

Unfolding Convex Polyhedra via Radially Monotone Cut Trees

Joseph O’Rourke*

July 26, 2016

Abstract

A notion of “radially monotone” cut paths is introduced as an effective choice for finding a non-overlapping edge-unfolding of a convex polyhedron. These paths have the property that the two sides of the cut avoid overlap locally as the cut is infinitesimally opened by the curvature at the vertices along the path. It is shown that a class of planar, triangulated convex domains always have a radially monotone spanning forest, a forest that can be found by an essentially greedy algorithm. This algorithm can be mimicked in 3D and applied to polyhedra inscribed in a sphere. Although the algorithm does not provably find a radially monotone cut tree, it in fact does find such a tree with high frequency, and after cutting unfolds without overlap. This performance of a greedy algorithm leads to the conjecture that spherical polyhedra always have a radially monotone cut tree and unfold without overlap.

1 Introduction

The question of whether or not every convex polyhedron in \mathbb{R}^3 has a non-overlapping edge-unfolding to the plane has been open for many years; see [DO07]. Here I introduce a notion of “radially monotone” cut paths, and show empirically that for polyhedra inscribed in a sphere (which I’ll call *spherical*¹), there exists with high frequency a spanning tree composed of radially monotone paths, and the corresponding unfolding avoids overlap. A 200-vertex example is shown in Fig. 1.² This “high frequency” claim contrasts with the near certainty of overlap for random spanning cut trees of random spherical polyhedra, as observed long ago: see Fig. 2, and Fig. 3 for an example of overlap.

Unfortunately, the “almost always” claim is not a theorem, but rather a conjecture supported by data. For example, one empirical exploration run found a radially monotone cut tree for 1,000 random spherical non-obtusely triangulated

*Department of Computer Science, Smith College, Northampton, MA 01063, USA.
orourke@cs.smith.edu.

¹“Inscribed” is often used in the literature.

² The quality of the figures had to be reduced for the arXiv.




Figure 1: Left: Polyhedron of 200 vertices inscribed in sphere, with radial monotone cut tree shown. (The role of the green edges will be explained later.) Right: Unfolding.

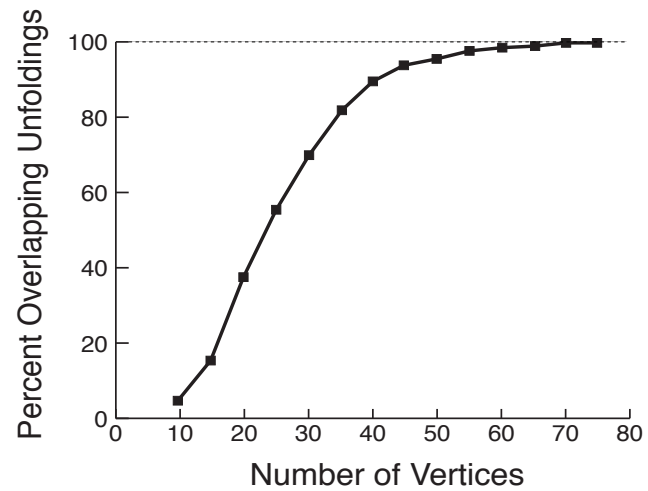


Figure 2: Each point represents 5,000 unfoldings of spherical polyhedra. Fig. 22.10 in [DO07, p.315].




Figure 3: Unfolding of a 50-vertex spherical polyhedron with several overlaps.

polyhedra, each of 100 vertices, and unfolded all 1,000 without overlap. (This claim will be hedged a bit in Section 5.3.)

What I have been able to prove is confined to 2D: Theorem 1 says that “round” convex polygons, meshed with non-obtuse triangles, always possess a radially monotone spanning forest. The 3D empirical results use an algorithm that mimics the 2D proof, and relies on a generalization of radial monotonicity to 3D governed by Theorem 2.

1.1 Definition of Radial Monotonicity in \mathbb{R}^2

We now define radial monotonicity in 2D; we will not return to 3D until Section 4.

1.2 Definition

Let C be a triangulated convex domain in \mathbb{R}^2 , with ∂C its boundary, a convex polygon. In general C contains many points in its interior that are the vertices of the triangulation. Let $Q = (v_0, v_1, v_2, \dots, v_k)$ be a simple (non-self-intersecting) directed path of edges of C connecting an interior vertex v_0 to a boundary vertex $v_k \in \partial C$.

We say that $Q = (v_0, v_1, \dots, v_k)$ is *radially monotone (rm)*³ w.r.t. v_0 if the distances from v_0 to all points of Q are (non-strictly) monotonically increasing.

³ I will also use “rm” to abbreviate “radially monotone” when convenient and unambiguous. Also, “w.r.t.” is an abbreviation for “with respect to.”

(Note that requiring the distance to just the vertices of Q to be monotonically increasing is not the same as requiring the distance to all points of Q be monotonically increasing.) We define path Q to be *radially monotone* (without qualification) if it is radially monotone w.r.t. each of its vertices: v_0, v_1, \dots, v_{k-1} . Before exploring this definition further, we discuss its intuitive motivation.

1.3 Motivating Intuition

A radial monotone path Q w.r.t. v_0 has the property that rigidly rotating Q and all its incident triangles about v_0 by a small angle avoids proper overlap between the triangles to the left and to the right of Q . One can imagine one triangle incident to v_0 reducing its angle at v_0 infinitesimally, as illustrated in Fig. 4(a). Fig. 4(b) shows that the path illustrated is not radially monotone




Figure 4: Path $Q = (v_0, v_1, \dots, v_6)$ is radially monotone w.r.t. v_0 , but not w.r.t. v_1 .

without qualification, because it is not w.r.t. v_1 .

The motivation behind the definition of radial monotonicity is as follows. Ultimately the path Q will be a path of edges on a convex polyhedron \mathcal{P} in \mathbb{R}^3 . At each vertex $v_i \in \mathcal{P}$, there will be some positive curvature $\omega_i > 0$, which represents the “angle gap” when the neighborhood of v_i is flattened to the plane. We can view ω_i as separating the left-half of the cut Q from the right-half at v_i .

