

Angle-monotone Paths in Non-obtuse Triangulations

Anna Lubiw*

Joseph O'Rourke†

Abstract

We reprove a result of Dehkordi, Frati, and Gudmundsson: every two vertices in a non-obtuse triangulation of a point set are connected by an angle-monotone path—an xy -monotone path in an appropriately rotated coordinate system. We show that this result cannot be extended to angle-monotone spanning trees, but can be extended to boundary-rooted spanning forests. The latter leads to a conjectural edge-unfolding of sufficiently shallow polyhedral convex caps.

1 Introduction

The central result of this paper is to offer an alternative—and we believe simpler—proof of a result of Dehkordi, Frati, and Gudmundsson [DFG15] (henceforth, DFG):

“Lemma 4. Let G be a Gabriel triangulation on a point set P . For every two points $s, t \in P$, there exists an angle θ such that G contains a θ -path from s to t .”

We first explain this result, using notation from [BBC⁺16], before detailing our other contributions. First, we use S for the point set and β instead of θ . A Gabriel triangulation as defined by DFG is a triangulation of S where “every angle of a triangle delimiting an internal face is acute.” Because neither they nor we need any of the various properties of Gabriel triangulations except the angle property, and we only need non-obtuse rather than strict acuteness, we define G to be a plane geometric graph that is a non-obtuse triangulation of S .

Define the *wedge* $W(\beta, v)$ to be the region of the plane bounded by rays at angles $\beta \pm 45^\circ$ emanating from v . W is closed along (i.e., includes) both rays, and has angular *width* of 90° . (Later we generalize to widths different from 90° .) A polygonal path (v_0, \dots, v_k) consisting of edges of G is called β -*monotone* (for short, a β -*path*) if the vector of every edge (v_i, v_{i+1}) lies in $W(\beta, v_0)$.

These are the θ -paths of DFG. Note that if $\beta = 45^\circ$, then a β -monotone path is both x - and y -monotone with respect to a Cartesian coordinate system. A path that is β -monotone for some β is called *angle-monotone*.

Our phrasing of the DFG result is:

Theorem 1 *In a non-obtuse triangulation G , every pair of vertices is connected by an angle-monotone path.*

Other Contributions. We extend Theorem 1 to wedges of any width γ —if a plane geometric graph that includes the convex hull of S has all angles at most γ , then there is an angle-monotone path of width γ between any two vertices. Of necessity, $60^\circ \leq \gamma < 180^\circ$. One significance of angle-monotone paths of width $\gamma < 180^\circ$ is that they have a spanning ratio of $1/\cos \frac{\gamma}{2}$ [BBC⁺16]. We do not pursue that aspect here. Instead, we investigate angle-monotone spanning trees. These were studied independently in [MS16], which addressed recognition and construction, but not existence—our focus. We show that Theorem 1 does not extend to angle-monotone spanning trees, but does extend to boundary-rooted spanning forests. Then, in Section 5 we make a novel connection to edge-unfolding polyhedra.

2 Proof

We prove Theorem 1 by showing that there is an angle-monotone path from an arbitrary fixed vertex s to every other vertex. The proof uses an angular sweep around s , which by convention we place at the origin. We first consider a fixed but arbitrary angle β and investigate which vertices are reached by β -paths from s . Let ∂G be the boundary of G , i.e., the convex hull of S . Our proof relies on two properties of vertices v not on ∂G : (1) the wedge $W(\beta, v)$ includes at least one edge incident to v ; (2) if the wedge has only one edge incident to v then that edge does not lie along a bounding ray of the wedge. In order to avoid dealing with boundary vertices as a separate case, we will augment G so that conditions (1) and (2) hold for boundary vertices as well. At every vertex v on ∂G add a finite set of rays that subdivide the exterior angle at v into angles of at most 90° . Call the result G^+ . By construction, properties (1) and (2) now hold for every vertex v if we consider both edges

*School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada. alubiw@uwaterloo.ca.

†Department of Computer Science, Smith College, Northampton, MA, USA. JORourke@smith.edu.

and rays incident to v . Note that no added ray crosses an edge of G . In the special case of widths $\geq 90^\circ$, the rays can be chosen so that they do not intersect one another. When we generalize to smaller widths, the rays will necessarily intersect each other but this will not influence our proof. In our figures, the rays are drawn short as a reminder that only their angles at the convex hull are relevant.

