

Development of Curves on Polyhedra via Conical Existence*

Joseph O'Rourke[†] Costin Vilcu[‡]

October 7, 2011

Abstract

We establish that certain classes of simple, closed, polygonal curves on the surface of a convex polyhedron develop in the plane without overlap. Our primary proof technique shows that such curves “live on a cone,” and then develops the curves by cutting the cone along a “generator” and flattening the cone in the plane. The conical existence results support a type of source unfolding of the surface of a polyhedron, described elsewhere.

1 Introduction

Nonoverlapping development of curves plays a role in unfolding polyhedra without overlap [2]. Any result on simple (non-self-intersecting) development of curves may help establishing nonoverlapping surface unfoldings. One of the earliest results in this regard is [8], which proved that the left development of a directed, simple, closed convex curve does not self-intersect. The proof used Cauchy's Arm Lemma. Here we extend this result to a wider class of curves without invoking Cauchy's lemma. Our results support a “source unfolding” based on these curves, described in [6].

Development. Let C be a simple, closed, polygonal curve on the surface of a convex polyhedron \mathcal{P} . For any point $p \in C$, let $L(p)$ be the total surface angle incident to p at the left side of C , and $R(p)$ the angle to the right side. The *left development* of C with respect to $x \in C$ is an isometric drawing \overline{C}_x of C in the plane, starting from x , such that the angle to the left of \overline{C}_x at every point in the plane is $L(p)$. The *right development* is defined analogously. The left and right developments of a curve are different if C passes through one or more vertices of \mathcal{P} . And in general the development depends upon the *cut point* x .

*A preliminary version of these results was reported in [9].

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

[‡]Institute of Mathematics “Simion Stoilow” of the Romanian Academy, P.O. Box 1-764, RO-014700 Bucharest, Romania. Costin.Vilcu@imar.ro.

Curve Classes. To describe our results, we introduce several different classes of curves on convex polyhedra, which exhibit different behavior with respect to living on a cone. Define a curve C to be *convex* (to the left) if the angle to the left is at most π at every point p : $L(p) \leq \pi$; and say that C is a *convex loop* if this condition holds for all but one exceptional *loop point* x , at which $L(x) > \pi$ is allowed. Analogously, define C to be a *reflex curve* if the angle to one side (we consistently use the right side) is at least π at every point p : $R(p) \geq \pi$; and say that C is a *reflex loop* if this condition holds for all but an exceptional loop point x , at which $R(x) < \pi$.

The loop versions of these curves arise naturally in some contexts. For example, extending a convex path on \mathcal{P} until it self-intersects leads to a convex loop.

Summary of Results.

1. Every convex curve C left-develops to $\overline{C_x}$ without intersection, for every cut point x . Here we give a new proof for the result in [8].
2. There are convex loops C such that, for some x , the left-development $\overline{C_x}$ self-intersects. However, for every convex loop, there exists a y for which $\overline{C_y}$ left-develops without overlap.
3. Every reflex curve C right-develops to $\overline{C_x}$ without intersection, for every cut point x .
4. Every reflex loop C whose other side is convex, right-develops to $\overline{C_x}$ without intersection, for every cut point x .

These results may be combined to reach conclusions about the left- and right-developments of the same curve: Every convex curve C that passes through at most one vertex, both left-develops, and right-develops without overlap, for every cut point x .

Living on a Cone. Our primary proof technique relies on the notion of a curve C “living on a cone,” which is based on neighborhoods of C . An open region N_L is a *vertex-free left neighborhood* of C to its left if it includes C as its right boundary, and it contains no vertices of \mathcal{P} . In general C will have many vertex-free left neighborhoods, and all will be equivalent for our purposes. We say that C *lives on a cone* to its left if there exists a cone Λ and a neighborhood N_L so that N_L may be embedded isometrically onto Λ , and encloses the cone apex a .

A *cone* is an unbounded developable surface with curvature zero everywhere except at one point, its *apex*, which has total incident surface angle, the *cone angle*, of at most 2π . Throughout, we will consider a cylinder as a cone whose apex is at infinity with cone angle 0, and a plane as a cone with apex angle 2π . We only care about the intrinsic properties of the cone’s surface; its shape in \mathbb{R}^3 is not relevant for our purposes. So one could view it as having a circular cross section, although we will often flatten it to the plane.

We should remark that the cone on which a curve C lives has no direct relationship (except in special cases) to the surface that results from extending the faces of \mathcal{P} crossed by C .

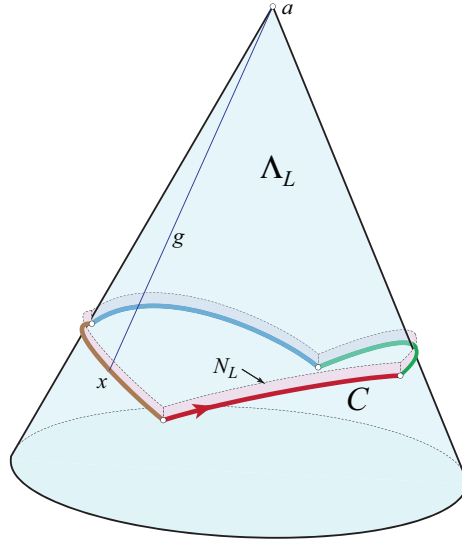


Figure 1: A 4-segment curve C that lives on cone Λ_L to its left. One possible N_L is shown, and a generator $g = ax$ is illustrated.

To say that N_L embeds isometrically into Λ means that we could cut out N_L (including its right boundary C) and paste it onto Λ with no wrinkles or tears: the distance between any two points of N_L on $N_L \cap \mathcal{P}$ is the same as it is on $N_L \cap \Lambda$. See Figure 1. We say that C lives on a cone to its right if N_R embeds isometrically on the cone, where N_R is a vertex-free right neighborhood of C such that the cone apex a is inside (the image of) C . We will call the cones to the left and right of C , Λ_L and Λ_R respectively. We will see that all four combinatorial possibilities occur: C may not live on a cone to either side, it may live on a cone to one side but not to the other, it may live on different cones to its two sides, or live on the same cone to both sides.

Cone Generators and Visibility. A *generator* of a cone Λ is a ray starting from the apex a and lying on Λ . A curve C that lives on Λ is *visible* from the apex if every generator meets C at one point. Although it is possible for a curve to live on a cone but not be visible from its apex, when we can establish visibility from the apex, then cutting C at any point $x \in C$ will develop $\overline{C_x}$ without overlap, because cutting Λ along g develops the cone and C simultaneously.

2 Preliminary Tools and Lemmas

C partitions \mathcal{P} into two *half-surfaces*. We call the left and right half-surfaces P_L and P_R respectively, or P if the distinction is irrelevant. We view each half-surface as closed, with boundary C .

Curvature. The *curvature* $\omega(p)$ at any point $p \in \mathcal{P}$ is the “angle deficit”: 2π minus the sum of the face angles incident to p . The curvature is only nonzero at vertices of \mathcal{P} ; at each vertex it is positive because \mathcal{P} is convex. The curvature at the apex of a cone is similarly 2π minus the cone angle.

Define a *corner* of curve C to be any point p at which either $L(p) \neq \pi$ or $R(p) \neq \pi$. Let c_1, c_2, \dots, c_m be the corners of C , which may or may not also be vertices of \mathcal{P} . C “turns” at each c_i , and is straight at any noncorner point. Let $\alpha_i = L(c_i)$ be the surface angle to the left side at c_i , and $\beta_i = R(c_i)$ the angle to the right side. Also let $\omega_i = \omega(c_i)$ to simplify notation. We have $\alpha_i + \beta_i + \omega_i = 2\pi$ by the definition of curvature. These definitions will be used to further detail the relationships among the curve classes in Section 5.

The Gauss-Bonnet Theorem. We will employ this theorem in two forms. The first is that the total curvature of \mathcal{P} is 4π : the sum of $\omega(v)$ for all vertices v of \mathcal{P} is 4π . It will be useful to partition the curvature into three pieces. Let $\Omega_L(C) = \Omega_L$ be the total curvature strictly interior to P_L , Ω_R the curvature to the right, and Ω_C the sum of the curvatures on C (which is nonzero only at vertices of \mathcal{P}). Then $\Omega_L + \Omega_C + \Omega_R = 4\pi$.

The second form of the Gauss-Bonnet theorem relies on the notion of the “turn” of a curve. Define $\tau_L(c_i) = \tau_i = \pi - \alpha_i$ as the left *turn* of curve C at corner c_i , and let $\tau_L(C) = \tau_L$ be the total (left) turn of C , i.e., the sum of τ_i over all corners of C . Thus a convex curve has nonnegative turn at each corner, and a reflex curve has nonpositive turn at each corner. Then $\tau_L + \Omega_L = 2\pi$, and defining the analogous term to the right of C , $\tau_R + \Omega_R = 2\pi$.

Alexandrov’s Gluing Theorem. In our proofs we use Alexandrov’s theorem [1, Thm. 1, p. 100] that gluing polygons to form a topological sphere in such a way that at most 2π angle is glued at any point, results in a unique convex polyhedron.