As a consequence, if a path Q is rm, then “opening” the path with sufficiently small curvatures ω_i at each v_i will avoid overlap between the two halves of the cut path. Whereas if a path is not rm, then there is some opening curvature assignments ω_i to the v_i that would cause overlap: assign a positive curvature $\omega_j > 0$ to the first vertex v_j at which radial monotonicity is violated, and assign the other vertices zero or negligible curvatures. Thus rm cut paths are

locally (infinitesimally) “safe,” and non-rm paths are potentially overlapping. This potential overlap may not be realizable, for the ω_i cannot be assigned arbitrarily, but must derive from a convex polyhedron. So the guarantee is one-way:

$$\boxed{\text{Radial monotonicity} \implies \text{safe (non-overlapping) infinitesimal opening.}}$$

1.4 Trees and Spanning Forests

Continuing to concentrate on a planar triangulated convex domain C , we extend the notion of radial monotonicity to trees and forests in the natural manner. A tree T rooted on a boundary vertex $v_k \in \partial C$ and containing no other boundary vertex is radially monotone if the path from every leaf of T to v_k is rm. A radially monotone spanning forest for C is a collection of boundary-rooted rm trees that span the interior vertices of C .

1.5 Spherical Caps

To presage why we concentrate on convex domains, we turn briefly again to 3D. Fig. 5 shows a convex cap composed of the faces of a spherical polyhedron whose normals are within 60° of the vertical. We will unfold convex caps via radially monotone cut forests. And ultimately we will show how to view a complete spherical polyhedron as a convex cap.




Figure 5: Left: Spherical cap of 147 vertices, with radial monotone cut forest marked. Right: Unfolding. Green edges are part of ∂C .

2 Properties of Radially Monotone Paths

We now embark on a rather lengthy description of properties of rm paths in 2D.

2.1 Radial Circles

The condition for Q to be rm w.r.t. v_0 can be interpreted as requiring Q to cross every circle centered on v_0 at most once; see Fig. 6. The concentric circles viewpoint makes it evident that infinitesimal rigid rotation of Q about v_0 to Q' ensures that $Q \cap Q' = \emptyset$.

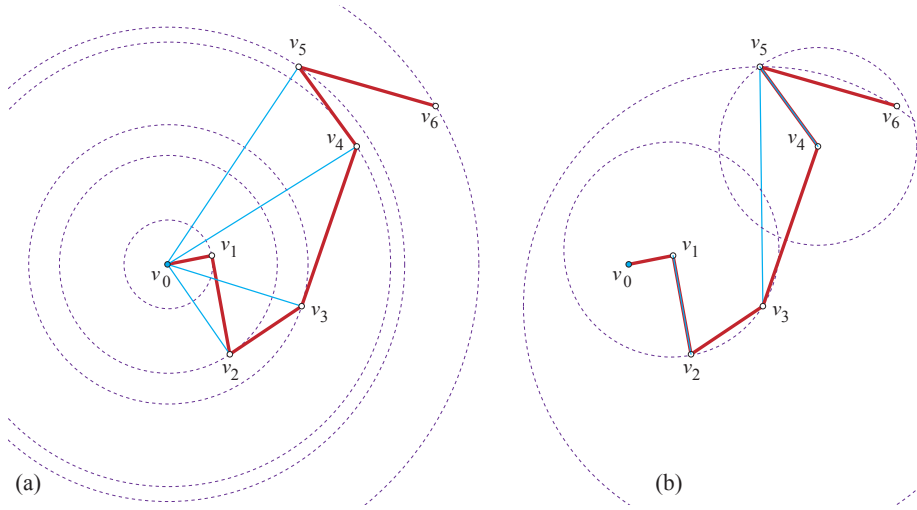


Figure 6: (a) A chain radially monotone w.r.t. v_0 . (b) The chain is not radially monotone w.r.t. v_1 (violation at v_2), or w.r.t. v_3 (violation at v_5), or w.r.t. v_4 (violation at v_5).

An equivalent definition is as follows. Let $\alpha_j(v_i) = \angle(x, v_j, v_{j+1})$. Then Q is rm w.r.t. v_i if $\alpha_j(v_i) \geq \pi/2$ for all $j > i$. For if $\alpha_j(v_i) < \pi/2$, Q violates monotonicity at v_j , and if $\alpha_j(v_i) \geq \pi/2$, then points along the segment (v_j, v_{j+1}) increase in distance from v_i .

2.2 Non-Properties of rm Paths

Let the *turn angle* τ_i of path Q at v_i be the signed angle between the vectors $v_i - v_{i-1}$ and $v_{i+1} - v_i$; so $\tau_i \in [-\pi, \pi]$, with $\tau_i = 0$ meaning that the joint at v_i is straightened.

It should be clear that no turn angle in a radially monotone path can exceed $\pi/2$: $\tau_i \leq \pi/2$ for all $i = 1, \dots, k-1$. Although this condition is necessary, it is not sufficient: if $Q = (v_0, v_1, v_2, v_3)$ has unit link lengths, and turn angles $\tau_1 = \tau_2 = \pi/2$, forming a \square -shape, then Q is not rm w.r.t. v_0 , violating rm at v_2 .

Note the definition of radial monotonicity considers paths directed from v_0 to v_k . A path radially monotone in one direction need not be rm when reversed: see Fig. 7.




Figure 7: $Q = (v_0, v_1, v_2, v_3)$ is radially monotone, but its reverse, (v_3, v_2, v_1, v_0) , is not rm. Shortcutting the former rm path to (v_0, v_1, v_3) destroys radial monotonicity.

Another property one might hope holds is that shortcutting a rm path $Q = (v_0, \dots, v_{i-1}, v_i, v_{i+1}, \dots, v_k)$ to $Q' = (v_0, \dots, v_{i-1}, v_{i+1}, \dots, v_k)$ would retain radial monotonicity, but this is false in general, as also illustrated in Fig. 7.

2.3 Searching for rm Paths

We will search for rm paths to form a rm cut forest to unfold a spherical cap. The search will be incremental, growing existing paths. So we next turn to properties that will allow a partial rm path to be extended, on either end. Again we concentrate on a planar convex domain C .

There are two basic strategies: Start from an interior vertex of C , and grow toward ∂C ; or start from a boundary vertex, and grow toward the interior. The former is attractive because, as Lemma 1 will show, a partial path can always be grown “forward.” Meanwhile, “backward” growth is not always possible, as will be shown in Section 2.5. However, so far I have only been able to prove that a rm spanning forest exists by growing from the boundary inward.

2.4 Forward Extension: O-cone Θ_j

For path $Q = (v_0, v_1, \dots, v_k)$, let Θ_j be the range of angles within which $(v_{j+1} - v_j)$ must lie for Q to be rm. We call Θ_j the *o-cone* at v_j , because its bounding rays are orthogonal to the *cone* at v_j : the smallest cone with apex at v_j within which all of Q lies. See Fig. 8.