A β -path starting at s is *maximal* if there is no edge or ray that can be added at the end of the path while keeping it a β -path. In particular, a path ending with a ray is maximal. DFG proved that every maximal β -path terminates on ∂G , which in our terms becomes:

Lemma 2 *Any maximal β -path ends with a ray.*

Proof. Consider a β -path that ends at a vertex v . The wedge $W(\beta, v)$ must include an edge or ray of G^+ by property (1) so the path can be extended. \square

We will see later (in Fig. 3(b)) that it is possible for a β -path to include edges of ∂G , return to interior edges, and then later again include edges of ∂G . For a fixed β , let $P(\beta)$ be the set of all maximal β -paths starting from s . Let $V(\beta)$ and $E(\beta)$ be the set of vertices and edges/rays in $P(\beta)$.

Let $U(\beta) \in P(\beta)$ be the *upper envelope* of $P(\beta)$, defined as follows. Starting from vertex $v = s$, grow $U(\beta)$ by always following the most counterclockwise edge/ray from v falling within $W(\beta, v)$. $U(\beta)$ is necessarily a maximal β -path, so it ends with a ray. By property (2) above, we have:

Observation 1 *$U(\beta)$ does not include any edge/ray along the lower ray of $W(\beta, s)$, at angle $\beta - 45^\circ$.*

$L(\beta)$ is similarly the *lower envelope*, the most clockwise path. Note that “upper” and “lower” are to be interpreted as most counterclockwise and most clockwise respectively, not in terms of y -coordinates.¹ Finally, define $R(\beta)$ to be the region of the plane whose boundary is $L(\beta), U(\beta)$. Fig. 1 illustrates these definitions.

Lemma 3 *Every vertex in $R(\beta)$ is in $V(\beta)$, i.e., every vertex in $R(\beta)$ can be reached from s via a β -path.*

Proof. Wlog assume $\beta = 45^\circ$, so that the wedge rays are at 0 and 90° . Let $v \in V(\beta)$ be the leftmost *inaccessible* vertex, i.e., the leftmost vertex not reached by a β -path from s . If there are ties for leftmost, let v be the lowest. Consider the backward wedge $\overline{W}(\beta, v)$ at v . Note that s lies in $\overline{W}(\beta, v)$. Consider the line segment sv . It lies in $\overline{W}(\beta, v)$ and inside the convex hull of S . Imagine rotating sv clockwise or counterclockwise about

¹These notions are analogous but not equivalent to DFG’s “high” and “low” paths.

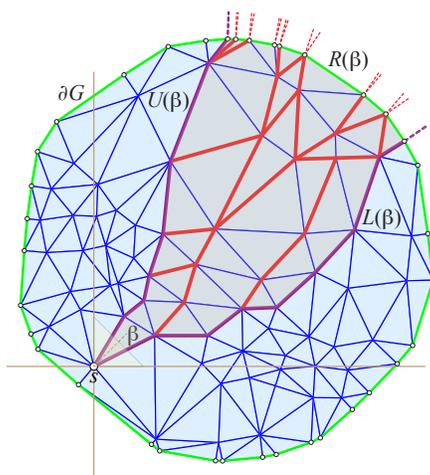


Figure 1: $P(\beta)$ edges are marked; $\beta = 45^\circ$. $L(\beta)$ and $U(\beta)$ delimit the region $R(\beta)$. Rays shown only for $V(\beta)$ hull vertices.

v while remaining inside $\overline{W}(\beta, v)$ and inside the convex hull of S in a small neighborhood of v . Since there are no obtuse angles at v , rotating in one direction or the other must result in an edge or ray in $\overline{W}(\beta, v)$. Furthermore, the result cannot be a ray since we never leave the convex hull. Thus we have identified an edge $e = (u, v)$ in $\overline{W}(\beta, v)$.