Vertex Merging. We now explain a technique used by Alexandrov, e.g., [1, p. 240]. Consider two vertices v_1 and v_2 of curvatures ω_1 and ω_2 on \mathcal{P} , with $\omega_1 + \omega_2 < 2\pi$, and cut \mathcal{P} along a shortest path $\gamma(v_1, v_2)$ joining v_1 to v_2 . Construct a planar triangle $T = \bar{v}'\bar{v}_1\bar{v}_2$ such that its base $\bar{v}_1\bar{v}_2$ has the same length as $\gamma(v_1, v_2)$, and the base angles are equal to $\frac{1}{2}\omega_1$ and respectively $\frac{1}{2}\omega_2$. Glue two copies of T along the corresponding lateral sides, and further glue the two bases of the copies to the two “banks” of the cut of \mathcal{P} along $\gamma(v_1, v_2)$. By Alexandrov’s Gluing Theorem, the result is a convex polyhedral surface \mathcal{P}' . On \mathcal{P}' , the points v_1 and v_2 are no longer vertices because exactly the angle deficit

at each has been sutured in; they have been replaced by a new vertex v' of curvature $\omega' = \omega_1 + \omega_2$ (preserving the total curvature). Figure 2(a) illustrates this. Here $\gamma(v_1, v_2) = v_1v_2$ is the top “roof line” of the house-shaped polyhedron \mathcal{P} . Because $\omega_1 = \omega_2 = \frac{1}{2}\pi$, T has base angles $\frac{1}{4}\pi$ and apex angle $\frac{1}{2}\pi$. Thus the curvature ω' at v' is π . (Other aspects of this figure will be discussed later.)

Note this vertex-merging procedure only works when $\omega_1 + \omega_2 < 2\pi$; otherwise the angle at the apex \bar{v}' of T would be greater than or equal to π .

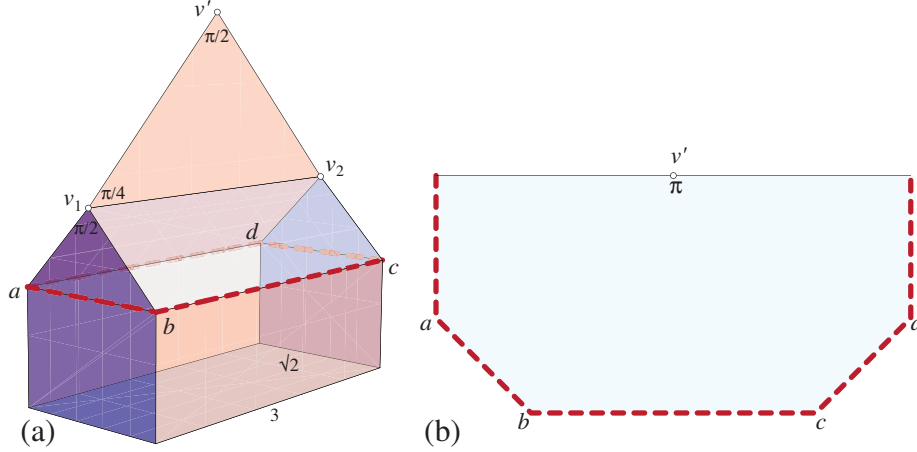


Figure 2: (a) $C = (a, b, c, d)$ is a convex curve with angle $\frac{3}{4}\pi$ to the left at each vertex. The curvature at v_1 and at v_2 is $\frac{1}{2}\pi$. (b) Cutting along the generator from v' through the midpoint of ad and developing C shows that it lives on a cone with apex angle π at v' . (Base of \mathcal{P} is $3 \times \sqrt{2}$.)

Lemma 1 *A curve C that lives on a cone Λ (say, to its left) uniquely determines that cone.*

Proof: Suppose that C lives on two cones Λ and Λ' . We will show that the regions of these two cones bounded by C are isometric. First note that the apex curvature of both Λ and Λ' is Ω_L , the total curvature inside and left of C . This follows from the Gauss-Bonnet theorem: $\tau_L + \Omega_L = 2\pi$, and τ_L is the same whether on Λ or Λ' . Let $x \in C$ be a point of C that has a tangent t to one side, and let x_1 be a point in the plane and t_1 a direction vector from x_1 . Roll Λ in the plane so that x and t coincide with x_1 and t_1 . Continue rolling until x is encountered again; call that point of the plane x_2 . The resulting positions of x_1 and x_2 are the same as would be produced by cutting the cone along a generator ax .

If $x_1 = x_2$, then both Λ and Λ' are planar and so isometric. So assume $x_1 \neq x_2$. If $\Omega_L \geq \pi$, then the cone angle $\alpha \leq \pi$, as in Figure 3(b). The segment x_1x_2 determines two isosceles triangles with apex angle α , only one of which can correspond to the left side of \bar{C} . Analogously, if $\Omega_L < \pi$, then x_1x_2 determines

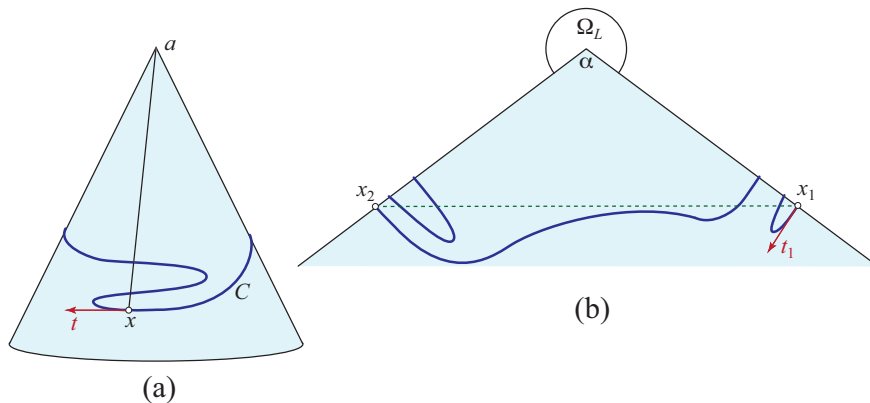


Figure 3: (a) Cone Λ on which C lives. (b) Positions of x_1 and x_2 after cutting open Λ along ax .

a unique isosceles triangle of apex angle Ω_L , the equal sides of which bound, together with \overline{C} , the region of Λ to the left of \overline{C} . Note that \overline{C} doesn't actually depend on the cones Λ and Λ' , but only on the left neighborhood of C in \mathcal{P} , and hence this development is the same for Λ and Λ' . So, up to planar isometries, the planar unfolding of the cone supporting C is unique, and thus the cone itself and the position of C on it are unique up to isometries. \square

Note that this lemma does not assume that C is convex; rather it holds for any closed curve C .

3 Convex Curves

Convexity of Half-Surfaces. In order to apply vertex merging, we use a lemma to guarantee the existence of a pair to merge. We first remark that it is not the case that every half-surface $P \subset \mathcal{P}$ bounded by a convex curve C is *convex* in the sense that, if $x, y \in P$, then a shortest path γ of \mathcal{P} connecting x and y lies in P .

Example 1. Let \mathcal{P} be defined as follows. Start with the top half of a regular octahedron, whose four equilateral triangle faces form a pyramid over a square base $abcd$. Remove the base and flex the pyramid by squeezing a toward c slightly while maintaining the four equilateral triangles, a motion that separates b from d . Define \mathcal{P} to be the convex hull of these four moved points $a'b'c'd'$ and the pyramid apex. Let $C = (a', b', c', d')$ and let P be the half-surface including the four equilateral triangles. Then a' and c' are in P , but the edge $a'c'$ of \mathcal{P} , which is the shortest path connecting those points, is not in P : it crosses the “bottom” of \mathcal{P} .

Although P may not be convex, P is *relatively convex* in the sense that it is isometric to a convex half-surface: there is some $\mathcal{P}^\#$ and a half-surface

$P^\# \subset \mathcal{P}^\#$ such that P is isometric to $P^\#$ and $P^\#$ is convex.

Lemma 2 *Every half-surface $P \subset \mathcal{P}$ bounded by a convex curve C is relatively convex, i.e., P is isometric to a half-surface that contains a shortest path γ between any two of its points x and y . More particularly, if neither x nor y is on C , then the shortest path γ contains no points of C . If exactly one of x or y is on C , then that is the only point of γ on C .*

Proof: We glue two copies of P along $\partial P = C$. Because C is convex, Alexandrov's Gluing Theorem says the resulting surface is isometric to a unique polyhedral surface, call it $\mathcal{P}^\#$. Because $\mathcal{P}^\#$ has intrinsic symmetry with respect to C , a lemma of Alexandrov [1, p. 214] applies to show that the polyhedron $\mathcal{P}^\#$ has a symmetry plane Π containing C .

Now consider the points x and y in the upper half P of $\mathcal{P}^\#$, at or above Π . If γ is a shortest path from x to y , then by the symmetry of $\mathcal{P}^\#$, so is its reflection γ' in Π . Because shortest paths on convex surfaces do not branch, γ must lie in the closed half-space above Π , and so lies on P .