For segment (v_j, v_{j+1}) to be rm w.r.t. v_i , that segment must fall outside the circle centered on v_i of radius $|v_j - v_i|$. Thus each v_i contributes a halfplane constraint at v_j , and it is the intersection of those halfplanes that determine the o-cone Θ_j . And clearly the extremes of the o-cone are orthogonal to the extremes of the cone.




Figure 8: The o-cones Θ_j of a radially monotone path. The cone at v_6 (illustrated) passes through v_3 and v_4 .

Let $|\Theta_j|$ be the measure of the o-cone angle. The following lemma is not used in the sequel, and is included just for intuition.

Lemma 1 *For any radially monotone path $Q = (v_0, \dots, v_k)$, $|\Theta_j| \in [\pi/2, \pi]$ for all j .*

Proof: Initially $|\Theta_1| = \pi$, and assume the claimed bound holds up to v_k . Now consider extending Q with v_{k+1} . The cone for Q encompasses the convex hull of Q , as illustrated in Fig. 9. Let α_k be the angle of the cone. Then the angle of the o-cone is $\pi - \alpha_k$. If v_{k+1} is close to v_k on the upper or lower boundary of Θ_k , then $|\Theta_{k+1}|$ approaches $\pi/2$ from above as the distance $|v_k - v_{k+1}| \rightarrow 0$. If v_{k+1} lies on the line containing (v_{k-1}, v_k) , and so extends the previous link collinearly, then $|\Theta_{k+1}|$ approaches $|\Theta_k|$ as $|v_k - v_{k+1}| \rightarrow 0$. There is a continuous variation in $|\Theta_{k+1}|$ between these extremes for v_{k+1} arbitrarily close to v_k .

Now consider moving v_{k+1} from near v_k along a line that falls within $|\Theta_k|$. The cone apex angle α_{k+1} narrows monotonically, and so $|\Theta_{k+1}|$ monotonically increases. So $|\Theta_{k+1}|$ is always $\geq \pi/2$. And clearly $\alpha_{k+1} \geq 0$ and so $|\Theta_{k+1}| \leq \pi$. \square

2.5 Backward Growth

Consider now a rm path $Q = (v_0, v_1, v_2, \dots, v_k)$ that we would like to extend “backwards,” with v' prior to v_0 , while retaining rm for $Q' = (v', v_0, v_1, \dots, v_k)$. Define $R = R(Q)$ to be the region of the plane within which v' can lie while retaining Q' rm. Each edge (v_i, v_{i+1}) of Q contributes a halfplane constraint to R : v' must lie in the halfplane whose boundary is orthogonal to (v_i, v_{i+1}) , passes through v_i , and excludes v_{i+1} . For otherwise, $\angle v', v_i, v_{i+1} < \pi/2$, and




Figure 9: $\Theta_{n+1} = \pi - \alpha_{n+1}$ varies from $\pi/2$ near v_n , to approaching π distant from v_n .

rm w.r.t. v' would be violated at v_i . R is then the intersection of all these halfplanes; see Fig. 10.




Figure 10: R is determined here by the halfplanes generated by (v_0, v_1) , (v_1, v_2) , and (v_3, v_4) .

It is possible that Q cannot be extended backwards at all: R could equal $\{v_0\}$, as illustrated in Fig. 11.

2.6 Logarithmic Spirals

Although there is a sense in which radially monotone paths are somewhat “straight,” in that they increase in distance from v_0 and never turn too sharply, the example in Fig. 11 shows that intuition cannot be pushed too far. In this section we identify the extreme rm path, which turns out to be an approximately 75° logarithmic spiral.




Figure 11: R is the small triangle with one corner v_0 .

A logarithmic spiral can be expressed in polar coordinates by the equation $r = ae^{b\theta}$. For our purposes we can take the scale factor a to be 1. $b = 1/\tan \phi$, where ϕ is the constant angle between the radial vector to a point p on the curve, and the tangent of the curve at p . This constant angle is the determining characteristic of such spirals. Fig. 12 shows a typical spiral. When $\phi = \pi/2$, the spiral becomes a circle. As $\theta \rightarrow -\infty$, $r \rightarrow 0$ and the spiral approaches the origin.

We can extend the definition of radial monotonicity to apply to smooth curves: a directed curve is radially monotone if it is rm w.r.t. to every point p on the curve, i.e., the distance from p to points beyond p on the curve (non-strictly) increases monotonically. It turns out that the spiral in Fig. 12 is not radially monotone, but a numerical calculation shows that log spirals with $\phi \leq 74.655^\circ$ are radially monotone. Fig. 13 shows the extreme spiral, what I will call a 75° -spiral.

3 Radially Monotone Forests Exist for Non-Obtuse Triangulations

In this section we derive our main 2D result, which is the inspiration for the 3D algorithm to follow.

3.1 Setting

Again we concentrate on be a triangulated domain C whose boundary ∂C is a convex polygon. Later will impose some conditions on the shape of C and on its




Figure 12: A logarithmic spiral with $\phi = 80^\circ$. Here θ is plotted out to 2π .




Figure 13: 74.655° spiral. The tangent at any point derived from polar angle θ on the curve makes a 90° angle w.r.t. the point at polar angle $\theta - 3\pi/2$.

triangulation. Our goal is to find a radially monotone spanning forest \mathcal{F} for C . \mathcal{F} is a collection of disjoint trees, each with exactly one vertex on ∂C (its root), following edges of the triangulation, and spanning all interior vertices (those not on ∂C). Note that each interior vertex has a unique path to ∂C , because each interior vertex lies in a unique tree rooted on ∂C .

As mentioned previously, we will eventually turn to 3D when C will be a spherical cap, and cutting \mathcal{F} will permit C to be flattened into the plane as a single piece. Our goal then will be to avoid overlap in this edge-unfolding, and a rm-forest guarantees an infinitesimally safe opening.

In 2D, however, C is already flat and cutting \mathcal{F} just has the property that the dual graph of C 's triangulation—with each triangle a node and two triangles adjacent across an uncut edge—is a tree, and so a single piece. The reason that cutting \mathcal{F} leaves C in one piece is the restriction that each tree only touch the boundary at one vertex. For if a tree touched two boundary vertices, then it would form a cycle with the portion of ∂C between those vertices, disconnecting a piece from the remainder. We insist that ∂C be convex because one can rather easily force overlap in the 3D situation by designing a nonconvex boundary appropriately. (However, as is clear from Fig. 5, a weaker condition than convexity may suffice.)

Before commencing, it is important to remark that not every triangulation with a convex boundary has an rm-forest. See Appendix 1 for a counterexample. Thus we do need to seek conditions that guarantee existence.

