

Edge-Unfolding Nearly Flat Convex Caps

Joseph O'Rourke*

August 31, 2017

Abstract

This paper proves a conjecture from [LO17]: A nearly flat, acutely triangulated convex cap \mathcal{C} has an edge-unfolding to a non-overlapping polygon in the plane. “Nearly flat” means that every face normal forms a sufficiently small angle with the z -axis. Although the result is not surprising, the proof relies on some recently developed concepts, angle-monotone and radially monotone curves.

1 Introduction

Let \mathcal{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* \mathcal{C} of angle Φ to be $\mathcal{C} = \mathcal{P} \cap H$ for some \mathcal{P} and H , such that $\phi(f) \leq \Phi$ for all f in \mathcal{C} . We will only consider $\Phi < 90^\circ$, which implies that the projection C of \mathcal{C} onto the xy -plane is one-to-one. Note that \mathcal{C} is not a closed polyhedron; it has no “bottom,” but rather a boundary $\partial\mathcal{C}$.

Say that a convex cap \mathcal{C} is *acutely triangulated* if every angle of every face is strictly acute, i.e., less than 90° . Note that \mathcal{P} being acutely triangulated does not always imply that $\mathcal{C} = \mathcal{P} \cap H$ is acutely triangulated. But any convex cap can be acutely triangulated; see Section 9.1. We allow flat (π) dihedral angles along edges, which could result from partitioning an obtuse triangle into several acute triangles.

An *edge unfolding* of a convex cap \mathcal{C} is a cutting of edges of \mathcal{C} that permits \mathcal{C} to be developed to the plane as a simple (non-self-intersecting) polygon, a “net.” The cut edges must form a boundary-rooted spanning forest \mathcal{F} : A forest of trees, each rooted on the boundary rim $\partial\mathcal{C}$, and spanning the internal vertices of \mathcal{C} .

Our main result is:

Theorem 1 *Every acutely triangulated convex cap \mathcal{C} with face normals bounded by a sufficiently small angle Φ from the vertical, has an edge unfolding to a non-overlapping polygon in the plane.*

*Department of Computer Science, Smith College, Northampton, MA, USA. jorourke@smith.edu.

An example is shown in Fig. 1, and another example later in Fig. 24. Loose bounds (Eq. 9) suggest that “sufficiently small” means that, e.g., for 86° acuteness, $\Phi < 5^\circ$ suffices.




Figure 1: (a) A convex cap of 98 vertices, $\Phi \approx 33^\circ$, with spanning forest \mathcal{F} marked. \mathcal{C} is non-obtusely triangulated (rather than acutely triangulated). (b) Edge unfolding by cutting \mathcal{F} . The quadrant lines are explained in Section 4.2.

It is a long standing open problem whether or not every convex polyhedron has a non-overlapping edge-unfolding, often called Dürer’s problem [DO07] [O’R13]. Theorem 1 can be viewed as an advance on a very narrow version of this problem. This theorem (without the acuteness assumption) has been a folk-conjecture for many years, but a specific line of attack was conjectured in [LO17], and it is that line I follow for the proof.

One might loosely view Theorem 1 as obtaining an edge unfolding for sufficiently flat polyhedra. In that sense, it could serve as a counterpart to Ghomi’s result that sufficiently thin polyhedra have edge unfoldings [Gho14].

2 Overview of Proof

The proof relies on two results from earlier work: the angle-monotone spanning forest result in [LO17], and a radially monotone unfolding result in [O’R16]. However, the former result needs generalization, and the latter is more general (and complex) than needed here. So those results are incorporated and explained as needed to allow this paper to stand alone. It is the use of angle-monotone and radially monotone curves and their properties that constitute the main novelties.

The approach is very roughly: project, lift, develop. The proof outline has these seven high-level steps:

1. Project \mathcal{C} to the plane containing its boundary rim, resulting in a triangulated convex region C . For sufficiently small Φ , C is again acutely triangulated.
2. Generalizing the result in [LO17], there is a θ -angle-monotone, boundary rooted spanning forest F of C , for $\theta < 90^\circ$. F lifts to a spanning forest \mathcal{F} of the convex cap \mathcal{C} .
3. For sufficiently small Φ , both sides L and R of each cut-path \mathcal{Q} of \mathcal{F} are θ -angle-monotone when developed in the plane, for some $\theta < 90^\circ$.
4. Any planar angle-monotone path for an angle $\leq 90^\circ$, is radially monotone, a concept from [O'R16].
5. Radial monotonicity of L and R , and sufficiently small Φ , imply that L and R do not cross in their planar development. This is a simplified version of a theorem from [O'R16], and here extended to trees.
6. Extending the cap \mathcal{C} to an unbounded polyhedron \mathcal{C}^∞ ensures that the non-crossing of each L and R extends arbitrarily far in the planar development.
7. The development of \mathcal{C} can be partitioned into θ -monotone “strips,” whose side-to-side development layout guarantees non-overlap in the plane.

We now proceed to detail the steps of the proof. I have decided to quantify steps even if they are in some sense obvious. Various quantities go to zero as $\Phi \rightarrow 0$. Quantifying by explicit calculation the dependence on Φ lengthens the proof. But I hope the details both solidify the proof and will assist exploring extensions suggested in Section 11. The reader willing to accept $\Phi \rightarrow 0$ arguments will be invited to skip these calculations.

2.1 Notation

We attempt to distinguish between objects in \mathbb{R}^3 , and planar projected versions of those objects, either by using calligraphy (\mathcal{C} in \mathbb{R}^3 vs. C in \mathbb{R}^2), or primes (γ in \mathbb{R}^3 vs. γ' in \mathbb{R}^2), and occasionally both (Q' vs. \mathcal{Q}). Sometimes this seems infeasible, in which case we use different symbols (u_i in \mathbb{R}^3 vs. v_i in \mathbb{R}^2). Sometimes we use \perp as a subscript to indicated projections or developments of lifted quantities.

We assume the plane containing $\partial\mathcal{C}$ is the xy -plane, $z = 0$.

3 Projection Angle Distortion

1. Project \mathcal{C} to the plane containing its boundary rim, resulting in a triangulated convex region C . For sufficiently small Φ , C is again acutely triangulated.

This first claim is obvious: since every triangle angle is strictly less than 90° , and the distortion due to projection to a plane goes to zero as \mathcal{C} becomes more flat, for some sufficiently small Φ , the acute triangles remain acute under projection.

In order to obtain a definite dependence on Φ , this section derives the following exact bound:

Lemma 1 *The maximum absolute value of the distortion Δ_\perp of any angle in \mathbb{R}^3 projected to the xy -plane, with respect to the tilt ϕ of the plane of that angle with respect to z , is given by:*

$$\Delta_\perp(\phi) = \cos^{-1} \left(\frac{\sin^2 \phi}{\sin^2 \phi - 2} \right) - \pi/2 . \quad (1)$$

For small Φ , for a convex cap with every normal satisfying $\phi \leq \Phi$, the expression becomes

$$\Delta_\perp(\Phi) \approx \Phi^2/2 - \Phi^4/12 + O(\Phi^5) , \quad (2)$$

and in particular, $\Delta_\perp(\Phi) \rightarrow 0$ as $\Phi \rightarrow 0$.

Readers willing to accept the import of this lemma may skip to Section 4.¹

3.1 Notation

The angle α in \mathbb{R}^3 is determined by two unit vectors a and b . The normal vector $\hat{n} = a \times b$ is tilted ϕ from the z -axis, and spun θ about that axis. See Fig. 2. In this figure, θ is chosen to bisect α , which we will see achieves the maximum distortion. Let primes indicate projections to the xy -plane. So a, b, α project to a', b', α' . Finally the distortion is $\Delta_\perp(\alpha, \phi, \theta) = |\alpha' - \alpha|$; it is about 8° in Fig. 2.

3.2 $\theta = \alpha/2$

Fixing α and ϕ , we first argue that the maximum distortion is achieved when θ bisects α (as it does in Fig. 2.)

Fig. 3 shows Δ_\perp as a function of θ . One can see it is a shifted and scaled sine wave with a period of π . This remains true over all α and all ϕ .

Proposition 1 *For fixed ϕ and α , the maximum distortion $\Delta_\perp(\alpha, \phi, \theta)$ is achieved with $\theta = \alpha/2$ (and $\alpha = \alpha/2 + \pi$).*

We leave this as a claim. An interesting aside is that for $\theta = \alpha/2 \pm \pi/4$, $\alpha' = \alpha$ and $\Delta_\perp = 0$.

So now we have reduced $\Delta_\perp(\alpha, \phi, \theta)$ to depending only on two variables: $\Delta_\perp(\alpha, \phi) = \Delta_\perp(\alpha, \phi, \alpha/2)$.

¹It seems likely that these calculations, and those in Section 6.1, are known, but I could not locate them in the literature.




Figure 2: $\phi = 30^\circ$, $\alpha = 70^\circ$ is determined by a and b . The projection to the xy -plane (green) results in a larger angle, $\alpha' = 78^\circ$.




Figure 3: For $\phi = 30^\circ$ and $\alpha = 70^\circ$, the maximum Δ_{\perp} is achieved at $\theta = \alpha/2 = 35^\circ$.

3.3 Right Angles Worst

Next we show that the maximum distortion $\Delta_{\perp}(\alpha, \phi)$ occurs when $\alpha = 90^{\circ}$. Fig. 4 plots Δ_{\perp} over the full range of α for a portion of the $\phi \in [0, 90^{\circ}]$ range of ϕ .




Figure 4: Δ_{\perp} as a function of α and $\phi \in [0, 30^{\circ}]$, showing the maximum distortion occurs at $\alpha = 90^{\circ}$.