If neither x nor y are on C , they are strictly above Π , and γ must be as well to avoid a shortest-path branch. If, say, $x \in C$ but $y \notin C$, and if γ touched C elsewhere, say at z , then from y to x we have a shortest path γ and another shortest path, composed of the arc of γ from y to z and the arc of γ' from z to x , hence we would have a shortest-path branch at z . If both x and y are on C , then either γ meets C in exactly those two points, or $\gamma \subset C$, for the same reason as above. \square

Lemma 3 *Let C be a convex curve on \mathcal{P} , convex to its left. Then C lives on a cone Λ_L to its left side, whose apex a has curvature Ω_L .*

Proof: By the Gauss-Bonnet theorem, $\tau_L + \Omega_L = 2\pi$. Because $\tau_L \geq 0$ for a convex curve, we must have $\Omega_L \leq 2\pi$. Let V be the set of vertices of the half-surface P_L not on C .

Suppose first that $\Omega_L < 2\pi$. If $|V| = 1$, then P_L is a pyramid, which is already a cone. So suppose $|V| \geq 2$, and let v_1 and v_2 be any two vertices in V . Lemma 2 guarantees that a shortest path γ between them is in $P_L^\#$ and disjoint from C . This shortest path corresponds to a geodesic γ in P_L . Perform vertex merging along γ , resulting in a new vertex v' whose curvature is the sum of that of v_1 and v_2 . Note that merging is always possible, because $\omega_1 + \omega_2 \leq \Omega_L < 2\pi$. Also note that v' is not on C , by Lemma 2. Let N_L be some small left neighborhood of C in P_L . Then N_L is unaffected by the vertex merging: neither v_1 nor v_2 is in N_L because it is vertex free, and N_L may be chosen narrow enough (by Lemma 2) so that no portion of γ is in N_L . Replace V by $(V \setminus \{v_1, v_2\}) \cup \{v'\}$.

Continue vertex merging in a like manner between vertices of V until $|V| = 1$, at which point we have C and N_L living on a cone, as claimed.

If $\Omega_L = 2\pi$, then the last step of vertex merging will not succeed. However, we can see that a slight altering of the two glued triangles so that $\Omega_L < 2\pi$ will result in the cone apex approaching infinity, as follows. Cut along a geodesic

between the two vertices, say v_i and v_{i+1} , and insert double triangles of base angles $\frac{1}{2}\omega_i$ and respectively $\frac{1}{2}\omega_{i+1} - \varepsilon_n$, with $\varepsilon_n > 0$ and $\lim_n \varepsilon_n = 0$. And so in this case, C and N_L live on a cylinder, which we earlier defined as a degenerate cone. \square

Lemma 4 *A convex curve C on \mathcal{P} is visible from the apex a of the unique cone Λ on which it lives to its convex side.*

Proof: Let z be a closest point of C to a . Then az must be orthogonal to C at z , by [4, Cor. 1] (repeated in [5, Lem. 1]). Now cut Λ along az , which clearly cannot intersect C except at z . Continue cutting around C , and call the result P . Insert an isosceles “curvature triangle” at the cut az with apex angle $\omega(a)$. This flattens P to a planar domain with a convex boundary, convex because the angles at the two images z_1 and z_2 of z are each less than π ; see Figure 4. Visibility of all of C from a follows. \square

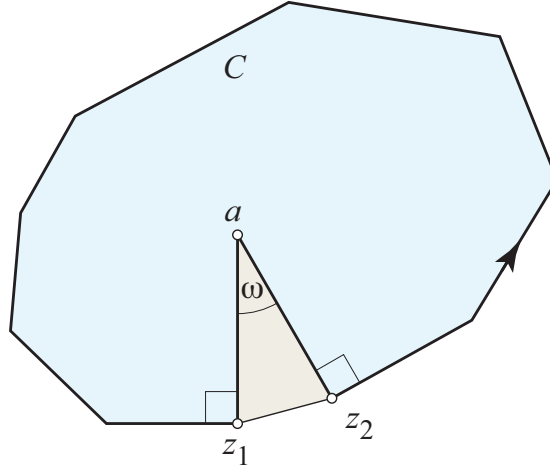


Figure 4: All of C is visible from a . Here z_1 and z_2 are images of z when the cone is cut along az .

A different proof of this lemma is given in [9, Lem. 4].

Example 2. In Figure 2, the two vertices inside C , of curvature $\frac{1}{2}\pi$ each, are merged to one of curvature π , which is then the apex of a cone on which C lives.

Example 3. Figure 5(a) shows an example with three vertices inside C . \mathcal{P} is a doubly covered flat pentagon, and $C = (v_4, v_5, v_4)$ is the closed curve consisting of a repetition of the segment v_4v_5 . C has π surface angle at every point to its left, and so is convex. The curvatures at the other vertices are $\omega_1 = \pi$ and $\omega_2 = \omega_3 = \frac{1}{2}\pi$. Thus $\Omega_L = 2\pi$, and the proof of Lemma 3 shows that C lives on a cylinder. Following the proof, merging v_1 and v_2 removes those vertices and creates a new vertex v_{12} of curvature $\frac{3}{2}\pi$; see (b) of the figure. Finally merging

v_{12} with v_3 creates a “vertex at infinity” v_{123} of curvature 2π . Thus C lives on a cylinder as claimed. If we first merged v_2 and v_3 to v_{23} , and then v_{23} to v_1 , the result is exactly the same, although not obviously so.

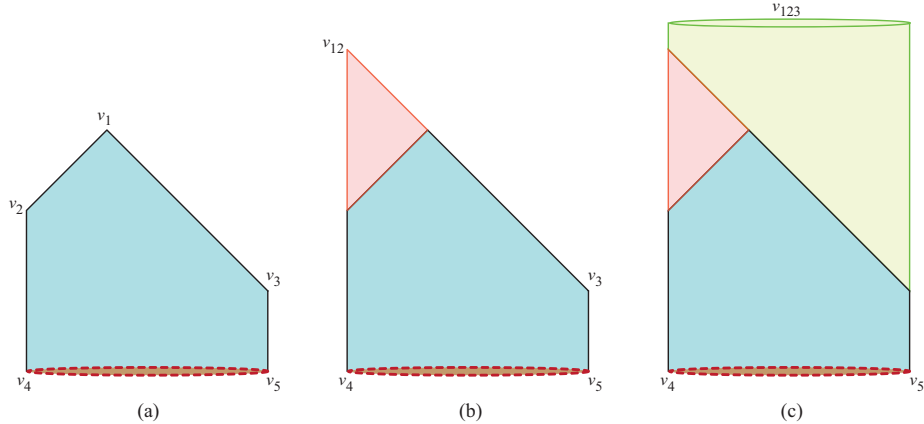


Figure 5: (a) A doubly covered flat pentagon. (b) After merging v_1 and v_2 . (c) After merging v_{12} and v_3 .

We summarize the preceding lemmas in a theorem:

Theorem 1 *Let C be a left-convex curve on \mathcal{P} . Then C lives on a unique cone Λ_L to its left side, whose apex a has curvature Ω_L , and so has cone apex angle $2\pi - \Omega_L$. C is visible from the apex a of Λ .*

4 Convex Loops

Convex Loops and Cones. We first show that the technique that proved successful for convex curves cannot apply to all convex loops: not every convex loop lives on a cone. Consider the polyhedron \mathcal{P} shown in Figure 6(a), which is a variation on the example from Figure 2(a). Here $C = (a, b, b', x, c', c, d)$ is a convex loop, with loop point x . The cone on which it should live is analogous to Figure 2(b): vertex merging of v_1 and v_2 again produces the cone apex v' whose curvature is π . But C does not “fit” on this cone, as Figure 6(b) shows; the apex $a = v'$ is not inside C .

Overlapping development of convex loop. In light of the preceding negative result, it is perhaps not surprising that there are convex loops C and points $x \in C$ such that $\overline{C_x}$ left-develops with overlap. Indeed Figure 7 shows an example where x is the loop point.

Despite the negative result illustrated above, we can show that there always exists some cut point y that develops a convex loop without overlap.

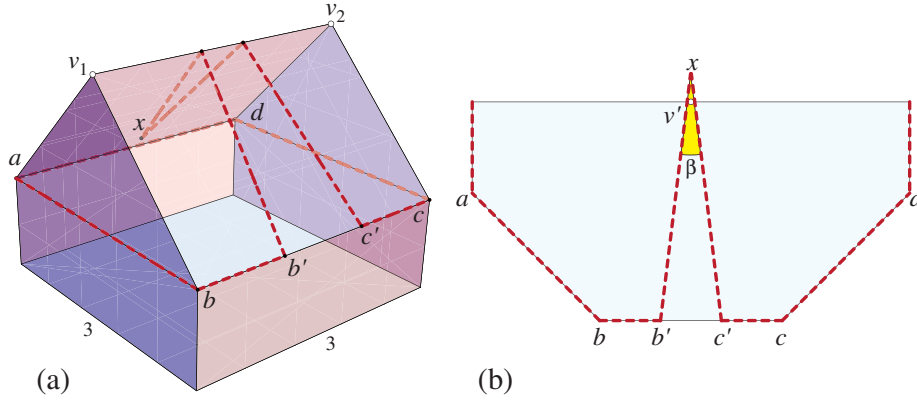


Figure 6: (a) A convex loop C that does not live on a cone. (b) A flattening of the cone on which it should live. (Base of \mathcal{P} is 3×3 .)