3.2 Overview of Algorithm

Although one could imagine a search over all possible spanning forests, instead we will implement a tightly regimented greedy search. First, C is placed in a minimum enclosing circle with center x . Second, the internal vertices are processed in order of their distance to x , with those closest to ∂C first, and the closest vertex to x last. So the rm trees are grown inward toward x . Third, we grow *hourglass paths*, which are a restricted class of radially monotone paths. Before describing the details, we sketch the algorithm at a high level in the boxed pseudocode below.

The order of growth of the trees in the forest \mathcal{F} is determined by the concentric-circle sorting w.r.t. x . The algorithm is greedy in the sense that among the options (incident triangulation edges) when connecting v_0 to some further-away $v_1 \in \mathcal{F}$, the “best connection” is selected. For a path of length at least 2, best is defined as the smallest worst turnangle τ_i , the angle from the vector $v_i - v_0$ to the vector $v_{i+1} - v_i$. When $\tau_i > \pi/2$, the path is not rm. Smaller τ_i means straighter paths, so this is a natural choice.

We need to define what the turnangle means when the path is of length 1, v_0v_1 with $v_1 \in \partial C$. We use the turnangle from the vector $v_1 - v_0$ to the tangent of the circle centered at x through v_1 . The reason for this choice (which ignores the orientation of ∂C at v_1) will be made clear below.

Note that, when a connection from v_0 to v_1 is explored for radial monotonicity, it is only necessary to check for rm w.r.t. v_0 , because by construction the

Algorithm 1: Find rm cut forest \mathcal{F} for planar C

```

Input : Convex triangulated domain  $C$ , with bounding circle center  $x$ 
Output: Radially monotone cut forest  $\mathcal{F}$ 

//  $\mathcal{F}$  is grown from  $\partial C$  inward
 $\mathcal{F} \leftarrow \emptyset$ 
Sort interior vertices by distance from center  $x$ , those nearest  $\partial C$  first.

// Grow  $\mathcal{F}$ :
foreach vertex  $v_0$  in sorted order do
    foreach vertex  $v_1$  already in  $\mathcal{F}$  or on  $\partial C$  do
        Let  $e = (v_0, v_1)$ , if  $e$  is a triangulation edge.
        Check if the path from  $v_0$  to  $\partial C$  in  $\mathcal{F} + e$  is radially monotone.
        If so, record its worst turnangle  $\tau$  (with  $\tau > 90^\circ$  not rm).
    end
    Choose the  $e^*$  that has the best (minimum)  $\tau$ .
     $\mathcal{F} \leftarrow \mathcal{F} + e^*$ 
end

```

path is already known to be rm w.r.t. v_i , $i > 0$, because it was earlier added to \mathcal{F} .

3.3 Hourglass Paths for Halfplanes

To provide intuition for hourglass paths, we first describe them for C a halfplane, bounded on the right by a vertical line $L = \partial C$. Imagine the halfplane meshed with a triangulation.

It will be occasionally more convenient to label the vertices of a path to increase from ∂C inward rather than the reverse. We will use u_i whenever indexing inward. So $Q = (v_0, v_1, \dots, v_k) = (u_k, u_{k-1}, \dots, u_0)$ with $u_0 \in \partial C$. So our rm paths will grow from u_0 inward to u_k (but the direction for radial monotonicity will still be interior-to-boundary, v_0 to v_k).

We define an *hourglass* H as a double cone bounded by two lines meeting at right angles. The *baseline* of an hourglass is the line passing through the apexes of the cones so that the bounding rays make angles of $\pi/4$ with the baseline. In our halfplane example, the baselines are all parallel to L .

We will now imagine growing a path Q from $u_0 \in L = \partial C$ inward. We place an hourglass H centered on each u_i , and call the cone of H pointing right, toward u_{i-1} and the halfplane boundary, the *out-cone*, and the cone of H pointing left, toward u_{i+1} and the interior, the *in-cone* of H . Finally, define an *hourglass path* $Q = (u_0, \dots, u_k)$ as one for which each edge $u_i u_{i+1}$ falls inside the out-cone of the hourglass at u_i and the in-cone of the hourglass at u_{i+1} . See Fig. 14. Note that, if $u_i u_{i+1}$ falls inside the out-cone of u_i , then it necessarily falls within the in-cone of u_{i+1} . We retain the definition as stated because this property relies on the parallel baselines of the hourglasses, which will not hold




Figure 14: The out-cone of the hourglass at u_4 falls inside the intersection of the orthogonal halfplanes (purple) through u_1, \dots, u_4 .

in more general situations.

Recall from Section 2.5 that we can extend Q “backwards” (inward) to $Q = (u_0, \dots, u_k, u_{k+1})$ only if the edge u_k, u_{k+1} falls within $R(Q)$, the region formed by halfplane intersections (and which could be a single point). The key point is that hourglass paths can always extend:

Lemma 2 *If $Q = (u_0, \dots, u_k)$ is an hourglass path in C a halfplane, with $u_0 \in L$, then (a) Q is a radially monotone path, and (b) $R(Q)$ contains the out-cone at u_k .*

Proof: The situation is as illustrated in Fig. 14. (a) Note that Q ’s direction for the purposes of radial monotonicity is from u_k to u_0 . Because each hourglass cone is angled 45° w.r.t. the vertical, it must be that the whole path Q falls within the out-cone at u_k . Then consider the angle $\angle u_k, u_i, u_{i-1}$ for any u_i . This angle must be at least 90° , exactly 90° if every edge of Q falls on (say) the lower cone boundaries. Since the argument holds for any pair of vertices of Q , Q is indeed rm.

(b) $R(Q)$ is the intersection of halfplanes orthogonal to $u_{i-1}u_i$ and through u_i . Falling within the hourglass in- and out-cones assures that each edge of Q makes an angle within $\pm 45^\circ$ of the horizontal, so the halfplanes comprising $R(Q)$ make angles within $\pm 45^\circ$ of the vertical. Thus their intersection includes the out-cone at u_k . \square

The import of this lemma is that an hourglass path can be extended inward with any edge in the in-cone of the hourglass at u_k , and remain radially monotone.

So far we have ignored the triangulation mesh of the halfplane C . A triangulation is called a *non-obtuse triangulation* if no triangle angle strictly exceeds $\pi/2$, i.e., is obtuse. If the halfplane C is non-obtusely triangulated, then the in-cone at u_k necessarily contains some triangulation edge (because the cone angle is 90°).

Now imagine executing Algorithm 1 on this halfplane C , except with the concentric circles replaced by the vertical lines through each vertex of the triangulation (effectively, infinite-radii concentric circles). Suppose we have grown a forest that connects every vertex strictly right of v to ∂C , and we seek to connect v to grow this forest. It should be clear that the out-cone of the hourglass at v must contain an edge of the non-obtuse triangulation, and so can indeed be connected by extending a radially monotone path.