Proposition 2 *The maximum distortion $\Delta_{\perp}(\alpha, \phi)$ occurs when $\alpha = 90^{\circ}$, for all $\phi \in [0, 90^{\circ}]$.*

3.4 Δ_{\perp} as a function of ϕ

Propositions 1 and 2 reduce Δ_{\perp} to a function of just ϕ , the tilt of the angle α in \mathbb{R}^3 . Those lemmas permit an explicit derivation² of this function:

$$\Delta_{\perp}(\phi) = \cos^{-1} \left(\frac{\sin^2 \phi}{\sin^2 \phi - 2} \right) - \pi/2 . \quad (1)$$

See Fig. 5. Note that $\Delta_{\perp}(0) = \cos^{-1}(0) - \pi/2 = 0$, as claimed earlier. Thus we have established Lemma 1 quoted at the beginning of this section.

For small Φ , the expression becomes

$$\Delta_{\perp}(\Phi) \approx \Phi^2/2 - \Phi^4/12 + O(\Phi^5) . \quad (2)$$

A few explicit values:

$$\begin{aligned} \Delta_{\perp}(10^{\circ}) &\approx 0.9^{\circ} \\ \Delta_{\perp}(20^{\circ}) &\approx 3.6^{\circ} \\ \Delta_{\perp}(30^{\circ}) &\approx 8.2^{\circ} \end{aligned}$$

²This derivation is not difficult but is likely of little interest, so it is not included.




Figure 5: $\Delta_{\perp}(\phi)$.

4 Angle-Monotone Spanning Forest

2. Generalizing the result in [LO17], there is a θ -angle-monotone, boundary rooted spanning forest F of C , for $\theta < 90^\circ$. F lifts to a spanning forest \mathcal{F} of the convex cap \mathcal{C} .

First we define angle-monotone paths, which originated in [DFG15] and were further developed in [BBC⁺16], and then turn to the spanning forests we need here.

4.1 Angle-Monotone Paths

Let C be a planar, triangulated convex domain, with ∂C its boundary, a convex polygon. Let G be the (geometric) graph of all the triangulation edges in C and on ∂C .

Define the θ -wedge $W(\beta, v)$ to be the region of the plane bounded by rays at angles β and $\beta + \theta$ emanating from v . W is closed along (i.e., includes) both rays, and has angular *width* of θ . A polygonal path $Q = (v_0, \dots, v_k)$ following edges of G is called θ -angle-monotone (or θ -monotone for short) if the vector of every edge (v_i, v_{i+1}) lies in $W(\beta, v_0)$ (and therefore $Q \subseteq W(\beta, v_0)$), for some β .³ Note that if $\beta = 0^\circ$ and $\theta = 90^\circ$, then a θ -monotone path is both x - and y -monotone, i.e., it meets every vertical, and every horizontal line in a point or a segment, or not at all.

³My notation here is slightly different from the notation in [LO17] and earlier papers.

4.2 Angle-Monotone Spanning Forest

It was proved in [LO17] that every nonobtuse triangulation G of a convex region C has a boundary-rooted spanning forest F of C , with all paths in F 90° -monotone. We describe the proof and simple construction algorithm before detailing the changes necessary for acute triangulations.

Some internal vertex q of G is selected, and the plane partitioned into four 90° -quadrants Q_0, Q_1, Q_2, Q_3 by orthogonal lines through v . Each quadrant is closed along one axis and open on its counterclockwise axis; q is considered in Q_0 and not in the others, so the quadrants partition the plane. It will simplify matters later if we orient the axes so that no vertex except for q lies on the axes, which is clearly always possible. Then paths are grown within each quadrant independently, as follows. A path is grown from any vertex $v \in Q_i$ not yet included in the forest F_i , stopping when it reaches either a vertex already in F_i , or ∂C . These paths never leave Q_i , and result in a forest F_i spanning the vertices in Q_i . No cycle can occur because a path is grown from v only when v is not already in F_i ; so v becomes a leaf of a tree in F_i . Then $F = F_1 \cup F_2 \cup F_3 \cup F_4$.

We cannot follow this construction exactly in our situation of an acute triangulation G , because the “quadrants” for θ -monotone paths for $\theta = 90^\circ - \Delta\theta < 90^\circ$ cannot cover the plane exactly: They leave a thin $4\Delta\theta$ angular gap; call the cone of this aperture g . We proceed as follows. Identify an internal vertex q of G so that it is possible to orient the cone-gap g , apexed at q , so that g contains no internal vertices of G . See Fig. 6 for an example. Then we proceed just as in [LO17]: paths are grown within each Q_i , forming four forests F_i , each composed of θ -monotone paths.

It remains to argue that there always is such a q at which to apex cone-gap g . Although it is natural to imagine q as centrally located (as in Fig. 6), it is possible that G is so dense with vertices that such a central location is not possible. However, it is clear that the vertex q that is closest to ∂C will suffice: aim g along the shortest path from q to ∂C . Then g might include several vertices on ∂C , but it cannot contain any internal vertices of G , as they would be closer to ∂C . Again we could rotate the axes slightly so that no vertex except for q lies on an axis.

We conclude this section with a lemma:

Lemma 2 *If G is an acute triangulation of a convex region C , then there exists a boundary-rooted spanning forest F of C , with all paths in F θ -angle-monotone, for $\theta = 90^\circ - \Delta\theta < 90^\circ$.*

That F lifts to a spanning forest \mathcal{F} of the convex cap \mathcal{C} is immediate. What is not straightforward is establishing the requisite properties of \mathcal{F} .

5 Curvature

For sufficiently small Φ , ...




Figure 6: Here the near-quadrants Q_i have width $\theta = 87^\circ$, so the gap g has angle $4\Delta\theta = 12^\circ$.

The “nearly flatness” of the convex cap \mathcal{C} is controlled by Φ , the maximum angle deviation of face normals from the z -axis. Let ω_i be the curvature at internal vertex $u_i \in \mathcal{C}$, and $\Omega = \sum_i \omega_i$ the total curvature. In this section we bound Ω as a function of Φ . We will see that the reverse is not possible: even a small Ω could be realized with large Φ . The bound is given in the following lemma:

Lemma 3 *The total curvature $\Omega = \sum_i \omega_i$ of \mathcal{C} satisfies*

$$\Omega \leq 2\pi(\cos \Phi - 1) \approx \pi\Phi^2 - \frac{\pi\Phi^4}{12} + O(\Phi^5).$$

The reader uninterested in the derivation may skip to Section 6.

One way to prove Lemma 3 is to argue that Ω is at most the curvature at the apex of a cone with lateral normal Φ .⁴ Here we opt for another approach which makes it clear why we cannot bound Φ in terms of Ω .

The proof depends on the *Gaussian sphere* representation, a graph G_S on a unit-radius sphere S with nodes corresponding to each face normal, and arcs corresponding to the dihedral angle of the edge shared by adjacent faces. An example is shown in Fig. 7. For a convex polyhedron (and so for a convex cap \mathcal{C}), each vertex v of \mathcal{C} maps to a convex spherical polygon $s(v)$ whose area is the

⁴This route was followed in the first version of this paper.




Figure 7: Gaussian sphere and G_S for a convex polyhedron of 500 vertices.

curvature at v . Each internal angle β at a face node f of $s(v)$ is $\pi - \alpha$, where α is the face angle incident to v for face f . These basic properties of G_S are well-known; see, e.g., [BLS07].

For a vertex v , the largest area of its spherical polygon $s(v)$ is achieved when that polygon approaches a circle. So an upperbound on the area, and so the curvature, is the area of a disk of radius Φ . This is the area of a spherical cap, which is $\Omega = 2\pi(\cos \Phi - 1)$. For small Φ , the area is nearly that of a flat disk, $\pi\Phi^2$. This establishes Lemma 3.

The reason that we cannot bound Φ in terms of Ω is that it is possible that the $s(v)$ is long and thin, as in Fig. 8. Then the area Ω can be small while the maximum Φ deviation is large. In the discussion in Section 11, we call such vertices “oblong.”




Figure 8: Spherical polygon of an “oblong” vertex: $|\phi(a) - \phi(b)|$ is large but Ω is small.

6 Curve Distortion

3. For sufficiently small Φ , both sides L and R of each cut-path Q of \mathcal{F} are θ -angle-monotone when developed in the plane, for some $\theta < 90^\circ$.

This step says, essentially, each θ -monotone path Q' in the planar projection

is not distorted much when lifted to \mathcal{Q} on \mathcal{C} .⁵ This is obviously true as $\Phi \rightarrow 0$, but it requires proof. We need to establish that the left and right sides of the cut \mathcal{Q} develop to the plane as still θ -monotone paths for some (different) $\theta \leq 90^\circ$.

For intuition only, we offer an example in Fig. 9. Here an 86° -monotone staircase Q' lies in the xy -plane. It is lifted to a sphere to \mathcal{Q} , with the sphere standing for the convex cap \mathcal{C} . Calculating what could be the angles incident to the left of \mathcal{Q} were the sphere a triangulated polyhedron, we develop \mathcal{Q} to \mathcal{Q}_\perp in the plane. \mathcal{Q}_\perp is distorted compared to Q' , but by at most 3.8° , so it remains θ -monotone for $\theta < 90^\circ$.




Figure 9: Q' is a (blue) path on the plane, θ -monotone for $\theta = 86^\circ$. Its lift to the sphere is \mathcal{Q} (red). \mathcal{Q}_\perp (green) is distorted, but remains acute.

Our proof uses the Gauss-Bonnet theorem, in the form $\tau + \omega = 2\pi$: the turn of a closed curve plus the curvature enclosed is 2π . See, for example, Lee's description [Lee06, Thm.9.3, p.164].⁶ To bound the curve distortion of Q' , we need to bound the distortion of pieces of a closed curve that includes Q' as a subpath. Our argument here is not straightforward, but the conclusion is that,

⁵We use Q' to emphasize it lies in the xy -plane.