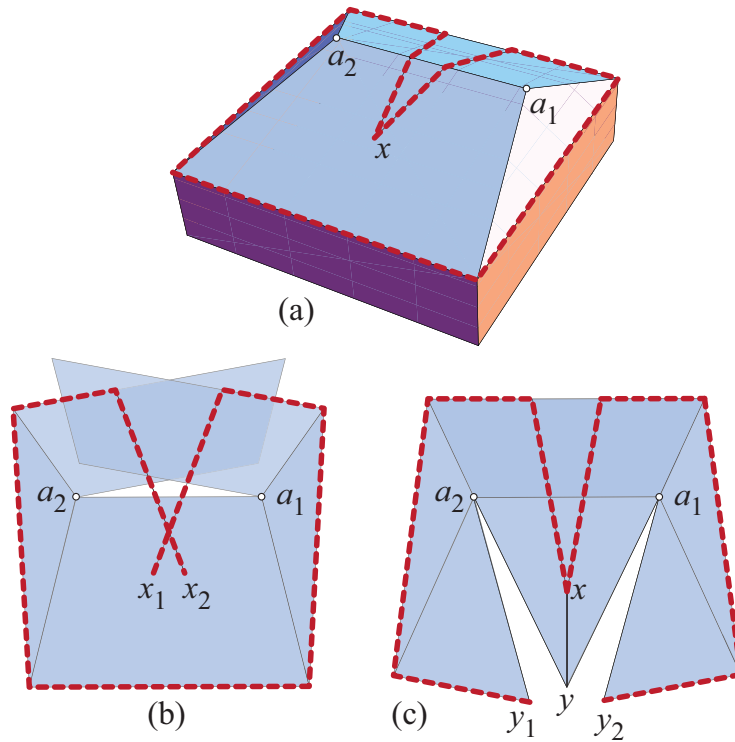


Figure 7: (a) \mathcal{P} with convex loop C . (b) $\overline{C_x}$ overlaps when cut at loop point x . (c) $\overline{C_y}$ does not overlap when cut at y .

Lemma 5 *Every convex loop C contains a point y different from its loop point x , such that \overline{C}_y left-develops without overlap.*

Proof: Let τ_1 and τ_2 be the tangent directions at x pointing into C .

Case 1. Assume first there exists a shortest path $\gamma = xy$ from x to some $y \in C$ whose tangent direction at x lies between τ_1 and τ_2 ; see Figure 8(a). Then γ splits $P = P_L$ into two convex regions P_i sharing a common boundary γ . We perform vertex merging within each P_i , just as in the proof of Lemma 3, producing two cones Λ_i (with apices a_i), sharing a common boundary γ .

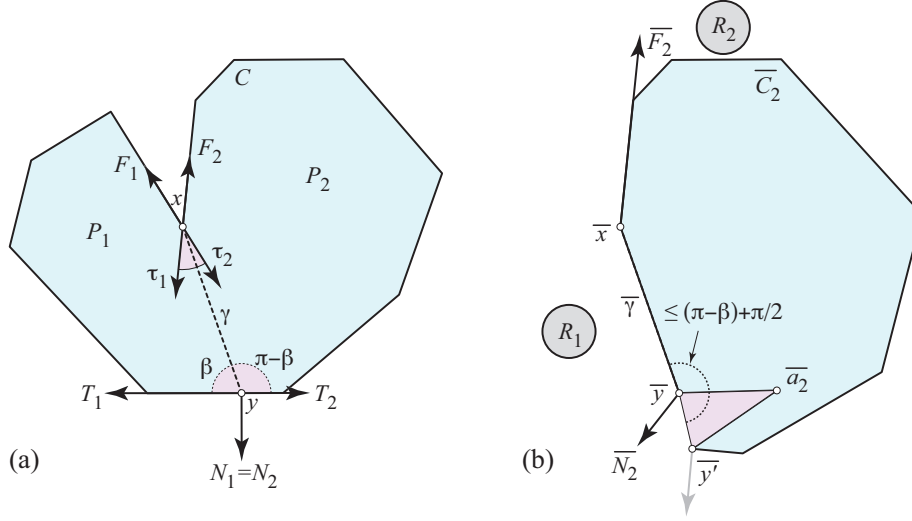


Figure 8: Case 1. (a) $P_L = P_1 \cup P_2$ on \mathcal{P} : $\gamma = xy$ is a shortest path. (b) Planar development of cone Λ_2 .

Now cut each cone along the generators $a_i y$, and unfold both cones, joining them along γ . We now describe the geometry of this planar layout \overline{P} , and show that \overline{C}_y is thereby developed without overlap.

Let T_1 and T_2 be rays tangent to C at y ; if y is not a corner of C , then T_1 and T_2 are collinearly opposing; we will assume this, as it is only easier if y is a corner. This situation is illustrated on \mathcal{P} in Figure 8(a). Now we describe the planar layout, using over-bars to represent elements embedded in the plane; see Figure 8(b).

The two cone unfoldings $\overline{\Lambda}_i$ are joined along $\overline{\gamma}$. Let \overline{N}_1 and \overline{N}_2 be rays from \overline{y} , making with $\overline{\gamma}$ angles $\beta + \pi/2$ and $(\pi - \beta) + \pi/2$, respectively. Informally these correspond to “normals” pointing to the reflex side of C at \overline{y} . We stress that \overline{N}_1 and \overline{N}_2 are defined uniquely by their angles with $\overline{\gamma}$, and not relative to \overline{T}_1 and \overline{T}_2 . We have $\overline{N}_1 = \overline{N}_2$ when y is not a corner of C . Let \overline{F}_2 be the ray tangent to \overline{P} at \overline{x} , directed opposite to $\overline{\tau}_1$, and define \overline{F}_1 similarly. Define R_i to be the regions of the plane bounded by $\overline{N}_i \cup \overline{\gamma} \cup \overline{F}_i$. The angle at \overline{y} in R_1 is $\beta + \pi/2$, and that in R_2 is $(\pi - \beta) + \pi/2$. Let \overline{y}' be the second image of y in Λ_2 that results

from cutting a_2y , so that $\triangle \overline{a_2\overline{y}y'}$ is an isosceles “curvature triangle” apexed at $\overline{a_2}$ with angle $\omega(a_2)$ there. First note that $\angle \overline{a_2y\overline{x}} < \pi - \beta$ because that is the angle at y formed between γ and T_2 on \mathcal{P} . The angle at the base of the isosceles triangle $\triangle \overline{a_2\overline{y}y'}$ is at most $\pi/2$. Therefore the angle $\angle \overline{y'y\overline{x}} < (\pi - \beta) + \pi/2$, as marked in Figure 8(b). This shows that $\overline{y'} \in R_2$. Thus the curve $\overline{C'_2} = \overline{C_2} \cup \overline{y'y'}$ remains in R_2 . This curve $\overline{C'_2}$ is itself either convex, or a convex loop (with loop point \overline{y}). In the former case, Corollary 4 in [5] shows that the flat surface it bounds is planar and so without overlappings. In the second case, we can split the flat surface it bounds into two flat, convex domains, each of which is planar, whence their join is planar. This implies that $\overline{C'_2}$ is without overlappings, and hence so is $\overline{C_2}$.

Applying analogous reasoning to $\overline{\Lambda_1}$ and $\overline{C_1}$ yields the claim that $\overline{C} = \overline{C_1} \cup \overline{C_2}$ does not overlap.

Case 2. Assume now that Case 1 does not hold. This means that all shortest paths falling between τ_1 and τ_2 do not reach C , i.e., they hit the cut locus with respect to x first. The *cut locus* $X = X(x)$ is the closure of the locus of points with more than one shortest path to x . X is a tree (it is also known as the “ridge tree”) with its leaves at the vertices of \mathcal{P} . The cut locus plays a role in related work, including our work in [5, 6], but here we only need its most basic properties. In particular, we established in [5] that the branch of X that is the target of the shortest paths between τ_1 and τ_2 meets C in a single point w . Then there are two shortest paths from x to w , enclosing that branch, which start at or outside of τ_1 and τ_2 . These two segments determine what we called a “fat digon” D , “fat” because it consumes all the potential γ segments that would keep us in Case 1 above. (A metrically accurate polyhedral example is provided in [5, Fig. 12].) Let the angle of D at x be α . See Figure 9(a).