3.4 Hourglass Paths for Convex C

Now we turn back to the less contrived situation of a convex domain C , but again non-obtusely triangulated. Every planar straight-line graph on n vertices has a conforming non-obtuse triangulation of $O(n^{2.5})$ triangles [Bis16]. We will follow Algorithm 1 just as in the halfplane case, but now the proofs are no longer straightforward. A sample result of applying the algorithm is shown in Fig. 15.




Figure 15: A 56-vertex example, with several concentric circles shown. The cut forest \mathcal{F} is marked (red).

3.5 Theorem Statement

Although I have no counterexample for arbitrary convex domains C , the proof below restricts C to those that “nicely” fit within the minimal bounding circle B . Define a convex domain C as *round* if the neighborhood of every vertex $v \in \partial C$ contains the in-cone of the hourglass H at v , where the baseline of H is tangent to the circle through v centered at the center x of B .

Theorem 1 *Let C be a non-obtusely triangulated round convex domain. Then C has an radially monotone spanning forest.*

3.6 Proof of Theorem 1

Let C be a round convex domain, with bounding circle B centered on x . We define the hourglass H at a vertex v to have baseline tangent to the circle through v centered on x . Let $Q = (u_0, \dots, u_k)$ be a path of edges in C , with $u_0 \in \partial C$. We say that Q is an *hourglass path* if each edge $u_i u_{i+1}$ falls within the in-cone at u_i and the out-cone at u_{i+1} . See Fig. 16. We need to prove three results to achieve Theorem 1:

1. An hourglass path is radially monotone: Lemma 4.
2. The in-cone at the last vertex u_k is inside the region $R(Q)$, so that Q can be extended inward: Lemma 5.




Figure 16: An hourglass path walks between concentric circles.

3. For the next vertex v to be processed, any connection within its out-cone to an earlier processed vertex v' falls within the in-cone of v' : Lemma 3.

Lemma 3 *Let v be a vertex at distance 1 from B 's center x . Then for every vertex v' within the out-cone of v at distance $r > 1$ from x , the in-cone of v' includes v .*

Proof: Refer to Fig. 17. The vertices within the out-cone of v each cover v with their in-cones. The extreme case occurs when $r \rightarrow 1$ and a vertex v' lies on the boundary of the out-cone of v . For all v' within the out-cone of v , the edge vv' lies in both the out-cone of v (by assumption) and within the in-cone of v' . \square The consequence of this lemma is that a connection to an hourglass path that terminates at v' by the edge vv' will maintain the grown path as an hourglass path.

The next two lemmas seem technically difficult, and my current arguments depend on numerical computations, but numerical arguments which I believe are convincing and could be converted to analytical proofs with effort. I will continue to call them “lemmas” with the understanding they are (justified) claims at this point.

Lemma 4 *Hourglass paths are radially monotone.*

Proof: Let Q be an hourglass path, ending at innermost vertex a , which lies on a circle of radius 1. Let b be an arbitrary point on Q . We need to prove that




Figure 17: v is covered by the in-cone of any vertex v' that falls within the out-cone of v .

the angle ab makes with the next edge bc of Q beyond (outward) of b is at most 90° . We know that edge lies in the out-cone of b , because Q is an hourglass path. The situation is illustrated in Fig. 18. From that figure, we need to show that $\beta \leq 45^\circ$ for any $r = |ab|$.

Fig. 19 plots the angle β as a function of r and α , which two parameters (together with the assumption w.l.o.g. that $|ax| = 1$) completely determine β . It is clear that in fact β can exceed 45° . However, the region of (r, α) values in which $\beta > 45^\circ$ cannot occur with an hourglass curve.

The extreme values of α occur when Q is a 45° -spiral, that is, when the path turns as much as possible clockwise within the confines of the hourglass constraints at each vertex. The maximum turn occurs in the smooth case, when the spacing between the vertices along Q approaches zero. It is easy to compute the (r, α) values achieved by this 45° -spiral. Fig. 20 shows that only at $r = 1$ and $\alpha = 45^\circ$ does the spiral values touch the boundary of the forbidden β region. Other less extreme Q avoid the forbidden region entirely. The figure also shows that other hourglass paths are less extreme than the 45° -spiral. \square

Finally we show that the in-cone of the last vertex of an hourglass curve is not “clipped,” which means it can be extended.

Lemma 5 *Let Q be an hourglass curve, ending at innermost vertex a . Then the region $R(Q)$ includes the in-cone of the hourglass at a as far as the center x .*

Proof: We assume the same setup as in the previous lemma: $|ax| = 1$. Again the issue is decided by the extreme curve, a 45° -spiral. Fig. 21 illustrates the




Figure 18: Lemma 4 is proved if $\beta \leq 45^\circ$.




Figure 19: The $\beta \geq 45^\circ$ region for all (r, α) combinations.

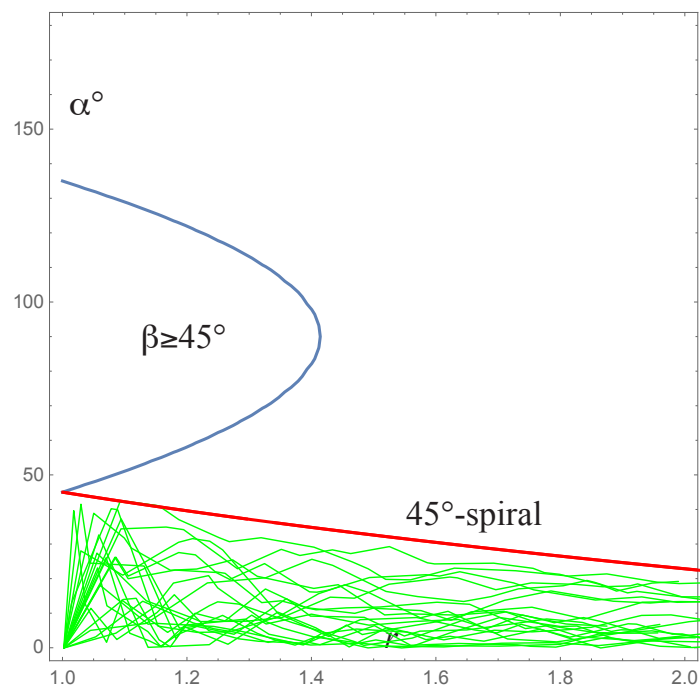


Figure 20: Hourglass paths never enter the $\beta \geq 45^\circ$ region. The green curves show (r, α) pairs for random hourglass paths.

challenge: each edge along the spiral generates a halfplane constraint that contributes to $R(Q)$. The triangle subset of the in-cone of a must not be clipped by any of these halfplanes—i.e., the halfplanes must include corner c of the triangle—or future inward growth of the hourglass path could be compromised.