⁶My τ is Lee's κ_N .

as $\Phi \rightarrow 0$, the distortion also goes to 0. The reader willing to accept this claim may skip to Lemma 5 below, which provides an explicit bound (Eq. 7).

The reason the proof is not straightforward is that Q' could have an arbitrarily large number n of vertices, so bounding the angle distortion at each by Δ_{\perp} would lead to arbitrarily large distortion $n\Delta_{\perp}$. The same holds for the rim, as we will see below in Section 6.2. So global arguments that do not cumulate errors seem necessary.

6.1 Angle Lifting to \mathcal{C}

We need to supplement the angle distortion calculations presented in Section 3 for a very specific bound on the total turn of the rims of C and of \mathcal{C} . In particular, here we are concerned with the sign of the distortion. Let $R' = \partial C$ and $R = \partial \mathcal{C}$ be the rims of the planar C and of the convex cap \mathcal{C} , respectively. Again we use primes to indicate projections to the plane. Note that $R' = R$ geometrically, but we will focus on the neighborhoods of these rims on C and \mathcal{C} , which are different.

Let $\Delta(a, b, c')$ be a triangle in the xy -plane, and c a point vertically above c' . We compare the angle $\alpha' = \angle c', a, b$ with angle $\alpha = \angle c, a, b$; see Fig. 10(a). (Note that, in contrast to the arbitrary-angle analysis in Section 3, here the two triangles share a side.) We start with the fact that the area of the projected $\Delta(c', a, b)$ is $\cos \phi$ times the area of the 3D $\Delta(c, a, b)$, where ϕ is the angle the normal to triangle $\Delta(c, a, b)$ makes with the z -axis. Defining $B = b - a$, $C = c - a$ and $C' = c' - a$, we have

$$\begin{aligned} B \cdot C' &= |B||C'| \cos \alpha' \\ B \cdot C &= |B||C| \cos \alpha \end{aligned}$$

so

$$\begin{aligned} |B||C'| \cos \alpha &= |B||C'| \cos \alpha' \cos \phi \\ \frac{\cos \alpha}{\cos \alpha'} &= \frac{|C'|}{|C|} \cos \phi \\ &\leq \frac{|C'|}{|C|} \cos \phi \\ &\leq 1 \\ \cos \alpha &\leq \cos \alpha' \\ \alpha &\geq \alpha' \text{ when } \alpha' \leq 90^\circ \\ \alpha &\leq \alpha' \text{ when } \alpha' \geq 90^\circ \end{aligned}$$

The last step follows because either both $\alpha', \alpha \leq 90^\circ$ or $\alpha', \alpha \geq 90^\circ$. The conclusion is that the 3D angle α is smaller for obtuse α' , and larger for acute α' . When $\alpha' = 90^\circ$, then $\alpha = 90^\circ$.

Now we turn to the general situation of three consecutive vertices (d, a, b) of the rim, and how the planar angle at a differs from the 3D angles of the

incident triangles of \mathcal{C} . We start assuming just two triangles are incident to a , sharing an edge ca . Note that, if there is no edge of $\mathcal{C} \setminus R$ incident to a , then the triangle $\triangle(d, a, b)$ is a face of \mathcal{C} , which implies (by convexity) that the cap \mathcal{C} is completely flat, and there is nothing to prove.

So we start with the situation depicted in Fig. 10(b). We seek to show that the 2D angle α' at a , $\angle d, a, b$, is always at most the 3D angle α , which is the sum of $\alpha_1 = \angle c, a, b$ and $\alpha_2 = \angle c, a, d$. Note that the projection of one of these angles could be smaller and the other larger than their planar counterparts, so it is not immediately obvious that the sum is always larger. But we can see that it as follows. If $\alpha' \leq 90^\circ$, then we know that both of the 3D angles α_1 and α_2 are larger by our previous analysis. If $\alpha' \geq 90^\circ$, then partition $\alpha' = 90^\circ + \beta$, where $\beta < 90^\circ$. Then we have that $\alpha_2 = 90^\circ$ and $\alpha_1 > \beta$, so again $\alpha_1 + \alpha_2 = \alpha \geq \alpha'$.

The general situation is that a vertex a on the rim of the cap \mathcal{C} will have several incident edges, rather than just the one ca that we used above. Continue to use the notation that d, a, b are consecutive vertices of R' , but now edges c_1, \dots, c_k of \mathcal{C} are incident to a . Consider two consecutive triangles $\triangle(c_{i-1}, a, c_i)$ and $\triangle(c_i, a, c_{i+1})$ of \mathcal{C} . These sit over a triangle $\triangle(c_{i-1}, a, c_{i+1})$ which is not a face of \mathcal{C} ; rather it is below \mathcal{C} (by convexity). The argument used above shows that the sum of the two triangle's angles at a are at least the internal triangle's angle at a . Repeating this argument shows that, in the general situation, the sum of all the incident face angles of \mathcal{C} is greater than or equal to the 2D angle $\alpha' = \angle d, a, b$.

We summarize in a lemma:⁷

Lemma 4 *The planar angle α' at a vertex of the rim lifts to 3D angles of the incident triangles of the cap \mathcal{C} , whose sum α satisfies $\alpha \geq \alpha'$.*




Figure 10: (a) Lifting α' to α according to a right tetrahedron. (b) Lifting the planar angle at a to the sum of two 3D angles.

6.2 Rim-Turn Bound

Now we use Lemma 4 to bound the total turn of the rim R of \mathcal{C} and R' of \mathcal{C}' . Although the rims are geometrically identical, their turns are not. The turn at

⁷It seems likely this is either known, or can be proved with a simpler argument.

each vertex a' of the planar rim R' is $2\pi - \alpha'$, where α' is the planar angle of C incident to a' . The turn at each vertex a of the 3D rim R is $2\pi - \alpha$, where α is the sum of the angles of the cap C incident to a . By Lemma 4, $\alpha \geq \alpha'$, so the turn at each vertex of the 3D rim R is smaller than or equal to the turn at each vertex of the 2D rim R' . Therefore the total turn of the 3D rim τ_R is smaller than or equal to the total turn of the 2D rim $\tau_{R'}$. And Gauss-Bonnet allows us to quantify this:

$$\begin{aligned}\tau_{R'} &= 2\pi \\ \tau_R + \Omega &= 2\pi \\ \tau_{R'} - \tau_R &= \Omega\end{aligned}$$

For any subportion of the rims $r' \subset R'$, $r \subset R$, Ω serves as an upper bound, because we know the sign of the difference is the same at every vertex of r' , r :

$$\tau_{r'} - \tau_r \leq \Omega \quad (3)$$

Note that the reason we could make this inference from R to $r \subset R$ is because we know the sign of the difference at every vertex of the rim: The planar rim turns more than the 3D rim at every vertex. Were the signs unknown, cancellation would have prevented this inference.

6.3 Turn Distortion of γ'

We first walk through a calculation that will serve as a “warm-up” for the calculation actually needed. I found matters complicated enough to warrant this approach.

Let γ' be a simple curve in the xy -plane. We aim to bound the total turn difference $\Delta\gamma$ between γ' and its lift γ to the cap \mathcal{C} . Let $r' \subset R'$ be the portion of the rim counterclockwise from b to a , so that $\gamma' \cup r'$ is a closed curve. Of course in the plane there is no curvature enclosed, and the total turn of this closed curve is 2π . We describe this total turn τ' in four pieces: the turn of γ' , the turn of r' , and the turn at the join points:

$$\begin{aligned}\tau' &= \tau_{\gamma'} + (\tau_{a'} + \tau_{b'}) + \tau_{r'} \\ &= 2\pi\end{aligned} \quad (4)$$

where $\tau_{a'}$ and $\tau_{b'}$ are the turn angles at a' and b' . See Fig. 11(a).

Now we turn to the convex cap \mathcal{C} , as illustrated in Fig. 11(b). We have a similar expression for τ , but now the Gauss-Bonnet theorem applies: $\tau + \omega = 2\pi$, where $\omega \leq \Omega$ is the total curvature inside the path $\gamma \cup r$:

$$\begin{aligned}\tau + \omega &= \tau_{\gamma} + (\tau_a + \tau_b) + \tau_r + \omega \\ &= 2\pi\end{aligned} \quad (5)$$

Combining Eqs. 5 and 6,

$$\begin{aligned}\tau_{\gamma'} + \tau_{a'} + \tau_{b'} + \tau_{r'} &= \tau_{\gamma} + \tau_a + \tau_b + \tau_r + \omega \\ \tau_{\gamma'} - \tau_{\gamma} &= (\tau_a - \tau_{a'}) + (\tau_b - \tau_{b'}) + (\tau_r - \tau_{r'}) + \omega\end{aligned} \quad (6)$$

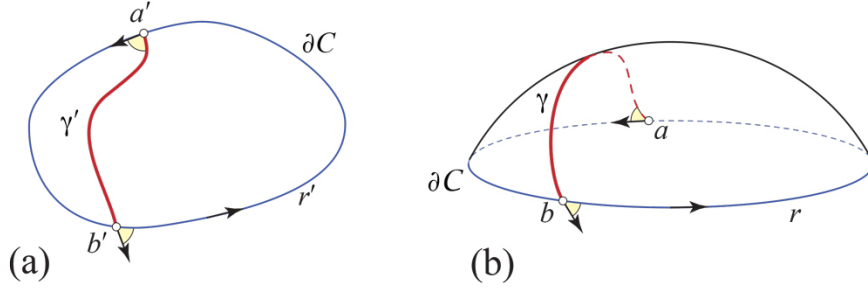


Figure 11: (a) C , the projection of the cap \mathcal{C} . (b) γ is the lift of γ' to \mathcal{C} .