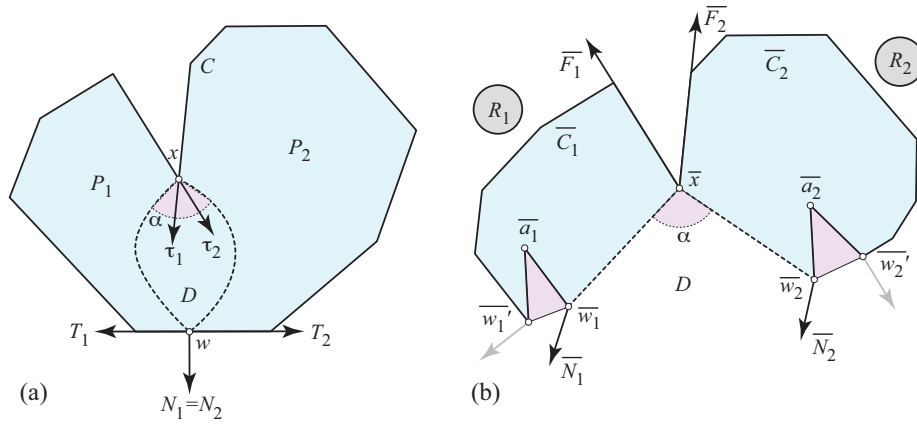


Figure 9: Case 2: (a) The digon D on \mathcal{P} . (b) Planar layout.

Call the convex regions that remain outside of D to either side P_i . Again we perform vertex merging within each P_i to obtain two cones Λ_i , with apices a_i , and we unfold each by cutting along a_iw . Now, in contrast to Case 1, here we

layout the cone unfoldings to share only the point \bar{x} , and such that the angle between $\bar{x}w_1$ and $\bar{x}w_2$, where \bar{w}_1 and \bar{w}_2 are the two images of w , is precisely α , the digon angle at x on \mathcal{P} . This guarantees we obtain a development of $\overline{C_w}$ in the neighborhood of \bar{x} . In analogy with Case 1, define F_1 and F_2 to oppose τ_2 and τ_1 respectively. See Figure 9(b).

The regions R_i bounded by $\overline{N_i} \cup \overline{xw_i} \cup \overline{F_i}$ contain $\overline{C_i}$, following the same logic as in Case 1: analyzing the angles at w_i shows that the second images w'_i are inside R_i , and the curves $\overline{C_i} \cup \overline{xw_i}$ are flat convex loops. Thus, $\overline{C_w} = \overline{C_1} \cup \overline{C_2}$, the development of C when cut at w , avoids overlap. \square

This result on convex loops is best possible in the sense that there are curves C that are convex except at two exceptional points—call them *convex 2-loops*—for which $\overline{C_x}$ overlaps for every x . The basic construction that illustrates this derives from a “sliver tetrahedron,” which has long been known to overlap from a particular edge unfolding. Figure 10 illustrates how doubling the cut tree leads to overlap.

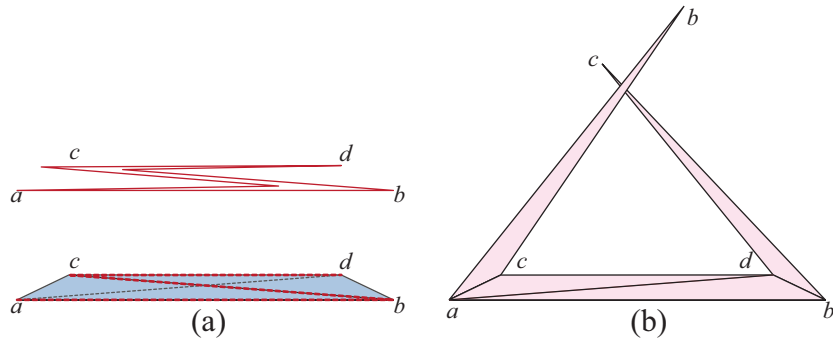


Figure 10: Overlap from a sliver tetrahedron by cutting a convex 2-loop. The curve is a doubling of the cut path (a, b, c, d) , nonconvex at one turn at b , and one turn at c .

The degeneracy of this example may leave it not entirely convincing, but it may be mimicked to be nondegenerate. Figure 11 shows an example of a curve that is convex except at two points, all of whose developments overlap.

5 Reflex Curves and Reflex Loops

Recall that, for each corner c_i of a curve C , $\alpha_i + \omega_i + \beta_i = 2\pi$, where α_i and β_i are the left and right angles at c_i respectively, and ω_i is the curvature at c_i . When C is vertex-free, $\omega_i = 0$ at all corners, and the relationships among the curve classes is simple and natural: the other side of a convex curve is reflex, the other side of a reflex curve is convex. The same holds for the loop versions: the other side of a convex loop is a reflex loop (because $\alpha_m \geq \pi$ implies $\beta_m \leq \pi$, where c_m is the loop point), and the other side of a reflex loop is a convex loop.

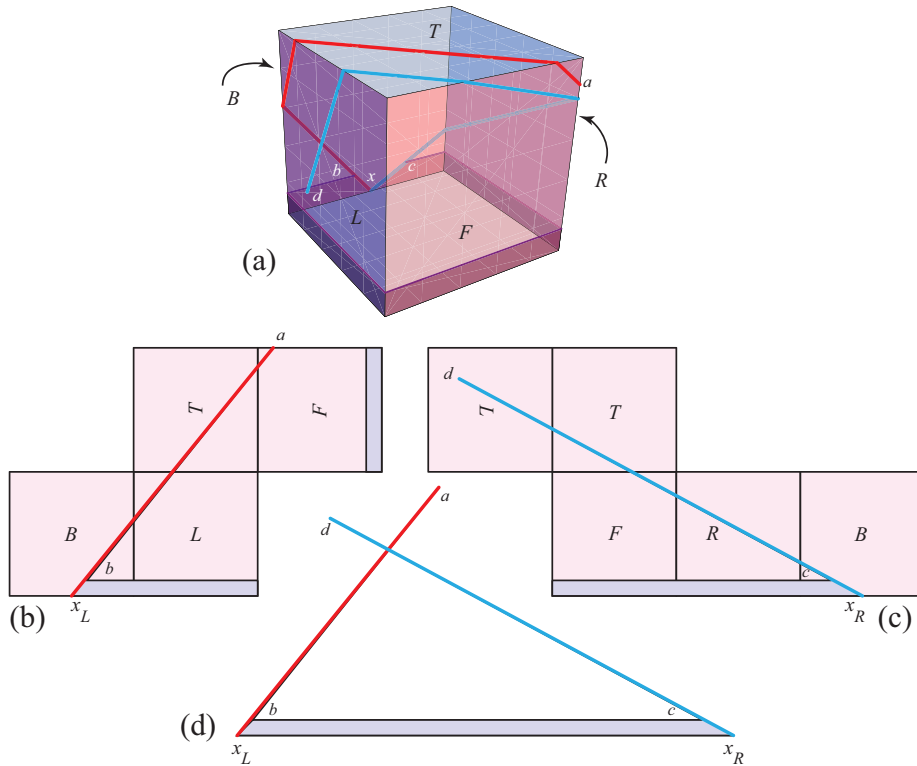


Figure 11: (a) A curve $C = (x_L, a, b, c, d, x_R)$, convex except at the two reflex corners c and d . Here x_L and x_R are slightly separated points near the midpoint x of the back bottom cube edge, and the two “spikes” are too thin for both sides to be distinguished at this resolution. The cube faces are labeled F, L, R, T, B for Front, Left, Right, Top, and Back, respectively. (b) The left portion (x_L, a, b) on an unfolding of the faces it crosses. (c) The right portion (c, d, x_R) . (d) Development of curve C . Note that because C encloses no vertices of the cube, it is isometric to a planar polygon. Thus its development is independent of a cut point: all of its developments are congruent.

When C includes vertices, the relationships between the curve classes are more complicated. The other side of a convex curve is reflex only if the curvatures at the vertices on C are small enough so that $\alpha_i + \omega_i \leq \pi$; C would still be convex even if it just included those vertices inside. The same holds for convex loops.

On the other hand, the other side of a reflex curve is always convex, because nonzero vertex curvatures only make the other side more convex. The other side of a reflex loop is a convex loop, and it is a convex curve if the curvature at the loop point c_m is large enough to force $\alpha_m \leq \pi$, i.e., if $\beta_m + \omega_m \geq \pi$.

This latter subclass of reflex loops—those whose other side is convex—especially interest us, because any convex curve that includes at most one vertex is a reflex loop of that type. All our results in this section hold for this class of curves.

Lemma 6 *Let C be a curve that is either reflex (to its right), or a reflex loop which is convex to the other (left) side, with $\beta_m < \pi$ at the loop point c_m . Then C lives on a cone Λ_R to its reflex side.*

Proof: Again let c_1, c_2, \dots, c_m be the corners of C , with c_m the loop point if C is a reflex loop. Because C is convex to its left, we have $\Omega_L \leq 2\pi$. Just as in Lemma 3, merge the vertices strictly in P_L to one vertex a . Let Λ_L be the cone with apex a on which C now lives. It will simplify subsequent notation to let $\Lambda = \Lambda_L$.