For a point b at polar coordinates (r, θ) , $r = se^{\pi/2 - \theta}$, we need the angle $\beta = \angle abc \leq 45^\circ$. (Note: this β is not the same as β in the previous lemma.) The angle β “consumes” a portion of the 45° between r and the halfplane. Thus only if $\beta > 45^\circ$ does the halfplane clip c . Fig. 22 shows that $\cos \beta \geq \sqrt{2}/2$, only equaling that value when $\theta = 90^\circ$, when $a = b$. This makes sense, as the first tangent to the spiral curve at point a is at 45° w.r.t. the horizontal, and then the halfplane passes through, but does not clip off, corner c . \square

That the same degeneracy $a = b$ occurs in both of the previous two lemmas suggests there may be a uniting viewpoint.




Figure 21: Circle tangents are green. Orthogonal halfplane boundaries black. The halfplanes do not clip corner c .

Now we have proven Theorem 1 which proves that Algorithm 1 works on round, non-obtusely triangulated convex domains: the algorithm always finds a spanning forest composed of hourglass paths, which are radially monotone. The restriction to round convex domains was imposed to enable the first edge of a path starting from a vertex on ∂C to fall within the hourglass in-cone there. It is likely that this roundness assumption is not necessary. Moreover, I believe even the convex assumption is not necessary: likely only star-shapedness from




Figure 22: When $\cos \beta > \sqrt{2}/2$, $\beta < 45^\circ$.

x is needed.

The insistence on a non-obtuse triangulation, however, is crucial, as the example in Appendix 1 demonstrates.

One might wonder whether the 45° -spiral, which plays a prominent role in the analysis, can actually occur in a non-obtuse triangulation. The answer is YES: see Fig. 23.

3.7 More Examples

Several more examples are shown in Fig. 24.

4 Radial Monotone Paths in 3D

Although the whole point of radial monotone paths is to cut open polyhedra via such paths, I have yet to define what the notion means on a polyhedron in \mathbb{R}^3 . There is more than one way to generalize the notion, but I have focused on one generalization that I find natural. We assume henceforth that C is a convex cap with boundary ∂C a topological circle, i.e., C is a simply connected subset of the faces of a convex polyhedron. The earlier Fig. 5 shows such a convex cap. All of our empirical explorations also assume that the polyhedron from which C is derived is spherical: all vertices on a sphere \mathcal{S} .

Under these circumstances, we define a path Q of edges of C , $Q = (v_0, \dots, v_k)$ with $v_k \in \partial C$ to be radially monotone if the *medial path* $M(Q)$ is radially monotone. The next section explains and justifies this definition.




Figure 23: Radially monotone 45° -spiral spanning forest in a non-obtuse triangulation.




Figure 24: Radially monotone spanning forests found by Algorithm 1.

4.1 LMR-Chains

Let ω_i be the curvature at v_i . As usual, we view v_0 as a leaf of a cut forest, which will then serve as the end of a cut path, and the “source” of opening that path.

Let λ_i be the angle at v_i left of Q , and ρ_i the angle right of Q there. So $\lambda_i + \omega_i + \rho_i = 2\pi$. Define L to be the planar path from the origin with left angles λ_i , R the path with right angles ρ_i , and M the *medial path* with left angles $\lambda_i - \omega_i/2$ (and therefore right angles $\rho_i + \omega_i/2$). (Each of these paths are understood to depend on Q : $L = L(Q)$ etc.) We label the vertices of the paths ℓ_i, m_i, r_i , with $m_0 m_1$ on the x -axis. See Fig. 25.




Figure 25: The three paths L, M, R . Here M is radially monotone but L is not.

The intuition behind the focus on $M(Q)$ is two-fold: (1) It is important to allow either or both of L and R to be non-rm; (2) One can prove that $L \cap R = \{v_0\}$ for much more than infinitesimal openings caused by the curvature encountered along Q .

Let $\Omega = \sum_i \omega_i$ along Q : The total curvature of all the vertices on the cut path. Let $\tau(M)$ be the maximum \pm turn of the edge vectors in M : the largest absolute value the vector $m_i - m_{i-1}$ makes with the x -axis.

Our goal is to prove this theorem:

Theorem 2 *If path M is radially monotone, and in addition, $\tau(M) \leq \pi/2$ and $\Omega \leq \pi$, then none of the three paths L, M , and R cross one another.*

The reason for the angle restrictions is that, without some restrictions, the claim could be false. Fig. 26 shows a 70° -spiral M with $\tau(M) > (3/2)\pi$ and $\Omega = 2\pi$, where L crosses R . Restricting $\tau(M) < \pi$ and $\Omega \leq \pi$ avoids crossing; see Fig. 27. I do not believe the angle restrictions are close to tight. For example, the theorem seems to hold for $\tau(M) < 2\pi$ and $\Omega \leq \pi$, and for $\tau(M) < \pi$ and $\Omega \leq 2\pi$. The severe angle restrictions are imposed to obtain a clean proof. And they seem to suffice for non-overlap of random spherical polyhedra.




Figure 26: A 70° -spiral M , with $\Omega = 2\pi$. L crosses R right-to-left.




Figure 27: $\Omega = \pi$ so the last tangents to L and R are 180° apart. Here $\tau(M) < \pi$.

We prove Theorem 2 via two lemmas.

Lemma 6 *Under the conditions stated, L cannot cross M left-to-right.*

Proof: (Crossing right-to-left will be excluded by the next lemma.) The proof is by induction. It clearly holds for $k = 1$, when M is just a single edge, and L is rotated ω_0 ccw of M .

Suppose the claim is true for k . So the situation is like that shown in Fig. 25. Imagine extending M prior to v_0 , to a vertex v' . Then the L path is rigidly rotated ccw by ω' . The new path L' is strictly further away from M along the concentric circles centered on v' . See Fig. 28. \square




Figure 28: Extending the L and M paths in Fig. 25 by one edge $v'v_0$. (R not shown.)

