractional cascading to support ray shooting in a polygonal chain. This takes $O(n \log n)$ preprocessing time and supports $O(\log n)$ time per query ray [CEG+94]. Next we construct the visibility polygon from each vertex of in overall $O(n^2)$ time [JS87]. From these visibility polygons, for each v_i we construct a *gravity diagram D_i*. This partitions all gravity vectors $\overrightarrow{g}$ into angular intervals labeled with the next vertex that B will roll to from v_i with tilt $\overrightarrow{g}$. For example, Fig. 4(a) shows the gravity diagram for v_4 in Fig. 1. Note that

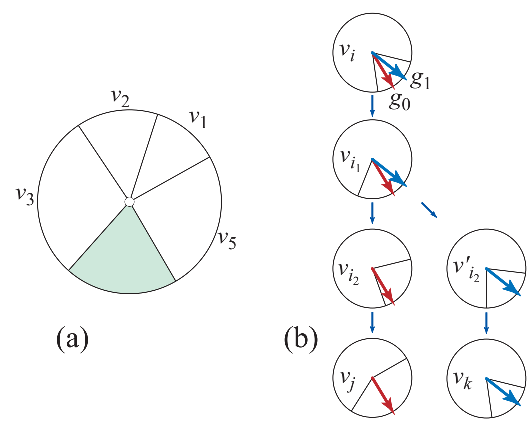

Figure 4: (a) Gravity diagram for v_4 in Fig. 1. (b) Gravity diagrams for paths ρ_0 and ρ_1.

D_i only records the next vertex encountered, not the ultimate destination. We maintain each diagram in a structure that permits any $\overrightarrow{g}$ to be located in $O(\log n)$ time.

We now argue that we can construct all paths with source v_i, and therefore all arcs of G_T leaving v_i, in

$O(n \log n)$ time. Compute the path ρ_0 for the cw extreme gravity vector $\overrightarrow{g_0}$ that leaves v_i (perpendicular to $v_i v_{i+1}$). This uses $O(n)$ ray-shooting queries for the fall segments of ρ_0, totaling $O(n \log n)$ time. Let $\rho_0 = (v_i, v_{i_1}, v_{i_2}, \ldots, v_j)$. During its construction, we locate $\overrightarrow{g_0}$ within each diagram D_{i_k}. Now we find the minimum angle between $\overrightarrow{g_0}$ and the next ccw event over all diagrams. This can be done in $O(\log n)$ time using a priority queue. Call the next event $\overrightarrow{g_1}$, and suppose it occurs at v_{i_k} in diagram D_{i_k}. We now construct the path ρ_1 from v_{i_k} onward, until it terminates at a new vertex, or rejoins ρ_0 (recall from Fig. 3 that paths might rejoin, i.e., the suffixes from v_{i_k} are not necessarily disjoint). In our "update" from ρ_0 to ρ_1, let V_0 be the set of vertices lost from ρ_0, and V_1 those gained in ρ_1. The priority queue of minima is updated by deleting those for V_0 and inserting those for V_1. The angular sweep about v_i continues in the same manner until the full gravity vector range about v_i is exhausted. Fig. 4(b) illustrates one step of this process, where v_{i_1} determines the transition event between g_0 and g_1, at which point the path changes from $\rho_0 = (v_i, v_{i_1}, v_{i_2}, v_j)$ to $\rho_1 = (v_i, v_{i_1}, v'_{i_2}, v_k)$.

By Lemma 5, each diagram abandoned in this sweep is never revisited. Thus the number of invocations of the minimum operation to find the next event is $O(n)$, or $O(n \log n)$ overall. Repeating for each v_i we obtain:

Lemma 8 *G can be constructed in $O(n^2 \log n)$ time.*

Theorem 9 *The locations of the minimum number of holes needed to drain P can be found in $O(n^2 \log n)$.*

We leave it as a claim that G_R can be constructed (and the holes located) in $O(n \log n)$ time for orthogonal polygons.

4 1-Drainable Shapes

Define a *k-drainable* polygon as one that can be drained with k holes but not with $k-1$ holes. For example, Fig. 1 is 1-drainable with a hole at v_6. We make a start here at exploring the 1-drainable shapes under each model. Note that these shapes do depend on the model: Fig. 2(a) is 1-drainable in the Tilt model but 3-drainable in the Rotation model.

Our definition of k-drainable polygons is inspired by the k-fillable polygons of [BvKT98], those mold shapes that can be filled with liquid metal poured into k holes. Despite the apparent inverse relationship between filling and draining, the two concepts are rather different. In particular, there are star-shaped polygons k-drainable in the rotation model (Proposition 13 below), but Theorem 7.2 of [BvKT98] shows that these are all "2-fillable with re-orientation." Also, there are 1-drainable polygons that are k-fillable (with or without reorientation).

Proposition 10 *Monotone polygons are 1-drainable.*

Let the *ccw roll* from v_i be the roll toward v_{i+1} in the Rotation model, or equivalently, the tilt according to $\overrightarrow{g}$ perpendicular to $v_i v_{i+1}$ in the Tilt model.

Lemma 11 (Kernel) *Let P be star-shaped with kernel K. Then for each arc $(v_i, v_j) \in G$ corresponding to the ccw roll path ρ from v_i, K is in $\overline{L(\rho)}$, i.e., K is on or to the left of ρ.*

A *fan* is a star-shaped polygon whose kernel includes a convex vertex.

Proposition 12 *Fans are 1-drainable.*

Proposition 12 cannot be extended to star-shaped polygons in the Rotation model:

Proposition 13 *For any $k > 1$, there is a k-drainable star-shaped n-gon in the Rotation model, with $k = \Omega(n)$.*

5 3D

Define the Tilt model for 3D polyhedra to permit departure from a vertex v at a direction vector lying in any of the faces of P incident to v. We do not see how to mimic the efficient construction of G previously described, so we content ourselves with showing (in the full version) that it can be accomplished in polynomial time: $O(n^7 \log n)$.

6 Open Problems

1. Can the upper bound of $\lceil n/4 \rceil$ in Theorem 1 be improved?
2. Are star-shaped polygons 1-drainable in the Tilt model? More generally, characterize 1-drainable polygons.
3. Suppose m balls are present in P at the start, and P is k-drainable. What is the computational complexity of finding an optimal schedule of rotations, say, in terms of the total absolute angle turn, or in terms of the number of angular reversals?

References

[BvKT98] P. Bose, M. van Kreveld, and Godfried T. Toussaint. Filling polyhedral molds. *Comput. Aided Design*, 30(4):245–254, April 1998.

[CEG+94] Bernard Chazelle, Herbert Edelsbrunner, Michelangelo Grigni, Leonidas J. Guibas, John Hershberger, Micha Sharir, and Jack Snoeyink. Ray shooting in polygons using geodesic triangulations. *Algorithmica*, 12:54–68, 1994.

[CLRS01] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 2nd edition, 2001.

[JS87] B. Joe and R. B. Simpson. Correction to Lee's visibility polygon algorithm. *BIT*, 27:458–473, 1987.