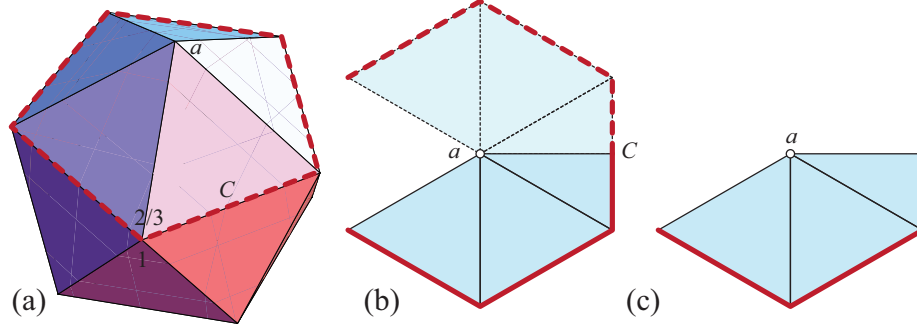


Figure 12: A convex curve C on an icosahedron, with $\alpha_i = \frac{2}{3}\pi$, $\beta_i = \pi$, and $\omega_i = \frac{1}{3}\pi$ at each corner. The cone Λ for C opened (b) and doubly covered (c).

Let N_R be a (small) right neighborhood of C , a neighborhood to the reflex side of C . For subsequent subscript embellishment, we use N to represent N_R . Its shape is irrelevant to the proof, as long as it is vertex free and its left boundary is C .

Join a to each corner c_i by a cone-generator g_i (a ray from a on Λ). Lemma 4 ensures this is possible. Cut along g_i beyond c_i into N . There are choices how to extend g_i beyond c_i , but the choice does not matter for our purposes. For example, one could choose a cut that bisects β_i at c_i . Insert along each cut into N a *curvature triangle*, that is, an isosceles triangle with two sides equal to the

cut length, and apex angle ω_i at c_i . (If c_i does not coincide with a vertex of \mathcal{P} , then $\omega_i = 0$ and no curvature triangle is inserted.) This flattens the surface at c_i , and “fattens” N to N' without altering C or the cone Λ left of C . Now N' lives on the same cone Λ that C and its left neighborhood N_L do.

From now on we view Λ and the subsequent cones we will construct as flattened into the plane, producing a doubly covered cone with half the apex angle. (Notice that here “doubly covered” above refers to a neighborhood of the cone apex, and not to the image of the curve C .) It is always possible to choose any generator ax for $x \in C$ and flatten so that ax is the leftmost extreme edge of the double cone. We start by selecting $x = c_1$, so that g_1 is the leftmost extreme; let h_1 be the rightmost extreme edge. We pause to illustrate the construction before proceeding.

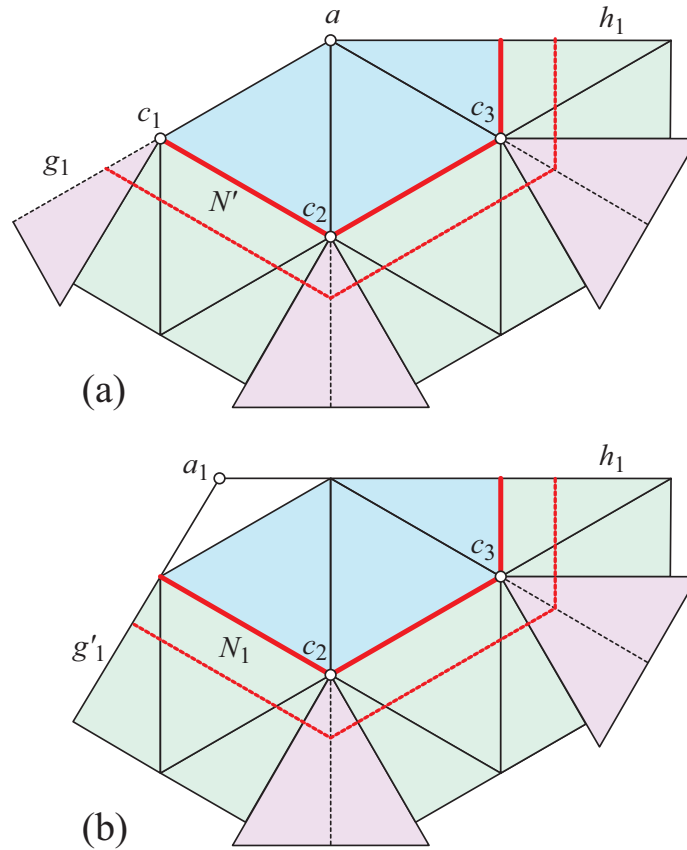


Figure 13: (a) After insertion of curvature triangles, N' lives on Λ . (b) Removing the doubly covered half curvature triangle at c_1 leads to a new cone Λ_1 . (In this and in Figure 14 we display the full icosahedron faces to the right of C , although only a small neighborhood is relevant to the proof.)

Let C be the curve on the icosahedron illustrated in Figure 12(a). This curve already lives on the pentagonal pyramid cone Λ without any vertex merging. Figure 12(b) shows the five equilateral triangles incident to the apex, and (c) shows the corresponding doubly covered cone. Figure 13(a) illustrates Λ after insertion of the curvature triangles to the right of C , each with apex angle $\omega_i = \frac{1}{3}\pi$. A possible neighborhood N' is outlined.

After insertion of all curvature triangles, we in some sense erase where they were inserted, and just treat N' as a band living on Λ . Now, with g_1 the leftmost extreme, we identify a half-curvature triangle on the front side, matched by a half-curvature triangle on the back side, incident to c_1 in N' . Each triangle has angle $\frac{1}{2}\omega_1$ at c_1 . See again Figure 13(a). Now rotate g_1 counterclockwise about c_1 by $\frac{1}{2}\omega_1$, and cut out the two half-curvature triangles from N' , regluing the front to the back along the cut segment. Extend the rotated line g'_1 to meet the extension of h_1 . Their intersection point is the apex a_1 of a new (doubly covered) cone Λ_1 , on which neither a nor c_1 are vertices. Note that the rotation of g_1 effectively removes an angle of measure ω_1 incident to c_1 from the N' side, and inserts it on the other side of C . See Figure 13(b). Call the new neighborhood N_1 , and the new convex curve C_1 . C_1 is the same as C except that the angle at c_1 is now $\alpha_1 + \omega_1$, which by the assumption of the lemma, is still convex because $\beta_1 \geq \pi$.

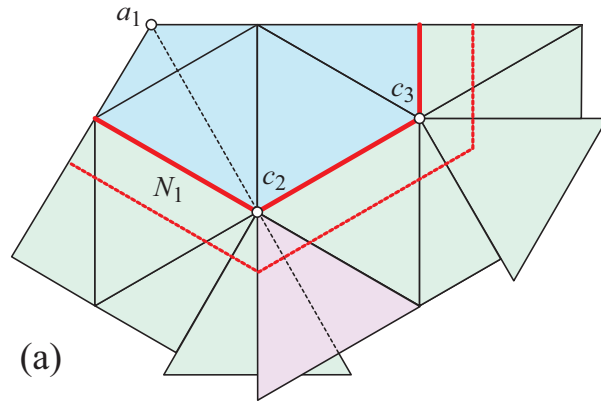
Now we argue that g'_1 does not intersect N_1 other than where it forms the leftmost boundary. For if g'_1 intersected N_1 elsewhere, then, taking N_1 to be smaller and smaller, tending to C_1 , we conclude that g'_1 must intersect C_1 at a point other than c_1 . But this contradicts the fact that either of the two planar images (from the two sides of Λ) of C_1 is convex. Indeed g'_1 is a supporting line at c_1 to the convex set constituted by Λ_1 up to C_1 .

Note that we have effectively merged vertices c_1 and a to form a_1 , in a manner similar to the vertex merging used in Lemma 3. The advantage of the process just described is that it does not rely on having a triangle half-angle no more than π at the new cone apex.

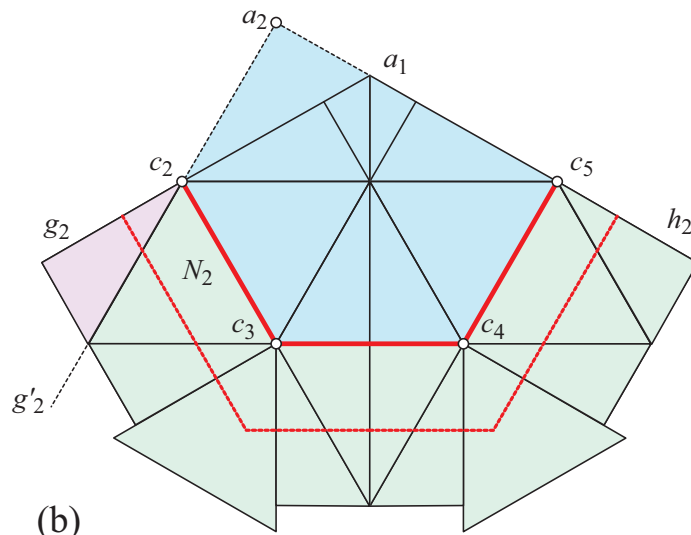
Next we eliminate the curvature triangle inserted at c_2 . Let g_2 be the generator from a_1 through c_2 (again, Lemma 4 applies). Identify a curvature triangle of apex angle ω_2 in N_1 bisected by g_2 ; see Figure 14(a). Now reflaten the cone Λ_1 so that g_2 is the left extreme, and let h_2 be the right extreme, as in (b) of the figure. Rotate g_2 by $\frac{1}{2}\omega_2$ about c_2 to produce g'_2 , cut out the half-curvature triangles on both the front and back of N_1 , and extend g'_2 to meet the extension of h_2 at a new apex a_2 . Now we have a new neighborhood N_2 , with left boundary the convex curve C_2 , living on a cone Λ_2 .