Suppose first that u is in $R(\beta)$. Because v is the leftmost lowest inaccessible vertex, u must be accessible (note that u must be at the same y -height or lower than v). But now v lies in $W(\beta, u)$, and so v is accessible after all, a contradiction. Instead suppose u lies outside $R(\beta)$. Then e must cross the boundary of $R(\beta)$. But that boundary is composed of edges/rays of G^+ , and e cannot cross an edge of G^+ without the two sharing a vertex, which would lie on the boundary of $R(\beta)$, not the exterior, again a contradiction. \square

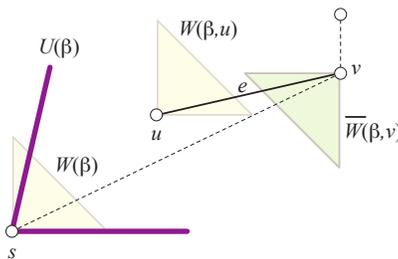


Figure 2: No vertex in $R(\beta)$ is inaccessible: all are reached by a β -path from s .

2.1 Critical angles β_i

We now analyze the relationships between $P(\beta_i)$ and $P(\beta_{i+1})$, where β_{i+1} is the next “relevant” angle after β_i , with the goal of showing that all vertices in G are “covered” over all β_i , i.e., belong to $P(\beta_i)$ for some β_i . Throughout, fix the source s , and let $W(\beta) = W(\beta, s)$. Define an angle β to be *critical* if $P(\beta + \varepsilon)$ or $P(\beta - \varepsilon)$ differs from $P(\beta)$, for an arbitrarily small $\varepsilon > 0$. At a critical angle β , one or both rays of $W(\beta)$ are parallel to one or more edges of $P(\beta)$. If $P(\beta + \varepsilon)$ differs from $P(\beta)$, one or more edges parallel to the $\beta - 45^\circ$ ray drop out of $P(\beta)$. If $P(\beta - \varepsilon)$ differs from $P(\beta)$, one or more edges parallel to the $\beta + 45^\circ$ ray enter $P(\beta)$. Fig. 3 illustrates two adjacent critical angles.

Let $\beta_1, \beta_2, \dots, \beta_i, \beta_{i+1}, \dots$ be the critical angles, sorted counterclockwise. For every β strictly between two adjacent critical angles, $\beta_i < \beta < \beta_{i+1}$, $P(\beta)$ is the same. We use the notation $P(\beta_{i'})$ to represent this intermediate set, which differs from $P(\beta_i)$ or $P(\beta_{i+1})$ or both.

In the transition from $P(\beta)$ to $P(\beta + \varepsilon)$, edges can drop out of $P(\beta)$. In particular, any edge $e = (u, v)$ that lies along the $\beta - 45^\circ$ -ray of $W(\beta, u)$ will drop out of $P(\beta)$. Furthermore, when edges drop out of $P(\beta)$ this may cause vertices to drop out of $P(\beta)$, and any edge originating at a dropped vertex also drops out of $P(\beta)$.

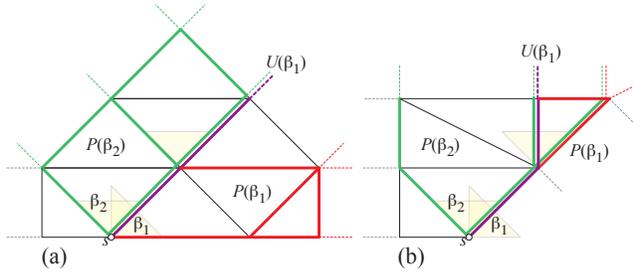


Figure 3: $P(\beta_1)$ (red), $U(\beta_1)$ (purple), and $P(\beta_2)$ (green). (a) $P(\beta_1)$ includes some but not all ∂G edges. (b) $U(\beta_1)$ includes an edge of ∂G before an internal edge.

The next lemma shows that no vertices fall strictly “between” $P(\beta_i)$ and $P(\beta_{i+1})$, where they would escape being spanned.

Lemma 4 $U(\beta_i)$ is a path in $P(\beta_{i+1})$.

Proof. Since edges may only enter, not leave, in the transition from $P(\beta_{i+1} - \varepsilon)$ to $P(\beta_{i+1})$, the lemma is equivalent to the claim that $U(\beta_i)$ is a path in $P(\beta_i + \varepsilon) = P(\beta_{i'})$. Let $\beta = \beta_i$, and $\beta' = \beta_{i'}$. Then we aim to prove that $U(\beta) \subseteq P(\beta')$. This requires showing that no edge $e \in U(\beta)$ drops out from $P(\beta)$ to $P(\beta')$, as β increases to β' . As usual, assume that $\beta = 45^\circ$.