Lemma 7 *Under the conditions stated, L cannot cross R right to left.*

Proof: With the convention that $m_1 - m_0$ aims along the $+x$ -axis, the restriction $\tau(M) \leq \pi/2$ implies that all the edge vectors of M point in the $+x$ -halfplane. The restriction $\Omega \leq \pi$ turns the vectors of M to point into the $+y$ -halfplane for L ; recall λ_i is turned ccw $\omega_i/2$ from the corresponding M angle. Similarly the vectors of M are turned cw $\omega_i/2$ and so point into the $-y$ -halfplane. Thus L lies entirely in the upper halfplane and R in the lower, so they cannot cross. \square

These two lemmas prove Theorem 2.

5 Algorithm 2: Finding a Radially Monotone Cut Tree

Using the definition in the previous section, we now suggest an “algorithm” to find a radially monotone cut tree for a convex polyhedron. The algorithm does

not provably succeed, even on spherical non-obtusely triangulated polyhedra. But its strong empirical performance suggests some variant might provably succeed. For now, I just describe the decisions made to achieve a definite algorithm.

There are five decisions that drive the algorithm:

1. Attention is restricted to spherical polyhedra \mathcal{P} , vertices inscribed in a sphere \mathcal{S} .
2. Latitude circles on \mathcal{S} are used to mimic the concentric circles employed in the 2D Algorithm 1.
3. The algorithm constructs trees from paths Q on \mathcal{P} whose planar medial paths $M(Q)$ are radially monotone.
4. The full \mathcal{P} is treated as a convex cap C with boundary ∂C .
5. Hourglass paths are abandoned.

We now remark on each of these decisions.

(1) Spherical polyhedra are special because of their relationship to Delaunay triangulations, a point to which we will return later. Outside of this class, the algorithm is far less successful. For example, it fails to find a rm cut tree for more than 50% of 100-vertex random polyhedra inscribed in an ellipsoid with axes $(1, 1, \frac{1}{4})$. See Fig. 29.




Figure 29: Overlap. The 100-vertex polyhedron is drawn from an ellipsoid with axes $(1, 1, \frac{1}{4})$. The three marked vertices could not be joined radially monotonically to the cut forest. Two of the non-rm cuts cause overlap in the unfolding.

(2) In order to mimic the successful 2D Algorithm 1, we process the vertices in order of their geodesic distance from the north pole N , furthest first; see Fig. 30. We view v as lying on a planar circle centered on N with radius γ .

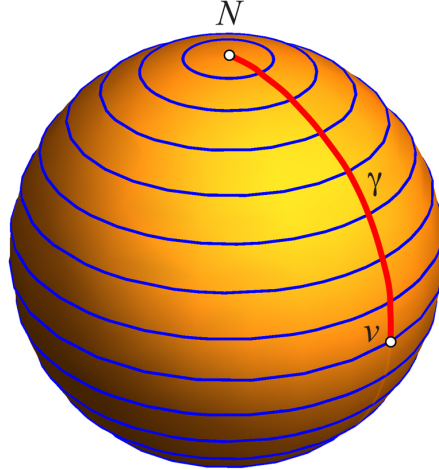


Figure 30: Concentric circles representing latitude lines on \mathcal{S} .

(3) We have already justified the concentration on medial rm cut paths in Section 4.1: the unfolded two sides of the cut do not intersect (Theorem 2).

(4) Fig. 5 showed a convex cap with a jagged boundary ∂C . One can extend a convex cap of polyhedron \mathcal{P} by removing just one triangle $\triangle abc$ from \mathcal{P} : then the convex “cap” is all but that one triangle. Even further, one can remove an infinitesimal slice around the boundary of that triangle, the nonconvex degenerate quadrilateral $\partial C = (a, b, c, b)$. Then C includes every face of the full \mathcal{P} .

Returning to our first figure, Fig. 1, the green V-shape represents that infinitely thin quadrilateral ∂C . The rm cut forest \mathcal{F} is grown from ∂C , and connected to a tree by the two edges ab and bc .

For our experiments with random spherical polyhedra, we added north- and south-pole points N and S , just for coding convenience. N plays the role of the 2D bounding circle center x in Algorithm 1, and S is a vertex of $\triangle abc$. This is illustrated more clearly in Fig. 31, which shows the three trees grown from ∂C .

(5) The reason I abandoned hourglass paths is that, under various generalizations, they did not always seem to exist. They were a specific tool to prove that Algorithm 1 finds rm forests in 2D, but seem less useful in 3D.




Figure 31: \mathcal{P} has $n = 50$ vertices. N and S label the north and south poles respectively.

5.1 More Examples

Figs. 32, 33, 34, and Figs. 35 and 36, show four more example unfoldings.




Figure 32: Cut tree and unfolding of a polyhedron of $n = 300$ vertices.

5.2 Algorithm 2

The algorithm is detailed below.

5.3 Empirical Results

Algorithm 1 only provably works on planar convex domains when they are non-obtusely triangulated, but the 3D examples we have so far provided contain obtuse triangles: They were generated simply as convex hulls of random points on a sphere \mathcal{S} , partly because it is by no means straightforward to generate non-obtuse triangulations, and partly because Algorithm 2 works with empirically high frequency on spherical polyhedra.

My experiments show that Algorithm 2 finds a radial monotone cut tree \mathcal{T} for more than 95% of random spherical polyhedra of $n = 100$ vertices. For example, in one run, 20 out of 1,000 polyhedra led to one (never more than one) vertex forced to select a non-rm connection. In all cases, the “cause” was an obtuse triangle, which if split, led to a rm-tree and a non-overlapping unfolding. Another run of 400 random 200-vertex spherical polyhedra found 21 single-vertex non-rm connections (95%).




Figure 33: Unfolding of a polyhedron of $n = 500$ vertices.




Figure 34: Unfolding of a polyhedron of $n = 1000$ vertices.




Figure 35: Spanning radial monotone cut tree for polyhedron of $n = 1500$ vertices. North pole marked on top; tree root not visible at south pole. Unfolding in Fig. 36.




Figure 36: Unfolding of polyhedron (Fig. 35) of $n = 1500$ vertices.

Algorithm 2: Find rm cut forest \mathcal{F} for polyhedron \mathcal{P}

Input : Convex polyhedron \mathcal{P} inscribed in sphere \mathcal{S}

Output: (Usually) Radially monotone cut tree \mathcal{T}

Find “bottommost” triangle $\triangle abc$ (most southward normal).

Set $\partial C = (a, b, c, b)$.

// Grow \mathcal{F} from ∂C upward/inward

Sort interior vertices by geodesic distance on \mathcal{S} from north-pole N , those nearest ∂C first.

// Grow \mathcal{F} :

$\mathcal{F} \leftarrow \emptyset$

foreach vertex v_0 in sorted (vertically ascending) order **do**

foreach vertex v_1 already in \mathcal{F} or on ∂C **do**

 // v_1 is below v_0 .