We apply this process through c_1, \dots, c_{m-1} . It could happen at some stage that g'_i and the h_i extension meet on the other side of C_i , in which case the cone apex is to the reflex side. (Or, they could be parallel and meet "at infinity," which is what occurs with the icosahedron example.) From the assumption of the lemma that $\beta_i \geq \pi$ for $i < m$, $\alpha_i + \omega_i \leq \pi$ and so the curves C_i remain convex throughout the process. So the argument above holds.

For the last, possibly exceptional corner c_m, C_{m-1} from the previous step is convex, but the final step could render C_m nonconvex (if $\alpha_m + \omega_m > \pi$). But as



(a)



(b)

Figure 14: (a) Generator g_2 from a_1 through c_2 into N_1 . (b) Reoriented so g_2 is left extreme.

there is no further processing, this nonconvexity does not affect the proof. \square

For the icosahedron example, five insertions of $\frac{1}{3}\pi$ curvature triangles, together with the original $\frac{1}{3}\pi$ curvature at a , produces a cylinder. And indeed, $\beta_i = \pi$ for the five c_i corners of C , and C forms a circle on a cylinder.

Lemma 7 *Let C be a curve satisfying the same conditions as for Lemma 6. Then C is visible from the apex a of the cone Λ on which it lives to its reflex side.*

Proof: Again letting c_1, \dots, c_m be the corners of C , with c_m the possibly exceptional vertex, we know that $\beta_i \geq \pi$ for $i = 1, \dots, m-1$, but it may be that $\beta_m < \pi$. Just as in the proof of Lemma 6, we flatten Λ into the plane, this time choosing c_m to lie on the leftmost extreme generator L_1 of Λ . Let b be the point of C that lies on the rightmost extreme generator L_2 in this flattening. Finally, let C_u be the portion of C on the upper surface of the flattened Λ , and C_l the portion on the lower surface. See Figure 15. Now that we have

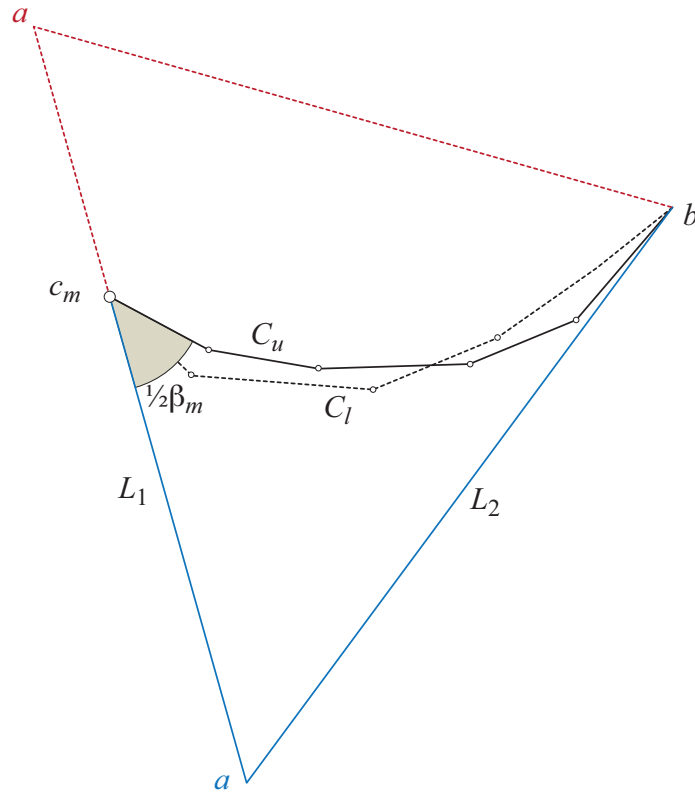


Figure 15: The apex a could lie either to the reflex or to the convex side of C .
 placed the one anomalous corner on the extreme boundary L_1 , both C_u and C_l

present a uniform aspect to the apex a , whether it is to the convex or reflex side of C : every corner of C_u and C_l is reflex (or flat) toward the reflex side, and convex (or flat) toward the convex side. In particular, $c_m b \cup C_u$ is a planar convex domain. Each line through a intersects $c_m b$ exactly once, and therefore intersects C_u exactly once; and similarly for C_l . \square

Just as we observed for convex loops, this visibility lemma does not hold for all reflex loops—the assumption that the other side is convex is essential to the proof.

We summarize this section in a theorem (recall that $\Omega_L + \Omega_C + \Omega_R = 4\pi$).

Theorem 2 *Let C be a curve that is either reflex (to its right), or a reflex loop which is convex to the other (left) side, with $\beta_m < \pi$ at the loop point c_m . Then C lives on a cone Λ_R to its reflex side, and is visible from its apex a . If $\Omega_R > 2\pi$, then the reflex neighborhood N_R is to the unbounded side of Λ_R , i.e., the apex of Λ_R is left of C ; if $\Omega_R < 2\pi$, then N_R is to the bounded side, i.e., the apex of Λ_R is to the right side of C . If $\Omega_R = 2\pi$, $C \cup N_R$ lives on a cylinder.*

Proof: The uniqueness follows from Lemma 1. The cone Λ_R constructed in the proof of Lemma 6 results in the cone apex to the convex side of C as long as $\Omega_L + \Omega_C \leq 2\pi$, when $\Omega_R \geq 2\pi$. Excluding the cylinder cases, this justifies the claims concerning on which side of Λ_R the neighborhood N_R resides. The apex curvature of Λ_R is $\min\{\Omega_L + \Omega_C, \Omega_R\}$. \square

Example 4. An example of a reflex loop that satisfies the hypotheses of Lemma 6 is shown in Figure 16(a). Here C has five corners, and is convex to one side at each. C passes through only one vertex of the cuboctahedron \mathcal{P} , and so it is reflex at the four non-vertex corners to its other side. Corner c_5 coincides with a vertex of \mathcal{P} , which has curvature $\omega_5 = \frac{1}{3}\pi$. Here $\alpha_5 = \beta_5 = \frac{5}{6}\pi$. Because $\beta_5 < \pi$, C is a reflex loop. We have $\Omega_L = \frac{2}{3}\pi$ because C includes two cuboctahedron vertices, u and v in the figure. $\Omega_C = \omega_5 = \frac{1}{3}\pi$. And therefore $\Omega_R = 3\pi$. The apex curvature of Λ_L is $\Omega_L = \frac{2}{3}\pi$, and the apex curvature of Λ_R is π . N_R lives on the unbounded side of this cone, which is shown shaded in Figure 16(b). Note the apex a is left of C , in accord with the lemma.

6 Discussion

We summarize the results claimed in the Introduction in a theorem:

Theorem 3 *On a convex polyhedron, every convex curve left-develops without overlap, and every reflex curve, and reflex loop whose other side is convex, right-develops without overlap, for every cut point. Every convex loop has some cut-point from which it left-develops without overlap.*

Proving that a curve on a convex polyhedron lives on a cone is a powerful technique for establishing that these polyhedron curves develop without overlap. Even when a curve—such as a convex loop—does not live on a cone, still the cone perspective can help prove nonoverlapping development (as it did in Lemma 5).

Many questions remain.

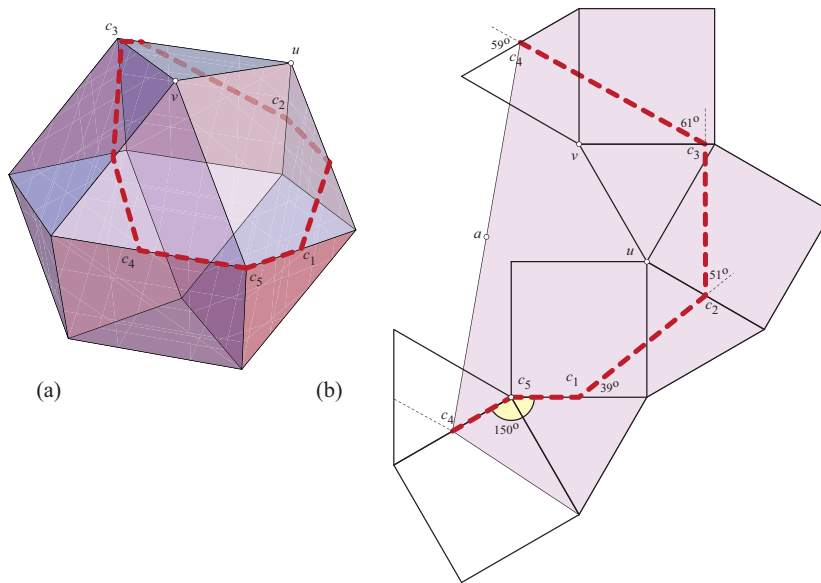


Figure 16: (a) A curve C of five corners passing through one polyhedron vertex. C is convex to one side, and a reflex loop to the other, with loop point c_5 , at which $\beta_5 = \frac{5}{6}\pi (= 150^\circ) < \pi$. (b) The cone Λ_R with apex a is shaded.