Our goal is to bound $\Delta\gamma = |\tau_{\gamma'} - \tau_\gamma|$, the total distortion of the turn of γ compared to that of γ' ; the sign of the distortion is not relevant.

The turn angles at a and b are both distorted by at most Δ_\perp :

$$\begin{aligned} |\tau_a - \tau_{a'}| &\leq \Delta_\perp \\ |\tau_b - \tau_{b'}| &\leq \Delta_\perp \end{aligned}$$

Note that the analysis in Section 6.1 shows that the sign of these angle changes could be positive or negative, depending on whether γ' meets R' in an acute or obtuse angle. So we bound the absolute magnitude.

From Eq. 3, we have $|\tau_r - \tau_{r'}| \leq \Omega$. Here we do know the sign of the difference, but we only use that sign to bound $r \subseteq R$.

Using these bounds in Eq. 7 leads to

$$\begin{aligned} \Delta\gamma &= |\tau_\gamma - \tau_{\gamma'}| \\ &\leq 2\Omega + 2\Delta_\perp \end{aligned}$$

Example. Before moving to the next calculation, we illustrate the preceding with a geometrically accurate example, the top of a regular icosahedron, shown in Fig. 12. Here $\Phi = 37.4^\circ$ and $\Omega = 60^\circ$. (Eq. ?? for this Φ yields $\Omega = 73.9^\circ$, an upperbound on the true Ω .) γ' is the (a, d) chord of the pentagon rim, which lifts to $\gamma = (a, b, c, d)$ on \mathcal{C} . Both the turn at each rim r' vertex, and $\tau_{a'} = \tau_{d'}$, is 72° . So the Gauss-Bonnet theorem for the planar circuit is

$$\begin{aligned} \tau_{\gamma'} + (\tau_{a'} + \tau_{d'}) + \tau_{r'} &= 2\pi \\ 0 + (72^\circ + 72^\circ) + 3(72^\circ) &= 360^\circ \end{aligned}$$

The 3D turns $\tau_a = \tau_d$ are slightly larger, 75.5° (consistent with the analysis in Section 6.1), and the turn at each rim r vertex is smaller, 60° (consistent with the analysis in Section 6.2). We use the Gauss-Bonnet theorem to solve for τ_γ :

$$\begin{aligned} \tau_\gamma + (\tau_a + \tau_d) + \tau_r + \omega &= 2\pi \\ \tau_\gamma + (75.5^\circ + 75.5^\circ) + 3(60^\circ) + 60^\circ &= 360^\circ \\ \tau_\gamma &= -31.0^\circ \end{aligned}$$




Figure 12: Icosahedron cap. $\gamma' = Q' = (a, d)$. $\gamma = Q = (a, b, c, d)$.

And indeed, $\tau_b = \tau_c = -15.5^\circ$. One can see that the final link of the developed chain γ_\perp has turned 31° with respect to the planar chord, much smaller than the crude bound of $2\Omega + 2\Delta_\perp$ derived in the previous section.

6.4 Turn Distortion of Q'

What we need is not $\Delta\gamma$ but ΔQ , where Q' is any prefix of an angle-monotone path in C . The reason for prefix here is that we want to bound the turn of any segment of Q , not just the last segment, whose turn is $\sum_i \tau_i$.

It is tempting to say that Q' is a subset of some γ' , and so the bound we derived above holds for Q' as well, say, by extending Q' until it hits ∂C , and so is indeed a chord of ∂C . But the problem is that there can be cancellations among the τ_i , as we have no guarantee that they are all the same sign. (It is a counterintuitive fact that the development of a “slice” curve formed by the intersection of a plane with a convex polyhedron, may turn both left and right: [O’R03].) So it is conceivable that Q turns a lot but is canceled out by the turn of the extension. Then the total turn is not a bound on ΔQ , the turn of Q . So we take a more complex approach.

Let Q' be a prefix of an angle-monotone path in C , Q its lift to the 3D cap \mathcal{C} . We seek to bound $\Delta Q = |\tau'_{Q'} - \tau_Q|$. The key idea is to reduce \mathcal{C} to a *half-cap* as follows; see Fig. 13. Let a' be the endpoint of Q' not on ∂C , interior to C' . Slice the cap \mathcal{C} with a vertical plane orthogonal to the clockwise ray of the θ -wedge containing Q' . This choice guarantees that Q' and therefore Q are to one side of the plane. We now study the closed curve $Q \cup r_1 \cup r_2$, where $r_1 \subset R$ is a portion of the original rim of \mathcal{C} , and r_2 is a portion of the rim created by the vertical plane slicing \mathcal{C} . These three curves meet at points a, b, c as illustrated.




Figure 13: A “half-cap” to bound ΔQ .

As before, we apply the Gauss-Bonnet theorem and note all the contributions:

$$\begin{aligned}\tau_{Q'} + (\tau_{a'} + \tau_{b'} + \tau_{c'}) + (\tau_{r_1'} + \tau_{r_2'}) &= 2\pi \\ \tau_Q + (\tau_a + \tau_b + \tau_c) + (\tau_{r_1} + \tau_{r_2}) + \omega &= 2\pi\end{aligned}$$

Again we use primes to indicate planar quantities. Note in particular that $\tau_{r_2'}$ is the turn of the rim portion r_2 in the vertical plane, whereas $\tau_{r_1'}$ is the turn of r_1 in the xy -plane as before. Combining these equations leads to

$$\begin{aligned}\tau_{Q'} - \tau_Q &= (\tau_a - \tau_{a'}) + (\tau_b - \tau_{b'}) + (\tau_c - \tau_{c'}) + (\tau_{r_1} - \tau_{r_1'}) + (\tau_{r_2} - \tau_{r_2'}) + \omega \\ \Delta Q &< 3\Delta_{\perp} + 2\Omega + \Omega = 3(\Delta_{\perp} + \Omega)\end{aligned}\quad (7)$$

The logic is as before: Each of the turn distortions at a, b, c is at most Δ_{\perp} , both rim distortions are at most Ω by Eq. 3, and $\omega < \Omega$.

Using the small- Φ bounds derived earlier in Eqs. 2 and ??:

$$\begin{aligned}|\Delta Q| &\leq 3(\Delta_{\perp} + \Omega) \\ &\approx 3\Phi^2(\pi + 1/2)\end{aligned}\quad (8)$$

Thus we have $\Delta Q \rightarrow 0$ as $\Phi \rightarrow 0$, as expected.

To summarize these calculations:

Lemma 5 *The difference in the total turn of any prefix of \mathcal{Q} from its planar projection \mathcal{Q}' is bounded by Eq. 7, which Eq. 8 shows to be a constant times Φ^2 (for small Φ). Therefore, this turn goes to zero as $\Phi \rightarrow 0$.*

These bounds are likely quite loose. A few explicit values:

$$\begin{aligned}\Phi = 1^\circ &\rightarrow 3\Phi^2(\pi + 1/2) \approx 0.2^\circ \\ \Phi = 2^\circ &\rightarrow 3\Phi^2(\pi + 1/2) \approx 0.8^\circ \\ \Phi = 3^\circ &\rightarrow 3\Phi^2(\pi + 1/2) \approx 1.7^\circ \\ \Phi = 5^\circ &\rightarrow 3\Phi^2(\pi + 1/2) \approx 4.8^\circ\end{aligned}$$

We finally return to the claim at the start of this section:

3. For sufficiently small Φ , both sides L and R of each path \mathcal{Q} of \mathcal{F} are θ -angle-monotone when developed in the plane, for some $\theta < 90^\circ$.

The turn at any vertex of \mathcal{Q} is determined by the incident face angles to the left following the orientation shown in Fig. 11, or to the right reversing that orientation (clearly the curvature enclosed by either curve is $\leq \Omega$). These incident angles determine the left and right planar developments, L and R , of \mathcal{Q} . Because we know that Q' is θ -angle-monotone for $\theta < 90^\circ$, there is some finite “slack” $\Delta\theta = 90^\circ - \theta$. Because Lemma 5 established a bound for any prefix of Q' , it bounds the turn distortion of each edge of \mathcal{Q} , which we can arrange to fit inside that slack. So the bound provided by Lemma 5 suffices to guarantee that:

Lemma 6 *For sufficiently small Φ , both L and R remain θ -angle-monotone for some (larger) θ , but still $\theta < 90^\circ$.*

Using the small- Φ approximation in Eq. 8 leads to

$$\Phi \leq \sqrt{\Delta\theta} \sqrt{\frac{2}{3 + 6\pi}} \approx 0.3\sqrt{\Delta\theta} \quad (9)$$

For example, if all triangles are acute by $\Delta\theta = 3^\circ$, then $\Phi \approx 4.0^\circ$ suffices.

7 Radially Monotone Paths

The next step in the proof is:

4. Any planar angle-monotone path for an angle $\leq 90^\circ$, is radially monotone, a concept from [O’R16].

To establish this claim, and remain independent of [O’R16], we repeat definitions in that report, sometimes directly quoting from that report’s Section 1.

Let C be a planar, triangulated convex domain, with ∂C its boundary, a convex polygon. Let⁸ $Q = (v_0, v_1, v_2, \dots, v_k)$ 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 offer three equivalent definitions of radial monotonicity.

(1). We say that $Q = (v_0, v_1, \dots, v_k)$ is *radially monotone* with respect to (w.r.t.) v_0 if the distances from v_0 to all points of Q are (non-strictly) monotonically increasing. (Note that requiring the distance to just the vertices of Q to be monotonically increasing is not the same as requiring the distance to

⁸In this section we dispense with the primes on symbols to indicate objects in the xy -plane, when there is little chance of ambiguity. We use v_i for vertices in the plane and u_i for their counterparts on the cap \mathcal{C} .