Suppose to the contrary that some edge drops out, and let $e = (u, v)$ be the leftmost lowest edge with $e \in U(\beta)$ but $e \notin P(\beta')$. Equivalently, e is the first edge of $U(\beta)$ that is not in $P(\beta')$. Because u is in $P(\beta')$, the only reason for e to drop out is that it lies along the lower, horizontal ray of the wedge $W(\beta, u)$. But by Observation (1), $U(\beta)$ does not include any edge along the lower ray of the wedge. \square

In analogy with the definition of $R(\beta)$, define $R(\beta_i, \beta_j)$ for $j > i$ to be the region bound by $L(\beta_i)$, $U(\beta_j)$, and the portion of ∂G between those lower and upper envelope endpoints.

Lemma 5 $R(\beta_i, \beta_{i+1}) = R(\beta_i) \cup R(\beta_{i+1})$. Informally, no vertices are “orphaned” between $P(\beta_i)$ and $P(\beta_{i+1})$.

Proof. The lemma essentially says that no vertices are “orphaned” between $P(\beta_i)$ and $P(\beta_{i+1})$, and this follows immediately from the fact that $U(\beta_i) \subseteq P(\beta_{i+1})$ as established in Lemma 4. \square

Now we can prove Theorem 1, the key result of [DFG15]:

Proof. [of Theorem 1] Fix s and construct $\bigcup_i P(\beta_i)$. By Lemma 5, this is a spanning graph of G , and so must include a path from s to v . \square

Our arguments extend to wedges of any width γ , thus proving that if a plane geometric graph that includes the convex hull of S has all internal angles at most γ , then there is an angle-monotone path of width γ between any two vertices. This answers a question raised in [BBC⁺16].

3 Spanning Tree

Now that Theorem 1 has established that there is a graph spanning all of G with angle-monotone paths from any source s , it is natural to wonder if the claim can be strengthened to the existence of an *angle-monotone spanning tree* for any s : a tree rooted at s with an angle-monotone path from s to any $v \in G$. The answer is NO, but we canvass a few positive results before detailing a counterexample for spanning trees. Throughout, we let s be an arbitrary vertex of G . First, within a fixed β region, $P(\beta)$ can be easily spanned:

Lemma 6 $P(\beta)$ includes an angle-monotone tree that spans the same vertices, $V(\beta)$.

Proof. By Lemma 3, $P(\beta)$ reaches every vertex in $V(\beta)$. For each vertex $v \in V(\beta)$, in any order, delete all but one incoming edge to v . Because an incoming edge remains to each v , v is spanned. Because eventually no v has more than one incoming edge, no cycles can remain. See Fig. 4. \square

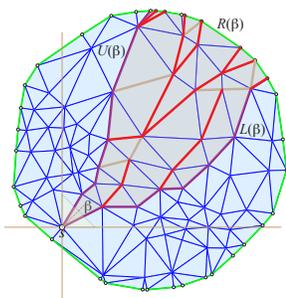


Figure 4: $P(\beta)$ spanning tree. Light-brown edges have been deleted.

We now consider triangulations with special angles.

Lemma 7 *Let G_{45° have edges only at multiples of 45° . Then there is an angle-monotone spanning tree rooted at any source vertex s .*

Proof. Let β_1 and $\beta_2 = \beta_1 + 45^\circ$ be two consecutive critical angles. We argue that $U(\beta_1)$ and $U(\beta_2)$ may share an initial portion of a path, but then diverge and do not rejoin before reaching their terminal rays.

By Observation (1), $U(\beta)$ only includes edges at angles β_i or $\beta_i + 45^\circ$. So the most counterclockwise edge in $U(\beta_1)$ is $\beta_1 + 45^\circ$, and the most clockwise edge in $U(\beta_2)$ is $\beta_2 = \beta_1 + 45^\circ$. Thus, $U(\beta_1)$ and $U(\beta_2)$ can have parallel edges, but once they separate, they can never rejoin.

Now it is easy to create a spanning tree between $U(\beta_1)$ and $U(\beta_2)$ that retains all edges in these two envelopes, by deleting all but one incoming edge to each vertex between the envelopes. \square

The problematic possibility avoided in such G_{45° graphs is $U(\beta_1)$ and $U(\beta_2)$ joining, separating, and rejoining. Already in a graph G_{30° that has edges only at multiples of 30° , the divergence of upper envelopes used in Lemma 7 is no longer guaranteed, and thwarts that proof.