 If v_0 connects by a triangulation edge to v_1 , set $e = (v_0, v_1)$.

 Check if the planar medial path $M(Q)$ from v_0 to ∂C in $\mathcal{F} + e$ is radially monotone.

 If so, record its worst turnangle τ (with $\tau > 90^\circ$ not rm).

end

 Choose the e^* that has the best (minimum) τ .

 Or: report failure to find a radially monotone connection & exit.

$\mathcal{F} \leftarrow \mathcal{F} + e^*$

end

Return $\mathcal{T} \leftarrow \mathcal{F} + ab + bc$

Fig. 37 shows a typical example. Fig. 37(a) shows the one vertex that makes a non-rm cut, in this case with turnangle 94° . This angle in fact causes (barely visible) overlap in the unfolding. Fig. 37(b) shows the same polyhedron with many obtuse angles split, splitting the most common situation: one obtuse triangle adjacent to a non-obtuse triangle. Now a different radially monotone cut tree is found, in particular, resolving the problematic overlap.




Figure 37: (a) \mathcal{P} of $n = 100$ vertices, with non-rm vertex cut marked. (b) The same \mathcal{P} with many obtuse angles split, now of $n = 203$ vertices.

Caveats. Although it is quite clear what constitutes a random spherical polyhedron—the convex hull of random points on a sphere—it is less clear what is a random non-obtusely triangulated spherical polyhedron. And in any case, the obtuse-splitting procedure I employed is ad hoc, and creates “vertices” of curvature 0 not touching \mathcal{S} . So I cannot make any justified empirical claims concerning the performance of Algorithm 2 on such polyhedra.

Moreover, the 1,000 examples mentioned in the Introduction were achieved by (a) running Algorithm 2 on spherical polyhedra, and then, for each case where a rm tree was not found, (b) rerunning it with obtuse triangles split as in Fig. 37.

However, I would like to emphasize that Algorithm 2 does not in any way “search” for a radially monotone cut tree: It uses whatever happens to be the bottommost triangle to form ∂C , and then it grows the spanning forest strictly in order of the geodesic circle radii illustrated in Fig. 30. It never backtracks or considers alternatives (aside from choosing the “best” rm connection among those available below). One could apply many heuristics to improve performance. In fact, when I included several such heuristics—selecting the “most

equilateral” triangle to become bottommost (to maximally separate the three spanning tree roots of \mathcal{F}), and not following the concentric circle ordering, but rather growing \mathcal{F} with the “best” connection (in any direction) at each stage—the heuristic-laden algorithm found random cut trees with such high frequency I could not find a random counterexample.

6 Questions & Conjectures

Algorithm 2 finds radially monotone cut trees for random spherical polyhedra with high frequency, and then (because of Theorem 2, whose bounds apparently suffice for random polyhedra), unfolds them without overlap. Triangulations of spherical polyhedra are special: They are intimately connected to Delaunay triangulations. This suggests:

Question 1 *Does every planar Delaunay triangulation have a radially monotone spanning forest?*

Note that the example in Appendix 1 (Fig. 39) is non-Delaunay. Also note that non-obtuse triangulations are automatically Delaunay. The 3D analog is:

Question 2 *Does every spherical (inscribed) polyhedron have a radially monotone cut tree?*

My empirical explorations only suggest that this may hold for random spherical polyhedra.

Question 3 *Does every non-obtusely triangulated convex polyhedron have a radially monotone cut tree?*

I have so far not explored non-spherical polyhedra enough to form an opinion on this question, but radial monotonicity seems intimately connected to non-obtuseness.

Question 4 *What is a natural definition of a random, spherical (inscribed in a sphere), non-obtusely triangulated polyhedron? And how could they be generated?*

Without an answer to this question, the following conjecture is vague, but nevertheless, I feel is justified:

Conjecture 1 *A random spherical, non-obtusely triangulated polyhedron has a radially monotone cut tree (as defined in Sec. 4) with high probability.*

More risky is the same conjecture without requiring the triangulation to be non-obtuse:

Conjecture 2 *A random spherical polyhedron has a radially monotone cut tree with high probability.*

One natural interpretation of “high probability” would be that, as the number of vertices $n \rightarrow \infty$, the probability goes to 1. This would essentially constitute an obverse of Fig. 2, which shows (empirically) that the probability of overlap from a random spanning cut tree goes to 1.

7 Appendix 1




Figure 38: A plane graph that cannot be spanned by an rm-forest. The circles indicate rm violations.

We first provide a counterexample for a plane graph, and then modify it to achieve a triangulation.

Lemma 8 *The plane graph G shown in Fig. 38 has no rm-forest. G has strictly convex faces, and a convex boundary ∂C .*

Proof: Let \mathcal{F} be an rm-forest for G , and let Q the unique path in a tree of \mathcal{F} from x to ∂G .

1. The four edges incident to x are symmetric, so we choose (x, ab) wlog. There are two choices from ab .
2. $Q = (x, ab, a, a')$ is non-monotonic at ab w.r.t. x : (ab, a) cuts into the circle centered at x passing through ab .
3. (x, ab, b) is radially monotone.
4. Extending this path with the up edge (b, ab') violates radial monotonicity at b w.r.t. ab . So it must be instead extended downward, to $Q = (x, ab, b, cd')$.
5. (x, ab, b, cd') cannot be extended to either $Q = (x, ab, b, cd', c')$ nor $Q = (x, ab, b, cd', b')$ monotonically: the former violates radial monotonicity at cd' w.r.t. x ; the latter violates at cd' w.r.t. b .

Because of the symmetries of G , this exhausts all possible paths from x to ∂G .

The pentagon face (a, ab, b, ab', a') and the quadrilateral face (b, ab', b', cd') are non-strictly convex faces, but can be made strictly convex by slight movements of ab and b respectively, without changing altering the monotonicity of any paths. \square

The claim in Lemma 8 can be strengthened to a triangulation:

Lemma 9 *The triangulated plane graph G_T shown in Fig. 39 has no rm-forest.*

Proof: *Sketch.* The structure of G_T is based on that of G in Lemma 8: G_T is a triangulation of G . Although there are now more paths from x to ∂G_T , they each still violate monotonicity, either w.r.t. x , or w.r.t. b . \square




Figure 39: A triangulated plane graph that cannot be spanned by an rm-forest.

References

- [Bis16] Christopher Bishop. Nonobtuse triangulations of PSLGs. *Discrete Comput. Geom.*, 56(1):43–92, 2016.
- [DO07] Erik D. Demaine and Joseph O’Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, July 2007. <http://www.gfalop.org>.