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} . It is an easy consequence of these definitions that, if Q is radially monotone, it is radially monotone w.r.t. any point p on Q , not only w.r.t. its vertices.

Before exploring this definition further, we discuss its intuitive motivation. If a path Q is radially monotone, 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 radially monotone, 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 radially monotone cut paths are locally (infinitesimally) opening “safe,” and non-radially monotone paths are potentially overlapping.

We now provide two more equivalent definitions of radial monotonicity to aid intuition.

(2). The condition for Q to be radially monotone w.r.t. v_0 can be interpreted as requiring Q to cross every circle centered on v_0 at most once; see Fig. 14. The concentric circles viewpoint makes it evident that infinitesimal rigid rotation of Q about v_0 to Q' ensures that $Q \cap Q' = \{v_0\}$, for each point of Q simply moves along its circle. Of course the concentric circles must be repeated, centered on every vertex v_i .

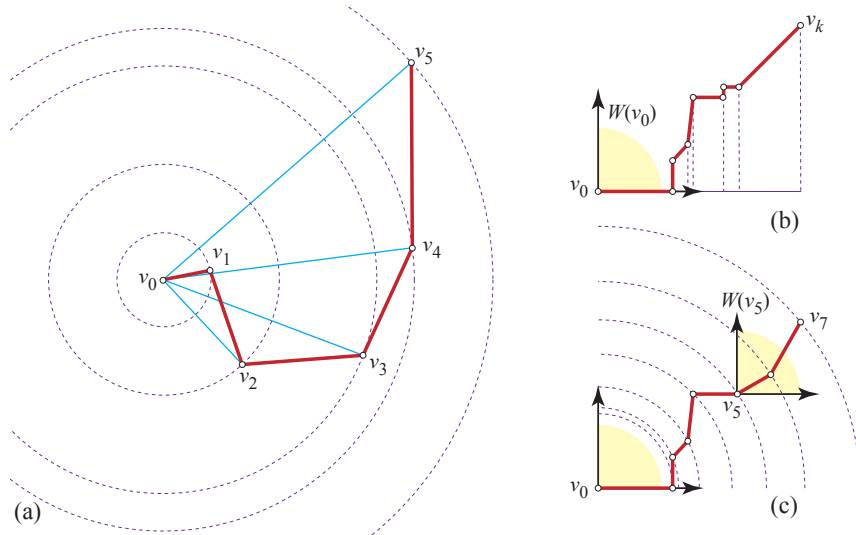


Figure 14: (a) A radially monotone chain, with its monotonicity w.r.t. v_0 illustrated. (b) A 90° -monotone chain, with x -monotonicity indicated. (c) Such a chain is also radially monotone.

(3). A third definition of radial monotonicity is as follows. Let $\alpha(v_i) = \angle(v_0, v_i, v_{i+1})$. Then Q is radially monotone w.r.t. v_0 if $\alpha(v_i) \geq \pi/2$ for all $i > 0$. For if $\alpha(v_i) < \pi/2$, Q violates monotonicity at v_i , and if $\alpha(v_i) \geq \pi/2$, then points along the segment (v_i, v_{i+1}) increase in distance from v_0 . Again this needs to hold for every vertex as the angle source, not just v_0 .

Radially monotone paths are the same⁹ as backwards “self-approaching curves,” introduced in [IKL99] and used for rather different reasons.

Many properties of radially monotone paths are derived in [O’R16], but here we need only the above definitions.

7.1 Angle-monotone chains are radially monotone

Recall the definition of an angle-monotone path from Section 4.1.

Fig. 14(c) illustrates why a θ -monotone chain Q , for any $\theta \leq 90^\circ$, is radially monotone: the vector of each edge of the chain points external to the quarter-circle passing through each v_i . And so the chain intersects the v_0 -centered circles at most once (definition (2)), and the angle $\alpha(v_i) \geq 90^\circ$ (definition (3)). Thus Q is radially monotone w.r.t. v_0 . But then the argument can be repeated for each v_i , for the wedge $W(v_i)$ is just a translation of $W(v_0)$.

It should be clear that these angle-monotone chains are very special cases of radially monotone chains. But we rely on the spanning-forest theorem in [LO17] to yield angle-monotone chains, and we rely on the unfolding properties of radially monotone chains from [O’R16] to establish non-overlap. We summarize in a lemma:

Lemma 7 *A θ -monotone chain Q , for any $\theta \leq 90^\circ$, is radially monotone.*

8 Noncrossing L & R Developments

5. Radial monotonicity of L and R , and sufficiently small Φ , imply that L and R do not cross in their planar development. This is a simplified version of a theorem from [O’R16], and here extended to trees.

We will use $\mathcal{Q} = (u_0, u_1, \dots, u_n)$ as a path of edges on \mathcal{C} , with each $u_i \in \mathbb{R}^3$ a vertex and each $u_i u_{i+1}$ an edge of \mathcal{C} . We use Q for a generic simple (non-self-intersecting) polygonal chain in the plane.

Let $Q = (v_0, v_1, \dots, v_n)$ be a chain. Define the *turn angle* τ_i at v_i to be the counterclockwise angle from $v_i - v_{i-1}$ to $v_{i+1} - v_i$. Thus $\tau_i = 0$ means that v_{i-1}, v_i, v_{i+1} are collinear, i.e., there is no turn at v_i . $\tau_i \in (-\pi, \pi)$; simplicity excludes $\tau_i = \pm\pi$.

Each turn of the chain Q sweeps out a sector of angles. We call the union of all these sectors $\Lambda(Q)$; this forms a cone such that, when apexed at v_0 , $Q \subseteq \Lambda(Q)$. The rays bounding $\Lambda(Q)$ are determined by the segments of Q

⁹Anna Lubiw, personal communication, July 2016.

at extreme angles; call these angles σ_{\max} and σ_{\min} . See ahead to Fig. 16 for examples. Let $|\Lambda(Q)|$ be the measure of the apex angle of the cone, $\sigma_{\max} - \sigma_{\min}$. We will assume that $|\Lambda(Q)| < \pi$ for our chains Q , although it is quite possible for radially monotone chains to have $|\Lambda(Q)| > \pi$. In our case, in fact $|\Lambda(Q)| < \pi/2$, but that tighter inequality is not needed for Theorem 2 below. The assumption $|\Lambda(Q)| < \pi$ guarantees that Q fits in a halfplane H_Q whose bounding line passes through v_0 .

Because σ_{\min} is turned to σ_{\max} , we have that the total absolute turn $\sum_i |\tau_i| \geq |\Lambda(Q)|$. But note that the sum of the turn angles $\sum_i \tau_i$ could be smaller than $|\Lambda(Q)|$ because of cancellations.

8.1 The left and right planar chains L & R

Again let $\mathcal{Q} = (u_0, u_1, \dots, u_n)$ with each u_i a vertex of \mathcal{C} . Let ω_i be the curvature at u_i . We view u_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 surface angle at u_i left of \mathcal{Q} , and ρ_i the surface angle right of \mathcal{Q} there. So $\lambda_i + \omega_i + \rho_i = 2\pi$, and $\omega_i \geq 0$. Define L to be the planar path from the origin with left angles λ_i , R the path with right angles ρ_i . These paths are the left and right planar developments of \mathcal{Q} . (Each of these paths are understood to depend on \mathcal{Q} : $L = L(\mathcal{Q})$ etc.) We label the vertices of the paths ℓ_i, r_i .

In [O’R16], the *medial path* M with left angles $\lambda_i - \omega_i/2$ (and therefore right angles $\rho_i + \omega_i/2$), played a significant role, in that a noncrossing result was proved when M is radially monotone, even if L and/or R is not radially monotone. Here we do not need M , and we make the stronger assumption that both L and R are radially monotone. In fact, both are θ -monotone for $\theta < 90^\circ$, but we do not use this fact to reach Theorem 2 below.

Define $\omega(\mathcal{Q}) = \sum_i \omega_i$, the total curvature along the path \mathcal{Q} . We will assume $\omega(\mathcal{Q}) < \pi$, a very loose constraint in our circumstances of a nearly flat convex cap \mathcal{C} . For example, with $\Phi = 30^\circ$, Ω for \mathcal{C} is $< \pi\Phi^2 \approx 49^\circ$, and $\omega(\mathcal{Q})$ can be at most Ω .

8.2 Left-of Definition

Let $A = (a_0, \dots, a_n)$ and $B = (b_0, \dots, b_n)$ be two radially monotone chains sharing $x = a_0 = b_0$. (Below, A and B will be the L and R chains.) Let $D(r)$ be the circle of radius r centered on x . Recall that $D(r)$ intersects any radially monotone chain in at most one point. Let a and b be two points on $D(r)$. Say that a is *left of* b , $a \preceq b$, if the counterclockwise arc from b to a is less than π . If $a = b$, then $a \preceq b$. Now we extend this relation to entire chains. Say that chain A is *left of* B , $A \preceq B$, if, for all $r > 0$, if $D(r)$ meets both A and B , in points a and b respectively, then $a \preceq b$. If $D(r)$ meets neither chain, or only one, no constraint is specified. Note that, if $A \preceq B$, A and B can touch but not properly cross.