3.1 Spanning Tree Counterexample

Fig. 5 shows a graph G that does not have an angle-monotone spanning tree rooted at s . The construction allows two angle-monotone paths to each of $\{C, D, E, F\}$, one of which is marked green in the figure. But vertices A and B are shifted slightly toward one another, which breaks the symmetry and, as we shall argue below, results in a unique angle-monotone path to each. The union of those two unique paths contains the cycle (s, a, x, b) . Thus there is no angle-monotone spanning tree from s .

We now argue that there is no angle-monotone path to A other than $saxA$. This is simply a matter of checking that any other path to A contains two *spread-apart* edges whose vectors do not lie in a 90° wedge. In particular, the path $sbxA$ contains spread-apart edges sb and xA , and the path $sawA$ contains spread-apart edges aw and wA . Other paths can be checked similarly. Similar reasoning constrains the (symmetric) paths to B .

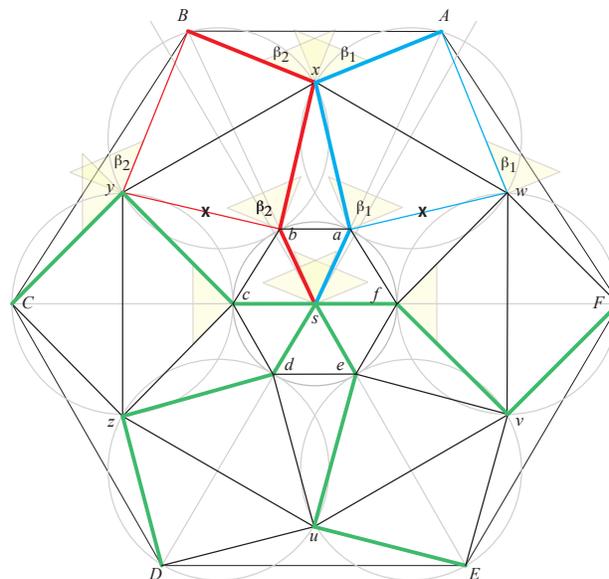


Figure 5: (s, a, x, A) is the unique angle-monotone path to A , and (s, b, x, B) is the unique angle-monotone path to B , forming the cycle (s, a, x, b) .

The outer ring of six circles in the construction make clear that various diameter-spanning angles are 90° , but points $\{a, A, \dots, f, F\}$ could be moved slightly exterior to those circles, rendering those angles $< 90^\circ$, while retaining the properties that force the (s, a, x, b) cycle. So the counterexample is “robust” in this sense.

4 Spanning Forest

For the unfolding application discussed in the next section, it is useful to span G by a boundary-rooted forest \mathcal{F} : A set of disjoint angle-monotone trees, each with its root on ∂G , and spanning every interior vertex of G . This can be achieved with β -monotone trees for just four β values.

With a Cartesian coordinate system centered on vertex s , define the quadrants Q_0, Q_1, Q_2, Q_3 as follows. Q_0 is the quadrant coincident with $W(\beta, s)$ when $\beta = 45^\circ$, closed along the x -axis and open along the y -axis, and includes the origin s . $Q_i, i > 0$ are defined analogously, except those quadrants do not include the origin. Thus

the quadrants are pairwise disjoint and together cover the plane.

We construct separate spanning forests for each quadrant, following Algorithm 1, which grows paths from vertices interior to G to ∂G . See Fig. 6.

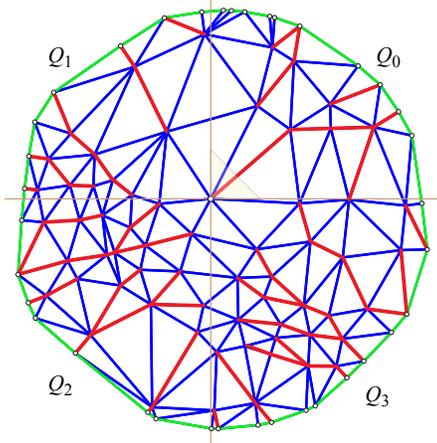


Figure 6: Spanning forest resulting from Algorithm 1.