Overlapping Developments. First, it is not the case that every curve that lives on a cone develops without overlap. Here we show that there exist C such that \overline{C}_x is nonsimple for every choice of x . We provide one specific example, but it can be generalized.

The cone Λ has apex angle $\alpha = \frac{3}{4}\pi$; it is shown cut open and flattened in two views in Figure 17(a,b). An open curve $C' = (p_1, p_2, p_3, p_4, p_5)$ is drawn on the cone. Directing C' in that order, it turns left by $\frac{3}{4}\pi$ at p_2, p_3 , and p_4 . From p_5 , we loop around the apex a with a segment $S = (p_5, p_6, p'_5)$, where p'_5 is a point near p_5 (not shown in the figure). Finally, we form a simple closed curve on Λ by then doubling C' at a slight separation (again not illustrated in the figure), so that from p_5 it returns in reverse order along that slightly displaced path to p_1 again. Note that $C = C \cup S \cup C'$ is closed and includes the apex a in its (left) interior.

Now, let x be any point on C from which we will start the development \overline{C}_x . Because C is essentially $C' \cup C'$, x must fall in one or the other copy of C' , or at their join at p_1 . Regardless of the location of x , at least one of the two copies of C' is unaffected. So \overline{C}_x must include $\overline{C'}$ as a subpath in the plane.

Finally, developing C' reveals that it self-intersects: Figure 17(c). Therefore, \overline{C}_x is not simple for any x . Moreover, it is easy to extend this example to force self-intersection for many values of α and analogous curves. The curve C' was selected only because its development is self-evident.

Slice Curves. There are curves already known to develop without overlap that are not known to live on a cone. One particular class we could not settle are the slice curves. A *slice curve* C is the intersection of \mathcal{P} with a plane. Slice curves in general are not convex. The intersection of \mathcal{P} with a plane is a convex polygon in that plane, but the surface angles of \mathcal{P} to either side along C could be greater or smaller than π at different points. Slice curves were proved to develop without intersection, to either side, in [7], so they are good candidates to live on cones. However, we have not been able to prove that they do.

Convex Loops. Although we have shown that there is some cut point from which every convex loop develops without overlap (Lemma 5), we have not determined all the cut points that enjoy this property.

Cone Curves. Finally, we have not obtained a complete classification of the curves on a cone that develop, for every cut point x , as simple curves in the plane. It would be equally interesting to identify the class of curves on cones for which there exists at least one cut-point that leads to simple development. Indeed, the same questions for curves on a sphere are also unresolved [3].

References

- [1] Aleksandr D. Alexandrov. *Convex Polyhedra*. Springer-Verlag, Berlin, 2005. Monographs in Mathematics. Translation of the 1950 Russian edition by N. S. Dairbekov, S. S. Kutateladze, and A. B. Sossinsky.
- [2] Erik D. Demaine and Joseph O’Rourke. *Geometric Folding Algorithms: Linkages, Origami, Polyhedra*. Cambridge University Press, July 2007. <http://www.gfalop.org>.
- [3] Erik D. Demaine and Joseph O’Rourke. Open problems from CCCG 2009. In *Proc. 22nd Canad. Conf. Comput. Geom.*, pages 83–86, 2010.
- [4] Kouki Ieiri, Jin-ichi Itoh, and Costin Vilcu. Quasigeodesics and farthest points on convex surfaces. To appear, *Adv. Geom.*, 2011.
- [5] Jin-ichi Itoh, Joseph O’Rourke, and Costin Vilcu. Star unfolding convex polyhedra via quasigeodesic loops. *Discrete Comput. Geom.*, 44:35–54, 2010.
- [6] Jin-ichi Itoh, Joseph O’Rourke, and Costin Vilcu. Source unfoldings of convex polyhedra via certain closed curves. Submitted, 2011.
- [7] Joseph O’Rourke. On the development of the intersection of a plane with a polytope. *Comput. Geom. Theory Appl.*, 24(1):3–10, 2003.
- [8] Joseph O’Rourke and Catherine Schevon. On the development of closed convex curves on 3-polytopes. *J. Geom.*, 13:152–157, 1989.

- [9] Joseph O'Rourke and Costin Vilcu. Conical existence of closed curves on convex polyhedra. <http://arxiv.org/abs/1102.0823>, February 2011.

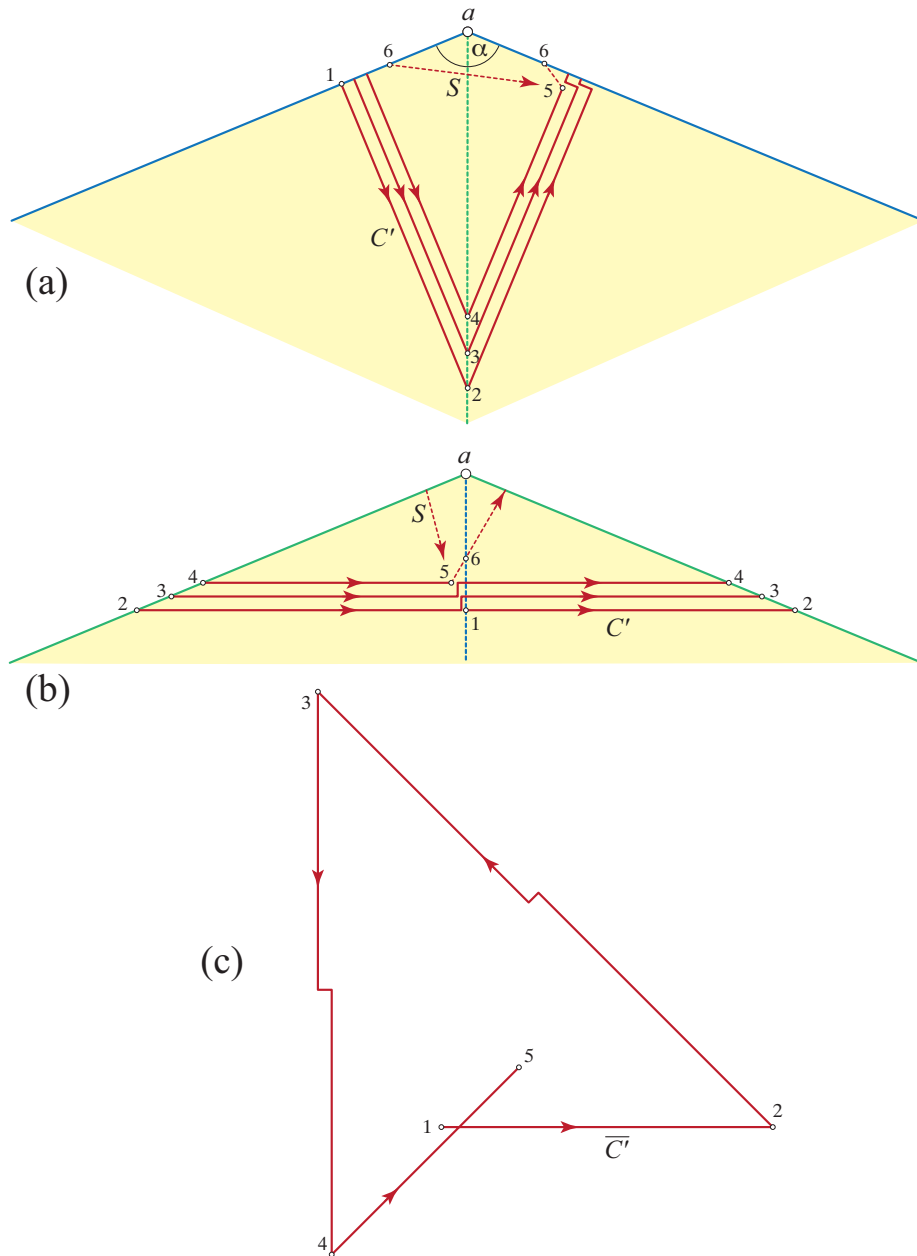


Figure 17: (a) Open curve $C' = (p_1, p_2, p_3, p_4, p_5)$ on cone of angle α , with cone opened. (b) A different opening of the same cone and curve. (c) Development of curve $\overline{C'}$ self-intersects.