8.3 Noncrossing Theorem

Theorem 2 *Let \mathcal{Q} be an edge cut-path on \mathcal{C} , and L and R the planar chains derived from \mathcal{Q} , as described above. Under the assumptions:*

1. *Both L and R are radially monotone,*
2. *The total curvature along \mathcal{Q} satisfies $\omega(\mathcal{Q}) < \pi$.*
3. *Both cone measures are less than π : $|\Lambda(L)| < \pi$ and $|\Lambda(R)| < \pi$,*

then $L \preceq R$: L and R may touch and share an initial chain from $\ell_0 = r_0$, but L and R do not properly cross, in either direction.

First we remark that the angle conditions (2) and (3) are necessary. Fig. 15(a) shows an example where they are violated and L crosses R from the right side of R to R 's left side. In this figure, $|\Lambda(L)| = |\Lambda(R)| = \pi$, because edge r_0r_1 points vertically upward and r_2r_3 points vertically downward, and similarly for L . Now suppose that $\omega_0 = \pi + \varepsilon$, and all other $\omega_i = 0$. So L is a rigid rotation of R about $\ell_0 = r_0$ by ω_0 , which allows L to cross R as illustrated. So some version of the angle conditions are necessary: we need that $|\Lambda(R)| + \omega(\mathcal{Q}) < 2\pi$ to prevent this type of “wrap-around” intersection, and conditions (2) and (3) meet this requirement. We address the sufficiency of these conditions in the proof below.

Proof: We first argue that L cannot wrap around as in Fig. 15(a) and cross R from its right side to its left side. Let ρ_{\max} be the counterclockwise bounding ray of $\Lambda(R)$. In order for L to enter the halfplane H_R containing $\Lambda(R)$, and intersect R from its right side, ρ_{\max} must turn to be oriented to enter H_R , a turn of $\geq \pi$. We can think of the effect of ω_i as augmenting R 's turn angles τ_i to L 's turn angles $\tau'_i = \tau_i + \omega_i$. Because $\omega_i \geq 0$ and $\omega(\mathcal{Q}) = \sum_i \omega_i < \pi$, the additional turn of the chain segments of R is $< \pi$, which is insufficient to rotate ρ_{\max} to aim into H_R . See Fig. 15(b). (Later (Section 9) we will see that we can assume L and R are arbitrarily long, so there is no possibility of L wrapping around the end of R and crossing R right-to-left.)

Next we show that L cannot cross R from left to right. We imagine \mathcal{Q} right developed in the plane, so that $Q = R$. We then view L as constructed from a fixed R by successively opening/turning the links of R by ω_i counterclockwise about r_i , with i running backwards from r_{n-1} to r_0 , the source vertex of R . Fig. 16(b) illustrates this process. Let $L_i = (\ell_i, \ell_{i+1}, \dots, \ell_n)$ be the resulting subchain of L after rotations $\omega_{n-1}, \dots, \omega_i$, and R_i the corresponding subchain of $R = (r_i, r_{i+1}, \dots, r_n)$, with $\ell_i = r_i$ the common source vertex. We prove $L_i \preceq R_i$ by induction.

$L_{n-1} \preceq R_{n-1}$ is immediate because $\omega_{n-1} \leq \omega(\mathcal{Q}) < \pi$; see Fig. 16(b). Assume now $L_{i+1} \preceq R_{i+1}$, and consider L_i ; refer to Fig. 17. Because both L_i and R_i are radially monotone, circles centered on $\ell_i = r_i$ intersect the chains in at most one point each. L_i is constructed by rotating L_{i+1} rigidly by ω_i counterclockwise about $\ell_i = r_i$; see Fig. 17(b). This only increases the arc distance between the intersections with those circles, because the circles must

pass through the gap representing $L_{i+1} \preceq R_{i+1}$, shaded in Fig. 17(a). And because we already established that L cannot enter the R halfplane H_R , we know these arcs are $< \pi$: for an arc of $\geq \pi$ could turn ρ_{\max} to aim into H_R . So $L_i \preceq R_i$. Repeating this argument back to $i = 0$ yields $L \preceq R$, establishing the theorem. \square




Figure 15: (a) Angle conditions are tight. (b) Turning ρ_{\max} .

8.4 From Paths to Trees

We have just settled in Theorem 2 the “safe” separation of a radially monotone path Q . Now we extend the result to trees, for our cut paths are (in general) leaf-to-root paths in some tree $\mathcal{T} \subseteq \mathcal{F}$ of the forest.¹⁰

Again let $Q = (u_0, \dots, u_n)$ be an edge cut-path on \mathcal{C} , with L and R the planar chains derived from Q , just as in Theorem 2. Before opening by curvatures, the vertices in the plane are $Q = (v_0, \dots, v_n)$. Assume there is path $Q' = (v'_0, \dots, v'_k = v_i)$ in the tree containing Q , which is incident to and joins the path at v_i from the left side. See Fig. 18(a).

We fix $Q = R$ and open Q to L as in Theorem 2, rigidly moving the unopened Q' attached to L at the same angle $\angle v_{i-1}v_i v'_{k-1}$ at the join; see Fig. 18(b). Now we apply the same procedure to Q' , but now rigidly moving the tail of L , $L_i = (\ell_i, \dots, \ell_n)$. The logic is that we have already opened that portion of the path, so the curvatures $\omega_i, \dots, \omega_n$ have already been expended. See Fig. 18(c).

We continue this left-expansion process for all the branches of the tree \mathcal{T} , stacking the openings one upon another. Rather than assume a curvature bound of $\omega(Q) < \pi$, we assume that bound summed over the whole tree \mathcal{T} : $\omega(\mathcal{T}) < \pi$.

¹⁰This extension was not described explicitly in [O'R16].




Figure 16: (a) $\omega_i = (17^\circ, 6^\circ, 7^\circ, 0^\circ, 5^\circ, 5^\circ, 7^\circ)$, $i = 0, \dots, 6$. (b) Steps in the induction proof.

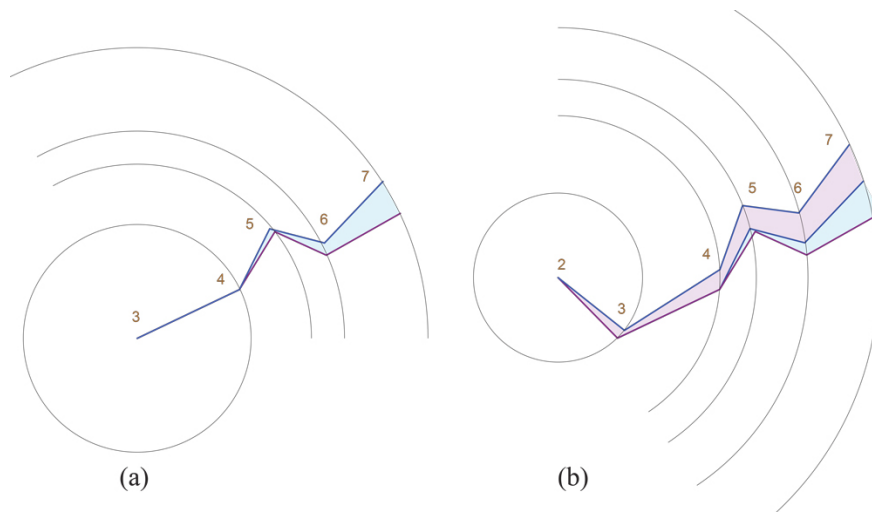


Figure 17: (a) $i + 1 = 3$, $L_3 \preceq R_3$. Note: $\omega_3 = 0$ so $l_3l_4 = r_3r_4$. (b) $i = 2$, $L_2 \preceq R_2$.




Figure 18: (a) Q' joins Q at $v'_3 = v_4$. (b) After opening Q to L and R . (c) After opening Q' .

For it is the total curvature in all descendants of one vertex u_i that rotates the next edge $u_i u_{i+1}$. And similarly, we assume $\Lambda(\mathcal{T}) < \pi$, where ρ_{\max} and ρ_{\min} range over all edges in \mathcal{T} . These reinterpretations of $\omega(\mathcal{T})$ and $\Lambda(\mathcal{T})$ retain the argument in Theorem 2 that the total turn of segments is less than π , and so avoids “wrap-around” crossing of L' with R right-to-left, for any such L' .

We can order the leaves of a tree, and their paths to the root, as they occur in an in-order depth-first search (DFS), so that the entire tree can be processed in this manner (this ordering will be used again in Section 10.1 below).

Corollary 1 *The $L \preceq R$ conclusion of Theorem 2 holds for all the paths in a tree \mathcal{T} : $L' \preceq R$, for any such L' .*

9 Extending \mathcal{C} to \mathcal{C}^∞

6. Extending the cap \mathcal{C} to an unbounded polyhedron \mathcal{C}^∞ ensures that the non-crossing of each L and R extends arbitrarily far in the planar development.

In order to establish nonoverlap of the unfolding, it will help to extend the convex cap \mathcal{C} to an unbounded polyhedron \mathcal{C}^∞ as follows. Define \mathcal{C}^∞ as the intersection of all the halfspaces determined by the faces of \mathcal{C} . Because we have assumed $\Phi < 90^\circ$, \mathcal{C}^∞ is unbounded. It will be convenient to define a “clipped,” bounded version of \mathcal{C}^∞ : let \mathcal{C}^Z be \mathcal{C}^∞ intersected with the halfspace $z \geq Z$. So $\lim_{Z \rightarrow -\infty} \mathcal{C}^Z = \mathcal{C}^\infty$.

We can imagine constructing \mathcal{C}^Z as follows. Let B be the set of *boundary faces* of \mathcal{C} , those that share an edge with $\partial\mathcal{C}$. See Fig. 19(b). Extend these faces downward. They intersect one another, and eventually the “surviving” faces extend to infinity. We will view \mathcal{C}^Z as $\mathcal{C} \cup \mathcal{E}^Z$, where \mathcal{E}^Z is the extension “skirt” of faces. See Fig. 20.