<p>Algorithm 1: Algorithm to construct spanning forest \mathcal{F}</p> <p>Input : Non-obtuse triangulation G</p> <p>Output: Spanning forest \mathcal{F} composed of β-monotone paths</p> <p>// Quadrants Q_j, each corresponding to $\beta_j = 45^\circ + j \cdot 90^\circ$, $j = 0, 1, 2, 3$.</p> <p>foreach Quadrant Q_j, $j = 0, 1, 2, 3$ do</p> <p style="padding-left: 1em;">$F_j \leftarrow \emptyset$</p> <p style="padding-left: 1em;">// Grow forest F_j inside Q_j.</p> <p style="padding-left: 1em;">foreach $v \in Q_j$ do</p> <p style="padding-left: 2em;">if $v \notin F_j$ then</p> <p style="padding-left: 3em;">Grow β_j-path p from v.</p> <p style="padding-left: 3em;">Stop when p reaches a vertex in F_j, or reaches ∂G.</p> <p style="padding-left: 2em;">end</p> <p style="padding-left: 1em;">$F_j \leftarrow F_j \cup p$</p> <p style="padding-left: 1em;">end</p> <p>end</p> <p>$\mathcal{F} = F_0 \cup F_1 \cup F_2 \cup F_3$</p> <p>return \mathcal{F}.</p>

Lemma 8 Algorithm 1 outputs a boundary-anchored spanning forest, each tree of which is composed of β -monotone paths, for four β 's: $\beta_j = 45^\circ + j \cdot 90^\circ$, $j = 0, 1, 2, 3$.

Proof. Observe that a β_j -path grown from $v \in Q_j$ remains in Q_j . So all the trees in F_j are composed of

β_j -paths. All vertices of each quadrant are spanned, because the inner loop of Algorithm 1 runs over all $v \in Q_j$. No cycles can be created because the algorithm only grows a path from v if v is not yet in F_j . So v becomes a leaf of some tree in F_j when its path reaches that tree. \square

5 Unfolding

Now we discuss an application of Algorithm 1 and Lemma 8 to edge-unfolding nearly flat convex caps. We only sketch the argument, as several steps need considerable elaboration, and other steps rely on definitions and details in an unpublished report [O'R16]. So this section will end with a conjecture rather than a theorem. At a high level, the construction depends on two claims: (1) angle-monotone paths are “radially monotone paths,” a concept introduced in [O'R16], but known before as backwards “self-approaching curves” [IKL99]. (2) Theorem 2 of [O'R16] concludes that the unfolding of a particular “medial” cut path M on a polyhedron is radially monotone and so does not self-cross when unfolded (if certain angle conditions are satisfied).

Let P be a convex polyhedron, and let $\phi(f)$ for a face f be the angle the normal to f makes with the z -axis. Let H be a halfspace whose bounding plane is orthogonal to the z -axis, and includes points vertically above that plane. Define a *convex cap* C of angle Φ to be $C = P \cap H$ for some P and H , such that $\phi(f) \leq \Phi$ for all f in C . We will only consider $\Phi < 90^\circ$, which implies that the projection C_\perp of C onto the xy -plane is one-to-one.

Say that a convex cap C is *acutely triangulated* if every angle of every face is strictly acute. Note that P being acutely triangulated does not always imply that $C = P \cap H$ is acutely triangulated, but it is known that any polyhedron can be acutely triangulated [Bis16]. We will need this lemma.

Lemma 9 Let a triangle in \mathbb{R}^3 , whose face normal makes angle ϕ with the z -axis, have one angle α , which projects to α_\perp on the xy -plane. Then the maximum value of $\Delta = |\alpha - \alpha_\perp|$ is a monotonically increasing function, as plotted in Fig. 7.

We only need that $\Delta \rightarrow 0$ as $\phi \rightarrow 0$, so we will not calculate the function explicitly. For example, for $\phi < 10^\circ$, $\Delta < 1^\circ$.

For a triangulated convex cap C , let α_{\max} be the maximum of any triangle angle. Using Lemma 9, we can guarantee that an acutely triangulated cap C will project to a non-obtuse plane graph C_\perp by choosing Φ so that $\Delta < 90^\circ - \alpha_{\max}$.