Note that Φ for \mathcal{C}^∞ is the same Φ for the original \mathcal{C} . \mathcal{C}^∞ will allow us to ignore the ends of our cuts, as they can be extended arbitrarily far.

9.1 Acute Triangulation

We will apply the spanning forest construction to \mathcal{C}^Z , and for that purpose, we need to acutely triangulate \mathcal{C}^Z . Fortunately, this is possible.

Very recently, Bishop proved that every PLSG (planar straight-line graph) of n vertices has a conforming acute triangulation, using $O(n^{2.5})$ triangles [Bis16].¹¹ “Conforming” means that each pair of triangles are either disjoint, share a single vertex, or share a whole edge. He cites Burago and Zalgaller as first establishing that acute triangulations exist [BZ80], without a bound in terms of n . For our purposes, the exact bound is not relevant, as long as there is a bound.

¹¹His main Theorem 1.1 is stated for nonobtuse triangulations, but he says later that “the theorem also holds with an acute triangulation, at the cost of a larger constant in the $O(n^{2.5})$.”

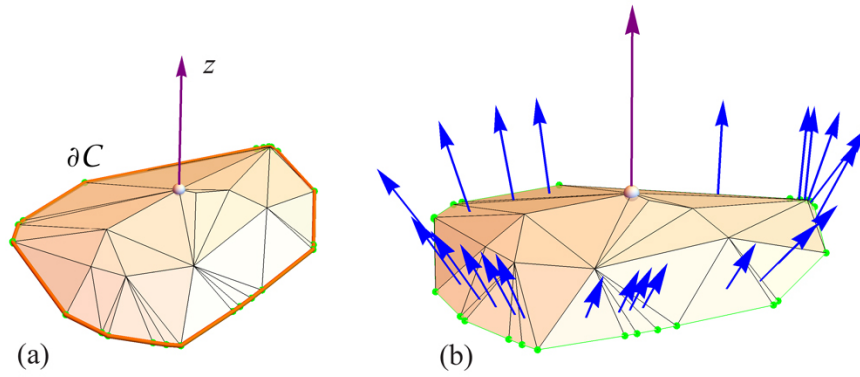


Figure 19: (a) A convex cap C . (b) Normals to the boundary faces B . Note: C is not yet acutely triangulated in this illustration.




Figure 20: Extension of the boundary faces in Fig. 19(b) to form C^Z (not yet acutely triangulated).

Recalling that $\mathcal{C}^Z = \mathcal{C} \cup \mathcal{E}^Z$, we need to acutely triangulate \mathcal{E}^Z . We apply Bishop’s algorithm, introducing (possibly many) new vertices of curvature zero on the extension skirt \mathcal{E}^Z . Then all of \mathcal{C}^Z is acutely triangulated. Note that Φ serves as a bound for both \mathcal{C} and \mathcal{C}^Z .

The consequence is that each cut path \mathcal{Q} can be viewed as extending arbitrarily far from its source on \mathcal{C} before reaching its root on the boundary of \mathcal{C}^Z . This permits us to “ignore” end effects as the cuts are developed in the next section.

10 Angle-Monotone Strips Partition

7. The development of \mathcal{C} can be partitioned into θ -monotone “strips,” whose side-to-side development layout guarantees non-overlap in the plane.

The final step of the proof is to partition the planar C (and so the cap \mathcal{C} by lifting) into strips that can be developed side-by-side to avoid overlap.¹² We return to the spanning forest F of C (graph G), as discussed in Section 4.2. Define an *angle-monotone strip* (or more specifically, a θ -monotone strip) S as a region of C bound by two angle-monotone paths L_S and R_S which emanate from the quadrant origin vertex $q \in L_S \cap R_S$, and whose interior is vertex-free. The strips we use connect from q to each leaf $\ell \in F$, and then follow to the tree’s root on ∂C . For ease of illustration, we will use $\theta = 90^\circ$, but we will see no substantive modifications are needed for $\theta < 90^\circ$. Although there are many ways to obtain such a partition, we describe one in particular, whose validity is easy to see. We describe the procedure in the Q_0 quadrant, with straightforward generalization to the other quadrants.

10.1 Waterfall Algorithm

Let T_0 be the set of leaves of F in Q_0 , with $|T_0| = n$. We describe an algorithm to connect each $\ell \in T_0$ to q via noncrossing θ -monotone paths. Unlike F , which is composed of edges of G , the paths we describe do not follow edges of G . Consult Fig. 21 throughout.

Center a circle of radius r on the origin q , with r smaller than the closest distance from a leaf to the quadrant axes. We may assume (Section 4) that no vertex aside from q lies on a quadrant axis, so $r > 0$. Mark off n “target points” c_i on the circle as in Fig. 21. Process the leaves in T_0 in the following order. Trees in F are processed in counterclockwise order of their root along ∂C . Within each tree, the leaves are ordered as they occur in an in-order depth-first search (DFS); again consult the figure.

Let y_i be the height of c_i , $y_1 > 0$ and $y_n < r$, and let ℓ_i be the i -th leaf. Connect ℓ_1 to c_1 by dropping vertically from ℓ_1 to height $y_1 > 0$, and then

¹²We should imagine \mathcal{C} replaced by \mathcal{C}^Z , which will make the strips arbitrarily long.




Figure 21: Waterfall algorithm in Q_0 .

horizontally to c_1 . So the connection has a “L-shape.” Then connect radially from c_1 to q . Define the path p_1 to be this 3-segment connection joined with the path in F from ℓ_1 to the root on ∂C .

For the i -th step, drop vertically from ℓ_i to path p_{i-1} , following just ε -above p_{i-1} until it reaches height y_i . Then connect horizontally to c_i and radially to q . We select ε to ensure noncrossing of the “stacked” paths. It suffices to use $1/(n+1)$ times the minimum of (a) the smallest vertical distance between a leaf and a point of F , and (b) the smallest horizontal distance between two leaves. We also use the same ε to separate c_{i-1} from c_i vertically around the circle.

To handle $\theta < 90^\circ$, all vertical drops instead drop at angle θ inclined with respect to the x -axis. We take it as clear that each path p_i is θ -monotone. They are noncrossing because p_i rides above p_{i-1} , and ε is small enough so that $n\varepsilon$ cannot bump into a later path p_j for $j > i$. Thus we have partitioned C into θ -monotone strips sharing vertex q . See Fig. 22 for a complete example.

Returning to Fig. 21, some strips will reach a tree junction before ∂C , such as S_3 reaching the junction between ℓ_3 and ℓ_4 . In that case, the strip continues with the path from that junction to ∂C , i.e., its “tail” is a zero-width path. Other strips, such as S_4 in Fig. 22, retain non-zero width to the boundary.

10.2 $L_\perp \preceq R_\perp$

Define S'_i as the strip counterclockwise of leaf ℓ_i in C' . and $L'_{S'_i}$ and $R'_{S'_i}$ as its right and left boundaries (which might coincide from some vertex onward). We reintroduce the primes to distinguish between objects in the planar C' , their




Figure 22: Waterfall strips partition. The S_4 strip highlighted.

lifts on the 3D cap \mathcal{C} , and the development from \mathcal{C} back to the plane. To ease notation, fix i , and let $S = S_i$, $L' = L'_{S_i}$, $R' = R'_{S_i}$. Let \mathcal{S} be the lift of S' to the cap \mathcal{C} . Let L_\perp and R_\perp be the developments of the boundaries of \mathcal{S} back into the plane: the left development of R and the right development of L , so that both developments are determined by the surface angles within \mathcal{S} . Our goal is to prove $L_\perp \preceq R_\perp$, with the left-of relation \preceq as defined in Section 8.2. We start with proving that $L' \preceq R'$.

By construction, both L' and R' are θ -monotone, and therefore radially monotone. So any circle D centered on q intersects each in at most one point, say $a' = D \cap L'$ and $b' = D \cap R'$. The arc from a' to b' lies inside the strip S ; or perhaps $a' = b'$ in the tail of S . To prove $L' \preceq R'$, we only need to show this $a'b'$ arc is at most π . This is obvious when S lies in one quadrant. Although it can be proved if S straddles two quadrants, it is easier to just use the quadrant boundaries to split a boundary-straddling strip into two halves, so that always $a'b'$ lies in one quadrant, and so is at most π .

Now we turn to the lift of S' to \mathcal{S} and the development of the boundaries L_\perp and R_\perp . Lemma 6 guarantees that L_\perp and R_\perp are still θ -monotone, for some $\theta < 90^\circ$. Thus the argument is the same as above, establishing that $L_\perp \preceq R_\perp$ for each strip \mathcal{S} .

10.3 Side-by-Side Layout

We now extend the \preceq relation to adjacent strips. We drop the \perp subscripts, and just let S_{i-1} and S_i be the developments in the plane of two adjacent strips. For $S_i \preceq S_{i-1}$ to hold, we require that every circle centered on their common source q intersects the right boundary of S_{i-1} at point a , the left boundary of S_i at point b counterclockwise of a , and intersects no other strip along the ab arc.

We just established that the boundaries of both S_{i-1} and S_i are θ -monotone, and so a circle D centered on q will indeed meet the extreme boundaries in one point each, a and b . Here we rely on the extension \mathcal{C}^∞ so that each strip extends arbitrarily long, effectively to ∞ . The two strips share a boundary from q to a leaf vertex v_i . Beyond that they deviate if $\omega_i > 0$; see Fig. 23(a). We established in Theorem 2 that the two sides of the “gap” at v_i , L_i , and R_i , satisfy $L_i \preceq R_i$ (with respect to v_i). This guarantees that there is no surface developed in the gap, which again we can imagine extending to ∞ . So the arc ab crosses S_{i-1} , then the gap, then S_i . So indeed $S_i \preceq S_{i-1}$. Here there is no worry about the length of the arc ab being so long that L_{S_i} could wraparound and cross $R_{S_{i-1}}$ from right-to-left, because each strip fits in a quadrant.