Now we apply Algorithm 1 and Lemma 8 to obtain an angular-monotone spanning forest \mathcal{F}_\perp of C_\perp . We then

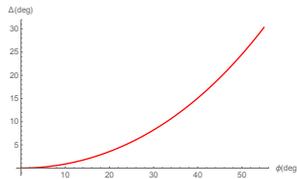


Figure 7: The maximum face angle change resulting from projection with normal at angle ϕ .

lift the trees in \mathcal{F}_\perp to cut trees \mathcal{F} on C in \mathbb{R}^3 . Again Lemma 9 ensures this can be accomplished without any turn angle in any path in any tree of \mathcal{F} exceeding 90° . Now finally we invoke a version of Theorem 2 as mentioned previously, which guarantees that the cut paths unfold without local overlap. We leave it a claim that the angle conditions for that theorem are satisfied by selecting Φ small enough. The conclusion is that the lifted paths are “radially monotone,” which is the condition that implies unfolding without overlap. The end result is this:

Conjecture 1 *For an acutely triangulated convex cap C with sufficiently small Φ bounding face normals, the spanning forest \mathcal{F}_\perp resulting from Algorithm 1 lifts to a cut forest \mathcal{F} that edge-unfolds C without overlap.*

We have implemented this construction. Fig. 8 shows a convex cap with $\Phi \approx 27^\circ$, and Fig. 9 shows the corresponding unfolding.²

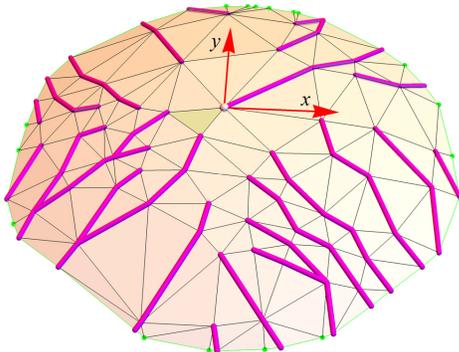


Figure 8: The cut forest \mathcal{F} resulting from lifting \mathcal{F}_\perp to the convex cap. (The marked face is the root of the dual unfolding tree.)

Acknowledgements. We thank Debajyoti Mondal for observing that our proof of Theorem 1 works for widths other than 90° .

²The forest in Fig. 8 is slightly different than that shown in Fig. 6, due to different ordering choices of $v \in Q_j$.

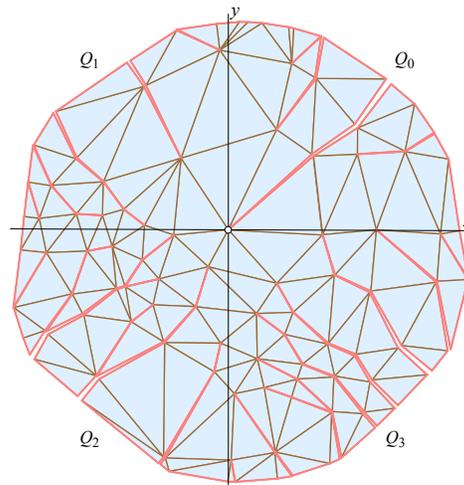


Figure 9: The edge-unfolded convex cap. The origin and quadrants used in Algorithm 1 are indicated.

References

- [Bis16] C.J. Bishop. Nonobtuse triangulations of PSLGs. *Discrete & Computational Geometry*, 56(1):43–92, 2016.
- [BBC⁺16] N. Bonichon, P. Bose, P. Carmi, I. Kostitsyna, A. Lubiw, and S. Verdonschot. Gabriel triangulations and angle-monotone graphs: Local routing and recognition. In: *24th Graph Drawing and Network Visualization (GD)*. Lect. Notes Comput. Sci., 9801. Springer, 2016, pages 519–531.
- [DFG15] H.R. Dehkordi, F. Frati, and J. Gudmundsson. Increasing-chord graphs on point sets. *J. Graph Algorithms Applications*, 19(2):761–778, 2015.
- [IKL99] C. Icking, R. Klein, and E. Langetepe. Self-approaching curves. *Math. Proc. Camb. Phil. Soc.*, 125:441–453, 1999.
- [MS16] K. Mastakas and A. Symvonis. Rooted uniform monotone minimum spanning trees. *Internat. Conf. Algorithms Complexity (CIAC)*, Athens, Greece, 2017. Lect. Notes Comput. Sci. 10236, Springer.
- [O’R16] J. O’Rourke. Unfolding convex polyhedra via radially monotone cut trees. arXiv:1607.07421 [cs.CG], 2016.