Now we lay out the strips according to their \preceq -order. We choose S_1 to be the strip left-adjacent to the gap of the cut of the forest to q (cf. Fig. 6), and proceed counterclockwise from there. The layout of S_i will not overlap with S_{i-1} because $S_i \preceq S_{i-1}$. Now we argue that nor can S_i wraparound and overlap S_1 .¹³

¹³Note here we cannot argue that wraparound intersection doesn't occur because of the

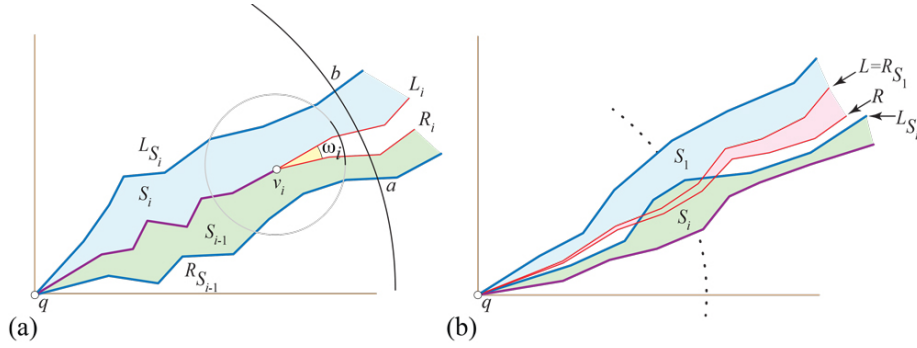


Figure 23: (a) Two adjacent strips. (b) S_i crosses into S_1 .

For suppose S_i crosses into S_1 from right-to-left, that is, L_{S_i} crosses R_{S_1} ; see Fig. 23(b). Then because $S_{i+1} \preceq S_i$, S_{i+1} must also cross into S_1 . Continuing, we conclude that S_n crosses into S_1 . Now S_1 and S_n are separated by the gap opening of the path in \mathcal{F} reaching q . The two sides of this gap are $R = L_{S_n}$ and $L = R_{S_1}$. But we know that $L \preceq R$ by Theorem 2. So we have reached a contradiction, and S_i cannot cross into S_1 .

We continue laying out strips until we layout S_n , which is right-adjacent to the q -gap. The strips together with that gap fill out the 360° neighborhood of q . Consequently, we have developed all of \mathcal{C} without overlap, and Theorem 1 is proved.

11 Discussion

Many assumptions made to reach Theorem 1 seem possibly unnecessary, only assumed to simplify the argument.

1. That the rim of \mathcal{C} lies in a plane is unlikely to be necessary. Perhaps some star-shaped assumption would suffice, but this would not greatly extend the reach of theorem. In Fig. 24, $\partial\mathcal{C}$ does not lie in a plane.
2. Although the proof requires “sufficiently small Φ ” ($\Phi < 0.3\sqrt{\Delta\theta}$ according to Eq. 9), limited empirical exploration suggests Φ need not be small, e.g., $\Phi = 53^\circ$ in Fig. 24. The proof assumes a worst-case scenario, with all the curvature concentrated on a single cut path. When the curvature is distributed over \mathcal{C} , likely larger Φ still will lead to non-overlap.
3. The assumption that \mathcal{C} is acutely triangulated seems overly cautious, especially since the results of [LO17] extend to θ -monotone paths for widths larger than 90° (indeed for any $\theta \geq 60^\circ$). However, note that a θ -monotone path for $\theta > 90^\circ$ need not be radially monotone, which is required in Theorem 2.

limitation on turning $< \pi$, which only excludes S_i wrapping around to S_{i-1} .

4. This suggests the question of whether the angle-monotone spanning forest could be replaced with a radially monotone spanning forest. There is empirical evidence that radially monotone spanning forests lead to edge-unfoldings of spherical polyhedra [O'R16]. There exist planar triangulations with no radially monotone spanning forest (Appendix of [O'R16]), but it is not clear they can be realized in \mathbb{R}^3 to force overlap.
5. It is natural to wonder if Theorem 1 leads to a “fewest nets” result.¹⁴ This asks for unfolding a convex polyhedron of n vertices into the fewest number of non-overlapping pieces (nets) [DO07, Open Prob. 22.12]. There is as yet no proof that this is possible for a constant number of pieces, independent of n . I have only been able to obtain the following weak result. Define a vertex v to be *oblong* if (a) the largest disk inscribed in v 's spherical polygon $s(v)$ (on the Gaussian sphere) has radius $< \Phi$, and (b) the aspect ratio of $s(v)$ is greater than 2:1. An example was shown in Fig. 8. Then, if an acutely triangulated convex polyhedron P has no (or a constant number of) oblong vertices, then there is a partition of its faces into a constant number of convex caps that edge-unfold to nets, with the constant depending on Φ (and not n). I leave this as a claim without providing a detailed proof.

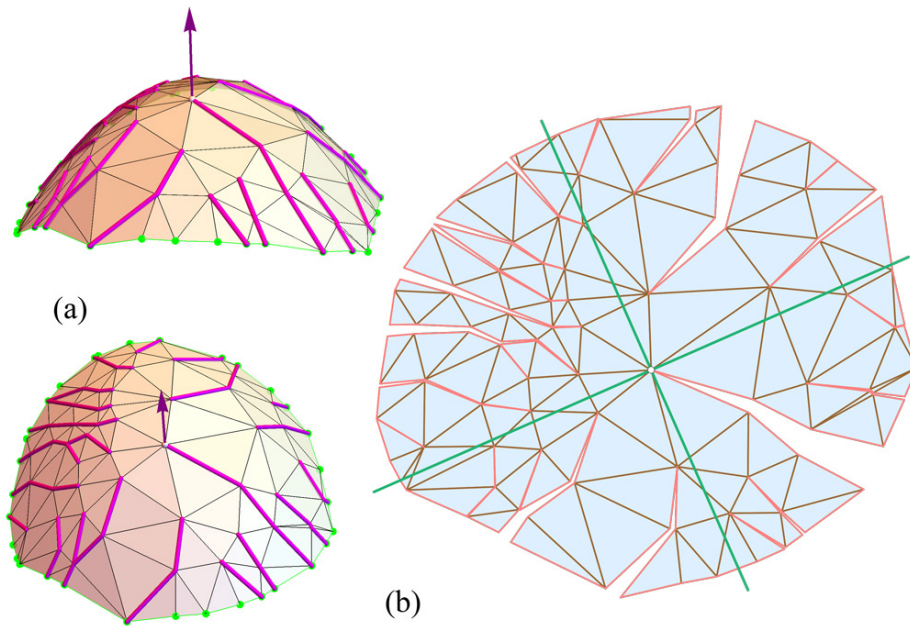


Figure 24: (a) Two views of a convex cap of 83 vertices with spanning forest \mathcal{F} marked. Here $\Phi \approx 53^\circ$. (b) Edge unfolding by cutting \mathcal{F} .

¹⁴Stefan Langerman, personal communication, August 2017.

References

- [BBC⁺16] Nicolas Bonichon, Prosenjit Bose, Paz Carmi, Irina Kostitsyna, Anna Lubiw, and Sander Verdonschot. Gabriel triangulations and angle-monotone graphs: Local routing and recognition. In *Graph Drawing*, 2016.
- [Bis16] Christopher J Bishop. Nonobtuse triangulations of pslgs. *Discrete & Comput. Geom.*, 56(1):43–92, 2016.
- [BLS07] Therese Biedl, Anna Lubiw, and Michael Spriggs. Cauchy’s theorem and edge lengths of convex polyhedra. *Algorithms and Data Structures*, pages 398–409, 2007.
- [BZ80] Y.D. Burago and V.A. Zalgaller. Polyhedral embedding of a net. 1980.
- [DFG15] Hooman Reisi Dehkordi, Fabrizio Frati, and Joachim Gudmundsson. Increasing-chord graphs on point sets. *J. Graph Algorithms Applications*, 19(2):761–778, 2015.
- [DO07] Erik D. Demaine and Joseph O’Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, July 2007. <http://www.gfalop.org>.
- [Gho14] Mohammad Ghomi. Affine unfoldings of convex polyhedra. *Geometry & Topology*, 18(5):3055–3090, 2014.
- [IKL99] Christian Icking, Rolf Klein, and Elmar Langetepe. Self-approaching curves. *Math. Proc. Camb. Phil. Soc.*, 125:441–453, 1999.
- [Lee06] John M. Lee. *Riemannian Manifolds: An Introduction to Curvature*, volume 176. Springer Science & Business Media, 2006.
- [LO17] Anna Lubiw and Joseph O’Rourke. Angle-monotone paths in non-obtuse triangulations. In *Proc. 29th Canad. Conf. Comput. Geom.*, August 2017. arXiv:1707.00219 [cs.CG]: <https://arxiv.org/abs/1707.00219>.
- [O’R03] Joseph O’Rourke. On the development of the intersection of a plane with a polytope. *Comput. Geom. Theory Appl.*, 24(1):3–10, 2003.
- [O’R13] Joseph O’Rourke. Dürer’s problem. In Marjorie Senechal, editor, *Shaping Space: Exploring Polyhedra in Nature, Art, and the Geometrical Imagination*, pages 77–86. Springer, 2013.
- [O’R16] Joseph O’Rourke. Unfolding convex polyhedra via radially monotone cut trees. arXiv:1607.07421 [cs.CG]: <https://arxiv.org/abs/1607.07421>, 2016.